

HP A5820X & A5800 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: 5998-1628
Software version: Release 1211
Document version: 5W100-20110430



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	1
Ethernet interface naming conventions	1
Management Ethernet interface	1
Switchable operating mode of Ethernet interfaces	2
General Ethernet interface configuration	2
Configuring basic settings of an Ethernet interface	2
Configuring the operating mode of an Ethernet interface	3
Configuring flow control on an Ethernet interface	4
Configuring link change suppression on an Ethernet interface	7
Configuring loopback testing on an Ethernet interface	8
Setting the statistics polling interval on an Ethernet interface	9
Configuring jumbo frame support on an Ethernet interface	9
Configuring link-up port auto power-down on an Ethernet interface	9
Configuring a Layer 2 Ethernet interface	10
Layer 2 Ethernet interface configuration task list	10
Configuring a port group	10
Enabling link-down port auto power-down	11
Setting speed options for auto negotiation on an Ethernet interface	12
Configuring traffic storm protection	13
Enabling single-port loopback detection on an Ethernet interface	15
Enabling multi-port loopback detection	16
Setting the MDI mode of an Ethernet interface	17
Enabling bridging on an Ethernet interface	18
Testing the cable connection of an Ethernet interface	18
Configuring the connection mode of an Ethernet interface	19
Configuring a Layer 3 Ethernet interface	19
Displaying and maintaining an Ethernet interface	19
PFC configuration example	21
Loopback and null interface configuration	23
Loopback interface	23
Configuring a loopback interface	24
Null interface	24
Configuring null 0 interface	25
Displaying and maintaining loopback and null interfaces	25
MAC address table configuration	26
How a MAC address table entry is created	26
Types of MAC address table entries	26
MAC address table-based frame forwarding	27
Configuring the MAC address table	27
Manually configuring MAC address table entries	28
Disabling MAC address learning on a VLAN	28
Configuring the aging timer for dynamic MAC address entries	29
Configuring the MAC learning limit on ports	29
Enabling MAC address roaming	30
Displaying and maintaining MAC address tables	31
MAC address table configuration example	32
MAC Information configuration	33
How MAC Information works	33
Configuring MAC Information	33

Enabling MAC Information globally	33
Enabling MAC Information on an interface	33
Configuring MAC Information mode	34
Configuring the interval for sending Syslog or trap messages	34
Configuring the MAC Information queue length	34
MAC Information configuration example	35
Ethernet link aggregation configuration	36
Basic concepts	36
Aggregating links in static mode	39
Aggregating links in dynamic mode	40
Load sharing criteria for link aggregation groups	42
Ethernet link aggregation configuration task list	42
Configuring an aggregation group	43
Configuring a static aggregation group	43
Configuring a dynamic aggregation group	44
Configuring an aggregate interface	45
Configuring the description of an aggregate interface	45
Enabling link state traps for an aggregate interface	45
Setting the minimum number of selected ports for an aggregation group	46
Shutting down an aggregate interface	46
Restoring the default settings for an aggregate interface	47
Configuring load sharing for link aggregation groups	47
Enabling local-first load sharing for link aggregation	48
Enabling link-aggregation traffic redirection	49
Displaying and maintaining Ethernet link aggregation	50
Ethernet link aggregation configuration examples	50
Layer 2 static aggregation configuration example	51
Layer 2 dynamic aggregation configuration example	53
Port isolation configuration	55
Configuring the isolation group	55
Displaying and maintaining isolation groups	56
Port isolation configuration example	56
MSTP configuration	58
Why STP	58
Protocol packets of STP	58
Basic concepts in STP	59
How STP works	60
Introduction to RSTP	66
Introduction to MSTP	66
Why MSTP	66
Basic concepts in MSTP	67
How MSTP works	70
Implementation of MSTP on devices	71
Protocols and standards	71
MSTP configuration task list	71
Configuring MSTP	73
Configuring an MST region	73
Configuring the root bridge or a secondary root bridge	73
Configuring the work mode of an MSTP device	74
Configuring the priority of a device	75
Configuring the maximum hops of an MST region	75
Configuring the network diameter of a switched network	76
Configuring timers of MSTP	76
Configuring the timeout factor	77

Configuring the maximum port rate	78
Configuring ports as edge ports	78
Configuring path costs of ports	79
Configuring port priority	81
Configuring the link type of ports	82
Configuring the mode a port uses to recognize/send MSTP packets	82
Enabling the output of port state transition information	83
Enabling the MSTP feature	83
Performing mCheck	84
Configuring digest snooping	85
Configuring no agreement check	87
Configuring TC snooping	89
Configuring protection functions	91
Displaying and maintaining MSTP	94
MSTP configuration example	94
BPDU tunneling configuration	99
BPDU tunneling implementation	100
Configuring BPDU tunneling	101
Configuration prerequisites	101
Enabling BPDU tunneling	101
Configuring destination multicast MAC address for BPDUs	102
BPDU tunneling configuration examples	103
BPDU tunneling for STP configuration example	103
BPDU tunneling for PVST configuration example	104
VLAN configuration	106
VLAN fundamentals	106
Types of VLAN	107
Configuring basic VLAN settings	108
Configuring basic settings of a VLAN interface	108
Port-based VLAN configuration	110
Introduction to port-based VLAN	110
Assigning an access port to a VLAN	111
Assigning a trunk port to a VLAN	113
Assigning a hybrid port to a VLAN	114
Port-based VLAN configuration example	115
MAC-based VLAN configuration	116
Introduction to MAC-based VLAN	116
Configuring MAC-based VLAN	118
MAC-based VLAN configuration example	120
Protocol-based VLAN configuration	123
Introduction to protocol-based VLAN	123
Configuring a protocol-based VLAN	124
Protocol-based VLAN configuration example	125
IP Subnet-based VLAN configuration	127
Configuring an IP subnet-based VLAN	128
Displaying and maintaining VLAN	129
Super VLAN configuration	130
Configuring super VLAN	130
Displaying and maintaining super VLAN	131
Super VLAN configuration example	132
Isolate-user-VLAN configuration	135
Configuring isolate-user-VLAN	136
Configuring an isolate-user-VLAN	136
Configuring secondary VLANs	137

Associating secondary VLANs with an isolate-user-VLAN	138
Displaying and maintaining isolate-user-VLAN	139
Isolate-user-VLAN configuration example	139
Voice VLAN configuration	142
OUI addresses	142
Voice VLAN assignment modes	142
Security mode and normal mode of voice VLANs	145
Configuring a voice VLAN	146
Configuration prerequisites	146
Configuring QoS priority settings for voice traffic on an interface	146
Configuring a port to operate in automatic voice VLAN assignment mode	147
Configuring a port to operate in manual voice VLAN assignment mode	147
Displaying and maintaining voice VLAN	148
Voice VLAN configuration examples	149
Automatic voice VLAN mode configuration example	149
Manual voice VLAN assignment mode configuration example	151
GVRP configuration	153
GARP	153
GVRP	156
Protocols and standards	156
GVRP configuration task list	157
Configuring GVRP functions	157
Configuring GARP timers	158
Displaying and maintaining GVRP	159
GVRP configuration examples	160
GVRP normal registration mode configuration example	160
GVRP fixed registration mode configuration example	161
GVRP forbidden registration mode configuration example	162
QinQ configuration	165
Background and benefits	165
How QinQ works	165
QinQ frame structure	166
Implementations of QinQ	167
Modifying the TPID in a VLAN tag	167
Protocols and standards	168
QinQ configuration task list	168
Configuring basic QinQ	169
Enabling basic QinQ	169
Configuring VLAN transparent transmission	169
Configuring selective QinQ	170
Configuring an outer VLAN tagging policy	170
Configuring an inner-outer VLAN 802.1p priority mapping	172
Configuring inner VLAN ID substitution	174
Configuring the TPID value in VLAN tags	175
Configuring the TPID value in the CVLAN tag	175
Configuring the TPID value in the SVLAN tag	175
QinQ configuration examples	175
Basic QinQ configuration example	175
Selective QinQ configuration example	178
VLAN mapping configuration	181
Application scenario of one-to-one and many-to-one VLAN mapping	181
Application scenario of two-to-two VLAN mapping	183
Concepts and terms	184
VLAN mapping implementations	184

Configuring VLAN mapping	186
Configuring one-to-one VLAN mapping	186
Configuring many-to-one VLAN mapping	188
Configuring two-to-two VLAN mapping	191
VLAN mapping configuration examples	195
One-to-one and many-to-one VLAN mapping configuration example	195
Two-to-two VLAN mapping configuration example	201
LLDP configuration	204
Basic concepts	204
How LLDP works	208
Protocols and standards	209
LLDP configuration task list	209
Performing basic LLDP configuration	210
Enabling LLDP	210
Setting the LLDP operating mode	210
Setting the LLDP re-initialization delay	211
Enabling LLDP polling	211
Configuring the advertisable TLVs	212
Configuring the management address and its encoding format	212
Setting other LLDP parameters	213
Setting an encapsulation format for LLDPDUs	214
Configuring CDP compatibility	215
Configuration prerequisites	215
Configuring CDP compatibility	215
Configuring DCBX	216
DCBX configuration task list	217
Enabling LLDP and DCBX TLV advertising	217
Configuring APP parameters	217
Configuring ETS parameters	219
Configuring PFC parameters	220
Configuring LLDP trapping	221
Displaying and maintaining LLDP	222
LLDP configuration examples	223
Basic LLDP configuration example	223
CDP-compatible LLDP configuration example	225
DCBX configuration example	227
Service loopback group configuration	232
Service types of service loopback groups	232
Requirements on service loopback ports	232
States of service loopback ports	232
Configuring a service loopback group	234
Displaying and maintaining service loopback groups	234
Service loopback group configuration example	235
Support and other resources	236
Contacting HP	236
Subscription service	236
Related information	236
Documents	236
Websites	236
Conventions	237
Index	239

Ethernet interface configuration

Ethernet interface naming conventions

The GE and 10-GE interfaces on the A5800 and A5820X switch series are named in the format of *interface-type A/B/C*, where the following definitions apply:

- A represents the ID of the switch in an IRF virtual device. If the switch is not assigned to any IRF virtual device, A takes 1.
- B represents a slot number on the switch. It takes 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.

For example, GigabitEthernet 1/1/2 represents a GE interface on expansion card 1 on a standalone switch.

The A5800-48G-PoE+ Switch (JC101A), A5800-48G-PoE+ TAA Switch (JG242A), A5820X-14XG-SFP+ Switch (JC106A), and A5820X-14XG-SFP+ TAA Switch (JG259A) provide an internal 10-GE interface named Ten-GigabitEthernet A/3/1 for OAA applications.

For more information about SFP and SFP+ interfaces, see the *HP A5800 Switches Series Installation Guide* or the *HP A5820X Switches Series Installation Guide*.

Management Ethernet interface

The switch provides one management Ethernet interface. This interface uses an RJ-45 connector. Connect it to a PC for software loading and system debugging.

To configure the management Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter management Ethernet interface view	interface M-GigabitEthernet <i>interface-number</i>	—
3. Change the description of the interface	description <i>text</i>	Optional. By default, the description of the management interface is M-GigabitEthernet0/0/0 Interface .
4. Shut down the interface	shutdown	Optional. By default, the management Ethernet interface is up.
5. Restore the default settings for the interface	default	Optional.

Switchable operating mode of Ethernet interfaces

Switch the operating mode of an Ethernet interface from Layer 2 (bridge) to Layer 3 (route) and vice versa, as needed. When operating in Layer 2 mode, the Ethernet interface acts as a Layer 2 interface. When operating in Layer 3 mode, the Ethernet interface acts as a Layer 3 interface. For more information, see [“Configuring the operating mode of an Ethernet interface.”](#)

General Ethernet interface configuration

This section describes the attributes and configurations common to Layer 2 and Layer 3 Ethernet interfaces. For specific attributes, see [“Configuring a Layer 2 Ethernet interface”](#) and [“Configuring a Layer 3 Ethernet interface.”](#)

Configuring basic settings of an Ethernet interface

Set an Ethernet interface to operate in one of the following duplex modes:

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

Similarly, you can set the speed of an Ethernet interface or enable it to negotiate a speed with its peer automatically. For a 1000-Mbps Layer 2 Ethernet interface, you can also set speed options for auto negotiation. The two ends can pick a speed only from the available options. For more information, see [“Setting speed options for auto negotiation on an Ethernet interface.”](#)

To configure an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	—
3. Change the description of the interface	description <i>text</i>	Optional. By default, the description of an interface is the interface name followed by the “Interface” string, GigabitEthernet1/0/1 Interface for example.
4. Set the duplex mode	duplex { auto full half }	Optional. By default, auto-negotiation mode applies. The half duplex mode is not available for SFP interfaces, 10G SFP+ interfaces, or Ethernet electrical interfaces operating at 1000 Mbps.

To do...	Use the command...	Remarks
5. Set the interface speed	<code>speed { 10 100 1000 10000 auto }</code>	<p>Optional.</p> <p>A 10G SFP+ port supports:</p> <ul style="list-style-type: none"> The 1000 and auto parameters if installed with a Gigabit SFP transceiver. The 10000 and auto parameters if installed with a 10G SFP+ transceiver. <p>A a Gigabit SFP port supports:</p> <ul style="list-style-type: none"> The 1000 and auto parameters if installed with a Gigabit SFP transceiver. The 100 and auto parameters if installed with a 100-Mbps SFP transceiver. <p>By default, the auto-negotiation mode applies.</p>
6. Shut down the Ethernet interface	<code>shutdown</code>	<p>Optional.</p> <p>By default, an Ethernet interface is in up state.</p> <p>To bring up an Ethernet interface, use the undo shutdown command.</p>
7. Restore the default settings for the Ethernet interface	<code>default</code>	Optional.

Configuring the operating mode of an Ethernet interface

△ CAUTION:

- To avoid packet loss caused by congestion, you must perform the same PFC configuration on all ports that the packets pass through.
- To configure DCBX on the port of an A5820X switch series, you must use the **priority-flow-control auto** command to enable PFC. Otherwise, PFC data cannot be advertised. For more information about DCBX, see the chapter “LLDP configuration.”
- To use the PFC feature on a link, ensure that it is enabled on both ends of the link. An interface processes PFC pause frames only when PFC is enabled on it.

The relationship between the PFC function and the generic flow control function is shown in [Table 2](#).

An Ethernet interface operates either in Layer 2 (bridge) or Layer 3 (route) mode. To meet networking requirements, you can use a command to set the operating mode of an Ethernet interface to bridge or route.

To change the operating mode of an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Change the operating mode of the Ethernet interface	port link-mode { bridge route }	Required. By default, Ethernet interfaces operate in bridge mode (as Layer 2 Ethernet interfaces).

After you change the operating mode of an Ethernet interface, all settings of the Ethernet interface are restored to their defaults under the new operating mode.

The relationship between the PFC function and the generic flow control function is shown in [Table 2](#).

Configuring flow control on an Ethernet interface

Avoid packet drops on a link by enabling flow control at both ends of the link. The flow control function enables the receiving end to require the sending end to suspend sending packets when congestion occurs.

The following flow control mechanisms are available:

- Generic flow control, which controls transmission on a link for all packets as a whole. Both the A5800 and the A5820X switch series support this mechanism.
- PFC, which controls transmission on a link based on 802.1p priority. Only the A5820X switch series support this mechanism.

The subsequent sections describe the two mechanisms and how to configure them.


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 1](#), 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 1 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.

To configure flow control on an interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Enable TxRx mode flow control	flow-control	Required. Use either command.
4. Enable Rx mode flow control	flow-control receive enable	By default, flow control is disabled on an Ethernet interface.

Configuring PFC on an Ethernet interface

PFC performs flow control based on 802.1p priorities. With PFC enabled, an interface requires its peer to suspend sending packets with certain 802.1p priorities when congestion occurs. By decreasing the transmission rate, PFC helps avoid packet loss.

You can enable PFC for certain 802.1p priorities at the two ends of a link. When network congestion occurs, the local device checks the PFC status for the 802.1p priority carried in each arriving packet.

- If PFC is enabled for the 802.1p priority, the local device accepts the packet and sends a PFC pause frame to the peer. The peer stops sending packets carrying this 802.1p priority for an interval as specified in the PFC pause frame. This process repeats until the congestion is removed.
- If PFC is disabled for the 802.1p priority, the local port drops the packet.

Each local precedence value corresponds to a queue. The 802.1p-to-local priority mapping is as shown in Table 1. Modify the 802.1p-to-local priority mapping table with the **qos map-table dot1p-lp** and **import import-value-list export export-value** commands. For more information about the two commands, see the *ACL and QoS Command Reference*.

Table 1 The default 802.1p-to-local priority mapping table

802.1p priority	Local precedence value
0	2
1	0
2	1
3	3
4	4
5	5
6	6
7	7

To configure PFC on an interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	The PFC feature is available only on 10-Gigabit interfaces.
3. Enable PFC on the interface	priority-flow-control { auto enable }	Required. Disabled by default.
4. Enable PFC for specific 802.1p priorities	priority-flow-control no-drop dot1p <i>dot1p-list</i>	Required. By default, PFC is disabled for all 802.1p priorities.
5. Configure the interface to use the 802.1p priorities carried in packets for priority mapping	qos trust dot1p	Required. By default, Ethernet interfaces do not trust the priorities carried in incoming packets, and the switch uses the port priority of packet receiving ports as the 802.1p priorities of incoming packets.

Only A5820X switch series support this function.

For more information about 802.1p priority, trusted packet priority type, and port priority, see the *ACL and QoS Configuration Guide*. For more information about the **qos trust dot1p** command, see the *ACL and QoS Command Reference*.

On an A5820X switch series, you can enable PFC by using either the **priority-flow-control enable** command or the **priority-flow-control auto** command.

HP recommends that you enable PFC for only one 802.1p priority to ensure lossless transmission of traffic, for example, traffic of a FCoE-based data center. Packet loss might still occur if you enable PFC for multiple 802.1p priorities.

Table 2 The relationship between the PFC function and the generic flow control function

flow-control	priority-flow-control enable	priority-flow-control no-drop dot1p	Remarks
Unconfigurable	Configured	Configured	You cannot enable generic flow control by using the flow-control command on a port where PFC is enabled and PFC is enabled for the specified 802.1p priority values.
Configured	Configurable	Unconfigurable	<ul style="list-style-type: none"> On a port configured with the flow-control command, you can enable PFC, but cannot enable PFC for specific 802.1p priorities. Enabling both generic flow control and PFC disables a port from sending common or PFC pause frames to inform its peer of congestion conditions. However, the port can still handle common and PFC pause frames from its peer.

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 with the **shutdown** or **undo shutdown** command, the interface reports the event to the upper layers immediately.

On an A5800 or A5820X 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

To enable an Ethernet interface to suppress link down events:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	—
3. Set a link down suppression interval	link-delay <i>delay-time</i>	Required. Link down suppression is disabled by default.

Configuring link up suppression

To configure link up suppression on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	—
3. Set a link up suppression interval	link-delay <i>delay-time</i> mode up	Required. Link up suppression is disabled by default.


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

Perform loopback testing on an Ethernet interface to check 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 2, internal loopback testing is performed on Port 1. During the internal loopback testing, the interface sends out a certain number of test packets, which are looped back to the interface over the self-loop created on the switching chip

Figure 2 Internal loopback testing



- External loopback testing, which tests the hardware of Ethernet interfaces. As shown in Figure 3, 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 the external loopback testing, the interface sends out 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 3 External loopback testing



To perform loopback testing on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface or ONU interface view	interface <i>interface-type interface-number</i>	—
3. Perform loopback testing	loopback { external internal }	Required

On an interface that is physically down, you can only perform internal loopback testing. On an interface shut down administratively, you can perform neither internal nor external loopback testing.

The **speed**, **duplex**, **mdi**, and **shutdown** commands are unavailable during loopback testing.

During loopback testing, an Ethernet interface works in full duplex mode. When you disable loopback testing, the original duplex setting of the interface restores.

Loopback testing is a one-time operation, and is not recorded in the configuration file.

Setting the statistics polling interval on an Ethernet interface

To set the statistics polling interval on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Set the statistics polling interval on the Ethernet interface	flow-interval <i>interval</i>	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.

Configuring jumbo frame support on an Ethernet interface

An Ethernet interface may receive some frames larger than the standard Ethernet frame size (called "jumbo frames") during high-throughput data exchanges such as file transfers. Usually, an Ethernet interface discards jumbo frames. With jumbo frame support enabled, the interface can process frames larger than the standard Ethernet frame size yet within the specified range.

To configure jumbo frame support:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure jumbo frame support	jumboframe enable [<i>value</i>]	Optional. By default, the device allows jumbo frames within 10000 bytes to pass through all Ethernet interfaces.

Configuring link-up port auto power-down on an Ethernet interface

To save power, enable the link-up port auto power-down function on interfaces. With this function, a link-up port enters the power save mode if it has not received any packet for a certain period of time (depending on the specifications of the chip, and not configurable). When a packet arrives later, the port enters the normal state.

To enable link-up port auto power-down for an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Enable link-up port auto power-down	eee enable	Required Disabled by default

This feature is only available on the A5800AF-48G Switch (JG225A).

When you configure the **eee enable** command on a link-up port, the port automatically goes down and up.

Configuring a Layer 2 Ethernet interface

Layer 2 Ethernet interface configuration task list

Complete these tasks to configure an Ethernet interface operating in bridge mode:

Task	Remarks
Configuring a port group	Optional
Enabling link-down port auto power-down	Optional
Setting speed options for auto negotiation on an Ethernet interface	Optional
Configuring traffic storm protection	Optional
Enabling single-port loopback detection on an Ethernet interface	Optional
Enabling multi-port loopback detection	Optional
Setting the MDI mode of an Ethernet interface	Optional
Enabling bridging on an Ethernet interface	Optional
Testing the cable connection of an Ethernet interface	Optional
Configuring the connection mode of an Ethernet interface	Optional

Configuring a port group

Some interfaces on your switch may 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 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 an interface basis rather than on a port group basis. You can only view the settings in the view of each interface by use the **display current-configuration** or **display this** command.

To configure a port group:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a port group and enter port group view	port-group manual <i>port-group-name</i>	Required.
3. Assign Ethernet interfaces to the port group	group-member <i>interface-list</i>	Required.
4. 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.
5. Configure jumbo frame support on all Ethernet interfaces in the port group	jumboframe enable [<i>value</i>]	Optional. By default, the device allows jumbo frames within 10,000 bytes to pass through all Layer 2 Ethernet interfaces.

Enabling link-down port auto power-down

To save power, enable the link-down port auto power-down function on Ethernet interfaces. The switch stops providing power to the interface if the interface is in the down state for a certain period of time (depends on the specifications of the chip, and not configurable). In this case, the interface enters the power save mode. When the interface goes up, the switch provides power to the interface.


To enable link-down port auto power-down on one or multiple Ethernet interfaces:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Ethernet interface view interface <i>interface-type interface-number</i>	Use either command. Configured in Ethernet interface view, the setting takes effect on the interface only.
	Enter port group view port-group manual <i>port-group-name</i>	Configured in port group view, the setting takes effect on all ports in the port group.
3. Enable link-down port auto power-down on an Ethernet interface	port auto-power-down	Required. Disabled by default.

Setting speed options for auto negotiation on an Ethernet interface

As shown in Figure 4, speed auto negotiation enables an Ethernet interface to negotiate with its peer for the highest speed supported by both ends by default. You can narrow down the speed option list for negotiation.

Figure 4 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.

To configure an auto-negotiation transmission rate:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Set speed options for auto negotiation	speed auto { 10 100 1000 } *	Optional

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 traffic storm protection

A traffic storm occurs when a large amount of broadcast, multicast, or unknown unicast packets congest a network. The A5800 and the A5820X switches provide the following storm protection approaches:

- Storm suppression, which enables you to limit the size of monitored traffic passing through an Ethernet interface by setting a traffic threshold. When the traffic threshold is exceeded, the interface discards all exceeding traffic.
- Storm control, which enables you to shut down Ethernet interfaces or block traffic when monitored traffic exceeds the traffic threshold. It also enables an interface to send trap or log messages when monitored traffic reaches a certain traffic threshold, depending on your configuration.

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

Configuring storm suppression on an Ethernet interface

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.

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

To configure storm suppression on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Ethernet interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Use either command. Configured in Ethernet interface view, the setting takes effect on the interface only. Configured in port group view, the setting takes effect on all ports in the port group.
3. Set a broadcast suppression threshold	broadcast-suppression { <i>ratio</i> pps <i>max-pps</i> kbps <i>max-kbps</i> }	Optional. By default, all broadcast traffic can pass through.
4. Set a multicast suppression threshold	multicast-suppression { <i>ratio</i> pps <i>max-pps</i> kbps <i>max-kbps</i> }	Optional. By default, all multicast traffic can pass through.
5. Set a unknown unicast suppression threshold	unicast-suppression { <i>ratio</i> pps <i>max-pps</i> kbps <i>max-kbps</i> }	Optional. By default, all unknown unicast traffic can pass through.

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 is detected dropping 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, perform the **undo shutdown** command or disable the storm constrain function.

To configure the storm constrain function on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Set the traffic polling interval of the storm control module	storm-constrain interval <i>seconds</i>	Optional. 10 seconds by default.
3. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
4. 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 } <i>max-pps-values min-pps-values</i>	Required. Disabled by default.
5. Set the control action to take when monitored traffic exceeds the upper threshold	storm-constrain control { block shutdown }	Optional. Disabled by default.
6. 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.
7. 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.

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. It takes an interface at least one polling interval and at most two polling interval to take a storm control action.

Enabling single-port loopback detection on an Ethernet interface

If an interface receives a packet that it sent out, a loop occurs. Loops may cause broadcast storms, degrading network performance. For example, 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 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 3](#).

Table 3 Actions to take upon detection of a loop condition

Port type	Actions	
	No protective action is configured	A protective action is configured
Access interface	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Perform the configured protective action. Create traps and log messages. Delete all MAC address entries of the interface.
Hybrid or trunk interface	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> 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.

To configure single-port loopback detection:

To do...		Use the command...	Remarks
1. Enter system view		system-view	—
2. Enable global loopback detection		loopback-detection enable	Required. Disabled by default.
3. Set the loopback detection interval		loopback-detection interval-time <i>time</i>	Optional. 30 seconds by default.
4. Enter Ethernet interface view or port group view	Enter Ethernet interface view	interface <i>interface-type interface-number</i>	Use either command. To configure loopback detection on one interface, enter Ethernet interface view.
	Enter port group view	port-group manual <i>port-group-name</i>	To configure loopback detection on a group of Ethernet interfaces, enter port group view.

To do...	Use the command...	Remarks
5. Enable loopback detection on the interface	loopback-detection enable	Required. Disabled by default.
6. Enable loopback detection control	loopback-detection control enable	Optional. Disabled by default. This command is available only on hybrid and trunk interfaces.
7. 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 default VLAN.
8. 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.

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 perform the **loopback-detection control enable** command on the interface.

When you change the link type of an Ethernet interface with 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 out 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 may 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. 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 “[Enabling single-port loopback detection on an Ethernet interface.](#)”

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.

To configure multi-port loopback detection:

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

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

Optical interfaces do not support the MDI mode setting.

Use both crossover and straight-through Ethernet cables to connect copper Ethernet interfaces. To accommodate these two types of cables, a copper Ethernet interface can operate in one of the following MDI modes:

- Across mode
- Normal mode
- Auto mode

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

- 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 the local 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 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.

To set the MDI mode of an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. 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.

To enable bridging on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Enable bridging on the Ethernet interface	port bridge enable	Required Disabled by default

Testing the cable connection of an Ethernet interface

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.

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.

To test the cable connection of an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Test the cable connected to the Ethernet interface	virtual-cable-test	Required

Configuring the connection mode of an Ethernet interface

This feature is available only on the internal 10-GE interfaces of the A5800-48G-PoE+ Switch(JC101A), A5800-48G-PoE+ TAA Switch(JG242A), A5820X-14XG-SFP+ Switch(JC106A), and A5820X-14XG-SFP+ TAA Switch(JG259A).

When configuring an OAA application, you must set the 10-GE interface connecting the switch and the OAA card to operate in extended connection mode for normal communication.

To configure the connection mode of an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter 10-GE interface view	interface <i>interface-type interface-number</i>	—
3. Set the connection mode of the 10-GE interface	port connection-mode { extend normal }	Required. The default is normal mode.

Configuring a Layer 3 Ethernet interface

When an Ethernet interface operates in route mode, you can set the MTU for it.

The value of MTU affects the fragmentation and re-assembly of IP packets.

To set the MTU for an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Set the MTU	mtu <i>size</i>	Optional 1500 bytes by default

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 [<i>interface-type</i> [<i>interface-number</i>]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view

To do...	Use the command...	Remarks
Display the summary of an interface	display interface [<i>interface-type</i> [<i>interface-number</i>]] brief [{ begin exclude include } <i>regular-expression</i>] display interface [<i>interface-type</i>] brief down [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the statistics on the packets passing through a specific type of interfaces	display counters { inbound outbound } interface [<i>interface-type</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the statistics on the rate of the packets passing 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 [<i>interface-type</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display information about discarded packets on an interface	display packet-drop interface [<i>interface-type</i> [<i>interface-number</i>]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display summary information about discarded packets on all interfaces	display packet-drop summary [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the information about a manual port group or all manual port groups	display port-group manual [all name <i>port-group-name</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the information about the loopback function	display loopback-detection [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display information about storm control on interfaces	display storm-constrain [broadcast multicast unicast] [interface <i>interface-type</i> <i>interface-number</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Clear the statistics of an interface	reset counters interface [<i>interface-type</i> [<i>interface-number</i>]]	Available in user view
Clear the statistics of discarded packets on an interface	reset packet-drop interface [<i>interface-type</i> [<i>interface-number</i>]]	Available in user view

PFC configuration example

Network requirements


As shown in [Figure 6](#),

- Switch A connects to the FCoE HBA of the DC server through Ten-GigabitEthernet 1/0/1, and connects to Switch B through Ten-GigabitEthernet 1/0/2.
- The 802.1p priority of FCoE packets is 3. Configure the switches to guarantee no packet loss for FCoE packets forwarded among Switch B, Switch A, and DC server.

Suppose Switch A and Switch B support PFC, and Switch A and DC server support DCBX.

You should also configure DCBX on the port that connects Switch A to the FCoE HBA. This example describes configuring DCBX to advertise the PFC data. For more information about the functions of DCBX, see the chapter “LLDP configuration.”

Figure 6 Network diagram for PFC configuration



Configuration procedure

1. Configure Switch A
 - Configure DCBX to advertise the PFC data.

Enable LLDP globally.

```
<SwitchA> system-view
[SwitchA] lldp enable
```

Enable LLDP and DCBX TLV advertising on interface Ten-GigabitEthernet 1/0/1.

```
[SwitchA] interface ten-gigabitethernet 1/0/1
[SwitchA-Ten-GigabitEthernet1/0/1] lldp enable
[SwitchA-Ten-GigabitEthernet1/0/1] lldp tlv-enable dot1-tlv dcbx
```

Enable PFC in auto mode on interface Ten-GigabitEthernet 1/0/1, and enable PFC for 802.1p priority value 3 on the interface.

In addition to enabling PFC on the local end, this configuration also enables synchronizing the PFC data on the local end to the FCoE HBA of the remote end.

```
[SwitchA-Ten-GigabitEthernet1/0/1] priority-flow-control auto
[SwitchA-Ten-GigabitEthernet1/0/1] priority-flow-control no-drop dot1p 3
[SwitchA-Ten-GigabitEthernet1/0/1] qos trust dot1p
[SwitchA-Ten-GigabitEthernet1/0/1] quit
```

- Enable PFC on port Ten-GigabitEthernet 1/0/2, which the FCoE packets pass through.

Enable PFC on interface Ten-GigabitEthernet 1/0/2, and enable PFC for packets carrying 802.1p priority value 3 on the interface.

```
[SwitchA] interface ten-gigabitethernet 1/0/2
[SwitchA-Ten-GigabitEthernet1/0/2] priority-flow-control enable
[SwitchA-Ten-GigabitEthernet1/0/2] priority-flow-control no-drop dot1p 3
[SwitchA-Ten-GigabitEthernet1/0/2] qos trust dot1p
[SwitchA-Ten-GigabitEthernet1/0/2] quit
```

2. Configure Switch B

Enable PFC on interface Ten-GigabitEthernet 1/0/1, which the FCoE packets pass through, and enable PFC for packets carrying 802.1p priority value 3 on the interface.

```
<SwitchB> system-view
[SwitchB] interface ten-gigabitethernet 1/0/1
[SwitchB-Ten-GigabitEthernet1/0/1] priority-flow-control enable
[SwitchB-Ten-GigabitEthernet1/0/1] priority-flow-control no-drop dot1p 3
[SwitchB-Ten-GigabitEthernet1/0/1] qos trust dot1p
[SwitchB-Ten-GigabitEthernet1/0/1] quit
```

Loopback and null interface configuration

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 shut down manually.
- Assign a loopback interface an IP address with an all-F mask to save the IP address resources. 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.

Because of the benefits mentioned above, loopback interfaces are widely used in the following scenarios.

- 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. Therefore, when you configure a rule on an authentication or security server to permit or deny packets generated by a switch, you can simplify the rule by configuring it to permit or deny packets carrying the loopback interface address identifying the switch.

Note that, when you use a loopback interface address as the source address of IP packets, make sure that the route from the loopback interface to the peer is reachable by performing routing configuration. All data packets sent to the loopback interface are considered as packets sent to the switch itself, so the switch does not forward these packets.

- Because a loopback interface is always up, it can be used in dynamic routing protocols. For example, if no router ID is configured for a dynamic routing protocol, the highest loopback interface IP address is selected as the router ID. In BGP, to avoid BGP sessions being interrupted by physical port failure, you can use a loopback interface as the source interface of BGP packets.

Configuring a loopback interface

To configure a loopback interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a Loopback interface and enter Loopback interface view	interface loopback <i>interface-number</i>	—
3. Set a description for the loopback interface	description <i>text</i>	Optional. By default, the description of an interface is the interface name followed by the "Interface" string.
4. Shut down the loopback interface	shutdown	Optional. By default, a loopback interface is up once created.
5. Restore the default settings for the interface	default	Optional.

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

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. 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 leading to null interface 0), you can have all packets destined to the network segment 92.101.0.0/16 discarded.

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

Configuring null 0 interface

To enter null interface view:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. 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.
3. Set a description for the null interface	description <i>text</i>	Optional. By default, the description of an interface is the interface name followed by the "Interface" string.
4. Restore the default settings of the interface	default	Optional.

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 } <i>regular-expression</i>] display interface loopback <i>interface-number</i> [brief] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display information about the null interface	display interface null [brief [down]] [{ begin exclude include } <i>regular-expression</i>] display interface null 0 [brief] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Clear the statistics on a loopback interface	reset counters interface [loopback [<i>interface-number</i>]]	Available in user view
Clear the statistics on the null interface	reset counters interface [null [0]]	Available in user view

MAC address table configuration

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 an entry is found, the switch forwards the frame out of the outgoing port in the entry. If no entry is found, the switch broadcasts the frame out of all but the incoming port.

How a MAC address table entry is created

The entries in the MAC address table come from two sources: automatically learned by the switch and manually added by the administrator.

MAC address learning

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

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

1. Checks the source MAC address (MAC-SOURCE for example) of the frame.
2. Looks up the MAC address in the MAC address table.
3. If an entry is found, updates the entry. If no entry is found, 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 learning the source MAC address of a frame, the switch looks up the destination MAC address in the MAC address table. If an entry is found for the MAC address, the switch 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 where 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 manually add MAC address entries into the MAC address table of the switch to bind specific user devices to the port.

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 learned and may 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 is learned 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.

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 the outgoing interface indicated by the MAC address table entry.
- Broadcast mode: If the switch receives a frame with the destination address being all ones, or no entry is available for the destination MAC address, the switch broadcasts the frame to all 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 on a VLAN](#)
- [Configuring the aging timer for dynamic MAC address entries](#)
- [Configuring the MAC learning limit on ports](#)
- [Enabling MAC address roaming](#)

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

The MAC address table can contain only Layer 2 Ethernet ports and Layer 2 aggregate interfaces.

This chapter covers only 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 fence off 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.

To add, modify, or remove entries in the MAC address table in system view:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. 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.
	Configure blackhole MAC Address Table Entries mac-address blackhole mac-address vlan vlan-id	Make sure that you have created the VLAN and assign the interface to the VLAN.

To add or modify a MAC address table entry in interface view:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view	interface interface-type interface-number	—
3. 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

When you configure a static MAC address entry on an interface that belongs to a specific isolate-user-VLAN, you only need to specify the isolate-user-VLAN, instead of any secondary VLANs associated with the isolate-user-VLAN. For more information about isolate-user-VLANs, see the chapter “Isolate-user-VLAN configuration.”

Disabling MAC address learning on a VLAN

You may need to disable MAC address learning sometimes 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.

You may disable MAC address learning on a per-VLAN basis.

To disable MAC address learning on a VLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter VLAN view	vlan <i>vlan-id</i>	—
3. Disable MAC address learning on the VLAN	mac-address mac-learning disable	Required Enabled by default

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

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 could timely update to accommodate latest network changes.

Set the aging timer appropriately. A long aging interval may 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 may result in the removal of valid entries and unnecessary broadcasts, which may affect device performance.

To configure the aging timer for dynamic MAC address entries:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the aging timer for dynamic MAC address entries	mac-address timer { aging <i>seconds</i> no-aging }	Optional 300 seconds by default

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 the chances for a data packet to reach unintended destinations are reduced.

Configuring the MAC learning limit on ports

As the MAC address table is growing, the forwarding performance of your device may degrade. To prevent the MAC address table from getting so large that the forwarding performance is affected, you can limit the number of MAC addresses that can be learned on a port.

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
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view interface <i>interface-type</i> <i>interface-number</i>	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 member ports in the port group.
	Enter port group view port-group manual <i>port-group-name</i>	
3. Configure the MAC learning limit on the interface or port group	mac-address max-mac-count <i>count</i>	Required. No MAC learning limit is configured by default.

Layer 2 aggregate interfaces do not support the **mac-address max-mac-count** command.

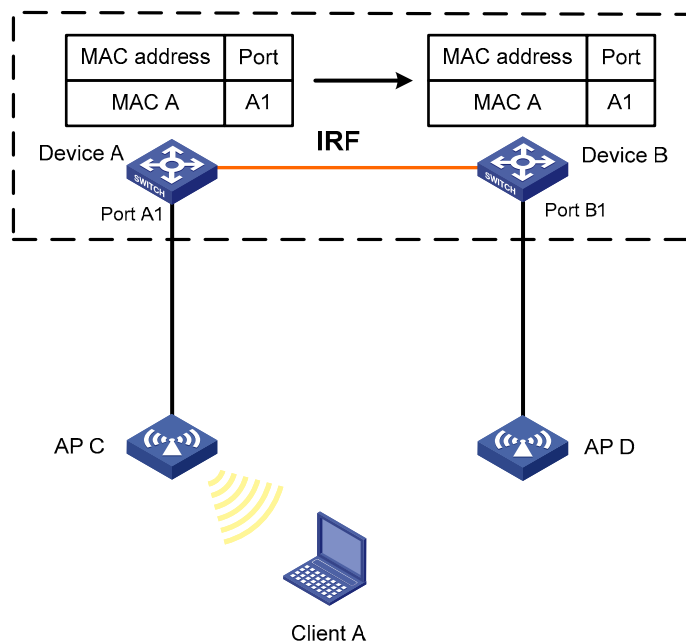
Do not configure the MAC learning limit on any member ports of an aggregation group. Otherwise, the member ports cannot be selected.

Enabling MAC address roaming

After you enable MAC address roaming on an IRF fabric, each member switch advertises learned MAC addresses to other member switches.

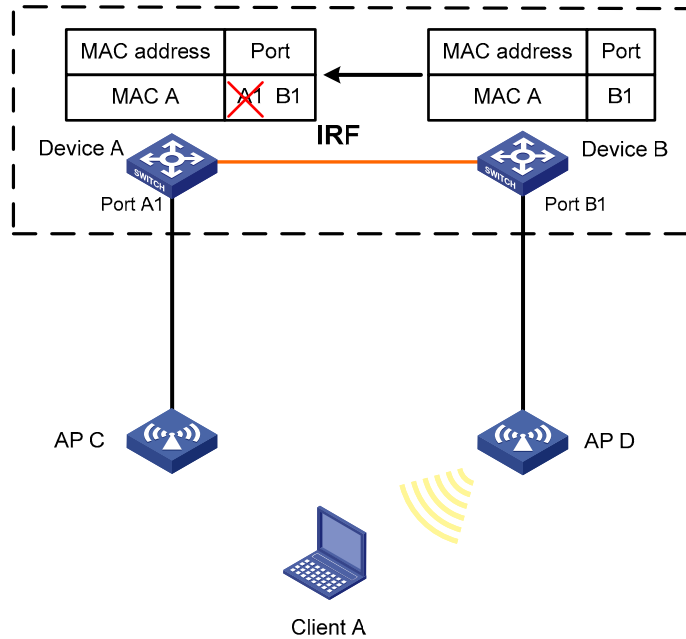
As Figure 7 shows, Device A and Device B form an IRF fabric enabled with MAC address roaming. They connect to AP C and AP D, respectively. When Client A associates with AP C, Device A learns the MAC address of Client A and advertises it to the member switch Device B.

Figure 7 MAC address tables of devices when Client A associates with AP C



If Client A roams to AP D, Device B learns the MAC address of Client A and advertises it to Device A to ensure service continuity for Client A, as shown in Figure 8.

Figure 8 MAC address tables of devices when Client A roams to AP D



To enable MAC address roaming:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable MAC address roaming	mac-address mac-roaming enable	Required Disabled by default

Displaying and maintaining MAC address tables

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

To monitor a network, you need to monitor users joining and leaving the network. Because a MAC address uniquely identifies a network user, you can monitor users 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 learn or delete MAC addresses. By analyzing these messages, the monitor end can monitor users accessing the network.

How MAC Information works

When a new MAC address is learned 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 is used up, the device sends the Syslog or trap messages to the monitor end immediately.

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

To enable MAC Information globally:

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

Enabling MAC Information on an interface

To enable MAC Information on an interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—

To do...	Use the command...	Remarks
3. Enable MAC Information on the interface	mac-address information enable { added deleted }	Required Disabled by default

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

Configuring MAC Information mode

To configure MAC Information mode:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. 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.

To set the interval for sending Syslog or trap messages:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Set the interval for sending Syslog or trap messages	mac-address information interval <i>interval-time</i>	Optional One second by default

Configuring the MAC Information queue length

To avoid losing user MAC address information, when the buffer storing user MAC address information is used up, 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.

To configure the MAC Information queue length:


To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the MAC Information queue length	mac-address information queue-length <i>value</i>	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 9 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

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 switched to other member ports automatically.

As shown in [Figure 10](#), 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 these three physical Ethernet links. At the same time, the three Ethernet links back up each other.

Figure 10 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 looks like 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.

Assign Layer 2 Ethernet interfaces only to a Layer 2 aggregation group.

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 may affect its aggregation state. Port configurations fall into the following classes:

- Port attribute configurations, including port rate, duplex mode, and link status (up/down), which are the most basic port configurations.
- Class-two configurations, as described in [Table 4](#). A member port can be placed in the selected state only if it has the same class-two configurations as the aggregate interface.

Table 4 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

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

Any class-two configuration change may affect the aggregation state of link aggregation member ports and ongoing traffic. To make 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 LACP enables dynamic aggregation of physical links. It uses LACPDU 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 5](#).

Table 5 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 above 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 TLV fields. This is how the LACP multi-active detection (MAD) mechanism of the IRF feature is implemented. A5800 and A5820X switch series can participate in LACP MAD either as an IRF member device or an intermediate device.

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

2. LACP priorities

LACP priorities have two types: system LACP priority and port LACP priority, as described in [Table 6](#).

Table 6 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 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 port LACP priority, the higher 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. Configure the LACP timeout interval either as the short timeout interval (1 second) or long timeout interval (30 seconds).

Link aggregation modes

Link aggregation has two modes: dynamic and static. Dynamic link aggregation uses LACP and static link aggregation does not. [Table 7](#) compares the two aggregation modes.

Table 7 A comparison between static and dynamic aggregation modes

Aggregation mode	LACP status on member ports	Pros	Cons
Static	Disabled	Aggregation is stable. The aggregation state of the member ports are not affected by the peer 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 instable. The aggregation state of member ports is susceptible to network changes.

In a dynamic link aggregation group:

- A selected port can receive and send LACPDU.
- An unselected port can receive and send LACPDU only if it is up and have 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

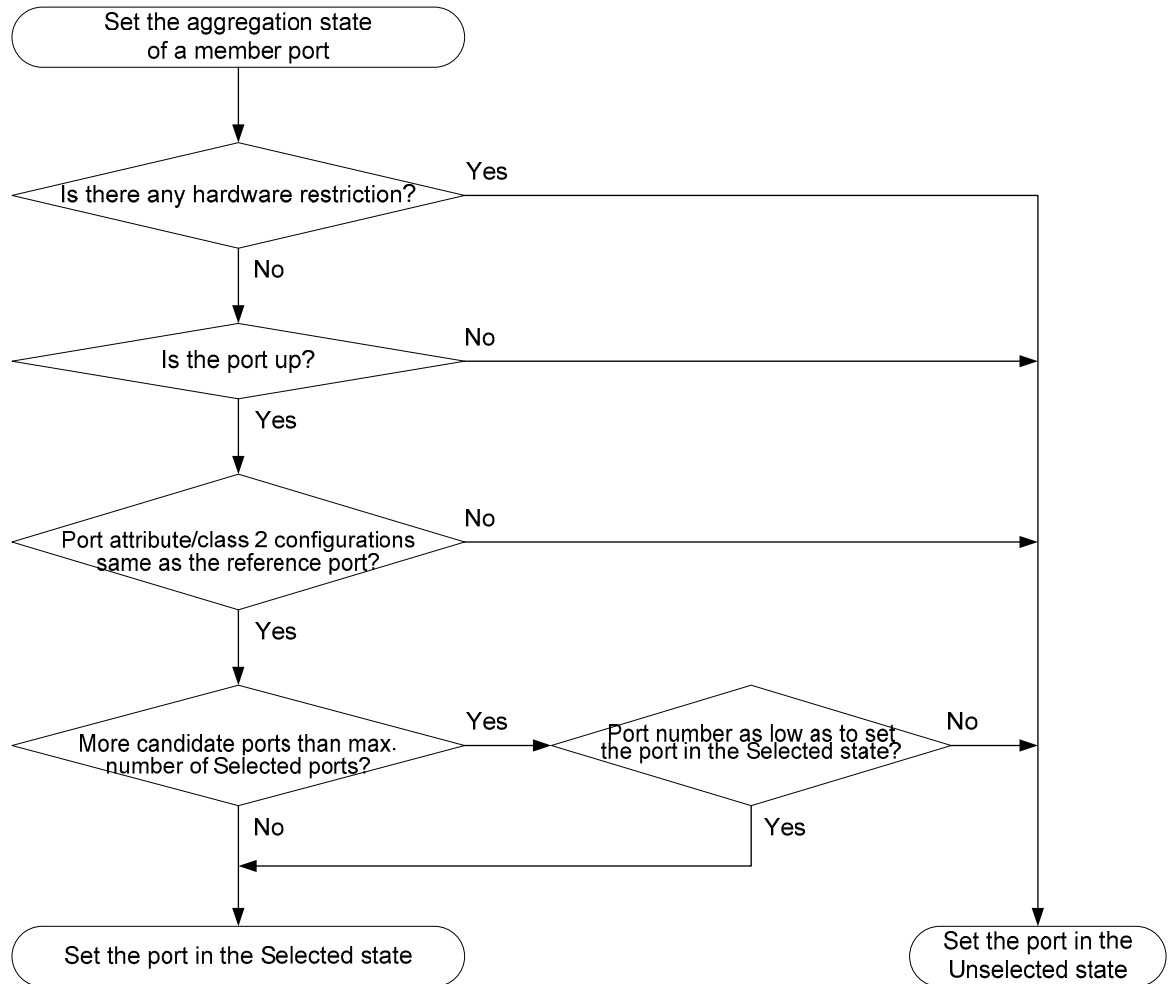
The system selects a reference port from the member ports that are in the up state and have the same class-two configurations as the aggregate interface.

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 out.

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, as shown in [Figure 11](#).

Figure 11 Set the aggregation state of a member port in a static aggregation group



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 joins the group is placed in the unselected state to avoid traffic interruption on the current selected ports. Avoid this situation, however, because it may cause the aggregation state of a port to change after a reboot.

Aggregating links in dynamic mode

LACP is enabled automatically 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

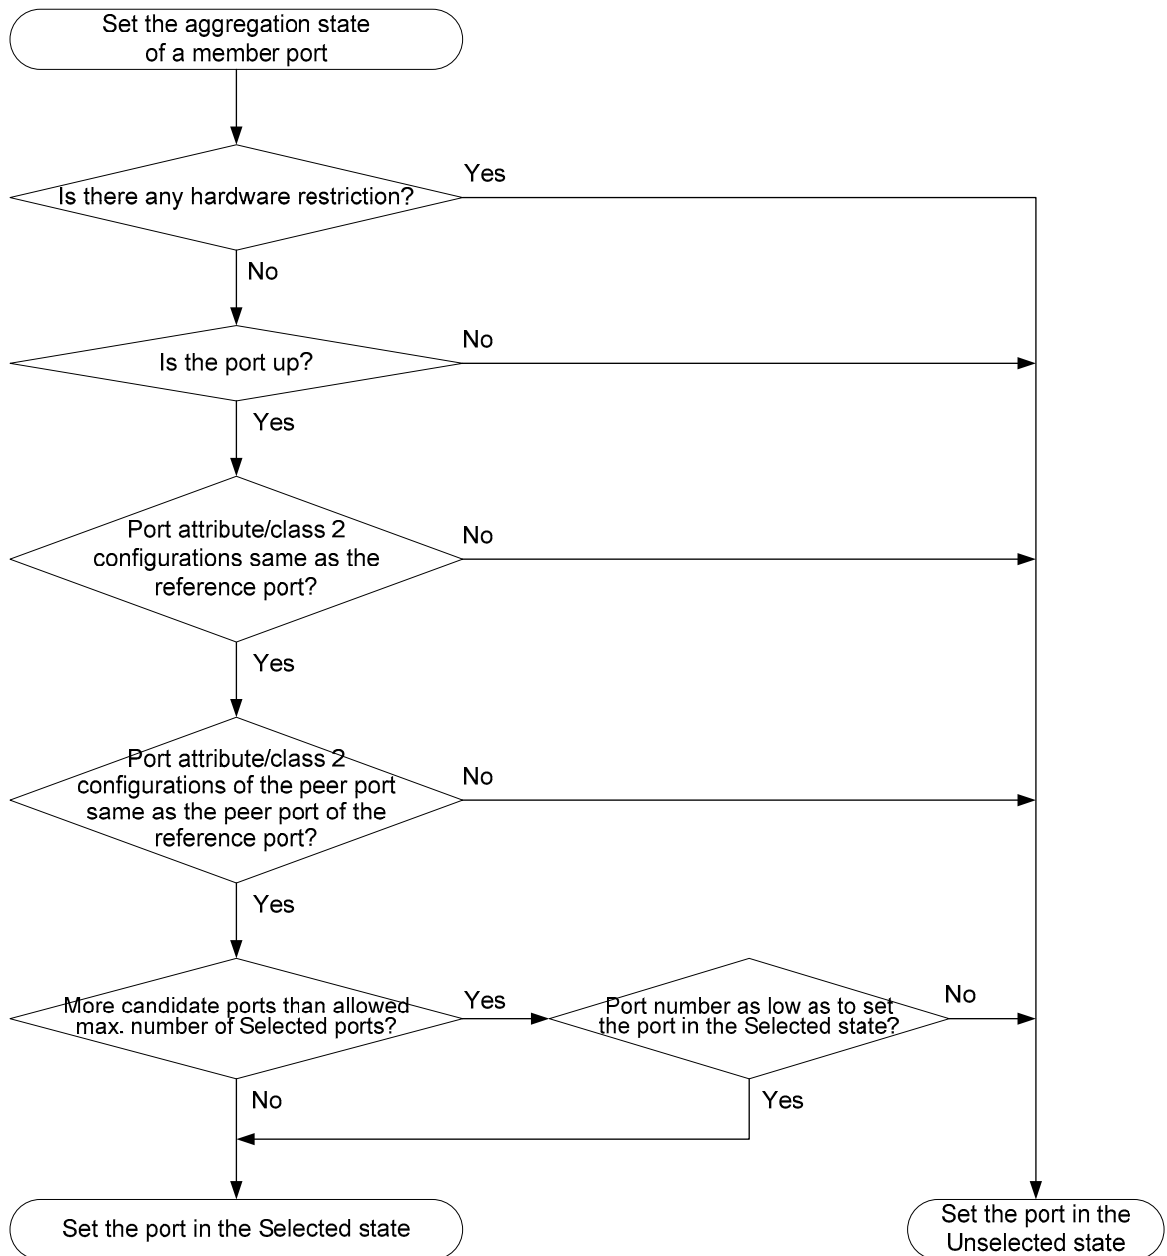
The local system (the actor) and the remote system (the partner) negotiate a reference port using the following workflow:

1. Compare the system ID (comprising the system LACP priority and the system MAC address). The system with the lower LACP priority value wins out. If they are the same, 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 out. 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 as shown in Figure 12.

Figure 12 Set the state of a member port in a dynamic aggregation group



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

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 also 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 may be load-shared across the selected member ports based on a set of criteria, depending on your configuration.

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
	Configuring a dynamic aggregation group
	Select either task
Configuring an aggregate interface	Configuring the description of an aggregate interface
	Enabling link state traps for an aggregate interface
	Setting the minimum number of selected ports for an aggregation group
	Shutting down an aggregate interface
	Restoring the default settings for an aggregate interface
Configuring load sharing for link aggregation groups	Configuring load sharing criteria for link aggregation groups
	Enabling local-first load sharing for link aggregation
Enabling link-aggregation traffic redirection	
	Optional

Configuring an aggregation group

△ CAUTION:

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

You cannot assign a port to a Layer 2 aggregation group if any of the features listed in [Table 8](#) is configured on the port.

Table 8 Features incompatible with Layer 2 aggregation groups

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

If a port is used as a reflector port for port mirroring, do not assign it to any aggregation group. For more information about reflector ports, see *Network Management and Monitoring Configuration Guide*.

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.

Configuring a static aggregation group

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

To configure a Layer 2 static aggregation group:

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

Configuring a dynamic aggregation group

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

To configure a Layer 2 dynamic aggregation group:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Set the system LACP priority	lacp system-priority <i>system-priority</i>	Optional. By default, the system LACP priority is 32,768. Changing the system LACP priority may affect the aggregation state of the ports in a dynamic aggregation group.
3. Create a Layer 2 aggregate interface and enter the Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	Required. When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
4. Configure the aggregation group to work in dynamic aggregation mode	link-aggregation mode dynamic	Required. By default, an aggregation group works in static aggregation mode.
5. Exit to system view	quit	—
6. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	Required.
7. Assign the Ethernet interface to the aggregation group	port link-aggregation group <i>number</i>	Repeat these two steps to assign more Layer 2 Ethernet interfaces to the aggregation group.
8. Assign the port an LACP priority	lacp port-priority <i>port-priority</i>	Optional. By default, the LACP priority of a port is 32,768. Changing the LACP priority of a port may affect the aggregation state of the ports in the dynamic aggregation group.
9. 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

Most 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

Configure the description of an aggregate interface for administration purposes such as describing the purpose of the interface.

To configure the description of an aggregate interface:

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

Enabling link state traps for an aggregate interface

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*.

To enable link state traps on an aggregate interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. 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.
3. Enter Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	—
4. Enable link state traps for the aggregate interface	enable snmp trap updown	Optional. Enabled by default.

Setting the minimum number of selected ports for an aggregation group

△ CAUTION:

- If you set this minimum threshold for a static aggregation group, you must also make the same setting for its peer aggregation group to guarantee correct aggregation.
- Configuring the minimum number of selected ports required to bring up an aggregation group may cause all member ports in the current aggregation group to become unselected.

The bandwidth of an aggregate link increases along with the number of selected member ports. To avoid congestion caused by insufficient selected ports on an aggregate link, you can set the minimum number of selected ports required for bringing up the specific aggregate interface.

This minimum threshold setting affects the aggregation state of both aggregation member ports and the aggregate interface in the following ways:

- All member ports change to the unselected state and the link of the aggregate interface goes down, when the number of member ports eligible for being selected is smaller than the minimum threshold.
- When the minimum threshold is reached, the eligible member ports change to the selected state, and the link of the aggregate interface goes up.

To set the minimum number of selected ports for an aggregation group:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	—
3. Set the minimum number of selected ports for the aggregation group	link-aggregation selected-port minimum <i>number</i>	Required Not specified 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.

To shut down an aggregate interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	—

To do...	Use the command...	Remarks
3. Shut down the aggregate interface	shutdown	Required. By default, aggregate interfaces are up.

Restoring the default settings for an aggregate interface

To restore the default settings for an aggregate interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter aggregate interface view	interface bridge-aggregation <i>interface-number</i>	—
3. Restore the default settings for the aggregate interface	default	Required

Configuring load sharing for link aggregation groups

△ CAUTION:

By default, an aggregation group uses the global link-aggregation load sharing criterion or criteria. 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.

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 is available, the group uses the global load sharing criteria.

Configuring the global link-aggregation load sharing criteria

To configure global link-aggregation load sharing criteria:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. 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. The default global aggregation load sharing criteria are ingress port, source MAC address and destination MAC address for Layer 2 packet types, such as ARP, and source and destination IP addresses for Layer 3 packet types, such as IP packet

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

To configure load sharing criteria for a link aggregation group:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter aggregate interface view	interface bridge-aggregation <i>interface-number</i>	—
3. 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.

Set the following group-specific load sharing criteria:

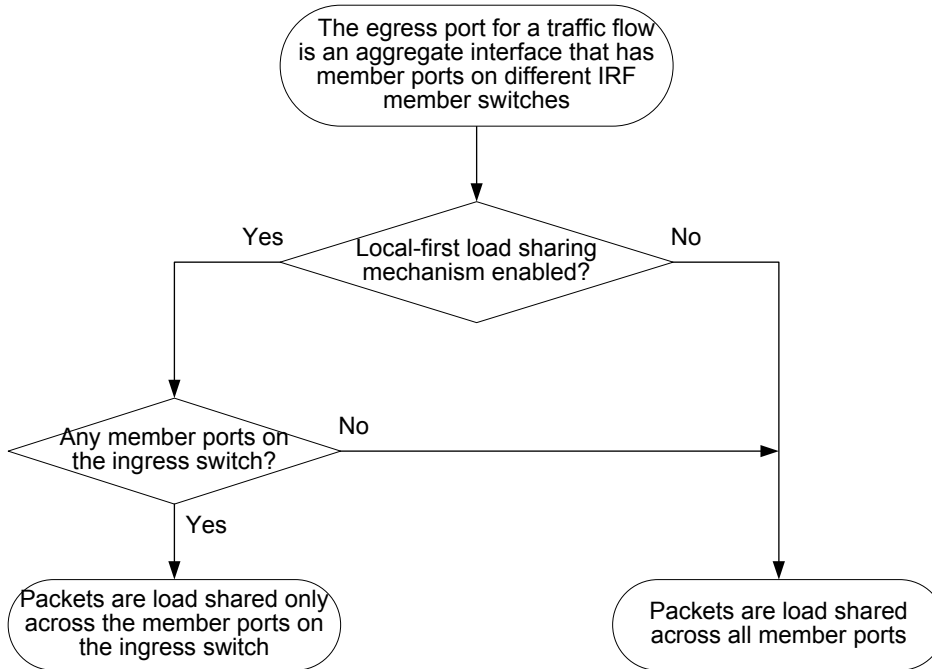
- Source IP address
- Destination IP address
- Source IP address and destination IP address
- Source MAC address
- Destination MAC address
- Destination MAC address and source MAC address

Enabling local-first load sharing for link aggregation

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

When you aggregate ports on different member switches in an IRF virtual device, you can use local-first load sharing to reduce traffic on IRF links, as shown in [Figure 13](#). For more information about IRF, see the *IRF Configuration Guide*.

Figure 13 Local-first link-aggregation load sharing



To enable local-first load sharing for link aggregation:

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

Enabling link-aggregation traffic redirection

△ 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 virtual device that adopts the ring connection, slight packet loss can occur when the IRF member device 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 devices (see [“Enabling local-first load sharing for link aggregation”](#)).

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

To enable link-aggregation traffic redirection:

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

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 } <i>regular-expression</i>] display interface bridge-aggregation <i>interface-number</i> [brief] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the local system ID	display lacp system-id [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the global or group-specific link-aggregation load sharing criteria	display link-aggregation load-sharing mode [interface [bridge-aggregation <i>interface-number</i>]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display detailed link aggregation information on link aggregation member ports	display link-aggregation member-port [<i>interface-list</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the summary of all aggregation groups	display link-aggregation summary [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display detailed information about a specific or all aggregation groups	display link-aggregation verbose [bridge-aggregation [<i>interface-number</i>]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Clear LACP statistics for a specific or all link aggregation member ports	reset lacp statistics [interface <i>interface-list</i>]	Available in user view
Clear statistics for a specific or all aggregate interfaces	reset counters interface [bridge-aggregation [<i>interface-number</i>]]	Available in user view

Ethernet link aggregation configuration examples

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 14:

- 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. Then 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 14 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.

This configuration automatically propagates to all 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 15:

- 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. Then 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 15 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.

This configuration automatically propagates to all 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.

Port isolation configuration

Assigning access ports to different VLANs is a typical way to isolate Layer 2 traffic for data privacy and security, but this approach is VLAN-resource demanding. 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 A5800 and the A5820X switch series support only one isolation group called isolation group 1. This isolation group is created automatically and cannot be deleted. There is no limit on the number of member ports.

Configuring the isolation group

To assign a port to the isolation group:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view or port group view	Enter Ethernet interface view interface <i>interface-type</i> <i>interface-number</i>	Required. Use one of the commands. <ul style="list-style-type: none">To assign an Ethernet port to the isolation group, enter Ethernet interface view.
	Enter Layer-2 aggregate interface view interface bridge-aggregation <i>interface-number</i>	<ul style="list-style-type: none">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.
	Enter port group view port-group manual <i>port-group-name</i>	<ul style="list-style-type: none">To assign multiple Ethernet ports to the isolation group in bulk, enter port group view.
3. Assign the port or ports to the isolation group	port-isolate enable	Required. The isolation group does not contain any ports by default.

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	<code>display port-isolate group [{ begin exclude include } regular-expression]</code>	Available in any view

Port isolation configuration example


Network requirements

As shown in [Figure 16](#):

- 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 16 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 STP eliminates Layer 2 loops by selectively blocking redundant links in a network, and in the mean time, allows for link redundancy.

Like many other protocols, STP evolves as the network grows. The later versions of STP are the RSTP and the MSTP. This chapter describes the features of STP, RSTP, and MSTP.

Why STP

STP was developed based on the 802.1d standard of IEEE to eliminate loops at the data link layer in a LAN. Devices running 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 decreased performance of network devices caused by duplicate packets received.

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 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 for calculating a spanning tree and maintaining the spanning tree topology.
- TCN BPDUs, used for notifying the concerned devices of network topology changes, if any.

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

- Root bridge ID: consisting of the priority and MAC address of the root bridge.
- Root path cost: the 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: the delay used by STP bridges to transition port state.

Basic concepts in STP

Root bridge

A tree network must have a root bridge.

There is only one root bridge in the entire network. The root bridge is not fixed, but can change along with changes of the network topology.

Upon initialization of a network, each device generates and sends out configuration BPDUs periodically with itself as the root bridge. After network convergence, only the root bridge generates and sends out configuration BPDUs at a certain interval, 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 9 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 17, 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 17 A schematic diagram of designated bridges and designated ports



Path cost

Path cost is a reference value used for link selection in STP. By calculating path costs, STP selects relatively robust links and blocks redundant links, and finally prunes the network into a loop-free tree.

How STP works

STP has the following workflow:

1. Initial state

Upon initialization of a device, each port generates a BPDU with itself as the root bridge, in which the root path cost is 0, 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 10 describes the process of selecting the root port and designated ports.

Table 10 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 11.
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 rest ports. <ul style="list-style-type: none">• 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 is to be defined, and acts depending on the comparison result: <ul style="list-style-type: none">• 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 out 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.

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

Table 11 Selection of the optimum configuration BPDU

Step	Actions
1	Upon receiving a configuration BPDU on a port, the device performs the following: <ul style="list-style-type: none">• 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 does not process the configuration BPDU of this port.• 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 ports and chooses the optimum configuration BPDU.

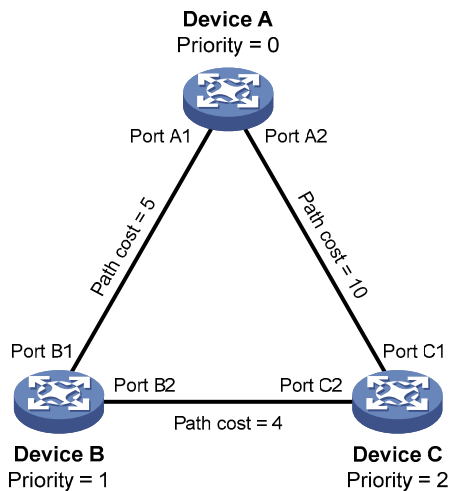
The following are the principles of configuration BPDU comparison:

- The configuration BPDU that has 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 containing the smallest ID wins out.

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

The following is an example of how the STP algorithm works.

Figure 18 Network diagram for the STP algorithm



As shown in [Figure 18](#), 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 12 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}

In Table 12, 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 13 Comparison process and result on each device

Device	Comparison process	Configuration BPDU on ports after comparison
Device A	<ul style="list-style-type: none"> 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 out configuration BPDUs. 	<ul style="list-style-type: none"> Port A1: {0, 0, 0, Port A1} Port A2: {0, 0, 0, Port A2}
Device B	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Port B1: {0, 0, 0, Port A1} Port B2: {1, 0, 1, Port B2}


Device	Comparison process	Configuration BPDU on ports after comparison
	<ul style="list-style-type: none"> • 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 out the calculated configuration BPDU. 	<ul style="list-style-type: none"> • Root port (Port B1): {0, 0, 0, Port A1} • Designated port (Port B2): {0, 5, 1, Port B2}

Device	Comparison process	Configuration BPDU on ports after comparison
Device C	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Port C1: {0, 0, 0, Port A2} Port C2: {1, 0, 1, Port B2}
	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Root port (Port C1): {0, 0, 0, Port A2} Designated port (Port C2): {0, 10, 2, Port C2}
	<ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Port C1: {0, 0, 0, Port A2} Port C2: {0, 5, 1, Port B2}
	<ul style="list-style-type: none"> 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 spanning tree calculation process is triggered by a new event, for example, the link between Device B and Device C is down. 	<ul style="list-style-type: none"> Blocked port (Port C1): {0, 0, 0, Port A2} Root port (Port C2): {0, 5, 1, Port B2}

In [Table 13](#), 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 13, a spanning tree with Device A as the root bridge is established, and the topology is shown in Figure 19.

Figure 19 Topology of the final calculated spanning tree



The spanning tree calculation process in this example is only a simplified process.

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 it is the root port that 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 out 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 out 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 out 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 may occur.

STP timers

STP calculation involves the following timing parameters.

- Forward delay: Specifies the delay time for device 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 right away, a temporary loop is likely to 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 transitioning 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 quicker under certain conditions than in STP.

In RSTP, a newly elected root port can enter the forwarding state rapidly if this condition is met: 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 this condition is met: the designated port is an edge port (a port 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 transitioning 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 the support for rapid network convergence, it allows data flows of different VLANs to be forwarded along separate paths, providing a better load sharing mechanism for redundant links. For more information about VLANs, see the chapter “VLAN configuration.”

MSTP includes 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 containing multiple spanning trees that are independent of one another.
- MSTP prunes a loop network into a loop-free tree avoiding proliferation and endless cycling of packets in a loop network. In addition, it provides multiple redundant paths for data forwarding supporting load balancing of VLAN data.
- MSTP is compatible with STP and RSTP.

Basic concepts in MSTP

Figure 20 Basic concepts in MSTP


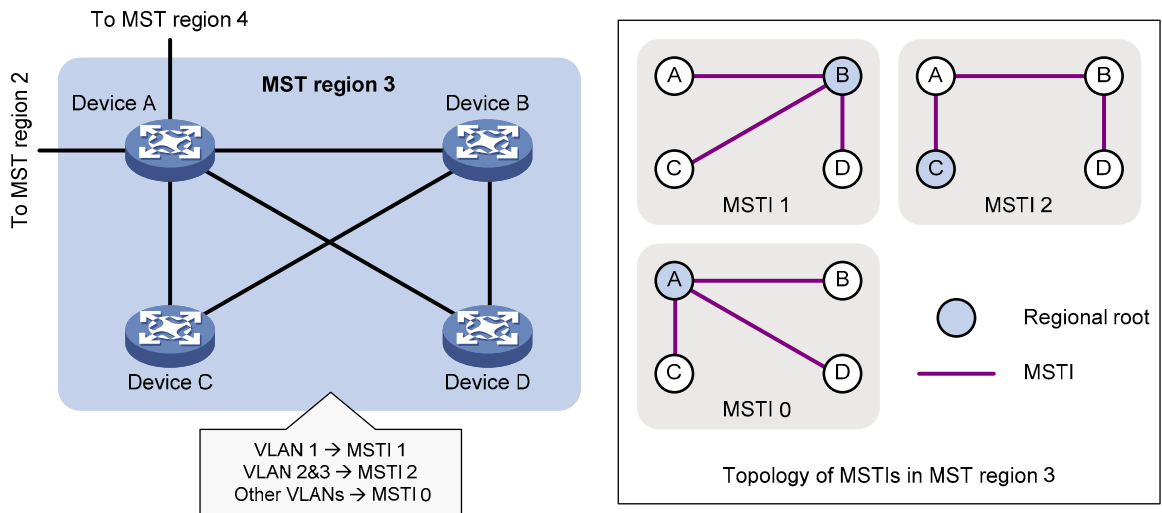


Figure 21 Network diagram and topology of MST region 3



As shown in Figure 20, a switched network comprises four MST regions, and each MST region comprises four devices running MSTP. Figure 21 shows the networking topology of MST region 3. This section describes some basic concepts of MSTP.

MST region

An 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 20](#), 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 spanning trees in an MST region, and each spanning tree is independent of another and maps to the specific VLANs. Each spanning tree is referred to as a “multiple spanning tree instance (MSTI).”

In [Figure 21](#), 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 21](#), 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 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 20](#) represent the CST.

IST

An 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 20](#), MSTI 0 is the IST in MST region 3.

CIST

Jointly constituted by ISTs and the CST, the 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 20](#), 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 may have different regional roots.

For example, in MST region 3 in [Figure 21](#), 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

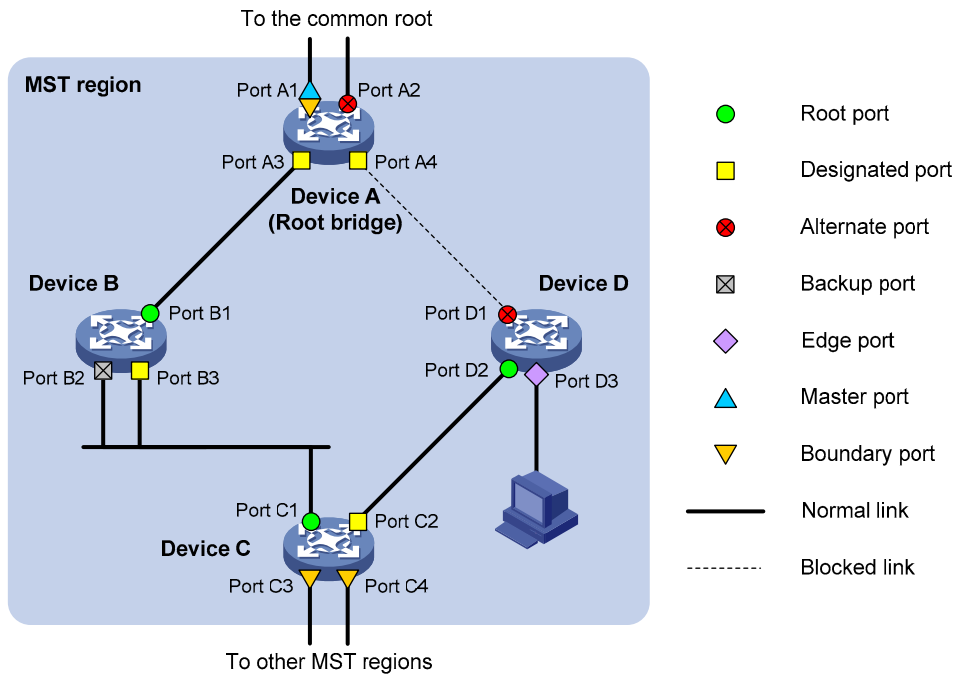
The common root bridge is the root bridge of the CIST.

In [Figure 20](#), 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 22](#), 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 22 Port roles



MSTP calculation involves these 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 may be in one of the following states:

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

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

A port state is not exclusively associated with a port role. [Table 14](#) lists the port states supported by each port role (“√” indicates that the port supports this state, and “—” indicates that the port does not support this state).

Table 14 Port states supported by different port roles

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

How MSTP works

MSTP divides an entire Layer 2 network into multiple MST regions, which are interconnected by a calculated CST. Inside an MST region, multiple spanning trees are calculated, each being an MSTI. Among these MSTIs, MSTI 0 is the IST. Similar to STP, MSTP uses configuration BPDUs to calculate spanning trees. The only difference between the two protocols is that an MSTP BPDU carries the MSTP configuration on the device from which this 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, and, 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. MSTP performs a separate calculation process, which is similar to spanning tree calculation in STP, for each spanning tree. 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. STP and RSTP protocol packets can be recognized by devices running MSTP and used for spanning tree calculation.

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:

- 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 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	
Configuring the leaf nodes	Enabling the output of port state transition information	Optional
	Enabling the MSTP feature	Required
	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 TC snooping	Optional	
Configuring protection functions	Optional	

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, make 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: service loopback, RRPP, Smart Link, and BPDU tunnel.

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 its MSTP configurations for participating 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.

To configure an MST region:

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

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, the same VLAN-to-instance mapping entries in the MST region and the same MST region revision level, and they are interconnected via a physical link.

The configuration of MST region-related parameters, especially the VLAN-to-instance mapping table, will cause MSTP to launch a new spanning tree calculation process, which may result in network topology instability. To reduce the possibility of topology instability caused by configuration, MSTP does not immediately launch a new spanning tree calculation process when 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 in the case that MSTP is disabled.

Configuring the root bridge or a secondary root bridge

MSTP can determine the root bridge of a spanning tree through MSTP calculation. Also, you have the option of specifying the current device as the root bridge or a secondary root bridge using the commands provided by the system.

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 being the root bridge or a secondary root bridge of another MSTI. However, the same 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 to be 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 secondary root bridge will not 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

To configure the current device as the root bridge of a specific spanning tree:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the current device as the root bridge of a specific spanning tree	stp [instance <i>instance-id</i>] 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

To configure the current device as a secondary root bridge of a specific spanning tree:

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

After specifying the current device as the root bridge or a secondary root bridge, you cannot change the priority of the device.

You can also 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 is unable to 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 out STP BPDUs,

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

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

To configure the MSTP work mode:

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

Configuring the priority of a device

△ CAUTION:

- After configuring a device as the root bridge or a secondary root bridge, you cannot change the priority of the device.
- 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.

Device priorities participate 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 value indicates a higher priority. By setting the priority of a device to a low value, you can 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.

To configure the priority of a device in a specified MSTI:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the priority of the current device in a specified MSTI	stp [instance <i>instance-id</i>] priority <i>priority</i>	Required 32,768 by default

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.

The regional root bridge always sends a configuration BPDU with 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 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 take part in spanning tree calculation, and the size of the MST region is confined.

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.

To configure the maximum number of hops of an MST region:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the maximum hops of the MST region	stp max-hops <i>hops</i>	Required 20 by default

Configuring the network diameter of a switched network

Any two terminal devices in a switched network are interconnected 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.

To configure the network diameter of a switched network:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the network diameter of the switched network	stp bridge-diameter <i>diameter</i>	Required 7 by default

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 also 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 timers of MSTP

STP calculation involves the following timing parameters.

- Forward delay: Determines the time interval of 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 determines whether a configuration BPDU received on a port has expired based on the max age timer. 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, the settings of the hello time, forward delay and max age timers must meet the following formulas:

- $2 \times (\text{forward delay} - 1 \text{ second}) \leq \text{max age}$
- $\text{Max age} \leq 2 \times (\text{hello time} + 1 \text{ second})$

HP does not recommend that you set the timers manually. 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.

To configure the timers of MSTP:

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

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 can be introduced. If the forward delay is too long, it may take a long time for the network to converge. HP recommends that you use the default setting.

An appropriate hello time enables the device to timely detect link failures on the network without using excessive network resources. If the hello time is set too long, the device will take 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 send repeated configuration BPDUs frequently, which adds to the device burden and causes waste of network resources. HP recommends that you use the default setting.

If the max age time is too short, the network devices will frequently launch spanning tree calculations and may take network congestion as a link failure. If the max age is too long, the network may fail to timely detect link failures and fail to timely launch 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 \times 3 \times 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 check 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 may 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 the network resources. In a very stable network, you can avoid such unwanted spanning tree calculations by setting the timeout factor to 5, 6, or 7.

To configure the timeout factor:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the timeout factor of the device	stp timer-factor <i>factor</i>	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.

To configure the maximum rate of a port or a group of ports:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	Required. Use either command.
	Enter port group view	
3. Configure the maximum rate of the ports	stp transmit-limit <i>limit</i>	Required. 10 by default.

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 instable. 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 know whether a port is directly connected to a terminal, you need to manually configure the port to be an edge port. After that, this port can transition rapidly from the blocked state to the forwarding state without delay.

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

To specify a port or a group of ports as edge port or ports:

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

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

With BPDU guard disabled, when a port set as an edge port receives a BPDU from another port, it will become a non-edge port again. 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 fast while ensuring network security. Otherwise, the port can be blocked and changes to the forwarding state after a period twice the forward delay. In the waiting period, service traffic is interrupted.

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

Configuring path costs of ports

△ CAUTION:

If you change the standard that the device uses in calculating the default path costs, you restore the path costs to the default. Also, when the path cost of a port changes, MSTP re-calculates the role of the port and initiates a state transition.

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.

The device can automatically calculate the default path cost; alternatively, you can also configure the path cost for ports.

Make the following configurations on the leaf nodes only.

Specifying a standard that the device uses when calculating the default path cost

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.

To specify a standard for the device to use when calculating the default path cost:

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

To do...	Use the command...	Remarks
2. Specify a standard for the device to use when calculating the default path costs of its ports	<code>stp pathcost-standard { dot1d-1998 dot1t legacy }</code>	Optional. By default, the device calculates the default path cost for ports based on a private standard.

Table 15 shows the mappings between the link speed and the path cost.

Table 15 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	—	65,535	200,000,000	200,000
10 Mbps	Single Port	100	2,000,000	2000
	Aggregate interface containing 2 selected ports		1,000,000	1800
	Aggregate interface containing 3 selected ports		666,666	1600
	Aggregate interface containing 4 selected ports		500,000	1400
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
	Aggregate interface containing 4 selected ports		500	1

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

To configure the path cost of ports:

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

Configuration example

Specify that the device use 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.

To configure the priority of a port or a group of ports:

To do...	Use the command...	Remarks	
1. Enter system view	system-view	—	
2. Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view	port-group manual <i>port-group-name</i>	Required.

To do...	Use the command...	Remarks
3. Configure the port priority	stp [instance <i>instance-id</i>] port priority <i>priority</i>	Required. 128 for all ports by default.

When the priority of a port is changed, MSTP will re-calculate the role of the port and initiate a state transition.

A lower priority value indicates a higher priority. If you configure the same priority value for all 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 directly connecting 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.

To configure the link type of a port or a group of ports:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view interface <i>interface-type</i> <i>interface-number</i>	Required. Use either command.
	Enter port group view port-group manual <i>port-group-name</i>	
3. 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.

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 a port is configured 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 may 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 it will send based on the recognized format.

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 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.

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

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

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 receives a packet in a format different from the specified type, the port will become a designated port and remain 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. The closed ports can be restored only by the network administrators.

Enabling the output of port state transition information

A large-scale, MSTP-enabled network can have many MSTIs, and ports may 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 so as to monitor the port states in real time.

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

To enable output of port state transition information:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable output of port state transition information	stp port-log { all instance <i>instance-id</i> }	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.

To enable the MSTP feature:

To do...		Use the command...	Remarks
1. Enter system view		system-view	—
2. Enable the MSTP feature globally		stp enable	Required. MSTP is globally disabled by default.
3. Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view	port-group manual <i>port-group-name</i>	
4. 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.

In system view, you can use the **stp enable** or **undo stp enable** command to enable or disable STP globally.

Use the **undo stp enable** command to disable the MSTP feature for certain ports so that they will not take part 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 running MSTP (or RSTP) connects to a device running STP, this port will migrate to the STP-compatible mode automatically. However, it will not be able to migrate automatically back to the MSTP (or RSTP) mode, but will remain working in the STP-compatible mode under the following circumstances:

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

By then, you can perform an mCheck operation to force the port to migrate to the MSTP (or RSTP) mode.

Perform mCheck on a port through the following two approaches, which lead to the same result.

Performing mCheck globally

To perform global mCheck:

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

Performing mCheck in interface view

To perform mCheck in interface view:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or Layer 2 aggregate interface view	interface <i>interface-type</i> <i>interface-number</i>	—
3. Perform mCheck	stp mcheck	Required

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, interconnected devices are in the same region only when the MST region-related configurations (region name, revision level, VLAN-to-instance mappings) on them are identical. An MSTP-enabled device identifies devices in the same MST region by checking the configuration ID in BPDU packets. The configuration ID includes the region name, revision level, configuration digest that is in 16-byte length and is the result calculated via the HMAC-MD5 algorithm based on VLAN-to-instance mappings.

Since MSTP implementations vary with vendors, the configuration digests calculated using private keys is different. The different vendors' devices in the same MST region cannot communicate with each other.

Enabling the digest snooping feature on the port connecting the local device to a third-party device in the same MST region can make the two devices communicate with each other.

Before enabling 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.

To configure digest snooping:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view interface <i>interface-type</i> <i>interface-number</i>	Required. Use either command.
	Enter port group view port-group manual <i>port-group-name</i>	
3. Enable digest snooping on the interface or port group	stp config-digest-snooping	Required. Disabled by default.
4. Return to system view	quit	—
5. Enable global digest snooping	stp config-digest-snooping	Required. Disabled by default.

With the digest snooping feature enabled, comparison of configuration digest is not needed for in-the-same-region check, 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 removing of the current region configuration using the **undo stp region-configuration** command are not allowed. You can only modify the region name and revision level.

You must enable digest snooping both globally and on associated ports to make it take effect. HP recommends that you enable digest snooping on all associated ports first and then globally, making the configuration take effect on all configured ports and reducing 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 traffic interruption, do not configure digest snooping when the network works well.

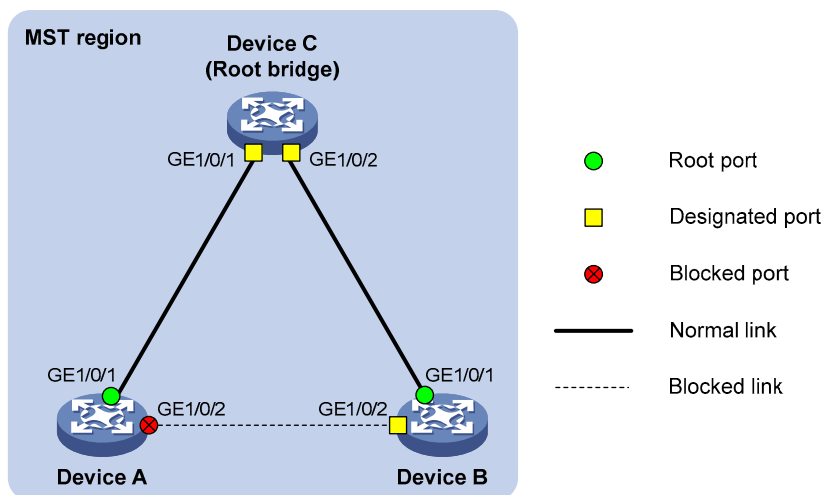
Digest snooping configuration example

1. Network requirements

As shown in Figure 23:

- 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 Device A's and Device B's ports that connect Device C, so that the three devices can communicate with one another.

Figure 23 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 downstream device's root port 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 24 shows the rapid state transition mechanism on MSTP designated ports.

Figure 24 Rapid state transition of an MSTP designated port

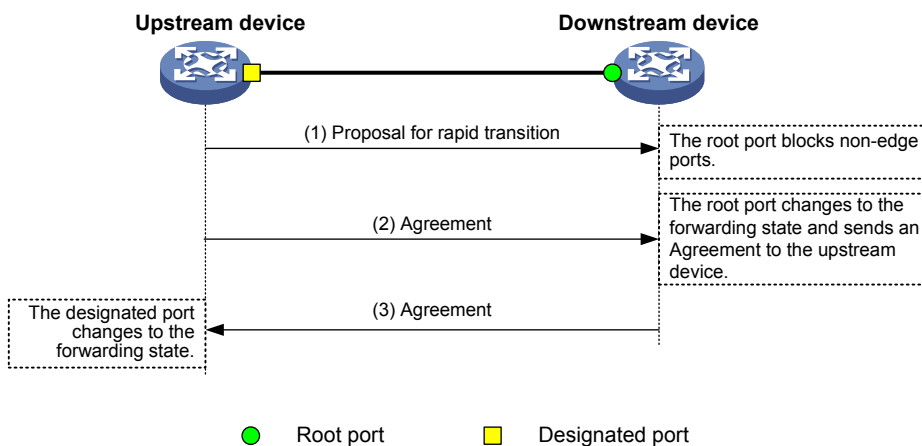
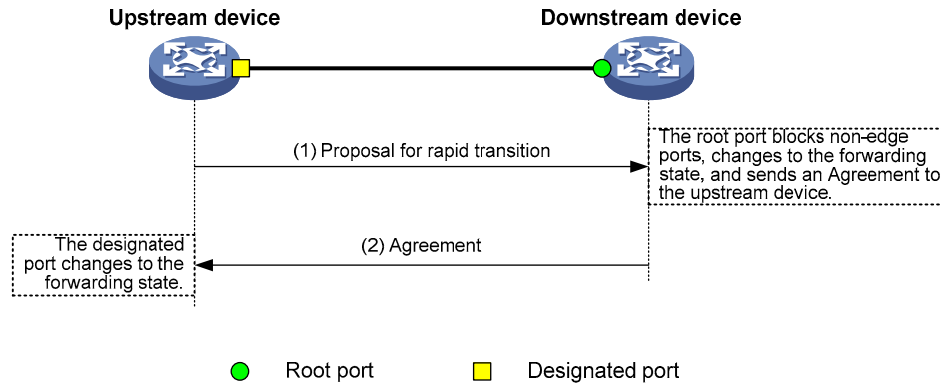


Figure 25 shows rapid state transition of an RSTP designated port.

Figure 25 Rapid state transition of an RSTP designated port



If the upstream device is a third-party device, the rapid state transition implementation may 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 to a third-party upstream device supporting 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.

To configure no agreement check:

To do...	Use the command...	Remarks	
1. Enter system view	system-view	—	
2. Enter interface or port group view	Enter Ethernet interface view or Layer 2 aggregate interface view	interface <i>interface-type</i> <i>interface-number</i>	Required. Use either command.
	Enter port group view	port-group manual <i>port-group-name</i>	
3. Enable no agreement check	stp no-agreement-check	Required. Disabled by default.	


No agreement check configuration example

1. Network requirements

As shown in [Figure 26](#):

- 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 26 No agreement check configuration



2. Configuration procedure

Enable no agreement check on GigabitEthernet 1/0/1 of Device A.

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

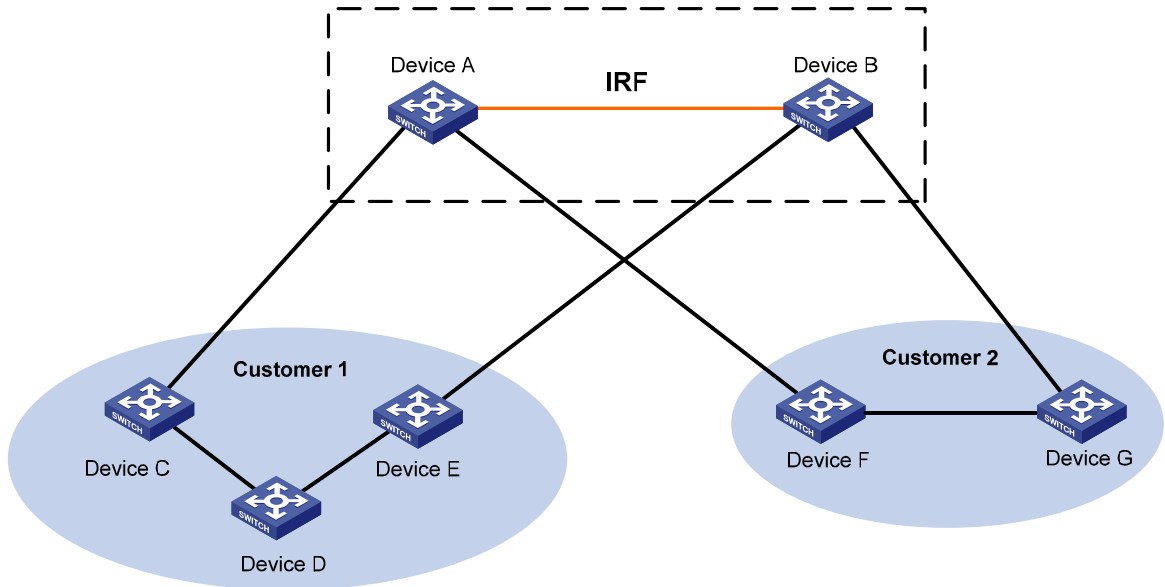
```
[DeviceA] interface gigabitethernet 1/0/1
```

```
[DeviceA-GigabitEthernet1/0/1] stp no-agreement-check
```

Configuring TC snooping

[Figure 27](#) shows a TC snooping application scenario. Device A and Device B are both IRF-enabled switches and form an IRF virtual device; they operate at the distribution layer and do not have STP enabled. The IRF virtual device formed by Device A and Device B connect to multiple access-layer customer networks, such as Customer 1 and Customer 2. Device C, Device D, and Device E in customer network Customer 1 are all enabled with STP. Customer 1 is dual-uplinked to the IRF virtual device for high availability. The IRF virtual device transparently transmits STP BPDUs from Customer 1 at Layer 2. Other customer networks (such as Customer 2) act the same as Customer 1.

Figure 27 TC snooping application scenario



In the network, the IRF virtual device transparently transmits the received STP BPDUs and does not participate in STP calculations. When a topology change occurs to the IRF virtual device or attached access-layer networks, the IRF virtual device may need a long time to learn the correct MAC address table entries and ARP entries, resulting in long network disruption. To avoid the network disruption, you can enable TC snooping on the IRF virtual device.

TC snooping enables a device to actively clear the MAC address table entries and ARP entries upon receiving TC-BPDUs and to re-learn the MAC address table entries and ARP entries, so that the device can normally forward the user traffic.

Configuration prerequisites

Disable STP globally.

Configuring TC snooping

Perform the TC snooping configuration on the IRF virtual device shown in Figure 27.

To configure TC snooping:

To do...	Use the command...	Description
1. Enter system view	system-view	—
2. Enable TC snooping	stp tc-snooping	Required Disabled by default

TC snooping and STP are mutually exclusive.

TC snooping does not take effect on the ports on which BPDU tunneling is enabled for STP. For more information about BPDU tunneling, see the chapter “BPDU tunneling configuration.”

For more information about IRF, see the *IRF Configuration Guide*.

For more information about ARP, see the *Layer 3—IP Services Configuration Guide*.

For more information about MAC address table entries, see the chapter “MAC address table configuration.”

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 will set these ports automatically as non-edge ports and start a new spanning tree calculation process. This will cause 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 will close these ports and notify the NMS that these ports have been closed by MSTP. The closed ports will be re-activated by the device 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.

To enable BPDU guard:

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

BPDU guard does not take effect on loopback test-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 may receive a configuration BPDU with a higher priority. The current legal root bridge will be superseded by another device, causing an undesired change of the network topology. As a result, the traffic that should go over high-speed links is switched to low-speed links, resulting in network congestion.

To prevent this situation from happening, 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. Once 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 will revert to its original state.

Make this configuration on a designated port.

To enable root guard:

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

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

Enabling loop guard

By keeping receiving 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 may 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.

To enable loop guard:

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

Do not enable loop guard on a port connecting 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 take effect on a port at the same time.

Enabling TC-BPDU guard

When receiving TC BPDUs (the BPDUs used to notify topology changes), a 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 certain period of time after receiving 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.

To enable TC-BPDU guard:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable the TC-BPDU guard function	stp tc-protection enable	Optional Enabled by default
3. 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 <i>number</i>	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 forge BPDUs to attack the network: By continuously sending forged BPDUs, they can make all 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.

To enable BPDU drop on an Ethernet interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. 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 } <i>regular-expression</i>]	Available in any view
Display BPDU statistics on ports	display stp bpdu-statistics [interface <i>interface-type interface-number</i> [instance <i>instance-id</i>]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display information about ports blocked by STP protection functions	display stp down-port [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the historical information of port role calculation for the specified MSTI or all MSTIs	display stp [instance <i>instance-id</i>] history [slot <i>slot-number</i>] [{ begin exclude include } <i>regular-expression</i>]	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 <i>instance-id</i>] tc [slot <i>slot-number</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the status and statistics of MSTP	display stp [instance <i>instance-id</i>] [interface <i>interface-list</i> slot <i>slot-number</i>] [brief] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the MST region configuration information that has taken effect	display stp region-configuration [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the root bridge information of all MSTIs	display stp root [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Clear the statistics of MSTP	reset stp [interface <i>interface-list</i>]	Available in user view

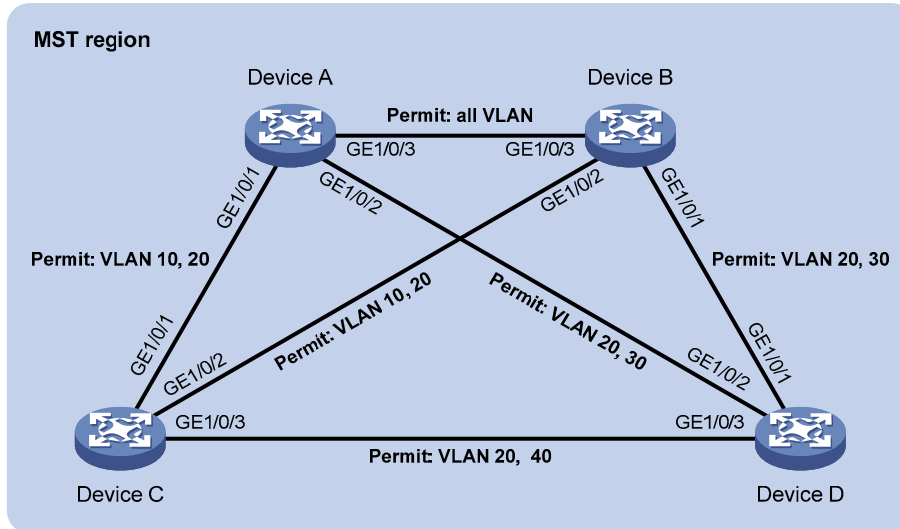
MSTP configuration example

Network requirements

As shown in [Figure 28](#):

- 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, so 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 28 Network diagram for MSTP configuration



Configuration procedure

1. VLAN and VLAN member port configuration

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. The detailed configuration procedure is omitted.

2. Configuration on 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, and 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. Configuration on 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. Configuration on 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. Configuration 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. Verifying the configurations

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

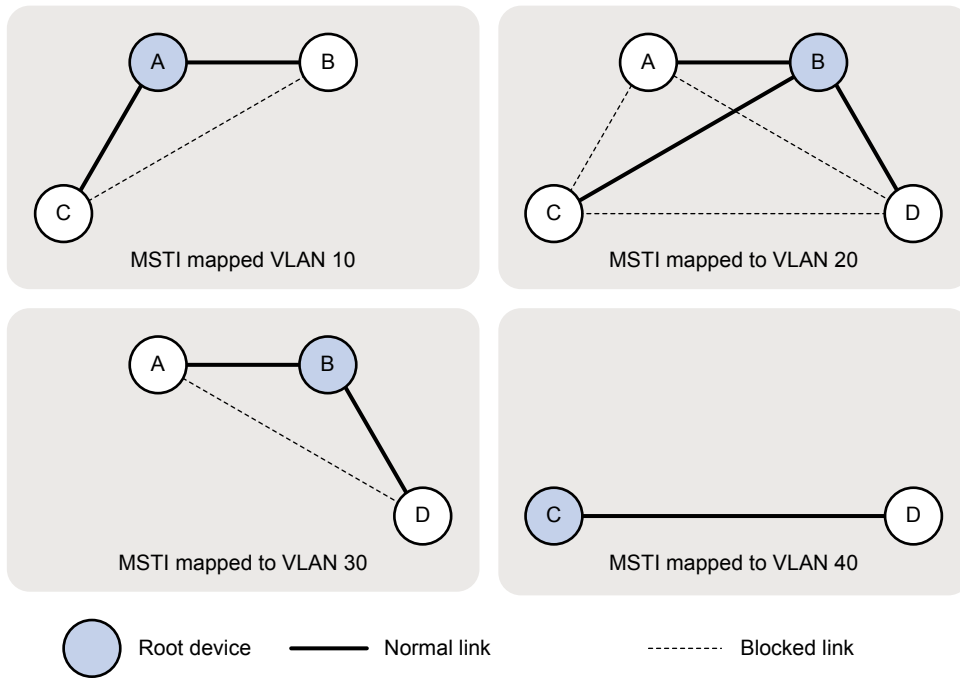
```
# Display brief spanning tree information on Device D.
```

```
[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 29](#).

Figure 29 MSTIs mapped to different VLANs




BPDU tunneling configuration

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.

Customers usually use dedicated lines in a service provider network to build their own Layer 2 networks. As a result, very often, a customer network is broken down into parts located at different sides of the service provider network. As shown in [Figure 30](#), User A has two devices: 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 30 BPDU tunneling application scenario



With BPDU tunneling, Layer 2 protocol packets from customer networks can be transparently transmitted in the service provider network:

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.

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

- CDP
- DLDAP
- EOAM
- GVRP
- HGMP
- LACP
- LLDAP

- PAGP
- PVST
- STP
- UDLD
- VTP

BPDU tunneling implementation

The BPDU tunneling implementations for different protocols are all similar. This section describes how BPDU tunneling is implemented by taking the STP as an example.

The term STP in this document 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, the devices at this 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.

To allow each network to calculate an independent spanning tree with STP, BPDU tunneling was introduced.

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. Each customer network can perform independent spanning tree calculation.

Figure 31 Network diagram for BPDU tunneling implementation



As shown in [Figure 31](#), 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, ensuring consistent spanning tree calculation throughout User A network, without affecting the spanning tree calculation of the service provider network.

Assume a BPDU is sent from User A network 1 to User A network 2. The BPDU is sent 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.

Make 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 connecting network devices in the service provider network as trunk ports allowing packets of any VLAN to pass through.

Enabling BPDU tunneling

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

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 enabling BPDU tunneling for DLDP, EOAM, GVRP, HGMP, LLDP, or STP on a port, disable the protocol on the port first. Because PVST is a special STP protocol, before enabling BPDU tunneling for PVST on a port, you must also disable STP and then enable BPDU tunneling for STP on the port first.

Before enabling BPDU tunneling for LACP on a member port of a link aggregation group, remove the port from the link aggregation group first.

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

To enable BPDU tunneling for a protocol in Ethernet interface view or port group view:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Ethernet interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Required. Use either command.
3. 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

To enable BPDU tunneling for a protocol in Layer 2 aggregate interface view:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	—
3. 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. Change it to 0x0100-0CCD-CDD0, 0x0100-0CCD-CDD1 or 0x0100-0CCD-CDD2 through the following configuration.

To configure destination multicast MAC address for BPDUs:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the destination multicast MAC address for BPDUs	bpdu-tunnel tunnel-dmac <i>mac-address</i>	Optional 0x010F-E200-0003 by default

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 32:

- 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 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 implement consistent spanning tree calculation across the service provider network, and that the destination multicast MAC address carried in BPDUs is 0x0100-0CCD-CDD0.

Figure 32 Network diagram for configuring BPDU tunneling for STP



Configuration procedure

1. Configuration on 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. Configuration on 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

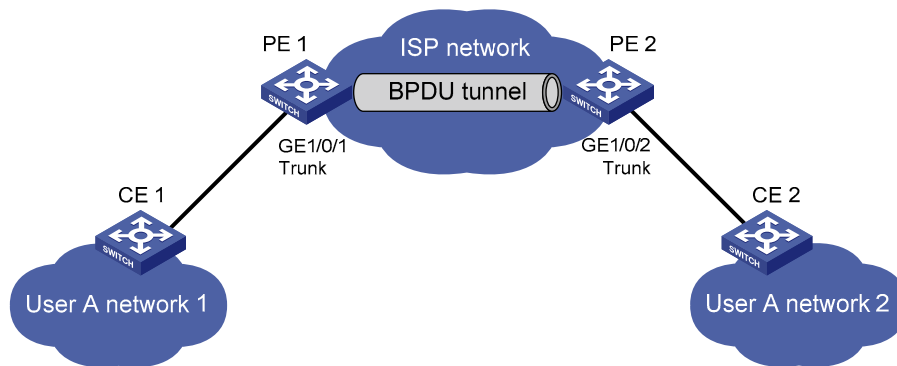
Network requirements

As shown in [Figure 33](#):

- 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 implement consistent PVST calculation across the service provider network, and that the destination multicast MAC address carried in BPDUs is 0x0100-0CCD-CDD0.

Figure 33 Network diagram for configuring BPDU tunneling for PVST



Configuration procedure

1. Configuration on 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. Configuration on 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

Ethernet is a network technology based on the CSMA/CD mechanism. As the medium is shared, collisions and excessive broadcasts are common on Ethernet networks. To address the issue, 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 all broadcast traffic is contained within it, as shown in [Figure 34](#).

Figure 34 A VLAN diagram



A VLAN is logically divided on an organizational basis rather than on a physical basis. For example, all workstations and servers used by a particular workgroup 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.
- Flexible virtual workgroup creation. As users from the same workgroup can be assigned to the same VLAN regardless of their physical locations, network construction and maintenance is 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 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 indicating the upper layer protocol type, as shown in [Figure 35](#).

Figure 35 The format of a traditional Ethernet frame



IEEE 802.1Q inserts a four-byte VLAN tag after the DA&SA field, as shown in [Figure 36](#).

Figure 36 The position and format of VLAN tag



A VLAN tag comprises the following fields: TPID, priority, 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 the frame belongs to. The VLAN ID range is 0 to 4095. As 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](#)."

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.

For a frame with multiple VLAN tags, the device handles it according to its outer-most VLAN tag and transmits its inner VLAN tags as payload.

Types of VLAN

Implement VLAN based on the following criteria:

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

This series Ethernet switches support 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.

Configure all four types of VLANs on a port at the same time. When determining to which VLAN a packet passing through the port should be assigned, the device 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

To configure basic VLAN settings:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create VLANs	vlan { <i>vlan-id1</i> [<i>to</i> <i>vlan-id2</i>] all }	Optional. Use this command to create VLANs in bulk.
3. Enter VLAN view	vlan <i>vlan-id</i>	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.
4. Configure a name for the current VLAN	name <i>text</i>	Optional. By default, the name of a VLAN is its VLAN ID, VLAN 0001 for example.
5. Configure the description of the current VLAN	description <i>text</i>	Optional. VLAN ID is used by default, VLAN 0001 for example.

As the default VLAN, VLAN 1 cannot be created or removed.

You cannot manually create or remove VLANs reserved for special purposes.

Dynamic VLANs cannot be removed with the **undo vlan** command.

A VLAN with a QoS policy applied cannot be removed.

A VLAN operating as a probe VLAN for remote port mirroring or an RRPP protected VLAN cannot be removed with the **undo vlan** command. To do that, remove the remote mirroring 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. To achieve this, VLAN interfaces are used.

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. 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.

To configure basic settings of a VLAN interface:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a VLAN interface and enter VLAN interface view	interface vlan-interface <i>vlan-interface-id</i>	Required. If the VLAN interface already exists, you enter its view directly.

To do...	Use the command...	Remarks
3. Assign an IP address to the VLAN interface	ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> } [sub]	Optional. No IP address is assigned to any VLAN interface by default.
4. Configure the description of the VLAN interface	description <i>text</i>	Optional. VLAN interface name is used by default, for example, Vlan-interface1 Interface .
5. Bring up the VLAN interface	undo shutdown	Optional. By default, a VLAN interface is in the up state. The VLAN interface is up so 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 however will be in the down state until you bring it up, regardless of how the state of the ports in the VLAN changes.

Before creating a VLAN interface for a VLAN, create the VLAN first.

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

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 recognize VLAN tagged-packets or when there is no need to separate different VLAN members.
- A trunk port can carry multiple VLANs to receive and send traffic for them. Except traffic of the default VLAN, traffic sent through a trunk port will be VLAN tagged. Usually, ports connecting 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.

Default VLAN

By default, VLAN 1 is the default VLAN for all ports. Configure the default VLAN for a port as required.

Use the following guidelines when configuring the default VLAN on a port:

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

Do not set the voice VLAN as the default VLAN 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 default VLAN ID for the local and remote ports.

Make sure that a port is assigned to its default VLAN. Otherwise, when the port receives frames tagged with the default VLAN ID 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 default VLAN tag.	<ul style="list-style-type: none"> Receive the frame if its VLAN ID is the same as the default VLAN ID. Drop the frame if its VLAN ID is different from the default VLAN ID. 	Remove the VLAN tag and send the frame.
Trunk	Check whether the default VLAN is permitted on the port: <ul style="list-style-type: none"> If yes, tag the frame with the default VLAN tag. If not, drop the frame. 	<ul style="list-style-type: none"> Receive the frame if its VLAN is carried on the port. Drop the frame if its VLAN is not carried on the port. 	<ul style="list-style-type: none"> Remove the tag and send the frame if the frame carries the default VLAN tag and the port belongs to the default VLAN. Send the frame without removing the tag if its VLAN is carried on the port but is different from the default one.
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 with the port hybrid vlan command. This is true of the default VLAN.

Assigning an access port to a VLAN

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.

To assign one or multiple access ports to a VLAN in VLAN view:

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

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
1. Enter system view		system-view	—
2. Enter interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view	Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	Required. Use either command.
	Enter Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	<ul style="list-style-type: none"> 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.
	Enter port group view	port-group manual <i>port-group-name</i>	<ul style="list-style-type: none"> In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.
3. 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.
4. Assign the current access ports to a VLAN		port access vlan <i>vlan-id</i>	Optional. By default, all access ports belong to VLAN 1.

Before assigning an access port to a VLAN, create the VLAN first.

In VLAN view, only Layer 2 Ethernet interfaces can be assigned 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. Assign it to a VLAN in interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view.

To assign a trunk port to one or multiple VLANs:

To do...		Use the command...	Remarks
1. Enter system view		system-view	—
2. Enter interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view	Enter Ethernet interface view	interface <i>interface-type interface-number</i>	Required. Use either command. <ul style="list-style-type: none"> 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 <i>interface-number</i>	
	Enter port group view	port-group manual <i>port-group-name</i>	
3. Configure the link type of the port or ports as trunk		port link-type trunk	Required.
4. Assign the trunk ports to the specified VLANs		port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port carries only VLAN 1.
5. Configure the default VLAN of the trunk ports		port trunk pvid vlan <i>vlan-id</i>	Optional. VLAN 1 is the default VLAN by default.

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 default VLAN for a trunk port, you must use the **port trunk permit vlan** command to configure the trunk port to allow packets from the default VLAN to pass through, so that the egress port can forward packets from the default VLAN.

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 with 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. Assign it to a VLAN in interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view.

To assign a hybrid port to one or multiple VLANs:

To do...	Use the command...	Remarks	
1. Enter system view	system-view	—	
2. Enter interface view (including Ethernet interface view, Layer 2 aggregate interface view) or port group view	Enter Ethernet interface view <hr/> Enter Layer 2 aggregate interface view <hr/> Enter port group view	interface <i>interface-type interface-number</i> <hr/> interface bridge-aggregation <i>interface-number</i> <hr/> port-group manual <i>port-group-name</i>	Required. Use either command. <ul style="list-style-type: none"> 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.
3. Configure the link type of the ports as hybrid	port link-type hybrid	Required.	
4. Assign the hybrid ports to the specified VLANs	port hybrid vlan <i>vlan-id-list</i> { tagged untagged }	Required. By default, a hybrid port allows only packets of VLAN 1 to pass through untagged.	
5. Configure the default VLAN of the hybrid port	port hybrid pvid vlan <i>vlan-id</i>	Optional. VLAN 1 is the default by default.	

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 with 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 assigning a hybrid port to a VLAN, create the VLAN first.

After configuring the default VLAN for a hybrid port, you must use the **port hybrid vlan** command to configure the hybrid port to allow packets from the default VLAN to pass through, so that the egress port can forward packets from the default 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.


Port-based VLAN configuration example

Network requirements

As shown in Figure 37:

- 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 37 Network diagram for port-based VLAN configuration



Configuration procedure

1. Configuration on 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, enabling 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. Check 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. The switch first performs a fuzzy match. In the fuzzy match, the switch searches the MAC address-to-VLAN entries whose masks are not all-Fs and performs a logical AND operation on the source MAC address and each mask. If the result of an AND operation matches the corresponding MAC address, the switch tags the frame with the corresponding VLAN ID. If the fuzzy match fails, the switch performs an exact match. In the

exact match, the switch searches the MAC address-to-VLAN entries whose masks are all-Fs. 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.

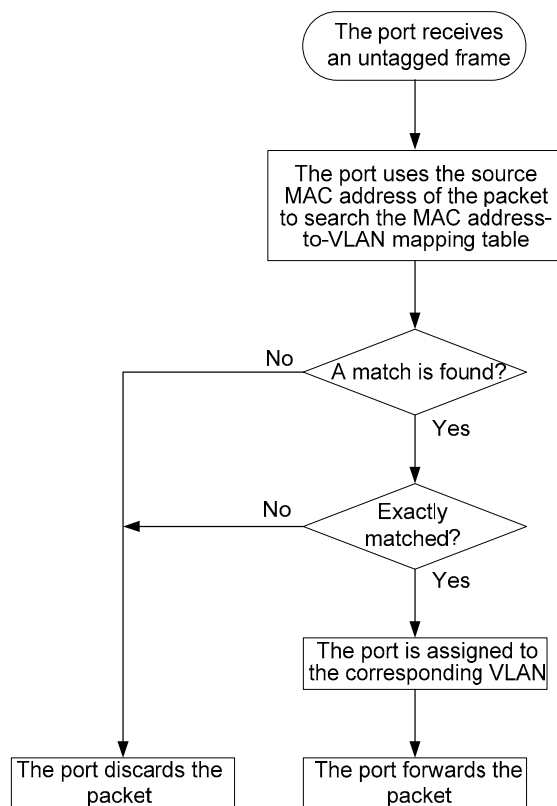
Approach 2: Dynamic MAC-based VLAN assignment

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 shown in Figure 38.

Figure 38 Flowchart for processing an untagged frame in dynamic MAC-based VLAN assignment



- 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.

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

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

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.

Configuring static MAC-based VLAN assignment

To configure static MAC-based VLAN assignment

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Associate MAC addresses with a VLAN	mac-vlan mac-address mac-address [mask mac-mask] vlan vlan-id [priority priority]	Required.
3. Enter Ethernet interface view or port group view	Enter Ethernet interface view interface interface-type interface-number <hr/> Enter port group view port-group manual port-group-name	Use either command. <ul style="list-style-type: none"> 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.
4. Configure the link type of the ports as hybrid	port link-type hybrid	Required.
5. Configure the hybrid ports 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.
6. Enable MAC-based VLAN	mac-vlan enable	Required. Disabled by default
7. 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

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 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

To configure dynamic MAC-based VLAN assignment:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Associate MAC addresses with a VLAN	mac-vlan mac-address mac-address vlan vlan-id [priority priority]	Required.
3. Enter Ethernet interface view or port group view	Enter Ethernet interface view interface interface-type interface-number Enter port group view port-group manual port-group-name	Use either command. <ul style="list-style-type: none"> 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.
4. Configure the link type of the ports as hybrid	port link-type hybrid	Required.
5. Enable MAC-based VLAN	mac-vlan enable	Required. Disabled by default.
6. Configure VLAN matching precedence	vlan precedence { mac-vlan ip-subnet-vlan }	Optional. By default, VLANs are preferably matched based on MAC addresses.
7. Enable dynamic MAC-based VLAN assignment	mac-vlan trigger enable	Required. Disabled by default.
8. Disable the default VLAN 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 default VLAN of the incoming port if they do not match any MAC address-to-VLAN mapping.

Configuring dynamic MAC-based VLAN

After enabling MAC-based VLAN on the switch, you must configure related authentication settings on the access authentication server. For more information about 802.1X authentication, see the *Security Configuration Guide*.

To configure dynamic MAC-based VLAN:


To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Ethernet interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Use either command. <ul style="list-style-type: none">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.
3. Configure the link type of the ports as hybrid	port link-type hybrid	Required.
4. Enable MAC-based VLAN	mac-vlan enable	Required. Disabled by default.

MAC-based VLAN configuration example

Network requirements

- As shown in [Figure 39](#), 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 meeting and may be used in any of the two meeting rooms.
- Laptop 1 and Laptop 2 are owned by different departments. The two departments use VLAN 100 and VLAN 200, respectively. Each laptop can 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 39 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 the MAC address of Laptop 2 with VLAN 200.

Configuration procedure

1. Configuration on 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 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. Configuration on Device B

Create VLANs 100 and 200. Assign GigabitEthernet 1/0/13 to VLAN 100, and 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. Configuration on 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

Protocol-based VLAN configuration applies to hybrid ports only.

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 AT. The encapsulation formats include Ethernet II, 802.3 raw, 802.2 LLC, and 802.2 SNAP.

A protocol type and an encapsulation format comprise a protocol template. Create multiple protocol templates for a protocol-based VLAN, and different protocol templates are assigned different *protocol-index* values. A protocol template can be uniquely identified by a protocol-based VLAN ID and a protocol index combined. 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 reaching 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 corresponding to the protocol template.
- If the packet matches no protocol templates, the packet will be tagged with the default VLAN ID 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 corresponding 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



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.

To configure a protocol-based VLAN:

To do...		Use the command...	Remarks
1. Enter system view		system-view	—
2. Enter VLAN view		vlan <i>vlan-id</i>	Required. If the specified VLAN does not exist, this command creates the VLAN first.
3. Create a protocol template for the VLAN		protocol-vlan [<i>protocol-index</i>] { 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 <i>etype-id</i> } }	Required.
4. Exit VLAN view		quit	Required.
5. Enter interface view or port group view	Enter Ethernet interface view	interface <i>interface-type interface-number</i>	Required. Use either command. <ul style="list-style-type: none"> • In Ethernet interface view, the subsequent configurations apply to the current port.
	Enter Layer 2 aggregate interface view	interface bridge-aggregation <i>interface-number</i>	<ul style="list-style-type: none"> • In port group view, the subsequent configurations apply to all ports in the port group.
	Enter port group view	port-group manual <i>port-group-name</i>	<ul style="list-style-type: none"> • In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.


To do...	Use the command...	Remarks
6. Configure the port link type as hybrid	port link-type hybrid	Required.
7. Configure current hybrid ports to permit the packets of the specified protocol-based VLANs to pass through	port hybrid vlan <i>vlan-id-list</i> { tagged untagged }	Required. By default, all hybrid ports permit packets of VLAN 1 to pass through only.
8. Associate the hybrid ports with the specified protocol-based VLAN	port hybrid protocol-vlan <i>vlan-id</i> { <i>protocol-index</i> [to <i>protocol-end</i>] all }	Required.

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 purpose. To avoid interference, isolate IPv4 traffic and IPv6 traffic at Layer 2.

Figure 40 Network diagram for protocol-based VLAN configuration



Configuration consideration

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

Configuration procedure

1. Configuration on 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 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/server in VLAN 200, and vice versa.
 2. Display protocol-based VLAN information on Device to check 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

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

This feature is only applicable on hybrid ports.

To configure an IP subnet-based VLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter VLAN view	vlan <i>vlan-id</i>	—
3. Associate an IP subnet with the current VLAN	ip-subnet-vlan [<i>ip-subnet-index</i>] ip <i>ip-address</i> [<i>mask</i>]	Required. The IP network segment or IP address to be associated with a VLAN cannot be a multicast network segment or a multicast address.
4. Return to system view	quit	—
5. Enter interface view or port group view	Enter Ethernet interface view interface <i>interface-type interface-number</i>	Required. Use either command. <ul style="list-style-type: none"> 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 <i>interface-number</i>	
	Enter port group view port-group manual <i>port-group-name</i>	
6. Configure port link type as hybrid	port link-type hybrid	Required.
7. Configure the hybrid ports to permit the specified IP subnet-based VLANs to pass through	port hybrid vlan <i>vlan-id-list</i> { tagged untagged }	Required.
8. Associate the hybrid ports with the specified IP subnet-based VLAN	port hybrid ip-subnet-vlan <i>vlan-id</i>	Required.

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

Super VLAN configuration

Super VLAN, also called “VLAN aggregation,” was introduced to save the IP address space.

A super VLAN is associated with multiple sub-VLANs. Create a VLAN interface for a super VLAN and assign an IP address for the VLAN interface. However, you cannot create a VLAN interface for a sub-VLAN. Assign a physical port to a sub-VLAN, but not to a super VLAN. All ports of a sub-VLAN use the VLAN interface IP address of the associated super VLAN. Packets cannot be forwarded between sub-VLANs at Layer 2.

To enable Layer 3 communication between sub-VLANs, create a super VLAN and the VLAN interface, and enable local proxy ARP or local proxy ND on the VLAN interface depending on the VLAN interface IP address type (IPv4 or IPv6) as follows:

- In an IPv4 network, enable local proxy ARP on the VLAN interface. The super VLAN can use local proxy ARP to forward and process ARP requests and replies.
- In an IPv6 network, enable local proxy ND on the VLAN interface. The super VLAN can use local proxy ND to forward and process the NS messages and NA messages.

Configuring super VLAN

To configure super VLAN, complete the following tasks:

1. Configure sub-VLANs.
2. Configure a super VLAN, and associate the super VLAN with the sub-VLANs configured earlier.
3. Configure a VLAN interface for the super VLAN. The VLAN interface enables communication among hosts and sub-VLANs.

Configuring sub-VLANs

To configure a sub-VLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a sub-VLAN and enter VLAN view	vlan <i>vlan-id</i>	Required. If the specified VLAN already exists, this command enters VLAN view only.

To configure more sub-VLANs, repeat these steps.

Configuring a super VLAN

To configure a super VLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter VLAN view	vlan <i>vlan-id</i>	Required. If the specified VLAN does not exist, this command creates the VLAN first, and then enters VLAN view.

To do...	Use the command...	Remarks
3. Configure the VLAN as a super VLAN	supervlan	Required.
4. Associate the super VLAN with the specified sub-VLANs	subvlan <i>vlan-list</i>	Required. VLANs specified by <i>vlan-list</i> must be the sub-VLANs configured earlier.

Configure a VLAN interface for the super VLAN

To configure a VLAN interface for the super VLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a VLAN interface, and enter VLAN interface view	interface <i>vlan-interface</i> <i>vlan-interface-id</i>	Required. The value of <i>vlan-interface-id</i> must be the ID of the super VLAN.
3. Configure the IP address of the VLAN interface	ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> } [sub] ipv6 address { <i>ipv6-address</i> { <i>prefix-length</i> link-local } <i>ipv6-address/prefix-length</i> [anycast eui-64] auto [link-local] }	Required. Use either command. By default, the IP address of a VLAN interface is not configured.
4. Enable local proxy ARP	Enable local proxy ARP local-proxy-arp enable	Required. Use either command.
	Enable local proxy ND local-proxy-nd enable	Disabled by default.

Configure the IP address of the VLAN interface with that of the corresponding super VLAN.

For more information about the **local-proxy-arp enable** command and the local proxy ARP function, see the *Layer 3—IP Services Configuration Guide*.

You cannot configure a super VLAN as the guest VLAN for a port, and vice versa. For more information about guest VLAN, see the *Security Configuration Guide*.

Configure Layer 2 multicast for a super VLAN. However, the configuration cannot take effect.

Configure DHCP, DHCPv6, Layer 3 multicast, dynamic routing, and NAT for the VLAN interface of a super VLAN. However, only DHCP and DHCPv6 take effect.

Configuring VRRP for the VLAN interface of a super VLAN affects network performance. HP does not recommend you to configure this function

Displaying and maintaining super VLAN

To do...	Use the command...	Remarks
Display the mapping between a super VLAN and its sub-VLANs	display supervlan [<i>supervlan-id</i>] [[{ begin exclude include } <i>regular-expression</i>]]	Available in any view

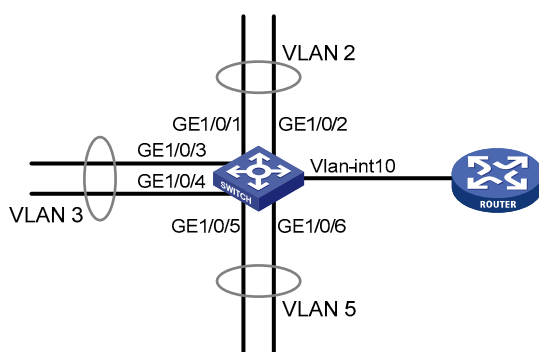
Super VLAN configuration example

Network requirements

As shown in Figure 41,

- Create super VLAN 10, and configure the IPv4 address and IPv6 address of its VLAN interface as 10.0.0.1/24 and 2001::1/64.
- Create the sub-VLANs VLAN 2, VLAN 3, and VLAN 5.
- Assign GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to VLAN 2, GigabitEthernet 1/0/3 and GigabitEthernet 1/0/4 to VLAN 3, and GigabitEthernet 1/0/5 and GigabitEthernet 1/0/6 to VLAN 5.
- The sub-VLANs are isolated at Layer 2 but connected at Layer 3.

Figure 41 Network diagram for super-VLAN configuration



Configuration procedure

```
# Create VLAN 10.
```

```
<Sysname> system-view
[Sysname] vlan 10
[Sysname-vlan10] quit
```

```
# Create VLAN-interface 10, and configure the IPv4 address and IPv6 address of VLAN-interface 10 as 10.0.0.1/24 and 2001::1/64.
```

```
[Sysname-vlan10] interface vlan-interface 10
[Sysname-Vlan-interface10] ip address 10.0.0.1 255.255.255.0
[Sysname-Vlan-interface10] ipv6 address 2001::1/64
```

```
# Enable local proxy ARP on VLAN interface 10, so that IPv4 packets can be exchanged between sub-VLANs at Layer 3.
```

```
[Sysname-Vlan-interface10] local-proxy-arp enable
```

```
# Enable local proxy ND on VLAN interface 10, so that IPv6 packets can be exchanged between sub-VLANs at Layer 3.
```

```
[Sysname-Vlan-interface10] local-proxy-nd enable
[Sysname-Vlan-interface10] quit
```

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

```
[Sysname] vlan 2
[Sysname-vlan2] port gigabitethernet 1/0/1 gigabitethernet 1/0/2
[Sysname-vlan2] quit
```

```

# Create VLAN 3, and assign GigabitEthernet 1/0/3 and GigabitEthernet 1/0/4 to it.
[Sysname] vlan 3
[Sysname-vlan3] port gigabitethernet 1/0/3 gigabitethernet 1/0/4
[Sysname-vlan3] quit

# Create VLAN 5, and assign GigabitEthernet 1/0/5 and GigabitEthernet 1/0/6 to it.
[Sysname] vlan 5
[Sysname-vlan5] port gigabitethernet 1/0/5 gigabitethernet 1/0/6
[Sysname-vlan5] quit

# Configure VLAN 10 as the super VLAN, and configure VLAN 2, VLAN 3, and VLAN 5 as its sub-
VLANs.
[Sysname] vlan 10
[Sysname-vlan10] supervlan
[Sysname-vlan10] subvlan 2 3 5
[Sysname-vlan10] quit
[Sysname] quit

```

Verification

Display information about VLAN 10, the super VLAN, to verify the configuration.

```

<Sysname> display supervlan
SuperVLAN ID : 10
SubVLAN ID : 2-3 5

VLAN ID: 10
VLAN Type: static
It is a Super VLAN.
Route Interface: configured
Ip Address: 10.0.0.1
Subnet Mask: 255.255.255.0
Description: VLAN 0010
Name: VLAN 0010
Tagged Ports: none
Untagged Ports: none

VLAN ID: 2
VLAN Type: static
It is a Sub VLAN.
Route Interface: configured
Ip Address: 10.0.0.1
Subnet Mask: 255.255.255.0
Description: VLAN 0002
Name: VLAN 0002
Tagged Ports: none
Untagged Ports:
    GigabitEthernet1/0/1    GigabitEthernet1/0/2

VLAN ID: 3
VLAN Type: static

```

It is a Sub VLAN.
Route Interface: configured
Ip Address: 10.0.0.1
Subnet Mask: 255.255.255.0
Description: VLAN 0003
Name: VLAN 0003
Tagged Ports: none
Untagged Ports:
 GigabitEthernet1/0/3 GigabitEthernet1/0/4

VLAN ID: 5
VLAN Type: static
It is a Sub VLAN.
Route Interface: configured
Ip Address: 10.0.0.1
Subnet Mask: 255.255.255.0
Description: VLAN 0005
Name: VLAN 0005
Tagged Ports: none
Untagged Ports:
 GigabitEthernet1/0/5 GigabitEthernet1/0/6

Isolate-user-VLAN configuration


An isolate-user-VLAN uses a two-tier VLAN structure. In this approach, two 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 is aware of only the isolate-user-VLAN but not the secondary VLANs, network configuration is simplified and VLAN resources are saved.
- Isolate the Layer 2 traffic of different users by assigning the ports connected to them to different secondary VLANs.

As shown in [Figure 42](#), the isolate-user-VLAN function is enabled on Switch B. VLAN 10 is the isolate-user-VLAN, and VLANs 2, 5, and 8 are secondary VLANs associated with VLAN 10 and are invisible to Switch A.

Figure 42 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 ports to the isolate-user-VLAN and configure these ports as upstream ports.
2. Configure the secondary VLANs.
Assign ports to each secondary VLAN and configure these ports as downstream ports.
3. Associate the isolate-user-VLAN with the specified secondary VLANs.
4. To enable users in the isolate-user-VLAN to communicate with other networks at Layer 3, configure VLAN interfaces for the isolate-user-VLAN and the secondary VLANs, and configure the gateway IP address for the isolate-user-VLAN interface (you do not need to configure IP addresses for the secondary VLAN interfaces).
5. 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 (for example, Switch A in Figure 42).

Configuring an isolate-user-VLAN

To configure an isolate-user-VLAN:

To do...	Use the command	Remarks
1. Enter system view	system-view	—
2. Create a VLAN and enter VLAN view	vlan <i>vlan-id</i>	—
3. Configure the VLAN as an isolate-user-VLAN	isolate-user-vlan enable	Required.
4. Return to system view	quit	—
5. Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	—
6. Configure the link type of the port	port link-type { access hybrid trunk }	—
7. Assign the port to the isolate-user-VLAN	Access port port access vlan <i>vlan-id</i>	Required. Use either approach.
	Hybrid port port hybrid vlan <i>vlan-id-list</i> { tagged untagged }	
	Trunk port port trunk permit vlan { <i>vlan-id-list</i> all }	
8. Configure the isolate-user-VLAN type of the port as upstream	port isolate-user-vlan <i>vlan-id</i> promiscuous	Required. By default, no isolate-user-VLAN type is specified for the port.
9. Return to system view	quit	—

To do...	Use the command	Remarks
10. Create the isolate-user-VLAN interface and enter the isolate-user-VLAN interface view	interface vlan-interface <i>vlan-interface-id</i>	<ul style="list-style-type: none"> This configuration is required when users in the isolate-user-VLAN must 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. <p>The <i>vlan-interface-id</i> argument must take the isolate-user-VLAN ID.</p>
11. Configure an IP address for the isolate-user-VLAN interface	ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> } [sub]	<ul style="list-style-type: none"> This configuration is required when users in the isolate-user-VLAN must 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. <p>By default, the isolate-user-VLAN ID is not configured with any IP address.</p>

Configuring secondary VLANs

To configure secondary VLANs:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create secondary VLANs	vlan { <i>vlan-id1</i> [to <i>vlan-id2</i>] all }	Required.
3. 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.
4. Return to system view	quit	—
5. Enter Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	—
6. Configure the link type of the port	port link-type { access hybrid trunk }	—

To do...	Use the command...	Remarks
	Access port port access vlan <i>vlan-id</i>	
7. Assign ports to the secondary VLAN	Hybrid port port hybrid vlan <i>vlan-id-list</i> { tagged untagged }	Required. Use either approach.
	Trunk port port trunk permit vlan { <i>vlan-id-list</i> all }	
8. Configure the isolate-user-VLAN type of the port as downstream	port isolate-user-vlan host	Required. By default, no isolate-user-VLAN type is specified for the port.
9. Return to system view	quit	—
10. Create a secondary VLAN interface and enter secondary VLAN interface view	interface vlan-interface <i>vlan-interface-id</i>	<ul style="list-style-type: none"> This configuration is required when users in the isolate-user-VLAN must 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. <p>The <i>vlan-interface-id</i> argument must take the secondary VLAN ID.</p>

Associating secondary VLANs with an isolate-user-VLAN

To associate secondary VLANs with an isolate-user-VLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Associate the specified secondary VLANs with the specified isolate-user-VLAN	isolate-user-vlan <i>isolate-user-vlan-id</i> secondary <i>secondary-vlan-id</i> [to <i>secondary-vlan-id</i>]	Required

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 [<i>isolate-user-vlan-id</i>] [[{ begin exclude include } <i>regular-expression</i>]]	Available in any view


Isolate-user-VLAN configuration example

Network requirements

As shown in Figure 43,

- 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 only has VLAN 5 and Device C has only VLAN 6.

Figure 43 Network diagram



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 in VLAN 5.

```
[DeviceB] interface gigabitethernet 1/0/5
```

```
[DeviceB-GigabitEthernet1/0/5] port isolate-user-vlan 5 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 in VLAN 6.

```
[DeviceC] interface gigabitethernet 1/0/5
```

```
[DeviceC-GigabitEthernet1/0/5] port isolate-user-vlan 6 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
```

Verifying the configurations

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

As voice communication technologies grow more mature, voice devices are more and more widely deployed, especially on broadband networks, where voice traffic and data traffic often co-exist. Usually, compared to data traffic, voice traffic is given a higher transmission priority for the purpose of reducing transmission delay and packet loss.

A voice VLAN is configured especially for voice traffic. After assigning the ports connecting to voice devices to a voice VLAN, the system automatically configures QoS parameters for voice traffic, improving the transmission priority of voice traffic and ensuring voice quality.

Common voice devices include IP phones and IADs. Only IP phones are used in the voice VLAN configuration examples in this chapter.

OUI addresses

A device determines whether a received packet is a voice packet by checking its source MAC address. A packet whose source MAC address complies with the voice device's OUI address is regarded as voice traffic.

Configure the OUI addresses of a device in advance or use the default OUI addresses. [Table 16](#) lists the default OUI address for each vendor's devices.

Table 16 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

In general, as the first 24 bits of a MAC address (in binary format), an OUI address is a globally unique identifier assigned to a vendor by IEEE. However, OUI addresses mentioned in this document are different from those in common sense. OUI addresses in this document are used by the system to determine whether a received packet is a voice packet. They are the results of the AND operation of the two arguments *mac-address* and *oui-mask* in the **voice vlan mac-address** command.

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 a match is found, the

system automatically assigns the receiving port to the voice VLAN, issues ACL rules and configures the packet precedence.

Configure 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. Assigning/removing ports to/from a voice VLAN are automatically performed by the system.

The 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 44](#). When the voice VLAN works normally, in case of a system reboot, 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, port assignment to the voice VLAN is not triggered by voice traffic streams.


Figure 44 PCs and IP phones connected in series access the network



- In manual mode, you need to 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 a match is found, the system issues ACL rules and configures the packet precedence. In this mode, assigning/removing ports to/from a voice VLAN are performed manually.

The manual mode is suitable for scenarios where only IP phones access the network through the device and ports on the device only transmit voice traffic, as shown in [Figure 45](#). 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 45 Only IP phones access the network



Both modes forward tagged packets according to their tags.

[Table 17](#) and [Table 18](#) list the required configurations on ports of different link types in order for these ports 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 17 Required configurations on ports of different links 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 default VLAN of the port, which cannot be the voice VLAN, and assign the port to its default VLAN.
	Manual		Make all configurations required for the automatic mode. In addition, assign the port to the voice VLAN.
Hybrid	Automatic	Yes	Configure the default VLAN of the port, which cannot be the voice VLAN, and configure the port to permit packets of its default VLAN to pass through untagged.
	Manual		Make all configurations required for the automatic mode. In addition, 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 18 Required configurations on ports of different links 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 default VLAN of the port as the voice VLAN.
Trunk	Automatic	No	—
	Manual	Yes	Configure the default VLAN of the port as the voice VLAN and assign the port to the voice VLAN.
Hybrid	Automatic	No	—
	Manual	Yes	Configure the default VLAN of the port as the voice VLAN and configure the port to permit packets of the voice VLAN to pass through untagged.

If an IP phone sends tagged voice traffic and its accessing port is configured with 802.1X authentication and guest VLAN, you should assign different VLAN IDs for the voice VLAN, the default VLAN 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 default VLAN of the IP phone's accessing port as the voice VLAN. As a result, 802.1X authentication cannot be implemented.

The default VLAN of a port is VLAN 1. Change the default VLAN 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 default VLAN 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:** In this mode, voice VLAN-enabled ports receive packets carrying the voice VLAN tag and forward packets in the voice VLAN without checking their source MAC addresses against the OUI addresses configured for the device. If the default VLAN 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, the voice VLANs are vulnerable to traffic attacks. Vicious users may forge a large amount of voice packets and send them to the device to consume the voice VLAN bandwidth, affecting normal voice communication.
- **Security mode:** In this 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 have to, ensure that the voice VLAN security mode is disabled.

Table 19 How a voice VLAN-enabled port processes packets in security/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 carrying the voice VLAN tag	
	Packets carrying 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 check the source MAC addresses of inbound packets. In this way, both voice traffic and non-voice traffic can be transmitted in the voice VLAN.
	Packets carrying the voice VLAN tag	
	Packets carrying 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.”](#)

A port can belong to only one voice VLAN at a time.

Voice VLAN cannot be enabled on a port with LACP 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 CoS and DSCP values, for voice traffic. Voice traffic carries its own QoS priority settings. Configure the device either to modify or not to modify the QoS priority settings carried by incoming voice traffic.

To configure QoS priority settings for voice traffic:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view	interface <i>interface-type</i> <i>interface-number</i>	—
3. 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 marked for voice VLAN traffic into 6 and 46, respectively.
4. Configure the interface to modify the CoS and DSCP values marked for incoming traffic of the voice VLAN into specified values	voice vlan qos <i>cos-value</i> <i>dscp-value</i>	The voice vlan qos command and the voice vlan qos trust command can overwrite each other, whichever is configured last.

Configure the QoS priority settings for voice traffic on an interface before enabling voice VLAN on the interface. If the configuration order is reversed, your priority trust setting will fail.

After configuring an interface enabled with voice VLAN to trust the QoS priority settings in incoming voice traffic, you still need to 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 a port to operate in automatic voice VLAN assignment mode

To set a port to operate in automatic voice VLAN assignment mode:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Set the voice VLAN aging time	voice vlan aging <i>minutes</i>	Optional. 1440 minutes by default. The voice VLAN aging time configuration is only applicable on ports in automatic voice VLAN assignment mode.
3. Enable the voice VLAN security mode	voice vlan security enable	Optional. Enabled by default.
4. Add a recognizable OUI address	voice vlan mac-address <i>oui mask</i> <i>oui-mask</i> [description <i>text</i>]	Optional. By default, each voice VLAN has default OUI addresses configured. For the default OUI addresses of different vendors, see Table 16 .
5. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
6. 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.
7. Enable voice VLAN on the port	voice vlan <i>vlan-id</i> enable	Required. Disabled by default.

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

To set a port to operate in manual voice VLAN assignment mode:

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

To do...	Use the command...	Remarks
2. Enable the voice VLAN security mode	voice vlan security enable	Optional. Enabled by default.
3. Add a recognizable OUI address	voice vlan mac-address <i>oui mask</i> <i>oui-mask</i> [description <i>text</i>]	Optional. By default, each voice VLAN has default OUI addresses configured. For the default OUI addresses of different vendors, see Table 16 .
4. Enter interface view	interface <i>interface-type interface-number</i>	—
5. Configure the port to operate in manual voice VLAN assignment mode	undo voice vlan mode auto	Required. Disabled by default.
6. 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 default VLAN of the port automatically.
7. Configure the voice VLAN as the default VLAN 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.
8. Enable voice VLAN on the port	voice vlan <i>vlan-id</i> enable	Required.

Configure different voice VLANs on different ports at the same time. However, one port can be configured with only one voice VLAN, and this voice VLAN must be a static VLAN that already exists on the device.

Voice VLAN cannot be enabled on a port with LACP 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 need to 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 } <i>regular-expression</i>]	Available in any view
Display the OUI addresses supported by system	display voice vlan oui [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Voice VLAN configuration examples


Automatic voice VLAN mode configuration example

Network requirements

As shown in Figure 46,

- 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 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 46 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 may 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: 128
```

```
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

- Create VLAN 2 and configure it as a voice VLAN permitting 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 47 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 default VLAN 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: 128
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 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 GVRP is a GARP application that registers and deregisters VLAN attributes. GVRP is based on the operating mechanism of GARP to maintain and propagate dynamic VLAN registration information for the GVRP devices on the network.

GARP

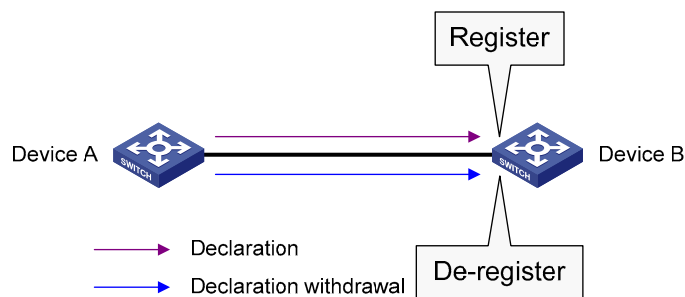
GARP provides a mechanism that allows participants in a GARP application to distribute, propagate, and register with other participants in a LAN the attributes specific to the GARP application, such as the VLAN or multicast address attributes.

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. As shown in Figure 48, 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 it receives.

Figure 48 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 de-registration 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 wants to register with other participants its attributes (including manually configured attributes), and when it receives Join messages from other participants. There are two types of Join messages, 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 receiving Leave messages from other GARP participants, and when attributes are manually deregistered on it. There are two types of Leave messages: 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. Upon the expiration of the LeaveAll timer, 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 four timers to control the sending of GARP messages:

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 21](#).

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 send 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 may 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.

All attributes share the same Join timer on a GARP participant. Set the Join timer long enough so that all attributes can be sent out 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 has not received a Join message for the attribute value before the timer expires, it deregisters the attribute value.

4. LeaveAll timer

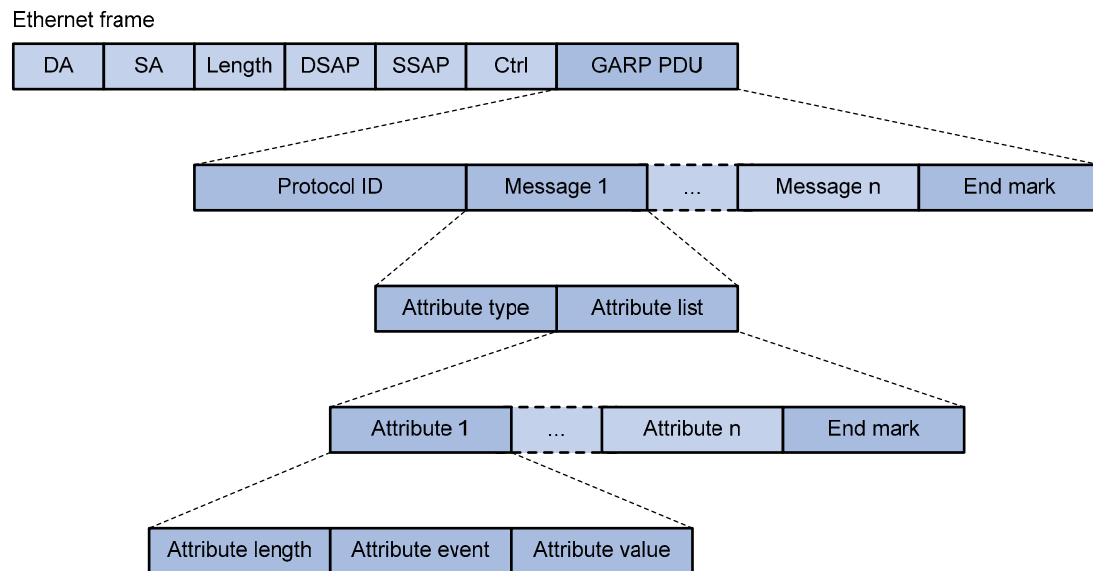
When a GARP application is enabled, a LeaveAll timer starts. The GARP participant sends a LeaveAll message when the timer expires. Then, the LeaveAll timer restarts to begin a new cycle. The LeaveAll timer and all other GARP timers also restart when the GARP participant receives a LeaveAll message.

Because a LeaveAll message deregisters all attributes in the entire network, do not set the LeaveAll timer too short. The LeaveAll timer should be at least longer than the Leave timer.

On a GARP-enabled network, a switch may send LeaveAll messages at the interval set by its LeaveAll timer or the LeaveAll timer on 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 49 GARP message format



As shown in [Figure 49](#), GARP messages use the IEEE 802.3 Ethernet frame format. [Table 20](#) describes the GARP message fields.

Table 20 Description on 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 containing an attribute type and an attribute list	—
End mark	Indicates the end of a GARP PDU	0x00
Attribute type	Defined by the GARP application	0x01 for GVRP, indicating the VLAN ID attribute
Attribute list	Contains one or multiple attributes	—

Field	Description	Value
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 described by the attribute	<ul style="list-style-type: none"> • 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 dynamically update the VLAN registration information from other devices to its local database about active VLAN members and through which port they can be reached. Thus, it 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

VLANs manually created are called static VLANs, and VLANs created by GVRP are called dynamic VLANs. GVRP provides the following registration modes on a port, including Normal, Fixed, and Forbidden. In different registration modes, a port handles static and dynamic VLANs differently.

- 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

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 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, GVRP can be configured only on trunk ports, and you must assign the involved trunk ports to all dynamic VLANs.

To configure GVRP functions on a trunk port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable GVRP globally	gvrp	Required. Globally disabled by default.
3. Enter Ethernet interface view, Layer 2 aggregate interface view, or port-group view	Enter Ethernet interface view or Layer 2 aggregate interface view interface <i>interface-type interface-number</i> Enter port-group view port-group manual <i>port-group-name</i>	Required. Perform either of the commands.
4. Configure the link type of the ports as trunk	port link-type trunk	Required. Access by default.
5. Assign the trunk ports to all VLANs	port trunk permit vlan all	Required. By default, a trunk port is assigned to VLAN 1 only.
6. Enable GVRP on the ports	gvrp	Required. Disabled by default.

To do...	Use the command...	Remarks
7. Configure the GVRP registration mode on the ports	gvrp registration { fixed forbidden normal }	Optional. normal . by default.

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/send GVRP packets.

Do not enable both GVRP and remote port mirroring. Otherwise, GVRP may register the remote probe VLAN to unexpected ports, resulting in undesired duplicates to be received by the monitor port. 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

Among the four GARP timers, the LeaveAll timer is configured in system view and takes effect on all ports, but the other three are configured on a port basis.

To configure GARP timers:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the GARP LeaveAll timer	garp timer leaveall <i>timer-value</i>	Optional. The default is 1000 centiseconds.
3. Enter Ethernet interface view, Layer 2 aggregate interface view, or port-group view	Enter Ethernet or Layer 2 aggregate interface view interface <i>interface-type interface-number</i> Enter port-group view port-group manual <i>port-group-name</i>	Required. Perform either of the commands. Depending on the view you accessed, the subsequent configuration takes effect on a port or all ports in a port-group.
4. Configure the Hold timer	garp timer hold <i>timer-value</i>	Optional. 10 centiseconds by default.
5. Configure the Join timer	garp timer join <i>timer-value</i>	Optional. 20 centiseconds by default.
6. Configure the Leave timer	garp timer leave <i>timer-value</i>	Optional. 60 centiseconds by default.

As shown in [Table 21](#), the value ranges for GARP timers are dependent on one another:

- If you want to set a value beyond the value range for a timer, you may 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, and then the Join, Leave, and LeaveAll timers.

Table 21 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	32765 centiseconds

Displaying and maintaining GVRP

To do...	Use the command...	Remarks
Display statistics about GARP on ports	display garp statistics [interface <i>interface-list</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display GARP timers on ports	display garp timer [interface <i>interface-list</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the local VLAN information maintained by GVRP on ports	display gvrp local-vlan interface <i>interface-type interface-number</i> [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the current GVRP state in the specified VLANs on ports	display gvrp state interface <i>interface-type interface-number vlan vlan-id</i> [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display GVRP statistics on ports	display gvrp statistics [interface <i>interface-list</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the global GVRP state	display gvrp status [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the information about dynamic VLAN operations on ports	display gvrp vlan-operation interface <i>interface-type interface-number</i> [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Clear the GARP statistics on ports	reset garp statistics [interface <i>interface-list</i>]	Available in user view

GVRP configuration examples


GVRP normal registration mode configuration example

Network requirements

As shown in [Figure 50](#),

- 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 50 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 maintained by GVRP on ports. For example:

```
# Display the local VLAN information maintained by GVRP 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 maintained by GVRP 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 51](#),

- 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 51 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 maintained by GVRP on ports. For example:

Display the local VLAN information maintained by GVRP 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 maintained by GVRP 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 52](#),

- 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 52 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 maintained by GVRP on ports. For example:

Display the local VLAN information maintained by GVRP 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 maintained by GVRP 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

Throughout this document, CVLANs, also called “*inner VLANs*,” refer to the VLANs that a customer uses on the private network. SVLANs, also called “*outer VLANs*,” refer to the VLANs that a service provider uses to carry VLAN tagged traffic for customers.

QinQ stands for 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 MANs.

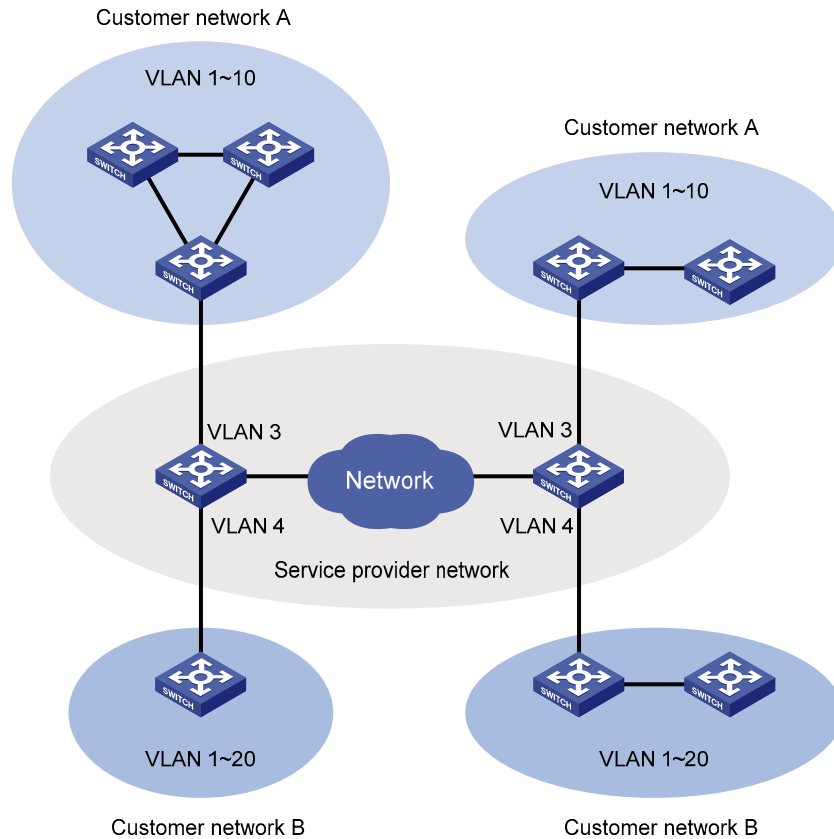
By tagging tagged frames, QinQ expands the available VLAN space from 4094 to 4094×4094 . QinQ delivers the following benefits:

- 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.
- Allows 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 learn 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 53 Typical QinQ application scenario

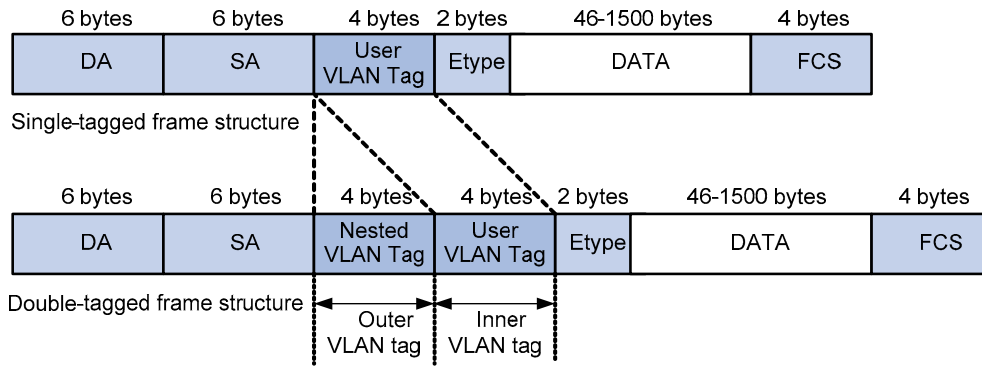


As shown in [Figure 53](#), customer network A has CVLANs 1 through 10, and customer network B has CVLANs 1 through 20. The service provider assigns SVLAN 3 for customer network A, and SVLAN 4 for customer network B. When a tagged Ethernet frame from customer network A arrives at the edge of the service 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 54 Single-tagged Ethernet frame header vs. double-tagged Ethernet frame header



The default MTU of an interface is 1500 bytes. The size of an outer VLAN tag is 4 bytes. H3C recommends you to increase the MTU of each interface on the service provider network to at least 1504 bytes. For more information about interface MTU configuration, see the chapter “Ethernet port configuration.”

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 default VLAN tag, 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 default VLAN tag.

2. Selective QinQ

Selective QinQ is more flexible than basic QinQ. In addition to all 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.
- Modifying the inner VLAN ID, in addition to tagging the frame with an outer VLAN tag.

Besides being able to separate the service provider network from the customer networks, selective QinQ provides abundant service features and allows more flexible networking.

Modifying the TPID in a VLAN tag

A VLAN tag uses the TPID field to identify the protocol type of the tag. The value of this field, as defined in IEEE 802.1Q, is 0x8100.

Figure 55 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.

In addition, the systems of different vendors may 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 to allow interoperability with the switches of that vendor.

The TPID in an Ethernet frame has the same position with 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 values in [Table 22](#).

Table 22 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	0xFFFF/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 inner VLAN ID substitution	Optional
Configuring the TPID value in VLAN tags		Optional

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 member ports in the aggregation group. Those made in port group view take effect on all member ports in the current port group.

On a port with basic QinQ enabled, you must configure the port to allow packets from its PVID to pass through. On a port with selective QinQ enabled, you must configure the port to allow packets from the outer VLANs of QinQ packets to pass through.

Configuring basic QinQ

Enabling basic QinQ

To enable basic QinQ:

To do...		Use the command...	Remarks
1. Enter system view		system-view	—
2. Enter interface view or port group view	Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view	interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view	port-group manual <i>port-group-name</i>	
3. Enable 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 default VLAN tag. However, by configuring the VLAN transparent transmission function on a port, you can specify the port not to add its default VLAN tag 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.

To configure VLAN transparent transmission:

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

To do...	Use the command...	Remarks
2. Enter interface view or port group view	Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view <hr/> Enter port group view interface <i>interface-type interface-number</i> <hr/> port-group manual <i>port-group-name</i>	Required. Use either command.
3. Configure the link type of the ports	port link-type { hybrid trunk }	—
4. Configure the ports to allow packets from its PVID and the transparent VLANs to pass through	<ul style="list-style-type: none"> When the ports are hybrid ports: port hybrid vlan <i>vlan-id-list</i> { tagged untagged } When the ports are trunk ports: port trunk permit vlan { <i>vlan-id-list</i> all } 	Required. Use either command.
5. Enable basic QinQ on the ports	qinq enable	Required. By default, basic QinQ is disabled on ports.
6. Configure VLAN transparent transmission on the ports	qinq transparent-vlan <i>vlan-list</i>	Required. By default, VLAN transparent transmission is not configured.

When configuring transparent transmission for a VLAN, you must configure all switches on the transmission path to permit packets of this VLAN to pass through.

For VLANs whose packets are to be transparently transmitted through a port, do not configure VLAN mapping for them on the port. For more information about VLAN mapping, see the chapter “VLAN mapping configuration.”

Configuring selective QinQ

Configuring an outer VLAN tagging policy

Basic QinQ can only tag received frames with the default VLAN tag of the receiving port. Selective QinQ allows adding different outer VLAN tags based on different inner VLAN tags.

The selective QinQ feature of the A5800&A5820X Switch Series is achieved through QoS policies. To enable the switch to tag tagged packets based on inner VLAN tags, follow these steps:

- Configure a class to match packets with certain tags;
- Configure a traffic behavior to tag packets with an outer VLAN tag;
- Create a QoS policy and associate the class with the behavior in the policy;
- Apply the QoS policy to the port that connects to the user.

To configure selective QinQ:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>classifier-name</i> [operator { and or }]	Required. By default, the operator of a class is AND.
3. Specify the inner VLAN IDs of matching frames	if-match customer-vlan-id <i>vlan-id-list</i>	Required.
4. Return to system view	quit	—
5. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
6. Specify an outer VLAN ID	nest top-most vlan-id <i>vlan-id</i>	Required.
7. Configure a traffic accounting action	accounting [byte packet]	Optional Configure a traffic accounting action to count QinQ packets that carry specific VLAN tags, keeping track of specific traffic changes, based on which you can design more effective transmission optimization policies. byte indicates that the traffic is counted in bytes, and packet indicates that the traffic is counted in packets.
8. Return to system view	quit	—
9. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
10. Associate the traffic class with the traffic behavior defined earlier	classifier <i>classifier-name</i> behavior <i>behavior-name</i>	Required.
11. Return to system view	quit	—
12. Enter the Ethernet port view of the customer network-side port	Enter Layer 2 Ethernet interface view <hr/> Enter port group view	Required. Use either command. <ul style="list-style-type: none"> 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.
13. Enable basic QinQ	qinq enable	Required.
14. Apply the QoS policy to the incoming traffic	qos apply policy <i>policy-name</i> inbound	Required.

△ CAUTION:

- Selective QinQ enjoys higher priority than basic QinQ. A received frame is tagged with an outer VLAN ID based on basic QinQ only after it fails to match the match criteria defined in the traffic class.
- Selective QinQ is achieved through QoS policies. For more information about QoS policies, see the *ACL and QoS Configuration Guide*.

Configuring an inner-outer VLAN 802.1p priority mapping

The A5800&A5820X switches series achieve the following inner-outer VLAN 802.1p priority mapping modes through QoS policies:

- 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.

If you set the trusted packet priority type to 802.1p priority on a port with QinQ or selective QinQ enabled, the port automatically copies the 802.1p priority from the inner VLAN tag to the outer VLAN tag when adding the outer VLAN tag to each packet. For more information about the trusted packet priority type, see the *ACL and QoS Configuration Guide*.

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
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>classifier-name</i> [operator { and or }]	Required. By default, the operator of a class is AND.
3. Configure a match criterion	Configure a match criterion to match the specified inner VLAN IDs if-match customer-vlan-id <i>vlan-id-list</i>	Use either command.
	Configure a match criterion to match the specified inner VLAN tag priorities if-match customer-dot1p <i>8021p-list</i>	
4. Return to system view	quit	—
5. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
6. Configure the action of specifying the outer VLAN tag priorities of	Configure the action of marking the 802.1p priorities in outer VLAN tags remark dot1p <i>8021p</i>	Use either command. Choose to configure inner-outer VLAN 802.1p priority mapping or copying as needed.

To do...		Use the command...	Remarks
packets	Configure the action of copying the 802.1p priorities in the inner VLAN tags to the outer VLAN tags	remark dot1p customer-dot1p-trust	
7.	Return to system view	quit	—
8.	Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
9.	Associate the traffic class with the traffic behavior defined earlier	classifier <i>classifier-name</i> behavior <i>behavior-name</i>	Required.
10.	Return to system view	quit	—
11.	Enter the Ethernet port view of the customer network-side port	interface <i>interface-type interface-number</i>	Required. Use either command. <ul style="list-style-type: none"> • 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.
	Enter port group view	port-group manual <i>port-group-name</i>	
12.	Enable basic QinQ	qinq enable	Required.
13.	Apply the QoS policy to the incoming traffic	qos apply policy <i>policy-name</i> inbound	Required.

Configuring inner VLAN ID substitution

Basic QinQ does not change the inner VLAN ID when tagging the customer VLAN frame with an outer VLAN tag. Selective QinQ can change the inner VLAN ID when tagging the customer VLAN frame with an outer VLAN tag according to the inner VLAN ID substitution function you configured.

To configure inner VLAN ID substitution:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>classifier-name</i> [operator { and or }]	Required. By default, the operator of a class is AND.
3. Configure a match criterion to match the specified inner VLAN IDs	if-match customer-vlan-id <i>vlan-id-list</i>	Required.
4. Return to system view	quit	—
5. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
6. Configure the action of marking the inner VLAN IDs	remark customer-vlan-id <i>vlan-id</i>	Required.
7. Return to system view	quit	—
8. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
9. Associate the traffic class with the traffic behavior defined earlier	classifier <i>classifier-name</i> behavior <i>behavior-name</i>	Required.
10. Return to system view	quit	—
11. Enter the Ethernet port view of the provider network-side port	Enter Layer 2 Ethernet port view	Required. Use either command.
	Enter port group view	<ul style="list-style-type: none"> 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.
12. Apply the QoS policy to the outgoing traffic	qos apply policy <i>policy-name</i> outbound	Required.

Configuring the TPID value in VLAN tags

Configuring the TPID value in the CVLAN tag

To configure the TPID value in the CVLAN tag:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Configure the TPID value in the CVLAN tag	qinq ethernet-type customer-tag <i>hex-value</i>	Optional 0x8100 by default

Configuring the TPID value in the SVLAN tag

To configure the TPID value in the SVLAN tag:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter interface view or port group view	Enter Layer 2 Ethernet port view or Layer 2 aggregate interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Required. Use either command.
3. Configure the TPID value in the SVLAN tag	qinq ethernet-type service-tag <i>hex-value</i>	Optional. 0x8100 by default.

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.

For more information about customer-side QinQ, see the chapter “VLAN mapping configuration.”

SVLAN TPID and QinQ cannot both be configured on a port at the same time.

QinQ configuration examples

Basic QinQ configuration example

Network requirements

As shown in [Figure 56](#),

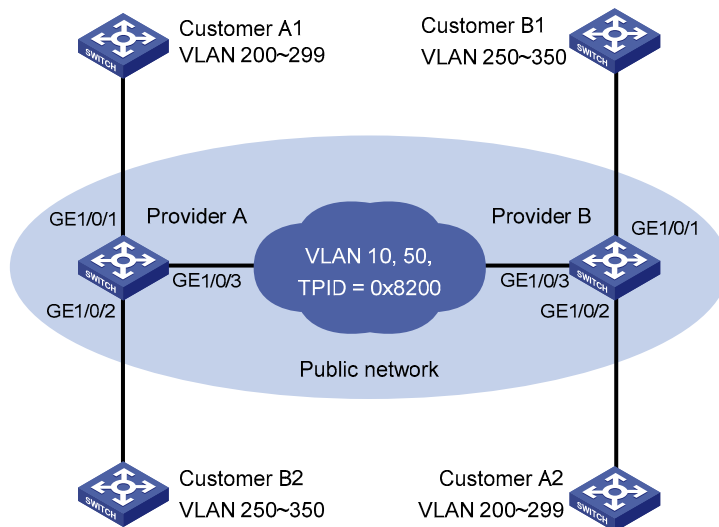
- Provider A and Provider B are edge switches on the service provider network and are interconnected through trunk ports. They belong to SVLAN 10 and 50.
- Customer A1, Customer A2, Customer B1 and Customer B2 are edge switches on the customer network.

- Third-party switches with a TPID value of 0x8200 are deployed between Provider A and Provider B.

Make 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 56 Network diagram for basic QinQ configuration



Configuration procedure

Make sure that the switches in the service provider network have been configured to allow QinQ packets to pass through.

1. Configuration on Provider A

- Configure GigabitEthernet 1/0/1

Configure VLAN 10 as the default VLAN 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 default VLAN 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] qinq ethernet-type service-tag 8200
```

```
[ProviderA-GigabitEthernet1/0/3] quit
```

2. Configuration on Provider B

- **Configure GigabitEthernet 1/0/1**

Configure VLAN 50 as the default VLAN 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 default VLAN 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] qinq ethernet-type service-tag 8200
```

```
[ProviderB-GigabitEthernet1/0/3] quit
```

3. Configuration on third-party switches

Configure the third-party switches between Provider A and Provider B as follows: configure the port connecting GigabitEthernet 1/0/3 of Provider A and that connecting 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

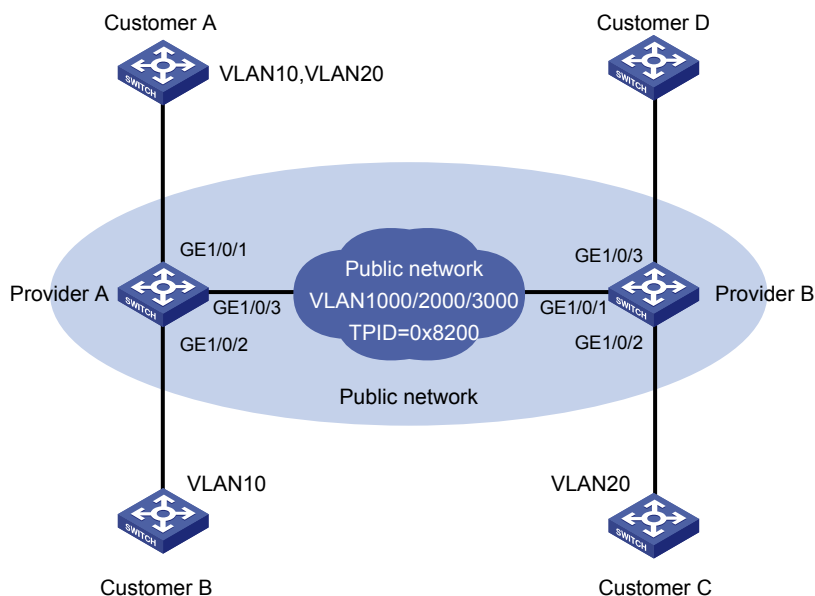
As shown in Figure 57,

- Provider A and Provider B are service provider network access switches.
- Customer A, Customer B, Customer C, and Customer D are customer network access switches.
- Provider A and Provider B are interconnected through a trunk port, which permits the frames of VLAN 1000, VLAN 2000, and VLAN 3000 to pass through.
- Third-party switches are deployed between Provider A and Provider B, with a TPID value of 0x8200.

Make configuration to satisfy the following requirements:

- VLAN 10 of Customer A and Customer B can intercommunicate across VLAN 1000 on the public network.
- VLAN 20 of Customer A and Customer C can intercommunicate across VLAN 2000 on the public network.
- Frames of the VLANs other than VLAN 10 and VLAN 20 of Customer A can be forwarded to Customer D across VLAN 3000 on the public network.

Figure 57 Network diagram for selective QinQ configuration



Configuration procedure

Make sure that the switches in the service provider network have been configured to allow QinQ packets to pass through.

1. Configuration on Provider A

Enter system view.

```
<ProviderA> system-view
```

- Configuration on GigabitEthernet 1/0/1

Configure the port as a hybrid port permitting frames of VLAN 1000, VLAN 2000, and VLAN 3000 to pass through with the outer VLAN tag removed.

```
[ProviderA] interface gigabitethernet 1/0/1
[ProviderA-GigabitEthernet1/0/1] port link-type hybrid
[ProviderA-GigabitEthernet1/0/1] port hybrid vlan 1000 2000 3000 untagged
```

Configure VLAN 3000 as the default VLAN of GigabitEthernet 1/0/1, and enable basic QinQ on GigabitEthernet 1/0/1. As a result, the frames received on the port are tagged with the outer VLAN tag 3000.

```
[ProviderA-GigabitEthernet1/0/1] port hybrid pvid vlan 3000
[ProviderA-GigabitEthernet1/0/1] qinq enable
[ProviderA-GigabitEthernet1/0/1] quit
```

Create a class **A10** to match frames of VLAN 10 of Customer A.

```
[ProviderA] traffic classifier A10
[ProviderA-classifier-A10] if-match customer-vlan-id 10
[ProviderA-classifier-A10] quit
```

Create a traffic behavior **P1000** and configure the action of tagging frames with the outer VLAN tag 1000 for the traffic behavior.

```
[ProviderA] traffic behavior P1000
[ProviderA-behavior-P1000] nest top-most vlan-id 1000
[ProviderA-behavior-P1000] quit
```

Create a class **A20** to match frames of VLAN 20 of Customer A.

```
[ProviderA] traffic classifier A20
[ProviderA-classifier-A20] if-match customer-vlan-id 20
[ProviderA-classifier-A20] quit
```

Create a traffic behavior **P2000** and configure the action of tagging frames with the outer VLAN tag 2000 for the traffic behavior.

```
[ProviderA] traffic behavior P2000
[ProviderA-behavior-P2000] nest top-most vlan-id 2000
[ProviderA-behavior-P2000] quit
```

Create a QoS policy **qinq**. Associate the class **A10** with the traffic behavior **P1000**, and associate the class **A20** with the traffic behavior **P2000** in the QoS policy **qinq**.

```
[ProviderA] qos policy qinq
[ProviderA-qospolicy-qinq] classifier A10 behavior P1000
[ProviderA-qospolicy-qinq] classifier A20 behavior P2000
[ProviderA-qospolicy-qinq] quit
```

Apply the QoS policy **qinq** to the incoming traffic of GigabitEthernet 1/0/1.

```
[ProviderA] interface GigabitEthernet 1/0/1
[ProviderA-GigabitEthernet1/0/1] qos apply policy qinq inbound
```

- Configuration on GigabitEthernet 1/0/2

Configure VLAN 1000 as the default VLAN.

```
[ProviderA] interface gigabitethernet 1/0/2
[ProviderA-GigabitEthernet1/0/2] port access vlan 1000
```

Enable basic QinQ. Tag frames from VLAN 10 with the outer VLAN tag 1000.

```
[ProviderA-GigabitEthernet1/0/2] qinq enable
[ProviderA-GigabitEthernet1/0/2] quit
```

- Configuration on GigabitEthernet 1/0/3.

Configure the port as a trunk port permitting frames of VLAN 1000, VLAN 2000 and VLAN 3000 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 1000 2000 3000
```

To enable interoperability with the third-party switches in the public network, set the TPID of the service provider network VLAN tags to 0x8200. The port tags the frames with the outer VLAN tag whose TPID is 0x8200.

```
[ProviderA-GigabitEthernet1/0/3] qinq ethernet-type service-tag 8200
```

2. Configuration on Provider B

- Configuration on GigabitEthernet 1/0/1

Configure the port as a trunk port permitting frames of VLAN 1000, VLAN 2000 and VLAN 3000 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 3000
```

To enable interoperability with the third-party switches in the public network, set the TPID of the service provider network VLAN tags to 0x8200. The port tags the received frames with the outer VLAN tag whose TPID is 0x8200.

```
[ProviderB-GigabitEthernet1/0/1] qinq ethernet-type service-tag 8200
[ProviderB-GigabitEthernet1/0/1] quit
```

- Configuration on GigabitEthernet 1/0/2

Configure VLAN 2000 as the default VLAN.

```
[ProviderB] interface GigabitEthernet 1/0/2
[ProviderB-GigabitEthernet1/0/2] port access vlan 2000
```

Enable basic QinQ. Tag frames from VLAN 20 with the outer VLAN tag 2000.

```
[ProviderB-GigabitEthernet1/0/2] qinq enable
[ProviderB-GigabitEthernet1/0/2] quit
```

- Configuration on GigabitEthernet 1/0/3

Configure VLAN 3000 as the default VLAN.

```
[ProviderB] interface GigabitEthernet 1/0/3
[ProviderB-GigabitEthernet1/0/3] port access vlan 3000
```

Enable basic QinQ to tag frames of all customer VLANs with the outer VLAN tag 3000.

```
[ProviderB-GigabitEthernet1/0/3] qinq enable
```

3. Configuration on switches on the public network

Third-party switches are deployed between Provider A and Provider B. This section describes only the basic configuration required on the switches. Configure that switch connecting with GigabitEthernet 1/0/3 of Provider A and the switch connecting with GigabitEthernet 1/0/1 of Provider B so that their ports send tagged frames of VLAN 1000, VLAN 2000 and VLAN 3000. The configuration steps are omitted here.

VLAN mapping configuration

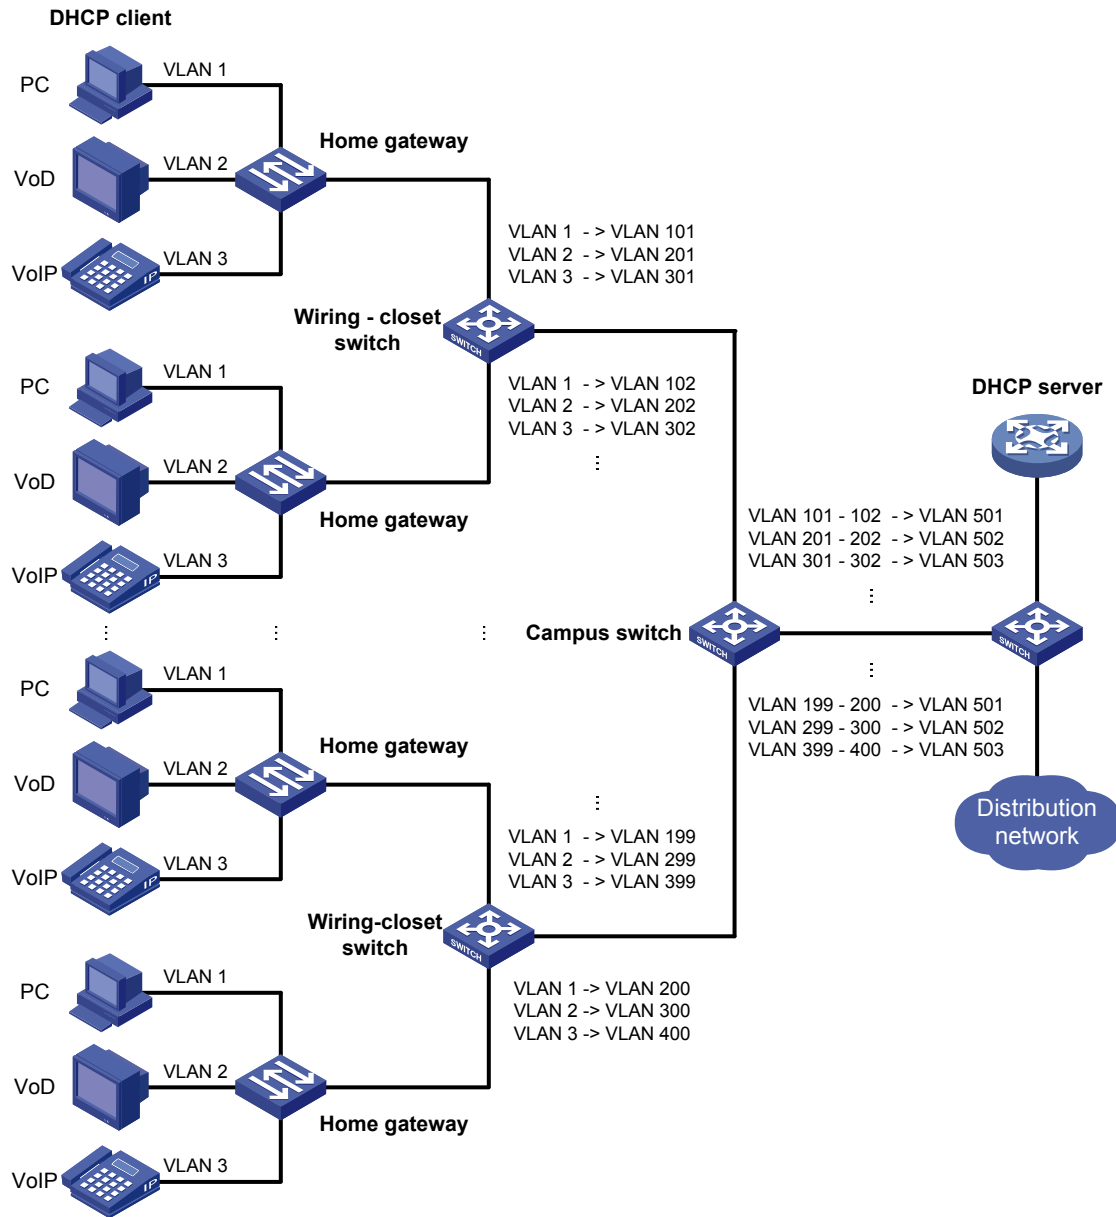
VLAN mapping re-marks VLAN tagged traffic with new VLAN IDs. A5800&A5820X series provides the following types of VLAN mapping:

- One-to-one VLAN mapping—Replaces one VLAN tag with another. Use one-to-one VLAN mapping to sub-classify traffic from a particular VLAN for granular QoS control.
- Many-to-one VLAN mapping—Replaces multiple VLAN tags with the same VLAN tag. Use many-to-one VLAN mapping to aggregate traffic from different VLANs to regulate the aggregate traffic as a whole. Many-to-one VLAN mapping is usually used together with one-to-one VLAN mapping.
- Two-to-two VLAN mapping—Replaces the outer and inner VLAN IDs of double tagged traffic with a new pair of VLAN IDs. Use two-to-two VLAN mapping to enable two remote sites in different VLANs to communicate at Layer 2 across two service provider networks that use different VLAN assignment schemes.

Application scenario of one-to-one and many-to-one VLAN mapping

Figure 58 shows a typical application scenario in which each home gateway uses different VLANs to transmit the PC, VoD, and VoIP services.

Figure 58 Application scenario of one-to-one and many-to-one VLAN mapping




To further sub-classify each type of traffic by customer, perform one-to-one VLAN mapping on the wiring-closet switches, assigning a separate VLAN for each type of traffic from each customer. The required total number of VLANs in the network can be very large. To prevent the maximum number of VLANs from being exceeded on the distribution layer device, perform many-to-one VLAN mapping on the campus switch to assign the same type of traffic from different customers to the same VLAN.

Application scenario of two-to-two VLAN mapping

Figure 59 shows a typical application scenario in which two remote sites of VPN A, Site 1 and Site 2, must communicate across two SP networks, SP 1 and SP 2.

Figure 59 Application scenario of two-to-two VLAN mapping



Site 1 and Site 2 are in VLAN 2 and VLAN 3, respectively. The VLAN assigned for VPN A is VLAN 10 in the SP 1 network and VLAN 20 in the SP 2 network.

If Site 1 sends a packet to Site 2, the packet is processed on the way to its destination using the following workflow:


1. When the packet tagged with VLAN 2 arrives at the edge of network SP 1, PE 1 tags the packet with outer VLAN 10 by using QinQ or selective QinQ.
2. When the double-tagged packet enters the SP 2 network, PE 3 replaces the outer VLAN tag (VLAN 10) with VLAN 20. Because the packet is destined for Site 2 in VLAN 3, PE 3 also replaces the inner tag (VLAN 2) of the packet with VLAN 3. This process is two-to-two VLAN mapping.
3. When PE4 receives the packet with the new VLAN tag pair, it removes the outer VLAN tag and forwards the packet to VLAN 3.

For more information about QinQ and selective QinQ, see the chapter "QinQ configuration."

Concepts and terms

Figure 60 shows a simplified network to help explain the concepts and terms that you may encounter when working with VLAN mapping.

Figure 60 Basic concepts of VLAN mapping



- Uplink traffic: Traffic transmitted from the customer network to the service provider network.
- Downlink traffic: Traffic transmitted from the service provider network to the customer network.
- Network-side port: A port connected to the service provider network.
- Customer-side port: A port connected to the customer network.
- Uplink policy: A QoS policy that defines VLAN mapping rules for uplink traffic.
- Downlink policy: A QoS policy that defines VLAN mapping rules for downlink traffic.
- CVLANs: VLANs assigned for customers.
- SVLANs: VLANs assigned for transmitting traffic across the service provider network.

For more information about QoS policies, see the *ACL and QoS Configuration Guide*.

VLAN mapping implementations

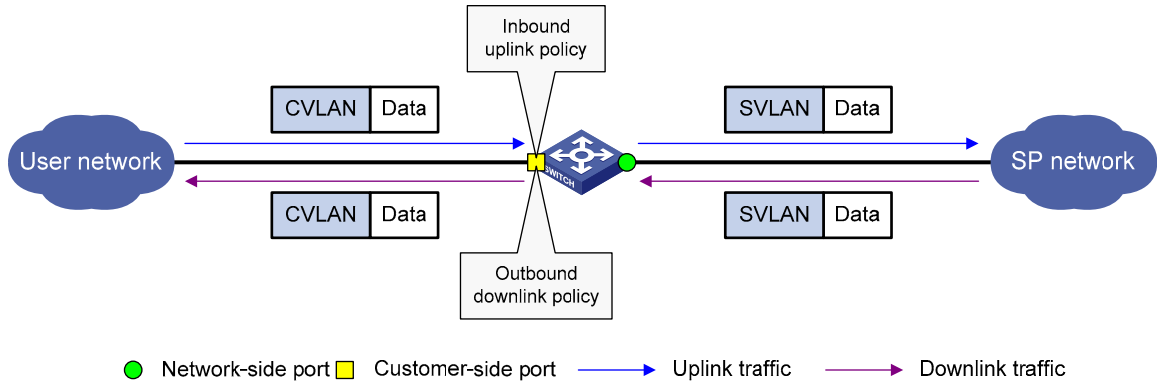
This section describes how VLAN mapping is implemented on your device.

One-to-one VLAN mapping

Implement one-to-one VLAN mapping on the customer-side port through the following configurations, as shown in Figure 61:

- Apply an uplink policy to the incoming traffic, mapping each CVLAN ID to a unique SVLAN ID. When a packet arrives, the switch replaces its CVLAN ID with the matching SVLAN ID.
- Apply a downlink policy to the outgoing traffic, mapping each SVLAN ID back to the corresponding CVLAN ID. When forwarding a packet out of the port, the switch replaces its SVLAN ID with the matching CVLAN ID.

Figure 61 One-to-one VLAN mapping implementation

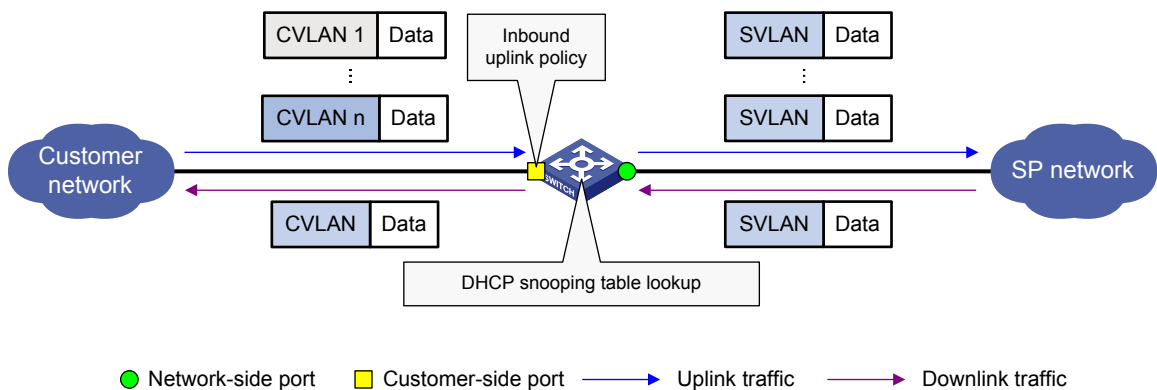


Many-to-one VLAN mapping

Implement many-to-one VLAN mapping through the following configurations, as shown in [Figure 62](#):

- Apply an uplink policy to the incoming traffic on the customer-side port to map different CVLAN IDs to one SVLAN ID. When a packet arrives, the switch replaces its CVLAN tag with the matching SVLAN tag.
- Configure the network-side port as a DHCP snooping trusted port. For downlink traffic, the switch looks through the DHCP snooping table, and replaces the SVLAN ID with the CVLAN ID found in the table.

Figure 62 Many-to-one VLAN mapping implementation



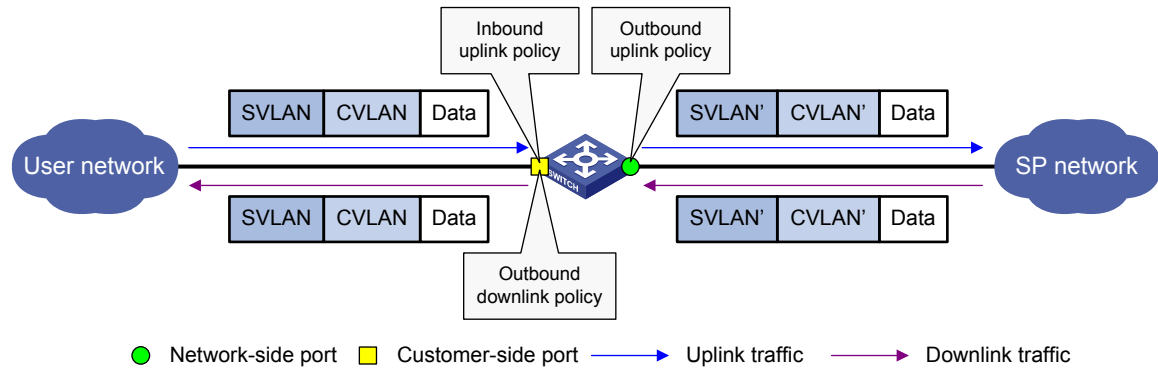
Each DHCP snooping entry contains information about one DHCP client, including its IP address, MAC address, and CVLAN. For more information about DHCP snooping, see the *Layer 3—IP Services Configuration Guide*.

Two-to-two VLAN mapping

Implement two-to-two VLAN mapping through the following configurations, as shown in [Figure 63](#).

- For uplink traffic, apply an inbound policy on the customer-side port to replace the SVLAN with a new SVLAN, and apply an outbound policy on the network-side port to replace the CVLAN with a new CVLAN.
- For downlink traffic, apply an outbound policy on the customer-side port to replace the double tags with the original VLAN tag pair.

Figure 63 Two-to-two VLAN mapping implementation



Configuring VLAN mapping

Use the VLAN mapping methods as appropriate to the roles of your switches in the network, as described in this table:

Task	Switch role
Configuring one-to-one VLAN mapping	Wiring-closet switch (see Figure 58)
Configuring many-to-one VLAN mapping	Campus switch (see Figure 58)
Configuring two-to-two VLAN mapping	Edge switch between SP networks, for example, PE 3 in Figure 59

Configuring one-to-one VLAN mapping

Perform one-to-one VLAN mapping on wiring-closet switches (see [Figure 58](#)) to isolate traffic by both user and traffic type.

Perform these tasks to configure one-to-one VLAN mapping:

Task	Description
Configuring an uplink policy	Creates CVLAN-to-SVLAN mappings (required).
Configuring a downlink policy	Creates SVLAN-to-CVLAN mappings (required).
Configuring the customer-side port	Configures settings required for one-to-one VLAN mapping (required).
Configuring the network-side port	Configures VLAN settings required for normal communication (required).

Configuration prerequisites

Create CVLANs and SVLANs, and plan CVLAN-SVLAN mappings.

Configuring an uplink policy

To configure an uplink policy to map each CVLAN to a unique SVLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>tcl-name</i> [operator { and or }]	Required.
3. Specify a CVLAN as the match criterion	if-match customer-vlan-id <i>vlan-id</i>	Repeat these steps to configure one class for each CVLAN.
4. Return to system view	quit	
5. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
6. Configure an SVLAN marking action	remark service-vlan-id <i>vlan-id</i>	Repeat these steps to configure one behavior for each SVLAN.
7. Return to system view	quit	
8. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
9. Associate the class with the behavior to map the CVLAN to the SVLAN	classifier <i>tcl-name</i> behavior <i>behavior-name</i>	Required. Repeat these steps to create other CVLAN-to-SVLAN mappings.

Configuring a downlink policy

To configure a downlink policy to map SVLANs back to CVLANs:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>tcl-name</i> [operator { and or }]	Required.
3. Configure an SVLAN as the match criterion	if-match service-vlan-id <i>vlan-id</i>	Repeat these steps to configure one class for each SVLAN.
4. Return to system view	quit	
5. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
6. Configure a CVLAN marking action	remark customer-vlan-id <i>vlan-id</i>	Repeat these steps to configure a behavior for each CVLAN.
7. Return to system view	quit	
8. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
9. Associate the class with the behavior to map the SVLAN to the CVLAN	classifier <i>tcl-name</i> behavior <i>behavior-name</i>	Required. Repeat these steps to create other CVLAN-to-SVLAN mappings.

Configuring the customer-side port

To configure the customer-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure the port as a trunk port	port link-type trunk	Required. The default link type of an Ethernet port is access.
4. Assign the port to CVLANs and SVLANs	port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port is in only VLAN 1.
5. Enable basic QinQ	qinq enable	Required. By default, basic QinQ is disabled.
6. Apply the uplink policy to the incoming traffic	qos apply policy <i>policy-name inbound</i>	Required.
7. Apply the downlink policy to the outgoing traffic	qos apply policy <i>policy-name outbound</i>	Required.

Configuring the network-side port

To configure the network-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure the port as a trunk port	port link-type trunk	Required. The default link type of ports is access.
4. Assign the port to SVLANs	port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port is in only VLAN 1.

Configuring many-to-one VLAN mapping

⚠ CAUTION:

Before changing VLAN mappings on a port, clear all DHCP snooping entries with the **reset dhcp-snooping** command (see *Layer 3—IP Services Command Reference*), or re-enable the dynamic address binding function of the IP Source Guard module on the port by using the **undo ip check source** command and then the **ip check source** command (see *Security Command Reference*).

Perform many-to-one VLAN mapping on campus switches (see [Figure 58](#)) to transmit the same type of traffic from different users in one VLAN.

Perform these tasks to configure many-to-one VLAN mapping:

Task	Description
Enabling DHCP snooping	Enables DHCP snooping globally (required).
Enabling ARP detection in SVLANs	Enables ARP detection in all SVLANs (required).
Configuring an uplink policy	Configures an uplink policy for the customer-side port (required).
Configuring the customer-side port	Configures VLAN and other settings required for many-to-one VLAN mapping (required).
Configuring the network-side port	Configures VLAN and other settings required for many-to-one VLAN mapping (required).

Configuration prerequisites

Before you configure many-to-one VLAN mapping, complete the following tasks:

- Make sure that all home users use DHCP to get IP addresses. For how to assign IP addresses through DHCP, see the *Layer 3—IP Services Configuration Guide*.
- Create CVLANs and SVLANs, and plan CVLANs-to-SVLAN mappings.

Enabling DHCP snooping

To enable DHCP snooping:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable DHCP snooping	dhcp-snooping	Required Disabled by default

Enabling ARP detection in SVLANs

The ARP detection function enables a switch to modify the VLAN attributes of ARP packets, which is impossible under the normal ARP packet processing procedure. For more information about ARP detection, see the *Security Configuration Guide*.

To enable ARP detection in all SVLANs:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter VLAN view	vlan <i>vlan-id</i>	—
3. Enable ARP detection	arp detection enable	Required Disabled by default

To defend against ARP attacks, enable ARP detection also in all CVLANs.

Configuring an uplink policy

To configure an uplink policy to map a group of CVLANs to one SVLAN:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>tcl-name</i> operator or	Required.
3. Configure multiple CVLANs as match criteria	if-match customer-vlan-id { <i>vlan-id-list</i> <i>vlan-id1 to vlan-id2</i> }	Repeat these steps to configure one class for each group of CVLANs.
4. Return to system view	quit	
5. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
6. Configure an SVLAN marking action	remark service-vlan-id <i>vlan-id</i>	Repeat these steps to configure one behavior for each SVLAN.
7. Return to system view	quit	
8. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
9. Map the CVLANs to the SVLAN by associating the class with the behavior	classifier <i>tcl-name</i> behavior <i>behavior-name</i> mode dot1q-tag-manipulation	Required. Repeat this step to create other CVLANs-to-SVLAN mappings.

Configuring the customer-side port

To configure the customer-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure the port as a trunk port	port link-type trunk	Required. The default link type of an Ethernet port is access.
4. Assign the port to CVLANs and SVLANs	port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port is in only VLAN 1.
5. Enable customer-side QinQ	qinq enable downlink	Required. By default, customer-side QinQ is disabled on all ports.
6. Apply the uplink policy to the incoming traffic	qos apply policy <i>policy-name</i> inbound	Required.

Before applying a QoS policy to the customer-side port, enable customer-side QinQ on the port. Before disabling customer-side QinQ on the customer-side port, remove the QoS policy from the port first.

Configuring the network-side port

To configure the network-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure the port as a trunk port	port link-type trunk	Required. The default link type of an Ethernet port is access.
4. Assign the port to SVLANs	port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port is in only VLAN 1.
5. Configure the port as a DHCP snooping trusted port	dhcp-snooping trust	Required. By default, all ports are DHCP snooping untrusted ports.
6. Configure the port as an ARP trusted port	arp detection trust	Required. By default, all ports are ARP untrusted ports.
7. Enable network-side QinQ	qinq enable uplink	Required. By default, network-side QinQ is disabled on all ports.

Configuring two-to-two VLAN mapping

Perform two-to-two VLAN mapping on an edge device that connects two SP networks, for example, on PE 3 in [Figure 59](#). Two-to-two VLAN mapping enables two remote sites in different VLANs to communicate at Layer 2 across two service provider networks that use different VLAN assignment schemes.

For the ease of description, the VLAN tags of the double-tagged frames that arrive at the customer-side port are called foreign CVLANs and SVLANs, and the VLAN tags marked by the edge device are called local CVLANs and SVLANs.

Perform these tasks to configure two-to-two VLAN mapping:

Task	Description
Configuring an uplink policy for the customer-side port	Replaces foreign SVLANs with local SVLANs for uplink traffic (required).
Configuring an uplink policy for the network-side port	Replaces foreign CVLANs with local CVLANs for uplink traffic (required).
Configuring a downlink policy for the customer-side port	Replaces local SVLANs and CVLANs with foreign SVLANs and CVLANs (required).
Configuring the customer-side port	Configures VLAN and other settings required for two-to-two VLAN mapping (required).
Configuring the network-side port	Configures VLAN and other settings required for two-to-two VLAN mapping (required).

Configuring an uplink policy for the customer-side port

The uplink policy on the customer-side port modifies the SVLAN ID of incoming traffic.

To configure an uplink policy for the customer-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>tcl-name</i> [operator and]	Required.
3. Specify a foreign CVLAN as a match criterion	if-match customer-vlan-id <i>vlan-id</i>	Repeat these steps to create one class for each foreign CVLAN and SVLAN pair.
4. Specify a foreign SVLAN as a match criterion	if-match service-vlan-id <i>vlan-id</i>	
5. Return to system view	quit	
6. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
7. Configure an SVLAN marking action to replace the foreign SVLAN ID with a local SVLAN ID	remark service-vlan-id <i>vlan-id</i>	Repeat these steps to configure one SVLAN marking action for each CVLAN and SVLAN pair.
8. Return to system view	quit	
9. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
10. Associate the class with the behavior	classifier <i>tcl-name</i> behavior <i>behavior-name</i>	Required. Repeat this step to create other class-behavior associations.

Configuring an uplink policy for the network-side port

The uplink policy on the network-side port modifies the CVLAN ID of incoming traffic.

To configure an uplink policy for the network-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>tcl-name</i> [operator and]	Required.
3. Specify a foreign CVLAN as a match criterion	if-match customer-vlan-id <i>vlan-id</i>	Repeat these steps to create one class for each local SVLAN and foreign CVLAN pair.
4. Specify a local SVLAN as a match criterion	if-match service-vlan-id <i>vlan-id</i>	
5. Return to system view	quit	
6. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
7. Configure a CVLAN marking action to replace the foreign CVLAN ID with a local CVLAN ID	remark customer-vlan-id <i>vlan-id</i>	Repeat these steps to configure one CVLAN marking action for each local SVLAN and foreign CVLAN pair.
8. Return to system view	quit	

To do...	Use the command...	Remarks
9. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
10. Associate the class with the behavior	classifier <i>tcl-name</i> behavior <i>behavior-name</i>	Required. Repeat this step to create other class-behavior associations.

Configuring a downlink policy for the customer-side port

The downlink policy on the customer-side port replaces local SVLAN and CVLAN pairs with foreign SVLAN and CVLAN pairs.

To configure a downlink policy for the customer-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a class and enter class view	traffic classifier <i>tcl-name</i> [operator and]	Required.
3. Specify a local CVLAN as a match criterion	if-match customer-vlan-id <i>vlan-id</i>	Repeat these steps to create one class for each local CVLAN and SVLAN pair.
4. Specify a local SVLAN as a match criterion	if-match service-vlan-id <i>vlan-id</i>	
5. Return to system view	quit	
6. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	
7. Configure a CVLAN marking action to replace the local CVLAN ID with a foreign CVLAN ID	remark customer-vlan-id <i>vlan-id</i>	Required. Repeat these steps to create one VLAN marking behavior for each local CVLAN and SVLAN pair.
8. Configure an SVLAN marking action to replace the local SVLAN ID with a foreign SVLAN ID	remark service-vlan-id <i>vlan-id</i>	
9. Return to system view	quit	
10. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
11. Associate the class with the behavior	classifier <i>tcl-name</i> behavior <i>behavior-name</i>	Required. Repeat this step to create other class-behavior associations.

Configuring the customer-side port

To configure the customer-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure the port as a trunk port	port link-type trunk	Required. The default link type of an Ethernet port is access.
4. Assign the port to the local SVLANs	port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port is in only VLAN 1.
5. Apply the uplink policy configured for the customer-side port to the incoming traffic	qos apply policy <i>policy-name</i> inbound	Required.
6. Apply the downlink policy configured for the customer-side port to the outgoing traffic	qos apply policy <i>policy-name</i> outbound	Required.

Configuring the network-side port

To configure the network-side port:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Configure the port as a trunk port	port link-type trunk	Required. The default link type of an Ethernet port is access.
4. Assign the port to local SVLANs	port trunk permit vlan { <i>vlan-id-list</i> all }	Required. By default, a trunk port is in only VLAN 1.
5. Apply the uplink policy for the network-side port to the outgoing traffic	qos apply policy <i>policy-name</i> outbound	Required.

VLAN mapping configuration examples

One-to-one and many-to-one VLAN mapping configuration example

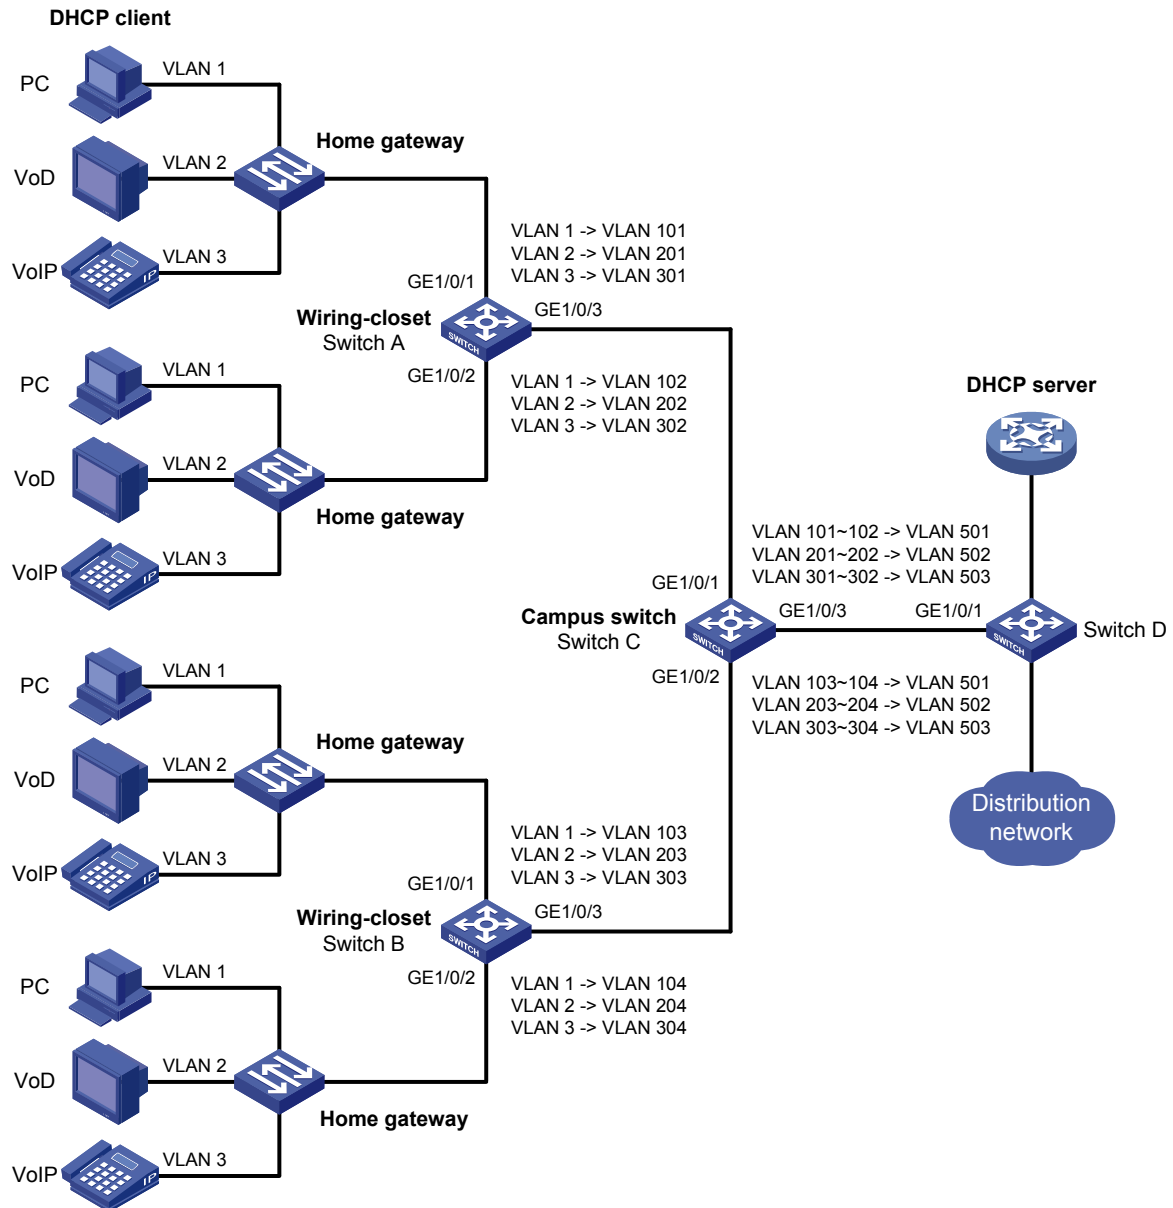
Network requirements

As shown in [Figure 64](#), assign one VLAN for each type of traffic from each user on the wiring-closet switches.

To prevent the maximum number of VLANs from being exceeded on Switch D, perform many-to-one VLAN mapping on Switch C to transmit the same type of traffic from different users in one VLAN: use VLAN 501 for PC traffic, VLAN 502 for VoD traffic, and VLAN 503 for VoIP traffic.

Because Switch C retains customer VLAN information, each type of traffic is still segregated by user, even though it appears to be transmitted in one VLAN.

Figure 64 Network diagram for one-to-one and many-to-one VLAN mapping configuration



Configuration procedure

1. Configuring Switch A

Create the CVLANs and the SVLANs.

```
<SwitchA> system-view
[SwitchA] vlan 2 to 3
[SwitchA] vlan 101 to 102
[SwitchA] vlan 201 to 202
[SwitchA] vlan 301 to 302
```

Configure uplink policies **p1** and **p2** to enable one SVLAN to transmit one service for one customer.

```
[SwitchA] traffic classifier c1
[SwitchA-classifier-c1] if-match customer-vlan-id 1
[SwitchA-classifier-c1] traffic classifier c2
```

```

[SwitchA-classifier-c2] if-match customer-vlan-id 2
[SwitchA-classifier-c2] traffic classifier c3
[SwitchA-classifier-c3] if-match customer-vlan-id 3
[SwitchA-classifier-c3] quit
[SwitchA] traffic behavior b1
[SwitchA-behavior-b1] remark service-vlan-id 101
[SwitchA-behavior-b1] traffic behavior b2
[SwitchA-behavior-b2] remark service-vlan-id 201
[SwitchA-behavior-b2] traffic behavior b3
[SwitchA-behavior-b3] remark service-vlan-id 301
[SwitchA-behavior-b3] traffic behavior b4
[SwitchA-behavior-b4] remark service-vlan-id 102
[SwitchA-behavior-b4] traffic behavior b5
[SwitchA-behavior-b5] remark service-vlan-id 202
[SwitchA-behavior-b5] traffic behavior b6
[SwitchA-behavior-b6] remark service-vlan-id 302
[SwitchA-behavior-b6] quit
[SwitchA] qos policy p1
[SwitchA-policy-p1] classifier c1 behavior b1
[SwitchA-policy-p1] classifier c2 behavior b2
[SwitchA-policy-p1] classifier c3 behavior b3
[SwitchA-policy-p1] quit
[SwitchA] qos policy p2
[SwitchA-policy-p2] classifier c1 behavior b4
[SwitchA-policy-p2] classifier c2 behavior b5
[SwitchA-policy-p2] classifier c3 behavior b6
[SwitchA-policy-p2] quit
# Configure downlink policies p11 and p22 to map the SVLANs back to the CVLANs.
[SwitchA] traffic classifier c11
[SwitchA-classifier-c11] if-match service-vlan-id 101
[SwitchA-classifier-c11] traffic classifier c22
[SwitchA-classifier-c22] if-match service-vlan-id 201
[SwitchA-classifier-c22] traffic classifier c33
[SwitchA-classifier-c33] if-match service-vlan-id 301
[SwitchA-classifier-c33] traffic classifier c44
[SwitchA-classifier-c44] if-match service-vlan-id 102
[SwitchA-classifier-c44] traffic classifier c55
[SwitchA-classifier-c55] if-match service-vlan-id 202
[SwitchA-classifier-c55] traffic classifier c66
[SwitchA-classifier-c66] if-match service-vlan-id 302
[SwitchA-classifier-c66] quit
[SwitchA] traffic behavior b11
[SwitchA-behavior-b11] remark customer-vlan-id 1
[SwitchA-behavior-b11] traffic behavior b22
[SwitchA-behavior-b22] remark customer-vlan-id 2
[SwitchA-behavior-b22] traffic behavior b33
[SwitchA-behavior-b33] remark customer-vlan-id 3
[SwitchA-behavior-b33] quit

```

```
[SwitchA] qos policy p11
[SwitchA-policy-p11] classifier c11 behavior b11
[SwitchA-policy-p11] classifier c22 behavior b22
[SwitchA-policy-p11] classifier c33 behavior b33
[SwitchA-policy-p11] quit
[SwitchA] qos policy p22
[SwitchA-policy-p22] classifier c44 behavior b11
[SwitchA-policy-p22] classifier c55 behavior b22
[SwitchA-policy-p22] classifier c66 behavior b33
[SwitchA-policy-p22] quit
```

Assign customer-side port GigabitEthernet 1/0/1 to CVLANs 1 to 3, and SVLANs 101, 201, and 301, and enable basic QinQ, and apply uplink policy **p1** to the incoming traffic and downlink policy **p11** to the outgoing traffic.

```
[SwitchA] interface gigabitethernet 1/0/1
[SwitchA-GigabitEthernet1/0/1] port link-type trunk
[SwitchA-GigabitEthernet1/0/1] port trunk permit vlan 1 2 3 101 201 301
[SwitchA-GigabitEthernet1/0/1] qinq enable
[SwitchA-GigabitEthernet1/0/1] qos apply policy p1 inbound
[SwitchA-GigabitEthernet1/0/1] qos apply policy p11 outbound
[SwitchA-GigabitEthernet1/0/1] quit
```

Assign customer-side port GigabitEthernet 1/0/2 to CVLANs 1 to 3, and SVLANs 102, 202, and 302, enable basic QinQ, and apply uplink policy **p2** to the incoming traffic and downlink policy **p22** to the outgoing traffic.

```
[SwitchA] interface gigabitethernet 1/0/2
[SwitchA-GigabitEthernet1/0/2] port link-type trunk
[SwitchA-GigabitEthernet1/0/2] port trunk permit vlan 1 2 3 102 202 302
[SwitchA-GigabitEthernet1/0/2] qinq enable
[SwitchA-GigabitEthernet1/0/2] qos apply policy p2 inbound
[SwitchA-GigabitEthernet1/0/2] qos apply policy p22 outbound
[SwitchA-GigabitEthernet1/0/2] quit
```

Assign network-side port GigabitEthernet 1/0/3 to all SVLANs.

```
[SwitchA] interface gigabitethernet 1/0/3
[SwitchA-GigabitEthernet1/0/3] port link-type trunk
[SwitchA-GigabitEthernet1/0/3] port trunk permit vlan 101 201 301 102 202 302
```

Packets sent by PCs are VLAN untagged. If a PC is directly attached to a wiring-closet switch, set the port on the switch as an access or trunk port:

- If the port is set as an access port, assign it to the SVLAN.
- If the port is set as a trunk port, specify the SVLAN as the default VLAN.

2. Configuring Switch B

Configure Switch B in the same procedure as on Switch A.

3. Configuring Switch C

Enable DHCP snooping.

```
<SwitchC> system-view
[SwitchC] dhcp-snooping
```

Create the CVLANs and SVLANs, and enable ARP detection in each VLAN.

```

[SwitchC] vlan 101
[SwitchC-vlan101] arp detection enable
[SwitchC-vlan101] vlan 201
[SwitchC-vlan201] arp detection enable
[SwitchC-vlan201] vlan 301
[SwitchC-vlan301] arp detection enable
[SwitchC-vlan301] vlan 102
[SwitchC-vlan102] arp detection enable
[SwitchC-vlan102] vlan 202
[SwitchC-vlan202] arp detection enable
[SwitchC-vlan202] vlan 302
[SwitchC-vlan302] arp detection enable
[SwitchC-vlan302] vlan 103
[SwitchC-vlan103] arp detection enable
[SwitchC-vlan103] vlan 203
[SwitchC-vlan203] arp detection enable
[SwitchC-vlan203] vlan 303
[SwitchC-vlan303] arp detection enable
[SwitchC-vlan303] vlan 104
[SwitchC-vlan104] arp detection enable
[SwitchC-vlan104] vlan 204
[SwitchC-vlan204] arp detection enable
[SwitchC-vlan204] vlan 304
[SwitchC-vlan304] arp detection enable
[SwitchC-vlan304] vlan 501
[SwitchC-vlan501] arp detection enable
[SwitchC-vlan501] vlan 502
[SwitchC-vlan502] arp detection enable
[SwitchC-vlan502] vlan 503
[SwitchC-vlan503] arp detection enable
[SwitchC-vlan503] quit

```

Configure uplink policies **p1** and **p2** to enable one SVLAN to transmit the same type of traffic from different customers.

```

[SwitchC] traffic classifier c1
[SwitchC-classifier-c1] if-match customer-vlan-id 101 to 102
[SwitchC-classifier-c1] traffic classifier c2
[SwitchC-classifier-c2] if-match customer-vlan-id 201 to 202
[SwitchC-classifier-c2] traffic classifier c3
[SwitchC-classifier-c3] if-match customer-vlan-id 301 to 302
[SwitchC-classifier-c3] traffic classifier c4
[SwitchC-classifier-c4] if-match customer-vlan-id 103 to 104
[SwitchC-classifier-c4] traffic classifier c5
[SwitchC-classifier-c5] if-match customer-vlan-id 203 to 204
[SwitchC-classifier-c5] traffic classifier c6
[SwitchC-classifier-c6] if-match customer-vlan-id 303 to 304
[SwitchC-classifier-c6] quit
[SwitchC] traffic behavior b1
[SwitchC-behavior-b1] remark service-vlan-id 501

```

```

[SwitchC-behavior-b1] traffic behavior b2
[SwitchC-behavior-b2] remark service-vlan-id 502
[SwitchC-behavior-b2] traffic behavior b3
[SwitchC-behavior-b3] remark service-vlan-id 503
[SwitchC-behavior-b3] quit
[SwitchC] qos policy p1
[SwitchC-policy-p1] classifier c1 behavior b1 mode dot1q-tag-manipulation
[SwitchC-policy-p1] classifier c2 behavior b2 mode dot1q-tag-manipulation
[SwitchC-policy-p1] classifier c3 behavior b3 mode dot1q-tag-manipulation
[SwitchC-policy-p1] quit
[SwitchC] qos policy p2
[SwitchC-policy-p2] classifier c4 behavior b1 mode dot1q-tag-manipulation
[SwitchC-policy-p2] classifier c5 behavior b2 mode dot1q-tag-manipulation
[SwitchC-policy-p2] classifier c6 behavior b3 mode dot1q-tag-manipulation
[SwitchC-policy-p2] quit

```

Assign customer-side port GigabitEthernet 1/0/1 to CVLANs 101, 201, 301, 102, 202, 302, and SVLANs 501 to 503. On this port, also enable customer-side QinQ, and apply uplink policy **p1** to the incoming traffic.

```

[SwitchC] interface gigabitethernet 1/0/1
[SwitchC-GigabitEthernet1/0/1] port link-type trunk
[SwitchC-GigabitEthernet1/0/1] port trunk permit vlan 101 201 301 102 202 302 501 502 503
[SwitchC-GigabitEthernet1/0/1] qinq enable downlink
[SwitchC-GigabitEthernet1/0/1] qos apply policy p1 inbound
[SwitchC-GigabitEthernet1/0/1] quit

```

Assign customer-side port GigabitEthernet 1/0/2 to CVLANs 103, 203, 303, 104, 204, 304, and SVLANs 501 to 503. On this port, also enable customer-side QinQ, and apply uplink policy **p2** to the incoming traffic.

```

[SwitchC] interface gigabitethernet 1/0/2
[SwitchC-GigabitEthernet1/0/2] port link-type trunk
[SwitchC-GigabitEthernet1/0/2] port trunk permit vlan 103 203 303 104 204 304 501 502 503
[SwitchC-GigabitEthernet1/0/2] qinq enable downlink
[SwitchC-GigabitEthernet1/0/2] qos apply policy p2 inbound
[SwitchC-GigabitEthernet1/0/2] quit

```

Assign network-side GigabitEthernet 1/0/3 to SVLANs 501 to 503, set the port as a DHCP and ARP trusted port, and enable network-side QinQ.

```

[SwitchC] interface gigabitethernet 1/0/3
[SwitchC-GigabitEthernet1/0/3] port link-type trunk
[SwitchC-GigabitEthernet1/0/3] port trunk permit vlan 501 502 503
[SwitchC-GigabitEthernet1/0/3] dhcp-snooping trust
[SwitchC-GigabitEthernet1/0/3] arp detection trust
[SwitchC-GigabitEthernet1/0/3] qinq enable uplink

```

4. Configuring Switch D

Enable DHCP snooping.

```

<SwitchD> system-view
[SwitchD] dhcp-snooping

```

Assign port GigabitEthernet 1/0/1 to SVLANs 501 to 503.

```
[SwitchD] interface gigabitethernet 1/0/1
[SwitchD-GigabitEthernet1/0/1] port link-type trunk
[SwitchD-GigabitEthernet1/0/1] port trunk permit vlan 501 502 503
```

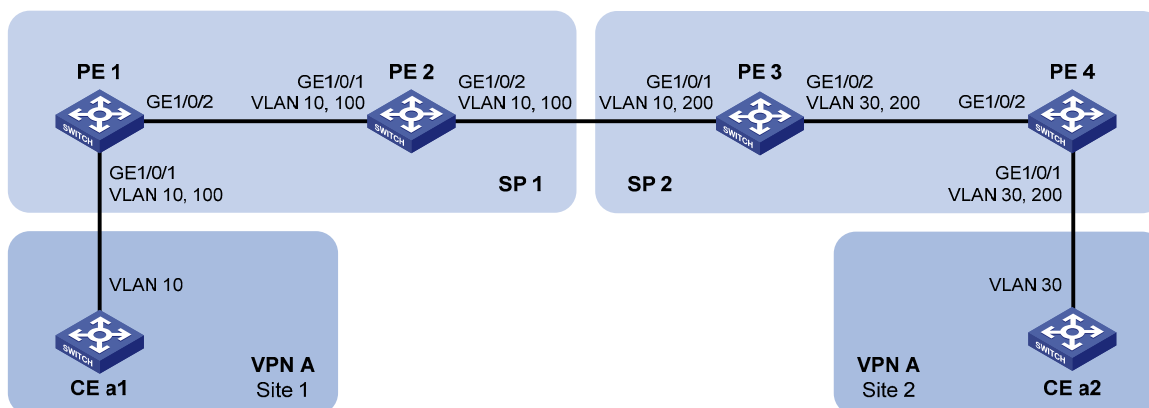
Two-to-two VLAN mapping configuration example

Network requirements

As shown in [Figure 65](#), two VPN A users, Site 1 and Site 2, are in VLAN 10 and VLAN 30, respectively. SP 1 assigns VLAN 100 for VPN A, and SP 2 assigns VLAN 200 for VPN A.

Configure QinQ and two-to-two VLAN mappings to enable the two sites to communicate across networks SP 1 and SP 2.

Figure 65 Network diagram for the two-to-two VLAN mapping application



Configuration procedure

1. Configuring PE 1

Configure QinQ function on GigabitEthernet 1/0/1 to add outer VLAN tag 100 to the traffic tagged with VLAN 10.

```
<PE1> system-view
[PE1] interface gigabitethernet 1/0/1
[PE1-GigabitEthernet1/0/1] port access vlan 100
[PE1-GigabitEthernet1/0/1] qinq enable
[PE1-GigabitEthernet1/0/1] quit
```

Configure the uplink port GigabitEthernet 1/0/2 to permit frames of VLAN 100 to pass through.

```
[PE1] interface gigabitethernet 1/0/2
[PE1-GigabitEthernet1/0/2] port link-type trunk
[PE1-GigabitEthernet1/0/2] port trunk permit vlan 100
```

2. Configuring PE 2

Set port GigabitEthernet 1/0/1 as a trunk port, and assign it to VLAN 100.

```
<PE2> system-view
[PE2] interface gigabitethernet 1/0/1
[PE2-GigabitEthernet1/0/1] port link-type trunk
[PE2-GigabitEthernet1/0/1] port trunk permit vlan 100
[PE2-GigabitEthernet1/0/1] quit
```


Set port GigabitEthernet 1/0/2 as a trunk port, and assign it to VLAN 100.

```
[PE2] interface gigabitethernet 1/0/2
[PE2-GigabitEthernet1/0/2] port link-type trunk
[PE2-GigabitEthernet1/0/2] port trunk permit vlan 100
```

3. Configuring PE 3

Configure an uplink policy **down_uplink** for customer-side port GigabitEthernet 1/0/1 to substitute SVLAN ID 200 for the SVLAN ID in the incoming traffic tagged with CVLAN 10 and SVLAN 100.

```
<PE3> system-view
[PE3] traffic classifier down_uplink
[PE3-classifier-down_uplink] if-match customer-vlan-id 10
[PE3-classifier-down_uplink] if-match service-vlan-id 100
[PE3-classifier-down_uplink] quit
[PE3] traffic behavior down_uplink
[PE3-behavior-down_uplink] remark service-vlan-id 200
[PE3-behavior-down_uplink] quit
[PE3] qos policy down_uplink
[PE3-qospolicy-down_uplink] classifier down_uplink behavior down_uplink
[PE3-qospolicy-down_uplink] quit
```

Configure a downlink policy **down_downlink** for customer-side port GigabitEthernet 1/0/1 to substitute CVLAN 10 and SVLAN 100 for traffic tagged with CVLAN 30 and SVLAN 200.

```
[PE3] traffic classifier down_downlink
[PE3-classifier-down_downlink] if-match customer-vlan-id 30
[PE3-classifier-down_downlink] if-match service-vlan-id 200
[PE3-classifier-down_downlink] quit
[PE3] traffic behavior down_downlink
[PE3-behavior-down_downlink] remark customer-vlan-id 10
[PE3-behavior-down_downlink] remark service-vlan-id 100
[PE3-behavior-down_downlink] quit
[PE3] qos policy down_downlink
[PE3-qospolicy-down_downlink] classifier down_downlink behavior down_downlink
[PE3-qospolicy-down_downlink] quit
```

Configure an uplink policy **up_uplink** for network-side port GigabitEthernet 1/0/2 to substitute CVLAN 30 for the CVLAN ID of the outgoing traffic tagged with CVLAN 10 and SVLAN 200.

```
[PE3] traffic classifier up_uplink
[PE3-classifier-up_uplink] if-match customer-vlan-id 10
[PE3-classifier-up_uplink] if-match service-vlan-id 200
[PE3-classifier-up_uplink] quit
[PE3] traffic behavior up_uplink
[PE3-behavior-up_uplink] remark customer-vlan-id 30
[PE3-behavior-up_uplink] quit
[PE3] qos policy up_uplink
[PE3-qospolicy-up_uplink] classifier up_uplink behavior up_uplink
[PE3-qospolicy-up_uplink] quit
```

Set customer-side port GigabitEthernet 1/0/1 as a trunk port, assign it to VLAN 200, and apply uplink policy **down_uplink** to the incoming traffic and downlink policy **down_downlink** to the outgoing traffic on the port.

```
[PE3] interface gigabitethernet 1/0/1
```

```
[PE3-GigabitEthernet1/0/1] port link-type trunk
[PE3-GigabitEthernet1/0/1] port trunk permit vlan 200
[PE3-GigabitEthernet1/0/1] qos apply policy down_uplink inbound
[PE3-GigabitEthernet1/0/1] qos apply policy down_downlink outbound
[PE3-GigabitEthernet1/0/1] quit
```

Set network-side port GigabitEthernet 1/0/2 as a trunk port, assign it to VLAN 200, and apply uplink policy **up_uplink** to the outgoing traffic on the port.

```
[PE3] interface gigabitethernet 1/0/2
[PE3-GigabitEthernet1/0/2] port link-type trunk
[PE3-GigabitEthernet1/0/2] port trunk permit vlan 200
[PE3-GigabitEthernet1/0/2] qos apply policy up_uplink outbound
[PE3-GigabitEthernet1/0/2] quit
```

4. Configuring PE 4

Configure QinQ function on GigabitEthernet 1/0/1 to add outer VLAN tag 200 to the traffic tagged with VLAN 30.

```
<DeviceD> system-view
[DeviceD] interface gigabitethernet 1/0/1
[DeviceD-GigabitEthernet1/0/1] port access vlan 200
[DeviceD-GigabitEthernet1/0/1] qinq enable
```

Configure GigabitEthernet 1/0/2 to permit frames of VLAN 200 to pass through.

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

LLDP configuration

In a heterogeneous network, it is important that different types of network devices from different vendors can discover one another and exchange configuration for interoperability and management sake. A standard configuration exchange platform was created.

The IETF drafted the 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 LLDPDUs to the directly connected devices, and at the same time, stores the device information received in LLDPDUs sent from the LLDP neighbors in a standard MIB. It allows a network management system to fast detect Layer-2 network topology change and identify what the change is.

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

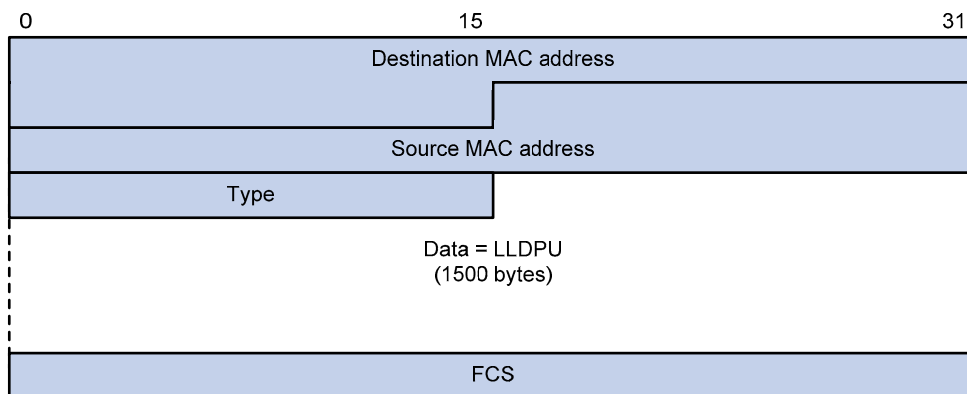
Basic concepts

LLDPDUs

LLDP sends device information in LLDPDUs. LLDPDUs are encapsulated in Ethernet II or SNAP frames.

1. Ethernet II-encapsulated LLDPDU format

Figure 66 Ethernet II-encapsulated LLDPDU format



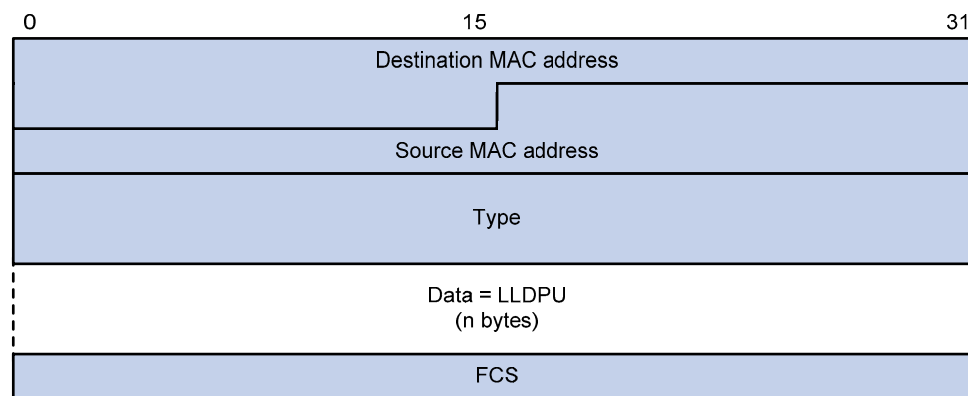
The fields in the frame are described in [Table 23](#):

Table 23 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	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 67 SNAP-encapsulated LLDPDU format



The fields in the frame are described in [Table 24](#):


Table 24 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

LLDPUs

LLDP uses LLDPUs to exchange information. An LLDPDU comprises multiple TLV sequences. Each carries a specific type of device information, as shown in [Figure 68](#).

Figure 68 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 these categories: 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 23](#) lists the basic management TLV types currently in use. Some of them must be included in every LLDPDU.

Table 25 Basic LLDP TLVs

Type	Description	Remarks
Chassis ID	Bridge MAC address of the sending device	
Port ID	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 in case the port does not have a MAC address. If the LLDPDU carries no LLDP-MED TLVs, the port ID TLV carries the port name.	Mandatory
Time To Live	Life of the transmitted information on the receiving device	
End of LLDPDU	Marks the end of the TLV sequence in the LLDPDU	
Port Description	Port description of the sending port	
System Name	Assigned name of the sending device	
System Description	Description of the sending device	
System Capabilities	Identifies the primary functions of the sending device and the enabled primary functions	Optional
Management Address	Management address, and the interface number and OID (object identifier) associated with the address	

2. IEEE 802.1 organizationally specific TLVs

Table 26 IEEE 802.1 organizationally specific TLVs

Type	Description
Port VLAN ID	Port's VLAN identifier. An LLDPDU carries only one TLV of this type.
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	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.
DCBX	Data center bridging exchange protocol

HP A5800&A5820X series Ethernet switches only support receiving protocol identity TLVs.

Layer 3 Ethernet ports do not support IEEE 802.1 organizationally specific TLVs.

3. IEEE 802.3 organizationally specific TLVs

Table 27 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 PoE type, which can be PSE or 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	Power state control configured on the sending port, including the power type of the PSE/PD, PoE sourcing/receiving priority, and PoE sourcing/receiving power.

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 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 28](#):

Table 28 LLDP-MED TLVs

Type	Description
LLDP-MED Capabilities	Allows a network device to advertise the LLDP-MED TLVs it supports
Network Policy	Allows a network device or terminal device to advertise VLAN ID of the specific port, VLAN type, and the Layer 2 and Layer 3 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 management address of a device is used by the network management system to identify and manage the device for topology maintenance and network management. The management address is encapsulated in the management address TLV.

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.

When the LLDP operating mode of a port changes, its LLDP protocol state machine re-initializes. To prevent LLDP from being initialized too frequently during times of frequent operating mode change, you can configure a re-initialization delay. With this delay configured, a port must wait for the specified interval before it can initialize LLDP 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 the network from being overwhelmed by LLDPDUs during times of frequent local device information change, 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 carrying device information new to the local device.
- The LLDP operating mode of the port changes from Disable/Rx to TxRx or Tx.

This is the fast sending mechanism of LLDP. This feature sends a specific number of LLDPDUs at the 1-second interval 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 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 TTL value in the Time To Live TLV carried in the LLDPDU. If the TTL value is zero, the information is aged 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*
- *DCB Capability Exchange Protocol Specification Rev 1.0*
- *DCB Capability Exchange Protocol Base Specification Rev 1.01*

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

Task	Remarks
Configuring CDP compatibility	Optional
Configuring DCBX	Optional
Configuring LLDP trapping	Optional

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.

The Layer 3 Ethernet interface is an Ethernet interface operating in route mode. Set an Ethernet port as a Layer 3 Ethernet interface by using the **port link-mode route** command (see chapter “Ethernet interface configuration”).

Performing basic LLDP configuration

Enabling LLDP

To make LLDP take effect on certain ports, you must enable LLDP both globally and on these ports.

To enable LLDP:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable LLDP globally	lldp enable	Required. By default, LLDP is globally enabled.
3. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Required. Use either command.
4. 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.

To set the LLDP operating mode:

To do...	Use the command...	Remarks	
1. Enter system view	system-view	—	
2. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view	interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view		
3. 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 LLDP operating mode changes on a port.

To set the LLDP re-initialization delay for ports:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Set the LLDP re-initialization delay	lldp timer reinit-delay <i>delay</i>	Optional 2 seconds by default

Enabling LLDP polling

With LLDP polling enabled, a device checks for local configuration changes periodically. Upon detecting a configuration change, the device sends LLDPDUs to inform the neighboring devices of the change.

To enable LLDP polling:

To do...	Use the command...	Remarks	
1. Enter system view	system-view	—	
2. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view	interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view		
3. Enable LLDP polling and set the polling interval	lldp check-change-interval <i>interval</i>	Required. Disabled by default.	

Configuring the advertisable TLVs

To configure the advertisable LLDPDU TLVs on the specified port or ports:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view Enter port group view interface <i>interface-type interface-number</i> port-group manual <i>port-group-name</i>	Required. Use either command.
3. Configure the advertisable TLVs (Layer 2 Ethernet interface view or port group view)	lldp tlv-enable { basic-tlv { all port-description system-capability system-description system-name } dot1-tlv { all dcbx port-vlan-id protocol-vlan-id [<i>vlan-id</i>] vlan-name [<i>vlan-id</i>] } dot3-tlv { all link-aggregation mac-physic max-frame-size power } med-tlv { all capability inventory location-id { civic-address <i>device-type country-code</i> { <i>ca-type ca-value</i> }&<1-10> elin-address <i>tel-number</i> } network-policy power-over-ethernet } }	Optional. By default, all types of LLDP TLVs except DCBX TLVs and location identification TLVs are advertisable on a Layer 2 Ethernet port.
4. Configure the advertisable TLVs (Layer 3 Ethernet interface view)	lldp tlv-enable { basic-tlv { all port-description system-capability system-description system-name } dot3-tlv { all link-aggregation mac-physic max-frame-size power } med-tlv { all capability inventory location-id { civic-address <i>device-type country-code</i> { <i>ca-type ca-value</i> }&<1-10> elin-address <i>tel-number</i> } } power-over-ethernet }	Optional. By default, all types of LLDP TLVs, except IEEE 802.1 organizationally specific TLVs, network policy TLVs, and location identification TLVs, are advertisable on a Layer 3 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.

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
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view <hr/> Enter port group view port-group manual <i>port-group-name</i>	Required. Use either command.
3. Allow LLDP to advertise the management address in LLDPDUs and configure the advertised management address	lldp management-address-tlv [<i>ip-address</i>]	Optional. By default, the management address is sent through LLDPDUs. <ul style="list-style-type: none"> 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. For a Layer 3 Ethernet port, the management address is its own IP address. If no IP address is configured for the Layer 3 Ethernet port, no management address will be advertised.
4. Configure the encoding format of the management address as character string	lldp management-address-format <i>string</i>	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:

$$\text{TTL} = \text{Min} (65,535, (\text{TTL multiplier} \times \text{LLDPDU transmit interval}))$$

As the expression shows, the TTL can be up to 65,535 seconds. TTLs greater than 65,535 will be rounded down to 65,535 seconds.

To change the TTL multiplier:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Set the TTL multiplier	lldp hold-multiplier <i>value</i>	Optional 4 by default
3. Set the LLDPDU transmit interval	lldp timer tx-interval <i>interval</i>	Optional 30 seconds by default
4. Set LLDPDU transmit delay	lldp timer tx-delay <i>delay</i>	Optional 2 seconds by default
5. Set the number of LLDPDUs sent each time fast LLDPDU transmission is triggered	lldp fast-count <i>count</i>	Optional 3 by default

To ensure that the LLDP neighbors can receive LLDPDUs to update information about the current device before it is aged 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 only processes incoming Ethernet II-encapsulated LLDPDUs.
- With SNAP encapsulation configured, an LLDP port sends LLDPDUs in SNAP frames and only processes incoming SNAP-encapsulated LLDPDUs.

By default, LLDPDUs are encapsulated in Ethernet II frames. If the neighbor devices encapsulate LLDPDUs in SNAP frames, configure the encapsulation format for LLDPDUs as SNAP to guarantee normal communication with the neighbors.

To set the encapsulation format for LLDPDUs to SNAP:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Required. Use either command.
3. Set the encapsulation format for LLDPDUs to SNAP	lldp encapsulation snap	Required. Ethernet II encapsulation format applies by default.

LLDP-CDP packets use only SNAP encapsulation.

Configuring CDP compatibility

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, disabling your device from differentiating 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. 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 follows modes:

- TxRx: The CDP packets can be transmitted and received.
- Disable: The CDP packets can neither be transmitted nor be received.

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.

To enable LLDP to be compatible with CDP:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable CDP compatibility globally	lldp compliance cdp	Required. Disabled by default.
3. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view port-group manual <i>port-group-name</i>	
4. Configure CDP-compatible LLDP to operate in TxRx mode	lldp compliance admin-status cdp txrx	Required. Disable mode by default

The maximum TTL value allowed by CDP is 255 seconds. To make CDP-compatible LLDP work properly with Cisco IP phones, make sure that the product of the TTL multiplier and the LLDPDU transmit interval is less than 255 seconds.

Configuring DCBX


DCE, also known as CEE, is enhancement and expansion of traditional Ethernet local area networks for use in data centers. DCE uses the DCBX to negotiate and remotely configure the bridge capability of network elements.

DCBX has two self-adaptable versions, DCB Capability Exchange Protocol Specification Rev 1.0 and DCB Capability Exchange Protocol Base Specification Rev 1.01. DCBX offers the following functions:

- Discovers the peer devices' capabilities and determines whether devices at both ends support these capabilities.
- Detects configuration errors on peer devices.
- Remotely configures the peer device, if the peer device accepts the configuration.

DCBX is supported only on the 10-GE ports of HP A5820X series Ethernet switches, and only the remote configuration function is supported.

Figure 69 DCBX application scenario



DCBX enables lossless packet transmission on DCE networks.

As shown in [Figure 69](#), DCBX applies to a FCoE based data center network, and operates on an access switch. DCBX enables the switch to control the server adapter, and simplifies the configuration and guarantees configuration consistency. DCBX extends LLDP by using the IEEE 802.1 organizationally specific TLVs (DCBX TLVs) to transmit DCBX data, including PFC, ETS, and APP

HP devices can send the three types of DCBX information to a server adapter supporting FCoE, but cannot receive these types of DCBX information.

DCBX configuration task list

Complete these tasks to configure DCBX:

Task	Remarks	
Enabling LLDP and DCBX TLV advertising	Required	
Configuring APP parameters	Required	
Configuring ETS parameters	Configuring the 802.1p-to-local priority mapping	Optional
	Configuring WRR queuing	Optional
Configuring PFC parameters	Required	

Enabling LLDP and DCBX TLV advertising

To enable the device to advertise APP, ETS, and PFC data through an interface, enable LLDP globally and enable LLDP and DCBX TLV advertising on the interface.

To enable LLDP and DCBX TLV advertising:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enable LLDP globally	lldp enable	Required. By default, LLDP is globally enabled.
3. Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view interface <i>interface-type interface-number</i>	Required. Use either command.
	Enter port group view port-group manual <i>port-group-name</i>	
4. Enable LLDP	lldp enable	Optional. By default, LLDP is enabled on an interface.
5. Enable the interface to advertise DCBX TLVs	lldp tlv-enable dot1-tlv dcbx	Optional. Disabled by default.

Configuring APP parameters

The device negotiates with the server adapter by using the APP parameters to control the 802.1p priority values of the protocol packets that the server adapter sends, and to identify traffic based on the 802.1p priority values. For example, the device can use the APP parameters to negotiate with the server adapter to set the 802.1p priority of all FCoE (protocol number 0x8906) and FIP (protocol number 0x8914) packets to 3. If the negotiation succeeds, all FCoE and FIP packets that the server adapter sends to the device carry the 802.1p priority 3.

To configure APP parameters:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create an Ethernet frame header ACL or an IPv4 advanced ACL and enter ACL view	acl number <i>acl-number</i> [name <i>acl-name</i>] [match-order { auto config }]	Required. An Ethernet frame header ACL number ranges from 4000 to 4999. An IPv4 advanced ACL number ranges from 3000 to 3999. DCBX Base Protocol Rev 1.0 supports only Ethernet frame header ACLs. DCBX Base Protocol Rev 1.01 supports both Ethernet frame header ACLs and IPv4 advanced ACLs.
3. Create a rule for the ACL	Create a rule for the Ethernet frame header ACL rule [<i>rule-id</i>] permit type <i>protocol-type</i> ffff	Required. Use either approach.
	Create a rule for the IPv4 advanced ACL rule [<i>rule-id</i>] permit { tcp udp } destination-port eq <i>port</i>	Create rules according to the type of the ACL previously created.
4. Return to system view	quit	—
5. Create a class and specify the operator of the class as OR	traffic classifier <i>tcl-name</i> operator or	Required.
6. Use the specified ACL as the match criterion of the class	if-match acl <i>acl-number</i>	Required.
7. Return to system view	quit	—
8. Create a traffic behavior and enter traffic behavior view	traffic behavior <i>behavior-name</i>	Required.
9. Configure the behavior to mark packets with the specific 802.1p priority	remark dot1p <i>8021p</i>	Required.
10. Return to system view	quit	—
11. Create a QoS policy and enter QoS policy view	qos policy <i>policy-name</i>	Required.
12. Associate the class with the traffic behavior in the QoS policy, and apply the association to DCBX	classifier <i>tcl-name</i> behavior <i>behavior-name</i> mode dcbx	Required.
13. Return to system view	quit	—
14. Apply the QoS policy to the outgoing packets globally	qos apply policy <i>policy-name</i> global outbound	Required. Use any approach.
15. Apply the QoS policy to the	Enter Layer 2 Ethernet interface view interface <i>interface-type</i> <i>interface-number</i>	The global configuration is effective for all interfaces. The

To do...		Use the command...	Remarks
outgoing packets of a Layer 2 Ethernet interface	Apply the policy to outgoing packets	qos apply policy <i>policy-name</i> outbound	configuration made on an interface is effective only for the interface.
16. Apply the QoS policy to the outgoing packets of a port group	Enter port group view Apply the policy to outgoing packets	port-group manual <i>port-group-name</i> qos apply policy <i>policy-name</i> outbound	

For more information about the **acl**, **classifier behavior**, **if-match**, **qos apply policy**, **qos apply policy global**, **qos policy**, **remark dot1p**, **rule**, **traffic behavior**, and **traffic classifier** commands, see the *ACL and QoS Command Reference*.

An Ethernet frame header ACL identifies application protocol packets by protocol number. An IPv4 advanced ACL identifies application protocol packets by IP port number. DCBX Base Protocol Rev 1.0 identifies only application protocol packets by protocol number and advertises TLVs with protocol number 0x8906 (FCoE) only. DCBX Base Protocol Rev 1.01 supports identifying application protocol packets by both protocol number and IP port number, does not restrict the protocol number or IP port number for advertising TLVs, and can advertise up to 77 TLVs according to the remaining length of the current packet.

Configuring ETS parameters

ETS provides committed bandwidth. The device uses ETS parameters to negotiate with the server adapter, controls the server adapter's transmission speed of the specific type of traffic, and guarantees that the transmission speed is within the committed bandwidth of the interface. In this way, no traffic loss occurs due to congestion.

To configure ETS parameters, you must configure the 802.1p-to-local priority mapping and WRR queuing.

Configuring the 802.1p-to-local priority mapping

To configure the 802.1p priority mapping:

To do...		Use the command...	Remarks
1. Enter system view		system-view	—
2. Enter interface view or port group view	Enter Layer 2 Ethernet interface view	interface <i>interface-type</i> <i>interface-number</i>	Required.
	Enter port group view	port-group manual <i>port-group-name</i>	Use either command.
3. Configure the interface or interfaces to use the 802.1p priority in incoming packets for priority mapping		qos trust dot1p	Required. By default, the switch does not use the priority in incoming packets, but the priority of the receiving port as the 802.1p priority of the incoming packets.

To do...	Use the command...	Remarks
4. Return to system view	quit	—
5. Enter priority mapping table view	qos map-table dot1p-lp	—
6. Configure the priority mapping table to map the specific 802.1p priority values to a local precedence value	import <i>import-value-list</i> export <i>export-value</i>	Optional. For information about the default priority mapping tables, see the <i>ACL and QoS Configuration Guide</i> .

For more information about the 802.1p priority, priority trust mode, and port priority, see the *ACL and QoS Configuration Guide*.

For more information about the **qos trust dot1p**, **qos map-table**, and **import** commands, see the *ACL and QoS Command Reference*.

Configuring WRR queuing

Configure WRR queuing to allocate bandwidth.

To configure WRR queuing:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Layer 2 Ethernet interface view interface <i>interface-type</i> <i>interface-number</i>	Required.
	Enter port group view port-group manual <i>port-group-name</i>	Use either command.
3. Enable byte-count WRR queuing	qos wrr byte-count	Optional. By default, byte-count WRR queuing is enabled on ports.
4. Add the specific queue to WRR priority group 1 and configure the scheduling weight for the queue	qos wrr <i>queue-id</i> group 1 byte-count <i>schedule-value</i>	Optional. By default, the scheduling weights of queues 0 through 7 are 1, 2, 3, 4, 5, 9, 13, and 15.

For more information about the **qos wrr** and **qos wrr byte-count** commands, see the *ACL and QoS Command Reference*.

Configuring PFC parameters

To avoid dropping packets with a certain 802.1p priority, you can enable PFC for the 802.1p priority and this feature helps reduce the sending rate of packets carrying this priority when network congestion occurs.

The device uses PFC parameters to negotiate with the server adapter, and to enable PFC for specific 802.1p priorities on the server adapter.

To configure PFC parameters:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—
3. Enable the Ethernet interface to automatically negotiate with its peer to decide whether to enable PFC	priority-flow-control auto	Required. Disabled by default.
4. Enable PFC for specific 802.1p priorities	priority-flow-control no-drop dot1p <i>dot1p-list</i>	Required. By default, all packets with an 802.1p priority are discarded when network congestion occurs.
5. Configure the switch to use the 802.1p priority in incoming packets for priority mapping	qos trust dot1p	Required. By default, the switch does not use the priority in incoming packets, but the priority of the receiving port as the 802.1p priority of the incoming packets.

To advertise the PFC data, you must enable PFC in auto-negotiation mode.

HP recommends you to enable PFC only for the 802.1p priority of the FCoE traffic. Packet loss might occur because of congestion, if you enable PFC for multiple 802.1p priorities.

For more information about the **priority-flow-control** and **priority-flow-control no-drop dot1p** commands, see the *Layer 2—LAN Switching Command Reference*.

For more information about the 802.1p priority, priority trust mode, and port priority, see the *ACL and QoS Configuration Guide*. For more information about the **qos trust dot1p** command, see the *ACL and QoS Command Reference*.

Configuring LLDP trapping

LLDP trapping notifies the NMS of events such as newly-detected neighboring devices and link malfunctions.

To prevent excessive LLDP traps from being sent when topology is unstable, you can set a minimum trap sending interval for LLDP.

To configure LLDP trapping:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Enter Ethernet interface view or port group view	Enter Layer 2/Layer 3 Ethernet interface view interface <i>interface-type interface-number</i> Enter port group view port-group manual <i>port-group-name</i>	Required. Use either command.
3. Enable LLDP trapping	lldp notification remote-change enable	Required. Disabled by default
4. Quit to system view	quit	—
5. Set the interval to send LLDP traps	lldp timer notification-interval <i>interval</i>	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 <i>interface-type interface-number</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the information contained in the LLDP TLVs sent from neighboring devices	display lldp neighbor-information [brief interface <i>interface-type interface-number</i> [brief] list [system-name system-name]] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display LLDP statistics	display lldp statistics [global interface <i>interface-type interface-number</i>] [[{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display LLDP status of a port	display lldp status [interface <i>interface-type interface-number</i>] [[{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display types of advertisable optional LLDP TLVs	display lldp tlv-config [interface <i>interface-type interface-number</i>] [[{ begin exclude include } <i>regular-expression</i>]	Available in any view

LLDP configuration examples


Basic LLDP configuration example

Network requirements

As shown in [Figure 70](#), 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 70 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), and 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), and 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 a 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.

Tear down 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 71](#):

- 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 to be isolated from other types of traffic.

Figure 71 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
Software 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
Software 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.

DCBX configuration example

Network requirements


As shown in [Figure 72](#), in a data center network:

- Interface Ten-GigabitEthernet 1/0/1 of the access switch (Switch A) connects to the FCoE adapter of the data center server (DC server).
- Configure Switch A to prioritize and implement lossless transmission for FCoE and FIP packets to DC server.

NOTE:

Suppose that both Switch A and DC server support DCBX Rev 1.01.

Figure 72 Network diagram for DCBX configuration



Configuration procedure

1. Enable LLDP and DCBX TLV advertising

Enable LLDP globally.

```
<SwitchA> system-view
[SwitchA] lldp enable
```

Enable LLDP and DCBX TLV advertising on interface Ten-GigabitEthernet 1/0/1.

```
[SwitchA] interface ten-gigabitethernet 1/0/1
[SwitchA-Ten-GigabitEthernet1/0/1] lldp enable
[SwitchA-Ten-GigabitEthernet1/0/1] lldp tlv-enable dot1-tlv dcbx
[SwitchA-Ten-GigabitEthernet1/0/1] quit
```

2. Configure APP parameters

Create Ethernet frame header ACL 4000 and configure the ACL to permit FCoE packets (whose protocol number is 0x8906) and FIP packets (whose protocol number is 0x8914) to pass through.

```
[SwitchA] acl number 4000
[SwitchA-acl-ethernetframe-4000] rule permit type 8906 ffff
[SwitchA-acl-ethernetframe-4000] rule permit type 8914 ffff
[SwitchA-acl-ethernetframe-4000] quit
```

Create a class named **app_c**, specify the operator of the class as OR, and use ACL 4000 as the match criterion of the class.

```
[SwitchA] traffic classifier app_c operator or
[SwitchA-classifier-app_c] if-match acl 4000
[SwitchA-classifier-app_c] quit
```

Create a traffic behavior named **app_b** and configure the traffic behavior to mark packets with 802.1p priority value 3.

```
[SwitchA] traffic behavior app_b
[SwitchA-behavior-app_b] remark dot1p 3
[SwitchA-behavior-app_b] quit
```

Create a QoS policy named **plcy**, associate class **app_c** with traffic behavior **app_b** in the QoS policy, and apply the association to DCBX.

```
[SwitchA] qos policy plcy
[SwitchA-qospolicy-plcy] classifier app_c behavior app_b mode dcbx
[SwitchA-qospolicy-plcy] quit
```

Apply the QoS policy **plcy** to the outgoing traffic of interface Ten-GigabitEthernet 1/0/1.

```
[SwitchA] interface ten-gigabitethernet 1/0/1
[SwitchA-Ten-GigabitEthernet1/0/1] qos apply policy plcy outbound
[SwitchA-Ten-GigabitEthernet1/0/1] quit
```

3. Configure ETS parameters

Map 802.1p priority value 3 to local precedence 3. (This is the default mapping table. Modify this configuration as needed.)

```
[SwitchA] qos map-table dot1p-lp
[SwitchA-maptbl-dot1p-lp] import 3 export 3
[SwitchA-maptbl-dot1p-lp] quit
```

Enable byte-count WRR queuing on interface Ten-GigabitEthernet 1/0/1, assign queue 3 to WRR priority group 1, and configure the scheduling weight for queue 3 as 15.

```
[SwitchA] interface ten-gigabitethernet 1/0/1
[SwitchA-Ten-GigabitEthernet1/0/1] qos wrr byte-count
[SwitchA-Ten-GigabitEthernet1/0/1] qos wrr 3 group 1 byte-count 15
```

4. Configure PFC parameters

Configure the Ethernet interface to enable PFC by automatically negotiating with its peer, and enable PFC for 802.1p priority 3.

```
[SwitchA-Ten-GigabitEthernet1/0/1] priority-flow-control auto
[SwitchA-Ten-GigabitEthernet1/0/1] priority-flow-control no-drop dot1p 3
```

Configure interface Ten-GigabitEthernet 1/0/1 to use the 802.1p priority carried in packets for priority mapping.

```
[SwitchA-Ten-GigabitEthernet1/0/1] qos trust dot1p
```

5. Verify the configuration

Through the specific menu on the DC server, you can see the data exchange procedure between the DC server and Switch A to verify the configuration. Take a Qlogic adapter on the DC server for example. The data exchange procedure is:

Wed Aug 25 15:08:56 CST 2010

DCBX TLV (Type-Length-Value) Data

=====

DCBX Parameter Type and Length

DCBX Parameter Length: 13

DCBX Parameter Type: 2

DCBX Parameter Information

Parameter Type: Current

Pad Byte Present: Yes

DCBX Parameter Valid: Yes

Reserved: 0

DCBX Parameter Data

Priority Group ID of Priority 1: 0

Priority Group ID of Priority 0: 2

Priority Group ID of Priority 3: 3

Priority Group ID of Priority 2: 1

Priority Group ID of Priority 5: 5

Priority Group ID of Priority 4: 4

Priority Group ID of Priority 7: 7

Priority Group ID of Priority 6: 6

Priority Group 0 Percentage: 1

Priority Group 1 Percentage: 3

Priority Group 2 Percentage: 5

Priority Group 3 Percentage: 24

// This value is calculated by the Qlogic adapter according to the local precedence, and is different from the configured weight value.

Priority Group 4 Percentage: 8

Priority Group 5 Percentage: 14

Priority Group 6 Percentage: 21

Priority Group 7 Percentage: 24

Number of Traffic Classes Supported: 8

DCBX Parameter Information

Parameter Type: Remote

Pad Byte Present: Yes

DCBX Parameter Valid: Yes
Reserved: 0

DCBX Parameter Data

Priority Group ID of Priority 1: 0
Priority Group ID of Priority 0: 2

Priority Group ID of Priority 3: 3
Priority Group ID of Priority 2: 1

Priority Group ID of Priority 5: 5
Priority Group ID of Priority 4: 4

Priority Group ID of Priority 7: 7
Priority Group ID of Priority 6: 6

Priority Group 0 Percentage: 1
Priority Group 1 Percentage: 3
Priority Group 2 Percentage: 5

Priority Group 3 Percentage: 24

// This value is calculated by the FCoE adapter according to the local precedence, and is different from the configured weight value.

Priority Group 4 Percentage: 8
Priority Group 5 Percentage: 14
Priority Group 6 Percentage: 21
Priority Group 7 Percentage: 24

Number of Traffic Classes Supported: 8

The output shows that the DC server maps 802.1p priority value 3 to local precedence value 3, and changes the weight of the queue holding packets with 802.1p priority value 3 to be the same as that of queue 7.

DCBX Parameter Type and Length

DCBX Parameter Length: 2
DCBX Parameter Type: 3

DCBX Parameter Information

Parameter Type: Current
Pad Byte Present: No
DCBX Parameter Valid: Yes
Reserved: 0

DCBX Parameter Data

PFC Enabled on Priority 0: No
PFC Enabled on Priority 1: No
PFC Enabled on Priority 2: No
PFC Enabled on Priority 3: Yes

PFC Enabled on Priority 4: No
PFC Enabled on Priority 5: No
PFC Enabled on Priority 6: No
PFC Enabled on Priority 7: No

Number of Traffic Classes Supported: 6

DCBX Parameter Information

Parameter Type: Remote
Pad Byte Present: No
DCBX Parameter Valid: Yes
Reserved: 0

DCBX Parameter Data

PFC Enabled on Priority 0: No
PFC Enabled on Priority 1: No
PFC Enabled on Priority 2: No
PFC Enabled on Priority 3: Yes
PFC Enabled on Priority 4: No
PFC Enabled on Priority 5: No
PFC Enabled on Priority 6: No
PFC Enabled on Priority 7: No

Number of Traffic Classes Supported: 6

DCBX Parameter Information

Parameter Type: Local
Pad Byte Present: No
DCBX Parameter Valid: Yes
Reserved: 0

DCBX Parameter Data

PFC Enabled on Priority 0: No
PFC Enabled on Priority 1: No
PFC Enabled on Priority 2: No
PFC Enabled on Priority 3: Yes
PFC Enabled on Priority 4: No
PFC Enabled on Priority 5: No
PFC Enabled on Priority 6: No
PFC Enabled on Priority 7: No

Number of Traffic Classes Supported: 1

The output shows that DC server performs PFC for packets carrying 802.1p priority 3 after negotiating with Switch A.

Service loopback group configuration

The service loopback function is implemented through service loopback groups. A service loopback group must contain at least one Ethernet port as its member port, called a service loopback port. To increase service redirecting throughput, you can assign multiple service loopback ports to a service loopback group. Similar to the Ethernet link aggregation function, a service loopback group can increase bandwidth and implement load sharing. For more information about Ethernet link aggregation, see the chapter “Ethernet link aggregation configuration.”

For example, by assigning three ports of the same device to a service loopback group, you can create a logical link whose bandwidth can be as high as the total bandwidth of these three ports. In addition, service traffic is load-balanced among these ports.

Service types of service loopback groups

Service loopback groups provide the following service types:

- IPv6, supporting IPv6 unicast traffic
- IPv6 multicast, supporting IPv6 multicast traffic
- Tunnel, supporting unicast tunnel traffic
- Multicast tunnel, supporting multicast tunnel traffic
- MPLS, supporting MPLS traffic

The switch only supports the service loopback group types of Tunnel and Multicast tunnel.

Requirements on service loopback ports

Before assigning a port to a service loopback group, ensure the port meets the following requirements.

- The port supports the services type or types of the service loopback group.
- The port is configured with only QoS and ACL settings, or physical settings such as rate and duplex mode.
- The port is not configured with MSTP, NDP, LLDP, 802.1X, MAC address authentication, port security mode, or IP source guard, or as the member port of an isolation group.
- The link type of the port is access.
- The port is not a member of any Ethernet link aggregation group or service loopback group.

States of service loopback ports

States of service loopback ports

A member port in a service loopback group is a service loopback port, which can be in one of the following states:

- Selected: a selected port can loop back user traffic.
- Unselected: an unselected port cannot loop back user traffic.

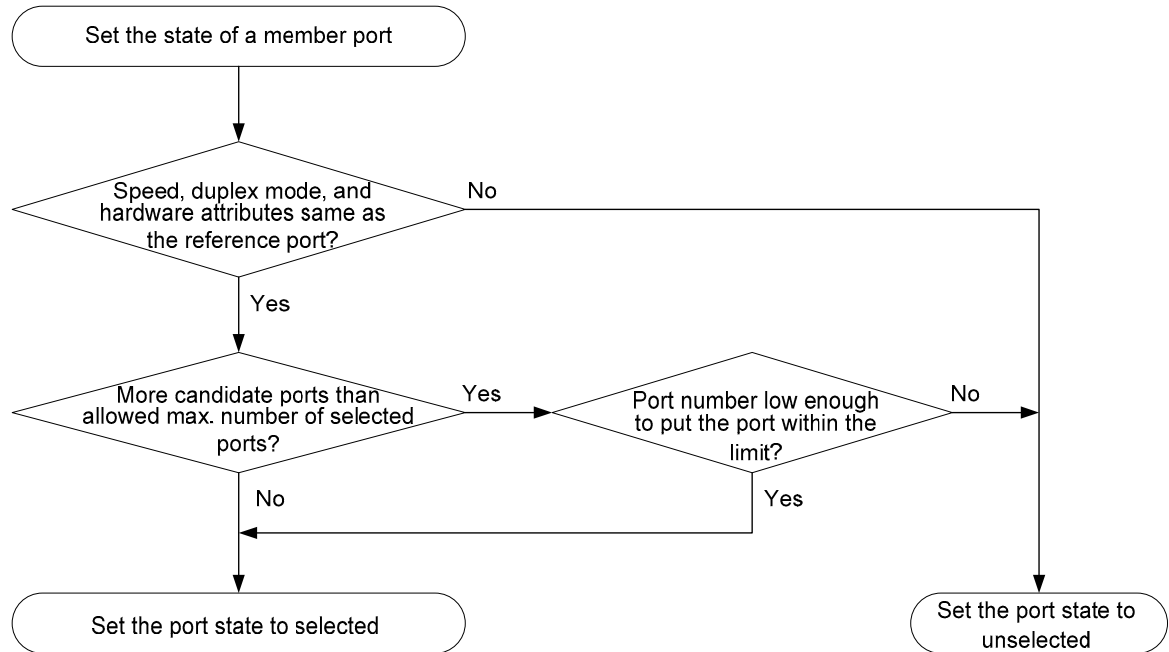
The number of selected ports is limited in a service loopback group.

Setting the state of service loopback ports

The system sets the state of each member port in a service loopback group to selected or unselected by using the following workflow.

1. Select the full-duplex port with the highest rate as the reference port. If two ports have the same duplex mode and speed, the one with the lower port number wins out.
2. Set the state of each member port in the service loopback group as shown in [Figure 73](#).

Figure 73 Set the state of each member port in a service loopback group



Every time a new port is assigned to a service loopback group, the system sets the state of the member ports in the service loopback group all over again according to the process above.

Configuring a service loopback group

To configure a service loopback group:

To do...	Use the command...	Remarks
1. Enter system view	system-view	—
2. Create a service loopback group and specify its service type	service-loopback group <i>number</i> type { multicast-tunnel tunnel } *	Required.
3. Enter Layer 2 Ethernet interface view	interface <i>interface-type interface-number</i>	—
4. Assign the Ethernet interface to the specified service loopback group	port service-loopback group <i>number</i>	Required. By default, a port does not belong to any service loopback group.

A service loopback group can process service only after it is referenced by other features.

A service loopback group can be referenced by multiple features at the same time, for example, by multiple tunnel interfaces. Assign more ports to a service loopback group that is likely to be referenced by multiple features at the same time, so that the service loopback group can provide sufficient bandwidth.

Change the service type of an existing service loopback group. For the change to be successful, you must ensure that the service group has not been referenced, The attributes of all member ports, if any, are not conflicting with the target service type, and no service loopback group has been created for the target service type because only one service loopback group is allowed for a service type.

You cannot remove a service loopback group that is referenced by other features.

Displaying and maintaining service loopback groups

To do...	Use the command...	Remarks
Display information about the specified service loopback group or all service loopback groups	display service-loopback group [<i>number</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Service loopback group configuration example

Network requirements

Ports of Switch A support the tunnel service. Assign GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 to a service loopback group to increase bandwidth and achieve load sharing.

Configuration procedure

Create service loopback group 1, and specify the service type as Tunnel (unicast tunnel service).

```
<SwitchA> system-view
```

```
[SwitchA] service-loopback group 1 type tunnel
```

Disable MSTP, NDP and LLDP on GigabitEthernet 1/0/1 through GigabitEthernet 1/0/3 and then assign them to service loopback group 1.

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

```
[SwitchA-GigabitEthernet1/0/1] undo stp enable
```

```
[SwitchA-GigabitEthernet1/0/1] undo ndp enable
```

```
[SwitchA-GigabitEthernet1/0/1] undo lldp enable
```

```
[SwitchA-GigabitEthernet1/0/1] port service-loopback group 1
```

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

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

```
[SwitchA-GigabitEthernet1/0/2] undo stp enable
```

```
[SwitchA-GigabitEthernet1/0/2] undo ndp enable
```

```
[SwitchA-GigabitEthernet1/0/2] undo lldp enable
```

```
[SwitchA-GigabitEthernet1/0/2] port service-loopback group 1
```

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

```
[SwitchA] interface gigabitethernet 1/0/3
```

```
[SwitchA-GigabitEthernet1/0/3] undo stp enable
```

```
[SwitchA-GigabitEthernet1/0/3] undo ndp enable
```

```
[SwitchA-GigabitEthernet1/0/3] undo lldp enable
```

```
[SwitchA-GigabitEthernet1/0/3] port service-loopback group 1
```

```
[SwitchA-GigabitEthernet1/0/3] quit
```

Create logical interface Tunnel 1 and reference service loopback group 1 on Tunnel 1.

```
[SwitchA] interface tunnel 1
```

```
[SwitchA-Tunnel1] service-loopback-group 1
```

Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/wwalerts>

After registering, you will receive email notification of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

Documents

To find related documents, browse to the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>

- For related documentation, navigate to the Networking section, and select a networking category.
- For a complete list of acronyms and their definitions, see *HP A-Series Acronyms*.

Websites

- HP.com <http://www.hp.com>
- HP Networking <http://www.hp.com/go/networking>
- HP manuals <http://www.hp.com/support/manuals>
- HP download drivers and software <http://www.hp.com/support/downloads>
- HP software depot <http://www.software.hp.com>

Conventions

This section describes the conventions used in this documentation set.





Command conventions

Convention	Description
Boldface	Bold text represents commands and keywords that you enter literally as shown.
<i>Italic</i>	<i>Italic</i> text represents arguments that you replace with actual values.
[]	Square brackets enclose syntax choices (keywords or arguments) that are optional.
{ x y ... }	Braces enclose a set of required syntax choices separated by vertical bars, from which you select one.
[x y ...]	Square brackets enclose a set of optional syntax choices separated by vertical bars, from which you select one or none.
{ x y ... } *	Asterisk-marked braces enclose a set of required syntax choices separated by vertical bars, from which you select at least one.
[x y ...] *	Asterisk-marked square brackets enclose optional syntax choices separated by vertical bars, from which you select one choice, multiple choices, or none.
&<1-n>	The argument or keyword and argument combination before the ampersand (&) sign can be entered 1 to n times.
#	A line that starts with a pound (#) sign is comments.

GUI conventions

Convention	Description
Boldface	Window names, button names, field names, and menu items are in bold text. For example, the New User window appears; click OK .
>	Multi-level menus are separated by angle brackets. For example, File > Create > Folder .

Symbols

Convention	Description
 WARNING	An alert that calls attention to important information that if not understood or followed can result in personal injury.
 CAUTION	An alert that calls attention to important information that if not understood or followed can result in data loss, data corruption, or damage to hardware or software.
 IMPORTANT	An alert that calls attention to essential information.
NOTE	An alert that contains additional or supplementary information.
 TIP	An alert that provides helpful information.

Network topology icons



Represents a generic network device, such as a router, switch, or firewall.



Represents a routing-capable device, such as a router or Layer 3 switch.



Represents a generic switch, such as a Layer 2 or Layer 3 switch, or a router that supports Layer 2 forwarding and other Layer 2 features.

Port numbering in examples

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

Index

- BPDU tunneling
 - configuration, 99
 - configuring, 101
 - enabling, 101
 - examples, 103
 - implementation, 100
 - prerequisites, 101
 - PVST configuration example, 104
 - STP configuration example, 103
- contacting HP, 236
- documentation
 - conventions used, 237
 - website, 236
- Ethernet interface
 - a Layer 2, 10
 - basic settings, 2
 - configuration, 1
 - connection mode, 19
 - displaying and maintaining, 19
 - enabling bridging, 18
 - enabling link-down port auto power-down, 11
 - enabling multi-port loopback detection, 16
 - enabling single-port loopback detection, 15
 - flow control, 4
 - General configuration, 2
 - jumbo frame support, 9
 - Layer 2 task list, 10
 - Layer 3, 19
 - link change suppression, 7
 - link-up port auto power-down, 9
 - loopback testing, 8
 - management, 1
 - naming conventions, 1
 - operating mode, 3
 - PFC configuration example, 21
 - port group, 10
 - setting speed options for auto negotiation, 12
 - setting the MDI mode, 17
 - setting the statistics polling interval, 9
 - switchable operating mode of Ethernet interfaces, 2
 - testing the cable connection, 18
 - traffic storm protection, 13
- Ethernet link aggregation
 - aggregate interface, 45
 - aggregating links in dynamic mode, 40
 - aggregating links in static mode, 39
 - aggregation group, 43
 - basic concepts, 36
 - configuration, 36
 - description of an aggregate interface, 45
 - displaying and maintaining, 50
 - dynamic aggregation group, 44
 - examples, 50
 - Layer 2 dynamic aggregation example, 53
 - Layer 2 static aggregation example, 51
 - link state traps for an aggregate interface, 45
 - link-aggregation traffic redirection, 49
 - load sharing criteria for link aggregation groups, 42
 - load sharing for link aggregation groups, 47
 - local-first load sharing for link aggregation, 48
 - minimum number of selected ports for an aggregation group, 46
 - restoring the default settings for an aggregate interface, 47
 - shutting down an aggregate interface, 46
 - static aggregation group, 43
 - task list, 42
- GVRP, 156
 - configuration, 153
 - displaying and maintaining, 159
 - examples, 160
 - fixed registration mode example, 161
 - forbidden registration mode example, 162
 - functions, 157
 - GARP, 153
 - GARP timers, 158
 - normal registration mode example, 160
 - protocols and standards, 156
 - task list, 157
- HP
 - customer support and resources, 236
 - document conventions, 237
 - documents and manuals, 236
 - icons used, 237
 - subscription service, 236
 - support contact information, 236

- symbols used, 237
- websites, 236
- icons, 237
- isolate-user-VLAN
 - associating secondary VLANs with, 138
 - configuration, 135
 - configuring, 136
 - displaying and maintaining, 139
 - example, 139
 - secondary VLANs, 137
- LLDP
 - advertisable TLVs, 212
 - APP parameters, 217
 - basic concepts, 204
 - basic LLDP example, 223
 - CDP compatibility, 215
 - CDP-compatible example, 225
 - configuration, 204
 - DCBX, 216
 - DCBX example, 227
 - DCBX task list, 217
 - displaying and maintaining, 222
 - enabling DCBX TLV advertising, 217
 - enabling LLDP advertising, 217
 - enabling polling, 211
 - ETS parameters, 219
 - examples, 223
 - management address and its encoding format, 212
 - operations, 208
 - performing basic LLDP configuration, 210
 - PFC parameters, 220
 - prerequisites, 215
 - protocols and standards, 209
 - setting an encapsulation format for LLDPDUs, 214
 - setting other parameters, 213
 - setting the operating mode, 210
 - setting the re-initialization delay, 211
 - task list, 209
 - trapping, 221
- loopback interface, 23
 - configuration, 23
 - configuring, 24
 - displaying and maintaining, 25
- MAC address table
 - aging timer for dynamic MAC address entries, 29
 - configuration, 26
 - configuring, 27
 - disabling MAC address learning on a VLAN, 28
 - displaying and maintaining, 31
 - example, 32
 - frame forwarding, 27
 - how to create an entry, 26
 - MAC address roaming, 30
 - MAC learning limit on ports, 29
 - manually configuring entries, 28
 - types of entries, 26
- MAC Information
 - configuration, 33
 - configuring, 33
 - enabling globally, 33
 - enabling on an interface, 33
 - example, 35
 - interval for sending Syslog messages, 34
 - interval for sending trap messages, 34
 - mode, 34
 - operations, 33
 - queue length, 34
- manuals, 236
- MSTP
 - basic concepts, 67
 - basic concepts in STP, 59
 - configuration, 58
 - configuring, 73
 - digest snooping, 85
 - displaying and maintaining, 94
 - enabling the MSTP feature, 83
 - example, 94
 - how STP works, 60
 - implementation on devices, 71
 - introduction, 66
 - introduction to RSTP, 66
 - link type of ports, 82
 - maximum hops of an MST region, 75
 - maximum port rate, 78
 - mode a port uses to recognize/send MSTP packets, 82
 - MST region, 73
 - network diameter of a switched network, 76
 - no agreement check, 87
 - operations, 70
 - output of port state transition information, 83
 - path costs of ports, 79
 - performing mCheck, 84
 - port priority, 81
 - ports as edge ports, 78
 - priority of a device, 75
 - protection functions, 91
 - protocol packets of STP, 58
 - protocols and standards, 71
 - root bridge, 73
 - secondary root bridge, 73

- task list, 71
- TC snooping, 89
- timeout factor, 77
- timers, 76
- why MSTP?, 66
- why STP?, 58
- work mode of an MSTP device, 74

null interface, 24

- configuration, 23
- displaying and maintaining, 25
- null 0 interface, 25

port isolation

- configuration, 55
- displaying and maintaining isolation groups, 56
- example, 56
- isolation group, 55

QinQ

- background and benefits, 165
- basic, 169
- basic QinQ example, 175
- configuration, 165
- examples, 175
- frame structure, 166
- implementations, 167
- inner VLAN ID substitution, 174
- inner-outer VLAN 802.1p priority mapping, 172
- modifying the TPID in a VLAN tag, 167
- operations, 165
- outer VLAN tagging policy, 170
- protocols and standards, 168
- selective, 170
- selective example, 178
- task list, 168
- TPID value in the CVLAN tag, 175
- TPID value in the SVLAN tag, 175
- TPID value in VLAN tags, 175
- VLAN transparent transmission, 169

service loopback group

- configuration, 232
- configuring, 234
- displaying and maintaining, 234
- example, 235
- requirements on service loopback ports, 232
- service types, 232
- states of service loopback ports, 232

subscription service, 236

super VLAN

- configuration, 130
- configuring, 130
- displaying and maintaining, 131
- example, 132

support and other resources, 236

symbols, 237

VLAN

- assigning a hybrid port, 114
- assigning a trunk port, 113
- assigning an access port, 111
- basic settings, 108
- basic settings of a VLAN interface, 108
- configuration, 106
- displaying and maintaining, 129
- fundamentals, 106
- introduction to MAC-based, 116
- introduction to port-based, 110
- introduction to protocol-based, 123
- IP subnet-based, 127, 128
- MAC-based, 116, 118
- MAC-based example, 120
- port-based, 110
- port-based example, 115
- protocol-based, 123, 124
- protocol-based example, 125
- types, 107

VLAN mapping

- many-to-one example, 195

VLAN mapping

- application scenario of one-to-one and many-to-one VLAN mapping, 181
- application scenario of two-to-two VLAN mapping, 183
- concepts and terms, 184
- configuration, 181
- configuring, 186
- examples, 195
- implementations, 184
- many-to-one, 188
- one-to-one, 186
- one-to-one example, 195
- two-to-two, 191
- two-to-two example, 201

voice VLAN

- assignment modes, 142
- automatic voice VLAN mode example, 149
- configuration, 142
- configuring, 146
- displaying and maintaining, 148
- examples, 149
- manual voice VLAN assignment mode example, 151
- normal mode, 145
- OUI addresses, 142

port to operate in automatic assignment mode,
147
port to operate in manual assignment mode, 147
prerequisites, 146

QoS priority settings for voice traffic on an
interface, 146
security mode, 145
websites, 236