

HP 5120 EI Switch Series

Layer 2—LAN Switching

Configuration Guide

Part number: 5998-1791

Software version: Release 2220

Document version: 6W100-20130810



Legal and notice information

© Copyright 2013 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

Configuring Ethernet interfaces	1
Ethernet interface naming conventions	1
Configuring a combo interface	1
Configuration prerequisites	1
Changing the active port of a combo interface	1
Configuring basic settings of an Ethernet interface	2
Shutting down an Ethernet interface	3
Setting speed options for auto negotiation on an Ethernet interface	3
Configuring flow control on an Ethernet interface	4
Configuring link change suppression on an Ethernet interface	5
Configuring link-down event suppression	6
Configuring link-up event suppression	6
Configuring loopback testing on an Ethernet interface	6
Configuration restrictions and guidelines	7
Configuration procedure	7
Configuring jumbo frame support	7
Configuring a port group	8
Enabling energy saving functions on an Ethernet interface	8
Enabling auto power-down	8
Configuring storm suppression	9
Setting the statistics polling interval	10
Enabling loopback detection on an Ethernet interface	10
Configuration restrictions and guidelines	11
Configuration procedure	12
Setting the MDI mode of an Ethernet interface	12
Enabling bridging on an Ethernet interface	13
Testing the cable connection of an Ethernet interface	14
Configuring storm control on an Ethernet interface	14
Configuration restrictions and guidelines	15
Configuration procedure	15
Displaying and maintaining an Ethernet interface	16
Configuring loopback and null interfaces	17
Configuring a loopback interface	17
Introduction to the loopback interface	17
Configuration procedure	17
Configuring a null interface	18
Introduction to the null interface	18
Configuration procedure	18
Displaying and maintaining loopback and null interfaces	18
Bulk configuring interfaces	20
Configuration guidelines	20
Configuration procedure	20
Configuring the MAC address table	21
Overview	21
How a MAC address table entry is created	21
Types of MAC address table entries	22
MAC address table-based frame forwarding	22

Configuring static, dynamic, and blackhole MAC address table entries	22
Configuring a static or dynamic MAC address table entry in system view	23
Configuring a static or dynamic MAC address table entry in interface view	23
Configuring a blackhole MAC address entry	23
Disabling MAC address learning	23
Disabling global MAC address learning	24
Disabling MAC address learning on ports	24
Configuring the aging timer for dynamic MAC address entries	24
Disabling MAC entry aging timer refresh based on destination MAC address	25
Application example	25
Configuring the MAC learning limit on ports	26
Enabling MAC address roaming	27
Enabling MAC address migration log notifying	28
Displaying and maintaining MAC address tables	29
MAC address table configuration example	29
Network requirements	29
Configuration procedure	30
Configuring MAC Information	31
Overview	31
Introduction to MAC Information	31
How MAC Information works	31
Enabling MAC Information globally	31
Enabling MAC Information on an interface	31
Configuring MAC Information mode	32
Configuring the interval for sending Syslog or trap messages	32
Configuring the MAC Information queue length	32
MAC Information configuration example	33
Network requirements	33
Configuration procedure	33
Configuring Ethernet link aggregation	34
Overview	34
Basic concepts	34
Aggregating links in static mode	37
Aggregating links in dynamic mode	38
Load-sharing criteria for link aggregation groups	40
Configuration restrictions and guidelines	40
Ethernet link aggregation configuration task list	40
Configuring an aggregation group	41
Configuration guidelines	41
Configuring a static aggregation group	41
Configuring a dynamic aggregation group	42
Configuring an aggregate interface	43
Configuring the description of an aggregate interface	43
Enabling link state traps for an aggregate interface	43
Limiting the number of Selected ports for an aggregation group	44
Shutting down an aggregate interface	45
Restoring the default settings for an aggregate interface	45
Configuring load sharing for link aggregation groups	46
Configuring load-sharing criteria for link aggregation groups	46
Enabling local-first load sharing for link aggregation	47
Enabling link-aggregation traffic redirection	48
Displaying and maintaining Ethernet link aggregation	49
Ethernet link aggregation configuration examples	50

Layer 2 static aggregation configuration example	50
Layer 2 dynamic aggregation configuration example	52
Configuring port isolation	55
Assigning a port to the isolation group	55
Displaying and maintaining the isolation group	55
Port isolation configuration example	56
Configuring spanning tree protocols	57
STP	57
STP protocol packets	57
Basic concepts in STP	58
Calculation process of the STP algorithm	59
RSTP	64
PVST	64
MSTP	64
STP, RSTP, and PVST limitations	64
MSTP features	64
MSTP basic concepts	65
How MSTP works	69
Implementation of MSTP on devices	69
Protocols and standards	69
Spanning tree configuration task list	70
Configuration restrictions and guidelines	70
STP configuration task list	70
RSTP configuration task list	71
PVST configuration task list	72
MSTP configuration task list	73
Setting the spanning tree mode	74
Configuring an MST region	75
Configuration restrictions and guidelines	75
Configuration procedure	75
Configuring the root bridge or a secondary root bridge	76
Configuration restrictions and guidelines	76
Configuring the current device as the root bridge of a specific spanning tree	76
Configuring the current device as a secondary root bridge of a specific spanning tree	77
Configuring the device priority	77
Configuring the maximum hops of an MST region	77
Configuring the network diameter of a switched network	78
Configuring spanning tree timers	78
Configuration restrictions and guidelines	79
Configuration procedure	79
Configuring the timeout factor	80
Configuring the maximum port rate	80
Configuring edge ports	81
Configuration restrictions and guidelines	81
Configuration procedure	81
Configuring path costs of ports	81
Specifying a standard for the device to use when it calculates the default path cost	82
Configuring path costs of ports	83
Configuration example	84
Configuring the port priority	84
Configuring the port link type	85
Configuration restrictions and guidelines	85
Configuration procedure	85

Configuring the mode a port uses to recognize/send MSTP packets	85
Enabling outputting port state transition information	86
Enabling the spanning tree feature	87
Configuration restrictions and guidelines	87
Enabling the spanning tree feature (in STP/RSTP/MSTP mode)	87
Enabling the spanning tree feature (in PVST mode)	87
Performing mCheck	88
Performing mCheck globally	88
Performing mCheck in interface view	88
Configuring Digest Snooping	89
Configuration restrictions and guidelines	89
Configuration procedure	89
Digest Snooping configuration example	90
Configuring No Agreement Check	91
Configuration prerequisites	92
Configuration procedure	92
No Agreement Check configuration example	93
Configuring TC snooping	93
Configuration restrictions and guidelines	94
Configuration procedure	94
Configuring protection functions	94
Configuration prerequisites	94
Enabling BPDU guard	94
Enabling root guard	95
Enabling loop guard	96
Enabling TC-BPDU guard	97
Enabling BPDU drop	97
Displaying and maintaining the spanning tree	98
Spanning tree configuration examples	98
MSTP configuration example	98
PVST configuration example	102
Configuring BPDU tunneling	106
Overview	106
Background	106
BPDU tunneling implementation	107
Enabling BPDU tunneling	108
Configuration prerequisites	108
Configuration restrictions and guidelines	109
Configuration procedure	109
Configuring destination multicast MAC address for BPDUs	109
BPDU tunneling configuration examples	110
BPDU tunneling for STP configuration example	110
BPDU tunneling for PVST configuration example	111
Configuring VLANs	113
Overview	113
VLAN fundamentals	113
VLAN types	114
Protocols and standards	115
Configuring basic VLAN settings	115
Configuration restrictions and guidelines	115
Configuration procedure	115
Configuring basic settings of a VLAN interface	116
Configuration procedure	116

VLAN interface configuration example.....	116
Configuring port-based VLANs.....	118
Introduction to port-based VLAN	118
Assigning an access port to a VLAN	119
Assigning a trunk port to a VLAN.....	120
Assigning a hybrid port to a VLAN.....	121
Port-based VLAN configuration example.....	122
Configuring MAC-based VLANs	124
Introduction to MAC-based VLAN	124
Configuration restrictions and guidelines	126
Configuration procedure	126
MAC-based VLAN configuration example	129
Configuring protocol-based VLANs	131
Introduction to protocol-based VLAN	131
Configuration restrictions and guidelines	131
Configuration procedure	132
Protocol-based VLAN configuration example.....	133
Configuring IP subnet-based VLANs	135
Configuration procedure	135
IP subnet-based VLAN configuration example	136
Displaying and maintaining VLAN	138
Configuring an isolate-user-VLAN.....	140
Overview.....	140
Configuration restrictions and guidelines	141
Configuration procedure	141
Displaying and maintaining isolate-user-VLAN.....	142
Isolate-user-VLAN configuration example.....	143
Network requirements.....	143
Configuration procedure	143
Verifying the configuration	144
Configuring a voice VLAN	146
Overview.....	146
Methods of identifying IP phones	146
OUI addresses	146
Automatically identifying IP phones through LLDP	147
Configuring a device to advertise voice VLAN information to IP phones.....	147
How a device advertises voice VLAN information	147
How a device obtains voice VLAN information.....	147
IP phone access methods	148
Connecting the host and the IP phone in series	148
Connecting the IP phone to the device separately	148
Configuring a voice VLAN on a port.....	148
Voice VLAN assignment modes.....	148
Security mode and normal mode of voice VLANs.....	150
Configuration prerequisites	151
Configuring QoS priority settings for voice traffic on an interface.....	151
Configuring a port to operate in automatic voice VLAN assignment mode	152
Configuring a port to operate in manual voice VLAN assignment mode.....	153
Enabling LLDP to automatically discover IP phones.....	154
Configuration prerequisites	154
Configuration procedure	155
Configuring LLDP to advertise a specific voice VLAN.....	155
Configuration guidelines	155

Configuration procedure	156
Dynamically advertising server-assigned VLANs through LLDP	157
Overview	157
Example for using 802.1X to authenticate IP phones	157
Displaying and maintaining voice VLAN	157
Voice VLAN configuration examples	158
Automatic voice VLAN mode configuration example	158
Manual voice VLAN assignment mode configuration example	160
Configuring GVRP	162
Overview	162
GARP	162
GVRP	165
Protocols and standards	165
GVRP configuration task list	165
Configuring GVRP functions	166
Configuration restrictions and guidelines	166
Configuration procedure	166
Configuring the GARP timers	167
Displaying and maintaining GVRP	168
GVRP configuration examples	168
GVRP normal registration mode configuration example	168
GVRP fixed registration mode configuration example	170
GVRP forbidden registration mode configuration example	171
Configuring QinQ	174
Overview	174
Background and benefits	174
How QinQ works	174
QinQ frame structure	175
Implementations of QinQ	176
Modifying the TPID in a VLAN tag	176
Protocols and standards	177
QinQ configuration task list	177
Configuring basic QinQ	178
Enabling basic QinQ	178
Configuring VLAN transparent transmission	178
Configuring selective QinQ	179
Configuring an outer VLAN tagging policy	179
Configuring an inner-outer VLAN 802.1p priority mapping	180
Configuring the TPID value in VLAN tags	181
QinQ configuration examples	181
Basic QinQ configuration example	181
Selective QinQ configuration example	184
Configuring LLDP	187
Overview	187
Background	187
Basic concepts	187
How LLDP works	191
Protocols and standards	192
LLDP configuration task list	192
Performing basic LLDP configuration	192
Enabling LLDP	192
Setting the LLDP operating mode	193
Setting the LLDP re-initialization delay	193

Enabling LLDP polling.....	193
Configuring the advertisable TLVs.....	194
Configuring the management address and its encoding format.....	194
Setting other LLDP parameters.....	195
Setting an encapsulation format for LLDPDUs	196
Configuring CDP compatibility.....	196
Configuration prerequisites.....	197
Configuration procedure	197
Configuring LLDP trapping	198
Displaying and maintaining LLDP.....	199
LLDP configuration examples	199
Basic LLDP configuration example.....	199
CDP-compatible LLDP configuration example.....	202
Configuring MVRP	204
Overview.....	204
Introduction to MRP	204
MVRP registration modes	206
Protocols and standards	207
MVRP configuration task list.....	207
Configuration prerequisites.....	207
Enabling MVRP.....	207
Configuration restrictions and guidelines	207
Configuration procedure	208
Configuring the MVRP registration mode.....	208
Configuring MRP timers.....	209
Enabling GVRP compatibility.....	210
Configuration restrictions and guidelines	210
Configuration procedure	210
Displaying and maintaining MVRP	210
Configuration example for MVRP in normal registration mode.....	211
Network requirements.....	211
Configuration procedure	212
Support and other resources	221
Contacting HP	221
Subscription service	221
Related information.....	221
Documents.....	221
Websites.....	221
Conventions	222
Index	224

Configuring Ethernet interfaces

Ethernet interface naming conventions

The GE and 10-GE interfaces on the HP 5120 EI switches 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 fabric. If the switch is not assigned to any IRF fabric, A uses 1.
- **B**—Represents a slot number on the switch. It uses 0 for fixed interfaces, 1 for interfaces on interface expansion card 1, and 2 for interfaces on interface expansion card 2.
- **C**—Represents the number of an interface on a slot.

Configuring a combo interface

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

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

Configuration prerequisites

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

- Use the **display port combo** command to identify the combo interfaces on your switch and identify the two physical ports that compose each combo interface.
- Use the **display interface** command to determine, of the two physical ports that compose a combo interface, which is the fiber combo port and which is the copper combo port. If the current port is the copper port, the output will include "Media type is twisted pair, Port hardware type is 1000_BASE_T". If the current port is the fiber port, the output will not include the information mentioned above.

Changing the active port of a combo interface

To change the active port of a double combo interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A

Step	Command	Remarks
3. Activate the current interface.	undo shutdown	Optional. By default, of the two ports that compose a combo interface, the one with a smaller port ID is active.

Configuring basic settings of an Ethernet interface

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

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

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

To configure an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Set the interface description.	description <i>text</i>	Optional. By default, the description of an interface is in the format of <i>interface-name</i> Interface . For example, GigabitEthernet1/0/1 Interface .
4. Set the duplex mode of the interface.	duplex { auto full half }	Optional. By default, the duplex mode is auto for Ethernet interfaces. The half keyword is not applicable to Ethernet copper ports that are configured with a 1000-Mbps port speed and fiber ports.
5. Set the port speed.	speed { 10 100 1000 auto }	Optional. By default, an Ethernet interface automatically negotiates a speed with the peer. GE (SFP) fiber ports do not support the 10 or 100 keyword. 10-GE fiber ports do not support this command.
6. Restore the default settings for the interface.	default	Optional.

NOTE:

Make sure that the fiber port speed matches the speed requirement of the inserted transceiver module. For example, after you insert a 1000-Mbps transceiver module into a fiber port, configure the port speed with the **speed 1000** or **speed auto** command.

Shutting down an Ethernet interface

⚠ CAUTION:

Use this feature with caution. After you manually shut down an Ethernet interface, the Ethernet interface cannot forward packets even if it is physically connected.

You might need to shut down and then bring up an Ethernet interface to activate some configuration changes, for example, the speed or duplex mode changes.


To shut down an Ethernet interface or a group of Ethernet interfaces:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>	Use any command. To shut down an Ethernet interface, enter Ethernet interface.
	<ul style="list-style-type: none">Enter port group view: port-group manual <i>port-group-name</i>	To shut down all Ethernet interfaces in a port group, enter port group view.
3. Shut down the Ethernet interface or interfaces.	shutdown	By default, Ethernet interfaces are up.

Setting speed options for auto negotiation on an Ethernet interface

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

Figure 1 Speed auto negotiation application scenario



As shown in Figure 1, all ports on Switch A 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 port GigabitEthernet 1/0/4, the port providing access to the Internet for the servers.

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

To set speed options for auto negotiation on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type interface-number</i>	N/A
3. Set speed options for auto negotiation.	speed auto { 10 100 1000 } *	Optional.

NOTE:

- This function is available only for copper GE ports that support speed auto negotiation.
- The **speed** and **speed auto** commands supersede each other, and whichever is configured last takes effect.

Configuring flow control on an Ethernet interface

To avoid packet drops on a link, you can enable flow control at both ends of the link. When traffic congestion occurs at the receiving end, the receiving end sends a flow control (Pause) frame to ask the sending end to suspend sending packets.


Flow control is implemented by receiving and sending Pause frames on ports. Flow control can operate in one of the following modes on an interface:

- **TxRx mode** (configured by using the **flow-control** command)—The interface can both send and receive flow control frames.

- **Rx mode** (configured by using the **flow-control receive enable** command)—The interface can receive, but not send flow control frames.

As shown in Figure 2, when both Port A and Port B forward packets at the rate of 1000 Mbps, Port C is congested. To avoid packet loss, enable flow control on Port A and Port B.

Figure 2 Flow control on ports



Configure flow control in TxRx mode on Port B and flow control in Rx mode on Port A:

- 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 handle unidirectional traffic congestion on a link, configure the **flow-control receive enable** command at one end, and the **flow-control** command at the other. To enable both ends of the link to handle traffic congestion, configure the **flow-control** command at both ends.

To enable flow control on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Enable flow control.	<ul style="list-style-type: none"> • Enable TxRx flow control: flow-control • Enable Rx flow control: flow-control receive enable 	Use either command. By default, Rx flow control is disabled on an Ethernet interface.

Configuring link change suppression on an Ethernet interface

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

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

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

Link-down event 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 event 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 event suppression

To enable an Ethernet interface to suppress link-down events:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Set a link-down event suppression interval.	link-delay <i>delay-time</i>	Link-down event suppression is disabled by default.

Configuring link-up event suppression

To configure link-up event suppression on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Set a link-up event suppression interval.	link-delay <i>delay-time</i> mode up	Link-up event suppression is disabled by default.

NOTE:

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

Configuring loopback testing on an Ethernet interface

If an Ethernet interface does not work correctly, you can enable loopback testing on it to identify the problem. Loopback testing has the following types:

- **Internal loopback testing**—Tests all on-chip functions related to Ethernet interfaces.

- **External loopback testing**—Tests hardware of Ethernet interfaces. To perform external loopback testing on an Ethernet interface, connect a loopback plug to the Ethernet interface. The switch sends test packets out of the interface, which are expected to loop over the plug and back to the interface. If the interface fails to receive any test packet, the hardware of the interface is faulty.

An Ethernet interface in a loopback test does not forward data traffic.

Configuration restrictions and guidelines

- On an interface administratively shut down, you can perform neither internal nor external loopback testing.
- During loopback testing, the Ethernet interface operates in full duplex mode. When you disable loopback testing, the interface returns to its duplex setting.
- Loopback testing is a one-time operation, and is not recorded in the configuration file.

Configuration procedure

To enable loopback testing on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Enable loopback testing.	loopback { external internal }	Optional. Disabled by default.

Configuring jumbo frame support

An Ethernet interface might 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.

In interface configuration mode (Ethernet interface view or port group view), you can set the length of jumbo frames that are allowed to pass through the Ethernet interface.

- If you execute the command in Ethernet interface view, the configuration takes effect only on the interface.
- If you execute the command in port group view, the configuration takes effect on all ports in the port group.

To configure jumbo frame support in interface view or port group view:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A

Step	Command	Remarks
3. Configure jumbo frame support.	jumboframe enable [<i>value</i>]	By default, the switch allows jumbo frames within 9216 bytes to pass through Ethernet interfaces. If you set the <i>value</i> argument multiple times, the latest configuration takes effect.

Configuring a port group

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

You create port groups manually. All settings made for a port group apply to all the member ports of the group.

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

To configure a manual port group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a manual port group and enter manual port group view.	port-group manual <i>port-group-name</i>	N/A
3. Assign Ethernet interfaces to the manual port group.	Group-member <i>interface-list</i>	If you use the group-member <i>interface-type interface-start-number to interface-type interface-end-number</i> command to add multiple ports in batch to the specified port group, make sure that the <i>interface-end-number</i> argument must be greater than the <i>interface-start-number</i> argument.

Enabling energy saving functions on an Ethernet interface

Enabling auto power-down

With the auto power-down function, the system automatically stops supplying power to an interface and the interface enters the power save mode if the interface is in down state for a certain period of time (which depends on the chip specifications and is not configurable). When the interface goes up, the system supplies power to the interface and the interface enters its normal state.

To enable auto power-down:

Step	Command	Remarks
1. Enter system view	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none"> Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> Enter port group view: port-group manual <i>port-group-name</i> 	<p>Use either command.</p> <p>To enable auto power-down on an Ethernet interface, enter Ethernet interface view.</p> <p>To enable auto power-down on a group of Ethernet interfaces, enter port group view.</p>
3. Enable auto power-down	port auto-power-down	Disabled by default.

NOTE:

When you connect an interface enabled with auto power-down to a device, if the link cannot go up correctly, disable auto power-down on the interface and try again.

Configuring storm suppression

In interface or port group view, you can set the maximum size of broadcast, multicast or unknown unicast traffic allowed to pass through an interface or each interface in a port group. When the broadcast, multicast, or unknown unicast traffic on the interface exceeds this threshold, the system discards packets until the traffic drops below this threshold.

The storm suppression thresholds configured for an Ethernet interface might become invalid if you enable the storm control function for the interface. For information about the storm control function, see "[Configuring storm control on an Ethernet interface](#)."

To set storm suppression thresholds on one or multiple Ethernet interfaces:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none"> Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> Enter port group view: port-group manual <i>port-group-name</i> 	<p>Use either command.</p> <p>To configure storm suppression on an Ethernet interface, enter Ethernet interface view.</p> <p>To configure storm suppression on a group of Ethernet interfaces, enter port group view.</p>
3. Set the broadcast suppression threshold ratio.	broadcast-suppression { <i>ratio</i> pps <i>max-pps</i> kbps <i>max-kbps</i> }	Optional. By default, all broadcast traffic is allowed to pass through.
4. Set the multicast suppression threshold ratio.	multicast-suppression { <i>ratio</i> pps <i>max-pps</i> kbps <i>max-kbps</i> }	Optional. By default, all multicast traffic is allowed to pass through.
5. Set the unknown unicast suppression threshold ratio.	unicast-suppression { <i>ratio</i> pps <i>max-pps</i> kbps <i>max-kbps</i> }	Optional. By default, all unknown unicast traffic is allowed to pass through.

NOTE:

For an Ethernet interface that belongs to a port group, if you set a traffic suppression threshold for the interface in both Ethernet interface view and port group view, the threshold configured last takes effect.

Setting the statistics polling interval

To set the statistics polling interval globally or on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
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.

Enabling loopback detection on an Ethernet interface

If a switch receives a packet that it sent, a loop has occurred to the switch. Loops might cause broadcast storms, which degrade network performance. You can use this feature to detect whether a loop has occurred.

Depending on whether the receiving interface is the same as the sending interface, loops fall into the following types:

- **Single-port loop**—Occurs when an interface receives a packet that it sent out and the receiving interface is the same as the sending interface, as shown in [Figure 3](#).
- **Multi-port loop**—Occurs when a switch receives a packet that it sent out but the receiving interface is not the sending interface. For example, as shown in [Figure 4](#), when Port 2 receives the packet sent out of Port 1, a loop occurs between Port 1 and Port 2, and Port 2 is the looped port.

Figure 3 Single-port loop





Figure 4 Multi-port loop



You can enable loopback detection to detect loops on an interface and, if the interface supports the **loopback-detection action** command, configure the protective action to take on the receiving interface when a loop is detected, for example, to shut down the interface. Depending on whether a protective action is configured, the switch takes the actions in [Table 1](#) to alleviate the impact of the loop condition.

Table 1 Actions to take upon detection of a loop condition

Port type	Actions	
	No protective action is configured	A protective action is configured
Access port	<ul style="list-style-type: none">Place the receiving interface in controlled mode. The interface drops the incoming packets and correctly sends packets.Generate traps and log messages.Delete all MAC address entries of the interface.	<ul style="list-style-type: none">Perform the configured protective action.Generate traps and log messages.Delete all MAC address entries of the interface.
Hybrid or trunk port	<ul style="list-style-type: none">Generate traps and log messages.If loopback detection control is enabled, place the receiving interface in controlled mode. The interface does not receive or send packets.Delete all MAC address entries of the interface.	<ul style="list-style-type: none">Generate 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.

Configuration restrictions and guidelines

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

Configuration procedure

To configure loopback detection:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable global loopback detection.	loopback-detection enable	Disabled by default.
3. Enable multi-port loopback detection.	loopback-detection multi-port-mode enable	Optional. By default, multi-port loopback detection is disabled, and the switch can only detect single-port loopback.
4. Set the loopback detection interval.	loopback-detection interval-time <i>time</i>	Optional. 30 seconds by default.
5. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>Enter port group view: port-group manual <i>port-group-name</i>	Use either command. To configure loopback detection on one interface, enter Ethernet interface view. To configure loopback detection on a group of Ethernet interfaces, enter port group view.
6. Enable loopback detection on the interface.	loopback-detection enable	Disabled by default.
7. Enable loopback detection control on a trunk port or a hybrid port.	loopback-detection control enable	Optional. Disabled by default.
8. Enable loopback detection in all the VLANs on the trunk or hybrid port.	loopback-detection per-vlan enable	Optional. By default, a trunk or hybrid port performs loopback detection only in its port VLAN ID (PVID).
9. Set the protective action to take on the interface when a loop is detected.	loopback-detection action { no-learning semi-block shutdown }	Optional. By default, a looped interface drops the incoming packets and correctly sends packets; the system generates traps and log messages, and deletes all MAC address entries of the looped interface. With the shutdown keyword specified, the switch shuts down the looped ports and set their physical state to Loop down. When a looped port recovers, you must use the undo shutdown command to restore its forwarding capability.

Setting the MDI mode of an Ethernet interface



IMPORTANT:

Fiber ports do not support the MDI mode setting.

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

- Across mode
- Normal mode
- Auto mode

A copper Ethernet interface uses an RJ-45 connector, which comprises eight pins, each 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 MDI modes as follows:

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

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

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

To set the MDI mode of an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
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 device 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, the device discards the packet.

To enable the device to forward such packets rather than drop them, enable the bridging function on the Ethernet interface.

To enable bridging on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Enable bridging on the Ethernet interface.	port bridge enable	Disabled by default.

Testing the cable connection of an Ethernet interface

! IMPORTANT:

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

You can test the cable connection of an Ethernet interface for a short or open circuit. The switch 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:

Step	Command
1. Enter system view.	system-view
2. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>
3. Test the cable connected to the Ethernet interface.	virtual-cable-test

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 drops below the lower threshold from the upper threshold.

When a particular type of traffic exceeds its upper threshold, the interface does either of the following, depending on your configuration:

- Blocks this type of traffic, while forwarding other types of traffic. Even though the interface does not forward the blocked traffic, it still counts the traffic. When the blocked traffic drops below the lower threshold, the port begins to forward the traffic.

- Shuts down automatically. The interface shuts down automatically and stops forwarding any traffic. When the blocked traffic drops below the lower threshold, the port does not forward the traffic. To bring up the interface, use the **undo shutdown** command or disable the storm control function.

Alternatively, you can configure the storm suppression function to control a specific type of traffic. Do not enable them both on an Ethernet interface at the same time because the storm suppression and storm control functions are mutually exclusive. For example, with an unknown unicast suppression threshold set on an Ethernet interface, do not enable storm control for unknown unicast traffic on the interface. For more information about storm suppression, see "[Configuring storm suppression](#)."

Configuration restrictions and guidelines

- 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 a port at least one polling interval and at most two polling intervals to take a storm control action.
- The storm control function allows you to set the upper and lower thresholds for all three types of packets respectively on the same interface.

Configuration procedure

To configure storm control on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
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</i> <i>interface-number</i>	N/A
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 } max-pps-values min-pps-values	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.

Step	Command	Remarks
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 drops below the lower threshold from the upper threshold.

Displaying and maintaining an Ethernet interface

Task	Command	Remarks
Display Ethernet interface information.	display interface [<i>interface-type</i>] brief [down] [[{ begin exclude include } <i>regular-expression</i>]] display interface <i>interface-type interface-number</i> [brief] [[{ begin exclude include } <i>regular-expression</i>]]	Available in any view
Display traffic statistics for the specified interfaces.	display counters { inbound outbound } interface [<i>interface-type</i>] [[{ begin exclude include } <i>regular-expression</i>]]	Available in any view
Display traffic rate statistics over the last 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 the specified interfaces.	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 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 information about the loopback function.	display loopback-detection [[{ begin exclude include } <i>regular-expression</i>]]	Available in any view
Display information about storm control.	display storm-constrain [broadcast multicast unicast] [interface <i>interface-type interface-number</i>] [[{ begin exclude include } <i>regular-expression</i>]]	Available in any view
Clear the interface statistics.	reset counters interface [<i>interface-type</i> [<i>interface-number</i>]]	Available in user view
Clear the statistics of discarded packets on the specified interfaces.	reset packet-drop interface [<i>interface-type</i> [<i>interface-number</i>]]	Available in user view

Configuring loopback and null interfaces

Configuring a loopback interface

Introduction to the loopback interface

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

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

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

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

Configuration procedure

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a loopback interface and enter loopback interface view.	interface loopback <i>interface-number</i>	N/A
3. Set the interface description.	description <i>text</i>	Optional By default, the description of a loopback interface is <i>interface name</i> Interface.
4. Shut down the loopback interface.	shutdown	Optional By default, a loopback interface is up.
5. Restore the default settings for the loopback interface.	default	Optional

NOTE:

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

Configuring a null interface

Introduction to the null interface

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

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

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

Configuration procedure

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter null interface view.	interface null 0	The Null 0 interface is the default null interface on your switch. It cannot be manually created or removed.
3. Set the interface description.	description <i>text</i>	Optional By default, the description of a null interface is <i>interface name</i> Interface.
4. Restore the default settings for the null interface.	default	Optional

Displaying and maintaining loopback and null interfaces

Task	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

Task	Command	Remarks
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

Bulk configuring interfaces

You can enter interface range view to bulk configure multiple interfaces with the same feature instead of configuring them one by one. For example, you can perform the **shutdown** command in interface range view to shut down a range of interfaces.

Failure of applying a command on one member interface does not affect the application of the command on the other member interfaces. If applying a command on one member interface fails, the system displays an error message and continues with the next member interface.

Configuration guidelines

When you bulk configure interfaces in interface range view, follow these restrictions and guidelines:

- In interface range view, only the commands supported by the first interface are available.
- Do not assign an aggregate interface and any of its member interfaces to an interface range at the same time. Some commands, after being executed on both an aggregate interface and its member interfaces, can break up the aggregation.
- No limit is set on the maximum number of interfaces in an interface range. The more interfaces in an interface range, the longer the command execution time.

Configuration procedure

To bulk configure interfaces:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface range view.	Approach 1: interface range { <i>interface-type</i> <i>interface-number</i> [to <i>interface-type</i> <i>interface-number</i>] } &<1-5> Approach 2: interface range name <i>name</i> [interface { <i>interface-type</i> <i>interface-number</i> [to <i>interface-type</i> <i>interface-number</i>] } &<1-5>]	Use either approach. In approach 2, you assign a name to an interface range and can specify this name rather than the interface range to enter the interface range view.
3. Display commands available for the first interface in the interface range.	Enter ? at the interface range prompt.	Optional.
4. Perform available commands to configure the interfaces.	Available commands vary by interface.	N/A
5. Verify the configuration.	display this	Optional.

Configuring the MAC address table

This feature covers only the unicast MAC address table. For information about configuring static multicast MAC address table entries for IGMP snooping and MLD snooping, see *IP Multicast Configuration Guide*.

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

The MAC address table configuration tasks are all optional and can be performed in any order.

Overview

To reduce single-destination packet flooding in a switched LAN, an Ethernet device uses a MAC address table for forwarding frames through unicast instead of broadcast. An MAC address table entry contains a destination MAC address, the outgoing interface corresponding to the MAC address, and the ID of the VLAN to which the outgoing interface belongs. When forwarding a single-destination frame, the device first looks up the MAC address of the frame in the MAC address table for a match. If the switch finds an entry, it forwards the frame out of the outgoing port in the entry. If the switch does not find an entry, it floods the frame out of all but the incoming port.

To view MAC address table information, use the **display mac-address** command, as follows:

```
<Sysname> display mac-address
```

MAC ADDR	VLAN ID	STATE	PORT INDEX	AGING TIME(s)
000f-e201-0101	1	Learned	GigabitEthernet1/0/1	AGING

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

How a MAC address table entry is created

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

MAC address learning

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

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

1. Verifies the source MAC address (for example, MAC-SOURCE) of the frame.
2. Looks up the source MAC address in the MAC address table.
 - If an entry is found, the device updates the entry.
 - If no entry is found, the device adds an entry for MAC-SOURCE and Port A.
3. After obtaining this source MAC address, when the device receives a frame destined for MAC-SOURCE, the device finds the MAC-SOURCE entry in the MAC address table and forwards the frame out of Port A.

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

Manually configuring MAC address entries

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

To improve port security, you can bind specific user devices to the port by manually adding MAC address entries to the MAC address table of the switch.

Types of MAC address table entries

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

- **Static entries**—Static entries are manually added in order to forward frames with specific destination MAC addresses out of their associated ports and never age out.
- **Dynamic entries**—Dynamic entries can be manually added or dynamically learned in order to forward frames with specific destination MAC addresses out of their associated ports and might age out.
- **Blackhole entries**—Blackhole entries are manually configured and never age out. Blackhole entries are configured for filtering out frames with specific source or 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 MAC address entry.

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

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 obtained or created, an aging time starts. If the entry has not updated when the aging timer expires, the device deletes the entry. If the entry has updated before the aging timer expires, the aging timer restarts.

MAC address table-based frame forwarding

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

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

Configuring static, dynamic, and blackhole MAC address table entries

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

Configuring a static or dynamic MAC address table entry in system view

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Add or modify a dynamic or static MAC address entry.	mac-address { dynamic static } <i>mac-address interface interface-type interface-number vlan vlan-id</i>	By default, no MAC address entry is configured. Make sure that you have created the VLAN and assigned the interface to the VLAN.

Configuring a static or dynamic MAC address table entry in interface view

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view.	interface interface-type interface-number	N/A
3. Add or modify a static or dynamic MAC address entry.	mac-address { dynamic static } <i>mac-address vlan vlan-id</i>	By default, no MAC address entry is configured. Make sure that you have created the VLAN and assigned the interface to the VLAN.

Configuring a blackhole MAC address entry

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Add or modify a blackhole MAC address entry.	mac-address blackhole mac-address vlan vlan-id	By default, no MAC address entry is configured. Make sure that you have created the VLAN.

Disabling MAC address learning

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

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

Disabling global MAC address learning

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

To disable MAC address learning:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Disable global MAC address learning.	mac-address mac-learning disable	Enabled by default.

Disabling MAC address learning on ports

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

To disable MAC address learning on an interface or a port group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable global MAC address learning.	undo mac-address mac-learning disable	Optional Enabled by default.
3. Enter interface view or port group view.	<ul style="list-style-type: none">Enter Layer 2 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>	Use either command. Settings in Layer 2 Ethernet interface view or Layer 2 aggregate interface view take effect on the interface only. Settings in port group view take effect on all member ports in the port group. For more information about configuring a port group, see "Configuring Ethernet interfaces."
4. Disable MAC address learning on the interface or all ports in the port group.	mac-address mac-learning disable	Enabled by default.

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 device deletes that entry. This aging mechanism ensures that the MAC address table can quickly update to accommodate the latest network changes.

Set the aging timer appropriately. Too long an aging interval might cause the MAC address table to retain outdated entries, exhaust the MAC address table resources, and fail to update its entries to

accommodate the latest network changes. Too short an interval might result in removal of valid entries, causing unnecessary flooding, which might affect device performance.

To configure the aging timer for dynamic MAC address entries:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the aging timer for dynamic MAC address entries.	mac-address timer { aging seconds no-aging }	Optional 300 seconds by default. The no-aging keyword disables the aging timer.

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

Disabling MAC entry aging timer refresh based on destination MAC address

To accommodate network changes, the MAC address table keeps updating. Each dynamic MAC address entry has an aging timer. When the device receives a packet with the source or destination MAC address matching a dynamic MAC address entry, it restarts the aging timer for the entry.

If you want the device to restart the aging timer of dynamic entries for only matching source MAC addresses, disable MAC entry aging timer refresh based on destination MAC address.


To disable MAC entry aging timer refresh based on destination MAC address:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Disable MAC entry aging timer refresh based on destination MAC address.	mac-address destination-hit disable	By default, MAC entry aging timer refresh based on destination MAC address is enabled.

Application example

Microsoft Network Load Balancing (NLB) is a load balancing technology for server clustering developed on Windows Server.

Figure 5 NLB cluster



NLB supports load sharing and redundancy among servers within a cluster. To implement fast failover, NLB requires that the switch forwards network traffic to all servers or specified servers in the cluster, and each server filters out unexpected traffic.

In NLB unicast mode, when a server joins the cluster or a failover occurs, a packet with a virtual source MAC address is sent within the cluster. The switch then adds the virtual MAC address to its MAC address table, and packets destined for the server use the virtual MAC address (although not used by the server) as their destination address. If the virtual MAC address never ages out, the switch forwards packets only through the port associated with the virtual MAC address rather than all ports connected to the servers within the cluster.

To address this issue, disable MAC entry aging timer refresh based on destination MAC address to age out the virtual MAC address, so that the switch can forward packets to all servers within the cluster.

Configuring the MAC learning limit on ports

To prevent the MAC address table from getting too large, you can limit the number of MAC addresses that a port can learn.

To configure the MAC learning limit on a Layer 2 Ethernet interface or all ports in a port group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none">Enter Layer 2 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. Settings in Layer 2 Ethernet interface view take effect on the interface only. Settings in port group view take effect on all member ports in the port group.
3. Configure the MAC learning limit on the interface or port group.	mac-address max-mac-count <i>count</i>	No MAC learning limit is configured by default. Layer 2 aggregate interfaces do not support this command.

NOTE:

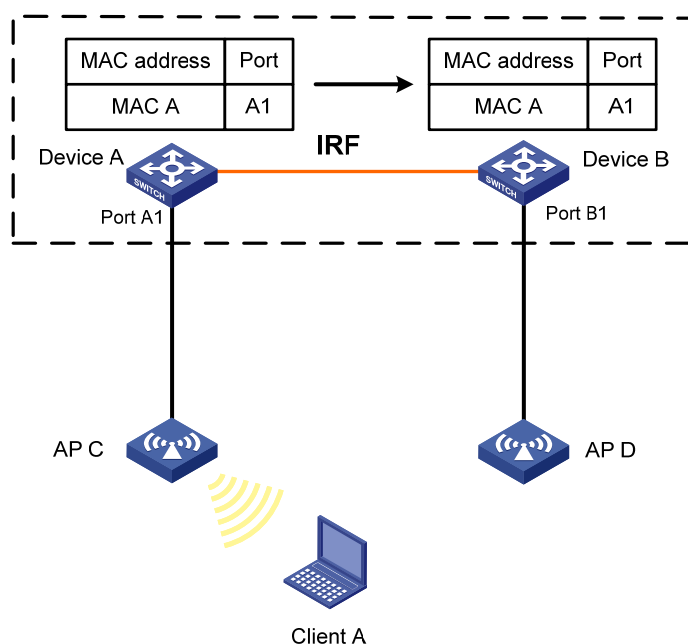
Do not configure the MAC learning limit on any member ports of an aggregation group. Otherwise, the member ports cannot be selected.

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.

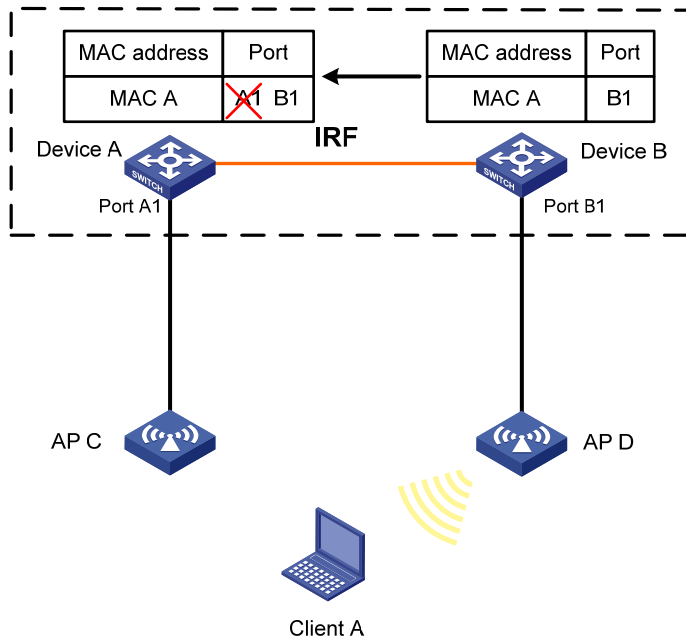
As shown in [Figure 6](#), Device A and Device B form an IRF fabric enabled with MAC address roaming. They connect to AP C and AP D, respectively. When Client A associates with AP C, Device A learns the MAC address of Client A and advertises it to the member switch Device B.

Figure 6 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 7](#).

Figure 7 MAC address tables of devices when Client A roams to AP D



To enable MAC address roaming:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable MAC address roaming.	mac-address mac-roaming enable	Disabled by default.

Enabling MAC address migration log notifying

This feature records and notifies MAC address migration information, including MAC addresses that migrate, IDs of VLANs to which MAC addresses belong, source interfaces from which MAC addresses migrate, and current interfaces with which MAC addresses associate, last migration time, and migration times in the last one minute.

MAC address migration refers to this process: a device learns a MAC address from an interface, Port A for example, and the device later learns the MAC address from another interface, Port B for example. If Port A and Port B belong to the same VLAN, the outgoing interface in the entry for the MAC address is changed to Port B from Port A, which means that the MAC address migrates from Port A to Port B.



TIP:

If a MAC address migrates between two specific interfaces frequently, a Layer 2 loop probably occurs in the network. To discover and locate Layer 2 loops, you can enable MAC address migration log notifying.

To enable MAC address migration log notifying:

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

Step	Command	Remarks
2. Enable MAC address migration log notifying.	mac-flapping notification enable	By default, MAC address migration log notifying is disabled.

The MAC address migration logs of the last one minute are displayed once every one minute.

Displaying and maintaining MAC address tables

Task	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</i> <i>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</i> <i>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

As shown in [Figure 8](#):

- The MAC address of Host A 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 for the host in the MAC address table of the device.
- The MAC address of Host B is 000f-e235-abcd and belongs to VLAN 1. For security, because this host once behaved suspiciously on the network, add a blackhole MAC address entry for the host MAC address, so all packets destined for the host are dropped.
- Set the aging timer for dynamic MAC address entries to 500 seconds.

Figure 8 Network diagram



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

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

Display information about the blackhole MAC address table.

```
[Sysname] display mac-address blackhole
```

MAC ADDR	VLAN ID	STATE	PORT INDEX	AGING TIME(s)
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
```

Configuring MAC Information

Overview

Introduction to MAC Information

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

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

How MAC Information works

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

The device writes information and sends messages only for the following MAC addresses: automatically learned source MAC addresses, MAC addresses that pass MAC authentication, MAC addresses that pass 802.1X authentication, MAC addresses matching OUI addresses in the voice VLAN feature, and secure MAC addresses. The device does not write information or send messages for blackhole MAC address, static MAC addresses, dynamic MAC addresses that are manually configured, multicast MAC addresses, and local MAC addresses.

For more information about MAC authentication, 802.1X, and secure MAC addresses in port security, see *Security Configuration Guide*. For more information about voice VLAN and OUI addresses, see "Voice VLAN configuration."

Enabling MAC Information globally

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable MAC Information globally.	mac-address information enable	Disabled by default.

Enabling MAC Information on an interface

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

To enable MAC Information on an interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Enable MAC Information on the interface.	mac-address information enable { added deleted }	Disabled by default.

Configuring MAC Information mode

Step	Command	Remarks
1. Enter system view.	system-view	N/A
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, change the interval for sending Syslog or trap messages.

To set the interval for sending Syslog or trap messages:

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

If the MAC Information queue length is 0, the device sends a Syslog or trap message immediately after learning or deleting a MAC address.

If it is not 0, the device stores MAC address changes in the queue:

- When the timer set for sending Syslog or trap messages does not expire, the device overwrites the last piece of information written into the queue with the new MAC address change if the queue has been exhausted.
- When the timer set for sending Syslog or trap messages expires, the device sends Syslog or trap messages regardless of whether the queue has been exhausted.

To configure the MAC Information queue length:

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


Step	Command	Remarks
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

As shown in [Figure 9](#), enable MAC Information on GigabitEthernet 1/0/1 on Device to send 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



Configuration procedure

1. Configure Device to send Syslog messages to Host B (see *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

```

Configuring Ethernet link aggregation

Overview

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

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

As shown in [Figure 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 the three physical Ethernet links. At the same time, the three Ethernet links back up one another.

Figure 10 Ethernet link aggregation



Basic concepts

Aggregation group, member port, and aggregate interface

Link aggregation is implemented by combining Ethernet interfaces into a link aggregation group. Each link aggregation group has one logical aggregate interface. To an upper layer entity that uses the link aggregation service, a link aggregation group appears to be a single logical link and data traffic is transmitted through the aggregate interface. 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](#)."

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 automatically created.

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 the operational key.

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

Configuration classes

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

- **Port attribute configurations**—Includes port rate, duplex mode, and link status (up/down). These are the most basic port configurations.
- **Class-two configurations**—A member port can be placed in Selected state only if it has the same class-two configurations as the aggregate interface. Class-two configurations made on an aggregate interface are automatically synchronized to all its member ports. These configurations are retained on the member ports even after the aggregate interface is removed.

Table 2 Class-two configurations

Feature	Considerations
Port isolation	Whether the port has joined an isolation group
QinQ	QinQ enable state (enable/disable), TPID for VLAN tags, outer VLAN tags to be added, inner-to-outer VLAN priority mappings, inner-to-outer VLAN tag mappings, inner VLAN ID substitution mappings
VLAN	Permitted VLANs, 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

NOTE:

Any class-two configuration change might affect the aggregation state of link aggregation member ports and ongoing traffic. To be sure that you are aware of the risk, the system displays a warning message every time you attempt to change a class-two configuration setting on a member port.

- **Class-one configurations**—Include settings that 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. The class-one configuration for a member port is effective only when the member port leaves the aggregation group.

Reference port

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

LACP

The IEEE 802.3ad Link Aggregation Control Protocol (LACP) enables dynamic aggregation of physical links. It uses link aggregation control protocol data units (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 3](#).

Table 3 Basic and extended LACP functions

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

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

2. LACP priorities

LACP priorities have the following types: system LACP priority and port aggregation priority.

Table 4 LACP priorities

Type	Description	Remarks
System LACP priority	Used by two peer devices (or systems) to determine which one is superior in link aggregation. In dynamic link aggregation, the system that has higher system LACP priority sets the Selected state of member ports on its side first, and then the system that has lower priority sets the port state accordingly.	The smaller the priority value, the higher the priority.
Port aggregation priority	Determines the likelihood of a member port to be selected on a system. The higher the port aggregation priority, the higher the likelihood.	

3. LACP timeout interval

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

Link aggregation modes

Link aggregation has the following modes: dynamic and static. Dynamic link aggregation uses LACP and static link aggregation does not. [Table 5](#) compares the two aggregation modes.

Table 5 A comparison between static and dynamic aggregation modes

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

The following points apply to a dynamic link aggregation group:

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

Aggregating links in static mode

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

The static link aggregation process comprises:

- [Selecting a reference port](#)
- [Setting the aggregation state of each member port](#)

Selecting a reference port

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

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

The candidate ports are sorted by aggregation priority, duplex, and speed in the following order:

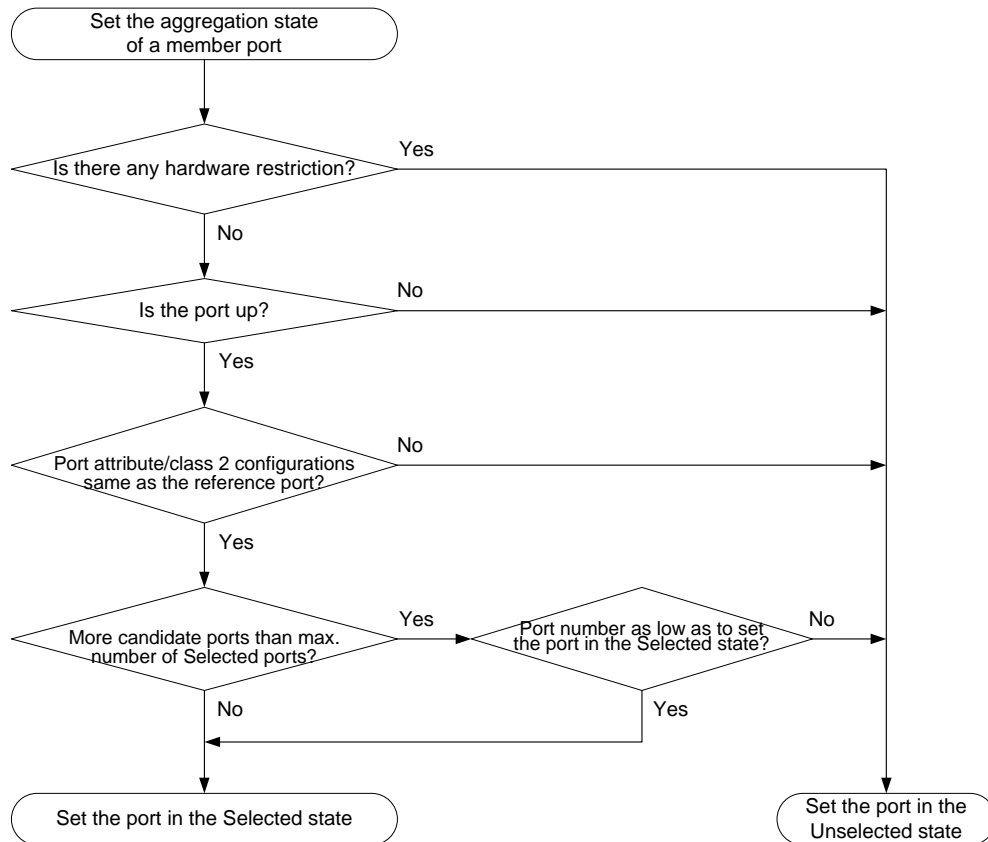
- Lowest aggregation priority value
- Full duplex/high speed
- Full duplex/low speed
- Half duplex/high speed
- Half duplex/low speed

The one at the top is selected as the reference port. If two ports have the same aggregation priority, duplex mode, and speed, the one with the lower port number wins.

Setting the aggregation state of each member port

After selecting the reference port, the static aggregation group sets the aggregation state of each member port, as shown in [Figure 11](#). After the static aggregation group has reached the limit on Selected ports, any port assigned to the group is placed in Unselected state to avoid traffic interruption on the current Selected ports.

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



Aggregating links in dynamic mode

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

The dynamic link aggregation process comprises:

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

Selecting a reference port

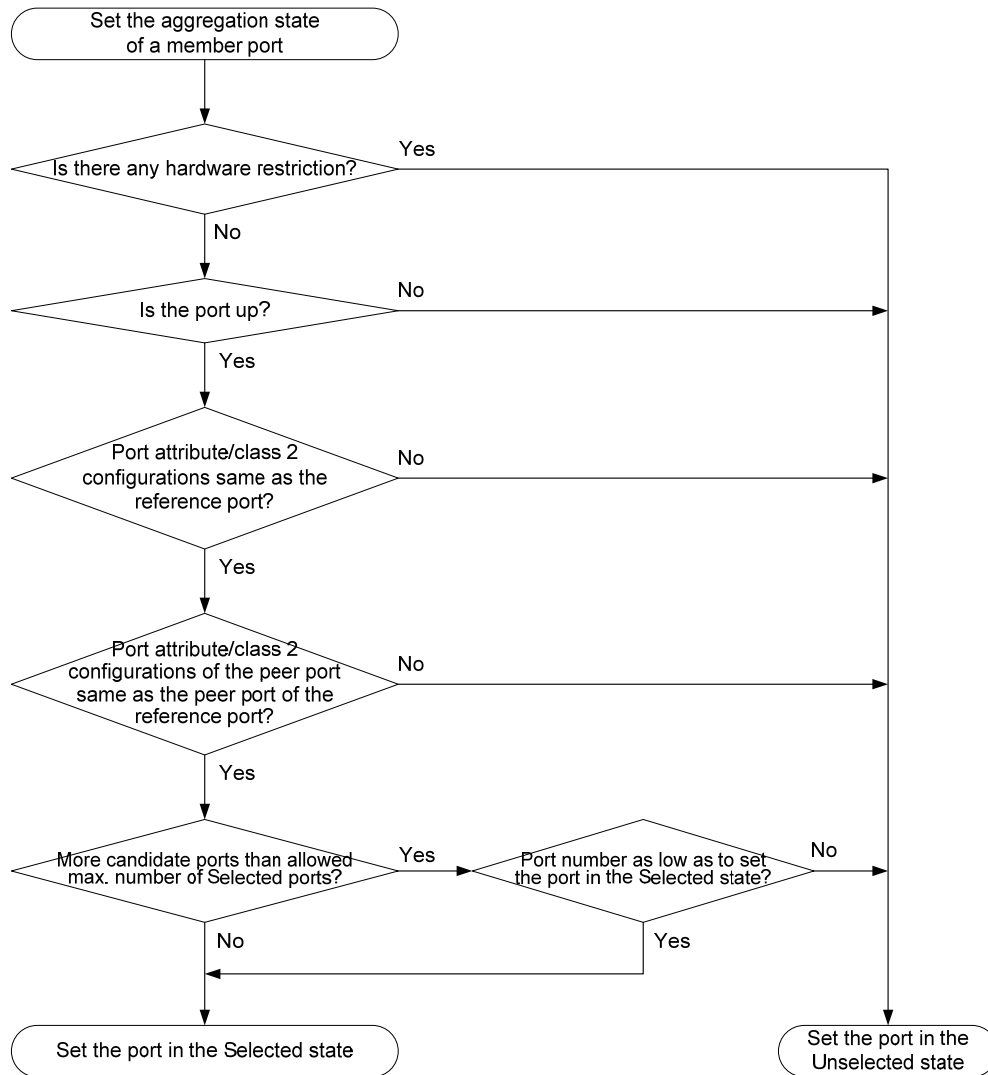
The local system (the actor) and the remote system (the partner) negotiate a reference port using the following workflow:

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

Setting the aggregation state of each member port

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

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



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

A dynamic link aggregation group preferably sets full-duplex ports as the Selected ports, and will set one, and only one, half-duplex port as a Selected port when none of the full-duplex ports can be selected or only half-duplex ports exist in the group.

When the aggregation state of a member port changes, the aggregation state of its peer port also changes.

After the Selected port limit has been reached, a port assigned to the dynamic aggregation group is placed in Selected state if it is more eligible for being selected than a current member port.

The port assigned to the dynamic aggregation group after the Selected port limit has been reached is placed in Selected state if it is more eligible for being selected than a current member port.

Load-sharing criteria for link aggregation groups

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

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

- Source/Destination MAC addresses
- Source/Destination service port numbers
- Ingress ports
- Source/Destination IP addresses

Configuration restrictions and guidelines

Follow these guidelines when you configure a link aggregation group:

- To ensure stable aggregation state and service continuity, do not change port attributes or class-two configurations on any member port. If you must, make sure you understand its impact on the live network. Any port attribute or class-two configuration change might affect the aggregation state of link aggregation member ports and ongoing traffic.

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

Ethernet link aggregation configuration task list

Task	Remarks
Configuring an aggregation group	Configuring a static aggregation group
	Configuring a dynamic aggregation group
Configuring an aggregate interface	Configuring the description of an aggregate interface
	Enabling link state traps for an aggregate interface
	Limiting the 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	

Configuring an aggregation group

Configuration guidelines

- You cannot assign a port to a Layer 2 aggregation group if any of the features listed in [Table 6](#) is configured on the port.

Table 6 Features incompatible with Layer 2 aggregation groups

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

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

Configuring a static aggregation group

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
3. Exit to system view.	quit	N/A
4. Assign an Ethernet interface to the aggregation group.	a. Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> b. Assign the Ethernet interface to the aggregation group: port link-aggregation group <i>number</i>	Repeat these two sub-steps to assign more Layer 2 Ethernet interfaces to the aggregation group.
5. Assign the port an aggregation priority.	link-aggregation port-priority <i>port-priority</i>	Optional. By default, the aggregation priority of a port is 32768. Changing the aggregation priority of a port might affect the aggregation state of the ports in the static aggregation group.

Configuring a dynamic aggregation group

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Set the system LACP priority.	lacp system-priority <i>system-priority</i>	Optional. By default, the system LACP priority is 32768. Changing the system LACP priority might affect the aggregation state of the ports in a dynamic aggregation group.
3. Create a Layer 2 aggregate interface and enter Layer 2 aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	When you create a Layer 2 aggregate interface, the system automatically creates a Layer 2 static aggregation group numbered the same.
4. Configure the aggregation group to operate in dynamic aggregation mode.	link-aggregation mode dynamic	By default, an aggregation group operates in static aggregation mode.
5. Exit to system view.	quit	N/A
6. Assign an Ethernet interface to the aggregation group.	a. Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> b. Assign the Ethernet interface to the aggregation group: port link-aggregation group <i>number</i>	Repeat these two sub-steps to assign more Layer 2 Ethernet interfaces to the aggregation group.
7. Assign the port an aggregation priority.	link-aggregation port-priority <i>port-priority</i>	Optional. By default, the aggregation priority of a port is 32768. Changing the aggregation priority of a port might affect the aggregation state of the ports in the dynamic aggregation group.
8. 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 of the configurations that can be performed on Layer 2 Ethernet interfaces can also be performed on Layer 2 interfaces.

Configuring the description of an aggregate interface

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

To configure the description of an aggregate interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	N/A
3. Configure the description of the aggregate interface.	description <i>text</i>	Optional By default, the description of an interface is in the format of <i>interface-name</i> Interface , such as Bridge-Aggregation1 Interface .

Enabling link state traps for an aggregate interface

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

To enable link state traps on an aggregate interface:

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

Step	Command	Remarks
3. Enter aggregate interface view.	<ul style="list-style-type: none"> interface bridge-aggregation <i>interface-number</i> 	N/A
4. Enable link state traps for the aggregate interface.	enable snmp trap updown	Optional. Enabled by default.

Limiting the number of Selected ports for an aggregation group

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.

By default, the maximum number of Selected ports allowed in an aggregation group depends on the hardware capabilities of the member ports. After you manually configure the maximum number of Selected ports in an aggregation group, the maximum number of Selected ports allowed in the aggregation group is the lower value of the two upper limits.

You can configure redundancy between two ports by assigning the two ports to an aggregation group and configuring the maximum number of Selected ports allowed in the aggregation group as 1. In this way, only one Selected port is allowed in the aggregation group at any point in time, while the Unselected port serves as a backup port.

Configuration guidelines

Follow these guidelines when you configure the port threshold settings:

- If you set a minimum threshold for a static aggregation group, also make the same setting for its peer aggregation group to guarantee correct aggregation.
- Make sure the two link aggregation ends have the same minimum and maximum numbers of selected ports.

Make sure you understand the following impacts of the port threshold settings:

- Configuring the minimum number of Selected ports required to bring up an aggregation group may cause all the member ports in the aggregation group to become unselected.
- Configuring the maximum number of Selected ports in an aggregation group may cause some of the selected member ports in the aggregation group to become unselected.

To limit the number of Selected ports for an aggregation group:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	<ul style="list-style-type: none">interface bridge-aggregation <i>interface-number</i>	N/A
3. Set the minimum number of Selected ports for the aggregation group.	link-aggregation selected-port minimum <i>number</i>	Not specified by default.
4. Set the maximum number of Selected ports for the aggregation group.	link-aggregation selected-port maximum <i>number</i>	By default, the maximum number of Selected ports allowed in an aggregation group depends on only the hardware capabilities of the member ports.

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

To shut down an aggregate interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	N/A
3. Shut down the aggregate interface.	shutdown	By default, aggregate interfaces are up.

Restoring the default settings for an aggregate interface

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

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

Configuring load sharing for link aggregation groups

Configuring load-sharing criteria for link aggregation groups

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

You can configure global or group-specific load-sharing criteria. A link aggregation group preferentially uses the group-specific load-sharing criteria. If no group-specific load-sharing criteria are available, the group uses the global load-sharing criteria.

NOTE:

The load sharing criteria configuration applies to only unicast packets, and can change the load sharing criteria for unicast packets. Broadcast packets and multicast packets always use the default load sharing criteria.

Configuring the global link-aggregation load-sharing criteria

Step	Command	Remarks
1. Enter system view.	system-view	N/A
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 } *	By default, the system selects the global load sharing criteria according to the packet type.

In system view, the switch supports the following load-sharing criteria and combinations:

- Load-sharing criteria automatically determined based on the packet type
- 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
- Destination IP address and destination port
- Source IP address, source port, destination IP address, and destination port
- Any combination of incoming port, source MAC address, and destination MAC address

Configuring group-specific load-sharing criteria

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	N/A
3. Configure the load-sharing criteria for the aggregation group.	link-aggregation load-sharing mode { destination-ip destination-mac source-ip source-mac } *	The default load-sharing criteria are the same as the global load-sharing criteria.

In Layer 2 aggregate interface view, the switch supports the following load-sharing criteria and combinations:


- Load-sharing criteria automatically determined based on the packet type
- Source IP address
- Destination IP address
- Source MAC address
- Destination MAC address
- Destination IP address and source IP 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-switch link aggregation scenario to distribute traffic preferentially across member ports on the ingress switch rather than all member ports.

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

Figure 13 Load sharing process for cross-switch link aggregation in an IRF fabric



To enable local-first load sharing for link aggregation:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable local-first load-sharing for link aggregation.	link-aggregation load-sharing mode local-first	Optional. Enabled by default. Local-first load sharing for link aggregation takes effect on only known unicast packets.

Enabling link-aggregation traffic redirection

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

Link-aggregation traffic redirection applies only to dynamic link aggregation groups and only to known unicast packets.

After link-aggregation traffic redirection is enabled, do not add an Ethernet interface configured with physical state change suppression to an aggregation group. Otherwise, Selected ports in the aggregation group might work improperly. For more information about physical state change suppression, see the **link-delay** command in Ethernet interface configuration commands.

To enable link-aggregation traffic redirection:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable link-aggregation traffic redirection.	link-aggregation lacp traffic-redirect-notification enable	Optional. Disabled by default.

CAUTION:

- 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.

Displaying and maintaining Ethernet link aggregation

Task	Command	Remarks
Display information about 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 for 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 summary information about 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. Make sure 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 aggregation group on both Device A and Device B. Enable VLAN 10 at one end of the aggregate link to communicate with VLAN 10 at the other end, 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 the source and destination MAC addresses.

Figure 14 Network diagram



Configuration procedure

1. Configure Device A:
Create VLAN 10, and assign port GigabitEthernet 1/0/4 to VLAN 10.

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


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

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

Create Layer 2 aggregate interface Bridge-Aggregation 1.

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

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

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

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

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

Configure Device A 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 in the same way as you configure Device A.

3. Verify the configurations:

Display 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 aggregation group on both Device A and Device B, enable VLAN 10 at one end of the aggregate link to communicate with VLAN 10 at the other end, 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



Configuration procedure

1. Configure Device A:

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

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

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

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

Create Layer 2 aggregate interface Bridge-Aggregation 1, and 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.

```
[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 in the same way as you configure Device A.

3. Verify the configurations:

Display 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.

Configuring port isolation

Port isolation enables isolating Layer 2 traffic for data privacy and security without using VLANs. You can also use this feature to isolate the hosts in a VLAN from one another.

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

The switch series supports only one isolation group "isolation group 1." The isolation group is automatically created and cannot be deleted. There is no limit on the number of member ports.

Assigning a port to the isolation group

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none">Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>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>	<p>Use one of the commands.</p> <ul style="list-style-type: none">In Ethernet interface view, the subsequent configurations apply to the current port.In Layer 2 aggregate interface view, the subsequent configurations apply to the Layer 2 aggregate interface and all its member ports.In port group view, the subsequent configurations apply to all ports in the port group.
3. Assign the port or ports to the isolation group as an isolated port or ports.	port-isolate enable	No ports are added to the isolation group by default.

Displaying and maintaining the isolation group

Task	Command	Remarks
Display isolation group information.	display port-isolate group [{ begin exclude include } <i>regular-expression</i>]	Available in any view


Port isolation configuration example

Network requirements

As shown in [Figure 16](#), Host A, Host B, and Host C are connected to GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, and GigabitEthernet 1/0/3 of Device, and Device is connected to the Internet through GigabitEthernet 1/0/4. All these ports are in the same VLAN.

Configure Device to provide Internet access for all the hosts and isolate them from one another.

Figure 16 Networking diagram



Configuration procedure

Add ports GigabitEthernet 1/0/1, GigabitEthernet 1/0/2, and GigabitEthernet 1/0/3 to the isolation group.

```
<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:
    GigabitEthernet 1/0/1    GigabitEthernet 1/0/2    GigabitEthernet 1/0/3
```

Configuring spanning tree protocols

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

The recent versions of STP include the Rapid Spanning Tree Protocol (RSTP), Per VLAN Spanning Tree (PVST), and the Multiple Spanning Tree Protocol (MSTP).

STP

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

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

STP protocol packets

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

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

STP uses the following types of BPDUs:

- **Configuration BPDUs**—Used by network devices to calculate a spanning tree and maintain the spanning tree topology.
- **Topology change notification (TCN) BPDUs**—Notify network devices of the network topology changes.

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

- **Root bridge ID**—Consisting of the priority and MAC address of the root bridge.
- **Root path cost**—Cost of the path to the root bridge denoted by the root identifier from the transmitting bridge.
- **Designated bridge ID**—Consisting of the priority and MAC address of the designated bridge.
- **Designated port ID**—Consisting of the priority and global port number of the designated port.
- **Message age**—Age of the configuration BPDU while it propagates in the network.
- **Max age**—Maximum age of the configuration BPDU stored on the switch.
- **Hello time**—Configuration BPDU transmission interval.
- **Forward delay**—Delay that STP bridges use to transit port state.

Basic concepts in STP

Root bridge

A tree network must have a root bridge. The entire network contains only one root bridge. The other bridges in the network are called "leaf nodes." The root bridge is not permanent, but can change with changes of the network topology.

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

Root port

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


Designated bridge and designated port

Table 7 Description of designated bridges and designated ports

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

As shown in Figure 17, Device B and Device C are directly connected to a 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 Designated bridges and designated ports



Path cost

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

Calculation process of the STP algorithm

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

The STP algorithm uses the following calculation process:

1. Initialize the state.

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

2. Select the root bridge.

Initially, each STP-enabled 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. Select the root ports and designated ports on non-root bridges.

Table 8 Selection of the root port and designated ports

Step	Description
1	A non-root-bridge device regards the port on which it received the optimum configuration BPDU as the root port. Table 9 describes how the optimum configuration BPDU is selected.
2	<p>Based on the configuration BPDU and the path cost of the root port, the device calculates a designated port configuration BPDU for each of the other ports.</p> <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	<p>The device compares the calculated configuration BPDU with the configuration BPDU on the port whose port role will be defined, and acts depending on the result of the comparison.</p> <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 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 cannot send BPDUs or forward data traffic.

NOTE:

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

Table 9 Selection of the optimum configuration BPDU

Step	Actions
1	<p>Upon receiving a configuration BPDU on a port, the device compares the priority of the received configuration BPDU with that of the configuration BPDU generated by the port, and:</p> <ul style="list-style-type: none"> • If the former priority is lower, the device discards the received configuration BPDU and keeps the configuration BPDU the port generated. • If the former priority is higher, the device replaces the content of the configuration BPDU generated by the port with the content of the received configuration BPDU.
2	The device compares the configuration BPDUs of all the ports and chooses the optimum configuration BPDU.


The following are the principles of configuration BPDU comparison:

- The configuration BPDU with the lowest root bridge ID has the highest priority.
- If 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 ports value, their designated bridge IDs, designated port IDs, and the IDs of the receiving ports are compared in sequence. The configuration BPDU that contains the smallest ID wins.

A tree-shape topology forms when the root bridge, root ports, and designated ports are selected.

Figure 18 describes with an example how the STP algorithm works. This example shows a simplified spanning tree calculation process.

Figure 18 The STP algorithm



As shown in Figure 18, the priority values of Device A, Device B, and Device C are 0, 1, and 2, and the path costs of links among the three devices are 5, 10, and 4, respectively.

4. Initial state of each device

Table 10 Initial state of each device

Device	Port name	Configuration BPDU on the port
Device A	Port A1	{0, 0, 0, Port A1}

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

NOTE:

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

5. Comparison process and result on each device

Table 11 Comparison process and result on each device

Device	Comparison process	Configuration BPDU on ports after comparison
Device A	<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 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 B	<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 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 new event triggers a spanning tree calculation process, 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}

NOTE:

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

After the comparison processes described in [Table 11](#), a spanning tree with Device A as the root bridge is established, and the topology is shown in [Figure 19](#).

Figure 19 The final calculated spanning tree



The configuration BPDU forwarding mechanism of STP

The configuration BPDUs of STP are forwarded following these guidelines:

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

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

STP timers

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

- Forward delay

Forward delay is the delay time for port state transition.

A path failure can cause spanning tree re-calculation to adapt the spanning tree structure to the change. However, the resulting new configuration BPDU cannot propagate throughout the network immediately. If the newly elected root ports and designated ports start to forward data immediately, a temporary loop will likely occur.

For this reason, as a mechanism for state transition in STP, the newly elected root ports or designated ports require twice the forward delay time before they transit to the forwarding state to make sure that the new configuration BPDU has propagated throughout the network.

- Hello time

The device sends hello packets at the hello time interval to the neighboring devices to make sure that the paths are fault-free.

- Max age

The device uses the max age to determine whether a stored configuration BPDU has expired and discards it if the max age is exceeded.

RSTP

RSTP achieves rapid network convergence by allowing a newly elected root port or designated port to enter the forwarding state much faster than STP.

A newly elected RSTP root port rapidly enters the forwarding state if the old root port on the device has stopped forwarding data and the upstream designated port has started forwarding data.

A newly elected RSTP designated port rapidly enters the forwarding state if it is an edge port (a port that directly connects to a user terminal rather than to another network device or a shared LAN segment) or it connects to a point-to-point link (to another device). Edge ports directly enter the forwarding state. Connecting to a point-to-point link, a designated port enters the forwarding state immediately after the device receives a handshake response from the directly connected device.

PVST

PVST was introduced to improve link bandwidth usage in network environments where multiple virtual LANs (VLANs) exist. Unlike STP and RSTP whose bridges in a LAN must forward their VLAN packets in the same spanning tree, PVST allows each VLAN to build a separate spanning tree.

PVST uses the following BPDUs:

- **STP BPDUs**—Sent by access ports according to the VLAN status, or by trunk ports and hybrid ports according to the status of VLAN 1.
- **PVST BPDUs**—Sent by trunk port and hybrid ports according to the status of permitted VLANs except VLAN 1.

MSTP

STP, RSTP, and PVST limitations

STP does not support rapid state transition of ports. A newly elected port must wait twice the forward delay time before it transits to the forwarding state, even if it connects to a point-to-point link or is 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, and the packets from all VLANs are forwarded along the same spanning tree, so redundant links cannot be blocked based on VLAN and traffic cannot be load-shared among VLANs.

The number of PVST BPDUs generated grows with that of permitted VLANs on trunk ports. When the status of a trunk port transits, network devices might be overloaded to re-calculate a large number of spanning trees.

MSTP features

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

MSTP provides the following features:

- MSTP divides a switched network into multiple regions, each of which contains multiple spanning trees that are independent of one another.
- 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 prunes a loop network into a loop-free tree, which avoids proliferation and endless cycling of packets in a loop network. In addition, it supports load balancing of VLAN data by providing multiple redundant paths for data forwarding.

MSTP basic concepts

Figure 20 shows a switched network that comprises four MST regions, each MST region comprising four MSTP devices. Figure 21 shows the networking topology of MST region 3.

Figure 20 Basic concepts in MSTP


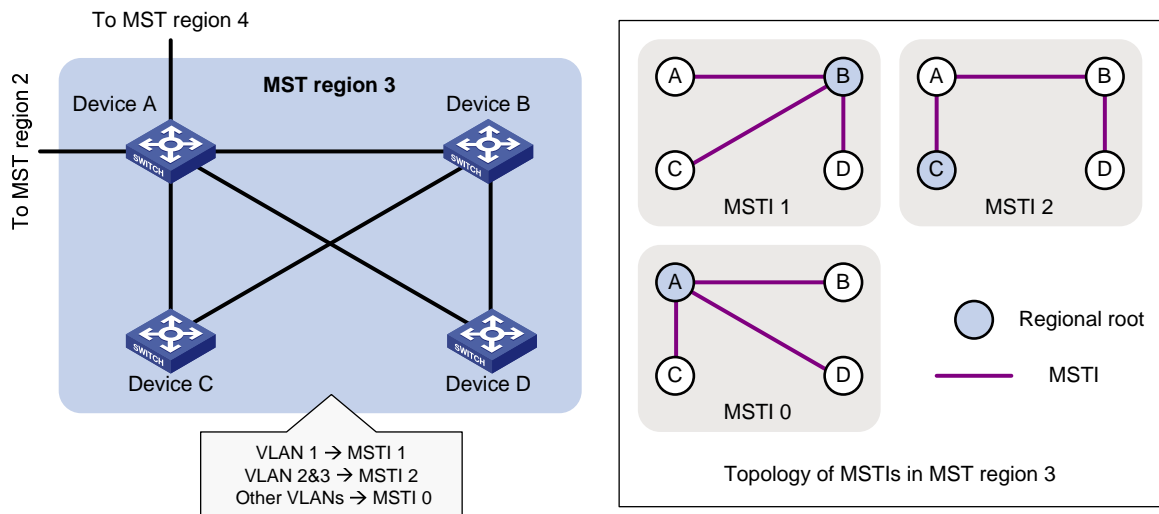


Figure 21 Network diagram and topology of MST region 3



MST region

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

- A spanning tree protocol enabled
- Same region name
- Same VLAN-to-instance mapping configuration
- Same MSTP revision level
- Physically linked together

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

In Figure 21, 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, the VLAN-to-instance mapping table of MST region 3 is: VLAN 1 to MSTI 1, VLAN 2 and VLAN 3 to MSTI 2, and other VLANs to MSTI 0. MSTP achieves load balancing by means of the VLAN-to-instance mapping table.

CST

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

The blue lines in Figure 20 represent the CST.

IST

An internal spanning tree (IST) is a spanning tree that runs in an MST region. It is also called MSTI 0, a special MSTI to which all VLANs are mapped by default.

In Figure 20, MSTI 0 is the IST in MST region 3.

CIST

The common and internal spanning tree (CIST) is a single spanning tree that connects all devices in a switched network. It consists of the ISTs in all MST regions and the CST.

In Figure 20, the ISTs (MSTI 0) in all MST regions plus the inter-region CST constitute the CIST of the entire network.

Regional root

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

For example, in MST region 3 in Figure 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

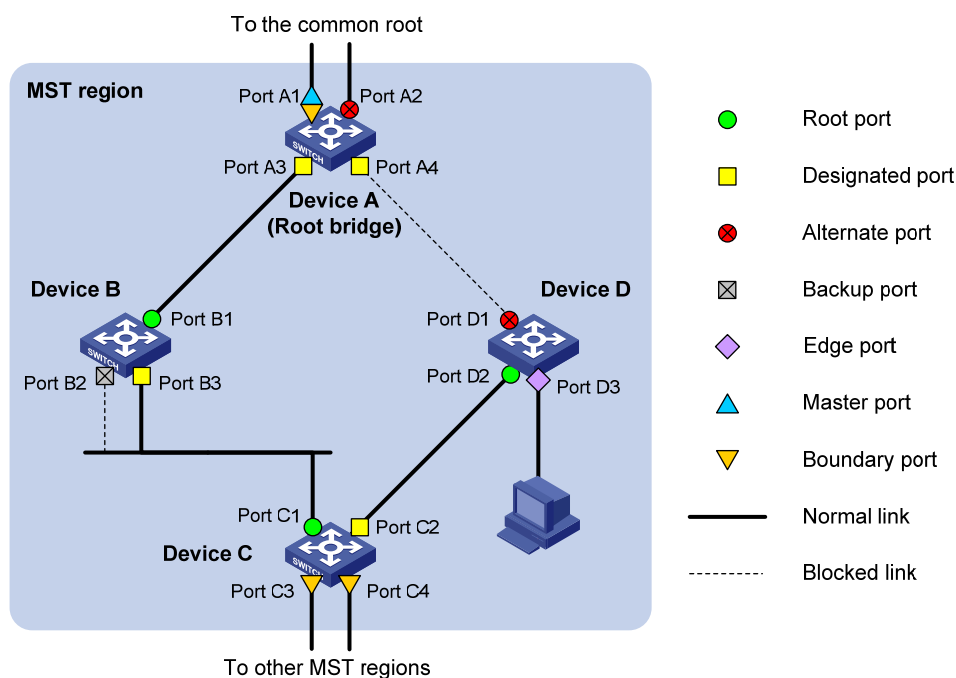
The common root bridge is the root bridge of the CIST.

In Figure 20, for example, the common root bridge is a device in MST region 1.

Port roles

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 the following port roles:

- **Root port**—Forwards data for a non-root bridge to the root bridge. The root bridge does not have any root port.
- **Designated port**—Forwards data to the downstream network segment or device.
- **Alternate port**—The backup port for a root port or master port. When the root port or master port is blocked, the alternate port takes over.
- **Backup port**—The backup port of a designated port. When the designated port is invalid, the backup port becomes the new designated port. A loop occurs when two ports of the same spanning tree device are interconnected, so the device blocks one of the ports. The blocked port acts as the backup.
- **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 not always located on the regional root. It is a root port on the IST or CIST and still a master port on the other MSTIs.
- **Boundary port**—Connects an MST region to another MST region or to an STP/RSTP-running device. In MSTP calculation, a boundary port's role on an MSTI is consistent with its role on the CIST. But that is not true with master ports. A master port on MSTIs is a root port on the CIST.

Port states

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

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

When in different MSTIs, a port can be in different states. A port state is not exclusively associated with a port role. [Table 12](#) lists the port states that each port role supports. (A check mark [✓] indicates that the port supports this state, while a dash [—] indicates that the port does not support this state.)

Table 12 Port states that different port roles support

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

How MSTP works

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

CIST calculation

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

MSTI calculation

Within an MST region, MSTP generates different MSTIs for different VLANs based on the VLAN-to-instance mappings. For each spanning tree, MSTP performs a separate calculation process similar to spanning tree calculation in STP. For more information, see "[Calculation process of the STP algorithm](#)."

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

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

Implementation of MSTP on devices

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

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

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

Protocols and standards

The spanning tree protocols are documented in the following standards:

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

Spanning tree configuration task list

Before configuring a spanning tree, you must determine the spanning tree protocol to be used (STP, RSTP, PVST, or MSTP) and plan the device roles (the root bridge or leaf node).

Configuration restrictions and guidelines

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

STP configuration task list

Task	Remarks
Configuring the root bridge	Required
	Setting the spanning tree mode
	Configure the device to operate in STP mode.
	Optional
	Configuring the root bridge or a secondary root bridge
	Optional
	Configuring the device priority
	Optional
	Configuring the network diameter of a switched network
	Optional
	Configuring spanning tree timers
	Optional
	Configuring the timeout factor
	Optional
	Configuring the maximum port rate
	Optional
	Configuring the mode a port uses to recognize/send MSTP packets
	Optional
	Enabling outputting port state transition information
	Optional
	Enabling the spanning tree feature
	Required

Task	Remarks
Configuring the leaf nodes	Setting the spanning tree mode Required Configure the device to operate in STP mode.
	Configuring the device priority Optional
	Configuring the timeout factor Optional
	Configuring the maximum port rate Optional
	Configuring path costs of ports Optional
	Configuring the port priority Optional
	Configuring the mode a port uses to recognize/send MSTP packets Optional
	Enabling outputting port state transition information Optional
	Enabling the spanning tree feature Required
	Configuring TC snooping Optional
Configuring protection functions Optional	

RSTP configuration task list

Task	Remarks
Configuring the root bridge	Setting the spanning tree mode Required Configure the device to operate in RSTP mode.
	Configuring the root bridge or a secondary root bridge Optional
	Configuring the device priority Optional
	Configuring the network diameter of a switched network Optional
	Configuring spanning tree timers Optional
	Configuring the timeout factor Optional
	Configuring the maximum port rate Optional
	Configuring edge ports Optional
	Configuring the port link type Optional
	Configuring the mode a port uses to recognize/send MSTP packets Optional
	Enabling outputting port state transition information Optional
	Enabling the spanning tree feature Required
Configuring the leaf nodes	Setting the spanning tree mode Required Configure the device to operate in RSTP mode.
	Configuring the device priority Optional
	Configuring the timeout factor Optional

Task	Remarks
Configuring the maximum port rate	Optional
Configuring edge ports	Optional
Configuring path costs of ports	Optional
Configuring the port priority	Optional
Configuring the port link type	Optional
Configuring the mode a port uses to recognize/send MSTP packets	Optional
Enabling outputting port state transition information	Optional
Enabling the spanning tree feature	Required
Performing mCheck	Optional
Configuring TC snooping	Optional
Configuring protection functions	Optional

PVST configuration task list

Task	Remarks
Configuring the root bridge	Required Configure the device to operate in PVST mode.
	Setting the spanning tree mode
	Configuring the root bridge or a secondary root bridge
	Configuring the device priority
	Configuring the network diameter of a switched network
	Configuring spanning tree timers
	Configuring the timeout factor
	Configuring the maximum port rate
	Configuring edge ports
	Configuring the port link type
	Enabling outputting port state transition information
Configuring the leaf nodes	Enabling the spanning tree feature
	Required
	Setting the spanning tree mode
	Required Configure the device to operate in PVST mode.
	Configuring the device priority
	Configuring the timeout factor
	Optional
	Configuring the maximum port rate
	Optional

Task	Remarks
Configuring edge ports	Optional
Configuring path costs of ports	Optional
Configuring the port priority	Optional
Configuring the port link type	Optional
Enabling outputting port state transition information	Optional
Enabling the spanning tree feature	Required
Performing mCheck	Optional
Configuring protection functions	Optional

MSTP configuration task list

Task	Remarks
Configuring the root bridge	Optional By default, the device operates in MSTP mode.
	Setting the spanning tree mode
	Configuring an MST region
	Configuring the root bridge or a secondary root bridge
	Configuring the device priority
	Configuring the maximum hops of an MST region
	Configuring the network diameter of a switched network
	Configuring spanning tree timers
	Configuring the timeout factor
	Configuring the maximum port rate
	Configuring edge ports
	Configuring the port link type
	Configuring the mode a port uses to recognize/send MSTP packets
	Enabling outputting port state transition information
	Enabling the spanning tree feature
Configuring the leaf nodes	Optional By default, the device operates in MSTP mode.
	Setting the spanning tree mode
	Configuring an MST region
	Configuring the device priority
	Configuring the timeout factor
	Configuring the maximum port rate
	Configuring edge ports

Task	Remarks
Configuring path costs of ports	Optional
Configuring the port priority	Optional
Configuring the port link type	Optional
Configuring the mode a port uses to recognize/send MSTP packets	Optional
Enabling outputting port state transition information	Optional
Enabling the spanning tree feature	Required
Performing mCheck	Optional
Configuring Digest Snooping	Optional
Configuring No Agreement Check	Optional
Configuring TC snooping	Optional
Configuring protection functions	Optional

Setting the spanning tree mode

The spanning tree modes include:

- **STP mode**—All ports of the device send STP BPDUs. Select this mode when the peer device of a port supports only STP.
- **RSTP mode**—All ports of the device send RSTP BPDUs. A port in this mode automatically transits to the STP mode when receiving STP BPDUs from the peer device, and a port in this mode does not transit to the MSTP mode when receiving MSTP BPDUs from the peer device.
- **MSTP mode**—All ports of the device send MSTP BPDUs. A port in this mode automatically transits to the STP mode when receiving STP BPDUs from the peer device, and a port in this mode does not transit to the RSTP mode when receiving RSTP BPDUs from the peer device.
- **PVST mode**—The device sends PVST BPDUs through all ports and maintains a spanning tree for each VLAN. The number of VLANs that PVST can maintain instances for depends on the switch model. Suppose the number is n , which is 32 on the HP 5120 EI Switch Series. When you configure PVST on devices of different models in a network, to avoid network failures, make sure that the number of VLANs for which PVST maintains instances does not exceed the lowest n . An HP device running PVST can communicate with third-party devices running PVST or Rapid PVST. When HP devices running PVST communicate with each other or an HP device running PVST communicates with a third-party device running Rapid PVST, the HP devices supports rapid convergence like that provided by RSTP.

The MSTP mode is compatible with the RSTP mode, and the RSTP mode is compatible with the STP mode. The PVST mode's compatibility with the other spanning tree mode varies by port type:

- On an access port, the PVST mode is compatible with any other spanning tree mode in any VLAN.
- On a trunk or hybrid port, the PVST mode is compatible with any other spanning tree mode in only VLAN 1.

To set the spanning tree mode:

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

Step	Command	Remarks
2.	Set the spanning tree mode. <code>stp mode { stp rstp mstp pvst }</code>	MSTP mode by default.

Configuring an MST region

Two or more spanning tree 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, MST region revision level, and the same VLAN-to-instance mapping entries in the MST region, and each two devices are connected by a physical link.

Configuration restrictions and guidelines

- The configuration of MST region-related parameters, especially the VLAN-to-instance mapping table, will result in a new spanning tree calculation. To reduce the possibility of topology instability, the MST region configuration takes effect only after you activate it by using the **active region-configuration** command, or enable a spanning tree protocol by using the **stp enable** command in the case that the spanning tree protocol is disabled.
- The device in PVST mode automatically maps VLANs to MSTIs, and supports more MSTIs than in MSTP mode. When you change the spanning tree mode from PVST to MSTP, exceeding VLAN-to-instance mappings (arranged in ascending order of MSTI IDs) are silently deleted and cannot be recovered even if you change the spanning tree mode back. To prevent loss of mappings, do not manually configure VLAN-to-instance mappings in PVST mode.

Configuration procedure

To configure an MST region:

Step	Command	Remarks
1.	Enter system view.	<code>system-view</code> N/A
2.	Enter MST region view.	<code>stp region-configuration</code> N/A
3.	Configure the MST region name.	<code>region-name name</code> Optional. The MST region name is the MAC address by default.
4.	Configure the VLAN-to-instance mapping table.	<ul style="list-style-type: none"> <code>instance instance-id vlan vlan-list</code> <code>vlan-mapping modulo modulo</code> Optional. Use either command. 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.	<code>revision-level level</code> Optional. 0 by default.
6.	Display the MST region configurations that are not activated yet.	<code>check region-configuration</code> Optional.
7.	Activate MST region configuration manually.	<code>active region-configuration</code> N/A

Step	Command	Remarks
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

Configuring the root bridge or a secondary root bridge

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

A device has independent roles in different spanning trees. It can act as the root bridge in one spanning tree and as a secondary root bridge in another. However, one device cannot be the root bridge and a secondary root bridge in the same spanning tree.

A spanning tree can have one root bridge only. If two or more devices are designated as the root bridge in a spanning tree at the same time, the device with the lowest MAC address wins.

When the root bridge of an instance fails or is shut down, the secondary root bridge (if you have specified one) can take over the role of the primary root bridge. However, if you specify a new primary root bridge for the instance then, the one you specify, not the secondary root bridge will become the root bridge. If you have specified multiple secondary root bridges for an instance, when the root bridge fails, the secondary root bridge with the lowest MAC address is selected as the new root bridge.

Configuration restrictions and guidelines

- You can specify one root bridge for each spanning tree, regardless of the device priority settings. Once you specify a device as the root bridge or a secondary root bridge, you cannot change its priority.
- You can configure the current device as the root bridge by setting the device priority to 0. For the device priority configuration, see "[Configuring the device priority](#)."

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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the current device as the root bridge.	<ul style="list-style-type: none"> In STP/RSTP mode: stp root primary In PVST mode: stp vlan <i>vlan-list</i> root primary In MSTP mode: stp [<i>instance instance-id</i>] root primary 	Use one of the commands. By default, a device does not function as the root bridge.

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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the current device as a secondary root bridge.	<ul style="list-style-type: none">In STP/RSTP mode: stp root secondaryIn PVST mode: stp vlan <i>vlan-list</i> root secondaryIn MSTP mode: stp [<i>instance instance-id</i>] root secondary	<p>Use one of the commands.</p> <p>By default, a device does not function as a secondary root bridge.</p>

Configuring the device priority

⚠ CAUTION:

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

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

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the priority of the current device.	<ul style="list-style-type: none">In STP/RSTP mode: stp priority <i>priority</i>In PVST mode: stp vlan <i>vlan-list</i> priority <i>priority</i>In MSTP mode: stp [<i>instance instance-id</i>] priority <i>priority</i>	<p>Use one of the commands.</p> <p>The default setting is 32768.</p>

Configuring the maximum hops of an MST region

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

Configuration BPDUs sent by the regional root bridge always have a hop count set to the maximum value. When a device receives this configuration BPDU, it decrements the hop count by 1, and uses the new hop count in the BPDUs that it propagates. When the hop count of a BPDU reaches 0, it is discarded by the

device that received it. This prevents devices beyond the reach of the maximum hop from participate in spanning tree calculation, so the size of the MST region is limited.

Make this configuration on the root bridge only. All other devices 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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the maximum hops of the MST region.	stp max-hops <i>hops</i>	20 by default.

Configuring the network diameter of a switched network

Any two terminal devices in a switched network are connected through a specific path composed of a series of devices. The network diameter is the number of devices on the path composed of the most devices. The network diameter is a parameter that indicates the network size. A bigger network diameter indicates a larger network size. Based on the network diameter you configured, the system automatically sets an optimal hello time, forward delay, and max age for the device.

To configure the network diameter of a switched network:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the network diameter of the switched network.	<ul style="list-style-type: none">In STP/RSTP/MSTP mode: stp bridge-diameter <i>diameter</i>In PVST mode: stp vlan <i>vlan-list</i> bridge-diameter <i>diameter</i>	<ul style="list-style-type: none">Use one of the commands.The default setting is 7.

NOTE:

- In STP/RSTP/MSTP mode, each MST region is considered as a device and the configured network diameter takes effect only on the CIST (or the common root bridge), but not on the MSTIs.
- In PVST mode, the network diameter configuration takes effect on the root bridge only.

Configuring spanning tree timers

The following timers are used for spanning tree calculation:

- Forward delay

It is the delay time for port state transition. To prevent temporary loops on a network, the spanning tree sets an intermediate port state, the learning state, before it transits from the discarding state to the forwarding state, and requires that the port transits its state after a forward delay timer to make sure that the state transition of the local port keeps synchronized with the peer.

- Hello time

The device detects whether a link failure has occurred with the hello time interval. The spanning tree sends a configuration BPDU every hello time interval. If the device receives no configuration BPDUs within the hello time interval, it recalculates the spanning tree.

- **Max age**

In the CIST of an MSTP network or each VLAN of a PVST network, the device uses the max age parameter to determine whether a configuration BPDU received by a port has expired. If a port receives a configuration BPDU that has expired, that MSTI must be re-calculated. The max age timer does not take effect on MSTIs.

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

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

HP does not recommend you to manually set the spanning tree timers. Instead, you can specify the network diameter and let spanning tree protocols automatically calculate the timers based on the network diameter. If the network diameter uses the default value, the timers also use their default values.

Configure the timers on the root bridge only, and the timer settings on the root bridge apply to all devices on the entire switched network.

Configuration restrictions and guidelines

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

Configuration procedure

To configure the spanning tree timers:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the forward delay timer.	<ul style="list-style-type: none">• In STP/RSTP/MSTP mode: stp timer forward-delay <i>time</i>• In PVST mode: stp vlan <i>vlan-list</i> timer forward-delay <i>time</i>	<p>Optional.</p> <p>Use one of the commands.</p> <p>The default setting is 15 seconds.</p>

Step	Command	Remarks
3. Configure the hello timer.	<ul style="list-style-type: none"> In STP/RSTP/MSTP mode: stp timer hello <i>time</i> In PVST mode: stp vlan <i>vlan-list</i> timer hello <i>time</i> 	Optional. Use one of the commands. The default setting is 2 seconds.
4. Configure the max age timer.	<ul style="list-style-type: none"> In STP/RSTP/MSTP mode: stp timer max-age <i>time</i> In PVST mode: stp vlan <i>vlan-list</i> timer max-age <i>time</i> 	Optional. Use one of the commands. The default setting is 20 seconds.

Configuring the timeout factor

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

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

Sometimes a device might fail to receive a BPDU from the upstream device because the upstream device is busy. If a spanning tree calculation occurs, the calculation can fail and also waste network resources. In a stable network, you can prevent undesired spanning tree calculations by setting the timeout factor to 5, 6, or 7.

To configure the timeout factor:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the timeout factor of the device.	stp timer-factor <i>factor</i>	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.

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 spanning tree protocols from using excessive network resources when the network becomes unstable. HP recommends you to use the default setting.

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

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

Step	Command	Remarks
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.
3. Configure the maximum rate of the ports.	stp transmit-limit <i>limit</i>	10 by default.

Configuring 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 network topology change occurs, an edge port will not cause a temporary loop. Because a device does not determine whether a port is directly connected to a terminal, you must manually configure the port as an edge port. After that, the port can transit rapidly from the blocked state to the forwarding state.

Configuration restrictions and guidelines

- If BPDU guard is disabled, a port set as an edge port will become a non-edge port again if it receives a BPDU from another port. To restore the edge port, re-enable it.
- If a port directly connects to a user terminal, configure it as an edge port and enable BPDU guard for it. This enables the port to transit to the forwarding state quickly while ensuring network security.
- You cannot configure edge port settings and loop guard on a port at the same time.

Configuration procedure

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.
3. Configure the current ports as edge ports.	stp edged-port enable	All ports are non-edge ports by default.

Configuring path costs of ports

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

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

Specifying a standard for the device to use when it calculates the default path cost

⚠ CAUTION:

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

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

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

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

Table 13 Mappings between the link speed and the path cost

Link speed	Port type	Path cost		
		IEEE 802.1d-1998	IEEE 802.1t	Private standard
0	N/A	65535	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

Link speed	Port type	Path cost		
		IEEE 802.1d-1998	IEEE 802.1t	Private standard
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

Configuration restrictions and guidelines

- When it calculates path cost for an aggregate interface, IEEE 802.1t takes into account the number of Selected ports in its aggregation group, but IEEE 802.1d-1998 does not. 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.
- When multiple ports operate at a rate higher than 10 Gbps and the standard for default path cost calculation is **dot1d-1998** or **legacy**, the path cost of a single port or an aggregate interface takes the smallest value. As a result, the forwarding path selected might not be optimal. To solve this problem, use **dot1t** as the standard for default path cost calculation, or manually set the path cost for a port (see "[Configuring path costs of ports](#)").

Configuration procedure

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Specify a standard for the device to use when it calculates the default path costs of its ports.	stp pathcost-standard { dot1d-1998 dot1t legacy }	Optional. legacy by default.

Configuring path costs of ports

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

To configure the path cost of ports:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.

Step	Command	Remarks
3. Configure the path cost of the ports.	<ul style="list-style-type: none"> In STP/RSTP mode: stp cost cost In PVST mode: stp vlan vlan-list cost cost In MSTP mode: stp [instance instance-id] cost cost 	<p>Use one of the commands.</p> <p>By default, the system automatically calculates the path cost of each port.</p>

Configuration example

In MSTP mode, specify the device to calculate the default path costs of its ports by using IEEE 802.1d-1998, and set the path cost of GigabitEthernet 1/0/3 to 200 on MSTI 2.

```
<Sysname> system-view
[Sysname] stp pathcost-standard dot1d-1998
[Sysname] interface gigabitethernet 1/0/3
[Sysname-GigabitEthernet1/0/3] stp instance 2 cost 200
```

In PVST mode, specify the device to calculate the default path costs of its ports by using IEEE 802.1d-1998, and set the path cost of GigabitEthernet 1/0/3 to 2000 on VLANs 20 through 30.

```
<Sysname> system-view
[Sysname] stp mode pvst
[Sysname] stp pathcost-standard dot1d-1998
[Sysname] interface gigabitethernet 1/0/3
[Sysname-GigabitEthernet1/0/3] stp vlan 20 to 30 cost 2000
```

Configuring the port priority

When the priority of a port changes, MSTP re-calculates the role of the port and initiates a state transition. 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 a spanning tree device, a port can have different priorities and play different roles in different spanning trees, 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.

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view: interface interface-type interface-number Enter port group view: port-group manual port-group-name 	Use one of the commands.

Step	Command	Remarks
3. Configure the port priority.	<ul style="list-style-type: none"> In STP/RSTP mode: stp port priority <i>priority</i> In PVST mode: stp vlan <i>vlan-list</i> port priority <i>priority</i> In MSTP mode: stp [instance <i>instance-id</i>] port priority <i>priority</i> 	<p>Use one of the commands.</p> <p>The default setting is 128.</p>

Configuring the port link type

A point-to-point link directly connects two devices. If two root ports or designated ports are connected over a point-to-point link, they can rapidly transit to the forwarding state after a proposal-agreement handshake process.

Configuration restrictions and guidelines

- You can configure the link type as point-to-point for a Layer 2 aggregate interface or a port that operates in full duplex mode. HP recommends you to use the default setting and let the device to automatically detect the port link type.
- The **stp point-to-point force-false** or **stp point-to-point force-true** command configured on a port in MSTP or PVST mode takes effect on all MSTIs or VLANs.
- If the physical link to which the port connects is not a point-to-point link but you set it to be one, the configuration might bring a temporary loop.

Configuration procedure

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.
3. Configure the port link type.	stp point-to-point { auto force-false force-true }	By default, the link type is auto where the port automatically detects the link type.

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

- **legacy**—Compatible format

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

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

MSTP provides MSTP packet format incompatibility guard. In MSTP mode, if a port is configured to recognize/send MSTP packets in a mode other than **auto**, and if it receives a packet in a format different from the specified type, the port becomes a designated port and remains in the discarding state to prevent the occurrence of a loop.

MSTP provides MSTP packet format frequent change guard. If a port receives MSTP packets of different formats frequently, the MSTP packet format configuration contains errors. If the port is operating in MSTP mode, it will be shut down for protection. Ports disabled in this way can be re-activated after a detection interval. For more information about the detection interval, see *Fundamentals Configuration Guide*.

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> • Enter Layer 2 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> 	Use one of the commands.
3. Configure the mode that the port uses to recognize/send MSTP packets.	stp compliance { auto dot1s legacy }	auto by default.

Enabling outputting port state transition information

In a large-scale spanning tree network, you can enable devices to output the port state transition information of all MSTIs or the specified MSTI in order to monitor the port states in real time.

To enable outputting port state transition information:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable outputting port state transition information.	<ul style="list-style-type: none"> • In STP/RSTP mode: stp port-log instance 0 • In PVST mode: stp port-log vlan <i>vlan-list</i> • In MSTP mode: stp port-log instance { instance-id all } 	Use one of the commands. Enabled by default.

Enabling the spanning tree feature

You must enable the spanning tree feature for the device before any other spanning tree related configurations can take effect.

Configuration restrictions and guidelines

- You can disable the spanning tree feature for certain ports with the **undo stp enable** command to exclude them from spanning tree calculation and save CPU resources of the device.
- In PVST mode, when you globally enable the spanning tree feature the device automatically enables the spanning tree feature for the first n (which is the number of PVST instances that the switch supports and is 32 for the HP 5120 EI switch) of the existing VLANs by default. To enable the spanning tree feature for other VLANs, you must first disable the spanning tree feature for certain VLANs. This guideline does not apply if the number of existing VLANs on the switch does not exceed n .

Enabling the spanning tree feature (in STP/RSTP/MSTP mode)

In STP/RSTP/MSTP mode, make sure that the spanning tree feature is enabled globally and on the desired ports.

To enable the spanning tree feature in STP/RSTP/MSTP mode:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable the spanning tree feature globally.	stp enable	By default, the spanning tree feature is disabled globally.
3. Enter interface view or port group view.	<ul style="list-style-type: none">Enter Layer 2 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>	Use either command.
4. Enable the spanning tree feature for the port or group of ports.	stp enable	Optional. By default, the spanning tree feature is enabled for all ports.

Enabling the spanning tree feature (in PVST mode)

In PVST mode, make sure that the spanning tree feature is enabled globally and on the desired VLANs and ports.

To enable the spanning tree feature in PVST mode:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Globally enable the spanning tree feature.	stp enable	By default, the spanning tree feature is disabled globally.

Step	Command	Remarks
3. Enable the spanning tree feature on specific VLANs.	stp vlan <i>vlan-list</i> enable	By default, the spanning tree feature is enabled on VLANs.
4. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
5. Enable the spanning tree feature for the port or group of ports.	stp enable	Optional. By default, the spanning tree feature is enabled on all ports.

Performing mCheck

If a port on a device that is running MSTP, RSTP, or PVST connects to an STP device, this port automatically transits to the STP mode. However, it cannot automatically transit back to the original mode under the following circumstances:

- The STP device is shut down or removed.
- The STP device transits to the MSTP, RSTP, or PVST mode.

Suppose Device A running STP, Device B with no spanning tree feature enabled, and Device C running MSTP, RSTP, or PVST are connected in order. Device B will transparently transmit the STP BPDUs, and the port on Device C and connecting to Device B will transit to the STP mode. After you enable the spanning tree feature on Device B, to run MSTP, RSTP, or PVST between Device B and Device C, you must perform an mCheck operation on the ports interconnecting Device B and Device C, in addition to configuring the spanning tree to operate in MSTP, RSTP, or PVST mode on Device B.

To forcibly transit the port to operate in the original mode, you can perform an mCheck operation.

The following methods for performing mCheck produce the same result.

Performing mCheck globally

Step	Command
1. Enter system view.	system-view
2. Perform mCheck.	stp mcheck

Performing mCheck in interface view

Step	Command
1. Enter system view.	system-view
2. Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>

Step	Command
3. Perform mCheck.	stp mcheck

NOTE:

An mCheck operation takes effect on a device that operates in MSTP, RSTP, or PVST mode.

Configuring Digest Snooping

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

Spanning tree implementations vary with vendors, and the configuration digests calculated using private keys is different, so devices of different vendors in the same MST region cannot communicate with each other.

To enable communication between an HP device and a third-party device, enable the Digest Snooping feature on the port that connects the HP device to the third-party device in the same MST region.

Configuration restrictions and guidelines

- Before you enable Digest Snooping, make sure that associated devices of different vendors are connected and run spanning tree protocols.
- With digest snooping enabled, in-the-same-region verification does not require comparison of configuration digest, so the VLAN-to-instance mappings must be the same on associated ports.
- With global Digest Snooping enabled, modification of VLAN-to-instance mappings and removal of the current region configuration by using the **undo stp region-configuration** command are not allowed. You can modify only the region name and revision level.
- To make Digest Snooping take effect, you must enable it both globally and on associated ports. To make the configuration effective on all configured ports and while reducing impact on the network, enable Digest Snooping on all associated ports first and then globally.
- To prevent loops, do not enable Digest Snooping on MST region edge ports.
- HP recommends you to enable Digest Snooping first and then the spanning tree feature. To avoid causing traffic interruption, do not configure Digest Snooping when the network is already working well.

Configuration procedure

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

To configure Digest Snooping:

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

Step	Command	Remarks
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
3. Enable Digest Snooping on the interface or port group.	stp config-digest-snooping	Disabled by default.
4. Return to system view.	quit	N/A
5. Enable global Digest Snooping.	stp config-digest-snooping	Disabled by default.


Digest Snooping configuration example

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 the ports of Device A and Device B that connect to Device C, so that the three devices can communicate with one another.

Figure 23 Digest Snooping configuration



Configuration procedure

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

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

```

[DeviceA] stp config-digest-snooping
# Enable Digest Snooping on GigabitEthernet 1/0/1 of Device B and enable global Digest
Snooping on Device B.
<DeviceB> system-view
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] stp config-digest-snooping
[DeviceB-GigabitEthernet1/0/1] quit
[DeviceB] stp config-digest-snooping

```

Configuring No Agreement Check

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

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

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

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

Figure 24 Rapid state transition of an MSTP designated port


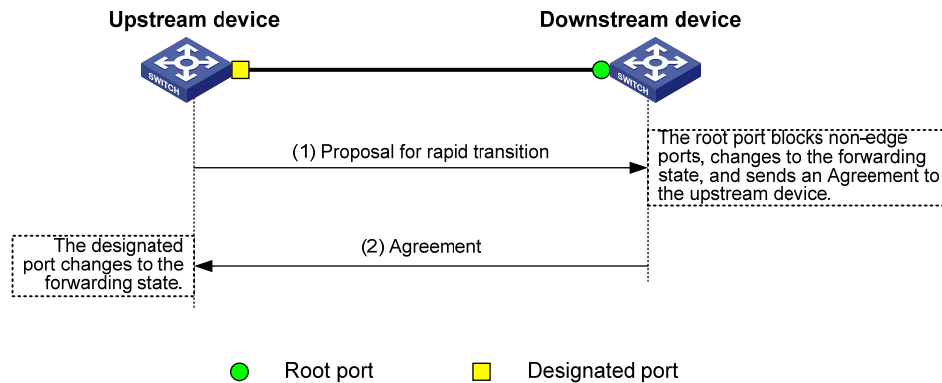


Figure 25 Rapid state transition of an RSTP designated port



If the upstream device is a third-party device, the rapid state transition implementation might be limited. For example, when the upstream device uses a rapid transition mechanism similar to that of RSTP, and the downstream device adopts MSTP and does not operate 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 that supports spanning tree protocols by using 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.

Configuration procedure

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

To configure No Agreement Check:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> • Enter Layer 2 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> 	Use either command.
3. Enable No Agreement Check.	stp no-agreement-check	Disabled by default.


No Agreement Check configuration example

Network requirements

As shown in Figure 26:

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

Figure 26 Network diagram



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 topology change (TC) snooping application scenario. Device A and Device B form an IRF fabric and do not have the spanning tree feature enabled. The IRF fabric connects to two user networks, in which all devices are enabled with the spanning tree feature. The user networks are dual-uplinked to the IRF fabric for high availability. The IRF fabric transparently transmits BPDUs in every user network.

Figure 27 TC snooping application scenario



In the network, the IRF fabric transparently transmits the received BPDUs and does not participate in spanning tree calculations. When a topology change occurs to the IRF fabric or user networks, the IRF fabric 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 fabric.

With TC snooping enabled, a device actively updates the MAC address table entries and ARP entries upon receiving TC-BPDUs, so that the device can normally forward the user traffic.

For more information about MAC address table entries, see "Configuring the MAC address table."

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

Configuration restrictions and guidelines

- TC snooping and STP are mutually exclusive. You must globally disable the spanning tree feature before enable TC snooping.
- TC snooping does not take effect on the ports on which BPDU tunneling is enabled for spanning tree protocols. For more information about BPDU tunneling, see "Configuring BPDU tunneling."
- TC snooping does not support PVST TC-BPDUs. As a result, TC snooping does not take effect on a PVST network.

Configuration procedure

To configure TC snooping:

Step	Command	Description
1. Enter system view.	system-view	N/A
2. Globally disable the spanning tree feature.	undo stp enable	By default, the spanning tree feature is disabled globally.
3. Enable TC snooping.	stp tc-snooping	Disabled by default.

Configuring protection functions

A spanning tree device supports the following protection functions:

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

Configuration prerequisites

The spanning tree feature has been correctly configured on the device.

Enabling BPDU guard

For access layer devices, the access ports can directly connect to the user terminals (such as PCs) or file servers. The access ports are configured as edge ports to allow rapid transition. When these ports

receive configuration BPDUs, the system automatically sets the ports as non-edge ports and starts a new spanning tree calculation process. This causes a change of network topology. Under normal conditions, these ports should not receive configuration BPDUs. However, if someone forges configuration BPDUs maliciously to attack the devices, the network will become unstable.

The spanning tree protocol 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, the system closes these ports and notifies the NMS that these ports have been closed by the spanning tree protocol. The device will reactivate the closed ports after a detection interval. For more information about this detection interval, see *Fundamentals Configuration Guide*.

Configure BPDU guard on a device with edge ports configured.

To enable BPDU guard:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable the BPDU guard function for the device.	stp bpdu-protection	Disabled by default.

NOTE:

BPDU guard does not take effect on loopback-testing-enabled ports. For more information about loopback testing, see "Configuring Ethernet interfaces."

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, due to possible configuration errors or malicious attacks in the network, the legal root bridge might receive a configuration BPDU with a higher priority. Another device will supersede the current legal root bridge, causing an undesired change of the network topology. The traffic that should go over high-speed links is switched to low-speed links, resulting in network congestion.

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

Configure root guard on a designated port.

To enable root guard:

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

Step	Command	Remarks
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
3. Enable the root guard function for the ports.	stp root-protection	Disabled by default.

NOTE:

You cannot configure root guard and loop guard on a port at the same time.

Enabling loop guard

A device that keeps receiving BPDUs from the upstream device can maintain the state of the root port and blocked ports. However, link congestion or unidirectional link failures might cause these ports to 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 transit 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, it transits its state. Otherwise, it stays in the discarding state to prevent temporary loops.

Configure loop guard on the root port and alternate ports of a device.

To enable loop guard:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
3. Enable the loop guard function for the ports.	stp loop-protection	Disabled by default.

NOTE:

- Do not enable loop guard on a port that connects user terminals. Otherwise, the port will stay in the discarding state in all MSTIs because it cannot receive BPDUs.
- You cannot configure edge port settings and loop guard, or configure root guard and loop guard on a port at the same time.

Enabling TC-BPDU guard

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

With the TC-BPDU guard function, you can set the maximum number of immediate forwarding address entry flushes that the device can perform every a specified period of time (10 seconds). For TC-BPDUs received in excess of the limit, the device performs a forwarding address entry flush when the time period expires. This prevents frequent flushing of forwarding address entries.

To enable TC-BPDU guard:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
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 every 10 seconds.	stp tc-protection threshold <i>number</i>	Optional. 6 by default.

NOTE:

HP does not recommend you disable this feature.

Enabling BPDU drop

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

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

To enable BPDU drop on an Ethernet interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Enable BPDU drop on the current interface.	bpdu-drop any	Disabled by default.

NOTE:

Because a port with BPDU drop enabled also drops the received 802.1X packets, do not enable BPDU drop and 802.1X on a port at the same time. For more information about 802.1X, see *Security Configuration Guide*.

Displaying and maintaining the spanning tree

Task	Command	Remarks
Display information about ports blocked by spanning tree protection functions.	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 shut down by spanning tree 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> vlan <i>vlan-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> vlan <i>vlan-id</i>] tc [slot <i>slot-number</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the spanning tree status and statistics.	display stp [instance <i>instance-id</i> vlan <i>vlan-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 spanning tree statistics.	reset stp [interface <i>interface-list</i>]	Available in user view

Spanning tree configuration examples

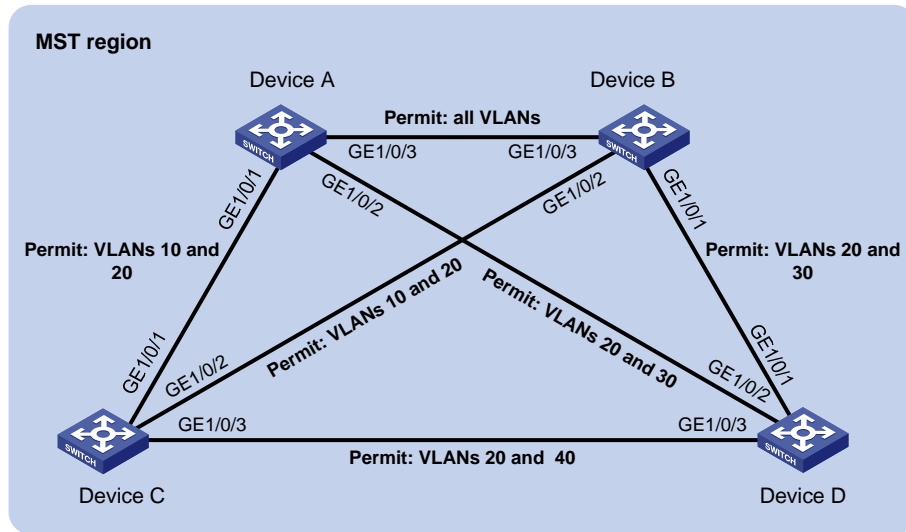
MSTP configuration example

Network requirements

As shown in [Figure 28](#):

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

Figure 28 Network diagram



Configuration procedure

- Configure VLANs and VLAN member ports (details not shown):
 - Create VLAN 10, VLAN 20, and VLAN 30 on Device A and Device B.
 - Create VLAN 10, VLAN 20, and VLAN 40 on Device C.
 - 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.
- Configure Device A:

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

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

Activate MST region configuration.

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

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

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

Enable the spanning tree feature globally.

```
[DeviceA] stp enable
```
- Configure Device B:

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

```
<DeviceB> system-view
```

```
[DeviceB] stp region-configuration
[DeviceB-mst-region] region-name example
[DeviceB-mst-region] instance 1 vlan 10
[DeviceB-mst-region] instance 3 vlan 30
[DeviceB-mst-region] instance 4 vlan 40
[DeviceB-mst-region] revision-level 0
# Activate MST region configuration.
[DeviceB-mst-region] active region-configuration
[DeviceB-mst-region] quit
# Specify the current device as the root bridge of MSTI 3.
[DeviceB] stp instance 3 root primary
# Enable the spanning tree feature globally.
[DeviceB] stp enable
```

4. Configure Device C:

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

```
<DeviceC> system-view
[DeviceC] stp region-configuration
[DeviceC-mst-region] region-name example
[DeviceC-mst-region] instance 1 vlan 10
[DeviceC-mst-region] instance 3 vlan 30
[DeviceC-mst-region] instance 4 vlan 40
[DeviceC-mst-region] revision-level 0
# Activate MST region configuration.
[DeviceC-mst-region] active region-configuration
[DeviceC-mst-region] quit
# Specify the current device as the root bridge of MSTI 4.
[DeviceC] stp instance 4 root primary
# Enable the spanning tree feature globally.
[DeviceC] stp enable
```

5. Configure 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 the spanning tree feature globally.
```

```
[DeviceD] stp enable
```

6. Verify the configurations:

In this example, suppose that Device B has the lowest root bridge ID. As a result, Device B is elected as the root bridge in MSTI 0.

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

Display brief spanning tree information on Device A.

```
[DeviceA] display stp brief
```

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet1/0/1	ALTE	DISCARDING	NONE
0	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
0	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE
1	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
1	GigabitEthernet1/0/3	DESI	FORWARDING	NONE
3	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
3	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE

Display brief spanning tree information on Device B.

```
[DeviceB] display stp brief
```

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
0	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
0	GigabitEthernet1/0/3	DESI	FORWARDING	NONE
1	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
1	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE
3	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
3	GigabitEthernet1/0/3	DESI	FORWARDING	NONE

Display brief spanning tree information on Device C.

```
[DeviceC] display stp brief
```

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
0	GigabitEthernet1/0/2	ROOT	FORWARDING	NONE
0	GigabitEthernet1/0/3	DESI	FORWARDING	NONE
1	GigabitEthernet1/0/1	ROOT	FORWARDING	NONE
1	GigabitEthernet1/0/2	ALTE	DISCARDING	NONE
4	GigabitEthernet1/0/3	DESI	FORWARDING	NONE


Display brief spanning tree information on Device D.

```
[DeviceD] display stp brief
```

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet1/0/1	ROOT	FORWARDING	NONE
0	GigabitEthernet1/0/2	ALTE	DISCARDING	NONE
0	GigabitEthernet1/0/3	ALTE	DISCARDING	NONE
3	GigabitEthernet1/0/1	ROOT	FORWARDING	NONE
3	GigabitEthernet1/0/2	ALTE	DISCARDING	NONE
4	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE

Based on the output, you can draw the MSTI mapped to each VLAN, as shown in [Figure 29](#).

Figure 29 MSTIs mapped to different VLANs



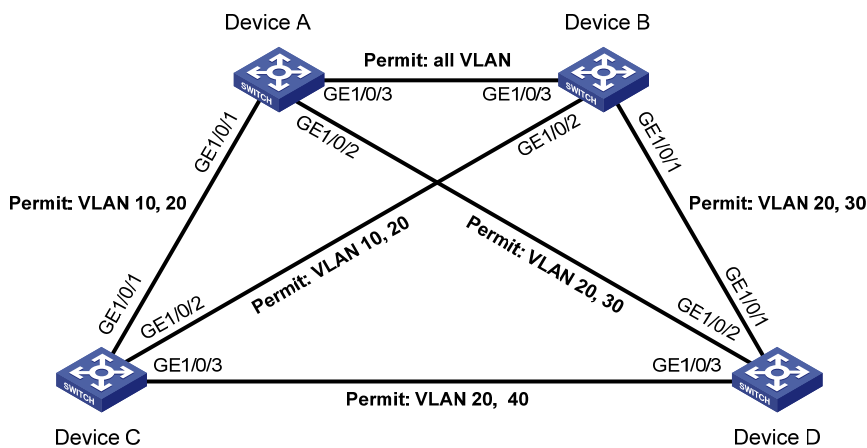
PVST configuration example

Network requirements

As shown in [Figure 30](#):

- Device A and Device B work at the distribution layer. Device C and Device D work at the access layer.
- Configure PVST so that packets of different VLANs are forwarded along different spanning trees.
- VLAN 10, VLAN 20, and VLAN 30 are terminated on the distribution layer devices, and VLAN 40 is terminated on the access layer devices. The root bridge of VLAN 10 and VLAN 20 is Device A, that of VLAN 30 is Device B, and that of VLAN 40 is Device C.

Figure 30 Network diagram



Configuration procedure

1. Configure VLANs and VLAN member ports (details not shown):
 - Create VLAN 10, VLAN 20, and VLAN 30 on Device A and Device B.
 - Create VLAN 10, VLAN 20, and VLAN 40 on Device C.
 - Create VLAN 20, VLAN 30, and VLAN 40 on Device D.
 - Configure the ports on these devices as trunk ports and assign them to related VLANs.
2. Configure Device A:
 - # Set the spanning tree mode to PVST.

```
<DeviceA> system-view
[DeviceA] stp mode pvst
```

 - # Specify the device as the root bridge of VLAN 10 and VLAN 20.

```
[DeviceA] stp vlan 10 20 root primary
```

 - # Enable the spanning tree feature globally and for VLANs 10, 20, and 30.

```
[DeviceA] stp enable
[DeviceA] stp vlan 10 20 30 enable
```
3. Configure Device B:
 - # Set the spanning tree mode to PVST.

```
<DeviceB> system-view
[DeviceB] stp mode pvst
```

 - # Specify the device as the root bridge of VLAN 30.

```
[DeviceB] stp vlan 30 root primary
```

 - # Enable the spanning tree feature globally and for VLANs 10, 20, and 30.

```
[DeviceB] stp enable
[DeviceB] stp vlan 10 20 30 enable
```
4. Configure Device C:
 - # Set the spanning tree mode to PVST.

```
<DeviceC> system-view
[DeviceC] stp mode pvst
```

 - # Specify the current device as the root bridge of VLAN 40.

```
[DeviceC] stp vlan 40 root primary
```

 - # Enable the spanning tree feature globally and for VLANs 10, 20, and 40.

```
[DeviceC] stp enable
[DeviceC] stp vlan 10 20 40 enable
```
5. Configure Device D:
 - # Set the spanning tree mode to PVST.

```
<DeviceD> system-view
[DeviceD] stp mode pvst
```

 - # Enable the spanning tree feature globally and for VLANs 20, 30, and 40.

```
[DeviceD] stp enable
[DeviceD] stp vlan 20 30 40 enable
```
6. Verify the configurations:

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

Display brief spanning tree information on Device A.

[DeviceA] display stp brief

VLAN	Port	Role	STP State	Protection
10	GigabitEthernet1/0/1	DESI	DISCARDING	NONE
10	GigabitEthernet1/0/3	DESI	FORWARDING	NONE
20	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
20	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
20	GigabitEthernet1/0/3	DESI	FORWARDING	NONE
30	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
30	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE

Display brief spanning tree information on Device B.

[DeviceB] display stp brief

VLAN	Port	Role	STP State	Protection
10	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
10	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE
20	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
20	GigabitEthernet1/0/2	DESI	FORWARDING	NONE
20	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE
30	GigabitEthernet1/0/1	DESI	FORWARDING	NONE
30	GigabitEthernet1/0/3	DESI	FORWARDING	NONE

Display brief spanning tree information on Device C.

[DeviceC] display stp brief

VLAN	Port	Role	STP State	Protection
10	GigabitEthernet1/0/1	ROOT	FORWARDING	NONE
10	GigabitEthernet1/0/2	ALTE	FORWARDING	NONE
20	GigabitEthernet1/0/1	ROOT	FORWARDING	NONE
20	GigabitEthernet1/0/2	ALTE	FORWARDING	NONE
20	GigabitEthernet1/0/3	DESI	DISCARDING	NONE
40	GigabitEthernet1/0/3	DESI	FORWARDING	NONE

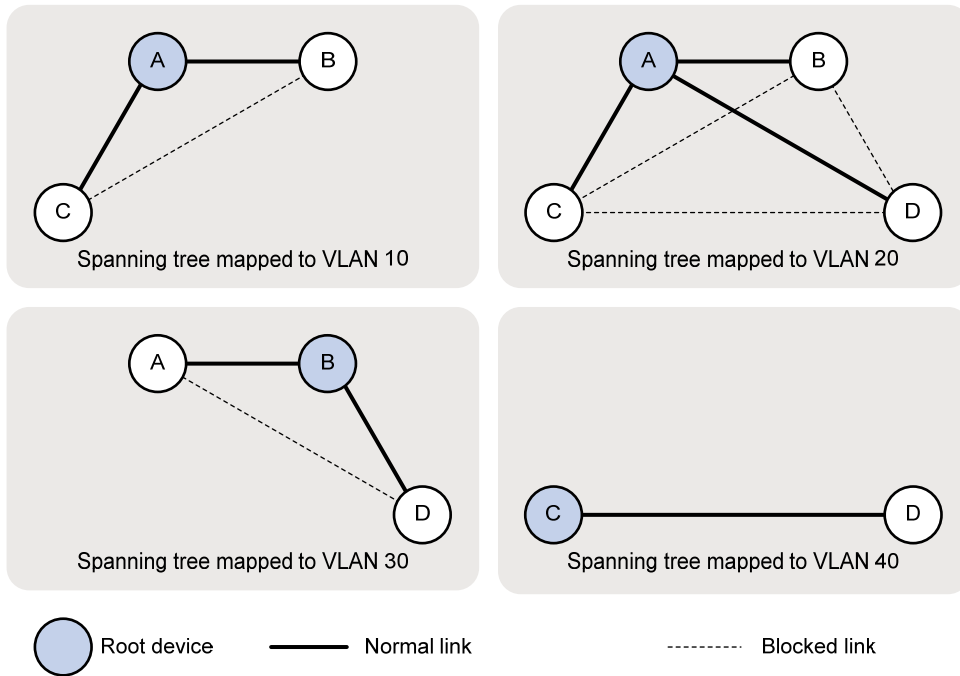
Display brief spanning tree information on Device D.

[DeviceD] display stp brief

VLAN	Port	Role	STP State	Protection
20	GigabitEthernet1/0/1	ALTE	FORWARDING	NONE
20	GigabitEthernet1/0/2	ROOT	DISCARDING	NONE
20	GigabitEthernet1/0/3	ALTE	DISCARDING	NONE
30	GigabitEthernet1/0/1	ROOT	FORWARDING	NONE
30	GigabitEthernet1/0/2	ALTE	DISCARDING	NONE
40	GigabitEthernet1/0/3	ROOT	FORWARDING	NONE

Based on the output, you can draw the spanning tree mapped to each VLAN, as shown in [Figure 31](#).

Figure 31 Spanning trees mapped to different VLANs



Configuring BPDU tunneling


Overview

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

Background

Dedicated lines are used in a service provider network to build user-specific Layer 2 networks. As a result, a user network consists of parts located at different sides of the service provider network. As shown in Figure 32, the devices for User A are CE 1 and CE 2, both of which belong to VLAN 100. User A's network is divided into network 1 and network 2, which are connected by the service provider network. When a Layer 2 protocol (for example, STP) runs on both network 1 and network 2, the Layer 2 protocol packets must be transmitted over the service provider network to implement Layer 2 protocol calculation (for example, spanning tree calculation). When receiving a Layer 2 protocol packet, the PEs cannot determine whether the packet is from the user network or the service provider network, and must deliver the packet to the CPU for processing. In this case, the Layer 2 protocol calculation in User A's network is mixed with that in the service provider network, and the user network cannot implement independent Layer 2 protocol calculation.

Figure 32 BPDU tunneling application scenario



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

1. After receiving a Layer 2 protocol packet from CE 1, PE 1 encapsulates the packet, replaces its destination MAC address with a specific multicast MAC address, and forwards the packet to the service provider network.
2. The encapsulated Layer 2 protocol packet (called bridge protocol data unit, BPDU) 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 CE 2.

HP devices support BPDU tunneling for the following protocols:

- Cisco Discovery Protocol (CDP)
- Device Link Detection Protocol (DLDP)

- Ethernet Operation, Administration and Maintenance (EOAM)
- GARP VLAN Registration Protocol (GVRP)
- HW Group Management Protocol (HGMP)
- Link Aggregation Control Protocol (LACP)
- Link Layer Discovery Protocol (LLDP)
- Port Aggregation Protocol (PAGP)
- Per VLAN Spanning Tree (PVST)
- Spanning Tree Protocol (STP)
- Unidirectional Link Direction (UDLD)
- VLAN Trunking Protocol (VTP)

BPDU tunneling implementation

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

This document uses the term *STP* in a broad sense. It includes STP, RSTP, and MSTP.

STP calculates the topology of a network by transmitting BPDUs among devices in the network. For more information, see "Configuring spanning tree protocols."


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

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

BPDU tunneling delivers the following benefits:

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

Figure 33 BPDU tunneling implementation



The upper section of [Figure 33](#) represents the service provider network (ISP network). The lower section, including User A network 1 and User A network 2, represents the customer networks. Enabling BPDU tunneling on 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 through the service provider network. This ensures consistent spanning tree calculation throughout User A network, without affecting the spanning tree calculation of the service provider network.

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

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

NOTE:

Through configuration, make sure 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.

Enabling BPDU tunneling

Configuration prerequisites

Before configuring BPDU tunneling for a protocol, perform the following tasks:

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

Configuration restrictions and guidelines

- Settings made in Layer 2 Ethernet interface view or Layer 2 aggregate interface view take effect only on the current port. Settings made in port group view take effect on all ports in the port group.
- Before you enable BPDU tunneling for DLDAP, EOAM, GVRP, HGMP, LLDP, or STP on a port, disable the protocol on the port first.
- Because PVST is a special STP protocol, you must do two things before you enable BPDU tunneling for PVST on a port: first, disable STP; second, enable BPDU tunneling for STP on the port.
- Do not enable BPDU tunneling for DLDAP, EOAM, LACP, LLDP, PAGP, or UDLD on the member port of a Layer 2 aggregation group.

Configuration procedure

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

Enabling BPDU tunneling for a protocol in Layer 2 Ethernet interface view or port group view

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">• Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>• Enter port group view: port-group manual <i>port-group-name</i>	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 }	Disabled by default.

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 aggregate interface view.	interface bridge-aggregation <i>interface-number</i>	N/A
3. Enable BPDU tunneling for a protocol on the Layer 2 aggregate interface.	bpdu-tunnel dot1q { cdp gvrp hgmp pvst stp vtp }	Disabled by default.

Configuring destination multicast MAC address for BPDUs

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

To configure destination multicast MAC address for BPDUs:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the destination multicast MAC address for BPDUs.	bpdu-tunnel tunnel-dmac <i>mac-address</i>	Optional. 0x010F-E200-0003 by default.

NOTE:

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

BPDU tunneling configuration examples

BPDU tunneling for STP configuration example


Network requirements

As shown in Figure 34:

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

Figure 34 Network diagram



Configuration procedure

1. Configure PE 1:

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

```
<PE1> system-view
```

```
[PE1] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

```
# Create VLAN 2 and assign GigabitEthernet 1/0/1 to VLAN 2.
[PE1] vlan 2
[PE1-vlan2] quit
[PE1] interface gigabitethernet 1/0/1
[PE1-GigabitEthernet1/0/1] port access vlan 2

# Disable STP on GigabitEthernet 1/0/1, and then enable BPDU tunneling for STP on it.
[PE1-GigabitEthernet1/0/1] undo stp enable
[PE1-GigabitEthernet1/0/1] bpdu-tunnel dot1q stp
```

2. Configure PE 2:

```
# Configure the destination multicast MAC address for BPDUs as 0x0100-0CCD-CDD0.
<PE2> system-view
[PE2] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0

# Create VLAN 2 and assign GigabitEthernet 1/0/2 to VLAN 2.
[PE2] vlan 2
[PE2-vlan2] quit
[PE2] interface gigabitethernet 1/0/2
[PE2-GigabitEthernet1/0/2] port access vlan 2

# Disable STP on GigabitEthernet 1/0/2, and then enable BPDU tunneling for STP on it.
[PE2-GigabitEthernet1/0/2] undo stp enable
[PE2-GigabitEthernet1/0/2] bpdu-tunnel dot1q stp
```

BPDU tunneling for PVST configuration example


Network requirements

As shown in [Figure 35](#):

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

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

Figure 35 Network diagram



Configuration procedure

1. Configure PE 1:

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

```
<PE1> system-view
```

```
[PE1] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

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

```
[PE1] interface gigabitethernet 1/0/1
```

```
[PE1-GigabitEthernet1/0/1] port link-type trunk
```

```
[PE1-GigabitEthernet1/0/1] port trunk permit vlan all
```

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

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

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

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

2. Configure PE 2:

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

```
<PE2> system-view
```

```
[PE2] bpdu-tunnel tunnel-dmac 0100-0ccd-cdd0
```

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

```
[PE2] interface gigabitethernet 1/0/2
```

```
[PE2-GigabitEthernet1/0/2] port link-type trunk
```

```
[PE2-GigabitEthernet1/0/2] port trunk permit vlan all
```

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

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

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


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

Configuring VLANs

Overview

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

Figure 36 A VLAN diagram



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

VLAN technology delivers the following benefits:

1. Confining broadcast traffic within individual VLANs. This reduces bandwidth waste and improves network performance.
2. 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.
3. Creating flexible virtual workgroups. Because users from the same workgroup can be assigned to the same VLAN regardless of their physical locations, network construction and maintenance are much easier and more flexible.

VLAN fundamentals


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

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

The Ethernet II encapsulation format is used here. Besides the Ethernet II encapsulation format, Ethernet also supports other encapsulation formats, including 802.2 LLC, 802.2 SNAP, and 802.3 raw. The VLAN tag fields are added to frames encapsulated in these formats for VLAN identification.


In the header of a traditional Ethernet data frame, the field after the destination MAC address and the source MAC address is the Type field, which indicates the upper layer protocol type, as shown in [Figure 37](#).

Figure 37 Traditional Ethernet frame format



IEEE 802.1Q inserts a four-byte VLAN tag after the DA&SA field, as shown in [Figure 38](#).

Figure 38 Position and format of VLAN tag



The fields of a VLAN tag are as follows:

- **TPID**—The 16-bit TPID field indicates whether a frame is VLAN-tagged. By default, the TPID value is 0x8100, which indicates that the frame is VLAN-tagged. Devices vendors can set the TPID to different values. For compatibility with these devices, modify the TPID value so that frames carry a TPID value identical to the value of a particular vendor, allowing interoperability with devices from that vendor. The device determines whether a received frame carries a VLAN tag by checking the TPID value. When the TPID value of a frame is the configured value or 0x8100, the frame is considered as a VLAN-tagged frame. For information about commands used to modify TPID values, see Layer 2—LAN Switching Command Reference.
- **Priority**—The 3-bit priority field indicates the 802.1p priority of the frame.
- **CFI**—The 1-bit CFI field indicates whether the MAC addresses are encapsulated in standard format when packets are transmitted across different media. A value of 0 indicates that MAC addresses are encapsulated in standard format. A value of 1 indicates that MAC addresses are encapsulated in a non-standard format. The value of this field is 0 by default.
- **VLAN ID**—The 12-bit VLAN ID field identifies the VLAN that the frame belongs to. The VLAN ID range is 0 to 4095. Because 0 and 4095 are reserved, a VLAN ID actually ranges from 1 to 4094.

A network device handles an incoming frame depending on whether the frame is VLAN tagged, and information about the VLAN tags, if any. For more information, see "[Introduction to port-based VLAN](#)."

NOTE:

When a frame carrying multiple VLAN tags passes through, the switch processes the frame according to its outer VLAN tag, and transmits the inner tags as payload.

VLAN types

You can implement VLANs based on the following criteria:

- Port
- MAC address
- Protocol

- IP subnet
- Policy
- Other criteria

This chapter covers port-based VLAN, MAC-based VLAN, protocol-based VLAN, and IP subnet-based VLAN. The port-based VLAN implementation is the basis of all other VLAN implementations. To use any other VLAN implementations, you must configure port-based VLAN settings.

You can configure all these types of VLANs on a port at the same time. When the switch is determining which VLAN a packet that passes through the port should be assigned to, it looks up the VLANs in the default order of MAC-based VLAN, IP sub-based VLAN, protocol-based VLAN, and port-based VLAN.

Protocols and standards

IEEE 802.1Q, *IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks*

Configuring basic VLAN settings

Configuration restrictions and guidelines

- As the default VLAN, VLAN 1 cannot be created or removed.
- You cannot manually create or remove VLANs reserved for special purposes.
- To delete a protocol reserved VLAN, voice VLAN, management VLAN, dynamic VLAN, VLAN with a QoS policy applied, control VLAN for a smart link group, control VLAN for an RRPP domain, remote probe VLAN for remote port mirroring, remove the configuration from the VLAN first, and execute the **undo vlan** command.

Configuration procedure

To configure basic VLAN settings:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a VLAN and enter its view, or create VLANs in batches.	vlan { <i>vlan-id1</i> [to <i>vlan-id2</i>] all }	Optional. By default, only the default VLAN (VLAN 1) exists in the system.
3. Enter VLAN view.	vlan <i>vlan-id</i>	Required only when you create VLANs in batches.
4. Configure a name for the VLAN.	name <i>text</i>	Optional. By default, the name of a VLAN is its VLAN ID (VLAN 0001 , for example).
5. Configure a description for the VLAN.	description <i>text</i>	Optional. The default description is VLAN <i>vlan-id</i> , which is the ID of the VLAN. For example, the description of VLAN 100 is VLAN 0100 by default.

Configuring basic settings of a VLAN interface

You can use VLAN interfaces to provide Layer 3 communication between hosts of different VLANs. VLAN interfaces are virtual interfaces used for Layer 3 communication between different VLANs. They do not exist as physical entities on devices. For each VLAN, you can create one VLAN interface. You can assign the VLAN interface an IP address and specify the IP address as the gateway address for the devices in the VLAN, so that traffic can be routed to other IP subnets.

Configuration procedure

To configure basic settings of a VLAN interface:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a VLAN interface and enter VLAN interface view.	interface vlan-interface <i>vlan-interface-id</i>	If the VLAN interface already exists, you enter its view directly.
3. Assign an IP address to the VLAN interface.	ip address <i>ip-address</i> { <i>mask</i> <i>mask-length</i> } [<i>sub</i>]	Optional. By default, no IP address is assigned to any VLAN interface.
4. Configure a description for the VLAN interface.	description <i>text</i>	Optional. By default, the description of a VLAN is the VLAN interface name. For example, Vlan-interface1 Interface .
5. Set the MTU for the VLAN interface.	mtu <i>size</i>	Optional. By default, the MTU is 1500 bytes.
6. Restore the default settings for the VLAN interface.	default	Optional.
7. Cancel the action of manually shutting down the VLAN interface.	undo shutdown	Optional. By default, a VLAN interface is not manually shut down. The VLAN interface is up if one or more ports in the VLAN is up, and goes down if all ports in the VLAN go down. A VLAN interface shut down with the shutdown command is in the DOWN (Administratively) state until you bring it up, regardless of how the state of the ports in the VLAN changes.

NOTE:

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


VLAN interface configuration example

Network requirements

As shown in [Figure 39](#), PC A is assigned to VLAN 5. PC B is assigned to VLAN 10. The PCs belong to different IP subnets and cannot communicate with each other.

Configure VLAN interfaces on Switch A and configure the PCs to enable Layer 3 communication between the PCs.

Figure 39 Network diagram



Configuration procedure

1. Configure Switch A:
 - # Create VLAN 5 and assign GigabitEthernet 1/0/1 to it.

```
<SwitchA> system-view
[SwitchA] vlan 5
[SwitchA-vlan5] port GigabitEthernet 1/0/1
```
 - # Create VLAN 10 and assign GigabitEthernet 1/0/2 to it.

```
[SwitchA-vlan5] vlan 10
[SwitchA-vlan10] port GigabitEthernet 1/0/2
[SwitchA-vlan10] quit
```
 - # Create VLAN-interface 5 and configure its IP address as 192.168.0.10/24.

```
[SwitchA] interface vlan-interface 5
[SwitchA-Vlan-interface5] ip address 192.168.0.10 24
[SwitchA-Vlan-interface5] quit
```
 - # Create VLAN-interface 10 and configure its IP address as 192.168.1.20/24.

```
[SwitchA] interface vlan-interface 10
[SwitchA-Vlan-interface10] ip address 192.168.1.20 24
[SwitchA-Vlan-interface10] return
```
2. Configure the default gateway of PC A as 192.168.0.10.
3. Configure the default gateway of PC B as 192.168.1.20.

Verifying the configurations

1. The PCs can ping each other.
2. Display brief information about Layer 3 interfaces on Switch A to verify the configuration.

```
<SwitchA> display ip interface brief
```

```
*down: administratively down
```

```
(s): spoofing
```

Interface	Physical	Protocol	IP Address	Description
Vlan-interface5	up	up	192.168.0.10	Vlan-inte...
Vlan-interface10	up	up	192.168.1.20	Vlan-inte...

Configuring port-based VLANs

Introduction to port-based VLAN

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

Port link type

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

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

PVID

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

When you configure the PVID on a port, use the following guidelines:

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

When you configure a PVID, follow these guidelines:

- Do not set the voice VLAN as the PVID of a port in automatic voice VLAN assignment mode. For information about voice VLAN, see "Configuring a voice VLAN."
- HP recommends that you set the same PVID ID for local and remote ports.
- Make sure that a port permits the traffic from its PVID to pass through. Otherwise, when the port receives frames tagged with the PVID or untagged frames, the port drops 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	

Port type	Actions (in the inbound direction)		Actions (in the outbound direction)
	Untagged frame	Tagged frame	
Access	Tags the frame with the PVID tag.	<ul style="list-style-type: none"> Receives the frame if its VLAN ID is the same as the PVID. Drops the frame if its VLAN ID is different from the PVID. 	Removes the VLAN tag and sends the frame.
Trunk	Checks whether the PVID is permitted on the port: <ul style="list-style-type: none"> If yes, tags the frame with the PVID tag. If not, drops the frame. 	<ul style="list-style-type: none"> Receives the frame if its VLAN is carried on the port. Drops the frame if its VLAN is not carried on the port. 	<ul style="list-style-type: none"> Removes the tag and send the frame if the frame carries the PVID tag and the port belongs to the PVID. Sends the frame without removing the tag if its VLAN is carried on the port but is different from the PVID.
Hybrid			Sends the frame if its VLAN is carried on the port. The frame is sent with the VLAN tag removed or intact depending on your configuration via the port hybrid vlan command. This is true of the PVID.

Assigning an access port to a VLAN

You can assign an access port to a VLAN in VLAN view, interface view (including Layer 2 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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter VLAN view.	vlan <i>vlan-id</i>	If the specified VLAN does not exist, this command creates the VLAN first.
3. Assign one or a group of access ports to the VLAN.	port <i>interface-list</i>	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:

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

Step	Command	Remarks
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> 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> 	<p>Use any command.</p> <ul style="list-style-type: none"> The configuration made in Layer 2 Ethernet interface view applies only to the port. The configuration made in port group view applies to all ports in the port group. The configuration made in Layer 2 aggregate interface view applies to the aggregate interface and its aggregation member ports. If the system fails to apply the configuration to the aggregate interface, it stops applying the configuration to aggregation member ports. If the system fails to apply the configuration to an aggregation member port, it skips the port and moves to the next member port.
3. Configure the link type of the ports as access.	port link-type access	Optional. By default, all ports are access ports.
4. Assign the access ports to a VLAN.	port access vlan <i>vlan-id</i>	Optional. By default, all access ports belong to VLAN 1.

NOTE:

- Before you assign an access port to a VLAN, create the VLAN.
- In VLAN view, you can assign only Layer 2 Ethernet interfaces to the VLAN.

Assigning a trunk port to a VLAN

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

To assign a trunk port to one or multiple VLANs:

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

Step	Command	Remarks
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> 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> 	<p>Use any command.</p> <ul style="list-style-type: none"> The configuration made in Layer 2 Ethernet interface view applies only to the port. The configuration made in port group view applies to all ports in the port group. The configuration made in Layer 2 aggregate interface view applies to the aggregate interface and its aggregation member ports. If the system fails to apply the configuration to the aggregate interface, it stops applying the configuration to aggregation member ports. If the system fails to apply the configuration to an aggregation member port, it skips the port and moves to the next member port.
3. Configure the link type of the ports as trunk.	port link-type trunk	<p>By default, all ports are access ports.</p> <p>To change the link type of a port from trunk to hybrid or vice versa, you must set the link type to access first.</p>
4. Assign the trunk ports to the specified VLANs.	port trunk permit vlan { <i>vlan-list</i> all }	By default, a trunk port carries only VLAN 1.
5. Configure the PVID of the trunk ports.	port trunk pvid vlan <i>vlan-id</i>	<p>Optional.</p> <p>By default, the PVID is VLAN 1.</p>

NOTE:

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

Assigning a hybrid port to a VLAN

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

To assign a hybrid port to one or multiple VLANs:

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

Step	Command	Remarks
		Use any command.
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> 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> 	<ul style="list-style-type: none"> The configuration made in Ethernet interface view applies only to the port. The configuration made in port group view applies to all ports in the port group. The configuration made in Layer 2 aggregate interface view applies to the aggregate interface and its aggregation member ports. If the system fails to apply the configuration to the aggregate interface, it stops applying the configuration to aggregation member ports. If the system fails to apply the configuration to an aggregation member port, it skips the port and moves to the next member port.
3. Configure the link type of the ports as hybrid.	port link-type hybrid	By default, all ports are access ports. To change the link type of a port from trunk to hybrid or vice versa, you must set the link type to access first.
4. Assign the hybrid ports to the specified VLANs.	port hybrid vlan <i>vlan-list</i> { tagged untagged }	By default, a hybrid port allows only packets of VLAN 1 to pass through untagged.
5. Configure the PVID of the hybrid ports.	port hybrid pvid vlan <i>vlan-id</i>	Optional. By default, the PVID is VLAN 1.

NOTE:

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


Port-based VLAN configuration example

Network requirements

As shown in [Figure 40](#):

- 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.
- Make sure 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 40 Network diagram



Configuration procedure

1. Configure Device A:

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

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

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

```
[DeviceA] vlan 200
[DeviceA-vlan200] port gigabitethernet 1/0/2
[DeviceA-vlan200] quit
```

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

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

2. Configure Device B as you configure Device A.

3. Configure Host A and Host C to be on the same IP subnet, 192.168.100.0/24, for example. Configure Host B and Host D to be on the same IP subnet, 192.168.200.0/24, for example.

Verifying the configurations

- 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.
- Determine whether the configuration is successful by displaying relevant VLAN information.

Display information about VLANs 100 and 200 on Device A.

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

Configuring MAC-based VLANs

Introduction to MAC-based VLAN

The MAC-based VLAN feature assigns hosts to a VLAN based on their MAC addresses. This feature is usually used in conjunction with security technologies such as 802.1X to provide secure, flexible network access for terminal devices.

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 device processes received frames by using the following guidelines:

- When the port receives an untagged frame, the device looks up the MAC address-to-VLAN map based on the source MAC address of the frame for a match.
 - a. If the MAC address of a MAC address-to-VLAN entry matches the source MAC address of the untagged frame, the device tags the frame with the corresponding VLAN ID and forwards the frame.
 - b. If no match is found, the device assigns a VLAN to the frame by using other criteria, such as IP subnet or protocol, and forwards the frame.
 - c. If no VLAN is available, the device tags the frame with the PVID of the receiving port and forwards the frame.
- 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.

Dynamic MAC-based VLAN assignment

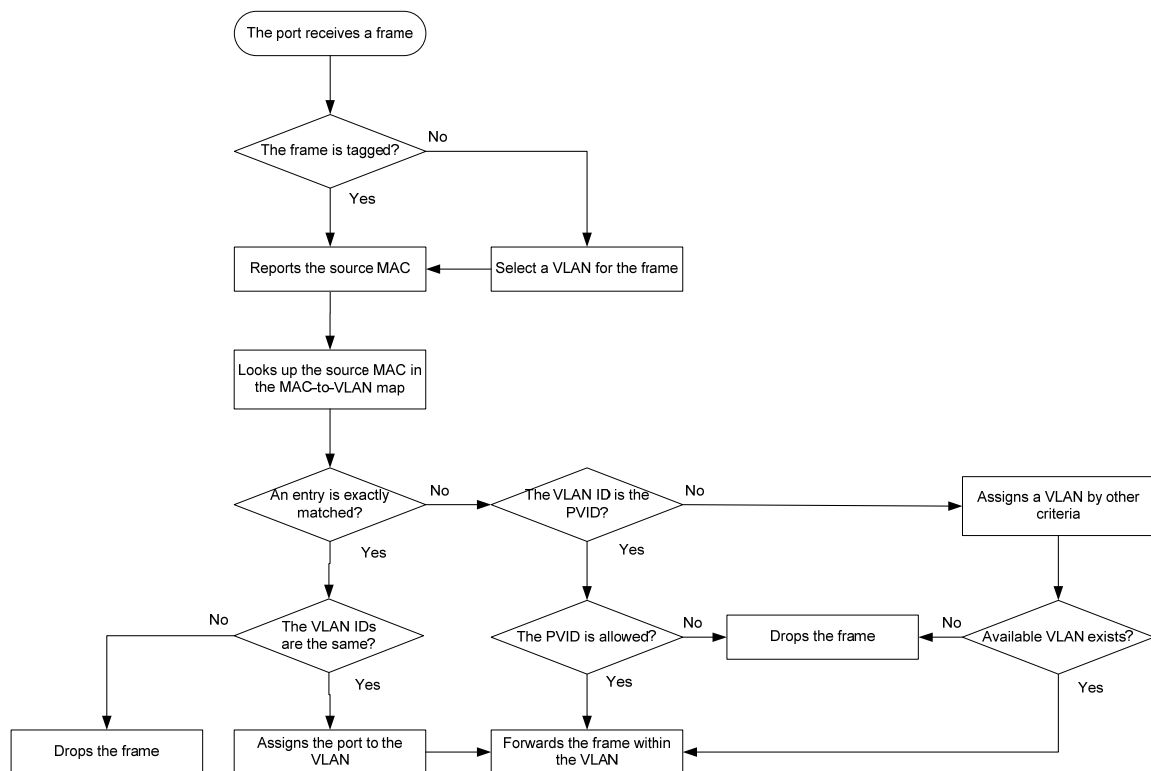
When you cannot determine the target MAC-based VLANs of a port, you can use dynamic MAC-based VLAN assignment on the port. To do that, you can create a MAC address-to-VLAN map containing multiple MAC address-to-VLAN entries, and enable the MAC-based VLAN feature and dynamic MAC-based VLAN assignment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

Dynamic MAC-based VLAN assignment uses the following workflows.

1. When the port receives a frame, the port first determines whether the frame is tagged.
- If yes, the port reports the source MAC address of the frame.

- If not, the port selects a VLAN for the frame in the order of MAC-based VLAN, IP subnet-based VLAN, protocol-based VLAN, and port-based VLAN, tags the untagged frame with the selected VLAN tag, and obtains the tag. Then, the port reports the source MAC address of the frame.
- 2. After reporting the source MAC address of the frame, the port looks up the source MAC address in the MAC-to-VLAN map, and processes the frame as follows:
 - If the source MAC address of the frame exactly matches a MAC address-to-VLAN entry configured on the port, the port checks whether the VLAN ID of the frame is the same as the VLAN in the MAC-to-VLAN entry.
 - a. If yes, the port dynamically joins the VLAN and forwards the frame.
 - b. If not, the port drops the frame.
 - If the source MAC address of the frame does not exactly match a MAC-to-VLAN entry, the port processes the frame depending on whether the VLAN ID of the frame is the PVID.
 - c. If yes, the port determines whether it allows PVID: if yes, the port forwards the frame in the PVID; if not, the port drops the frame.
 - d. If not, the port assigns a VLAN to the frame by using other criteria, such as IP subnet or protocol, and forwards the frame. If no VLAN is available, the port drops the frame.

Figure 41 Flowchart for processing a frame in dynamic MAC-based VLAN assignment



When you configure dynamic MAC-based VLAN assignment, follow these guidelines:

- When a port is assigned to the corresponding VLAN in a MAC address-to-VLAN entry, but has not been assigned to the VLAN by using the **port hybrid vlan** command, the port sends packets from the VLAN with VLAN tags removed.
- If you configure both static and dynamic MAC-based VLAN assignment on the same port, dynamic MAC-based VLAN assignment applies.

- A port forwards frames matching MAC-to-VLAN entries according to the 802.1p priorities of the MAC-based VLANs.

Dynamic MAC-based VLAN

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

When a user passes authentication of the access authentication server, the device 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 device automatically deletes the MAC address-to-VLAN entry, and removes the port from the MAC-based VLAN. For more information about 802.1X, MAC, and portal authentication, see *Security Configuration Guide*.

Configuration restrictions and guidelines

When you configure a MAC-based VLAN, follow these guidelines:

- MAC-based VLANs are available only on hybrid ports.
- 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.
- 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.
- With MSTP enabled, if a port is blocked in the MST instance (MSTI) of the target MAC-based VLAN, the port drops the received packets, instead of delivering them to the CPU. As a result, the receiving port will not be dynamically assigned to the corresponding VLAN. Do not configure dynamic MAC-based VLAN assignment together with MSTP, because the former is mainly configured on the access side.
- When PVST is enabled, if the VLAN to which a port is to be assigned is not allowed by the port, the port is blocked. In this case, the port drops received packets instead of delivering them to the CPU, failing to complete dynamic MAC-based VLAN assignment. Do not configure dynamic MAC-based VLAN assignment together with PVST, because the former is mainly configured on the access side.
- When you configure MAC-to-VLAN entries, if you specify the 802.1p priority for the VLAN of a MAC address, you must configure the **qos trust dot1p** command on the corresponding port, so that the port trusts the 802.1p priority of incoming packets and your configuration takes effect. For more information about the **qos trust dot1p** command, see *ACL and QoS Command Reference*.

Configuration procedure

To configure static MAC-based VLAN assignment:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Associate a specific MAC address with a VLAN.	mac-vlan mac-address <i>mac-address</i> vlan <i>vlan-id</i> [priority <i>priority</i>]	N/A
3. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>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 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	By default, all ports are access ports.
5. Configure the hybrid ports to permit packets from specific MAC-based VLANs to pass through.	port hybrid vlan <i>vlan-list</i> { tagged untagged }	By default, a hybrid port only permits the packets from VLAN 1 to pass through.
6. Enable the MAC-based VLAN feature.	mac-vlan enable	Disabled by default.
7. Configure VLAN matching precedence.	vlan precedence { mac-vlan ip-subnet-vlan }	Optional. By default, VLANs are preferably matched based on MAC addresses.

To configure dynamic MAC-based VLAN assignment:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Associate MAC addresses with a VLAN.	mac-vlan mac-address <i>mac-address</i> vlan <i>vlan-id</i> [priority <i>priority</i>]	With dynamic MAC-based VLAN assignment enabled, a port is automatically assigned to the VLAN in the MAC address-to-VLAN entry that is exactly matched by the source MAC address of the packet received on the port.
3. Enter Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
4. Configure the link type of the port as hybrid.	port link-type hybrid	By default, all ports are access ports.
5. Enable the MAC-based VLAN feature.	mac-vlan enable	Disabled by default.

Step	Command	Remarks
6. Enable dynamic MAC-based VLAN assignment.	mac-vlan trigger enable	<p>By default, dynamic MAC-based VLAN assignment is disabled.</p> <p>When you use the mac-vlan trigger enable command to enable dynamic MAC-based VLAN assignment, HP recommends that you configure the vlan precedence mac-vlan command, so that VLANs are assigned based on single MAC addresses preferentially. When dynamic MAC-based VLAN assignment is enabled, HP does not recommend configuring the vlan precedence ip-subnet-vlan command, which will make the system assign VLANs based on IP subnets, because the configuration does not take effect.</p>
7. Configure VLAN matching precedence.	vlan precedence mac-vlan	<p>Optional.</p> <p>By default, VLANs are preferentially matched based on MAC addresses.</p>
8. Disable the PVID of the port from forwarding packets with unknown source MAC addresses that do not match any MAC address-to-VLAN entry.	port pvid disable	<p>Optional.</p> <p>By default, when a port receives a packet with an unknown source MAC address that does not match to any MAC address-to-VLAN entry, it forwards the packet in its PVID.</p>

To configure dynamic MAC-based VLAN:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type interface-number</i> Enter port group view: port-group manual <i>port-group-name</i> 	<p>Use either command.</p> <ul style="list-style-type: none"> The configuration made in Ethernet interface view applies only to the 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	By default, all ports are access ports.
4. Configure the hybrid ports to permit packets from specific MAC-based VLANs to pass through.	port hybrid vlan <i>vlan-list</i> { tagged untagged }	By default, a hybrid port only permits the packets of VLAN 1 to pass through.
5. Enable the MAC-based VLAN feature.	mac-vlan enable	Disabled by default.
6. Configure 802.1X/MAC/port authentication or any combination.	For more information, see <i>Security Command Reference</i> .	N/A


MAC-based VLAN configuration example

Network requirements

As shown in Figure 42:

- GigabitEthernet 1/0/1 of Device A and Device C are each connected to a meeting room. Laptop 1 and Laptop 2 are used for meetings and might be used in either of the two meeting rooms.
- Different departments own Laptop 1 and Laptop 2. The two departments use VLAN 100 and VLAN 200, respectively. Each laptop must be able to access only its own department server, no matter which meeting room it is used in.
- The MAC address of Laptop 1 is 000D-88F8-4E71, and that of Laptop 2 is 0014-222C-AA69.

Figure 42 Network diagram



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. Assign the uplink ports of Device B to VLANs 100 and 200.
- Associate the MAC address of Laptop 1 with VLAN 100, and associate the MAC address of Laptop 2 with VLAN 200.

Configuration procedure

1. Configure Device A:
Create VLANs 100 and 200.

```
<DeviceA> system-view  
[DeviceA] vlan 100
```

```
[DeviceA-vlan100] quit
[DeviceA] vlan 200
[DeviceA-vlan200] quit

# Associate the MAC address of Laptop 1 with VLAN 100, and associate the MAC address of
Laptop 2 with VLAN 200.
[DeviceA] mac-vlan mac-address 000d-88f8-4e71 vlan 100
[DeviceA] mac-vlan mac-address 0014-222c-aa69 vlan 200

# Configure Laptop 1 and Laptop 2 to access the network through GigabitEthernet 1/0/1.
Configure GigabitEthernet 1/0/1 as a hybrid port that sends packets of VLANs 100 and 200
untagged, and enable the MAC-based VLAN feature on it.
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] port link-type hybrid
[DeviceA-GigabitEthernet1/0/1] port hybrid vlan 100 200 untagged
Please wait... Done.
[DeviceA-GigabitEthernet1/0/1] mac-vlan enable
[DeviceA-GigabitEthernet1/0/1] quit

# To enable the laptops to access Server 1 and Server 2, configure the uplink port GigabitEthernet
1/0/2 as a trunk port, and assign it to VLANs 100 and 200.
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-type trunk
[DeviceA-GigabitEthernet1/0/2] port trunk permit vlan 100 200
[DeviceA-GigabitEthernet1/0/2] quit
```

2. Configure Device B:

```
# Create VLANs 100 and 200. Assign GigabitEthernet 1/0/13 to VLAN 100, and assign
GigabitEthernet 1/0/14 to VLAN 200.
<DeviceB> system-view
[DeviceB] vlan 100
[DeviceB-vlan100] port gigabitethernet 1/0/13
[DeviceB-vlan100] quit
[DeviceB] vlan 200
[DeviceB-vlan200] port gigabitethernet 1/0/14
[DeviceB-vlan200] quit

# Configure GigabitEthernet 1/0/3 and GigabitEthernet 1/0/4 as trunk ports, and assign them
to VLANs 100 and 200.
[DeviceB] interface gigabitethernet 1/0/3
[DeviceB-GigabitEthernet1/0/3] port link-type trunk
[DeviceB-GigabitEthernet1/0/3] port trunk permit vlan 100 200
[DeviceB-GigabitEthernet1/0/3] quit
[DeviceB] interface gigabitethernet 1/0/4
[DeviceB-GigabitEthernet1/0/4] port link-type trunk
[DeviceB-GigabitEthernet1/0/4] port trunk permit vlan 100 200
[DeviceB-GigabitEthernet1/0/4] quit
```

3. Configure Device C as you configure Device A.

Verifying the configuration

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 usually configured on the downlink ports of access layer devices, and cannot be configured together with the link aggregation function.

Configuring protocol-based VLANs

Introduction to protocol-based VLAN

You use the protocol-based VLAN feature to assign packets to VLANs by their application type.

The protocol-based VLAN feature assigns inbound packets to different VLANs based on their protocol type and encapsulation format. The protocols available for VLAN assignment include IP, IPX, and AppleTalk (AT), and the encapsulation formats include Ethernet II, 802.3 raw, 802.2 LLC, and 802.2 SNAP.

A protocol template defines a protocol type and an encapsulation format. A protocol-based VLAN ID and a protocol index, combined, can uniquely identify a protocol template. You can assign multiple protocol templates to a protocol-based VLAN.

Protocol-based VLAN assignment is available only on hybrid ports, and a protocol template applies only to untagged packets.

When an untagged packet arrives, a protocol-based VLAN assignment enabled hybrid port processes the packet by using the following workflow:

- If the protocol type and encapsulation format in the packet matches a protocol template, the packet is tagged with the VLAN tag specific to the protocol template.
- If no protocol template is matched, the packet is tagged with the PVID of the port.

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

- If the port is in the same VLAN as the packet, it forwards the packet.
- If not, the port drops the packet.

Configuration restrictions and guidelines

- 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 processes only untagged inbound packets, whereas the voice VLAN in automatic mode processes only tagged voice traffic. Do not configure a VLAN as both a protocol-based VLAN and a voice VLAN. For more information, see "[Configuring a voice VLAN](#)."

Configuration procedure

To configure a protocol-based VLAN:


Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter VLAN view.	vlan <i>vlan-id</i>	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 <i>etype-id</i> llc { dsap <i>dsap-id</i> [ssap <i>ssap-id</i>] ssap <i>ssap-id</i> } snap etype <i>etype-id</i> } }	Not configured by default.
4. Exit VLAN view.	quit	N/A
5. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type interface-number</i> 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> 	Use any command. <ul style="list-style-type: none"> The configuration made in Ethernet interface view applies only to the port. The configuration made in port group view applies to all ports in the port group. The configuration made in Layer 2 aggregate interface view applies to the aggregate interface and its aggregation member ports. If the system fails to apply the configuration to the aggregate interface, it stops applying the configuration to aggregation member ports. If the system fails to apply the configuration to an aggregation member port, it skips the port and moves to the next member port.
6. Configure the port link type as hybrid.	port link-type hybrid	By default, all ports are access ports.
7. Assign the hybrid port to the specified protocol-based VLANs.	port hybrid vlan <i>vlan-list</i> { tagged untagged }	By default, a hybrid port is only in VLAN 1.
8. Assign the protocol template you have created to the hybrid port.	port hybrid protocol-vlan <i>vlan-id</i> { <i>protocol-index</i> [to <i>protocol-end</i>] all }	N/A

Protocol-based VLAN configuration example

Network requirements

In a lab environment, as shown in Figure 43, most hosts run the IPv4 protocol, and the rest of the hosts run the IPv6 protocol for teaching purposes. To avoid interference, isolate IPv4 traffic and IPv6 traffic at Layer 2.

Figure 43 Network diagram



Configuration consideration

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

Configuration procedure

1. Configure Device:

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

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

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

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

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

```
[Device-vlan200] protocol-vlan 1 ipv6
[Device-vlan200] quit
```

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

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

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

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

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

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

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

2. Keep the default settings of L2 Switch A and L2 Switch B.
3. Configure IPv4 Host A, IPv4 Host B, and IPv4 Server to be on the same IP subnet (192.168.100.0/24, for example), and configure IPv6 Host A, IPv6 Host B, and IPv6 Server to be on the same IP subnet (2001::1/64, for example).

Verifying the configurations

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 or server in VLAN 100 cannot ping the hosts and server in VLAN 200, and vice versa.
2. Display protocol-based VLAN information on Device to determine whether the configurations have become valid.

Display protocol-based VLAN configuration on Device.

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

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

```
[Device-GigabitEthernet1/0/2] display protocol-vlan interface all
Interface: GigabitEthernet 1/0/1
```

VLAN ID	Protocol Index	Protocol Type
100	1	ipv4
200	1	ipv6
Interface: GigabitEthernet 1/0/2		
VLAN ID	Protocol Index	Protocol Type
100	1	ipv4
200	1	ipv6

Configuration guidelines

Protocol-based VLAN configuration applies only to hybrid ports.

Configuring IP subnet-based VLANs

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 IP subnet or IP address to a specific VLAN.

Configuration procedure

ⓘ IMPORTANT:

This feature is applicable only on hybrid ports.

To configure an IP subnet-based VLAN:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter VLAN view.	vlan <i>vlan-id</i>	N/A
3. Associate an IP subnet with the VLAN.	ip-subnet-vlan [<i>ip-subnet-index</i>] ip <i>ip-address</i> [<i>mask</i>]	The IP subnet or IP address to be associated with a VLAN cannot be a multicast subnet or a multicast address.
4. Return to system view.	quit	N/A

Step	Command	Remarks
5. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> 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> 	<p>Use any command.</p> <ul style="list-style-type: none"> The configuration made in Ethernet interface view applies only to the port. The configuration made in port group view applies to all ports in the port group. The configuration made in Layer 2 aggregate interface view applies to the aggregate interface and its aggregation member ports. If the system fails to apply the configuration to the aggregate interface, it stops applying the configuration to aggregation member ports. If the system fails to apply the configuration to an aggregation member port, it skips the port and moves to the next member port.
6. Configure port link type as hybrid.	port link-type hybrid	By default, all ports are access ports.
7. Configure the hybrid ports to permit the specified IP subnet-based VLANs to pass through.	port hybrid vlan <i>vlan-list</i> { tagged untagged }	By default, a hybrid port allows only packets from VLAN 1 to pass through untagged.
8. Associate the hybrid ports with the specified IP subnet-based VLAN.	port hybrid ip-subnet-vlan <i>vlan</i> <i>vlan-id</i>	Not configured by default.


IP subnet-based VLAN configuration example

Network requirements

As shown in [Figure 44](#), the hosts in the office belong to different IP subnets 192.168.5.0/24 and 192.168.50.0/24.

Configure Device C to transmit packets over separate VLANs based on their source IP addresses.

Figure 44 Network diagram



Configuration consideration

- Create VLANs 100 and 200.
- Associate IP subnets with the VLANs.
- Assign ports to the VLANs.

Configuration procedure

Associate IP subnet 192.168.5.0/24 with VLAN 100.

```
<DeviceC> system-view
[DeviceC] vlan 100
[DeviceC-vlan100] ip-subnet-vlan ip 192.168.5.0 255.255.255.0
[DeviceC-vlan100] quit
```

Associate IP subnet 192.168.50.0/24 with VLAN 200.

```
[DeviceC] vlan 200
[DeviceC-vlan200] ip-subnet-vlan ip 192.168.50.0 255.255.255.0
[DeviceC-vlan200] quit
```

Configure interface GigabitEthernet 1/0/11 to permit packets of VLAN 100 to pass through.

```
[DeviceC] interface gigabitethernet 1/0/11
[DeviceC-GigabitEthernet1/0/11] port link-type hybrid
[DeviceC-GigabitEthernet1/0/11] port hybrid vlan 100 tagged
Please wait... Done.
[DeviceC-GigabitEthernet1/0/11] quit
```

Configure interface GigabitEthernet 1/0/12 to permit packets of VLAN 200 to pass through.

```
[DeviceC] interface gigabitethernet 1/0/12
[DeviceC-GigabitEthernet1/0/12] port link-type hybrid
[DeviceC-GigabitEthernet1/0/12] port hybrid vlan 200 tagged
Please wait... Done.
```

```
[DeviceC-GigabitEthernet1/0/12] quit

# Associate interface GigabitEthernet 1/0/1 with IP subnet-based VLANs 100 and 200.
[DeviceC] interface gigabitethernet 1/0/1
[DeviceC-GigabitEthernet1/0/1] port link-type hybrid
[DeviceC-GigabitEthernet1/0/1] port hybrid vlan 100 200 untagged
Please wait... Done.
[DeviceC-GigabitEthernet1/0/1] port hybrid ip-subnet-vlan vlan 100
[DeviceC-GigabitEthernet1/0/1] port hybrid ip-subnet-vlan vlan 200
[DeviceC-GigabitEthernet1/0/1] return
```

Verifying the configurations

```
# Display the IP subnet information for all VLANs.
<Device C> display ip-subnet-vlan vlan all
VLAN ID: 100
Subnet Index      IP Address      Subnet Mask
=====
0                 192.168.5.0    255.255.255.0
VLAN ID: 200
Subnet Index      IP Address      Subnet Mask
=====
0                 192.168.50.0   255.255.255.0

# Display the IP subnet-based VLAN information on GigabitEthernet 1/0/1.
<DeviceC> display ip-subnet-vlan interface gigabitethernet 1/0/1
Interface: GigabitEthernet1/0/1
VLAN ID   Subnet-Index   IP ADDRESS      NET MASK
=====
100       0              192.168.5.0     255.255.255.0
200       0              192.168.50.0    255.255.255.0
```

Configuration guidelines

The IP subnet-based VLAN configurations are only effective on hybrid ports.

Displaying and maintaining VLAN

Task	Command	Remarks
Display VLAN information.	display vlan [<i>vlan-id1</i> [<i>to</i> <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 [<i>vlan-interface</i>] [brief [down]] [[{ begin exclude include } <i>regular-expression</i>]] display interface <i>vlan-interface</i> <i>vlan-interface-id</i> [brief] [[{ 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

Task	Command	Remarks
Display MAC address-to-VLAN entries.	display mac-vlan { all dynamic mac-address <i>mac-address</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-type</i> <i>interface-number1</i> [to <i>interface-type</i> <i>interface-number2</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

Configuring an isolate-user-VLAN

Overview


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

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

- Isolate-user-VLANs are mainly used for upstream data exchange. An isolate-user-VLAN can be associated with multiple secondary VLANs. As the upstream device identifies only the isolate-user-VLAN and not the secondary VLANs, network configuration is simplified and VLAN resources are saved.
- You can isolate the Layer 2 traffic of different users by assigning the ports connected to them to different secondary VLANs. To enable communication between secondary VLANs associated with the same isolate-user-VLAN, you can enable local proxy ARP on the upstream device (for example, Device A in [Figure 45](#)) to realize Layer 3 communication between the secondary VLANs.

As shown in [Figure 45](#), the isolate-user-VLAN function is enabled on Device B. VLAN 10 is the isolate-user-VLAN. VLAN 2, VLAN 5, and VLAN 8 are secondary VLANs associated with VLAN 10 and are invisible to Device A.

Figure 45 An isolate-user-VLAN example



To configure an isolate-user-VLAN, complete the following tasks:

1. Configure the isolate-user-VLAN.
2. Configure the secondary VLANs.
3. Associate the isolate-user-VLAN with the specified secondary VLANs.
4. Configure uplink and downlink ports:
 - Configure the uplink ports, for example, the port connecting Device B to Device A in [Figure 45](#), to operate in **promiscuous** mode in the specified VLAN, so that the uplink ports can be added to the specified isolate-user-VLAN and the secondary VLANs associated with the isolate-user-VLAN automatically.

- Configure the downlink ports, for example, the ports connecting Device B to hosts in [Figure 45](#), to operate in host mode, so that the downlink ports can be added to the isolate-user-VLAN associated with the secondary VLAN automatically.

For more information about the promiscuous and host mode commands, see *Layer 2—LAN Switching Command Reference*.

Configuration restrictions and guidelines

- To enable users in the isolate-user-VLAN to communicate with other networks at Layer 3, follow these steps:
 - a. 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).
 - b. You must configure the **isolated-vlan enable** command for at least one secondary VLAN to isolate the ports in the secondary VLAN.
- The dynamic MAC addresses entries learned in the isolate-user-VLAN are automatically synchronized to all the secondary VLANs, and the dynamic MAC address entries learned in a secondary VLAN are automatically synchronized to the isolate-user-VLAN. Static MAC address entries cannot be automatically synchronized. If you have configured static MAC address entries in the isolate-user-VLAN, you should also configure the same static MAC address entries in the secondary VLANs to avoid broadcasts, and vice versa.

Configuration procedure

To configure an isolate-user-VLAN:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a VLAN and enter VLAN view.	vlan <i>vlan-id</i>	N/A
3. Configure the VLAN as an isolate-user-VLAN.	isolate-user-vlan enable	Not configured by default.
4. Return to system view.	quit	N/A
5. Create secondary VLANs.	vlan { <i>vlan-id1</i> [to <i>vlan-id2</i>] all }	N/A
6. Configure Layer 2 isolation between ports in the same secondary VLAN.	isolated-vlan enable	Optional. By default, ports in the same secondary VLAN can communicate with one another at Layer 2. This configuration takes effect only after you configure all ports in the same secondary VLAN to operate in host mode and associate the secondary VLANs with an isolate-user-VLAN.
7. Return to system view.	quit	N/A

Step	Command	Remarks
8. Associate the isolate-user-VLAN with the specified secondary VLANs.	isolate-user-vlan <i>isolate-user-vlan-id</i> secondary <i>secondary-vlan-id</i> [to <i>secondary-vlan-id</i>]	Not configured by default.
9. Configure the uplink port for the isolate-user-VLAN.	<ul style="list-style-type: none"> a. Enter Layer 2 Ethernet or aggregate interface view: interface <i>interface-type</i> <i>interface-number</i> Or interface bridge-aggregation <i>interface-number</i> b. Configure the port to operate in promiscuous mode in a specific VLAN: port isolate-user-vlan <i>vlan-id</i> promiscuous 	By default, a port does not operate in promiscuous mode or host mode in a VLAN.
10. Return to system view.	quit	N/A
11. Configure a downlink port for the isolate-user-VLAN.	<ul style="list-style-type: none"> a. Enter Layer 2 Ethernet or aggregate interface view: interface <i>interface-type</i> <i>interface-number</i> Or interface bridge-aggregation <i>interface-number</i> b. Configure the link type of the port: port link-type { access hybrid trunk } c. Assign the downlink port to the secondary VLAN according to its link type: port access vlan <i>vlan-id</i> Or port hybrid vlan <i>vlan-list</i> { tagged untagged } Or port trunk permit vlan { <i>vlan-list</i> all } d. Configure the port to operate in host mode: port isolate-user-vlan host 	By default, a port does not operate in host mode or promiscuous mode.
12. Return to system view.	quit	N/A

Displaying and maintaining isolate-user-VLAN

Task	Command	Remarks
Display the mapping between an isolate-user-VLAN and its secondary 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 46:

- 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 only has VLAN 6.

Figure 46 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
```

 - # Create secondary VLANs.

```
[DeviceB] vlan 2 to 3
```



```

# Associate the isolate-user-VLAN with the secondary VLANs.
[DeviceB] isolate-user-vlan 5 secondary 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 and GigabitEthernet 1/0/2 to VLAN 3 and
VLAN 2, respectively, 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

```

2. Configure Device C:

```

# Configure the isolate-user-VLAN.
<DeviceC> system-view
[DeviceC] vlan 6
[DeviceC-vlan6] isolate-user-vlan enable
[DeviceC-vlan6] quit
# Create secondary VLANs.
[DeviceC] vlan 3 to 4
# Associate the isolate-user-VLAN with the secondary VLANs.
[DeviceC] isolate-user-vlan 6 secondary 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 and GigabitEthernet 1/0/4 to VLAN 3 and
VLAN 4, respectively, 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

```

Verifying the configuration

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

Configuring a voice VLAN

Overview

A voice VLAN is configured for voice traffic. After assigning the ports that connect to voice devices to a voice VLAN, the system automatically configures quality of service (QoS) parameters for voice traffic, to improve the transmission priority of voice traffic and ensure voice quality.

Common voice devices include IP phones and integrated access devices (IADs). Only IP phones are used in the voice VLAN configuration examples in this document.

When an IP phone accesses a device, the device must have completed the following tasks:

- Identifying the IP phone and obtaining the MAC address of the IP phone, so that the device can perform security authentication for the IP phone and increase the priority for voice traffic.
- Advertising the voice VLAN information to the IP phone, so that the IP phone can automatically configure the voice VLAN according to the received information and the voice packets sent out of the IP phone can be transmitted within the voice VLAN.

Methods of identifying IP phones

OUI addresses

A device determines whether a received packet is a voice packet by evaluating its source MAC address. A packet whose source MAC address complies with the Organizationally Unique Identifier (OUI) address of the voice device is regarded as voice traffic.

You can remove the default OUI address of a device manually and then add new ones manually. You can configure the OUI addresses of a device in advance or use the default OUI addresses. [Table 14](#) lists the default OUI address for each vendor's devices. The switch supports 16 OUI addresses.

Table 14 The default OUI addresses of different vendors

Number	OUI address	Vendor
1	0001-E300-0000	Siemens phone
2	0003-6B00-0000	Cisco phone
3	0004-0D00-0000	Avaya phone
4	00D0-1E00-0000	Pingtel phone
5	0060-B900-0000	Philips/NEC phone
6	00E0-7500-0000	Polycom phone
7	00E0-BB00-0000	3Com phone

In general, as the first 24 bits of a MAC address (in binary format), an OUI address is a globally unique identifier that IEEE assigns to a vendor. In this document, however, OUI addresses are addresses that the system uses to determine whether a received packet is a voice packet. They are the results of the AND operation of the arguments *mac-address* and *oui-mask* in the **voice vlan mac-address** command.

Automatically identifying IP phones through LLDP

When you use OUI addresses to identify IP phones, the number of OUI addresses that can be configured is limited. Additionally, when there are plenty of IP phones in the network, you must configure many OUI addresses. If your IP phones support LLDP, you can configure LLDP to automatically identify IP phones. For more information, see "[Enabling LLDP to automatically discover IP phones.](#)"

Configuring a device to advertise voice VLAN information to IP phones

How a device advertises voice VLAN information

When an IP phone supports LLDP, the device can advertise the voice VLAN information to IP phones through the LLDP-MED TLVs.

When an IP phone supports CDP rather than LLDP, you can configure CDP compatibility to enable the device to advertise the voice VLAN information to IP phones through the CDP packets.

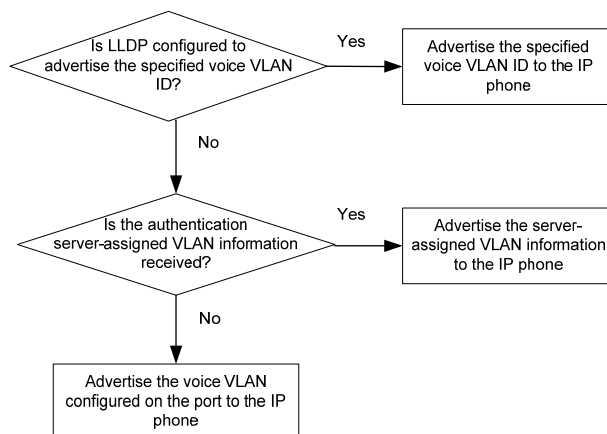
For more information about LLDP and CDP compatibility, see "Configuring LLDP."

How a device obtains voice VLAN information

- Configuring a voice VLAN on ports. For more information, see "[Configuring a voice VLAN on a port.](#)"
- When an IP phone cooperates with the authentication function, the server-assigned VLAN information can be advertised to the IP phone. For more information, see "[Dynamically advertising server-assigned VLANs through LLDP.](#)"
- Configuring LLDP to advertise the specific voice VLAN information. For more information, see "[Configuring LLDP to advertise a specific voice VLAN.](#)"

Figure 47 shows the order of the three methods.

Figure 47 Workflow of advertising voice VLAN information to IP phones




IP phone access methods

Connecting the host and the IP phone in series

As shown in Figure 48, the host is connected to the IP phone, and the IP phone is connected to the device. When the host and the IP phone are connected in series, the host and the IP phone must be assigned to different VLANs, and the IP phone must be able to send out VLAN-tagged packets, so that the data traffic and the voice traffic can be distinguished. Also, you must configure the port to allow the voice VLAN and the PVID.


Figure 48 Connecting the host and IP phone in series



Connecting the IP phone to the device separately

As shown in Figure 49, the IP phone is separately connected to the device. This connection method applies when the IP phone sends out untagged voice packets. In this case, you must configure the voice VLAN as the PVID of the port and configure the port to allow the PVID.

Figure 49 Connecting the IP phone to the device separately



Configuring a voice VLAN on a port

Voice VLAN assignment modes

A port can be assigned to a voice VLAN in automatic mode or manual mode.

Automatic mode

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, as shown in Figure 48.

The system matches the source MAC address carried in the protocol packets sent when an IP phone is powered on against the device's OUI addresses. If the system finds a match, it automatically assigns the

receiving port to the voice VLAN, issues ACL rules, and configures the packet precedence. You can configure a voice VLAN aging time on the device. The system will remove a port from the voice VLAN if no packet is received from the port during the aging time. The system automatically assigns ports to, or removes ports from, a voice VLAN. When the voice VLAN works correctly, when the system reboots, the system reassigns ports in automatic voice VLAN assignment mode to the voice VLAN after the reboot, ensuring that existing voice connections can work correctly. In this case, voice traffic streams do not trigger port assignment to the voice VLAN.

Manual mode

Manual mode is suitable for scenarios where only IP phones access the network through the device and ports on the device transmit only voice traffic, as shown in [Figure 49](#). 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.

You must manually assign an IP phone accessing port to a voice VLAN. Then, the system matches the source MAC addresses carried in the packets against the device's OUI addresses. If the system finds a match, it issues ACL rules and configures the packet precedence. In this mode, you must manually assign ports to, or remove ports from, a voice VLAN.

Cooperation of voice VLAN assignment modes and IP phones

Both modes forward tagged packets sent out of IP phones according to their tags. Some IP phones can send out VLAN-tagged packets, and some IP phones can send out only untagged packets. [Table 15](#) and [Table 16](#) list the configurations required for ports of different link types to support tagged or untagged voice traffic sent from IP phones when different voice VLAN assignment modes are configured.

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 15 Required configurations on ports of different link types for them to support tagged voice traffic

Port link type	Voice VLAN assignment mode	Support for tagged voice traffic	Configuration requirements
Access	Automatic	No	N/A
	Manual		
Trunk	Automatic	Yes	The PVID of the port cannot be the voice VLAN.
	Manual		The PVID of the port cannot be the voice VLAN. Configure the port to permit packets of the voice VLAN to pass through.
Hybrid	Automatic	Yes	The PVID of the port cannot be the voice VLAN.
	Manual		The PVID of the port cannot be the voice VLAN. Configure the port to permit packets of the voice VLAN to pass through tagged.

Table 16 Required configurations on ports of different link types for them to support tagged voice traffic

Port link type	Voice VLAN assignment mode	Support for untagged voice traffic	Configuration requirements
Access	Automatic	No	N/A
	Manual	Yes	Configure the PVID of the port as the voice VLAN.
Trunk	Automatic	No	N/A

Port link type	Voice VLAN assignment mode	Support for untagged voice traffic	Configuration requirements
Hybrid	Manual	Yes	Configure the PVID of the port as the voice VLAN and assign the port to the voice VLAN.
	Automatic	No	N/A
	Manual	Yes	Configure the PVID of the port as the voice VLAN and configure the port to permit packets of the voice VLAN to pass through untagged.

When you configure the voice VLAN assignment modes, follow these guidelines:

- If an IP phone sends tagged voice traffic and its accessing port is configured with 802.1X authentication and any of the guest VLAN, Auth-Fail VLAN, and critical VLAN features, assign different VLAN IDs to the voice VLAN, PVID of the connecting port, and 802.1X guest, Auth-Fail, or critical VLAN.
- If an IP phone sends untagged voice traffic, to implement the voice VLAN feature, you must configure the PVID of the IP phone's accessing port as the voice VLAN. As a result, you cannot implement 802.1X authentication.

Security mode and normal mode of voice VLANs

Depending on their inbound packet filtering mechanisms, voice VLAN-enabled ports operate in the following modes:

- **Normal mode**—Voice VLAN-enabled ports receive packets that carry the voice VLAN tag, and forward packets in the voice VLAN without comparing their source MAC addresses against the OUI addresses configured for the device. If the PVID of the port is the voice VLAN and the port operates in manual VLAN assignment mode, the port forwards all received untagged packets in the voice VLAN. In normal mode, voice VLANs are vulnerable to traffic attacks. Malicious users might send large quantities of forged voice VLAN-tagged or untagged packets to consume the voice VLAN bandwidth, affecting normal voice communication.
- **Security mode**—Only voice packets whose source MAC addresses match the recognizable OUI addresses can pass through the voice VLAN-enabled inbound port, but 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 transmitting both voice traffic and non-voice traffic in a voice VLAN. If you must transmit both voice traffic and non-voice traffic, make sure that the voice VLAN security mode is disabled.

Table 17 How a voice VLAN-enabled port processes packets in security and normal mode

Voice VLAN mode	Packet type	Packet processing mode
Security mode	Untagged packets	If the source MAC address of a packet matches an OUI address configured for the device, it is forwarded in the voice VLAN; otherwise, it is dropped.
	Packets that carry the voice VLAN tag	
	Packets that carry other tags	Forwarded or dropped depending on whether the port allows packets of these VLANs to pass through
Normal mode	Untagged packets	The port does not determine the source MAC addresses of inbound packets. In this way, both voice traffic and non-voice traffic can be transmitted in the voice VLAN.
	Packets that carry the voice VLAN tag	
	Packets that carry other tags	Forwarded or dropped depending on whether the port allows packets of these VLANs to pass through

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 you enable 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](#)."

Configuring QoS priority settings for voice traffic on an interface

In voice VLAN applications, you can improve the quality of voice traffic by configuring the appropriate QoS priority settings, including the Class of Service (CoS) and Differentiated Services Code Point (DSCP) values, for voice traffic. Voice traffic carries its own QoS priority settings. You can configure the device either to modify or not to modify the QoS priority settings carried by incoming voice traffic.

Configuration restrictions and guidelines

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

Configuration procedure

To configure QoS priority settings for voice traffic:

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

Step	Command	Remarks
2. Enter Layer 2 Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
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	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.

Configuring a port to operate in automatic voice VLAN assignment mode

Configuration restrictions and guidelines

- 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 "Configuring VLANs."
- Do not configure automatic voice VLAN assignment together with MSTP, because the former is mainly configured on the access side. With MSTP enabled, if a port is blocked in the MST instance (MSTI) of the target voice 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 automatic voice VLAN assignment together with PVST, because the former is mainly configured on the access side. With PVST enabled, if the target voice VLAN is not permitted on a port, the port is placed in the blocked state and 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.

Configuration procedure

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Set the voice VLAN aging time.	voice vlan aging <i>minutes</i>	Optional. By default, the aging time of a voice VLAN is 1440 minutes. The voice VLAN aging time configuration is only applicable on ports in automatic voice VLAN assignment mode.

Step	Command	Remarks
3. Enable the voice VLAN security mode.	voice vlan security enable	Optional. By default, the voice VLAN security mode is enabled.
4. Add a recognizable OUI address.	voice vlan mac-address <i>oui</i> mask <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 14 .
5. Enter Layer 2 Ethernet interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
6. Configure the link type of the Ethernet interface.	<ul style="list-style-type: none"> port link-type trunk port link-type hybrid 	Use one of the commands.
7. Configure the port to operate in automatic voice VLAN assignment mode.	voice vlan mode auto	Optional. By default, the automatic voice VLAN assignment mode is enabled. The voice VLAN assignment modes on different ports are independent of one another.
8. Enable the voice VLAN feature.	voice vlan <i>vlan-id</i> enable	By default, the voice VLAN feature is disabled.

Configuring a port to operate in manual voice VLAN assignment mode

Configuration restrictions and guidelines

- You can configure different voice VLANs on different ports at the same time. However, you can configure one port with only one voice VLAN, and this voice VLAN must be a static VLAN that already exists on the device.
- You cannot enable voice VLAN on the member ports of a link aggregation group. For more information about the member ports, see "Configuring Ethernet link aggregation."
- To make voice VLAN take effect on a port that is enabled with voice VLAN and operates in manual voice VLAN assignment mode, you must assign the port to the voice VLAN manually.

Configuration procedure

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable the voice VLAN security mode.	voice vlan security enable	Optional. By default, the voice VLAN security mode is enabled.

Step	Command	Remarks
3. Add a recognizable OUI address.	voice vlan mac-address <i>oui</i> mask <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 14 .
4. Enter interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
5. Configure the port to operate in manual voice VLAN assignment mode.	undo voice vlan mode auto	By default, the manual voice VLAN assignment mode is disabled.
6. Assign the access, trunk, or hybrid port in manual voice VLAN assignment mode to the voice VLAN.	For the configuration procedure, see "Configuring VLANs."	After you assign an access port to the voice VLAN, the voice VLAN becomes the PVID of the port automatically.
7. Configure the voice VLAN as the PVID of the trunk or hybrid port.	For the configuration procedure, see "Configuring VLANs."	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	Disabled by default.

Enabling LLDP to automatically discover IP phones

In a traditional voice VLAN network, the switch maps the source MAC addresses of IP phones to a limited number of OUI addresses to allow them to access the network. This method restricts the types of IP phones on the network, if the IP phones with the source MAC addresses match the same OUI address are categorized as a type.

To break the restriction, you can enable the switch to automatically discover IP phones through LLDP. With this function, the switch can automatically discover the peer, and exchange LLDP TLVs with the peer. If the LLDP System Capabilities TLV received on a port shows that the peer is phone capable, the switch determines that the peer is an IP phone and sends an LLDP TLV carrying the voice VLAN configuration to the peer.

When the IP phone discovery process is complete, the port will automatically join the voice VLAN and improve the transmission priority of the voice traffic for the IP phone. To ensure that the IP phone can pass authentication, the switch will add the MAC address of the IP phone to the MAC address table.

Configuration prerequisites

Before you enable the switch to automatically discover IP phones through LLDP, complete the following tasks:

- Enable LLDP globally and on ports.
- Configure voice VLANs.

Configuration procedure

To enable LLDP to automatically discover IP phones:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable LLDP to automatically discover IP phones.	voice vlan track lldp	Disabled by default.

⚠ IMPORTANT:

- When the switch is enabled to automatically discover IP phones through LLDP, you can connect at most five IP phones to each port of the switch.
- You cannot use this function together with CDP compatibility.

Configuring LLDP to advertise a specific voice VLAN

Voice VLAN advertisement through LLDP is available only for LLDP-enabled IP phones. If CDP-compatibility is enabled, this feature is also available for CDP-enabled IP phones. For more information about CDP compatibility, see "Configuring LLDP."

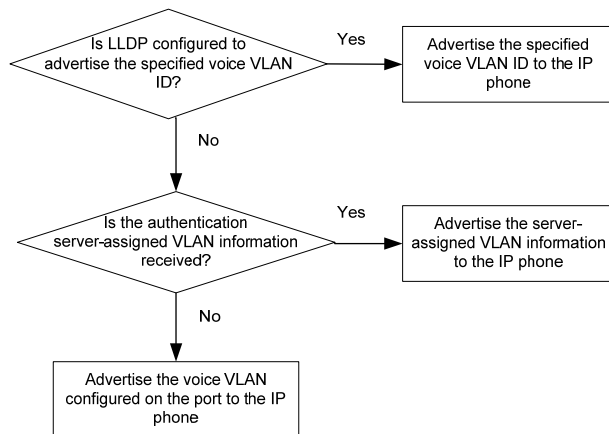
Configuration guidelines

Use this feature in one of the following scenarios:

- Decrease the voice VLAN processing delay in an IRF fabric.
On an LLDP-enabled port, LLDP advertises the voice VLAN information to the IP phone connected to the port. When a packet arrives on the port, the switch compares the source MAC address against its voice device OUI list. If a match is found, the switch learns the MAC address in the voice VLAN, and promotes the forwarding priority for the packet. Because this process is completed in software, in an IRF fabric, MAC address learning and synchronization of the learned MAC address entry to all member devices introduces an undesirable delay. Directly specifying the voice VLAN to be advertised by LLDP enables the IRF fabric to learn and synchronize MAC address entries faster in hardware.
- Avoid configuring the voice VLAN function on a port.

Figure 50 shows the procedure of voice VLAN advertisement through LLDP.

Figure 50 Voice VLAN advertisement through LLDP



With the received voice VLAN information, the IP phone automatically completes the voice VLAN configuration, including the voice VLAN ID, tagging status, and priority. This voice VLAN can be the voice VLAN directly specified for LLDP advertisement, the voice VLAN configured on the port, or the voice VLAN assigned by a server, depending on your configuration.

To identify the voice VLAN advertised by LLDP, execute the **display lldp local-information** command, and examine the MED information fields in the command output.

The LLDP packets that the device send to IP phones carry the priority information, but the CDP packets that the device send to IP phones do not carry the priority information.

Configuration procedure

To configure LLDP to advertise a specific voice VLAN:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none">Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>Enter port group view: port-group manual <i>port-group-name</i>	Use one of the commands.
3. Configure LLDP to advertise a specific voice VLAN.	lldp voice-vlan <i>vlan-id</i>	By default, LLDP advertises the voice VLAN configured on the port.

Dynamically advertising server-assigned VLANs through LLDP

Overview

Dynamic advertisement of server-assigned VLANs through LLDP must work with 802.1X or MAC authentication, and is available only for LLDP-enabled IP phones. If 802.1X authentication is used, make sure the IP phones also support 802.1X authentication.

To implement this function for an IP phone, perform the following configuration tasks:

- Enable LLDP globally and on the port connected to the IP phone.
- Configure 802.1X or MAC authentication to make sure the IP phone can pass security authentication. For more information about 802.1X authentication, MAC authentication, and VLAN assignment by servers, see *Security Configuration Guide*.
- Configure VLAN authorization for the IP phone on the authentication server.

After the IP phone passes authentication, LLDP advertises the server-assigned VLAN in the LLDP-MED Network Policy TLV to the IP phone. The IP phone will send its traffic tagged with the assigned VLAN. Also, the port connected to the IP phone will be added to the server-assigned VLAN.


Example for using 802.1X to authenticate IP phones

As shown in [Figure 51](#), configure 802.1X on the device to authenticate the host and the IP phone (which must support 802.1X). Configure the authentication server to assign an untagged VLAN to the host and assign a tagged VLAN to the IP phone. After the host and the IP phone pass the authentication, the port connected to the IP phone is added to the VLAN assigned to the IP phone as an tagged member and added to the VLAN assigned to the host as a untagged member. Also, the LLDP-MED TLVs that the device sends to the IP phone carry information about the VLAN assigned to the IP phone, so that the voice packets sent out of the IP phone can be forwarded in the server-assigned VLAN with tags.

The EAPOL packets defined in the 802.1X protocol do not carry VLAN tags. When the server is configured to assign a tagged VLAN to the IP phone, you must configure the port connected to the IP phone to send 802.1X protocol packets without tags by using the **dot1x eapol untag** command.

Only 802.1X supports assigning tagged VLANs.

Figure 51 Using 802.1X to authenticate an IP phone



Displaying and maintaining voice VLAN

Task	Command	Remarks
Display the voice VLAN state.	display voice vlan state [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Task	Command	Remarks
Display the OUI addresses that the system supports.	display voice vlan oui [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Voice VLAN configuration examples

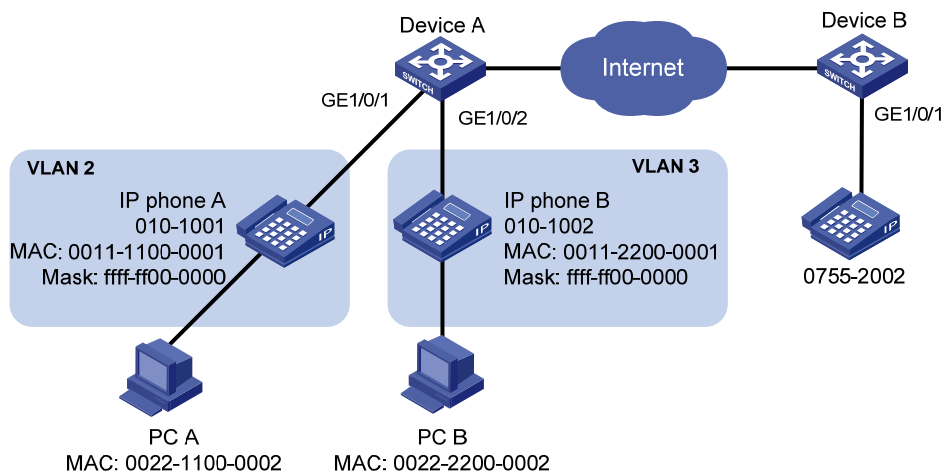
Automatic voice VLAN mode configuration example

Network requirements

As shown in Figure 52,

- The MAC address of IP phone A is 0011-1100-0001. The phone connects to a downstream device named PC A whose MAC address is 0022-1100-0002 and to GigabitEthernet 1/0/1 on an upstream device named Device A.
- The MAC address of IP phone B is 0011-2200-0001. The phone connects to a downstream device named PC B whose MAC address is 0022-2200-0002 and to GigabitEthernet 1/0/2 on Device A.
- Device A uses voice VLAN 2 to transmit voice packets for IP phone A and uses voice VLAN 3 to transmit voice packets for IP phone B.
- Configure GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to operate 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 52 Network diagram



Configuration procedure

Create VLAN 2 and VLAN 3.

```

<DeviceA> system-view
[DeviceA] vlan 2 to 3
Please wait... Done.
  
```

Set the voice VLAN aging time to 30 minutes.

```

[DeviceA] voice vlan aging 30
  
```

Since GigabitEthernet 1/0/1 might receive both voice traffic and data traffic at the same time, to ensure the quality of voice packets and effective bandwidth use, configure voice VLANs to operate in security mode. Configure the voice VLANs to transmit only voice packets. By default, voice VLANs operate 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. In this way, 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
```

Verifying the configurations

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
0060-b900-0000	ffff-ff00-0000	Philips/NEC phone
00d0-1e00-0000	ffff-ff00-0000	Pingtel phone
00e0-7500-0000	ffff-ff00-0000	Polycom phone
00e0-bb00-0000	ffff-ff00-0000	3com phone

Display the states of voice VLANs.

```
<DeviceA> display voice vlan state
```

```
Maximum of Voice VLANs: 8
```

```
Current Voice VLANs: 2
```

```
Voice VLAN security mode: Security
```

```
Voice VLAN aging time: 30 minutes
```

```
Voice VLAN enabled port and its mode:
```


PORT	VLAN	MODE	COS	DSCP
GigabitEthernet1/0/1	2	AUTO	6	46
GigabitEthernet1/0/2	3	AUTO	6	46


Manual voice VLAN assignment mode configuration example

Network requirements

As shown in Figure 53,

- Create VLAN 2 and configure it as a voice VLAN that permits only voice traffic to pass through.
- The IP phones send untagged voice traffic. Configure GigabitEthernet 1/0/1 as a hybrid port.
- Configure GigabitEthernet 1/0/1 to operate in manual voice VLAN assignment mode. Configure GigabitEthernet 1/0/1 to allow voice traffic with an OUI address of 0011-2200-0000, a mask of ffff-ff00-0000, and a description string of **test** to be forwarded in the voice VLAN.

Figure 53 Network diagram



Configuration procedure

Configure the voice VLAN to operate in security mode. A voice VLAN operates in security mode by default. (Optional)

```
<DeviceA> system-view
[DeviceA] voice vlan security enable
```

Add a recognizable OUI address 0011-2200-0000.

```
[DeviceA] voice vlan mac-address 0011-2200-0000 mask ffff-ff00-0000 description test
```

Create VLAN 2.

```
[DeviceA] vlan 2
[DeviceA-vlan2] quit
```

Configure GigabitEthernet 1/0/1 to operate in manual voice VLAN assignment mode.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] undo voice vlan mode auto
```

Configure GigabitEthernet 1/0/1 as a hybrid port.

```
[DeviceA-GigabitEthernet1/0/1] port link-type hybrid
```

Configure the voice VLAN (VLAN 2) as the PVID of GigabitEthernet 1/0/1 and configure GigabitEthernet 1/0/1 to permit the voice traffic of VLAN 2 to pass through untagged.

```
[DeviceA-GigabitEthernet1/0/1] port hybrid pvid vlan 2
[DeviceA-GigabitEthernet1/0/1] port hybrid vlan 2 untagged
```

Enable voice VLAN on GigabitEthernet 1/0/1.

```
[DeviceA-GigabitEthernet1/0/1] voice vlan 2 enable
```

Verifying the configurations

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 states of voice VLANs.

```
<DeviceA> display voice vlan state
```

Maximum of Voice VLANs: 8

Current Voice VLANs: 1

Voice VLAN security mode: Security

Voice VLAN aging time: 1440 minutes

Voice VLAN enabled port and its mode:

PORT	VLAN	MODE	COS	DSCP

GigabitEthernet1/0/1	2	MANUAL	6	46

Configuring GVRP

The Generic Attribute Registration Protocol (GARP) provides a generic framework for devices in a switched LAN, such as end stations and switches, to register and deregister attribute values. The GARP VLAN Registration Protocol (GVRP) is a GARP application that registers and deregisters VLAN attributes. GVRP uses the operating mechanism of GARP to maintain and propagate dynamic VLAN registration information for GVRP devices on the network.

Overview

GARP

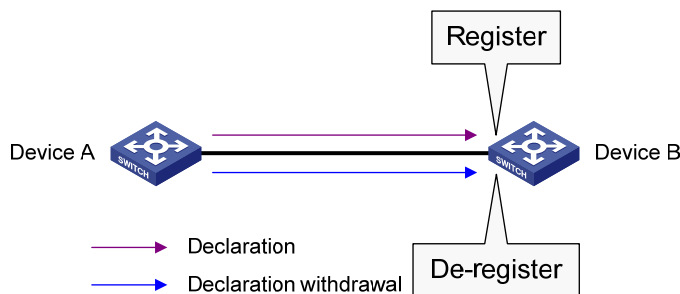
GARP provides a mechanism that allows participants in a GARP application to distribute, propagate, and register—with other participants in a LAN—the attributes specific to the GARP application, such as VLAN or multicast address attributes.

How GARP works

Each port that participates in a GARP application (GVRP, for example) is a GARP participant.

GARP enables GARP participants to propagate attribute values throughout the switched LAN. As shown in [Figure 54](#), a GARP participant registers and deregisters its attribute values with other GARP participants by sending and withdrawing declarations, and registers and deregisters the attribute values of other participants according to the declarations and withdrawals that it has received.

Figure 54 How GARP works



For example, a GVRP-enabled port registers and deregisters VLAN in the following cases.

- When the port receives a VLAN attribute declaration, it registers the VLAN attribute and joins the VLAN.
- When the port receives a VLAN withdrawal, it deregisters the VLAN and leaves the VLAN.

GARP messages

A GARP participant exchanges information with other GARP participants by sending GARP messages: Join, Leave, and LeaveAll. As a GARP application, GVRP also uses GARP messages for information exchange.

- Join messages

A GARP participant sends Join messages when it wishes to declare its attribute values or receives Join messages from other GARP participants.

Join messages fall into JoinEmpty and JoinIn. A GARP participant sends JoinEmpty messages to declare attribute values that it has not registered. It sends JoinIn messages to declare attribute values that it has registered.

- Leave messages

A GARP participant sends Leave messages when it wishes to withdraw declarations of its attribute values (because, for example, it has deregistered its attribute values), or receives Leave messages from other participants.

Leave messages fall into LeaveEmpty and LeaveIn. A GARP participant sends LeaveEmpty messages to withdraw declarations of the attribute values that it has not registered. It sends LeaveIn messages to withdraw declarations of the attribute values that it has registered.

- LeaveAll messages

A GARP participant sends a LeaveAll message when it declares that it is deregistering all attribute values or receives LeaveAll messages from other participants. If any participants want to maintain the registration for a particular attribute value, they must send a Join message.

GARP timers

HP's implementation of GARP uses the following timers to control GARP message transmission:

- 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.

- Join timer

A GARP participant might declare an attribute twice to ensure reliable transmission. The Join timer sets the interval between the two declarations.

A GARP participant starts a Join timer when it declares an attribute value or receives a JoinIn message for the attribute value. If the GARP participant does not receive any declaration for the attribute value when the Join timer expires, it re-declares the attribute value.

Because all attributes of a GARP participant share the same Join timer, you must set the Join timer long enough so that all attributes can be sent out in one declaration.

- Leave timer

A GARP participant starts a Leave timer when it receives a Leave message for an attribute value. If the GARP participant receives no Join message for the attribute value before the timer expires, it deregisters the attribute value.

- 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.

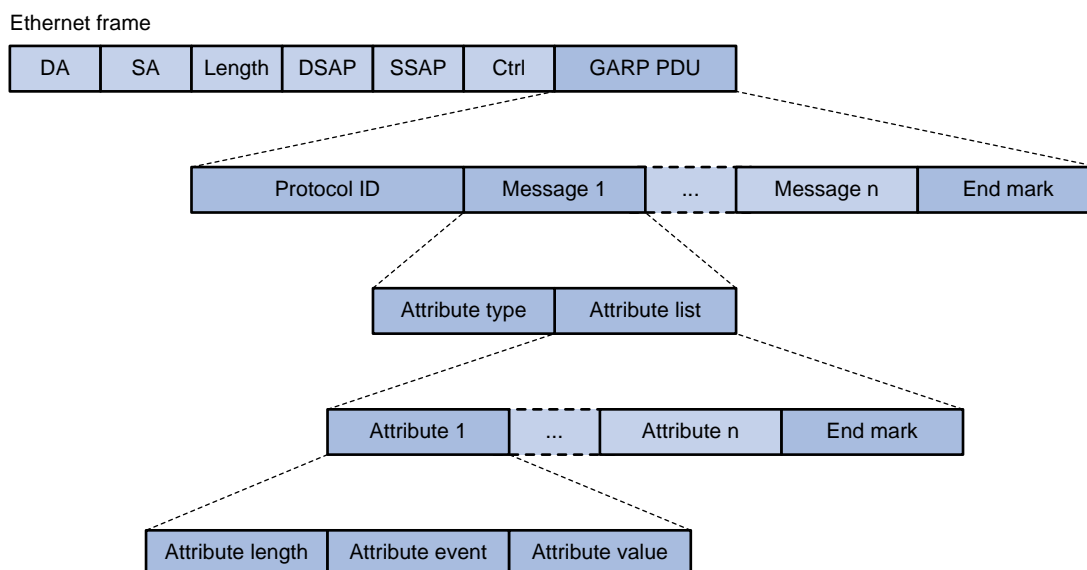
When you configure GARP timers, follow these guidelines:

- The settings of GARP timers apply to all GARP applications, such as GVRP, on a LAN.

- On a GARP-enabled network, each port maintains its own Hold, Join, and Leave timers, but only one LeaveAll timer is maintained on each device. This LeaveAll timer applies to all ports on the device.
- The value ranges for the Hold, Join, Leave, and LeaveAll timers are dependent on one another. See [Table 19](#) for their dependencies.
- Set the LeaveAll timer greater than any Leave timer and not smaller than its default value, 1000 centiseconds. Each time a LeaveAll timer expires, a network-wide re-join occurs.
- A device can send LeaveAll messages at the interval set by its LeaveAll timer or the LeaveAll timer of another device on the network, whichever is smaller. This is because each time a device on the network receives a LeaveAll message, it resets its LeaveAll timer.

GARP PDU format

Figure 55 GARP PDU format



As shown in [Figure 55](#), GARP PDUs are encapsulated in IEEE 802.3 Ethernet frames.

Table 18 GARP PDU fields

Field	Description	Value
Protocol ID	Protocol identifier for GARP	0x0001
Message	One or multiple messages, each of which contains an attribute type and an attribute list	N/A
End mark	Indicates the end of a GARP PDU	0x00
Attribute type	Defined by the GARP application	0x01 for GVRP, which indicates the VLAN ID attribute
Attribute list	Contains one or multiple attributes	N/A
Attribute	Consists of an attribute length, an attribute event, and an attribute value	N/A
Attribute length	Length of an attribute, inclusive of the attribute length field	2 to 255 (in bytes)

Field	Description	Value
Attribute event	Event that the attribute describes	<ul style="list-style-type: none"> • 0x00—LeaveAll event • 0x01—JoinEmpty event • 0x02—JoinIn event • 0x03—LeaveEmpty event • 0x04—LeaveIn event • 0x05—Empty event
		VLAN ID for GVRP
Attribute value	Attribute value	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.

GVRP

GVRP overview

As a GARP application, GVRP uses the operating mechanism of GARP to maintain and propagate dynamic VLAN registrations throughout a switched LAN.

In a switched LAN, each GVRP-enabled switch sends and receives VLAN declarations and withdrawals from other GVRP-enabled switches, and dynamically updates its local database, including active VLAN members and through which port each VLAN member can be reached. This makes sure that all GVRP-enabled switches in a LAN maintain the same VLAN information.

The VLAN information propagated by GVRP includes not only manually configured static VLAN information but also dynamic VLAN information from other switches.

GVRP registration modes

GVRP is available on trunk ports. It provides the following registration modes:

- **Normal mode**—Performs dynamic VLAN registrations and deregistrations on the trunk port, and sends declarations and withdrawals for dynamic and static VLANs. VLANs manually configured are called static VLANs, and VLANs created by GVRP are called dynamic VLANs.
- **Fixed mode**—Disables the trunk port to register or withdraw dynamic VLAN information, but allows the port to send declarations for static VLANs. A trunk port in this mode carries only static VLANs, even if it has been assigned to all VLANs.
- **Forbidden mode**—Disables the trunk port to register or withdraw dynamic VLAN information, and allows the port to send declarations only for VLAN 1. A trunk port in this mode carries only VLAN 1 even if it has been assigned to any other VLANs.

Protocols and standards

IEEE 802.1Q, *Virtual Bridged Local Area Networks*

GVRP configuration task list

When you configure GVRP, follow these guidelines:

- GVRP configuration made in Ethernet interface view or Layer 2 aggregate interface view takes effect on the current interface only; GVRP configuration made in port group view takes effect on all the member ports in the group.
- GVRP configuration made on a member port in an aggregation group takes effect only after the port is removed from the aggregation group.

Complete these tasks to configure GVRP:

Task	Remarks
Configuring GVRP functions	Required
Configuring the GARP timers	Optional

Configuring GVRP functions

Before enabling GVRP on a port, you must enable GVRP globally. In addition, you can configure GVRP only on trunk ports, and you must assign the involved trunk ports to all dynamic VLANs.

Configuration restrictions and guidelines

- GVRP can work with STP, RSTP, or MSTP CIST but not PVST. When GVRP runs on the CIST, blocked ports on the CIST cannot receive or send GVRP packets. For more information about STP, RSTP, MSTP CIST, and PVST, see "Configuring spanning tree protocols."
- 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 *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.

Configuration procedure

To configure GVRP functions on a trunk port:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable GVRP globally.	gvrp	Globally disabled by default.
3. Enter Ethernet interface view, Layer 2 aggregate interface view, or port group view.	<ul style="list-style-type: none"> • 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> 	Use either command.

Step	Command	Remarks
4. Configure the link type of the ports as trunk.	port link-type trunk	Access by default. For more information about the port link-type trunk command, see <i>Layer 2—LAN Switching Command Reference</i> .
5. Assign the trunk ports to all VLANs.	port trunk permit vlan all	By default, a trunk port is assigned to VLAN 1 only. For more information about the port trunk permit vlan all command, see <i>Layer 2—LAN Switching Command Reference</i> .
6. Enable GVRP on the ports.	gvrp	Disabled by default.
7. Configure the GVRP registration mode on the port.	gvrp registration { fixed forbidden normal }	Optional. normal by default.

Configuring the GARP timers

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the GARP LeaveAll timer.	garp timer leaveall timer-value	Optional. 1000 centiseconds by default. The LeaveAll timer applies to all ports.
3. Enter Ethernet interface view, Layer 2 aggregate interface view, or port group view.	<ul style="list-style-type: none"> Enter Ethernet interface view or Layer 2 aggregate interface view: interface interface-type interface-number Enter port group view: port-group manual port-group-name 	Use either command.
4. Configure the Hold timer.	garp timer hold timer-value	Optional. 10 centiseconds by default.
5. Configure the Join timer.	garp timer join timer-value	Optional. 20 centiseconds by default.
6. Configure the Leave timer.	garp timer leave timer-value	Optional. 60 centiseconds by default.

As shown in [Table 19](#), the value ranges for GARP timers are dependent on one another; use the following guidelines to configure GARP timers:

- If you want to set a value beyond the value range for a timer, you can change the value range by tuning the value of another related timer.

- If you want to restore the default settings of the timers, restore the Hold timer first, followed by the Join, Leave, and LeaveAll timers.

Table 19 Dependencies of the GARP timers

Timer	Lower limit	Upper limit
Hold	10 centiseconds	No greater than half of the Join timer
Join	No less than twice the Hold timer	Less than half of the Leave timer
Leave	Greater than twice the Join timer	Less than the LeaveAll timer
LeaveAll	Greater than the Leave timer	32,765 centiseconds

NOTE:

To keep the dynamic VLANs learned through GVRP stable, do not set the LeaveAll timer smaller than its default value, 1000 centiseconds.

Displaying and maintaining GVRP

Task	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 that GVRP maintains 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</i> vlan <i>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 56](#):

- Device A and Device B are connected through their ports GigabitEthernet 1/0/1.
- 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 56 Network diagram



Configuration procedure

1. Configure Device A:

Enable GVRP globally.

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

```
[DeviceA] gvrp
```

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

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

```
[DeviceA-GigabitEthernet1/0/1] port link-type trunk
```

```
[DeviceA-GigabitEthernet1/0/1] port trunk permit vlan all
```

Enable GVRP on trunk port GigabitEthernet 1/0/1.

```
[DeviceA-GigabitEthernet1/0/1] gvrp
```

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

Create VLAN 2 (a static VLAN).

```
[DeviceA] vlan 2
```

```
[DeviceA-vlan2] quit
```

2. Configure Device B:

Enable GVRP globally.

```
<DeviceB> system-view
```

```
[DeviceB] gvrp
```

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

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

```
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
```

```
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan all
```

Enable GVRP on trunk port GigabitEthernet 1/0/1.

```
[DeviceB-GigabitEthernet1/0/1] gvrp
```

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

Create VLAN 3 (a static VLAN).

```
[DeviceB] vlan 3
```

```
[DeviceB-vlan3] quit
```

3. Verify the configuration:

Use the **display gvrp local-vlan** command to display the local VLAN information that GVRP maintains on ports. For example:

Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device A.

```
[DeviceA] display gvrp local-vlan interface gigabitethernet 1/0/1
```

```
Following VLANs exist in GVRP local database:
```

```
1(default),2-3
```

According to the output, information about VLAN 1, static VLAN information of VLAN 2 on the local device, and dynamic VLAN information of VLAN 3 on Device B are all registered through GVRP.

Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device B.

```
[DeviceB] display gvrp local-vlan interface gigabitethernet 1/0/1
```

Following VLANs exist in GVRP local database:

```
1(default),2-3
```

According to the output, information about VLAN 1, static VLAN information of VLAN 3 on the local device, and dynamic VLAN information of VLAN 2 on Device A are all registered through GVRP.


GVRP fixed registration mode configuration example

Network requirements

As shown in Figure 57:

- Device A and Device B are connected through their ports GigabitEthernet 1/0/1.
- 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 57 Network diagram



Configuration procedure

1. Configure Device A:

Enable GVRP globally.

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

```
[DeviceA] gvrp
```

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

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

```
[DeviceA-GigabitEthernet1/0/1] port link-type trunk
```

```
[DeviceA-GigabitEthernet1/0/1] port trunk permit vlan all
```

Enable GVRP on GigabitEthernet 1/0/1 and set the GVRP registration mode to fixed on the port.

```
[DeviceA-GigabitEthernet1/0/1] gvrp
```

```
[DeviceA-GigabitEthernet1/0/1] gvrp registration fixed
```

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

Create VLAN 2 (a static VLAN).

```
[DeviceA] vlan 2
```

```
[DeviceA-vlan2] quit
```

2. Configure Device B:

Enable GVRP globally.

```
<DeviceB> system-view
```

```
[DeviceB] gvrp
# Configure port GigabitEthernet 1/0/1 as a trunk port, and assign it to all VLANs.
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan all
# Enable GVRP on GigabitEthernet 1/0/1, and set the GVRP registration mode to fixed on the port.
[DeviceB-GigabitEthernet1/0/1] gvrp
[DeviceB-GigabitEthernet1/0/1] gvrp registration fixed
[DeviceB-GigabitEthernet1/0/1] quit
# Create VLAN 3 (a static VLAN).
[DeviceB] vlan 3
[DeviceB-vlan3] quit
```

3. Verify the configuration:

Use the **display gvrp local-vlan** command to display the local VLAN information that GVRP maintains on ports. For example:

Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device A.

```
[DeviceA] display gvrp local-vlan interface gigabitethernet 1/0/1
Following VLANs exist in GVRP local database:
1(default), 2
```

According to the output, information about VLAN 1 and static VLAN information of VLAN 2 on the local device are registered through GVRP, but dynamic VLAN information of VLAN 3 on Device B is not.

Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device B.

```
[DeviceB] display gvrp local-vlan interface gigabitethernet 1/0/1
Following VLANs exist in GVRP local database:
1(default), 3
```

According to the output, information about VLAN 1 and static VLAN information of VLAN 3 on the local device are registered through GVRP, but dynamic VLAN information of VLAN 2 on Device A is not.


GVRP forbidden registration mode configuration example

Network requirements

As shown in [Figure 58](#):

- Device A and Device B are connected through their ports GigabitEthernet 1/0/1.
- 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 58 Network diagram



Configuration procedure

1. Configure Device A:

Enable GVRP globally.

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

```
[DeviceA] gvrp
```

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

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

```
[DeviceA-GigabitEthernet1/0/1] port link-type trunk
```

```
[DeviceA-GigabitEthernet1/0/1] port trunk permit vlan all
```

Enable GVRP on GigabitEthernet 1/0/1, and set the GVRP registration mode to forbidden on the port.

```
[DeviceA-GigabitEthernet1/0/1] gvrp
```

```
[DeviceA-GigabitEthernet1/0/1] gvrp registration forbidden
```

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

Create VLAN 2 (a static VLAN).

```
[DeviceA] vlan 2
```

```
[DeviceA-vlan2] quit
```

2. Configure Device B:

Enable GVRP globally.

```
<DeviceB> system-view
```

```
[DeviceB] gvrp
```

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

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

```
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
```

```
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan all
```

Enable GVRP on GigabitEthernet 1/0/1, and set the GVRP registration mode to forbidden on the port.

```
[DeviceB-GigabitEthernet1/0/1] gvrp
```

```
[DeviceB-GigabitEthernet1/0/1] gvrp registration forbidden
```

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

Create VLAN 3 (a static VLAN).

```
[DeviceB] vlan 3
```

```
[DeviceB-vlan3] quit
```

3. Verify the configuration:

Use the **display gvrp local-vlan** command to display the local VLAN information that GVRP maintains on ports. For example:

Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device A.

```
[DeviceA] display gvrp local-vlan interface gigabitethernet 1/0/1
```

Following VLANs exist in GVRP local database:

```
1(default)
```

According to the output, information about VLAN 1 is registered through GVRP, but static VLAN information of VLAN 2 on the local device and dynamic VLAN information of VLAN 3 on Device B are not.

Display the local VLAN information that GVRP maintains on port GigabitEthernet 1/0/1 of Device B.

```
[DeviceB] display gvrp local-vlan interface gigabitethernet 1/0/1
```

Following VLANs exist in GVRP local database:

```
1(default)
```

According to the output, information about VLAN 1 is registered through GVRP, but static VLAN information of VLAN 3 on the local device and dynamic VLAN information of VLAN 2 on Device A are not.

Configuring QinQ

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

Overview

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 device 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 device supports a maximum of 4094 VLANs. This is far from enough for isolating users in actual networks, especially in metropolitan area networks (MANs).


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

- Releases the stress on the SVLAN resource.
- Enables customers to plan their CVLANs without conflicting with SVLANs.
- Provides an easy-to-implement Layer 2 VPN solution for small-sized MANs or intranets.
- Enables the customers to keep their VLAN assignment schemes unchanged when the service provider upgrades the service provider network.

How QinQ works

The devices in the public network forward a frame only according to its outer VLAN tag and obtain its source MAC address into the MAC address table of the outer VLAN. The inner VLAN tag of the frame is transmitted as the payload.

Figure 59 Typical QinQ application scenario



As shown in [Figure 59](#), customer network A has CVLANs 1 through 10, and customer network B has CVLANs 1 through 20. The service provider assigns SVLAN 3 for customer network A, and assigns SVLAN 4 for customer network B.

When a tagged Ethernet frame from customer network A arrives at a provider edge device (PE), the PE tags the frame with outer VLAN 3. When a tagged Ethernet frame from customer network B arrives at a PE, the PE tags the frame with outer VLAN 4. There is no overlap of VLAN IDs among customers, and traffic from different customers can be identified separately.

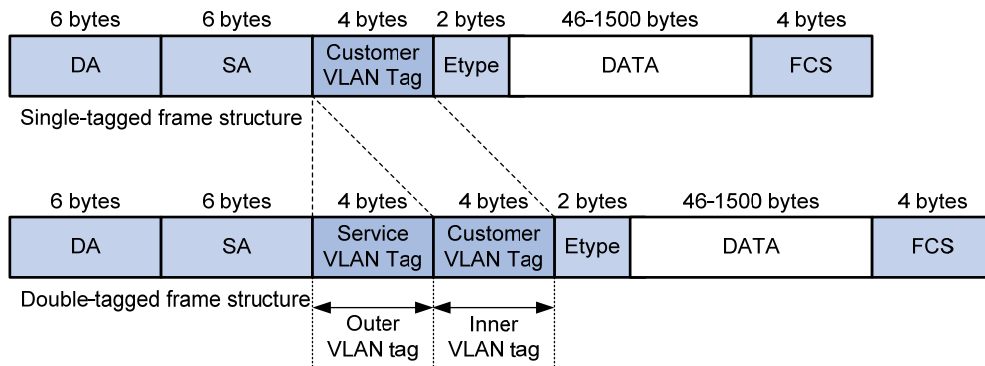
The double-tagged Ethernet frame is then transmitted over the service provider network and arrives at the other PE. The PE removes the SVLAN of the frame before sending it to the target customer edge device (CE).

QinQ frame structure

A QinQ frame is transmitted double-tagged over the service provider network. As shown in [Figure 60](#), 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.

QinQ uses CVLAN tags to transmit frames over the private network, and uses SVLAN tags to transmit frames over the public network. When a QinQ frame is transmitted over the public network, its CVLAN tag is transmitted as the payload.

Figure 60 Single-tagged Ethernet frame header and double-tagged Ethernet frame header



The default maximum transmission unit (MTU) of an interface is 1500 bytes. The size of an outer VLAN tag is 4 bytes. HP recommends you to increase the MTU of each interface on the service provider network to at least 1504 bytes.

Implementations of QinQ

HP provides the following QinQ implementations: basic QinQ and selective QinQ.


- **Basic QinQ**
Basic QinQ enables a port to tag any incoming frames with its port VLAN ID (PVID) 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 PVID tag.
- **Selective QinQ**
Selective QinQ is more flexible than basic QinQ. In addition to all the functions of basic QinQ, selective QinQ enables a port to perform the following per-CVLAN actions for incoming frames:
 - Tag frames from different CVLANs with different SVLAN tags.
 - Mark the outer VLAN 802.1p priority based on the existing inner VLAN 802.1p priority.
 Besides being able to separate the service provider network from the customer networks, selective QinQ provides abundant service features and enables more flexible networking.

Modifying the TPID in a VLAN tag

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

Figure 61 shows the 802.1Q-defined tag structure of an Ethernet frame.

Figure 61 VLAN tag structure of an Ethernet frame



The switch determines whether a received frame carries a VLAN tag by checking the TPID value. For example, if a frame carries a VLAN tag with TPID value 0x8100, but the configured TPID value is 0x9100, the switch considers that the frame does not carry any VLAN tag.

Devices of different vendors may set the TPID of the outer VLAN tag of QinQ frames to different values. For compatibility with these devices, modify the TPID value so that the QinQ frames, when sent to the public network, carry the TPID value identical to the value of a particular vendor, allowing interoperability with the devices of that vendor.

The TPID in an Ethernet frame has the same position as the protocol type field in a frame without a VLAN tag. To avoid problems in packet forwarding and handling in the network, do not set the TPID value to any of the reserved values.

Table 20 Reserved protocol type values

Protocol type	Value
ARP	0x0806
PUP	0x0200
RARP	0x8035
IP	0x0800
IPv6	0x86DD
PPPoE	0x8863/0x8864
MPLS	0x8847/0x8848
IPX/SPX	0x8137
IS-IS	0x8000
LACP	0x8809
802.1X	0x888E
Cluster	0x88A7
Reserved	0xFFFD/0xFFFE/0xFFFF

Protocols and standards

IEEE 802.1Q: *IEEE standard for local and metropolitan area networks: Virtual Bridged Local Area Networks*

QinQ configuration task list

When you configure QinQ, follow these guidelines:

- QinQ requires configurations only on the service provider network.
- QinQ configurations made in Ethernet interface view take effect on the current interface only. Those made in Layer 2 aggregate interface view take effect on the current aggregate interface and all the member ports in the aggregation group. Those made in port group view take effect on all member ports in the current port group.
- Do not configure QinQ on a reflector port. For more information about reflector ports, see *Network Management and Monitoring Configuration Guide*.

- On a port with QinQ enabled, you must configure the port to allow packets from the inner and outer VLANs of QinQ packets to pass through.

Complete the follows tasks to configure QinQ:

Task		Remarks
Configuring basic QinQ	Enabling basic QinQ	Required
	Configuring VLAN transparent transmission	Optional
Configuring selective QinQ	Configuring an outer VLAN tagging policy	Required
	Configuring an inner-outer VLAN 802.1p priority mapping	Perform at least one of these tasks.
Configuring the TPID value in VLAN tags		Optional

Configuring basic QinQ

Enabling basic QinQ

A basic QinQ-enabled port tags an incoming packet with its PVID tag.

To enable basic QinQ:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
3. Enable basic QinQ.	qinq enable	Disabled by default.

Configuring VLAN transparent transmission

When basic QinQ is enabled on a port, all packets passing through the port are tagged with the port's PVID tag. However, by configuring the VLAN transparent transmission function on a port, you can specify the port not to add its PVID 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.

Configuration restrictions and guidelines

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

Configuration procedure

To configure VLAN transparent transmission:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
3. Configure the link type of the ports.	port link-type { hybrid trunk }	N/A
4. Configure the ports to allow packets from inner and outer VLANs of QinQ packets 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 } 	Use either command.
5. Enable basic QinQ on the ports.	qinq enable	By default, basic QinQ is disabled on ports.
6. Configure VLAN transparent transmission on the ports.	qinq transparent-vlan <i>vlan-list</i>	By default, VLAN transparent transmission is not configured.

Configuring selective QinQ

Configuring an outer VLAN tagging policy

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

Before enabling selective QinQ on a port, enable basic QinQ on the port first. When both features are enabled on the port, frames that meet the selective QinQ condition are handled with selective QinQ on this port first, and the left frames are handled with basic QinQ.

To configure an outer VLAN tagging policy by using the port-based method:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use either command.
3. Enable basic QinQ.	qinq enable	Disabled by default.
4. Enter QinQ view and configure the SVLAN tag for the port to add.	qinq vid <i>vlan-id</i>	By default, the SVLAN tag to be added is the PVID tag of the receiving port.

Step	Command	Remarks
5. Tag frames of the specified CVLANs with the current SVLAN.	raw-vlan-id inbound { all <i>vlan-list</i> }	N/A

NOTE:

- An inner VLAN tag corresponds to only one outer VLAN tag.
- To change an outer VLAN tag, you must delete the old outer VLAN tag configuration and configure a new outer VLAN tag.

Configuring an inner-outer VLAN 802.1p priority mapping

Through QoS policies, the HP 5120 EI switches achieve the following inner-outer VLAN 802.1p priority mapping modes:

- Marking the 802.1p priorities in outer VLAN tags according to the inner VLAN IDs or the 802.1p priorities in the inner VLAN tags.
- Copying the 802.1p priority in the inner VLAN tags to the outer VLAN tags.

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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Create a class and enter class view.	traffic classifier <i>classifier-name</i> [operator { and or }]	By default, the operator of a class is AND.
3. Configure a match criterion.	<ul style="list-style-type: none"> • Match the specified inner VLAN IDs: if-match customer-vlan-id <i>vlan-id-list</i> • Match the specified inner VLAN tag priorities: if-match customer-dot1p <i>8021p-list</i> 	Use either command.
4. Return to system view.	quit	N/A
5. Create a traffic behavior and enter traffic behavior view.	traffic behavior <i>behavior-name</i>	N/A
6. Configure a marking action or an inner-to-outer tag priority copying action.	<ul style="list-style-type: none"> • Mark the 802.1p priorities in outer VLAN tags: remark dot1p <i>8021p</i> • Copy the inner 802.1p priorities to outer 802.1p priorities: remark dot1p customer-dot1p-trust 	Use either command.
7. Return to system view.	quit	N/A
8. Create a QoS policy and enter QoS policy view.	qos policy <i>policy-name</i>	N/A
9. Associate the traffic class with the traffic behavior defined earlier.	classifier <i>classifier-name</i> behavior <i>behavior-name</i>	N/A
10. Return to system view.	quit	N/A

Step	Command	Remarks
11. Enter Ethernet interface view or port group view of the customer network-side port.	<ul style="list-style-type: none"> Enter Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> Enter port group view: port-group manual <i>port-group-name</i> 	Use either command.
12. Enable basic QinQ.	qinq enable	N/A
13. Apply the QoS policy to the incoming traffic.	qos apply policy <i>policy-name</i> inbound	N/A

Configuring the TPID value in VLAN tags

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Configure the TPID value.	qinq ethernet-type <i>hex-value</i>	<p>Optional.</p> <p>By default, the TPID value is 0x8100.</p> <p>The configuration applies to all ports.</p>

NOTE:

The TPID value configured on the HP 5120 EI Switch Series applies to both the CVLAN tags and the SVLAN tags.

QinQ configuration examples

Basic QinQ configuration example

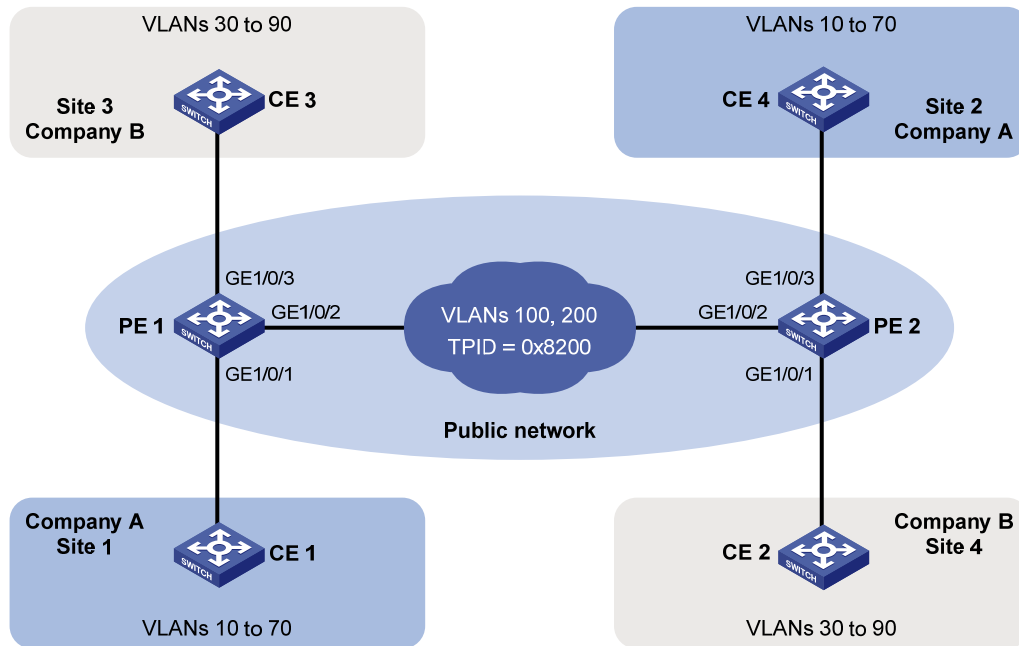
Network requirements

As shown in [Figure 62](#):

- The two branches of Company A, Site 1 and Site 2, are connected through the service provider network and use CVLANs 10 through 70. The two branches of Company B, Site 3 and Site 4, are connected through the service provider network and use CVLANs 30 through 90.
- PE 1 and PE 2 are edge devices on the service provider network and are connected through third-party devices with a TPID value of 0x8200.

Configure the edge and third-party devices to enable communication between the branches of Company A through SVLAN 100, and communication between the branches of Company B through SVLAN 200.

Figure 62 Network diagram



Configuration procedure



IMPORTANT:

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

1. Configure PE 1:

- Configure GigabitEthernet 1/0/1:

Configure GigabitEthernet 1/0/1 as a trunk port and assign it to VLAN 100 and VLANs 10 through 70..

```
<PE1> system-view
[PE1] interface gigabitethernet 1/0/1
[PE1-GigabitEthernet1/0/1] port link-type trunk
[PE1-GigabitEthernet1/0/1] port trunk permit vlan 100 10 to 70
```

Configure VLAN 100 as the PVID for the port.

```
[PE1-GigabitEthernet1/0/1] port trunk pvid vlan 100
```

Enable basic QinQ on the port.

```
[PE1-GigabitEthernet1/0/1] qinq enable
[PE1-GigabitEthernet1/0/1] quit
```

- Configure GigabitEthernet 1/0/2:

Configure GigabitEthernet 1/0/2 as a trunk port and assign it to VLAN 100 and VLAN 200.

```
[PE1] interface gigabitethernet 1/0/2
[PE1-GigabitEthernet1/0/2] port link-type trunk
[PE1-GigabitEthernet1/0/2] port trunk permit vlan 100 200
[PE1-GigabitEthernet1/0/2] quit
```

Set the TPID value in the outer VLAN tag to 0x8200 on the port.

```
[PE1] qinq ethernet-type 8200
```

- Configure GigabitEthernet 1/0/3:
 # Configure GigabitEthernet 1/0/3 as a trunk port and assign it to VLAN 200 and VLANs 30 through 90.

```
[PE1] interface gigabitethernet 1/0/3
[PE1-GigabitEthernet1/0/3] port link-type trunk
[PE1-GigabitEthernet1/0/3] port trunk permit vlan 200 30 to 90
```

 # Configure VLAN 200 as the PVID for the port.

```
[PE1-GigabitEthernet1/0/3] port trunk pvid vlan 200
```

 # Enable basic QinQ on the port.

```
[PE1-GigabitEthernet1/0/3] qinq enable
[PE1-GigabitEthernet1/0/3] quit
```

2. Configure PE 2:

- Configure GigabitEthernet 1/0/1:
 # Configure GigabitEthernet 1/0/1 as a trunk port and assign it to VLAN 200 and VLANs 30 through 90.

```
<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 200 30 to 90
```

 # Configure VLAN 200 as the PVID for the port.

```
[PE2-GigabitEthernet1/0/1] port trunk pvid vlan 200
```

 # Enable basic QinQ on the port.

```
[PE2-GigabitEthernet1/0/1] qinq enable
[PE2-GigabitEthernet1/0/1] quit
```
- Configure GigabitEthernet 1/0/2:
 # Configure GigabitEthernet 1/0/2 as a trunk port and assign it to VLAN 100 and VLAN 200.

```
[PE2] interface gigabitethernet 1/0/2
[PE2-GigabitEthernet1/0/2] port link-type trunk
[PE2-GigabitEthernet1/0/2] port trunk permit vlan 100 200
[PE2-GigabitEthernet1/0/2] quit
```

 # Set the TPID value in the outer VLAN tag to 0x8200 on the port.

```
[PE2] qinq ethernet-type 8200
```
- Configure GigabitEthernet 1/0/3:
 # Configure GigabitEthernet 1/0/3 as a trunk port and assign it to VLAN 100 and VLANs 10 through 70.

```
[PE2] interface gigabitethernet 1/0/3
[PE2-GigabitEthernet1/0/3] port link-type trunk
[PE2-GigabitEthernet1/0/3] port trunk permit vlan 100 10 to 70
```

 # Configure VLAN 100 as the PVID for the port.

```
[PE2-GigabitEthernet1/0/3] port trunk pvid vlan 100
```

 # Enable basic QinQ on the port.

```
[PE2-GigabitEthernet1/0/3] qinq enable
[PE2-GigabitEthernet1/0/3] quit
```


3. On the third-party devices between PE 1 and PE 2, configure the port that connects to PE 1 and that connecting to PE 2 to allow tagged frames of VLAN 100 and VLAN 200 to pass through. (Details not shown.)

Selective QinQ configuration example

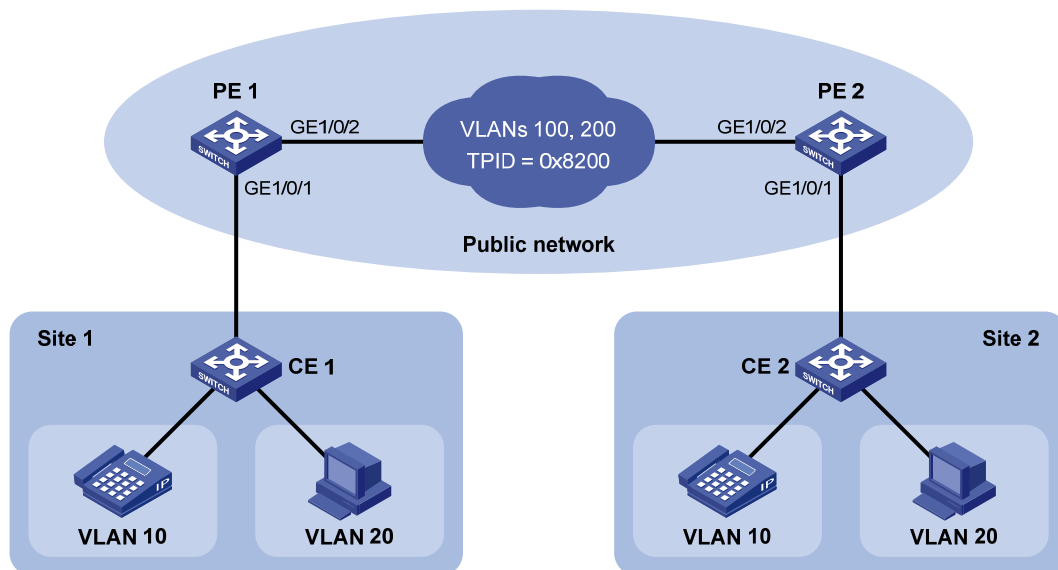
Network requirements

As shown in Figure 63:

- The two branches of a company, Site 1 and Site 2, are connected through the service provider network and use CVLAN 10 and CVLAN 20 to transmit voice traffic and data traffic separately.
- PE 1 and PE 2 are edge devices on the service provider network and are connected through third-party devices with a TPID value of 0x8200.

Configure the edge and third-party devices to allow frames from CVLAN 10 to be transmitted between the branches via SVLAN 100 and frames from CVLAN 20 to be transmitted between the branches via SVLAN 200.

Figure 63 Network diagram



Configuration procedure

! IMPORTANT:

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

1. Configure PE 1:
 - Configure GigabitEthernet 1/0/1:
Configure GigabitEthernet 1/0/1 as a hybrid port to permit frames of VLAN 10 and VLAN 20 to pass through tagged, and frames of VLAN 100 and VLAN 200 to pass through untagged.

```
<PE1> system-view
[PE1] interface gigabitethernet 1/0/1
[PE1-GigabitEthernet1/0/1] port link-type hybrid
```

```
[PE1-GigabitEthernet1/0/1] port hybrid vlan 10 20 tagged
[PE1-GigabitEthernet1/0/1] port hybrid vlan 100 200 untagged
# Enable basic QinQ on the port.
[PE1-GigabitEthernet1/0/1] qinq enable
# Configure the port to tag VLAN 10 frames with outer VLAN ID 100.
[PE1-GigabitEthernet1/0/1] qinq vid 100
[PE1-GigabitEthernet1/0/1-vid-100] raw-vlan-id inbound 10
[PE1-GigabitEthernet1/0/1-vid-100] quit
# Configure the port to tag VLAN 20 frames with outer VLAN ID 200.
[PE1-GigabitEthernet1/0/1] qinq vid 200
[PE1-GigabitEthernet1/0/1-vid-200] raw-vlan-id inbound 20
[PE1-GigabitEthernet1/0/1-vid-200] quit
[PE1-GigabitEthernet1/0/1] quit
```

- Configure GigabitEthernet 1/0/2:

```
# Configure GigabitEthernet 1/0/2 as a trunk port and assign it to VLAN 100 and VLAN 200.
[PE1] interface gigabitethernet 1/0/2
[PE1-GigabitEthernet1/0/2] port link-type trunk
[PE1-GigabitEthernet1/0/2] port trunk permit vlan 100 200
[PE1-GigabitEthernet1/0/2] quit
# Set the TPID in the outer VLAN tags to 0x8200.
[PE1] qinq ethernet-type 8200
```

2. Configure PE 2:

- Configure GigabitEthernet 1/0/1:

```
# Configure GigabitEthernet 1/0/1 as a hybrid port to permit frames of VLAN 10 and VLAN 20 to pass through tagged, and frames of VLAN 100 and VLAN 200 to pass through untagged.
```

```
[PE2] interface gigabitethernet 1/0/1
[PE2-GigabitEthernet1/0/1] port link-type hybrid
[PE2-GigabitEthernet1/0/1] port hybrid vlan 10 20 tagged
[PE2-GigabitEthernet1/0/1] port hybrid vlan 100 200 untagged
```

```
# Enable basic QinQ on the port.
```

```
[PE2-GigabitEthernet1/0/1] qinq enable
```

```
# Configure the port to tag VLAN 10 frames with outer VLAN ID 100.
```

```
[PE2-GigabitEthernet1/0/1] qinq vid 100
[PE2-GigabitEthernet1/0/1-vid-100] raw-vlan-id inbound 10
[PE2-GigabitEthernet1/0/1-vid-100] quit
```

```
# Configure the port to tag VLAN 20 frames with outer VLAN ID 200.
```

```
[PE2-GigabitEthernet1/0/1] qinq vid 200
[PE2-GigabitEthernet1/0/1-vid-200] raw-vlan-id inbound 20
[PE2-GigabitEthernet1/0/1-vid-200] quit
[PE2-GigabitEthernet1/0/1] quit
```

- Configure GigabitEthernet 1/0/2:

```
# Configure GigabitEthernet 1/0/2 as a trunk port and assign it to VLAN 100 and VLAN 200.
```

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

```
[PE2-GigabitEthernet1/0/2] port trunk permit vlan 100 200
[PE2-GigabitEthernet1/0/2] quit
# Set the TPID in the outer VLAN tags to 0x8200.
[PE2] qinq ethernet-type 8200
```

3. On the third-party devices between PE 1 and PE 2, configure the port that connects to PE 1 and that connecting to PE 2 to allow tagged frames of VLAN 100 and VLAN 200 to pass through. (Details not shown.)

Configuring LLDP

Overview

Background

In a heterogeneous network, a standard configuration exchange platform ensures that different types of network devices from different vendors can discover one another and exchange configuration for the sake of interoperability and management.

The Link Layer Discovery Protocol (LLDP) is specified 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 LLDP Data Units (LLDPDUs) to the directly connected devices. At the same time, the device stores the device information received in LLDPDUs sent from the LLDP neighbors in a standard management information base (MIB). For more information about MIBs, see *Network Management and Monitoring Configuration Guide*. LLDP enables a network management system to quickly detect and identify Layer 2 network topology changes.

Basic concepts

LLDPDU formats

LLDP sends device information in LLDPDUs. LLDPDUs are encapsulated in Ethernet II or Subnetwork Access Protocol (SNAP) frames.

1. Ethernet II-encapsulated LLDPDU format

Figure 64 Ethernet II-encapsulated LLDPDU format

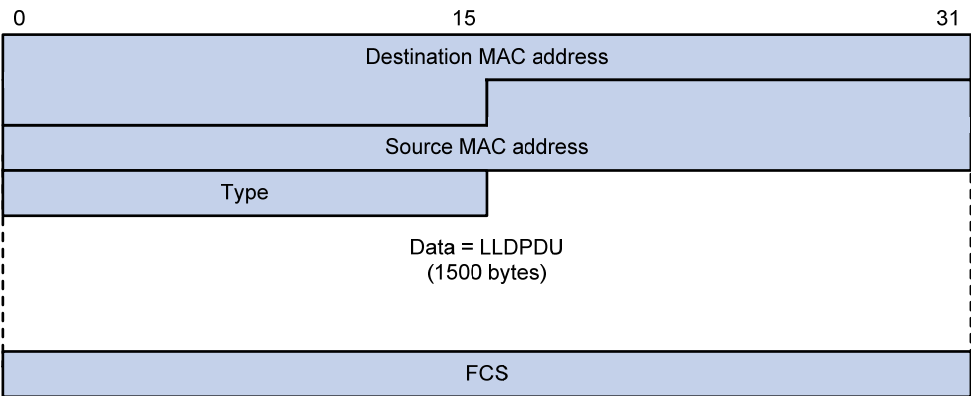


Table 21 Fields in an Ethernet II-encapsulated LLDPDU

Field	Description
Destination MAC address	MAC address to which the LLDPDU is advertised. It is fixed to 0x0180-C200-000E, a multicast MAC address.

Field	Description
Source MAC address	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	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 65 SNAP-encapsulated LLDPDU format

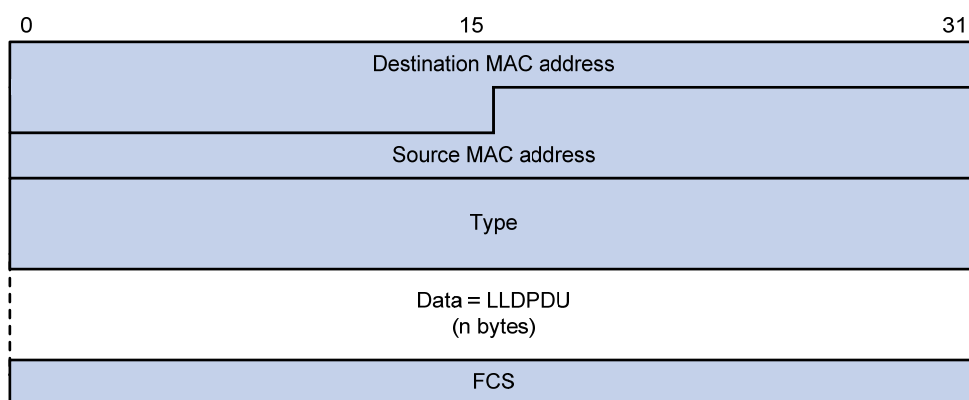



Table 22 Fields in a SNAP-encapsulated LLDPDU

Field	Description
Destination MAC address	MAC address to which the LLDPDU is advertised. It is fixed at 0x0180-C200-000E, a multicast MAC address.
Source MAC address	MAC address of the sending port.
Type	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.

LLDPDU

LLDP uses LLDPDUs to exchange information. An LLDPDU comprises multiple TLV sequences. Each TLV carries a type of device information, as shown in [Figure 66](#).

Figure 66 LLDPDU encapsulation format



An LLDPDU can carry up to 28 types of TLVs. Mandatory TLVs include Chassis ID TLV, Port ID TLV, Time to Live TLV, and End of LLDPDU TLV. Other TLVs are optional.

TLVs

TLVs are type, length, and value sequences that carry information elements. The type field identifies the type of information, the length field measures the length of the information field in octets, and the value field contains the information itself.

LLDPDU TLVs fall into the following categories:

- Basic management TLVs
- Organizationally (IEEE 802.1 and IEEE 802.3) specific TLVs
- 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. Some of them are mandatory to LLDPDUs, that is, must be included in every LLDPDU.

Table 23 Basic management TLVs

Type	Description	Remarks
Chassis ID	Specifies the bridge MAC address of the sending device	Mandatory
Port ID	Specifies the ID of the sending port If the LLDPDU carries LLDP-MED TLVs, the port ID TLV carries the MAC address of the sending port. If the LLDPDU carries no LLDP-MED TLVs, the port ID TLV carries the port name.	
Time To Live	Specifies the life of the transmitted information on the receiving device	
End of LLDPDU	Marks the end of the TLV sequence in the LLDPDU	
Port Description	Specifies the port description of the sending port	Optional
System Name	Specifies the assigned name of the sending device	
System Description	Specifies the description of the sending device	
System Capabilities	Identifies the primary functions of the sending device and the enabled primary functions	
Management Address	Specifies the management address, and the interface number and object identifier (OID) associated with the address	

2. IEEE 802.1 organizationally specific TLVs

Table 24 IEEE 802.1 organizationally specific TLVs

Type	Description
Port VLAN ID	Specifies the port's VLAN identifier (PVID). An LLDPDU carries only one TLV of this type.
Port And Protocol VLAN ID	Indicates whether the device supports protocol VLANs and, if so, what VLAN IDs these protocols will be associated with. An LLDPDU can carry multiple different TLVs of this type.
VLAN Name	Specifies the textual name of any VLAN to which the port belongs. An LLDPDU can carry multiple different TLVs of this type.

Type	Description
Protocol Identity	Indicates protocols supported on the port. An LLDPDU can carry multiple different TLVs of this type.

NOTE:

HP devices support only receiving protocol identity TLVs.

3. IEEE 802.3 organizationally specific TLVs

Table 25 IEEE 802.3 organizationally specific TLVs

Type	Description
MAC/PHY Configuration/Status	Contains the bit-rate and duplex capabilities of the sending port, support for auto negotiation, enabling status of auto negotiation, and the current rate and duplex mode.
Power Via MDI	Contains the power supply capability of the port, including the Power over Ethernet (PoE) type, which can be Power Sourcing Equipment (PSE) or Powered Device (PD), PoE mode, whether PSE power supply is supported, whether PSE power supply is enabled, and whether the PoE mode is controllable.
Link Aggregation	Indicates the aggregation capability of the port (whether the link is capable of being aggregated), and the aggregation status (whether the link is in an aggregation).
Maximum Frame Size	Indicates the supported maximum frame size. It is now the maximum transmission unit (MTU) of the port.
Power Stateful Control	Indicates the power state control configured on the sending port, including the power type of the PSE or PD, PoE sourcing and receiving priority, and PoE sourcing and receiving power.

NOTE:

The Power Stateful Control TLV is defined in IEEE P802.3at D1.0. The later versions no longer support this TLV. HP devices send this type of TLVs only after receiving them.

LLDP-MED TLVs

LLDP-MED TLVs provide multiple advanced applications for voice over IP (VoIP), such as basic configuration, network policy configuration, and address and directory management. LLDP-MED TLVs provide a cost-effective and easy-to-use solution for deploying voice devices in Ethernet. LLDP-MED TLVs are shown in [Table 26](#).

Table 26 LLDP-MED TLVs

Type	Description
LLDP-MED Capabilities	Allows a network device to advertise the LLDP-MED TLVs that it supports.
Network Policy	Allows a network device or terminal device to advertise the VLAN ID of the specific port, the VLAN type, and the Layer 2 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.

Type	Description
Hardware Revision	Allows a terminal device to advertise its hardware version.
Firmware Revision	Allows a terminal device to advertise its firmware version.
Software Revision	Allows a terminal device to advertise its software version.
Serial Number	Allows a terminal device to advertise its serial number.
Manufacturer Name	Allows a terminal device to advertise its vendor name.
Model Name	Allows a terminal device to advertise its model name.
Asset ID	Allows a terminal device to advertise its asset ID. The typical case is that the user specifies the asset ID for the endpoint to facilitate directory management and asset tracking.
Location Identification	Allows a network device to advertise the appropriate location identifier information for a terminal device to use in the context of location-based applications.

Management address

The network management system uses the management address of a device to identify and manage the device for topology maintenance and network management. The management address TLV encapsulates the management address.

How LLDP works

Operating modes of LLDP

LLDP can operate in one of the following modes:

- **TxRx mode**—A port in this mode sends and receives LLDPDUs.
- **Tx mode**—A port in this mode only sends LLDPDUs.
- **Rx mode**—A port in this mode only receives LLDPDUs.
- **Disable mode**—A port in this mode does not send or receive LLDPDUs.

Each time the LLDP operating mode of a port changes, its LLDP protocol state machine re-initializes. A re-initialization delay, which is user configurable, prevents LLDP from being initialized too frequently at times of frequent changes to the operating mode. With this delay configured, before a port can initialize LLDP, it must wait for the specified interval after the LLDP operating mode changes.

Transmitting LLDPDUs

An LLDP-enabled port operating in TxRx mode or Tx mode sends LLDPDUs to its directly connected devices both periodically and when the local configuration changes. To prevent LLDPDUs from overwhelming the network during times of frequent changes to local device information, an interval is introduced between two successive LLDPDUs.

This interval is shortened to 1 second in either of the following cases:

- A new neighbor is discovered. A new LLDPDU is received and carries device information new to the local device.
- The LLDP operating mode of the port changes from Disable or Rx to TxRx or Tx.

This is the fast sending mechanism of LLDP. With this mechanism, a specific number of LLDPDUs are sent successively at 1-second intervals, to help LLDP neighbors discover the local device as soon as possible. Then, the normal LLDPDU transmit interval resumes.

Receiving LLDPDUs

An LLDP-enabled port that is operating in TxRx mode or Rx mode checks the validity of TLVs carried in every received LLDPDU. If valid, the information is saved and an aging timer is set for it based on the time to live (TTL) value in the Time to Live TLV carried in the LLDPDU. If the TTL value is zero, the information ages out immediately.

Protocols and standards

- IEEE 802.1AB-2005, *Station and Media Access Control Connectivity Discovery*
- ANSI/TIA-1057, *Link Layer Discovery Protocol for Media Endpoint Devices*

LLDP configuration task list

LLDP-related configurations made in Layer 2 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.

Complete these tasks to configure LLDP:

Task	Remarks
Enabling LLDP	Required
Setting the LLDP operating mode	Optional
Setting the LLDP re-initialization delay	Optional
Performing basic LLDP configuration	Enabling LLDP polling
	Configuring the advertisable TLVs
	Configuring the management address and its encoding format
	Setting other LLDP parameters
	Setting an encapsulation format for LLDPDUs
Configuring CDP compatibility	Optional
Configuring LLDP trapping	Optional

Performing basic LLDP configuration

Enabling LLDP

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

To enable LLDP:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable LLDP globally.	lldp enable	By default, LLDP is globally enabled.

Step	Command	Remarks
3. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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.
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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 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.
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 changes to the LLDP operating mode on a port.

To set the LLDP re-initialization delay for ports:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
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 periodically searches for local configuration changes. On detecting a configuration change, the device sends LLDPDUs to inform neighboring devices of the change.

To enable LLDP polling:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">Enter Layer 2 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.
3. Enable LLDP polling and set the polling interval.	lldp check-change-interval <i>interval</i>	Disabled by default.

Configuring the advertisable TLVs

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

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">Enter Layer 2 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.
3. Configure the advertisable TLVs.	lldp tlv-enable { basic-tlv { all port-description system-capability system-description system-name } dot1-tlv { all port-vlan-id protocol-vlan-id [<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</i> <i>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 location identification TLVs are advertisable on a Layer 2 Ethernet port.

Configuring the management address and its encoding format

LLDP encodes management addresses in numeric or character string format in management address TLVs.

By default, management addresses are encoded in numeric format. If a neighbor encoded its management address in character string format, you must configure the encoding format of the management address as string on the connecting port to guarantee normal communication with the neighbor.

To configure a management address to be advertised and its encoding format on a port or group of ports:

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

Step	Command	Remarks
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none"> Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i> Enter port group view: port-group manual <i>port-group-name</i> 	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. 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.
4. Configure the encoding format of the management address as character string.	lldp management-address-format string	Optional. By default, the management address is encapsulated in the numeric format.

Setting other LLDP parameters

The Time to Live TLV carried in an LLDPDU determines how long the device information carried in the LLDPDU can be saved on a recipient device.

By setting the TTL multiplier, you can configure the TTL of locally sent LLDPDUs, which determines how long information about the local device can be saved on a neighboring device. The TTL is expressed by using the following formula:

$$\text{TTL} = \text{Min} (65535, (\text{TTL multiplier} \times \text{LLDPDU transmit interval}))$$

As the expression shows, the TTL can be up to 65535 seconds. TTLs greater than 65535 will be rounded down to 65535 seconds.

Configuration restrictions and guidelines

- To make sure that LLDP neighbors can receive LLDPDUs to update information about the current device before it ages out, configure both the LLDPDU transmit interval and delay to be less than the TTL.
- It is a good practice to set the LLDPDU transmit interval to be no less than four times the LLDPDU transmit delay.
- If the LLDPDU transmit delay is greater than the LLDPDU transmit interval, the device uses the LLDPDUs transmit delay as the transmit interval.

Configuration procedure

To change the TTL multiplier:

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

Step	Command	Remarks
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 the 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.

Setting an encapsulation format for LLDPDUs

LLDPDUs can be encapsulated in the following formats: Ethernet II or SNAP frames.

- With Ethernet II encapsulation configured, an LLDP port sends LLDPDUs in Ethernet II frames and processes only incoming, Ethernet II encapsulated LLDPDUs.
- With SNAP encapsulation configured, an LLDP port sends LLDPDUs in SNAP frames and processes only incoming, SNAP encapsulated LLDPDUs.

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

To set the encapsulation format for LLDPDUs to SNAP:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none"> • Enter Layer 2 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.
3. Set the encapsulation format for LLDPDUs to SNAP.	lldp encapsulation snap	Ethernet II encapsulation format applies by default.

NOTE:

LLDP-CDP (Cisco Discovery Protocol) packets use only SNAP encapsulation.

Configuring CDP compatibility

When the switch is directly connected to a Cisco device that supports only CDP rather than LLDP, you can enable CDP compatibility to enable the switch to exchange information with the directly-connected device.

With CDP compatibility enabled on the switch, the switch can use LLDP to receive and recognize the CDP packets received from the directly-connected device and send CDP packets to the directly-connected device. The packets that the switch sends to the neighboring CDP device carry the fields in [Table 27](#).

Table 27 Fields in CDP packets

Field	Description
Device ID	Device ID, which is the bridge MAC address of the device.
Addresses	IPv4 address of the interface.
	The port IPv4 address is the main IP address of the VLAN interface that is in up state and whose corresponding VLAN ID is the lowest among the VLANs permitted on the port. If none of the VLAN interfaces of the permitted VLANs is assigned an IP address or all VLAN interfaces are down, no port IP address will be advertised.
Port ID	Port ID.
Capabilities	Device capability, which is Switch.
Software Version	Software version of the device.
Platform	Device model.
Duplex	Duplex mode of the port.
MTU	Maximum transmit unit.
System Name	System name.
Native VLAN	PVID of a port.
Voice VLAN	VLAN specified by using the lldp voice-vlan command or the voice VLAN configured on the port.

The CDP neighbor-information-related fields in the output of the **display lldp neighbor-information** command show the CDP neighboring device information that can be recognized by the switch. For more information about the **display lldp neighbor-information** command, see *Layer 2—LAN Switching Command Reference*.

To make your device work with Cisco IP phones, you must enable CDP compatibility.

If your LLDP-enabled device cannot recognize CDP packets, it does not respond to the requests of Cisco IP phones for the voice VLAN ID configured on the device. As a result, a requesting Cisco IP phone sends voice traffic without any tag to your device, and, as a result, your device cannot differentiate the voice traffic from other types of traffic.

With CDP compatibility enabled, your device can receive and recognize CDP packets from a Cisco IP phone and respond with CDP packets, which carry the voice VLAN configuration TLVs. According to the voice VLAN configuration TLVs, the IP phone automatically configures the voice VLAN. As a result, the voice traffic is confined in the configured voice VLAN, and differentiated from other types of traffic.

For more information about voice VLANs, see "Configuring a voice VLAN."

Configuration prerequisites

Before you configure CDP compatibility, complete the following tasks:

- Globally enable LLDP.
- Enable LLDP on the port connecting to a device supporting CDP, and configure the port to operate in TxRx mode.

Configuration procedure

△ CAUTION:

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

CDP-compatible LLDP operates in one of the follows modes:

- **TxRx**—CDP packets can be transmitted and received.
- **Disable**—CDP packets cannot be transmitted or received.

To make CDP-compatible LLDP take effect on specific 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:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable CDP compatibility globally.	lldp compliance cdp	Disabled by default.
3. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">• Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>• Enter port group view: port-group manual <i>port-group-name</i>	Use either command.
4. Configure CDP-compatible LLDP to operate in TxRx mode.	lldp compliance admin-status cdp txrx	Disable mode by default.

Configuring LLDP trapping

LLDP trapping notifies the network management system (NMS) of events such as newly-detected neighboring devices and link malfunctions.

LLDP traps are sent periodically, and the interval is configurable. To prevent excessive LLDP traps from being sent when the topology is unstable, set a trap transmit interval for LLDP.

To configure LLDP trapping:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Ethernet interface view or port group view.	<ul style="list-style-type: none">• Enter Layer 2 Ethernet interface view: interface <i>interface-type</i> <i>interface-number</i>• Enter port group view: port-group manual <i>port-group-name</i>	Use either command.
3. Enable LLDP trapping.	lldp notification remote-change enable	Disabled by default.
4. Return to system view.	quit	N/A
5. Set the LLDP trap transmit interval.	lldp timer notification-interval <i>interval</i>	Optional. 5 seconds by default.

Displaying and maintaining LLDP

Task	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 <i>system-name</i>]] [[{ 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

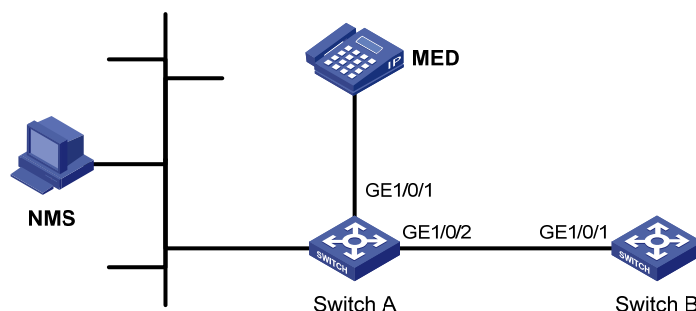
Basic LLDP configuration example

Network requirements

As shown in Figure 67, 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 67 Network diagram



Configuration procedure

1. Configure Switch A:
Enable LLDP globally.


```
<SwitchA> system-view
```

```
[SwitchA] lldp enable
```

Enable LLDP on GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2. (You can skip this step because LLDP is enabled on ports by default.) Set the LLDP operating mode to Rx.

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

```
[SwitchA-GigabitEthernet1/0/1] lldp enable
```

```
[SwitchA-GigabitEthernet1/0/1] lldp admin-status rx
```

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

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

```
[SwitchA-GigabitEthernet1/0/2] lldp enable
```

```
[SwitchA-GigabitEthernet1/0/2] lldp admin-status rx
```

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

2. Configure Switch B:

Enable LLDP globally.

```
<SwitchB> system-view
```

```
[SwitchB] lldp enable
```

Enable LLDP on GigabitEthernet1/0/1. (You can skip this step because LLDP is enabled on ports by default.) Set the LLDP operating mode to Tx.

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

```
[SwitchB-GigabitEthernet1/0/1] lldp enable
```

```
[SwitchB-GigabitEthernet1/0/1] lldp admin-status tx
```

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

3. Verify the configuration:

Display the global LLDP status and port LLDP status on Switch A.

```
[SwitchA] display lldp status
```

```
Global status of LLDP: Enable
```

```
The current number of LLDP neighbors: 2
```

```
The current number of CDP neighbors: 0
```

```
LLDP neighbor information last changed time: 0 days,0 hours,4 minutes,40 seconds
```

```
Transmit interval          : 30s
```

```
Hold multiplier            : 4
```

```
Reinit delay               : 2s
```

```
Transmit delay             : 2s
```

```
Trap interval              : 5s
```

```
Fast start times           : 3
```

```
Port 1 [GigabitEthernet1/0/1]:
```

```
Port status of LLDP       : Enable
```

```
Admin status              : Rx_Only
```

```
Trap flag                  : No
```

```
Polling interval          : 0s
```

```
Number of neighbors:      1
```

```
Number of MED neighbors   : 1
```

```
Number of CDP neighbors   : 0
```

```
Number of sent optional TLV : 0
```

```
Number of received unknown TLV : 0
```

```

Port 2 [GigabitEthernet1/0/2]:
Port status of LLDP           : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s

```

```

Number of neighbors:          1
Number of MED neighbors       : 0
Number of CDP neighbors       : 0
Number of sent optional TLV   : 0
Number of received unknown TLV : 3

```

As the sample output shows, GigabitEthernet 1/0/1 of Switch A connects to an MED device, and GigabitEthernet 1/0/2 of Switch A connects to a non-MED device. Both ports operate in Rx mode, and they only receive LLDPDUs.

Remove the link between Switch A and Switch B and then display the global LLDP status and port LLDP status on Switch A.

```

[SwitchA] display lldp status
Global status of LLDP: Enable
The current number of LLDP neighbors: 1
The current number of CDP neighbors: 0
LLDP neighbor information last changed time: 0 days,0 hours,5 minutes,20 seconds
Transmit interval           : 30s
Hold multiplier              : 4
Reinit delay                 : 2s
Transmit delay               : 2s
Trap interval                : 5s
Fast start times             : 3

```

```

Port 1 [GigabitEthernet1/0/1]:
Port status of LLDP           : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s

```

```

Number of neighbors          : 1
Number of MED neighbors       : 1
Number of CDP neighbors       : 0
Number of sent optional TLV   : 0
Number of received unknown TLV : 5

```

```

Port 2 [GigabitEthernet1/0/2]:
Port status of LLDP           : Enable
Admin status                   : Rx_Only
Trap flag                      : No
Polling interval               : 0s

```

```

Number of neighbors          : 0

```

```
Number of MED neighbors      : 0
Number of CDP neighbors      : 0
Number of sent optional TLV  : 0
Number of received unknown TLV : 0
```

As the sample output shows, GigabitEthernet 1/0/2 of Switch A does not connect to any neighboring devices.


CDP-compatible LLDP configuration example

Network requirements

As shown in [Figure 68](#), GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 of Switch A are each connected to a Cisco IP phone. The two IP phones send out tagged voice traffic.

Configure voice VLAN 2 on Switch A. Enable CDP compatibility of LLDP on Switch A to allow the Cisco IP phones to automatically configure the voice VLAN, confining their voice traffic within the voice VLAN and isolating the voice traffic from other types of traffic.

Figure 68 Network diagram



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 by displaying 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.

Configuring MVRP

Overview

Multiple Registration Protocol (MRP) is an attribute registration protocol and transmits attribute messages. Multiple VLAN Registration Protocol (MVRP) is a typical MRP application. MVRP propagates and learns VLAN configuration among devices. MVRP enables a device to propagate the local VLAN configuration to the other devices, receive VLAN configuration from other devices, and dynamically update the local VLAN configuration (including the active VLANs and the ports through which a VLAN can be reached). MVRP makes sure that all MVRP-enabled devices in a LAN maintain the same VLAN configuration, and reduces the VLAN configuration workload. When the network topology changes, MVRP can propagate and learn VLAN configuration information again according to the new topology, and real-time synchronize the network topology.

MRP is an enhanced version of Generic Attribute Registration Protocol (GARP) and improves the declaration efficiency. MVRP is an enhanced version of GARP VLAN Registration Protocol (GVRP). MVRP delivers the following benefits over GVRP:

- GVRP does not support the multiple spanning tree instance (MSTI). MVRP runs on a per-MSTI basis, and implements per-VLAN redundant link calculation and load sharing.
- MVRP decreases the number of packets transmitted for the same amount of VLAN configuration, and improves the declaration efficiency.

For more information about GVRP, see "Configuring GVRP." For more information about MSTI, see "Configuring spanning tree protocols."

Introduction to MRP

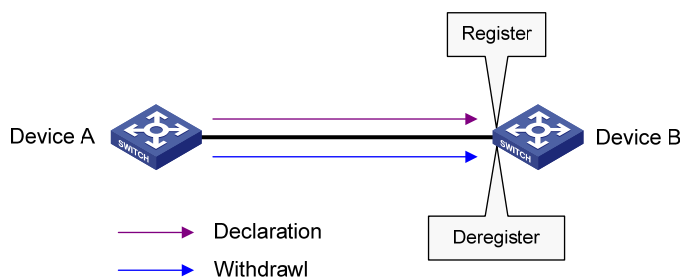
MRP allows participants in the same LAN to declare, propagate, and register information (for example, VLAN information) on a per Multiple Spanning Tree Instance (MSTI) basis.

MRP implementation

Each port that participates in an MRP application (for example, MVRP) is called an "MRP participant". Similarly, a port that participates in an MVRP application is called an "MVRP participant."

As shown in [Figure 69](#), an MRP participant registers and deregisters its attribute values on other MRP participants by sending declarations and withdrawals, and registers and deregisters the attribute values of other participants according to the received declarations and withdrawals. MRP rapidly propagates the configuration information of an MRP participant throughout the LAN.

Figure 69 MRP implementation



MVRP registers and deregisters VLAN attributes as follows:

- When a port receives the declaration of a VLAN attribute, the port registers the VLAN and joins the VLAN.
- When a port receives the withdrawal of a VLAN attribute, the port deregisters the VLAN and leaves the VLAN.

Figure 69 shows a simple MVRP implementation on an MSTI. In a network with multiple MSTIs, VLAN registration and deregistration are performed on a per-MSTI basis.

MRP messages

MRP exchanges information among MRP participants by advertising MRP messages, including Join, New, Leave, and LeaveAll. Join and New messages are declarations, and Leave and LeaveAll messages are withdrawals.

- Join message
 - An MRP participant sends Join messages when it wishes to declare the attribute values configured on it and receives Join messages from other MRP participants.
 - When receiving a Join message, an MRP participant sends a Join message to all participants except the sender.

Join messages fall into the following types:

- **JoinEmpty**—An MRP participant sends JoinEmpty messages to declare attribute values that it has not registered. For example, when a static VLAN exists on a device, the attribute of the VLAN on the device is not changed even if the device learns the VLAN again through MRP. In this case, the Join message for the VLAN attribute is a JoinEmpty message, because the VLAN attribute is not registered.
- **JoinIn**—An MRP participant sends JoinIn messages to declare attribute values that it has registered. For example, when the device learns a VLAN through MRP messages, and dynamically creates the VLAN, the Join message for the VLAN attribute is a JoinIn message.
- New message

Similar to a Join message, a New message enables MRP participants to register attributes.

 - When the Multiple Spanning Tree Protocol (MSTP) topology changes, an MRP participant sends New messages to declare the topology change.
 - On receiving a New message, an MRP participant sends a New message out of each port except the receiving port.
- Leave message
 - An MRP participant sends Leave messages when it wishes other participants to deregister the attributes that it has deregistered.
 - When receiving a Leave message, an MRP participant sends a Leave message to all participants except the sender.
- LeaveAll message
 - Each MRP participant is configured with an individual LeaveAll timer. When the timer expires, the MRP participant sends LeaveAll messages to the remote participants, so that the local participant and the remote participants deregister all attributes and re-register all attributes. This process periodically clears the useless attributes in the network.
 - On receiving a LeaveAll message, MRP determines whether to send a Join message to request the sender to re-register these attributes according to attribute status.

MRP timers

The implementation of MRP uses the following timers to control MRP message transmission.

- Periodic timer

On startup, an MRP participant starts its own Periodic timer to control MRP message transmission. The MRP participant collects the MRP messages to be sent before the Periodic timer expires, and sends the MRP messages in as few packets as possible when the Periodic timer expires and meanwhile restarts the Periodic timer. This mechanism reduces the number of MRP protocol packets periodically sent.

You can enable or disable the Periodic timer at the CLI. When you disable the Periodic timer, MRP will not periodically send MRP messages, and MRP messages are sent only when the LeaveAll timer expires or the local participant receives LeaveAll messages from a remote participant.

- Join timer

The Join timer controls the transmission of Join messages. To make sure that Join messages can be reliably transmitted to other participants, an MRP participant waits for a period of the Join timer after sending a Join message. If the participant receives JoinIn messages from other participants and the attributes in the JoinIn messages are the same as the sent Join messages before the Join timer expires, the participant does not re-send the Join message. When both the Join timer and the Periodic timer expire, the participant re-sends the Join message.

- Leave timer

The Leave timer controls the deregistration of attributes. When an MRP participant wishes other participants to deregister its attributes, it sends a Leave message. On receiving a Leave message, MRP starts the Leave timer, and deregisters the attributes if it does not receive any Join message for the attributes before the Leave timer expires. When an MRP participant sends or receives LeaveAll messages, it starts the Leave timer. MRP deregisters the attributes in the LeaveAll messages if it does not receive any Join message for the attributes before the Leave timer expires.

- LeaveAll timer

On startup, an MRP participant starts its own LeaveAll timer. When the LeaveAll timer expires, MRP sends out a LeaveAll message and restarts the LeaveAll timer. On receiving the LeaveAll message, other participants re-register all the attributes and re-start their LeaveAll timer.

When the LeaveAll timer of an MRP participant expires, the MRP participant sends LeaveAll messages to the remote participants. On receiving a LeaveAll message, a remote participant restarts its LeaveAll timer, and stops sending out LeaveAll messages. This mechanism effectively reduces the number of LeaveAll messages in the network.

To avoid the case that the LeaveAll timer of a fixed participant always first expires, the switch randomly changes the LeaveAll timer within a certain range when the MRP participant restarts its LeaveAll timer.

MVRP registration modes

The VLAN information propagated by MVRP includes not only locally, manually configured static VLAN information but also dynamic VLAN information from other devices.

VLANs created manually, locally are called "static VLANs", and VLANs learned through MVRP are called "dynamic VLANs". The following MVRP registration modes are available.

- Normal

An MVRP participant in normal registration mode performs dynamic VLAN registrations and deregistrations, and sends declarations and withdrawals for dynamic and static VLANs.

- **Fixed**
An MVRP participant in fixed registration mode disables deregistering dynamic VLANs, sends declarations for dynamic VLANs and static VLANs, and drops received MVRP protocol packets. As a result, an MVRP participant port in fixed registration mode does not deregister or register dynamic VLANs.
- **Forbidden**
An MVRP participant in forbidden registration mode disables registering dynamic VLANs, sends declarations for dynamic VLANs and static VLANs, and drops received MVRP protocol packets. As a result, an MVRP participant in forbidden registration mode does not register dynamic VLANs, and does not re-register a dynamic VLAN when the VLAN is deregistered.

Protocols and standards

IEEE 802.1ak *IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks – Amendment 07: Multiple Registration Protocol*

MVRP configuration task list

Task	Remarks
Enabling MVRP	Required.
Configuring the MVRP registration mode	Optional.
Configuring MRP timers	Optional.
Enabling GVRP compatibility	Optional.

Configuration prerequisites

Before configuring MVRP, perform the following tasks:

- Make sure that all MSTIs in the network are effective and each MSTI is mapped to an existing VLAN on each device in the network, because MVRP runs on a per-MSTI basis.
- Configure the involved ports as trunk ports, because MVRP is available only on trunk ports.

Enabling MVRP

Configuration restrictions and guidelines

- MVRP can work with STP, RSTP, or MSTP, but not other link layer topology protocols, including PVST, RRPP, and Smart Link. Ports blocked by STP, RSTP, or MSTP can receive and send MVRP protocol packets. For more information about STP, RSTP, MSTP, and PVST, see "Configuring spanning tree protocols." For more information about RRPP and Smart Link, see *High Availability Configuration Guide*.
- Do not enable both MVRP and remote port mirroring on a port. Otherwise, MVRP may register the remote probe VLAN to incorrect ports, which would cause the monitor port to receive undesired duplicates. For more information about port mirroring, see *Network Management and Monitoring Configuration Guide*.

- Enabling MVRP on a Layer 2 aggregate interface enables both the aggregate interface and all Selected member ports in the link aggregation group to participate in dynamic VLAN registration and deregistration.

Configuration procedure

To enable MVRP:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable MVRP globally.	mvrp global enable	By default, MVRP is globally disabled. To enable MVRP on a port, first enable MVRP globally.
3. Enter interface view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.
4. Configure the port to permit the specified VLANs.	port trunk permit vlan { <i>vlan-list</i> all }	By default, a trunk port permits only VLAN 1. Make sure that the trunk port permits all registered VLANs. For more information about the port trunk permit vlan command, see <i>Layer 2—LAN Switching Command Reference</i> .
5. Enable MVRP on the port.	mvrp enable	By default, MVRP is disabled on a port.

Configuring the MVRP registration mode

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.

Step	Command	Remarks
3. Configure the MVRP registration mode.	mvrp registration { fixed forbidden normal }	Optional. The default setting is normal registration mode.

Configuring MRP timers

⚠ CAUTION:

The MRP timers apply to all MRP applications, for example, MVRP, on a port. To avoid frequent VLAN registrations and deregistrations, use the same MRP timers throughout the network.

Each port maintains its own Periodic, Join, and LeaveAll timers, and each attribute of a port maintains a Leave timer.

To configure MRP timers:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter interface view.	<ul style="list-style-type: none"> Enter Layer 2 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> 	Use one of the commands.
3. Configure the LeaveAll timer.	mrp timer leaveall <i>timer-value</i>	Optional. The default setting is 1000 centiseconds.
4. Configure the Join timer.	mrp timer join <i>timer-value</i>	Optional. The default setting is 20 centiseconds.
5. Configure the Leave timer.	mrp timer leave <i>timer-value</i>	Optional. The default setting is 60 centiseconds.
6. Configure the Periodic timer.	mrp timer periodic <i>timer-value</i>	Optional. The default setting is 100 centiseconds.

Table 28 shows the value ranges for Join, Leave, and LeaveAll timers and their dependencies.

- If you set a timer to a value beyond the allowed value range, your configuration will fail. To do that, you can change the allowed value range by tuning the value of another related timer.
- To restore the default settings of the timers, restore the Join timer first, followed by the Leave and LeaveAll timers.

Table 28 Dependencies of the Join, Leave, and LeaveAll timers

Timer	Lower limit	Upper limit
Join	20 centiseconds	Half the Leave timer
Leave	Twice the Join timer	LeaveAll timer
LeaveAll	Leave timer on each port	32760 centiseconds

NOTE:

You can restore the Periodic timer to the default at any time.

Enabling GVRP compatibility

MVRP can be compatible with GVRP. When the peer device supports GVRP, you can enable GVRP compatibility on the local end, so that the local end can receive and send MVRP and GVRP protocol packets at the same time.

Configuration restrictions and guidelines

- MVRP with GVRP compatibility enabled can work together with STP or RSTP, but cannot work together with MSTP. When MVRP with GVRP compatibility enabled works with MSTP, the network might operate improperly.
- When GVRP compatibility is enabled for MVRP, HP recommends disabling the Period timer. Otherwise, the VLAN status might frequently change when the system is busy.

Configuration procedure

To enable GVRP compatibility:

Step	Command	Remarks
1. Enter system view	system-view	N/A
2. Enable GVRP compatibility	mvrp gvrp-compliance enable	By default, GVRP compatibility is disabled.

Displaying and maintaining MVRP

Task	Command	Remarks
Display the MVRP status of the specified port and each MVRP interface in the specified VLAN.	display mvrp state interface <i>interface-type interface-number</i> vlan <i>vlan-id</i> [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the MVRP running status.	display mvrp running-status [interface <i>interface-list</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view

Task	Command	Remarks
Display the MVRP statistics.	display mvrp statistics [interface <i>interface-list</i>] [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Display the dynamic VLAN operation information of the specified port.	display mvrp vlan-operation interface <i>interface-type interface-number</i> [{ begin exclude include } <i>regular-expression</i>]	Available in any view
Clear the MVRP statistics of the specified ports.	reset mvrp statistics [interface <i>interface-list</i>]	Available in user view

Configuration example for MVRP in normal registration mode

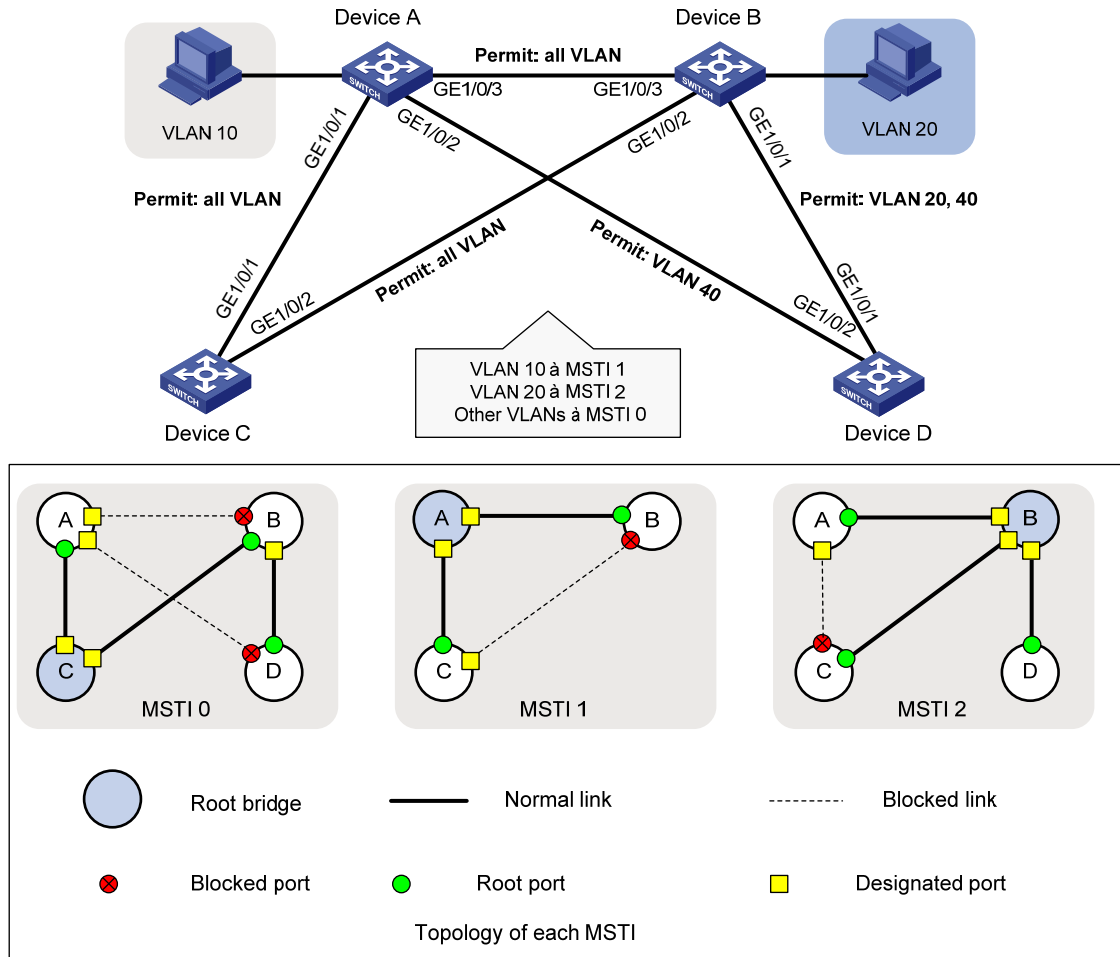
Network requirements

As shown in [Figure 70](#), configure MSTP, map VLAN 10 to MSTI 1, map VLAN 20 MST 2, and map the other VLANs to MSTI 0.

Configure MVRP and set the MVRP registration mode to normal, so that Device A, Device B, Device C, and Device D can register and deregister dynamic and static VLANs and keep identical VLAN configuration for each MSTI.

When the network is stable, set the MVRP registration mode to fixed on the port that connecting Device B to Device A, so that the dynamic VLANs on Device B are not de-registered.

Figure 70 Network diagram



Configuration procedure

Configuring Device A

Enter MST region view.

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

```
[DeviceA] stp region-configuration
```

Configure the MST region name, VLAN-to-instance mappings, and revision level.

```
[DeviceA-mst-region] region-name example
```

```
[DeviceA-mst-region] instance 1 vlan 10
```

```
[DeviceA-mst-region] instance 2 vlan 20
```

```
[DeviceA-mst-region] revision-level 0
```

Manually activate the MST region configuration.

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

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

Configure Device A as the primary root bridge of MSTI 1.

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

Globally enable the spanning tree feature.

```
[DeviceA] stp enable
```

```

# Globally enable MVRP.
[DeviceA] mvrp global enable

# Configure port GigabitEthernet 1/0/1 as a trunk port, and configure it to permit 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 MVRP on port GigabitEthernet 1/0/1.
[DeviceA-GigabitEthernet1/0/1] mvrp enable
[DeviceA-GigabitEthernet1/0/1] quit

# Configure port GigabitEthernet1/0/2 as a trunk port, and configure it to permit VLAN 40.
[DeviceA] interface gigabitethernet 1/0/2
[DeviceA-GigabitEthernet1/0/2] port link-type trunk
[DeviceA-GigabitEthernet1/0/2] port trunk permit vlan 40

# Enable MVRP on port GigabitEthernet1/0/2.
[DeviceA-GigabitEthernet1/0/2] mvrp enable
[DeviceA-GigabitEthernet1/0/2] quit

# Configure port GigabitEthernet 1/0/3 as a trunk port, and configure it to permit all VLANs.
[DeviceA] interface gigabitethernet 1/0/3
[DeviceA-GigabitEthernet1/0/3] port link-type trunk
[DeviceA-GigabitEthernet1/0/3] port trunk permit vlan all

# Enable MVRP on port GigabitEthernet 1/0/3.
[DeviceA-GigabitEthernet1/0/3] mvrp enable
[DeviceA-GigabitEthernet1/0/3] quit

# Create VLAN 10.
[DeviceA] vlan 10
[DeviceA-vlan10] quit

```

Configuring Device B

```

# Enter MST region view.
<DeviceB> system-view
[DeviceB] stp region-configuration

# Configure the MST region name, VLAN-to-instance mappings, and revision level.
[DeviceB-mst-region] region-name example
[DeviceB-mst-region] instance 1 vlan 10
[DeviceB-mst-region] instance 2 vlan 20
[DeviceB-mst-region] revision-level 0

# Manually activate the MST region configuration.
[DeviceB-mst-region] active region-configuration
[DeviceB-mst-region] quit

# Configure Device B as the primary root bridge of MSTI 2.
[DeviceB] stp instance 2 root primary

# Globally enable the spanning tree feature.
[DeviceB] stp enable

# Globally enable MVRP.

```

```

[DeviceB] mvrp global enable

# Configure port GigabitEthernet 1/0/1 as a trunk port, and configure it to permit VLANs 20 and 40.
[DeviceB] interface gigabitethernet 1/0/1
[DeviceB-GigabitEthernet1/0/1] port link-type trunk
[DeviceB-GigabitEthernet1/0/1] port trunk permit vlan 20 40

# Enable MVRP on port GigabitEthernet 1/0/1.
[DeviceB-GigabitEthernet1/0/1] mvrp enable
[DeviceB-GigabitEthernet1/0/1] quit

# Configure port GigabitEthernet1/0/2 as a trunk port, and configure it to permit all VLANs.
[DeviceB] interface gigabitethernet 1/0/2
[DeviceB-GigabitEthernet1/0/2] port link-type trunk
[DeviceB-GigabitEthernet1/0/2] port trunk permit vlan all

# Enable MVRP on port GigabitEthernet1/0/2.
[DeviceB-GigabitEthernet1/0/2] mvrp enable
[DeviceB-GigabitEthernet1/0/2] quit

# Configure port GigabitEthernet 1/0/3 as a trunk port, and configure it to permit all VLANs.
[DeviceB] interface gigabitethernet 1/0/3
[DeviceB-GigabitEthernet1/0/3] port link-type trunk
[DeviceB-GigabitEthernet1/0/3] port trunk permit vlan all

# Enable MVRP on port GigabitEthernet 1/0/3.
[DeviceB-GigabitEthernet1/0/3] mvrp enable
[DeviceB-GigabitEthernet1/0/3] quit

# Create VLAN 20.
[DeviceB] vlan 20
[DeviceB-vlan20] quit

```

Configuring Device C

```

# Enter MST region view.
<DeviceC> system-view
[DeviceC] stp region-configuration

# Configure the MST region name, VLAN-to-instance mappings, and revision level.
[DeviceC-mst-region] region-name example
[DeviceC-mst-region] instance 1 vlan 10
[DeviceC-mst-region] instance 2 vlan 20
[DeviceC-mst-region] revision-level 0

# Manually activate the MST region configuration.
[DeviceC-mst-region] active region-configuration
[DeviceC-mst-region] quit

# Configure Device C as the root bridge of MSTI 0.
[DeviceC] stp instance 0 root primary

# Globally enable the spanning tree feature.
[DeviceC] stp enable

# Globally enable MVRP.
[DeviceC] mvrp global enable

```

Configure port GigabitEthernet 1/0/1 as a trunk port, and configure it to permit all VLANs.

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

Enable MVRP on port GigabitEthernet 1/0/1.

```
[DeviceC-GigabitEthernet1/0/1] mvrp enable
[DeviceC-GigabitEthernet1/0/1] quit
```

Configure port GigabitEthernet1/0/2 as a trunk port, and configure it to permit all VLANs.

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

Enable MVRP on port GigabitEthernet1/0/2.

```
[DeviceC-GigabitEthernet1/0/2] mvrp enable
[DeviceC-GigabitEthernet1/0/2] quit
```

Configuring Device D

Enter MST region view.

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

Configure the MST region name, VLAN-to-instance mappings, and revision level.

```
[DeviceD-mst-region] region-name example
[DeviceD-mst-region] instance 1 vlan 10
[DeviceD-mst-region] instance 2 vlan 20
[DeviceD-mst-region] revision-level 0
```

Manually activate the MST region configuration.

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

Globally enable the spanning tree feature.

```
[DeviceD] stp enable
```

Globally enable MVRP.

```
[DeviceD] mvrp global enable
```

Configure port GigabitEthernet 1/0/1 as a trunk port, and configure it to permit VLANs 20 and 40.

```
[DeviceD] interface gigabitethernet 1/0/1
[DeviceD-GigabitEthernet1/0/1] port link-type trunk
[DeviceD-GigabitEthernet1/0/1] port trunk permit vlan 20 40
```

Enable MVRP on port GigabitEthernet 1/0/1.

```
[DeviceD-GigabitEthernet1/0/1] mvrp enable
[DeviceD-GigabitEthernet1/0/1] quit
```

Configure port GigabitEthernet1/0/2 as a trunk port, and configure it to permit VLAN 40.

```
[DeviceD] interface gigabitethernet 1/0/2
[DeviceD-GigabitEthernet1/0/2] port link-type trunk
[DeviceD-GigabitEthernet1/0/2] port trunk permit vlan 40
```

Enable MVRP on port GigabitEthernet1/0/2.

```
[DeviceD-GigabitEthernet1/0/2] mvrp enable
```



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

Verifying the configuration

1. Verify the normal registration mode configuration:

Use the **display mvrp running-status** command to display the local MVRP VLAN information to verify whether the configuration takes effect.

Check the local VLAN information on Device A.

```
[DeviceA] display mvrp running-status
-----[MVRP Global Info]-----
Global Status      : Enabled
Compliance-GVRP   : False

----[GigabitEthernet1/0/1]----
Config Status      : Enabled
Running Status     : Enabled
Join Timer         : 20 (centiseconds)
Leave Timer         : 60 (centiseconds)
Periodic Timer     : 100 (centiseconds)
LeaveAll Timer      : 1000 (centiseconds)
Registration Type   : Normal
Local VLANs :
  1(default),

----[GigabitEthernet1/0/2]----
Config Status      : Enabled
Running Status     : Enabled
Join Timer         : 20 (centiseconds)
Leave Timer         : 60 (centiseconds)
Periodic Timer     : 100 (centiseconds)
LeaveAll Timer      : 1000 (centiseconds)
Registration Type   : Normal
Local VLANs :
  1(default),

----[GigabitEthernet1/0/3]----
Config Status      : Enabled
Running Status     : Enabled
Join Timer         : 20 (centiseconds)
Leave Timer         : 60 (centiseconds)
Periodic Timer     : 100 (centiseconds)
LeaveAll Timer      : 1000 (centiseconds)
Registration Type   : Normal
Local VLANs :
  1(default), 20,
```

The output shows that:

- Ports GigabitEthernet 1/0/1 and GigabitEthernet1/0/2 have learned only VLAN 1 through MVRP.

- Port GigabitEthernet 1/0/3 has learned VLAN 1 and dynamic VLAN 20 created on Device B through MVRP.

Check the local VLAN information on Device B.

```
[DeviceB] display mvrp running-status
-----[MVRP Global Info]-----
Global Status      : Enabled
Compliance-GVRP    : False

----[GigabitEthernet1/0/1]----
Config Status      : Enabled
Running Status      : Enabled
Join Timer          : 20 (centiseconds)
Leave Timer          : 60 (centiseconds)
Periodic Timer      : 100 (centiseconds)
LeaveAll Timer       : 1000 (centiseconds)
Registration Type    : Normal
Local VLANs :
  1(default),

----[GigabitEthernet1/0/2]----
Config Status      : Enabled
Running Status      : Enabled
Join Timer          : 20 (centiseconds)
Leave Timer          : 60 (centiseconds)
Periodic Timer      : 100 (centiseconds)
LeaveAll Timer       : 1000 (centiseconds)
Registration Type    : Normal
Local VLANs :
  1(default), 10,

----[GigabitEthernet1/0/3]----
Config Status      : Enabled
Running Status      : Enabled
Join Timer          : 20 (centiseconds)
Leave Timer          : 60 (centiseconds)
Periodic Timer      : 100 (centiseconds)
LeaveAll Timer       : 1000 (centiseconds)
Registration Type    : Normal
Local VLANs :
  1(default), 10,
```

The output shows that:

- Port GigabitEthernet 1/0/1 has learned only VLAN 1 through MVRP.
- Ports GigabitEthernet 1/0/2 and GigabitEthernet 1/0/3 have learned VLAN 1 and dynamic VLAN 10 created on Device A through MVRP.

Check the local VLAN information on Device C.

```
[DeviceC] display mvrp running-status
-----[MVRP Global Info]-----
```

```
Global Status      : Enabled
Compliance-GVRP    : False
```

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

```
Config Status      : Enabled
Running Status      : Enabled
Join Timer          : 20 (centiseconds)
Leave Timer          : 60 (centiseconds)
Periodic Timer      : 100 (centiseconds)
LeaveAll Timer       : 1000 (centiseconds)
Registration Type    : Normal
Local VLANs :
  1(default), 10, 20,
```

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

```
Config Status      : Enabled
Running Status      : Enabled
Join Timer          : 20 (centiseconds)
Leave Timer          : 60 (centiseconds)
Periodic Timer      : 100 (centiseconds)
LeaveAll Timer       : 1000 (centiseconds)
Registration Type    : Normal
Local VLANs :
  1(default), 20,
```

The output shows that:

- Port GigabitEthernet 1/0/1 has learned VLAN 1, dynamic VLAN 10 created on Device A, and dynamic VLAN 20 created on Device B through MVRP.
- Port GigabitEthernet1/0/2 has learned VLAN 1 and dynamic VLAN 20 created on Device B through MVRP.

Check the local VLAN information on Device D.

```
[DeviceD] display mvrp running-status
```

```
-----[MVRP Global Info]-----
```

```
Global Status      : Enabled
Compliance-GVRP    : False
```

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

```
Config Status      : Enabled
Running Status      : Enabled
Join Timer          : 20 (centiseconds)
Leave Timer          : 60 (centiseconds)
Periodic Timer      : 100 (centiseconds)
LeaveAll Timer       : 1000 (centiseconds)
Registration Type    : Normal
Local VLANs :
  1(default), 20,
```

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

```
Config Status      : Enabled
```

```

Running Status                : Enabled
Join Timer                    : 20 (centiseconds)
Leave Timer                    : 60 (centiseconds)
Periodic Timer                : 100 (centiseconds)
LeaveAll Timer                 : 1000 (centiseconds)
Registration Type              : Normal
Local VLANs :
    1(default),

```

The output shows that:

- Port GigabitEthernet 1/0/1 has learned VLAN 1 and dynamic VLAN 20 created on Device B through MVRP.
- Port GigabitEthernet1/0/2 has learned only VLAN 1 through MVRP.

2. Change the registration mode and verify the configuration:

Set the MVRP registration mode to fixed on GigabitEthernet 1/0/3 of Device B, so that the dynamic VLANs that Device B learns in VLAN 1 are not de-registered.

Set the MVRP registration mode to fixed on GigabitEthernet 1/0/3.

```

[DeviceB] interface gigabitethernet 1/0/3
[DeviceB-GigabitEthernet1/0/3] mvrp registration fixed
[DeviceB-GigabitEthernet1/0/3] quit

```

Display the local MVRP VLAN information on GigabitEthernet 1/0/3.

```

[DeviceB] display mvrp running-status interface gigabitethernet 1/0/3
-----[MVRP Global Info]-----
Global Status      : Enabled
Compliance-GVRP    : False

----[GigabitEthernet1/0/3]----
Config  Status                : Enabled
Running Status            : Enabled
Join Timer                : 20 (centiseconds)
Leave Timer                : 60 (centiseconds)
Periodic Timer            : 100 (centiseconds)
LeaveAll Timer             : 1000 (centiseconds)
Registration Type          : Fixed
Local VLANs :
    1(default), 10,

```

The output shows that the VLAN information on GigabitEthernet 1/0/3 is not changed after you set the MVRP registration mode to fixed on GigabitEthernet 1/0/3.

Delete VLAN 10 on Device A.

```

[DeviceA] undo vlan 10

```

Display the local MVRP VLAN information on GigabitEthernet 1/0/3.

```

[DeviceB] display mvrp running-status interface gigabitethernet 1/0/3
-----[MVRP Global Info]-----
Global Status      : Enabled
Compliance-GVRP    : False

----[GigabitEthernet1/0/3]----

```

```
Config Status           : Enabled
Running Status          : Enabled
Join Timer               : 20 (centiseconds)
Leave Timer               : 60 (centiseconds)
Periodic Timer           : 100 (centiseconds)
LeaveAll Timer            : 1000 (centiseconds)
Registration Type        : Fixed
Local VLANs :
  1(default), 10,
```

The output shows that the dynamic VLAN information on GigabitEthernet 1/0/3 is not changed after you set the MVRP registration mode to fixed on GigabitEthernet 1/0/3.

Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/wwalerts>

After registering, you will receive email notification of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

Documents

To find related documents, browse to the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>

- For related documentation, navigate to the Networking section, and select a networking category.
- For a complete list of acronyms and their definitions, see *HP FlexNetwork Technology Acronyms*.

Websites

- HP.com <http://www.hp.com>
- HP Networking <http://www.hp.com/go/networking>
- HP manuals <http://www.hp.com/support/manuals>
- HP download drivers and software <http://www.hp.com/support/downloads>
- HP software depot <http://www.software.hp.com>
- HP Education <http://www.hp.com/learn>

Conventions

This section describes the conventions used in this documentation set.





Command conventions

Convention	Description
Boldface	Bold text represents commands and keywords that you enter literally as shown.
<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.
	Represents an access controller, a unified wired-WLAN module, or the switching engine on a unified wired-WLAN switch.
	Represents an access point.
	Represents a security product, such as a firewall, a UTM, or a load-balancing or security card that is installed in a device.
	Represents a security card, such as a firewall card, a load-balancing card, or a NetStream card.

Port numbering in examples

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

Index

[A](#) [B](#) [C](#) [D](#) [E](#) [G](#) [I](#) [L](#) [M](#) [O](#) [P](#) [Q](#) [R](#) [S](#) [T](#) [V](#)

A

Assigning a port to the isolation group, [55](#)

B

BPDU tunneling configuration examples, [110](#)

C

Configuration example for MVRP in normal registration mode, [211](#)

Configuring a combo interface, [1](#)

Configuring a device to advertise voice VLAN information to IP phones, [147](#)

Configuring a loopback interface, [17](#)

Configuring a null interface, [18](#)

Configuring a port group, [8](#)

Configuring a voice VLAN on a port, [148](#)

Configuring an aggregate interface, [43](#)

Configuring an aggregation group, [41](#)

Configuring an MST region, [75](#)

Configuring basic QinQ, [178](#)

Configuring basic settings of a VLAN interface, [116](#)

Configuring basic settings of an Ethernet interface, [2](#)

Configuring basic VLAN settings, [115](#)

Configuring CDP compatibility, [196](#)

Configuring destination multicast MAC address for BPDUs, [109](#)

Configuring Digest Snooping, [89](#)

Configuring edge ports, [81](#)

Configuring flow control on an Ethernet interface, [4](#)

Configuring GVRP functions, [166](#)

Configuring IP subnet-based VLANs, [135](#)

Configuring jumbo frame support, [7](#)

Configuring link change suppression on an Ethernet interface, [5](#)

Configuring LLDP to advertise a specific voice VLAN, [155](#)

Configuring LLDP trapping, [198](#)

Configuring load sharing for link aggregation groups, [46](#)

Configuring loopback testing on an Ethernet interface, [6](#)

Configuring MAC Information mode, [32](#)

Configuring MAC-based VLANs, [124](#)

Configuring MRP timers, [209](#)

Configuring No Agreement Check, [91](#)

Configuring path costs of ports, [81](#)

Configuring port-based VLANs, [118](#)

Configuring protection functions, [94](#)

Configuring protocol-based VLANs, [131](#)

Configuring selective QinQ, [179](#)

Configuring spanning tree timers, [78](#)

Configuring static, dynamic, and blackhole MAC address table entries, [22](#)

Configuring storm control on an Ethernet interface, [14](#)

Configuring storm suppression, [9](#)

Configuring TC snooping, [93](#)

Configuring the aging timer for dynamic MAC address entries, [24](#)

Configuring the device priority, [77](#)

Configuring the GARP timers, [167](#)

Configuring the interval for sending Syslog or trap messages, [32](#)

Configuring the MAC Information queue length, [32](#)

Configuring the MAC learning limit on ports, [26](#)

Configuring the maximum hops of an MST region, [77](#)

Configuring the maximum port rate, [80](#)

Configuring the mode a port uses to recognize/send MSTP packets, [85](#)

Configuring the MVRP registration mode, [208](#)

Configuring the network diameter of a switched network, [78](#)

Configuring the port link type, [85](#)

Configuring the port priority, [84](#)

Configuring the root bridge or a secondary root bridge, [76](#)

Configuring the timeout factor, [80](#)

Configuring the TPID value in VLAN tags, [181](#)

D

- Disabling MAC address learning, [23](#)
- Disabling MAC entry aging timer refresh based on destination MAC address, [25](#)
- Displaying and maintaining an Ethernet interface, [16](#)
- Displaying and maintaining Ethernet link aggregation, [49](#)
- Displaying and maintaining GVRP, [168](#)
- Displaying and maintaining isolate-user-VLAN, [142](#)
- Displaying and maintaining LLDP, [199](#)
- Displaying and maintaining loopback and null interfaces, [18](#)
- Displaying and maintaining MAC address tables, [29](#)
- Displaying and maintaining MVRP, [210](#)
- Displaying and maintaining the isolation group, [55](#)
- Displaying and maintaining the spanning tree, [98](#)
- Displaying and maintaining VLAN, [138](#)
- Displaying and maintaining voice VLAN, [157](#)
- Dynamically advertising server-assigned VLANs through LLDP, [157](#)

E

- Enabling BPDU tunneling, [108](#)
- Enabling bridging on an Ethernet interface, [13](#)
- Enabling energy saving functions on an Ethernet interface, [8](#)
- Enabling GVRP compatibility, [210](#)
- Enabling link-aggregation traffic redirection, [48](#)
- Enabling LLDP to automatically discover IP phones, [154](#)
- Enabling loopback detection on an Ethernet interface, [10](#)
- Enabling MAC address migration log notifying, [28](#)
- Enabling MAC address roaming, [27](#)
- Enabling MAC Information globally, [31](#)
- Enabling MAC Information on an interface, [31](#)
- Enabling MVRP, [207](#)
- Enabling outputting port state transition information, [86](#)
- Enabling the spanning tree feature, [87](#)
- Ethernet interface naming conventions, [1](#)
- Ethernet link aggregation configuration examples, [50](#)
- Ethernet link aggregation configuration task list, [40](#)

G

- GVRP configuration examples, [168](#)
- GVRP configuration task list, [165](#)

I

- IP phone access methods, [148](#)

- Isolate-user-VLAN configuration example, [143](#)

L

- LLDP configuration examples, [199](#)
- LLDP configuration task list, [192](#)

M

- MAC address table configuration example, [29](#)
- MAC Information configuration example, [33](#)
- Methods of identifying IP phones, [146](#)
- MSTP, [64](#)
- MVRP configuration task list, [207](#)

O

- Overview(Isolate-user-VLAN), [140](#)
- Overview(MVRP), [204](#)
- Overview(BPDU tunneling), [106](#)
- Overview(GVRP), [162](#)
- Overview(QinQ), [174](#)
- Overview(LLDP), [187](#)
- Overview(VLAN), [113](#)
- Overview(Voice VLAN), [146](#)
- Overview(MAC address table), [21](#)
- Overview(MAC Information), [31](#)
- Overview(Ethernet link aggregation), [34](#)

P

- Performing basic LLDP configuration, [192](#)
- Performing mCheck, [88](#)
- Port isolation configuration example, [56](#)
- Protocols and standards, [69](#)
- PVST, [64](#)

Q

- QinQ configuration examples, [181](#)
- QinQ configuration task list, [177](#)

R

- RSTP, [64](#)

S

- Setting speed options for auto negotiation on an Ethernet interface, [3](#)
- Setting the MDI mode of an Ethernet interface, [12](#)
- Setting the spanning tree mode, [74](#)
- Setting the statistics polling interval, [10](#)
- Shutting down an Ethernet interface, [3](#)
- Spanning tree configuration examples, [98](#)

Spanning tree configuration task list, [70](#)
STP, [57](#)

T

Testing the cable connection of an Ethernet interface, [14](#)

V

Voice VLAN configuration examples, [158](#)