Contents

Configuring Ethernet interfaces

About Ethernet interface

The Switch Series supports Ethernet interfaces, management Ethernet interfaces, Console interfaces, and USB interfaces. For the interface types and the number of interfaces supported by a switch model, see the installation guide.

This chapter describes how to configure management Ethernet interfaces and Ethernet interfaces.

Configuring a management Ethernet interface

About a management interface

A management interface uses an RJ-45 connector. You can connect the interface to a PC for software loading and system debugging, or connect it to a remote NMS for remote system management.

Each member device in an IRF system has a management Ethernet interface. For management link backup, perform the following tasks:

1. Connect your PC to the management Ethernet interface on the master device.
2. Connect the PC to a management Ethernet interface with the same interface number on a subordinate device.

The two management Ethernet interfaces operate as follows:

- When the IRF system has multiple management Ethernet interfaces, only the management Ethernet interface on the master device processes management traffic.
- When the management Ethernet interface on the master device fails, the management Ethernet interface on the subordinate device takes over to process management traffic.
- When the management Ethernet interface on the master device recovers, it takes over to process management traffic again.

Procedure

1. Enter system view.
   system-view
2. Enter management Ethernet interface view.
   interface M-GigabitEthernet *interface-number*
3. (Optional.) Set the interface description.
   description *text*
   The default setting is **M-GigabitEthernet0/0/0 Interface**.
4. (Optional.) Set the duplex mode.
   duplex { **auto** | **full** | **half** }
   By default, the duplex mode is **auto** for a management Ethernet interface.
5. (Optional.) Set the speed.
   speed { **10** | **100** | **1000** | **auto** }
   By default, a management Ethernet interface negotiates a speed with its peer.
6. (Optional.) Shut down the interface.
   shutdown

By default, the management Ethernet interface is up.

Ethernet interface naming conventions

The Ethernet interfaces are named in the format of **interface type A/B/C**. The letters that follow the interface type represent the following elements:

- **A**—IRF member ID. If the switch is not in an IRF fabric, A is 1 by default.
- **B**—Card slot number. **0** indicates the interface is a fixed interface of the switch.
- **C**—Port index.

Configuring common Ethernet interface settings

This section describes the settings common to Layer 2 Ethernet interfaces, Layer 3 Ethernet interfaces, and Layer 3 Ethernet subinterfaces. For more information about the settings specific to Layer 2 Ethernet interfaces, see "Configuring a Layer 2 Ethernet interface." For more information about the settings specific to Layer 3 Ethernet interfaces or subinterfaces, see "Configuring a Layer 3 Ethernet interface or subinterface."

Configuring basic settings of an Ethernet interface

About Ethernet interface basic settings

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

- **Full-duplex mode**—The interface can send and receive packets simultaneously.
- **Half-duplex mode**—The interface can only send or receive packets at a given time.
- **Autonegotiation mode**—The interface negotiates a duplex mode with its peer.

You can set the speed of an Ethernet interface or enable it to automatically negotiate a speed with its peer.

Restrictions and guidelines

The **shutdown**, **port up-mode**, and **loopback** commands are mutually exclusive.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Set the description for the Ethernet interface.
   description *text*
   The default setting is *interface-name* **Interface**. For example, **Ten-GigabitEthernet1/0/1 Interface**.
4. Set the duplex mode for the Ethernet interface.
   duplex { **auto** | **full** | **half** }
   By default, the duplex mode is **auto** for Ethernet interfaces.
   Ethernet copper ports that operate in 1000 Mbps or 10000 Mbps and fiber ports do not support the **half** keyword.
5. Set the speed for the Ethernet interface.
   speed { **1000** | **10000** | **40000** | **100000** | **auto** }

By default, an Ethernet interface negotiates a speed with its peer.

6. Set the expected bandwidth for the Ethernet interface.

   bandwidth *bandwidth-value*

   By default, the expected bandwidth (in kbps) is the interface baud rate divided by 1000.

7. Bring up the Ethernet interface.

   undo shutdown

   By default, Ethernet interfaces are in down state.

Configuring basic settings of an Ethernet subinterface

Restrictions and guidelines for Ethernet subinterface basic settings

- Before creating a Layer 3 Ethernet subinterface, do not reserve a resource for the VLAN interface whose interface number is the subinterface number. After you reserve a VLAN interface resource, do not create a Layer 3 Ethernet subinterface whose subinterface number is the VLAN interface number. A Layer 3 Ethernet subinterface uses the VLAN interface resource in processing tagged packets whose VLAN ID matches the subinterface number. For more information about reserving resources for VLAN interfaces, see *Layer 2—LAN Switching Configuration Guide*.
- The **shutdown**, **port up-mode**, and **loopback** commands are mutually exclusive.

Procedure

1. Enter system view.

   system-view

2. Create an Ethernet subinterface.

   interface *interface-type interface-number.subnumber*

3. Set the description for the Ethernet subinterface.

   description *text*

   The default setting is *interface-name* **Interface**. For example, **Ten-GigabitEthernet1/0/1.1 Interface**.

4. Set the expected bandwidth for the Ethernet subinterface.

   bandwidth *bandwidth-value*

   By default, the expected bandwidth (in kbps) is the interface baud rate divided by 1000.

5. Bring up the Ethernet subinterface.

   undo shutdown

   By default, Ethernet subinterfaces are in up state.

Configuring the link mode of an Ethernet interface

About the link mode of an Ethernet interface

Interfaces on the device can operate either as Layer 2 or Layer 3 Ethernet interfaces. You can use commands to set the link mode to bridge or route.

Restrictions and guidelines

After you change the link mode of an Ethernet interface, all commands (except the **shutdown**) on the Ethernet interface are restored to their defaults in the new link mode.

Procedure

1. Enter system view.

system-view

2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Configure the link mode of the Ethernet interface.
   port link-mode { **bridge** | **route** }
   By default, all Ethernet interfaces on the device operate in bridge mode.

Configuring jumbo frame support

About jumbo frame

Jumbo frames are frames larger than 1536 bytes and are typically received by an Ethernet interface during high-throughput data exchanges, such as file transfers.

The Ethernet interface processes jumbo frames in the following ways:

- When the Ethernet interface is configured to deny jumbo frames (by using the **undo jumboframe enable** command), the Ethernet interface discards jumbo frames.
- When the Ethernet interface is configured with jumbo frame support, the Ethernet interface performs the following operations:
  - Processes jumbo frames within the specified length.
  - Discards jumbo frames that exceed the specified length.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Configure jumbo frame support.
   jumboframe enable [*size*]
   By default, the device allows jumbo frames within 12288 bytes to pass through.
   If you set the *size* argument multiple times, the most recent configuration takes effect.

Configuring physical state change suppression on an Ethernet interface

About physical state change suppression

The physical link state of an Ethernet interface is either up or down. Each time the physical link of an interface comes up or goes down, the interface immediately reports the change to the CPU. The CPU then performs the following operations:

- Notifies the upper-layer protocol modules (such as routing and forwarding modules) of the change for guiding packet forwarding.
- Automatically generates traps and logs to inform users to take the correct actions.

To prevent frequent physical link flapping from affecting system performance, configure physical state change suppression. You can configure this feature to suppress only link-down events, only link-up events, or both. If an event of the specified type still exists when the suppression interval expires, the system reports the event to the CPU.

Restrictions and guidelines

Do not enable this feature on an interface that has RRPP, spanning tree protocols, or Smart Link enabled.

You can configure different suppression intervals for link-up and link-down events.

If you configure this command multiple times for link-up or link-down events on an Ethernet interface, the most recent configuration takes effect.

The **link-delay**, **dampening**, and **port link-flap protect enable** commands are mutually exclusive on an Ethernet interface.

On an interface, you can configure different suppression intervals for link-up and link-down events. If you configure the **link-delay** command multiple times for link-up or link-down events, the most recent configuration takes effect.

Procedure

1. Enter system view.

   system-view

2. Enter Ethernet interface view.

   interface *interface-type interface-number*

3. Configure physical state change suppression.

   link-delay [**msec**] *delay-time* [**mode** { **up** | **updown** }]

   By default, each time the physical link of a port goes up or comes down, the interface immediately reports the change to the CPU.

   To suppress only link-down events, do not specify the **mode** keyword. To suppress only link-up events, specify the **mode up** keywords. To suppress both link-down and link-up events, specify the **mode updown** keywords.

Configuring dampening on an Ethernet interface

About dampening

The interface dampening feature uses an exponential decay mechanism to prevent excessive interface flapping events from adversely affecting routing protocols and routing tables in the network. Suppressing interface state change events protects the system resources.

If an interface is not dampened, its state changes are reported. For each state change, the system also generates an SNMP trap and log message.

After a flapping interface is dampened, it does not report its state changes to the CPU. For state change events, the interface only generates SNMP trap and log messages.

Parameters

- **Penalty**—The interface has an initial penalty of 0. When the interface flaps, the penalty increases by 1000 for each down event until the ceiling is reached. It does not increase for up events. When the interface stops flapping, the penalty decreases by half each time the half-life timer expires until the penalty drops to the reuse threshold.
- **Ceiling**—The penalty stops increasing when it reaches the ceiling.
- **Suppress-limit**—The accumulated penalty that triggers the device to dampen the interface. In dampened state, the interface does not report its state changes to the CPU. For state change events, the interface only generates SNMP traps and log messages.
- **Reuse-limit**—When the accumulated penalty decreases to this reuse threshold, the interface is not dampened. Interface state changes are reported to the upper layers. For each state change, the system also generates an SNMP trap and log message.
- **Decay**—The amount of time (in seconds) after which a penalty is decreased.

- **Max-suppress-time**—The maximum amount of time the interface can be dampened. If the penalty is still higher than the reuse threshold when this timer expires, the penalty stops increasing for down events. The penalty starts to decrease until it drops below the reuse threshold.

When configuring the **dampening** command, follow these rules to set the values mentioned above:

- The ceiling is equal to $2^{(\text{Max-suppress-time/Decay})}$ × reuse-limit. It is not user configurable.
- The configured suppress limit is lower than or equal to the ceiling.
- The ceiling is lower than or equal to the maximum suppress limit supported.

Figure 1 shows the change rule of the penalty value. The lines t_0 and t_2 indicate the start time and end time of the suppression, respectively. The period from t_0 to t_2 indicates the suppression period, t_0 to t_1 indicates the max-suppress-time, and t_1 to t_2 indicates the complete decay period.

Figure 1 Change rule of the penalty value

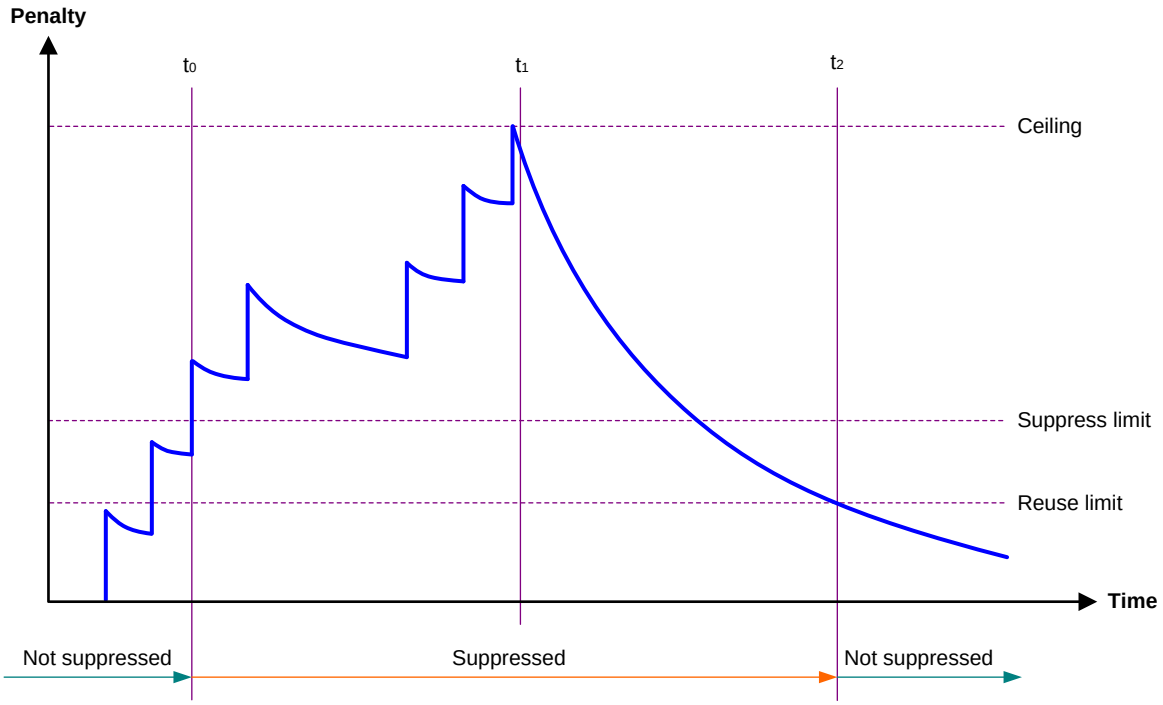

Restrictions and guidelines

- The **dampening**, **link-delay**, and **port link-flap protect enable** commands are mutually exclusive on an interface.
- The **dampening** command does not take effect on the administratively down events. When you execute the **shutdown** command, the penalty restores to 0, and the interface reports the down event to the upper-layer protocols.
- Do not enable the dampening feature on an interface with RRPP, MSTP, or Smart Link enabled.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Enable dampening on the interface.
   dampening [*half-life reuse suppress max-suppress-time*]

By default, interface dampening is disabled on Ethernet interfaces.

Enabling link flapping protection on an interface

About link flapping protection

Link flapping on an interface changes network topology and increases the system overhead. For example, in an active/standby link scenario, when interface status on the active link changes between **UP** and **DOWN**, traffic switches between active and standby links. To solve this problem, configure this feature on the interface.

With this feature enabled on an interface, when the interface goes down, the system enables link flapping detection. During the link flapping detection interval, if the number of detected flaps reaches or exceeds the link flapping detection threshold, the system shuts down the interface.

Restrictions and guidelines

This feature takes effect only if it is configured in both the system view and interface view.

The **dampening**, **link-delay**, and **port link-flap protect enable** commands are mutually exclusive on an Ethernet interface.

To bring up an interface that has been shut down by link flapping protection, execute the **undo shutdown** command.

In the **display interface** command output, the **Link-Flap DOWN** value of the **Current state** field indicates that the interface has been shut down by link flapping protection.

Procedure

1. Enter system view.
   system-view
2. Enable link flapping protection on all interfaces.
   link-flap protect enable
   By default, link flapping protection is disabled on all interfaces.
3. Enter Ethernet interface view.
   interface *interface-type interface-number*
4. Enable link flapping protection on the Ethernet interface.
   port link-flap protect enable [**interval** *interval* | **threshold** *threshold*] *
   By default, link flapping protection is disabled on an Ethernet interface.

Configuring storm suppression

About storm suppression

The storm suppression feature ensures that the size of a particular type of traffic (broadcast, multicast, or unknown unicast traffic) does not exceed the threshold on an interface. 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.

Both storm suppression and storm control can suppress storms on an interface. Storm suppression uses the chip to suppress traffic. Storm suppression has less impact on the device performance than storm control, which uses software to suppress traffic.

Restrictions and guidelines

- For the traffic suppression result to be determined, do not configure storm control together with storm suppression for the same type of traffic. For more information about storm control, see "Configuring storm control on an Ethernet interface."
- When you configure the suppression threshold in kbps, the actual suppression threshold might be different from the configured one as follows:
  - If the configured value is smaller than 64, the value of 64 takes effect.
  - If the configured value is greater than 64 but not an integer multiple of 64, the integer multiple of 64 that is greater than and closest to the configured value takes effect.

  For the suppression threshold that takes effect, see the prompt on the device.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Enable broadcast suppression and set the broadcast suppression threshold.
   broadcast-suppression { *ratio* | **pps** *max-pps* | **kbps** *max-kbps* }
   By default, broadcast suppression is disabled.
4. Enable multicast suppression and set the multicast suppression threshold.
   multicast-suppression { *ratio* | **pps** *max-pps* | **kbps** *max-kbps* }
   By default, multicast suppression is disabled.
5. Enable unknown unicast suppression and set the unknown unicast suppression threshold.
   unicast-suppression { *ratio* | **pps** *max-pps* | **kbps** *max-kbps* }
   By default, unknown unicast suppression is disabled.

Configuring generic flow control on an Ethernet interface

About generic flow control

To avoid dropping packets on a link, you can enable generic 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. Generic flow control includes the following types:

- **TxRx-mode generic flow control**—Enabled by using the **flow-control** command. With TxRx-mode generic flow control enabled, an interface can both send and receive flow control frames:
  - When congestion occurs, the interface sends a flow control frame to its peer.
  - When the interface receives a flow control frame from its peer, it suspends sending packets to its peer.
- **Rx-mode generic flow control**—Enabled by using the **flow-control receive enable** command. With Rx-mode generic flow control enabled, an interface can receive flow control frames, but it cannot send flow control frames:
  - When congestion occurs, the interface cannot send flow control frames to its peer.
  - When the interface receives a flow control frame from its peer, it suspends sending packets to its peer.

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 end. To enable both ends of a link to handle traffic congestion, configure the **flow-control** command at both ends.

Restrictions and guidelines

The generic flow control and PFC features are mutually exclusive.

Procedure

1. Enter system view.
   `system-view`
2. Enter Ethernet interface view.
   `interface` *`interface-type interface-number`*
3. Enable generic flow control.
   - Enable TxRx-mode generic flow control.
     `flow-control`
   - Enable Rx-mode generic flow control.
     `flow-control receive enable`

   By default, generic flow control is disabled on an Ethernet interface.

Configuring PFC on an Ethernet interface

About PFC

When congestion occurs in the network, the local device notifies the peer to stop sending packets carrying the specified 802.1p priority if all of the following conditions exist:

- Both the local end and the remote end have priority-based flow control (PFC) enabled.
- Both the local end and the remote end have the **`priority-flow-control no-drop dot1p`** command configured.
- The specified 802.1p priority is in the 802.1p priority list specified by the *`dot1p-list`* argument.
- The local end receives a packet carrying the specified 802.1p priority.

The state of the PFC feature is determined by the PFC configuration on the local end and on the peer end. In Table 1:

- The first row lists the PFC configuration on the local interface.
- The first column lists the PFC configuration on the peer.
- The **Enabled** and **Disabled** fields in other cells are possible negotiation results.

Make sure all interfaces that a data flow passes through have the same PFC configuration.

Table 1 PFC configurations and negotiation results

Local (right) Peer (below)	enable	auto	Default
enable	Enabled	Enabled.	Disabled
auto	Enabled	• Enabled if negotiation succeeds. • Disabled if negotiation fails.	Disabled
Default	Disabled	Disabled.	Disabled

Restrictions and guidelines

- For IRF and other protocols to operate correctly, as a best practice, do not enable PFC for 802.1p priorities 0, 6, and 7.

- To perform PFC in an overlay network, execute the **qos trust tunnel-dot1p** command. For information about the overlay network, see *VXLAN Configuration Guide*. For information about the **qos trust tunnel-dot1p** command, see *ACL and QoS Command Reference*.
- To avoid packet loss, apply the same PFC configuration to all interfaces that the packets pass through.
- If you do not enable PFC on an interface, the interface can receive but cannot process PFC pause frames. To make PFC take effect, you must enable PFC on both ends.

The PFC and generic flow control features are mutually exclusive.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Enable PFC in auto mode or forcibly on the Ethernet interface.
   priority-flow-control { **auto** | **enable** }
   By default, PFC is disabled.
4. Enable PFC for 802.1p priorities.
   priority-flow-control no-drop dot1p *dot1p-list*
   By default, PFC is disabled for all 802.1p priorities.

Setting the statistics polling interval

About statistics polling interval

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

Setting the statistics polling interval in Ethernet interface view

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Set the statistics polling interval for the Ethernet interface.
   flow-interval *interval*
   By default, the statistics polling interval is 300 seconds.

Enabling loopback testing on an Ethernet interface

About loopback testing

Perform this task to determine whether an Ethernet link works correctly.

Loopback testing includes the following types:

- **Internal loopback testing**—Tests the device where the Ethernet interface resides. The Ethernet interface sends outgoing packets back to the local device. If the device fails to receive the packets, the device fails.

- **External loopback testing**—Tests the inter-device link. The Ethernet interface sends incoming packets back to the remote device. If the remote device fails to receive the packets, the inter-device link fails.

Restrictions and guidelines

- After you enable this feature on an Ethernet interface, the interface does not forward data traffic.
- The **shutdown**, **port up-mode**, and **loopback** commands are mutually exclusive.
- After you enable this feature on an Ethernet interface, the Ethernet interface switches to full duplex mode. After you disable this feature, the Ethernet interface restores to its duplex setting.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Enable loopback testing.
   loopback { **external** | **internal** }

Restoring the default settings for an interface

Restrictions and guidelines

CAUTION:
This feature might interrupt ongoing network services. Make sure you are fully aware of the impacts of this feature when you use it in a live network.

This feature might fail to restore the default settings for some commands because of command dependencies or system restrictions. You can use the **display this** command in interface view to check for these commands and perform their **undo** forms or follow the command reference to restore their default settings. If your restoration attempt still fails, follow the error message to resolve the problem.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view or Ethernet subinterface view.
   interface *interface-type* { *interface-number* | *interface-number.subnumber* }
3. Restore the default settings for the interface.
   default

Configuring a Layer 2 Ethernet interface

Setting the MDIX mode of an Ethernet interface

IMPORTANT:
Fiber ports do not support the MDIX mode setting.

About MDIX mode

A physical Ethernet interface has eight pins, each of which plays a dedicated role. For example, pins 1 and 2 transmit signals, and pins 3 and 6 receive signals. 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-Crossover (MDIX) modes:

- **MDIX mode**—Pins 1 and 2 are receive pins and pins 3 and 6 are transmit pins.
- **MDI mode**—Pins 1 and 2 are transmit pins and pins 3 and 6 are receive pins.
- **AutoMDIX mode**—The interface negotiates pin roles with its peer.

NOTE:

This feature does not take effect on pins 4, 5, 7, and 8 of physical Ethernet interfaces.

- Pins 4, 5, 7, and 8 of interfaces operating at 10 Mbps or 100 Mbps do not receive or transmit signals.
- Pins 4, 5, 7, and 8 of interfaces operating at 1000 Mbps or higher rates receive and transmit signals.

Restrictions and guidelines

To enable a copper Ethernet interface to communicate with its peer, set the MDIX mode of the interface by following these guidelines:

- Typically, set the MDIX mode of the interface to AutoMDIX. Set the MDIX mode of the interface to MDI or MDIX only when the device cannot determine the cable type.
- When a straight-through cable is used, configure the interface to operate in an MDIX mode different than its peer.
- When a crossover cable is used, perform one of the following tasks:
  - Configure the interface to operate in the same MDIX mode as its peer.
  - Configure either end to operate in AutoMDIX mode.

Procedure

1. Enter system view.

   system-view
2. Enter Ethernet interface view.

   interface *interface-type interface-number*
3. Set the MDIX mode of the Ethernet interface.

   mdix-mode { **automdix** | **mdi** | **mdix** }

   By default, a copper Ethernet interface operates in auto mode to negotiate pin roles with its peer.

Configuring storm control on an Ethernet interface

About storm control

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 an upper threshold.

Depending on your configuration, when a particular type of traffic exceeds its upper threshold, the interface performs either of the following operations:

- **Blocks this type of traffic and forwards 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 interface begins to forward the traffic.
- **Goes down automatically**—The interface goes down automatically and stops forwarding any traffic. When the blocked traffic drops below the lower threshold, the interface does not automatically come up. To bring up the interface, use the **undo shutdown** command or disable the storm control feature.

You can configure an Ethernet interface to output threshold event traps and log messages when monitored traffic meets one of the following conditions:

- Exceeds the upper threshold.
- Drops below the lower threshold.

Both storm suppression and storm control can suppress storms on an interface. Storm suppression uses the chip to suppress traffic. Storm suppression has less impact on the device performance than storm control, which uses software to suppress traffic. For more information about storm suppression, see "Configuring storm suppression."

Storm control uses a complete polling cycle to collect traffic data, and analyzes the data in the next cycle. An interface takes one to two polling intervals to take a storm control action.

Restrictions and guidelines

For the traffic suppression result to be determined, do not configure storm control together with storm suppression for the same type of traffic.

Procedure

1. Enter system view.

   system-view

2. (Optional.) Set the statistics polling interval of the storm control module.

   storm-constrain interval *interval*

   The default setting is 10 seconds.

   For network stability, use the default or set a longer statistics polling interval.

3. Enter Ethernet interface view.

   interface *interface-type interface-number*

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** } *upperlimit lowerlimit*

   By default, storm control is disabled.

5. Set the control action to take when monitored traffic exceeds the upper threshold.

   storm-constrain control { **block** | **shutdown** }

   By default, storm control is disabled.

6. Enable the Ethernet interface to output log messages when it detects storm control threshold events.

   storm-constrain enable log

   By default, the Ethernet interface outputs log messages when monitored traffic exceeds the upper threshold or drops below the lower threshold from a value above the upper threshold.

7. Enable the Ethernet interface to send storm control threshold event traps.

   storm-constrain enable trap

   By default, the Ethernet interface sends traps when monitored traffic exceeds the upper threshold or drops below the lower threshold from the upper threshold from a value above the upper threshold.

Forcibly bringing up a fiber port

About forcibly bringing up a fiber port

As shown in Figure 2, a fiber port uses separate fibers for transmitting and receiving packets. The physical state of the fiber port is up only when both transmit and receive fibers are physically connected. If one of the fibers is disconnected, the fiber port does not work.

To enable a fiber port to forward traffic over a single link, you can use the **port up-mode** command. This command forcibly brings up a fiber port, even when no fiber links or transceiver modules are present for the fiber port. When one fiber link is present and up, the fiber port can forward packets over the link unidirectionally.

Figure 2 Forcibly bring up a fiber port

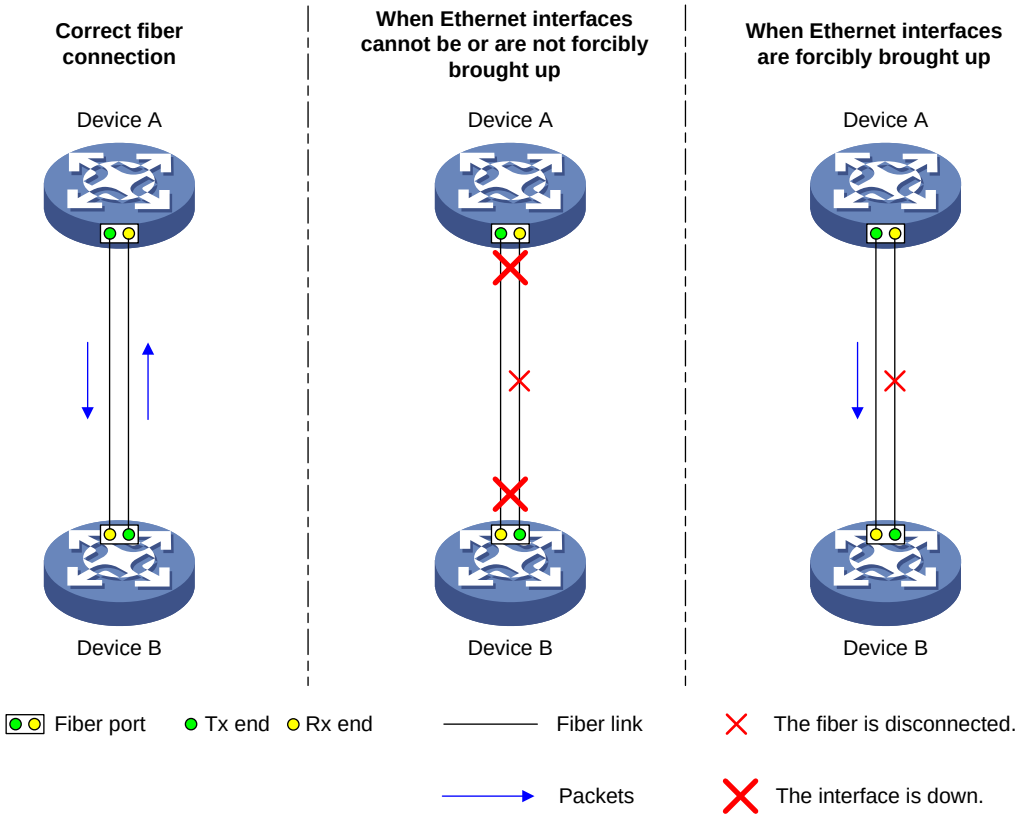

Restrictions and guidelines

- Copper ports do not support this feature.
- The **port up-mode**, **shutdown**, and **loopback** commands are mutually exclusive.
- A fiber port forcibly brought up stays physically up whether or not a transceiver module or a fiber link is present for the port.

Procedure

1. Enter system view.
   system-view
2. Enter Ethernet interface view.
   interface *interface-type interface-number*
3. Forcibly bring up the fiber port.
   port up-mode

By default, a fiber port is not forcibly brought up, and the physical state of a fiber port depends on the physical state of the fibers.

Configuring a Layer 3 Ethernet interface or subinterface

Setting the MTU for an Ethernet interface or subinterface

Restrictions and guidelines

The maximum transmission unit (MTU) of an Ethernet interface affects the fragmentation and reassembly of IP packets on the interface. Typically, you do not need to modify the MTU of an interface.

Procedure

1. Enter system view.
   system-view
2. Enter interface view.
   interface *interface-type* { *interface-number* | *interface-number.subnumber* }
3. Set the interface MTU.
   mtu *size*
   The default setting is 1500 bytes.

Setting the MAC address of an Ethernet interface or subinterface

About Layer 3 Ethernet interface MAC address

In a network, when the Layer 3 Ethernet interfaces or subinterfaces of different devices have the same MAC address, the devices might fail to communicate correctly. To eliminate the MAC address conflicts, use the **mac-address** command to modify the MAC addresses of Layer 3 Ethernet interfaces or subinterfaces.

Restrictions and guidelines

When you set a MAC address, make sure the following requirements are met:

- The MAC address must have the same highest 36 bits as the base MAC address.
- The MAC address must be no lower than the base MAC address plus 128 (decimal).

For more information about the base MAC address, see MAC address table in *Layer 2—LAN Switching Configuration Guide*.

Procedure

1. Enter system view.
   system-view
2. Enter interface view.
   interface *interface-type* { *interface-number* | *interface-number.subnumber* }
3. Set the interface MAC address.
   mac-address *mac-address*

By default, no MAC address is set for an Ethernet interface.

As a best practice, do not set a MAC address in the VRRP-reserved MAC address range for a Layer 3 Ethernet subinterface.

Display and maintenance commands

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

Task	Command
Display interface traffic statistics.	**display counters** { **inbound** | **outbound** } **interface** [*interface-type* [*interface-number*]]
Display traffic rate statistics of interfaces in up state over the last statistics polling interval.	**display counters rate** { **inbound** | **outbound** } **interface** [*interface-type* [*interface-number*]]
Display the Ethernet module statistics.	**display ethernet statistics slot** *slot-number*
Display the operational and status information of the specified interfaces.	**display interface** [*interface-type* [*interface-number* | *interface-number.subnumber*]] [**brief** [**description** | **down**]]
Display information about link flapping protection on interfaces.	**display link-flap protection** [**interface** *interface-type* [*interface-number*]]
Display information about dropped packets on the specified interfaces.	**display packet-drop** { **interface** [*interface-type* [*interface-number*]] | **summary** }
Display PFC information on the specified interfaces.	**display priority-flow-control interface** [*interface-type* [*interface-number*]]
Display information about storm control on the specified interfaces.	**display storm-constrain** [**broadcast** | **multicast** | **unicast**] [**interface** *interface-type interface-number*]
Clear interface statistics.	**reset counters interface** [*interface-type* [*interface-number* | *interface-number.subnumber*]]
Clear the Ethernet module statistics.	**reset ethernet statistics** [**slot** *slot-number*]
Clear the statistics of dropped packets on the specified interfaces.	**reset packet-drop interface** [*interface-type* [*interface-number*]]