7 IP Unicast Routing Commands

7.1 Static Route Configuration Commands

7.1.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.1.2 ip route recursive-lookup bgp-vpnv4-route

Function

The **ip route recursive-lookup bgp-vpnv4-route enable** command enables a device to recurse static routes to cross routes on the remote VPN.

The **ip route recursive-lookup bgp-vpnv4-route disable** command disables a device from recursing static routes to cross routes on the remote VPN.

The **undo ip route recursive-lookup bgp-vpnv4-route disable** command enables a device to recurse static routes to cross routes on the remote VPN.

By default, a static route can recurse to a cross route on the remote VPN.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

ip route recursive-lookup bgp-vpnv4-route { **enable** | **disable** }

undo ip route recursive-lookup bgp-vpnv4-route disable

Parameters

Parameter	Description	Value
enable	Enables a device to recurse static routes to cross routes on the remote VPN.	-
disable	Disables a device from recursing static routes to cross routes on the remote VPN.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Recursion is required if next hops of static routes are not directly reachable. You can run the **ip route recursive-lookup bgp-vpnv4-route enable** command to enable the device to recurse a static route to a cross route on the remote VPN. The static route can then inherit the Label and Token information about the cross route. The device can forward data through the tunnel to which the static route recurses.

Precautions

The **ip route recursive-lookup bgp-vpnv4-route enable** command and the **undo ip route recursive-lookup bgp-vpnv4-route disable** command have the same function and do not generate configurations in the configuration file.

This command allows IPv4 static routes to recurse to cross routes on the remote VPN, not applicable to IPv6 static routes.

Example

Enable a device to recurse static routes to a cross route on the remote VPN.

```
<HUAWEI> system-view
[HUAWEI] ip route recursive-lookup bgp-vpnv4-route enable
```

Disable a device from recursing static routes to a cross route on the remote VPN.

```
<HUAWEI> system-view
[HUAWEI] ip route recursive-lookup bgp-vpnv4-route disable
```

Enable a device to recurse static routes to a cross route on the remote VPN.

```
<HUAWEI> system-view
[HUAWEI] undo ip route recursive-lookup bgp-vpnv4-route disable
```

7.1.3 ip route recursive-lookup blackhole protocol static

Function

The **ip route recursive-lookup blackhole protocol static disable** command disables a device from recursing static routes to blackhole routes.

The **ip route recursive-lookup blackhole protocol static enable** command allows a device to recurse static routes to blackhole routes.

The **undo ip route recursive-lookup blackhole protocol static disable** command allows a device to recurse static routes to blackhole routes.

By default, static routes can recurse to blackhole routes.

Format

ip route recursive-lookup blackhole protocol static { **disable** | **enable** }

undo ip route recursive-lookup blackhole protocol static disable

Parameters

Parameter	Description	Value
disable	Disables a device from recursing static routes to blackhole routes.	-
enable	Allows a device to recurse static routes to blackhole routes.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If IGP, static, and blackhole routes exist on a network and a link fault occurs, services may be interrupted because static routes may recurse to blackhole routes and remain active. To ensure uninterrupted services, run the **ip route recursive-lookup blackhole protocol static disable** command. This command will prevent static routes from recursing to blackhole routes so that static routes will become inactive and IGP routes will be selected.

Precautions

The **ip route recursive-lookup blackhole protocol static enable** and **undo ip route recursive-lookup blackhole protocol static disable** commands have the same function. You only need to run one of them.

Example

Disable a device from recursing static routes to blackhole routes.

```
<HUAWEI> system-view
[HUAWEI] ip route recursive-lookup blackhole protocol static disable
```

Allow a device to recurse static routes to blackhole routes.

```
<HUAWEI> system-view
[HUAWEI] ip route recursive-lookup blackhole protocol static enable
```

Allow a device to recurse static routes to blackhole routes.

```
<HUAWEI> system-view
[HUAWEI] undo ip route recursive-lookup blackhole protocol static disable
```

7.1.4 ip route recursive-lookup inherit-label-route enable

Function

The **ip route recursive-lookup inherit-label-route enable** command allows a device to recurse routes to cross routes on the remote VPN.

The **undo ip route recursive-lookup inherit-label-route enable** command disables a device from recursing routes to cross routes on the remote VPN.

By default, routes except static routes cannot recurse to cross routes on the remote VPN.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

ip route recursive-lookup inherit-label-route enable

undo ip route recursive-lookup inherit-label-route enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Recursion is required if the next hop of a route is unreachable. To allow routes to recurse to cross routes on the remote VPN, run the **ip route recursive-lookup inherit-label-route enable** command. After the command is run, routes inherit the labels and tokens of the cross routes on the remote VPN, ensuring correct traffic forwarding. This command mainly applies to BGP/MPLS IP VPN scenarios.

In **Figure 7-1**, PE1 and PE2 are IBGP peers, and CE1 and PE1 are IGP neighbors. PE2 learns the IP address of CE1's loopback interface, and an EBGP peer relationship is established between CE1 and PE2 using loopback interface IP addresses. By default, the BGP routes that PE2 learns from CE1 through the EBGP peer connection cannot recurse to cross routes on the remote VPN. As a result, traffic fails to be forwarded. To address this problem, run the **ip route recursive-lookup inherit-label-route enable** command to allow the BGP routes to recurse to cross routes on the remote VPN. After the BGP routes inherit the labels and tokens of the cross routes, traffic can be correctly forwarded.

Figure 7-1 Basic BGP/MPLS IP VPN networking

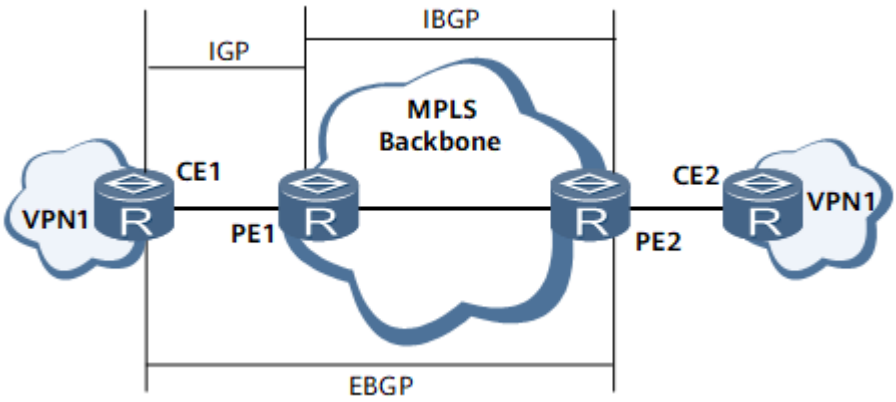

Example

Allow a device to recurse routes to cross routes on the remote VPN.


```
<HUAWEI> system-view
[HUAWEI] ip route recursive-lookup inherit-label-route enable
```

7.1.5 ip route-static

Function

The **ip route-static** command configures a unicast static route.

The **undo ip route-static** command deletes a unicast static route.

By default, no unicast static route is configured.

Format

ip route-static *ip-address* { *mask* | *mask-length* } { *nexthop-address* | *interface-type interface-number* [*nexthop-address*] | **vpn-instance** *vpn-instance-name nexthop-address* } [**preference** *preference* | **tag** *tag*] * [**track** { **route-monitor-group** *route-monitor-group-name* | **bfd-session** *cfg-name* | **efm-state** *interface-type interface-number* | **nqa** *admin-name test-name* } | **permanent**] [**description** *text*]

ip route-static *ip-address* { *mask* | *mask-length* } **vpn-instance** *vpn-instance-name* [**preference** *preference* | **tag** *tag*] * [**description** *text*] (supported only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

undo ip route-static *ip-address* { *mask* | *mask-length* } [*nexthop-address* | *interface-type interface-number* [*nexthop-address*]] [**preference** *preference* | **tag** *tag*] * [**track** { **bfd-session** | **efm-state** } | **permanent**]

undo ip route-static [**track bfd-session**] **all**

undo ip route-static *ip-address* { *mask* | *mask-length* } **vpn-instance** *vpn-instance-name* (supported only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

Parameters

Parameter	Description	Value
ip-address	Specifies a destination IP address.	The value is in dotted decimal notation.
mask	Specifies a subnet mask.	The value is in dotted decimal notation.
mask-length	Specifies the mask length. The 1s in each 32-bit mask must be consecutive. Therefore, a mask in dotted decimal notation can be presented by a mask length.	The value is an integer in the range from 0 to 32.

Parameter	Description	Value
nexthop-address	Specifies the next-hop address.	The value is in dotted decimal notation.
interface-type interface-number	Specifies the type and number of an interface that forwards packets.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance. If a VPN instance name is specified, the device searches the routing table of the VPN instance for a static route according to *nexthop-address*. If the *nexthop-address* parameter is not specified, the device searches a specified VPN instance forwarding table if the device fails to find a forwarding path in the public network forwarding table.	The value is a string of 1 to 31 case-sensitive characters with spaces. The string can contain spaces if it is enclosed with double quotation marks (").
preference *preference*	Specifies the preference of a static route. A smaller value indicates a higher preference.	The value is an integer that ranges from 1 to 255. The default value is 60.
tag *tag*	Specifies the tag value of a static route. By configuring different tag values, you can classify static routes to implement different routing policies. For example, routing protocols can import static routes with specified tag values through routing policies.	The value is an integer that ranges from 1 to 4294967295. The default value is 0.
track route-monitor-group *route-monitor-group-name*	Binds a static route to a route monitoring group for fast fault detection.	The value is a string of 1 to 31 case-sensitive characters.
track bfd-session *cfg-name*	Binds a static BFD session to a static route for fast fault detection.	*cfg-name* is a string of 1 to 15 characters without spaces. The string can contain spaces if it is enclosed with double quotation marks (").

Parameter	Description	Value
track efm-state *interface-type interface-number*	Detects the EFM OAM state of a specified interface.	-
track nqa *admin-name test-name*	Specifies the administrator name of an NQA test instance bound to a static route, and the name of the test instance.	The values of *admin-name* and *test-name* are both a string of 1 to 32 case-sensitive characters without spaces.
permanent	Configures permanent advertisement of static routes.	-
description *text*	Configures a description for a static route.	The value is a string of 1 to 80 characters that can contain spaces.
all	Deletes all IPv4 unicast static routes.	-

NOTE

Only the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **track bfd-session** *cfg-name* parameter.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a simple network, you can configure static routes to ensure network connectivity. If a switch cannot use dynamic routing protocols for network connectivity, static routes can be configured on the switch.

Precautions

When configuring unicast static routes, note the following:

- Only public routes support tunnel recursion.
- If no preference is set for a static route, the static route uses the default preference 60.

- When both the destination IP address and the mask are set to 0.0.0.0, a default route is configured. If the switch cannot find a route in the routing table to forward packets, the switch uses the default route to forward packets.
- By setting different preferences, you can implement different routing management policies. For example, if multiple routes to the same destination are configured with the same preference, load balancing can be implemented. If multiple routes to the same destination are configured with different preferences, route backup can be implemented.
- You can specify the **description** *text* parameter to add a description for a static route so that the administrator can easily check and maintain the static route. You can run the **display this** command in the system view or run the **display current-configuration** command to view the descriptions of static routes.
- When configuring static routes, you can specify outbound interfaces or next-hop addresses based on the outbound interface type. For example, you can specify only outbound interfaces for static routes on P2P interfaces and specify only next hops for static routes on NBMA interfaces, and you must specify next hops for static routes on broadcast interfaces.
- Static routes with the next hop associated with DHCP cannot work in load balancing mode with static routes with the same prefix and not associated with DHCP.
- In some cases, for example, the link layer protocol is PPP and the peer IP address is unknown, you can also specify outbound interfaces when configuring a switch. In this manner, you do not need to modify the switch configuration when the peer IP address changes.
- When a static route to be configured is a supernet route, you need to specify both the outbound interface and next-hop address for it. Otherwise, this route cannot be activated. A route that meets either of the following conditions is a supernet route:
  - When bitwise AND operations are performed on the destination address mask with the destination address and next hop address, the two obtained network addresses are the same, and the destination address mask is greater than or equal to the next hop address mask.
  - When bitwise AND operations are performed on the destination address mask with the destination address and next hop address, the two obtained network addresses are different. When bitwise AND operations are performed on the next hop address mask with the destination address and next hop address, however, the two obtained network addresses are the same.

Example

Set a static route with the destination address and mask are both 0.0.0.0 and the next-hop address 172.16.0.1.

```
<HUAWEI> system-view
[HUAWEI] ip route-static 0.0.0.0 0.0.0.0 172.16.0.1
```


7.1.6 ip route-static default-preference

Function

The **ip route-static default-preference** command sets a default preference for an IPv4 static route.

The **undo ip route-static default-preference** command restores the default preference of IPv4 static routes.

By default, the default preference of IPv4 static routes is 60.

Format

ip route-static default-preference *preference*

undo ip route-static default-preference

Parameters

Parameter	Description	Value
preference	Specifies a default preference for IPv4 static routes. A smaller value indicates a higher preference.	The value is an integer that ranges from 1 to 255.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can change the status of a route by changing the preference of the route. You can use the **ip route-static default-preference** command to change the default preference of all the new IPv4 static routes.

Configuration Impact

After the **ip route-static default-preference** command is used, all the new IPv4 static routes that use the default preference are restored to a new default preference.

Precautions

After a default preference is specified, the new default preference is valid for subsequent rather than existing IPv4 static routes.

Example

Set the default preference of IPv4 static routes to 70.

```
<HUAWEI> system-view
[HUAWEI] ip route-static default-preference 70
```

7.1.7 ip route-static track bfd-session admindown invalid

Function

The **ip route-static track bfd-session admindown invalid** command disables a switch from selecting a static route if the BFD session associated with the route is in the AdminDown state.

The **undo ip route-static track bfd-session admindown invalid** command restores the default configuration.

By default, a static route can participate in route selection even if its bound BFD session is in the AdminDown state.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ip route-static track bfd-session session-name *bfd-name* **admindown invalid**

undo ip route-static track bfd-session session-name *bfd-name* **admindown invalid**

undo ip route-static track bfd-session admindown invalid

Parameters

Parameter	Description	Value
session-name *bfd-name*	Specifies a BFD session bound to a static route.	The value is a string of 1 to 15 case-insensitive characters without spaces. If the string is enclosed within double quotation marks ("), the string can contain spaces.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, if a BFD session associated with a static route is in the AdminDown state, the route can participate in route selection on Huawei switches, but is not allowed for route selection on some non-Huawei devices. As a result, Huawei switches cannot communicate with such non-Huawei devices.

To address this problem, run the **ip route-static track bfd-session admindown invalid** command to configure the device not to select the static route if its bound BFD session is in the AdminDown state.

Prerequisites

The BFD session specified by *bfd-name* has been created.

Example

Disable a switch from selecting a static route if the BFD session associated with the route is in AdminDown state.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] bfd bfda bind peer-ip 1.1.1.1
[HUAWEI-bfd-session-bfda] discriminator local 20
[HUAWEI-bfd-session-bfda] discriminator remote 10
[HUAWEI-bfd-session-bfda] commit
[HUAWEI-bfd-session-bfda] quit
[HUAWEI] ip route-static track bfd-session session-name bfda admindown invalid
```

7.1.8 ip route-static vpn-instance

Function

The **ip route-static vpn-instance** command configures a static route for a VPN instance.

The **undo ip route-static vpn-instance** command deletes static routes from a VPN instance.

By default, no static route is configured for a VPN instance.

Format

ip route-static vpn-instance *vpn-source-name destination-address* { *mask* | *mask-length* } { *nexthop-address* [**public**] | *interface-type interface-number* [*nexthop-address*] | **vpn-instance** *vpn-destination-name nexthop-address* } [**preference** *preference* | **tag** *tag*] * [**track** { **bfd-session** *cfg-name* | **efm-state** *interface-type interface-number* | **nqa** *admin-name test-name* } | **permanent**] [**description** *text*]

ip route-static vpn-instance *vpn-source-name destination-address* { *mask* | *mask-length* } { **public** | **vpn-instance** *vpn-destination-name* } [**preference** *preference* | **tag** *tag*] * [**description** *text*] (supported only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

undo ip route-static vpn-instance *vpn-source-name destination-address* { *mask* | *mask-length* } [*nexthop-address* | *interface-type interface-number* [*nexthop-address*]] [**preference** *preference* | **tag** *tag*] * [**track** { **bfd-session** | **efm-state** } | **permanent**]

undo ip route-static vpn-instance *vpn-source-name* **all**

undo ip route-static vpn-instance *vpn-source-name destination-address* { *mask* | *mask-length* } { **public** | **vpn-instance** *vpn-destination-name* } (supported only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

Parameters

Parameter	Description	Value
vpn-instance *vpn-source-name*	Specifies the name of the source VPN instance. Each VPN instance has its own routing table. The configured static routes are added to the routing table of the specified VPN instance.	The value must be an existing VPN instance name.
destination-address	Specifies the destination IP address of a static route.	The value is in dotted decimal notation.
mask	Specifies the subnet mask of the destination IP address.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of the destination IP address.	The value is an integer that ranges from 0 to 32.

Parameter	Description	Value
interface-type interface-number	Specifies the type and number of the outbound interface of a static route.	-
nexthop-address	Specifies the next-hop address of a static route. If the outbound interface of a static route is a broadcast interface, the next-hop address must be specified for the static route.	The value is in dotted decimal notation.
public	Specifies the next-hop address as a public network address but not an address in the source VPN instance. If the *nexthop-address* parameter is not specified, a device searches another outbound interface in the public IP routing table if the device fails to find a forwarding path in the current VPN instance forwarding table.	-
preference *preference*	Specifies the preference of a static route.	The value is an integer that ranges from 1 to 255. If this parameter is not specified, the default preference of a static route is 60.
tag *tag*	Specifies the tag value of a static route. By configuring different tag values, you can classify static routes to implement different routing policies. For example, routing protocols can import static routes with specified tag values through routing policies.	The value is an integer that ranges from 1 to 4294967295. The default value is 0.
track bfd-session *cfg-name*	Binds a static BFD session to a static route for fast fault detection.	*cfg-name* is a string of 1 to 15 characters without spaces.
track efm-state *interface-type interface-number*	Detects the EFM OAM state of a specified interface.	-

Parameter	Description	Value
track nqa *admin-name test-name*	Specifies the administrator name of an NQA test instance bound to a static route, and the name of the test instance.	The values of *admin-name* and *test-name* are both a string of 1 to 32 case-sensitive characters without spaces.
permanent	Configures permanent advertisement of static routes.	-
description *text*	Specifies the description of a static route. You can specify **description** to add a description for static routes, facilitating static route query and maintenance.	The value is a string of 1 to 80 characters that can contain spaces.
vpn-instance *vpn-destination-name*	Specifies the name of the destination VPN instance. The system searches the routing table of the VPN instance for the outbound interface of a static route according to the configured next-hop address. If the *nexthop-address* parameter is not specified, a device searches another specified VPN instance forwarding table if the device fails to find a forwarding path in the current VPN instance forwarding table. If **vpn-instance** *vpn-destination-name* is specified, devices can communicate with each other between VPNs.	The value is a string of 1 to 31 case-sensitive characters with spaces. If the string is enclosed in double quotation marks (" "), the string can contain spaces.
all	Deletes all static routes from a VPN instance.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You need to configure static routes for a VPN instance on a PE in the following scenarios:

- The PE connects to a CE.
- The specified VPN connects to the Internet.

Follow-up Procedure

After a static route is configured, import the static route to BGP running between PEs if the static route needs to be sent to the peer PE.

Precautions

- During the configuration of a static route for a VPN instance, the next-hop address can belong to the VPN instance or a public network. When the next-hop address of the static route is a public network address, specify **public** after *nexthop-address*.
- If the outbound interface of a static route is a broadcast interface, you must specify the next-hop address for the static route.
- You can specify **description** *text* to add a description for a static route so that the administrator can easily query and maintain the static route. To view the description of static routes, run the **display this** command in the system view or run the **display current-configuration** command.
- If two static routes with the same destination IP address are configured, one of the routes is specified with a next hop IP address or outbound interface, and the other static route is configured with a next-hop VPN instance, without an outbound interface or next hop IP address, the latest configuration overwrites the previous one.

Example

Configure a static route for VPN instance named **vpn1**, and set the destination address of the static route to 10.1.1.0/24 and next-hop address to 192.168.1.2 (IP address of VPN instance **vpn2**).

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vpn1
[HUAWEI-vpn-instance-vpn1] route-distinguisher 100:200
[HUAWEI-vpn-instance-vpn1-af-ipv4] quit
[HUAWEI-vpn-instance-vpn1] quit
[HUAWEI] ip vpn-instance vpn2
[HUAWEI-vpn-instance-vpn2] route-distinguisher 300:400
[HUAWEI-vpn-instance-vpn2-af-ipv4] quit
[HUAWEI-vpn-instance-vpn2] quit
[HUAWEI] ip route-static vpn-instance vpn1 10.1.1.0 24 vpn-instance vpn2 192.168.1.2
```

7.1.9 ipv6 route recursive-lookup inherit-label-route enable

Function

The **ipv6 route recursive-lookup inherit-label-route enable** command allows a device to recurse IPv6 routes to VPN remote cross routes.

The **undo ipv6 route recursive-lookup inherit-label-route enable** command disables a device from recursing IPv6 routes to VPN remote cross routes.

By default, IPv6 routes cannot recurse to VPN remote cross routes.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

ipv6 route recursive-lookup inherit-label-route enable

undo ipv6 route recursive-lookup inherit-label-route enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Recursion is required if the next hop of a route is unreachable. To allow IPv6 routes to recurse to VPN remote cross routes, run the **ipv6 route recursive-lookup inherit-label-route enable** command. After the command is run, IPv6 routes inherit the labels and tokens of the VPN remote cross routes, ensuring correct traffic forwarding. This command mainly applies to BGP/MPLS IPv6 VPN scenarios.

In **Figure 7-2**, PE1 and PE2 are IBGP peers, and CE1 and PE1 are IGP neighbors. PE2 learns the IP address of CE1's loopback interface, and an EBGP peer relationship is established between CE1 and PE2 using loopback interface IP addresses. By default, the BGP routes that PE2 learns from CE1 through the EBGP peer connection cannot recurse to VPN remote cross routes. As a result, traffic fails to be forwarded. To address this problem, run the **ipv6 route recursive-lookup inherit-label-route enable** command to allow the BGP4+ routes to recurse to VPN remote cross routes. After the BGP4+ routes inherit the labels and tokens of the VPN remote cross routes, traffic can be correctly forwarded.

Figure 7-2 Basic BGP/MPLS IPv6 VPN networking

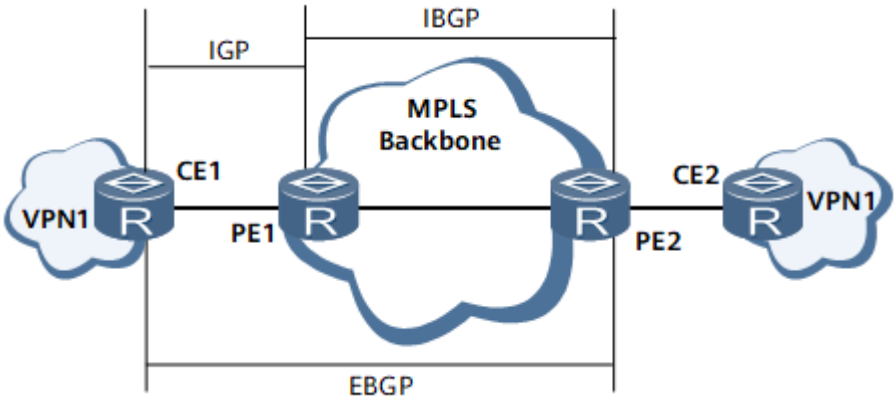

Example

Allow a device to recurse IPv6 routes to VPN remote cross routes.

```
<HUAWEI> system-view
[HUAWEI] ipv6 route recursive-lookup inherit-label-route enable
```

7.1.10 ipv6 route-static

Function

The **ipv6 route-static** command configures an IPv6 static route.

The **undo ipv6 route-static** command deletes configured IPv6 static routes.

By default, no IPv6 static route is configured.

Format

ipv6 route-static *dest-ipv6-address prefix-length* { *interface-type interface-number* [*nexthop-ipv6-address*] | *nexthop-ipv6-address* | **vpn-instance** *vpn-destination-name nexthop-ipv6-address* } [**preference** *preference* | **tag** *tag*] * [**track** { **bfd-session** *cfg-name* | **nqa** *admin-name test-name* }] [**description** *text*]

ipv6 route-static *dest-ipv6-address prefix-length* **vpn-instance** *vpn-destination-name* [**preference** *preference* | **tag** *tag*] * [**description** *text*] (support only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

undo ipv6 route-static *dest-ipv6-address prefix-length* [*interface-type interface-number* [*nexthop-ipv6-address*] | *nexthop-ipv6-address*] [**preference** *preference* | **tag** *tag*] * [**track bfd-session**]

undo ipv6 route-static [**track bfd-session**] **all**

undo ipv6 route-static *dest-ipv6-address prefix-length* **vpn-instance** *vpn-destination-name* [**preference** *preference* | **tag** *tag*] * (support only by the

S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

Parameters

Parameter	Description	Value
dest-ipv6-address	Specifies a destination IPv6 address.	The value is a 32-bit hexadecimal number, in X:X:X:X:X:X:X:X format.
prefix-length	Specifies the length of an IPv6 prefix.	The value is an integer that ranges from 0 to 128.
interface-type interface-number	Specifies the type and number of an outbound interface.	-
vpn-instance *vpn-destination-name*	Specifies the name of a VPN instance. If the VPN instance name is specified, a switch searches the routing table of the VPN instance for the outbound interface of a static route according to the configured next-hop address. If the *nexthop-ipv6-address* parameter is not specified, a device searches a specified VPN instance forwarding table to find a forwarding path.	The value is a string of 1 to 31 case-sensitive characters with spaces. If the string is enclosed in double quotation marks (" "), the string can contain spaces.
nexthop-ipv6-address	Specifies the next-hop IPv6 address.	The value is a 32-bit hexadecimal number, in X:X:X:X:X:X:X:X format.
preference *preference*	Specifies the preference of an IPv6 static route.	The value is an integer that ranges from 1 to 255. The default value is 60.
tag *tag*	Specifies the tag value of a static route. By configuring different tag values, you can classify static routes to implement different routing policies. For example, routing protocols can import static routes with specified tag values through routing policies.	The value is an integer that ranges from 1 to 4294967295. The default value is 0.

Parameter	Description	Value
track bfd-session *cfg-name*	Binds a static BFD session to a static route for fast fault detection.	*cfg-name* is a string of 1 to 15 characters without spaces. The string can contain spaces if it is enclosed with double quotation marks (").
track nqa *admin-name test-name*	Specifies the administrator name of an NQA test instance bound to a static route, and the name of the test instance.	The values of *admin-name* and *test-name* are both a string of 1 to 32 case-sensitive characters without spaces.
description *text*	Specifies the description of a static route.	The value is a string of 1 to 80 characters that can contain spaces.
all	Deletes all IPv6 static routes.	-

NOTE

Only the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S support **track bfd-session** *cfg-name* parameters.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a simple network, you can configure static routes to ensure network connectivity. If a device cannot use dynamic routing protocols for network connectivity, static routes can also be configured on the device.

You can specify **description** *text* to add a description for a static route so that the administrator can easily query and maintain the static route. To view the description of static routes, run the **display this** command in the system view or run the **display current-configuration** command.

Precautions

When configuring unicast static routes, note the following:

- If no preference is set for a static route, the static route uses the default preference 60.

- If the destination address and mask of a static route are all 0s, the static route is a default route.
- For an IPv6 static route, the next hop can be set to either an IP address/ outbound interface or a VPN instance. Only the most recent configuration takes effect.

Example

Configure an IPv6 static route on a public network.

```
<HUAWEI> system-view
[HUAWEI] ipv6 route-static fc00:0:0:2001::1 128 fc00:0:0:2002::2
```

7.1.11 ipv6 route-static default-preference

Function

The **ipv6 route-static default-preference** command sets a default preference for an IPv6 static route.

The **undo ipv6 route-static default-preference** command restores the default preference of IPv6 static routes.

By default, the default preference of IPv6 static routes is 60.

Format

ipv6 route-static default-preference *preference*

undo ipv6 route-static default-preference

Parameters

Parameter	Description	Value
preference	Specifies the default preference of IPv6 static routes. A smaller value indicates a higher preference.	The value is an integer that ranges from 1 to 255.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can change the activation status of routes by changing their preference. To change the default preference of a new IPv6 static route, run the **ipv6 route-static default-preference** command.

Configuration Impact

After the **ipv6 route-static default-preference** command is used, the default preference of all the new IPv6 static routes is changed.

The new default preference is valid only for new IPv6 static routes but not for existing IPv6 static routes.

Example

Set the default preference of IPv6 static routes to 70.

```
<HUAWEI> system-view
[HUAWEI] ipv6 route-static default-preference 70
```

7.1.12 ipv6 route-static track bfd-session admindown invalid

Function

The **ipv6 route-static track bfd-session admindown invalid** command disables a switch from selecting an IPv6 static route if the BFD session associated with the route is in the AdminDown state.

The **undo ipv6 route-static track bfd-session admindown invalid** command restores the default configuration.

By default, an IPv6 static route can participate in route selection even if the BFD session associated with it is in the AdminDown state.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ipv6 route-static track bfd-session session-name *bfd-name* **admindown invalid**

undo ipv6 route-static track bfd-session [**session-name** *bfd-name*] **admindown invalid**

Parameters

Parameter	Description	Value
session-name *bfd-name*	Specifies a BFD session bound to an IPv6 static route.	The value is a string of 1 to 15 case-insensitive characters without spaces. The string can contain spaces if it is enclosed with double quotation marks (").

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, if a BFD session associated with an IPv6 static route is in the AdminDown state, the route can participate in route selection on Huawei switches, but is not allowed for route selection on some non-Huawei devices. As a result, Huawei switches cannot communicate with such non-Huawei devices.

To address this problem, run the **ipv6 route-static track bfd-session admindown invalid** command to configure the switch not to select the IPv6 static route if the BFD session associated with it is in the AdminDown state.

Prerequisites

The BFD session specified by *bfd-name* has been created.

Precautions

The **undo ipv6 route-static track bfd-session admindown invalid** command enables all IPv6 static routes to participate in route selection even if the BFD sessions associated with them are in the AdminDown state.

Example

Disable a switch from selecting an IPv6 static route if the BFD session associated with the route is in the AdminDown state.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] bfd bfda bind peer-ipv6 2001:db8:1::1
[HUAWEI-bfd-session-bfda] discriminator local 20
[HUAWEI-bfd-session-bfda] discriminator remote 10
[HUAWEI-bfd-session-bfda] commit
[HUAWEI-bfd-session-bfda] quit
[HUAWEI] ipv6 route-static track bfd-session session-name bfda admindown invalid
```


7.1.13 ipv6 route-static vpn-instance

Function

The **ipv6 route-static vpn-instance** command configures an IPv6 static route for a VPN instance.

The **undo ipv6 route-static vpn-instance** command deletes IPv6 static routes from a VPN instance.

By default, no IPv6 static routes are configured for VPN instances.

Format

ipv6 route-static vpn-instance *vpn-instance-name dest-ipv6-address prefix-length* { [*interface-type interface-number* [*nexthop-ipv6-address*]] | *nexthop-ipv6-address* [**public**] | **vpn-instance** *vpn-destination-name nexthop-ipv6-address* } [**preference** *preference* | **tag** *tag*] * [**track** { **bfd-session** *cfg-name* | **nqa** *admin-name test-name* }] [**description** *text*]

ipv6 route-static vpn-instance *vpn-instance-name dest-ipv6-address prefix-length* { **public** | **vpn-instance** *vpn-destination-name* } [**preference** *preference* | **tag** *tag*] * [**description** *text*] (support only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

undo ipv6 route-static vpn-instance *vpn-instance-name dest-ipv6-address prefix-length* [*interface-type interface-number*] [*nexthop-ipv6-address*] [**preference** *preference* | **tag** *tag*] *[**track bfd-session**]

undo ipv6 route-static vpn-instance *vpn-instance-name* [**track bfd-session**] **all**

undo ipv6 route-static vpn-instance *vpn-instance-name dest-ipv6-address prefix-length* { **public** | **vpn-instance** *vpn-destination-name* } [**preference** *preference* | **tag** *tag*] * (support only by the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S)

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance. Each VPN instance has its own unicast routing table. The configured static routes are added to the routing table of the specified VPN instance.	The value must be an existing VPN instance name.
dest-ipv6-address	Specifies a destination IPv6 address.	The value is a 32-bit hexadecimal number, in X:X:X:X:X:X:X:X format.

Parameter	Description	Value
prefix-length	Specifies the length of an IPv6 prefix. The length equals the number of consecutive 1s in the mask.	The value is an integer that ranges from 0 to 128.
interface-type	Specifies the type of an interface.	-
interface-number	Specifies the number of an interface.	-
nexthop-ipv6-address	Specifies the next-hop IPv6 address.	The value is a 32-bit hexadecimal number, in X:X:X:X:X:X:X:X format.
vpn-instance *vpn-destination-name*	Specifies the name of the destination VPN instance. If the name of the destination VPN instance is specified, a static route searches the destination VPN instance for an outbound interface according to the configured next-hop address. If the *nexthop-ipv6-address* parameter is not specified, a device searches the destination VPN instance forwarding table to find a forwarding path. If **vpn-instance** *vpn-destination-name* is specified, devices can communicate with each other between VPNs.	The value is a string of 1 to 31 case-sensitive characters with spaces. If the string is enclosed in double quotation marks (" "), the string can contain spaces.
public	Specifies a public network address as the next-hop address.	-
preference *preference*	Specifies the preference of a static route.	The value is an integer that ranges from 1 to 255.
tag *tag*	Specifies the tag value of a static route. By configuring different tag values, you can classify static routes to implement different routing policies. For example, routing protocols can import static routes with specified tag values through routing policies.	The value is an integer that ranges from 1 to 4294967295. The default value is 0.

Parameter	Description	Value
track bfd-session *cfg-name*	Binds a static BFD session to a static route for fast fault detection.	*cfg-name* is a string of 1 to 15 characters without spaces. The string can contain spaces if it is enclosed with double quotation marks (").
track nqa *admin-name test-name*	Specifies the administrator name of an NQA test instance bound to a static route, and the name of the test instance.	The values of *admin-name* and *test-name* are both a string of 1 to 32 case-sensitive characters without spaces.
description *text*	Specifies the description of a static route.	The value is a string of 1 to 80 characters that can contain spaces.
all	Deletes all the IPv6 static routes configured for the specified VPN instance.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a simple IPv6 VPN network, you can run the **ipv6 route-static vpn-instance** command to configure static routes for this VPN.

- To configure VPN users to access a public network, you can perform network address translation (NAT) or run the **ipv6 route-static vpn-instance** command with the keyword **public** specified to configure a public network address as the VPN route's next-hop address.
- You can specify **description** *text* to add a description for a static route so that the administrator can easily query and maintain the static route. To view the description of static routes, run the **display this** command in the system view or run the **display current-configuration** command.

Precautions

If a network fault occurs or the network topology changes, static routes cannot automatically change. Therefore, configure static routes with caution.

If the destination address and the prefix length are set to all 0s, a default route is configured.

Example

Configure a default route with next-hop address FC00:0:0:2001::1.

```
<HUAWEI> system-view
[HUAWEI] ipv6 route-static vpn-instance vpn1 :: 0 fc00:0:0:2001::1
```

7.2 RIP Configuration Commands

7.2.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.2.2 bfd all-interfaces enable(RIP)

Function

The **bfd all-interfaces enable** command enables BFD on all the interfaces in a RIP process.

The **undo bfd all-interfaces enable** command disables BFD on all the interfaces in a RIP process.

By default, BFD is disabled in a RIP process.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

bfd all-interfaces enable

undo bfd all-interfaces enable

Parameters

None

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

If BFD parameters are configured, but the **bfd all-interfaces enable** command is not configured, BFD sessions cannot be established.

NOTE

- The BFD priority on an interface is higher than the BFD priority on a RIP process.
- If the **rip bfd block** command is configured on an interface, the **bfd all-interfaces enable** command cannot be used to enable BFD on the interface.

Example

Enable BFD on all interfaces in a RIP process.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] rip
[HUAWEI-rip-1] bfd all-interfaces enable
```

7.2.3 bfd all-interfaces(RIP)

Function

The **bfd all-interfaces** command sets BFD session parameters.

The **undo bfd all-interfaces** command restores BFD session parameters to default values.

By default, the minimum intervals for receiving and sending BFD packets are 1000 ms and the detection time multiplier is 3.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

bfd all-interfaces { **min-rx-interval** *min-receive-value* | **min-tx-interval** *min-transmit-value* | **detect-multiplier** *detect-multiplier-value* } *

undo bfd all-interfaces { **min-rx-interval** [*min-receive-value*] | **min-tx-interval** [*min-transmit-value*] | **detect-multiplier** [*detect-multiplier-value*] } *

Parameters

Parameter	Description	Value
min-rx-interval *min-receive-value*	Specifies the minimum interval for receiving BFD packets from the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *min-transmit-value*	Specifies the minimum interval for sending BFD packets to the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *detect-multiplier-value*	Indicates the local detection multiplier.	The value is an integer that ranges from 3 to 50. The default value is 3.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The priority of BFD configured on an interface is higher than that configured for a RIP process. If BFD is enabled on an interface, BFD sessions are established using the BFD parameters set on the interface.

Actual interval at which BFD packets are transmitted on the local device = Max { transmit-value on the local device, receive-value on the peer}.

Actual interval at which BFD packets are received on the local device = Max { transmit-value on the peer, receive-value on the local device}.

Local BFD detection time = Actual interval at which BFD packets are received on the local device x detect-multiplier-value on the peer.

Prerequisites

BFD has been enabled on all the interfaces using the **bfd all-interfaces enable** command in the RIP view.

Example

Configure BFD for a RIP process and set the minimum interval for sending BFD packets to 500 ms.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] rip
[HUAWEI-rip-1] bfd all-interfaces enable
[HUAWEI-rip-1] bfd all-interfaces min-tx-interval 500
```

7.2.4 checkzero (RIP)

Function

The **checkzero** command enables the device to check the zero fields in RIP-1 packets.

The **undo checkzero** command disables the device from checking the zero fields in RIP-1 packets.

By default, the device checks the zero fields in RIP-1 packets.

Format

checkzero

undo checkzero

Parameters

None

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a RIP-1 packet, the values of some fields must be zero. These fields are zero fields. After zero field check is enabled, the device checks the zero fields in the RIP-1 packets and discards the packets in which the zero field values are not 0.

The **checkzero** command is valid only for RIP-1 packets.

Run the **undo checkzero** command to reduce CPU usage if there is no zero field to check (all neighbors are trusted).

Configuration Impact

By default, devices reject all RIP-1 packets in which the zero fields are not 0.

RIP-1 implementations vary with device manufacturers. If the third-party peer device allows zero fields in RIP-1 packets to carry non-zero values, run the **undo checkzero** command on the local device. The **undo checkzero** command causes potential security risks to network, so it is not recommended.

Example

Enable the device to check zero fields in RIP-1 packets.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] checkzero
```

7.2.5 default-cost (RIP)

Function

The **default-cost** command changes the default metric of imported routes.

The **undo default-cost** command restores the default metric of imported routes.

By default, the metric of imported routes is 0.

Format

default-cost *cost*

undo default-cost

Parameters

Parameter	Description	Value
cost	Specifies the default metric of imported routes.	The value is an integer that ranges from 0 to 15. The default value is 0.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

You can run one of the following commands to set the metric of imported routes. The following commands are listed in descending order of priorities.

- Run the **apply cost** command to set the metric of imported routes.
- Run the **import-route** (RIP) command to set the metric of imported routes.
- Run the **default-cost** (RIP) command to set the default metric of imported routes.

Example

Set the default metric for imported routes to 2.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] default-cost 2
```

7.2.6 default-route originate

Function

The **default-route originate** command configures RIP to originate a default route or advertise the default route in the routing table to neighbors.

The **undo default-route originate** command restores the default configuration.

By default, the device does not advertise the default route to neighbors.

Format

default-route originate [**cost** *cost* | { **match default** | **route-policy** *route-policy-name* } [**avoid-learning**]]*

undo default-route originate

Parameters

Parameter	Description	Value
cost *cost*	Specifies the metric of the default route.	The value is an integer that ranges from 0 to 15. The default value is 0.
match default	Indicates that the device advertises the default route imported from another routing protocol or RIP process in the routing table to neighbors.	-
avoid-learning	Disables RIP from importing default routes. If a default route in active state already exists in the routing table, the default route is set to inactive after this parameter is used.	-
route-policy *route-policy-name*	Specifies the name of the routing policy. RIP will originate the default route only if route permitted by the route policy is present as active in the routing table.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a routing table, the destination address and mask of a default route are all 0s. If the destination address of a packet does not match any entry in the routing table of the switch, the switch sends the packet along the default route.

If no default route exists and the destination address of the packet does not match any entry in the routing table, the switch discards the packet and sends an Internet Control Message Protocol (ICMP) packet, informing the originating host that the destination address or network is unreachable.

Default routes are originated unconditionally, or based on the configured routing policy or **match default**. When default routes are originated based on the configuration, default routes learned from RIP neighbors will be deleted.

- When the **default-route originate** command is configured without any parameter, the default route is originated unconditionally regardless of whether the default route exists in the IP routing table.
- When the **default-route originate** command is configured with **route-policy**, RIP will originate the default route only if the route that matches **route-policy** is active in the IP routing table.
- When the **default-route originate** command is configured with **match default**, RIP will originate the default route only if a default route learned by another routing protocol or RIP process is present in the IP routing table.
- When the **default-route originate** command is configured with **avoid-learning**, RIP will not learn the default route advertised by RIP peers.

Prerequisites

A RIP process has been created, and the RIP view is displayed using the **rip** command.

Example

Set the metric for the default route to 2.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] default-route originate cost 2
```

7.2.7 description (RIP)

Function

The **description** command configures a description for a RIP process.

The **undo description** command deletes the description of a RIP process.

By default, there is no description for a RIP process.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Specifies a description for a RIP process.	It is a string of 1 to 80 case-sensitive characters. Spaces are supported.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By configuring descriptions for RIP processes, you can easily identify RIP processes.

Configuration Impact

If the **description** command is run multiple times, only the latest configuration takes effect.

Example

Configure a description for RIP process 100.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] description this process configure the poison reverse process
```

7.2.8 display default-parameter rip

Function

The **display default-parameter rip** command displays the default RIP configuration.

Format

display default-parameter rip

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

After the default RIP configuration is modified, running the **display default-parameter rip** command still displays the default configuration during RIP initialization.

Example

Display the default RIP configuration.

```
<HUAWEI> display default-parameter rip
```



```
-----------------------------------------------
 Protocol Level Default Configurations
-----------------------------------------------
      RIP version   : 1
      Preference    : 100
      Checkzero     : Enabled
      Default-cost  : 0
      Auto Summary  : Enabled
      Host-route    : Enabled
      Maximum Balanced Paths : 16
      Update time   : 30 sec             Age time : 180 sec
      Garbage-collect time : 120 sec
      Default-route     : Disabled
      Verify-source     : Enabled
      Graceful restart  : Disabled
-----------------------------------------------
 Interface Level Default Configurations
-----------------------------------------------
  Metricin               : 0
  Metricout              : 1
  Input Packet Processing : Enabled
  Output Packet Processing: Enabled
  Poison Reverse         : Disabled
  Replay Protect         : Disabled
  Split Horizon
   For Broadcast and P2P Interfaces :  Enabled
   For NBMA Interfaces              :  Disabled
  Packet Transmit Interval         :  200 msecs
  Packet Transmit Number           :  50
  RIP Protocol Version             :  RIPv1 Compatible (Non-Standard)
```

Table 7-1 Description of the display default-parameter rip command output

Item	Description
Protocol Level Default Configurations	Default RIP protocol-level configuration.
RIP version	Default global RIP version.
Preference	Default RIP preference.
Checkzero	Whether zero fields in RIP-1 packets are checked by default.
Default-cost	Default metric of routes imported from other routing protocols.
Auto Summary	Whether RIP summarization is enabled by default.
Host-route	Whether host routes are added to routing tables by default.
Maximum Balanced Paths	Default maximum number of equal-cost routes for load balancing.
Update time	Default interval for sending Update packets.
Age time	Default aging time of RIP routes.

Item	Description
Garbage-collect time	Default Garbage-collect time of RIP routes.
Default-route	Whether the default route is used if no entries in the routing table match packets.
Verify-source	Whether source address check is enabled by default.
Graceful restart	Whether RIP GR is enabled by default.
Interface Level Default Configurations	Default RIP configuration on the interface.
Metricin	Default metric added to routes when RIP packets are received.
Metricout	Default metric added to routes when RIP packets are sent.
Input Packet Processing	Whether an interface is allowed to receive RIP packets by default.
Output Packet Processing	Whether an interface is allowed to send RIP packets by default.
Poison Reverse	Whether poison reverse is enabled by default.
Replay Protect	Whether replay-protect is enabled by default.
Split Horizon	Whether split horizon is enabled for the following interfaces: • Broadcast and P2P interfaces • NBMA interfaces
Packet Transmit Interval	Default interval for forwarding packets, in milliseconds.
Packet Transmit Number	Default number of forwarded packets.
RIP Protocol Version	Default RIP version on the interface.

7.2.9 display rip

Function

The **display rip** command displays the status and configuration of a RIP process.

Format

display rip [*process-id* | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
process-id	Displays the status and configuration of a specified RIP process.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Displays the status and configuration of a RIP process with the specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display rip** command to check the status and configuration of a RIP process.

Example

Display the status and configuration of a RIP process. The command output shows that two VPN instances are running. One is a public instance and another one is a VPN instance named **VPN-Instance-1**.

```
<HUAWEI> display rip
Public VPN-instance
  RIP process : 1
    RIP version   : 1
    Preference    : 100
    Checkzero     : Enabled
    Default-cost  : 0
    Summary       : Enabled
    Host-route    : Enabled
    Maximum number of balanced paths : 8
    Update time   : 30 sec            Age time : 180 sec
    Garbage-collect time : 120 sec
    Graceful restart  : Disabled
    BFD               : Disabled
    Silent-interfaces : None
    Default-route     : Disabled
    Verify-source     : Enabled
    Networks          : None
    Configured peers          : None
    Number of routes in database : 0
    Number of interfaces enabled : 0
    Triggered updates sent       : 0
```



```
    Number of route changes      : 0
    Number of replies to queries : 0
    Number of routes in ADV DB   : 0

Private VPN-instance name : VPN-Instance-1
  RIP process : 100
    RIP version   : 1
    Preference    : 100
    Checkzero     : Enabled
    Default-cost  : 0
    Summary       : Enabled
    Host-route    : Enabled
    Maximum number of balanced paths : 8
    Update time   : 30 sec             Age time : 180 sec
    Garbage-collect time : 120 sec
    Graceful restart  : Disabled
    BFD               : Disabled
    Silent-interfaces : None
    Default-route     : Disabled
    Verify-source     : Enabled
    Networks          : None
    Configured peers            : None
    Number of routes in database : 0
    Number of interfaces enabled : 0
    Triggered updates sent       : 0
    Number of route changes      : 0
    Number of replies to queries : 0
    Number of routes in ADV DB   : 0

  Total count for 2 process :
    Number of routes in database : 0
    Number of interfaces enabled : 0
    Number of routes sendable in a periodic update : 0
    Number of routes sent in last periodic update : 0
```

Table 7-2 Description of the display rip command output

Item	Description
RIP process	Running RIP process. To create a RIP process, run the **rip** command.
RIP version	Global RIP version. To set the RIP version, run the **version** command.
Preference	RIP preference. To set the RIP preference, run the **preference** command.
Checkzero	Whether zero fields in RIP-1 packets are checked. To check zero fields in RIP-1 packets, run the **checkzero** command.
Default-cost	Default metric for the routes imported from other routing protocols. To set the default metric, run the **default-cost** command.
Summary	Whether RIP summarization is enabled. To enable RIP summarization, run the **summary** command.
Host-route	Whether host routes are added to routing tables. To add host routes to routing tables, run the **host-route** command.

Item	Description
Maximum number of balanced paths	Maximum number of equal-cost routes for load balancing. To set the maximum number of equal-cost routes, run the **maximum load-balancing** command.
Update time	Interval for sending RIP Update packets. To set the interval for sending RIP Update packets, run the **timers rip** command.
Age time	Aging time of RIP routes. To set the RIP aging time, run the **timers rip** command.
Garbage-collect time	Garbage-collect time of RIP routes. To set the Garbage-collect time, run the **timers rip** command.
Graceful restart	Whether graceful restart is enabled for RIP routes.
BFD	Whether RIP is associated with BFD.
Silent-interfaces	Suppressed RIP interface, which only receives but does not send RIP packets. To specify the suppressed interface, run the **silent-interface** command.
Default-route	Whether RIP originates default routes. To configure RIP to originate default routes, run the **default-route originate** command.
Verify-source	Whether the source addresses of received RIP Update packets are checked. To check the source addresses of received RIP Update packets, run the **verify-source** command.
Networks	Network segment running RIP. To enable RIP for a network segment, run the **network** command.
Configured peers	RIP neighbors. To configure a RIP neighbor, run the **peer** command.
Number of routes in database	Number of routes in the RIP database.
Number of interfaces enabled	Number of interfaces running RIP.
Triggered updates sent	Number of RIP packets triggering update.
Number of route changes	Number of changed routes in the RIP database.
Number of replies to queries	Number of response packets to RIP requests.
Number of routes in ADV DB	Number of routes in each Update packet.

Item	Description
Number of routes sendable in a periodic update	Number of routes sent in each Update interval.
Number of routes sent in last periodic update	Number of routes sent in last Update interval.

7.2.10 display rip bfd session

Function

The **display rip bfd session** command displays information about BFD sessions.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

display rip *process-id* **bfd session** { **interface** *interface-type interface-number* | *neighbor-id* | **all** }

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.
interface *interface-type interface-number*	Specifies the type and number of the interface on which the BFD session is configured.	-
neighbor-id	Specifies the RIP neighbor ID on which the BFD session is configured.	The value is in dotted decimal notation.

Parameter	Description	Value
all	Displays information about BFD sessions established on all BFD-enabled interfaces in the RIP process.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display rip bfd session** command can be used to view the following information in a RIP process:

- BFD sessions of a specified process
- BFD sessions on a specified interface
- BFD sessions on a specified neighbor

Example

Display information about BFD sessions in all RIP processes with IDs of 1.

```
<HUAWEI> display rip 1 bfd session all
 LocalIp      :10.1.0.1         RemoteIp :10.1.0.2         BFDState  :Up
 TX           :1000             RX       :1000             Multiplier:3
 BFD Local Dis:8192             Interface:Vlanif10
 Diagnostic Info: No diagnostic information
 LocalIp      :10.2.0.1         RemoteIp :10.2.0.2         BFDState  :Up
 TX           :1000             RX       :1000             Multiplier:3
 BFD Local Dis:8193             Interface:Vlanif20
 Diagnostic Info: No diagnostic information
```

Table 7-3 Description of the display rip bfd session all command output

Item	Description
LocalIp	Local IP address
RemoteIp	Remote IP address
BFDState	A BFD session status can be either Up or Down.
TX	Minimum interval at which BFD packets are sent
RX	Minimum interval at which BFD packets are received
Multiplier	Remote detection multiplier

Item	Description
BFD Local Dis	Local discriminator dynamically assigned by BFD
Interface	Local interfaces on which BFD sessions are configured
Diagnostic Info	Diagnostic information

7.2.11 display rip database

Function

The **display rip database** command displays all the active routes in the RIP database. These routes are sent in RIP Update packets.

Format

display rip *process-id* **database** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.
verbose	Displays detailed information about the routes in the RIP database.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display rip database** command to check all the active routes in the RIP database. These routes are sent in RIP Update packets.

Example

Display the RIP database.

```
<HUAWEI> display rip 100 database
------------------------------------------------------
Advertisement State : [A] - Advertised
                      [I] - Not Advertised/Withdraw
------------------------------------------------------
```



```
10.0.0.0/8, cost 0, ClassfulSumm
    10.1.1.0/24, cost 0, [A], Imported
    10.10.10.0/24, cost 0, [A], Rip-interface
    10.137.220.0/23, cost 1, [A], nexthop 10.10.10.2
```

Table 7-4 Description of the display rip database command output

Item	Description
Advertisement State	Route status. • [A]: Advertised • [I]: Not Advertised/Withdraw
ClassfulSumm	Classful route summarization.
Imported	Routes imported from other routing protocols.
Rip-interface	Routes learned on the RIP interface.
nexthop	Next-hop address.

7.2.12 display rip graceful-restart

Function

The **display rip graceful-restart** command displays the GR status of RIP.

Format

display rip *process-id* **graceful-restart** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.
verbose	Displays detailed information about GR status.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After receiving the GR abort notification from the helper or detecting a network topology change, the restarter stops the GR process.

Example

Display the GR status in RIP process 1.

```
<HUAWEI> display rip 1 graceful-restart
Restart mode     : Restarting
Restart status   : In Progress - Waiting for updates
Last complete reason : None
Update progress summary:
------------------------
Restart capable peers  : 0
 Completed: 0    Inprogress: 0
Restart incapable peers: 1
 Completed: 0    Inprogress: 1
Update period finishes in 293 seconds
```

Table 7-5 Description of the display rip graceful-restart command output

Item	Description
Restart mode	Role of a switch: • Restarting: indicates the restarter. • Helper: indicates the helper. • None: indicates a common switch that does not participate in GR.
Restart status	GR status: • In Progress - waiting for updates: indicates that the switch is waiting for Update packets. • In Progress - calculating active routes: indicates that the switch is calculating RIP routes. • Aborted - calculating active routes: indicates that the route exits from GR abnormally. • GR completed: indicates that GR is complete.
Last complete reason	Reason that the switch exits from GR for the last time: • Unknown: indicates an unknown reason. • Abort: indicates that the switch exits from GR abnormally. • Period Expired: indicates that GR period expires. • Successful: indicates that the switch exits from GR normally. • None: indicates that the switch does not perform GR.
Update progress summary	Title bar indicating the updated RIP progress summary

Item	Description
Restart capable peers	Number of neighbors that help the local switch to perform GR
Completed	Number of neighbors that have sent all Update packets
Inprogress	Number of neighbors that have not sent all Update packets
Restart incapable peers	Number of neighbors that cannot help the local switch to perform GR
Update period finishes in 293 seconds	Time to finish GR: 293 seconds

7.2.13 display rip interface

Function

The **display rip interface** command displays information about RIP interfaces.

Format

display rip *process-id* **interface** [*interface-type interface-number*] [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies the type and the number of an interface.	-
verbose	Displays detailed information about a RIP interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display rip interface** command output displays the configuration and operating status of RIP, which facilitate fault location and configuration verification.

Example

Display information about RIP on VLANIF100.

```
<HUAWEI> display rip 1 interface vlanif 100
 ------------------------------------------------------------------------------
 Interface               IP Address      State    Protocol            MTU
 ------------------------------------------------------------------------------
 Vlanif 100              10.1.1.2        UP       RIPv1 Compatible    500
```

Table 7-6 Description of the display rip interface command output

Item	Description
Interface	RIP-enabled interfaces
IP Address	IP address of the interface
State	Status of the interface: • UP • DOWN
Protocol	Version of RIP running on the interface: • RIPv1 Compatible • RIPv2 Multicast • RIPv2 Broadcast To set the version of RIP, run the **rip version** command.
MTU	MTU value of the link

Display detailed information about RIP on VLANIF100.

```
<HUAWEI> display rip 1 interface vlanif 100 verbose
 Vlanif100 (10.1.1.1)
  State    :   UP            MTU : 500
  Metricin : 0
  Metricout : 1
  Input    : Enabled         Output : Enabled
  Protocol : RIPv1 Compatible (Non-Standard)
  Send version     : RIPv1 Packets
  Receive version  : RIPv1 Packets, RIPv2 Multicast and Broadcast Packets
  Poison-reverse               : Disabled
  Split-Horizon                : Enabled
  Authentication type          : None
  Replay Protection   : Disabled
  BFD                 : Enabled (Static)
  Transmit-Interval   : 1000 ms
  Receive-Interval    : 1000 ms
  Detect-Multiplier   : 3
  Summary Address (es):  10.1.0.0/16
```


Table 7-7 Description of the display rip interface verbose command output

Item	Description
State	Status of the interface: • UP • DOWN
MTU	Maximum Transmission Unit
Metricin	Metric that is added to the route when the interface receives a RIP packet. To set the metric, run the **rip metricin** command.
Metricout	Metric that is added to the route when the interface sends a RIP packet. To set the metric, run the **rip metricout** command.
Input	Whether receiving packets is enabled. To enable the specified interface to receive RIP packets, run the **rip input** command.
Output	Whether sending packets is enabled. To enable the specified interface to send RIP packets, run the **rip output** command.
Protocol	Protocol running on the interface: • RIPv1 Compatible (Non-Standard) • RIPv1 • RIPv2 Multicast • RIPv2 Broadcast To set the version of RIP, run the **rip version** command.
Send version	Type of packets sent on the interface: • RIPv1 packets • RIPv2 Multicast Packets • RIPv2 Broadcast Packets
Receive version	Type of packets received on the interface: • RIPv1 packets • RIPv2 Multicast and Broadcast Packets
Poison-reverse	Whether poison reverse is enabled on the interface. To enable poison reverse, run the **rip poison-reverse** command.

Item	Description
Split-Horizon	Whether split horizon is enabled on the interface. To enable split horizon, run the **rip split-horizon** command.
Authentication type	Authentication type configured on the interface. To set the authentication mode and authentication parameters, run the **rip authentication-mode** command.
Replay Protection	Whether replay-protect is enabled on the interface. To enable the replay-protect function, run the **rip replay-protect** command.
BFD	Whether BFD is enabled on the interface: • Enabled (Dynamic): indicates dynamic BFD is enabled on the interface. To enable BFD on the specified interface, run the **rip bfd enable** command. • Blocked: indicates BFD is blocked on the interface. To block BFD on a specified interface, run the **rip bfd block** command.
Transmit-Interval	Interval for sending BFD packets to the peer on the interface. To set BFD session parameters on the specified interface, run the **rip bfd** command.
Receive-Interval	Interval for receiving BFD packets from the peer on the interface. To set BFD session parameters on the specified interface, run the **rip bfd** command.
Detect-Multiplier	BFD detect multiplier value configured on the interface. To set BFD session parameters on the specified interface, run the **rip bfd** command.
Summary Address (es)	Summary address

7.2.14 display rip neighbor

Function

The **display rip neighbor** command displays information about RIP neighbors.

Format

display rip *process-id* **neighbor** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.
verbose	Displays detailed information about a RIP neighbor.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display rip neighbor** command to check information about RIP neighbors.

Example

Display information about the neighbor in RIP process 1.

```
<HUAWEI> display rip 1 neighbor
------------------------------------------------------------------------
 IP Address      Interface        Type   Last-Heard-Time
------------------------------------------------------------------------
 10.1.1.1        Vlanif100        RIP    0:0:7
 Number of RIP routes:1
```

Table 7-8 Description of the display rip neighbor command output

Item	Description
IP Address	IP address of the neighboring interface
Interface	Interface type
Type	Protocol used to establish adjacencies with neighbors
Last-Heard-Time	Time since the last time packets are received from neighbors
Number of RIP routes	Number of RIP routes

Display detailed information about the neighbor in RIP process 1.


```
<HUAWEI> display rip 1 neighbor verbose
------------------------------------------------------------------------
 IP Address      Interface        Type   Last-Heard-Time
------------------------------------------------------------------------
 10.1.1.1        Vlanif100    RIP    0:0:17
   Number of Active routes       : 1
   Number of routes in garbage   : 0
   Last Received Sequence Number : 0x0
```

Table 7-9 Description of the display rip neighbor verbose command output

Item	Description
Number of Active routes	Number of routes in the active state
Number of routes in garbage	Number of routes in the garbage state
Last Received Sequence Number	Sequence number of last received packet from neighbor

7.2.15 display rip neighbor last-nbr-down

Function

The **display rip neighbor last-nbr-down** command displays information about the last neighbor that goes Down in a RIP process.

Format

display rip *process-id* **neighbor last-nbr-down**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The **display rip neighbor last-nbr-down** command can be used to view information about the last 10 RIP neighbor that goes down and the reason why the neighbor goes down.

Example

View information about the last neighbor that goes Down in a RIP process.

```
<HUAWEI> display rip 1 neighbor last-nbr-down
 Neighbor IP Address             : 10.2.2.2
 Interface                       : Vlanif10
 Reason for Neighbor down        : Interface Down
 Time at which neighbor went down : 2011-06-02 11:53:50
```

Table 7-10 Description of the display rip neighbor last-nbr-down command output

Item	Description
Neighbor IP Address	IP address of a neighbor
Interface	Interface connected to a neighbor
Reason for Neighbor down	Reason that a neighbor goes Down: • Unknown • Interface Down • Configuration Change • Time Out (Normal) • Time Out (Message Processing Failed) • BFD Session Down • Invalid Packet Received • Received Worst Metric Routes • Received NQA Down Notification
Time at which neighbor went down	Time when a neighbor goes Down

7.2.16 display rip route

Function

The **display rip route** command displays all the RIP routes that are learned from other switches and view the values of different timers related to each route.

Format

display rip *process-id* **route**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After a RIP process is enabled on all devices that need to run RIP and the **network** command is run to enable RIP on the related network segment, the **display rip route** command can be used to view RIP routing information.

Example

Display all the RIP routes and the values of different timers related to each route.

```
<HUAWEI> display rip 1 route
 Route Flags: R - RIP
              A - Aging, G - Garbage-collect
 ------------------------------------------------------------------------------
 Peer 192.168.5.1  on Vlanif100
      Destination/Mask        Nexthop     Cost   Tag     Flags   Sec
       172.16.0.0/16         192.168.5.1    1    0       RA      15
     192.168.14.0/24          192.168.5.1    2    0       RA      15
```

Table 7-11 Description of the display rip route command output

Item	Description
Route Flags	Route flags. First character indicating that the route is a RIP or TRIP route, Second character indicating the status of the route: • RA: indicates that a RIP route is active. • RG: indicates that a RIP route is inactive and that the Garbage-Collect timer has been started.
Destination/ Mask	Destination IP Address and its mask value
Nexthop	Next hop of the route
Cost	Metric value of the route

Item	Description
Tag	Tag that is used to differentiate internal RIP routes from external routes ● 0: indicates that a RIP route is internal route. ● 1: indicates that a RIP route is external route.
Sec	Time during which a route remains in a specific state

7.2.17 display rip statistics interface

Function

The **display rip statistics interface** command displays statistics on a RIP interface, including the number of packets sent and received on the interface.

Format

display rip *process-id* **statistics interface** { **all** | *interface-type interface-number* [**verbose** | **neighbor** *neighbor-ip-address*] }

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535.
all	Displays the statistics of all interfaces.	-
interface-type interface-number	Specifies the type and number of an interface.	-
verbose	Displays detail information about interface statistics.	-
neighbor *neighbor-ip-address*	Specifies the IP address of a neighbor.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display rip statistics interface** command to check packet statistics on a RIP interface, including the number of packets sent and received on the interface.

Example

Display statistics on VLANIF100 in RIP process 1.

```
<HUAWEI> display rip 1 statistics interface vlanif 100
Vlanif100(10.0.0.11)
Statistical information          Last min     Last 5 min    Total
------------------------------------------------------------------------
Periodic updates sent            5            23            259
Triggered updates sent           5            30            408
Response packets sent            10           34            434
Response packets received        15           38            467
Response packets ignored         0            0             0
Request packets sent             1            3             8
Request packets received         4            20            40
Request packets ignored          0            0             0
Bad packets received             0            0             0
Routes received                  0            0             0
Routes sent                      0            0             0
Bad routes received              0            0             0
Packet authentication failed     0            0             0
Packet send failed               0            0             0
```

Table 7-12 Description of the display rip statistics interface command output

Item	Description
Statistical information	Packet type
Last min	Statistics within the last 1 minute
Last 5 min	Statistics within the last 5 minutes
Total	Total number of packets
Periodic updates sent	Number of periodic Update packets that are sent on the interface
Triggered updates sent	Number of triggered Update packets that are sent on the interface
Response packets sent	Number of RIP Response packets that are sent on the interface
Response packets received	Number of RIP Response packets that are received on the interface
Response packets ignored	Number of RIP Response packets that are ignored on the interface
Request packets sent	Number of RIP Request packets that are sent on the interface
Request packets received	Number of RIP Request packets that are received on the interface

Item	Description
Request packets ignored	Number of RIP Request packets that are ignored on the interface
Bad packets received	Number of received packets that cannot be parsed correctly
Routes received	Number of received routes
Routes sent	Number of sent routes
Bad routes received	Number of received routes that cannot be parsed correctly
Packet authentication failed	Number of packets that fail to pass authentication
Packet send failed	Number of RIP packets that fail to be sent

7.2.18 filter-policy export (RIP)

Function

The **filter-policy export** command configures egress filtering policy for RIP routes.

The **undo filter-policy export** command deletes the filtering policy.

By default, no filtering policy is configured.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } **export** [*protocol* [*process-id*] | *interface-type interface-number*]

undo filter-policy [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name*] **export** [*protocol* [*process-id*] | *interface-type interface-number*]

Parameters

Parameter	Description	Value
acl-number	Specifies the number of the basic ACL that is used to filter the destination addresses of routes.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the ACL filtering route destination addresses.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
ip-prefix *ip-prefix-name*	Specifies the name of the IP prefix list that is used to filter the destination addresses of routes.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
protocol process-id	Specifies the name of the routing protocol that exports the routes and the protocol process ID. The imported route type can be **static**, **direct**, **rip**, **ospf**, **bgp**, **unr** and **isis**. The process ID must be specified when routes are imported from **isis**, **rip** or **ospf**.	The value is an integer that ranges from 1 to 65535.
interface-type interface-number	Specifies type and number of the interface based on which routes are filtered.	-

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Only filtered routes can be added to routing tables and advertised in Update packets.

When you remove an interface-based filtering policy, you must specify *interface-type interface-number* in the **undo filter-policy export** command. The policy on only one interface can be deleted each time.

This command runs in the RIP view. If the filtering policy is based on interfaces or protocols, each interface or protocol can be configured with only one filtering policy. If the interface or protocol is not specified, a global filtering rule is configured. Only one global filtering rule can be configured each time. If you run the command again when a global rule exists, the new policy overwrites the old one.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Example

Filter the imported static routes based on the IP prefix list named abc. Then, only the routes that match the filtering policy are added into the RIP routing table and then sent through RIP Update packets.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] filter-policy ip-prefix abc export static
```

Filter the routes imported from IS-IS process 1 based on ACL 2002. Only the routes that match the filtering policy are added into the RIP routing table and then sent through RIP Update packets.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] filter-policy 2002 export isis 1
```

7.2.19 filter-policy import (RIP)

Function

The **filter-policy import** command filters the received RIP routes.

The **undo filter-policy import** command cancels route filtering.

By default, no filtering policy is configured.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* [**gateway** *ip-prefix-name*] } **import** [*interface-type interface-number*]

filter-policy gateway *ip-prefix-name* **import**

undo filter-policy [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* [**gateway** *ip-prefix-name*]] **import** [*interface-type interface-number*]

undo filter-policy [**gateway** *ip-prefix-name*] **import**

Parameters

Parameter	Description	Value
acl-number	Specifies the number of the basic ACL that is used to filter the destination addresses of routes.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the ACL filtering route destination addresses.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
interface-type interface-number	Indicates the type and the number of the interface based on which routes are filtered.	-
ip-prefix	Filters routes by using the IP prefix list.	-
ip-prefix-name	Specifies the address prefix list used to filter route destination addresses.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
gateway	Filters routes based on the gateway.	-

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

The **filter-policy import** command filters received RIP routes, including:

- Filtering a specified route in the packet
- Not monitoring the Update packets from the device

To cancel route filtering, run the **undo filter-policy import** command. When you remove an interface-based filtering policy, you must specify *interface-type interface-number*. The policy on only one interface can be deleted each time.

This command runs in the RIP view. If the filtering policy is based on interfaces, each interface can be configured with only one filtering policy. If the interface is not specified, a global filtering policy is configured. Only one global filtering policy can be configured. If you run the command again when a global policy exists, the new policy overwrites the old one.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule takes effective only when the source address range is specified by the **source** parameter and the time period is specified by the **time-range** parameter.

Example

Filter the RIP Update packets that are received from all the interfaces according to the IP prefix list named **abc**.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] filter-policy ip-prefix abc import
```

7.2.20 graceful-restart (RIP)

Function

The **graceful-restart** command enables RIP GR on the restarter.

The **undo graceful-restart** command disables RIP GR on the restarter.

By default, RIP GR is disabled.

Format

graceful-restart [**period** *period* | **wait-time** *time* | **planned-only** *time*] *

undo graceful-restart

Parameters

Parameter	Description	Value
period *period*	Specifies the period of GR.	It is an integer ranging from 30 to 3600, in seconds. The default value is 180 seconds.
wait-time *time*	Specifies the period of GR performed between the restarter and the neighbors that do not support GR. The value of **wait-time** *time* cannot be greater than the value of **period** *period*.	It is an integer ranging from 1 to 3600, in seconds. The default value is 45 seconds.
planned-only *time*	Specifies the period of GR performed between the restarter and the neighbors that support GR.	It is an integer ranging from 5 to 3600, in seconds. The default value is 60 seconds.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Planned GR indicates the master/slave switchover triggered through the command. Unplanned GR indicates the master/slave switchover triggered because of a fault.

When most switches on a network do not support RIP GR, you are advised to set **wait-time** *time* to a larger value. This ensures that the restarter has enough time to learn correct routes.

If the restart switch completes GR before the period specified by **period** *period* expires, the restart switch automatically exits from GR. If the restart switch does not complete GR after the period specified by **period** *period* expires, the restart switch is forced to exit from GR.

If the period of GR performed between the restarter and the neighbors that do not support GR depends on **wait-time**, the restarter does not exit from GR regardless of whether GR is finished within **wait-time**. After **wait-time** expires, however, the restarter is forced to exit from GR.

Example

Enable RIP GR and set the GR period to 200s.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] graceful-restart period 200
```

7.2.21 host-route

Function

The **host-route** command adds host routes with 32 bits into the routing table.

The **undo host-route** command disables host routes with 32 bits from being added into the routing table.

By default, host routes can be added into the routing table.

Format

host-route

undo host-route

Parameters

None

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

This command applies to RIP-1. In RIP-2, host routes are added to the routing table, regardless of whether the command is run.

Example

Add host routes into the routing table.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] host-route
```

7.2.22 import-route (RIP)

Function

The **import-route** command configures RIP to import routes from other routing protocols.

The **undo import-route** command disables RIP from importing routes from other routing protocols.

By default, RIP does not import routes from other routing protocols.

Format

import-route bgp [**permit-ibgp**] [**cost** { *cost* | **transparent** } | **route-policy** *route-policy-name*] *

import-route { { **static** | **direct** | **unr** } | { { **rip** | **ospf** | **isis** } [*process-id*] } } [**cost** *cost* | **route-policy** *route-policy-name*] *

undo import-route { { **static** | **direct** | **bgp** | **unr** } | { { **rip** | **ospf** | **isis** } [*process-id*] } }

Parameters

Parameter	Description	Value
bgp | **static** | **direct** | **rip** | **ospf** | **isis** | **unr**	Specifies the routing protocol from which RIP imports routes.	-
permit-ibgp	Imports routes from IBGP to RIP.	-
process-id	Specifies a process ID. The process ID can be specified when you import routes from RIP, OSPF, or IS-IS to RIP.	The value is an integer that ranges from 1 to 65535.
cost *cost*	Specifies the metric for imported routes.	The value is an integer that ranges from 0 to 15.

Parameter	Description	Value
cost transparent	This parameter is valid only when RIP imports routes from BGP. The metric of the imported routes is the MED of the BGP routes.	-
route-policy *route-policy-name*	Specifies the routing policy for importing routes.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

After you run the **import-route** command to set a routing policy, RIP can import only the specified routes and sets the route attributes.

To import routes from RIP, or IS-IS, you can specify the process ID. If the process ID is specified, RIP imports routes only from the specified process. If no process ID is specified, RIP imports routes from all processes. When a static or direct route is imported, the process ID cannot be specified.

You can run one of the following commands to set the metric of the imported routes. The following commands are listed in the descending order of priorities.

- Run the **apply cost** command to set route metric.
- Run the **import-route** (RIP) command to set the metric for imported routes.
- Run the **default-cost (RIP)** command to set the default metric for routes.

NOTE

- RIP defines a 16-bit tag, while other routing protocols define 32-bit tags. If the routes of other protocols are imported to RIP and the tag is used in the routing policy, the tag value cannot exceed 65535. If the tag value exceeds 65535, the routing policy becomes invalid or the matching result is incorrect.
- After the **import-route direct** command is executed, routes to the network segment where the IP address of the management interface belongs are also imported in the RIP routing table. Therefore, use this command with caution.
- The **permit-ibgp** parameter cannot be set for the RIP VPN instance.
- Importing IBGP routes to RIP may cause routing loops.

Example

Import IBGP routes, which are filtered by policy abc, to RIP process 1 and set the metric to 5.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] import-route bgp permit-ibgp cost 5 route-policy abc
```

7.2.23 maximum load-balancing (RIP)

Function

The **maximum load-balancing** command configures the maximum number of equal-cost routes for load balancing.

The **undo maximum load-balancing** command restores the default setting.

By default, the maximum number of equal-cost routes on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S6720S-S, and S5736-S is 8, and the maximum number of equal-cost routes on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S is 16.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

maximum load-balancing *number*

undo maximum load-balancing

Parameters

Parameter	Description	Value
number	Specifies the number of equal-cost routes.	The value is an integer that ranges from 1 to 8 on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S and S6720S-S. The value ranges from 1 to 16 on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

If a network has multiple redundant links, the maximum number of equal-cost routes can be configured to implement load balancing. With load balancing, network resources are fully utilized, situations where some links are overloaded while others are idle can be avoided, and delay in packet transmissions is shortened.

Example

Set the maximum number of equal-cost routes to 4.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] maximum load-balancing 4
```

7.2.24 network (RIP)

Function

The **network** command enables RIP for the interface on the specified network segment.

The **undo network** command disables RIP for the interface on the specified network segment.

By default, RIP is disabled for the interface on the specified network segment.

Format

network *network-address*

undo network *network-address*

Parameters

Parameter	Description	Value
network-address	Specifies the network address on which RIP is enabled. It must be the address of the natural network segment.	The value is in dotted decimal notation.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

An interface can be associated with only one RIP process.

Example

Enable RIP routes for the interface on the specified network segment.

```
<HUAWEI> system-view
[HUAWEI] rip
[HUAWEI-rip-1] network 10.0.0.0
```

7.2.25 peer (RIP)

Function

The **peer** command specifies the IP addresses of the RIP neighbors. After this command is configured, Update packets are sent to the peer in unicast instead of multicast or broadcast mode.

The **undo peer** command deletes the specified neighbor IP address.

By default, no RIP neighbor address is specified.

Format

peer *ip-address*

undo peer *ip-address*

Parameters

Parameter	Description	Value
ip-address	Specifies the IP address of a neighbor.	The value is in dotted decimal notation.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Generally, you are advised to use the **peer** command because an interface may receive the same packets both in multicast (or broadcast) and unicast mode. You are advised to change the mode of the interface to the silent mode when configuring the **peer** command.

Example

Specify the IP address of the neighbor to 10.0.0.1.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] peer 10.0.0.1
```

7.2.26 preference (RIP)

Function

The **preference** command sets the preference for RIP routes.

The **undo preference** command restores the default preference of RIP routes.

The default preference for RIP routes is 100.

Format

preference { *preference* | **route-policy** *route-policy-name* } *

undo preference

Parameters

Parameter	Description	Value
preference	Specifies the preference for routes.	The value is an integer ranging from 1 to 255. By default, it is 100.
route-policy *route-policy-name*	Specifies the routing policy that sets preference for the routes meeting conditions.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

A small value indicates a high preference. To enable RIP routes to have a higher preference than the routes learned by other IGP protocols, you need to configure a smaller preference value for the RIP routes. The preference determines the algorithm through which the optimal route is obtained among the routes in the IP routing table.

Example

Set the preference of RIP routes to 120.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] preference 120
```

Set the preference of RIP routes that match the routing policy named **rt-policy1** to 120.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] preference 120 route-policy rt-policy1
```

7.2.27 reset rip configuration

Function

The **reset rip configuration** command resets system parameters for a specified RIP process. When a RIP process starts, all the parameters of the process retain the default values.

Format

reset rip *process-id* **configuration**

Parameters

Parameter	Description	Value
process-id	Specifies a RIP process ID.	The value is an integer that ranges from 1 to 65535.

Views

User view

Default Level

3: Management level

Usage Guidelines

NOTICE

Restarting RIP processes may interrupt services. Exercise caution when you run this command.

After the command is executed, RIP neighbor relationships are set up again and learned routes are deleted.

Example

Reset the parameters of RIP process 100.

```
<HUAWEI> reset rip 100 configuration
```

7.2.28 reset rip statistics

Function

The **reset rip statistics** command resets the counter that is maintained by a particular RIP process. This command allows you to repeatedly record statistics during debugging.

Format

reset rip *process-id* **statistics** [**interface** { **all** | *interface-type interface-number* [**neighbor** *neighbor-ip-address*] }]

Parameters

Parameter	Description	Value
process-id	Specifies a RIP process ID.	The value is an integer ranging from 1 to 65535.
all	Clears statistics of all RIP processes.	-
interface all	Clears statistics on all the interfaces of a specified RIP process.	-
interface *interface-type interface-number*	Clears statistics on the interface with specified number and type.	-
neighbor *neighbor-ip-address*	Clears statistics on the RIP process between the local RIP interface and the specified neighbor.	The value is in dotted decimal notation.

Views

User view

Default Level

3: Management level

Usage Guidelines

Statistics cannot be restored after being cleared. Exercise caution when you run this command.

Example

Clear statistics in RIP process 100.

```
<HUAWEI> reset rip 100 statistics
```

Clear statistics on all interfaces in RIP process 100.

```
<HUAWEI> reset rip 100 statistics interface all
```

7.2.29 rip

Function

The **rip** command starts the specified RIP in the system view.

The **undo rip** command stops the specified RIP process.

By default, no RIP process is started.

Format

rip [*process-id*] [**vpn-instance** *vpn-instance-name*]

undo rip *process-id*

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer that ranges from 1 to 65535. The default value is 1.
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

If no VPN instance is specified, a RIP process is run globally or in the default VPN instance. Before configuring each global RIP parameter, you need to start RIP. This restriction is not applicable when you configure interface-related parameters.

Example

Start RIP process 1.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1]
```

Start RIP process 100 in the VPN instance named **abc**.

```
<HUAWEI> system-view
[HUAWEI] rip 100 vpn-instance abc
[HUAWEI-rip-100]
```

7.2.30 rip authentication-mode

Function

The **rip authentication-mode** command sets the RIP-2 authentication mode and authentication parameters. Only one authentication password is used for each authentication. If multiple authentication passwords are configured, the latest one takes effect.

The **undo rip authentication-mode** command cancels authentication.

By default, no authentication is configured.

Format

rip authentication-mode simple { **plain** *plain-text* | [**cipher**] *password-key* }

rip authentication-mode keychain *keychain-name*

rip authentication-mode md5 usual { **plain** *plain-text* | [**cipher**] *password-key* }

rip authentication-mode md5 nonstandard { **keychain** *keychain-name* | { **plain** *plain-text* | [**cipher**] *password-key* } *key-id* }

rip authentication-mode hmac-sha256 { **plain** *plain-text* | [**cipher**] *password-key* } *key-id*

undo rip authentication-mode

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
simple	Specifies simple authentication. **NOTE** Simple authentication mode carries potential risks. HMAC-SHA256 ciphertext authentication is recommended.	-
md5	Specifies MD5 authentication. **NOTE** MD5 ciphertext authentication carries potential risks. HMAC-SHA256 ciphertext authentication is recommended.	-
usual	Indicates that MD5 ciphertext authentication packets use the universal format (private standard).	-
nonstandard	Indicates that MD5 ciphertext authentication packets use non-standard packet format (IETF standard).	-
plain	Indicates that only plain text can be entered and only plain text is displayed when the configuration file is viewed. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. You are advised to select **cipher** to save the password in cipher text.	-

Parameter	Description	Value
plain-text	Specifies the authentication password that is displayed in plain text.	The value is a string of case-sensitive characters. It contains letters and digits without spaces. When the authentication mode is **simple** or **md5 usual**, the password consists of 1 to 16 characters. When the authentication mode is **md5 nonstandard** or **hmac-sha256**, the password consists of 1 to 255 characters.
cipher	Indicates that either plain text or cipher text can be entered and cipher text is displayed when the configuration file is viewed.	-
password-key	Specifies the authentication password that is displayed in cipher text.	The value is a string of case-sensitive characters. It contains letters and digits without spaces. • When the authentication mode is **simple** or **md5 usual**, the password is in plain text (1 to 16 characters) or in cipher text (24 or 32 or 48 characters). If the source version supports a ciphertext password which is a string of 24 or 32 characters, the target version also supports this type of password. • When the authentication mode is **md5 nonstandard** or **hmac-sha256**, the password is in plain text (1 to 255 characters) or in cipher text (20 to 392 characters).
keychain *keychain-name*	Specifies keychain authentication.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.

Parameter	Description	Value
key-id	Specifies the identifier of Cryptographic authentication.	The value is an integer that ranges from 1 to 255.
hmac-sha256	Indicates Keyed-Hash Message Authentication Code (HMAC) for Secure Hash Algorithm 256 (SHA256).	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Keychain authentication improves UDP connection security. Keychain authentication must be configured on both ends of a link. Encryption algorithms and passwords configured on both ends must be the same; otherwise, the UDP connection cannot be set up and RIP messages cannot be transmitted.

Example

Set HMAC-SHA256 authentication on VLANIF100, with the authentication password **YsHsjx_202206** and key-id **255**.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip authentication-mode hmac-sha256 cipher YsHsjx_202206 255
```

Set HMAC-SHA256 authentication on GE0/0/1, with the authentication password **YsHsjx_202206** and key-id **255**.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip authentication-mode hmac-sha256 cipher YsHsjx_202206 255
```

7.2.31 rip bfd

Function

The **rip bfd** command sets BFD session parameters on the specified interface.

The **undo rip bfd** command restores BFD session parameters set on the specified interface.

By default, the minimum intervals for receiving and sending BFD packets are 1000 ms and the detection time multiplier is 3.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

rip bfd { **min-rx-interval** *min-receive-value* | **min-tx-interval** *min-transmit-value* | **detect-multiplier** *detect-multiplier-value* } *

undo rip bfd { **min-rx-interval** [*min-receive-value*] | **min-tx-interval** [*min-transmit-value*] | **detect-multiplier** [*detect-multiplier-value*] } *

Parameters

Parameter	Description	Value
min-rx-interval *min-receive-value*	Specifies the minimum interval at which BFD packets are received from the remote device.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *min-transmit-value*	Specifies the minimum interval at which BFD packets are sent to the remote device.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.

Parameter	Description	Value
detect-multiplier *detect-multiplier-value*	Specifies the local detection multiplier.	The value is an integer ranging from 3 to 50.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The *min-receive-value* value is obtained after the local **min-rx-interval** value and a neighbor's **min-tx-interval** value are negotiated. For detailed negotiation policies, see **bfd all-interfaces**. If the switch does not receive BFD packets from the neighbor at the interval of *min-receive-value* x *detect-multiplier-value*, the neighbor will go Down.

Prerequisites

Enable BFD globally and run the **rip bfd enable** command before establishing a BFD session.

Precautions

BFD session parameters configured in a RIP process take effect only after BFD is enabled on the interface.

The BFD priority configured on an interface is higher than the BFD priority configured in a RIP process. If BFD session parameters are configured on an interface, establish a BFD session based on the configured parameters.

Example

Enable BFD on VLANIF100 and set the minimum sending interval to 600 ms and local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip bfd enable
[HUAWEI-Vlanif100] rip bfd min-tx-interval 600 detect-multiplier 4
```

Enable BFD on GE0/0/1 and set the minimum sending interval to 600 ms and local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
```



```
[HUAWEI-GigabitEthernet0/0/1] rip bfd enable
[HUAWEI-GigabitEthernet0/0/1] rip bfd min-tx-interval 600 detect-multiplier 4
```

7.2.32 rip bfd block

Function

The **rip bfd block** command blocks BFD on a specified interface.

The **undo rip bfd block** command disables the blocking function.

By default, the blocking function is disabled.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

rip bfd block

undo rip bfd block

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

This command can be used to block the specified interface from enabling BFD on some links where no BFD session is needed.

Precautions

The **rip bfd block** and **rip bfd enable** commands are mutually exclusive. If both of them are configured, only the later configured one takes effect.

Example

Block BFD on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip bfd block
```

Block BFD on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip bfd block
```

7.2.33 rip bfd enable

Function

The **rip bfd enable** command enables BFD on the specified interface to establish a BFD session with default parameters.

The **undo rip bfd enable** command disables BFD on the specified interface.

By default, BFD is disabled on a RIP interface.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

rip bfd enable

undo rip bfd enable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

If BFD is not enabled globally, you can set BFD parameters on the specified interface but you cannot establish a BFD session on this interface.

The **rip bfd block** and **rip bfd enable** commands are mutually exclusive. If both of them are configured, only the later configured one takes effect.

Example

Enable BFD on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip bfd enable
```

Disable BFD on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] undo rip bfd enable
```

Enable BFD on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip bfd enable
```

Disable BFD on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] undo rip bfd enable
```

7.2.34 rip bfd static

Function

The **rip bfd static** command enables static BFD on a specified RIP interface.

The **undo rip bfd static** command disables static BFD on a specified RIP interface.

By default, static BFD is disabled on a RIP interface.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

rip bfd static

undo rip bfd static

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a network deployed with high-speed data services, if a fault occurs on a link, a large amount of data is lost because it takes too long time for RIP to detect the fault. Deploying BFD for RIP to accelerate fault detection is necessary.

The **rip bfd static** command is used to enable BFD for RIP on a specified link to rapidly detect the fault on the link.

In addition, because many devices do not support BFD on the live network, this command can be also used to implement BFD between a BFD-capable device and a BFD-incapable device.

Prerequisites

BFD has been enabled globally using the **bfd** command.

Precautions

If **rip bfd static**, **rip bfd enable**, and **rip bfd block** are simultaneously configured, the latest configuration overrides the previous ones.

Example

Enable static BFD on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] bfd
```



```
[HUAWEI-bfd] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip bfd static
```

Enable static BFD on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip bfd static
```

7.2.35 rip input

Function

The **rip input** command enables the specified interface to receive RIP packets.

The **undo rip input** command disables the specified interface from receiving RIP packets.

By default, interfaces can receive RIP packets.

Format

rip input

undo rip input

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The priority of **silent-interface** is higher than the priority of **rip input** or **rip output** configured in the interface view. By default, an interface does not work in the silent state.

Example

Enable the specified interface VLANIF100 to receive RIP packets.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip input
```

Enable the specified interface GE0/0/1 to receive RIP packets.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
```



```
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip input
```

7.2.36 rip metricin

Function

The **rip metricin** command sets the metric that is added to the route when an interface receives a RIP packet.

The **undo rip metricin** command restores the additional metric to the default value.

By default, there is no metric added to the route when an interface receives a RIP packet.

Format

rip metricin { *value* | { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } *value1* }

undo rip metricin

Parameters

Parameter	Description	Value
value	Specifies the metric that is added to the received route.	The value is an integer ranging from 0 to 15. By default, it is 0.
acl-number	Specifies the basic ACL number.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the ACL name.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
value1	Specifies the metric that is added to the route that passes the filtering of the ACL or IP prefix list.	The value is an integer ranging from 1 to 15.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

When an interface receives a route, RIP adds the additional metric of the interface to the route, and then installs the route into the routing table. Therefore, increasing the metric of an interface also increases the metric of the RIP route received by the interface.

Adjust RIP route selection by increasing the metrics of received routes.

Example

Set the additional metric to 12 when VLANIF100 receives RIP routes.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip metricin 12
```

Set the additional metric to 12 when GE0/0/1 receives RIP routes.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip metricin 12
```

Set the additional metric to 12 using acl-name.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip metricin acl-name abcd 12
```

Set the additional metric to 12 using ip-prefix.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip metricin ip-prefix ip1 12
```

7.2.37 rip metricout

Function

The **rip metricout** command sets the metric that is added to the route when an interface sends a RIP packet.

The **undo rip metricout** command restores the additional metric to the default value.

By default, the metric that is added to the route when an interface sends a RIP packet is 1.

Format

rip metricout { *value* | { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } *value1* }

undo rip metricout

Parameters

Parameter	Description	Value
value	Specifies the metric that is added to the sent route.	The value is an integer ranging from 1 to 15. By default, it is 1.
acl-number	Specifies the number of a basic ACL.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the name of an ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of the IP prefix list. The name must be exclusive.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
value1	Specifies the metric that is added to the route that passes the filtering of the ACL or IP prefix list.	The value is an integer ranging from 2 to 15.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

When a RIP route is advertised, the additional metric is added to the route. Therefore, increasing the metric of an interface also increases the metric of the RIP route sent on the interface. However, the metric of the route in the routing table remains unchanged.

You can specify the metric to be added to the RIP route that passes the filtering of the ACL or IP prefix list by specifying *value1*. If a RIP route does not pass the filtering, its metric is increased by 1.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Example

Set the metric that is added when the interface sends RIP routes to 12.


```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip metricout 12
```

Set the metric that is added when the interface sends RIP routes to 12.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip metricout 12
```

Increase the metric of a RIP route that passes the filtering of ACL 2050 by 12.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip metricout 2050 12
```

Increase the metric of a RIP route that passes the filtering of ACL 2050 by 12.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip metricout 2050 12
```

Increase the metric of a RIP route that passes the filtering of the IP prefix list named **p1** by 12.

```
<HUAWEI> system-view
[HUAWEI] ip ip-prefix p1 permit 10.10.10.1 24
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip metricout ip-prefix p1 12
```

Increase the metric of a RIP route that passes the filtering of the IP prefix list named **p1** by 12.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip metricout ip-prefix p1 12
```

7.2.38 rip mib-binding

Function

The **rip mib-binding** command sets the binding between the Management Information Base (MIB) and RIP process ID, and specifies the ID of the RIP process that receives SNMP requests.

The **undo rip mib-binding** command cancels the binding.

By default, there is no binding between the Management Information Base and RIP process ID.

Format

rip mib-binding *process-id*

undo rip mib-binding

Parameters

Parameter	Description	Value
process-id	Specifies the ID of a RIP process.	The value is an integer ranging from 1 to 65535. The default value is 1.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

All SNMP requests are sent to the bound RIP process.

Prerequisites

A RIP process has been created using the **rip** command.

Example

Configure RIP process 100 to receive SNMP requests.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] quit
[HUAWEI] rip mib-binding 100
```

7.2.39 rip output

Function

The **rip output** command enables an interface to send RIP packets.

The **undo rip output** command disables an interface from sending RIP packets.

By default, an interface can send RIP packets.

Format

rip output

undo rip output

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

When a device running RIP is connected to a network running other routing protocols, you can run the **rip output** command on the interface that connects the device to the network to prevent the interface from sending useless packets to the network.

The priority of **silent-interface** is higher than the priority of **rip input** or **rip output** that is configured in the interface view. By default, an interface does not work in the silent state.

Example

Enable the interface VLANIF100 to send RIP packets.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip output
```

Enable the interface GE0/0/1 to send RIP packets.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip output
```

7.2.40 rip pkt-transmit

Function

The **rip pkt-transmit** command sets the interval for sending Update packets and the number of packets sent each time on the specified interface.

The **undo rip pkt-transmit** command restores the default values on an interface.

By default, the interval for sending RIP Update packets is 200 ms and 50 packets are sent each time.

Format

rip pkt-transmit { **interval** *interval* | **number** *pkt-count* } *

undo rip pkt-transmit

Parameters

Parameter	Description	Value
interval *interval*	Specifies the interval for sending Update packets.	The value is an integer ranging from 50 to 500, in milliseconds. The default value is 200.
number *pkt-count*	Specifies the number of packets sent each time.	The value is an integer and ranges from 25 to 100. The default value is 50.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The **rip pkt-transmit** command can be used on an interface to control the interval for sending Update packets and the number of sent packets. This improves RIP performance.

Example

Set the interval for sending packets on VLANIF100 to 100 milliseconds and the number of the packets sent each time to 50.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip pkt-transmit interval 100 number 50
```

Set the interval for sending packets on GE0/0/1 to 100 milliseconds and the number of the packets sent each time to 50.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip pkt-transmit interval 100 number 50
```

7.2.41 rip poison-reverse

Function

The **rip poison-reverse** command enables poison reverse.

The **undo rip poison-reverse** command disables poison reverse.

By default, poison reverse is disabled.

Format

rip poison-reverse

undo rip poison-reverse

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

After poison reverse is enabled, RIP sets the cost of the routes learned from a specified interface to 16 (indicating unreachable), and then sends the routes to neighbors through the same interface.

When both split horizon and poison reverse are configured, only poison reverse takes effect.

Example

Enable poison reverse.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip poison-reverse
```

Enable poison reverse on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip poison-reverse
```

7.2.42 rip replay-protect

Function

The **rip replay-protect** command enables the replay-protect function.

The **undo rip replay-protect** command disables the replay-protect function.

By default, the replay-protect function is disabled.

Format

rip replay-protect [*window-range*]

undo rip replay-protect

Parameters

Parameter	Description	Value
window-range	Specifies the size of the connection-oriented transmission buffer.	It is an integer ranging from 50 to 50000. The default value is 50.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

If the Identification field in the last RIP packet sent before a RIP interface goes Down is X, after the interface goes Up, the Identification field in the RIP packet sent by this interface becomes 0. If the remote end does not receive the RIP packet with the Identification field being 0, subsequent RIP packets will be discarded until the remote end receives the RIP packet with the Identification field being X+1. As a result, RIP routing information of both ends is inconsistent. To solve this problem, you need to configure the **rip replay-protect** command so that RIP can obtain the Identification field in the RIP packet sent before the RIP packet goes Down and increases the Identification field of the subsequently sent RIP packet by one.

Before configuring the **rip replay-protect** command, you need to configure the **rip authentication-mode** command in the RIP interface view to configure the authentication mode and authentication parameters of RIP-2.

NOTE

If you configure the **rip replay-protect** command in the same view for multiple times, only the last configuration takes effect.

Example

Enable the replay-protect function on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip replay-protect
```

Enable the replay-protect function on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip replay-protect
```


7.2.43 rip split-horizon

Function

The **rip split-horizon** command enables split horizon.

The **undo rip split-horizon** command disables split horizon.

By default, split horizon is enabled.

Format

rip split-horizon

undo rip split-horizon

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

If a route is advertised through the interface from which the route is learned, the route is blocked by split horizon. Split horizon is used to avoid routing loops between neighboring devices.

Generally, it is not recommended that you disable split horizon.

If split horizon is enabled on the interface that is configured with secondary IP addresses, RIP Update packets may not be sent by each secondary address. An Update packet does not regard every network as the source unless split horizon is disabled.

If an interface is connected to a Non Broadcast Multiple Access (NBMA) network, split horizon on the interface is disabled by default.

If both poison reverse and split horizon are configured, simple split horizon (the route is suppressed by the interface through which the route is learned) is substituted by poison reverse. Here, simple split horizon means that the route is suppressed when it is advertised through the interface from which it is learned.

Example

Enable split horizon.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip split-horizon
```


Enable split horizon.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip split-horizon
```

7.2.44 rip summary-address

Function

The **rip summary-address** command configures a RIP switch to advertise a local summarized IP address.

The **undo rip summary-address** command disables a RIP switch from advertising a local summarized IP address.

By default, a RIP switch does not advertise local summarized IP addresses.

Format

rip summary-address *ip-address mask* [**avoid-feedback**]

undo rip summary-address *ip-address mask*

Parameters

Parameter	Description	Value
ip-address	Specifies the network IP address to be summarized.	The value is in dotted decimal notation.
mask	Specifies the network mask.	The value is in dotted decimal notation.
avoid-feedback	Specifies that learning the same summarized route from an interface is not allowed.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

After the keyword **avoid-feedback** is specified, an interface no longer learns the summarized route with the same IP address as the advertised summarized IP address. This avoids routing loops.

Example

Configure the switch to advertise a local summarized IP address.


```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip summary-address 10.0.0.0 255.0.0.0
```

Configure the switch to advertise a local summarized IP address.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip summary-address 10.0.0.0 255.0.0.0
```

7.2.45 rip valid-ttl-hops

Function

The **rip valid-ttl-hops** command enables the RIP GTSM functions and sets the TTL value to be detected.

The **undo rip valid-ttl-hops** command cancels the function.

By default, the RIP GTSM functions are disabled.

Format

rip valid-ttl-hops *valid-ttl-hops-value* [**vpn-instance** *vpn-instance-name*]

undo rip valid-ttl-hops [*valid-ttl-hops-value*] [**vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
valid-ttl-hops-value	Specifies the number of TTL hops to be detected. The valid TTL range of the detected packets is [255 -*valid-ttl-hops-value* + 1, 255].	The value is an integer ranging from 1 to 255.
vpn-instance *vpn-instance-name*	Specifies the name of the VPN instance. If this parameter is used, you need only to specify the TTL value to be detected by the VPN instance.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a network demanding higher security, you can enable GTSM to improve the security of the RIP network. GTSM defends against attacks by checking the TTL value. If an attacker simulates RIP unicast packets and keeps sending them to a switch, a switch receives the packets and directly sends them to the main control board for RIP processing, without checking the validity of the packets. In this case, the switch is busy processing these packets, causing high usage of the CPU. GTSM protects the routers and enhances the system security by checking whether the TTL value in the IP packet header is in a pre-defined range.

The **rip valid-ttl-hops** command is used to enable RIP GTSM.

Precautions

GTSM configurations must be symmetrical. That is, GTSM must be enabled on devices at both ends.

If GTSM is enabled on a device, after the device receives a RIP packet, it checks whether the TTL value in the packet is in a pre-defined range. If the TTL value is beyond the pre-defined range, the device considers the packet as an attack packet and discards it.

Example

Enable the RIP GTSM functions, and configure the maximum number of TTL hops to 5 for the packets that a switch is allowed to receive.

```
<HUAWEI> system-view
[HUAWEI] rip valid-ttl-hops 5
```

7.2.46 rip version

Function

The **rip version** command sets the RIP version of an interface.

The **undo rip version** command restores the default setting.

By default, an interface sends only RIP-1 packets, but it can receive both RIP-1 and RIP-2 packets.

Format

rip version { **1** | **2** [**broadcast** | **multicast**] }

undo rip version

Parameters

Parameter	Description	Value
1	Indicates RIP-1 packets.	-
2	Indicates RIP-2 packets.	-

Parameter	Description	Value
broadcast	Indicates that RIP-2 packets are sent in broadcast mode.	-
multicast	Indicates that RIP-2 packets are sent in multicast mode. **NOTE** By default, RIP-2 packets are sent in multicast mode.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The versions of the request packets and response packets vary with the configured RIP version.

- If the version of RIP is not set, a device sends RIP-1 packets in broadcast mode and receives the RIP-1 and RIP-2 packets that are sent in broadcast mode.
- If the RIP version is set to RIP-1, a device sends only RIP-1 packets in broadcast mode and receives the RIP-1 packets that are sent in broadcast mode.
- If the RIP version is set to RIP-2, a device sends only RIP-2 packets in multicast mode and receives RIP-2 packets that are sent in multicast or broadcast mode.
- If the RIP version is set to multicast RIP-2, a device sends RIP-2 packets in multicast mode and receives RIP-2 packets that are sent in multicast mode.
- If the RIP version is set to broadcast RIP-2, a device sends RIP-2 packets in broadcast mode and receives RIP-1 and RIP-2 packets.

Precautions

You can also set a RIP version in a RIP process, but the RIP version that is set on an interface has a higher priority.

Example

Send RIP-2 packets in broadcast mode.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] rip version 2 broadcast
```

Send RIP-2 packets in broadcast mode.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] rip version 2 broadcast
```


7.2.47 silent-interface (RIP)

Function

The **silent-interface** command controls an interface only to receive packets to update its routing table and prevent it from sending RIP packets.

The **undo silent-interface** command enables an RIP interface to send Update packets.

The **silent-interface disable** command enables an RIP interface to send Update packets.

The **undo silent-interface disable** command controls an interface only to receive packets to update its routing table and prevent it from sending RIP packets.

By default, silent interface not enabled.

Format

silent-interface { **all** | *interface-type interface-number* }

undo silent-interface { **all** | *interface-type interface-number* }

silent-interface disable *interface-type interface-number*

undo silent-interface disable *interface-type interface-number*

Parameters

Parameter	Description	Value
all	Indicates that all interfaces are suppressed.	-
interface-type interface-number	Specifies the type and the number of the interface.	-
disable	Disables the suppression of the RIP interface so that the interface can send Update packets.	-

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a switch running RIP is connected to a network running other routing protocols, you can run the **silent-interface** command on the interface that connects the switch to the network to prevent the interface from sending useless packets to the network.

The **silent-interface** command is used together with the **peer (RIP)** command to advertise routes to the specified equipment.

Configuration Impact

If an interface is suppressed, the direct routes of the network segment where the interface resides can still be advertised to other interfaces.

Precautions

When the **silent-interface** command is used to suppress the specified interface, the priority of the **silent-interface** command is higher than the priority of the **rip input** and **rip output** command that is configured in the interface view.

NOTE

After you configure all interfaces as silent interfaces, you can run the **silent-interface disable** *interface-type interface-number* command to activate a specified silent interface.

Example

Configure all the interfaces as silent interfaces.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] silent-interface all
```

Configure the RIP interface VLANIF100 as a silent interface and enable it to send routes to the neighbor with the IP address of 10.1.1.1/24.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] silent-interface vlanif 100
[HUAWEI-rip-100] peer 10.1.1.1
```

Configure all interfaces as silent interfaces, and then activate the interface VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] silent-interface disable vlanif 100
```

7.2.48 summary (RIP)

Function

The **summary** command enables RIP classful summarization. The summarized routes are advertised using natural masks.

The **undo summary** command disables classful summarization so that routing between subnets can be performed. The subnet information is then advertised. Route summarization reduces the routing table size.

By default, classful summarization is enabled for RIP-2.

Format

summary [**always**]

undo summary

Parameters

Parameter	Description	Value
always	Enables classful summarization no matter whether split horizon is configured.	-

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

RIP-1 does not support classful summarization. When using RIP-2, you can use the **undo summary** command to disable classful summarization.

RIP-2 route summarization improves scalability and efficiency on large networks. IP address summarization means that there is no sub-routing entry in the routing table. That is, there is no routing entry composed of single IP address. In addition to reducing the routing table size, route summarization enables the switch to handle more routes.

When classful summarization is enabled, the switch summarizes subnet addresses to the natural network segment border while advertising routes to the natural network segment border. When split horizon or poison reverse is enabled, route summarization will be invalid if the **always** parameter is not specified. When summarized routes are sent outside the natural network segment, split horizon or poison reverse must be disabled.

NOTE

- By default, classful summarization is enabled for RIP-2. If split horizon or poison reverse has been configured, classful summarization is invalid. When summarized routes are sent to the network border, split horizon and poison reverse must be disabled.
- The **summary always** command can enable classful summarization no matter whether split horizon or poison reverse is enabled.
- The summarization preference on interfaces is higher than the summarization preference in RIP processes. That is, the preference of **rip summary-address** is higher than the preference of **summary**. When summarization is configured on both interface and RIP process, the summarized route is advertised only when a few specific routes are beyond the summarization range configured on the interface.

Example

Enable RIP-2 classful summarization.

```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] version 2
[HUAWEI-rip-1] summary
```

Enable RIP-2 classful summarization when split horizon is enabled.


```
<HUAWEI> system-view
[HUAWEI] rip 1
[HUAWEI-rip-1] summary always
```

7.2.49 timers rip

Function

The **timers rip** command sets the values of RIP timers.

The **undo timers rip** command restores the values of RIP timers to the default value.

By default, the interval for sending Update packets is 30s, the time for aging routes is 180s, the time for deleting a route from the routing table is 120s.

Format

timers rip *update age garbage-collect*

undo timers rip

Parameters

Parameter	Description	Value
update	Specifies the interval for sending Update packets.	The value is an integer ranging from 1 to 86400, in seconds.
age	Specifies the time for aging routes.	The value is an integer ranging from 1 to 86400, in seconds.
garbage-collect	Specifies the time for deleting a route from the routing table, that is, the standard garbage collection time.	The value is an integer ranging from 1 to 86400, in seconds.

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

By adjusting the RIP timers, you can improve routing protocol performance to meet network requirements. If the values of the preceding three timers are set improperly, route flapping occurs.

The relationship of the values of the three timers is: *update* < *age* and *update* < *garbage-collect*. For example, if the update time is longer than the aging time,

switches cannot inform neighbors on time if RIP routes change during the update time.

NOTE

Generally, the default values of the timers do not need to be changed, and thus the **timers rip** command must be used with caution.

Example

Set values for RIP timers.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] timers rip 35 170 240
```

7.2.50 verify-source (RIP)

Function

The **verify-source** command enables check on source IP addresses in RIP Update packets.

The **undo verify-source** command disables check on source IP addresses in RIP Update packets.

By default, the source IP addresses in RIP Update packets are checked.

Format

verify-source

undo verify-source

Parameters

None

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **verify-source** command is executed, RIP checks whether the IP addresses of interfaces that send and receive the Update packets are on the same network segment. If not, the device does not process the packets.

NOTE

It is not recommended that you disable source address check.

Example

Enable source address check in RIP 100.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] verify-source
```

7.2.51 version (RIP)

Function

The **version** command specifies a global RIP version.

The **undo version** command restores the default global RIP version.

By default, an interface sends only RIP-1 packets, and receives both RIP-1 and RIP-2 packets.

Format

version { **1** | **2** }

undo version

Parameters

Parameter	Description	Value
1	Sets the global RIP version to RIP-1.	-
2	Sets the global RIP version to RIP-2.	-

Views

RIP view

Default Level

2: Configuration level

Usage Guidelines

Two versions are available for RIP: RIP-1 and RIP-2. RIP-2 is an extension to RIP-1. You can run the **version** command to specify a global RIP version.

Example

Send and receive RIP-2 packets.

```
<HUAWEI> system-view
[HUAWEI] rip 100
[HUAWEI-rip-100] version 2
```


7.3 RIPng Configuration Commands

7.3.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.3.2 checkzero (RIPng)

Function

The **checkzero** command enables zero field check for RIPng packets.

The **undo checkzero** command disables zero field check for RIPng packets.

By default, zero field check is enabled for RIPng packets.

Format

checkzero

undo checkzero

Parameters

None

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

After zero field check is enabled, the switch refuses to process the RIPng packets in which zero fields are not 0. If all RIPng packets are reliable and no zero field needs to be checked, run the **undo checkzero** command to reduce CPU usage.

Example

Disable zero field check for RIPng packets.

```
<HUAWEI> system-view
[HUAWEI] ripng 1
[HUAWEI-ripng-1] undo checkzero
```


7.3.3 default-cost (RIPng)

Function

The **default-cost** command sets the default cost for imported routes.

The **undo default-cost** command restores the default value 0.

By default, the default cost of RIPng routes is 0.

Format

default-cost *cost*

undo default-cost

Parameters

Parameter	Description	Value
cost	Sets the default cost of imported routes.	The value is an integer that ranges from 0 to 15.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

You can run one of the following commands to set the cost of imported routes. The following commands are listed in descending order of priorities.

- Run the **apply cost** command to set the cost of routes.
- Run the **import-route** command to set the cost of imported routes.
- Run the **default-cost** (RIPng) command to set the default cost of imported routes.

Example

Set the default cost of imported routes to 2.

```
<HUAWEI> system-view
[HUAWEI] ripng 100
[HUAWEI-ripng-100] default-cost 2
```

7.3.4 description (RIPng)

Function

The **description** command configures a description for a RIPng process.

The **undo description** command deletes the configured description of a RIPng process.

By default, there is no description for a RIPng process.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Configures a description for a RIPng process.	The description is a string of 1 to 80 case-sensitive characters that can contain spaces.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Configuring descriptions for RIPng processes helps identify and configure RIPng processes.

Precautions

If you run the **description** command multiple times, only the latest configuration takes effect.

Example

Configure a description for RIPng process 100.

```
<HUAWEI> system-view
[HUAWEI] ripng 100
[HUAWEI-ripng-100] description this process configure the poison reverse process
```

7.3.5 display default-parameter ripng

Function

The **display default-parameter ripng** command displays the default RIPng configuration.

Format

display default-parameter ripng

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

After the default RIPng configuration is modified, you can use the **display default-parameter ripng** command to view the default RIPng configuration used when RIPng is initialized.

Example

Display the default RIPng configuration.

```
<HUAWEI> display default-parameter ripng
-----------------------------------------------
 Protocol Level Default Configurations
-----------------------------------------------
      Preference    : 100
      Checkzero     : Enabled
      Default-cost  : 0
      Maximum Balanced Paths : 8
      Update time   : 30 sec     Age time   : 180 sec
      Garbage-collect time           : 120 sec
-----------------------------------------------
 Interface Level Default Configurations
-----------------------------------------------
   Metricin               : 0
   Metricout              : 1
   Input Packet Processing : Enabled
   Output Packet Processing: Enabled
   Poison Reverse         : Disabled
   Split Horizon
    For Broadcast and P2P Interfaces     :  Enabled
    For NBMA Interfaces and LoopBack     :  Disabled
   Default Route                        :  Disabled
   Packet Transmit Interval             :  200 msecs
   Packet Transmit Number               :  30
```

Table 7-13 Description of the display default-parameter ripng command output

Item	Description
Protocol Level Default Configurations	Default RIPng process configuration.
Preference	RIPng route preference.

Item	Description
Checkzero	Whether zero field check is enabled for RIPng packets.
Default Cost	Default cost of routes imported by RIPng from other routing protocols.
Maximum Balanced Paths	Maximum number of equal-cost routes for load balancing.
Update time	Interval for sending Update packets.
Age time	Aging time of RIPng routes.
Garbage-Collect time	Time from when a route is marked invalid until the route is removed from the routing table.
Interface Level Default Configurations	Default RIPng configuration on an interface.
Metricin	Metric that is added to a route when a RIPng packet is received.
Metricout	Metric that is added to a route when a RIPng packet is sent.
Poison Reverse	Whether poison reverse is enabled.
Split Horizon	Whether split horizon is enabled: • For Broadcast and P2P Interfaces: broadcast and P2P interfaces • For NBMA Interfaces and LoopBack: NBMA and loopback interfaces
Default-route	Default route. This route is used when no matching entry can be found for a packet in the routing table.
Packet Transmit Interval	Interval for forwarding packets, in milliseconds.
Packet Transmit Number	Number of forwarded packets.
Input Packet Processing	Whether this interface is enabled to receive RIPng packets
Output Packet Processing	Whether this interface is enabled to send RIPng packets

7.3.6 display ripng

Function

The **display ripng** command displays the current operating status and configuration of RIPng processes.

Format

display ripng [*process-id* | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
process-id	Displays the current operating status and configuration of the specified RIPng process. If this parameter is not specified, the configurations of all RIPng processes are displayed.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Displays the status and configuration of a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ripng** command to check the current operating status and configuration of RIPng processes.

Example

Display information about a specified RIPng process.

```
<HUAWEI> display ripng 100
Public vpn-instance
   RIPng process : 100
     Preference : 100
     Checkzero : Enabled
     Default-cost : 0
     Maximum number of balanced paths : 8
     Update time   : 30 sec   Age time  : 180 sec
     Garbage-collect time : 120 sec
     Number of periodic updates sent : 0
     Number of trigger updates sent : 1
     Number of routes in database : 1
     Number of interfaces enabled : 1
     Total number of routes : 3
```



```
    Total number of routes in ADV DB is : 0
```

Table 7-14 Description of the display ripng command output

Item	Description
Public vpn-instance	Public VPN.
RIPng process	RIPng process ID.
Preference	RIPng process preference. To set the preference for RIPng routes, run the **preference (RIPng)** command.
Checkzero	Whether zero field check is enabled for RIPng packets.
Default-cost	Default cost of routes imported by RIPng from other routing protocols.
Maximum number of balanced paths	Maximum number of equal-cost routes for load balancing. To set the preference for RIPng routes, run the **maximum load-balancing (RIPng)** command.
Update time	Interval for sending Update packets.
Age time	Aging time of RIPng routes.
Garbage-collect time	Time from when a route is marked invalid until the route is removed from the routing table.
Number of periodic updates sent	Number of RIPng Update packets sent periodically.
Number of trigger updates sent	Number of triggered RIPng Update packets.
Number of routes in database	Number of routes in the RIPng database.
Number of interfaces enabled	Number of RIPng interfaces.
Total number of routes	Total number of routes of a RIPng process.
Total number of routes in ADV DB is	Total number of advertised routes in the database.

7.3.7 display ripng database

Function

The **display ripng database** command displays all the active routes in the RIPng database.

Format

display ripng *process-id* **database** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Displays all the active routes of the specified RIPng process.	The value is an integer that ranges from 1 to 65535.
verbose	Displays detailed information about the routes in the RIPng database.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ripng database** command to check all the active routes in the RIPng database.

Example

Display routes in the RIPng database.

```
<HUAWEI> display ripng 100 database
1::/64,
      via FE80::82FB:6FF:FE35:45B6, Vlanif200, cost 0, Imported
  2::/64,
      Vlanif100, cost 0, RIPng-interface
  3::/64,
      via FE80::82FB:6FF:FE35:45B6, Vlanif200, cost 16
  3::/64,
      Vlanif200, cost 0, RIPng-interface
  4::/64,
      via FE80::82FB:6FF:FE35:45B6, Vlanif200, cost 1
```

Table 7-15 Description of the display ripng database command output

Item	Description
1::/64	Destination IPv6 address of a route.
via	Next-hop link-local address.
cost	Cost of a route.
Imported	Route imported from other routing protocols.

Item	Description
RIPng-interface	Route generated by RIPng.

7.3.8 display ripng interface

Function

The **display ripng interface** command displays information about RIPng interfaces.

Format

display ripng *process-id* **interface** [*interface-type interface-number*] [**verbose**]

Parameters

Parameter	Description	Value
process-id	Displays interface information in the specified RIPng process.	The value is an integer that ranges from 1 to 65535.
interface-type interface-number	Displays information on the interface with the specified type and number.	-
verbose	Displays detailed information about RIPng interfaces.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

This command displays the configuration and operating status of RIPng. The information is helpful in fault location and configuration verification.

Example

Display information about VLANIF100 of the specified RIPng process.

```
<HUAWEI> display ripng 1 interface vlanif 100
 Vlanif100
    FE80::A0A:200:1
    State : UP, Protocol : RIPNG, MTU : 1440
```


Table 7-16 Description of the display ripng interface command output

Item	Description
FE80::A0A:200:1	Link-local address of the interface.
State	Status of the interface: • UP • DOWN
Protocol	Routing protocol running on the interface.
MTU	Link MTU.

Display detailed information about VLANIF100 of the specified RIPng process.

```
<HUAWEI> display ripng 1 interface vlanif100 verbose
 VLANIF100
    FE80::A0A:200:1
    State : UP, Protocol : RIPNG, MTU : 1440
    Metricin       : 0
    Metricout      : 1
    Input          : Enabled     Output : Enabled
    Default Route  : Disabled
    Poison Reverse : Disabled
    Split Horizon  : Enabled
```

Table 7-17 Description of the display ripng interface verbose command output

Item	Description
FE80::A0A:200:1	Link-local address of the interface.
State	Status of the interface: • UP • DOWN
Protocol	Routing protocol running on the interface.
MTU	Link MTU.
Metricin	Metric added to a received route. To set the metric, run the **ripng metricin** command.
Metricout	Metric added to a sent route. To set the metric, run the **ripng metricout** command.
Input	Indicates whether this interface is enabled to receive RIPng packets. To enable the specified interface to receive RIPng packets, run the **ripng input** command.
Output	Indicates whether this interface is enabled to send RIPng packets. To enable an interface to send RIPng packets, run the **ripng output** command.

Item	Description
Default Route	Whether the interface is enabled to advertise default routes. To generate a default route to the RIPng routing domain, run the **ripng default-route** command.
Poison Reverse	Whether poison reverse is enabled on the interface. To enable poison reverse for RIPng, run the **ripng poison-reverse** command.
Split Horizon	Whether split horizon is enabled on the interface. To enable split horizon for RIPng, run the **ripng split-horizon** command.

7.3.9 display ripng neighbor

Function

The **display ripng neighbor** command displays information about RIPng neighbors.

Format

display ripng *process-id* **neighbor** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Displays neighbor information in the specified RIPng process.	The value is an integer that ranges from 1 to 65535.
verbose	Displays detailed information about RIPng neighbors.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ripng neighbor** command to check information about RIPng neighbors.

Example

Display neighbor information in RIPng process 1.

```
<HUAWEI> display ripng 1 neighbor
 Neighbor : FE80::A0A:201:1 Vlanif100
   Protocol : RIPNG
```

Table 7-18 Description of the display ripng neighbor command output

Item	Description
Neighbor	IPv6 address and interface type of the neighbor interface.
Protocol	Routing protocol.

Display detailed neighbor information in RIPng process 1.

```
<HUAWEI> display ripng 1 neighbor verbose
 Neighbor : FE80::A0A:201:1 VLANIF100
   Protocol : RIPNG
   Number of Active routes      : 1
   Number of routes in garbage  : 0
```

Table 7-19 Description of the display ripng neighbor verbose command output

Item	Description
Number of Active routes	Number of active routes.
Number of routes in garbage	Number of garbage routes.

7.3.10 display ripng neighbor last-nbr-down

Function

The **display ripng neighbor last-nbr-down** command displays information about the last neighbor that goes Down in a RIPng process.

Format

display ripng *process-id* **neighbor last-nbr-down**

Parameters

Parameter	Description	Value
process-id	Displays information about the last neighbor that goes Down in the specified RIPng process.	The value is an integer that ranges from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display ripng neighbor last-nbr-down** command can be used to view information about the last neighbor that goes Down in a RIPng process and the reason why the neighbor goes Down.

Example

Display information about the last neighbor that goes Down in a RIPng process.

```
<HUAWEI> display ripng 1 neighbor last-nbr-down
 Neighbor down index              : 1
 Neighbor Link local Address      : FE80::2E0:B7FF:FE5B:8242
 Interface                        : vlanif 100
 Reason for Neighbor down         : Interface Down
 Time at which neighbor went down : 2011-06-02 12:03:40+03:00 DST
```

Table 7-20 Description of the display ripng neighbor last-nbr-down command output

Item	Description
Neighbor down index	Index of the neighbor that goes Down.
Neighbor link local Address	IPv6 link-local address of the neighbor interface.
Interface	Interface connecting to the neighbor.
Reason for Neighbor down	Reason that a neighbor goes Down • Interface Down • Configuration Change • Time Out (Normal) • Time Out (Authentication Failed)
Time at which neighbor went down	Time the neighbor goes Down.

7.3.11 display ripng route

Function

The **display ripng route** command displays all the RIPng routes that the switch learns from other switches.

Format

display ripng *process-id* **route**

Parameters

Parameter	Description	Value
process-id	Displays the RIPng routes of the specified RIPng process that the switch learns from other switches.	The value is an integer that ranges from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display ripng route** command is useful for debugging and testing protocols.

Example

Display all active and inactive RIPng routes and timers of each route.

```
<HUAWEI> display ripng 100 route
  Route Flags: R - RIPng
               A - Aging, G - Garbage-collect
 ------------------------------------------------------------------------

 Peer 2000::1  on vlanif100
 Dest 2001:DB8:1::1/32,
    via 2000::1, cost  2, tag 0, RA, 6 Sec
 Dest 2001:DB8:1::2/64,
    via 2000::1, cost  2, tag 0, RA, 6 Sec
 Dest 2001:DB8:1::3/64,
    via 2000::1, cost  2, tag 0, RA, 6 Sec
 Dest 2001:DB8:1::4/64,
    via 2000::1, cost  2, tag 0, RA, 6 Sec
```


Table 7-21 Description of the display ripng route command output

Item	Description
Route Flags	Route tag. The first character indicates the type of the route, namely, a RIPng route; the second character indicates the status of the route: • RA: The RIPng route is in aging state. • RG: The RIPng route is in garbage collection state.
Peer	Neighbor connecting to the interface.
Dest	Destination IPv6 address.
via	Next-hop IPv6 address.
cost	Cost of a route.
tag	Route tag.
Sec	Time during which a route remains in the state indicated by the route tag.

7.3.12 display ripng statistics interface

Function

The **display ripng statistics interface** command displays statistics about packets received and sent by interfaces in a RIPng process.

Format

display ripng *process-id* **statistics interface** { **all** | *interface-type interface-number* [**verbose** | **neighbor** *neighbor-ipv6-address*] }

Parameters

Parameter	Description	Value
process-id	Displays statistics on the interface of the specified RIPng process.	The value is an integer that ranges from 1 to 65535.
all	Displays statistics on all interfaces.	-
interface-type interface-number	Displays statistics on the specified interface.	-
verbose	Displays detailed statistics.	-

Parameter	Description	Value
neighbor *neighbor-ipv6-address*	Displays statistics about the specified neighbor.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

This command displays the configuration and operating status of RIPng. The information is helpful in fault location and configuration verification.

Example

Display statistics about VLANIF100 of RIPng process 1.

```
<HUAWEI> display ripng 1 statistics interface vlanif 100
Vlanif100(FE80::2E0:B9FF:FE7F:1A01)
Statistical information          Last min     Last 5 min    Total
----------------------------------------------------------------
Periodic updates sent             5             23            259
Triggered updates sent            5             30            408
Response packets sent             10            34            434
Response packets received         15            38            467
Response packets ignored          0             0             0
Request packets sent              1             3             8
Request packets received          4             20            40
Request packets ignored           0             0             0
Bad packets received              0             0             0
Routes received                   2             10            0
Routes sent                       1             4             4
Bad routes received               0             0             0
Packets send failed               0             0             0
```

Table 7-22 Description of the display ripng statistics interface command output

Item	Description
Statistical information	Packet type.
Last min	Statistics within the last minute.
Last 5 min	Statistics within the last 5 minutes.
Total	Total statistics.
Periodic updates sent	Number of periodically sent Update packets.
Triggered updates sent	Number of triggered sent Update packets.
Response packets sent	Number of sent RIPng Response packets.

Item	Description
Response packets received	Number of received RIPng Response packets.
Response packets ignored	Number of ignored RIPng Response packets.
Request packets sent	Number of sent RIPng Request packets.
Request packets received	Number of received RIPng Request packets.
Request packets ignored	Number of ignored RIPng Request packets.
Bad packets received	Number of received packets that cannot be parsed.
Bad routes received	Number of received routes that cannot be parsed.
Routes received	Number of received routes.
Routes sent	Number of sent routes.
Packets send failed	Number of RIPng packets that fail to be sent.

7.3.13 filter-policy export (RIPng)

Function

The **filter-policy export** command specifies the filtering policy for sent routes.

The **undo filter-policy export** command cancels the configuration.

By default, no filtering policy is configured.

Format

filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } **export** [*protocol* [*process-id*]]

undo filter-policy [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name*] **export** [*protocol* [*process-id*]]

Parameters

Parameter	Description	Value
acl6-number	Specifies the number of a basic IPv6 ACL.	The value is an integer ranging from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of a named IPv6 ACL, which is case sensitive.	The name is a string of 1 to 64 case-sensitive characters without spaces.

Parameter	Description	Value
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of the routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
protocol	Specifies the name of a routing protocol from which routes are imported.	The value can be **bgp**, **direct**, **isis**, **unr**, **ripng**, **static** or **ospfv3**.
process-id	Specifies the process ID of **IS-IS**, **RIPng** or **OSPFv3**.	The value is an integer that ranges from 1 to 65535.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

In some cases, you need to precisely control the advertisement of RIPng routes to meet complicated network requirements. RIPng can use an IPv6 prefix list, route policy and ACL to filter imported routes, allowing only the routes matching the IPv6 prefix list, route policy and ACL to be advertised to RIPng neighbors.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Example

Use IPv6 prefix list **Filter2** to filter the RIPng Update packets that need to be advertised.

```
<HUAWEI> system-view
[HUAWEI] ripng 100
[HUAWEI-ripng-100] filter-policy ipv6-prefix Filter2 export
```


7.3.14 filter-policy import (RIPng)

Function

The **filter-policy import** command specifies the filtering policy for received routes. Only the routes that match the filtering policy can be received.

The **undo filter-policy import** command cancels the configuration.

By default, no filtering policy is configured.

Format

filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } **import**

undo filter-policy [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name*] **import**

Parameters

Parameter	Description	Value
acl6-number	Specifies the number of a basic IPv6 ACL.	The value is an integer ranging from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of a named IPv6 ACL, which is case sensitive.	The name is a string of 1 to 64 case-sensitive characters without spaces.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of the routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

In some cases, you need to precisely control the receiving of RIPng routes to meet complicated network requirements. RIPng can use an IPv6 prefix list, route policy

and ACL to filter received RIPng routes, allowing only the routes matching the IPv6 prefix list, route policy and ACL to be added to RIPng routing tables.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Example

Use IPv6 prefix list **Filter1** to filter received RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ripng 1
[HUAWEI-ripng-1] filter-policy ipv6-prefix Filter1 import
```

7.3.15 import-route (RIPng)

Function

The **import-route** command imports routes.

The **undo import-route** command cancels importing routes.

By default, no route is imported.

Format

import-route { { **ripng** | **isis** | **ospfv3** } [*process-id*] | **bgp** [**permit-ibgp**] | **unr** | **direct** | **static** } [[**cost** *cost* | **inherit-cost**] | **route-policy** *route-policy-name*] *

undo import-route *protocol* [*process-id*]

Parameters

Parameter	Description	Value
protocol	Specifies the routing protocol from which routes are imported.	The value can be **direct**, **static**, **ripng**, **isis**, **unr**, **bgp**, or **ospfv3**.
ripng	Imports RIPng routes.	-
isis	Imports IS-IS routes.	-
ospfv3	Imports OSPFv3 routes.	-
process-id	Specifies the process ID of imported routes.	The value is an integer that ranges from 1 to 65535.
bgp	Imports BGP routes.	-

Parameter	Description	Value
permit-ibgp	Imports IBGP routes in a public network instance. **NOTE** • Import of IBGP routes in RIPng process can lead to routing loops. Administrators should take care of routing loops before configuring **permit-ibgp**.	-
unr	Imports UNR routes.	-
direct	Imports direct routes.	-
static	Imports static routes.	-
cost *cost*	Specifies the cost of imported routes. If no cost is specified, the default cost configured using the **default-cost (RIPng)** command is used.	The value is an integer that ranges from 0 to 15.
inherit-cost	Specifies to retain the original cost of the imported route.	-
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

You can configure a route-policy to allow RIPng to import only the routes that match the rout-policy and set attributes for imported routes.

You can run one of the following commands to set the cost of imported routes. The following commands are listed in descending order of priorities.

- Run the **apply cost** command to set the cost of routes.
- Run the **import-route** (RIPng view) command to set the cost of imported routes.
- Run the **default-cost (RIPng)** command to set the default cost of imported routes.

NOTE

After the **import-route direct** command is executed, routes to the network segment where the IPv6 address of the management interface belongs are also imported in the RIPng routing table. Therefore, use this command with caution.

Example

Import routes from IS-IS process 7 and set the cost of imported routes to 7.

```
<HUAWEI> system-view
[HUAWEI] ripng 1
[HUAWEI-ripng-1] import-route isis 7 cost 7
```

Import the IBGP routes matching route-policy **abc** to RIPng process 1 and set the cost of imported routes to 5.

```
<HUAWEI> system-view
[HUAWEI] ripng 1
[HUAWEI-ripng-1] import-route bgp permit-ibgp cost 5 route-policy abc
```

7.3.16 ipsec sa (RIPng)

Function

The **ipsec sa** command enables Internet Protocol Security (IPSec) authentication in a RIPng process.

The **undo ipsec sa** command disables IPSec authentication in a RIPng process.

By default, IPSec authentication is disabled in a RIPng process.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ipsec sa *sa-name*

undo ipsec sa

Parameters

Parameter	Description	Value
sa-name	Specifies the name of a Security Association (SA)	The value is an existing SA name.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ipsec sa** command enables IPSec authentication in a RIPng process. IPSec authenticates received and sent RIPng packets by using the specified SA (including the security algorithm and key). This improves the security of the RIPng network.

If the **ipsec sa** command is run for a RIPng process, all packets of the process will be authenticated by using the SA specified in the command. This means that the IPSec authentication configuration takes effect on all interfaces in the RIPng process.

If IPSec authentication needs to be enabled only on a certain RIPng interface, run the **ripng ipsec sa** command in the view of the interface.

Prerequisites

An IPSec SA has been configured.

Precaution

The **ripng ipsec sa** command takes precedence over the **ipsec sa** command. If both commands are run in respective views and different SA names are specified, only the configuration of the **ripng ipsec sa** command takes effect.

Example

Enable IPSec in a RIPng process and specify sa1 as the SA name.

```
<HUAWEI> system-view
[HUAWEI] ripng 1
[HUAWEI-ripng-1] ipsec sa sa1
```

7.3.17 maximum load-balancing (RIPng)

Function

The **maximum load-balancing** command configures the maximum number of equal-cost routes for load balancing.

The **undo maximum load-balancing** command restores the default setting.

By default, the maximum number of equal-cost routes is 8.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

maximum load-balancing *number*

undo maximum load-balancing

Parameters

Parameter	Description	Value
number	Specifies the maximum number of equal-cost routes for load balancing.	The value is an integer that ranges from 1 to 8. The default value is 8.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

If there are multiple links on the network, you can set the maximum number of equal-cost routes for load balancing. This prevents some links from being heavily loaded and improves network resource efficiency.

Example

Set the maximum number of equal-cost routes to 4.

```
<HUAWEI> system-view
[HUAWEI] ripng 1
[HUAWEI-ripng-1] maximum load-balancing 4
```


7.3.18 preference (RIPng)

Function

The **preference** command specifies the preference for RIPng routes.

The **undo preference** command restores the default setting.

By default, the preference of RIPng routes is 100.

Format

preference { *preference* | **route-policy** *route-policy-name* } *

undo preference

Parameters

Parameter	Description	Value
preference	Specifies the preference of RIPng routes.	The value is an integer that ranges from 1 to 255.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

A smaller value indicates a higher preference. To make RIPng routes take precedence over other IGP routes, set a smaller preference value for RIPng routes. The preference values also determine which routing protocol's routes are optimal routes in an IPv6 routing table.

Example

Set the preference of RIPng routes to 120.

```
<HUAWEI> system-view
[HUAWEI] ripng 100
[HUAWEI-ripng-100] preference 120
```

Set the preference of the RIPng routes matching route-policy **policy1** to 60.

```
<HUAWEI> system-view
```



```
[HUAWEI] ripng 100
[HUAWEI-ripng-100] preference route-policy policy1 60
```

7.3.19 reset ripng statistics

Function

The **reset ripng statistics** command clears statistics in a specified RIPng process.

Format

reset ripng *process-id* **statistics** [**interface** { **all** | *interface-type interface-number* [**neighbor** *neighbor-ipv6-address*] }]

Parameters

Parameter	Description	Value
process-id	Specifies a RIPng process ID.	The value is an integer that ranges from 1 to 65535.
interface all	Clears statistics on all interfaces.	-
interface *interface-type interface-number*	Clears statistics on the specified interface.	-
neighbor *neighbor-ipv6-address*	Clears statistics in the RIPng process with the specified neighbor.	The value is a 32-digit hexadecimal number, in the common format of X:X:X:X:X:X:X:X.

Views

User view

Default Level

3: Management level

Usage Guidelines

You can use the **reset ripng statistics** command in the user view to clear statistics in a specified RIPng process and re-collect statistics for debugging.

NOTICE

Statistics about a RIPng process cannot be restored after being cleared. So, exercise caution when using this command.

Example

Clear statistics on all the interfaces in RIPng process 100.

```
<HUAWEI> reset ripng 100 statistics interface all
Warning:The RIPNG statistics will be reset. Continue? [Y/N] y
```

7.3.20 ripng

Function

The **ripng** command creates a RIPng process.

The **undo ripng** command deletes a RIPng process.

By default, no RIPng process is created.

Format

ripng [*process-id*] [**vpn-instance** *vpn-instance-name*]

undo ripng *process-id*

Parameters

Parameter	Description	Value
process-id	Specifies a RIPng process ID.	The value is an integer that ranges from 1 to 65535. The default value is 1.
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Before configuring RIPng functions, run the **ripng** command to create a RIPng process.

When a RIPng process is deleted, the **ripng enable** command needs to be run again on interfaces.

Example

Create a specified RIPng process.

```
<HUAWEI> system-view
[HUAWEI] ripng 100
```



```
[HUAWEI-ripng-100]
```

7.3.21 ripng default-route

Function

The **ripng default-route** command generates a default route to the RIPng routing domain.

The **undo ripng default-route** command disables advertising RIPng default routes and forwarding IPv6 default routes.

By default, there is no default route in the RIPng routing domain.

Format

ripng default-route { **only** | **originate** } [**cost** *cost*]

undo ripng default-route

Parameters

Parameter	Description	Value
only	Advertises only IPv6 default routes (::/0) and suppresses the advertisement of other routes.	-
originate	Advertises IPv6 default routes (::/0) without affecting the advertisement of other routes.	-
cost *cost*	Specifies the cost of default routes.	The value is an integer that ranges from 0 to 15. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The generated default RIPng route is forcibly advertised in an Update packet through the specified interface, regardless of whether this route already exists in the IPv6 routing table.

This command can take effect only after IPv6 is enabled for the interface by the **ipv6 enable** command.

Example

Advertise only default routes in Update packets through an interface.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng default-route only
```

Advertise only default routes in Update packets through an interface.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng default-route only
```

Advertise default routes together with other routes in Update packets through an interface.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng default-route originate
```

Advertise default routes together with other routes in Update packets through an interface.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng default-route originate
```

7.3.22 ripng enable

Function

The **ripng enable** command enables RIPng on an interface.

The **undo ripng** command disables RIPng on an interface.

By default, RIPng process is disabled.

Format

ripng *process-id* **enable**

undo ripng

Parameters

Parameter	Description	Value
process-id	Specifies a RIPng process ID.	The value is an integer that ranges from 1 to 65535.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

This command can take effect only after IPv6 is enabled for the RIPng process and interfaces.

Example

Enable RIPng process 100 on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ripng 100
[HUAWEI-ripng-100] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ipv6 address fc00:0:0:1::1/64
[HUAWEI-Vlanif100] ripng 100 enable
```

Enable RIPng process 100 on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ripng 100
[HUAWEI-ripng-100] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ipv6 address fc00:0:0:1::1/64
[HUAWEI-GigabitEthernet0/0/1] ripng 100 enable
```

7.3.23 ripng input

Function

The **ripng input** command enables the specified interface to receive RIPng packets.

The **undo ripng input** command disables the specified interface from receiving RIPng packets.

By default, interfaces can receive RIPng packets.

Format

ripng input

undo ripng input

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

When a device running RIPng is connected to a network running other routing protocols, you can run the **undo ripng input** command on the interface that connects the device to the network to prevent the interface from receiving RIPng packets from the network.

Example

Disable VLANIF100 interface to receive RIPng packets.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] undo ripng input
```

7.3.24 ripng ipsec sa

Function

The **ripng ipsec sa** command enables Internet Protocol Security (IPSec) authentication on a RIPng interface.

The **undo ripng ipsec sa** command disables IPSec authentication on a RIPng interface.

By default, IPSec authentication is disabled on a RIPng interface.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ripng ipsec sa *sa-name*

undo ripng ipsec sa

Parameters

Parameter	Description	Value
sa-name	Specifies the name of a Security Association (SA).	The value is an existing SA name.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ripng ipsec sa** command enables IPSec on a RIPng interface. IPSec authenticates received and sent RIPng packets by using the specified SA (including the security algorithm and key). This improves the security of the RIPng network.

The command can take effect only after IPv6 is enabled for the interface by the **ipv6 enable** command.

If the **ripng ipsec sa** command is run on an interface, all RIPng packets received and sent by the interface will be authenticated by using the SA specified in the command.

Prerequisites

An IPSec SA has been configured.

Precaution

The **ripng ipsec sa** command takes precedence over the **ipsec sa** command. If both commands are run in respective views and different SA names are specified, only the configuration of the **ripng ipsec sa** command takes effect.

Example

Enable IPSec authentication on a RIPng interface and specify sa3 as the SA name.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ripng ipsec sa sa3
```

7.3.25 ripng metricin

Function

The **ripng metricin** command sets the metric that is added to the RIPng route received by an interface.

The **undo ripng metricin** command restores the default metric added to a received RIPng route.

By default, an interface does not add the metric to a received RIPng route.

Format

ripng metricin *value*

undo ripng metricin

Parameters

Parameter	Description	Value
value	Specifies the value of the metric to be added to a received RIPng route.	The value is an integer that ranges from 0 to 15. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

When an interface receives a valid RIPng route, the interface adds a metric to the route before adding the route to the routing table. The metric of this route is changed in the routing table. That is, increasing the metric of an interface also increases the metric of the RIPng route received by the interface.

Before running the **ripng metricin** command on an interface, run the **ipv6 enable** command to enable IPv6 on the interface.

Example

Set the metric that is added to a received RIPng route to 12.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng metricin 12
```

Set the metric that is added to a received RIPng route to 12.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng metricin 12
```


7.3.26 ripng metricout

Function

The **ripng metricout** command sets the metric that is added to the RIPng route sent by an interface.

The **undo ripng metricout** command restores the default setting.

By default, the metric that is added to the RIPng route sent by an interface is 1.

Format

ripng metricout { *value* | { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } *value1* }

undo ripng metricout

Parameters

Parameter	Description	Value
value	Specifies the metric added to a sent route.	The value is an integer that ranges from 1 to 15. The default value is 1.
acl6-number	Specifies the number of a basic IPv6 ACL.	The value is an integer ranging from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of a named IPv6 ACL.	The name is a string of 1 to 64 case-sensitive characters without spaces.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
value1	Specifies the metric added to the route matching the IPv6 prefix list.	The value is an integer that ranges from 2 to 15.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

An interface adds a metric to a route before advertising the route. Therefore, increasing the metric of an interface also increases the metric of the RIPng routes

sent by the interface. The metric of the route in the routing table, however, remains unchanged.

Before running the **ripng metricout** command on an interface, run the **ipv6 enable** command to enable IPv6 on the interface.

Example

Set the metric that is added to a route sent on an interface to 12.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng metricout 12
```

Set the metric that is added to a route sent on an interface to 12.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng metricout 12
```

Set the metric that is added to the RIPng route matching IPv6 prefix list **p1** to 12.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ip ipv6-prefix p1 permit fc00:0:0:1::1 128
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng metricout ipv6-prefix p1 12
```

7.3.27 ripng output

Function

The **ripng output** command enables an interface to send RIPng packets.

The **undo ripng output** command disables an interface from sending RIPng packets.

By default, an interface can send RIPng packets.

Format

ripng output

undo ripng output

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

When a device running RIPng is connected to a network running other routing protocols, you can run the **undo ripng output** command on the interface that connects the device to the network to prevent the interface from sending RIPng packets to the network.

Example

Disable VLANIF100 interface to send RIPng packets.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] undo ripng output
```

7.3.28 ripng pkt-transmit

Function

The **ripng pkt-transmit** command sets the interval for sending Update packets and the number of packets sent each time on a specified interface.

The **undo ripng pkt-transmit** command restores the default value.

By default, the interval for sending Update packets is 200 ms and the number of packets sent each time is 30 on the RIPng interface.

Format

ripng pkt-transmit { **interval** *interval* | **number** *pkt-count* } *

undo ripng pkt-transmit

Parameters

Parameter	Description	Value
interval *interval*	Specifies the interval for sending Update packets.	The value is an integer that ranges from 50 to 500, in milliseconds. The default value is 200.
number *pkt-count*	Specifies the number of packets sent each time.	The value is an integer that ranges from 25 to 100. The default value is 30.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Running the **ripng pkt-transmit** command in the interface view as required can accurately control the interval for sending packets and the number of sent packets. RIPng performance is then improved.

Before running the **ripng pkt-transmit** command on an interface, run the **ipv6 enable** command to enable IPv6 on the interface.

Example

Set the interval for sending Update packets on VLANIF100 to 100 ms and the number of packets sent each time to 50.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng pkt-transmit interval 100 number 50
```

Set the interval for sending Update packets on GE0/0/1 to 100 ms and the number of packets sent each time to 50.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng pkt-transmit interval 100 number 50
```

7.3.29 ripng poison-reverse

Function

The **ripng poison-reverse** command enables poison reverse for RIPng.

The **undo ripng poison-reverse** command disables poison reverse for RIPng.

By default, poison reverse is disabled for RIPng.

Format

ripng poison-reverse

undo ripng poison-reverse

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

After poison reverse is configured, the route learned from an interface is also sent to the neighboring device through the same interface. The metric of this route is set to 16, indicating that the route is unreachable.

When split horizon and poison reverse are configured on the same interface, only poison reverse takes effect.

Before running the **ripng poison-reverse** command on an interface, run the **ipv6 enable** command to enable IPv6 on the interface.

Example

Enable poison reverse for RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng poison-reverse
```

Enable poison reverse for RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng poison-reverse
```

Disable poison reverse for RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] undo ripng poison-reverse
```

7.3.30 ripng split-horizon

Function

The **ripng split-horizon** command enables split horizon for RIPng.

The **undo ripng split-horizon** command disables split horizon for RIPng.

By default, split horizon is enabled except on the NBMA network.

Format

ripng split-horizon

undo ripng split-horizon

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

After split horizon is enabled, a route is blocked if the route is advertised through the interface that learns the route. The split horizon mechanism avoids routing loops between adjacent neighbors.

When poison reverse and split horizon are configured on the same interface, only poison reverse takes effect.

Before running the **ripng split-horizon** command on an interface, run the **ipv6 enable** command to enable IPv6 on the interface.

Example

Enable split horizon for RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ripng split-horizon
```

Enable split horizon for RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ripng split-horizon
```

Disable split horizon for RIPng Update packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] undo ripng split-horizon
```

7.3.31 ripng summary-address

Function

The **ripng summary-address** command configures a RIPng device to advertise summarized IPv6 addresses on an interface.

The **undo ripng summary-address** command deletes this configuration.

By default, a RIPng switch does not advertise summarized IPv6 addresses.

Format

ripng summary-address *ipv6-address prefix-length* [**avoid-feedback**]

undo ripng summary-address *ipv6-address prefix-length*

Parameters

Parameter	Description	Value
ipv6-address	Specifies a summarized IPv6 address.	The value is a 32-digit hexadecimal number, in the X:X:X:X:X:X:X:X format.
prefix-length	Specifies the prefix length of an IPv6 address.	The value is an integer that ranges from 0 to 128.
avoid-feedback	Disables an interface from learning the same summarized route.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

If the prefix and prefix length of a route match the defined IPv6 prefix, the defined IPv6 prefix is advertised instead of the route. As a result, multiple routes are replaced by a single route that has a lower metric.

After the keyword **avoid-feedback** is specified, an interface no longer learns the summarized route with the same IP address as the advertised summarized IP address. This avoids routing loops.

This command can take effect only after IPv6 is enabled for the RIPng process and interfaces.

Example

Configure IPv6 address FC00:200::3EFF:FE11:6770 for VLANIF100, set the prefix length to 64 bits, and summarize the configured IPv6 address as an IPv6 address with prefix FC00:200::/35 in RIPng process 100.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ripng 100
[HUAWEI-ripng-100] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ipv6 address fc00:200::3EFF:FE11:6770/64
```



```
[HUAWEI-Vlanif100] ripng 100 enable
[HUAWEI-Vlanif100] ripng summary-address fc00:200:: 35
```

Configure IPv6 address FC00:200::3EFF:FE11:6770 for GE0/0/1, set the prefix length to 64 bits, and summarize the configured IPv6 address as an IPv6 address with prefix FC00:200::/35 in RIPng process 100.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ripng 100
[HUAWEI-ripng-100] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ipv6 address fc00:200::3EFF:FE11:6770/64
[HUAWEI-GigabitEthernet0/0/1] ripng 100 enable
[HUAWEI-GigabitEthernet0/0/1] ripng summary-address fc00:200:: 35
```

7.3.32 timers ripng

Function

The **timers ripng** command configures RIPng timers.

The **undo timers ripng** command restores the default values of timers.

By default, the interval for sending RIPng Update packets is 30s, the aging time of routes is 180s, and the time for deleting routes from an IPv6 routing table is 120s.

Format

timers ripng *update age garbage-collect*

undo timers ripng

Parameters

Parameter	Description	Value
update	Specifies the interval for sending Update packets. This parameter is the basic timing parameter of routing protocols.	The value is an integer that ranges from 1 to 86400, in seconds. The default value is 30 seconds.
age	Specifies the aging time of routes. The value of *age* must be at least three times the value of *update*. If no Update packet of a route is received within the aging time, the route becomes invalid (unreachable).	The value is an integer that ranges from 1 to 86400, in seconds. The default value is 180 seconds.
garbage-collect	Specifies the time from when a route is found invalid until it is deleted from the routing table.	The value is an integer that ranges from 1 to 86400, in seconds. The default value is 120 seconds.

Views

RIPng view

Default Level

2: Configuration level

Usage Guidelines

The basic timing parameters of RIPng can be adjusted because RIPng adopts a distributed asynchronous routing algorithm. On the network, these parameters on switches need to be consistent with those on access servers.

The command can take effect only after the global IPv6 is enabled.

If the values of the preceding four timers are set improperly, route flapping occurs. The relationship of the values of the four timers is: *update* < *age* and *update* < *garbage-collect*. For example, if the update time is longer than the aging time, switches cannot inform neighbors on time if RIPng routes change during the update time.

Example

Set values for RIPng timers.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ripng 100
[HUAWEI-ripng-100] timers ripng 5 15 30
```

7.4 OSPF Configuration Commands

7.4.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.4.2 abr-summary (OSPF area)

Function

The **abr-summary** command configures route summarization on an area border router (ABR).

The **undo abr-summary** command disables route summarization on an ABR.

By default, route summarization is not configured on ABRs.

Format

abr-summary *ip-address mask* [**cost** { *cost* | **inherit-minimum** } | [**advertise** [**generate-null0-route**] | **not-advertise** | **generate-null0-route** [**advertise**]]]*

undo abr-summary *ip-address mask*

Parameters

Parameter	Description	Value
ip-address	Specifies the IP address of a summarized route.	The value is in dotted decimal notation.
mask	Specifies the mask of the IP address of the summarized route.	The value is in dotted decimal notation.
advertise | **not-advertise**	Indicates whether to advertise the summarized route. By default, the summarized route is advertised.	-
cost *cost*	Specifies the cost of the summarized route. By default, the highest cost of specific routes is used as the cost of the summarized route.	The value is an integer that ranges from 0 to 16777214.
inherit-minimum	Indicates that the smallest cost of specific routes is used as the cost of the summarized route.	-
generate-null0-route	Generates a blackhole route to prevent routing loops.	-

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a large-scale OSPF network, route search speed may decrease due to a large routing table size. Route summarization can be configured to reduce the size of the routing table and simplify management.

Route summarization aggregates multiple routes with the same IP prefix into one. If a link connected to a device within a summarized IP address range alternates between Up and Down states, the link status change is not advertised to the devices outside that IP address range. This prevents route flapping and improves network stability.

When the ABR sends routing information to other areas, it originates Type 3 LSAs for each network segment. If any contiguous segments exist in this area, run the **abr-summary** command to summarize these segments into one. The ABR then sends just one summarized LSA, and no LSAs that belong to the summarized network segment specified by the command. Therefore, the routing table size is reduced, and switch performance is improved.

Prerequisites

The **network** command has been run to specify the segments that need to be summarized before configuring route summarization.

Precautions

- This command applies only to ABRs for intra-area route summarization. The **asbr-summary** command configures AS Boundary Routers (ASBRs) to summarize the routes imported by OSPF.
- Route summarization cannot be configured on ABRs in different areas of the same process.

Example

In OSPF 100 area 1, summarize routes in two network segments, 10.42.10.0 and 10.42.110.0, into one route 10.42.0.0, and advertise the summarized route to other areas.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 1
[HUAWEI-ospf-100-area-0.0.0.1] network 10.42.10.0 0.0.0.255
[HUAWEI-ospf-100-area-0.0.0.1] network 10.42.110.0 0.0.0.255
[HUAWEI-ospf-100-area-0.0.0.1] abr-summary 10.42.0.0 255.255.0.0
```

7.4.3 advertise mpls-lsr-id

Function

The **advertise mpls-lsr-id** command configures OSPF to advertise MPLS LSR IDs to multiple areas as intra-area routes.

The **undo advertise mpls-lsr-id** command cancels the configuration.

By default, OSPF does not advertise MPLS LSR IDs to multiple areas as intra-area routes.

NOTE

Only the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H support this command.

Format

advertise mpls-lsr-id [**cost** *cost*]

undo advertise mpls-lsr-id

Parameters

Parameter	Description	Value
cost *cost*	Specifies the cost of the advertised route.	The value is an integer that ranges from 0 to 65535. By default, it is 0.

Views

OSPF view

Usage Guidelines

Usage Scenario

The prerequisite of a valid tunnel is that an intra-area route to the egress is reachable. When an ABR serves as the egress of tunnels in two areas, OSPF considers that only one tunnel is valid, because there is only one intra-area route to the egress, namely the IP address of the ABR's loopback interface (used as an MPLS LSR ID for tunnel establishment). To the other areas, this route is an inter-area route. To allow the tunnels in both areas to take effect, run the **advertise mpls-lsr-id** command to configure OSPF to advertise the intra-area route to the MPLS LSR ID to all areas connected to this device.

Prerequisites

Before running this command, the **mpls te** command has been run to enable MPLS TE globally.

Configuration Impact

Running this command may have the following impacts:

- If the OSPF area where the local switch resides does not have the interface whose IP address is the MPLS LSR ID, the type of the OSPF routes (with the destination address being the MPLS LSR ID) on the other devices in this OSPF area will be changed from inter-area to intra-area, and the cost will also be changed.
- If an OSPF NSSA has the interface whose IP address is the MPLS LSR ID and the forwarding address of the NSSA LSAs advertised in this area is the MPLS LSR ID, the forwarding address of these NSSA LSAs will be changed and the NSSA LSAs will be re-advertised, thus causing route calculation.

Example

Configure OSPF to advertise MPLS LSR IDs to multiple areas as intra-area routes.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] advertise mpls-lsr-id
```


7.4.4 area (OSPF)

Function

The **area** command creates an OSPF area and displays the OSPF area view.

The **undo area** command deletes a specified area.

By default, the system does not creates any OSPF area.

Format

area *area-id*

undo area *area-id*

Parameters

Parameter	Description	Value
area-id	Specifies an area ID. The area with the *area-id* of 0 is the backbone area.	The value can be a decimal integer or in dotted decimal notation. When the value is an integer, the value ranges from 0 to 4294967295.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The number of devices increases with the expansion of a network. This leads to a large LSDB on every OSPF-enabled device on a large-scale network. Route flapping frequently occurs and as such, a large number of OSPF packets are transmitted on the network. This wastes bandwidth resources.

OSPF resolves this problem by partitioning an AS into different areas. An area is regarded as a logical group, and each group is identified by a unique area ID.

Configuration Impact

After OSPF partitions the AS into different areas, the functions of multiple devices in the same area, such as the timer, filter, and summarization, can be planned and configured uniformly in the area. Therefore, the size of the LSDB is reduced, and network performance is improved.

Prerequisites

An OSPF process has been started using the **ospf** command.

Precautions

- At the border of an area resides a switch instead of a link.
- A network segment or a link belongs to only one area. Specify the area to which each OSPF interface belongs.
- The backbone area is responsible for forwarding inter-area routing information. The routing information between the non-backbone areas must be forwarded through the backbone area.
- All non-backbone areas must maintain connectivity with the backbone area. The backbone area must also maintain connectivity within itself.

Example

Enter the view of an OSPF area.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 0
[HUAWEI-ospf-100-area-0.0.0.0]
```

7.4.5 asbr-summary (OSPF)

Function

The **asbr-summary** command configures an AS Boundary Router (ASBR) to summarize the routes imported by OSPF.

The **undo asbr-summary** command disables an ASBR from summarizing the routes imported by OSPF.

By default, ASBRs do not summarize the routes imported by OSPF.

Format

asbr-summary *ip-address mask* [[**not-advertise** | **generate-null0-route**] | **tag** *tag* | **cost** *cost* | **distribute-delay** *interval*] *

asbr-summary type nssa-trans-type-reference [**cost nssa-trans-cost-reference**]

undo asbr-summary type

undo asbr-summary *ip-address mask*

Parameters

Parameter	Description	Value
ip-address	Specifies the IP address of a summarized route.	The value is in dotted decimal notation.

Parameter	Description	Value
mask	Specifies the mask of the IP address of the summarized route.	The value is in dotted decimal notation.
not-advertise	Indicates that the summarized route is not advertised. If this parameter is not specified, the summarized route is advertised.	-
generate-null0-route	Generates a blackhole route to prevent routing loops.	-
tag *tag*	Specifies the tag of the summarized route.	The value is an integer that ranges from 0 to 4294967295. The default value is 1.
cost *cost*	Specifies the cost of the summarized route. By default, for Type 1 external routes, the cost of the summarized route is the highest cost of specific routes; for Type 2 external routes, the cost of the summarized route equals the highest cost of specific routes plus 1.	The value is an integer that ranges from 0 to 16777214.
distribute-delay *interval*	Specifies the delay in advertising the summarized route.	The value is an integer that ranges from 1 to 65535, in seconds.
type nssa-trans-type-reference	Enables OSPF to consider Type 5 LSAs that have been translated from Type 7 LSAs when it sets types for summary routes on ASBRs. By default, when OSPF sets types for summary routes on ASBRs, OSPF does not consider Type 5 LSAs that have been translated from Type 7 LSAs.	-

Parameter	Description	Value
cost nssa-trans-cost-reference	Enables OSPF to consider Type 5 LSAs that have been translated from Type 7 LSAs when it sets costs for summary routes on ASBRs. If the **asbr-summary type nssa-trans-type-reference cost nssa-trans-cost-reference** command is not run, OSPF does not consider Type 5 LSAs that have been translated from Type 7 LSAs when setting types and costs for summary routes on ASBRs.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a large-scale OSPF network, route search speed may decrease due to a large routing table size. Configure route summarization to reduce the routing table size and simplify management.

Route summarization aggregates multiple routes with the same IP prefix into one. If a link connected to a device within a summarized IP address range alternates between Up and Down states, the link status change is not advertised to the devices outside the IP address range. This prevents route flapping and improves network stability.

Imported routes with the same prefix can be summarized into one and advertised as one route using the **asbr-summary** command. Route summarization reduces routing information and routing table size, improving device performance.

After route summarization is implemented, if the local device:

- If a local router is an ASBR in an NSSA, the local router summarizes all imported Type 5 LSAs within the summary address range.
- If the local router is an ASBR in an NSSA, the local router summarizes all imported Type 7 LSAs within the summary address range.
- If the local router functions as both an ASBR and an ABR in an NSSA, the local router summarizes all imported Type 7 LSAs within the address range. It also summarizes the Type 5 LSAs that are translated from Type 7 LSAs.

Precautions

When a large number of routes are summarized, specify the **distribute-delay** parameter to set a delay in advertising the summarized routes. This ensures that

the advertised summarized routes contain more valid routes and avoids network flapping and incorrect routing information.

Example

Configure route summarization for the imported routes.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] asbr-summary 10.2.0.0 255.255.0.0 not-advertise tag 2 cost 100
```

Cancel route summarization for the imported routes.

```
[HUAWEI-ospf-100] undo asbr-summary 10.2.0.0 255.255.0.0
```

7.4.6 authentication-mode (OSPF area)

Function

The **authentication-mode** command sets an authentication mode and a password for an OSPF area.

The **undo authentication-mode** command cancels the authentication mode configured for an OSPF area.

By default, no authentication mode is configured.

Format

authentication-mode simple [**plain** *plain-text* | [**cipher**] *cipher-text*]

authentication-mode { **md5** | **hmac-md5** | **hmac-sha256** } [*key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* }]

authentication-mode keychain *keychain-name*

undo authentication-mode

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
simple	Sets simple authentication. In simple authentication, the password type is **cipher** by default. **NOTICE** Simple authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
plain	Sets a plaintext password. If this parameter is specified, you can only enter a plaintext password, which is then displayed in plain text when the configuration file is viewed. **NOTICE** If **plain** is specified, the password is saved in the configuration file in plain text. This carries security risks. You are advised to specify **cipher** to save the password in cipher text.	-
plain-text	Sets a plaintext password.	The value is a string of case-sensitive characters that can be letters or digits without spaces. In **simple** authentication, the value is a string of 1 to 8 characters. In **md5**, **hmac-md5** or **hmac-sha256** authentication, the value is a string of 1 to 255 characters.
cipher	Sets a ciphertext password. Either a plaintext or ciphertext password can be entered, and cipher text is displayed when the configuration file is viewed.	When **cipher** is configured, the password can only be entered in cipher text. Then, the password is displayed in cipher text in configuration files. MD5 authentication, HMAC-SHA256 authentication or HMAC-MD5 authentication defaults to use the password in cipher text.

Parameter	Description	Value
cipher-text	Specifies the ciphertext password.	The value is a string of case-sensitive characters that can be letters or digits without spaces. In **simple** authentication, the value is a string of 1 to 8 characters in plain text, or a string of 24 or 32 or 48 characters in cipher text. In **md5**, **hmac-sha256** or **hmac-md5** authentication, the value is a string of 1 to 255 characters in plain text, or a string of 20 to 392 characters in cipher text.
md5	Indicates MD5 authentication using the ciphertext password. **NOTICE** MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-md5	Indicates HMAC MD5 authentication using the ciphertext password. **NOTICE** HMAC-MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-sha256	Indicates HMAC-SHA256 authentication.	-
key-id	Specifies authentication key ID of the interface's cipher authentication. The key ID must be consistent with that of the peer.	The value is an integer that ranges from 1 to 255.

Parameter	Description	Value
keychain	Indicates keychain authentication. **NOTE** Before configuring this parameter, run the **keychain** command to create a keychain. Then, run the **key-id**, **key-string**, and **algorithm** commands to configure a key ID, a password, and an authentication algorithm for this keychain. Otherwise, OSPF authentication will fail. Currently, only the **HMAC-MD5**, **SM3**, and **HMAC-SHA256** algorithms can be used in OSPF.	-
keychain-name	Specifies the keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

OSPF authentication can be configured to improve network security to meet high security demands. When area authentication is used, interfaces on all devices in an area must have the same area authentication mode and the password.

Precautions

The priority of area authentication is lower than the priority of interface identification. The **ospf authentication-mode** command can be used to change the priority of interface authentication.

Example

Configure HMAC-SHA256 authentication for OSPF area 0.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
```



```
[HUAWEI-ospf-100] area 0
[HUAWEI-ospf-100-area-0.0.0.0] authentication-mode hmac-sha256
```

7.4.7 bandwidth-config enable

Function

The **bandwidth-config enable** command enables a device to calculate the cost of an OSPF interface based on the configured interface bandwidth.

The **undo bandwidth-config enable** command disables a device from calculating the cost of an OSPF interface based on the configured interface bandwidth.

By default, a device is not enabled to calculate the cost of an OSPF interface based on the configured interface bandwidth.

Format

bandwidth-config enable

undo bandwidth-config enable

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

External factors may affect the physical bandwidth of a link and change the physical bandwidth of an interface, thereby deteriorating network performance. To enable a device to adjust and optimize the route selection rules, run the **bandwidth** *bandwidth* command in the OSPF interface to adjust the interface bandwidth and run the **bandwidth-config enable** command to enable the device to calculate the cost of an OSPF interface based on the configured interface bandwidth. The calculation formula is as follows: Cost of the interface = Bandwidth reference value/Interface bandwidth.

Precautions

- If the **bandwidth** command is not configured, the OSPF interface cost is calculated based on the physical bandwidth of the interface.
- If the **bandwidth-config enable** command is not configured, the OSPF interface cost is calculated based on the physical bandwidth of the interface.

Example

Enable a device to calculate the cost of an OSPF interface based on the configured interface bandwidth.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] bandwidth-config enable
```

7.4.8 bandwidth-reference (OSPF)

Function

The **bandwidth-reference** command sets a bandwidth reference value that is used to calculate interface costs.

The **undo bandwidth-reference** command restores the default bandwidth reference value.

The default bandwidth reference value is 100 Mbit/s.

Format

bandwidth-reference *value*

undo bandwidth-reference

Parameters

Parameter	Description	Value
value	Specifies a bandwidth reference value for link cost calculation.	The value is an integer ranging from 1 to 2147483648, in Mbit/s. The default value is 100 Mbit/s.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The calculation formula is as follows: Interface cost = Bandwidth reference value/Interface bandwidth. The integer of the calculated result is the interface cost. If the calculated result is smaller than 1, the interface cost value is 1. If the **bandwidth-reference** command is executed to configure a new bandwidth reference value, the interface cost will be changed. As a result, OSPF will re-select routes.

The default bandwidth reference value is 100 Mbit/s. The interface cost value is 100000000 divided by the interface bandwidth value.

Precautions

- After the **bandwidth-reference** command is configured in a process view, bandwidth reference values of all interfaces in the process are changed to the specified value.
- If the **bandwidth-reference** command is run on an Eth-Trunk interface, the bandwidth of the Eth-Trunk interface is equal to the total bandwidth of all its member interfaces.

Example

Set the bandwidth reference value to 1000 Mbit/s.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] bandwidth-reference 1000
```

7.4.9 bfd all-interfaces (OSPF)

Function

The **bfd all-interfaces** command enables bidirectional forwarding detection (BFD) in an OSPF process and sets the parameter values of a BFD session.

The **undo bfd all-interfaces** command disables BFD in an OSPF process and restores the default parameter values of a BFD session.

By default, BFD is disabled in an OSPF process.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

bfd all-interfaces enable

undo bfd all-interfaces enable

bfd all-interfaces { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* | **frr-binding** } *

undo bfd all-interfaces { **min-rx-interval** | **min-tx-interval** | **detect-multiplier**| **frr-binding** } *

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **frr-binding** parameter.

Parameters

Parameter	Description	Value
min-rx-interval *receive-interval*	Indicates the minimum interval at which BFD packets are received from the remote end.	The value is an integer that ranges from 100 to 1000, in milliseconds. ● After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. ● After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *transmit-interval*	Indicates the minimum interval at which BFD packets are sent to the remote end.	The value is an integer that ranges from 100 to 1000, in milliseconds. ● After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. ● After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier-value*	Indicates the local detection multiplier.	The value is an integer ranging from 3 to 50. By default, it is 3.
frr-binding	Binds the BFD status to the link status of an interface.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a device communicates with its neighbors, BFD can fast detect faults to minimize the fault impact on services.

You can bind a BFD session to an interface or an OSPF process. If a BFD session is bound to both an interface and an OSPF process, the BFD session bound to an interface takes precedence over that bound to an OSPF process.

Precautions

- *receive-interval* is negotiated based on the **min-rx-interval** value on the local end and the **min-tx-interval** value on the remote end. The local end uses the larger one between the **min-rx-interval** value on the local end and the **min-tx-interval** value on the remote end as the remote **min-tx-interval**. If the local end does not receive any BFD packets within the interval of *receive-interval* × *multiplier-value* (local detection multiplier), it declares that the remote end is unreachable.
- An OSPF device sets up BFD sessions with only the neighbors in Exstart state. Two ends can set up a BFD session only when the **bfd** command is run on both ends to configure BFD globally and the **bfd all-interfaces enable** command is run.
- The **bfd all-interfaces** command and the **ospf bfd block** command are mutually exclusive.

Example

Configure BFD in an OSPF process and set the minimum interval for sending BFD packets to 400 ms.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] bfd all-interfaces enable
[HUAWEI-ospf-1] bfd all-interfaces min-tx-interval 400
```

7.4.10 default (OSPF)

Function

The **default** command configures default parameters for OSPF-imported external routes. The parameters include the cost, type (Type 1 or Type 2), tag, and number of imported routes.

The **undo default** command restores the default setting.

By default, the default cost of the imported external routes is 1; the upper limit of the imported external routes is 2147483647; the type of the imported external routes is Type 2; the default tag value is 1.

Format

default { **cost** { *cost-value* | **inherit-metric** } | **limit** *limit* | **tag** *tag* | **type** *type* } *

undo default { **cost** | **limit** | **tag** | **type** } *

Parameters

Parameter	Description	Value
cost *cost-value*	Specifies the default cost of the external routes imported by OSPF.	The value is an integer ranging from 0 to 16777214. By default, it is 1.
inherit-metric	Indicates that the cost of the imported route is the cost that the route itself carries. If no cost is specified, the default cost set using the **default** command is used.	-
limit *limit*	Specifies the default upper limit of the external routes to be imported within a given period.	The value is an integer ranging from 1 to 2147483647.
tag *tag*	Specifies the tag of the external routes.	The value is an integer ranging from 0 to 4294967295. By default, it is 1.
type *type*	Specifies the type of the external routes.	The value is an integer ranging from 1 to 2. By default, it is 2. • 1: Type 1 external route • 2: Type 2 external route

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The imported external routes carry various parameters that can change the priorities and next hops of those routes in the OSPF routing table.

By setting default parameters for OSPF-imported external routes, you can change OSPF routing policies.

The route tag is used to identify protocol-related information. For example, it can be used to differentiate AS numbers when OSPF receives BGP routes. It also allows you to apply OSPF routing policies to tagged routes.

Follow-up Procedure

The priority of the **default (OSPF)** command is the lowest. Thus, ensure that no other commands are configured when configuring this command. Otherwise, this command cannot take effect.

Precautions

You can run any of the following commands to set a cost for an imported route. The following commands are listed in the descending order of priority.

- **apply cost**
- **import-route (OSPF)**
- **default (OSPF)**

Example

Set the default values for the cost, type, and tag of imported routes.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] default cost 10 tag 100 type 2
```

7.4.11 default-cost (OSPF Area)

Function

The **default-cost** command sets a cost for the Type3 default route that is transmitted to a stub or NSSA by OSPF.

The **undo default-cost** command restores the default setting.

By default, the cost of the Type3 default route transmitted to a stub or NSSA is 1.

Format

default-cost *cost*

undo default-cost

Parameters

Parameter	Description	Value
cost	Specifies the cost of the Type3 default route transmitted to a stub or NSSA by OSPF.	The value is an integer ranging from 0 to 16777214.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Configuring a cost for a default route can change OSPF route selection, improving networking flexibility.

This command is applicable to area border routers (ABRs) connected to stubs or NSSAs.

Prerequisites

A default route exists in the local routing table.

Example

Set Area 1 as a stub area, and set the cost of the Type 3 default route transmitted to this stub area to 20.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 1
[HUAWEI-ospf-100-area-0.0.0.1] stub
[HUAWEI-ospf-100-area-0.0.0.1] default-cost 20
```

7.4.12 default-route-advertise (OSPF)

Function

The **default-route-advertise** command advertises default routes to a common OSPF area.

The **undo default-route-advertise** command disables advertisement of default routes to a common OSPF area.

By default, OSPF devices in a common OSPF area do not generate default routes.

Format

default-route-advertise [[**always** | **permit-calculate-other**] | **cost** *cost* | **type** *type* | **route-policy** *route-policy-name* [**match-any**]] *

default-route-advertise summary cost *cost*

undo default-route-advertise

Parameters

Parameter	Description	Value
always	Generates and advertises an LSA that describes the default route, regardless of whether there are active default routes of other OSPF processes in the routing table of the host. ● If **always** is configured, the switch does not calculate the default routes from other switches. ● If **always** is not configured, an LSA for advertising a default route can be generated only when there are active default routes of other OSPF processes in the routing table of the local device.	-
permit-calculate-other	Generates and advertises an ASE LSA that describes the default route only when there are active default routes of other OSPF processes in the routing table of the local device. The device still calculates the default routes from other devices. **NOTE** If neither **always** nor **permit-calculate-other** is configured, ● When there are active default routes of other OSPF processes in the routing table of the local device, the device does not calculate the default routes from other devices. ● When there are no active default routes of other OSPF processes in the routing table of the local device, the device still calculates the default routes from other devices.	-
cost *cost*	Specifies the cost of the ASE LSA.	The value is an integer that ranges from 0 to 16777214. The default value is 1.

Parameter	Description	Value
type *type*	Specifies the type of the external routes.	The value is 1 or 2. The default value is 2. • 1: Type 1 external route • 2: Type 2 external route
route-policy *route-policy-name*	Specifies the name of a routing policy. The device advertises default routes according to the parameters of the configured routing policy when there are matched default routes of other OSPF processes in the routing table of the device.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
summary	Advertises the Type 3 summary LSA of the specified default route. Before specifying this parameter, ensure that a VPN is enabled. Otherwise, routes cannot be advertised.	-
match-any	Indicates that a device matches the routing entry in the routing table against a routing policy and then advertises the default route according to the parameters set through the routing policy.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **import-route (OSPF)** command cannot be used to import a default route from another AS. Running the **default-route-advertise** command on an ASBR can advertise a non-OSPF default route in a common OSPF area.

If the ASBR has a default route, the **default-route-advertise** command enables the ASBR to advertise the default route 0.0.0.0 to the OSPF area.

If the ASBR has no default route, the **default-route-advertise always** command or the **default-route-advertise** command can be used:

- With **always** configured: The ASBR can advertise the default route 0.0.0.0 even if there is no default route. This allows the default route to remain in the routing table and prevents the ASBR from using a default route sent by another device.
- Without **always** configured: The ASBR generates an LSA. It describes a default route only when the local routing table contains an activated non-OSPF default route (except BGP route).

  If the local routing table contains an activated default BGP route, the default route can be injected to the OSPF routing table based on the following situations:

  – Whether the **ospf** *process-id* **vpn-instance** *vpn-instance-name* command is run together with the **vpn-instance-capability simple** command:

    ▪ If the two commands are executed together, to inject an activated default EBGP route into the OSPF routing table, run the **default-route-advertise** command. To inject an activated default IBGP route to the OSPF routing table, run the **import-route bgp permit-ibgp** command before you run the **default-route-advertise** command.

    ▪ If only the **ospf** *process-id* **vpn-instance** *vpn-instance-name* command is run, to inject an activated default EBGP or IBGP route into the OSPF routing table, run the **default-route-advertise** command.

  – If the **ospf** *process-id* **vpn-instance** *vpn-instance-name* command is not run, to inject an activated default EBGP route into the OSPF routing table, run the **default-route-advertise** command. To inject an activated default IBGP route to the OSPF routing table, run the **import-route bgp permit-ibgp** command before you run the **default-route-advertise** command.

NOTICE

Injecting an IBGP route into the OSPF routing table may cause a routing loop. Exercise caution when you perform this step.

If a routing policy is configured with **match-any**, and multiple routes match the policy, a default LSA will be generated based on the optimal route. The principles for optimal route selection are as follows:

1. A route configured with **type** takes precedence over that not configured with **type**. A route configured with a smaller **type** value takes precedence over that configured with a larger **type** value.
2. A route configured with **cost** takes precedence over that not configured with **cost**. A route configured with a smaller **cost** value takes precedence over that configured with a larger **cost** value.
3. A route configured with **tag** takes precedence over that not configured with **tag**. A route configured with a smaller **tag** value takes precedence over that configured with a larger **tag** value.

Prerequisites

Before advertising a default route, OSPF compares the priorities of default routes in an OSPF area and then advertises a default route with the highest priority. If a static default route is configured on an OSPF device, check the priority of the static default route. The priority must be lower than that of the default route to be advertised by OSPF. This ensures that the default route advertised by OSPF will be added to the routing table of the OSPF device.

Configuration Impact

After the **default-route-advertise** command is configured on the ASBR, the ASBR will generate a Type 5 ASE LSA with a link state ID of 0.0.0.0 and mask of 0.0.0.0. In addition, it will advertise the ASE LSA in an entire OSPF area.

If a routing policy is configured, default routes are advertised based on the following principles:

- If a default route matches the routing policy, a default route is generated on an OSPF device based on the parameters configured in the routing policy. Parameters such as **cost**, **tag**, and **type** can be configured in the routing policy.
- If the default route does not match the routing policy, and **always** is configured, the default route is still advertised.
  - If **always** is not configured, the OSPF device will not advertise the default route.
  - If **always** is configured, OSPF devices will advertise the default route. In addition, only **always** configured in the **default-route-advertise** command takes effect on advertisement of default routes.

Precautions

In different OSPF areas, OSPF advertises default routes using different modes. This **default-route-advertise** command can be used to advertise default routes to a common OSPF area. In a stub, totally stub, or totally NSSA area, default routes are advertised automatically. In an NSSA, the **nssa default-route-advertise** command is used to advertise default routes.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the device advertises the default route as long as a default route that is not generated by the current OSPF process exists in the local routing table.

If the **default-route-advertise command** is used in multiple OSPF processes to generate default routes, routing loops tend to occur. To prevent this issue, you are advised to configure route-policies to filter routes so that default routes will not be learned from these OSPF processes.

Example

Configure an ASBR to advertise the ASE LSA of the default route to common OSPF areas when it has no default route.

```
<HUAWEI> system-view
```



```
[HUAWEI] ospf 1
[HUAWEI-ospf-1] default-route-advertise always
```

7.4.13 description (OSPF)

Function

The **description** command configures a description for an OSPF process.

The **undo description** command deletes the description.

By default, there is no description for any OSPF process.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Specifies the description of an OSPF process.	The value is a string of 1 to 80 case-sensitive characters, spaces supported.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An OSPF process description helps identify special processes, facilitating network maintenance.

Configuration Impact

If the **description** command is run multiple times, the latest configuration overrides the previous one.

Example

Configure a description for an OSPF process.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] description this process contains 3 areas
```


7.4.14 description (OSPF Area)

Function

The **description** command configures a description for an OSPF area.

The **undo description** command deletes the description.

By default, there is no description for any OSPF area.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Specifies the description of an OSPF area.	The value is a string of 1 to 80 case-sensitive characters, spaces supported.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

An OSPF area description helps identify special areas, facilitating network maintenance.

Example

Configure a description for OSPF area 1.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 1
[HUAWEI-ospf-100-area-0.0.0.1] description this is a stub area
```

7.4.15 display default-parameter ospf

Function

The **display default-parameter ospf** command displays the default OSPF configuration.

Format

display default-parameter ospf

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display default-parameter ospf** command to check the default OSPF configuration during OSPF initialization. After the OSPF configuration is modified, this command still displays the default OSPF configuration during OSPF initialization.

Example

Display the default OSPF configuration.

```
<HUAWEI> display default-parameter ospf
 Process View:
 ------------------------------------------------------------
  Default Metric                          : 1
  Default Tag                             : 1
  Default Type                            : 2
  SPF Intelligent-timer Max-interval(msec)     : 10000
  SPF Intelligent-timer Start-interval(msec)   : 500
  SPF Intelligent-timer Hold-interval(msec)    : 1000
  Lsa Maxage (sec)                        : 3600
  Lsa Refresh Time(sec)                   : 1800
  Lsa Maxagediff Interval (sec)           : 900
  Minimum Lsa Arrival Interval(sec)       : 1
  Minimum Lsa Originate Interval(sec)     : 5
  Sham Link Cost                          : 1
  VPN Domain ID                           : 0
  VPN Router Tag                          : 0
  Route Preference for Internal Routes    : 10
  Route Preference for External Routes    : 150
 ------------------------------------------------------------

 Area View:
 ------------------------------------------------------------
  Default Stub Cost                       : 1
 ------------------------------------------------------------

 Interface View:
 ------------------------------------------------------------
  P2P&Broadcast Hello Interval(sec)       : 10
  P2MP&NBMA Hello Interval(sec)           : 30
  P2P&Broadcast Dead Interval(sec)        : 40
  P2MP&NBMA Dead Interval(sec)            : 120
  Poll Interval(sec)                      : 120
  Router DR Priority                      : 1
  Retransmit Interval(sec)                : 5
```



```
 Transmit Delay(sec)                 : 1
-----------------------------------------------------------
```

Table 7-23 Description of the **display default-parameter ospf** command output

Item	Description
Process View	Process view.
Default Metric	Default metric of the imported external route.
Default Tag	Default tag value of the imported external route.
Default Type	Default type of the imported external route.
SPF Intelligent-timer Max-interval(msec)	Default maximum interval of SPF calculation.
SPF Intelligent-timer Start-interval(msec)	Default start interval of SPF calculation.
SPF Intelligent-timer Hold-interval(msec)	Default hold interval of SPF calculation.
Lsa Maxage(sec)	Default maximum age of the LSA.
Lsa Refresh Time(sec)	Default maximum interval for generating an LSA. If the LS age of the LSAs generated by the device reaches the LSA Refresh Time, a new instance must be generated for the LSAs.
Lsa Maxagediff Interval(sec)	Default value difference in the MaxAge fields of LSAs. If the value difference in the MaxAge fields of two LSAs is greater than MaxAgeDiff Interval, the two LSAs are considered to belong to different instances of the same LSA.
Minimum Lsa Arrival Interval(sec)	Default minimum interval for receiving the same LSA.
Minimum Lsa Originate Interval(sec)	Default minimum interval for sending the same LSA.
Sham Link Cost	Default cost of the sham link.
VPN Domain ID	Default domain ID of the VPN.
VPN Router Tag	Default router tag of the VPN.
Route Preference for Internal Routes	Default preference of the internal route.
Route Preference for External Routes	Default preference of the external route.
Area View	Area view.
Default Stub Cost	Default cost of a route in the stub area.

Item	Description
Interface View	Interface view.
P2P&Broadcast Hello Interval(sec)	Default interval for sending Hello packets on a P2P or broadcast network.
P2MP&NBMA Hello Interval(sec)	Default interval for sending Hello packets on a P2MP or NBMA network.
P2P&Broadcast Dead Interval(sec)	Default interval for declaring a neighbor to be Down after no Hello packets are received on a P2P or broadcast network.
P2MP&NBMA Dead Interval(sec)	Default interval for declaring a neighbor to be Down after no Hello packets are received on a P2MP or NBMA network.
Poll Interval(sec)	Default interval for the local device to send Hello packets to a neighbor in the Down state on the NBMA network. The value of Poll Interval is greater than the value of Hello Interval.
Router DR Priority	Default priority of the DR.
Retransmit Interval(sec)	Default interval for retransmitting packets.
Transmit Delay(sec)	Default estimated time for transmitting an LSU packet over this interface. LSAs in the LSU packet must have their age incremented by this amount before transmission.

7.4.16 display gtsm statistics

Function

The **display gtsm statistics** command displays GTSM statistics on a device.

Format

display gtsm statistics all

Parameters

Parameter	Description	Value
all	Displays GTSM statistics on a device.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display gtsm statistics** command to view GTSM statistics, including the total number of BGP, BGPv6, OSPF, LDP, OSPFv3, and RIP packets, the number of packets that have passed, and the number of discarded packets on a device.

Example

Display GTSM statistics on the device.

```
<HUAWEI> display gtsm statistics all
GTSM Statistics Table
------------------------------------------------------------------
SlotId  Protocol  Total Counters  Drop Counters  Pass Counters
------------------------------------------------------------------
 0      BGP       0               0              0
 0      BGPv6     0               0              0
 0      OSPF      0               0              0
 0      LDP       0               0              0
 0      OSPFv3    0               0              0
 0      RIP       0               0              0
------------------------------------------------------------------
```

Table 7-24 Description of the **display gtsm statistics** command output

Item	Description
SlotId	Slot ID.
Protocol	Protocol type: • Software-based forwarding: protocol differentiated, displaying BGP, BGPv6, OSPF, LDP, OSPFv3, or RIP • Hardware-based forwarding: protocol undifferentiated, displaying -----
Total Counters	Total number of packets.
Drop Counters	Total number of dropped packets.
Pass Counters	Total number of packets that have passed.

7.4.17 display ospf abr-asbr

Function

The **display ospf abr-asbr** command displays information about the ABRs and ASBRs of OSPF.

Format

display ospf [*process-id*] **abr-asbr** [*router-id*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
router-id	Specifies the router ID of an ABR or ASBR.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

An area border router (ABR) can belong to two or more areas. One of the areas must be a backbone area. An ABR is used to connect the backbone area and non-backbone areas. The connection with the backbone area can be physically or logically.

An AS boundary router (ASBR) exchanges routing information with other ASs. An ASBR may not reside at the boundary of an AS. It can be an internal device or an ABR. If an OSPF device imports external routes, the device is an ASBR.

This command can view information about the ABRs and ASBRs of OSPF.

Example

Display information about the ABRs and ASBRs of OSPF.

```
<HUAWEI> display ospf abr-asbr

        OSPF Process 1 with Router ID 1.1.1.1
              Routing Table to ABR and ASBR

 RtType      Destination      Area      Cost  Nexthop         Type
 Intra-area  10.10.10.11      0.0.0.0   1     10.2.0.3        ABR
```

Table 7-25 Description of the **display ospf abr-asbr** command output

Item	Description
RtType	Intra-area or inter-area router.
Destination	Router ID of the ABR or ASBR.
Area	Area ID.

Item	Description
Cost	Cost of the route from the local device to the ABR or ASBR.
NextHop	Next hop address through which packets are transmitted to the ABR or ASBR.
Type	ABR or ASBR.

7.4.18 display ospf asbr-summary

Function

The **display ospf asbr-summary** command displays information about OSPF route summarization.

Format

display ospf [*process-id*] **asbr-summary** [*ip-address mask*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
ip-address	Specifies the summary IP address.	The value is in dotted decimal notation.
mask	Specifies the mask of the summary IP address. If no IP address or mask is specified, summarization information of all the imported routes is displayed.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After an AS is divided into areas, configuring route aggregation can reduce routing information transmitted between areas. This reduces the size of the routing table and improves route performance. After the **asbr-summary (OSPF)** command is

run to configure OSPF route summarization on an ASBR, you can run the **display ospf asbr-summary** command to view information about OSPF route summarization on the ASBR.

Example

Display summarization information about all the imported OSPF routes.

```
<HUAWEI> display ospf asbr-summary
        OSPF Process 1 with Router ID 192.168.1.2
                Summary Addresses
 Total summary address count: 1
                Summary Address
 net         : 10.0.0.0
 mask        : 255.0.0.0
 tag         : 10
 status      : Advertise
 Cost        : 0 (Not Configured)
 delay       : 30 (Configured)
 The Count of Route is : 2
 Destination     Net Mask        Proto     Process  Type     Metric
 10.1.0.0        255.255.0.0     Static    1        2        10
 10.2.0.0        255.255.0.0     Static    1        2        10
```

Table 7-26 Description of the **display ospf asbr-summary** command output

Item	Description
Total summary address count	Number of routes that are being summarized through the **asbr-summary** command.
net	Network address of the summary route.
mask	Network mask of the summary route.
tag	Tag of the summary route.
status	Advertisement status of the summary route: • Advertise: indicates that the summary route is advertised. • DoNotAdvertise: indicates that the summary route is not advertised.
Cost	Cost of the summarized route.
delay	Delay for advertising the summary route.
The Count of Route is	Number of routes that are being summarized.
Destination	Destination address of the routes that are being summarized.
Net Mask	Mask of the routes that are being summarized.
Proto	Protocol of the routes that are being summarized.
Process	Process ID.

Item	Description
Type	Type of the imported AS external route, which can be Type 1 or Type 2.
Metric	Metric of the routes that are being summarized.

7.4.19 display ospf bfd session

Function

The **display ospf bfd session** command displays information about the BFD-enabled neighbor.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

display ospf [*process-id*] **bfd session** *interface-type interface-number* [*router-id*]

display ospf [*process-id*] **bfd session** { *router-id* | **all** }

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies the type and number of an interface.	-

Parameter	Description	Value
router-id	Specifies the router ID of the neighbor.	The value is in dotted decimal notation.
all	Indicates all the OSPF-enabled interfaces in the OSPF process.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

A BFD session can be associated with OSPF to fast detect a link fault and notifies OSPF of the fault. This speeds up OSPF's response to the change of the network topology.

The **display ospf bfd session** command displays information about the BFD-enabled neighbor.

Example

Display information about the BFD-enabled neighbor.

```
<HUAWEI> display ospf bfd session all
         OSPF Process 1 with Router ID 10.3.3.3
 Area 0.0.0.0 interface 192.168.1.1(Vlanif100)'s BFD Sessions
NeighborId:10.2.2.2         AreaId:0.0.0.0          Interface:Vlanif100
BFDState:up                 rx    :1000             tx       :1000
Multiplier:3                BFD Local Dis:8198      LocalIpAdd:10.1.1.1
RemoteIpAdd:10.1.1.2        Diagnostic Info:No diagnostic information
```

Table 7-27 Description of the **display ospf bfd session all** command output

Item	Description
NeighborId	Router ID of the neighbor.
AreaId	Area ID.
Interface	Interface through which the local device establishes a BFD session with the neighbor.
BFDState	BFD status: • up • down • unknown

Item	Description
rx	Negotiated minimum interval for receiving BFD packets.
tx	Negotiated minimum interval for sending BFD packets.
Multiplier	Remote detection multiplier.
BFD Local Dis	Local discriminator dynamically assigned by BFD.
LocalIpAdd	Local IP address.
RemoteIpAdd	Remote IP address.
Diagnostic Info	Diagnostic information: ● Init: indicates that the BFD session is in the Initiate state. ● Admin down: indicates that the **shutdown** command is run on the local BFD session. ● BFD global disable: indicates that BFD is not enabled globally. ● BFD session number exceed: indicates that the number of BFD sessions exceeds the limit. ● Detect down: indicates that the local link becomes Down. ● Receive admin down: indicates that the remote link becomes Down. ● BFD is in rearranging: indicates that the board of the BFD session changes, and the data of the BFD session is transferred to another board. ● No diagnostic information: indicates that no diagnostic information exists. ● No BFD packets were received: indicates that detection expires because no BFD packets are received within a specified period. ● administrator down event received: indicates that the device receives an event that the BFD session is set Down by the network administrator.

7.4.20 display ospf brief

Function

The **display ospf brief** command displays OSPF brief information.

Format

display ospf [*process-id*] **brief**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If no OSPF process ID is specified, brief information about all the OSPF processes is displayed.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In addition to OSPF brief information, running this command also displays the following information:

- Configured Router ID
- Configured VPN domain ID

When locating OSPF faults, you can run the **display ospf brief** command to obtain OSPF brief information. You can then analyze OSPF faults according to the brief information.

Example

Display OSPF brief information.

```
<HUAWEI> display ospf brief
        OSPF Process 1 with Router ID 10.5.5.5
              OSPF Protocol Information

RouterID: 10.5.5.5        Border Router:  AREA
Multi-VPN-Instance is not enabled
Opaque Capable
Global DS-TE Mode: Non-Standard IETF Mode
Graceful-restart capability: disabled
Helper support capability  : enabled
      filter capability   : disabled
```



```
      policy capability   : strict lsa check, planned and un-planned
Applications Supported: MPLS Traffic-Engineering
Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
Default ASE parameters: Metric: 1 Tag: 1 Type: 2
Route Preference: 10
ASE Route Preference: 150
SPF Computation Count: 99
RFC 1583 Compatible
Retransmission limitation is disabled
Import routes limitation is enabled, the maximal limitation value: 4294967295
Area Count: 3   Nssa Area Count: 0
ExChange/Loading Neighbors: 0
Process total up interface count: 1
Process valid up interface count: 0
Flush protect mode: false

Area: 0.0.0.0          (MPLS TE not enabled)
Authtype: None   Area flag: Normal
SPF scheduled Count: 94
ExChange/Loading Neighbors: 0
Router ID conflict state: Normal
Area interface up count: 1

Interface: 172.16.16.5 (Vlanif1001)
Cost: 1       State: BDR       Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 172.16.16.2
Backup Designated Router: 172.16.16.5
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.1
Authtype: Simple   Area flag: Normal
SPF scheduled Count: 83
ExChange/Loading Neighbors: 0
Router ID conflict state: Normal
Area interface up count: 0

Area: 0.0.0.2
Authtype: None   Area flag: Normal
SPF scheduled Count: 81
ExChange/Loading Neighbors: 0
Router ID conflict state: Normal
Area interface up count: 0

Interface: 10.100.100.100 (LoopBack100)
Cost: 0       State: P-2-P     Type: P2P       MTU: 1500
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

Table 7-28 Description of the **display ospf brief** command output

Item	Description
RouterID	Current router ID.
Border Router	Border router: • AS: autonomous system border router (ASBR). • AREA: area border router (ABR). • NSSA: ABR of an NSSA.
Multi-VPN-Instance is not enabled	The current process does not support multi-VPN-instance.

Item	Description
Opaque Capable	Opaque-LSA capability is enabled. To enable the Opaque-LSA capability, run the **opaque-capability enable** command.
Global DS-TE Mode	Globally configured DS-TE mode: ● Non-Standard IETF Mode: The IETF mode is not supported. ● Standard IETF Mode: The IETF mode is supported.
Graceful-restart capability	Whether graceful restart is enabled: ● disabled: Graceful restart is disabled. ● planned only: Planned-GR is supported. ● un-planned: Unplanned-GR is supported. ● totally: Totally GR is supported. ● planned and un-planned: Planned-GR and unplanned-GR are supported. To enable the GR function, run the **graceful-restart (OSPF)** command.
Helper support capability	Whether the Helper mode is enabled: ● enabled: The Helper mode is enabled. ● not configured: The Helper mode is disabled. To configure a device as a GR helper, run the **graceful-restart helper-role (OSPF)** command.
filter capability	Whether the filtering rule of the Helper mode is enabled: ● ip-prefix: The IP prefix filtering rule of the Helper mode is enabled. ● acl-number: The basic ACL filtering rule of the Helper mode is enabled. ● acl-name: The named ACL filtering rule of the Helper mode is enabled. ● disabled: The filtering rule of the Helper mode is disabled.

Item	Description
policy capability	Whether a policy is configured for the Helper mode: ● strict lsa check: The Helper checks all LSAs. To prevent the Helper from checking AS external LSAs, run the **graceful-restart helper-role ignore-external-lsa** command. ● ignore external lsa check: The Helper does not check AS external LSAs. ● planned and un-planned: The Helper supports planned-GR and unplanned-GR. To configure the Helper to support only planned-GR, run the **graceful-restart helper-role planned-only** command. ● planned: The Helper supports only planned-GR.
Applications Supported: MPLS Traffic-Engineering	OSPF supports Traffic Engineering (TE).
Spf-schedule-interval	Interval for performing SPF calculation. To set the interval for OSPF to calculate routes, run the **spf-schedule-interval** command.
Route Preference	Preference of the default route.
ASE Route Preference	Priority of the external route.
Default ASE parameters	Default parameters of the external LSA. ● Metric: default metric of the external LSA. ● Tag: default tag of the external LSA. ● Type: default type of the external LSA.
SPF Computation Count	Number of times that SPF calculation is performed.
RFC 1583 Compatible	Whether RFC 1583 compatibility is enabled. To convert rules defined in RFC 2328 into rules defined in RFC 1583, run the **rfc1583 compatible** command.
Retransmission limitation is disabled	Retransmission limit is disabled. To enable retransmission limit and set the maximum number of retransmissions, run the **retransmission-limit** command.

Item	Description
Import routes limitation is enabled, the maximal limitation value	Maximum number of imported routes is restricted and that the maximum value is displayed.
Area Count	Number of areas in the current process.
Nssa Area Count	Number of NSSAs in the current process.
Process total up interface count	Number of interfaces that are up.
Process valid up interface count	Number of interfaces that are valid.
Flush protect mode	Whether master/slave board switching triggered by abnormal OSPF LSA aging is enabled: • false: Master/Slave board switching triggered by abnormal OSPF LSA aging is disabled. • true: Master/Slave board switching triggered by abnormal OSPF LSA aging is enabled.
ExChange/Loading Neighbors	Number of neighbors in the ExChange/Loading state.
Area	Information about each area in the current process, including ID of the current area in dotted decimal notation.
Authtype	Area authentication type, including none-authentication, simple authentication, MD5 authentication, HMAC-SHA256 authentication, and HMAC-MD5 authentication.
Area flag	Flag used to describe the area attributes, including Transit/Vlink/Stub/Nssa/Normal.
SPF scheduled Count	Number of times that SPF calculation is performed.
Interface	Interface information in the area.
Cost	Cost of an OSPF interface. To set the cost for an OSPF on an interface, run the **ospf cost** command.

Item	Description
State	Interface status: • Down • Waiting • Loopback • P-2-P • DR • BDR • DROTHER DR and BDR DROTHER exist in only broadcast and NBMA networks, and P-2-P exists in only P2P and P2MP Vlinks.
Type	Interface type, including P2P, broadcast, NBMA, and P2MP.
MTU	MTU value of the interface.
Priority	Interface priority.
Designated Router	(Optional) the current interface is not the DR.
Backup Designated Router	(Optional) the current interface is not the BDR.
Timers	Interval of the timer.
Hello	Interval of the Hello timer. To set the interval for sending Hello packets on an interface, run the **ospf timer hello** command.
Dead	Interval of the Dead timer. To set the dead interval after which an interface considers its OSPF neighbor invalid, run the **ospf timer dead** command.
Poll	Interval of the Poll timer. To set the poll interval for sending Hello packets on NBMA networks, run the **ospf timer poll** command.
Retransmit	Interval of the Retransmit timer. To set the interval for retransmitting LSA on an interface, run the **ospf timer retransmit** command.

Item	Description
Transmit Delay	(Optional) delay for transmitting LSAs on the interface. To add the transmission delay to LSAs before they are sent by an interface, run the **ospf trans-delay** command.
Router ID conflict state	Status of the automatic recovery function. The value can be one of the following: • Normal: The automatic recovery function is properly detecting router ID conflict. • Wait select: The automatic recovery function delays defining a new router ID if the device starts after an unexpected delay (two hours by default). • Selecting: The automatic recovery function restarts the OSPF process with the router ID and waits for the restarted OSPF process to take effect. • RtrId Changed: The automatic recovery function determines whether router ID conflict occurs after the new router ID takes effect and returns to the Normal state if no new router ID conflict is detected. • Suspend: If the maximum number of conflict times is reached, automatic recovery function does not define a new router ID any longer. The maximum number of conflict times is three by default.
Area interface up count	The number of interfaces up in the area.

7.4.21 display ospf cumulative

Function

The **display ospf cumulative** command displays OSPF statistics.

Format

display ospf [*process-id*] **cumulative**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If no OSPF process ID is specified, statistics of all the OSPF processes are displayed.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command output helps you troubleshoot OSPF faults.

Example

Display OSPF statistics.

```
<HUAWEI> display ospf cumulative
        OSPF Process 1 with Router ID 10.1.1.1
              Cumulations

 IO Statistics
          Type       Input     Output
         Hello          26         62
  DB Description         3          2
  Link-State Req         1          1
Link-State Update        5          5
  Link-State Ack         3          3
 ASE: 2  Checksum Sum: 233779
 LSAs originated by this router

 Router: 1
 Network: 0
 Sum-Net: 0
 Sum-Asbr: 0
 External: 3
 NSSA: 0
 Opq-Link: 0
 Opq-Area: 0
 Opq-As: 0
 LSAs Originated: 4  LSAs Received: 14

 Routing Table:
   Intra Area: 2  Inter Area: 0  ASE: 2

 Up Interface Cumulate: 2

    Neighbor Cumulate:
==========================================================

    Neighbor cumulative data. (Process 1)
----------------------------------------------------------
Down:       0 Init:       0 Attempt:    0 2-Way:    0
```



```
Exstart:    0 Exchange:    0 Loading:    0 Full:     1
Retransmit Count: 0

   Neighbor cumulative data. (Total)
---------------------------------------------------------
Down:       0 Init:       0 Attempt:    0 2-Way:    0
Exstart:    0 Exchange:    0 Loading:    0 Full:     1
Retransmit Count: 0
```

Table 7-29 Description of the **display ospf cumulative** command output

Item	Description
IO Statistics	Statistics of the transmitted packets and LSAs.
Type	OSPF packet type.
Input	Number of received packets.
Output	Number of sent packets.
Hello	OSPF Hello packet.
DB Description	OSPF Database Description packet.
Link-State Req	OSPF Link State Request packet.
Link-State Update	OSPF Link State Update packet.
Link-State Ack	OSPF Link State Acknowledgement packet.
Checksum Sum	Checksum of the AS external LSA.
ASE	Number of ASE routes (If there are no ASE routes, **Disabled** is displayed.).
LSAs originated by this router	Detailed statistics of the transmitted LSAs.
Router	Router LSA.
Network	Network LSA.
Sum-Net	Type 3 summary LSA.
Sum-Asbr	Type 4 summary LSA.
External	AS external LSA.
NSSA	NSSA.
Opq-Link	Number of Type 9 Opaque LSAs.
Opq-Area	Number of Type 10 Opaque LSAs.
Opq-As	Number of Type 11 Opaque LSAs
LSAs Originated	Generated LSAs.
LSAs Received	Received LSAs.

Item	Description
Routing Table	Routing table.
Intra Area	Number of intra-area routes.
Inter Area	Number of inter-area routes.
Up Interface Cumulate	Statistics of up state interface.
Neighbor Cumulate	Statistics of neighbors.
Neighbor cumulative data	Detailed statistics of neighbors: • Down • Init • Attempt • 2-Way • Exstart • Exchange • Loading • Full
Retransmit Count	Total number of nodes in the retransmission list.

7.4.22 display ospf error

Function

The **display ospf error** command displays OSPF error information.

Format

display ospf [*process-id*] **error** [**lsa** | **interface** *interface-type interface-number*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If no OSPF process ID is specified, error information of all OSPF processes is displayed.	The value is an integer ranging from 1 to 65535.
lsa	Displays the OSPF LSA errors.	-

Parameter	Description	Value
interface *interface-type interface-number*	Specifies the type and number of an interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

When locating OSPF faults, you can run the **display ospf error** command to obtain OSPF error information. You can then analyze OSPF faults according to the OSPF error information.

Example

Display OSPF error information.

```
<HUAWEI> display ospf error
        OSPF Process 1 with Router ID 10.1.1.1
                OSPF error statistics

General packet errors:
 0     : IP: received my own packet     0     : Bad packet
 0     : Bad version                    0     : Bad checksum
 0     : Bad area id                    0     : Drop on unnumbered interface
 1     : Bad virtual link               0     : Bad authentication type
 0     : Bad authentication key         0     : Packet too small
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor
 0     : Bad net segment                0     : Extern option mismatch
 0     : Router id confusion            0     : Bad authentication sequence number

HELLO packet errors:
 0     : Netmask mismatch               0     : Hello timer mismatch
 0     : Dead timer mismatch            0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown          0     : Invalid Source Address

DD packet errors:
 0     : Neighbor state low             0     : Unknown LSA type
 0     : MTU option mismatch

LS ACK packet errors:
 0     : Neighbor state low             0     : Unknown LSA type

LS REQ packet errors:
 0     : Neighbor state low             0     : Empty request
 0     : Bad request

LS UPD packet errors:
 0     : Neighbor state low             0     : Newer self-generate LSA
 0     : LSA checksum bad               0     : Received less recent LSA
 0     : Unknown LSA type

Opaque errors:
```



```
 0    : 9-out of flooding scope       0    : 10-out of flooding scope
 0    : 11-out of flooding scope      0    : Unknown TLV type

Retransmission for packet over Limitation errors:
 0    : Number for DD Packet          0    : Number for Update Packet
 0    : Number for Request Packet

Receive Grace LSA errors:
 0    : Number of invalid LSAs        0    : Number of policy failed LSAs
 0    : Number of wrong period LSAs

Configuration errors:
0    : Tunnel cost mistake
```

Table 7-30 Description of the **display ospf error** command output

Item	Description
General packet errors	Indicates general packet errors.
IP: received my own packet	Indicates that the packet sent by its own interface is received and therefore the packet is not processed.
Bad packet	Indicates that the parsed packet is incorrect, including the checksum of the length field.
Bad version	Indicates that the OSPF version is incorrect, that is, it is not version 2.
Bad checksum	Indicates that the OSPF checksum is incorrect.
Bad area id	Indicates that the area ID in the received packet does not match the local area ID. (Vlink can receive packets from only Area 0 and its own area.)
Drop on unnumbered interface	Indicates that the unnumbered rather than P2P interface receives packets (the interface must be of the P2P type).
Bad virtual link	Indicates that the Vlink receives invalid packets.
Bad authentication type	Indicates that packet authentication is incorrect. If the value of this field keeps increasing, the OSPF authentication types of the two devices that establish the neighbor relationship are inconsistent. In this case, run the **area-authentication-mode** command to configure the same authentication type for the two devices.
Bad authentication key	Packet authentication key is incorrect.
Packet too small	Indicates that the length of the received packet does not equal the sum of the IP header length and the packet length.
Packet size > ip length	Indicates that the length of the OSPF packet is greater than the permitted length of the IP packet.

Item	Description
Transmit error	Indicates that sending packets to the socket fails.
Interface down	Indicates the number of times that the OSPF interface goes Down.
Unknown neighbor	Indicates that OSPF packets are received from non-OSPF neighbors on NBMA networks, virtual links, and sham links.
HELLO packet errors	Indicates Hello packet errors.
Netmask mismatch	Indicates that the address mask does not match the local address mask.
Hello timer mismatch	Indicates that the Hello intervals on the two ends are inconsistent. If the value of this field keeps increasing, the value of the Hello timers on the two devices that establish the neighbor relationship are inconsistent. In this case, check the interface configurations of the two devices and run the **ospf timer hello** command to set the same value for the Hello timers.
Dead timer mismatch	Indicates that the Dead intervals on the two ends are inconsistent. If the value of this field keeps increasing, the values of the dead timers on the two devices that establish the neighbor relationship are inconsistent. In this case, check the interface configurations of the two devices and run the **ospf timer dead** command to set the same value for the dead timers.
Extern option mismatch	Indicates that the extension attributes of the Hello packets on the two ends are inconsistent. If the value of this field keeps increasing, the area types of the two devices that establish the neighbor relationship are inconsistent (the area type of one device is common area, and the area type of the other device is stub area or NSSA). In this case, configure the same area type for the two devices (in the OSPF area view, the **stub** command indicates the area type is stub and the **stub** command indicates the area type is nssa).
Bad net segment	The source address of received packets is not on the same network segment as the IP address of the interface that receives packets.
Router id confusion	Indicates that the router IDs on the two ends are the same.

Item	Description
Bad authentication sequence number	Indicates bad authentication sequence number errors.
Virtual neighbor unknown	Indicates that the router ID of the packet is inconsistent with that of the neighbor that is configured by the virtual link.
NBMA neighbor unknown	Indicates that the status of the NBMA neighbor is not active.
Invalid Source Address	Indicates that the source address of LSA is invalid.
DD packet errors	Indicates DD packet errors.
Neighbor state low	Indicates the following situations: • A DD packet is received but its neighbor status is lower than 2-way. • An LSR packet is received but its neighbor status is lower than Exchange. • An LSU packet is received but its neighbor status is lower than Exchange. • An LSAck packet is received but its neighbor status is lower than Exchange.
Unknown LSA type	Indicates the unknown LSA type.
MTU option mismatch	Indicates that the MTU check of the OSPF interface is enabled and the MTU of the DD packet received by the interface is greater than the MTU of the interface.
LS ACK packet errors	Indicates LSAck packet errors.
Bad ack	Indicates the number of times that incorrect LSAck packets are received.
Duplicate ack	Indicates the number of times that duplicate LSAck packets are received.
LS REQ packet errors	Indicates LSR packet errors.
Empty request	Indicates empty LSR packets.
Bad request	Indicates the BadRequest event in the protocol.
LS UPD packet errors	Indicates LSU packet errors.
Newer self-generate LSA	Indicates the number of new self-generated LSAs. This field is reserved for future use.
LSA checksum bad	Indicates that the LSA checksum is incorrect.
Received less recent LSA	Indicates that the LSA older than the local LSA is received.

Item	Description
Opaque errors	Indicates opaque errors.
9-out of flooding scope	Indicates the number of Type 9 LSAs that exceed the flooding scope.
10-out of flooding scope	Indicates the number of Type 10 LSAs that exceed the flooding scope.
11-out of flooding scope	Indicates the number of Type 11 LSAs that exceed the flooding scope.
Unknown TLV type	Indicates the unknown TLV type.
Retransmission for packet over Limitation errors	Indicates the number of times that retransmitting packets expires.
Number for DD Packet	Indicates the number of times that retransmitting DD packets expires.
Number for Update Packet	Indicates the number of times that retransmitting LSU packets expires.
Number for Request Packet	Indicates the number of times that retransmitting LSR packets expires.
Receive Grace LSA errors	Indicates the number of received incorrect Grace LSAs.
Number of invalid LSAs	Indicates the total number of invalid LSAs.
Number of policy failed LSAs	Indicates the total number of policy failed LSAs.
Number of wrong period LSAs	Indicates the total number of wrong period LSAs.
Configuration errors	Indicates configuration errors.
Tunnel cost mistake	Indicates the number of times that the cost of the OSPF tunnel interface is smaller than 1. This count increases by one each time the cost of the OSPF tunnel interface is smaller than one. If the cost is smaller than one, the cost is calculated as one.

Display the OSPF LSA errors.

```
<HUAWEI> display ospf error lsa
 OSPF Process  1  with Router ID 10.1.1.14

 Last Received Bad LSA Header
   LS Age            : 36
   Link State Type    : 0x0008
   Link State ID      : 0.0.1.66
   Advertising Router : 10.10.10.22
   LS Sequence Number : 0x80000002
   LS Checksum        : 0x00bd2e
   Length             : 96
```



```
Interface        : Vlanif100
Recv Time        : 2011-05-27 14:37:17
```

Table 7-31 Description of the display ospf error lsa command output

Item	Description
Last Received Bad LSA Header	Bad LSA information received last time.
LS Age	Aging time of the LSA.
Link State Type	LSA type.
Link State ID	LSA state ID.
Advertising Router	Router that advertises or generates LSAs.
LS Sequence Number	Sequence number in the LSA header.
LS Checksum	LSA checksum.
Length	Size of the LSA.
Interface	LSA receiving interface.
Recv Time	LSA receiving time.

7.4.23 display ospf global-statistics

Function

The **display ospf global-statistics** command displays global OSPF statistics. If no OSPF process ID is specified, brief information about all the OSPF processes is displayed.

Format

display ospf global-statistics { **process** *process-id* | **vpn-instance** *vpn-instance-name* | **public-instance** | **timewheel** | **brief** }

Parameters

Parameter	Description	Value
process *process-id*	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Parameter	Description	Value
public-instance	Displays the statistics of all the public network instances.	-
timewheel	Displays the number of updated or aged LSAs in different periods.	-
brief	Displays brief information.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospf global-statistics** command to check global OSPF statistics for fault location.

Example

Display global OSPF statistics, including statistics of neighbors, processes, and instances.

```
<HUAWEI> display ospf global-statistics brief
  Neighbor cumulative data (OSPF total statistics):
  ------------------------------------------------------------
  Down:              0       Init:              0      Attempt:     0    2-Way:     0
  Exstart:           0       Exchange:          0      Loading:     0    Full:   2400
  Instance Numer:    1       Process Number:    2
  HighSocketExpire:  27      HighSocketEmpty:  526275
  Total Neighbor Number:                0
  Total Press:   LOW
```

Table 7-32 Description of the **display ospf global-statistics brief** command output

Item	Description
Neighbor cumulative data (OSPF total statistics)	Neighbor statistics.
Instance Number	Number of instances, including public network instances and VPN instances.
Process Number	Number of OSPF processes running on the device.
HighSocketExpire	Number of unprocessed messages in high-priority queues.
HighSocketEmpty	Number of processed messages in high-priority queues.

Item	Description
Total Neighbor Number	Total number of OSPF neighbors.
Total Press	Total pressure of the current service. In most cases, the Total Press field is **LOW**. If the number of retransmission timers or the number of LSAs being flooded in the OSPF processes exceeds a certain value (default 300), or the number of OSPF neighbors exceeds 300, the Total Press field is **HIGH**. If the number of retransmission timers and the number of LSAs being flooded in the OSPF processes fall below a certain value (default 300), or the number of OSPF neighbors falls below 300, the Total Press field is **LOW**.

Display global statistics of OSPF process 1.

```
<HUAWEI> display ospf global-statistics process 1
  OSPF 1 statistics data:
  ----------------------------------------
  LSA NUM of Flood cache:         0
  Packet NUM of FloodUpdt Hash:         0
  Packet NUM of Flood Queue:         0
```

Table 7-33 Description of the **display ospf global-statistics process 1** command output

Item	Description
LSA NUM of Flood cache	Number of LSAs being flooded in the OSPF process.
Packet NUM of FloodUpdt Hash	Number of Update packets waiting to be flooded in the OSPF process, and the total size of the Update packets does not reach the MTU.
Packet NUM of Flood Queue	Number of Update packets waiting to be flooded in the OSPF process, and the total size of the Update packets reaches the MTU.

Display OSPF time wheel information.

```
<HUAWEI> display ospf global-statistics timewheel
 ================== TimeWheel Info Begin ==================
 TimeWheel current index is 2845, datanode count is 8
 Bucket Number: 3596, Expiry time: 1306(s), Datanode Count: 0
 High expiry time: 0, Low expiry time 527742(s)
```



```
Bucket Number: 3597, Expiry time: 1307(s), Datanode Count: 0
High expiry time: 0, Low expiry time 527743(s)
Bucket Number: 3598, Expiry time: 1308(s), Datanode Count: 0
```

Table 7-34 Description of the **display ospf global-statistics timewheel** command output

Item	Description
TimeWheel current index	Current index of the time wheel.
datanode count	Total number of nodes in the time wheel.
Bucket Number	Total number of indexes of the time wheel.
Expiry time	Expiry time of the index.
High expiry time	High expiry time corresponding to the index in the time wheel after the system starts.
Low expiry time	Low expiry time corresponding to the index in the time wheel after the system starts.

7.4.24 display ospf graceful-restart

Function

The **display ospf graceful-restart** command displays the status of OSPF GR.

Format

display ospf [*process-id*] **graceful-restart** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
verbose	Displays detailed information about OSPF GR.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospf graceful-restart** command to view details and statistics of the GR.

Example

Display OSPF GR information.

```
<HUAWEI> display ospf graceful-restart

        OSPF Process 1 with Router ID 10.1.1.1
 Graceful-restart capability     : enabled
 Graceful-restart support        : planned and un-planned, totally
 Helper-policy support           : planned and un-planned, strict lsa check
 Current GR state                : normal
 Graceful-restart period         : 120 seconds

 Number of neighbors under helper:
  Normal neighbors    : 0
  Virtual neighbors   : 0
  Sham-link neighbors : 0
  Total neighbors     : 0

 Number of restarting neighbors : 0

 Last exit reason:
  On graceful restart : successful exit
  On Helper           : none
```

Display detailed information about OSPF GR.

```
<HUAWEI> display ospf graceful-restart verbose

        OSPF Process 1 with Router ID 10.1.1.1
 Graceful-restart capability     : enabled
 Graceful-restart support        : planned and un-planned, totally
 Helper-policy support           : planned and un-planned, strict lsa check
 Current GR state                : normal
 Graceful-restart period         : 120 seconds

 Number of neighbors under helper:
  Normal neighbors    : 0
  Virtual neighbors   : 0
  Sham-link neighbors : 0
  Total neighbors     : 0
 Number of restarting neighbors : 0

 Last exit reason:
  On graceful restart : successful exit
  On Helper           : none

 All area count      : 1

              Area ID    : 0.0.0.0

 Authtype       : None     Area flag : Normal

 Normal interface count: 1

 Interface: 10.1.1.1 (Vlanif100)
 GR state : normal                State: P-2-P          Type: P2P
 Last Helper Exit reason: none
 Neighbor count of this interface : 1

 Neighbor         IP address       GR state    Last Helper Exit reason
 10.2.2.2         10.1.1.2         Normal      none
```


Table 7-35 Description of the **display ospf graceful-restart** command output

Item	Description
Graceful-restart capability	Whether IETF GR is enabled: • enabled • disabled
Graceful-restart support	IETF GR mode that is supported currently: • planned: indicates that only planned-GR is supported. • planned and un-planned: indicates that both planned GR and unplanned GR are supported. • totally: indicates that totally GR is supported. • partial: indicates that partial GR is supported.
Helper-policy support	Policy that supports the Helper: • planned: indicates that the Helper supports only planned GR. • planned and un-planned: indicates the Helper supports both planned GR and unplanned GR. • strict lsa check: indicates that the Helper supports strict external LSA check. • ignore external lsa check: indicates that the Helper does not check external LSAs. • never: indicates that the device does not support the Helper mode.
Current GR state	Current GR status: • Normal: indicates that GR is in the Normal state. • Under GR: indicates that the device enters the Restarter mode. • Under Helper: indicates that the device enters the Helper mode.
Graceful-restart period	GR period.

Item	Description
Number of neighbors under helper	Number of neighbors in the Helper state: • Normal neighbors: indicates the number of normal neighbors. • Virtual neighbors: indicates the number of virtual neighbors. • Sham-link neighbors: indicates the number of sham link neighbors. • Total neighbors: indicates the total number of neighbors.
Number of restarting neighbors	Number of restarted devices displayed on the Helper.
Last exit reason	Reason why a device exits from GR last time: • On graceful restart: indicates the reason that the Restarter exits from GR. • On Helper: indicates the reason that the Helper exits from GR.
On graceful restart	Reason why the Restarter exits from GR: • 1-way hello received: indicates that 1-way Hello packets are received. • back-link check failed: indicates that the back link check fails. • DR election fail: indicates that DR election fails. • grace period expired: indicates that the GR period expires. • interface state change: indicates that the interface state machine changes. • none: indicates that the device never performs GR since startup. • successful exit: indicates that the Restarter successfully exits from GR. • two Grace-LSAs received: indicates that two Grace LSAs are received.

Item	Description
On Helper	Reason why the Helper exits from GR: • flooding change LSA: indicates that the changed LSAs are received. • grace period expired: indicates that the GR period expires. • graceful restart unconfigured at process level: indicates that the GR function of the OSPF process is disabled. • interface state change: indicates that the interface state machine changes. • policy check failed for received grace LSA: indicates the policy that mismatches the Helper. • received 1-way hello packet: indicates that 1-way Hello packets are received. • received flushed grace LSA: indicates that flushed Grace LSAs are received. • received multiple grace LSA: indicates that multiple Grace LSAs are received. • neighbor reset: indicates that GR is disabled when neighbors are in the Helper mode. • none: indicates that the device never enters the Helper mode since startup. • successful exit: indicates that the Helper successfully exits from GR.
All area count	Number of areas in the process.
Area ID	Area ID.
Authtype	Authentication type.
Area flag	Area attributes: • Normal • NSSA • Stub
Normal interface count	Number of interfaces in the area.
Interface	IP address of the interface.

Item	Description
GR state	GR status of the interface: • Normal: indicates that the device is in the Normal state. • Restarter: indicates that the device enters the Restarter mode. • Helper: indicates that the device enters the Helper mode.
State	Interface status: • P-2-P • DR • BDR • DROther • Waiting • Down
Type	Interface type: • P2P • P2MP • NBMA • Broadcast

Item	Description
Last Helper Exit reason	Reason why the neighbor exits from the Helper mode the last time: • none: indicates that the device never enters the Helper mode since startup. • successful exit: indicates that the neighbor successfully exits from the Helper mode. • grace period expired: indicates that the GR period expires. • received flushed grace LSA: indicates that flushed Grace LSAs are received. • flooding change LSA: indicates that the changed LSAs are received. • received multiple grace LSA: indicates that multiple Grace LSAs are received. • received 1-way hello packet: indicates that 1-way Hello packets are received. • policy check failed for received grace LSA: indicates the policy that mismatches the Helper. • neighbor reset: indicates that the network topology changes after the **reset** command is run on the neighbor of the Helper. • interface state change: indicates that the interface state machine changes. • graceful restart unconfigured at process level: indicates that the neighbor is not configured with the GR function.
Neighbor count of this interface	Total number of neighbors of this interface.
Neighbor	Router ID of the neighbor.
IP address	IP address of the neighboring interface.
GR state	GR status of the neighbor: • Normal: indicates that the neighbor is in the Normal state. • Restarter: indicates that the neighbor enters the Restarter mode. • Helper: indicates that the neighbor enters the Helper mode.

7.4.25 display ospf interface

Function

The **display ospf interface** command displays information about OSPF interfaces.

Format

display ospf [*process-id*] **interface** [**all** | *interface-type interface-number*] [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
all	Displays information about all OSPF interfaces.	-
interface-type interface-number	Specifies the interface type and the interface number.	-
verbose	Displays verbose configuration information.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display ospf interface** command output displays the configuration and operating status of OSPF, which facilitate fault location and configuration verification.

Example

Display OSPF interface information.

```
<HUAWEI> display ospf interface

        OSPF Process 1 with Router ID 192.168.1.1
              Interfaces

 Area: 0.0.0.0       (MPLS TE not enabled)
 IP Address      Type        State    Cost  Pri   DR            BDR
 192.168.1.2     Broadcast   DR       1     1     192.168.1.2   192.168.1.3

 Area: 0.0.0.1       (MPLS TE not enabled)
```

```
IP Address      Type         State    Cost  Pri  DR             BDR
172.16.0.1      Broadcast    DR       1     1    172.16.0.1     172.16.0.2
```

Table 7-36 Description of the **display ospf interface** command output

Item	Description
Area	ID of the area to which the interface belongs.
IP Address	IP address of the interface (regardless of whether TE is enabled on the interface).
Type	Interface type: P2P, PTMP, broadcast, or NBMA.
State	Status of the interface, which is determined by the OSPF interface state machines: ● Down: The status of the interface is Down. If an interface is Down, the interface is unavailable and cannot be used to transmit traffic. ● Loopback: The interface connecting to the network on the device is in the Loopback state. The loopback interface cannot be used to transmit data but can collect interface information by performing the ICMP ping operation or bit error detection. ● Waiting: The device is determining the DR and BDR on the network. The DR or BDR election mechanism should not be implemented until the waiting period ends. This prevents unnecessary changes in the DR and BDR roles. ● P-2-P: The interface is connected to a P2P network or a virtual link. ● DROther: The device itself is not elected as the DR or BDR. Instead, another device connecting to the broadcast network or NBMA network is elected as the DR. The device starts to set up adjacency with the DR and BDR (if existing). ● BDR: The device functions as the BDR on the network, and will turn into a DR when the current DR fails. The device sets up adjacency with other devices that access the network. ● DR: The device functions as the DR on the network. The device sets up adjacency with other devices that access the network.
Cost	Cost of the interface.

Item	Description
Pri	Priority of the device interface during the DR and BDR election The greater the value, the higher the priority.
DR	DR of the network where the interface resides.
BDR	BDR of the network where the interface resides.
Timer Hello	Interval for sending Hello packets.

Display detailed information about an OSPF interface.

```
<HUAWEI> display ospf interface Vlanif 501 verbose

        OSPF Process 1 with Router ID 192.168.2.1
              Interfaces


 Interface: 192.168.100.2 (Vlanif501)
 Cost: 1       State: BDR       Type: Broadcast    MTU: 1500
 Priority: 1
 Designated Router: 192.168.100.1
 Backup Designated Router: 192.168.100.2
 Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
  IO Statistics
            Type        Input     Output
           Hello          11         10
    DB Description         3          2
    Link-State Req         1          1
Link-State Update          4          3
    Link-State Ack         2          3
 ALLSPF GROUP
 ALLDR GROUP
 OpaqueId: 0    PrevState: Waiting
 Effective cost: 1, enabled by OSPF Protocol
 Suppress flapping peer: enable(flapping-count: 0, threshold: 10)
```

Table 7-37 Description of the **display ospf interface verbose** command output

Item	Description
IO Statistics	Statistics about received and sent OSPF packets.
Type	OSPF packet type.
Input	Number of OSPF packets that the interface receives.
Output	Number of OSPF packets sent by the interface.
DB Description	Statistics about received and sent OSPF DD packets.

Item	Description
Link-State Req	Statistics about received and sent OSPF LSR packets.
Link-State Update	Statistics about received and sent OSPF LSU packets.
Link-State Ack	Statistics about received and sent OSPF LSAck packets.
ALLSPF GROUP	ALLSPF GROUP that the interface joins.
ALLDR GROUP	ALLDR GROUP that the interface joins.
OpaqueId	Opaque ID of the interface.
PrevState	Previous state of the interface.
Effective cost	Effective cost of the interface: • enabled by OSPF Protocol: default value or the one configured using the **ospf cost** command. • enabled by RUI: RUI route cost. • enabled by IGP_LDP: interface cost configured in LDP. • enabled by BGP_IGP: interface cost configured in BGP. • enabled by Tunnel: cost generated after a TE Tunnel is configured.

Item	Description

Suppress flapping peer	Status of OSPF neighbor relationship flapping suppression:
	● enable: OSPF neighbor relationship flapping suppression is enabled.
	– flapping-count: number of valid flapping_events
	If the interval between two successive neighbor status changes from Full to a non-Full state is shorter than *detecting-interval*, a valid flapping_event is recorded, and the flapping_count is incremented by 1. To change *detecting-interval*, run the **ospf suppress-flapping peer detecting-interval** *detecting-interval* command.
	– threshold: flapping suppression threshold
	When the flapping_count reaches or exceeds **threshold**, flapping suppression takes effect. To configure the threshold, run the **ospf suppress-flapping peer threshold** *threshold* command.
	● disable: OSPF neighbor relationship flapping suppression is disabled. In this case, the following information is displayed, without flapping-count or threshold:
	`Suppress flapping peer: disable`
	● hold-down: OSPF neighbor relationship flapping suppression works in Hold-down mode. In this case, an example of the displayed information is as follows:
	`Suppress flapping peer: hold-down(start: 2016-01-02 09:58:41, remain-interval: 476 sec)`
	– start: time when the flapping suppression started
	– remain-interval: remaining time of the flapping suppression
	● hold-max-cost: OSPF neighbor relationship flapping suppression works in Hold-max-cost mode. In this case, an example of the displayed information is as follows:
	`Suppress flapping peer: hold-max-cost(start: 2016-01-02 09:58:41, remain-interval: 476 sec)`

7.4.26 display ospf lsdb

Function

The **display ospf lsdb** command displays an OSPF Link State Database (LSDB).

Format

display ospf [*process-id*] **lsdb** [**brief**]

display ospf [*process-id*] **lsdb** [{ **router** | **network** | **summary** | **asbr** | **ase** | **nssa** | **opaque-link** | **opaque-area** | **opaque-as** } [*link-state-id*]] [**originate-router** [*advertising-router-id*] | **self-originate**] [**age** { **min-value** *min-age-value* | **max-value** *max-age-value* } *]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
brief	Displays brief information about the LSDB.	-
router	Displays information about the router LSA.	-
network	Displays information about the network LSA.	-
summary	Displays information about the network summary LSA.	-
asbr	Displays information about the ASBR summary LSA.	-
ase	Displays information about the AS external LSA.	-
nssa	Displays information about the status of external links in the NSSA.	-
opaque-link	Displays information about the opaque link LSA.	-
opaque-area	Displays information about the opaque area LSA.	-

Parameter	Description	Value
opaque-as	Displays information about the opaque AS LSA.	-
originate-router	Displays the LSA of the advertising router.	-
link-state-id	Specifies the ID of an LSA.	The value is an IP address in dotted decimal notation.
advertising-router-id	Specifies the router ID of the device that advertises the LSA.	The value is an IP address in dotted decimal notation.
self-originate	Displays information about the self-originated LSA.	-
age	Displays the LSAs that meet the age filtering rule.	-
min-value *min-age-value*	Displays information about only LSAs with the **age** value greater than or equal to the *min-age-value* value.	The value is an integer ranging from 0 to 3600.
max-value *max-age-value*	Displays information only about LSAs with the **age** value less than or equal to the *max-age-value* value.	The value is an integer ranging from 0 to 3600.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To check information about the LSDB, run the **display ospf lsdb** command. You can choose to view:

- Brief information about the LSDB
- LSAs of a specified type
- LSAs of the originating device
- Locally originated LSAs

The command output helps you troubleshoot OSPF faults.

Example

Display information about the OSPF LSDB.

```
<HUAWEI> display ospf lsdb
         OSPF Process 1 with Router ID 10.1.1.1
                 Link State Database

                         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter        Age  Len   Sequence   Metric
 Router    10.1.1.1        10.1.1.1         1348  48    800007C9    10000
 Router    10.1.1.2        10.1.1.2         1078  60    80000654       0
 Network   192.168.10.1    192.168.10.1     1349  32    80000060       0

                 AS External Database
 Type      LinkState ID    AdvRouter        Age  Len   Sequence   Metric
 External  0.0.0.0         10.1.1.2         1743  36    800005FE       1
```

Table 7-38 Description of the **display ospf lsdb** command output

Item	Description
Area	Area whose LSDB information needs to be displayed.
Type	LSA type: Router, Network, Sum-Net, Sum-Asbr, NSSA, External, Opq-Link, Opq-Area, and Opq-As.
LinkState ID	Link State ID in the LSA header.
AdvRouter	Device that advertises or generates LSAs.
Age	Aging time of the LSA.
Len	Size of the LSA.
Sequence	Sequence number in the LSA header.
Metric	Metric.
AS External Database	LSDB that contains external LSAs.

Display brief information about the OSPF LSDB.

```
<HUAWEI> display ospf lsdb brief
        OSPF Process 1 with Router ID 10.1.1.1
                LS Database Statistics

 Area ID        Stub   Router   Network  S-Net   S-ASBR   Type-7   | Subtotal
 0.0.0.0         0      3        2        0       0        0       | 5
 Total           0      3        2        0       0        0       |
 ----------------------------------------------------------------------+---------
 Area ID        Opq-9  Opq-10                                        | Subtotal
 0.0.0.0         0      0                                           | 0
 Total           0      0                                           |
 ----------------------------------------------------------------------+---------
              ASE    Opq-11                                          | Subtotal
 Total           0      0                                           | 0
 ----------------------------------------------------------------------+---------
                                                                    | Total
                                                                    | 5
```


Table 7-39 Description of the **display ospf lsdb brief** command output

Item	Description
Area ID	Area to which the device belongs.
Stub	Number of stub networks.
Router	Number of Router LSAs.
Network	Number of network LSAs.
S-Net	Number of Sum-Net LSAs.
S-ASBR	Number of Sum-Asbr LSAs.
Type-7	Number of Type 7 LSAs.
Opq-9	Number of Opq-Link LSAs.
Opq-10	Number of Opq-Area LSAs.
ASE	Number of ASE routes.
Opq-11	Number of Opq-As LSAs.
Subtotal	Total number of sub-items.
Total	Total quantity.

Display information about router LSAs in the LSDB.

```
<HUAWEI> display ospf lsdb router
        OSPF Process 1 with Router ID 10.1.1.1
                    Area: 0.0.0.0
                Link State Database

 Type      : Router
 Ls id     : 10.2.2.2
 Adv rtr   : 10.2.2.2
 Ls age    : 52
 Len       : 48
 Options   :  E
 seq#      : 80000006
 chksum    : 0xbf5a
 Link count: 2
   Link ID: 10.2.1.2
   Data   : 10.2.1.2
   Link Type: TransNet
   Metric : 1
   Link ID: 10.2.2.2
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 0
   Priority : Medium
```

Table 7-40 Description of the **display ospf lsdb router** command output

Item	Description
Ls id	Link State ID in the LSA header.

Item	Description
Adv rtr	Device that advertises or generates LSAs.
Ls age	Aging time of the LSA.
Len	Size of the LSA.
Options	Options field: • E: allows Flood AS-external-LSAs. • MC: forwards IP multicast packets. • N/P: processes Type-7 LSAs. • DC: processes the required links.
seq#	Sequence number, which is used to check the order of LSAs.
chksum	LSA checksum.
Link count	Number of links.
Link ID	Link ID of the router LSA, which is classified according to the link type. • If the type of link is **Point-to-Point**, **Link ID** indicates the router ID of a neighbor. • If the type of link is **TransNet**, **Link ID** indicates IP address of DR. • If the type of link is **Stub**, **Link ID** indicates IP address. • If the type of link is **Virtual Link**, **Link ID** indicates the router ID of a neighbor.
Data	Link data of the router LSA. • If the type of link is Point-to-Point, TransNet, or Virtual Link, **Data** indicates IP address. • If the type of link is Stub, **Data** indicates mask of IP address.
Link Type	Link type of the router LSA: P-2-P, TransNet, StubNet, or Virtual.
Metric	Link metric of the router LSA.

Item	Description
Priority	OSPF convergence priorities: • Critical: indicates that the convergence priority of OSPF routes is critical. • High: indicates that the convergence priority of OSPF routes is high. • Medium: indicates that the convergence priority of OSPF routes is medium. • Low: indicates that the convergence priority of OSPF routes is low.

Display information about network LSAs in the LSDB.

```
<HUAWEI> display ospf 1 lsdb network 10.1.1.1
        OSPF Process 1 with Router ID 10.1.1.1
                    Area: 0.0.0.0
              Link State Database

 Type      : Network
 Ls id     : 10.1.1.1
 Adv rtr   : 10.1.1.1
 Ls age    : 167
 Len       : 32
 Options   :  E
 seq#      : 80000002
 chksum    : 0x3408
 Net mask  : 255.255.255.0

 Attached Router    10.2.2.2
 Attached Router    10.1.1.1
```

Table 7-41 Description of the **display ospf lsdb network** command output

Item	Description
Net mask	Network mask of the network LSA.
Attached Router	Device that is connected to the network.

Display information about network summary LSAs in the LSDB.

```
<HUAWEI> display ospf 1 lsdb summary 10.20.1.0
        OSPF Process 1 with Router ID 10.1.1.1
                  Area: 0.0.0.0
              Link State Database

 Type      : Sum-Net
 Ls id     : 10.1.1.0
 Adv rtr   : 10.2.2.2
 Ls age    : 419
 Len       : 28
 Options   :  E
 seq#      : 80000001
 chksum    : 0x1d21
```



```
Net mask  : 255.255.255.0
Tos 0  metric: 1
Priority  : Medium
```

Table 7-42 Description of the **display ospf lsdb summary** command output

Item	Description
Net mask	Network mask of the network summary LSA.
Tos	Type of service of the network summary LSA.
Metric	Metric or cost of the route from the advertising router to the network, which is carried in the network summary LSA.
Priority	OSPF convergence priorities: ● Critical: indicates that the convergence priority of OSPF routes is critical. ● High: indicates that the convergence priority of OSPF routes is high. ● Medium: indicates that the convergence priority of OSPF routes is medium. ● Low: indicates that the convergence priority of OSPF routes is low.

Display information about ASBR summary LSAs in the LSDB.

```
<HUAWEI> display ospf 1 lsdb asbr 10.2.2.2
        OSPF Process 1 with Router ID 10.1.1.1
                Area: 0.0.0.2
              Link State Database

 Type      : Sum-Asbr
 Ls id     : 10.2.2.2
 Adv rtr   : 10.1.1.1
 Ls age    : 90
 Len       : 28
 Options   :  E
 seq#      : 80000001
 chksum    : 0xec62
 Tos 0  metric: 1
```

Display information about AS external LSAs in the LSDB.

```
<HUAWEI> display ospf 100 lsdb ase 10.1.1.0
OSPF Process 1 with Router ID 10.1.1.1
Link State Database

 Type      : External
 Ls id     : 10.1.1.0
 Adv rtr   : 10.2.2.2
 Ls age    : 569
 Len       : 36
 Options   :  E
 seq#      : 80000002
 chksum    : 0x90d0
 Net mask  : 255.255.255.0
```



```
Tos 0  Metric: 1
E type    : 2
Forwarding Address : 0.0.0.0
Tag       : 1
Priority  : Medium
```

Table 7-43 Description of the **display ospf lsdb ase** command output

Item	Description
Net mask	Network mask of the ASE or NSSA LSA.
Tos	Type of service of the ASE or NSSA LSA.
Metric	Metric or cost of the route from the advertising router to the network, which is carried in the ASE or NSSA LSA.
E type	E type of the ASE or NSSA LSA.
Forwarding Address	Forwarding address of the ASE or NSSA LSA.
Tag	32-bit tag, which is carried in Type 5 and Type 7 LSAs to avoid routing loops
Priority	OSPF convergence priorities: ● Critical: indicates that the convergence priority of OSPF routes is critical. ● High: indicates that the convergence priority of OSPF routes is high. ● Medium: indicates that the convergence priority of OSPF routes is medium. ● Low: indicates that the convergence priority of OSPF routes is low.

Display information about NSSA external LSAs in the LSDB.

```
<HUAWEI> display ospf 1 lsdb nssa 192.168.1.0
        OSPF Process 1 with Router ID 10.1.1.1
                Area: 0.0.0.1
             Link State Database

 Type      : NSSA
 Ls id     : 10.1.1.0
 Adv rtr   : 10.2.2.2
 Ls age    : 521
 Len       : 36
 Options   : None
 seq#      : 80000005
 chksum    : 0x9ea7
 Net mask  : 255.255.255.0
 Tos 0  Metric: 1
 E type    : 2
 Forwarding Address : 10.1.1.2
 Tag       : 1
 Priority  : Medium
```


Display information about Opaque-link LSAs in the LSDB.

```
<HUAWEI> display ospf 1 lsdb opaque-link
         OSPF Process 1 with Router ID 10.1.1.1
Area: 0.0.0.0                                          Link State Database
 Link State Database for interface 10.1.1.1 (Vlanif200) Type: Broadcast

  Type      : Opq-Link
  Ls id     : 10.0.0.0
  Adv rtr   : 10.2.2.2
  Ls age    : 12
  Len       : 44
  Options   :  E
  seq#      : 80000001
  chksum    : 0x9579
        Opaque type : 3, Opaque ID : 0
        Grace LSA TLV information:
        Grace Period        : 1800
        GR reason           : 1
        IP address          : 10.1.1.2
```

Table 7-44 Description of the **display ospf lsdb opaque-link** command output

Item	Description
Opaque type	Opaque-link LSA.
Opaque ID	Opaque ID of an Opaque-link LSA (Link state ID in the LSA header consists of Opaque type and Opaque ID.).
Grace LSA TLV information:	GR information.
Grace Period	GR waiting period.
GR reason	Cause of GR: ● 0: unknown ● 1: software ● 2: upgrade ● 3: switchover
IP address	Address of the interface that performs GR on the switch.

Display information about Opaque-area LSAs in the LSDB.

```
<HUAWEI> display ospf 1 lsdb opaque-area
OSPF Process 1 with Router ID 10.1.1.1
Area: 0.0.0.0                                          Link State Database
  Type      : Opq-Area
  Ls id     : 10.0.0.1
  Adv rtr   : 10.1.1.1
  Ls age    : 639
  Len       : 200
  Options   :  E
  seq#      : 80000001
  chksum    : 0x2175
  Opaque Type: 1
  Opaque Id: 1
  Opaque lsa information:
```



```
00 02 00 b0 00 01 00 01 02 00 00 00 00 02 00 04
0a 01 01 01 00 03 00 04 0a 01 01 01 00 04 00 04
00 00 00 00 00 05 00 04 00 00 00 01 80 02 00 04
00 00 00 01 00 06 00 04 00 00 00 00 00 07 00 04
00 00 00 00 80 00 00 04 00 00 00 00 00 09 00 04
00 00 00 00 00 08 00 20 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 80 01 00 20 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 0a 00 09
00 00 00 00 00 00 00 00 00 00 00 00 00 0c 00 04
00 01 00 01
```

7.4.27 display ospf mesh-group

Function

The **display ospf mesh-group** command displays brief information about OSPF mesh groups.

Format

display ospf [*process-id*] **mesh-group** [**brief**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If no process ID is specified, brief information about mesh groups in all OSPF processes is displayed.	The value is an integer ranging from 1 to 65535.
brief	Displays brief information about mesh groups in each OSPF area.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

When concurrent links exist between a switch and its neighbor, run the **mesh-group enable** command to enable the mesh group function and identify the mesh group using the router ID of the switch. When receiving LSAs, the switch will select a primary link to flood LSAs. The link interface does not flood in reverse. This avoids repeated flooding, reducing link load and saving resources.

Running the **display ospf mesh-group** command allows you to check brief information about mesh groups in a specified process in an OSPF area.

Example

Display brief information about OSPF mesh groups.

```
<HUAWEI> display ospf mesh-group

   OSPF Process 1 with Router ID 10.1.1.1
            Mesh-Groups

Area 0.0.0.0
     Mesh-Group ID:10.2.2.2

   Interface              IP Address/Mask     Nbr State
   Vlanif100               10.11.10.202/24    Exchange
   Vlanif200               10.22.20.202/24    Loading
   Vlanif300               10.33.30.202/24    Full
   Count of Interface in this Mesh-Group: 1
   Count of Mesh-Group in this Area: 1
```

Table 7-45 Description of the **display ospf mesh-group** command output

Item	Description
Area	OSPF area.
Mesh-Group ID	Key ID of a mesh group, namely, the router ID of a neighbor.
Interface	Interface enabled with the mesh-group feature.
IP Address/Mask	IP address and mask of the interface enabled with the mesh-group feature.
Nbr State	Neighbor status on the interface.
Count of Interface in this Mesh-Group	Number of interfaces in the mesh group.
Count of Mesh-Group in this Area	Number of mesh groups in the area.

7.4.28 display ospf migp-routing

Function

The **display ospf migp-routing** command displays OSPF Multicast IGP (MIGP) routing information.

Format

display ospf [*process-id*] **migp-routing** [*ip-address* [*mask* | *mask-length*]] [**interface** *interface-type interface-number*] [**nexthop** *nexthop-address*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
ip-address	Specifies an IP address.	The value is in dotted decimal notation.
mask	Specifies a subnet mask.	The value is in dotted decimal notation.
mask-length	Specifies the mask length.	The value is an integer ranging from 0 to 32.
interface-type interface-number	Specifies the type and number of an interface.	-
nexthop *nexthop-address*	Displays the route with a specified next hop IP address.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

When both multicast and MPLS TE tunnels are configured on a network and the TE tunnels are configured with IGP Shortcut, the outbound interface that an IGP calculates for a route may be not a physical interface but a TE tunnel interface. Based on a unicast route to a multicast source address, a switch can send a Join message through a TE tunnel interface. In this case, devices spanned by the TE tunnel cannot detect the Join message so that they do not create any multicast forwarding entry.

To solve this problem, you can run the **local-mt enable** command to configure local multicast topology (MT). MT allows multicast routing entries to be created correctly.

If local MT is enabled and the outbound interface of the calculated route is a TE tunnel interface of IGP Shortcut type, the route management (RM) module creates a separate Multicast IGP (MIGP) routing table for the multicast protocol, calculates the physical outbound interface of the route, and adds the IP address of the physical outbound interface to the MIGP routing table.

Running the **display ospf migp-routing** command allows you to view the physical outbound interface of the route with the outbound interface of the TE tunnel interface. The physical outbound interface is OSPF routing information in the MIGP routing table.

Precautions

- Local MT supports only OSPF processes of public network instances.
- Local MT does not support forwarding adjacency (FA).

Example

Display OSPF MIGP routing information.

```
<HUAWEI> display ospf migp-routing
        OSPF Process 1 with Router ID 10.2.2.2
                MIGP Routing Tables

Routing for Network
Destination        Cost  Type       NextHop         AdvRouter       Area
192.168.3.0/24     4     Stub       10.0.1.1        10.5.5.5        0.0.0.0
10.0.3.0/24        3     Transit    10.0.1.1        10.5.5.5        0.0.0.0

Total Nets: 4
Intra Area: 4  Inter Area: 0  ASE: 0  NSSA: 0
```

Table 7-46 Description of the **display ospf migp-routing** command output

Item	Description
Destination	Destination IP address.
Cost	Cost of the route to the destination address.
Type	Router types are as follows: • Inter-area: Indicates routes between areas. • Stub: indicates routes within an area, including routes advertised based on router LSAs, and direct routes to non-broadcast and non-NBMA networks. • Transit: indicates routes within an area, including routes advertised based on network LSAs.
NextHop	Next hop address to the destination address.
AdvRouter	Advertising router.
Area	Area ID.
Total Nets	Total number of networks in an area, between areas, in ASE areas, and in NSSAs.
Intra Area	Total number of intra-area routes (that is, stub routes and transit routes).
Inter Area	Total number of inter-area routes.
ASE	Total number of routes in the ASE area.

Item	Description
NSSA	Total number of routes in the NSSA.

Display OSPF MIGP routing information with the specified next hop address.

```
<HUAWEI> display ospf migp-routing nexthop 10.0.1.1
        OSPF Process 1 with Router ID 10.2.2.2

 Destination : 192.168.3.0/24
 AdverRouter : 10.5.5.5              Area      : 0.0.0.0
 Cost        : 4                     Type      : Stub
 NextHop     : 10.0.1.1              Interface : Vlanif10
 Priority    : Low

 Destination : 10.4.4.4/32
 AdverRouter : 10.4.4.4              Area      : 0.0.0.0
 Cost        : 3                     Type      : Stub
 NextHop     : 10.0.1.1              Interface : Vlanif20
 Priority    : Medium
```

Table 7-47 Description of the **display ospf migp-routing nexthop** command output

Item	Description
Priority	OSPF convergence priorities: • Critical: indicates that the convergence priority of OSPF routes is critical. • High: indicates that the convergence priority of OSPF routes is high. • Medium: indicates that the convergence priority of OSPF routes is medium. • Low: indicates that the convergence priority of OSPF routes is low.

7.4.29 display ospf nexthop

Function

The **display ospf nexthop** command displays OSPF next hop information.

Format

display ospf [*process-id*] **nexthop**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command output can display information about all the OSPF next hops, which helps you troubleshoot OSPF faults.

Example

Display OSPF next hop information.

```
<HUAWEI> display ospf nexthop
            OSPF Process 100 with Router ID 10.0.0.1
                 Routing Nexthop information
Next hops:
Address         Type      Refcount  IntfAddr       Intf Name
------------------------------------------------------------------------------
10.0.0.1        Local     3         10.0.0.1       Vlanif100
10.0.0.2        Local     5         10.0.0.1       Vlanif100
```

Table 7-48 Description of the **display ospf nexthop** command output

Item	Description
Next hops	Detailed information about the next hop.
Address	Address of the next hop.
Type	Type of the route passing through the next hop. Local indicates that the route is destined for the local network segment.
Refcount	Number of OSPF routes that use the next hop.
IntfAddr	IP address of the interface.
Intf Name	Name of the interface.

7.4.30 display ospf peer

Function

The **display ospf peer** command displays information about neighbors in each OSPF area.

Format

display ospf [*process-id*] **peer** [[*interface-type interface-number*] [*neighbor-id*] | **brief** | **last-nbr-down**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies the interface type and the interface number.	-
neighbor-id	Specifies the neighbor's router ID.	It is in dotted decimal notation.
brief	Displays brief information about neighbors in each OSPF area.	-
last-nbr-down	Displays brief information about the last neighbor that went Down in the OSPF area.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command output can display information about OSPF neighbors, and help you troubleshoot OSPF faults, verify the configurations of OSPF neighbors, and check whether the neighbor performs Graceful Restart (GR).

Example

Display information about the OSPF neighbor.


```
<HUAWEI> display ospf peer
        OSPF Process 1 with Router ID 10.1.1.2
                Neighbors

 Area 0.0.0.0 interface 10.1.1.2(Vlanif100)'s neighbors
 Router ID: 10.1.1.1          Address: 10.1.1.1          GR State: Normal
   State: Full  Mode:Nbr is Slave  Priority: 1
   DR: 10.1.1.2  BDR: 10.1.1.1    MTU: 0
   Dead timer due in 35  sec
   Retrans timer interval: 5
   Neighbor is up for 00:00:05
   Authentication Sequence: [ 0 ]
```

Table 7-49 Description of the **display ospf peer** command output

Item	Description
Area	Area to which the neighbor belongs.
interface	Interface that connects to the neighbor.
Router ID	Router ID of the neighbor.
Address	Address of the neighboring interface.
GR State	GR status after OSPF GR is enabled: • Normal: indicates that a switch is in the Normal state and does not perform GR. • Doing GR: indicates that a switch is performing GR. • Complete GR: indicates that a switch finishes GR. • Helper: indicates that the neighbor is the Helper when a switch is performing GR.

Item	Description
State	Neighbor status: • Down: It is the initial status of the neighbor, indicating that the neighbor does not receive any information. On an NBMA network, when the neighbor is in the Down state, Hello packets can still be transmitted at the poll interval, which is longer than the Hello interval. • Attempt: It exists only on an NBMA network, indicating that two ends are attempting to establish the neighbor relationship. The interval for sending Hello packets is the Hello interval, which is shorter than the poll interval. • Init: It indicates that the Hello packet has been received from the neighbor. • 2-Way: It indicates that the Hello packet has been received from the neighbor, and the neighbor list of the Hello packet contains the local Router ID. That is, the two ends can interwork. • ExStart: It is the first step of establishing adjacencies. In this step, the master and slave relationship and DD sequence number are negotiated. • Exchange: It indicates that the LSDBs start to be synchronized. In this process, DD packets, LSR packets, and LSU packets are exchanged. • Loading: It indicates that the LSDBs are being synchronized. In this process, LSR packets and LSU packets are exchanged. • Full: It indicates that the LSDB of the neighbor is already synchronized, and the Full adjacency is established between both ends.
Mode	Master or slave in the process of exchanging DD packets: • Nbr is Master: indicates that the neighbor is the master and actively sends DD packets. • Nbr is Slave: indicates that the neighbor is the slave and cooperates with the master to send DD packets.
Priority	Priority of the neighboring device.
DR	Designated router.
BDR	Backup designated router.
MTU	MTU value of the neighboring interface.

Item	Description
Dead timer due in 35 sec	The dead timer due in 35 seconds.
Retrans timer interval	Interval for retransmitting LSAs, in seconds.
Neighbor is up for	Time during which the neighbor remains Up.
Authentication Sequence	Authentication sequence number.

Display brief information about OSPF neighbors.

```
<HUAWEI> display ospf 1 peer brief
        OSPF Process 1 with Router ID 10.10.10.1
                  Peer Statistic Information
 --------------------------------------------------------------------------
 Area Id          Interface                Neighbor id      State
 0.0.0.0          Vlanif10                 10.10.10.3       Full
 --------------------------------------------------------------------------
 Total Peer(s):     1
```

Table 7-50 Description of the **display ospf peer brief** command output

Item	Description
Area Id	Area to which the neighbor belongs.
Interface	Interface that connects to the neighbor.
Neighbor id	Router ID of the neighbor.
Total Peer(s)	Number of neighbors.

Display information about the OSPF neighbor that went Down for the last time.

```
<HUAWEI> display ospf 1 peer last-nbr-down
         OSPF Process 1 with Router ID 10.1.1.1

                     Last Down OSPF Peer

        Neighbor Ip Address : 10.2.1.2
        Neighbor Area   Id  : 0.0.0.0
        Neighbor Router Id  : 2.2.2.2
        Interface           : Vlanif100
        Immediate Reason    : Neighbor Down Due to Kill Neighbor
        Primary Reason      : Logical Interface State Change
        Down Time           : 2012-09-14 17:17:7
```

Table 7-51 Description of the **display ospf peer last-nbr-down** command output

Item	Description
Neighbor Ip Address	Address of the neighboring interface.

Item	Description
Neighbor Area Id	Area to which the neighbor belongs.
Neighbor Router Id	Router ID of the neighbor.
Interface	Interface that connects to the neighbor.
Immediate Reason	Immediate reason that the neighbor went Down: • Neighbor Down Due to Inactivity: indicates that the inactivity timer times out. • Neighbor Down Due to LL Down: indicates that the link is Down. For example, the interface went Down from Up or the IP address of the link is deleted. • Neighbor Down Due to Kill Neighbor: indicates that the kill neighbor event is generated on the neighbor state machine. • Neighbor Down Due to 1-Wayhello: indicates that the neighbor went Down because it receives a 1-way packet. • Received: indicates that the AdjOK? event is generated on this interface. • Neighbor Down Due to SequenceNum Mismatch: indicates that the SequenceNum Mismatch event is generated on the neighbor state machine. • Neighbor Down Due to BadLSreq: indicates that the BadLSreq event is generated on the neighbor state machine.

Item	Description
Primary Reason	Primary reason that the neighbor went Down: ● Hello Not Seen: indicates that no Hello packet is received. ● Interface Parameter Mismatch: indicates that the parameters set on both ends of the link do not match. ● Logical Interface State Change: indicates that the status of the logical interface changes. ● Physical Interface State Change: indicates that the status of the physical interface changes. ● OSPF Process Reset: indicates that the OSPF process restarts. ● Area reset: indicates that the area restarts because the area type changes. ● Area Option Mis-match: indicates that the area options of the interfaces on both ends of the link do not match. ● Vlink Peer Not Reachable: indicates that the neighbor on the virtual link is not reachable. ● Sham-Link Unreachable: indicates that the neighbor on the sham link is not reachable. ● Undo Network Command: indicates that the **network** command is deleted. ● Undo NBMA Peer: indicates that the neighbor configuration on the NBMA interface is deleted. ● Passive Interface Down: indicates that the neighbor relationship went Down because the **silent-interface** command is configured on the local interface. ● Opaque Capability Enabled: indicates that Opaque capability is enabled. ● Opaque Capability Disabled: indicates that Opaque capability is disabled. ● Virtual Interface State Change: indicates that the status of a virtual link interface changes. ● BFD Session Down: indicates that the BFD session went Down. ● Down Retransmission Limit Exceed: indicates that the number of retransmission times reaches the limit. ● 1-Wayhello Received: indicates that the device receives 1-way hello packets. ● Router State Change from DR or BDR to DROTHER: indicates that the interface state machine changes to DROTHER from DR or BDR.

Item	Description
	• Neighbor State Change from DR or BDR to DROTHER: indicates that the neighbor state machine changes to DROTHER from DR or BDR. • NSSA Area Configure Change: indicates that the configuration of the NSSA area changes. • Stub Area Configure Change: indicates that the configuration of the Stub area changes. • Received Invalid DD Packet: indicates that invalid DD packets are received. • Not Received DD during RouterDeadInterval: indicates that no DD packet is received during the time when the Dead timer starts. • M,I,MS bit or SequenceNum Incorrect: indicates that the M, I, and MS bits do not comply with specifications in the protocol. • Unable Opaque Capability,Find 9,10,11 Type Lsa: indicates that Type9, Type10, and Type11 LSAs are received and Opaque capability is disabled. • Not NSSA,Find 7 Type Lsa in Summary List: indicates that this area is not an NSSA area and Type7 LSAs are found in the summary table. • LSrequest Packet,Unknown Reason: indicates that LSR packets are received with the reason unknown. • NSSA or STUB Area,Find 5 ,11 Type Lsa: indicates that this area is an NSSA or Stub area and Type5 and Type11 LSAs are found. • LSrequest Packet,Request Lsa is Not in the Lsdb: indicates that the neighbor sends an LSR to this process or area to request an LSA and this LSA does not exist in the LSDB of this process. • LSrequest Packet, exist same lsa in the Lsdb: indicates that this process receives an LSA that is same as that in the LSDB and the LSA is found in the request list of the neighbor. • LSrequest Packet, exist newer lsa in the Lsdb: indicates that this process receives a new LSA that exists in the local LSDB and the LSA is found in the request list of the neighbor. • Neighbor state was not full when LSDB overflow: indicates that the LSDB overflows and the neighbor state machine is not Full. • Filter LSA configuration change: indicates that the configuration of LSA filter changes. • ACL changed for Filter LSA: indicates that the ACL configuration of LSA filter changes.

Item	Description
	• Reset Ospf Peer: indicates that the OSPF neighbor is restarted.
Down Time	Time when the neighbor went Down.

7.4.31 display ospf request-queue

Function

The **display ospf request-queue** command displays the OSPF request list.

Format

display ospf [*process-id*] **request-queue** [*interface-type interface-number*] [*neighbor-id*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies the interface type and number.	-
neighbor-id	Specifies the neighbor's router ID.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command output helps you troubleshoot OSPF faults.

Example

Display the OSPF request list.

```
<HUAWEI> display ospf request-queue
 OSPF Process 1 with Router ID 10.1.1.1
 OSPF Request List
  The Router's Neighbor is Router ID 10.4.4.4         Address 172.16.4.2
  Interface 172.16.4.1        Area 0.0.0.2
  Request list:
```



```
    Type       LinkState ID     AdvRouter        Sequence   Age
    Router     10.1.1.1         10.1.1.1         8000001b   677
```

Table 7-52 Description of the **display ospf request-queue** command output

Item	Description
The Router's Neighbor is Router ID	Router ID of the neighbor.
Address	IP address of the neighboring interface.
Interface	IP address of the interface.
Area	Area to which the local device belongs.
Request list	Request list.
Type	LSA type: Router LSA, network LSA, network summary LSA, ASBR summary LSA, AS external LSA, NSSA LSA, and opaque LSA
LinkState ID	Link state ID in the LSA header.
AdvRouter	Advertising router in the LSA header.
Sequence	Sequence number in the LSA header.
Age	Aging time in the LSA header.

7.4.32 display ospf retrans-queue

Function

The **display ospf retrans-queue** command displays the OSPF retransmission list.

Format

display ospf [*process-id*] **retrans-queue** [*interface-type interface-number*] [*neighbor-id*] [**low-level-of-retrans-times-range** *min-time*] [**high-level-of-retrans-times-range** *max-time*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies the interface type and number.	-
neighbor-id	Specifies the neighbor's router ID.	It is in dotted decimal notation.

Parameter	Description	Value
low-level-of-retrans-times-range *min-time*	Specifies the minimum number of allowed LSA retransmission.	The value is an integer ranging from 0 to 65535.
high-level-of-retrans-times-range *max-time*	Specifies the maximum number of allowed LSA retransmission.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command output helps you troubleshoot OSPF faults.

Example

Display the OSPF retransmission list.

```
<HUAWEI> display ospf retrans-queue

        OSPF Process 1 with Router ID 10.1.1.1
              OSPF Retransmit List

 The Router's Neighbor is Router ID 10.2.2.2  Address 192.168.0.2
 Interface 192.168.0.1      Area 0.0.0.0
 Retransmit list:
     Type       LinkState ID      AdvRouter        Sequence   Age
     Router     10.1.1.1          10.1.1.1         80000002   533
```

Table 7-53 Description of the **display ospf retrans-queue** command output

Item	Description
The Router's Neighbor	Basic information about the neighboring switch.
Router ID	Router ID of the neighbor.
Address	IP address of the neighboring interface.
Interface	IP address of the interface.
Area	Area ID.
Retransmit List	Retransmission list.

Item	Description
Type	LSA type: Router LSA, network LSA, network summary LSA, ASBR summary LSA, AS external LSA, NSSA LSA, and opaque LSA
LinkState ID	Link state ID in the LSA header.
AdvRouter	Advertising router in the LSA header.
Sequence	Sequence number in the LSA header.
Age	Aging time in the LSA header.

7.4.33 display ospf routing

Function

The **display ospf routing** command displays an OSPF routing table.

Format

display ospf [*process-id*] **routing router-id** [*router-id*]

display ospf [*process-id*] **routing** [*ip-address* [*mask* | *mask-length*]] [**interface** *interface-type interface-number*] [**nexthop** *nexthop-address*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
router-id *router-id*	Specifies the router ID of a destination router.	The value is in dotted decimal notation.
ip-address	Specifies an IP address.	The value is in dotted decimal notation.
mask	Specifies a subnet mask.	The value is in dotted decimal notation.
mask-length	Specifies the mask length.	The value is an integer ranging from 0 to 32.

Parameter	Description	Value
interface *interface-type interface-number*	Specifies the type and number of an interface.	-
nexthop *nexthop-address*	Displays the route with a specified next hop IP address.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

By using this command with different parameters, you can view the routes of a specified interface or next hop.

The command output helps you troubleshoot OSPF faults.

Example

Display the OSPF routing table.

```
<HUAWEI> display ospf routing
        OSPF Process 1 with Router ID 10.2.2.9
              Routing Tables

 Routing for Network
 Destination        Cost  Type       NextHop         AdvRouter       Area
 10.12.12.0/24      1     Transit    10.12.12.10     10.2.2.9        0.0.0.1
 10.13.13.0/24      1     Stub       10.13.13.1      10.2.2.9        0.0.0.0
 10.11.11.0/24      2     Transit    10.12.12.11     10.0.0.1        0.0.0.1

 Routing for ASEs
 Destination        Cost      Type       Tag      NextHop         AdvRouter
 10.0.0.0/8         1         Type2      1        10.12.12.11     10.0.0.1

 Total Nets: 4
 Intra Area: 3  Inter Area: 0  ASE: 1  NSSA: 0
```

Table 7-54 Description of the **display ospf routing** command output

Item	Description
Destination	Destination network.
Cost	Cost of the route to the destination address.

Item	Description
Type	• Type 1 external route: When the cost of external routes equals that of AS internal routes, and can be compared with the cost of OSPF routes, these external routes have a high reliability and can be configured as Type 1 external routes. • Type 2 external route: When the cost of the routes from an ASBR to the destination outside an AS is much greater than the cost of the internal routes to the ASBR, these external routes have a low reliability and can be configured as Type 2 external routes. Type of the destination network: • Inter-area: indicates inter-area routes. • Intra-area: indicates intra-area routes. – Stub: indicates the routes advertised by router LSAs, which correspond to the direct routes of non-broadcast and non-NBMA networks. – Transit: indicates the routes advertised by network LSAs.
NextHop	Next hop address to the destination address.
AdvRouter	Device that advertises LSAs.
Area	Area ID.
Tag	Tag of the external route.
Total Nets	Total number of networks in an area, between areas, in ASE areas, and in NSSAs.
Intra Area	Total number of intra-area networks (that is, stub networks and transit networks).
Inter Area	Total number of inter-area networks.
ASE	Total number of networks in the ASE area.
NSSA	Total number of networks in the NSSA.

7.4.34 display ospf sham-link

Function

The **display ospf sham-link** command displays the sham links of an OSPF area.

If no OSPF process ID or area ID is specified, all sham links are displayed.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

display ospf [*process-id*] **sham-link** [**area** *area-id*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the ID of an OSPF area.	The value is an integer ranging from 0 to 4294967295 or in the IPv4 address format.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospf sham-link** command to check information about the sham links of an OSPF area. The information helps determine the establishment of sham links.

Example

Display all OSPF sham links.

```
<HUAWEI> display ospf sham-link
        OSPF Process 100 with Router ID 10.1.1.2
 Sham Link:
 Area       NeighborId     Source-IP      Destination-IP   State   Cost
 0.0.0.1    10.1.1.2       10.3.3.3       10.5.5.5         P-2-P   10
```

Display the OSPF sham link of Area 1.


```
<HUAWEI> display ospf sham-link area 1
        OSPF Process 100 with Router ID 10.1.1.2
 Sham-Link: 10.3.3.3 --> 10.5.5.5
 Neighbor ID: 10.1.1.2,    State: Full
 Area: 0.0.0.1
 Cost: 10  State: P-2-P, Type: Sham
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
```

Table 7-55 Description of the **display ospf sham-link** command output

Item	Description
Area	OSPF area that the sham link belongs to.
NeighborId	Neighbor ID of the switch.
Source-IP	Source IP address of the sham link.
Destination-IP	Destination IP address of the sham link.
State	Interface status of the sham link. P-2-P indicates the point-to-point link.
Cost	Cost of the sham link.
Type	Connection type.
Timers	Information about the following items: the interval for sending Hello messages, Dead time, retransmission interval, and transmission delay on the interface.

7.4.35 display ospf spf-statistics

Function

The **display ospf spf-statistics** command displays route calculation statistics in OSPF processes.

Format

display ospf [*process-id*] **spf-statistics** [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If no OSPF process ID is specified, brief information about route calculation statistics in all processes is displayed.	The value is an integer ranging from 1 to 65535.

Parameter	Description	Value
verbose	Displays detailed information about route calculation statistics.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display ospf spf-statistics** command displays route calculation statistics in OSPF processes, including the time when route calculation occurs, cause of route calculation, and number of changed routes.

When identifying the cause of OSPF route flapping, you can run this command to obtain OSPF route calculation statistics, and then identify the cause according to the command output.

Example

Display brief information about route calculation statistics in OSPF process 1.

```
<HUAWEI> display ospf 1 spf-statistics

          OSPF Process 1 with Router ID 10.2.2.2
Routing table change statistics:
Date         Time          Intra      Inter    External    Reason
2008-08-14   10:17:16       17         17       17          LSA
2008-08-14   09:16:47       77         62       127         Other
2008-08-14   08:16:37       0          0        0          LSA
2008-08-14   07:04:40       24         230      108         LSA
2008-08-14   06:03:15       204        230      18          Other
2008-08-14   05:02:55       34         236      128         LSA
2008-08-14   04:01:49       54         130      158         LSA
2008-08-14   03:01:48       44         220      138         LSA
2008-08-14   02:01:43       22         233      158         LSA
2008-08-14   01:00:53       977        897      907         LSA
```

Table 7-56 Description of the **display ospf spf-statistics** command output

Item	Description
Date	Date when route calculation occurs.
Time	Time when route calculation occurs.
Intra	Number of intra-area routes in the routing table, which are changed because of route calculation.
Inter	Number of inter-area routes in the routing table, which are changed because of route calculation.

Item	Description
External	Number of external routes in the routing table, which are changed because of route calculation.
Reason	Cause of route calculation: • LSA: indicates that route calculation is caused by LSAs. • Other: indicates that route calculation is caused by other causes. For example, the configuration changes; or the interface goes Down.

Display detailed information about route calculation statistics in OSPF process 1.

```
<HUAWEI> display ospf 1 spf-statistics verbose

        OSPF Process 1 with Router ID 10.10.10.2
Routing table change statistics:
Index: 1
        Time     : 2008-11-29,17:36:59
        Intra    : 0 Added,0 Deleted, 0 Modified
        Inter    : 0 Added,0 Deleted, 0 Modified
        External : 10 Added,0 Deleted, 0 Modified
        The reason of calculation is:LSA
        NO.     Type          LS ID           Adv Router
        1       External      10.1.5.0        10.10.10.1
        2       External      10.1.3.0        10.10.10.1
        3       External      10.1.9.0        10.10.10.1
        4       External      10.1.4.0        10.10.10.1
        5       External      10.1.2.0        10.10.10.1
        6       External      10.1.8.0        10.10.10.1
        7       External      10.1.7.0        10.10.10.1
        8       External      10.1.6.0        10.10.10.1
        9       External      10.1.10.0       10.10.10.1
        10      External      10.1.1.0        10.10.10.1
```

Table 7-57 Description of the **display ospf spf-statistics verbose** command output

Item	Description
Time	Date and time when route calculation occurs.
Intra	Number of intra-area routes in the routing table, which are added and deleted because of route calculation.
Inter	Number of inter-area routes in the routing table, which are added and deleted because of route calculation.
External	Number of external routes in the routing table, which are added and deleted because of route calculation.

Item	Description
The reason of calculation is	Cause of route calculation: • LSA: indicates that route calculation is caused by LSAs. • Other: indicates that route calculation is caused by other causes. For example, the configuration changes; or the interface goes Down.
No.	Sequence number of the LSA that causes route calculation, which ranges from 1 to 10.
Type	Type of the LSA that causes route calculation, including Router, Network, Sum-Net, External, and NSSA.
LS ID	Link state ID of the LSA that causes route calculation.
Adv Router	Router ID of the switch that generates the LSA that causes route calculation.

7.4.36 display ospf statistics updated-lsa

Function

The **display ospf statistics updated-lsa** command displays the frequent updates of the LSAs that the LSDB receives.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

display ospf [*process-id*] **statistics updated-lsa** [**originate-router** *advertising-router-id* | **history**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If no OSPF process ID is specified, this command displays brief information about the route calculation statistics in all processes.	The value is an integer ranging from 1 to 65535.
originate-router	Specifies the link status of the advertising switch.	-
advertising-router-id	Specifies the ID of the advertising switch.	The value is in dotted decimal notation
history	Specifies the change history of LSAs that the LSDB receives.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The **display ospf statistics updated-lsa** command displays the frequent updates of LSAs, facilitating the location of the cause for a routing flapping.

- If the **history** parameter is not configured, the **display ospf statistics updated-lsa** command only displays the updates of LSAs within the latest hour.
- If the **history** parameter is configured, the **display ospf statistics updated-lsa** command displays the change history of LSAs within the last 24 hours.

Precautions

- If you run the **reset ospf** command to restart the OSPF process, the real-time and historical statistics on the process will be cleared.
- The **display ospf statistics updated-lsa** command is used only to display the frequent updates of LSAs. The updated LSAs are compared with the LSAs in the local LSDB, and those with the **age** greater than 900 will not be displayed except those with the **age** of 3600.

Example

Display the LSA updates within the last one hour.

```
<HUAWEI> display ospf statistics updated-lsa
        OSPF Process 1 with Router ID 10.1.1.1
            Statistics of Received LSAs

 Begin time: 2011-04-25 11:37:32

 AdvRouter          Total       Updated at
 10.1.1.1           18          11:37:40/2011/04/25
 10.2.2.2           5           11:37:40/2011/04/25
 10.3.3.3           5           11:37:41/2011/04/25
 10.4.4.4           5           11:37:41/2011/04/25
 10.5.5.5           2           11:37:40/2011/04/25
 10.6.6.6           3           11:37:40/2011/04/25
 10.7.7.7           5           11:37:40/2011/04/25
 10.8.8.8           6           11:37:41/2011/04/25
```

Table 7-58 Description of the **display ospf statistics updated-lsa** command output

Item	Description
Begin time	Start time of collecting statistics.
AdvRouter	Advertising switch.
Total	Total update times of LSAs.
Updated at	Latest update time.

Display the LSA updates of the specified advertising switch.

```
<HUAWEI> display ospf statistics updated-lsa originate-router 10.1.1.1
        OSPF Process 1 with Router ID 10.2.2.2
            Statistics of Received LSAs

 Begin time: 2011-04-25 11:37:32

 AdvRouter      : 10.1.1.1
 Total          : 6              Updated at      : 2011-04-25 11:37:41
 Router(1)      : 3                Network(2)      : 2
 Summary-Net(3) : 0                Summary-Asbr(4) : 0
 External(5)    : 1                NSSA(7)         : 0
 Opaque-link(9) : 0                Opaque-area(10) : 0
 Opaque-AS(11)  : 0
```

Table 7-59 Description of the **display ospf statistics updated-lsa originate-router** command output

Item	Description
Router(1)	Update times of Router LSAs.
Network(2)	Update times of Network LSAs.
Summary-Net(3)	Update times of Network Summary LSAs.

Item	Description
Summary-Asbr(4)	Update times of ASBR Summary LSAs.
External(5)	Update times of AS External LSAs.
nssa(7)	Update times of Type7 LSAs.
Opaque-link(9)	Update times of Type9 LSAs.
Opaque-area(10)	Update times of Type10 LSAs.
Opaque-AS(11)	Update times of Type11 LSAs.

Display the change history of LSAs.

```
<HUAWEI> display ospf statistics updated-lsa history
        OSPF Process 1 with Router ID 10.1.1.1
        History Information for Received LSAs

Record  1:
Begin time: 2011-04-25 11:39:32
End   time: 2011-04-25 11:41:32

no Record

Record  2:
Begin time: 2011-04-25 11:37:32
End   time: 2011-04-25 11:39:32

AdvRouter      : 10.1.1.1         Total          : 18
Router(1)      : 0                Network(2)     : 0
Summary-Net(3) : 0                Summary-Asbr(4) : 18
External(5)    : 0                NSSA(7)        : 0
Opaque-link(9) : 0                Opaque-area(10) : 0
Opaque-AS(11)  : 0

AdvRouter      : 10.2.2.2         Total          : 5
Router(1)      : 3                Network(2)     : 2
Summary-Net(3) : 0                Summary-Asbr(4) : 0
External(5)    : 0                NSSA(7)        : 0
Opaque-link(9) : 0                Opaque-area(10) : 0
Opaque-AS(11)  : 0

AdvRouter      : 10.3.3.3         Total          : 5
Router(1)      : 3                Network(2)     : 2
Summary-Net(3) : 0                Summary-Asbr(4) : 0
External(5)    : 0                NSSA(7)        : 0
Opaque-link(9) : 0                Opaque-area(10) : 0
Opaque-AS(11)  : 0

AdvRouter      : 10.4.4.4         Total          : 5
Router(1)      : 2                Network(2)     : 2
Summary-Net(3) : 0                Summary-Asbr(4) : 0
External(5)    : 1                NSSA(7)        : 0
Opaque-link(9) : 0                Opaque-area(10) : 0
Opaque-AS(11)  : 0
```


Table 7-60 Description of the **display ospf statistics updated-lsa history** command output

Item	Description
Record	Record number.
End time	End time of collecting statistics.

7.4.37 display ospf vlink

Function

The **display ospf vlink** command displays OSPF virtual links.

Format

display ospf [*process-id*] **vlink**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command output helps you troubleshoot OSPF faults.

Example

Display OSPF virtual links.

```
<HUAWEI> display ospf vlink
        OSPF Process 1 with Router ID 10.1.1.1
                Virtual Links
 Virtual-link Neighbor-id  -> 10.2.2.2, Neighbor-State: Full
 Interface: 10.1.1.1 (Vlanif10)
 Cost: 1  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```


Table 7-61 Description of the **display ospf vlink** command output

Item	Description
Virtual-link Neighbor-id	ID of the neighboring switch that is connected through the virtual link.
Neighbor-State	Neighbor status, such as Down, Init, 2-Way, ExStart, Exchange, Loading, and Full.
Interface	Information about interfaces in the area, that is, IP address and name of the primary interface (If the interface is a serial interface, **Unknown** is displayed.)
Cost	Cost.
State	Interface status.
Type	Interface type.
Transit Area	Transit area ID if the current interface is a virtual link.
Timers	Information about the following items: the interval for sending Hello messages, Dead time, retransmission interval, and transmission delay on the interface.
GR State	GR status: • Normal: indicates that a switch is in the Normal state and does not perform GR. • Doing GR: indicates that a switch is performing GR. • Complete GR: indicates that a switch finishes GR. • Helper: indicates that the neighbor is the Helper when a switch is performing GR.

7.4.38 dn-bit-set

Function

The **dn-bit-set disable** command disables OSPF from setting the DN bit in LSAs.

The **undo dn-bit-set disable** command enables OSPF to set the DN bit in LSAs.

By default, OSPF is enabled to set the DN bit in LSAs.

Format

dn-bit-set disable { **summary** | **ase** | **nssa** }

undo dn-bit-set disable { **summary** | **ase** | **nssa** }

Parameters

Parameter	Description	Value
summary	Specifies that the DN bit is not set in summary LSAs.	-
ase	Specifies that the DN bit is not set in ASE LSAs.	-
nssa	Specifies that the DN bit is not set in NSSA LSAs.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **dn-bit-set disable** command can be used in the following scenarios:

- In the VPN option A scenario, the local PE imports BGP routes to generate LSAs and advertise the generated LSAs to the peer PE. According to RFC 4577, setting of the DN bit is restricted. The peer PE may fail to calculate a route. In this situation, the **dn-bit-set disable** command can be used to set the DN bit on or remove the setting of the DN bit from the local PE.
- When a PE is connected to an MCE, the MCE needs to calculate routes advertised by the PE. By default, the MCE does not check the DN bit. In this situation, the **dn-bit-set disable** command can be used to set the DN bit on or remove the setting of the DN bit from the local PE.

NOTE

To prevent routing loops, the OSPF multi-instance process uses a bit as a flag. The bit is called DN bit.

Configuration Impact

When the **dn-bit-set disable** command is used to disable OSPF from setting the DN bit in LSAs, routing loops may occur. If the parameter **ase** or **nssa** is specified, the DN bit in ASE LSAs or NSSA LSAs is not set. You can use the **route-tag** command to set the same tag value to prevent routing loops. Therefore, it is recommended that the **dn-bit-set disable** command be used in only the scenarios specified in **Usage Scenario**.

If the **dn-bit-set disable ase** command is configured, the DN bit is not set in type 5 LSAs that are converted from type 7 LSAs even if the DN bit is set in type 7 LSAs.

Precautions

The **dn-bit-set disable** command can be configured for only private OSPF processes. The configuration of this command takes effect only on the PEs.

The **dn-bit-check disable** command can be used to control whether OSPF running on the peer PE checks the DN bit when calculating routes.

Example

Disable OSPF from setting the DN bit in ASE LSAs.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance test
[HUAWEI-vpn-instance-test] route-distinguisher 100:1
[HUAWEI-vpn-instance-test-af-ipv4] quit
[HUAWEI-vpn-instance-test] quit
[HUAWEI] ospf 100 vpn-instance test
[HUAWEI-ospf-100] dn-bit-set disable ase
```

7.4.39 dn-bit-check

Function

The **dn-bit-check disable** command disables OSPF from checking the DN bit in LSAs.

The **undo dn-bit-check disable** command enables OSPF to check the DN bit in LSAs.

By default, OSPF is enabled to check the DN bit in LSAs.

Format

dn-bit-check disable { **summary** [**router-id** *router-id*] | **ase** | **nssa** }

undo dn-bit-check disable { **summary** [**router-id** *router-id*] | **ase** | **nssa** }

Parameters

Parameter	Description	Value
summary	Specifies that the DN bit in summary LSAs is not checked.	-
router-id *router-id*	Specifies the ID of a device on which the DN bit in summary LSAs is checked.	-
ase	Specifies that the DN bit in ASE LSAs is not checked.	-
nssa	Specifies that the DN bit in NSSA LSAs is not checked.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In the VPN option A scenario, the local PE imports BGP routes to generate LSAs and advertise the generated LSAs to the peer PE. According to RFC 4577, setting of the DN bit is restricted. The peer PE may fail to calculate a route. In this situation, you need to use the **dn-bit-check disable** command to disable OSPF from checking the DN bit in LSAs.

NOTE

To prevent routing loops, the OSPF multi-instance process uses a bit as a flag. The bit is called DN bit.

Configuration Impact

When the **dn-bit-check disable** command is run, routing loops may occur. If the parameter **ase** or **nssa** is specified, the DN bit in ASE LSAs or NSSA LSAs is not checked. You can use the **route-tag** command to set the same tag value to prevent routing loops. Therefore, run the **dn-bit-check disable** command only in the scenario specified in **Usage Scenario**.

Precautions

When a PE is connected to an MCE, the MCE does not check the DN bit by default.

The **dn-bit-check disable** command can be configured only for private OSPF processes. The configuration of this command takes effect only on the PEs.

In this scenario, you can run the **dn-bit-set disable** command to set the DN bit on or remove the setting of the DN bit from the local PE.

Example

Disable OSPF from checking the DN bit in summary LSAs.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance test
[HUAWEI-vpn-instance-test] route-distinguisher 100:1
[HUAWEI-vpn-instance-test-af-ipv4] quit
[HUAWEI-vpn-instance-test] quit
[HUAWEI] ospf 100 vpn-instance test
[HUAWEI-ospf-100] dn-bit-check disable summary router-id 10.1.1.1
```

7.4.40 domain-id (OSPF)

Function

The **domain-id** command sets an ID for an OSPF domain.

The **undo domain-id** command restores the default setting.

By default, the domain ID is null.

Format

domain-id { **null** | *domain-id* [**type** *type* **value** *value* | **secondary**] * }

undo domain-id [*domain-id* [**type** *type* **value** *value*]]

Parameters

Parameter	Description	Value
domain-id	Specifies the ID of an OSPF domain.	The value can be an integer or in dotted decimal notation. • If it is an integer, the value ranges from 0 to 4294967295, and it is converted to dotted decimal notation when the ID is displayed. • If it is in dotted decimal notation, it is displayed as entered.
null	Indicates that the OSPF domain ID is null.	-
type *type*	Specifies the type of the OSPF domain ID.	It can be 0005, 0105, 0205, or 8005. By default, it is 0005.
value *value*	Specifies the value of the type of the OSPF domain ID.	The value is a hexadecimal number that ranges from 0x0 to 0xffff, and the default value is 0x0.
secondary	Indicates the ID of a secondary domain.	The maximum number of **domain-id secondary** in each OSPF process is 1000.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Domain IDs are used to identify domains.

If the local OSPF area and an OSPF area of a remote VPN attempts to exchange Type 3 LSAs, the two areas must be in the same OSPF domain. You can run the **domain-id** command to configure the same domain ID for the two OSPF areas.

The routes that are imported from a PE switch are advertised using External-LSAs. The routes destined for different nodes in the same OSPF domain are advertised based on Type 3 LSAs. This requires that the nodes in the same OSPF domain be configured with the same domain ID.

If the **undo domain-id** command without any parameter is executed, the primary domain ID will be deleted.

NOTE

OSPF direct routes to the PE do not carry the domain ID, while BGP direct routes to the PE carry the domain ID.

Configuration Impact

Before sending routes to a remote CE switch, a PE switch sends Type 3 LSAs or Type 5 LSAs to the CE based on domain ID. If local domain IDs are the same as or compatible with remote domain IDs in BGP routes, the PE advertises Type 3 routes. If local domain IDs are different from or incompatible with remote domain IDs in BGP routes, the PE advertises Type 5 routes.

Precautions

- Each OSPF domain has one or multiple domain IDs. One of them is a primary ID and the others are secondary IDs.
- If an OSPF instance does not have a specific domain ID, its ID is considered as null.
- If the value of the domain ID is 0, **secondary** cannot be configured.
- The maximum number of **domain-id secondary** items configured in an OSPF process is 1000.
- The **domain-id** command is forbidden in public networks.

Example

Set an OSPF domain ID.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] domain-id 234
```

7.4.41 eca-route-type compatible

Function

The **eca-route-type compatible** command sets the route type of the extended community attribute of OSPF VPN to 0x8000.

The **undo eca-route-type compatible** command restores the route type of the extended community attribute of OSPF VPN to 0x0306.

By default, the route type of the extended community attribute of OSPF VPN is 0x0306.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

eca-route-type compatible

undo eca-route-type compatible

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **eca-route-type compatible** command is used in OSPF VPN scenarios.

- For the switch supporting RFC 4577, you can set the route type of the extended community attribute of OSPF VPN to 0x0306 and configure the switch to identify both 0x0306 and 0x8000 route types.
- For the switch that does not support RFC 4577, you can set the route type of the extended community attribute of OSPF VPN to 0x8000 and configure the switch to identify only the 0x8000 route type.

The **eca-route-type compatible** command enables different switches to communicate with each other and avoid the failure in parsing the route type because the route type of the extended community attribute of OSPF VPN is unrecognized.

Precautions

The **eca-route-type compatible** command is forbidden in public network.

Example

Set the route type of the extended community attribute of OSPF VPN to 0x8000.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance test
[HUAWEI-vpn-instance-test] route-distinguisher 100:1
[HUAWEI-vpn-instance-test-af-ipv4] quit
[HUAWEI-vpn-instance-test] quit
[HUAWEI] ospf 1 vpn-instance test
[HUAWEI-ospf-1] eca-route-type compatible
```

7.4.42 enable log

Function

The **enable log** command enables the logging function.

The **undo enable log** command disables the logging function.

By default, the logging function is disabled.

Format

enable log [**config** | **error** | **state** | **snmp-trap**]

undo enable log [**config** | **error** | **state** | **snmp-trap**]

Parameters

Parameter	Description	Value
config	Enables the configuration log.	-
state	Enables the state log.	-
error	Enables the error log.	-
snmp-trap	Enables the SNMP trap log.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Running the **enable log** command can enable the logging function. You can view running status of devices based on logs, facilitating device maintenance.

If the **undo enable log** command is executed, the logging function will be disabled. Then, running status of devices will be not displayed. This is inconvenient for network maintenance.

Precautions

Configuring different parameters in the **enable log** command can display different log information. If no parameter is not specified, the command output displays all log information.

Example

Enable the OSPF logging function.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] enable log
```

7.4.43 filter export (OSPF Area)

Function

The **filter export** command filters the outgoing Type3 LSAs of the local area.

The **undo filter export** command restores the default setting.

By default, the outgoing Type3 LSAs of the local area are not filtered.

Format

filter { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **export**

undo filter [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name*] **export**

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The command can be used to filter out invalid LSAs sent to neighbors to reduce the size of the LSDB and speed up network convergence.

Configuration Impact

After filtering conditions are set for the outgoing summary LSAs to be advertised using the **filter export** command, only the outgoing Type3 LSAs that pass the filtering can be advertised.

Precautions

- The command can be configured only on an ABR.
- For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule takes effective only when the source address range is specified by the **source** parameter and the time period is specified by the **time-range** parameter.
- Run the **filter import** command to set filtering conditions for the incoming Type3 LSAs to be advertised.
- Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, OSPF advertises all Type 3 LSAs.
- Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, OSPF advertises all Type 3 LSAs.

Example

Configure OSPF to filter outgoing Type 3 LSAs.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] area 1
[HUAWEI-ospf-1-area-0.0.0.1] filter 2000 export
```

7.4.44 filter import (OSPF Area)

Function

The **filter import** command filters the incoming Type3 LSAs of the local area.

The **undo filter import** command restores the default setting.

By default, the incoming Type3 LSAs are not filtered.

Format

filter { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **import**

undo filter [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name*] **import**

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the name of an ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After filtering conditions are set for the incoming summary LSAs to be advertised using the **filter import** command, only the incoming Type3 LSAs that pass the filtering can be received.

Configuration Impact

The command can be used to filter out invalid LSAs sent to neighbors to reduce the size of the LSDB and speed up network convergence.

Precautions

- The command can be configured only on an ABR.
- For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule takes effect only when the source address range is specified by the **source** parameter and the time period is specified by the **time-range** parameter.
- Run the **filter export** command to set filtering conditions for the outgoing Type3 LSAs to be advertised.
- Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, OSPF receives all Type 3 LSAs.
- Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, OSPF receives all Type 3 LSAs.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all the routes learned by the specified protocol are imported to the OSPF routing table.

Example

Configure an ABR to filter incoming Type 3 LSAs of the area where the ABR resides.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 1
[HUAWEI-ospf-100-area-0.0.0.1] filter ip-prefix my-prefix-list import
```

7.4.45 filter-lsa-out peer

Function

The **filter-lsa-out peer** command configures a switch to filter the LSAs that are sent by specified neighbors on a P2MP network.

The **undo filter-lsa-out peer** command cancels the configuration.

By default, the LSAs that are sent by specified neighbors on a P2MP network are not filtered.

Format

filter-lsa-out peer *ip-address* { **all** | { **summary** [**acl** { *acl-number* | *acl-name* }] | **ase** [**acl** { *acl-number* | *acl-name* }] | **nssa** [**acl** { *acl-number* | *acl-name* }] } * }

undo filter-lsa-out peer *ip-address*

Parameters

Parameter	Description	Value
ip-address	Specifies the IP address of a P2MP neighbor. This parameter is configured by the neighbor using the **ip address** command.	The value is in dotted decimal notation.
all	Filters all the outgoing LSAs except Grace LSAs.	-
summary	Filters the outgoing network summary LSAs (Type 3).	-
acl *acl-number*	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ase	Filters the outgoing AS external LSAs (Type 5).	-
nssa	Filters the outgoing NSSA LSAs (Type 7).	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a P2MP network, when multiple P2MP links exist between two devices, you can configure the local device to filter the outgoing LSAs on the specified link. This reduces unnecessary LSA retransmission attempts and saves bandwidth resources.

For an ACL, when the **rule** command is used to configure filtering rules, only the source address range that is specified by the **source** parameter and the period of time that is specified by the **time-range** parameter take effect.

Prerequisites

OSPF does not regard a network as a P2MP network by default regardless of any link layer protocol. A P2MP network is forcibly changed from the network of another type using the **ospf network-type p2mp** command.

Configuration Impact

This command is valid for all the interfaces of the OSPF process.

Follow-up Procedure

Configure a device to filter the outgoing LSAs on the specified OSPF interface by using the **ospf filter-lsa-out** command.

Example

On a P2MP network, configure a switch to filter all the LSAs (except Grace LSAs) sent to neighbor 10.1.1.1.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] filter-lsa-out peer 10.1.1.1 all
```

7.4.46 filter-policy export (OSPF)

Function

The **filter-policy export** command filters the imported routes when these routes are advertised based on a filtering policy.

The **undo filter-policy export** command restores the default setting.

By default, the imported routes to be advertised are not filtered.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **export** [*protocol* [*process-id*]]

undo filter-policy [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name*] **export** [*protocol* [*process-id*]]

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
protocol process-id	Filters imported routes of a specified protocol. The value can be **direct**, **rip**, **isis**, **bgp**, **ospf**, **unr**, or **static**. When the routing protocol is RIP, IS-IS, or OSPF, you can specify a process ID.	The value is an integer that ranges from 1 to 65535. The default value is 1.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After OSPF imports external routes using the **import-route** command, you can use the **filter-policy export** command to filter the imported routes to be advertised. Only the external routes that pass the filtering can be converted into AS-external LSAs and advertised.

The *protocol* or *process-id* parameter can be specified to determine a specified protocol or process. If the *protocol* or *process-id* parameter is not specified, OSPF filters all imported routes.

Precautions

- This command can be configured only on ASBRs because AS-external-LSAs are generated by ASBRs.
- For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule takes effective only when the source address range is specified

by the **source** parameter and the time period is specified by the **time-range** parameter.

- Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all external routes imported to OSPF are converted to Type 5 LSAs (AS-external-LSAs) or Type 7 LSAs (NSSA-external-LSAs) and then are advertised to neighbors.
- Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all external routes imported to OSPF are converted to Type 5 LSAs (AS-external-LSAs) or Type 7 LSAs (NSSA-external-LSAs) and then are advertised to neighbors.

Example

Filter the routes that are imported from RIP and advertised by OSPF based on a filtering policy.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] import-route rip
[HUAWEI-ospf-100] filter-policy 2000 export
```

7.4.47 filter-policy import (OSPF)

Function

The **filter-policy import** command configures a filtering policy to filter routes received by OSPF.

The **undo filter-policy import** command restores the default setting.

By default, OSPF does not filter received routes.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* [**secondary**] } **import**

undo filter-policy [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* [**secondary**]] **import**

Parameters

Parameter	Description	Value
acl-number	Specifies the basic ACL number.	The value is an integer ranging from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
ip-prefix *ip-prefix-name*	Specifies the name of an address prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of the route policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
secondary	Selects a secondary route.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **filter-policy import** command is used to set a filtering policy for received routes. Only the routes that pass the filtering can be added to the routing table. The routes that fail to pass the filtering cannot be added to the routing table but can be advertised.

The OSPF routing information is recorded in the LSDB. Instead of filtering the received or sent LSAs, the device filters routes calculated by OSPF using the **filter-policy import** command.

Precautions

For an ACL configured using the **acl** command, when the **rule** command is used to configure a filtering rule, the filtering rule takes effect only when the source address range is specified by the **source** parameter and the time period is specified by the **time-range** parameter.

Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all routes received by OSPF are delivered to the IP routing table.

Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all routes received by OSPF are delivered to the IP routing table.

Example

Configure OSPF to filter received routes.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] filter-policy 2000 import
```

7.4.48 flooding-control

Function

The **flooding-control** command restricts the flooding of updated LSAs.

The **undo flooding-control** command cancels the restriction on the flooding of updated LSAs.

By default, this function is enabled when the number of neighbors exceeds 256.

Format

flooding-control [**number** *transmit-number* | **timer-interval** *transmit-interval*] *

undo flooding-control [**number** | **timer-interval**] *

Parameters

Parameter	Description	Value
number *transmit-number*	Sets the number of updated LSAs to be flooded each time.	The value is an integer ranging from 1 to 1000. By default, the value is 50.
timer-interval *transmit-interval*	Sets the interval for flooding updated LSAs.	The value is an integer that ranges from 30 to 100000, in milliseconds. By default, the value is 30.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When multiple neighbors are configured or a large number of updated LSAs are flooded, a switch may receive a large number of updated LSAs in a short period. If the switch is busy processing these update packets and discards the Hello packets that maintain neighbor relationships, neighbor relationships may be interrupted. During the reestablishment of neighbor relationships, more packets need to be exchanged, which deteriorates the processing of packets.

To avoid the preceding problem, you can run the **flooding-control** command to restrict the flooding of updated LSAs to keep stable neighbor relationships.

Configuration Impact

After the **flooding-control** command is run, the flooding of updated LSAs is immediately restricted.

Precautions

By default, a switch spends 50 ms in flooding updated LSAs each time. If not all the updated LSAs are flooded within 50 ms, the switch floods the remaining LSAs after the time specified by *transmit-interval*.

Example

Set the number of updated LSAs to be flooded each time to 100.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] flooding-control number 100
```

7.4.49 frr (OSPF)

Function

The **frr** command creates and then displays the OSPF FRR view.

The **undo frr** command deletes the OSPF FRR view.

By default, the OSPF FRR view does not exist.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

frr

undo frr

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

OSPF IP FRR allows devices to fast switch traffic from faulty links to back up links without interrupting traffic. This function protects traffic and greatly improves the reliability of OSPF networks. OSPF IP FRR must be configured in the OSPF FRR view. The **frr** command run in the OSPF view creates and displays the OSPF FRR view.

Follow-up Procedure

The **frr** command run in the OSPF view creates and displays the OSPF FRR view only, but cannot enable the OSPF IP FRR function. Run the **loop-free-alternate** command in the OSPF view to enable OSPF IP FRR to create the loop-free backup route.

Example

Create and display the OSPF FRR view.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] frr
[HUAWEI-ospf-1-frr]
```

7.4.50 frr-policy route (OSPF)

Function

The **frr-policy route** command configures a filtering policy for the OSPF IP FRR backup routes. The filtering policy determines what kind of OSPF backup route can be added to the routing table.

The **undo frr-policy route** command cancels the filtering function.

By default, the filtering function is disabled.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

frr-policy route route-policy *route-policy-name*

undo frr-policy route

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of the policy used to filter OSPF backup routes.	The value must be an existing route-policy.

Views

OSPF FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

OSPF IP FRR allows devices to fast switch traffic on fault links to back up links without interrupting traffic. This protects traffic and greatly improves the reliability of OSPF networks.

After the filtering policy is configured using the **frr-policy route** command, only the OSPF backup route that satisfies filtering rules can be delivered to the forwarding table.

Configuration Impact

To protect the traffic over a specific OSPF route, you can configure a filtering policy *route-policy-name* that the OSPF route matches to ensure that the backup route can be added to the forwarding table. When this route fails, OSPF can fast switch the traffic to a backup route.

Precautions

The **frr-policy route** command is cyclic in nature, and only the latest configuration takes effect.

If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, OSPF adds all backup routes to the IP routing table.

Example

Configure OSPF to add the OSPF backup routes that match the named ACL **abc** to the IP routing table.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] frr
[HUAWEI-ospf-1-frr] loop-free-alternate
[HUAWEI-ospf-1-frr] frr-policy route route-policy abc
```

7.4.51 frr-priority static low

Function

The **frr-priority static low** command enables dynamic backup links to take preference over static backup links so that the LFA algorithm is used to calculate the nexthop and outbound interface.

The **undo frr-priority static** command disables this function.

By default, this function is disabled, static backup links take preference over dynamic backup links during route selection.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

frr-priority static low

undo frr-priority static

Parameters

None

Views

OSPF FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The nexthop and outbound interface of an OSPF loop-free backup link can be obtained using either of the following methods:

- For a static backup link, after IP FRR is enabled using the **ip frr** command in the system view or VPN instance view, configure a nexthop and an outbound interface for the static backup link.
- For a dynamic backup link, after OSPF IP FRR is enabled using the **loop-free-alternate** command, enable the device to use the LFA algorithm to calculate the nexthop and outbound interface for the dynamic backup link.

By default, static backup links take preference over dynamic backup links during route selection. However, static backup links are less flexible than dynamic backup links. If a link failure occurs, static backup links cannot update automatically, but dynamic backup links can. Therefore, to ensure automatic link updates, run the **frr-priority static low** command to enable dynamic backup links to take preference over static backup links so that the LFA algorithm is used to calculate the nexthop and outbound interface.

Prerequisites

The OSPF IP FRR view has been displayed using the **frr** command.

Example

Enable the device to use the LFA algorithm to calculate the backup nexthop and outbound interface.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] frr
[HUAWEI-ospf-1-frr] frr-priority static low
```

7.4.52 graceful-restart (OSPF)

Function

The **graceful-restart** command enables the GR function.

The **undo graceful-restart** command disables the GR function.

By default, OSPF GR is disabled.

Format

graceful-restart [**period** *period* | **planned-only** | **partial**] *

undo graceful-restart [**period** | **planned-only** | **partial**] *

Parameters

Parameter	Description	Value
period *period*	Specifies the duration of GR.	It is an integer ranging from 1s to 1800s. The default value is 120s.
planned-only	Indicates that the switch supports only the planned GR. By default, the switch supports both the planned GR and unplanned GR.	-
partial	Indicates that the switch partially supports the GR. By default, the switch totally supports the GR.	-

NOTE

Planned GR: indicates that a device manually restarts or performs a master/slave device switchover by using the command. Before the device restarts or performs a master/slave device switchover, Restarter will send a grace LSA.

Unplanned GR: indicates that a device restarts or performs a master/slave device switchover because of faults. A device directly performs a master/slave device switchover without sending a grace LSA, and then enters GR after the slave device goes Up. It is different from the planned GR.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To prevent traffic interruption caused by route flapping after an active/standby switchover, you can run the **graceful-restart** command to enable OSPF GR. After the **graceful-restart** command is run to enable GR for a switch, the Helper function is also enabled.

Prerequisites

Opaque LSAs provide a generic mechanism for OSPF extension:

- OSPF supports GR using Type 9 LSAs.
- OSPF supports TE using Type 10 LSAs.

Before configuring OSPF GR, you must enable opaque LSA capability running the **opaque-capability enable** command.

Configuration Impact

After an OSPF process is restarted using GR, the Restarter switch and the Helper switch reestablish the neighbor relationship, exchange routing information, synchronize the LSDB, and update the routing table and forwarding table. This implements OSPF fast convergence, prevents traffic interruption, and stabilizes the network topology.

Precautions

If there are special requirements on the GR Helper, run the **graceful-restart helper-role** command to configure the requirements.

You are advised not to enable OSPF GR if BFD is enabled for a stack of fixed switches. This is because the time taken for smooth data switching during a master/standby switchover in the stack may be longer than the BFD timeout interval, in which situation BFD sessions go Down and OSPF GR fails.

Example

Enable OSPF GR and set the GR period to 200s.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] opaque-capability enable
[HUAWEI-ospf-1] graceful-restart period 200
```

7.4.53 graceful-restart helper-role (OSPF)

Function

The **graceful-restart helper-role** command configures a device as a GR helper.

The **undo graceful-restart helper-role** command cancels the configuration.

By default, the device does not function as a GR helper.

Format

graceful-restart [**period** *period* | **partial** | **planned-only**] * **helper-role** { [{ **ip-prefix** *ip-prefix-name* | **acl-number** *acl-number* | **acl-name** *acl-name* } | **ignore-external-lsa** | **planned-only**] * | **never** }

undo graceful-restart [**period** | **partial** | **planned-only**] * **helper-role** [[{ **ip-prefix** | **acl-number** | **acl-name** } | **ignore-external-lsa** | **planned-only**] * | **never**]

Parameters

Parameter	Description	Value
period *period*	Specifies the duration of GR.	The value is an integer that ranges from 1 to 1800, in seconds. The default value is 120.
planned-only	Configures the device to support only planned GR. By default, the device supports both the planned GR and unplanned GR.	-
partial	Configures the device to support partial GR. By default, the device supports totally GR.	-
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list. The name is a string.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
acl-number *acl-number*	Specifies the basic ACL number.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ignore-external-lsa	Indicates that the device does not check Type 5 and Type 7 LSAs.	-
planned-only	Indicates that the device supports only planned GR.	By default, the device supports both planned GR and unplanned GR.
never	Indicates that the device does not support the Helper mode.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If an OSPF neighbor supports GR and can be a restarter, the current device needs to be set to helper mode, and the neighbor is assisted to complete the GR process.

After the **graceful-restart** command is run to enable GR for a switch, the Helper function is also enabled.

Prerequisites

OSPF supports GR using Type 9 LSAs. So before configuring GR, run the **opaque-capability enable** command to enable opaque LSA capability.

Example

Configure a device as an OSPF GR helper and configure the helper to support only planned GR.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] graceful-restart helper-role planned-only
```

7.4.54 gtsm default-action

Function

The **gtsm default-action** command sets the default action that is performed on the packets that do not match the GTSM policies.

The **undo gtsm default-action drop** command restores the default setting.

By default, the packets that do not match the GTSM policies can pass the filtering.

Format

gtsm default-action { **drop** | **pass** }

undo gtsm default-action drop

Parameters

Parameter	Description	Value
drop	Indicates that the packets that do not match the GTSM policies cannot pass the filtering. The packets are dropped.	-
pass	Indicates that the packets that do not match the GTSM policies can pass the filtering.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

For a network demanding high security, you can configure Generalized TTL Security Mechanism (GTSM) to improve the security of the OSPF network. GTSM defends against attacks by checking the Time-to-Live (TTL) value. If an attacker simulates real OSPF packets and keeps sending them to a switch, a switch receives the packets and directly sends them to the main control board for OSPF processing, without checking the validity of the packets. In this case, the switch is busy in processing these packets, causing high usage of the CPU. GTSM function protects the switch by checking whether the TTL value in the IP packet header is in a pre-defined range to improve the system security.

GTSM only checks the TTL values of the packets that match the GTSM policy. The packets that do not match the GTSM policy can pass the filtering using the **undo gtsm default-action drop** command or using the **gtsm default-action** command to set the **pass** parameter, or be dropped after the **gtsm default-action** command is run to set the **drop** parameter.

Configuration Impact

If the default action to be taken on GTSM packets is drop, the connection cannot be established by the switch. Therefore, GTSM improves security but reduces the ease of use.

Precautions

You can enable the log function by using the **gtsm log drop-packet** command to record the information about dropped packets for further fault location.

If you configure the default action by using the **gtsm default-action** command but not configure GTSM policy (the **drop** or **pass** parameter), GTSM does not take effect.

Example

Set the default action performed on the packets that do not match the GTSM policies to **drop**.

```
<HUAWEI> system-view
[HUAWEI] gtsm default-action drop
```

Set the default action performed on the packets that do not match the GTSM policy to pass the filtering.

```
<HUAWEI> system-view
[HUAWEI] undo gtsm default-action drop
```


7.4.55 gtsm log drop-packet all

Function

The **gtsm log drop-packet** command configures a GTSM-capable switch to record logs when it drops packets.

The **undo gtsm log drop-packet** command configures a GTSM-capable switch not to record logs when it drops packets.

By default, a GTSM-capable switch does not record logs when dropping packets.

Format

gtsm log drop-packet all

undo gtsm log drop-packet all

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

For a network demanding higher security, you can configure Generalized TTL Security Mechanism (GTSM) to improve the security of the OSPF network. GTSM defends against attacks by checking the Time-to-Live (TTL) value. If an attacker simulates real OSPF packets and keeps sending them to a device, an interface board on the device receives the packets and directly sends them to the main control board for OSPF processing, without checking the validity of the packets. In this case, the device is busy in processing these packets, causing high usage of the CPU. GTSM function protects the device by checking whether the TTL value in the IP packet header is in a pre-defined range to improve the system security.

GTSM only checks the TTL values of the packets that match the GTSM policy. The packets that do not match the GTSM policy can be allowed or dropped by using the **gtsm default-action** command.

You can also enable the logging function by using the **gtsm log drop-packet** command to record the information about dropped packets for further fault location.

Prerequisites

The **gtsm default-action drop** command has been run.

Example

Enable all GTSM-capable boards to record logs when they drop packets.

```
<HUAWEI> system-view
[HUAWEI] gtsm default-action drop
[HUAWEI] gtsm log drop-packet all
```

7.4.56 import-route (OSPF)

Function

The **import-route** command imports routes learned by other protocols.

The **undo import-route** command cancels the configuration.

By default, routes learned by other protocols are not imported.

Format

import-route { **limit** *limit-number* | { **bgp** [**permit-ibgp**] | **direct** | **unr** | **rip** [*process-id-rip*] | **static** | **isis** [*process-id-isis*] | **ospf** [*process-id-ospf*] } [**cost** *cost* | **type** *type* | **tag** *tag* | **route-policy** *route-policy-name*] * }

undo import-route { **limit** | **bgp** | **direct** | **unr** | **rip** [*process-id-rip*] | **static** | **isis** [*process-id-isis*] | **ospf** [*process-id-ospf*] }

NOTE

Only the S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **bgp** [**permit-ibgp**] and **isis** [*process-id-isis*] parameter.

Parameters

Parameter	Description	Value
bgp	Specifies the BGP protocol whose routes are imported. **NOTE** If this parameter is specified, only EBGP routes can be imported; IBGP routes cannot be imported.	-
permit-ibgp	IBGP routes that are permitted to be imported. **NOTE** The import of IBGP routes may cause route loops. Therefore, this command must not be configured unless it is necessary.	-

Parameter	Description	Value
direct	Specifies the direct protocol whose routes are imported.	-
unr	Specifies the imported source routing protocol as **unr**. User Network Route (UNR) is allocated if dynamic routing protocols cannot be used when users are getting online.	-
rip	Specifies the RIP protocol whose routes are imported.	-
process-id-rip	Specifies the process ID of the protocol whose routes are imported.	The value is an integer ranging from 1 to 65535. The default value is 1.
static	Specifies the static protocol whose routes are imported.	-
isis	Specifies the IS-IS protocol whose routes are imported.	-
process-id-isis	Specifies the process ID of the protocol whose routes are imported.	The value is an integer ranging from 1 to 65535. The default value is 1.
ospf	Specifies the OSPF protocol whose routes are imported.	-
process-id-ospf	Specifies the process ID of the protocol whose routes are imported.	The value is an integer ranging from 1 to 65535. The default value is 1.
limit *limit-number*	The maximum number of external routes that can be imported into an OSPF process.	The value is an integer that ranges from 1 to 4294967295.
cost *cost*	Indicates the route cost.	The value is an integer ranging from 0 to 16777214. For details about the default value, see **default (OSPF)**.

Parameter	Description	Value
route-policy *route-policy-name*	Imports only the route that meets the requirements of the specified route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
tag *tag*	Specifies the tag of the external LSA.	The value is an integer ranging from 0 to 4294967295. For details about the default value, see **default (OSPF)**.
type *type*	Specifies the type of the external routes.	The value is an integer ranging from 1 to 2. For details about the default value, see **default (OSPF)**. • 1: Type 1 external route • 2: Type 2 external route

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Importing routes discovered by other routing protocols can enrich OSPF routing information.

OSPF routes are classified into the following types in the descending order of priorities:

- Intra-area routes: refer to the routes in an area within an autonomous system.
- Inter-area routes: refer to the routes between different areas of the same AS. Intra-area routes and area external routes are internal routes of an AS.
- Type 1 external routes: When the cost of external routes equals that of AS internal routes, and can be compared with the cost of OSPF routes, these external routes have a high reliability and can be configured as Type 1 external routes.
- Type 2 external routes: When the cost of the routes from an ASBR to the destination outside an AS is much greater than the cost of the internal routes

to the ASBR, these external routes have a low reliability and can be configured as Type 2 external routes.

On a non-PE device, only EBGP routes are imported after the **import-route bgp** command is configured. IBGP routes are also imported after the **import-route bgp permit-ibgp** command is configured. If IBGP routes are imported, routing loops may occur. To prevent loops, run the **preference (OSPF)** and **preference (BGP)** commands to specify preferences for OSPF and BGP routes. If IBGP routes need to be imported, run the **import-route bgp permit-ibgp** command, and run the **preference (OSPF)** and **preference (BGP)** commands to set the preference of OSPF ASE routes lower than that of IBGP routes (preference value of OSPF ASE routes larger than that of IBGP routes).

On a PE, configuring the **import-route bgp** command imports both EBGP routes and IBGP routes, regardless of whether the **import-route bgp permit-ibgp** command is configured or not. If the **import-route bgp permit-ibgp** command and the **default-route-advertise (OSPF)** command are both configured, the active IBGP default routes can be imported into OSPF.

Prerequisites

To import certain external routes using a route-policy, a route-policy must have been created using the **route-policy** command before running the **import-route** command.

Configuration Impact

After a route-policy is configured, the OSPF process imports only routes that satisfy certain conditions. This prevents devices from receiving unrequired routes.

Precautions

You can use the **default (OSPF)** command to configure default parameters for external routes imported by OSPF, including the cost, type (Type 1 or Type 2), tag, and number of routes.

NOTE

The **import-route (OSPF)** command cannot import the default route of an external protocol. To enable a device to advertise the default route of an external protocol it learns when updating the OSPF routing table to other devices within the area, run the **default-route-advertise (OSPF)** command.

After the **import-route direct** command is executed, routes to the network segment where the IP address of the management interface belongs are also imported in the OSPF routing table. Therefore, use this command with caution.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all the routes learned by the specified protocol are imported to the OSPF routing table.

Example

Import Type 2 RIP routes, with the process 40, the tag being 33 and cost being 50.

```
<HUAWEI> system-view
```



```
[HUAWEI] ospf 100
[HUAWEI-ospf-100] import-route rip 40 type 2 tag 33 cost 50
```

7.4.57 local-mt enable (OSPF)

Function

The **local-mt enable** command enables OSPF Local Multicast-Topology (MT).

The **undo local-mt enable** command disables OSPF local MT.

By default, OSPF local MT is disabled.

NOTE

Only the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-Hsupport this command.

Format

local-mt enable

undo local-mt enable

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When multicast and an MPLS TE tunnel are configured on a network and the TE tunnel is configured with IGP Shortcut, the outbound interface of the route calculated by an IGP is not an actual physical interface but a TE tunnel interface. The device uses a unicast route to the multicast source address to send multicast Join messages from the TE tunnel interface. The device spanned by the TE tunnel, however, cannot sense the multicast Join messages. As a result, multicast forwarding entries will not be created on these devices.

To avoid the preceding problem, you can run the **local-mt enable** command to enable OSPF local Multicast-Topology (MT). MT allows the multicast routing entries to be correct created to guide the forwarding of multicast traffic.

After OSPF local MT is enabled, if the outbound interface of the calculated route is a TE tunnel interface of the IGP Shortcut type, the route management (RM) module creates a separate MIGP routing table for the multicast protocol,

calculates the actual physical outbound interface for the route, and adds the physical interface to the MIGP routing table for multicast forwarding.

Prerequisites

IGP Shortcut feature has been enabled using the **enable traffic-adjustment** command.

Precautions

- To control the number of entries in the MIGP routing table and speed up the MIGP routing table search, you can run the **local-mt filter-policy** command to configure a policy for filtering multicast source address. Then only the multicast source address that matches the policy to be added to the MIGP routing table. Configure the routing policy before enabling local MT. This can prevent excessive routes to non-multicast source address from being added to the MIGP routing table and keeps the number of routes in the MIGP routing table within the upper limit.
- Local MT is applicable only to the OSPF process of a public network instance.
- OSPF local MT does not support forwarding adjacency (FA).

Example

Enable local MT on OSPF process 1.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] local-mt enable
```

7.4.58 local-mt filter-policy (OSPF)

Function

The **local-mt filter-policy** command configures the filtering policy of OSPF Local Multicast-Topology (MT).

The **undo local-mt filter-policy** command deletes the filtering policy of OSPF local MT.

By default, the filtering policy of OSPF local MT is not configured.

Format

local-mt filter-policy { **acl** { *acl-number* | *acl-name* } | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* }

undo local-mt filter-policy

Parameters

Parameter	Description	Value
acl *acl-number*	Specifies the number of the basic access control list (ACL).	The value is an integer ranging from 2000 to 2999.

Parameter	Description	Value
acl *acl-name*	Specifies the name of a Named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of the IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of the route policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When multicast and an MPLS TE tunnel are configured on a network and the TE tunnel is configured with IGP Shortcut, the outbound interface of the route calculated by an IGP is not an actual physical interface but a TE tunnel interface. The device uses a unicast route to the multicast source address to send multicast Join messages from the TE tunnel interface. The device through which the TE tunnel passes, however, cannot sense the multicast Join messages. As a result, multicast forwarding entries will not be created.

To avoid the preceding problem, you can create a proper multicast routing table to guide the forwarding of multicast packets by configuring local MT. After OSPF local MT is enabled, if the outbound interface of the calculated route is a TE tunnel interface of the IGP Shortcut type, the route management (RM) module creates a separate MIGP routing table for the multicast protocol, calculates the actual physical outbound interface for the route, and adds the physical interface to the MIGP routing table for multicast forwarding.

To control the number of entries in the MIGP routing table and speed up the MIGP routing table search, you can configure filtering conditions by using the **local-mt filter-policy** command to allow only the matching routes to the multicast source address to be added to the MIGP routing table.

Prerequisites

OSPF local MT has been enabled using the **local-mt enable** command.

Precautions

You are recommended to configure the routing policy before enabling local MT. This can prevent the excessive routes to the not-multicast source address from being added to the MIGP routing table and keeps the number of routes in the MIGP routing table within the upper limit.

When the rule command is used to configure filtering rules for an ACL configured using the **acl** command, only the source address range that is specified by the **source** parameter and the period of time that is specified by the **time-range** parameter take effect.

Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all OSPF routes are added to the MIGP routing table.

Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all OSPF routes are added to the MIGP routing table.

Example

Configure the filtering policy of OSPF local MT.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] local-mt enable
[HUAWEI-ospf-1] local-mt filter-policy acl 2000
```

7.4.59 loop-free-alternate (OSPF)

Function

The **loop-free-alternate** command enables OSPF IP FRR to enable the device to use the LFA algorithm to calculate the nexthop and outbound interface for the dynamic backup link.

The **undo loop-free-alternate** command disables OSPF IP FRR.

By default, this function is disabled.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

loop-free-alternate

undo loop-free-alternate

Parameters

None

Views

OSPF FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **frr** command run in the OSPF view creates and displays the OSPF FRR view only, but cannot enable the OSPF IP FRR function. You must run the **loop-free-alternate** command in the OSPF view to enable OSPF IP FRR use the LFA algorithm to calculate the nexthop and outbound interface for the dynamic backup link to create the loop-free backup route.

Follow-up Procedure

Run the **frr-policy route** command in the OSPF view to configure a filtering policy for OSPF IP FRR. Only the OSPF backup route that satisfies specific rules can be delivered to the forwarding table.

Example

Enable OSPF IP FRR to create the loop-free backup route.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] frr
[HUAWEI-ospf-1-frr] loop-free-alternate
```

7.4.60 lsa-arrival-interval

Function

The **lsa-arrival-interval** command sets an interval for receiving LSAs.

The **undo lsa-arrival-interval** command restores the default interval for receiving LSAs.

By default, the interval for receiving LSA packets is one second.

Format

lsa-arrival-interval { *interval* | **intelligent-timer** *max-interval start-interval hold-interval* }

undo lsa-arrival-interval

Parameters

Parameter	Description	Value
interval	Specifies the interval for receiving LSAs.	The value is an integer ranging from 0 to 10000, in ms.
intelligent-timer	Enables an intelligent timer to receive LSAs.	-
max-interval	Specifies the maximum interval for receiving LSAs.	The value is an integer ranging from 1 to 120000, in ms. The default value is 1000.
start-interval	Specifies the initial interval for receiving LSAs.	The value is an integer ranging from 0 to 60000, in ms. The default value is 500.
hold-interval	Specifies the Holdtime interval for receiving LSAs.	The value is an integer ranging from 1 to 60000, in ms. The default value is 500.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To prevent frequent LSA updates caused by network connection or frequent route flapping, you can set the interval for receiving LSAs using this command. A router receives LSAs only when it meets the set interval.

In a stable network where routes need to be fast converged, you can set the interval for receiving LSAs at millisecond level to receive and update LSAs in time. The change of the topology or the route, therefore, can be immediately detected, and LSDB synchronization speed is improved.

If there is no special requirement for the network, using the default value is recommended.

Configuration Impact

After this command is configured, the interval for receiving LSAs is as follows:

1. The initial interval for receiving LSAs is specified by the parameter *start-interval*.
2. The interval for receiving LSAs for the nth (n≥2) time is equal to *hold-interval* × $2^{(n-2)}$.
3. When the interval specified by *hold-interval* x $2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPF performs SPF calculation at the maximum interval until *max-interval* expires without flapping or the OSPF process is restarted.

Precautions

You are advised to set the receiving interval specified by **lsa-arrival-interval** to be a value smaller than or equal to the Holdtime interval specified by**lsa-originate-interval**.

Example

Set the interval for receiving LSAs to 0 milliseconds.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] lsa-arrival-interval 0
```

7.4.61 lsa-originate-interval

Function

The **lsa-originate-interval** command sets an interval for updating LSAs.

The **undo lsa-originate-interval** command restores the default interval for updating LSAs.

By default, the interval for updating LSA packets is 5 seconds.

Format

lsa-originate-interval { **0** | { **intelligent-timer** *max-interval start-interval hold-interval* | **other-type** *interval* } * }

undo lsa-originate-interval

Parameters

Parameter	Description	Value
0	Sets an interval for updating LSAs to 0 ms, that is, deletes the initial interval (5000 ms) for updating LSAs.	-

Parameter	Description	Value
intelligent-timer	Enables an intelligent timer to update OSPF router LSAs and network LSAs.	-
max-interval	Specifies the maximum interval for updating OSPF LSAs.	The value is an integer ranging from 1 to 120000, in ms. The default value is 5000.
start-interval	Specifies the initial interval for updating OSPF LSAs.	The value is an integer that ranges from 0 to 60000, in ms. The default value is 500.
hold-interval	Specifies the Holdtime interval for updating OSPF LSAs.	The value is an integer ranging from 1 to 60000, in ms. The default value is 1000.
other-type	Sets the interval for updating the LSAs other than the OSPF router LSAs and network LSAs.	-
interval	Specifies the interval for updating LSAs.	The value is an integer ranging from 0 to 10, in seconds. The default value is 5.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To prevent frequent LSA updates caused by network connection or frequent route flapping, you can set the interval for receiving LSAs using this command. A router receives LSAs only when it meets the set interval.

In a stable network where routes need to be fast converged, you can set the interval for receiving LSAs at millisecond level to update LSAs in time. The change of the topology or the route, therefore, can be immediately detected, and LSDB synchronization speed is improved.

If there is no special requirement for the network, using the default value is recommended.

Configuration Impact

After this command is configured, the interval for updating LSAs is as follows:

1. The initial interval for updating LSAs is specified by the parameter *start-interval*.
2. The interval for updating LSAs for the nth (n≥2) time is equal to *hold-interval* × $2^{(n-2)}$.
3. When the interval specified by *hold-interval* x $2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPF performs SPF calculation at the maximum interval until *max-interval* expires without flapping or the OSPF process is restarted.

Precautions

You are advised to set the updating interval specified by **lsa-originate-interval** to be a value longer than or equal to the Holdtime interval specified by **lsa-arrival-interval**.

Example

Set the interval for updating LSAs to 0 milliseconds.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] lsa-originate-interval 0
```

7.4.62 lsdb-overflow-limit

Function

The **lsdb-overflow-limit** command sets the maximum number of external LSAs in an OSPF LSDB.

The **undo lsdb-overflow-limit** command restores the default configuration.

By default, the maximum number of external LSAs is not set.

Format

lsdb-overflow-limit *number*

undo lsdb-overflow-limit

Parameters

Parameter	Description	Value
number	Specifies the maximum number of external LSAs in an LSDB.	The value is an integer ranging from 1 to 1000000.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the number of external LSAs (Type 5 and Type 7 LSAs) imported by OSPF exceeds the limit, the excessive external LSAs cannot be processed properly and are discarded.

To avoid the preceding problem, you can set the maximum number of external LSAs in the LSDB to adjust and optimize OSPF networks.

Prerequisites

If **OSPF is in LSDB overflow status** is displayed in the **display ospf lsdb brief** command output, you must run the **lsdb-overflow-limit** command to set the maximum number of external LSAs imported by OSPF.

Precautions

The configuration of this command must be consistent in the entire AS.

Example

Set a maximum number of 400000 external LSAs allowed in an LSDB.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] lsdb-overflow-limit 400000
```

7.4.63 maximum load-balancing (OSPF)

Function

The **maximum load-balancing** command sets the maximum number of equal-cost routes for load balancing.

The **undo maximum load-balancing** command restores the default setting.

By default, the maximum number of equal-cost routes on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S6720S-S, and S5736-S is 8, and the maximum number of equal-cost routes on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S is 16.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

maximum load-balancing *number*

undo maximum load-balancing

Parameters

Parameter	Description	Value
number	Specifies the maximum number of equal-cost routes.	The value is an integer that ranges from 1 to 8 on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S and S6720S-S. The value ranges from 1 to 16 on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If a routing protocol discovers multiple routes with the same cost to a single destination, traffic can be balanced among these routes. The **maximum load-balancing** command sets the maximum number of equal-cost routes that can carry out load balancing. This optimizes the routing policy and ensures traffic forwarding on a complex network.

Configuration Impact

Packets will be balanced among multiple equal-cost routes to a single destination.

Follow-up Procedure

If more existing equal-cost OSPFv3 routes than the value set using the **maximum load-balancing** command are available, valid routes are selected for load balancing based on the following criteria:

- Route preference: Routes with lower preferences are selected for load balancing.
- Interface index: If routes have the same priorities, routes with higher interface index values are selected for load balancing.
- Next hop IP address: If routes have the same priorities and interface index values, routes with larger IP address are selected for load balancing.

The **nexthop** command allows routes with a specified weight to carry out load balancing.

Precautions

To disable load balancing, set the value of *number* to 1.

Example

Set the maximum number of the equal-cost routes.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] maximum load-balancing 2
```

Restore the default maximum number of equal-cost routes for carrying out load balancing.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] undo maximum load-balancing
```

7.4.64 maximum-routes

Function

The **maximum-routes** command sets the maximum number of routes of different types that OSPF supports.

The **undo maximum-routes** command restores the default maximum number of routes of different types that OSPF supports.

Format

maximum-routes { **external** | **inter** | **intra** } *number*

undo maximum-routes { **external** | **inter** | **intra** }

Parameters

Parameter	Description	Value
external	Indicates the maximum number of AS external routes that OSPF supports.	-
inter	Indicates the maximum number of inter-area routes that OSPF supports.	-
intra	Indicates the maximum number of intra-area routes that OSPF supports.	-
number	Specifies the maximum number of routes of different types.	The value is an integer. • When the **external** parameter is specified, the value of *number* is an integer ranging from 100 to 5000000. • When the **inter** parameter is specified, the value of *number* is an integer ranging from 100 to 1000000. • When the **intra** parameter is specified, the value of *number* is an integer ranging from 100 to 100000.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Based on the real world situation of a network, such as the convergence speed, you can run the **maximum-routes** command to configure the maximum number of routes of different types. This ensures full use of network resources and improves network performance.

Precautions

The maximum number of routes supported by OSPF cannot exceed the maximum number of all routes supported by the switch.

The **maximum-routes** command configuration limits the maximum number of routes that can be locally calculated. Therefore, the command configuration

affects the maximum number of routes that can be sent, but does not affect the maximum number of LSAs that can be received.

Example

Set the maximum number of routes that OSPF supports.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] maximum-routes intra 500
```

7.4.65 mesh-group enable

Function

The **mesh-group enable** command enables the mesh-group function.

The **undo mesh-group enable** command disables the mesh-group function.

By default, the mesh-group function is disabled.

Format

mesh-group enable

undo mesh-group enable

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When concurrent links exist between the device and its neighbor, the **mesh-group enable** command enables the mesh-group function. The router ID of a neighboring device uniquely identifies a mesh group. After LSAs are received, the device selects a primary link to flood the received LSAs, without performing reverse flooding. This prevents repeated flooding, reduces the load on the links, and saves system resources.

Device interfaces that meet the following conditions can form a mesh group:

- The interfaces belong to the same area and OSPF process.
- The neighbor status is Exchange or Full.
- Each interface is connected only to one neighbor.

Precautions

After the device is enabled with the mesh-group function, if the router IDs of the directly connected neighbor are the same, the LSDBs of the entire network cannot be synchronized and routes cannot be calculated correctly. In this case, you need to reconfigure the router ID of the neighbors, and then restart the neighbors to validate the configured router ID.

Example

Enable the mesh-group function on the switch.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] mesh-group enable
```

7.4.66 network (OSPF area)

Function

The **network** command specifies the interface that runs OSPF and the area to which the interface belongs.

The **undo network** command deletes the interface that runs OSPF.

By default, an interface does not belong to any area.

Format

network *network-address wildcard-mask* [**description** *text*]

undo network *network-address wildcard-mask*

Parameters

Parameter	Description	Value
network-address	Specifies the address of the network segment where the interface resides.	The value is in dotted decimal notation.
wildcard-mask	Specifies the wildcard mask of an IP address, which is the reverse form of the mask of the IP address. For example, 0.0.0.255 indicates that the mask length is 24 bits. **NOTE** When configuring a short mask, ensure that the area does not contain many interface addresses and secondary addresses. Otherwise, the neighbor relationship may fail to be established due to excessively long router-LSAs.	The value is in dotted decimal notation.

Parameter	Description	Value
description *text*	Specifies the description of the specified OSPF network segment.	The value is a string of 1 to 80 case-sensitive characters with spaces supported.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After creating an OSPF process, you can run the **network** command to configure the network segments in an area and specify *network-address* and *wildcard-mask* to configure one or multiple interfaces in an area. To run OSPF on an interface, ensure that the primary IP address of this interface is in the network segment range specified in the **network** (OSPF) command. If the secondary IP address of the main interface is in the network segment range specified in this command, the main interface cannot run OSPF.

OSPF can run on an interface only when the following two conditions are met:

- The mask length of the interface's IP address is not less than that specified in the **network** command. OSPF uses a reverse mask. For example, 0.0.0.255 indicates that the mask length is 24 bits.

  NOTE

  When the *wildcard-mask* parameter in the **network** command is set to all 0s, OSPF runs on the interface if its IP address is the IP address specified in the **network** *network-address* command.

- The primary address of the interface must be within the network segment range specified in the **network** command.

Precautions

- OSPF neighbor relationships cannot be established using the secondary IP addresses of interfaces.
- After the **network** 0.0.0.0 0.0.0.0 command is configured, the device automatically changes the command to **network** 0.0.0.0 255.255.255.255. That is, all interfaces (including the management interface) will run OSPF. Therefore, exercise caution when configuring this command.
- For the same **network** *address wildcard-mask*, the last description configured by **description** takes effect.

- On a loopback interface, by default, OSPF advertises its IP address in the form of a 32-bit host route, independent of the mask length of the IP address on the interface.
- To advertise the network segment route of a loopback interface, you need to run the **ospf network-type** command to set the network type to broadcast or NBMA.
- When an OSPF sham link is configured, the local address cannot be advertised through the OSPF process of a private network.
- Two areas that overlap cannot be configured between different processes in the same instance, or between different areas in the same process.
- The **ospf enable** command configuration takes precedence over the **network** command configuration.

Example

Configure the primary IP address of the interface that runs OSPF to be in the network segment of 192.168.1.0/24, set the ID of the OSPF area where the interface resides to 2, and configure the description for the network segment.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 2
[HUAWEI-ospf-100-area-0.0.0.2] network 192.168.1.0 0.0.0.255 description this network is connected to Beijing
```

7.4.67 nexthop (OSPF)

Function

The **nexthop** command sets a preference for equal-cost routes. After OSPF calculates the equal-cost routes, the next hop is chosen from these equal-cost routes based on the value of weights. A smaller value indicates a higher preference.

The **undo nexthop** command cancels the preference of these equal-cost routes.

By default, the value of weight is 255. Equal-cost routes have no preference, and they forward packets at the same time. Load balancing is performed among them.

Format

nexthop *ip-address* **weight** *value*

undo nexthop *ip-address*

Parameters

Parameter	Description	Value
ip-address	Indicates the IP address of the next hop.	The value is in dotted decimal notation.

Parameter	Description	Value
weight *value*	Indicates the weight of the next hop. A smaller value indicates a higher preference.	It is an integer that ranges from 1 to 254.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

You can run the **nexthop** command to set a high preference for specified valid routes if both of the following conditions are met:

- The number of equal-cost routes on a network is larger than the value configured using the **maximum load-balancing** command.
- Valid routes must be specified for load balancing.

Example

Set a preference for equal-cost routes in OSPF.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] nexthop 10.0.0.3 weight 1
```

7.4.68 nssa (OSPF Area)

Function

The **nssa** command configures an NSSA.

The **undo nssa** command cancels the configuration of an NSSA.

By default, no OSPF area is configured as an NSSA.

Format

nssa [{ **default-route-advertise** [**backbone-peer-ignore**] | **suppress-default-route** } | **flush-waiting-timer** *interval-value* | **no-import-route** | **no-summary** | **set-n-bit** | **suppress-forwarding-address** | **translator-always** | **translator-interval** *interval-value* | **zero-address-forwarding** | **translator-strict**] *

undo nssa [**flush-waiting-timer** *interval-value*]

Parameters

Parameter	Description	Value
default-route-advertise	Generates default Type7 LSAs on the ASBR and then advertises them to the NSSA. **NOTE** If the backbone-peer-ignore parameter is not configured, the ABR automatically generates a default NSSA LSA (Type7 LSA) and advertises it in the NSSA when an interface in the backbone area is in Up state and an OSPF neighbor relationship in Full state exists on the ABR. Type 7 LSAs carrying the default route will be generated only when the default route 0.0.0.0/0 exists in the routing table on the ASBR.	-
backbone-peer-ignore	Prevents the ABR from checking the neighbor status when the ABR generates default Type 7 LSAs and advertises them to the NSSA. Specifically, the ABR generates default Type 7 LSAs and advertises them to the NSSA as long as an interface that is Up exists in the backbone area.	-
suppress-default-route	Generates default Type-7 LSAs on the ASBR or ABR and then not advertises them to the NSSA.	-
flush-waiting-timer *interval-value*	Indicates the interval for an ASBR to send aged Type 5 LSAs. The parameter takes effect only when it is set.	The value is an integer that ranges from 1 to 40, in seconds.
no-import-route	Indicates that no external route is imported to an NSSA.	-
no-summary	Indicates that an ABR is prohibited from sending summary LSAs to the NSSA.	-
set-n-bit	Sets the N-bit in DD packets.	-
suppress-forwarding-address	Sets the FA of the Type 5 LSAs translated from Type 7 LSAs by the NSSA ABR to 0.0.0.0.	-
translator-always	Specifies an ABR in an NSSA as an all-the-time translator. Multiple ABRs in an NSSA can be configured as translators.	-

Parameter	Description	Value
translator-interval *interval-value*	Specifies the timeout period of a translator.	The value is an integer ranging from 1 to 120, in seconds. The default value is 40.
zero-address-forwarding	Sets the FA of the generated NSSA LSAs to 0.0.0.0 when external routes are imported by the ABR in an NSSA.	-
translator-strict	Configures the translator to perform strict check on the P-bit flag. The translator determines whether to translate Type 7 LSAs into Type 5 LSAs based on the P-bit flag.	-

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An NSSA is configured in the scenario where AS external routes are to be imported but not forwarded to save system resources. AS external routes can be imported to an NSSA and transmitted to the entire NSSA.

All switches in the NSSA must be configured with NSSA attributes using the **nssa** command.

The **nssa** command is applicable to the following scenarios:

- The **default-route-advertise** parameter is configured to advertise Type 7 LSAs carrying the default route on the ASBR to the NSSA.

  Regardless of whether the default route 0.0.0.0/0 exists in the routing table on the ABR, Type 7 LSAs carrying the default route will be generated. However, Type 7 LSAs carrying the default route will be generated on the ASBR only when the default route 0.0.0.0/0 exists in the routing table on the ASBR.
- When the area to which the ASBR belongs is configured as an NSSA, invalid Type 5 LSAs from other switches in the area where LSAs are flooded will be reserved. These LSAs will be deleted only when the aging time reaches 3600s. The switch performance is affected because the forwarding of a large number of LSAs consumes the memory resources. To resolve such a problem, you can

set the parameter **flush-waiting-timer** to the maximum value 3600s for Type 5 LSAs. This ensures that the invalid Type 5 LSAs from other switches can be deleted in time.

NOTE

- When the LS age field value (aging time) in the header of an LSA reaches 3600s, the LSA is deleted.
- If an ASBR also functions as an ABR, **flush-waiting-timer** does not take effect. This prevents Type 5 LSAs in the non-NSSAs from being deleted.

- If an ASBR also functions as an ABR, the **no-import-route** parameter is configured to prevent external routes imported using the **import-route** command from being advertised to the NSSA.
- The **no-summary** parameter is configured on an ABR to reduce the number of LSAs that are transmitted to the NSSA. This implementation prevents the ABR from transmitting Type 3 LSAs to the NSSA.

  NOTE

  After the **nssa default-route-advertise backbone-peer-ignore no-summary** command is run, the ABR generates default Type 7 and Type 3 LSAs as long as an interface that is Up exists in the backbone area. The default Type 3 LSAs preferentially take effect.
- After the **set-n-bit** parameter is configured, the N-bit is set in the database description (DD) packets during the synchronization between the switch and neighboring switches.
- If multiple ABRs are deployed in the NSSA, the system automatically selects an ABR (generally the switch with the largest router ID) as a translator to convert Type 7 LSAs into Type 5 LSAs. You can configure the **translator-always** parameter on an ABR to specify the ABR as an all-the-time translator. To specify two ABRs for load balancing, configure the **translator-always** parameter on the chosen ABRs to specify the ABRs as all-the-time translators. You can use this command to pre-configure a fixed translator to prevent LSA flooding caused by translator role changes.
- The **translator-interval** parameter is used to ensure uninterrupted services when translator roles change. The value of *interval-value* must be greater than the flooding period.

Configuration Impact

Configuring or deleting NSSA attributes will trigger routing updates in the area. A second configuration of NSSA attributes can be implemented or canceled only after a routing update is complete.

Precautions

It is recommended that a loopback address be configured for a device in the NSSA and be enabled in the corresponding OSPF area, so that the loopback address can be automatically selected as the FA. If other switches have routes of the same cost to the switch in the NSSA, load balancing is performed.

When the last ordinary area (other than a stub area or NSSA) under an OSPF process is deleted, useless Type 5 LSAs originated by the local switch in the area where LSAs are flooded will be deleted immediately. The local switch still reserves useless Type 5 LSAs from other switches. These useless Type 5 LSAs will be deleted only when the aging time reaches 3600s.

Example

Configure area 1 as an NSSA.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] area 1
[HUAWEI-ospf-1-area-0.0.0.1] nssa
```

7.4.69 opaque-capability enable

Function

The **opaque-capability enable** command enables the Opaque-LSA capability so that an OSPF process can generate Opaque LSAs and receive Opaque LSAs from neighbors.

The **undo opaque-capability** command disables the Opaque-LSA capability.

By default, the Opaque-LSA capability is disabled.

Format

opaque-capability enable

undo opaque-capability

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Opaque LSAs provide a generic mechanism for OSPF extension:

- OSPF supports GR using Type 9 LSAs.
- OSPF supports TE using Type 10 LSAs.

Before configuring OSPF GR or OSPF TE, you must enable opaque LSA capability running the **opaque-capability enable** command.

Configuration Impact

Enabling or disabling the opaque LSA function may delete and re-establish all sessions and instances.

Example

Enable OSPF opaque-lsa.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] opaque-capability enable
```

7.4.70 ospf

Function

The **ospf** command creates and runs an OSPF process.

The **undo ospf** command terminates an OSPF process.

By default, OSPF is disabled, that is, no OSPF process runs.

Format

ospf [*process-id* | **router-id** *router-id* | **vpn-instance** *vpn-instance-name*] *

undo ospf *process-id* [**flush-waiting-timer** *time*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535. By default, it is 1.
router-id *router-id*	Specifies a router ID.	The value is in dotted decimal notation.
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
flush-waiting-timer *time*	Indicates the interval for generating aged LSAs. The parameter takes effect only when it is set.	The value is an integer that ranges from 1 to 40, in seconds.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can set OSPF parameters only after creating an OSPF process.

OSPF supports multi-process. More than one OSPF process can run on the same device and are independent of each other. Route interaction between different OSPF processes is similar to route interaction between different routing protocols.

An OSPF router ID can be manually configured or automatically selected. To ensure OSPF stability, you are advised to manually configure the IP address of a loopback interface as the router ID.

If no OSPF router ID is manually configured for a router, the router selects the global router ID as its OSPF router ID. If both the OSPF router ID and global router ID are not configured, the router selects a router ID based on IP addresses of current interfaces:

- If there are loopback interfaces that have IP addresses configured, the device selects the largest IP address among loopback interface addresses as the global router ID.
- If no loopback interface is configured or loopback interfaces do not have IP addresses configured, the device selects the largest IP address among interface addresses as the global router ID without considering the Up/Down state of interfaces.

In any of the following situations, the router ID is re-selected:

- The **ospf** command is run to re-configure an OSPF router ID, and the OSPF process is restarted.
- The global router ID is re-configured, and the OSPF process is restarted.
- The IP address of the original global router ID is deleted, and the OSPF process is restarted.

Configuration Impact

After an OSPF process is disabled using **undo ospf** command, the receive end still maintains the LSAs generated by this OSPF process. These invalid LSAs occupy the system memory and are deleted only when the LS age field (aging time) reaches 3600 seconds. When the **undo ospf** *process-id* **flush-waiting-timer** *time* command is used to delete an OSPF process, the switch regenerates an LSA in the set time and sets the LS age field to 3600 seconds. After other switches receive the LSA with the LS age field as 3600 seconds, they delete the LSA immediately. If the host does not send all the LSAs in the set time, other switches still reserve invalid LSAs.

NOTE

When the LS age field (aging time) in the LSA header reaches 3600 seconds, this LSA is deleted.

Precautions

An interface on a device belongs to only one OSPF process.

If a VPN instance is specified, the OSPF process specified in this command belongs to this VPN instance. If no VPN instance is specified, the OSPF process specified in

this command belongs to the global VPN instance. *vpn-instance-name* cannot be changed after being specified.

NOTE

The router ID of each OSPF process must be unique on the entire network; otherwise, the OSPF neighbor relationship cannot be set up and routing information is incorrect. Configuring a unique router ID for each OSPF process on each OSPF device is recommended.

Example

Run an OSPF process.

```
<HUAWEI> system-view
[HUAWEI] ospf 100 router-id 10.10.10.1 vpn-instance test
```

7.4.71 ospf authentication-mode

Function

The **ospf authentication-mode** command sets an authentication mode and password used between neighboring nodes.

The **ospf authentication-mode null** command configures the null authentication mode on an interface.

The **undo ospf authentication-mode** command deletes the authentication mode on an interface.

By default, an interface does not authenticate OSPF packets.

Format

ospf authentication-mode { **simple** [**plain** *plain-text* | [**cipher**] *cipher-text*] | **null** }

ospf authentication-mode { **md5** | **hmac-md5** | **hmac-sha256** } [*key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* }]

ospf authentication-mode keychain *keychain-name*

undo ospf authentication-mode

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
simple	Indicates simple authentication. **NOTICE** Simple authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	In simple authentication, the password type is **cipher** by default.
plain	Indicates plain authentication. Only plain text can be entered, and only plain text is displayed when the configuration file is viewed. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This carries security risks. Selecting **cipher** to save the password in cipher text is recommended.	-
plain-text	Specifies a plain text password.	*plain-text* is a string of 1 to 8 characters without spaces when **simple** is configured, and is a string of 1 to 255 characters without spaces when **md5**, **hmac-md5** or **hmac-sha256** is configured.
cipher	Indicates cipher authentication. Either plain text or cipher text can be entered, and cipher text is displayed when the configuration file is viewed.	When **cipher** is configured, enter only the password in cipher text. Then, the password is displayed in cipher text in configuration files. MD5 authentication, HMAC-SHA256 authentication or HMAC-MD5 authentication uses the password in cipher text by default.

Parameter	Description	Value
cipher-text	Specifies a cipher text password.	The value is a string of characters without spaces. In simple authentication, a plain text password is a string of 1 to 8 characters and a cipher text password is a string of 24 or 32 or 48 characters. In MD5 authentication, HMAC-SHA256 authentication or HMAC-MD5 authentication, a plain text password is a string of 1 to 255 characters and a cipher text password is a string of 20 to 392 characters.
md5	Indicates MD5 authentication. **NOTICE** MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-md5	Indicates HMAC-MD5 authentication. **NOTICE** HMAC-MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-sha256	Indicates HMAC-SHA256 authentication.	-
key-id	Specifies the authentication key ID of the interface's cipher authentication. The key ID must be consistent with that of the peer.	The value is an integer that ranges from 1 to 255.

Parameter	Description	Value
keychain	Indicates keychain authentication. **NOTE** Before configuring this parameter, run the **keychain** command to create a keychain. Then, run the **key-id**, **key-string**, and **algorithm** commands to configure a key ID, a password, and an authentication algorithm for this keychain. Otherwise, OSPF authentication will fail. Currently, OSPF supports only HMAC-MD5 and HMAC-SHA256 algorithms.	-
keychain-name	Specifies the keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.
null	Indicates null authentication.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Due to the defects and non-strict implementation of the TCP/IP protocol suite and increasing attacks on TCP/IP networks, the impact generated by attacks on the network may become more serious. Attacks on network devices may lead to a network crash. To improve OSPF network security, configure authentication.

Configuration Impact

Interface authentication is used to set the authentication mode and password used between neighboring devices. It takes precedence over area authentication.

Precautions

Null authentication is an authentication method. It does not indicate that no authentication is configured.

The authentication mode and password configured for interfaces on the same network segment must be the same.

OSPF does not support the configuration on a null interface.

An authentication password cannot contain spaces.

Example

Configure OSPF HMAC-SHA256 authentication on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf authentication-mode hmac-sha256
```

Configure OSPF HMAC-SHA256 authentication on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf authentication-mode hmac-sha256
```

7.4.72 ospf bfd

Function

The **ospf bfd** command enables BFD on an OSPF interface or sets parameter values for a BFD session.

The **undo ospf bfd** command deletes BFD on an OSPF interface or restores the default parameter values of a BFD session.

By default, BFD is not enabled on any OSPF interfaces.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ospf bfd enable

undo ospf bfd enable

ospf bfd { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* | **frr-binding** } *

undo ospf bfd { **min-rx-interval** [*receive-interval*] | **min-tx-interval** [*transmit-interval*] | **detect-multiplier** [*multiplier-value*] | **frr-binding** } *

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **frr-binding** parameter.

Parameters

Parameter	Description	Value
enable	Enables BFD.	-
min-rx-interval *receive-interval*	Indicates the minimum interval at which BFD packets are received from the remote end.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *transmit-interval*	Indicates the minimum interval at which BFD packets are sent to the remote end.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier-value*	Specifies the local detection multiplier.	The value is an integer ranging from 3 to 50. By default, it is 3.

Parameter	Description	Value
frr-binding	Binds the BFD status to the link status of an interface. That is, when the BFD status goes Down, the link status of the interface also goes Down. This enables traffic to be switched to the backup path.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A link failure or topology change causes devices to recalculate routes. Therefore, the convergence of routing protocols must be sped up to improve network performance.

Link faults are unavoidable. Therefore, a feasible solution is required to detect faults faster and notify routing protocols of the faults immediately. If BFD is associated with routing protocols and a link fault occurs, BFD can speed up the convergence of routing protocols.

Prerequisites

The configured parameters of the BFD session are valid on an interface only when BFD is enabled on the interface.

Procedure

The *receive-interval* is obtained through the negotiation between the local end and peer end by comparing the values of the local **min-rx-interval** and the peer **min-tx-interval**. If the local end fails to receive a BFD packet from the peer end within an interval of *receive-interval* x *multiplier-value*, it considers that the neighbor is Down.

Configuration Impact

If global BFD is not enabled, you can enable BFD on an interface but cannot set up BFD sessions. Similarly, if only parameters of a BFD session are set but the **ospf bfd enable** command is not used, the BFD session cannot be set up.

BFD configured on an interface takes precedence over BFD configured in a process. If BFD is enabled on an interface, the BFD parameters on the interface are used to establish BFD sessions.

Precautions

- After BFD is enabled, BFD sessions can be created only between the two ends that have set up an OSPF neighbor relationship and the relationship is in the Exstart state.
- The **ospf bfd enable** command and the **ospf bfd block** command are mutually exclusive.
- After BFD is disabled from an interface through the **undo ospf bfd enable** command, the parameters for setting up BFD sessions remain on this interface but do not take effect.

Example

Enable BFD on VLANIF100 and set the minimum interval for receiving BFD packets to 400 ms and the local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf bfd enable
[HUAWEI-Vlanif100] ospf bfd min-rx-interval 400 detect-multiplier 4
```

Enable BFD on GE0/0/1 and set the minimum interval for receiving BFD packets to 400 ms and the local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf bfd enable
[HUAWEI-GigabitEthernet0/0/1] ospf bfd min-rx-interval 400 detect-multiplier 4
```

7.4.73 ospf bfd block

Function

The **ospf bfd block** command prevents an interface from dynamically setting up a BFD session.

The **undo ospf bfd block** command cancels the configuration.

By default, the device does not prevent an interface from dynamically setting up a BFD session.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ospf bfd block

undo ospf bfd block

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **bfd all-interfaces enable** command is used for an OSPF process, BFD sessions are created on all OSPF interfaces whose neighbor status is Full. If you do not want some interfaces to have BFD sessions created, run the **ospf bfd block** command to prevent these interfaces from dynamically setting up BFD sessions.

Prerequisites

BFD has been enabled on interfaces.

Precautions

The **ospf bfd enable** command and the **ospf bfd block** command are mutually exclusive.

Example

Prevent VLANIF100 from dynamically setting up a BFD session.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf bfd block
```

Prevent GE0/0/1 from dynamically setting up a BFD session.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf bfd block
```

7.4.74 ospf cost

Function

The **ospf cost** command sets a cost for an OSPF on an interface.

The **undo ospf cost** command restores the default cost for OSPF.

By default, OSPF automatically calculates its cost based on its interface bandwidth. The default cost of the loopback interface is 0.

Format

ospf cost *cost*

undo ospf cost

Parameters

Parameter	Description	Value
cost	Specifies the cost of an OSPF-enabled interface.	The value is an integer ranging from 1 to 65535. By default, it is 1.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, OSPF automatically calculates the link cost for an interface based on the bandwidth of the interface. You can modify the interface cost using the **ospf cost** command. For example, when you configure load balancing, you can run this command to modify the interface cost of the routes if conditions are satisfied for load balancing.

If no cost is set for an OSPF interface using the **ospf cost** command, OSPF automatically calculates its cost based on the interface bandwidth. The calculation formula is as follows: Cost of the interface = Bandwidth reference value/Interface bandwidth. The integer of the calculated result is the cost of the interface. If the calculated result is smaller than 1, the cost is 1. Changing the bandwidth reference value or interface bandwidth change the cost of an interface:

- To use the bandwidth reference value to determine the interface cost, run the **bandwidth-reference** command.
- To use the interface bandwidth to determine the interface cost, run the **bandwidth** command to set configuration bandwidth for an interface, and then run the **bandwidth-config enable** command to enable the device to calculate the cost for the interface based on the configuration bandwidth of the interface.
- To use both the bandwidth reference value and interface bandwidth to determine the interface cost, run the preceding three commands.

Precautions

- The **ospf cost** command cannot run on null interfaces.
- No default cost is configured for trunk interfaces, because a trunk interface has multiple member interfaces that are in constant change.

Example

Set the cost of VLANIF100 that runs OSPF to 65.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf cost 65
```

Set the cost of GE0/0/1 that runs OSPF to 65.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf cost 65
```

7.4.75 ospf dr-priority

Function

The **ospf dr-priority** command sets a priority for an interface that participates in the DR election.

The **undo ospf dr-priority** command restores the default setting.

By default, the priority is 1.

Format

ospf dr-priority *priority*

undo ospf dr-priority

Parameters

Parameter	Description	Value
priority	Specifies the priority of an interface that participates in the DR or BDR election. A larger value indicates a higher priority.	The value is an integer ranging from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The DR priority of an interface determines whether it is qualified to be a DR. The interface with the highest DR priority is elected as the DR. If the DR priority of an interface on a device is 0, the device cannot be elected as a DR or a BDR. On a broadcast or NBMA network, you can set a DR priority for an interface to determine whether it is qualified to be a DR or a BDR.

Configuration Impact

When the DR and BDR are elected on a network segment, they send DD packets to all neighboring nodes and set up adjacencies with all neighboring nodes.

Precautions

NOTICE

Restarting or shutting down an interface will interrupt the OSPF adjacency between devices. Therefore, perform the operation with caution.

If the DR priority of a device is re-configured, the DR or BDR on the network will not be re-elected. You can re-elect a DR or a BDR by using either of the following methods. This, however, will interrupt the OSPF adjacency between devices. Therefore, use the following methods with caution.

- Restart the OSPF processes on all devices.
- Run the **shutdown** and then **undo shutdown** commands on the interfaces where OSPF adjacencies are set up.

In OSPF, the DR priority cannot be configured for null interfaces.

Example

Set the priority of VLANIF100 that participates in the DR election to 8.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf dr-priority 8
```

Set the priority of GE0/0/1 that participates in the DR election to 8.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf dr-priority 8
```

7.4.76 ospf enable

Function

The **ospf enable** command enables OSPF on an interface.

The **undo ospf enable** command disables OSPF on an interface.

By default, the interface does not run OSPF.

Format

ospf enable [*process-id*] **area** *area-id*

undo ospf enable [*process-id*] **area** *area-id*

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535. The default value is 1.
area *area-id*	Specifies an area ID.	The value can be a decimal integer or an IP address. When the value is an integer, the value ranges from 0 to 4294967295.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ospf enable** command configuration takes precedence over the **network** command configuration.

After the **undo ospf enable** command is run to disable OSPF on an interface, the **network** configuration takes effect on the interface automatically.

Configuration Impact

The interface will alternate between Up and Down when the **ospf enable** command and the **network** command are run on the interface repeatedly.

Precautions

An interface can be configured with only one OSPF process.

The configured interface and the OSPF process must be in the same VPN.

- The **ospf enable** command can be configured on an interface before an OSPF process is created. The interface specified by the **ospf enable** command and the created OSPF process must be in the same VPN.
- If a process is created before the **ospf enable** command is run on an interface, the process of the interface and existing process must belong to the same VPN. Otherwise, the **ospf enable** command cannot be run.
- If no OSPF process is created, interfaces that belong to different VPN instances cannot be added to the same OSPF process.

Example

Enable VLANIF100 in the specified OSPF area.

```
<HUAWEI> system-view
```



```
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf enable 1 area 0
```

Enable GE0/0/1 in the specified OSPF area.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf enable 1 area 0
```

7.4.77 ospf filter-lsa-out

Function

The **ospf filter-lsa-out** command configures an OSPF interface to filter outgoing LSAs.

The **undo ospf filter-lsa-out** command disables an OSPF interface from filtering outgoing LSAs.

By default, outgoing LSAs are not filtered.

Format

ospf filter-lsa-out { **all** | { **summary** [**acl** { *acl-number* | *acl-name* }] | **ase** [**acl** { *acl-number* | *acl-name* }] | **nssa** [**acl** { *acl-number* | *acl-name* }] } * }

undo ospf filter-lsa-out

Parameters

Parameter	Description	Value
all	Filters all outgoing LSAs except grace LSAs.	-
summary	Filters outgoing network summary LSAs (Type3).	-
ase	Filters outgoing AS external LSAs (Type5).	-
nssa	Filters outgoing NSSA LSAs (Type7).	-
acl *acl-number*	Specifies the number of the basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When multiple links exist between two switches, you can run the **ospf filter-lsa-out** command to configure the local switch to filter the outgoing LSAs based on the filtering policy before the device sends these LSAs along specified links. This reduces the unnecessary retransmission of LSAs and saves bandwidth resources.

Configuration Impact

Filtering the outgoing LSAs on the specified OSPF interface can prevent useless LSAs from being sent to neighbors. This can reduce the size of the LSDB of neighbors and speed up the network convergence.

NOTE

- After the command is configured on an interface, the OSPF neighbor relationship of the interface will automatically re-establish.
- When the rule in the ACL changes, the OSPF neighbor relationship of the interface will automatically re-establish.

Precautions

When the **rule** command is used to configure the filtering rules for an ACL configured using the **acl** command, only the source address range that is specified by the **source** parameter and the period of time that is specified by the **time-range** parameter take effect.

Grace LSAs are used to inform the neighbor of the Graceful Restart (GR) time, cause, and interface instance ID when GR starts and ends. The command is not used to filter the grace LSAs.

Example

Configure VLANIF100 to filter all outgoing LSAs except grace LSAs.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf filter-lsa-out all
```

Configure GE0/0/1 to filter all outgoing LSAs except grace LSAs.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf filter-lsa-out all
```


7.4.78 ospf frr block

Function

The **ospf frr block** command disables the OSPF IP FRR function on a specified interface.

The **undo ospf frr block** command restores the OSPF IP FRR function on the specified interface.

By default, the OSPF IP FRR function is enabled on a specific interface.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ospf frr block

undo ospf frr block

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

OSPF IP FRR can be disabled using the **ospf frr block** command on an interface of a specific device that is running important services and resides on an FRR backup link. This setting prevents the device connected to this interface from being a part of a backup link and being burdened after FRR switches traffic to the backup link.

Example

Disable the OSPF IP FRR function on the interface VLANIF 10.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ospf frr block
```

7.4.79 ospf maxage-lsa auto-protect disable

Function

The **ospf maxage-lsa auto-protect disable** command disables master/slave board switching triggered by abnormal OSPF LSA aging.

The **undo ospf maxage-lsa auto-protect disable** command enables master/slave board switching triggered by abnormal OSPF LSA aging.

By default, master/slave board switching triggered by abnormal OSPF LSA aging is enabled.

Format

ospf maxage-lsa auto-protect disable

undo ospf maxage-lsa auto-protect disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

If the local device clock is faster than usual and the aging timer expires abnormally, the local device incorrectly clears all Router LSAs from the peer device, which causes route flapping and service interruptions. To resolve this issue, master/slave board switching triggered by abnormal OSPF LSA aging is automatically enabled. Master/Slave board switching is triggered to restore network connections and service traffic when the following condition is met:

(Number of incorrectly cleared Router LSAs/Total number of Router LSAs) x 100% ≥ 80% (Router LSAs are those sent by the peer device to the local device)

Example

Disable master/slave board switching triggered by abnormal OSPF LSA aging.

```
<HUAWEI> system-view
```



```
[HUAWEI] ospf maxage-lsa auto-protect disable
```

7.4.80 ospf mib-binding

Function

The **ospf mib-binding** command binds an OSPF process to SNMP and enables OSPF to respond to SNMP requests.

The **undo ospf mib-binding** command cancels the binding.

By default, OSPF processes are not bound to SNMP.

Format

ospf mib-binding *process-id*

undo ospf mib-binding

Parameters

Parameter	Description	Value
process-id	Specifies an OSPF process ID.	The value is an integer ranging from 1 to 65535.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The OSPF MIB is a virtual database of the device status maintained by the managed devices.

Before performing read and write operations on the OSPF MIB, you must run the **ospf mib-binding** *process-id* command to bind a specified OSPF process to SNMP.

Prerequisites

An OSPF process has been created using the **ospf** command in the system view.

Example

Bind OSPF process 100 to SNMP.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] quit
```



```
[HUAWEI] ospf mib-binding 100
```

Cancel the binding.

```
<HUAWEI> system-view
[HUAWEI] undo ospf mib-binding
```

7.4.81 ospf mtu-enable

Function

The **ospf mtu-enable** command enables an interface to add its actual MTU in DD packets to be sent and check whether the MTU in a received DD packet is greater than the local MTU.

The **undo ospf mtu-enable** command restores the default settings.

By default, an interface adds the MTU 0 (not the actual MTU) in DD packets to be sent and does not check the MTUs in received DD packets.

Format

ospf mtu-enable

undo ospf mtu-enable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To improve compatibility with a non-Huawei device, an OSPF-enabled Huawei device adds the MTU 0 in DD packets to be sent and does not check the MTUs in received DD packets, allowing an OSPF neighbor relationship to be set up even if the two ends have different MTU settings.

However, under the default configuration, the non-Huawei device may discard an OSPF packet received from the Huawei device if the packet's actual MTU is greater than the MTU of the non-Huawei device. If the discarded packet is an LSU, an OSPF neighbor relationship can still be set up, but the route carried in the LSU fails to be learned, causing service interruptions.

To resolve this issue, run the **ospf mtu-enable** command to configure an interface to add the actual MTU in DD packets to be sent and check whether the MTU in a received DD packet is greater than the local MTU. If the interface MTU settings of

the local and remote ends are different, an OSPF neighbor relationship cannot enter the Full state. In this manner, MTU inconsistency can then be identified in time.

Precautions

OSPF does not support the configuration on a null interface.

After the command is configured, the system automatically restarts the OSPF process.

Example

Set Vlanif100 to fill in the MTU field when sending DD packets.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf mtu-enable
```

Set GigabitEthernet0/0/1 to fill in the MTU field when sending DD packets.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf mtu-enable
```

7.4.82 ospf network-type

Function

The **ospf network-type** command sets a network type for an OSPF interface.

The **undo ospf network-type** command restores the default network type of the OSPF interface.

By default, the network type of an interface is determined by the physical interface. The network type of Ethernet interfaces is **broadcast**.

Format

ospf network-type { **broadcast** | **nbma** | **p2mp** | **p2p** [**peer-ip-ignore**] }

undo ospf network-type

Parameters

Parameter	Description	Value
broadcast	Indicates that the network type of the interface is changed to broadcast.	-
nbma	Indicates that the network type of the interface is changed to NBMA.	-
p2mp	Indicates that the network type of the interface is changed to point-to-multipoint.	-

Parameter	Description	Value
p2p	Indicates that the network type of the interface is changed to point-to-point.	-
peer-ip-ignore	Disables network segment check when IP address unnumbered is not configured for a P2P interface changed from a broadcast interface and the interface tries to establish an OSPF neighbor relationship. By default, if **peer-ip-ignore** is not specified in the command, OSPF checks the network segment of the two ends during which an OSPF neighbor relationship is to be established. Specifically, OSPF performs an AND operation on the local subnet mask and the local IP address, and on the local subnet mask and the remote IP address. An OSPF neighbor relationship can be established only when the results on the two ends are the same.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When link layer protocols remain unchanged, you can change network types and configure OSPF features to flexibly build networks.

- In the broadcast network, if there is a device that does not support multicast address, you can change the network type of the interface to NBMA.
- If the network type of the interface is NBMA, when the interface type is changed to broadcast, neighbor relationships do not need to be configured.

The condition for changing an NBMA network to broadcast network is that there must be a direct virtual circuit between any two devices, and the network must be a full mesh network.

If a network does not meet the preceding conditions, you must change the type of network to point-to-multipoint. In this manner, two indirect devices can communicate with the help of one device directly connected to the two devices. Instead of configuring the neighboring device, you can change the network type of the interface to point-to-multipoint.

If there are only two devices that run OSPF in the same network segment, the network type of an interface can be changed to p2p.

Precautions

- OSPF does not support the configuration on a null interface.
- When the network type of an interface is NBMA or the network type of an interface is changed to NBMA manually, you must run the **peer** command to configure the neighbor.
- If the network type of an OSPF interface is NBMA, OSPF does not advertise the interface's information to RSVP-TE, and TE tunnels passing through this interface fail to go Up.
- Generally, the network type of two OSPF interfaces on both ends of the link must be identical. Otherwise, the two interfaces cannot set up the neighbor relationship.

  Only when the network type of one OSPF interface is broadcast and the network type of the other OSPF interface is P2P or P2MP, the two interfaces can still set up the neighbor relationship, but cannot learn the OSPF routing information each other.

Example

Set network type of VLANIF100 to NBMA.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf network-type nbma
```

Set network type of GE0/0/1 to NBMA.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf network-type nbma
```

7.4.83 ospf p2mp-mask-ignore

Function

The **ospf p2mp-mask-ignore** command disables a device from checking the network mask on a Point-to-Multipoint (P2MP) network.

The **undo ospf p2mp-mask-ignore** command configures the device to check the network mask on a P2MP network.

By default, devices on a P2MP network do not check the network mask.

Format

ospf p2mp-mask-ignore

undo ospf p2mp-mask-ignore

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

OSPF checks the network masks carried in Hello packets. If the network mask carried in a received Hello packet is not the same as the network mask of the local device, the Hello packet is discarded.

On a P2MP network, when the mask lengths of devices are different, you can use the **ospf p2mp-mask-ignore** command on devices to disable them from checking the network mask in Hello packets. In this manner, the OSPF neighbor relationship can be established.

Prerequisites

Because P2MP is not a link layer protocol, each P2MP network is forcibly changed from a network of another type. A common P2MP network is changed from a non-fully connected non-broadcast multiple access (NBMA) network through the **ospf network-type p2mp** command.

Example

Disable a device from checking the network mask on a P2MP network.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf network-type p2mp
[HUAWEI-Vlanif100] ospf p2mp-mask-ignore
```

Disable a device from checking the network mask on a P2MP network.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf network-type p2mp
[HUAWEI-GigabitEthernet0/0/1] ospf p2mp-mask-ignore
```

7.4.84 ospf router-id auto-recover disable

Function

The **ospf router-id auto-recover disable** command disables a device from performing automatic recovery after detecting router ID conflict.

The **undo ospf router-id auto-recover disable** command enables a device to perform automatic recovery after detecting router ID conflict.

By default, a device performs automatic recovery after detecting router ID conflict.

Format

ospf router-id auto-recover disable

undo ospf router-id auto-recover disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

If router ID conflict occurs in an OSPF area, the system can define a new router ID, preventing route flapping and reducing route calculation operations. Other protocols will not go Down when the CPU usage is controlled.

NOTE

- If the automatic recovery function is enabled and a router ID conflict occurs between indirectly connected routers in one OSPF area, the system replaces the conflicted router ID with a newly calculated one. The automatic recovery function takes effect on both configured and automatically generated router IDs.
- The system can replace a router ID in a maximum of three attempts in case the router ID conflict persists.

Example

Disable a device from performing automatic recovery after detecting router ID conflict.

```
<HUAWEI> system-view
[HUAWEI] ospf router-id auto-recover disable
```

7.4.85 ospf smart-discover

Function

The **ospf smart-discover** command enables smart-discover on an interface.

The **undo ospf smart-discover** command disables smart-discover on an interface.

By default, smart-discover is disabled on interfaces.

Format

ospf smart-discover

undo ospf smart-discover

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In normal situations, devices periodically send Hello packets through OSPF interfaces. By sending Hello packets, devices set up and maintain neighbor relationships, and elect the DR and BDR on the multi-access network (broadcast or NBMA network). When setting up neighbor relationships or electing the DR and BDR on the multi-access network, interfaces send Hello packets only when the Hello timer expires. This slows down the establishment of neighbor relationships and election of the DR and BDR.

After smart-discover is configured, when the status of the neighbor relationship changes or the DR and BDR on the multi-access network changes, the device sends Hello packets to its neighbor immediately without waiting for the expiration of the Hello timer.

Procedure

On broadcast and NBMA networks, neighbor relationships can be rapidly set up and a DR and a BDR can be rapidly elected.

- When the neighbor status becomes 2-way for the first time or returns to Init from the 2-way or higher state, the smart-discover-enabled interface sends Hello packets to a neighbor without waiting for the expiration of the Hello timer when detecting that the neighbor status changes.
- When the status of the interface functioning as the DR or BDR on the multi-access network changes, the smart-discover-enabled interface actively sends Hello packets on the network segment and then participates in the DR or BDR election.

The principle of setting up adjacencies rapidly on P2P or P2MP networks is the same as that on broadcast and NBMA networks.

Configuration Impact

The interval for sending Hello packets on an interface is determined by the interval for sending Hello packets set on the interface.

Precautions

The default interval for sending Hello packets varies with the network type.

Example

Enable smart-discover on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf smart-discover
```


Enable smart-discover on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf smart-discover
```

7.4.86 ospf suppress-flapping peer

Function

The **ospf suppress-flapping peer** command configures detection parameters for OSPF neighbor relationship flapping suppression.

The **undo ospf suppress-flapping peer** command restores the default detection parameters.

By default, the detection interval of OSPF neighbor relationship flapping suppression is 60s, the suppression threshold is 10, and the interval for exiting from suppression is 120s.

Format

ospf suppress-flapping peer { **detecting-interval** *detecting-interval* | **threshold** *threshold* | **resume-interval** *resume-interval* } *

undo ospf suppress-flapping peer { **detecting-interval** *detecting-interval* | **threshold** *threshold* | **resume-interval** *resume-interval* } *

Parameters

Parameter	Description	Value
detecting-interval *detecting-interval*	Specifies a detection interval for OSPF neighbor relationship flapping suppression. Each OSPF interface on which OSPF neighbor relationship flapping suppression is enabled starts a flapping counter. If the interval between two successive neighbor status changes from Full to a non-Full state is shorter than *detecting-interval*, a valid flapping_event is recorded, and the flapping_count is incremented by 1.	The value is an integer ranging from 1 to 300, in seconds. The default value is 60s.
threshold *threshold*	Specifies the threshold of OSPF neighbor relationship flapping suppression. When the flapping-count reaches or exceeds *threshold*, flapping suppression takes effect.	The value is an integer ranging from 1 to 1000. The default value is 10.

Parameter	Description	Value
resume-interval *resume-interval*	Specifies an interval for exiting from OSPF neighbor relationship flapping suppression. If the interval between two successive neighbor status changes from Full to a non-Full state is longer than *resume-interval*, the flapping-count is reset. **NOTE** The value of *resume-interval* must be greater than that of *detecting-interval*.	The value is an integer ranging from 2 to 1000, in seconds. The default value is 120s.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To configure detection parameters for OSPF neighbor relationship flapping suppression on an interface, run the **ospf suppress-flapping peer** command. However, keeping the default configurations is recommended.

Prerequisites

OSPF neighbor relationship flapping suppression must have been enabled globally before you configure detection parameters for it. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it before you configure the detection parameters.

Example

Set the detection interval of OSPF neighbor relationship flapping suppression to 5s, the suppression threshold to 40, and the interval for exiting from suppression to 20s on VLANIF 100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf suppress-flapping peer detecting-interval 5 threshold 40 resume-interval 20
```

7.4.87 ospf suppress-flapping peer disable

Function

The **ospf suppress-flapping peer disable** command disables OSPF neighbor relationship flapping suppression on an interface.

The **undo ospf suppress-flapping peer disable** command enables OSPF neighbor relationship flapping suppression on an interface.

By default, OSPF neighbor relationship flapping suppression is enabled on all interfaces.

Format

ospf suppress-flapping peer disable

undo ospf suppress-flapping peer disable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, OSPF neighbor relationship flapping suppression is enabled on all interfaces in the same OSPF process. To disable the function on one of the interfaces, run the **ospf suppress-flapping peer disable** command.

NOTE

When an interface enters the flapping suppression state, all neighbor relationships on the interface enter the state accordingly.

Prerequisites

OSPF neighbor relationship flapping suppression must have been enabled globally before you enable the function on an interface using the **undo ospf suppress-flapping peer disable** command. By default, the function is enabled globally. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it.

Example

Disable OSPF neighbor relationship flapping suppression on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf suppress-flapping peer disable
```

7.4.88 ospf suppress-flapping peer hold-down

Function

The **ospf suppress-flapping peer hold-down** command configures the Hold-down mode and sets a duration for this mode.

The **undo ospf suppress-flapping peer hold-down** command cancels the Hold-down mode.

By default, the Hold-down mode is disabled.

Format

ospf suppress-flapping peer hold-down *interval*

undo ospf suppress-flapping peer hold-down [*interval*]

Parameters

Parameter	Description	Value
interval	Specifies a duration for the Hold-down mode.	The value is an integer ranging from 1 to 600, in seconds. The default value is 60.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Flapping suppression works in either Hold-down or Hold-max-cost mode.

- Hold-down mode: In the case of frequent flooding and topology changes during neighbor relationship establishment, interfaces prevent neighbor relationship reestablishment during Hold-down suppression, which minimizes LSDB synchronization attempts and packet exchanges.
- Hold-max-cost mode: If the traffic forwarding path changes frequently, interfaces use 65535 as the cost of the flapping link during Hold-max-cost suppression, which prevents traffic from passing through the flapping link.

Flapping suppression can also work first in Hold-down mode and then in Hold-max-cost mode.

By default, the Hold-max-cost mode takes effect. To configure the Hold-down mode and set a duration for this mode, run the **ospf suppress-flapping peer hold-down** command.

Prerequisites

OSPF neighbor relationship flapping suppression must have been enabled globally before you configure the Hold-down mode and set a duration for this mode. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it.

Example

Configure the Hold-down mode and set its duration to 200s on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf suppress-flapping peer hold-down 200
```

7.4.89 ospf suppress-flapping peer hold-max-cost disable

Function

The **ospf suppress-flapping peer hold-max-cost disable** command disables the Hold-max-cost mode.

The **undo ospf suppress-flapping peer hold-max-cost disable** command enables the Hold-max-cost mode.

By default, the Hold-max-cost mode is enabled.

Format

ospf suppress-flapping peer hold-max-cost disable

undo ospf suppress-flapping peer hold-max-cost disable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Flapping suppression works in either Hold-down or Hold-max-cost mode.

- Hold-down mode: In the case of frequent flooding and topology changes during neighbor relationship establishment, interfaces prevent neighbor relationship reestablishment during Hold-down suppression, which minimizes LSDB synchronization attempts and packet exchanges.
- Hold-max-cost mode: If the traffic forwarding path changes frequently, interfaces use 65535 as the cost of the flapping link during Hold-max-cost suppression, which prevents traffic from passing through the flapping link.

Flapping suppression can also work first in Hold-down mode and then in Hold-max-cost mode.

By default, the Hold-max-cost mode takes effect. To configure the Hold-down mode and set a duration for this mode, run the **ospf suppress-flapping peer hold-down** *interval* command.

Prerequisites

OSPF neighbor relationship flapping suppression must have been enabled globally before you configure duration for the Hold-max-cost mode. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it.

Precautions

The Hold-max-cost mode takes effect only unidirectionally. If a remote device does not support OSPF neighbor relationship flapping suppression, bidirectional traffic between the local and remote devices may travel along different paths.

Example

Disable the Hold-max-cost mode on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf suppress-flapping peer hold-max-cost disable
```

7.4.90 ospf suppress-reachability

Function

The **ospf suppress-reachability** command suppresses the advertisement of the IP addresses of a specified interface.

The **undo ospf suppress-reachability** command restores the default configuration.

By default, all OSPF interfaces advertise their IP addresses.

Format

ospf suppress-reachability [**disable**]

undo ospf suppress-reachability [**disable**]

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To suppress the advertisement of the IP addresses of a specified interface so that the IP addresses can be reused, run the **ospf suppress-reachability** command.

Configuration Impact

Suppressing the advertisement of the IP addresses of all interfaces in an OSPF process reduces the number of IP prefixes carried in LSAs and the number of OSPF routes on the network, and improves network security and OSPF performance.

Precautions

If the advertisement of the IP addresses of virtual-link interfaces is suppressed, involved virtual-link neighbors cannot reach the Full state.

The interface IP addresses whose advertisement is suppressed are not used as the forwarding addresses (FAs) of the routes imported by OSPF. As a result, external routes may fail to be calculated.

Running the **ospf suppress-reachability** command suppresses the advertisement of primary and secondary IP addresses, without affecting the advertisement of loopback interfaces' IP addresses.

Example

Suppress the advertisement of IP addresses on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf suppress-reachability
```

7.4.91 ospf timer dead

Function

The **ospf timer dead** command sets a dead interval after which an interface considers its OSPF neighbor invalid.

The **undo ospf timer dead** command restores the default dead interval of the neighbor.

By default, for a P2P or broadcast interface, the dead interval for OSPF neighbors is 40 seconds; for an NBMA or P2MP interface, it is 120 seconds.

Format

ospf timer dead *interval*

undo ospf timer dead

Parameters

Parameter	Description	Value
interval	Specifies the dead interval of the OSPF neighbors.	The value is an integer ranging from 1 to 235926000, in seconds. Setting the dead interval of an OSPF neighbor to be longer than 20s is recommended.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If an interface does not receive any Hello packet from its neighbor within the dead interval, the interface considers its neighbor invalid. The dead interval on an OSPF interface must be greater than the transmission interval of Hello messages. In addition, the dead intervals of devices on the same network segment must be the same.

By default, the dead interval of OSPF neighbors is four times the transmission interval of Hello messages.

Precautions

OSPF does not support the configuration on a null interface.

If the dead interval of an OSPF neighbor is shorter than 20s, the session may be closed. Therefore, if **dead** *interval* is shorter than 20s, the actual dead interval of an OSPF neighbor is not shorter than 20s.

Example

Set the dead interval on VLANIF100 to 60 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf timer dead 60
```

Set the dead interval on GE0/0/1 to 60 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf timer dead 60
```


7.4.92 ospf timer hello

Function

The **ospf timer hello** command sets an interval for sending Hello packets on an interface.

The **undo ospf timer hello** command restores the default value of the interval.

By default, for a P2P or broadcast interface, the interval for sending Hello packets is 10 seconds; for an NBMA or P2MP interface, it is 30 seconds.

Format

ospf timer hello *interval* [**conservative**]

undo ospf timer hello

Parameters

Parameter	Description	Value
interval	Specifies an interval for sending the Hello packet on an interface.	The value is an integer ranging from 1 to 65535, in seconds. Setting *interval* to be longer than or equal to 5s is recommended.
conservative	Indicates the conservative mode of the dead timer. If the conservative mode is configured, the value configured for the dead timer using the **ospf timer dead** command takes effect even when the value is less than 20s.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Hello packets are periodically sent on OSPF interfaces to establish and maintain neighbor relationships. A Hello packet contains information about timers, DRs, BDRs, and known neighbors.

The smaller the **hello** *interval* is, the faster the changing speed of the network topology is. The cost of routes, however, becomes greater. Ensure that the parameters of this interface and the adjacent routers are consistent.

Precautions

OSPF does not support the configuration on a null interface.

If **hello** *interval* is set but a dead interval is not set using the **ospf timer dead** command, the dead interval of an OSPF neighbor is four times the value of **hello** *interval*. If the dead interval of an OSPF neighbor is shorter than 20s, the session may be closed. Therefore, if **hello** *interval* is shorter than 5s, the actual dead interval of an OSPF neighbor is not shorter than 20s.

Example

Set the interval for sending Hello packets on VLANIF100 to 20 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf timer hello 20
```

Set the interval for sending Hello packets on GE0/0/1 to 20 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf timer hello 20
```

7.4.93 ospf timer poll

Function

The **ospf timer poll** command sets a poll interval for sending Hello packets on NBMA networks.

The **undo ospf timer poll** command restores the default poll interval.

By default, it is 120 seconds.

Format

ospf timer poll *interval*

undo ospf timer poll

Parameters

Parameter	Description	Value
interval	Specifies the poll interval for sending Hello packets.	The value is an integer ranging from 1 to 3600, in seconds.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On an NBMA network, after a device's neighbor becomes invalid, the device periodically sends Hello packets to the neighbor according to the poll interval set in the **ospf timer poll** command. The poll interval should be at least 4 times that of the **Hello** interval.

Precautions

OSPF does not support the configuration on a null interface.

Example

Set the poll interval for sending Hello packets on VLANIF100 to 130 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf timer poll 130
```

Set the poll interval for sending Hello packets on GE0/0/1 to 130 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf timer poll 130
```

7.4.94 ospf timer retransmit

Function

The **ospf timer retransmit** command sets an interval for retransmitting LSA on an interface.

The **undo ospf timer retransmit** command restores the default interval for retransmitting LSA on an interface.

By default, the interval time is 5 seconds.

Format

ospf timer retransmit *interval*

undo ospf timer retransmit

Parameters

Parameter	Description	Value
interval	Specifies interval for retransmitting LSA on an interface.	The value is an integer ranging from 1 to 3600, in seconds.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a device transmits an LSA to its neighbor, it has to wait for the ACK packet from the neighbor. If no ACK packet is received from the neighbor in the LSA retransmission interval, this LSA is retransmitted.

Do not set LSA retransmission intervals too short between adjacent routers. Otherwise, it leads to unnecessary retransmission.

Precautions

OSPF does not support the configuration on a null interface.

Example

Specify the interval for retransmitting LSAs between VLANIF100 and the adjacent device to 8 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf timer retransmit 8
```

Specify the interval for retransmitting LSAs between GE0/0/1 and the adjacent device to 8 seconds.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf timer retransmit 8
```

7.4.95 ospf trans-delay

Function

The **ospf trans-delay** command adds a transmission delay to LSAs before they are sent by an interface.

The **undo ospf trans-delay** command restores the default transmission delay of LSAs on an interface.

By default, the transmission delay is 1 second.

Format

ospf trans-delay *interval*

undo ospf trans-delay

Parameters

Parameter	Description	Value
interval	Specifies a transmission delay to be added to LSAs before they are sent by an interface.	The value is an integer ranging from 1 to 500, in seconds.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

LSAs age (increase by 1 each second) in LSDBs but do not age during transmission in the network. To ensure that the aging time of LSAs transmitted to the remote device is correct, run the **ospf trans-delay** command to add a transmission delay to LSAs before they are sent by a local interface. This configuration is especially important to low-speed networks.

Precautions

OSPF does not support the configuration on a null interface.

Example

Add a transmission delay of 3 seconds to LSAs before they are sent by a VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospf trans-delay 3
```

Add a transmission delay of 3 seconds to LSAs before they are sent by a GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ospf trans-delay 3
```

7.4.96 ospf valid-ttl-hops

Function

The **ospf valid-ttl-hops** command enables OSPF GTSM and sets a TTL value to be checked.

The **undo ospf valid-ttl-hops** command disables OSPF GTSM.

By default, OSPF GTSM is disabled.

Format

ospf valid-ttl-hops *hops* [**nonstandard-multicast**] [**vpn-instance** *vpn-instance-name*]

undo ospf valid-ttl-hops [*hops* [**nonstandard-multicast**]] [**vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
hops	Specifies a TTL value to be checked.	The value is an integer that ranges from 1 to 255. The default value is 255.
nonstandard-multicast	Specifies the GTSM configuration is also valid for multicast packets. When the **nonstandard-multicast** parameter is configured: • The TTL values of the multicast packets which will be sent are set as 255. • The received multicast packets will be checked for the TTL value 1 or in the range of [255-*hops*+1, 255].	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance. If the parameter is specified, only the TTL value of the packets in the specified VPN instance needs to be checked.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a network demanding higher security, you can enable GTSM to improve the security of the OSPF network. GTSM defends against attacks by checking the TTL value. If an attacker simulates and keeps sending OSPF unicast packets to a switch, the switch receives and directly sends the packets to the main control board for OSPF processing, without checking the validity of the packets. In this case, the switch is busy processing these packets, causing high CPU usage. GTSM protects the switches and enhances the system security by checking whether the TTL value in the IP packet header is within a pre-defined range.

The **ospf valid-ttl-hops** command is used to enable OSPF GTSM. To check the TTL value of packets that match the GTSM policy, the **vpn-instance** parameter must be specified in the command.

For example, running the **ospf valid-ttl-hops** command enables OSPF GTSM on both the public network and the private network. If you run the **ospf valid-ttl-hops 5 vpn-instance** *vpn1* command:

- OSPF GTSM is enabled on both the public network and the private network.
- The TTL value of OSPF packets in the VPN instance named *vpn1* is detected.
- The default action is performed for packets that are from the public network and other VPN instances and do not match the GTSM policy.

Precautions

- If a VPN instance is specified in the **ospf valid-ttl-hops** command and the interface is bound to the VPN instance, all the unicast packets sent to this interface are dropped when the set number of TTL hops is smaller than the actual number of hops on the network.
- If a virtual link or sham link is configured, the actual TTL value and the configured TTL value must be the same. That means that the number of virtual links or sham links that pass through the switch is calculated. Otherwise, packets sent from neighbors of a virtual link or a sham link will be dropped.
- GTSM only checks the TTL values of the packets that match the GTSM policy. If the packets do not match the GTSM policy, you can set the **pass** parameter or **drop** parameter in the **gtsm default-action** command to pass or drop these packets.
- If only a private or public network policy is configured, run the **gtsm default-action** command to set the default behavior for processing the packets unmatched with the GTSM policy to **pass** to prevent the OSPF packets of other instances from being discarded.

Example

Enable OSPF GTSM. Set the maximum number of TTL hops to 5 for the packets that can be received from the public network.

```
<HUAWEI> system-view
[HUAWEI] ospf valid-ttl-hops 5
```

7.4.97 p2mp-peer

Function

The **p2mp-peer** command sets a cost for an OSPF route to a specified neighbor on a P2MP network.

The **undo p2mp-peer** command restores the default value.

By default, the cost of an OSPF route to a neighbor in a P2MP network equals the interface cost.

Format

p2mp-peer *ip-address* **cost** *cost*

undo p2mp-peer *ip-address*

Parameters

Parameter	Description	Value
ip-address	Specifies the IP address of a neighbor on a P2MP network.	The value is in dotted decimal notation.
cost *cost*	Sets a cost for an OSPF route to a specified neighbor on a P2MP network.	The value is an integer ranging from 1 to 65535.

Views

OSPF process view

Default Level

2: Configuration level

Usage Guidelines

By default, the cost of an OSPF route to a neighbor equals the interface cost on a P2MP network. If this default cost is not suitable, run the **p2mp-peer** command to set a desired cost.

The interface cost calculation formula is as follows: Cost of an interface = Bandwidth reference value/Interface bandwidth. The integer of the calculated result is used as the interface cost, and 1 is used as the interface cost if the calculated result is smaller than 1. The bandwidth reference value is set using the **bandwidth-reference** command, with the default value being 100 Mbit/s.

Example

Set the cost to 100 for the OSPF route to the neighbor 10.1.1.1 in OSPF process 100 on a P2MP network.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] p2mp-peer 10.1.1.1 cost 100
```

7.4.98 peer (OSPF)

Function

The **peer** command sets an IP address and a DR priority for the adjacent switch on an NBMA network.

The **undo peer** command cancels the IP address of the adjacent switch on an NBMA network.

By default, the IP address and DR priority for the adjacent switch on an NBMA network are not set.

Format

peer *ip-address* [**dr-priority** *priority*]

undo peer *ip-address*

Parameters

Parameter	Description	Value
ip-address	Specifies an IP address for the adjacent switch.	The value is in dotted decimal notation.
dr-priority *priority*	Sets a DR priority for the adjacent switch.	The value of the priority is an integer that ranges from 0 to 255. By default, it is 1.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

On an NBMA network (such as an X.25 or FR network), the entire network can be fully meshed based on the mapping configuration. This means that there is a virtual link between any two devices. In this case, the network running OSPF can be considered as a broadcast network where a DR or a BDR can be selected. However, you must run the **peer** command to manually specify an IP address and a DR priority for an adjacent switch. This is necessary because it is impossible to find the adjacent devices dynamically by broadcasting Hello packets.

Example

Set the IP address of the adjacent switch to 10.1.1.1 on an NBMA network.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] peer 10.1.1.1
```

7.4.99 preference (OSPF)

Function

The **preference** command sets a preference for an OSPF route.

The **undo preference** command restores the default preference of OSPF routes.

By default, the preference of OSPF routes is 10 and that of ASE routes is 150.

Format

preference [**ase**] { *preference* | **route-policy** *route-policy-name* } *

undo preference [**ase**]

Parameters

Parameter	Description	Value
ase	Indicates the preference of an AS external route.	-
preference	Specifies the preference of an OSPF route. The smaller the preference value, the higher the preference.	The value is an integer ranging from 1 to 255.
route-policy *route-policy-name*	Specifies the name of a route policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Multiple dynamic routing protocols can run on a switch at the same time. This creates a problem of route sharing and selecting among routing protocols. The system sets a default preference for each routing protocol. If different protocols have routes to the same destination, the protocol with a higher preference is selected to forward IP packets. To set a preference for an OSPF route, run the **preference** command.

You can create a route-policy to set a preference for a specific route by setting the **route-policy** parameter in the **preference** command:

- If the **apply preference** clause is configured for the route-policy, route preference is determined as follows:
  - Route matching the route-policy: Its preference is determined by the **apply** clause.

  - Route that does not match the route-policy: Its preference is determined by the **preference** command.

  In the following example, the preference of the route matching the route-policy **abc** is set to 50 and the preference of the route that does not match the route-policy is set to 30.

```
#
route-policy abc permit node 1
 if-match cost 20
 apply preference 50
#
ospf 1
 preference 30 route-policy abc
```

- If the **apply preference** clause is not included in the route-policy, the preference of routes is set by the **preference** command.

  In the above example, if the **apply preference 50** clause is not included in the policy **abc**, the preference of all routes is set to 30.

Configuration Impact

When there are routes discovered by multiple routing protocols on the same switch, you can make the switch prefer OSPF routes by setting a high preference for them.

Precautions

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the configured priority applies to all OSPF routes.

Example

Set the preference of routes in OSPF process 100 to 150.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] preference 150
```

Set the preference of external routes in OSPF process 200 to 130.

```
<HUAWEI> system-view
[HUAWEI] ospf 200
[HUAWEI-ospf-200] preference ase 130
```

7.4.100 prefix-priority (OSPF)

Function

The **prefix-priority** command sets a convergence priority for OSPF routes.

The **undo prefix-priority** command restores the default convergence priority of OSPF routes.

By default, the convergence priority of public 32-bit host routes is **medium**, and the convergence of other OSPF routes is **low**.

Format

prefix-priority { **critical** | **high** | **medium** } **ip-prefix** *ip-prefix-name*

undo prefix-priority { **critical** | **high** | **medium** }

Parameters

Parameter	Description	Value
critical	Sets the convergence priority of OSPF routes to critical.	-
high	Sets the convergence priority of OSPF routes to high.	-
medium	Sets the convergence priority of OSPF routes to medium.	-
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **prefix-priority** command sets a convergence priority for OSPF routes according to a specified IP prefix list name. It takes effect on the public network only.

After the **prefix-priority** command is run in the OSPF view, OSPF route calculation, link-state advertisement (LSA) flooding, and LSDB synchronization can be implemented according to the configured priority, which accelerates route convergence.

Prerequisites

An IP prefix list has been created using the **ip ip-prefix** command in the system view.

Configuration Impact

When an LSA meets multiple priorities, the highest priority takes effect.

With the **prefix-priority** command, OSPF can calculate and flood LSAs, and synchronize LSDBs according to priorities. This speeds up route convergence. OSPF calculates LSAs in the sequence of intra-area routes, inter-area routes, and AS external routes. This command makes OSPF calculate the three types of routes separately according to the specified route calculation priorities. Convergence priorities are critical, high, medium, and low. To speed up the processing of LSAs with the higher priority, during LSA flooding, the LSAs need to be placed into the corresponding critical, high, medium, and low queues according to priorities.

Precautions

By default, the convergence priorities of public OSPF host routes, direct routes, static routes, and other protocol (such as BGP and RIP) routes are **medium**, **high**, **medium**, and **low** respectively. In the public network, OSPF 32-bit host routes are uniformly identified as **medium**.

Example

Set the convergence priority of OSPF routes of 10.0.0.0/8 to critical.

```
<HUAWEI> system-view
[HUAWEI] ip ip-prefix critical-prefix index 10 permit 10.0.0.0 8
[HUAWEI] ospf 1
[HUAWEI-ospf-1] prefix-priority critical ip-prefix critical-prefix
```

7.4.101 reset gtsm statistics

Function

The **reset gtsm statistics** command clears GTSM statistics.

Format

reset gtsm statistics all

Parameters

Parameter	Description	Value
all	Clears GTSM statistics.	-

Views

User view

Default Level

3: Management level

Usage Guidelines

Before collecting the GTSM statistic within a certain period, you need to clear the existing statistics.

Example

Clear GTSM statistics.

```
<HUAWEI> reset gtsm statistics all
```

7.4.102 reset ospf counters

Function

The **reset ospf counters** command resets OSPF counters.

Format

reset ospf [*process-id*] **counters** [**neighbor** [*interface-type interface-number*] [*router-id*]]

Parameters

Parameter	Description	Value
process-id	Indicates an OSPF process ID. If the parameter is not specified, all OSPF processes are restarted.	The value is an integer ranging from 1 to 65535.
neighbor	Specifies neighbor information of an interface.	-
interface-type interface-number	Specifies the type and the number of an interface.	-
router-id	Specifies the router ID of the neighbor.	The value is in dotted decimal notation.

Views

User view

Default Level

3: Management level

Usage Guidelines

When the switch restarts an OSPF process, the neighboring switch always reserves invalid LSAs. This occupies the memory of the system. These LSAs are deleted only when they expire, which happens when the LS age field in the LSA reaches 3600 seconds. After the **reset ospf** is used to restart an OSPF process and only when the router ID is changed, the switch generates an LSA in the set time and sets the LS age field to 3600 seconds. After receiving the LSA, other switches delete the LSA immediately from their LSDBs. If a switch does not send all the LSAs within the set time, other neighboring switches still store some invalid LSAs.

Clearing OSPF statistics does not affect the normal operation of OSPF services.

NOTICE

Once deleted, statistics cannot be restored. Therefore, exercise caution when deleting statistics.

Example

Reset OSPF counters.

```
<HUAWEI> reset ospf counters
```

7.4.103 reset ospf process

Function

The **reset ospf process** command restarts an OSPF process.

Format

reset ospf [*process-id*] **process** [**flush-waiting-timer** *time* | **graceful-restart**]

Parameters

Parameter	Description	Value
process-id	Indicates an OSPF process ID. If the parameter is not specified, all OSPF processes are restarted.	The value is an integer ranging from 1 to 65535.
flush-waiting-timer *time*	Specified the time when the LSA is generated. The parameter takes effect only when it is set.	The value is an integer that ranges from 1 to 40, in seconds.
graceful-restart	Indicates that graceful restart is enabled.	-

Views

User view

Default Level

3: Management level

Usage Guidelines

Usage Scenario

The **reset ospf process** command can be used to clear OSPF information to reset the board.

NOTICE

If OSPF connections are reset, OSPF neighbor relationships will be interrupted and the original information cannot be restored. Exercise caution before running the **reset ospf process** command.

Configuration Impact

After the **reset ospf process** command is used to restart OSPF, the following situations may occur:

- If a router ID is changed, a new router ID will take effect after the command is run.
- Re-elect a DR and BDR.
- OSPF configuration will not be lost after OSPF restarts.

Precautions

Configuring the **flush-waiting-timer** parameter when OSPF starts allows you to clear invalid LSAs within the set time before LSAs time out.

Whether all invalid LSAs on other switches can be deleted depends on the set time.

When the switch restarts an OSPF process, the neighboring switch always reserves invalid LSAs. This occupies the memory of the system. These LSAs are deleted only when they expire, which happens when the LS age field in the LSA reaches 3600 seconds. After the **reset ospf** is used to restart an OSPF process and only when the router ID is changed, the switch generates an LSA in the set time and sets the LS age field to 3600 seconds. After receiving the LSA, other switches delete the LSA immediately from their LSDBs. If a device does not send all the LSAs within the set time, other neighboring switches still store some invalid LSAs.

Example

Restart all OSPF processes.

```
<HUAWEI> reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]: y
```

7.4.104 reset ospf redistribution

Function

The **reset ospf redistribution** command re-imports OSPF routes.

Format

reset ospf [*process-id*] **redistribution**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process. If this parameter is not specified, all OSPF processes are restarted.	The value is an integer that ranges from 1 to 65535.

Views

User view

Default Level

3: Management level

Usage Guidelines

The **reset ospf redistribution** command re-imports OSPF routes to generate Type 5 or Type 7 LSAs.

Example

Re-import OSPF routes in OSPF process 1.

```
<HUAWEI> reset ospf 1 redistribution
```

7.4.105 reset ospf suppress-flapping peer

Function

The **reset ospf suppress-flapping peer** command forces an interface to exit from OSPF neighbor relationship flapping suppression.

Format

reset ospf *process-id* **suppress-flapping peer** [*interface-type interface-number*] [**notify-peer**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPF process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies an interface type and number.	-
notify-peer	Instructs neighbors to exit from OSPF neighbor relationship flapping suppression too.	-

Views

User view

Default Level

3: Management level

Usage Guidelines

Interfaces exit from flapping suppression in any of the following scenarios:

- The suppression timer expires.
- The corresponding OSPF process is reset.
- The **reset ospf suppress-flapping peer** command is run.
- OSPF neighbor relationship flapping suppression is disabled globally using the **suppress-flapping peer disable** command in the OSPF view.

If **notify-peer** is specified when the **reset ospf suppress-flapping peer** command is run on a device, the device sends Hello packets in which **HelloInterval** and **RouterDeadInterval** are 0s to its neighbors to instruct the neighbors to exit from OSPF neighbor relationship flapping suppression too. If the neighbors fail to receive such Hello packets, the function of **notify-peer** does not take effect. To force the neighbors to exit from OSPF neighbor relationship flapping suppression, run the **reset ospf suppress-flapping peer** command on them.

Example

Force interfaces to exit from OSPF neighbor relationship flapping suppression.

```
<HUAWEI> reset ospf 1 suppress-flapping peer
```

7.4.106 retransmission-limit

Function

The **retransmission-limit** command enables retransmission limit and sets the maximum number of retransmissions.

The **undo retransmission-limit** command disables retransmission limit.

By default, retransmission limit is disabled.

Format

retransmission-limit [*max-number*]

undo retransmission-limit

Parameters

Parameter	Description	Value
max-number	Indicates the maximum number of retransmissions.	The value is an integer that ranges from 2 to 255. The default value is 30.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **retransmission-limit** command can be used to enable Retransmission Limitation for OSPF (RL-for OSPF) to prevent dead loops caused by repeated transmissions if neighbors cannot receive packets.

Configuration Impact

The OSPF retransmission limit can be used in the following packets:

- DD packets
- LSU packets
- LSR packets

Limit the maximum number of retransmissions. If the preceding three types of packets cannot be responded within the allowed retransmission times, disconnect neighbor relationships.

Example

Enable OSPF retransmission limit and set the maximum number of retransmissions to 40.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] retransmission-limit 40
```

7.4.107 route-tag

Function

The **route-tag** command sets a tag value for imported VPN routes.

The **undo route-tag** command restores the default setting.

By default, the tag value of a VPN route is calculated based on the AS number of BGP. If no BGP is configured, the default tag value is 0.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

route-tag { *tag* | **disable** }

undo route-tag

Parameters

Parameter	Description	Value
tag	Specifies the tag value of the imported VPN routes.	The value is an integer ranging from 0 to 4294967295.
disable	Disables a device from using a VPN route tag to detect loops.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **route-tag** command can be used only in VPN scenarios to prevent loops on Type-5 and Type-7 LSAs when a CE is dual-homed to two PEs.

In a networking where a CE is dual-homed to two PEs, PE1 generates Type-5 and Type-7 LSAs based on the imported BGP routes and sends the LSAs to the CE, then to PE2. Because OSPF routes have higher priorities over BGP routes, OSPF routes will replace BGP routes on PE2, causing loops. After the **route-tag** command is

run, if the tag value of the PE and an LSA are the same, the PE will neglect the LSA, and a loop is prevented.

By default, the first two bytes of the tag value are fixed as 0xD000, while the last two bytes are the AS number of the local BGP. If a BGP AS number is greater than 65535, the default tag 0 is used. You can use the command to change the tag in this case.

Precautions

- Configuring the same VPN route tag on the PEs within the same area is recommended.
- Different OSPF processes can be configured with the same VPN route tag.

The tags set by the **route-tag** command or other commands are different only in preference.

1. The preference of the tag configured using the **import-route** command is the highest.
2. The preference of the tag configured using the **route-tag** command is medium.
3. The preference of the tag configured using the **default tag** command is the lowest.

Example

Set the route tag for OSPF process 100 to 100 in a VPN instance named **test**.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance test
[HUAWEI-vpn-instance-test] route-distinguisher 100:1
[HUAWEI-vpn-instance-test-af-ipv4] quit
[HUAWEI-vpn-instance-test] quit
[HUAWEI] ospf 100 vpn-instance test
[HUAWEI-ospf-100] route-tag 100
```

7.4.108 rfc1583 compatible

Function

The **rfc1583 compatible** command converts rules defined in RFC 2328 into rules defined in RFC 1583.

The **undo rfc1583 compatible** command converts rules defined in RFC 1583 into rules defined in RFC 2328.

By default, OSPF supports the routing rule of RFC 1583.

Format

rfc1583 compatible

undo rfc1583 compatible

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

RFC 2328 and RFC 1583 define different OSPF route selection rules. When enabling OSPF, configure the same route selection rules on all devices in the same OSPF area. For example, in scenarios where an OSPF device uses the default route selection rules defined in RFC 1583, if the other switches in the same OSPF area use route selection rules defined in RFC 2328, you need to run the **undo rfc1583 compatible** command.

Example

Converts rules defined in RFC 1583 into rules defined in RFC 2328.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] undo rfc1583 compatible
```

7.4.109 sham-hello enable (OSPF)

Function

The **sham-hello enable** command enables the sham-hello function of OSPF.

The **undo sham-hello** command disables the sham-hello function.

By default, the sham-hello feature is disabled.

Format

sham-hello enable

undo sham-hello

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

After this command is executed, the device maintains neighbor relationships through both the Hello packets and also all OSPF protocol packets. This accurately senses the existence of OSPF neighbors and stabilizes neighboring relations.

Example

Enable the sham-hello function.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] sham-hello enable
```

7.4.110 sham-link (OSPF Area)

Function

The **sham-link** command configures a sham link or sets parameter values for a sham link.

The **undo sham-link** command deletes a sham link or restores the default parameter values of a sham link.

By default, no sham link is configured for OSPF.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

sham-link *source-ip-address destination-ip-address* [[**simple** [**plain** *plain-text* | [**cipher**] *cipher-text*] | { **md5** | **hmac-md5** | **hmac-sha256** } [*key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* }]] | **authentication-null** | **keychain** *keychain-name*] | **smart-discover** | **cost** *cost* | **dead** *dead-interval* | **hello** *hello-interval* | **retransmit** *retransmit-interval* | **trans-delay** *trans-delay-interval*] *

undo sham-link *source-ip-address destination-ip-address* [[**simple** | **md5** | **hmac-md5** | **hmac-sha256** | **authentication-null** | **keychain**] | **smart-discover** | **cost** | **dead** | **hello** | **retransmit** | **trans-delay**] *

Parameters

Parameter	Description	Value
source-ip-address	Specifies the source IP address.	The value is in dotted decimal notation.
destination-ip-address	Specifies the destination IP address.	The value is in dotted decimal notation.
smart-discover	Indicates that Hello packets are sent automatically and immediately.	-
simple	Indicates simple authentication. In simple authentication, the password type is **cipher** by default. **NOTICE** Simple authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
plain	Indicates plain authentication. Only plain text can be entered, and only plain text is displayed when the configuration file is viewed. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This carries security risks. You are advised to select **cipher** to save the password in cipher text.	-
plain-text	Specifies a plain text password.	• In **simple** authentication, the value is a string of 1 to 8 characters without spaces. • In **md5**, **hmac-sha256** or **hmac-md5** authentication, the value is a string of 1 to 255 characters.

Parameter	Description	Value
cipher	Indicates cipher authentication. Either plain text or cipher text can be entered, and cipher text is displayed when the configuration file is viewed.	-
cipher-text	Specifies a cipher text password.	• In **simple** authentication, the value is a string of 1 to 8 plain text characters and 48 cipher text characters. • In **md5**, **hmac-sha256** or **hmac-md5** authentication, the value is a string of 1 to 255 plain text characters and 20 to 392 cipher text characters.
md5	Indicates MD5 authentication. **NOTICE** MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-md5	Indicates **hmac-md5** authentication. **NOTICE** HMAC-MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-sha256	Indicates **HMAC-SHA256** authentication.	-
key-id	Specifies the authentication key ID of the interface's cipher authentication. The key ID must be consistent with that of the peer.	The value is an integer that ranges from 1 to 255.
authentication-null	Indicates that no authentication is used.	-

Parameter	Description	Value
keychain	Indicates keychain authentication. **NOTE** Before configuring this parameter, run the **keychain** command to create a keychain. Then, run the **key-id**, **key-string**, and **algorithm** commands to configure a key ID, a password, and an authentication algorithm for this keychain. Otherwise, OSPF authentication will fail. Currently, only the **HMAC-MD5**, **SM3**, and **HMAC-SHA256** algorithms can be used in OSPF.	-
keychain-name	Specifies the keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.
cost *cost*	Specifies the cost of the sham link.	The value of the cost is an integer that ranges from 1 to 65535. The default value is 1.
dead *dead-interval*	Specifies the dead interval. This value must be equal to *dead-interval* of the switch that sets up a virtual link with the local switch, and must be at least four times that of *hello-interval*.	The value of the interval is an integer that ranges from 1 to 235926000, in seconds.
hello *hello-interval*	Specifies an interval for transmitting Hello packets on an interface. This value must be equal to *hello-interval* of the switch that sets up a virtual link with the local switch.	The value is an integer that ranges from 1 to 65535, in seconds.

Parameter	Description	Value
retransmit *retransmit-interval*	Specifies an interval for retransmitting the LSA packets on an interface.	The value is an integer that ranges from 1 to 3600, in seconds.
trans-delay *trans-delay-interval*	Specifies the delay in transmitting LSA packets on an interface.	The value is an integer that ranges from 1 to 3600, in seconds.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **sham-link** command can be used only in VPN scenarios.

This command can create a sham link to allow VPN traffic to be preferentially forwarded through routes within the backbone area. This prevents traffic from the same VPN in the same OSPF area from being forwarded through intra-area OSPF routes.

Before enabling neighbors of a sham link to set up adjacencies quickly, configure the **smart-discover** parameter to actively send Hello packets immediately.

Configuration Impact

After a sham link is configured between two PEs, the sham link is considered as an intra-area OSPF route. This configuration enables a route passing through an MPLS VPN backbone network to become an intra-area OSPF route, preventing VPN traffic from being transmitted through this route. A 32-bit loopback interface address is specified as the source and destination addresses of the sham link. The loopback interface must be bound to a VPN instance and advertised using BGP.

Precautions

The route to the endpoint address of a sham link cannot be advertised to the remote PE using an OSPF process in a private network. Otherwise, two routes to the endpoint address of the sham link exist on the remote PE. One route is learned from the OSPF process and the other is learned using MP-BGP. OSPF routes have higher priorities over BGP routes. As such, the remote PE selects an incorrect OSPF route. As a result, the sham link cannot be created.

Example

Create a sham link with the source address 10.1.1.1 and destination address 10.2.2.2 in a VPN instance named **test**.


```
<HUAWEI> system-view
[HUAWEI] ospf 1 vpn-instance test
[HUAWEI-ospf-1] area 1
[HUAWEI-ospf-1-area-0.0.0.1] sham-link 10.1.1.1 10.2.2.2
```

7.4.111 silent-interface (OSPF)

Function

The **silent-interface** command disables an interface from receiving and sending OSPF packets.

The **undo silent-interface** command restores the default setting.

By default, the interface is permitted to receive or send OSPF packet.

Format

silent-interface { **all** | *interface-type interface-number* }

undo silent-interface { **all** | *interface-type* [*interface-number*] }

Parameters

Parameter	Description	Value
all	Indicates all interfaces in a specified process.	-
interface-type interface-number	Specifies the interface type and interface number.	-

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To ensure that OSPF routing information is not obtained by the devices of a certain network and the local device does not receive routing update information advertised by other devices, you can run the **silent-interface** command to disable an interface from receiving and sending OSPF packets, thereby preventing routing loops.

Configuration Impact

After an OSPF interface is set to be in the silent state, the interface can still advertise its direct routes. Hello packets on the interface, however, will be blocked and no neighbor relationship can be established on the interface. This can enhance the networking adaptability of OSPF and reduce the consumption of system resources.

Example

Disable VLANIF200 from sending or receiving OSPF packets.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] silent-interface vlanif 200
```

7.4.112 spf-schedule-interval

Function

The **spf-schedule-interval** command sets an interval for OSPF to calculate routes.

The **undo spf-schedule-interval** command restores the default setting.

By default, the intelligent timer is enabled. The interval for the SPF calculation is expressed in milliseconds. The maximum interval for the SPF calculation is 10000 ms, the initial interval is 500 ms, and the Holdtime interval is 1000 ms.

Format

spf-schedule-interval { *interval1* | **intelligent-timer** *max-interval start-interval hold-interval* | **millisecond** *interval2* }

undo spf-schedule-interval

Parameters

Parameter	Description	Value
interval1	Specifies an interval for OSPF to perform the SPF calculation.	The value is an integer ranging from 1 to 10, in seconds.
intelligent-timer	Sets an interval for OSPF SPF calculation through an intelligent timer.	-
max-interval	Specifies a maximum interval for OSPF to perform the SPF calculation.	The value is an integer ranging from 1 to 120000, in milliseconds. The default value is 10000.
start-interval	Specifies the initial interval for OSPF to perform the SPF calculation.	The value is an integer ranging from 1 to 60000, in milliseconds. The default value is 500.

Parameter	Description	Value
hold-interval	Specifies the Holdtime interval for OSPF to perform the SPF calculation.	The value is an integer ranging from 1 to 60000, in milliseconds. The default value is 1000.
millisecond *interval2*	Specifies an interval for OSPF to perform the SPF calculation.	The value is an integer ranging from 1 to 10000, in milliseconds.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Based on the LSDB, the device that runs OSPF calculates the SPT with itself as the root based on the SPF arithmetic, and determines the next hop to the destination network according to the SPT. When the OSPF LSDB changes, the shortest path needs to be recalculated. Frequent network changes and continual calculation of the shortest path consume many system resources and affect the efficiency of the devices. You can configure an intelligent timer and set a proper interval for the SPF calculation to prevent the excessive consumption of device memory and bandwidth resources.

On a network where the convergence time of routes is required to be shorter, set millisecond as the unit of interval to increase the frequency of calculating routes. This increases route convergence. In other networking environments, the default value is recommended.

Configuration Impact

After this command is configured, the interval for the SPF calculation is as follows:

1. The initial interval for the SPF calculation is specified by the parameter *start-interval*.
2. The interval for the SPF calculation for the nth (n≥2) time is equal to *hold-interval* x $2^{(n-2)}$.
3. When the interval specified by *hold-interval* x $2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPF performs SPF calculation at the maximum interval until *max-interval* expires without flapping or the OSPF process is restarted.

Example

Set the interval for OSPF to calculate routes to 6 seconds.


```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] spf-schedule-interval 6
```

7.4.113 stub (OSPF area)

Function

The **stub** command sets an area to a stub area.

The **undo stub** command cancels the settings.

By default, no area is set to a stub area.

Format

stub [**no-summary** | **default-route-advertise backbone-peer-ignore**] *

undo stub

Parameters

Parameter	Description	Value
no-summary	Disables an ABR from sending non-default Type 3 LSAs to the stub area. This parameter applies only to ABRs of stub areas. If the parameter is configured on an ABR, the ABR advertises only one default Type 3 LSA (no other Type 3 LSAs) to the stub area. In this case, this area is also called a totally stub area.	-
default-route-advertise	Configures an ABR to advertise a default Type 3 LSA to the stub area. This parameter applies only to ABRs of stub areas.	-
backbone-peer-ignore	Prevents an ABR from checking the neighbor status when the ABR generates a default Type 3 LSA and advertises it to the stub area. This parameter applies only to ABRs of stub areas. If the parameter is configured on an ABR, the ABR generates a default Type 3 LSA and advertises it to the stub area as long as an interface is Up in the backbone area.	-

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

There are two configuration commands for stub areas: **stub** and **default-cost**. The stub area attribute must be configured on all switches in a stub area using the **stub** command.

The **default-cost** command takes effect only when it is configured on an ABR. The command is used to specify the cost of the default summary route transmitted by the ABR to the stub area.

- If the **stub** command is run on an ABR, without parameters specified, the ABR advertises a default Type 3 LSA and advertises it to the stub area as long as an OSPF neighbor in the full state exists in the backbone area.
- If the **stub no-summary** command is run on an ABR, the ABR is disabled from sending non-default Type 3 LSAs to the stub area connected to the ABR.
- If the **stub default-route-advertise backbone-peer-ignore** command is run on an ABR, the ABR generates a default Type 3 LSA and advertises it to the stub area as long as an interface is Up in the backbone area.

Precautions

The backbone area (area 0) cannot be configured as a stub area.

Example

Set OSPF area 1 as a stub area.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 1
[HUAWEI-ospf-100-area-0.0.0.1] stub
```

7.4.114 stub-router (OSPF)

Function

The **stub-router** command configures a stub router.

The **undo stub-router** command restores the default configuration.

By default, no device is configured as a stub router.

Format

stub-router [**on-startup** [*interval*]]

undo stub-router

Parameters

Parameter	Description	Value
on-startup [*interval*]	Specifies the interval during which a device acts as a stub router when the device is restarted or faulty. • If **on-startup** is not specified, the device is always a stub router, even if the cost of all routes advertised by the device is 65535. • If **on-startup** is specified, the device works as a stub router only when it restarts or is faulty. The hold time of the stub router state is determined by *interval* parameter. If the *interval* parameter is not configured, the default interval (500 seconds) is used.	The value is an integer that ranges from 5 to 65535, in seconds. By default, the value is 500 seconds.

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

After the **stub-router** command is configured on a device, the device informs other devices not to use this stub router to forward data by increasing the metric to 65535 of the links in the LSA that is generated by the device. As the metric is not infinite, there can still be a route to this stub router. The metric of the LSA links that are generated by the stub router is very high.

Example

Configure the device as a stub router.

```
<HUAWEI> system-view
[HUAWEI] ospf 1
[HUAWEI-ospf-1] stub-router
```

7.4.115 suppress-flapping peer disable (OSPF)

Function

The **suppress-flapping peer disable** command disables OSPF neighbor relationship flapping suppression globally.

The **undo suppress-flapping peer disable** command enables OSPF neighbor relationship flapping suppression globally.

By default, OSPF neighbor relationship flapping suppression is enabled globally.

Format

suppress-flapping peer disable

undo suppress-flapping peer disable

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

If an OSPF interface alternates between Up and Down, OSPF neighbor relationship flapping occurs on the interface. During the flapping, OSPF frequently sends Hello packets to reestablish the neighbor relationship, synchronizes LSDBs, and recalculates routes. In this process, a large number of packets are exchanged, adversely affecting neighbor relationship stability, OSPF services, and other OSPF-dependent services, such as LDP and BGP. OSPF neighbor relationship flapping suppression can address this problem by delaying OSPF neighbor relationship reestablishment or preventing service traffic from passing through flapping links.

By default, OSPF neighbor relationship flapping suppression is enabled globally. To disable this function globally, run the **suppress-flapping peer disable** command.

Example

Disable neighbor relationship flapping suppression globally.

```
<HUAWEI> system-view
[HUAWEI] ospf
[HUAWEI-ospf-1] suppress-flapping peer disable
```

7.4.116 suppress-reachability (OSPF)

Function

The **suppress-reachability** command suppresses the advertisement of the IP addresses of all interfaces in an OSPF process.

The **undo suppress-reachability** command restores the default configuration.

By default, all interfaces in an OSPF process advertise their IP addresses.

Format

suppress-reachability

undo suppress-reachability

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To suppress the advertisement of the IP addresses of all interfaces in an OSPF process so that the IP addresses can be reused, run the **suppress-reachability** command.

Configuration Impact

Suppressing the advertisement of the IP addresses of all interfaces in an OSPF process reduces the number of IP prefixes carried in LSAs and the number of OSPF routes on the network, and improves network security and OSPF performance.

Precautions

If the advertisement of the IP addresses of virtual-link interfaces is suppressed, involved virtual-link neighbors cannot reach the Full state.

The interface IP addresses whose advertisement is suppressed are not used as the forwarding addresses (FAs) of the routes imported by OSPF. As a result, external routes may fail to be calculated.

Running the **suppress-reachability** command suppresses the advertisement of primary and secondary IP addresses, without affecting the advertisement of loopback interfaces' IP addresses.

Example

Suppress the advertisement of the IP addresses of all interfaces in OSPF process 100.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] suppress-reachability
```

7.4.117 vlink-peer (OSPF area)

Function

The **vlink-peer** command creates and configures a virtual link.

The **undo vlink-peer** command deletes the virtual link or restores the default setting.

By default, no virtual link is configured for OSPF.

Format

vlink-peer *router-id* [**dead** *dead-interval* | **hello** *hello-interval* | **retransmit** *retransmit-interval* | **smart-discover** | **trans-delay** *trans-delay-interval* | [**simple** [**plain** *plain-text* | [**cipher**] *cipher-text*] | { **md5** | **hmac-md5** | **hmac-sha256** } [*key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* }] | **authentication-null** | **keychain** *keychain-name*]] *

undo vlink-peer *router-id* [**dead** | **hello** | **retransmit** | **smart-discover** | **trans-delay** | **simple** | **md5** | **hmac-md5** | **hmac-sha256** | **authentication-null** | **keychain**]

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
router-id	Specifies the switch ID of virtual link neighbor.	-
dead *dead-interval*	Specifies a dead interval. This value must be equal to *dead-interval* of the Switch that sets up a virtual link with the interface and must be at least 4 times that of *hello-interval*.	The value is an integer that ranges from 1 to 235926000, in seconds. The default value is 40 seconds.
hello *hello-interval*	Specifies an interval for transmitting Hello packets on an interface. This value must be equal to *hello-interval* value of the Switch that sets up a virtual link with the interface. The default value is 10 seconds.	The value ranges from 1 to 65535 seconds. The default value is 10 seconds.
retransmit *retransmit-interval*	Specifies an interval for retransmitting the LSA packets on an interface.	The value is an integer that ranges from 1 to 3600, in seconds. The default value is 5 seconds.
smart-discover	Automatically sends Hello packets	-

Parameter	Description	Value
trans-delay *trans-delay-interval*	Specifies the delay in transmitting LSA packets on an interface.	The value is an integer that ranges from 1 to 3600, in seconds. The default value is 1 second.
simple	Indicates simple authentication. In simple authentication, the password type is **cipher** by default. **NOTICE** Simple authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
plain	Indicates plain authentication. Only plain text can be entered, and only plain text is displayed when the configuration file is viewed. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This carries security risks. Select **cipher** to save the password in cipher text for increased security.	-
plain-text	Specifies a plain text password.	• In **simple** mode, the value is a string of 1 to 8 characters without spaces. • In **md5**, **hmac-sha256** or **hmac-md5** mode, the value is a string of 1 to 255 characters without spaces.
cipher	Indicates cipher authentication. Either plain text or cipher text can be entered, and cipher text is displayed when the configuration file is viewed.	-

Parameter	Description	Value
cipher-text	Specifies a cipher text password.	• In **simple** mode, the value is a string of 1 to 8 characters (plaintext password) or 48 characters (ciphertext password) without spaces. • In **md5**, **hmac-sha256** or **hmac-md5** mode, the value is a string of 1 to 255 (plain text password) and 20 to 392 characters (cipher text password) without spaces.
md5	Indicates MD5 authentication. In MD5 authentication, the password type is **cipher** by default. **NOTICE** MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-md5	Indicates HMAC-MD5 authentication. In HMAC-MD5 authentication, the password type is **cipher** by default. **NOTICE** HMAC-MD5 authentication carries potential security risks. As such, HMAC-SHA256 authentication is recommended.	-
hmac-sha256	Indicates HMAC-SHA256 authentication. In HMAC-SHA256 authentication, the password type is **cipher** by default.	-
key-id	Specifies the authentication key ID of the interface's cipher authentication. The key ID must be consistent with that of the peer.	The value is an integer that ranges from 1 to 255.

Parameter	Description	Value
authentication-null	Indicates that no authentication is used.	-
keychain	Indicates keychain authentication. **NOTE** Before configuring this parameter, run the **keychain** command to create a keychain. Then, run the **key-id**, **key-string**, and **algorithm** commands to configure a key ID, a password, and an authentication algorithm for this keychain. Otherwise, OSPF authentication will fail. Currently, only the **HMAC-MD5**, **SM3**, and **HMAC-SHA256** algorithms can be used in OSPF.	-
keychain-name	Specifies the keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.

Views

OSPF area view

Default Level

2: Configuration level

Usage Guidelines

Usage Guidelines

After OSPF areas are defined, OSPF route updates between non-backbone areas are transmitted through a backbone area. Therefore, OSPF requires that all non-backbone areas be directly connected to the backbone area and devices within the backbone area keep connected as well. However, these requirements may not be met due to various limitations. OSPF virtual links can be configured to solve the problem.

Follow-up Procedure

After virtual links are established, devices provided by different vendors may use different default MTUs. To ensure consistent MTUs on the devices, run the **undo**

ospf mtu-enable command to set the default MTU in DD packets sent by interfaces to 0.

NOTICE

Configuring the MTU in DD packets will cause the neighbor relationship to be re-established.

Precautions

When configuring parameters, pay attention to the following:

- A smaller **hello** value indicates faster detection of changes in network topology and higher network resource usage.
- A **retransmit** value that is too small leads to unnecessary retransmission of LSAs. On a low-speed network, set a large **retransmit** value.
- The authentication mode of a virtual link must be the same as that in the backbone area.

Example

Configure a virtual link with the peer device's router ID 10.1.1.1.

```
<HUAWEI> system-view
[HUAWEI] ospf 100
[HUAWEI-ospf-100] area 2
[HUAWEI-ospf-100-area-0.0.0.2] vlink-peer 10.1.1.1
```

7.4.118 vpn-instance-capability simple (OSPF)

Function

The **vpn-instance-capability simple** command directly calculates a route instead of conducting the routing loop detection.

The **undo vpn-instance-capability** command detects the DN bit to avoid routing loops.

By default, the routing-loop check is enabled.

Format

vpn-instance-capability simple

undo vpn-instance-capability

Parameters

None

Views

OSPF view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If OSPF VPN multi-instance is deployed on a Multi-VPN-Instance CE (MCE), routes cannot be calculated based on Type3, Type5, or Type7 LSAs with the DN bit. Therefore, OSPF detects routing loops when calculating routes. In this case, the **vpn-instance-capability simple** command is used to disable OSPF routing loop detection and enable OSPF to calculate routes based on received LSAs without checking the DN bit and route-tag in the LSAs. The route-tag is restored to the default value 1.

Prerequisites

OSPF VPN multi-instance has been deployed on an MCE using the **ospf** *process-id* **vpn-instance** *vpn-instance-name* command.

Configuration Impact

- If there is no ABR and the **vpn-instance-capability simple** command is run on an MCE, the MEC cannot become an ABR.
- After the **vpn-instance-capability simple** command is run, OSPF routes that have been imported by BGP do not carry any OSPF domain ID, route tag, or router ID.
- By default, when BGP imports an OSPF route, it uses the cost of the OSPF route plus 1 as the MED value. The MED in BGP is similar to the cost in an IGP in terms of functions. After the **vpn-instance-capability simple** command is run, BGP uses the cost of an OSPF route as the MED when it imports the OSPF route. Therefore, MED values change after the command is run, which may affect the route selection result.

Precautions

NOTE

The **undo vpn-instance-capability** command cannot be used to enable OSPF routing loop detection in a scenario without MCEs.

Example

Disable the OSPF routing-loop check.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vrf1
[HUAWEI-vpn-instance-vrf1] route-distinguisher 100:1
[HUAWEI-vpn-instance-vrf1-af-ipv4] vpn-target 3:3 export-extcommunity
[HUAWEI-vpn-instance-vrf1-af-ipv4] vpn-target 4:4 import-extcommunity
[HUAWEI-vpn-instance-vrf1-af-ipv4] quit
[HUAWEI-vpn-instance-vrf1] quit
[HUAWEI] ospf 100 vpn-instance vrf1
[HUAWEI-ospf-100] vpn-instance-capability simple
```


7.5 OSPFv3 Configuration Commands

7.5.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.5.2 abr-summary (OSPFv3 Area)

Function

The **abr-summary** command configures IPv6 route summarization on an ABR.

The **undo abr-summary** command cancels IPv6 route summarization on an ABR.

By default, IPv6 route summarization is not configured on ABRs.

Format

abr-summary *ipv6-address prefix-length* [**cost** *cost* | **not-advertise**] *

undo abr-summary *ipv6-address prefix-length*

Parameters

Parameter	Description	Value
ipv6-address	Specifies the address range of IPv6 routes to be summarized.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the prefix length.	The value is an integer that ranges from 1 to 128.
cost *cost*	Specifies the cost of a summarized route. By default, the cost of a summarized route is the maximum cost among those of routes that are summarized.	The value is an integer that ranges from 1 to 16777214.
not-advertise	Indicates that the summarized IPv6 routes that match a specified IPv6 address prefix or prefix length are not advertised.	-

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

This command applies only to the ABR of the current area. The ABR transmits only one summarized route to other areas.

An area can be configured with multiple summarized network segments. Thus OSPFv3 can summarize multiple network segments.

Example

Summarize two routes FC00:0:0:1::/64 and FC00:0:0:2::/64 of OSPFv3 area 1 into route FC00:0:0::/48 with the cost being 400, and advertise it to other areas.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 1
[HUAWEI-ospfv3-1-area-0.0.0.1] abr-summary fc00:0:0:: 48 cost 400
```

7.5.3 area (OSPFv3)

Function

The **area** command displays the OSPFv3 area view.

Format

area *area-id*

Parameters

Parameter	Description	Value
area-id	Specifies an OSPFv3 area ID.	The value can be an integer in the decimal format or in the IPv4 address format. If the value is a decimal integer, it ranges from 0 to 4294967295. If the value is an IPv4 address, it specifies the matched IP address in dotted decimal notation.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The number of devices increases with the expansion of the network. This leads to a large LSDB on every OSPFv3-enabled device on a large-scale network. Consequently, route flapping frequently occurs and thus a large number of OSPF packets are transmitted on the network, which wastes bandwidth resources.

OSPFv3 addresses the preceding problem by logically partitioning an AS into different areas.

Prerequisites

The OSPFv3 process view has been entered.

Precautions

After an AS is partitioned into different areas, not all areas are equal. The area with ID 0 is a backbone area. The backbone area is responsible for forwarding inter-area routes. In addition, the routing information between non-backbone areas must be forwarded through the backbone area.

The first time the **area** command is run, an OSPFv3 area is created, and the OSPFv3 area view is displayed. Running the **area** command later enters the OSPFv3 area view only.

Example

Enter the OSPFv3 area 0 view.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 0
[HUAWEI-ospfv3-1-area-0.0.0.0]
```

7.5.4 asbr-summary (OSPFv3)

Function

The **asbr-summary** command enables an autonomous system border router (ASBR) to summarize the routes imported by OSPFv3.

The **undo asbr-summary** command disables an ASBR from summarizing the routes imported by OSPFv3.

By default, ASBRs do not summarize the routes imported by OSPFv3.

Format

asbr-summary *ipv6-address prefix-length* [**cost** *cost* | **tag** *tag* | **distribute-delay** *dist-delay-interval* | **not-advertise**] *

undo asbr-summary *ipv6-address prefix-length*

Parameters

Parameter	Description	Value
ipv6-address prefix-length	Specifies the IPv6 address and prefix length of a summary IPv6 route.	• The IPv6 address is a 32-digit hexadecimal number in the X:X:X:X:X:X:X:X format. • The prefix length ranges from 1 to 128.
cost *cost*	Specifies the cost of a summary route. By default, for Type 1 external routes, the cost of the summarized route is the highest cost of specific routes; for Type 2 external routes, the cost of the summarized route equals the highest cost of specific routes plus 1.	The value is an integer in the range from 1 to 16777214.
tag *tag*	Indicates the tag used to control route advertisement based on routing policies.	The value is an integer in the range from 0 to 4294967295. The default value is 1.
distribute-delay *dist-delay-interval*	Specifies the delay in advertising summary routes.	The value is an integer in the range from 1 to 65535, in seconds.
not-advertise	Indicates that the summary IPv6 routes that match a specified IPv6 address prefix or prefix length are not advertised.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

An ASBR connects to more than one autonomous system (AS) and generates AS external LSAs.

You can run the **asbr-summary** command to enable an ASBR to summarize the imported routes with the same prefix into a single summary route and advertise the summary route.

Multiple summary routes can be configured on an ASBR. Therefore, multiple network segments can be summarized in OSPFv3.

Example

Enable an ASBR to summarize the routes imported by OSPFv3.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] asbr-summary fc00::1 64 cost 20 tag 100
```

7.5.5 authentication-mode (OSPFv3)

Function

The **authentication-mode** command configures an authentication mode and a password for an OSPFv3 process or area.

The **undo authentication-mode** command deletes the authentication mode and password configured for an OSPFv3 process or area.

By default, no authentication mode or password is configured for any OSPFv3 process or area.

Format

authentication-mode hmac-sha256 key-id *key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* }

authentication-mode keychain *keychain-name*

undo authentication-mode hmac-sha256 key-id *key-id*

undo authentication-mode keychain

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
hmac-sha256	Configures HMAC-SHA256 authentication.	-
key-id *key-id*	Specifies the key ID for authentication, which must be the same as the one configured at the other end.	The value is an integer that ranges from 1 to 65535.

Parameter	Description	Value
plain	Configures the plaintext password type. Only a plaintext password can be entered, and the password is displayed in plaintext in the configuration file. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	-
plain-text	Specifies a plaintext password.	The value is a string of 1 to 255 characters without spaces.
cipher	Configures the ciphertext password type. You can enter either a plaintext or ciphertext password, but the password is displayed in ciphertext in the configuration file.	-
cipher-text	Specifies a ciphertext password.	The value can be a string of 1 to 255 characters for plaintext passwords and 20 to 392 characters for ciphertext passwords without spaces.
keychain	Configures keychain authentication. **NOTE** Before configuring this parameter, you must run the **keychain** command to create a keychain. Then, run the **key-id**, **key-string**, and **algorithm** commands to configure a key ID, a password, and an authentication algorithm for this keychain. Otherwise, the OSPF authentication will fail.	-
keychain-name	Specifies a keychain name.	The value is a string of 1 to 47 case-insensitive characters without spaces.

Views

OSPFv3 view, OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Due to inherent defects and flawed implementation of the TCP/IP protocol suite, there are an increasing number of attacks, which poses greater threats on TCP/IP networks than ever before. The attacks on network devices may lead to network failures. To configure an authentication mode and a password for an OSPFv3 process or area to improve OSPFv3 network security, run the **authentication-mode** command.

Precautions

If you use area authentication, the authentication and password configurations on the interfaces of all the routers in the area must be the same.

OSPFv3 area authentication has a lower priority than OSPFv3 interface authentication.

To configure OSPFv3 interface authentication, run the **ospfv3 authentication-mode** command.

Example

Configure HMAC-SHA256 authentication for OSPFv3 process 100.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 100
[HUAWEI-ospfv3-100] authentication-mode hmac-sha256 key-id 10 cipher test
```

Configure HMAC-SHA256 authentication for OSPFv3 area 0.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 100
[HUAWEI-ospfv3-100] area 0
[HUAWEI-ospfv3-100-area-0.0.0.0] authentication-mode hmac-sha256 key-id 10 cipher test
```

7.5.6 bandwidth-reference (OSPFv3)

Function

The **bandwidth-reference** command sets a bandwidth reference value for link cost calculation.

The **undo bandwidth-reference** command restores the default setting.

By default, the bandwidth reference value of the link cost is 100 Mbit/s. That is, the link cost is 100000000/bandwidth.

Format

bandwidth-reference *value*

undo bandwidth-reference

Parameters

Parameter	Description	Value
value	Specifies a bandwidth reference value for link cost calculation.	The value is an integer that ranges from 1 to 2147483648, in Mbit/s. The default value is 100 Mbit/s.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

To set a cost for an interface that runs OSPFv3, you can run the **ospfv3 cost** command. If the link cost is not set, OSPFv3 calculates the link cost according to the link bandwidth (Cost = Reference Value/Bandwidth). The priority of the cost set using the **ospfv3 cost** command is higher than that set using the **bandwidth-reference** command.

The **bandwidth-reference** command has the same function on Eth-Trunk interfaces and physical interfaces. If the command is run on an Eth-Trunk interface, the bandwidth of the Eth-Trunk interface is the total bandwidth of all its member interfaces.

The **bandwidth** *bandwidth* command can only set an interface bandwidth obtained by the NMS from the MIB. It cannot change an interface actual bandwidth and interface cost.

Example

Set the bandwidth reference value of the link cost to 1000.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] bandwidth-reference 1000
```

7.5.7 bfd all-interfaces (OSPFv3)

Function

The **bfd all-interfaces** command enables BFD for OSPFv3 or configures BFD parameters for OSPFv3.

The **undo bfd all-interfaces** command disables BFD for OSPFv3 or deletes the configured BFD parameters for OSPFv3.

By default, BFD is not enabled or configured for OSPFv3 processes.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

bfd all-interfaces { **enable** | { **min-transmit-interval** *min-transmit-value* | **min-receive-interval** *min-receive-value* | **detect-multiplier** *detect-multiplier-value* } * }

undo bfd all-interfaces { **enable** | { **min-transmit-interval** [*min-transmit-value*] | **min-receive-interval** [*min-receive-value*] | **detect-multiplier** [*detect-multiplier-value*] } * }

Parameters

Parameter	Description	Value
enable	Enables BFD for OSPFv3.	-
min-transmit-interval *min-transmit-value*	Specifies the minimum interval for sending BFD packets to the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.

Parameter	Description	Value
min-receive-interval *min-receive-value*	Specifies the minimum interval for receiving BFD packets from the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. ● After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. ● After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *detect-multiplier-value*	Indicates the local detection multiplier.	The value is an integer in the range from 3 to 50. The default value is 3.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

- To allow OSPFv3 to quickly detect device faults and recalculate routes, run the **bfd all-interfaces** command to enable BFD and configure BFD session parameters for OSPFv3. This speeds up OSPFv3 route convergence. If an OSPFv3 neighbor relationship goes Down, the BFD session between the OSPFv3 neighbors is dynamically deleted.

After OSPFv3 establishes a BFD session, the default parameter values are used. BFD session parameters can also be modified as required.

- Actual interval at which BFD packets are transmitted on the local device = Max { configured interval *min-transmit-value* at which BFD packets are transmitted on the local device, configured interval *min-receive-value* at which BFD packets are received on the peer device }
- Actual interval at which BFD packets are received on the local device = Max { configured interval *min-transmit-value* at which BFD packets are transmitted on the peer device, configured interval *min-receive-value* at which BFD packets are received on the local device }
- Actual period for BFD detection on the local device = Actual interval at which BFD packets are received on the local device x Detection multiplier *detect-multiplier-value* configured on the peer device

For example:

- On the local device, the configured interval at which BFD packets are transmitted is 200 ms; the interval at which BFD packets are received is 300 ms; the detection multiplier is 4.
- On the peer device, the configured interval at which BFD packets are transmitted is 100 ms; the interval at which BFD packets are received is 600 ms; the detection multiplier is 5.

Then:

- On the local device, the actual interval at which BFD packets are transmitted is 600 ms calculated by using the formula max {200 ms, 600 ms}; the interval at which BFD packets are received is 300 ms calculated by using the formula max {100 ms, 300 ms}; the detection period is 1500 ms calculated by multiplying 300 ms by 5.
- On the peer device, the actual interval at which BFD packets are transmitted is 300 ms calculated by using the formula max {100 ms, 300 ms}, the actual interval at which BFD packets are received is 600 ms calculated by using the formula max {200 ms, 600 ms}, and the detection period is 2400 ms calculated by multiplying 600 ms by 4.

Prerequisites

BFD has been enabled globally using the **bfd** command.

After BFD in the OSPFv3 process is enabled using the **bfd all-interfaces enable** command, OSPFv3 establish BFD sessions only with neighbors whose status is Full.

Follow-up Procedure

- Configure a BFD session on a specified interface.

  If you need to configure BFD on a specified interface, or if global BFD for OSPFv3 is configured but you want a specific interface to detect link faults faster, run the **ospfv3 bfd** command to configure BFD on the specified interface.
- Prevent an interface from dynamically creating a BFD session.

  If you do not want to enable BFD on a specified interface, run the **ospfv3 bfd block** command to disable the interface from dynamically creating a BFD session.

Precautions

The **bfd all-interfaces** command and the **ospfv3 bfd block** command are mutually exclusive.

BFD cannot be enabled on VBDIF interfaces.

Example

Enable BFD for OSPFv3 process 1.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] bfd all-interfaces enable
```


7.5.8 default (OSPFv3)

Function

The **default** command sets default parameters for the external routes that are imported by OSPFv3. The default parameters include the cost, tag and type (Type 1 or Type 2).

The **undo default** command restores the default values of the parameters of the external routes that are imported by OSPFv3.

By default, the cost, tag value, and type of the external routes is 1, 1, and Type 2, respectively.

Format

default { **cost** *cost* | **tag** *tag* | **type** *type* } *

undo default { **cost** [*cost*] | **tag** [*tag*] | **type** [*type*] } *

Parameters

Parameter	Description	Value
cost *cost*	Specifies the default cost of the external route that is imported by OSPFv3.	The value is an integer that ranges from 1 to 16777214. By default, the value is 1.
tag *tag*	Specifies the tag value of the external route that is imported by OSPFv3.	The value is an integer that ranges from 0 to 4294967295. By default, the value is 1.
type *type*	Specifies the type of the external route that is imported by OSPFv3.	The value is an integer that ranges from 1 to 2. By default, it is 2. • 1: Type 1 external route • 2: Type 2 external route

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Since OSPFv3 can import external routes and advertise them to the entire AS, it is necessary to specify a default cost for these external routes.

If multiple OSPFv3 processes are enabled, the **default** command takes effect only on the present process.

Example

Set the default cost of the external routes that are imported by OSPFv3 to 10.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] default cost 10
```

7.5.9 default-cost (OSPFv3 Area)

Function

The **default-cost** command specifies a cost for the default route that is sent to the stub area by OSPFv3.

The **undo default-cost** command restores the default cost of the default route.

By default, the cost of the Type 3 default route transmitted to the stub area is 1.

Format

default-cost *cost*

undo default-cost

Parameters

Parameter	Description	Value
cost	Specifies the cost of the default route that is sent to the stub area by OSPFv3.	The value ranges from 0 to 16777214. By default, the value is 1.

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

The **default-cost** command applies only to the ABR that is connected to the stub area.

Before setting the cost of the default route that is transmitted to the stub area by OSPFv3, run the **stub** command to configure the local area as a stub area.

If multiple OSPFv3 processes are enabled, the **default-cost** command takes effect only on the present process.

Example

Set area 1 as the stub area and the cost of the default route that is sent to the stub area to 60.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 1
[HUAWEI-ospfv3-1-area-0.0.0.1] stub
[HUAWEI-ospfv3-1-area-0.0.0.1] default-cost 60
```

7.5.10 default-route-advertise (OSPFv3)

Function

The **default-route-advertise** command advertises default routes into an OSPFv3 routing domain.

The **undo default-route-advertise** command cancels advertising default routes.

Format

default-route-advertise [**always** | **cost** *cost* | **type** *type* | **tag** *tag* | **route-policy** *route-policy-name* [**match-any**]]*

undo default-route-advertise [**always** | **cost** [*cost*] | **type** [*type*] | **tag** [*tag*] | **route-policy** [*route-policy-name*] [**match-any**]]*

Parameters

Parameter	Description	Value
always	Generates and advertises an LSA that describes a default route, regardless of whether the local device has active default routes from processes other than the local OSPFv3 process. The switch that has **always** configured does not calculate default routes of other switches. If **always** is not specified, an LSA is generated only when there are active non-OSPFv3 default routes in the routing table of the local device.	-
cost *cost*	Specifies the cost of an advertised default route.	The value is an integer that ranges from 1 to 16777214. The default value is 1.
type *type*	Specifies the type of an external route.	The value is 1 or 2. The default value is 2. • 1: Type 1 external route • 2: Type 2 external route

Parameter	Description	Value
tag *tag*	Specifies the tag value of an advertised default route.	The value is an integer that ranges from 0 to 4294967295. The default value is 1.
route-policy *route-policy-name*	Specifies the name of a route-policy. The device advertises default routes according to the parameters configured in the route-policy when there are matched non-OSPFv3 default routes in the routing table.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
match-any	Indicates that through a route-policy, the device advertises default routes according to the parameters configured in the route-policy when there are matched routes in the routing table.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If a static default route has been configured on an OSPFv3-capable switch, OSPFv3 compares the preferences of default routes before advertising a default route. To add the default route advertised by OSPFv3 to the current routing table, ensure that the preference of the static default route is lower than that of the default route advertised by OSPFv3.

Default routes that are imported using the **import-route** command cannot be advertised. To advertise default routes, run the **default-route-advertise** command. The cost of a generated default route is determined by the **cost** and **type** parameters in the **default-route-advertise** command, and the default **tag** value is 1. If the local routing table does not contain non-OSPFv3 default routes, specify the **always** parameter to generate an ASE LSA of a default route.

A route-policy is only used to filter existing active non-OSPFv3 default routes of the local device, and match routes according to the following rules:

- If a default route matches the policy, the local device generates a default route according to the parameters configured in the route-policy, such as the cost, tag, and type.

- If a default route does not match the policy and **always** is not specified, the local device does not generate any default route.
- If a default route does not match the policy and **always** is specified, the switch generates a default route. The parameters of the generated default route are determined by the **default-route-advertise** [**always** | **cost** *cost* | **type** *type* | **tag** *tag* | **route-policy** *route-policy-name*]* command.

NOTE

The preferences of parameters of the default route configured in a route-policy are higher than those of parameters configured in the **default-route-advertise** [**always** | **cost** *cost* | **type** *type* | **tag** *tag* | **route-policy** *route-policy-name*] * command.

Precautions

Creating a route-policy before it is referenced is recommended. If a nonexistent route-policy is referenced using the command, the device advertises the default IPv6 route as long as a default IPv6 route that is not generated by the current OSPFv3 process exists in the local routing table.

Example

Advertise the ASE LSA that describes a default route into an OSPFv3 routing domain. The local device does not have a default route.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] default-route-advertise always
```

7.5.11 description (OSPFv3 Area)

Function

The **description** command configures a description for an OSPFv3 area.

The **undo description** command deletes the configured description of an OSPFv3 area.

By default, there is no description for any OSPFv3 areas.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Specifies a description for an OSPFv3 area.	The value is a string of 1 to 80 case-sensitive characters, spaces supported.

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

By configuring a description for an OSPFv3 area, you can identify different OSPFv3 areas easily.

Example

Configure a description for OSPFv3 area 0.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 0
[HUAWEI-ospfv3-1-area-0.0.0.0] description main area
```

7.5.12 description (OSPFv3)

Function

The **description** command configures a description for an OSPFv3 process.

The **undo description** command deletes the configured description of an OSPFv3 process.

By default, there is no description for any OSPFv3 processes.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Specifies a description for an OSPFv3 process.	The value is a string of 1 to 80 case-sensitive characters, spaces supported.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

By configuring a description for an OSPFv3 process, you can identify different OSPFv3 processes easily.

Example

Configure a description for OSPFv3 process 1.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] description main process
```

7.5.13 display default-parameter ospfv3

Function

The **display default-parameter ospfv3** command displays the default OSPFv3 configuration.

Format

display default-parameter ospfv3

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

This command displays only the default configurations, regardless of whether the configurations about OSPFv3 parameters are changed.

Example

Display the default OSPFv3 configuration.

```
<HUAWEI> display default-parameter ospfv3

OSPFv3 Default Values:

Process View:
-------------------------------------------------------
 Maximum ECMP Count                 : 8
 SPF Delay Time(sec)                : 5
 SPF Hold Time(sec)                 : 10
 SPF Max Time(millisec)             : 10000
 SPF Start Time(millisec)           : 500
 SPF ExpHold Time(millisec)         : 2000
```



```
  LSA Reorig Time(sec)                    : 5
  Wait Interval for stub-router on-startup   : 500
  Grace Period(sec)                       : 120
  Retransmit-Interval for Grace LSAs(sec)    : 5
  Ack wait-time for Grace LSAs(sec)          : 15
  Helper Max Grace Period Support(sec)       : 1800
  Bandwidth-Reference(Mbps)                  : 100
  Sham Link Cost                          : 1
  VPN Domain ID                           : 0
  VPN Router Tag                          : 0
  Default Metric                          : 1
  Default Tag                             : 1
  Default Type                            : 2
  Route Preference for Internal Routes       : 10
  Route Preference for External Routes       : 150
  Suppress flapping peer                  : Enable
-------------------------------------------------------

Area View:
-------------------------------------------------------
  Default Stub Cost                       : 1
  NSSA Translator Stability Index(sec)       : 40
-------------------------------------------------------

Interface View:
-------------------------------------------------------
  Hello Interval(sec)                     : 10
  NBMA Hello Interval(sec)                   : 30
  Dead Interval(sec)                      : 40
  NBMA Dead Interval(sec)                    : 120
  Poll Interval(sec)                      : 120
  Retransmit Interval(sec)                   : 5
  Transmit Delay(sec)                     : 1
  Router DR Priority                      : 1
  Suppress flapping peer                     : Enable
  Suppress flapping detect interval          : 60
  Suppress flapping detect threshold         : 10
  Suppress flapping Resume interval          : 120
  Suppress flapping mechanism                : Hold-max-cost
  Hold-max-cost interval                  : 120
```

Table 7-62 Description of the **display default-parameter ospfv3** command output

Item	Description
OSPFv3 Default Values	Default values of OSPFv3
Process View	Process view
Area View	Area view
Interface View	Interface view
Maximum ECMP Count	Maximum number of equal-cost routes
SPF Delay Time(sec)	Default delay for SPF calculation
SPF Hold Time(sec)	Hold time interval between two consecutive SPF calculations
SPF Max Time(millisec)	Maximum delay time between two consecutive SPF calculations
SPF Start Time(millisec)	Initial SPF schedule delay

Item	Description
SPF ExpHold Time(millisec)	Minimum hold time between two consecutive SPF calculations
LSA Reorig Time(sec)	Interval at which LSAs are updated
Wait Interval for stub-router on-startup	Default period during which a router remains to be a stub router during master/slave switchover
Grace Period(sec)	GR period
Retransmit-Interval for Grace LSAs(sec)	Default interval for retransmitting Grace LSAs, in seconds
Ack wait-time for Grace LSAs(sec)	Default time during which a router waits to reply the received Grace LSA with an LSAck message If the switch does not receive any Grace LSA within the default waiting time, the switch does not send the LSAck message.
Helper Max Grace Period Support(sec)	Default GR period of the helper router
Bandwidth-Reference(Mbps)	Default bandwidth reference value used to calculate the link cost, in Mbit/s
Default Metric	Default metric of the imported external route
Default Type	Default type of the imported external route
Route Preference for Internal Routes	Default preference of the internal route
Route Preference for External Routes	Default preference of the external route
Default Stub Cost	Default cost of a stub area
NSSA Translator Stability Index(sec)	Default dead time of a translator
Hello Interval(sec)	Default interval for sending Hello packets on a P2P or a broadcast network
NBMA Hello Interval(sec)	Default interval for sending Hello packets on an NBMA network
Dead Interval(sec)	Default interval for declaring a neighbor to be Down after no Hello packets are received on a P2P or broadcast network

Item	Description
NBMA Dead Interval(sec)	Default interval for declaring a neighbor to be Down after no Hello packets are received on an NBMA network
Poll Interval(sec)	Default poll interval for sending Hello packets on an NBMA network
Retransmit Interval(sec)	Default interval for retransmitting packets
Transmit Delay(sec)	Default estimated time for transmitting an LSU over this interface Before packet transmission, the estimated time needs to be added to the time limit of each LSA. The estimated time should be added to the transmission delay of the inbound interface.
Router DR Priority	Default priority of the DR
Suppress flapping peer	Whether OSPFv3 neighbor relationship flapping suppression is enabled
Suppress flapping detect interval	Default detection interval of OSPFv3 neighbor relationship flapping suppression
Suppress flapping detect threshold	Default threshold of OSPFv3 neighbor relationship flapping suppression
Suppress flapping Resume interval	Default interval for exiting from OSPFv3 neighbor relationship flapping suppression
Suppress flapping mechanism	Default mode of OSPFv3 neighbor relationship flapping suppression
Hold-max-cost interval	Default duration of the Hold-max-cost mode

7.5.14 display ospfv3

Function

The **display ospfv3** command displays brief information about OSPFv3.

Format

display ospfv3 [*process-id*]

Parameters

Parameter	Description	Value
process-id	Identifies the OSPFv3 process ID. If process-id is not specified, the brief information of all OSPFv3 processes is displayed in the ascending order of the process ID.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3** command to check brief information about OSPFv3. The information helps learn about the OSPFv3 configuration and status.

Example

Display brief information about OSPFv3.

```
<HUAWEI> display ospfv3

 Routing Process "OSPFv3 (1)" with ID 0.0.0.0
 Route Tag: 0
 Multi-VPN-Instance is not enabled
 SPF schedule delay 5 secs, Hold time between SPFs 10 secs
 LSA Origination interval 5 secs
 LSA Arrival interval 1 sec
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Stub router capability: enabled
 Number of AS-External LSA 0. AS-External LSA's Checksum Sum 0x0000
 Number of AS-Scoped Unknown LSA 0. AS-Scoped Unknown LSA's Checksum Sum 0x0000
 Number of FULL neighbors 0
 Number of Exchange and Loading neighbors 0
 Number of LSA originated 0
 Number of LSA received 0
 SPF Count          : 0
 Non Refresh LSA    : 0
 Non Full Nbr Count : 0
 Number of areas in this router is 0
```

Table 7-63 Description of the **display ospfv3** command output

Item	Description
Route Tag	Tag value for the imported VPN route.
Multi-VPN-Instance is not enabled	The current process does not support Multi-VPN instance.

Item	Description
SPF schedule delay	Delay of SPF calculation after the SPF calculation is scheduled.
Hold time between SPFs	Interval between two SPF calculation.
LSA Origination interval	Minimum interval for producing the same LSA.
LSA Arrival interval	Minimum interval for receiving the same LSA.
Default ASE parameters	Default parameters for AS external LSA: ● Metric: cost of the route ● Tag: tag value ● Type: type of LSA (1 or 2)
Stub router capability	Stub router capability. If **on-startup** is not specified when the **stub-router** command is run, information about this field is always displayed. If **on-startup** is specified when the stub-router command is run, information about this field is displayed only in the valid period during which the device keeps serving as a stub router after being restarted.
Number of AS-External LSA	Number of LSA originated for routes outside an AS.
AS-External LSA's Checksum Sum	Checksum of LSA originated for routes outside an AS.
Number of AS-Scoped Unknown LSA	Number of unknown LSA with the flooding range of an entire AS.
AS-Scoped Unknown LSA's Checksum Sum	Checksum of unknown LSA with the flooding range of an entire AS.
Number of FULL neighbors	Number of Full neighbors.
Number of Exchange and Loading neighbors	Number of neighbors in the Exchange state and neighbors in the Loading state.
Number of LSA originated	Number of generated LSAs.
Number of LSA received	Number of received LSAs.
SPF Count	Number of the events triggered by the SPF calculation.
Non Refresh LSA	Number of LSAs that are not updated in the retransmission list during GR.

Item	Description
Non Full Nbr Count	Number of neighbors in the Exchange state and neighbors in the Loading state during GR.
Number of areas in this router is	Number of areas on a switch.

7.5.15 display ospfv3 abr-summary-list

Function

The **display ospfv3 abr-summary-list** command displays information about the summarization of the routes imported by OSPFv3.

If the IP address and mask are not specified, information about the summarization of all the routes imported by OSPFv3 is displayed.

Format

display ospfv3 [*process-id*] **abr-summary-list** [*ipv6-address prefix-length*]

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
ipv6-address	Specifies the matched IPv6 address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the length of the route prefix.	The value is an integer ranging from 1 to 128.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 abr-summary-list** command to check information about the summarization of the routes imported by an ABR. If the IP address and mask are not specified, information about the summarization of all the routes imported by the ABR is displayed.

Example

Display information about the summarization of all the routes imported by OSPFv3.

```
<HUAWEI> display ospfv3 abr-summary-list
OSPFv3 Process (1)
Area ID  :  0.0.0.1
Prefix          Prefix-Len  Matched       Status
FC00:0:0:2001::  16          2[Active]     Advertise
FC00:0:0:3001::  16          0[NotActive]  Advertise
FC00:0:0:4001::  16          0[NotActive]  Advertise
```

Table 7-64 Description of the **display ospfv3 abr-summary-list** command output

Item	Description
OSPFv3 Process (1)	OSPFv3 process ID
Area ID	Area ID
Prefix	IPv6 prefix
Prefix-Len	Length of the prefix
Matched	Status of the summarized routes
Status	Advertisement status of the summarized routes: • Advertise: Advertise after the summarization. • NotAdvertise: Do not advertise after the summarization.

7.5.16 display ospfv3 area

Function

The **display ospfv3 area** command displays OSPFv3 area information.

Format

display ospfv3 [*process-id*] **area** [*area-id*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID for an OSPFv3 process.	The value is an integer ranging from 1 to 65535.

Parameter	Description	Value
area-id	Specifies the area ID for the area to be displayed.	If the value is a decimal integer, it ranges from 0 to 4294967295. If the value is an IPv4 address, it specifies the matched IP address in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 area** command to check OSPFv3 area information. The information helps determine the status and parameters of each area.

Example

Display OSPFv3 area information.

```
<HUAWEI> display ospfv3 area
OSPFv3 Process (1)

Area 0.0.0.0(BACKBONE)  InActive
Number of interfaces in this area is 1
 SPF algorithm executed 0 times
 Number of LSA 0. Checksum Sum 0x0000
 Number of Unknown LSA 0
 Area Bdr Router count: 0
 Area ASBdr Router count: 0
 Next SPF Trigger Time 500 millisecs
 Router ID conflict state : Normal
```

Table 7-65 Description of the display ospfv3 area command output

Item	Description
Number of interfaces in this area	Number of active interfaces.
SPF algorithm executed	Number of SPF calculations.
Number of LSA	Number of LSAs in this area.
Checksum	Area checksum.
Number of Unknown LSA	Number of unknown LSAs in the area.
Area Bdr Router count	Number of ABRs in the area.
Area ASBdr Router count	Number of ASBRs in the area.
Next SPF Trigger Time	Interval for SPF calculations in the area.

Item	Description
Router ID conflict state	Status of the automatic recovery function. The value can be one of the following: • Normal: The automatic recovery function is properly detecting router ID conflict. • Wait select: The automatic recovery function delays defining a new router ID if the device starts after an unexpected delay (two hours by default). • Selecting: The automatic recovery function restarts the OSPF process with the router ID and waits for the restarted OSPF process to take effect. • RtrId Changed: The automatic recovery function determines whether router ID conflict occurs after the new router ID takes effect and returns to the Normal state if no new router ID conflict is detected. • Suspend: If the maximum number of conflict times is reached, automatic recovery function does not define a new router ID any longer. The maximum number of conflict times is three by default.

7.5.17 display ospfv3 asbr-summary

Function

The **display ospfv3 asbr-summary** command displays OSPFv3 route summarization information.

Format

display ospfv3 [*process-id*] **asbr-summary** [*ipv6-address prefix-length*] [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies an OSPFVv3 process ID.	The value is an integer in the range from 1 to 65535.

Parameter	Description	Value
ipv6-address prefix-length	Specifies an IPv6 address and route prefix length. If both *ipv6-address* and *prefix-length* are not specified, all OSPFv3 route summarization information is displayed.	• The IPv6 address is a 32-digit hexadecimal number in the X:X:X:X:X:X:X:X format. • The prefix length ranges from 1 to 128.
verbose	Displays detailed information about OSPFv3 route summarization.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To view OSPFv3 route summarization information, run the **display ospfv3 asbr-summary** command.

Example

Display detailed information about OSPFv3 route summarization.

```
<HUAWEI> display ospfv3 asbr-summary verbose

      OSPFv3 Process (1)
 Total summary address count :1

          Summary Address

 Prefix        : 10::
 Prefix length : 32
 Tag           : 0 (Not Configured)
 Status        : Advertised
 Cost          : 10 (Configured)
 Delay         : (Not Configured)
 Type          : 2 (Larger than any link state path)
 The Count of Route is: 3

Destination       Mask        Protocol      Proc       Type     Metric
   10::0          64          Static        0            2        1
   10::2          128         Static        0            2        1
   10::3          128         Static        0            2        1
```

Table 7-66 Description of the **display ospfv3 asbr-summary** command output

Item	Description
OSPFv3 Process	OSPFv3 process ID.

Item	Description
Total Summary address count	Number of summary routes. To configure summary routes, run the **asbr-summary** command.
Summary Address	Detailed information about summary routes.
Prefix	Prefix of a summary route.
Prefix length	Prefix length.
Tag	Tag of a summary route.
Status	Whether summary routes are advertised: • Advertised: Summary routes will be advertised. • NotAdvertised: Summary routes will not be advertised.
Cost	Cost of a summary route.
Delay	Delay in advertising summary routes.
Type	AS external route type: • 1: Type 1 external route • 2: Type 2 external route
The Count of Route is	Number of routes that are summarized.
Destination	Destination address of a summary route.
Mask	Mask of a summary route.
Protocol	Protocol of a summary route.
Proc	Protocol process ID of a summary route.
Metric	Cost of a summary route.

7.5.18 display ospfv3 bfd session

Function

The **display ospfv3 bfd session** command displays bidirectional forwarding detection (BFD) session information of all OSPFv3 processes.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

display ospfv3 [*process-id*] **bfd session** [*interface-type interface-number*] [*neighbor-id*] [**verbose** | **all**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID for an OSPFv3 process. If the process number is not specified, information about all OSPFv3 processes is displayed in an ascending order.	The value is an integer that ranges from 1 to 65535.
interface-type interface-number	Specifies the name, type and the number of the interface.	-
neighbor-id	Specifies router ID of the neighbor.	-
verbose	Displays detailed information of the neighbor.	-
all	Displays all the BFD sessions configured in OSPFv3.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

BFD can provide millisecond-level fault detection. BFD can work with OSPFv3 to rapidly detect faults on neighboring devices, and instruct OSPFv3 to recalculate routes for correct packet forwarding. You can run the command to check BFD session parameters in OSPFv3 processes.

Example

Display BFD session information of all OSPFv3 processes.

```
<HUAWEI> display ospfv3 bfd session
```



```
* - STALE

OSPFv3 Process (1)
 Neighbor-Id      Interface                BFD-Status
 10.3.3.3         GE0/0/1                  Up
 10.1.1.1         GE0/0/2                  Up
```

Table 7-67 Description of the **display ospfv3 bfd session** command output

Item	Description
* - STALE	Stale BFD session
OSPFv3 Process	ID of the OSPFv3 process
Neighbor-Id	Neighbor ID
Interface	Interface to which this session is associated
BFD-Status	BFD session status: • Up • Down

Display the BFD session information for process 1.

```
<HUAWEI> display ospfv3 1 bfd session verbose

* - STALE

OSPFv3 Process (1)

 Neighbor-Id: 10.3.3.3
 BFD Status: Up
 Interface: GE0/0/1
 IPv6-Local-Address: FE80::201:FF:FE01:1
 IPv6-Remote-Address: FE80::225:9EFF:FEFB:BFF1
 BFD Module preferred timer values
    Transmit-Interval(ms): 100
    Receive-Interval(ms): 100
    Detect-Multiplier: 4
 OSPFv3 Module preferred timer values
    Transmit-Interval(ms): 100
    Receive-Interval(ms): 100
    Detect-Multiplier: 4
 Configured timer values
    Transmit-Interval(ms): 100
    Receive-Interval(ms): 100
    Detect-Multiplier: 4
```

Table 7-68 Description of the **display ospfv3 bfd session** command output

Item	Description
IPv6-Local-Address	Local IPv6 address to which this session is associated

Item	Description
IPv6-Remote-Address	Remote IPv6 address to which this session is associated
BFD Module preferred timer values	Actual BFD parameter values which are negotiated by BFD Module for the established session
Transmit-Interval(ms)	Minimum interval for sending BFD packets
Receive-Interval(ms)	Minimum interval between received BFD packets
Detect-Multiplier	Detection multiplier
OSPFv3 Module preferred timer values	Preferred BFD parameter values across the processes sharing the session
Configured timer values	Configured BFD parameter values for the neighbor

Display all the BFD sessions in OSPFv3.

```
<HUAWEI> display ospfv3 bfd session all

* - STALE

BFD Session (1)
 BFD Status: Up
 Interface: GE0/0/1
 IPv6-Local-Address: FE80::201:FF:FE01:1
 IPv6-Remote-Address: FE80::225:9EFF:FEFB:BFF1
 Process-Id  Neighbor-Id
     1       10.3.3.3

BFD Session (2)
 BFD Status: Up
 Interface: GE0/0/2
 IPv6-Local-Address: FE80::201:FF:FE01:1
 IPv6-Remote-Address: FE80::200:13FF:FE82:4569
 Process-Id  Neighbor-Id
     1       10.1.1.1
```

Table 7-69 Description of the **display ospfv3 bfd session all** command output

Item	Description
BFD Session	BFD session number configured in OSPFv3
Process-Id	Configured OSPFv3 process ID
Neighbor-Id	Neighbor ID

7.5.19 display ospfv3 black-box neighbor-down

Function

The **display ospfv3 black-box neighbor-down** command displays detailed information about OSPFv3 neighbors in the Down state.

Format

display ospfv3 black-box neighbor-down [*number*]

Parameters

Parameter	Description	Value
number	Specifies the number of OSPFv3 neighbors in the Down state to be displayed.	The value is an integer ranging from 1 to 10.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If *number* is not specified, the **display ospfv3 black-box neighbor-down** command displays detailed information about the latest 10 OSPFv3 neighbors in the Down state.

Example

Display the latest 2 OSPFv3 neighbors in the Down state.

```
<HUAWEI> display ospfv3 black-box neighbor-down 2
 Neighbor Router-id                : 10.11.11.99
 Interface                       : Vlanif10
 Process ID                       : 99
 Instance ID                      : 1
 Reason for Neighbor down           : Interface down
 Time at which neighbor went down    : 2010-12-28 12:45:20+08:00
 Neighbor Router-id                : 10.11.11.99
 Interface                       : Vlanif10
 Process ID                       : 99
 Instance ID                      : 1
 Reason for Neighbor down           : One-way received
 Time at which neighbor went down    : 2010-12-28 12:45:20+08:00
```


Table 7-70 Description of the **display ospfv3 black-box neighbor-down** command output

Item	Description
Neighbor Router-id	Router ID of a neighbor whose neighbor relationship becomes Down
Interface	Interface of a device involved in a neighbor relationship
Process ID	ID of the OSPFv3 process to which the neighbor relationship belongs
Instance ID	ID of the OSPFv3 instance to which the neighbor relationship belongs
Reason for Neighbor down	Reason why a neighbor relationship becomes Down: • Interface down • Dead timer expired • Sequence number mismatch • One-way received • Forced down • Bad LS request • DR or BDR has changed
Time at which neighbor went down	Date and time at which the neighbor went Down.

7.5.20 display ospfv3 error

Function

The **display ospfv3 error** command displays the OSPFv3 errors.

Format

display ospfv3 [*process-id*] **error** [**lsa**]

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID. If no OSPFv3 process number is specified, errors of all OSPFv3 processes are displayed.	The value ranges from 1 to 65535.
lsa	Display the OSPFv3 LSA errors	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

While analyzing the faults of OSPFv3, you can obtain information of errors using the command. You can then diagnose the faults of OSPFv3 according to information.

Example

Display the errors of OSPFv3.

```
<HUAWEI> display ospfv3 error
        OSPFv3 Process 1 with Router ID 0.0.0.0
              OSPFv3 Error Statistics

General Packet Errors :
 0          : Bad packet                 0          : Bad version
 0          : Bad area ID                0          : Bad virtual link
 0          : Packet too small           0          : Transmit error
 0          : Bad Instance ID            0          : Interface down
 0          : Unknown neighbor           0          : Authentication failure
 0          : AuthSeqNum mismatch        0          : AuthKeyId mismatch
 0          : IPv6: received my own packet  0       : Bad Checksum

Hello Packet Errors :
 0          : Hello timer mismatch       0          : Dead timer mismatch
 0          : Extern option mismatch     0          : Router ID confusion
 0          : Virtual neighbor unknown

DD Packet Errors :
 0          : Neighbor state low         0          : Extern option mismatch
 0          : Unknown LSA Type           0          : MTU option mismatch
 0          : Router ID confusion

LS ACK packet Errors :
 0          : Neighbor state low         0          : Bad ACK
 0          : Duplicate ACK              0          : Unknown LSA type

LS REQ packet Errors :
 0          : Neighbor state low         0          : Bad request
 0          : Empty Request

LS UPD packet Errors :
 0          : Neighbor state low         0          : LSA checksum bad
 0          : Received less recent LSA   0          : Unknown LSA type
 0          : Newer self-generate LSA
 0          : Received LSA within LSA Arrival interval

Receive Grace LSA errors :
 0          : Number of invalid LSAs
 0          : Number of policy failed LSAs
 0          : Number of wrong period LSAs

License Errors :
 0          : Max external routes reached  0          : Max NBRs reached
```


Table 7-71 Description of the **display ospfv3 error** command output

Item	Description
General Packet Errors	General packet errors.
Bad packet	The analyzed packet errors, including the verification of the length field.
Bad version	OSPFv3 version errors, that is, it is not version 3.
Bad area ID	Area ID in the received packet does not match. That is, when the area IDs in the received packets are different from the local area ID, except area 0, other area IDs are considered incorrect.
Bad virtual link	V-link receives illegal packets.
Packet too small	Length of the received packet is less than the sum of length of IP header and the length field of the packet.
Transmit error	Transmitting packets to socket fails.
Bad Instance ID	The Instance ID errors
Interface down	Times that the OSPFv3 interface goes Down.
Unknown neighbor	For NBMA, sham-link, and V-link networks, the device receives OSPFv3 packets from neighbors that are not enabled with OSPFv3.
Authentication failure	Times that the authentication fails.
AuthSeqNum mismatch	Number of packets with a sequence number different from the local one
AuthKeyId mismatch	Number of packets with key IDs different from the local one
IPv6: received my own packet	Number of IPv6 packets received from self.
Bad Checksum	Number of packets that fail in the check and calculation
Hello Packet Errors	Hello packet errors.
Hello timer mismatch	The Hello interval is not consistent.
Dead timer mismatch	The Dead interval is not consistent.
Extern option mismatch	The extension attributes of Hello packets are not consistent.
Router ID confusion	Router id is configured to the same.

Item	Description
Virtual neighbor unknown	Router id of the packet is inconsistent with that of the neighbor that is configured by the V-link.
DD Packet Errors	DD packet errors.
Neighbor state low	Cases of error: • DD packet is received but its neighbor status is lower than 2-way. • LSR packet is received but its neighbor status is lower than Exchange. • LSU packet is received but its neighbor status is lower than Exchange. • ACK packet is received but its neighbor status is lower than Exchange.
Extern option mismatch	Option of the DD packets is not matched.
Unknown LSA Type	The router receives unknown LSA type from neighbors that are not enabled with OSPFv3.
MTU option mismatch	MTU check of the interface on which OSPFv3 is enabled. The MTU value of the DD packet that is received on the interface is greater than that of this interface.
LS ACK packet Errors	LS ACK packet errors.
Bad ACK	The count of receiving incorrect ACK packets.
Duplicate ACK	The count of receiving repeat ACK packets.
LS REQ packet Errors	Title bar: LS REQ packet errors.
Bad request	BadRequest event in the protocol.
Empty Request	Number of Empty Request.
LS UPD packet Errors	LS UPD packet error.
LSA checksum bad	LSA checksum not proper
Received less recent LSA	Receiving older LSA than the local copy.
Newer self-generate LSA	LSA recently generated by self.
Received LSA within LSA Arrival interval	Received LSA within arrival interval.
Receive Grace LSA errors	Title bar: Grace LSA error.
Number of invalid LSAs	Number of LSAs that are invalid.
Number of policy failed LSAs	Number of LSAs with policy failure.

Item	Description
Number of wrong period LSAs	Number of LSAs with wrong period.
License Errors	License errors.
Max external routes reached	Number of external routes reached the maximum value.
Max NBRs reached	Number of neighbors reached the maximum value.

7.5.21 display ospfv3 graceful-restart-information

Function

The **display ospfv3 graceful-restart-information** command displays the status of OSPFv3 GR.

Format

display ospfv3 [*process-id*] **graceful-restart-information**

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process number. If the process number is not specified, brief information about all OSPFv3 processes is displayed in an ascending order.	The value is an integer that ranges from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 graceful-restart-information** command to check the status of OSPFv3 GR.

Example

Display the status of OSPFv3 GR.

```
<HUAWEI> display ospfv3 graceful-restart-information
```



```
        OSPFv3 Router with ID (10.0.0.0) (Process 1)

Graceful-restart capability      : enabled
Graceful-restart support         : planned and unplanned, strict lsa check
Grace-Period Configured          : 120 Sec
Last Restart-exit Reason         : none

Helper capability              : enabled
  Last Helper-exit Reason          : none
```

Table 7-72 Description of the **display ospfv3 graceful-restart-information** command output

Item	Description
Graceful-restart capability	Whether GR is enabled: • enabled • disabled
Graceful-restart support	GR configuration feature. • planned: indicates that only the planned GR is supported. • unplanned: indicates that only the un-planned GR is supported. • strict lsa check: indicates that the Helper supports strict external LSA check.
Grace-Period Configured	GR period.

Item	Description
Last Restart-exit Reason	Reason for the Restart device exiting from GR last time: ● none: indicates that no GR is performed. ● successful: indicates that the device is successfully restarted through GR, and then exits from GR. ● topology change: indicates that the network topology changes. ● grace lsa not recv: indicates that no Grace-LSAs are received. ● grace period exp: indicates that the GR period expires. ● process reset: indicates that a parameter of the OSPFv3 process is reset during GR. ● hold timer expire: indicates that the hold timer expires. ● interface address deleted: indicates that the interface address is deleted. ● interface down: indicates that the interface goes Down. ● DR/BDR change on network: indicates that the DR or BDR on the network changes. ● Received route uninstall notification: indicates that the message indicating that OSPFv3 routes are not installed to the RM is received.
Helper capability	Whether Helper is enabled: ● enabled ● disabled

Item	Description
Last Helper-exit Reason	Reason for the Helper device exiting from GR last time: • none: indicates that the switch never enters the Helper mode since it starts. • successful: indicates that the Helper successfully exits from GR. • topology change: indicates that the network topology changes. • interface id change: indicates that the interface ID changes. • grace period exp: indicates that GR period expires. • hold timer expire: indicates that the hold timer expires. • interface address deleted: indicates that the interface address is deleted.

7.5.22 display ospfv3 interface

Function

The **display ospfv3 interface** command displays information about OSPFv3 interfaces.

Format

display ospfv3 [*process-id*] **interface** [**area** *area-id*] [*interface-type interface-number*]

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the OSPFv3 area ID. The value of the area ID can be a decimal integer or in the IPv4 address format.	If the value is a decimal integer, it ranges from 0 to 4294967295. If the value is an IPv4 address, it specifies the matched IP address in dotted decimal notation.

Parameter	Description	Value
interface-type interface-number	Specifies the type and the number of the interface. If the interface is not specified, information of all OSPFv3 interfaces is displayed.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 interface** command to check information about OSPFv3 interfaces.

Example

Display information about OSPFv3 VLANIF 10.

```
<HUAWEI> display ospfv3 interface vlanif 10
 Vlanif10 is up, line protocol is up
 Interface ID 0x45
 Interface MTU 1500
 IPv6 Prefixes
 FE80::2E0:13FF:FE96:1600 (Link-Local Address)
 FC00:0:0:4000::1/64
 OSPFv3 Process (1), Area 0.0.0.0, Instance ID 0
     Router ID 10.5.5.5, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State Backup, Priority 1
  No designated router on this link
  No backup designated router on this link
  Timer interval configured, Hello 10, Dead 40, Wait 40, Retransmit 5
     Hello due in 00:00:08
  Neighbor Count is 1, Adjacent neighbor count is 1
  Interface Event: 2, Lsa Count: 2, Lsa Checksum: 0x9b74
  Interface Physical BandwidthHigh 0, BandwidthLow 1000000000
```

Display information about VLANIF 10 after OSPFv3 neighbor relationship flapping suppression is configured.

```
<HUAWEI> display ospfv3 interface Vlanif 10
 Vlanif10 is up, line protocol is up
  Interface ID 0x7
  Interface MTU 1500
  IPv6 Prefixes
    FE80::2E0:74FF:FE00:8201 (Link-Local Address)
    FC00:3::1/64
  OSPFv3 Process (1), Area 0.0.0.0, Instance ID 0
    Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
    Transmit Delay is 1 sec, State Backup, Priority 1
    Designated Router (ID) 4.4.4.4
    Interface Address FE80::2E0:4FF:FE09:8201
    Backup Designated Router (ID) 1.1.1.1
    Interface Address FE80::2E0:74FF:FE00:8201
    Timer interval configured, Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello due in 00:00:09
```



```
Neighbor Count is 1, Adjacent neighbor count is 1
Interface Event 3, Lsa Count 2, Lsa Checksum 0x15084
Interface Physical BandwidthHigh 0, BandwidthLow 100000000
Suppress flapping peer: enable(flapping-count: 0, threshold: 3)
```

Table 7-73 Description of the **display ospfv3 interface** command output

Item	Description
Interface ID	Interface ID
Interface MTU	MTU value
IPv6 Prefixes	IPv6 prefixes
Link-Local Address	Address of local link
OSPFv3 Process	OSPFv3 process ID
Area	Area that the interface belongs to
Instance ID	ID of the instance that the interface belongs to
Router ID	Router ID
Network Type	Network type of the interface that can be displayed as Point-to-Point, NBMA, Point-To-Multipoint and Broadcast
Cost	Cost of the interface. To set the cost of the interface in different instances, run the **ospfv3 cost** command.
Transmit Delay	Transmission delay
State	State of the interface
Priority	Priority of the interface
Designated Router (ID)	Designated router on the link
Backup Designated Router (ID)	Backup designated router on the link

Item	Description
Timer interval configured	Configured interval • Hello: Interval for sending Hello packets. To set the interval for sending Hello packets on the interface instance, run the **ospfv3 timer hello** command. • Dead: Interval of the Dead timer. To set the dead interval of the OSPFv3 neighbor of the instance on the interface, run the **ospfv3 timer dead** command. • Wait: Interval of the Wait timer • Retransmit: Retransmission interval. To set interval for retransmitting the LSA on the interface instance, run the **ospfv3 timer retransmit** command.
Hello due in	Time remaining to send Hello packet
Neighbor Count	Number of neighbors
Adjacent neighbor count	Number of adjacency
Interface Event	Number of events on this interface
Lsa Count	Total number of LSA in the interface scope database
Lsa Checksum	Checksum of all the LSAs in the database
Interface Physical BandwidthHigh	Maximum physical bandwidth on the interface
BandwidthLow	Minimum physical bandwidth on the interface

Item	Description

Suppress flapping peer	Status of OSPFv3 neighbor relationship flapping suppression:
	● enable: OSPFv3 neighbor relationship flapping suppression is enabled.
	– flapping-count: number of valid flapping_events
	If the interval between two successive neighbor status changes from Full to a non-Full state is shorter than *detecting-interval*, a valid flapping_event is recorded, and the flapping_count is incremented by 1. To change *detecting-interval*, run the **ospfv3 suppress-flapping peer detecting-interval** *detecting-interval* command.
	– threshold: flapping suppression threshold
	When the flapping_count reaches or exceeds **threshold**, flapping suppression takes effect. To configure the threshold, run the **ospfv3 suppress-flapping peer threshold** *threshold* command.
	● disable: OSPFv3 neighbor relationship flapping suppression is disabled. In this case, the following information is displayed, without flapping-count or threshold:
	`Suppress flapping peer: disable`
	● hold-down: OSPFv3 neighbor relationship flapping suppression works in Hold-down mode. In this case, an example of the displayed information is as follows:
	`Suppress flapping peer: hold-down(start: 2016-01-29 12:09:36-08:00, remain-interval: 500s)`
	– start: time when the flapping suppression started
	– remain-interval: remaining time of the flapping suppression
	● hold-max-cost: OSPFv3 neighbor relationship flapping suppression works in Hold-max-cost mode. In this case, an example of the displayed information is as follows:
	`Suppress flapping peer: hold-max-cost(start: 2016-01-29 12:09:36-08:00, remain-interval: 476s)`

7.5.23 display ospfv3 lsdb

Function

The **display ospfv3 lsdb** command displays the OSPFv3 LSDB.

Format

display ospfv3 [*process-id*] **lsdb** [**area** *area-id*] [**originate-router** *advertising-router-id* | **self-originate**] [{ **router** | **network** | **inter-router** [**asbr-router** *asbr-router-id*] | { **inter-prefix** | **nssa** } [*ipv6-address prefix-length*] | **link** | **intra-prefix** | **grace** } [*link-state-id*]]

display ospfv3 [*process-id*] **lsdb** [**originate-router** *advertising-router-id* | **self-originate**] **external** [*ipv6-address prefix-length*] [*link-state-id*]

display ospfv3 [*process-id*] **lsdb statistics**

Parameters

Parameter	Description	Value
process-id	Identifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the OSPFv3 area ID. The value of the area ID can be a decimal integer or in the IPv4 address format.	If the value is a decimal integer, it ranges from 0 to 4294967295. If the value is an IPv4 address, it specifies the matched IP address in dotted decimal notation.
external	Displays the AS-external LSA in the LSDB. If the **area** parameter is selected, the **external** parameter cannot be selected.	-
grace	Displays the Grace-LSA in the LSDB.	-
inter-prefix	Displays the Inter-area-prefix LSA in the LSDB.	-
nssa	Displays the NSSA in the LSDB.	-
inter-router	Displays the Inter-area-router LSA in the LSDB.	-

Parameter	Description	Value
intra-prefix	Displays the Intra-area-prefix LSA in the LSDB.	-
link	Displays the Link-local LSA in the LSDB.	-
network	Displays the Network-LSA in the LSDB.	-
router	Displays the Router-LSA in the LSDB.	-
link-state-id	Indicates the link state ID.	The value is in dotted decimal notation.
originate-router *advertising-router-id*	Specifies the Router ID of the switch that advertises LSA packet.	The value is in dotted decimal notation.
asbr-router *asbr-router-id*	Specifies the router ID of the ASBR.	The value is in dotted decimal notation.
self-originate	Displays LSAs in the LSDB that are advertised by the local switch.	-
ipv6-address prefix-length	Specifies the IPv6 destination address and the length of the prefix.	*ipv6-address* is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X. *prefix-length* is an integer ranging from 0 to 128.
statistics	Indicates the statistics of the LSDB.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 lsdb** command to check the OSPFv3 LSDB.

Example

Display the OSPFv3 LSDB.

```
<HUAWEI> display ospfv3 lsdb
* indicates STALE LSA

         OSPFv3 Router with ID (10.5.5.5) (Process 1)
            Link-LSA (Interface Vlanif1002)

Link State ID   Origin Router    Age   Seq#       CkSum  Prefix
0.0.0.71        10.2.2.2         0995  0x8000000c 0xa44c      1
0.0.0.69        10.5.5.5         0621  0x8000000f 0x20ff      1

            Router-LSA (Area 0.0.0.0)

Link State ID   Origin Router    Age   Seq#       CkSum    Link
0.0.0.0         10.2.2.2         0565  0x8000008b 0xa54f     1
0.0.0.0         10.5.5.5         0616  0x80000014 0x243c     1

            Network-LSA (Area 0.0.0.0)

Link State ID   Origin Router    Age   Seq#       CkSum
0.0.0.71        10.2.2.2         0565  0x8000000a 0x585c

            Inter-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   Origin Router    Age   Seq#       CkSum
0.0.0.1         10.2.2.2         0320  0x8000000c 0x6a0f
0.0.0.2         10.2.2.2         1637  0x8000000b 0xe1a3
0.0.0.1         10.5.5.5         0616  0x8000000d 0x81f7

            Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   Origin Router    Age   Seq#       CkSum  Prefix  Reference
0.0.0.1         10.2.2.2         0564  0x8000000e 0x79a6      1  Network-LSA

            Router-LSA (Area 0.0.0.1)

Link State ID   Origin Router    Age   Seq#       CkSum    Link
0.0.0.0         10.5.5.5         0621  0x8000000b 0x8c84     0

            Inter-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   Origin Router    Age   Seq#       CkSum
0.0.0.1         10.5.5.5         0621  0x8000000a 0x8f08
0.0.0.2         10.5.5.5         0612  0x8000000a 0x105d
0.0.0.3         10.5.5.5         0612  0x8000000a 0x85f2

            Intra-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   Origin Router    Age   Seq#       CkSum  Prefix  Reference
0.0.0.1         10.5.5.5         0621  0x8000000b 0xab89      1  Router-LSA

            Router-LSA (Area 0.0.0.3)

Link State ID   Origin Router    Age   Seq#       CkSum    Link
0.0.0.0         10.5.5.5         0621  0x8000000b 0x8092     0

            Inter-Area-Prefix-LSA (Area 0.0.0.3)

Link State ID   Origin Router    Age   Seq#       CkSum
0.0.0.1         10.5.5.5         0617  0x8000000b 0x85f5
0.0.0.2         10.5.5.5         0621  0x8000000a 0x8511
0.0.0.3         10.5.5.5         0612  0x8000000a 0x0666
0.0.0.4         10.5.5.5         0612  0x8000000a 0x7bfb

            NSSA-external-LSA (Area 0.0.0.3)
```



```
Link State ID   Origin Router    Age   Seq#       CkSum  Type
0.0.0.1         10.5.5.5         0626  0x8000000a 0xe72a E2
```

Table 7-74 Description of the **display ospfv3 lsdb** command output

Item	Description
Link-LSA	Description of the Link-LSA
Router-LSA	Description of the Router-LSA
Network-LSA	Description of the Network-LSA
Inter-Area-Prefix-LSA	Description of the Inter-Area-Prefix-LSA
NSSA-external-LSA	Description of the NSSA-external-LSA
Intra-Area-Prefix-LSA	Description of the Intra-Area-Prefix-LSA
AS-External-LSA	Description of the AS-external-LSA
Link State ID	Link state ID in LSA header
Origin Router	switch that generates LSA
Age	Aging time of LSA
Seq#	Sequence number of LSA (from LSA header)
CkSum	Checksum of LSA
Link	Number of links in router LSA
Prefix	Number of prefixes in Link LSA
Reference	The Reference for Intra Area Prefix LSA, including Router-LSA and Network-LSA
Type	The type of the external routes: • E1: Type 1 external route • E2: Type 2 external route

Display the Link-local LSA in the LSDB.

```
<HUAWEI> display ospfv3 lsdb link
        OSPFv3 Router with ID (10.2.2.2) (Process 1)

            Link-LSA (Interface Vlanif10)
 LS age: 11
 LS Type: Link-LSA
 Link State ID: 0.0.2.6
 Originating Router: 10.2.2.2
 LS Seq Number: 0x80000002
 Retransmit Count: 0
 Checksum: 0xEFFA
 Length: 56
 Priority: 1
 Options: 0x000013 (-|R|-|-|E|V6)
 Link-Local Address: FE80::1441:0:E213:1
```



```
Number of Prefixes: 1

 Prefix: 2000:1::/64
  Prefix Options: 0 (-|-|-|-|-)
```

Table 7-75 Description of the **display ospfv3 lsdb link** command output

Item	Description
LS age	Aging time of LSA
LS Type	Type of LSA: • Router-LSA • Network-LSA • Inter-Area-Prefix-LSA • Inter-Area-Router-LSA • AS-external-LSA • Link-LSA • Intra-Area-Prefix-LSA • NSSA-LSA
Link State ID	Link state ID in LSA header
Originating Router	Router that generates LSA
LS Seq Number	Sequence number of LSA (from LSA header)
Retransmit Count	Total number of nodes in the retransmission list
Checksum	Checksum of LSA
Length	Length of the LSA
Priority	Priority of the interface that corresponds to the link
Options	Value of options of the link
Link-Local Address	Link local address
Number of Prefixes	Number of IPv6 prefixes in the LSA
Prefix	IPv6 prefix
Prefix Options	Value of options of the prefix

7.5.24 display ospfv3 next-hop

Function

The **display ospfv3 next-hop** command displays the next-hop routing table of OSPFv3.

Format

display ospfv3 [*process-id*] **next-hop**

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 next-hop** command to check the next-hop routing table of OSPFv3.

Example

Display the next-hop routing table of OSPFv3 process 1.

```
<HUAWEI> display ospfv3 1 next-hop
OSPFv3 Process (1)
Neighbor-Id     Next-Hop                        Interface    RefCount
10.3.3.9        FE80::2E0:FCFF:FE01:814F          Vlanif10     1
```

Table 7-76 Description of the **display ospfv3 next-hop** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process
Neighbor-Id	Neighbor ID
Next-Hop	Next hop address
Interface	Egress
RefCount	Number of routes using the next hop

7.5.25 display ospfv3 path

Function

The **display ospfv3 path** command displays the paths to a destination address.

Format

display ospfv3 [*process-id*] **path**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPFv3 process.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the process ID is not specified, paths in all processes are displayed.

Example

Display all OSPFv3 paths.

```
<HUAWEI> display ospfv3 path
OSPFv3 Process (1)
Type       AreaId/LSACost   PathCost
Intra      0.0.0.1          2
Inter      0.0.0.1          2
External   0.0.0.10          1
```

Table 7-77 Description of the **display ospfv3 path** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process
Type	Path type: • Intra: indicates the intra-area routes. • Inter: indicates the inter-area routes. • External: indicates the AS-external routes.
AreaId/LSACost	• Area ID in case of an intra-area or inter-area path. • LSA cost in case the path type is External. • Cost in case of a path to a destination outside an AS
PathCost	Cost of the path

7.5.26 display ospfv3 peer

Function

The **display ospfv3 peer** command displays the qualified OSPFv3 neighbor.

Format

display ospfv3 [*process-id*] [**area** *area-id*] **peer** [*interface-type interface-number* | *neighbor-id*] [**verbose**]

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the neighbor of the specific area.	*area-id* can be a decimal integer or in IPv4 address format. The integer ranges from 0 to 4294967295.
interface-type interface-number	Indicates the type and the number of the interface.	-
neighbor-id	Specifies router ID of the neighbor.	The value is in dotted decimal notation.
verbose	Displays detailed information of the neighbor of each area.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 peer** command to check qualified OSPFv3 neighbor information.

Example

Display neighbor information about OSPFv3 process 1 on Vlanif10.

```
<HUAWEI> display ospfv3 1 peer vlanif 10
```



```
OSPFv3 Process (1)
OSPFv3 Area (0.0.0.0)
Neighbor ID     Pri  State           Dead Time Interface          Instance ID
10.1.1.1          1  Full/DR         00:00:30  Vlanif10                 0
```

Display the detailed neighbor information of OSPFv3 process 100 on Vlanif10.

```
<HUAWEI> display ospfv3 100 peer vlanif 10 verbose
OSPFv3 Process (100)
 Neighbor 10.1.1.1 is Full, interface address FE80::3D43:0:8C14:1
   In the area 0.0.0.1 via interface Vlanif10
   DR Priority is 1 DR is 0.0.0.0 BDR is 0.0.0.0
   Options is 0x000013 (-|R|-|-|E|V6)
   Dead timer due in 00:00:29
   Neighbour is up for 18:40:40
   Database Summary Packets List 0
   Link State Request List 0
   Link State Retransmission List 0
   Neighbour Event: 6
   Neighbour If Id: 0x1017
```

Table 7-78 Description of the **display ospfv3 peer** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process.
OSPFv3 Area	ID of the OSPFv3 area.
Neighbor ID	Neighbor ID.
Pri	Priority of the neighbor switch.
State	State of the neighbor / DR or BDR. State of the neighbor: includes Down, Attempt, Init, 2-Way, ExStart, Exchange, Loading, and Full. DR or BDR: includes DR and Backup.
Dead Time	Dead time of the neighbor.
Interface	Interface that connects with the neighbor.
Instance ID	ID of the instance that the neighbor belongs to.
Neighbor	Neighbor ID.
interface address	Link address of the neighbor interface.
In the area 0.0.0.1 via interface Vlanif10	Vlanif10 of the neighbor is in Area 1.
DR Priority	Priority of DR.
DR	Designated router.
BDR	Backup designated router.
Options	Options of the neighbor.
Dead timer due in 00:00:29	The time before the neighbor dies.

Item	Description
Neighbour is up for 18:40:40	The up time of the neighbour.
Database Summary Packets List	Number of packets in the database summary list of the neighbor.
Link State Request List	Number of LSAs in the request list of the neighbor.
Link State Retransmission List	Number of LSAs in the retransmission list of the neighbor.
Neighbour Event	Number of times NFSM (Neighbor Finite State Machine) state changed.
Neighbour If Id	Interface ID of the neighbor.

7.5.27 display ospfv3 request-list

Function

The **display ospfv3 request-list** command displays the statistics of the request list of OSPFv3.

Format

display ospfv3 [*process-id*] **request-list** [[**area** *area-id* | **interface** *interface-type interface-number* | **peer** *router-id*]* | **statistics**]

Parameters

Parameter	Description	Value
process-id	Specifies the process ID of OSPFv3.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the OSPFv3 area ID.	The integer ranges from 0 to 4294967295. The value of the area ID can be a decimal integer or in the IPv4 address format.
interface *interface-type interface-number*	Specifies the type and number of the interface.	-
peer *router-id*	Indicates the Router ID of a neighbor.	It is in dotted decimal notation.

Parameter	Description	Value
statistics	Displays the statistics of the request list of OSPFv3.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The display of the command is used to diagnose faults.

Example

Display the statistics of the request list in OSPFv3.

```
<HUAWEI> display ospfv3 request-list
   OSPFv3 Process (1)
   Interface  FC00:0:0:1::2        Area-ID   0.0.0.1
   Nbr-ID   10.254.255.255
   LS-Type              LS-ID          AdvRouter       SeqNum       Age
   Link-LSA             0.0.2.2        10.254.255.255  0x80000001   151
   Router-LSA           0.0.0.0        10.254.255.255  0x80000005   46
   Intra-Area-Prefix-LSA 0.0.0.1       10.254.255.255  0x80000001   45
   AS-External-LSA      0.0.0.1        10.254.255.255  0x80000001   113
   AS-External-LSA      0.0.0.2        10.254.255.255  0x80000001   113
```

Table 7-79 Description of the **display ospfv3 request-list** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process
Interface	Interface that connects with the neighbor
Area-ID	ID of the area that the interface belongs to
Nbr-ID	Router ID of the neighbor

Item	Description
LS-Type	LSA type: • Link-LSA: is generated by the switch for each link and is advertised in the local link. • Router-LSA: is generated by each switch and advertised in the area whether the switch resides. It describes the link state and cost of the switch. • Intra-Area-Prefix-LSA: – An LSA generated on a router describes the IPv6 address prefix associated with the Router LSA. – An LSA generated on a DR describes the IPv6 address prefix associated with the network LSA. • AS-External-LSA: is generated by the ASBR and is advertised to the entire AS (except the stub area). It describes the routes that reach other ASs.
LS-ID	Link State ID (from LSA header)
AdvRouter	Advertising router (from LSA header)
SeqNum	Sequence number (from LSA header)
Age	Aging time (from LSA header)

7.5.28 display ospfv3 retrans-list

Function

The **display ospfv3 retrans-list** command displays the statistics of the retransmission list in OSPFv3.

Format

display ospfv3 [*process-id*] **retrans-list** [[**area** *area-id* | **interface** *interface-type interface-number* | **peer** *router-id*]* | **statistics**]

Parameters

Parameter	Description	Value
process-id	Specifies the process ID of OSPFv3.	The value is an integer ranging from 1 to 65535.

Parameter	Description	Value
area *area-id*	Specifies the OSPFv3 area ID.	The integer ranges from 0 to 4294967295. The value of the area ID can be a decimal integer or in the IPv4 address format.
interface *interface-type interface-number*	Specifies the type and number of the interface.	-
peer *router-id*	Indicates the Router ID of a neighbor.	It is in dotted decimal notation.
statistics	Displays the statistics of the retransmission list in OSPFv3.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The display of the command is used to diagnose faults.

Example

Display the statistics of the retransmission list in OSPFv3.

```
<HUAWEI> display ospfv3 retrans-list

    OSPFv3 Process (1)
    Interface  FC00:0:0:1::2        Area-ID   0.0.0.1
    Nbr-ID   10.22.2.255
    LS-Type              LS-ID          AdvRouter       SeqNum       Age
    Network-LSA          0.0.2.2        10.254.255.255  0x80000002   1
    Intra-Area-Prefix-LSA 0.0.0.2       10.254.255.255  0x80000002   1
```

Table 7-80 Description of the **display ospfv3 retrans-list** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process
Interface	Interface that connects with the neighbor
Area-ID	ID of the area that the interface belongs to

Item	Description
Nbr-ID	Router ID of the neighbor
LS-Type	LSA type: • Network-LSA: is generated by the DR of a broadcast network or an NBMA network and is advertised in the area where the DR resides. It describes the link state of the network segment. • Intra-Area-Prefix-LSA: – An LSA generated on a router describes the IPv6 address prefix associated with the Router LSA. – An LSA generated on a DR describes the IPv6 address prefix associated with the network LSA.
LS-ID	Link State ID (from the LSA header)
AdvRouter	Advertising router (from LSA header)
SeqNum	Sequence number (from LSA header)
Age	Aging time (from LSA header)

7.5.29 display ospfv3 routing

Function

The **display ospfv3 routing** command displays the OSPFv3 routing table.

Format

display ospfv3 [*process-id*] **routing** [*ipv6-address prefix-length* | **abr-routes** | **asbr-routes** | **intra-routes** | **inter-routes** | **ase-routes** | **nssa-routes** | [**statistics**] [**uninstalled**]]

Parameters

Parameter	Description	Value
process-id	Identifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
ipv6-address	Indicates the IPv6 address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
prefix-length	Specifies the length of the prefix.	The value is an integer ranging from 0 to 128.
abr-routes	Displays the routing table of ABR in OSPFv3.	-
asbr-routes	Displays routing table of ASBR in OSPFv3.	-
intra-routes	Displays the statistics of OSPFv3 routes in an area.	-
inter-routes	Displays the statistics of OSPFv3 routes between areas.	-
ase-routes	Displays the statistics of OSPFv3 routes outside an AS.	-
nssa-routes	Displays the statistics of OSPFv3 routes in an NSSA.	-
statistics	Displays the statistics of all routing tables.	-
uninstalled	Displays the routes are not installed into the RM.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

By using this command with different parameters, you can view the routes of a specified interface or next hop.

The command output can help you troubleshoot OSPFv3 faults.

Example

Display the OSPFv3 routing table.

```
<HUAWEI> display ospfv3 routing
Codes : E2 - Type 2 External, E1 - Type 1 External, IA - Inter-Area,
        N - NSSA, U - Uninstalled, D - Denied by Import Policy
```



```
OSPFv3 Process (1)
  Destination                                  Metric
    Next-hop
    FC00:0:0:3::1/128                             0
     directly connected, LoopBack1
 IA FC00:0:0:4::1/128                             1
      via FE80::200:FF:FE00:DC00, Vlanif1002
 IA FC00:0:0:5::1/128                             1
      via FE80::200:FF:FE00:DC00, Vlanif1002
    FC00:0:0:6::/64                                1
     directly connected, Vlanif1002
```

Table 7-81 Description of the **display ospfv3 routing** command output

Item	Description
Codes	The following information provides the full spelling and explanation: • E2: Type 2 external route • E1: Type 1 external route • IA: Inter-area route • N: NSSA route • U: OSPFv3 routes that are not advertised to the IPv6 routing table • D: OSPFv3 routes denied by import route policy
OSPFv3 Process	OSPFv3 process
Destination	Destination address
Metric	Metric to the destination
Next-hop	Next hop to the destination

Displays the routes are not installed into the RM.

```
<HUAWEI> display ospfv3 routing uninstalled
Codes : E2 - Type 2 External, E1 - Type 1 External, IA - Inter-Area,N - NSSA, U - Uninstalled

OSPFv3 Process (1)
  Destination                                  Metric
    Next-hop
 N    via FE80:1000:2000:3000:4000:5000:6000:7000, Vlanif10
 U E2 FC00:0:0:1::1/128                            5
```

Table 7-82 Description of the **display ospfv3 routing uninstalled** command output

Item	Description
N	NSSA

Display statistics about the OSPFv3 routes that are not installed into the RM.


```
<HUAWEI> display ospfv3 routing statistics uninstalled
OSPFv3 Process (1) UNINSTALLED ROUTES STATISTICS
     Intra-area-routes   :  2
     Inter-area-routes   :  0
     External-routes     :  1
```

Table 7-83 Description of the **display ospfv3 routing statistics uninstalled** command output

Item	Description
Intra-area-routes	Number of intra-area routes
Inter-area-routes	Number of inter-area routes
External-routes	Number of external routes

7.5.30 display ospfv3 sham-link

Function

The **display ospfv3 sham-link** command displays information about sham links in an OSPFv3 area.

Format

display ospfv3 [*process-id*] **sham-link** [**area** *area-id*]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPFv3 process.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the OSPFv3 area ID.	The integer ranges from 0 to 4294967295. The value of the area ID can be a decimal integer or in the IPv4 address format.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the process ID and area ID are not specified, information about all the configured sham links is displayed.

Example

Display all sham link information of OSPFv3.

```
<HUAWEI> display ospfv3 sham-link
 OSPFv3 Process (10)
 Sham Link SHAM-LINK1 to router 0.0.0.0 is down
 Area 0.0.0.1, via Interface *, Instance ID 0, cost 1
 Source address 2001:db8:1::1
 Destination address 2001:db8:2::1
 Interface ID 0x80000001
 Sham-Link Interface Events: 24
 Sham-Link Interface LsaCount: 0
 Sham-Link Interface Lsa Checksum: 0x0
 Transmit Delay is 1 sec, State Down
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:00
 Adjacency state Down
 IP security association configured: sa1
 IP security association applied: sa1
```

Table 7-84 Description of the **display ospfv3 sham-link** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process
Area	ID of the OSPFv3 area to which the sham link belongs
via Interface	Interface of the sham link
Instance ID	ID of the instance
cost	Cost of the sham link
Source address	Source IPv6 address of the sham link
Destination address	Destination IPv6 address of the sham link
Interface ID	Interface ID
Sham-Link Interface Events	Number of interface events of the sham-link
Sham-Link Interface LsaCount	Number of LSAs sent from all sham-link interfaces
Sham-Link Interface Lsa Checksum	Checksum of the LSA sent from the sham-link interface
Transmit Delay	Delay for an interface to send Update packets
State	Interface status of the sham link • Down • Point-To-Point

Item	Description
Timer intervals configured	Values of the following timers: • Hello • Dead • Retransmit • Wait
Adjacency state	Adjacency of the sham link • Down • Init • Exch • Start • Exchange • Loading • Full
IP security association configured	IP security association configured on a sham link
IP security association applied	IP security association applied on a sham link

7.5.31 display ospfv3 topology

Function

The **display ospfv3 topology** command displays the OSPFv3 area topology.

Format

display ospfv3 [*process-id*] **topology** [**area** *area-id*]

Parameters

Parameter	Description	Value
process-id	Identifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the area ID.	The value of the area ID can be a decimal integer or in the IPv4 address format. The integer ranges from 0 to 4294967295.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 topology** command to check the OSPFv3 area topology.

Example

Display the topology of OSPFv3 area 1.

```
<HUAWEI> display ospfv3 topology area 1
OSPFv3 Process (1)
OSPFv3 Area (0.0.0.1) topology
Type  ID(If-Index)       Bits      Metric    Next-Hop        Interface
Rtr   1.1.1.1                      --
Rtr   10.2.2.2                     1         10.2.2.2        Vlanif10
Net   10.2.2.2(268435842)          1         0.0.0.0         Vlanif10
```

Table 7-85 Description of the **display ospfv3 topology** command output

Item	Description
OSPFv3 Process	ID of the OSPFv3 process
OSPFv3 Area	ID of the OSPFv3 area
Type	SPF type: • Rtr: Router • Net: Network
ID (If-Index)	Router ID (If-Index) is the interface ID of the DR in the network type.
Bits	B bit and E bit indicate the switch type: • B bit: ABR • E bit: ASBR
Metric	Metric of path to the node
Next-Hop	Next hop of the path to the node
Interface	Egress of the path to the node

7.5.32 display ospfv3 vlink

Function

The **display ospfv3 vlink** command displays the OSPFv3 virtual links.

Format

display ospfv3 [*process-id*] **vlink**

Parameters

Parameter	Description	Value
process-id	Identifies the OSPFv3 process ID.	The value ranges from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ospfv3 vlink** command to check the OSPFv3 virtual links.

Example

Display OSPFv3 virtual links.

```
<HUAWEI> display ospfv3 vlink
OSPFv3 Process (100)

Virtual Link VLINK1 to router 10.2.2.2 is up
 Transit area 0.0.0.1 via interface Vlanif10, instance ID 0
 Local address FC00:0:0:2000::1
 Remote address FC00:0:0:1001::2
 Interface ID 0x80000001
 VirtualInterface Event: 1
 VirtualInterface LsaCount: 0
 VirtualInterface Lsa Checksum: 0x0
 Transmit Delay is 1 sec, State Point-To-Point, Cost 1562
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
     Hello due in 00:00:04
  Adjacency state Init
```

Table 7-86 Description of the **display ospfv3 vlink** command output

Item	Description
Transit area	If the current interface is the V-link interface, the transfer area ID is displayed
via interface	Interface of the V-link
instance ID	Instance ID
Local address	Local address
Remote address	Remote address
Interface ID	Interface ID

Item	Description
VirtualInterface Event	Number of interface events of the V-link
VirtualInterface LsaCount	Number of LSAs sent from all V-link interfaces
VirtualInterface Lsa Checksum	Checksum of the LSA sent from the V-link interface
Transmit Delay	Transmission delay of the interface
State	Interface state: • Down • Point-To-Point
Cost	Cost
Timer intervals configured	Timers: • Hello: Hello interval • Dead: Interval of Dead timer • Wait: Interval of Wait timer. It is pointless to use a virtual link. • Retransmit: Retransmission interval
Adjacency state	State of the V-link: • Down • Init • Exch • Start • Exchange • Loading • Full

7.5.33 domain-id (OSPFv3)

Function

The **domain-id** command sets an OSPFv3 domain ID.

The **undo domain-id** command restores the default setting.

By default, the domain ID is null.

Format

domain-id { **null** | *domain-id* [**type** *type* **value** *value* | **secondary**] * }

undo domain-id { *domain-id* [**type** *type* **value** *value*] }

Parameters

Parameter	Description	Value
null	Specifies that the OSPFv3 domain ID is null.	-
domain-id	Specifies an OSPFv3 domain ID.	The value can be an integer or expressed in dotted decimal notation. If the domain ID is an integer, its value ranges from 0 to 4294967295; it is converted to the dotted decimal notation and carried 1 with 256 when being displayed. If the domain ID is in dotted decimal notation format, it is displayed as originally entered.
type *type*	Specifies the type of the OSPFv3 domain ID.	The value is an integer that can be 0005, 0105, 0205, or 8005. The default value is 0005.
value *value*	Specifies the value of the OSPFv3 domain ID type.	The value is a hexadecimal number that ranges from 0 to FFFF. The default value is 0.
secondary	Specifies the ID of a secondary domain.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Generally, the routes that are imported from a PE are advertised as External-LSAs. However, such routes that belong to different nodes of the same OSPFv3 domain should be advertised as Type 3 LSAs (intra-domain routes). This requires that different nodes in the same OSPFv3 domain have the same domain ID.

Values 0 and null indicate different meanings.

The maximum number of secondary domain IDs in each OSPFv3 process is 10. The maximum number of secondary domain IDs may vary with products.

The parameter **secondary** can be configured only when the primary domain ID is configured. When the **undo domain-id** command is run, if no parameter is specified, the primary domain ID is deleted.

Example

Set the VPN domain ID in OSPFv3 VPN extension.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance test
[HUAWEI-vpn-instance-test] ipv6-family
[HUAWEI-vpn-instance-test-af-ipv6] route-distinguisher 100:1
[HUAWEI-vpn-instance-test-af-ipv6] quit
[HUAWEI-vpn-instance-test] quit
[HUAWEI] ospfv3 1 vpn-instance test
[HUAWEI-ospfv3-1] domain-id 234
```

7.5.34 filter export (OSPFv3)

Function

The **filter export** command filters the outgoing Type3 LSAs of the local area.

The **undo filter export** command restores the default setting.

By default, the outgoing Type3 LSAs of the local area are not filtered.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

filter { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } **export**

undo filter [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name*] **export**

Parameters

Parameter	Description	Value
acl6-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The command can be used to filter out invalid LSAs sent to neighbors to reduce the size of neighbors' LSDBs and speed up network convergence.

Configuration Impact

After filtering conditions are set for the outgoing Type3 LSAs to be advertised using the **filter export** command, only the outgoing Type3 LSAs that pass the filtering can be advertised.

Precautions

- The command can be configured only on ABRs.
- When the **rule** command is used to configure a filtering rule for an ACL, only the source address range specified by the parameter **source** and the time range specified by the parameter **time-range** in the **rule** command take effect.
- To set filtering conditions for the incoming Type3 LSAs to be advertised, run the **filter import (OSPFv3)** command.
- Creating an ACL6, IPv6 prefix list, or route-policy before it is referenced is recommended. If a nonexistent ACL6, IPv6 prefix list, or route-policy is referenced using the command, OSPFv3 advertises all Type 3 LSAs.

Example

Configure OSPFv3 to filter the outgoing Type3 LSAs.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 100
[HUAWEI-ospfv3-100] area 1
```


[HUAWEI-ospfv3-100-area-0.0.0.1] **filter 2000 export**

7.5.35 filter import (OSPFv3)

Function

The **filter import** command filters the incoming Type3 LSAs of the local area.

The **undo filter import** command restores the default setting.

By default, the incoming Type3 LSAs are not filtered.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

filter { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } **import**

undo filter [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name*] **import**

Parameters

Parameter	Description	Value
acl6-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **filter import** command can be used to filter out invalid LSAs sent to neighbors to reduce the size of the neighbors' LSDBs and speed up network convergence.

Configuration Impact

After filtering conditions are set for the incoming Type3 LSAs to be advertised using the **filter import** command, only the incoming Type3 LSAs that pass the filtering can be received and advertised.

Precautions

- The command can be configured only on ABRs.
- When the **rule** command is used to configure a filtering rule for an ACL, only the source address range specified by the parameter **source** and the time range specified by the parameter **time-range** in the **rule** command take effect.
- To set filtering conditions for the outgoing Type3 LSAs to be advertised, run the **filter export (OSPFv3)** command.
- Creating an ACL6, IPv6 prefix list, or route-policy before it is referenced is recommended. If a nonexistent ACL6, IPv6 prefix list, or route-policy is referenced using the command, OSPFv3 receives all Type 3 LSAs.

Example

Configure OSPFv3 to filter the incoming Type3 LSAs within the local area.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 100
[HUAWEI-ospfv3-100] area 1
[HUAWEI-ospfv3-100-area-0.0.0.1] filter ipv6-prefix my-prefix-list import
```


7.5.36 filter-policy export (OSPFv3)

Function

The **filter-policy export** command filters imported routes.

The **undo filter-policy export** command restores the default setting.

By default, imported routes are not filtered.

Format

filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } **export** [*protocol* [*process-id*]]

undo filter-policy [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name*] **export** [*protocol* [*process-id*]]

Parameters

Parameter	Description	Value
acl6-number	Specifies the ACL6 number.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of an IPv6 named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of the IPv6 address prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
protocol	Specifies the protocol from which the route is imported. You can use this parameter to filter the imported routes of specific protocols. If this parameter is not specified, all imported routes are filtered.	The value may be **bgp, direct**, **isis, ospfv3, ripng**, **static**, or **unr**. The specific value varies depending on the routing protocol supported by the device.
process-id	Specifies the process ID of the protocol.	The value is an integer ranging from 1 to 65535.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **filter-policy export** command is only valid for the routes that are imported using the **import-route** command.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Precautions

Creating an ACL6 or IPv6 prefix list before it is referenced is recommended. If a nonexistent ACL6 or IPv6 prefix list is referenced using the command, all external routes imported to OSPFv3 are converted to Type 5 LSAs (AS-external-LSAs) or Type 7 LSAs (NSSA-external-LSAs) and then are advertised to neighbors.

Example

Filter the imported routes based on the rule defined by ACL6 2002.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] filter-policy 2002 export
```

7.5.37 filter-policy import (OSPFv3)

Function

The **filter-policy import** command filters received routes.

The **undo filter-policy import** command restores the default setting.

By default, received routes are not filtered.

Format

filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } **import**

undo filter-policy [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name*] **import**

Parameters

Parameter	Description	Value
acl6-number	Specifies the ACL6 number.	The value is an integer ranging from 2000 to 2999.

Parameter	Description	Value
acl6-name *acl6-name*	Specifies the name of an IPv6 named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of the IPv6 address prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **filter-policy import** command only filters the routes calculated by OSPFv3. Routes that do not pass the filtering are neither added to the OSPFv3 routing table nor advertised.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Precautions

Creating an ACL6 or IPv6 prefix list before it is referenced is recommended. If a nonexistent ACL6 or IPv6 prefix list is referenced using the command, all IPv6 routes received by OSPFv3 are delivered to the IPv6 routing table.

Example

Filter the received routes based on the IPv6 address prefix list named **abc**.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] filter-policy ipv6-prefix abc import
```

7.5.38 frr (OSPFv3)

Function

The **frr** command creates and then displays the OSPFv3 FRR view.

The **undo frr** command deletes the OSPFv3 FRR view.

By default, the OSPFv3 FRR view does not exist.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

frr

undo frr

Parameters

None

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

OSPFv3 IP FRR allows devices to fast switch traffic from faulty links to back up links without interrupting traffic. This function protects traffic and greatly improves the reliability of OSPFv3 networks. OSPFv3 IP FRR must be configured in the OSPFv3 FRR view. The **frr** command run in the OSPFv3 view creates and displays the OSPFv3 FRR view.

Follow-up Procedure

The **frr** command run in the OSPFv3 view creates and displays the OSPFv3 FRR view only, but cannot enable the OSPFv3 IP FRR function. Run the **loop-free-alternate** command in the OSPFv3 view to enable OSPFv3 IP FRR to create the loop-free backup route.

Example

Create and display the OSPFv3 FRR view.


```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] frr
[HUAWEI-ospfv3-1-frr]
```

7.5.39 frr-policy route (OSPFv3)

Function

The **frr-policy route** command configures a filtering policy for the OSPFv3 IP FRR backup routes. The filtering policy determines what kind of OSPFv3 backup route can be added to the routing table.

The **undo frr-policy route** command cancels the filtering function.

By default, the filtering function is disabled.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

frr-policy route route-policy *route-policy-name*

undo frr-policy route

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of the policy used to filter OSPFv3 backup routes.	The value must be an existing route-policy.

Views

OSPFv3 FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

OSPFv3 IP FRR allows devices to fast switch traffic on fault links to back up links without interrupting traffic. This protects traffic and greatly improves the reliability of OSPFv3 networks.

After the filtering policy is configured using the **frr-policy route** command, only the OSPFv3 backup route that satisfies filtering rules can be delivered to the forwarding table.

Configuration Impact

To protect the traffic over a specific OSPFv3 route, you can configure a filtering policy *route-policy-name* that the OSPFv3 route matches to ensure that the backup route can be added to the forwarding table. When this route fails, OSPFv3 can fast switch the traffic to a backup route.

Precautions

The **frr-policy route** command is cyclic in nature, and only the latest configuration takes effect.

If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, OSPFv3 adds all backup routes to the IPv6 routing table.

Example

Configure OSPFv3 to add the OSPFv3 backup routes that match the named ACL **abc** to the IPv6 routing table.

```
<HUAWEI> system-view
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] frr
[HUAWEI-ospfv3-1-frr] loop-free-alternate
[HUAWEI-ospfv3-1-frr] frr-policy route route-policy abc
```

7.5.40 graceful-restart (OSPFv3)

Function

The **graceful-restart** command enables GR of the OSPFv3 process.

The **undo graceful-restart** command disables GR of the OSPFv3 process.

By default, OSPFv3 GR is not enabled.

Format

graceful-restart [**period** *period* | **ack-time** *time* | **retransmit-interval** *interval* | **lsa-checking-ignore** | **planned-only**] *

undo graceful-restart [**period** [*period*] | **ack-time** [*time*] | **retransmit-interval** [*interval*] | **lsa-checking-ignore** | **planned-only**] *

Parameters

Parameter	Description	Value
period *period*	Specifies the GR period.	The value is an integer ranging from 1 to 1800, expressed in seconds. By default, the value is 120 seconds.
ack-time *time*	Specifies the period during which a switch waits for the Grace-LSA Ack message from its neighbor.	The value is an integer ranging from 1 to 30, expressed in seconds. By default, the value is 15 seconds.
retransmit-interval *interval*	Specifies the interval for retransmitting Grace-LSAs.	The value is an integer ranging from 1 to 5, expressed in seconds. By default, the value is 5 seconds.
lsa-checking-ignore	Indicates that strict check is not performed on LSAs. By default, switches perform strict check on received LSAs.	-
planned-only	Indicates that switches support Planned-GR only. By default, switches support Planned-GR and Unplanned-GR.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

ack-time is an optional parameter. After the parameter is specified, the restarter can discover more neighbors in *ack-time* period.

Example

Enable GR of the OSPFv3 process.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] graceful-restart
```


7.5.41 helper-role (OSPFv3)

Function

The **helper-role** command configures a switch to support the GR Helper mode.

The **undo helper-role** command cancels GR Helper mode.

By default, the switch is not enabled to support GR Helper mode.

Format

helper-role [{ **ip-prefix** *ip-prefix-name* | **acl-number** *acl-number* | **acl-name** *acl-name* } | **max-grace-period** *period* | **planned-only** | **lsa-checking-ignore**]*

undo helper-role [[**ip-prefix** [*ip-prefix-name*] | **acl-number** [*acl-number*] | **acl-name** [*acl-name*]] | **max-grace-period** *period* | **planned-only** | **lsa-checking-ignore**]*

Parameters

Parameter	Description	Value
ip-prefix *ip-prefix-name*	Specifies the name of the IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
acl-number *acl-number*	Specifies the number of the basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a IPv6 Named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
max-grace-period *period*	Specifies the permitted maximum restart period.	The value is an integer ranging from 1 to 1800, expressed in seconds. By default, the value is 1800 seconds.
planned-only	Indicates that switches support Planned-GR only. By default, switches support Planned-GR and Unplanned-GR.	-

Parameter	Description	Value
lsa-checking-ignore	Indicates that strict check is not performed on LSAs. By default, switches perform strict check on received LSAs.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Switches cannot use the helper mode during the GR process.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Example

Configure a switch to support the helper mode.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] helper-role acl-number 2001 max-grace-period 250 planned-only lsa-checking-ignore
```

7.5.42 import-route (OSPFv3)

Function

The **import-route** command imports an external route. Before the routes are imported, the OSPFv3 process of the route must be active.

The **undo import-route** command stops importing external routes.

By default, no route of other protocols is imported.

Format

import-route { **bgp** [**permit-ibgp**] | **unr** | **direct** | **ripng** *help-process-id* | **static** | **isis** *help-process-id* | **ospfv3** *help-process-id* } [{ **cost** *cost* | **inherit-cost** } | **type** *type* | **tag** *tag* | **route-policy** *route-policy-name*]*

undo import-route { **bgp** | **unr** | **direct** | **ripng** *help-process-id* | **static** | **isis** *help-process-id* | **ospfv3** *help-process-id* }

Parameters

Parameter	Description	Value
bgp	Specifies the protocol from which routes are imported, as **bgp**.	-
permit-ibgp	IBGP routes that are permitted to be imported. The import of IBGP routes may cause route loops. Therefore, this command must not be configured unless it is necessary.	-
unr	Specifies the imported source routing protocol as **unr**. User Network Route (UNR) is allocated if dynamic routing protocols cannot be used when users are getting online.	-
direct	Specifies the imported source routing protocol as **direct**.	-
ripng	Specifies the protocol from which routes are imported, as **ripng**.	-
static	Specifies the imported source routing protocol as **static**.	-
isis	Specifies the protocol from which routes are imported, as **isis**.	-
ospfv3	Specifies the protocol from which routes are imported, as **ospfv3**.	-
help-process-id	Specifies the process ID of the imported source protocol.	The value is an integer ranging from 1 to 65535.
cost *cost*	Indicates the cost of the imported route.	The value is an integer ranging from 1 to 16777214. For details about the default value, see **default (OSPFv3)**.
inherit-cost	Indicates the original cost of the imported routes.	-

Parameter	Description	Value
type *type*	Specifies the type of the external routes.	The value is an integer ranging from 1 to 2. For details about the default value, see **default (OSPFv3)**. • 1: Type 1 external route • 2: Type 2 external route
tag *tag*	Specifies the tag value of the imported routes.	The value is an integer that ranges from 0 to 4294967295. For details about the default value, see **default (OSPFv3)**.
route-policy *route-policy-name*	Specifies the name of the routing policy. Only the routes that match the *route-policy-name* are imported.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a non-PE device, only EBGP routes are imported after the **import-route bgp** command is configured. IBGP routes are also imported after the **import-route bgp permit-ibgp** command is configured. If IBGP routes are imported, routing loops may occur. In this case, run the **preference (OSPFv3)** and **preference (BGP)** commands to prevent loops by specifying preferences of OSPFv3 and BGP routes. If IBGP routes need to be imported, configure the **import-route bgp permit-ibgp** command, and run the **preference (OSPFv3)** and **preference (BGP)** commands to set the preference of OSPFv3 ASE routes lower than that of IBGP routes (preference value of OSPFv3 ASE routes larger than that of IBGP routes).

On a PE, configuring the **import-route bgp** command imports both EBGP routes and IBGP routes, no matter whether the **import-route bgp permit-ibgp** command is configured or not. If the **import-route bgp permit-ibgp** command and the **default-route-advertise (OSPFv3)** command are both configured, the active IBGP default routes can be imported into OSPFv3.

NOTE

After the **import-route direct** command is executed, routes to the network segment where the IPv6 address of the management interface belongs are also imported in the OSPFv3 routing table. Therefore, use this command with caution.

Precautions

Creating a route-policy before it is referenced is recommended. If a nonexistent route-policy is referenced using the command, all the routes learned by the specified protocol are imported to the OSPFv3 routing table.

Example

Import the RIPng route of type 2. The cost of the route is 50.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] import-route ripng 1 type 2 cost 50
```

OSPFv3 process 100 imports the routes discovered by OSPFv3 process 160.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 100
[HUAWEI-ospfv3-100] import-route ospfv3 160
```

7.5.43 ipsec sa (OSPFv3)

Function

The **ipsec sa** command configures a Security Association (SA) in the OSPFv3 area or OSPFv3 process.

The **undo ipsec sa** command deletes the SA configured in the OSPFv3 area or OSPFv3 process.

By default, no SA is configured in the OSPFv3 area or OSPFv3 process.

Format

ipsec sa *sa-name*

undo ipsec sa

Parameters

Parameter	Description	Value
sa-name	Specifies the name of an SA.	The value is an existing SA name.

Views

OSPFv3 view or OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An SA defines a set of security algorithms and keys to ensure IP security. Incoming and outgoing OSPFv3 packets are authenticated and encrypted based on the rules defined by the SA.

The **ipsec sa** *sa-name* command run in the OSPFv3 process view is used to authenticate packets of the OSPFv3 process. An SA applied in an OSPFv3 process is used to authenticate the packets of the process. After an OSPFv3 process is associated with an OSPFv3 area, the SA applied in the OSPFv3 process is also applied to the OSPFv3 area.

The **ipsec sa** *sa-name* command run in the OSPFv3 area view is used to authenticate packets of the OSPFv3 area.

NOTE

- The SA applied in an OSPFv3 area takes precedence over that applied in an OSPFv3 process.
- If the SA applied in the OSPFv3 area is deleted, the SA applied in the OSPFv3 process is used to authenticate packets.

Example

Configure an SA named **sa1** in the OSPFv3 process. (This SA has been created.)

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] ipsec sa sa1
```

Configure an SA named **sa2** in the OSPFv3 area. (This SA has been created.)

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 10.0.0.0
[HUAWEI-ospfv3-1-area-10.0.0.0] ipsec sa sa2
```

7.5.44 loop-free-alternate (OSPFv3)

Function

The **loop-free-alternate** command enables OSPFv3 IP FRR to enable the device to use the LFA algorithm to calculate the nexthop and outbound interface for the dynamic backup link.

The **undo loop-free-alternate** command disables OSPFv3 IP FRR.

By default, this function is disabled.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

loop-free-alternate

undo loop-free-alternate

Parameters

None

Views

OSPFv3 FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **frr** command run in the OSPFv3 view creates and displays the OSPFv3 FRR view only, but cannot enable the OSPFv3 IP FRR function. You must run the **loop-free-alternate** command in the OSPFv3 view to enable OSPFv3 IP FRR use the LFA algorithm to calculate the nexthop and outbound interface for the dynamic backup link to create the loop-free backup route.

Follow-up Procedure

Run the **frr-policy route** command in the OSPFv3 view to configure a filtering policy for OSPFv3 IP FRR. Only the OSPFv3 backup route that satisfies specific rules can be delivered to the forwarding table.

Example

Enable OSPFv3 IP FRR to create the loop-free backup route.

```
<HUAWEI> system-view
[HUAWEI] ospfv3
```



```
[HUAWEI-ospfv3-1] frr
[HUAWEI-ospfv3-1-frr] loop-free-alternate
```

7.5.45 lsa-forwarding-address

Function

The **lsa-forwarding-address** command enables the OSPFv3 forwarding address (FA) function.

The **undo lsa-forwarding-address** command disables the OSPFv3 FA function.

By default, the OSPFv3 FA function is disabled.

Format

lsa-forwarding-address { **standard** | **zero-translate** }

undo lsa-forwarding-address

Parameters

Parameter	Description	Value
standard	Indicates that the function is compatible with RFC 3101.	-
zero-translate	Allows the Type 7 LSAs with the P-bit set and the FA being 0 to be translated to Type 5 LSAs.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

For an OSPFv3 Type 5 or Type 7 LSA:

- When there is no FA, other switches consider how to reach the ASBR (the switch that generates this LSA) and then reach an external network through this ASBR.
- When there is an FA, other switches consider how to reach the device to which this FA belongs and then reach an external network through this device.

In this situation, if no FA is available, an additional next hop may be generated when a switch within an OSPF domain first reaches an ASBR before reaching an external network.

Therefore, the FA in an OSPFv3 Type 5 or Type 7 LSA provides a faster next hop to reach an external network.

Precautions

If the FA in an OSPFv3 Type 5 or Type 7 LSA is not 0, the corresponding route is calculated only when the **lsa-forwarding-address** command is configured.

Example

Enable the OSPFv3 FA function that is compatible with RFC 3101.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] lsa-forwarding-address standard
```

7.5.46 maximum load-balancing (OSPFv3)

Function

The **maximum load-balancing** command sets the maximum number of equal-cost routes for carrying out load balancing.

The **undo maximum load-balancing** command restores the default setting.

The default maximum number of equal-cost OSPFv3 routes for carrying out load balancing is 8.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

maximum load-balancing *number*

undo maximum load-balancing

Parameters

Parameter	Description	Value
number	Specifies the maximum number of equal-cost routes.	The value is an integer ranging from 1 to 8.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Applicable Environment

If a routing protocol discovers multiple routes with the same cost to a single destination, they can carry out load balancing. The **maximum load-balancing** command sets the maximum number of equal-cost routes that can carry out load balancing. This optimizes the routing policy and ensures traffic forwarding on a complex network.

Effect

Packets will be load-balanced by multiple equal-cost routes to a single destination.

NOTE

Load balancing is performed in either per-flow or per-packet mode. By default, packets are load-balanced in per-flow mode.

Follow-up Procedure

If more existing equal-cost OSPFv3 routes than the value set using the **maximum load-balancing** command are available, valid routes are selected for load balancing based on the following criteria:

- Route preference: Routes with lower preferences are selected for load balancing.
- Interface index: If routes have the same priorities, routes with higher interface index values are selected for load balancing.
- Next hop IP address: If routes have the same priorities and interface index values, routes with larger IP address are selected for load balancing.

The **nexthop** command allows routes with a specified weight to carry out load balancing.

Precautions

To disable load balancing, set the value of *number* to 1.

Example

Set the maximum number of equal-cost routes that OSPFv3 can support to 3.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] maximum load-balancing 3
```

Restore the default maximum number of equal-cost routes for carrying out load balancing.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] undo maximum load-balancing
```

7.5.47 nexthop (OSPFv3)

Function

The **nexthop** command sets the preference for equal-cost routes.

The **undo nexthop** command cancels the preference of equal-cost routes.

By default, the preference is not set for equal-cost routes. That is, equal-cost routes forward packets at the same time for load balancing.

Format

nexthop *router-id interface-type interface-number* **weight** *value*

undo nexthop *router-id interface-type interface-number*

Parameters

Parameter	Description	Value
router-id	Specifies the router ID of a neighbor.	It is in dotted decimal notation.
interface-type interface-number	Specifies the type and number of an interface.	-
weight *value*	Specifies the weight of the next hop. The smaller the weight is, the higher the route preference is.	It is an integer ranging from 1 to 254.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

After OSPFv3 calculates equal-cost routes, you can run the **nexthop** command to select the route with the highest preference from the equal-cost routes as the next

hop. OSPFv3 selects a next hop from these equal-cost routes according to the weight. The smaller the weight is, the higher the route preference is.

Example

Set the preference for equal-cost routes in OSPFv3.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] quit
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] nexthop 10.2.2.2 vlanif 10 weight 240
```

7.5.48 nssa (OSPFv3 Area)

Function

The **nssa** command configures an OSPFv3 area as an NSSA.

The **undo nssa** command cancels the configuration.

By default, no OSPFv3 area is configured as an NSSA.

Format

nssa [**default-route-advertise** [**cost** *cost* | **type** *type* | **tag** *tag*] * | **no-import-route** | **no-summary** | **translator-always** | **translator-interval** *translator-interval* | **set-n-bit** | **suppress-forwarding-address**] *

undo nssa

Parameters

Parameter	Description	Value
default-route-advertise	Generates default Type7 LSAs on the ASBR and then advertises them to the NSSA. **NOTE** The ABR generates a default NSSA LSA (Type7 LSA) automatically and advertises it in the NSSA.	-
cost *cost*	Specifies the default cost of Type 7 LSAs.	It is an integer ranging from 1 to 16777214. The default value is 1.

Parameter	Description	Value
type *type*	Specifies the type of the external routes.	The value is an integer ranging from 1 to 2. By default, it is 2. • 1: Type 1 external route • 2: Type 2 external route
tag *tag*	Specifies the tag value of the OSPFv3 route imported to an NSSA.	It is an integer ranging from 0 to 4294967295. The default value is 0.
no-import-route	Indicates that no external routes are imported to NSSAs.	-
no-summary	Disables ABRs from sending summary LSAs to NSSAs.	-
translator-always	Specifies the ABR in an NSSA as the translator. Multiple ABRs in an NSSA can be configured as translators.	-
translator-interval *translator-interval*	Specifies the timeout period of a translator.	It is an integer ranging from 1 to 120, in seconds. The default value is 40 seconds.
set-n-bit	Indicates that the N-bit is set in DD packets.	-
suppress-forwarding-address	Enables the device to suppress the forwarding address of a Type 5 LSA translated from a Type 7 LSA.	-

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

If an area is configured to the nssa area, all switches of the area must be configured with the NSSA attribute.

The area will be updated after NSSA attributes are configured or deleted. Thus, the NSSA attributes can be re-configured or deleted only after the last update of NSSA attributes is complete.

The keyword **default-route-advertise** is used to generate default Type 7 LSAs. Regardless of whether there is route ::/0 in the routing table on an ABR, a Type 7 LSA default route is generated. A Type 7 LSA default route can be generated only when there is route ::/0 in the routing table on an ASBR.

When an ASBR is also an ABR, you can specify **no-import-route** to disable OSPFv3 from advertising the external routes imported through the **import-route** command to NSSAs. To reduce the number of LSAs to be transmitted to NSSAs, you can specify **no-summary** on an ABR to disable the ABR from transmitting summary LSAs (Type 3 LSAs) to NSSAs.

NOTE

- When the LS age field (aging time) in the LSA header reaches 3600 seconds, this LSA is deleted.
- After the keyword **set-n-bit** is set, a switch re-establishes neighbor relationships with switches directly connected to it in the NSSA.

Example

Configure Area 1 as an NSSA, and set the timeout period of the translator to 20s.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 1
[HUAWEI-ospfv3-1-area-0.0.0.1] nssa translator-interval 20
```

7.5.49 ospfv3

Function

The **ospfv3** command enables an OSPFv3 process.

The **undo ospfv3** command disables the OSPFv3 process.

By default, the system does not support OSPFv3.

Format

ospfv3 [*process-id*] [**vpn-instance** *vpn-instance-name*]

undo ospfv3 *process-id*

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535. By default, it is 1.
vpn-instance *vpn-instance-name*	Specifies the name of the VPN instance.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Only when you configure the Route ID in the OSPFv3 view, the OSPFv3 process can run normally. Otherwise, the OSPFv3 process can be viewed but cannot generate LSAs.

If a VPN instance is specified, the OSPFv3 process belongs to the VPN instance. If no VPN instance is specified, the OSPFv3 process is a global instance.

NOTE

Before configuring OSPFv3-related parameters, you must enable OSPFv3.

Example

Enable the OSPFv3 protocol.

```
<HUAWEI> system-view
[HUAWEI] ospfv3
```

7.5.50 ospfv3 area

Function

The **ospfv3 area** command enables the OSPFv3 process on an interface and specifies the area the process belongs to.

The **undo ospfv3 area** command disables the OSPFv3 process on the interface.

By default, the OSPFv3 protocol is not enabled on the interface.

Format

ospfv3 *process-id* **area** *area-id* [**instance** *instance-id*]

undo ospfv3 *process-id* **area** *area-id* [**instance** *instance-id*]

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value is an integer ranging from 1 to 65535.
area *area-id*	Specifies the OSPFv3 area ID.	The value of the area ID can be a decimal integer or in the IPv4 address format. The integer ranges from 0 to 4294967295.
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Before running this command on an interface, run the **ipv6 enable** command to enable IPv6 in the interface view.

Example

Enable OSPFv3 on an interface and enable OSPFv3 in area 1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
```

Enable OSPFv3 in GE0/0/1, and enable OSPFv3 in area 1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
```


7.5.51 ospfv3 authentication-mode

Function

The **ospfv3 authentication-mode** command configures an authentication mode and a password for an OSPFv3 interface.

The **undo ospfv3 authentication-mode** command deletes the authentication mode and password configured for an OSPFv3 interface.

By default, no authentication mode or password is configured for any OSPFv3 interface.

Format

ospfv3 authentication-mode hmac-sha256 key-id *key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* } [**instance** *instance-id*]

ospfv3 authentication-mode keychain *keychain-name* [**instance** *instance-id*]

undo ospfv3 authentication-mode hmac-sha256 key-id *key-id* [**plain** *plain-text* | **cipher** *cipher-text*] [**instance** *instance-id*]

undo ospfv3 authentication-mode keychain [*keychain-name*] [**instance** *instance-id*]

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
hmac-sha256	Configures the HMAC-SHA256 authentication mode.	N/A
key-id *key-id*	Specifies the key ID for authentication, which must be the same as the one configured at the other end.	The value is an integer ranging from 1 to 65535.

Parameter	Description	Value
plain	Configures the plaintext password type. Only a plaintext password can be entered, and the password is displayed in plaintext in the configuration file. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	N/A
plain-text	Specifies a plaintext password.	The value is a string of 1 to 255 characters, spaces not supported.
cipher	Configures the ciphertext password type. You can enter either a plaintext or ciphertext password, but the password is displayed in ciphertext in the configuration file.	N/A
cipher-text	Specifies a ciphertext password.	The value can be a string of 1 to 255 characters for plaintext passwords and 20 to 392 characters for ciphertext passwords, spaces not supported.
keychain	Configures keychain authentication. **NOTE** Before you configure keychain authentication, run the **keychain** command to configure a keychain, the **key-id** command to configure a key ID, the **key-string** command to configure a password, and the **algorithm** command to configure an algorithm. If these commands are not run, OSPFv3 authentication fails.	N/A

Parameter	Description	Value
keychain-name	Specifies a keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.
instance *instance-id*	Specifies the ID of a VLAN to which the specified interface belongs.	The value ranges from 0 to 255, with default value 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Due to inherent defects and flawed implementation of the TCP/IP protocol suite, there are an increasing number of attacks, which poses greater threats on TCP/IP networks than ever before. The attacks on network devices may lead to network failures. To configure an authentication mode and a password for an OSPFv3 interface to improve OSPFv3 network security, run the **ospfv3 authentication-mode** command.

Precautions

OSPFv3 interface authentication takes precedence over OSPFv3 area authentication.

To configure OSPFv3 area authentication, run the **authentication-mode** command.

Example

Configure OSPFv3 HMAC-SHA256 authentication on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] ospfv3 1 area 0
```



```
[HUAWEI-Vlanif100] ospfv3 authentication-mode hmac-sha256 key-id 10 cipher test
```

Configure OSPFv3 HMAC-SHA256 authentication on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 0
[HUAWEI-GigabitEthernet0/0/1] ospfv3 authentication-mode hmac-sha256 key-id 10 cipher test
```

7.5.52 ospfv3 bfd

Function

The **ospfv3 bfd** command enables the bidirectional forwarding detection (BFD) on the specified interface enabled with OSPFv3, or sets the parameter values of a BFD session.

The **undo ospfv3 bfd** command deletes the BFD on the specified interface, or restores the default parameter values of a BFD session.

By default, BFD is not enabled on OSPFv3 interfaces.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

ospfv3 bfd enable [**instance** *instance-id*]

ospfv3 bfd { **min-transmit-interval** *min-transmit-value* | **min-receive-interval** *min-receive-value* | **detect-multiplier** *detect-multiplier-value* } * [**instance** *instance-id*]

undo ospfv3 bfd enable [**instance** *instance-id*]

undo ospfv3 bfd { **min-transmit-interval** [*min-transmit-value*] | **min-receive-interval** [*min-receive-value*] | **detect-multiplier** [*detect-multiplier-value*] } * [**instance** *instance-id*]

Parameters

Parameter	Description	Value
enable	Enables BFD for OSPFv3 on the specified interface.	-
min-transmit-interval *min-transmit-value*	Specifies the minimum interval for sending BFD packets to the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. ● After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. ● After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-receive-interval *min-receive-value*	Specifies the minimum interval for receiving BFD packets from the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. ● After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. ● After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *detect-multiplier-value*	Indicates the local detection multiplier.	The value is an integer in the range from 3 to 50. The default value is 3.
instance *instance-id*	Specifies an interface instance ID.	The value is an integer in the range from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A link fault or a topology change causes devices to recalculate routes. Therefore, the convergence of routing protocols must be as fast as possible to improve network performance.

Link faults are inevitable. Therefore, quickly detecting faults and notifying routing protocols of the faults is a feasible solution to immediately rectify link faults. If BFD is associated with routing protocols, BFD can speed up the convergence of routing protocols when a link fault occurs.

After OSPFv3 establishes a BFD session, the default parameter values are used. BFD session parameters can also be modified as required.

- Actual interval at which BFD packets are transmitted on the local device = Max { configured interval *min-transmit-value* at which BFD packets are transmitted on the local device, configured interval *min-receive-value* at which BFD packets are received on the peer device }
- Actual interval at which BFD packets are received on the local device = Max { configured interval *min-transmit-value* at which BFD packets are transmitted on the peer device, configured interval *min-receive-value* at which BFD packets are received on the local device }
- Actual period for BFD detection on the local device = Actual interval at which BFD packets are received on the local device x Detection multiplier *detect-multiplier-value* configured on the peer device

For example:

- On the local device, the configured interval at which BFD packets are transmitted is 200 ms; the interval at which BFD packets are received is set to 300 ms; the detection multiplier is 4.
- On the peer device, the configured interval at which BFD packets are transmitted is 100 ms; the interval at which BFD packets are received is 600 ms; the detection multiplier is 5.

Then:

- On the local device, the actual interval at which BFD packets are transmitted is 600 ms calculated by using the formula max {200 ms, 600 ms}; the interval at which BFD packets are received is 300 ms calculated by using the formula max {100 ms, 300 ms}; the detection period is 1500 ms calculated by multiplying 300 ms by 5.
- On the peer device, the actual interval at which BFD packets are transmitted is 300 ms calculated by using the formula max {100 ms, 300 ms}, the actual interval at which BFD packets are received is 600 ms calculated by using the formula max {200 ms, 600 ms}, and the detection period is 2400 ms calculated by multiplying 600 ms by 4.

Prerequisites

The set BFD session parameters take effect only when BFD is enabled on the interface.

Configuration Impact

If the global BFD is not configured, BFD on the interface can be configured but the BFD session cannot be set up. Similarly, if the parameters of a BFD session are set but the **ospfv3 bfd enable** command is not configured, the BFD session cannot be set up.

The priority of BFD configured on an interface is higher than that configured in a process. If the BFD is enabled on the interface, the BFD session is set up by using the parameters of the BFD configured on the interface.

Precautions

- After BFD is configured, OSPFv3 establishes BFD sessions only with neighbors.
- The **ospfv3 bfd enable** command and the **ospfv3 bfd block** command are mutually exclusive.
- After BFD is disabled on the interface using the **undo ospfv3 bfd enable** command, the configurations of the parameters of the BFD session on the interface still exist.
- BFD cannot be enabled on VBDIF interfaces.

Example

Enable BFD for OSPFv3 on interface Vlanif10.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ospfv3 bfd enable instance 1
```

Configure BFD parameters for OSPFv3 on interface Vlanif10.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ospfv3 bfd min-transmit-interval 100 min-receive-interval 100 detect-multiplier 3 instance 1
```

7.5.53 ospfv3 bfd block

Function

The **ospfv3 bfd block** command blocks the bidirectional forwarding detection (BFD) dynamically created by an interface.

The **undo ospfv3 bfd block** command cancels the feature.

By default, BFD is not blocked.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ospfv3 bfd block [**instance** *instance-id*]

undo ospfv3 bfd block [**instance** *instance-id*]

Parameters

Parameter	Description	Value
instance *instance-id*	Indicates the interface instance id.	The value is an integer ranging from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If the **bfd all-interfaces enable** command is run in the OSPFv3 process, all neighbors that enable OSPFv3 and the neighbor relationship is in Full state create BFD sessions. To prevent specific interfaces from being enabled with BFD, you can disable their interfaces from dynamically creating BFD sessions.

Prerequisites

BFD has been enabled on these interfaces.

Before configuring this command, you must run the **ipv6 enable** command in the interface view to enable IPv6.

Precautions

The **ospfv3 bfd enable** command and the **ospfv3 bfd block** command are mutually exclusive.

Example

Block the BFD dynamically created by Vlanif100 when configured the BFD for OSPF3 of all interfaces in the OSPF3 process.


```
<HUAWEI> system-view
[HUAWEI] interface Vlanif100
[HUAWEI-Vlanif100] ospfv3 bfd block instance 1
```

7.5.54 ospfv3 cost

Function

The **ospfv3 cost** command sets the cost of the interface in different instances.

The **undo ospfv3 cost** command restores the default cost of the interface in different instances.

By default, the cost of an interface running OSPF is calculated using the following formula: **Interface cost = Bandwidth reference value/Interface bandwidth** where, the bandwidth reference value can be changed using the **bandwidth-reference** command.

Format

ospfv3 cost *cost* [**instance** *instance-id*]

undo ospfv3 cost [*cost*] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
cost	Specifies the cost for running OSPFv3.	The value is an integer ranging from 1 to 65535.
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Set the cost for public instances to run OSPFv3 on VLANIF 10 to 33.

```
<HUAWEI> system-view
[HUAWEI] ipv6
```



```
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 cost 33
```

Set the cost for public instances to run OSPFv3 on GE0/0/1 to 33.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 cost 33
```

7.5.55 ospfv3 dr-priority

Function

The **ospfv3 dr-priority** command sets the priority of an interface that takes part in the DR or BDR election.

The **undo ospfv3 dr-priority** command restores the default value.

By default, the priority of an interface that takes part in the DR or BDR election is 1.

Format

ospfv3 dr-priority *priority* [**instance** *instance-id*]

undo ospfv3 dr-priority [*priority*] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
priority	Specifies the priority of the interface that candidates for DR or BDR.	The value is an integer ranging from 0 to 255. The default value is 1.
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The priority of the interface determines the qualification of the interface when electing DR or BDR. The interface with the higher priority is preferred. The router whose priority is 0 cannot be elected as a DR or a BDR.

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Set the priority of public instance for electing DR or BDR to 8 on Vlanif10.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 dr-priority 8
```

Set the priority of public instance for electing DR or BDR to 8 on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 dr-priority 8
```

7.5.56 ospfv3 frr block

Function

The **ospfv3 frr block** command disables the OSPFv3 IP FRR function on a specified interface.

The **undo ospfv3 frr block** command restores the OSPFv3 IP FRR function on the specified interface.

By default, the OSPFv3 IP FRR function is enabled on a specific interface.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ospfv3 frr block [**instance** *instance-id*]

undo ospfv3 frr block [**instance** *instance-id*]

Parameters

Parameter	Description	Value
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

OSPFv3 IP FRR can be disabled using the **ospfv3 frr block** command on an interface of a specific device that is running important services and resides on an FRR backup link. This setting prevents the device connected to this interface from being a part of a backup link and being burdened after FRR switches traffic to the backup link.

Example

Disable the OSPFv3 IP FRR function on the interface VLANIF 10.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ospfv3 frr block
```

7.5.57 ospfv3 ipsec sa

Function

The **ospfv3 ipsec sa** command configures an SA in the OSPFv3 interface.

The **undo ospfv3 ipsec sa** command deletes the SA configured in the OSPFv3 interface.

By default, no SA is configured in the OSPFv3 interface.

Format

ospfv3 ipsec sa *sa-name*

undo ospfv3 ipsec sa

Parameters

Parameter	Description	Value
sa-name	Specifies the name of an SA.	The value is an existing SA name.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If an SA is configured on an interface, OSPFv3 uses the SA to authenticate and encrypt packets sent and received by the interface.

If no SA is configured on the interface or the configured SA is deleted from the interface, OSPFv3 uses the SA configured in the process or the area where the interface resides to authenticate packets sent and received by the interface.

NOTE

- The SA configured in the interface view takes precedence over that configured in the OSPFv3 area view or the OSPFv3 process view.
- The **ospfv3 ipsec sa** command can be used on all the OSPFv3 instances of an interface.

Example

Configure an SA named **sa3** for the interface VLANIF10. (This SA has been created.)

```
<HUAWEI> system-view
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 0
[HUAWEI-Vlanif10] ospfv3 ipsec sa sa3
```


7.5.58 ospfv3 mib-binding

Function

The **ospfv3 mib-binding** command binds an OSPFv3 process to SNMP and makes OSPFv3 respond to SNMP requests.

The **undo ospfv3 mib-binding** command disables the binding.

By default, there is no binding between the OSPFv3 process and SNMP.

Format

ospfv3 mib-binding *process-id*

undo ospfv3 mib-binding

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process ID.	The value ranges from 1 to 65535.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

When multiple OSPFv3 processes are enabled, you can bind OSPFv3 MIB to a process so that this process can be processed.

Before running this command, run the **ospfv3** command in the system view to create an OSPFv3 process.

NOTE

Only after **ospfv3 mib-binding** *process-id* command is run on a device, can the trap associated with the specified process ID be triggered on the device.

Example

Bind OSPFv3 process to SNMP.

```
<HUAWEI> system-view
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] quit
[HUAWEI] ospfv3 mib-binding 100
```

Disable the binding.

```
<HUAWEI> system-view
```



```
[HUAWEI] undo ospfv3 mib-binding
```

7.5.59 ospfv3 mtu-ignore

Function

The **ospfv3 mtu-ignore** command disables the MTU check on DD packets.

The **undo ospfv3 mtu-ignore** command restores the default value.

By default, the MTU check on DD packets is enabled.

Format

ospfv3 mtu-ignore [**instance** *instance-id*]

undo ospfv3 mtu-ignore [**instance** *instance-id*]

Parameters

Parameter	Description	Value
instance *instance-id*	Specifies the interface instance ID.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Disable OSPFv3 from checking the MTU of the DD packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 mtu-ignore
```

Disable OSPFv3 from checking the MTU of the DD packets.

```
<HUAWEI> system-view
[HUAWEI] ipv6
```



```
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 mtu-ignore
```

7.5.60 ospfv3 network-type

Function

The **ospfv3 network-type** command sets the network type of the OSPFv3 interface.

The **undo ospfv3 network-type** command restores the default network type of the OSPFv3 interface.

By default, the network type of an interface is determined by the physical interface. The network type of Ethernet interface is **broadcast**.

Format

ospfv3 network-type { **broadcast** | **nbma** | **p2mp** [**non-broadcast**] | **p2p** } [**instance** *instance-id*]

undo ospfv3 network-type [**broadcast** | **nbma** | **p2mp** [**non-broadcast**] | **p2p**] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
broadcast	Indicates that the network type of the interface is changed to broadcast.	-
nbma	Indicates that the network type of the interface is changed to NBMA.	-
p2mp	Indicates that the network type of the interface is changed to point-to-multipoint.	-
non-broadcast	Indicates that the network type of the interface is changed to non-broadcast point-to-multipoint.	-
p2p	Indicates that the network type of the interface is changed to point-to-point.	-
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

None

Example

Set network type of Vlanif10 to NBMA.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 network-type nbma
```

Set network type of GE0/0/1 to NBMA.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 network-type nbma
```

7.5.61 ospfv3 peer router-id

Function

The **ospfv3 peer router-id** command sets the DR priority and IPv6 addresses for adjacent switches on an NBMA network.

The **undo ospfv3 peer router-id** command cancels the setting.

By default, IPv6 addresses are not configured for the adjacent switches on an NBMA network.

Format

ospfv3 peer router-id *router-id ipv6-address* [**dr-eligible** | **cost** *cost* | **poll** *interval* | **instance** *instance-id*]*

undo ospfv3 peer router-id *router-id* [*ipv6-address*] [**dr-eligible** | **cost** [*cost*] | **poll** [*interval*] | **instance** *instance-id*]*

Parameters

Parameter	Description	Value
router-id	Specifies the router ID of the adjacent switch.	The value is in dotted decimal notation.
ipv6-address	Specifies the link-local address of the adjacent node.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
dr-eligible	Specifies the neighbor switch that can take part in DR election on the NBMA network.	-
cost *cost*	Specifies the cost of the interface on the neighbor switch of the NBMA network.	The value is an integer ranging from 1 to 65535.
poll *interval*	Specifies the interval for sending polling Hello packets on the neighbor switch on the NBMA network.	The value is an integer ranging from 1 to 65535.
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

You can run the **ospfv3 peer router-id** command to set the DR priority and IPv6 addresses for adjacent switches on an NBMA network.

Example

On Vlanif10, specify the peer end to take part in DR selection.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
```



```
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 peer router-id 10.1.1.1 FE80::1 dr-eligible
```

On GE0/0/1, specify the peer end to take part in DR selection.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 peer router-id 10.1.1.1 FE80::1 dr-eligible
```

7.5.62 ospfv3 router-id auto-recover disable

Function

Using the **ospfv3 router-id auto-recover disable** command, you can disable automatic recovery that will take effect after router ID conflict is detected.

Using the **undo ospfv3 router-id auto-recover disable** command, you can enable automatic recovery that will take effect after router ID conflict is detected.

By default, automatic recovery takes effect after router ID conflict occurs.

Format

ospfv3 router-id auto-recover disable

undo ospfv3 router-id auto-recover disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

If router ID conflict occurs in an OSPFv3 area, the system can define a new router ID, preventing route flapping and reducing route calculation operations. Other protocols will not go Down when the CPU usage is controlled.

NOTE

- After the automatic recovery function is restored and router ID conflict occurs in an OSPFv3 area, the switch changes the conflicted router ID to its own IP address. Even the router ID that is manually configured will be changed.
- After the router ID is changed, the system defines a new router ID for a maximum of three times by default if router ID conflict persists.

Example

Disable automatic recovery that will take effect after router ID conflict is detected.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 router-id auto-recover disable
```

7.5.63 ospfv3 suppress-flapping peer

Function

The **ospfv3 suppress-flapping peer** command configures detection parameters for OSPFv3 neighbor relationship flapping suppression.

The **undo ospfv3 suppress-flapping peer** command restores the default detection parameters.

By default, the detection interval of OSPFv3 neighbor relationship flapping suppression is 60s, the suppression threshold is 10, and the interval for exiting from suppression is 120s.

Format

ospfv3 suppress-flapping peer { **detecting-interval** *detecting-interval* | **threshold** *threshold* | **resume-interval** *resume-interval* } * [**instance** *instance-id*]

undo ospfv3 suppress-flapping peer { **detecting-interval** *detecting-interval* | **threshold** *threshold* | **resume-interval** *resume-interval* } * [**instance** *instance-id*]

Parameters

Parameter	Description	Value
detecting-interval *detecting-interval*	Specifies the detection interval of OSPFv3 neighbor relationship flapping suppression. Each OSPFv3 interface on which OSPFv3 neighbor relationship flapping suppression is enabled starts a flapping counter. If the interval between two successive neighbor status changes from Full to a non-Full state is shorter than *detecting-interval*, a valid flapping_event is recorded, and the flapping_count is incremented by 1.	The value is an integer ranging from 1 to 300, in seconds. The default value is 60s.

Parameter	Description	Value
threshold *threshold*	Specifies the threshold of OSPFv3 neighbor relationship flapping suppression. When the flapping_count reaches or exceeds *threshold*, flapping suppression takes effect.	The value is an integer ranging from 1 to 1000. The default value is 10.
resume-interval *resume-interval*	Specifies the interval for exiting from OSPFv3 neighbor relationship flapping suppression. If the interval between two successive neighbor status changes from Full to a non-Full state is longer than *resume-interval*, the flapping_count is reset. **NOTE** The value of *resume-interval* must be greater than that of *detecting-interval*.	The value is an integer ranging from 2 to 1000, in seconds. The default value is 120s
instance *instance-id*	Specifies the ID of the instance to which an interface belongs.	The value is an integer ranging from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To configure detection parameters for OSPFv3 neighbor relationship flapping suppression on an interface, run the **ospfv3 suppress-flapping peer** command. However, keeping the default configurations is recommended.

Prerequisites

OSPFv3 neighbor relationship flapping suppression must have been enabled globally before you configure detection parameters for it. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it before you configure the detection parameters.

Example

Set the detection interval of OSPFv3 neighbor relationship flapping suppression to 5s, the suppression threshold to 40, and the interval for exiting from suppression to 20s on VLANIF 100.


```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospfv3 suppress-flapping peer detecting-interval 5 threshold 40 resume-interval
20
```

7.5.64 ospfv3 suppress-flapping peer disable

Function

The **ospfv3 suppress-flapping peer disable** command disables OSPFv3 neighbor relationship flapping suppression from an interface.

The **undo ospfv3 suppress-flapping peer disable** command enables OSPFv3 neighbor relationship flapping suppression on an interface.

By default, OSPFv3 neighbor relationship flapping suppression is enabled on all interfaces.

Format

ospfv3 suppress-flapping peer disable [**instance** *instance-id*]

undo ospfv3 suppress-flapping peer disable [**instance** *instance-id*]

Parameters

Parameter	Description	Value
instance *instance-id*	Specifies the ID of the instance to which an interface belongs.	The value is an integer ranging from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, OSPFv3 neighbor relationship flapping suppression is enabled on all interfaces in the same OSPFv3 process. To disable the function from one of the interfaces, run the **ospfv3 suppress-flapping peer disable** command.

NOTE

When an interface enters the flapping suppression state, all neighbor relationships on the interface enter the state accordingly.

Prerequisites

OSPFv3 neighbor relationship flapping suppression must have been enabled globally before you enable the function on an interface using the **undo ospfv3**

suppress-flapping peer disable command. By default, the function is enabled globally. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it first.

Example

Disable OSPFv3 neighbor relationship flapping suppression from VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospfv3 suppress-flapping peer disable
```

7.5.65 ospfv3 suppress-flapping peer hold-down

Function

The **ospfv3 suppress-flapping peer hold-down** command configures the Hold-down mode and sets duration for this mode.

The **undo ospfv3 suppress-flapping peer hold-down** command cancels the Hold-down mode.

By default, the Hold-down mode is disabled.

Format

ospfv3 suppress-flapping peer hold-down *interval* [**instance** *instance-id*]

undo ospfv3 suppress-flapping peer hold-down *interval* [**instance** *instance-id*]

Parameters

Parameter	Description	Value
interval	Specifies the duration of the Hold-down mode.	The value is an integer ranging from 1 to 600, in seconds.
instance *instance-id*	Specifies the ID of the instance to which an interface belongs.	The value is an integer ranging from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Flapping suppression works in either Hold-down or Hold-max-cost mode.

- Hold-down mode: In the case of frequent flooding and topology changes during neighbor relationship establishment, interfaces prevent neighbor relationship reestablishment during Hold-down suppression, which minimizes LSDB synchronization attempts and packet exchanges.
- Hold-max-cost mode: If the traffic forwarding path changes frequently, interfaces use 65535 as the cost of the flapping link during Hold-max-cost suppression, which prevents traffic from passing through the flapping link.

Flapping suppression can also work first in Hold-down mode and then in Hold-max-cost mode.

By default, the Hold-max-cost mode takes effect. To configure the Hold-down mode and set duration for this mode, run the **ospfv3 suppress-flapping peer hold-down** *interval* command.

Prerequisites

OSPFv3 neighbor relationship flapping suppression must have been enabled globally before you configure the Hold-down mode and set duration for this mode. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it before you configure the Hold-down mode and set duration for this mode.

Example

Configure the Hold-down mode and set its duration to 200s on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospfv3 suppress-flapping peer hold-down 200
```

7.5.66 ospfv3 suppress-flapping peer hold-max-cost disable

Function

The **ospfv3 suppress-flapping peer hold-max-cost disable** command disables the Hold-max-cost mode.

The **undo ospfv3 suppress-flapping peer hold-max-cost disable** command enables the Hold-max-cost mode.

By default, the Hold-max-cost mode is enabled.

Format

ospfv3 suppress-flapping peer hold-max-cost disable [**instance** *instance-id*]

undo ospfv3 suppress-flapping peer hold-max-cost disable [**instance** *instance-id*]

Parameters

Parameter	Description	Value
instance *instance-id*	Specifies the ID of the instance to which an interface belongs.	The value is an integer ranging from 0 to 255.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Flapping suppression works in either Hold-down or Hold-max-cost mode.

- Hold-down mode: In the case of frequent flooding and topology changes during neighbor relationship establishment, interfaces prevent neighbor relationship reestablishment during Hold-down suppression, which minimizes LSDB synchronization attempts and packet exchanges.
- Hold-max-cost mode: If the traffic forwarding path changes frequently, interfaces use 65535 as the cost of the flapping link during Hold-max-cost suppression, which prevents traffic from passing through the flapping link.

Flapping suppression can also work first in Hold-down mode and then in Hold-max-cost mode.

By default, the Hold-max-cost mode takes effect. To configure the Hold-down mode and set duration for this mode, run the **ospfv3 suppress-flapping peer hold-down** *interval* command.

NOTE

The Hold-max-cost mode can prevent a device from being isolated from the network. If a device on a key path is isolated from the network due to OSPFv3 neighbor relationship flapping, the network is separated into two isolated parts. To prevent this problem, use the Hold-max-cost mode on the key path.

Prerequisites

OSPFv3 neighbor relationship flapping suppression must have been enabled globally before you configure duration for the Hold-max-cost mode. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it before you configure duration for the Hold-max-cost mode.

Precautions

The Hold-max-cost mode takes effect only unidirectionally. If a remote device does not support OSPFv3 neighbor relationship flapping suppression, bidirectional traffic between the local and remote devices may travel along different paths.

Example

Disable the Hold-max-cost mode on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ospfv3 suppress-flapping peer hold-max-cost disable
```

7.5.67 ospfv3 timer dead

Function

The **ospfv3 timer dead** command sets the dead interval of the OSPFv3 neighbor of the instance on the interface.

The **undo ospfv3 timer dead** command restores the default value.

By default, the dead interval of OSPFv3 neighbor is 40 seconds for the interface of P2P or Broadcast type. The dead interval of OSPFv3 neighbor is 120 seconds for the interface of P2MP or NBMA type.

Format

ospfv3 timer dead *interval* [**instance** *instance-id*]

undo ospfv3 timer dead [*interval*] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
interval	Indicates the dead interval of OSPFv3.	The value is an integer ranging from 1 to 65535, in seconds.
instance *instance-id*	Specifies the interface instance ID.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Dead interval of OSPFv3 neighbor refers to that the neighbor is invalid if the neighbor does not receive the Hello packets from its neighbor in the interval. The dead interval of switches in the same network segment must be consistent.

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

By default, the dead interval of OSPF neighbors is four times the Hello packet interval. After the Hello packet interval is configured in the instance of the interface using the **ospfv3 timer hello** command, the default dead interval of the OSPFv3 neighbor of the instance on the interface is changed to be four times the Hello packet interval.

Example

Set the dead interval of the neighbor to 80 seconds for Vlanif10.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 timer dead 80
```

Set the dead interval of the neighbor to 80 seconds for GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 timer dead 80
```

7.5.68 ospfv3 timer hello

Function

The **ospfv3 timer hello** command specifies the interval for sending Hello packets on the interface instance.

The **undo ospfv3 timer hello** command restores the default interval.

By default, for the interface of the P2P and broadcast type, the interval for sending Hello packets is 10 seconds. For the interface of the P2MP and NBMA type, the interval for sending Hello packets is 30 seconds.

Format

ospfv3 timer hello *interval* [**conservative**] [**instance** *instance-id*]

undo ospfv3 timer hello [*interval* [**conservative**]] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
interval	Indicates the interval for an interface to send the Hello packets.	The value is an integer ranging from 1 to 65535, in seconds.
conservative	Indicates the conservative mode of the dead timer. If the conservative mode is configured, the value configured for the dead timer using the **ospfv3 timer dead** command takes effect even when the value is less than 10s.	-
instance *instance-id*	Specifies the interface instance ID.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Set the interval for sending Hello packets on Vlanif10 to 20 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 timer hello 20
```

Set the interval for sending Hello packets on GE0/0/1 to 20 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
```



```
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 timer hello 20
```

7.5.69 ospfv3 timer poll

Function

The **ospfv3 timer poll** command sets the polling interval for sending Hello packets on NBMA network.

The **undo ospfv3 timer poll** command restores the default polling interval.

By default, the polling interval for sending Hello packets on NBMA network is 120 seconds.

Format

ospfv3 timer poll *interval* [**instance** *instance-id*]

undo ospfv3 timer poll [*interval*] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
interval	Specifies the polling interval of Hello packets.	The value is an integer ranging from 1 to 65535, in seconds.
instance *instance-id*	Specifies the instance ID of the interface.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

On the NBMA network, if a neighbor is invalid, you can configure switches to send Hello packet based on the polling interval that is set by the **ospfv3 timer poll** command. The polling interval should be at least 4 times that of the Hello interval.

OSPFv3 does not support the configuration on a null interface.

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Set the polling interval for sending Hello packets on the Vlanif10 to 130 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 timer poll 130
```

Set the polling interval for sending Hello packets on the GE0/0/1 to 130 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 timer poll 130
```

7.5.70 ospfv3 timer retransmit

Function

The **ospfv3 timer retransmit** command specifies the interval for retransmitting the LSA on the interface instance.

The **undo ospfv3 timer retransmit** command restores the default value.

By default, the interval for retransmitting LSAs is 5 seconds.

Format

ospfv3 timer retransmit *interval* [**instance** *instance-id*]

undo ospfv3 timer retransmit [*interval*] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
interval	Indicates the interval for retransmitting LSAs.	The value is an integer ranging from 1 to 3600, in seconds.
instance *instance-id*	Specifies the interface instance ID.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

After sending an LSA to its neighbor, a switch needs to wait for the Ack packet from the neighbor. If the switch does not receive the Ack packet in the retransmission interval, the switch retransmits the LSA.

The interval for retransmitting LSA between the neighboring switches should not be set too short. Otherwise, it leads to unnecessary retransmission.

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Set the interval for retransmitting the LSA between VLANIF 10 and its neighboring switch to 12 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 timer retransmit 12
```

Set the interval for retransmitting the LSA between GE0/0/1 and its neighboring switch to 12 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 timer retransmit 12
```

7.5.71 ospfv3 trans-delay

Function

The **ospfv3 trans-delay** command sets the delay for transmitting LSA on an interface instance.

The **undo ospfv3 trans-delay** command restores the default value.

By default, the delay for transmitting LSA is 1 second.

Format

ospfv3 trans-delay *interval* [**instance** *instance-id*]

undo ospfv3 trans-delay [*interval*] [**instance** *instance-id*]

Parameters

Parameter	Description	Value
interval	Indicates the delay for transmitting LSA on an interface.	The value is an integer ranging from 1 to 800, in seconds. By default, it is 1 second.
instance *instance-id*	Specifies the interface instance ID.	The value is an integer ranging from 0 to 255. The default value is 0.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The LSA in the LSDB of the local switch ages with time. The aging of the LSA increases by 1 per second, but the aging of the LSA does not change in the transmission process. Therefore, it is necessary to add the delay to the aging time of the LSA before the LSA is sent.

Before running this command on an interface, run the **ospfv3 area** command on the interface to enable OSPFv3.

Example

Set the delay for transmitting LSA on VLANIF 10 to 3 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ospfv3 1 area 1
[HUAWEI-Vlanif10] ospfv3 trans-delay 3
```

Set the delay for transmitting LSA on GE0/0/1 to 3 seconds.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] ospfv3
[HUAWEI-ospfv3-1] router-id 10.1.1.1
[HUAWEI-ospfv3-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ospfv3 1 area 1
[HUAWEI-GigabitEthernet0/0/1] ospfv3 trans-delay 3
```


7.5.72 ospfv3 valid-ttl-hops

Function

The **ospfv3 valid-ttl-hops** command enables OSPFv3 GTSM and sets a TTL value.

The **undo ospfv3 valid-ttl-hops** command disables OSPFv3 GTSM.

By default, OSPFv3 GTSM is disabled.

Format

ospfv3 valid-ttl-hops *valid-ttl-hops-value* [**vpn-instance** *vpn-instance-name*]

undo ospfv3 valid-ttl-hops [*valid-ttl-hops-value*] [**vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
valid-ttl-hops-value	Specifies a TTL.	The value is an integer ranging from 1 to 255.
vpn-instance *vpn-instance-name*	Indicates the name of a VPN instance. If this parameter is specified, only the TTL values of the packets in the specified VPN instance are checked.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If an OSPFv3 network requires high security, you can configure OSPFv3 GTSM and an authentication mode to improve network security. During network attacks, attackers may simulate OSPFv3 packets and continuously send them to a device. If the packets are destined for the device, it directly sends them to the control plane for processing without validating them. As a result, the increased processing workload on the control plane results in high CPU usage. GTSM protects devices against potential attacks and improves system security by checking whether the TTL value in each IP packet header is within a pre-defined range.

To enable OSPFv3 GTSM, run the **ospfv3 valid-ttl-hops**. To check the TTL values of packets that match a GTSM policy in a specified VPN instance, specify **vpn-instance** in the command.

The **ospfv3 valid-ttl-hops** command enables OSPFv3 GTSM on both the public network and VPNs. For example, if you run the **ospfv3 valid-ttl-hops 5 vpn-instance** *vpn1* command, OSPFv3 GTSM takes effect on both the public network and VPN 1, the TTL values of the OSPFv3 packets in the VPN 1 are checked, and the default action is performed on the packets that fail to match the GTSM policy.

Follow-up Procedure

GTSM checks the TTL values of only the packets that match a GTSM policy. For packets that do not match the GTSM policy, you can specify **pass** in the **gtsm default-action** command to allow these packets to pass the filtering or specify **drop** in the command to discard them.

Precautions

- If a VPN instance is specified in the **ospfv3 valid-ttl-hops** command and an interface is bound to the VPN instance, the interface discards all received unicast packets if the set TTL value is less than the actual hop count on the network.
- If a virtual link or sham link is deployed, configure a TTL value based on the actual hop count on the network (the number of routers through which the virtual link or sham link passes) to prevent the local switch from discarding packets from neighbors over the virtual link or sham link.
- Therefore, if you want to apply the configured TTL value to packets only in a VPN or the public network, specify **pass** in the **gtsm default-action** command to prevent the OSPFv3 packets in other instances from being discarded incorrectly.

Example

Enable OSPFv3 GTSM and set the maximum number of TTL hops to 5 for the packets that can be received from the public network.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 valid-ttl-hops 5
```

7.5.73 preference (OSPFv3)

Function

The **preference** command sets the preference for an OSPFv3 route.

The **undo preference** command restores the default setting.

By default, the preference of the OSPFv3 route is 10. When ASE is specified, the default value is 150.

Format

preference [**ase**] { *preference* | **route-policy** *route-policy-name* }*

undo preference [**ase**]

Parameters

Parameter	Description	Value
ase	Sets the preference for an AS external route or an NSSA route.	-
preference	Specifies the preference for OSPFv3 routes.	It is an integer ranging from 1 to 255.
route-policy *route-policy-name*	Specifies the name of a routing policy and sets the preference for specified routes.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The smaller the preference value, the higher the preference.

Multiple dynamic routing protocols may be run on a switch at the same time. Consequently, there is a problem of sharing and selecting routing information among routing protocols. Therefore, the system sets a default reference for each routing protocol. If different routing protocols find multiple routes to the same destination, the route discovered by the routing protocol with a higher preference is selected to forward IP packets.

Precautions

Creating a route-policy before it is referenced is recommended. If a nonexistent route-policy is referenced using the command, the configured priority applies to all OSPFv3 routes.

Example

Set the preference for OSPFv3 routes.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] preference 5
```


7.5.74 reset ospfv3

Function

The **reset ospfv3** command restarts the OSPFv3 process.

The **reset ospfv3 counters** command resets OSPFv3 statistics.

Format

reset ospfv3 { *process-id* | **all** } [**graceful-restart** [**extend-period** *period*]]

reset ospfv3 { *process-id* | **all** } **counters** [**neighbor** [*interface-type interface-number*] [*router-id*]]

Parameters

Parameter	Description	Value
process-id	Specifies the OSPFv3 process number.	The value is an integer that ranges from 1 to 65535.
all	Restarts all OSPFv3 processes.	-
graceful-restart	Restarts the OSPFv3 process in GR mode.	-
extend-period *period*	Prolongs the current GR period.	The value is an integer that ranges from 1 to 1800, expressed in seconds.
neighbor	Resets neighbor statistics.	-
interface-type interface-number	Specifies the type and number of the interface.	-
router-id	Specifies the router ID of a neighbor.	It is in dotted decimal notation.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can choose general restart or graceful restart through the parameters.

By setting **extend-period**, you can prolong the current GR period. The prolonged GR period becomes invalid after the process restarts. This parameter takes effect for one time. That is, after the switch performs GR next time, the value of the period returns to the one set by using **graceful-restart** [**extend-period** *period*].

NOTICE

If OSPFv3 connections are reset, OSPFv3 neighbor relationships will be interrupted and the original information cannot be restored. Exercise caution before running the **reset ospfv3** command.

Example

Reset OSPFv3 statistics.

```
<HUAWEI> reset ospfv3 1 counters
```

7.5.75 reset ospfv3 suppress-flapping peer

Function

The **reset ospfv3 suppress-flapping peer** command forces an interface to exit from OSPFv3 neighbor relationship flapping suppression.

Format

reset ospfv3 *process-id* **suppress-flapping peer** [*interface-type interface-number*] [**notify-peer**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an OSPFv3 process.	The value is an integer ranging from 1 to 65535.
interface-type interface-number	Specifies an interface type and number.	-
notify-peer	Instructs neighbors to exit from OSPFv3 neighbor relationship flapping suppression.	-

Views

User view

Default Level

3: Management level

Usage Guidelines

Interfaces exit from flapping suppression in the following scenarios:

- The suppression timer expires.
- The corresponding OSPFv3 process is reset.
- An OSPFv3 neighbor is reset using the **reset ospfv3** command.
- The **reset ospfv3 suppress-flapping peer** command is run.
- OSPFv3 neighbor relationship flapping suppression is disabled globally using the **suppress-flapping peer disable (OSPFv3)** command in the OSPFv3 view.

If **notify-peer** is specified when the **reset ospfv3 suppress-flapping peer** command is run on a device, the device sends Hello packets in which **HelloInterval** and **RouterDeadInterval** are 0s to its neighbors to instruct the neighbors to exit from OSPFv3 neighbor relationship flapping suppression too. If the neighbors fail to receive such Hello packets, the function of **notify-peer** does not take effect. To force the neighbors to exit from OSPFv3 neighbor relationship flapping suppression, run the **reset ospfv3 suppress-flapping peer** command on them.

Example

Force interfaces to exit from OSPFv3 neighbor relationship flapping suppression.

```
<HUAWEI> reset ospfv3 1 suppress-flapping peer
```

7.5.76 rfc1583 compatible (OSPFv3)

Function

The **rfc1583 compatible** command converts rules defined in RFC 5340 into rules defined in RFC 1583.

The **undo rfc1583 compatible** command converts rules defined in RFC 1583 into rules defined in RFC 5340.

By default, OSPFv3 supports the routing rule of RFC 5340.

Format

rfc1583 compatible

undo rfc1583 compatible

Parameters

None

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

RFC 5340 and RFC 1583 define different OSPFv3 route selection rules. When enabling OSPFv3, configure the same route selection rules on all devices in the same OSPFv3 area. For example, an OSPFv3 device supports route selection rules defined in RFC 5340 by default. If the other routers in the same OSPFv3 area support route selection rules defined in RFC 1583, run the **rfc1583 compatible** command.

Example

Converts rules defined in RFC 5340 into rules defined in RFC 1583.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] rfc1583 compatible
```

7.5.77 route-tag (OSPFv3)

Function

The **route-tag** command sets the tag value for imported VPN routes.

The **undo route-tag** command restores the default setting.

By default, the first two bytes of the tag value are fixed as 0xD000, and the last two bytes are the local BGP AS number. For example, if the local BGP AS number is 100, the default tag value in decimal notation is 3489661028.

Format

route-tag *tag-value*

undo route-tag

Parameters

Parameter	Description	Value
tag-value	Specifies the tag value for the imported VPN route.	It is an integer ranging from 0 to 4294967295.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

On a PE, the VPN instance associated with an OSPFv3 instance is configured with a VPN route tag. This route tag must be carried in Type 5 or Type 7 LSAs. You are advised to configure the same route tag for PEs in the same area. Not transmitted in BGP extended community attributes, the VPN route tag is configured and takes effect only on the PEs that receive BGP routes and generate OSPFv3 LSAs. Different OSPFv3 processes can be configured with the same route tag.

If a BGP AS number is greater than 65535, the default tag 0 is used. You can use the command to change the tag in this case.

The only difference between the tag value set through the **route-tag** command and the tag value set through other commands is the preference:

- The preference of the tag value set through the **import-route** command is the highest.
- The preference of the tag value set through the **route-tag** command is medium.
- The preference of the tag value set through the **default tag** *tag* command is the lowest.

If a Type 5 or Type 7 LSA whose tag is the same as the local tag is received, this LSA is ignored during route calculation.

Example

Set the tag value of the imported VPN routes in OSPFv3 process 100 to 100.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance test
[HUAWEI-vpn-instance-test] ipv6-family
[HUAWEI-vpn-instance-test-af-ipv6] route-distinguisher 100:1
[HUAWEI-vpn-instance-test-af-ipv6] quit
[HUAWEI-vpn-instance-test] quit
[HUAWEI] ospfv3 100 vpn-instance test
[HUAWEI-ospfv3-100] route-tag 100
```

7.5.78 router-id (OSPFv3)

Function

The **router-id** command sets the Router ID for the switch that runs OSPFv3.

The **undo router-id** command deletes the Router ID that is set for OSPFv3 switch.

By default, there is no Router ID for the switch that runs OSPFv3.

Format

router-id *router-id*

undo router-id

Parameters

Parameter	Description	Value
router-id	Indicates the switch ID.	It is in dotted decimal notation.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

The Router ID of an OSPFv3 process is unique in an AS. If no Router ID is set, the OSPFv3 process does not run.

When you set a switch ID, ensure that the Router IDs of any two processes are different in an AS.

NOTE

Multiple OSPFv3 processes can run on a switch by setting the different process ID. In this case, you need to specify different Router IDs for different processes.

Example

Set the Router ID of the OSPFv3 process 1 to 10.1.1.3.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] router-id 10.1.1.3
```

7.5.79 sham-hello enable (OSPFv3)

Function

The **sham-hello enable** command enables OSPFv3 to regard the LSU packets and the LSAck packets the same as the Hello packets.

The **undo sham-hello** command disables the operation.

By default, the sham-hello feature is disabled.

Format

sham-hello enable

undo sham-hello

Parameters

none

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

The **sham-hello enable** command, OSPFv3 can regard the LSU packets and the LSAck packets the same as the Hello packets. After receiving this kind of packets, OSPFv3 refreshes the timeout timer of the neighbors to maintain the integrity of their relationship.

Example

Enable sham-hello.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] sham-hello enable
```

7.5.80 sham-link (OSPFv3)

Function

The **sham-link** command configures a sham link.

The **undo sham-link** command deletes a sham link or restores the default setting. If no optional parameters are specified, a sham link is deleted; if optional parameters are specified, the default values of the parameters are restored.

Format

sham-link *source-address destination-address* [**cost** *cost* | **dead** *dead-interval* | **hello** *hello-interval* | **instance** *instance-id* | **retransmit** *retransmit-interval* | **trans-delay** *trans-delay-interval* | { **authentication-mode** { **hmac-sha256 key-id** *key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* } | **keychain** *keychain-name* } | **ipsec sa** *sa-name* }] *

undo sham-link *source-address destination-address* [**cost** [*cost*] | **dead** [*dead-interval*] | **hello** [*hello-interval*] | **retransmit** [*retransmit-interval*] | **trans-delay** [*trans-delay-interval*] | { **authentication-mode** { **hmac-sha256 key-id** *key-id* | **keychain** } | **ipsec sa** [*sa-name*] }] *

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
source-address	Specifies the source IPv6 address.	-
destination-address	specifies the destination IPv6 address.	-
cost *cost*	Specifies the cost of a sham link	The value is an integer ranging from 1 to 65535. By default, it is 1.
dead *dead-interval*	Specifies the dead interval. This value must be equal to the *dead-interval* of the switch that sets up a sham link with the local switch, and must be at least four times that of *hello-interval*.	The value is an integer ranging from 1 to 65535, in seconds.
hello *hello-interval*	Specifies the interval for sending Hello packets on an interface. This value must be equal to the *hello-interval* of the switch that sets up the sham link with the local switch.	The value is an integer ranging from 1 to 65535, in seconds.
instance *instance-id*	Specifies the instance ID of a sham link.	The value is an integer ranging from 0 to 255.
retransmit *retransmit-interval*	Specifies the interval for retransmitting LSAs on an interface.	The value is an integer ranging from 1 to 3600, in seconds.
trans-delay *trans-delay-interval*	Specifies the delay for sending LSAs on an interface.	The value is an integer ranging 1 to 800, in seconds.
authentication-mode	Indicates the authentication mode over the sham link.	-
hmac-sha256	Sets the HMAC-SHA256 authentication mode.	-
key-id *key-id*	Specifies the key ID for authentication, which must be the same as the one configured at the other end.	The value is an integer ranging from 1 to 65535.

Parameter	Description	Value
plain	Configures the plaintext password type. Only a plaintext password can be entered, and the password is displayed in plaintext in the configuration file. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	-
plain-text	Specifies a plaintext password.	The value is a string of 1 to 255 characters, spaces not supported.
cipher	Configures the ciphertext password type. You can enter either a plaintext or ciphertext password, but the password is displayed in ciphertext in the configuration file.	-
cipher-text	Specifies a ciphertext password.	The value can be a string of 1 to 255 characters for plaintext passwords and 20 to 392 characters for ciphertext passwords, spaces not supported.
keychain	Configures keychain authentication. **NOTE** Before you configure keychain authentication, run the **keychain** command to configure a keychain, the **key-id** command to configure a key ID, the **key-string** command to configure a password, and the **algorithm** command to configure an algorithm. If these commands are not run, OSPFv3 authentication fails.	-

Parameter	Description	Value
keychain-name	Specifies a keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.
ipsec sa *sa-name*	Specifies the name of an SA configured for an OSPFv3 sham link.	The value is an existing SA name.

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

The **sham-link** command can be configured only in the OSPFv3 VPN process. If two PEs belong to the same area and have an intra-area route, you can set up a sham link between the two PEs so that the VPN backbone route is preferred over the intra-area route.

Example

Create an OSPFv3 sham link with the source address being FC00:0:0:1001::1 and destination address being FC00:0:0:2001::1.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1 vpn-instance vrf1
[HUAWEI-ospfv3-1] area 1
[HUAWEI-ospfv3-1-area-0.0.0.1] sham-link fc00:0:0:1001::1 fc00:0:0:2001::1
```

7.5.81 silent-interface (OSPFv3)

Function

The **silent-interface** command suppresses the specified interface from sending and receiving OSPFv3 packets.

The **undo silent-interface** command restores the default setting.

By default, the interface is permitted to send or receive OSPFv3 packets.

Format

silent-interface { **all** | *interface-type interface-number* }

undo silent-interface { **all** | *interface-type interface-number* }

Parameters

Parameter	Description	Value
all	Indicates all interfaces in a process.	-
interface-type interface-number	Indicates the interface type and interface number.	-

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

To prevent OSPFv3 routing information of a switch from being obtained by other switches on a certain network and prevent the switch from receiving routing information from other switches, you can disable an OSPFv3 interface on the switch from sending or receiving OSPFv3 packets.

After an interface on a switch is disabled from sending or receiving OSPFv3 packets, the direct route of the interface can still be advertised by the switch through an Intra-Area-Prefix-LSA, but no OSPFv3 neighbor relationship will be established on the interface. This enhances OSPFv3 adaptability.

In different processes, you can suppress the same interface from sending and receiving OSPFv3 packets. However, the **silent-interface** command takes effect only in the specified process of an interface, not all processes associated with the interface.

Example

Suppress Vlanif10 from sending and receiving OSPFv3 packets.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] silent-interface vlanif 10
```

7.5.82 spf timers

Function

The **spf timers** command sets the interval for calculating OSPFv3 routes by a common SPF timer.

The **undo spf timers** command restores the default setting.

By default, the interval for calculating OSPFv3 routes by a common SPF timer is not set.

Format

spf timers *delay-interval hold-interval*

undo spf timers

Parameters

Parameter	Description	Value
delay-interval	Specifies the delay for OSPFv3 from receiving network changes to performing SPF calculation.	The value is an integer ranging from 0 to 65535 seconds.
hold-interval	Specifies the holding interval between two consecutive SPF calculations.	The value is an integer ranging from 0 to 65535 seconds.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

According to the local LSDB, the switch that runs OSPFv3 can calculate an SPT with itself as the root and determines the next hop to the destination network according to the SPT. Adjusting SPF calculation interval can avoid the exhaustion of bandwidth and switch sources caused by frequent change of the network.

Precautions

- By default, the intelligent SPF timer takes effect.
- The **spf timers** *delay-interval hold-interval* and **spf-schedule-interval** *delay-interval hold-interval* commands have the same functions.
- The configurations of the **spf timers** *delay-interval hold-interval*, **spf-schedule-interval** *delay-interval hold-interval*, and **spf-schedule-interval intelligent-timer** *max-interval start-interval hold-interval-1* commands will override each other, and only the configuration of the last executed command takes effect.

Example

Set both the OSPFv3 route calculation interval and the hold interval to 6 seconds.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] spf timers 6 6
```

7.5.83 spf-schedule-interval (OSPFv3)

Function

The **spf-schedule-interval** command sets the interval for OSPFv3 to calculate routes.

The **undo spf-schedule-interval** command restores the default setting.

By default, the intelligent timer is enabled. The interval for the SPF calculation is expressed in milliseconds. The maximum interval for the SPF calculation is 10000 ms, the initial interval is 500 ms, and the Holdtime interval is 2000 ms.

Format

spf-schedule-interval { *delay-interval hold-interval* | **intelligent-timer** *max-interval start-interval hold-interval-1* }

undo spf-schedule-interval

Parameters

Parameter	Description	Value
delay-interval	Specifies the delay from the time when OSPFv3 receives a route change to the time when the SPF calculation is performed.	The value is an integer ranging from 0 to 65535, in seconds.
hold-interval	Specifies the hold interval between two consecutive SPF calculations.	The value is an integer ranging from 0 to 65535, in seconds.
intelligent-timer	Specifies the SPF calculation interval set through an intelligent timer.	-
max-interval	Specifies the maximum interval for performing the SPF calculation.	The value is an integer ranging from 1 to 20000, in milliseconds. The default value is 10000.
start-interval	Specifies the initial interval for performing the OSPFv3 SPF calculation.	The value is an integer ranging from 1 to 1000, in milliseconds. The default value is 500.

Parameter	Description	Value
hold-interval-1	Specifies the hold interval for performing the OSPFv3 SPF calculation.	The value is an integer ranging from 1 to 5000, in milliseconds. The default value is 2000.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Based on the local LSDB, a device that runs OSPFv3 uses the SPF algorithm to calculate the shortest path tree with itself as the root, and determines the next hop to the destination network based on the shortest path tree. Adjusting the interval at which the SPF calculation is performed can prevent too many bandwidth resources and device resources from being consumed due to frequent network changes.

Precautions

- The **spf timers** *delay-interval hold-interval* and **spf-schedule-interval** *delay-interval hold-interval* commands have the same functions.
- The configurations of the **spf timers** *delay-interval hold-interval*, **spf-schedule-interval** *delay-interval hold-interval*, and **spf-schedule-interval intelligent-timer** *max-interval start-interval hold-interval-1* commands will override each other, and only the configuration of the last executed command takes effect.

Configuration Impact

After the **spf-schedule-interval intelligent-timer** command is configured, the interval for the SPF calculation is as follows:

1. The initial interval for the SPF calculation is specified by the parameter *start-interval*.
2. The interval for the SPF calculation for the nth (n≥2) time is equal to *hold-interval-1*×$2^{(n-2)}$.
3. When the interval specified by *hold-interval-1*×$2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPFv3 performs SPF calculation at the maximum interval until *max-interval* expires without flapping or the OSPF process is restarted.

Example

Set the delay interval at which OSPFv3 route calculation is performed to 5s and the hold interval at which OSPFv3 route calculation is performed to 6s.


```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] spf-schedule-interval 5 6
```

Set the maximum interval for performing the SPF calculation to 10000 ms, the initial interval to 700 ms, and the hold interval to 4000 ms.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] spf-schedule-interval intelligent-timer 10000 700 4000
```

7.5.84 stub (OSPFv3 Area)

Function

The **stub** command sets an OSPFv3 area to the stub area.

The **undo stub** command cancels the settings.

By default, no area is set to the stub area.

Format

stub [**no-summary** | **default-route-advertise backbone-peer-ignore**] *

undo stub

Parameters

Parameter	Description	Value
no-summary	Disables an ABR from sending non-default Type 3 LSAs to the stub area. This parameter applies only to ABRs of stub areas. If the parameter is configured on an ABR, the ABR advertises only one default Type 3 LSA (no other Type 3 LSAs) to the stub area. In this case, this area is also called a totally stub area.	-
default-route-advertise	Configures an ABR to advertise a default Type 3 LSA to the stub area. This parameter applies only to ABRs of stub areas.	-
backbone-peer-ignore	Prevents an ABR from checking the neighbor status when the ABR generates a default Type 3 LSA and advertises it to the stub area. This parameter applies only to ABRs of stub areas. If the parameter is configured on an ABR, the ABR generates a default Type 3 LSA and advertises it to the stub area as long as an interface is Up in the backbone area.	-

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The stub area attribute must be configured on all switches in a stub area.

- If the **stub** command is run on an ABR, without parameters specified, the ABR advertises a default Type 3 LSA and advertises it to the stub area as long as an OSPFv3 neighbor in the full state exists in the backbone area.
- If the **stub no-summary** command is run on an ABR, the ABR is disabled from sending non-default Type 3 LSAs to the stub area connected to the ABR.
- If the **stub default-route-advertise backbone-peer-ignore** command is run on an ABR, the ABR generates a default Type 3 LSA and advertises it to the stub area as long as an interface is Up in the backbone area.

Precautions

The backbone area cannot be configured as a stub area.

Example

Set the OSPFv3 area 1 to the stub area.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] area 1
[HUAWEI-ospfv3-1-area-0.0.0.1] stub
```

7.5.85 stub-router (OSPFv3)

Function

The **stub-router** command configures a device as stub router.

The **undo stub-router** command restores the default configuration.

By default, no stub router exists.

Format

stub-router [**on-startup** [*interval*]]

undo stub-router

Parameters

Parameter	Description	Value
on-startup [*interval*]	Specifies the interval during which a device acts as a stub router when the device is restarted or faulty. • If **on-startup** is not specified, the device is always a stub router, even if the cost of all routes advertised by the device is 65535. • If **on-startup** is specified, the device works as a stub router only when it restarts or is faulty. The hold time of the stub router state is determined by *interval* parameter. If the *interval* parameter is not configured, the default interval (500 seconds) is used.	The value is an integer ranging from 5 to 65535, in seconds. By default, it is 500 seconds.

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

After the **stub-router** command is configured on a switch, the stub router informs other OSPFv3 devices not to use this stub router for data forwarding by increasing the metrics of the links in the Router-LSA generated by the switch. Since the metric is not infinite, the devices to this stub router still exist. The metrics of all links in the device Router-LSAs generated by the stub router are set to 65535.

Example

Configure the device as a stub router.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] stub-router
```

7.5.86 suppress-flapping peer disable (OSPFv3)

Function

The **suppress-flapping peer disable** command disables OSPFv3 neighbor relationship flapping suppression globally.

The **undo suppress-flapping peer disable** command enables OSPFv3 neighbor relationship flapping suppression globally.

By default, OSPFv3 neighbor relationship flapping suppression is enabled globally.

Format

suppress-flapping peer disable

undo suppress-flapping peer disable

Parameters

None

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

If an interface carrying OSPFv3 services alternates between Up and Down, OSPFv3 neighbor relationship flapping occurs on the interface. During the flapping, OSPFv3 frequently sends Hello packets to reestablish the neighbor relationship, synchronizes LSDBs, and recalculates routes. In this process, a large number of packets are exchanged, adversely affecting neighbor relationship stability, OSPFv3 services, and other OSPFv3-dependent services, such as LDP and BGP. OSPFv3 neighbor relationship flapping suppression can address this problem by delaying OSPFv3 neighbor relationship reestablishment or preventing service traffic from passing through flapping links.

By default, OSPFv3 neighbor relationship flapping suppression is enabled globally. To disable this function globally, run the **suppress-flapping peer disable** command.

Example

Disable neighbor relationship flapping suppression globally.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
[HUAWEI-ospfv3-1] suppress-flapping peer disable
```

7.5.87 vlink-peer (OSPFv3 Area)

Function

The **vlink-peer** command creates and configures a virtual link.

The **undo vlink-peer** command removes the existing virtual link.

By default, no virtual link is configured on OSPFv3.

Format

vlink-peer *router-id* [**hello** *hello-interval* | **retransmit** *retransmit-interval* | **trans-delay** *trans-delay-interval* | **dead** *dead-interval* | **instance** *instance-id* |

{ **authentication-mode** { **hmac-sha256 key-id** *key-id* { **plain** *plain-text* | [**cipher**] *cipher-text* } | **keychain** *keychain-name* } | **ipsec sa** *sa-name* }] *

undo vlink-peer *router-id* [**hello** [*hello-interval*] | **retransmit** [*retransmit-interval*] | **trans-delay** [*trans-delay-interval*] | **dead** [*dead-interval*] | { **authentication-mode** { **hmac-sha256 key-id** *key-id* | **keychain** } | **ipsec sa** [*sa-name*] }] *

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
router-id	Specifies the switch ID of a virtual link neighbor.	The value is in dotted decimal notation.
hello *hello-interval*	Specifies the interval for sending Hello packets on an interface. This value must be equal to the *hello-interval* value of the switch that sets up a virtual link with the interface.	The value is an integer ranging from 1 to 65535 seconds. By default, it is 10 seconds.
retransmit *retransmit-interval*	Specifies the interval for retransmitting the LSA packets on an interface.	The value is an integer ranging from 1 to 3600 seconds. By default, it is 5 seconds.
trans-delay *trans-delay-interval*	Specifies the delay for sending LSA packets on an interface.	The value is an integer ranging from 1 to 800 seconds. By default, it is 1 second.
dead *dead-interval*	Specifies the dead interval of the neighbor. This value must be equal to *dead-interval* of the switch that sets up the virtual link with it and must be at least four times that of the *hello-interval*.	The value is an integer ranging from 1 to 65535 seconds. By default, it is 40 seconds.
instance *instance-id*	Specifies the instance ID of the virtual link.	The value is an integer ranging from 0 to 255. By default, it is 0.
authentication-mode	Indicates the authentication mode over the virtual link.	N/A

Parameter	Description	Value
hmac-sha256	Sets the HMAC-SHA256 authentication mode.	N/A
key-id *key-id*	Specifies the key ID for authentication, which must be the same as the one configured at the other end.	The value is an integer ranging from 1 to 65535.
plain	Configures the plaintext password type. Only a plaintext password can be entered, and the password is displayed in plaintext in the configuration file. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	N/A
plain-text	Specifies a plaintext password.	The value is a string of 1 to 255 characters, spaces not supported.
cipher	Configures the ciphertext password type. You can enter either a plaintext or ciphertext password, but the password is displayed in ciphertext in the configuration file.	N/A
cipher-text	Specifies a ciphertext password.	The value can be a string of 1 to 255 characters for plaintext passwords and 20 to 392 characters for ciphertext passwords, spaces not supported.

Parameter	Description	Value
keychain	Configures keychain authentication. **NOTE** Before you configure keychain authentication, run the **keychain** command to configure a keychain, the **key-id** command to configure a key ID, the **key-string** command to configure a password, and the **algorithm** command to configure an algorithm. If these commands are not run, OSPFv3 authentication fails.	N/A
keychain-name	Specifies a keychain name.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.
ipsec sa *sa-name*	Specifies the name of an SA configured for an OSPFv3 virtual link.	The value is an existing SA name.

Views

OSPFv3 area view

Default Level

2: Configuration level

Usage Guidelines

You can use the **vlink-peer** command to set up a logical connection for non-backbone area which does not connect to the backbone area directly or discontinuous backbone area.

The virtual link can be regarded as a common interface on which OSPFv3 is enabled, because the principles of **hello**, **retransmit,** and **trans-delay** that are configured on the virtual link are similar.

Example

Create an OSPFv3 virtual link to 10.110.0.3.

```
<HUAWEI> system-view
[HUAWEI] ospfv3 1
```



```
[HUAWEI-ospfv3-1] area 10.0.0.0
[HUAWEI-ospfv3-1-area-10.0.0.0] vlink-peer 10.110.0.3
```

7.5.88 vpn-instance-capability simple (OSPFv3)

Function

The **vpn-instance-capability simple** command disables loop detection and calculates routes directly.

The **undo vpn-instance-capability** command enables the DN-bit check to prevent routing loops.

By default, loop detection is enabled.

Format

vpn-instance-capability simple

undo vpn-instance-capability

Parameters

None

Views

OSPFv3 view

Default Level

2: Configuration level

Usage Guidelines

If a Multi-VPN-Instance CE (MCE) needs to support the VPN multi-instance, loop detection needs to be disabled. The **vpn-instance-capability simple** command takes effect only in the OSPFv3 VPN instance.

Example

Disable loop detection.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vpn1
[HUAWEI-vpn1] ipv6-family
[HUAWEI-vpn1-af-ipv6] route-distinguisher 100:1
[HUAWEI-vpn1-af-ipv6] quit
[HUAWEI] ospfv3 100 vpn-instance vpn1
[HUAWEI-ospfv3-100] vpn-instance-capability simple
```

7.6 IPv4 IS-IS Configuration Commands

7.6.1 Command Support

Only the following switch models support IPv4 IS-IS:

S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S

7.6.2 adjacency-strict-check enable

Function

The **adjacency-strict-check enable** command enables IS-IS adjacency strict-check.

The **undo adjacency-strict-check enable** command disables IS-IS adjacency strict-check.

The **adjacency-strict-check disable** command disables IS-IS adjacency strict-check.

By default, IS-IS adjacency strict-check is disabled when IS-IS establishes neighbor relationships.

Format

adjacency-strict-check enable

undo adjacency-strict-check [**enable**]

adjacency-strict-check disable

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

During the establishment of IS-IS neighbor relationships, if both IPv4 and IPv6 are configured at both ends, both IPv4 and IPv6 neighbors are established. By default, IPv4 and IPv6 share the standard topology. When the faulty link is restored, IPv4 goes Up faster than IPv6. After the IS-IS router receives a message indicating that IPv4 goes Up, it considers that both IPv4 and IPv6 neighbors are established. If IPv6 packets are being transmitted at that time, some of these IPv6 packets are discarded.

To resolve this problem, run the **adjacency-strict-check enable** command to enable IS-IS adjacency strict-check to ensure that an IS-IS neighbor is established only when both IPv4 and IPv6 go Up.

Prerequisites

You have created an IS-IS process and entered the IS-IS view using the **isis** command.

Configuration Impact

After you run the **adjacency-strict-check enable** command on a broadcast network, the basic topology becomes Down if the IP protocol enabled on the local router is different from that on its neighbors.

After you run the **adjacency-strict-check enable** command on a P2P network, neighbor relationships cannot be established if the IP protocol enabled on the local router is different from that on its neighbors and only the basic topology is available for the local router.

Precautions

undo adjacency-strict-check [**enable**] and **undo adjacency-strict-check** commands have the same function, **enable** is just provide convenience for users.

Example

Enable adjacency strict-check on IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] adjacency-strict-check enable
```

7.6.3 area-authentication-mode

Function

The **area-authentication-mode** command configures an IS-IS area to authenticate received Level-1 packets (LSPs and SNPs) using the specified authentication mode and password or adds authentication information to Level-1 packets to be sent.

The **undo area-authentication-mode** command restores the default configuration.

By default, the system neither encapsulates generated Level-1 packets with authentication information nor authenticates received Level-1 packets.

Format

area-authentication-mode { { **simple** { **plain** *plain-text* | [**cipher**] *plain-cipher-text* } | **md5** { [**cipher**] *plain-cipher-text* | **plain** *plain-text* } } [**ip** | **osi**] | { **keychain** *keychain-name* } } [**snp-packet** { **authentication-avoid** | **send-only** } | **all-send-only**]

area-authentication-mode hmac-sha256 key-id *key-id* { **plain** *plain-text* | [**cipher**] *plain-cipher-text* } [**snp-packet** { **authentication-avoid** | **send-only** } | **all-send-only**]

undo area-authentication-mode

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
simple	Transmits the password in plain text. **NOTICE** Simple authentication has potential risks. HMAC-SHA256 ciphertext authentication is recommended.	-
plain *plain-text*	Specifies the authentication password in plain text. You can enter only the password in plain text. When you view the configuration file, the password is displayed in plain text. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	The value is a string of case-sensitive characters without spaces. The value contains digits and letters. When the authentication mode is **simple**, the value is a string of 1 to 16 characters. When the authentication mode is **md5** or **hmac-sha256**, the value is a string of 1 to 255 characters.
cipher *plain-cipher-text*	Specifies the authentication password in cipher text. You can enter the password in plain or cipher text. When you view the configuration file, the password is displayed in cipher text. By default, the password is in cipher text.	The value is a string of case-sensitive characters without spaces. The value contains digits and letters. When the authentication mode is **simple**, the value is a string of 1 to 16 characters in plain text or a string of 32 or 48 characters in cipher text. When the authentication mode is **md5** or **hmac-sha256**, the value is a string of 1 to 255 characters in plain text or a string of 20 to 392 characters in cipher text.

Parameter	Description	Value
md5	Transmits the password that is encrypted using MD5. **NOTICE** MD5 authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.	-
ip	Indicates the IP authentication password. This parameter cannot be configured when keychain authentication is used.	-
osi	Indicates the OSI authentication password. This parameter cannot be configured when keychain authentication is used. When neither **osi** nor **ip** is specified, the default parameter **osi** is used.	-
keychain *keychain-name*	Indicates the keychain that changes with time and is encrypted using MD5. This parameter takes effect only when *keychain-name* is set using the **keychain** command. Currently, IS-IS supports only HMAC-MD5 and HMAC-SHA256 algorithms.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.
snp-packet	Authenticates SNPs.	-
authentication-avoid	Encapsulates generated LSPs but not SNPs with authentication information and authenticates received LSPs but not SNPs.	-
send-only	Encapsulates generated LSPs and SNPs with authentication information, and authenticates received LSPs but not SNPs.	-

Parameter	Description	Value
all-send-only	Encapsulates generated LSPs and SNPs with authentication information, but does not authenticate received LSPs and SNPs.	-
hmac-sha256	Encapsulates generated packets with the HMAC-SHA256 authentication and a password encrypted by the HMAC-SHA256 algorithm and authenticates received packets.	-
key-id *key-id*	Indicates key ID of the HMAC-SHA256 algorithm.	The value is an integer ranging from 0 to 65535.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, the IS-IS packets to be sent are not encapsulated with authentication information, and the received packets are not authenticated. If a user sends malicious packets to attack a network, information on the entire network may be stolen. Therefore, you can configure IS-IS authentication to improve the network security.

The area authentication password is encapsulated into Level-1 IS-IS packets. Only the packets that pass the area authentication can be accepted. Therefore, you can configure IS-IS area authentication to authenticate the Level-1 area.

Precautions

The **area-authentication-mode** command is valid only on Level-1 or Level-1-2 routers.

By using this command enables IS-IS to discard all the Level-1 LSPs and SNPs whose area authentication passwords are not consistent with the one set by this command. At the same time, IS-IS inserts the configured area authentication password into all the Level-1 routing packets (LSPs and SNPs) sent from the local node. The establishment of the Level-1 neighbor relationship is not affected, regardless of whether the packets pass the area authentication.

The authentication takes effect only on the peer configured with authentication. The peer with no authentication configured can still receive the LSP and SNP packet with the password.

Example

Set the area authentication mode to keychain and keychain name to **test**.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] area-authentication-mode keychain test
```

7.6.4 attached-bit advertise

Function

The **attached-bit advertise** command configures a rule for setting the ATT bit in Link state Protocol Data Units (LSPs).

The **undo attached-bit advertise** command restores the default setting rule.

By default, the Level-1-2 device sets the ATT bit in LSPs using the default rule.

Format

attached-bit advertise { **always** | **never** }

undo attached-bit advertise

Parameters

Parameter	Description	Value
always	Indicates that the ATT bit is set to 1.	-
never	Indicates that the ATT bit is set to 0.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An ATT bit is a field in link state packets. An ATT bit identifies whether Level-1 area is attached to other areas. A Level-1-2 device sets the ATT bit to 1 to notify the devices in Level-1 area of its attachment to a Level-2 backbone area. After a Level-1 device receives the LSPs carrying the ATT bit 1 from the Level-1-2 device, the Level-1 device generates a route with the same destination address as the default route of the Level-1-2 device. Traffic can be forwarded along this route.

The preceding rule is the default rule for setting the ATT bit in the Intermediate System to Intermediate System (IS-IS) protocol. You can run the **attached-bit advertise** command to configure a rule for setting the ATT bit as required.

- If you want the ATT bit in LSPs to be set to 1, run the **attached-bit advertise always** command.
- If you want to disable the Level-1 device connected to the Level-1-2 device from advertising default routes when the ATT bit is set to 1, run either of the following commands:
  - Run the **attached-bit advertise never** command on the Level-1-2 device to disable the device from sending LSPs with the ATT bit 1.
  - Run the **attached-bit avoid-learning** command on the Level-1 device that is connected to the Level-1-2 device.

  Running the **attached-bit advertise never** command applies to all Level-1 devices that receive LSPs with the ATT bit 0. Therefore, it is more convenient.

Precautions

The ATT bit is defined in both Level-1 and Level-2 LSPs, but it can be set to 1 only in Level-1 LSPs. Additionally, only the Level-1-2 device can set the ATT bit. Therefore, running the **attached-bit avoid-learning** command takes effect only on the Level-1-2 device.

Example

Enable the Level-1-2 device to set the ATT bit to 1 in the LSPs sent in IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] attached-bit advertise always
```

7.6.5 attached-bit avoid-learning

Function

The **attached-bit avoid-learning** command disables a device from advertising default routes to a routing table when the ATT bit is set to 1.

The **undo attached-bit avoid-learning** command restores the device to generate default routes when the ATT bit is set to 1.

By default, a device advertises default routes when the ATT bit is set to 1.

Format

attached-bit avoid-learning

undo attached-bit avoid-learning

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

An ATT bit is a field in link state packets. An ATT bit identifies whether Level-1 area is attached to other areas. A Level-1-2 device sets the ATT bit to 1 to notify the devices in Level-1 area of its attachment to a Level-2 backbone area. After a Level-1 device receives the LSPs carrying the ATT bit 1 from the Level-1-2 device, the Level-1 device generates a route with the same destination address as the default route of the Level-1-2 device. Traffic can be forwarded along this route.

To prevent the Level-1 device from advertising the default route when the ATT bit is set to 1, run the **attached-bit avoid-learning** command.

To prevent the Level-1 device from advertising default routes to the routing table, run either of the following commands:

- Run the **attached-bit advertise never** command on the Level-1-2 device to disable the device from sending LSPs with the ATT bit 1.
- Run the **attached-bit avoid-learning** command on the Level-1 device that is connected to the Level-1-2 device.

The difference between the preceding commands lies in that the **attached-bit avoid-learning** command applies to a specified device.

Example

Disable IS-IS-enabled device from advertising default routes to a routing table when the ATT bit is set to 1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] attached-bit avoid-learning
```

7.6.6 auto-cost enable

Function

The **auto-cost enable** command enables IS-IS to automatically calculate the interface cost based on the interface bandwidth.

The **undo auto-cost enable** command disables IS-IS from automatically calculating the interface cost based on the interface bandwidth.

By default, IS-IS is disabled from automatically calculating the interface cost based on the interface bandwidth.

Format

auto-cost enable [**compatible**]

undo auto-cost enable

Parameters

Parameter	Description	Value
compatible	Specifies the IS-IS to calculate the cost of an interface based on the bandwidth of the interface automatically in compatible.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After this command is run, if the cost of an IS-IS interface is not set in the interface view and no cost is set in the IS-IS view, the system automatically calculates a cost for the interface.

If the cost style of the system is wide or wide-compatible:

When **auto-cost enable** command is configured, Interface cost = (Bandwidth-reference/Link-bandwidth) x 10.

The cost style is set by the **cost-style** command. The Bandwidth-reference is set by the **bandwidth-reference** command. The Link-bandwidth is the interface bandwidth.

If the cost style of the system is narrow, narrow-compatible or compatible, the cost of each interface is based on the following table.

Table 7-87 Relationship between the IS-IS interface cost and the bandwidth

Cost	Range of the Interface Bandwidth
60	Interface bandwidth =< 10 Mbit/s
50	10 Mbit/s < Interface bandwidth ≤ 100 Mbit/s
40	100 Mbit/s < Interface bandwidth ≤ 155 Mbit/s
30	155 Mbit/s < Interface bandwidth ≤ 622 Mbit/s

Cost	Range of the Interface Bandwidth
20	622 Mbit/s < Interface bandwidth ≤ 2.5 Gbit/s
10	2.5 Gbit/s < Interface bandwidth

Precautions

The priority of the cost value of the global configured by the **circuit-cost** command is higher than the auto cost value.

The **auto-cost enable** command cannot change the cost of the loopback interface.

Example

Enable IS-IS to automatically calculate the interface cost.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] auto-cost enable
```

7.6.7 bandwidth-reference (IS-IS)

Function

The **bandwidth-reference** command sets the bandwidth reference value that is used in automatic IS-IS interface cost calculation.

The **undo bandwidth-reference** command restores the default bandwidth reference value that is used in automatic IS-IS interface cost calculation.

By default, the bandwidth reference value is 100 Mbit/s.

Format

bandwidth-reference *value*

undo bandwidth-reference

Parameters

Parameter	Description	Value
value	Specifies the bandwidth reference value used in automatic IS-IS interface cost calculation.	The value is an integer that ranges from 1 to 2147483648, in Mbit/s.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To make the set link cost of IS-IS routes reflect the actual link cost, configure a proper link cost for all interfaces or enable automatic interface cost calculation. To enable automatic interface cost calculation, set a proper bandwidth reference value.

You can run the **bandwidth-reference** command to set a proper bandwidth reference value. After automatic interface cost calculation is enabled, and the cost type is wide or wide-compatible, the system calculates the cost of each interface with the following formula: Interface cost = (Bandwidth-reference/Interface bandwidth) x 10. If the interface cost calculated through this formula is greater than 16777214, 16777214 is used as the interface cost for route calculation. That is, the interface cost will never be greater than 16777214.

The **bandwidth** *bandwidth* command can only set an interface bandwidth obtained by the NMS from the MIB. It cannot change an interface actual bandwidth and interface cost.

After the **auto-cost enable** command enables automatic IS-IS interface cost calculation, the system automatically calculates the interface cost based on the bandwidth reference value set using the **bandwidth-reference** command.

The **bandwidth-reference** command can be operated on Eth-Trunk interfaces as same with on physical interfaces. If the command is run on an Eth-Trunk interface, the bandwidth of the Eth-Trunk interface is equal to the total bandwidth of all its member interfaces.

Precautions

Rules for automatically calculating the IS-IS interface cost vary according to the cost style of the IS-IS interface:

- The bandwidth reference value set using the **bandwidth-reference** command takes effect only when the cost style is wide or wide-compatible. In this case, the interface cost is calculated using the following formula:

  Interface cost = (Bandwidth-reference/Interface bandwidth) × 10.
- If the cost style of the system is narrow, narrow-compatible or compatible, the cost of each interface is based on the following table.

Table 7-88 Relationship between the IS-IS interface cost and the bandwidth

Cost	Range of the Interface Bandwidth
60	Interface bandwidth =< 10 Mbit/s
50	10 Mbit/s < Interface bandwidth =< 100 Mbit/s

Cost	Range of the Interface Bandwidth
40	100 Mbit/s < Interface bandwidth =< 155 Mbit/s
30	155 Mbit/s < Interface bandwidth =< 622 Mbit/s
20	622 Mbit/s < Interface bandwidth =< 2.5 Gbit/s
10	2.5 Gbit/s < Interface bandwidth

Example

Set the reference value of the system bandwidth to 1000 Mbit/s.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] bandwidth-reference 1000
```

7.6.8 bfd all-interfaces (IS-IS)

Function

The **bfd all-interfaces** command sets values for BFD session parameters used in BFD for IS-IS.

The **undo bfd all-interfaces** command restores the default values of BFD session parameters used in BFD for IS-IS.

By default, the minimum intervals for receiving and sending BFD packets are 1000 ms and the detection time multiplier is 3.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

bfd all-interfaces { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* | **frr-binding** } *

undo bfd all-interfaces { **min-rx-interval** [*receive-interval*] | **min-tx-interval** [*transmit-interval*] | **detect-multiplier** [*multiplier-value*] | **frr-binding** } *

Parameters

Parameter	Description	Value
min-rx-interval *receive-interval*	Specifies the minimum interval for receiving BFD packets from the peer end. The interval for receiving BFD packets between determines the BFD session detection time. On an unstable link, a smaller receive interval may result in BFD session flapping. To prevent BFD session flapping, increase the receive interval.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *transmit-interval*	Specifies the minimum interval for transmitting BFD packets to the peer end. The interval for transmitting BFD packets determines the BFD session detection time. On a stable link, you can increase the transmit interval to prevent frequent link status detection.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier-value*	Specifies the local detection multiplier. On a stable link, you can increase the BFD detection multiplier to prevent frequent link status detection.	The value is an integer that ranges from 3 to 50. The default value is 3.

Parameter	Description	Value
frr-binding	Binds the BFD session status to IS-IS Auto FRR. When BFD detects a link fault on an interface, the BFD session goes Down, triggering FRR on the interface. Then traffic is switched from the faulty link to the backup link to ensure traffic forwarding.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

During the establishment of a BFD session, you can adjust the interval for transmitting BFD packets, interval for receiving BFD packets, and local detection multiplier according to network and performance requirements.

Precautions

The value of *receive-interval* is obtained by negotiating the local **min-rx-interval** value and the remote **min-tx-interval** value. If a router does not receive any BFD packet from its peer in the *receive-interval* × *multiplier-value* period, the router advertises that its peer is Down.

If only parameters of a BFD session are set but the **bfd all-interfaces enable** command is not used, the BFD session cannot be set up.

NOTE

The BFD priority of the interface is higher than the BFD priority of the process. If BFD of the interface is enabled, the BFD session is set up based on the BFD parameters on the interface.

Example

Enable BFD in an IS-IS process and set the minimum interval for transmitting BFD packets to 600 ms.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] isis
[HUAWEI-isis-1] bfd all-interfaces enable
[HUAWEI-isis-1] bfd all-interfaces min-tx-interval 600
```


7.6.9 bfd all-interfaces enable

Function

The **bfd all-interfaces enable** command enables BFD in an IS-IS process.

The **undo bfd all-interfaces enable** command disables BFD in an IS-IS process.

By default, BFD is disabled in an IS-IS process.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

bfd all-interfaces enable

undo bfd all-interfaces enable

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If global BFD is not enabled, you can configure IS-IS BFD but cannot set up the BFD session. Before establishing BFD sessions, you need to enable BFD on the two ends.

Precautions

When BFD is enabled in the process using the **bfd all-interfaces enable** command, the interface does not set up a BFD session in the following cases:

- The **isis bfd block** command is used on the interface. The command suppresses the BFD capability of the interface. To set up a session on the interface, run the **undo isis bfd block** command.
- When the **isis bfd static** command is used on the interface, the interface does not set up the BFD session. To set up a session on the interface, run the **undo isis bfd static** command.

NOTE

The BFD priority of the interface is higher than the BFD priority of the process. If BFD of the interface is enabled, a BFD session is set up based on the BFD parameters on the interface.

Example

Enable BFD in an IS-IS process.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] isis
[HUAWEI-isis-1] bfd all-interfaces enable
```

7.6.10 circuit-cost

Function

The **circuit-cost** command sets the link cost for all IS-IS interfaces during SPF calculation.

The **undo circuit-cost** command deletes the configured link cost of all IS-IS interfaces.

By default, no link cost is set for IS-IS interfaces.

Format

circuit-cost { *cost* | **maximum** } [**level-1** | **level-2**]

undo circuit-cost [*cost* | **maximum**] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
cost	Specifies the interface cost used in SPF calculation.	If the IS-IS cost style is wide or wide-compatible, the cost of imported routes ranges from 1 to 1677721416777214; otherwise, the cost ranges from 1 to 63.
maximum	Sets the link cost of IS-IS interfaces to 16777215. **NOTE** You can configure this parameter only when the IS-IS cost style is wide or wide-compatible. After the interface cost is set to 16777215, the neighbor TLV generated on the link can only be used to transmit TE information but cannot be used for route calculation.	-
level-1	Specifies the link cost for all Level-1 interfaces. If no level is specified, the link cost is set for Level-1-2 interfaces.	-

Parameter	Description	Value
level-2	Specifies the link cost for all Level-2 interfaces. If no level is specified, the link cost is set for Level-1-2 interfaces.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a large network, there are often multiple valid routes to the same destination. IS-IS can use the SPF algorithm to calculate the optimal route to forward traffic. This may cause the following problems:

- All traffic is forwarded along the optimal path, resulting in load unbalance.
- If the optimal path is intermittently disconnected, traffic is still forwarded along the optimal path, resulting in traffic loss.

To solve the preceding problems, run the **circuit-cost** command to change the cost of all IS-IS interfaces so that traffic can be forwarded along different physical links.

Precautions

The **isis cost** command takes precedence over the **circuit-cost** command. The **circuit-cost** command cannot change the cost of loopback interfaces.

Changing the link cost of interfaces will cause routes of the entire network to be recalculated and change the forwarding path of traffic.

Example

Set the default cost of all Level-1-2 interfaces to 30.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] circuit-cost 30
```

Set the link cost of all IS-IS interfaces to 16777215.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] circuit-cost maximum
```

7.6.11 circuit default-tag

Function

The **circuit default-tag** command sets the administrative tag value for all interfaces in an IS-IS process.

The **undo circuit default-tag** command restores the default administrative tag value for all interfaces in an IS-IS process.

By default, the administrative tag value of all IS-IS interfaces is 0.

Format

circuit default-tag *tag* [**level-1** | **level-2**]

undo circuit default-tag [*tag*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
tag	Specifies the administrative tag value for all interfaces in an IS-IS process.	The value is an integer that ranges from 1 to 4294967295.
level-1	Specifies the administrative tag value for all Level-1 interfaces in an IS-IS process. If no interface level is specified, the administrative tag values of Level-1 and Level-2 interfaces are set.	-
level-2	Specifies the administrative tag value for all Level-2 interfaces in an IS-IS process. If no interface level is specified, the administrative tag values of Level-1 and Level-2 interfaces are set.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The administrative tag carries management information about IP address prefixes to control the import of routes of different levels and areas. You can use the administrative tag to control the advertisement of IP address prefixes in an IS-IS routing domain to simplify route management.

The **circuit default-tag** command allows you to set the administrative tag value for all routes in an IS-IS process. You can filter routes based on the configured administrative tag value.

The administrative tag for an interface is advertised along with routing information.

- When the cost style of IS-IS is narrow or narrow-compatible, the administrative tag value is neither advertised nor takes effect in the LSP.
- When the cost style of IS-IS is wide, wide-compatible, or compatible, the administrative tag value is advertised in the LSP.

Precautions

The value of the administrative tag configured through the**circuit default-tag** command is the global administrative tag. The priority of the interface administrative tag configured by the **isis tag-value** command is higher than the priority of the global administrative tag.

Example

Set the administrative tag value of Level-1 interfaces to 30.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] circuit default-tag 30 level-1
```

7.6.12 cost-style

Function

The **cost-style** command sets the cost style of routes received and sent by an IS-IS device.

The **undo cost-style** command restores the default cost style of routes received and sent by an IS-IS device.

By default, the cost style of routes received and sent by an IS-IS device is narrow.

Format

cost-style { **narrow** | **wide** | **wide-compatible** | { **compatible** | **narrow-compatible** } [**relax-spf-limit**] }

undo cost-style

Parameters

Parameter	Description	Value
narrow	Configures an IS-IS device to receive and send only routes with cost style narrow. When the cost style is narrow, the cost of routes ranges from 1 to 63.	-
wide	Configures an IS-IS device to receive and send only routes with cost style wide. When the cost style is wide, the cost of routes ranges from 1 to 16777215.	-
wide-compatible	Configures an IS-IS device to receive routes with cost style narrow or wide and sent only routes with cost style wide.	-
compatible	Configures an IS-IS device to receive and send routes with cost style narrow or wide.	-

Parameter	Description	Value
narrow-compatible	Configures an IS-IS device to receive routes with cost style narrow or wide and sent only routes with cost style narrow.	-
relax-spf-limit	Configures an IS-IS device to receive routes with cost higher than 1023. If this parameter is specified, there is no restriction on the link costs of interfaces or route costs. The cost of a received route is the actual cost. If this parameter is not specified, the following situations occur: ● If the cost of a route is smaller than or equal to 1023, and the link cost of each interface through which the route passes is smaller than or equal to 63: The cost of the route is the sum of link costs of all interfaces through which the route passes. ● If the cost of a route is smaller than or equal to 1023, and the link costs of certain interfaces through which the route passes are larger than 63: An IS-IS device can learn only the routes imported by the interfaces and direct routes of the network segment where the interfaces reside. The costs of these routes are the actual cost. ● If the cost of a route is larger than 1023: The device can receive only the routes of the network segment where the interface with link cost smaller than 1023 resides. If the cost of a route is larger than 1023, the cost is calculated as 1023.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, the cost style of IS-IS routes is narrow, indicating that only routes with the cost ranging from 1 to 63 can be received and sent. LSPs carrying tag information to be filtered by a routing policy cannot be flooded when the cost style is narrow. The **cost-style** command can be used to change the cost style of IS-IS routes so that they can be transmitted successfully.

To implement extended IS-IS functions, setting the cost style of IS-IS routes to **wide** is recommended.

Precautions

When the cost style of a route is changed from wide to narrow, transmission of the route may be interrupted.

If you want to change the cost style of IS-IS routes, running the command while configuring basic IS-IS functions is recommended. If the cost style of IS-IS routes is changed during network operation, the IS-IS process is restarted and neighbors are disconnected.

Example

Configure an IS-IS device to send only the packets with cost style narrow and to receive the packets with cost style narrow or wide.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] cost-style narrow-compatible
```

7.6.13 default-route-advertise (IS-IS)

Function

The **default-route-advertise** command configures IS-IS devices to generate default routes.

The **undo default-route-advertise** command disables IS-IS devices from generating default routes.

By default, IS-IS devices do not generate default routes.

Format

default-route-advertise [**always** | **match default** | **route-policy** *route-policy-name*] [**cost** *cost* | **tag** *tag* | [**level-1** | **level-1-2** | **level-2**]] * [**avoid-learning**]

undo default-route-advertise

Parameters

Parameter	Description	Value
always	Configures an IS-IS device to unconditionally advertise default routes with itself as the next hop.	-
match default	Advertises a default route through an LSP if the default route is generated by other routing protocols or other IS-IS processes in the routing table. If this default route is deleted from the routing table, it is not advertised through an LSP.	-

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a route-policy. A Level-1-2 device advertises default routes to the IS-IS routing domain only when there are external routes matching the route-policy in the routing table of the device. This prevents routing blackhole when link faults make some important external routes unavailable but default routes are still advertised. This route-policy does not affect external route import in IS-IS.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
cost *cost*	Specifies the cost of default routes.	The value is an integer. The value range varies according to the cost style. When the cost style is narrow, narrow-compatible, or compatible, the value ranges from 0 to 63. When the cost style is wide or wide-compatible, the value ranges from 0 to 4261412864.
tag *tag*	Specifies the tag value of advertised default routes. Advertised LSPs carry the tag value only when the IS-IS cost style is wide, wide-compatible, or compatible.	The value is an integer that ranges from 1 to 4294967295.
level-1	Sets the level of default routes to Level-1. If no level is specified, Level-2 default routes are generated by default.	-
level-2	Sets the level of default routes to Level-2. If no level is specified, Level-2 default routes are generated by default.	-
level-1-2	Sets the level of default routes to Level-1-2. If no level is specified, Level-2 default routes are generated by default.	-
avoid-learning	Prevents an IS-IS process from learning default routes or adding them to the routing table. If existing default routes in the routing table are active, set the default route that needs to be added to the routing table to inactive.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When IS-IS and other routing protocols are running on a network, use either of the following methods to allow the traffic within the IS-IS domain to be transmitted to other routing domains:

- Configure a boundary router in the IS-IS domain to advertise default routes to the IS-IS domain.
- Configure a boundary router in the IS-IS domain to import routes from the other routing domains to the IS-IS domain.

The first method is simpler, because the routers in the IS-IS domain do not need to learn routes imported form the other routing protocols.

Precautions

You can specify **always** to allow default routes to be advertised unconditionally; in this case, the device still calculates default routes from other devices.

NOTE

If **always** is configured on multiple devices within the same area, a routing loop may occur.

After this command is run on an IS-IS router, all traffic in an IS-IS domain will be forwarded by this IS-IS router to a destination outside the domain. Compared with configuring a static default route on each router in an IS-IS domain, running this command simplifies configurations, because this command only needs to be run on a boundary router in the IS-IS domain. In addition, you can specify different parameters to allow default routes to be advertised in different ways.

If this command is run on a Level-1 router, the router advertises default routes only to the Level-1 area.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, an ABR advertises the default route to the IS-IS domain as long as the local routing table contains external routes.

Example

Configure the current IS-IS device to advertise the IPv4 default routes that match route-policy **filter** and set the cost of these default routes to 15.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] default-route-advertise route-policy filter cost 15
```


7.6.14 description (IS-IS)

Function

The **description** command configures a description for an IS-IS process.

The **undo description** command deletes the description of an IS-IS process.

By default, no description is configured for an IS-IS process.

Format

description *description*

undo description

Parameters

Parameter	Description	Value
description	Configures a description for an IS-IS process.	The value is a string of 1 to 80 case-sensitive characters. Spaces are supported.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Configuring descriptions for IS-IS processes helps identify and maintain IS-IS processes.

IS-IS process description that is configured through the **description** command is not advertised in an LSP.

IS-IS process description that is configured through the **is-name** command is advertised in an LSP.

Precautions

If you run the **description** command multiple times, only the latest configuration takes effect.

Example

Configure description for IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
```



```
[HUAWEI-isis-1] description this process configure the area-authentication-mode
```

7.6.15 display default-parameter isis

Function

The **display default-parameter isis** command displays the default IS-IS configuration.

Format

display default-parameter isis

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display default-parameter isis** command to check the default IS-IS configuration in IS-IS initialization.

Example

Display the default IS-IS configuration.

```
<HUAWEI> display default-parameter isis

                  Default Configurations For Process
                  ----------------------------------
 Cost-Style                          : Narrow
 Circuit-Cost <IPv4>                   : L1 10 L2 10
 Circuit-Cost <IPv6>                   : L1 10 L2 10
 IS-Level                           : L12
 LSP-Originate-Length                  : 1497
 LSP-Receive-Length                  : 1497
 LSP-Max-Age <s>                     : 1200
 LSP-Generation-IntelliTimer <s,ms,ms> : L1 Max 2 Init 0 Incr 0
                                 L2 Max 2 Init 0 Incr 0
 LSP-Refresh-Interval <s>              : 900
 Preference                          : IPv4 15 IPv6 15
 SPF-IntelliTimer <s,ms,ms>            : Max 5 Init 50 Incr 200

                Default Configurations For Interfaces
                -------------------------------------
 Circuit-Level                       : L12
 CSNP-Interval <s>                     : L1 10 L2 10
 Cost <IPv4>                         : L1 10 L2 10
 Cost <IPv6>                         : L1 10 L2 10
 DIS-Priority                        : L1 64 L2 64
 Hello-Interval <s>                    : L1 10 L2 10
```



```
Holding-Multiplier                : L1 3 L2 3
LSP-Retransmit <s>                 : 5
LSP-Throttle <ms>                 : 50 count 10
PPP-Negotiation                   : 3-Way
Peer Suppress-Flapping <s>          : Detect-Interval 60
                                  Resume-Interval 120
                                  Threshold       10
```

Table 7-89 Description of the display default-parameter isis command output

Item	Description
Default Configurations For Process	Default configuration of an IS-IS process.
Cost-Style	IS-IS cost style: ● narrow ● wide ● wide-compatible ● compatible ● narrow-compatible To set the IS-IS cost style, run the **cost-style** command.
Circuit-Cost <IPv4>	Cost of all IS-IS IPv4 interfaces. To set the cost of IS-IS IPv4 interfaces, run the **circuit-cost** command.
Circuit-Cost <IPv6>	Cost of all IS-IS IPv6 interfaces. To set the cost of IS-IS IPv6 interfaces, run the **ipv6 circuit-cost** command.
IS-Level	Level of an IS-IS device: ● L1: Level-1 device ● L2: Level-2 device ● L12: Level-1-2 device To set the level of an IS-IS device, run the **is-level** command.
LSP-Originate-Length	Maximum length of an LSP generated by IS-IS. To set the maximum length of an LSP generated by IS-IS, run the **lsp-length** command.
LSP-Receive-Length	Maximum length of an LSP received by IS-IS. To set the maximum length of an LSP received by IS-IS, run the **lsp-length** command.

Item	Description
LSP-Max-Age <s>	Maximum lifetime of an LSP generated by an IS-IS process, in seconds. To set the maximum lifetime of an LSP generated by an IS-IS process, run the **timer lsp-max-age** command.
LSP-Generation-IntelliTimer <s,ms,ms>	Delay in generating LSPs: ● L1: delay for a Level-1 router to generate LSPs ● L2: delay for a Level-2 router to generate LSPs ● Max: maximum delay in generating LSPs, in seconds ● Init: initial delay in generating LSPs, in milliseconds ● Incr: incremental delay in generating LSPs, in milliseconds To set the delay in generating LSPs, run the **timer lsp-generation** command.
LSP-Refresh-Interval <s>	Interval for refreshing LSPs, in seconds. To set the interval for refreshing LSPs, run the **timer lsp-refresh** command.
Preference	Protocol preference of IS-IS routes. To set the protocol preference of IS-IS routes, run the **preference** command.
SPF-IntelliTimer <s,ms,ms>	Delay in SPF calculation: ● Max: maximum delay in SPF calculation, in seconds ● Init: initial delay in SPF calculation, in milliseconds ● Incr: incremental delay in SPF calculation, in milliseconds To set the delay in SPF calculation, run the **timer spf** command.
Default Configurations For Interfaces	Default configuration of IS-IS interfaces.
Circuit-Level	IS-IS interface level. To set the IS-IS interface level, run the **isis circuit-level** command.
CSNP-Interval <s>	Interval for sending CSNPs, in seconds. To set the interval for sending CSNPs, run the **isis timer csnp** command.

Item	Description
Cost <IPv4>	Cost of IS-IS interfaces in an IPv4 topology. To set the cost of IS-IS interfaces in an IPv4 topology, run the **isis cost** command.
Cost <IPv6>	Cost of IS-IS interfaces in an IPv6 topology. To set the cost of IS-IS interfaces in an IPv6 topology, run the **isis ipv6 cost** command.
DIS-Priority	Priority for DIS election. To set the priority for DIS election, run the **isis dis-priority** command.
Hello-Interval <s>	Interval for sending Hello packets, in seconds. To set the interval for sending Hello packets, run the **isis timer hello** command.
Holding-Multiplier	Number of consecutive Hello packets that are not received before the neighbor is declared Down. To set this value, run the **isis timer holding-multiplier** command.
LSP-Retransmit <s>	Interval for retransmitting LSPs on a P2P link, in seconds. To set the retransmit interval, run the **isis timer lsp-retransmit** command.
LSP-Throttle <ms>	Interval for sending LSPs or CSNPs, in milliseconds, and number of LSPs or CSNPs that are sent each time. To set the two values, run the **isis timer lsp-throttle** command.
PPP-Negotiation	PPP negotiation type during the establishment of IS-IS adjacencies: • 2-way • 3-way • only To set the PPP negotiation type during the establishment of IS-IS adjacencies, run the **isis ppp-negotiation** command.

Item	Description
Peer Suppress-Flapping <s>	Detection parameters for IS-IS neighbor relationship flapping suppression on an interface: • Detect-Interval: detection interval of IS-IS neighbor relationship flapping suppression • Resume-Interval: interval for exiting from IS-IS neighbor relationship flapping suppression • Threshold: threshold of IS-IS neighbor relationship flapping suppression

7.6.16 display isis cost interface

Function

The **display isis cost interface** command displays costs of an interface and how they are generated.

Format

display isis cost interface *interface-type interface-number*

Parameters

Parameter	Description	Value
interface-type interface-number	Specifies the type and number of an interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis cost interface** command to check costs of an IS-IS interface, including the link cost and route prefix cost that the interface uses in each topology, and how the costs are generated.

Example

Display costs of VLANIF 100 and how they are generated.

```
<HUAWEI> display isis cost interface vlanif 100
Interface: Vlanif 100
                    Level-2 interface cost
 Topology base(0):
  Link effective cost: 10(A)                    enabled by circuit cost
  IP prefix effective cost:
   10.10.7.0/24               cost: 10         enabled by circuit cost
   10.10.8.0/24               cost: 10         enabled by circuit cost
   10.10.9.0/24               cost: 10         enabled by circuit cost
 Topology red(7):
  Link effective cost: 87(A)                    enabled by circuit cost
  IP prefix effective cost:
   10.10.7.0/24               cost: 87         enabled by circuit cost
   10.10.8.0/24               cost: 87         enabled by circuit cost
   10.10.9.0/24               cost: 87         enabled by circuit cost

             Flags: R-Relative cost  A-Absolute cost
```

Table 7-90 Description of the **display isis cost interface** command output

Item	Description
Interface	Type and ID of an interface on which IS-IS is enabled
Link effective cost	Link cost
IP prefix effective cost	Cost of an IPv4 route
cost	Interface cost
enabled by circuit cost	Default cost or cost configured in the **isis cost** command
enabled by auto cost	Cost configured in the **auto-cost enable** command
enabled by global cost	Cost configured in the **circuit cost** command
enabled by igp ldp	The interface cost is configured by LDP
enabled by rui cost	Cost that inherits the cost of the RUI route
enabled by silent cost	Cost configured in the **isis silent** command in the interface view
enabled by tunnel cost	Cost generated after the TE Tunnel is configured
IPv6 prefix effective cost	Cost of an IPv6 route
Flags	Cost flag ● R-Relative cost: relative cost ● A-Absolute cost: absolute cost

7.6.17 display isis bfd interface

Function

The **display isis bfd interface** command displays information about BFD-enabled IS-IS interfaces.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display isis [*process-id*] **bfd interface**

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

BFD can provide millisecond-level fault detection. It can work with IS-IS to fast detect faults on neighboring devices and instruct IS-IS to recalculate routes for correct packet forwarding. You can run the **display isis bfd interface** command to check information about BFD-enabled IS-IS interfaces.

Example

Display information about BFD-enabled IS-IS interfaces.

```
<HUAWEI> display isis 1 bfd interface

                BFD information of interface for ISIS(1)
               ----------------------------------------
Interface               BFD.State      Min-Tx       Min-Rx       Mul
Vlanif101                 enable        1000         1000         3
Total interfaces: 1                 Total bfd enabled interfaces: 1
```

Table 7-91 Description of the display isis bfd interface command output

Item	Description
Interface	BFD-enabled IS-IS interface. To enable BFD on an IS-IS interface, run the **bfd all-interfaces enable** or **isis bfd enable** command.
BFD.State	BFD status on the IS-IS interface: • enable: BFD is enabled on the interface. • disable: BFD is disabled on the interface. To enable BFD on an IS-IS interface, run the **bfd all-interfaces enable** or **isis bfd enable** command.
Min-Tx	Minimum interval for transmitting BFD packets. To set the minimum interval, run the **bfd all-interfaces** or **isis bfd** command.
Min-Rx	Minimum interval for receiving BFD packets. To set the minimum interval, run the **bfd all-interfaces** or **isis bfd** command.
Mul	Local detection multiplier. To set the local detection multiplier, run the **bfd all-interfaces** or **isis bfd** command.

7.6.18 display isis bfd session

Function

The **display isis bfd session** command displays information about dynamic BFD sessions.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display isis [*process-id* | **vpn-instance** *vpn-instance-name*] **bfd session** { **all** | **peer** *ip-address* | **interface** *interface-type interface-number* }

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.
all	Specifies all IS-IS interfaces in the IS-IS process.	-
peer *ip-address*	Specifies the IP address of a neighbor.	The value is in dotted decimal notation.
interface *interface-type interface-number*	Specifies the interface on which BFD session statistics need to be collected.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

BFD can provide millisecond-level fault detection. It can work with IS-IS to fast detect faults on neighboring devices and instruct IS-IS to recalculate routes for correct packet forwarding. You can run the **display isis bfd session** command to check information about dynamic BFD sessions.

Example

Display information about dynamic BFD sessions.

```
<HUAWEI> display isis bfd session all
                BFD session information for ISIS(1)
                -----------------------------------
Peer System ID : 0000.0000.0002          Interface : Vlanif10
TX : 1000          BFD State : up        Peer IP Address : 10.1.1.2
RX : 1000          LocDis : 8194         Local IP Address: 10.1.1.1
Multiplier : 3     RemDis : 8197          Type : L1
Diag : No diagnostic information
Total BFD session(s) : 1
```

Table 7-92 Description of the display isis bfd session all command output

Item	Description
Peer System ID	System ID of the neighbor.

Item	Description
Interface	Local IS-IS interface connected to the neighbor.
TX	Negotiated minimum interval for transmitting BFD packets.
BFD State	BFD session status: Up or Down.
Peer IP Address	IP address of the remote IS-IS interface.
RX	Negotiated minimum interval for receiving BFD packets.
LocDis	Local discriminator dynamically assigned by BFD.
Local IP Address	IP address of the local IS-IS interface.
Multiplier	Remote detection multiplier.
RemDis	Remote identifier dynamically assigned by BFD.
Type	Level of the neighbor: • L1: Level-1 • L2: Level-2 • L1L2: Level-1-2
Diag	Diagnostic information. "No diagnostic information" indicates that BFD runs properly and no diagnostic information is displayed.
Total BFD session (s)	Total number of BFD sessions.

7.6.19 display isis brief

Function

The **display isis brief** command displays brief information about IS-IS.

Format

display isis brief [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis [*process-id*] **brief**

Parameters

Parameter	Description	Value
process-id	Display brief information about a specified IS-IS process.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Display brief information about the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis brief** command to check brief information about IS-IS.

Example

Display brief information about IS-IS.

```
<HUAWEI> display isis brief

                ISIS Protocol Information for ISIS(1)
                -------------------------------------
SystemId: 0000.0000.0001      System Level: L1
Area-Authentication-mode: NULL
Domain-Authentication-mode: NULL
Ipv6 is not enabled
ISIS is in starting status
ISIS is in protocol hot standby state: Real-Time Backup
Flush Protect Mode: False

Interface: (Vlanif10)
Cost: L1 10        L2 10               Ipv6 Cost: L1 10   L2 10
State: IPV4 Down                       IPV6 Down
Type: BROADCAST                          MTU: 1497
Priority: L1 64   L2 64
Timers:     Csnp: L1 10    L2 10    ,Retransmit: L12 5   , Hello: L1 10 L2 10  ,

Hello Multiplier: L1 3    L2 3     , LSP-Throttle Timer: L12 50
```

Table 7-93 Description of the display isis brief command output

Item	Description
ISIS Protocol Information for ISIS(1)	Information about IS-IS process 1. To create an IS-IS process, run the **isis** command.

Item	Description
SystemId	System ID of the current IS-IS device. To set the system ID, run the **network-entity** command.
System Level	Level of an IS-IS device: ● L1: Level-1 ● L2: Level-2 ● L12: Level-1-2 To set the level of an IS-IS device, run the **is-level** command.
Area-Authentication-mode	IS-IS area authentication mode. To set the IS-IS area authentication mode and password, run the **area-authentication-mode** command.
Domain-Authentication-mode	IS-IS routing domain authentication mode. To set the IS-IS domain authentication mode and password, run the **domain-authentication-mode** command.
Ipv6 is not enabled	IPv6 is not enabled in the IS-IS process.
ISIS is in starting status	IS-IS is running.
ISIS is in protocol hot standby state: Real-Time Backup	IS-IS hot standby is real-time backup.
Flush Protect Mode	Indicates whether the process works in purge LSP protection mode.
Interface	Interface on which IS-IS is enabled. To enable IS-IS on an interface, run the **isis enable** command.
Cost	Cost of IS-IS IPv4 interfaces. To set the cost of IS-IS IPv4 interfaces, run the **circuit-cost** or **isis cost** command.
Ipv6 Cost	Cost of IS-IS IPv6 interfaces. To set the cost of IS-IS IPv6 interfaces, run the **ipv6 circuit-cost** or **isis ipv6 cost** command.
State	Status of the IS-IS interface.
Type	Network type of the IS-IS interface: P2P or broadcast. To set the network type of an IS-IS interface, run the **isis circuit-type** command.
MTU	MTU of the IS-IS interface.

Item	Description
Priority	Priority of the IS-IS interface. **NOTE** When the current interface is a virtual link interface, the priority field is not displayed.
Timers	IS-IS timer.
Csnp	Interval for sending CSNPs. To set the interval for sending CSNPs, run the **isis timer csnp** command.
Retransmit	Interval for retransmitting LSPs on a P2P link. To set the retransmit interval, run the **isis timer lsp-retransmit** command.
Hello	Interval for sending Hello packets. To set the interval for sending Hello packets, run the **isis timer hello** command.
Hello Multiplier	Number of consecutive Hello packets that are not received before the neighbor is declared Down. To set this value, run the **isis timer holding-multiplier** command.
LSP-Throttle Timer	Interval for sending LSPs or CSNPs and number of LSPs or CSNPs that are sent each time. To set the two values, run the **isis timer lsp-throttle** command.

7.6.20 display isis debug-switches

Function

The **display isis debug-switches** command displays the current IS-IS debugging status.

Format

display isis debug-switches [*process-id*]

display isis [*process-id*] **debug-switches**

Parameters

Parameter	Description	Value
process-id	Displays the current status of IS-IS debugging of specified IS-IS process ID. If *process-id* is not specified, display the current status of IS-IS debugging of all IS-IS process.	The value is an integer that ranges from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If no IS-IS process is enabled or IS-IS debugging is not enabled, the **display isis debug-switches** command output is empty.

Example

Display the IS-IS debugging status.

```
<HUAWEI> display isis debug-switches 1
ISIS-1 SPF events related debugging switch is on
```

7.6.21 display isis error

Function

The **display isis error** command displays statistics about error LSPs and Hello packets that are received by IS-IS interfaces or processes.

Format

display isis error [{ *process-id* | **vpn-instance** *vpn-instance-name* } [**interface**]]

display isis error interface *interface-type interface-number*

Parameters

Parameter	Description	Value
process-id	Displays statistics about error packets received by a specified IS-IS process.	The value is an integer that ranges from 1 to 65535.

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays statistics about error packets received by the IS-IS process in a specified VPN instance.	The value must be an existing VPN instance name.
interface	Displays statistics about error packets received on all interfaces in a specified IS-IS process or VPN instance.	-
interface-type interface-number	Displays statistics about error packets received by a specified interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis error** command to check statistics about error LSPs and Hello packets that are received by IS-IS interfaces or processes.

Example

Display statistics about error LSPs and Hello packets received on all interfaces in IS-IS process 1.

```
<HUAWEI> display isis error 1
                Statistics of error packets for ISIS(1)
                ---------------------------------------
LSP packet errors:
Longer LSP             : 0          Smaller LSP            : 0
Mismatched Level       : 0          Invalid Sysid          : 0
Zero Sequence Number   : 0          Illegal IS Type        : 0
Zero Checksum          : 0          Incorrect Checksum     : 0
Bad Authentication     : 0          Bad Auth Count         : 0
More Protocol TLV      : 0          Bad Nbr TLV            : 0
Bad Extended IS TLV    : 0          Bad IF Addr TLV        : 0
Bad Reach TLV          : 0          Bad Inter Domain TLV   : 0
Mismatched Area Id(L1) : 0          Bad TLV Length         : 0
Bad Alias TLV          : 0          Bad Area TLV           : 0
Bad SRLG TLV           : 0          Unknown Adjacency      : 0
Bad Protocol ID        : 0          Bad Version            : 0
Zero Lifetime          : 0          Bad Ext Reach TLV      : 0
Bad TE Router ID TLV   : 0          Bad TE Sub TLV         : 0

Hello packet errors:
Bad Packet Length      : 0          Reserved CircType      : 0
Repeated System ID     : 0          Bad Circuit Type       : 0
Longer packet          : 0          More Area Addr         : 0
Longer Area Addr       : 0          Bad Area Addr TLV      : 0
More IF Addr           : 0          Bad Formatted IF TLV   : 0
More Nbr SNPA(LAN)     : 0          Invalid Sysid          : 0
Bad TLV Length         : 0          Zero HoldingTime       : 0
```



```
Unusable IP Addr        : 0          Repeated IPv4 Addr      : 0
Mismatched Area Addr(L1): 0          Mismatched Proto        : 0
SNPA Conflicted(LAN)    : 0          Mismatched Level        : 0
Mismatched Max Area Addr: 0          Bad Authentication      : 0
More Auth TLV           : 0          3-Way Option Error(P2P) : 0
No Area Addr TLV        : 0          Bad Protocol ID         : 0
Bad Version             : 0          Invalid IPv6 Addr       : 0
More IPv6 IF Addr       : 0          Duplicate IPv6 Addr     : 0
More Optional Checksum  : 0          Bad Optional Checksum   : 0
------------------------------------------------------------------------------
```

Table 7-94 Description of the display isis error command output

Item	Description
LSP packet errors	LSP errors.
Longer LSP	The LSP length is greater than the value set using the **lsp-length receive** command.
Smaller LSP	The LSP header length is smaller than the fixed length.
Mismatched Level	The level of received LSPs mismatches the local IS-IS level.
Invalid Sysid	The system ID in an LSP is invalid.
Zero Sequence Number	The sequence number of an LSP is 0.
Illegal IS Type	The IS type is invalid.
Zero Checksum	The checksum of an LSP is 0.
Incorrect Checksum	The checksum of an LSP is incorrect.
Bad Authentication	The authentication field of an LSP is incorrect.
Bad Auth Count	The number of authentication fields carried in an LSP is incorrect, that is, it is greater than 1.
More Protocol TLV	The number of protocol TLVs in an LSP is greater than 1.
Bad Nbr TLV	The neighbor TLV is incorrect.
Bad Extended IS TLV	The extended IS TLV is incorrect.
Bad IF Addr TLV	The interface address TLV is incorrect.
Bad Reach TLV	The reachability TLV is incorrect.
Bad Inter Domain TLV	The inter-domain TLV is incorrect. The correct TLV is 0x83.
Mismatched Area Id(L1)	Mismatched Level-1 area ID.
Bad TLV Length	The TLV length is incorrect.
Bad Alias TLV	The Alias TLV is incorrect.

Item	Description
Bad Area TLV	The area TLV is incorrect.
Bad SRLG TLV	The SRLG TLV is incorrect.
Unknown Adjacency	LSPs are received from unknown adjacency.
Bad Protocol ID	The protocol ID is incorrect.
Bad Version	The version is incorrect.
Zero Lifetime	The remaining lifetime of an LSP is 0.
Bad Ext Reach TLV	The Ext Reach TLV is incorrect.
Bad TE Router ID TLV	The TE Router ID TLV is incorrect.
Bad TE Sub TLV	The TE sub-TLV is incorrect.
Hello packet errors	Hello packet errors.
Bad Packet Length	The Hello packet length is incorrect.
Reserved CircType	The reserved link type is incorrect.
Repeated System ID	The system ID is repeated.
Bad Circuit Type	The link type is incorrect.
Longer packet	The Hello packet length is greater than the larger value between the interface MTU and the value set using the **lsp-length originate** command.
More Area Addr	Area addresses are superfluous.
Longer Area Addr	The area address is too long.
Bad Area Addr TLV	The area address TLV is incorrect.
More IF Addr	Interface addresses are superfluous.
Bad Formatted IF TLV	The format of the interface TLV is incorrect.
More Nbr SNPA(LAN)	Sub-network Points of Attachment (SNPAs) of a neighbor on a broadcast network are superfluous.
Invalid Sysid	The system ID length field is not 0 or 6.
Bad TLV Length	The TLV length is incorrect.
Zero HoldingTime	The neighbor holdtime is 0.
Unusable IP Addr	The IP address is on a different network segment than the peer end.
Repeated IPv4 Addr	The IPv4 address is repeated.
Mismatched Area Addr(L1)	Mismatched Level-1 area address.
Mismatched Proto	Mismatched protocol.

Item	Description
SNPA Conflicted(LAN)	Conflicting SNPA on a broadcast network.
Mismatched Level	The level of received Hello packets mismatches the local IS-IS level.
Mismatched Max Area Addr	The maximum area address is incorrect.
Bad Authentication	The authentication field of a Hello packet is incorrect.
More Auth TLV	Authentication TLVs are superfluous.
3-Way Option Error(P2P)	3-way information is incorrect.
No Area Addr TLV	The received Hello packet has no area address TLV.
Bad Protocol ID	The protocol ID is incorrect.
Bad Version	The version is incorrect.
Invalid IPv6 Addr	An IPv6 address is invalid.
More IPv6 IF Addr	A Hello packet carries more than 11 IPv6 addresses.
Duplicate IPv6 Addr	The IPv6 address is repeated.
More Optional Checksum	More than one optional checksum TLV is contained in a packet.
Bad Optional Checksum	Error optional checksum TLV.

7.6.22 display isis graceful-restart status

Function

The **display isis graceful-restart status** command displays the IS-IS GR status.

Format

display isis *process-id* **graceful-restart status** [**level-1** | **level-2**]

display isis graceful-restart status [**level-1** | **level-2**] [*process-id* | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
level-1	Displays the Level-1 IS-IS GR status.	-
level-2	Displays the Level-2 IS-IS GR status.	-

Parameter	Description	Value
process-id	Displays the IS-IS GR status of a specified IS-IS process.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Displays the IS-IS GR status of the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis graceful-restart status** command to check the IS-IS GR status.

Example

Display the IS-IS GR status.

```
<HUAWEI> display isis graceful-restart status
                  Restart information for ISIS(1)
                    ------------------------------
IS-IS(1) Level-1 Restart Status
Restart Interval: 300
SA Bit Supported
  Total Number of Interfaces = 1
  Restart Status: RESTART COMPLETE
IS-IS(1) Level-2 Restart Status
Restart Interval: 300
SA Bit Supported
  Total Number of Interfaces = 1
  Restart Status: RESTART COMPLETE
```

Table 7-95 Description of the display isis graceful-restart status command output

Item	Description
Restart Interval	Expected restart time of the device. This parameter is configured using the **graceful-restart interval** command.
SA Bit Supported	Whether the device supports SA.
Total Number of Interfaces = 1	Number of interfaces on which IS-IS is enabled. IS-IS is enabled using the **isis enable** command on an interface.

Item	Description
Restart Status:	Restart status of the current device. RESTART COMPLETE indicates that the restart is complete.

7.6.23 display isis interface

Function

The **display isis interface** command displays information about IS-IS interfaces.

Format

display isis interface [[**verbose** | **traffic-eng**] * | **te-tunnel**] [*process-id* | **vpn-instance** *vpn-instance-name*] (supported only by the S5731-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display isis *process-id* **interface** [**te-tunnel** | [**traffic-eng** | **verbose**] *] (supported only by the S5731-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display isis interface *interface-type interface-number* [**verbose**]

display isis interface *interface-type interface-number* [**verbose**]

display isis interface

display isis interface *process-id*

display isis *process-id* **interface**

display isis interface vpn-instance *vpn-instance-name*

Parameters

Parameter	Description	Value
verbose	Displays detailed information about IS-IS interfaces.	-
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
traffic-eng	Displays Traffic Engineering information of IS-IS.	-

Parameter	Description	Value
te-tunnel	Displays information on the MPLS TE tunnel of IS-IS.	-
vpn-instance *vpn-instance-name*	Display IS-IS interface information in the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.
interface-type interface-number	Specifies the type and number of an interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use this command to view IS-IS interface information, including the interface name, IP address, link status of the interface. If **verbose** is specified, the command output also includes the configuration of IS-IS parameters for the interface, such as the CSNP broadcast interval, Hello packet broadcast interval, and number of IS-IS Hello packets that fail to be received before IS-IS declares the neighbor Down.

Example

Display information about IS-IS interfaces.

```
<HUAWEI> display isis interface

                  Interface information for ISIS(1)
                  ----------------------------------
Interface      Id     IPV4.State          IPV6.State      MTU  Type  DIS
Vlanif10       001        Up                Down        1497 L1/L2 No/No
```

Display detailed information about IS-IS interfaces.

```
<HUAWEI> display isis interface verbose
                  Interface information for ISIS(1)
                  ----------------------------------
Interface      Id     IPV4.State          IPV6.State      MTU  Type  DIS
Vlanif10       001        Up                Down        1497 L1/L2 No/No
 Circuit MT State           : Standard
 Description                :
 SNPA Address               : 0000-1382-4569
 IP Address                 : 10.1.1.5
 IPV6 Link Local Address    :
 IPV6 Global Address (es)   :
 Csnp Timer Value           : L1    10  L2    10
 Hello Timer Value          : L1    10  L2    10
 DIS Hello Timer Value      : L1     3  L2     3
 Hello Multiplier Value     : L1     3  L2     3
```



```
LSP-Throttle Timer          : L12    50
Cost                    : L1    10  L2    10
Ipv6 Cost               : L1    10  L2    10
Priority                : L1    64  L2    64
Retransmit Timer Value      : L12    5
Bandwidth-Value             : Low 1000000000  High          0
Static Bfd              : NO
Dynamic Bfd             : NO
Dynamic IPv6 Bfd            : NO
Fast-Sense Rpr          : NO
Suppress flapping peer      : enable (flapping-count: 0, threshold: 10)
```

Table 7-96 Description of the display isis interface command output

Item	Description
Interface	IS-IS interface type and number.
Id	Link ID.
IPV4.State	IPv4 link status.
MTU	Interface MTU. An IS-IS neighbor relationship can be established only when two interfaces on both ends of a link have the same MTU.
Type	Interface type: • L1: Level-1 interface • L2: Level-2 interface • L1/L2: Level-1-2 interface
DIS	Whether the interface is a DIS. **NOTE** Only an interface with network type broadcast can be selected as the DIS. If an interface is a P2P interface, "--" is displayed in this field.

Item	Description
Circuit MT State	Topology status on the interface: • Standard: – IPv4 is enabled on the interface, but IPv6 is not. – IPv6 is enabled on the interface, but IPv4 is not. In addition, the topology type is standard. – Both IPv4 and IPv6 are enabled on the interface. In addition, the topology type is standard. • IPv6: IPv6 is enabled on the interface, but IPv4 is not. In addition, the topology type is IPv6. • Standard IPv6: – IPv6 is enabled on the interface, but IPv4 is not. In addition, the topology type is compatible. – Both IPv4 and IPv6 are enabled on the interface. In addition, the topology type is compatible or IPv6. To configure the IPv6 topology type, run the **ipv6 enable** command in the IS-IS view.
SNPA Address	MAC address.
IP Address	IPv4 address of the interface.
IPV6 Link Local Address	IPv6 link-local address.
IPV6 Global Address(es)	IPv6 global address.
Csnp Timer Value	Interval for sending CSNPs. To set the interval for sending CSNPs, run the **isis timer csnp** command.
Hello Timer Value	Interval for sending Hello packets. To set the interval for sending Hello packets, run the **isis timer hello** command.
DIS Hello Timer Value	Interval at which the DIS sends Hello packets, which is one third of **Hello Timer Value**. The field takes effect only when the interface functions as the DIS.
Hello Multiplier Value	Number of consecutive Hello packets that are not received before the neighbor is declared Down. To set this value, run the **isis timer holding-multiplier** command.

Item	Description
LSP-Throttle Timer	Interval for sending LSPs or CSNPs and number of LSPs or CSNPs that are sent each time. To set the two values, run the **isis timer lsp-throttle** command.
Cost	Cost of an IPv4 interface. This field value affects route selection.
Ipv6 Cost	Cost of an IPv6 interface. This field value affects route selection in IPv6 topologies.
Priority	Priority for DIS election. To set the priority for DIS election, run the **isis dis-priority** command.
Retransmit Timer Value	Interval for retransmitting LSPs on a P2P link. To set the retransmit interval, run the **isis timer lsp-retransmit** command.
Bandwidth-Value	Physical bandwidth of an interface. The value can be calculated using the following formula: **Bandwidth-Value** = 4294967296 x **high** + **low**. For example, if the value of **high** is 1 and the value of **low** is 705032704, Bandwidth-Value = 1 x 4294967296 + 705032704 = 5000000000. ● low: low bandwidth with the maximum value of 4294967295 ● high: high bandwidth with the maximum value of 4294967295
Static Bfd	Static BFD status: ● NO: Static BFD is not enabled. ● YES: Static BFD is enabled. To enable BFD on an interface, run the **isis bfd static** command.
Dynamic Bfd	Dynamic BFD status: ● NO: Dynamic BFD is not enabled. ● YES: Dynamic BFD is enabled. To enable dynamic BFD on an interface, run the **bfd all-interfaces enable** or **isis bfd enable** command.

Item	Description
Dynamic IPv6 Bfd	Status of dynamic IPv6 BFD: • NO: Dynamic BFD is not enabled on the interface. • YES: Dynamic BFD is enabled on the interface. To enable dynamic IPv6 BFD on an interface, run the **ipv6 bfd all-interfaces enable** or **isis ipv6 bfd enable** command.
Fast-Sense Rpr	Whether the fast sense RPR function is enabled.
Suppress flapping peer	Status of IS-IS neighbor relationship flapping suppression: • NO: IS-IS neighbor relationship flapping suppression is disabled. • YES (flapping-count: 3, threshold: 10): IS-IS neighbor relationship flapping suppression is enabled and IS-IS neighbor relationship is not suppressed. The flapping-count and threshold value. • Hold-down(start: UTC 10:20:30, remain-interval: 7): IS-IS neighbor relationship flapping suppression works in Hold-down mode. The starting and remaining time of the flapping suppression. • Hold-max-cost(start: UTC 10:20:30, remain-interval: 7): IS-IS neighbor relationship flapping suppression works in Hold-max-cost mode. The starting and remaining time of the flapping suppression. **NOTE** In the flapping suppression state, *remain-interval* is reset each time when the device detects a valid neighbor flapping event.

7.6.24 display isis last-peer-change

Function

The **display isis last-peer-change** command displays changes in IS-IS neighbor relationships.

Format

display isis *process-id* **last-peer-change**

display isis last-peer-change [*process-id* | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use the **display isis last-peer-change** command to view a maximum of 100 latest records in each IS-IS process.

Example

Display changes in IS-IS neighbor relationships.

```
<HUAWEI> display isis last-peer-change
                  Peer change information for ISIS(1)
                  -----------------------------------

Time      : 2011-03-28 09:30:34
System Id : 0000.0000.0002
Type      : L1 LAN
Interface : vlanif10
State     : IPv4 -> IPv4/IPv6
Details   : Protocol change
```

Table 7-97 Description of the display isis last-peer-change command output

Item	Description
Time	Time a neighbor relationship changes.
System Id	System ID of a neighbor.
Type	Neighbor type: • L1 LAN • L2 LAN • P2P
Interface	Interface connected to the neighbor.

Item	Description
State	Changes in the IS-IS neighbor relationship status: ● Init -> Up ● Up -> Down ● Up -> Init ● Down -> Up ● IPv4/IPv6 -> IPv4 ● IPv4/IPv6 -> IPv6 ● IPv4 -> IPv4/IPv6 ● IPv6 -> IPv4/IPv6 ● IPv4 -> IPv6 ● IPv6 -> IPv4 ● L12 -> L2 ● L2 -> L12
Details	Causes for IS-IS neighbor relationship changes: ● Circuit down ● BFD down ● Hold timer expired ● New adjacency created ● Clear neighbor ● Multiple P2P adjacency ● Adjacency usage mismatch ● Internal error ● Memory low ● P2P circuit ID conflict ● Area mismatch ● Three way down ● Three way init ● Three way up ● Protocol change ● Mt usage mismatch ● Peer state change ● Peer level change

7.6.25 display isis lsdb

Function

The **display isis lsdb** command displays information about IS-IS LSDBs.

Format

display isis *process-id* **lsdb** [{ **level-1** | **level-2** } | **verbose** | { **local** | *lsp-id* | **is-name** *symbolic-name* }] *

display isis lsdb [{ **level-1** | **level-2** } | **verbose** | { **local** | *lsp-id* | **is-name** *symbolic-name* }] * [*process-id* | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
level-1	Displays Level-1 LSDBs. If no level is specified, Level-1 and Level-2 LSDBs are displayed.	-
level-2	Displays Level-2 LSDBs. If no level is specified, Level-1 and Level-2 LSDBs are displayed.	-
verbose	Displays detailed information about IS-IS LSDBs.	-
local	Displays information about the local LSDB.	-
lsp-id	Specifies an LSP ID.	The value is in dotted decimal notation. The value ranges from 16 to 20 in ####.####.####.##-## format, such as 0050.0500.5004.00-00.
is-name *symbolic-name*	Specifies the dynamic host name in is-name or is-name.##-## format. ##-## indicates the pseudonode ID-LSP fragment ID.	The value is a string of 1 to 70 characters.

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays the LSDB of the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis lsdb** command to check information about IS-IS LSDBs.

Example

Display IS-IS LSDB information.

```
<HUAWEI> display isis lsdb
                  Database information for ISIS(1)
                  --------------------------------
                    Level-2 Link State Database
LSPID              Seq Num      Checksum      Holdtime      Length  ATT/P/OL
-------------------------------------------------------------------------------
0000.0000.0001.00-00  0x0000017a   0xa21c        846           84      0/0/0
2222.2222.2222.00-00* 0x000001ce   0xbdcc        845           111     0/0/0
3333.3333.3333.00-00  0x00000013   0x8847        1004          84      0/0/0
3333.3333.3333.01-00  0x0000000b   0x95bc        1004          55      0/0/0
Total LSP(s): 4
   *(In TLV)-Leaking Route, *(By LSPID)-Self LSP, +-Self LSP(Extended),
         ATT-Attached, P-Partition, OL-Overload
```

Table 7-98 Description of the display isis lsdb command output

Item	Description
LSPID	LSP ID.
Seq Num	LSP sequence number.
Checksum	LSP checksum.
Holdtime	LSP holdtime.
Length	LSP length.
ATT/P/OL	• ATT: Attach bit • P: partition bit • OL: overload bit
Total LSP(s)	Number of LSPs.

Item	Description
*/+	• *(In TLV): Leaking route. • *(By LSPID): Locally generated LSP. • +: Locally generated extended LSP.

Display detailed information about IS-IS LSDB.

```
<HUAWEI> display isis lsdb verbose

                  Database information for ISIS(1)
                  --------------------------------

                   Level-1 Link State Database

LSPID              Seq Num      Checksum      Holdtime      Length  ATT/P/OL
-------------------------------------------------------------------------------
0000.0000.0001.00-00  0x00000010   0xc776        0 (724)       36      0/0/0
 SOURCE       0000.0000.0001.00
 NLPID        IPV4
 AREA ADDR    10
 INTF ADDR    10.22.21.1
 NBR  ID      0000.0000.0002.02  COST: 10
 IP-Internal  10.22.21.0      255.255.255.0    COST: 10

0000.0000.0003.00-00* 0x00000026   0x6fb4        1145          350     0/0/0
 SOURCE       0000.0000.0003.00
 HOST NAME    RouterA
 NLPID        IPV4
 NLPID        IPV6
 AREA ADDR    10
 INTF ADDR    10.1.1.2
 INTF ADDR    10.1.2.2
 INTF ADDR V6 FC00:1::1
 Topology     Standard
 NBR  ID      0000.0000.0001.00  COST: 10
+NBR  ID      0000.0000.0001.00  COST: 10
 IP-Internal  10.1.1.0       255.255.255.0    COST: 10
 IP-Internal  10.1.2.0       255.255.255.0    COST: 10
+IP-Extended  10.1.1.0       255.255.255.0    COST: 10 +IP-Extended  10.1.2.0       255.255.255.0    COST: 10
 IPV6         FC00:1::/64                     COST: 10
Total LSP(s): 2
   *(In TLV)-Leaking Route, *(By LSPID)-Self LSP, +-Self LSP(Extended),
         ATT-Attached, P-Partition, OL-Overload


                   Level-2 Link State Database

LSPID              Seq Num      Checksum      Holdtime      Length  ATT/P/OL
-------------------------------------------------------------------------------
0000.0000.0003.00-00* 0x0000001c   0x9b0d        1131          233     0/0/0
 SOURCE       RouterA.00 ORG ID: RouterA.00 ORG ID: 0000.0000.0003.00
 HOST NAME    RouterA
 NLPID        IPV4
 NLPID        IPV6
 AREA ADDR    10
 INTF ADDR    10.1.1.2
 INTF ADDR    10.1.2.2
 INTF ADDR V6 FC00:1::1
 Topology     Standard
 NBR  ID      0000.0000.0004.00  COST: 10
+NBR  ID      0000.0000.0004.00  COST: 10 +MT NBR ID    0000.0000.0004.00  COST: 10
 IP-Internal  10.1.1.0       255.255.255.0    COST: 10
 IP-Internal  10.1.2.0       255.255.255.0    COST: 10
+IP-Extended  10.1.1.0       255.255.255.0    COST: 10 +IP-Extended  10.1.2.0       255.255.255.0    COST: 10
```



```
 IPV6         FC00:1::/64                      COST: 10

0000.0000.0003.00-01* 0x00000005   0x5ec1        1129          70      0/0/0
 SOURCE        RouterA.00
Auth: *****      Len: --      Type: MD5
 IP-External  10.1.3.0        255.255.255.0     COST: 0
+IP-Extended  10.1.3.0        255.255.255.0     COST: 0 Total LSP(s): 2
   *(In TLV)-Leaking Route, *(By LSPID)-Self LSP, +-Self LSP(Extended),
         ATT-Attached, P-Partition, OL-Overload
```

Table 7-99 Description of the display isis lsdb verbose command output

Item	Description
SOURCE	System ID of the source node.
ORG ID	Source system ID of the virtual system.
HOST NAME	Dynamic host name.
Auth	Authentication password.
Len	Authentication password length.
Type	Authentication type: • Plain Text: indicates that the password is sent in plain text. • MD5: indicates that the password is encrypted using the MD5 algorithm. • CRYPTO_AUTH: The password is encrypted with the HMAC-SHA256 encryption algorithm.
NLPID	Supported network protocol. • IPV4: supported the IPv4 network protocol. • IPV6: supported the IPv6 network protocol.
AREA ADDR	Area address.
INTF ADDR	Interface IPv4 address.
INTF ADDR V6	Interface IPv6 address.
Topology	Topology type.
NBR ID	System ID of a neighbor.
+NBR ID	System ID of a neighbor, which can carry TE information.
+MT NBR ID	Neighbor ID in the MT topology.
COST	Route cost.
IP-Internal	Internal IPv4 routing information.
IP-External	External IPv4 routing information.

Item	Description
+IP-Extended	Extended IP routing information, which can carry TE information.
IPV6	Internal IPv6 routing information.

7.6.26 display isis mesh-group

Function

The **display isis mesh-group** command displays the configuration of the IS-IS mesh-group.

Format

display isis mesh-group [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **mesh-group**

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis mesh-group** command to check the configuration of the IS-IS mesh-group.

Example

Display information about the configured IS-IS mesh-group.

```
<HUAWEI> display isis mesh-group
          Mesh Group information for ISIS(1)
--------------------------------------------------------------------------
Interface         Status
```



```
Vlanif10        100
Vlanif20        100
```

Table 7-100 Description of the display isis mesh group command output

Item	Description
Interface	Type and number of the interface in the mesh-group.
Status	Mesh-group number.

7.6.27 display isis name-table

Function

The **display isis name-table** command displays the mapping between the dynamic hostnames of IS-IS routers and the system IDs.

Format

display isis name-table [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **name-table**

Parameters

Parameter	Description	Value
process-id	Displays the mapping between dynamic hostnames of IS-IS routers and system IDs in a specified IS-IS process.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Displays the mapping between dynamic hostnames of IS-IS routers and system IDs in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After dynamic hostnames are configured, you can run the **display isis name-table** command to check the mapping between the dynamic hostnames and system IDs.

Example

Configure a name for the local IS-IS system.

```
[Switch] isis
[Switch-isis-1] is-name RUTA
```

Configure a static name mapping for the remote IS-IS system (0000.0000.0041).

```
[Switch-isis-1] is-name map 0000.0000.0041 RUTB
```

Display the IS-IS system name table.

```
<Switch> display isis name-table
          Name table information for ISIS(1)
System ID            Hostname      Type
------------------------------------------------------------
6789.0000.0001       RUTA          DYNAMIC
0000.0000.0041       RUTB          STATIC
```

Table 7-101 Description of the display isis name-table command output

Item	Description
System ID	System ID of the current IS-IS device. To set the system ID, run the **network-entity** command.
Hostname	Host name of the device. To set a host name for the local or remote device, run the **is-name** or **is-name map** command.
Type	Type of the mapping between the host name and system ID: static or dynamic.

7.6.28 display isis peer

Function

The **display isis peer** command displays IS-IS neighbor information.

Format

display isis *process-id* **peer** [**verbose**]

display isis peer [**verbose**] [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis peer interface *interface-type interface-number* [**verbose**]

Parameters

Parameter	Description	Value
verbose	Displays detailed IS-IS neighbor information, including the area address, period during which the neighbor remains Up, and IP address of its directly connected interface.	-
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.
interface *interface-type interface-number*	Specifies the type and number of an interface.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

In an IS-IS area, if you need to check whether two devices can communicate properly, run the **display isis peer** command to view information about the neighbor, including the neighbor status, period during which the neighbor remains Up, and neighbor type.

Example

Display information about IS-IS neighbors.

```
<HUAWEI> display isis peer
                    Peer information for ISIS(1)

 System Id     Interface         Circuit Id      State HoldTime Type     PRI
------------------------------------------------------------------------------
0000.0000.0001  Vlanif10         0000.0000.0001.01  Up    24s   L1(L1L2)  64
0000.0000.0001  Vlanif10         0000.0000.0001.01  Up    24s   L2(L1L2)  64

Total Peer(s): 2
```


Table 7-102 Description of the display isis peer command output

Item	Description
System Id	System ID of the neighbor.
Interface	Type and number of the interface.
Circuit Id	Circuit ID.
State	Neighbor status: • Up: indicates that the neighbor is Up and the two ends can communicate. • Init: indicates that the local end can receive packets from the remote end but the remote end cannot receive packets from the local end. This status is displayed when authentication is configured on the remote end. • Down: indicates that the neighbor is Down. This is the initial status, indicating that no message is received from the neighbor. In most cases, this status is not displayed.
HoldTime	Neighbor holdtime.
Type	Neighbor type: • L1: indicates that the neighbor type is Level-1 and interfaces on both ends are Level-1 interfaces. • L2: indicates that the neighbor type is Level-2 and interfaces on both ends are Level-2 interfaces. • L1(L1L2): indicates that the neighbor type is Level-1 and interfaces on both ends are Level-1-2 interfaces. • L2(L1L2): indicates that the neighbor type is Level-2 and interfaces on both ends are Level-1-2 interfaces.
PRI	Priority of the neighbor in DIS election.

Display detailed information about IS-IS neighbor.

```
<HUAWEI> display isis peer verbose
                    Peer information for ISIS(1)

 System Id     Interface         Circuit Id         State  HoldTime   Type     PRI
------------------------------------------------------------------------------
0000.0000.0001  Vlanif10          0000.0000.0001.01     Up     26s     L1(L1L2)   64
  MT IDs supported     : 0(UP)    Local MT IDs        : 0
  Area Address(es)     : 10
  Peer IP Address(es)  : 10.10.10.1
```



```
  Peer IPv6 Address(es): FE80:2000:57::7
  Uptime              : 00:00:19
  Adj Protocol        : IPV4  IPv6
  Restart Capable     : YES
  Suppressed Adj      : NO
  Peer System Id      : 0000.0000.0001

0000.0000.0001  Vlanif10          0000.0000.0001.01     Up    27s    L2(L1L2)   64
  MT IDs supported    : 0(UP)    Local MT IDs        : 0
  Area Address(es)    : 10
  Peer IP Address(es) : 10.10.10.1
  Peer IPv6 Address(es): FE80:2000:56::6
  Uptime              : 00:00:19
  Adj Protocol        : IPV4  IPv6
  Restart Capable     : YES
  Suppressed Adj      : NO
  Peer System Id      : 0000.0000.0001

Total Peer(s): 2
```

Table 7-103 Description of the **display isis peer verbose** command output

Item	Description
MT IDs supported	IDs of topology instances supported by the remote interface.
Local MT IDs	IDs of topology instances supported by the local interface.
Area Address(es)	Area addresses of the neighbor.
Peer IP Address(es)	IP address of the remote interface.
Peer IPv6 Address(es)	IPv6 address of the remote interface.
Uptime	Period during which the neighbor remains Up.
Adj Protocol	Protocol used for establishing the adjacency.
Restart Capable	Whether GR is supported: • YES: indicates that GR is supported. • NO: indicates that GR is not supported.
Suppressed Adj	Whether neighbor suppression is supported: • YES: indicates that neighbor suppression is supported. • NO: indicates that neighbor suppression is not supported.
Peer System Id	The System ID of the IS-IS peer.
Total Peer(s)	Number of neighbors.

7.6.29 display isis purge packet

Function

The **display isis purge packet** command displays statistics about received IS-IS purge LSPs carrying the POI TLV.

Format

display isis *process-id* **purge packet** [*packet-number*]

display isis purge packet *process-id* [*packet-number*]

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer ranging from 1 to 65535.
packet-number	Specifies the number of purge LSPs whose statistics are to be displayed.	The value is an integer ranging from 1 to 20.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To check statistics about received IS-IS purge LSPs carrying the POI TLV, run the **display isis purge packet** command. You can specify *packet-number* in the command to check statistics about a specified number of purge LSPs, and the maximum value of *packet-number* is 20. If its value is set to 20:

- The first two displayed purge LSPs are sent by the local source system, with fragment number 0, one of which is Level-1, and the other is Level-2.
- The first 10 displayed purge LSPs are generated locally, and the next 10 purge LSPs are from neighbors. The purge LSPs are displayed in reverse order of the time when they were received. If the LSP ID is followed by an asterisk (*), the purge LSP is generated locally.

Example

Display statistics about received IS-IS purge LSPs carrying the POI TLV.

```
<HUAWEI> display isis purge packet 1 10

                Purge LSP packet for ISIS(1)
-------------------------------------------------------------
Packet information(Index 1):
```



```
-------------------------------------------------------
Received LSPID      :  0000.0000.0027.00-00*
Source Interface    :  Vlanif100
Time                :  2015-1-22 13:55:06
Level               :  Level-2
PDU Type            :  20(Level-2 LSP)
PDU Length          :  55
Sequence Number     :  0x00000015
Checksum            :  0xc891
POI NAME            :  0000.0000.0004
POI NAME(Neighbor)  :  0000.0000.0005
HOST NAME           :  RT4-Pro1
Auth Type           :  **
0010: 83 1b 01 06 12 01 00 03 00 43 00 00 00 00 00 00
0020: 00 27 00 00 00 01 6a 6c bd ed 01 0d 07 01 00 00
0030: 00 00 00 01 89 0a 52 54 31 2d 50 72 6f 2d 30 31
0040: ** ** ** ** ** ** ** ** ** ** ** ** ** ** ** **
0050: ** ** **
-------------------------------------------------------
   **-Authentication TLV, *(By LSPID)-Self LSP
```

Table 7-104 Description of the **display isis purge packet** command output

Item	Description
Received LSPID	ID of a received purge LSP
Source Interface	Source interface of a received purge LSP
Time	Time when a purge LSP was received
Level	Level of a received purge LSP
PDU Type	Type of received purge LSP
PDU Length	Length of a received purge LSP
Sequence Number	Sequence number (LSN) of a received purge LSP
Checksum	Checksum of a received purge LSP
POI NAME	POI TLV carried in a received purge LSP
POI NAME(Neighbor)	Neighbor system ID in the POI TLV carried in a received purge LSP
HOST NAME	Dynamic hostname carried in a received purge LSP
Auth Type	Authentication information carried in a received purge LSP

7.6.30 display isis route

Function

The **display isis route** command displays IS-IS routing information.

Format

display isis route [*process-id* | **vpn-instance** *vpn-instance-name*] [**ipv4**] [**verbose** | [**level-1** | **level-2**] | *ip-address* [*mask* | *mask-length*]] *

display isis route [*process-id* | **vpn-instance** *vpn-instance-name*] **ipv6** [**verbose** | [**level-1** | **level-2**] | *ipv6-address* [*prefix-length*]] *

display isis *process-id* **route** [**ipv4**] [**verbose** | [**level-1** | **level-2**] | *ip-address* [*mask* | *mask-length*]] *

display isis *process-id* **route ipv6** [**verbose** | [**level-1** | **level-2**] | *ipv6-address* [*prefix-length*]] *

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.
ipv4	Displays IPv4 routes. If this parameter is not specified, both IPv4 and IPv6 IS-IS routes are displayed.	-
verbose	Displays detailed routing information.	-
level-1	Displays Level-1 IS-IS routes.	-
level-2	Displays Level-2 IS-IS routes.	-
ip-address	Displays the routes with the specified IPv4 destination address.	The value is in dotted decimal notation.
mask	Specifies the mask of an IP address.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of an IP address.	The value is an integer that ranges from 0 or 32.
ipv6	Displays IPv6 routes. If this parameter is not specified, both IPv4 and IPv6 IS-IS routes are displayed.	-
ipv6-address	Displays the routes with the specified IPv6 destination address.	The value is a 32-digit hexadecimal number, in X:X:X:X:X:X:X:X format.
prefix-length	Specifies the prefix length of an IPv6 address.	The value is an integer that ranges from 0 to 128.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis route** command to check IS-IS routing information.

Example

Display all IS-IS routes.

```
<HUAWEI> display isis route

                  Route information for ISIS(1)
                  -----------------------------

                 ISIS(1) Level-1 Forwarding Table
                 --------------------------------

IPV4 Destination     IntCost    ExtCost ExitInterface   NextHop         Flags
-------------------------------------------------------------------------------
10.2.1.0/24          10         NULL    Vlanif10        Direct          D/-/L/-
172.18.0.0/16        20         NULL    Vlanif10        10.2.1.1        A/-/L/-

 IPV6 Dest.          ExitInterface   NextHop                    Cost       Flags
-------------------------------------------------------------------------------
 FC00:0:0:200::/64   Vlanif10        Direct                     10         D/L/-

    Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                         U-Up/Down Bit Set


                 ISIS(1) Level-2 Forwarding Table
                 --------------------------------

IPV4 Destination     IntCost    ExtCost ExitInterface   NextHop         Flags
-------------------------------------------------------------------------------
10.2.1.0/24          10         NULL    Vlanif10        Direct          D/-/L/-
172.18.0.0/16        20         NULL    Vlanif10        10.2.1.1        A/-/L/-

 IPV6 Dest.          ExitInterface   NextHop                    Cost       Flags
-------------------------------------------------------------------------------
 FC00:0:0:200::/64   Vlanif10        Direct                     10         D/L/-

    Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                         U-Up/Down Bit Set
```

Table 7-105 Description of the display isis route command output

Item	Description
IPV4 Destination	IPv4 destination address and mask.
IPV6 Dest.	IPv6 destination address and mask.

Item	Description
IntCost	IPv4 internal cost or IS-IS route cost. To change the IS-IS IPv4 route cost, run the **circuit-cost** or **isis cost** command to set the link cost of an IS-IS IPv4 interface.
ExtCost	IPv4 external cost or cost of routes imported from other routing protocols. To set the cost of imported routes, run the **import-route** command. **NOTE** The costs of external routes imported using the **import-route cost-type external** command are displayed in this field to differentiate the costs of IS-IS routes.
Cost	IPv6 route cost. To change the IS-IS IPv6 route cost, run the **ipv6 circuit-cost** or **isis ipv6 cost** command to set the link cost of an IS-IS IPv6 interface.
ExitInterface	Outbound interface of a route.
NextHop	Next hop address of a route. This field displays Direct if the destination network segment is the direct network segment.
Flags	Route flag: ● D-Direct: indicates a direct route. ● A-Added to URT: indicates that a route is added to the unicast routing table. ● L-Advertised in LSPs: indicates that a route is advertised through an LSP. ● S-IGP Shortcut: indicates that the interface on which IGP-Shortcut is enabled exists on the path to the destination. ● U-Up/Down Bit Set: indicates the Up/Down bit.

Display IS-IS IPv4 routes.

```
<HUAWEI> display isis route ipv4

                  Route information for ISIS(1)
                  -----------------------------

                 ISIS(1) Level-1 Forwarding Table
                 --------------------------------

IPV4 Destination     IntCost    ExtCost ExitInterface   NextHop         Flags
-------------------------------------------------------------------------------
10.1.1.0/24          10         NULL    Loop1           Direct          D/-/L/-
10.2.2.0/24          10         NULL    Loop2           Direct          D/-/L/-
```



```
    Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                        U-Up/Down Bit Set


                    ISIS(1) Level-2 Forwarding Table
                    --------------------------------

IPV4 Destination     IntCost    ExtCost ExitInterface   NextHop         Flags
-------------------------------------------------------------------------------
10.1.1.0/24          10         NULL    Loop1           Direct          D/-/L/-
10.2.2.0/24          10         NULL    Loop2           Direct          D/-/L/-
    Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                        U-Up/Down Bit Set
```

Display detailed IS-IS routing information.

```
<HUAWEI> display isis route verbose

                    Route information for ISIS(1)
                    -----------------------------

                    ISIS(1) Level-1 Forwarding Table
                    --------------------------------

 IPV4 Dest  : 10.0.0.0/24       Int. Cost : 10          Ext. Cost : NULL
 Admin Tag  : -                 Src Count : 2           Flags     : D/L/-
 Priority   : -
 NextHop    :                   Interface :             ExitIndex :
   Direct                         Vlanif10                0x00000000

 IPV4 Dest  : 10.1.0.0/24       Int. Cost : 10          Ext. Cost : NULL
 Admin Tag  : -                 Src Count : 2           Flags     : D/L/-
 Priority   : -
 NextHop    :                   Interface :             ExitIndex :
   Direct                         Vlanif20                0x00000000

 IPV4 Dest  : 10.2.0.0/24       Int. Cost : 20          Ext. Cost : NULL
 Admin Tag  : -                 Src Count : 2           Flags     : A/L/-
 Priority   : Low
 NextHop    :                   Interface :             ExitIndex :
   10.2.0.2                       Vlanif20                0x00000003
   10.0.0.2                       Vlanif10                0x00000005

 IPV4 Dest  : 10.4.1.1/32       Int. Cost : 10          Ext. Cost : NULL
 Admin Tag  : -                 Src Count : 1           Flags     : A/L/-
 Priority   : Medium
 NextHop    :                   Interface :             ExitIndex :
   10.0.0.2                       Vlanif10                0x00000005
(B)10.2.0.2                       Vlanif20                0x00000003

    Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                        U-Up/Down Bit Set


                    ISIS(1) Level-2 Forwarding Table
                    --------------------------------

 IPV4 Dest  : 10.0.0.0/24       Int. Cost : 10          Ext. Cost : NULL
 Admin Tag  : -                 Src Count : 3           Flags     : D/L/-
 Priority   : -
 NextHop    :                   Interface :             ExitIndex :
   Direct                         Vlanif10                0x00000000

 IPV4 Dest  : 10.1.0.0/24       Int. Cost : 10          Ext. Cost : NULL
 Admin Tag  : -                 Src Count : 3           Flags     : D/L/-
 Priority   : -
 NextHop    :                   Interface :             ExitIndex :
   Direct                         Vlanif20                0x00000000
```



```
IPV4 Dest  : 10.2.0.0/24       Int. Cost : 20          Ext. Cost : NULL
Admin Tag  : -                 Src Count : 2           Flags     : -/-/-
Priority   : Low

IPV4 Dest  : 10.3.1.1/32       Int. Cost : 10          Ext. Cost : NULL
Admin Tag  : -                 Src Count : 2           Flags     : -/-/-
Priority   : Medium

  Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                   U-Up/Down Bit Set
```

Table 7-106 Description of the display isis route verbose command output

Item	Description
IPV4 Dest	IPv4 destination address and mask.
Int.Cost	IPv4 internal cost, the cost of an IS-IS route.
Ext.Cost	IPv4 external cost, the cost of other protocol route imported by IS-IS. **NOTE** The costs of non-IS-IS routes imported by the **import-route cost-type external** command are displayed in this field so that the costs of non-IS-IS routes are kept independently from the costs of IS-IS routes.
Admin Tag	Administrative tag.
Src Count	Number of source addresses to the same destination.
Priority	Convergence priority of IS-IS routes: • Critical • High • Medium • Low To set the convergence priority of IS-IS routes, run the **prefix-priority** command.
NextHop	Next hop of a route. The next hop (B) tag of a route is a backup route tag.
Interface	Outbound interface of a route.
ExitIndex	Index of the outbound interface.
Flags	Flags of routing information

7.6.31 display isis spf-log

Function

The **display isis spf-log** command displays IS-IS SPF logs.

Format

display isis spf-log [*process-id* | **vpn-instance** *vpn-instance-name*] [[**level-1** | **level-2**] | **ipv6** | **verbose**]*

display isis *process-id* **spf-log** [[**level-1** | **level-2**] | **ipv6** | **verbose**]*

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.
ipv6	Displays SPF logs in IPv6 topologies.	-
level-1	Displays Level-1 SPF logs.	-
level-2	Displays Level-2 SPF logs.	-
verbose	Displays detailed information about SPF logs.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the SPT calculated by IS-IS is incorrect, you can run the **display isis spf-log** command to diagnose the fault. The command displays information about SPF calculation, such as the start time and duration of the SPF calculation, number of nodes, and events that trigger the SPF calculation. You can determine whether the events are the cause of the SPF calculation fault based on the start time of the SPF calculation.

Example

Display IS-IS SPF logs.


```
<HUAWEI> display isis spf-log

                SPF Log information for ISIS(10)
                --------------------------------

                   ISIS(10) Level-1 SPF Log
                   ------------------------

 StartTime  Duration  Nodes  Count  Last Trigger LSP      Trigger Event
-------------------------------------------------------------------------------
2013-09-05 10:14:48+00:00
         0        1      1      NULL                ISPF_CIRC_DR_CHANGE
2013-09-05 10:13:23+00:00
         1        1      2      NULL                FULL_SPF

                   ISIS(10) Level-2 SPF Log
                   ------------------------

 StartTime  Duration  Nodes  Count  Last Trigger LSP      Trigger Event
-------------------------------------------------------------------------------
2013-09-05 10:14:48+00:00
         1        1      1      NULL                ISPF_CIRC_DR_CHANGE
2013-09-05 10:13:23+00:00
         1        1      2      NULL                FULL_SPF
```

Table 7-107 Description of the display isis spf-log command output

Item	Description
StartTime	Start time of SPF calculation.
Duration	Duration of SPF calculation.
Nodes	Number of nodes in SPF calculation.
Count	Number of times SPF calculated is triggered.
Last Trigger LSP	LSP that triggers the last SPF calculation: • NULL indicates that the LSP is local. • LSP ID (specific ID) indicates that the LSP is not local.

Item	Description
Trigger Event	Event that triggers the last SPF calculation: • NEWAREA indicates that a new NET is configured. • TUNNEL_ADJ indicates that a tunnel neighbor relationship is established. • ADJDOWN indicates that a tunnel neighbor goes Down. • NEWLSP indicates that the router receives an LSP of a new process. • LSPCHANGE indicates that the received LSP is different from local LSPs. • RST_T2_CANCEL indicates that the GR T2 Timer is disabled. • RST_T3_EXPIRE indicates that the GR T3 Timer expires. • RESTART_COMPLETE indicates that the GR is over. • CIRC_VLINK_CHANGE indicates that the status of the VLink interface changes. • PRC_IPV4_PREFIX_ADD indicates that an IPv4 route is added. • PRC_IPV4_PREFIX_DEL indicates that an IPv4 route is deleted. • PRC_IPV4_PREFIX_MODIFY indicates that an IPv4 route is modified. • PRC_IPV4_PREFIX_MIGP_ADD indicates that an MIGP IPv4 route is added. • PRC_IPV4_PREFIX_MIGP_DEL indicates that an MIGP IPv4 route is deleted. • PRC_IPV4_PREFIX_MIGP_MODIFY indicates that an MIGP IPv4 route is modified. • PRC_IPV6_PREFIX_ADD indicates that an IPv6 route is added. • PRC_IPV6_PREFIX_DEL indicates that an IPv6 route is deleted. • PRC_IPV6_PREFIX_MODIFY indicates that an IPv6 route is modified. • ISPF_ADJ_STATE_CHANGE indicates that the IS-IS neighbor status changes. • ISPF_ADJ_USAGE_CHANGE indicates that the IS-IS neighbor level changes.

Item	Description
	• ISPF_ADJ_PROT_USAGE_CHANGE indicates that the protocol used by the IS-IS neighbor changes. • ISPF_ADJ_NEXTHOP_CHANGE indicates that the next hop of the IS-IS neighbor changes. • ISPF_CIRC_METRIC_CHANGE indicates that the cost on the interface changes. • ISPF_CIRC_DR_CHANGE indicates that the DIS changes. • ISPF_NODE_DEL indicates that the nodes in the SPT change. • ISPF_NODE_OLOAD_CHANGE indicates that the overload bit of the system changes. • ISPF_LINK_ADD indicates that a new link is added to the SPT. • ISPF_LINK_DEL indicates that a link is deleted from the SPT. • ISPF_LINK_METRIC_CHANGE indicates that the cost of a link changes. • FULL_SPF indicates that a full SPF calculation is triggered. • AREA_CFG_CHANGE indicates that the area of the process changes. • AREA_LEARNT_CHANGE indicates that the area of the IS-IS neighbor changes. • ISPF_TUNNEL_TYPE_CHANGE indicates that the tunnel type changes. • PRC_LEAVE_MAN_OVERLOAD indicates that the system exits the **set-overload** state that is set manually. • PRC_IPV4_SELFLSP_CHANGE indicates that the IPv4 route in the local LSP changes. • PRC_IPV6_SELFLSP_CHANGE indicates that the IPv6 route in the local LSP changes. • PRC_ALIAS_TLV_CHANGE indicates that the **ALIAS TLV** in the LSP changes. • LFA_LINK_ADD indicates that a link is added to the routes for FRR. • LFA_LINK_DEL indicates that a link is deleted from the routes for FRR.

Item	Description
	• LFA_LINK_CHANGE indicates that the links of the routes for FRR change. • LFA_NODE_ADD indicates that a node is added along the links of the routes for FRR. • LFA_NODE_DEL indicates that a node is deleted from the links of the routes for FRR. • LFA_NODE_CHANGE indicates that the nodes along the links of the routes for FRR change. • KEY_RESRORE indicates that the local device receives an indirect next hop key.

7.6.32 display isis spf-tree

Function

The **display isis spf-tree** command displays the topology of the SPF tree that is generated by IS-IS.

Format

display isis spf-tree [**systemid** *systemid* | **dname** *dname*] [[**level-1** | **level-2**] | **ipv6** | **verbose**] * [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **spf-tree** [**systemid** *systemid* | **dname** *dname*] [[**level-1** | **level-2**] | **ipv6** | **verbose**] *

display isis *process-id* **spf-tree statistics** [[**level-1** | **level-2**] | **ipv6**] *

display isis spf-tree statistics [[**level-1** | **level-2**] | **ipv6**] * [*process-id* | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
systemid *systemid*	Displays the SPF tree of the switch with a specified system ID or pseudonode ID.	The value is in XXXX.XXXX.XXXX[.XX] format.

Parameter	Description	Value
dname *dname*	Displays the SPF tree of the switch with a specified dynamic name.	The value is a string of 1 to 64 characters without spaces. When double quotation marks are used around the string, spaces are allowed in the string.
level-1	Displays Level-1 SPF tree.	-
level-2	Displays Level-2 SPF tree.	-
ipv6	Displays SPF tree in IPv6 topologies.	-
verbose	Displays detailed information about SPF tree.	-
process-id	Display SPF tree in a specified IS-IS process.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Displays SPF tree of the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.
statistics	Displays SPF tree statistics, including information about links and nodes on the SPF tree.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify **statistics** to view the current status of the SPF tree. The generated SPF tree can be used to calculate routes only when the SPF calculation status is Completed, that is, incremental SPF (ISPF) calculation is complete.

Example

Display information about the SPF tree.

```
<HUAWEI> display isis spf-tree
                Shortest Path Tree for ISIS(1)
                ------------------------------
```



```
      Flags:  T-System is on SPF TREE  R-System is directly reachable
              O-System is Overload     D-System or Link is to be deleted
              C-Neighbor is child      P-Neighbor is parent
              G-Cost gets greater      L-Cost gets lower
              H-Nexthop is changed     U-Protocol usage is changed
              V-Link is involved       N-Link is a new path
              S-Link is IGP Shortcut   *-Relative cost
                  ISIS(1) Level-1 Shortest Path Tree
                  ----------------------------------

 SpfNode          NodeFlags  NeighbourNode     LinkCost   LinkFlags
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
 0000.0000.0001.00 T/-/-/-   >0000.0000.0002.01       10   P/-/-/-/-/-/-
>0000.0000.0002.00 T/-/-/-    0000.0000.0003.01       10   C/-/-/-/-/-/-
                     >0000.0000.0002.01       10   C/-/-/-/-/-/-
                      0000.0000.0004.00      -10*  C/-/-/-/S/-/-
>0000.0000.0002.01 T/R/-/-    0000.0000.0001.00        0   C/-/-/-/-/-/-
                     >0000.0000.0002.00        0   P/-/-/-/-/-/-
 0000.0000.0003.00 T/-/-/-    0000.0000.0004.02       10   C/-/-/-/-/-/-
                      0000.0000.0003.01       10   P/-/-/-/-/-/-
 0000.0000.0003.01 T/R/-/-    0000.0000.0003.00        0   C/-/-/-/-/-/-
                     >0000.0000.0002.00        0   P/-/-/-/-/-/-
 0000.0000.0004.00 T/-/-/-    0000.0000.0004.02       10   P/-/-/-/-/-/-
                      0000.0000.0004.01       10   C/-/-/-/-/-/-
 0000.0000.0004.01 T/-/-/-    0000.0000.0004.00        0   P/-/-/-/-/-/-
                      0000.0000.0005.00        0   C/-/-/-/-/-/-
 0000.0000.0004.02 T/-/-/-    0000.0000.0004.00        0   C/-/-/-/-/-/-
                      0000.0000.0003.00        0   P/-/-/-/-/-/-
 0000.0000.0005.00 T/-/-/-    0000.0000.0004.01       10   P/-/-/-/-/-/-
                  ISIS(1) Level-2 Shortest Path Tree
                  ----------------------------------

 SpfNode          NodeFlags  NeighbourNode     LinkCost   LinkFlags
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
 0000.0000.0001.00 T/-/-/-   >0000.0000.0002.01       10   P/-/-/-/-/-/-
>0000.0000.0002.00 T/-/-/-    0000.0000.0003.01       10   C/-/-/-/-/-/-
                     >0000.0000.0002.01       10   C/-/-/-/-/-/-
                      0000.0000.0004.00      -10*  C/-/-/-/S/-/-
>0000.0000.0002.01 T/R/-/-    0000.0000.0001.00        0   C/-/-/-/-/-/-
                     >0000.0000.0002.00        0   P/-/-/-/-/-/-
 0000.0000.0003.00 T/-/-/-    0000.0000.0004.02       10   C/-/-/-/-/-/-
                      0000.0000.0003.01       10   P/-/-/-/-/-/-
 0000.0000.0003.01 T/R/-/-    0000.0000.0003.00        0   C/-/-/-/-/-/-
                     >0000.0000.0002.00        0   P/-/-/-/-/-/-
 0000.0000.0004.00 T/-/-/-    0000.0000.0004.02       10   P/-/-/-/-/-/-
                      0000.0000.0004.01       10   C/-/-/-/-/-/-
 0000.0000.0004.01 T/-/-/-    0000.0000.0004.00        0   P/-/-/-/-/-/-
                      0000.0000.0005.00        0   C/-/-/-/-/-/-
 0000.0000.0004.02 T/-/-/-    0000.0000.0004.00        0   C/-/-/-/-/-/-
                      0000.0000.0003.00        0   P/-/-/-/-/-/-
 0000.0000.0005.00 T/-/-/-    0000.0000.0004.01       10   P/-/-/-/-/-/-
```

Table 7-108 Description of the display isis spf-tree command output

Item	Description
SpfNode	Node ID in the network topology.

Item	Description
NodeFlags	Node flag: • T: The node is on the SPF tree. • R: The node is directly reachable. • O: The node is overloaded. • D: The node is to be deleted.
NeighbourNode	ID of the neighboring node.
LinkCost	Link cost.
LinkFlags	Link flag: • C: Neighbor is child • P: Neighbor is parent • G: Cost gets greater • L: Cost gets lower • H: Nexthop is changed • U: Protocol usage is changed • V: Link is involved • N: Link is a new path • S: Link is IGP Shortcut • *: Relative cost
>	Mark of the local node.

Display detailed information about SPF tree.

```
<HUAWEI> display isis spf-tree verbose
                    Shortest Path Tree for ISIS(1)
                    ------------------------------

                  ISIS(1) Level-1 Shortest Path Tree
                  ----------------------------------
>0000.0000.0001.00
  Distance           : 0
  Distance-URT       : 0
  Flags              : SPT
  IPv4 Nexthops-URT   : 0
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:2 Parents:0 Others:0)
          (1)   0000.0000.0002.01
                Cost  : 10
                Flags : Child

          (2)  >0000.0000.0001.01
                Cost  : 10
                Flags : Child

>0000.0000.0001.01
  Distance           : 10
  Distance-URT       : 10
  Flags              : SPT/Direct/Isolate/V6_Islt
  IPv4 Nexthops-URT   : 0
```



```
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
          (1)   0000.0000.0003.00
                Cost  : 0
                      C:0 I:vlanif 10
                Flags : Adj/Child

          (2)  >0000.0000.0001.00
                Cost  : 0
                Flags : Parent

0000.0000.0002.00
  Distance          : 10
  Distance-URT      : 10
  Flags             : SPT/V6_Islt
  IPv4 Nexthops-URT   : 1
          (1) 10.1.0.2        IF:vlanif 10  NBR:0000.0000.0002.00
          (B) 10.0.0.2        IF:vlanif 20  NBR:0000.0000.0003.00
                          TYPE:LOOP-FREE  PROTECT:LINK
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
          (1)   0000.0000.0002.02
                Cost  : 10
                Flags : Child

          (2)   0000.0000.0002.01
                Cost  : 10
                Flags : Parent

0000.0000.0002.01
  Distance          : 10
  Distance-URT      : 10
  Flags             : SPT/Direct/Isolate/V6_Islt
  IPv4 Nexthops-URT   : 0
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
          (1)   0000.0000.0002.00
                Cost  : 0
                      C:0 I:vlanif 10
                Flags : Adj/Child

          (2)  >0000.0000.0001.00
                Cost  : 0
                Flags : Parent

0000.0000.0002.02
  Distance          : 20
  Distance-URT      : 20
  Flags             : SPT/V6_Islt
  IPv4 Nexthops-URT   : 2
          (1) 10.1.0.2        IF:vlanif 10  NBR:0000.0000.0002.00
          (2) 10.0.0.2        IF:vlanif 20  NBR:0000.0000.0003.00
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:0 Parents:2 Others:0)
          (1)   0000.0000.0002.00
                Cost  : 0
                Flags : Parent

          (2)   0000.0000.0003.00
                Cost  : 0
                Flags : Parent

0000.0000.0003.00
  Distance          : 10
  Distance-URT      : 10
```



```
  Flags              : SPT/V6_Islt
  IPv4 Nexthops-URT   : 1
         (1) 10.0.0.2         IF:vlanif 10  NBR:0000.0000.0003.00
         (B) 10.1.0.2         IF:vlanif 20  NBR:0000.0000.0002.00
                           TYPE:LOOP-FREE  PROTECT:LINK
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
         (1)   0000.0000.0002.02
               Cost  : 10
               Flags : Child

         (2)  >0000.0000.0001.01
               Cost  : 10
               Flags : Parent


                ISIS(1) Level-2 Shortest Path Tree
                -----------------------------------
>0000.0000.0001.00
  Distance           : 0
  Distance-URT       : 0
  Flags              : SPT
  IPv4 Nexthops-URT   : 0
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:2 Parents:0 Others:0)
         (1)   0000.0000.0002.01
               Cost  : 10
               Flags : Child

         (2)  >0000.0000.0001.01
               Cost  : 10
               Flags : Child

>0000.0000.0001.01
  Distance           : 10
  Distance-URT       : 10
  Flags              : SPT/Direct/Isolate/V6_Islt
  IPv4 Nexthops-URT   : 0
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
         (1)   0000.0000.0003.00
               Cost  : 0
                     C:0 I:vlanif 10
               Flags : Adj/Child

         (2)  >0000.0000.0001.00
               Cost  : 0
               Flags : Parent

 0000.0000.0002.00
  Distance           : 10
  Distance-URT       : 10
  Flags              : SPT/V6_Islt
  IPv4 Nexthops-URT   : 1
         (1) 10.1.0.2         IF:Meth0/0/1  NBR:0000.0000.0002.00
         (B) 10.0.0.2         IF:vlanif 20  NBR:0000.0000.0003.00
                           TYPE:LOOP-FREE  PROTECT:LINK
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
         (1)   0000.0000.0002.02
               Cost  : 10
               Flags : Child

         (2)   0000.0000.0002.01
               Cost  : 10
```



```
                Flags : Parent

0000.0000.0002.01
  Distance           : 10
  Distance-URT       : 10
  Flags              : SPT/Direct/Isolate/V6_Islt
  IPv4 Nexthops-URT   : 0
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
          (1)   0000.0000.0002.00
                Cost  : 0
                      C:0 I:vlanif 10
                Flags : Adj/Child

          (2)  >0000.0000.0001.00
                Cost  : 0
                Flags : Parent

0000.0000.0002.02
  Distance           : 20
  Distance-URT       : 20
  Flags              : SPT/V6_Islt
  IPv4 Nexthops-URT   : 2
          (1) 10.1.0.2        IF:vlanif 10  NBR:0000.0000.0002.00
          (2) 10.0.0.2        IF:vlanif 20  NBR:0000.0000.0003.00
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:0 Parents:2 Others:0)
          (1)   0000.0000.0002.00
                Cost  : 0
                Flags : Parent

          (2)   0000.0000.0003.00
                Cost  : 0
                Flags : Parent

0000.0000.0003.00
  Distance           : 10
  Distance-URT       : 10
  Flags              : SPT/V6_Islt
  IPv4 Nexthops-URT   : 1
          (1) 10.0.0.2        IF:vlanif 10  NBR:0000.0000.0003.00
          (B) 10.1.0.2        IF:vlanif 20  NBR:0000.0000.0002.00
                          TYPE:LOOP-FREE  PROTECT:LINK
  IPv4 Nexthops-MIGP  : 0
  IPv6 Nexthops       : 0
  Neighbors: 2 (Children:1 Parents:1 Others:0)
          (1)   0000.0000.0002.02
                Cost  : 10
                Flags : Child

          (2)  >0000.0000.0001.01
                Cost  : 10
                Flags : Parent
```

Table 7-109 Description of the display isis spf-tree verbose command output

Item	Description
Distance	Cost of the shortest path from the root node to the destination node, excluding the TE tunnel link.
Distance-URT	Cost of the shortest path from the root node to the destination node, including the TE tunnel link.

Item	Description
Flags	Flag: ● SPT: The node is in the tree. ● Direct: The node is a direct node. ● Shortcut: The node is on a shortcut link. ● Oload: The node is in Overload state. ● Isolate: The node is not in the tree. ● MIGP_Islt: The node is not in the MIGP tree. ● V6_Islt: The node is not in the IPv6 tree. ● Del: The node will be deleted. ● Remote: The node is on a tunnel.
IPv4 Nexthops-URT	IPv4 next hop of the node in the unicast routing table.
IPv6 Nexthops	IPv6 next hop of the node in the unicast routing table.
Neighbors	Information about all the neighbors of this node.
Cost	Cost of the link from the root node to this node.
Flags	Relationship with the neighbor: ● Parent: The neighbor is a parent node. ● Child: The neighbor is a child node.
IF	Name of the outbound interface.
NBR	System ID of the next hop.
TYPE	Traffic protection type: ● LOOP-FREE: loop-free protection for the backup next hop
PROTECT	Traffic protection type of IS-IS Auto FRR: ● NONE: indicates no protection. ● LINK: indicates link protection. ● NODE: indicates node protection. ● LINK-NODE: indicates link and node protection.

Display the current status statistics of the SPF tree.

```
<HUAWEI> display isis spf-tree statistics
                Statistics information of SPT for ISIS(100)
               ------------------------------------------


                          Level-1 Statistics
                          ------------------
Nodes information:
       Total:                          1
       Count of nodes in SPT:          1
       Count of isolate nodes in SPT:      0
       Count of IPv6 isolate nodes in SPT:  0
```



```
      Max Distance of nodes in SPT:        0
Links information:
      Total:                       0
      Count of links from Parent to Son:   0
      Count of links from Son to Parent:   0
      Count of links that just changed:    0
Status of SPF:                     Completed
```

Table 7-110 Description of the display isis spf-tree statistics command output

Item	Description
Total (Nodes information)	Total number of nodes.
Count of nodes in SPT	Number of nodes in the SPF tree.
Count of isolate nodes in SPT	Number of isolated nodes in the network topology.
Count of IPv6 isolate nodes in SPT	Number of isolated IPv6 nodes in the network topology.
Max Distance of nodes in SPT	Maximum distance from the root node to the leaf node.
Total (Links information)	Total number of links.
Count of links from Parent to Son	Number of links from the parent node to the child node.
Count of links from Son to Parent	Number of links from the child node to the parent node.
Count of links that just changed	Number of links that just changed.
Status of SPF	Status of SPF calculation: • Completed: indicates that SPF calculation is complete. • Scheduled: indicates that SPF calculation is triggered but does not start. • Running: indicates that SPF calculation is being performed.

7.6.33 display isis statistics

Function

The **display isis statistics** command displays IS-IS process statistics.

Format

display isis statistics packet [**interface** [*interface-type interface-number*]]

display isis statistics [**updated-lsp** [**history**]] [**level-1** | **level-2** | **level-1-2**] [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **statistics** [[[**updated-lsp** [**history**]] [**level-1** | **level-2** | **level-1-2**]] | **packet**]

NOTE

Only the S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support **updated-lsp** and **history** parameters.

Parameters

Parameter	Description	Value
packet	Displays IS-IS packet statistics.	-
interface *interface-type interface-number*	Displays IS-IS packet statistics on a specified interface.	-
updated-lsp	Displays real-time data of received LSPs. By default, statistics about the LSPs received within the last 1 hour are displayed.	-
history	Displays historical data of received LSPs. By default, statistics about the LSPs received within the last 24 hours are displayed.	-
level-1	Displays IS-IS Level-1 statistics.	-
level-2	Displays IS-IS Level-2 statistics.	-
level-1-2	Displays IS-IS Level-1-2 statistics.	-
process-id	Displays statistics in a specified IS-IS process.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Displays statistics about the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use the **display isis statistics** command to view IS-IS process statistics, including the number of routes learned from other IS-IS devices, number of routes imported from other routing protocols, number of locally generated LSPs, and convergence priorities of IS-IS routes.

Example

Display IS-IS process statistics.

```
<HUAWEI> display isis statistics
                  Statistics information for ISIS(1)
                  ----------------------------------
                        Level-1 Statistics
                        ------------------

Forwarding routes information:
       Total IPV4 Learnt Routes: 0
                    Critical: 0
                    High    : 0
                    Medium  : 0
                    Low     : 0
       Total IPv4 Forwarding Routes: 1
Total IPv6 Learnt Routes: 0
                   Critical: 0
                   High    : 0
                   Medium  : 0
                   Low     : 0
       Total IPv6 Forwarding Routes: 1

Imported routes information:
       IPv4 Imported Routes:
                    Static: 0       Direct: 0
                    ISIS:   0       BGP:    0
                    RIP:    0       OSPF:   0
       IPv6 Imported Routes:
                    Static: 0       Direct: 0
                    ISIS:   0       BGP:    0
                    RIPng:  0       OSPFv3: 0
Number of advertised imported routes:
      IPv4 Imported Routes:   0
      IPv6 Imported Routes:   0

Lsp information:
              LSP Source ID:          No. of used LSPs
              0000.0000.0022                  002

                        Level-2 Statistics
                        ------------------

Forwarding routes information:
       Total IPV4 Learnt Routes: 0
                    Critical: 0
                    High    : 0
                    Medium  : 0
                    Low     : 0
       Total IPv4 Forwarding Routes: 1
Total IPv6 Learnt Routes: 0
                   Critical: 0
                   High    : 0
```



```
                  Medium  : 0
                  Low     : 0
       Total IPv6 Forwarding Routes: 1

Imported routes information:
       IPv4 Imported Routes:
                  Static: 0       Direct: 0
                  ISIS:   0       BGP:    0
                  RIP:    0       OSPF:   0
       IPv6 Imported Routes:
                  Static: 0       Direct: 0
                  ISIS:   0       BGP:    0
                  RIPng:  0       OSPFv3: 0
Number of advertised imported routes:
      IPv4 Imported Routes:   0
      IPv6 Imported Routes:   0

Lsp information:
            LSP Source ID:          No. of used LSPs
            0000.0000.0002                  001
```

Table 7-111 Description of the display isis statistics command output

Item	Description
Total IPV4 Learnt Routes • Critical • High • Medium • Low	Number of IPv4 routes learned by IS-IS: • Critical: indicates the number of IPv4 routes with convergence priority critical. • High: indicates the number of IPv4 routes with convergence priority high. • Medium: indicates the number of IPv4 routes with convergence priority medium. • Low: indicates the number of IPv4 routes with convergence priority low.
Total IPv4 Forwarding Routes	Number of IPv4 routes counted based on the destination address. If there are multiple routes to the same destination, these routes are counted as one route.
Total IPv6 Learnt Routes • Critical • High • Medium • Low	Number of IPv6 routes learned by IS-IS: • Critical: indicates the number of IPv6 routes with convergence priority critical. • High: indicates the number of IPv6 routes with convergence priority high. • Medium: indicates the number of IPv6 routes with convergence priority medium. • Low: indicates the number of IPv6 routes with convergence priority low.
Total IPv6 Forwarding Routes	Number of IPv6 routes counted based on the destination address. If there are multiple routes to the same destination, these routes are counted as one route.
IPv4 Imported Routes	Imported IPv4 routes.

Item	Description
Static	Number of imported static routes.
Direct	Number of imported direct routes.
ISIS	Number of imported IS-IS routes.
BGP	Number of imported BGP routes.
RIP	Number of imported RIP routes.
OSPF	Number of imported OSPF routes.
IPv6 Imported Routes	Imported IPv6 routes.
RIPng	Number of imported RIPng routes.
OSPFv3	Number of imported OSPFv3 routes.
LSP Source ID	System ID of the switch that generates the LSP.
No. of used LSPs	Number of used LSPs.
Number of advertised imported routes	Number of imported external routes that are advertised.

Display IS-IS packet statistics.

```
<HUAWEI> display isis statistics packet

                  PDUs information for ISIS(1)
                  ----------------------------
      Sent packets:
      PDU type              Total(packets)
      L1 IIH                2516
      L1 LSP                  84
      L1 CSNP                  8
      L1 PSNP                 44
      L2 IIH                5028
      L2 LSP                  80
      L2 CSNP                  8
      L2 PSNP                 46

      Received packets:
      PDU type              Total(packets)
      L1 IIH               12943
      L1 LSP                 216
      L1 CSNP               3911
      L1 PSNP                 40
      L2 IIH               14907
      L2 LSP                 206
      L2 CSNP               3900
      L2 PSNP                 41
```

Table 7-112 Description of the display isis statistics packet command output

Item	Description
PDU type	Packet type.

Item	Description
Total(packets)	Total number of packets.
L1 IIH	Level-1 Hello packets.
L1 LSP	Level-1 LSPs.
L1 CSNP	Level-1 CSNPs.
L1 PSNP	Level-1 PSNPs.
L2 IIH	Level-2 Hello packets.
L2 LSP	Level-2 LSPs.
L2 CSNP	Level-2 CSNPs.
L2 PSNP	Level-2 PSNPs.

7.6.34 display isis traffic-eng advertisements

Function

The **display isis traffic-eng advertisements** command displays latest advertised traffic engineering (TE) information.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

display isis traffic-eng advertisements [{ **level-1** | **level-2** | **level-1-2** } | { *lsp-id* | **local** }] * [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **traffic-eng advertisements** [{ **level-1** | **level-2** | **level-1-2** } | { *lsp-id* | **local** }] *

Parameters

Parameter	Description	Value
level-1	Displays TE information in a Level-1 LSDB.	-
level-2	Displays TE information in a Level-2 LSDB.	-

Parameter	Description	Value
level-1-2	Displays TE information in a Level-1 or Level-2 LSDB based on the local node type. That is, if the local node is a Level-1 or Level-2 node, TE information in a Level-1 or Level-2 LSDB is displayed. If the local node is a Level-1-2 node, TE information in Level-1 and Level-2 LSDBs is displayed.	-
lsp-id	Displays TE information in the specified LSP.	The value is in dotted decimal notation. The value ranges from 16 to 20 in ####.####.####.##-## format, such as 0050.0500.5004.00-00.
local	Displays locally advertised TE information.	-
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis traffic-eng advertisements** command to check latest advertised TE information.

Example

Display advertised TE information.

```
<HUAWEI> display isis traffic-eng advertisements
                    TE information for ISIS(1)
                    --------------------------
Level-1 Link State Database
----------------------------
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0000.0001.00-00* 0x00000001   0x3f57        534           0/0/0
 NLPID      : IPV4
 AREA ADDR  : 00.0005
Level-2 Link State Database
```



```
----------------------------
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0000.0001.00-00* 0x0000001c   0xf1ec        687           0/0/0
 NLPID     : IPV4
 AREA ADDR  : 00.0005
 Router ID  : 10.1.1.9
+NBR        : 0000.0000.0002.02  COST: 10
   Affinity: 0x00000000
   Interface IP Address:  10.1.1.1
   Physical BW  :      12500 Bytes/sec
   Reservable BW:       6250 Bytes/sec
   Unreserved BW:
   BW Unresrv[0] :      6250 Bytes/sec BW Unresrv[1] :      6250 Bytes/sec
   BW Unresrv[2] :      6250 Bytes/sec BW Unresrv[3] :      6250 Bytes/sec
   BW Unresrv[4] :      6250 Bytes/sec BW Unresrv[5] :      6250 Bytes/sec
   BW Unresrv[6] :      6250 Bytes/sec BW Unresrv[7] :      6250 Bytes/sec
   TE Cost     :        10
   Sub Unreserved BW for Class Type 1:
   BW Unresrv[0] :         0 Bytes/sec BW Unresrv[1] :         0 Bytes/sec
   BW Unresrv[2] :         0 Bytes/sec BW Unresrv[3] :         0 Bytes/sec
   BW Unresrv[4] :         0 Bytes/sec BW Unresrv[5] :         0 Bytes/sec
   BW Unresrv[6] :         0 Bytes/sec BW Unresrv[7] :         0 Bytes/sec
   Bandwidth Constraint Model: Russian Doll
   Bandwidth Constraints:
   BC[0] :      6250 Bytes/sec BC[1] :         0 Bytes/sec
   Local Overbooking Multiplier:
   LOM[0] :       100 %        LOM[1] :       100 %
+NBR        : 0000.0000.0004.00  COST: 10
   Affinity: 0x00000000
   Interface IP Address:  10.3.1.1
   Peer IP Address     :  10.3.1.2
   Physical BW  :      12500 Bytes/sec
   Reservable BW:       6250 Bytes/sec
   Unreserved BW:
   BW Unresrv[0] :      6250 Bytes/sec BW Unresrv[1] :      6250 Bytes/sec
   BW Unresrv[2] :      6250 Bytes/sec BW Unresrv[3] :      6250 Bytes/sec
   BW Unresrv[4] :      6250 Bytes/sec BW Unresrv[5] :      6250 Bytes/sec
   BW Unresrv[6] :      6250 Bytes/sec BW Unresrv[7] :      6250 Bytes/sec
   TE Cost     :        10
   Sub Unreserved BW for Class Type 1:
   BW Unresrv[0] :         0 Bytes/sec BW Unresrv[1] :         0 Bytes/sec
   BW Unresrv[2] :         0 Bytes/sec BW Unresrv[3] :         0 Bytes/sec
   BW Unresrv[4] :         0 Bytes/sec BW Unresrv[5] :         0 Bytes/sec
   BW Unresrv[6] :         0 Bytes/sec BW Unresrv[7] :         0 Bytes/sec
   Bandwidth Constraint Model: Russian Doll
   Bandwidth Constraints:
   BC[0] :      6250 Bytes/sec BC[1] :         0 Bytes/sec
   Local Overbooking Multiplier:
   LOM[0] :       100 %        LOM[1] :       100 %
+SRLG NBR ID: 0000.0000.0004.00
   Interface IP Address: 10.3.1.1
   Neighbor IP Address : 10.3.1.2
   Shared Risk Link Group: 10,20
```

Table 7-113 Description of the display isis traffic-eng advertisements command output

Item	Description
LSPID	LSP ID.
LSP Seq Num	LSP sequence number.
LSP Checksum	LSP checksum.
LSP Holdtime	LSP holdtime.

Item	Description
ATT/P/OL	• ATT: Attach bit • P: partition bit • OL: overload bit
NLPID	Network protocol.
AREA ADDR	Area address.
Router ID	Router ID of the switch.
+NBR	System ID of the neighbor with the cost style wide, wide-compatible, or compatible.
COST	Cost.
Affinity:	Affinity attribute of a link.
Interface IP Address	IP address of an interface.
Peer IP Address	Peer IP address of an interface.
Physical BW	Physical bandwidth of a link.
Reservable BW	Reservable bandwidth of a link.
Unreserved BW	Unreserved bandwidth.
Sub Unreserved BW for Class Type 1	Unreserved bandwidth for CT1.
BW Unresrv [x]	Unreserved bandwidth for the link with priority x.
TE Cost	TE cost of a link.
Bandwidth Constraint Model	Bandwidth constraint model used by a link.
Bandwidth Constraints	Bandwidth constraint.
BC[x]	Bandwidth constraint of BCx.
Local Overbooking Multiplier	Local overbooking multiplier.
LOM[x]	Local overbooking multiplier of BCx.
+SRLG NBR ID	System ID of a neighbor in a shared risk link group (SRLG).
Neighbor IP Address	IP address of the neighbor.
Shared Risk Link Group	Shared risk link group. **NOTE** When SRLG is enabled on an interface, TE information contains SRLG information.

7.6.35 display isis traffic-eng link

Function

The **display isis traffic-eng link** command displays IS-IS TE link information.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

display isis traffic-eng link [{ **level-1** | **level-2** | **level-1-2** } | **verbose**] * [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **traffic-eng link** [{ **level-1** | **level-2** | **level-1-2** } | **verbose**] *

Parameters

Parameter	Description	Value
level-1	Displays TE link information on a Level-1 switch.	-
level-2	Displays TE link information on a Level-2 switch.	-
level-1-2	Displays TE link information on a Level-1-2switch.	-
verbose	Displays detailed TE link information.	-
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the IS level is not specified, Level-1-2 TE link information is displayed.

Example

Display IS-IS TE link information.

```
<HUAWEI> display isis traffic-eng link
                    TE information for ISIS(1)
                    ---------------------------
Level-2 Link Information
------------------------
 0000.0000.0001.00--> 0000.0000.0001.01    Type: MULACC  LinkID: 10.1.1.1
 0000.0000.0002.00--> 0000.0000.0003.00    Type: P2P     LinkID: 10.3.3.9
 0000.0000.0002.00--> 0000.0000.0001.01    Type: MULACC  LinkID: 10.1.1.1
 0000.0000.0003.00--> 0000.0000.0002.00    Type: P2P     LinkID: 10.2.2.9
 0000.0000.0003.00--> 0000.0000.0004.01    Type: MULACC  LinkID: 10.1.1.2
 0000.0000.0004.00--> 0000.0000.0004.01    Type: MULACC  LinkID: 10.1.1.2
 Total Number of TE Links in Level-2 Area: 6, Num Active: 6
```

Table 7-114 Description of the display isis traffic-eng link command output

Item	Description
Type	Link type: • MULACC: multi-point access link • P2P: point-to-point link
LinkID	TE link ID.
Total Number of TE Links in Level-2 Area	Total number of TE links in a Level-2 area.
Num Active	Number of active TE links.

Display detailed IS-IS TE link information.

```
<HUAWEI> display isis traffic-eng link verbose
                    TE information for ISIS(1)
                    ---------------------------

Level-2 Link Information
------------------------
 Link Status In CSPF : INACTIVE
 0000.0000.0001.00-->0000.0000.0002.00        Type: P2P     LinkID: 10.2.2.2
        Process ID       :  1          Router ID       :  10.2.2.1
        Area ID          :  2          Admin Group     :  0
        Te Cost          :  10         Igp Cost        :  10
        Max Bandwidth    :  0          Max Res Bandwidth:  0
        Shared Risk Link Group: 20
        BC Model         :  RDM
        DS-TE Mode       :  Non-standard IETF DS-TE Mode

 Link Status In CSPF : INACTIVE
 0000.0000.0002.00-->0000.0000.0001.00        Type: P2P     LinkID: 10.2.2.1
        Process ID       :  1          Router ID       :  10.2.2.2
        Area ID          :  2          Admin Group     :  0
        Te Cost          :  10         Igp Cost        :  10
        Max Bandwidth    :  0          Max Res Bandwidth:  0
        BC Model         :  RDM
        DS-TE Mode       :  Non-standard IETF DS-TE Mode

 Total Number of TE Links in Level-2 Area: 2, Num Active: 0
```


Table 7-115 Description of the display isis traffic-eng link verbose command output

Item	Description
Link Status In CSPF	Link status in CSPF calculation.
Process ID	IS-IS process ID.
Router ID	Router ID.
Area ID	Area ID.
Admin Group	Administrative group.
Te Cost	TE cost.
Igp Cost	IGP cost.
Max Bandwidth	Maximum bandwidth.
Max Res Bandwidth	Maximum reservable bandwidth.
Shared Risk Link Group	Shared risk link group.
BC Model	Bandwidth constraint model.
DS-TE Mode	DS-TE mode, including standard and non-standard DS-TE modes.

7.6.36 display isis traffic-eng network

Function

The **display isis traffic-eng network** command displays IS-IS TE network information.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

display isis traffic-eng network [**level-1** | **level-2** | **level-1-2**] [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **traffic-eng network** [**level-1** | **level-2** | **level-1-2**]

Parameters

Parameter	Description	Value
level-1	Displays TE network information on a Level-1 router.	-
level-2	Displays TE network information on a Level-2 router.	-
level-1-2	Displays TE network information on a Level-1-2 router.	-
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

This command does not take effect on P2P networks.

If the IS level is not specified, Level-1-2 TE network information is displayed.

Example

Display IS-IS TE network information.

```
<HUAWEI> display isis traffic-eng network
                    TE information for ISIS(1)
                    ---------------------------
Level-2 Network Information
----------------------------
   DIS Router ID   : 10.1.1.9     DIS''s Ip Address          : 10.1.1.1
   Status In CSPF : ACTIVE      Attached Router Count    : 2
   List of Attached Routers
 RouterId : 10.1.1.9   Nbr : 0000.0000.0001.00  Link State : 1
 RouterId : 10.2.2.9   Nbr : 0000.0000.0002.00  Link State : 1
   DIS Router ID   : 10.4.4.9     DIS''s Ip Address          : 10.1.1.2
   Status In CSPF : ACTIVE      Attached Router Count    : 2
   List of Attached Routers
 RouterId : 10.4.4.9   Nbr : 0000.0000.0004.00  Link State : 1
 RouterId : 10.3.3.9   Nbr : 0000.0000.0003.00  Link State : 1
 Total Number of TE Networks in Level-2 Area: 2, Num Active: 2
```


Table 7-116 Description of the display isis traffic-eng network command output

Item	Description
DIS Router ID	Router ID of the DIS.
DIS's Ip Address	IP address of the DIS.
Status In CSPF	Whether CSPF is enabled on the switch.
Attached Router Count	Number of connected switches, including the local switch.
List of Attached Routers	List of connected switches.
RouterId	Router ID of the local switch.
Nbr	IP address of the neighbor.
Link State	Link status.
Total Number of TE Networks in Level-2 Area	Total number of networks on which TE is enabled in a Level-2 area.

7.6.37 display isis traffic-eng statistics

Function

The **display isis traffic-eng statistics** command displays IS-IS TE statistics.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

display isis traffic-eng statistics [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **traffic-eng statistics**

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID. If no process ID is specified, TE statistics about all IS-IS processes are displayed.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use this command to view IS-IS TE statistics on the switch, including the IS level, cost style, IS-IS TE level, and IS-IS router ID.

Example

Display IS-IS TE statistics.

```
<HUAWEI> display isis traffic-eng statistics
                    TE information for ISIS(1)
                    ---------------------------
TE Statistics Information
-------------------------
 IS-IS System Type                      : Level-1-2
 IS-IS Cost Style Status                : Wide
 IS-IS Level-1 Traffic Engineering Status   : Disabled
 IS-IS Level-2 Traffic Engineering Status   : Enabled
 IS-IS Router ID                        : 10.1.1.9
```

Table 7-117 Description of the display isis traffic-eng statistics command output

Item	Description
IS-IS System Type	IS level.
IS-IS Cost Style Status	Cost style of the switch.
IS-IS Level-1 Traffic Engineering Status	TE status of a Level-1 router.
IS-IS Level-2 Traffic Engineering Status	TE status of a Level-2 router.
IS-IS Router ID	Router ID of an IS-IS router.

7.6.38 display isis traffic-eng sub-tlvs

Function

The **display isis traffic-eng sub-tlvs** command displays the types of sub-TLVs carrying DS-TE parameters.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

display isis traffic-eng sub-tlvs [*process-id* | **vpn-instance** *vpn-instance-name*]

display isis *process-id* **traffic-eng sub-tlvs**

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display isis traffic-eng sub-tlvs** command to check the types of sub-TLVs carrying DS-TE parameters.

Example

Display the types of sub-TLVs carrying DS-TE parameters.

```
<HUAWEI> display isis traffic-eng sub-tlvs
IS-IS(1) SubTlv Information
---------------------------------
 Unreserved sub-pool bandwidth sub-tlv value : 251
 Bandwidth constraint sub-tlv value          : 252
 LO multiplier sub-tlv value                 : 253
```

Table 7-118 Description of the display isis traffic-eng sub-tlvs command output

Item	Description
IS-IS(1) SubTlv Information	Sub-TLV information about IS-IS process 1.
Unreserved sub-pool bandwidth sub-tlv value	Sub-TLV of the unreserved sub-pool bandwidth.
Bandwidth constraint sub-tlv value	Sub-TLV of the bandwidth constraint.
LO multiplier sub-tlv value	Sub-TLV of the local overbooking multiplier.

7.6.39 domain-authentication-mode

Function

The **domain-authentication-mode** command configures an IS-IS routing domain to authenticate received Level-2 packets using the specified authentication mode and password and adds authentication information to Level-2 packets to be sent.

The **undo domain-authentication-mode** command cancels authenticating Level-2 packets and deletes the added authentication information from Level-2 packets.

By default, the system neither encapsulates generated Level-2 packets with authentication information nor authenticates received Level-2 packets.

Format

domain-authentication-mode { **simple** { **plain** *plain-text* | [**cipher**] *plain-cipher-text* } | **md5** { [**cipher**] *plain-cipher-text* | **plain** *plain-text* } } [**ip** | **osi**] [**snp-packet** { **authentication-avoid** | **send-only** } | **all-send-only**]

domain-authentication-mode keychain *keychain-name* [**snp-packet** { **authentication-avoid** | **send-only** } | **all-send-only**]

domain-authentication-mode hmac-sha256 key-id *key-id* { **plain** *plain-text* | [**cipher**] *plain-cipher-text* } [**snp-packet** { **authentication-avoid** | **send-only** } | **all-send-only**]

undo domain-authentication-mode

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
simple	Transmits the password in plain text. **NOTICE** Simple authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.	-

Parameter	Description	Value
plain *plain-text*	Specifies the authentication password in plain text. You can enter only the password in plain text. When you view the configuration file, the password is displayed in plain text. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	The value is a string of case-sensitive characters without spaces. The value contains digits and letters. When the authentication mode is **simple**, the value is a string of 1 to 16 characters. When the authentication mode is **md5** or **hmac-sha256**, the value is a string of 1 to 255 characters.
cipher *plain-cipher-text*	Specifies the authentication password in cipher text. You can enter the password in plain or cipher text. When you view the configuration file, the password is displayed in cipher text. By default, the password is in cipher text.	The value is a string of case-sensitive characters without spaces. The value contains digits and letters. When the authentication mode is **simple**, the value is a string of 1 to 16 characters in plain text or a string of 32 or 48 characters in cipher text. When the authentication mode is **md5** or **hmac-sha256**, the value is a string of 1 to 255 characters in plain text or a string of 20 to 392 characters in cipher text.
md5	Transmits the password that is encrypted using MD5. **NOTICE** MD5 authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.	-
keychain *keychain-name*	Specifies the keychain that changes with time. Currently, IS-IS supports only HMAC-MD5 and HMAC-SHA256 algorithms.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.

Parameter	Description	Value
ip	Indicates the IP authentication password. When neither **ip** nor **osi** is specified, the default parameter **osi** is used.	-
osi	Indicates the OSI authentication password. When neither **ip** nor **osi** is specified, the default parameter **osi** is used.	-
snp-packet	Authenticates SNPs.	-
authentication-avoid	Encapsulates generated LSPs but not SNPs with authentication information and authenticates received LSPs but not SNPs.	-
send-only	Encapsulates generated LSPs and SNPs with authentication information, and authenticates received LSPs but not SNPs.	-
all-send-only	Encapsulates generated LSPs and SNPs with authentication information, but does not authenticate received LSPs and SNPs.	-
hmac-sha256	Encapsulates generated packets with the HMAC-SHA256 authentication and a password encrypted by the HMAC-SHA256 algorithm and authenticates received packets.	-
key-id *key-id*	Indicates key ID of the HMAC-SHA256 algorithm.	It is an integer ranging from 0 to 65535.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, the IS-IS packets to be sent are not encapsulated with authentication information, and the received packets are not authenticated. If a user sends malicious packets to attack a network, information on the entire network may be stolen. Therefore, you can configure IS-IS authentication to improve the network security.

The domain authentication password is encapsulated into Level-2 IS-IS packets. Only the packets that pass the domain authentication can be accepted. Therefore, you can configure IS-IS domain authentication to authenticate Level-2 area.

Precautions

This command is valid in all the topologies in the specified IS-IS process and is only valid for Level-2 or Level-1-2 routers.

By using this command, you can discard all the Level-2 packets whose domain authentication password does not contain the one set through this command. At the same time, IS-IS adds the configured domain authentication password in all the Level-2 packets carrying routing information sent from the local node.

The authentication takes effect on the interface with the password. The port without the password can still receive the LSP and SNP with password.

Example

Set the domain authentication mode to HMAC-SHA256, authentication password to **YsHsjx_202206**, and key ID to **33** to authenticate Level-2 packets.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] domain-authentication-mode hmac-sha256 key-id 33 cipher YsHsjx_202206
```

7.6.40 filter-policy export (IS-IS)

Function

The **filter-policy export** command configures a filtering policy to allow IS-IS to filter the imported routes to be advertised.

The **undo filter-policy export** command cancels the filtering function.

By default, IS-IS does not filter the imported routes to be advertised.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **export** [*protocol* [*process-id*]]

undo filter-policy [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name*] **export** [*protocol* [*process-id*]]

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix-list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy to filter routes based on tag and other protocol parameters.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
protocol	Specifies the imported routes that need to be filtered when the routes are advertised. If this parameter is not specified, all the imported routes to be advertised are filtered.	The value can be **direct**, **static**, **unr**, **rip**, **bgp, ospf**, or another **isis** process.
process-id	Specifies the process ID if *protocol* is **rip**, **ospf**, or another **isis** process.	The value is an integer that ranges from 1 to 65535.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When IS-IS and other routing protocols are running on the network, and a boundary router in the IS-IS routing domain has imported routes of other routing

protocols, the boundary router will advertise all the imported routes to its IS-IS neighbors by default. To advertise some of the imported routes to neighbors, use the **filter-policy export** command.

Precautions

Running the **filter-policy export** command does not affect the routes on the local device, but only advertises specific imported routes to IS-IS neighbors.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all external routes of the specified routing domain that are imported by IS-IS are advertised to the specified neighbor.

Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all external routes of the specified routing domain that are imported by IS-IS are advertised to the specified neighbor.

Example

Configure IS-IS to filter the imported routes using ACL 2000 before advertising the routes to other switches.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] filter-policy 2000 export
```

7.6.41 filter-policy import (IS-IS)

Function

The **filter-policy import** command configures a filtering policy to allow IS-IS to filter the received routes to be added to the IP routing table.

The **undo filter-policy import** command cancels the filtering function.

By default, IS-IS does not filter the received routes to be added to the IP routing table.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **import**

undo filter-policy [*acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name*] **import**

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy to filter routes based on tag and other protocol parameters.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS routing entries need to be added to an IP routing table to guide IP packet forwarding. If an IS-IS routing table has routes destined for a specific network segment, but these routes are not expected to be added to an IP routing table, run the **filter-policy import** command with specified parameters to allow only the needed IS-IS routes to be added to the IP routing table.

Precautions

Running the **filter-policy import** command on a router does not affect LSP flooding and LSDB synchronization on the router, but affects the local IP routing table.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all routes received by IS-IS are delivered to the IP routing table.

Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all routes received by IS-IS are delivered to the IP routing table.

Example

Configure IS-IS to filter the received routes using ACL 2000 and add the routes matching the filtering conditions to the IP routing table.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] filter-policy 2000 import
```

7.6.42 flash-flood

Function

The **flash-flood** command enables LSP fast flooding to speed up IS-IS network convergence.

The **undo flash-flood** command disables LSP fast flooding.

By default, LSP fast flooding is disabled.

Format

flash-flood [*lsp-count* | **max-timer-interval** *interval* | [**level-1** | **level-2**]] *

undo flash-flood [*lsp-count* | **max-timer-interval** *interval* | [**level-1** | **level-2**]] *

Parameters

Parameter	Description	Value
lsp-count	Specifies the maximum number of LSPs to be flooded at a time on an interface.	The value is an integer that ranges from 1 to 15. The default value is 5.
max-timer-interval *interval*	Specifies the maximum interval for LSP flooding.	The value is an integer that ranges from 10 to 50000, in milliseconds. The default value is 10 ms.
level-1	Enables LSP fast flooding in Level-1. If no level is specified, by default, LSP fast flooding is enabled in both Level-1 and Level-2.	-

Parameter	Description	Value
level-2	Enables LSP fast flooding in Level-2. If no level is specified, by default, LSP fast flooding is enabled in both Level-1 and Level-2.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In most cases, when an IS-IS router receives new LSPs from other routers, it updates the LSPs in its LSDB and periodically floods the updated LSPs according to a timer.

LSP fast flooding speeds up LSDB synchronization because it allows a device to flood fewer LSPs than the specified number before route calculation when the device receives one or more new LSPs. This mechanism also speeds up network convergence.

Precautions

You can specify the number of LSPs to be flooded each time. The number is valid for all IS-IS interfaces. If the number of LSPs to be flooded is greater than the specified value, the *lsp-count* number of the LSPs are sent. If the timer is configured and does not time out before route calculation, LSPs are flooded immediately; otherwise, LSPs are flooded when the timer times out.

Example

Enable LSP fast flooding, configure each interface to send a maximum of six LSPs at a time, and set the maximum interval for sending LSPs to 100 ms.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] flash-flood 6 max-timer-interval 100
```

7.6.43 frr (IS-IS)

Function

The **frr** command enables the IS-IS Auto FRR function and displays the IS-IS Auto FRR view.

The **undo frr** command disables the IS-IS Auto FRR function.

By default, the IS-IS Auto FRR function is disabled.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

frr

undo frr

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage scenario

With the development of networks, the Voice over IP (VoIP) and on-line video services require high-quality real-time transmission. Nevertheless, if an IS-IS fault occurs, multiple processes, including fault detection, LSP update, LSP flooding, route calculation, and FIB entry delivery, must be performed to switch traffic to a new link. This results in a lengthy traffic interruption, which cannot meet the requirement for real-time services.

IS-IS Auto FRR can fast switch traffic to a backup link, ensuring millisecond-level traffic interruption. This protects traffic and improves IS-IS network reliability.

Precautions

After running this command, run the **loop-free-alternate** command to calculate a loop-free backup route.

Example

Enter the IS-IS FRR view.


```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] frr
[HUAWEI-isis-1-frr] loop-free-alternate
```

7.6.44 frr-policy route

Function

The **frr-policy route** command configures a filtering policy to allow IS-IS to filter the IS-IS backup routes to be added in the IP routing table.

The **undo frr-policy route** command cancels the filtering function.

By default, IS-IS does not filter the IS-IS backup routes to be added in the IP routing table.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

frr-policy route route-policy *route-policy-name*

undo frr-policy route

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a route-policy to filter IS-IS backup routes.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

IS-IS FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The filtering policy can be configured as required. In this case, the IS-IS backup route that satisfies specified rules can be added to the IP routing table and delivered to the forwarding table. When a fault occurs on the route, the system can fast switch the forwarded traffic to the IS-IS backup route to protect traffic.

You can use IP prefix lists or ACLs to filter IS-IS backup routes.

Precautions

If you run the **frr-policy route** command multiple times, only the latest configuration takes effect.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, IS-IS adds all backup routes to the IP routing table.

Example

Configure IS-IS to filter IS-IS backup routes using route-policy and add the routes matching the route-policy **abc** to the IP routing table.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] frr
[HUAWEI-isis-1-frr] frr-policy route route-policy abc
```

7.6.45 graceful-restart (IS-IS)

Function

The **graceful-restart** command enables the graceful restart (GR) function for an IS-IS process.

The **undo graceful-restart** command disables the GR function for an IS-IS process.

By default, the GR function is disabled for an IS-IS process.

Format

graceful-restart

undo graceful-restart

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

This command is valid for all topologies in an IS-IS process.

When an IS-IS process on a device is restarted, neighbors delete neighbor relationships with the device and delete LSPs from the device. As a result, routes of neighbors are calculated incorrectly and packets are lost. Consequently, the network is temporarily interrupted.

To solve this problem, you can enable GR of the IS-IS process by running the **graceful-restart** command.

After the **graceful-restart** command is run, the device can notify its restart status to neighbors and permit neighbors to maintain neighbor relationships. In this manner, nonstop packet forwarding is implemented.

Configuration Impact

If IS-IS GR is enabled on a router, the holdtimes of this router's neighbors are automatically changed to 60s if they are smaller than 60s, and the holdtimes of this router's neighbors are kept unchanged if they are equal to or greater than 60s. If a router is faulty in non-GR scenarios, its neighbors need to take 60 seconds to detect the fault. A large number of packets may be discarded within the 60 seconds.

To resolve this problem, run the **graceful-restart no-impact-holdtime** command to configure the holdtimes of the neighbors to remain unchanged after IS-IS GR is enabled.

Example

Enable the GR function for IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] graceful-restart
```

7.6.46 graceful-restart no-impact-holdtime

Function

The **graceful-restart no-impact-holdtime** command configures the holdtime of an IS-IS neighbor to remain unchanged after IS-IS GR is enabled.

The **undo graceful-restart no-impact-holdtime** command cancels the configuration.

By default, after IS-IS GR is enabled, the holdtime of an IS-IS neighbor is automatically changed to 60s if it is smaller than 60s, and the holdtime of an IS-IS neighbor is kept unchanged if it is equal to or greater than 60s.

Format

graceful-restart no-impact-holdtime

undo graceful-restart no-impact-holdtime

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If IS-IS GR is enabled on a router, the holdtimes of this router's neighbors are automatically changed to 60s if they are smaller than 60s, and the holdtimes of this router's neighbors are kept unchanged if they are equal to or greater than 60s. If a router is faulty in non-GR scenarios, its neighbors need to take 60 seconds to detect the fault. A large number of packets may be discarded within the 60 seconds.

To resolve this problem, run the **graceful-restart no-impact-holdtime** command to configure the holdtimes of the neighbors to remain unchanged after IS-IS GR is enabled. After you run this command, the router can still fast detect neighbor status, implementing rapid network convergence.

Prerequisites

You have run the **graceful-restart (IS-IS)** command in the IS-IS view.

Example

Configure the holdtime of an IS-IS neighbor to remain unchanged after IS-IS GR is enabled.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] graceful-restart
[HUAWEI-isis-1] graceful-restart no-impact-holdtime
```


7.6.47 graceful-restart interval

Function

The **graceful-restart interval** command sets the GR T3 timer.

The **undo graceful-restart interval** command restores the default T3 timer.

By default, the GR T3 timer is 300 seconds.

Format

graceful-restart interval *interval-value*

undo graceful-restart interval

Parameters

Parameter	Description	Value
interval-value	Specifies the GR T3 timer.	The value is an integer that ranges from 30 to 1800, in seconds. Setting a value greater than that of the GR T2 timer specified using the **graceful-restart t2-interval** command is recommended. If the value is smaller than that of the GR T2 timer, the GR may fail.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS GR can be classified into restarting GR and starting GR based on the restart type. The restarting GR is triggered by a master/slave main control board switchover or an IS-IS process restart, while the starting GR is triggered by a restart of the IS-IS-enabled router.

During a restarting GR, the restarter restarts the protocol and starts the T1, T2, and T3 timers at the same time. The value of the T1 timer indicates the longest time during which the GR restarter waits for the LSP from the GR helper. The value of the T2 timer indicates the longest time during which the system waits for the LSDB synchronization. The value of the T3 timer indicates the longest time that a GR lasts. A router disables the T3 timer after the LSDB synchronization ends in all areas. If LSDBs are not synchronized yet when the T3 timer expires, the GR fails.

You can run the **graceful-restart interval** command to adjust the value of the T3 timer so that the LSDB synchronization can end before the T3 timer expires, which prevents the GR failure.

Prerequisites

The GR of the IS-IS process has been enabled using the **graceful-restart** command.

Configuration Impact

If the **graceful-restart interval** command is run on an IS-IS-enabled router, *interval-value* is used as the holdtime of its neighbor during the GR.

Precautions

The **graceful-restart interval** command is applicable only to restarting GRs.

Example

Set the GR interval for IS-IS process 1 to 120 seconds (2 minutes).

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] graceful-restart
[HUAWEI-isis-1] graceful-restart interval 120
```

7.6.48 graceful-restart suppress-sa

Function

The **graceful-restart suppress-sa** command suppresses the suppress-advertisement (SA) bit of the restart TLV.

The **undo graceful-restart suppress-sa** command restores the default setting.

By default, the SA bit of the restart TLV is not suppressed.

Format

graceful-restart suppress-sa

undo graceful-restart suppress-sa

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The switch that starts for the first time (not including GR) does not maintain the forwarding status. If the switch does not start for the first time, the LSPs generated when the switch run last time may exist in the LSP database of other switches in the network.

The sequence number of the LSP fragment is reinitialized when the switch starts, so the LSPs stored in the LSP database of other switches seem to be newer than the LSPs generated after the switch starts. This leads to the black hole in the network, and the black hole lasts until the switch regenerates its LSPs and advertises the LSPs with the highest sequence number.

If the neighbor suppresses the advertisement of the adjacency relationship to this switch during the switch starting until the switch advertises the updated LSPs, the preceding case can be avoided.

Prerequisites

GR has been enabled for the IS-IS process using the **graceful-restart** command.

Example

Suppress the SA bit in the restart TLV of IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] graceful-restart
[HUAWEI-isis-1] graceful-restart suppress-sa
```

7.6.49 graceful-restart t2-interval

Function

The **graceful-restart t2-interval** command configures a value for the T2 timer during an IS-IS GR.

The **undo graceful-restart t2-interval** command restores the default value of the T2 timer.

By default, the GR T2 timer is 60 seconds.

Format

graceful-restart t2-interval *interval-value*

undo graceful-restart t2-interval

Parameters

Parameter	Description	Value
interval-value	Specifies the value of the T2 timer during the IS-IS GR.	The value is an integer that ranges from 30 to 1800, in seconds. Setting a value smaller than that of the GR T3 timer specified using the **graceful-restart interval** command is recommended.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS GR can be classified into restarting GR and starting GR based on the restart type. The restarting GR is triggered by a master/slave main control board switchover or an IS-IS process restart, while the starting GR is triggered by a restart of the IS-IS-enabled router.

A GR restarter starts the T1 and T2 timers during all GRs. The value of the T1 timer indicates the longest time during which the GR restarter waits for the LSP from the GR helper. The value of the T2 timer indicates the longest time during which the system waits for the LSDB synchronization. The T2 timer is disabled after the LSDB synchronization among Level-1 or Level-2 routers ends. If LSDBs are not synchronized yet when the T2 timer expires, the GR fails.

You can run the **graceful-restart t2-interval** command to adjust the value of the T2 timer so that the LSDB synchronization can end before the T2 timer expires, which prevents the GR failure.

Prerequisites

The GR of the IS-IS process has been enabled using the **graceful-restart** command.

Precautions

If the value of the T2 timer is too small, a GR may fail. Keeping the default value is recommended. If you want to adjust it, configure a value to meet the conditions on the live network.

Example

Set the GR T2 timer of IS-IS process 1 to 120 seconds (2 minutes).

```
<HUAWEI> system-view
```



```
[HUAWEI] isis 1
[HUAWEI-isis-1] graceful-restart
[HUAWEI-isis-1] graceful-restart t2-interval 120
```

7.6.50 import-route (IS-IS)

Function

The **import-route** command configures IS-IS to import routes from other routing protocols.

The **undo import-route** command restores the default setting.

By default, IS-IS does not import routes from other routing protocols.

Format

import-route { { **rip** | **isis** | **ospf** } [*process-id*] | **static** | **direct** | **unr** | **bgp** [**permit-ibgp**] } [**cost-type** { **external** | **internal** } | **cost** *cost* | **tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

import-route { { **rip** | **isis** | **ospf** } [*process-id*] | **direct** | **unr** | **bgp** [**permit-ibgp**] } **inherit-cost** [**tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

import-route limit *limit-number* [**threshold-alarm upper-limit** *upper-limit-value* **lower-limit** *lower-limit-value*] { **level-1** | **level-2** | **level-1-2** }

undo import-route { { **rip** | **isis** | **ospf** } [*process-id*] | **static** | **direct** | **unr** | **bgp** [**permit-ibgp**] } [**cost-type** { **external** | **internal** } | **cost** *cost* | **tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

undo import-route { { **rip** | **isis** | **ospf** } [*process-id*] | **direct** | **unr** | **bgp** [**permit-ibgp**] } **inherit-cost** [**tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

undo import-route limit [*limit-number*] [**threshold-alarm upper-limit** *upper-limit-value* **lower-limit** *lower-limit-value*] { **level-1** | **level-2** | **level-1-2** }

Parameters

Parameter	Description	Value
rip	Indicates that the routing protocol from which routes are imported is RIP.	-
isis	Indicates that the routing protocol from which routes are imported is IS-IS.	-
ospf	Indicates that the routing protocol from which routes are imported is OSPF.	-
process-id	Specifies a process ID. When *protocol* is **rip**, **ospf**, or **isis**, a process ID needs to be specified. The default process ID is 1.	The value is an integer that ranges from 1 to 65535.

Parameter	Description	Value
static	Indicates that the imported routes are active static routes.	-
direct	Indicates that the imported routes are direct routes.	-
unr	Specifies the imported source routing protocol as **unr**. User Network Route (UNR) is allocated if dynamic routing protocols cannot be used when users are getting online.	-
bgp	Indicates that the routing protocol from which routes are imported is BGP.	-
permit-ibgp	Specifies the imported source route as an IBGP route. If you do not configure this parameter, only the EBGP routes are imported.	-
cost-type { **external** | **internal** }	Indicates the cost type of the imported routes. By default, the cost type is **external**. The configuration of this parameter will affect the costs of imported routes. ● If the cost type of an imported route is configured as **external**, the cost of the imported route equals the cost of the original route plus 64. ● If the cost type of an imported route is configured as **internal**, the imported route inherits the cost of the original route. **NOTE** If the cost style of the switch is wide, compatible, or wide-compatible, the cost types of imported routes are not differentiated between **external** and **internal**.	-
cost *cost*	Specifies the cost value of imported routes.	If the cost style of the switch is wide or wide-compatible, the cost value of imported routes ranges from 0 to 4261412864. Otherwise, the value ranges from 0 to 63. The default value is 0.

Parameter	Description	Value
tag *tag*	Specifies the administrative tag of imported routes.	The value is an integer that ranges from 1 to 4294967295.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
level-1	Indicates that routes are imported into Level-1 routing tables. If no level is specified, routes are imported to Level-2 routing tables by default.	-
level-2	Indicates that routes are imported into Level-2 routing tables. If no level is specified, routes are imported to Level-2 routing tables by default.	-
level-1-2	Indicates that routes are imported into Level-1 and Level-2 routing tables. If no level is specified, routes are imported to Level-2 routing tables by default.	-
inherit-cost	Indicates that the original cost value of imported external routes is retained. When IS-IS is configured to retain the original cost value of imported routes, the cost style and cost value cannot be set for the imported routes.	-
limit *limit-number*	Specifies the maximum number of external routes allowed to be imported to the IS-IS area.	The value is an integer ranging from 1 to 10000000.
threshold-alarm	Specifies the alarm threshold for imported routes.	-
upper-limit *upper-limit-value*	Specifies the upper alarm threshold for imported routes.	The value is an integer ranging from 1 to 100. The default value is 80.
lower-limit *lower-limit-value*	Specifies the lower alarm threshold for imported routes.	The value is an integer ranging from 1 to 100. The default value is 70.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When IS-IS and other routing protocols are deployed on a network, you can enable the traffic within an IS-IS routing domain to reach a destination outside the IS-IS routing domain using either of the following methods:

- Configure boundary devices in the IS-IS routing domain to advertise default routes to the IS-IS routing domain.
- Configure boundary devices in the IS-IS routing domain to import routes from other routing domains into the IS-IS routing domain.

If there are multiple boundary devices in the IS-IS routing domain, optimal routes destined for another routing domain need to be selected. This requires all devices in the IS-IS routing domain learn all or some external routes. Configure boundary devices in the IS-IS routing domain to import routes from other routing domains into the IS-IS routing domain. Alternatively, run the **route-policy** *route-policy-name* command to import some external routes from other routing domains.

Precautions

When the routes of the other protocols are imported, you can set the cost value and cost style for the imported route. You can also configure IS-IS to retain the original cost value of the imported external route. During route advertisement and route calculation, the original cost values of these routes are used. In this case, the cost style and cost value of the imported routes cannot be set, and static routes cannot be imported.

After the **import-route direct** command is executed, routes to the network segment where the IP address of the management interface belongs are also imported in the ISIS routing table. Therefore, use this command with caution.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all routes of the specified routing domain are imported to the IS-IS routing table.

Example

Configure IS-IS to import static routes and set the cost value of the routes to 15.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] import-route static cost 15
```

Configure IS-IS to import OSPF routes and retain the original cost value of the routes.


```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] import-route ospf inherit-cost
```

7.6.51 import-route isis level-1 into level-2

Function

The **import-route isis level-1 into level-2** command configures route leaking from Level-1 areas to Level-2 areas.

The **undo import-route isis level-1 into level-2** command prohibits route leaking from Level-1 areas to Level-2 areas.

By default, all Level-1 routing information, excluding information about default routes, is leaked to Level-2 areas.

Format

import-route isis level-1 into level-2 [**filter-policy** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct allow-filter-policy**] *

undo import-route isis level-1 into level-2 [**filter-policy** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct allow-filter-policy**] *

Parameters

Parameter	Description	Value
filter-policy	Specifies a filtering policy.	-
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list. Only the routes that match the IP prefix can be leaked to Level-2 areas.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
tag *tag*	Specifies the administrative tag value of imported routes.	The value is an integer that ranges from 1 to 4294967295.
direct allow-filter-policy	Specifies the filtering policy to filter the direct routes. Only the IS-IS Level-1 area direct routing information that matches the filtering policy can be shared with the Level-2 area with this parameter, and all Level-1 area direct routing information will be shared with the Level-2 area without this parameter.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS is designed for a hierarchical network. This means that routes of Level-1 areas are leaked to Level-2 areas, whereas routes of Level-2 areas are not leaked to Level-1 areas.

The import-route isis level-1 into level-2 command can be run only on Level-1-2 routers to allow some or no Level-1 routes to be leaked to Level-2 areas. For example, there are two Level-1-2 routers in a Level-1 area. You can run the import-route isis level-1 into level-2 command on one Level-1-2 router to allow some Level-1 routes to be leaked to the Level-2 area, and run the import-route isis level-1 into level-2 command on the other Level-1-2 router to allow the remaining Level-1 routes to be leaked to the Level-2 area. Then, traffic that is sent from the Level-2 area and destined for different network segments in Level-1 area will be forwarded to different Level-1-2 routers. This allows route selection to be controlled.

Precautions

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all routes in the Level-1 area leak to the Level-2 area.

Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all routes in the Level-1 area leak to the Level-2 area.

Example

Control route leaking from Level-1 areas to Level-2 areas on Level-1-2 routers using filtering policy 2000.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] import-route isis level-1 into level-2 filter-policy 2000
```

7.6.52 import-route isis level-2 into level-1

Function

The **import-route isis level-2 into level-1** command configures route leaking from Level-2 areas to Level-1 areas.

The **undo import-route isis level-2 into level-1** command prohibits route leaking from Level-2 areas to Level-1 areas.

By default, Level-2 routing information is not leaked to Level-1 areas.

Format

import-route isis level-2 into level-1 [**filter-policy** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct** { **allow-filter-policy** | **allow-up-down-bit** } *] *

undo import-route isis level-2 into level-1 [**filter-policy** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct** { **allow-filter-policy** | **allow-up-down-bit** } *] *

Parameters

Parameter	Description	Value
filter-policy	Specifies a filtering policy.	-
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list. Only the routes that match the specified IP prefix can be imported.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
tag *tag*	Specifies the administrative tag value of imported routes.	The value is an integer that ranges from 1 to 4294967295.
direct allow-filter-policy	Indicates the filtering policy to filter the direct routes. Only the IS-IS Level-1 area direct routing information that matches the filtering policy can leak to the Level-2 area with this parameter, and all Level-1 area direct routing information is available to the Level-2 area without this parameter.	-
direct allow-up-down-bit	Indicates that the Up or Down bit is used during the leak of directs routes. If **direct allow-up-down-bit** is specified, the direct routes that have already leaked to the Level-1 area have the lowest priority and cannot leak back.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When multiple Level-1-2 switches in a Level-1 area are connected to a Level-2 area, Level-1 switches are unable to know the network topology of the Level-2 area, but forward traffic to the nearest Level-1-2 switch (based on the route with the smallest cost). For Level-2 switches, however, this route may not be the optimal one. Therefore, you need to allow some Level-2 routes to be leaked to the Level-1 area to help Level-1 switches select the optimal route for forwarding traffic to the Level 2 area.

To solve the preceding problem, IS-IS provides the route leaking function to enable Level-1 devices to choose the best path for traffic forwarding.

The **import-route isis level-2 into level-1** command can be run only on Level-1-2 routers to allow all or some Level-2 routes to be leaked to the Level-1 area.

Precautions

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all routes in the Level-2 area leak to the Level-1 area.

Creating an IP prefix list or route-policy before it is referenced is recommended. By default, nonexistent IP prefix lists or route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list or route-policy is referenced using the current command, all routes in the Level-1 area leak to the Level-2 area.

Example

Configure IS-IS to perform route leaking from a Level-2 area to a Level-1 area using filtering policy 2000.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] import-route isis level-2 into level-1 filter-policy 2000
```

7.6.53 isis

Function

The **isis** command starts an IS-IS process and a specified VPN instance, and displays the IS-IS view.

The **undo isis** command deletes a specified IS-IS process.

By default, no IS-IS instance exists on the network.

Format

isis [*process-id*] [**vpn-instance** *vpn-instance-name*]

undo isis *process-id*

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Before you configure IS-IS functions and interface-related features, run the **isis** command to create an IS-IS process and enable IS-IS on the interface.

On a large-scale network, if a large number of switches run IS-IS, there will be a huge number of routes, increasing maintenance costs, slowing down route convergence, and affecting network stability. To resolve the problem, you can run the **isis** *process-id* command to start multi-processes to reduce the number of routes to be maintained.

In addition, to ensure that different services are forwarded properly on the network, you can run the **isis vpn-instance** *vpn-instance-name* command to start multiple IS-IS processes on one device to isolate these services.

Follow-up Procedure

After the **isis** command is used to enable an IS-IS process, run the **network-entity** command to set a NET for the switch, and run the **isis enable** command to enable IS-IS on each interface that needs to run IS-IS. You can start IS-IS only when these configurations are completed.

Precautions

One IS-IS process can be bound to only one VPN instance. Multiple IS-IS interfaces can be bound to one VPN instance.

If a VPN instance is deleted, the IS-IS process bound to the VPN instance is deleted.

When creating an IS-IS process, bind it to a VPN instance. An existing IS-IS process cannot be bound to any VPN instance.

Example

Bind IS-IS process 2 to VPN instance **vpn1**.

```
<HUAWEI> system-view
[HUAWEI] isis 2 vpn-instance vpn1
```

7.6.54 isis authentication-mode

Function

The **isis authentication-mode** command configures an IS-IS interface to authenticate Hello packets using the specified mode and password.

The **undo isis authentication-mode** command cancels the authentication and deletes the authentication information in Hello packets.

By default, no authentication information is added to Hello packets and no authentication is performed on received Hello packets.

Format

isis authentication-mode { **simple** | **md5** } { **plain** *plain-text* | [**cipher**] *plain-cipher-text* } [**level-1** | **level-2**] [**ip** | **osi**] [**send-only**]

isis authentication-mode keychain *keychain-name* [**level-1** | **level-2**] [**send-only**]

isis authentication-mode hmac-sha256 key-id *key-id* { **plain** *plain-text* | [**cipher**] *plain-cipher-text* } [**level-1** | **level-2**] [**send-only**]

undo isis authentication-mode [**level-1** | **level-2**]

undo isis authentication-mode keychain *keychain-name* [**level-1** | **level-2**] [**send-only**]

undo isis authentication-mode { **simple** { **plain** *plain-text* | **cipher** *plain-cipher-text* } | **md5** { **cipher** *plain-cipher-text* | **plain** *plain-text* } } [**level-1** | **level-2**] [**ip** | **osi**] [**send-only**]

undo isis authentication-mode hmac-sha256 key-id *key-id* { **plain** *plain-text* | **cipher** *plain-cipher-text* } [**level-1** | **level-2**] [**send-only**]

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **keychain** *keychain-name* parameter.

Parameters

Parameter	Description	Value
simple	Indicates that the password is transmitted in plain text. **NOTICE** Simple authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.	-
plain *plain-text*	Indicates that the password is in plain text. Only a plain-text password can be entered. The password in the configuration file is displayed in plain text. **NOTICE** If **plain** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	The value is a string of case-sensitive characters. It contains letters and digits without spaces. In **simple** authentication mode, the value is a string of 1 to 16 characters. In **md5** or **hmac-sha256** authentication mode, the value is a string of 1 to 255 characters.
cipher *plain-cipher-text*	Indicates that the password is in cipher text. A plain-text or cipher-text password can be entered. The password in the configuration file is displayed in cipher text. By default, the password is in cipher text.	The value is a string of case-sensitive characters. It contains letters and digits without spaces. In **simple** authentication mode, the value is a string of 1 to 16 characters in plain text or a string of 32 characters in cipher text. In **md5** or **hmac-sha256** authentication mode, the value is a string of 1 to 255 characters in plain text or a string of 20 to 392 characters in cipher text.
md5	Indicates that the password to be transmitted is encrypted using MD5. **NOTICE** MD5 authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.	-

Parameter	Description	Value
level-1	Indicates Level-1 authentication. When the link type of an IS-IS interface is Level-1-2, if **level-1** and **level-2** are not specified, both Level-1 and Level-2 Hello packets are configured with the authentication mode and password.	-
level-2	Indicates Level-2 authentication. When the link type of an IS-IS interface is Level-1-2, if **level-1** and **level-2** are not specified, both Level-1 and Level-2 Hello packets are configured with the authentication mode and password.	-
ip	Indicates the IP authentication password. This parameter cannot be configured in keychain authentication mode. If parameters **ip** and **osi** are not specified, the parameter **osi** is used by default.	-
osi	Indicates the OSI authentication password. This parameter cannot be configured in **keychain** authentication mode. If parameters **ip** and **osi** are not specified, the parameter **osi** is used by default.	-
send-only	Encapsulates sent Hello packets with authentication information but does not authenticate received Hello packets.	-
keychain *keychain-name*	Indicates that the password is a keychain that changes with time. This parameter takes effect only when *keychain-name* is set using the **keychain** command. Currently, IS-IS supports only HMAC-MD5 and HMAC-SHA256 algorithms.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.

Parameter	Description	Value
hmac-sha256	Encapsulates generated packets with the HMAC-SHA256 authentication and a password encrypted by the HMAC-SHA256 algorithm and authenticates received packets.	-
key-id *key-id*	Indicates key ID of the HMAC-SHA256 algorithm.	It is an integer ranging from 0 to 65535.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To improve network security, authenticate received packets or encapsulate sent packets with authentication information. Only the packets that pass the authentication can be transmitted on the network.

You can use the **isis authentication-mode** command to discard the Hello packets whose authentication passwords are different from the authentication password configured using this command. At the same time, IS-IS adds the configured interface authentication password into all the Hello packets sent from the local node.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

If a broadcast interface is emulated as a P2P interface using the **isis circuit-type** command and then restored to the broadcast interface using the **undo isis circuit-type** command, the authentication configuration of the IS-IS area is restored to the default setting.

Example

Set HMAC-SHA256 authentication password **YsHsjx_202206** key id **33** on VLANIF100

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis authentication-mode hmac-sha256 key-id  33 cipher YsHsjx_202206
```


Set HMAC-SHA256 authentication password **YsHsjx_202206** key id **33** GE0/0/1

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis authentication-mode hmac-sha256 key-id  33 cipher YsHsjx_202206
```

7.6.55 isis bfd

Function

The **isis bfd** command sets values of BFD session parameters on a specified IS-IS interface.

The **undo isis bfd** command restores the default values of BFD session parameters on a specified IS-IS interface.

By default, BFD session parameters use default values.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis bfd { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* | **frr-binding** } *

undo isis bfd { **min-rx-interval** [*receive-interval*] | **min-tx-interval** [*transmit-interval*] | **detect-multiplier** [*multiplier-value*] | **frr-binding** } *

Parameters

Parameter	Description	Value
min-rx-interval *receive-interval*	Specifies the minimum interval for receiving BFD packets from the peer end.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.

Parameter	Description	Value
min-tx-interval *transmit-interval*	Specifies the minimum interval for transmitting BFD packets to the peer end.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier-value*	Specifies the local detection multiplier.	The value is an integer that ranges from 3 to 50. The default value is 3.
frr-binding	Binds the BFD session status to IS-IS Auto FRR. When BFD detects the link fault on an interface, the BFD session goes Down, triggering FRR on the interface. After that, the traffic is switched from the faulty link to the backup link, which protects the traffic.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BFD can quickly detect faults on links. Configure proper parameters based on the actual network situation to improve the IS-IS convergence speed when network faults occur. You can use the **isis bfd** command to change BFD session parameters such as the minimum interval for sending BFD packets, minimum interval for receiving BFD packets, and local detection multiplier.

In an IS-IS process, after IS-IS establishes a BFD session, the value of *receive-interval* is obtained after the negotiation of the local **min-rx-interval** value and the remote **min-tx-interval** value. If no BFD packet is received from the peer end within the specified period (*receive-interval* x *multiplier-value*), the neighbor is considered Down.

Negotiation principle: Actual interval for the local device to receive BFD packets = MAX {local **min-rx-interval** value, remote **min-tx-interval** value}

Prerequisites

BFD has been enabled globally. In the interface view, IS-IS has been enabled and BFD has been enabled on the interface using the **isis bfd enable** command.

Precautions

The BFD priority of the interface is higher than that of the process. If BFD of the interface is enabled, the BFD session is set up based on the BFD parameters on the interface.

Example

Enable BFD on VLANIF100 and set the minimum receive interval to 600 ms and local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis bfd enable
[HUAWEI-Vlanif100] isis bfd min-rx-interval 600 detect-multiplier 4
```

Enable BFD on GE0/0/1 and set the minimum receive interval to 600 ms and local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis bfd enable
[HUAWEI-GigabitEthernet0/0/1] isis bfd min-rx-interval 600 detect-multiplier 4
```

7.6.56 isis bfd block

Function

The **isis bfd block** command prevents an IS-IS interface from dynamically establishing a BFD session.

The **undo isis bfd block** command restores the default setting.

By default, an interface can dynamically establish a BFD session.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis bfd block

undo isis bfd block

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BFD can provide millisecond-level fault detection, help IS-IS rapidly detect the faults that occur on neighboring devices or links, and instruct IS-IS to recalculate routes for correct packet forwarding. If the network contains unstable links that do not require high reliability and BFD has been enabled, a link cannot transmit data normally when it flaps. You can use the **isis bfd block** command to prevent specified interfaces from dynamically establishing BFD sessions.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

If the **isis bfd block**, **isis bfd enable**, and **isis bfd static** commands are executed, only the last command takes effect.

Example

Prevent VLANIF100 from dynamically establishing a BFD session.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis bfd block
```

Prevent GE0/0/1 from dynamically establishing a BFD session.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis bfd block
```


7.6.57 isis bfd enable

Function

The **isis bfd enable** command enables BFD on a specified IS-IS interface.

The **undo isis bfd enable** command disables BFD on a specified IS-IS interface.

By default, BFD is not enabled on an IS-IS interface.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis bfd enable

undo isis bfd enable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BFD can provide millisecond-level fault detection, help IS-IS to detect the faults that occur on neighboring devices or links more rapidly, and instruct IS-IS to recalculate routes for correct packet forwarding. The **isis bfd enable** command can be used to enable BFD on a specified IS-IS interface and establish BFD sessions by using default parameters.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

If global BFD is not enabled, you can configure BFD parameters on an interface but cannot establish a BFD session.

The BFD priority of the interface is higher than the BFD priority of the process. If BFD of the interface is enabled, the BFD session is set up based on the BFD parameters on the interface.

If the **isis bfd block**, **isis bfd enable**, and **isis bfd static** commands are run, only the last command takes effect.

Example

Enable BFD on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis bfd enable
```

Enable BFD on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis bfd enable
```

7.6.58 isis bfd static

Function

The **isis bfd static** command enables static BFD on a specified IS-IS interface.

The **undo isis bfd static** command disables static BFD on a specified IS-IS interface.

By default, static BFD is disabled on an IS-IS interface.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis bfd static

undo isis bfd static

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In static BFD, BFD session parameters are set manually and BFD session establishment requests are delivered manually. Then static BFD can quickly detect faults on links. You can run the **isis bfd static** command to enable static BFD on a specified interface to establish static BFD sessions on specified links.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

If the **isis bfd block**, **isis bfd enable**, and **isis bfd static** commands are run, only the last command takes effect.

Example

Enable static BFD on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis bfd static
```

Enable static BFD on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis bfd static
```

7.6.59 isis circuit-level

Function

The **isis circuit-level** command sets the link type of an interface on a Level-1-2 router.

The **undo isis circuit-level** command restores the default link type of an interface on a Level-1-2 router.

By default, the link type of an interface on a Level-1-2 router is Level-1-2, and both Level-1 and Level-2 neighbor relationships can be established on the interface.

Format

isis circuit-level [**level-1** | **level-1-2** | **level-2**]

undo isis circuit-level

Parameters

Parameter	Description	Value
level-1	Specifies the Level-1 link type. That is, only Level-1 neighbor relationship can be established on the interface.	-
level-1-2	Specifies the Level-2 link type. That is, both Level-1 and Level-2 neighbor relationships can be established on the interface.	-
level-2	Specifies the Level-2 link type. That is, only Level-2 neighbor relationship can be established on the interface.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When an adjacency is established between a Level-1-2 router and a remote device, the Level-1-2 router sends and receives both Level-1 and Level-2 Hello packets, wasting bandwidth and memory resources. To solve this problem, run the **isis circuit-level** command to set a specified link type for an interface.

Prerequisites

IS-IS has been enabled using the **isis enable** command in the interface view.

Precautions

Network flapping may occur if the link type of an IS-IS interface is changed during network operation. Therefore, setting a link type for an IS-IS interface on the switch when configuring IS-IS is recommended.

The configuration of the **isis circuit-level** command takes effect only when the IS-IS system type is Level-1-2. Otherwise, the level configured using the **is-level** command is used as the link type.

Example

If VLANIF100 is connected to a non-backbone router in the same area, set this interface to Level-1, and prohibit the interface from sending and receiving Level-2 Hello packets.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis circuit-level level-1
```


If GE0/0/1 is connected to a non-backbone router in the same area, set this interface to Level-1, and prohibit the interface from sending and receiving Level-2 Hello packets.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis circuit-level level-1
```

7.6.60 isis circuit-type

Function

The **isis circuit-type** command simulates the network type of an IS-IS broadcast interface to a P2P interface.

The **undo isis circuit-type** command restores the default network type of an IS-IS interface.

By default, the network type of an interface is determined by the physical type of the interface.

Format

isis circuit-type p2p [**strict-snpa-check**]

undo isis circuit-type

Parameters

Parameter	Description	Value
p2p	Sets the network type of an IS-IS interface to P2P.	-
strict-snpa-check	Enables IS-IS to check the SNPA address of each received LSP or SNP.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The network type of IS-IS interfaces at the ends of the link must be identical. Otherwise, the two interfaces cannot set up the neighbor relationship. In most cases, the network types of interfaces on a broadcast network and a P2P network are Ethernet and P2P respectively.

The **isis circuit-type** command sets the network type of an interface to simulate a P2P interface so that the network type of IS-IS interfaces at the ends of the link is identical and the neighbor relationship can be established between them.

When an IS-IS neighbor relationship is established between a P2P interface and a simulated P2P interface and the simulated P2P interface has direct neighbors, the P2P interface may receive unneeded packets from these direct neighbors. To prevent the P2P interface from accepting these unneeded packets, specify **strict-snpa-check** in the **isis circuit-type** command to enable IS-IS to check the SNPA address of each received LSP or SNP. After the command is run, the P2P interface accepts only the packets whose SNPA addresses are included in the local neighbor address list, which improves network security.

Precautions

For an interface enabled with IS-IS using the **isis enable** command, when the network type of an interface changes, the corresponding configurations change. Details are as follows:

- After a broadcast interface is simulated as a P2P interface using the **isis circuit-type** command, the interval for sending Hello packets, the number of Hello packets that IS-IS does not receive from a neighbor before the neighbor is declared Down, interval for resending LSP packets on a P2P link, and various IS-IS authentication modes are restored to the default settings; other configurations such as the DIS priority, DIS name, and interval for sending CSNP packets on a broadcast network become invalid.
- After the **undo isis circuit-type** command is run to restore the network type of an IS-IS interface, the interval for sending Hello packets, the number of Hello packets that IS-IS does not receive from a neighbor before the neighbor is declared Down, interval for resending LSP packets on a P2P link, various IS-IS authentication modes, DIS priority, and interval for sending CSNP packets on a broadcast network are restored to the default settings.

Example

Set the network type of VLANIF100 to P2P.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis circuit-type p2p
```

Set the network type of GE0/0/1 to P2P.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis circuit-type p2p
```


7.6.61 isis cost

Function

The **isis cost** command sets the link cost value of an IS-IS interface.

The **undo isis cost** command restores the default link cost value of an IS-IS interface.

By default, the link cost value of an IS-IS interface is 10.

Format

isis cost { *cost* | **maximum** } [**level-1** | **level-2**]

undo isis cost [*cost* | **maximum**] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
cost	Specifies the link cost value of an interface.	The value is an integer that varies according to the cost style. • When the cost style is **narrow**, **narrow-compatible**, or **compatible**, the value ranges from 1 to 63. • When the cost style is **wide** or **wide-compatible**, the value ranges from 1 to 16777214. The default value is 10.
maximum	Sets the link cost of IS-IS interfaces to 16777215. **NOTE** You can configure this parameter only when the IS-IS cost style is wide or wide-compatible. After the interface cost is set to 16777215, the neighbor TLV generated on the link can only be used to transmit TE information but cannot be used for route calculation.	-
level-1	Specifies the link cost value of a Level-1 interface. If the interface level is not specified, link cost values of Level-1 and Level-2 interfaces are set.	-

Parameter	Description	Value
level-2	Specifies the link cost value of a Level-2 interface. If the interface level is not specified, link cost values of Level-1 and Level-2 interfaces are set.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On large networks, there may be multiple valid routes to the same destination. IS-IS uses the SPF algorithm to calculate an optimal route and transmits traffic over it, which brings the following problems:

- All traffic is transmitted over the optimal route, causing load imbalance.
- If the optimal route is faulty, traffic will get lost.

To solve the preceding problems, run the **isis cost** command to set a link cost for interfaces so that traffic can be transmitted over different physical links.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Configuration Impact

If the link cost of an interface is changed, routes will be re-calculated on the whole network, causing the changes in traffic forwarding paths.

Precautions

The priority of the **circuit-cost** command is lower than that of the **isis cost** command.

Example

Set the Level-2 link cost of VLANIF100 to 5.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis cost 5 level-2
```

Set the Level-2 link cost of GE0/0/1 to 5.

```
<HUAWEI> system-view
[HUAWEI] isis
```



```
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis cost 5 level-2
```

7.6.62 isis delay-peer

Function

The **isis delay-peer** command configures a delay for an interface so that the interface delays establishing neighbor relationships after the neighbor relationship with the last neighbor goes Down due to packet timeout.

The **undo isis delay-peer** command deletes the configured delay and restores the default configurations.

By default, if the neighbor relationship goes Down due to packet timeout, the interface re-establishes the neighbor relationship after it receives a new Hello packet.

Format

isis delay-peer track last-peer-expired [**delay-time** *delay-interval*]

undo isis delay-peer [**track last-peer-expired** [**delay-time** *delay-interval*]]

Parameters

Parameter	Description	Value
track	Tracks the mode of the neighbor relationship establishment delay.	-
last-peer-expired	Indicates that neighbor relationship establishment is delayed after the neighbor relationship with the last neighbor goes Down due to packet timeout.	-
delay-time *delay-interval*	Specifies the neighbor relationship establishment delay.	The value is an integer ranging from 1 to 3600, in seconds. The default value is 60 seconds.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On IS-IS networks, devices at both ends of a link establish a neighbor relationship by exchanging Hello packets. After the neighbor relationship is established, both devices need to send Hello packets at a specified interval to maintain the neighbor relationship. If an IS-IS device does not receive any Hello packets from the neighbor at the other end within the specified period (Holddown time), the local device considers the neighbor Down and re-establishes the neighbor relationship after it receives a new Hello packet. If links are unstable or some Hello packets are lost or incorrect due to network transmission delay or poor transmission, neighbor relationships may alternate between Up and Down frequently, which causes a route flapping.

To address this issue, run the **isis delay-peer** command to configure a neighbor relationship establishment delay after the neighbor relationship goes Down.

Pre-configuration Tasks

IS-IS has been enabled using the **isis enable** command in the interface view.

Precautions

If a new *delay-interval* is configured and it is less than the remaining time of the ongoing delay, the new *delay-interval* takes effect immediately; if the new *delay-interval* is greater than the remaining time of the ongoing delay, the ongoing delay continues until the new *delay-interval* takes effect at the next delay.

Example

Set the delay to 100s on VLANIF100 so that the neighbor relationship establishment is delayed after the neighbor relationship with the last neighbor goes Down due to packet timeout.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis delay-peer track last-peer-expired delay-time 100
```

Set the delay to 100s on GE0/0/1 so that the neighbor relationship establishment is delayed after the neighbor relationship with the last neighbor goes Down due to packet timeout.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis delay-peer track last-peer-expired delay-time 100
```


7.6.63 isis dis-name

Function

The **isis dis-name** command configures a host name for the DIS.

The **undo isis dis-name** command deletes the host name configured for the DIS.

By default, no host name is configured for the DIS.

Format

isis dis-name *symbolic-name*

undo isis dis-name

Parameters

Parameter	Description	Value
symbolic-name	Specifies a host name for the DIS.	The value is a string of 1 to 64 characters without spaces. It is case sensitive. **NOTE** When double quotation marks are used around the string, spaces are allowed in the string.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

You can configure the name for the DIS on an interface only after you run the **isis enable** command to enable IS-IS on the interface. The DIS name is advertised in a pseudo-node LSP. In this manner, the configured DIS name is associated with the system ID of the DIS.

NOTE

This command takes effect only on the DIS in the broadcast network.

If the **isis circuit-type** command is run to emulate the interface as a P2P interface, the **isis dis-name** command becomes invalid on the interface; after the **undo isis circuit-type** command is run to restore the broadcast interface, the **isis dis-name** command becomes valid on the interface.

Example

Configure a host name for the DIS on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
```



```
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis dis-name LOCALAREA
```

Configure a host name for the DIS on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis dis-name LOCALAREA
```

7.6.64 isis dis-priority

Function

The **isis dis-priority** command sets the priority of the IS-IS interface that is a candidate for the DIS at a specified level.

The **undo isis dis-priority** command restores the default priority.

By default, the DIS priority of broadcast IS-IS interfaces at Level-1 and Level-2 is 64.

Format

isis dis-priority *priority* [**level-1** | **level-2**]

undo isis dis-priority [*priority*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
priority	Specifies the priority for DIS election.	The value is an integer that ranges from 0 to 127. The default value is 64. A larger value indicates a higher DIS priority.
level-1	Indicates the DIS priority of interfaces at Level 1. If the level is not specified, the same priority is set for Level-1 and Level-2 interfaces.	-
level-2	Indicates the DIS priority of interfaces at Level 2. If the level is not specified, the same priority is set for Level-1 and Level-2 interfaces.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The configuration is useful only for the broadcast network.

The DIS priority is advertised through Hello packets. The switch with the highest priority is elected as the DIS. In the case of the same priority, the switch with the largest MAC address is elected as the DIS.

NOTE

IS-IS has been enabled on the interface using the **isis enable** command.

If the **isis circuit-type** command is run to emulate the interface as a P2P interface, the **isis dis-priority** command becomes invalid on the interface; after the **undo isis circuit-type** command is run to restore the broadcast interface, the **isis dis-priority** command becomes valid on the interface.

Example

Set the Level-2 priority for electing the DIS to 127 on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis dis-priority 127 level-2
```

Set the Level-2 priority for electing the DIS to 127 on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis dis-priority 127 level-2
```

7.6.65 isis enable

Function

The **isis enable** command enables IS-IS on an interface and specifies the ID of the IS-IS process to be associated with the interface.

The **undo isis enable** command disables IS-IS on an interface and cancels the ID of the IS-IS process associated with the interface.

By default, IS-IS is disabled on an interface.

Format

isis enable [*process-id*]

undo isis enable

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535. The default value is 1.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After an IS-IS process is configured in the system view, to make the IS-IS protocol function normally, enable IS-IS on an interface to associate the interface with the IS-IS process.

Prerequisites

An IS-IS process has been enabled using the **isis** command in the system view.

Precautions

An interface can be associated with only one IS-IS process.

Example

Create IS-IS process 1 and activate the process on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] network-entity 10.0001.1010.1020.1030.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
```

Create IS-IS process 1 and activate the process on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] network-entity 10.0001.1010.1020.1030.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
```

7.6.66 isis lfa-backup

Function

The **isis lfa-backup** command enables an IS-IS interface to participate in loop-free alternate (LFA) calculation so that the interface can be a candidate for a backup interface.

The **undo isis lfa-backup** command disables an IS-IS interface from participating in LFA calculation.

By default, an IS-IS interface can participate in LFA calculation.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

isis lfa-backup [**level-1** | **level-2** | **level-1-2**]

undo isis lfa-backup [**level-1** | **level-2** | **level-1-2**]

Parameters

Parameter	Description	Value
level-1	Indicates that an interface can be a backup interface in Level-1 areas.	-
level-2	Indicates that an interface can be a backup interface in Level-2 areas.	-
level-1-2	Indicates that an interface can be a backup interface in both Level-1 and Level-2 areas.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

To facilitate traffic management and avoid the uncertainty in the traffic forwarding path in the case the primary link fails, run the **undo isis lfa-backup** command on some interfaces to disable them from participating in LFA calculation.

IS-IS has been enabled on the interface using the **isis enable** command.

Example

Disable VLANIF100 from becoming a backup interface for IS-IS Auto FRR.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] undo isis lfa-backup
```

Disable GE0/0/1 from becoming a backup interface for IS-IS Auto FRR.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] undo isis lfa-backup
```

7.6.67 isis lsp seq-overflow auto-recover disable

Function

The **isis lsp seq-overflow auto-recover disable** command prevents an IS-IS system from changing its system ID when it receives a locally generated LSP with the maximum sequence number (0xFFFFFFFF).

The **undo isis lsp seq-overflow auto-recover disable** command restores the default configuration.

By default, an IS-IS system changes its system ID when it receives a locally generated LSP with the maximum sequence number.

Format

isis lsp seq-overflow auto-recover disable

undo isis lsp seq-overflow auto-recover disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On an IS-IS network, if a device receives a locally generated LSP with the sequence number greater than that of the corresponding LSP stored locally, the device adds 1 to the sequence number of the received LSP and floods it. An attacker may send an IS-IS LSP with the maximum sequence number (0xFFFFFFFF) and the system ID of a target device. Upon receipt of the LSP, the target device considers it a locally generated LSP because it carries the local system ID and adds 1 to the sequence number because the sequence number is greater than that of the corresponding LSP stored locally. Consequently, the sequence number exceeds the maximum number, causing the target device to enter the hibernation state. The state can last 18 hours and 1 minute, affecting network operation. To prevent this problem, an IS-IS system changes its system ID when it receives a locally generated LSP with the sequence number of 0xFFFFFFFF. However, if the IS-IS system has changed its system ID for three times within 24 hours when it receives one more such an LSP, it directly enters the hibernation state.

The preceding function also applies to CSNPs and PSNPs.

By default, an IS-IS system changes its system ID when it receives a locally generated LSP with the maximum sequence number. To disable this function, run the **isis lsp seq-overflow auto-recover disable** command.

Example

Prevent an IS-IS system from changing its system ID when it receives a locally generated LSP with the maximum sequence number (0xFFFFFFFF).

```
<HUAWEI> system-view
[HUAWEI] isis lsp seq-overflow auto-recover disable
```

7.6.68 isis mesh-group

Function

The **isis mesh-group** command adds an IS-IS interface to a specified mesh group.

The **undo isis mesh-group** command deletes an IS-IS interface from a specified mesh group.

By default, an IS-IS interface does not belong to any mesh group and floods LSPs normally.

Format

isis mesh-group { *mesh-group-number* | **mesh-blocked** }

undo isis mesh-group

Parameters

Parameter	Description	Value
mesh-group-number	Specifies the mesh group number.	The value is an integer that ranges from 1 to 4294967295.

Parameter	Description	Value
mesh-blocked	Blocks an interface to prevent it from flooding received LSPs to other interfaces.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When receiving LSPs, the interface, not in the mesh group, floods the LSPs to other interfaces following the normal procedure. For the NBMA network that is with higher connectivity and several P2P links, this process causes repeated LSP flooding and wastes bandwidth.

After receiving LSPs, the interface that joins a mesh group only floods the LSPs to the interfaces that are not in the local mesh group.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

When adding interfaces to mesh groups or blocking interfaces, keep certain interfaces from being configured with the isis mesh-group command. This can prevent link faults from affecting the normal spreading of LSPs.

Example

Add VLANIF100 to mesh group 3.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis mesh-group 3
```

Add GE0/0/1 to mesh group 3.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis mesh-group 3
```


7.6.69 isis padding-hello

Function

The **isis padding-hello** command configures an IS-IS interface to send Hello packets with the padding field.

The **undo isis padding-hello** command restores the default setting.

By default, an IS-IS interface is not configured to send Hello packets with the padding field.

Format

isis padding-hello

undo isis padding-hello

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The **isis padding-hello** and **isis small-hello** commands are mutually exclusive and cannot be configured on the same interface simultaneously.

If an interface that is not configured with the two commands, it sends Hello packets based on the following rules:

- For a P2P interface
  - Before the P2P neighbor relationship is established, the P2P interface sends Hello packets with the padding field.
  - After the P2P neighbor relationship is established, the P2P interface sends Hello packets without the padding field.

  NOTE

  For a P2P interface, the length of the padding field is equal to the length of LSP packets generated by the local IS.
- For a broadcast interface

  It sends Hello packets with the padding field.

  NOTE

  For a broadcast interface, the length of padding field is equal to the length of MTU.

IS-IS has been enabled on the interface using the **isis enable** command.

Example

Configure VLANIF100 to send Hello packets with the padding field.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis padding-hello
```

Configure GE0/0/1 to send Hello packets with the padding field.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis padding-hello
```

7.6.70 isis peer-ip-ignore

Function

The **isis peer-ip-ignore** command configures IS-IS not to check the IP address of Hello packets received by an interface.

The **undo isis peer-ip-ignore** command restores the default setting.

By default, IS-IS checks the IP address of received Hello packets.

Format

isis peer-ip-ignore

undo isis peer-ip-ignore

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, IS-IS checks the IP address carried in the Hello packet. The neighbor relationship can be set up only when the local interface address and the interface

address carried in the received packets belong to the same network segment. When the two interface IP addresses are not in the same network segment, if the **isis peer-ip-ignore** command is used, the check on the peer IP address is ignored, and the two IS-IS interfaces can set up normal neighbor relationship. Routes of this two network segments exist in the routing table, but cannot ping through each other.

Precautions

This command is valid for P2P interfaces, NBMA interfaces and the interface with the network type configured to P2P by using the **isis circuit-type p2p** command. The command takes effect only when the command is used on the two ends of the link.

NOTE

- For the broadcast interface, you can run the **isis peer-ip-ignore** command only after you run the **isis circuit-type p2p** command. The **isis circuit-type p2p** command is valid only for the broadcast interface.
- For the P2P and NBMA interfaces, you can run the **isis peer-ip-ignore** command without running the **isis circuit-type p2p** command.

Example

Configure VLANIF100 not to check the IP address of the Hello packets sent by the peer.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis circuit-type p2p
[HUAWEI-Vlanif100] isis peer-ip-ignore
```

Configure GE0/0/1 not to check the IP address of the Hello packets sent by the peer.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis circuit-type p2p
[HUAWEI-GigabitEthernet0/0/1] isis peer-ip-ignore
```

7.6.71 isis ppp-negotiation

Function

The **isis ppp-negotiation** command specifies the PPP negotiation mode for establishing neighbor relationships.

The **undo isis ppp-negotiation** command restores the default negotiation mode.

By default, the 3-way handshake mode is used.

Format

isis ppp-negotiation { **2-way** | **3-way** [**only**] }

undo isis ppp-negotiation

Parameters

Parameter	Description	Value
2-way	Establishes the neighbor relationship using the 2-way handshake negotiation mode.	-
3-way	Establishes the neighbor relationship using the 3-way handshake negotiation mode.	-
only	Establishes the neighbor relationship using the 3-way handshake negotiation mode that is not backward compatible.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The 3-way handshake negotiation mode is backward compatible. If the neighbor only supports 2-way handshake, use the 2-way handshake negotiation mode to establish a neighbor relationship.

The command is applicable only for P2P interfaces. It can be used on the broadcast interfaces only after the circuit type is set to P2P.

Example

Establish the neighbor relationship on VLANIF100 using the 2-way handshake negotiation mode.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis circuit-type p2p
[HUAWEI-Vlanif100] isis ppp-negotiation 2-way
```

Establish the neighbor relationship on GE0/0/1 using the 2-way handshake negotiation mode.

```
<HUAWEI> system-view
[HUAWEI] isis
```



```
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis circuit-type p2p
[HUAWEI-GigabitEthernet0/0/1] isis ppp-negotiation 2-way
```

7.6.72 isis ppp-osicp-check

Function

The **isis ppp-osicp-check** command enables OSICP negotiation check on a PPP interface. The negotiation status can affect IS-IS interface status.

The **undo isis ppp-osicp-check** command restores the default setting.

By default, IS-IS does not check OSICP status of PPP.

Format

isis ppp-osicp-check

undo isis ppp-osicp-check

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, OSCIP status of PPP does not affect IS-IS interface status.

After this command is configured, OSI network negotiation status of PPP can affect IS-IS interface status. When PPP senses that the OSI network fails, the link status of the IS-IS interface turns Down and the route to the network segment where the interface resides is not advertised through LSP.

Precautions

This command applies to only PPP interfaces. For other point-to-point interfaces, this command is invalid.

Example

Configure VLANIF100 to check OSICP status of PPP.

```
<HUAWEI> system-view
[HUAWEI] isis
```



```
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis circuit-type p2p
[HUAWEI-Vlanif100] isis ppp-osicp-check
```

Configure GE0/0/1 to check OSICP status of PPP.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis circuit-type p2p
[HUAWEI-GigabitEthernet0/0/1] isis ppp-osicp-check
```

7.6.73 isis purge-lsp auto-protect disable

Function

The **isis purge-lsp auto-protect disable** command disables IS-IS purge LSPs from triggering master/slave main control board switchovers.

The **undo isis purge-lsp auto-protect disable** command restores the default configuration.

By default, IS-IS purge LSPs trigger master/slave main control board switchovers.

Format

isis purge-lsp auto-protect disable

undo isis purge-lsp auto-protect disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

When an IS-IS process on a device proactively sends a purge LSP, the device deletes the corresponding LSP and floods it to the network. In most cases, before the device sends a purge LSP, the end that generated the corresponding LSP sends an updated LSP. If the clock on the device runs fast, the device frequently floods purge LSPs to devices on the entire network, causing network flapping. If the device generates more than five purge LSPs for 80% or more non-pseudonode LSPs with a non-zero fragment number in the local LSDB within 6500s, a master/

slave main control board switchover is performed if the device has two main control boards, or the device is restarted if it has only one main control board. The switchover or restart prevents network flapping.

By default, IS-IS purge LSPs trigger master/slave main control board switchovers. To disable IS-IS purge LSPs from triggering master/slave main control board switchovers, run the **isis purge-lsp auto-protect disable** command.

Example

Disable IS-IS purge LSPs from triggering master/slave main control board switchovers.

```
<HUAWEI> system-view
[HUAWEI] isis purge-lsp auto-protect disable
```

7.6.74 isis silent

Function

The **isis silent** command configures an IS-IS interface as a silent interface. That is, the interface is suppressed from sending and receiving IS-IS packets, but routes of the network segment on which the interface resides can be advertised.

The **undo isis silent** command restores the default setting.

By default, no IS-IS interface is configured as a silent interface.

Format

isis silent [**advertise-zero-cost**]

undo isis silent

Parameters

Parameter	Description	Value
advertise-zero-cost	Specifies the route cost as 0. The default cost of IS-IS routes is 10.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When an IS-IS network is connected to other ASs, it is required to enable IS-IS on the interfaces that connect the IS-IS network to other ASs so that the switch on the IS-IS network can learn the routes to other ASs. This interface, however, unnecessarily advertises IS-IS Hello packets on its network segment. In this case, you can run the **isis silent** command to suppress the IS-IS interface.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Example

Configure VLANIF100 as a silent interface.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis silent
```

Configure GE0/0/1 as a silent interface.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis silent
```

7.6.75 isis small-hello

Function

The **isis small-hello** command configures an IS-IS interface to send Hello packets without the padding field.

The **undo isis small-hello** command restores the default setting.

By default, an IS-IS interface is not configured to send Hello packets without the padding field.

Format

isis small-hello

undo isis small-hello

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The command can simplify the operations of sending and receiving Hello packets.

The **isis small-hello** and **isis padding-hello** commands are mutually exclusive and cannot be configured on the same interface simultaneously.

If an interface is not configured with the two commands, it sends Hello packets based on the following rules:

- P2P interface
  - Before the P2P neighbor relationship is established, the P2P interface sends Hello packets with the padding field.
  - After the P2P neighbor relationship is established, the P2P interface sends Hello packets without the padding field.

  NOTE

  For a P2P interface, the length of the padding field is equal to the length of LSP packets generated by the local IS.
- Broadcast interface

  The interface sends Hello packets with the padding field.

  NOTE

  For a broadcast interface, the length of padding field is equal to the length of MTU.

Before running this command on an interface, run the **isis enable** command on the interface to enable IS-IS.

Example

Configure VLANIF100 to send Hello packets without the padding field.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis small-hello
```

Configure GE0/0/1 to send Hello packets without the padding field.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis small-hello
```


7.6.76 isis suppress-flapping peer

Function

The **isis suppress-flapping peer** command configures detection parameters for IS-IS neighbor relationship flapping suppression.

The **undo isis suppress-flapping peer** command restores the default detection parameters.

By default, the detection interval of IS-IS neighbor relationship flapping suppression is 60s, the suppression threshold is 10, and the interval for exiting from suppression is 120s.

Format

isis suppress-flapping peer { **detecting-interval** *detecting-interval* | **threshold** *threshold* | **resume-interval** *resume-interval* } *

undo isis suppress-flapping peer { **detecting-interval** [*detecting-interval*] | **threshold** [*threshold*] | **resume-interval** [*resume-interval*] } *

Parameters

Parameter	Description	Value
detecting-interval *detecting-interval*	Specifies the detection interval of IS-IS neighbor relationship flapping suppression. Each IS-IS interface on which IS-IS neighbor relationship flapping suppression is enabled starts a flapping counter. If the interval between two successive neighbor status changes from Full to a non-Full state is shorter than *detecting-interval*, a valid flapping_event is recorded, and the flapping_count is incremented by 1.	The value is an integer ranging from 1 to 300, in seconds. The default value is 60s.
threshold *threshold*	Specifies the threshold of IS-IS neighbor relationship flapping suppression. When the flapping-count reaches or exceeds *threshold*, flapping suppression takes effect.	The value is an integer ranging from 1 to 1000. The default value is 10.

Parameter	Description	Value
resume-interval *resume-interval*	• Specifies the interval for exiting from IS-IS neighbor relationship flapping suppression. If the interval between two successive neighbor status changes from Full to a non-Full state is longer than *resume-interval*, the flapping-count is reset. • If IS-IS neighbor relationship flapping suppression works in hold-max-cost mode, *resume-interval* indicates the duration of this mode. **NOTE** The value of *resume-interval* must be greater than that of *detecting-interval*.	The value is an integer ranging from 2 to 1000, in seconds. The default value is 120s.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To configure detection parameters for IS-IS neighbor relationship flapping suppression on an interface, run the **isis suppress-flapping peer** command. However, keeping the default configurations is recommended.

Prerequisites

IS-IS neighbor relationship flapping suppression must have been enabled globally before you configure detection parameters for it. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it before you configure the detection parameters.

Example

Set the detection interval of IS-IS neighbor relationship flapping suppression to 5s, the suppression threshold to 40, and the interval for exiting from suppression to 20s on VLANIF 100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis suppress-flapping peer detecting-interval 5 threshold 40 resume-interval 20
```


7.6.77 isis suppress-flapping peer disable

Function

The **isis suppress-flapping peer disable** command disables IS-IS neighbor relationship flapping suppression from an interface.

The **undo isis suppress-flapping peer disable** command enables IS-IS neighbor relationship flapping suppression on an interface.

By default, IS-IS neighbor relationship flapping suppression is enabled on all interfaces.

Format

isis suppress-flapping peer disable

undo isis suppress-flapping peer disable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, IS-IS neighbor relationship flapping suppression is enabled on all interfaces in the same IS-IS process. To disable the function from one of the interfaces, run the **isis suppress-flapping peer disable** command.

NOTE

When an interface enters the flapping suppression state, all neighbor relationships on the interface enter the state accordingly.

Prerequisites

IS-IS neighbor relationship flapping suppression must have been enabled globally before you enable the function on an interface using the **undo isis suppress-flapping peer disable** command. By default, the function is enabled globally. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it first.

Example

Disable IS-IS neighbor relationship flapping suppression from VLANIF100.


```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis suppress-flapping peer disable
```

7.6.78 isis suppress-flapping peer hold-down

Function

The **isis suppress-flapping peer hold-down** command configures the Hold-down mode and sets duration for this mode.

The **undo isis suppress-flapping peer hold-down** command cancels the Hold-down mode.

By default, the Hold-down mode is disabled.

Format

isis suppress-flapping peer hold-down *interval*

undo isis suppress-flapping peer hold-down [*interval*]

Parameters

Parameter	Description	Value
interval	Specifies the duration of the Hold-down mode.	The value is an integer ranging from 1 to 600, in seconds.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Flapping suppression works in either Hold-down or Hold-max-cost mode.

- Hold-down mode: In the case of frequent flooding and topology changes during neighbor relationship establishment, interfaces prevent neighbor relationship reestablishment during Hold-down suppression, which minimizes LSDB synchronization attempts and packet exchanges.
- Hold-max-cost mode: If the traffic forwarding path changes frequently, interfaces use the maximum value (16777214 for the wide mode and 63 for the narrow mode) as the cost of the flapping link during Hold-max-cost suppression, which prevents traffic from passing through the flapping link.

Flapping suppression can also work first in Hold-down mode and then in Hold-max-cost mode.

By default, the Hold-max-cost mode takes effect. To configure the Hold-down mode and set duration for this mode, run the **isis suppress-flapping peer hold-down** *interval* command.

Prerequisites

IS-IS neighbor relationship flapping suppression must have been enabled globally before you configure the Hold-down mode and set duration for this mode. By default, the function is enabled. If it is disabled, run the **undo suppress-flapping peer disable** command to enable it before you configure the Hold-down mode and set duration for this mode.

Example

Configure the Hold-down mode and set its duration to 200s on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis suppress-flapping peer hold-down 200
```

7.6.79 isis suppress-flapping peer hold-max-cost disable

Function

The **isis suppress-flapping peer hold-max-cost disable** command disables the Hold-max-cost mode.

The **undo isis suppress-flapping peer hold-max-cost disable** command enables the Hold-max-cost mode.

By default, the Hold-max-cost mode is enabled.

Format

isis suppress-flapping peer hold-max-cost disable

undo isis suppress-flapping peer hold-max-cost disable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Flapping suppression works in either Hold-down or Hold-max-cost mode.

- Hold-down mode: In the case of frequent flooding and topology changes during neighbor relationship establishment, interfaces prevent neighbor relationship reestablishment during Hold-down suppression, which minimizes synchronization attempts and packet exchanges.
- Hold-max-cost mode: If the traffic forwarding path changes frequently, interfaces use the maximum value (16777214 for the wide mode and 63 for the narrow mode) as the cost of the flapping link during Hold-max-cost suppression, which prevents traffic from passing through the flapping link. If a device on a key path is isolated from the network due to IS-IS neighbor relationship flapping, the network is separated into two isolated parts. To prevent this problem, use the Hold-max-cost mode on the key path.

If the neighbor relationship does not go Down within successive **resume-interval**s, or the interval between two successive neighbor Down events is greater than or equal to **resume-interval**, the suppression in Hold-max-cost mode exits.

Flapping suppression can also work first in Hold-down mode and then in Hold-max-cost mode.

By default, the Hold-max-cost mode takes effect. To configure the Hold-down mode and set duration for this mode, run the **isis suppress-flapping peer hold-down** *interval* command.

Precautions

The Hold-max-cost mode takes effect only unidirectionally. If a remote device does not support IS-IS neighbor relationship flapping suppression, bidirectional traffic between the local and remote devices may travel along different paths.

Example

Disable the Hold-max-cost mode on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis suppress-flapping peer hold-max-cost disable
```

7.6.80 isis suppress-reachability

Function

The **isis suppress-reachability** command suppresses the advertisement of direct routes on an IS-IS interface in a specified topology.

The **undo isis suppress-reachability** command restores the default setting.

By default, direct routes on an IS-IS interface are advertised.

Format

isis suppress-reachability [**level-1** | **level-1-2** | **level-2**]

undo isis suppress-reachability

Parameters

Parameter	Description	Value
level-1	Indicates that the advertisement of IPv4 addresses is suppressed on Level-1 interfaces. If the level is not specified, the advertisement of IPv6 addresses on Level-1 and Level-2 interfaces is suppressed.	-
level-1-2	Indicates that the advertisement of IPv4 addresses is suppressed on Level-1 and Level-2 interfaces.	-
level-2	Indicates that the advertisement of IPv4 addresses is suppressed on Level-2 interfaces. If the level is not specified, the advertisement of IPv4 addresses is suppressed on Level-1 and Level-2 interfaces.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Before running this command on an interface, run the **isis enable** command on the interface to enable IS-IS.

Example

Suppress the advertisement of IPv4 addresses on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis suppress-reachability
```

Suppress the advertisement of IPv4 addresses on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis suppress-reachability
```


7.6.81 isis system-id auto-recover disable

Function

The **isis system-id auto-recover disable** command disables the system from automatically resolving IS-IS system ID conflicts.

The **undo isis system-id auto-recover disable** command enables the system to automatically resolve IS-IS system ID conflicts.

By default, if the system detects an IS-IS system ID conflict, it automatically changes the local system ID to resolve the conflict.

Format

isis system-id auto-recover disable

undo isis system-id auto-recover disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A system ID uniquely identifies an IS-IS device. If the same system ID is configured for more than one device on the network, a routing loop may occur. By default, if the system detects an IS-IS system ID conflict, it automatically changes the local system ID to resolve the conflict.

To disable the system from automatically resolving IS-IS system ID conflicts, run the **isis system-id auto-recover disable** command. After the command is run, IS-IS system ID conflicts need to be manually resolved.

NOTE

The first two bytes of the system ID automatically changed by the system are Fs, and the last four bytes are randomly generated. For example, FFFF:1234:5678 is such a system ID.

Precautions

If an IS-IS system ID conflict occurs between two directly connected devices, a neighbor relationship fails to be established only between the two devices, without affecting the entire network. As a result, the conflict is not automatically resolved in this case.

On broadcast networks, the system ID generated automatically is not recorded in the configuration file. If the device is restarted, the system restores this system ID to the originally configured one and then generates another one, which may be different from the one last generated automatically. If the conflict persists after the system automatically generates three system IDs, the system no longer resolves this conflict.

Example

Disable the system to automatically resolve IS-IS system ID conflicts.

```
<HUAWEI> system-view
[HUAWEI] isis system-id auto-recover disable
```

7.6.82 isis tag-value

Function

The **isis tag-value** command sets the administrative tag value of an IS-IS interface.

The **undo isis tag-value** command deletes the administrative tag value on an IS-IS interface.

By default, an IS-IS interface has no administrative tag value.

Format

isis tag-value *tag* [**level-1** | **level-2**]

undo isis tag-value [*tag*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
tag	Specifies the administrative tag of an IS-IS interface.	The value is an integer that ranges from 1 to 4294967295.
level-1	Indicates the administrative tag value of a Level-1 interface. If the interface level is not specified, the administrative tag value is set for Level-1 and Level-2 interfaces.	-
level-2	Indicates the administrative tag value of a Level-2 interface. If the interface level is not specified, the administrative tag value is set for Level-1 and Level-2 interfaces.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An administrative tag carries administrative information about IP address prefixes. The tag can be used to import routes of different levels and different areas. Administrative tags can advertise IP address prefixes in an IS-IS routing domain to control routes and simplify management.

Using the **isis tag-value** command, you can set the administrative tag value for all routes of a specified IS-IS process. The tag can be used as a filtering condition of a route-policy to filter routes.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

The advertised LSPs contain the administrative tag value only when the IS-IS cost style is wide, wide-compatible, or compatible.

The administrative tag value set using the **isis tag-value** command has a higher priority than the administrative tag value set using the **circuit default-tag** command.

Example

Set the administrative tag value of VLANIF100 to 77.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis tag-value 77
```

Set the administrative tag value of GE0/0/1 to 77.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis tag-value 77
```

7.6.83 isis timer csnp

Function

The **isis timer csnp** command sets the interval for sending CSNPs on a broadcast network.

The **undo isis timer csnp** command restores the default setting.

By default, the interval for sending CSNPs on a broadcast network is 10 seconds.

Format

isis timer csnp *csnp-interval* [**level-1** | **level-2**]

undo isis timer csnp [*csnp-interval*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
csnp-interval	Specifies the interval for sending CSNPs on a broadcast network.	The value is an integer that ranges from 1 to 65535, in seconds. The default value is 10 seconds.
level-1	Indicates the interval for sending Level-1 CSNPs. If no level is specified, the interval for the IS-IS process of the current level to send CSNP packets is set by default.	-
level-2	Indicates the interval for sending Level-2 CSNPs. If no level is specified, the interval for the IS-IS process of the current level to send CSNP packets is set by default.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a broadcast network, a DIS sends CSNPs periodically to enable all devices to synchronize LSDBs with one another. If a device finds that the local LSDB does not have a specific LSP or an existing LSPs is not the latest one after the device has received a CSNP, the device will send a PSNP to request the corresponding LSP. Only a DIS sends CSNPs periodically. Therefore, the **isis timer csnp** command will take effect only on a broadcast interface of the DIS. This command can be used to set an interval for sending CSNPs in an area at a specified level. A router may be elected as a DIS in both Level-1 and Level-2 areas. Therefore, you can set different intervals at which the DIS sends CSNPs in Level-1 and Level-2 areas.

Precautions

The IS-IS route convergence speed depends on the LSDB synchronization speed. Therefore, reducing the interval for sending CSNPs can speed up LSDB

synchronization and IS-IS route convergence. If the interval is set too small, however, the DIS will send CSNPs frequently. This causes high CPU, memory, and network bandwidth usage and affects services.

NOTE

If the **isis circuit-type** command is run to emulate the interface as a P2P interface, the **isis timer csnp** command becomes invalid on the interface; after the **undo isis circuit-type** command is run to restore the broadcast interface, the interval for sending CSNPs is restored to the default setting.

Example

Set the interval for sending Level-2 CSNPs to 15 seconds on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis timer csnp 15 level-2
```

Set the interval for sending Level-2 CSNPs to 15 seconds on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis timer csnp 15 level-2
```

7.6.84 isis timer hello

Function

The **isis timer hello** command sets the interval for sending Hello packets on an IS-IS interface.

The **undo isis timer hello** command restores the default setting.

By default, the interval for sending Hello packets 10 seconds on an interface.

Format

isis timer hello *hello-interval* [**level-1** | **level-2**] [**conservative**]

undo isis timer hello [*hello-interval*] [**level-1** | **level-2**] [**conservative**]

Parameters

Parameter	Description	Value
hello-interval	Specifies the interval for sending Hello packets.	The value is an integer that ranges from 3 to 255, in seconds. The default value is 10 seconds.
level-1	Indicates the interval for sending Level-1 Hello packets. If the level is not specified, the interval for sending Level-1 and Level-2 Hello packets is set by default.	-
level-2	Indicates the interval for sending Level-2 Hello packets. If the level is not specified, the interval for sending Level-1 and Level-2 Hello packets is set by default. **NOTE** Parameters **level-1** and **level-2** are configured only on a broadcast interface that is enabled with IS-IS. Level-1 and Level-2 Hello packets are sent separately and their intervals must be set respectively. There is only one Hello packet on a point-to-point link. Therefore, level-1 and level-2 parameters are not used.	-
conservative	Indicates the conservative mode of the dead timer. If the conservative mode is configured, the value configured for the dead timer takes effect even when the value is less than 10s.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS maintains neighbor relationships between neighbors by sending and receiving Hello packets. If the local device does not receive Hello packets from its neighbor within a specified period, the device considers the neighbor Down.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

The shorter the interval, the more system resources used to send Hello packets. The interval should therefore be set according to the actual conditions.

If a broadcast interface is emulated as a P2P interface through the **isis circuit-type** command or then restored to the broadcast interface through the **undo isis circuit-type** command, the interval for sending IS-IS Hello packets is restored to the default value.

Example

Set the interval for sending Level-2 Hello packets to 20 seconds on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis timer hello 20 level-2
```

Set the interval for sending Level-2 Hello packets to 20 seconds on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis timer hello 20 level-2
```

7.6.85 isis timer holding-multiplier

Function

The **isis timer holding-multiplier** command sets the multiplier of the interval for sending Hello packets to change the holdtime of IS-IS neighbor relationship.

The **undo isis timer holding-multiplier** command restores the default setting.

By default, the multiplier of the interval for sending Hello packets is 3, that is, the neighbor holdtime is three times the interval for sending Hello packets.

Format

isis timer holding-multiplier *number* [**level-1** | **level-2**]

undo isis timer holding-multiplier [*number*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
number	Indicates that the neighbor holdtime is a multiplier of the interval for sending Hello packets.	The value is an integer that ranges from 3 to 1000. The default value is 3.
level-1	Indicates the holdtime of Level-1 neighbors. If the level is not specified, the holdtime is set for both Level-1 and Level-2 neighbors.	-
level-2	Indicates the holdtime of Level-2 neighbors. If the level is not specified, the holdtime is set for both Level-1 and Level-2 neighbors. **NOTE** Parameters **level-1** and **level-2** are configured only on a broadcast interface that is enabled with IS-IS. Level-1 and Level-2 Hello packets are sent separately and their intervals must be set respectively. There is only one Hello packet on a point-to-point link. Therefore, level-1 and level-2 parameters are not used.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Devices at both ends of a link establish a neighbor relationship by sending Hello packets to each other. After the neighbor relationship is established, both devices need to send Hello packets at a specified interval to maintain the neighbor relationship. If a device does not receive any Hello packet from its neighbor within a specified period of time, the device considers the neighbor to be Down. The specified time period is known as the neighbor holdtime.

For example, run the **isis timer hello** command on a local device to set the interval for sending Hello packets to 20s. Then, run the **isis timer holding-multiplier** 4 command. The holdtime is 80s (four times the interval for sending Hello packets). If the interval for sending Hello packets is changed using the **isis timer hello** 20 command, the holdtime will be changed accordingly.

Prerequisites

IS-IS has been enabled on the interface using the **isis enable** command.

Precautions

If the *number* value is set too large, the local device needs to wait a long time before detecting that the remote device has gone Down. This slows down IS-IS route convergence. If the value of *number* is set too small, the neighbor relationship will alternate between **Up** and **Down** when some Hello packets are lost due to transmission delays and errors on the network. This causes route flapping on the IS-IS network. Therefore, exercise caution when setting the value of *number*. Set the same interval for sending Hello packets and the same neighbor holdtime for all devices on the IS-IS network is recommended. This is to ensure that all devices detect link failures at the same time and guarantee timely IS-IS route convergence.

If a broadcast interface is emulated as a P2P interface through the **isis circuit-type** command or then restored to the broadcast interface through the **undo isis circuit-type** command, the number of Hello packets that IS-IS does not receive from a neighbor before the neighbor is declared Down is restored to the default value.

Example

Set the number of Level-2 Hello packets that IS-IS does not receive from a neighbor before the neighbor is declared Down to 6 on VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis timer holding-multiplier 6 level-2
```

Set the number of Level-2 Hello packets that IS-IS does not receive from a neighbor before the neighbor is declared Down to 6 on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis timer holding-multiplier 6 level-2
```

7.6.86 isis timer lsp-retransmit

Function

The **isis timer lsp-retransmit** command sets the interval for retransmitting LSPs over a P2P link.

The **undo isis timer lsp-retransmit** command restores the default value.

By default, the interval for retransmitting LSPs over a P2P link is 5 seconds.

Format

isis timer lsp-retransmit *retransmit-interval*

undo isis timer lsp-retransmit

Parameters

Parameter	Description	Value
retransmit-interval	Specifies the interval for retransmitting LSPs.	The value is an integer that ranges from 1 to 300.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a point-to-point network, devices at both ends of a link synchronize LSDBs with each other by flooding LSPs. The device at one end of the link sends an LSP. If the device at the other end receives this LSP, it replies with a PSNP. If the device that has sent an LSP does not receive a PSNP from the other end in a period of time, the device will retransmit the LSP.

The **isis timer lsp-retransmit** command is used to set an interval for retransmitting LSPs. Only the devices on a point-to-point network send PSNPs. Therefore, the **isis timer lsp-retransmit** command will take effect only when it is run on point-to-point interfaces.

Precautions

After the **isis timer lsp-retransmit** command is run on a device, the device will wait *retransmit-interval* after having sent an LSP. If the device receives a PSNP from the other end, the device will not retransmit the LSP. Otherwise, the device will retransmit the LSP.

If the value of *retransmit-interval* is set too small, an LSP will be retransmitted even though it is not necessary, causing high CPU, memory, and network bandwidth usage.

NOTE

If a broadcast interface is emulated as a P2P interface through the **isis circuit-type** command or restored a P2P interface to the broadcast interface through the **undo isis circuit-type** command, the interval for retransmitting LSP packets on a P2P link is restored to the default value.

Example

Set the interval for retransmitting LSPs on VLANIF100 to 10 seconds.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
```



```
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis circuit-type p2p
[HUAWEI-Vlanif100] isis timer lsp-retransmit 10
```

Set the interval for retransmitting LSPs on GE0/0/1 to 10 seconds.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis circuit-type p2p
[HUAWEI-GigabitEthernet0/0/1] isis timer lsp-retransmit 10
```

7.6.87 isis timer lsp-throttle

Function

The **isis timer lsp-throttle** command sets the minimum interval for sending LSPs on an IS-IS interface and the maximum number of LSPs sent within the interval.

The **undo isis timer lsp-throttle** command restores the default setting.

By default, the minimum interval for sending LSPs on an IS-IS interface is 50 milliseconds and the maximum number of LSPs sent within the interval is 10.

Format

isis timer lsp-throttle *throttle-interval* [**count** *count*]

undo isis timer lsp-throttle

Parameters

Parameter	Description	Value
throttle-interval	Specifies the minimum interval for sending LSPs.	The value is an integer that ranges from 1 to 10000, in milliseconds.
count *count*	Specifies the maximum number of LSPs that are sent within the interval specified by *throttle-interval*.	The value is an integer that ranges from 1 to 1000.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

The value specified by *throttle-interval* is the interval between two consecutive LSPs and is also the interval for sending multiple fragments of a CSNP.

Before running this command on an interface, run the **isis enable** command on the interface to enable IS-IS.

Example

Set the interval for retransmitting LSPs on VLANIF100 to 500 ms.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] isis enable 1
[HUAWEI-Vlanif100] isis timer lsp-throttle 500
```

Set the interval for retransmitting LSPs on GE0/0/1 to 500 ms.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 01.0000.0000.0001.00
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] isis enable 1
[HUAWEI-GigabitEthernet0/0/1] isis timer lsp-throttle 500
```

7.6.88 is-level

Function

The **is-level** command sets the level of an IS-IS switch.

The **undo is-level** command restores the default setting.

By default, the level of an IS-IS switch is Level-1-2.

Format

is-level { **level-1** | **level-1-2** | **level-2** }

undo is-level

Parameters

Parameter	Description	Value
level-1	Indicates that the level of the switch is Level-1. The switch calculates only intra-area routes and maintains the Level-1 LSDB.	-

Parameter	Description	Value
level-1-2	Indicates that the level of the switch is Level-1-2. The switch calculates Level-1 and Level-2 routes and maintains Level-1 and Level-2 LSDBs.	-
level-2	Indicates that the level of the switch is Level-2. The switch calculates only Level-2 routes, and maintains the Level-2 LSDB.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To support a network with a large number of routes, IS-IS uses a two-level structure in an IS-IS routing domain. IS-IS switches are classified into the following levels:

- Level-1 switch: A Level-1 switch manages intra-area routing. It establishes neighbor relationships with only Level-1 and Level-1-2 switches in the same area. Level-1 switches can be connected to other areas through Level-1-2 switches only.
- Level-2 switch: A Level-2 switch manages intra-area routing. It establishes neighbor relationships with Level-2 switches in the same area and Level-1-2 switches in other areas only. All Level-2 switches form the backbone network of the routing domain. They are responsible for communications between areas. The Level-2 switches in the routing domain must be reachable, and no switch of other levels is deployed between every two Level-2 switches.
- Level-1-2 switch: A Level-1-2 switch can establish Level-1 neighbor relationships with Level-1 switches and Level-1-2 switches in the same area. It can also establish Level-2 neighbor relationships with Level-2 switches and Level-1-2 switches in other areas.

In most cases, Level-1 switches are located within an area, Level-2 switches are located between areas, and Level-1-2 switches are located between Level-1 switches and Level-2 switches.

The level of an IS-IS switch and of an interface determine the level of a neighbor relationship. By default, neighbor relationships between two Level-1-2 switches are Level-1 and Level-2. To establish a Level-1 or Level-2 neighbor relationship, run the **isis circuit-level** command to modify the level of interfaces.

If only one area exists, setting the level of switches to Level-1 or Level-2 is recommended to prevent devices from maintaining two LSDBs that are the same. On an IP network, setting the level of all switches to Level-2 for future extension is recommended.

Precautions

- If the levels of IS-IS switches are changed during network operation, the IS-IS process will be restarted and IS-IS neighbor relationships will be disconnected. Setting the levels of switches when configuring IS-IS is recommended.
- If the **is-level** command is not configured for a switch, the switch works at Level-1-2. That is, the switch calculates Level-1 and Level-2 routes and maintains Level-1 and Level-2 LSDBs simultaneously.
- When both Level-1 and Level-2 switches have routes destined for the same destination address, the route on the Level-1 switch is preferred.

Example

Set the level of the current switch to Level-1.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] is-level level-1
```

7.6.89 is-name

Function

The **is-name** command enables the capability of identifying the host name in an LSP and configures the dynamic host name for the IS-IS system of the local switch. The dynamic hostname is advertised in an LSP packet.

The **undo is-name** command deletes the dynamic host name configured for the IS-IS system of the local switch.

By default, the IS-IS system of the local switch has no dynamic host name.

Format

is-name *symbolic-name*

undo is-name [*symbolic-name*]

Parameters

Parameter	Description	Value
symbolic-name	Specifies a dynamic host name.	The value is a string of 1 to 64 case-sensitive characters without spaces. When double quotation marks are used around the string, spaces are allowed in the string.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

In most cases, to check information about IS-IS neighbors and LSDBs on an IS-IS switch, you need to use a system ID of a 12-digit hexadecimal number, for example, aaaa.eeee.1234. This representation, however, is complicated and not easy to use. The dynamic host name mapping mechanism is introduced to facilitate maintenance and management of IS-IS networks. The **is-name** command is used to configure a simple host name for a local switch and enables the switch to advertise it in an LSP.

After the configuration is completed, you can run the **display isis name-table** command to check the configured hostname.

Example

Configure host name **RUTA** for the local IS-IS system.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] is-name RUTA
```

7.6.90 is-name map

Function

The **is-name map** command enables the local switch to identify the host name in LSPs, and sets a static host name for a remote IS-IS system on the local switch.

The **undo is-name map** command disables the local switch from identifying the host name in LSPs, and deletes the static host name for a remote IS-IS system set by the local switch.

By default, the local switch does not identify the host name in LSPs, and has no static host name of a remote IS-IS system.

Format

is-name map *system-id symbolic-name*

undo is-name map *system-id* [*symbolic-name*]

Parameters

Parameter	Description	Value
system-id	Specifies the ID of the remote mapped IS-IS system or pseudonode ID.	The format is XXXX.XXXX.XXXX[.XX].

Parameter	Description	Value
symbolic-name	Specifies the static host name of the remote mapped IS-IS system.	The value is a string of 1 to 64 characters without spaces. It is case sensitive. When double quotation marks are used around the string, spaces are allowed in the string.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When checking information about IS-IS neighbors and LSDBs on the switch that runs IS-IS, each switch in the IS-IS routing domain is represented by a system ID with a 12-bit hexadecimal number, for example, aaaa.eeee.1234. This representation is complicated and not easy to use. To manage and maintain IS-IS networks conveniently, IS-IS uses the dynamic host name exchange mechanism. The **is-name map** command can configure a host name for the remote switch and does not advertise the host name through an LSP.

After running the **is-name map** command to map the remote switch to the host name, you can find that the system ID of the remote switch is replaced by the host name configured using the **display isis name-table** command.

Precautions

This configuration is static configuration and takes effect only on the local device. Therefore, the configured host name *symbolic-name* is not advertised through an LSP. If dynamic host name mapping is configured on an IS-IS device, dynamic host name mapping takes precedence over static host name mapping.

Example

Configure static host name mapping for the IS-IS system on the remote switch as 0000.0000.0041.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] is-name map 0000.0000.0041 RUTB
```

7.6.91 log-peer-change

Function

The **log-peer-change** command enables the output of IS-IS adjacency changes.

The **undo log-peer-change** command disables the output of IS-IS adjacency changes.

By default, the output of IS-IS adjacency changes is disabled.

Format

log-peer-change [**topology**]

undo log-peer-change

Parameters

Parameter	Description	Value
topology	Enables the output of the IS-IS adjacency changes in an IPv6 topology.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On an IS-IS network, IS-IS adjacency changes pose a threat to network stability, causing frequent route convergence which consumes a lot of memory resources and possible traffic loss. Therefore, diagnose the fault immediately when the IS-IS adjacency changes.

To locate the fault, run the **log-peer-change** command to enable the output of IS-IS adjacency changes and record the changes in the log.

Precautions

IS-IS adjacency changes are recorded only when the **terminal monitor** and **log-peer-change** commands are run.

Example

Enable the output of IS-IS adjacency changes on the current switch.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] log-peer-change
```


7.6.92 loop-free-alternate

Function

The **loop-free-alternate** command enables IS-IS Auto FRR to calculate loop-free backup routes using the loop-free alternate (LFA) algorithm.

The **undo loop-free-alternate** command disables IS-IS Auto FRR from calculating loop-free backup routes using the LFA algorithm.

By default, IS-IS Auto FRR is disabled from calculating loop-free backup routes using the LFA algorithm.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

loop-free-alternate [**level-1** | **level-2** | **level-1-2**]

undo loop-free-alternate [**level-1** | **level-2** | **level-1-2**]

Parameters

Parameter	Description	Value
level-1	Enables Level-1 IS-IS Auto FRR to generate loop-free backup routes. If the IS level is not specified, Level-1 and Level-2 IS-IS Auto FRR is enabled to generate backup routes.	-
level-2	Enables Level-2 IS-IS Auto FRR to generate loop-free backup routes. If the IS level is not specified, Level-1 and Level-2 IS-IS Auto FRR is enabled to generate backup routes.	-
level-1-2	Enables Level-1 and Level-2 IS-IS Auto FRR to generate loop-free backup routes.	-

Views

IS-IS FRR view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS Auto FRR pre-computes a backup link by using the Loop-Free Alternate (LFA) algorithm, and then adds the backup link and the primary link to the forwarding table. In the case of an IS-IS network failure, IS-IS Auto FRR can fast switch traffic to the backup link before routes on the control plane converge.

Precautions

To make IS-IS Auto FRR to take effect, enter the IS-IS FRR view and then run the **loop-free-alternate** command.

IS-IS can generate loop-free backup routes only when the IS-IS Auto FRR traffic protection inequality is met.

Example

Enable Level-2 IS-IS Auto FRR to generate loop-free backup routes.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] frr
[HUAWEI-isis-1-frr] loop-free-alternate level-2
```

7.6.93 lsp-fragments-extend

Function

The **lsp-fragments-extend** command enables LSP fragment extension on the IS-IS switch.

The **undo lsp-fragments-extend** command disables LSP fragment extension on the IS-IS switch.

By default, LSP fragment extension is disabled on the IS-IS switch.

Format

lsp-fragments-extend [[**level-1** | **level-2** | **level-1-2**] | [**mode-1** | **mode-2**]] *

undo lsp-fragments-extend [**mode-1** | **mode-2**] [**level-1** | **level-2** | **level-1-2**]

Parameters

Parameter	Description	Value
level-1	Enables LSP fragment extension on a Level-1 switch.	-
level-2	Enables LSP fragment extension on a Level-2 switch.	-

Parameter	Description	Value
level-1-2	Enables LSP fragment extension on a Level-1-2 switch.	-
mode-1	Indicates that the switch supporting LSP fragment extension is compatible with the switch with an earlier version that does not support LSP fragment extension.	-
mode-2	Requires that all the switches support LSP fragment extension. **NOTE** By default, **mode-1** and **level-1-2** are used.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

To enable the IS-IS switch to generate extended LSP fragments, run the **virtual-system** command to configure at least one virtual system ID.

After the **lsp-fragments-extend** command is run, run the **reset isis all** command to restart the IS-IS process to make LSP fragment extension take effect.

Example

Enable Level-2 LSP fragment extension in mode 1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] lsp-fragments-extend mode-1 level-2
```

7.6.94 lsp-length

Function

The **lsp-length** command sets the length of the LSP that is generated and received by the current IS-IS switch.

The **undo lsp-length** command restores the default setting.

By default, the IS-IS switch generates and receives 1497-byte LSPs.

Format

lsp-length { **originate** | **receive** } *max-size*

undo lsp-length { **originate** | **receive** }

Parameters

Parameter	Description	Value
originate	Indicates the maximum length of generated LSPs.	-
receive	Indicates the maximum length of received LSPs.	-
max-size	Specifies the maximum length of LSPs. The *max-size* of a generated LSP must be smaller than or equal to the *max-size* of a received LSP.	The value is an integer that ranges from 512 to 16384, in bytes. The default value is 1497. **NOTE** Because the maximum MTU supported by interfaces is 9600 bytes and the LSP header length is 3 bytes, the allowed maximum LSP length is 9597 (9600 – 3) bytes to ensure that two devices on both ends communicate properly.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By controlling the length of LSPs generated and received by IS-IS switches, you can adjust the LSDB synchronization speed and change the network convergence speed.

Precautions

The *max-size* set using the **lsp-length** command must meet the following requirements:

- The MTU of an Ethernet interface must be greater than or equal to the sum of the value of *max-size* and 3.
- The MTU of a P2P interface must be greater than or equal to the value of *max-size*.

Example

Set the maximum length of an LSP generated by the IS-IS switch to 1024 bytes.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] lsp-length originate 1024
```


7.6.95 maximum load-balancing (IS-IS)

Function

The **maximum load-balancing** command sets the maximum number of equal-cost routes for load balancing. You need to set the maximum value according to memory capacity.

The **undo maximum load-balancing** command restores the default maximum number of equal-cost routes for load balancing.

By default, the maximum number of equal-cost routes on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S6720S-S, and S5736-S is 8, and the maximum number of equal-cost routes on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S is 16.

Format

maximum load-balancing *number*

undo maximum load-balancing [*number*]

Parameters

Parameter	Description	Value
number	Specifies the maximum number of equal-cost routes for load balancing.	The value is an integer that ranges from 1 to 8 on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S and S6720S-S. The value ranges from 1 to 16 on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If multiple routes are available on an IS-IS network, multiple equal-cost routes may exist. Use the load balancing method to allocate traffic evenly to each link. This method improves link efficiency on the network and reduces the congestion possibility caused by overloaded links.

Specify the *number* parameter to limit the number of equal-cost routes for load balancing.

Precautions

To cancel load balancing, set the *number* parameter to 1 or run the **nexthop** command to set the preferences of equal-cost routes.

Example

Set the maximum number of equal-cost routes for load balancing to 2.

```
<HUAWEI> system-view
[HUAWEI] isis 100
[HUAWEI-isis-100] maximum load-balancing 2
```

Restore the default value.

```
<HUAWEI> system-view
[HUAWEI] isis 100
[HUAWEI-isis-100] undo maximum load-balancing
```

7.6.96 network-entity

Function

The **network-entity** command configures a network entity title (NET) for an IS-IS process.

The **undo network-entity** command deletes the NET of an IS-IS process.

By default, no NET is configured for an IS-IS process.

Format

network-entity *net*

undo network-entity *net*

Parameters

Parameter	Description	Value
net	Specifies a NET.	The format is X...X.XXXX.XXXX.XXXX.00, in which X...X indicates an area address, the 12 Xs in the middle indicate the system ID of the switch, and the last 00 indicates the selector (SEL).

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

NET is the special form of the network service access point (NSAP). After the IS-IS view is displayed, IS-IS can start only when a NET is configured for an IS-IS process.

A NET consists of the following parts:

- Area ID: is 1 to 13 bytes in length.
- System ID: is 6 bytes in length.
- SEL: is 1 byte in length and fixed as 00.

Generally, you only need to configure one NET for an IS-IS process. When an area needs to be redefined, for example, the area needs to be merged with other areas or divided into sub-areas, configure multiple NETs to ensure route correctness.

Precautions

An area address uniquely identifies an area in a routing domain. All the switches in the same Level-1 area must have the same area address. The switches in a Level-2 area can have different area addresses.

The switches in an area or a backbone area must have the same system ID.

A maximum of three area addresses can be configured for an IS-IS process. Therefore, a maximum of three NETs can be configured for an IS-IS process. When configuring multiple NETs, ensure that their system IDs are the same.

Example

Set the NET to 10.0001.1010.1020.1030.00, in which the system ID is 1010.1020.1030 and the area ID is 10.0001.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] network-entity 10.0001.1010.1020.1030.00
```

7.6.97 nexthop (IS-IS)

Function

The **nexthop** command sets the preference of equal-cost routes.

The **undo nexthop** command cancels the preference of equal-cost routes.

By default, no preference is set for equal-cost routes.

Format

nexthop *ip-address* **weight** *value*

undo nexthop *ip-address*

Parameters

Parameter	Description	Value
ip-address	Specifies a next-hop IP address.	The value is in dotted decimal notation.

Parameter	Description	Value
weight *value*	Specifies the next-hop weight.	The value is an integer that ranges from 1 to 254.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When there are multiple redundant links on an IS-IS network, there may be multiple equal-cost routes to the same destination network segment.

You can use this command to set the next-hop preference to determine the next hop when the interface cost remains unchanged. A smaller preference value indicates a higher preference.

Precautions

After this command is used, an IS-IS device does not load balance traffic when forwarding the traffic. Instead, the device forwards the traffic to the next hop with the highest preference.

Example

Set the preference of IS-IS equal-cost routes.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] nexthop 10.0.0.3 weight 1
```

7.6.98 optional-checksum enable

Function

The **optional-checksum enable** command enables IS-IS to configure Hello packets and SNP packets to carry optional checksum TLVs and to check received IS-IS packets and SNP packets.

The **undo optional-checksum enable** command restores IS-IS packets to default settings.

By default, Hello packets and SNPs do not carry checksum TLVs and the receiver does not check received packets.

Format

optional-checksum enable

undo optional-checksum enable

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To prevent the attack of malicious packets and to ensure packets are correctly received on an IS-IS network, you can configure the **optional-checksum enable** command to enable IS-IS routers to send SNP packets and Hello packets carrying optional checksum TLVs. After the peer device receives the packets, it checks whether the carried optional checksum TLVs are correct. If the TLVs are not correct, the peer device rejects the packets.

Prerequisites

You have run the **isis** command to create an IS-IS process and entered the IS-IS view.

Precautions

If MD5 authentication or Keychain authentication with valid MD5 authentication is configured on an IS-IS interface or area, IS-IS routers send Hello packets and SNP packets carrying no checksum TLVs and verify the checksum of the received packets.

Example

Configure IS-IS to add optional checksum TLVs to Hello packets and SNPs.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] optional-checksum enable
```

7.6.99 preference (IS-IS)

Function

The **preference** command sets the IS-IS preference.

The **undo preference** command restores the default IS-IS preference.

The default IS-IS preference is 15.

Format

preference { *preference* | **route-policy** *route-policy-name* } *

undo preference

Parameters

Parameter	Description	Value
preference	Specifies the IS-IS preference. A smaller value indicates a higher preference.	The value is an integer that ranges from 1 to 255.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A device may run multiple dynamic routing protocols and there may be multiple routes of different routing protocols to the same destination. Therefore, routing protocols may share and select routing information. To solve this problem, the system sets a preference for each routing protocol. If different protocols discover routes to the same destination, the system selects the route with a higher protocol preference to forward traffic.

The **preference** command can set the preference of IS-IS routes to affect routing information sharing and selection:

- The **preference** *preference* command can set the preference for all IS-IS routes.
- The **preference** *preference* **route-policy** *route-policy-name* command can set different preferences for matched and unmatched routes.
- The **preference route-policy** *route-policy-name preference* command can set the preference for matched routes without affecting the preference of other IS-IS routes.

Precautions

You can use route-policies to set the preference for specified routes. If the **apply preference** clause is configured in the **route-policy** command, the route preference is as follows:

- The preference of matched routes is set using the **apply** clause.
- The preference of unmatched routes is set using the **preference** command.

As shown in the following example, the preference of the routes that pass route-policy **abc** is set to 50 and the preference of the routes that do not pass the route-policy is set to 30.

```
#
route-policy abc permit node 1
 if-match cost 20
 apply preference 50
#
isis 1
 preference 30 route-policy abc
```

If the **apply preference 50** clause is not configured in route-policy **abc**, the preference of all routes is set to 30.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the configured priority applies to all IS-IS routes.

Example

Set the IPv4 IS-IS preference to 25.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] preference 25
```

7.6.100 prefix-priority (IS-IS)

Function

The **prefix-priority** command sets the convergence priority of IS-IS routes.

The **undo prefix-priority** command restores the default convergence priority of IS-IS routes.

By default, the convergence priority of IS-IS host routes and default routes is medium, and the convergence priority of other IS-IS routes is low.

Format

prefix-priority [**level-1** | **level-2**] { **critical** | **high** | **medium** } { **ip-prefix** *prefix-name* | **tag** *tag-value* }

undo prefix-priority [**level-1** | **level-2**] { **critical** | **high** | **medium** }

Parameters

Parameter	Description	Value
level-1	Specifies the convergence priority of Level-1 IS-IS routes.	-

Parameter	Description	Value
level-2	Specifies the convergence priority of Level-2 IS-IS routes.	-
critical	Sets the convergence priority of IS-IS routes to critical.	-
high	Sets the convergence priority of IS-IS routes to high.	-
medium	Sets the convergence priority of IS-IS routes to medium.	-
ip-prefix *prefix-name*	Sets the convergence priority of the IS-IS routes matching the specified IP prefix.	The value is a string of 1 to 169 case-sensitive characters without spaces.
tag *tag-value*	Sets the convergence priority of the IS-IS routes with the specified tag value. To use the tag value to filter the IPv4 routes for which convergence priorities need to be set, ensure that the IS-IS cost style of sent packets is not narrow and the routes carry the tag value.	The value is an integer that ranges from 1 to 4294967295.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a routing table has a large number of routing entries on an IS-IS device, the device needs to spend much time performing SPF calculation. To enable some key routes to be calculated first, run the **prefix-priority** command to set the convergence priority for these routes.

The convergence priorities of IS-IS routes are classified as critical, high, medium, and low in descending order.

The rules for applying convergence priorities to IS-IS routes are as follows:

- Existing IS-IS routes are converged based on the priorities set using the **prefix-priority** command.
- New IS-IS routes are converged based on the priorities set using the **prefix-priority** command.
- If an IS-IS route conforms to the matching rules of multiple convergence priorities, the highest convergence priority is used.
- If the route level is not specified, the **prefix-priority** command configuration takes effect for both Level-1 and Level-2 IS-IS routes.

Precautions

The **prefix-priority** command applies only to the public network.

After the **prefix-priority** command is run to set the convergence priority for IS-IS routes (including IS-IS host routes and default routes), the convergence priority of all the IS-IS routes that meet the matching rules is changed according to the command configuration, and the convergence priority of the IS-IS routes that do not meet the matching rules is changed to low.

Example

Set the convergence priority of Level-1 routes with tag value 3 to critical.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] prefix-priority level-1 critical tag 3
```

Set the convergence priority of the routes matching IP prefix **p1** to medium.

```
<HUAWEI> system-view
[HUAWEI] ip ip-prefix p1 permit 192.168.0.1 24
[HUAWEI] isis 1
[HUAWEI-isis-1] prefix-priority medium ip-prefix p1
```

7.6.101 purge-originator-identification enable

Function

The **purge-originator-identification enable** command configures IS-IS to add purge-originator-identification (POI) TLV to Purge packets. If a dynamic hostname has been configured for the local device, the hostname TLV is also added to the Purge packets.

The **undo purge-originator-identification enable** command deletes POI TLV and hostname TLV from Purge packets.

By default, Purge packets do not carry POI TLV or hostname TLV.

Format

purge-originator-identification enable [**always**]

undo purge-originator-identification enable [**always**]

Parameters

Parameter	Description
always	

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When **Remaining Lifetime** of an LSP is 0, the LSP is invalid, and this invalid packet is a Purge packet. In most cases, a Purge packet does not carry any information about the router that generated the Purge packet. Without such information, troubleshooting is difficult if a network problem occurs.

To address this issue, you can run the **purge-originator-identification enable** command to configure IS-IS to add POI TLV to Purge packets. If a dynamic hostname has been configured for the local device, the hostname TLV is also added to the Purge packets, which facilitates troubleshooting.

- If the **purge-originator-identification enable** command is run and an authentication mode is configured, generated Purge LSPs do not carry the POI TLV or hostname TLV.
- If the **purge-originator-identification enable** command is run and no authentication is configured, generated Purge LSPs will carry the POI TLV or hostname TLV.
- If the **purge-originator-identification enable always** command is run, generated Purge LSPs carry the POI TLV and hostname TLV, regardless of whether authentication is configured.

Example

Configure IS-IS to add POI TLV to Purge packets.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] purge-originator-identification enable
```

7.6.102 reset isis all

Function

The **reset isis all** command restarts IS-IS processes.

Format

reset isis all [[*process-id* | **vpn-instance** *vpn-instance-name*] | **graceful-restart**] *

reset isis *process-id* **all** [**graceful-restart**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an IS-IS process to be restarted.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance. *vpn-instance-name* specifies the name of a VPN instance.	The value must be an existing VPN instance name.
graceful-restart	Restarts an IS-IS process in GR mode.	-

Views

User view

Default Level

3: Management level

Usage Guidelines

Usage Scenario

The **reset isis all** command is used when LSPs need to be updated immediately. For example, after the **area-authentication-mode** and **domain-authentication-mode** commands are used, you can run the **reset isis all** command to clear old LSPs.

Precautions

NOTICE

Restarting IS-IS processes may interrupt services. Therefore, confirm the action before running the **reset isis all** command.

Example

Restart IS-IS processes.

```
<HUAWEI> reset isis all
Warning: The ISIS process(es) will be reset. Continue?[Y/N]y
```

7.6.103 reset isis peer

Function

The **reset isis peer** command resets a specified IS-IS neighbor relationship.

Format

reset isis peer *system-id* [*process-id* | **vpn-instance** *vpn-instance-name*]

reset isis *process-id* **peer** *system-id*

Parameters

Parameter	Description	Value
system-id	Specifies the system ID of an IS-IS neighbor.	The value is 6 bytes in length and in XXXX.XXXX.XXXX format.
process-id	Specifies the IS-IS process ID for which the neighbor relationship needs to be reset.	The value is an integer that ranges from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the IS-IS multi-instance process in a specified VPN instance. This parameter is optional for a default VPN instance.	The value must be an existing VPN instance name.

Views

User view

Default Level

3: Management level

Usage Guidelines

The **reset isis peer** command is used when a device needs to re-establish the neighbor relationship with a specified neighbor.

NOTICE

Exercise caution when using the **reset isis peer** command. This command will re-establish neighbor relationships with a specified neighbor, which may result in route flapping.

Example

Reset the IS-IS neighbor with system ID 0000.0c11.1111.

```
<HUAWEI> reset isis peer 0000.0c11.1111
```


7.6.104 reset isis suppress-flapping peer

Function

The **reset isis suppress-flapping peer** command forces an interface to exit from IS-IS neighbor relationship flapping suppression.

Format

reset isis *process-id* **suppress-flapping peer** [**interface** *interface-type interface-number*] [**notify-peer**]

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an IS-IS process.	The value is an integer ranging from 1 to 65535.
interface *interface-type interface-number*	Specifies an interface type and number.	-
notify-peer	Instructs neighbors to exit from IS-IS neighbor relationship flapping suppression too.	-

Views

User view

Default Level

3: Management level

Usage Guidelines

Interfaces exit from flapping suppression in the following scenarios:

- The suppression timer expires.
- The corresponding IS-IS process is reset.
- An IS-IS neighbor is reset using the **reset isis peer** command.
- IS-IS neighbor relationship flapping suppression is disabled globally using the **suppress-flapping peer disable (IS-IS)** command in the IS-IS view.
- The **reset isis suppress-flapping peer** command is run.
- Suppression is aborted forcibly using the **reset isis** *process-id* **suppress-flapping peer** [**interface** *interface-type interface-number*] **notify-peer** command on the remote device.

Example

Force interfaces to exit from IS-IS neighbor relationship flapping suppression.

```
<HUAWEI> reset isis 1 suppress-flapping peer
```

7.6.105 set-overload

Function

The **set-overload** command sets the overload bit for non-pseudonode LSPs.

The **undo set-overload** command removes the overload bit of non-pseudonode LSPs.

By default, no overload bit is set for non-pseudonode LSPs.

Format

set-overload [**on-startup** [*timeout1* | **start-from-nbr** *system-id* [*timeout1* [*timeout2*]] | **wait-for-bgp** [*timeout1*]] [**send-sa-bit** [*timeout3*]]] [**allow** { **interlevel** | **external** }*]

undo set-overload

Parameters

Parameter	Description	Value
on-startup	Indicates that the overload bit remains set within the specified period when the switch restarts or is faulty.	-
timeout1	Specifies the duration of the overload bit after the system starts.	The value is an integer that ranges from 5 to 86400, in seconds. The default value is 600 seconds.
start-from-nbr *system-id*	Specifies the duration of the system overload bit according to the status of a specified neighbor.	The value is in XXXX.XXXX.XXXX format.

Parameter	Description	Value
timeout1 [*timeout2*]	Specifies the period during which the overload bit remains set, which is related to the neighbor status.	• If the specified neighbor does not go Up before *timeout2* expires, the duration of the system overload bit is *timeout2*. *timeout2* ranges from 5 to 86400, in seconds. The default value is 1200 seconds (20 minutes). • If the specified neighbor goes Up before *timeout2* expires, the duration of the system overload bit is *timeout1*. *timeout1* ranges from 5 to 86400, in seconds. The default value is 600 seconds (10 minutes).
wait-for-bgp	Specifies the period during which the overload bit of the system remains set according to the BGP convergence status.	-
send-sa-bit	Specifies the overload bit remains in Hello packets after the device is started.	-
timeout3	Specifies the period during which the overload bit remains in Hello packets after the device is started.	The value is an integer ranging from 5 to 120, in seconds. The default value is 30 seconds.
allow	Allows advertising IP prefixes. By default, the system is prohibited from advertising IP prefixes when the system enters the overload state.	-
interlevel	Allows advertising the IP prefixes learned from IS-IS of different levels when **allow** is specified.	-
external	Allows advertising the IP prefixes learned from other protocols when **allow** is specified.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Although LSPs with the overload bit are flooded on the network, the LSPs are ignored in the calculation of the routes that pass through the switch in overload state. That is, after the overload bit is set on the switch, other switches ignore the switch when performing SPF calculation. Direct routes of the switch are not ignored during SPF calculation.

To prevent the local switch from being used by other switches to perform SPF calculation, run the **set-overload** command on the local switch without specifying the **on-startup** keyword. Then the system immediately sets the overload bit in the LSP to be sent. To remove the overload bit, run the **undo set-overload** command.

When the local switch restarts or is faulty, to prevent the local switch from being used by other switches to perform SPF calculation, run the **set-overload** command on the local switch and specify the **on-startup** keyword.

When the switch is experiencing memory shortage, the system automatically sets the overload bit in the sent LSPs regardless of whether the **set-overload** command is run.

You can set the overload bit to solve the problem of network traffic loss caused by the difference between the BGP convergence speed and IGP convergence speed.

Precautions

If a TE LSP uses the local device as a transit node before the **set-overload** command is run, the TE LSP is not torn down and re-established and still uses the local device as a transit node after the **set-overload** command is run; if the local device is restarted after the command is run and fast convergence is not configured on the ingress of the RSVP-LSP, TE LSP forwarding fails, and services are affected. Therefore, the **mpls te path-selection overload** command needs to be run on the ingress of the RSVP-LSP before the device is restarted.

Example

Set the overload bit for IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] set-overload
Warning: The IS-IS process overload state will be set. Continue?[Y/N]
```


7.6.106 spf-priority

Function

The **spf-priority** command sets the priority of SPF calculation.

The **undo spf-priority** command restores the default priority of SPF calculation.

By default, the priority of SPF calculation in a base topology is 64.

Format

spf-priority *priority-value*

undo spf-priority

Parameters

Parameter	Description	Value
priority-value	Specifies the priority of SPF calculation. A larger value indicates a higher priority.	The value is an integer that ranges from 1 to 127.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

When multiple IS-IS topology instances are deployed on a network, the **spf-priority** command can be used to set a priority for the SPF calculation in each IS-IS topology instance.

If a higher priority is set for a topology, routes of the topology instance are converged first, ensuring that important services are running properly.

Example

Set the priority of SPF calculation in an IPv4 base topology to 30.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] spf-priority 30
```


7.6.107 spf-slice-size

Function

The **spf-slice-size** command sets the maximum duration for IS-IS route calculation.

The **undo spf-slice-size** command restores the default setting.

By default, IS-IS route calculation lasts for a maximum of 2 ms at a time.

Format

spf-slice-size *duration-time*

undo spf-slice-size

Parameters

Parameter	Description	Value
duration-time	Specifies the maximum duration for IS-IS route calculation.	The value is an integer that ranges from 1 to 5000, in milliseconds. The default value is 2.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

When a routing table contains a large number of routes, run the **spf-slice-size** command to specify the maximum duration of SPF calculation to prevent route calculation from using system resources for a long period.

Example

Set the maximum duration of IS-IS route calculation to 50 ms.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] spf-slice-size 50
```

7.6.108 summary (IS-IS)

Function

The **summary** command configures IS-IS to generate summarized routes.

The **undo summary** command disables IS-IS from generating summarized routes.

By default, IS-IS does not generate summarized routes.

Format

summary *ip-address mask* [**avoid-feedback** | **generate_null0_route** | **tag** *tag* | [**level-1** | **level-1-2** | **level-2**]] *

undo summary *ip-address mask* [**level-1** | **level-1-2** | **level-2**]

Parameters

Parameter	Description	Value
ip-address	Specifies the IPv4 address of a summarized route.	The value is in dotted decimal notation.
mask	Specifies the mask of the IP address of a summarized route.	The value is in dotted decimal notation.
avoid-feedback	Prevents summarized routes from being learned through SPF calculation.	-
generate_null0_route	Generates null0 routes to prevent routing loops.	-
tag *tag*	Assigns administrative tags to advertised routes.	The value is an integer that ranges from 1 to 4294967295.
level-1	Summarizes only the routes imported into Level-1 areas. If no level is specified, Level-2 is used by default.	-
level-1-2	Summarizes the routes imported into Level-1 areas and the backbone network. If no level is specified, Level-2 is used by default.	-
level-2	Summarizes only the routes imported into the backbone network. If no level is specified, Level-2 is used by default.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Route summarization on a large-scale IS-IS network efficiently reduces routing entries. This minimizes system resource consumption and maintains system performance.

If a link frequently alternates between Up and Down, the links not involved in the route summarization will not be affected. Therefore, route summarization prevents route flapping on the network and improves network stability.

Precautions

Configuring route summarization does not affect the routing table of the local device because each specific route is still displayed in the routing table. Route summarization reduces the flooding of LSPs. Routing tables of other devices that receive the LSP from the local device contain only one summarized routes but not specific routes. Specific routes are still displayed in the routing table of the device directly connected to the summarized network segment.

Both IS-IS routes and routes imported from other protocols can be summarized. The lowest cost among that of the routes that are summarized is used as the cost of the summarized route.

Example

Generate a summarized route with destination subnet 202.0.0.0/8.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] summary 10.0.0.0 255.0.0.0
```

7.6.109 suppress-flapping peer disable (IS-IS)

Function

The **suppress-flapping peer disable** command disables IS-IS neighbor relationship flapping suppression globally.

The **undo suppress-flapping peer disable** command enables IS-IS neighbor relationship flapping suppression globally.

By default, IS-IS neighbor relationship flapping suppression is enabled globally.

Format

suppress-flapping peer disable

undo suppress-flapping peer disable

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

If an interface carrying IS-IS services alternates between Up and Down, IS-IS neighbor relationship flapping occurs on the interface. During the flapping, IS-IS reestablishes the neighbor relationship and recalculates routes. In this process, a large number of packets are exchanged, adversely affecting neighbor relationship stability, IS-IS services, and other IS-IS-dependent services, such as LDP and BGP. IS-IS neighbor relationship flapping suppression can address this problem by delaying IS-IS neighbor relationship reestablishment or preventing service traffic from passing through flapping links.

By default, IS-IS neighbor relationship flapping suppression is enabled globally. To disable this function globally, run the **suppress-flapping peer disable** command.

Example

Disable neighbor relationship flapping suppression globally.

```
<HUAWEI> system-view
[HUAWEI] isis 1
[HUAWEI-isis-1] suppress-flapping peer disable
```

7.6.110 timer lsp-generation

Function

The **timer lsp-generation** command configures an intelligent timer for generating LSPs.

The **undo timer lsp-generation** command cancels the intelligent timer for generating LSPs.

By default, no intelligent timer is configured to generate LSPs.

Format

timer lsp-generation *max-interval* [*init-interval* [*incr-interval*]] [**level-1** | **level-2**]

undo timer lsp-generation [*max-interval* [*init-interval* [*incr-interval*]]] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
max-interval	Specifies the maximum delay in generating LSPs with the same LSP ID.	The value is an integer that ranges from 1 to 120, in seconds. The default value is 2.
init-interval	Specifies the initial delay in generating LSPs.	The value is an integer that ranges from 1 to 60000, in milliseconds. By default, this delay is not used.
incr-interval	Specifies the incremental delay in generating two LSPs with the same LSP ID.	The value is an integer that ranges from 1 to 60000, in milliseconds. By default, this delay is not used.
level-1	Specifies the delay in generating Level-1 LSPs. If no level is specified, the delay in generating Level-1 and Level-2 LSPs is set.	-
level-2	Specifies the delay in generating Level-2 LSPs. If no level is specified, the delay in generating Level-1 and Level-2 LSPs is set.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On an IS-IS network, if the local routing information changes, a device needs to generate a new LSP to notify this change. If the local routing information changes frequently, a large number of new LSPs are generated, which occupies a lot of system resources and decreases system performance.

To speed up network convergence and prevent the system from being affected, run the **timer lsp-generation** command to configure an intelligent timer for generating LSPs. This timer can adjust the delay in generating LSPs based on the routing information change frequency. This command can implement different functions when different parameters are specified:

- When only *max-interval* is specified, the intelligent timer functions as an ordinary one-time triggering timer.

- When both *init-interval* and *incr-interval* are specified, the delay in generating an LSP for the first time is determined by *init-interval*, and the delay in generating an LSP with the same LSP ID for the second time is determined by *incr-interval*. Subsequently, each time routes change, the delay in generating an LSP doubles the last delay until the delay reaches the value specified by *max-interval*. After the delay remains at the value specified by *max-interval* for three times or the IS-IS process is restarted, the delay decreases to the value specified by *init-interval*.
- When *init-interval* is specified but *incr-interval* is not, the delay in generating an LSP for the first time is determined by *init-interval*, and the delay in generating subsequent LSPs is determined by *max-interval*. After the delay remains at the value specified by *max-interval* for three times or the IS-IS process is restarted, the delay decreases to the value specified by *init-interval*.
- The initial delay for generating an LSP is *init-interval*; the delay for generating the LSP with the same LSP ID is *incr-interval*. From the third time on, the delay for generating the LSP doubles each time until the delay reaches *max-interval*. After *max-interval* expires without flapping or the IS-IS process is restarted, the delay decreases to *init-interval*.

Precautions

If the delay in generating an LSP is too long, a device cannot advertise routing information changes to neighbors in time. This slows down network convergence.

Example

Set the delay in generating LSPs to 5s.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] timer lsp-generation 5
```

Set the maximum delay in generating LSPs to 20s, initial delay to 50 ms, and incremental delay to 2000 ms.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] timer lsp-generation 20 50 2000
```

7.6.111 timer lsp-max-age

Function

The **timer lsp-max-age** command sets the maximum lifetime for the LSPs generated by an IS-IS process.

The **undo timer lsp-max-age** command restores the default setting.

By default, the maximum lifetime of LSPs is 1200 seconds.

Format

timer lsp-max-age *age-time*

undo timer lsp-max-age

Parameters

Parameter	Description	Value
age-time	Specifies the maximum lifetime for LSPs.	The value is an integer that ranges from 2 to 65535, in seconds.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the switch generates the system LSP, the system fills in the LSP with its maximum lifetime. After other switches receive this LSP, the lifetime of the LSP is reduced accordingly. If the switch does not receive an updated LSP, the lifetime of the LSP is reduced to 0, and the switch retains the LSP for another 60 seconds. If the switch still does not receive any updated LSP within 60 seconds, the switch deletes the LSP.

You can run the **timer lsp-max-age** command to adjust the maximum lifetime of LSPs to ensure the validity of existing LSPs before new LSPs are received.

Precautions

The maximum lifetime of LSPs must be longer than the interval for updating LSPs to ensure that LSPs can be updated before being deleted.

Example

Set the maximum lifetime of the LSPs generated by the current IS-IS process to 25 minutes (1500 seconds).

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] timer lsp-max-age 1500
```

7.6.112 timer lsp-refresh

Function

The **timer lsp-refresh** command sets the interval for updating LSPs.

The **undo timer lsp-refresh** command restores the default setting.

By default, the interval for updating LSPs is 900 seconds.

Format

timer lsp-refresh *refresh-time*

undo timer lsp-refresh

Parameters

Parameter	Description	Value
refresh-time	Specifies the interval for updating LSPs.	The value is an integer that ranges from 1 to 65534, in seconds.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On an IS-IS network, LSDB synchronization is implemented through LSP flooding. During LSP flooding, a device sends an LSP to its neighbors and then the neighbors send the received LSP to their respective neighbors except the device that first sends the LSP. In this manner, the LSP is flooded among the devices of the same level. LSP flooding allows each device of the same level to have the same LSP information and synchronize its LSDB with each other.

LSPs need to be periodically flooded. When the system generates the system LSP, the system fills in the LSP with its maximum lifetime. After other devices receive this LSP, the lifetime of the LSP is reduced accordingly. If the device does not receive an updated LSP, the lifetime of the LSP is reduced to 0, and the device retains the LSP for another 60 seconds. If the device still does not receive any updated LSP within 60 seconds, the device deletes the LSP.

You can run the **timer lsp-refresh** command to set the interval for updating LSPs to control network convergence speed.

Precautions

If the interval for updating LSPs is small, LSPs may be updated frequently, which occupies a lot of bandwidth resources. If the interval is long, routing information changes cannot be notified in a timely manner.

Example

Set the interval for updating LSPs to 1200 seconds.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] timer lsp-refresh 1200
```


7.6.113 timer spf

Function

The **timer spf** command sets the delay in SPF calculation.

The **undo timer spf** command restores the default setting.

By default, the maximum delay in SPF calculation is 5 seconds.

Format

timer spf *max-interval* [*init-interval* [*incr-interval*]]

undo timer spf

Parameters

Parameter	Description	Value
max-interval	Specifies the maximum delay in SPF calculation.	The value is an integer that ranges from 1 to 120, in seconds. The default value is 5s.
init-interval	Specifies the initial delay in SPF calculation. If the *init-interval* parameter is not specified, the intelligent timer functions as an ordinary one-time triggering timer.	The value is an integer that ranges from 1 to 60000, in milliseconds. By default, it is 50 milliseconds.
incr-interval	Specifies the incremental delay in two SPF calculations. If the *incr-interval* parameter is not specified, the delay in SPF calculation for the first time is determined by *init-interval*. From the second time on, the delay in SPF calculation is determined by *max-interval*.	The value is an integer that ranges from 1 to 60000, in milliseconds. The default value is 200 milliseconds.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In IS-IS, route calculation is required when an LSDB changes. Frequent route calculations, however, consume lots of system resources and decrease system performance. Delaying SPF calculation can improve route calculation efficiency. However, a long delay in SPF calculation slows down route convergence.

To speed up route convergence without affecting the efficiency of switches, configure an intelligent timer for SPF calculation. This timer automatically adjusts the delay in SPF calculation based on the LSDB change frequency.

When SPF calculation is performed for the first time, the interval is *init-interval*. Each time a route changes, the interval is added by *incr-interval* until the interval reaches *max-interval*. After *max-interval* expires without flapping or the IS-IS process is restarted, the delay decreases to *init-interval*.

When only *max-interval* is used, the intelligent timer becomes a one-short triggered timer.

Example

Set the delay in SPF calculation to 5s.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] timer spf 5
```

Set the maximum delay in SPF calculation to 15s, initial delay to 10 ms, and incremental delay to 5000 ms.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] timer spf 15 10 5000
```

7.6.114 virtual-system

Function

The **virtual-system** command sets the virtual system ID for an IS-IS process. If no virtual system ID is set, no extended LSP can be generated. A maximum of 50 virtual system IDs can be configured for an IS-IS process.

The **undo virtual-system** command deletes the virtual system ID of an IS-IS process.

By default, no virtual system ID is configured for an IS-IS process.

Format

virtual-system *virtual-system-id*

undo virtual-system *virtual-system-id*

Parameters

Parameter	Description	Value
virtual-system-id	Specifies a virtual system ID for an IS-IS process.	The value is 6 bytes in length and in XXXX.XXXX.XXXX format.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Like a common system ID, a virtual system ID must be unique in a routing domain.

The **virtual-system** command is used together with the **lsp-fragments-extend** command. The configured virtual system ID takes effect only when LSP fragment extension is enabled and the IS-IS process is restarted using the **reset isis all** command.

If LSP fragment extension is not enabled and the IS-IS process is not restarted, you can configure the virtual system ID but the configured virtual system ID does not take effect.

Before running this command, run the **network-entity** command in the IS-IS view to configure a network entity title for the IS-IS process.

Example

Set the virtual system ID of IS-IS process 1 to 2222.2222.2222.

```
<HUAWEI> system-view
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 10.0000.0000.0001.00
[HUAWEI-isis-1] virtual-system 2222.2222.2222
```

7.7 IPv6 IS-IS Configuration Commands

7.7.1 Command Support

Only the following switch models support IPv6 IS-IS:

S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S

7.7.2 display isis ipv6 bfd interface

Function

The **display isis ipv6 bfd interface** command displays information about the interface enabled with IPv6 BFD for IS-IS.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display isis [*process-id*] **ipv6 bfd interface**

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an IS-IS process.	The value is an integer ranging from 1 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

BFD can provide millisecond-level fault detection. It can work with IS-IS to fast detect faults on neighboring devices and instruct IS-IS to recalculate routes for correct packet forwarding. You can run the **display isis ipv6 bfd interface** command to check information about the interface enabled with IPv6 BFD for IS-IS.

Example

Display information about the interface enabled with IPv6 BFD.

```
<HUAWEI> display isis 1 ipv6 bfd interface
             IPV6 BFD information of interface for ISIS(1)
            ------------------------------------------------
Interface                BFD6.State     Min-Tx       Min-Rx        Mul
GE0/0/1                   enable        150          150           3
GE0/0/2                   enable        150          150           3
Total interfaces: 2              Total IPv6 bfd enabled interfaces: 2
```

Table 7-119 Description of the **display isis ipv6 bfd interface** command output

Item	Description
Interface	Interface enabled with IPv6 BFD for IS-IS
BFD6.State	IPv6 BFD status on the interface: • enable • disable To enable IPv6 BFD for IS-IS, run the **ipv6 bfd all-interfaces enable (IS-IS)** or **isis ipv6 bfd enable** command.

Item	Description
Min-Tx	Configured minimum interval for sending IPv6 BFD packets. To the parameters of IPv6 BFD sessions, run the **ipv6 bfd all-interfaces (IS-IS)** or **isis ipv6 bfd** command.
Min-Rx	Configured minimum interval for receiving IPv6 BFD packets. To the parameters of IPv6 BFD sessions, run the **ipv6 bfd all-interfaces (IS-IS)** or **isis ipv6 bfd** command.
Mul	Local detection multipliers of IPv6 BFD packets. To the parameters of IPv6 BFD sessions, run the **ipv6 bfd all-interfaces (IS-IS)** or **isis ipv6 bfd** command.
Total interfaces	Number of interfaces enabled with the IS-IS process
Total IPv6 bfd enabled interfaces	Number of interfaces enabled with IPv6 BFD for IS-IS

7.7.3 display isis ipv6 bfd session

Function

The **display isis ipv6 bfd session** command displays information about the IPv6 BFD session for IS-IS.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display isis [*process-id* | **vpn-instance** *vpn-instance-name*] **ipv6 bfd session** { **all** | **peer** *ipv6-address* | **interface** *interface-type interface-number* }

Parameters

Parameter	Description	Value
process-id	Specifies the ID of an IS-IS process.	The value is an integer ranging from 1 to 65535.
vpn-instance *vpn-instance-name*	Specifies the name of an IPv6 VPN instance, which is case sensitive.	The value must be an existing VPN instance name.
all	Indicates all the interfaces that run the IS-IS process.	-

Parameter	Description	Value
peer *ipv6-address*	Specifies the IPv6 link-local address of a peer.	The prefix is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
interface *interface-type interface-number*	Specifies the interface whose IPv6 BFD session statistics need to be collected.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

BFD can provide millisecond-level fault detection. It can work with IS-IS to fast detect faults on neighboring devices and instruct IS-IS to recalculate routes for correct packet forwarding. You can run the **display isis ipv6 bfd session** command to check information about the IPv6 BFD session for IS-IS.

Example

Display information about the IPv6 BFD session for IS-IS.

```
<HUAWEI> display isis 1 ipv6 bfd session all

              IPv6 BFD session information for ISIS(1)
              ----------------------------------------

Peer System ID : 0000.0000.0002   Type : L2
Interface : GE0/0/1
IPv6 BFD State : up       TX : 150   RX : 150   Multiplier : 3
LocDis : 8195         Local IPv6 Address : FE80::200:13FF:FE82:4569
RemDis : 8194         Peer  IPv6 Address : FE80::225:9EFF:FEFB:BFF1
Diag : No diagnostic information

Peer System ID : 0000.0000.0003   Type : L2
Interface : GE0/0/2
IPv6 BFD State : up       TX : 150   RX : 150   Multiplier : 3
LocDis : 8196         Local IPv6 Address : FE80::200:13FF:FE82:4569
RemDis : 8205         Peer  IPv6 Address : FE80::201:FF:FE01:1
Diag : No diagnostic information

Total IPv6 BFD session(s): 2
```

Display information about the IPv6 BFD session for IS-IS of the peer with the specified IPv6 address.

```
<HUAWEI> display isis 1 ipv6 bfd session peer FE80::225:9EFF:FEFB:BFF1

              IPv6 BFD session information for ISIS(1)
              ----------------------------------------

Peer System ID : 0000.0000.0002   Type : L2
```



```
Interface : GE0/0/2
IPv6 BFD State : up       TX : 150   RX : 150   Multiplier : 3
LocDis : 8195         Local IPv6 Address : FE80::200:13FF:FE82:4569
RemDis : 8194         Peer  IPv6 Address : FE80::225:9EFF:FEFB:BFF1
Diag : No diagnostic information

Total IPv6 BFD session(s): 1
```

Display information about the IPv6 BFD session for IS-IS of the specified interface.

```
<HUAWEI> display isis 1 ipv6 bfd session interface GigabitEthernet 0/0/2

              IPv6 BFD session information for ISIS(10)
              -----------------------------------------

Peer System ID : 0000.0000.0002   Type : L2
Interface : GE0/0/2
IPv6 BFD State : up       TX : 150   RX : 150   Multiplier : 3
LocDis : 8195         Local IPv6 Address : FE80::200:13FF:FE82:4569
RemDis : 8194         Peer  IPv6 Address : FE80::225:9EFF:FEFB:BFF1
Diag : No diagnostic information

Total IPv6 BFD session(s): 1
```

Table 7-120 Description of the **display isis ipv6 bfd session** command output

Item	Description
Peer system ID	System ID of the peer
Interface	Local interface name
Type	Level of the peer: • L1: Level-1 • L2: Level-2 • L12: Level-1-2
IPv6 BFD State	IPv6 BFD status: • up: indicates that the IPv6 BFD session is established successfully. • down: indicates that the IPv6 BFD fails to be established.
TX	Negotiated minimum interval for sending IPv6 BFD packets
RX	Negotiated minimum interval for receiving IPv6 BFD packets
Multiplier	Remote detection multiplier
LocDis	Local identifier dynamically assigned by IPv6 BFD
Local IPv6 Address	IPv6 link-local address
RemDis	Remote identifier dynamically assigned by IPv6 BFD
Peer IPv6 Address	IPv6 link-local address of the peer

Item	Description
Diag	Diagnostic information: • No diagnostic information: IPv6 BFD runs properly, and no diagnostic information is displayed. • Administrator down: The session is set Down by the network administrator. • Global IPv6 BFD is not enabled: Global IPv6 BFD is disabled. • IPv6 BFD session number reaches the MAX BFD: The number of IPv6 BFD sessions reaches the upper limit. • No IPv6 BFD packets were received: The device does not receive any IPv6 BFD packet. • Neighbour is down: The neighbor goes Down. • Administrator down event received: The device receives an event indicating that the IPv6 BFD session is set Down by the network administrator.
Total IPv6 BFD session(s)	Total number of IPv6 BFD sessions

7.7.4 ipv6 auto-cost enable

Function

The **ipv6 auto-cost enable** command enables IS-IS to automatically calculate the interface cost on an IPv6 network according to the interface bandwidth.

The **undo ipv6 auto-cost enable** command disables IS-IS from automatically calculating the interface cost on an IPv6 network according to the interface bandwidth.

By default, the automatic interface cost calculation function is disabled on an IPv6 network.

Format

ipv6 auto-cost enable [**compatible**]

undo ipv6 auto-cost enable

Parameters

Parameter	Description	Value
compatible	Specifies the IS-IS to calculate the cost of an interface on an IPv6 network based on the bandwidth of the interface automatically in compatible.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the function is enabled, for a certain IS-IS interface, if the cost of the interface is not configured in the interface view, and the global cost is not configured in the IS-IS view, the system calculates the cost for the interface automatically.

If the cost style of the system is wide or wide-compatible:

When **auto-cost enable** command is configured, Interface cost = (Bandwidth-reference/Link-bandwidth) x 10.

The cost style is set by the **cost-style** command. The Bandwidth-reference is set by the **ipv6 bandwidth-reference** command. The Link-bandwidth is the interface bandwidth.

If the cost-style of the system is narrow, narrow-compatible, or compatible, the cost of each interface is determined according to the following table.

Table 7-121 Relationship of the IS-IS interface cost and the bandwidth

Cost	Range of Interface Bandwidth
60	Interface bandwidth ≤ 10 Mbit/s
50	10 Mbit/s < Interface bandwidth ≤ 100 Mbit/s
40	100 Mbit/s < Interface bandwidth ≤ 155 Mbit/s
30	155 Mbit/s < Interface bandwidth ≤ 622 Mbit/s
20	622 Mbit/s < Interface bandwidth ≤ 2.5 Gbit/s
10	2.5 Gbit/s < Interface bandwidth

Precautions

The priority of the cost value of the global IPv6 configured by the **ipv6 circuit-cost** command is higher than the auto cost value.

This command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

The **ipv6 auto-cost enable** command cannot change the cost of the loopback interface.

Example

Enable IS-IS to automatically calculate the interface cost on an IPv6 network according to the interface bandwidth.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 auto-cost enable
```

7.7.5 ipv6 bandwidth-reference

Function

The **ipv6 bandwidth-reference** command sets the bandwidth reference value used in automatic interface cost calculation on an IPv6 network.

The **undo ipv6 bandwidth-reference** command restores the default setting.

By default, the bandwidth reference value is 100 Mbit/s.

Format

ipv6 bandwidth-reference *bandwidth-reference-value*

undo ipv6 bandwidth-reference

Parameters

Parameter	Description	Value
bandwidth-reference-value	Specifies the bandwidth reference value used in automatic interface cost calculation.	The value is an integer that ranges from 1 to 2147483648, in Mbit/s.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To make the link cost of IS-IS routes reflect the actual link cost, configure a proper link cost for all interfaces or enable the automatic interface cost calculation

function. To enable automatic interface cost calculation, set a proper bandwidth reference value.

You can run the **ipv6 bandwidth-reference** command to set a proper reference bandwidth value. After the automatic interface cost calculation function is enabled, the system automatically calculates the interface cost according to the bandwidth reference value set using the command.

The **bandwidth** *bandwidth* command can only set an interface bandwidth obtained by the NMS from the MIB. It cannot change an interface actual bandwidth and interface cost.

Precautions

The **ipv6 bandwidth-reference** command takes effect only after IPv6 is enabled for the IS-IS process using the **ipv6 enable** command.

Example

Set the IPv6 reference bandwidth value used in automatic interface cost calculation to 200 Mbit/s.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 bandwidth-reference 200
```

7.7.6 ipv6 bfd all-interfaces enable (IS-IS)

Function

The **ipv6 bfd all-interfaces enable** command enables IPv6 BFD for IS-IS globally in the IS-IS process.

The **undo ipv6 bfd all-interfaces enable** command disables IPv6 BFD for IS-IS globally in the IS-IS process.

By default, IPv6 BFD for IS-IS is disabled in the IS-IS process.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

ipv6 bfd all-interfaces enable

undo ipv6 bfd all-interfaces enable

Parameters

None

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IPv6 BFD for IS-IS can provide millisecond level fault detection, help IS-IS to detect faults on neighboring devices or links rapidly, and instruct IS-IS to recalculate routes for correct packet forwarding.

Prerequisites

The **ipv6 bfd all-interfaces enable** command can take effect only after global BFD is enabled by the **bfd** and IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Configuration Impact

After the **ipv6 bfd all-interfaces enable** command is run, all the interfaces in the IS-IS process are automatically enabled with IPv6 BFD for IS-IS. IS-IS establishes IPv6 BFD sessions on the interfaces whose IS-IS IPv6 neighbor status is Up (the DIS is Up on the broadcast network) by using the default IPv6 BFD parameters.

After the **ipv6 bfd all-interfaces enable** command is run, if you start the IS-IS process on the IS-IS interface that is not enabled with the IS-IS process, the interface is also automatically enabled with IPv6 BFD for IS-IS.

After the **ipv6 bfd all-interfaces enable** command is run, check whether all the interfaces need to be enabled with IPv6 BFD for IS-IS according to network planning. If some interfaces do not need to be enabled with IPv6 BFD for IS-IS, run the **isis ipv6 bfd block** command on these interfaces to block IPv6 BFD for IS-IS.

Precautions

If IPv6 BFD for IS-IS is neither enabled globally nor enabled on interfaces, you can configure IPv6 BFD parameters of IS-IS; however, IPv6 BFD sessions cannot be established.

Example

Configure IPv6 BFD for an IS-IS process.

```
<HUAWEI> system-view
[HUAWEI] bfd
[HUAWEI-bfd] quit
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 bfd all-interfaces enable
```

7.7.7 ipv6 bfd all-interfaces(IS-IS)

Function

The **ipv6 bfd all-interfaces** command configures the parameters of IPv6 BFD sessions.

The **undo ipv6 bfd all-interfaces** command restores the default parameters of IPv6 BFD sessions.

By default, the minimum interval for sending or receiving IPv6 BFD packets is 1000 milliseconds, and the local IPv6 BFD detection multiplier is 3.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

ipv6 bfd all-interfaces { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* } *

undo ipv6 bfd all-interfaces { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* } *

Parameters

Parameter	Description	Value
min-rx-interval *receive-interval*	Specifies the minimum interval for receiving BFD packets from the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *transmit-interval*	Specifies the minimum interval for sending BFD packets to the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier-value*	Indicates the local detection multiplier.	The value is an integer that ranges from 3 to 50. The default value is 3.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IPv6 BFD for IS-IS can provide millisecond level fault detection, help IS-IS to detect faults on neighboring devices or links rapidly, and instruct IS-IS to recalculate routes for correct packet forwarding. The **isis ipv6 bfd** command can be used to configure IPv6 BFD detection parameters on a specified interface.

Prerequisites

The **ipv6 bfd all-interfaces** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Precautions

If only IPv6 BFD session parameters are configured but the **ipv6 bfd all-interfaces enable** command is not run, no IPv6 BFD sessions can be established.

After the local *min-rx-interval* is compared with the remote *min-tx-interval*, the larger value is chosen as the actual local receiving interval.

NOTE

To ensure that the actual local receiving interval is the same as the local configured value, you are recommended to configure the same value for the local *min-rx-interval* and the remote *min-tx-interval*.

Similarly, to ensure that the actual local IPv6 BFD detection multiplier is the same as the local configured value, you are recommended to configure the same value for the local *multiplier-value* and the remote *multiplier-value*.

If no IPv6 BFD packets are received from the peer within the actual local detection time, the IPv6 BFD session status is set to Down. The actual local detection time is calculated with the following formula: Actual local detection time = Actual local receiving interval x Local detection multiplier.

If IPv6 BFD parameters are configured in both an IS-IS process and an IS-IS interface, IPv6 BFD parameters configured on the interface take effect.

Example

Configure IPv6 BFD for an IS-IS process, and specify the minimum interval for sending IPv6 BFD packets to 300 ms.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 bfd all-interfaces enable
[HUAWEI-isis-1] ipv6 bfd all-interface min-tx-interval 300
```


7.7.8 ipv6 circuit-cost

Function

The **ipv6 circuit-cost** command sets the cost value of all IPv6 interfaces in IS-IS SPF calculation, that is, the global IPv6 cost value.

The **undo ipv6 circuit-cost** command deletes the configured global IPv6 cost value.

By default, the global IPv6 cost value is not set.

Format

ipv6 circuit-cost { *cost* | **maximum** } [**level-1** | **level-2**]

undo ipv6 circuit-cost [*cost* | **maximum**] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
cost	Specifies the cost value of IPv6 interfaces in IS-IS SPF calculation.	IF the IS-IS cost style is **wide** or **wide-compatible**, the cost value of imported routes ranges from 1 to 16777214. Otherwise, the value ranges from 1 to 63.
maximum	Sets the link cost value of an interface to 16777215. **NOTE** This parameter can be configured only when the IS-IS cost style is **wide** or **wide-compatible**. When the link cost value of an interface is set to 16777215, the neighbor TLV generated on the link cannot be used for route calculation and can only be used to transmit TE information.	-
level-1	Indicates the cost value of a Level-1 interface. If the interface level is not specified, the cost value is set for Level-1 and Level-2 interfaces.	-
level-2	Indicates the cost value of a Level-2 interface. If the interface level is not specified, the cost value is set for Level-1 and Level-2 interfaces.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, multiple valid routes to the same destination are available on large IPv6 networks. IS-IS calculates the optimal routes using SPF calculation and chooses the optimal routes for traffic forwarding. This process often causes the following two problems:

- All the traffic is forwarded through the optimal route, which may cause unbalanced load.
- If the optimal route on a network is disconnected intermittently, traffic is still forwarded through the optimal route, which causes traffic loss.

To solve the preceding problems, run the **ipv6 circuit-cost** command to change the cost value of all interfaces at a time and traffic can be forwarded through different physical links.

Precautions

The priority of the IPv6 cost value on the interface is higher than the priority of the global IPv6 cost value.

The **ipv6 circuit-cost** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Example

Set the global IPv6 cost value in SPF calculation to 20.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 circuit-cost 20
```

7.7.9 ipv6 circuit default-tag

Function

The **ipv6 circuit default-tag** command sets the administrative tag value of all interfaces in IS-IS processes on an IPv6 network.

The **undo ipv6 circuit default-tag** command restores the default tag value.

By default, the administrative tag value of all interfaces in IS-IS processes on an IPv6 network is 0.

Format

ipv6 circuit default-tag *tag* [**level-1** | **level-2**]

undo ipv6 circuit default-tag [*tag*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
tag	Specifies the administrative tag value of an interface in an IS-IS process on an IPv6 network.	The value is an integer that ranges from 1 to 4294967295.
level-1	Indicates the administrative tag value of all Level-1 interfaces in IS-IS processes on an IPv6 network. If the interface level is not specified, the administrative tag value is set for Level-1 and Level-2 interfaces.	-
level-2	Indicates the administrative tag value of all Level-2 interfaces in IS-IS processes on an IPv6 network. If the interface level is not specified, the administrative tag value is set for Level-1 and Level-2 interfaces.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Using the **ipv6 circuit default-tag** command, you can set the administrative tag value for all routes of a specified IS-IS process and use the tag to filter routes.

The administrative tag value is advertised together with the routing information.

- When the cost style of IS-IS is narrow or narrow-compatible, the administrative tag value is not advertised in the LSP and does not takes effect.
- When the cost style of IS-IS is wide, wide-compatible, or compatible, the administrative tag value is advertised in the LSP.

Precautions

The **ipv6 circuit default-tag** command takes effect only after IPv6 is enabled for the IS-IS process using the **ipv6 enable** command.

Example

Set the administrative tag value for a Level-1 IPv6 interface to 30.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 circuit default-tag 30 level-1
```


7.7.10 ipv6 default-route-advertise

Function

The **ipv6 default-route-advertise** command configures the switch to advertise default IPv6 routes.

The **undo ipv6 default-route-advertise** disables the switch from advertising default IPv6 routes.

By default, the switch does not advertise default IPv6 routes.

Format

ipv6 default-route-advertise [**always** | **match default** | **route-policy** *route-policy-name*] [**cost** *cost* | **tag** *tag* | [**level-1** | **level-2** | **level-1-2**]] * [**avoid-learning**]

undo ipv6 default-route-advertise

Parameters

Parameter	Description	Value
always	Indicates that the switch always advertises default IPv6 routes.	-
match default	Indicates that if the routing table contains a default IPv6 route generated by other routing protocols but not IS-IS, the default route is advertised through an LSP.	-
route-policy *route-policy-name*	Specifies the name of a route-policy. A Level-1-2 device advertises default routes to the IS-IS routing domain only when there are external routes matching the route-policy in the routing table of the device. This prevents routing blackhole when link faults make some important external routes unavailable but default routes are still advertised. This route-policy does not affect external route import in IS-IS.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
cost *cost*	Specifies the cost value of an advertised default route.	The value is an integer that ranges from 0 to 4261412864.
tag *tag*	Specifies the tag value of an advertised default route.	The value is an integer that ranges from 1 to 4294967295.

Parameter	Description	Value
level-1	Sets the level of advertised default IPv6 routes to Level-1. If the level is not specified, Level-2 default routes are generated by default.	-
level-2	Sets the level of advertised default IPv6 routes to Level-2. If the level is not specified, Level-2 default routes are generated by default.	-
level-1-2	Sets the level of advertised default IPv6 routes to Level-1-2. If the level is not specified, Level-2 default routes are generated by default.	-
avoid-learning	Prevents an IS-IS process from learning default routes and adding them to the routing table. If there has been an active default IPv6 route in the routing table, the default IPv6 route becomes inactive after this parameter is specified.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, if other routing protocols are configured on an IS-IS network, to forward traffic in the IS-IS routing domain outside the domain, use the following two methods:

- Configure IS-IS on a Level-1-2 device to advertise default routes into the IS-IS routing domain.
- Configure IS-IS on a Level-1-2 device to import routes of other routing domains into the IS-IS routing domain.

Advertising default routes is easy to configure and does not require an IS-IS process to learn external routes. You can run the **ipv6 default-route-advertise** command to configure IS-IS devices to advertise default routes into the IS-IS routing domain.

Precautions

Using the route-policy, you can force IS-IS to generate default routes only when there are matched routes in the routing table.

- If the **apply isis level-1** command is used in the route-policy view, IS-ISv6 can generate default routes in Level-1 LSPs. If the **apply isis level-1** command is

used in the route-policy view, IS-IS can generate default routes in Level-1 LSPs.

- If the **apply isis level-2** command is used in the route-policy view, IS-ISv6 can generate default routes in Level-2 LSPs.
- If the **apply isis level-1-2** command is used in the route-policy view, IS-ISv6 can generate default routes in Level-1 and Level-2 LSPs.

The **ipv6 default-route-advertise** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, an ABR advertises the default IPv6 route to the IS-IS domain as long as the local routing table contains external routes.

Example

Configure the current switch to generate default IPv6 routes in Level-1-2 LSPs.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 default-route-advertise level-1-2
```

7.7.11 ipv6 enable(IS-IS)

Function

The **ipv6 enable** command enables the IPv6 capability of an IS-IS process.

The **undo ipv6 enable** command disables the IPv6 capability of an IS-IS process.

By default, the IPv6 capability of an IS-IS process is disabled.

Format

ipv6 enable [**topology** { **compatible** [**enable-mt-spf**] | **ipv6** | **standard** }]

undo ipv6 enable

Parameters

Parameter	Description	Value
topology	Specifies the network topology type.	-

Parameter	Description	Value
compatible	Sets the topology type to compatible. That is, the standard topology and the IPv6 topology are compatible. After IS-IS IPv6 adjacency relationship is established, IS-IS advertises IPv6 topology links and standard topology links. SPF, however, runs only in a standard topology. The compatible mode is helpful in a transition from the standard topology to the IPv6 topology.	-
enable-mt-spf	Indicates that SPF runs in the IPv6 topology in compatible mode.	-
ipv6	Sets the topology type to IPv6. That is, the IPv6 capability of an IS-IS process can be enabled in an IPv6 topology. Links on the network can be configured as IPv4 or IPv6 links. SPF calculation is performed separately in IPv4 and IPv6 topologies.	-
standard	Sets the topology type to standard. That is, the IPv6 capability of an IS-IS process can be enabled in the integrated topology. Network administrators must ensure that all links on the network support the same topology mode. If the IPv6 capability of an IS-IS process is enabled, **standard** is used by default.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

To enable the IPv6 capability of an IS-IS process, perform the following operations:

1. Run the **ipv6** command in the system view to enable IPv6 globally.
2. Run the **isis** command to enable an IS-IS process.
3. Run the **network-entity** command to set a NET for IS-IS.
4. Run the **ipv6 enable** command in the IS-IS view to enable the IPv6 capability of an IS-IS process.
5. On each interface that needs to run an IS-IS process, run the **ipv6 enable** command to enable the IPv6 capability and configure an IPv6 address.
6. Run the **isis ipv6 enable** command to enable IS-IS for IPv6 on each interface.

Example

Enable the IPv6 capability in the IPv6 topology in IS-IS process 1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
```



```
[HUAWEI-isis-1] network-entity 10.0001.1010.1020.1030.00
[HUAWEI-isis-1] ipv6 enable topology ipv6
```

7.7.12 ipv6 filter-policy export

Function

The **ipv6 filter-policy export** command configures a filtering policy to allow IS-IS to filter the imported IPv6 routes to be advertised.

The **undo ipv6 filter-policy export** command cancels the filtering function.

By default, IS-IS does not filter imported IPv6 routes to be advertised.

Format

ipv6 filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } **export** [*protocol* [*process-id*] | **unr**]

undo ipv6 filter-policy [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name*] **export** [*protocol* [*process-id*] | **unr**]

Parameters

Parameter	Description	Value
acl6-number	Specifies the number of a basic ACL6.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of an IPv6 named ACL.	The value of *acl6-name* is a string of 1 to 64 case-sensitive characters without spaces.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy to filter routes based on tag and other protocol parameters.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
protocol	Specifies the imported routes that need to be filtered when the routes are advertised. Currently, the *protocol* can be **direct**, **static**, **ripng**, **bgp**, **unr**, **ospfv3**, or another **isis** process. If the *protocol* parameter is not specified, all imported routes are filtered.	-
process-id	Specifies the process ID if *protocol* is **ripng**, **ospfv3**, or another **isis** process.	The value is an integer that ranges from 1 to 65535.
unr	Indicates the imported UNR routes that need to be filtered when the routes are advertised.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ipv6 filter-policy export** command together with the **ipv6 import-route** command can be used to filter only the imported routes that are advertised to other switches. If *protocol* is not specified, the routes that are imported from all other protocols are filtered. If *protocol* is specified, the routes that are imported from the specific protocol are filtered.

The **ipv6 filter-policy export** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Precautions

Creating an ACL6 or IPv6 prefix list before it is referenced is recommended. If a nonexistent ACL6 or IPv6 prefix list is referenced using the command, all external IPv6 routes of the specified routing domain that are imported by IS-IS are advertised to the specified neighbor.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all external IPv6 routes of the specified routing domain that are imported by IS-IS are advertised to the specified neighbor.

Example

Configure IS-IS to filter the imported IPv6 routes using ACL6 2002 before advertising the IPv6 routes to other switches.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 filter-policy 2002 export
```

7.7.13 ipv6 filter-policy import

Function

The **ipv6 filter-policy import** command configures a filtering policy to allow IS-IS to filter the received routes to be added to the IPv6 routing table.

The **undo ipv6 filter-policy import** command cancels the filtering function.

By default, IS-IS does not filter the received routes to be added to the IPv6 routing table.

Format

ipv6 filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } **import**

undo ipv6 filter-policy [*acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name*] **import**

Parameters

Parameter	Description	Value
acl6-number	Specifies the number of a basic ACL6.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of an IPv6 named ACL.	The value of *acl6-name* is a string of 1 to 64 case-sensitive characters without spaces.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a route-policy to filter routes based on tag and other protocol parameters.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IS-IS routing entries need to be added to an IPv6 routing table to guide IPv6 packet forwarding. If an IS-IS routing table has routes destined for a specified network segment, but these routes do not need to be sent to the IPv6 routing table, run the **ipv6 filter-policy import** command and specify an IPv6 prefix list or a route-policy to allow only the required IS-IS routes to be added to the IPv6 routing table.

Precautions

The **ipv6 filter-policy import** command can take effect only after IPv6 is enabled for the IS-IS process by the ipv6 enable command.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Creating an ACL6 or IPv6 prefix list before it is referenced is recommended. If a nonexistent ACL6 or IPv6 prefix list is referenced using the command, all IPv6 routes received by IS-IS are delivered to the IPv6 routing table.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all IPv6 routes received by IS-IS are delivered to the IPv6 routing table.

Example

Filter received IPv6 routes using a route-policy.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 filter-policy route-policy test import
```


7.7.14 ipv6 import-route

Function

The **ipv6 import-route** command configures IS-IS to import IPv6 routes from other routing protocols.

The **undo ipv6 import-route** command disables IS-IS form importing IPv6 routes from other routing protocols.

By default, IS-IS does not import IPv6 routes from other routing protocols.

Format

ipv6 import-route { **direct** | **unr** | { **ospfv3** | **ripng** | **isis** } [*process-id*] | **bgp** [**permit-ibgp**] } **inherit-cost** [**tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

ipv6 import-route { **static** | **direct** | **unr** | { **ospfv3** | **ripng** | **isis** } [*process-id*] | **bgp** [**permit-ibgp**] } [**cost** *cost* | **tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

ipv6 import-route limit *limit-number* [**threshold-alarm upper-limit** *upper-limit-value* **lower-limit** *lower-limit-value*] { **level-1** | **level-2** | **level-1-2** }

undo ipv6 import-route { **direct** | **unr** | { **ospfv3** | **ripng** | **isis** } [*process-id*] | **bgp** [**permit-ibgp**] } **inherit-cost** [**tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

undo ipv6 import-route { **static** | **direct** | **unr** | { **ospfv3** | **ripng** | **isis** } [*process-id*] | **bgp** [**permit-ibgp**] } [**cost** *cost* | **tag** *tag* | **route-policy** *route-policy-name* | [**level-1** | **level-2** | **level-1-2**]] *

undo ipv6 import-route limit [*limit-number*] [**threshold-alarm upper-limit** *upper-limit-value* **lower-limit** *lower-limit-value*] { **level-1** | **level-2** | **level-1-2** }

Parameters

Parameter	Description	Value
direct	Indicates that the imported routes are direct routes.	-
static	Indicates that the imported routes are activated static routes.	-

Parameter	Description	Value
unr	Indicates that the routing protocol from which routes are imported is UNR. The user network route (UNR) is used to allocate routes to user traffic when users cannot use dynamic routing protocols during logins.	-
ospfv3	Indicates that the routing protocol from which routes are imported is OSPFv3.	-
ripng	Indicates that the routing protocol from which routes are imported is RIPng.	-
isis	Indicates that the routing protocol from which routes are imported is IS-IS.	-
process-id	When the protocol is **ospfv3**, **ripng**, or **isis**, a process ID needs to be specified.	The value is an integer that ranges from 1 to 65535.
bgp	Indicates that the routing protocol from which routes are imported is BGP.	-
permit-ibgp	Indicates that the imported routes are IBGP routes. If this parameter is not set, only EBGP routes can be imported.	-
inherit-cost	Indicates that the original cost value of imported external routes is retained.	-
cost *cost*	Specifies the cost value of imported routes.	The value is an integer that ranges from 0 to 4261412864. The default value is 0.
tag *tag*	Specifies the administrative tag of imported routes.	The value is an integer that ranges from 1 to 4294967295.

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
level-1	Indicates that routes are imported into Level-1 routing tables. If no level is specified, routes are imported to Level-2 routing tables by default.	-
level-2	Indicates that routes are imported into Level-2 routing tables. If no level is specified, routes are imported to Level-2 routing tables by default.	-
level-1-2	Indicates that routes are imported into Level-1 and Level-2 routing tables. If no level is specified, routes are imported to Level-2 routing tables by default.	-
limit *limit-number*	Specifies the maximum number of external IPv6 routes allowed to be imported to the IS-IS area.	The value is an integer ranging from 1 to 10000000.
threshold-alarm	Specifies the alarm threshold for imported routes.	-
upper-limit *upper-limit-value*	Specifies the upper alarm threshold for imported routes.	The value is an integer ranging from 1 to 100. The default value is 80.
lower-limit *lower-limit-value*	Specifies the lower alarm threshold for imported routes.	The value is an integer ranging from 1 to 100. The default value is 70.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If other routing protocols are configured on an IS-IS network, to forward traffic in the IS-IS routing domain outside the domain, use either of the following methods:

- Configure IS-IS on a Level-1-2 device to advertise default routes into the IS-IS routing domain.
- Configure IS-IS on a Level-1-2 device to import routes of other routing domains into the IS-IS routing domain.

If there are multiple Level-1-2 devices in the IS-IS routing domain, optimal routes destined for another routing domain need to be selected. This requires all devices in the IS-IS routing domain learn all or some external routes.

Configure IS-IS on a Level-1-2 device to import routes of other routing domains into the IS-IS routing domain. Alternatively, run the **route-policy** *route-policy-name* command to import some external routes from other routing domains.

Prerequisites

IS-IS has been enabled using the **isis** command and the IS-IS view has been displayed. IPv6 has been enabled for the IS-IS process using the **ipv6 enable** command.

Precautions

When routes of other routing protocols are imported, you can set the cost value and cost style of the imported routes. You can also configure IS-IS to retain the original cost value of the imported external routes. During route advertisement and route calculation, the original cost values of these routes are used. In this case, the cost style and cost value of the imported routes cannot be set, and static routes cannot be imported.

After the **ipv6 import-route direct** command is executed, routes to the network segment where the IPv6 address of the management interface belongs are also imported in the IS-IS routing table. Therefore, use this command with caution.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all IPv6 routes of the specified routing domain are imported to the IS-IS routing table.

Example

Configure IS-IS to import static IPv6 routes and set the cost value of the routes to 15.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 import-route static cost 15
```

Configure IS-IS to import OSPFv3 routes and retain the original cost value of the routes.


```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 import-route ospfv3 1 inherit-cost
```

7.7.15 ipv6 import-route isis level-1 into level-2

Function

The **ipv6 import-route isis level-1 into level-2** command configures IPv6 route leaking from Level-1 areas to Level-2 areas.

The **undo ipv6 import-route isis level-1 into level-2** command prohibits IPv6 route leaking from Level-1 areas to Level-2 areas.

By default, all Level-1 routing information, excluding information about default routes, is leaked to Level-2 areas.

Format

ipv6 import-route isis level-1 into level-2 [**filter-policy** { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct allow-filter-policy**] *

undo ipv6 import-route isis level-1 into level-2 [**filter-policy** { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct allow-filter-policy**] *

Parameters

Parameter	Description	Value
filter-policy	Specifies a filtering policy for IPv6 routes.	-
acl6-number	Specifies the number of a basic ACL6.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of an IPv6 named ACL.	The value of *acl6-name* is a string of 1 to 64 case-sensitive characters without spaces.
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list. Only the routes that match the IPv6 prefix can be imported to Level-2 areas.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
tag *tag*	Specifies the administrative tag of imported routes.	The value is an integer that ranges from 1 to 4294967295.
direct allow-filter-policy	Specifies the filtering policy to filter the direct routes. Only the IS-IS Level-1 area direct routing information that matches the filtering policy can be shared with the Level-2 area with this parameter, and all Level-1 area direct routing information will be shared with the Level-2 area without this parameter.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ipv6 import-route isis level-1 into level-2** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Precautions

Creating an ACL6 or IPv6 prefix list before it is referenced is recommended. If a nonexistent ACL6 or IPv6 prefix list is referenced using the command, all IPv6 routes in the Level-1 area leak to the Level-2 area.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and

a nonexistent route-policy is referenced using the current command, all IPv6 routes in the Level-1 area leak to the Level-2 area.

Example

Control IPv6 route leaking from Level-1 areas to Level-2 areas on Level-1-2 routers using IPv6 prefix list **list**.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 import-route isis level-1 into level-2 filter-policy ipv6-prefix list
```

7.7.16 ipv6 import-route isis level-2 into level-1

Function

The **ipv6 import-route isis level-2 into level-1** command configures IPv6 route leaking from Level-2 areas to Level-1 areas.

The **undo ipv6 import-route isis level-2 into level-1** command prohibits IPv6 route leaking from Level-2 areas to Level-1 areas.

By default, Level-2 routing information is not leaked to Level-1 areas.

Format

ipv6 import-route isis level-2 into level-1 [**filter-policy** { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct** { **allow-filter-policy** | **allow-up-down-bit** } *] *

undo ipv6 import-route isis level-2 into level-1 [**filter-policy** { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* | **route-policy** *route-policy-name* } | **tag** *tag* | **direct** { **allow-filter-policy** | **allow-up-down-bit** } *] *

Parameters

Parameter	Description	Value
filter-policy	Specifies a filtering policy for IPv6 routes.	-
acl6-number	Specifies the number of a basic ACL6.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of an IPv6 named ACL. The value is a string of case-sensitive characters.	The value of *acl6-name* is a string of 1 to 64 case-sensitive characters without spaces.

Parameter	Description	Value
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
tag *tag*	Specifies the administrative tag of imported routes.	The value is an integer that ranges from 1 to 4294967295.
direct allow-filter-policy	Specifies the filtering policy to filter the direct routes. Only the IS-IS Level-1 area direct routing information that matches the filtering policy can be shared with the Level-2 area with this parameter, and all Level-1 area direct routing information will be shared with the Level-2 area without this parameter.	-
direct allow-up-down-bit	Indicates that the Up or Down bit is used during the leak of directs routes. If **direct allow-up-down-bit** is specified, the direct routes that have already leaked to the Level-1 area have the lowest priority and cannot leak back.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ipv6 import-route isis level-2 into level-1** command must be configured on Level-1-2 routers at the area border.

The **ipv6 import-route isis level-2 into level-1** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Precautions

Creating an ACL6 or IPv6 prefix list before it is referenced is recommended. If a nonexistent ACL6 or IPv6 prefix list is referenced using the command, all IPv6 routes in the Level-2 area leak to the Level-1 area.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all IPv6 routes in the Level-2 area leak to the Level-1 area.

Example

Configure the switch to perform IPv6 route leaking from a Level-2 area to a Level-1 area using filtering policy 2002.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 import-route isis level-2 into level-1 filter-policy 2002
```

7.7.17 ipv6 maximum load-balancing

Function

The **ipv6 maximum load-balancing** command sets the maximum number of IS-IS IPv6 equal-cost routes for load balancing.

The **undo ipv6 maximum load-balancing** command restores the default number of equal-cost routes for load balancing.

By default, IS-IS supports IPv6 load balancing and the maximum number of IS-IS IPv6 equal-cost routes is 8.

Format

ipv6 maximum load-balancing *number*

undo ipv6 maximum load-balancing [*number*]

Parameters

Parameter	Description	Value
number	Specifies the number of IPv6 equal-cost routes for load balancing.	The value is an integer that ranges from 1 to 8.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

The **ipv6 maximum load-balancing** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Example

Set the maximum number of IS-IS IPv6 equal-cost routes for load balancing to 2.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 100
[HUAWEI-isis-100] ipv6 enable
[HUAWEI-isis-100] ipv6 maximum load-balancing 2
```

7.7.18 ipv6 preference

Function

The **ipv6 preference** command configures the preference of IPv6 routes generated by the IS-IS protocol.

The **undo ipv6 preference** command restores the default reference of IPv6 routes generated by the IS-IS protocol.

By default, the preference of IPv6 routes generated by the IS-IS protocol is 15.

Format

ipv6 preference { **route-policy** *route-policy-name* | *preference* } *

undo ipv6 preference

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
preference	Specifies the preference of IPv6 routes generated by the IS-IS. A smaller value indicates a higher priority.	The value is an integer that ranges from 1 to 255.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The routing policy can be used to set the preference for the filtered routes. If the preference is set, only the routes that pass the route-policy can apply the preference specified by this command; otherwise, routes apply the default value of the preference.

Multiple dynamic routing protocols can be run on a switch at the same time. To properly select routes, the system sets a default preference for each routing protocol. If different protocols find routes to the same destination, the route of the protocol with a higher priority is preferred.

The **ipv6 preference** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Precautions

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the configured priority applies to all IS-IS IPv6 routes.

Example

Set the IS-IS protocol preference to 20.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
```



```
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 preference 20
```

7.7.19 ipv6 prefix-priority (IS-IS)

Function

The **ipv6 prefix-priority** command configures the convergence priority of IS-IS IPv6 routes.

The **undo ipv6 prefix-priority** command restores the default convergence priority of IS-IS IPv6 routes.

By default, the convergence priority of IS-IS IPv6 host routes and default routes is medium, and the convergence priority of other IS-IS IPv6 routes is low.

Format

ipv6 prefix-priority [**level-1** | **level-2**] { **critical** | **high** | **medium** } { **ipv6-prefix** *prefix-name* | **tag** *tag-value* }

undo ipv6 prefix-priority [**level-1** | **level-2**] { **critical** | **high** | **medium** }

Parameters

Parameter	Description	Value
level-1	Specifies the convergence priority of Level-1 IS-IS IPv6 routes.	-
level-2	Specifies the convergence priority of Level-2 IS-IS IPv6 routes.	-
critical	Sets the convergence priority of IS-IS IPv6 routes to critical.	-
high	Sets the convergence priority of IS-IS IPv6 routes to high.	-
medium	Sets the convergence priority of IS-IS IPv6 routes to medium.	-
ipv6-prefix *prefix-name*	Specifies the prefix name of IPv6 addresses for filtering routes to set the convergence priority for IS-IS IPv6 routes that match the specified IPv6 address prefix.	The value of *prefix-name* is a string of 1 to 169 case-sensitive characters without spaces.

Parameter	Description	Value
tag *tag-value*	Specifies the tag value used for filtering routes to set the convergence priority for IS-IS IPv6 routes that carry the specified tag value. To specify the tag value to filter the IPv6 routes that need to be set with the convergence priority, ensure that IS-IS IPv6 routes carry the tag value.	The value is an integer that ranges from 1 to 4294967295.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

This command takes effect only in IPv6 base topologies but not in VPNs.

The convergence priorities of IS-IS IPv6 routes are classified as critical, high, medium, and low in descending order.

After the **ipv6 prefix-priority** command is run to set the convergence priority for IS-IS IPv6 routes, the following situations occur:

- The convergence priority of existing IS-IS IPv6 routes is re-set according to the configuration of the **ipv6 prefix-priority** command.
- The convergence priority of new IS-IS IPv6 routes is set according to the filtering result of the **ipv6 prefix-priority** command.
- If an IS-IS IPv6 route meets the matching rules specified in multiple commands that are used to configure convergence priorities, this IS-IS IPv6 route is of top convergence priority among the set convergence priorities.
- no level is specified, the convergence priority is set for both Level-1 routes and Level-2 routes according to the configuration of the **ipv6 prefix-priority** command.

NOTE

The **ipv6 prefix-priority** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

After the **ipv6 prefix-priority** command is run to set the convergence priority for IS-IS IPv6 routes (including IS-IS IPv6 host routes and default routes), the convergence priority of all the IS-IS IPv6 routes that meet the matching rules is changed according to the configuration of the **ipv6 prefix-priority** command, and the convergence priority of the IS-IS IPv6 routes that do not meet the matching rules is changed to low.

Example

In an IPv6 base topology, set the convergence priority of IS-IS Level-1 routes with tag value 3 to critical.


```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 prefix-priority level-1 critical tag 3
```

7.7.20 ipv6 spf-priority

Function

The **ipv6 spf-priority** command configures a priority for SPF calculation in an IPv6 base topology.

The **undo ipv6 spf-priority** command restores the default priority for SPF calculation in an IPv6 base topology.

By default, the priority of SPF calculation in an IPv6 base topology is 64.

Format

ipv6 spf-priority *priority-value*

undo ipv6 spf-priority

Parameters

Parameter	Description	Value
priority-value	Specifies the priority for SPF calculation in an IPv6 base topology. A larger value indicates a higher priority.	The value is an integer that ranges from 1 to 127.

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

A higher priority of SPF calculation in an IPv6 base topology means that the routes in the IPv6 base topology are converged preferentially, which ensures the proper operation of important services.

The **ipv6 spf-priority** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Example

Set the priority of SPF calculation in an IPv6 base topology to 30.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
```



```
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 spf-priority 30
```

7.7.21 ipv6 summary

Function

The **ipv6 summary** command configures IS-IS to generate IPv6 summarized routes.

The **undo ipv6 summary** command disables IS-IS from generating IPv6 summarized routes.

By default, IS-IS does not generate IPv6 summarized routes.

Format

ipv6 summary *ipv6-address prefix-length* [**avoid-feedback** | **generate_null0_route** | **tag** *tag* | [**level-1** | **level-1-2** | **level-2**]] *

undo ipv6 summary *ipv6-address prefix-length* [**level-1** | **level-1-2** | **level-2**]

Parameters

Parameter	Description	Value
ipv6-address	Specifies the address range of IPv6 routes to be summarized.	-
prefix-length	Specifies the prefix length of an IPv6 route.	The value is an integer that ranges from 0 to 128.
avoid-feedback	Avoids learning summarized routes through SPF calculation.	-
generate_null0_route	Generates NULL0 routes to avoid routing loop.	-
tag *tag*	Specifies the administrative tag value.	The value is an integer that ranges from 1 to 4294967295.
level-1	Indicates that only routes in Level-1 areas are summarized. If the level is not specified, **level-2** is used by default.	-
level-1-2	Indicates that routes in Level-1 and Level-2 areas are summarized. If the level is not specified, **level-2** is used by default.	-

Parameter	Description	Value
level-2	Indicates that only routes in Level-2 areas are summarized. If the level is not specified, **level-2** is used by default.	-

Views

IS-IS view

Default Level

2: Configuration level

Usage Guidelines

The routes with the same prefix can be aggregated into one route. This can reduce the scale of the routing table, the size of the LSPs generated by the local switch, and the scale of the LSDB. The aggregated routes can be the routes discovered by IS-IS or the imported IPv6 routes. After the aggregation, the smallest cost among those of the routes that are aggregated is used as the cost of the aggregated IPv6 route.

The **ipv6 summary** command can take effect only after IPv6 is enabled for the IS-IS process by the **ipv6 enable** command.

Example

Set a summarized route fc00::/32.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] ipv6 summary fc00:: 32
```

7.7.22 isis ipv6 bfd

Function

The **isis ipv6 bfd** command configures IPv6 BFD detection parameters on a specified interface.

The **undo isis ipv6 bfd** command restores the default values of IPv6 BFD detection parameters on a specified interface.

By default, the minimum interval for sending or receiving IPv6 BFD packets is 1000 milliseconds, and the local IPv6 BFD detection multiplier is 3.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis ipv6 bfd { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* } *

undo isis ipv6 bfd { **min-rx-interval** [*receive-interval*] | **min-tx-interval** [*transmit-interval*] | **detect-multiplier** [*multiplier-value*] } *

Parameters

Parameter	Description	Value
min-rx-interval *receive-interval*	Specifies the minimum interval for receiving BFD packets from the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-tx-interval *transmit-interval*	Specifies the minimum interval for sending BFD packets to the peer.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier-value*	Indicates the local detection multiplier.	The value is an integer that ranges from 3 to 50. The default value is 3.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IPv6 BFD for IS-IS can provide millisecond level fault detection, help IS-IS to detect faults on neighboring devices or links rapidly, and instruct IS-IS to recalculate routes for correct packet forwarding. The **isis ipv6 bfd** command can be used to configure IPv6 BFD detection parameters on a specified interface.

Prerequisites

Before running the **isis ipv6 bfd** command, you need to run the **ipv6 enable** command in the IS-IS process and configure IS-IS on the interface.

The configured IPv6 BFD detection parameters take effect on an interface only when the **isis ipv6 bfd enable** command is configured on the interface.

Precautions

After the local *receive-interval* is compared with the remote *transmit-interval*, the larger value is chosen as the actual local receiving interval. If no IPv6 BFD packets are received from the peer within the actual local detection time, the IPv6 BFD session status is set to Down. The actual local detection time is calculated with the following formula: Actual local detection time = Actual local receiving interval x Local detection multiplier.

NOTE

To ensure that the actual local receiving interval is the same as the local configured value, you are recommended to configure the same value for the local *min-rx-interval* and the remote *min-tx-interval*.

Similarly, to ensure that the actual local IPv6 BFD detection multiplier is the same as the local configured value, you are recommended to configure the same value for the local *multiplier-value* and the remote *multiplier-value*.

The **isis ipv6 bfd** command takes priority over the **ipv6 bfd all-interfaces** command.

Example

Enable IPv6 BFD for IS-IS on Vlanif10, and specify the minimum interval for receiving IPv6 BFD packets to 400 ms and the local detection multiplier to 4.

```
<HUAWEI> system-view
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] isis ipv6 bfd enable
[HUAWEI-Vlanif10] isis ipv6 bfd min-rx-interval 400 detect-multiplier 4
```

7.7.23 isis ipv6 bfd block

Function

The **isis ipv6 bfd block** command disables an interface from dynamically establishing an IPv6 BFD session.

The **undo isis ipv6 bfd block** command cancels the configuration.

By default, no interface is disabled from dynamically establishing a BFD session.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis ipv6 bfd block

undo isis ipv6 bfd block

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If the **ipv6 bfd all-interfaces enable** command is run in an IS-IS process, IPv6 BFD for IS-IS is enabled on all IS-IS interfaces that are associated with this IS-IS process. You can run the **isis ipv6 bfd block** command to disable an interface from dynamically establishing an IPv6 BFD session.

Prerequisites

To make the **isis ipv6 bfd block** command take effect, do as follows first:

1. Enable IPv6 globally in the system view by using the **ipv6** command.
2. Enable an IS-IS process by using the **isis** command.
3. Enable IPv6 capability of the IS-IS process by using the **ipv6 enable** command in the IS-IS view.
4. On each interface that runs an IS-IS process, enable IPv6 capability by using the **ipv6 enable** command.
5. Enable IS-IS on each interface by using the **isis ipv6 enable** command.

Precautions

The **isis ipv6 bfd block** command and the **isis ipv6 bfd enable** command are mutually exclusive and cannot be run on the same interface.

The function of the **isis ipv6 bfd block** command is similar to the function of the **undo isis ipv6 bfd enable** command, except that:

- After the **undo isis ipv6 bfd enable** command is configured on an IS-IS interface, if the **ipv6 bfd all-interfaces enable** command is reconfigured in the IS-IS view, the interface still has IPv6 BFD for IS-IS capabilities.

- After the **isis ipv6 bfd block** command is configured on an IS-IS interface, the interface does not have IPv6 BFD for IS-IS capabilities, even if the **ipv6 bfd all-interfaces enable** command is reconfigured in the IS-IS view.

NOTE

- To enable IPv6 BFD for IS-IS on most IS-IS interfaces, you need to configure the **isis ipv6 bfd block** command on the IS-IS interfaces that do not need to be enabled with IPv6 BFD for IS-IS, and then configure the **ipv6 bfd all-interfaces enable** command in the IS-IS view.
- To enable IPv6 BFD for IS-IS only on a few IS-IS interfaces, configure the **isis ipv6 bfd enable** command on these interfaces. In this situation, you can configure the **undo isis ipv6 bfd enable** or **isis ipv6 bfd block** command to disable IPv6 BFD for IS-IS on these interfaces.

Example

Disable Vlanif10 from dynamically establishing an IPv6 BFD session.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 bfd block
```

7.7.24 isis ipv6 bfd enable

Function

The **isis ipv6 bfd enable** command enables IPv6 BFD for IS-IS on the IS-IS interface and uses the default parameters to establish an IPv6 BFD session.

The **undo isis ipv6 bfd enable** command disables IPv6 BFD for IS-IS on the IS-IS interface.

By default, if IPv6 BFD for IS-IS is enabled in the IS-IS view, the IS-IS interface is also enabled with IPv6 BFD for IS-IS. If IPv6 BFD for IS-IS is not enabled in the IS-IS view, the IS-IS interface is not enabled with IPv6 BFD for IS-IS either.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

isis ipv6 bfd enable

undo isis ipv6 bfd enable

Parameters

None

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

IPv6 BFD for IS-IS can provide millisecond level fault detection, help IS-IS to detect faults on neighboring devices or links rapidly, and instruct IS-IS to recalculate routes for correct packet forwarding. The **isis ipv6 bfd enable** command can be used to enable IPv6 BFD for IS-IS on a specified IS-IS interface.

If IPv6 BFD for IS-IS is not enabled globally, you can run the **isis ipv6 bfd enable** command on an interface to enable IPv6 BFD for IS-IS on this interface.

Prerequisites

To make the **isis ipv6 bfd enable** command take effect, do as follows:

1. Enable IPv6 globally in the system view by using the **ipv6** command.
2. Enable an IS-IS process by using the **isis** command.
3. Enable IPv6 capability of the IS-IS process by using the **ipv6 enable** command in the IS-IS view.
4. On each interface that runs an IS-IS process, enable IPv6 capability by using the **ipv6 enable** command.
5. Enable IS-IS on each interface by using the **isis ipv6 enable** command.

Example

Enable IPv6 BFD on Vlanif10.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 bfd enable
```

7.7.25 isis ipv6 cost

Function

The **isis ipv6 cost** command sets the link cost value of an IS-IS interface in an IPv6 topology.

The **undo isis ipv6 cost** command restores the default link cost value of an IS-IS interface in an IPv6 topology.

By default, the link cost value of an IS-IS interface in an IPv6 topology is 10.

Format

isis ipv6 cost { *cost* | **maximum** } [**level-1** | **level-2**]

undo isis ipv6 cost [*cost* | **maximum**] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
cost	Specifies the IPv6 link cost value of an interface.	The value is an integer that varies according to the cost style. ● When the cost style is **narrow**, **narrow-compatible** or **compatible**, the value ranges from 1 to 63. ● When the cost style is **wide** or **wide-compatible**, the value ranges from 1 to 16777214. The default value is 10. To set the cost style, use the **cost-style** command.
maximum	Sets the link cost value of an interface to 16777215. **NOTE** This parameter can be configured only when the IS-IS cost style is **wide** or **wide-compatible**. When the link cost value of an interface is set to 16777215, the neighbor TLV generated on the link cannot be used for route calculation and can only be used to transmit TE information.	-
level-1	Specifies the link cost value of a Level-1 interface. If the interface level is not specified, link cost values of Level-1 and Level-2 interfaces are set.	-
level-2	Specifies the link cost value of a Level-2 interface. If the interface level is not specified, link cost values of Level-1 and Level-2 interfaces are set.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a large-scale IPv6 network, there may be multiple valid routes to the same destination. IS-IS uses the SPF algorithm to calculate an optimal route and transmits traffic over it, which brings the following problems:

- All traffic is transmitted over the optimal route, causing load imbalance.
- If the optimal route is faulty, traffic will get lost.

To solve the preceding problems, run the **isis ipv6 cost** command to set IPv6 link costs for interfaces so that traffic can be transmitted over different physical links.

Prerequisites

IS-IS has been enabled on a specified interface using the **isis ipv6 enable** [*process-id*] command in the interface view.

Configuration Impact

If the link cost of an interface is changed, routes will be re-calculated on the whole network, causing the changes in traffic forwarding paths.

The priority of the **ipv6 circuit-cost** command is lower than that of the **isis ipv6 cost** command.

If the IPv6 topology type is **compatible** or **standard**, the IPv6 and IPv4 link costs are the same on an interface even though the IPv6 link cost set using the **isis ipv6 cost** command is different from the IPv4 link cost. The IPv6 link cost set using the **isis ipv6 cost** command takes effect only when the IPv6 topology type is **compatible enable-mt-spf** or **ipv6**.

NOTE

The IPv6 topology type can be set for an IS-IS process using the **ipv6 enable(IS-IS)** command.

Example

Set the IPv6 link cost of VLANIF 10 to 50.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 cost 50
```


Set the IPv6 link cost of GE0/0/1 to 50.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 cost 50
```

7.7.26 isis ipv6 enable

Function

The **isis ipv6 enable** command enables the IS-IS IPv6 capability for an interface and specifies the ID of the IS-IS process to be associated with the interface.

The **undo isis ipv6 enable** command disables the IS-IS IPv6 capability of an interface and disassociates an IS-IS process from the interface.

By default, the IS-IS IPv6 capability is disabled on an interface.

Format

isis ipv6 enable [*process-id*]

undo isis ipv6 enable

Parameters

Parameter	Description	Value
process-id	Specifies an IS-IS process ID.	The value is an integer that ranges from 1 to 65535. The default value is 1.

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After completing the configuration of an IS-IS process in the system view, enable IS-IS IPv6 on the interface that runs IS-IS and associate the interface with the IS-IS process using the **isis ipv6 enable** command.

Prerequisites

The following steps have been performed:

1. Enable IPv6 globally using the **ipv6** command in the system view.
2. Enable an IS-IS process using the **isis** command and configure a network entity title (NET) for the device running IS-IS using the **network-entity** command.
3. Enable IPv6 for the IS-IS process using the **ipv6 enable (IS-IS)** command in the IS-IS view.
4. Enable IPv6 and configure an IPv6 address using the **ipv6 enable** command on each interface that runs the IS-IS process.

Precautions

An interface can be associated with only one IS-IS process.

NOTE

To perform IPv6-related IS-IS configurations, you must enable IPv6 first.

Example

Create IS-IS process 1, enable the IPv6 capability, and activate the IPv6 capability on VLANIF 10.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] network-entity 10.0001.1010.1020.1030.00
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] ipv6 address fc00:0:0:2::1/64
[HUAWEI-Vlanif10] isis ipv6 enable 1
```

Create IS-IS process 1, enable the IPv6 capability, and activate the IPv6 capability on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] network-entity 10.0001.1010.1020.1030.00
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] ipv6 address fc00:0:0:2::1/64
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 enable 1
```

7.7.27 isis ipv6 suppress-reachability

Function

The **isis ipv6 suppress-reachability** command suppresses the advertisement of direct routes on an IPv6 interface in a specified topology.

The **undo isis ipv6 suppress-reachability** command restores the default configuration of an IPv6 interface.

By default, the advertisement of IPv6 addresses on an interface is not suppressed.

Format

isis ipv6 suppress-reachability [**level-1** | **level-1-2** | **level-2**]

undo isis ipv6 suppress-reachability

Parameters

Parameter	Description	Value
level-1	Suppresses the advertisement of IPv6 addresses is suppressed on Level-1 interfaces. If the level is not specified, Level-1-2 is used by default.	-
level-1-2	Indicates that the advertisement of IPv6 addresses is suppressed on Level-1 and Level-2 interfaces. If the level is not specified, Level-1-2 is used by default.	-
level-2	Indicates that the advertisement of IPv6 addresses is suppressed on Level-2 interfaces. If the level is not specified, Level-1-2 is used by default.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

You must enable the IPv6 capability for an IS-IS process before running this command. For details, see the **isis ipv6 enable** command.

Example

Suppress the advertisement of IPv6 addresses on VLANIF 10.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 10
[HUAWEI-Vlanif10] ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 enable
[HUAWEI-Vlanif10] isis ipv6 suppress-reachability
```

Suppress the advertisement of IPv6 addresses on GE0/0/1.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
```



```
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 suppress-reachability
```

7.7.28 isis ipv6 tag-value

Function

The **isis ipv6 tag-value** command configures the IPv6 administrative tag value on an IS-IS interface.

The **undo ipv6 isis tag-value** command deletes the IPv6 administrative tag value on an IS-IS interface.

By default, an IS-IS interface has no IPv6 administrative tag value.

Format

isis ipv6 tag-value *tag* [**level-1** | **level-2**]

undo isis ipv6 tag-value [*tag*] [**level-1** | **level-2**]

Parameters

Parameter	Description	Value
tag	Specifies the administrative tag value.	The value is an integer that ranges from 1 to 4294967295.
level-1	Indicates the administrative tag value of a Level-1 interface. If the interface level is not specified, the administrative tag value is set for Level-1 and Level-2 interfaces.	-
level-2	Indicates the administrative tag value of a Level-2 interface. If the interface level is not specified, the administrative tag value is set for Level-1 and Level-2 interfaces.	-

Views

Interface view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Using the **isis ipv6 tag-value** command, you can set the administrative tag value for all routes of a specified IS-IS process. The tag can be used as a filtering condition of a route-policy to filter routes.

Precautions

You must enable the IPv6 capability for an IS-IS process before running this command.

If the IS-IS cost style is wide, wide-compatible, or compatible, the administrative tag for an interface is advertised using an LSP.

Example

Set the IPv6 administrative tag value of VLANIF100 to 77.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface vlanif 100
[HUAWEI-Vlanif100] ipv6 enable
[HUAWEI-Vlanif100] isis ipv6 enable
[HUAWEI-Vlanif100] isis ipv6 tag-value 77
```

Set the IPv6 administrative tag value of GE0/0/1 tp 77.

```
<HUAWEI> system-view
[HUAWEI] ipv6
[HUAWEI] isis 1
[HUAWEI-isis-1] ipv6 enable
[HUAWEI-isis-1] quit
[HUAWEI] interface gigabitethernet 0/0/1
[HUAWEI-GigabitEthernet0/0/1] undo portswitch
[HUAWEI-GigabitEthernet0/0/1] ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 enable
[HUAWEI-GigabitEthernet0/0/1] isis ipv6 tag-value 77
```

7.8 BGP Configuration Commands

7.8.1 Command Support

Only the following switch models support BGP:

S5720I-SI, S500, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S

NOTE

The S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, and S6720S-S support only the BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, and BGP-IPv6 unicast address family view.

The S5731-S, S5731S-S, S6730-S, and S6730S-S support only the BGP-IPv4 multicast address family view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, and BGP-VPN instance IPv6 address family view.

7.8.2 active-route-advertise

Function

The **active-route-advertise** command enables BGP to advertise only the preferred routes in the IP routing table.

The **undo active-route-advertise** command restores the default setting.

By default, BGP advertises all preferred routes in the BGP routing table to neighbors.

Format

active-route-advertise

undo active-route-advertise

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, BGP advertises all preferred routes in the BGP routing table to neighbors. After the **active-route-advertise** is configured, only the routes preferred by BGP and also active at the routing management layer are advertised to neighbors.

Precautions

The **active-route-advertise** command and the **routing-table rib-only** command are mutually exclusive.

Example

Enable BGP to advertise only the preferred routes in the IP routing table to neighbors.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] active-route-advertise
```


7.8.3 aggregate (BGP)

Function

The **aggregate** command creates a summary route in the BGP routing table.

The **undo aggregate** command deletes a summary route from the BGP routing table.

By default, no summary route is created in the BGP routing table.

Format

aggregate *ipv4-address* { *mask* | *mask-length* } [**as-set** | **attribute-policy** *route-policy-name1* | **detail-suppressed** | **origin-policy** *route-policy-name2* | **suppress-policy** *route-policy-name3*] *

aggregate *ipv6-address prefix-length* [**as-set** | **attribute-policy** *route-policy-name1* | **detail-suppressed** | **origin-policy** *route-policy-name2* | **suppress-policy** *route-policy-name3*] *

undo aggregate *ipv4-address* { *mask* | *mask-length* } [**as-set** | **attribute-policy** *route-policy-name1* | **detail-suppressed** | **origin-policy** *route-policy-name2* | **suppress-policy** *route-policy-name3*] *

undo aggregate *ipv6-address prefix-length* [**as-set** | **attribute-policy** *route-policy-name1* | **detail-suppressed** | **origin-policy** *route-policy-name2* | **suppress-policy** *route-policy-name3*] *

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a summary route.	The address is in dotted decimal notation.
mask	Specifies the network mask of a summary route.	The mask is in dotted decimal notation.
mask-length	Specifies the network mask length of a summary route.	The value is an integer that ranges from 0 to 32.
ipv6-address	Specifies the IPv6 address of a summary route.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the length of the prefix of an IPv6 summary route.	The value is an integer that ranges from 0 to 128.

Parameter	Description	Value
as-set	Generates a route with the AS-SET.	-
attribute-policy *route-policy-name1*	Specifies the name of an attribute policy for summary routes.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
detail-suppressed	Advertises only the summary route.	-
origin-policy *route-policy-name2*	Specifies the name of a policy that allows route summarization.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
suppress-policy *route-policy-name3*	Specifies the name of a policy for suppressing the advertisement of specific routes.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP route summarization is classified into manual summarization and automatic summarization. The **aggregate** command can be used to configure manual

summarization. This command can summarize routes in the local BGP routing table. Manually-summarized routes take precedence over automatically-summarized routes. The outbound interface of the summary route on the local switch is NULL0. When receiving the summary route, other switches automatically add the outbound interface.

The generated summary route inherits the Origin attribute from the specific routes if they have the same Origin attribute. If the specific routes have different Origin attributes, the summary route selects one as its own from these Origin attributes in descending order of preference: incomplete > egp > igp. If specific routes are not suppressed, the summary route carries the community attribute of every specific route.

The **aggregate** command adds a summary route to the BGP routing table.

- The parameter **as-set** is used to create a summary route whose AS_Path attribute contains AS_Path information of specific routes. Exercise caution when using this parameter if many AS_Path attributes need to be summarized because frequent changes in routes may cause route flapping.
- The parameter **detail-suppressed** is used to suppress the advertisement of specific routes. After **detail-suppressed** is configured, only summary routes are advertised. Summary routes carry the atomic-aggregate attribute, not the community attributes of specific routes.
- The parameter **suppress-policy** is used to suppress the advertisement of specific routes. The **if-match** clause of **route-policy** can be used to filter the routes to be suppressed. This means that only the routes matching the policy will be suppressed, and the other routes will still be advertised. The **peer route-policy** command can also be used to filter out the routes not to be advertised to peers.
- After the parameter **origin-policy** is used, only the routes matching **route-policy** are summarized.
- The parameter **attribute-policy** is used to set attributes for a summary route. If the AS_Path attribute is set in the policy using the **apply as-path** command and **as-set** is set in the **aggregate** command, the AS_Path attribute in the policy does not take effect. The **peer route-policy** command can also be used to set attributes for a summary route.

Prerequisites

Before the **aggregate** command is run, BGP must be enabled.

Configuration Impact

Summary routes are generated after you run the **aggregate** command. If **detail-suppressed** is configured in the command, the advertisement of specific routes will be suppressed. If **suppress-policy** is configured in the command, the advertisement of specific routes that match the policy will be suppressed.

Precautions

When the **undo aggregate** command is run, the system matches routes based on the configured parameter **attribute-policy**, **origin-policy**, **suppress-policy**, **as-set**, or **detail-suppressed**. If none of these parameters is configured, the **undo aggregate** command will fail to be executed.

Example

Create a summary route. The path that is used to advertise this route is an AS-SET consisting of all summarized paths.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] aggregate 172.16.0.0 255.255.0.0 as-set
```

7.8.4 as-notation plain

Function

The **as-notation plain** command configures a BGP 4-byte AS number to be displayed as an integer.

The **undo as-notation plain** command configures a BGP 4-byte AS number to be displayed in dotted notation.

By default, a BGP 4-byte AS number is displayed in dotted notation (that is, in the format of x.y).

Format

as-notation plain

undo as-notation plain

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, display commands such as the **display bgp peer** command display 4-byte AS numbers in dotted notation.

To display a 4-byte AS number as an integer, run the **as-notation plain** command. After the **as-notation plain** command is run, display commands display all 4-byte AS numbers as integers. These 4-byte AS numbers include:

- Independent 4-byte AS number
- 4-byte AS number in an extended community attribute
- 4-byte AS number in a route distinguisher (RD)

- 4-byte AS number in a VPN target (also called a route target)
- 4-byte AS number in the Site-of-Origin attribute

Assume that a 4-byte AS number in dotted notation is *x.y*. Following is the conversion relationship between an integral 4-byte AS number and a 4-byte AS number in dotted notation:

Integral 4-byte AS number = *x* x 65536 + *y*

For example, if a 4-byte AS number in dotted notation is 2.3, the corresponding integral 4-byte AS number is 131075 (2 x 65536 + 3).

Precautions

After the **as-notation plain** command is run, the formats of 4-byte AS numbers in configuration information generated by the system do not change.

- If integral 4-byte AS numbers are configured, 4-byte AS numbers in configuration information generated by the system are also displayed as integers.
- If 4-byte AS numbers in dotted notation are configured, 4-byte AS numbers in configuration information generated by the system are also displayed in dotted notation.

Changing the format of 4-byte AS numbers will affect matching results of AS_Path regular expressions and extended community attribute filters. Therefore, if the system is using an AS_Path regular expression or an extended community attribute filter as an import or export policy, you must reconfigure an AS_Path regular expression using the **ip as-path-filter** command or an extended community attribute filter using the **ip extcommunity-filter** command after changing the format of 4-byte AS numbers. This reconfiguration ensures that routes match the import or export policy.

- If integral 4-byte AS numbers are configured, you must change 4-byte AS numbers in AS_Path regular expressions and extended community attribute filters to integral 4-byte AS numbers.
- If 4-byte AS numbers in dotted notation are configured, you must change 4-byte AS numbers in AS_Path regular expressions and extended community attribute filters to 4-byte AS numbers in dotted notation.

Example

Configure a BGP 4-byte AS number to be displayed as an integer.

```
<HUAWEI> system-view
[HUAWEI] as-notation plain
Warning: If the configuration takes effect, the regular expression of the filter for 4-byte AS path should be changed to the asplain format, otherwise match will fail. Continue? [Y/N]y
```

Configure a BGP 4-byte AS number to be displayed in dotted notation.

```
<HUAWEI> system-view
[HUAWEI] undo as-notation plain
Warning: If the configuration takes effect, the regular expression of the filter for 4-byte AS path should be changed to the asdot format, otherwise match will fail. Continue? [Y/N]y
```


7.8.5 as-path-limit

Function

The **as-path-limit** command sets the maximum number of AS numbers in the AS_Path attribute.

The **undo as-path-limit** command restores the default setting.

By default, no limit is configured on the maximum number of AS numbers in the AS_Path attribute, but the maximum number of AS numbers carried in the AS_Path attribute is limited by the BGP message length.

Format

as-path-limit [*as-path-limit-num*]

undo as-path-limit

Parameters

Parameter	Description	Value
as-path-limit-num	Specifies the maximum number of AS numbers in the AS-Path attribute.	The value is an integer that ranges from 1 to 2000. **NOTE** • The maximum value of *as-path-limit-num* for the 2-byte and 4-byte AS numbers is the same. • If *as-path-limit-num* is not specified, the maximum number of AS numbers in the AS_Path attribute is 255.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

This command is used to restrict the maximum number of AS numbers in the AS_Path attribute only when the AS_Path attribute is reconstructed or the aggregated routes are generated.

Precautions

After the **as-path-limit** command is configured, a switch checks whether the number of AS numbers in the AS-Path attribute of each received or advertised

route exceeds the maximum value. If the number of AS numbers exceeds the maximum value, the local switch discards the route. Therefore, if the maximum number of AS numbers in the AS-Path attribute is set too small, routes are lost.

Example

Set the maximum number of AS numbers in the AS-Path attribute to 200.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] as-path-limit 200
```

7.8.6 auto-frr

Function

The **auto-frr** command enables BGP Auto FRR.

The **undo auto-frr** command disables BGP Auto FRR.

By default, BGP Auto FRR is disabled.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

auto-frr

undo auto-frr

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

This function is applicable to networks that require a low packet loss rate and a short delay.

Using BGP Auto FRR together with BFD is recommended. They can rapidly detect a link fault and switch traffic to a standby link if a fault occurs.

Precautions

If both the **ip frr** command and the **auto-frr** command are configured, the **ip frr** command takes precedence over the **auto-frr** command. If a route fails to match the routing policy specified in the **ip frr** command, the **auto-frr** command takes effect.

NOTE

On the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the **auto-frr** command can implement fast active/standby link switchovers of the forwarding plane and fast route convergence of the control plane. On other switch models, this command can only implement fast route convergence of the control plane.

Example

Enable BGP Auto FRR for unicast routes.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] auto-frr
```

7.8.7 bestroute as-path-ignore

Function

The **bestroute as-path-ignore** command configures BGP to ignore the AS_Path attribute when it selects the optimal route.

The **undo bestroute as-path-ignore** command restores the default configuration.

By default, BGP uses the AS_Path attribute as one of route selection rules.

Format

bestroute as-path-ignore

undo bestroute as-path-ignore

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **bestroute as-path-ignore** command is used, BGP does not compare the AS path lengths of routes. By default, the route with the shortest AS path is preferred.

Precautions

After the **bestroute as-path-ignore** command is run, the AS_Path attribute is not used as one of the BGP route selection rules.

Example

Configure BGP to ignore the AS_Path attribute when selecting the optimal route.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] bestroute as-path-ignore
```

7.8.8 bestroute igp-metric-ignore

Function

The **bestroute igp-metric-ignore** command configures BGP to ignore the metric value of the next-hop IGP route when selecting the optimal route.

The **undo bestroute igp-metric-ignore** command restores the default setting.

By default, BGP selects a route with the smallest metric value of the next-hop IGP route as the optimal route.

Format

bestroute igp-metric-ignore

undo bestroute igp-metric-ignore

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a BGP network, the BGP device always receives multiple routes with the same prefix but to different paths from neighbors. BGP must select the optimal route to

the prefix to guide packet forwarding. By default, BGP compares the next-hop IGP route metric values of these routes and selects the route with the smallest metric value as the optimal route.

To customize route selection policies, you can run the **bestroute igp-metric-ignore** command to configure BGP to ignore the metric value of the next-hop IGP route when selecting the optimal route.

Example

Configure BGP to ignore the IGP metric when selecting the optimal route.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] bestroute igp-metric-ignore
```

7.8.9 bestroute l2vpn-preference vpls

Function

The **bestroute l2vpn-preference vpls** command configures BGP to compare the L2VPN preferences of routes when selecting the optimal route.

The **undo bestroute l2vpn-preference vpls** command restores the default setting.

By default, BGP does not compare the L2VPN preferences when selecting the optimal route.

Format

bestroute l2vpn-preference vpls

undo bestroute l2vpn-preference vpls

Parameters

None

Views

BGP-L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP often receives routes with the same prefix but different next hops from multiple peers. In this case, BGP needs to select the optimal route and deliver it to the routing table to guide packet forwarding.

To configure BGP to compare the L2VPN preferences of routes when selecting the optimal route, run the **bestroute l2vpn-preference vpls** command.

Precautions

This command is valid only for Kompella VPLS.

After this command is configured, routes with the same prefix are preferentially selected.

When you run the **multi-homing-preference** command to configure multi-homing preferences for VSIs, ensure that the preference configured on the local device and that configured on the remote device are different. Otherwise, the **bestroute l2vpn-preference vpls** configuration does not take effect.

Example

Configure BGP to compare the L2VPN preferences of routes when selecting the optimal route.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] l2vpn-ad-family
[HUAWEI-bgp-af-l2vpn-ad] bestroute l2vpn-preference vpls
```

7.8.10 bestroute med-confederation (BGP)

Function

The **bestroute med-confederation** command enables BGP to compare the Multi Exit Discriminator (MED) values of routes in a confederation when BGP selects the optimal route.

The **undo bestroute med-confederation** command restores the default settings.

By default, BGP compares the MED values of the routes that are from the same AS only.

Format

bestroute med-confederation

undo bestroute med-confederation

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-IPv6 unicast address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, BGP compares only MED values of routes from the same AS, not including sub-ASs in a confederation. To enable BGP to compare MED values of routes in a confederation when selecting the optimal route, run the **bestroute med-confederation** command.

After the **bestroute med-confederation** command is configured, BGP compares MED values only when AS_Path does not contain the external AS (AS that is not in the confederation) number. Otherwise, BGP does not compare MED values.

For example, ASs 65000, 65001, 65002, and 65004 belong the same confederation. Routes to the same destination are listed as follows:

- path1: AS_Path=65000 65004, med=2
- path2: AS_Path=65001 65004, med=3
- path3: AS_Path=65002 65004, med=4
- path4: AS_Path=65003 65004, med=1

After the **bestroute med-confederation** command is run, the AS_Path attributes of paths 1, 2, and 3 does not contain the numbers of ASs that belong to other confederations. Therefore, when selecting routes based on MED values, BGP compares the MED values of paths 1, 2, and 3 only. This is because the AS_Path attribute of path 4 contains the number of an AS that does not belong to this confederation.

Example

Configure BGP to compare the MED values of routes only in the confederation when selecting the optimal route.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] bestroute med-confederation
```

7.8.11 bestroute med-none-as-maximum

Function

The **bestroute med-none-as-maximum** command configures BGP to assign the maximum Multi Exit Discriminator (MED), 4294967295, to a route without an MED in route selection.

The **undo bestroute med-none-as-maximum** command restores the default configuration.

By default, BGP assigns 0 to a route without an MED in route selection.

Format

bestroute med-none-as-maximum

undo bestroute med-none-as-maximum

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **bestroute med-none-as-maximum** command takes effect in BGP route selection and is used only when the MED is null in the route attributes. If the MED is null and the **bestroute med-none-as-maximum** command is not run, the system probably cannot select an optimal route.

Configuration Impact

During BGP route selection, if the **bestroute med-none-as-maximum** command is run, a route without any MED is assigned the maximum MED of 4294967295.

Example

Configure BGP to assign the maximum MED of 4294967295 to a route without an MED in route selection.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] bestroute med-none-as-maximum
```

7.8.12 bestroute nexthop-resolved

Function

The **bestroute nexthop-resolved** command configures the condition of the route selection of labeled BGP IPv4 routes, VPNv4 routes, or VPNv6 routes for next hop recursion.

The **undo bestroute nexthop-resolved** command restores the default configuration.

By default, labeled BGP IPv4 routes, VPNv4 routes, or VPNv6 routes participate in route selection only when their next hops recurse to IP addresses.

Product	Support
S5731-S, S5731S-S, S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

bestroute nexthop-resolved { **ip** | **tunnel** }

undo bestroute nexthop-resolved

Parameters

Parameter	Description	Value
ip	Allows labeled routes that recurse to IP addresses to participate in route selection.	-
tunnel	Allows labeled routes that recurse to MPLS tunnels to participate in route selection.	-

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In Route Reflector (RR) networking where the next hop of a labeled BGP IPv4 route is changed or non-RR networking, BGP needs to check whether there is an LSP destined for the next hop. If the LSP is not found, BGP needs to establish an LSP in advance. If the default setting is adopted, a labeled BGP IPv4 unicast route may be selected prior to LSP establishment. As a result, services are incorrectly switched before LSP establishment and service stability is affected. The **bestroute nexthop-resolved tunnel** command can be run to allow route selection only after the labeled BGP IPv4 route recurses to an LSP. This ensures service stability.

By default, a VPNv4 route or VPNv6 routes can participate in route selection if its next hop recurses to an IP address. The system adds the optimal VPNv4 route or VPNv6 route to the forwarding table for traffic forwarding. However, traffic cannot

be forwarded over the route if the route does not recurse to an LSP. In this situation, run the **bestroute nexthop-resolved tunnel** command to enable VPNv4 routes or VPNv6 routes to be available in route selection only when their next hops recurse to LSPs to ensure uninterrupted traffic forwarding.

Precautions

The **bestroute nexthop-resolved ip** and **bestroute nexthop-resolved tunnel** commands are mutually exclusive.

Example

In the BGP VPNv4 view, configure BGP VPNv4 routes that recurse to LSP tunnels to participate in route selection.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family vpnv4
[HUAWEI-bgp-af-vpnv4] bestroute nexthop-resolved tunnel
```

In the BGP VPNv6 view, configure BGP VPNv6 routes that recurse to LSP tunnels to participate in route selection.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv6-family vpnv6
[HUAWEI-bgp-af-vpnv6] bestroute nexthop-resolved tunnel
```

7.8.13 bestroute routerid-prior-clusterlist

Function

The **bestroute routerid-prior-clusterlist** command enables Originator_ID to take precedence over Cluster_List during BGP route selection.

The **undo bestroute routerid-prior-clusterlist** command restores the default configurations.

By default, Cluster_List takes precedence over Originator_ID during BGP route selection.

Format

bestroute routerid-prior-clusterlist

undo bestroute routerid-prior-clusterlist

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view, BGP-IPv4 multicast address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-L2VPN address family view

Default Level

2: Configuration level

Usage Guidelines

On a BGP network, after a device receives multiple routes with the same prefix but different paths from different peers, the device needs to select an optimal route from these routes to forward packets. By default, Cluster_List takes precedence over Originator_ID during BGP route selection. To enable Originator_ID to take precedence over Cluster_List during BGP route selection, run the **bestroute routerid-prior-clusterlist** command.

Example

Enable Router ID to take precedence over Cluster_List during BGP route selection.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] bestroute routerid-prior-clusterlist
```

7.8.14 bgp

Function

The **bgp** command enables BGP and displays the BGP view.

The **undo bgp** command disables BGP.

By default, BGP is disabled.

Format

bgp { *as-number-plain* | *as-number-dot* }

undo bgp [*as-number-plain* | *as-number-dot*]

Parameters

Parameter	Description	Value
as-number-plain	Specifies the number of the AS, in integer format.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of the AS, in dotted notation.	The value is in the x.y format. Here, "x" and "y" are integers that range from 1 to 65535 and 0 to 65535 respectively.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP is an inter-AS dynamic routing protocol. BGP running within an AS is called Internal BGP (IBGP), and BGP running between ASs is called External BGP (EBGP).

Whether to apply BGP to routing information transmission between ASs depends on the following conditions:

- If at least one of the following conditions is met, BGP can be used:
  - An AS allows data packets to pass through on their way to another AS.
  - Multiple connections to external ISPs and the Internet exist in ASs.
  - Data flows entering or leaving ASs must be controlled.
- If one of the following conditions is met, BGP does not need to be used:
  - Users are connected to only one ISP network.
  - The ISP does not need to provide Internet access services for users.
  - ASs adopt default routes between each other.

Precautions

- After the **bgp** command is run, BGP is enabled.
- Each device runs in only one AS; therefore, each device can be specified with only one local AS number.
- If the BGP AS number is changed, the route calculation result of the local or remote OSPF VPN process may be affected. Because the tag value of the LSAs of the OSPF VPN process is calculated based on the BGP AS number, the local tag value of the OSPF VPN process and the tag value carried in the LSA advertised by the OSPF VPN process change. After receiving an LSA, the local or remote OSPF process checks the tag of the LSA. If the local tag of the OSPF process is the same as the tag of the received LSA, the LSA is not used for OSPF route calculation.

NOTICE

After the **undo bgp** [*as-number-plain* | *as-number-dot*] command is run, BGP services may be interrupted, and all BGP configurations on the device are cleared. Therefore, exercise caution when you run this command.

Example

Enable BGP and enter the BGP view.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp]
```


7.8.15 bgp dynamic-session-limit

Function

The **bgp dynamic-session-limit** command configures a maximum number for dynamic BGP peer sessions.

The **undo bgp dynamic-session-limit** command restores the default configuration.

By default, a maximum of 100 dynamic BGP peer sessions can be established after the dynamic BGP peer function is enabled.

Format

bgp dynamic-session-limit *limit-value*

undo bgp dynamic-session-limit [*limit-value*]

Parameters

Parameter	Description	Value
limit-value	Specifies the maximum number of dynamic BGP peer sessions allowed.	The value is an integer that ranges from 1 to 256.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the dynamic BGP peer function is enabled using the **peer listen-net** command, BGP listens to BGP connection requests from the network segment specified in the command and establish BGP peer relationships dynamically. If a large number of dynamic BGP peer sessions are established, excessive system resources will be consumed. To prevent this problem, run the **bgp dynamic-session-limit** command to configure a maximum number for dynamic BGP peer sessions as required.

Precautions

This command does not apply to static BGP peer sessions.

Example

Configure the maximum number of dynamic BGP peer sessions to 128.


```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] bgp dynamic-session-limit 128
```

7.8.16 bgp fast-refresh enable

Function

The **bgp fast-refresh enable** command enables the BGP Prefix Independent Convergence (PIC) function.

The **undo bgp fast-refresh enable** command disables the BGP PIC function.

By default, the BGP PIC function is enabled on the switch.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

bgp fast-refresh enable

undo bgp fast-refresh enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When there are primary and backup paths for a BGP route, and a fault occurs on the primary path, the route must be updated and traffic must be switched to the

backup path as soon as possible. This process is called BGP route convergence. As the network expands in scale and route prefixes, the BGP route convergence speed is affected by the number of route prefixes. More route prefixes indicate slower route convergence speed. If a network fault occurs, many packets are lost due to the slow route convergence speed. This can have a large impact on some services that are sensitive to packet loss. BGP PIC is also called BGP fast refresh. It indicates that the BGP route convergence speed is independent of the number of route prefixes. After this function is enabled, the BGP route convergence speed becomes faster without being affected by the number of route prefixes. This function reduces data loss during route convergence. When this function becomes unavailable, you can run the **undo bgp fast-refresh enable** command to disable it. BGP routes are then converged as usual; however, the route convergence speed is slow.

Precautions

- Only IBGP routes support the BGP PIC function, and VPN routes do not support this function.
- If the BGP load balancing (ECMP) or BGP Auto FRR function has been configured on the device or the next hop of a BGP route recurses to a tunnel interface, the BGP PIC function does not take effect.
- Disabling the BGP PIC function requires the device to be restarted after the configuration is saved or requires all BGP ASs to be **reset**. Otherwise, disabling the BGP PIC function takes effect only for the subsequently generated routing tables.

  Enabling the BGP PIC function requires the device to be restarted.

Example

Disable the BGP PIC function on the switch.

```
<HUAWEI> system-view
[HUAWEI] undo bgp fast-refresh enable
Warning: The operation will take effect only on routing tables generated later.
Please save configuration and reboot this device or reset all BGP ASs after this
 operation.
```

7.8.17 check-first-as

Function

The **check-first-as** command enables the function to check the first AS number in the AS_Path list that is carried in the Update message sent by the EBGP peer.

The **undo check-first-as** command disables the function.

By default, the function is enabled.

Format

check-first-as

undo check-first-as

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, BGP checks the first AS number in the AS_Path list that is carried in the Update message sent by the EBGP peer. If only the first AS number indicates the AS where the EBGP peer locates, the Update message is accepted. Otherwise, the Update message is denied, and the EBGP connection goes Down.

Precautions

The **check-first-as** command is not listed in the configuration file.

After the **undo check-first-as** command is configured, loops have a greater chance to occur. Therefore, use the command with caution.

Follow-up Procedure

After the configuration is complete, run the **refresh bgp** command if you want to check the received routes again.

Example

Check the first AS number in the AS_Path list that is carried in the Update message sent by the EBGP peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] check-first-as
```

7.8.18 compare-different-as-med

Function

The **compare-different-as-med** command enables BGP to compare the MEDs in the routes of peers in different ASs.

The **undo compare-different-as-med** command restores the default configuration.

By default, BGP does not compare the MEDs in the routes of peers in different ASs.

Format

compare-different-as-med

undo compare-different-as-med

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The command is used to control the MEDs to change the policy of BGP route selection. If the **compare-different-as-med** command is run, BGP will compare the MEDs of the routes from different ASs. If there are multiple reachable paths to the same destination, BGP prefers the route with the smallest MED.

Precautions

Do not use this command unless different ASs use the same IGP and route selection mode.

Example

Enable BGP to compare the MEDs in the routes of peers in different ASs.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] compare-different-as-med
```

7.8.19 confederation id

Function

The **confederation id** command configures a BGP confederation and specifies a confederation ID for the BGP confederation.

The **undo confederation id** command removes the specified BGP confederation.

By default, no BGP confederation is configured.

Format

confederation id { *as-number-plain* | *as-number-dot* }

undo confederation id

Parameters

Parameter	Description	Value
as-number-plain	Specifies the number of the AS, in integer format.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of the AS, in dotted notation.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A large AS may contain a huge number of fully meshed IBGP peer relationships. In this situation, configure a confederation.

Divide the AS into multiple small sub-ASs and group them into a confederation. Then establish EBGP connections between sub-ASs and establish fully meshed IBGP connections within the same sub-AS. Some key attributes of routes, such as the next hop, MED, and local preference are not discarded when these routes pass through sub-ASs. This can reduce the number of fully meshed IBGP peer relationships in an AS and keep the integrity of the original AS.

Precautions

The confederation ID is equal to the AS number. An external AS must specify the confederation ID when specifying the AS number of the peer. All the sub-ASs in the same confederation must be configured with the same confederation ID that must be different from the number of any sub-AS.

Example

Configure a confederation ID. An AS is divided into sub-ASs 65001, 65002, 65003, and 65004, and their confederation ID is 9. Peer 10.2.3.4 is a member of the AS confederation. Peer 10.4.5.6 is a member outside the AS confederation. For the external members, confederation 9 is a complete AS.

```
<HUAWEI> system-view
[HUAWEI] bgp 65001
[HUAWEI-bgp] confederation id 9
[HUAWEI-bgp] confederation peer-as 65002 65003 65004
[HUAWEI-bgp] peer 10.2.3.4 as-number 65002
```



```
[HUAWEI-bgp] peer 10.4.5.6 as-number 65005
```

7.8.20 confederation nonstandard

Function

The **confederation nonstandard** command configures standard devices (in RFC 3065) in a confederation to communicate with nonstandard devices.

The **undo confederation nonstandard** command configures standard devices in a confederation to communicate only with standard devices.

By default, only standard devices in a confederation can communicate with each other.

Format

confederation nonstandard

undo confederation nonstandard

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To communicate with nonstandard devices, configure this command on all switches in a confederation.

Precautions

After the confederation ID is configured, running the **confederation nonstandard** command will disconnect the sessions between a router and its IBGP peers as well as its confederation EBGP peers. Then, new connections are reestablished.

Example

Enable the switch to communicate with nonstandard routers. AS 100 contains two sub-ASs, 64000 and 65000.

```
<HUAWEI> system-view
[HUAWEI] bgp 64000
[HUAWEI-bgp] confederation id 100
[HUAWEI-bgp] confederation peer-as 65000
[HUAWEI-bgp] confederation nonstandard
```


7.8.21 confederation peer-as

Function

The **confederation peer-as** command configures the number of each sub-AS of the same confederation.

The **undo confederation peer-as** command removes the specified sub-AS from the confederation.

By default, no sub-AS number of a confederation is configured.

Format

confederation peer-as { *as-number-plain* | *as-number-dot* } &<1-32>

undo confederation peer-as { *as-number-plain* | *as-number-dot* } &<1-32>

Parameters

Parameter	Description	Value
as-number-plain	Specifies the number of an AS, which is an integer.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of an AS, which is in dotted notation.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A large AS may contain a huge number of fully meshed IBGP peer relationships. In this situation, configure a confederation.

Prerequisites

The **confederation id** command specifies the confederation ID of each sub-AS. If the confederation ID is not configured, this configuration is invalid.

Precautions

The sub-ASs configured in this command belong to the same confederation, and each sub-AS uses fully meshed network.

Example

Configure sub-ASs for a confederation.

```
<HUAWEI> system-view
[HUAWEI] bgp 1090
[HUAWEI-bgp] confederation id 100
[HUAWEI-bgp] confederation peer-as 1091 1092 1093
```

7.8.22 confederation route unicast-to-label disable

Function

The **confederation route unicast-to-label disable** command prevents non-RRs from advertising the IPv6 non-labeled routes learned from peers in the local confederation as labeled routes carrying a label of all Fs to the BGP peers in another confederation.

The **undo confederation route unicast-to-label disable** command restores the default configuration.

By default, non-RRs can advertise the IPv6 non-labeled routes learned from peers in the local confederation as labeled routes carrying a label of all Fs to the BGP peers in another confederation.

Product	Support
S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5731-S, S5731S-S, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

confederation route unicast-to-label disable

undo confederation route unicast-to-label disable

Parameters

None

Views

BGP-IPv6 unicast address family view

Default Level

2: Configuration level

Usage Guidelines

If two PEs that are not RRs reside in two confederations, the PEs can advertise labeled routes carrying a label of all Fs to each other by default. However, non-Huawei devices may consider the label of all Fs invalid, and if these non-Huawei devices receive such labeled routes, BGP peer relationships may be disconnected. To prevent this problem, run the **confederation route unicast-to-label disable** command to prevent the PEs from advertising labeled routes carrying a label of all Fs.

Example

Prevent non-RRs from advertising the IPv6 non-labeled routes learned from peers in the local confederation as labeled routes carrying a label of all Fs to the BGP peers in another confederation.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] confederation id 100
[HUAWEI-bgp] confederation peer-as 38 39
[HUAWEI-bgp] peer 192.168.1.1 as-number 38
[HUAWEI-bgp] peer fc00:0:0:191::1 as-number 39
[HUAWEI-bgp] ipv6-family unicast
[HUAWEI-bgp-af-ipv6] peer 192.168.1.1 enable
[HUAWEI-bgp-af-ipv6] peer 192.168.1.1 label-route-capability
[HUAWEI-bgp-af-ipv6] peer fc00:0:0:191::1 enable
[HUAWEI-bgp-af-ipv6] confederation route unicast-to-label disable
```

7.8.23 dampening

Function

The **dampening** command enables BGP route flap suppression and modifies BGP route flap suppression parameters.

The **undo dampening** command restores the default configuration.

By default, BGP route flap suppression is disabled.

Format

dampening [ibgp] [*half-life-reach reuse suppress ceiling* | **route-policy** *route-policy-name*] *

undo dampening [ibgp]

Parameters

Parameter	Description	Value
half-life-reach	Specifies the half life of a reachable route.	The value is an integer that ranges from 1 to 45, in minutes. The default value is 15.
reuse	Specifies the threshold for the route to be unsuppressed. If the penalty of the route falls below the threshold, the route is reused.	The value is an integer that ranges from 1 to 20000. The default value is 750.
suppress	Specifies the threshold for the route to be suppressed. If the penalty value of the route exceeds the threshold, the route is not used.	The value is an integer that ranges from 1 to 20000, which must be greater than the value of *reuse*. The default value is 2000.
ceiling	Specifies the penalty ceiling.	The value is an integer that ranges from 1001 to 20000. The configured value must be greater than that of *suppress*. The default value is 16000.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ibgp	Specifies the route type as IBGP route. If this parameter is not specified, the route type is EBGP route. **NOTE** This parameter takes effect only in the BGP-VPNv4 address family view.	-

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view, BGP-IPv4 multicast address family view, BGP-VPNv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP route flap dampening measures the stability of a route by using a penalty value. The greater the penalty value, the less stable the route. Each time route flapping occurs, that is, when the switch receives a Withdraw packet or an Update packet for a route, BGP adds a certain penalty value (1000) for the route.

When the penalty value of the route exceeds the suppression threshold, the route is suppressed. The switch does not add the route to the IP routing table or advertise any Update packet to other BGP peers. If the route is marked with the d flag, the last packet received by the switch is an Update packet; if the route is marked with the h flag, the last packet received by the switch is a Withdraw packet. After the penalty value reaches a certain limit, it does not increase any more. The limit is called the penalty ceiling.

After the route is suppressed for a certain period, the penalty value is reduced by half. If the penalty value of a route marked with the d flag decreases to the reuse threshold, the d flag is removed, and the route becomes available and is selected preferentially. After that, the route can be added to the IP routing table and used to send Update packets to other BGP peers. If the penalty value of a route marked with the h flag decreases to 0, the route is deleted from the BGP routing table.

You can run the **display bgp routing-table label** command to check the d flag and h flag indicating the routes that have been dampened and were dampened respectively.

Precautions

If the **dampening** command is run multiple times, the latest configuration overrides the previous one.

After the **dampening** command is run, the system suppresses unstable routes. This means that the system does not add unstable routes to the BGP routing table or advertise them to other BGP peers.

Note the following items when configuring BGP route flap dampening:

- The value of *suppress* must be greater than that of *reuse* and smaller than that of *ceiling*.
- If MaxSuppressTime obtained by using the formula of MaxSuppressTime = *half-life-reach* x 60 x (ln(*ceiling*/*reuse*)/ln(2)) is smaller than 1, suppression cannot be performed. You need to ensure that the value of MaxSuppressTime is equal to or greater than 1. This means that the value of *ceiling*/*reuse* must be great enough.

NOTE

The **dampening** command is valid only for EBGP routes.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-**

policy nonexistent-config-check disable command is run in the system view and a nonexistent route-policy is referenced using the current command, the configured dampening parameters apply to all routes; if no dampening parameters are configured, the default dampening parameters apply to the routes.

Example

Enable EBGP route dampening and modify EBGP route damping parameters.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] dampening 10 1000 2000 5000
```

7.8.24 default ipv4-unicast

Function

The **default ipv4-unicast** command enables the IPv4 unicast address family for BGP peers by default.

The **undo default ipv4-unicast** command disables the IPv4 unicast address family for BGP peers by default.

By default, the IPv4 unicast address family is enabled for BGP peers.

Format

default ipv4-unicast

undo default ipv4-unicast

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If the IPv4 unicast address family needs to be enabled by default for created BGP peers, the **default ipv4-unicast** command can be used to configure the default address family of BGP to the IPv4 unicast address family.

If the IPv4 unicast address family does not need to be enabled by default for created BGP peers, the **undo default ipv4-unicast** command can be used to disable the IPv4 unicast address family for all peers.

Precautions

After the **undo default ipv4-unicast** command is run, the **peer enable** command needs to be run if the created BGP peer needs to be enabled with the IPv4 unicast address family.

Example

Enable the IPv4 unicast address family for all peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] default ipv4-unicast
```

7.8.25 default local-preference

Function

The **default local-preference** command sets the default local preference for BGP routes.

The **undo default local-preference** command restores the default configuration.

By default, the local preference for BGP routes is 100.

Format

default local-preference *local-preference*

undo default local-preference

Parameters

Parameter	Description	Value
local-preference	Specifies the local preference for BGP routes. The greater the value, the higher the preference.	The value is an integer that ranges from 0 to 4294967295.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The Local_Pref attribute is used to determine an optimal route for traffic before traffic leaves an AS. When the switch that runs BGP has multiple routes to the

same destination, the switch selects the route with the highest local preference as the optimal route.

Precautions

If the switch is already configured with a default local preference for BGP routes, the configuration of a new default local preference will override the previous configuration.

The local preference is exchanged only between IBGP peers and is not advertised to other ASs.

Example

Set the default local preference for BGP routes to 200.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] default local-preference 200
```

7.8.26 default med

Function

The **default med** command sets the default multi-exit-discriminator (MED) for BGP routes.

The **undo default med** command restores the default configuration.

By default, the MED is 0.

Format

default med *med*

undo default med

Parameters

Parameter	Description	Value
med	Specifies the MED for BGP routes.	The value is an integer that ranges from 0 to 4294967295.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **default med** command is valid only for the routes imported by using the **import-route (BGP)** command and the routes summarized by BGP on the local device.

The **default med** command sets the default MED and controls route selection for the traffic entering an AS. If the other attributes of routes to the same destination are the same, BGP will select the route with the smallest MED as the optimal route.

Precautions

If a default MED is configured on a device, configuring a new default MED will override the previous configuration and the new default MED will overwrite the previous one.

The MED is transmitted between two neighboring ASs only. Devices in an AS do not advertise the received MED to peers in other ASs.

Example

Set the default MED of a BGP route to 10.

```
<HUAWEI> system-view
[HUAWEI] bgp 1
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] default med 10
```

7.8.27 default-route imported

Function

The **default-route imported** command enables the import of default routes in the local IP routing table to the BGP routing table.

The **undo default-route imported** command disables the import of the default routes in the local IP routing table to the BGP routing table.

By default, BGP does not add the default route to the BGP routing table.

Format

default-route imported

undo default-route imported

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

This configuration reduces the number of routes on the network and minimizes the routing table size. If a default route exists in the local IP routing table and the **default-route imported** command is run, the system imports the default route to the BGP routing table. After the default route is imported, the packets can be transmitted over the default route when no matched routing entry is found in the routing table, preventing packet loss.

Precautions

To import default routes, you need to run both the **default-route imported** command and the **import-route (BGP)** command. If only the **import-route (BGP)** command is used, default routes cannot be imported. In addition, the **default-route imported** command is used to import only the default routes that exist in the local routing table.

When a device needs to advertise default routes to a peer (group) and no default route exists in the local routing table, the **peer default-route-advertise** command needs to be used.

Example

Import default routes to the BGP routing table.

```
<HUAWEI> system-view
[HUAWEI] bgp 1
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] default-route imported
[HUAWEI-bgp-af-ipv4] import-route ospf 1
```

7.8.28 deterministic-med (BGP)

Function

The **deterministic-med** command enables the BGP deterministic-MED function so that routes with the same leftmost AS number are first compared during route selection.

The **undo deterministic-med** command disables the BGP deterministic-MED function so that routes are compared against each other according to the sequence in which they are received.

By default, the BGP deterministic-MED function is disabled.

Format

deterministic-med

undo deterministic-med

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

After the BGP deterministic-MED function is enabled, when an optimal route is to be selected from routes that are received from different ASs and carry the same prefix, routes are first grouped based on the leftmost AS number in the AS_Path attribute. Routes with the same leftmost AS number are grouped together, and after comparison, an optimal route is selected for the group. The group optimal route is then compared with optimal routes from other groups to determine the final optimal route. This mode of route selection ensures that the sequence in which routes are received is irrelevant to the route selection result.

If the BGP deterministic-MED function is disabled, routes are compared against each other according to the sequence in which they are received. In this manner, the sequence in which routes are received is relevant to the result of route selection.

For example: Assume that the following BGP routes are available on the switch.

- Route A1: AS(PATH) 12, med 100, igp metric 13, internal, rid 4.4.4.4
- Route A2: AS(PATH) 12, med 150, igp metric 11, internal, rid 5.5.5.5
- Route B: AS(PATH) 3, med 0, igp metric 12, internal, rid 6.6.6.6

If Route A1, Route A2, and Route B are received in turn, Route A1 and Route A2 are first compared. The leftmost AS number of Route A1 is the same as the leftmost AS number of Route A2, and therefore Route A1 is selected because its MED is smaller. After that, Route A1 and Route B are compared. Because the leftmost AS numbers of the two routes are different, the optimal route cannot be selected by comparing the MEDs of the two routes unless the **compare-different-as-med** command is configured. As a result, Route B is selected because its IGP metric is smaller.

If Route A2, Route B, and Route A1 are received in turn, Route A2 and Route B are first compared. Because leftmost AS number of Route A2 is different from the leftmost AS number of Route B, the optimal route cannot be selected by comparing the MEDs of the two routes unless the **compare-different-as-med** command is configured. As a result, Route A2 is selected because its IGP metric is smaller. After that, Route A2 and Route A1 are compared. The leftmost AS number of Route A1 is the same as the leftmost AS number of Route A2, and therefore Route A1 is selected because its MED is smaller.

Judging from the preceding route selection procedure, when the BGP deterministic-MED function is disabled, the sequence in which routes are received

is relevant to the result of route selection. After the BGP deterministic-MED function is enabled, the sequence in which routes are received is no longer relevant to the result of route selection. Route A1 and Route A2 have the same leftmost AS number, Route A1 and Route A2 are compared first regardless of the sequence in which routes are received.

Example

Enable the deterministic-MED function in the BGP view.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] deterministic-med
```

7.8.29 display bgp bfd session

Function

The **display bgp bfd session** command displays information about BFD sessions between BGP peers.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp bfd session { [**vpnv4 vpn-instance** *vpn-instance-name*] **peer** *ipv4-address* | **all** }

Parameters

Parameter	Description	Value
vpnv4 vpn-instance *vpn-instance-name*	Displays information about the BFD session between BGP peers with the specified IPv4 VPN instance name.	The value must be an existing VPN instance name.
peer *ipv4-address*	Displays information about the BFD session of the BGP peer with the specified IPv4 address.	The value is in dotted decimal notation.
all	Displays all BFD sessions between BGP peers.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The BFD session set up by BGP helps BGP quickly discover faults between BGP peers and informs BGP to recalculate routes for packet forwarding.

Run the **display bgp bfd session** command to display information about the BFD session set up by BGP in the following situations:

- Check whether the BFD session is successfully set up or view detailed information about the BFD session that is successfully configured.
- Check whether the BFD session set up by BGP is successfully deleted after running the **undo peer bfd enable** command.
- Verify the configuration after running the **undo peer bfd block** command to prevent a peer from inheriting the BFD function of the peer group.
- Verify the configuration after running the **peer bfd** command to set BFD parameters.

The information about the BFD session of a specified BGP peer can be displayed by specifying different parameters.

- Run the **display bgp bfd session vpnv4 vpn-instance** *vpn-instance-name* **peer** *ipv4-address* command to display information about the BFD session of a specified BGP peer in a specified VRF.
- Run the **display bgp bfd session peer** *ipv4-address* command to display information about the BFD session of a specified BGP peer on the public network.
- Run the **display bgp bfd session all** command to display information about the BFD sessions of all BGP peers.

Prerequisites

The BFD session has been set up using the **peer bfd enable** command. If the BFD session has not been set up by BGP, no information is displayed after running the **display bgp bfd session** command.

Example

Display all BFD sessions between BGP peers.

```
<HUAWEI> display bgp bfd session all
  Local_Address      Peer_Address       LD/RD       Interface
  10.1.1.2           10.1.1.1           8192/8193   Unknown
  Tx-interval(ms)    Rx-interval(ms)    Multiplier  Session-State
  1000               1000               3           Up
  Wtr-interval(m)
0
```

Table 7-122 Description of the display bgp bfd session command output

Item	Description
Local_Address	Local address
Peer_Address	Peer address

Item	Description
LD/RD	Local/remote discriminator
Interface	Interface on which the BFD session is set up **NOTE** Information about the interface on which the BFD session is set up only when the directly connected interface is used to set up the EBGP neighbor relationship. In other cases, information about the interface is displayed as Unknown.
Tx-interval (ms)	Interval for sending BFD packets, in milliseconds
Rx-interval (ms)	Interval for receiving BFD packets, in milliseconds
Multiplier	Local detection multiple
Session-State	BFD status • Admin down: The BFD session is closed on the local end. • BFD global disable: BFD is disabled globally. • BFD session number exceed: The number of BFD sessions exceeds the maximum limit. • Detect down: BFD detects a link status fault and interrupts the connection. • Init: The BFD session is in the initialized state. • Neighbor down: The peer end detects that the BFD session goes Down and informs the local end of the change, and the local end then sets the neighbor status to Down. • Receive admin down: The BFD session is closed on the peer end (for example, the BFD session is disabled on the peer end). • Up: The BFD session is set up.
Wtr-interval(m)	Interval for flap dampening, in minutes

7.8.30 display bgp error

Function

The **display bgp error** command displays BGP errors.

Format

display bgp error

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

When a BGP fault occurs, run the **display bgp error** command to display BGP error information. BGP error information includes peer error information, route error information, and resource threshold-crossing error information.

Precautions

If no BGP fault occurs, no information is displayed after running the **display bgp error** command.

Example

Display BGP errors.

```
<HUAWEI> display bgp error
 Error Type    : Peer Error
 Date/Time     : 2010/03/22 11:40:39 UTC-08:00
 Peer Address  : 10.1.1.2
 VRF Name      : Public
 Error Info    : Router-ID conflict

 Error Type    : Peer Error
 Date/Time     : 2010/03/22 11:40:39 UTC-08:00
 Peer Address  : 10.1.1.2
 VRF Name      : Public
 Error Info    : Incorrect remote AS

 Error Type    : Route Error
 AddressFamily : IPv4-UNC
 InstanceID    : 0
 Discard count : 20

 Error Type    : Resource exceed limit
 Date/Time     : 2010/03/22 11:40:39 UTC-08:00
 Limit info    : Route number limit

 Error Type    : Resource exceed limit
 Date/Time     : 2010/03/22 11:40:39 UTC-08:00
 Limit info    : Label number limit
```


Table 7-123 Description of the display bgp error command output

Item	Description
Error Type	Error type: • Peer Error: indicates neighbor errors. • Route Error: indicates route errors. • Resource exceed limit: indicates that resources exceed the limit.
Date/Time	Date and Time when an error occurs
Peer Address	Address of a peer
VRF Name	VPN Instance name
Error Info	Error information: • Router-ID conflict: indicates that router IDs conflict. • Incorrect remote AS: indicates an incorrect remote AS number.
AddressFamily	Address family
Discard count	Number of discarded routes
Limit info	Information indicating that resources exceed the limit: • Memory shortage: indicates that memory exceeds the limit. • Route number limit: indicates that the number of routes exceeds the limit. • Label number limit: indicates that the number of labels exceeds the limit.

7.8.31 display bgp error discard

Function

The **display bgp error discard** command displays the information about the discarded error BGP packets.

Format

display bgp error discard [**peer** { *ipv4-address* | *ipv6-address* }]

Parameters

Parameter	Description	Value
peer	Displays errors on a specified peer.	-
ipv4-address	Displays errors on a peer with the specified IPv4 address.	It is in dotted decimal notation.
ipv6-address	Displays errors on a peer with the specified IPv6 address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

When a BGP fault occurs, the **display bgp error discard** command can be run to check the information about discarded error packets. The errors include cluster ID conflicts and the threshold overflow of AS-Path attributes.

By specifying the **peer** parameter, you can check the information about the error packets discarded by a specified BGP peer.

Precautions

The **display bgp error discard** command can be used to check only the error routing. To check the error routing among BGP peers, run the **display bgp error** command.

Example

Display information about discarded BGP error packets.

```
<HUAWEI> display bgp error discard
BGP Discard Info Counts:
Routes received with cluster ID loop          : 0
Routes received with as path count over limit  : 0
Routes advertised with as path count over limit: 0

 No discard record.
```

Table 7-124 Description of the display bgp error discard command output

Item	Description
BGP Discard Info Counts	Number of discarded BGP routes

Item	Description
Routes received with cluster ID loop	Number of discarded BGP routes with a duplicate cluster ID
Routes received with as path count over limit	Number of received BGP routes discarded due to the number of AS-Paths exceeding the upper threshold
Routes advertised with as path count over limit	Number of sent BGP routes discarded due to the number of AS-Paths exceeding the upper threshold
No discard record	No record about packet discarding

7.8.32 display bgp group

Function

The **display bgp group** command displays information about BGP peer groups.

Format

display bgp group [*group-name*]

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **group** [*group-name*]

display bgp vpls group [*group-name*]

display bgp ipv6 group [*group-name*]

display bgp vpnv6 all group [*group-name*]

display bgp vpnv6 vpn-instance *vpn-instance-name* **group** [*group-name*]

display bgp l2vpn-ad group [*group-name*]

display bgp { **mdt** | **mvpn** } **all group** [*group-name*]

NOTE

The **mdt,mvpn** parameter is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
vpnv4	Displays information about BGP VPNv4 peer groups.	-
all	Displays information about all BGP VPNv4 or VPNv6 peer groups.	-
vpn-instance *vpn-instance-name*	Displays information about BGP peer groups in a specified VPN instance.	The value must be an existing VPN instance name.
vpls	Displays information about BGP peer groups of VPLS.	- **NOTE** The **vpls** parameter is supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H.
ipv6	Displays information about IPv6 peer groups.	-
vpnv6	Displays information about BGP VPNv6 peer groups.	- **NOTE** The **vpnv6** parameter is supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H.
l2vpn-ad	Displays information about BGP peer groups of L2VPN-AD.	- **NOTE** The **l2vpn-ad** parameter is supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

Configuring BGP peer groups simplifies BGP network configuration and improves the route advertisement efficiency.

The **display bgp group** command displays information about the peer group, including peers in the peer group and configuration information about the peer group. This command is used in the following scenarios:

- Verify the configuration after running the **group** command to configure the peer group.
- Verify the configuration after running the **peer group** command to add a peer to the peer group.
- Verify the configuration after running the **undo peer group** command to delete a peer from the peer group.
- Verify the configuration after modifying the configuration of the peer group.

Precautions

BGP has multiple address families, each of which is configured independently. Information about peer groups in address families can be displayed by specifying different parameters.

If no parameter is specified, the **display bgp group** command displays default information about peer groups in the IPv4 unicast address family.

Example

Display information about a peer group named **my-peer**.

```
<HUAWEI> display bgp group my-peer

 BGP peer-group: my-peer
 Remote AS: 200
 listen-net: 10.1.1.0 24
 Authentication type configured: None
 Group's BFD has been enabled
 Type : internal
 Configured hold timer value: 180
 Keepalive timer value: 60
 Minimum route advertisement interval is 15 seconds
 listen-only has been configured
 PeerSession Members:
   10.2.2.2

 Maximum allowed route limit: 8000  Threshold: 75%, Parameter:  always connect-retry(default)
 Status codes: * - Dynamic
 Peer Preferred Value: 0
 No routing policy is configured
 Peer Members:
  Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv
  10.1.1.2        4   200        0        0     0 00:00:47      Active       0
```


Table 7-125 Description of the **display bgp group** command output

Item	Description
BGP peer-group	Name of a BGP peer group.
Remote AS	Number of the AS where a peer group resides.
listen-net	Network segment from which BGP listens to BGP connection requests.
Authentication type configured	Configured BGP authentication type. The value can be: • MD5 • None: indicates that no BGP authentication is configured.
Group's BFD has been enabled	BFD has been enabled for a peer group.
Type	Type of a peer group: • internal: indicates that the peer group is an IBGP peer group. • external: indicates that the peer group is an EBGP peer group.
Configured hold timer value	Value of the Hold timer.
Keepalive timer value	Value of the Keepalive timer.
Minimum route advertisement interval	Minimum interval between route advertisements.
listen-only has been configured	The peer or peer group only detects connection requests, and does not initiate any connection.
PeerSession Members	Peers that set up sessions.
Maximum allowed route limit	Maximum number of allowed BGP routes.
Threshold	Threshold for the quantity of received BGP routes (in percentage) out of the maximum number of routes that can be received.

Item	Description
Parameter	If **peer route-limit** command is configured, this parameter will be displayed: ● always connect-retry(default): By default, BGP always attempts to re-establish a connection after the BGP session is terminated due to the number of routes exceeding the threshold. ● alert-only: When the number of routes reaches the threshold, the switch generates only log information without terminating the BGP session. ● idle-forever: After the BGP session is terminated due to the number of routes exceeding the threshold, BGP does not re-establish a connection if the **reset bgp** command is not executed. ● idle-timeout: After the BGP session is terminated due to the number of routes exceeding the threshold, BGP automatically attempts to re-establish a connection if the timer expires.
Status codes: * - Dynamic	Status code. If the value starts with an asterisk (*), the peer is a dynamic peer. Currently, the value can only be *** - Dynamic**.
Peer Preferred Value	Preferred value of a peer.
Peer Members	Information about peers.
Peer	IP address of a peer.
V	BGP version.
AS	Number of the AS where a member of a peer group resides.
MsgRcvd	Number of received messages.
MsgSent	Number of sent messages.
OutQ	Number of messages to be sent to peers.
Up/Down	Period of time during which a BGP session keeps the current state.

Item	Description
State	BGP state mechanism: • Idle: indicates that BGP denies any request of entering. This is the initiatory status of BGP. Upon receiving a Start event, BGP initiates a TCP connection to the remote BGP peer, starts the ConnectRetry Timer with the initial value, detects a TCP connection initiated by the remote BGP peer, and changes its state to Connect. • Idle(Admin): indicates that the peer relationship is shut down initiatively and no attempt is made to establish the neighbor relationship. If the **peer ignore** command is configured or the peer is set to the Down state through the MIB, the neighbor is in the Idle (Admin) state. • Idle(Ovlmt): indicates that the peer relationship is interrupted because the number of routes exceeds the upper threshold. After a BGP peer relationship is interrupted due to the running of the **peer route-limit** command, the status of the BGP peer relationship is displayed as Idle(Ovlmt). If the **reset bgp** command is not run, the BGP peer relationship will not be reestablished. • Connect: indicates that BGP waits for the TCP connection to be set up before it determines whether to perform other operations. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer

Item	Description
	with the initial value, initiates a TCP connection to the remote BGP peer, and stays in the Connect state. ● Active: indicates that BGP tries to set up a TCP connection. This is the intermediate status of BGP. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value and changes its state to Connect. – If BGP initiates a TCP connection with an unknown IP address, the TCP connection fails. When this occurs, BGP restarts the ConnectRetry Timer with the initial value and stays in the Active state. ● OpenSent: indicates that BGP has sent one Open message to its peer and waits for an Open message from the peer. – If there are no errors in the Open message received, BGP changes its state to OpenConfirm. – If there are errors in the Open message received, BGP sends a Notification message to the remote peer and changes its state to Idle. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. ● OpenConfirm: indicates that BGP waits for a Notification message or a Keepalive message. – If BGP receives a Notification message, or the TCP connection fails, BGP changes its state to Idle. – If BGP receives a Keepalive message, BGP changes its state to Established.

Item	Description
	• Established: indicates that BGP peers can exchange Update, Notification and Keepalive packets. – If BGP receives an Update or a Keepalive message, its state stays in Established. – If BGP receives a Notification message, BGP changes its state to Idle.
PrefRcv	Indicates the number of route prefixes received by the local peer from the remote peer.

Display information about all BGP VPNv4 peer groups.

```
<HUAWEI> display bgp vpnv4 all group
Group in VPNV4:

 BGP peer-group: aa
 Remote AS number isn't specified
 Type : external
 PeerSession Members:
   10.3.3.3

 Peer Members:
   10.3.3.3
**********************

 BGP peer-group: bb
 Remote AS 100
 Type : internal
 PeerSession Members:
  NONE

 Peer Members:
   10.4.4.4

Group in VPN-Instance:

 BGP peer-group: cc
 Remote AS number isn't specified
 VPN-Instance(IPv4-family): vpn1

 Type : external
 PeerSession Members:
   10.2.2.1

 Peer Members:
   10.2.2.1
```

Table 7-126 Description of the **display bgp vpnv4 all group** command output

Item	Description
Group in VPNV4	Information about all BGP peer groups in the VPNv4 address family view.

Item	Description
Remote AS number isn't specified	This item is displayed when the peer group is a mixed EBGP peer group.
Group in VPN-Instance	Information about peer groups in a VPN instance.
VPN-Instance	Name of a VPN instance.

Display information about a BGP VPNv4 peer group named **rr1**.

```
<HUAWEI> display bgp vpnv4 all group rr1
Group in VPNV4:
 No such a peer-group

Group in VPN-Instance:

 BGP peer-group: rr1
 Remote AS number isn't specified
 VPN-Instance: 1

 Type : external
 Configured hold timer value: 180
 Keepalive timer value: 60
 Minimum route advertisement interval is 30 seconds
 PeerSession Members:
  NONE

 Peer Preferred Value: 0
 No routing policy is configured
 Peer Members:
  No Peer Exists
```

Display information about all BGP IPv6 peer groups.

```
<HUAWEI> display bgp vpnv6 all group
Group in VPNV6:

 BGP peer-group: 123
 Remote AS number isn't specified
 Type : external
 PeerSession Members:
   FC00:0:0:1::1

 Peer Members:
  No Peer Exists
**********************

 BGP peer-group: 222
 Remote AS 200
 Type : internal
 PeerSession Members:
   FC00:0:0:2::2

 Peer Members:
  No Peer Exists
**********************

 BGP peer-group: 333
 Remote AS 400
 Type : external
 PeerSession Members:
   FC00:0:0:3::3
```



```
 Peer Members:
  No Peer Exists


Group in VPN-Instance:

 BGP peer-group: 55
 Remote AS number isn't specified
 VPN-Instance(IPv6-family): vpn1

 Type : external
 PeerSession Members:
   FC00:0:0:4::4

 Peer Members:
   FC00:0:0:4::4
```

Table 7-127 Description of the **display bgp vpnv6 all group** command output

Item	Description
Group in VPNV6	Information about all BGP peer groups in a VPNv6 address family.
Group in VPN-Instance	Information about peer groups in an IPv6 VPN instance.
VPN-Instance	Name of an IPv6 VPN instance.

Display information about all BGP IPv6 peer groups.

```
<HUAWEI> display bgp ipv6 group
BGP peer-group is in
 Remote AS 100
 Type : internal
 PeerSession Members:
   FC00:0:0:2::1

 Peer Members:
   FC00:0:0:1::1              FC00:0:0:2::1
 ************************

 BGP peer-group is ex
 Remote AS number not specified
 Type : external
 PeerSession Members:
   FC00:0:0:20::1
 Peer Members:
   FC00:0:0:10::1              FC00:0:0:20::1
```

Display information about a BGP VPNv6 peer group named **rr1**.

```
<HUAWEI> display bgp vpnv6 all group rr1
Group in VPNV6:
 No such a peer-group

Group in VPN-Instance:
 BGP peer-group: rr1
 Remote AS number isn't specified
 VPN-Instance: 1
 Type : external
 Configured hold timer value: 180
 Keepalive timer value: 60
 Minimum route advertisement interval is 30 seconds
```



```
PeerSession Members:
 NONE
Peer Preferred Value: 0
No routing policy is configured
Peer Members:
 No Peer Exists
```

Display information about all peer groups in an IPv6 VPN instance named **vpn6** on the local device.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn6 group
BGP peer-group: g1
 Remote AS 65410
 Type : external
 PeerSession Members:
  FC00:0:0:2000::2

 Peer Members:
  FC00:0:0:2000::2
```

Display information about a peer group named **g1** in an IPv6 VPN instance named **vpn6** on the local device.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn6 group g1
BGP peer-group: g1
 Remote AS 65410
 Type : external
 Configured hold timer value: 180
 Keepalive timer value: 60
 Minimum time between advertisement runs is 30 seconds
 PeerSession Members:
  FC00:0:0:2000::2

 Peer Preferred Value: 0
 No routing policy is configured
 Peer Members:
  Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  FC00:0:0:2000::2 4 65410      103       90     0 01:20:55 Established       0
```

7.8.33 display bgp ipv6 bfd session

Function

The **display bgp ipv6 bfd session** command displays information about the BFD session set up by BGP.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp ipv6 bfd session { [**vpnv6 vpn-instance** *vpn-instance-name*] **peer** *ipv6-address* | **all** }

Parameters

Parameter	Description	Value
vpnv6	Displays the BFD session of VPNv6.	-
vpn-instance *vpn-instance-name*	Specifies the name of the VPN instance.	The value must be an existing VPN instance name.
peer *ipv6-address*	Specifies the IPv6 address of the peer.	-
all	Displays all BFD sessions set up by BGP.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The BFD session set up by BGP4+ helps BGP4+ quickly discover faults between BGP4+ peers and inform BGP4+ to recalculate routes for packet forwarding.

Run the **display bgp ipv6 bfd session** command to display information about the BFD session set up by BGP4+ in the following situations:

- Check whether the BFD session is successfully set up or view detailed information about the BFD session that is successfully configured.
- Check whether the BFD session set up by BGP4+ is successfully deleted after running the **undo peer bfd enable** command.
- Verify the configuration after running the **undo peer bfd block** command to restore a peer inheriting the BFD function of the peer group.
- Verify the configuration after running the **peer bfd** command to set BFD parameters.

The information about the BFD session of a specified peer can be displayed by specifying different parameters.

- Run the **display bgp ipv6 bfd session vpnv6 vpn-instance** *vpn-instance-name* **peer** *ipv6-address* command to display information about the BFD session of a specified peer in a VPN instance.
- Run the **display bgp ipv6 bfd session peer** *ipv6-address* command to display information about the BFD session of a specified peer on the public network.

- Run the **display bgp ipv6 bfd session all** command to display information about the BFD sessions of all BGP4+ peers.

Prerequisites

The BFD session has been set up using the **peer bfd enable** command. If the BFD session has not been set up by BGP4+, no information is displayed after running the **display bgp ipv6 bfd session** command.

Example

Display all BFD sessions set up by BGP.

```
<HUAWEI> display bgp ipv6 bfd session all
  Local_Address  : FC00:3::1
  Peer_Address   : FC00:3::2
  Tx-interval(ms): 100          Rx-interval(ms): 100
  Multiplier     : 4            Interface      : GigabitEthernet0/0/1
  LD/RD          : 8193/8207    Session-State  : Up
  Wtr-interval(m): 0
```

Table 7-128 Description of the **display bgp ipv6 bfd session** command output

Item	Description
Local_Address	Local address
Peer_Address	Peer address
Tx-interval (ms)	Interval for sending BFD packets, in milliseconds
Rx-interval (ms)	Interval for receiving BFD packets, in milliseconds
Multiplier	Local detection multiple
Interface	Interface on which the BFD session is set up **NOTE** Information about the interface on which the BFD session is set up only when the directly connected interface is used to set up the EBGP neighbor relationship. In other cases, information about the interface is displayed as Unknown.
LD/RD	Local/remote discriminator

Item	Description
Session-State	BFD status • Admin down: The BFD session is closed on the local end. • BFD global disable: BFD is disabled globally. • BFD session number exceed: The number of BFD sessions exceeds the maximum limit. • Detect down: BFD detects a link status fault and interrupts the connection. • Init: The BFD session is in the initialized state. • Neighbor down: The peer end detects that the BFD session goes Down and informs the local end of the change, and the local end then sets the neighbor status to Down. • Receive admin down: The BFD session is closed on the peer end (for example, the BFD session is disabled on the peer end). • Up: The BFD session is set up.
Wtr-interval(m)	Interval for flap dampening, in minutes

7.8.34 display bgp ipv6 routing-table

Function

The **display bgp ipv6 routing-table** command displays BGP IPv6 routes.

Format

display bgp ipv6 routing-table [**verbose**]

display bgp ipv6 routing-table *ipv6-address* [*prefix-length*]

display bgp ipv6 routing-table as-path-filter { *as-path-filter-number* | *as-path-filter-name* }

display bgp ipv6 routing-table community [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp ipv6 routing-table community-filter { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp ipv6 routing-table different-origin-as

display bgp ipv6 routing-table regular-expression *as-regular-expression*

display bgp ipv6 routing-table peer *ipv6-address* { **accepted-routes** | **advertised-routes** [*dest-ipv6-address* [*prefix-length*]] | **received-routes** [**active**] }

display bgp ipv6 routing-table peer *ipv6-address* **received-routes** *dest-ipv6-address* [*prefix-length* [**original-attributes**]]

display bgp ipv6 routing-table peer *ipv4-address* **received-routes** *dest-ipv6-address* [*prefix-length* [**original-attributes**]]

display bgp ipv6 routing-table time-range *start-time end-time*

Parameters

Parameter	Description	Value
verbose	Displays detailed information about BGP4+ public network routes.	-
ipv6-address	Specifies the IPv6 address of the peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the length of the prefix.	The value is an integer that ranges from 0 to 128.
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matching AS-Path filter.	It is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS-Path filter.	The name is a string of 1 to 51 characters without any space. It is case-sensitive.
community	Displays the routing information of the specified BGP community attribute in the routing table.	-
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.

Parameter	Description	Value
aa:nn	Specifies the community attribute number.	Both *aa* and *nn* are integers ranging from 0 to 65535.
internet	Displays the BGP routes with Internet community attribute.	-
no-advertise	Displays the BGP routes with No-Advertise community attribute.	-
no-export	Displays the BGP routes with the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes with the No-Export-Subconfed community attribute.	-
whole-match	Indicates the exact matching.	-
community-filter	Displays the routes that match the specified BGP community filter.	-
community-filter-name	Specifies the name of the community filter.	The name is a string of 1 to 51 characters. The string cannot be all numbers.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
different-origin-as	Displays routes that have the same destination address but different source ASs.	-

Parameter	Description	Value
peer	Displays the routing information for the specified BGP peer.	-
ipv4-address	Specifies the IPv4 address of the peer. **NOTE** The parameter is not supported on S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, and S5736-S.	The value is in dotted decimal notation.
advertised-routes	Displays the routes advertised to the specified peer.	-
dest-ipv6-address	Specifies the destination IPv6 address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
received-routes	Displays the routes received from the specified peer.	-
active	Displays the active routes received from the specified peer.	-
original-attributes	Displays the original attributes of a public route from a specified BGP peer before the route is filtered by the local import policy. To display such attributes, the **peer keep-all-routes** command must have been run.	-
regular-expression *as-regular-expression*	Specifies the matched AS regular expression.	The value is a string of 1 to 80 characters.

Parameter	Description	Value
accepted-routes	Displays the routes that are received from the peer and filtered through a routing policy.	-
time-range *start-time end-time*	Displays IPv6 BGP routes that flap within the specified time period. For example, the value 0d0h5m0s of *start-time* indicates 5 minutes before the current time. The value 0d0h10m0s of *end-time* indicates 10 minutes before the current time. All IPv6 BGP routes with the Keepalive time in the range of 5 to 10 minutes are displayed.	The value ranges of *start-time* and *end-time* both are 0d0h0m0s-10000d23h59m59s.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify different parameters to view the specific routing information.

You can run the **display bgp ipv6 routing-table time-range** *start-time end-time* command to view BGP4+ routes that flap within the specified time period. For example, if service traffic is abnormal or CPU usage of the device remains high within a certain time period, you can run this command to check whether route flapping occurs within the specified time period. The faulty route can be viewed in the command output, facilitating fault location.

Example

Display BGP IPv6 routes.

```
<HUAWEI> display bgp ipv6 routing-table

BGP Local router ID is 10.1.1.1
Status codes: * - valid, > - best, d - damped, x - best external,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete
```



```
Total Number of Routes: 1
*>  Network  : FC00:0:0:2000::                PrefixLen : 64
    NextHop  : FC00:0:0:2000::1               LocPrf    :
    MED      : 0                           PrefVal   : 0
    Label    :
    Path/Ogn : i
```

Displays BGP IPv6 routes that flap within the specified time period.

```
<HUAWEI> display bgp ipv6 routing-table time-range 0d5h0m0s 1d5h0m0s
BGP Local router ID is 10.1.1.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 *>i Network  : FC00:0:0:100::100             PrefixLen : 128
     NextHop  : FC00:0:0:12::1                Duration  : 05h46m16s
     Peer     : FC00:0:0:12::1
     Path/Ogn : ?
 * i
     NextHop  : FC00:0:0:112::1               Duration  : 05h46m16s
     Peer     : FC00:0:0:112::1
     Path/Ogn : ?
 *>i Network  : FC00:0:0:111::111             PrefixLen : 128
     NextHop  : FC00:0:0:12::1                Duration  : 05h46m16s
     Peer     : FC00:0:0:12::1
     Path/Ogn : ?
 * i
     NextHop  : FC00:0:0:112::1               Duration  : 05h46m16s
     Peer     : FC00:0:0:112::1
     Path/Ogn : ?
```

Table 7-129 Description of the display bgp ipv6 routing-table command output

Item	Description
BGP Local router ID	Indicates the router ID of the local BGP device.
Network	Indicates the network address in the BGP routing table.
PrefixLen	Indicates the prefix length.
NextHop	Indicates the next-hop address of the packet.
LocPrf	Indicates the local preference.
MED	Indicates the MED of the route.
PrefVal	Indicates the preferred value.
Label	Indicates the label value.
Duration	Route duration
Peer	Peer IP address
Path/Ogn	Indicates the AS_Path number and the Origin attribute.

Display detailed information about the specified invalid BGP4+ routes.

```
<HUAWEI> display bgp ipv6 routing-table FC00::1
 BGP local router ID : 10.1.1.1
 Local AS number : 100
 Paths:   1 available, 0 best, 0 select
 BGP routing table entry information of FC00::1/128
 From: 10::2 (10.1.1.2)
 Route Duration: 00h17m46s
 Relay IP Nexthop: ::
 Relay IP Out-Interface:
 Original nexthop: FC00::2
 AS-path 200, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for IP
unreachable
 Not advertised to any peer yet
```

Table 7-130 Description of the display bgp ipv6 routing-table command output

Item	Description
BGP local router ID	ID of the local BGP device. The format is the same as the IPv4 address.
Local AS number	Local AS number.
Paths	Information about paths of BGP routes
BGP routing table entry information of FC00::1/128	The following information is about FC00::1/128 routing entries.
From	IP address of the router that sends the route. 10.1.1.2 is the source interface IP address of the peer with which the BGP connection is established, and 10::2 is the router ID of the peer.
Route Duration	Duration of routes.
Relay IP Nexthop	Recursive next hop.
Relay IP Out-Interface	Recursive outbound interface.
Original nexthop	Original next hop.
AS-path 200	AS_Path attribute.

Item	Description
origin incomplete	Well-known mandatory property. This property defines the origin of a path and records how a route turns to a BGP route. The property has the following three values: ● IGP: The priority of this value is the highest. The origin property of the routes that are added to the BGP routing table by using the **network (BGP)** command is IGP. ● EGP: The priority of this value is second to that of IGP. The origin property of the routes imported from EGP is EGP. ● Incomplete: The priority of this value is the lowest. The value indicates the origin of a route is unknown. The origin property of the routes that are added to the BGP routing table by using the **import-route (BGP)** command is Incomplete.
MED	Multi-Exit discriminator of route.
localpref	Local priority.
pref-val	Value preferred by the protocol.
internal	The BGP route is an internal route.
pre 255	The priority of the BGP route is 255.
invalid for IP unreachable	Reason why a route is invalid: ● invalid for route-policy not pass: The route does not match the route-policy. ● invalid for supernet route: The route is a supernet route. ● invalid for IP unreachable: The route fails to recurse to another route. ● invalid for supernet route not advertise: No supernet routes are advertised. ● invalid for supernet label route not advertise: No supernet labeled routes are advertised. ● invalid for next-hop unreachable: The next-hop IP address is unreachable. ● invalid for tunnel unreachable: The route fails to recurse to a tunnel.

Item	Description
Not advertised to any peer yet	The BGP route has not been advertised to any peer yet.

7.8.35 display bgp ipv6 routing-table statistics

Function

The **display bgp ipv6 routing-table statistics** command displays statistics about BGP IPv6 routes.

Format

display bgp ipv6 routing-table statistics

display bgp ipv6 routing-table statistics as-path-filter { *as-path-filter-number* | *as-path-filter-name* }

display bgp ipv6 routing-table statistics community [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp ipv6 routing-table statistics community-filter { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp ipv6 routing-table peer *ipv6-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

display bgp ipv6 routing-table statistics regular-expression *as-regular-expression*

display bgp ipv6 routing-table statistics different-origin-as

Parameters

Parameter	Description	Value
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matching AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The name is a string of 1 to 51 characters without any space. It is case-sensitive.
community	Displays the routing information of the specified BGP community attribute in the routing table.	-

Parameter	Description	Value
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
aa:nn	Specifies the community attribute number.	Both *aa* and *nn* are integers ranging from 0 to 65535.
internet	Displays the BGP routes with Internet community attribute.	-
no-advertise	Displays the BGP routes with the No-Advertise community attribute.	-
no-export	Displays the BGP routes with the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes with the No-Export-Subconfed community attribute.	-
whole-match	Indicates exact matching.	-
community-filter	Displays the routing information that matches the specified BGP community filter.	-
community-filter-name	Specifies the name of the community filter.	The name is a string of 1 to 51 characters. The string cannot be all numbers.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
peer	Displays the routing information for the specified BGP peer.	-
ipv6-address	Specifies the IPv6 address of the peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
advertised-routes	Displays the routes advertised to the specified peer.	-
received-routes	Displays the routes received from the specified peer.	-
active	Displays the active routes received from the specified peer.	-

Parameter	Description	Value
regular-expression *as-regular-expression*	Specifies the matched AS regular expression.	The value is a string of 1 to 80 characters.
different-origin-as	Displays routes that have the same destination address but different source ASs.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp ipv6 routing-table statistics** command to check statistics about BGP IPv6 routes.

Example

Display statistics of BGP IPv6 routes.

```
<HUAWEI> display bgp ipv6 routing-table statistics

 Total Number of Routes: 4
```

Table 7-131 Description of the **display bgp ipv6 routing-table statistics** command output

Item	Description
Total Number of Routes	Total number of routes in the routing table.

7.8.36 display bgp mdt brief

Function

The **display bgp mdt brief** command displays brief information about VPN instances in BGP MDT address family.

Format

display bgp mdt { **all** | **vpn-instance** *vpn-instance-name* } **brief**

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Displays information about MDT and all VPN instances.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display bgp mdt brief** command is used to display brief information about VPN instances in BGP MDT address family.

Example

Display brief information about MDT and all VPN instances.

```
<HUAWEI> display bgp mdt all brief

MDT:
  Rd Num           Peer Num         Route Num
  0                1                0

VPN-Instance(IPv4-MDT-family):
  VPN-Instance Name   Peer Num         Route Num
  vrf0                0                0
  vrf1                0                0
  vrf11               0                0
  vrf12               0                0
  vrf13               0                0
  vrf14               0                0
  vrf2                0                20
  vrf3                0                20
  vrf4                0                24
  vrf5                0                24
  vrf6                0                0
  vrf7                0                0
  vrf8                0                20
```


Table 7-132 Description of the **display bgp mdt all brief** command output

Item	Description
Rd Num	Number of Route Distinguishers (RDs).
Peer Num	Number of peers.
Route Num	Number of routes.
VPN-Instance Name	Name of a VPN instance.

7.8.37 display bgp mdt routing-table

Function

The **display bgp mdt routing-table** command displays the information of BGP MDT routes.

Format

display bgp mdt { **all** | **vpn-instance** *vpn-instance-name* } **routing-table** [*network* [{ *mask* | *mask-length* } [**longer-prefixes**]]]

display bgp mdt route-distinguisher *route-distinguisher* **routing-table** [*network* [*mask* | *mask-length*]]

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table regular-expression** *as-regular-expression*

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table community** [*aa:nn* | *community-number*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp mdt all routing-table peer *ipv4-address* [**advertised-routes** [**statistics** | *ipv4-address*] | **received-routes** [**active**] [**statistics**]]

display bgp mdt all routing-table peer *ipv4-address* **received-routes** *network*

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Display all the BGP MDT routing information.	-
route-distinguisher *route-distinguisher*	Displays the BGP MDT routing information of the specified Route Distinguisher (RD).	-
vpn-instance *vpn-instance-name*	Displays the BGP MDT routing information of the specified VPN instance.	The value must be an existing VPN instance name.
network	Specifies the IPv4 network address.	It is in dotted decimal notation.
mask | *mask-length*	Specifies mask in dotted decimal notation or mask-length.	The value of mask-length is an integer that ranges from 0 to 32.
longer-prefixes	Matches according to the mask longer than the specified length.	-
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matching AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals.
cidr	Displays the BGP MDT routing information about the Classless Inter-Domain Routing (CIDR).	-
different-origin-as	Displays routes that have the same destination address but different source ASs.	-

Parameter	Description	Value
regular-expression *as-regular-expression*	Specifies the matched AS regular expression.	The value is a string of 1 to 80 characters.
community-filter	Displays the BGP MDT routing information that matches the specified BGP community filter.	-
community-filter-name	Specifies the name of the community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
community	Displays the BGP MDT routing information of the specified BGP community attribute in the routing table.	-
aa:nn	Specifies the community attribute number.	Both aa and nn are integers ranging from 0 to 65535. You can set a maximum of 29 community numbers.
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
internet	Displays the BGP routes with Internet community attribute.	-
no-advertise	Displays the BGP routes with No-Advertise community attribute.	-
no-export	Displays the BGP routes with the No-Export community attribute.	-

Parameter	Description	Value
no-export-subconfed	Displays the BGP routes with the No-Export-Subconfed community attribute.	-
whole-match	Indicates the exact matching.	-
peer *ipv4-address*	Displays the BGP MDT routing information for the specified BGP peer.	-
advertised-routes	Displays the BGP MDT routing information advertised to the specified peer.	-
statistics	Display the statistics of the BGP MDT routes.	-
ipv4-address	Specifies the IPv4 network address.	It is in dotted decimal notation.
received-routes	Displays the BGP MDT routing information received from the specified peer.	-
active	Displays the BGP MDT routing information received from the specified peer.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify different parameters to view the specific routing information of BGP MDT routes. If the length of the destination address mask of an IPv4 route is the same as that of its natural mask, the mask length is not displayed.

Example

Display all the BGP MDT routing information.

```
<HUAWEI> display bgp mdt all routing-table
```



```
BGP Local router ID is 192.168.7.1
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

Total number of routes from all PE: 3
Route Distinguisher: 100:1

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>   1.1.1.1/32         0.0.0.0        0                    0      ?

Route Distinguisher: 200:1

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>i  1.1.1.1/32         192.168.100.10  0         100        0      33 55?
*>i  2.2.2.2/32         192.168.100.10  0         100        0      33 55?

Total number of routes of IPv4-MDT-family for vpn-instance vrf1: 3
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>   1.1.1.1/32         0.0.0.0        0                    0      ?
* i                   192.168.100.10  0         100        0      33 55?
*>i  2.2.2.2/32         192.168.100.10  0         100        0      33 55?
```

Display all the BGP MDT routing information of the VPN instance named vpna.

```
<HUAWEI> display bgp mdt vpn-instance vpna routing-table

Total Number of Routes: 2

BGP Local router ID is 2.2.2.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>i  10.1.1.0/24        1.1.1.9        0          100        0      ?
*>i  10.2.1.0/24        3.3.3.9        0          100        0      ?
```

Display the BGP MDT routing information of the specified RD.

```
<HUAWEI> display bgp mdt route-distinguisher 100:1 routing-table

BGP Local router ID is 192.168.7.1
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

Route Distinguisher: 100:1

Total Number of Routes: 1

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>   1.1.1.1/32         0.0.0.0        0                    0      ?

Total number of routes of IPv4-MDT-family for vpn-instance vrf1: 3
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>   1.1.1.1/32         0.0.0.0        0                    0      ?
* i                   192.168.100.10  0         100        0      33 55?
*>i  2.2.2.2/32         192.168.100.10  0         100        0      33 55?
```

Display all BGP MDT routes of community 1000:100.

```
<HUAWEI> display bgp mdt all routing-table community 1000:100

BGP Local router ID is 2.2.2.2
```



```
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

Total number of routes from all PE: 1

Route Distinguisher: 100:1

     Network          NextHop       MED       LocPrf    PrefVal Community

*>i  10.1.1.0/24      1.1.1.1       0         100       0     <1000:100>

Total number of routes of IPv4-MDT-family for vpn-instance vpna: 1
     Network          NextHop       MED       LocPrf    PrefVal Community

*>i  10.1.1.0/24      1.1.1.1       0         100       0     <1000:100>
```

Table 7-133 Description of the **display bgp mdt routing-table** command output

Item	Description
Network	Network address in the BGP routing table.
Next Hop	Next Hop address through which the packet has to be sent.
MED	Multi-Exit discriminator.
LocPrf	Local priority.
PrefVal	Value preferred by the protocol.
Path/Ogn	AS_Path number and the attributes of Origin.
Community	Community attributes.

Display all BGP MDT routes of community 1000:100 with the Internet community attribute.

```
<HUAWEI> display bgp mdt all routing-table community 1000:100 internet

BGP Local router ID is 2.2.2.2
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

Total number of routes from all PE: 1

Route Distinguisher: 100:1


     Network          NextHop       MED       LocPrf    PrefVal Community

*>i  10.1.1.0/24      1.1.1.1       0         100       0     <1000:100>

Total number of routes of IPv4-MDT-family for vpn-instance vpna: 4
     Network          NextHop       MED       LocPrf    PrefVal Community

*>i  10.1.1.0/24      1.1.1.1       0         100       0     <1000:100>
*>   10.2.1.0/24      0.0.0.0       0                   0
*                   10.2.1.1       0                   0
*>   10.2.1.2/32      0.0.0.0
```

Display routes sent by 2.2.2.2.


```
<HUAWEI> display bgp mdt all routing-table peer 2.2.2.2 received-routes

 BGP Local router ID is 1.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1

 Route Distinguisher: 2:2


      Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  2.2.2.2/32       2.2.2.2         0          100        0      ?
```

Display detailed information about the route with prefix 2.2.2.2 sent by 2.2.2.2.

```
<HUAWEI> display bgp mdt all routing-table peer 2.2.2.2 received-routes 2.2.2.2

 BGP local router ID : 1.1.1.1
 Local AS number : 100
 Route Distinguisher: 2:2

 BGP routing table entry information of 2.2.2.2/32:
 From: 2.2.2.2 (2.2.2.9)
 Route Duration: 00h03m42s
 Relay IP Nexthop: 10.1.1.2
 Relay IP Out-Interface: Ethernet0/0/1
 Original nexthop: 2.2.2.2
 Qos information : 0x0
 Ext-Community:RT <1 : 1>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255, IGP
cost 1
 MDT group address: 232.1.1.1
 Not advertised to any peer yet
```

Table 7-134 Description of the **display bgp mdt all routing-table peer received-routes** command output

Item	Description
BGP local router ID	ID of the local BGP device. The format is the same as the IPv4 address.
Local AS number	Local AS number.
Route Distinguisher	Route distinguisher.
BGP routing table entry information of 2.2.2.2/32	The following information is about 2.2.2.2/32 routing entries.
From	IP address of the router that sends the route. 2.2.2.2 is the IP address of the source interface of the peer with which the BGP connection is established, and 2.2.2.9 is the Router ID of the peer.
Route Duration	Duration of routes.
Relay IP Nexthop	Recursive next hop.
Relay IP Out-Interface	Recursive outbound interface.

Item	Description
Original nexthop	Original next hop.
Qos information	QoS information.
Ext-Community	Extended community attribute.
AS-path Nil	AS_Path attribute, with Nil indicating that the attribute value is null.
origin incomplete	Well-known mandatory property. This property defines the origin of a path and records how a route turns to a BGP route. The property has the following three values: • IGP: The priority of this value is the highest. The origin property of the routes that are added to the BGP routing table by using the **network (BGP)** command is IGP. • EGP: The priority of this value is second to that of IGP. The origin property of the routes imported from EGP is EGP. • Incomplete: The priority of this value is the lowest. The value indicates the origin of a route is unknown. The origin property of the routes that are added to the BGP routing table by using the **import-route (BGP)** command is Incomplete.
localpref	Local priority.
pref-val	Value preferred by the protocol.
valid	The BGP route is a valid route.
internal	The BGP route is an internal route.
best	The BGP route is an optimal route.
select	The BGP route is a preferred route.
pre 255	The priority of the BGP route is 255.
IGP cost	Cost of the relied next hop route of BGP route.
MDT group address	Group address of MDT.
Not advertised to any peer yet	The BGP route has not been advertised to any peer yet.

7.8.38 display bgp mdt routing-table statistics

Function

The **display bgp mdt routing-table statistics** command displays the statistics of the BGP MDT routes.

Format

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics**

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics regular-expression** *as-regular-expression*

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp mdt { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics community** [*aa:nn* | *community-number*] & <1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Displays all the statistics of the BGP MDT routes.	-
route-distinguisher *route-distinguisher*	Displays the BGP MDT routing statistics of the specified Route Distinguisher (RD).	-
vpn-instance *vpn-instance-name*	Displays the BGP MDT routing statistics of the specified VPN instance.	The value must be an existing VPN instance name.

Parameter	Description	Value
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matched AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals.
cidr	Displays the BGP MDT routing statistics about the Classless Inter-Domain Routing (CIDR).	-
different-origin-as	Displays the routes that have the same destination address but different source AS number.	-
regular-expression *as-regular-expression*	Indicates the matched AS regular expression.	The value is a string of 1 to 80 characters.
community-filter	Displays the BGP MDT routing statistics that matches the specified BGP community filter.	-
community-filter-name	Specifies the name of the community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
whole-match	Indicates exact matching.	-
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.

Parameter	Description	Value
community	Displays the BGP MDT routing statistics of the specified BGP community attribute in the routing table.	-
aa:nn	Specifies the community number.	Both aa and nn are integers ranging from 0 to 65535. You can set a maximum of 29 community numbers.
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
internet	Displays the BGP routes with Internet community attribute.	-
no-advertise	Displays the BGP routes with the No-Advertise community attribute.	-
no-export	Displays the BGP routes with the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes with the No-Export-Subconfed community attribute.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

None

Example

Displays all the statistics of the BGP MDT routes.

```
<HUAWEI> display bgp mdt all routing-table statistics

 Total number of routes from all PE: 20
```



```
Total number of routes of IPv4-MDT-family for vpn-instance vpn1: 12
```

Table 7-135 Description of the **display bgp mdt routing-table statistics** command output

Item	Description
Total number of routes from all PE	The number of the routes received from PEs in the BGP MDT routing table.
Total number of routes of IPv4-MDT-family for vpn-instance	The number of the routes of the specified VPN instance in the BGP MDT routing table.

7.8.39 display bgp multicast group

Function

The **display bgp multicast group** command displays the information about an MBGP peer group.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast group [*group-name*]

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the peer group is specified, the detailed information on the specified peer group is displayed.

If the peer group is not specified, the information on all MBGP peer groups is displayed.

Example

Display the information about peer group **my-peer**.

```
<HUAWEI> display bgp multicast group my-peer
BGP peer-group: my-peer
 Remote AS: 1
 Authentication type configured: None
 Type : external
 Configured hold timer value: 180
 Keepalive timer value: 60
 Connect-retry timer value: 32
 Minimum route advertisement interval is 30 seconds
 PeerSession Members:
   172.16.14.1
 It's route-reflector-client
 Peer Preferred Value: 0
 No routing policy is configured
 Peer Members:
  Peer         V  AS  MsgRcvd  MsgSent  OutQ  Up/Down    State   PrefRcv
  172.16.14.1  4   1       43       29     0  00:03:03   Established   21
```

Table 7-136 Description of the display bgp multicast group command output

Item	Description
BGP peer-group: my-peer	Peer group: name of the group
Remote AS	The number of the AS where the peer resides
Authentication type configured	Indicates the configured BGP authentication type: • MD5 • Keychain (kk), in which kk indicates the name of the configured keychain authentication • None, which indicates no BGP authentication is configured
Type	Indicates the types of peers: • internal: The type of the peer group is IBGP • external: The type of the peer group is EBGP
Configured hold timer value	Value of the holdtime timer
Keepalive timer value	Value of the Keepalive timer

Item	Description
Connect-retry timer value	Value of the Connect-retry timer
Minimum route advertisement interval	Shortest interval for advertising routes
PeerSession Members	Indicates peers that set up session connections
It's route-reflector-client	The local device is a route reflector client
Peer Preferred Value	Preferred value of peers
Peer	IP address of the peer
V	MBGP version
AS	The number of the AS where the peer resides
MsgRcvd	Number of messages received
MsgSent	Number of messages sent
OutQ	Number of messages to be sent to the peer
Up/Down	Time during which the MBGP session is in the current state
State	Indicates the MBGP state mechanism: ● Idle: indicates that MBGP denies the connection request. This is the initiate status of MBGP. ● Active: indicates that MBGP tries to set up TCP connection. This is the intermediate status of MBGP. ● Established: In the status, MBGP peers can exchange Update, Notification and Keepalive packets. ● Connect: indicates that MBGP performs other actions after the TCP connection is set up. ● OpenSent: indicates that MBGP waits for an Open message from the peer. ● OpenConfirm: indicates that MBGP waits for a Notification message or a Keepalive message.
PrefRcv	The number of prefixes received by the local peer from the remote peer

7.8.40 display bgp multicast network

Function

The **display bgp multicast network** command displays the routes to be advertised by MBGP through the **network** command.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast network

Parameters

None.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

This command is used to display all the configurations of the **network (BGP)** command in the IPv4 multicast address family view. Routes can be imported and then advertised by MBGP only when the route prefix satisfies the following conditions:

- It is specified in the **network** command.
- It already exists in the IP routing table.
- It is active.

Example

Display the routing information of a network segment advertised by MBGP.

```
<HUAWEI> display bgp multicast network
 BGP Local Router ID is 10.2.2.9
 Local AS Number is 100(Multicast)
 Network          Mask            Route-policy

 10.1.1.1         255.255.255.0
 10.2.2.2         255.255.255.0
```


Table 7-137 Description of the display bgp multicast network command output

Item	Description
BGP Local Router ID	ID of the local MBGP router
Local AS Number	Number of the local AS
Network	Network address locally advertised
Mask	Mask of the network address
Route-policy	Used routing policy

7.8.41 display bgp multicast paths

Function

The **display bgp multicast paths** command displays AS_Path information of MBGP.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast paths [*as-regular-expression*]

Parameters

Parameter	Description	Value
as-regular-expression	Displays the AS_Path regular expression.	The value is a string of 1 to 80 characters.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast paths** command to check AS_Path information of MBGP.

Example

Display information about MBGP paths.


```
<HUAWEI> display bgp multicast paths 1

Total Number of Routes: 17
Total Number of Paths: 1


  Address         Refcount   MED        Path/Origin
  0x54169A4       17         0          1?
```

Table 7-138 Description of the display bgp multicast paths command output

Item	Description
Total Number of Routes	Total number of routes
Total Number of Paths	Total number of AS paths
Address	Address of the path attribute node in the local database in the hexadecimal format
Refcount	Number of times that the route is referenced
MED	Multi-Exit discriminator
Path	AS_Path list of the route
Origin	Origin of the route

7.8.42 display bgp multicast peer

Function

The **display bgp multicast peer** command displays information about a specified MBGP peer. If *peer-address* is not specified, information about all MBGP peers is displayed.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast peer [[*peer-address*] **verbose**]

Parameters

Parameter	Description	Value
peer-address	Specifies the address of an MBGP peer.	The value is in dotted decimal notation.
verbose	Specifies detailed information about an MBGP peer.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast peer** command to check information about a specified MBGP peer. If *peer-address* is not specified, information about all MBGP peers is displayed.

Example

Display information about all MBGP peers.

```
<HUAWEI> display bgp multicast peer

BGP local router ID : 10.13.13.9
Local AS number : 1
Total number of peers : 1                 Peers in established state : 1

 Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

 10.2.1.2        4     2       36       37     0 00:15:35 Established      24
```

Table 7-139 Description of the display bgp multicast peer command output

Item	Description
BGP local router ID	ID of the local MBGP router
Local AS number	Local AS number
Total number of peers	Total number of peers
Peers in established state	Number of peers in established state
Peer	IP address of peers
V	MBGP version of peers
AS	AS number
MsgRcvd	Number of messages received
MsgSent	Number of messages sent
OutQ	Messages to be sent to the specified peers
Up/Down	Period during which the MBGP session is in the current state
State	Status of the peers
PrefRcv	Number of prefixes received by the local peer from the remote peer

7.8.43 display bgp multicast routing-table

Function

The **display bgp multicast routing-table** command displays the MBGP routing information of a specified network in the MBGP routing table.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table [*ip-address* [*mask-length* [**longer-prefixes**] | *mask* [**longer-prefixes**]]]

Parameters

Parameter	Description	Value
ip-address	Specifies an IPv4 address.	The address is in dotted decimal notation.
mask-length	Specifies the mask length of the IPv4 address.	The value is an integer that ranges from 0 to 32.
mask	Specifies the mask of the IPv4 address.	The value is in dotted decimal notation.
longer-prefixes	Matches routes whose masks are shorter than the specified mask length.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If *ip-address* is not specified, this command will display all the information in the MBGP routing table.

Example

Display MBGP routing information.

```
<HUAWEI> display bgp multicast routing-table

BGP local router ID is 10.13.13.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete
Total Number of Routes: 14
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>   10.1.2.0/24        10.2.1.2                    0         2?
*>   10.1.13.0/24       0.0.0.0        0            0         ?
*>   10.1.13.2/32       0.0.0.0        0            0         ?
*>   10.2.1.0/24        0.0.0.0        0            0         ?
*>   10.2.1.0/24        10.2.1.2       0            0         2?
*>   10.2.1.1/32        0.0.0.0        0            0         ?
*>   10.2.13.0/24       0.0.0.0        0            0         ?
*>   10.2.13.1/32       0.0.0.0        0            0         ?
*>   10.2.13.2/32       0.0.0.0        0            0         ?
```

Table 7-140 Description of the display bgp multicast routing-table command output

Item	Description
BGP local router ID	BGP local router ID.
Status codes	Status of a routing entry.
Total Number of Routes	Total number of routes in the routing table.
Network	Indicates the network address in the MBGP routing table.
NextHop	Indicates the next-hop address for the packet.
MED	Indicates the MED of the route.
LocPrf	Indicates the local preference.
PrefVal	Indicates the preferred value.
Path/Ogn	Indicates the AS_Path number and the Origin attribute.

Display detailed information about the specified invalid MBGP routes.

```
<HUAWEI> display bgp multicast routing-table 192.168.1.1
 BGP local router ID : 10.1.1.1
 Local AS number : 100
 Paths:   2 available, 0 best, 0 select
 BGP routing table entry information of 192.168.1.1/32:
 From: 10.2.2.2 (10.1.1.2)
 Route Duration: 00h01m15s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface:
 Original nexthop: 172.16.1.2
```



```
 Qos information : 0x0
 AS-path 200, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for IP
unreachable
 Not advertised to any peer yet

 BGP routing table entry information of 192.168.1.1/32:
 From: 10.1.1.2 (10.1.1.2)
 Route Duration: 00h01m15s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface:
 Original nexthop: 172.16.1.2
 Qos information : 0x0
 AS-path 200, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for IP
unreachable
 Not advertised to any peer yet
Not advertised to any peers yet
```

Table 7-141 Description of the display bgp multicast routing-table command output

Item	Description
BGP local router ID	ID of the local BGP device. The format is the same as the IPv4 address.
Local AS number	Local AS number.
Paths	Information about paths of BGP routes
BGP routing table entry information of 192.168.1.1/32	The following information is about 192.168.1.1/32 routing entries.
From	IP address of the router that sends the route. 10.1.1.2 is the IP address of the source interface of the peer with which the BGP connection is established, and 10.2.2.2 is the Router ID of the peer.
Route Duration	Duration of routes.
Relay IP Nexthop	Recursive next hop.
Relay IP Out-Interface	Recursive outbound interface.
Original nexthop	Original next hop.

Item	Description
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *precedence* command has been run and the *precedence* varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.
AS-path 200	AS_Path attribute.
origin incomplete	Well-known mandatory property. This property defines the origin of a path and records how a route turns to a BGP route. The property has the following three values: • IGP: The priority of this value is the highest. The origin property of the routes that are added to the BGP routing table by using the **network (BGP)** command is IGP. • EGP: The priority of this value is second to that of IGP. The origin property of the routes imported from EGP is EGP. • Incomplete: The priority of this value is the lowest. The value indicates the origin of a route is unknown. The origin property of the routes that are added to the BGP routing table by using the **import-route (BGP)** command is Incomplete.
MED	Multi-Exit discriminator of route.
localpref	Local priority.
pref-val	Value preferred by the protocol.
internal	The BGP route is an internal route.
pre 255	The BGP route preference is 255.

Item	Description
invalid for IP unreachable	Reason why a route is invalid: • invalid for route-policy not pass: The route does not match the route-policy. • invalid for supernet route: The route is a supernet route. • invalid for IP unreachable: The route fails to recurse to another route. • invalid for supernet route not advertise: No supernet routes are advertised. • invalid for supernet label route not advertise: No supernet labeled routes are advertised. • invalid for next-hop unreachable: The next-hop IP address is unreachable. • invalid for tunnel unreachable: The route fails to recurse to a tunnel.
Not advertised to any peer yet	The BGP route has not been advertised to any peer yet.

7.8.44 display bgp multicast routing-table as-path-filter

Function

The **display bgp multicast routing-table as-path-filter** command displays information about the routes that match the AS_Path filter.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table as-path-filter { *as-path-filter-number* | *as-path-filter-name* }

Parameters

Parameter	Description	Value
as-path-filter-number	Specifies the number of the matching AS_Path filter.	The value is an integer that ranges from 1 to 256.

Parameter	Description	Value
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The value is a string of 1 to 51 case-sensitive characters without spaces.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table as-path-filter** command to check information about the routes that match the AS_Path filter.

Example

Display information about the routes that match the AS_Path filter 1.

```
<HUAWEI> display bgp multicast routing-table as-path-filter 1
BGP local router ID is 10.14.14.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 15
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
 *>i  10.1.13.0/24       10.2.1.1       0          100       0       1?
 *>i  10.2.1.0/24        10.2.1.1       0          100       0       1?
 *>i  10.2.13.0/24       10.2.1.1       0          100       0       1?
 *>i  10.2.13.1/32       10.2.1.1       0          100       0       1?
```

Table 7-142 Description of the display bgp multicast routing-table as-path-filter command output

Item	Description
BGP local router ID	BGP local router ID
Network	Network address in the MBGP routing table
NextHop	Next hop address for forwarding packets
MED	Multi-Exit discriminator
LocPrf	Local preference
PrefVal	Preferred value of protocols
Path/Ogn	AS path and Origin attribute

7.8.45 display bgp multicast routing-table cidr

Function

The **display bgp multicast routing-table cidr** command displays routing information of classless inter-domain routing (CIDR).

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table cidr

Parameters

None.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table cidr** command to check routing information of CIDR.

Example

Display routing information about CIDR.

```
<HUAWEI> display bgp multicast routing-table cidr
 BGP Local router ID is 10.13.13.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 9
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
 *>   10.1.2.0/24        10.2.1.2                       0      2?
 *>   10.2.1.0/24        10.2.1.2       0               0      2?
 *>   10.7.1.0/24        10.2.1.2                       0      2?
 *>   10.7.1.1/32        10.2.1.2                       0      2?
 *>   10.7.1.2/32        10.2.1.2       0               0      2?
 *>   10.11.12.0/24      10.2.1.2       0               0      2?
 *>   10.11.12.1/32      10.2.1.2       0               0      2?
```


Table 7-143 Description of the display bgp multicast routing-table cidr command output

Item	Description
BGP local router ID	BGP local router ID
Status codes	Status of a routing entry
Total Number of Routes	Total number of routes in the routing table
Network	Network address in the MBGP routing table
NextHop	Next hop address for forwarding packets
MED	Multi-Exit discriminator
LocPrf	Local preference
PrefVal	Preferred value of protocols
Path/Ogn	AS path and Origin attribute

7.8.46 display bgp multicast routing-table community

Function

The **display bgp multicast routing-table community** command displays routing information of a specified MBGP community.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table community [*aa:nn* | *community-number*] & <1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

Parameters

Parameter	Description	Value
aa:nn	Specifies the MBGP community number. You can specify a maximum of 29 MBGP communities.	Both aa and nn are integers ranging from 0 to 65535.
community-number	Specifies the MBGP community number. You can specify a maximum of 29 community numbers.	The value is an integer that ranges from 0 to 4294967295.

Parameter	Description	Value
internet	Displays the MBGP routes with the Internet community attribute. The Internet community attribute indicates the matching routes sent to all remote peers.	-
no-advertise	Displays MBGP routes with the no-advertise community attribute. The no-advertise community attribute indicates that the matching routes are not sent to any peer.	-
no-export	Indicates MBGP routes with the no-export community attribute. The no-expert community attribute indicates that the matching routes are not advertised to other ASs but to other sub-ASs in the confederation.	-
no-export-subconfed	Displays MBGP routes with the no-export-subconfed community attribute. The no-export-subconfed community attribute indicates that the matching routes are not advertised outside the local AS.	-
whole-match	Indicates precise matching.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table community** command to check routing information of a specified MBGP community.

Example

Display MBGP routing information of Community 100:100.

```
<HUAWEI> display bgp multicast routing-table community 100:100
 BGP local router ID is 10.12.12.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 1
      Network         NextHop      MED    LocPrf    PrefVal    Community
 *>   10.1.1.0/24     10.2.1.1     0                0          <100:100>
```

Display MBGP routing information with the internet community attribute or that of Community 100:100.

```
<HUAWEI> display bgp multicast routing-table community 100:100 internet
```



```
BGP local router ID is 10.12.12.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

Total Number of Routes: 5
     Network          NextHop      MED      LocPrf    PrefVal   Community

*>   10.1.1.1/32      10.1.1.2     0                0         <100:100>,internet
*>   10.3.3.3/32      10.2.1.1     0                0         <100:100>,internet
*>   10.1.1.0/24      10.2.1.1     0                0         <100:100>,internet
*>   10.2.13.0/24     10.2.1.1      0                0         <100:100>,internet
```

Table 7-144 Description of the display bgp multicast routing-table community command output

Item	Description
BGP local router ID	BGP local router ID
Status codes	Status of a routing entry.
Total Number of Routes	Total number of routes in the routing table.
Network	Network address in the MBGP routing table
NextHop	Next hop address for forwarding packets
MED	Multi-Exit discriminator
LocPrf	Local preference
PrefVal	Preferred value of protocols
Community	Community attribute

7.8.47 display bgp multicast routing-table community-filter

Function

The **display bgp multicast routing-table community-filter** command displays the multicast routing information that matches a specified MBGP community list.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table community-filter { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

Parameters

Parameter	Description	Value
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
community-filter-name	Specifies the name of the community filter.	The value is a string of 1 to 51 case-sensitive characters without spaces.
whole-match	Indicates exact matching.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table community-filter** command to check the multicast routing information that matches a specified MBGP community list.

Example

Display the multicast routing information that matches a specified MBGP community list.

```
<HUAWEI> display bgp multicast routing-table community-filter 1
BGP local router ID is 10.12.12.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal    Community
*>   10.10.10.0/24      10.2.1.1       0                  0         <100:100>
```

Table 7-145 Description of the display bgp multicast routing-table community-filter command output

Item	Description
BGP local router ID	BGP local router ID
Network	Network address in the MBGP routing table
NextHop	Next hop address of packets

Item	Description
MED	Multi-Exit discriminator
LocPrf	Local preference
PrefVal	Preferred value
Community	Community attribute

7.8.48 display bgp multicast routing-table dampened

Function

The **display bgp multicast routing-table dampened** command displays dampened MBGP routes.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table [**statistics**] **dampened**

Parameters

Parameter	Description	Value
statistics	Displays statistics of dampened routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table dampened** command to check dampened MBGP routes.

Example

Display dampened MBGP routes.

```
<HUAWEI> display bgp multicast routing-table dampened
 BGP local router ID is 10.12.12.9
 Status codes: * - valid, > - best, d - damped,
```



```
            h - history,  i - internal, s - suppressed, S - Stale
            Origin : i - IGP, e - EGP, ? - incomplete
Total Number of Routes: 1
   Network          From            Reuse      Path/Ogn
d  10.1.0.0/16      10.2.1.1        00:45:05   1?
```

Table 7-146 Description of the display bgp multicast routing-table dampened command output

Item	Description
BGP local router ID	BGP local router ID
Network	Network address in the MBGP routing table
From	IP address of MBGP peer from which the route is received
Reuse	Reuse value
Path/Ogn	AS_Path number and Origin attribute

7.8.49 display bgp multicast routing-table dampening parameter

Function

The **display bgp multicast routing-table dampening parameter** command displays information about MBGP dampening parameters.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table dampening parameter

Parameters

None.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table dampening parameter** command to check information about MBGP dampening parameters.

Example

Display MBGP dampening parameters.

```
<HUAWEI> display bgp multicast routing-table dampening parameter
 Maximum Suppress Time(in second) : 3973
 Ceiling Value                    : 16000
 Reuse Value                      : 750
 HalfLife Time(in second)         : 900
 Suppress-Limit                   : 2000
```

Table 7-147 Description of the display bgp multicast routing-table dampening parameter command output

Item	Description
Maximum Suppress Time (in second)	Maximum suppression time (in seconds)
Ceiling Value	Ceiling value of the penalty
Reuse Value	Threshold for routes leaving the suppression state
HalfLife Time (in second)	Half life time of the reachable route
Suppress-Limit	Threshold for routes entering the suppression state

7.8.50 display bgp multicast routing-table different-origin-as

Function

The **display bgp multicast routing-table different-origin-as** command displays the routes with the same destination but different source AS numbers.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table different-origin-as

Parameters

None.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table different-origin-as** command to check the routes with the same destination but different source AS numbers.

Example

Display the routes with the same destination but different source AS numbers.

```
<HUAWEI> display bgp multicast routing-table different-origin-as
BGP local router ID is 10.13.13.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 8
      Network            NextHop        MED     LocPrf    PrefVal Path/Ogn
 *>   10.2.1.0/24        10.2.1.2       0                 0       2?
 *>                      0.0.0.0        0                 0       ?
 *>   10.2.0.0/16        10.2.1.2                         0       2?
 *>                      0.0.0.0        0                 0       ?
 *>   192.168.14.0       10.2.1.2                         0       2?
 *>                      0.0.0.0        0                 0       ?
 *>   192.168.14.1/32    10.2.1.2                         0       2?
 *>                      0.0.0.0        0                 0       ?
```

Table 7-148 Description of the display bgp multicast routing-table different-origin-as command output

Item	Description
BGP local router ID	BGP local router ID
Network	Network address in the MBGP routing table
NextHop	Next hop address for forwarding packets
MED	Multi-Exit discriminator
LocPrf	Local preference
PrefVal	Preferred value of protocols
Path/Ogn	AS path and Origin attribute

7.8.51 display bgp multicast routing-table flap-info

Function

The **display bgp multicast routing-table flap-info** command displays information about MBGP route flapping.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table flap-info [*ip-address* [*mask* [**longer-match**] | *mask-length* [**longer-match**]] | **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* }]

display bgp multicast routing-table flap-info regular-expression *as-regular-expression*

Parameters

Parameter	Description	Value
ip-address	Specifies an IPv4 network address.	The address is in dotted decimal notation.
mask	Specifies the mask of an IPv4 network address.	The address is in dotted decimal notation.
mask-length	Specifies the mask length of the IPv4 network address.	The value is an integer that ranges from 0 to 32.
longer-match	Indicates the longest prefix matching rules.	-
regular-expression *as-regular-expression*	Displays the statistics of route flapping that matches the AS_Path regular expression. *as-regular-expression* specifies the AS path regular expression.	The value is a string of 1 to 80 characters.
as-path-filter *as-path-filter-number*	Displays the statistics of the route flapping for the specified AS path list. *as-path-filter-number* indicates the matched AS path list number.	The value is an integer that ranges from 1 to 256.

Parameter	Description	Value
as-path-filter *as-path-filter-name*	Displays the statistics of the route flapping for the specified AS path list. *as-path-filter-name* indicates the matched AS path list name.	The value is a string of 1 to 51 case-sensitive characters without spaces.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table flap-info** command to check information about MBGP route flapping.

Example

Display information about MBGP route flapping.

```
<HUAWEI> display bgp multicast routing-table flap-info
BGP Local router ID is 10.12.12.9
 Status codes: * - valid, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
      Network          From            Flaps  Duration     Reuse  Path/Ogn
 *>   10.1.1.0/24      10.2.1.1        1      00:00:29           1?
```

Table 7-149 Description of the display bgp multicast routing-table flap-info command output

Item	Description
BGP Local router ID	BGP local router ID
Network	Network address in the MBGP routing table
From	IP address of MBGP peer from which the route is received
Flaps	Count of route flapping
Duration	Duration of route flapping
Reuse	Reuse value
Path/Ogn	AS_Path number and Origin attribute

7.8.52 display bgp multicast routing-table peer

Function

The **display bgp multicast routing-table peer** command displays the routes received from or sent to a specified MBGP peer.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table peer *peer-address* { **advertised-routes** [*network* [{ *mask* | *mask-length* } [**longer-prefixes**]]] | **received-routes** [**active**] | **accepted-routes** }

display bgp multicast routing-table peer *peer-address* **received-routes** *network* [{ *mask* | *mask-length* } [**longer-prefixes** | **original-attributes**]]

Parameters

Parameter	Description	Value
peer-address	Specifies the address of an MBGP peer.	The value is in dotted decimal notation.
advertised-routes	Indicates the routes sent to a specified peer.	-
network	Specifies the IPv4 network address.	The value is in dotted decimal notation.
mask	Specifies the address mask.	The value is in dotted decimal notation.
mask-length	Specifies the mask length.	The value is an integer that ranges from 0 to 32.
longer-prefixes	Uses the longest match rule to select routes.	-
received-routes	Indicates the routes received from a specified peer.	-
active	Displays the active routes received from the specified peer.	-
accepted-routes	Displays routing information that is received from neighbors and matches the filter policy.	-

Parameter	Description	Value
original-attributes	Displays the original attributes of a public route from a specified BGP peer before the route is filtered by the local import policy. To display such attributes, the **peer keep-all-routes** command must have been run.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table peer** command to check routes received of a specified MBGP peer.

Example

Display the routes sent to multicast peer 10.10.1.11.

```
<HUAWEI> display bgp multicast routing-table peer 10.10.1.11 advertised-routes
 BGP local router ID is 10.12.12.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 9
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
 *>i  10.1.2.0/24        10.7.1.2       0          100        0      ?
 *>   10.2.1.0/24        0.0.0.0        0                     0      ?
 *>   10.7.1.0/24        0.0.0.0        0                     0      ?
 *>i  10.7.1.0/24        10.7.1.2       0          100        0      ?
 *>i  10.7.1.1/32        10.7.1.2       0          100        0      ?
 *>   10.7.1.2/32        0.0.0.0        0                     0      ?
 *>   10.10.1.0/24       0.0.0.0        0                     0      ?
```

Table 7-150 Description of the display bgp multicast routing-table peer command output

Item	Description
BGP local router ID	BGP local router ID
Total Number of Routes	Total number of routes
Network	Network address in the MBGP routing table
NextHop	Next hop address for forwarding packets
MED	Multi-Exit discriminator
LocPrf	Local preference

Item	Description
PrefVal	Preferred value of protocols
Path/Ogn	AS path and Origin attribute

7.8.53 display bgp multicast routing-table regular-expression

Function

The **display bgp multicast routing-table regular-expression** command displays the routes that match the specified AS_Path regular expression.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table regular-expression *as-regular-expression*

Parameters

Parameter	Description	Value
as-regular-expression	Specifies the AS_Path regular expression.	The value is a string of 1 to 80 case-sensitive characters without spaces.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table regular-expression** command to check the routes that match the specified AS_Path regular expression.

Example

Display the routes that match the AS_Path regular expression 2.

```
<HUAWEI> display bgp multicast routing-table regular-expression 2
BGP local router ID is 10.13.13.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
```



```
          Origin : i - IGP, e - EGP, ? - incomplete
Total Number of Routes: 13
     Network          NextHop        MED    LocPrf    PrefVal Path/Ogn
*>   10.1.2.0/24      10.2.1.2                0         2?
*>   10.2.1.0/24      10.2.1.2       0        0         2?
*>   10.7.1.0/24      10.2.1.2                0         2?
*>   10.7.1.1/32      10.2.1.2                0         2?
*>   10.7.1.2/32      10.2.1.2       0        0         2?
*>   10.11.12.0/24    10.2.1.2       0        0         2?
*>   10.11.12.1/32    10.2.1.2       0        0         2?
```

Table 7-151 Description of the display bgp multicast routing-table regular-expression command output

Item	Description
BGP local router ID	BGP local router ID.
Network	Network address in the MBGP routing table.
NextHop	Next-hop address for packets.
MED	MED of the route.
LocPrf	Local preference.
PrefVal	Preferred value.
Path/Ogn	AS_Path number and Origin attribute.

7.8.54 display bgp multicast routing-table statistics

Function

The **display bgp multicast routing-table statistics** command displays statistics about routes in the MBGP routing table.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast routing-table statistics

display bgp multicast routing-table peer *ipv4-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

Parameters

Parameter	Description	Value
peer *ipv4-address*	Displays the number of routes with a specified peer address.	The value is in dotted decimal notation.
advertised-routes	Displays the number of routes advertised to a specified peer.	-
received-routes	Displays the number of routes received from a specified peer.	-
active	Displays the number of active routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp multicast routing-table statistics** command to check statistics about routes in the MBGP routing table.

Example

Display statistics about routes in the MBGP routing table.

```
<HUAWEI> display bgp multicast routing-table statistics
Total Number of Routes: 50
```

Table 7-152 Description of the display bgp multicast routing-table statistics command output

Item	Description
Total Number of Routes	Total number of routes in the MBGP routing table

7.8.55 display bgp multicast update-peer-group

Function

The **display bgp multicast update-peer-group** command displays information about MBGP update-groups.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

display bgp multicast update-peer-group [**index** *update-group-index*]

Parameters

Parameter	Description	Value
index *update-group-index*	Specifies the index of an update-group.	The value is an integer that ranges from 0 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the index of an update-group is specified, the command will display detailed information about the specified update-group.

Example

Display information about MBGP update-groups.

```
<HUAWEI> display bgp multicast update-peer-group
 The Multicast instance's update peer group number : 1
 Keep buffer update peer group number : 0
 BGP Version : 4

 Group ID : 0
 Group Type : external
 Addr Family : IPv4-MLC
 AdvMinTimeVal : 30
 Total Peers : 1
 Leader Peer : 192.168.1.2
 Peers List : 192.168.1.2
```

Table 7-153 Description of the display bgp multicast update-peer-group command output

Item	Description
The Multicast instance's update peer group number	Indicates the number of update-groups in the instance.
Keep buffer update peer group number	Number of packets in update-groups saved in the batch buffer.

Item	Description
BGP Version	Indicates the BGP version.
Group ID	Indicates the ID of the update-group.
Group Type	Indicates the type of the update-group, which can be one of the following: • external: indicates an EBGP update-group. • internal: indicates an IBGP update-group. • external-confed: indicates an EBGP update-group in the confederation. • internal-confed: indicates an IBGP update-group in the confederation. • unknown: indicates an update-group of an unknown type.
Addr Family	Indicates the address family.
AdvMinTimeVal	Indicates the minimum interval for sending Update packets with the same route prefix.
Total Peers	Indicates the total number of peers in an update-group.
Leader Peer	Indicates the representative of an update-group.
Peers List	Indicates a list of peers.

7.8.56 display bgp mvpn brief

Function

The **display bgp mvpn brief** command displays brief information about VPN instances in BGP MVPN address family.

Format

display bgp mvpn { **all** | **vpn-instance** *vpn-instance-name* } **brief**

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Displays brief information about MVPN and all VPN instances.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display bgp mvpn brief** command is used to display brief information about VPN instances in BGP MVPN address family.

Example

Display brief information about MVPN and all VPN instances.

```
<HUAWEI> display bgp mvpn all brief

MVPN:
 Rd Num           Peer Num         Route Num
 0                1                0

VPN-Instance(IPv4-MVPN-family):
 VPN-Instance Name   Peer Num         Route Num
 vrf0                0                0
 vrf1                0                0
 vrf11               0                0
 vrf12               0                0
 vrf13               0                0
 vrf14               0                0
 vrf2                0                20
 vrf3                0                20
 vrf4                0                24
 vrf5                0                24
 vrf6                0                0
 vrf7                0                0
 vrf8                0                20
```

Table 7-154 Description of the **display bgp mvpn all brief** command output

Item	Description
Rd Num	Number of Route Distinguishers (RDs).
Peer Num	Number of peers.

Item	Description
Route Num	Number of routes.
VPN-Instance Name	Name of a VPN instance.

7.8.57 display bgp mvpn routing-table

Function

The **display bgp mvpn routing-table** command displays the information of BGP MVPN routes.

Format

display bgp mvpn { **all** | **vpn-instance** *vpn-instance-name* } **routing-table** [*network* [{ *mask* | *mask-length* } [**longer-prefixes**]]]

display bgp mvpn route-distinguisher *route-distinguisher* **routing-table** [*network* [*mask* | *mask-length*]]

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table regular-expression** *as-regular-expression*

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table community** [*aa:nn* | *community-number*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp mvpn all routing-table peer *ipv4-address* { **advertised-routes** [*network*] | **received-routes** [**active**] }

display bgp mvpn all routing-table peer *ipv4-address* **received-routes** *network*

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Display all the BGP MVPN routing information.	-
vpn-instance *vpn-instance-name*	Displays the BGP MVPN routing information of the specified VPN instance.	The value must be an existing VPN instance name.
network	Specifies the IPv4 network address.	It is in dotted decimal notation.
mask | *mask-length*	Specifies mask in dotted decimal notation or mask-length.	The value of mask-length is an integer that ranges from 0 to 32.
longer-prefixes	Matches according to the mask longer than the specified length.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays BGP routing information of the specified Route Distinguisher (RD).	The RD formats are divided into the following types: • 2-byte AS number:32-bit user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. • Integral 4-byte AS number: 2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of x.y, where x and y are integers that range from 1 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • 32-bit IP address:16-bit user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.

Parameter	Description	Value
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matching AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals.
cidr	Displays the BGP MVPN routing information about the Classless Inter-Domain Routing (CIDR).	-
different-origin-as	Displays routes that have the same destination address but different source ASs.	-
regular-expression *as-regular-expression*	Specifies the matched AS regular expression.	The value is a string of 1 to 80 characters.
community-filter	Displays the BGP MVPN routing information that matches the specified BGP community filter.	-
community-filter-name	Specifies the name of the community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
whole-match	Indicates the exact matching.	-
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.

Parameter	Description	Value
community	Displays the BGP MVPN routing information of the specified BGP community attribute in the routing table.	-
aa:nn	Specifies the community attribute number.	Both aa and nn are integers ranging from 0 to 65535. You can set a maximum of 29 community numbers.
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
internet	Displays the BGP routes with Internet community attribute.	-
no-advertise	Displays the BGP routes with No-Advertise community attribute.	-
no-export	Displays the BGP routes with the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes with the No-Export-Subconfed community attribute.	-
peer *ipv4-address*	Displays the BGP MVPN routing information for the specified BGP peer.	It is in dotted decimal notation.
advertised-routes	Displays the BGP MVPN routing information advertised to the specified peer.	-
received-routes	Displays the BGP MVPN routing information received from the specified peer.	-

Parameter	Description	Value
active	Displays the active BGP MVPN routing information received from the specified peer.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify different parameters to view the specific routing information of BGP MVPN routes. If the length of the destination address mask of an IPv4 route is the same as that of its natural mask, the mask length is not displayed.

Example

Display all the BGP MVPN routing information.

```
<HUAWEI> display bgp mvpn all routing-table

 BGP Local router ID is 192.168.7.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PE: 3
 Route Distinguisher: 100:1

     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>   1.1.1.1/32       0.0.0.0        0                   0      ?

 Route Distinguisher: 200:1

     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>i  1.1.1.1/32       192.168.100.10  0         100       0      33 55?
 *>i  2.2.2.2/32       192.168.100.10  0         100       0      33 55?

 Total number of routes of IPv4-MVPN-family for vpn-instance vrf1: 3
     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>   1.1.1.1/32       0.0.0.0        0                   0      ?
 * i                  192.168.100.10  0         100       0      33 55?
 *>i  2.2.2.2/32       192.168.100.10  0         100       0      33 55?
```

Display all the BGP MVPN routing information of the VPN instance named vpna.

```
<HUAWEI> display bgp mvpn vpn-instance vpna routing-table

Total Number of Routes: 2

 BGP Local router ID is 2.2.2.9
 Status codes: * - valid, > - best, d - damped,
```



```
          h - history,  i - internal, s - suppressed, S - Stale
          Origin : i - IGP, e - EGP, ? - incomplete
    Network         NextHop       MED      LocPrf    PrefVal Path/Ogn

*>i  10.1.1.0/24     1.1.1.9       0        100       0     ?
*>i  10.2.1.0/24     3.3.3.9       0        100       0     ?
```

Display the BGP MVPN routing information of the specified RD.

```
<HUAWEI> display bgp mvpn route-distinguisher 100:1 routing-table

 BGP Local router ID is 192.168.7.1
 Status codes: * - valid, > - best, d - damped,
          h - history,  i - internal, s - suppressed, S - Stale
          Origin : i - IGP, e - EGP, ? - incomplete

 Route Distinguisher: 100:1

 Total Number of Routes: 1

     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>   1.1.1.1/32       0.0.0.0       0                   0     ?

 Total number of routes of IPv4-MVPN-family for vpn-instance vrf1: 3
     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>   1.1.1.1/32       0.0.0.0       0                   0     ?
 * i                 192.168.100.10  0         100       0     33 55?
 *>i  2.2.2.2/32       192.168.100.10  0         100       0     33 55?
```

Display all BGP MVPN routes of community 1000:100.

```
<HUAWEI> display bgp mvpn all routing-table community 1000:100

 BGP Local router ID is 2.2.2.2
 Status codes: * - valid, > - best, d - damped,
          h - history,  i - internal, s - suppressed, S - Stale
          Origin : i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PE: 1

 Route Distinguisher: 100:1

     Network          NextHop       MED       LocPrf    PrefVal Community

 *>i  10.1.1.0/24     1.1.1.1       0         100       0     <1000:100>

 Total number of routes of IPv4-MVPN-family for vpn-instance vpna: 1
     Network          NextHop       MED       LocPrf    PrefVal Community

 *>i  10.1.1.0/24     1.1.1.1       0         100       0     <1000:100>
```

Table 7-155 Description of the **display bgp mvpn routing-table** command output

Item	Description
Network	Network address in the BGP routing table.
Next Hop	Next Hop address through which the packet has to be sent.
MED	Multi_Exit discriminator.
LocPrf	Local preference.

Item	Description
PrefVal	Value preferred by the protocol.
Path/Ogn	AS_Path number and the attributes of Origin.
Community	Community attributes.

Display all BGP MVPN routes of community 1000:100 with the Internet community attribute.

```
<HUAWEI> display bgp mvpn all routing-table community 1000:100 internet

 BGP Local router ID is 2.2.2.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PE: 1

 Route Distinguisher: 100:1


     Network          NextHop       MED       LocPrf    PrefVal Community

 *>i  10.1.1.0/24       1.1.1.1         0         100        0     <1000:100>

 Total number of routes of IPv4-MVPN-family for vpn-instance vpna: 4
     Network          NextHop       MED       LocPrf    PrefVal Community

 *>i  10.1.1.0/24       1.1.1.1         0         100        0     <1000:100>
 *>   10.2.1.0/24       0.0.0.0         0                    0
 *                    10.2.1.1        0                    0
 *>   10.2.1.2/32       0.0.0.0
```

Display routes sent by 2.2.2.2.

```
<HUAWEI> display bgp mvpn all routing-table peer 2.2.2.2 received-routes

 BGP Local router ID is 1.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1

 Route Distinguisher: 2:2


     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>i  2.2.2.2/32        2.2.2.2         0         100        0     ?
```

Display detailed information about the BGP MVPN route with prefix 3.3.3.3.

```
<HUAWEI> display bgp mvpn all routing-table 3.3.3.3


 BGP local router ID : 192.168.1.108
 Local AS number : 100

 Total routes of Route Distinguisher(1:1): 2
 BGP routing table entry information of 3.3.3.3/32:
 From: 1.1.1.1 (192.168.1.104)
```



```
 Route Duration: 00h00m10s
 Relay IP Nexthop: 10.1.1.4
 Relay IP Out-Interface: Ethernet0/0/1
 Original nexthop: 1.1.1.1
 Qos information : 0x0
 Ext-Community:RT <111 : 1>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255, IGP
cost 10,
 PMSI: Flags 0, PIM-SM, label 0:0:0(0), Sender address: 3.3.3.3, Group address: 232.1.1.1
 Not advertised to any peer yet

 BGP routing table entry information of 3.3.3.3/32:
 From: 10.1.1.4 (192.168.1.104)
 Route Duration: 00h05m44s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface: Ethernet0/0/1
 Original nexthop: 10.1.1.4
 Qos information : 0x0
 Ext-Community:RT <111 : 1>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, pre 255, IGP cost 0, not
preferred for peer address,
 PMSI: Flags 0, PIM-SM, Label 0:0:0(0), Sender address: 3.3.3.3, Group address: 232.1.1.1
 Not advertised to any peer yet


 Total number of routes of IPv4-MVPN-family for vpn-instance 1: 2
 BGP routing table entry information of 3.3.3.3/32:
 From: 1.1.1.1 (192.168.1.104)
 Route Duration: 00h00m10s
 Original nexthop: 1.1.1.1
 Qos information : 0x0
 Ext-Community:RT <111 : 1>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255,
 PMSI: Flags 0, PIM-SM, Label 0:0:0(0), Sender address: 3.3.3.3, Group address: 232.1.1.1
 Not advertised to any peer yet

 BGP routing table entry information of 3.3.3.3/32:
 From: 10.1.1.4 (192.168.1.104)
 Route Duration: 00h05m45s
 Original nexthop: 10.1.1.4
 Qos information : 0x0
 Ext-Community:RT <111 : 1>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, pre 255, not preferred for
peer address,
 PMSI: Flags 0, PIM-SM, Label 0:0:0(0), Sender address: 3.3.3.3, Group address: 232.1.1.1
 Not advertised to any peer yet
```

Display detailed information about the route with prefix 2.2.2.2 sent by 2.2.2.2.

```
<HUAWEI> display bgp mvpn all routing-table peer 2.2.2.2 received-routes 2.2.2.2

 BGP local router ID : 1.1.1.1
 Local AS number : 100
 Route Distinguisher: 2:2

 BGP routing table entry information of 2.2.2.2/32:
 From: 2.2.2.2 (2.2.2.9)
 Route Duration: 18h27m03s
 Relay IP Nexthop: 10.1.1.2
 Relay IP Out-Interface: Ethernet0/0/1
 Original nexthop: 2.2.2.2
 Qos information : 0x0
 Ext-Community:RT <1 : 1>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255, IGP
cost 1
 PMSI: Flags 0, PIM-SM, label 0:0:0, Sender address: 2.2.2.2, Group address: 232.1.1.1
 Not advertised to any peer yet
```

Table 7-156 Description of the display bgp mvpn all routing-table peer received-routes command output

Item	Description
BGP local router ID	ID of the local BGP device. The format is the same as the IPv4 address.
Local AS number	Local AS number.
Route Distinguisher	Route distinguisher.
BGP routing table entry information of 2.2.2.2/32	The following information is about 2.2.2.2/32 routing entries.
From	IP address of the router that sends the route. 2.2.2.2 is the IP address of the source interface of the peer with which the BGP connection is established, and 2.2.2.9 is the Router ID of the peer.
Route Duration	Duration of routes.
Relay IP Nexthop	Recursive next hop.
Relay IP Out-Interface	Recursive outbound interface.
Original nexthop	Original next hop.
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *ip-precedence* command has been run and the *ip-precedence* varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.
Ext-Community	Extended community attribute.
AS-path Nil	AS_Path attribute, with Nil indicating that the attribute value is null.

Item	Description
origin incomplete	Well-known mandatory property. This property defines the origin of a path and records how a route turns to a BGP route. The property has the following three values: • IGP: The priority of this value is the highest. The origin property of the routes that are added to the BGP routing table by using the **network (BGP)** command is IGP. • EGP: The priority of this value is second to that of IGP. The origin property of the routes imported from EGP is EGP. • Incomplete: The priority of this value is the lowest. The value indicates the origin of a route is unknown. The origin property of the routes that are added to the BGP routing table by using the **import-route (BGP)** command is Incomplete.
localpref	Local priority.
pref-val	Value preferred by the protocol.
valid	The BGP route is a valid route.
internal	The BGP route is an internal route.
best	The BGP route is an optimal route.
select	The BGP route is a preferred route.
pre 255	The priority of the BGP route is 255.
IGP cost	Cost of the relied route.
PMSI: Flags 0, PIM-SM, label 0:0:0, Sender address: 2.2.2.2, Group address: 232.1.1.1	In NG MVPNs, P-Multicast Service Interfaces (PMSIs) are logical channels that transmit multicast VPN services over the public network. • Flags: Point-to-Multipoint (P2MP) tunnel flags. • PIM-SM: The PIM-SM protocol. • label: Label information about the PMSI. • Sender address: IP address of the sender PE. • Group address: IP address of the multicast group.

Item	Description
Not advertised to any peer yet	The BGP route has not been advertised to any peer yet.

7.8.58 display bgp mvpn routing-table statistics

Function

The **display bgp mvpn routing-table statistics** command displays the statistics of the BGP MVPN routes.

Format

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics**

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics regular-expression** *as-regular-expression*

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp mvpn { **all** | **route-distinguisher** *route-distinguisher* | **vpn-instance** *vpn-instance-name* } **routing-table statistics community** [*aa:nn* | *community-number*] & <1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp mvpn all routing-table peer *ipv4-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Displays all the statistics of the BGP MVPN routes.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays the BGP MVPN routing statistics of the specified Route Distinguisher (RD).	-
vpn-instance *vpn-instance-name*	Displays the BGP MVPN routing statistics of the specified VPN instance.	The value must be an existing VPN instance name.
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matched AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals.
cidr	Displays the BGP MVPN routing statistics about the Classless Inter-Domain Routing (CIDR).	-
different-origin-as	Displays the routes that have the same destination address but different source AS number.	-
regular-expression *as-regular-expression*	Indicates the matched AS regular expression.	The value is a string of 1 to 80 characters.
community-filter	Displays the BGP MVPN routing statistics that matches the specified BGP community filter.	-
community-filter-name	Specifies the name of the community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
whole-match	Indicates exact matching.	-

Parameter	Description	Value
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
community	Displays the BGP MVPN routing statistics of the specified BGP community attribute in the routing table.	-
aa:nn	Specifies the community number.	Both aa and nn are integers ranging from 0 to 65535. You can set a maximum of 29 community numbers.
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
internet	Displays the BGP routes with Internet community attribute.	-
no-advertise	Displays the BGP routes with the No-Advertise community attribute.	-
no-export	Displays the BGP routes with the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes with the No-Export-Subconfed community attribute.	-
peer *ipv4-address*	Displays the BGP MVPN routing statistics for the specified BGP peer.	It is in dotted decimal notation.
advertised-routes	Displays the BGP MVPN routing statistics advertised to the specified peer.	-
received-routes	Displays the BGP MVPN routing statistics from the specified peer.	-

Parameter	Description	Value
active	Displays the active BGP MVPN routing statistics received from the specified peer.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

None.

Example

Display all the statistics of the BGP MVPN routing table.

```
<HUAWEI> display bgp mvpn all routing-table statistics

 Total number of routes from all PE: 20

 Total number of routes of IPv4-MVPN-family for vpn-instance vpn1: 12
```

Table 7-157 Description of the **display bgp mvpn routing-table statistics** command output

Item	Description
Total number of routes from all PE	The number of the routes received from PEs in the BGP MVPN routing table
Total number of routes of IPv4-MVPN-family for vpn-instance	The number of the routes of the specified VPN instance in the BGP MVPN routing table

7.8.59 display bgp network

Function

The **display bgp network** command displays the routes imported into the BGP routing table by using the **network** command.

Format

display bgp network

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **network**

display bgp ipv6 network

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **network** (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
vpnv4	Displays the VPNv4 routes that are advertised by using the **network** command.	-
vpnv6	Displays the VPNv6 routes that are advertised by using the **network** command.	-
all	Displays all the VPNv4 routes that are advertised by using the **network** command.	-
vpn-instance *vpn-instance-name*	Displays the routes of a specified VPN instance that are advertised by using the **network** command.	The value must be an existing VPN instance name.
ipv6	Displays the IPv6 routes advertised by BGP.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

BGP cannot discover routes by itself. Run the **network (BGP)** command and the **import-route (BGP)** command to import routes from other protocols, such as IGP routes or static routes, to the BGP routing table.

The **display bgp network** command displays the routing information advertised by BGP through the **network (BGP)** command.

Precautions

BGP has multiple address families and the routing information about each address family is isolated from each other. The routing information about address families

advertised by BGP through the **network (BGP)** command can be displayed by specifying different parameters.

If no parameter is specified, the **display bgp network** command displays the routing information about the IPv4 unicast address family advertised by BGP through the **network (BGP)** command.

Example

Display information about routes that are imported using the **network** command.

```
<HUAWEI> display bgp network
 BGP Local Router ID is 10.1.1.9
 Local AS Number is 10(Public)
 Network          Mask            Route-policy

 10.2.0.0         255.255.0.0
 10.0.0.0         255.0.0.0       Policy1
 10.4.4.0         255.255.255.0
```

Table 7-158 Description of the display bgp network command output

Item	Description
Local AS Number	Indicates the local AS number.
Network	Indicates the locally-advertised network address.
Mask	Indicates the mask of the network address.
Route-policy	Indicates the used routing policy.

Display information about BGP VPNv4 routes that are imported using the **network** command.

```
<HUAWEI> display bgp vpnv4 all network
 BGP Local Router ID is 10.2.2.9
 Local AS Number is 100
VPN-Instance vrf1, Router ID 10.2.2.9:
 Network          Mask            Route-policy
 10.4.4.4         255.255.255.255
VPN-Instance vrf2, Router ID 2.2.2.9:
 Network          Mask            Route-policy
 10.5.5.5         255.255.255.255
```

Table 7-159 Description of the display bgp vpnv4 all network command output

Item	Description
VPN-Instance	Name of the VPN instance
Router ID	Router ID of the local BGP router.

Display IPv6 routes advertised by BGP.


```
<HUAWEI> display bgp ipv6 network
 BGP Local Router ID is 10.5.5.5
 Local AS Number is 100(PublicV6)
 Network          Prefix          Route-policy
 FC00:0:0:100::   64
 FC00:0:0:200::   64
```

Display VPNv6 routes locally advertised by BGP (by using the **network** command) for a specified VPN instance.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 network
 BGP Local Router ID is 10.1.1.1
 Local AS Number is 100
 Route Distinguisher: 100:1
 (vpn1)
 Network          Prefix          Route-policy
 FC00:0:0:2000::  100              policy1
```

Table 7-160 Description of the display bgp vpnv6 network command output

Item	Description
BGP Local Router ID	Indicates the ID of the local BGP device. The ID is in the same format as an IPv4 address.
Local AS Number	Indicates the local AS number.
Route Distinguisher	Indicates the route distinguisher for the VPN instance.
Prefix	Indicates the network address mask advertised by the local BGP device.

Display the VPNv6 routes locally advertised by BGP (by using the **network** command).

```
<HUAWEI> display bgp vpnv6 all network
 BGP Local Router ID is 10.2.2.9
 Local AS Number is 100
 Route Distinguisher: 100:4
 (vpn1)
 Network          Prefix          Route-policy
 FC00:0:0:1::1    128
 FC00:0:0:2::2    128
 Route Distinguisher: 100:5
 (vrf1)
 Network          Prefix          Route-policy
 FC00:0:0:3::3    128
 Route Distinguisher: 100:9
 (vrf2)
 Network          Prefix          Route-policy
 FC00:0:0:8::9    128
```

7.8.60 display bgp paths

Function

The **display bgp paths** command displays the AS_Path information of BGP routes.

Format

display bgp paths [*as-regular-expression*]

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **paths** [*as-regular-expression*]

display bgp ipv6 paths [*as-regular-expression*]

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **paths** [*as-regular-expression*] (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp { **mdt** | **mvpn** } { **all** | **vpn-instance** *vpn-instance-name* } **paths** [*as-regular-expression*]

NOTE

The **mdt,mvpn** parameter is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.

Parameters

Parameter	Description	Value
as-regular-expression	Specifies the regular expression used to match the AS_Path information.	The name is a string of 1 to 80 characters.
vpnv4	Displays the AS_Path information of the routes of a VPNv4 instance.	-
vpnv6	Displays the AS_Path information of the routes of a VPNv6 instance.	-
all	Displays the AS_Path information of all VPNv4 routes.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ipv6	Displays the path attributes of IPv6-BGP routes in the local BGP database.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The **display bgp paths** command displays the path attributes of BGP stored in the system.

BGP usually has a large number of path attributes. The **display bgp paths** command displays a lot of BGP path information for viewing. After configuring *as-regular-expression*, the **display bgp paths** command displays only the path attributes of BGP that match *as-regular-expression*. For details on a regular expression, see Filtering Output Information Based on the Regular Expression in "CLI Overview" in the *S300, S500, S2700, S5700, and S6700 V200R023C00 Configuration Guide - Basic Configuration*.

Precautions

BGP has a number of address families and the path attributes of BGP in each address family is stored independently. By default, the **display bgp paths** command displays only the path attributes of BGP in the IPv4 unicast address family. The path attributes of BGP in other address families can be displayed by specifying address family parameters.

Example

Display the AS_Path information.

```
<HUAWEI> display bgp paths

Total number of routes of IPv4-family for vpn-instance _public_: 6
Total Number of Paths: 1

Address         Refcount        MED  Path/Origin
1282430404             6          0  N?
```

Table 7-161 Description of the display bgp paths command output

Item	Description
Address	Indicates the address of the route in the local database, in hexadecimal notation.
Refcount	Indicates the number of times that the route is referenced.
MED	Indicates the MED of the route.
Path	Indicates the list of ASs through which the packet has to pass through.
Origin	Indicates the origin of the route.

7.8.61 display bgp peer

Function

The **display bgp peer** command displays information about BGP peers.

Format

display bgp [**vpnv4 vpn-instance** *vpn-instance-name*] **peer** [{ *group-name* | *ipv4-address* } **log-info** | [*ipv4-address*] **verbose**]

display bgp vpnv4 all peer [[*ipv4-address*] **verbose**]

display bgp vpls peer [[*ipv4-address*] **verbose**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp ipv6 peer [**verbose**]

display bgp ipv6 peer *ipv6-address* { **log-info** | **verbose** }

display bgp ipv6 peer *ipv4-address* **verbose**

display bgp vpnv6 all peer [[*ipv4-address*] **verbose**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv6 vpn-instance *vpn-instance-name* **peer** [{ *group-name* | *ipv6-address* } **log-info** | [*ipv6-address*] **verbose**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp l2vpn-ad peer [[*ipv4-address*] **verbose**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp l2vpn peer [[*ipv4-address*] **verbose**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp { **mdt** | **mvpn** } **all peer** [[*ipv4-address*] **verbose**](supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ipv4-address	Specifies the IPv4 address of a peer to be displayed.	It is in dotted decimal notation.

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
log-info	Indicates to display log information of the specified peer.	-
verbose	Indicates to display detailed peer information.	-
vpnv4	Indicates to display information about peers in a VPNv4 instance.	-
all	Indicates to display information about peers in all VPNv4 or VPNv6 instances.	-
vpls	Displays information about peers in the BGP-VPLS view.	-
ipv6-address	Specifies the IPv6 address of a peer to be displayed.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
vpnv6	Indicates to display BGP VPNv6 peer information.	-
l2vpn-ad	Displays information about peers in the BGP L2VPN-AD view.	-
l2vpn	Displays information about peers in the BGP L2VPN view.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The **display bgp peer** command displays information about a BGP peer. You can implement the following operations based on the command output:

- To check the status of BGP connections
- To check information about a BGP peer
- To check whether a BGP peer is successfully configured using the **peer as-number** command
- To check whether a BGP peer is successfully deleted using the **undo peer as-number** command

Precautions

BGP has multiple address families. By default, the **display bgp peer** command displays information about BGP peers in IPv4 unicast address family only. If you want to view information about BGP peers in another address family, you need to specify its address family parameter.

To view detailed information about a BGP peer, such as information about BGP timers, the number of sent and received routes, capacities supported, the number of sent and received BGP messages, and enabled functions, specify **verbose** in the command.

If **log-info** is specified in the command, log information about a BGP peer is displayed, including information about BGP peer flapping.

Example

Display peer information.

```
<HUAWEI> display bgp peer

 Status codes: * - Dynamic

 BGP Local router ID : 10.2.3.4
 local AS number : 10
 Total number of peers : 2                 Peers in established state : 1
 Total number of dynamic peers : 0

  Peer        V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.1.1.1    4   100        0        0     0 00:00:07        Idle       0
  10.2.5.6    4   200       32       35     0 00:17:49 Established       0
```

Table 7-162 Description of the display bgp peer command output

Item	Description
Status codes: * - Dynamic	Status code. If the value starts with an asterisk (*), the peer is a dynamic peer. Currently, the value can only be *** - Dynamic**.

Item	Description
BGP Local router ID	Indicates the ID of the BGP local router. **NOTE** If two ends have the same BGP local router ID, no BGP peer relationship can be established between them. In this situation, run the **router id** command in the BGP view on either end to change the BGP local router ID. Changing it to the IP address of a loopback interface is recommended.
local AS number	Indicates the local AS number.
Total number of peers	Indicates the number of peers.
Peers in established state	Indicates the number of peers in established state.
Total number of dynamic peers	Total number of dynamic BGP peers.
Peer	Indicates the IP address of the peer.
V	Indicates the BGP version used on the peer.
AS	Indicates the AS number.
MsgRcvd	Indicates the number of received messages.
MsgSent	Indicates the number of sent messages.
OutQ	Indicates the message to be sent to the specified peer.
Up/Down	Indicates the period of time during which a BGP session keeps the current state.

Item	Description
State	Status of the peer: ● Idle: indicates that BGP denies any request of entering. This is the initiatory status of BGP. Upon receiving a Start event, BGP initiates a TCP connection to the remote BGP peer, starts the ConnectRetry Timer with the initial value, detects a TCP connection initiated by the remote BGP peer, and changes its state to Connect. ● Idle(Admin): indicates that the peer relationship is shut down initiatively and no attempt is made to establish the peer relationship. If the **peer ignore** command is configured or the peer is set to the Down state through the MIB, the neighbor is in the Idle (Admin) state. ● No neg: The capabilities of the BGP peer's address family have not been negotiated. ● Idle(Ovlmt): indicates that the peer relationship is interrupted because the number of routes exceeds the upper threshold. After a BGP peer relationship is interrupted due to the running of the **peer route-limit** command, the status of the BGP peer relationship is displayed as Idle(Ovlmt). If the **reset bgp** command is not run, the BGP peer relationship will not be reestablished. ● Connect: indicates that BGP waits for the TCP connection to be set up before it determines whether to perform other operations. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the TCP connection fails, BGP restarts the ConnectRetry Timer

Item	Description
	with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value, initiates a TCP connection to the remote BGP peer, and stays in the Connect state. ● Active: indicates that BGP tries to set up a TCP connection. This is the intermediate status of BGP. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value and changes its state to Connect. – If BGP initiates a TCP connection with an unknown IP address, the TCP connection fails. When this occurs, BGP restarts the ConnectRetry Timer with the initial value and stays in the Active state. ● OpenSent: indicates that BGP has sent one Open message to its peer and waits for an Open message from the peer. – If there are no errors in the Open message received, BGP changes its state to OpenConfirm. – If there are errors in the Open message received, BGP sends a Notification message to the remote peer and changes its state to Idle. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues

Item	Description
	to detect a TCP connection initiated by the remote peer, and changes its state to Active. • OpenConfirm: indicates that BGP waits for a Notification message or a Keepalive message. – If BGP receives a Notification message, or the TCP connection fails, BGP changes its state to Idle. – If BGP receives a Keepalive message, BGP changes its state to Established. • Established: indicates that BGP peers can exchange Update, Notification and Keepalive packets. – If BGP receives an Update or a Keepalive message, its state stays in Established. – If BGP receives a Notification message, BGP changes its state to Idle.
PrefRcv	Indicates the number of route prefixes sent from the peer.

Display detailed information about the peer 10.2.2.9.

```
<HUAWEI> display bgp peer 10.2.2.9 verbose

  BGP Peer is 10.2.2.9,  remote AS 100
  Type: IBGP link
  BGP version 4, Remote router ID 10.1.1.1
  Update-group ID: 1
  BGP current state: Established, Up for 00h57m53s
  BGP current event: RecvKeepalive
  BGP last state: Established
  BGP Peer Up count: 1
  Received total routes: 0
  Received active routes total: 0
  Received mac routes: 0
  Advertised total routes: 2
  Port: Local - 42796        Remote - 179
  Configured: Connect-retry Time: 32 sec
  Configured: Min Hold Time: 0 sec
  Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
  Received  : Active Hold Time: 180 sec
  Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
  Peer optional capabilities:
  Peer supports bgp multi-protocol extension
  Peer supports bgp route refresh capability
  Peer supports bgp 4-byte-as capability
  Address family IPv4 Unicast: advertised and received
  Address family IPv6 Unicast: received
```



```
Received: Total 60 messages
              Update messages              1
              Open messages                1
              KeepAlive messages           58
              Notification messages        0
              Refresh messages             0
Sent: Total 61 messages
              Update messages              2
              Open messages                1
              KeepAlive messages           58
              Notification messages        0
              Refresh messages             0
Authentication type configured: None
 Last keepalive received: 2012/03/06 19:17:37  Last keepalive sent    : 2012/03/06 19:17:37  Last update
received: 2012/03/06 19:17:43  Last update    sent    : 2012/03/06 19:17:37

 Minimum route advertisement interval is 15 seconds
 Optional capabilities:
 Route refresh capability has been enabled
 4-byte-as capability has been enabled
 Listen-only has been configured
 Peer's BFD has been enabled
 Peer Preferred Value: 0
 Routing policy configured:
 No routing policy is configured
```

Table 7-163 Description of the display bgp peer verbose command output

Item	Description
Type	Indicates the BGP link type, which can be IBGP or EBGP.
BGP version	Indicates the BGP version.
remote router ID	Indicates the router ID of the peer.
Update-group ID	Indicates the ID of the update-group to which the peer belongs.

Item	Description
BGP current state	Current state of BGP: • Idle: indicates that BGP denies any request of entering. This is the initiatory status of BGP. Upon receiving a Start event, BGP initiates a TCP connection to the remote BGP peer, starts the ConnectRetry Timer with the initial value, detects a TCP connection initiated by the remote BGP peer, and changes its state to Connect. • Idle(Admin): indicates that the peer relationship is shut down initiatively and no attempt is made to establish the peer relationship. If the **peer ignore** command is configured or the peer is set to the Down state through the MIB, the neighbor is in the Idle (Admin) state. • No neg: The capabilities of the BGP peer's address family have not been negotiated. • Idle(Ovlmt): indicates that the peer relationship is interrupted because the number of routes exceeds the upper threshold. After a BGP peer relationship is interrupted due to the running of the **peer route-limit** command, the status of the BGP peer relationship is displayed as Idle(Ovlmt). If the **reset bgp** command is not run, the BGP peer relationship will not be reestablished. • Connect: indicates that BGP waits for the TCP connection to be set up before it determines whether to perform other operations. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the TCP connection fails, BGP restarts the ConnectRetry Timer

Item	Description
	with the initial value, continues to detect a TCP connection initiated by the remote peer, and changes its state to Active. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value, initiates a TCP connection to the remote BGP peer, and stays in the Connect state. ● Active: indicates that BGP tries to set up a TCP connection. This is the intermediate status of BGP. – If the TCP connection succeeds, BGP stops the ConnectRetry Timer, sends an Open message to the remote peer, and changes its state to OpenSent. – If the ConnectRetry Timer has expired before a TCP connection is established, BGP restarts the timer with the initial value and changes its state to Connect. – If BGP initiates a TCP connection with an unknown IP address, the TCP connection fails. When this occurs, BGP restarts the ConnectRetry Timer with the initial value and stays in the Active state. ● OpenSent: indicates that BGP has sent one Open message to its peer and waits for an Open message from the peer. – If there are no errors in the Open message received, BGP changes its state to OpenConfirm. – If there are errors in the Open message received, BGP sends a Notification message to the remote peer and changes its state to Idle. – If the TCP connection fails, BGP restarts the ConnectRetry Timer with the initial value, continues

Item	Description
	to detect a TCP connection initiated by the remote peer, and changes its state to Active. • OpenConfirm: indicates that BGP waits for a Notification message or a Keepalive message. – If BGP receives a Notification message, or the TCP connection fails, BGP changes its state to Idle. – If BGP receives a Keepalive message, BGP changes its state to Established. • Established: indicates that BGP peers can exchange Update, Notification and Keepalive packets. – If BGP receives an Update or a Keepalive message, its state stays in Established. – If BGP receives a Notification message, BGP changes its state to Idle.
BGP current event	Indicates the current BGP event.
BGP last state	Indicates the last BGP status, which may be Idle, Connect, Active, OpenSent, OpenConfirm, or Established.
BGP Peer Up count	Indicates the flapping count of a BGP peer in a specified period of time.
Received total routes	Indicates the number of received route prefixes.
Received active routes total	Indicates the number of received active route prefixes.
Received mac routes	Number of MAC routes received.
Advertised total routes	Indicates the number of sent route prefixes.
Port	Indicates the port number. • Local: indicates the local port number, which is always 179. BGP uses TCP at the transport layer. • Remote: indicates the port number used on the peer.

Item	Description
Configured	Indicates locally configured timers. • Connect-retry Time: indicates the ConnectRetry interval for a BGP peer or peer group, in seconds. When BGP initiates a TCP connection, the ConnectRetry timer is stopped if the TCP connection is established successfully. If the first attempt to establish a TCP connection fails, BGP tries again to establish the TCP connection after the ConnectRetry timer expires. • Min Hold Time: indicates the minimum Holdtime configured on the local device, in seconds. • Active Hold Time: indicates the hold time. If BGP does not receive any Keepalive message from the peer in the hold time, BGP considers that the peer is Down and then instructs other peers to remove the routes that are sent from the peer. • Keep Alive Time: indicates the interval for sending Keepalive messages to the peer. BGP peers send Keepalive messages to each other periodically to maintain their relationships.
Received : Active Hold Time	Indicates the hold time on the peer.
Negotiated : Active Hold Time	Indicates the hold time agreed between the BGP peers after capability negotiation.
Address family IPv4 Unicast	Indicates the IPv4 unicast address family.
Address family IPv6 Unicast	Indicates the IPv6 unicast address family.

Item	Description
Received	Indicates the number of packets received from a peer. • Total: indicates the total number of messages received from a peer. • Update messages: indicates the number of Update messages received from a peer. • Open messages: indicates the number of Open messages received from a peer. • KeepAlive messages: indicates the number of Keepalive messages received from a peer. • Notification messages: indicates the number of Notification messages received from a peer. • Refresh messages: indicates the number of route-refresh messages received from a peer.
Sent	Indicates the number of messages sent to a peer. • Total: indicates the total number of messages sent to a peer. • Update messages: indicates the number of Update messages sent to a peer. • Open messages: indicates the number of Open messages sent to a peer. • KeepAlive messages: indicates the number of Keepalive messages sent to a peer. • Notification messages: indicates the number of Notification messages sent to a peer. • Refresh messages: indicates the number of route-refresh messages sent to a peer.
Authentication type configured	Indicates the authentication type configured.

Item	Description
Last keepalive received	Indicates the time when the Keepalive packet is received last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is set through the **clock timezone** command; DST indicates that the daylight saving time is set through the **clock daylight-saving-time** command.
Last keepalive sent	Indicates the time when the Keepalive packet is sent last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is set through the **clock timezone** command; DST indicates that the daylight saving time is set through the **clock daylight-saving-time** command.

Item	Description
Last update received	Indicates the time when the Update packet is received last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is set through the **clock timezone** command; DST indicates that the daylight saving time is set through the **clock daylight-saving-time** command.
Last update sent	Indicates the time when the Update packet is sent last time. It can be in the following formats: ● YYYY/MM/DD HH:MM:SS ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM DST ● YYYY/MM/DD HH:MM:SS UTC ±HH:MM ● YYYY/MM/DD HH:MM:SS DST UTC±HH:MM indicates that a time zone is set through the **clock timezone** command; DST indicates that the daylight saving time is set through the **clock daylight-saving-time** command.
Minimum route advertisement interval is 15 seconds	Indicates the minimum interval between route advertisements. ● The minimum interval for advertising EBGP routes is 30 seconds. ● The minimum interval for advertising IBGP routes is 15 seconds.
Optional capabilities	(Optional) Indicates the peer-supported capabilities.
Route refresh capability has been enabled	Indicates that route refreshing has been enabled.

Item	Description
4-byte-as capability has been enabled	Indicates that the 4-byte AS number capability is enabled.
Listen-only has been configured	Indicates that only connection requests are snooped and no connections will be initiated proactively.
Peer Preferred Value	Indicates the preferred value of the peer.
Routing policy configured	Indicates the configured routing policy.
Peer's BFD has been enabled	Indicates that BFD has been enabled on the peer.

Display log information on BGP peer 10.1.1.2.

```
<HUAWEI> display bgp peer 10.1.1.2 log-info
Peer : 10.1.1.2
 Date/Time     : 2011/13/06 11:53:21
 State         : Up
 Date/Time     : 2011/13/06 11:53:09
 State         : Down
 Error Code    : 6(CEASE)
 Error Subcode : 4(Administrative Reset)
 Notification  : Receive Notification
 Date/Time     : 2011/13/06 10:34:05
 State         : Up
```

Table 7-164 Description of the display bgp peer 1.1.1.2 log-info command output

Item	Description
Error Code	Error code.
Error Subcode	Error subcode.
Notification	Notification packet sent or received by a peer.

Display information about the BGP peer of the VPN instance vrf1.

```
<HUAWEI> display bgp vpnv4 vpn-instance vrf1 peer

 Status codes: * - Dynamic

 BGP local router ID : 10.1.1.9
 Local AS number : 100
 VPN-Instance vrf1, router ID 10.1.1.1:
 Total number of peers : 1                 Peers in established state : 1
 Total number of dynamic peers : 0

  Peer        V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.1.1.1    4 65410      207      192     0 02:59:49 Established       1
```

Display detailed information about the BGP peer of the VPN instance vpna.


```
<HUAWEI> display bgp vpnv4 vpn-instance vpna peer verbose

  BGP Peer is 10.1.1.1,  remote AS 200
  Type: EBGP link
  BGP version 4, remote router ID 10.1.1.1

  Update-group ID: 1
  BGP current state: Established, Up for   03h01m22s
  BGP current event: RecvKeepalive
  BGP last state: OpenConfirm
  BGP Peer Up count: 1            Received total routes: 0            Received active routes total: 0
Received mac routes: 0            Advertised total routes: 3
  Port:  Local - 3722     Remote - 179
  Configured: Connect-retry Time: 32 sec
  Configured: Min Hold Time: 0 sec
  Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
  Received  : Active Hold Time: 180 sec
  Negotiated: Active Hold Time: 180 sec
  Peer optional capabilities:
  Peer supports bgp multi-protocol extension
  Peer supports bgp route refresh capability
  Peer supports bgp 4-byte-as capability
  Address family IPv4 Unicast: advertised and received

 Received: Total 76 messages
           Update messages              0
           Open messages                5
           KeepAlive messages           71
           Notification messages        0
           Refresh messages             0

 Sent: Total 91 messages
           Update messages              0
           Open messages                10
           KeepAlive messages           77
           Notification messages        4
           Refresh messages             0
Authentication type configured: None
Last keepalive received: 2012/03/06 19:17:37
Last keepalive sent    : 2012/03/06 19:17:37
Last update    received: 2012/03/06 19:17:43
Last update    sent    : 2012/03/06 19:17:37
Maximum allowed prefix number: 150000  Threshold: 75%
Minimum route advertisement interval is 30 seconds
Optional capabilities:
Route refresh capability has been enabled
4-byte-as capability has been enabled
Peer Preferred Value: 0
Routing policy configured:
No routing policy is configured
```

Display information about IPv6 peers.

```
<HUAWEI> display bgp ipv6 peer

 BGP Local router ID : 10.0.0.1
 local AS number : 100
  Total number of peers : 1                 Peers in established state : 1

  Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  FC00:0:0:20::21 4   200       17       19     0 00:09:59 Established       3
```

Display information about IPv6 peers.

```
<HUAWEI> display bgp ipv6 peer fc00:0:0:2001::1 verbose
  BGP Peer is FC00:0:0:2001::1,  remote AS 1
  Type: EBGP link
  BGP version 4, remote router ID 10.1.1.1
```



```
    Update-group ID: 1
    BGP current state: Established, Up for 00h00m59s
    BGP current event: RecvKeepalive
    BGP last state: OpenConfirm
    BGP Peer Up count: 2
    Received total routes: 0
    Received active routes total: 0
    Received mac routes: 0
    Advertised total routes: 0
    Port:  Local - 179      Remote - 49153
    Configured: Connect-retry Time: 32 sec
    Configured: Min Hold Time: 0 sec
    Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
    Received  : Active Hold Time: 180 sec
    Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
    Peer optional capabilities:
    Peer supports bgp multi-protocol extension
    Peer supports bgp route refresh capability
    Peer supports bgp 4-byte-as capability
    Address family IPv6 Unicast: advertised and received

 Received: Total 76 messages
             Update messages              0
             Open messages                5
             KeepAlive messages           71
             Notification messages        0
             Refresh messages             0

 Sent: Total 91 messages
             Update messages              0
             Open messages                10
             KeepAlive messages           77
             Notification messages        4
             Refresh messages             0
 Authentication type configured: None
 Last keepalive received: 2012/03/06 19:17:37
 Last keepalive sent    : 2012/03/06 19:17:37
 Last update    received: 2012/03/06 19:17:43
 Last update    sent    : 2012/03/06 19:17:37
 Minimum route advertisement interval is 30 seconds
 Optional capabilities:
 Route refresh capability has been enabled
 4-byte-as capability has been enabled
 listen-only has been configured
 Peer Preferred Value: 0
 Routing policy configured:
 No routing policy is configured
```

Display brief information about a VPNv6 peer.

```
<HUAWEI> display bgp vpnv6 all peer

 BGP local router ID : 10.1.1.1
 Local AS number : 100
 Total number of peers : 2                 Peers in established state : 2

  Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.2.2.2        4   100      210      220     0   02:42:55    Established    1

  Peer of IPv6-family for vpn instance :

  VPN-Instance vpn1, Router ID 10.4.4.4 :
  FC00:0:0:200::2 4 65410      205      178     0 02:42:53     Established    0
```

Display detailed information about the VPNv6 peer whose IPv4 address is 10.2.2.2.

```
<HUAWEI> display bgp vpnv6 all peer 10.2.2.2 verbose
    BGP Peer is 10.2.2.2,  remote AS 200,
```



```
    Type: IBGP link
    BGP version 4, remote router ID 10.2.2.2

Group ID : 0
Peer Local Interface Name: Vlanif10
Local Ifnet Tunnel: 0xb0010000
    BGP current state: Established, Up for 02h43m52s
    BGP current event: RecvKeepalive
    BGP last state: OpenConfirm
    BGP Peer Up count: 2
    Received total routes: 0
    Received active routes total: 0
    Received mac routes: 0
    Advertised total routes: 0
    Port:  Local - 53308    Remote - 179
    Configured: Connect-retry Time: 32 sec
    Configured: Min Hold Time: 0 sec
    Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
    Received  : Active Hold Time: 180 sec
    Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
    Peer optional capabilities:
    Peer supports bgp multi-protocol extension
    Peer supports bgp route refresh capability
    Peer supports bgp 4-byte-as capability
    Address family IPv4 Unicast: advertised and received
    Address family VPNv6: advertised and received

 Received: Total 76 messages
             Update messages              0
             Open messages                5
             KeepAlive messages           71
             Notification messages        0
             Refresh messages             0

 Sent: Total 91 messages
             Update messages              0
             Open messages                10
             KeepAlive messages           77
             Notification messages        4
             Refresh messages             0
 Last keepalive received: 2009-09-30 09:52:41
 Minimum route advertisement interval is 15 seconds
 Optional capabilities:
 Route refresh capability has been enabled
 4-byte-as capability has been enabled
 Connect-interface has been configured
 Peer Preferred Value: 0
 Routing policy configured:
 No routing policy is configured
```

Display detailed information about the BGP peer of an IPv6 address family-enabled VPN instance whose IPv6 address is FC00:0:0:2000::2.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 peer fc00:0:0:2000::2 verbose

    BGP Peer is FC00:0:0:2000::2,  remote AS 65410,
    Type: EBGP link
    BGP version 4, remote router ID 10.1.1.1

Group ID : 0
Peer Local Interface Name: Vlanif10
Local Ifnet Tunnel: 0xb0010000
    BGP current state: Established, Up for 02h39m36s
    BGP current event: KATimerExpired
    BGP last state: OpenConfirm
    BGP Peer Up count: 2
    Received total routes: 0
    Received active routes total: 0
    Received mac routes: 0
    Advertised total routes: 0
```



```
      Port:  Local - 49177    Remote - 179
      Configured: Connect-retry Time: 32 sec
      Configured: Min Hold Time: 0 sec
      Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
      Received  : Active Hold Time: 180 sec
      Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
      Peer optional capabilities:
      Peer supports bgp multi-protocol extension
      Peer supports bgp route refresh capability
      Peer supports bgp 4-byte-as capability
      Address family IPv6 Unicast: advertised and received

Received: Total 76 messages
              Update messages              0
              Open messages                5
              KeepAlive messages           71
              Notification messages        0
              Refresh messages             0

Sent: Total 91 messages
              Update messages              0
              Open messages                10
              KeepAlive messages           77
              Notification messages        4
              Refresh messages             0
Last keepalive received: 2009-09-30 09:51:52
Minimum route advertisement interval is 30 seconds
Optional capabilities:
Route refresh capability has been enabled
4-byte-as capability has been enabled
Peer Preferred Value: 0
Routing policy configured:
No routing policy is configured
```

7.8.62 display bgp resource

Function

The **display bgp resource** command displays statistics about BGP specification information.

Format

display bgp resource

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To check statistics about BGP specification information, run the **display bgp resource** command. This command displays only the total number of supported BGP peer sessions, including the numbers of various configured sessions.

Example

Display statistics about BGP specification information.

```
<HUAWEI> display bgp resource

 BGP Peer session support total number : 64

  Capacity Statistics     Used Number

  BGP Peer Session          3
  IBGP Peer Session         3
  EBGP Peer Session         0
```

Table 7-165 Description of the **display bgp resource** command output

Item	Description
BGP Peer session support total number	Maximum number of BGP peer sessions that are supported.
Capacity Statistics	BGP peer session type.
Used Number	Number of currently configured peer sessions.
BGP Peer Session	Total number of currently configured BGP peer sessions.
IBGP Peer Session	Number of IBGP peer sessions, including confederation IBGP and common IBGP peer sessions.
EBGP Peer Session	Number of EBGP peer sessions, including confederation EBGP and common EBGP peer sessions.

7.8.63 display bgp routing-table

Function

The **display bgp routing-table** command displays information about BGP routes. Information about specified routes can be displayed by specifying different parameters.

Format

display bgp routing-table [**verbose** | *ipv4-address* [{ *mask* | *mask-length* } [**longer-prefixes**]]]

display bgp routing-table [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp routing-table regular-expression *as-regular-expression*

display bgp routing-table community-filter { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp routing-table community [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp routing-table peer *ipv4-address* **received-routes** *ipv4-address* [{ *mask* | *mask-length* } [**longer-prefixes** | **original-attributes**]]

display bgp routing-table peer *ipv4-address* { **advertised-routes** [*ipv4-address* [{ *mask* | *mask-length* } [**longer-prefixes**]]] | **received-routes** [**active**] }

display bgp routing-table peer *ipv4-address* **accepted-routes**

display bgp routing-table time-range *start-time end-time*

Parameters

Parameter	Description	Value
verbose	Displays detailed information about BGP public network routes.	-
ipv4-address	Specifies an IPv4 network address.	The value is in dotted decimal notation.
mask | *mask-length*	Specifies a mask in dotted decimal notation or the mask length.	-
longer-prefixes	Matches any route whose prefix mask is longer than the specified length.	-
as-path-filter	Displays the routes that match a specified filter.	-
as-path-filter-number	Specifies the number of the matching AS_Path filter.	The value is an integer in the range from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The value is a string of 1 to 51 case-sensitive characters. It cannot contain spaces.

Parameter	Description	Value
cidr	Displays Classless InterDomain Routing (CIDR) information.	-
regular-expression *as-regular-expression*	Specifies the regular expression used to match the AS_Path information.	The value is a string of 1 to 80 characters.
different-origin-as	Displays routes that have the same destination address but different source AS numbers.	-
community-filter	Displays the routes that match a specified BGP community filter.	-
community-filter-name	Specifies the name of a community filter.	The name is a string of 1 to 51 characters. The string cannot contain only digits.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer in the range from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer in the range from 100 to 199.
community	Displays the routes carrying the specified BGP community attribute in the routing table.	-
community-number	Specifies a community number.	The value is an integer in the range from 0 to 4294967295.
aa:*nn*	Specifies a community number. A maximum of 29 community numbers can be set.	The values of *aa* and *nn* are both an integer in the range from 0 to 65535.
internet	Displays the BGP routes carrying the Internet community attribute.	-
no-advertise	Displays the BGP routes carrying the No-Advertise community attribute.	-

Parameter	Description	Value
no-export	Displays the BGP routes carrying the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes carrying the No-Export-Subconfed community attribute.	-
whole-match	Indicates exact matching.	-
peer *ipv4-address*	Displays routes of a specified peer.	-
advertised-routes	Displays the routes advertised to a specified peer.	-
received-routes	Displays the routes received from a specified peer.	-
active	Displays the active routes received from a specified peer.	-
original-attributes	Displays the original attributes of a public route from a specified BGP peer before the route is filtered by the local import policy. To display such attributes, the **peer keep-all-routes** command must have been run.	-
accepted-routes	Displays the routes that are received from the peer and filtered through a routing policy.	-

Parameter	Description	Value
time-range *start-time end-time*	Displays BGP routes that flap within the specified time period. For example, the value **0d0h5m0s** of *start-time* indicates 5 minutes before the current time. The value **0d0h10m0s** of *end-time* indicates 10 minutes before the current time. All BGP routes with the keepalive time in the range of 5 minutes to 10 minutes are displayed.	The value ranges of *start-time* and *end-time* are both **0d0h0m0s** to **10000d23h59m59s**.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

You can check information about specified routes by specifying different parameters.

The **display bgp routing-table** command is used to display active and inactive BGP routes on the public network.

Example

Display information about all BGP routes.

```
<HUAWEI> display bgp routing-table
 BGP Local router ID is 10.1.1.2
 Status codes: * - valid, > - best, d - damped, h - history,
               i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 4
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
*    10.1.1.0/24        10.1.1.1       0                    0         100?
*    10.1.1.2/32        10.1.1.1       0                    0         100?
*>   10.1.1.0/24        10.1.1.1       0                    0         100?
*>   10.1.1.0/24        10.1.1.1       0                    0         100?
```

Display routes with the community attribute.

```
<HUAWEI> display bgp routing-table community
 BGP Local router ID is 10.1.1.2
 Status codes: * - valid, > - best, d - damped, h - history,
               i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
 Total Number of Routes: 4
      Network            NextHop        MED        LocPrf    PrefVal Community
```



```
*    10.1.1.0/24        10.1.1.1        0                0      no-export
*    10.1.1.2/32        10.1.1.1        0                0      no-export
*>   10.1.1.0/24        10.1.1.1        0                0      no-export
*>   10.1.1.0/24        10.1.1.1        0                0      no-export
```

Table 7-166 Description of the **display bgp routing-table** command output

Item	Description
BGP Local router ID	Router ID of the local BGP device.
Network	Network address in the BGP routing table.
NextHop	Next-hop address for packets.
MED	MED of a route.
LocPrf	Local preference.
PrefVal	Preferred value (PrefVal).
Path/Ogn	AS_Path number and Origin attribute.
Community	Community attribute information.

Display detailed information about a specified route.

```
<HUAWEI> display bgp routing-table 10.1.1.1
 BGP local router ID : 10.2.3.107
 Local AS number : 100
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.1.1.1/32:
 Imported route.
 From: 10.1.1.2 (10.1.1.2)
 Route Duration: 0d00h01m33s
 Direct Out-interface: Vlanif10
 Original nexthop: 10.1.1.2
 Qos information : 0x0
 AS-path 200, origin incomplete, MED 0, pref-val 0, valid, external, best, select, pre 255
 Advertised to such 1 peers:
    10.1.1.2
```

Table 7-167 Description of the **display bgp routing-table** command output

Item	Description
Local AS number	Local AS number.
Paths	Route selection result.
BGP routing table entry information of	Routing entry information.
Imported route	Route imported into the BGP routing table using the **import-route** command.
From	IP address of the device that advertises the route.

Item	Description
Route Duration	Route duration.
Direct Out-interface	Directly connected outbound interface.
Original nexthop	Original next hop. **NOTE** If no routes to the BGP original next hop are available in the IP routing table, BGP routes cannot be advertised. In this case, find out why there is no route to the original next hop (this fault is generally associated with IGP or static routes).
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* value varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *precedence* command has been run and the *precedence* value varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.
AS-path	AS_Path attribute, with **Nil** indicating that the attribute value is null.
origin	Origin attribute of a BGP route. The value can be: • IGP: If a route is added to the BGP routing table using the **network** command, its origin is IGP. • EGP: indicates that the route is obtained through EGP. • Incomplete: indicates that the origin of the route is unknown. For example, the origin attribute of the routes imported using the **import-route** command is Incomplete.
valid	Valid route.
external	External route.
best	Optimal route.
select	Preferred route.
pre	BGP route preference.

Display detailed information about a specified invalid BGP route.

```
<HUAWEI> display bgp routing-table 192.168.1.1
 BGP local router ID : 10.1.1.1
 Local AS number : 100
 Paths:   2 available, 0 best, 0 select
 BGP routing table entry information of 192.168.1.1/32:
 From: 10.1.1.2 (10.1.1.2)
 Route Duration: 00h01m31s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface:
 Original nexthop: 172.16.1.2
 Qos information : 0x0
 AS-path 200, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for IP
unreachable
 Not advertised to any peer yet
```

Table 7-168 Description of the **display bgp routing-table** command output

Item	Description
BGP local router ID	ID of the local BGP router. The format is the same as an IPv4 address.
Local AS number	Local AS number.
Paths	Information about paths of BGP routes.
BGP routing table entry information of 192.168.1.1/32	The following information is about the route to 192.168.1.1/32.
From	IP address of the device that advertises the route.
Route Duration	Route duration.
Relay IP Nexthop	IP address of the recursive next hop.
Relay IP Out-Interface	Outbound interface obtained when the route recurses to another route.
Original nexthop	Original next hop.
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* value varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *precedence* command has been run and the *precedence* value varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.

Item	Description
AS-path 200	AS_Path attribute.
origin incomplete	Origin attribute of a route. The value can be: • IGP: If a route is added to the BGP routing table using the **network (BGP)** command, its origin is IGP. • EGP: indicates that the route is obtained through EGP. • Incomplete: indicates that the origin of the route is unknown. For example, the origin of the routes that are added to the BGP routing table using the **import-route (BGP)** command is Incomplete.
MED	MED of a route.
localpref	Local preference.
pref-val	Preferred value (PrefVal).
internal	Internal route.
pre 255	The BGP route preference is 255.
invalid for IP unreachable	Reason why a route is invalid: • invalid for route-policy not pass: The route does not match the route-policy. • invalid for supernet route: The route is a supernet route. • invalid for IP unreachable: The route fails to recurse to another route. • invalid for supernet route not advertise: No supernet routes are advertised. • invalid for supernet label route not advertise: No supernet labeled routes are advertised. • invalid for next-hop unreachable: The next-hop IP address is unreachable. • invalid for tunnel unreachable: The route fails to recurse to a tunnel.
Not advertised to any peer yet	Flag indicating that the BGP route has not been advertised to any peer yet.

7.8.64 display bgp routing-table dampened

Function

The **display bgp routing-table dampened** command displays BGP dampened routes.

Format

display bgp [**vpnv4 vpn-instance** *vpn-instance-name*] **routing-table** [**statistics**] **dampened**

display bgp ipv6 routing-table [**statistics**] **dampened**

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table** [**statistics**] **dampened** (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table** [**statistics**] **dampened** (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
vpnv4	Displays the BGP routes of a VPNv4 instance.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
statistics	Displays the statistics of dampened routes.	-
ipv6	Displays all dampened IPv6 routes.	-
vpnv6	Displays the BGP routes of a VPNv6 instance.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays the Route Distinguisher dampened BGP routes.	The RD formats are divided into the following types: • 2-byte AS number:4-byte user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot be 0s at the same time. That is, an RD cannot be 0:0. • Integral 4-byte AS number:2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • IPv4-address:2-byte user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the length of the destination address mask of an IPv4 route is the same as that of its natural mask, the mask length is not displayed after the command is run.

Example

Display BGP dampened routes.

```
<HUAWEI> display bgp routing-table dampened

 BGP Local router ID is 10.1.41.102
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 8
      Network            From            Reuse      Path/Origin
 d  10.6.244.0/23      10.10.41.247      01:06:25  65534 4837 174 11096 6356i
 d  10.17.79.0/24      10.10.41.247      01:06:25  65534 837 3356 23504 29777i
 d  10.17.110.0/24     10.10.41.247      01:06:25  65534 837 3356 23504 29777i
 d  10.57.144.0/20     10.10.41.247      01:06:25  65534 4837 10026 9924 18429,18429i
 d  10.76.216.0/24     10.10.41.247      01:06:25  65534 4837 701 26959i
 d  10.78.142.0/24     10.10.41.247      01:06:25  65534 4837 701 26959i
 d  10.115.136.0/23    10.10.41.247      01:06:25  65534 4837 701 26956i
 d  10.243.170.0/24    10.10.41.247      01:06:25  65534 4837 701 26959i
```

Table 7-169 Description of the display bgp routing-table dampened command output

Item	Description
Network	Indicates the network address in the BGP routing table.
From	Indicates the IP address of the peer that receives the routes.
Reuse	Indicates the reuse value (in seconds).
Path/Origin	Indicates the AS_Path number and the Origin attribute.

Display IPv6 dampened routes in the BGP routing table.

```
<HUAWEI> display bgp ipv6 routing-table dampened



 BGP Local router ID is 10.0.0.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale,
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 4
  d  Network : FC00:0:0:1:1::                PrefixLen : 48
     From    : FC00:0:0:1::2:2               Reuse     : 01:06:26
     Path/Ogn: 65001?

  d  Network : FC00:0:0:1:2::                PrefixLen : 48
     From    : FC00:0:0:1::2:2               Reuse     : 01:06:26
     Path/Ogn: 65001?

  d  Network : FC00:0:0:1:3::                PrefixLen : 48
     From    : FC00:0:0:1::2:2               Reuse     : 01:06:26
     Path/Ogn: 65001?

  d  Network : FC00:0:0:1:4::                PrefixLen : 48
```



```
From    : FC00:0:0:1::2:2                Reuse     : 01:06:26
Path/Ogn: 65001?
```

Table 7-170 Description of the **display bgp ipv6 routing-table dampened** command output

Item	Description
PrefixLen	Prefix length

7.8.65 display bgp routing-table dampening parameter

Function

The **display bgp routing-table dampening parameter** command displays configured BGP route dampening parameters.

Format

display bgp [**vpnv4 vpn-instance** *vpn-instance-name*] **routing-table dampening parameter**

display bgp ipv6 routing-table dampening parameter

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table dampening parameter** (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 all routing-table dampening parameter (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
vpnv4	Displays the BGP route dampening parameters of a VPNv4 instance.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ipv6	Displays configured BGP4+ route dampening parameters.	-
vpnv6	Displays the BGP route dampening parameters of a VPNv6 instance.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp routing-table dampening parameter** command to check configured BGP route dampening parameters.

Example

Display BGP route dampening parameters.

```
<HUAWEI> display bgp routing-table dampening parameter
 Maximum Suppress Time(in second) : 3973
 Ceiling Value                    : 16000
 Reuse Value                      : 750
 HalfLife Time(in  second)        : 900
 Suppress-Limit                   : 2000
```

Table 7-171 Description of the display bgp routing-table dampening parameter command output

Item	Description
Maximum Suppress Time (in second)	Maximum time taken for route suppression, in seconds.
Ceiling Value	Penalty ceiling.
Reuse Value	Threshold for a route to be unsuppressed.
HalfLife Time(in second)	Half life of a reachable route, in seconds.
Suppress-Limit	Threshold for a route to be suppressed.

Display BGP4+ route dampening parameters.

```
<HUAWEI> display bgp ipv6 routing-table dampening parameter
 Maximum Suppress Time(in second) : 3069
 Ceiling Value                    : 16000
 Reuse Value                      : 750
 HalfLife Time(in  second)        : 900
 Suppress-Limit                   : 2000
```

7.8.66 display bgp routing-table flap-info

Function

The **display bgp routing-table flap-info** command displays statistics about BGP route flapping.

Format

display bgp [**vpnv4 vpn-instance** *vpn-instance-name*] **routing-table flap-info** [**regular-expression** *as-regular-expression*]

display bgp [**vpnv4 vpn-instance** *vpn-instance-name*] **routing-table flap-info** { **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *network-address* [{ *mask* | *mask-length* } [**longer-match**]] }

display bgp ipv6 routing-table flap-info [**regular-expression** *as-regular-expression*]

display bgp ipv6 routing-table flap-info { **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *network-address* [*prefix-length* [**longer-match**]] }

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table flap-info** [**regular-expression** *as-regular-expression*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table flap-info** { **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *network-address* [*prefix-length* [**longer-match**]] } (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table flap-info** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *ipv4-address* [{ *mask* | *mask-length* } [**longer-match**]]] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table flap-info** [**regular-expression** *as-regular-expression*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
vpnv4	Displays statistics about route flapping of a VPNv4 instance.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
regular-expression *as-regular-expression*	Displays statistics about the flapping routes that match the AS_Path regular expression.	The value is a string of 1 to 80 characters.
as-path-filter *as-path-filter-number*	Specifies the number of an AS_Path filter.	It is an integer that ranges from 1 to 256.

Parameter	Description	Value
as-path-filter *as-path-filter-name*	Specifies the name of an AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. When double quotation marks are used around the string, spaces are allowed in the string.
network-address	Specifies the network address related to dampened routes.	-
mask	Specifies the network mask.	It is in dotted decimal notation.
mask-length	Specifies the mask length.	The value is an integer that ranges from 0 to 32.
longer-match	Matches a route with a longer prefix.	-
prefix-length	Specifies the prefix length.	The value is an integer that ranges from 0 to 128.
vpnv6	Displays statistics about route flapping of a VPNv6 instance.	-
ipv4-address	Displays statistics about the flapping routes that match the IPv4 prefix.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays the Route Distinguisher dampened BGP routes.	The RD formats are divided into the following types: • 2-byte AS number:4-byte user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot be 0s at the same time. That is, an RD cannot be 0:0. • Integral 4-byte AS number:2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0:0. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • IPv4-address:2-byte user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp routing-table flap-info** command to check statistics about BGP route flapping.

Example

Display statistics about BGP route flapping.

```
<HUAWEI> display bgp routing-table flap-info


 BGP Local router ID is 10.20.200.201
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed,
               Origin codes: i - IGP, e - EGP, ? - incomplete

Total Number of Routes: 7
     Network         From           Flaps  Duration     Reuse  Path/Ogn

 d  10.1.1.0        10.20.200.200   5      00:00:36  00:40:47  600i
 *>  10.1.1.0       10.20.200.202   1      00:04:07            100?
 d  10.1.2.0        10.20.200.200   5      00:00:36  00:40:47  600i
 *>  10.1.2.0       10.20.200.202   1      00:04:07            100?
 d  10.1.3.0        10.20.200.200   5      00:00:36  00:40:47  600i
 d  10.1.4.0        10.20.200.200   5      00:00:36  00:40:47  600i
 d  10.1.5.0        10.20.200.200   5      00:00:36  00:40:47  600i
```

Table 7-172 Description of the display bgp routing-table flap command output

Item	Description
Network	Network address in the BGP routing table.
From	IP address of the peer that receives the routes.
Flaps	Total number of times of route flapping.
Duration	Total time length of flapping.
Reuse	Reuse value.
Path/Ogn	AS_Path number and Origin attribute.

Display statistics about BGP4+ route flapping.

```
<HUAWEI> display bgp ipv6 routing-table flap-info


 BGP Local router ID is 10.53.53.53
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 3
  d  Network : FC00:0:0:1::                    PrefixLen : 96
     From    : FC00:0:0:1::1                   Flaps     : 8
     Duration: 00:02:11                        Reuse     : 00:49:21
     Path/Ogn: 100?

  d  Network : FC00:0:0:2::2                   PrefixLen : 128
```



```
   From    : FC00:0:0:1::1                 Flaps     : 5
   Duration: 00:00:18                      Reuse     : 00:41:06
   Path/Ogn: 100?

d  Network : FC00:0:0:2::3                   PrefixLen : 128
   From    : FC00:0:0:1::1                 Flaps     : 5
   Duration: 00:00:18                      Reuse     : 00:41:06
   Path/Ogn: 100?
```

7.8.67 display bgp routing-table label

Function

The **display bgp routing-table label** command displays labeled routes in the BGP routing table.

Product	Support
S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5731-S, S5731S-S, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table label include-mask**

display bgp routing-table label [**statistics** | **include-mask**]

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table** [**statistics**] **label**

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table** [**statistics**] **label**

display bgp ipv6 routing-table [**statistics**] **label**

Parameters

Parameter	Description	Value
include-mask	Displays labeled routes carrying masks.	-

Parameter	Description	Value
statistics	Indicates statistics about labeled routes.	-
vpnv4	Displays VPNv4 labeled routes.	-
all	Displays the labeled routes of all VPN instances.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ipv6	Displays BGP4+ labeled routes.	-
vpnv6	Displays VPNv6 labeled routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp routing-table label** command to check labeled routes in the BGP routing table.

Example

Display BGP labeled routes of all VPN instances.

```
<HUAWEI> display bgp vpnv4 all routing-table label
 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total number of routes from all PE: 3
 Route Distinguisher: 100:1

      Network          NextHop          In/Out Label
 *>i   10.22.22.22      10.3.3.9         NULL/1036

 Route Distinguisher: 100:4

      Network          NextHop          In/Out Label
 *>    10.1.2.0         10.1.2.1         1037/NULL
 *>    10.11.11.11      127.0.0.1        1038/NULL

 VPN-Instance vpn1, router ID 1.1.1.9:
 Total Number of Routes: 1
```



```
     Network          NextHop          In/Out Label
*>i   10.22.22.22      10.3.3.9         NULL/1036

VPN-Instance vpn2, router ID 10.1.1.9:
Total Number of Routes: 0
```

Table 7-173 Description of the display bgp vpnv4 all routing-table label command output

Item	Description
Network	Network address in the BGP routing table.
NextHop	IP address of the reachable next hop.
In/Out Label	Incoming label and outgoing label.

Display BGP4+ labeled routes.

```
<HUAWEI> display bgp ipv6 routing-table label
Total Number of Routes: 3

 Network        Prefix    NextHop          Label
 FC00:0:0:1::4  128       FC00:0:0:2::2    1024
 FC00:0:0:1::5  128       FC00:0:0:2::2    1025
 FC00:0:0:1::6  128       FC00:0:0:2::2    1026
```

Table 7-174 Description of the display bgp ipv6 routing-table label command output

Item	Description
Prefix	IP prefix.

7.8.68 display bgp routing-table peer no-advertise

Function

The **display bgp routing-table peer no-advertise** command displays the routes that a device is prevented from advertising to a specified peer in different address families.

Format

display bgp routing-table peer *ipv4-address* **no-advertise** *network* [*mask* | *mask-length*]

display bgp ipv6 routing-table peer *ipv6-address* **no-advertise** *ipv6–network* [*prefix-length*]

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table peer** *ipv4-address* **no-advertise** *network* [*mask* | *mask-length*]

display bgp vpnv4 all routing-table peer *ipv4-address* **no-advertise** *network* [*mask* | *mask-length*] (supported only by the S5731-H, S5731S-H, S5731-S,

S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp ipv6 routing-table peer *ipv4-address* **no-advertise** *ipv6–network* [*prefix-length*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp multicast routing-table peer *ipv4-address* **no-advertise** *network* [*mask* | *mask-length*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv6 all routing-table peer *ipv4-address* **no-advertise** *ipv6–network* [*prefix-length*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table peer** *ipv6-address* **no-advertise** *ipv6–network* [*prefix-length*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp l2vpn-ad routing-table peer *ipv4-address* **no-advertise** *network* (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
peer *ipv4-address*	Displays the routes that a device is prevented from advertising to the specified IPv4 peer.	The value is in dotted decimal notation.
network	Displays the routes with the specified destination IPv4 address that are prevented from being advertised.	The value is in dotted decimal notation.
mask	Specifies the subnet mask of the specified IPv4 address.	The value is in dotted decimal notation.
mask-length	Specifies the subnet mask length of the specified IPv4 address.	The value is an integer that ranges from 0 to 32.
vpnv4	Displays the routes that a device is prevented from advertising to a specified peer in the VPNv4 address family.	-

Parameter	Description	Value
all	Displays the routes that a device is prevented from advertising to a specified peer in all VPN instances of the current address family.	-
vpn-instance *vpn-instance-name*	Displays the routes that a device is prevented from advertising to a specified peer in the specified VPN instance of the current address family.	The value must be an existing VPN instance name.
multicast	Displays the routes that a device is prevented from advertising to a specified peer in the IPv4 multicast address family.	-
ipv6	Displays the routes that a device is prevented from advertising to a specified peer in the IPv6 unicast address family.	-
peer *ipv6-address*	Displays the routes that a device is prevented from advertising to the specified IPv6 peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
ipv6–network	Displays the routes with the specified destination IPv6 address that are prevented from being advertised.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the prefix length of an IPv6 address.	The value is an integer that ranges from 0 to 128.
vpnv6	Displays the routes that a device is prevented from advertising to a specified peer in the VPNv6 address family.	-
l2vpn-ad	Displays the routes that a device is prevented from advertising to a specified peer in the L2VPN-AD address family.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To check the routes that a device is prevented from advertising to a specified peer in different address families, run the **display bgp routing-table peer no-advertise** command. The command output also includes the reasons why the routes are not advertised.

You can specify an address family to check the routes that a device is prevented from advertising to a specified peer in the address family. If you do not specify any address family, the command output shows the routes that a device is prevented from advertising to a specified peer in the IPv4 unicast address family.

Example

Display the route 10.3.3.4/32 that a device is prevented from advertising to peer with IP address 192.168.1.1 in the IPv4 unicast address family.

```
<HUAWEI> display bgp routing-table peer 192.168.1.1 no-advertise 10.3.3.4 32
 BGP routing table entry information of 10.3.3.4/32:
 From: 192.168.1.1 (10.1.1.1)
 Route Duration: 22h36m17s
 Relay IP Nexthop: 10.1.1.1
 Relay IP Out-Interface: Vlanif100
 Original nexthop: 192.168.1.1
 Qos information : 0x0
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, active, pre 255,
IGP cost 1
 No advertise without config reflect-client
```

Table 7-175 Description of the **display bgp routing-table peer no-advertise** command output

Item	Description
From	Peer from which the route was received.
Route Duration	Duration of the route.
Relay IP Nexthop	Next-hop IP address to which the route recurses.
Relay IP Out-Interface	Outbound interface to which the route recurses.
Original nexthop	Original next-hop IP address of the route.

Item	Description
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *precedence* command has been run and the *precedence* varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.
AS-path	AS_Path attribute. **Nil** indicates that the attribute value is null.
origin	Origin attribute: • IGP: indicates that the route is added to the BGP routing table using the **network (BGP)** command. • EGP: indicates that the route is learned through the EGP protocol. • Incomplete: indicates that the origin of the route cannot be identified. For example, if a route is imported using the **import-route (BGP)** command, its origin is **Incomplete**.
MED	MED of the route.
localpref	Local_Pref of the BGP route.
pref-val	PrefVal of the route.
valid	Valid route.
internal	Internal route.
best	Optimal route.
select	Selected route.
active	Active route.
pre 255	Priority of the route (255 in this example).
IGP cost 1	IGP cost of the route (1 in this example).

Item	Description
No advertise without config reflect-client	Reason why the route is not advertised: • No advertise with no-adv flag: The route carries the No-Advertise flag. • No advertise with peer default-route-advertise in vt family: The local device is configured to advertise only default routes to the peer in the VT address family. • No advertise with no-adv flag in community attr: The route carries the No-Advertise community attribute. • No advertise with no-export flag in community attr: The route carries the No-Export community attribute. • No advertise with no-export-subconfed flag in community attr: The route carries the No_Export_Subconfed community attribute. • No advertise by detail route with config aggregate: Specific routes for summarization are not advertised. • No advertise with config active-route-advertise: Only the active routes in the IP routing table are advertised. • No advertise with config peer upe and no export policy: The specified peer is a UPE, and the route fails to match the export policy. • No advertise without config reflect-client: No RR is configured. • No advertise with config reflect-client but without config reflect between-clients: Route reflection among RR clients is disabled. • No advertise for other reason: Other reasons prevent the route from being advertised.

7.8.69 display bgp routing-table peer statistics

Function

The **display bgp routing-table peer statistics** command displays statistics about received and advertised BGP routes.

Format

display bgp [**multicast** | **ipv6**] **routing-table peer statistics**

display bgp [**l2vpn** | **vpls** | **l2vpn-ad**] **routing-table peer statistics**

display bgp { **vpnv4** | **vpnv6** } { **all** | **vpn-instance** *vpn-instance-name* } **routing-table peer statistics**

display bgp l2vpn-ad routing-table [**all** | **vpls-ad** | **vpws** | **vpls**] **peer statistics**

NOTE

The **l2vpn**, **l2vpn-ad**, **vpls**, **vpls-ad**, **vpnv4**, and **vpnv6** parameters are supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H.

Parameters

Parameter	Description	Value
multicast	Displays statistics about routes in the IP multicast routing table.	-
l2vpn	Displays statistics about routes in the BGP-L2VPN address family.	-
vpls	Displays statistics about routes in the BGP VPLS address family.	-
ipv6	Displays statistics about routes in the BGP IPv6 unicast address family.	-
l2vpn-ad	Displays statistics about routes in the BGP L2VPN-AD address family.	-
vpnv4	Displays statistics about routes in the BGP-VPNv4 address family.	-
vpnv6	Displays statistics about routes in the BGP-VPNv6 address family.	-
all	Displays statistics about all types of routes.	-
vpn-instance *vpn-instance-name*	Specifies a VPN instance name.	The value must be an existing VPN instance name.
vpls-ad	Displays statistics about VPLS-AD routes.	-

Parameter	Description	Value
vpws	Displays statistics about virtual private wire service (VPWS) routes.	-
vpls	Displays statistics about virtual private LAN service (VPLS) routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

To check statistics of received and advertised BGP routes, run the **display bgp routing-table peer statistics** command. You can specify parameters as needed. If no parameter is specified in the command, statistics about BGP IPv4 unicast routes are displayed.

Example

Display statistics about received and advertised BGP IPv4 unicast routes.

```
<HUAWEI> display bgp routing-table peer statistics
 BGP local router ID : 10.1.1.1
 Local AS number : 100
 Total number of peers : 3                 Number of Peers in established state : 3

  Peer            Received routes         Advertised routes

  10.2.2.2             1                       1
  10.2.2.3             1                       1
  10.3.3.3             2                       0
```

Table 7-176 Description of the **display bgp routing-table peer statistics** command output

Item	Description
BGP local router ID	Local router ID
Local AS number	AS number
Total number of peers	Total number of peers
Number of Peers in established state	Number of peers in the established state
Peer	IP address of a peer

Item	Description
Received routes	Total number of routes received from the peer
Advertised routes	Total number of routes advertised to the peer

7.8.70 display bgp routing-table statistics

Function

The **display bgp routing-table statistics** command displays statistics about BGP routes.

Format

display bgp routing-table statistics

display bgp routing-table statistics as-path-filter { *as-path-filter-number* | *as-path-filter-name* }

display bgp routing-table statistics cidr

display bgp routing-table statistics community [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp routing-table statistics community-filter { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp routing-table statistics dampened

display bgp routing-table statistics different-origin-as

display bgp routing-table statistics regular-expression *as-regular-expression*

display bgp routing-table peer *ipv4-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

Parameters

Parameter	Description	Value
as-path-filter *as-path-filter-number*	Displays the routes that match an AS-Path filter with the specified number.	It is an integer that ranges from 1 to 256.

Parameter	Description	Value
as-path-filter *as-path-filter-name*	Displays the routes that match an AS-Path filter with the specified name.	The name is a string of 1 to 51 characters without any space. It is case-sensitive.
cidr	Displays CIDR information.	-
community	Displays the routes carrying the specified BGP community attribute in the routing table.	-
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
aa:*nn*	Specifies the community attribute number.	Both *aa* and *nn* are integers ranging from 0 to 65535.
internet	Displays the matching routes that can be sent to any peer.	-
no-advertise	Displays the BGP routes carrying the No-Advertise community attribute.	-
no-export	Displays the BGP routes carrying the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes carrying the No-Export-Subconfed community attribute.	-
whole-match	Indicates exact matching.	-
community-filter	Displays the routes that match a specified BGP community filter.	-
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.

Parameter	Description	Value
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
community-filter-name	Specifies the name of a community filter.	The name is a string of 1 to 51 characters. The string cannot be all numbers.
dampened	Displays the statistics of BGP dampened roues.	-
different-origin-as	Displays routes that have the same destination address but different source AS numbers.	-
regular-expression *as-regular-expression*	Specifies the regular expression used to match the AS_Path information.	The value is a string of 1 to 80 characters.
peer *ipv4-address*	Displays the routing information for the specified BGP peer.	It is in dotted decimal notation.
advertised-routes	Displays the routes advertised to the specified peer.	-
received-routes	Displays the routes received from the specified peer.	-
active	Specifies the number of active routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display bgp routing-table statistics** command is used to display statistics about BGP routes on the public network.

The **display bgp routing-table statistics as-path-filter** command is used to display statistics about the BGP routes that match the specified AS_Path filter on the public network.

The **display bgp routing-table statistics cidr** command is used to display statistics about the BGP CIDR information of the public network.

The **display bgp routing-table statistics community** command is used to display statistics about the BGP routes carrying the specified community attribute on the public network.

The **display bgp routing-table statistics community-filter** command is used to display statistics about the BGP routes that match the specified community filter on the public network.

The **display bgp routing-table statistics dampened** command is used to display statistics about BGP dampened routes on the public network.

The **display bgp routing-table statistics different-origin-as** command is used to display statistics about the BGP routes with the same destination address but different source AS numbers on the public network.

The **display bgp routing-table statistics regular-expression** command is used to display statistics about the BGP routes whose AS_Path information matches the AS_Path regular expression on the public network.

Example

Display route statistics.

```
<HUAWEI> display bgp routing-table statistics
 Total Number of Routes: 4
```

Table 7-177 Description of the display bgp routing-table statistics command output

Item	Description
Total Number of Routes	Total number of routes in the routing table

7.8.71 display bgp update-peer-group

Function

The **display bgp update-group** command displays information about update-groups. By setting **index** in the command displays detailed information about a specified update-group. If no address family is specified, information about the update-group of the IPv4 unicast address family is displayed by default.

Format

display bgp update-peer-group [**index** *update-group-index*]

display bgp ipv6 update-peer-group [**index** *update-group-index*]

display bgp vpnv4 all update-peer-group [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **update-peer-group** [**index** *update-group-index*]

display bgp vpnv6 all update-peer-group [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv6 vpn-instance *vpn-instance-name* **update-peer-group** [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp l2vpn all update-peer-group [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpls all update-peer-group [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp l2vpn-ad update-peer-group [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp { **mdt** | **mvpn** } **all update-peer-group** [**index** *update-group-index*] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
index *update-group-index*	Specifies the index of an update-group.	The value is an integer that ranges from 0 to 65535.
ipv6	Displays information about the BGP update-groups of IPv6 routes.	-
vpnv4	Displays information about the BGP update-groups of a VPNv4 instance.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
vpnv6	Displays information about the BGP update-groups of a VPNv6 instance.	-

Parameter	Description	Value
all	Displays information about all the update-groups in a VPNv6 address family.	-
l2vpn	Displays information about L2VPN BGP update-groups.	-
vpls	Displays information about VPLS BGP update-groups.	-
l2vpn-ad	Displays information about L2VPN-AD BGP update-groups.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify the index of an update-group to view detailed information about the specified update-group. If no address family is specified, information about the update peer-groups in the IPv4 unicast address family is displayed by default.

Example

Display information about the update peer-groups in the IPv4 unicast address family.

```
<HUAWEI> display bgp update-peer-group
 The Public instance's update peer group number : 1
 Keep buffer update peer group number : 0
 BGP Version : 4
 Group ID : 0
 Group Type : external
 Addr Family : IPv4-UNC
 AdvMinTimeVal : 30
 Total Peers : 1
 Leader Peer : 192.168.1.2
 Peers List : 192.168.1.2
```

Table 7-178 Description of the display bgp update-peer-group command output

Item	Description
The Public instance's update peer group number	Number of update-groups in the instance

Item	Description
Keep buffer update peer group number	Number of packets in update-groups saved in the batch buffer
BGP Version	Indicates the BGP version.
Group ID	Indicates the ID of the update-group.
Group Type	Indicates the type of the update-group, which can be one of the following: • external: indicates that this is an EBGP peer group. • internal: indicates that this is an IBGP peer group. • unknown: indicates that the type is unknown.
Addr Family	Indicates the address family.
AdvMinTimeVal	Indicates the minimum interval for sending Update packets with the same route prefix.
Total Peers	Indicates the total number of peers in an update-group.
Leader Peer	Indicates the representative of an update-group.
Peers List	Indicates a list of peers.

Display information about the update-group with a specified index.

```
<HUAWEI> display bgp update-peer-group index 0
 Group ID : 0
 BGP Version : 4
 Group Type : external
 Addr Family : IPv4-UNC
 AdvMinTimeVal : 30
 Total Peers : 1
  Leader Peer : 192.168.1.2

 Total format packet number : 3
 Total send packet number : 3
 Total replicate packet number : 0
 The replication percentages(%) : 0

 Peers List : 192.168.1.2
```

Table 7-179 Description of the display bgp update-peer-group index command output

Item	Description
Total format packet number	Indicates the total number of formatted packets.
Total send packet number	Indicates the total number of sent packets.

Item	Description
Total replicate packet number	Indicates the total number of replicate packets, which equals the total number of sent packets minus the total number of formatted packets.
The replication percentages(%)	Indicates the replication percentage, which is obtained with the formula: (Total number of sent packets - Total number of formatted packets) x 100/Total number of formatted packets.

7.8.72 display bgp vpnv4 brief

Function

The **display bgp vpnv4 brief** command displays brief information about VPNv4 and VPN instances (IPv4 address family).

Format

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **brief**

Parameters

Parameter	Description	Value
all	Displays information about all VPNv4 and VPN instances (IPv4 address family).	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After the **display bgp vpnv4 brief** command is used to display information about VPNv4 and VPN instances (IPv4 address family), the VPN instances are displayed and arranged alphabetically by name.

Example

Display brief information about VPNv4 instances and all VPN instances (IPv4 address family).

```
<HUAWEI> display bgp vpnv4 all brief

VPNV4 :
  Rd Num           Peer Num         Route Num
  0                1                0

VPN-Instance(IPv4-family):
  VPN-Instance Name   Peer Num         Route Num
  vrf0                0                0
  vrf1                0                0
  vrf11               0                0
  vrf12               0                0
  vrf13               0                0
  vrf14               0                0
  vrf2                0                20
  vrf3                0                20
  vrf4                0                24
  vrf5                0                24
  vrf6                0                0
  vrf7                0                0
  vrf8                0                20
```

Table 7-180 Description of the display bgp vpnv4 brief command output

Item	Description
Rd Num	Indicates the number of RDs.
Peer Num	Indicates the number of peers.
Route Num	Indicates the number of routes.
VPN-Instance Name	Indicates the name of a VPN instance.

7.8.73 display bgp vpnv4 routing-table

Function

The **display bgp vpnv4 routing-table** command displays BGP routes of VPNv4 address family and the private networks.

Format

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table regular-expression** *as-regular-expression* (supported only by the S5731-H,

S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table regular-expression** *as-regular-expression*

display bgp vpnv4 route-distinguisher *route-distinguisher* **routing-table** [**verbose** | *ipv4-address* [*mask* | *mask-length*]] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table** [**verbose** | *ipv4-address* [*mask* [**longer-prefixes**] | *mask-length* [**longer-prefixes**]]]

display bgp vpnv4 all routing-table peer *ipv4-address* { **advertised-routes** [*network* [{ *mask* | *mask-length* } [**longer-prefixes**]]] | **received-routes** [**active**] } (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table peer** *ipv4-address* { **advertised-routes** [*network* [{ *mask* | *mask-length* } [**longer-prefixes**]]] | **received-routes** [**active**] }

display bgp vpnv4 all routing-table peer *ipv4-address* **received-routes** *network* [{ *mask* | *mask-length* } [**longer-prefixes**]] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table peer** *ipv4-address* **received-routes** *network* [{ *mask* | *mask-length* } [**longer-prefixes** | **original-attributes**]]

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* } (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table peer** *ipv4-address* **accepted-routes**

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table time-range** *start-time end-time* (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table time-range** *start-time end-time*

Parameters

Parameter	Description	Value
all	Displays all BGP VPNv4 routes.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays the BGP routes with the specified RD.	The RD formats are divided into the following types: • 2-byte AS number:4-byte user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot be 0s at the same time. That is, an RD cannot be 0:0. • Integral 4-byte AS number: 2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • IPv4-address:2-byte user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.
vpn-instance *vpn-instance-name*	Displays the BGP routes of a specified VPN instance.	The value must be an existing VPN instance name.

Parameter	Description	Value
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matching AS-Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS-Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.
cidr	Displays CIDR information.	-
different-origin-as	Displays the routes that have the same destination address but different source AS numbers.	-
regular-expression *as-regular-expression*	Specifies the regular expression used to match the AS_Path information.	The value is a string of 1 to 80 characters.
verbose	Displays detailed information about BGP VPNv4 routes.	-
ipv4-address	Specifies the destination address.	-
mask	Specifies a mask in dotted decimal notation.	The value is in dotted decimal notation.
mask-length	Specifies the mask length.	The value is an integer that ranges from 0 to 32.
peer *ipv4-address*	Displays routes of a specified peer.	-
advertised-routes	Displays the routes advertised to a specified peer.	-

Parameter	Description	Value
network	Specifies the IPv4 network address.	The value is in dotted decimal notation.
longer-prefixes	Matches any route whose prefix mask is longer than the specified length.	-
received-routes	Displays the routes received from a specified peer.	-
active	Displays the active routes received from a specified peer.	-
original-attributes	Displays the original attributes of a BGP route from a specified BGP peer before the route is filtered by the local import policy. To display such attributes, the **peer keep-all-routes** command must have been run.	-
community-filter	Displays the routes that match a specified BGP community filter.	-
community-filter-name	Specifies the name of a community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
whole-match	Indicates exact matching.	-
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
community	Displays the routes carrying the specified BGP community attribute in the routing table.	-

Parameter	Description	Value
community-number	Specifies the number of a community.	-
aa:*nn*	Specifies the community number. A maximum of 29 community numbers can be set.	Both *aa* and *nn* are integers ranging from 0 to 65535.
internet	Displays the matching routes that can be sent to any peer.	-
no-advertise	Displays the BGP routes carrying the No-Advertise community attribute.	-
no-export	Displays the BGP routes carrying the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes carrying the No-Export-Subconfed community attribute.	-
accepted-routes	Displays the routes that are received from a neighbor and accepted by a routing policy.	-
time-range *start-time end-time*	Displays information about BGP VPNv4 routes that have undergone status flapping during the specified period. For example, when *start-time* is set to 0d0h5m0s and *end-time* is set to 0d0h10m0s, information about all BGP VPNv4 routes whose lifetime ranges from 5 to 10 minutes is displayed.	In the values of *start-time* and *end-time*, d ranges from 0 to 10000, h ranges from 0 to 23, m ranges from 0 to 59, and s ranges from 0 to 59.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify different parameters to view the specific routing information.

When BGP routing table is displayed, if the length of the destination address mask of an IPv4 route is the same as that of its natural mask, the mask length is not displayed.

You can run the **display bgp vpnv4** { **all** | **vpn-instance** *vpn-instance-name* } **routing-table time-range** *start-time end-time* command to view BGP VPNv4 routes that flap within the specified time period. For example, if service traffic is abnormal or CPU usage of the device remains high within a certain time period, you can run this command to check whether route flapping occurs within the specified time period. The faulty route can be viewed in the command output, facilitating fault location.

Example

Display all the BGP routing information of the VPN instance named vpn1.

```
<HUAWEI> display bgp vpnv4 vpn-instance vpn1 routing-table

 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 VPN-Instance vpn1, Router ID 10.1.1.9:

 Total Number of Routes: 4
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.2.0/24        0.0.0.0         0                    0      ?
 *>   10.1.2.1/32        0.0.0.0         0                    0      ?
 *>   10.11.11.11/32     0.0.0.0         0                      0      ?
 *>i  10.22.22.22/32     10.3.3.9        0          100         0      ?
```

Display all the BGP VPNv4 routing information.

```
<HUAWEI> display bgp vpnv4 all routing-table


 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete



 Total number of routes from all PE: 4
 Route Distinguisher: 100:1


      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.22.22.22/32     10.3.3.9        0          100         0      ?

 Route Distinguisher: 100:4
```



```
     Network          NextHop       MED        LocPrf    PrefVal Path/Ogn

*>   10.1.2.0/24      0.0.0.0       0                    0      ?
*>   10.1.2.1/32      0.0.0.0       0                    0      ?
*>   10.11.11.11/32   0.0.0.0       0                    0      ?

VPN-Instance vpn1, Router ID 10.1.1.9:

Total Number of Routes: 4
     Network          NextHop       MED        LocPrf    PrefVal Path/Ogn

*>   10.1.2.0/24      0.0.0.0       0                    0      ?
*>   10.1.2.1/32      0.0.0.0       0                    0      ?
*>   10.11.11.11/32   0.0.0.0       0                    0      ?
*>i  10.22.22.22/32   10.3.3.9      0          100       0      ?
```

Display the VPNv4 routing information of the specified RD.

```
<HUAWEI> display bgp vpnv4 route-distinguisher 100:1 routing-table

BGP Local router ID is 10.1.1.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

Route Distinguisher: 100:1


Total Number of Routes: 3
     Network          NextHop       MED        LocPrf    PrefVal Path/Ogn

*>   10.1.1.0/24      0.0.0.0       0                    0      ?
*                     10.1.1.1      0                    0      65410?
*>   10.1.1.2/32      0.0.0.0       0                    0      ?

VPN-Instance vpna, Router ID 10.1.1.9:

Total Number of Routes: 7
     Network          NextHop       MED        LocPrf    PrefVal Path/Ogn

*>   10.1.1.0/24      0.0.0.0       0                    0      ?
                      10.1.1.1      0                    0      65410?
*>   10.1.1.2/32      0.0.0.0       0                    0      ?
*>   10.2.1.0/24      10.2.1.2      0                    0      ?
* i                   10.3.3.9      0          100       0      ?
*>   10.2.1.2/32      127.0.0.1     0                    0      ?
*>i  10.4.1.0/24      10.4.4.9      0          100       0      ?
```

Display all BGP VPNv4 routes of community 1000:100.

```
<HUAWEI> display bgp vpnv4 all routing-table community 1000:100

BGP Local router ID is 10.1.1.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


Total number of routes from all PE: 1

Route Distinguisher: 100:1


     Network          NextHop       MED        LocPrf    PrefVal Community

*>i  10.22.22.22/32   10.3.3.9      0          100       0      <1000:100>

VPN-Instance vpn1, Router ID 10.1.1.9:

Total Number of Routes: 1
```

```
     Network            NextHop        MED        LocPrf    PrefVal Community

*>i  10.22.22.22/32     10.3.3.9       0          100       0       <1000:100>

VPN-Instance vpn2, Router ID 10.1.1.9:

Total Number of Routes: 0
```

Displays BGP VPNv4 routes that flap within the specified time period.

```
<HUAWEI> display bgp vpnv4 all routing-table time-range 0d5h0m0s 1d5h0m0s

BGP Local router ID is 192.168.1.250
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


Route Distinguisher: 2:2

     Network            NextHop        Peer           Duration       Path/Ogn

*>i  0.0.0.0            10.2.2.2       10.2.2.2       05h30m11s      i
*>   10.2.3.4/32        0.0.0.0        0.0.0.0        05h30m42s      ?

VPN-Instance 1, Router ID 192.168.1.250:
     Network            NextHop        Peer           Duration       Path/Ogn

*>   10.2.3.4/32        0.0.0.0        0.0.0.0        05h30m42s      ?

VPN-Instance 2, Router ID 192.168.1.250:
     Network            NextHop        Peer           Duration       Path/Ogn

*>i  0.0.0.0            10.2.2.2       10.2.2.2       05h30m11s      i
*>   10.2.3.4/32        0.0.0.0        0.0.0.0        05h30m42s      ?
```

Table 7-181 Description of the display bgp vpnv4 routing-table command output

Item	Description
BGP Local router ID	ID of the local BGP device.
Network	Network address in the BGP routing table.
NextHop	Next Hop address through which the packet has to be sent
MED	MED of the route.
LocPrf	Local preference.
PrefVal	Preferred value.
Peer	Peer IP address
Duration	Route duration
Path/Ogn	AS_Path number and the origin attribute.
Community	Community attribute information.

Display information about a specified VPNv4 route.

```
<HUAWEI> display bgp vpnv4 all routing-table 192.168.2.0


 BGP local router ID : 10.2.2.9
 Local AS number : 100

 Total routes of Route Distinguisher(100:10): 2
 BGP routing table entry information of 192.168.2.0/24:
 Imported route.
 Label information (Received/Applied): NULL/1025
 From: 0.0.0.0 (0.0.0.0)
 Route Duration: 00h50m09s
 Direct Out-interface: Vlanif30
 Original nexthop: 192.168.2.1
 Qos information : 0x0
 Ext-Community:RT <1 : 1>
 AS-path Nil, origin incomplete, MED 0, pref-val 0, valid, local, best, select,
pre 255
 Advertised to such 1 peers:
    10.3.3.3
 BGP routing table entry information of 192.168.2.0/24:
 From: 192.168.2.2 (1.1.1.1)
 Route Duration: 00h48m50s
 Direct Out-interface: Vlanif30
 Original nexthop: 192.168.2.2
 Qos information : 0x0
 Ext-Community:RT <1 : 1>
 AS-path 10, origin incomplete, MED 0, pref-val 0, valid, external, pre 255, not
 preferred for route type
 Not advertised to any peer yet


 VPN-Instance vpna, Router ID 10.2.2.9:

 Total Number of Routes: 2
 BGP routing table entry information of 192.168.2.0/24:
 Imported route.
 From: 0.0.0.0 (0.0.0.0)
 Route Duration: 00h50m09s
 Direct Out-interface: Vlanif30
 Original nexthop: 192.168.2.1
 Qos information : 0x0
 AS-path Nil, origin incomplete, MED 0, pref-val 0, valid, local, best, select,
pre 0
 Advertised to such 1 peers:
    192.168.2.2
 BGP routing table entry information of 192.168.2.0/24:
 From: 192.168.2.2 (1.1.1.1)
 Route Duration: 00h48m51s
 Direct Out-interface: Vlanif30
 Original nexthop: 192.168.2.2
 Qos information : 0x0
 AS-path 10, origin incomplete, MED 0, pref-val 0, external, pre 255
 Not advertised to any peer yet
```

Table 7-182 Description of the display bgp vpnv4 routing-table command output

Item	Description
Local AS number	Local AS number.
Total routes of Route Distinguisher	Total number of VPNv4 routes with a specified RD.

Item	Description
BGP routing table entry information of x.x.x.x/x	The following information is about a specified BGP routing entry.
Imported route	Routes imported to BGP using the **import-route** command.
Label information (Received/ Applied)	Information about labels, including received and sent labels.
From	IP address of the route originator.
Route Duration	Route duration.
Direct Out-interface	Direct outbound interface.
Original nexthop	Original next hop.
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *precedence* command has been run and the *precedence* varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.
Ext-Community	Extended community attribute of BGP.
AS-path	AS_Path attribute.
origin	Origin attribute of the BGP route. • IGP: indicates that the origin attribute of a route added to the BGP routing table by using the **network** command is IGP. • EGP: indicates that the origin attribute of a route obtained by using EGP is EGP. • Incomplete: indicates that the origin attribute of a route whose source is unknown is Incomplete. For example, the origin attribute of the routes imported by using the **import-route** command is Incomplete.
pref-val	Preferred value.
valid	Valid BGP route.

Item	Description
external	The BGP route is learned from the EBGP peer.
best	The BGP route is the optimal route.
select	The BGP route is a preferred route.
pre 255	The preference of the BGP route is 255.
Advertised to such 1 peers	The BGP route has been advertised to one peer.
Not advertised to any peer yet	The BGP route has not been advertised to any peer.
VPN-Instance vpna, Router ID 10.2.2.9	The VPN instance is vpna, the Route ID is 10.2.2.9
Total Number of Routes	Number of routes in VPN instance vpna.

Display detailed information about the specified invalid BGP VPNv4 routes.

```
<HUAWEI> display bgp vpnv4 all routing-table 192.168.4.4
 BGP local router ID : 10.1.1.1
 Local AS number : 100

 Total routes of Route Distinguisher(10:10): 1
 BGP routing table entry information of 192.168.4.4/32:
 Label information (Received/Applied): 1025/NULL
 From: 10.1.1.2 (10.1.1.2)
 Route Duration: 00h12m26s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Relay Tunnel Out-Interface:
 Relay token: 0x0
 Original nexthop: 10.1.1.2
 Qos information : 0x0
 Ext-Community:RT <10 : 10>, RT <20 : 20>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255
 Not advertised to any peer yet

 VPN-Instance vrf1, Router ID 10.1.1.1:

 Total Number of Routes: 1
 BGP routing table entry information of 192.168.4.4/32:
 Label information (Received/Applied): 1025/NULL
 From: 10.1.1.2 (10.1.1.2)
 Route Duration: 00h12m26s
 Relay Tunnel Out-Interface:
 Relay token: 0x0
 Original nexthop: 10.1.1.2
 Qos information : 0x0
 Ext-Community:RT <10 : 10>, RT <20 : 20>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for tunnel
unreachable
 Not advertised to any peer yet
```


Table 7-183 Description of the display bgp vpnv4 routing-table command output

Item	Description
BGP local router ID	ID of the local BGP device. The format is the same as the IPv4 address.
Local AS number	Local AS number.
Total routes of Route Distinguisher(10:10)	Total number of BGP VPNv4 routes of the specified RD.
BGP routing table entry information of 192.168.4.4/32	The following information is about 192.168.4.4/32 routing entries.
Label information (Received/ Applied)	Label information (received or sent).
From	IP address of the device that sends the route. 10.1.1.2 is the IP address of the source interface of the peer with which the BGP connection is established, and 10.1.1.2 is the Router ID of the peer.
Route Duration	Duration of routes.
Relay IP Nexthop	Recursive next hop.
Relay IP Out-Interface	Recursive outbound interface.
Relay Tunnel Out-Interface	Tunnel recursed outbound interface.
Relay token	Recursive token value used for MPLS forwarding, which is a part of tunnel ID and is assigned by the system.
Original nexthop	Original next hop.
Qos information	QoS information. • 0x20000000: indicates that the **apply behavior** command has been run. • 0x40000001–0x40000FFF: indicates that the **apply qos-local-id** *qos-local-id* command has been run and the *qos-local-id* varies from 1 to 4095. • 0x80000001–0x80000007: indicates that the **apply ip-precedence** *precedence* command has been run and the *precedence* varies from 1 to 7. • 0x0: indicates that the preceding QoS configurations are not performed.
Ext-Community	Extended community attribute.
AS-path Nil	AS_Path attribute, with Nil indicating that the attribute value is null.

Item	Description
origin incomplete	Well-known mandatory property. This property defines the origin of a path and records how a route turns to a BGP route. The property has the following three values: • IGP: The priority of this value is the highest. The origin property of the routes that are added to the BGP routing table by using the **network (BGP)** command is IGP. • EGP: The priority of this value is second to that of IGP. The origin property of the routes imported from EGP is EGP. • Incomplete: The priority of this value is the lowest. The value indicates the origin of a route is unknown. The origin property of the routes that are added to the BGP routing table by using the **import-route (BGP)** command is Incomplete.
MED	Multi-Exit discriminator of route.
localpref	Local priority.
pref-val	Value preferred by the protocol.
valid	The BGP route is a valid route.
internal	The BGP route is an internal route.
best	The BGP route is an optimal route.
select	The BGP route is a preferred route.
pre 255	The priority of the BGP route is 255.

Item	Description
invalid for tunnel unreachable	Reason why a route is invalid: • invalid for route-policy not pass: The route does not match the route-policy. • invalid for supernet route: The route is a supernet route. • invalid for IP unreachable: The route fails to recurse to another route. • invalid for supernet route not advertise: No supernet routes are advertised. • invalid for supernet label route not advertise: No supernet labeled routes are advertised. • invalid for next-hop unreachable: The next-hop IP address is unreachable. • invalid for tunnel unreachable: The route fails to recurse to a tunnel.
Not advertised to any peer yet	The BGP route has not been advertised to any peer yet.
VPN-Instance vrf1, Router ID 10.1.1.1	The local VPN instance is vrf1, and its router ID is 10.1.1.1.
Total Number of Routes	Total number of BGP VPNv4 routes that match 192.168.4.4/32 in VPN instance vrf1.

7.8.74 display bgp vpnv4 routing-table statistics

Function

The **display bgp vpnv4 routing-table statistics** command displays statistics about BGP VPNv4 routes.

Format

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table statistics** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table statistics**

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table statistics** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **cidr** | **different-origin-as**]

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table statistics regular-expression** *as-regular-expression* (supported only by the S5731-

H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table statistics regular-expression** *as-regular-expression*

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table statistics community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table statistics community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp vpnv4 { **all** | **route-distinguisher** *route-distinguisher* } **routing-table statistics community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* } (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table statistics community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table statistics dampened**

display bgp vpnv4 all routing-table peer *ipv4-address* { **advertised-routes** | **received-routes** [**active**] } **statistics** (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

display bgp vpnv4 vpn-instance *vpn-instance-name* **routing-table peer** *ipv4-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

Parameters

Parameter	Description	Value
all	Displays all the statistics of BGP VPNv4 routes.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays statistics about the BGP routes with a specified RD.	The RD formats are divided into the following types: • 2-byte AS number:4-byte user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot be 0s at the same time. That is, an RD cannot be 0:0. • Integral 4-byte AS number: 2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • IPv4-address:2-byte user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.
vpn-instance *vpn-instance-name*	Displays statistics about the BGP routes of a specified VPN instance.	The value must be an existing VPN instance name.
as-path-filter	Displays the routes that match the specified filter.	-

Parameter	Description	Value
as-path-filter-number	Specifies the number of the matching AS-Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS-Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.
cidr	Displays CIDR statistics.	-
different-origin-as	Displays statistics about the routes that have the same destination address but different source AS numbers.	-
regular-expression *as-regular-expression*	Specifies the regular expression used to match the AS_Path information.	The value is a string of 1 to 80 characters.
community	Displays statistics about the routes carrying the specified BGP community attribute in the routing table.	-
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
aa:*nn*	Specifies a community attribute value.	Both *aa* and *nn* are integers ranging from 0 to 65535.
internet	Displays statistics about the matching routes that can be sent to any peer.	-
no-advertise	Displays statistics about the BGP routes carrying the No-Advertise community attribute.	-

Parameter	Description	Value
no-export	Displays statistics about the BGP routes carrying the No-Export community attribute.	-
no-export-subconfed	Displays statistics about the BGP routes carrying the No-Export-Subconfed community attribute.	-
whole-match	Indicates exact matching.	-
community-filter	Displays statistics about the routes that match a specified BGP community filter.	-
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
community-filter-name	Specifies the name of a community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits.
dampened	Displays the statistics of BGP dampened routes.	-
active	Displays statistics about the routes that have the same destination address but different source AS numbers.	-
peer *ipv4-address*	Displays statistics about routes of a specified peer.	It is in dotted decimal notation.
advertised-routes	Displays statistics about the routes advertised to a specified peer.	-

Parameter	Description	Value
received-routes	Displays statistics about the routes received from a specified peer.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp vpnv4 routing-table statistics** command to check statistics about BGP VPNv4 routes.

Example

Display statistics about BGP VPNv4 routes.

```
<HUAWEI> display bgp vpnv4 all routing-table statistics

Total number of routes from all PE: 2

VPN-Instance vpn1, Router ID 10.1.1.9:

Total Number of Routes: 2
```

Table 7-184 Description of the display bgp vpnv4 routing-table statistics command output

Item	Description
Total number of routes from all PE	Total number of VPNv4 routes
VPN-Instance vpn1	Indicating the name of VPN instance is vpn1.
Router ID 10.1.1.9	Indicating the router ID is 10.1.1.9.
Total Number of Routes	Total number of routes of the VPN instance.
Route Distinguisher	RD of the VPN instance IPv4 address family.

7.8.75 display bgp vpnv6 brief

Function

The **display bgp vpnv6 brief** command displays brief information about VPNv6 and VPN instances (IPv6 address family).

Product	Support
S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5731-S, S5731S-S, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **brief**

Parameters

Parameter	Description	Value
all	Displays information about all VPNv6 and VPN instances (IPv6 address family).	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPNv6 instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After the **display bgp vpnv6 brief** command is used to display information about VPNv6 and VPN instances (IPv6 address family), the VPN instances are displayed and arranged alphabetically by name.

Example

Display brief information about VPNv6 and VPN instances (IPv6 address family).

```
<HUAWEI> display bgp vpnv6 all brief
VPNV6:
  Rd Num              Peer Num           Route Num
    2                 1                  2

VPN-Instance(IPv6-family):
  VPN-Instance Name   Peer Num           Route Num
    vrf0              1                  2
    vrf1              0                  0
    vrf11             0                  0
    vrf12             0                  0
    vrf13             0                  0
    vrf14             0                  0
    vrf2              0                  20
    vrf3              0                  20
    vrf4              0                  24
    vrf5              0                  24
    vrf6              0                  0
```

Table 7-185 Description of the display bgp vpnv6 all brief command output

Item	Description
Rd Num	Number of RDs.
Peer Num	Number of peers.
Route Num	Number of routes.
VPN-Instance Name	Name of a VPN instance.

7.8.76 display bgp vpnv6 routing-table

Function

The **display bgp vpnv6 routing-table** command displays BGP VPNv6 routes.

Product	Support
S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5731-S, S5731S-S, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table** [**verbose** | *ipv6-address* [*prefix-length*]]

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table** [**verbose** | *ipv6-address* [*prefix-length*]]

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* | **route-distinguisher** *route-distinguisher* } **routing-table as-path-filter** { *as-path-filter-number* | *as-path-filter-name* }

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table as-path-filter** { *as-path-filter-number* | *as-path-filter-name* }

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table as-path-filter** { *as-path-filter-number* | *as-path-filter-name* }

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table different-origin-as**

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table different-origin-as**

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table regular-expression** *as-regular-expression*

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table regular-expression** *as-regular-expression*

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table peer** *ipv6-address* { **advertised-routes** [*dest-ipv6-address* [*prefix-length*]] | **received-routes** [**active**] }

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table peer** *ipv6-address* **received-routes** *dest-ipv6-address* [*prefix-length* [**original-attributes**]]

display bgp vpnv6 all routing-table peer *ipv4-address* **received-routes** *dest-ipv6-address* [*prefix-length*]

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table peer** *ipv6-address* **accepted-routes**

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table time-range** *start-time end-time*

display bgp vpnv6 all routing-table peer *ipv4-address* { **advertised-routes** [*dest-ipv6-address* [*prefix-length*]] | **received-routes** [**active**] }

Parameters

Parameter	Description	Value
all	Displays statistics about all BGP VPNv6 routes.	-
vpn-instance *vpn-instance-name*	Displays the BGP routes of a specified an IPv6 address family-enabled VPN instance on the local end.	The value must be an existing VPN instance name.

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays the BGP routes with the specified RD.	The RD formats are divided into the following types: • 2-byte AS number:4-byte user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot be 0s at the same time. That is, an RD cannot be 0:0. • Integral 4-byte AS number:2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • IPv4-address:2-byte user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.

Parameter	Description	Value
ipv6-address	Specifies the IPv6 address of a peer to be displayed.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
verbose	Displays detailed information about BGP VPNv6 routes.	-
prefix-length	Specifies the prefix length of an IPv6 address.	-
as-path-filter	Displays the routes that match the specified filter.	-
as-path-filter-number	Specifies the number of the matching AS-Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS-Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.
community	Displays the routes carrying the specified BGP community attribute in the routing table.	-
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
aa:nn	Specifies the community number. A maximum of 29 community numbers can be set.	-
internet	Displays the BGP routes carrying the Internet community attribute.	-
no-advertise	Displays the BGP routes carrying the No-Advertise community attribute.	-

Parameter	Description	Value
no-export	Displays the BGP routes carrying the No-Export community attribute.	-
no-export-subconfed	Displays the BGP routes carrying the No-Export-Subconfed community attribute.	-
whole-match	Indicates exact matching.	-
community-filter	Displays the routes that match a specified BGP community filter.	-
community-filter-name	Specifies the name of a community filter.	-
basic-community-filter-number	Specifies the number of a basic community filter.	-
advanced-community-filter-number	Specifies the number of an advanced community filter.	-
different-origin-as	Displays the routes that have the same destination address but different source AS numbers.	-
regular-expression *as-regular-expression*	Specifies the regular expression used to match the AS_Path information.	The value is a string of 1 to 80 characters.
peer *ipv6-address*	Displays the BGP routes of a specified peer.	-
ipv4-address	Specifies the IPv4 address of the peer.	The value is in dotted decimal notation.
advertised-routes	Displays the routes advertised to a specified peer.	-

Parameter	Description	Value
dest-ipv6–address	Specifies the destination IPv6 address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
received-routes	Displays the routes received from a specified peer.	-
active	Displays the active routes received from a specified peer.	-
original-attributes	Displays the original attributes of a BGP route from a specified BGP peer before the route is filtered by the local import policy. To display such attributes, the **peer keep-all-routes** command must have been run.	-
accepted-routes	Displays the routes that are received from a neighbor and accepted by a routing policy.	-
time-range *start-time end-time*	Displays information about BGP VPNv6 routes that have undergone status flapping during the specified period. For example, when *start-time* is set to 0d0h5m0s and *end-time* is set to 0d0h10m0s, information about all BGP VPNv6 routes whose lifetime ranges from 5 to 10 minutes is displayed.	In the values of *start-time* and *end-time*, d ranges from 0 to 10000, h ranges from 0 to 23, m ranges from 0 to 59, and s ranges from 0 to 59.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Information about specified routes can be displayed by specifying different parameters.

To view information about BGP VPNv6 route flapping during a specified period, you can run the **display bgp vpnv6** { **all** | **vpn-instance** *vpn-instance-name* } **routing-table time-range** *start-time end-time* command. When service traffic is abnormal during a period of time, you can run this command to check whether route flapping occurs. When the CPU usage is high during a period of time, you can run this command to check whether a large number of routes have undergone status flapping. This command allows you to find the flapping routes, which facilitates fault location.

Example

Display all BGP VPNv6 routes.

```
<HUAWEI> display bgp vpnv6 all routing-table


 BGP Local router ID is 10.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total number of routes from all PE: 3
 Route Distinguisher: 100:1


 *>i Network  : FC00:0:0:2001::                PrefixLen : 64
     NextHop  : FC00:0:0:2001::1               LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    : NULL
     Path/Ogn :  ?
 *>  Network  : FC00:0:0:2001::2              PrefixLen : 128
     NextHop  : ::                         LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    : NULL
     Path/Ogn :  ?
 *>  Network  : FE80::                        PrefixLen : 10
     NextHop  : ::                         LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    : NULL
     Path/Ogn :  ?

 VPN-Instance whm1, Router ID 10.4.4.4 :

 Total Number of Routes: 3
 *>  Network  : FC00:0:0:2001::              PrefixLen : 64
     NextHop  : ::                         LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    : NULL
     Path/Ogn :  ?
 *>  Network  : FC00:0:0:2001::2              PrefixLen : 128
     NextHop  : ::                         LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    : NULL
     Path/Ogn :  ?
 *>  Network  : FE80::                        PrefixLen : 10
     NextHop  : ::                         LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    : NULL
     Path/Ogn :  ?
```

Display information about BGP VPNv6 route flapping of a specified VPN instance during the specified period.

```
<HUAWEI> display bgp vpnv6 vpn-instance 1 routing-table time-range 0d5h0m0s 1d5h0m0s
BGP Local router ID is 10.1.1.1
 Status codes: * - valid, > - best, d - damped,
```



```
          h - history,  i - internal, s - suppressed, S - Stale
          Origin : i - IGP, e - EGP, ? - incomplete

Route Distinguisher: 300:1

*>  Network  : FC00:0:0:1991::                PrefixLen : 64
    NextHop  : ::                          Duration  : 16h32m17s
    Peer     : ::
    Path/Ogn : ?
*>  Network  : FC00:0:0:2004::                PrefixLen : 64
    NextHop  : ::                          Duration  : 16h34m02s
    Peer     : ::
    Path/Ogn : ?

Route Distinguisher: 10011:1

*>i Network  : FC00:0:0:1998::                 PrefixLen : 32
    NextHop  : FC00:0:0:1::9                 Duration  : 16h38m16s
    Peer     : 10.1.1.9
    Path/Ogn : 65410  ?
*>i Network  : FC00:0:0:1998::                 PrefixLen : 64
    NextHop  : FC00:0:0:1::9                 Duration  : 16h37m01s
    Peer     : 10.1.1.9
    Path/Ogn : ?
*>i Network  : FC00:0:0:2001::                 PrefixLen : 64
    NextHop  : FC00:0:0:1::9                 Duration  : 16h47m31s
    Peer     : 10.1.1.9
    Path/Ogn : ?
*>i Network  : FC00:0:0:3001::                 PrefixLen : 64
    NextHop  : FC00:0:0:1::9                 Duration  : 16h45m40s
    Peer     : 10.1.1.9
    Path/Ogn : 65410  ?

VPN-Instance vpna, Router ID 10.4.4.4 :
*>  Network  : FC00:0:0:1991::                PrefixLen : 64
    NextHop  : ::                          Duration  : 16h32m17s
    Peer     : ::
    Path/Ogn : ?
*>i Network  : FC00:0:0:1998::                 PrefixLen : 32
    NextHop  : FC00:0:0:1::9                 Duration  : 16h38m16s
    Peer     : 10.1.1.9
    Path/Ogn : 65410  ?
*>i Network  : FC00:0:0:1998::                 PrefixLen : 64
    NextHop  : FC00:0:0:1::9                 Duration  : 16h37m01s
    Peer     : 10.1.1.9
    Path/Ogn : ?
*>i Network  : FC00:0:0:2001::                 PrefixLen : 64
    NextHop  : FC00:0:0:1::9                 Duration  : 16h47m31s
    Peer     : 10.1.1.9
    Path/Ogn : ?
*>  Network  : FC00:0:0:2004::                PrefixLen : 64
    NextHop  : ::                          Duration  : 16h34m02s
    Peer     : ::
    Path/Ogn : ?
*>i Network  : FC00:0:0:3001::                 PrefixLen : 64
    NextHop  : FC00:0:0:1::9                 Duration  : 16h45m40s
    Peer     : 10.1.1.9
    Path/Ogn : 65410  ?
```

Table 7-186 Description of the display bgp vpnv6 all routing-table command output

Item	Description
BGP Local router ID	ID of the local BGP router. The ID is in the same format as an IPv4 address.

Item	Description
Total number of routes from all PE	Total number of BGP VPNv6 routes received by the switch from its peer PEs.
Network	Destination network or host address of the route.
PrefixLen	Prefix length of the destination network or host address of the route.
NextHop	IPv6 address of the next hop.
LocPrf	Local preference of the BGP route. The default value is 100.
MED	MED of the route. The default value is 0.
PrefVal	Preferred value of the route.
Label	Label carried by the data packet destined for the destination network or host address of the route.
Duration	Route duration.
Peer	IP addresses of the peer.
Path/Ogn	AS_Path number and Origin attribute of the route.

Display the routes of an IPv6 address family-enabled VPN instance named vpn1 on the local device.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table

 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
 *>i Network  : FC00:0:0:1::                 PrefixLen : 64
     NextHop  : FC00:0:0:3::1                LocPrf    :
     MED      : 0                          PrefVal   : 0
     Label    :
     Path/Ogn : 65410  ?
 *>i Network  : FC00:0:0:1::                 PrefixLen : 64
     NextHop  : FC00:0:0:2::1                LocPrf    : 100
     MED      : 0                          PrefVal   : 0
     Label    : 1037/NULL
     Path/Ogn : ?
```

Display the BGP routes with a specified destination address of an IPv6 address family-enabled VPN instance.

```
<HUAWEI> display bgp vpnv6 vpn-instance vrf1 routing-table fc00:0:0:1::

 BGP local router ID : 10.1.1.1
 Local AS number : 100
```



```
 Paths:   2 available, 1 best, 1 select
 BGP routing table entry information of FC00:0:0:1::/64:
 Imported route.
 From: :: (0.0.0.0)
 Route Duration: 1d03h46m24s
 Direct Out-interface: Vlanif100
 Original nexthop: ::
 AS-path Nil, origin incomplete, MED 0, pref-val 0, valid, local, best, select,
pre 0
 Advertised to such 1 peers:
    FC00:0:0:2::2
 BGP routing table entry information of 2001::/64:
 From: FC00:0:0:1::1 (10.10.10.10)
 Route Duration: 02h39m43s
 Direct Out-interface: Vlanif100
 Original nexthop: FC00:0:0:1::1
 AS-path 65410, origin incomplete, MED 0, pref-val 0, external, pre 255
 Not advertised to any peer yet
```

Display BGP VPNv6 routes of the VPN instance named **vpn1** whose AS_Path attribute contains 65420.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table as-path-filter 1

 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 VPN-Instance vpn1, Router ID 10.4.4.4 :

 Total Number of Routes: 1
     Network  : FC00:0:0:1::2001::                PrefixLen : 64
     NextHop  : FC00:0:0:1::2001::1               LocPrf    :
     MED      : 0                           PrefVal   : 0
     Label    :
     Path/Ogn : 65420  ?
```

Display BGP VPNv6 routes that the local switch advertises to the peer at 10.3.3.3.

```
<HUAWEI> display bgp vpnv6 all routing-table peer 10.3.3.3 advertised-routes

 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1

 Route Distinguisher: 100:1


 *>  Network  : FC00:0:0:1::                  PrefixLen : 64
     NextHop  : FC00:0:0:2::2                 LocPrf    :
     MED      : 0                           PrefVal   : 0
     Label    : NULL
     Path/Ogn : 65410 ?
```

Display BGP VPNv6 routes that the local switch receives from the peer at 10.3.3.3.

```
<HUAWEI> display bgp vpnv6 all routing-table peer 10.3.3.3 received-routes


 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
```



```
            h - history,  i - internal, s - suppressed, S - Stale
            Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1

 Route Distinguisher: 100:2


 *>i Network  : FC00:0:0:1::2002::                PrefixLen : 64
    NextHop  : FC00:0:0:1::2001::1                LocPrf    : 100
    MED      : 0                              PrefVal   : 0
    Label    : 1037
    Path/Ogn : ?
```

Display all BGP VPNv6 routes that match the BGP community filter 1.

```
<HUAWEI> display bgp vpnv6 all routing-table community-filter 1 whole-match

 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
            h - history,  i - internal, s - suppressed, S - Stale
            Origin : i - IGP, e - EGP, ? - incomplete


 Total number of routes from all PE: 2

 Route Distinguisher: 100:1


 *>  Network  : FC00:0:0:2001::                  PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                LocPrf    :
    MED      : 0                              PrefVal   : 0
    Label    : NULL

 Route Distinguisher: 100:2


 *>i Network  : FC00:0:0:2002::                   PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                LocPrf    : 100
    MED      : 0                              PrefVal   : 0
    Label    : 1037

 VPN-Instance vpn1, Router ID 10.4.4.4, Router ID 10.4.4.4 :

 Total Number of Routes: 2
    Network  : FC00:0:0:2001::                  PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                LocPrf    :
    MED      : 0                              PrefVal   : 0
    Label    :
 *>i Network  : FC00:0:0:2002::                   PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                LocPrf    : 100
    MED      : 0                              PrefVal   : 0
    Label    : 1037/NULL
```

Display BGP4+ routes of the VPN instance named **vpn1** and matching the BGP community filter 1.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table community-filter 1 whole-match

 BGP Local router ID is 10.1.1.9
 Status codes: * - valid, > - best, d - damped,
            h - history,  i - internal, s - suppressed, S - Stale
            Origin : i - IGP, e - EGP, ? - incomplete


 VPN-Instance vpn1 :

 Total Number of Routes: 2
    Network  : FC00:0:0:2001::                  PrefixLen : 64
```



```
    NextHop  : FC00:0:0:2001::1                 LocPrf    :
    MED      : 0                              PrefVal   : 0
    Label    :
*>i Network  : FC00:0:0:2002::                  PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                 LocPrf    : 100
    MED      : 0                              PrefVal   : 0
    Label    : 1037/NULL
```

Display all BGP4+ routes of the VPN instance named **vpn1** and matching the AS regular expression.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table regular-expression ^65420


BGP Local router ID is 10.1.1.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


VPN-Instance vpn1, Router ID 10.4.4.4, Router ID 10.4.4.4 :
    Network  : FC00:0:0:2001::                  PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                 LocPrf    :
    MED      : 0                              PrefVal   : 0
    Label    :
    Path/Ogn : 65420  ?
```

Display all BGP4+ routes of the VPN instance named **vpn1** that are received from the peer at FC00:0:0:2001::1.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table peer fc00:0:0:2001::1 received-routes

BGP Local router ID is 10.1.1.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


Total Number of Routes: 1
    Network  : FC00:0:0:2001::                  PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                 LocPrf    :
    MED      : 0                              PrefVal   : 0
    Label    :
    Path/Ogn : 65410  ?
```

Display BGP4+ routes sent to the peer at FC00:0:0:2001::1.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table peer fc00:0:0:2001::1 advertised-routes

BGP Local router ID is 10.1.1.9
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


Total Number of Routes: 1
*>i Network  : FC00:0:0:2002::                  PrefixLen : 64
    NextHop  : FC00:0:0:2001::1                 LocPrf    : 100
    MED      : 0                              PrefVal   : 0
    Label    : 1037/NULL
    Path/Ogn : ?
```

Display detailed information about the specified invalid BGP VPNv6 routes.

```
<HUAWEI> display bgp vpnv6 vpn-instance vrf1 routing-table fc00:0:0:2001::5
BGP local router ID : 10.1.1.1
Local AS number : 100

VPN-Instance vrf1, Router ID 10.1.1.1:
Paths:   1 available, 0 best, 0 select
```



```
BGP routing table entry information of fc00:0:0:2001::5/128:
Label information (Received/Applied): 1027/NULL
From: 10.1.1.2 (10.1.1.2)
Route Duration: 00h01m22s
Relay Tunnel Out-Interface:
Relay token: 0x0
Original nexthop: ::FFFF:10.1.1.2
Ext-Community:RT <100 : 100>
AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for tunnel
unreachable
Not advertised to any peer yet
```

Table 7-187 Description of the display bgp vpnv6 routing-table command output

Item	Description
BGP local router ID	Router ID of the local BGP device. The format is the same as the IPv4 address.
Local AS number	Local AS number.
VPN-Instance vrf1, Router ID 10.1.1.1	The local VPN instance is vrf1, and its router ID is 10.1.1.1.
Paths	Information about paths of BGP routes
BGP routing table entry information of fc00:0:0:2001::5/128	The following information is about fc00:0:0:2001::5/128 routing entries.
Label information (Received/ Applied)	Label information (received or sent).
From	IP address of the router that sends the route. 10.1.1.2 is the source interface IP address of the peer with which the BGP connection is established, and 10.1.1.2 is the router ID of the peer.
Route Duration	Duration of routes.
Relay Tunnel Out-Interface	Tunnel recursed outbound interface.
Relay token	Recursive token value used for MPLS forwarding, which is a part of tunnel ID and is assigned by the system.
Original nexthop	Original next hop.
Ext-Community	Extended community attribute.
AS-path Nil	AS_Path attribute, with Nil indicating that the attribute value is null.

Item	Description
origin	Well-known mandatory property. This property defines the origin of a path and records how a route turns to a BGP route. The property has the following three values: ● IGP: The priority of this value is the highest. The origin property of the routes that are added to the BGP routing table by using the **network (BGP)** command is IGP. ● EGP: The priority of this value is second to that of IGP. The origin property of the routes imported from EGP is EGP. ● Incomplete: The priority of this value is the lowest. The value indicates the origin of a route is unknown. The origin property of the routes that are added to the BGP routing table by using the **import-route (BGP)** command is Incomplete.
MED	Multi-Exit discriminator of route.
localpref	Local priority.
pref-val	Preferred value of the protocol.
internal	The BGP route is an internal route.
pre 255	The priority of the BGP route is 255.
invalid for tunnel unreachable	Reason why a route is invalid: ● invalid for route-policy not pass: The route does not match the route-policy. ● invalid for supernet route: The route is a supernet route. ● invalid for IP unreachable: The route fails to recurse to another route. ● invalid for supernet route not advertise: No supernet routes are advertised. ● invalid for supernet label route not advertise: No supernet labeled routes are advertised. ● invalid for next-hop unreachable: The next-hop IP address is unreachable. ● invalid for tunnel unreachable: The route fails to recurse to a tunnel.

Item	Description
Not advertised to any peer yet	The BGP route has not been advertised to any peer yet.

7.8.77 display bgp vpnv6 routing-table statistics

Function

The **display bgp vpnv6 routing-table statistics** command displays statistics about BGP VPNv6 routes.

Product	Support
S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5731-S, S5731S-S, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table statistics** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **different-origin-as**]

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table statistics** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | **different-origin-as**]

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table statistics regular-expression** *as-regular-expression*

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table statistics regular-expression** *as-regular-expression*

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table statistics community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name*} **routing-table statistics community** [*community-number* | *aa:nn*] &<1-29> [**internet** | **no-advertise** | **no-export** | **no-export-subconfed**] * [**whole-match**]

display bgp vpnv6 route-distinguisher *route-distinguisher* **routing-table statistics community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp vpnv6 { **all** | **vpn-instance** *vpn-instance-name* } **routing-table statistics community-filter** { { *community-filter-name* | *basic-community-filter-number* } [**whole-match**] | *advanced-community-filter-number* }

display bgp vpnv6 all routing-table peer *ipv4-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

display bgp vpnv6 vpn-instance *vpn-instance-name* **routing-table peer** *ipv6-address* { **advertised-routes** | **received-routes** [**active**] } **statistics**

Parameters

Parameter	Description	Value
all	Displays statistics about all BGP VPNv6 routes.	-

Parameter	Description	Value
route-distinguisher *route-distinguisher*	Displays BGP routes with the specified RD.	The RD formats are divided into the following types: • 2-byte AS number:4-byte user-defined number, for example, 101:3. An AS number ranges from 0 to 65535. A user-defined number ranges from 0 to 4294967295. The AS number and the user-defined number cannot be 0s at the same time. That is, an RD cannot be 0:0. • Integral 4-byte AS number:2-byte user-defined number, for example, 65537:3. An AS number ranges from 65536 to 4294967295. A user-defined number ranges from 0 to 65535. • 4-byte AS number in dotted notation:2-byte user-defined number, for example, 0.0:3 or 0.1:0. A 4-byte AS number in dotted notation is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 and from 0 to 65535, respectively. A user-defined number ranges from 0 to 65535. The AS number and user-defined number cannot be both 0s. That is, an RD cannot be 0.0:0. • IPv4-address:2-byte user-defined number, for example, 192.168.122.15:1. An IP address ranges from 0.0.0.0 to 255.255.255.255. A user-defined number ranges from 0 to 65535.
vpn-instance *vpn-instance-name*	Displays statistics about the BGP routes of a specified VPN instance.	The value must be an existing VPN instance name.
as-path-filter	Displays the routes that match the specified filter.	-

Parameter	Description	Value
as-path-filter-number	Specifies the number of the matching AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter-name	Specifies the name of the matching AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.
different-origin-as	Displays statistics about the routes that have the same destination address but different source AS numbers.	-
regular-expression *as-regular-expression*	Specifies the regular expression used to match the AS_Path information.	The value is a string of 1 to 80 case-sensitive characters. It cannot contain spaces.
community	Displays statistics about the routes carrying the specified BGP community attribute in the routing table.	-
community-number	Specifies the community number.	The value is an integer that ranges from 0 to 4294967295.
aa:nn	Specifies the community number.	Both *aa* and *nn* are integers ranging from 0 to 65535.
internet	Displays statistics about the BGP routes carrying the Internet community attribute.	-
no-advertise	Displays statistics about the BGP routes carrying the No-Advertise community attribute.	-

Parameter	Description	Value
no-export	Displays statistics about the BGP routes carrying the No-Export community attribute.	-
no-export-subconfed	Displays statistics about the BGP routes carrying the No-Export-Subconfed community attribute.	-
whole-match	Indicates exact matching.	-
community-filter	Displays statistics about the routes that match a specified BGP community filter.	-
community-filter-name	Specifies the name of a community filter.	The name is a string of 1 to 51 case-sensitive characters without spaces.
basic-community-filter-number	Specifies the number of a basic community filter.	The value is an integer that ranges from 1 to 99.
advanced-community-filter-number	Specifies the number of an advanced community filter.	The value is an integer that ranges from 100 to 199.
peer *ipv6-address*	Displays statistics about the BGP routes of a specified peer.	-
advertised-routes	Displays statistics about the routes advertised to a specified peer.	-
received-routes	Displays statistics about the routes received from a specified peer.	-
active	Specifies the number of active routes.	-
peer *ipv4-address*	Displays the routing information for the specified BGP peer.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display bgp vpnv6 routing-table statistics** command to check statistics about BGP VPNv6 routes.

Example

Display statistics about BGP VPNv6 routes on the local device.

```
<HUAWEI> display bgp vpnv6 all routing-table statistics

 Total number of routes from all PE: 2

 VPN-Instance vpn1, Router ID 10.4.4.4 :

 Total Number of Routes: 2
```

Display statistics about the BGP4+ routes of an IPv6 address family-enabled VPN instance named **vpn1** on the local device.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table statistics

 Total Number of Routes: 5
```

Display statistics about the BGP VPNv6 routes received from peer 10.1.1.1.

```
<HUAWEI> display bgp vpnv6 all routing-table peer 10.1.1.1 received-routes statistics
  Received routes total: 1
```

Display statistics about the VPNv6 routes that have the Internet community attribute and specified RD on the local device.

```
<HUAWEI> display bgp vpnv6 route-distinguisher 100:1 routing-table statistics community internet

 Total number of routes from all PE: 1

 VPN-Instance vpn1, Router ID 10.4.4.4 :

 Total Number of Routes: 2
```

Display statistics about the BGP VPNv6 routes that match a specified community filter on the local device.

```
<HUAWEI> display bgp vpnv6 all routing-table statistics community-filter 1

 Total number of routes from all PE: 1

 VPN-Instance vpn1, Router ID 10.4.4.4 :

 Total Number of Routes: 2
```

Display statistics about the BGP VPNv6 routes that match a specified AS_Path regular expression.

```
<HUAWEI> display bgp vpnv6 all routing-table statistics regular-expression 65420*

 Total number of routes from all PE: 1
```



```
 VPN-Instance vpn1, Router ID 10.4.4.4 :

 Total Number of Routes: 1
```

Display statistics of BGP routes sent by the local device to peer 2001:db8:1::1 of the IPv6 VPN instance named **vpn1**.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table peer 2001:db8:1::1 received-routes statistics
  Received routes total: 2
```

Display statistics about the IPv6 routes sent by the local device to peer 2001:db8:1::1 in a VPN instance named **vpn1**.

```
<HUAWEI> display bgp vpnv6 vpn-instance vpn1 routing-table peer 2001:db8:1::1 advertised-routes statistics
  Advertised routes total: 2

  Default originated : 0
```

7.8.78 display default-parameter bgp

Function

The **display default-parameter bgp** command displays default configurations in BGP initialization.

Format

display default-parameter bgp

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display default-parameter bgp** command to check default configurations in BGP initialization.

Example

Display default configurations in BGP initialization.

```
<HUAWEI> display default-parameter bgp
BGP version               : 4
EBGP preference           : 255
IBGP preference           : 255
Local preference          : 255
BGP connect-retry         : 32s
```



```
BGP holdtime               : 180s
BGP keepAlive              : 60s
EBGP route-update-interval  : 30s
IBGP route-update-interval  : 15s
Default local-preference    : 100
Default MED                 : 0
Default TRACKING TIMER      : 9
IPv4-family unicast         : enable
EBGP-interface-sensitive    : enable
Reflect between-clients     : enable
Check-first-as              : enable
Synchronization             : disable
Nexthop-resolved rules      :
    IPv4-family             : unicast(ip)
                              label-route(ip)
                              multicast(ip)
                              vpn-instance(tunnel)
                              vpnv4(ip)
    IPv6-family             : unicast(ip)
                              vpn-instance(tunnel)
                              vpnv6(ip)
                              6PE(tunnel)
    L2VPN-AD-family         : ip
 Routing-table limit max-alarm upper limit : 100
 Routing-table limit max-alarm lower limit : 95
 Routing-table limit threshold-alarm upper limit : 80
 Routing-table limit threshold-alarm lower limit : 70
```

Table 7-188 Description of the **display default-parameter bgp** command output

Item	Description
BGP version	BGP version number.
EBGP preference	EBGP route preference.
IBGP preference	IBGP route preference.
Local preference	Local route preference.
BGP connect-retry	BGP ConnectRetry interval.
BGP holdtime	BGP holdtime interval.
BGP keepAlive	BGP keepalive interval.
EBGP route-update-interval	Minimum interval for sending EBGP Update messages.
IBGP route-update-interval	Minimum interval for sending IBGP Update messages.
Default local-preference	Local preference of a BGP route.
Default MED	MED of a BGP route.
Default TRACKING TIMER	The interval between peer unreachability discovery and connection interruption.
IPv4-family unicast	BGP-IPv4 unicast address family view.

Item	Description
EBGP-interface-sensitive	The BGP session between the directly connected peer and an interface is deleted immediately when the interface becomes Down.
Reflect between-clients	Route reflection between clients.
Check-first-as	The first AS number in the AS_Path list that is carried in the Update message sent by the EBGP peer is checked.
Synchronization	Synchronization between IBGP and IGP.
Nexthop-resolved rules	Default recursion mode of preferred routes.
Routing-table limit max-alarm upper limit	Upper alarm threshold for the number of BGP routes.
Routing-table limit max-alarm lower limit	Lower alarm threshold for the number of BGP routes.
Routing-table limit threshold-alarm upper limit	Default upper alarm threshold for the number of BGP routes.
Routing-table limit threshold-alarm lower limit	Default lower alarm threshold for the number of BGP routes.

7.8.79 display mbgp routing-table

Function

The **display mbgp routing-table** command displays MBGP routes.

Format

display mbgp routing-table [*ipv4-address* [*mask* | *mask-length*]] [**verbose**]

Parameters

Parameter	Description	Value
ipv4-address	Specifies the destination IP address in dotted decimal notation.	-
mask | *mask-length*	Specifies mask in dotted decimal notation or mask-length.	-

Parameter	Description	Value
verbose	Displays detailed information about active and inactive routes. If the parameter **verbose** is not specified, detailed information about active routes is displayed.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can specify different parameters to view the specific routing information.

Example

Display information about all MBGP routes.

```
<HUAWEI> display mbgp routing-table
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
Routing Tables: MBGP
         Destinations : 1        Routes : 1

Destination/Mask   Proto   Pre  Cost      Flags NextHop         Interface

 10.5.5.1/32       MBGP    255  0           R   10.1.1.1        Vlanif10
```

Table 7-189 Description of the display mbgp routing-table command output

Item	Description
Route Flags	Flag of a route: • R: indicates that the route is a recursive route. • D: indicates that the route is delivered to the FIB table. • T: indicates a route whose next hop belongs to a VPN instance.
Routing Tables: MBGP	Indicates an MBGP routing table.
Destinations	Indicates the total number of destination networks or hosts.
Routes	Indicates the total number of routes.
Destination/Mask	Indicates the address and mask length of the destination network or host.

Item	Description
Proto	Indicates the protocol through which routes are learned.
Pre	Indicates the preference.
Cost	Indicates the route cost.
Flags	Indicates the route flag, that is, Route Flags in the header of the routing table.
NextHop	Indicates the next hop.
Interface	Indicates the outbound interface through which the next hop is reachable.

Display the detailed information of the specified routes.

```
<HUAWEI> display mbgp routing-table 5.5.5.1 verbose
Routing Table : MBGP
Summary Count : 1

 Destination: 10.5.5.1/32
    Protocol: MBGP              Process ID: 0
  Preference: 255                   Cost: 0
     NextHop: 10.1.1.1          Neighbour: 0.0.0.0
       State: Active Adv GotQ         Age: 00h43m25s
         Tag: 0                 Priority: 0
       Label: NULL                QoSInfo: 0x0
 RelayNextHop: 0.0.0.0            Interface: Vlanif 10
    TunnelID: 0x0                 Flags: R
```

Table 7-190 Description of the display mbgp routing-table verbose command output

Item	Description
Destination	Indicates the address and mask length of the destination network or host.
Protocol	Indicates the routing protocol.
Process ID	Indicates the process ID of the routing protocol.
Preference	Indicates the preference of the route.
Cost	Indicates the route cost.
NextHop	Indicates the next hop.
Neighbour	Indicates the neighbor.

Item	Description
State	Indicates the status of routes: • Active: indicates active routes. • Invalid: indicates invalid routes. • Inactive: indicates inactive routes. • NoAdv: indicates the routes that cannot be advertised. • Adv: indicates the routes that can be advertised. • Del: indicates the routes to be deleted. • GotQ: indicates the route that finds the next hop and outbound interface or the route that finds the tunnel. • WaitQ: indicates the route that does not find the next hop or outbound interface or the route that does not find the tunnel. • Stale: indicates the routes with the stale flag. The routes are used in GR.
Age	Indicates the lifetime of the route.
Tag	Indicates the administrative tag for routes.
Priority	Indicates the priority.
Label	Indicates the allocated MPLS label.
QoSInfo	Indicates QoS information.
RelayNextHop	Indicates the relay next hop.
Interface	Indicates the outbound interface through which the next hop is reachable.
Tunnel ID	Indicates the tunnel ID.
Flags	Indicates the route flag, that is, Route Flags in the header of the routing table.

7.8.80 display mbgp routing-table statistics

Function

The **display mbgp routing-table statistics** command displays the statistics of the MBGP routes.

Format

display mbgp routing-table statistics

Parameters

None.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Route statistics include:

- Total number of routes that are added or deleted by the protocol
- Number of active routes or inactive routes that are labeled for deletion but are not actually deleted

Example

Display statistics of the MBGP routing table.

```
<HUAWEI> display mbgp routing-table statistics
Proto     total      active      added       deleted      freed
          routes     routes      routes      routes       routes
MBGP      6          4           10          0            0
```

Table 7-191 Description of the display mbgp routing-table statistics command output

Item	Description
Proto	Routing protocol
total routes	Total number of routes in the routing table
active routes	Number of active routes in the routing table
added routes	Number of active and inactive routes added in the routing table
deleted routes	Number of routes to be deleted from the routing table
freed routes	Number of routes that are permanently deleted from the routing table

7.8.81 ebgp-interface-sensitive

Function

The **ebgp-interface-sensitive** command immediately resets a BGP session on an interface that is directly connected to an external peer when the interface goes Down.

The **undo ebgp-interface-sensitive** command disables the function.

By default, this function is enabled.

Format

ebgp-interface-sensitive

undo ebgp-interface-sensitive

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If the **ebgp-interface-sensitive** command is not configured, the system does not immediately select the sub-optimal route for packet transmission when an interface goes Down. Instead, the system waits for a period of time (defaulting to 180 seconds) before checking whether another interface can be used to send packets to the same destination address. This will interrupt services for a period of time. If the **ebgp-interface-sensitive** command is run, BGP can quickly detect EBGP link faults and use another interface to establish a BGP peer relationship with the remote peer.

When the interface used for a BGP connection alternates between Up and Down states, running the **undo ebgp-interface-sensitive** command can prevent the repeated reestablishment and deletion of the BGP session in the event of route flapping. This reduces the use of network bandwidth.

Precautions

If the interface used for a BGP connection alternates between Up and Down states, it would be better not to run the **ebgp-interface-sensitive** command to prevent route flapping.

Example

Enable the function that automatically resets a BGP session.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ebgp-interface-sensitive
```

7.8.82 ext-community-change enable

Function

The **ext-community-change enable** command enables a device to change extended community attributes based on a Route-Policy.

The **undo ext-community-change enable** command disables a device from changing extended community attributes based on a Route-Policy.

By default, extended community attributes cannot be changed based on a Route-Policy.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

ext-community-change enable

undo ext-community-change enable

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, BGP prevents a device from changing extended community attributes from a peer or peer group based on an import policy or changing extended community attributes to be advertised to a peer or peer group based on an export policy. To enable a device to change extended community attributes from a peer or peer group based on an import policy and change extended community attributes to be advertised to a peer or peer group based on an export policy before advertising them, run the **ext-community-change enable** command.

Precautions

- The **ext-community-change enable** and **peer route-policy import** commands must both be run so that the device can change extended community attributes from a peer or peer group based on an import policy.
- The **ext-community-change enable** command must be run with either of the following commands so that the extended community attributes changed based on an export policy can be advertised to a peer or peer group.
  - **peer advertise-ext-community**
  - **peer route-policy export**

Example

Enable the device to change extended community attributes of the BGP routes received from peer 10.1.1.1 based on a Route-Policy named **policy1**.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy1 permit node 10
[HUAWEI-route-policy] if-match as-path-filter 2
[HUAWEI-route-policy] apply extcommunity rt 10.1.1.1:1 additive
[HUAWEI-route-policy] quit
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] ext-community-change enable
[HUAWEI-bgp-af-ipv4] peer 10.1.1.1 route-policy policy1 import
```

7.8.83 filter-policy export (BGP)

Function

The **filter-policy export** command configures a device to filter the routes to be advertised. BGP advertises only the routes that pass filtering.

The **undo filter-policy export** command restores the default configuration.

By default, the routes to be advertised are not filtered.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } **export** [*protocol* [*process-id*]]

filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } **export** [*protocol* [*process-id*]]

undo filter-policy export [*protocol* [*process-id*]]

undo filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } **export** [*protocol* [*process-id*]]

undo filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } **export** [*protocol* [*process-id*]]

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IPv4 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
protocol	Specifies the name of a routing protocol.	The BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, and BGP-VPN instance IPv4 address family view support **direct**, **isis**, **ospf**, **rip**, **static**, and **unr**. The BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view support **direct**, **isis**, **ospfv3**, **ripng**, **static**, and **unr**.
process-id	Specifies the number of a process that needs to perform matching. If *protocol* is **direct**, **static**, or **unr**, no process ID is required.	The value is an integer that ranges from 1 to 65535.
acl6-number	Specifies the number of a basic ACL6.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
acl6-name *acl6-name*	Specifies the name of a named ACL6.	The value is a string of 1 to 64 case-sensitive characters without spaces. The name should start with a letter and can contain numbers, hyphens (-), or underscores (_).
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

NOTE

- *protocol* [*process-id*] is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, and BGP-VPN instance IPv6 address family view.
- **acl-name** *acl-name*, *acl-number*, and **ip-prefix** *ip-prefix-name* are valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, and BGP-VPNv4 address family view.
- **acl6-name** *acl6-name*, *acl6-number*, and **ipv6-prefix** *ipv6-prefix-name* are valid only in the BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, and BGP-VPNv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-MDT address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

The **filter-policy export** command affects the routes advertised by BGP. After the command is run, BGP filters the routes that are imported by using the **import-route** command. Only the routes that pass the filtering can be added to the local BGP routing table and advertised by BGP.

If *protocol* is specified, only the routes imported from the specified protocol will be filtered. If *protocol* is not specified, the routes imported from all protocols will be filtered.

Precautions

Creating an IP prefix list before it is referenced is recommended. By default, nonexistent IP prefix lists cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list is referenced using the current command, all routes are advertised to the specified peer.

Example

Use ACL 2000 to filter all the routes to be advertised by BGP.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] filter-policy 2000 export
```

7.8.84 filter-policy import (BGP)

Function

The **filter-policy import** command configures a device to filter received routes.

The **undo filter-policy import** command restores the default configuration.

By default, received routes are not filtered.

Format

filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } **import**

filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } **import**

undo filter-policy import

undo filter-policy { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* } **import**

undo filter-policy { *acl6-number* | **acl6-name** *acl6-name* | **ipv6-prefix** *ipv6-prefix-name* } **import**

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.

Parameter	Description	Value
acl-name *acl-name*	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IPv4 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
acl6-number	Specifies the number of a basic ACL6.	-
acl6-name *acl6-name*	Specifies the name of a named ACL6.	The value is a string of 1 to 64 case-sensitive characters without spaces. The name should start with a letter and can contain numbers, hyphens (-), or underscores (_).
ipv6-prefix *ipv6-prefix-name*	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

NOTE

- **acl-name** *acl-name*, *acl-number*, and **ip-prefix** *ip-prefix-name* are valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, and BGP-VPNv4 address family view.
- **acl6-name** *acl6-name*, *acl6-number*, and **ipv6-prefix** *ipv6-prefix-name* are valid only in the BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, and BGP-VPNv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **filter-policy import** command is used to filter the routes received by BGP to determine whether to add them to the BGP routing table.

When the **rule** command is run to configure rules for an ACL, only the source address range specified by **source** and the time period specified by **time-range** are valid as the rules.

When the **rule (basic ACL6 view)** or **rule (advanced ACL6 view)** command is run to configure rules for an ACL6, only the source address range specified by **source** and the time period specified by **time-range** are valid as the rules.

Precautions

Creating an IP prefix list before it is referenced is recommended. By default, nonexistent IP prefix lists cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list is referenced using the current command, all routes advertised by the specified peer are accepted.

Example

Use ACL 2000 to filter the routes received by BGP.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] filter-policy 2000 import
```

7.8.85 graceful-restart (BGP)

Function

The **graceful-restart** command enables GR for the BGP speaker.

The **undo graceful-restart** command restores the default configuration.

By default, GR is disabled.

Format

graceful-restart

undo graceful-restart

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **graceful-restart** command is configured to enable GR for the switch, to prevent traffic interruption, another GR-capable device can assist the switch to perform GR when BGP on the switch restarts due to an active/standby switchover, and the switch can assist other GR-capable devices to perform GR when BGP on other devices restarts.

Follow-up Procedure

After running the **graceful-restart** command, run the **graceful-restart timer wait-for-rib** command to set the time for waiting for the End-Of-RIB flag.

Precautions

Enabling or disabling GR may delete and reestablish all BGP sessions and instances.

If the **graceful-restart timer wait-for-rib** command has been configured, using the **undo graceful-restart** command will delete the **graceful-restart timer wait-for-rib** command configuration.

Example

Enable GR for the speaker in BGP process 100.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] graceful-restart
```

7.8.86 graceful-restart peer-reset

Function

The **graceful-restart peer-reset** command enables a device to reset a BGP connection in GR mode.

The **undo graceful-restart peer-reset** command restores the default configuration.

By default, a device is not enabled to reset a BGP connection in GR mode.

Format

graceful-restart peer-reset

undo graceful-restart peer-reset

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Currently, BGP does not support dynamic capability negotiation. Therefore, each time a new BGP capability is enabled on a BGP speaker, the BGP speaker tears down existing sessions with its peer and renegotiates BGP capabilities.

With the GR reset function configured, when you enable a new BGP capability on the BGP speaker, the BGP speaker enters the GR state, resets the BGP session, and renegotiates BGP capabilities with the peer. In the whole process, the BGP speaker re-establishes the existing sessions but does not delete the routing entries for the existing sessions, so that the existing services are not interrupted.

Prerequisites

GR has been enabled by running the **graceful-restart** command. If this prerequisite is not met, the system does not allow you to configure the **graceful-restart peer-reset** command.

Precautions

After you run the **undo graceful-restart** command to disable GR, the **graceful-restart peer-reset** command configuration will be deleted automatically.

Example

Enable the switch to reset a BGP connection in GR mode.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] graceful-restart
[HUAWEI-bgp] graceful-restart peer-reset
```

7.8.87 graceful-restart timer restart

Function

The **graceful-restart timer restart** command sets the maximum period from the time when the peer finds that the local peer restarts to the time when the BGP session is reestablished.

The **undo graceful-restart timer restart** command deletes the setting.

By default, the maximum period from the time when the peer finds that the local peer restarts to the time when the BGP session is reestablished is 150 seconds.

Format

graceful-restart timer restart *time*

undo graceful-restart timer restart

Parameters

Parameter	Description	Value
time	Specifies the maximum period from the time when the peer finds that the local peer restarts to the time when the BGP session is reestablished.	The value ranges from 3 to 4095, in seconds.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Modifying the maximum period for reestablishing the BGP session leads to the reestablishment of the BGP peer relationship.

NOTE

The **graceful-restart timer restart** command can be run only after the **graceful-restart** command is run.

Example

Set the maximum period for reestablishing the BGP session to 250 seconds.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] graceful-restart
[HUAWEI-bgp] graceful-restart timer restart 250
```

7.8.88 graceful-restart timer wait-for-rib

Function

The **graceful-restart timer wait-for-rib** command sets the length of time that the BGP restarter waits for the End-Of-RIB flag.

The **undo graceful-restart timer wait-for-rib** command deletes the configured length of time that the BGP restarter waits for the End-Of-RIB flag.

By default, the time that the BGP restarter waits for the End-Of-RIB flag is 600 seconds.

Format

graceful-restart timer wait-for-rib *time*

undo graceful-restart timer wait-for-rib

Parameters

Parameter	Description	Value
time	Specifies the length of time that the BGP restarter waits for the End-Of-RIB flag.	The value is an integer that ranges from 3 to 3000, in seconds.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the BGP session is set up or reestablished, the BGP restarter must receive the End-Of-RIB flag in the period set by using this command. If the BGP restarter does not receive the End-Of-RIB flag, ensure that the switch can exit from the GR process.

Prerequisites

The **graceful-restart timer wait-for-rib** command can be run only after the **graceful-restart** command is run.

Configuration Impact

The latest configuration overrides the previous one.

Precautions

If there are a large number of routes, set a larger value for the *time* parameter to ensure that all routes will be updated.

Example

Set the time that the BGP restarter waits for the End-Of-RIB flag to 100 seconds.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] graceful-restart timer wait-for-rib 100
```

7.8.89 group

Function

The **group** command creates a peer group.

The **undo group** command deletes a peer group.

By default, no peer group is created.

Format

group *group-name* [**external** | **internal**]

undo group *group-name*

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
external	Creates an EBGP peer group.	-
internal	Creates an IBGP peer group.	-

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A peer group is a group of peers with the same configurations. After a peer is added to a peer group, it inherits the configurations of this peer group. Peers in a peer group inherit the configurations of the peer group. When the configurations of the peer group are changed, the configurations of these peers are changed accordingly.

On a large-scale BGP network, there are a large number of peers and many of them have the same configurations. To configure these peers, you have to repeatedly use some commands. In such a case, configuring peer groups can simplify configurations. If the configurations for several peers are the same, these peers can be added to a created and configured peer group. The peers in the peer group then inherit the configurations of the peer group.

Precautions

If the **group** command is run multiple times, the latest configuration does not override the previous one.

If the type (IBGP or EBGP) of a peer group is not specified, an IBGP peer group is created by default.

The configuration of a peer takes precedence over that of the peer group to which the peer belongs.

If an attribute configuration of a BGP peer in a peer group differs from that of the peer group, you can disable the attribute configuration of the peer by using an **undo** command; then the peer inherits the attribute configuration of the peer group.

NOTICE

Deleting a peer group closes the connections on the peers that have no AS numbers in the peer group. Before deleting a peer group, you are recommended to delete these peers or configure AS numbers for these peers.

Example

Create an IBGP peer group named in.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group in internal
```

Create an EBGP peer group named ex, and set its AS number to 500.1.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group ex external
[HUAWEI-bgp] peer ex as-number 500.1
```

7.8.90 import-route (BGP)

Function

The **import-route** command configures BGP to import routes of other routing protocols and types.

The **undo import-route** command restores the default setting.

By default, BGP does not import routes.

Format

import-route *protocol* [*process-id*] [**med** *med* | **route-policy** *route-policy-name*] *

undo import-route *protocol* [*process-id*]

Parameters

Parameter	Description	Value
protocol	Specifies the routing protocol type and route type.	The BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, and BGP-VPN instance IPv4 address family view support **direct**, **isis**, **ospf**, **rip**, **static**, and **unr**. The BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view support **direct**, **isis**, **ospfv3**, **ripng**, **static**, and **unr**.
process-id	Specifies a process ID if BGP is configured to import routes. If *protocol* is **direct**, **static**, or **unr**, no process ID is required.	The value is an integer that ranges from 1 to 65535.
med *med*	Specifies the MED.	The value is an integer that ranges from 0 to 4294967295.
route-policy *route-policy-name*	Indicates that routes are filtered and route attributes are modified by using the Route-Policy specified by the parameter when these routes are imported from other protocols.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP can import routes by using the **import-route** command or **network** command:

- In import mode, BGP imports IGP routes, including RIP, OSPF, and IS-IS routes, into the BGP routing table based on protocol type. To ensure the validity of imported IGP routes, BGP can also import static routes and direct routes in import mode.

- In network mode, BGP imports the routes in the IP routing table one by one into the BGP routing table. The network mode is more accurate than the import mode.

Precautions

If the **default-route imported** command has not been used, BGP cannot import default routes when you run the **import-route** command to import routes from other protocols.

After the **import-route direct** command is executed, routes to the network segment where the IP address of the management interface belongs are also imported in the BGP routing table. Therefore, use this command with caution.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all routes of the specified protocol are imported to the BGP routing table.

Example

Import routes from RIP process 1.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] import-route rip 1
```

7.8.91 ipv4-family

Function

The **ipv4-family** command enables the IPv4 address family of BGP, and then displays the address family view.

The **undo ipv4-family** command deletes the configurations in the IPv4 address family.

By default, the BGP-IPv4 unicast address family view is displayed.

Format

ipv4-family unicast

ipv4-family multicast

ipv4-family vpnv4 [**unicast**] (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

ipv4-family vpn-instance *vpn-instance-name*

ipv4-family { **mdt** | **mvpn** }

undo ipv4-family vpnv4 [**unicast**](supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

undo ipv4-family { **multicast** | **vpn-instance** *vpn-instance-name* }

undo ipv4-family { **mdt** | **mvpn** }

NOTE

The **mdt, mvpn, multicast** parameter is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.

Parameters

Parameter	Description	Value
unicast	Displays the unicast address family view.	-
multicast	Displays the multicast address family view.	-
vpnv4	Displays the BGP-VPNv4 address family view.	-
vpn-instance *vpn-instance-name*	Associates a specified VPN instance with the IPv4 address family. You can enter the BGP-VPN instance IPv4 address family view by using the parameter.	The value must be an existing VPN instance name.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Before performing BGP configurations in an IPv4 address family, you need to run the **ipv4-family** command in the BGP view to enable the IPv4 address family, and then enter the address family view. By default, BGP uses the IPv4 unicast address family.

Precautions

To disable the IPv4 unicast address family from being the default BGP route, run the **undo default ipv4-unicast** command.

Example

Enter the BGP-IPv4 unicast address family view.

```
<HUAWEI> system-view
```



```
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4]
```

7.8.92 ipv6-family

Function

The **ipv6-family** command enters the IPv6 address family view of BGP.

The **undo ipv6-family** command quits the IPv6 address family view and deletes the configurations in the view.

Format

ipv6-family [**vpnv6**](supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

ipv6-family [**unicast** | **vpn-instance** *vpn-instance-name*]

undo ipv6-family [**vpnv6**] (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

undo ipv6-family [**unicast** | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
unicast	Displays the unicast address family view.	-
vpnv6	Displays the BGP-VPNv6 address family view.	-
vpn-instance *vpn-instance-name*	Associates a specified VPN instance with the IPv6 address family. You can enter the BGP-VPN instance IPv6 address family view by using the parameter. **NOTE** The parameter is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.	The value must be an existing VPN instance name.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If no parameter is specified, the IPv6 address family view is displayed by default. The **undo ipv6-family** command without any parameter is used to delete configurations in the BGP-IPv6 unicast address family view.

Precautions

The **undo ipv6-family** command without any parameters deletes all IPv6 unicast address family configurations.

Example

Enter the BGP-IPv6 unicast address family view.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv6-family
[HUAWEI-bgp-af-ipv6]
```

7.8.93 load-balancing as-path-ignore

Function

The **load-balancing as-path-ignore** command configures a router not to compare the AS_Path attributes of routes that are to be used for load balancing.

The **undo load-balancing as-path-ignore** command configures a router to compare the AS_Path attributes of routes that are to be used for load balancing.

By default, a router compares the AS-Path attributes of routes that are to be used for load balancing.

Format

load-balancing as-path-ignore

undo load-balancing as-path-ignore

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **load-balancing as-path-ignore** command is run, a router does not compare the AS_Path attributes of the routes (including the AS_Path length and content) that are to be used for load balancing. This command applies to the scenarios where EBGP and IBGP routes carry out load balancing. Exercise caution when using the command because the execution of this command will change the conditions of load balancing.

Precautions

The **load-balancing as-path-ignore** command and the **bestroute as-path-ignore** command are mutually exclusive. This means that if the **bestroute as-path-ignore** command is configured, the **load-balancing as-path-ignore** command cannot be configured.

Example

Configure a router not to compare the AS-Path attributes of the routes that are to be used for load balancing.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] load-balancing as-path-ignore
```

7.8.94 maximum load-balancing (BGP)

Function

The **maximum load-balancing** command configures the maximum number of equal-cost routes for load balancing.

The **undo maximum load-balancing** command restores the default value.

By default, the maximum number of equal-cost routes is 1, indicating that load balancing is not implemented.

Format

maximum load-balancing [**ebgp** | **ibgp**] *number* [**ecmp-nexthop-changed**]

undo maximum load-balancing [**ebgp** | **ibgp**]

Parameters

Parameter	Description	Value
ebgp	Indicates that only EBGP routes implement load balancing.	-

Parameter	Description	Value
ibgp	Indicates that only IBGP routes implement load balancing.	-
number	Specifies the maximum number of equal-cost routes in the BGP routing table.	BGP: The value is an integer that ranges from 1 to 8 on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S and S6720S-S. The value ranges from 1 to 16 on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S BGP4+: The value is an integer that ranges from 1 to 8.
ecmp-nexthop-changed	Configures a BGP device to change the next-hop addresses of only the routes that participate in load balancing to its address.	-

NOTE

ecmp-nexthop-changed is valid only in the BGP view, BGP IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After BGP load balancing is configured, BGP routes that meet the following conditions and have the same AS_Path attribute will become equal-cost routes and implement load balancing:

- PrefVal values are the same.
- Local_Pref attributes are the same.

- All BGP routes are aggregated or non-aggregated routes.
- Accumulated Interior Gateway Protocol metric (AIGP) values are the same.
- Lengths of the AS_path attributes are the same.
- Origin types (IGP, EGP, or Incomplete) are the same.
- Multi_Exit Discriminator (MED) values are the same.
- All BGP routes are EBGP or IBGP routes.
- The IGP metric values within an AS are the same.

Configuring BGP load balancing better utilizes network resources.

After the **maximum load-balancing ebgp** *number* command is run, only EBGP routes take part in load balancing. After the **maximum load-balancing ibgp** *number* command is run, only IBGP routes take part in load balancing. If [**ebgp** | **ibgp**] is not set, both EBGP routes and IBGP routes take part in load balancing, and the number of EBGP routes involved in load balancing is the same as that of IBGP routes involved in load balancing.

If you run the **maximum load-balancing** *number* command, the device changes the next-hop addresses of the routes to be advertised to a local address no matter whether the routes are used for load balancing. However, in RR or BGP confederation scenarios, the device does not change the next-hop addresses of non-local routes to be advertised to a local address.

If you run the **maximum load-balancing** { **ebgp** | **ibgp** } *number* command, the device does not change the next-hop addresses of the routes to be advertised to a local address no matter whether the routes are used for load balancing.

If you run the **maximum load-balancing** [**ebgp** | **ibgp**] *number* **ecmp-nexthop-changed** command, the device changes the next-hop addresses of the routes to be advertised to a local address only when the routes are used for load balancing.

Configuration Impact

If the **maximum load-balancing** command is run for multiple times, the latest configuration overrides the previous one.

Precautions

The **maximum load-balancing** *number* command cannot be configured together with the **maximum load-balancing ebgp** *number* or **maximum load-balancing ibgp** *number* command.

Example

Set two equal-cost routes to a specified destination.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] maximum load-balancing 2
```


7.8.95 maximum load-balancing eibgp

Function

The **maximum load-balancing eibgp** command configures the maximum number of EBGP and IBGP routes for load balancing.

The **undo maximum load-balancing eibgp** command deletes the configured maximum number of EBGP and IBGP routes for load balancing.

By default, the maximum number of EBGP and IBGP routes for load balancing is not configured.

Format

maximum load-balancing eibgp *number* [**ecmp-nexthop-changed**]

undo maximum load-balancing eibgp

Parameters

Parameter	Description	Value
number	Specifies the maximum number of equal-cost EBGP and IBGP routes.	The value is an integer that ranges from 1 to 8 on the S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S and S6720S-S. The value ranges from 1 to 16 on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S
ecmp-nexthop-changed	Configures a BGP device to change the next-hop addresses of only the routes that participate in load balancing to its address.	-

Views

BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **maximum load-balancing eibgp** command is used in a VPN where a CE is dual-homed to two PEs. When the CE and one PE belong to an AS and the CE and

the other PE belong to a different AS, you can set the number of EBGP and IBGP routes to be used for load balancing. This allows VPN traffic to be balanced among EBGP and IBGP routes.

After BGP load balancing is configured, BGP routes that meet the following conditions and have the same AS_Path attribute will become equal-cost routes and implement load balancing:

- PrefVal values are the same.
- Local_Pref attributes are the same.
- All BGP routes are summarized or non-summarized routes.
- Accumulated Interior Gateway Protocol metric (AIGP) values are the same.
- Lengths of the AS_path attributes are the same.
- Origin types (IGP, EGP, or Incomplete) are the same.
- Multi_Exit Discriminator (MED) values are the same.
- Protocol priorities are the same. By default, EBGP and IBGP routes have the same protocol priority (255). If the EBGP or IBGP route protocol priority is changed using a route-policy, load balancing cannot be implemented.
- Load balancing cannot be implemented between black-hole routes and non-black-hole routes.
- Load balancing cannot be implemented between labeled routes and non-labeled routes.
- Load balancing cannot be implemented between local routes and non-local routes.

Configuring BGP load balancing better utilizes network resources.

If you run the **maximum load-balancing eibgp** *number* command, the device changes the next-hop addresses of the routes to be advertised to a local address no matter whether the routes are used for load balancing. However, in RR or BGP confederation scenarios, the device does not change the next-hop addresses of non-local routes to be advertised to a local address.

If you run the **maximum load-balancing eibgp** *number* **ecmp-nexthop-changed** command, the device changes the next-hop addresses of the routes to be advertised to a local address only when the routes are used for load balancing.

Configuration Impact

If the **maximum load-balancing eibgp** command is run for multiple times, the latest configuration overrides the previous one.

Load balancing cannot be implemented between crossed and non-crossed routes.

Example

Set the maximum number of EBGP and IBGP routes for load balancing to 3.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vpna
[HUAWEI-vpn-instance-vpna] route-distinguisher 100:1
[HUAWEI-vpn-instance-vpna-af-ipv4] quit
[HUAWEI-vpn-instance-vpna] quit
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family vpn-instance vpna
[HUAWEI-bgp-vpna] maximum load-balancing eibgp 3
```


7.8.96 network (BGP)

Function

The **network** command configures BGP to statically add routes in the IP routing table to the BGP routing table and advertise these routes to peers.

The **undo network** command deletes the routes statically added to the BGP routing table.

By default, BGP does not statically add routes in the IP routing table to the BGP routing table.

Format

network { *ipv4-address* [*mask* | *mask-length*] | *ipv6-address prefix-length* } [**route-policy** *route-policy-name*]

undo network { *ipv4-address* [*mask* | *mask-length*] | *ipv6-address prefix-length* }

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of the route to be imported by BGP.	It is in dotted decimal notation.
mask	Specifies the IP address mask. If no mask is specified, the IP address is considered as a classful address.	It is in dotted decimal notation.
mask-length	Specifies the mask length of the IP address. If no mask length is specified, the IP address is considered as a classful address.	The value is an integer that ranges from 0 to 32.
ipv6-address	Specifies the IPv6 address of the route to be imported by BGP.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the prefix length of the IPv6 network address advertised by BGP.	The value is an integer that ranges from 0 to 128.
route-policy *route-policy-name*	Specifies the name of the route-policy that is used for route import.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP itself cannot discover routes. Instead, it imports routes discovered by other protocols such as IGP or static routes into the BGP routing table. These imported routes then are transmitted within an AS or between ASs. Before adding routes to the BGP routing table, BGP filters these routes by the routing protocol. If routes in the local IP routing table need to be statically added to the BGP routing table and then advertised, you can use the **network** command.

The Origin attribute of the routes imported into the BGP routing table by using the **network** command is IGP.

If a route with a specific prefix or mask is added to the BGP routing table by using the **network** command, this route is the optimal route selected from all types of protocol routes. Unlike the **network** command, the **import-route (BGP)** command is used to add all routes of a particular protocol such as RIP, OSPF, IS-IS, static route, or direct route to the BGP routing table.

Precautions

The **network** command is used to import exactly-matching routes. This means that only the routes in the local IP routing table that exactly match the specified destination address and prefix length can be added to the BGP routing table. If *mask* is not specified, routes are exactly matched against the natural network mask.

When using the **undo network** command to delete the existing configuration, specify a correct mask.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, routes in the local routing table are added to the BGP routing table.

Example

Configure BGP to import the local route 10.0.0.0/16.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] network 10.0.0.0 255.255.0.0
```

7.8.97 nexthop recursive-lookup (BGP)

Function

The **nexthop recursive-lookup** command configures BGP to recurse the next hop based on a routing policy.

The **undo nexthop recursive-lookup** command restores the default setting.

By default, BGP does not recurse the next hop based on a routing policy.

Format

nexthop recursive-lookup route-policy *route-policy-name*

undo nexthop recursive-lookup route-policy

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Indicates the name of a routing policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP needs route recursion in case of indirect next hops. If an indirect next hop is not recursed based on a routing policy, a BGP route may recurse to an incorrect forwarding path, causing traffic loss. Therefore, next hop recursion should be performed according to certain conditions.

The **nexthop recursive-lookup route-policy** *route-policy-name* command can be run to control next hop recursion based on a routing policy. If a recursive route is filtered out by the routing policy, the route is considered unreachable. This prevents BGP routes from recursing to incorrect forwarding paths.

Prerequisites

The next hop to which a BGP route recurses has been determined and a routing policy has been configured.

NOTICE

Before configuring a routing policy, ensure that all desired recursive routes will not be filtered out by the routing policy.

Precautions

The command does not apply to the routes received from directly connected EBGP peers or LinkLocal peers and default routes.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, route recursion is not restricted.

Example

Configure next hop recursion based on the routing policy **rp_nexthop**.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] nexthop recursive-lookup route-policy rp_nexthop
```

7.8.98 nexthop recursive-lookup delay

Function

The **nexthop recursive-lookup delay** command configures the delay in responding to changes of the next hop.

The **undo nexthop recursive-lookup delay** command restores the default setting.

By default, the delay in responding to changes of the next hop is not configured.

Format

nexthop recursive-lookup delay [*delay-time*]

undo nexthop recursive-lookup delay

Parameters

Parameter	Description	Value
delay-time	Specifies the delay in responding to changes of the next hop.	The value is an integer that ranges from 1 to 100, in seconds. The default value is 5 seconds.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-IPv4 multicast address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the route path on the upstream of a PE connected to an RR changes, if the PE detects that the recursive next hop becomes unreachable before the RR instructs the PE to switch the route, the PE withdraws the original optimal route advertised to its connected CE. After the RR re-advertises the switched route to the PE, the PE re-advertises an optimal route to the CE after route selection. During the route switchover, a huge volume of traffic will be dropped. If the **nexthop recursive-lookup delay** command is run on the PE to delay responding to the next hop unreachable event and to respond to this event only after the RR advertises the switched route, the volume of lost traffic will be reduced during route switchover.

Example

Set the delay in responding to changes of the next hop to 10s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] nexthop recursive-lookup delay 10
```

7.8.99 nexthop recursive-lookup non-critical-event delay

Function

The **nexthop recursive-lookup non-critical-event delay** command enables a device to respond to non-urgent BGP next-hop recursion changes after a specified delay time.

The **undo nexthop recursive-lookup non-critical-event delay** command enables a device to immediately respond to non-urgent BGP next-hop recursion changes.

By default, a device immediately responds to non-urgent BGP next-hop recursion changes.

Format

nexthop recursive-lookup non-critical-event delay [*delay-time*]

undo nexthop recursive-lookup non-critical-event delay

Parameters

Parameter	Description	Value
delay-time	Specifies the delay time.	The value is an integer that ranges from 1 to 100, in seconds. The default value is 5.

Views

BGP view, BGP-IPv4 unicast address family view, or BGP-VPN instance IPv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If repeated recursion causes the BGP next hop to frequently change on a device, the device frequently selects and advertises routes. To prevent unwanted re-routing and route advertisement, run the **nexthop recursive-lookup delay** command to enable the device to respond to BGP next-hop recursion changes after a specified delay time. The recursion results are as follows:

- Urgent recursion result change: The recursive next hop is changed, and BGP route reachability is also changed. For example, if a fault occurs on a network, a device finds no next-hop route or tunnel to which a BGP route recurses. As a result, traffic is interrupted.
- Non-urgent recursion result change: The recursive next hop is changed, and BGP route reachability is not affected. For example, after the interface or type of a tunnel to which the next hop of a BGP route recurses is changed, traffic keeps traveling over the BGP route.

To delay a response to the non-urgent recursion change, run the **nexthop recursive-lookup non-critical-event delay** command, not the **nexthop recursive-lookup delay** command.

Configuration Impact

After the **nexthop recursive-lookup delay** command is run, the device delays responses to all next-hop recursion changes. After the **nexthop recursive-lookup non-critical-event delay** command is run, the device delays responses only to non-urgent BGP next-hop recursion changes. If both commands are run, the

nexthop recursive-lookup non-critical-event delay command takes precedence over the **nexthop recursive-lookup non-critical-event delay** command. **Table 7-192** provides examples of the two commands.

Table 7-192 Functions and their descriptions

Example	Description
[HUAWEI-bgp] **nexthop recursive-lookup delay**	The device responds to all BGP next-hop recursion changes after a 5-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup non-critical-event delay**	The device immediately responds to urgent BGP next-hop recursion changes and responds to non-urgent BGP next-hop recursion changes after a 5-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup delay 3**	The device responds to all BGP next-hop recursion changes after a 3-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup non-critical-event delay 6**	The device immediately responds to urgent BGP next-hop recursion changes and responds to non-urgent BGP next-hop recursion changes after a 6-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup delay** [HUAWEI-bgp] **nexthop recursive-lookup non-critical-event delay**	The device responds to all BGP next-hop recursion changes after a 5-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup delay 3** [HUAWEI-bgp] **nexthop recursive-lookup non-critical-event delay**	The device responds to urgent BGP next-hop recursion changes after a 3-second delay and responds to non-urgent BGP next-hop recursion changes after a 5-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup delay 3** [HUAWEI-bgp] **nexthop recursive-lookup non-critical-event delay 6**	The device responds to urgent BGP next-hop recursion changes after a 3-second delay and responds to non-urgent BGP next-hop recursion changes after a 6-second delay.
[HUAWEI-bgp] **nexthop recursive-lookup delay** [HUAWEI-bgp] **nexthop recursive-lookup non-critical-event delay 6**	The device responds to urgent BGP next-hop recursion changes after a 5-second delay and responds to non-urgent BGP next-hop recursion changes after a 6-second delay.

Precautions

The delay time specified in the **nexthop recursive-lookup non-critical-event delay** command must be greater than or equal to that specified in the **nexthop recursive-lookup delay** command if both commands are run.

Example

Enable the device to respond to non-urgent BGP next-hop recursion changes after a **10**-second delay.


```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] nexthop recursive-lookup non-critical-event delay 10
```

7.8.100 out-delay

Function

The **out-delay** command configures a delay for sending Update packets to all BGP peers.

The **undo out-delay** command deletes the configured delay value.

The default delay value is 0, indicating that the intermediate device on the primary path sends Update packets without a delay.

Format

out-delay *delay-value*

undo out-delay

Parameters

Parameter	Description	Value
delay-value	Specifies the delay for sending Update packets.	The value is an integer that ranges from 0 to 3600, in seconds.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a scenario with both primary and backup routes, traffic may be lost after it switches back to the primary path. Use a VPN FRR scenario as an example. PE3 and CE2 connect both to PE1 and PE2. The primary path is PE3 -> PE1 -> CE2, and the backup path is PE3 -> PE2 -> CE2. CE2 uses BGP to communicate with PE1 and PE2. FRR is configured on PE3. If PE1 restarts or the link between PE3 and PE1 fails, traffic switches from the primary path to the backup path. After the primary path recovers, traffic switches back to the primary path. If PE3 completes refreshing forwarding entries before PE1 does, PE1 may temporarily fail to forward traffic from PE3, and packet loss may occur. The severity of packet loss is proportional to the number of routes stored on PE1.

To solve this problem, run the **out-delay** command on PE1 to configure a delay for sending Update packets. An appropriate delay ensures that traffic switches back to the primary path after PE1 completes refreshing forwarding entries.

To configure a delay for sending Update packets to all BGP peers, run the **out-delay** command. To configure a delay for sending Update packets to a specified BGP peer or peer group, run the **peer out-delay** command.

Precautions

If you run the **peer out-delay** command more than once, the latest configuration overrides the previous one.

If the **out-delay** and **peer route-update-interval** commands are both configured, only the **out-delay** command takes effect.

If a network has high route convergence requirements, do no use the **out-delay** command.

Example

Set the delay for sending Update packets to all BGP peers to 300s in the BGP view.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] out-delay 300
```

7.8.101 peer advertise-community

Function

The **peer advertise-community** command configures a device to advertise a community attribute to its peer or peer group.

The **undo peer advertise-community** command restores the default setting.

By default, a device advertises no community attribute to its peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **advertise-community**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **advertise-community**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.

Parameter	Description	Value
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view, BPG-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP L2VPN-AD address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer advertise-community** command is used to configure a device to advertise a community attribute to its peer or peer group. If a device advertises a community attribute to its peer group, all the members of the peer group will inherit the configuration. This simplifies the application of routing policies and facilitates route maintenance and management.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

A specific community attribute has been defined in a routing policy.

Example

Configure a device to advertise a community attribute to its peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 advertise-community
```


7.8.102 peer advertise-ext-community

Function

The **peer advertise-ext-community** command enables a device to advertise an extended community attribute to its peer or peer group.

The **undo peer advertise-ext-community** command restores the default setting.

By default, a device advertises no extended community attribute to its peer or peer group.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

peer { *group-name* | *ipv6-address* } **advertise-ext-community**

undo peer { *group-name* | *ipv6-address* } **advertise-ext-community**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Views

BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer advertise-ext-community** command is used to enable a device to advertise an extended community attribute to a specified peer or peer group. If a device advertises an extended community attribute to a specified peer group, all the members of the peer group will inherit the configuration. This simplifies the application of routing policies and facilitates route maintenance and management.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

A specific extended community attribute has been defined in a routing policy.

Example

Configure a device to advertise an extended community attribute to its peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 2:3::4:5 as-number 200
[HUAWEI-bgp] ipv6-family vpn-instance abc
[HUAWEI-bgp6-abc] peer 2:3::4:5 as-number 200
[HUAWEI-bgp6-abc] peer 2:3::4:5 advertise-ext-community
```

7.8.103 peer allow-as-loop

Function

The **peer allow-as-loop** command sets the number of local AS number repetitions.

The **undo peer allow-as-loop** command restores the default setting.

By default, the local AS number cannot be repeated.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **allow-as-loop** [*number*] [**global-as** [**vpn-as**]]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **allow-as-loop** [*number*] [**global-as** [**vpn-as**]]

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a peer.	The value is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
number	Specifies the number of local AS number repetitions.	The value is an integer in the range from 1 to 10. The default value is 1.
global-as	Specifies a global AS number of BGP.	-
vpn-as	Specifies the AS number of a VPN instance.	-

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, and BGP L2VPN-AD address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.
- The **global-as** and **vpn-as** parameters apply only to BGP VPN instance IPv4 and BGP VPN instance IPv6 address family views.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view , BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP uses AS numbers to detect routing loops. The AS numbers in the AS_Path of each received route are matched against the local AS number configured using the **bgp** command, the fake AS number configured using the **peer fake-as**

command, and the VPN instance AS number configured using the **as-number** command. The largest number of times any of the configured AS numbers is repeated is considered as the maximum number. In the Hub and Spoke networking, if EBGP runs between a Hub-PE and a Hub-CE on a Hub site, the route sent from the Hub-PE to the Hub-CE carries the AS number of the Hub-PE. If the Hub-CE sends a routing update to the Hub-PE, the Hub-PE will deny it because the routing update contains the AS number of the Hub-PE.

To ensure proper route transmission in the Hub and Spoke networking, configure all the BGP peers on the path, along which the Hub-CE advertises private network routes to the Spoke-CE, to accept the routes in which the AS number repeats once.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

If the **peer allow-as-loop** command is run for a peer or peer group multiple times, the latest configuration overrides the previous one.

Example

Set the number of local AS number repetitions to 2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 allow-as-loop 2
```

7.8.104 peer as-number

Function

The **peer as-number** command creates a peer or configures an AS number for a specified peer group.

The **undo peer as-number** command deletes a specified peer or the AS number of a specified peer group.

By default, no BGP peer is configured, and no AS number is specified for a peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **as-number** { *as-number-plain* | *as-number-dot* }

peer *group-name* **as-number** { *as-number-plain* | *as-number-dot* } [**optional-as** { *optional-as-number-plain* | *optional-as-number-dot* } &<1-5>]

undo peer { *group-name* **as-number** | *ipv4-address* | *ipv6-address* }

undo peer *group-name* **as-number** { *as-number-plain* | *as-number-dot* } [**optional-as** { *optional-as-number-plain* | *optional-as-number-dot* } &<1-5>]

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer. The IPv4 address can be the IP address of an interface that is directly connected to the peer or the IP address of a loopback interface of the reachable peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
as-number-plain	Specifies the number of the AS, in integer format.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of the AS, in dotted notation.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.
optional-as	Specifies an optional AS number.	-
optional-as-number-plain	Specifies an integral optional AS number.	The value is an integer that ranges from 1 to 4294967295.
optional-as-number-dot	Specifies an optional AS number in dotted notation.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer** { *ipv4-address* | *ipv6-address* } **as-number** { *as-number-plain* | *as-number-dot* } command is used to create a BGP peer.

The **peer** *group-name* **as-number** { *as-number-plain* | *as-number-dot* } command is used to configure an AS number for a specified peer group.

Dynamic peers on the same network segment must be added to the same peer group. If the dynamic peers reside in different ASs, run the **peer** *group-name* **as-number** { *as-number-plain* | *as-number-dot* } **optional-as** { *optional-as-number-plain* | *optional-as-number-dot* } &<1-5> command to configure an optional AS number for the peer group.

Precautions

If a peer does not join any peer group or the peer group to which a peer belongs is not configured with an AS number, deleting the AS number of the peer will reset the peer relationship.

If a peer in a peer group is not configured with an AS number, deleting the AS number of the peer group will interrupt the connection on the peer.

The AS number of an external peer group must be different from the local AS number.

Example

Set the AS number to 100 for IPv4 peer 10.1.1.1.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.1 as-number 100
```

Set the AS number to 100 for a peer group named test.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group test
[HUAWEI-bgp] peer test as-number 100
```


Set the peer AS number and optional AS number for the peer group named **dynamic-group** to 65546 and 65547, respectively.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group dynamic-group external
[HUAWEI-bgp] peer dynamic-group as-number 65546 optional-as 65547
```

7.8.105 peer as-path-filter

Function

The **peer as-path-filter** command configures a BGP route filtering policy based on an AS_Path list for a peer or peer group.

The **undo peer as-path-filter** command cancels the existing configuration.

By default, no route filtering policy based on an AS_Path list is configured for a peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } { **import** | **export** }

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } { **import** | **export** }

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
as-path-filter-number	Specifies the number of an AS_Path filter.	The value is an integer that ranges from 1 to 256.

Parameter	Description	Value
as-path-filter-name	Specifies the name of an AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters. It cannot be all numbers. When double quotation marks are used around the string, spaces are allowed in the string.
import	Applies a route filtering policy to received routes.	-
export	Applies a route filtering policy to routes to be advertised.	-

NOTE

- *ipv4-address* is valid only in the BGP view, BPG-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **peer as-path-filter** command is used to apply a route filtering policy based on an AS_Path list to BGP routes, the routes that do not match the policy are filtered out.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

The **ip as-path-filter** command has been run to define an AS-Path filter.

Precautions

Only one AS_Path filter can be used to filter the routes received from the same peer. Similarly, only one AS_Path filter can be used to filter routes to be advertised to the same peer.

Creating an AS_Path filter before it is referenced is recommended. By default, the command in an IPv4 address family cannot reference a non-existent AS_Path filter, but the command in an IPv6 address family can reference a non-existent AS_Path filter. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent AS_Path filter is referenced using the current command, all routes are advertised to the specified peer, or all routes advertised by the specified peer are accepted.

Example

Configure an AS_Path filter for a peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 as-path-filter 3 export
```

7.8.106 peer bfd

Function

The **peer bfd** command sets BFD detection parameters for a peer or peer group.

The **undo peer bfd** command restores default BFD detection parameter values.

By default, the interval for sending BFD packets is 1000 ms, the interval for receiving BFD packets is 1000 ms, and the local detection multiplier is 3.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

peer { *group-name* | *ipv4-address* } **bfd** { **min-tx-interval** *min-tx-interval* | **min-rx-interval** *min-rx-interval* | **detect-multiplier** *multiplier* } *

peer { *group-name* | *ipv6-address* } **bfd** { **min-tx-interval** *min-tx-interval* | **min-rx-interval** *min-rx-interval* | **detect-multiplier** *multiplier* } *

undo peer { *group-name* | *ipv4-address* } **bfd** { **min-tx-interval** | **min-rx-interval** | **detect-multiplier** } *

undo peer { *group-name* | *ipv6-address* } **bfd** { **min-tx-interval** | **min-rx-interval** | **detect-multiplier** } *

undo peer { *group-name* | *ipv4-address* } **bfd** { **min-tx-interval** *min-tx-interval* | **min-rx-interval** *min-rx-interval* | **detect-multiplier** *multiplier* } *

undo peer { *group-name* | *ipv6-address* } **bfd** { **min-tx-interval** *min-tx-interval* | **min-rx-interval** *min-rx-interval* | **detect-multiplier** *multiplier* } *

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
min-tx-interval *min-tx-interval*	Specifies the interval at which BFD packets are sent.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
min-rx-interval *min-rx-interval*	Specifies the interval at which BFD packets are received.	The value is an integer that ranges from 100 to 1000, in milliseconds. • After the **set service-mode enhanced** command is configured on the S5731-H, S5731-S, S5731S-H, and S5731S-S, the value ranges from 3 to 1000. • After the **set service-mode enhanced-bfd** command is configured on the S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S, the value ranges from 3 to 1000.
detect-multiplier *multiplier*	Specifies the local detection time multiplier.	The value is an integer that ranges from 3 to 50. By default, the value is 3.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BFD provides millisecond-level fault detection. It helps BGP to detect faults in neighboring devices or links more quickly, and instructs BGP to recalculate routes for correct packet forwarding. The **peer bfd** command can be used to set the values of BFD session parameters on a specified interface.

The BFD configuration of a peer takes precedence over that of the peer group to which the peer belongs. If BFD is not configured on a peer and the peer group to which the peer belongs is enabled with BFD, the peer will inherit the BFD configurations of the peer group.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

A BFD session can be established only when the corresponding BGP session is in the Established state.

Precautions

If the **peer bfd** command is run multiple times, the latest configuration overwrites the previous one. The BFD session uses the latest parameters as the detection parameters.

Assume that BFD is configured on a peer group. If the **peer bfd block** command is not run on members of the peer group, the members will establish BFD sessions.

If BFD parameters are set on a peer, a BFD session will be established by using the BFD parameters on the peer.

Example

Configure BFD and set detection parameters on peer 10.2.2.9.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.2.2.9 as-number 200
[HUAWEI-bgp] peer 10.2.2.9 bfd min-tx-interval 100 min-rx-interval 100 detect-multiplier 5
```

7.8.107 peer bfd block

Function

The **peer bfd block** command prevents a peer from inheriting the BFD function of its peer group.

The **undo peer bfd block** command restores the default configuration.

By default, a peer inherits the BFD function from its peer group.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

peer { *ipv4-address* | *ipv6-address* } **bfd block**

undo peer { *ipv4-address* | *ipv6-address* } **bfd block**

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BFD provides millisecond-level fault detection. It helps BGP to detect faults in neighboring devices or links more quickly, and instructs BGP to recalculate routes for correct packet forwarding. If a peer group is configured with BFD, all members of the peer group will establish BFD sessions. The **peer bfd block** command can be used to prevent a peer from inheriting the BFD function from its peer group.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

A BFD session has been established.

Precautions

After the **peer bfd block** command is run on a peer, the corresponding BFD session on the peer is deleted. As a result, fast link fault detection cannot be implemented.

The **peer bfd block** command and the **peer bfd enable** command are mutually exclusive. After the **peer bfd block** command is run, related BFD sessions are automatically deleted.

Example

Prevent peer 10.2.2.9 from inheriting the BFD function of its peer group.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.2.2.9 as-number 200
[HUAWEI-bgp] peer 10.2.2.9 bfd block
```

7.8.108 peer bfd enable

Function

The **peer bfd enable** command enables BFD for peers or a peer group.

The **undo peer bfd enable** command disables BFD for peers or a peer group.

By default, BFD is disabled for peers or a peer group.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **bfd enable** [**single-hop-prefer**]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **bfd enable**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
single-hop-prefer	Indicates that single-hop BFD sessions are preferentially established. **single-hop-prefer** takes effect only on IBGP peers. By default, **single-hop-prefer** is not specified, and multi-hop sessions are established between direct IBGP peers (Huawei devices). To interconnect a Huawei device and a non-Huawei device that defaults the sessions between IBGP peers to single-hop, configure **single-hop-prefer** in the command.	-

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP uses BFD to quickly detect faults in links between BGP4 peers. This accelerates network convergence. The **peer bfd enable** command is used to configure a device to establish a BFD session with its peer or peer group by using default detection parameter values.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

A BFD session can be established only when the corresponding BGP peer relationship is in the Established state.

Precautions

The BFD configuration of a peer takes precedence over that of the peer group to which the peer belongs. If BFD is not configured on a peer and the peer group to which the peer belongs is enabled with BFD, the peer inherits the BFD configurations from the peer group.

Before enabling BFD on a BGP peer, enable BFD in the system view. If no BFD detection parameter is specified, a BFD session is established by using default parameter values.

NOTE

The **peer bfd block** command and the **peer bfd enable** command are mutually exclusive.

Example

Configure BFD for peer 10.2.2.9.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.2.2.9 as-number 200
[HUAWEI-bgp] peer 10.2.2.9 bfd enable
```

7.8.109 peer capability-advertise

Function

The **peer capability-advertise** command enables a BGP device to advertise optional BGP capabilities to its peer.

The **undo peer capability-advertise** command restores the default setting.

By default, a BGP device advertises route-refresh and 4-byte AS number capabilities to its peer.

Format

peer { *group-name* | *ipv4-address* } **capability-advertise** { **4-byte-as** | **route-refresh** | **conventional** }

peer *ipv6-address* **capability-advertise** { **4-byte-as** | **route-refresh** }

undo peer { *group-name* | *ipv4-address* } **capability-advertise** { **4-byte-as** | **route-refresh** | **conventional** }

undo peer *ipv6-address* **capability-advertise** { **4-byte-as** | **route-refresh** }

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
4-byte-as	Indicates the 4-byte AS number capability.	-
route-refresh	Indicates the route-refresh capability.	-

Parameter	Description	Value
conventional	Indicates the regular router capability.	-

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer capability-advertise route-refresh** command is used to configure a device to advertise the route-refresh capability to its peer.

The **peer capability-advertise conventional** command is used to configure a device to advertise the regular router capability to its peer. If **conventional** is specified in the command, the router does not have all extension functions such as route-refresh capability, GR capability, and multi-address family negotiation. This allows the router to be compatible with routers of earlier versions.

The **peer capability-advertise 4-byte-as** command is used to configure a device to advertise the 4-byte AS number capability to its peer. If AS number resources are used up, this command can be used to configure devices to use 4-byte AS numbers.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

The corresponding BGP session must be in the Established state.

Precautions

If you enable or disable the route-refresh, general router, or 4-byte AS number function, the BGP peer relationship will be re-established, which can lead to a temporary network interruption. Therefore, exercise caution when running the related commands.

Example

Configure a BGP device to advertise the route-refresh capability to its peer.


```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.89.2.33 as-number 100
[HUAWEI-bgp] peer 10.89.2.33 capability-advertise route-refresh
```

7.8.110 peer connect-interface

Function

The **peer connect-interface** command specifies a source interface from which BGP packets are sent, and a source address used for initiating a connection.

The **undo peer connect-interface** command restores the default setting.

By default, the outbound interface of a BGP packet serves as the source interface of the BGP packet.

Format

peer { *group-name* | *ipv4-address* } **connect-interface** *interface-type interface-number* [*ipv4-source-address*]

peer { *group-name* | *ipv6-address* } **connect-interface** *interface-type interface-number* [*ipv6-source-address*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **connect-interface**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
interface-type interface-number	Specifies the type and number of an interface.	-
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
ipv4-source-address	Specifies an IPv4 source address used for establishing a connection.	It is in dotted decimal notation.
ipv6-source-address	Specifies an IPv6 source address used for establishing a connection.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The command is used in the following scenarios:

- To establish a BGP peer connection between two indirectly connected physical interfaces, the **peer connect-interface** command must be run on both sides.

  To establish a BGP peer connection between a local loopback interface and a remote interface, the loopback interface must be specified as the source interface in the command. To establish a BGP peer connection between a local physical interface and a remote interface, the physical interface must be specified as the source interface in the command.
- If multiple IP addresses are configured for the physical interface that needs a BGP peer connection established with a remote interface, the **peer connect-interface** command must be run, with *ipv4-source-address* or *ipv6-source-address* set to the source IP address.
- If two devices need multiple BGP peer connections established through different interfaces, the **peer connect-interface** command must be run for each BGP peer connection, with the source interface specified.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

Running the **peer connect-interface** command causes the teardown and re-establishment of peer relationships.

NOTE

Because the BGP peer relationships in various address families on the same device share one TCP connection, **connect-interface** configured in the BGP view can be inherited in the IPv4 unicast address family.

To enable a device to send BGP packets even if its physical interface fails, you can configure the device to use a loopback interface as the source interface of the BGP packets. When configuring a device to use a loopback interface as the source interface of BGP packets, note the following points:

- The loopback interface of the device's BGP peer must be reachable.
- In the case of an EBGP connection, the **peer ebgp-max-hop** command needs to be run to enable the two devices to establish an indirect peer relationship.

NOTE

If the specified interface borrows the IP address of another interface and then the IP address of the specified interface is changed, BGP still uses the borrowed IP address to keep the connection if no connection reestablishment is triggered, and data receiving and sending is not affected; if connection reestablishment is triggered, BGP uses the new IP address to reestablish the connection.

Example

Specify a source interface for sending BGP packets and a source address for initiating a connection.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.16.2.3 as-number 100
[HUAWEI-bgp] peer 10.16.2.3 connect-interface vlanif 100
```

7.8.111 peer connected-check-ignore

Function

The **peer connected-check-ignore** command configures a device not to check the hop count when establishing a one-hop EBGP peer relationship using loopback interface addresses.

The **undo peer connected-check-ignore** command cancels the configuration.

By default, a device checks the hop count when establishing a one-hop EBGP peer relationship using loopback interface addresses.

Format

peer{ *group-name* | *ipv4-address* | *ipv6-address* } **connected-check-ignore**

undo peer{ *group-name* | *ipv4-address* | *ipv6-address* } **connected-check-ignore**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The value is a string of 1 to 47 case-sensitive characters, spaces not supported.
ipv4-address	Specifies the IPv4 address of a peer.	-
ipv6-address	Specifies the IPv6 address of a peer.	-

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, or BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To allow one-hop EBGP peer relationships to be established using loopback interface addresses, run the **peer connected-check-ignore** command or the **peer ebgp-max-hop** command (in which *hop-count* is greater than or equal to 2).

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

If the **peer connected-check-ignore** command is used on one end of an EBGP connection, it must also be used on the other end.

Example

Configure the device not to check the hop count when establishing a one-hop EBGP peer relationship using loopback interface addresses.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 connected-check-ignore
```

7.8.112 peer default-route-advertise

Function

The **peer default-route-advertise** command configures a BGP device to advertise a default route to its peer or peer group.

The **undo peer default-route-advertise** command restores the default setting.

By default, a BGP device does not advertise default routes to its peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **default-route-advertise** [**route-policy** *route-policy-name*] [**conditional-route-match-all** { *ipv4-address1* { *mask1* | *mask-length1* } } &<1-4> | **conditional-route-match-any** { *ipv4-address2* { *mask2* | *mask-length2* } } &<1-4>]

peer { *group-name* | *ipv4-address* | *ipv6-address* } **default-route-advertise** [**route-policy** *route-policy-name*] [**conditional-route-match-all** { *ipv6-address1 prefix-length1* } &<1-4> | **conditional-route-match-any** { *ipv6-address2 prefix-length2* } &<1-4>]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **default-route-advertise**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	The value is in dotted decimal notation.
prefix-length1	Specifies the IPv6 prefix range using the mask length.	The value is an integer that ranges from 0 to 128. If ::0 less-equal 128 is used, all IPv6 addresses will be matched.
prefix-length2	Specifies the IPv6 prefix range using the mask length.	The value is an integer that ranges from 0 to 128. If ::0 less-equal 128 is used, all IPv6 addresses will be matched.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
conditional-route-match-all	Configures the BGP device to send default routes to a peer or peer group when the routing table contains all conditional routes.	-

Parameter	Description	Value
ipv4-address1	Specifies the IPv4 address of conditional routes.	The value is in dotted decimal notation.
ipv6-address1	Specifies the IPv6 address of conditional routes.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
mask1	Specifies the mask of conditional routes.	The value is in dotted decimal notation.
mask-length1	Specifies the mask length of conditional routes.	The value is an integer that ranges from 0 to 32.
conditional-route-match-any	Configures the BGP device to send default routes to a peer or peer group when the routing table contains any conditional route.	-
ipv4-address2	Specifies the IPv4 address of conditional routes.	The value is in dotted decimal notation.
ipv6-address2	Specifies the IPv6 address of conditional routes.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
mask2	Specifies the mask of conditional routes.	The value is in dotted decimal notation.
mask-length2	Specifies the mask length of conditional routes.	The value is an integer that ranges from 0 to 32.

NOTE

- *ipv4-address* is only valid in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-IPv4 multicast address family view, and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is only valid in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.
- **conditional-route-match-all** and **conditional-route-match-any** are valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-IPv6 unicast address family view, and BGP-VPN instance IPv4 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Default routes are commonly used on a network that meets the following conditions:

- Each device has multiple EBGP peers and receives all routes on the network from each EBGP peer.
- There are multiple route reflectors (RRs), and each RR receives all routes on the network.

If load balancing is not implemented on the network, a BGP peer receives at most one copy of active routes on the network. If load balancing is implemented on the network, the number of active routes received by a BGP peer will be increased by multiple times, causing the number of routes on the network to sharply increase. To greatly reduce the number of routes on such a network, configure a BGP device to advertise only default routes to its BGP peer and use default routes for traffic load balancing.

Prerequisites

BGP peer relationships have been established using the **peer as-number** command.

Precautions

After this command is run, a BGP device sends a default route with the next hop as itself to its peer or peer group regardless of whether default routes exist in the routing table.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the attributes of the default route to be advertised to the specified peer or peer group are not changed.

Example

Configure a BGP device to advertise a default route to its peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 default-route-advertise
```

7.8.113 peer description (BGP)

Function

The **peer description** command configures a description for a peer or peer group.

The **undo peer description** command deletes the description of a peer or peer group.

By default, no description is configured for a peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **description** *description-text*

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **description**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
description-text	Specifies a description, which can be letters and digits.	The value is a string of 1 to 80 characters, with spaces supported.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer description** can be used to configure a description for a peer or peer group, which facilitates management of peers or peer groups. The description records information about the peer or peer group, such as the VPN instance to which the peer or peer group belongs, and devices that establish peer relationships with the peer or peer group.

Precautions

The description configured by using the **peer description** command for a peer is displayed from the first non-space character, and a maximum of 80 characters can be displayed.

Follow-up Procedure

You can run **display bgp peer** *ipv4-address* **verbose** command can be used to view the description of a peer.

Example

Configure a description for a peer group named group1.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group group1
[HUAWEI-bgp] peer group1 description ISP1
```

7.8.114 peer ebgp-max-hop

Function

The **peer ebgp-max-hop** command configures a BGP device to establish an EBGP peer relationship with a peer on an indirectly-connected network and set the maximum number of hops between the two devices.

The **undo peer ebgp-max-hop** command restores the default setting.

By default, an EBGP connection can be set up only on a directly-connected physical link.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **ebgp-max-hop** [*hop-count*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **ebgp-max-hop**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
hop-count	Specifies the maximum number of hops.	The value is an integer that ranges from 1 to 255. The default value is 255. If the maximum number of hops is 1, a device cannot establish an EBGP connection with a peer on an indirectly-connected network.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A directly-connected physical link must be available between EBGP peers. Otherwise, the **peer ebgp-max-hop** command must be used to allow EBGP peers to establish a TCP connection over multiple hops.

If loopback interfaces are used to establish an EBGP peer relationship, the **peer ebgp-max-hop** command (*hop-count* ≥ 2) must be run; otherwise, the peer relationship cannot be established. If a one-hop EBGP peer relationship is established using loopback interface addresses, you can also run the **peer connected-check-ignore** command to establish an EBGP peer relationship.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

If the **peer ebgp-max-hop** command is used on one end of an EBGP connection, it must also be used on the other end.

The configurations of GTSM and EBGP-MAX-HOP affect the TTL values of sent BGP packets, and the configurations of the two functions are mutually exclusive.

Example

Allow an indirectly connected EBGP peer with the IP address of 10.1.1.2 to establish a connection with the local device.

```
<HUAWEI> system-view
```



```
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 ebgp-max-hop
```

7.8.115 peer enable (BGP)

Function

The **peer enable** command enables a BGP device to exchange routes with a specified peer or peer group in the address family view.

The **undo peer enable** command disables a BGP device from exchanging routes with a specified peer or peer group.

By default, only the peer in the BGP IPv4 unicast address family view is automatically enabled.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **enable**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **enable**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view,BGP-MDT address family view, BGP-MVPN address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, and BGP L2VPN-AD address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address

family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, only the peer in the BGP IPv4 unicast address family view is automatically enabled. After the **peer as-number** command is used in the BGP view, the system automatically runs the **peer enable** command to enable a peer. In other address family views, however, a peer must be enabled manually.

After the **undo default ipv4-unicast** command is run, the **peer enable** command needs to be run in the BGP view or the BGP-IPv4 unicast address family view to enable the IPv4 unicast address family for the created BGP peer.

Precautions

Enabling or disabling a BGP peer in an address family, for example, running the **peer enable** command or the **undo peer enable** command in a VPNv4 address family, causes teardown and re-establishment of the BGP connection of the peer in other address families.

Example

Disable a BGP device from exchanging IPv4 routes with a specified peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] undo peer 10.1.1.2 enable
```

7.8.116 peer fake-as

Function

The **peer fake-as** command specifies a fake AS number for a local peer.

The **undo peer fake-as** command restores the default setting.

By default, a peer uses the actual local AS number.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **fake-as** { *as-number-plain* | *as-number-dot* } [**prepend-global-as**]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **fake-as**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
as-number-plain	Specifies the number of the AS, in integer format.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of the AS, in dotted notation.	The value is in the x.y format. Here, "x" and "y" are integers that range from 1 to 65535 and 0 to 65535 respectively.
prepend-global-as	Indicates that the actual AS number is added to packets to be sent.	-

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer fake-as** command is often used for network deployment changing of carriers' networks. For example, if carrier A purchases the network of carrier B and the networks of the two carriers belong to different ASs, the ASs of the two carriers need to be combined and the AS number of carrier B needs to be changed to the AS number of carrier A. BGP peers (in another AS) of devices on carrier B's network may not be willing to have their BGP configurations changed or changed

immediately during network combination. As a result, the connections with these BGP peers will be interrupted.

To ensure that the ASs are combined properly, you can run the **peer fake-as** command on the ASBR on carrier's B network to set the AS number of carrier B to the fake AS number of carrier A. This setting enables the BGP peers of devices in carrier B's network to use the fake AS number to set up connections.

NOTE

If **prepend-global-as** is specified in the command and the local end establishes an EBGP peer relationship with the remote end using the fake AS number, the local end adds the actual AS number to the packets to be sent to the remote end.

Prerequisites

Peers have been created by using the **peer as-number** command.

Precautions

If the **peer fake-as** command is run several times for a peer or a peer group, the latest configuration will overwrite the previous one.

The **peer fake-as** command takes effect only for EBGP peers, and the configured fake AS number cannot be the same as the peer AS number.

After the 4-byte AS number capability is disabled on a peer, configuring a 4-byte fake AS number for the peer may cause a failure to establish a BGP session.

Example

Set a 2-byte fake AS number for a peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 fake-as 99
```

Set a 4-byte fake AS number for a peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 fake-as 100.200
```

7.8.117 peer filter-policy

Function

The **peer filter-policy** command configures a filtering policy to advertise routes to or receive routes from peers or peer group.

The **undo peer filter-policy** command deletes the filtering policy used to advertise routes to or receive routes from peers or peer group.

By default, no filtering policy is configured to advertise routes to or receive routes from peers or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **filter-policy** { *acl-number* | **acl-name** *acl-name* | *acl6-number* | **acl6-name** *acl6-name* } { **import** | **export** }

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **filter-policy** { *acl-number* | **acl-name** *acl-name* | *acl6-number* | **acl6-name** *acl6-name* } { **import** | **export** }

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
acl-number	Specifies the number of a basic ACL.	The value is an integer that ranges from 2000 to 2999.
acl-name *acl-name*	Specifies the name of an ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
acl6-number	Specifies the number of a basic IPv6 ACL.	The value is an integer that ranges from 2000 to 2999.
acl6-name *acl6-name*	Specifies the name of an IPv6 ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The name should start with a letter and can contain numbers, hyphens (-), or underscores (_).
import	Filters received routes.	-
export	Filters routes to be advertised.	-

NOTE

- *ipv4-address*, **acl-name** *acl-name*, and *acl-number* are valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address*, **acl6-name** *acl6-name*, and *acl6-number* are valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer filter-policy** command is used to configure an ACL-based policy for filtering routes received from peers or routes to be advertised to peers.

A proper basic ACL needs to be selected based on the address family of a peer or peer group.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

A basic ACL needs to be configured.

Precautions

If the **peer filter-policy** command is run multiple times, the latest configuration overwrites the previous one. For example, if the **peer 10.1.1.2 filter-policy 2600 import** command is run and then the **peer 10.1.1.2 filter-policy 2400 import** command is run, the configuration of the **peer 10.1.1.2 filter-policy 2400 import** command overwrites that of the **peer 10.1.1.2 filter-policy 2600 import** command.

When the **rule** command is run to configure rules for an ACL, only the source address range specified by **source** and the time period specified by **time-range** take effect.

Example

Set the IPv4 filtering policy for peers.

```
<HUAWEI> system-view
[HUAWEI] acl 2000
[HUAWEI-acl-basic-2000] rule permit
[HUAWEI-acl-basic-2000] quit
[HUAWEI] bgp 100
```



```
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 filter-policy 2000 import
```

Set the IPv6 filtering policy for peers.

```
<HUAWEI> system-view
[HUAWEI] acl ipv6 2001
[HUAWEI-acl6-basic-2001] rule permit
[HUAWEI-acl6-basic-2001] quit
[HUAWEI] bgp 100
[HUAWEI-bgp] peer fc00:0:0:2::3 as-number 200
[HUAWEI-bgp] ipv6-family unicast
[HUAWEI-bgp-af-ipv6] peer fc00:0:0:2::3 filter-policy 2000 import
```

7.8.118 peer group

Function

The **peer group** command adds a peer to a peer group.

The **undo peer group** command deletes a peer from a peer group and all configurations of the peer.

By default, no peer is in peer group.

Format

peer { *ipv4-address* | *ipv6-address* } **group** *group-name*

undo peer { *ipv4-address* | *ipv6-address* } **group** *group-name*

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view, and BGP-EVPN address family view.
- *ipv6-address* is valid only in the BGP view, BGP-IPv6 unicast address family view, and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view, BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On a large-scale BGP network, there are a large number of peers and many of them have the same routing policies. To configure these peers, you have to repeatedly use some commands. In such a case, configuring peer groups can simplify configurations. If you intend to perform the same configuration on several peers, create and configure a peer group. Then, add the peers to the peer group. The peers will inherit the configurations of the peer group.

Precautions

- Peers with different AS numbers can be added to the same peer group. If a peer has an AS number, the peer keeps its own AS number after being added to a peer group. If a peer has no AS number but the peer group to which the peer will be added has an AS number, the peer inherits the AS number of the peer group after being added to the peer group.
- The members of a peer group can be configured with different route receiving and advertising policies.
- The **undo peer group** command has the same function with the **undo peer** command and the **undo peer enable** command.
- By default, only the IPv4 peers configured in the BGP-IPv4 unicast address family view are automatically added to a peer group. That is, after the **peer** *ipv4-address* **group** *group-name* command is configured in the BGP view, the system configures the **peer** *ipv4-address* **group** *group-name* command in the BGP-IPv4 unicast address family view. In other address family views, peers must be manually added to a peer group.

Example

Create an IBGP peer group named test, and then add a peer to it.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group test internal
[HUAWEI-bgp] peer 10.1.1.1 group test
```

Create an EBGP peer group named **test** and add an IPv6 peer to this peer group.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group test external
```



```
[HUAWEI-bgp] peer 2001:db8:1::1 as-number 200
[HUAWEI-bgp] peer 2001:db8:1::1 group test
[HUAWEI-bgp] ipv6-family unicast
[HUAWEI-bgp-af-ipv6] peer test enable
[HUAWEI-bgp-af-ipv6] peer 2001:db8:1::1 group test
```

7.8.119 peer ignore

Function

The **peer ignore** command prevents a BGP device from establishing a session with a peer or peer group.

The **undo peer ignore** command permits a BGP device from establishing a session with a peer or peer group.

By default, the device is permitted to set up the session with the BGP peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **ignore**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **ignore**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-VPN instance IPv4 address family view, and BGP-VPN instance IPv6 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a BGP device needs to transitorily close the session with a peer and reconfiguring the peer requires too much workload, the **peer ignore** command can be used to avoid the reconfiguration of the peer. For example, if the peer relationship frequently alternates between Up and Down because of the upgrade of a peer or the adjustment of the link, the BGP peer relationship needs to be interrupted temporarily. In this case, the **peer ignore** command can be used on the stabler end to prevent frequent flapping of the route or peer relationship.

The **peer ignore** command is used to tear down sessions between a BGP device and its peer or peer group and delete all related routing information. In the case of a peer group, a large number of sessions are suddenly torn down.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

After the **peer ignore** command is run on a device, the session between the device and its peer is closed and all the related routing information is cleared.

After a BGP session is successfully established, running the **peer ignore** command interrupts the BGP session. The interrupted BGP session cannot be established again, and the status of the corresponding BGP peer relationship is displayed as Idle.

Running the **peer ignore** command together with the **peer enable** command equals running of the **reset bgp** command. Both methods can be used to configure a device to re-establish a session.

Example

Prohibit a device from establishing any session with peer 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 ignore
```

7.8.120 peer ip-prefix

Function

The **peer ip-prefix** command configures a route filtering policy based on an IP address prefix list for a peer or peer group.

The **undo peer ip-prefix** command cancels the route filtering policy based on an IP address prefix list of a peer or peer group.

By default, no route filtering policy based on an IP address prefix list is configured for a peer or peer group.

Format

peer { *group-name* | *ipv4-address* } **ip-prefix** *ip-prefix-name* { **import** | **export** }

peer { *group-name* | *ipv4-address* | *ipv6-address* } **ipv6-prefix** *ipv6-prefix-name* { **import** | **export** }

undo peer { *group-name* | *ipv4-address* } **ip-prefix** [*ip-prefix-name*] { **import** | **export** }

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **ipv6-prefix** [*ipv6-prefix-name*] { **import** | **export** }

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ip-prefix *ip-prefix-name*	Indicates the filtering policy that is based on the IPv4 prefix list of the peer or peer group.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
import	Applies a filtering policy to the routes received from a peer or peer group.	-
export	Applies a filtering policy to the routes to be advertised to a peer or peer group.	-
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
ipv6-prefix *ipv6-prefix-name*	Indicates the filtering policy that is based on the IPv6 prefix list of the peer or peer group.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view, and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer ip-prefix** command can be used to configure a route filtering policy that is based on an IP prefix list to filter routes received or routes to be advertised, implementing route control.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

If the **peer ip-prefix** command specifies an IP prefix list that does not exist for a peer or peer group, use the **ip ip-prefix** or **ip ipv6-prefix** command to create an IP prefix list.

Precautions

If an IP prefix list is specified for a peer group, all the members of the peer group inherit the configuration.

After an IP prefix list is specified for a peer or peer group, the peer or peers in the peer group filter the routes to be advertised to or received from other peers based on the IP prefix list. Only the routes that pass the filtering of the IP prefix list can be advertised or received.

Creating an IP prefix list before it is referenced is recommended. By default, nonexistent IP prefix lists cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list is referenced using the current command, all routes are advertised to the specified peer, or all routes advertised by the specified peer are accepted.

Example

Configure a route filtering policy based on an IP prefix list.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 ip-prefix list1 import
```

7.8.121 peer keychain (BGP)

Function

The **peer keychain** command configures the keychain authentication for establishing the TCP connection between BGP peers.

The **undo peer keychain** command restores the default setting.

By default, the keychain authentication is not configured for BGP peers.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

peer { *group-name* | *ipv4-address* } **keychain** *keychain-name*

undo peer { *group-name* | *ipv4-address* } **keychain**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a BGP peer.	It is in dotted decimal notation.
keychain-name	Specifies the name of the keychain authentication.	The value is a string of 1 to 47 case-insensitive characters. Except the question mark (?) and space. However, when double quotation marks (") are used around the string, spaces are allowed in the string.

Views

BGP view, BGP-VPN instance IPv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Configuring keychain authentication improves the security of the TCP connection. You must configure keychain authentication specified for TCP-based applications on both BGP peers. Note that encryption algorithms and passwords configured for the keychain authentication on both peers must be the same; otherwise, the TCP connection cannot be set up between BGP peers and BGP messages cannot be transmitted.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Before configuring the BGP keychain authentication, a keychain in accordance with the configured *keychain-name* must be configured first. For keychain configuration details, see Keychain Configuration in the *S300, S500, S2700, S5700, and S6700 V200R023C00 Configuration Guide - Security*.

Precautions

The **peer keychain** command and the **peer password** command are mutually exclusive. SHA256 and HMAC-SHA256 encryption algorithm are recommended in keychain authentication.

Example

Configure the keychain authentication named test for BGP peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 keychain test
```

7.8.122 peer keep-all-routes

Function

The **peer keep-all-routes** command saves all the BGP routing updates from the specified peer or the peer group after the BGP connection is set up, even though those routes do not pass the configured ingress policy.

The **undo peer keep-all-routes** command restores the default setting.

By default, only the BGP routing updates received from the peers and passing the configured ingress policy are saved.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **keep-all-routes**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **keep-all-routes**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, and BGP-IPv6 unicast address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After changing a BGP import policy, you can reset BGP connections for the new import policy to take effect immediately, interrupting these BGP connections temporarily. If a device's peer does not support route-refresh, the **peer keep-all-routes** command can be used on the device to remain all routing updates received from the peer so that the device can refresh its routing table without closing the connection with the peer.

Precautions

If the switch does not support the route-refresh capability, the **peer keep-all-routes** command needs to be run on the switch and its peer. If the **peer keep-all-routes** command is run on a device for the first time, the sessions between the device and its peers will be reestablished.

If the switch supports the route-refresh capability, running this command does not result in re-establishment of the sessions between the switch and its peers. After the **refresh bgp** command is run on the switch, however, the switch does not refresh its routing table.

NOTE

If the switch supports the route-refresh capability, the **peer keep-all-routes** command does not need to be run on it.

Example

Configure a device to store all BGP routing updates received from its IPv4 peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 keep-all-routes
```

7.8.123 peer label-route-capability (BGP)

Function

The **peer label-route-capability** command enables a BGP device to send labeled routes to a specified peer or peer group.

The **undo peer label-route-capability** command disables this function.

By default, this function is disabled.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

BGP view and BGP-IPv4 unicast address family view:

peer { *group-name* | *ipv4-address* } **label-route-capability** [**check-tunnel-reachable**]

undo peer { *group-name* | *ipv4-address* } **label-route-capability**

BGP-IPv6 unicast address family view:

peer *ipv4-address* **label-route-capability**

undo peer *ipv4-address* **label-route-capability**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
check-tunnel-reachable	Checks tunnel reachability when imported routes are sent as labeled routes.	-

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view

Default Level

2: Configuration level

Usage Guidelines

When the capability of sending labeled routes is enabled or disabled, BGP connection is automatically closed and the capability of the neighbor is re-negotiated.

To configure IPv6 provider edge (6PE), specify the peer as an IPv4 unicast address in the IPv6 unicast address family view.

Tunnel reachability checking can only be used to check tunnels on IPv4 public networks.

- If tunnel reachability checking is enabled, BGP advertises IPv4 unicast routes to peers when routed tunnels are unreachable or advertises labeled routes to peers when routed tunnels are reachable. This eliminates the risk of establishing an MP-EBGP peer relationship between PEs over a faulty label switched path (LSP) because this will cause data forwarding failures.
- If tunnel reachability checking is disabled, BGP advertises labeled routes to peers whether the tunnels for imported routes are reachable or not.

NOTE

To disable tunnel reachability checking, run the **peer** { *group-name* | *ipv4-address* } **label-route-capability** command, not the **undo peer label-route-capability** command.

Before you run the **peer label-route-capability** command, the **peer as-number** command should be used to create a peer or peer group.

Example

Enable BGP to send the labeled routes.

```
<HUAWEI> system-view
[HUAWEI] bgp 200
[HUAWEI-bgp] peer 10.2.3.4 as-number 200
[HUAWEI-bgp] ipv6-family
[HUAWEI-bgp-af-ipv6] peer 10.2.3.4 enable
[HUAWEI-bgp-af-ipv6] peer 10.2.3.4 label-route-capability
```

7.8.124 peer listen-net

Function

The **peer listen-net** command configures BGP to listen to BGP connection requests from a specified network segment and establish BGP peer relationships dynamically.

The **undo peer listen-net** command restores the default configuration.

By default, BGP does not listen to BGP connection requests from any network segment.

Format

peer *group-name* **listen-net** *network* { *mask* | *mask-length* }

undo peer *group-name* **listen-net** *network* { *mask* | *mask-length* }

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
network	Specifies the address of the network segment from which BGP listens to BGP connection requests.	The value is in dotted decimal notation.
mask	Specifies the mask of the network segment from which BGP listens to BGP connection requests.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of the network segment from which BGP listens to BGP connection requests.	The value is an integer ranging from 0 to 32.

Views

BGP or BGP-VPN instance IPv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When multiple BGP peers reside on the same network segment, and the number of the peers may change, you can run the **peer listen-net** command to configure BGP to listen to BGP connection requests from the network segment, establish BGP peer relationships dynamically, and add the peers to a peer group. This spares the local device from adding or deleting BGP peer configurations in response to each change in the peer number, which reduces the maintenance workload.

Prerequisites

A peer group has been configured using the **group** *group-name* [**external** | **internal**] command.

Precautions

If you run the command multiple times, BGP listens to BGP connection requests from multiple network segments.

Example

Configure BGP to listen to BGP connection requests from network segment 10.1.1.0/24 and establish BGP peer relationships dynamically.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group test internal
[HUAWEI-bgp] peer test listen-net 10.1.1.0 24
```

7.8.125 peer listen-only

Function

The **peer listen-only** command configures a peer or peer group to only detect connection requests and not to initiatively send connection requests.

The **undo peer listen-only** command restores the default setting.

By default, a peer or peer group detects and sends connection requests.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **listen-only**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **listen-only**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer listen-only** command is used to configure a peer or peer group to only detect connection requests and not to initiatively send connection requests.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

If the **peer listen-only** command is run multiple times, the latest configuration overwrites the previous one.

After the **peer listen-only** command is executed, the local end does not initiate any connection request to a specified peer.

NOTICE

The **peer listen-only** command can be run at only one end of a peer relationship. If this command is run at both ends of a peer relationship, the ends fail to establish a connection.

Example

Configure peer 10.1.1.1 to only detect connection requests from the remote peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.1 as-number 200
[HUAWEI-bgp] peer 10.1.1.1 listen-only
```

7.8.126 peer log-change

Function

The **peer log-change** command enables a BGP device to log the session status and events of a specified peer or a peer group.

The **undo peer log-change** command disables a BGP device to log the session status and events of a specified peer or a peer group.

By default, a BGP device is enabled to log the session status and events of a specified peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **log-change**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **log-change**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, or BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

The **peer log-change** command can be used to enable a device to log the session status and events of a specified peer or peer group, facilitating service management.

Peer relationships have been established using the **peer as-number** command.

Example

Configure a BGP device to log the session status and events of peer 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 log-change
```


7.8.127 peer next-hop-invariable

Function

The **peer next-hop-invariable** command configures PEs in different ASs not to change next hops of routes when the PEs advertise them to their EBGP peers, and configures the PEs to use the next hops of imported IGP routes when the PEs advertise them to their IBGP peers.

The **undo peer next-hop-invariable** command restores the default setting.

By default, when advertising routes to its EBGP peers and imported IGP routes to IBGP peers, a BGP speaker changes the next hop to its interface address.

Format

peer { *ipv4-address* | *group-name* } **next-hop-invariable**

undo peer { *ipv4-address* | *group-name* } **next-hop-invariable**

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-L2VPN address family view, BGP-VPLS address family view, BGP L2VPN-AD address family view, BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **peer next-hop-invariable** command is used on a device, the device does not change the next hop of an imported IGP route when advertising the route to its IBGP peers, and does not change the next hop of a route when advertising it to its EBGP peers.

In the inter-AS VPN option C networking where an RR is used, the **peer next-hop-invariable** command needs to be run to configure the RR not to change the next-hop address of a route when advertising the route to an EBGP peer. This ensures that the remote PE recurses a route to the BGP LSP destined for the local PE during traffic transmission.

Prerequisites

The **peer as-number** command has been used to create a peer or peer group.

Configuration Impact

After the **peer next-hop-invariable** command is used on a device, the device does not change the next hop of an imported IGP route when advertising the route to its IBGP peers, and does not change the next hop of a route when advertising it to its EBGP peers.

Precautions

The **peer next-hop-invariable** command configured on EBGP peers takes effect only for VPNv4 and VPNv6 routes.

If a device needs to advertise routes to its IBGP peer or peer group, the **peer next-hop-invariable** and **peer next-hop-local** commands are mutually exclusive on the device.

Example

Configure a BGP device not to change the next hop of a route when the BGP device advertises it to its EBGP peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family vpnv4
[HUAWEI-bgp-af-vpnv4] peer 10.1.1.2 enable
[HUAWEI-bgp-af-vpnv4] peer 10.1.1.2 next-hop-invariable
```

Use the next hop of an IGP route when advertising the IGP route to IBGP peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 next-hop-invariable
```

7.8.128 peer next-hop-local

Function

The **peer next-hop-local** command configures a BGP device to set its IP address as the next hop of routes when the BGP device advertises routes to an IBGP peer or peer group.

The **undo peer next-hop-local** command restores the default setting.

By default:

- A BGP router sets its interface IP address as the next-hop address of routes when advertising these routes to an EBGP peer.

- A BGP router does not change the next-hop address of non-labeled routes if the routes are from an EBGP peer and are to be sent to an IBGP peer. The router sets its interface IP address as the next-hop address of labeled routes if the routes are from an EBGP peer and are to be sent to an IBGP peer.
- A BGP router does not change the next-hop address of routes if the routes are from an IBGP peer and are to be sent to an IBGP peer.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **next-hop-local**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **next-hop-local**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view, BPG-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address family view, BGP-L2VPN address family view, and BGP L2VPN-AD address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view, and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer next-hop-local** command is usually run on an autonomous system boundary router (ASBR). By default, when an ASBR forwards a route learned from an EBGP peer to its IBGP peers, the ASBR does not change the next hop of the route. The next-hop address of a route advertised by an EBGP peer is the address of the EBGP peer. After being forwarded to the IBGP peers, the route cannot become an active route because of the unreachable next hop. The **peer next-hop-local** command needs to be run to configure the ASBR to change the next hop of the route to its IP address when the ASBR advertises the route to an IBGP peer. As an IGP runs within the AS, the next hop of the route is reachable. As such, the route received by the IBGP peer is active.

Prerequisites

The **peer as-number** command has been used to create a peer or peer group.

Precautions

The **peer next-hop-local** command is applicable to IBGP peers.

Regarding non-labeled BGP routes on a public network, running the **peer next-hop-local** command on a route reflector to change the next hop of BGP routes does not take effect.

Regarding labeled BGP routes, running this command on a route reflector to change the next hop of BGP routes takes effect.

Running this command on a local device to change the next hop of routes imported or aggregated by the local device does not take effect.

If a device needs to advertise routes to its IBGP peer or peer group, the **peer next-hop-local** and **peer next-hop-invariable** commands are mutually exclusive on the device.

The **peer next-hop-local** command can be configured in the L2VPN-AD address family view, but the configuration does not take effect.

This command is not supported on RRs in BGP VPNv6 scenarios. If it is incorrectly configured on an RR in a BGP VPNv6 scenario, traffic will be interrupted.

Example

Configure a BGP device to set its IP address as the next hop of routes when advertising the routes to peer 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 next-hop-local
```

7.8.129 peer out-delay

Function

The **peer out-delay** command configures a delay for sending Update packets.

The **undo peer out-delay** command deletes the delay for sending Update packets.

The default delay is 0, indicating that Update packets are sent without a delay.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **out-delay** *delay-value*

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **out-delay**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	The value is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
delay-value	Specifies the delay for sending Update packets.	The value is an integer that ranges from 0 to 3600, in seconds.

NOTE

- The *ipv4-address* parameter applies to the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-IPv6 unicast address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-VPN instance IPv4 address family view, BGP L2VPN-AD address family view and BGP-EVPN address family view.
- The *ipv6-address* parameter applies to the BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-IPv4 multicast address family view, BGP L2VPN-AD address family view, BGP-EVPN address family view.

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a scenario in which primary and backup routes exist, packets may get lost after traffic switches back to the primary path from the backup path. Use a VPN FRR scenario as an example. PE3 and CE2 connect both to PE1 and PE2. The primary path is PE3 -> PE1 -> CE2, and the backup path is PE3 -> PE2 -> CE2. CE2 uses BGP to communicate with PE1 and PE2. FRR is configured on PE3. If PE1 restarts or the link between PE3 and PE1 is disconnected, traffic switches from the primary path to the backup path. After the primary path recovers, traffic switches back to the primary path. If PE3 completes refreshing forwarding entries before PE1 does so, PE1 may temporarily fail to forward traffic received from PE3, and packet loss may occur. The severity of packet loss is proportional to the number of routes stored on PE1.

To solve this problem, run the **peer out-delay** command on the PE1 to configure a delay for sending Update packets. An appropriate delay ensures that traffic switches back to the primary path after PE1 completes refreshing forwarding entries.

Precautions

If you run the **peer out-delay** command repeatedly, the latest configuration overrides the previous configurations.

If the **peer out-delay** and **peer route-update-interval** commands are both configured, only the **peer out-delay** command takes effect.

If a network has high route convergence requirements, do no use the **peer out-delay** command.

Example

In the BGP view, configure the delay for sending Update packets to 10.1.1.1 as 300s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.1 out-delay 300
```

7.8.130 peer password

Function

The **peer password** command enables a BGP device to implement MD5 authentication for BGP messages exchanged during the establishment of a TCP connection with a peer.

The **undo peer password** command restores the default setting.

By default, a BGP device does not perform MD5 authentication for BGP messages exchanged during the establishment of a TCP connection with a peer.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **password** { **cipher** *cipher-password* | **simple** *simple-password* }

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **password**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
cipher *cipher-password*	Specifies a cipher text password.	The value is a string of case-sensitive characters without spaces. When the value is displayed in plaintext, its length ranges from 1 to 255. When the value is displayed in ciphertext, its length ranges from 20 to 392.
simple *simple-password*	Specifies a plain text password. **NOTICE** If **simple** is selected, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.	The value is a string of 1 to 255 case-sensitive characters, without spaces.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-VPN instance IPv4 address family view, and BGP-VPN instance IPv6 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP uses TCP as the transport layer protocol. To enhance BGP security, MD5 authentication can be implemented for BGP packets exchanged during the establishment of a TCP connection. MD5 authentication sets the MD5 authentication password for the TCP connection, and the authentication is performed by TCP.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

After the **peer password** command is run, if the MD5 authentication fails, no TCP connection is established.

MD5 authentication and keychain authentication are mutually exclusive on a peer.

After the **peer password** command is run on a device to enable MD5 authentication, the device will re-establish the peer relationship with its peer.

Example

Configure authentication for the TCP connection between a device and peer 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 password cipher test
```

7.8.131 peer preferred-value

Function

The **peer preferred-value** command sets a preferred value for the routes that a BGP device learns from its peer.

The **undo peer preferred-value** command restores the default preferred value for the routes that a BGP device learns from its peer.

By default, the preferred value of a route learned from a BGP peer is 0.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **preferred-value** *value*

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **preferred-value**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
value	Specifies the preferred value of the routes that a BGP device learns from its peer.	The value is an integer that ranges from 0 to 65535.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP L2VPN-AD address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the **peer preferred-value** command is run, all the routes learned from the specified peer have the preferred value. If there are multiple routes to the same address prefix, the route with the highest preferred value is preferred.

Prerequisites

A BGP peer has been configured. If the **peer preferred-value** command is used but no BGP peer exists, a message is displayed, indicating that the peer does not exist.

Precautions

If a preferred value is set for the routes that a BGP device learns from a peer group, all members of the peer group inherit the configuration.

Example

Set the preferred value to 50 for the routes that a BGP device learns from a specified peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 preferred-value 50
```

7.8.132 peer public-as-only

Function

The **peer public-as-only** command configures the AS-Path attribute in a BGP Update packet not to carry the private AS number. Only the public AS number is contained in the update packets.

The **undo peer public-as-only** command restores the default setting.

By default, the AS-Path attribute in a BGP Update packet is allowed to carry private AS numbers.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **public-as-only** [**force**]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **public-as-only**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.

Parameter	Description	Value
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
force	In the following two cases, BGP does not delete the private AS number after the command is used: • The AS_Path of a route contains the AS number of the peer. In this case, deleting the private AS numbers may lead to a routing loop. • The AS_Path list contains both public network AS numbers and private AS numbers, indicating that the route has passed through the public network. Deleting the private AS numbers may lead to a forwarding error. To enable the device to delete the private AS numbers from the AS_Path attribute before sending update packets even in the preceding scenarios, specify **force** in the command.	-

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP L2VPN-AD address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In general, an AS number ranges from 1 to 4294967295. The public AS number ranges from 1 to 64511, and 65536 (1.0 in the format of x.y) to 4294967295 (65535.65535 in the format of x.y), and the private AS number ranges from 64512 to 65534. 65535 is used as the reserved AS number in certain circumstances.

Public AS numbers can be used on the Internet. Private AS numbers cannot be advertised to the Internet, and they are used only within ASs. If private AS numbers are advertised to the Internet, a routing loop may occur. After this command is configured, if the AS path attribute contains only private AS numbers, BGP deletes the private AS numbers and then advertises these update routes.

BGP does not delete private AS numbers in either of the following scenarios if the **peer public-as-only** command is run, without any parameter following **public-as-only** specified:

- The AS_Path attribute of a route carries the AS number of the remote peer. In this case, deleting private AS numbers may lead to a routing loop.
- The AS_Path attribute carries both public and private AS numbers, which indicates that the route has passed through the public network. In this case, deleting private AS numbers may lead to a traffic forwarding error.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Example

Configure a device to remove all private AS numbers from the AS_Path attribute when sending a BGP Update packet to its peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 public-as-only
```

7.8.133 peer reflect-client

Function

The **peer reflect-client** command configures the local device as the route reflector and the peer or peer group as the client of the route reflector.

The **undo peer reflect-client** command restores the default setting.

By default, the route reflector and its client are not configured.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **reflect-client**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **reflect-client**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view and BGP-EVPN address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view, BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Full-meshed connections need to be established between IBGP peers to ensure the connectivity between the IBGP peers. If there are n switches in an AS, n*(n-1)/2 IBGP connections need to be established. A large number of IBGP peer use a lot of network and CPU resources. An RR can be used to solve the problem.

In an AS, one switch functions as an RR and other switches function as clients. The clients establish IBGP connections with the RR. The RR and its clients form a cluster. The RR reflects routes among the clients, and therefore the clients do not need to establish any IBGP connection. Assume that an AS has n devices. If one of

the devices functions as a RR, and other devices function as clients, the number of IBGP connections to be established is n-1. This means that network and CPU resources are greatly reduced.

An RR is easy to configure, because it needs to be configured only on the device that functions as a reflector and clients do not need to know that they are clients.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

The device where the **peer reflect-client** command is run serves as the RR and a specified peer or peer group serves as the client of the RR.

The **peer reflect-client** command can be only used between IBGP peers or IBGP peer groups.

reflect-client configured in an address family is valid in this family address and cannot be inherited by other address families. Configuring **reflect-client** in a specified address family is recommended.

Example

Configure a peer as a client of an RR.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 reflect-client
```

7.8.134 peer route-limit

Function

The **peer route-limit** command sets the maximum number of routes that can be received from a peer.

The **undo peer route-limit** command restores the default setting.

By default, there is no limit on the number of routes that can be received from a peer.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **route-limit** *limit* [*percentage*] [**alert-only** | **idle-forever** | **idle-timeout** *minutes*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **route-limit**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
limit	Specifies the maximum number of routes that can be received from a peer.	-
percentage	Specifies the percentage of received routes when the switch starts to generate alarms.	The value is an integer that ranges from 1 to 100. The default value is 75.
alert-only	Indicates that if the number of received routes exceeds the limit, an alarm will be generated and no additional routes will be accepted.	-
idle-forever	Indicates that after the number of routes exceeds the limit, no connection is established automatically until the **reset bgp** command is run.	-
idle-timeout *minutes*	Specifies the value of the timeout timer. The connection, which is closed because the number of routes exceeds the threshold, is automatically reestablished after the timeout timer expires. Before the timer expires, the **reset bgp** command can be used to re-establish a connection.	The value is an integer that ranges from 1 to 1200, in minutes.

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view,BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP L2VPN-AD address family, BGP-L2VPN address family view, BGP-VPLS address family view, BGP-IPv6 unicast address family view, and BGP-VPNv6 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP L2VPN-AD address family view, BGP-L2VPN address family view, BGP-VPLS address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **peer route-limit** command is used to set the maximum number of routes that a BGP switch is allowed to receive from its peer. This provides a mechanism for controlling the routes received from peers in addition to distribution lists, filtering lists, and route mappings.

If a peer relationship between two devices is in the Established state, the following situations occur:

- If the number of routes received by the switch exceeds the upper limit and the **peer route-limit** command is used for the first time, the switch and its peer re-establish the peer relationship, regardless of whether **alert-only** is set.
- If the upper limit set on the switch is increased to be greater than the number of received routes, the switch sends Refresh packets to receive routes again. If the switch does not support the route-refresh capability, the switch needs to re-establish the connection with its peer.
- If the upper limit set on the switch is reduced but is still greater than the number of received routes, only configuration parameters need to be modified.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

If the **peer route-limit** command is run for a peer group, the peers of the peer group inherit the configuration.

If the peer relationship is in the Idle state because the number of received routes exceeds the upper limit and **idle-forever** or **idle-timeout** is set, the **reset bgp** command can be used to re-establish the peer relationship.

Assume that none of **alert-only**, **idle-forever**, and **idle-timeout** is configured. If the number of routes exceeds the upper limit, an alarm is generated and recorded in the log. Then, the peer relationship is disconnected. The devices try to re-establish the peer relationship after 30 seconds.

Example

Set the maximum number of routes that can be received from a peer to 5000.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 route-limit 5000
```

7.8.135 peer route-policy

Function

The **peer route-policy** command specifies a route-policy for filtering routes received from a peer or peer group or routes to be advertised to a peer or peer group.

The **undo peer route-policy** command restores the default setting.

By default, no route-policy is configured for filtering routes received from a peer or peer group or routes to be advertised to a peer or peer group.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **route-policy** *route-policy-name* { **import** | **export** }

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **route-policy** *route-policy-name* { **import** | **export** }

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
route-policy-name	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
import	Applies a route-policy to routes to be imported from a peer or peer group.	-
export	Applies a route-policy to routes to be advertised to a peer or peer group.	-

NOTE

- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPNv6 address family view, and BGP L2VPN-AD address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After a route-policy is created, the **peer route-policy** command is used to apply a Route-Policy to a peer or a peer group so that the routes advertised to or received from the peer or peer group can be controlled. To be specific, only the necessary routes are received from or advertised to the peer or peer group. In this manner, route management is implemented, the scale of the routing table is reduced, and fewer network resources are consumed.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

If the **peer route-policy** command specifies a route-policy that does not exist, use the **route-policy** command to create the route-policy.

Precautions

Currently, the **peer route-policy** command cannot be used to apply a route-policy to set preferences for BGP routes.

The **preference** command cannot configure the preference for routes imported using the **network** command or the **import-route** command. The preference for these routes is 0.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the configured priority is set for all routes; if no priority is configured, the default priority is set for the routes.

Example

Apply a route-policy named test-policy to the routes received from a peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer 10.1.1.2 route-policy test-policy import
```

7.8.136 peer route-update-interval

Function

The **peer route-update-interval** command sets the interval at which a device sends routing updates carrying the same prefix to a peer or peer group.

The **undo peer route-update-interval** command restores the default setting.

By default, the interval at which routing updates are sent to IBGP peers is 15s, and the interval at which routing updates are sent to EBGP peers is 30s.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **route-update-interval** *interval*

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **route-update-interval**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
interval	Specifies the minimum interval at which routing updates are sent.	The value is an integer that ranges from 0 to 600, in seconds.

NOTE

- *group-name* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, and BGP-VPN instance IPv6 address family view.
- *ipv4-address* is valid only in the BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP-IPv6 unicast address family view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When routes change, the switch sends routing updates to notify its peers. If a route changes frequently, the **peer route-update-interval** command can be used to adjust the interval at which Update packets are sent for changes of this route. This frees the switch from sending Update packets for every route change.

Prerequisites

If the **peer route-update-interval** command is used but no peer exists, a message is displayed, indicating that the peer does not exist.

Precautions

If a route is withdrawn, the switch immediately sends an Update message to its peers, regardless of the **peer route-update-interval** configuration. If a route is added and the interval between the last route addition time and the current route addition time is greater than the interval configured using the **peer route-update-interval** command, the switch immediately sends an Update message to its peers. If a route is added and the interval between the last route addition time and the current route addition time is less than the interval configured using the **peer route-update-interval** command, the switch sends an Update message to its peers only after the configured interval expires.

Example

Set the interval at which routing updates are sent to a peer to 10s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 route-update-interval 10
```

7.8.137 peer timer

Function

The **peer timer** command sets the Keepalive timer and Hold timer for a peer or peer group.

The **undo peer timer** command restores the default values of the Keepalive timer and Hold timer.

By default, the value of a Keepalive timer is 60s and the value of a Hold timer is 180s.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **timer keepalive** *keepalive-time* **hold** *hold-time* [**min-holdtime** *min-holdtime*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **timer keepalive** *keepalive-time* **hold** *hold-time* [**min-holdtime** *min-holdtime*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **timer keepalive hold** [**min-holdtime**]

Parameters

Parameter	Description	Value
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
keepalive-time	Specifies the Keepalive period.	The value is an integer that ranges from 0 to 21845, in seconds.
hold-time	Specifies the holdtime.	The value is an integer that can be 0, or ranges from 3 to 65535, in seconds. **NOTE** Setting the hold interval of a BGP peer or peer group to be longer than 20s is recommended. If the hold interval of a BGP peer or peer group is shorter than 20s, the session may be closed.
min-holdtime *min-holdtime*	Specifies the minimum Holdtime configured on the local device. **NOTE** The value of *min-holdtime* configured cannot exceed the value of *hold-time*.	The value is an integer that ranges from 20 to 65535, in seconds.

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After establishing a BGP connection, two peers send Keepalive messages periodically to each other to detect the status of the BGP connection. If a device receives no Keepalive message or any other types of packets from its peer within the holdtime, the device considers the BGP connection closed, and it closes the BGP connection.

When the peers set up a connection, the values of *keepalive-time* and *hold-time* are determined by negotiations between the peers. Each of the two peers sends the other an Open packet containing *hold-time*. The smaller of the *hold-time* values contained in the Open packets from both peers is used. The smaller of the locally configured *keepalive-time* value and one third of the negotiated *hold-time* value is used as the actual *keepalive-time* value.

The **peer timer** command is used to set the Keepalive period and holdtime.

- If short Keepalive period and holdtime are set, BGP can detect a link fault quickly and implement link switching. The number of Keepalive messages on the network, however, is increased. This increases device loads and consumption of network bandwidth resources.
- If long Keepalive period and holdtime are set, the number of Keepalive messages on the network is reduced. This reduces device loads. If, however, the Keepalive period is too long, BGP is unable to detect link status changes in a timely manner, causing many packets to be lost.

If the local device establishes BGP peer relationships with many devices, it needs to process huge BGP messages. If *hold-time* negotiated among BGP peers is small, the timer may expire before the local device processes the Keepalive messages sent from other BGP peers. The peer relationships are then interrupted, and routes flap. To solve the preceding problem, you can configure an appropriate value for **min-holdtime** *min-holdtime* based on the CPU processing capability of the local device.

If the value of *min-holdtime* is changed, but the values of *keepalive-time* and *hold-time* negotiated between two BGP peers remain unchanged, the established peer relationship is not affected. Only when the local device attempts to re-establish a relationship with a remote device, the value of *min-holdtime* configured on the local device takes effect. The local device compares *min-holdtime* with *hold-time* sent from the remote device. If the value of *min-holdtime* exceeds that of *hold-time*, *hold-time* negotiation fails, and the peer relationship fails to be established.

NOTE

If *min-holdtime* is configured on the local device, and the value of *hold-time* sent from the remote device is 0, *hold-time* negotiation between the two devices succeeds. The negotiated value of *hold-time* is 0, and the peer relationship is established. The value 0 of *hold-time* indicates that the peer relationship never expires.

Prerequisites

Peer relationships have been established using the **peer as-number** command.

Precautions

NOTICE

If the value of a timer changes, the BGP peer relationship between devices will be disconnected. This is because the devices need to re-negotiate the values of *keepalive-time* and *hold-time*. Therefore, exercise caution before changing the value of a timer.

The Keepalive period must be at least three times of the holdtime.

When setting the values of *keepalive-time* and *hold-time*, note the following points:

- The values of *keepalive-time* and *hold-time* cannot both be set to 0. This renders the BGP timers invalid. BGP is unable to detect link faults using the timers.
- The *hold-time* value cannot be significantly greater than the *keepalive-time* value. A setting of **timer keepalive 1 hold 65535**, for example, would be improper. If the holdtime is too long, link faults cannot be detected in a timely manner.

The Keepalive period and Holdtime can be configured globally, or on a particular peer or peer group. The Keepalive period and Holdtime configured on a specific peer or peer group take precedence over the global Keepalive period and Holdtime. Using this command can still change the Keepalive period and Holdtime configured on a peer or peer group, although they were globally configured through the **timer** command.

Example

Set the Keepalive timer and Hold timer for peer 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 timer keepalive 10 hold 30
```

7.8.138 peer tnl-policy

Function

The **peer tnl-policy** command applies the tunnel policy to the specified IPv4 peer.

The **undo peer tnl-policy** command removes the tunnel policy applied to the peer.

By default, no tunnel policy is applied to the peer.

NOTE

Only the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H support this command.

Format

peer *ipv4-address* **tnl-policy** *tnl-policy-name*

undo peer *ipv4-address* **tnl-policy**

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IPv4 address of the peer.	-
tnl-policy-name	Specifies the name of the tunnel policy.	The name is a string of 1 to 39 characters. It is case-sensitive.

Views

BGP-IPv6 unicast address family view

Default Level

2: Configuration level

Usage Guidelines

When you configure the 6PE to support the tunnel, you need to configure the tunnel policy and apply the tunnel policy by using the command.

Before you configure the **peer tnl-policy** command, the **peer as-number** command should be used to create a peer or peer group.

Example

Apply the tunnel policy to the specified peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 100
[HUAWEI-bgp] ipv6-family unicast
[HUAWEI-bgp-af-ipv6] peer 10.1.1.2 enable
[HUAWEI-bgp-af-ipv6] peer 10.1.1.2 tnl-policy policy-a
```

7.8.139 peer timer connect-retry

Function

The **peer timer connect-retry** command sets a ConnectRetry interval for a peer or peer group.

The **undo peer timer connect-retry** command restores the default setting.

By default, the ConnectRetry interval is 32s.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **timer connect-retry** *connect-retry-time*

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **timer connect-retry**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies an IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies an IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
connect-retry-time	Specifies a ConnectRetry interval.	The value ranges from 1 to 65535, in seconds.

NOTE

- *ipv4-address* can be set only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view.

Views

BGP view, BGP-VPN instance IPv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When BGP initiates a TCP connection, the ConnectRetry timer is stopped if the TCP connection is established successfully. If the first attempt to establish a TCP connection fails, BGP tries again to establish the TCP connection after the ConnectRetry timer expires. The ConnectRetry interval can be adjusted as needed.

- The ConnectRetry interval can be reduced in order to lessen the time BGP waits to retry establishing a TCP connection after the first attempt fails.
- To suppress route flapping caused by constant peer flapping, the ConnectRetry interval can be increased to accelerate route convergence.

Prerequisites

The **peer as-number** command has been used to create a peer or peer group.

Precautions

A ConnectRetry interval can be configured globally, or on a particular peer or peer group. A ConnectRetry interval configured on a specific peer or peer group takes precedence over a global ConnectRetry interval.

- If both the **peer** { *ipv4-address* | *ipv6-address* } **timer connect-retry** *connect-retry-time* command and the **peer** *group-name* **timer connect-retry** *connect-retry-time* command are run on a device, the configuration of the **peer** { *ipv4-address* | *ipv6-address* }**timer connect-retry** *connect-retry-time* command takes effect, but the configuration of the **peer** *group-name* **timer connect-retry** *connect-retry-time* command does not.
- If both the **peer** { *group-name* | *ipv4-address* | *ipv6-address* } **timer connect-retry** *connect-retry-time* command and the **timer connect-retry** *connect-retry-time* command are run on a device, the configuration of the **peer** { *group-name* | *ipv4-address* | *ipv6-address* } **timer connect-retry** *connect-retry-time* command takes effect, but the configuration of the **timer connect-retry** *connect-retry-time* command does not.

Example

Set the ConnectRetry interval to 60s for peer 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 timer connect-retry 60
```

7.8.140 peer tracking

Function

The **peer tracking** command enables BGP peer tracking.

The **undo peer tracking** The command restores the interval between peer unreachability discovery and connection interruption to the default value (9s).

The **peer tracking disable** command disables BGP peer tracking.

The **undo peer tracking disable** command enables BGP peer tracking. The default interval between peer unreachability discovery and connection interruption is 9s.

By default, BGP peer tracking is enabled, and the interval between peer unreachability discovery and connection interruption is 9s.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **tracking** [**delay** *delay-time*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **tracking**

peer { *group-name* | *ipv4-address* | *ipv6-address* } **tracking disable**

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **tracking disable**

Parameters

Parameter	Description	Value
group-name	Specifies the name of a BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a BGP peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a BGP peer.	The address is a 32-digit hexadecimal number in the X:X:X:X:X:X:X:X format.
delay *delay-time*	Indicates the interval between when BGP detects the peer unreachable and when BGP tears down the corresponding connection.	The value is an integer that ranges from 0 to 65535, in seconds. The default value is 9 seconds.

NOTE

- *ipv4-address* can be set only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* can be set only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a network where BFD is unsuitable to be deployed, you can configure BGP peer tracking on the local device to implement fast network convergence by rapidly detecting the unreachable state of the peer.

A proper value of *delay-time* can ensure network stability when a peer is detected unreachable.

- If *delay-time* is set to 0, BGP immediately tears down the connection between the local device and its peer after the peer is detected unreachable.

- If IGP route flapping occurs and *delay-time* for an IBGP peer is set to 0, the peer relationship between the local device and the peer alternates between Up and Down. Therefore, *delay-time* for an IBGP peer should be set to a value greater than the actual IGP route convergence time.
- When BGP neighbors successfully perform the GR negotiation, the active/standby switchover occurs on the BGP neighbors, to prevent the failure of GR, *delay-time* should be set to a value greater than GR convergence time. If *delay-time* is set to be smaller than the GR convergence time, the connection between the local device and the BGP peer will be torn down, which leads to the failure of GR.

Prerequisites

The **peer as-number** command has been used to create a peer or peer group.

Precautions

IGP is configured with GR, and the BGP neighbor relationship is established based on IGP routes. In such a situation, when a node fails on the network and the master/slave switchover occurs on the node, IGP does not delete the routes from the node, and BGP neighbors cannot sense the fault on the node. Therefore, the BGP peer tracking function does not take effect.

Example

Configure BGP peer tracking.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 tracking delay 20
```

7.8.141 peer valid-ttl-hops

Function

The **peer valid-ttl-hops** command applies the GTSM function on the peer or peer group.

The **undo peer valid-ttl-hops** command cancels the application of the GTSM function on the peer or peer group.

By default, the GTSM function on the peer or peer group is not configured.

Format

peer { *group-name* | *ipv4-address* | *ipv6-address* } **valid-ttl-hops** [*hops*]

undo peer { *group-name* | *ipv4-address* | *ipv6-address* } **valid-ttl-hops**

Parameters

Parameter	Description	Value
group-name	Specifies the name of the peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
ipv4-address	Specifies the IPv4 address of a peer.	It is in dotted decimal notation.
ipv6-address	Specifies the IPv6 address of a peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
hops	Specifies the number of TTL hops to be checked.	The value is an integer that ranges from 1 to 255. The default value is 255. If the value is configured as *hops*, the valid TTL range of the detected packet is [255 - *hops* + 1, 255].

NOTE

- *ipv4-address* is valid only in the BGP view and BGP-VPN instance IPv4 address family view.
- *ipv6-address* is valid only in the BGP view and BGP-VPN instance IPv6 address family view.

Views

BGP view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To protect a device against the attacks of forged BGP packets, you can configure GTSM to check whether the TTL value in the IP packet header is within the specified range.

Prerequisites

Before configuring GTSM for a peer group, run the **peer group** command to add peers to the peer group.

Precautions

When the **undo peer valid-ttl-hops** command is run and no parameter is specified, all the GTSM configurations on a peer or a peer group are deleted.

The configuration in the BGP view is also valid for the extension of MP-BGP. This is because they use the same TCP connection.

The GTSM configurations are symmetrical, that is, GTSM must be enabled on both ends of the BGP connection at the same time.

NOTE

- GTSM and EBGP-MAX-HOP are mutually exclusive because both of them affect the TTL of the sent BGP packet. Therefore, the two functions cannot be enabled on a peer or peer group simultaneously.
- If GTSM is enabled on two directly connected EBGP peers, the fast sensing function is invalid on the interfaces directly connecting the two EBGP peers. This is because BGP considers the EBGP peers indirectly connected when GTSM is enabled on the EBGP peers.

Example

Configure the GTSM function for the peer.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] peer 10.1.1.2 as-number 200
[HUAWEI-bgp] peer 10.1.1.2 valid-ttl-hops 1
```

Configure the GTSM function for the peer group.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group gtsm-group external
[HUAWEI-bgp] peer gtsm-group valid-ttl-hops 1
```

7.8.142 preference (BGP)

Function

The **preference** command configures the protocol preferences for external routes, internal routes, aggregated routes, and crossed routes.

The **undo preference** command restores the default protocol preferences.

By default, the protocol preferences of external routes, internal routes, aggregated routes, and crossed routes are all 255.

Format

preference { *external internal local* | **route-policy** *route-policy-name* }

preference *external internal local* **route-policy** *route-policy-name*

undo preference

Parameters

Parameter	Description	Value
external	Specifies the protocol preference for external routes. An external route is the optimal route learned from an EBGP peer outside the local AS.	The value is an integer that ranges from 1 to 255. The smaller the value is, the higher the preference is.

Parameter	Description	Value
internal	Specifies the protocol preference for internal routes. An internal route is a route learned from an IBGP peer inside the local AS.	The value is an integer that ranges from 1 to 255. The smaller the value is, the higher the preference is.
local	Specifies the protocol preference for aggregated and crossed routes. This parameter takes effect for the following routes: • Manually summarized routes generated using the **aggregate (BGP)** command • Automatically summarized routes generated using the **summary automatic** command • Routes generated through remote route cross • Routes generated through local route cross For details about these routes, see **Precautions**.	The value is an integer that ranges from 1 to 255. The smaller the value is, the higher the preference is.
route-policy *route-policy-name*	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Running the **preference** command to set protocol preferences for BGP routes affects route selection among BGP routes and routes of other routing protocols.

After a route-policy is configured on a device, the device sets preferences only for the routes received from peers, which meet the matching rules. The routes that do not meet the rules use the default preference.

The smaller the preference value, the higher the preference.

Different protocol preferences can be configured for BGP routes in different address family views.

If both *external internal local* and **route-policy** *route-policy-name* are specified in the command, the priority of the routes that match the route-policy is set based on the route-policy, and the priorities of other routes are set based on the *external internal local* configuration.

Prerequisites

Create the route-policy first if the **preference** command uses the **route-policy** to set preferences.

Perform the following steps when the route-policy is used to set preferences:

- Use the **route-policy** command to create the route-policy, and enter the route-policy view.
- Configure the **if-match** clause to set the matching rules for routes. The relationship between the **if-match** clauses in a node of a route-policy is "AND". A route must match all the rules before the action defined by the **apply** clause is taken. If no **if-match** clause is specified, all routes will match the node in the route-policy.
- Use the **apply preference** command to set preferences for routes that pass the filtering.

Precautions

- Currently, the **peer route-policy** command cannot be used to apply a route-policy to set preferences for BGP routes.
- The **preference** command cannot configure the preference for routes imported using the **network** command or the **import-route** command. The preference for these routes is 0.
- Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, the configured priority is set for all routes; if no priority is configured, the default priority is set for the routes.
- In this command, the *local* parameter specifies the protocol preference for aggregated routes and crossed routes. This parameter takes effect for the following routes:
  - Manually aggregated routes generated using the **aggregate (BGP)** command. The **Route type** field for this type of route displays **Aggregated route**. For example:

```
<HUAWEI> display bgp routing-table 10.0.0.0
 BGP local router ID : 192.168.2.4
 Local AS number : 200
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.0.0.0/8:
 Aggregated route
 Route Duration: 04h50m46s
 Direct Out-interface: NULL0
 Original nexthop: 127.0.0.1
 Qos information : 0x0
 AS-path {65001 10 100}, origin incomplete, pref-val 0, valid, local, best, select, active, pre 255
```



```
 Aggregator: AS 200, Aggregator ID 192.168.2.4, Atomic-aggregate
 Advertised to such 3 peers:
    10.1.7.2
    172.16.1.2
    192.168.1.2
...
```

- Automatically aggregated routes generated using the **summary automatic** command. The route type for these routes displays **Summary automatic route**. For example:

```
<HUAWEI> display bgp routing-table 10.0.0.0
 BGP local router ID : 192.168.2.4
 Local AS number : 200
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.0.0.0/8:
 Summary automatic route
 Route Duration: 04h50m46s
 Direct Out-interface: NULL0
 Original nexthop: 127.0.0.1
 Qos information : 0x0
 AS-path {65001 10 100}, origin incomplete, pref-val 0, valid, local, best, select, active, pre 255
 Aggregator: AS 200, Aggregator ID 192.168.2.4, Atomic-aggregate
 Advertised to such 3 peers:
    10.1.7.2
    172.16.1.2
    192.168.1.2
...
```

- Routes generated through remote route cross. The route type for these routes displays **Remote-Cross route**. For example:

```
[HUAWEI-diagnose] display bgp vpnv4 vpn-instance vrf routing-table 172.16.17.161

BGP local router ID : 10.17.0.17
Local AS number : 100

VPN-Instance vrf, Router ID 10.17.0.17:
Paths:   1 available, 1 best, 1 select
BGP routing table entry information of 172.16.17.161/32:
Remote-Cross route
Label information (Received/Applied): 155665/NULL
From: 10.216.0.1 (10.17.0.9)
Route Duration: 3d04h26m28s
 Relay Tunnel Out-Interface: Vlanif100
Relay token: 0x24e7
Relay Tunnel Key: 7
Original nexthop: 10.216.0.1
Qos information : 0x0
Ext-Community:RT <10 : 10>
AS-path 100, origin incomplete, localpref 150, pref-val 1, valid, internal, best, select, active, pre
255, IGP cost 1
...
```

- Routes generated through local route cross. The route type for these routes displays **Local-Cross route**. For example:

```
[HUAWEI-diagnose] display bgp vpnv4 vpn-instance vrf1 routing-table 10.2.2.0

 BGP local router ID :
10.1.1.1
 Local AS number :
100

 VPN-Instance vrf1, Router ID
10.1.1.1:
 Paths:   1 available, 1 best, 1
select
 BGP routing table entry information of
10.2.2.0/24:
 Local-Cross route(via VPN-Instance
vrf2)
```



```
 Route Duration:
00h26m00s
 Direct Out-interface:
Vlanif40
 Original nexthop:
10.2.2.2
 Qos information :
0x0
 Ext-Community:RT <100 :
1>
 AS-path Nil, origin incomplete, MED 0, pref-val 0, valid, local, best, select, active, pre 255
...
```

Example

Set the protocol preference to 2 for external routes, 2 for internal routes, and 20 for aggregated routes and crossed routes.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] preference 2 2 20
```

7.8.143 reflect between-clients

Function

The **reflect between-clients** command enables route reflection among clients.

The **undo reflect between-clients** command disables route reflection among clients.

By default, route reflection among clients is enabled.

Format

reflect between-clients

undo reflect between-clients

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view, BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

On some networks, if the clients of an RR establish full-mesh connections with each other, they can directly exchange routing information. Route reflection among clients is unnecessary. The **undo reflect between-clients** command can be used to prohibit the clients from reflecting routes to each other to reduce costs.

Prerequisites

An RR has been configured.

Precautions

The **reflect between-clients** command is run only on RRs.

Example

Disable route reflection among fully-meshed clients.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] group rr-client internal
[HUAWEI-bgp] peer 10.1.2.1 group rr-client
[HUAWEI-bgp] peer 10.1.3.1 group rr-client
[HUAWEI-bgp] peer 10.1.4.1 group rr-client
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] peer rr-client reflect-client
[HUAWEI-bgp-af-ipv4] undo reflect between-clients
```

7.8.144 reflect change-path-attribute

Function

The **reflect change-path-attribute** command enables a route-reflector (RR) to modify the route attributes of BGP routes using the export policy.

The **undo reflect change-path-attribute** command disables an RR from modifying the route attributes of BGP routes using the export policy.

By default, an RR is disabled from modifying route attributes of BGP routes using the export policy.

Format

reflect change-path-attribute

undo reflect change-path-attribute

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN

instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Using an RR to modify route attributes may cause route loops. According to RFC 4456, you cannot enable an RR to modify the route attributes using the export policy. So an RR is disabled from modifying route attributes using the export policy by default.

To re-plan the network traffic, run the **reflect change-path-attribute** command to enable the RR to modify the route attributes using the export policy. After the command is run, the following configurations take effect.

- Run the **apply as-path** command to modify the AS-Path attributes of BGP routes.
- Run the **apply comm-filter delete** command to delete all community attributes from BGP routes.
- Run the **apply community** command to modify the community attributes of BGP routes.
- Run the **apply cost** command to modify the cost of BGP routes, that is to modify its MED.
- Run the **apply ip-address next-hop** command to modify the next hop of BGP routes.
- Run the **apply ipv6 next-hop** command to modify the next hop of BGP4+ routes.
- Run the **apply local-preference** command to modify the local preference of BGP routes.
- Run the **apply origin** command to modify the Origin attributes of BGP routes.
- Run the **apply extcommunity** command to modify the extended community attributes of BGP routes.

After the **undo reflect change-path-attribute** command is used, the previous configurations on the RR do not take effect.

Precautions

Export policies on the RR do not take effect before the **reflect change-path-attribute** command is run. After the **reflect change-path-attribute** command is run, these configurations may take effect and affect BGP route selection. Exercise caution when using this command.

For example, peer relationships are established between Switch A(10.1.1.1) and Switch B(10.1.1.2), and Switch A functions as an RR. If two configurations are on SwitchA:

1. [HUAWEI] **bgp 65001**


```
[HUAWEI-bgp] peer 10.1.1.2 next-hop-local
```

2.
```
[HUAWEI] route-policy aa permit node 10
[HUAWEI-route-policy] apply ip-address next-hop 10.3.3.3
[HUAWEI-route-policy] quit
[HUAWEI] bgp 65001
[HUAWEI-bgp] peer 10.1.1.2 route-policy aa export
```

Then:

- Before the **reflect change-path-attribute** command is run, the former configuration takes effect but the latter does not. After A receives routes information from its IBGP peers, it reflects the information to Switch B, and changes the next hop to 10.1.1.1.
- After the **reflect change-path-attribute** command is run, the latter configuration takes effect but the former does not. After A receives routes information from its IBGP peers, it reflects the information to Switch B, and changes the next hop to 10.3.3.3.

NOTE

After you enable the **reflect change-path-attribute** command on the RR, the **peer route-policy export** command takes precedence over the **peer next-hop-invariable** and **peer next-hop-local**.

Example

Enable the RR to modify the route attributes of the BGP routes by using the export policy.

```
<HUAWEI> system-view
[HUAWEI] bgp 65001
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] reflect change-path-attribute
```

7.8.145 reflector cluster-id

Function

The **reflector cluster-id** command sets a cluster ID for an RR.

The **undo reflector cluster-id** command restores the default setting.

By default, each RR uses its router ID as the cluster ID.

Format

reflector cluster-id *cluster-id*

undo reflector cluster-id

Parameters

Parameter	Description	Value
cluster-id	Specifies the cluster ID of an RR.	The value can be an integer that ranges from 1 to 4294967295 or in the format of an IPv4 address.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-VPNv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv6 address family view, BGP-VPLS address family view, BGP-L2VPN address family view, BGP L2VPN-AD address family view, BGP-EVPN address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Sometimes, more than one RR needs to be configured in a cluster to improve network reliability and prevent single-point failures. If a cluster has more than one RR, the **reflector cluster-id** command needs to be used to set the same cluster ID for the RRs. This helps to identify the cluster and avoid routing loops.

Configuring an RR allows IBGP peers to advertise routes learned in the local AS to each other. The Cluster_List attribute is introduced to avoid loops within an AS. The Cluster_List is composed of a series of Cluster_IDs. It records all the RRs through which a route passes.

Precautions

If the **reflector cluster-id** command is run several times, the latest configuration overrides the previous one.

The **reflector cluster-id** command is run only on RRs.

To enable clients to receive routes reflected by RRs, ensure that the cluster ID of the RRs is different from the router ID of any client. If the cluster ID of the RRs is the same as the router ID of a client, the client will discard received routes.

Example

Configure the cluster ID to 50 for the switch, which is an RR in a cluster.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] reflector cluster-id 50
```


7.8.146 refresh bgp

Function

The **refresh bgp** command softly resets a BGP connection.

Format

refresh bgp [**vpn-instance** *vpn-instance-name* **ipv4-family** | **vpnv4**] { **all** | *ipv4-address* | **group** *group-name* | **external** | **internal** } { **export** | **import** }

refresh bgp ipv6 { **all** | **group** *group-name* | *ipv4-address* | *ipv6-address* | **external** | **internal** } { **export** | **import** } (The *ipv4-address* is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S)

refresh bgp vpnv6 { **all** | *ipv4-address* | **group** *group-name* | **external** | **internal** } { **export** | **import** }

refresh bgp vpn-instance *vpn-instance-name* **ipv6-family** { **all** | *ipv6-address* | **group** *group-name* | **external** | **internal** } { **export** | **import** } (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

refresh bgp l2vpn-ad { *ipv4-address* | **all** | **group** *group-name* | **internal** | **external** } { **import** | **export** }

refresh bgp { **mdt** | **mvpn** } { *ipv4-address* | **all** | **group** *group-name* | **internal** | **external** } { **import** | **export** }

NOTE

The **mdt, mvpn, l2vpn-ad**, **vpnv4**, and **vpnv6** parameter are only supported by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H.

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name* **ipv4-family**	Softly resets the connection of a specified VPN instance enabled with an IPv4 address family.	The value must be an existing VPN instance name.
vpnv4	Softly resets the BGP connections in a VPNv4 address family.	-
all	Softly resets all the BGP IPv4 connections.	-
ipv4-address	Specifies the IPv4 address of a BGP peer.	It is in dotted decimal notation.

Parameter	Description	Value
group *group-name*	Specifies the name of a BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
external	Softly resets EBGP connections.	-
internal	Softly resets IBGP connections.	-
export	Triggers outbound soft resetting.	-
import	Triggers inbound soft resetting.	-
ipv6	Softly resets BGP4+ connections.	-
ipv6-address	Specifies the IPv6 address of a BGP4+ peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
vpnv6	Softly resets the BGP connections in a VPNv6 address family.	-
vpn-instance *vpn-instance-name* **ipv6-family**	Resets the connection of a specified VPN instance enabled with an IPv6 address family.	The value must be an existing VPN instance name.
l2vpn-ad	Softly resets the BGP connections related to L2VPN-AD.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If a device's peer supports route-refresh, the **refresh bgp** command can be used on the device to softly reset the BGP connection with the peer. BGP soft resetting can be used to refresh the BGP routing table and apply new routing policies, without closing any BGP connection.

Prerequisites

Configuring BGP soft resetting requires that the peers support the route-refresh capability.

Precautions

After the **refresh bgp** command is run on a device configured with the **peer keep-all-routes** command, the device does not refresh its routing table.

Example

Perform inbound soft resetting for all BGP connections to make new configurations take effect.

```
<HUAWEI> refresh bgp all import
```

7.8.147 refresh bgp multicast

Function

The **refresh bgp multicast** command softly resets an MBGP connection. MBGP soft resetting can be used to refresh the MBGP routing table and apply new routing policies, without closing any MBGP connection.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

refresh bgp multicast { *peer-address* | **all** | **group** *group-name* } { **import** | **export** }

Parameters

Parameter	Description	Value
peer-address	Softly resets the MBGP connections with a specified peer.	The peer address is in dotted decimal notation.
all	Softly resets all MBGP connections.	-

Parameter	Description	Value
group-name	Specifies the name of an MBGP peer group. If the parameter is set, the MBGP connections between a device and the members of the specified MBGP peer group are softly reset.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
import	Triggers inbound soft resetting.	-
export	Triggers outbound soft resetting.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

MBGP soft resetting requires that all MBGP devices on a network support the route-refresh capability. The **reset bgp** command can be used on a device that does not support the route-refresh capability to reset the connections between the device and its peer and enable the device to refresh its routing table.

Example

Perform inbound soft resetting for all the MBGP connections to make new configurations take effect.

```
<HUAWEI> refresh bgp multicast all import
```

7.8.148 refresh bgp multicast external

Function

The **refresh bgp multicast external** command softly resets the connections between multicast EBGP peers.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

refresh bgp multicast external { **import** | **export** }

Parameters

Parameter	Description	Value
import	Triggers inbound soft resetting.	-
export	Triggers outbound soft resetting.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can run the **refresh bgp multicast external** command to softly reset the connections between multicast EBGP peers.

Example

Trigger inbound soft resetting for all multicast EBGP connections.

```
<HUAWEI> refresh bgp multicast external import
```

7.8.149 refresh bgp multicast internal

Function

The **refresh bgp multicast internal** command softly resets the connections between multicast IBGP peers.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

refresh bgp multicast internal { **import** | **export** }

Parameters

Parameter	Description	Value
import	Triggers inbound soft resetting.	-
export	Triggers outbound soft resetting.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can run the **refresh bgp multicast internal** command to softly reset the connections between multicast IBGP peers.

Example

Trigger inbound soft resetting for all multicast IBGP connections.

```
<HUAWEI> refresh bgp multicast internal import
```

7.8.150 reset bgp

Function

The **reset bgp** command resets specified BGP connections.

Format

reset bgp [**vpn-instance** *vpn-instance-name* **ipv4-family** | **vpnv4**] { **all** | *as-number-plain* | *as-number-dot* | *ipv4-address* | **group** *group-name* | **external** | **internal** } [**graceful**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

reset bgp ipv4 all [**graceful**]

reset bgp ipv6 { **all** | *as-number-plain* | *as-number-dot* | **group** *group-name* | *ipv6-address* | *ipv4-address* | **external** | **internal** } [**graceful**] (The *ipv4-address* is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.)

reset bgp vpls { **all** | *as-number-plain* | *as-number-dot* | *ipv4-address* | **external** | **internal** } (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

reset bgp vpnv6 { **all** | *as-number-plain* | *as-number-dot* | *ipv4-address* | **group** *group-name* | **external** | **internal** } [**graceful**] (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

reset bgp vpn-instance *vpn-instance-name* **ipv6-family** { **all** | *as-number-plain* | *as-number-dot* | *ipv6-address* | **group** *group-name* | **external** } [**graceful**] (supported only by the S5731-H, S5731S-H, S5731-S, S5731S-S, S6730-S, S6730S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name* **ipv4-family**	Resets the connection of a specified VPN instance enabled with an IPv4 address family.	The value must be an existing VPN instance name.
vpnv4	Resets BGP connections associated with VPNv4.	-
all	Resets all BGP connections.	-
as-number-plain	Specifies the number of the AS, in integer format.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of the AS, in dotted notation.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.
ipv4-address	Resets the BGP connection with a specified peer.	It is in dotted decimal notation.
group *group-name*	Resets the BGP connection with a specified peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
external	Resets all EBGP connections.	-
internal	Resets all IBGP connections.	-
ipv4	Resets BGP IPv4 connections.	-
ipv6	Resets BGP IPv6 connections.	-
ipv6-address	Resets the TCP connection with a specified BGP4+ peer (all the routes learned by using the connection are deleted).	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Parameter	Description	Value
vpls	Resets BGP connections associated with VPLS.	-
vpnv6	Resets BGP connections associated with VPNv6.	-
vpn-instance *vpn-instance-name* **ipv6-family**	Resets the connection of a specified VPN instance enabled with an IPv6 address family.	The value must be an existing VPN instance name.
graceful	Resets BGP connections in GR mode.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **reset bgp** command is used to make new BGP configurations take effect.

If a BGP routing policy is configured on the device that does not support Route-Refresh, the **reset bgp** command can be used to make the new routing policy to take effect.

To reset a BGP connection in GR mode, run the **reset bgp** command with the **graceful** parameter specified and run the **graceful-restart peer-reset** command. If the **graceful** parameter is not specified in the **reset bgp** command or if the **graceful-restart peer-reset** command is not run, the GR reset mode does not take effect, so that routing entries will be deleted for existing sessions, interrupting services. The services will be restored after the BGP peer relationship is reestablished.

The **reset bgp ipv4 all** command resets all public-network BGP IPv4 connections.

Precautions

After the **reset bgp** command is run on a device, the TCP connection established by the BGP device and the corresponding peer is reestablished. Exercise caution when running this command.

Example

Reset all BGP connections.


```
<HUAWEI> reset bgp all
```

Reset BGP connections with the peer 2001:DB8:1::1.

```
<HUAWEI> reset bgp ipv6 2001:DB8:1::1
```

Reset all BGP connections with VPNv6.

```
<HUAWEI> reset bgp vpnv6 all
```

Reset BGP sessions with BGP peers within a specified 2-byte AS number.

```
<HUAWEI> reset bgp vpnv6 100
```

Reset BGP sessions with BGP peers within a specified 4-byte AS number.

```
<HUAWEI> reset bgp vpnv4 200.300
```

7.8.151 reset bgp dampening

Function

The **reset bgp dampening** command clears BGP route dampening information and releases the suppressed routes.

Format

reset bgp [**vpn-instance** *vpn-instance-name* **ipv4-family**] **dampening** [*ipv4-address* [*mask* | *mask-length*]]

reset bgp { **ipv6** | **vpn-instance** *vpn-instance-name* **ipv6-family** } **dampening** [*ipv6-address prefix-length*]

reset bgp vpnv4 dampening [*ipv4-address* [*mask* | *mask-length*]] (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

NOTE

The **vpn-instance** *vpn-instance-name* **ipv6-family** is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Clears the route dampening information of the specified VPN instance.	The value must be an existing VPN instance name.
ipv4-family	Indicates an IPv4 unicast address family.	-
ipv4-address	Specifies an IPv4 network address.	It is in dotted decimal notation.

Parameter	Description	Value
mask	Specifies the network mask in dotted decimal notation. If neither of the mask and mask length is specified, the address is considered as a classful address.	It is in dotted decimal notation.
mask-length	Specifies the network mask length. If neither of the mask and mask length is specified, the address is considered as a classful address.	The value is an integer that ranges from 0 to 32.
ipv6	Clears IPv6 route dampening information and releases the suppressed routes.	-
ipv6-family	Indicates an IPv6 unicast address family.	-
ipv6-address	Specifies the IPv6 network address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the length of an IPv6 prefix in decimal notation. It specifies the number of bits in the network address.	The value is an integer that ranges from 0 to 128.
vpnv4	Clears the route dampening information of the BGP VPNv4 routes and releases the suppressed routes.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Route dampening is enabled to solve the problem of route instability. In most situations, BGP is applied to complex networks where routes change frequently. Route dampening is then used to suppress instable routes.

The **reset bgp dampening** command is used to clear the dampening information about specified routes on the public network and release specified suppressed routes. If no IP address is specified in the command, the dampening information about all the routes on the public network is cleared and all suppressed routes are released.

Prerequisites

You can use **display bgp routing-table dampened** command to view the information about suppressed routes.

Precautions

After the **reset bgp dampening** command is run, the suppressed routes are released. If the status of some routes still changes frequently, route flapping may occur. Routing flapping consumes a large number of bandwidth and CPU resources.

When *ipv6-address prefix-length* is not specified, after you run the **reset bgp ipv6 dampening** command, IPv6 route dampening information in the whole BGP routing table is cleared.

Example

Clear the dampening information about routes to network segment 10.1.0.0 and release suppressed routes.

```
<HUAWEI> reset bgp dampening 10.1.0.0 255.255.0.0
```

Clear the dampening information about the IPv6 routes to network segment FC00:0:0:1:: and release suppressed routes.

```
<HUAWEI> reset bgp ipv6 dampening fc00:0:0:1:: 64
```

Clear the dampening information about the routes of IPv6 VPN instance named vpn1.

```
<HUAWEI> reset bgp vpn-instance vpn1 ipv6-family dampening
```

7.8.152 reset bgp flap-info

Function

The **reset bgp flap-info** command clears route flapping statistics.

Format

reset bgp [**vpn-instance** *vpn-instance-name* **ipv4-family**] **flap-info** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *network-address* [*mask* | *mask-length*] | **regexp** *as-path-regexp*]

reset bgp { **ipv6** | **vpn-instance** *vpn-instance-name* **ipv6-family** } **flap-info** [**as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *network-ipv6-address prefix-length* | **regexp** *as-path-regexp*]

reset bgp [**vpn-instance** *vpn-instance-name* **ipv4-family**] *ipv4-address* **flap-info**

reset bgp { **ipv6** | **vpn-instance** *vpn-instance-name* **ipv6-family** } *ipv6-address* **flap-info**

reset bgp vpnv4 flap-info [**regexp** *as-path-regexp* | **as-path-filter** { *as-path-filter-number* | *as-path-filter-name* } | *ipv4-address* [*mask* | *mask-length*]] (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

reset bgp vpnv4 *ipv4-address* **flap-info** (supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6730-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H)

NOTE

The **vpn-instance** *vpn-instance-name* **ipv6-family** is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name* **ipv4-family**	Clears the route flapping information of a specified VPN instance enabled with an IPv4 address family.	The value must be an existing VPN instance name.
as-path-filter *as-path-filter-number*	Clears route flapping statistics based on the number of a specified AS_Path filter.	It is an integer that ranges from 1 to 256.
as-path-filter *as-path-filter-name*	Clears route flapping statistics based on the name of a specified AS_Path filter.	The value of *as-path-filter-name* is a string of 1 to 51 case-sensitive characters. **NOTE** When double quotation marks are used around the string, spaces are allowed in the string.
network-address	Specifies the IPv4 prefix address that is used to filter the BGP IPv4 routes.	It is in dotted decimal notation.
mask	Specifies the network mask that is used to filter the BGP IPv4 routes. If neither the mask nor the mask length is specified, the address is processed as a classful address.	It is in dotted decimal notation.

Parameter	Description	Value
mask-length	Specifies the network mask length that is used to filter the BGP IPv4 routes. If neither the mask nor the mask length is specified, the address is processed as a classful address.	The value is an integer that ranges from 0 to 32.
regexp *as-path-regexp*	Clears statistics about the flapping routes that match the AS_Path regular expression.	-
ipv6	Clears the route flapping statistics on all IPv6 peers.	-
vpn-instance *vpn-instance-name* **ipv6-family**	Clears the route flapping information of a specified VPN instance enabled with an IPv6 address family.	The value must be an existing VPN instance name.
network-ipv6-address	Specifies the IPv6 prefix address that is used to filter the BGP IPv6 routes.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the network mask length that is used to filter the BGP IPv6 routes.	The value is an integer that ranges from 0 to 128.
ipv4-address	Specifies the network address of an IPv4 peer.	It is in dotted decimal notation.
ipv6-address	Specifies the network address of an IPv6 peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
vpnv4	Resets route flapping statistics of BGP VPNv4 routes.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The process of adding a route to and then deleting the route from a routing table is called route flapping.

When route flapping occurs, the routing protocol sends Update packets to neighbors. The neighbors that receive the Update packets need to recalculate routes and modify the routing table. Therefore, frequent route flapping consumes great bandwidth and CPU resources and even seriously affects network operations.

The **reset bgp flap-info** command is used to clear the flapping information about routes. This allows the switch to re-collect statistics about flapping routes and helps to monitor route changes and locate network problems.

Prerequisites

You can use **display bgp routing-table flap-info** command to view the information about BGP route flapping.

If there are a large number of flapping routes, define an AS_Path filter or an AS_Path regular expression to be referenced in the **reset bgp flap-info** command. The flapping statistics of the routes matching the AS_Path filter or the AS_Path regular expression are then cleared.

Precautions

After the **reset bgp flap-info** command is run, the flapping statistics of routes are reset and cannot be displayed.

Follow-up Procedure

After the flapping statistics of routes are cleared, run the **display bgp routing-table flap-info** command again to display the flapping statistics about BGP routes in order to locate problems.

Example

Clear the flapping statistics about the routes that match AS_Path filter 10.

```
<HUAWEI> reset bgp flap-info as-path-filter 10
```

Clear the flapping statistics about the BGP4+ routes of the VPN instance named vpn1.

```
<HUAWEI> reset bgp vpn-instance vpn1 ipv6-family flap-info
```

7.8.153 reset bgp flapping-count

Function

The **reset bgp flapping-count** command resets the flapping count of a specified BGP peer.

Format

reset bgp [**vpn-instance** *vpn-instance-name* **ipv4-family**] *ipv4-address* **flapping-count**

reset bgp { **ipv6** | **vpn-instance** *vpn-instance-name* **ipv6-family** } *ipv6-address* **flapping-count**

NOTE

The **vpn-instance** *vpn-instance-name* **ipv6-family** is only supported on S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S.

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name* **ipv4-family**	Specifies the name of VPN instance enabled with an IPv4 address family.	The value must be an existing VPN instance name.
ipv4-address	Specifies the IPv4 address of a BGP peer.	It is in dotted decimal notation.
ipv6	Clears the flapping count of a specified BGP IPv6 peer.	-
vpn-instance *vpn-instance-name* **ipv6-family**	Specifies the name of VPN instance enabled with an IPv6 address family.	The value must be an existing VPN instance name.
ipv6-address	Specifies the IPv6 address of a BGP peer.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP peer flapping affects the stability of a BGP network and BGP route convergence.

The **reset bgp flapping-count** command can be used to clear the flapping account of a BGP peer. This allows the switch to re-collect the flapping statistics of a peer to locate BGP network problems.

Precautions

After the **reset bgp flapping-count** command is run, the flapping statistics of BGP peers are reset and cannot be displayed.

Follow-up Procedure

After the reset bgp flapping-count command is used to clear the statistics count of a specified BGP peer, run the **display bgp peer** command to display the flapping count of BGP peers and locate BGP network problems.

Example

Clear the flapping count of a specified BGP peer.

```
<HUAWEI> reset bgp 10.116.10.2 flapping-count
```

7.8.154 reset bgp mdt

Function

The **reset bgp mdt** command resets the BGP connections of a BGP MDT address family.

Format

reset bgp mdt { **all** | *ipv4-address* | **group** *group-name* | *as-number-plain* | *as-number-dot* } [**graceful**]

reset bgp mdt { **internal** | **external** } [**graceful**]

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Resets all BGP connections of a BGP MDT address family.	-
as-number-plain	Specifies an integral AS number.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies an AS number in dotted notation.	The value is in the format of x.y, where x and y are integers that range from 1 to 65535 and from 0 to 65535, respectively.
ipv4-address	Resets the connection with a specified BGP peer.	The value is in dotted decimal notation.

Parameter	Description	Value
group *group-name*	Resets the connection with a specified BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
internal	Resets all IBGP connections.	-
external	Resets all EBGP connections.	-
graceful	Resets BGP connections in GR mode.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Applicable Environment

When the BGP MDT configuration is changed, you can run the **reset bgp mdt** command to make the new configuration take effect immediately.

NOTICE

After the command is run, the TCP connection established by the BGP device is reset and the corresponding peer relationship is re-established. Therefore, exercise caution before you run this command.

Example

Reset all BGP connections of a BGP MDT address family.

```
<HUAWEI> reset bgp mdt all
```


7.8.155 reset bgp multicast

Function

The **reset bgp multicast** command resets the connections between a device and specified MBGP peers.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

reset bgp multicast { **all** | *peer-address* | **group** *group-name* | *as-number-plain* | *as-number-dot* | **external** | **internal** } [**graceful**]

Parameters

Parameter	Description	Value
all	Resets all MBGP connections.	-
peer-address	Resets the MBGP connection with a specified peer.	The peer address is in dotted decimal notation.
group-name	Specifies the name of a peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
as-number-plain	Specifies the number of the AS, in integer format.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies the number of the AS, in dotted notation.	The value is in the x.y format. Here, "x" and "y" are integers that range from 1 to 65535 and 0 to 65535 respectively.
external	Reset all EBGP connections.	-
internal	Reset all IBGP connections.	-
graceful	Specifies to reset MBGP connections in GR mode.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

NOTE

After the **reset bgp multicast** command is run on a device, the TCP connection established by the device and the corresponding peer is reset and then reestablished. Exercise caution when running this command.

Example

Reset the MBGP connections of all address families.

```
<HUAWEI> reset bgp multicast all
```

7.8.156 reset bgp multicast dampening

Function

The **reset bgp multicast dampening** command clears dampening information about routes in the MBGP routing table.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

reset bgp multicast dampening [*ip-address* [*mask* | *mask-length*]]

Parameters

Parameter	Description	Value
ip-address	Specifies a network address.	The value is in dotted decimal notation.
mask	Specifies the mask of a network address.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of a network address.	The value is an integer that ranges from 0 to 32.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

If a network address is specified in the **reset bgp multicast dampening** command, running this command clears the dampening information about a specific route and releases the suppressed route.

Example

Clear the dampening information about routes to network segment 10.1.0.0 and release suppressed routes.

```
<HUAWEI> reset bgp multicast dampening 10.1.0.0 255.255.0.0
```

7.8.157 reset bgp multicast external

Function

The **reset bgp multicast external** command resets all multicast EBGP connections.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

reset bgp multicast external [**graceful**]

Parameters

Parameter	Description	Value
graceful	Resets MBGP connections in GR mode.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can run the **reset bgp multicast external** command to reset all multicast EBGP connections.

Example

Reset all multicast EBGP connections.

```
<HUAWEI> reset bgp multicast external
```

7.8.158 reset bgp multicast flap-info

Function

The **reset bgp multicast flap-info** command clears route flapping statistics.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

reset bgp multicast flap-info [*ip-address* [*mask* | *mask-length*] | **as-path-filter** { *as-path-list-number* | *as-path-list-name* } | **regexp** *regexp*]

reset bgp multicast *ip-address* **flap-info**

Parameters

Parameter	Description	Value
ip-address	Specifies a network address.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of a network address.	The value is an integer that ranges from 0 to 32.
mask	Specifies the mask of a network address.	The value is in dotted decimal notation.
as-path-filter *as-path-list-number*	Clears route flapping statistics based on the number of a specified AS_Path filter.	The value is an integer that ranges from 1 to 256.
as-path-filter *as-path-list-name*	Clears route flapping statistics based on the name of a specified AS_Path filter.	The value is a string of 1 to 51 case-sensitive characters without spaces.
regexp *regexp*	Clears flapping statistics about the routes that match the AS_Path regular expression.	The value is a string of 1 to 80 characters.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can run the **reset bgp multicast flap-info** command to clear route flapping statistics.

Example

Clear the flapping statistics about the routes that match AS_Path filter 10.

```
<HUAWEI> reset bgp multicast flap-info as-path-filter 10
```

7.8.159 reset bgp multicast internal

Function

The **reset bgp multicast internal** command resets the multicast IBGP connections in an AS.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

reset bgp multicast internal [**graceful**]

Parameters

Parameter	Description	Value
graceful	Specifies to reset MBGP connections in GR mode.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

You can run the **reset bgp multicast internal** command to reset the multicast IBGP connections in an AS.

Example

Reset all multicast IBGP connections.

```
<HUAWEI> reset bgp multicast internal
```

7.8.160 reset bgp mvpn

Function

The **reset bgp mvpn** command resets the BGP connections of a BGP MVPN address family.

Format

reset bgp mvpn { **all** | *ipv4-address* | **group** *group-name* | *as-number-plain* | *as-number-dot* } [**graceful**]

reset bgp mvpn { **internal** | **external** } [**graceful**]

NOTE

Only the following switch models support this command:

S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H

Parameters

Parameter	Description	Value
all	Resets all BGP connections of a BGP MVPN address family.	-
as-number-plain	Specifies an integral AS number.	The value is an integer that ranges from 1 to 4294967295.
as-number-dot	Specifies an AS number in dotted notation.	The value is in the format of x.y, where x and y are integers that range from 1 to 65535 and from 0 to 65535, respectively.
ipv4-address	Resets the connection with a specified BGP peer.	The value is in dotted decimal notation.
group *group-name*	Resets the connection with a specified BGP peer group.	The name is a string of 1 to 47 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
internal	Resets all IBGP connections.	-
external	Resets all EBGP connections.	-
graceful	Resets BGP connections in GR mode.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Applicable Environment

When the BGP MVPN configuration is changed, you can run the **reset bgp mvpn** command to make the new configuration take effect immediately.

NOTICE

After the command is run, the TCP connection established by the BGP device is reset and the corresponding peer relationship is re-established. Therefore, exercise caution before you run this command.

Example

Reset all BGP connections of a BGP MVPN address family.

```
<HUAWEI> reset bgp mvpn all
```

7.8.161 router-id (BGP)

Function

The **router-id** command configures a Router ID for the switch.

The **undo router-id** command deletes the Router ID configured for the switch.

By default, no BGP Router ID is configured, and the global Router ID configured through the **router id** command is used.

Format

router-id { *ipv4-address* | **vpn-instance auto-select** }

undo router-id [**vpn-instance auto-select**]

Parameters

Parameter	Description	Value
ipv4-address	Specifies a Router ID.	It is in dotted decimal notation.
vpn-instance auto-select	Configures automatic Router ID selection for all BGP-VPN instance IPv4 or IPv6 address families. If a Router ID is manually specified for a BGP-VPN instance IPv4 or IPv6 address family, the manually specified Router ID takes precedence over the automatically selected Router ID.	-

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **router-id** command is used to configure a Router ID for the switch. Each Router ID uniquely identifies one BGP switch in an AS.

By configuring automatic Router ID selection for a BGP-VPN instanceIPv4 or IPv6 address family, you can differentiate the configured Router ID of the BGP-VPN instance IPv4 or IPv6 address family from the BGP Router ID. For more information about the Router ID of a BGP-VPN instance IPv4 or IPv6 address family, see the **router-id (BGP-VPN Instance View)** command.

Prerequisites

The **bgp** command is run to enable BGP.

Precautions

Changing or deleting a configured Router ID in the BGP view resets a BGP session. If a BGP session has been established in a BGP-VPN instance IPv4 address family, deleting the configured Router ID resets the BGP session. Exercise caution when changing or deleting a Router ID.

By default, the switch that is not configured with any interface uses the Router ID of 0.0.0.0 assigned by routing management.

Example

Configure a Router ID for the switch.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] router-id 10.1.1.1
```

Configure automatic Router ID selection for all BGP-VPN instance IPv4 address families.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] router-id vpn-instance auto-select
```

7.8.162 router-id (BGP-VPN Instance View)

Function

The **router-id** command configures router ID for BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view.

The **undo router-id** command deletes the router ID configured for BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view.

By default, no router ID is configured for BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view, and the BGP router ID is used as the router ID.

Format

router-id { *ipv4-address* | **auto-select** }

undo router-id

Parameters

Parameter	Description	Value
ipv4-address	Specifies the router ID of a BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view. The router ID is expressed in the IPv4 address format.	It is in dotted decimal notation.
auto-select	Configures automatic route ID selection for the current BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view.	-

Views

BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By configuring router ID for BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view, you can differentiate the configured router ID of BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view from the BGP router ID.

For example, if two VPN instances named **vrf1** and **vrf2** are configured on a PE, and a BGP session needs to be established between the interfaces bound to the two VPN instances, you need to configure different router IDs for the two VPN instance IPv4 address families. If no router ID is configured for the two VPN instance IPv4 address families, no BGP session can be established because the two VPN instance IPv4 address families have the same router ID, which is consistent with the BGP router ID.

Rules for automatically selecting a router ID for a BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view are as follows:

- If loopback interfaces configured with IP addresses are bound to the VPN instance, the largest IP address among the IP addresses of the loopback interfaces is selected as the router ID.
- If no loopback interfaces configured with IP addresses are bound to the VPN instance, the largest IP address among the IP addresses of other interfaces bound to the VPN instance is selected as the router ID, regardless of whether the interface is Up or Down.

Precautions

If a BGP session has been established in a BGP VPN instance IPv4 address family or BGP-VPN instance IPv6 address family view, changing or deleting the configured router ID resets the BGP session. So, confirm the action before you use the **router-id** command.

Example

Configure a router ID for a BGP VPN instance IPv4 address family.

```
<HUAWEI> system-view
[HUAWEI] ip vpn-instance vrf1
[HUAWEI-vpn-instance-vrf1] route-distinguisher 100:1
[HUAWEI-vpn-instance-vrf1-af-ipv4] quit
[HUAWEI-vpn-instance-vrf1] quit
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family vpn-instance vrf1
[HUAWEI-bgp-vrf1] router-id 192.168.100.1
```

7.8.163 route-select delay

Function

The **route-select delay** command configures a delay for selecting routes.

The **undo route-select delay** command deletes the delay for selecting routes.

The default delay is 0, indicating that routes are selected without a delay.

Format

route-select delay *delay-value*

undo route-select delay

Parameters

Parameter	Description	Value
delay-value	Specifies the delay for selecting routes.	The value is an integer that ranges from 0 to 3600, in seconds.

Views

BGP view, BGP-IPv4 unicast address family view, ,BGP-MDT address family view, BGP-MVPN address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv6 address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-IPv4 multicast address family view, BGP L2VPN-AD address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In a scenario in which primary and backup routes exist, packets may get lost after traffic switches back to the primary path from the backup path. **Figure 7-3** illustrates a VPN FRR scenario. PE3 and CE2 connect both to PE1 and PE2. The primary path is PE3 -> PE1 -> CE2, and the backup path is PE3 -> PE2 -> CE2. CE2 uses BGP to communicate with PE1 and PE2. FRR is configured on PE3. If PE1 restarts or the link between CE2 and PE1 is disconnected, traffic switches from the primary path to the backup path. After the primary path recovers, traffic switches back to the primary path. If PE3 completes refreshing forwarding entries before PE1 does so, PE1 may temporarily fail to forward traffic after a switchback. As a result, packet loss may occur. The severity of packet loss is proportional to the number of routes stored on PE1.

To solve this problem, run the **route-select delay** command on PE3 to configure a delay for selecting a route to PE1. An appropriate delay ensures that traffic switches back to the primary path after PE1 completes refreshing forwarding entries.

Figure 7-3 VPN FRR networking

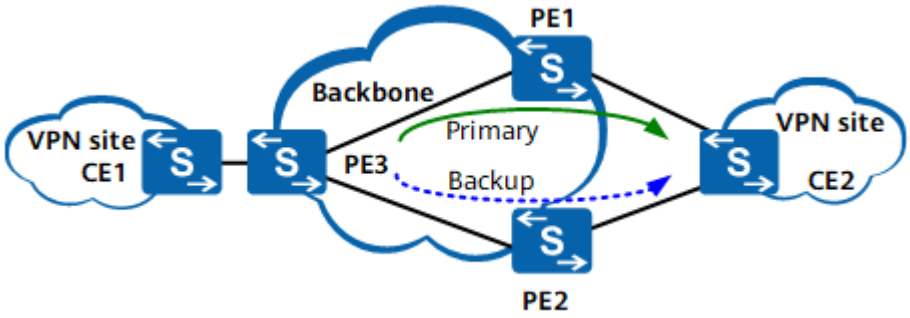

Precautions

If you run the **route-select delay** command repeatedly, the latest configuration overrides the previous configurations. If a route selection delay timer has started when you configure a new route select delay, the new route selection delay takes effect since the next route selection.

Example

Configure the delay for selecting routes as 300s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] route-select delay 300
```

7.8.164 routing-table limit threshold-alarm

Function

The **routing-table limit threshold-alarm** command configures alarm and alarm clear thresholds for the number of BGP routes.

The **undo routing-table limit threshold-alarm** command restores the default settings.

By default, the alarm threshold is 80%, and the alarm clear threshold is 70%.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

routing-table limit threshold-alarm upper-limit *upper-limit-value* **lower-limit** *lower-limit-value*

undo routing-table limit threshold-alarm [**upper-limit** *upper-limit-value* **lower-limit** *lower-limit-value*]

Parameters

Parameter	Description	Value
upper-limit *upper-limit-value*	Specifies an alarm threshold for the number of BGP routes.	The value is an integer that ranges from 1 to 100, in percentage. The default value is 80.
lower-limit *lower-limit-value*	Specifies an alarm clear threshold for the number of BGP routes.	The value is an integer that ranges from 1 to 100, in percentage. The default value is 70. *lower-limit-value* must be smaller than *upper-limit-value*; otherwise, alarms are generated and cleared repeatedly if route flapping occurs.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The number of BGP routes that can be added to a routing table is limited. If the value exceeds the limit, new routes cannot be added to the routing table, causing service interruptions. To address this problem, run the **routing-table limit threshold-alarm** command to configure alarm and alarm clear thresholds for the number of BGP routes. Alarms are then generated and cleared as expected. The alarms act as a prompt for checking whether an exception occurs and to take preventive measures.

- When the ratio of BGP routes to the maximum value exceeds *upper-limit-value*, an alarm is generated. New routes can still be accepted until the number of BGP routes reaches the maximum value.
- When the ratio falls below *lower-limit-value*, the alarm is cleared.

Configuration Impact

If the **routing-table limit threshold-alarm** command is run multiple times, the latest configuration takes effect.

Precautions

In addition to the **routing-table limit threshold-alarm** command, the **snmp-agent trap enable feature-name bgp trap-name** { **hwBgpRouteThresholdExceed** | **hwBgpRouteThresholdClear** } command must be run to enable the alarm and alarm clear functions; otherwise, alarms cannot be generated and cleared as expected.

Example

Configure alarm and alarm clear thresholds for the number of BGP routes.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] routing-table limit threshold-alarm upper-limit 70 lower-limit 60
```

7.8.165 routing-table rib-only

Function

The **routing-table rib-only** command prevents BGP routes from being added into the IP routing table.

The **undo routing-table rib-only** command restores the default setting.

By default, the preferred BGP routes are added to the IP routing table.

Format

routing-table rib-only [**route-policy** *route-policy-name*]

undo routing-table rib-only

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of a Route-Policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In BGP/MPLS IP VPN networking, if the BGP routing table has large numbers of VPN routes, these routes will consume large numbers of memory resources after being delivered to the IP VPN routing table. If these routes are not used in traffic forwarding, you can run the **routing-table rib-only** command to prevent these routes from being added to the IP VPN routing table. If some of these routes are not used in traffic forwarding, you can run the **routing-table rib-only route-policy** command to prevent this part of routes from being added to the IP VPN routing table.

If a route reflector (RR) is used and preferred BGP routes do not need to be used during the forwarding, the **routing-table rib-only** command can be used to make BGP routes unable to be added to the IP routing table or the forwarding layer. This improves forwarding efficiency and the system capacity.

When **route-policy** *route-policy-name* is specified in the command, the routes matching the policy are not added to the IP routing table, and the routes not matching the policy are added to the IP routing table with the route attributes unchanged.

Configuration Impact

After the **routing-table rib-only** command is run, the routes preferred by BGP are not added to the IP routing table.

The **routing-table rib-only** command does not take effect on the labeled routes.

Precautions

The **routing-table rib-only** command and the **active-route-advertise** command are mutually exclusive.

Creating a route-policy before it is referenced is recommended. By default, nonexistent route-policies cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent route-policy is referenced using the current command, all routes are not delivered to the IP routing table.

Example

Configure the routing policy named **ribonly** to prevent certain BGP routes from being added into the IP routing table.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] routing-table rib-only route-policy ribonly
```

7.8.166 shutdown (BGP)

Function

The **shutdown** command terminates all sessions between a device and its BGP peers.

The **undo shutdown** command restores the default setting.

By default, the function of closing all sessions between a device and its BGP peers is disabled.

Format

shutdown

undo shutdown

Parameters

None

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

NOTICE

Running the **shutdown** command closes all sessions between a device and its peers. Exercise caution when using this command.

During system upgrade or maintenance, the sessions between a device and its BGP peers need to be closed to minimize the impact of BGP route flapping on the network. If a large number of BGP peers exist, the **shutdown** command can be run in the BGP view to close all sessions with BGP peers. This frees you from running the **peer ignore** command repeatedly to close the sessions one by one.

Example

Close all sessions between a device and its BGP peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] shutdown
Warning: All BGP peer sessions will be interrupted. Continue? [Y/N]:y
```

7.8.167 slow-peer detection disable

Function

The **slow-peer detection disable** command disables slow peer detection.

The **undo slow-peer detection disable** command restores the default configuration.

By default, slow peer detection is enabled.

Format

slow-peer detection disable

undo slow-peer detection disable

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, or BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An update peer-group may consist of multiple BGP peers. If a network problem (congestion for example) occurs and slows down the speed at which the local device advertises routes to a BGP peer in the update peer-group, the speed at which the local device advertises routes to other BGP peers in the update peer-group is affected. To address this problem, slow peer detection is enabled by default.

When slow peer detection is enabled, the local device identifies the BGP peer to which routes are sent the slowest based on the time taken to send 100 packets to each BGP peer. If this time is greater than the period threshold for slow peer detection plus the average time taken to send 100 packets to BGP peers (excluding the longest and shortest times), the local device considers the peer a slow peer and removes it from the update peer-group. Slow peer detection prevents this slow peer from affecting route advertisement to other peers in the update peer-group.

To disable slow peer detection, run this command.

Example

Disable slow peer detection.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] slow-peer detection disable
```


7.8.168 slow-peer detection threshold

Function

The **slow-peer detection threshold** command sets a period threshold for slow peer detection.

The **undo slow-peer detection** command restores the default configuration.

By default, the period threshold for slow peer detection is 300s.

Format

slow-peer detection threshold *threshold-value*

undo slow-peer detection [**threshold** *threshold-value*]

Parameters

Parameter	Description	Value
threshold *threshold-value*	Specifies a period threshold for slow peer detection.	The value is an integer that ranges from 120 to 3600, in seconds.

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv6 unicast address family view, BGP-VPNv4 address family view, BGP-VPNv6 address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view, or BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When slow peer detection is enabled, the local device identifies the BGP peer to which routes are sent the slowest based on the time taken to send 100 packets to each BGP peer. If this time is greater than the period threshold for slow peer detection plus the average time taken to send 100 packets to BGP peers (excluding the longest and shortest times), the local device considers the peer a slow peer and removes it from the update peer-group. Slow peer detection prevents this slow peer from affecting route advertisement to other peers in the update peer-group.

By default, slow peer detection is enabled and the period threshold for slow peer detection is 300s. To adjust the period threshold for slow peer detection, run this command.

Configuration Impact

If the command is run more than once, the latest configuration overrides the previous one.

Precautions

After a slow peer is removed from the update peer-group, the peer relationship between the local device and the slow peer is reestablished.

Example

Set the period threshold for slow peer detection to 200s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] slow-peer detection threshold 200
```

7.8.169 summary automatic

Function

The **summary automatic** command enables automatic aggregation for the locally imported routes.

The **undo summary automatic** command disables automatic aggregation for the locally imported routes.

By default, automatic aggregation is disabled for the locally imported routes.

Format

summary automatic

undo summary automatic

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-VPN instance IPv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In the BGP-IPv4 unicast address family view, the **summary automatic** command is used to configure a BGP device to automatically aggregate locally imported routes on the public network.

In the BGP-VPN instance IPv4 address family view, the **summary automatic** command is used to configure a BGP device to automatically aggregate locally imported routes on a private network.

The **summary automatic** command is used to aggregate the routes imported by BGP. These routes can be direct routes, static routes, RIP routes, OSPF routes, or IS-IS routes. After this command is run on a BGP device, the BGP device aggregates routes based on the natural network segment (for example, 10.1.1.0/24 and 10.2.1.0/24 are aggregated to 10.0.0.0/8, a Class A address), and sends only the aggregated route to its peers. This reduces the number of routes.

Precautions

BGP route aggregation is classified into manual aggregation and automatic aggregation. The command is used to implement automatic aggregation. Manual aggregation takes precedence over automatic aggregation.

The **summary automatic** command is invalid for the routes imported by using the **network** command.

Example

Enable automatic aggregation for imported routes.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] summary automatic
```

7.8.170 supernet unicast advertise

Function

The **supernet unicast advertise enable** command configures a BGP device to advertise BGP supernet unicast routes to its peers.

The **undo supernet unicast advertise enable** or **supernet unicast advertise disable** command restores the default configuration.

By default, BGP supernet unicast routes are considered invalid and cannot be advertised to BGP peers or delivered to the IP routing table.

Format

supernet unicast advertise enable

supernet unicast advertise disable

undo supernet unicast advertise enable

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-VPN instance IPv4 address family view, BGP-IPv6 unicast address family view, or BGP-VPN instance IPv6 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A BGP supernet route has the same destination address and next-hop address or has a more detailed destination address than the next-hop address. Any route that meets one of the following conditions is a BGP supernet route.

- If you perform bitwise AND operations on the destination address mask with the destination address and next-hop address, respectively, the calculated network addresses are the same, and the destination address mask is greater than or equal to the next-hop address mask.
- If you perform bitwise AND operations on the destination address mask with the destination address and next-hop address, respectively, the calculated network addresses are different. However, if you perform bitwise AND operations on the next-hop address mask with the destination address and next-hop address, respectively, the calculated network addresses are the same.

For example, the route destined for 10.6.6.6 in the following command output is a BGP supernet route.

```
<HUAWEI> display bgp routing-table
 BGP Local router ID is 10.1.1.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
  *>i 10.6.6.6/32        10.6.6.6       0          100       0       ?
```

BGP supernet routes include BGP supernet labeled routes and BGP supernet unicast routes. To allow a Huawei device to advertise BGP supernet unicast routes that it receives from a connected non-Huawei device to its BGP peers, run the **supernet unicast advertise enable** command on the Huawei device.

Example

Configure a BGP device to advertise BGP supernet unicast routes to its peers.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] supernet unicast advertise enable
```


7.8.171 timer (BGP)

Function

The **timer** command sets the values for the Keepalive timer and Hold timer.

The **undo timer** command restores the default values of the Keepalive timer and Hold timer.

By default, the value of a Keepalive timer is 60s and the value of a Hold timer is 180s.

Format

timer keepalive *keepalive-time* **hold** *hold-time* [**min-holdtime** *min-holdtime*]

undo timer keepalive *keepalive-time* **hold** *hold-time* [**min-holdtime** *min-holdtime*]

undo timer keepalive hold [**min-holdtime**]

Parameters

Parameter	Description	Value
keepalive *keepalive-time*	Specifies the Keepalive period.	The value is an integer that ranges from 0 to 21845, in seconds.
hold *hold-time*	Specifies the holdtime.	The value is an integer that can be 0, or ranges from 3 to 65535, in seconds.
min-holdtime *min-holdtime*	Specifies the minimum holdtime configured on the local device. **NOTE** The value of *min-holdtime* configured cannot exceed the value of *hold-time*.	The value is an integer that ranges from 20 to 65535, in seconds.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After a connection is established between peers, the values of *keepalive-time* and *hold-time* are negotiated by the peers.

- The smaller of the *hold-time* values carried by Open messages of both peers is taken as the *hold-time* value.
- The smaller of one third of the *hold-time* value and the locally configured *keepalive-time* value is taken as the *keepalive-time* value.

If the local device establishes BGP peer relationships with many devices, it needs to process huge BGP messages. If *hold-time* negotiated among BGP peers is small, the timer may expire before the local device processes the Keepalive messages sent from other BGP peers. The peer relationships are then interrupted, and routes flap. To solve the preceding problem, you can configure an appropriate value for **min-holdtime** *min-holdtime* based on the CPU processing capability of the local device.

If the value of *min-holdtime* is changed, but the values of *keepalive-time* and *hold-time* negotiated between two BGP peers remain unchanged, the established peer relationship is not affected. Only when the local device attempts to re-establish a relationship with a remote device, the value of *min-holdtime* configured on the local device takes effect. The local device compares *min-holdtime* with *hold-time* sent from the remote device. If the value of *min-holdtime* exceeds that of *hold-time*, *hold-time* negotiation fails, and the peer relationship fails to be established.

NOTE

If *min-holdtime* is configured on the local device, and the value of *hold-time* sent from the remote device is 0, *hold-time* negotiation between the two devices succeeds. The negotiated value of *hold-time* is 0, and the peer relationship is established. The value 0 of *hold-time* indicates that the peer relationship never expires.

Precautions

The timers configured for a specific peer or peer group by using the **peer timer** command override the timers configured for all BGP peers by using the **timer** command.

If the value of a timer changes, the BGP peer relationship between devices is disconnected. This is because the devices need to re-negotiate the values of *keepalive-time* and *hold-time*. Therefore, exercise caution before changing the value of a timer.

Setting the Hold timer value to at least three times the Keepalive timer value is recommended. When setting the values of *keepalive-time* and *hold-time*, note the following points:

- The values of *keepalive-time* and *hold-time* cannot both be set to 0. This renders the BGP timers become invalid. This means that BGP is unable to detect link faults using the timers.
- The *hold-time* value cannot be significantly greater than the *keepalive-time* value. A setting of **timer keepalive 1 hold 65535**, for example, would be improper. If the holdtime is too long, link faults cannot be detected in a timely manner.

Example

On a BGP device, set the value of the Keepalive timer to 30s and the value of the Hold timer to 90s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] timer keepalive 30 hold 90
```

7.8.172 timer connect-retry

Function

The **timer connect-retry** command sets a global ConnectRetry interval.

The **undo timer connect-retry** command restores the default setting.

By default, the ConnectRetry interval is 32s.

Format

timer connect-retry *connect-retry-time*

undo timer connect-retry

Parameters

Parameter	Description	Value
connect-retry-time	Specifies a ConnectRetry interval.	The value is an integer that ranges from 1 to 65535, in seconds.

Views

BGP view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When BGP initiates a TCP connection, the ConnectRetry timer is stopped if the TCP connection is established successfully. If the first attempt to establish a TCP connection fails, BGP tries again to establish the TCP connection after the ConnectRetry timer expires. The ConnectRetry interval can be adjusted as needed.

- The ConnectRetry interval can be reduced in order to lessen the time BGP waits to retry establishing a TCP connection after the first attempt fails.
- To suppress route flapping caused by constant peer flapping, the ConnectRetry interval can be increased to accelerate route convergence.

Precautions

A ConnectRetry interval can be configured globally, or on a particular peer or peer group. A ConnectRetry interval configured on a specific peer or peer group takes precedence over a global ConnectRetry interval.

If both the **peer** { *group-name* | *ipv4-address* | *ipv6-address* } **timer connect-retry** *connect-retry-time* command and the **timer connect-retry** *connect-retry-time* command are run on a device, the configuration of the **peer** { *group-name* | *ipv4-address* | *ipv6-address* } **timer connect-retry** *connect-retry-time* command takes effect, but the configuration of the **timer connect-retry** *connect-retry-time* command does not.

Example

Set a global BGP ConnectRetry interval to 60s.

```
<HUAWEI> system-view
[HUAWEI] bgp 100
[HUAWEI-bgp] timer connect-retry 60
```

7.8.173 undo synchronization (BGP)

Function

The **undo synchronization** command disables synchronization between BGP and an IGP.

By default, synchronization between BGP and an IGP is disabled.

Format

undo synchronization

Parameters

None

Views

BGP view, BGP-IPv4 unicast address family view, BGP-IPv4 multicast address family view, BGP-IPv6 unicast address family view

Default Level

2: Configuration level

Usage Guidelines

You can run the **undo synchronization** command to disable synchronization between BGP and an IGP.

Example

Disable synchronization between BGP and an IGP.

```
<HUAWEI> system-view
```



```
[HUAWEI] bgp 100
[HUAWEI-bgp] ipv4-family unicast
[HUAWEI-bgp-af-ipv4] undo synchronization
```

7.9 Routing Policy Configuration Commands

7.9.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.9.2 apply as-path

Function

The **apply as-path** command sets the action for changing the AS_Path attribute of BGP routes in a routing policy.

The **undo apply as-path** command restores the default setting.

By default, the action for changing the AS_Path attribute of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply as-path { { *as-number-plain* | *as-number-dot* } &<1-10> { **additive** | **overwrite** } | **none overwrite** }

undo apply as-path

Parameters

Parameter	Description	Value
as-number-dot	Specifies an AS number in dotted notation to be added to the AS_Path list or to replace the existing AS_Path list. A maximum of 10 AS numbers can be specified in one command.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.
additive	Adds the specified AS number to the original AS_Path attribute.	-
overwrite	Replaces the original AS_Path with the specified AS number.	-
none	Clears the original AS_Path list.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To change the AS_Path attribute of BGP routes BGP for selecting the optimal route, you can apply a routing policy containing the **apply as-path** command.

AS_Path is a private attribute of BGP and records all ASs that a route passes through from the local end to the destination address. Using the AS_Path attribute controls route selection and prevents routing loops. If multiple routes are destined for the same destination address, BGP compares the AS_Path lists of these routes and considers the route with the shortest AS_Path list as the optimal route.

After this command is configured, the AS_Path list for matched BGP routes will change. Assume that the original AS-Path is (30, 40, 50) and the BGP route matching condition is met. In this case:

- If the **apply as-path 60 70 80 additive** command is run, the AS-Path list is changed to (60, 70, 80, 30, 40, 50). This configuration change is generally used to make the BGP route not preferentially selected.
- If the **apply as-path 60 70 80 overwrite** command is run, the AS-Path list is changed to (60, 70, 80). There are many application scenarios for changing the AS-Path list, and the major application scenarios are as follows:
  - Hide the real path information of routes. For example, after the AS-Path list is changed to (60, 70, 80), the AS-Path information of the route (30, 40, 50) is lost.

  - Implement load balancing. For example, a router receives two routes with the same destination IP address 10.1.0.0/16. The AS_Path list of one route is (60, 70, 80) and that of the other route is (30, 40, 50). In this case, you can change the AS_Path list (30, 40, 50) to (60, 70, 80), and load balancing then may be implemented on the two routes.
  - Shorten the AS-Path list to prevent the route from being discarded. If the **as-path-limit** command is configured, whether the number of AS numbers in the AS-Path list of the incoming route exceeds the maximum value needs to be checked. If the number exceeds the maximum value, the route is discarded. Therefore, before receiving a route with a long AS-Path list, replace the AS-Path list with a shorter AS-Path list. For example, if the original AS-Path list is (60, 70, 80, 65001, 65002, 65003), run the **apply as-path 60 70 80 overwrite** command to change the AS-Path list to (60, 70, 80). In this manner, the length of the AS-Path is shortened, preventing the route from being discarded.
  - Shorten the AS-Path list to make the route to be preferentially selected and traffic to be directed to the local AS.
- If the **apply as-path none overwrite** command is run, the AS-Path list is changed to be vacant. In BGP route selection, if the AS-Path list is vacant, the length of the AS-Path list is considered as 0. Therefore, clearing the AS-Path list cannot only hide the real path information, but also make the route to be preferentially selected and traffic to be directed to the local AS because the AS-Path list is shortened.

Prerequisites

The **apply as-path** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it affects BGP route selection.

Running the apply as-path command changes the path through which network traffic passes. Use this command only when you are familiar with the network topology and impact of the command on services.

Example

Change the AS number in the original AS_Path attribute to 200, 10.10.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply as-path 200 10.10 additive
```

7.9.3 apply backup-interface

Function

The **apply backup-interface** command configures a backup outbound interface in a routing policy.

The **undo apply backup-interface** command restores the default setting.

By default, the backup outbound interface is not configured in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply backup-interface *interface-type interface-number*

undo apply backup-interface

Parameters

Parameter	Description	Value
interface-type interface-number	Specifies the type and number of the backup outbound interface.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply backup-interface** command is used in IP FRR scenarios to configure a backup outbound interface for a route. After IP FRR is enabled, data traffic can be quickly switched to the backup outbound interface if the primary link fails.

Prerequisites

if-match clauses can be used to configure matching rules such as IP prefix lists, and ACLs before a backup outbound interface is configured.

Follow-up Procedure

Reference a configured route-policy in the **ip frr (system view)** command or the **ip frr (VPN instance view)** command to configure IP FRR on a public network or VPN.

The **apply backup-interface** command is usually used together with the **apply backup-nexthop** command.

Precautions

For P2P links, a backup next hop is not necessary. For non-P2P links, a backup next hop is necessary.

Example

Configure the backup outbound interface and the backup next hop in the route-policy named **ip_frr_rp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy ip_frr_rp permit node 10
[HUAWEI-route-policy] apply backup-interface vlanif10
[HUAWEI-route-policy] apply backup-nexthop 192.168.20.2
```

Delete the configured backup outbound interface from the route-policy named **ip_frr_rp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy ip_frr_rp permit node 10
[HUAWEI-route-policy] undo apply backup-interface
```

7.9.4 apply backup-nexthop

Function

The **apply backup-nexthop** command configures a backup next hop in a routing policy.

The **undo apply backup-nexthop** command deletes the configured backup next hop.

By default, the backup next hop is not configured in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply backup-nexthop { *ipv4-address* | **auto** }

undo apply backup-nexthop

Parameters

Parameter	Description	Value
ipv4-address	Specifies the IP address of a backup next hop.	It is in dotted decimal notation.
auto	Automatically searches for the backup next hop.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply backup-nexthop** command is used to configure a backup next hop for a route in IP FRR and VPN FRR scenarios. After IP FRR is enabled, data traffic can be quickly switched to the backup next hop if the primary link fails.

Prerequisites

if-match clauses can be used to set matching rules such as interfaces, IP prefix lists, and ACLs before a backup next hop is configured.

Follow-up Procedure

The configured route-policy can be used in the **ip frr (system view)** command or the **ip frr (VPN instance view)** command that is run to configure IP FRR for a public or private network. It can also be used in the **vpn frr** command that is run to enable VPN FRR.

In a VPN FRR scenario, you only need to run the **apply backup-nexthop** command to configure a backup next hop.

In an IP FRR scenario, you need to run both the **apply backup-nexthop** and **apply backup-interface** commands.

Precautions

For P2P links, a backup next hop is not necessary. For non-P2P links, a backup next hop is necessary.

NOTE

Example

Configure the backup interface and the backup next hop 192.168.20.2 in the route-policy named **ip_frr_rp**.


```
<HUAWEI> system-view
[HUAWEI] route-policy ip_frr_rp permit node 10
[HUAWEI-route-policy] apply backup-interface vlanif10
[HUAWEI-route-policy] apply backup-nexthop 192.168.20.2
```

Delete the configured backup next hop from the route-policy named **ip_frr_rp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy ip_frr_rp permit node 10
[HUAWEI-route-policy] undo apply backup-nexthop
```

7.9.5 apply behavior

Function

The **apply behavior** command configures a QoS traffic behavior for routes.

The **undo apply behavior** command restores the default setting.

By default, no QoS traffic behavior is configured.

Format

apply behavior *behavior-name*

undo apply behavior

Parameters

Parameter	Description	Value
behavior-name	Specifies the name of a QoS traffic behavior.	The value is a string of 1 to 31 case-sensitive characters without spaces, and must start with a letter.

Views

Route-policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When a large and complex network is deployed, traffic flows of different types, such as voice, video, and data, need to be managed separately, with different bandwidth and delays assigned to these traffic flows. In this case, class-based QoS needs to be configured and complex traffic classification needs to be performed. The **apply behavior** command is used to associate filtered route with a specific traffic behavior.

Prerequisites

Before running the **apply behavior** *behavior-name* command, you need to create a traffic behavior.

Meanwhile, certain matching conditions need to be configured to classify routes, such as the AS-Path lost, community attribute list, address prefix list, and route cost.

Precautions

The **apply behavior** command is mutually exclusive with the **apply ip-precedence** command and the **apply qos-local-id** command, and only one of these commands can be configured on a node of a routing policy. For example, if the **apply behavior** command is configured in the view created by the **route-policy test permit node 10** command, configuring the **apply qos-local-id** command replaces **apply behavior** command.

Example

Configure the behavior named **example** in the system view, and then apply this QoS traffic behavior in the route-policy view.

```
<HUAWEI> system-view
[HUAWEI] traffic behavior example
[HUAWEI-behavior-example] quit
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] apply behavior example
```

7.9.6 apply comm-filter delete

Function

The **apply comm-filter delete** command sets the action for deleting community attributes of a specified community filter in a routing policy.

The **undo apply comm-filter** command restores the default setting.

By default, the action for deleting community attributes of a specified community filter is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply comm-filter { *basic-comm-filter-number* | *adv-comm-filter-number* | *comm-filter-name* } **delete**

undo apply comm-filter

Parameters

Parameter	Description	Value
basic-comm-filter-number	Specifies the number of a basic community filter.	The value is an integer ranging from 1 to 99.
adv-comm-filter-number	Specifies the number of an advanced community filter.	The value is an integer ranging from 100 to 199.
comm-filter-name	Specifies the name of a community filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To delete the community attributes, you can run the **ip community-filter** command several times to configure community attributes one by one, and apply the routing policy containing the **apply comm-filter delete** command to delete these community attributes.

The community attribute is a private attribute of BGP. The **apply comm-filter delete** command takes effect only for BGP routes.

Prerequisites

The **apply comm-filter delete** command can be used only after the **route-policy** command is used.

Precautions

After routes meet the filtering conditions, the specified community attributes of these routes are deleted.

1. When the delete operation is configured on a specified community attribute list, only one community attribute can be configured for the specified community attribute list. To delete multiple community attributes, you need to configure multiple community attribute lists. For example, if community attribute list 1 is used to delete 100:100 200:200 from the community attribute 100:100 200:200 carried in a route, you need to perform the following configurations on community attribute list 1:

   ```
   [HUAWEI] ip community-filter 1 permit 100:100
   [HUAWEI] ip community-filter 1 permit 200:200
   [HUAWEI] display ip community-filter
   Community filter Number: 1
   permit 100:100
   permit 200:200
   [HUAWEI] route-policy RP1 permit node 10
   [HUAWEI-route-policy] apply comm-filter 1 delete
   ```

   If multiple community attributes are configured in the same community filter, the **apply comm-filter delete** command cannot delete these community attributes. To delete the community attributes, you can run the **ip community-filter** command several times to configure community attributes one by one, and apply the routing policy containing the **apply comm-filter delete** command to delete these community attributes. For example, the following command cannot delete the community attribute 100:100 200:200 of the route:

   ```
   [HUAWEI] ip community-filter 1 permit 100:100 200:200
   [HUAWEI] display ip community-filter
   Community filter Number: 1
   permit 100:100 200:200
   [HUAWEI] route-policy RP1 permit node 10
   [HUAWEI-route-policy] apply comm-filter 1 delete
   ```

2. When the **apply community** and **apply comm-filter delete** commands are run on the same node in a routing policy, the system performs the delete operation before the set operation regardless of the sequence in which the two commands are run.

   ```
   [HUAWEI] display route-policy
   Route-policy : 123a
     permit : 10
   Match clauses:
   Apply clauses: a
   apply community 999:9 additive
   apply comm-filter 1 delete
   ```

   The following command output shows that community attribute 111:1 of the corresponding BGP route is deleted and community attribute 999:9 is added.

   ```
   [HUAWEI] display ip community-filter
   Community filter Number: 1
   permit 111:1
   permit 999:9
   ```

Example

Delete the specified BGP route community attributes 1:200, 2:200, and 3:200 from the community filter.

```
<HUAWEI> system-view
[HUAWEI] ip community-filter 1 permit 1:200
[HUAWEI] ip community-filter 1 permit 2:200
[HUAWEI] ip community-filter 1 permit 3:200
```



```
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] apply comm-filter 1 delete
```

7.9.7 apply community

Function

The **apply community** command sets the action for changing the community attribute of BGP routes in a routing policy.

The **undo apply community** command restores the default setting.

By default, the action for changing the community attribute of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply community none

apply community { *community-number* | *aa:nn* | **internet** | **no-advertise** | **no-export** | **no-export-subconfed** } &<1-32> [**additive**]

undo apply community

Parameters

Parameter	Description	Value
none	Indicates that all the community attributes of routes are deleted.	-

Parameter	Description	Value
community-number | *aa:nn*	Specifies the community number. A maximum of 32 community numbers can be configured in the **apply community** command. ● If you do not configure any one of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 32 *community-number* and *aa:nn* together. ● If you configure one of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 31 *community-number* and *aa:nn* together. ● If you configure two of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 30*community-number* and *aa:nn* together. ● If you configure three of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 29 *community-number* and *aa:nn* together. ● If you configure all of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 28 *community-number* and *aa:nn* together.	The value of *community-number* is an integer ranging from 0 to 4294967295. The value of *aa* or *nn* ranges from 0 to 65535.
internet	Indicates that matching routes are sent to any peer. By default, all routes belong to the Internet community.	-
no-advertise	Indicates that matching routes are not sent to any peer. That is, after a router receives a route with this attribute, it does not advertise the route to other BGP peers.	-
no-export	Indicates that matching routes are sent to other sub-ASs but not to other ASs. That is, after a router receives a route with this attribute, it does not advertise the route outside the local AS.	-
no-export-subconfed	Indicates that matching routes are neither sent to other sub-ASs nor to other ASs. That is, after a router receives a route with this attribute, it does not advertise the route to other sub-ASs.	-
additive	Indicates that community attributes are added to matching routes.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To identify the BGP routes, you can apply a routing policy containing the **apply community** command to configure the community attribute of matched BGP routes.

The community attribute is a private attribute of BGP. It simplifies the application of routing policies and facilitates route maintenance and management. A community is a set of destination addresses with the same characteristics. These addresses have no physical boundary and are independent of their ASs. They share one or multiple community attributes, which can be changed or set by running the **apply community** command.

Prerequisites

The **apply community** command can be used only after the **route-policy** command is used.

Precautions

If the **apply community** command is configured in a routing policy, the community attributes of the BGP routes that match the routing policy are changed according to the configurations in the routing policy.

Assume that the original community attribute of a BGP route is 30. If this BGP route matches a certain routing policy, the AS number is replaced or added on the basis of the routing policy. For example:

- If the **apply community 100** command is run, the community attribute is changed to 100.
- If the **apply community 100 150** command is run, the community attribute is changed to 100, 150.
- If the **apply community 100 150 additive** command is run, the community attribute is changed to 30, 100, 150.
- If the **apply community none** command is run, the community attribute of the BGP route is deleted.

Example

Configure a routing policy named **setcommunity**, match the route with the AS_Path filter being 8, and change its community attribute to **no-export**.

```
<HUAWEI> system-view
[HUAWEI] route-policy setcommunity permit node 16
[HUAWEI-route-policy] if-match as-path-filter 8
[HUAWEI-route-policy] apply community no-export
```


7.9.8 apply cost

Function

The **apply cost** command sets the action for changing the cost of routes in a routing policy.

The **undo apply cost** command restores the default setting.

By default, the action for changing the cost of routes is not set in a routing policy.

Format

apply cost [**+** | **-**] *cost*

undo apply cost

Parameters

Parameter	Description	Value
+	Increases the route cost.	If the MED of BGP routes or cost of non-BGP routes is greater than the maximum value (4294967295) after the adjustment, 4294967295 takes effect.
-	Reduces the route cost.	If the MED of BGP routes or cost of non-BGP routes is less than the minimum value (0) after the adjustment, 0 takes effect.
cost	Specifies the route cost. To control route selection, you can adjust the route cost to prevent routing loops.	The value is an integer ranging from 0 to 4294967295.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When the filtering conditions specified by **if-match** clauses are met, you can run the **apply cost** command to change the route MED or cost to control route selection. After setting the MED or cost, the MED or cost of the routes that are imported using the route-policy is changed accordingly.

Prerequisites

A route-policy has been configured by **route-policy**.

Configuration Impact

The costs of routes that match the route-policy are changed. BGP routes do not have costs, and instead, they have MEDs. If the **apply cost** command is run to configure an **apply** clause for a route-policy that is designed for BGP routes, the MEDs of BGP routes that match the route-policy are changed.

Precautions

The MEDs or costs of imported routes are independent of the route-policy after the **undo apply cost** command is used to cancel the route MED or cost.

Example

Define an **apply** clause to set the route cost to 120.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply cost 120
```

7.9.9 apply cost-type

Function

The **apply cost-type** command sets the action for changing the cost type of routes in a routing policy.

The **undo apply cost-type** command restores the default setting.

By default, the action for changing the cost type of routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply cost-type { **external** | **internal** | **type-1** | **type-2** }

undo apply cost-type

Parameters

Parameter	Description	Value
external	Sets the cost type of IS-IS external routes.	-
internal	Sets the cost type of IS-IS internal routes or sets the MED value of BGP routes as the IGP cost of the next hop.	-
type-1	Sets Type 1 external routes of OSPF.	-
type-2	Sets Type 2 external routes of OSPF.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply cost-type** { **external** | **internal** } command sets the cost type of IS-IS routes. The cost of an internal route imported to IS-IS remains unchanged and the cost of an external route imported to IS-IS is increased by 64.

NOTE

When the cost-style of an IS-IS device is wide, compatible, or wide-compatible, the cost types of external or internal are classified. When the cost-style of an IS-IS device is narrow or narrow-compatible, the imported route classifies the cost type of external or internal. In other words, the **apply cost-type** command is valid on IS-IS devices only in the narrow or narrow-compatible modes.

The **apply cost-type** { **type-1** | **type-2** } command modifies the type of OSPF routes. During route import, OSPF modifies the type but not the cost value of the original route. When OSPF advertises the imported route with the cost and type information to a peer, the peer device will recalculate the cost value of the imported route based on the received information.

When the filtering conditions specified by **if-match** clauses are met, you can change the cost type of routes to set the imported external routes all to Type-1 or all to Type-2 by using the **apply cost-type** command. After the cost type of the routes that match the route-policy is set, the cost type of the routes that are imported by using the route-policy is the set cost type.

Prerequisites

Before running the **apply cost-type** command, you need to configure a route-policy by **route-policy**.

Configuration Impact

After routes match the route-policy, the cost type of the routes is changed.

Precautions

Different operations are performed when the **apply cost-type internal** command is applied to IS-IS routes and BGP routes:

- When the **apply cost-type internal** command is applied to IS-IS routes:

  Routes are configured as IS-IS internal routes.

- When the **apply cost-type internal** command is applied to BGP routes:

  When a switch advertises a route learned from an IBGP peer to an EBGP peer, if the **apply cost-type internal** command is run, the switch sets the MED value of the route to be advertised to the EBGP peer as the IGP cost of the next hop of the route.

Example

Set the cost type to OSPF external Type-1.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply cost-type type-1
```

7.9.10 apply dampening

Function

The **apply dampening** command sets the action for changing the dampening parameters of EBGP routes in a routing policy.

The **undo apply dampening** command restores the default setting.

By default, the action for changing the dampening parameters of EBGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply dampening *half-life-reach reuse suppress ceiling*

undo apply dampening

Parameters

Parameter	Description	Value
half-life-reach	Specifies the half-life of a reachable route.	The value is an integer ranging from 1 to 45, in minutes.
reuse	Specifies the threshold for routes to be released from the dampening state. When the penalty value falls below the threshold, routes are reused.	The value is an integer ranging from 1 to 20000.
suppress	Specifies the threshold for routes to enter the dampening state. When the penalty value exceeds the threshold, routes are suppressed.	The value is an integer ranging from 1 to 20000. The configured value of *suppress* must be greater than the value of *reuse*.
ceiling	Specifies the upper limit of the penalty value of routes.	The value is an integer ranging from 1001 to 20000. The configured value of *ceiling* must be greater than the value of *suppress*.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply dampening** command, which is mostly used in BGP, is used to prevent frequent route dampening from affecting routers on the network.

You can configure different route dampening parameters for different nodes in the same routing policy. When route flapping occurs, BGP can use different route dampening parameters to suppress the routes that match the routing policy.

Procedure

If the **apply dampening** command is run multiple times, the latest configuration overwrites the previous one.

Prerequisites

The **route-policy** command has been configured.

Precautions

The parameters in this command do not have default values and must be set. The values of *reuse*, *suppress*, and *ceiling* are listed in ascending order: *reuse* < *suppress* <*ceiling*. According to the formula, MaxSuppressTime = *half-life-reach* x 60 x (ln (*ceiling*/*reuse*)/ln (2)), routes are unsuppressed if the value of MaxSuppressTime is less than 1. Therefore, the value of the *ceiling*/*reuse* must be great enough so that the value of MaxSuppressTime can be equal to or greater than 1.

Example

Set dampening parameters for EBGP routes.

```
<HUAWEI> system-view
[HUAWEI] route-policy aa permit node 10
[HUAWEI-route-policy] apply dampening 20 2000 10000 16000
```

7.9.11 apply extcommunity

Function

The **apply extcommunity** command sets the action for changing the extended community attribute of BGP routes in a routing policy.

The **undo apply extcommunity** command restores the default setting.

By default, the action for changing the extended community attribute of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply extcommunity { **rt** { *as-number:nn* | *4as-number:nn* | *ipv4-address:nn* } } &<1-16> [**additive**]

undo apply extcommunity

Parameters

Parameter	Description	Value
rt	Indicates the route-target extended community. A maximum of 16 route targets can be configured.	-
as-number	Specifies the AS number.	The value is an integer ranging from 0 to 65535.
4as-number	Specifies a 4-byte AS number.	A 4-byte AS number is divided into the following types: • It is an integer ranging from 65536 to 4294967295. • It is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535 respectively
ipv4-address	Specifies the IPv4 address.	It is in dotted decimal notation.
nn	Specifies an integer.	• When the value of *as-number* is a 2-byte AS number, the value of *nn* ranges from 0 to 4294967295. • When the value of *4as-number* is a 4-byte AS number, the value of *nn* ranges from 0 to 65535. • For *ipv4-address*, the value of *nn* ranges from 0 to 65535.
additive	Indicates that existing community attributes can be added to routes.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

When controlling inter-AS VPN route receiving and advertising, apply the routing policy that contains the **apply extcommunity** command to change the RT extended community attribute of matched routes. Currently, only the RT extended community attribute is supported. This command cannot specify an extended community attribute for public routes.

Prerequisites

The **apply extcommunity** command can be used only after the **route-policy** command is used.

Precautions

When the routing policy that contains the action is used in the BGP view, BGP IPv4 unicast address view, or BGP IPv6 unicast address view, the action does not take effect.

When a routing policy takes effect, it affects inter-AS VPN route receiving and advertising.

If the keyword **additive** is not set in the **apply extcommunity** command, the original extended community attribute is replaced.

Example

Add 100:2, 10.1.1.1:22, 100.100:100 to the VPN route-target extended community attribute of BGP.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply extcommunity rt 100:2 rt 10.1.1.1:22 rt 100.100:100 additive
```

7.9.12 apply ip-address next-hop (Route-Policy view)

Function

The **apply ip-address next-hop** command sets the action for changing the next hop address of BGP routes in a routing policy.

The **undo apply ip-address next-hop** command restores the default setting.

By default, the action for changing the next hop address of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply ip-address next-hop { *ipv4-address* | **peer-address** }

undo apply ip-address next-hop { *ipv4-address* | **peer-address** }

Parameters

Parameter	Description	Value
ipv4-address	Specifies the next hop address.	It is in dotted decimal notation.
peer-address	Sets the next hop address to the local address when the **apply** clause is used by an export policy. Sets the next hop address to the peer address when the **apply** clause is used by an import policy.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To change the next hop address of BGP routes for selecting the optimal route, you can apply a routing policy containing the **apply ip-address next-hop** command.

The next hop address of a BGP route is set using the policy in the following situations:

- IBGP: Configure the import or export policy for the IBGP peer. If the next hop address configured in the routing policy is unreachable, the IBGP peer adds the corresponding route to the BGP routing table. However, this route is invalid.
- EBGP: Configure the import policy for the EBGP peer. If an export policy is configured, the route destined for the EBGP peer is discarded because the next hop address is unreachable.

Prerequisites

The **apply ip-address next-hop** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it affects BGP route selection.

When a routing policy is specified in the **import-route** and **network** commands, the **apply ip-address next-hop** clause in the routing policy does not take effect.

The command sets a next hop IP address for the routes that match the relevant route-policy, which may change the service forwarding path. Therefore, exercise caution when running this command.

Example

Define an **apply** clause to set the next hop address as 192.168.1.8.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply ip-address next-hop 192.168.1.8
```

7.9.13 apply ip-precedence

Function

The **apply ip-precedence** command sets the QoS parameter *ip-precedence* for routes.

The **undo apply ip-precedence** command restores the configuration.

By default, no IP preference is set.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply ip-precedence *ip-precedence*

undo apply ip-precedence

Parameters

Parameter	Description	Value
ip-precedence	IP precedence	The value can be a preference value or a keyword: • The value is an integer ranging from 0 to 7. • The preference keyword can be Routine, Priority, Immediate, Flash, Flash-override, Critical, Internet, or Network. **Table 7-193** shows the relationship between preference values and keywords.

Table 7-193 Relationship between preference values and keywords

Value	Keyword
0	Routine
1	Priority
2	Immediate
3	Flash
4	Flash-override
5	Critical
6	Internet
7	Network

Views

Route-policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After receiving routes, a BGP route receiver matches the attributes of the BGP routes based on the import route-policy, sets the IP precedence, delivers the BGP routes together with the associated QoS parameters, and applies QoS traffic policies to the classified data. In this case, the BGP route receiver can apply QoS policies to the data sent to the destination network segment based on the IP precedence. This applies QoS policies in BGP.

Prerequisites

The **apply ip-precedence** command can be used only after the **route-policy** command is used.

Configuration Impact

If a route matches a route-policy, you can change the value of the Precedence field in the IP header. The Precedence field is the first three bits of the Type of Service (ToS) field in the IP header.

Precautions

If an integer is used to specify *ip-precedence*, the preference is saved as an integer in the configuration file. If a keyword is used to specify *ip-precedence*, the preference is saved as a keyword in the configuration file.

Example

Set the IP precedence in the route-policy named **test**.

```
<HUAWEI> system-view
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] apply ip-precedence internet
```

7.9.14 apply ipv6 backup-interface

Function

The **apply ipv6 backup-interface** command configures a backup outbound interface in a routing policy.

The **undo apply ipv6 backup-interface** command restores the default setting.

By default, the backup outbound interface is not configured in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply ipv6 backup-interface *interface-type interface-number*

undo apply ipv6 backup-interface

Parameters

Parameter	Description	Value
interface-type interface-number	Specifies the type and number of the backup outbound interface.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply ipv6 backup-interface** command is used in IPv6 FRR scenarios to configure a backup outbound interface for a route. After IPv6 FRR is enabled, data traffic can be quickly switched to the backup outbound interface if the primary link fails.

Prerequisites

if-match clauses can be used to configure matching rules such as IPv6 prefix lists, and ACLs before a backup outbound interface is configured.

Follow-up Procedure

Reference a configured route-policy in the **ipv6 frr (system view)** command or the **ipv6 frr (VPN instance view)** command to configure IPv6 FRR on a public network or VPN.

The **apply ipv6 backup-interface** command is usually used together with the **apply ipv6 backup-nexthop** command.

Precautions

For P2P links, a backup next hop is not necessary. For non-P2P links, a backup next hop is necessary.

Example

Configure the backup outbound interface and the backup next hop in the route-policy named **ipv6_frr_rp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy ipv6_frr_rp permit node 10
[HUAWEI-route-policy] apply ipv6 backup-interface vlanif10
[HUAWEI-route-policy] apply ipv6 backup-nexthop 2001:db8:1::1
```

Delete the configured backup outbound interface from the route-policy named **ipv6_frr_rp**.


```
<HUAWEI> system-view
[HUAWEI] route-policy ipv6_frr_rp permit node 10
[HUAWEI-route-policy] undo apply ipv6 backup-interface
```

7.9.15 apply ipv6 backup-nexthop

Function

The **apply ipv6 backup-nexthop** command configures a backup next hop in a routing policy.

The **undo apply ipv6 backup-nexthop** command deletes the configured backup next hop.

By default, the backup next hop is not configured in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply ipv6 backup-nexthop { *ipv6-address* | **auto** }

undo ipv6 apply backup-nexthop

Parameters

Parameter	Description	Value
ipv6-address	Specifies the IPv6 address of a backup next hop.	The address is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
auto	Automatically searches for the backup next hop.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply ipv6 backup-nexthop** command is used to configure a backup next hop for a route in IPv6 FRR and VPN FRR scenarios. After IPv6 FRR is enabled, data traffic can be quickly switched to the backup next hop if the primary link fails.

Prerequisites

if-match clauses can be used to set matching rules such as interfaces, IPv6 prefix lists, and ACLs before a backup next hop is configured.

Follow-up Procedure

The configured route-policy can be used in the **ipv6 frr (system view)** command or the **ipv6 frr (VPN instance view)** command that is run to configure IPv6 FRR for a public or private network. It can also be used in the **vpn frr** command that is run to enable VPN FRR.

In a VPN FRR scenario, you only need to run the **apply ipv6 backup-nexthop** command to configure a backup next hop.

In an IPv6 FRR scenario, you need to run both the **apply ipv6 backup-nexthop** and **apply ipv6 backup-interface** commands.

Precautions

For P2P links, a backup next hop is not necessary. For non-P2P links, a backup next hop is necessary.

Example

Configure the backup interface and the backup next hop in the route-policy named **ipv6_frr_rp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy ip_frr_rp permit node 10
[HUAWEI-route-policy] apply ipv6 backup-interface vlanif10
[HUAWEI-route-policy] apply ipv6 backup-nexthop 2001:db8:1::1
```

Delete the configured backup next hop from the route-policy named **ipv6_frr_rp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy ipv6_frr_rp permit node 10
[HUAWEI-route-policy] undo apply ipv6 backup-nexthop
```

7.9.16 apply ipv6 next-hop

Function

The **apply ipv6 next-hop** command sets the action for changing an IPv6 next hop address of a BGP route in a route-policy.

The **undo apply ipv6 next-hop** command restores the default setting.

By default, the action for changing the IPv6 next hop addresses of BGP routes are not configured in a route-policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply ipv6 next-hop { **peer-address** | *ipv6-address* }

undo apply ipv6 next-hop { **peer-address** | *ipv6-address* }

Parameters

Parameter	Description	Value
ipv6-address	Specifies the IPv6 next hop address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
peer-address	Specifies the peer address as the next hop.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply ipv6 next-hop** command configures an IPv6 next hop address for a BGP route.

In BGP, the next hop address of a route can be set through the route-policy in the following situations:

- IBGP

  For an IBGP peer, the configured inbound and outbound policies can take effect. If the next hop address configured in the policy is unreachable, the IBGP peer still adds the route to the BGP routing table, but the route is not valid.

- EBGP

  For an EBGP peer, when the policy is used to modify the next hop address of a route, the inbound policy is configured. If the outbound policy is configured, the route is discarded because its next hop is unreachable.

Prerequisites

The **apply ipv6 next-hop** command can be used only after the **route-policy** command is used.

Configuration Impact

After a BGP route matches a route-policy, you can change the IPv6 next hop address of the BGP route.

Precautions

When a route-policy is being applied in the **import-route** and **network** commands, the **apply ipv6 next-hop** clause in the route-policy does not take effect.

Example

Set FC00:0:0:6::1 as the next hop address.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply ipv6 next-hop fc00:0:0:6::1
```

7.9.17 apply isis

Function

The **apply isis** command sets the action for changing the level of routes imported to IS-IS in a routing policy.

The **undo apply isis** command restores the default setting.

By default, the action for changing the level of routes imported to IS-IS is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply isis { **level-1** | **level-1-2** | **level-2** }

undo apply isis

Parameters

Parameter	Description	Value
level-1	Indicates IS-IS Level-1 routes.	-
level-1-2	Indicates IS-IS Level-1 and Level-2 routes.	-
level-2	Indicates IS-IS Level-2 routes.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A large number of external routes can be imported to IS-IS, which causes extra burdens on IS-IS-enabled devices. To solve this problem, run the **apply isis** command to set the level of the routes to be imported to IS-IS.

Prerequisites

The **apply isis** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it affects route receiving and advertising in IS-IS.

Example

Set the level of the routes imported to IS-IS.

```
<HUAWEI> system-view
```



```
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply isis level-1
```

7.9.18 apply local-preference

Function

The **apply local-preference** command sets the action for changing the local preference of BGP routes in a routing policy.

The **undo apply local-preference** command restores the default setting.

By default, the action for changing the local preference of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply local-preference *preference*

undo apply local-preference

Parameters

Parameter	Description	Value
preference	Specifies the local preference of BGP routes.	The value is an integer ranging from 0 to 4294967295.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The Local-Pref attribute is a private attribute of BGP. The **apply local-preference** command sets only the local preference for BGP routes. The Local_Pref attribute is used to determine the optimal route when traffic leaves an AS, and its default value is 100. When a BGP router obtains multiple routes to the same destination address but with different next hops through IBGP peers, the route with the largest Local_Pref value is selected.

Prerequisites

After a BGP route matches a routing policy, you can change the local preference of the BGP route.

Precautions

- When a routing policy takes effect, it affects BGP route selection.
- The Local_Pref attribute applies to the routing within an AS rather than be advertised to the outside of the AS. In this case, the **apply local-preference** command does not take effect when EBGP neighbor relationships are set up.
- After the **apply local-preference** command is run in a route-policy or its configuration is changed, route updates are triggered.

Example

Set the local preference of BGP routes to 130.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply local-preference 130
```

7.9.19 apply mpls-label

Function

The **apply mpls-label** command sets the action for allocating MPLS labels to public routes in a routing policy.

The **undo apply mpls-label** command restores the default setting.

By default, the action for allocating MPLS labels to public routes is not set in a routing policy.

NOTE

Only the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H support this command.

Format

apply mpls-label

undo apply mpls-label

Parameters

None

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In the scenario where inter-AS VPN Option C or Carrier Support Carrier (CSC) is deployed, you can use the **apply mpls-label** command to allocate labels to public routes.

Prerequisites

The **apply mpls-label** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it allocates MPLS labels to public routes.

Example

Assign MPLS labels to the routes that match the routing policy.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply mpls-label
```

7.9.20 apply origin

Function

The **apply origin** command sets the action for changing the Origin attribute of BGP routes in a routing policy.

The **undo apply origin** command restores the default setting.

By default, the action for changing the Origin attribute of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply origin { **egp** { *as-number-plain* | *as-number-dot* } | **igp** | **incomplete** }

undo apply origin

Parameters

Parameter	Description	Value
egp *as-number-plain*	Sets the origin of BGP routes as EGP. The parameter *as-number-plain* specifies the Integral AS number of an external route. An AS number uniquely identifies an AS. *as-number-plain* is required when you need to change the origin of BGP routes as EGP. EGP has the secondary highest priority. The Origin attribute of the routes obtained through EGP is EGP.	The value is an integer ranging from 1 to 4294967295.
egp *as-number-dot*	Sets the origin of BGP routes as EGP. The parameter *as-number-dot* specifies the AS number in dotted notation of an external route. An AS number uniquely identifies an AS. *as-number-dot* is required when you need to change the origin of BGP routes as EGP. EGP has the secondary highest priority. The Origin attribute of the routes obtained through EGP is EGP.	The value is in the format of *x.y*, where *x* and *y* are integers that range from 1 to 65535 and from 0 to 65535, respectively.
igp	Sets the origin of BGP routes as IGP. IGP has the highest priority. The Origin attribute of the routes obtained through an IGP of the AS that originates the routes, such as the routes imported to the BGP routing table through the **network** command, is IGP.	-

Parameter	Description	Value
incomplete	Sets the origin code of BGP routes as unknown. Incomplete has the lowest priority. The Origin attribute of the routes learned through other methods, such as the routes imported by BGP through the **import-route** command, is Incomplete.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To change the Origin attribute of routes for selecting the optimal route, you can apply a routing policy containing the **apply origin** command. The Origin attribute is a private attribute of BGP and defines the origin of a route.

Prerequisites

The **apply origin** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it affects BGP route selection.

Example

Set the origin of BGP routes to IGP.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply origin igp
```

7.9.21 apply ospf

Function

The **apply ospf** command sets the action performed for configuring an OSPF area to which the route is imported in a routing policy.

The **undo apply ospf** command restores the default setting.

By default, the action performed for configuring an OSPF area to which the route is imported is not set in a routing policy.

Format

apply ospf { **backbone** | **stub-area** }

undo apply ospf

Parameters

Parameter	Description	Value
backbone	Imports routes to the OSPF backbone area.	-
stub-area	Imports routes to an OSPF NSSA.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **apply ospf** command can be used to specify the OSPF backbone area or NSSA area to which routes are imported. This can prevent OSPF from importing too many external routes, which brings heavy burden on OSPF devices.

Prerequisites

The **apply ospf** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, routes are imported to the specified OSPF area.

Example

Import routes to the OSPF backbone area.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply ospf backbone
```

7.9.22 apply preference

Function

The **apply preference** command sets the action for changing the preference of routes in a routing policy.

The **undo apply preference** command restores the default setting.

By default, the action for changing the preference of routes is not set in a routing policy.

Format

apply preference *preference*

undo apply preference

Parameters

Parameter	Description	Value
preference	Specifies the route precedence. Route sharing and route selection are difficult because multiple routing protocols can run on the device at the same time; therefore, a default preference needs to be specified for each routing protocol. When different protocols discover multiple routes to the same destination, the route discovered by the protocol with a higher preference is selected to forward IP packets. The smaller the preference value, the higher the preference.	The value is an integer ranging from 1 to 255.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To change the preference of routes for selecting the optimal route, you can apply a routing policy containing the **apply preference** command.

Prerequisites

The **apply preference** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it affects route selection.

Example

Set the preference for routes.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply preference 90
```


7.9.23 apply preferred-value

Function

The **apply preferred-value** command sets the action for changing the preferred value of BGP routes in a routing policy.

The **undo apply preferred-value** command restores the default setting.

By default, the action for changing the preferred value of BGP routes is not set in a routing policy.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply preferred-value *preferred-value*

undo apply preferred-value

Parameters

Parameter	Description	Value
preferred-value	Specifies the preferred value of BGP routes. In route selection, the BGP route with the largest preferred value is preferred.	The value is an integer ranging from 0 to 65535.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To change the preferred value of BGP routes for selecting the optimal route, you can apply a routing policy containing the **apply preferred-value** command.

Prerequisites

The **apply preferred-value** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effect, it affects BGP route selection.

The preferred value of a route indicates the weight of the route in BGP routing. The preferred value is not a standard RFC-defined attribute and is valid only on local devices. The preferred value is inapplicable to export policies of BGP.

Example

Set the preferred value for BGP routes.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply preferred-value 66
```

7.9.24 apply qos-local-id

Function

The **apply qos-local-id** command sets the QoS local ID.

The **undo apply qos-local-id** command cancels the configuration.

By default, no QoS local ID is set.

Format

apply qos-local-id *qos-local-id*

undo apply qos-local-id

Parameters

Parameter	Description	Value
qos-local-id	Specifies the QoS local ID.	The value is an integer ranging from 1 to 4095.

Views

Route-policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The QoS local ID is a local identifier of QoS. In actual applications, you can set the QoS local ID in the route-policy, and add the command that matches the QoS local ID in the QoS policy. The QoS local ID set in the route-policy is delivered to the FIB table. During packet forwarding, the system obtains the QoS local ID from the FIB table and applies the related QoS policy according to the QoS local ID.

Configuration Impact

The **apply qos-local-id** command is mutually exclusive with the **apply behavior** and **apply ip-precedence** commands, and only one of these commands can be configured on a node of a routing policy. For example, if the **apply qos-local-id** command is configured in the view created by the **route-policy test permit node 10** command, configuring the **apply ip-precedence** command replaces **apply qos-local-id** command.

Example

Set the QoS local ID in the route-policy named **test**.

```
<HUAWEI> system-view
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] apply qos-local-id 10
```

7.9.25 apply tag

Function

The **apply tag** command sets the action for changing the tag of routes in a routing policy.

The **undo apply tag** command restores the default setting.

By default, the action for changing the tag of routes is not set in a routing policy.

Format

apply tag *tag*

undo apply tag

Parameters

Parameter	Description	Value
tag	Specifies the tag of routes. Routes can be tagged as required. You can set the same tag for the same type of route. Routes can be flexibly controlled and managed through tags in the routing policy.	The value is an integer ranging from 0 to 4294967295.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

To identify the routes, you can apply a routing policy containing the **apply tag** command to add the same tag to the matched routes.

Prerequisites

The **apply tag** command can be used only after the **route-policy** command is used.

Precautions

When a routing policy takes effects, routes will be matched by routing policies related to the tag.

BGP routes do not support tags. The **apply tag** command sets the tag for only IGP routes.

Example

Set the tag of routes to 100.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] apply tag 100
```

7.9.26 apply traffic-index

Function

The **apply traffic-index** command sets the BGP traffic index.

The **undo apply traffic-index** command cancels the configuration.

By default, no BGP traffic index is set.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.

Product	Support
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

apply traffic-index *traffic-index*

undo apply traffic-index

Parameters

Parameter	Description	Value
traffic-index	Specifies the index of BGP traffic.	The value is an integer ranging from 1 to 64.

Views

Route-policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

BGP accounting uses different BGP traffic indexes in BGP community attributes to identify routes and charge the traffic accordingly.

The sending end of BGP routes can set attributes for BGP routes by using the route-policy. The receiving end of BGP routes can set the BGP traffic index for BGP routes according to the BGP community filter, BGP AS_Path filter, ACL, and IP prefix list. The BGP traffic index together with routing information is delivered to the FIB table. After BGP accounting is enabled on an interface, the traffic-index-based traffic collection table can be generated for the interface.

During packet forwarding, traffic statistics can be collected according to the traffic index on each interface. Traffic statistics can be collected according to either the destination address in the inbound direction or the source address in the outbound direction.

Prerequisites

The **route-policy** command has been run.

Example

Configure the BGP traffic index.

```
<HUAWEI> system-view
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] apply traffic-index 10
```

7.9.27 description (Route-Policy view)

Function

The **description** command configures the description of a route-policy.

The **undo description** command deletes the description of a route-policy.

By default, no description is configured for the route-policy.

Format

description *text*

undo description

Parameters

Parameter	Description	Value
text	Specifies the description of a route-policy.	The description is a string of 1 to 80 case-sensitive characters that can contain spaces.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **description** command can be used to configure a description for a created route-policy. If many route-policies have been configured, configuring descriptions for the policies will facilitate policy management.

Prerequisites

A route-policy has been created by using **route-policy** command.

Example

Configure the description of the route-policy named **temp**.

```
<HUAWEI> system-view
[HUAWEI] route-policy temp permit node 10
[HUAWEI-route-policy] description This policy-name is temp
```

7.9.28 display ip as-path-filter

Function

display ip as-path-filter command displays the configuration of the AS_Path filter.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

display ip as-path-filter [*as-path-filter-number* | *as-path-filter-name*]

Parameters

Parameter	Description	Value
as-path-filter-number	Displays the configuration of an AS_Path filter with a specified number.	It is an integer that ranges from 1 to 256.
as-path-filter-name	Displays the configuration of an AS_Path filter with a specified name.	The name is a string of 1 to 51 characters without any space. It is case-sensitive.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The AS-Path attribute is a BGP-specific attribute. An AS-Path filter is used to filter BGP routes.

You can run the **display ip as-path-filter** command to:

- View detailed information about a configured AS path filter.
- Check whether an AS-Path filter is deleted successfully after running the **undo ip as-path-filter** command.

Precautions

The **display ip as-path-filter** command:

- Displays the configuration information about a specified AS-Path filter, if the number or name of the AS-Path filter is specified.
- Displays the configuration information about all AS-Path filters, if neither the number nor name of the AS-Path filter is specified.
- Does not display any information, if the AS-Path filter does not exist in the system or the AS-Path filter that is queried does not exist.

Example

Display the configured AS_Path filter.

```
<HUAWEI> display ip as-path-filter
As path filter number: 1
        permit    1.1 100,200
As path filter name: abc
        deny      2.2 200,400
```

Table 7-194 Description of the display ip as-path-filter command output

Item	Description
As path filter number	AS-Path filter number.
As path filter name	AS-Path filter name.
permit	Matching mode is permit.
1.1 100,200	Content of the regular expression.
deny	Matching mode is deny.

7.9.29 display ip community-filter

Function

The **display ip community-filter** command displays the configuration of the community filter.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

display ip community-filter [*basic-comm-filter-num* | *adv-comm-filter-num* | *comm-filter-name*]

Parameters

Parameter	Description	Value
basic-comm-filter-num	Displays the configuration of a basic community filter with a specified number.	The value is an integer ranging from 1 to 99.
adv-comm-filter-num	Displays the configuration of an advanced community filter with a specified number.	The value is an integer ranging from 100 to 199.
comm-filter-name	Displays the configuration of a community filter with a specified name.	The name is a string of 1 to 51 characters. The string cannot be all numerals.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The community attribute is a BGP-specific attribute. A community filter is used to filter BGP routes.

You can run the **display ip community-filter** command to:

- View detailed information about a configured community filter.
- Check whether a community filter is successfully deleted after running the **undo ip community-filter** command.

Precautions

The **display ip community-filter** command:

- Displays the configuration information about a specified community filter, if the number or name of the community filter is specified.
- Displays the configuration information about all community filters, if neither the number nor name of the community filter is specified.
- Does not display any information, if the community filter does not exist in the system or the community filter that is queried does not exist.

Example

Display all community filters.

```
<HUAWEI> display ip community-filter
Community filter Number: 10
        deny  no-export
Community filter Number: 110
        permit 110:110
Named Community basic filter: aa (ListID = 200)
        permit  1 internet
Named Community advanced filter: bb (ListID = 700)
        permit ^20
```

Table 7-195 Description of the display ip community-filter command output

Item	Description
Community filter Number	Indicates the number of a community filter.
permit	Indicates that the matching mode is permit.
deny	Indicates that the matching mode is deny.
Named Community basic filter	Indicates the name of a basic community filter.
Named Community advanced filter	Indicates the name of an advanced community filter.

7.9.30 display ip extcommunity-filter

Function

display ip extcommunity-filter command displays the configuration of the extended community filter.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

display ip extcommunity-filter [*basic-extcomm-filter-num* | *advanced-extcomm-filter-num* | *extcomm-filter-name*]

Parameters

Parameter	Description	Value
basic-extcomm-filter-num	Specifies the basic extended community filter number.	It is an integer that ranges from 1 to 199.
advanced-extcomm-filter-num	Specifies the advanced extended community filter number.	It is an integer that ranges from 200 to 399.
extcomm-filter-name	Displays the configuration of an extended community filter with a specified name.	The name is a string of 1 to 51 characters without any space. It is case-sensitive.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The extended community attribute is a BGP-specific attribute. An extended community filter is used to filter VPN routes.

You can run the **display ip extcommunity-filter** command to:

- View detailed information about a configured extended community filter.

- Check whether an extended community filter is successfully deleted after running the **undo ip excommunity-filter** command.

Precautions

The **display ip extcommunity-filter** command:

- Displays the configuration information about a specified extended community filter, if the number or name of the extended community filter is specified.
- Displays the configuration information about all extended community filters, if neither the number nor name of the extended community filter is specified.
- Does not display any information, if the extended community filter does not exist in the system or the extended community filter that is queried does not exist.

Example

Display information about the extended community filter.

```
<HUAWEI> display ip extcommunity-filter
Extended Community filter Number 10
      permit rt : 100:10
Extended Community filter Number 280
      permit rt 100:65
Extended Community filter basic filter: bas-abc
      permit rt : 200:10
Extended Community filter advanced filter: adv-abc
      deny 1.1.1.1:10
```

Table 7-196 Description of the display ip extcommunity-filter command output

Item	Description
Extended Community filter Number	Indicates the number of an extended community filter.
Extended Community filter basic filter	Basic extended community filter name.
Extended Community filter advanced filter	Advanced extended community filter name.
permit	Indicates that the matching mode is permit.
deny	Indicates that the matching mode is deny.
rt	Indicates the extended community attribute of the specified RT.

7.9.31 display ip ip-prefix

Function

The **display ip ip-prefix** command displays the configuration of IPv4 prefix lists.

Format

display ip ip-prefix [*ip-prefix-name*]

Parameters

Parameter	Description	Value
ip-prefix-name	Displays the configuration of an IP prefix list with a specified name.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

An IPv4 prefix list is used to filter IPv4 addresses. To achieve the following purposes, run the **display ip ip-prefix** command:

- View detailed configuration of a configured IPv4 prefix list.
- Check whether an IPv4 prefix list is deleted after running the **undo ip ip-prefix** command.
- View the number of routes that do or do not match the route-policy in an IPv4 prefix list.

Precautions

The **display ip ip-prefix** command:

- Displays the configuration of a specified IPv4 prefix list if the name of the IPv4 prefix list is specified.
- Displays the configuration of all IPv4 prefix lists if no IPv4 prefix list name is specified.
- Does not display information if no IPv4 prefix list exists in the system or the queried IPv4 prefix list does not exist.

Before collecting the number of routes that do or do not match the route-policy in an IPv4 prefix list within a certain period, run the **reset ip ip-prefix** command to clear existing statistics.

NOTE

If **The specified filter list does not exist** is displayed in the command output, the specified IPv4 prefix list failed to be configured. To re-configure it, run the **ip ip-prefix** command in the system view.

Example

Display the configuration of the IP prefix list named **p1**.

```
<HUAWEI> display ip ip-prefix p1
Prefix-list pl
Permitted 0
  Description prefixok
Denied 0
       index: 10      permit  192.168.0.0/16         ge  17  le  18
```

Table 7-197 Description of the display ip ip-prefix command output

Item	Description
Prefix-list	Name of an IPv4 prefix list.
Permitted	Number of routes that match a route-policy.
Description	Description of an IPv4 prefix list. This field is displayed only after a description is configured using the **ip ip-prefix** *ip-prefix-name* **description** *text* command.
Denied	Number of routes that do not match the route-policy.
index	Index of the entry in the IPv4 prefix list.
permit	Contents of the entry in the IPv4 prefix list.
ge 17	The mask is greater than or equal to 17.
le 18	The mask is less than or equal to 18.

7.9.32 display ip ipv6-prefix

Function

display ip ipv6-prefix displays the configuration of IPv6 prefix lists.

Format

display ip ipv6-prefix [*ipv6-prefix-name*]

Parameters

Parameter	Description	Value
ipv6-prefix-name	Displays the configuration of an IP prefix list with a specified name. If *ipv6-prefix-name* is not specified, the configuration of all the configured IPv6 prefix lists is displayed.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

An IPv6 prefix list is used to filter IPv6 addresses. To achieve the following purposes, run the **display ip ipv6-prefix** command:

- View detailed configuration of a configured IPv6 prefix list.
- Check whether an IPv6 prefix list is deleted after running the **undo ip ipv6-prefix** command.
- View the number of routes that do or do not match the route-policy in an IPv6 prefix list.

Precautions

The **display ip ipv6-prefix** command:

- Displays the configuration of a specified IPv6 prefix list if the name of the IPv6 prefix list is specified.
- Displays the configuration of all IPv6 prefix lists if no IPv6 prefix list name is specified.
- Does not display information if no IPv6 prefix list exists in the system or the queried IPv6 prefix list does not exist.

Before collecting the number of routes that do or do not match the route-policy in an IPv6 prefix list within a certain period, run the **reset ip ipv6-prefix** command to clear existing statistics.

Example

Display the configuration of all the IPv6 prefix lists.

```
<HUAWEI> display ip ipv6-prefix
 Prefix-list6 abc
 Description prefixok
 Permitted 0
 Denied 0
 index:   10          permit  ::/0
 index:   20          permit  ::/1          ge  1   le  128
```

Table 7-198 Description of the display ip ipv6-prefix command output

Item	Description
Prefix-list6	Name of an IPv6 prefix list.
Description	Description of an IPv6 prefix list. This field is displayed only after a description is configured using the **ip ipv6-prefix** *ipv6-prefix-name* **description** *text* command.

Item	Description
Permitted	Number of routes that match a route-policy.
Denied	Number of routes that do not match a route-policy.
index	Index of the entry in the IPv6 prefix list.
permit	Contents of the entry in the IPv6 prefix list.
ge	Greater than or equal to.
le	Less than or equal to.

7.9.33 display ip rd-filter

Function

The **display ip rd-filter** command displays the configuration of the route distinguisher (RD) filter.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

display ip rd-filter [*rd-filter-number*]

Parameters

Parameter	Description	Value
rd-filter-number	Displays the configuration of an RD filter with a specified number.	The value is an integer ranging from 1 to 255.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

The RD attribute is carried in VPN routes. An RD filter is used to filter VPN routes.

You can run the **display ip rd-filter** command to:

- View detailed information about a configured RD filter.
- Check whether an RD filter is successfully deleted after running the **undo ip rd-filter** command.

Precautions

The **display ip rd-filter** command:

- Displays the configuration information about a specified RD filter, if the number of an RD filter is specified.
- Displays the configuration information about all RD filters, if the number of no RD filter is specified.
- Does not display any information, if the RD filter does not exist in the system or the RD filter that is queried does not exist.

Example

Display the configured RD filter.

```
<HUAWEI> display ip rd-filter
Route Distinguisher Filter 1
      permit 10.1.1.1:1 10.2.2.2:* 100:1 200:*
Route Distinguisher Filter 2
      deny 1:1 2:2
      permit 1:* 2:*
```

Table 7-199 Description of the display ip rd-filter command output

Item	Description
Route Distinguisher Filter	Number of the RD filter
permit	Matching mode: permit
deny	Matching mode: deny

7.9.34 display route-policy

Function

The **display route-policy** command displays the configuration of the route-policy.

Format

display route-policy [*route-policy-name*]

Parameters

Parameter	Description	Value
route-policy-name	Displays the configuration of a route-policy with a specified name.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display route-policy** command to check detailed configuration of a route-policy.

Example

Display the configuration of the route-policy named **policy1**.

```
<HUAWEI> display route-policy policy1
Route-policy : policy1
  permit : 10 (matched counts: 2)
    Match clauses :
      if-match acl 2000
    Apply clauses :
      apply cost 100
      apply tag 100
```

Table 7-200 Description of the display route-policy command output

Item	Description
Route-policy	Name of the routing policy
permit	Matching mode and node index of the routing policy
matched counts: 2	Number of nodes that routes are matched in a routing policy
Match clauses	Matching condition list
Apply clauses	Apply clause list

7.9.35 goto next-node

Function

The **goto next-node** command further matches routes against a specified node after the routes match the current node.

The **undo goto next-node** command restores the default configuration.

By default, if a route matches the current node, it matches the route-policy and is no longer matched against other nodes.

Product	Support
S5731-S, S5731S-S, S5731-H, S5731S-H, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730S-H, and S6730-H	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S, S6730-S, and S6730S-S	Not supported

Format

goto next-node [*node*]

undo goto next-node

Parameters

Parameter	Description	Value
node	Specifies the index of a node against which routes are further matched.	The value is an integer ranging from 1 to 65535 and must be greater than the index of the current node.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The relationship among the matching rules of nodes in the same route-policy is OR. Specifically, if a route matches a node, it matches the route-policy and is no longer matched against other nodes. If you want the route to be matched against two or more nodes, run the **goto next-node** command so that the route is further matched against a specified node after the route matches the current node.

- If *node* is not specified in the command, the route will be further matched against the next node of the current node by default.
- If the node specified in the command does not exist, the route will be further matched against the next node of the specified node by default. If the next node of the specified node does not exist either, the route fails to match the route-policy, and no **apply** clause will be applied to the route.

Precautions

- If the **goto next-node** command is run in the route-policy view and a route matches all the specified nodes, the **apply** clauses of these nodes will be applied to the route.
- If the route fails to match one node, the route is matched against the next node until it succeeds in matching a node, and then the **apply** clauses of the nodes that the route matches will be applied to the route. If the route fails to match all nodes, no **apply** clauses will be applied to the route.
- In the same route-policy, a maximum of 50 nodes can be specified using the command.
- In the same route-policy, if this command is configured to match routes against multiple nodes, a maximum of 10 AS numbers, 16 RT extended community attributes, 16 color extended community attributes, and 36 community attributes can be added to the routes that match the nodes; if multiple community filters are used in the same route-policy to delete community attributes from the routes that match filtering conditions, only the first community filter takes effect.

Example

Configure a route-policy named **test** to further match routes against node 20 after the routes match node 10.

```
<HUAWEI> system-view
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] if-match tag 123
[HUAWEI-route-policy] apply cost 10
[HUAWEI-route-policy] goto next-node 20
```

7.9.36 if-match acl (Route-Policy view)

Function

The **if-match acl** command sets a matching rule that is based on the Access Control List (ACL).

The **undo if-match acl** command deletes the matching rule based on the specified ACL.

By default, no matching rule based on the ACL is configured.

Format

if-match acl { *acl-number* | *acl-name* }

undo if-match acl { *acl-number* | *acl-name* }

Parameters

Parameter	Description	Value
acl-number	Specifies the number of a basic ACL.	The value is an integer ranging from 2000 to 2999.
acl-name	Specifies the name of a named ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can run the **if-match acl** command to set a matching rule based on the ACL to match IPv4 prefixes.

Prerequisites

The **if-match acl** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes using the ACL. Routes that match the ACL will be checked by other **if-match** clauses of this node. Routes that do not match the ACL will be checked by the next node.

An ACL name is a character string that starts with a letter. For example, **2a** is an invalid ACL name.

The **if-match acl** command and the **if-match ip-prefix** command are mutually exclusive. If you run the **if-match ip-prefix** command after running the **if-match acl** command, the configuration of the **if-match ip-prefix** command overrides the configuration of the **if-match acl** command.

For an ACL, when the **rule** command is used to configure a filtering rule, the filtering rule is effective only with the source address range that is specified by the **source** parameter and with the time period that is specified by the **time-range** parameter.

Example

Set a matching rule that is based on ACL 2000.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match acl 2000
```

7.9.37 if-match as-path-filter

Function

The **if-match as-path-filter** command creates a matching rule based on the AS_Path filter.

The **undo if-match as-path-filter** command deletes a matching rule based on the specified AS_Path filter.

By default, no matching rule based on the AS_Path filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

if-match as-path-filter { *as-path-filter-number* &<1-16> | *as-path-filter-name* }

undo if-match as-path-filter [*as-path-filter-number* &<1-16> | *as-path-filter-name*]

Parameters

Parameter	Description	Value
as-path-filter-number	Specifies the number of an AS_Path filter. A maximum of 16 AS_Path filters can be specified.	The value is an integer ranging from 1 to 256.

Parameter	Description	Value
as-path-filter-name	Specifies the name of the AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The value cannot contain only numerals. **NOTE** When double quotation marks are used around the string, spaces are allowed in the string.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The AS_Path attribute is the private attribute of BGP. The **if-match as-path-filter** command is applicable to only BGP routes. The **ip as-path-filter** command must be used to define an AS_Path filter so that the matching rule based on this AS_Path filter can take effect. For example:

- If the **if-match as-path-filter 1** command is used but AS_Path filter 1 is not configured, all routes are permitted, that is, all routes match the matching rule.
- If the **if-match as-path-filter 1** command after the **ip as-path-filter 1 permit *20** command is used, the BGP routes with the AS_Path attribute being 20 are permitted.

Multiple **if-match as-path-filter** clauses can be specified. The relationship between **if-match as-path-filter** clauses is "OR". The relationship between **if-match** clauses is "AND".

Prerequisites

Before running the **if-match as-path-filter** command, run the **ip as-path-filter** command to configure an AS_Path filter.

Precautions

The routing policy matches routes using the AS-Path filter. Routes that match the AS-Path filter will be checked by other **if-match** clauses of this node. Routes that do not match the AS-Path filter will be checked by the next node.

A maximum of 16 AS_Path filters can be specified. The relationship between these AS_Path filters is OR. Specifically, if a route matches one of these AS_Path filters, it matches the matching rules of the command.

Creating an AS_Path filter before it is referenced is recommended. By default, nonexistent AS_Path filters cannot be referenced using the command. If the **route-**

policy nonexistent-config-check disable command is run in the system view and a nonexistent AS_Path filter is referenced using the current command, all routes match the AS_Path filter.

Example

Configure AS_Path filter 2 to permit AS200 and AS300. Create a routing policy named **test**, and define AS_Path filter 2 in an **if-match** clause for node 10 of the routing policy.

```
<HUAWEI> system-view
[HUAWEI] ip as-path-filter 2 permit _200_300
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] if-match as-path-filter 2
```

7.9.38 if-match community-filter

Function

The **if-match community-filter** command creates a matching rule based on the community filter.

The **undo if-match community-filter** command deletes the matching rule based on the specified community filter.

By default, no matching rule based on the community filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

if-match community-filter { *basic-comm-filter-num* [**whole-match**] | *adv-comm-filter-num* } &<1-16>

if-match community-filter *comm-filter-name* [**whole-match**]

undo if-match community-filter [*basic-comm-filter-num* | *adv-comm-filter-num*] &<1-16>

undo if-match community-filter *comm-filter-name*

Parameters

Parameter	Description	Value
basic-comm-filter-num	Specifies the number of a basic community filter.	The value is an integer ranging from 1 to 99.
adv-comm-filter-num	Specifies the number of an advanced community filter.	The value is an integer ranging from 100 to 199.
comm-filter-name	Specifies the name of a community filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.
whole-match	Indicates complete matching. That is, all the communities in the command must be matched. Complete matching is valid only for the basic community filter.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The community attribute is a private attribute of BGP. The **if-match community-filter** command is applicable to only BGP routes. The **ip community-filter** command must be used to define a community filter so that the matching rule based on this community filter can take effect. For example:

- If the **if-match community-filter 1** command is used but community filter 1 is not configured, all routes are permitted, that is, all routes can match the matching rule.
- If the **if-match community-filter 1** command is used after the **ip community-filter 1 permit 1:1** command is used, the BGP routes with the community attribute being 1:1 are permitted.

Multiple **if-match community-filter** clauses can be specified. The relationship between **if-match community-filter** clauses is "OR". The relationship between **if-match** clauses is "AND".

Prerequisites

Before using the **if-match community-filter** command, you must use the **ip community-filter** command to configure a community filter.

The **if-match community-filter** command can be used only after a routing policy is configured.

Precautions

The routing policy matches routes using the community filter. Routes that match the community filter will be checked by other **if-match** clauses of this node. Routes that do not match the community filter will be checked by the next node.

A maximum of 16 community filters can be configured in the **if-match community-filter** command. The relationship between these community-filters is OR. Specifically, if a route matches one of these community-filters, it matches the matching rules of the command.

The parameter **whole-match** is valid only for its front community filter number. If multiple community filters are specified in the **if-match community-filter** command and packets are required to completely match each filter, you need to specify the parameter **whole-match** behind each community filter and it is valid to only the basic community filter.

The name of a community filter cannot be all numerals.

Creating a community attribute filter before it is referenced is recommended. By default, nonexistent community attribute filters cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent community attribute filter is referenced using the current command, all routes match the community filter.

Example

Set a matching rule that is based on the community filter 1.

```
<HUAWEI> system-view
[HUAWEI] ip community-filter 1 permit 100:200
[HUAWEI] route-policy test permit node 10
[HUAWEI-route-policy] if-match community-filter 1
```

Set the complete matching rule for community attribute filters 1 and 2.

```
<HUAWEI> system-view
[HUAWEI] route-policy test permit node 11
[HUAWEI-route-policy] if-match community-filter 1 whole-match 2 whole-match
```

Set a matching rule that is based on the community filter named **aa**.

```
<HUAWEI> system-view
[HUAWEI] route-policy test permit node 12
[HUAWEI-route-policy] if-match community-filter aa
```

7.9.39 if-match cost

Function

The **if-match cost** command creates a matching rule based on the route cost.

The **undo if-match cost** command deletes the matching rule based on the specified route cost.

By default, no matching rule based on the route cost is configured.

Format

if-match cost { *cost* | **greater-equal** *greater-equal-value* [**less-equal** *less-equal-value*] | **less-equal** *less-equal-value* }

undo if-match cost

Parameters

Parameter	Description	Value
cost	Specifies the route cost. Route costs can be changed to prevent routing loops.	The value is an integer ranging from 0 to 4294967295.
greater-equal *greater-equal-value*	Specifies the minimum value of route cost.	The value is an integer ranging from 0 to 4294967294.
less-equal *less-equal-value*	Specifies the maximum value of route cost. *less-equal-value* is demanded to be greater than *greater-equal-value*.	The value is an integer ranging from 1 to 4294967295.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can use the **if-match cost** command to configure a node to filter routes based on the route costs. After such a matching rule is configured, you can apply the **apply** clauses to change the attributes of the routes that match the matching rule.

Prerequisites

The **if-match cost** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes based on the route cost. Routes that match the route cost will be checked by other **if-match** clauses of this node. Routes that do not match the route cost will be checked by the next node.

Example

Match the route with the cost 8.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match cost 8
```

7.9.40 if-match extcommunity-filter

Function

The **if-match extcommunity-filter** command sets a matching rule that is based on the extended community filter.

The **undo if-match extcommunity-filter** command deletes the matching rule based on the specified extended community filter.

By default, no matching rule based on the extended community filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

if-match extcommunity-filter { { *basic-extcomm-filter-num* | *adv-extcomm-filter-num* } &<1-16> | *extcomm-filter-name* }

undo if-match extcommunity-filter [[*basic-extcomm-filter-num* | *adv-extcomm-filter-num*] &<1-16> | *extcomm-filter-name*]

Parameters

Parameter	Description	Value
basic-extcomm-filter-num	Specifies the number of a basic extended community filter.	It is an integer ranging from 1 to 199.
adv-extcomm-filter-num	Specifies the number of an advanced extended community filter.	It is an integer ranging from 200 to 399.

Parameter	Description	Value
extcomm-filter-name	Specifies the name of an extended community filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The extended community attributes help flexibly control the routing policy. You can use the **if-match extcommunity-filter** command to configure a node to filter routes based on the extended community filter.

The **if-match extcommunity-filter** command is applicable to only BGP routes and must work in conjunction with the **ip extcommunity-filter** command. For example:

- If the **if-match extcommunity-filter 1** command is used but the extended community filter 1 is not configured, all routes are permitted, that is, all routes can match the matching rule.
- If the **if-match extcommunity-filter 1** command is used after the **ip extcommunity-filter 1 permit rt 1:1** command is used, the BGP routes with the extended community attribute being 1:1 are permitted.

Multiple **if-match extcommunity-filter** clauses can be specified. The relationship between **if-match extcommunity-filter** clauses is "OR". The relationship between **if-match** clauses is "AND".

Prerequisites

Before using the **if-match extcommunity-filter** command, you must use the **ip extcommunity-filter** command to configure an extended community filter.

Precautions

The routing policy matches routes using the extended community filter. Routes that match the extended community filter will be checked by other **if-match** clauses of this node. Routes that do not match the extended community filter will be checked by the next node.

A maximum of 16 extended community filters can be configured in the **if-match extcommunity-filter** command. The relationship between these extended community filters is OR. Specifically, if a route matches one of these extended community filters, it matches the matching rules of the command.

Example

Define a rule to match the routes of the specified extended community filter.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match extcommunity-filter 100
```

7.9.41 if-match interface

Function

The **if-match interface** command creates a matching rule based on the outbound interface.

The **undo if-match interface** command deletes the matching rule based on the specified outbound interface.

By default, no matching rule based on the outbound interface is configured.

Format

if-match interface { *interface-type interface-number* } &<1-16>

undo if-match interface [*interface-type interface-number*] &<1-16>

Parameters

Parameter	Description	Value
interface-type interface-number	Specifies the type and number of the outbound interface. A maximum of 16 outbound interfaces can be specified in the **if-match interface** command.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **if-match interface** command is used to filter routes based on the outbound interfaces.

A maximum of 16 outbound interfaces can be configured in this command.

If a node contains multiple **if-match interface** clauses, the relationship between the **if-match interface** clauses is OR. If a node contains both **if-match interface** clauses and other **if-match** clauses with different matching rules, the relationship between the **if-match interface** clauses and other **if-match** clauses is AND. For

example, if a node contains **if-match interface GE0/0/1**, **if-match interface GE0/0/2**, and **if-match acl 2000** clauses, **if-match interface GE0/0/1** and **if-match interface GE0/0/2** are ORed, whereas **if-match interface GE0/0/1** and **if-match acl 2000** are ANDed.

Prerequisites

The **if-match interface** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes based on outbound interface information. Routes that match the outbound interface information will be checked by other **if-match** clauses of this node. Routes that do not match the outbound interface information will be checked by the next node.

Example

Define a rule to match the routes with the outbound interface VLANIF100.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match interface vlanif 100
```

7.9.42 if-match ip

Function

The **if-match ip** command creates a matching rule based on IP information.

The **undo if-match ip** command deletes the matching rule based on specified IP information.

By default, no matching rule based on IP information is configured.

Format

if-match ip { **next-hop** | **route-source** | **group-address** } { **acl** { *acl-number* | *acl-name* } | **ip-prefix** *ip-prefix-name* }

undo if-match ip { **next-hop** | **route-source** | **group-address** } [**acl** { *acl-number* | *acl-name* } | **ip-prefix** *ip-prefix-name*]

Parameters

Parameter	Description	Value
next-hop	Specifies the next hop address.	-
route-source	Specifies the source address of routes.	-
group-address	Indicates the IP address of the multicast group.	The value is in dotted decimal notation.

Parameter	Description	Value
acl	Indicates route filtering using the ACL.	-
acl-number	Specifies the number of a basic ACL.	The value is an integer ranging from 2000 to 2999.
acl-name	Specifies the name of a basic ACL.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ip-prefix *ip-prefix-name*	Specifies the name of an IP prefix list that is used to filter routes.	The value is a string of case-sensitive characters without space and ranges from 1 to 169.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An ACL or IP prefix must be configured before running the **if-match ip** command so that the matching rule can take effect. For example:

- If the **if-match ip next-hop ip-prefix aa** command is used but the IP prefix **aa** is not configured, all routes are permitted, that is, all routes match the matching rule. This rule also applies to ACL.
- If the **if-match ip next-hop ip-prefix aa** and **ip ip-prefix aa permit 10.1.1.1 32** commands are used, the routes with the next hop being 10.1.1.1 is permitted. This rule also applies to ACL.

Prerequisites

The **if-match ip** command can be used only after the **route-policy** command is used.

Before running the **if-match ip** command, configure an **ACL** or an **IP prefix**.

Precautions

- The routing policy matches routes based on the next hop address or source address. Routes that match the next hop address or source address will be checked by other **if-match** clauses of this node. Routes that do not match the next hop address or source address will be checked by the next node.
- If the next hop address or source address of a route to be filtered is 0.0.0.0, by default, the system considers the mask length as 0 and matches the route.

If the next hop address or source address of a route to be filtered is not 0.0.0.0, by default, the system considers the mask length as 32 and matches the route.

- When you run the **rule** command to configure a filtering rule in an ACL, only the **source** and **time-range** parameters are valid for the filtering rule.
- Creating an ACL before it is referenced is recommended. If a nonexistent ACL is referenced using the command, all routes match the ACL.
- Creating an IP prefix list before it is referenced is recommended. By default, nonexistent IP prefix lists cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list is referenced using the current command, all routes match the IP prefix list.

Example

Set an IP prefix list named **p1** to filter routes.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match ip next-hop ip-prefix p1
```

Set a rule that source addresses of routes match ACL 2000 to filter routes.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match ip route-source acl 2000
```

7.9.43 if-match ip-prefix

Function

The **if-match ip-prefix** command creates a matching rule based on the IP prefix list.

The **undo if-match ip-prefix** command deletes the matching rule based on the specified IP prefix list.

By default, no matching rule based on the IP prefix list is configured in the routing policy.

Format

if-match ip-prefix *ip-prefix-name*

undo if-match ip-prefix *ip-prefix-name*

Parameters

Parameter	Description	Value
ip-prefix-name	Specifies the name of an IP address prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The routing policy matches routes using the IP prefix list. Routes are either permitted or denied.

The **ip ip-prefix** command must be used so that the matching rule can take effect. For example:

- If the **if-match ip-prefix aa** command is used but the IP prefix **aa** is not configured, all routes are permitted, that is, all routes match the matching rule.
- If the **if-match ip-prefix aa** and **ip ip-prefix aa permit 10.1.1.1 32** commands are used, the routes with the IP prefix being 10.1.1.1 and mask being 32 are permitted.

Prerequisites

The **if-match ip-prefix** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes based on IP prefix information. Routes that match the IP prefix information will be checked by other **if-match** clauses of this node. Routes that do not match the IP prefix information will be checked by the next node.

The **if-match acl** and **if-match ip-prefix** commands cannot be used together in the same node of a routing policy, because the latest configuration will override the previous one.

Creating an IP prefix list before it is referenced is recommended. By default, nonexistent IP prefix lists cannot be referenced using the command. If the **route-policy nonexistent-config-check disable** command is run in the system view and a nonexistent IP prefix list is referenced using the current command, all routes match the IP prefix list.

Example

Set an IP prefix list named **p1** to filter routes.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match ip-prefix p1
```


7.9.44 if-match ipv6

Function

The **if-match ipv6** command creates a matching rule based on IPv6 information.

The **undo if-match ipv6** command deletes a specified matching rule based on IPv6 information.

By default, no matching rule based on IPv6 information is configured in a routing policy.

Format

if-match ipv6 { **address** | **next-hop** | **route-source** } **prefix-list** *ipv6-prefix-name*

undo if-match ipv6 { **address** | **next-hop** | **route-source** } **prefix-list** *ipv6-prefix-name*

Parameters

Parameter	Description	Value
address	Matches the destination addresses of IPv6 routes.	-
next-hop	Matches the next hops of IPv6 routes. **NOTE** The S500 series switches do not support this parameter.	-
route-source	Matches the source addresses of the advertised IPv6 routes.	-
prefix-list	Indicates the IP prefix list.	-
ipv6-prefix-name	Specifies the name of the IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

Route-policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **ip ipv6-prefix** command must be used to configure an IPv6 prefix so that the matching rule (based on the destination addresses, next hop addresses, or source addresses of IPv6 routes) configured using the **if-match ipv6** command can take effect. If the **ip ipv6-prefix** command is not configured, all routes are permitted.

Prerequisites

The **if-match ipv6** command can be configured only after the **route-policy** command is configured.

Before running the **if-match ipv6** command, you must run the **ip ipv6-prefix** command to configure an IPv6 prefix.

Precautions

- When you filter routes based on the destination addresses, next hop addresses, or source addresses of IPv6 routes, the routes that match the matching rule are permitted and the routes that do not match the matching rule are denied.
- If the next hop address or source address of a route to be filtered is 0::0, by default, the system matches the route and considers that its mask length is 0.

  If the next hop address or source address of a route to be filtered is not 0::0, by default, the system matches the route and considers that its mask length is 128.
- Creating an IPv6 prefix list before it is referenced is recommended. If a nonexistent IPv6 prefix list is referenced using the command, all routes match the IPv6 prefix list.

Example

Define an if-match clause to match the related IPv6 routing information.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match ipv6 address prefix-list p1
[HUAWEI-route-policy] if-match ipv6 next-hop prefix-list p1
[HUAWEI-route-policy] if-match ipv6 route-source prefix-list p1
```

7.9.45 if-match mpls-label

Function

The **if-match mpls-label** command creates a matching rule based on the MPLS label.

The **undo if-match mpls-label** command deletes the matching rule based on the specified MPLS label.

By default, no matching rule based on the MPLS label is configured.

NOTE

Only the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H support this command.

Format

if-match mpls-label

undo if-match mpls-label

Parameters

None

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In the scenario where inter-AS VPN Option C or Carrier Support Carrier (CSC) is deployed, you can use the **if-match mpls-label** command to allocate labels to public routes.

Prerequisites

The **if-match mpls-label** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes based on the MPLS label. Routes that match the MPLS label will be checked by other **if-match** clauses of this node. Routes that do not match the MPLS label will be checked by the next node.

Example

Assign MPLS labels to the routes that match the routing policy.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match mpls-label
```

7.9.46 if-match rd-filter

Function

The **if-match rd-filter** command creates a matching rule based on the RD filter.

The **undo if-match rd-filter** command deletes the matching rule based on the specified RD filter.

By default, no matching rule based on the RD filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

if-match rd-filter *rd-filter-number*

undo if-match rd-filter

Parameters

Parameter	Description	Value
rd-filter-number	Specifies the number of an RD filter.	The value is an integer ranging from 1 to 255.

Views

Route-policy view, Tunnel selector view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The **if-match rd-filter** command and the **ip rd-filter** command work together to filter routes based on RD attributes. For example:

- If **if-match rd-filter 1** is configured, but **rd-filter 1** is not configured, then all current routes will be permitted.
- If **if-match rd-filter 1** is configured, and **ip rd-filter 1 permit 1:1** has been configured, then routes with RD 1:1 will be permitted.

Prerequisites

The **if-match rd-filter** command must be run after the **route-policy** command is run.

Precautions

The routing policy matches routes using the RD filter. Routes that match the RD filter will be checked by other **if-match** clauses of this node. Routes that do not match the RD filter will be checked by the next node.

Example

Define a matching rule to match an RD filter.

```
<HUAWEI> system-view
[HUAWEI] route-policy abc permit node 10
[HUAWEI-route-policy] if-match rd-filter 1
```

7.9.47 if-match route-type

Function

The **if-match route-type** command sets a matching rule that is based on the route type.

The **undo if-match route-type** command deletes the matching rule based on the specified route type.

By default, no matching rule based on the route type is configured.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

if-match route-type { **external-type1** | **external-type1or2** | **external-type2** | **internal** | **is-is-level-1** | **is-is-level-2** | **nssa-external-type1** | **nssa-external-type1or2** | **nssa-external-type2** }

undo if-match route-type { **external-type1** | **external-type1or2** | **external-type2** | **internal** | **is-is-level-1** | **is-is-level-2** | **nssa-external-type1** | **nssa-external-type1or2** | **nssa-external-type2** }

Parameters

Parameter	Description	Value
external-type1	Indicates OSPF external Type 1 routes.	-
external-type1or2	Indicates OSPF external routes.	-
external-type2	Indicates OSPF external Type 2 routes.	-
internal	Indicates internal routes, including OSPF inter-area routes and intra-area routes.	-
is-is-level-1	Indicates IS-IS Level-1 routes.	-
is-is-level-2	Indicates IS-IS Level-2 routes.	-
nssa-external-type1	Indicates NSSA external Type 1 routes.	-
nssa-external-type1or2	Indicates NSSA external routes.	-
nssa-external-type2	Indicates NSSA external Type 2 routes.	-

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can run the **if-match route-type** command to filter OSPF or IS-IS routes based on the route type.

Multiple **if-match route-type** clauses can be specified. The relationship between **if-match route-type** clauses is "OR". The relationship between **if-match** clauses is "AND".

Prerequisites

The **if-match route-type** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes based on the route type. Routes that match the route type will be checked by other **if-match** clauses of this node. Routes that do not match the route type will be checked by the next node.

For the same node in a routing policy, if two **if-match route-type** clauses are the same, the latter **if-match route-type** will not override the previous **if-match route-type**. After the latter clause is configured, both clauses take effect simultaneously. The relationship between **if-match route-type** clauses is "OR". That is, the actions defined by **apply** clauses can be performed on a route as long as the route meets one of the matching rules. For example, if both the **if-match route-type is-is-level-1** and **if-match route-type external-type1or2** commands are configured on the same node of a route policy, both IS-IS Level-1 routes and OSPF external routes can match the route policy.

NOTE

external-type1or2 refers to **external-type1** or **external-type2**. For the same node in a route policy, configuring both the **if-match route-type external-type1** and **if-match route-type external-type2** is equivalent to configuring the **if-match route-type external-type1or2** command. The two operations generate the same configuration file.

Similarly, **nssa-external-type1or2** refers to **nssa-external-type1** or **nssa-external-type2**. For the same node in a route policy, configuring both the **if-match route-type nssa-external-type1** and **if-match route-type nssa-external-type2** commands is equivalent to configuring the **if-match route-type nssa-external-type1or2** command. The two operations generate the same configuration file.

Example

Define a rule to match the routes of the specified type.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match route-type nssa-external-type1
```

7.9.48 if-match tag

Function

The **if-match tag** command sets a matching rule that is based on the route tag.

The **undo if-match tag** command deletes the matching rule based on the specified route tag.

By default, no matching rule based on the route tag is configured.

Format

if-match tag *tag*

undo if-match tag

Parameters

Parameter	Description	Value
tag	Indicates the tag value. Route tags classify routes as required. The same type of routes has the same tags. Routes are managed and controlled based on the tag by using the routing policy.	The value is an integer ranging from 0 to 4294967295.

Views

Route-Policy view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can run the **if-match tag** command to filter routes based on the tags.

Prerequisites

The **if-match tag** command can be used only after the **route-policy** command is used.

Precautions

The routing policy matches routes based on the route tag. Routes that match the route tag will be checked by other **if-match** clauses of this node. Routes that do not match the route tag will be checked by the next node.

Example

Define a rule to match the OSPF routes with the tag value 8.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy permit node 10
[HUAWEI-route-policy] if-match tag 8
```

7.9.49 ip as-path-filter

Function

The **ip as-path-filter** command creates an AS_Path filter.

The **undo ip as-path-filter** command deletes a specified AS_Path filter.

By default, no AS_Path filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

ip as-path-filter { *as-path-filter-number* | *as-path-filter-name* } { **deny** | **permit** } *regular-expression*

undo ip as-path-filter { *as-path-filter-number* | *as-path-filter-name* } [{ **deny** | **permit** } *regular-expression*]

Parameters

Parameter	Description	Value
as-path-filter-number	Specifies the number of an AS_Path filter.	The value is an integer ranging from 1 to 256.
as-path-filter-name	Specifies the name of an AS_Path filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The string cannot be all numerals. When double quotation marks are used around the string, spaces are allowed in the string.
deny	Sets the matching mode of the AS_Path filter to deny.	-
permit	Sets the matching mode of the AS_Path filter to permit.	-
regular-expression	Specifies the AS_Path regular expression.	The value is a string of 1 to 255 characters, with spaces supported.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An AS_Path filter uses the regular expression to define matching rules. After an AS_Path filter is set, the RM module immediately instructs each protocol to apply the filter by default.

The AS_Path attribute is a private attribute of BGP, and is used to filter BGP routes.

- The filter can be directly applied by using a command such as **peer as-path-filter**.
- The filter can be used as a matching condition of a routing policy by using a command such as **if-match as-path-filter zz**.

Configuration Impact

Multiple rules (permit or deny) can be specified in a filter.

By default, AS_Path filters work in **deny** mode. If all matching rules in a filter are configured to work in **deny** mode, all routes are denied by the filter; to prevent this problem, configure one matching rule in **permit** mode after one or multiple matching rules in **deny** mode so that the routes except for those denied by preceding matching rules are permitted by the filter.

Before you run the **undo ip as-path-filter** command to delete an AS_Path filter that is referenced by another command, delete the reference configuration.

Follow-up Procedure

To view detailed configurations of the AS_Path filter, run the **display ip as-path-filter** command.

Example

Create the AS_Path filter with the sequence number being 1, and permit routes that begin with 10 in the AS_Path to pass.

```
<HUAWEI> system-view
[HUAWEI] ip as-path-filter 1 permit ^10_
```

Create the AS_Path filter 2, and permit routes that contain 20 in the AS_Path to pass through.

```
<HUAWEI> system-view
[HUAWEI] ip as-path-filter 2 permit _20_
```

Create the AS_Path filter 3, and prohibit routes that contain 30 in the AS_Path from passing through.

```
<HUAWEI> system-view
```



```
[HUAWEI] ip as-path-filter 3 deny _30_
[HUAWEI] ip as-path-filter 3 permit .*
```

7.9.50 ip community-filter

Function

The **ip community-filter** command creates a community filter.

The **undo ip community-filter** command deletes a community filter.

By default, no community filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

ip community-filter { **basic** *comm-filter-name* | *basic-comm-filter-num* } { **permit** | **deny** } [*community-number* | *aa:nn* | **internet** | **no-export-subconfed** | **no-advertise** | **no-export**] &<1-20>

ip community-filter { **advanced** *comm-filter-name* | *adv-comm-filter-num* } { **permit** | **deny** } *regular-expression*

undo ip community-filter { **basic** *comm-filter-name* | *basic-comm-filter-num* } [**permit** | **deny**] [*community-number* | *aa:nn* | **internet** | **no-export-subconfed** | **no-advertise** | **no-export**] &<1-20>

undo ip community-filter { **advanced** *comm-filter-name* | *adv-comm-filter-num* } [**permit** | **deny**] [*regular-expression*]

Parameters

Parameter	Description	Value
basic *comm-filter-name*	Specifies the name of a basic community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits. **NOTE** When double quotation marks are used around the string, spaces are allowed in the string.
basic-comm-filter-num	Specifies the number of a basic community filter.	The value is an integer ranging from 1 to 99.
deny	Sets the matching mode of the community filter to deny.	-
permit	Sets the matching mode of the community filter to permit.	-
community-number	Specifies the community number.	The value is an integer ranging from 0 to 4294967295.

Parameter	Description	Value
aa:nn	Specifies the community number. You can configure a maximum of 20 community numbers once. ● If you do not configure any one of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 20 *community-number* and *aa:nn* together. ● If you configure one of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 19 *community-number* and *aa:nn* together. ● If you configure two of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 18 *community-number* and *aa:nn* together. ● If you configure three of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 17 *community-number* and *aa:nn* together. ● If you configure all of **internet**, **no-export-subconfed**, **no-advertise**, and **no-export**, you can specify 16 *community-number* and *aa:nn* together.	*aa* and *nn* are integers ranging from 0 to 65535.
internet	Indicates that the matching routes can be sent to any peer.	-
no-export-subconfed	Indicates that routes are not advertised outside an AS. If an AS confederation is used, routes are not advertised to any other sub-ASs in the AS confederation.	-
no-advertise	Indicates that routes are not advertised to other peers.	-
no-export	Indicates that routes are not advertised outside an AS. If an AS confederation is used, routes are not advertised outside the AS confederation, but to other sub-ASs.	-

Parameter	Description	Value
advanced *comm-filter-name*	Specifies the name of an advanced community filter.	The value is a string of 1 to 51 case-sensitive characters. The string cannot be all digits. **NOTE** When double quotation marks are used around the string, spaces are allowed in the string.
adv-comm-filter-num	Specifies the number of an advanced community filter.	The value is an integer ranging from 100 to 199.
regular-expression	Specifies the regular expression used to match the community information.	The value is a string of 1 to 255 case-sensitive characters, with spaces supported.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The community attribute is a private attribute of BGP, and can be used only to filter BGP routes. The community attribute can be used as a matching rule of a routing policy by using the **ip community-filter** and **if-match community-filter** commands together.

Precautions

Only the community number or known community attribute can be specified for a basic community filter. The regular expression can be used as a matching rule in an advanced community filter.

- The **ip community-filter basic** *comm-filter-name* command or the **ip community-filter** *basic-comm-filter-num* command can be used to configure a basic community filter. **basic** *comm-filter-name* specifies the name of a basic community filter, and the name cannot be all digits. A maximum of 20 community numbers can be configured in one command. *basic-comm-filter-num* specifies only the basic community filter with the number ranging from 1 to 99. A maximum of 20 community numbers can be configured in one command.

- The **ip community-filter advanced** *comm-filter-name* command or the **ip community-filter** *adv-comm-filter-num* command can be used to configure an advanced community filter. **advanced** *comm-filter-name* specifies the name of an advanced community filter, and the name cannot be all digits. *adv-comm-filter-num* specifies only the advanced community filter with the number ranging from 100 to 199.

The relationship between the rules of the community filter is "AND". This is different from the route distinguisher (RD) filter. This is because each route has only one RD but can have multiple communities.

For example, the community filters in the following formats have different matching results:

Format 1:

```
ip community-filter 1 permit 100:1 200:1 300:1
```

Format 2:

```
ip community-filter 1 permit 100:1
ip community-filter 1 permit 200:1 300:1
```

In the preceding configuration of the community filter, the community defined in each rule must be a sub-set of route communities so that the rule can be matched.

The RD filters in the following formats have the same matching results:

Format 1:

```
ip rd-filter 100 permit 100:1 200:1 2.2.2.2:1 3.3.3.3:1
```

Format 2:

```
ip rd-filter 100 permit 100:1 200:1
ip rd-filter 100 permit 2.2.2.2:1
ip rd-filter 100 permit 3.3.3.3:1
```

The **apply comm-filter delete** command run in the Route-Policy view deletes the specified community attribute from routes. An **ip community-filter** command can be used to specify community attributes but one such command specifies only one community attribute each time. To delete more than one community attribute, run the **ip community-filter** command multiple times. If multiple community attributes are specified in one filter, none of them can be deleted. For information about examples, see **apply comm-filter delete**.

By default, Community filters work in **deny** mode. If all matching rules in a filter are configured to work in **deny** mode, all routes are denied by the filter; to prevent this problem, configure one matching rule in **permit** mode after one or multiple matching rules in **deny** mode so that the routes except for those denied by preceding matching rules are permitted by the filter.

Before you run the **undo ip community-filter** command to delete a community attribute filter that is referenced by another command, delete the reference configuration.

Follow-up Procedure

By default, the Route Management (RM) module will instruct all protocols to apply this community filter. To delay the effective time, run the **route-policy-change notify-delay** command.

Run the **display ip community-filter** command to view detailed configuration for the community filter.

Example

Configure a basic community filter of which the sequence number is 1 to prevent matching routes from being advertised to any peer.

```
<HUAWEI> system-view
[HUAWEI] ip community-filter 1 deny internet
```

Configure an advanced community filter of which the sequence number is 100 to permit all the routes that match the AS 65001.

```
<HUAWEI> system-view
[HUAWEI] ip community-filter advanced 100 permit 65001:[0-9]+
```

7.9.51 ip extcommunity-filter

Function

The **ip extcommunity-filter** command creates an extended community filter.

The **undo ip extcommunity-filter** command deletes an extended community filter.

By default, no extended community filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

ip extcommunity-filter { *basic-extcomm-filter-num* | **basic** *basic-extcomm-filter-name* } { **deny** | **permit** } { **rt** { *as-number:nn* | *4as-number:nn* | *ipv4-address:nn* } } &<1-16>

ip extcommunity-filter { *advanced-extcomm-filter-num* | **advanced** *advanced-extcomm-filter-name* } { **deny** | **permit** } *regular-expression*

undo ip extcommunity-filter { *basic-extcomm-filter-num* | **basic** *basic-extcomm-filter-name* } [{ **deny** | **permit** } { **rt** { *as-number:nn* | *4as-number:nn* | *ipv4-address:nn* } } &<1-16>]

undo ip extcommunity-filter { *advanced-extcomm-filter-num* | **advanced** *advanced-extcomm-filter-name* } [*regular-expression*]

Parameters

Parameter	Description	Value
deny	Sets the matching mode of the extended community filter to deny.	-
permit	Sets the matching mode of the extended community filter to permit.	-
rt	Sets the extended community filter type to RT.	-
as-number	Specifies the AS number.	The value is an integer ranging from 0 to 65535.
4as-number	Specifies a 4-byte AS number.	A 4-byte AS number is divided into the following types: • It is an integer ranging from 65536 to 4294967295. • It is in the format of *x.y*, where *x* and *y* are integers that range from 0 to 65535.
ipv4-address	Specifies an IPv4 address.	The value is in dotted decimal notation.
nn	Specifies an integer.	• When the value of *as-number* is a 2-byte AS number, the value of *nn* ranges from 0 to 4294967295. • When the value of *4as-number* is a 4-byte AS number, the value of *nn* ranges from 0 to 65535. • For *ipv4-address*, the value of *nn* ranges from 0 to 65535.
basic-extcomm-filter-num	Specifies the number of a basic extended community filter.	The value is an integer ranging from 1 to 199.

Parameter	Description	Value
basic *basic-extcomm-filter-name*	Specifies the name of a basic extended community filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The value cannot contain only numerals. When double quotation marks are used around the string, spaces are allowed in the string.
advanced-extcomm-filter-num	Specifies the number of an advanced extended community filter.	The value is an integer ranging from 200 to 399.
advanced *advanced-extcomm-filter-name*	Specifies the name of an advanced extended community filter.	The name is a string of 1 to 51 case-sensitive characters without spaces. The value cannot contain only numerals. When double quotation marks are used around the string, spaces are allowed in the string.
regular-expression	Specifies the regular expression used to match the extended community information.	It is a string of 1 to 255 space-tolerant characters.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An extended community filter can be used as a matching condition of a route-policy by using a command such as **if-match extcommunity-filter zz**.

Only the extended community number can be specified for a basic extended community filter. The regular expression can be used as a matching rule in an advanced extended community filter.

- The **ip extcommunity-filter basic** *extcomm-filter-name* command or the **ip extcommunity-filter** *basic-extcomm-filter-num* command can be used to configure a basic extended community filter. **basic** *extcomm-filter-name* specifies the name of a basic extended community filter, and the name cannot be all digits. *basic-extcomm-filter-num* specifies only the basic extended community filter with the number ranging from 1 to 199. A

maximum of 16 extended community numbers can be configured using one command.

- The **ip extcommunity-filter advanced** *extcomm-filter-name* command or the **ip extcommunity-filter** *adv-extcomm-filter-num* command can be used to configure an advanced extended community filter. **advanced** *extcomm-filter-name* specifies the name of an advanced extended community filter, and the name cannot be all digits. *adv-extcomm-filter-num* specifies only the advanced extended community filter with the number ranging from 200 to 399.

The relationship between the rules of the extended community filter is "OR".

For example, the extended community filters in the following formats have the same matching results:

Format 1:

```
ip extcommunity-filter 1 permit rt 100:1 200:1 300:1
```

Format 2:

```
ip extcommunity-filter 1 permit rt 100:1
ip extcommunity-filter 1 permit rt 200:1 300:1
```

After the extended community filter is configured, if the policy application delay is set by using the **route-policy-change notify-delay** command, the Route Management (RM) module will instruct each protocol to apply this filter after the delay expires. By default, the RM module instructs each protocol to immediately apply this filter.

The **undo ip extcommunity-filter** command is used to delete a specified extended community filter.

The **display ip extcommunity-filter** command is used to display the detailed configurations of the extended community filter.

Configuration Impact

The **ip extcommunity-filter** command is used to filter routes based on the RT attributes of the routes. The routes that pass the filtering are permitted to pass through and the routes that fail to pass the filtering are denied.

Precautions

The extended community attributes of a route include VPN-target and Site-of-Origin (SoO). Only VPN-target, however, is supported by the policy.

By default, extended community filters work in **deny** mode. If all matching rules in a filter are configured to work in **deny** mode, all routes are denied by the filter; to prevent this problem, configure one matching rule in **permit** mode after one or multiple matching rules in **deny** mode so that the routes except for those denied by preceding matching rules are permitted by the filter.

Example

Configure an RT extended community filter of which the sequence number is 1.

```
<HUAWEI> system-view
[HUAWEI] ip extcommunity-filter 1 deny rt 200:200
```


7.9.52 ip ip-prefix

Function

The **ip ip-prefix** command creates an IPv4 prefix list or an entry in an IPv4 prefix list.

The **undo ip ip-prefix** command deletes an IPv4 prefix list or an entry from an IPv4 prefix list.

By default, no IPv4 prefix list is created.

Format

ip ip-prefix *ip-prefix-name* [**index** *index-number*] { **permit** | **deny** } *ipv4-address mask-length* [**match-network**] [**greater-equal** *greater-equal-value*] [**less-equal** *less-equal-value*]

undo ip ip-prefix *ip-prefix-name* [**index** *index-number*]

ip ip-prefix *ip-prefix-name* **description** *text*

undo ip ip-prefix *ip-prefix-name* **description** [*text*]

Parameters

Parameter	Description	Value
ip-prefix-name	Specifies the name of an IPv4 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
index *index-number*	Specifies the sequence number of an entry in the IPv4 prefix list.	The value is an integer that ranges from 1 to 4294967295. By default, the sequence number increases by 10 according to the configuration order, and the first sequence number is 10. **NOTE** A maximum of 65535 entries can be configured in an IP prefix list.

Parameter	Description	Value
permit	Specifies the matching mode of the IP prefix list as permit. In permit mode, if the IP address to be filtered is within the defined prefix range, the IP address matches the routing policy and does not continue to match the next entry. Otherwise, the IP address continues to match the next entry.	-
deny	Specifies the matching mode of the IP prefix list as deny. In deny mode, if the IP address to be filtered is within the defined prefix range, the IP address fails to match the routing policy and cannot match the next entry. Otherwise, the IP address continues to match the next entry.	-
ipv4-address	Specifies an IP address.	The value is in dotted decimal notation.
mask-length	Specifies the mask length.	The value is an integer that ranges from 0 to 32.
match-network	Matches the network address. The **match-network** parameter can be configured only when the IP address generated after *ipv4-address* is ANDed with *mask-length* is 0.0.0.0. This parameter is mainly used to match routes with a specified network address. For example, the **ip ip-prefix prefix1 permit 0.0.0.0 8** command filters all routes with mask length 8, while the **ip ip-prefix prefix1 permit 0.0.0.0 8 match-network** command filters all routes to the IP address range from 0.0.0.1 to 0.255.255.255.	-
greater-equal *greater-equal-value*	Specifies the lower threshold of the mask length. If **greater-equal** *greater-equal-value* and **less-equal** *less-equal-value* are not specified, the value of *mask-length* is the mask length.	*greater-equal-value* must meet the following requirement: *mask-length* ≤ *greater-equal-value* ≤ *less-equal-value* ≤ 32. If **greater-equal** is configured, the mask ranges from *greater-equal-value* to 32.

Parameter	Description	Value
less-equal *less-equal-value*	Specifies the upper threshold of the mask length. If **greater-equal** *greater-equal-value* and **less-equal** *less-equal-value* are not specified, the value of *mask-length* is the mask length.	*less-equal-value* must meet the following requirement: *mask-length* ≤ *greater-equal-value* ≤ *less-equal-value* ≤ 32. If **less-equal** is configured, the mask ranges from *mask-length* to *less-equal-value*.
description *text*	Specifies the description of the IP prefix list.	The value is a string of 1 to 80 case-sensitive characters without spaces. If the string is enclosed within double quotation marks ("), the string can contain spaces.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

An IP prefix list can be used as a filter or as matching conditions of a routing policy when it is used together with the **if-match** command.

Each entry in an IP prefix list can be used as a filtering rule. When a route to be filtered matches an entry, whether the route matches the IP prefix list is determined by the matching mode. A route to be filtered matches an entry or entries based on the following rules:

- Sequential matching: The route has to match the entries in the IP prefix list in ascending order of their *index-number* values. Therefore, specifying *index-number* in a required sequence is recommended.
- One-time matching: If a route matches one entry, the route matches the IP prefix list and will not be matched against the next entry.
- Matching failure by default: If a route fails to match any of the entries, it fails to match the IP prefix list.

The following example shows how different IP prefix lists take effect on the routes 1.1.1.1/24, 1.1.1.1/32, 1.1.1.1/26, 2.2.2.2/24, and 1.1.1.2/16.

Table 7-201 Matching results of IP prefix lists

Case	Commands	Matching result	Note
1	ip ip-prefix aa index 10 permit 1.1.1.1 24	Only the route 1.1.1.1/24 is permitted, and the other routes are denied.	This is a single-node accurate matching case, which indicates that only the route whose destination IP address and mask are the same as those specified by the entry meets the matching conditions. In addition, **permit** is configured as the matching mode. Therefore, the route 1.1.1.1/24 is permitted, and other routes are denied because they fail to meet the matching conditions.
2	ip ip-prefix aa index 10 deny 1.1.1.1 24	All routes are denied.	This is also a single-node accurate matching case. **deny** is configured as the matching mode. Therefore, the route 1.1.1.1/24 is denied, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.
3	ip ip-prefix aa index 10 permit 1.1.1.1 24 less-equal 32	The routes 1.1.1.1/24, 1.1.1.1/32, and 1.1.1.1/26 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, and **less-equal** is set to 32. Therefore, the routes with 1.1.1.0 as the prefix and the mask ranging from 24 to 32 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.
4	ip ip-prefix aa index 10 permit 1.1.1.0 24 greater-equal 24 less-equal 32	The routes 1.1.1.1/24, 1.1.1.1/32, and 1.1.1.1/26 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, **greater-equal** is set to 24, and **less-equal** is set to 32. Therefore, the routes with 1.1.1.0 as the prefix and the mask ranging from 24 to 32 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions. This case is similar to case 3 in terms of the matching result.

Case	Commands	Matching result	Note
5	ip ip-prefix aa index 10 permit 1.1.1.1 24 greater-equal 26	The routes 1.1.1.1/32 and 1.1.1.1/26 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, and **greater-equal** is set to 26. Therefore, the routes with 1.1.1.0 as the prefix and the mask ranging from 26 to 32 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.
6	ip ip-prefix aa index 10 permit 1.1.1.1 24 greater-equal 26 less-equal 32	The routes 1.1.1.1/32 and 1.1.1.1/26 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, **greater-equal** is set to 26, and **less-equal** is set to 32. Therefore, the routes with 1.1.1.0 as the prefix and the mask ranging from 26 to 32 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions. This case is similar to case 5 in terms of the matching result.
7	ip ip-prefix aa index 10 deny 1.1.1.1 24 ip ip-prefix aa index 20 permit 1.1.1.1 32	The route 1.1.1.1/32 is permitted, and the other routes are denied.	This is a multi-node accurate matching case. **deny** is configured as the matching mode of the matching entry indexed 10, and therefore the route 1.1.1.1/24 is denied by the matching entry indexed 10 based on the rule of one-time matching. The route 1.1.1.1/32 fails to match the matching conditions, and it is then matched against the entry indexed 20 for which **permit** is configured as the matching mode. Consequently, the route 1.1.1.1/32 matches the matching conditions of the entry indexed 20. The other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.

Case	Commands	Matching result	Note
8	ip ip-prefix aa index 10 permit 0.0.0.0 8 less-equal 32	The routes 1.1.1.1/24, 1.1.1.1/32, 1.1.1.1/26, 2.2.2.2/24, and 1.1.1.2/16 are all permitted.	If the IP prefix is 0.0.0.0 and you specify a mask and a mask length range after this IP prefix, all routes with the mask length within the specified mask length range are denied or permitted, regardless of the mask. The mask length range is from 8 to 32, 0.0.0.0 is specified as the IP address, and **permit** is configured as the matching mode. Therefore, all routes with the mask length within the range are permitted.
9	ip ip-prefix aa index 10 deny 0.0.0.0 24 less-equal 32 ip ip-prefix aa index 20 permit 0.0.0.0 0 less-equal 32	The route 1.1.1.2/16 is permitted, and the other routes are denied.	Note: For the entry indexed 10, the mask length range is from 24 to 32, 0.0.0.0 is specified as the IP address, and **deny** is configured as the matching mode. Therefore, all routes with the mask length within the range are denied, and the route 1.1.1.2/16 that fails to match its matching conditions is then matched against the entry indexed 20. For the entry indexed 20, the mask length range is from 0 to 32, 0.0.0.0 is specified as the IP address, and **permit** is configured as the matching mode. Therefore, the route 1.1.1.2/16 is permitted by the entry indexed 20.
9	ip ip-prefix aa index 10 deny 2.2.2.2 24 ip ip-prefix aa index 20 permit 0.0.0.0 0 less-equal 32	All routes except the route 2.2.2.2/24 are permitted.	For the entry indexed 10, **deny** is configured as the matching mode. Therefore, the route 2.2.2.2/24 that matches its matching conditions is denied, and the other routes that fail to match the matching conditions are then matched against the entry indexed 20. For the entry indexed 20, the mask length range is from 0 to 32, 0.0.0.0 is specified as the IP address, and **permit** is configured as the matching mode. Therefore, all routes except the route 2.2.2.2/24 are permitted by the entry indexed 20.

Configuration Impact

If you create an entry whose *index-number* has existed in the same IP prefix list but has different filtering rules, the new entry overwrites the existing one.

Precautions

- Because of the matching failure by default, if one or more than one entry with **deny** as the matching mode is created, create an entry using the **ip ip-prefix** *ip-prefix-name* [**index** *index-number*] **permit 0.0.0.0 0 less-equal 32** command so that all IPv4 routes may match the IP prefix list.
- If *ipv4-address mask-length* is specified as **0.0.0.0 0**, only default routes are matched.
- If *ipv4-address mask-length* is set to **0.0.0.0 0 less-equal 32**, all routes are matched.
- Before you run the **undo ip ip-prefix** command to delete an IP prefix list that is referenced by another command, delete the reference configuration.
- After a configuration is delivered, the device checks the validity of the parameters in the configuration and processes these parameters. After the processing, the generated configuration is the result of the AND calculation between the specified *ipv4-address* and *mask-length*. For example, if the specified *ipv4-address* and *mask-length* are 1.1.1.1 and 24, respectively, the generated configuration is 1.1.1.0 24.

  If the *ipv4-address* in the generated configuration is 0.0.0.0, the configuration matches all IPv4 addresses. In this case, routes are filtered based on the following rules.

  NOTE

  If the specified *ipv4-address* is not 0.0.0.0, the *mask-length* must not be 0.

Table 7-202 Route filtering rules

Whether *greater-equal* and *less-equal* Exist in the Post-Processing Configuration	Condition	Matching Result	Example
Neither *greater-equal* nor *less-equal* exists.	The post-processing *ipv4-address* and *mask-length* are 0.0.0.0 and X (non-0 value), respectively.	Matches all routes with the mask length of X.	Pre-processing: ip ip-prefix aa index 10 permit 0.0.1.1 16 Post-processing: ip ip-prefix aa index 10 permit 0.0.0.0 16 Matching result: The routes with the mask length of 16 are permitted.
greater-equal exists, but *less-equal* does not.	The post-processing *ipv4-address* and *mask-length* are 0.0.0.0 and X (non-0 value), respectively.	Matches all the routes whose mask length is within the range from *greater-equal* to 32.	Pre-processing: ip ip-prefix aa index 10 permit 0.0.1.1 16 greater-equal 20 Post-processing: ip ip-prefix aa index 10 permit 0.0.0.0 16 greater-equal 20 less-equal 32 Matching result: The routes whose mask length is within the range from 20 to 32 are permitted.
greater-equal does not exist, but *less-equal* does.	The post-processing *ipv4-address* and *mask-length* are 0.0.0.0 and X (non-0 value), respectively.	Matches all the routes whose mask length is within the range from X to *less-equal*.	Pre-processing: ip ip-prefix aa index 10 permit 0.0.1.1 16 less-equal 30 Post-processing: ip ip-prefix aa index 10 permit 0.0.0.0 16 greater-equal 16 less-equal 30 Matching result: The routes whose mask length is within the range from 16 to 30 are permitted.

Whether *greater-equal* and *less-equal* Exist in the Post-Processing Configuration	Condition	Matching Result	Example
Both *greater-equal* and *less-equal* exist.	The post-processing *ipv4-address* and *mask-length* are 0.0.0.0 and X (non-0 value), respectively.	Matches all the routes whose mask length is within the range from *greater-equal* to *less-equal*.	Pre-processing: ip ip-prefix aa index 10 permit 0.0.1.1 16 greater-equal 20 less-equal 30 Post-processing: ip ip-prefix aa index 10 permit 0.0.0.0 16 greater-equal 20 less-equal 30 Matching result: The routes whose mask length is within the range from 20 to 30 are permitted.

Follow-up Procedure

In a scenario in which a routing policy is being modified, after an IP prefix is configured, the RM module notifies protocols of applying the changed routing policy immediately by default. However, in some cases, multiple commands need to be run to modify a routing policy. If other commands need to be run after an IP prefix is configured, protocols may apply the routing policy whose modification is not complete yet. To solve this problem, run the **route-policy-change notify-delay** command to configure a delay for protocols to apply the changed routing policy.

Example

Configure the IP prefix list named **p1** to permit only the routes with the mask length ranging from 17 to 18 on the network segment 10.0.0.0/8.

```
<HUAWEI> system-view
[HUAWEI] ip ip-prefix p1 permit 10.0.0.0 8 greater-equal 17 less-equal 18
```

Configure the IP prefix list named **p3** to deny the routes to the IP address ranging from 0.0.0.1 to 0.255.255.255.

```
<HUAWEI> system-view
[HUAWEI] ip ip-prefix p3 index 10 deny 0.0.0.0 8 match-network
[HUAWEI] ip ip-prefix p3 index 20 permit 0.0.0.0 0 less-equal 32
```


7.9.53 ip ipv6-prefix

Function

The **ip ipv6-prefix** command configures an IPv6 prefix list or an entry in an IPv6 prefix list.

The **undo ip ipv6-prefix** command deletes an IPv6 prefix list or an entry from an IPv6 prefix list.

By default, no IPv6 prefix list is created.

Format

ip ipv6-prefix *ipv6-prefix-name* [**index** *index-number*] { **deny** | **permit** } *ipv6-address prefix-length* [**match-network**] [**greater-equal** *greater-equal-value*] [**less-equal** *less-equal-value*]

undo ip ipv6-prefix *ipv6-prefix-name* [**index** *index-number*]

ip ipv6-prefix *ipv6-prefix-name* **description** *text*

undo ip ipv6-prefix *ipv6-prefix-name* **description** [*text*]

Parameters

Parameter	Description	Value
ipv6-prefix-name	Specifies the name of an IPv6 prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.
index *index-number*	Specifies the sequence number of an entry in the IPv6 prefix list.	The value is an integer that ranges from 1 to 4294967295. By default, the sequence number increases by 10 according to the configuration order, and the first sequence number is 10. **NOTE** A maximum of 65535 entries can be configured in an IPv6 prefix list.

Parameter	Description	Value
permit	Specifies the matching mode of the IPv6 prefix list as permit. In permit mode, if the IPv6 address to be filtered is within the defined prefix range, the IPv6 address matches the routing policy and does not continue to match the next entry. Otherwise, the IPv6 address continues to match the next entry.	-
deny	Specifies the matching mode of the IPv6 prefix list as deny. In deny mode, if the IPv6 address to be filtered is within the defined prefix range, the IPv6 address fails to match the routing policy and cannot match the next entry. Otherwise, the IPv6 address continues to match the next entry.	-
ipv6-address	Specifies the IPv6 prefix range in the form of an IPv6 address. If :: is specified, the address 0::0 is matched.	-
prefix-length	Specifies the IPv6 prefix range using the mask length.	The value is an integer that ranges from 0 to 128. If ::0 less-equal 128 is used, all IPv6 addresses are matched.
match-network	Matches the network address. The **match-network** parameter can be configured only when the IP address generated after *ipv6-address* is ANDed with *prefix-length* is ::. For example, the **ip ipv6-prefix prefix1 permit :: 96** command filters all IPv6 routes with mask length 96, while the **ip ipv6-prefix prefix1 permit :: 96 match-network** command filters all routes to the IPv6 address range from ::1 to ::FFFF:FFFF.	-
greater-equal *greater-equal-value*	Specifies the lower threshold of the mask length.	*greater-equal-value* must meet the following requirement: *prefix-length* ≤ *greater-equal-value* ≤ *less-equal-value* ≤ 128.

Parameter	Description	Value
less-equal *less-equal-value*	Specifies the upper threshold of the mask length.	*less-equal-value* must meet the following requirement: *prefix-length* ≤ *greater-equal-value* ≤ *less-equal-value* ≤ 128.
description *text*	Specifies the description of the IPv6 prefix list.	The value is a string of 1 to 80 case-sensitive characters without spaces. If the string is enclosed within double quotation marks ("), the string can contain spaces.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The IPv6 prefix list can be used by the protocols as a prefix filter, or used with the **if-match ipv6** command as a matching condition of routing policies.

Each entry in the IPv6 prefix list can be considered as a filter rule. When a route to be filtered matches one of the entries, the route can be determined whether it is permitted the IPv6 prefix list based on the matching mode. Entries in the IPv6 prefix list can be matched with routes to be filtered based on the following rules:

- Sequence match: Each entry in the IP prefix list is matched in ascending order of the index number. When setting index numbers for entries, you can arrange your expected matching order.
- Unique match: When the route to be filtered matches one entry, the route stops to match the other entries.
- Default deny: By default, if routes to be filtered do not match any entry in the IP prefix list, the routes are denied by IP prefix list.

The following example shows how different IPv6 prefix lists take effect on the routes 1::1/96, 1::1/128, 1::1/100, 2::2/96, and 1::2/64.

Table 7-203 Matching results of IPv6 prefix lists

Case	Commands	Matching result	Note
1	ip ipv6-prefix aa index 10 permit 1::1 96	Only the route 1::1/96 is permitted, and the other routes are denied.	This is a single-node accurate matching case, which indicates that only the route whose destination IPv6 address and mask are the same as those specified by the entry meets the matching conditions. In addition, **permit** is configured as the matching mode. Therefore, the route 1::1/96 is permitted, and other routes are denied because they fail to meet the matching conditions.
2	ip ipv6-prefix aa index 10 deny 1::1 96	All routes are denied.	This is also a single-node accurate matching case. **deny** is configured as the matching mode. Therefore, the route 1::1/96 is denied, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.
3	ip ipv6-prefix aa index 10 permit 1::1 96 less-equal 128	The routes 1::1/96, 1::1/128, and 1::1/100 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, and **less-equal** is set to 128. Therefore, the routes with 1::1 as the prefix and the mask ranging from 96 to 128 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.
4	ip ipv6-prefix aa index 10 permit 1::1 96 greater-equal 96 less-equal 128	The routes 1::1/96, 1::1/128, and 1::1/100 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, **greater-equal** is set to 96, and **less-equal** is set to 128. Therefore, the routes with 1::1 as the prefix and the mask ranging from 96 to 128 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions. This case is similar to case 3 in terms of the matching result.

Case	Commands	Matching result	Note
5	ip ipv6-prefix aa index 10 permit 1::1 96 greater-equal 100	The routes 1::1/128 and 1::1/100 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, and **greater-equal** is set to 100. Therefore, the routes with 1::1 as the prefix and the mask ranging from 100 to 128 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.
6	ip ipv6-prefix aa index 10 permit 1::1 96 greater-equal 100 less-equal 128	The routes 1::1/128 and 1::1/100 are permitted, and the other routes are denied.	This is also a single-node accurate matching case. **permit** is configured as the matching mode, **greater-equal** is set to 100, and **less-equal** is set to 128. Therefore, the routes with 1::1 as the prefix and the mask ranging from 100 to 128 can be permitted, and the other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions. This case is similar to case 5 in terms of the matching result.
7	ip ipv6-prefix aa index 10 deny 1::1 96 ip ipv6-prefix aa index 20 permit 1::1 128	The route 1::1/128 is permitted, and the other routes are denied.	This is a multi-node accurate matching case. **deny** is configured as the matching mode of the matching entry indexed 10, and therefore the route 1::1/96 is denied by the matching entry indexed 10 based on the rule of one-time matching. The route 1::1/128 fails to match the matching conditions, and it is then matched against the entry indexed 20 for which **permit** is configured as the matching mode. Consequently, the route 1::1/128 matches the matching conditions of the entry indexed 20. The other routes are denied based on the rule of matching failure by default because they fail to meet the matching conditions.

Case	Commands	Matching result	Note
8	ip ipv6-prefix aa index 10 permit :: 64 less-equal 128	The routes 1::1/96, 1::1/128, 1::1/100, 2::2/96, and 1::2/64 are all permitted.	If the IPv6 prefix is :: and you specify a mask and a mask length range after this IPv6 prefix, all routes with the mask length within the specified mask length range are denied or permitted, regardless of the mask. The mask length range is from 64 to 128, :: is specified as the IPv6 address, and **permit** is configured as the matching mode. Therefore, all routes with the mask length within the range are permitted.
9	ip ipv6-prefix aa index 10 deny :: 96 less-equal 128 ip ipv6-prefix aa index 20 permit :: 0 less-equal 128	The route 1::2/64 is permitted, and the other routes are denied.	For the entry indexed 10, the mask length range is from 96 to 128, :: is specified as the IPv6 address, and **deny** is configured as the matching mode. Therefore, all routes with the mask length within the range are denied, and the route 1::2/64 that fails to match its matching conditions is then matched against the entry indexed 20. For the entry indexed 20, the mask length range is from 0 to 128, :: is specified as the IPv6 address, and **permit** is configured as the matching mode. Therefore, the route 1::2/64 is permitted by the entry indexed 20.
10	ip ipv6-prefix aa index 10 deny 2::2 96 ip ipv6-prefix aa index 20 permit :: 0 less-equal 128	All routes except the route 2::2/96 are permitted.	Note: For the entry indexed 10, **deny** is configured as the matching mode. Therefore, the route 2::2/96 that matches its matching conditions is denied, and the other routes that fail to match the matching conditions are then matched against the entry indexed 20. For the entry indexed 20, the mask length range is from 0 to 128, :: is specified as the IPv6 address, and **permit** is configured as the matching mode. Therefore, all routes except the route 2::2/96 are permitted by the entry indexed 20.

Configuration Impact

If you create an entry whose index number is the same as an existing entry in the IPv6 prefix list, the created entry will replace the existing entry.

Precautions

- The IPv6 prefix list adopts default deny as the matching mode. If you have created one or multiple entries in **deny** mode, but no entry in the **permit** mode, you must create an entry **permit :: 0 less-equal 128** to permit IPv6 routes which do not match the entries in **deny** mode.
- If you specify *ipv6-address prefix-length* to be :: 0, only IPv6 default routes are matched.
- If *ipv6-address prefix-length* is set to **:: 0 less-equal 128**, all routes will be matched.
- Before you run the **undo ip ipv6-prefix** command to delete an IPv6 prefix list that is referenced by another command, delete the reference configuration.
- After a configuration is delivered, the device checks the validity of the parameters in the configuration and processes these parameters. After the processing, the generated configuration is the result of the AND calculation between the specified *ipv6-address* and *prefix-length*. For example, if the specified *ipv6-address* and *prefix-length* are 1::1 and 64, respectively, the generated configuration is 1:: 64.

  If the *ipv6-address* in the generated configuration is ::, the configuration matches all IPv6 addresses. In this case, routes are filtered based on the following rules.

  NOTE

  If the specified *ipv6-address* is not ::, the *prefix-length* must not be 0.

Table 7-204 Route filtering rules

Whether *greater-equal* and *less-equal* Exist in the Post-Processing Configuration	Condition	Matching Result	Example
Neither *greater-equal* nor *less-equal* exists.	The post-processing *ipv6-address* and *prefix-length* are :: and X (non-0 value), respectively.	Matches all IPv6 routes with the prefix length of X.	Pre-processing: ip ipv6-prefix aa index 10 permit ::1:1 96 Post-processing: ip ipv6-prefix aa index 10 permit :: 96 Matching result: The IPv6 routes with the prefix length of 96 are permitted.
greater-equal exists, but *less-equal* does not.	The post-processing *ipv6-address* and *prefix-length* are :: and X (non-0 value), respectively.	Matches all the IPv6 routes whose prefix length is within the range from *greater-equal* to 128.	Pre-processing: ip ipv6-prefix aa index 10 permit ::1:1 96 greater-equal 120 Post-processing: ip ipv6-prefix aa index 10 permit :: 96 greater-equal 120 less-equal 128 Matching result: The IPv6 routes whose prefix length is within the range from 120 to 128 are permitted.
greater-equal does not exist, but *less-equal* does.	The post-processing *ipv6-address* and *prefix-length* are :: and X (non-0 value), respectively.	Matches all the IPv6 routes whose prefix length is within the range from X to *less-equal*.	Pre-processing: ip ipv6-prefix aa index 10 permit ::1:1 96 less-equal 120 Post-processing: ip ipv6-prefix aa index 10 permit :: 96 greater-equal 96 less-equal 120 Matching result: The IPv6 routes whose prefix length is within the range from 96 to 120 are permitted.

Whether *greater-equal* and *less-equal* Exist in the Post-Processing Configuration	Condition	Matching Result	Example
Both *greater-equal* and *less-equal* exist.	The post-processing *ipv6-address* and *prefix-length* are :: and X (non-0 value), respectively.	Matches all the IPv6 routes whose prefix length is within the range from *greater-equal* to *less-equal*.	Pre-processing: ip ipv6-prefix aa index 10 permit ::1:1 96 greater-equal 120 less-equal 124 Post-processing: ip ipv6-prefix aa index 10 permit :: 96 greater-equal 120 less-equal 124 Matching result: The IPv6 routes whose prefix length is within the range from 120 to 124 are permitted.

Follow-up Procedure

In a scenario in which a used routing policy is being modified, after you configure the IPv6 prefix list, RM immediately notifies the protocols of re-applying the routing policy. However, you must run several commands to modify the routing policy. To prevent the protocols from repeatedly re-applying the routing policy which is being modified, you can run the **route-policy-change notify-delay** command to configure delay time for re-applying the routing policy, after you configure the IPv6 prefix list.

Example

Permit the routes with the mask length ranging from 32 to 64 bits.

```
<HUAWEI> system-view
[HUAWEI] ip ipv6-prefix abc permit :: 0 greater-equal 32 less-equal 64
```

Deny the routes with the IP prefix FC00:0:0:D00::/32 and with the prefix longer than 32 bits, and permit the other IPv6 routes.

```
<HUAWEI> system-view
[HUAWEI] ip ipv6-prefix abc deny fc00:0:0:d00:: 32 less-equal 128
[HUAWEI] ip ipv6-prefix abc permit :: 0 less-equal 128
```

Configure the IPv6 prefix list named **p3** to deny the routes to the IPv6 address ranging from ::1 to ::FFFF:FFFF.

```
<HUAWEI> system-view
[HUAWEI] ip ipv6-prefix p3 index 10 deny :: 96 match-network
```



```
[HUAWEI] ip ipv6-prefix p3 index 20 permit :: 0 less-equal 128
```

7.9.54 ip rd-filter

Function

The **ip rd-filter** command creates an RD filter.

The **undo ip rd-filter** command deletes an RD filter.

By default, no RD filter is configured.

Product	Support
S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported.
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported.

Format

ip rd-filter *rd-filter-number* { **deny** | **permit** } *route-distinguisher* &<1-10>

undo ip rd-filter *rd-filter-number* [{ **deny** | **permit** } *route-distinguisher* &<1-10>]

Parameters

Parameter	Description	Value
rd-filter-number	Specifies the number of an RD filter.	The value is an integer ranging from 1 to 255.
permit	Permits a route to match the rules if its RD matches the rules.	-
deny	Denied a route if its RD matches the rules.	-

Parameter	Description	Value
route-distinguisher	Specifies the RD to aa:nn or ipv4-address:nn. You can set a maximum of 10 RDs. The switch support RDs in the following formats: • ipv4-address:nn, such as 10.1.1.1:200 • aa:nn, such as 100:1 • aa.aa:nn, such as 100.100:1 • ipv4-address:* in the wildcard format, such as 10.1.1.1:*, indicating that the RD begins with 10.1.1.1 • aa:* in the wildcard format, such as 100:*, indicating that the RD begins with 100 • aa.aa:* in the wildcard format, such as 100.100:*, indicating that the RD begins with 100.100	• The IPv4 address is in dotted decimal notation. • The *nn* in ipv4-address:nn is an integer ranging from 0 to 65535. • In aa:nn, the *aa* is an integer ranging from 0 to 65535, and *nn* is an integer ranging from 0 to 4294967295. • The *aa* and *nn* in aa:*, aa.aa:*, and aa.aa:nn are both integers ranging from 0 to 65535.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

The **ip rd-filter** command is used together with the **if-match rd-filter** command. First use the **ip rd-filter** command to configure an RD filter, and use the **if-match rd-filter** command to configure a matching rule based on the RD filter in a routing policy. The routing policy is used to filter routes that are received and advertised.

The RD filter has the following rules:

- If the RD filter is not configured but is used to filter routes, the matching result is **permit**.

  For example, the RD filter 100 is not configured but is used by the routing policy:

  ```
  route-policy test permit node 10
  if-match rd-filter 100
  ```

  When the routing policy is used to filter routes, the routes match this **if-match** clause, and the routes match the node 10 in the routing policy named **test**.

- If the RD filter is configured but the RD of routes does not match any RD defined in the RD filter, the default matching result is **deny**.

  For example, the RD of routes is 100:1, and the configuration of the RD filter is as follows:

  ```
  ip rd-filter 100 permit 10.1.1.1:100
  ```

  When the RD filter is used to filter routes, the matching result is **deny**.
- The relationship between the rules of the RD filter is "OR". This is different from the community filter. This is because each route has only one RD but can have multiple communities.

  For example, the RD filters in the following formats have the same matching results:

  Format 1:

  ```
  ip rd-filter 100 permit 100:1 200:1 10.2.2.2:1 10.3.3.3:1
  ```

  Format 2:

  ```
  ip rd-filter 100 permit 100:1 200:1
  ip rd-filter 100 permit 10.2.2.2:1
  ip rd-filter 100 permit 10.3.3.3:1
  ```

  The community filters in the following formats have different matching results:

  Format 1:

  ```
  ip community-filter 1 permit 100:1 200:1 300:1
  ```

  Format 2:

  ```
  ip community-filter 1 permit 100:1
  ip community-filter 1 permit 200:1 300:1
  ```

  In the preceding configuration of the community filter, the community defined in each rule must be a sub-set of route communities so that the rule can be matched.
- Routes are filtered according to the configuration order of multiple rules. For example:

  ```
  ip rd-filter 100 deny 200:1 10.5.5.5:1
  ip rd-filter 100 permit 200:* 10.5.5.5:*
  ```

  In this situation, the route with the RD 200:1 or 5.5.5.5:1 is denied. If the configuration order of multiple rules is reversed as follows:

  ```
  ip rd-filter 100 permit 200:* 10.5.5.5:*
  ip rd-filter 100 deny 200:1 10.5.5.5:1
  ```

  In this situation, the route with the RD 200:1 or 10.5.5.5:1 is permitted.
- Each RD filter can be configured with a maximum of 255 rules.

Example

Configure an RD filter.

```
<HUAWEI> system-view
[HUAWEI] ip rd-filter 1 permit 100:1
```


7.9.55 reset ip ip-prefix

Function

The **reset ip ip-prefix** command resets the statistics of the specified IPv4 prefix list.

Format

reset ip ip-prefix [*ip-prefix-name*]

Parameters

Parameter	Description	Value
ip-prefix-name	Specifies the name of an IPv4 prefix list. If *ip-prefix-name* is not specified, you can reset the statistics of all the IPv4 prefix lists.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The IPv4 prefix list can be used to filter IPv4 addresses. When filtering IPv4 addresses, the system records the numbers of prefixes that are permitted and denied by the IPv4 prefix list. You can run the **display ip ip-prefix** command to view the numbers.

To view the number of IPv4 prefixes that are permitted and denied by the IPv4 prefix list, run the **reset ip ip-prefix** command to clear statistics about permitted and denied routes in the IPv4 prefix list, and then run the **display ip ip-prefix** command to display the number of IPv4 prefixes since the previous operation.

Configuration Impact

The **reset ip ip-prefix** command clears statistics about the IPv4 prefix list. After that, the previous statistics cannot be shown.

Precautions

The **reset ip ip-prefix** command:

- Clears statistics in a specified IPv4 prefix list, if the name of the IPv4 prefix list is specified using *ip-prefix-name*.
- Clears statistics in all IPv4 prefix lists, if the name of the IPv4 prefix list is not specified using *ip-prefix-name*.

Example

Reset the statistics of the specified IPv4 prefix list.

```
<HUAWEI> reset ip ip-prefix abc
```

7.9.56 reset ip ipv6-prefix

Function

The **reset ip ipv6-prefix** command resets the timer of a specified IPv6 prefix list.

Format

reset ip ipv6-prefix [*ipv6-prefix-name*]

Parameters

Parameter	Description	Value
ipv6-prefix-name	Specifies the name of an IP prefix list.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The IPv6 prefix list can be used to filter IPv6 addresses. When filtering IPv6 addresses, the system records the numbers of prefixes that are permitted and denied by the IPv6 prefix list. You can run the **display ip ipv6-prefix** command to view the numbers.

To view the number of IPv6 prefixes that are permitted and denied by the IPv6 prefix list, run the **reset ip ipv6-prefix** command to clear statistics about permitted and denied routes in the IPv6 prefix list, and then run the **display ip ipv6-prefix** command to display the number of IPv6 prefixes since the previous operation.

Configuration Impact

The **reset ip ipv6-prefix** command clears statistics about the IPv6 prefix list. After that, the previous statistics cannot be shown.

Precautions

The **reset ip ipv6-prefix** command:

- Clears statistics in a specified IPv6 prefix list, if the name of the IPv6 prefix list is specified using *ipv6-prefix-name*.
- Clears statistics in all IPv6 prefix lists, if the name of the IPv6 prefix list is not specified using *ipv6-prefix-name*.

Example

Resets the timer of the IPv6 prefix list named **abc**.

```
<HUAWEI> reset ip ipv6-prefix abc
```

7.9.57 reset route-policy counters

Function

The **reset route-policy counters** command resets route-policy counters.

Format

reset route-policy *route-policy-name* **counters**

Parameters

Parameter	Description	Value
route-policy-name	Specifies the name of a route-policy.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The route-policy is used to filter routes and set the attributes of a route that matches a route-policy. When a route-policy filters routes, the system records the

number of routes that match the route-policy nodes. You can run the **display route-policy** to view the numbers.

The **reset route-policy counters** command clears the number of routes which match or do not match the route-policy. You can run both the **reset route-policy counters** command and the **display route-policy** command to instruct whether to record the number of routes matching a specified route-policy.

Configuration Impact

The **reset route-policy counters** command clears the number of routes which match or do not match the route-policy. After that, the number cannot be restored.

Example

Reset the counters of a route-policy named **policy1**.

```
<HUAWEI> reset route-policy policy1 counters
```

7.9.58 route-policy

Function

The **route-policy** command creates a routing policy and displays the Route-Policy view.

The **undo route-policy** command deletes a specified routing policy.

By default, no routing policy is configured.

Format

route-policy *route-policy-name* { **permit** | **deny** } **node** *node*

undo route-policy *route-policy-name* [**node** *node*]

Parameters

Parameter	Description	Value
route-policy-name	Specifies the name of a routing policy. If the routing policy does no exist, create a routing policy and enter its Route-Policy view. If the routing policy exists, enter its Route-Policy view.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Parameter	Description	Value
permit	Specifies the matching mode of the routing policy as permit. In permit mode, a route matches all the **if-match** clauses, the route matches the routing policy and the actions defined by the **apply** clause are performed on the route. Otherwise, the route continues to match the next entry.	-
deny	Specifies the matching mode of the routing policy as deny. In deny mode, if a route matches all the **if-match** clauses, the route fails to match the routing policy and cannot match the next node.	-
node *node*	Specifies the index of the node in the routing policy. When the routing policy is used to filter routes, the node with the smaller value of *node* is matched first.	The value is an integer ranging from 0 to 65535.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A routing policy is used to filter routes and set route attributes for the routes that match the routing policy. A routing policy consists of multiple nodes. One node can be configured with multiple **if-match** and **apply** clauses.

The **if-match** clauses define matching rules for this node, and the **apply** clauses define behaviors for the routes that match the rules. The relationship between **if-match** clauses is "AND". That is, a route must match all the **if-match** clauses. The relationship between the nodes of a routing policy is "OR". That is, if a route matches one node, the route matches the routing policy. If the route does not match any node, the route fails to match the routing policy.

Procedure

After a routing policy is created, the system prompts "Info: New Sequence of this List !" and displays the Route-Policy view. The system displays no prompt when a routing policy is deleted.

Precautions

After a route-policy is configured, by default, the RM immediately notifies each protocol to apply the route-policy to filter routes. To delay applying a route-policy,

you need to run the **route-policy-change notify-delay** command to set the delay for applying the route-policy.

You can run the **display route-policy** command to view the number of routes that match and do not match the route-policy.

A *route-policy-name* must have been configured using the **route-policy** command before the *route-policy-name* is referenced by another command.

Before you run the **undo route-policy** command to delete a route-policy that is referenced by another command, delete the reference configuration.

If an **if-match** clause of a route-policy defines an **ip-prefix**-based filtering rule, the filtering rule applies to IPv4 prefixes, not to IPv6 prefixes, and IPv6 prefixes match the filtering rule by default. If IPv6 prefixes also need to be filtered, add an **ipv6-prefix**-based **if-match** clause. Similarly, if an **if-match** clause of a route-policy defines an **ipv6-prefix**-based filtering rule, the filtering rule applies to IPv6 prefixes, not to IPv4 prefixes, and IPv4 prefixes match the filtering rule by default. If IPv4 prefixes also need to be filtered, add an **ip-prefix**-based **if-match** clause.

The configuration of a peer takes precedence over that of the peer group to which the peer belongs. That is, when a routing policy is used to control BGP route advertisement or receiving, the priority of this configuration on a BGP peer is higher than that of the configuration on the peer group.

Example

Configure the routing policy named **policy1** whose node number is 10 and the matching mode is **permit**.

```
<HUAWEI> system-view
[HUAWEI] route-policy policy1 permit node 10
[HUAWEI-route-policy]
```

7.9.59 route-policy-change notify-delay

Function

The **route-policy-change notify-delay** command sets the delay before the RM to notify each protocol of applying a new policy after the original route-policy changes.

The **undo route-policy-change notify-delay** command restores the default setting.

By default, this command is not configured, and the delay time is 0s.

Format

route-policy-change notify-delay *delay-time*

undo route-policy-change notify-delay

Parameters

Parameter	Description	Value
delay-time	Specifies the delay for applying a new policy after the original route-policy changes.	The value is an integer ranging from 1 to 180, in seconds.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

The switch process the changes of a route-policy according to the following rules.

After the configuration of a route-policy changes, by default, the RM immediately notifies the protocol of performing related operations. To delay processing the changes of the route-policy, you can run the **route-policy-change notify-delay** command to set the delay for change processing. The new policy is applied after the timer expires:

- If the configuration of the route-policy changes again within the delay, the RM resets the timer.
- If the new policy is configured for BGP, the **refresh bgp all** command can be used within the delay set by the **route-policy-change notify-delay** command to trigger BGP to immediately use the new policy.

The following commands are related to the timer:

- **route-policy**
- **ip ip-prefix**
- **ip ipv6-prefix**
- **ip as-path-filter**
- **ip community-filter**
- **ip extcommunity-filter**
- **ip rd-filter**
- **acl**

Example

Set the delay before the RM to notify each protocol of applying a new policy after the original route-policy changes.

```
<HUAWEI> system-view
[HUAWEI] route-policy-change notify-delay 20
```


7.9.60 route-policy nonexistent-config-check

Function

The **route-policy nonexistent-config-check** command controls whether the system allows a nonexistent route-policy to be specified in a command.

The **undo route-policy nonexistent-config-check disable** command forbids a nonexistent route-policy to be specified in a command.

By default, the system does not allow a nonexistent route-policy to be specified in a command.

Format

route-policy nonexistent-config-check { **disable** | **enable** }

undo route-policy nonexistent-config-check disable

Parameters

Parameter	Description	Value
disable	Indicates that the system allows a nonexistent route-policy to be specified in a command.	-
enable	Indicates that the system does not allow a nonexistent route-policy to be specified in a command.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

By default, if you specify a nonexistent route-policy in a command, the command does not take effect. To enable the system to allow a nonexistent route-policy to be specified in a command, run the **route-policy nonexistent-config-check disable** command.

Example

Enable the system to allow a nonexistent route-policy to be specified in a command.

```
<HUAWEI> system-view
[HUAWEI] route-policy nonexistent-config-check disable
```


7.10 IP Routing Table Management Commands

7.10.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.10.2 arp vlink-direct-route advertise

Function

The **arp vlink-direct-route advertise** command configures the advertisement of IPv4 ARP Vlink direct routes.

The **undo arp vlink-direct-route advertise** command cancels advertising IPv4 ARP Vlink direct routes.

By default, a device does not advertise IPv4 ARP Vlink direct routes.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

arp vlink-direct-route advertise [**route-policy** *route-policy-name*]

undo arp vlink-direct-route advertise

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Specifies the name of the route-policy used to filter IPv4 ARP Vlink direct routes.	The name is a string of 1 to 40 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

System view, VPN instance view, VPN instance IPv4 address family view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, IPv4 ARP Vlink direct routes are only used to guide local traffic forwarding. To control the routing table size and maintain routing table stability, devices are prohibited from importing these direct routes into dynamic routing protocols for advertisement. In some situations, however, devices need to perform operations based on specific routes. For example, a device needs to direct remote traffic for different users using different export policies. Subsequently, the device needs to import IPv4 ARP Vlink direct routes into dynamic routing protocols and advertises them to the remote end.

If you only need to visit several users, you can specify the **route-policy** *route-policy-name* parameter to filter IPv4 ARP Vlink direct routes to be advertised. This operation can control the routing table size and implement precise control over data traffic.

Configuration Impact

After you specify the **route-policy** *route-policy-name* parameter to filter the IPv4 ARP Vlink direct routes to be advertised, only the routes that match the specified route-policy are advertised.

Follow-up Procedure

IPv4 ARP Vlink direct routes need to be imported into the routing tables of dynamic routing protocols on the device running dynamic routing protocols before being advertised by dynamic routing protocols.

Precautions

You need to run the **route-policy** command to configure a route-policy. If the specified route-policy does not exist, IPv4 ARP Vlink direct routes cannot be filtered.

Example

Configure the advertisement of IPv4 ARP Vlink direct routes and configure a route-policy **policy1** to allow route 10.1.1.4/32 to be advertised.

```
<HUAWEI> system-view
[HUAWEI] ip ip-prefix prefix1 permit 10.1.1.4 32
[HUAWEI] route-policy policy1 permit node 10
[HUAWEI-route-policy] if-match ip-prefix prefix1
[HUAWEI-route-policy] quit
[HUAWEI] arp vlink-direct-route advertise route-policy policy1
```

7.10.3 display fib

Function

The **display fib** command displays information about the FIB table.

Format

display fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] [**verbose**]

Parameters

Parameter	Description	Value
slot-id	Displays information about the FIB table with a specified slot ID.	The value is an integer and the value range depends on the device configuration.
vpn-instance *vpn-instance-name*	Displays information about the FIB table of a specified VPN instance.	The value must be an existing VPN instance name.
verbose	Displays detailed information about the FIB table.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display fib** command displays information about the FIB table. Each row represents a route.

NOTE

If there are lots of routes, using wildcard (|, **begin**, **exclude**, **include**, *regular-expression*) to display information or details lasts a long time. You can press Ctrl+C to terminate information display.

Example

Display brief information about the FIB table.

```
<HUAWEI> display fib
Route Flags: G - Gateway Route, H - Host Route,    U - Up Route
             S - Static Route,  D - Dynamic Route, B - Black Hole Route
             L - Vlink Route
------------------------------------------------------------------------------
 FIB Table:
 Total number of Routes : 5

Destination/Mask   Nexthop         Flag  TimeStamp     Interface       TunnelID
192.168.150.111/32 127.0.0.1       HU    t[117]        InLoop0         0x0
127.0.0.1/32       127.0.0.1       HU    t[54]         InLoop0         0x0
127.0.0.0/8        127.0.0.1       U     t[54]         InLoop0         0x0
192.168.150.0/24   192.168.150.111 U     t[117]        MEth0/0/1       0x0
0.0.0.0/0          192.168.150.1   GSU   t[117]        MEth0/0/1       0x0
```

Display information starting from the line that contains 169.254.0.0

```
<HUAWEI> display fib | begin 10.254.0.0
Destination/Mask  Nexthop   Flag  TimeStamp  Interface   TunnelID
10.254.0.0/16     10.2.1.1  U     t[0]       Vlanif10    0x0
10.12.0.0/16      10.2.1.1  U     t[0]       Vlanif10    0x0
127.0.0.0/8       127.0.0.1 U     t[0]       InLoop0     0x0
```

Display all lines that contain Vlanif10.

```
<HUAWEI> display fib | include Vlanif10
Destination/Mask  Nexthop   Flag  TimeStamp  Interface  TunnelID
10.254.0.0/16     10.2.1.1  U     t[0]       Vlanif10   0x0
10.12.0.0/16      10.2.1.1  U     t[0]       Vlanif10   0x0
```

Display all lines that do not include 192.168.150.0.

```
<HUAWEI> display fib | exclude 192.168.150.0

Destination/Mask   Nexthop         Flag  TimeStamp     Interface       TunnelID
192.168.150.111/32 127.0.0.1       HU    t[117]        InLoop0         0x0
127.0.0.1/32       127.0.0.1       HU    t[54]         InLoop0         0x0
127.0.0.0/8        127.0.0.1       U     t[54]         InLoop0         0x0
0.0.0.0/0          192.168.150.1   GSU   t[117]        MEth0/0/1       0x0
```

Table 7-205 Description of the display fib command output

Item	Description
Route Flags	Flag of a route.
Destination/Mask	Destination address or mask length.
Nexthop	Next hop.

Item	Description
Flag	Current flag, which is the combination of G, H, U, S, D, B and L. ● G (gateway route): indicates that the next hop is a gateway. ● H (host route): indicates that the route is a host route. ● U (available route): indicates that the route status is Up. ● S (static route): indicates that the route is manually configured. ● D (dynamic route): indicates that the route is generated based on the routing algorithm. ● B (blackhole route): indicates that the next hop is a null interface. ● L: indicates a Vlink route.
TimeStamp	Timestamp, which indicates the time equal to FIB entry generation time minus system startup time.
Interface	Outbound interface to the destination address.
TunnelID	Index of a forwarding entry. It is used in packet forwarding between the upstream and downstream boards. If the value is not 0x0, packets matching the entry are forwarded through the MPLS tunnel. If the value is 0x0, packets matching the entry are not forwarded through the MPLS tunnel.

Display detailed information about the FIB table.

```
<HUAWEI> display fib verbose
Route Flags: G - Gateway Route, H - Host Route,    U - Up Route
             S - Static Route,  D - Dynamic Route, B - Black Hole Route
             L - Vlink Route
------------------------------------------------------------------------------
 FIB Table:
 Total number of Routes : 3

 Destination: 127.0.0.1          Mask     : 255.255.255.255
 Nexthop    : 127.0.0.1          OutIf    : InLoopBack0
 LocalAddr  : 127.0.0.1          LocalMask: 0.0.0.0
 Flags      : HU                 Age      : 354280sec
 ATIndex    : 0                  Slot     : 0
 LspFwdFlag : 0                    LspToken : 0x0
 InLabel    : NULL                 OriginAs : 0
 BGPNextHop : 0.0.0.0                PeerAs   : 0
 QosInfo    : 0x0                OriginQos: 0x0
 NexthopBak : 0.0.0.0            OutIfBak : [No Intf]
```



```
LspTokenBak: 0x0                InLabelBak : NULL
LspToken_ForInLabelBak : 0x0
EntryRefCount : 0
VlanId : 0x0
BgpKey : 0
BgpKeyBak : 0
LspType        : 0            Label_ForLspTokenBak   : 0
MplsMtu        : 0            Gateway_ForLspTokenBak : 0.0.0.0
NextToken      : 0x0          IfIndex_ForLspTokenBak : 0
Label_NextToken : 0             Label : 0
LspBfdState    : 0

Destination: 127.255.255.255     Mask     : 255.255.255.255
Nexthop    : 127.0.0.1          OutIf    : InLoopBack0
LocalAddr  : 127.0.0.1          LocalMask: 0.0.0.0
Flags      : HU                 Age      : 354280sec
ATIndex    : 0                  Slot     : 0
LspFwdFlag : 0                    LspToken : 0x0
InLabel    : NULL               OriginAs : 0
BGPNextHop : 0.0.0.0               PeerAs   : 0
QosInfo    : 0x0                OriginQos: 0x0
NexthopBak : 0.0.0.0              OutIfBak : [No Intf]
LspTokenBak: 0x0                InLabelBak : NULL
LspToken_ForInLabelBak : 0x0
EntryRefCount : 0
VlanId : 0x0
BgpKey : 0
BgpKeyBak : 0
LspType        : 0            Label_ForLspTokenBak   : 0
MplsMtu        : 0            Gateway_ForLspTokenBak : 0.0.0.0
NextToken      : 0x0          IfIndex_ForLspTokenBak : 0
Label_NextToken : 0             Label : 0
LspBfdState    : 0

Destination: 255.255.255.255     Mask     : 255.255.255.255
Nexthop    : 127.0.0.1          OutIf    : InLoopBack0
LocalAddr  : 127.0.0.1          LocalMask: 0.0.0.0
Flags      : HU                 Age      : 354280sec
ATIndex    : 0                  Slot     : 0
LspFwdFlag : 0                    LspToken : 0x0
InLabel    : NULL               OriginAs : 0
BGPNextHop : 0.0.0.0               PeerAs   : 0
QosInfo    : 0x0                OriginQos: 0x0
NexthopBak : 0.0.0.0              OutIfBak : [No Intf]
LspTokenBak: 0x0                InLabelBak : NULL
LspToken_ForInLabelBak : 0x0
EntryRefCount : 0
VlanId : 0x0
BgpKey : 0
BgpKeyBak : 0
LspType        : 0            Label_ForLspTokenBak   : 0
MplsMtu        : 0            Gateway_ForLspTokenBak : 0.0.0.0
NextToken      : 0x0          IfIndex_ForLspTokenBak : 0
Label_NextToken : 0             Label : 0
LspBfdState    : 0
```

Table 7-206 Description of the display fib verbose command output

Item	Description
Destination	Destination address.
Mask	Mask.
Nexthop	Next hop.
OutIf	Outbound interface.

Item	Description
LocalAddr	Local IP address.
LocalMask	Mask of the local IP address.
Flags	Current flag, which is the combination of G, H, U, S, D, and B. • G (gateway route): indicates that the next hop is a gateway. • H (host route): indicates that the route is a host route. • U (available route): indicates that the route status is Up. • S (static route): indicates that the route is manually configured. • D (dynamic route): indicates that the route is generated based on the routing algorithm. • B (blackhole route): indicates that the next hop is a null interface. • L: indicates a Vlink route.
Age	Lifetime of a route, in seconds.
ATIndex	Index of the virtual link connecting the local end and the gateway.
Slot	Slot ID of the outbound interface.
LspFwdFlag	Forwarding flag of an LSP.
LspToken	Forwarding ID of an LSP.
InLabel	Inner tag of a VPN LSP.
OriginAs	Original AS number.
BGPNextHop	Address of the BGP next hop.
PeerAs	Neighbor AS number.
QosInfo	QoS information.
OriginQos	Original QoS information.
NexthopBak	Backup of the next hop.
OutIfBak	Backup of the outbound interface.
VlanId	VLAN ID.
BgpKey	Key value of the BGP route.
BgpKeyBak	Backup key value of the BGP route.

7.10.4 display fib acl

Function

The **display fib acl** command displays information about FIB entries that match a specified ACL rule.

Format

display fib [**vpn-instance** *vpn-instance-name*] **acl** *acl-number* [**verbose**]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays information about the FIB table of a specified VPN instance.	The value must be an existing VPN instance name.
acl-number	Displays information about the FIB table with a specified ACL rule.	The value is an integer that ranges from 2000 to 2999.
verbose	Displays detailed information about the FIB table.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display fib acl** command to check information about FIB entries that match a specified ACL rule.

Example

Display FIB entries that match an ACL rule.

```
<HUAWEI> display fib acl 2010
Route Entry matched by access-list 2010
 Summary Counts: 1
Destination/Mask Nexthop  Flag TimeStamp Interface     TunnelID
127.0.0.0/8      127.0.0.1 U    t[0]      InLoop0       0x0
```


Table 7-207 Description of the display fib acl command output

Item	Description
Route Entry matched by access-list	FIB entries that match an ACL rule.
Summary Counts	Total number of the FIB entries.
Destination/Mask	Destination address or mask length.
Nexthop	Next hop.
Flag	Current flag, which is the combination of G, H, U, S, D, B and T. • G (gateway route): indicates that the next hop is a gateway. • H (host route): indicates that the route is a host route. • U (available route): indicates that the route status is Up. • S: indicates a static route. • D: indicates a dynamic route. • B (blackhole route): indicates that the next hop is a null interface. • T: indicates ingress TOKEN_SETBYLSPM nodes.
TimeStamp	Timestamp, which indicates the lifetime of an entry, in seconds.
Interface	Outbound interface to the destination address.
TunnelID	Index of a forwarding entry. It is used in packet forwarding between the upstream and downstream boards. If the value is not 0, packets matching the entry are forwarded through the MPLS tunnel. If the value is 0, packets matching the entry are not forwarded through the MPLS tunnel.

7.10.5 display fib interface

Function

The **display fib interface** command displays information about FIB entries with a specified outbound interface.

Format

display fib [**vpn-instance** *vpn-instance-name*] **interface** *interface-type interface-number*

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays information about the FIB table of a specified VPN instance.	The value must be an existing VPN instance name.
interface-type interface-number	Specifies the type and number of the outbound interface to a specified destination address.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display fib interface** command to check information about FIB entries with a specified outbound interface.

Example

Display FIB entries with a specified outbound interface.

```
<HUAWEI> display fib interface vlanif10
Destination/Mask   Nexthop    Flag TimeStamp   Interface       TunnelID
10.1.1.0/24        10.1.1.2   U    t[115]      Vlanif10        0x0
```

Table 7-208 Description of the display fib interface command output

Item	Description
Destination/Mask	Destination address or mask length.
Nexthop	Next hop.

Item	Description
Flag	Current flag, which is the combination of G, H, U, S, D, and B. • G (gateway route): indicates that the next hop is a gateway. • H (host route): indicates that the route is a host route. • U (available route): indicates that the route status is Up. • S: indicates a static route. • D: indicates a dynamic route. • B (blackhole route): indicates that the next hop is a null interface.
TimeStamp	Timestamp, which indicates the lifetime of an entry, in seconds.
Interface	Outbound interface to the destination address.
TunnelID	Index of a forwarding entry. It is used in packet forwarding between the upstream and downstream boards. If the value is not 0, packets matching the entry are forwarded through the MPLS tunnel. If the value is 0, packets matching the entry are not forwarded through the MPLS tunnel.

7.10.6 display fib ip-prefix

Function

The **display fib ip-prefix** command displays information about the FIB table.

Format

display fib [**vpn-instance** *vpn-instance-name*] **ip-prefix** *prefix-name* [**verbose**]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays the FIB table of a specified VPN instance.	The value must be an existing VPN instance name.

Parameter	Description	Value
prefix-name	Specifies the name of an IP prefix list.	The value is a string of 1 to 169 characters.
verbose	Displays detailed information about the FIB table.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command displays FIB entries that match a specified IP prefix list.

If 0 FIB entries match, the command output displays that zero FIB entries match.

If the name of an IP prefix list is not found, all FIB entries are displayed.

Example

Display FIB entries matching the IP prefix list **abc0**.

```
<HUAWEI> display fib ip-prefix abc0
Route Entry matched by prefix-list abc0
 Summary Counts: 4
Destination/Mask  Nexthop   Flag  TimeStamp  Interface   TunnelID
127.0.0.0/8       127.0.0.1 U     t[0]       InLoop0       0x0
127.0.0.1/32      127.0.0.1 U     t[0]       InLoop0       0x0
172.16.0.0/8      10.1.1.1  SU    t[0]       Vlanif10      0x0
172.16.0.0/15     10.1.1.1  SU    t[0]       Vlanif10      0x0
```

Table 7-209 Description of the display fib ip-prefix command output

Item	Description
Destination/Mask	Destination address or mask length.
Nexthop	Next hop.

Item	Description
Flag	Current flag, which is the combination of G, H, U, S, D, and B. • G (gateway route): indicates that the next hop is a gateway. • H (host route): indicates that the route is a host route. • U (available route): indicates that the route status is Up. • S: indicates a static route. • D: indicates a dynamic route. • B (blackhole route): indicates that the next hop is a null interface.
TimeStamp	Timestamp, which indicates the lifetime of an entry, in seconds.
Interface	Outbound interface to the destination address.
TunnelID	Index of a forwarding entry. It is used in packet forwarding between the upstream and downstream boards. If the value is not 0, packets matching the entry are forwarded through the tunnel. If the value is 0, packets matching the entry are not forwarded through the tunnel.

7.10.7 display fib longer

Function

The **display fib longer** command displays FIB entries that match a specified parameter.

The **display fib** [*slot-id*] *destination-address* command displays FIB entries that match a specified destination IP address. If the specified destination IP address matches an FIB entry in the natural mask range, all the subnets are displayed. Otherwise, FIB entries are displayed based on the longest matching principle.

The **display fib** [*slot-id*] *destination-address destination-mask* command displays FIB entries that accurately match the destination address and mask.

The **display fib** [*slot-id*] *destination-address* **longer** command displays all FIB entries that match destination IP addresses in the natural mask range.

The **display fib** [*slot-id*] *destination-address* destination-mask **longer** command displays all FIB entries that match destination IP addresses in a specified mask range.

The **display fib** [*slot-id*] *destination-address1 destination-mask1 destination-address2 destination-mask2* command displays FIB entries that match destination

IP addresses in the range of *destination-address1 destination-mask1* to *destination-address2 destination-mask2*.

If *slot-id* is specified, FIB entries on the device are displayed.

Format

display fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] *destination-address1* [*destination-mask1*] [**longer**] [**verbose**]

display fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] *destination-address1 destination-mask1 destination-address2 destination-mask2* [**verbose**]

Parameters

Parameter	Description	Value
slot-id	Displays information about the FIB table with a specified slot ID.	The value is an integer and the value range depends on the device configuration.
vpn-instance *vpn-instance-name*	Displays information about the FIB table of a specified VPN instance.	The value must be an existing VPN instance name.
destination-address1	Indicates destination IP address 1.	The value is in dotted decimal notation.
destination-mask1	Indicates subnet mask 1.	The value is in dotted decimal notation or an integer that ranges from 0 to 32.
destination-address2	Indicates destination IP address 2.	The value is in dotted decimal notation.
destination-mask2	Indicates subnet mask 2.	The value is in dotted decimal notation or an integer that ranges from 0 to 32.
longer	Displays FIB entries matching a specified network segment or mask.	-
verbose	Displays detailed information about the FIB table.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display fib longer** command to check FIB entries that match a specified condition.

Example

Display FIB entries that match destination address 172.16.0.0 in the natural mask range or based on the longest match principle.

```
<HUAWEI> display fib 172.16.0.0
  Route Entry Count: 1
Destination/Mask  Nexthop  Flag  TimeStamp  Interface  TunnelID
172.16.0.0/16     10.1.1.1 U     t[0]       Vlanif10     0x0
```

Display the FIB entries with the destination addresses in the range of 172.16.0.0/16 to 172.16.0.6/16.

```
<HUAWEI> display fib 172.16.0.0 255.255.0.0 172.16.0.6 255.255.0.0
  Route Entry Count: 1
Destination/Mask  Nexthop  Flag  TimeStamp  Interface  TunnelID
172.16.0.1/8      10.1.1.1  U     t[0]       Vlanif10     0x0
```

Table 7-210 Description of the display fib longer command output

Item	Description
Destination/Mask	Destination address or mask length.
Nexthop	Next hop.
Flag	Current flag, which is the combination of G, H, U, S, D, and B. • G (gateway route): indicates that the next hop is a gateway. • H (host route): indicates that the route is a host route. • U (available route): indicates that the route status is Up. • S: indicates a static route. • D: indicates a dynamic route. • B (blackhole route): indicates that the next hop is a null interface.
TimeStamp	Timestamp, which indicates the lifetime of an entry, in seconds.
Interface	Outbound interface to the destination address.

Item	Description
TunnelID	Index of a forwarding entry. It is used in packet forwarding between the upstream and downstream boards. If the value is not 0, packets matching the entry are forwarded through the tunnel. If the value is 0, packets matching the entry are not forwarded through the tunnel.

7.10.8 display fib next-hop

Function

The **display fib next-hop** command displays FIB entries that match a specified next-hop IP address.

Format

display fib [**vpn-instance** *vpn-instance-name*] **next-hop** *ip-address*

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays information about the FIB table of a specified VPN instance.	The value must be an existing VPN instance name.
ip-address	Specifies the next-hop IP address.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display fib next-hop** command to check FIB entries that match a specified next-hop IP address.

Example

Display FIB entries that match a specified next-hop IP address.

```
<HUAWEI> display fib next-hop 10.1.1.1
Destination/Mask   Nexthop        Flag TimeStamp     Interface      TunnelID
```



```
10.1.1.1/32      10.1.1.1        HU   t[115]      Vlanif10        0x0
```

Table 7-211 Description of the display fib next-hop command output

Item	Description
Destination/Mask	Destination address or mask length.
Nexthop	Next hop.
Flag	Current flag, which is the combination of G, H, U, S, D, and B. ● G (gateway route): indicates that the next hop is a gateway. ● H (host route): indicates that the route is a host route. ● U (available route): indicates that the route status is Up. ● S: indicates a static route. ● D: indicates a dynamic route. ● B (blackhole route): indicates that the next hop is a null interface.
TimeStamp	Timestamp, which indicates the lifetime of an entry, in seconds.
Interface	Outbound interface to the destination address.
TunnelID	Index of a forwarding entry. It is used in packet forwarding between the upstream and downstream boards. If the value is not 0, packets matching the entry are forwarded through the MPLS tunnel. If the value is 0, packets matching the entry are not forwarded through the MPLS tunnel.

7.10.9 display fib statistics

Function

The **display fib statistics** command displays the total number of IPv4 FIB entries.

Format

display fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] **statistics**

Parameters

Parameter	Description	Value
slot-id	Displays status information about a specified IPv4 FIB module.	The value is an integer and the value range depends on the device configuration.
vpn-instance *vpn-instance-name*	Specifies the name of an IPv4 VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display fib statistics** command to check the total number of IPv4 FIB entries.

Example

Display the total number of IPv4 FIB entries.

```
<HUAWEI> display fib statistics
Route Entry Count : 30
Route Prefix Count : 30
```

Table 7-212 Description of the display fib statistics command output

Item	Description
Route Entry Count	Total number of FIB entries.
Route Prefix Count	Total number of route prefix entries.

7.10.10 display fib statistics all

Function

The **display fib statistics all** command displays IPv4 FIB entry statistics.

Format

display fib [*slot-id*] **statistics all**

Parameters

Parameter	Description	Value
slot-id	Specifies the slot ID of the switch on which IPv4 FIB entry statistics are displayed.	The value is an integer and the value range depends on the device configuration.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

When running the **display fib statistics all** command, you can view the IPv4 FIB route prefix capacity and the number of IPv4 FIB forwarding entries on the device.

Example

Display statistics about IPv4 FIB entries on the switch in slot 0.

```
<HUAWEI> display fib 0 statistics all
IPv4 FIB Route Prefix Capacity : 320000
IPv4 FIB Total Route Prefix Count :2 ; Entry Count : 2
IPv4 FIB Public Route Prefix Count :2 ; Entry Count : 2
IPv4 FIB VPN-instance vpna Route Prefix Count :0 ; Entry Count : 0
```

Table 7-213 Description of the display fib statistics all command output

Item	Description
IPv4 FIB Route Prefix Capacity:	IPv4 FIB route prefix capacity on the local switch.
IPv4 FIB Total Route Prefix Count: Entry Count:	Total number of IPv4 FIB route prefixes and forwarding entries on the local switch.
IPv4 FIB Public Route Prefix Count: Entry Count:	Total number of IPv4 route prefixes and forwarding entries on a public network.
IPv4 FIB VPN-instance vpna Route Prefix Count: Entry Count:	Total number of route prefixes and forwarding entries of IPv4 VPN instance **vpna**.

7.10.11 display ip routing-table

Function

The **display ip routing-table** command displays information about an IPv4 routing table.

Format

display ip routing-table [**vpn-instance** *vpn-instance-name*] [**verbose**]

display ip routing-table [**vpn-instance** *vpn-instance-name*] *ip-address* [*mask* | *mask-length*] [**longer-match**] [**verbose**]

display ip routing-table [**vpn-instance** *vpn-instance-name*] *ip-address* { *mask* | *mask-length* } **nexthop** *nexthop-address* [**verbose**]

display ip routing-table [**vpn-instance** *vpn-instance-name*] *ip-address1* { *mask1* | *mask-length1* } *ip-address2* { *mask2* | *mask-length2* } [**verbose**]

display ip routing-table [**vpn-instance** *vpn-instance-name*] **acl** { *acl-number* | *acl-name* } [**verbose**]

display ip routing-table [**vpn-instance** *vpn-instance-name*] **ip-prefix** *ip-prefix-name* [**verbose**]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance of an enabled IPv4 address family.	The value must be an existing VPN instance name.
verbose	Displays detailed information about active routes and inactive routes. If this keyword is not specified, only brief information about active routes is displayed.	-
ip-address	Displays the routes with the specified destination address.	The value is in dotted decimal notation.
nexthop *nexthop-address*	Displays the routes with the specified next-hop address.	The value is in dotted decimal notation.
longer-match	Displays the routes with the specified destination address and mask.	-

Parameter	Description	Value
mask	Specifies the mask of the destination IP address.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of the destination IP address.	The value is an integer that ranges from 0 to 32.
ip-address1	Specifies the start IP address in an IP address range. *ip-address1* and *ip-address2* determine an IP address range. Only the routes in the IP address range are displayed.	The value is in dotted decimal notation.
ip-address2	Specifies the end IP address in an IP address range. *ip-address1* and *ip-address2* determine an IP address range. Only the routes in the IP address range are displayed.	The value is in dotted decimal notation.
mask1	Specifies the subnet mask of the start IP address.	The value is in dotted decimal notation.
mask-length1	Specifies the mask length of the start IP address.	The value is an integer that ranges from 0 to 32.
mask2	Specifies the subnet mask of the end IP address.	The value is in dotted decimal notation.
mask-length2	Specifies the mask length of the end IP address.	The value is an integer that ranges from 0 to 32.
acl *acl-number*	Displays the routes that match the ACL with the specified ACL number. If the specified ACL does not exist, information about all active routes is displayed.	The value is an integer that ranges from 2000 to 2999.
acl *acl-name*	Displays the routes that match the ACL with the specified ACL name. If the specified ACL does not exist, information about all active routes is displayed.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.

Parameter	Description	Value
ip-prefix *ip-prefix-name*	Displays the routes that match the specified IP prefix list. If the specified IP prefix list does not exist, information about all active routes is displayed.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the **verbose** keyword is not specified in the **display ip routing-table** command, each line in the command output indicates a route, including the destination address, mask length, protocol, preference, cost, flag, next hop, and outbound interface of the route.

NOTE

A recursive route is counted as one route regardless of how many outbound interfaces and next hops the route finds.

Matching rules vary with parameters in the command:

- If the **display ip routing-table** *ip-address* command is used, the routes that longest match the destination address are displayed.
- If the **display ip routing-table** *ip-address mask* command is used, the routes that accurately match the destination address and mask are displayed.
- If the **display ip routing-table** *ip-address* **longer-match** command is used, all the routes that match the IP address specified by the destination address and the natural mask are displayed.
- If the **display ip routing-table** *ip-address mask* **longer-match** command is used, all the routes that match the IP address specified by the destination address and the input mask are displayed.
- If the **display ip routing-table** *ip-address1 mask1 ip-address2 mask2* command is used, the routes whose destination address ranges from *ip-address1 mask1* to *ip-address2 mask2* are displayed.

Example

Display brief information about the current IPv4 routing table. The command output shows that there are two static routes with the same destination address 10.1.1.1/32 but different next hops.


```
<HUAWEI> display ip routing-table
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 5        Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.1.1/32   Static  60   0           D   0.0.0.0         NULL0
                    Static  60   0           D   192.168.0.2     Vlanif100
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
  192.168.150.0/24  Direct  0    0           D   192.168.150.22  Vlanif4094
 192.168.150.22/32  Direct  0    0           D   127.0.0.1       Vlanif4094
```

Table 7-214 Description of the display ip routing-table command output

Item	Description
Route Flags	Flag of a route: ● R: indicates that the route is a recursive route. ● D: indicates that the route is delivered to the FIB table. ● T: indicates a route whose next hop belongs to a VPN instance.
Routing Tables: Public	The routing table is a public routing table. If the routing table is a private routing table, a private network name is displayed, for example, Routing Tables: ABC.
Destinations	Total number of destination networks or hosts.
Routes	Total number of routes.
Destination/Mask	Address and mask length of the destination network or host.
Proto	Routing protocol that learns a route.
Pre	Preference of a route.
Cost	Cost of a route.
Flags	Route flags in the heading of the routing table.
NextHop	Next-hop address of a route.
Interface	Outbound interface through which the next hop of a route is reachable.

Display the summary of the routing table of the IPv4 VPN instance named **vpn1**.

```
<HUAWEI> display ip routing-table vpn-instance vpn1
Route Flags: R - relay, D - download to fib, T - to vpn-instance
```



```
------------------------------------------------------------------------------
Routing Table: vpn1
         Destinations : 3        Routes : 3

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

        10.1.1.0/24  Direct 0    0           D  10.1.1.1        Vlanif10
        10.1.1.1/32  Direct 0    0           D  127.0.0.1       Vlanif10
        10.5.5.0/24  Static 60   0          RD  10.1.1.2        Vlanif10
```

Table 7-215 Description of the display ip routing-table vpn-instance command output

Item	Description
Route Flags	Route flag: • R: indicates a recursive route. • D: indicates that the route is downloaded to the FIB table. • T: indicates a route whose next hop belongs to a VPN instance.
Routing Tables: vpn1	VPN routing table named vpn1.
Destinations	Total number of destination networks or hosts.
Routes	Total number of routes.
Destination/Mask	Address and mask length of the destination network or host.
Proto	Routing protocol.
Pre	Preference.
Cost	Route cost.
Flags	Route flag, that is, Route Flags in the heading of the routing table.
NextHop	Next hop.
Interface	Outbound interface through which the next hop is reachable.

Display brief information about the current routing table. Route 10.2.2.2/32 is a static route with next hop 10.1.1.1. This route is a recursive route and has two outbound interfaces because 10.1.1.1 has two outbound interfaces.

```
<HUAWEI> display ip routing-table
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 6        Routes : 7

  Destination/Mask  Proto   Pre  Cost Flags  NextHop    Interface

      10.1.1.1/32   Static  60   0      D    0.0.0.0    NULL0
                    Static  60   0      D    10.10.0.2  Vlanif100
```



```
      10.2.2.2/32   Static  60   0     RD   10.1.1.1    NULL0
                    Static  60   0     RD   10.1.1.1    Vlanif100
     10.10.0.0/24   Direct   0   0      D    10.10.0.1  Vlanif100
     10.10.0.1/32   Direct   0   0      D    127.0.0.1  Vlanif100
     127.0.0.0/8    Direct   0   0      D    127.0.0.1  InLoopBack0
     127.0.0.1/32   Direct   0   0      D    127.0.0.1  InLoopBack0
```

Display detailed information about the IPv4 routing table.

```
<HUAWEI> display ip routing-table verbose
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------ Routing Tables: Public
         Destinations : 5        Routes : 5


Destination: 10.0.0.36/32
    Protocol: Direct            Process ID: 0
  Preference: 0                       Cost: 0
     NextHop: 127.0.0.1         Neighbour: 0.0.0.0
       State: Active Adv              Age: 01h03m56s
         Tag: 0                 Priority: high
       Label: NULL                 QoSInfo: 0x0
  IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: InLoopBack0
    TunnelID: 0x0                    Flags: D

Destination: 10.10.36.0/24
    Protocol: Direct            Process ID: 0
  Preference: 0                       Cost: 0
     NextHop: 10.10.36.2        Neighbour: 0.0.0.0
       State: Active Adv              Age: 00h26m36s
         Tag: 0                 Priority: high
       Label: NULL                 QoSInfo: 0x0
  IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: Vlanif10
    TunnelID: 0x0                    Flags: D

Destination: 10.10.36.2/32
    Protocol: Direct            Process ID: 0
  Preference: 0                       Cost: 0
     NextHop: 127.0.0.1         Neighbour: 0.0.0.0
       State: Active Adv              Age: 00h26m46s
         Tag: 0                 Priority: high
       Label: NULL                 QoSInfo: 0x0
  IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: Vlanif10
    TunnelID: 0x0                    Flags: D

Destination: 127.0.0.0/8
    Protocol: Direct            Process ID: 0
  Preference: 0                       Cost: 0
     NextHop: 127.0.0.1         Neighbour: 0.0.0.0
       State: Active NoAdv            Age: 3d01h20m39s
         Tag: 0                 Priority: high
       Label: NULL                 QoSInfo: 0x0
  IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: InLoopBack0
    TunnelID: 0x0                    Flags: D

Destination: 127.0.0.1/32
    Protocol: Direct            Process ID: 0
  Preference: 0                       Cost: 0
     NextHop: 127.0.0.1         Neighbour: 0.0.0.0
       State: Active NoAdv            Age: 3d01h20m39s
         Tag: 0                 Priority: high
       Label: NULL                 QoSInfo: 0x0
  IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: InLoopBack0
    TunnelID: 0x0                    Flags: D
```


Table 7-216 Description of the display ip routing-table verbose command output

Item	Description
Route Flags	Flag of a route: • R: indicates that the route is a recursive route. • D: indicates that the route is delivered to the FIB table. • T: indicates a route whose next hop belongs to a VPN instance.
Destinations	Total number of destination networks or hosts.
Routes	Total number of active routes and inactive routes.
Destination	Address and mask length of the destination network or host.
Protocol	Routing protocol of a route.
Process ID	Routing protocol process ID of a route.
Preference	Preference of a route.
Cost	Cost of a route.
NextHop	Next-hop address of a route.
Neighbour	Neighbor address of a route.
State	Status of a route: • Active: active route • Invalid: invalid route • Inactive: inactive route • NoAdv: route that cannot be advertised • Adv: route that can be advertised • Del: route to be deleted • Relied: route that recurses to an outbound interface and a next hop or that recurses to a tunnel • Stale: route that is marked Stale and used in GR
Age	Lifetime of a route.
Tag	Routing management tag. The value is an integer that ranges from 0 to 4294967295.

Item	Description
Priority	Convergence priority of route: • low • medium • high • critical
Label	Label allocated by MPLS.
QoSInfo	QoS information. The value 0x0 indicates that QoS information is empty.
IndirectID	ID of indirect next hop.
RelayNextHop	Recursive next-hop address.
Interface	Outbound interface.
TunnelID	Tunnel ID.
Flags	Route flags in the heading of the routing table.

7.10.12 display ip routing-table limit

Function

The **display ip routing-table limit** command displays the maximum number of routes and prefixes.

Format

display ip routing-table limit [**all-vpn-instance** | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
all-vpn-instance	Indicates all IPv4 VPN instances.	-
vpn-instance *vpn-instance-name*	Specifies the name of a IPv4 VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command displays limits on the number of IPv4 routes and prefixes.

- The **display ip routing-table limit** command displays limits on the number of IPv4 public routes and prefixes.
- The **display ip routing-table limit all-vpn-instance** command displays the limit on the number of routes and prefixes of all IPv4 VPN instances.
- The **display ip routing-table limit vpn-instance** *vpn-instance-name* command displays the limit on the number of routes and prefixes of a specified IPv4 VPN instance.

Example

Display the limits on the number of IPv4 public routes and prefixes.

```
<HUAWEI> display ip routing-table limit
Public Instance:
Limit-Object Limit-Type     Upper-Limit  Warning    Current    Log-Interval
Route        Default        -            -          9          5
Prefix       Default        -            -          9          5
```

Display limits on the numbers of routes and prefixes of all IPv4 VPN instances.

```
<HUAWEI> display ip routing-table limit all-vpn-instance
Limit-Object     Limit-Type    Upper-Limit     Warning     Current    Log-Interval
------------------------------------------------------------------------------
VPN Instance Name: vpn1
Route            Simply-Alert  5000           -           4223       5
Prefix           Alert-Percent 1000           800         760        5
------------------------------------------------------------------------------
VPN Instance Name: vpn2
Route            Alert-Percent 2000           1000        823        5
Prefix           Default       -              -           760        5
```

Display limits on the numbers of routes and prefixes of the IPv4 VPN instance named **vpn1**.

```
<HUAWEI> display ip routing-table limit vpn-instance vpn1
VPN Instance Name: vpn1
Limit-Object     Limit-Type     Upper-Limit    Warning    Current    Log-Interval
Route            Simply-Alert   5000           -          4223       5
Prefix           Alert-Percent  1000           800        760        5
```

Table 7-217 Description of the display ip routing-table limit command output

Item	Description
Limit-Object	Object whose total number is limited: • Prefix • Route

Item	Description
Limit-Type	Limit type for the routes and prefixes in the routing table: • Simply-Alert: indicates that only alarms are generated when the number of routes or prefixes exceeds the upper limit. • Alert-Percent: indicates the percentage of the alarm threshold of routes. • Default: indicates that the number of routes or prefixes is not limited by default.
Upper-Limit	Upper limit of routes or prefixes in the current routing table.
Warning	Alarm threshold of routes or prefixes in the current routing table.
Current	Number of routes or prefixes in the current routing table.
Log-Interval	Frequency of displaying logs when the number of routes or prefixes in the current routing table exceeds the upper limit, in seconds. The default value is 5s.

7.10.13 display ip routing-table protocol

Function

The **display ip routing-table protocol** command displays routing information about a specified routing protocol.

Format

display ip routing-table [**vpn-instance** *vpn-instance-name*] **protocol** *protocol* [**inactive** | **verbose**]

Parameters

Parameter	Description	Value
vpn-instance-name	Specifies the name of a VPN instance of an enabled IPv4 address family.	The value must be an existing VPN instance name.

Parameter	Description	Value
protocol	Displays routing information of a specified routing protocol.	The value may be **bgp, direct**, **isis, ospf, rip**, **static**, or **unr**. The specific value varies depending on the routing protocol supported by the device.
inactive	Displays brief information about inactive routes only.	-
verbose	Displays detailed information about active routes and inactive routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If neither **verbose** nor **inactive** is specified, brief information about all the routes of each routing protocol is displayed.

Example

Display brief information about all direct routes.

```
<HUAWEI> display ip routing-table protocol direct
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
Public routing table : Direct
         Destinations : 4        Routes : 4

Direct routing table status : <Active>
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

   10.137.216.0/23  Direct  0    0           D   10.137.217.210  Vlanif100
 10.137.217.210/32  Direct  0    0           D   127.0.0.1       Vlanif100
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0

Direct routing table status : <Inactive>
         Destinations : 0        Routes : 0
```


Table 7-218 Description of the display ip routing-table protocol command output

Item	Description
Route Flags	Flag of a route: ● R: indicates that the route is a recursive route. ● D: indicates that the route is delivered to the FIB table. ● T: indicates a route whose next hop belongs to a VPN instance.
Public routing table	Contents of a public routing table.
Direct routing table status	Status of direct routes: ● **Inactive**: inactive routes in the routing table ● **Active**: active routes in the routing table
Destinations	Total number of destination addresses.
Routes	Total number of routes in the routing table.
Destination/Mask	Destination address or mask length.
Proto	Routing protocol of a route.
Pre	Routing protocol preference of a route.
Cost	Cost of a route.
Flags	Route flags in the heading of the routing table.
Nexthop	Next-hop address of a route.
Interface	Outbound interface of a route.

Display all the direct routes of the VPN instance named **vpn1**.

```
<HUAWEI> display ip routing-table vpn-instance vpn1 protocol direct
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
vpn4 routing table : Direct
         Destinations : 2          Routes : 2

Direct routing table status : <Active>
         Destinations : 2          Routes : 2

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

       10.1.1.0/24  Direct 0    0           D  10.1.1.1        Vlanif30
       10.1.1.1/32  Direct 0    0           D  127.0.0.1       Vlanif30

Direct routing table status : <Inactive>
         Destinations : 0          Routes : 0
```


Table 7-219 Description of the display ip routing-table vpn-instance protocol command output

Item	Description
Active	Active routes.
Inactive	Inactive routes.

7.10.14 display ip routing-table statistics

Function

The **display ip routing-table statistics** command displays statistics about routes in an IPv4 routing table.

Format

display ip routing-table [**vpn-instance** *vpn-instance-name* | **all-vpn-instance**] **statistics**

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays statistics about routes in the routing table of a specified VPN instance. If neither **vpn-instance** *vpn-instance-name* nor **all-vpn-instance** is specified, statistics about routes in a public routing table are displayed.	The value must be an existing VPN instance name.
all-vpn-instance	Displays statistics about routes in the routing tables of all VPN instances. If neither **vpn-instance** *vpn-instance-name* nor **all-vpn-instance** is specified, statistics about routes in a public routing table are displayed.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Route statistics include:

- Total number of routes that are added or deleted through routing protocols
- Number of active or inactive routes that are labeled for deletion but are not deleted

Example

Display statistics about routes in an IPv4 routing table.

```
<HUAWEI> display ip routing-table statistics
Summary Prefixes: 9
Proto   total   active  added     deleted    freed
        routes  routes  routes    routes     routes
DIRECT  6       6       42        36         36
STATIC  3       3       21        18         18
RIP     0       0       0         0          0
OSPF    0       0       0         0          0
IS-IS   0       0       0         0          0
BGP     0       0       0         0          0
UNR     0       0       0         0          0
Total   9       9       63        54         54
```

Table 7-220 Description of the display ip routing-table statistics command output

Item	Description
Summary Prefixes	Total number of prefixes in the current routing table.
Proto	Routing protocol type
total routes	Total number of routes that a routing protocol learns in the routing table, including active and inactive routes.
active routes	Number of active routes that a routing protocol learns in the routing table.
added routes	Number of active and inactive routes that are added to the routing table through a routing protocol.
deleted routes	Number of routes to be deleted from the routing table.
freed routes	Number of routes that are permanently deleted from the routing table.

7.10.15 display ip routing-table time-range

Function

The **display ip routing-table time-range** command displays information about routes generated in a specified time range in an IP routing table.

Format

display ip routing-table [**vpn-instance** *vpn-instance-name*] **time-range** *min-age max-age* [**verbose**]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays information about routes generated in a specified time range in the IP routing table of the specified VPN instance. If you do not specify this parameter, the **display ip routing-table time-range** command displays information about routes generated in a specified time range in the public IP routing table.	The value must be an existing VPN instance name.
min-age	Specifies the end time of the period when routes are generated.	The value is in xxdxxhxxmxxs format. • d indicates days. The value is an integer that ranges from 0 to 10000. • h indicates hours. The value is an integer that ranges from 0 to 23. • m indicates minutes. The value is an integer that ranges from 0 to 59. • s indicates seconds. The value is an integer that ranges from 0 to 59. For example, you can enter 5d4h30m20s to specify 5 days, 4 hours, 30 minutes, and 20 seconds. **NOTE** If the value of d is 10000, the values of h, m, and s can only be 0.

Parameter	Description	Value
max-age	Specifies the start time of the period when routes are generated.	The value is in xxdxxhxxmxxs format. • d indicates days. The value is an integer that ranges from 0 to 10000. • h indicates hours. The value is an integer that ranges from 0 to 23. • m indicates minutes. The value is an integer that ranges from 0 to 59. • s indicates seconds. The value is an integer that ranges from 0 to 59. For example, you can enter 5d4h30m20s to specify 5 days, 4 hours, 30 minutes, and 20 seconds. **NOTE** If the value of d is 10000, the values of h, m, and s can only be 0.
verbose	Displays detailed information about active and inactive routes. If you do not specify this parameter, the **display ip routing-table time-range** command displays only brief information about active routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

If route flapping occurs on a network, you can run the **display ip routing-table time-range** command and specify a small time range for the command. You can then find the flapping route rapidly and accelerate fault location.

Precautions

max-age must be larger than *min-age*. Otherwise, the **display ip routing-table time-range** command does not display any information.

If *max-age* is larger than *min-age* but no route was generated within this time range, the **display ip routing-table time-range** command displays only the table heading.

Example

Display information about public network routes generated in the last 20 minutes.

```
<HUAWEI> display ip routing-table time-range 0 20m
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 2        Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.3.3.3/32  Direct  0    0           D   127.0.0.1       LoopBack1
        10.4.4.0/24  Static  60   0           D   0.0.0.0         NULL0
```

Table 7-221 Description of the display ip routing-table time-range command output

Item	Description
Route Flags	Route flags: • R: The route is a recursive route. • D: The route is sent to the FIB table. • T: Indicates a route whose next hop belongs to a VPN instance.
Routing Tables: Public	The routing table is a public routing table.
Destinations	Number of destination networks and hosts.
Routes	Number of routes.
Destination/Mask	Address and mask length of the destination network and host.
Proto	Protocol used to learn routes.
Pre	Route preference.
Cost	Route cost.
Flags	Route flags in the heading of the routing table.
NextHop	Next hop.
Interface	Outbound interface in a route with a reachable next hop.

Display detailed information about public network routes generated in the last 20 minutes.

```
<HUAWEI> display ip routing-table time-range 0 20m verbose
Route Flags: R - relay, D - download to fib, T - to vpn-instance
------------------------------------------------------------------------------
Routing Tables: Public
```



```
         Destinations : 3        Routes : 3


Destination: 10.3.3.3/32
     Protocol: Direct           Process ID: 0
   Preference: 0                      Cost: 0
      NextHop: 127.0.0.1         Neighbour: 0.0.0.0
        State: Active Adv             Age: 00h14m06s
          Tag: 0                 Priority: high
        Label: NULL                QoSInfo: 0x0
   IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: LoopBack1
    TunnelID: 0x0                 Flags:  D

Destination: 10.4.4.0/24
     Protocol: Static           Process ID: 0
   Preference: 60                     Cost: 0
      NextHop: 0.0.0.0           Neighbour: 0.0.0.0
        State: Active Adv             Age: 00h01m38s
          Tag: 0                 Priority: medium
        Label: NULL                QoSInfo: 0x0
   IndirectID: 0x0
 RelayNextHop: 0.0.0.0            Interface: NULL0
    TunnelID: 0x0                 Flags:  D

Destination: 10.4.4.4/32
     Protocol: Static           Process ID: 0
   Preference: 60                     Cost: 0
      NextHop: 10.4.4.0          Neighbour: 0.0.0.0
        State: Invalid Adv Relied      Age: 00h01m38s
          Tag: 0                 Priority: medium
        Label: NULL                QoSInfo: 0x0
   IndirectID: 0x80000002
 RelayNextHop: 0.0.0.0            Interface: NULL0
    TunnelID: 0x0                 Flags: R
```

Table 7-222 Description of the display ip routing-table time-range verbose command output

Item	Description
Process ID	Process ID of the routing protocol.
Preference	Route preference.
Cost	Route cost.
NextHop	Next hop.
Neighbour	IP address of a neighbor. 0.0.0.0 indicates that the route is generated by a local device.

Item	Description
State	Route status: • Active: an active route • Invalid: an invalid route • Inactive: an inactive route • NoAdv: a route that cannot be advertised • Adv: a route that can be advertised • Del: a route to be deleted • Relied: a route that recurses to the next hop and outbound interface or that recurses to a tunnel • Stale: a route with the Stale flag and used in GR
Age	Time after the route is generated.
Tag	Administrative tag of the route. The value is an integer that ranges from 0 to 4294967295.
Priority	Convergence priority of the route: • low: The convergence priority of a route is low. • medium: The convergence priority of a route is medium. • high: The convergence priority of a route is high. • critical: The convergence priority of a route is critical.
Label	Label allocated by MPLS.
QoSInfo	QoS information. 0x0 indicates QoS information is empty.
IndirectID	ID of the indirect next hop.
RelayNextHop	Recursive next hop.
Interface	Recursive outbound interface.
Tunnel ID	Tunnel ID: • The value 0x0 indicates that the route does not use a tunnel or the tunnel fails to be set up. • If the value is not 0x0, the route recurses to a tunnel.

Item	Description
Flags	Route flags in the heading of the routing table.

7.10.16 display ipv6 fib

Function

The **display ipv6 fib** command displays FIB entries on the device.

Format

display ipv6 fib [*slot-id*] **statistics** [**all**]

display ipv6 fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] **statistics**

display ipv6 fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] [*ipv6-address* [*prefix-length*]] [**verbose**]

display ipv6 fib [*slot-id*] [**vpn-instance** *vpn-instance-name*] **verbose statistics**

Parameters

Parameter	Description	Value
slot-id	Specifies the slot ID.	-
ipv6-address	Specifies the prefix of an IPv6 address.	The value is a 32-digit hexadecimal number, in the format X:X:X:X:X:X:X:X.
statistics	Displays the statistics of FIB entries.	-
all	Displays all FIB entries.	-
vpn-instance *vpn-instance-name*	Displays FIB entries of the specified VPN instance.	The value must be an existing VPN instance name.
prefix-length	Specifies the prefix length of an IPv6 address.	The value is an integer that ranges from 1 to 128.
verbose	Displays detailed information about FIB entries.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ipv6 fib** command to check IPv6 FIB entries on the device.

Example

Display all FIB entries on the device.

```
<HUAWEI> display ipv6 fib
 IPv6 FIB Table:
 Total number of Routes : 1

Destination:  ::1                          PrefixLength : 128
Nexthop    :  ::1                          Flag         : HU
Interface  :  InLoopBack0                      Tunnel ID    : 0x0
TimeStamp  :  2008-11-03 13:47:52
```

Table 7-223 Description of the display ipv6 fib command output

Item	Description
Total number of Routes	Number of IPv6 routes.
Destination	Destination IPv6 address of a route.
PrefixLength	Prefix length of the destination IPv6 address.
Nexthop	Next-hop router that forwards packets to the destination address.
Flag	Description of route characteristics using S/U/G/H/B/D: • S: static • U: Up • G: gateway • H: host • B. blackhole • D: dynamic
Tunnel ID	ID of the tunnel. If packets are forwarded to the next hop using MPLS forwarding, the tunnel ID of the packets cannot be 0. If packets are forwarded using IP forwarding, the tunnel ID of the packets is 0.
TimeStamp	Time taken to generate a route in the FIB table.
Interface	Outbound interface through which packets are forwarded.

7.10.17 display ipv6 routing-table

Function

The **display ipv6 routing-table** command displays information about an IPv6 routing table.

Format

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] [**verbose** | **brief**]

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] *ipv6-address* [*prefix-length*] [**longer-match**] [**verbose** | **brief**]

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] *ipv6-address1* [*prefix-length1*] *ipv6-address2 prefix-length2* [**verbose** | **brief**]

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] **acl** { *acl6-number* | *acl6-name* } [**verbose** | **brief**]

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] **ipv6-prefix** *ipv6-prefix-name* [**verbose** | **brief**]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance of an enabled IPv6 address family.	The value must be an existing VPN instance name.
verbose	Displays detailed information about all the routes in the current routing table, including active and inactive routes.	-
brief	Displays brief information about active routes in the current routing table.	-
ipv6-address	Displays the routes with the specified IPv6 destination address.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length	Specifies the prefix length of an IPv6 destination address.	The value is an integer that ranges from 0 to 128.
longer-match	Displays the routes with the specified destination address and mask.	-

Parameter	Description	Value
ipv6-address1	Specifies the start IPv6 address in an IP address range. *ipv6-address1* and *ipv6-address2* determine an IP address range. Only the routes in the IP address range are displayed.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
ipv6-address2	Specifies the end IPv6 address in an IP address range. *ipv6-address1* and *ipv6-address2* determine an IP address range. Only the routes in the IP address range are displayed.	The value is a 32-digit hexadecimal number, in the format of X:X:X:X:X:X:X:X.
prefix-length1	Specifies the mask length of the start IPv6 address.	The value is an integer that ranges from 0 to 128.
prefix-length2	Specifies the mask length of the end IPv6 address.	The value is an integer that ranges from 0 to 128.
acl *acl6-number*	Displays the routes that match the ACL6 with the specified ACL number. If the specified ACL6 does not exist, information about all active routes is displayed.	The value is an integer that ranges from 2000 to 2999.
acl *acl6-name*	Displays the routes that match the ACL6 with the specified ACL name. If the specified ACL6 does not exist, information about all active routes is displayed.	The value is a string of 1 to 64 case-sensitive characters without spaces. The value must start with a letter.
ipv6-prefix *ipv6-prefix-name*	Displays the routes that match the specified IPv6 prefix list. If the specified IPv6 prefix list does not exist, information about all active routes is displayed.	The name is a string of 1 to 169 case-sensitive characters, with spaces not supported. When double quotation marks are used around the string, spaces are allowed in the string.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

If the **verbose** keyword is not specified, the command output includes the destination address, prefix length, protocol type, preference, cost, next hop, outbound interface, tunnel ID, flag, and status of a route.

NOTE

A recursive route is counted as one route regardless of how many outbound interfaces and next hops the route finds.

Example

Display brief information about the current IPv6 routing table.

```
<HUAWEI> display ipv6 routing-table
Routing Table : Public
         Destinations : 4        Routes : 4

 Destination  : ::1                        PrefixLength : 128
 NextHop      : ::1                        Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : InLoopBack0                Flags        : D

 Destination  : FC00:0:0:112::             PrefixLength : 64
 NextHop      : FC00:0:0:112::2            Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : Vlanif10                   Flags        : D

 Destination  : FC00:0:0:112::2                     PrefixLength : 128
 NextHop      : ::1                        Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : Vlanif10                   Flags        : D

 Destination  : FE80::                     PrefixLength : 10
 NextHop      : ::                         Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : NULL0                      Flags        : D
```

Table 7-224 Description of the display ipv6 routing-table command output

Item	Description
Routing Tables : Public	The routing table is a public routing table.
Destinations	Total number of destination networks or hosts.
Routes	Total number of routes.
Destination	IP address of the destination network or host of a route.
PrefixLength	Prefix length of a route.
NextHop	Next-hop IPv6 address of a route.
Preference	Preference of a route.

Item	Description
Cost	Cost of a route.
Protocol	Routing protocol of a route.
RelayNextHop	Recursive next-hop address.
TunnelID	Tunnel ID. The value 0x0 indicates that no tunnel is used or the tunnel fails to be established.
Interface	Outbound interface through which the next hop of a route can be reached.
Flags	Flag of a route: ● R: The route is a recursive route. ● D: The route is delivered to the FIB table.

Display detailed information about the IPv6 routing table.

```
<HUAWEI> display ipv6 routing-table verbose
Routing Table : Public
         Destinations : 2        Routes : 2

 Destination  : ::1                      PrefixLength : 128
 NextHop      : ::1                      Preference   : 0
 Neighbour    : ::                       ProcessID    : 0
 Label        : NULL                     Protocol     : Direct
 State        : Active NoAdv             Cost         : 0
 Entry ID     : 1                        EntryFlags   : 0x80010050
 Reference Cnt: 1                        Tag          : 0
 Priority     : high                     Age          : 84410sec
 IndirectID   : 0x0
 RelayNextHop : ::                       TunnelID     : 0x0
 Interface    : InLoopBack0              Flags        : D

 Destination  : FC00:0:0:1::1            PrefixLength : 128
 NextHop      : FC00:0:0:2::2            Preference   : 60
 Neighbour    : ::                       ProcessID    : 0
 Label        : NULL                     Protocol     : Static
 State        : Active Adv Relied        Cost         : 0
 Entry ID     : 2                        EntryFlags   : 0x80020140
 Reference Cnt: 1                        Tag          : 0
 Priority     : high                     Age          : 79036sec
 IndirectID   : 0x80000001
 RelayNextHop : ::                       TunnelID     : 0x0
 Interface    : NULL0                    Flags        : RD
```

Table 7-225 Description of the display ipv6 routing-table verbose command output

Item	Description
Neighbour	IPv6 address of a neighbor interface.
ProcessID	Routing protocol process ID of a route.
Label	Label carried in a route.

Item	Description
State	Status of a route: • Active: active route • Invalid: invalid route • Inactive: inactive route • NoAdv: route that cannot be advertised • Adv: route that can be advertised • Del: route to be deleted • Relied: route that recurses to an outbound interface and a next hop or that recurses to a tunnel • Stale: route that is marked Stale and used in GR
Entry ID	ID of a routing entry in the routing table.
EntryFlags	Flag of a routing entry.
Reference Cnt	Number of times a route is referenced.
Tag	Routing management tag. The value is an integer that ranges from 0 to 4294967295.
Priority	Convergence priority of a route: • low • medium • high • critical
IndirectID	ID of indirect next hop.
Age	Time a route is generated.

Display brief information about the active routes that match ACL6 2000.

```
<HUAWEI> display ipv6 routing-table acl 2000
Routes Matched by Access list  2000  :
Summary Count : 2

 Destination  : ::1                        PrefixLength : 128
 NextHop      : ::1                        Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : InLoopBack0                Flags        : D

 Destination  : FC00:0:0:111::             PrefixLength : 64
 NextHop      : FC00:0:0:111::2            Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : Vlanif10                   Flags        : D
```

Display brief information about the routes with the specified IPv6 destination address.


```
<HUAWEI> display ipv6 routing-table fc00:0:0:111::1 64
Routing Table :Public
Summary Count : 1

 Destination  : FC00:0:0:111::              PrefixLength : 64
 NextHop      : FC00:0:0:111::2             Preference   : 0
 Cost         : 0                       Protocol     : Direct
 RelayNextHop : ::                          TunnelID     : 0x0
 Interface    : Vlanif10                    Flags        : D
```

Display the routes within the specified IPv6 address range.

```
<HUAWEI> display ipv6 routing-table fc00:0:0:111::1 64 fc00:0:0:111::2 128
Routing Table :
Summary Count : 2

 Destination  : FC00:0:0:111::              PrefixLength : 64
 NextHop      : FC00:0:0:111::2             Preference   : 0
 Cost         : 0                       Protocol     : Direct
 RelayNextHop : ::                          TunnelID     : 0x0
 Interface    : Vlanif10                    Flags        : D

 Destination  : FC00:0:0:111::2            PrefixLength : 128
 NextHop      : ::1                       Preference   : 0
 Cost         : 0                       Protocol     : Direct
 RelayNextHop : ::                          TunnelID     : 0x0
 Interface    : Vlanif10                    Flags        : D
```

Display brief information about the active routes that match IPv6 prefix list **abc2**.

```
<HUAWEI> display ipv6 routing-table ipv6-prefix abc2
Routes Matched by Prefix-list abc2 :
Summary Count: 3

 Destination  : ::1                       PrefixLength : 128
 NextHop      : ::1                       Preference   : 0
 Cost         : 0                       Protocol     : Direct
 RelayNextHop : ::                          TunnelID     : 0x0
 Interface    : InLoopBack0                 Flags        : D

 Destination  : FC00:0:0:112::              PrefixLength : 64
 NextHop      : FC00:0:0:112::1             Preference   : 0
 Cost         : 0                       Protocol     : Direct
 RelayNextHop : ::                          TunnelID     : 0x0
 Interface    : Vlanif20                    Flags        : D

 Destination  : FC00:0:0:112::1             PrefixLength : 128
 NextHop      : ::1                       Preference   : 0
 Cost         : 0                       Protocol     : Direct
 RelayNextHop : ::                          TunnelID     : 0x0
 Interface    : Vlanif20                    Flags        : D
```

7.10.18 display ipv6 routing-table limit

Function

The **display ipv6 routing-table limit** command displays limits on the numbers of routes and prefixes.

Format

display ipv6 routing-table limit [**all-vpn-instance** | **vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
all-vpn-instance	Indicates all IPv6 VPN instances.	-
vpn-instance *vpn-instance-name*	Specifies the name of an IPv6 VPN instance.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The command displays limits on the number of IPv6 routes and prefixes.

- The **display ipv6 routing-table limit** command displays limits on the number of IPv6 public routes and prefixes.
- The **display ipv6 routing-table limit all-vpn-instance** command displays the limit on the number of routes and prefixes of all IPv6 VPN instances.
- The **display ipv6 routing-table limit vpn-instance** *vpn-instance-name* command displays the limit on the number of routes and prefixes of a specified IPv6 VPN instance.

Example

Display the limits on the number of IPv6 public routes and prefixes.

```
<HUAWEI> display ipv6 routing-table limit
Public Instance:
Limit-Object Limit-Type     Upper-Limit  Warning    Current    Log-Interval
Route        Default        -            -          1          5
Prefix       Default        -            -          1          5
```

Table 7-226 Description of the **display ipv6 routing-table limit** command output

Item	Description
Limit-Object	Indicates the object whose total number is limited: • Prefix • Route

Item	Description
Limit-Type	Indicates the limit mode for the routes and prefixes in the current routing table: • Simply-Alert: indicates that only alarms are generated after the number of routes or prefixes exceeds the upper limit. • Alert-Percent: indicates the percentage of the alarm threshold of routes. • Default: indicates that the number of routes or prefixes is not limited by default.
Upper-Limit	Indicates the upper limit of routes or prefixes in the current routing table.
Warning	Indicates the alarm threshold of routes or prefixes in the current routing table.
Current	Indicates the number of routes or prefixes in the current routing table.
Log-Interval	Indicates the frequency of displaying logs when the number of routes or prefixes in the current routing table exceeds the upper limit, in seconds.

7.10.19 display ipv6 routing-table protocol

Function

The **display ipv6 routing-table protocol** command displays routing information about a specified IPv6 routing protocol.

Format

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] **protocol** *protocol* [**inactive** | **verbose** | **brief**]

Parameters

Parameter	Description	Value
vpn-instance-name	Specifies the name of a VPN instance of an enabled IPv6 address family.	The value must be an existing VPN instance name.

Parameter	Description	Value
protocol	Displays routing information of a specified routing protocol.	The value may be **bgp, direct**, **isis, ospfv3, ripng**, **static**, or **unr**. The specific value varies depending on the routing protocol supported by the device.
inactive	Displays information about inactive routes. If this parameter is not specified, information about all active and inactive routes is displayed.	-
verbose	Displays detailed information about active and inactive routes. If this keyword is not specified, only brief information about active routes is displayed.	-
brief	Displays brief information about active and inactive routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can run the **display ipv6 routing-table protocol** command to check routing information about a specified IPv6 routing protocol.

Example

Display brief information about all IPv6 direct routes.

```
<HUAWEI> display ipv6 routing-table protocol direct
Public Routing Table : Direct
Summary Count : 2

Direct Routing Table's Status : < Active >
Summary Count : 2

 Destination  : ::1                        PrefixLength : 128
 NextHop      : ::1                        Preference   : 0
 Cost         : 0                          Protocol     : Direct
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : InLoopBack0                Flags        : D
```



```
Destination  : FE80::                    PrefixLength : 10
NextHop      : ::                       Preference   : 0
Cost         : 0                        Protocol     : Direct
RelayNextHop : ::                       TunnelID     : 0x0
Interface    : NULL0                    Flags        : D

Direct Routing Table's Status : < Inactive >
Summary Count : 0
```

Table 7-227 Description of the display ipv6 routing-table protocol command output

Item	Description
Public Routing Table	Contents of a public routing table: ● Direct: direct IPv6 route ● Static: static IPv6 route ● bgp: BGP4+ route ● ripng: RIPng route ● isis: IS-IS IPv6 route ● ospfv3: OSPFv3 route ● unr: user network route
Summary Count	Number of prefixes of routes.
Direct Routing Table's Status	Status of direct routes: ● active: information about active routes ● inactive: information about inactive routes
Destination	IP address of the destination network or host of a route.
PrefixLength	Prefix length of a route.
NextHop	Next-hop address of a route.
Preference	Routing protocol preference of a route.
Cost	Cost of a route.
Protocol	Routing protocol that learns a route.
RelayNextHop	Recursive next-hop address.
TunnelID	Tunnel ID. The value 0x0 indicates that no tunnel is used or the tunnel fails to be established.
Interface	Outbound interface through which the next hop of a route can be reached.
Flags	Flag of a route: ● R: The route is a recursive route. ● D: The route is delivered to the FIB table.

7.10.20 display ipv6 routing-table statistics

Function

The **display ipv6 routing-table statistics** command displays statistics about routes in an IPv6 routing table.

Format

display ipv6 routing-table [**vpn-instance** *vpn-instance-name* | **all-vpn-instance**] **statistics**

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays statistics about routes in the routing table of a specified VPN instance. If neither *vpn-instance-name* nor **all-vpn-instance** is specified, statistics about routes in a public routing table are displayed.	The value must be an existing VPN instance name.
all-vpn-instance	Displays statistics about routes in the routing tables of all VPN instances. If neither *vpn-instance-name* nor **all-vpn-instance** is specified, statistics about routes in a public routing table are displayed.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Route statistics include the total number of routes, number of added routes, and number of deleted routes.

Example

Display statistics about routes in an IPv6 routing table.

```
<HUAWEI> display ipv6 routing-table statistics
Summary Prefixes : 0
Protocol   route       active      added       deleted     freed
DIRECT     0           0           0           0           0
STATIC     0           0           0           0           0
RIPng      0           0           0           0           0
OSPFv3     0           0           0           0           0
```

```
IS-IS      0          0          0          0          0
BGP        0          0          0          0          0
UNR        0          0          0          0          0
Total      0          0          0          0          0
```

Table 7-228 Description of the display ipv6 routing-table statistics command output

Item	Description
Summary Prefixes	Total number of prefixes in the current routing table.
Protocol	Routing protocol type:
route	Total number of routes that a routing protocol learns in the routing table, including active and inactive routes.
active	Number of active routes of a routing protocol in the routing table.
added	Number of active and inactive routes that are added to the routing table through a routing protocol.
deleted	Number of routes that are deleted from the routing table.
freed	Number of routes that are permanently deleted from the routing table.

7.10.21 display ipv6 routing-table time-range

Function

The **display ipv6 routing-table time-range** command displays information about routes generated in a specified time range in an IPv6 routing table.

Format

display ipv6 routing-table [**vpn-instance** *vpn-instance-name*] **time-range** *min-age max-age* [**verbose** | **brief**]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays information about routes generated in a specified time range in the IPv6 routing table of the specified VPN instance. If this parameter is not specified, the **display ipv6 routing-table time-range** command displays information about routes generated in a specified time range in the public IPv6 routing table.	The value must be an existing VPN instance name.
min-age	Specifies the end time of the period when routes are generated.	The value is in xxdxxhxxmxxs format. ● d indicates days. The value is an integer that ranges from 0 to 10000. ● h indicates hours. The value is an integer that ranges from 0 to 23. ● m indicates minutes. The value is an integer that ranges from 0 to 59. ● s indicates seconds. The value is an integer that ranges from 0 to 59. For example, enter 5d4h30m20s to specify 5 days, 4 hours, 30 minutes, and 20 seconds. **NOTE** If the value of d is 10000, the values of h, m, and s can only be 0.

Parameter	Description	Value
max-age	Specifies the start time of the period when routes are generated.	The value is in xxdxxhxxmxxs format. ● d indicates days. The value is an integer that ranges from 0 to 10000. ● h indicates hours. The value is an integer that ranges from 0 to 23. ● m indicates minutes. The value is an integer that ranges from 0 to 59. ● s indicates seconds. The value is an integer that ranges from 0 to 59. For example, enter 5d4h30m20s to specify 5 days, 4 hours, 30 minutes, and 20 seconds. **NOTE** If the value of d is 10000, the values of h, m, and s can only be 0.
verbose	Displays detailed information about active and inactive routes. If this parameter is not specified, the **display ipv6 routing-table time-range** command displays only brief information about active routes.	-
brief	Displays brief information about active and inactive routes.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Usage Scenario

If route flapping occurs on a network, to find the flapping route rapidly and accelerate fault location, run the **display ipv6 routing-table time-range** command and specify a small time range.

Precautions

max-age must be larger than *min-age*. Otherwise, the **display ipv6 routing-table time-range** command does not display any information.

If *max-age* is larger than *min-age* but no route is generated within this time range, the **display ipv6 routing-table time-range** command displays only the table heading.

Example

Display information about public network IPv6 routes generated in the last 2 hours, 20 minutes, and 10 seconds.

```
<Switch> display ipv6 routing-table time-range 0 2h20m10s
Routing Table : Public
        Destinations : 4        Routes : 4

Destination  : ::1                         PrefixLength : 128
NextHop      : ::1                         Preference   : 0
Cost         : 0                           Protocol     : Direct
RelayNextHop : ::                          TunnelID     : 0x0
Interface    : InLoopBack0                 Flags        : D

Destination  : FC00:0:0:1::1                PrefixLength : 128
NextHop      : ::1                         Preference   : 0
Cost         : 0                           Protocol     : Direct
RelayNextHop : ::                          TunnelID     : 0x0
Interface    : LoopBack2                   Flags        : D

Destination  : FC00:0:0:1::2                PrefixLength : 128
NextHop      : ::                          Preference   : 60
Cost         : 0                           Protocol     : Static
RelayNextHop : ::                          TunnelID     : 0x0
Interface    : NULL0                       Flags        : D

Destination  : FE80::                       PrefixLength : 10
NextHop      : ::                          Preference   : 0
Cost         : 0                           Protocol     : Direct
RelayNextHop : ::                          TunnelID     : 0x0
Interface    : NULL0                       Flags        : D
```

Table 7-229 Description of the display ipv6 routing-table time-range command output

Item	Description
Routing Tables: Public	The routing table is a public routing table.
Destinations	Number of destination networks or hosts.
Routes	Number of routes.
Destination	IP addresses of the destination network or host.
PrefixLength	Prefix length of a route.
NextHop	Next-hop IPv6 address.
Preference	Route preference.
Cost	Route cost.

Item	Description
Protocol	Name of the routing protocol.
RelayNextHop	Recursive next hop.
TunnelID	Tunnel ID: • The value 0x0 indicates that the route does not use a tunnel or the tunnel fails to be set up. • If the value is not 0x0, the route recurses to a tunnel.
Interface	Outbound interface through which the next hop of a route can be reached.
Flags	Route flag: • R: The route is a recursive route. • D: The route is delivered to the FIB table.

Display detailed information about public network routes generated in the last 20 minutes.

```
<Switch> display ipv6 routing-table time-range 0 20m verbose
Routing Table : Public
         Destinations : 4        Routes : 5

 Destination  : ::1                        PrefixLength : 128
 NextHop      : ::1                        Preference   : 0
 Neighbour    : ::                         ProcessID    : 0
 Label        : NULL                       Protocol     : Direct
 State        : Active NoAdv               Cost         : 0
 Entry ID     : 262551588                  EntryFlags   : 0x80010050
 Reference Cnt: 2                          Tag          : 0
 Priority     : high                       Age          : 7694sec
 IndirectID   : 0x0
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : InLoopBack0                Flags        : D

 Destination  : FC00:0:0:1::1              PrefixLength : 128
 NextHop      : ::1                        Preference   : 0
 Neighbour    : ::                         ProcessID    : 0
 Label        : NULL                       Protocol     : Direct
 State        : Active Adv                 Cost         : 0
 Entry ID     : 262552100                  EntryFlags   : 0x80010040
 Reference Cnt: 2                          Tag          : 0
 Priority     : high                       Age          : 101sec
 IndirectID   : 0x0
 RelayNextHop : ::                         TunnelID     : 0x0
 Interface    : LoopBack2                  Flags        : D

 Destination  : FC00:0:0:1::1              PrefixLength : 128
 NextHop      : FC00:0:0:1::2              Preference   : 60
 Neighbour    : ::                         ProcessID    : 0
 Label        : NULL                       Protocol     : Static
 State        : Inactive Adv Relied        Cost         : 0
 Entry ID     : 262551716                  EntryFlags   : 0x20140
 Reference Cnt: 2                          Tag          : 0
 Priority     : medium                     Age          : 7668sec
 IndirectID   : 0x80000001
 RelayNextHop : ::                         TunnelID     : 0x0
```



```
Interface    : NULL0                    Flags       : R

Destination  : FC00:0:0:1::2              PrefixLength : 128
NextHop      : ::                       Preference   : 60
Neighbour    : ::                       ProcessID    : 0
Label        : NULL                     Protocol     : Static
State        : Active Adv               Cost         : 0
Entry ID     : 262551844                EntryFlags   : 0x80000040
Reference Cnt: 2                        Tag          : 0
Priority     : medium                   Age          : 7670sec
IndirectID   : 0x0
RelayNextHop : ::                       TunnelID     : 0x0
Interface    : NULL0                    Flags        : D

Destination  : FE80::                   PrefixLength : 10
NextHop      : ::                       Preference   : 0
Neighbour    : ::                       ProcessID    : 0
Label        : NULL                     Protocol     : Direct
State        : Active NoAdv             Cost         : 0
Entry ID     : 262551972                EntryFlags   : 0x80010010
Reference Cnt: 2                        Tag          : 0
Priority     : high                     Age          : 104sec
IndirectID   : 0x0
RelayNextHop : ::                       TunnelID     : 0x0
Interface    : NULL0                    Flags        : D
```

Table 7-230 Description of the display ipv6 routing-table time-range verbose command output

Item	Description
Neighbour	IP address of a neighboring interface. :: indicates that the route is generated by a local device.
ProcessID	Process ID of the routing protocol.
Label	Label carried in a route.
State	Status of the route: ● Active: an active route ● Invalid: an invalid route ● Inactive: an inactive route ● NoAdv: a route that cannot be advertised ● Adv: a route that can be advertised ● Del: a route to be deleted ● Relied: a route that recurses to the next hop and outbound interface or that recurses to a tunnel ● Stale: a route with the Stale flag and used in GR
Entry ID	ID of a route in the routing table.
EntryFlags	Flag of a route.

Item	Description
Reference Cnt	Number of times the route is referenced.
Tag	Administrative tag of the route. The value is an integer that ranges from 0 to 4294967295.
Priority	Convergence priority of the route: • low • medium • high • critical
IndirectID	ID of the indirect next hop.
Age	Time when the route is generated.

7.10.22 display rm bfd-session

Function

The **display rm bfd-session** command displays the BFD session configuration in routing management (RM).

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

display rm bfd-session [**vpn-instance** *vpn-instance-name*] [**destination** *destination-address*] [**source** *source-address*] [**interface** *interface-type interface-number*] [**protocol** { **bgp** | **isis-l1** | **isis-l2** | **isis-l1l2** | **ospf** | **rip** | **pim** }]

display rm bfd-session all

Parameters

Parameter	Description	Value
all	Displays all the BFD session configurations in RM, including the BFD session configurations of the public and private networks. If **all** is not specified, only the BFD session configuration on the public network is displayed.	-
vpn-instance *vpn-instance-name*	Displays the configurations of BFD sessions of the specified VPN instance in RM.	The value must be an existing VPN instance name.
destination *destination-address*	Displays the configurations of BFD sessions with the specified destination address in RM.	The value is in dotted decimal notation.
source *source-address*	Displays the configurations of BFD sessions with the specified source address in RM.	The value is in dotted decimal notation.
interface *interface-type interface-number*	Displays the configurations of BFD sessions with the specified outbound interface type and number in RM.	-
protocol	Displays the configurations of BFD sessions of the specified routing protocol in RM: • **bgp**: BGP • **isis-l1**: IS-IS Level-1 • **isis-l2**: IS-IS Level-2 • **isis-l1l2**: IS-IS Level-1-2 • **ospf**: OSPF • **rip**: RIP • **pim**: PIM	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use this command to view the configuration of BFD sessions dynamically created by routing protocols. The command output includes the global BFD status, number of BFD sessions, and BFD session configuration, such as VPN instance, destination address, source address, interface, and status.

Example

Display the BFD configuration in RM.

```
<HUAWEI> display rm bfd-session vpn-instance vrf1 destination 10.8.1.2 source 10.8.1.1 interface
Vlanif10
BFD Global Status: ON
BFD Session Total Number: 1

BFD Session: 1
  Interface: Vlanif10    VRF Name: vrf1
  Destination     Source          Session-State Create-State
  10.8.1.2        10.8.1.1        Up            Established
  Route-Protocol  Rx-Interval(ms) Tx-Interval(ms) Multiplier
  BGP             1000            1000           5
```

Table 7-231 Description of the display rm bfd-session command output

Item	Description
BFD Global Status	Global BFD status: • ON: BFD is enabled globally. • OFF: BFD is disabled globally. To set the global BFD status, run the **bfd** command.
BFD Session Total Number	Total number of BFD sessions.
BFD Session	Number of the BFD session.
Interface	Local physical interface bound to the BFD session.
VRF Name	Name of a VPN instance to which the BFD session is bound to.
Destination	Peer IP address bound to the BFD session.
Source	Source IP address bound to the BFD session.
Session-State	Status of the BFD session: • AdminDown: indicates that the BFD session is in AdminDown state when the **shutdown** command is run. • Down: indicates that the BFD session in Down state. • Init: indicates that the BFD session is in Init state. • Up: indicates that the BFD session is in Up state.

Item	Description
Create-State	BFD session setup status in RM: • Creating: indicates that the routing protocol notifies RM of the session, but RM does not instruct BFD to set up a session. • Established: indicates that RM notifies BFD to set up a BFD session. You can view the Session-State field to check whether a BFD session is set up.
Route-Protocol	Routing protocol enabled with BFD: • BGP • IS-IS Level-1 • IS-IS Level-2 • IS-IS Level-1-2 • OSPF • RIP • PIM
Rx-Interval(ms)	Configured receiving interval, in milliseconds.
Tx-Interval(ms)	Configured sending interval, in milliseconds.
Multiplier	Configured local detection multiplier.

7.10.23 display rm interface

Function

The **display rm interface** command displays routing management (RM) information on an interface.

Format

display rm interface [*interface-type interface-number*]

display rm interface [**vpn-instance** *vpn-instance-name*] [**ip-address** *ip-address*]

Parameters

Parameter	Description	Value
interface-type interface-number	Displays RM information on the specified interface.	-
vpn-instance *vpn-instance-name*	Displays RM information on the interface of the specified VPN instance.	The value must be an existing VPN instance name.
ip-address *ip-address*	Displays RM information on the interface with the specified destination address.	The value is in dotted decimal notation.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use this command to view IPv4 RM information on an interface, including information about physical and logical interfaces of the interface. This information helps locate routing problems.

Example

Display RM information on interfaces.

```
<HUAWEI> display rm interface

Name: InLoopBack0
Physical IF Info:
 IfnetIndex: 0x1
 State: UP LOOP MULT
 Slot: 0(Logic Slot: 9)
 IntType: 26, PriLog: 1, MTU: 1500, Reference Count: 7
 Bandwidth: 0, 0
 Baudrate: 0, 0
 Delay: 0, Reliability: 0, Load: 0, Quality: 0

 LDP-ISIS sync capability: disabled
 LDP-OSPF sync capability: disabled
 InstanceID: 0, Instance Name: Public
 Age: 75455sec
Logical IF Info:
 IfnetIndex: 0x1, PhyIndex: 1 Logical Index : 1,
 Dest: 127.0.0.1, Mask: 255.0.0.0
 State: UP LOOP PRM MULT , Reference Count 6
 Age: 75455sec

Name: NULL0
Physical IF Info:
```



```
 IfnetIndex: 0x2
 State: UP NBMA MULT
 Slot: 0(Logic Slot: 9)
 IntType: 27, PriLog: 2, MTU: 1500, Reference Count 1
 Bandwidth: 0, 0
 Baudrate: 0, 0
 Delay: 0, Reliability: 0, Load: 0, Quality: 0

 LDP-ISIS sync capability: disabled
 LDP-OSPF sync capability: disabled
 InstanceID: 0, Instance Name: Public
 Age: 75455sec
Logical IF Info:
 IfnetIndex: 0x2, PhyIndex: 2 Logical Index : 2,
 Dest: 0.0.0.0, Mask: 255.255.255.255
 State: UP PRM NBMA MULT , Reference Count: 0
 Age: 75456sec

Name: MEth0/0/1                                          Physical IF
Info:
 IfnetIndex: 0x4
 State: UP BCA MULT OSI
 Hardware Address: 000e0-fc1-2123
 Slot: 0(Logic Slot: 9)
 IntType: 1, PriLog: 6, MTU: 1500, Reference Count: 4
 Bandwidth: 0, 0
 Baudrate: 0, 0
 Delay: 0, Reliability: 0, Load: 0, Quality: 0

 LDP-ISIS sync capability: disabled
 LDP-OSPF sync capability: disabled
 InstanceID: 0, Instance Name: Public
 Age: 75423sec
Logical IF Info:
 IfnetIndex: 0x4, PhyIndex: 3 Logical Index : 6,
 Dest: 10.137.217.207, Mask: 255.255.254.0
 State: UP PRM BCA MULT , Reference Count: 3
 Age: 75447sec
```

Display RM information on VLANIF100.

```
<HUAWEI> display rm interface vlanif 100
Name: vlanif100
Physical IF Info:
 IfnetIndex: 0x23
 State: DOWN BCA MULT
 Hardware Address: 00E0-FCCB-7543
 Slot: 0(Logic Slot: 9)
 IntType: 38, PriLog: 3, MTU: 1500, Reference Count: 1
 Bandwidth: 0, 1000000000
 Baudrate: 0, 1000000000
 Delay: 0, Reliability: 0, Load: 0, Quality: 0
 LDP-ISIS sync capability: disabled
 LDP-OSPF sync capability: disabled
 InstanceID: 0, Instance Name: Public
 Age: 2152167sec
Logical IF Info:
 IfnetIndex: 0x23, PhyIndex: 6, Logical Index: 3
 Dest: 10.20.10.2, Mask: 255.255.0.0
 State: DOWN PRM BCA MULT , Reference Count:
0
 Age: 2152166sec
```


Table 7-232 Description of the display rm interface command output

Item	Description
Name	Name of the interface.
Physical IF Info	Physical interface information.
IfnetIndex	Network segment index of the interface.
State	Current interface status.
Hardware Address	MAC address of the interface.
Slot	Slot ID of the interface.
IntType	Type of the interface.
PriLog	Index of the primary logical interface.
MTU	Maximum transmission unit (MTU) of the interface.
Reference Count	Number of times the interface is referenced.
Bandwidth	Bandwidth of the interface.
Baudrate	Baud rate of the interface.
Delay	Link delay.
Reliability	Link reliability.
Load	Link load.
Quality	Link quality. Currently: • 0: indicates the link quality is "GOOD". • 4: indicates the link quality is "LOW".
LDP-ISIS sync capability	Whether synchronization between LDP and IS-IS is enabled: • enabled: indicates that synchronization is enabled. • disabled: indicates that synchronization is disabled.
LDP-OSPF sync capability	Whether synchronization between LDP and OSPF is enabled: • enabled: indicates that synchronization is enabled. • disabled: indicates that synchronization is disabled.
InstanceID	Instance ID.
Instance Name	Instance name.

Item	Description
Logical IF Info	Logical interface information.
PhyIndex	Physical interface index.
Logical Index	Logical interface index.
Dest	Destination address.
Mask	Mask of the destination address.
Age	Lifetime of the displayed information.

7.10.24 display rm ipv6 bfd-session

Function

The **display rm ipv6 bfd-session** command displays the configurations of BFD sessions in Route Management (RM).

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

display rm ipv6 bfd-session [**all** | [[**vpn-instance** *vpn-instance-name*] [**destination** *destination-address*] [**source** *source-address*] [**interface** *interface-type interface-number*] [**protocol** { **bgp** | **isis-l1** | **isis-l2** | **isis-l1l2** | **ospfv3** | **pim** }]]]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
destination *destination-address*	Specifies the remote destination address.	-
source *source-address*	Specifies the local source address.	-
interface *interface-type interface-number*	Specifies the type and number of an outbound interface.	-
protocol	Indicates a routing protocol. • **bgp**: indicates the Border Gateway Protocol (BGP). • **isis-l1**: indicates IS-IS Level-1. • **isis-l2**: indicates IS-IS Level-2. • **Isis-l1l2**: indicates IS-IS Level-1-2. • **ospfv3**: indicates the Open Shortest Path First (OSPF) version 3 protocol. • **pim**: indicates the Protocol Independent Multicast (PIM) protocol.	-
all	Displays all the configurations of BFD sessions in RM, including the BFD sessions of the public network and VPN. If **all** is not specified, only the configurations of BFD sessions in the public network are displayed.	-

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use the command to view the BFD sessions that are dynamically created by the routing protocol. The configurations include the global status of BFD, number of BFD sessions, and VPN instance, destination address, source address, interface, and session status of each BFD session. To modify the parameters of a BFD session, refer to the **bfd all-interfaces**, **ospfv3 bfd**, **isis ipv6 bfd**, and **peer bfd**.

Example

Display the configurations of BFD sessions in RM.

```
<HUAWEI> display rm ipv6 bfd-session vpn-instance a destination fc00:0:0::1 source fc00:0:0::2
interface GigabitEthernet0/0/1


BFD Global Status: ON
RM IPv6 BFD Session Total Number: 1
Destination: fc00:0:0::1
Source     : fc00:0:0::2
  Session-State      Local-Discr      Interface             VPN-Name
  Up                 8192             GigabitEthernet0/0/1  a
  Route-Protocol     Rx-Interval(ms)  Tx-Interval(ms)       Multiplier
  BGP                1000             1000                  7
```

Table 7-233 Description of the **display rm ipv6 bfd-session** command output

Item	Description
BFD Global Status	Global status of BFD: ● On: indicates that BFD is enabled globally. ● Off: indicates that BFD is disabled globally.
RM IPv6 BFD Session Total Number	Total number of BFD sessions.
Destination	Destination IP address bound to the BFD session.
Source	Source IP address bound to the BFD session.
Session-State	Current status of a BFD session: ● AdmDown: If the **shutdown (BFD session view)** command is used, the BFD session enters the AdmDown state. ● Down: the BFD session is in the Down state. ● Init: the BFD session is in the Init state. ● Up: the BFD session is in the Up state.
Local-Discr	Local discriminator of the BFD session. ● If the value is 0, it indicates that the RM is notified of the need to set up of a BFD session, but the RM has not instructed BFD to set up the session. ● If the value is not 0, it indicates that RM has instructed BFD to set up a BFD session. If you need to learn the establishment of the BFD session, you can view the Session-State.
Interface	Local physical interface bound to the BFD session.

Item	Description
Route-Protocol	Routing protocol enabled with BFD: • BGP • IS-IS Level-1 • IS-IS Level-2 • IS-IS Level-1-2 • OSPFv3 • PIM
Rx-Interval (ms)	Configured receiving interval in milliseconds.
Tx-Interval (ms)	Configured sending interval in milliseconds.
Multiplier	Configured local detection multiplier.
VPN-Name	Name of the VPN instance.

7.10.25 display rm ipv6 interface

Function

The **display rm ipv6 interface** command displays IPv6 routing management (RM) information on an interface.

Format

display rm ipv6 interface [*interface-type interface-number*]

display rm ipv6 interface [**vpn-instance** *vpn-instance-name*] [**ipv6-address** *ipv6-address*]

Parameters

Parameter	Description	Value
interface-type interface-number	Displays IPv6 RM information on the specified interface.	-
vpn-instance *vpn-instance-name*	Displays IPv6 RM information in a specified VPN instance.	The value must be an existing VPN instance name.
ipv6-address *ipv6-address*	Displays IPv6 RM information with the specified destination IPv6 address.	The value is a 32-digit hexadecimal number, in the X:X:X:X:X:X:X:X format.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

You can use this command to view IPv6 RM information on an interface, including information about physical and logical interfaces of the interface. This information helps locate routing problems.

Example

Display IPv6 RM information on all interfaces.

```
<HUAWEI> display rm ipv6 interface

Name: InLoopBack0
Physical IF Info:
 IfnetIndex: 0x1
 State: UP LOOP MULT
 Slot: 0(Logic Slot: 0)
 IntType: 26, PriLog: 1, MTU: 1500, Ref: 1
 Bandwidth: 0, 0
 Baudrate: 0, 0
 Delay: 0, Reliability: 0, Load: 0
 Age: 1635sec

Logical IF Info:
 IfnetIndex: 0x1, PhyIndex: 1, LogiIndex: 1
 Dest: ::1, Mask: 128
 State: UP LOOP PRM MULT , Ref 1
 Age: 1635sec

Name: NULL0
Physical IF Info:
 IfnetIndex: 0x2
 State: UP NBMA MULT
 Slot: 0(Logic Slot: 0)
 IntType: 27, PriLog: 2, MTU: 1500, Ref: 0
 Bandwidth: 0, 0
 Baudrate: 0, 0
 Delay: 0, Reliability: 0, Load: 0
 Age: 1635sec

Logical IF Info:
 IfnetIndex: 0x2, PhyIndex: 2, LogiIndex: 2
 Dest: ::, Mask: 128
 State: UP PRM NBMA MULT , Ref: 0
 Age: 1635sec
```

Table 7-234 Description of the display rm ipv6 interface command output

Item	Description
Name	Name of an interface.
Physical IF Info	Physical interface information.

Item	Description
IfnetIndex	Network segment index of the interface.
State	Interface status.
IntType	Interface type.
Slot	Slot ID.
PriLog	Index of the primary logical interface.
MTU	Maximum transmission unit (MTU) of the interface.
Ref	Number of times the interface is referenced.
Bandwidth	Bandwidth of the interface.
Baudrate	Baud rate of the interface.
Delay	Link delay.
Reliability	Link reliability.
Load	Link load.
Age	Lifetime of the displayed information.
Logical IF Info	Logical interface information.
PhyIndex	Physical index.
LogiIndex	Logical index.
Dest	Destination IPv6 address.
Mask	Prefix length of the destination IPv6 address.

7.10.26 display router id

Function

The **display router id** command displays the configured router ID.

Format

display router id [**vpn-instance** *vpn-instance-name*]

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Displays the configured router ID of the specified VPN instance. If this parameter is not specified, the configured router ID of the public network is displayed.	The value must be an existing VPN instance name.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

Some dynamic routing protocols require a router ID to uniquely identify a device. If no router ID is specified when these routing protocols are enabled, the global router ID is used by default. You can run the **display router id** command to check the global router ID.

Example

Display the configured router ID.

```
<HUAWEI> display router id
RouterID:10.1.1.1
```

7.10.27 display route resource

Function

The **display route resource** command displays the total number of IPv4 and IPv6 route prefixes in the IP routing table on the device.

Format

display route resource

Parameters

None

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

The **display route resource** command displays the total number of IPv4 and IPv6 route prefixes in the IP routing table on the device.

Example

Display the total number of IPv4 and IPv6 route prefixes in the IP routing table on the device.

```
<HUAWEI> display route resource
 Total number of IPv4 Prefixes: 16
 Total number of IPv6 Prefixes: 13
```

Table 7-235 Description of the **display route resource** command output

Item	Description
Total number of IPv4 Prefixes	Current number of IPv4 route prefixes.
Total number of IPv6 Prefixes	Current number of IPv6 route prefixes.

7.10.28 ecmp load-balance

Function

The **ecmp load-balance** command sets the ECMP load balancing mode.

The **undo ecmp load-balance** command deletes the configured ECMP load balancing mode.

By default, ECMP load balancing is performed on packets based on the source IP address, destination IP address, transport-layer source port number and destination port number.

Product	Support
S5731-S, S5731S-S, S5731-H, S5731S-H, S5732-H, S5735-S, S5735S-S, S5735-S-I, S6735-S, S6720-EI, S6720S-EI, S6730S-H, S6730-H, S6730-S, and S6730S-S	Supported
SS1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735S-H, S5736-S and S6720S-S	Not supported

Format

ecmp load-balance sip [**dip**] [**port**]

undo ecmp load-balance sip [**dip**] [**port**]

Parameters

Parameter	Description	Value
sip	Configures ECMP load balancing based on the source IP address.	-
dip	Configures ECMP load balancing based on the destination IP address.	-
port	Configures ECMP load balancing based on the transport-layer source port number and destination port number.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

If you run the **ecmp load-balance** command in the system view multiple times, only the latest configuration takes effect.

Example

Configure ECMP load balancing based on the source IP address.

```
<HUAWEI> system-view
[HUAWEI] ecmp load-balance sip
```

7.10.29 ecmp load-balance diffluence

Function

The **ecmp load-balance diffluence** command configures ECMP-based distribution of packets with the same source address and destination address.

The **undo ecmp load-balance diffluence** command cancels ECMP-based distribution of packets with the same source address and destination address.

By default, ECMP-based distribution of packets with the same source address and destination address is not configured.

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6730-H, S6730S-H, S6730-S, and S6730S-S support this command.

Format

ecmp load-balance diffluence

undo ecmp load-balance diffluence

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

In addition to analyzing the unidirectional traffic of the two communication parties, a packet analysis device needs to analyze traffic between the two communication parties so that the traffic information can be fully analyzed. In this case, the packets of the two communication parties need to be distributed to the same traffic distribution server. The traffic distribution device is required to support the algorithm based on the same source and destination in the specific forwarding procedure. Identical source and destination indicates that bidirectional data packets of a network connection must be sent out from the same outbound interface. To configure ECMP-based traffic distribution of packets with the same source addresses and destination addresses, run the **ecmp load-balance diffluence** command.

Precautions

This function may cause uneven load balancing among IPv4 and IPv6 routes.

In a VXLAN scenario, ECMP-based distribution of packets with the same source address and destination address does not take effect; instead, the configured ECMP load balancing mode takes effect..

Example

Configure ECMP-based distribution of packets with the same source addresses and destination addresses.

```
<HUAWEI> system-view
[HUAWEI] ecmp load-balance diffluence
```


7.10.30 ecmp local-preference disable

Function

The **ecmp local-preference disable** command disables ECMP local preferential forwarding.

The **undo ecmp local-preference disable** command enables ECMP local preferential forwarding.

By default, ECMP local preferential forwarding is enabled.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ecmp local-preference disable

undo ecmp local-preference disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Generally, you can use an Eth-Trunk or equal-cost routes to implement traffic load balancing. In a stack, when an inter-chassis Eth-Trunk is unavailable, you can use equal-cost routes for traffic load balancing. Because traffic will be forwarded between member switches, during a stack switchover, traffic switching depends on route convergence, causing some traffic to be lost. In this case, use ECMP local

preferential forwarding to improve route convergence performance during a stack switchover and reduce traffic loss. In ECMP local preferential forwarding, traffic reaching the local switch is preferentially forwarded through a local interface. If the local outbound interface fails, traffic is forwarded through an interface on another member switch.

In **Figure 7-4**, SwitchA and SwitchB set up a stack and implement ECMP load balancing with upstream devices. If local preferential forwarding is not configured, traffic reaching SwitchA is load balanced based on ECMP, and some traffic is forwarded through stack cables and sent out from a physical interface on SwitchB. After local preferential forwarding is configured, traffic reaching SwitchA is forwarded through a local physical interface instead of a physical interface on SwitchB.

Figure 7-4 ECMP local preferential forwarding

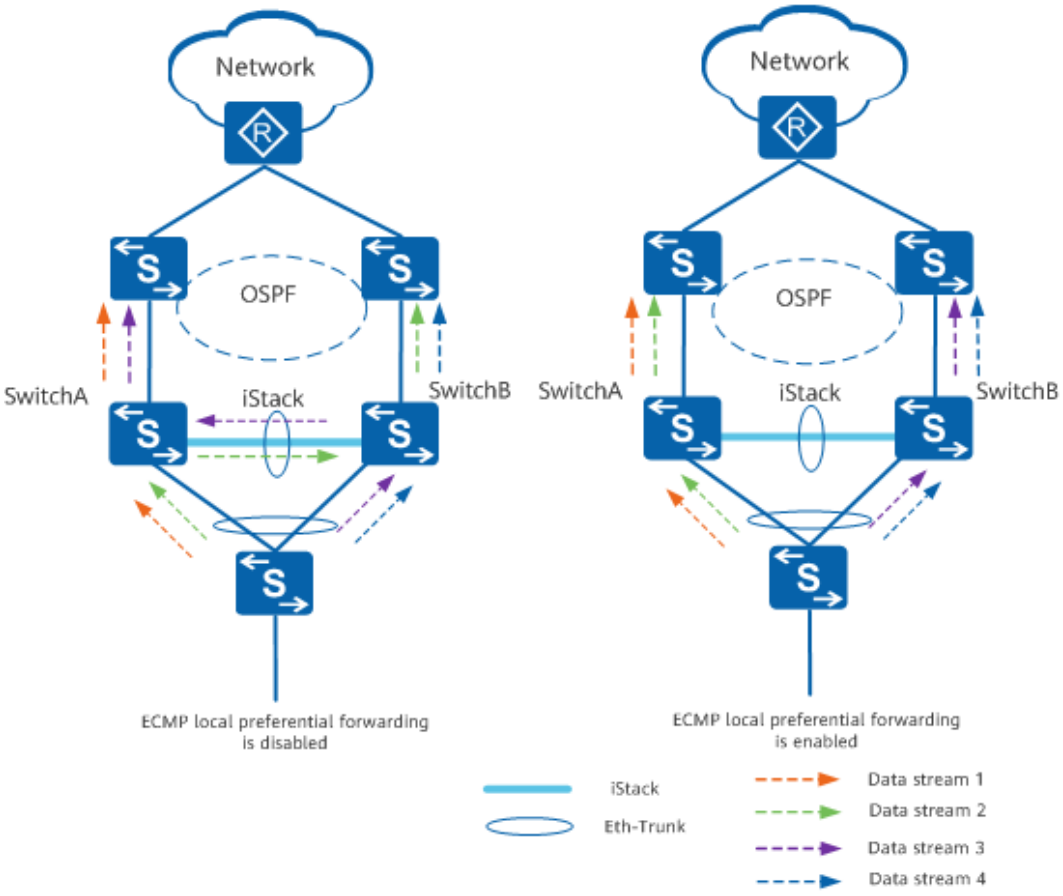

Configuration Impact

After ECMP local preferential forwarding is enabled in a stack, traffic entering from the local device is preferentially forwarded through a local interface. Therefore, even load balancing of the traffic cannot be guaranteed. To address this issue, you can run the **ecmp local-preference disable** command to disable ECMP local preferential forwarding.

Precautions

- Only a stack supports ECMP local preferential forwarding.
- After the system software of a switch is upgraded to V200R010C00 or later, the **ecmp local-preference disable** configuration is automatically generated.
- Enabling or disabling ECMP local preferential forwarding takes effect immediately for newly added routes and takes effect for existing routes only after these routes are re-advertised.
- ECMP local preferential forwarding does not take effect in the following scenarios:
  - No ARP or ND entry for the next hop of a route is found.
  - The next hop of a route is a blackhole.
  - The next-hop outbound interface of a route is an Eth-Trunk interface.

Example

Disable ECMP local preferential forwarding.

```
<HUAWEI> system-view
[HUAWEI] ecmp local-preference disable
Warning: The operation will take effect only on routing tables generated later.
```

7.10.31 fib regularly-refresh

Function

The **fib regularly-refresh** command configures the entire FIB entry update interval, FIB entry update interval per cycle, and the number of FIB entries updated per cycle.

The **undo fib regularly-refresh** command restores the default entire FIB entry update interval, FIB entry update interval per cycle, and the number of FIB entries updated per cycle.

By default, the FIB entry update interval per cycle is 1 second, and the number of FIB entries updated per cycle is 50. The entire FIB entry update interval is 1 minute on the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H, and is fixed to 1440 minutes on other devices.

Format

fib regularly-refresh { **interval** *interval* [**entry-number** *entry-number*] | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* }

undo fib regularly-refresh { **interval** *interval* [**entry-number** *entry-number*] | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* }

Parameters

Parameter	Description	Value
interval *interval*	Specifies an interval at which FIB entries are refreshed.	The value is an integer ranging from 1 to 300, in seconds.

Parameter	Description	Value
entry-number *entry-number*	Specifies the number of FIB entries refreshed per circle.	The value is an integer ranging from 20 to 2000.
cycle-interval *cycle-interval*	Specifies the entire interval at which FIB entries are updated.	The value is an integer ranging from 1 to 1440, in minutes. **NOTE** This parameter is supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Regular FIB entry updating occupies system resources, including CPU and memory resources. A longer update interval and a smaller number of updated entries result in less impact on other services. You can flexibly configure the update interval and the number of entries to be updated based on system resources.

To allow underlying data to be forwarded in real-time and effective ways, you can configure a device to update FIB entries in cycles. The *cycle-interval* parameter specifies the interval for FIB entry update per cycle. Note that the *cycle-interval*, *interval*, and *entry-number* parameters together determine the FIB entry update process. As an example, assume that a device has 500 FIB entries.

- If *interval* is set to 1 second, *entry-number* is set to 50, and *cycle-interval* is set to 1 minute, the 500 FIB entries are updated in 10 cycles (with 50 FIB entries updated in each cycle), and the entry update can be completed within 10s (1s for each cycle). The device stays idle from 10s to 1 minute (entire FIB entry update interval).
- If *interval* is set to 10 seconds, *entry-number* is set to 50, and *cycle-interval* is set to 1 minute, the 500 FIB entries are still updated in 10 cycles, taking a total of 100 seconds. In this case, the update period exceeds 1 minute, and the device cycles update from the next minute (a total of 2 minutes at this point). That is, the device updates FIB entries at an interval of 2 minutes.

Prerequisites

Regular FIB entry updating has been enabled using the **undo fib regularly-refresh disable** command in the system view. This function is enabled by default.

Example

Set an interval at which FIB entries are updated to 5 seconds and the number of FIB entries updated per circle to 200.

```
<HUAWEI> system-view
[HUAWEI] fib regularly-refresh interval 5 entry-number 200
```

7.10.32 fib regularly-refresh disable

Function

The **fib regularly-refresh disable** command disables the FIB entry update function.

The **undo fib regularly-refresh disable** command enables the FIB entry update function.

By default, the FIB entry update function is enabled.

Format

fib regularly-refresh disable

undo fib regularly-refresh disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, a device starts FIB entry update 30 minutes after a device is powered on, and the update is performed regularly at an interval of 1 minute. Updating FIB entries occupy system resources. To avoid a high CPU usage within a short period, run the **fib regularly-refresh** { **interval** *interval* [**entry-number** *entry-number*] | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* } command to set the entire FIB entry update interval or the FIB entry update interval per cycle to a larger value or reduce the number of FIB entries updated per cycle. If the CPU usage problem still persists, run the **fib regularly-refresh disable** command to disable the FIB entry update function.

Precautions

Disabling the FIB entry update function may cause a failure to rectify certain hardware faults in time. Exercise caution when running the **fib regularly-refresh disable** command.

Example

Disable the FIB from being refreshed.

```
<HUAWEI> system-view
[HUAWEI] fib regularly-refresh disable
```

7.10.33 fib threshold-alarm

Function

The **fib threshold-alarm** command enables the alarm function for IPv4 route prefix usage and sets the alarm threshold.

The **undo fib threshold-alarm** command restores the default alarm threshold.

By default, the alarm function for IPv4 route prefix usage is enabled. The upper alarm threshold is 85% and the lower alarm threshold is 75%.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

fib threshold-alarm upper-limit *upper-limit-value* **lower-limit** *lower-limit-value*

undo fib threshold-alarm

Parameters

Parameter	Description	Value
upper-limit *upper-limit-value*	Sets the upper alarm threshold for IPv4 route prefix usage.	The value is an integer that ranges from 1 to 100, in percentage.

Parameter	Description	Value
lower-limit *lower-limit-value*	Sets the lower alarm threshold for IPv4 route prefix usage.	The value is an integer that ranges from 1 to 100, in percentage.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Overload of the IPv4 route prefix usage will cause entries unable to be delivered and result in a forwarding failure.

- When the upper alarm threshold for IPv4 route prefix usage is smaller than 100%, an alarm is generated when the resource usage exceeds the upper alarm threshold or reaches 100%.
- When the upper alarm threshold for IPv4 route prefix usage is set to 100%, an alarm is generated when the resource usage reaches 100%.

Precautions

In versions earlier than V200R019C00SPC300, the upper alarm threshold for the IPv4 route prefix usage can be greater than or equal to the lower alarm threshold. In V200R019C00SPC300 and later versions, the upper alarm threshold for the IPv4 route prefix usage must be greater than the lower alarm threshold. If the upper alarm threshold for the IPv4 route prefix usage is configured to be equal to the lower alarm threshold on a switch running a version earlier than V200R019C00SPC300, the configuration will be lost after the switch is upgraded to V200R019C00SPC300 or a later version, and the upper and lower alarm thresholds will be restored to the default settings.

Example

Enable the alarm function for IPv4 route prefix usage and set the upper and lower alarm thresholds to 85% and 60% respectively.

```
<HUAWEI> system-view
[HUAWEI] fib threshold-alarm upper-limit 85 lower-limit 60
```

7.10.34 fib trap lpm-fail enable

Function

The **fib trap lpm-fail enable** command enables the trap function for the event that FIB entries fail to be delivered.

The **undo fib trap lpm-fail enable** command disables the trap function for the event that FIB entries fail to be delivered.

By default, the trap function is enabled for the event that FIB entries fail to be delivered.

NOTE

Only the S5735S-H, S5736-S and S6720S-S support this command.

Format

fib trap lpm-fail enable

undo fib trap lpm-fail enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

FIB entries may fail to be delivered because of an LPM algorithm error. For fault location, you can run the **fib trap lpm-fail enable** command to enable the trap function for the event that FIB entries fail to be delivered and check trap function. After this trap function is disabled, no trap is generated when FIB entries fail to be delivered because of an LPM algorithm error.

Precautions

Traps are generated for the event that FIB entries fail to be delivered only when this trap function is enabled. To set the trap interval, run the **fib trap lpm-fail interval** command. To set the maximum number of traps generated in a trap interval, run the **fib trap lpm-fail history** command.

Example

Enable the trap function for the event that FIB entries fail to be delivered because of an LPM algorithm error.

```
<HUAWEI> system-view
[HUAWEI] fib trap lpm-fail enable
```


7.10.35 fib trap lpm-fail history

Function

The **fib trap lpm-fail history** command sets the maximum number of traps generated when FIB entries fail to be delivered.

The **undo fib trap lpm-fail history** command restores the default maximum number of traps generated when FIB entries fail to be delivered.

By default, a maximum of 10 traps are generated each time FIB entries fail to be delivered.

NOTE

Only the S5735S-H, S5736-S and S6720S-S support this command.

Format

fib trap lpm-fail history *history-number*

undo fib trap lpm-fail history

Parameters

Parameter	Description	Value
history-number	Sets the maximum number of traps generated each time FIB entries fail to be delivered.	The value is an integer that ranges from 10 to 30.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

FIB entries may fail to be delivered because of an LPM algorithm error. By default, a maximum of 10 traps are generated each time FIB entries fail to be delivered. To change the maximum number of traps generated each time FIB entries fail to be delivered, run the **fib trap lpm-fail history** *history-number* command.

Example

Set the maximum number of traps generated each time FIB entries fail to be delivered to 30.

```
<HUAWEI> system-view
```


[HUAWEI] **fib trap lpm-fail history 30**

7.10.36 fib trap lpm-fail interval

Function

The **fib trap lpm-fail interval** command sets the interval for generating traps when FIB entries fail to be delivered.

The **undo fib trap lpm-fail interval** command restores the default interval for generating traps when FIB entries fail to be delivered.

By default, traps are generated at an interval of 60s when FIB entries fail to be delivered.

NOTE

Only the S5735S-H, S5736-S and S6720S-S support this command.

Format

fib trap lpm-fail interval *interval-number*

undo fib trap lpm-fail interval

Parameters

Parameter	Description	Value
interval-number	Sets the interval for generating traps.	The value is an integer that ranges from 30 to 3600, in seconds. The default value is 60.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

FIB entries may fail to be delivered because of an LPM algorithm error. By default, traps are generated at an interval of 60s when FIB entries fail to be delivered because of an LPM algorithm error. To change the interval, run the **fib trap lpm-fail interval** *interval-number* command.

Example

Set the interval for generating traps when FIB entries fail to be delivered to 90s.


```
<HUAWEI> system-view
[HUAWEI] fib trap lpm-fail interval 90
```

7.10.37 ip frr (system view)

Function

The **ip frr** command enables IP FRR for public routes.

The **undo ip frr** command disables IP FRR for public routes.

By default, IP FRR is disabled for public routes.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ip frr route-policy *route-policy-name*

undo ip frr [**all**]

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Enables IP FRR for public routes matching the specified route-policy.	The value must be an existing route-policy.
all	Disables IP FRR for all public and private routes. If this parameter is not specified, IP FRR is disabled for all public routes.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

With the development of the network, services such as audio, online video, and finance have more requirements for real time. Generally, active/standby links are deployed on the network to ensure service stability.

However, under traditional forwarding modes, when multiple routes to the same destination exist, switch selects the optimal route, which is delivered to FIB table to direct data forwarding. When the optimal link is faulty, switch waits for the completion of route convergence, then selects another optimal route, and then deliver the route to the FIB table. Then the service is recovered. This process leads to a long-time service interruption and cannot meet service requirements.

Using the **ip frr** command, you can enable IP FRR of the public network. IP FRR can specify a backup next hop and a backup interface and set backup forwarding information for IPv4 routes. When the active link is faulty, the system can switch the traffic immediately to the backup link. This process is irrelevant to route convergence and therefore services are interrupted for short time.

Pre-configuration Tasks

The **ip frr** command should be used with the **apply backup-interface** command and the **apply backup-nexthop** command. It is required to use the **route-policy** command to create route-policy at first, in which the **apply backup-interface** command and the **apply backup-nexthop** command are used to set a backup outbound interface and a backup next hop for IPv4 route of the public network.

Precautions

Only one policy can be used at one time. New configuration will replace the previous one if another policy is configured. Configuration in the system view and that in the VPN instance IPv4 address family view will not interfere each other.

Using the **undo ip frr all** command, IP FRR of all the public network and the private networks is disabled. Use it with caution.

Example

Specify a backup outbound interface and a backup next hop in route-policy **ip_frr_rp** and enable IP FRR for public routes in the system view.

```
<HUAWEI> system-view
[HUAWEI] route-policy ip_frr_rp permit node 10
[HUAWEI-route-policy] apply backup-interface vlanif 100
[HUAWEI-route-policy] apply backup-nexthop 192.168.20.2
[HUAWEI-route-policy] quit
[HUAWEI] ip frr route-policy ip_frr_rp
```

7.10.38 ip prefix-limit

Function

The **ip prefix-limit** command configures a limit on the number of IPv4 public route prefixes.

The **undo ip prefix-limit** command restores the default configuration.

By default, the maximum number of IPv4 public route prefixes is not limited.

Format

ip prefix-limit *number* { *alert-percent* [**route-unchanged**] | **simply-alert** }

undo ip prefix-limit

Parameters

Parameter	Description	Value
number	Specifies the maximum number of IPv4 public route prefixes.	The value is an integer, and the minimum value is 1. The maximum number is determined by the license file.
alert-percent	Specifies the percentage of the maximum number of IPv4 public route prefixes. If you specify *alert-percent* in the command, when the number of IPv4 public route prefixes exceeds the value calculated by (*number* x *alert-percent*) / 100, an alarm is generated. Additional IPv4 public route prefixes can still be added to the routing table until the number of IPv4 public route prefixes reaches *number*. Subsequent route prefixes are discarded.	The value is an integer ranging from 1 to 100.
route-unchanged	Indicates that the routing table remains unchanged. If you decrease *alert-percent* after the number of IPv4 public route prefixes exceeds *number*, whether the routing table remains unchanged is determined by **route-unchanged**. • If you specify **route-unchanged** in the command, the routing table remains unchanged. • If you do not specify **route-unchanged** in the command, the system deletes the routes from the routing table and re-adds routes. By default, the system deletes the routes from the routing table and re-adds routes.	-

Parameter	Description	Value
simply-alert	Indicates the following function: If you specify *simply-alert* in the command, new IPv4 public route prefixes can still be added to the routing table and only an alarm is generated after the number of IPv4 public route prefixes exceeds *number*. However, when the total number of private and public route prefixes reaches the limit on the number of unicast route prefixes specified in the PAF file, subsequent IPv4 public route prefixes are discarded.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If the switch imports a large number of routes, system performance may be affected when processing services because the routes consume a lot of system resources. To improve system security and reliability, you can run the **ip prefix-limit** command to configure a limit on the number of IPv4 public route prefixes. When the number of IPv4 public route prefixes exceeds the limit, an alarm is generated, prompting you to check whether unneeded IPv4 public route prefixes exist.

Configuration Impact

After the **ip prefix-limit** command is run, the switch may discard unneeded IPv4 public route prefixes.

- If the number of IPv4 public route prefixes exceeds the value calculated from *number* x *alert-percent*/100, an alarm (RM_1.3.6.1.4.1.2011.5.25.145.19.1.3 hwPublicIpv4PrefixThresholdExceed) is generated.
- If the number of IPv4 public route prefixes exceeds *number*, an alarm (RM_1.3.6.1.4.1.2011.5.25.145.19.1.1 hwPublicIpv4PrefixExceed) is generated.
- If the number of IPv4 public route prefixes falls below the value calculated from (*number* x (*alert-percent* – 5))/100, a clear alarm (RM_1.3.6.1.4.1.2011.5.25.145.19.1.4 hwPublicIpv4PrefixThresholdExceedClear) is generated.
- If the number of IPv4 public route prefixes exceeds *number*, a clear alarm (RM_1.3.6.1.4.1.2011.5.25.145.19.1.2 hwPublicIpv4PrefixExceedClear) is generated.

Precautions

If you run the **ip prefix-limit** command for several times, the last configuration overrides previous configurations.

After the number of IPv4 public route prefixes exceeds the limit, note the following rules:

- If you run the **ip prefix-limit** command to increase *number* or the **undo ip prefix-limit** command to delete the limit, the switch relearns IPv4 public route prefixes.
- Direct and static routes can still be added to the IP routing table.

Example

Configure **simply-alert** so that only an alarm is generated when the switch imports more than 10000 IPv4 public route prefixes.

```
<HUAWEI> system-view
[HUAWEI] ip prefix-limit 10000 simply-alert
```

7.10.39 ip prefix-limit log-interval

Function

The **ip prefix-limit log-interval** command configures an interval at which logs are generated after the number of IPv4 public route prefixes exceeds the limit.

The **undo ip prefix-limit log-interval** command restores the default configuration.

By default, the system generates logs at an interval of 5s after the number of IPv4 public route prefixes exceeds the limit.

Format

ip prefix-limit log-interval *interval*

undo ip prefix-limit log-interval

Parameters

Parameter	Description	Value
interval	Specifies an interval at which logs are generated after the number of IPv4 public route prefixes exceeds the limit.	The value is an integer ranging from 1 to 60, in seconds.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The system generates logs at an interval of 5s after the number of IPv4 public route prefixes exceeds the limit. You can run the **ip prefix-limit log-interval** command to set a larger value for the interval to decrease the frequency at which these logs are generated.

The maximum number of IPv4 public route prefixes supported by the routing table can be adjusted using the **ip prefix-limit** command.

Precautions

If a log is generated to record the event that the number of IPv4 public route prefixes reaches the limit, no more routes can be added to the routing table, and subsequent routes are discarded.

Example

Set the interval at which logs are generated after the number of IPv4 public route prefixes exceeds the limit to 30s.

```
<HUAWEI> system-view
[HUAWEI] ip prefix-limit log-interval 30
```

7.10.40 ip route recursive-lookup default-route protocol

Function

The **ip route recursive-lookup default-route protocol** command sets specified routes to recurse to the default route.

The **undo ip route prefix-priority-scheduler protocol** command cancels specified routes from recursing to the default route.

By default, routes cannot recurse to the default route.

Format

ip route recursive-lookup default-route protocol { **static** | **bgp** | **msr** }

undo ip route recursive-lookup default-route protocol { **static** | **bgp** | **msr** }

NOTE

- Only the S5720I-SI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S support **bgp** parameter.
- Only the S5720-LI, S5720S-LI, S5720I-SI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S support **msr** parameter.

Parameters

Parameter	Description	Value
static	Static route	-
bgp	BGP Route	-
msr	Multicast Static Route	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Traffic cannot be transmitted along static, BGP, or multicast static routes with an unreachable next hop. In this situation, route recursion is required. Run the **ip route recursive-lookup default-route protocol** command to recurse these routes to the default route.

Precautions

After route recursion, the actual forwarding route may change.

Example

Set static routes to recurse to the default route.

```
<HUAWEI> system-view
[HUAWEI] ip route recursive-lookup default-route protocol static
```

7.10.41 ip route trace unicast-route-change disable

Function

The **ip route trace unicast-route-change disable** command disables recording of unicast route changes.

The **undo ip route trace unicast-route-change disable** command enables recording of unicast route changes.

By default, recording of unicast route changes is enabled.

Format

ip route trace unicast-route-change disable

undo ip route trace unicast-route-change disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

If recording of unicast route changes is enabled, information about changes in direct, static, OSPF, IS-IS, and BGP routes that are delivered to the routing table is recorded.

Example

Disable recording of unicast route changes.

```
<HUAWEI> system-view
[HUAWEI] ip route trace unicast-route-change disable
```

7.10.42 ipv6 fib-loadbalance-type hash-based

Function

The **ipv6 fib-loadbalance-type hash-based** command configures packets sent from a specified source to a specified destination to be forwarded using the same route.

The **undo ipv6 fib-loadbalance-type hash-based** command restores the default configuration.

By default, flow-based load balancing is used.

Format

ipv6 fib-loadbalance-type hash-based

undo ipv6 fib-loadbalance-type hash-based

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

This command configures packets sent from a specified source to a specified destination to be forwarded using the same route. This command takes effect only when the FIB cache fails.

Example

Configure packets sent from a specified source to a specified destination to be forwarded using the same route.

```
<HUAWEI> system-view
[HUAWEI] ipv6 fib-loadbalance-type hash-based
```

7.10.43 ipv6 fib regularly-refresh

Function

The **ipv6 fib regularly-refresh** command configures the entire IPv6 FIB entry update interval, IPv6 FIB entry update interval per cycle, and the number of IPv6 FIB entries updated per cycle.

The **undo ipv6 fib regularly-refresh** command restores the default entire IPv6 FIB entry update interval, IPv6 FIB entry update interval per cycle, and the number of IPv6 FIB entries updated per cycle.

By default, the IPv6 FIB entry update interval per cycle is 1 second, and the number of IPv6 FIB entries updated per cycle is 50. The entire IPv6 FIB entry update interval is 1 minute on the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H, and is fixed to 1440 minutes on other devices.

Format

ipv6 fib regularly-refresh { **interval** *interval* | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* }

undo ipv6 fib regularly-refresh { **interval** *interval* | **entry-number** *entry-number* | **cycle-interval** *cycle-interval* }

Parameters

Parameter	Description	Value
interval *interval*	Specifies an interval at which IPv6 FIB entries are refreshed.	The value is an integer ranging from 1 to 300, in seconds.

Parameter	Description	Value
entry-number *entry-number*	Specifies the number of IPv6 FIB entries refreshed per circle.	The value is an integer ranging from 20 to 2000.
cycle-interval *cycle-interval*	Specifies the entire interval at which IPv6 FIB entries are updated.	The value is an integer ranging from 1 to 1440, in minutes. **NOTE** This parameter is supported only by the S5731-S, S5731-H, S5731S-H, S5732-H, S6720-EI, S6720S-EI, S6730-S, S6730S-H, and S6730-H.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

Regular IPv6 FIB entry updating occupies system resources, including CPU and memory resources. A longer update interval and a smaller number of updated entries result in less impact on other services. You can flexibly configure the update interval and the number of entries to be updated based on system resources.

To allow underlying data to be forwarded in real-time and effective ways, you can configure a device to update IPv6 FIB entries in cycles. The *cycle-interval* parameter specifies the interval for IPv6 FIB entry update per cycle. Note that the *cycle-interval*, *interval*, and *entry-number* parameters together determine the IPv6 FIB entry update process. As an example, assume that a device has 500 IPv6 FIB entries.

- If *interval* is set to 1 second, *entry-number* is set to 50, and *cycle-interval* is set to 1 minute, the 500 FIB entries are updated in 10 cycles (with 50 FIB entries updated in each cycle), and the entry update can be completed within 10s (1s for each cycle). The device stays idle from 10s to 1 minute (entire FIB entry update interval).
- If *interval* is set to 10 seconds, *entry-number* is set to 50, and *cycle-interval* is set to 1 minute, the 500 FIB entries are still updated in 10 cycles, taking a total of 100 seconds. In this case, the update period exceeds 1 minute, and the device cycles update from the next minute (a total of 2 minutes at this point). That is, the device updates FIB entries at an interval of 2 minutes.

Prerequisites

Regular IPv6 FIB entry updating has been enabled using the **undo ipv6 fib regularly-refresh disable** command in the system view. This function is enabled by default.

Example

Set an interval at which IPv6 FIB entries are updated to 5 seconds and the number of FIB entries updated per circle to 200.

```
<HUAWEI> system-view
[HUAWEI] ipv6 fib regularly-refresh interval 5 entry-number 200
```

7.10.44 ipv6 fib regularly-refresh disable

Function

The **ipv6 fib regularly-refresh disable** command disables IPv6 FIB regularly-refresh.

The **undo ipv6 fib regularly-refresh disable** command enables IPv6 FIB regularly-refresh.

By default, IPv6 FIB regularly-refresh is enabled.

Format

ipv6 fib regularly-refresh disable

undo ipv6 fib regularly-refresh disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

During peak hours, refreshing the IPv6 FIB can cause over-high CPU usage and affect service reliability. To resolve this problem, run the **ipv6 fib regularly-refresh disable** command to disable IPv6 FIB regularly-refresh. After the CPU usage restores to a normal state, run the **undo ipv6 fib regularly-refresh disable** command to restore IPv6 FIB regularly-refresh.

Precautions

Disabling IPv6 FIB regularly-refresh may prevent the rectification of certain hardware forwarding entry faults. Exercise caution when running the **ipv6 fib regularly-refresh disable** command.

Example

Disable IPv6 FIB regularly-refresh.

```
<HUAWEI> system-view
[HUAWEI] ipv6 fib regularly-refresh disable
Warning: This operation will disable the function of the FIB regularly refresh,Confirm?[Y/N]:y
Info: The function of the FIB regularly refresh is disabled.
```

7.10.45 ipv6 prefix-limit

Function

The **ipv6 prefix-limit** command configures a limit on the number of IPv6 public route prefixes.

The **undo ipv6 prefix-limit** command restores the default configuration.

By default, the maximum number of IPv6 public route prefixes is not limited.

Format

ipv6 prefix-limit *number* { *alert-percent* [**route-unchanged**] | **simply-alert** }

undo ipv6 prefix-limit

Parameters

Parameter	Description	Value
number	Specifies the maximum number of IPv6 public route prefixes.	The value is an integer, and the minimum value is 1. The maximum number is determined by the license file.
alert-percent	Specifies the percentage of the maximum number of IPv6 public route prefixes. If you specify *alert-percent* in the command, when the number of IPv6 public route prefixes exceeds the value calculated by (*number* x *alert-percent*) / 100, an alarm is generated. Additional IPv6 public route prefixes can still be added to the routing table until the number of IPv6 public route prefixes reaches *number*. Subsequent route prefixes are discarded.	The value is an integer ranging from 1 to 100.

Parameter	Description	Value
route-unchanged	Indicates that the routing table remains unchanged. If you decrease *alert-percent* after the number of IPv6 public route prefixes exceeds *number*, whether the routing table remains unchanged is determined by **route-unchanged**. • If you specify **route-unchanged** in the command, the routing table remains unchanged. • If you do not specify **route-unchanged** in the command, the system deletes the routes from the routing table and re-adds routes. By default, the system deletes the routes from the routing table and re-adds routes.	-
simply-alert	Indicates the following function: If you specify *simply-alert* in the command, new IPv6 public route prefixes can still be added to the routing table and only an alarm is generated after the number of IPv6 public route prefixes exceeds *number*. However, when the total number of private and public route prefixes reaches the limit on the number of unicast route prefixes specified in the PAF file, subsequent IPv6 public route prefixes are discarded.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If the switch imports a large number of routes, system performance may be affected when processing services because the routes consume a lot of system resources. To improve system security and reliability, you can run the **ipv6 prefix-limit** command to configure a limit on the number of IPv6 public route prefixes. When the number of IPv6 public route prefixes exceeds the limit, an alarm is generated, prompting you to check whether unneeded IPv6 public route prefixes exist.

Configuration Impact

After the **ipv6 prefix-limit** command is run, the switch may discard unneeded IPv6 public route prefixes.

Precautions

If you run the **ipv6 prefix-limit** command for several times, the last configuration overrides previous configurations.

After the number of IPv6 public route prefixes exceeds the limit, note the following rules:

- If you run the **ipv6 prefix-limit** command to increase *number* or the **undo ipv6 prefix-limit** command to delete the limit, the switch relearns IPv6 public route prefixes.
- Direct and static routes can still be added to the IPv6 routing table.

Example

Configure **simply-alert** so that only an alarm is generated when the switch imports more than 10000 IPv6 public route prefixes.

```
<HUAWEI> system-view
[HUAWEI] ipv6 prefix-limit 10000 simply-alert
```

7.10.46 ipv6 prefix-limit log-interval

Function

The **ipv6 prefix-limit log-interval** command configures an interval at which logs are generated after the number of IPv6 public route prefixes exceeds the limit.

The **undo ipv6 prefix-limit log-interval** command restores the default configuration.

By default, the system generates logs at an interval of 5s after the number of IPv6 public route prefixes exceeds the limit.

Format

ipv6 prefix-limit log-interval *interval*

undo ipv6 prefix-limit log-interval

Parameters

Parameter	Description	Value
interval	Specifies an interval at which logs are generated after the number of IPv6 public route prefixes exceeds the limit.	The value is an integer ranging from 1 to 60, in seconds.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The system generates logs at an interval of 5s after the number of IPv6 public route prefixes exceeds the limit. You can run the **ip prefix-limit log-interval** command to set a larger value for the interval to decrease the frequency at which these logs are generated.

The maximum number of IPv6 public route prefixes supported by the routing table can be adjusted using the **ipv6 prefix-limit** command.

Precautions

If a log is generated to record the event that the number of IPv6 public route prefixes reaches the limit, no more routes can be added to the routing table, and subsequent routes are discarded.

Example

Set the interval at which logs are generated after the number of IPv6 public route prefixes exceeds the limit to 30s.

```
<HUAWEI> system-view
[HUAWEI] ipv6 prefix-limit log-interval 30
```

7.10.47 ipv6 frr (system view)

Function

The **ipv6 frr** command enables IPv6 FRR for public routes.

The **undo ipv6 frr** command disables IPv6 FRR for public routes.

By default, IPv6 FRR is disabled for public routes.

Product	Support
S5720I-SI, S5735-S, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6720S-S, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S500, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, and S5720S-LI	Not supported

Format

ipv6 frr route-policy *route-policy-name*

undo ipv6 frr [**all**]

Parameters

Parameter	Description	Value
route-policy *route-policy-name*	Enables IPv6 FRR for public routes matching the specified route-policy.	The value must be an existing route-policy.
all	Disables IPv6 FRR for all public and private routes. If this parameter is not specified, IPv6 FRR is disabled for all public routes.	-

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

With the development of the network, services such as audio, online video, and finance have more requirements for real time. Generally, active/standby links are deployed on the network to ensure service stability.

However, under traditional forwarding modes, when multiple routes to the same destination exist, switch selects the optimal route, which is delivered to FIB table to direct data forwarding. When the optimal link is faulty, switch waits for the completion of route convergence, then selects another optimal route, and then deliver the route to the FIB table. Then the service is recovered. This process leads to a long-time service interruption and cannot meet service requirements.

Using the **ipv6 frr** command, you can enable IPv6 FRR of the public network. IPv6 FRR can specify a backup next hop and a backup interface and set backup forwarding information for IPv6 routes. When the active link is faulty, the system can switch the traffic immediately to the backup link. This process is irrelevant to route convergence and therefore services are interrupted for short time.

Pre-configuration Tasks

The **ipv6 frr** command should be used with the **apply ipv6 backup-interface** command and the **apply ipv6 backup-nexthop** command. It is required to use the **route-policy** command to create route-policy at first, in which the **apply ipv6 backup-interface** command and the **apply ipv6 backup-nexthop** command are used to set a backup outbound interface and a backup next hop for IPv6 route of the public network.

Precautions

Only one policy can be used at one time. New configuration will replace the previous one if another policy is configured. Configuration in the system view and that in the VPN instance IPv6 address family view will not interfere each other.

Using the **undo ipv6 frr all** command, IPv6 FRR of all the public network and the private networks is disabled. Use it with caution.

Example

Specify a backup outbound interface and a backup next hop in route-policy **ipv6_frr_rp** and enable IPv6 FRR for public routes in the system view.

```
<HUAWEI> system-view
[HUAWEI] route-policy ipv6_frr_rp permit node 10
[HUAWEI-route-policy] apply ipv6 backup-interface vlanif 100
[HUAWEI-route-policy] apply ipv6 backup-nexthop 2001:db8:1::1
[HUAWEI-route-policy] quit
[HUAWEI] ipv6 frr route-policy ipv6_frr_rp
```

7.10.48 l3-forward-entry aging

Function

The **l3-forward-entry aging enable** command enables the l3-forward-entry aging function.

The **undo l3-forward-entry aging enable** command disables the l3-forward-entry aging function.

By default, the l3-forward-entry aging function is disabled.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

l3-forward-entry aging enable

undo l3-forward-entry aging enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

The **l3-forward-entry aging enable** command must be used together with the **l3-forward-entry regularly-check** command. After the l3-forward-entry aging function is enabled, the system reads the periodic l3-forward-entry check result everyday between 2 am and 5 am. When residual Layer 3 forwarding entries exist for more than 48 hours, the system triggers resource recycling and entry deletion.

Example

Enable the l3-forward-entry aging function.

```
<HUAWEI> system-view
[HUAWEI] l3-forward-entry aging enable
```

Disable the l3-forward-entry aging function.

```
<HUAWEI> system-view
[HUAWEI] undo l3-forward-entry aging enable
```

7.10.49 l3-forward-entry regularly-check

Function

The **l3-forward-entry regularly-check** command enables periodic l3-forward-entry check and sets the periodic l3-forward-entry check interval, scanning time for each round of check, and number of entries checked in each round of check.

The **undo l3-forward-entry regularly-check** command restores the default settings of periodic l3-forward-entry check.

By default, periodic l3-forward-entry check is enabled, the periodic l3-forward-entry check interval is 1 minute, the scanning time for each round of check is 1 second, and the number of entries checked in each round of check is 200.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported

Product	Support
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

l3-forward-entry regularly-check { **enable** | **cycle-interval** *cycle-interval-value* | **interval** *interval-value* | **entry-number** *number* }

undo l3-forward-entry regularly-check { **enable** | **cycle-interval** | **interval** | **entry-number** }

Parameters

Parameter	Description	Value
enable	Enables periodic l3-forward-entry check	-
cycle-interval *cycle-interval-value*	Sets the periodic l3-forward-entry check interval.	The value is an integer that ranges from 1 to 1440, in minutes.
interval *interval-value*	Sets the scanning time for each round of l3-forward-entry check.	The value is an integer that ranges from 1 to 300, in seconds.
entry-number *number*	Sets the number of entries checked in each round of check.	The value is an integer that ranges from 20 to 2000.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After periodic l3-forward-entry check is configured, the device periodically checks l3-forward-entries to ensure realtime and valid bottom-layer forwarded data and

records the check result. The check result can help locate problems of Layer 3 forwarding software resource leak and residual Layer 3 forwarding hardware entries.

The three parameters, *cycle-interval*, *interval*, and *entry-number*, together determine the periodic l3-forward-entry check process. For example, the device has a total of 2,000 l3-forward-entries:

- If the three parameters use their default values, that is, 2,000 entries need to be checked in 10 rounds of check and can be checked within 10 seconds (with the scanning time of each round of check as 1 second), the system is idle during the period between 10 seconds to 1 minutes (periodic check interval).
- If *cycle-interval* is set to 1 minute, *interval* is set to 10 seconds, and *entry-number* uses the default value, 2,000 entries still need to be updated for 10 rounds of check, and a total of 100 seconds is required to finish updating these entries. Then the update time will exceed 1 minute, so the device needs to start periodic update within the next 1 minute. In this situation, periodic update is performed at an interval of 2 minutes.

Precautions

Periodic l3-forward-entry check occupies system resources, including CPU and memory resources. A longer check interval and fewer entries indicate a smaller impact on other services. You can flexibly configure the l3-forward-entry check interval and number of checked entries based on system resources.

Example

Enable periodic l3-forward-entry check.

```
<HUAWEI> system-view
[HUAWEI] l3-forward-entry regularly-check enable
```

Set the periodic l3-forward-entry check interval to 10 minutes.

```
<HUAWEI> system-view
[HUAWEI] l3-forward-entry regularly-check cycle-interval 10
```

Set the scanning time for each round of l3-forward-entry check to 5 seconds.

```
<HUAWEI> system-view
[HUAWEI] l3-forward-entry regularly-check interval 5
```

Set the number of entries checked in each round of check to 500.

```
<HUAWEI> system-view
[HUAWEI] l3-forward-entry regularly-check entry-number 500
```

Disable periodic l3-forward-entry check.

```
<HUAWEI> system-view
[HUAWEI] undo l3-forward-entry regularly-check enable
```

7.10.50 l3-forward-entry self-healing enable

Function

The **l3-forward-entry self-healing enable** command enables l3-forward-entry self-healing.

The **undo l3-forward-entry self-healing enable** command disables l3-forward-entry self-healing.

By default, the l3-forward-entry self-healing function is disabled.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

l3-forward-entry self-healing enable

undo l3-forward-entry self-healing enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

After the **l3-forward-entry regularly-check** command is configured on a switch, the switch periodically checks l3-forward-entries to ensure real-time and effective forwarding of data at the bottom layer. In this situation, if the **l3-forward-entry self-healing enable** command is configured to enable l3-forward-entry self-healing, the switch can repair l3-forward-entries for self-healing based on the check result to ensure correct packet forwarding.

Example

Enable l3-forward-entry self-healing.

```
<HUAWEI> system-view
[HUAWEI] l3-forward-entry self-healing enable
```


7.10.51 l3-interface regularly-refresh

Function

The **l3-interface regularly-refresh** command enables periodic l3-interface entry update and configures the periodic l3-interface entry update interval, scanning time for each round of update, and number of entries updated in each round of update.

The **undo l3-interface regularly-refresh** command restores the default configuration.

By default, periodic l3-interface entry update is enabled, the periodic l3-interface entry update interval is 1 minute, the scanning time for each round of update is 1 second, and the number of entries updated in each round of update is 50.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

l3-interface regularly-refresh { **enable** | **cycle-interval** *cycle-interval-value* | **scan-interval** *scan-interval-value* | **entry-number** *entry-number* }

undo l3-interface regularly-refresh { **enable** | **cycle-interval** | **scan-interval** | **entry-number** }

Parameters

Parameter	Description	Value
enable	Enables periodic l3-interface entry update.	-
cycle-interval *cycle-interval-value*	Sets the periodic l3-interface entry update interval.	The value is an integer that ranges from 1 to 1440, in minutes.

Parameter	Description	Value
scan-interval *scan-interval-value*	Sets the scanning time for each round of l3-interface entry update.	The value is an integer that ranges from 1 to 300, in seconds.
entry-number *entry-number*	Sets the number of entries updated in each round of l3-interface entry update.	The value is an integer that ranges from 1 to 100.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the periodic l3-interface entry update function is configured on a switch, the switch periodically updates l3-interface entries to ensure real-time and effective forwarding of data at the bottom layer and records revision information.

The three parameters, *cycle-interval*, *scan-interval*, and *entry-number*, together determine the periodic l3-interface entry update process. For example, the switch has a total of 500 l3-interface entries:

- If the three parameters retain their default values, that is, 500 entries need to be updated in 10 rounds and can be checked within 10 second (with the scanning time of each round of update as 1 second), the system is idle during the period between 10 seconds to 1 minute (periodic update interval).
- If *cycle-interval* is set to 1 minute, *scan-interval* is set to 10 seconds, and *entry-number* retains the default value, 500 entries still need to be updated for 10 rounds, and a total of 100 seconds is required to finish updating these entries. Then the update time will exceed 1 minute, so the switch needs to start periodic update within the next 1 minute. In this situation, periodic update is performed at an interval of 2 minutes.

Precautions

Periodic l3-interface entry update occupies system resources, including CPU and memory resources. A longer check interval and fewer entries indicate a smaller impact on other services. Flexibly configure the l3-interface entry update interval and number of l3-interface entries updated each time based on system resources.

Example

Enable periodic l3-interface entry update.

```
<HUAWEI> system-view
[HUAWEI] l3-interface regularly-refresh enable
```


Set the periodic l3-interface entry update interval to 10 minutes.

```
<HUAWEI> system-view
[HUAWEI] l3-interface regularly-refresh cycle-interval 10
```

Set the scanning time for each round of l3-interface entry update to 5 seconds.

```
<HUAWEI> system-view
[HUAWEI] l3-interface regularly-refresh scan-interval 5
```

Set the number of entries checked in each round of update to 100.

```
<HUAWEI> system-view
[HUAWEI] l3-interface regularly-refresh entry-number 100
```

Disable periodic l3-interface entry update.

```
<HUAWEI> system-view
[HUAWEI] undo l3-interface regularly-refresh enable
```

7.10.52 l3-interface periodic-refresh disable

Function

The **l3-interface periodic-refresh disable** command disables the function of updating l3-interface entries every 10 hours.

The **undo l3-interface periodic-refresh disable** command enables the function of updating l3-interface entries every 10 hours.

By default, l3-interface entries are updated every 10 hours.

Format

l3-interface periodic-refresh disable

undo l3-interface periodic-refresh disable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

By default, a device updates l3-interface entries every 10 hours. To reduce the load on the device, you can run this command to disable the device from updating l3-interface entries every 10 hours.

Precautions

To disable all functions of periodically updating l3-interface entries on the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S, you need to run both the **l3-interface periodic-refresh disable** and **undo l3-interface regularly-refresh** commands.

Example

Disable the function of updating l3-interface entries every 10 hours.

```
<HUAWEI> system-view
[HUAWEI] l3-interface periodic-refresh disable
```

7.10.53 refresh fib

Function

The **refresh fib** command delivers IPv4 FIB policies and forwarding entries.

Format

refresh fib slot *slot-id*

Parameters

Parameter	Description	Value
slot *slot-id*	Specifies the slot ID of a switch that needs to update IPv4 FIB policies.	The value is determined based on the device configuration.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

After the IPv4 FIB policy changes, you must run the **refresh fib** command to update the switch to make the configuration take effect.

After the **refresh fib** command is executed, you can view the FIB table to check whether the update succeeds.

Precautions

If the FIB module is in overload state and the switch works in overload suspension mode, the IPv4 and IPv6 modules need to be updated.

Example

Refresh IPv4 FIB module on the switch in slot 2.

```
<HUAWEI> system-view
[HUAWEI] refresh fib slot 2
```

7.10.54 refresh ipv6 fib

Function

The **refresh ipv6 fib** command delivers IPv6 FIB policies and forwarding entries.

Format

refresh ipv6 fib slot *slot-id*

Parameters

Parameter	Description	Value
slot *slot-id*	Specifies the slot ID of a switch that needs to update IPv6 FIB policies.	The value is determined based on the device configuration.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

After the IPv6 FIB policy changes, you must run the **refresh ipv6 fib** command to update the switch to make the configuration take effect.

After the **refresh ipv6 fib** command is executed, you can view the FIB table to check whether the update succeeds.

NOTE

If the FIB module is in overload state and the switch works in overload suspension mode, the IPv4 and IPv6 modules need to be updated.

Example

Update the IPv6 FIB module on the switch in slot 2.

```
<HUAWEI> system-view
[HUAWEI] refresh ipv6 fib slot 2
```


7.10.55 reset ip routing-table statistics protocol

Function

The **reset ip routing-table statistics protocol** command clears route statistics in an IPv4 routing table.

Format

reset ip routing-table statistics protocol [**vpn-instance** *vpn-instance-name*] { **all** | *protocol* }

reset ip routing-table all-vpn-instance statistics protocol { **all** | *protocol* }

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Clears route statistics in an IPv4 routing table of the specified VPN instance. If this parameter is not specified, route statistics in an IPv4 routing table of the public network are cleared.	The value must be an existing VPN instance name.
all	Clear route statistics of all routing protocols in an IPv4 routing table.	-
protocol	Clears route statistics of the specified routing protocol in an IPv4 routing table. This parameter can be the following keywords. For details, see the routing protocols supported by the device: **direct**, **bgp**, **isis**, **ospf**, **rip**, **static**, and **unr**.	-
all-vpn-instance	Clears the statistics in the IPv4 routing tables of all VPN instances.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can use this command to clear route statistics in an IPv4 routing table, including statistics about the routes added, deleted, and released by each routing

protocol. Subsequently, the system can re-collect route statistics of each routing protocol again to monitor route changes and locate network faults.

Precautions

IPv4 route statistics cannot be restored after being cleared. Exercise caution when you use this command.

Example

Clear route statistics of all routing protocols in an IPv4 routing table.

```
<HUAWEI> reset ip routing-table statistics protocol all
```

7.10.56 reset ipv6 routing-table statistics protocol

Function

The **reset ipv6 routing-table statistics protocol** command clears route statistics in an IPv6 routing table.

Format

reset ipv6 routing-table [**vpn-instance** *vpn-instance-name*] **statistics protocol** { **all** | *protocol* }

reset ipv6 routing-table all-vpn-instance statistics protocol { **all** | *protocol* }

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of an IPv6 VPN instance.	The value must be an existing VPN instance name.
all	Clear route statistics of all routing protocols in an IPv6 routing table.	-
protocol	Clears route statistics of the specified routing protocol in an IPv6 routing table.	The value may be **bgp**, **direct**, **isis, ospfv3, ripng**, **static**, or **unr**. The specific value varies depending on the routing protocol supported by the device.
all-vpn-instance	Clears the statistics in the IPv6 routing tables of all VPN instances.	-

Views

User view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can use this command to clear route statistics of each routing protocol in an IPv6 routing table. Subsequently, the switch can re-collect route statistics of each routing protocol to monitor route changes and locate network faults.

Precautions

IPv6 route statistics cannot be restored after being cleared. Exercise caution when you use this command.

Example

Clear route statistics of all routing protocols in an IPv6 routing table.

```
<HUAWEI> reset ipv6 routing-table statistics protocol all
```

7.10.57 router id

Function

The **router id** command sets the global router ID.

The **undo router id** command deletes the configured global router ID.

By default, the global router ID is 0.0.0.0 when no IPv4 interface address is configured.

Format

router id *router-id*

undo router id

Parameters

Parameter	Description	Value
router-id	Sets a router ID in the IPv4 address format.	The value is in dotted decimal notation.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A router ID is a 32-bit integer used to uniquely identify a device that uses a specific dynamic routing protocol for route exchange. A router ID uses the same format as an IP address. Router IDs are classified into global router IDs of devices and router IDs of dynamic routing protocols.

The rules for selecting a global router ID are as follows:

1. If you set the global router ID using the **router id** command, the configured router ID is used.
2. If no global router ID is set, the router selects a router ID based on IP addresses of current interfaces:
   - If there are loopback interfaces that have IP addresses configured, the device selects the largest IP address among loopback interface addresses as the global router ID.
   - If no loopback interface is configured or loopback interfaces do not have IP addresses configured, the device selects the largest IP address among interface addresses as the global router ID without considering the Up/Down state of interfaces.

     NOTE

     The global router ID is reselected only when the interface address that is selected as the global router ID is deleted or changed. The global router ID is not reselected in any of the following situations:

     - The interface is Down.
     - A loopback interface is configured when the IP address of a non-loopback interface is selected as the router ID.
     - A larger interface address is configured later.

     You can run the **display router id** command to view the global router ID.
3. Each VPN instance selects the router ID from the IP addresses of the interfaces in the VPN instance based on rules 1 and 2.

After the global router ID is changed, manually run the **reset** command for each routing protocol to make the new global router ID take effect.

Precautions

When a device is being initialized and there is no interface address on the device, the first configured IP address on the device, which may be a loopback interface address or another interface address, will become the global router ID of the device. This global router ID remains unchanged even if there is a larger interface IP address on the device.

The IP address of the interface whose configuration is restored the first is used as the global router ID after the restart.

NOTE

To enhance network reliability, configure the address of a loopback interface as the global router ID.

Example

Set the global router ID to 10.10.10.1.

```
<HUAWEI> system-view
[HUAWEI] router id 10.10.10.1
```

7.10.58 route low-priority enable

Function

The **route low-priority enable** command sets the priority of routes on the forwarding plane to be lower than the ARP/ND priority.

The **undo route low-priority enable** command sets the priority of routes on the forwarding plane to be higher than the ARP/ND priority.

By default, on a device running a version earlier than V200R019C00SPC210, the priority of routes (excluding BGP routes and blackhole routes) on the forwarding plane is lower than the ARP/ND priority. On a device running V200R019C00SPC210 or later, the priority of routes on the forwarding plane is higher than the ARP/ND priority.

Format

route low-priority enable

undo route low-priority enable

Parameters

None

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

The control plane and forwarding plane of the routing system are separated and cooperate with each other to forward data. The control plane learns and exchanges routing protocols, calculates routes, generates routing entries, and delivers the entries to the forwarding plane for packet forwarding. The forwarding plane receives and sends packets, encapsulates and decapsulates packets, queries routing entries, and forwards packets.

If the control plane has host routes and ARP entries with the same destination IP address as the host routes, the forwarding plane forwards packets based on the routes because the route priority on the control plane is higher than the ARP/ND priority. On the forwarding plane, packets are forwarded based on routes or ARP/ND entries, depending on which ones have higher priority. For example, if the route priority on the forwarding plane is lower than the ARP/ND priority, packets are forwarded based on ARP/ND entries. In this case, if the outbound interface of an ARP/ND entry is different from that of a route, for example, the outbound interface is manually specified in a static ARP entry, different forwarding behaviors will be performed.

Therefore, you need to run the **route low-priority enable** command to adjust the route priority on the forwarding plane to control whether packets are forwarded based on routes or ARP/ND entries.

Precautions

- In actual deployment, if there is no special requirement, to ensure consistent forwarding behaviors on the forwarding plane and control plane, it is recommended that the route priority on the forwarding plane be higher than the ARP/ND priority.
- This command takes effect only for newly generated routes. For existing routes, you need to run the **refresh fib** and **refresh ipv6 fib** commands to manually refresh routes after running this command. During route refresh, path switching may cause temporary traffic interruptions.
- In an upgrade scenario, after a device running V200R019C00SPC200 or an earlier version is upgraded to V200R019C00SPC210 or a later version, the route priority on the forwarding plane remains unchanged. That is, the route priority on the forwarding plane is lower than the ARP/ND priority. This conflicts with the default setting on switches running V200R019C00SPC210 or a later version. This is because the route priority on the forwarding plane is higher than the ARP/ND priority on these switches by default. To resolve the conflict, the **route low-priority enable** command is automatically generated in the configuration file of a switch.

Example

Set the route priority on the forwarding plane to be higher than the ARP/ND priority.

```
<HUAWEI> system-view
[HUAWEI] undo route low-priority enable
```

7.11 PBR Configuration Commands

7.11.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.11.2 redirect ip-multihop

Function

The **redirect ip-multihop** command configures an action of redirecting packets to multiple next hop IP addresses in a traffic behavior.

The **undo redirect** command deletes the redirection configuration.

By default, an action of redirecting packets to multiple next hop IP addresses is not configured in a traffic behavior.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

redirect [**vpn-instance** *vpn-instance-name*] **ip-multihop** { **nexthop** *ip-address* } &<2-4>

redirect [**vpn-instance** *vpn-instance-name*] **ip-multihop** *acl-ip-pool-name*

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
nexthop *ip-address*	Specifies the IP address of the next hop.	The value is in dotted decimal notation and in X.X.X.X format.
acl-ip-pool-name	Specifies the name of an ACL IP address pool.	The ACL IP address pool name must exist.

Views

Traffic behavior view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If multiple next hop IP addresses are specified, the device redirects packets through equal-cost routes in load balancing mode.

If the outbound interface corresponding to a next hop IP address becomes Down or a route changes, the device switches traffic to the outbound interface corresponding to an available next hop. If the specified next hops are unavailable, the device forwards the packets to the original destination.

Follow-up Procedure

Run the **traffic policy** command to create a traffic policy and run the **classifier behavior** command in the traffic policy view to bind the traffic classifier to the traffic behavior containing redirection to multiple next hop IP addresses.

Precautions

- In a traffic behavior, the **redirect ip-multihop** command cannot be used together with any of the following commands: **remark 8021p**, **add-tag vlan-id**, **remark cvlan-id** and **remark vlan-id**.
- A traffic policy containing the redirection action can be only used globally, on an interface, or in a VLAN in the inbound direction.
- If you configure only two identical IP addresses when executing the **redirect ip-multihop** command, the system displays the following error message:
  ```
  Error: Redirect ip-multihop should configure at least two address.
  ```
- If the device has no ARP entry matching the specified next hop IP address, the **redirect ip-multihop** command can be used but redirection does not take effect. The device still forwards packets to the original destination until the device has the corresponding ARP entry.

Example

Configure three next hop IP addresses in the traffic behavior **b1**: 10.1.42.1, 10.2.12.3, and 10.1.1.2.

```
<HUAWEI> system-view
[HUAWEI] traffic behavior b1
[HUAWEI-behavior-b1] redirect ip-multihop nexthop 10.1.42.1 nexthop 10.2.12.3 nexthop 10.1.1.2
```

In the traffic behavior **b2**, configure the action of redirecting packets to five next hop IP addresses in the ACL IP address pool **abc**.

```
<HUAWEI> system-view
[HUAWEI] acl ip-pool abc
[HUAWEI-acl-ip-pool-abc] ip address 192.168.10.1 32
[HUAWEI-acl-ip-pool-abc] ip address 192.168.20.1 32
[HUAWEI-acl-ip-pool-abc] ip address 192.168.30.1 32
[HUAWEI-acl-ip-pool-abc] ip address 192.168.40.1 32
[HUAWEI-acl-ip-pool-abc] ip address 192.168.50.1 32
[HUAWEI-acl-ip-pool-abc] quit
[HUAWEI] traffic behavior b2
[HUAWEI-behavior-b2] redirect ip-multihop abc
```


7.11.3 redirect ip-nexthop

Function

The **redirect ip-nexthop** command configures the action of redirecting packets to a next-hop IP address in a traffic behavior.

The **undo redirect** command deletes the redirection configuration.

By default, the action of redirecting packets to a next-hop IP address is not configured in a traffic behavior.

Format

redirect [**remote**] [**vpn-instance** *vpn-instance-name*] **ip-nexthop** { *ip-address* [**track-nqa** *admin-name test-name*] } &<1-4> [**forced** | **low-precedence**] *

undo redirect

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support **vpn-instance** *vpn-instance-name* and **low-precedence**.

Parameters

Parameter	Description	Value
remote	Redirects packets to a remote next hop.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ip-address	Specifies a next-hop IP address.	The value is in dotted decimal notation and in X.X.X.X format.
track-nqa	Specifies an NQA test instance.	-
admin-name	Specifies the administrator of an NQA test instance.	The value is a string of 1 to 32 case-sensitive characters, excluding question marks (?), hyphens (-), and a single or consecutive quotation marks (").

Parameter	Description	Value
test-name	Specifies the name of an NQA test instance.	The value is a string of 1 to 32 case-sensitive characters, excluding question marks (?), hyphens (-), and a single or consecutive quotation marks (").
forced	Specifies that packets are discarded if the next hop is unavailable.	-
low-precedence	Specifies a low-priority next hop. **NOTE** If this parameter is specified, redirection-based PBR has a lower priority than routes generated through dynamic routing protocols or statically configured routes. If this parameter is not specified, the former has a higher priority than the latter. In enhanced-ipv4 or ipv4-ipv6 6:1 resource assignment mode, the S6720-EI does not support redirection to low-priority next hops. You can run the **display resource-mode configuration** command to view the resource assignment mode on the device, and run the **assign resource-mode** { **enhanced-mac** | **enhanced-arp** | **enhanced-ipv4** | **ipv4-ipv6 6:1** } [**slot** *slot-id* | **all**] command to change the resource assignment mode.	-

Views

Traffic behavior view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If multiple next-hop IP addresses are configured, the device redirects packets in active/standby mode. A maximum of four next-hop IP addresses can be configured in a traffic behavior. The device determines the primary and backup paths according to the sequence in which next-hop IP addresses are configured. The next-hop IP address that is configured first has the highest priority and this next hop is used as the primary path. Other next hops are used as backup paths. When the primary link becomes Down, a next hop with a higher priority is used as the primary link.

Follow-up Procedure

Run the **traffic policy** command to create a traffic policy and run the **classifier behavior** command in the traffic policy view to bind the traffic classifier to the traffic behavior containing redirection to a next-hop IP address.

Precautions

- In a traffic behavior, the **redirect ip-nexthop** command cannot be used together with any of the following commands: **remark 8021p**, **add-tag vlan-id**, **remark cvlan-id** and **remark vlan-id**.
- A traffic policy containing the redirection action can be applied globally, as well as on an interface and in a VLAN in the inbound direction.
- If no ARP entry matches the next-hop address on the device, the device triggers ARP learning. If the ARP entry cannot be learned, redirection does not take effect and packets are forwarded along the original path.
- In V200R021C00SPC600 and later versions, when the action of redirecting packets to a remote next hop, the configuration does not take effect if the next-hop address is reachable only through tunnel forwarding.

Example

Configure the action of redirecting packets to next hop 10.0.0.1 in the traffic behavior **b1**.

```
<HUAWEI> system-view
[HUAWEI] traffic behavior b1
[HUAWEI-behavior-b1] redirect ip-nexthop 10.0.0.1
```

7.11.4 redirect ipv6-multihop

Function

The **redirect ipv6-multihop** command configures an action of redirecting packets to multiple next hop IPv6 addresses in a traffic behavior.

The **undo redirect** command deletes the redirection configuration.

By default, an action of redirecting packets to multiple next hop IPv6 addresses is not configured in a traffic behavior.

Product	Support
S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S	Supported
S1720GW-E, S1720GWR-E, S5720I-SI, S5720-LI, S2730S-S, S5735-L-I, S5735-L1,S300, S5735-L, S5735S-L, S5735S-L1, S5735S-L-M, S5720S-LI, S5735-S, S500, S5735S-S, S5735-S-I, S5735S-H, S5736-S, S6720S-S	Not supported

Format

redirect [**vpn-instance** *vpn-instance-name*] **ipv6-multihop** { *ipv6-address* | **link-local** *link-local-address* **interface** *interface-type interface-number* } & <2-4>

redirect [**vpn-instance** *vpn-instance-name*] **ipv6-multihop** *acl6-ip-pool-name*

undo redirect

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ipv6-address	Specifies a next hop IPv6 address.	The address is a 32-digit hexadecimal number, in X:X:X:X:X:X:X:X format.
link-local *link-local-address*	Specifies a link-local address. The prefix of the specified IPv6 address must match FE80::/10.	The address is a 32-digit hexadecimal number, in X:X:X:X:X:X:X:X format.
interface *interface-type interface-number*	Specifies the interface corresponding to the link-local address. • *interface-type* specifies the interface type. • *interface-number* specifies the interface number.	-

Parameter	Description	Value
acl6-ip-pool-name	Specifies the name of the ACL IPv6 address pool to be created.	The value is a string of 1 to 32 case-sensitive characters without spaces and starting with a letter.

Views

Traffic behavior view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

If multiple next hop IPv6 addresses are specified, the device redirects packets through equal-cost routes in load balancing mode.

If the outbound interface corresponding to a next hop IPv6 address becomes Down or a route changes, the device switches traffic to the outbound interface corresponding to an available next hop. If the specified next hops are unavailable, the device forwards the packets to the original destination.

Follow-up Procedure

Run the **traffic policy** command to create a traffic policy and run the **classifier behavior** command in the traffic policy view to bind the traffic classifier to the traffic behavior containing redirection to multihop IPv6 addresses.

Precautions

- In a traffic behavior, the **redirect ipv6-multithop** command cannot be used together with any of the following commands: **remark 8021p**, **add-tag vlan-id**, **remark cvlan-id** and **remark vlan-id**.
- When the next hop IPv6 address is a local-link address, you must configure a VLANIF interface. The address is only valid in the VLAN corresponding to the VLANIF interface.
- The traffic policy that contains the redirection action can only be applied to the system, an interface, or a VLAN.
- When running the **redirect ipv6-multithop** command to configure multiple next hop IPv6 addresses, you can specify 2 to 4 next hop IPv6 addresses. When configuring multiple next hop IPv6 addresses through an ACL IPv6-POOL, you can specify a maximum of 64 next hop IPv6 addresses. To create and configure an ACL IPv6-POOL, run the **acl ipv6 ip-pool** and **ipv6 address (ACL IPv6-POOL view)** commands.
- If the device does not match the neighbor entry corresponding to the next hop IPv6 address, the device sends NS packets to check whether the neighbor

is reachable. If the neighbor is unreachable, packets are forwarded based on the original path and redirection does not take effect.

- If you run the **redirect ipv6-multithop** command in the same traffic classifier view multiple times, only the latest configuration takes effect.

Example

Configure two next hop IPv6 addresses in the traffic behavior **b1**: FC00:0:0:2000::1 and FC00:0:0:3000::1.

```
<HUAWEI> system-view
[HUAWEI] traffic behavior b1
[HUAWEI-behavior-b1] redirect ipv6-multihop fc00:0:0:2000::1 fc00:0:0:3000::1
```

In the traffic behavior **b2**, configure an action of redirecting packets to five next hop IPv6 addresses in the ACL IPv6-POOL **abc**.

```
<HUAWEI> system-view
[HUAWEI] acl ipv6 ip-pool abc
[HUAWEI-acl6-ip-pool-abc] ipv6 address 2001:db8::1 128
[HUAWEI-acl6-ip-pool-abc] ipv6 address 2001:db8::2 128
[HUAWEI-acl6-ip-pool-abc] ipv6 address 2001:db8::3 128
[HUAWEI-acl6-ip-pool-abc] ipv6 address 2001:db8::4 128
[HUAWEI-acl6-ip-pool-abc] ipv6 address 2001:db8::5 128
[HUAWEI-acl6-ip-pool-abc] quit
[HUAWEI] traffic behavior b2
[HUAWEI-behavior-b2] redirect ipv6-multihop abc
```

7.11.5 redirect ipv6-nexthop

Function

The **redirect ipv6-nexthop** command configures the action of redirecting packets to a next-hop IPv6 address in a traffic behavior.

The **undo redirect** command deletes the redirection configuration.

By default, the action of redirecting packets to a next-hop IPv6 address is not configured in a traffic behavior.

Format

redirect [**remote**] [**vpn-instance** *vpn-instance-name*] **ipv6-nexthop** { *ipv6-address* | **link-local** *link-local-address* **interface** *interface-type interface-number* } & <1-4> [**forced**]

undo redirect

NOTE

Only the S5731-H, S5731-S, S5731S-H, S5731S-S, S5732-H, S6735-S, S6720-EI, S6720S-EI, S6730-H, S6730S-H, S6730-S, and S6730S-S support the **vpn-instance** *vpn-instance-name* parameter.

Parameters

Parameter	Description	Value
remote	Redirects packets to a remote next hop.	-
vpn-instance *vpn-instance-name*	Specifies the name of a VPN instance.	The value must be an existing VPN instance name.
ipv6-address	Specifies a next-hop IPv6 address.	The address is a 32-digit hexadecimal number, in the format X:X:X:X:X:X:X:X.
link-local *link-local-address*	Specifies a link-local IPv6 address. The prefix of the specified IPv6 address must match FE80::/10.	The address is a 32-digit hexadecimal number, in the format X:X:X:X:X:X:X:X.
interface *interface-type interface-number*	Specifies the interface corresponding to the link-local address. ● *interface-type* specifies the interface type. ● *interface-number* specifies the interface number.	-
forced	Indicates that packets are discarded if the next hop is unavailable.	-

Views

Traffic behavior view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

You can configure the action of redirecting packets to a specified next-hop IPv6 address in a traffic behavior to implement PBR.

Follow-up Procedure

Run the **traffic policy** command to create a traffic policy and run the **classifier behavior** command in the traffic policy view to bind the traffic classifier to the

traffic behavior containing the action of redirecting packets to a next-hop IPv6 address.

Precautions

- In a traffic behavior, the **redirect ipv6-nexthop** command cannot be used together with any of the following commands: **remark 8021p**, **add-tag vlan-id**, **remark cvlan-id** and **remark vlan-id**.
- If no neighbor entry matches the next-hop IPv6 address, the device sends NS packets to check whether the neighbor is reachable. If the neighbor is unreachable, redirection does not take effect and packets are forwarded along the original path.
- If multiple next-hop IPv6 addresses are configured, the device redirects packets in active/standby mode. A maximum of four next-hop IPv6 addresses can be configured in a traffic behavior. The device determines the primary and backup paths according to the sequence in which next-hop IPv6 addresses are configured. The next-hop IPv6 address that is configured first has the highest priority and this next hop is used as the primary path. Other next hops are used as backup paths. When the primary link becomes Down, a next hop with a higher priority is used as the primary link.
- When the next-hop IPv6 address is a local-link address, you must configure a VLANIF interface. The address is valid only in the VLAN corresponding to the VLANIF interface.
- If you run the **redirect ipv6-nexthop** command in the same traffic classifier view multiple times, only the latest configuration takes effect.
- In V200R021C00SPC600 and later versions, when the action of redirecting packets to a remote next hop, the configuration does not take effect if the next-hop address is reachable only through tunnel forwarding.

Example

Configure the action of redirecting packets to the next-hop IPv6 address FC00:0:0:2001::1 in the traffic behavior **b1**.

```
<HUAWEI> system-view
[HUAWEI] traffic behavior b1
[HUAWEI-behavior-b1] redirect ipv6-nexthop fc00:0:0:2001::1
```

7.12 Route Monitoring Group Configuration Commands

7.12.1 Command Support

Commands provided in this section and all the parameters in the commands are supported by all switch models (except the S5731-L and S5731S-L), unless otherwise specified. For details, see specific commands.

7.12.2 display ip route-monitor-group

Function

The **display ip route-monitor-group** command displays information about a route monitoring group or all route monitoring groups.

Format

display ip route-monitor-group [*group-name*]

Parameters

Parameter	Description	Value
group-name	Displays information about a specified route monitoring group. If *group-name* is specified, detailed information about the specified route monitoring group is displayed. If *group-name* is not specified, a summary of all route monitoring groups is displayed.	The value is a string of 1 to 31 case-sensitive characters.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

After a route monitoring group is configured, you can run the **display ip route-monitor-group** command to check information about the route monitoring group, such as the status and information about routes in the route monitoring group.

Example

Display brief information about all route monitoring groups.

```
<HUAWEI> display ip route-monitor-group
 Route monitor group number : 1
 Route monitor group             State
 uplink                          Enable
```

Display detailed information about the route monitoring group **uplink**.

```
<HUAWEI> display ip route-monitor-group uplink
route monitor group uplink
 State       : Enabled
-------------------------------------------------------
 route monitor group tracking route number : 1
-------------------------------------------------------
```



```
VPN name     : Public
Destination  : 10.2.2.2
Mask         : 255.255.255.255
State        : Inactive
```

Table 7-236 Description of the **display ip route-monitor-group** command output

Item	Description
Route monitor group number	Number of route monitor groups.
Route monitor group	Name of a route monitoring group.
State	Status of a route monitoring group: ● Enabled: The route monitoring group is enabled. ● Disabled: The route monitoring group is disabled. ● Up: The route monitoring group is Up. ● Down: The route monitoring group is Down. You can run the **monitor enable** command to enable a route monitoring group.
route monitor group tracking route number	Number of routes in a route monitoring group.
VPN name	VPN instance to which routes belongs. The value **Public** indicates a public network.
Destination	Destination IP address of a route.
Mask	Mask of the destination IP address of a route.
State	Status of a route: ● Active ● Inactive

7.12.3 display ip route-monitor-group track-route

Function

The **display ip route-monitor-group track-route** command displays information about all route monitoring groups to which a specified route is added.

Format

display ip route-monitor-group track-route [**vpn-instance** *vpn-instance-name*] *dest-address* { *mask* | *mask-length* }

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the VPN instance to which a route belongs. If this parameter is specified, information about all route monitoring groups to which a VPN route in the VPN instance is added is displayed. If this parameter is not specified, information about all route monitoring groups to which a public network route is added is displayed.	The value is a string of 1 to 31 case-sensitive characters without spaces. If the string is enclosed in double quotation marks (" "), the string can contain spaces.
dest-address	Specifies the destination IP address of a route.	The value is in dotted decimal notation.
mask	Specifies the mask of the destination IP address of a route.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of the destination IP address of a route.	The value is an integer in the range from 0 to 32.

Views

All views

Default Level

1: Monitoring level

Usage Guidelines

A route can be added to multiple route monitoring groups to monitor the route status change. You can run the **display ip route-monitor-group track-route** command to display information about all route monitoring groups to which a specified route is added.

Example

Display information about all route monitoring groups to which the route to 192.168.1.0/24 is added.

```
<HUAWEI> display ip route-monitor-group track-route 192.168.1.0 24
 Route monitor group            State
 uplink                         Enabled
```


Table 7-237 Description of the **display ip route-monitor-group track-route** command output

Item	Description
Route monitor group	Name of a route monitoring group.
State	Status of a route monitoring group: • Enabled: The route monitoring group is enabled. • Disabled: The route monitoring group is disabled. • Up: The route monitoring group is Up. • Down: The route monitoring group is Down. You can run the **monitor enable** command to enable a route monitoring group.

7.12.4 ip route-monitor-group

Function

The **ip route-monitor-group** command creates a route monitoring group and displays the view of this route monitoring group, or displays the view of an existing route monitoring group.

The **undo ip route-monitor-group** command deletes a route monitoring group.

By default, no route monitoring group is created on a switch.

Format

ip route-monitor-group *group-name*

undo ip route-monitor-group *group-name*

Parameters

Parameter	Description	Value
group-name	Specifies the name of a route monitoring group.	The value is a string of 1 to 31 case-sensitive characters.

Views

System view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

All network-side routes of the same type can be added to the same route monitoring group, which monitors a certain number of routes. In a hot standby scenario, you can add network-side routes to a route monitoring group and associate access-side service modules with it so that the service modules can perform active/standby link switchovers upon route changes in the group. This mechanism prevents network congestion and packet loss.

Follow-up Procedure

Run the **track ip route** command to add a route to the route monitoring group.

Example

Create a route monitoring group named **uplink**.

```
<HUAWEI> system-view
[HUAWEI] ip route-monitor-group uplink
[HUAWEI-route-monitor-group-uplink]
```

7.12.5 monitor enable

Function

The **monitor enable** command enables a route monitoring group.

The **undo monitor enable** command disables a route monitoring group.

By default, a route monitoring group is disabled.

Format

monitor enable

undo monitor enable

Parameters

None

Views

Route monitoring group view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A route monitoring group can monitor status changes of network-side routes. After the routes to be monitored are added to a route monitoring group, run the

monitor enable command to enable the route monitoring group so that it can take effect immediately.

When a large number of routes are being added to or deleted from a route monitoring group, the status of the route monitoring group changes frequently, which in turn leads to service flapping of the service modules associated with the route monitoring group. To prevent such service flapping, run the **undo monitor enable** command to disable the route monitoring group so that it is dissociated from all service modules. After the routes are added or deleted, run the **monitor enable** command to re-associate the route monitoring group with all service modules.

Prerequisites

A route has been added to the route monitoring group using the **track ip route** command.

Example

Enable the route monitoring group **uplink**.

```
<HUAWEI> system-view
[HUAWEI] ip route-monitor-group uplink
[HUAWEI-route-monitor-group-uplink] track ip route 10.1.1.0 24
[HUAWEI-route-monitor-group-uplink] monitor enable
```

7.12.6 operator and

Function

The **operator and** command sets the status of routes in a route monitoring group to be of the AND relationship.

The **undo operator and** command restores the default configuration.

By default, the status of routes in a route monitoring group is of the OR relationship.

Format

operator and

undo operator and

Parameters

None

Views

Route monitoring group view

Default Level

2: Configuration level

Usage Guidelines

You can run the **operator and** command to set the status of routes in a route monitoring group to be of the AND relationship. That is, the status of the route monitoring group is Up only when all routes in the route monitoring group are Up. The status of the monitoring group is Down as long as one route in the group is Down. By default, the status of routes in a route monitoring group is of the OR relationship. That is, the status of the route monitoring group is Up as long as one route in the group is Up. The status of the monitoring group is Down only when all routes in the group are Down.

Example

Set the status of routes in a route monitoring group to be of the AND relationship.

```
<HUAWEI> system-view
[HUAWEI] ip route-monitor-group uplink
[HUAWEI-route-monitor-group-uplink] operator and
```

7.12.7 trigger-down-delay

Function

The **trigger-down-delay** command configures a delay after which the route management (RM) module instructs the associated service modules to perform link switchovers.

The **undo trigger-down-delay** command restores the default configuration.

By default, the delay is 0s.

Format

trigger-down-delay *delay-value*

undo trigger-down-delay [*delay-value*]

Parameters

Parameter	Description	Value
delay-value	Specifies a delay after which the RM module instructs the associated service modules to perform link switchovers.	The value is an integer in the range from 0 to 1000, in seconds.

Views

Route monitoring group view

Default Level

2: Configuration level

Usage Guidelines

In a hot standby scenario, service modules can be associated with a route monitoring group to trigger link switchovers. When the status of the route monitoring group goes Down, the RM module instructs the associated service modules to perform link switchovers. In addition, the RM module re-delivers routes to the forwarding table and establishes forwarding entries for the routes, which takes some time. Packet loss may occur if the RM module instructs the associated service modules to perform link switchovers immediately. To resolve this problem, run the **trigger-down-delay** command to configure a delay after which the RM module instructs the associated service modules to perform link switchovers.

Example

Set the delay to 10s after which the RM module instructs the associated service modules to perform link switchovers.

```
<HUAWEI> system-view
[HUAWEI] ip route-monitor-group uplink
[HUAWEI-route-monitor-group-uplink] trigger-down-delay 10
```

7.12.8 trigger-up-delay

Function

The **trigger-up-delay** command configures a delay after which the RM module instructs the associated service modules to perform link switchbacks.

The **undo trigger-up-delay** command restores the default configuration.

By default, the delay is 5s.

Format

trigger-up-delay *delay-value*

undo trigger-up-delay

Parameters

Parameter	Description	Value
delay-value	Specifies a delay after which the RM module instructs the associated service modules to perform link switchbacks.	The value is an integer in the range from 0 to 1000, in seconds.

Views

Route monitoring group view

Default Level

2: Configuration level

Usage Guidelines

In a hot standby scenario, service modules can be associated with a route monitoring group to trigger link switchovers. When the status of the route monitoring group goes Up, the RM module instructs the associated service modules to perform link switchbacks. In addition, the RM module re-delivers routes to the forwarding table and establishes forwarding entries for the routes, which takes some time. Packet loss may occur if the RM module instructs the associated service modules to perform link switchovers immediately. To resolve this problem, run the **trigger-up-delay** command to configure a delay after which the RM module instructs the associated service modules to perform link switchbacks.

Example

Set the delay to 10s after which the RM module instructs the associated service modules to perform link switchbacks.

```
<HUAWEI> system-view
[HUAWEI] ip route-monitor-group uplink
[HUAWEI-route-monitor-group-uplink] trigger-up-delay 10
```

7.12.9 track ip route

Function

The **track ip route** command adds a route to a route monitoring group.

The **undo track ip route** command deletes a route from a route monitoring group.

By default, no route is added to a route monitoring group.

Format

track ip route [**vpn-instance** *vpn-instance-name*] *dest-address* { *mask* | *mask-length* }

undo track ip route [**vpn-instance** *vpn-instance-name*] *dest-address* { *mask* | *mask-length* }

Parameters

Parameter	Description	Value
vpn-instance *vpn-instance-name*	Specifies the VPN instance to which a route belongs. If this parameter is specified, a VPN route in the VPN instance is added to the route monitoring group. If this parameter is not specified, a public network route is added to the route monitoring group.	The value is a string of 1 to 31 case-sensitive characters without spaces. If the string is enclosed in double quotation marks (" "), the string can contain spaces.
dest-address	Specifies the destination IP address of a route.	The value is in dotted decimal notation.
mask	Specifies the mask of the destination IP address of a route.	The value is in dotted decimal notation.
mask-length	Specifies the mask length of the destination IP address of a route.	The value is an integer in the range from 0 to 32.

Views

Route monitoring group view

Default Level

2: Configuration level

Usage Guidelines

Usage Scenario

A route monitoring group can monitor the status of its member routes to prevent network congestion and packet loss. You can run the **track ip route** command to add a route to be monitored to a route monitoring group. All network-side routes of the same type can be added to the same route monitoring group.

Follow-up Procedure

Run the **monitor enable** command to enable the route monitoring group.

Precautions

A maximum of 16 routes can be added to a route monitoring group. A route can be added to multiple route monitoring groups.

Example

Add the route to 10.1.1.0/24 to the route monitoring group **uplink**.


```
<HUAWEI> system-view
[HUAWEI] ip route-monitor-group uplink
[HUAWEI-route-monitor-group-uplink] track ip route 10.1.1.0 24
```

