

IP Routing Configuration

Оглавление

1 CONFIGURING ROUTING POLICIES	10
1.1 Overview	10
1.2 Applications	10
1.2.1 Route Filtering	10
1.2.2 Route Re-distribution	11
1.2.3 PBR	12
1.3 Features	13
1.3.1 Filtering List	13
1.3.2 Route Map	15
1.4 Configuration	20
1.4.1 Configuring a Route Map	21
1.4.2 Configuring a Filtering List	44
1.5 Monitoring	62
2 CONFIGURING KEYS	63
2.1 Overview	63
2.2 Applications	63
2.2.1 RIP Authentication	63
2.3 Features	64
2.3.1 Key Chain	64
2.4 Configuration	65
2.4.1 Configuring a Key Chain	65
2.5 Monitoring	70
3 CONFIGURING RIP	71
3.1 Overview	71
3.2 Applications	71
3.2.1 Basic RIP Application	71
3.2.2 Interworking Between RIP and BGP	72
3.3 Features	73
3.3.1 RIPv1 and RIPv2	74
3.3.2 Exchanging Routing Information	75
3.3.3 Routing Algorithm	78
3.3.4 Avoiding Route Loops	80

3.3.5	Security Measures	82
3.3.6	Reliability Measures	83
3.3.7	Multiple Instances	85
3.4	Configuration	85
3.4.1	Configuring RIP Basic Functions	89
3.4.2	Controlling Interaction of RIP Packets	96
3.4.3	Enabling Triggered Updates	101
3.4.4	Enabling Source Address Verification	104
3.4.5	Enabling Authentication	106
3.4.6	Enabling Route Summarization	110
3.4.7	Enabling Supernetting Routes	113
3.4.8	Advertising the Default Route or External Routes	115
3.4.9	Setting Route Filtering Rules	119
3.4.10	Modifying Route Selection Parameters	123
3.4.11	Modifying Timers	126
3.4.12	Enabling BFD Correlation	129
3.4.13	Enabling Fast Reroute	132
3.4.14	Enabling GR	135
3.4.15	Enabling Multiple Instances	137
3.5	Monitoring	141
4	CONFIGURING OSPFV2	143
4.1	Overview	143
4.2	Applications	144
4.2.1	Intra-Domain Interworking	145
4.2.2	Inter-Domain Interworking	145
4.3	Features	146
4.3.1	Link-State Routing Protocols	151
4.3.2	OSPF Route Management	156
4.3.3	Enhanced Security and Reliability	162
4.3.4	Network Management	167
4.4	Configuration	168
4.4.1	Configuring OSPF Basic Functions	175
4.4.2	Setting the Network Type	183
4.4.3	Configuring Route Redistribution and Default Route	187
4.4.4	Configuring Stub Area and NSSA Area	193
4.4.5	Configuring Route Summarization	200
4.4.6	Configuring Route Filtering	204

4.4.7	Modifying Route Cost and AD	208
4.4.8	Enabling Authentication	215
4.4.9	Enabling Overflow	220
4.4.10	Modifying the Maximum Number of Concurrent Neighbors	224
4.4.11	Disabling Source Address Verification	226
4.4.12	Disabling MTU Verification	228
4.4.13	Enabling Two-Way Maintenance	230
4.4.14	Enabling GR	232
4.4.15	Enabling NSR	236
4.4.16	Correlating OSPF with BFD	238
4.4.17	Enabling Fast Reroute	240
4.4.18	Enabling iSPF	244
4.4.19	Configuring the Network Management Function	248
4.4.20	Modifying Protocol Control Parameters	252
4.5	Monitoring	260
5	CONFIGURING OSPFV3	262
5.1	Overview	262
5.2	Applications	263
5.2.1	Intra-Domain Interworking	263
5.2.2	Inter-Domain Interworking	264
5.3	Features	265
5.3.1	Link-State Routing Protocols	270
5.3.2	OSPF Route Management	275
5.3.3	Enhanced Security and Reliability	280
5.3.4	Network Management Functions	284
5.4	Configuration	285
5.4.1	Configuring OSPF Basic Functions	290
5.4.2	Setting the Network Type	298
5.4.3	Configuring Route Redistribution and Default Route	302
5.4.4	Configuring the Stub Area and NSSA Area	309
5.4.5	Configuring Route Summarization	315
5.4.6	Configuring Route Filtering	319
5.4.7	Modifying the Route Cost and AD	323
5.4.8	Enabling Authentication	328
5.4.9	Modifying the Maximum Number of Concurrent Neighbors	333
5.4.10	Disabling MTU Verification	335
5.4.11	Enabling Two-Way Maintenance	337

5.4.12	Correlating OSPF with BFD	338
5.4.13	Enabling GR	341
5.4.14	Configuring Network Management Functions	345
5.4.15	Modifying Protocol Control Parameters	348
5.5	Monitoring	357
6	CONFIGURING IS-IS	359
6.1	Overview	359
6.2	Applications	359
6.2.1	Planar Topology	360
6.2.2	Hierarchical Topology	360
6.3	Features	361
6.3.1	IS-IS Network Hierarchy	364
6.3.2	IS-IS Address Coding Mode	365
6.3.3	IS-IS Packet Types	366
6.3.4	DIS Election	368
6.3.5	IS-IS Supported TLV Types	369
6.3.6	LSP Fragment Extension	371
6.3.7	IS-IS VRF	372
6.3.8	IS-IS MTR	374
6.3.9	IS-IS Neighbor	376
6.4	Configuration	377
6.4.1	Enabling IS-IS	387
6.4.2	Configuring IS-IS Hello Packets	393
6.4.3	Configuring IS-IS LSPs	396
6.4.4	Configuring IS-IS SNPs	406
6.4.5	Configuring the IS-IS Level Type	407
6.4.6	Configuring IS-IS Authentication	413
6.4.7	Configuring IS-IS GR	423
6.4.8	Configuring IS-IS NSR	427
6.4.9	Configuring BFD Support for IS-IS	429
6.4.10	Setting the IS-IS Overload Bit	432
6.4.11	Configuring IS-IS VRF	435
6.4.12	Configuring IS-IS MTR	438
6.4.13	Configuring SNMP for IS-IS	445
6.4.14	Configuring IS-IS Two-way Maintenance	448
6.4.15	Configuring Other IS-IS Parameters	450
6.5	Monitoring	470

7	CONFIGURING BGP	472
7.1	Overview	472
7.2	Applications	473
7.2.1	Inter-AS Route Advertisement	473
7.2.2	Intra-AS Route Reflection	474
7.3	Features	475
7.3.1	Creating a BGP Neighbor	479
7.3.2	Configuring a BGP Route Reflector	481
7.3.3	Configuring a BGP Alliance	483
7.3.4	Re-distributing Local AS Network Information to BGP	483
7.3.5	Controlling Route Exchange Between BGP Peers	484
7.3.6	Obtaining Accessible Networks of Other ASs from BGP	486
7.3.7	Configuring Synchronization Between BGP and IGP	487
7.3.8	Configuring BGP Soft Reset	488
7.3.9	Configuring the Route Attributes of BGP	489
7.3.10	Configuring BGP Route Aggregation	493
7.3.11	Configuring BGP Route Dampening	493
7.3.12	Configuring the Management Distance of BGP	495
7.3.13	Configuring Multi-path Load Balancing of BGP	496
7.3.14	Configuring BGP FRR	497
7.3.15	Configuring Fast Withdrawal of Specified BGP Routes	498
7.3.16	Configuring BGP Multi-path Bypass Protection	499
7.3.17	Allowing the BGP Route Reflector to Modify Route Attributes	500
7.3.18	Configuring the BGP Device to Advertise Routes with the Lowest Priority upon Restart	501
7.3.19	Configuring BGP Timers	501
7.3.20	Configuring BGP Route Update Mechanisms	503
7.3.21	Configuring the Next-Hop Triggering Update Function of BGP	504
7.3.22	Configuring BGP LOCAL AS	504
7.3.23	Configuring BGP Capacity Protection	505
7.3.24	Configuring BGP GR	507
7.3.25	Configuring 4-Byte AS Numbers of BGP	510
7.3.26	Configuring a Regular Expression	511
7.3.27	Configuring BGP Session Retention	512
7.3.28	Configuring BGP Delayed Advertisement upon System Restart	513
7.3.29	Configuring BGP Delayed Advertisement for First Routes	514
7.3.30	Configuring BGP NSR	515
7.3.31	Configuring BGP Routes to Be Recursive Only to Host Routes	516

7.3.32	Configuring Outbound Loop Detection for a BGP Neighbor	517
7.3.33	Shutting Down BGP Connections Gracefully	518
7.3.34	Configuring Enhanced VPN Route Import	519
7.3.35	BGP Route Update Group	519
7.3.36	Configuring BGP EVPN	520
7.3.37	Other Related Configurations	521
7.4	Configuration	521
7.4.1	Configuring a BGP Peer (Group)	523
7.4.2	Configuring MD5 Authentication	534
7.4.3	Configuring a Route Reflector	542
7.4.4	Configuring an AS Alliance	548
7.4.5	Configuring Multi-path Load Balancing of BGP	558
7.4.6	Configuring EBGp FRR	564
7.4.7	Configuring Fast Withdrawal of Specified BGP Routes	570
7.4.8	Configuring Local ASs	573
7.4.9	Configuring BGP GR	576
7.4.10	Configuring a BGP IPv6 Address Family	581
7.4.11	Configuring a BGP EVPN	587
7.4.12	Configuring Interconnection with Devices Supporting Only 2-Byte AS Numbers	598
7.4.13	Using Local IPv6 Link Addresses to Establish BGP Neighborships	601
7.4.14	Configuring BGP NSR	607
7.4.15	Configuring BGP Routes to Be Recursive Only to Host Routes	611
7.4.16	Configuring Outbound Loop Detection for a BGP Neighbor	613
7.4.17	Shutting Down BGP Connections Gracefully	616
7.4.18	Configuring EBGp Multi-path Bypass Protection	620
7.4.19	Configuring Inter-VRF Multi-Path Route Import	627
7.5	Monitoring	634
8	CONFIGURING RIPNG	636
8.1	Overview	636
8.2	Application	636
8.3	Features	637
8.3.1	RIPng and RIP	637
8.3.2	Exchanging Routing Information	638
8.3.3	Routing Algorithm	640
8.3.4	Avoiding Route Loops	642
8.4	Configuration	644
8.4.1	Configuring RIPng Basic Functions	645

8.4.2	Advertising the Default Route or External Routes	652
8.4.3	Setting Route Filtering Rules	655
8.4.4	Modifying Route Selection Parameters	657
8.4.5	Modifying Timers	660
8.5	Monitoring	662
9	CONFIGURING PBR	664
9.1	Overview	664
9.2	Applications	664
9.2.1	Selecting an ISP by Using PBR	664
9.2.2	Implementing Traffic Classification by Using PBR	665
9.3	Features	667
9.3.1	Configuring a Policy	667
9.3.2	Configuring PBR	671
9.4	Limitations	673
9.5	Configuration	673
9.5.1	Configuring Basic Functions of PBR	674
9.5.2	Setting Redundant Backup or Load Balancing	692
9.5.3	Configuring Source-Address-Based PBR	699
9.6	Monitoring	708
10	MANAGING ROUTES	711
10.1	Overview	711
10.2	Applications	711
10.2.1	Basic Functions of the Static Route	712
10.2.2	Floating Static Route	712
10.2.3	Load Balancing Static Route	713
10.2.4	Correlation of Static Routes with Track, BFD or ARP	714
10.2.5	Fast Reroute of Static Routes	715
10.2.6	ARP-to-host Conversion	716
10.3	Features	717
10.3.1	Route Computation	717
10.3.2	Optimal Route Selection	718
10.3.3	Default Route	719
10.3.4	Route Reliability	719
10.4	Limitations	720
10.5	Configuration	720
10.5.1	Configuring a Static Route	722

10.5.2	Configuring a Default Route	732
10.5.3	Configuring Route Limitations	738
10.5.4	Correlating a Static Route with BFD	742
10.5.5	Configure Static Fast Reroute	746
10.5.6	Configuring ARP-to-host Conversion	749
10.6	Monitoring	752
11	CONFIGURING VRF	754
11.1	Overview	754
11.2	Applications	754
11.2.1	Local Inter-VPN Access	754
11.2.2	VRF only on PEs	755
11.2.3	VRF on CEs and PEs (MCE Application)	756
11.3	Features	757
11.3.1	VPN Instance	757
11.3.2	VPN Route	758
11.3.3	VPN Route Attribute	759
11.4	Configuration	759
11.4.1	Configuring a Single-Protocol VRF Table	761
11.4.2	Configuring a Multiprotocol VRF Table	769
11.5	Monitoring	781
12	CONFIGURING L3MAC	783
12.1	Overview	783
12.2	Applications	783
12.2.1	Data Center Scenario	783
12.3	Configuration	784
12.3.1	Adding a Layer-3 MAC Address	784
12.4	Monitoring	786
13	CONFIGURING HARDWARE CAPACITY QUERY	787
13.1	Overview	787
13.2	Features	787
13.3	Monitoring	787

1 CONFIGURING ROUTING POLICIES

1.1 Overview

Routing policies are a policy set for changing the packet forwarding path or routing information and are often implemented by a filtering list and a route map. Routing policies are flexibly and widely applied in the following methods:

- Use a filtering list in a routing protocol to filter or modify routing information.
- Use a route map in a routing protocol to filter or modify routing information. Where, the route map can further use a filtering list.
- Use a route map in policy-based routing (PBR) to control packet forwarding or modify packet fields.

1.2 Applications


Application	Description
Route Filtering	Use a filtering list in a routing protocol to filter the routing information sent or received by the protocol.
Route Re-distribution	Use a route map in a routing protocol to filter or modify routing information and re-distribute RIP routes to OSPF. Only RIP routes with 4 hops can be re-distributed.
PBR	Use a route map in PBR to control packet forwarding or modify packet fields and specify optimum output interfaces for packets from different subnets.

1.2.1 Route Filtering

By default, a routing protocol advertises and learns all routing information. When a filtering list is used, the routing protocol advertises only required routes or receives only required routing information.

Scenario

Figure 1-1



As shown in Figure 1-1, router A has routes to 3 networks: 10.0.0.0, 20.0.0.0 and 30.0.0.0.

Configure a filtering list on the routers to achieve the following purposes:

- Filter the sent routing information on router A to filter routes that router A does not need to send.
- Filter the received routing information on router B to filter routes that router B does not need to learn.

Deployment

- Filter the sent routing information 30.0.0.0 on router A.
- Filter the received routing information 20.0.0.0 on router B to ensure that router B learns only routing information 10.0.0.0.

1.2.2 Route Re-distribution

By default, route re-distribution will re-distribute all routing information in a routing protocol to another routing protocol. All routing attributes will also be inherited. You can use a route map to perform conditional control for re-distribution between two routing protocols, including:

- Specify the range for re-distributing routes and re-distribute only routing information that meets certain rules.
- Set the attributes of routes generated by re-distribution.

Scenario

Figure 1-2



As shown in Figure 1-2, configure route re-distribution on the devices to achieve the following purposes:

- Re-distribute only RIP routes with 4 hops to OSPF.
- In the OSPF routing domain, the initial metric of this route is 40, the route type is the external route type-1 and the route tag value is set to 40.

Deployment

- Configure a route with 4 hops in the route map `rip_to_ospf`: match, and set the initial metric of this route to 40, the route type to the external route type-1 and the route tag value to 40.
- Configure route re-distribution to re-distribute RIP routes to OSPF and use the route map `rip_to_ospf`.


1.2.3 PBR

PBR is implemented by applying a route map including policies to interfaces and devices.

Similar to static routing, PBR is also manually configured, where recursive routing supports automatic update with network changes. As compared with static and dynamic routing, PBR is more flexible. Static and dynamic routing can forward packets only based on destination addresses. PBR can forward packets based on the source and destination addresses, packet length and input interface.

Scenario

Figure 1-3



Configure PBR on the layer-3 device DEV1 to achieve the following purposes:

- Packets from subnet 1 (200.24.16.0/24) are sent from GE0/1 first.
- Packets from subnet 2 (200.24.17.0/24) are sent from GE0/2 first.

Deployment

- Configure two different ACLs to match packets from subnets 1 and 2 respectively.

- Configure the route map RM_FOR_PBR: policy 10 is used to ensure that "packets from subnet 1 are sent from GE0/1 first"; policy 20 is used to ensure that "packets from subnet 2 are sent from GE0/2 first".
- Perform PBR for packets received from GE0/3 and use the route map RM_FOR_PBR.

1.3 Features

Overview

Feature	Description
Filtering List	Define a group of lists based on a route attribute, which can be used by a routing protocol for route filtering.
Route Map	A policy defines "if certain conditions are matched, you can perform certain processing actions".

1.3.1 Filtering List

Filtering lists are a group of lists defined based on a routing attribute and are a tool for filtering routing policies. Independent filtering lists are meaningless and can be used to filter routes only when they are applied in a routing protocol.

Working Principle

Based on different routing attributes, filtering lists are classified into the following types:

Access Control List (ACL)

ACLs comprise IPv4 and IPv6 ACLs. When defining ACLs, you can specify IPv4/IPv6 addresses and masks to match the destination network segment or next-hop addresses of routing information.

For description about ACLs, see the *ACL Configuration Guide*.

Address Prefix List (prefix-list)

Similar to ACLs, prefix-lists, including IPv4 prefix-lists and IPv6 prefix-lists, are used to match destination network segments of routing information during route filtering.

AS-Path List

AS-path lists are used only for BGP. They are used to match AS paths during BGP route filtering.

Community Attribute Filtering List (Community-List)

Community-lists are used only for BGP. They are used to match community attributes during BGP route filtering.

Extended Community Attribute Filtering List (Extcommunity-List)

Extcommunity-lists are used only for BGP. They are used to match extended community attributes during BGP route filtering.

Related Configuration

Creating an ACL

By default, no ACL is configured and no policy is set.

In the global configuration mode, run the **ip access-list** { **extended** | **standard** } { *id* | *name* } command to create an IPv4 ACL.

You can set multiple policies in an ACL, sorted by their sequence numbers. Policies have two working modes: permit and deny.

Creating a Prefix-List

By default, no prefix-list is configured and no entry is set.

In the global configuration mode, run the **ip prefix-list** *prefix-list-name* [**seq** *seq-number*] { **deny** | **permit** } *ip-prefix* [**ge** *minimum-prefix-length*] [**le** *maximum-prefix-length*] command to create an IPv4 prefix-list and add a prefix entry to the list.

You can set multiple entries in the prefix-list, sorted by their sequence numbers. Entries have two working modes: permit and deny.

Run the **ip prefix-list** *prefix-list-name* **description** *description-text* command to add description to the prefix-list.

Run the **ip prefix-list sequence-number** command to enable the sorting function for the prefix-list.

Creating an AS-Path List

By default, no AS-path list is configured and no entry is set.

In the global configuration mode, run the **ip as-path access-list** *path-list-num* { **permit** | **deny** } *regular-expression* command to create an AS-path list and add an entry to the list.

You can set multiple entries in the AS-path list. Entries have two working modes: permit and deny.

Creating a Community-List

By default, no community-list is configured and no entry is set.

In the global configuration mode, run the **ip community-list** { { **standard** | **expanded** } *community-list-name* | *community-list-number* } { **permit** | **deny** } [*community-number..*] command to create a community-list and add an entry to the list.

You can set multiple entries in the community-list. Entries have two working modes: permit and deny.

Creating an Extcommunity-List

By default, no excommunity-list is configured and no entry is set.

In the global configuration mode, run the **ip extcommunity-list** *{standard-list | standard list-name}* **{ permit | deny }** [*rt value*] [*soo value*] command to create a standard extcommunity list and add an entry to the list.

Run the **ip extcommunity-list** *{expanded-list | expanded list-name}* **{ permit | deny }** [*regular-expression*] command to create an extcommunity list and add an entry to the list.

You can also run the **ip extcommunity-list** *{expanded-list | expanded list-name | standard-list | standard list-name}* command to create an extcommunity list and enter the configuration mode of **ip extcommunity-list** to add entries.

You can set multiple entries in the extcommunity-list. Entries have two working modes: permit and deny.

1.3.2 Route Map

A policy is a "match ..., set..." statement, which indicates that "if certain conditions are matched, you can perform some processing actions".

Working Principle

Executing policies

A route map may contain multiple policies. Each policy has a corresponding sequence number. A smaller sequence number means a higher priority. Policies are executed based on their sequence numbers. Once the matching condition of a policy is met, the processing action for this policy needs to be performed and the route map exits. If no matching condition of any policy is met, no processing action will be performed.

Working Modes Of Policies

Policies have two working modes:

- permit: When the matching condition of a policy is met, the processing action for this policy will be performed and the route map will exit.
- deny: When the matching condition of a policy is met, the processing action for this policy will not be performed and the route map will exit.

Matching Conditions Of Policies

The matching condition of a policy may contain 0, 1 or more match rules.

- If the matching condition contains 0 match rule, no packet will be matched.
- If the matching condition contains one or more match rules, all rules must be matched.

Processing Action for a Policy

The processing action of a policy may contain 0, 1 or more set rules.

- If the processing action contains 0 set rule, no processing action will be performed and the route map will directly exit.
- If the processing action contains one or more set rules, all processing actions will be performed and then the route map will exit.

 If set rules have different priorities, the set rule with the highest priority will take effect.

Related Configuration

Creating a Route Map (Policy)

By default, no route map is configured and no policy is set.

In the global configuration mode, you can run the **route-map** *route-map-name* [**permit** | **deny**] [*sequence-number*] command to create a route map and add a policy to the route map.

You can set multiple policies in a route map. Each policy uses different sequence numbers.

Setting Matching Conditions of a Policy

By default, no match rule is set (that is, the matching condition of a policy contains 0 match rule).

In the route map mode, run the **match** command to set match rules. One **match** command is mapped to one match rule.

RGOS provides abundant **match** commands for setting flexible matching conditions.

Command	Description
match as-path	Uses the AS_PATH attribute of a BGP route as the matching condition.
match community	Uses the community attribute of a BGP route as the matching condition.
match extcommunity	Uses the extended community attribute of a BGP route as the matching condition.
match interface	Uses the output interface of a route as the matching condition.

match ip address	Uses the destination IPv4 address of a route as the matching condition.
match ip next-hop	Uses the next-hop IPv4 address of a route as the matching condition.
match ip route-source	Uses the source IPv4 address of a route as the matching condition.
match ipv6 address	Uses the destination IPv6 address of a route as the matching condition.
match ipv6 next-hop	Uses the next-hop IPv6 address of a route as the matching condition.
match ipv6 route-source	Uses the source IPv6 address of a route as the matching condition.
match metric	Uses the metric of a route as the matching condition.
match origin	Uses the source of a route as the matching condition.
match route-type	Uses the type of a route as the matching condition.
match tag	Uses the tag value of a route as the matching condition.

Setting the Processing Actions of a Policy

By default, no set rule is configured (that is, the processing action of a policy contains 0 set rule).

In the route map mode, run the **set** command to configure set rules. One **set** command is mapped to one set rule.

RGOS provides abundant **set** commands for setting flexible processing actions.

Command	Description
set aggregator as	Modifies the AS attribute value of a route aggregator.
set as-path prepend	Adds a specified as-path attribute value.
set as-path replace	Replace the specified as-path attribute value.
set atomic-aggregate	Sets the atomic-aggregate attribute of a route.

set comm-list delete	Deletes all community attribute values from the community attribute list for a route matching the match rules.
set community	Sets the community attribute value of a route.
set dampening	Sets the flapping parameters of a route.
set extcomm-list delete	Deletes all extended community attribute values from the extcommunity attribute list for a route matching the match rules.
set extcommunity	Sets the extended community attribute value of a route.
set fast-reroute	Sets the backup output interface and next hop of a fast reroute.
set ip default nexthop	Specifies the default next hop of a route. This command has a lower priority than a common route and a higher priority than set default interface .
set ip dscp	Modifies the dscp field of an IP packet.
set ip global next-hop	Specifies the next hop of a route.
set ip global default next-hop	Specifies the default next hop of a route.
set ip nexthop	Specifies the next hop of a route. This command has a higher priority than set interface .
set ip next-hop recursive	Specifies the recursive next-hop IP address of a route.
set ip next-hop verify-availability	Specifies the next-hop IP address of a route and checks the accessibility of the next hop by using BFD.
set ip precedence	Modifies the precedence field of an IP packet.
set ip tos	Modifies the tos field of an IP packet.
set ipv6 default next-hop	Specifies the default next hop of a route. This command has a lower priority than a common route and a higher priority than the default route.

set ipv6 global next-hop	Specifies the IPv6 next hop of a route.
set ipv6 global default next-hop	Specifies the default IPv6 next hop of a route.
set ipv6 next-hop	Specifies the IPv6 next hop of a route. This command has a higher priority than a common route.
set ipv6 next-hop verify-availability	Specifies the next-hop IP address of a route and checks the accessibility of the next hop by using BFD.
set ipv6 precedence	Sets the priority of an IPv6 packet header.
set level	Sets the destination area type to which a route will be directed.
set local-preference	Sets the local-preference attribute value of a route.
set metric	Modifies the metric value of a route.
set metric-type	Sets the metric type of a route.
set next-hop	Sets the next-hop IP address of a route.
set ip next-hop verify-availability	Specifies the next-hop IP address of a route and checks the accessibility of the next hop by using a third-party protocol.
set origin	Sets the source attribute of a route.
set originator-id	Sets the originator IP address of a route.
set tag	Sets the tag value of a route.
set distance	Sets the administrative distance.
set weight	Sets the weight value of a route.

1.4 Configuration

Configuration	Description and Command	
Configuring a Route Map	⚠ (Optional) It is used to define a policy.	
	route-map	Creates a policy (route map).
	match	Sets the matching conditions of the policy.
	set	Sets the processing actions of the policy.
Configuring a Filtering List	⚠ (Optional) It is used to define a filtering list.	
	ip as-path	Defines AS path filtering rules.
	ip community-list	Defines a community list.
	ip extcommunity-list	Defines an extcommunity list.
	ip prefix-list	Creates a prefix-list.
	ip prefix-list description	Adds description to a prefix-list.
	ip prefix-list sequence-number	Enables the sorting function for a prefix-list.
	Ipv6 prefix-list	Creates an IPv6 prefix-list.
	ipv6 prefix-list description	Adds description to an IPv6 prefix-list.
ipv6 prefix-list sequence-number	Enables the sorting function for an IPv6 prefix-list.	

1.4.1 Configuring a Route Map

Configuration

Effect

- Define a set of routing policies to be used by routing protocols or PBR.

Notes

- If a **match** command uses an ACL to define packet matching conditions, the ACL must be configured.
- The following **match** commands cannot be configured at the same time:

The Following match Commands	Cannot Be Configured with the Following match Commands At the Same Time
match ip address	match ip prefix-list
match ipv6 address	match ipv6 prefix-list
match ip next-hop	match ip next-hop prefix-list
match ipv6 next-hop	match ipv6 next-hop prefix-list
match ip route-source	match ip route-source prefix-list
match ipv6 route-source	match ipv6 route-source prefix-list

- The following **set** commands cannot be configured at the same time:

The Following set Commands	Cannot Be Configured with the Following set Commands At the Same Time
set ip next-hop	set ip next-hop verify-availability
set ip dscp	set ip tos
set ip dscp	set ip precedence

Configuration

Steps

Creating a Policy (Route Map)

- Mandatory.
- Perform this configuration on a device to which a policy needs to be applied.

Setting Matching Conditions of a Policy

- Optional.
- If no match rule is configured, no packet will be matched.
- If multiple match rules are configured, all the match rules must be matched.
- Perform this configuration on a device to which a policy needs to be applied.

Setting the Processing Actions of a Policy

- Optional.
- If no set rule is configured, no processing action will be performed.
- If multiple set rules are configured, all set rules must be executed (if the set rules have different priorities, the set rule with the highest priority takes effect).
- Perform this configuration on a device to which a policy needs to be applied.

Verification

- Check the configurations of the route map.

Related Commands

Creating a Policy (Route Map)

Command	<code>route-map route-map-name [{ permit deny } sequence]</code>
Parameter Description	<p><i>route-map-name</i>: Indicates the name of a route map, comprising not more than 32 characters.</p> <p><i>permit</i>: Specifies the working mode of this policy as permit, which is the default mode.</p> <p><i>deny</i>: Specifies the working mode of this policy as deny. The default mode is permit.</p> <p><i>sequence</i>: Specifies the sequence number of this policy. A smaller value means a higher priority. The default value is 10.</p>
Command Mode	Global configuration mode
Usage Guide	<p>If this route map is unavailable, this command will create a route map and add a policy to the route map.</p> <p>If this route map is available, this command will add a policy to the route map.</p>

Setting Matching Conditions of a Policy

Command	match as-path <i>as-path-acl-list-number</i> [<i>as-path-acl-list-number</i>....]
Parameter Description	<i>as-path-acl-list-number</i> : Indicates the AS-PATH list number, ranging from 1 to 500.
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the AS-PATH attribute of a BGP route. Run the ip as-path access-list <i>path-list-num</i> { permit deny } <i>regular-expression</i> command to configure the AS-PATH list.

Command	match community { <i>community-list-number</i> <i>community-list-name</i> } [exact-match] [{ <i>community-list-number</i> <i>community-list-name</i> } [exact-match] ...]
Parameter Description	<i>community-list-number</i> : Indicates the community list number. For a standard community list, the value ranges from 1 to 99. For an extcommunity list, the value ranges from 100 to 199. <i>community-list-name</i> : Indicates the community list name, comprising not more than 80 characters. exact-match : Indicates the exact match list. It is a non-exact match list by default, that is, the match rule is met as long as the routing attributes contain the attributes specified by a community list.
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the community attribute specified in a community list.

Command	match extcommunity { <i>standard-list-number</i> <i>standard-list-name</i> <i>expanded-list-num</i> <i>expanded-list-name</i> }
Parameter Description	<i>standard-list-number</i> : Indicates an ID, ranging from 1 to 99. It is used to identify a standard extcommunity list. One extcommunity list may contain multiple extcommunity values.

	<p><i>standard-list-name</i>: Indicates the name of a standard extcommunity list. It is used to identify the name of a standard extcommunity list. One extcommunity list may contain multiple extcommunity values.</p> <p><i>expanded-list-num</i>: Indicates an ID, ranging from 100 to 199. It is used to identify an extcommunity list. One extcommunity list may contain multiple extcommunity values.</p> <p><i>expanded-list-name</i>: Indicates the name of an extcommunity. It is used to identify the name of an extcommunity list. One extcommunity list may contain multiple extcommunity values.</p>
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the extended community attribute specified in an extcommunity list.

Command	match interface <i>interface-type interface-number</i> [...<i>interface-type interface-number</i>]
Parameter Description	<i>interface-type interface-number</i> : Indicates the interface type and interface number.
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the next-hop output interface of a route or a packet.

Command	match ip address { <i>access-list-number</i> [<i>access-list-number...</i> <i>access-list-name...</i>] <i>access-list-name</i> [<i>access-list-number...</i> <i>access-list-name</i>] prefix-list <i>prefix-list-name</i> [<i>prefix-list-name...</i>] }
Parameter Description	<p><i>access-list-number</i>: Indicates the access list number. For a standard access list, the value ranges are 1 to 99 and 1300 to 1999. For an extended access list, the value ranges are 100 to 199 and 2000 to 2699.</p> <p><i>access-list-name</i>: Indicates the access list name.</p> <p>prefix-list <i>prefix-list-name</i>: Indicates the name of a prefix-list to be matched.</p>

Command Mode	Route map configuration mode
Usage Guide	This match rule matches the destination IPv4 address of a packet or route by using an ACL or a prefix-list. An ACL and a prefix-list cannot be configured at the same time.

Command	match ip next-hop { <i>access-list-number</i> [<i>access-list-number...</i> <i>access-list-name...</i>] <i>access-list-name</i> [<i>access-list-number...</i> <i>access-list-name</i>] prefix-list <i>prefix-list-name</i> [<i>prefix-list-name...</i>] }
Parameter Description	<i>access-list-number</i> : Indicates the access list number. For a standard access list, the value ranges are 1 to 99 and 1300 to 1999. For an extended access list, the value ranges are 100 to 199 and 2000 to 2699. <i>access-list-name</i> : Indicates the access list name. prefix-list <i>prefix-list-name</i> : Indicates the name of a prefix-list to be matched.
Command Mode	Route map configuration mode
Usage Guide	This match rule matches the next-hop IPv4 address of a route by using an ACL or a prefix-list. An ACL and a prefix-list cannot be configured at the same time.

Command	match ip route-source { <i>access-list-number</i> [<i>access-list-number...</i> <i>access-list-name...</i>] <i>access-list-name</i> [<i>access-list-number...</i> <i>access-list-name</i>] prefix-list <i>prefix-list-name</i> [<i>prefix-list-name...</i>] }
Parameter Description	<i>access-list-number</i> : Indicates the access list number. For a standard access list, the value ranges are 1 to 99 and 1300 to 1999. For an extended access list, the value ranges are 100 to 199 and 2000 to 2699. <i>access-list-name</i> : Indicates the access list name. prefix-list <i>prefix-list-name</i> : Indicates the name of a prefix-list to be matched.
Command Mode	Route map configuration mode

Usage Guide	This match rule matches the source IPv4 address of a route by using an ACL or a prefix-list. An ACL and a prefix-list cannot be configured at the same time.
-------------	--

Command	match ipv6 address { <i>access-list-name</i> prefix-list <i>prefix-list-name</i> }
Parameter Description	<i>access-list-name</i> : Indicates the access list name. <i>prefix-list prefix-list-name</i> : Indicates the name of an IPv6 prefix-list to be matched.
Command Mode	Route map configuration mode
Usage Guide	This match rule matches the destination IPv6 address of a packet or route by using an ACL or a prefix-list. An ACL and a prefix list cannot be configured at the same time.

Command	match ipv6 next-hop { <i>access-list-name</i> prefix-list <i>prefix-list-name</i> }
Parameter Description	<i>access-list-name</i> : Indicates the access list name. <i>prefix-list prefix-list-name</i> : Indicates the name of an IPv6 prefix-list to be matched.
Command Mode	Route map configuration mode
Usage Guide	This match rule matches the next-hop IPv6 address of a route by using an ACL or a prefix-list. An ACL and a prefix-list cannot be configured at the same time.

Command	match ipv6 route-source { <i>access-list-name</i> prefix-list <i>prefix-list-name</i> }
Parameter Description	<i>access-list-name</i> : Indicates the access list name. <i>prefix-list prefix-list-name</i> : Indicates the name of an IPv6 prefix-list to be matched.
Command Mode	Route map configuration mode

Usage Guide	This match rule matches the source IPv6 address of a route by using an ACL or a prefix-list. An ACL and a prefix-list cannot be configured at the same time.
-------------	--

Command	match metric <i>metric</i>
Parameter Description	<i>metric</i> : Indicates the metric value of a route, ranging from 0 to 4,294,967,295.
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the metric value of a route.

Command	match origin { <i>egp</i> <i>igp</i> <i>incomplete</i> }
Parameter Description	<i>egp</i> : Indicates the source is remote EGP. <i>igp</i> : Indicates the source is local IGP. <i>incomplete</i> : Indicates that the source is an incomplete type.
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the source of a route.

Command	match route-type { <i>local</i> <i>internal</i> <i>external</i> [<i>type-1</i> <i>type-2</i>] <i>level-1</i> <i>level-2</i> }
Parameter Description	<i>local</i> : Indicates a route locally generated. <i>Internal</i> : Indicates an internal OSPF route. <i>external</i> : Indicates an external route (that of BGP or OSPF). <i>type-1</i> <i>type-2</i> : Indicates type-1 or type-2 external route of OSPF. <i>level-1</i> <i>level-2</i> : Indicates level-1 or level-2 route of ISIS.

Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the type of a route.

Command	match tag <i>tag</i> [...<i>tag</i>]
Parameter Description	<i>tag</i> : Indicates the tag value of a route.
Command Mode	Route map configuration mode
Usage Guide	This match rule is used to match the tag value of a route.

Setting the Processing Actions of a Policy

Command	set aggregator as <i>as-number</i> <i>ip-address</i>
Parameter Description	<i>as-number</i> : Indicates the AS number of an aggregator. The AS number ranges from 1 to 4,294,967,295, which can be indicated by 1 to 65535.65535 in the dot mode. <i>ip-address</i> : Indicates the address of an aggregator.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to modify the AS attribute value of a route's aggregator.

Command	set as-path prepend <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates the AS number to be added to the AS_PATH attribute. The AS number ranges from 1 to 4,294,967,295, which can be indicated by 1 to 65535.65535 in the dot mode.

Command Mode	Route map configuration mode
Usage Guide	This set rule is used to add a specified as-path attribute value.

Command	set as-path replace as-number
Parameter Description	<i>as-number</i> : Indicates the AS number which will replace the AS_PATH attribute. The AS number ranges from 1 to 4,294,967,295, which can be indicated by 1 to 65535.65535 in the dot mode.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to replace the specified as-path attribute value.

Command	set atomic-aggregate
Parameter Description	-
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the atomic-aggregate attribute of a route.

Command	set comm-list { <i>community-list-number</i> <i>community-list-name</i> } delete
Parameter Description	<p><i>community-list-number</i>: Indicates the community list number. For a standard community list, the value ranges from 1 to 99. For an extcommunity list, the value ranges from 100 to 199.</p> <p><i>community-list-name</i>: Indicates the community list name, comprising not more than 80 characters.</p>

Command Mode	Route map configuration mode
Usage Guide	This rule is used to delete all community attribute values from the community list for a route matching the match rules.

Command	set community { <i>community-number</i> [<i>community-number ...</i>] additive none }
Parameter Description	<i>community-number</i> : Indicates the community attribute value. additive : Adds a number based on the original community attribute. none : Keeps the community attribute empty.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the community attribute value of a route.

Command	set dampening <i>half-life reuse suppress max-suppress-time</i>
Parameter Description	<i>half-life</i> : half-life when a route is accessible or not accessible, ranging from 1 to 45 minutes. The default value is 15 minutes. <i>reuse</i> : When the penalty value of a route is smaller than this value, route suppression will be canceled. The value ranges from 1 to 20,000 and the default value is 750. <i>suppress</i> : When the penalty value of a route is greater than this value, the route will be suppressed. The value ranges from 1 to 20,000 and the default value is 2,000. <i>max-suppress-time</i> : Indicates the longest time that a route can be suppressed, ranging from 1 to 255 minutes. The default value is 4 x half-life .
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the flapping parameters of a route.

Command	set extcomm-list { <i>extcommunity-list-number</i> <i>extcommunity-list-name</i> } delete
Parameter Description	<p><i>extcommunity-list-number</i>: Indicates the extcommunity list number. For a standard extcommunity list, the value ranges from 1 to 99. For an extended extcommunity list, the value ranges from 100 to 199.</p> <p><i>extcommunity-list-name</i>: Indicates the extcommunity list name, comprising not more than 80 characters.</p>
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to delete all extended community attribute values from the extcommunity attribute list for a route matching the match rules.

Command	set extcommunity { rt <i>extend-community-value</i> soo <i>extend-community-value</i> }
Parameter Description	<p>rt: Sets the RT attribute value of a route.</p> <p>soo: Sets the SOO attribute value of a route.</p> <p><i>extend-community-value</i>: Indicates the value of an extended community.</p>
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the extended community attribute value of a route.

Command	set fast-reroute backup-interface <i>interface-type interface-number</i> [backup-nexthop <i>ip-address</i>]
Parameter Description	<p><i>interface-type interface-number</i>: Specifies a backup output interface.</p> <p>backup-nexthop ip-address: Specifies a backup next hop. For a non-point-to-point interface, a backup next hop must be specified.</p>
Command Mode	Route map configuration mode

Usage Guide	This set rule is used to set the backup output interface and next hop of a fast reroute.
-------------	--

Command	set ip default next-hop <i>ip-address</i> [<i>weight</i>] [...<i>ip-address</i> [<i>weight</i>]]
Parameter Description	<i>ip-address</i> : Indicates the next-hop IP address. <i>weight</i> : Indicates the weight of this next hop.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to specify the default next hop of a route.

Command	set ip dscp <i>dscp_value</i>
Parameter Description	<i>dscp_value</i> : Sets the DSCP value in the IP header of an IP packet.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to modify the dscp field of an IP packet.

Command	set ip next-hop recursive <i>ip-address</i>
Parameter Description	<i>ip-address</i> : Indicates the recursive next-hop IP address.
Command Mode	Route map configuration mode

Usage Guide	<p>This command is used only for PBR configuration.</p> <p>This set rule is used to specify the recursive next hop of a route. An IP address can recur to a static or dynamic route that has an output interface and a next-hop IP address. A maximum of 32 next hops are supported. If a recursive route is a static route, only one next hop is supported for the static recursive route.</p>
-------------	---

Command	set ip next-hop verify-availability <i>ip-address</i> [track <i>track-obj-number</i> bfd <i>interface-type interface-number gateway</i>]
Parameter Description	<p><i>ip-address</i>: Indicates the next-hop IP address.</p> <p>track: Judges whether the next hop is effective by using <i>Track</i>.</p> <p>bfd: Indicates that BFD is used for neighbor detection.</p> <p><i>interface-type</i>: Configures the interface type.</p> <p><i>interface-number</i>: Configures the interface number.</p> <p><i>gateway</i>: Configures the gateway IP address, which is the neighbor IP address of BFD. If the next hop is configured as the neighbor, BFD will be used to detect the accessibility of the forwarding path.</p>
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to specify the next hop of a route, and BFD is used to fast detect the effectiveness of the next hop.

Command	set ip precedence { <i>number</i> critical flash flash-override immediate internet network priority routine }
Parameter Description	<p><i>number</i>: Indicates the priority of the IP header with a number, ranging from 0 to 7.</p> <p>7: critical</p> <p>6: flash</p> <p>5: flash-override</p> <p>4: immediate</p> <p>3: internet</p> <p>2: network</p>

	<p>1: priority 0: routine</p> <p>critical flash flash-override immediate internet network priority routine: priority of an IP header.</p>
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to modify the precedence field of an IP packet header.

Command	set ip tos { <i>number</i> max-reliability max-throughput min-delay min-monetary-cost normal }
Parameter Description	<p><i>number</i>: Indicates the TOS value of an IP header with a number, ranging from 0 to 15.</p> <p>2: max-reliability 4: max-throughput 8: min-delay 1: min-monetary-cost 0: normal</p> <p>max-reliability max-throughput min-delay min-monetary-cost normal: priority of an IP header.</p>
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to modify the tos field of an IP packet.

Command	set ipv6 default next-hop <i>global-ipv6-address</i> [<i>weight</i>] [<i>global-ipv6-address</i> [<i>weight</i>] ...]
Parameter Description	<p><i>global-ipv6-address</i>: Indicates the next-hop IPv6 address for packet forwarding. The next-hop router must be a neighbor router.</p> <p><i>weight</i>: Indicates the weight in the load balancing mode, ranging from 1 to 8. A larger value means larger packet traffic to be shared by the next hop.</p>

Command Mode	Route map configuration mode
Usage Guide	This set rule is used to specify the default next hop IPv6 address of a route.

Command	set ipv6 next-hop <i>global-ipv6-address</i> [<i>weight</i>] [<i>global-ipv6-address</i> [<i>weight</i>] ...]
Parameter Description	<i>global-ipv6-address</i> : Indicates the next-hop IPv6 address for packet forwarding. The next-hop router must be a neighbor router. <i>weight</i> : Indicates the weight in the load balancing mode, ranging from 1 to 8. A larger value means larger packet traffic to be shared by the next hop.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to specify the next hop IPv6 address of a route.

Command	set ipv6 next-hop verify-availability <i>global-ipv6-address</i> bfd <i>interface-type</i> <i>interface-number</i> <i>gateway</i>
Parameter Description	<i>global-ipv6-address</i> : Indicates the next-hop IPv6 address. bfd: Indicates that BFD is used for neighbor detection. <i>interface-type</i> : Configures the interface type. <i>interface-number</i> : Configures the interface number. <i>gateway</i> : Configures the gateway IPv6 address, which is the neighbor IPv6 address of BFD. If the next hop is configured as the neighbor, BFD will be used to detect the accessibility of the forwarding path.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to specify the next hop of a route and BFD is used to fast detect the effectiveness of the next hop.

Command	set ipv6 precedence { <i>number</i> critical flash flash-override immediate internet network priority routine }
Parameter Description	<p><i>number</i>: Indicates the priority of the IP header with a number, ranging from 0 to 7.</p> <p>7: critical</p> <p>6: flash</p> <p>5: flash-override</p> <p>4: immediate</p> <p>3: internet</p> <p>2: network</p> <p>1: priority</p> <p>0: routine</p> <p>critical flash flash-override immediate internet network priority routine: priority of an IP header.</p>
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the priority of an IPv6 packet header.

Command	set level { level-1 level-2 level-1-2 stub-area backbone }
Parameter Description	<p>level-1: Indicates that the re-distribution route is advertised to ISIS Level 1.</p> <p>level-2: Indicates that the re-distribution route is advertised to ISIS Level 2.</p> <p>level-1-2: Indicates that the re-distribution route is advertised to ISIS Level 1 and Level 2.</p> <p>stub-area: Indicates that the re-distribution route is advertised to OSPF Stub Area.</p> <p>backbone: Indicates that the re-distribution route is advertised to the OSPF backbone area.</p>
Command Mode	Route map configuration mode

Usage Guide	This set rule is used to set the destination area type to which a route will be redirected.
-------------	---

Command	set local-preference <i>number</i>
Parameter Description	<i>number</i> : Indicates the metric value of a local priority, ranging from 0 to 4,294,967,295. A larger value means a higher priority.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the local-preference attribute value of a route.

Command	set metric [+ <i>metric-value</i> - <i>metric-value</i> <i>metric-value</i>]
Parameter Description	+: Increases (based on the metric value of the original route). -: Decreases (based on the metric value of the original route). <i>metric-value</i> : Sets the metric value of a re-distribution route. A larger value means a lower priority.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to modify the metric value of a route.

Command	set metric-type <i>type</i>
Parameter Description	<i>type</i> : Sets the type of a re-distribution route. The default type of an OSPF re-distribution route is type-2.
Command Mode	Route map configuration mode

Usage Guide	This set rule is used to set the metric type.
-------------	---

Command	set next-hop <i>ip-address</i>
Parameter Description	<i>ip-address</i> : Indicates the next-hop IP address.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the next-hop IP address.

Command	set origin { <i>egp</i> <i>igp</i> <i>incomplete</i> }
Parameter Description	<i>egp</i> : Indicates the source is remote EGP. <i>igp</i> : Indicates the source is local IGP. <i>incomplete</i> : Indicates that the source is the incomplete type and generally refers to a route generated due to re-distribution.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the source attribute of a route.

Command	set originator-id <i>ip-address</i>
Parameter Description	<i>ip-address</i> : Indicates the address of an originator.
Command Mode	Route map configuration mode

Usage Guide	This set rule is used to set the originator IP address of a route.
-------------	--

Command	set tag <i>tag</i>
Parameter Description	<i>tag</i> : Sets the tag of a re-distribution route.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the tag value of a route.

Command	set distance <i>number</i>
Parameter Description	<i>number</i> : The greater the administrative value, the lower the priority.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the administrative distance.

Command	set weight <i>number</i>
Parameter Description	<i>number</i> : Sets the weight of a route, ranging from 0 to 65,535. A larger value means a higher priority.
Command Mode	Route map configuration mode
Usage Guide	This set rule is used to set the weight of a route.


Displaying the Configurations of a Route Map

Command	show route-map [<i>name</i>]
Parameter Description	<i>name</i> : Specifies a route map.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	Run the show route-map command to display the configurations of a route map. If an ACL is used when a route map is configured, you can run the show access-list command to display the configurations of the ACL.

Configuration

Example

Using a Route Map in Route Re-distribution to Filter and Modify Routing Information

Scenario Figure 1-4	<p>As shown in Figure 1-4, a device is connected to both an OSPF routing domain and RIP routing domain.</p> 
	<ul style="list-style-type: none"> Re-distribute only RIP routes with 4 hops to OSPF. In the OSPF route domain, if the route type is the external route type-1, set the tag value of the route to 40. Re-distribute only OSPF routes with the tag value 10 to RIP. In the RIP route domain, set the initial metric value of this route to 10.
Configuration Steps	<ul style="list-style-type: none"> Configure the route map redrip: Match a route with 4 hops, set the initial metric value of the route to 40, set the route type to the external route type-1, and set the tag value of the route to 40. Configure the route map redospf: match a route with the tag value 10 and set the initial metric value of the route to 10. Configure re-distribution of the RIP route to OSPF and apply the route map redrip. Configure re-distribution of the OSPF route to RIP and apply the route map redospf.
	<pre>QTECH(config)# route-map redrip permit 10 QTECH(config-route-map)# match metric 4 QTECH(config-route-map)# set metric-type type-1</pre>

	<pre> QTECH(config-route-map)# set tag 40 QTECH(config-route-map)# exit QTECH(config)# route-map redospf permit 10 QTECH(config-route-map)# match tag 10 QTECH(config-route-map)# set metric 10 QTECH(config-route-map)# exit QTECH(config)# router ospf 1 QTECH(config-router)# redistribute rip subnets route-map redrip QTECH(config-router)# exit QTECH(config)# router rip QTECH(config-router)# redistribute ospf 1 route-map redospf QTECH(config-router)# exit </pre>
Verification	<ul style="list-style-type: none"> ▪ Check the configurations of the route map to verify the policy rules. ▪ Check the OSPF routing information library to verify that the rules matching the policy rules are re-distributed.
	<pre> QTECH# show route-map route-map redrip, permit, sequence 10 Match clauses: metric 4 Set clauses: metric 40 metric-type type-1 tag 40 route-map redospf, permit, sequence 10 Match clauses: tag 10 Set clauses: metric 10 </pre>
	<pre> QTECH# show ip ospf database external OSPF Router with ID (192.100.1.9) (Process ID 1) </pre>

AS External Link States

LS age: 5

Options: 0x2 (-|-|-|-|-|E|-)

LS Type: AS-external-LSA

Link State ID: 192.168.199.0 (External Network Number)

Advertising Router: 192.100.1.9

LS Seq Number: 80000001

Checksum: 0x554d

Length: 36

Network Mask: /24

Metric Type: 1

TOS: 0


Metric: 4

Forward Address: 0.0.0.0

External Route Tag: 40

Applying a Route Map in PBR

Scenario
Figure 1-5



Configure PBR on the device DEV1 to achieve the following purposes:

- Packets from subnet 1 (200.24.16.0/24) are sent from GE0/1 first.
- Packets from subnet 2 (200.24.17.0/24) are sent from GE0/2 first.
- The two output links work in the mutual backup mode.

Configuration Steps	<ul style="list-style-type: none"> ▪ Configure two different ACLs to match packets from subnets 1 and 2 respectively. ▪ Configure the route map RM_FOR_PBR: policy 10 is used to ensure that "packets from subnet 1 are sent from GE0/1 first"; policy 20 is used to ensure that "packets from subnet 2 are sent from GE0/2 first". ▪ Configure PBR for packets received from GE0/3 and apply the route map RM_FOR_PBR. ▪ Set PBR to implement redundant backup among multiple next hops. <p>⚠ In the redundant backup mode, the sequence of multiple set next hops is the sequence of the priorities for taking effect.</p>
	<pre> QTECH(config)# access-list 1 permit 200.24.16.0 0.0.0.255 QTECH(config)# access-list 2 permit 200.24.17.0 0.0.0.255 QTECH(config)# route-map RM_FOR_PBR 10 QTECH(config-route-map)# match ip address 1 QTECH(config-route-map)# set ip next-hop 200.24.18.1 QTECH(config-route-map)# set ip next-hop 200.24.19.1 QTECH(config-route-map)# exit QTECH(config)# route-map RM_FOR_PBR 20 QTECH(config-route-map)# match ip address 2 QTECH(config-route-map)# set ip next-hop 200.24.19.1 QTECH(config-route-map)# set ip next-hop 200.24.18.1 QTECH(config-route-map)# exit QTECH(config)# interface GigabitEthernet 0/3 QTECH(config-if)# ip policy route-map RM_FOR_PBR QTECH(config)# ip policy redundancy </pre>
Verification	<ul style="list-style-type: none"> ▪ Check the configurations of PBR to verify that the route map is applied to the interfaces. ▪ Check the configurations of the route map to verify the policy rules. ▪ Check the ACL configurations to verify the packet filtering rules.
	<pre> QTECH# show ip policy Balance mode: redundancy Interface Route map GigabitEthernet 0/3 RM_FOR_PBR ! </pre>

	<pre>QTECH# show route-map route-map RM_FOR_PBR, permit, sequence 10 Match clauses: ip address 1 Set clauses: ip next-hop 200.24.18.1 ip next-hop 200.24.19.1 route-map RM_FOR_PBR, permit, sequence 20 Match clauses: ip address 2 Set clauses: ip next-hop 200.24.19.1 ip next-hop 200.24.18.1</pre>
	<pre>QTECH# show access-lists ip access-list standard 1 10 permit 200.24.16.0 0.0.0.255 10 permit 200.24.16.0 0.0.0.255 ip access-list standard 2 10 permit 200.24.17.0 0.0.0.255</pre>

Common Errors

- After matching of ACLs and prefix-lists is configured, the corresponding ACLs and prefix lists are not defined.

1.4.2 Configuring a Filtering List

Configuration Effect

- Define a set of route filtering rules to be used by routing protocols.

Notes

- A configured filtering list can take effect only after it is associated with a routing protocol.

Configuration Steps

Configuring a Prefix-List

- To filter address prefixes, you should perform this configuration.
- If there is no special requirement, you should perform this configuration on a route for which filtering based on a prefix-list needs to be performed.

Configuring an AS Path List

- To filter address prefixes, you should perform this configuration.
- If there is no special requirement, you should perform this configuration on a route for which filtering based on an AS path needs to be performed.

Configuring a Community List

- To filter community attributes, you should perform this configuration.
- If there is no special requirement, you should perform this configuration on a route for which community attributes need to be filtered.

Configuring an Extcommunity List

- To filter extended community attributes, you should perform this configuration.
- If there is no special requirement, you should perform this configuration on a route for which extended community attributes need to be filtered.

Verification

- Check whether the filtering list is correctly configured.
- Check the routing table to verify that routes can be correctly filtered.

Related Commands

Defining AS Path Filtering Rules

Command	<code>ip as-path access-list <i>path-list-num</i> { permit deny } <i>regular-expression</i></code>
Parameter Description	<p><i>path-list-num</i>: Indicates an AS-path ACL name based on a regular expression and is an AS path list identifier, ranging from 1 to 500.</p> <p>permit: Permits access.</p> <p>deny: Denies access.</p> <p><i>regular-expression</i>: Indicates a regular expression, ranging from 1 to 255.</p>

Command Mode	Global configuration mode
Usage Guide	-

Defining a Community List

Command	ip community-list { { standard expanded } <i>community-list-name</i> <i>community-list-number</i> } { permit deny } [<i>community-number..</i>]
Parameter Description	<p>standard: Indicates a standard community list.</p> <p>expanded: Indicates an extended community list.</p> <p><i>community-list-name:</i> Indicates the community list name, comprising not more than 80 characters.</p> <p><i>community-list-number:</i> Indicates the community list number. For a standard community list, the value ranges from 1 to 99. For an extended community list, the value ranges from 100 to 199.</p> <p>permit: Permits access.</p> <p>deny: Denies access.</p> <p><i>community-number:</i> Indicates the community attribute value.</p>
Command Mode	Global configuration mode
Usage Guide	Use this command to define a community list used for BGP.

Defining an Extcommunity List

Command	ip extcommunity-list {<i>expanded-list</i> expanded <i>list-name</i> } { permit deny } [<i>regular-expression</i>]
Parameter Description	<i>expand-list:</i> Indicates an extended extcommunity list, ranging from 100 to 199. One extcommunity list may contain multiple rules.

	<p>standard-list: Indicates a standard extcommunity list, ranging from 1 to 99. One extcommunity list may contain multiple rules.</p> <p>expanded list-name: Indicates the name of an extended extcommunity, comprising not more than 32 characters. When using this parameter, you enter the extcommunity list configuration mode.</p> <p>standard list-name: Indicates the name of a standard extcommunity list, comprising not more than 32 characters. When using this parameter, you enter the extcommunity list configuration mode.</p> <p>permit: Defines an extcommunity rule for permitting.</p> <p>deny: Defines an extcommunity rule for denying.</p> <p>regular-expression: (optional) Defines a matching template that is used to match an extcommunity.</p> <p>sequence-number: (Optional) Defines the sequence number of a rule, ranging from 1 to 2,147,483,647. If no sequence number is specified, the sequence number automatically increases by 10 when a rule is added by default. The initial number is 10.</p> <p>rt: (Optional) Sets the RT attribute value. This command can be used only for the standard extcommunity configuration, but not for the extended extcommunity configuration.</p> <p>soo: (Optional) Sets the SOO attribute value. This command can be used only for the standard extcommunity configuration, but not for the extended extcommunity configuration.</p> <p>value: Indicates the value of an extended community (extend_community_value).</p>
Command Mode	Global configuration mode and ip extcommunity-list configuration mode
Usage Guide	-

Creating a Prefix-List

Command	ip prefix-list <i>prefix-list-name</i> [seq <i>seq-number</i>] { deny permit } <i>ip-prefix</i> [ge <i>minimum-prefix-length</i>] [le <i>maximum-prefix-length</i>]
Parameter Description	<p>prefix-list-name: Indicates the prefix-list name.</p> <p>seq-number: Assigns a sequence number to an prefix-list entry, ranging from 1 to 2,147,483,647. If this command does not contain the sequence number, the system will assign a default sequence number to the prefix-list entry. The default sequence number of the first entry is 5. Subsequently, the default sequence number of each</p>

	<p>entry not assigned with a value is the first multiple of 5 greater than the previous sequence number.</p> <p>deny: Denies access when certain conditions are matched.</p> <p>permit: Permits access when certain conditions are matched.</p> <p>ip-prefix: Configures the IP address and mask, ranging from 0 to 32 digits.</p> <p>minimum-prefix-length: Specifies the minimum range (namely, the start length of a range).</p> <p>maximum-prefix-length: Specifies the maximum range (namely, the end length of a range).</p>
Command Mode	Global configuration mode
Usage Guide	-

Adding Description to a Prefix-List

Command	ip prefix-list <i>prefix-list-name</i> description <i>description-text</i>
Parameter Description	<p><i>prefix-list-name</i>: Indicates the prefix-list name.</p> <p><i>description-text</i>: Describes the prefix-list.</p>
Command Mode	Global configuration mode
Usage Guide	-

Enabling the Sorting Function for a Prefix-List

Command	ip prefix-list sequence-number
Parameter Description	-
Command Mode	Global configuration mode

Usage Guide	-
-------------	---

Creating an IPv6 Prefix-List

Command	<code>ipv6 prefix-list <i>prefix-list-name</i> [seq <i>seq-number</i>] { deny permit } <i>ipv6-prefix</i> [ge <i>minimum-prefix-length</i>] [le <i>maximum-prefix-length</i>]</code>
Parameter Description	<p><i>prefix-list-name</i>: Indicates the prefix-list name.</p> <p><i>seq-number</i>: Assigns a sequence number to an prefix-list entry, ranging from 1 to 2,147,483,647. If this command does not contain the sequence number, the system will assign a default sequence number to the prefix-list entry. The default sequence number of the first entry is 5. Subsequently, the default sequence number of each entry not assigned with a value is the first multiple of 5 greater than the previous sequence number.</p> <p>deny: Denies access when certain conditions are matched.</p> <p>permit: Permits access when certain conditions are matched.</p> <p><i>ipv6-prefix</i>: Configures the IP address and mask, ranging from 0 to 128 digits.</p> <p><i>minimum-prefix-length</i>: Specifies the minimum range (namely, the start length of a range).</p> <p><i>maximum-prefix-length</i>: Specifies the maximum range (namely, the end length of a range).</p>
Command Mode	Global configuration mode
Usage Guide	-

Adding Description to an IPv6 Prefix List

Command	<code>ipv6 prefix-list <i>prefix-list-name</i> description <i>description-text</i></code>
Parameter Description	<p><i>prefix-list-name</i>: Indicates the prefix list name.</p> <p><i>description-text</i>: Describes the prefix list.</p>
Command Mode	Global configuration mode

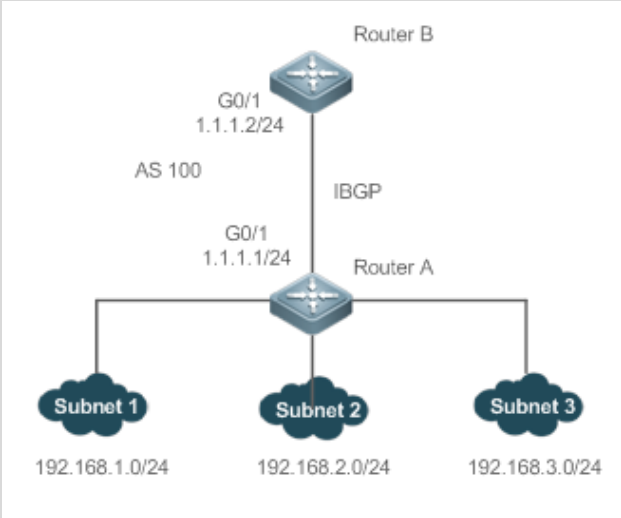
Usage Guide	-
-------------	---

Enabling the Sorting Function for an IPv6 Prefix-List

Command	ipv6 prefix-list sequence-number
Parameter Description	-
Command Mode	Global configuration mode
Usage Guide	-

Configuration Example

Configuring a Prefix-List

<p>Scenario Figure 1-6</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure an IBGP neighbor and advertise the neighbor to the three connected subnets. ▪ Configure a prefix-list. ▪ Associate a prefix-list with A to filter sent routes.

A	<pre>A# configure terminal A(config)# ip prefix-list pre1 permit 192.168.1.0/24 A(config)# router bgp 100 A(config-router)# neighbor 1.1.1.2 prefix-list pre1 out A(config-router)# end</pre>
Verification	<ul style="list-style-type: none"> ▪ Run the show command to display the prefix-list. ▪ Run the show command to display the BGP routing table to check whether the filtering behavior is correct.
A	<pre>A# show ip prefix-list ip prefix-list pre1: 1 entries seq 5 permit 192.168.1.0/24 A# show ip bgp BGP table version is 2, local router ID is 1.1.1.1 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S Stale, b - backup entry Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 192.168.1.0 0.0.0.0 0 32768 i *> 192.168.2.0 0.0.0.0 0 32768 i *> 192.168.3.0 0.0.0.0 0 32768 i Total number of prefixes 3</pre>
B	<pre>B# show ip bgp BGP table version is 4, local router ID is 1.1.1.2 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S Stale Origin codes: i - IGP, e - EGP, ? - incomplete</pre>

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i192.168.1.0	1.1.1.1	0	100	0	i
Total number of prefixes 1					

Configuring an AS Path List

<p>Scenario Figure 1-7</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Create an AS-path filtering rule to match path information including only AS 200. ▪ Establish EBGP neighborship on A with B and C. ▪ Associate an AS-path list with A to filter the routes received from B and C.
<p>A</p>	<pre>A(config)# ip as-path access-list 123 permit ^200\$ A(config)# router bgp 100 A(config)# neighbor 192.168.1.2 filter-list 123 in A(config)# neighbor 192.168.2.2 filter-list 123 in</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Run the show command to display the AS-path list. ▪ Run the show command to display the BGP routing table to check whether the filtering behavior is correct.
<p>A</p>	<pre>A# show ip as-path-access-list AS path access list 123 permit ^200\$</pre>

//When no AS-path list is associated with A, run the **show** command to check the BGP routing table.

```
A(config)# show ip bgp
```

```
BGP table version is 1, local router ID is 1.1.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
S Stale, b - backup entry
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight Path
*> 10.0.0.0/24	192.168.1.2	0		0 200 i
*> 20.0.0.0/24	192.168.2.2	0		0 300 i

```
Total number of prefixes 2
```

//When an AS-path list is associated with A, run the **show** command to display the BGP routing table and check whether the filtering behavior is correct.

```
A(config)# show ip bgp
```

```
BGP table version is 1, local router ID is 1.1.1.1
```

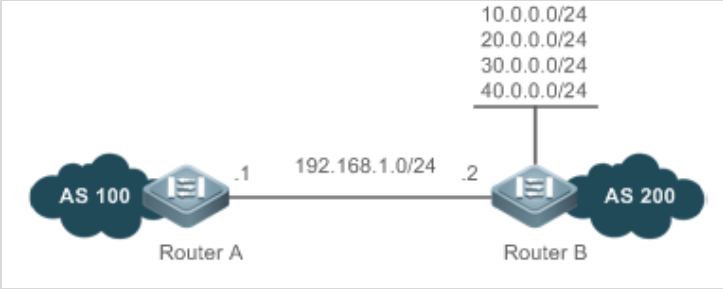
```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
S Stale, b - backup entry
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight Path
*> 10.0.0.0/24	192.168.1.2	0		0 200 i

```
Total number of prefixes 1
```

Configuring a Community List

<p>Scenario Figure 1-8</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Define a standard community list to match the community attribute 100: 20. ▪ Establish EBGP neighborhood between A and B. ▪ Advertise a route with the community attribute on B. ▪ Associate the community list on A (BGP can be applied only through a route map) to filter routes received on B.
<p>A</p>	<pre>A(config)# ip community-list standard test permit 100:20 A(config)# route-map COM A(config-route-map)# match community test A(config-route-map)# exit A(config)# router bgp 100 A(config-router)# neighbor 192.168.1.2 route-map COM in</pre>
<p>B</p>	<pre>B(config)# route-map comm1 B(config-route-map)# set community 100:20 200:20 B(config-route-map)# route-map comm2 B(config-route-map)# set community 100:20 B(config-route-map)# route-map comm3 B(config-route-map)# set community 200:20 B(config-route-map)# exit B(config)# router bgp 200 B(config-router)# neighbor 192.168.1.1 send-community B(config-router)# network 10.0.0.0 mask 255.255.255.0 route-map comm1 B(config-router)# network 20.0.0.0 mask 255.255.255.0 route-map comm2 B(config-router)# network 30.0.0.0 mask 255.255.255.0 route-map comm3 B(config-router)# network 40.0.0.0 mask 255.255.255.0</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Run the show command to display the community list.

- Run the **show** command to display the BGP routing table to check whether the filtering behavior is correct.

A

```
A# show ip community-list
Named Community standard list test
permit 100:20
```

//When no community list is associated with A, run the **show** command to check the BGP routing table.

```
A# show ip bgp
```

```
BGP table version is 1, local router ID is 192.168.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
          S Stale, b - backup entry
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight Path
*> 10.0.0.0/24	192.168.1.2	0		0 200 i
*> 20.0.0.0/24	192.168.1.2	0		0 200 i
*> 30.0.0.0/24	192.168.1.2	0		0 200 i
*> 40.0.0.0/24	192.168.1.2	0		0 200 i

```
Total number of prefixes 4
```

```
A# show ip bgp 10.0.0.0
```

```
BGP routing table entry for 10.0.0.0/24
```

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
```

```
Not advertised to any peer
```

```
200
```

```
192.168.1.2 from 192.168.1.2 (192.168.1.2)
```

```
Origin IGP, metric 0, localpref 100, valid, external, best
```

```
Community: 100:20 200:20
```

```
Last update: Wed Nov 6 18:58:18 2013
```

```
A# show ip bgp 20.0.0.0
```

```
BGP routing table entry for 20.0.0.0/24
```

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
```

```
Not advertised to any peer
200
  192.168.1.2 from 192.168.1.2 (192.168.1.2)
  Origin IGP, metric 0, localpref 100, valid, external, best
  Community: 100:20
  Last update: Wed Nov 6 18:58:18 2013
A# show ip bgp 30.0.0.0
BGP routing table entry for 30.0.0.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200
  192.168.1.2 from 192.168.1.2 (192.168.1.2)
  Origin IGP, metric 0, localpref 100, valid, external, best
  Community: 200:20
  Last update: Wed Nov 6 18:58:18 2013
A# show ip bgp 40.0.0.0
BGP routing table entry for 40.0.0.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200
  192.168.1.2 from 192.168.1.2 (192.168.1.2)
  Origin IGP, metric 0, localpref 100, valid, external, best
  Last update: Wed Nov 6 18:58:18 2013
```

//When a community list is associated with A, run the **show** command to display the BGP routing table and check whether the filtering behavior is correct.

```
A# show ip bgp
BGP table version is 1, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               S Stale, b - backup entry
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight Path
*> 10.0.0.0/24	192.168.1.2	0		0 200 i
*> 20.0.0.0/24	192.168.1.2	0		0 200 i

Total number of prefixes 2

A#

A# show ip bgp 10.0.0.0

BGP routing table entry for 10.0.0.0/24

Paths: (1 available, best #1, table Default-IP-Routing-Table)

Not advertised to any peer

200

192.168.1.2 from 192.168.1.2 (192.168.1.2)

Origin IGP, metric 0, localpref 100, valid, external, best

Community: 100:20 200:20

Last update: Wed Nov 6 19:02:49 2013

A# show ip bgp 20.0.0.0

BGP routing table entry for 20.0.0.0/24

Paths: (1 available, best #1, table Default-IP-Routing-Table)

Not advertised to any peer

200

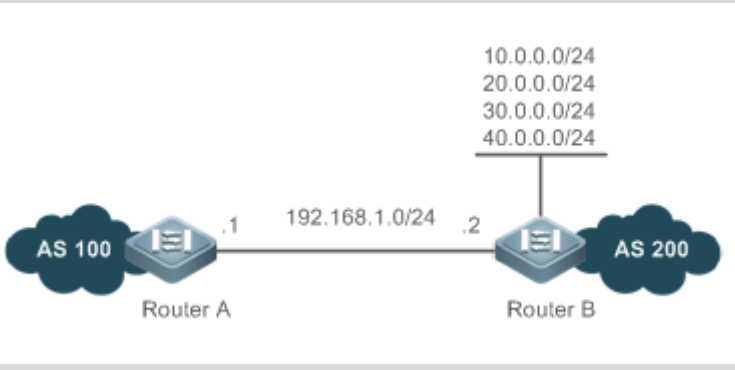
192.168.1.2 from 192.168.1.2 (192.168.1.2)

Origin IGP, metric 0, localpref 100, valid, external, best

Community: 100:20

Last update: Wed Nov 6 19:02:49 2013

Configuring an Extcommunity List

<p>Scenario Figure 1-9</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Define an extcommunity list to match the extcommunity attribute RT 1: 100. ▪ Establish EBGP neighborship between A and B. ▪ Advertise a route with the extcommunity attribute on B. ▪ Associate the extcommunity list with A (BGP can be applied only through a route map) to filter routes received on B.
<p>A</p>	<pre>A(config)# ip extcommunity-list 10 permit rt 1:100 A(config)# route-map EXTCOM A(config-route-map)# match extcommunity 10 A(config-route-map)# exit A(config)# router bgp 100 A(config-router)# neighbor 192.168.1.2 route-map EXTCOM in</pre>
<p>B</p>	<pre>B(config)# route-map ecomm1 B(config-route-map)# set extcommunity rt 1:100 2:200 B(config-route-map)# route-map ecomm2 B(config-route-map)# set extcommunity rt 1:100 B(config-route-map)# route-map ecomm3 B(config-route-map)# set extcommunity rt 2:200 B(config-route-map)# exit B(config)# router bgp 200 B(config-router)# neighbor 192.168.1.1 send-community both B(config-router)# network 10.0.0.0 mask 255.255.255.0 route-map ecomm1 B(config-router)# network 20.0.0.0 mask 255.255.255.0 route-map ecomm2 B(config-router)# network 30.0.0.0 mask 255.255.255.0 route-map ecomm3 B(config-router)# network 40.0.0.0 mask 255.255.255.0</pre>

Verification	<ul style="list-style-type: none"> Run the show command to display the extcommunity list. Run the show command to display the BGP routing table to check whether the filtering behavior is correct. 																									
A	<pre>EG1000M(config)#show ip extcommunity-list Extended community standard list 10 10 permit RT:1:100</pre>																									
	<p>//When no extcommunity list is associated with A, run the show command to check the BGP routing table.</p> <pre>A# show ip bgp BGP table version is 1, local router ID is 192.168.1.1 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S Stale, b - backup entry Origin codes: i - IGP, e - EGP, ? - incomplete</pre> <table border="1"> <thead> <tr> <th>Network</th> <th>Next Hop</th> <th>Metric</th> <th>LocPrf</th> <th>Weight Path</th> </tr> </thead> <tbody> <tr> <td>*> 10.0.0.0/24</td> <td>192.168.1.2</td> <td>0</td> <td></td> <td>0 200 i</td> </tr> <tr> <td>*> 20.0.0.0/24</td> <td>192.168.1.2</td> <td>0</td> <td></td> <td>0 200 i</td> </tr> <tr> <td>*> 30.0.0.0/24</td> <td>192.168.1.2</td> <td>0</td> <td></td> <td>0 200 i</td> </tr> <tr> <td>*> 40.0.0.0/24</td> <td>192.168.1.2</td> <td>0</td> <td></td> <td>0 200 i</td> </tr> </tbody> </table> <p>Total number of prefixes 4</p> <pre>A# A# show ip bgp 10.0.0.0 BGP routing table entry for 10.0.0.0/24 Paths: (1 available, best #1, table Default-IP-Routing-Table) Not advertised to any peer 200 192.168.1.2 from 192.168.1.2 (192.168.1.2) Origin IGP, metric 0, localpref 100, valid, external, best Extended Community: RT:1:100 RT:2:200 Last update: Wed Nov 6 19:15:12 2013</pre>	Network	Next Hop	Metric	LocPrf	Weight Path	*> 10.0.0.0/24	192.168.1.2	0		0 200 i	*> 20.0.0.0/24	192.168.1.2	0		0 200 i	*> 30.0.0.0/24	192.168.1.2	0		0 200 i	*> 40.0.0.0/24	192.168.1.2	0		0 200 i
Network	Next Hop	Metric	LocPrf	Weight Path																						
*> 10.0.0.0/24	192.168.1.2	0		0 200 i																						
*> 20.0.0.0/24	192.168.1.2	0		0 200 i																						
*> 30.0.0.0/24	192.168.1.2	0		0 200 i																						
*> 40.0.0.0/24	192.168.1.2	0		0 200 i																						

```
A# show ip bgp 20.0.0.0
BGP routing table entry for 20.0.0.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Not advertised to any peer
200
  192.168.1.2 from 192.168.1.2 (192.168.1.2)
    Origin IGP, metric 0, localpref 100, valid, external, best
    Extended Community: RT:1:100
    Last update: Wed Nov 6 19:15:12 2013
```

```
A# show ip bgp 30.0.0.0
BGP routing table entry for 30.0.0.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Not advertised to any peer
200
  192.168.1.2 from 192.168.1.2 (192.168.1.2)
    Origin IGP, metric 0, localpref 100, valid, external, best
    Extended Community: RT:2:200
    Last update: Wed Nov 6 19:15:12 2013
```

```
A# show ip bgp 40.0.0.0
BGP routing table entry for 40.0.0.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Not advertised to any peer
200
  192.168.1.2 from 192.168.1.2 (192.168.1.2)
    Origin IGP, metric 0, localpref 100, valid, external, best
    Last update: Wed Nov 6 19:15:12 2013
```


//When an extcommunity list is associated with A, run the **show** command to display the BGP routing table and check whether the filtering behavior is correct.

A# show ip bgp

BGP table version is 1, local router ID is 192.168.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S Stale, b - backup entry

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight Path
*> 10.0.0.0/24	192.168.1.2	0		0 200 i
*> 20.0.0.0/24	192.168.1.2	0		0 200 i

Total number of prefixes 2

A#

A# show ip bgp 10.0.0.0

BGP routing table entry for 10.0.0.0/24

Paths: (1 available, best #1, table Default-IP-Routing-Table)

Not advertised to any peer

200

192.168.1.2 from 192.168.1.2 (192.168.1.2)

Origin IGP, metric 0, localpref 100, valid, external, best

Extended Community: RT:1:100 RT:2:200

Last update: Wed Nov 6 19:17:04 2013

A# show ip bgp 20.0.0.0

BGP routing table entry for 20.0.0.0/24

Paths: (1 available, best #1, table Default-IP-Routing-Table)

Not advertised to any peer

200

192.168.1.2 from 192.168.1.2 (192.168.1.2)

Origin IGP, metric 0, localpref 100, valid, external, best

Extended Community: RT:1:100

Common

Errors

- A filtering list is configured but is not correctly applied in a routing protocol, which causes that the filtering list cannot take effect.

1.5 Monitoring

Displaying

Description	Command
Displays the configurations of a route map.	show route-map [<i>route-map-name</i>]
Displays the configurations of an ACL.	show access-lists [<i>id</i> <i>name</i>]
Displays the configurations of an IPv4 prefix-list.	show ip prefix-list [<i>prefix-name</i>]
Displays the configurations of an IPv6 prefix-list.	show ipv6 prefix-list [<i>prefix-name</i>]
Displays the configurations of an AS-path list.	show ip as-path-access-list [<i>num</i>]
Displays the configurations of a community list.	show ip community-list [<i>community-list-number</i> <i>community-list-name</i>]
Displays the configurations of an extcommunity list.	show ip extcommunity-list [<i>extcommunity-list-num</i> <i>extcommunity-list-name</i>]

2 CONFIGURING KEYS

2.1 Overview

Keys are a kind of parameters that are used in algorithms for conversion from plain text to cipher text or from cipher text to plain text.

Plain text and cipher text authentication are supported for packet authentication in a routing protocol, during which keys need to be used.

i At present, keys are used only for RIP and ISIS packet authentication.

2.2 Applications


Application	Description
RIP Authentication	RIP uses keys for packet authentication.

2.2.1 RIP Authentication

Scenario

Network devices run RIP and use the MD5 authentication mode to increase the protocol security.

Figure 2-1



Deployment

- Configure a key chain on A. Configure RIP to enable packet authentication and use the key chain.
- Configure a key chain on B. Configure RIP to enable packet authentication and use the key chain.

2.3 Features

Overview

Feature	Description
Key Chain	Provide a tool for authentication in a routing protocol.

2.3.1 Key Chain

Working

Principle

A key chain may contain multiple different keys. Each key contains the following attributes:

- Key ID: Identifies a key. In the current key chain, keys and IDs are mapped in the one-to-one manner.
- Authentication string: Indicates a set of key characters used for verifying the consistency of authentication strings in a routing protocol.
- Lifetime: Specifies the lifetime of the current key for sending or receiving packets. Different authentication keys can be used in different periods.

Related

Configuration

Creating a Key Chain and a Key

In the global configuration mode, run the **key chain** *key-chain-name* command to define a key chain and enter the key chain configuration mode.

In the key chain configuration mode, run the **key** *key-id* command to define a key and enter the key chain key configuration mode.

Configuring an Authentication String

In the key chain key configuration mode, run the **key-string** *[0|7] text* command to specify an authentication string.


- A plain text authentication string is configured by default. The value **0** indicates that a plain text authentication key is configured.
- The value **7** indicates that a cipher text authentication string is configured.
- The encryption authentication service is disabled by default. You can run the **service password-encryption** command to enable the encryption service to forcibly convert plain text authentication into cipher text.

Configuring Lifetime

In the key chain key configuration mode, you can configure the lifetime of a key chain in the receiving and sending directions.

- **accept-lifetime** *start-time* { **infinite** | *end-time* | **duration seconds** }: Configures the lifetime of a key chain in the receiving direction.
- **send-lifetime** *start-time* { **infinite** | *end-time* | **duration seconds** }: Configures the lifetime of a key chain in the sending direction.

2.4 Configuration

Configuration	Description and Command	
Configuring a Key Chain	 (Mandatory) It is used to create a key.	
	key chain	Creates a key chain.
	key	Configures a key ID.
	key-string	Configures a key string.
	accept-lifetime	Configures the lifetime in the receiving direction.
	send-lifetime	Configures the lifetime in the sending direction.

2.4.1 Configuring a Key Chain

Configuration

Effect

- Define a key chain to be used by a routing protocol.

Notes

- A key chain can take effect only after it is associated with a routing protocol.

Configuration

Steps

Creating a Key Chain

- This configuration is mandatory if a key chain needs to be used.

- If there is no special requirement, you should perform this configuration on all routers for which routing protocol authentication needs to be performed.

Configuring a Key ID

- This configuration is mandatory if a key chain needs to be used.
- If there is no special requirement, you should perform this configuration on all routers for which routing protocol authentication needs to be performed.

Configuring a Key String

- This configuration is mandatory if a key chain needs to be used.
- If there is no special requirement, you should perform this configuration on all routers for which routing protocol authentication needs to be performed.

Configure the Lifetime in the Receiving Direction

- Optional.
- If the lifetime in the sending direction is not configured, the key chain will be always effective.

Configure the Lifetime in the Sending Direction

- Optional.
- If the lifetime in the sending direction is not configured, the key chain will be always effective.

Verification

- Use keys in a routing protocol and observe the neighborhood established by the routing protocol. If the keys are inconsistent, the neighborhood fails to be established.

Related Commands

Configuring a Key Chain

Command	key chain <i>key-chain-name</i>
Parameter Description	<i>key-chain-name</i> : Indicates the name of a key chain.
Command Mode	Global configuration mode
Usage Guide	To make a key chain take effect, you must configure at least one key.

Configuring a Key ID

Command	key <i>key-id</i>
Parameter Description	<i>key-id</i> : Indicates the authentication key ID in a key chain, ranging from 0 to 2,147,483,647.
Command Mode	Key chain configuration mode.
Usage Guide	-

Configuring a Key Authentication String

Command	key-string [0 7] <i>text</i>
Parameter Description	0 : Specifies that the key is displayed in plain text. 7 : Specifies that the key is displayed in cipher text. <i>text</i> : Specifies the authentication string characters.
Command Mode	Key chain key configuration mode.
Usage Guide	-

Configuring the Lifetime in the Sending Direction

Command	send-lifetime <i>start-time</i> {infinite <i>end-time</i> duration <i>seconds</i>}
Parameter Description	<i>start-time</i> : Indicates the start time of the lifetime. infinite : Indicates that the key is always effective. <i>end-time</i> : Indicates the end time of the lifetime, which must be later than <i>start-time</i> . duration <i>seconds</i> : Specifies the duration from the start time to the end time, ranging from 1 to 2,147,483,646.
Command Mode	Key chain key configuration mode.

Usage Guide	Run this command to define the lifetime of the key in the sending direction.
-------------	--

Configuring the Lifetime in the Receiving Direction

Command	<code>accept-lifetime <i>start-time</i> {infinite <i>end-time</i> duration <i>seconds</i>}</code>
Parameter Description	<p><i>start-time</i>: Indicates the start time of the lifetime.</p> <p><i>infinite</i>: Indicates that the key is always effective.</p> <p><i>end-time</i>: Indicates the end time of the lifetime, which must be later than <i>start-time</i>.</p> <p>duration <i>seconds</i>: Specifies the duration from the start time to the end time, ranging from 1 to 2,147,483,646.</p>
Command Mode	Key chain key configuration mode.
Usage Guide	Run this command to define the lifetime of the key in the receiving direction.

Configuration Example

Configuring a Key Chain and Using the Key Chain in RIP Packet Authentication

Scenario Figure 2-2	
Configuration Steps	<ul style="list-style-type: none"> Configure a key on all routers. Configure RIP on all routers. Enable RIP authentication on all routers.
A	<pre>A>enable A#configure terminal A(config)#key chain ripchain</pre>


```
A(config-keychain)#key 1
A(config-keychain-key)#key-string Hello
A(config-keychain-key)#accept-lifetime 16:30:00 Oct 1 2013 duration 43200
A(config-keychain-key)#send-lifetime 16:30:00 Oct 1 2013 duration 43200
A(config-keychain-key)#exit
A(config-keychain)#key 2
A(config-keychain-key)#key-string World
A(config-keychain-key)#accept-lifetime 04:00:00 Oct 2 2013 infinite
A(config-keychain-key)#send-lifetime 04:00:00 Oct 2 2013 infinite
A(config-keychain-key)#exit
A(config)#interface gigabitEthernet 0/1
A(config-if)#ip address 192.168.27.1 255.255.255.0
A(config-if)#ip rip authentication key-chain ripchain
A(config-if)#ip rip authentication mode md5
A(config-if)#exit
A(config)#router rip
A(config-router)#version 2
A(config-router)#network 192.168.27.0
```

B

```
B>enable
B#configure terminal
B(config)#key chain ripchain
B(config-keychain)#key 1
B(config-keychain-key)#key-string Hello
B(config-keychain-key)#accept-lifetime 16:30:00 Oct 1 2013 duration 43200
B(config-keychain-key)#send-lifetime 16:30:00 Oct 1 2013 duration 43200
B(config-keychain-key)#exit
B(config-keychain)#key 2
B(config-keychain-key)#key-string World
B(config-keychain-key)#accept-lifetime 04:00:00 Oct 2 2013 infinite
B(config-keychain-key)#send-lifetime 04:00:00 Oct 2 2013 infinite
B(config-keychain-key)#exit
B(config)#interface gigabitEthernet 0/1
```

	<pre> B(config-if)#ip address 192.168.27.2 255.255.255.0 B(config-if)#ip rip authentication key-chain ripchain B(config-if)#ip rip authentication mode md5 B(config-if)#exit B(config)#router rip B(config-router)#version 2 B(config-router)#network 192.168.27.0 B(config-router)#redistribute static </pre>
Verification	Run the show ip route rip command to check whether router A can receive an RIP route from router B.
A	<pre> A(config)#show ip route rip R 172.168.0.0/16 [120/1] via 192.168.27.2, 00:05:16, GigabitEthernet 0/1 </pre>

Common Errors

- A key is not correctly associated with a routing protocol, which causes that authentication does not take effect.
- The keys configured on multiple routers are not consistent, which causes authentication failure.

2.5 Monitoring

Displaying

Description	Command
Displays the configurations of a key chain.	show key chain [<i>key-chain-name</i>]

3 CONFIGURING RIP

3.1 Overview

Routing Information Protocol (RIP) is a unicast routing protocol applied on IPv4 networks. RIP-enabled routers exchange routing information to obtain routes to remote networks.

As an Interior Gateway Protocol (IGP), RIP can run only within the autonomous system (AS) and is applicable to small-sized networks whose longest path involves less than 16 hops.

Protocols and Standards

- RFC1058: Defines RIPv1.
- RFC2453: Defines RIPv2.

3.2 Applications

Application	Description
Basic RIP Application	The routing information is automatically maintained through RIP on a small-sized network.
Interworking Between RIP and BGP	Several ASs are interconnected. RIP runs within each AS, and Border Gateway Protocol (BGP) runs between ASs.


3.2.1 Basic RIP Application

Scenario

On a network with a simple structure, you can configure RIP to implement network interworking. Configuring RIP is simpler than configuring other IGP protocols like Open Shortest Path First (OSPF). Compared with static routes, RIP can dynamically adapt to the network structure changes and is easier to maintain.

As shown in Figure 3-1, to implement interworking between PC1, PC2, and PC3, you can configure RIP routes on R1, R2, and R3.

Figure 3-1



Deployment

- Configure IP addresses and gateways on three PCs.
- Configure IP addresses and subnet masks on three routers.
- Configure RIP on three routers.


3.2.2 Interworking Between RIP and BGP

Scenario

Several ASs are interconnected. RIP runs within each AS, and BGP runs between ASs. Generally, RIP and BGP learn the routing information from each other.

As shown in Figure 3-2, unicast routing is implemented within AS 100 and AS 200 using RIP, and between the two ASs using BGP.

Figure 3-2 Interworking Between RIP and BGP



Remarks	RIP and BGP run concurrently on Router A and Router D.
----------------	---

Deployment

- RIP runs within AS 100 and AS 200 to implement unicast routing.
- BGP runs between the two ASs to implement unicast routing.

3.3 Features

Basic Concepts

IGP and EGP

IGP runs within an AS. For example, RIP is a type of IGP.

Exterior Gateway Protocol (EGP) runs between ASs. For example, BGP is a type of EGP.

Classful Routing Protocol and Classless Routing Protocol

Protocols can be classified based on the type of routes supported:

- Classful routing protocol: It supports classful routes. For example, RIPv1 is a classful routing protocol.
- Classless routing protocol: It supports classless routes. For example, RIPv2 is a classless routing protocol.

Overview

Feature	Description
---------	-------------

RIPv1 and RIPv2	RIP is available in two versions: RIPv1 and RIPv2.
Exchanging Routing Information	By exchanging routing information, RIP-enabled devices can automatically obtain routes to a remote network and update the routes in real time.
Routing Algorithm	RIP is a protocol based on the distance-vector algorithm. It uses the vector addition method to compute the routing information.
Avoiding Route Loops	RIP uses functions, such as split horizon and poison reverse, to avoid route loops.
Security Measures	RIP uses functions, such as authentication and source address verification, to ensure protocol security.
Reliability Measures	RIP uses functions, such as bidirectional forwarding detection (BFD) correlation, fast reroute, and graceful restart (GR), to enhance reliability of the protocol.
Multiple Instances	RIP supports multiple instances and VPN applications.

3.3.1 RIPv1 and RIPv2

Two RIP versions are available: RIPv1 and RIPv2.

[Working Principle](#)

RIPv1

RIPv1 packets are broadcast. The broadcast address is 255.255.255.255, and the UDP port ID is 520. RIPv1 cannot identify the subnet mask, and supports only classful routes.

RIPv2

RIPv2 packets are multicast. The multicast address is 224.0.0.9, and the UDP port ID is 520. RIPv2 can identify the subnet mask, and supports classless routes, summarized route, and supernetting routes. RIPv2 supports plain text authentication and message digest 5 (MD5) authentication.

[Related Configuration](#)

Enabling the RIP Process

The RIP process is disabled by default.

Run the **router rip** command to enable the RIP process.

You must enable the RIP process on a device; otherwise, all functions related to RIP cannot take effect.

Running RIP on an Interface

By default, RIP does not run on an interface.

Run the **network** command to define an address range. RIP runs on interfaces that belong to this address range.

After RIP runs on an interface, RIP packets can be exchanged on the interface and RIP can learn routes to the network segments directly connected to the device.

Defining the RIP Version

By default, an interface receives RIPv1 and RIPv2 packets, and sends RIPv1 packets.

Run the **version** command to define the version of RIP packets sent or received on all interfaces.

Run the **ip rip send version** command to define the version of RIP packets sent on an interface.

Run the **ip rip receive version** command to define the version of RIP packets received on an interface.

 If the versions of RIP running on adjacent routers are different, the RIPv1-enabled router will learn incorrect routes.

Preventing an Interface from Sending or Receiving Packets

By default, a RIP-enabled interface is allowed to send and receive RIP packets.

Run the **no ip rip receive enable** command to prevent an interface from receiving RIP packets.

Run the **no ip rip send enable** command to prevent an interface from sending RIP packets.

Run the **passive-interface** command to prevent an interface from sending broadcast or multicast RIP packets.

Configuring the Mode for Sending RIP Packets

By default, broadcast RIPv1 packets and multicast RIPv2 are sent.

Run the **ip rip v2-broadcast** command to send broadcast RIPv2 packets on an interface.

Run the **neighbor** command to send unicast RIP packets to a specified neighbor router.

3.3.2 Exchanging Routing Information

Compared with static routing, the dynamic routing protocol has a significant advantage, that is, by exchanging routing information, devices can automatically obtain routes to a remote network and update the routes in real time.

Working Principle

Initialization

After RIP is enabled on a router, the router sends a request packet to its neighbor router, requesting for all routing information, that is, the routing table. After receiving the request message, the neighbor router returns a response packet containing the local routing table. After receiving the response packet, the router updates the local routing table, and sends an update packet to the neighbor router, informing the neighbor router of the route update information. After receiving the update packet, the neighbor router updates the local routing table, and sends the update packet to other adjacent routers. After a series of updates, all routers can obtain and retain the latest routing information.

Periodical Update

By default, periodical update is enabled for RIP. Adjacent routers exchange complete routing information with each other every 30s (update timer), that is, the entire routing table is sent to neighbor routers. One update packet contains at most 25 routes. Therefore, a lot of update packets may be required to send the entire routing table. You can set the sending delay between update packets to avoid loss of routing information.

-
- i** For every non-local route, if the route is not updated within 180s (invalid timer), the metric of the route is changed to 16 (unreachable). If the route is still not updated in the next 120s (flush timer), the route is deleted from the routing table.
-

Triggered Updates

After the triggered updates function is enabled, periodical update is automatically disabled. When routing information changes on a router, the router immediately sends routes related to the change (instead of the complete routing table) to the neighbor router, and use the acknowledgment and retransmission mechanisms to ensure that the neighbor router receives the routes successfully. Compared with periodical update, triggered updates help reduce flooding and accelerates route convergence.

Events that can trigger update include router startup, interface status change, changes in routing information (such as the metric), and reception of a request packet.

Route Summarization

When sending routing information to a neighbor router, the RIP-enabled router summarizes subnet routes that belong to the same classful network into a route, and sends the route to the neighbor router. For example, summarize 80.1.1.0/24 (metric=2) and 80.1.2.0/24 (metric=3) into 80.0.0.0/8 (metric=2), and set the metric of the summarized route to the optimum metric.

Only RIPv2 supports route summarization. Route summarization can reduce the size of the routing table and improve the efficiency of routing information exchange.

Supernetting Route

If the subnet mask length of a route is smaller than the natural mask length, this route is called supernetting route. For example, in the 80.0.0.0/6 route, as 80.0.0.0 is a Class A network address and the natural mask is 8 bits, 80.0.0.0/6 route is a supernetting route.

Only RIPv2 supports supernetting routes.

Default Route

In the routing table, a route to the destination network 0.0.0.0/0 is called default route.

The default route can be learned from a neighbor router, or sent to a neighbor router.

Route Redistribution

For RIP, other types of routes (such as direct routes, static routes, and routes of other routing protocols) are called external routes.

External routes (excluding the default route) can be redistributed to RIP and advertised to neighbors.

Route Filtering

Filtering conditions can be configured to limit the routing information exchanged between adjacent routers. Only the routing information that meets filtering conditions can be sent or received.

Related Configuration

Sending Delay Between Update Packets

By default, the update packets are sent continuously without any delay.

Run the **output-delay** command to set the sending delay between update packets.

RIP Timers

By default, the update timer is 30s, the invalid timer is 180s, and the flush timer is 120s.

Run the **timers basic** command to modify durations of the RIP timers.

Increasing the duration of the flush timer can reduce the route flapping. Decreasing the duration of the flush timer helps accelerate route convergence.

The durations of RIP timers must be consistent on adjacent routers. Unless otherwise required, you are advised not to modify the RIP timers.

Triggered Updates

By default, periodical update is enabled.

Run the **ip rip triggered** command to enable triggered updates on the interface and disable periodical update.

Run the **ip rip triggered retransmit-timer** command to modify the retransmission interval of update packets. The default value is 5s.

Run the **ip rip triggered retransmit-count** command to modify the maximum retransmission times of update packets. The default value is 36.

Route Summarization

By default, route summarization is automatically enabled if an interface is allowed to send RIPv2 packets.

Run the **no auto-summary** command to disable route summarization.

Run the **ip rip summary-address** command to configure route summarization on an interface.

Supernetting Route

By default, supernetting routes can be sent if an interface is allowed to send RIPv2 packets.

Run the **no ip rip send supernet-routes** command to prevent the sending of supernetting routes.

Default Route

Run the **ip rip default-information** command to advertise the default route to neighbors on an interface.

Run the **default-information originate** command to advertise the default route to neighbors from all interfaces.

Route Redistribution

Run the **redistribute** command to redistribute external routes (excluding the default route) to RIP and advertise them to neighbors.

Route Filtering

Run the **distribute-list out** command to set filtering rules to limit the routing information sent by the device.

Run the **distribute-list in** command to set filtering rules to limit the routing information received by the device.

3.3.3 Routing Algorithm

RIP is a protocol based on the distance-vector algorithm. It uses the vector addition method to compute the routing information.

Working Principle


Distance-Vector Algorithm

RIP is a protocol based on the distance-vector algorithm. The distance-vector algorithm treats a route as a vector that consists of the destination network and distance (metric). The router obtains a route from its neighbor and adds the distance vector from itself to the neighbor to the route to form its own route.

RIP uses the hop count to evaluate the distance (metric) to the destination network. By default, the hop count from a router to its directly connected network is 0, the hop count from a router to a network that can be reached through the router is 1, and so on. That is, the metric is equal to the number of routers from the local network to the destination network. To restrict the convergence time, RIP stipulates that the metric must be an integer between 0 and 15. If the metric is equal to or greater than 16, the destination network or host is unreachable. For this reason, RIP cannot be applied on a large-scale network.

As shown in Figure 3-3, Router A is connected to the network 10.0.0.0. Router B obtains the route (10.0.0.0,0) from Router A and adds the metric 1 to the route to obtain its own route ((10.0.0.0,1), and the next hop points to Router A.

Figure 3-3




Selecting the Optimum Route

RIP selects an optimum route based on the following principle: If multiple routes to the same destination network is available, a router preferentially selects the route with the smallest metric.

As shown in Figure 3-4, Router A is connected to the network 10.0.0.0. Router C obtains the route (10.0.0.0,0) from Router A and the route (10.0.0.0,1) from Router B. Router C will select the route that is obtained from Router A and add metric 1 to this route to form its own route (10.0.0.0,1), and the next hop points to Router A.

Figure 3-4



i When routes coming from different sources exist on a router, the route with the smallest distance is preferentially selected.

Route Source	Default Distance
Directly-connected network	0
Static route	1
OSPF route	110
IS-IS route	115
RIP route	120
Unreachable route	255

[Related Configuration](#)

Modifying the Distance

By default, the distance of a RIP route is 120.

Run the **distance** command to modify the distance of a RIP route.

Modifying the Metric

For a RIP route that is proactively discovered by a device, the default metric is equal to the number of hops from the local network to the destination network. For a RIP router that is manually configured (default route or redistributed route), the default metric is 1.

Run the **offset-list in** command to increase the metric of a received RIP route.

Run the **offset-list out** command to increase the metric of a sent RIP route.

Run the **default-metric** command to modify the default metric of a redistributed route.

Run the **redistribute** command to modify the metric of a route when the route is redistributed.

Run the **default-information originate** command to modify the metric of a default route when the default route is introduced.

Run the **ip rip default-information** command to modify the metric of a default route when the default route is created.

3.3.4 Avoiding Route Loops

RIP uses functions, such as split horizon and poison reverse, to avoid route loops.

Working Principle

Route Loop

A RIP route loop occurs due to inherent defects of the distance-vector algorithm.

As shown in Figure 3-5, Router A is connected to the network 10.0.0.0, and sends an update packet every 30s. Router B receives the route 10.0.0.0 from Router A every 30s. If Router A is disconnected from 10.0.0.0, the route to 10.0.0.0 will be deleted from the routing table on Router A. Next time, the update packet sent by Router A no longer contains this route. As Router B does not receive an update packet related to 10.0.0.0, Router B determines that the route to 10.0.0.0 is valid within 180s and uses the Update packet to send this route to Router A. As the route to 10.0.0.0 does not exist on Router A, the route learned from Router B is added to the routing table. Router B determines that data can reach 10.0.0.0 through Router A, and Router A determines that data can reach 10.0.0.0 through Router B. In this way, a route loop is formed.

Figure 3-5



Split Horizon

Split horizon can prevent route loops. After split horizon is enabled on an interface, a route received on this interface will not be sent out from this interface.

As shown in Figure 3-6, after split horizon is enabled on the interface between Router A and Router B, Router B will not send the route 10.0.0.0 back to Router A. Router B will learn 180s later that 10.0.0.0 is not reachable.

Figure 3-6




Poison Reverse

Poison reverse can also prevent route loops. Compared with split horizon, poison reverse is more reliable, but brings more protocol packets, which makes network congestion more severe.

After poison reverse is enabled on an interface, a route received from this interface will be sent out from this interface again, but the metric of this router will be changed to 16 (unreachable).

As shown in Figure 3-7, after learning the route 10.0.0.0 from Router A, Router B sets the metric of this route to 16 and sends the route back to Router A. After this route becomes invalid, Router B advertises the route 10.0.0.0 (metric = 16) to Router A to accelerate the process of deleting the route from the routing table.

Figure 3-7



Related Configuration

Split Horizon

By default, split horizon is enabled.

Run the **no ip rip split-horizon** command to disable split horizon.

Poison Reverse

By default, poison reverse is disabled.

Run the **ip rip split-horizon poisoned-reverse** command to enable poison reverse. (After poison reverse is enabled, split horizon is automatically disabled.)

3.3.5 Security Measures

RIP uses functions, such as authentication and source address verification, to ensure protocol security.

Working Principle

Authentication

RIPv2 supports authentication, but RIPv1 does not.

After authentication is enabled on an interface, the routing information cannot be exchanged between adjacent devices if authentication fails. The authentication function is used to prevent unauthorized devices from accessing the RIP routing domain.

RIPv2 supports plain text authentication and MD5 authentication.

Source Address Verification

When a RIP-enabled device receives an Update packet, it checks whether the source IP address in the packet and the IP address of the inbound interface are in the same network segment. If not, the device drops the packet. Source address verification is used to ensure that RIP routing information is exchanged only between adjacent routing devices.

- i** On an unnumbered IP interface, source address verification is not performed (not configurable).
- i** If the triggered updates function is enabled, source address verification is automatically enabled (not configurable).
- i** If split horizon is disabled, source address verification is automatically enabled (not configurable).

Related Configuration

Authentication

By default, authentication is disabled.

Run the **ip rip authentication mode text** command to enable plain text authentication on an interface.

Run the **ip rip authentication mode md5** command to enable MD5 authentication on an interface.

Run the **ip rip authentication text-password** command to set the password for plain text authentication on an interface.

Run the **ip rip authentication key-chain** command to reference the key in the configured key chain as the authentication key on an interface.

Source Address Verification

By default, source address verification is enabled.

Run the **no validate-update-source** command to disable source address verification.

3.3.6 Reliability Measures

RIP uses functions, such as BFD correlation, fast reroute, and GR to enhance reliability of the protocol.


Working Principle

BFD Correlation and Fast Reroute

When a link or a device is faulty on the network, packets transmitted through this route will be lost until the route is converged again.

As shown in Figure 3-8, after the link between Router A and Router S is faulty, Router B may wait 180s before it can detect the failure of the route (Destination network: 10.0.0.0; Next hop: Router A). Later, Router B may need to wait 30s to re-obtain the route (Destination network: 10.0.0.0; Next hop: Router C) from Router C. Therefore, the traffic is interrupted for 210s.

Figure 3-8



Quick detection of a route failure or fast switchover to the standby route helps shorten the traffic interruption time.

- A BFD session can be set up between Router A and Router B, and correlated with RIP. BFD can quickly test the connectivity between adjacent routers. Once a link is faulty, RIP can detect the route failure within 1s.
- The fast reroute function can be enabled. A standby route (Destination network: 10.0.0.0; Next hop: Router C) can be configured on Router B in advance. Once RIP detects a route failure, the standby route is immediately enabled.
- GR

GR ensures uninterrupted data transmission when the protocol is restarted. If RIP is restarted on a GR-enabled device, the forwarding table before restart will be retained and a request packet will be sent to the neighbor so that the route can be learned again. During the GR period, RIP completes re-convergence of the route. After the GR period expires, RIP updates the forwarding entry and advertises the routing table to the neighbor.

Related Configuration

BFD Correlation

By default, RIP is not correlated with BFD.

Run the **bfd all-interfaces** command to set up the correlation between RIP and BFD. This configuration takes effect on all interfaces.

Run the **ip rip bfd** command to set up the correlation between RIP and BFD on the current interface.

Fast Reroute

By default, fast reroute is disabled.

Run the **fast-reroute route-map** command to enable fast reroute and reference the route map.

Run the **set fast-reroute backup-interface backup-nexthop** command to configure a standby route in the route map.

3.3.7 Multiple Instances

[Working Principle](#)

Multiple VPN instances may exist on a device.

RIP supports multiple instances. You can enable the RIP process in VPN routing and forwarding (VRF) address family mode to run RIP on VPN instances. One VRF address family is mapped to one VPN instance.

VPN instances cannot be distinguished from each other when you perform RIP operations using SNMP. You must bind the management information base (MIB) of RIP with a VPN instance before the SNMP operations take effect on the VPN instance.

[Related Configuration](#)

VRF Address Family

By default, the RIP process runs on a public network instance.

Run the **address-family** command to create a VRF address family and enter VRF address family mode.

Run the **exit-address-family** command to exit from VRF address family mode.


Run the **no address-family** command to delete a VRF address family.

MIB Binding

By default, the RIP MIB is bound with a public network instance.

Run the **enable mib-binding** command to bind the RIP MIB with a VPN instance.

3.4 Configuration

Configuration	Description and Command
Configuring RIP Basic Functions	 (Mandatory) It is used to build a RIP routing domain.
	router rip Enables a RIP routing process and enters routing process configuration mode.
	network Runs RIP on interfaces in the specified address range.
	version Defines the RIP version.

	ip rip split-horizon	Enables split horizon or poison reverse on an interface.
	passive-interface	Configures a passive interface.
Controlling Interaction of RIP Packets	⚠ (Optional) This configuration is required if you wish to change the default mechanism for sending or receiving RIP packets.	
	neighbor	Sends unicast RIP packets to a specified neighbor.
	ip rip v2-broadcast	Sends broadcast RIPv2 packets on an interface.
	ip rip receive enable	Allows the interface to receive RIP packets.
	ip rip send enable	Allows the interface to send RIP packets.
	ip rip send version	Defines the version of RIP packets sent on an interface.
	ip rip receive version	Defines the version of RIP packets received on an interface.
Enabling Triggered Updates	⚠ Optional.	
	ip rip triggered	Enables triggered updates on an interface.
Enabling Source Address Verification	⚠ Optional.	
	validate-update-source	Enables source address verification.
Enabling Authentication	⚠ (Optional) Only RIPv2 supports authentication.	
	ip rip authentication mode	Enables authentication and sets the authentication mode on an interface.

	ip rip authentication text-password	Configures the password for plain text authentication on an interface.
	ip rip authentication key-chain	Configures the authentication key chain on an interface.
Enabling Route Summarization	⚠ (Optional) Only RIPv2 supports route summarization.	
	auto-summary	Enables automatic summarization of RIP routes.
	ip rip summary-address	Configures route summarization on an interface.
Enabling Supernetting Routes	⚠ (Optional) Only RIPv2 supports supernetting routes.	
	ip rip send supernet-routes	Enables advertisement of RIP supernetting routes on an interface
Advertising the Default Route or External Routes	⚠ Optional.	
	ip rip default-information	Advertises the default route to neighbors on an interface.
	default-information originate	Advertises the default route to neighbors.
	redistribute	Redistributes routes and advertises external routes to neighbors.
Setting Route Filtering Rules	⚠ Optional.	
	distribute-list in	Filters the received RIP routing information.
	distribute-list out	Filters the sent RIP routing information.
	⚠ Optional.	

Modifying Route Selection Parameters	distance	Modifies the administrative distance (AD) of a RIP route.
	offset-list	Increases the metric of a received or sent RIP route.
	default-metric	Configures the default metric of an external route redistributed to RIP.
Modifying Timers	⚠ Optional.	
	timers basic	Modifies the update timer, invalid timer, and flush timer.
	output-delay	Sets the sending delay between RIP route update packets.
Enabling BFD Correlation	⚠ Optional.	
	bfd all-interfaces	Correlates RIP with BFD on all interfaces.
	ip rip bfd	Correlates RIP with BFD on an interface.
Enabling Fast Reroute	⚠ Optional.	
	fast-reroute route-map	Enables fast reroute and references the route map.
	set fast-reroute backup-interface backup-nexthop	Configures the standby interface and standby next hop for fast reroute in the route map.
Enabling GR	⚠ Optional.	
	graceful-restart	Configures the GR restarter capability.
	⚠ (Optional) It is used to run RIP on VPN instances.	

Enabling Multiple Instances	address-family ipv4 vrf	Creates a VRF address family and enters IPv4 VRF address family mode.
	exit-address-family	Exits from an IPv4 VRF address family.
	enable mib-binding	Binds RIP MIB with a VPN instance.

3.4.1 Configuring RIP Basic Functions

Configuration

Effect

- Build a RIP routing domain on the network.
- Routers in the domain obtain routes to a remote network through RIP.

Notes

- IPv4 addresses must be configured.
- IPv4 unicast routes must be enabled.

Configuration

Steps

Enabling a RIP Routing Process

- Mandatory.
- Unless otherwise required, this configuration must be performed on every router in the RIP routing domain.

Associating with the Local Network

- Mandatory.
- Unless otherwise required, this configuration must be performed on every router in the RIP routing domain.
- Unless otherwise required, the local network associated with RIP should cover network segments of all L3 interfaces.

Defining the RIP Version

- If RIPv2 functions (such as the variable length subnet mask and authentication) are required, enable the RIPv2.
- Unless otherwise required, you must define the same RIP version on every router.

Enabling Split Horizon or Poison Reverse

- By default, split horizon is enabled and poison reverse is disabled.
- Unless otherwise required, enable split horizon on every interface connected to the broadcast network, such as the Ethernet. (Retain the default setting.)

- Unless otherwise required, enable split horizon on every interface connected to the point-to-point (P2P) network, such as the PPP and HDLC. (Retain the default setting.)
- It is recommended that split horizon and poison reverse be disabled on an interface connected to a non-broadcast multi-access (NBMA) network, such as FR and X.25; otherwise, some devices may fail to learn the complete routing information.
- If the secondary IP address is configured for an interface connected to a non-broadcast, it is recommended that split horizon and poison reverse be disabled.

Configuring a Passive Interface

- If you want to suppress Update packets on a RIP interface, configure the interface as a passive interface.
- Use the passive interface to set the boundary of the RIP routing domain. The network segment of the passive interface belongs to the RIP routing domain, but RIP packets cannot sent over the passive interface.
- If RIP routes need to be exchanged on an interface (such as the router interconnect interface) in the RIP routing domain, this interface cannot be configured as a passive interface.

Verification

- Check the routing table on a router to verify that the route to a remote network can be obtained through RIP.

Related Commands

Enabling a RIP Routing Process

Command Syntax	<code>router rip</code>
Parameter Description	N/A
Command Mode	Global configuration mode
Configuration Usage	This command is used to create a RIP routing process and enter routing process configuration mode.

Associating with the Local Network

Command Syntax	network <i>network-number</i> [<i>wildcard</i>]
Parameter Description	<i>network-number</i> : Indicates the number of a network. <i>wildcard</i> : Defines the IP address comparison bit. 0 indicates accurate matching, and 1 indicates that no comparison is performed.
Command Mode	Routing process configuration mode
Configuration Usage	RIP can run and learn direct routes and RIP packets can be exchanged only on an interface covered by network . If network 0.0.0.0 255.255.255.255 is configured, all interfaces are covered. If <i>wildcard</i> is not configured, the classful address range is used by default, that is, the interfaces whose addresses fall into the classful address range participate in RIP operations.

Defining the RIP Version

Command Syntax	version { 1 2 }
Parameter Description	1 : Indicates RIPv1. 2 : Indicates RIPv2.
Command Mode	Global configuration mode
Configuration Usage	This command takes effect on the entire router. You can run this command to define the version of RIP packets sent or received on all interfaces.

Enabling Split Horizon

Command Syntax	ip rip split-horizon [<i>poisoned-reverse</i>]
Parameter Description	poisoned-reverse : Indicates poison reverse.

Command Mode	Interface configuration mode
Configuration Usage	After poison reverse is enabled, split horizon is automatically disabled.

Configuring a Passive Interface

Command Syntax	passive-interface { default <i>interface-type interface-num</i> }
Parameter Description	default: Indicates all interfaces. <i>interface-type interface-num:</i> Specifies an interface.
Command Mode	Routing process configuration mode
Configuration Usage	First, run the passive-interface default command to configure all interfaces as passive interfaces. Then, run the no passive-interface <i>interface-type interface-num</i> command to cancel the interfaces used for interconnection between routers in the domain.

Configuration Example

Building a RIP Routing Domain

<p>Scenario Figure 3-9</p>	
<p>Remarks</p>	<p>The interface IP addresses are as follows: A: GE0/1 110.11.2.1/24 GE0/2 155.10.1.1/24</p>

	<p>B: GE0/1 110.11.2.2/24 GE0/2 196.38.165.1/24</p> <p>C: GE0/1 110.11.2.3/24 GE0/2 117.102.0.1/16</p>
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. Configure the RIP basic functions on all routers.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip address 110.11.2.1 255.255.255.0 A(config-if-GigabitEthernet 0/1)# exit A(config)# interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)# ip address 155.10.1.1 255.255.255.0 A(config)# router rip A(config-router)# version 2 A(config-router)# network 0.0.0.0 255.255.255.255 A(config-router)# passive-interface default A(config-router)# no passive-interface GigabitEthernet 0/1</pre>
B	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip address 110.11.2.2 255.255.255.0 B(config-if-GigabitEthernet 0/1)# exit B(config)# interface GigabitEthernet 0/2 B(config-if-GigabitEthernet 0/2)# ip address 196.38.165.1 255.255.255.0 B(config-if-GigabitEthernet 0/2)# exit B(config)# router rip B(config-router)# version 2 B(config-router)# network 0.0.0.0 255.255.255.255 B(config-router)# passive-interface default B(config-router)# no passive-interface GigabitEthernet 0/1</pre>
C	<pre>C# configure terminal C(config)# interface GigabitEthernet 0/1</pre>

	<pre> C(config-if-GigabitEthernet 0/1)# ip address 110.11.2.3 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# interface GigabitEthernet 0/2 C(config-if-GigabitEthernet 0/2)# ip address 117.102.0.1 255.255.0.0 C(config-if-GigabitEthernet 0/2)# exit C(config)# router rip C(config-router)# version 2 C(config-router)#no auto-summary C(config-router)# network 0.0.0.0 255.255.255.255 C(config-router)# passive-interface default C(config-router)# no passive-interface GigabitEthernet 0/1 </pre>
Verification	<p>Check the routing tables on Router A, Router B, and Router C. Verify that RIP learns the routes to remote networks (contents marked in blue).</p>
A	<pre> A# show ip route Codes: C - connected, S - static, R - RIP, B - BGP O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default Gateway of last resort is no set C 110.11.2.0/24 is directly connected, GigabitEthernet 0/1 C 110.11.2.1/32 is local host. R 117.0.0.0/8 [120/1] via 110.11.2.2, 00:00:47, GigabitEthernet 0/1 C 155.10.1.0/24 is directly connected, GigabitEthernet 0/2 C 155.10.1.1/32 is local host. C 192.168.217.0/24 is directly connected, VLAN 1 C 192.168.217.233/32 is local host. R 196.38.165.0/24 [120/1] via 110.11.2.3, 00:19:18, GigabitEthernet 0/1 </pre>

<p>B</p>	<pre> B# show ip route Codes: C - connected, S - static, R - RIP, B - BGP O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default Gateway of last resort is no set C 110.11.2.0/24 is directly connected, GigabitEthernet 0/1 C 110.11.2.2/32 is local host. R 155.10.0.0/16 [120/1] via 110.11.2.1, 00:15:21, GigabitEthernet 0/1 C 196.38.165.0/24 is directly connected, GigabitEthernet 0/2 C 196.38.165.1/32 is local host. R 117.0.0.0/8 [120/1] via 110.11.2.2, 00:00:47, GigabitEthernet 0/1 </pre>
<p>C</p>	<pre> C# show ip route Codes: C - connected, S - static, R - RIP, B - BGP O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default Gateway of last resort is no set C 110.11.2.0/24 is directly connected, GigabitEthernet 0/1 C 110.11.2.3/32 is local host. C 117.102.0.0/16 is directly connected, GigabitEthernet 0/2 C 117.102.0.1/32 is local host. R 155.10.0.0/16 [120/1] via 110.11.2.1, 00:20:55, GigabitEthernet 0/1 R 196.38.165.0/24 [120/1] via 110.11.2.3, 00:19:18, GigabitEthernet 0/1 </pre>

Common Errors

- The IPv4 address is not configured on an interface.
- The RIP version is not defined on a device, or the RIP version on the device is different from that on other routers.
- The address range configured by the **network** command does not cover a specific interface.
- The **wildcard** parameter in the **network** command is not correctly configured. **0** indicates accurate matching, and **1** indicates that no comparison is performed.
- The interface used for interconnection between devices is configured as a passive interface.

3.4.2 Controlling Interaction of RIP Packets

Configuration Effect

Change the default running mechanism of RIP through configuration and manually control the interaction mode of RIP packets, including:

- Allowing or prohibiting the sending of unicast RIP packets to a specified neighbor on an interface
- Allowing or prohibiting the sending of unicast RIPv2 packets instead of broadcast packets to a specified neighbor on an interface
- Allowing or prohibiting the receiving of RIP packets on an interface
- Allowing or prohibiting the sending of RIP packets on an interface
- Allowing or prohibiting the receiving of RIP packets of a specified version on an interface
- Allowing or prohibiting the sending of RIP packets of a specified version on an interface

Notes

- The RIP basic functions must be configured.
- On an interface connecting to a neighbor device, the configured version of sent RIP packets must be the same as the version of received RIP packets.

Configuration Steps

Sending Unicast RIP Route Update Packets to a Specified Neighbor

- Configure this function if you wish that only some of devices connected to an interface can receive the updated routing information.
- By default, RIPv1 uses the IP broadcast address (255.255.255.255) to advertise the routing information, whereas RIPv2 uses the multicast address (224.0.0.9) to advertise the routing information. If you do not wish all devices on the broadcast network or NBMA network to receive routing information, configure the related interface as the passive interface and specify the

neighbors that can receive the routing information. This command does not affect the receiving of RIP packets. RIPv2 packets are broadcast on an interface.

- Unless otherwise required, this function must be enabled on a router that sends the unicast Update packets.

Broadcasting RIPv2 Packets on an Interface

- This function must be configured if the neighbor router does not support the receiving of multicast RIPv2 packets.
- Unless otherwise required, this function must be configured on every router interface that broadcasts RIPv2 packets.

Allowing an Interface to Receive RIP Packets

- This function is enabled by default, and must be disabled if an interface is not allowed to receive RIP packets.
- Unless otherwise required, this function must be configured on every router interface that is not allowed to receive RIP packets.

Allowing an Interface to Send RIP Packets

- This function is enabled by default, and must be disabled if an interface is not allowed to send RIP packets.
- Unless otherwise required, this function must be configured on every router interface that is not allowed to send RIP packets.

Allowing an Interface to Send RIP Packets of a Specified Version

- This function must be configured if the version of RIP packets that can be sent on an interface is required to be different from the global configuration.
- Unless otherwise required, this function must be configured on every router interface that is allowed to send RIP packets of a specified version.

Allowing an Interface to Receive RIP Packets of a Specified Version

- This function must be configured if the version of RIP packets that can be received on an interface is required to be different from the global configuration.
- Unless otherwise required, this function must be configured on every router interface that is allowed to receive RIP packets of a specified version.

Verification

Run the **debug ip rip packet** command to verify the packet sending result and packet type.

Related

Commands

Sending Unicast RIP Route Update Packets to a Specified Neighbor

Command Syntax	neighbor <i>ip-address</i>
Parameter Description	<i>ip-address</i> : Indicates the IP address of the neighbor. It should be the address of the network directly connected to the local device.
Command Mode	Routing process configuration mode
Configurati on Usage	Generally, you can first run the passive-interface command in routing process configuration mode to configure the related interface as a passive interface, and then specify the neighbors that can receive the routing information. This command does not affect the receiving of RIP packets. After an interface is configured as a passive interface, the interface does not send the request packets even after the device is restarted.

Broadcasting RIPv2 Packets on an Interface

Command Syntax	ip rip v2-broadcast
Parameter Description	N/A
Command Mode	Interface configuration mode
Configurati on Usage	The default behavior is determined by the configuration of the version command. The configuration result of this command can overwrite the default configuration of the version command. This command affects the behavior of sending RIP packets on the current interface, and the interface is allowed to send RIPv1 and RIPv2 packets simultaneously. If this command does not contain any parameter, the behavior of receiving RIP packets is determined by the configuration of the version command.

Allowing an Interface to Receive RIP Packets

Command Syntax	ip rip receive enable
Parameter Description	N/A

Command Mode	Interface configuration mode
Configuration Usage	To prohibit the receiving of RIP packets on an interface, use the no form of this command. This command takes effect only on the current interface. You can use the default form of the command to restore the default setting, that is, allowing the interface to receive RIP packets.

Allowing an Interface to Send RIP Packets

Command Syntax	ip rip send enable
Parameter Description	N/A
Command Mode	Interface configuration mode
Configuration Usage	To prohibit the sending of RIP packets on an interface, use the no form of this command in interface configuration mode. This command takes effect only on the current interface. You can use the default form of the command to restore the default setting, that is, allowing the interface to send RIP packets.

Allowing an Interface to Send RIP Packets of a Specified Version

Command Syntax	ip rip send version [1] [2]
Parameter Description	1: Indicates that only RIPv1 packets are sent. 2: Indicates that only RIPv2 packets are sent.
Command Mode	Interface configuration mode
Configuration Usage	The default behavior is determined by the configuration of the version command. The configuration result of this command can overwrite the default configuration of the version command. This command affects the behavior of sending RIP packets on the current interface, and the interface is allowed to send RIPv1 and RIPv2 packets simultaneously. If this command does not contain any parameter, the behavior of receiving RIP packets is determined by the configuration of the version command.

Allowing an Interface to Receive RIP Packets of a Specified Version

Command Syntax	ip rip receive version [1] [2]
Parameter Description	1: Indicates that only RIPv1 packets are received. 2: Indicates that only RIPv2 packets are received.
Command Mode	Interface configuration mode
Configuration Usage	The default behavior is determined by the configuration of the version command. The configuration result of this command can overwrite the default configuration of the version command. This command affects the behavior of receiving RIP packets on the current interface, and the interface is allowed to receive RIPv1 and RIPv2 packets simultaneously. If this command does not contain any parameter, the behavior of receiving RIP packets is determined by the configuration of the version command.

Configuration Example

Prohibiting an Interface from Sending RIP Packets

Scenario Figure 3-10	
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Prohibit the sending of RIP packets on an interface of Router A.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 QTECH(config-if-GigabitEthernet 0/1)# no ip rip send enable</pre>

Verification	Run the debug ip rip packet send command on Router A, and verify that packets cannot be sent.
A	<pre>A# debug ip rip packet recv *Nov 4 08:19:31: %RIP-7-DEBUG: [RIP] Prepare to send BROADCAST response... *Nov 4 08:19:31: %RIP-7-DEBUG: [RIP] Building update entries on GigabitEthernet 0/1 *Nov 4 08:19:31: %RIP-7-DEBUG: 117.0.0.0/8 via 0.0.0.0 metric 1 tag 0 *Nov 4 08:19:31: %RIP-7-DEBUG: [RIP] Interface GigabitEthernet 0/1 is disabled to send RIP packet!</pre>

Common Errors

A compatibility error occurs because the RIP version configured on the neighbor is different from that configured on the local device.

3.4.3 Enabling Triggered Updates

Configuration Effect

- Enable the RIP triggered updates function, after which RIP does not periodically send the route update packets.

Notes

- The RIP basic functions must be configured.
- It is recommended that split horizon with poisoned reverse be enabled; otherwise, invalid routing information may exist.
- This function cannot be enabled together with the function of correlating RIP with BFD.
- Ensure that the triggered updates function is enabled on every router on the same link; otherwise, the routing information cannot be exchanged properly.

Configuration Steps

Enabling Triggered Updates

- This function must be enabled if demand circuits are configured on the WAN interface.
- The triggered updates function can be enabled in either of the following cases: (1) The interface has only one neighbor; (2) The interface has multiple neighbors but the device interacts with these neighbors in unicast mode.

- It is recommended that triggered updates be enabled on a WAN interface (running the PPP, Frame Relay, or X.25 link layer protocol) to meet the requirements of demand circuits.
- If the triggered updates function is enabled on an interface, source address verification is performed no matter whether the source address verification function is enabled by the **validate-update-source** command.
- Unless otherwise required, triggered updates must be enabled on demand circuits of every router.

Verification

When the RIP triggered updates function is enabled, RIP cannot periodically send the route update packets. RIP sends the route update packets to the WAN interface only in one of the following cases:

- A route request packet is received.
- The RIP routing information changes.
- The interface state changes.
- The router is started.

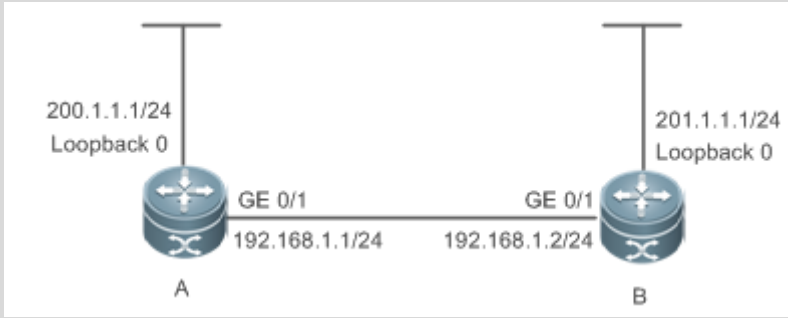
Related Commands

Enabling Triggered Updates

Command Syntax	ip rip triggered { retransmit-timer <i>timer</i> / retransmit-count <i>count</i> }
Parameter Description	<p>retransmit-timer <i>timer</i>: Configures the interval at which the update request or update response packet is retransmitted. The default value is 5s. The value ranges from 1 to 3,600.</p> <p>retransmit-count <i>count</i>: Configures the maximum retransmission times of the update request or update response packet. The default value is 36. The value ranges from 1 to 3,600.</p>
Command Mode	Interface configuration mode
Configuration Usage	<p>You can run the ip rip triggered command to enable the RIP triggering function.</p> <p>When this function is enabled, the RIP periodical update function is automatically disabled. Therefore, the acknowledgment and retransmission mechanisms must be used to ensure that the Update packets are successfully sent or received on the WAN.</p> <p>You can use the retransmit-timer and retransmit-count parameters to specify the retransmission interval and maximum retransmission times of the request and update packets.</p>

Configuration Example

Enabling Triggered Updates

Scenario Figure 3-11	 <p>The diagram illustrates a network topology with two routers, A and B, connected via their GigabitEthernet 0/1 interfaces. Router A has a loopback interface 200.1.1.1/24 and a physical interface 192.168.1.1/24. Router B has a loopback interface 201.1.1.1/24 and a physical interface 192.168.1.2/24.</p>
Configuration Steps	<ul style="list-style-type: none">▪ Configure the interface IP addresses on all routers. (Omitted)▪ Configure the RIP basic functions on all routers. (Omitted)▪ On Router A, enable the RIP triggered updates function, and set the retransmission interval and maximum retransmission times of the request and update packets to 10s and 18, respectively.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# encapsulation ppp A(config-if-GigabitEthernet 0/1)# ip rip triggered A(config-if-GigabitEthernet 0/1)# ip rip triggered retransmit-timer 10 A(config-if-GigabitEthernet 0/1)# ip rip triggered retransmit-count 18 A(config-if-GigabitEthernet 0/1)# ip rip split-horizon poisoned-reverse A(config)# router rip A(config-router)# network 192.168.1.0 A(config-router)# network 200.1.1.0</pre>
B	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# encapsulation ppp B(config-if-GigabitEthernet 0/1)# ip rip triggered B(config-if-GigabitEthernet 0/1)# ip rip split-horizon poisoned-reverse B(config)# router rip B(config-router)# network 192.168.1.0</pre>

	<pre>B(config-router)# network 201.1.1.0</pre>
Verification	On Router A and Router B, check the RIP database and verify that the corresponding routes are permanent.
A	<pre>A# sho ip rip database 201.1.1.0/24 auto-summary 201.1.1.0/24 [1] via 192.168.12.2 GigabitEthernet 0/1 06:25 permanent</pre>
B	<pre>B# sho ip rip database 200.1.1.0/24 auto-summary 200.1.1.0/24 [1] via 192.168.12.1 GigabitEthernet 0/1 06:25 permanent</pre>

Common Errors

- The triggered updates function is enabled when the RIP configurations at both ends of the link are consistent.
- Both the triggered updates and BFD functions are enabled.
- The triggered updates function is not enabled on all routers on the same link.

3.4.4 Enabling Source Address Verification

Configuration Effect

- The source address of the received RIP route update packet is verified.

Notes

- The RIP basic functions must be configured.

Configuration Steps

Enabling Source Address Verification

- This function is enabled by default, and must be disabled when source address verification is not required.

- After split horizon is disabled on an interface, the RIP routing process will perform source address verification on the Update packet no matter whether the **validate-update-source** command is executed in routing process configuration mode.
- For an IP unnumbered interface, the RIP routing process does not perform source address verification on the Update packet no matter whether the **validate-update-source** command is executed in routing process configuration mode.
- Unless otherwise required, this function must be disabled on every router that does not requires source address verification.

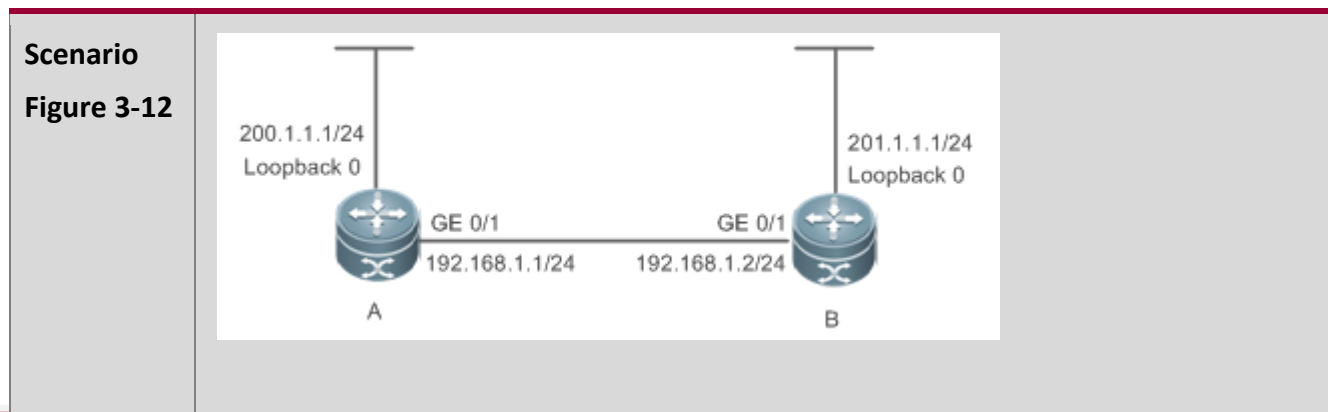
Verification

Only the route update packets coming from the same IP subnet neighbor are received.

Related Commands

Command Syntax	validate-update-source
Parameter Description	N/A
Command Mode	Routing process configuration mode
Configurati on Usage	Source address verification of the Update packet is enabled by default. After this function is enabled, the source address of the RIP route update packet is verified. The purpose is to ensure that the RIP routing process receives only the route update packets coming from the same IP subnet neighbor.

Configuration Example



Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Disable source address verification of Update packets on all routers.
A	<pre>A# configure terminal A(config)# router rip A(config-router)# no validate-update-source</pre>
B	<pre>B# configure terminal B(config)# router rip B(config-router)# no validate-update-source</pre>
Verification	<ul style="list-style-type: none"> On Router A, check the routing table and verify that the entry 201.1.1.0/24 is loaded. On Router B, check the routing table and verify that the entry 200.1.1.0/24 is loaded.
A	<pre>A# show ip route rip R 201.1.1.0/24 [120/1] via 192.168.2.2, 00:06:11, GigabitEthernet 0/1</pre>
B	<pre>B# show ip route rip R 200.1.1.0/24 [120/1] via 192.168.1.1, 00:06:11, GigabitEthernet 0/1</pre>

3.4.5 Enabling Authentication

Configuration

Effect

- Prevent learning unauthenticated and invalid routes and advertising valid routes to unauthorized devices, ensuring stability of the system and protecting the system against intrusions.

Notes

- The RIP basic functions must be configured.
- Only RIPv2 supports authentication of RIP packets, and RIPv1 does not.

Configuration

Steps

Enabling Authentication and Specifying the Key Chain Used for RIP Authentication

- This configuration is mandatory if authentication must be enabled.

- If the key chain is already specified in the interface configuration, run the **key chain** command in global configuration mode to define the key chain; otherwise, authentication of RIP packets may fail.
- Unless otherwise required, this configuration must be performed on every router that requires authentication.

Defining the RIP Authentication Mode

- This configuration is mandatory if authentication must be enabled.
- The RIP authentication modes configured on all devices that need to directly exchange RIP routing information must be the same; otherwise, RIP packets may fail to be exchanged.
- If plain text authentication is used, but the key chain for plain text authentication is not configured or associated, authentication is not performed. Similarly, if MD5 authentication is used, but the key chain is not configured or associated, authentication is not performed.
- Unless otherwise required, this configuration must be performed on every router that requires authentication.

Enabling RIP Plain Text Authentication and Configuring the Key Chain

- This configuration is mandatory if authentication must be enabled.
- If RIP plain text authentication should be enabled, use this command to configure the key chain for plain text authentication. Alternatively, you can obtain the key chain for plain text authentication by associating the key chain. The key chain obtained using the second method takes precedence over that obtained using the first method.
- Unless otherwise required, this configuration must be performed on every router that requires authentication.

Verification

- RIP plain text authentication provides only limited security because the password transferred through the packet is visible.
- RIP MD5 authentication can provide higher security because the password transferred through the packet is encrypted using the MD5 algorithm.
- Routes can be learned properly if the correct authentication parameters are configured.
- Routes cannot be learned if the incorrect authentication parameters are configured.

Related Commands

Enabling Source Address Verification

Command Syntax	ip rip authentication key-chain <i>name-of-keychain</i>
-----------------------	--

Parameter Description	<i>name-of-keychain</i> : Specifies the name of the key chain used for RIP authentication.
Command Mode	Interface configuration mode
Configuration Usage	The specified key chain must be defined by the key chain command in global configuration mode in advance.

Defining the RIP Authentication Mode

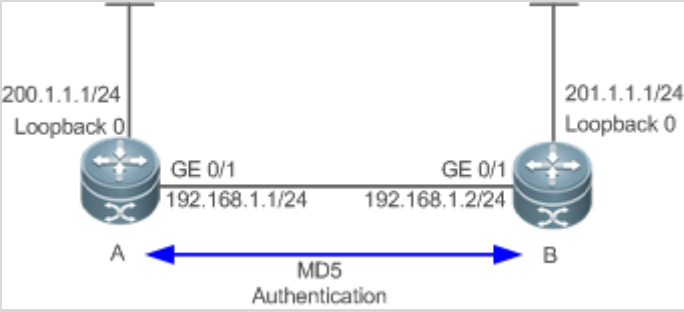
Command Syntax	ip rip authentication mode { text md5 }
Parameter Description	text: Indicates that the RIP authentication mode is plain text authentication. md5: Indicates that the RIP authentication mode is MD5 authentication.
Command Mode	Interface configuration mode
Configuration Usage	For all devices that need to directly exchange the RIP routing information, the RIP authentication mode of these devices must be the same.

Enabling RIP Plain Text Authentication and Configuring the Key Chain

Command Syntax	ip rip authentication text-password [0 7] <i>password-string</i>
Parameter Description	0: Indicates that the key is displayed in plain text. 7: Indicates that the key is displayed in cipher text. <i>password-string</i> : Indicates the key chain used for plain text authentication. The key chain is a string of 1 to 16 bytes.
Command Mode	Interface configuration mode
Configuration Usage	This commands takes effect only in plain text authentication mode.

Configuration Example

Configuring RIP Basic Functions and Enabling MD5 Authentication

<p>Scenario Figure 3-13</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the RIP basic functions on all routers. (Omitted) ▪ Configure the authentication type and MD5 authentication key on all routers.
<p>A</p>	<pre>A# configure terminal A(config)# key chain hello A(config-keychain)# key 1 A(config-keychain-key)# key-string world A(config-keychain-key)# exit A(config-keychain)# exit A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip rip authentication mode md5 A(config-if-GigabitEthernet 0/1)# ip rip authentication key-chain hello</pre>
<p>B</p>	<pre>B# configure terminal B(config)# key chain hello B(config-keychain)# key 1 B(config-keychain-key)# key-string world B(config-keychain-key)# exit B(config-keychain)# exit B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip rip authentication mode md5 B(config-if-GigabitEthernet 0/1)# ip rip authentication key-chain hello</pre>

Verification	<ul style="list-style-type: none"> On Router A, check the routing table and verify that the entry 201.1.1.0/24 is loaded. On Router B, check the routing table and verify that the entry 200.1.1.0/24 is loaded.
A	<pre>A# show ip route rip R 201.1.1.0/24 [120/1] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1</pre>
B	<pre>A# show ip route rip R 200.1.1.0/24 [120/1] via 192.168.1.1, 00:06:11, GigabitEthernet 0/1</pre>

Common Errors

- The keys configured on routers that need to exchange RIP routing information are different.
- The authentication modes configured on routers that need to exchange RIP routing information are different.

3.4.6 Enabling Route Summarization

Configuration Effect

Reduce the size of the routing table, improve the routing efficiency, avoid route flapping to some extent, and improve scalability and effectiveness of the network.

- i** If a summarized route exists, subroutes included by the summarized route cannot be seen in the routing table, which greatly reduces the size of the routing table.
- i** Advertising a summarized route is more efficient than advertising individual routes because: (1) A summarized route is processed first when RIP looks through the database; (2) All subroutes are ignored when RIP looks through the database, which reduces the processing time required.

Notes

- The RIP basic functions must be configured.
- The range of supernetting routes is larger than that of the classful network. Therefore, the automatic route summarization function is invalid for supernetting routes.
- RIPv1 always performs automatic route summarization. If the detailed routes should be advertised, you must set the RIP version to RIPv2.

Configuration Steps

Enabling Automatic Route Summarization

- This function is enabled by default.
- To learn specific subnet routes instead of summarized network routes, you must disable automatic route summarization.
- You can disable automatic route summarization only in RIPv2. RIPv1 always performs automatic route summarization.

Configuring RIP Route Summarization on an Interface

- This function must be configured if it is required to summarize classful subnets.
- The **ip rip summary-address** command is used to summarize an address or a subnet under a specified interface. RIP automatically summarizes to the classful network boundary. Each classful subnet can be configured only in the **ip rip summary-address** command.
- The summary range configured in this command cannot be supernetting routes, that is, the configured subnet mask length cannot be smaller than the natural mask length of the network.
- Unless otherwise required, this configuration should be performed on a router that requires classful subnet summarization.

Verification

Verify that the routes are summarized in the routing table of the peer end.

Related Commands

Enabling Automatic Route Summarization

Command Syntax	auto-summary
Parameter Description	N/A
Command Mode	Routing process configuration mode
Configuration Usage	Route summarization is enabled by default for RIPv1 and RIPv2. You can disable automatic route summarization only in RIPv2. RIPv1 always performs automatic route summarization.

Configuring RIP Route Summarization on an Interface

Command Syntax	ip rip summary-address <i>ip-address ip-network-mask</i>
-----------------------	---

Parameter Description	<i>ip-address</i> : Indicates the IP address to be summarized. <i>ip-network-mask</i> : Indicates the subnet mask of the IP address to be summarized.
Command Mode	Interface configuration mode
Configuration Usage	This command is used to summarize an address or a subnet under a specified interface.

Configuration Example

Configuring Route Summarization

Scenario Figure 3-14	
Remark	<p>The interface IP addresses are as follows:</p> <p>A: GE0/1 192.168.1.1</p> <p>B: GE0/1 192.168.1.2 GE0/2 172.16.2.1 GE0/3 172.16.3.1</p> <p>C: GE0/2 172.16.2.2 GE0/3 172.16.4.2</p> <p>D: GE0/2 172.16.3.2 GE0/3 172.16.5.2</p>
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Configure route summarization on Router B.
	B# configure terminal

	<pre>B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip rip summary-address 172.16.0.0 255.255.0.0 B(config)# router rip B(config-router)# version 2 B(config-router)# no auto-summary</pre>
Verification	Check the routing table on Router A, and verify that the entry 172.16.0.0/16 is generated.
	<pre>A# show ip route rip R 172.16.0.0/16 [120/2] via 192.168.1.2, 00:01:04, GigabitEthernet 0/1</pre>

Common Errors

- RIP basic functions are not configured or fail to be configured.

3.4.7 Enabling Supernetting Routes

Configuration Effect

- Allow RIP to send RIP supernetting routes on a specified interface.

Notes

- The RIP basic functions must be configured.

Configuration Steps

Enabling Supernetting Routes

- If a supernetting route is detected when a RIPv1-enabled router monitors the RIPv2 route response packets, the router will learn an incorrect route because RIPv1 ignores the subnet mask in the routing information of the packet. In this case, the **no** form of the command must be used on the RIPv2-enabled router to prohibit advertisement of supernetting routes on the related interface. This command takes effect only on the current interface.
- The command is effective only when RIPv2 packets are sent on the interface, and is used to control the sending of supernetting routes.

Verification

Verify that the peer router cannot learn the supernetting route.

Related Commands

Command Syntax	<code>ip rip send supernet-routes</code>
Parameter Description	N/A
Command Mode	Interface configuration mode
Configuration Usage	By default, an interface is allowed to send RIP supernetting routes.

Configuration Example

Disabling Supernetting Routes

Scenario Figure 3-15	<p>The diagram shows two routers, A and B, connected via GigabitEthernet 0/1 interfaces. Router A has IP 192.168.1.1/24 and Router B has IP 192.168.1.2/24. Router B is connected to a static network with IP 207.0.0.0/8 and 208.1.1.0/24.</p>
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Prohibit the sending of RIP supernetting routes on the GigabitEthernet 0/1 interface of Router B.
	<pre> B# configure terminal B(config)# ip route 207.0.0.0 255.0.0.0 Null 0 B(config)# ip route 208.1.1.0 255.255.255.0 Null 0 B(config)# router rip B(config-router)# redistribute static B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# no ip rip send supernet-routes </pre>

Verification	Check the routing table on Router A, and verify that Router A can learn only the non-supernetting route 208.1.1.0/24, but not the supernetting route 207.0.0.0/8.
	<pre>A#show ip route rip R 208.1.1.0/24 [120/1] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1</pre>

3.4.8 Advertising the Default Route or External Routes

Configuration

Effect

- In the RIP domain, introduce a unicast route of another AS so that the unicast routing service to this AS can be provided for users in the RIP domain.
- In the RIP domain, inject a default route to another AS so that the unicast routing service to this AS can be provided for users in the RIP domain.

Notes

- The RIP basic functions must be configured.
- Route redistribution cannot introduce default routes of other protocols to the RIP routing domain.

Configuration

Steps

Advertising the Default Route to Neighbors

This function must be enabled if it is required to advertise the default route to neighbors.

By default, a default route is not generated, and the metric of the default route is 1.

If the RIP process can generate a default route using this command, RIP does not learn the default route advertised by the neighbor.

Unless otherwise required, this configuration should be performed on a router that needs to advertise the default route.

Advertising the Default Route to Neighbors on an Interface

This function must be enabled if it is required to advertise the default route to neighbors on a specified interface.

By default, a default route is not configured and the metric of the default route is 1.

After this command is configured on an interface, a default route is generated and advertised through this interface.

Unless otherwise required, this configuration should be performed on a router that needs to advertise the default route.

Redistributes Routes and Advertises External Routes to Neighbors

This function must be enabled if routes of other protocols need to be redistributed.

By default,

- If OSPF redistribution is configured, redistribute the routes of all sub-types of the OSPF process.
- If IS-IS redistribution is configured, redistribute the level-2 routes of the IS-IS process.
- In other cases, redistribute all external routes.
- The metric of a redistributed route is 1 by default.
- The route map is not associated by default.

During route redistribution, it is not necessary to convert the metric of one routing protocol to the metric of another routing protocol because different routing protocols use completely different metric measurement methods. RIP measures the metric based on the hop count, and OSPF measures the metric based on the bandwidth. Therefore, the computed metrics cannot be compared with each other. During route redistribution, however, it is necessary to configure a symbolic metric; otherwise, route redistribution fails.

Unless otherwise required, this configuration should be performed on a router that needs to redistribute routes.

Verification

- On a neighbor device, verify that a default route exists in the RIP routing table.
- On the local and neighbor devices, verify that external routes (routes to other ASs) exist in the RIP routing table.

Related Commands

Advertising the Default Route to Neighbors

Command Syntax	<code>default-information originate [always] [metric <i>metric-value</i>] [route-map <i>map-name</i>]</code>
Parameter Description	<p><code>always</code>: Enables RIP to generate a default route no matter whether the local router has a default route.</p> <p><code>metric <i>metric-value</i></code>: Indicates the initial metric of the default route. The value ranges from 1 to 15.</p> <p><code>route-map <i>map-name</i></code>: Indicates the associated route map name. By default, no route map is associated.</p>
Command Mode	Routing process configuration mode

Configurati on Usage	<p>If a default route exists in the routing table of a router, RIP does not advertise the default route to external entities by default. You need to run the default-information originate command in routing process configuration mode to advertise the default route to neighbors.</p> <p>If the always parameter is selected, the RIP routing process advertises a default route to neighbors no matter the default route exists, but this default route is not displayed in the local routing table. To check whether the default route is generated, run the show ip rip database command to check the RIP routing information database.</p> <p>To further control the behavior of advertising the RIP default route, use the route-map parameter. For example, run the set metric rule to set the metric of the default route.</p> <p>You can use the metric parameter to set the metric of the advertised default value, but the priority of this configuration is lower than that of the set metric rule of the route-map parameter. If the metric parameter is not configured, the default route uses the default metric configured for RIP.</p> <p>You still need to run the default-information originate command to introduce the default route generated by ip default-network to RIP.</p>
-------------------------	--

Advertising the Default Route to Neighbors on an Interface

Command Syntax	ip rip default-information { only originate } [metric <i>metric-value</i>]
Parameter Description	<p>only: Indicates that only the default route is advertised.</p> <p>originate: Indicates that the default route and other routes are advertised.</p> <p>metric <i>metric-value</i>: Indicates the metric of the default route. The value ranges from 1 to 15.</p>
Command Mode	Interface configuration mode
Configurati on Usage	<p>If you configure the ip rip default-information command for the interface, and the default-information originate command for the RIP process, only the default route configured for the interface is advertised.</p> <p>So far as ip rip default-information is configured for one interface, RIP does not learn the default route advertised by the neighbor.</p>

Redistributes Routes and Advertises External Routes to Neighbors

Command Syntax	<code>redistribute { bgp connected isis [<i>area-tag</i>] ospf <i>process-id</i> static } [{ level-1 level-1-2 level-2 }] [match { internal external [1 2] nssa-external [1 2] }] [metric <i>metric-value</i>] [route-map <i>route-map-name</i>]</code>
Parameter Description	<p><code>bgp</code>: Indicates redistribution from BGP.</p> <p><code>connected</code>: Indicates redistribution from direct routes.</p> <p><code>isis <i>area-tag</i></code>: Indicates redistribution from IS-IS. <i>area-tag</i> indicates the IS-IS process ID.</p> <p><code>ospf <i>process-id</i></code>: Indicates redistribution from OSPF. <i>process-id</i> indicates the OSPF process ID. The value ranges from 1 to 65535.</p> <p><code>static</code>: Indicates redistribution from static routes.</p> <p><code>level-1 level-1-2 level-2</code>: Used only when IS-IS routes are redistributed. Only the routes of the specified level are redistributed.</p> <p><code>match</code>: Used only when OSPF routes are redistributed. Only the routes that match the filtering conditions are redistributed.</p> <p><code>metric <i>metric-value</i></code>: Sets the metric of the redistributed route. The value ranges from 1 to 16.</p> <p><code>route-map <i>route-map-name</i></code>: Sets the redistribution filtering rules.</p>
Command Mode	Routing process configuration mode
Configuration Usage	<p>When you configure redistribution of IS-IS routes without specifying the level parameter, only level-2 routes can be redistributed by default. If you specify the level parameter during initial configuration of redistribution, routes of the specified level can be redistributed. If both level-1 and level-2 are configured, the two levels are combined and saved as level-1-2 for the convenience sake.</p> <p>If you configure redistribution of OSPF routes without specifying the match parameter, OSPF routes of all sub-types can be distributed by default. The latest setting of the match parameter is used as the initial match parameter. Only routes that match the sub-types can be redistributed. You can use the no form of the command to restore the default value of match.</p> <p>The configuration rules for the no form of the redistribute command are as follows:</p> <ol style="list-style-type: none">1. If some parameters are specified in the no form of the command, default values of these parameters will be restored.

2. If no parameter is specified in the no form of the command, the entire command will be deleted.


For example, if redistribute isis 112 level-2 is configured, you can run the no redistribute isis 112 level-2 command to restore the default value of level-2. As level-2 itself is the default value of the parameter, the configuration saved is still redistribute isis 112 level-2 after the preceding no form of the command is executed.

To delete the entire command, run the no redistribute isis 112 command.

Configuration

Example

Redistributing Routes and Advertising External Routes to Neighbors

<p>Scenario Figure 3-16</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the RIP basic functions on all routers. (Omitted) ▪ On Router B, configure redistribution of static routes.
<p>B</p>	<pre>B# configure terminal B(config)# router rip B(config-router)# redistribute static</pre>
<p>Verification</p>	<p>On Router A, check the routing table and verify that the entry 172.10.10.0/24 is loaded.</p>
	<pre>A# show ip route rip R 172.10.10.0/24 [120/1] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1</pre>

3.4.9 Setting Route Filtering Rules

Configuration

Effect

- Routes that do not meet filtering criteria cannot be loaded to the routing table, or advertised to neighbors. In this way, users within the network can be prevented from accessing specified destination networks.

Notes

- The RIP basic functions must be configured.
- In regard to the filtering rules of sent routes, you must configure route redistribution first, and then filter the redistributed routes.

Configuration

Steps

Filtering the Received RIP Routing Information

- This function must be configured if it is required to filter received routing information.
- To refuse receiving some specified routes, you can configure the route distribution control list to process all the received route update packets. If no interface is specified, route update packets received on all interfaces will be processed.
- Unless otherwise required, this configuration should be performed on a router that requires route filtering.

Filtering the Sent RIP Routing Information

- This function must be configured if it is required to filter the redistributed routing information that is sent.
- If this command does not contain any optional parameter, route update advertisement control takes effect on all interfaces. If the command contains the interface parameter, route update advertisement control takes effect only on the specified interface. If the command contains other routing process parameters, route update advertisement control takes effect only on the specified routing process.
- Unless otherwise required, this configuration should be performed on a router that requires route filtering.

Verification

- Run the **show ip route rip** command to verify that the routes that have been filtered out are not loaded to the routing table.

Related

Commands

Filtering the Received RIP Routing Information

Command Syntax	<code>distribute-list { [<i>access-list-number</i> <i>name</i>] prefix <i>prefix-list-name</i> [<i>gateway prefix-list-name</i>] } in [<i>interface-type interface-number</i>]</code>
Parameter Description	<i>access-list-number</i> <i>name</i> : Specifies the access list. Only routes permitted by the access list can be received. <i>prefix prefix-list-name</i> : Uses the prefix list to filter routes.

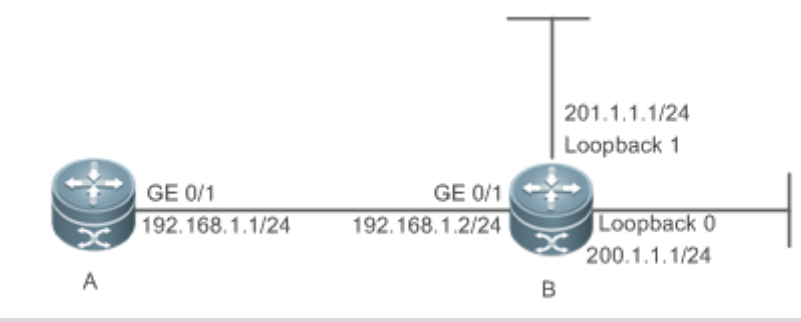
	<p><i>gateway prefix-list-name</i>: Uses the prefix list to filter the route sources.</p> <p><i>interface-type interface-number</i>: Indicates that the distribution list is applied to the specified interface.</p>
Command Mode	Routing process configuration mode
Configuration Usage	N/A

Filtering the Sent RIP Routing Information

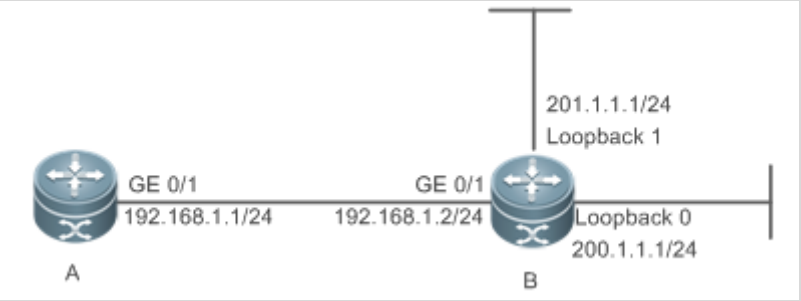
Command Syntax	distribute-list { [<i>access-list-number</i> <i>name</i>] prefix <i>prefix-list-name</i> } out [<i>interface</i> [<i>bgp</i> <i>connected</i> <i>isis</i> [<i>area-tag</i>] <i>ospf process-id</i> <i>rip</i> <i>static</i>]]
Parameter Description	<p><i>access-list-number</i> <i>name</i>: Specifies the access list. Only routes permitted by the access list can be sent.</p> <p>prefix <i>prefix-list-name</i>: Uses the prefix list to filter routes.</p> <p><i>Interface</i>: Applies route update advertisement control only on the specified interface.</p> <p>bgp: Applies route update advertisement control only on the routes introduced from BGP.</p> <p>connected: Applies route update advertisement control only on direct routes introduced through redistribution.</p> <p>isis [<i>area-tag</i>]: Applies route update advertisement control only on the routes introduced from IS-IS. <i>area-tag</i> specifies an IS-IS process.</p> <p>ospf <i>process-id</i>: Applies route update advertisement control only on the routes introduced from OSPF. <i>process-id</i> specifies an OSPF process.</p> <p>rip: Applies route update advertisement control only on RIP routes.</p> <p>static: Applies route update advertisement control only on static routes introduced through redistribution.</p>
Command Mode	Routing process configuration mode
Configuration Usage	N/A

Configuration Example

Filtering the Received RIP Routing Information

<p>Scenario Figure 3-17</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Enable the RIP routing process to control routes received over the GigabitEthernet 0/1 port and receive only the route 200.1.1.0.
<p>A</p>	<pre>A# configure terminal A(config)# router rip A(config-router)# distribute-list 10 in GigabitEthernet 0/1 A(config-router)# no auto-summary A(config)# access-list 10 permit 200.1.1.0 0.0.0.255</pre>
<p>Verification</p>	<p>On Router A, check the routing table and verify that only the entry 200.1.1.0/24 exists.</p>
<p>A</p>	<pre>A# show ip route rip R 200.1.1.0/24 [120/1] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1</pre>

Filtering the Sent RIP Routing Information

<p>Scenario Figure 3-18</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Enable the RIP routing process to advertise only the route 200.1.1.0/24.

B	<pre> B# configure terminal B(config)# router rip B(config-router)# redistribute connected B(config-router)# distribute-list 10 out B(config-router)# version 2 B(config)# access-list 10 permit 200.1.1.0 0.0.0.255 </pre>
Verification	Check the routing table on Router A, and verify that route in the 200.1.1.0 network segment exists.
A	<pre> A# show ip route rip R 200.1.1.0/24 [120/1] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1 </pre>

Common Errors

- Filtering fails because the filtering rules of the access list are not properly configured.

3.4.10 Modifying Route Selection Parameters

Configuration Effect

- Change the RIP routes to enable the traffic pass through specified nodes or avoid passing through specified nodes.
- Change the sequence that a router selects various types of routes so as to change the priorities of RIP routes.

Notes

- The RIP basic functions must be configured.

Configuration Steps

Modifying the Administrative Distance of a RIP Route

- Optional.
- This configuration is mandatory if you wish to change the priorities of RIP routes on a router that runs multiple unicast routing protocols.

Increasing the Metric of a Received or Sent RIP Route

- Optional.

- Unless otherwise required, this configuration should be performed on a router where the metrics of routes need to be adjusted.

Configuring the Default Metric of an External Route Redistributed to RIP

- Optional.
- Unless otherwise required, this configuration must be performed on an ASBR to which external routes are introduced.

Verification

Run the **show ip rip** command to display the administrative distance currently configured. Run the **show ip rip data** command to display the metrics of redistributed routes to verify that the configuration takes effect.

Related Commands

Modifying the Administrative Distance of a RIP Route

Command Syntax	distance <i>distance</i> [<i>ip-address wildcard</i>]
Parameter Description	<i>distance</i> : Sets the administrative distance of a RIP route. The value is an integer ranging from 1 to 255. <i>ip-address</i> : Indicates the prefix of the source IP address of the route. <i>wildcard</i> : Defines the IP address comparison bit. 0 indicates accurate matching, and 1 indicates that no comparison is performed.
Command Mode	Routing process configuration mode
Configuration Usage	Run this command to configure the administrative distance of a RIP route.

Increasing the Metric of a Received or Sent RIP Route

Command Syntax	offset-list { <i>access-list-number</i> <i>name</i> } { in out } <i>offset</i> [<i>interface-type interface-number</i>]
Parameter Description	<i>access-list-number</i> <i>name</i> : Specifies the access list. In : Uses the ACL to modify the metric of a received route. out : Uses the ACL to modify the metric of a sent route.

	<p><i>offset</i>: Indicates the offset of the modified metric. The value ranges from 0 to 16.</p> <p><i>interface-type</i>: Uses the ACL on the specified interface.</p> <p><i>interface-number</i>: Specifies the interface number.</p>
Command Mode	Routing process configuration mode
Configuration Usage	Run this command to increase the metric of a received or sent RIP route. If the interface is specified, the configuration takes effect only on the specified interface; otherwise, the configuration takes effect globally.

Configuring the Default Metric of an External Route Redistributed to RIP

Command Syntax	default-metric <i>metric-value</i>
Parameter Description	<i>metric-value</i> : Indicates the default metric. The valid value ranges from 1 to 16. If the value is equal to or greater than 16, the RGOS determines that this route is unreachable.
Command Mode	Routing process configuration mode
Configuration Usage	This command must be used together with the routing protocol configuration command redistribute .

Configuration Example

Increasing the Metric of a Received or Sent RIP Route

Scenario Figure 3-19	<p>The diagram illustrates a network topology with two routers, A and B, connected via their GE 0/1 interfaces. Router A has a loopback interface Loopback 0 with IP 200.1.1.1/24. Router B has a loopback interface Loopback 0 with IP 201.1.1.1/24. The link between the two routers has IP addresses 192.168.1.1/24 on Router A and 192.168.1.2/24 on Router B.</p>
-------------------------	--

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the RIP basic functions on all routers. (Omitted) ▪ Increase by 7 the metric of each RIP route in the range specified by ACL 7. ▪ Increase by 7 the metric of each learned RIP route in the range specified by ACL 8.
A	<pre>A# configure terminal A(config)# access-list 7 permit host 200.1.1.0 A(config)# access-list 8 permit host 201.1.1.0 A(config)# router rip A(config-router)# offset-list 7 out 7 A(config-router)# offset-list 8 in 7</pre>
Verification	Check the routing table on Router A and Router B to verify that the metrics of RIP routes are 8.
A	<pre>A# show ip route rip R 201.1.1.0/24 [120/8] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1</pre>
B	<pre>B# show ip route rip R 200.1.1.0/24 [120/8] via 192.168.1.1, 00:06:11, GigabitEthernet 0/1</pre>

3.4.11 Modifying Timers

Configuration

Effect

- Change the duration of RIP timers to accelerate or slow down the change of the protocol state or occurrence of an event.

Notes

- The RIP basic functions must be configured.
- Modifying the protocol control parameters may result in protocol running failures. Therefore, you are advised not to modify the timers.

Configuration

Steps

Modifying the Update Timer, Invalid Timer, and Flush Timer

This configuration must be performed if you need to adjust the RIP timers.

By adjusting the timers, you can reduce the convergence time and fault rectification time of the routing protocol. For routers connected to the same network, values of the three RIP timers must be the same. Generally, you are advised not to modify the RIP timers unless otherwise required.

Setting timers to small values on a low-speed link brings risks because a lot of Update packets consume the bandwidth. You can set timers to small values generally on the Ethernet or a 2 Mbps (or above) link to reduce the convergence time of network routes.

Unless otherwise required, this configuration should be performed on a router where RIP timers need to be modified.

Setting the Sending Delay Between RIP Route Update Packets

This configuration must be performed if you need to adjust the sending delay between RIP Update packets.

Run the **output-delay** command to increase the sending delay between packets on a high-speed device so that a low-speed device can receive and process all Update packets.

Unless otherwise required, this configuration should be performed on a router where the sending delay needs to be adjusted.

Verification

Run the **show ip rip** command to display the current settings of RIP timers.

Related Commands

Modifying the Update Timer, Invalid Timer, and Flush Timer

Command Syntax	timers basic update invalid flush
Parameter Description	<p><i>update</i>: Indicates the route update time in second. It defines the interval at which the device sends the route update packet. Each time an Update packet is received, the invalid timer and flush timer are reset. By default, a routing update packet is sent every 30s.</p> <p><i>invalid</i>: Indicates the route invalid time in second, counted from the last time when a valid update packet is received. It defines the time after which the route in the routing list becomes invalid because the route is not updated. The duration of the invalid timer must be at least three times the duration of the update timer. If no Update packet is received before the invalid timer expires, the corresponding route enters the invalid state. If the Update packet is received before the invalid timer expires, the timer is reset. The default duration of the invalid timer is 180s.</p> <p><i>flush</i>: Indicates the route flushing time in second, counted from the time when the RIP route enters the invalid state. When the flush timer expires, the route in the invalid</p>

	state will be deleted from the routing table. The default duration of the flush timer is 120s.
Command Mode	Routing process configuration mode
Configuration Usage	By default, the update timer is 30s, the invalid timer is 180s, and the flush timer is 120s.

Setting the Sending Delay Between RIP Route Update Packets

Command Syntax	output-delay <i>delay</i>
Parameter Description	<i>delay</i> : Sets the sending delay between packets in ms. The value ranges from 8 to 50.
Command Mode	Interface configuration mode
Configuration Usage	<p>Normally, a RIP route update packet is 512 bytes long and can contain 25 routes. If the number of routes to be updated exceeds 25, more than one update packet will be sent as fast as possible.</p> <p>When a high-speed device sends a lot of update packets to a low-speed device, the low-speed device may not be able to process all update packets in time, causing a loss of routing information. In this case, you need to run the output-delay command to increase the sending delay between packets on a high-speed device so that a low-speed device can receive and process all update packets.</p>

Configuration

Example

Setting the Sending Delay Between RIP Route Update Packets

Scenario Figure 3-20	
Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the RIP basic functions on all routers. (Omitted) ▪ Configure the sending delay of update packets on Router A.
A	<pre>A# configure terminal A(config)# router rip A(config-router)# output-delay 30</pre>
Verification	<p>Capture packets on Router A and compare the sending time of update packets before and after the configuration, and verify that a delay of 30 ms is introduced.</p>

Common Errors

For routers connected to the same network, values of the three RIP timers are not the same.

3.4.12 Enabling BFD Correlation

Configuration Effect

- Once a link is faulty, RIP can quickly detect the failure of the route. This configuration helps shorten the traffic interruption time.

Notes

- The RIP basic functions must be configured.
- The BFD correlation configured in interface configuration mode takes precedence over the global configuration.

Configuration Steps

Correlating RIP with BFD on All Interfaces

- This configuration must be performed if you need to enable BFD correlation.
- After BFD is enabled on RIP, a BFD session will be set up for the RIP routing information source (that is, the source address of RIP route update packets). Once the BFD neighbor fails, the corresponding RIP route directly enters the invalid state and is not forwarded.
- You can also run the **ip ospf bfd [disable]** command in interface configuration mode to enable or disable the BFD function on a specified interface, and this configuration takes precedence over the **bfd all-interfaces** command used in routing process configuration mode.
- Unless otherwise required, this configuration should be performed on every router.

Correlating RIP with BFD on an Interface

- This configuration must be performed if you need to enable or disable BFD correlation on a specified interface.
- The interface-based configuration takes precedence over the **bfd all-interfaces** command used in routing process configuration mode.
- Based on the actual environment, you can run the **ip ospf bfd** command to enable BFD on a specified interface for link detection, or run the **bfd all-interfaces** command in RIP process configuration mode to enable BFD on all interface of the OSPF process, or run the **ospf bfd disable** command to disable BFD on a specified interface.
- Unless otherwise required, configure this function on a router interface where BFD correlation should be configured separately.

Verification

- Verify that the BFD session is properly set up with RIP.
- After a link fails, the RIP route can quickly converges.

Related Commands

Correlating RIP with BFD on All Interfaces

Command Syntax	bfd all-interfaces
Parameter Description	N/A
Command Mode	Routing process configuration mode

Configurati on Usage	N/A
-------------------------	-----

Correlating RIP with BFD on an Interface

Command Syntax	ip rip bfd [disable]
Parameter Description	disable: Disables BFD for link detection on a specified RIP-enabled interface.
Command Mode	Interface configuration mode
Configurati on Usage	By default, BFD correlation is not configured for a specified interface, and the configuration is subject to that configured in routing process configuration mode.

Configuration Example

Enabling BFD Correlation with RIP

Scenario Figure 3-21	
Configurati on Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) Configure the BFD parameters for interfaces of all routers. Correlate RIP with BFD on all routers.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# bfd interval 200 min_rx 200 multiplier 5 A(config)# router rip A(config-router)# bfd all-interfaces</pre>
B	<pre>B# configure terminal</pre>

	<pre>B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# bfd interval 200 min_rx 200 multiplier 5 B(config)# router rip B(config-router)# bfd all-interfaces</pre>
Verification	<ul style="list-style-type: none"> On routers A and B, verify that the BFD session is in Up state. Disconnect Router B from the switch, and verify that the RIP route is deleted on Router A.
A	<pre>A# show ip rip peer Peer 192.168.1.2: Local address: 192.168.1.1 Input interface: GigabitEthernet 0/1 Peer version: RIPv2 Received bad packets: 0 Received bad routes: 0 BFD session state up</pre>
B	<pre>A# show ip rip peer Peer 192.168.1.1: Local address: 192.168.1.2 Input interface: GigabitEthernet 0/1 Peer version: RIPv2 Received bad packets: 0 Received bad routes: 0 BFD session state up</pre>

Common Errors

- The preceding two commands are executed in RIP before the BFD function is enabled.

3.4.13 Enabling Fast Reroute

Configuration Effect

- Once RIP detects a route failure, the router can immediately switch to the second-best route. This configuration helps shorten the traffic interruption time.

Notes

- The RIP basic functions must be configured.
- The route map and the standby next hop must be configured.
- To accelerate the convergence, set carrier-delay of the interface to 0 and enable BFD correlation with RIP.

Configuration Steps

Enabling Fast Reroute and Referencing the Route Map

This configuration must be performed if you need to enable fast reroute.

If **route-map** is configured, a standby path can be specified for a successfully matched route through the route map.

When the RIP fast reroute function is used, it is recommended that BFD be enabled at the same time so that the device can quickly detect any link failure and therefore shorten the forwarding interruption time. If the interface is up or down, to shorten the forwarding interruption time during RIP fast reroute, you can configure **carrier-delay 0** in interface configuration mode to achieve the fastest switchover speed.

Unless otherwise required, this configuration should be performed on every router.

Verification

- The standby route can be correctly computed and generated.
- When the active link fails, the data can be quickly switch over to the standby link for forwarding.

Related Commands

Enabling Fast Reroute and Referencing the Route Map

Command Syntax	fast-reroute route-map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Specifies a standby path through the route map.
Command Mode	Routing process configuration mode

Configurati on Usage	Currently, the RIP fast reroute function is subject to the following constraints: (1) Only one standby next hop can be generated for one route; (2) No standby next hop can be generated for equal and equal-cost multi-path routing (ECMP).
-------------------------	--

Configuration Example

Enabling Fast Reroute and Referencing the Route Map

<p>Scenario Figure 3-22</p>	
	<p>Remarks</p> <p>The interface IP addresses are as follows: A: GE0/1 192.168.1.1 GE0/2 192.168.2.1] B: GE0/1 192.168.1.2 GE0/2 192.168.3.1 GE0/3 192.168.4.1 C: GE0/1 192.168.3.2 GE 0/2 192.168.2.2</p>
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the RIP basic functions on all routers. (Omitted) ▪ Configure fast re-route on Router A. ▪ Configure carrier-delay 0 for the interface on Router A.
<p>A</p>	<pre>A# configure terminal A(config)# route-map fast-reroute A(config-route-map)# match interface GigabitEthernet 0/2 A(config-route-map)# set fast-reroute backup-interface GigabitEthernet 0/1 backup- nextrhop 192.168.1.1 A(config)# router rip A(config-router)# fast-reroute route-map fast-reroute A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# carrier-delay 0 A(config-if-GigabitEthernet 0/1)# exit</pre>

	<pre>A(config)# interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)# carrier-delay 0</pre>
Verification	On Router A, check the routing table and verify that a standby route exists for the entry 192.168.4.0/24.
A	<pre>A# show ip route fast-reroute begin 192.168.4.0 R 192.168.4.0/24 [ma] via 192.168.1.2, 00:39:28, GigabitEthernet 0/1 [b] via 192.168.2.2, 00:39:28, GigabitEthernet 0/2</pre>

Common Errors

- The standby next hop is not properly configured for the route map.
- The carrier-delay is not configured for the interface or BFD correlation is not configured. Consequently, the switchover speed of the forwarding line is slow.

3.4.14 Enabling GR

Configuration Effect

- When a distributed route switches services from the active board to the standby board, traffic forwarding continues and is not interrupted.
- When the RIP process is being restarted, traffic forwarding continues and is not interrupted.

Notes

- The RIP basic functions must be configured.
- The GR period is at least twice the RIP route update period.
- During the RIP GR process, ensure that the network environment is stable.

Configuration Steps

Configuring the GR Restarter Capability

This configuration must be performed if RIP needs to be gracefully restarted to ensure data forwarding during hot standby switchover.

The GR function is configured based on the RIP process. You can configure different parameters for different RIP processes based on the actual conditions.

The GR period is the maximum time from restart of the RIP process to completion of GR. During this period, the forwarding table before the restart is retained, and the RIP route is restored so as to restore

the RIP state before the restart. After the restart period expires, RIP exits from the GR state and performs common RIP operations.

Unless otherwise required, this configuration should be performed on every router that needs to be gracefully restarted.

Verification

- Run the **show ip rip** command to display the GR state and configured time.
- Trigger a hot standby switchover, and verify that data forwarding is not interrupted.

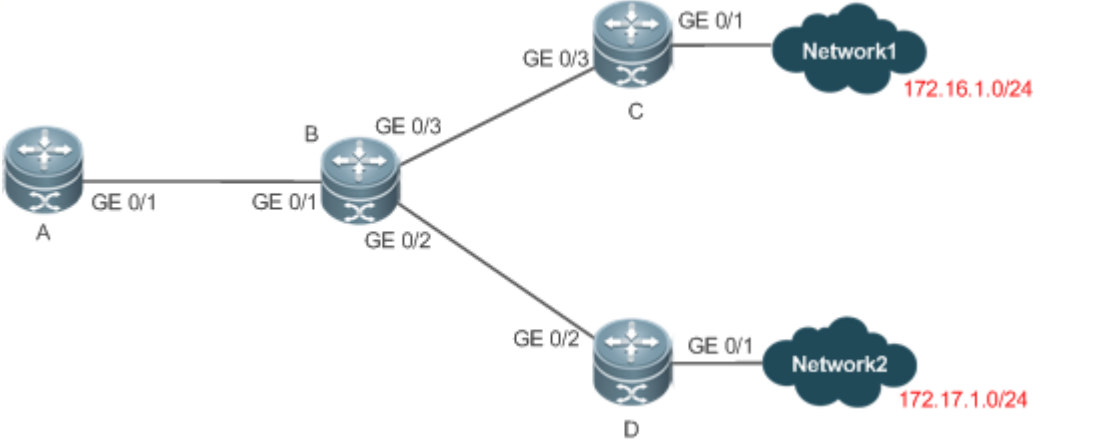
Related Commands

Configuring the GR Restarter Capability

Command Syntax	graceful-restart [grace-period <i>grace-period</i>]
Parameter Description	graceful-restart: Enables the GR function. grace-period: Explicitly configures the grace period. <i>grace-period:</i> Indicates the GR period. The value ranges from 1s to 1800s. The default value is twice the update time or 60s, whichever is the smaller.
Command Mode	Routing process configuration mode
Configuration Usage	This command allows you to explicitly modify the GR period. Note that GR must be completed after the update timer of the RIP route expires and before the invalid timer of the RIP route expires. An inappropriate GR period cannot ensure uninterrupted data forwarding during the GR process. A typical case is as follows: If the GR period is longer than the duration of the invalid timer, GR is not completed when the invalid timer expires. The route is not re-advertised to the neighbor, and forwarding of the route of the neighbor stops after the invalid timer expires, causing interruption of data forwarding on the network. Unless otherwise required, you are advised not to adjust the GR period. If it is necessary to adjust the GR period, ensure that the GR period is longer than the duration of the update timer but shorter than the duration of the invalid timer based on the configuration of the timers basic command.

Configuration Example

Configuring the GR Restarter Capability

<p>Scenario Figure 3-23</p>	 <table border="1" data-bbox="349 651 1445 913"> <tr> <td>Remark</td> <td>The interface IP addresses are as follows:</td> </tr> <tr> <td>S</td> <td>A: GE 0/1 192.168.1.1</td> </tr> <tr> <td></td> <td>B: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1 GE 0/3 192.168.3.1</td> </tr> <tr> <td></td> <td>C: GE 0/1 192.168.4.2 GE 0/3 192.168.3.2</td> </tr> <tr> <td></td> <td>D: GE 0/1 192.168.5.2 GE 0/2 192.168.2.2</td> </tr> </table>	Remark	The interface IP addresses are as follows:	S	A: GE 0/1 192.168.1.1		B: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1 GE 0/3 192.168.3.1		C: GE 0/1 192.168.4.2 GE 0/3 192.168.3.2		D: GE 0/1 192.168.5.2 GE 0/2 192.168.2.2
Remark	The interface IP addresses are as follows:										
S	A: GE 0/1 192.168.1.1										
	B: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1 GE 0/3 192.168.3.1										
	C: GE 0/1 192.168.4.2 GE 0/3 192.168.3.2										
	D: GE 0/1 192.168.5.2 GE 0/2 192.168.2.2										
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the RIP basic functions on all routers. (Omitted) On Router B, enable the GR function. 										
	<pre>B# configure terminal B(config)# router rip B(config-router)# graceful-restart grace-period 90</pre>										
<p>Verification</p>	<ul style="list-style-type: none"> Trigger a hot standby switchover on Router B, and verify that the routing tables of destination Network 1 and Network 2 remain unchanged on Router A during the switchover. Trigger a hot standby switchover on Router B, ping destination Network 1 from Router A, and verify that traffic forwarding is not interrupted during the switchover. 										

3.4.15 Enabling Multiple Instances

Configuration

Effect

- Run RIP on VPN instances.

Notes

- The RIP basic functions (with the VRF parameter) must be configured.

Configuration Steps

Creating a VRF Instance and Entering the IPv4 VRF Address Family

- This configuration must be performed if you need to configure RIP multiple instances and associate these RIP instances with VRF.
- Unless otherwise required, this configuration should be performed on every router that requires the RIP multiple instances.

Binding the RIP MIB with a VPN Instance

- This configuration must be performed if you configure RIP multiple instances and wish to manage non-default RIP instances using the MIB.
- The RIP MIB does not have the RIP instance information. Therefore, you must perform operations only on one instance through SNMP. By default, the RIP MIB is bound with the RIP instance of the default VRF, and all user operations take effect on this instance.
- If you wish to perform operations on a specified RIP instance through SNMP, run this command to bind the MIB with the instance.
- Unless otherwise required, this configuration should be performed on a router where the instance is managed using the MIB.

Verification

- Check the VRF routing table on a router to verify that the route to a remote network can be obtained through RIP.
- Use the MIB management software to manage the bound instance.

Related Commands

Creating a VRF Instance and Entering the IPv4 VRF Address Family

Command Syntax	<code>address-family ipv4 vrf <i>vrf-name</i></code>
Parameter Description	vrf <i>vrf-name</i> : Specifies the name of the VRF associated with the address family configuration sub-mode.
Command Mode	Routing process configuration mode

Configurati on Usage	<p>Run the address-family command to enter address family configuration sub-mode, the prompt of which is (config-router-af)#. When the VRF associated with the address family configuration sub-mode is specified for the first time, the RIP instance corresponding to the VRF will be created. In this submode, you can configure the RIP routing information for the related VRF.</p> <p>To exit from address family configuration sub-mode and return routing process configuration mode, run the exit-address-family or exit command.</p>
-------------------------	---

Exiting From an IPv4 VRF Address Family

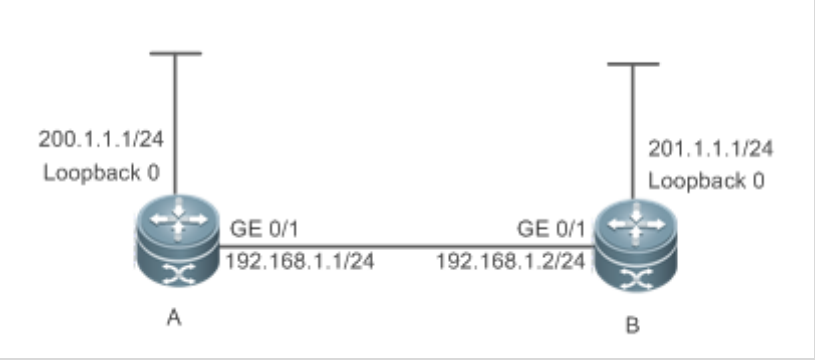
Command Syntax	exit-address-family
Parameter Description	N/A
Command Mode	Address family configuration mode
Configurati on Usage	<p>Run this command in address family configuration mode to exit from this configuration mode.</p> <p>This command can be abbreviated as exit.</p>

Binding the RIP MIB with a VPN Instance

Command Syntax	enable mib-binding
Parameter Description	N/A
Command Mode	Routing process configuration mode
Configurati on Usage	N/A

Configuration Example

Creating a VRF Instance and Enabling Network Management of This Instance


<p>Scenario Figure 3-24</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the RIP basic functions on all routers. (Omitted) ▪ Create a VRF named "vpn1" and create a RIP instance for this VRF. ▪ On Router A, bind the MIB with the RIP vpn1 instance.
	<pre>A# configure terminal A(config)# snmp-server community public rw A(config)# ip vrf vpn1 A(config-vrf)# exit A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet0/1)# ip vrf forwarding vpn1 A(config-if-GigabitEthernet0/1)# ip address 192.168.1.1 255.255.255.0 A(config)# router rip A(config-router)# address-family ipv4 vrf vpn1 A(config-router)# enable mib-binding A(config-router-af)# network 192.168.1.0 A(config-router-af)# exit-address-family</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Check the routing table on Router A, and verify that the VRF route 201.1.1.0/24 can be learned. ▪ Read and configure parameters of the RIP vpn1 instance using the MIB tool.
	<pre>A# show ip route vrf vpn1 rip R 201.1.1.0/24 [120/1] via 192.168.1.2, 00:06:11, GigabitEthernet 0/1</pre>

3.5 Monitoring

Displaying

Description	Command
Displays the basic information about a RIP process.	show ip rip
Displays the RIP routing table.	show ip rip database [vrf vrf-name] [network-number network-mask] [count]
Displays information about external routes redistributed by RIP.	show ip rip external [bgp connected isis [process-id] ospf process-id static] [vrf vrf-name]
Displays the RIP interface information.	show ip rip interface [vrf vrf-name] [interface-type interface-number]
Displays the RIP neighbor information.	show ip rip peer [ip-address] [vrf vrf-name]

Debugging

 System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Debugs events that occur when the RIP process is running.	debug ip rip event
Debugs interaction with the NSM process.	debug ip rip nsm
Debugs the sent and received packets.	debug ip rip packet [interface interface-type interface-number recv send]
Debugs the RIP GR process.	debug ip rip restart

Debugs the route changes of the RIP process.

debug ip rip route

4 CONFIGURING OSPFV2

4.1 Overview

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) that is used within the Autonomous System (AS) to allow routers to obtain a route to a remote network.

- i** OSPF Version 2 (OSPFv2) is applicable to IPv4, and OSPF Version 3 (OSPFv3) is applicable to IPv6. The protocol running mechanism and most configurations are the same.

OSPF has the following characteristics:

- Wide scope of application: OSPF is applicable to a larger-scale network that supports hundreds of routers.
- Fast convergence: Once the network topology changes, notifications can be quickly sent between routers to update routes.
- No self-loop: Only the link status information is synchronized between routers. Each router computes routes independently, and a self-loop will not occur.
- Area division: A large routing domain is divided into multiple small areas to save system resources and network bandwidth and ensure stability and reliability of routes.
- Route classification: Routes are classified into several types to support flexible control.
- Equivalent routes: OSPF supports equivalent routes.
- Authentication: OSPF supports packet authentication to ensure security of protocol interaction.
- Multicast transmission: Protocol packets are sent using the multicast address to avoid interfering with irrelevant entities and save system resources.

- i** In this chapter, the term "router" refers to any network device that supports the routing function. These network devices can be L3 switches, routers, or firewall.

- i** Unless otherwise specified, "OSPF" in the following descriptions refers to OSPFv2.

Protocols and Standards

RFC2328	This memo documents version 2 of the OSPF protocol. OSPF is a link-state routing protocol.
RFC 2370	This memo defines enhancements to the OSPF protocol to support a new class of link-state advertisements (LSA) called Opaque LSAs. Opaque LSAs provide a generalized mechanism to allow for the future extensibility of OSPF.

RFC3137	This memo describes a backward-compatible technique that may be used by OSPF (Open Shortest Path First) implementations to advertise unavailability to forward transit traffic or to lower the preference level for the paths through such a router.
RFC3623	This memo documents an enhancement to the OSPF routing protocol, whereby an OSPF router can stay on the forwarding path even as its OSPF software is restarted.
RFC3630	This document describes extensions to the OSPF protocol version 2 to support intra-area Traffic Engineering (TE), using Opaque Link State Advertisements.
RFC3682	The use of a packet's Time to Live (TTL) (IPv4) or Hop Limit (IPv6) to protect a protocol stack from CPU-utilization based attacks has been proposed in many settings.
RFC3906	This document describes how conventional hop-by-hop link-state routing protocols interact with new Traffic Engineering capabilities to create Interior Gateway Protocol (IGP) shortcuts.
RFC4576	This document specifies the necessary procedure, using one of the options bits in the LSA (Link State Advertisements) to indicate that an LSA has already been forwarded by a PE and should be ignored by any other PEs that see it.
RFC4577	This document extends that specification by allowing the routing protocol on the PE/CE interface to be the OSPF protocol.
RFC4750	This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based Internets. In particular, it defines objects for managing version 2 of the Open Shortest Path First Routing Protocol. Version 2 of the OSPF protocol is specific to the IPv4 address family.

4.2 Applications


Application	Description
Intra-Domain Interworking	OSPF runs within the AS, which is divided into several areas.
Inter-Domain Interworking	Several ASs are interconnected. OSPF runs within each AS, and Border Gateway Protocol (BGP) runs between ASs.

4.2.1 Intra-Domain Interworking

Scenario

OSPF runs within the AS. If the number of routers exceeds 40, it is recommended that the AS be divided into several areas. Generally, high-end devices featuring reliable performance and fast processing speed are deployed in a backbone area, and low-end or medium-range devices with relatively lower performance can be deployed in a normal area. All normal areas must be connected to the backbone area. It is recommended that a normal area allocated on the stub be configured as a stub area. As shown in Figure 4-1, the network is divided into four areas. Communication between these areas must go through the backbone area, that is area 0.

Figure 4-1 Division of the OSPF Areas



Remark
S

**A, B, C, D, E, and H are located in the backbone area, and are backbone routers.
Area 3 is configured as a stub area.**

Deployment

- OSPF runs on all routers within the AS to implement unicast routing.


4.2.2 Inter-Domain Interworking

Scenario

Several ASs are interconnected. OSPF runs within each AS, and BGP runs between ASs. Generally, OSPF and BGP learn the routing information from each other.

As shown in Figure 4-2, unicast routing is implemented within AS 100 and AS 200 using OSPF, and between the two ASs using BGP.

Figure 4-2 Interworking Between OSPF and BGP



Remarks	OSPF and BGP run concurrently on Router A and Router D.
----------------	---

Deployment

- OSPF runs within AS 100 and AS 200 to implement unicast routing.
- BGP runs between the two ASs to implement unicast routing.

4.3 Features

Basic Concepts

Routing Domain

All routers in an AS must be interconnected and use the same routing protocol. Therefore, the AS is also called routing domain.

An AS on which OSPF runs is also called OSPF routing domain, or OSPF domain for short.

OSPF Process

OSPF supports multiple instances, and each instance corresponds to an OSPF process.

One or more OSPF processes can be started on a router. Each OSPF process runs OSPF independently, and the processes are mutually isolated.

The process ID takes effect only on the local router, and does not affect exchange of OSPF packets on adjacent interfaces.

RouterID

The router ID uniquely identifies a router in an OSPF domain. Router IDs of any two routers cannot be the same.

If multiple OSPF processes exist on a router, each OSPF process uses one router ID. Router IDs of any two OSPF processes cannot be the same.


Area

OSPF supports multiple areas. An OSPF domain is divided into multiple areas to ease the computing pressure of a large-scale network.

An area is a logical group of routers, and each group is identified by an area ID. The border between areas is a router. A router may belong to one area or multiple areas. One network segment (link) can belong to only one area, or each OSPF-enabled interface must belong to a specified area.

Area 0 is the backbone area, and other areas are normal areas. Normal areas must be directly connected to the backbone area.

Figure 4-3 Division of the OSPF Areas



OSPF Router


The following types of routers are defined in OSPF, and assigned with different responsibilities:

- Internal router
All interface of an interval router belong to the same OSPF area. As shown in Figure 1-3, A, C, F, G, I, M, J, K, and L are internal routers.
- Area border router (ABR)
An ABR is used to connect the backbone area with a normal area. An ABR belongs to two or more areas, and one of the areas must be the backbone area. As shown in Figure 1-3, B, D, E, and H are ABRs.
- Backbone router
A backbone router has at least one interface that belongs to the backbone area. All ABRs and all routers in area 0 are backbone routers. As shown in Figure 4-3, A, B, C, D, E, and H are backbone routers.
- AS boundary router (ASBR)
An ASBR is used to exchange routing information with other ASs. An ASBR is not necessarily located on the border of an AS. It may be a router inside an area, or an ABR. As shown in Figure 1-3, A is an ASBR.

Virtual Link


OSPF supports virtual links. A virtual link is a logical link that belongs to the backbone area. It is used to resolve the problems such as a discontinuous backbone area or a failure to directly connect a normal area to the backbone area on the physical network. A virtual link supports traversal of only one normal area, and this area is called transit area. Routers on both ends of a virtual link are ABRs.

Figure 4-4 Discontinuous Backbone Area on the Physical Network



As shown in Figure 4-4, a virtual link is set up between A and B to connect two separated area 0s. Area 1 is a transit area, and A and B are ABRs of area 1.

Figure 4-5 Failure to Directly Connect a Normal Area to the Backbone Area on the Physical Network



As shown in Figure 4-4, a virtual link is set up between A and B to extend area 0 to B so that area 0 can be directly connected to area 2 on B. Area 1 is a transit area, A is an ABR of area 1, and B is an ABR of area 0 and area 2.

LSA

OSPF describes the routing information by means of Link State Advertisement (LSA).

LSA Type	Description
Router-LSA(Type 1)	This LSA is originated by every router. It describes the link state and cost of the router, and is advertised only within the area where the originating router is located.
Network-LSA(Type 2)	This LSA is originated by a designated routers (DR) on the NBMA network. It describes the link state in the current network segment, and is advertised only within the area where the DR is located.
Network-summary-LSA(Type 3)	This LSA is originated by an ABR. It describes a route to another area, and is advertised to areas except totally stub areas or Not-So-Stubby Area (NSSA) areas.
ASBR-summary-LSA(Type 4)	This LSA is originated by an ABR. It describes a route to an ASBR, and is advertised to areas except areas where the ASBR is located.
AS-external-LSA(Type 5)	This LSA is originated by an ABR. It describes a route to a destination outside the AS, and is advertised to all areas except the stub and NSSA areas.
NSSA LSA(Type 7)	This LSA is originated by an ABR. It describes a route to a destination outside the AS, and is advertised only within the NASSA areas.

Opaque LSA(Type 9/Type 10/Type 11)	<p>Opaque LSAs provide a generalized mechanism to allow for the future extensibility of OSPF, wherein,</p> <ul style="list-style-type: none"> ▪ Type 9 LSAs are only advertised within the network segment where interfaces resides. The Grace LSA used to support graceful restart (GR) is one of Type 9 LSAs. ▪ Type 10 LSAs are advertised within an area. The LSA used to support Traffic Engineering (TE) is one of Type 10 LSAs. ▪ Type 11 LSAs are advertised within an AS. At present, there are no application examples of Type 11 LSAs.
------------------------------------	--

i Stub areas, NSSA areas, totally stub areas, and totally NSSA areas are special forms of normal areas and help reduce the load of routers and enhance reliability of OSPF routes.

OSPF Packet

The following table lists the protocol packets used by OSPF. These OSPF packets are encapsulated in IP packets and transmitted in multicast or unicast mode.

Packet Type	Description
Hello	Hello packets are sent periodically to discover and maintain OSPF neighbor relationships.
Database Description (DD)	DD packets carry brief information about the local Link-State Database (LSDB) and are used to synchronize the LSDBs between OSPF neighbors.
Link State Request (LSR)	LSR packets are used to request the required LSAs from neighbors. LSR packets are sent only after DD packets are exchanged successfully between OSPF neighbors.
Link State Update (LSU)	LSU packets are used to send the required LSAs to peers.
Link State Acknowledgment (LSAck)	LSAck packets are used to acknowledge the received LSAs.

Overview

Feature	Description
Link-State Routing Protocols	Run OSPF on the router to obtain routes to different destinations on the network.

OSPF Route Management	Plan or optimize OSPF routes through manual configuration to implement management of OSPF routes.
Enhanced Security and Reliability	Use functions such as authentication and bidirectional forwarding detection (BFD) correlation to enhance security, stability, and reliability of OSPF.
Network Management	Use functions such as the management information base (MIB) and Syslog to facilitate OSPF management.

4.3.1 Link-State Routing Protocols

OSPF is a type of link-state routing protocols. Its working process is as follows:

- Neighbor discovery → Bidirectional communication
An OSPF neighbor relationship is set up between adjacent routers, and bidirectional communication is maintained.
- Database synchronization → Full adjacency
A router uses LSAs to advertise all its link states. LSAs are exchanged between neighbors and the link state database (LSDB) is synchronized to achieve full adjacency.
- Shortest Path Tree (SPT) computation → Formation of a routing table
The router computes the shortest path to each destination network based on the LSDB and forms an OSPF routing table.

Working Principle

Neighbor Discovery → Bidirectional Communication

Routers send Hello packets through all OSPF-enabled interfaces (or virtual links). If Hello packets can be exchanged between two routers, and parameters carried in the Hello packets can be successfully negotiated, the two routers become neighbors. Routers that are mutually neighbors find their own router IDs from Hello packets sent from neighbors, and bidirectional communication is set up.

A Hello packet includes, but is not limited to, the following information:

- Router ID of the originating router
- Area ID of the originating router interface (or virtual link)
- Subnet mask of the originating router interface (or virtual link)
- Authentication information of the originating router interface (or virtual link)
- Hello interval of the originating router interface (or virtual link)
- Neighbor dead interval of the originating router interface (or virtual link)
- Priority of the originating router interface (used for DR/BDR election)
- IP addresses of the DR and Backup Designated Router (BDR)

- Router ID of the neighbor of the originating router

Database Synchronization → Full Adjacency

After bidirectional communication is set up between neighbor routers, the DD, LSR, LSU, and LSAck packets are used to exchange LSAs and set up the adjacency. The brief process is as follows:

- A router generates an LSA to describe all link states on the router.
- The LSA is exchanged between neighbors. When a router receives the LSA from its neighbor, it copies the LSA and saves the copy in the local LSDB, and then advertises the LSA to other neighbors.
- When the router and its neighbors obtain the same LSDB, full adjacency is achieved.

i OSPF will be very quiet without changes in link costs or network addition or deletion. If any change takes place, the changed link states are advertised to quickly synchronize the LSDB.

SPT Computation → Formation of a Routing Table

After the complete LSDB is obtained from the router, the Dijkstra algorithm is run to generate an SPT from the local router to each destination network. The SPT records the destination networks, next-hop addresses, and costs. OSPF generates a routing table based on the SPT.

If changes in link costs or network addition or deletion take place, the LSDB will be updated. The router again runs the Dijkstra algorithm, generates a new SPT, and updates the routing table.

i The Dijkstra algorithm is used to find a shortest path from a vertex to other vertices in a weighted directed graph.

OSPF Network Types

A router does not necessarily need to exchange LSAs with every neighbor and set up an adjacency with every neighbor. To improve efficiency, OSPF classifies networks that use various link layer protocols into five types so that LSAs are exchanged in different ways to set up an adjacency:

- **Broadcast**
Neighbors are discovered, and the DR and BDR are elected. The DR (or BDR) exchanges LSAs with all other routers to set up an adjacency. Except the DR and BDR, all other routers do not exchange LSAs with each other, and the adjacency is not set up. Ethernet and fiber distributed data interface (FDDI) belong to the broadcast network type by default.
- **Non-broadcast multiple access (NBMA)**
Neighbors are manually configured, and the DR and BDR are elected. The DR (or BDR) exchanges LSAs with all other routers to set up an adjacency. Except the DR and BDR, all other routers do not exchange LSAs with each other, and the adjacency is not set up. X.25, frame relay, and ATM belong to NBMA networks by default.
- **Point-to-point (P2P)**
Neighbors are automatically discovered, and the DR or BDR is not elected. LSAs are exchanged between routers at both ends of the link, and the adjacency is set up.

PPP, HDLC, and LAPB belongs to the P2P network type by default.

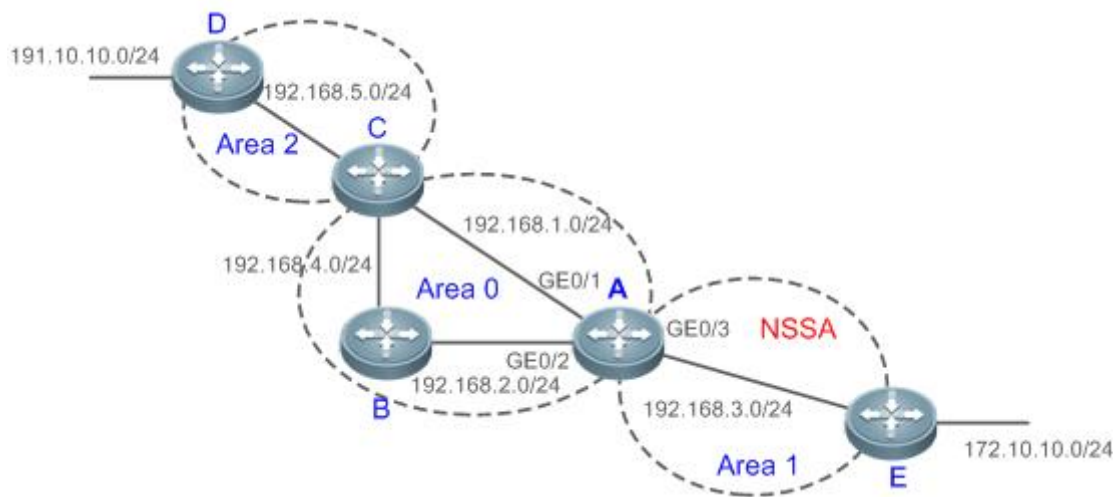
- Point-to-multipoint(P2MP)

Neighbors are automatically discovered, and the DR or BDR is not elected. LSAs are exchanged between any two routers, and the adjacency is set up. Networks without any link layer protocol belong to the P2MP network type by default. P2MP broadcast

Neighbors are manually configured, and the DR or BDR is not elected. LSAs are exchanged between any two routers, and the adjacency is set up. Networks without any link layer protocol belong to the P2MP network type by default.

OSPF Route Types

Figure 4-6



Display the OSPF routes (marked in red) in the routing table of Router A.

```
A#show ip route
```

Codes: C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default

Gateway of last resort is no set

```
O N2 172.10.10.0/24 [110/20] via 192.168.3.2, 00:01:00,GigabitEthernet 0/3
```

```
O E2 191.10.10.0/24 [110/20] via 192.168.1.2, 01:11:26,GigabitEthernet 0/1
```

- C 192.168.1.0/24 is directly connected,GigabitEthernet 0/1
- C 192.168.1.1/32 is local host.
- C 192.168.2.0/24 is directly connected,GigabitEthernet 0/2
- C 192.168.2.1/32 is local host.
- C 192.168.3.0/24 is directly connected,GigabitEthernet 0/3
- C 192.168.3.1/32 is local host.
- O 192.168.4.0/24 [110/2] via 192.168.2.2, 00:00:02,GigabitEthernet 0/2
- O IA 192.168.5.0/24 [110/3] via 192.168.1.2, 00:01:02,GigabitEthernet 0/1

A mark is displayed in front of each OSPF route to indicate the type of the route. There are six types of OSPF routes:

- O: Intra-area route
This type of route describes how to arrive at a destination network in the local area. The cost of this type of route is equal to the cost of the route from the local router to the destination network.
 - IA: Inter-area route
This type of route describes how to arrive at a destination network in another area. The cost of this type of route is equal to the cost of the route from the local router to the destination network.
 - E1: Type 1 external route
This type of route describes how to arrive at a destination network outside the AS. The cost of this type of route is equal to the cost of the route from the local router to the ASBR plus the cost of the route from the ASBR to the destination network. This type of route does not exist on routers in the stub or NSSA area.
 - E2: Type 2 external route
This type of route describes how to arrive at a destination network outside the AS. The cost of this type of route is equal to the cost of the route from the ASBR to the destination network. This type of route does not exist on routers in the stub or NSSA area.
 - N1: Type 1 external route of the NSSA area
This type of route describes how to arrive at a destination network outside the AS through the ASBR in the NSSA area. The cost of this type of route is equal to the cost of the route from the local router to the ASBR plus the cost of the route from the ASBR to the destination network. This type of route exists only on routers in the NSSA area.
 - N2: Type 2 external route of the NSSA area
This type of route describes how to arrive at a destination network outside the AS through the ASBR in the NSSA area. The cost of this type of route is equal to the cost of the route from the ASBR to the destination network. This type of route exists only on routers in the NSSA area.
- i** Reliability of E2 and N2 routes is poor. OSPF believes that the cost of the route from the ASBR to a destination outside an AS is far greater than the cost of the route to the ASBR within the AS. Therefore, when the route cost is computed, only the cost of the route from the ASBR to a destination outside an AS is considered.

Related Configuration

Enabling OSPF

OSPF is disabled by default.

Run the **router ospf 1** command to create an OSPF process on the router.

Run the **network area** command to enable OSPF on the interface and specify the area ID.

Run the **area virtual-link** command to create a virtual link on the router. The virtual link can be treated as a logical interface.

Router ID

By default, the OSPF process elects the largest IP address among the IP addresses of all the loopback interfaces as the router ID. If the loopback interfaces configured with IP addresses are not available, the OSPF process elects the largest IP address among the IP addresses of all the loopback interfaces as the router ID.

Alternatively, you can run the **router-id** command to manually specify the router ID.

Protocol Control Parameters

Run the **ip ospf hello-interval** command to modify the Hello interval on the interface. The default value is 10s (or 30s for NBMA networks).

Run the **ip ospf dead-interval** command to modify the neighbor dead interval on the interface. The default value is four times the Hello interval.

Use the **poll-interval** parameter in the **neighbor** command to modify the neighbor polling interval on the NBMA interface. The default value is 120s.

Run the **ip ospf transmit-delay** command to modify the LSU packet transmission delay on the interface. The default value is 1s.

Run the **ip ospf retransmit-interval** command to modify the LSU packet retransmission interval on the interface. The default value is 5s.

Use the **hello-interval** parameter in the **area virtual-link** command to modify the Hello interval on the virtual link. The default value is 10s.

Use the **dead-interval** parameter in the **area virtual-link** command to modify the neighbor dead interval on the virtual link. The default value is four times the Hello interval.

Use the **transmit-delay** parameter in the **area virtual-link** command to modify the LSU packet transmission delay on the virtual link. The default value is 1s.

Use the **retransmit-interval** parameter in the **area virtual-link** command to modify the LSU packet retransmission interval on the virtual link. The default value is 5s.

Run the **timers throttle lsa all** command to modify parameters of the exponential backoff algorithm that generates LSAs. The default values of these parameters are 0 ms, 5000 ms, and 5000 ms.

Run the **timers pacing lsa-group** command to modify the LSA group update interval. The default value is 30s.

Run the **timers pacing lsa-transmit** command to modify the LS-UPD packet sending interval and the number of sent LS-UPD packets. The default values are 40 ms and 1.

Run the **timers lsa arrival** command to modify the delay after which the same LSA is received. The default value is 1000 ms.

Run the **timers throttle spf** command to modify the SPT computation delay, minimum interval between two SPT computations, and maximum interval between two SPT computations. The default values are 1000 ms, 5000 ms, and 10000 ms.

OSPF Network Types

By default, Ethernet and FDDI belong to the broadcast type, X.25, frame relay, and ATM belong to the NBMA type, and PPP, HDLC, and LAPB belong to the P2P type.

Run the **ip ospf network** command to manually specify the network type of an interface.

Run the **neighbor** command to manually specify a neighbor. For the NBMA and P2MP non-broadcast types, you must manually specify neighbors.

Run the **ip ospf priority** command to adjust the priorities of interfaces, which are used for DR/BDR election. The DR/BDR election is required for the broadcast and NBMA types. The router with the highest priority wins in the election, and the router with the priority of 0 does not participate in the election. The default value is 1.

4.3.2 OSPF Route Management

Plan or optimize OSPF routes through manual configuration to implement management of OSPF routes.

Working Principle

(Totally) Stub Area and (Totally)NSSA Area

The (totally) stub and (totally)NSSA areas help reduce the protocol interaction load and the size of the routing table.

- If an appropriate area is configured as a (totally) stub or NSSA area, advertisement of a large number of Type 5 and Type 3 LSAs can be avoided within the area.

Area	Type1 and Type2 LSAs	Type 3 LSA	Type 4 LSA	Type 5 LSA	Type 7 LSA
------	----------------------	------------	------------	------------	------------

Non (totally) stub area and NSSA area	Allowed	Allowed	Allowed	Allowed	Not allowed
Stub area	Allowed	Allowed (containing one default route)	Not allowed	Not allowed	Not allowed
Totally stub area	Allowed	Only one default route is allowed.	Not allowed	Not allowed	Not allowed
NSSA area	Allowed	Allowed (containing one default route)	Allowed	Not allowed	Allowed
Totally NSSA area	Allowed	Only one default route is allowed.	Allowed	Not allowed	Allowed

- i The ABR uses Type 3LSAs to advertise a default route to the (totally) stub or NSSA area.
- i The ABR converts Type 7 LSAs in the totally NSSA area to Type 5LSAs, and advertise Type 5LSAs to the backbone area.
- If an area is appropriately configured as a (totally) stub area or an NSSA area, a large number of E1, E2, and IA routes will not be added to the routing table of a router in the area.

Area	Routes Available in the Routing Table of a Router Inside the Area
Non (totally) stub area and NSSA area	O: a route to a destination network in the local area IA: a route to a destination network in another area E1 or E2: a route or default route to a destination network segment outside the AS (via any ASBR in the AS)
Stub area	O: a route to a destination network in the local area IA: a route or a default route to a destination network in another area
Totally stub area	O: a route to a destination network in the local area IA: a default route
NSSA area	O: a route to a destination network in the local area IA: a route or a default route to a destination network in another area N1 or N2: a route or default route to a destination network segment outside the AS (via any ASBR in the local area)

Totally NSSA area	<p>O: a route to a destination network in the local area</p> <p>IA: a default route</p> <p>N1 or N2: a route or default route to a destination network segment outside the AS (via any ASBR in the local area)</p>
-------------------	--

Route Redistribution

Route redistribution refers to the process of introducing routes of other routing protocols, routes of other OSPF processes, static routes, and direct routes that exist on the device to an OSPF process so that these routes can be advertised to neighbors using Type 5 and Type 7 LSAs. A default route cannot be introduced during route redistribution.

Route redistribution is often used for interworking between ASs. You can configure route redistribution on an ASBR to advertise routes outside an AS to the interior of the AS, or routes inside an AS to the exterior of the AS.

Default Route Introduction

By configuring a command on an ASBR, you can introduce a default route to an OSPF process so that the route can be advertised to neighbors using Type 5 and Type 7 LSAs.

Default route introduction is often used for interworking between ASs. One default route is used to replace all the routes outside an AS.

Route Summarization

Route summarization is a process of summarizing routing information with the same prefix into one route, and advertising the summarized route (replacing a large number of individual routes) to neighbors. Route summarization helps reduce the protocol interaction load and the size of the routing table.

By default, the ABR advertises inter-area routing information by using Type3 LSAs within a network segment, and advertises redistributed routing information by using Type 5 and Type 7 LSAs. If continuous network segments exist, it is recommended that you configure route summarization.

When configuring route summarization, the summarization range may exceed the actual network scope of routes. If data is sent to a network beyond the summarization range, a routing loop may be formed and the router processing load may increase. To prevent these problems, the ABR or ASBR automatically adds a discard route to the routing table. This route will not be advertised.

Route Filtering

OSPF supports route filtering to ensure security and facilitate control when the routing information is being learned, exchanged, or used.

Using configuration commands, you can configure route filtering for the following items:

- Interface: The interface is prevented from sending routing information (any LSAs) or exchanging routing information (any LSAs) with neighbors.
- Routing information advertised between areas: Only the routing information that meets the filtering conditions can be advertised to another area (Type 3 LSAs).
- Routing information outside an AS: Only the routing information that meets the filtering conditions can be redistributed to the OSPF process (Type 5 and Type 7 LSAs).
- LSAs received by a router: In the OSPF routing table, only the routes that are computed based on the LSAs meeting the filtering conditions can be advertised.

Route Cost

If redundancy links or devices exist on the network, multiple paths may exist from the local device to the destination network. OSPF selects the path with the minimum total cost to form an OSPF route. The total cost of a path is equal to the sum of the costs of individual links along the path. The total cost of a path can be minimized by modifying the costs of individual links along the path. In this way, OSPF selects this path to form a route.

Using configuration commands, you can modify the link costs:

- Cost from an interface to a directly connected network segment and cost from the interface to a neighbor
- Cost from an ABR to the inter-area summarization network segment and cost from the ABR to the default network segment
- Cost from an ASBR to an external network segment and cost from the ASBR to the default network segment

i Both the cost and the metric indicate the cost and are not differentiated from each other.

OSPF Administrative Distance

The administrative distance (AD) evaluates reliability of a route, and the value is an integer ranging from 0 to 255. A smaller AD value indicates that the route is more trustworthy. If multiples exist to the same destination, the route preferentially selects a route with a smaller AD value. The route with a greater AD value becomes a floating route, that is, a standby route of the optimum route.

By default, the route coming from one source corresponds to an AD value. The AD value is a local concept. Modifying the AD value affects route selection only on the current router.

Route Source	Directly-Connected Network	Static Route	EBGP Route	OSPF Route	IS-IS Route	RIP Route	IBGP Route	Unreachable Route
Default AD	0	1	20	110	115	120	200	255

Related Configuration

Stub Area and NSSA Area

No stub or NSSA area is configured by default.

Run the **area stub** command to configure a specified area as a stub area.

Run the **area nssa** command to configure a specified area as an NSSA area.

- i** The backbone area cannot be configured as a stub or an NSSA area.
 - i** A transit area (with virtual links going through) cannot be configured as a stub or an NSSA area.
 - i** An area containing an ASBR cannot be configured as a stub area.
-

Route Redistribution and Default Route Introduction

By default, routes are not redistributed and the default route is not introduced.

Run the **redistribute** command to configure route redistribution.

Run the **default-information originate** command to introduce the default route.

After configuring route redistribution and default route introduction, the route automatically becomes an ASBR.

Route Summarization

By default, routes are not summarized. If route summarization is configured, a discard route will be automatically added.

Run the **arearange** command to summarize routes distributed between areas (Type 3 LSA) on the ABR.

Run the **summary-address** command to summarize redistributed routes (Type 5 and Type 7 LSAs) on the ASBR.

Run the **discard-route** command to add a discard route to the routing table.

Route Filtering

By default, routes are not filtered.

Run the **passive-interface** command to configure a passive interface. Routing information (any LSAs) cannot be exchanged on a passive interface.

Run the **ip ospfdatabase-filter all out** command to prohibit an interface from sending routing information (any LSAs).

Run the **area filter-list** command to filter routing information advertised between areas on the ABR. Only the routing information that meets the filtering conditions can be advertised to another area (Type 3 LSAs).

Use the **route-map** parameter in the **redistribute** command, or use the **distribute-list out** command to filter the external routing information of the AS on the ASBR. Only the routing information that meets the filtering conditions can be redistributed to the OSPF process (Type 5 and Type 7 LSAs).

Run the **distribute-list in** command to filter LSAs received by the router. In the OSPF routing table, only the routes that are computed based on the LSAs meeting the filtering conditions can be advertised.

Route Cost

- Cost from the interface to the directly-connected network segment (cost on the interface)
The default value is the auto cost. Auto cost = Reference bandwidth/Interface bandwidth
Run the **auto-costreference-bandwidth** command to set the reference bandwidth of auto cost. The default value is 100 Mbps.
Run the **ip ospf cost** command to manually set the cost of the interface. The configuration priority of this item is higher than that of the auto cost.
- Cost from the interface to a specified neighbor (that is, cost from the local device to a specified neighbor)
The default value is the auto cost.
Use the **cost** parameter in the **neighbor** command to modify the cost from the interface to a specified neighbor. The configuration priority of this item is higher than that of the cost of the interface.
This configuration item is applicable only to P2MP-type interfaces.
- Cost from the ABR to the inter-area summarization network segment (that is, the cost of the summarized inter-area route)
If OSPF routing is compatible with RFC1583, the default value is the minimum cost among all costs of the summarized links; otherwise, the default value is the maximum cost among all costs of the summarized links.
Run the **compatible rfc1583** command to make OSPF routing compatible with RFC1583. By default, OSPF routing is compatible with RFC1583.
Use the **cost** parameter in the **area range** command to modify the cost of inter-area route summarization.
- Cost from the ABR to the default network segment (that is, the cost of the default route that is automatically advertised by the ABR to the stub or NSSA areas)
The default value is 1.
Run the **area default-cost** command to modify the cost of the default route that the ABR automatically advertise to the stub or NSSA areas.
- Cost from the ASBR to an external network segment (that is, the metric of an external route)
By default, the metric of a redistributed BGP route is 1, the metric of other types of redistributed routes is 20, and the route type is Type 2 External.
Run the **default-metric** command to modify the default metric of the external route.
Use the **metric**, **metric-type** and **route-map** parameters in the **redistribute** command to modify the metric and route type of the external route.

- Cost from the ASBR to the default network segment (that is, the metric of the default route that is manually introduced)
By default, the metric is 1, and the route type is Type 2 External. Use the **metric**, **metric-type** and **route-map** parameters in the **default-information originate** command to modify the metric and route type of the default route that is manually introduced. Use the **metric** and **metric-type** parameters of **default-information originate** in the **area nssa** command to modify the metric and type of the default route that is manually introduced to the NSSA area.
- Run the **max-metric router-lsa** command to set metrics of all routes advertised on the router to the maximum value. In this way, the total cost of any path that passes through this router will become very large, and the path can hardly become the shortest path.

OSPF Administrative Distance

By default, the OSPF AD is 110.

Run the **distance** command to set the AD of an OSPF route.

4.3.3 Enhanced Security and Reliability

Use functions such as authentication and BFD correlation to enhance security, stability, and reliability of OSPF.

Working Principle

Authentication

Authentication prevents routers that illegally access the network and hosts that forge OSPF packet from participating in the OSPF process. OSPF packets received on the OSPF interface (or at both ends of the virtual link) are authenticated. If authentication fails, the packets are discarded and the adjacency cannot be set up.

Enabling authentication can avoid learning unauthenticated or invalid routes, thus preventing advertising valid routes to unauthenticated devices. In the broadcast-type network, authentication also prevents unauthenticated devices from becoming designated devices, ensuring stability of the routing system and protecting the routing system against intrusions.

MTU Verification

On receiving a DD packet, OSPF checks whether the MTU of the neighbor interface is the same as the MTU of the local interface. If the MTU of the interface specified in the received DD packet is greater than the MTU of the interface that receives the packet, the adjacency cannot be set up. Disabling MTU verification can avoid this problem.

Source Address Verification

Generally, the source address of a packet received by OSPF is in the same network segment as the receiving interface. The addresses at both ends of a P2P link are configured separately and are not

necessarily in the same network segment. In this scenario, as the peer address information will be notified during the P2P link negotiation process, OSPF checks whether the source address of the packet is the address advertised by the peer during negotiation. If not, OSPF determines that the packet is invalid and discards this packet. In particular, OSPF does not verify the address of an unnumbered interface.

In some scenarios, the source address of a packet received by OSPF may not be in the same network segment as the receiving interface, and therefore OSPF address verification fails. For example, the negotiated peer address cannot be obtained on a P2P link. In this scenario, source address verification must be disabled to ensure that the OSPF adjacency can be properly set up.

Two-Way Maintenance

OSPF routers periodically send Hello packets to each other to maintain the adjacency. On a large network, a lot of packets may be sent or received, occupying too much CPU and memory. As a result, some packets are delayed or discarded. If the processing time of Hello packets exceeds the dead interval, the adjacency will be destroyed.

If the two-way maintenance function is enabled, in addition to the Hello packets, the DD, LSU, LSR, and LSAck packets can also be used to maintain the bidirectional communication between neighbors, which makes the adjacency more stable.

Concurrent Neighbor Interaction Restriction

When a router simultaneously exchanges data with multiple neighbors, its performance may be affected. If the maximum number of neighbors that concurrently initiate or accept interaction with the OSPF process, the router can interact with neighbors by batches, which ensures data forwarding and other key services.

Overflow

OSPF requires that routers in the same area store the same LSDB. The number of routers keeps increasing on the network. Some routers, however, cannot store so much routing information due to the limited system resources. The large amount of routing information may exhaust the system resources of routers, causing failures of the routers.

The overflow function limits the number of external routes in the LSDB to control the size of the LSDB.

When the number of external routes on a router exceeds the upper limit, the router enters the overflow state. The router deletes the external routes generated by itself from the LSDB, and does not generate new external routes. In addition, the router discards the newly received external routes. After the overflow state timer (5s) expires, if the number of external routes is lower than the upper limit, the normal state is restored.

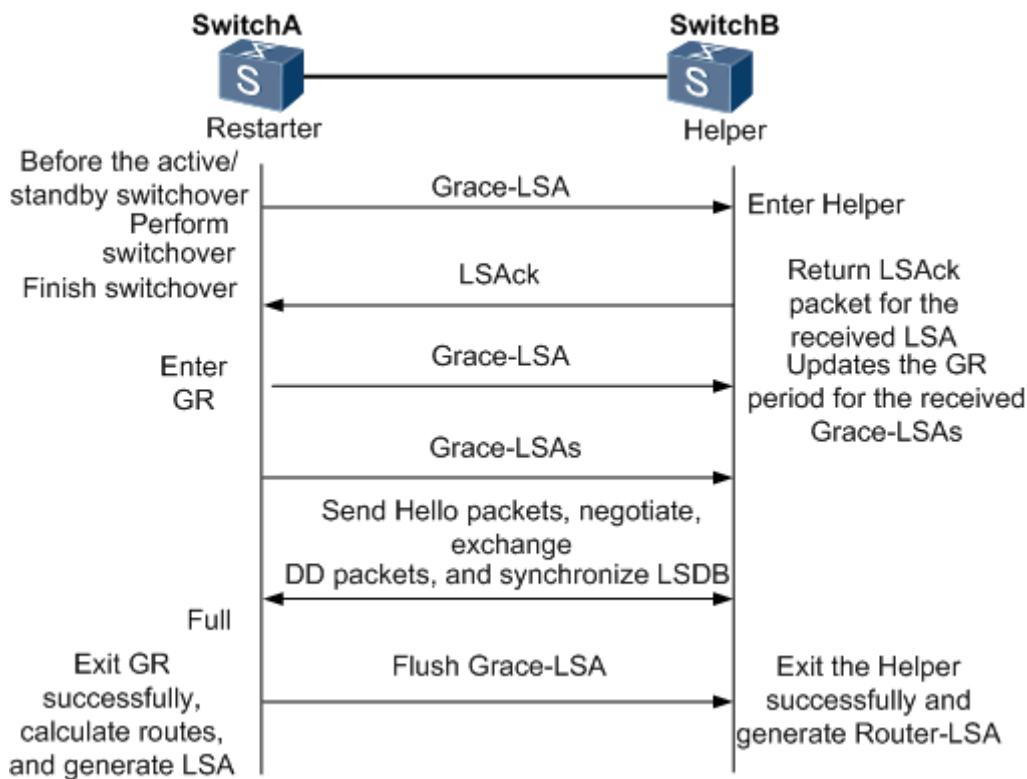
GR

The control and forwarding separated technology is widely used among routers. On a relatively stable network topology, when a GR-enabled router is restarted on the control plane, data forwarding can continue on the forwarding plane. In addition, actions (such as adjacency re-forming and route

computation) performed on the control plane do not affect functions of the forwarding plane. In this way, service interruption caused by route flapping can be avoided, thus enhancing reliability of the entire network.

Currently, the GR function is used only during active/standby switchover and system upgrade.

Figure 4-7 Normal OSPF GR Process



- The GR process requires collaboration between the restarter and the helper. The restarter is the router where GR occurs. The helper is a neighbor of the restarter.
- When entering or exiting the GR process, the restarter sends a Grace-LSA to the neighbor, notifying the neighbor to enter or exit the helper state.
- When the adjacency between the restarter and the helper reaches the Full state, the router can exit the GR process successfully.

NSR

During nonstop routing (NSR), OSPF-related information is backed up from the active supervisor module of a distributed device to the standby supervisor module. In this way, the device can automatically recover the link state and re-generate routes without the help of the neighbor devices

during the active/standby switchover. Information that should be backed up includes the adjacency and link state.

Fast Hello, BFD Correlation, and Fast Reroute

After a link fault occurs, OSPF senses the death of the neighbor only after a period of time (about 40s). Then, OSPF advertises the information and re-computes the SPT. During this period, traffic is interrupted.

- After the fast Hello function is enabled (that is, the neighbor dead interval is set to 1s), OSPF can sense the death of a neighbor within 1s once a link is faulty. This greatly accelerates route convergence and prevents traffic interruption.
- BFD is used to test connectivity between devices. A link fault can be detected in as short as 150 ms. After OSPF is correlated with BFD, OSPF can sense the death of a neighbor in as short as 150 ms once a link is faulty. This greatly accelerates route convergence and prevents traffic interruption.
- Fast reroute prepares a standby route for OSPF. Once the OSPF senses the death of a neighbor, the traffic is immediately switched over to the standby route, thus preventing traffic interruption.

iSPF

- The OSPF topology is area based. The SPF algorithm is run for independent computation in each area. The standard SPF algorithm re-computes the topology of the entire area each time even if only the leaf nodes change in the area topology.
- When computing the network topology, the incremental SPF (iSPF) corrects only the nodes on the SPT that are affected by the topological changes, and does not re-build the entire SPT. This can effectively ease the pressure on the router processors on a large network, especially when the network is not stable.

Related Configuration

OSPF Packet Authentication

By default, authentication is disabled.

- Run the **area authentication** command to enable the authentication function in the entire area so that the function takes effect on all interfaces in this area. If authentication is enabled in area 0, the function takes effect on the virtual link.
- Run the **ip ospf authentication** command to enable authentication on an interface. This configuration takes precedence over the area-based configuration.
- Run the **ip ospf authentication-key** command to set the text authentication key on an interface.
- Run the **ip ospf message-digest-key** command to set the message digest 5 (MD5) authentication key on an interface.
- Use the **authentication** parameter in the **area virtual-link** command to enable authentication at both ends of a virtual link. This configuration takes precedence over the area-based configuration.

- Use the **authentication-key** parameter in the **area virtual-link** command to set the text authentication key at both ends of a virtual link.
- Use the **message-digest-key** parameter in the **area virtual-link** command to set the MD5 authentication key at both ends of a virtual link.

MTU Verification

By default, MTU verification is disabled.

Run the **ip ospf mtu-ignore** command to disable MTU verification on an interface.

Source address verification

By default, source address verification is enabled on a P2P interface.

Run the **ip ospf source-check-ignore** command to disable source address verification on an interface.

Two-Way Maintenance

By default, bidirectional maintenance is enabled.

Run the **two-way-maintain** command to enable two-way maintenance.

Concurrent neighbor Interaction Restriction

Run the **max-concurrent-dd** command to modify the maximum number of neighbors that are concurrently interacting with the current OSPF process. The default value is 5.

Run the **ipv6 router ospf max-concurrent-dd** command to modify the maximum number of neighbors that are concurrently interacting with all OSPF processes on the router. The default value is 10.

Overflow

Run the **overflow memory-lack** command to allow the router to enter the overflow state when the memory is insufficient. By default, the router is allowed to enter the overflow state when the memory is insufficient.

Run the **overflow database** command to allow the router to enter the overflow state when the number of LSAs is too large. By default, the router is not allowed to enter the overflow state when the number of LSAs is too large.

Run the **overflow database external** command to allow the router to enter the overflow state when the number of external LSAs is too large. By default, the router is not allowed to enter the overflow state when the number of external-LSAs is too large.

GR

By default, the restarter function is disabled, and the helper function is enabled.

Run the **graceful-restart** command to configure the restarter function.

Run the **graceful-restart helper** command to configure the helper function.

NSR

By default, NSR is disabled.

Run the **nsr** command to enable NSR on the current OSPF process.

Fast Hello

By default, the neighbor dead interval on the interface is 40s.

Run the **ip ospf dead-intervalminimal hello-multiplier** command to enable the Fast Hello function on an interface, that is, the neighbor dead interval is 1s.

Correlating OSPFwith BFD

By default, OSPF is not correlated with BFD.

Run the **bfd interval min_rx multiplier** command to set the BFD parameters.

Run the **bfd all-interfaces** command to correlate OSPF with BFD on all interfaces.

Run the **ip ospf bfd** command to correlate OSPF with BFD on the current interface.

Fast Reroute

By default, fast reroute is disabled.

Run the **fast-reroute route-map** command to enable fast reroute on an OSPF process so that the standby route defined in the route map can be used.

Run the **fast-reroute lfa** command to enable fast reroute on an OSPF process so that the standby route can be computed by using the loop-free standby path.

Run the **fast-reroute lfdownstream-paths** command to enable fast reroute on an OSPF process so that the standby route can be computed by using the downstream path.

Run the **set fast-reroute backup-interfacebackup-nexthop** command to define a standby route in the route map.

Run the **ip ospf fast-reroute protection** command to specify theloop-freealternate(LFA) protection mode of an interface.

Run the **ip ospf fast-reroute no-eligible-backup** command to prevent an interface from becoming a standby interface.

iSPF

By default, iSPF is disabled.

Run the **ispf enable** command to enable iSPF on the OSPF process.

4.3.4 Network Management

Use functions such as the MIB and Syslog to facilitate OSPF management.

Working Principle

MIB

MIB is the device status information set maintained by a device. You can use the management program to view and set the MIB node.

Multiple OSPF processes can be simultaneously started on a router, but the OSPF MIB can be bound with only one OSPF process.

Trap

A Trap message is a notification generated when the system detects a fault. This message contains the related fault information.

If the Trap function is enabled, the router can proactively send the Trap messages to the network management device.

Syslog

The Syslog records the operations (such as command configuration) performed by users on routers and specific events (such as network connection failures).

If the Syslog is allowed to record the adjacency changes, the network administrator can view the logs to learn the entire process that the OSPF adjacency is set up and maintained.

[Related Configuration](#)

MIB

By default, the MIB is bound with the OSPF process with the smallest process ID.

Run the **enable mib-binding** command to bind the MIB with the current OSPF process.

Trap

By default, all traps are disabled, and the device is not allowed to send OSPF traps.

Run the **enable traps** command to enable a specified trap for an OSPF process.


Run the **snmp-server enable traps ospf** command to allow the device to send OSPF traps.





SYSLOG



By default, the Syslog is allowed to record the adjacency changes.




Run the **log-adj-changes** command to allow the Syslog to record the adjacency changes.

4.4 Configuration




Configuration	Description and Command
	 (Mandatory) It is used to build an OSPF routing domain.




Configuring OSPF Basic Functions	routerospf	Creates an OSPF process.
	router-id	Configures a router ID.
	network area	Enables OSPF on an interface and specifies an area ID.
	area virtual-link	Creates a virtual link.
Setting the Network Type	 (Optional) The configurations are mandatory if the physical network is the X.25, frame relay, or ATM network.	
	ip ospf network	Defines the network type.
	neighbor	Specifies a neighbor.
	ip ospf priority	Configures the DR priority.
Configuring Route Redistribution and Default Route	 (Optional) The configurations are recommended if the OSPF routing domain is connected with an external network.	
	redistribute	Configures route redistribution.
	default-information originate	Introduces a default route.
Configuring Stub Area and NSSA Area	 (Optional) It is used to reduce interaction of routing information and the size of routing table, and enhance stability of routes.	
	areastub	Configures a stub area.
	areanssa	Configures an NSSA area.
Configuring Route Summarization	 (Optional) It is used to reduce interaction of routing information and the size of routing table, and enhance stability of routes.	
	area range	Summarizes routes that are advertised between areas.
	summary-address	Summarizes routes that are introduced through redistribution.

	discard-route	Adds a discard route to the routing table.
Configuring Route Filtering	 (Optional) It is used to manually control interaction of routing information and filter available OSPF routes.	
	passive-interface	Configures a passive interface.
	ip ospfdatabase-filter all out	Prohibits an interface from sending LSAs.
	area filter-list	Filters routes that are advertised between areas..
	distribute-list out	Filters routes that are introduced through redistribution.
	distribute-listin	Filters routes that are calculated based on the received LSAs.
Modifying Route Cost and AD	 (Optional) It is used to manually control the shortest route computed by OSPF and determine whether to select an OSPF route preferentially.	
	auto-costreference-bandwidth	Modifies the reference bandwidth of the auto cost.
	ip ospf cost	Modifies the cost in the outbound direction of an interface.
	areadefault-cost	Modifies the cost of the default route in a stub or an NSSA area.
	default-metric	Modifies the default metric of a redistributed route.
	max-metric router-lsa	Configures the maximum metric.
	compatible rfc1583	Enables the routing rules to be compatible with RFC1583.

	distance	Modifies the OSPF AD.
Enabling Authentication	 (Optional) It is used to prevent routers that illegally access the network and hosts that forge OSPF packets from participating in the OSPF protocol process.	
	areaauthentication	Enables authentication and sets the authentication mode in an area.
	ip ospf authentication	Enables authentication and sets the authentication mode on an interface.
	ip ospf authentication-key	Sets the text authentication key on an interface.
	ip ospfmessage-digest-keymd5	Sets the MD5 authentication key on an interface.
Enabling Overflow	 (Optional) It is used to prevent the problem that OSPF processes stop running due to over-consumption of the memory.	
	overflow memory-lack	Allows the router to enter the overflow state when the memory is insufficient.
	overflow database	Allows the router to enter the overflow state when the number of LSAs exceeds the preset limit.
	overflow database external	Allows the router to enter the overflow state when the number of external LSAs exceeds the preset limit.
Modifying the Maximum Number	 (Optional) It is used to prevent the problem of performance deterioration caused by over-consumption of the CPU.	

of Concurrent Neighbors	max-concurrent-dd	Modifies the maximum number of concurrent neighbors on the current OSPF process.
	router ospf max-concurrent-dd	Modifies the maximum number of concurrent neighbors on all OSPF processes.
Disabling Source Address Verification	⚠ (Optional) It is used to prevent the problem that the adjacency cannot be set up due to the failure to obtain the peer address.	
	ip ospf source-check-ignore	Disables source address verification on an interface.
Disabling MTU Verification	⚠ (Optional) It is used to prevent the problem that the adjacency cannot be set up due to MTU inconsistency on the neighbor interface.	
	ip ospf mtu-ignore	Disables MTU verification on an interface.
Enabling Two-Way Maintenance	⚠ (Optional) It is used to prevent termination of the adjacency due to the delay or loss of Hello packets.	
	two-way-maintain	Enables two-way maintenance.
Enabling GR	⚠ (Optional) It is used to retain OSPF routing forwarding during restart or active/standby switchover of the OSPF processes to prevent traffic interruption.	
	graceful-restart	Configures the restarter function.
	graceful-restart helper	Configures the helper function.
Enabling NSR	⚠ (Optional) It is used to retain OSPF routing forwarding during active/standby switchover of the OSPF processes to prevent traffic interruption.	
	nsr	Enables NSR.

Enabling Fast Hello	<p> (Optional) It is used to quickly discover the death of a neighbor to prevent traffic interruption when a link is faulty.</p>	
	ip ospf dead-intervalminimal hello-multiplier	Enabling Fast Hello.
Correlating OSPF with BFD	<p> (Optional) It is used to quickly discover the death of a neighbor to prevent traffic interruption when a link is faulty.</p>	
	bfd interval min_rx multiplier	Sets BFD parameters.
	bfd all-interfaces	Correlates OSPF with BFD on all interfaces.
	ip ospf bfd	Correlates OSPF with BFD on the current interface.
Enabling Fast Reroute	<p> (Optional) It is used to quickly switch over services to the standby route to prevent traffic interruption.</p>	
	fast-reroute route-map	Enables fast reroute on the OSPF process so that the standby route defined in the route map can be used.
	fast-reroute lfa	Enables fast reroute on an OSPF process so that the standby route can be computed by using the loop-free standby path.
	fast-reroute lfdownstream-paths	Enables fast reroute on an OSPF process so that the standby route can be computed by using the downstream path.
	set fast-reroute backup-interface backup-nexthop	Defines a standby route in the route map.
	ip ospf fast-reroute protection	Specifies the LFA protection mode of an interface.

	ip ospf fast-reroute no-eligible-backup	Prevents an interface from becoming a standby interface.
Enabling iSPF	 (Optional) It is used to enable the incremental topology computation to ease the pressure on the processor.	
	ispf enable	Enables iSPF on an OSPF process.
Configuring the Network Management Function	 (Optional) The configurations enable users to use the SNMP network management software to manage OSPF.	
	enable mib-binding	Binds the MIB with the current OSPF process.
	enable traps	Enables a specified trap for an OSPF process.
	snmp-server enable traps ospf	Allows the device to send OSPF traps.
	log-adj-changes	Allows the Syslog to record the adjacency changes.
Modifying Protocol Control Parameters	 (Optional) You are advised not to modify protocol control parameters unless necessary.	
	ip ospf hello-interval	Modifies the Hello interval.
	ip ospf dead-interval	Modifies the neighbor death interval.
	timers throttle lsa all	Modifies parameters of the exponential backoff algorithm that generates LSAs.
	timers throttle route inter-area	Modifies the inter-area route computation delay.

	timers throttle route ase	Modifies the external route computation delay.
	timers pacing lsa-group	Modifies the LSA group update interval.
	timers pacing lsa-transmit	Modifies the LS-UPD packet sending interval.
	ip ospf transmit-delay	Modifies the LSU packet transmission delay.
	ip ospf retransmit-interval	Modifies the LSU packet retransmission interval.
	timers lsa arrival	Modifies the delay after which the same LSA is received.
	timers throttlespf	Modifies the SPT computation timer.

4.4.1 Configuring OSPF Basic Functions

Configuration

Effect

- Set up an OSPF routing domain on the network to provide IPv4 unicast routing service for users on the network.

Notes

- Ensure that the IP unicast routing function is enabled, that is, **ip routing** is not disabled; otherwise, OSPF cannot be enabled.
- It is strongly recommended that you manually configure the router ID.
- After **ip ospf disable all** is configured, the interface neither sends or receives any OSPF packet, nor participates in OSPF computation even if the interface belongs to the network.

Configuration

Steps

Creating an OSPF Process

- Mandatory.
- The configuration is mandatory for every router.

Configuring a Router ID

- (Optional) It is strongly recommended that you manually configure the router ID.
- If the router ID is not configured, OSPF selects an interface IP address. If the IP address is not configured for any interface, or the configured IP addresses have been used by other OSPF instances, you must manually configure the router ID.

Enabling OSPF on an Interface and Specifying an Area ID

- Mandatory.
- The configuration is mandatory for every router.

Verification

- Run the **show ip route ospf** command to verify that the entries of the OSPF routing table are correctly loaded.
- Run the **ping** command to verify that the IPv4 unicast service is correctly configured.

Related Commands

Creating an OSPF Process

Command	router ospf <i>process-id</i> [vrf <i>vrf-name</i>]
Parameter Description	<i>process-id</i> : Indicates the OSPF process ID. If the process ID is not specified, the process ID is 1. <i>vrf-name</i> : Specifies the VPN routing and forwarding (VRF) to which the OSPF process belongs.
Command Mode	Global configuration mode
Usage Guide	Different OSPF processes are independent of each other, and can be treated as different routing protocols that run independently.

Configuring a Router ID

Command	router-id <i>router-id</i>
Parameter Description	<i>router-id</i> : Indicates the router ID to be configured. It is expressed in the IP address.

Command Mode	OSPF routing process configuration mode
Usage Guide	Different OSPF processes are independent of each other, and can be treated as different routing protocols that run independently. Each OSPF process uses a unique router ID.

Enabling OSPF on an Interface and Specifying an Area ID

Command	<code>network ip-address wildcard area area-id</code>
Parameter Description	<i>ip-address</i> : Indicates the IP address of the interface. <i>wildcard</i> : Indicates the IP address comparison mode. 0 indicates accurate matching, and 1 indicates that no comparison is performed. <i>area-id</i> : Indicates the ID of an OSPF area. An OSPF area is always associated with an address range. To facilitate management, you can use a subnet as the ID of an OSPF area.
Command Mode	OSPF routing process configuration mode
Usage Guide	By defining <i>ip-address</i> and <i>wildcard</i> , you can use one command to associate multiple interfaces with one OSPF area. To run OSPF on one interface, you must include the primary IP address of the interface in the IP address range defined by network area . If the IP address range defined by network area contains only the secondary IP address of the interface, OSPF does not run on this interface. If the interface address matches the IP address ranges defined in the network commands of multiple OSPF processes, the OSPF process that the interface is associated with is determined based on the best match method.

Creating a Virtual Link

Command	<code>area area-id virtual-link router-id [authentication [message-digest null]] [dead-interval { seconds / minimal hello-multiplier multiplier }] [hello-interval seconds] [retransmit-interval seconds] [transmit-delay seconds] [[authentication-key [0 7] key] [message-digest-key key-id md5 [0 7] key]]</code>
Parameter Description	<i>area-id</i> : Indicates the ID of the OSPF transit area. The area ID can be a decimal integer or an IP address. <i>router-id</i> : Indicates the ID of a neighbor router on the virtual link.

	<p>dead-interval<i>seconds</i>: Indicates the time that the neighbor is declared lost. The unit is second. The value ranges from 0 to 2,147,483,647. The setting of this parameter must be consistent with that on a neighbor.</p> <p>minimal: Indicates that the Fast Hello function is enabled to set the dead interval to 1s.</p> <p>hello-multiplier: Indicates the result of the dead interval multiple by the Hello interval in the Fast Hello function.</p> <p><i>multiplier</i>: Indicates the number of Hello packets sent per second in the Fast Hello function. The value ranges from 3 to 20.</p> <p>hello-interval <i>seconds</i>: Indicates the interval at which OSPF sends the Hello packet to the virtual link. The unit is second. The value ranges from 1 to 65,535. The setting of this parameter must be consistent with that on a neighbor.</p> <p>retransmit-interval <i>seconds</i>: Indicates the OSPF LSA retransmission time. The unit is second. The value ranges from 1 to 65,535.</p> <p>transmit-delay <i>seconds</i>: Indicates the delay after which OSPF sends the LSA. The unit is second. The value ranges from 1 to 65,535.</p> <p>authentication-key [0 7]<i>key</i>: Defines the key for OSPF plain text authentication.</p> <p>message-digest-key <i>key-id</i>md5 [0 7]<i>key</i>: Defines the key ID and key for OSPF MD5 authentication.</p> <p>authentication: Sets the authentication type to plain text authentication.</p> <p>message-digest: Sets the authentication type to MD5 authentication.</p> <p>null: Indicates that authentication is disabled.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>In the OSPF routing domain, all areas must be connected to the backbone area. If the backbone area is disconnected, a virtual link must be configured to connect to the backbone area; otherwise, network communication problems will occur. A virtual link must be created between two ABRs, and the area to which both ABRs belong is the transit area. A stub area or an NSSA area cannot be used as a transit area. A virtual link can also be used to connect other non-backbone areas.</p> <p>router-id is the ID of an OSPF neighbor router. If you are sure about the value of router-id, run the show ip ospf neighbor command to confirm the value. You can configure the loopback address as the router ID.</p> <p>The area virtual-link command defines only the authentication key of the virtual link. To enable OSPF packet authentication in the areas connected to the virtual link, you must run the area authentication command.</p>

OSPF supports the Fast Hello function.

After the OSPF Fast Hello function is enabled, OSPF finds neighbors and detects neighbor failures faster. You can enable the OSPF Fast Hello function by specifying the **minimal** and **hello-multiplier** keywords and the **multiplier** parameter. The **minimal** keyword indicates that the death interval is set to 1s, and **hello-multiplier** indicates the number of Hello packets sent per second. In this way, the interval at which the Hello packet is sent decreases to less than 1s.

If the Fast Hello function is configured for a virtual link, the Hello interval field of the Hello packet advertised on the virtual link is set to 0, and the Hello interval field of the Hello packet received on this virtual link is ignored.

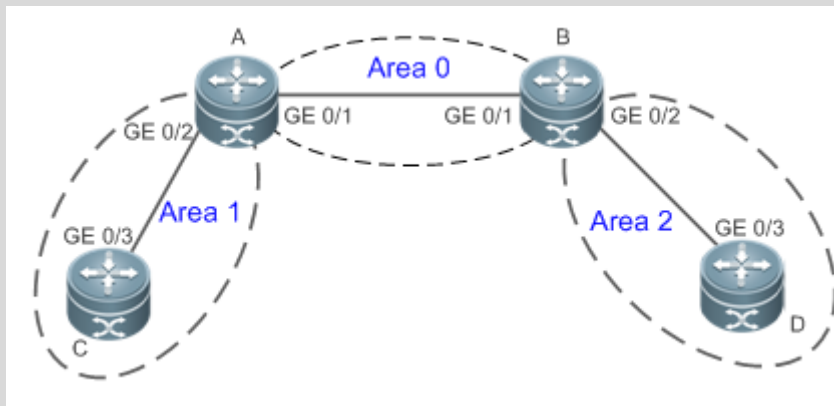
No matter whether the Fast Hello function is enabled, the death interval must be consistent and the **hello-multiplier** values can be inconsistent on routers at both ends of the virtual link. Ensure that at least one Hello packet can be received within the death interval.

Run the **show ip ospf virtual-links** command to monitor the death interval and Fast Hello interval configured for the virtual link.

The **dead-interval minimal hello-multiplier** and **hello-interval** parameters introduced for the Fast Hello function cannot be configured simultaneously.

Configuration Example

Scenario
Figure 4-8



Remark
S The interface IP addresses are as follows:
A: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1
B: GE 0/1 192.168.1.2 GE 0/2 192.168.3.1
C: GE 0/3 192.168.2.2
D: GE 0/3 192.168.3.2

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. ▪ Enable the IPv4 unicast routing function on all routers. (This function is enabled by default.) ▪ Configure the OSPF instances and router IDs on all routers. ▪ Enable OSPF on the interfaces configured on all routers.
A	<pre>A#configure terminal A(config)#interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#ip address 192.168.1.1 255.255.255.0 A(config-if-GigabitEthernet 0/1)#exit A(config)#interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)#ip address 192.168.2.1 255.255.255.0 A(config-if-GigabitEthernet 0/2)#exit A(config)#router ospf 1 A(config-router)#router-id192.168.1.1 A(config-router)#network 192.168.1.0 0.0.0.255 area 0 A(config-router)#network 192.168.2.0 0.0.0.255 area 1</pre>
B	<pre>B#configure terminal B(config)#interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)#ip address 192.168.1.2 255.255.255.0 B(config-if-GigabitEthernet 0/1)#exit B(config)#interface GigabitEthernet 0/2 B(config-if-GigabitEthernet 0/2)#ip address 192.168.3.1 255.255.255.0 B(config-if-GigabitEthernet 0/2)#exit B(config)#router ospf 1 B(config-router)#router-id192.168.1.2 B(config-router)#network 192.168.1.0 0.0.0.255 area 0 B(config-router)#network 192.168.3.0 0.0.0.255 area 2</pre>
C	<pre>C#configure terminal C(config)#interface GigabitEthernet 0/3 C(config-if-GigabitEthernet 0/3)#ip address 192.168.2.2 255.255.255.0 C(config-if-GigabitEthernet 0/3)#exit</pre>

	<pre>C(config)#router ospf 1 C(config-router)#router-id192.168.2.2 C(config-router)#network 192.168.2.0 0.0.0.255 area 1</pre>
D	<pre>D#configure terminal D(config)#interface GigabitEthernet 0/3 D(config-if-GigabitEthernet 0/3)#ip address 192.168.3.2 255.255.255.0 D(config-if-GigabitEthernet 0/3)#exit D(config)#router ospf 1 D(config-router)#router-id192.168.3.2 D(config-router)#network 192.168.3.0 0.0.0.255 area 2</pre>
Verification	<ul style="list-style-type: none"> ▪ Verify that the OSPF neighbors are correct on all routers. ▪ Verify that the routing table is correctly loaded on all routers. ▪ On Router D, verify that the IP address 192.168.2.2 can be pinged successfully.
A	<pre>A# show ip ospf neighbor OSPF process 1, 2 Neighbors, 2 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.2 1 Full/DR 00:00:40192.168.1.2 GigabitEthernet 0/1 192.168.2.2 1 Full/BDR00:00:34 192.168.2.2 GigabitEthernet 0/2 A# show ip route ospf O IA 192.168.3.0/24 [110/2] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1</pre>
B	<pre>B# show ip ospf neighbor OSPF process 1, 2 Neighbors, 2 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.1 1 Full/BDR 00:00:32 192.168.1.1 GigabitEthernet 0/1 192.168.3.2 1 Full/BDR00:00:30 192.168.3.2 GigabitEthernet 0/2 B# show ip route ospf O IA 192.168.2.0/24 [110/2] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1</pre>
C	<pre>C# show ip ospf neighbor</pre>

	<pre> OSPF process 1,1 Neighbors,1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.1 1 Full/BDR 00:00:32 192.168.2.1 GigabitEthernet 0/3 C# show ip route ospf O IA 192.168.1.0/24 [110/2] via 192.168.2.1, 00:19:05, GigabitEthernet 0/3 O IA 192.168.3.0/24 [110/3] via 192.168.2.1, 00:19:05, GigabitEthernet 0/3 </pre>
D	<pre> D# show ip ospf neighbor OSPF process 1,1 Neighbors,1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.21 Full/BDR00:00:30 192.168.3.1 GigabitEthernet 0/3 D# show ip route ospf O IA 192.168.1.0/24 [110/2] via 192.168.3.1, 00:19:05, GigabitEthernet 0/3 O IA 192.168.2.0/24 [110/3] via 192.168.3.1, 00:19:05, GigabitEthernet 0/3 D# ping 192.168.2.2 Sending 5, 100-byte ICMP Echoes to 192.168.2.2, timeout is 2 seconds: < press Ctrl+C to break > !!!! Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms. </pre>

Common Errors

- OSPF cannot be enabled because the IP unicast routing function is disabled.
- The network segment configured by the **network** command does not include the interface IP addresses.
- The area IDs enabled on adjacent interfaces are inconsistent.
- The same router ID is configured on multiple routers, resulting in a router ID conflict.
- The same interface IP address is configured on multiple routers, resulting in a running error of the OSPF network.

4.4.2 Setting the Network Type

Configuration

Effect

- Run OSPF to provide the IPv4 unicast routing service if the physical network is X.25, frame relay, or ATM.

Notes

- The OSPF basic functions must be configured.
- The broadcast network sends OSPF packets in multicast mode. Neighbors are automatically discovered, and the DR/BDR election is required.
- The P2P network sends OSPF packets in multicast mode. Neighbors are automatically discovered.
- The NBMA network sends OSPF packets in unicast mode. Neighbors must be manually specified, and the DR/BDR election is required.
- The P2MP network (without the **non-broadcast** parameter) sends OSPF packets in multicast mode. Neighbors are automatically discovered.
- The P2MP network (with the **non-broadcast** parameter) sends OSPF packets in unicast mode. Neighbors must be manually specified.

Configuration

Steps

Configuring the Interface Network Type

- Optional.
- The configuration is required on routers at both ends of the link.

Configuring Neighbors

- (Optional) If the interface network type is set to NBMA or P2MP (with the **non-broadcast** parameter), neighbors must be configured.
- Neighbors are configured on routers at both ends of the NBMA or P2MP (with the **non-broadcast** parameter) network.

Configuring the Interface Priority

- (Optional) You must configure the interface priority if a router must be specified as a DR, or a router cannot be specified as a DR.
- Configure the interface priority on a router that must be specified as a DR, or cannot be specified as a DR.

Verification

- Run the **show ip ospf interface** command to verify that the network type of each interface is correct.

Related

Commands

Configuring the Interface Network Type

Command	ip ospf network { broadcast non-broadcast point-to-multipoint [non-broadcast] point-to-point }
Parameter Description	<p>broadcast: Sets the interface network type to broadcast.</p> <p>non-broadcast: Sets the interface network type to non-broadcast.</p> <p>point-to-multipoint [non-broadcast]: Sets the interface network type to P2MP. If the interface does not have the broadcast capability, the non-broadcast parameter must be available.</p> <p>point-to-point: Sets the interface network type to P2P.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>The broadcast type requires that the interface must have the broadcast capability.</p> <p>The P2P type requires that the interfaces are interconnected in one-to-one manner.</p> <p>The NBMA type requires full-meshed connections, and all interconnected routers can directly communicate with each other.</p> <p>The P2MP type does not raise any requirement.</p>

Configuring Neighbors

Command	neighbor ip-address [poll-intervalseconds] [prioritypriority] [cost cost]
Parameter Description	<p><i>ip-address</i>: Indicates the IP address of the neighbor interface.</p> <p>poll-intervalseconds: Indicates the neighbor polling interval. The unit is second. The value ranges from 0 to 2,147,483,647. This parameter is applicable only to the NBMA interface.</p> <p>prioritypriority: Indicates the neighbor priority. The value ranges from 0 to 255. This parameter is applicable only to the NBMA interface.</p> <p>costcost: Indicates the cost required to reach each neighbor. There is no default value. The value ranges from 0 to 65,535. This parameter is applicable only to the P2MP interface.</p>
Command Mode	OSPF routing process configuration mode

Usage Guide	<p>Neighbors must be specified for the NBMA or P2MP (non-broadcast) interfaces. The neighbor IP address must be the primary IP address of this neighbor interface.</p> <p>If a neighbor router becomes inactive on the NBMA network, OSPF still sends Hello packets to this neighbor even if no Hello packet is received within the router death time. The interval at which the Hello packet is sent is called polling interval. When running for the first time, OSPF sends Hello packets only to neighbors whose priorities are not 0. In this way, neighbors with priorities set to 0 do not participate in the DR/BDR election. After a DR/BDR is elected, the DR/BDR sends the Hello packets to all neighbors to set up the adjacency.</p> <p>The P2MP (non-broadcast) network cannot dynamically discover neighbors because it does not have the broadcast capability. Therefore, you must use this command to manually configure neighbors for the P2MP (non-broadcast) network. In addition, you can use the cost parameter to specify the cost to reach each neighbor on the P2MP network.</p>
-------------	---

Configuring the Interface Priority

Command	ip ospf priority <i>priority</i>
Parameter Description	<i>priority</i> : Indicates the OSPF priority of an interface. The value ranges from 0 to 255.
Command Mode	Interface configuration mode
Usage Guide	<p>The OSPF interface priority is contained in the Hello packet. When the DR/BDR election occurs on the OSPF broadcast network, the router with the highest priority becomes the DR or BDR. If the priorities are the same, the router with the largest router ID becomes the DR or BDR. A router with the priority set to 0 does not participate in the DR/BDR election.</p> <p>This command is applicable only to the OSPF broadcast and NBMA interfaces.</p>

Configuration Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Setting the Interface Network Type to P2MP

<p>Scenario Figure 4-9</p>	<div data-bbox="349 210 885 598" data-label="Diagram"> </div> <div data-bbox="349 619 1453 840" data-label="Text"> <p>Remark S The interface IP addresses are as follows: A: S1/0 192.168.1.2 B: S1/0 192.168.1.3 C: S1/0 192.168.1.4</p> </div>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Set the interface network type to P2MP on all routers.
<p>A</p>	<pre>A#configure terminal A(config)# interface Serial1/0 A(config-Serial1/0)# encapsulation frame-relay A(config-Serial1/0)# ip ospf network point-to-multipoint</pre>
<p>B</p>	<pre>B#configure terminal B(config)# interface Serial1/0 B(config-Serial1/0)# encapsulation frame-relay B(config-Serial1/0)# ip ospf network point-to-multipoint</pre>
<p>C</p>	<pre>C#configure terminal C(config)# interface Serial1/0 C(config-Serial1/0)# encapsulation frame-relay C(config-Serial1/0)# ip ospf network point-to-multipoint</pre>
<p>Verification</p>	<p>Verify that the interface network type is P2MP.</p> <pre>A# show ip ospf interface Serial1/0 Serial1/0 is up, line protocol is up</pre>

```
Internet Address 192.168.1.2/24, Ifindex 2, Area 0.0.0.1, MTU 1500
Matching network config: 192.168.1.0/24
Process ID 1, Router ID 192.168.1.2, Network Type POINTOMULTIPOINT, Cost: 1
Transmit Delay is 1 sec, State Point-To-Point
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:02
Neighbor Count is 1, Adjacent neighbor count is 0
Crypt Sequence Number is 4787
Hello received 465 sent 466, DD received 8 sent 8
LS-Req received 2 sent 2, LS-Upd received 8 sent 21
LS-Ack received 14 sent 7, Discarded 3
```

Common Errors

- The network types configured on interfaces at two ends are inconsistent, causing abnormal route learning.
- The network type is set to NBMA or P2MP (with the **non-broadcast** parameter), but neighbors are not specified.

4.4.3 Configuring Route Redistribution and Default Route

Configuration Effect

- In the OSPF domain, introduce a unicast route to other AS domains so that the unicast routing service to other AS domains can be provided for users in the OSPF domain.
- In the OSPF domain, inject a default route to other AS domains so that the unicast routing service to other AS domains can be provided for users in the OSPF domain.

Notes

- The OSPF basic functions must be configured.

Configuration Steps

Configuring External Route Redistribution

- (Optional) This configuration is required if external routes of the OSPF domain should be introduced to an ASBR.

- This configuration is performed on an ASBR.

Generating a Default Route

- (Optional) This configuration is required if the default route should be introduced to an ASBR so that other routers in the OSPF domain access other AS domains through this ASBR by default.
- This configuration is performed on an ASBR.

Verification

- On a router inside the OSPF domain, run the **show ip route** command to verify that the unicast routes to other AS domains are loaded.
- On a router inside the OSPF domain, run the **show ip route** command to verify that the default route to the ASBR is loaded.
- Run the **ping** command to verify that the IPv4 unicast service to other AS domains is correct.

Related Commands

Configuring External Route Redistribution

Command	<code>redistribute {bgp connected isis[<i>area-tag</i>] ospfprocess-id rip arp-host static}[{ level-1 level-1-2 level-2}] [match {internal external [1 2]}] [metric <i>metric-value</i>] [metric-type {1 2}] [route-map <i>route-map-name</i>] [subnets] [tagtag-value]</code>
Parameter Description	<p>bgp: Indicates redistribution from BGP.</p> <p>connected: Indicates redistribution from direct routes.</p> <p>isis [<i>area-tag</i>]: Indicates redistribution from IS-IS. <i>area-tag</i> specifies the IS-IS instance.</p> <p>ospf <i>process-id</i>: Indicates redistribution from OSPF. <i>process-id</i> specifies an OSPF process. The value ranges from 1 to 65,535.</p> <p>rip: Indicates redistribution from RIP.</p> <p>static: Indicates redistribution from static routes.</p> <p>arp-host: Indicates redistribution from arp routes.</p> <p>level-1 level-1-2 level-2: Used only when IS-IS routes are redistributed. Only the routes of the specified level are redistributed. By default, only level-2 IS-IS routes can be redistributed.</p> <p>match: Used only when OSPF routes are redistributed. Only the routes meeting the filtering conditions are redistributed. By default, all OSPF routes can be redistributed.</p> <p>metric <i>metric-value</i>: Specifies the metric of the OSPF external LSA. <i>metric-value</i> specifies the size of the metric. The value ranges from 0 to 16,777,214.</p> <p>metric-type { 1 2 }: Sets the external route type, which can be E-1 or E-2.</p>

	<p>route-map <i>route-map-name</i>: Sets the redistribution filtering rules.</p> <p>subnets: Specifies the non-standard networks for redistribution.</p> <p>tag <i>tag-value</i>: Specifies the tag value of the route that is redistributed into the OSPF routing domain. The value ranges from 0 to 4,294,967,295.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>After this command is configured, the router becomes an ASBR, imports related routing information to the OSPF domain, and advertises the routing information as Type 5 LSAs to other OSPF routers in the domain.</p> <p>If you configure redistribution of IS-IS routes without specifying the level parameter, only level-2 routes can be redistributed by default. If you specify the level parameter during initial configuration of redistribution, routes of the specified level can be redistributed. If both level-1 and level-2 are configured, the two levels are combined and saved as level-1-2. For details, see the configuration example.</p> <p>If you configure redistribution of OSPF routes without specifying the match parameter, OSPF routes of all sub-types can be distributed by default. The latest setting of the match parameter is used as the initial match parameter. Only routes that match the sub-types can be redistributed. You can use the no form of the command to restore the default value of match. For details, see the configuration example.</p> <p>If route-map is specified, the filtering rules specified in route-map are applicable to original parameters of redistribution. For redistribution of OSPF or IS-IS routes, the routemap is used for filtering only when the redistributed routes meet criteria specified by match or level.</p> <p>The set metric value associated with route-map should fall into the range of 0 to 16,777,214. If the value exceeds this range, routes cannot be introduced.</p> <p>The configuration rules for the no form of the redistribute command are as follows:</p> <ol style="list-style-type: none"> 1. If some parameters are specified in the no form of the command, default values of these parameters will be restored. 2. If no parameter is specified in the no form of the command, the entire command will be deleted. <p>For example, if redistribute isis 112 level-2 is configured, you can run the no redistribute isis 112 level-2 command to restore the default value of level-2.</p> <p>As level-2 itself is the default value of the parameter, the configuration saved is still redistribute isis 112 level-2 after the preceding no form of the command is executed. To delete the entire command, run the no redistribute isis 112 command.</p>

Introducing a Default Route

Command	default-information originate [always] [metric <i>metric</i>] [metric-type <i>type</i>] [route-map <i>map-name</i>]
Parameter Description	<p>always: Enables OSPF to generate a default route regardless of whether the local router has a default route.</p> <p>metric <i>metric</i>: Indicates the initial metric of the default route. The value ranges from 0 to 16,777,214.</p> <p>metric-type <i>type</i>: Indicates the type of the default route. OSPF external routes are classified into two types: Type 1: The metric varies with routers; Type 2: The metric is the same for all routers. Type 1 external routes are more trustworthy than Type 2 external routes.</p> <p>route-map <i>map-name</i>: Indicates the associated route-map name. By default, no route-map is associated.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>When the redistribute or default-information command is executed, the OSPF router automatically becomes an ASBR. The ASBR, however, does not automatically generate or advertise a default route to all routers in the OSPF routing domain. To have the ASBR generate a default route, configure the default-information originate command.</p> <p>If always is specified, the OSPF routing process advertises an external default route to neighbors regardless of whether a default route exists. This default route, however, is not displayed on the local router. To confirm whether the default route is generated, run the show ip ospf database command to display the OSPF link status database. The external link with the ID 0.0.0.0 describes the default route. On an OSPF neighbor, you can run the show ip route command to see the default route.</p> <p>The metric of the external default route can only be defined in the default-information originate command, instead of the default-metric command.</p> <p>OSPF has two types of external routes. The metric of the Type 1 external route changes, but the metric of the Type 2 external route is fixed. If two parallel paths to the same destination have the same route metric, the priority of the Type 1 route is higher than that of the Type 2 route. Therefore, the show ip route command displays only the Type 1 route.</p> <p>A router in the stub area cannot generate an external default route.</p> <p>The set metric value associated with route-map should fall into the range of 0 to 16,777,214. If the value exceeds this range, routes cannot be introduced.</p>

Configuration

Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring Static Route Redistribution

<p>Scenario Figure 4-10</p>	
	<p>Remarks The interface IP addresses are as follows: A: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1 B: GE 0/1 192.168.1.2 GE 0/2 192.168.3.1 C: GE 0/2 192.168.2.2 D: GE 0/1 192.168.6.2 GE 0/2 192.168.3.2</p>
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Introduce an external static route to Router D.
<p>D</p>	<pre>D# configure terminal D(config)# ip route 172.10.10.0 255.255.255.0 192.168.6.3 D(config)#router ospf 1 D(config-router)# redistribute staticsubnets</pre>
<p>Verification</p>	<ul style="list-style-type: none"> On Router D, run the show ip ospf database external brief command to verify that an LSA corresponding to an external route is generated. On Router C, run the show ip route ospf command to verify that the external static route has been introduced.
<p>D</p>	<pre>D# show ip ospf database external brief OSPF Router with ID (192.168.22.30) (Process ID 1) AS External Link States</pre>

	Link ID	ADV Router	Age	Seq#	CkSum	Route	Tag
	172.10.10.0	192.168.22.30	11	0x80000001	0xa4bb E2	172.10.10.0/24	0
C	C# show ip route ospf						
	O E2 172.10.10.0/24 [110/20] via 192.168.2.1, 00:18:03, GigabitEthernet 0/2						

Configuring the Default Route

Scenario Figure 4-11	
Remarks	<p>The interface IP addresses are as follows:</p> <p>A: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1</p> <p>B: GE 0/1 192.168.1.2 GE 0/2 192.168.3.1</p> <p>C: GE 0/2 192.168.2.2</p> <p>D: GE 0/2 192.168.3.2</p>
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Configure the default route on Router D.
D	<pre>D# configure terminal D(config)#router ospf 1 D(config-router)#default-information originate always</pre>
Verification	<ul style="list-style-type: none"> On Router D, run the show ip ospf database external brief command to verify that an LSA corresponding to the default route is generated. On Router C, run the show ip route ospf command to verify that the OSPF default route exists.
D	<pre>D#show ip ospf database external brief OSPF Router with ID (192.168.22.30) (Process ID 1)</pre>

AS External Link States							
Link ID	ADV Router	Age	Seq#	CkSum	Route	Tag	
0.0.0.0	192.168.22.30	565	0x80000002	0xa190 E2	0.0.0.0/0	1	

C	<pre>C# show ip route ospf O E20.0.0.0/0 [110/20] via 192.168.2.1, 00:18:03, GigabitEthernet 0/2</pre>
---	--

Common Errors

- The subnet route is not introduced because the **subnets** parameter in the **redistribute** command is not configured.
- A routing loop is formed because the **default-information originate always** command is configured on multiple routers.
- Routes cannot be introduced because route redistribution is configured on a router in the stub area.

4.4.4 Configuring Stub Area and NSSA Area

Configuration Effect

- Configure an area located on the stub as a stub area to reduce interaction of routing information and the size of routing table, and enhance stability of routes.

Notes

- The OSPF basic functions must be configured.
- A backbone or transit area cannot be configured as a stub or an NSSA area.
- A router in the stub area cannot introduce external routes, but a router in the NSSA area can introduce external routes.

Configuration Steps

Configuring a Stub Area

- (Optional) This configuration is required if you wish to reduce the size of the routing table on routers in the area.
- The area must be configured as a stub area on all routers in this area.

Configuring an NSSA Area

- (Optional) This configuration is required if you wish to reduce the size of the routing table on routers in the area and introduce OSPF external routes to the area.

- The area must be configured as an NSSA area on all routers in this area.

Verification

Verifying the Stub Area

- On a router in the stub area, run the **show ip route** command to verify that the router is not loaded with any external routes.

Verifying the NSSA Area

- On a router in the NSSA area, run the **show ip ospf database** command to verify that the introduced external route generates Type 7 LSAs.
- On a router in the backbone area, run the **show ip route** command to verify that the router is loaded with external routes introduced from the NSSA area.

Related

Commands

Configuring a Stub Area

Command	area <i>area-id</i> stub [no-summary]
Parameter Description	<i>area-id</i> : Indicates the ID of the stub area. no-summary : Prohibits the ABR from sending network summary LSAs. At this time, the stub can be called totally stub area. This parameter is configured only when the router is an ABR.
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>You must run the area stub command on all routers in the OSPF stub area. The ABR sends only three types of LSAs to the stub area: (1) Type 1: Router LSA; (2) Type 2: Network LSA; (3) Type 3: Network Summary LSA. From the routing table point of view, a router in the stub area can learn only the internal routes of the OSPF routing domain, including the internal default route generated by an ABR. A router in the stub area cannot learn external routes of the OSPF routing domain.</p> <p>To configure a totally stub area, add the no-summary keyword when running the area stub command on the ABR. A router in the totally stub area can learn only the internal routes of the local area, including the internal default route generated by an ABR.</p> <p>You can run either the area stub or area default-cost command to configure an OSPF area as a stub area. If area stub is used, you must configure this command on all routers connected to the stub area. If area default-cost is used, run this command only on the ABR in the stub area. The area default-cost command defines the initial cost (metric) of the internal default route.</p>

Configuring an NSSA Area

Command	<code>area area-id nssa [no-redistribution] [default-information-originate [metric-value] [metric-type type]] [no-summary] [translator [stability-interval seconds always]]</code>
Parameter Description	<p><i>area-id</i>: Indicates the ID of the NSSA area.</p> <p>no-redistribution: Select this option if the router is an NSSA ABR and you want to use only the redistribute command to introduce the routing information into a common area instead of an NSSA area.</p> <p>default-information-originate: Indicates that a default Type 7 LSA is generated and introduced to the NSSA area. This option takes effect only on an NSSA ABR or ASBR.</p> <p>metric-value: Specifies the metric of the generated default LSA. The value ranges from 0 to 16,777,214. The default value is 1.</p> <p>metric-type: Specifies the route type of the generated default LSA. The values include 1 and 2. 1 represents N-1, and 2 represents N-2. The default value is 2.</p> <p>no-summary: Prohibits the ABR in the NSSA area from sending summary LSAs (Type-3 LSA).</p> <p>translator: Indicates that the NSSA ABR is a translator.</p> <p>stability-interval seconds: Indicates the stability interval after the NSSA ABR is changed from a translator to a non-translator. The unit is second. The default value is 40. The value ranges from 0 to 2,147,483,647.</p> <p>always: Indicates that the current NSSA ABR always acts as a translator. The default value is the standby translator.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The default-information-originate parameter is used to generate a default Type 7 LSA. This parameter has different functions on the ABR and the ASBR in the NSSA area. On the ABR, a Type 7 LSA default route is generated regardless of whether the default route exists in the routing table. On the ASBR (not an ABR), a Type 7 LSA default route is generated only when the default route exists in the routing table.</p> <p>If the no-redistribution parameter is configured on the ASBR, other external routes introduced by OSPF through the redistribute command cannot be advertised to the NSSA area. This parameter is generally used when a router in the NSSA area acts both as the ASBR and the ABR. It prevents external routing information from entering the NSSA area.</p> <p>To further reduce the number of LSAs sent to the NSSA area, you can configure the no-summary parameter on the ABR to prevent the ABR from sending the summary LSAs (Type 3 LSA) to the NSSA area.</p>

area default-cost is used on an ABR or ASBR connected to the NSSA area. This command configures the cost of the default route sent from the ABR/ASBR to the NSSA area. By default, the cost of the default route sent to the NSSA area is 1.

If an NSSA area has two or more ABRs, the ABR with the largest router ID is elected by default as the translator for converting Type 7 LSAs into Type 5 LSAs. If the current device is always the translator ABR for converting Type 7 LSAs into Type 5 LSAs, use the **translator always** parameter.

If the translator role of the current device is replaced by another ABR, the conversion capability is retained during the time specified by **stability-interval**. If the router does not become a translator again during **stability-interval**, LSAs that are converted from Type 7 to Type 5 will be deleted from the AS after **stability-interval** expires.

To prevent a routing loop, LSAs that are converted from Type 7 to Type 5 will be deleted from the AS immediately after the current device loses the translator role even if **stability-interval** does not expire.

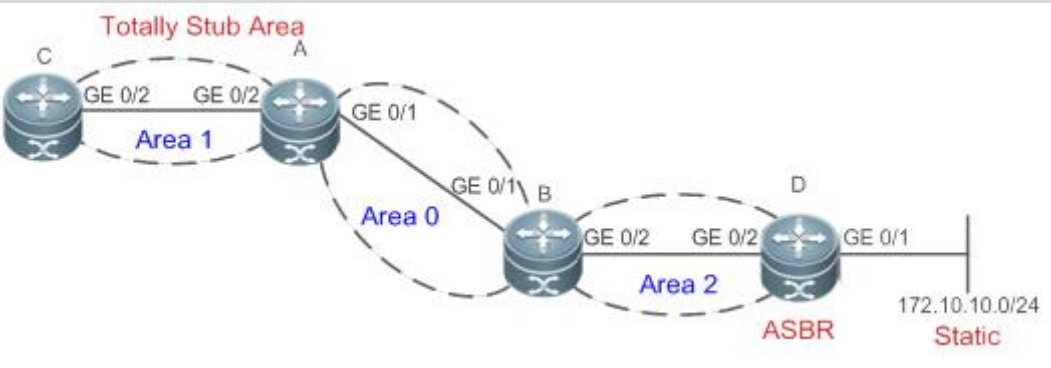
In the same NSSA area, it is recommended that **translator always** be configured on only one ABR.

Configuration

Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring a Stub Area

<p>Scenario Figure 4-12</p>	 <p>Remarks The interface IP addresses are as follows: A: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1 B: GE 0/1 192.168.1.2 GE 0/2 192.168.3.1 C: GE 0/2 192.168.2.2 D: GE 0/1 192.168.6.2 GE 0/2 192.168.3.2</p>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Introduce an external static route to Router D.

	<ul style="list-style-type: none"> Configure area 1 as the stub area on Router A and Router C.
D	<pre>D# configure terminal D(config)#router ospf 1 D(config-router)# redistribute staticsubnets</pre>
A	<pre>A# configure terminal A(config)#router ospf 1 A(config-router)#area 1 stubno-summary</pre>
C	<pre>C# configure terminal C(config)#router ospf 1 C(config-router)#area 1 stub</pre>
Verification	<p>On Router C, run the show ip route ospf command to display the routing table. Verify that there is only one default inter-area route, and no external static route is introduced from Router D.</p>
	<pre>C#show ip route ospf O*IA 0.0.0.0/0 [110/2] via 192.168.2.1, 00:30:53, GigabitEthernet 0/2</pre>

Configuring an NSSA Area

Scenario
Figure 4-13

Remarks **The interface IP addresses are as follows:**

A: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1

B: GE 0/1 192.168.1.2 GE 0/2 192.168.3.1

C: GE 0/2 192.168.2.2

D: GE 0/1 192.168.6.2 GE 0/2 192.168.3.2

<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Introduce an external static route to Router D. ▪ Configure area 2 as the NSSA area on Router B and Router D.
<p>B</p>	<pre>B# configure terminal B(config)#router ospf 1 B(config-router)#area 2 nssa</pre>
<p>D</p>	<pre>D# configure terminal D(config)#ip route 172.10.10.0 255.255.255.0 192.168.6.2 D(config)#router ospf 1 D(config-router)#redistribute static subnets D(config-router)#area 2 nssa</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ On Router D, verify that the Type 7 LSA, 172.10.10.0/24, is generated. ▪ On Router B, verify that Type 5 and Type 7 LSAs coexist on 172.10.10.0/24. ▪ On Router B, verify that the N-2 route of 172.10.10.0/24 is generated.
<p>D</p>	<pre>D# show ip ospf database nssa-external OSPF Router with ID (192.168.6.2) (Process ID 1) NSSA-external Link States (Area 0.0.0.1 [NSSA]) LS age: 61 Options: 0x8 (- - - N/P - -) LS Type: AS-NSSA-LSA Link State ID: 172.10.10.0 (External Network Number For NSSA) Advertising Router: 192.168.6.2 LS Seq Number: 80000001 Checksum: 0xc8f8 Length: 36 Network Mask: /24 Metric Type: 2 (Larger than any link state path) TOS: 0 Metric: 20 NSSA: Forward Address: 192.168.6.2</pre>

External Route Tag: 0

B

```
B# show ip ospf database nssa-external
      OSPF Router with ID (192.168.3.1) (Process ID 1)
        NSSA-external Link States (Area 0.0.0.1 [NSSA])
LS age: 314
Options: 0x8 (-|-|-|-|N/P|-|-|-)
LS Type: AS-NSSA-LSA
Link State ID: 172.10.10.0 (External Network Number For NSSA)
Advertising Router: 192.168.6.2
LS Seq Number: 80000001
Checksum: 0xc8f8
Length: 36
Network Mask: /24
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 20
      NSSA: Forward Address: 192.168.6.2
      External Route Tag: 0

B# show ip ospf database external
      OSPF Router with ID (192.168.3.1) (Process ID 1)
        AS External Link States
LS age: 875
Options: 0x2 (-|-|-|-|-|E|-)
LS Type: AS-external-LSA
Link State ID: 172.10.10.0 (External Network Number)
Advertising Router: 192.168.3.1
LS Seq Number: 80000001
Checksum: 0xd0d3
Length: 36
Network Mask: /24
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
```

```
Metric: 20
Forward Address: 192.168.6.2
External Route Tag: 0
```

```
B# show ip route ospf
O N2 172.10.10.0/24 [110/20] via 192.168.3.2, 00:06:53, GigabitEthernet 0/2
```

Common Errors

- Configurations of the area type are inconsistent on routers in the same area.
- External routes cannot be introduced because route redistribution is configured on a router in the stub area.

4.4.5 Configuring Route Summarization

Configuration Effect

- Summarize routes to reduce interaction of routing information and the size of routing table, and enhance stability of routes.
- Shield or filter routes.

Notes

- The OSPF basic functions must be configured.
- The address range of summarized routes may exceed the actual network range in the routing table. If data is sent to a network beyond the summarization range, a routing loop may be formed and the router processing load may increase. To prevent these problems, a discard route must be added to the routing table or shield or filter routes.

Configuration Steps

Configuring Inter-Area Route Summarization

- (Optional) This configuration is required when routes of the OSPF area need to be summarized.
- Unless otherwise required, this configuration should be performed on an ABR in the area where routes to be summarized are located.

Configuring External Route Summarization

- (Optional) This configuration is required when routes external to the OSPF domain need to be summarized.

- Unless otherwise required, this configuration should be performed on an ASBR to which routes to be summarized are introduced.

Verification

Run the **show ip route ospf** command to verify that individual routes do not exist and only the summarized route exists.

Related Commands

Configuring Inter-Area Route Summarization

Command	<code>area <i>area-idrange</i> <i>ip-address net-mask</i> [advertise not-advertise] [cost <i>cost</i>]</code>
Parameter Description	<p><i>area-id</i>: Specifies the ID of the OSPF area to which the summarized route should be injected. The area ID can be a decimal integer or an IP address.</p> <p><i>ip-address net-mask</i>: Defines the network segment of the summarized route.</p> <p>advertise not-advertise: Specifies whether the summarized route should be advertised.</p> <p>cost <i>cost</i>: Indicates the metric of the summarized route. The value ranges from 0 to 16777215.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>This command can be executed only on the ABR. It is used to combine or summarize multiple routes of an area into one route, and advertise the route to other areas. Combination of the routing information occurs only on the boundary of an area. Routers inside the area can learn specific routing information, whereas routers in other areas can learn only one summarized route. In addition, you can set advertise or not-advertise to determine whether to advertise the summarized route to shield and filter routes. By default, the summarized route is advertised. You can use the cost parameter to set the metric of the summarized route.</p> <p>You can configure route summarization commands for multiple areas. This simplifies routes in the entire OSPF routing domain, and improve the network forwarding performance, especially for a large-sized network.</p> <p>When multiple route summarization commands are configured and have the inclusive relationship with each other, the area range to be summarized is determined based on the maximum match principle.</p>

Configuring External Route Summarization

Command	summary-address <i>ip-address net-mask</i> [not-advertise tag value]
Parameter Description	<p><i>ip-address</i>: Indicates the IP address of the summarized route.</p> <p><i>net-mask</i>: Indicates the subnet mask of the summarized route.</p> <p>not-advertise: Indicates that the summarized route is not advertised. If this parameter is not specified, the summarized route is advertised.</p> <p>tagvalue: Indicates the tag of the summarized route. The value ranges from 0 to 4,294,967,295.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>When routes are redistributed from other routing processes and injected to the OSPF routing process, each route is advertised to the OSPF routers using an external LSA. If the injected routes are a continuous address space, the ABR can advertised only one summarized route to significantly reduce the size of the routing table.</p> <p>area range summarizes the routes between OSPF routes, whereas summary-address summarizes external routes of the OSPF routing domain.</p> <p>When configured on the NSSA ABR translator, summary-address summarizes redistributed routes and routes obtained based on the LSAs that are converted from Type 7 to Type 5. When configured on the ASBR (not an NSSA ABR translator), summary-address summarizes only redistributed routes.</p>

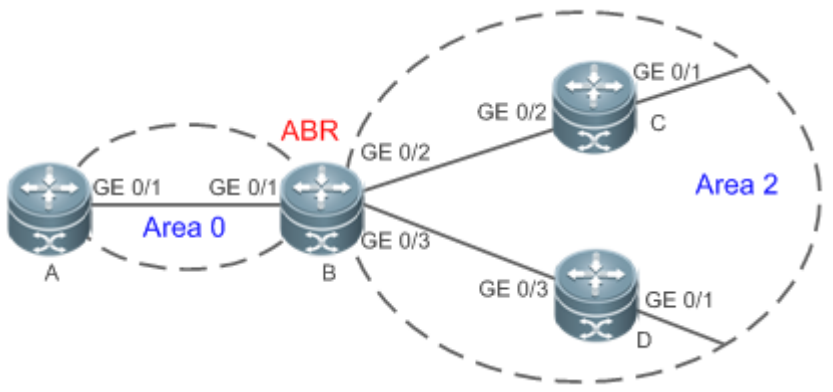
Configuring a Discard Route

Command	discard-route { internal external }
Parameter Description	<p>internal: Indicates that the discard route generated by the area range command can be added.</p> <p>external: Indicates that the discard route generated by the summary-address command can be added.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	The address range of summarized routes may exceed the actual network range in the routing table. If data is sent to a network beyond the summarization range, a routing loop may be formed and the router processing load may increase. To prevent these

problems, a discard route must be added to the routing table on the ABR or ASBR. This route is automatically generated, and is not advertised.

Configuration Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

<p>Scenario Figure 4-14</p>	 <p>Remark S</p> <p>The interface IP addresses are as follows:</p> <p>A: GE0/1 192.168.1.1</p> <p>B: GE0/1 192.168.1.2 GE0/2 172.16.2.1 GE0/3 172.16.3.1</p> <p>C: GE0/2 172.16.2.2 GE0/1 172.16.4.2</p> <p>D: GE0/2 172.16.3.2 GE0/1 172.16.5.2</p>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Summarize routes of area 2 on Router B.
<p>B</p>	<pre>B# configure terminal B(config)#router ospf 1 B(config-router)#area 2 range 172.16.0.0 255.255.0.0</pre>
<p>Verification</p>	<p>On Router A, verify that the entry 172.16.0.0/16 is added to the routing table.</p>
<p>A</p>	<pre>A#show ip route ospf O IA 172.16.0.0/16 [110/2] via 192.168.1.2, 00:01:04, GigabitEthernet 0/1</pre>

Common Errors

- Inter-area route summarization cannot be implemented because the **area range** command is configured on a non-ABR device.

4.4.6 Configuring Route Filtering

Configuration Effect

- Routes that do not meet filtering conditions cannot be loaded to the routing table, or advertised to neighbors. Network users cannot access specified destination network.

Notes

- The OSPF basic functions must be configured.
- Filtering routes by using the **distribute-list in** command affects forwarding of local routes, but does not affect route computation based on LSAs. Therefore, if route filtering is configured on the ABR, Type 3 LSAs will still be generated and advertised to other areas because routes can still be computed based on LSAs. As a result, black-hole routes are generated. In this case, you can run the **area filter-list** or **area range** (containing the **not-advertise** parameter) command on the ABR to prevent generation of black-hole routes.

Configuration Steps

Configuring Inter-Area Route Filtering

- (Optional) This configuration is recommended if users should be restricted from accessing the network in a certain OSPF area.
- Unless otherwise required, this configuration should be performed on an ABR in the area where filtered routes are located.

Configuring Redistributed Route Filtering

- (Optional) This configuration is required if external routes introduced by the ASBR need to be filtered.
- Unless otherwise required, this configuration should be performed on an ASBR to which filtered routes are introduced.

Configuring Learned Route Filtering

- (Optional) This configuration is required if users should be restricted from accessing a specified destination network.
- Unless otherwise required, this configuration should be performed on a router that requires route filtering.

Verification

- Run the **show ip route** command to verify that the router is not loaded with routes that have been filtered out.
- Run the **ping** command to verify that the specified destination network cannot be accessed.

Related Commands

Configuring a Passive Interface

Command	<code>passive-interface{ default <i>interface-type interface-number</i> <i>interface-type interface-number ip-address</i>}</code>
Parameter Description	<i>interface-type interface-number</i> : Indicates the interface that should be configured as a passive interface. default : Indicates that all interface will be configured as passive interfaces. <i>interface-type interface-number ip-address</i> : Specifies an address of the interface as the passive address.
Command Mode	OSPF routing process configuration mode
Usage Guide	To prevent other routers on the network from learning the routing information of the local router, you can configure a specified network interface of the local router as the passive interface, or a specified IP address of a network interface as the passive address.

Configuring the LSA Update Packet Filtering

Command	<code>ip ospf database-filter all out</code>
Parameter Description	N/A
Command Mode	Interface configuration mode
Usage Guide	Enable this function on an interface to prevent sending the LSA update packet on this interface. After this function is enabled, the local router does not advertise the LSA update packet to neighbors, but still sets up the adjacency with neighbors and receives LSAs from neighbors.

Configuring Inter-Area Route Filtering

Command	<code>area <i>area-id</i> filter-list { <i>access acl-name</i> prefix <i>prefix-name</i> } {in out}</code>
Parameter Description	<i>area-id</i> : Indicates the area ID. access <i>acl-name</i> : Indicates the associated ACL. prefix <i>prefix-name</i> : Indicates the associated prefix list. in out : Filters routes that are received by or sent from the area.
Command Mode	OSPF routing process configuration mode
Usage Guide	This command can be configured only on an ABR. Use this command when it is required to configure filtering conditions for inter-area routes on the ABR.

Configuring Redistributed Route Filtering

Command	<code>distribute-list { [<i>access-list-number</i> <i>name</i>] prefix <i>prefix-list-name</i> } out [<i>bgp</i> <i>connected</i> <i>isis</i>[<i>area-tag</i>] <i>ospf process-id</i> <i>rip</i> <i>arp-host</i> <i>static</i>]</code>
Parameter Description	<i>access-list-number</i> / <i>name</i> : Uses the ACL for filtering. prefix <i>prefix-list-name</i> : Uses the prefixlist for filtering. <i>bgp</i> <i>connected</i> <i>isis</i> [<i>area-tag</i>] <i>ospf process-id</i> <i>rip</i> <i>arp-host</i> <i>static</i> : Indicates the source of routes to be filtered.
Command Mode	OSPF routing process configuration mode
Usage Guide	distribute-list out is similar to redistribute route-map , and is used to filter routes that are redistributed from other protocols to OSPF. The distribute-list out command itself does not redistribute routes, and is generally used together with the redistribute command. The ACL and the prefixlist filtering rules are mutually exclusive in the configuration. That is, if the ACL is used for filtering routes coming from a certain source, the prefixlist cannot be configured to filter the same routes.


Configuring Learned Route Filtering

Command	distribute-list {[<i>access-list-number</i> <i>name</i>] prefix <i>prefix-list-name</i> [<i>gateway prefix-list-name</i>] route-map <i>route-map-name</i> } in [<i>interface-type</i> <i>interface-number</i>]
Parameter Description	<i>access-list-number</i> <i>name</i> : Uses the ACL for filtering. gateway <i>prefix-list-name</i> : Uses the gateway for filtering. prefix <i>prefix-list-name</i> : Uses the prefixlist for filtering. route-map <i>route-map-name</i> : Uses the route map for filtering. <i>interface-type interface-number</i> : Specifies the interface for which LSA routes are filtered.
Command Mode	OSPF routing process configuration mode
Usage Guide	Filter routes that are computed based on received LSAs. Only routes meeting the filtering conditions can be forwarded. The command does not affect the LSDB or the routing tables of neighbors. The ACL, prefix list, and route map filtering rules are mutually exclusive in the configuration. That is, if the ACL is used for filtering routes of a specified interface, the prefix list or router map cannot be configured for filtering routes of the same interface.

Configuration Example

i The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Scenario
Figure 4-15



Remarks The interface IP addresses are as follows:
A: GE0/1 192.168.1.1
B: GE0/1 192.168.1.2 GE0/2 172.16.2.1 GE0/3 172.16.3.1

	<pre>C: GE0/2 172.16.2.2 GE0/3 172.16.4.2 D: GE0/2 172.16.3.2 GE0/3 172.16.5.2</pre>
Configurati on Steps	<p>Configure the interface IP addresses on all routers. (Omitted)</p> <p>Configure the OSPF basic functions on all routers. (Omitted)</p> <p>On Router A, configure route filtering.</p>
A	<pre>A# configure terminal A(config)#access-list 3 permit host 172.16.5.0 A(config)#router ospf 1 A(config-router)#distribute-list 3 in GigabitEthernet 0/1</pre>
Verification	<p>On Router A, check the routing table. Verify that only the entry 172.16.5.0/24 is loaded.</p>
A	<pre>A# show ip route ospf O 172.16.5.0/24 [110/2] via 192.168.1.2, 10:39:40, GigabitEthernet 0/1</pre>

Common Errors

- Filtering routes by using the **distribute-list in** command affects forwarding of local routes, but does not affect route computation based on LSAs. Therefore, if route filtering is configured on the ABR, Type 3 LSAs will still be generated and advertised to other areas because routes can still be computed based on LSAs. As a result, black-hole routes are generated.

4.4.7 Modifying Route Cost and AD

Configuration Effect

- Change the OSPF routes to enable the traffic pass through specified nodes or avoid passing through specified nodes.
- Change the sequence that a router selects routes so as to change the priorities of OSPF routes.

Notes

- The OSPF basic functions must be configured.
- If you run the **ip ospf cost** command to configure the cost of an interface, the configured cost will automatically overwrite the cost that is computed based on the auto cost.

Configuration Steps

Configuring the Reference Bandwidth

- Optional.
- A router is connected with lines with different bandwidths. This configuration is recommended if you wish to preferentially select the line with a larger bandwidth.

Configuring the Cost of an Interface

- Optional.
- A router is connected with multiple lines. This configuration is recommended if you wish to manually specify a preferential line.

Configuring the Default Metric for Redistribution

- Optional.
- This configuration is mandatory if the cost of external routes of the OSPF domain should be specified when external routes are introduced to an ASBR.

Configuring the Maximum Metric

- Optional.
- A router may be unstable during the restart process or a period of time after the router is restarted, and users do not want to forward data through this router. In this case, this configuration is recommended.

Configuring the AD

- Optional.
- This configuration is mandatory if you wish to change the priorities of OSPF routes on a router that runs multiple unicast routing protocols.

Verification

- Run the **show ip ospf interface** command to verify that the costs of interfaces are correct.
- Run the **show ip route** command to verify that the costs of external routes introduced to the ASBR are correct.
- Restart the router. Within a specified period of time, data is not forwarded through the restarted router.

Related Commands

Configuring the Reference Bandwidth

Command
<code>auto-costreference-bandwidth ref-bw</code>

Parameter Description	<i>ref-bw</i> : Indicates the reference bandwidth. The unit is Mbps. The value ranges from 1 to 4,294,967.
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>By default, the cost of an OSPF interface is equal to the reference value of the auto cost divided by the interface bandwidth.</p> <p>Run the auto-cost command to obtain the reference value of the auto cost. The default value is 100 Mbps.</p> <p>Run the bandwidth command to set the interface bandwidth.</p> <p>The costs of OSPF interfaces on several typical lines are as follows:</p> <p>64Kbps serial line: The cost is 1562.</p> <p>E1 line: The cost is 48.</p> <p>10M Ethernet: The cost is 10.</p> <p>100M Ethernet: The cost is 1.</p> <p>If you run the ip ospf cost command to configure the cost of an interface, the configured cost will automatically overwrite the cost that is computed based on the auto cost.</p>

Configuring the Cost of an Interface

Command	ip ospf costcost
Parameter Description	<i>cost</i> : Indicates the cost of an OSPF interface. The value ranges from 0 to 65,535.
Command Mode	Interface configuration mode
Usage Guide	<p>By default, the cost of an OSPF interface is equal to the reference value of the auto cost divided by the interface bandwidth.</p> <p>Run the auto-cost command to obtain the reference value of the auto cost. The default value is 100 Mbps.</p> <p>Run the bandwidth command to set the interface bandwidth.</p> <p>The costs of OSPF interfaces on several typical lines are as follows:</p> <p>64Kbps serial line: The cost is 1562.</p>

	<p>E1 line: The cost is 48.</p> <p>10M Ethernet: The cost is 10.</p> <p>100M Ethernet: The cost is 1.</p> <p>If you run the ip ospf cost command to configure the cost of an interface, the configured cost will automatically overwrite the cost that is computed based on the auto cost.</p>
--	---

Configuring the Cost of the Default Route in a Stub or an NSSA Area

Command	area <i>area-id</i> default-cost <i>cost</i>
Parameter Description	<p><i>area-id</i>: Indicates the ID of the stub or NSSA area.</p> <p><i>cost</i>: Indicates the cost of the default summarized route injected to the stub or NSSA area. The value ranges from 0 to 16,777,215.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>This command takes effect only on an ABR in a stub area or an ABR/ASBR in an NSSA area.</p> <p>An ABR in a stub area or an ABR/ASBR in an NSSA area is allowed to advertise an LSA indicating the default route in the stub or NSSA area. You can run the area default-cost command to modify the cost of the advertised LSA.</p>

Configuring the Default Metric for Redistribution

Command	default-metric <i>metric</i>
Parameter Description	<p><i>metric</i>: Indicates the default metric of the OSPF redistributed route. The value ranges from 1 to 16,777,214.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The default-metric command must be used together with the redistribute command to modify the initial metrics of all redistributed routes.</p>

The **default-metric** command does not take effect on external routes that are injected to the OSPF routing domain by the **default-information originate** command.

Configuring the Maximum Metric

Command	max-metric router-lsa [external-lsa [<i>max-metric-value</i>]] [include-stub] [on-startup[<i>seconds</i>]] [summary-lsa [<i>max-metric-value</i>]]
Parameter Description	<p>router-lsa: Sets the metrics of non-stub links in the Router LSA to the maximum value (0xFFFF).</p> <p>external-lsa: Allows a router to replace the metrics of external LSAs (including Type 5 and Type 7 LSAs) with the maximum metric.</p> <p><i>max-metric-value:</i> Indicates the maximum metric of the LSA. The default value is 16711680. The value ranges from 1 to 16,777,215.</p> <p>include-stub: Sets the metrics of stub links in the Router LSA advertised by the router to the maximum value.</p> <p>on-startup: Allows a router to advertise the maximum metric when started.</p> <p><i>seconds:</i> Indicates the interval at which the maximum metric is advertised. The default value is 600s. The value ranges from 5 to 86,400.</p> <p>summary-lsa: Allows a router to replace the metrics of summary LSAs (including Type 3 and Type 4 LSAs) with the maximum metric.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>After the max-metric router-lsa command is executed, the metrics of the non-stub links in the Router LSAs generated by the router will be set to the maximum value (0xFFFF). If you cancel this configuration or the timer expires, the normal metrics of the links are restored.</p> <p>By default, if the max-metric router-lsa command is executed, the stub links still advertise common metrics, that is, the costs of outbound interfaces. If the include-stub parameter is configured, the stub links will advertise the maximum metric.</p> <p>If an ABR does not wish to transfer inter-area traffic, use the summary-lsa parameter to set the metric of the Summary LSA to the maximum metric.</p> <p>If an ASBR does not wish to transfer external traffic, use the external-lsa parameter to set the metric of the external LSA to the maximum metric.</p> <p>The max-metric router-lsa command is generally used in the following scenarios:</p> <p>Restart a device. After the device is restarted, IGP generally converges faster, and other devices attempt to forward traffic through the restarted device. If the current device is</p>

still building the BGP routing table and some BGP routes are not learned yet, packets sent these networks will be discarded. In this case, you can use the **on-startup** parameter to set a delay after which the restarted device acts as the transmission mode.

- Add a device to the network but the device is not used to transfer traffic. The device is added to the network. If a candidate path exists, the current device is not used to transfer traffic. If a candidate path does not exist, the current device is still used to transfer traffic.
- Delete a device gracefully from the network. After the **max-metric router-lsa** command is executed, the current device advertises the maximum metric among all metrics of routes. In this way, other devices on the network can select the standby path for data transmission before the device is shut down.

In the earlier OSPF version (RFC1247 or earlier), the links with the maximum metric (0xFFFF) in the LSAs do not participate in the SPF computation, that is, no traffic is sent to routers that generate these LSAs.

Configuring RFC1583Compatibility

Command	compatible rfc1583
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	When there are multiple paths to an ASBR or the forwarding address of an external route, RFC1583 and RFC2328 define different routing rules. If RFC1583 compatibility is configured, a path in the backbone area or an inter-area path is preferentially selected. If RFC1583 compatibility is not configured, a path in a non-backbone area is preferentially selected.

Configuring the AD

Command	distance { <i>distance</i> [route-map <i>map-name</i>] ospf { intra-area <i>distance</i> [route-map <i>map-name</i>] inter-area <i>distance</i> [route-map <i>map-name</i>] external <i>distance</i> [route-map <i>map-name</i>] } }
Parameter Description	<i>distance</i> : Indicates the AD of a route. The value ranges from 1 to 255. intra-area <i>distance</i> : Indicates the AD of an intra-area route. The value ranges from 1 to 255.

	<p>inter-area distance: Indicates the AD of an inter-area route. The value ranges from 1 to 255.</p> <p>external distance: Indicates the AD of an external route. The value ranges from 1 to 255.</p> <p>route-map map-name: Indicates the AD of a specified route by policy.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>Use this command to specify different ADs for different types of OSPF routes.</p> <p>Configure the route-map parameter to set the AD for a specified route by policy. If route-map has configured set distance,</p> <p>through matched route: AD is set by set distance</p> <p>through unmatched route: AD is set by the distance command</p>

Configuration Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring the Cost of an Interface

<p>Scenario</p> <p>Figure 4-16</p>	
<p>Remarks</p>	<p>The interface IP addresses are as follows:</p> <p>A: GE0/1 192.168.1.1 GE0/2 192.168.2.1</p> <p>B: GE0/1 192.168.1.2 GE0/2 192.168.3.2</p> <p>C: GE0/1 192.168.4.2 GE0/2 192.168.2.2</p>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) On Router A, configure the cost of each interface.

A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip ospf cost 10 A(config)# interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)# ip ospf cost 20</pre>
Verification	On Router A, check the routing table. The next hop of the optimum path to 172.16.1.0/24 is Router B.
A	<pre>A# show ip route ospf O E2172.16.1.0/0 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1</pre>

Common Errors

- If the cost of an interface is set to 0 in the **ip ospf cost** command, a route computation error may occur. For example, a routing loop is obtained.

4.4.8 Enabling Authentication

Configuration Effect

- All routers connected to the OSPF network must be authenticated to ensure stability of OSPF and protect OSPF against intrusions.

Notes

- The OSPF basic functions must be configured.
- If authentication is configured for an area, the configuration takes effect on all interfaces that belong to this area.
- If authentication is configured for both an interface and the area to which the interface belongs, the configuration for the interface takes effect preferentially.

Configuration Steps

Configuring the Authentication Type of an Area

- (Optional) This configuration is recommended if the same authentication type should be used on all interfaces in the same area.
- This configuration is required if a router accesses a network that requires authentication.

Configuring the Authentication Type of an Interface

- (Optional) This configuration is recommended if the different authentication types should be used on different interfaces in the same area.
- This configuration is required if a router accesses a network that requires authentication.

Configuring a Plain Text Authentication Key for an Interface

- Optional.
- This configuration is required if a router accesses a network that requires plain text authentication.

Configuring an MD5 Authentication Key for an Interface

- (Optional) MD5 authentication features a high security, and therefore is recommended. You must configure either plain text authentication or MD5 authentication.
- This configuration is required if a router accesses a network that requires MD5 authentication.

Verification

- If routers are configured with different authentication keys, run the **show ip ospf neighbor** command to verify that there is no OSPF neighbor.
- If routers are configured with the same authentication key, run the **show ip ospf neighbor** command to verify that there are OSPF neighbors.

Related Commands

Configuring the Authentication Type of an Area

Command	<code>area <i>area-id</i> authentication [message-digest]</code>
Parameter Description	<i>area-id</i> : Indicates the ID of the area where OSPF authentication is enabled. The area ID can be a decimal integer or an IP address. message-digest : Enables MD5 authentication.
Command Mode	OSPF routing process configuration mode
Usage Guide	The RGOS supports three authentication types: (1) Type 0: No authentication is required. If this command is not configured to enable OSPF authentication, the authentication type in the OSPF data packet is 0. (2) Type 1: The authentication type is plain text authentication if this command is configured but does not contain the message-digest parameter.

(3) Type 3: The authentication type is MD5 authentication if this command is configured and contains the **message-digest** parameter.

All routers in the same OSPF area must use the same authentication type. If authentication is enabled, the authentication key must be configured on interfaces that are connected to neighbors. You can run the interface configuration command **ip ospf authentication-key** to configure the plain text authentication key, or **ip ospf message-digest-key** to configure the MD5 authentication key.

Configuring the Authentication Type of an Interface

Command	ip ospfauthentication [message-digest null]
Parameter Description	message-digest: Indicates that MD5 authentication is enabled on the current interface. null: Indicates that authentication is disabled.
Command Mode	Interface configuration mode
Usage Guide	If the ip ospfauthentication command does not contain any option, it indicates that plain text authentication is enabled. If you use the no form of the command to restore the default authentication mode, whether authentication is enabled is determined by the authentication type that is configured in the area to which the interface belongs. If the authentication type is set to null, authentication is disabled forcibly. When authentication is configured for both an interface and the area to which the interface belongs, the authentication type configured for the interface is used preferentially.

Configuring a Plain Text Authentication Key for an Interface

Command	ip ospf authentication-key[0 7]key
Parameter Description	0: Indicates that the key is displayed in plain text. 7: Indicates that the key is displayed in cipher text. key: Indicates the key. The key is a string of up to eight characters.
Command Mode	Interface configuration mode

Usage Guide	<p>The key configured by the ip ospf authentication-key command will be inserted to the headers of all OSPF packets. If the keys are inconsistent, two directly connected devices cannot set up the OSPF adjacency and therefore cannot exchange the routing information.</p> <p>Different keys can be configured for different interface, but all routers connected to the same physical network segment must be configured with the same key.</p> <p>You can enable or disable authentication in an OSPF area by running the areaauthentication command in OSPF routing process configuration mode.</p> <p>You can also enable authentication on an individual interface by running the ip ospf authentication command in interface configuration mode. When authentication is configured for both an interface and the area to which the interface belongs, the authentication type configured for the interface is used preferentially.</p>
-------------	--

Configuring an MD5 Authentication Key for an Interface

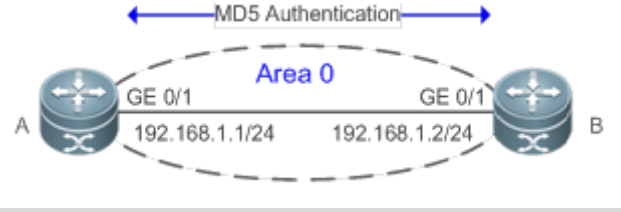
Command	ip ospf message-digest-key <i>key-id</i> md5[0 7]<i>key</i>
Parameter Description	<p><i>key-id</i>: Indicates the key ID. The value ranges from 1 to 255.</p> <p>0: Indicates that the key is displayed in plain text.</p> <p>7: Indicates that the key is displayed in cipher text.</p> <p><i>key</i>: Indicates the key. The key is a string of up to 16 characters.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>The key configured by the ip ospf message-digest-key command will be inserted to the headers of all OSPF packets. If the keys are inconsistent, two directly connected devices cannot set up the OSPF adjacency and therefore cannot exchange the routing information.</p> <p>Different keys can be configured for different interface, but all routers connected to the same physical network segment must be configured with the same key. The same key ID on neighbor routers must correspond to the same key.</p> <p>You can enable or disable authentication in an OSPF area by running the area authentication command in OSPF routing process configuration mode. You can also enable authentication on an individual interface by running the ip ospf authentication command in interface configuration mode. When authentication is configured for both</p>

an interface and the area to which the interface belongs, the authentication type configured for the interface is used preferentially.

The RGOS software supports smooth modification of the MD5 authentication key. A new MD5 authentication key must be first added before the old key can be deleted. When an OSPF MD5 authentication key is added to a router, the router determines that other routers do not use the new key yet and therefore uses different keys to send multiple OSPF packets until it confirms that the new key has been configured on neighbors. After configuring the new key all routers, you can delete the old key.

Configuration Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

<p>Scenario Figure 4-17</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Configure the authentication type and MD5 authentication key on all routers.
<p>A</p>	<pre>A# configure terminal A(config)#router ospf 1 A(config-router)#area 0 authentication message-digest A(config-router)#exit A(config)#interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#ip ospf message-digest-key 1 md5 hello</pre>
<p>B</p>	<pre>B# configure terminal B(config)#router ospf 1 B(config-router)#area 0 authentication message-digest B(config-router)#exit B(config)#interface GigabitEthernet 0/3 B(config-if-GigabitEthernet 0/3)#ip ospf message-digest-key 1 md5 hello</pre>

Verification	On Router A and Router B, verify that the OSPF neighbor status is correct.
A	<pre>A#show ip ospf neighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.2 1 Full/DR 00:00:32 192.168.1.2 GigabitEthernet 0/1</pre>
B	<pre>A#show ip ospf neighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.1 1 Full/DR 00:00:32 192.168.1.1 GigabitEthernet 0/1</pre>

Common Errors

- The authentication modes configured on routers are inconsistent.
- The authentication keys configured on routers are inconsistent.

4.4.9 Enabling Overflow

Configuration Effect

- New routes are not loaded to routers when the router memory is insufficient.
- New routes are not loaded to routers when the usage of the database space reaches the upper limit.

Notes

- The OSPF basic functions must be configured.
- After a router enters the overflow state, you can run the **clear ip ospf process** command, or stop and then restart the OSPF to exit the overflow state.

Configuration Steps

Configuring the Memory Overflow Function

- Optional.
- This configuration is recommended if a large number of routes exist in the domain and may cause insufficiency of the router memory.

Configuring the Database Overflow Function

- Optional.

- This configuration is recommended if a large number of routes exist in the domain and may cause insufficiency of the router memory.

Configuring the External LSA Database Overflow Function

- Optional.
- This configuration is recommended if the ASBR introduces a large number of external routes and the router memory may be insufficient.

Verification

- After the memory becomes insufficient, add new routers to the network, and run the **show ip route** command to verify that new routes are not loaded.
- After the usage of the database space reaches the upper limit, add new routers to the network, and run the **show ip route** command to verify that new routes are not loaded.

Related Commands

Configuring the Memory Overflow Function

Command	overflow memory-lack
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The OSPF process enters the overflow state to discard newly-learned external routes. This behavior can effectively ensure that the memory usage does not increase.</p> <p>After the overflow function is enabled, the OSPF process enters the overflow state and discards newly-learned external routes, which may cause a routing loop on the entire network. To reduce the occurrence probability of this problem, OSPF generates a default route to the null interface, and this route always exists in the overflow state.</p> <p>You can run the clear ip ospf process command to reset the OSPF process so that the OSPF process can exit the overflow state. You can use the no form of the command to prevent the OSPF process from entering the overflow state when the memory is insufficient. This, however, may lead to over-consumption of the memory resource, after which the OSPF process will stop and delete all the learned routes.</p>

Configuring the Database Overflow Function

Command	overflow databasenumbers [hard soft]
---------	---

Parameter Description	<p><i>number</i>: Indicates the maximum number of LSAs. The value ranges from 1 to 4,294,967,294.</p> <p>hard: Indicates that the OSPF process will be stopped if the number of LSAs exceeds the limit.</p> <p>soft: Indicates that a warning will be generated if the number of LSAs exceeds the limit.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	If the number of LSAs exceeds the limit, use the hard parameter if the OSPF process should be stopped, and use the soft parameter if a warning should be generated without stopping the OSPF process.

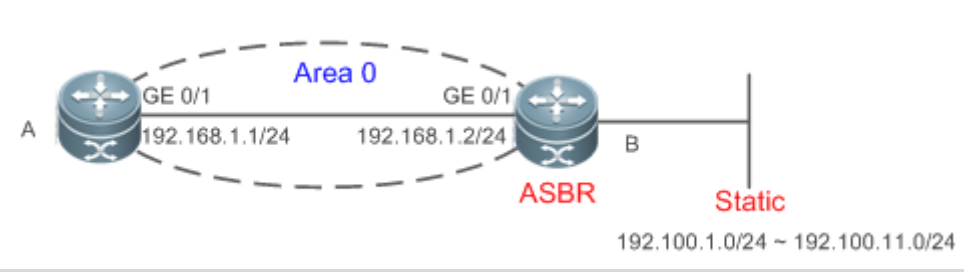
Configuring the External LSA Database Overflow Function

Command	overflow database external <i>max-dbsize wait-time</i>
Parameter Description	<p><i>max-dbsize</i>: Indicates the maximum number of external LSAs. This value must be the same on all routers in the same AS. The value ranges from 0 to 2,147,483,647.</p> <p><i>wait-time</i>: Indicates the waiting time after a router in overflow state attempts to restore the normal state. The value ranges from 0 to 2,147,483,647.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>When the number of external LSAs of a router exceeds the configured max-dbsize, the router enters the overflow state. In this state, the router no longer loads external LSAs and deletes external LSAs that are generated locally. After <i>wait-time</i> elapses, the device restores the normal state, and loads external LSAs again. When using the overflow function, ensure that the same max-dbsize is configured on all routers in the OSPF backbone area and common areas; otherwise, the following problems may occur:</p> <ul style="list-style-type: none"> Inconsistent LSDBs throughout network are inconsistent, and the failure to achieve the full adjacency Incorrect routes, including routing loops Frequent retransmission of AS external LSAs

Configuration Example

i The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring the External LSA Database Overflow Function

<p>Scenario Figure 4-18</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ On Router B, configure redistribution and introduce external static routes. ▪ On Router B, configure the maximum number of external LSAs.
<p>B</p>	<pre>B# configure terminal B(config)# router ospf 1 B(config-router)# redistribute static subnets</pre>
<p>A</p>	<pre>A# configure terminal A(config)# router ospf 1 A(config-router)# overflow database external 10 3</pre>
<p>Verification</p>	<p>On Router B, configure 11 static routes (192.100.1.0/24 to 192.100.11.0/24). On Router A, verify that only 10 static routes are loaded.</p>
<p>A</p>	<pre>A# show ip route ospf O E2 192.100.1.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.2.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.3.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.4.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.5.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.6.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.7.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1 O E2 192.100.8.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1</pre>

```
O E2 192.100.9.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1
O E2 192.100.10.0/24 [110/20] via 192.168.1.2, 00:18:03, GigabitEthernet 0/1
```

Common

Errors

- The OSPF adjacency is abnormal because the maximum number of LSAs is inconsistent on different routers.

4.4.10 Modifying the Maximum Number of Concurrent Neighbors

Configuration

Effect

- Control the maximum number of concurrent neighbors on the OSPF process to ease the pressure on the device.

Notes

- The OSPF basic functions must be configured.

Configuration

Steps

Configuring the Maximum Number of Concurrent Neighbors on the OSPF Process

- (Optional) This configuration is recommended if you wish to set up the OSPF adjacency more quickly when a router is connected with a lot of other routers.
- This configuration is performed on a core router.

Verification

- Run the **show ip ospf neighbor** command to display the number of neighbors that are concurrently interacting with the OSPF process.

Related

Commands

Configuring the Maximum Number of Concurrent Neighbors on the Current Process

Command	<code>max-concurrent-ddnumber</code>
Parameter Description	<i>number</i> : Specifies the maximum number of neighbors that are concurrently interacting with the OSPF process. The value ranges from 1 to 65,535.

Command Mode	OSPF routing process configuration mode
Usage Guide	When the performance of a router is affected because the router exchanges data with multiple neighbors, you can configure this command to restrict the maximum of neighbors with which one OSPF process can concurrently initiate or accept interaction.

Configuring the Maximum Number of Concurrent Neighbors on All Processes

Command	<i>router ospf max-concurrent-ddnumber</i>
Parameter Description	<i>number</i> : Specifies the maximum number of neighbors that are concurrently interacting with the OSPF process. The value ranges from 1 to 65,535.
Command Mode	Global configuration mode
Usage Guide	When the performance of a router is affected because the router exchanges data with multiple neighbors, you can configure this command to restrict the maximum of neighbors with which all OSPF processes can concurrently initiate or accept interaction.

Configuration Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring the Maximum Number of Concurrent Neighbors on the OSPF Process

<p>Scenario Figure 4-19</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ On the router Core, set the maximum number of concurrent neighbors to 4.
<p>Core</p>	<pre>Core# configure terminal Core(config)# router ospf max-concurrent-dd 4</pre>
<p>Verification</p>	<p>On the router Core, check the neighbor status and verify that at most eight neighbors concurrently interact with the OSPF process.</p>

4.4.11 Disabling Source Address Verification

Configuration

Effect

- The unicast routing service can be provided even if the interface IP addresses of neighbor routers are not in the same network segment.

Notes

- The OSPF basic functions must be configured.
- Source address verification cannot be disabled on a broadcast or NBMA network.

Configuration

Steps

Disabling Source Address Verification

- (Optional) This configuration is mandatory if an adjacency should be set up between routers with interface IP addresses in different network segments.

- This configuration is performed on routers with interface IP addresses in different network segments.

Verification

- An adjacency can be set up between routers in different network segments.

Related Commands

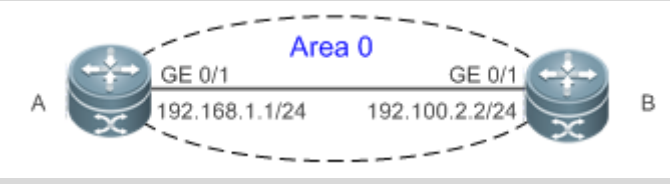
Disabling Source Address Verification

Command	<code>ip ospf source-check-ignore</code>
Parameter Description	N/A
Command Mode	Interface configuration mode
Usage Guide	Generally, the source address of a packet received by OSPF is in the same network segment as the receiving interface. The addresses at both ends of a P2P link are configured separately and are not necessarily in the same network segment. In this scenario, as the peer address information will be notified during the P2P link negotiation process, OSPF checks whether the source address of the packet is the address advertised by the peer during negotiation. If not, OSPF determines that the packet is invalid and discards this packet. In particular, OSPF does not verify the address of an unnumbered interface. In some scenarios, the source address may not meet the preceding requirement, and therefore OSPF address verification fails. For example, the negotiated peer address cannot be obtained on a P2P link. In this scenario, source address verification must be disabled to ensure that the OSPF adjacency can be properly set up.

Configuration Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Disabling Source Address Verification

Scenario Figure 4-20	
Configuration Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Set the network types of interfaces on all routers to P2P. ▪ Disable source address verification on all routers.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip ospf network point-to-point A(config-if-GigabitEthernet 0/1)# ip ospf source-check-ignore</pre>
B	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip ospf network point-to-point B(config-if-GigabitEthernet 0/1)# ip ospf source-check-ignore</pre>
Verification	<p>On Router A, verify that the OSPF neighbor information is correct.</p>
A	<pre>A# show ip ospfneighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.100.2.2 1 Full/- 00:00:34 192.100.2.2 GigabitEthernet 0/1</pre>

4.4.12 Disabling MTU Verification

Configuration

Effect

- The unicast routing service can be provided even if the MTUs of interfaces on neighbor routers are different.

Notes

- The OSPF basic functions must be configured.

Configuration Steps

Disabling MTU Verification

- (Optional) MTU verification is disabled by default. You are advised to retain the default configuration.
- This configuration is performed on two routers with different interface MTUs.

Verification

The adjacency can be set up between routers with different MTUs.

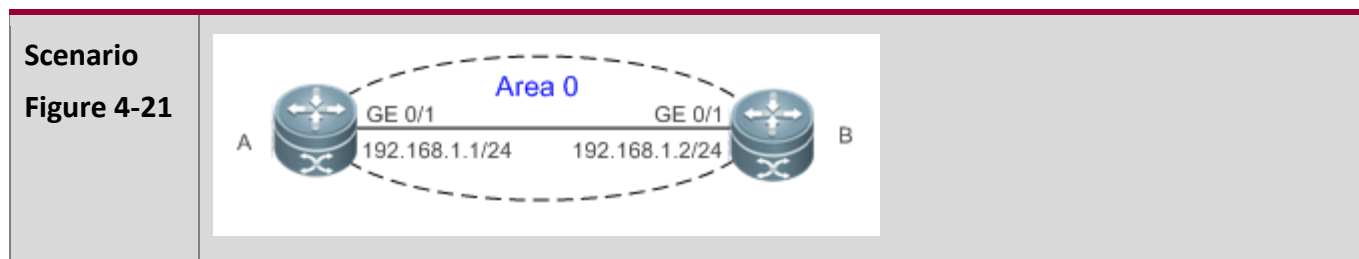
Related Commands

Disabling MTU Verification

Command	<code>ip ospf mtu-ignore</code>
Parameter Description	N/A
Command Mode	Interface configuration mode
Usage Guide	On receiving the database description packet, OSPF checks whether the MTU of the interface on the neighbor is the same as the MTU of its own interface. If the interface MTU specified in the received database description packet is greater than the MTU of the local interface, the adjacency cannot be set up. To resolve this problem, you can disable MTU verification.

Configuration Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."



Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Configure different MTUs for interfaces on two routers. ▪ Disable MTU verification on all routers. (By default, the function of disabling MTU verification is enabled.)
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip mtu 1400 A(config-if-GigabitEthernet 0/1)# ip ospf mtu-ignore</pre>
B	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip mtu 1600 B(config-if-GigabitEthernet 0/1)# ip ospf mtu-ignore</pre>
Verification	<ul style="list-style-type: none"> ▪ On Router A, verify that the OSPF neighbor information is correct.
A	<pre>A# show ip ospfneighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.2 1 Full/DR 00:00:34 192.168.1.2 GigabitEthernet 0/1</pre>

4.4.13 Enabling Two-Way Maintenance

Configuration

Effect

- Non-Hello packets can also be used to maintain the adjacency.

Notes

- The OSPF basic functions must be configured.

Configuration

Steps

Enabling Two-Way Maintenance

- (Optional) This function is enabled by default. You are advised to retain the default configuration.
- This configuration is performed on all routers.

Verification

Non-Hello packets can also be used to maintain the adjacency.

Related Commands

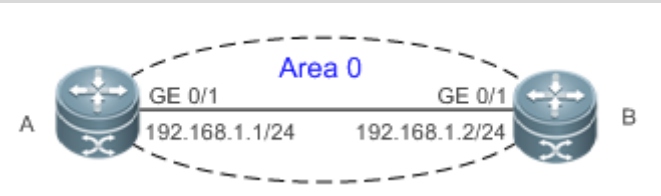
Enabling Two-Way Maintenance

Command	two-way-maintain
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	On a large network, a lot of packets may be sent or received, occupying too much CPU and memory. As a result, some packets are delayed or discarded. If the processing time of Hello packets exceeds the dead interval, the adjacency will be destroyed due to timeout. If the two-way maintenance function is enabled, in addition to the Hello packets, the DD, LSU, LSR, and LSack packets can also be used to maintain the bidirectional communication between neighbors when a large number of packets exist on the network. This prevents termination of the adjacency caused by delayed or discarded Hello packets.

Configuration

Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Scenario Figure 4-22	
Configuration Steps	<ul style="list-style-type: none">▪ Configure the interface IP addresses on all routers. (Omitted)▪ Configure the OSPF basic functions on all routers. (Omitted)▪ On Router A, enable the two-way maintenance function. (This function is enabled by default.)

A	<pre>A# configure terminal A(config)#routerospf 1 A(config-router)#two-way-maintain</pre>
Verification	When the adjacency is being set up, Router A checks the neighbor dead interval and updates the dead interval without waiting for Router B to send a Hello packet.
A	<pre>A# show ip ospfneighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Address Interface 192.168.1.2 1 Full/BDR 00:00:40 192.168.1.2 GigabitEthernet 0/1</pre>

4.4.14 Enabling GR

Configuration Effect

- When a distributed router switches services from the active board to the standby board, data forwarding continues and is not interrupted.
- When the OSPF process is being restarted, data forwarding continues and is not interrupted.

Notes

- The OSPF basic functions must be configured.
- The neighbor router must support the GR helper function.
- The grace period cannot be shorter than the neighbor dead time of the neighbor router.

Configuration Steps

Configuring the OSPF GR Function

- (Optional) This function is enabled by default. You are advised to retain the default configuration.
- This configuration is performed on all routers.

Configuring the OSPF GR Helper Function

- (Optional) This function is enabled by default. You are advised to retain the default configuration.
- This configuration is performed on all routers.

Verification

- When a distributed router switches services from the active board to the standby board, data forwarding continues and is not interrupted.
- When the OSPF process is being restarted, data forwarding continues and is not interrupted.

Related Commands

Configuring the OSPF GR Function

Command	<code>graceful-restart [grace-period <i>grace-period</i> inconsistent-lsa-checking]</code>
Parameter Description	<p>grace-period <i>grace-period</i>: Indicates the grace period, which is the maximum time from occurrence of an OSPF failure to completion of the OSPF GR. The value of the graceperiod varies from 1s to 1800s. The default value is 120s.</p> <p>inconsistent-lsa-checking: Enables topological change detection. If any topological change is detected, OSPF exits the GR process to complete convergence. After GR is enabled, topological change detection is enabled by default.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The GR function is configured based on the OSPF process. You can configure different parameters for different OSPF processes based on the actual conditions. This command is used to configure the GR restarter capability of a device. The grace period is the maximum time of the entire GR process, during which link status is rebuilt so that the original state of the OSPF process is restored. After the grace period expires, OSPF exits the GR state and performs common OSPF operations.</p> <p>Run the graceful-restart command to set the grace period to 120s. The graceful-restart grace-period command allows you to modify the grace period explicitly.</p> <p>The precondition for successful execution of GR and uninterrupted forwarding is that the topology remains stable. If the topology changes, OSPF quickly converges without waiting for further execution of GR, thus avoiding long-time forwarding black-hole.</p> <p>Disabling topology detection: If OSPF cannot converge in time when the topology changes during the hot standby process, forwarding black-hole may appear in a long time.</p> <p>Enabling topology detection: Forwarding may be interrupted when topology detection is enabled, but the interruption time is far shorter than that when topology detection is disabled.</p> <p>In most cases, it is recommended that topology detection be enabled. In special scenarios, topology detection can be disabled if the topology changes after the hot standby process, but it can be ensured that the forwarding black-hole will not appear in</p>

a long time. This can minimize the forwarding interruption time during the hot standby process.

If the Fast Hello function is enabled, the GR function cannot be enabled.

Configuring the OSPF GR Helper Function


Command	<code>graceful-restart helper { disable strict-lsa-checking internal-lsa-checking }</code>
Parameter Description	<p>disable: Prohibits a device from acting as a GR helper for another device.</p> <p>strict-lsa-checking: Indicates that changes in Type 1 to Type 5 and Type 7 LSAs will be checked during the period that the device acts as a GR helper to determine whether the network changes. If the network changes, the device will stop acting as the GR helper.</p> <p>internal-lsa-checking: Indicates that changes in Type 1 to Type 3 LSAs will be checked during the period that the device acts as a GR helper to determine whether the network changes. If the network changes, the device will stop acting as the GR helper.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>This command is used to configure the GR helper capability of a router. When a neighbor router implements GR, it sends a Grace-LSA to notify all neighbor routers. If the GR helper function is enabled on the local router, the local router becomes the GR helper on receiving the Grace-LSA, and helps the neighbor to complete GR. The disable option indicates that GR helper is not provided for any device that implements GR.</p> <p>After a device becomes the GR helper, the network changes are not detected by default. If any change takes place on the network, the network topology converges after GR is completed. If you wish that network changes can be quickly detected during the GR process, you can configure strict-lsa-checking to check Type 1 to 5 and Type 7 LSAs that indicate the network information or internal-lsa-checking to check Type 1 to 3 LSAs that indicate internal routes of the AS domain. When the network scale is large, it is recommended that you disable the LSA checking options (strict-lsa-checking and internal-lsa-checking) because regional network changes may trigger termination of GR and consequently reduce the convergence of the entire network.</p>

Configuration

Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Scenario
Figure 4-23



Remark The interface IP addresses are as follows:
S A: GE 0/1 192.168.1.1
 B: GE 0/1 192.168.1.1 GE 0/2 192.168.2.1 GE 0/3 192.168.3.1
 C: GE 0/1 192.168.4.2 GE 0/3 192.168.3.2
 D: GE 0/1 192.168.5.2 GE 0/2 192.168.2.2

Configuration Steps

- Configure the interface IP addresses on all routers. (Omitted)
- Configure the OSPF basic functions on all routers. (Omitted)
- On Router A, Router C, and Router D, enable the GR helper function. (This function is enabled by default.)
- On Router B, enable the GR function.

B

```
B# configure terminal
B(config)# router ospf1
B(config-router)# graceful-restart
```

Verification

- Trigger a hot standby switchover on Router B, and verify that the routing tables of destination networks 1 and 2 remain unchanged on Router A during the switchover.
- Trigger a hot standby switchover on Router B, ping destination network 1 from Router A, and verify that data forwarding is not interrupted during the switchover.

Common Errors

- Traffic forwarding is interrupted during the GR process because the configured grace period is shorter than the neighbor dead time of the neighbor router.

4.4.15 Enabling NSR

Configuration

Effect

- During the active/standby switchover of a distributed router, data forwarding continues and is not interrupted.

Notes

- The OSPF basic functions must be configured.

Configuration

Steps

Enabling the OSPF NSR Function

- (Optional) This function is disabled by default and enabled only when the function needs to be used.

Verification

- During the active/standby switchover of a distributed router, data forwarding continues and is not interrupted.

Related

Commands

Enabling NSR

Command	nsr
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>This command is used to enable the NSR function. Enable either NSR or GR for the same OSPF process. That is, when GR is enabled, NSR is automatically disabled. When NSR is enabled, GR is automatically disabled, but the GR helper capability is not affected.</p> <p>The switchover of a distributed router takes some time. If the OSPF neighbor dead time is shorter than the switchover time, the OSPF adjacency will be destroyed, causing service interruption during the switchover. Therefore, when enabling the NSR function, you are advised to configure an OSPF neighbor dead time that is equal to or greater than</p>

the default value. When the Fast Hello function is enabled, the OSPF neighbor dead time is shorter than 1s, and therefore it is recommended that the NSR function be disabled.

Configuration Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

<p>Scenario Figure 4-24</p>	
<p>Remarks</p>	<p>The interface IP addresses are as follows:</p> <p>A: GE 0/1 192.168.1.1</p> <p>B: GE 0/1 192.168.1.2 GE 0/2 192.168.2.1 GE 0/3 192.168.3.1</p> <p>C: GE 0/1 192.168.4.2 GE 0/3 192.168.3.2</p> <p>D: GE 0/1 192.168.5.2 GE 0/2 192.168.2.2</p>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) On Router B, enable the NSR function.
<p>B</p>	<pre>B# configure terminal B(config)# router ospf1 B(config-router)# nsr</pre>
<p>Verification</p>	<ul style="list-style-type: none"> Trigger a hot standby switchover on Router B, and verify that the routing tables of destination networks 1 and 2 remain unchanged on Router A during the switchover. Trigger a hot standby switchover on Router B, ping destination network 1 from Router A, and verify that data forwarding is not interrupted during the switchover.

Common Errors

- The configured OSPF neighbor dead interval is too short. If the Fast Hello function is enabled, the OSPF adjacency will be destroyed during the switchover, causing interruption of data forwarding.

4.4.16 Correlating OSPF with BFD

Configuration Effect

- Once a link is faulty, OSPF can quickly detect the failure of the route. This configuration helps shorten the traffic interruption time.

Notes

- The OSPF basic functions must be configured.
- The BFD parameters must be configured for the interface in advance.
- If BFD is configured for both a process and an interface, the configuration for the interface takes effect preferentially.

Configuration Steps

Correlating OSPF with BFD

- (Optional) This configuration is required if you wish to accelerate OSPF network convergence.
- The configuration must be performed on routers at both ends of the link.

Verification

- Run the **show bfd neighbor** command to verify that the BFD neighbors are normal.

Related Commands

Correlating an OSPF Interface with BFD

Command	ip ospf bfd [disable]
Parameter Description	disable: Disables BFD for link detection on a specified OSPF-enabled interface.
Command Mode	Interface configuration mode

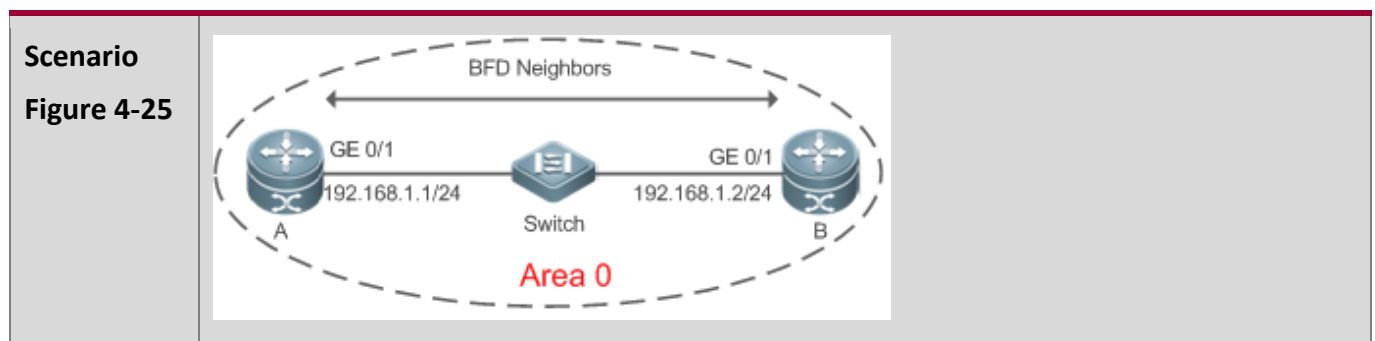
Usage Guide	<p>The interface-based configuration takes precedence over the bfd all-interfaces command used in process configuration mode.</p> <p>Based on the actual environment, you can run the ip ospf bfd command to enable BFD on a specified interface for link detection, or run the bfd all-interfaces command in OSPF process configuration mode to enable BFD on all interface of the OSPF process, or run the ospf bfd disable command to disable BFD on a specified interface.</p>
-------------	--

Correlatingan OSPF Process with BFD

Command	bfd all-interfaces
Parameter Description	N/A
Command Mode	OSPF process configuration mode
Usage Guide	<p>OSPF dynamically discovers neighbors through the Hello packets. After OSPF enables the BFD function, a BFD session will be set up to achieve the full adjacency, and use the BFD mechanism to detect the neighbor status. Once a neighbor failure is detected through BFD, OSPF performs network convergence immediately.</p> <p>You can also run the ip ospf bfd [disable] command in interface configuration mode to enable or disable the BFD function on a specified interface, and this configuration takes precedence over the bfd all-interfaces command used in OSPF process configuration mode.</p>

Configuration Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."



Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. (Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Configure the BFD parameters for interfaces of all routers. ▪ Correlate OSPF with BFD on all routers.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#bfd interval 200 min_rx 200 multiplier 5 A(config)# router ospf 1 A(config-router)#bfd all-interfaces</pre>
B	<pre>B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 2/1)#bfd interval 200 min_rx 200 multiplier 5 B(config)# router ospf 1 B(config-router)#bfd all-interfaces</pre>
Verification	<ul style="list-style-type: none"> ▪ On Router A and Router B, verify that the BFD state is Up. ▪ Disconnect Router A from the switch. On Router A, verify that a neighbor is found disconnected during BFD, and the corresponding OSPF route is deleted.
A	<pre>A# show ip ospf neighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State BFD State Dead Time Address Interface 192.168.1.2 1 Full/BDR Up 00:00:40 192.168.1.2 GigabitEthernet 0/1</pre>
B	<pre>B# show ip ospf neighbor OSPF process 1, 1 Neighbors, 1 is Full: Neighbor ID Pri State BFD State Dead Time Address Interface 192.168.1.1 1 Full/BDR Up 00:00:40 192.168.1.1 GigabitEthernet 0/1</pre>

4.4.17 Enabling Fast Reroute

Configuration

Effect

- Once OSPF detects a route failure, the router can immediately switch to the second-best route. This configuration helps shorten the traffic interruption time.

Notes

- The OSPF basic functions must be configured.
- The LAF configuration for fast reroute is mutually exclusive with the virtual link configuration.
- You must set **carrier-delay** of an interface to 0.

Configuration Steps

Configuring Fast Reroute

- (Optional) This configuration is required if you wish to increase the OSPF network convergence speed to the millisecond level.
- This configuration is performed on a router that has multiple paths to a destination network.

Preventing an Interface From Becoming a Standby Interface

- (Optional) This configuration is mandatory if you wish that data traffic is not switched over to a specified path after the best path fails. After the best path fails, the traffic will be switched over another second-best path, but a new best path will be selected based on the interface costs after OSPF converges again.
- This configuration is performed on a device where fast reroute is enabled.

Verification

Run the **show ip route fast-reroute** command to verify that both the best and second-best paths exist.

Related Commands

Configuring Fast Reroute

Command	fast-reroute { lfa [downstream-paths] route-map <i>route-map-name</i> }
Parameter Description	lfa : Enables computation of the loop-free standby path. downstream-paths : Enables computation of the downstream path. route-map <i>route-map-name</i> : Specifies a standby path through the route map.
Command Mode	OSPF routing process configuration mode
Usage Guide	If the lfa parameter is configured, computation of the loop-free standby path is enabled. In this case, you can use the interface mode command to specify the path protection mode of the interface.

It is recommended that computation of the loop-free standby path be disabled if any of the following case exists on the network:

1. Virtual links exist.
2. Alternative ABRs exist.
3. An ASBR is also an ABR.
4. Multiple ASBRs advertise the same external route.

If both **lfa** and **downstream-paths** are configured, computation of the downstream path is enabled.

If **route-map** is configured, a standby path can be specified for a matched route through the route-map.

When the OSPF fast reroute function is used, it is recommended that BFD be enabled at the same time so that the device can quickly detect any link failure and therefore shorten the forwarding interruption time. If the interface is up or down, to shorten the forwarding interruption time during OSPF fast reroute, you can configure **carrier-delay 0** in L3 interface configuration mode to achieve the fastest switchover speed.

Configuring the Interface LFA Protection

Command	ip ospf fast-reroute protection { node link-node disable}
Parameter Description	node: Enables the LFA node protection. link-node: Enables the LFA link node protection. disable: Disables LFA protection.
Command Mode	Interface configuration mode
Usage Guide	<p>If the fast-reroutelfa command is executed in OSPF route process configuration mode, the OSPF fast reroute computation function will be generated, and a standby route will be generated for the active route based on the LFA protection mode specified in interface configuration mode. Link protection is enabled by default for each OSPF interface. Under this protection mode, the failure of the active link does not affect data forwarding on the standby route.</p> <p>Use the node parameter to enable node protection for the interface, that is, data forwarding on the standby route will not be affected by the failure of a neighbor node corresponding to the active link.</p> <p>Use the link-node parameter to protect both the link and neighbor node corresponding to the active link.</p>

Use the **disable** parameter to disable the LFA protection function of the interface, that is, not to generate a standby entry for the route whose next hop is the interface.
This command does not take effect if **fast-reroute route-map** is configured.

Preventing an Interface From Becoming a Standby Interface


Command	ip ospf fast-reroute no-eligible-backup
Parameter Description	N/A
Command Mode	Interface configuration mode
Usage Guide	If the remaining bandwidth of an interface is small or if the interface and its active interface may fail at the same time, the interface cannot be used as a standby interface. Therefore, you need to run this command in interface configuration mode to prevent this interface from becoming a standby interface during OSPF fast reroute computation. After this command is executed, the standby interface is selected from other interface. This command does not take effect if fast-reroute route-map is configured.

Configuration Example

- The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring Fast Reroute

Scenario
Figure 4-26



	<p>Remark</p> <p>S</p> <p>The interface IP addresses are as follows:</p> <p>A: GE0/1 192.168.1.1 GE0/2 192.168.2.1</p> <p>B: GE0/1 192.168.1.2 GE0/2 192.168.3.1 GE0/3 192.168.4.1</p> <p>C: GE0/1 192.168.3.2 GE 0/2 192.168.2.2</p>
Configurati on Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Configure fast reroute on Router A. Configure carrier-delay 0 for the interface on Router A.
A	<pre>A# configure terminal A(config)# router ospf 1 A(config-router)# fast-reroute lfa A(config-router)# exit A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#carrier-delay 0 A(config-if-GigabitEthernet 0/1)# exit A(config)# interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)#carrier-delay 0</pre>
Verification	<p>On Router A, check the routing table and verify that a standby route exists for the entry 192.168.4.0/24.</p>
	<pre>A# show ip route fast-reroute begin 192.168.4.0 O 192.168.4.0/24 [ma] via 192.168.1.2, 00:39:28, GigabitEthernet 0/1 [b] via 192.168.2.2, 00:39:28, GigabitEthernet 0/2</pre>

4.4.18 Enabling iSPF

Configuration

Effect

- OSPF adopts the iSPF algorithm to compute the network topology.

Notes

- The OSPF basic functions must be configured.

Configuration

Steps

Configuring iSPF

- (Optional) This configuration is recommended if you wish to accelerate route convergence in a single area with more than 100 routers.
- This configuration is performed on all routers in the area.

Verification

Run the **show ip ospf** command to verify that iSPF is enabled.

Related Commands

Configuring iSPF


Command	ispf enable
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>After iSPF is enabled, OSPF will use the iSPF algorithm to compute the network topology. That is, after the network topology changes, OSPF corrects only the nodes affected by the topological change, instead of re-building the entire SPT.</p> <p>The iSPF function is generally used on a large-sized network to ease the pressure on router processors.</p>

Configuration Example

- ❗ The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring iSPF

Scenario
Figure 4-27



Remarks The interface IP addresses are as follows:
 Core1: GE0/1 192.168.1.1 GE0/2 192.168.2.1
 Core2: GE0/1 192.168.3.1 GE0/2 192.168.4.1
 Access1: GE0/1 192.168.1.2 GE 0/2 192.168.3.2
 Access2: GE0/1 192.168.4.2 GE 0/2 192.168.2.2

**Configurati
on Steps**

- Configure the interface IP addresses on all routers. (Omitted)
- Configure the OSPF basic functions on all routers. (Omitted)
- Configure iSPF on all routers.

Core1

```
Core1# configure terminal
Core1(config)# router ospf 1
Core1(config-router)# ispf enable
```

Core2

```
Core2# configure terminal
Core2(config)# router ospf 1
Core2(config-router)# ispf enable
```

Access1

```
Access1# configure terminal
Access1(config)# router ospf 1
Access1(config-router)# ispf enable
```

Access2	<pre> Access2# configure terminal Access2(config)# router ospf 1 Access2(config-router)# ispf enable </pre>
Verification	<p>On router Core1, verify that iSPF is enabled.</p>
	<pre> Core1# show ip ospf Routing Process "ospf 1" with ID 1.1.1.1 Process uptime is 17 hours 48 minutes Process bound to VRF default Memory Overflow is enabled. Router is not in overflow state now. Conforms to RFC2328, and RFC1583Compatibility flag is enabled Supports only single TOS(TOS0) routes Supports opaque LSA Enable two-way-maintain Enable ispf Initial SPF schedule delay 1000 msec Minimum hold time between two consecutive SPFs 5000 msec Maximum wait time between two consecutive SPFs 10000 msec Initial LSA throttle delay 0 msec Minimum hold time for LSA throttle 5000 msec Maximum wait time for LSA throttle 5000 msec Lsa Transmit Pacing timer 40 msec, 1 LS-Upd Minimum LSA arrival 1000 msec Pacing lsa-group: 30 sec Number of incoming current DD exchange neighbors 0/5 Number of outgoing current DD exchange neighbors 0/5 Number of external LSA 0. Checksum 0x000000 Number of opaque AS LSA 0. Checksum 0x000000 Number of non-default external LSA 0 External LSA database is unlimited. Number of LSA originated 2 </pre>

```
Number of LSA received 93
Log Neighbor Adjacency Changes : Enabled
Graceful-restart disabled
Graceful-restart helper support enabled
Number of areas attached to this router: 1: 1 normal 0 stub 0 nssa
Area 1
  Number of interfaces in this area is 1(1)
  Number of fully adjacent neighbors in this area is 0
  Number of fully adjacent virtual neighbors through this area is 0
  Area has no authentication
  SPF algorithm executed 0 times
  iSPF algorithm last executed 00:04:14.534 ago
  iSPF algorithm executed 12 times
  Number of LSA 1. Checksum 0x0029b3
```

4.4.19 Configuring the Network Management Function

Configuration

Effect

- Use the network management software to manage OSPF parameters and monitor the OSPF running status.

Notes

- The OSPF basic functions must be configured.
- You must enable the MIB function of the SNMP-Server before enabling the OSPF MIB function.
- You must enable the Trap function of the SNMP-Server before enabling the OSPF Trap function.
- You must enable the logging function of the device before outputting the OSPF logs.

Configuration

Steps

Binding the MIB with the OSPF Process

- (Optional) This configuration is required if you want to use the network management software to manage parameters of a specified OSPF process.
- This configuration is performed on all routers.

Enabling the Trap Function

- (Optional) This configuration is required if you want to use the network management software to monitor the OSPF running status.
- This configuration is performed on all routers.

Configuring the Logging Function

- (Optional) This function is enabled by default. You are advised to retain the default configuration. If you want to reduce the log output, disable this function.
- This configuration is performed on all routers.

Verification

- Use the network management software to manage the OSPF parameters.
- Use the network management software to monitor the OSPF running status.

Related

Commands

Binding the MIB with the OSPF Process

Command	enable mib-binding
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The OSPFv2 MIB does not have the OSPFv2 process information. Therefore, you must perform operations on a single OSPFv2 process through SNMP. By default, the OSPFv2 MIB is bound with the OSPFv2 process with the smallest process ID, and all user operations take effect on this process.</p> <p>If you wish to perform operations on a specified OSPFv2 through SNMP, run this command to bind the MIB with the process.</p>

Enabling the Trap Function

Command	enable traps[error [IfAuthFailure IfConfigError IfRxBadPacket VirtIfAuthFailure VirtIfConfigError VirtIfRxBadPacket] lsa [LsdbApproachOverflow LsdbOverflow MaxAgeLsa OriginateLsa] retransmit [IfTxRetransmit VirtIfTxRetransmit] state-change[IfStateChange NbrRestartHelperStatusChange NbrStateChange NssaTranslatorStatusChange RestartStatusChange VirtIfStateChange VirtNbrRestartHelperStatusChange VirtNbrStateChange]]
----------------	--

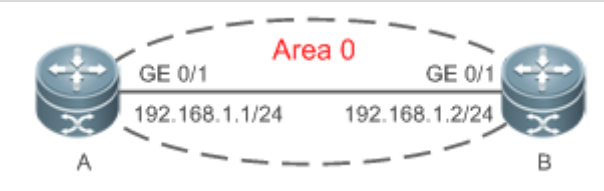
Parameter Description	<p>IfAuthFailure: Indicates that an interface authentication failure occurs.</p> <p>IfConfigError: Indicates that an interface parameter configuration error occurs.</p> <p>IfRxBadPacket: Indicates that the interface receives a bad packet.</p> <p>IfRxBadPacket: Indicates that the interface receives a bad packet.</p> <p>VirtIfAuthFailure: Indicates that a virtual interface authentication failure occurs.</p> <p>VirtIfConfigError: Indicates that a virtual interface parameter configuration error occurs.</p> <p>VirtIfRxBadPacket: Indicates that the virtual interface receives a bad packet.</p> <p>LsdbApproachOverflow: Indicates that the number of external LSAs has reached 90% of the upper limit.</p> <p>LsdbOverflow: Indicates that the number of external LSAs has reached the upper limit.</p> <p>MaxAgeLsa: Indicates that the LSA aging timer expires.</p> <p>OriginateLsa: Indicates that a new LSA is generated.</p> <p>IfTxRetransmit: Indicates that a packet is retransmitted on the interface.</p> <p>VirtIfTxRetransmit: Indicates that a packet is retransmitted on the virtual interface.</p> <p>IfStateChange: Indicates that interface state changes.</p> <p>NbrRestartHelperStatusChange: Indicates that the state of the neighbor GR process changes.</p> <p>NbrStateChange: Indicates that the neighbor state changes.</p> <p>NssaTranslatorStatusChange: Indicates that the NSSA translation state changes.</p> <p>RestartStatusChange: Indicates that the GR state of the local device changes.</p> <p>VirtIfStateChange: Indicates that the virtual interface state changes.</p> <p>VirtNbrRestartHelperStatusChange: Indicates that the GR state of the virtual neighbor changes.</p> <p>VirtNbrStateChange: Indicates that the virtual neighbor state changes.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The function configured by this command is restricted by the snmp-server command. You can configure snmp-server enable traps ospf and then enable traps command before the corresponding OSPF traps can be correctly sent out.</p> <p>This command is not restricted by the MIB bound with the process. The trap function can be enabled concurrently for different processes.</p>

Configuring the Logging Function

Command	log-adj-changes[detail]
Parameter Description	detail: Records all status change information.
Command Mode	OSPF routing process configuration mode
Usage Guide	N/A

Configuration Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Scenario Figure 4-28	
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Bind the MIB with the OSPF process on Router A. Enable the trap function on Router A.
A	<pre>A# configure terminal A(config)# snmp-server host 192.168.2.2 traps version 2c public A(config)# snmp-server community public rw A(config)# snmp-server enable traps A(config)# router ospf 10 A(config-router)# enable mib-binding A(config-router)# enable traps</pre>
Verification	Use the MIB tool to read and set the OSPF parameters and display the OSPF running status.

Common Errors

Configurations on the SNMP-Server are incorrect. For example, the MIB or trap function is not enabled.

4.4.20 Modifying Protocol Control Parameters

Configuration Effect

Modify protocol control parameters to change the protocol running status.

Notes

- The OSPF basic functions must be configured.
- The neighbor dead time cannot be shorter than the Hello interval.

Configuration Steps

Configuring the Hello Interval

- (Optional) You are advised to retain the default configuration.
- This configuration is performed on routers at both end of a link.

Configuring the Dead Interval

- (Optional) You are advised to retain the default configuration. This configuration can be adjusted if you wish to accelerate OSPF convergence when a link fails.
- This configuration is performed on routers at both end of a link.

Configuring LSU Retransmission Interval

- (Optional) You are advised to adjust this configuration if a lot of routes exist in the user environment and network congestion is serious.

Configuring the LSA Generation Time

- (Optional) You are advised to retain the default configuration.

Configuring the LSA Group Refresh Time

- (Optional) You are advised to retain the default configuration. This configuration can be adjusted if a lot of routes exist in the user environment.
- This configuration is performed on an ASBR or ABR.

Configuring LSA Repeated Receiving Delay

- (Optional) You are advised to retain the default configuration.

Configuring the SPF Computation Delay

- (Optional) This configuration can be adjusted if network flapping frequently occurs.

Configuring the Inter-Area Route Computation Delay

- (Optional) You are advised to retain the default configuration.
- This configuration is performed on all routers.

Configuring the External Route Computation Delay

- (Optional) You are advised to retain the default configuration.
- This configuration is performed on all routers.

Verification

Run the **show ip ospf** and **show ip ospf neighbor** commands to display the protocol running parameters and status.

Related Commands

Configuring the Hello Interval

Command	ip ospf hello-interval <i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the interval at which OSPF sends the Hello packet. The unit is second. The value ranges from 1 to 65,535.
Command Mode	Interface configuration mode
Usage Guide	The Hello interval is contained in the Hello packet. A shorter Hello interval indicates that OSPF can detect topological changes more quickly, but the network traffic increases. The Hello interval must be the same on all routers in the same network segment. If you want to manually modify the neighbor dead interval, ensure that the neighbor dead interval is longer than the Hello interval.

Configuring the Dead Interval

Command	ip ospf dead-interval <i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the time that the neighbor is declared lost. The unit is second. The value ranges from 0 to 2,147,483,647.
Command Mode	Interface configuration mode

Usage Guide	<p>The OSPF dead interval is contained in the Hello packet. If OSPF does not receive a Hello packet from a neighbor within the dead interval, it declares that the neighbor is invalid and deletes this neighbor record from the neighbor list. By default, the dead interval is four times the Hello interval. If the Hello interval is modified, the dead interval is modified automatically.</p> <p>When using this command to manually modify the dead interval, pay attention to the following issues:</p> <ol style="list-style-type: none"> 1. The dead interval cannot be shorter than the Hello interval. 2. The dead interval must be the same on all routers in the same network segment.
-------------	--

Configuring the LSU Transmission Delay

Command	ip ospf transmit-delayseconds
Parameter Description	<i>seconds</i> : Indicates the LSU transmission delay on the OSPF interface. The unit is second. The value ranges from 1 to 65,535.
Command Mode	Interface configuration mode
Usage Guide	<p>Before an LSU packet is transmitted, the Age fields in all LSAs in this packet will increase based on the amount specified by the ip ospf transmit-delay command. Considering the transmit and line propagation delays on the interface, you need to set the LSU transmission delay to a greater value for a low-speed line or interface. The LSU transmission delay of a virtual link is defined by the transmit-delay parameter in the area virtual-link command.</p> <p>If the value of the Age field of an LSA reaches 3600, the packet will be retransmitted or a retransmission will be requested. If the LSA is not updated in time, the expired LSA will be deleted from the LSDB.</p>

Configuring LSU Retransmission Interval

Command	ip ospf retransmit-intervalseconds
Parameter Description	<i>seconds</i> : Indicates the LSU retransmission interval. The unit is second. The value ranges from 0 to 65,535. This interval must be longer than the round-trip transmission delay of data packets between two neighbors.
Command Mode	Interface configuration mode

Usage Guide	<p>After a router finishes sending an LSU packet, this packet is still kept in the transmit buffer queue. If an acknowledgment from the neighbor is not received within the time defined by the ip ospf retransmit-interval command, the router retransmits the LSU packet.</p> <p>The retransmission delay can be set to a greater value on a serial line or virtual link to prevent unnecessary retransmission. The LSU retransmission delay of a virtual link is defined by the retransmit-interval parameter in the area virtual-link command.</p>
-------------	---

Configuring the LSA Generation Time

Command	timers throttle lsa all <i>delay-time hold-time max-wait-time</i>
Parameter Description	<p><i>delay-time</i>: Indicates the minimum delay for LSA generation. The first LSA in the database is always generated instantly. The value ranges from 0 to 600,000. The unit is ms.</p> <p><i>hold-time</i>: Indicates the minimum interval between the first LSA update and the second LSA update. The value ranges from 1 to 600,000. The unit is ms.</p> <p><i>max-wait-time</i>: Indicates the maximum interval between two LSA updates when the LSA is updated continuously. This interval is also used to determine whether the LSA is updated continuously. The value ranges from 1 to 600,000. The unit is ms.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>If a high convergence requirement is raised when a link changes, you can set delay-time to a smaller value. You can also appropriately increase values of the preceding parameters to reduce the CPU usage.</p> <p>When configuring this command, the value of hold-time cannot be smaller than the value of delay-time, and the value of max-wait-time cannot be smaller than the value of hold-time.</p>

Configuring the LSA Group Refresh Time

Command	timers pacinglsa-group <i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the LSA group pacing interval. The value ranges from 10 to 1,800. The unit is second.
Command Mode	OSPF routing process configuration mode

Usage Guide	<p>Every LSA has a time to live (LSA age). When the LSA age reaches 1800s, a refreshment is needed to prevent LSAs from being cleared because their ages reaching the maximum. If LSA update and aging computation are performed for every LSA, the device will consume a lot of CPU resources. In order to use CPU resources effectively, you can refresh LSAs by group on the device. The interval of group refreshment is called group pacing interval. The group refreshment operation is to organize the LSAs generated within a group pacing interval into a group and refresh the group as a whole.</p> <p>If the total number of LSAs does not change, a larger group pacing interval indicates that more LSAs need to be processed after timeout. To maintain the CPU stability, the number of LSAs processes upon each timeout cannot be too large. If the number of LSAs is large, you are advised to reduce the group pacing interval. For example, if there are 1000 LSAs in the database, you can reduce the pacing interval; if there are 40 to 100 LSAs, you can set the pacing interval to 10-20 minutes.</p>
-------------	--

Configuring the LSA Group Refresh Interval

Command	timers pacing lsa-transmit <i>transmit-time transmit-count</i>
Parameter Description	<p><i>transmit-time</i>: Indicates the LSA group transmission interval. The value ranges from 10 to 1,000. The unit is ms.</p> <p><i>transmit-count</i>: Indicates the number of LS-UPD packets in a group. The value ranges from 1 to 200.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>If the number of LSAs is large and the device load is heavy in an environment, properly configuring transmit-time and transmit-count can limit the number of LS-UPD packets flooded on a network.</p> <p>If the CPU usage is not high and the network bandwidth load is not heavy, reducing the value of transmit-time and increasing the value of transmit-count can accelerate the environment convergence.</p>

Configuring LSA Repeated Receiving Delay

Command	timers lsa arrival <i>arrival-time</i>
Parameter Description	<i>arrival-time</i> : Indicates the delay after which the same LSA is received. The value ranges from 0 to 600,000. The unit is ms.

Command Mode	OSPF routing process configuration mode
Usage Guide	No processing is performed if the same LSA is received within the specified time.

Configuring the Inter-Area Route Computation Delay

Command	timers throttle route inter-area <i>ia-delay</i>
Parameter Description	<i>ia-delay</i> : Indicates the inter-area route computation delay. The unit is ms. The value ranges from 0 to 600,000.
Command Mode	OSPF routing process configuration mode
Usage Guide	This delay cannot be modified if strict requirements are raised for the network convergence time.

Configuring the External Route Computation Delay

Command	timers throttle route ase <i>ase-delay</i>
Parameter Description	<i>ase-delay</i> : Indicates the external route computation delay. The unit is ms. The value ranges from 0 to 600,000.
Command Mode	OSPF routing process configuration mode
Usage Guide	This delay cannot be modified if strict requirements are raised for the network convergence time.

Configuring the SPF Computation Delay

Command	timers throttle spf <i>spf-delay spf-holdtime spf-max-waittime</i>
Parameter Description	<p><i>spf-delay</i>: Indicates the SPF computation delay. The unit is ms. The value ranges from 1 to 600,000. When detecting a topological change, the OSPF routing process triggers the SPF computation at least after spf-delay elapses.</p> <p><i>spf-holdtime</i>: Indicates the minimum interval between two SPF computations. The unit is ms. The value ranges from 1 to 600,000.</p>

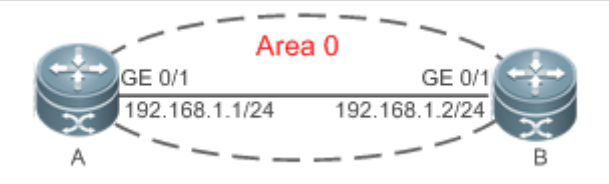
	<p><i>spf-max-waittime</i>: Indicates the maximum interval between two SPF computations. The unit is ms. The value ranges from 1 to 600,000.</p> <p><i>number</i>: indicates the metric of the summarized route.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>spf-delay indicates the minimum time between the occurrence of the topological change and the start of SPF computation. spf-holdtime indicates the minimum interval between the first SPF computation and the second SPF computation. After that, the interval between two SPF computations must be at least twice of the previous interval. When the interval reaches spf-max-waittime, the interval cannot increase again. If the interval between two SPF computations already exceeds the required minimum value, the interval is computed by starting from spf-holdtime.</p> <p>You can set spf-delay and spf-holdtime to smaller values to accelerate topology convergence, and set spf-max-waittime to a larger value to reduce SPF computation. Flexible settings can be used based on stability of the network topology.</p> <p>Compared with the timers spf command, this command supports more flexible settings to accelerate the convergence speed of SPF computation and further reduce the system resources consumed by SPF computation when the topology continuously changes. Therefore, you are advised to use the timers throttle spf command for configuration.</p> <p>The value of spf-holdtime cannot be smaller than the value of spf-delay; otherwise, spf-holdtime will be automatically set to the value of spf-delay.</p> <p>The value of spf-max-waittime cannot be smaller than the value of spf-holdtime; otherwise, spf-max-waittime will be automatically set to the value of spf-holdtime.</p> <p>The configurations of timers throttle spf and timers spf are mutually overwritten.</p> <p>When both timers throttle spf and timers spf are not configured, the default values of timers throttle spf prevail.</p>

Configuration

Example

- i** The following configuration examples assume that the OSPF basic functions have been configured. For details about the OSPF basic functions, see section 4.4.1 "Configuring OSPF Basic Functions."

Configuring the Hello Interval and Dead Interval

Scenario Figure 4-29	
Configuration Steps	Configure the interface IP addresses on all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Configure the Hello interval and dead interval on all routers.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip ospf hello-interval 15 A(config-if-GigabitEthernet 0/1)# ip ospf dead-interval 50</pre>
B	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip ospf hello-interval 15 A(config-if-GigabitEthernet 0/1)# ip ospf dead-interval 50</pre>
Verification	Check the interface parameters on Router A. Verify that the Hello interval is 10s and the dead interval is 50s.
A	<pre>A# show ip ospf interface GigabitEthernet 0/1 is up, line protocol is up Internet Address 192.168.1.1/24, Ifindex 2, Area 0.0.0.0, MTU 1500 Matching network config: 192.168.1.0/24 Process ID 1, Router ID 192.168.1.2, Network Type POINTOMULTIPOINT, Cost: 1 Transmit Delay is 1 sec, State Point-To-Point Timer intervals configured, Hello 15, Dead 50, Wait 40, Retransmit 5 Hello due in 00:00:02 Neighbor Count is 1, Adjacent neighbor count is 0 Crypt Sequence Number is 4787 Hello received 465 sent 466, DD received 8 sent 8 LS-Req received 2 sent 2, LS-Upd received 8 sent 21</pre>

LS-Ack received 14 sent 7, Discarded 3

Common Errors

- The configured neighbor dead time is shorter than the Hello interval.

4.5 Monitoring

Clearing

! Running the **clear** commands may lose vital information and thus interrupt services.

Description	Command
Clears and resets an OSPF process.	clear ip ospf [process-id] process

Displaying

Description	Command
Displays the OSPF process configurations.	show ip ospf [process-id]
Displays the OSPF internal routing table, including routes to ABRs and ASBRs.	show ip ospf [process-id] border-routers
Displays information about the OSPF LSDB.	show ip ospf [process-id area-id] database [{ asbr-summary external network nssa-external opaque-area opaque-as opaque-link router summary }] [{ adv-router ip-address self-originate } link-state-id brief] [database-summary max-age detail]
Displays OSPF-enabled interfaces.	show ip ospf [process-id] interface [interface-type interface-number brief]
Displays the OSPF neighbor list.	show ip ospf [process-id] neighbor [detail] [interface-type interface-number] [neighbor-id]

Displays the OSPF routing table.	show ip ospf [process-id] route[count]
Displays the number of times SPT is computed in the OSPF area.	show ip ospf [process-id] spf
Displays the summarized route of OSPF redistributed routes.	show ip ospf[process-id] summary-address
Displays the OSPF network topology information.	show ip ospf [process-id[area-id]] topology[adv-routeradv-router-id [router-id] self-originate[router-id]]
Displays OSPF virtual links.	show ip ospf [process-id] virtual-links [ip-address]

Debugging

! System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Debugs OSPF events.	debug ip ospf events [abr asbr lsa nssa os restart router slink vlink]
Debugs OSPF interfaces.	debug ip ospf ifsm [events status timers]
Debugs OSPF neighbors.	debug ip ospf nfm [events status timers]
Debugs the OSPF NSM.	debug ip ospf nsm [interface redistribute route]
Debugs OSPF LSAs.	debug ip ospf lsa [flooding generate install maxage refresh]
Debugs OSPF packets.	debug ip ospf packet [dd detail hello ls-ack ls-request ls-update recv send]
Debugs OSPF routes.	debug ip ospf route [ase ia install spf time]

5 CONFIGURING OSPFV3

5.1 Overview

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) that is used within the Autonomous System (AS) to allow routers to obtain a route to a remote network.

- i** OSPF Version 2 (OSPFv2) is applicable to IPv4, and OSPF Version 3 (OSPFv3) is applicable to IPv6. The protocol running mechanism and most configurations are the same.

OSPF has the following characteristics:

- Wide scope of application: OSPF is applicable to a larger-scale network that supports hundreds of routers.
- Fast convergence: Once the network topology changes, notifications can be quickly sent between routers to update routes.
- No self-loop: Only the link status information is synchronized between routers. Each router computes routes independently, and a self-loop will not occur.
- Area division: A large routing domain is divided into multiple small areas to save system resources and network bandwidth and ensure stability and reliability of routes.
- Route classification: Routes are classified into several types to support flexible control.
- Equivalent routes: OSPF supports equivalent routes.
- Authentication: OSPF supports packet authentication to ensure security of protocol interaction.
- Multicast transmission: Protocol packets are sent using the multicast address to avoid interfering with irrelevant entities and save system resources.

- i** In this chapter, the term "router" refers to any network device that supports the routing function. These network devices can be L3 switches, routers, or firewall.

- i** Unless otherwise specified, "OSPF" in the following descriptions refers to OSPFv3.

Protocols and Standards

RFC2740	This document describes the modifications to OSPF to support version 6 of the Internet Protocol (IPv6).
draft-ietf-ospf-ospfv3-graceful-restart	This document describes the OSPFv3 graceful restart. The OSPFv3 graceful restart is identical to OSPFv2 except for the differences described in this document. These differences include the format of the grace Link State Advertisements (LSA) and other considerations.

draft-ietf-ospf-ospfv3-mib-11	This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in IPv6-based internets. In particular, it defines objects for managing the Open Shortest Path First Routing Protocol for IPv6.

5.2 Applications


Application	Description
Intra-Domain Interworking	OSPF runs within the AS, which is divided into several areas.
Inter-Domain Interworking	Several ASs are interconnected. OSPF runs within each AS, and BGP runs between ASs.

5.2.1 Intra-Domain Interworking

Scenario

OSPF runs within the AS. If the number of routers exceeds 40, it is recommended that the AS be divided into several areas. Generally, high-end devices featuring reliable performance and fast processing speed are deployed in a backbone area, and low-end or medium-range devices with relatively lower performance can be deployed in a normal area. All normal areas must be connected to the backbone area. It is recommended that a normal area located on the stub be configured as a stub area. As shown in Figure 5-1, the network is divided into four areas. Communication between these areas must go through the backbone area, that is, area 0.

Figure 5-1 Division of the OSPF Areas



Remark S	A, B, C, D, E, and H are located in the backbone area, and are backbone routers. Area 3 is configured as a stub area.
-------------	--

Deployment

- OSPF runs on all routers within the AS to implement unicast routing.


5.2.2 Inter-Domain Interworking

Scenario

Several ASs are interconnected. OSPF runs within each AS, and BGP runs between ASs. Generally, OSPF and BGP learn the routing information from each other.

As shown in Figure 5-2, unicast routing is implemented within AS 100 and AS 200 using OSPF, and between the two ASs using BGP.

Figure 5-2 Interworking Between OSPF and BGP



Remarks

OSPF and BGP run concurrently on Router A and Router D.

Deployment

- OSPF runs within AS 100 and AS 200 to implement unicast routing.
- BGP runs between the two ASs to implement unicast routing.

5.3 Features

Basic

Concepts

Routing Domain

All routers in an AS must be interconnected and use the same routing protocol. Therefore, an AS is also called a routing domain.

An AS on which OSPF runs is also called OSPF routing domain, or OSPF domain for short.

OSPF Process

OSPF supports multiple instances, and each instance corresponds to an OSPF process.

One or more OSPF processes can be started on a router. Each OSPF process runs OSPF independently, and the processes are mutually isolated.

An OSPF packet header contains the Instance ID field, and multiple OSPF instances can run concurrently on a single link. The process ID is valid only on the local device.

RouterID

The router ID uniquely identifies a router in an OSPF domain. Router IDs of any two routers cannot be the same.

If multiple OSPF processes exist on a router, each OSPF process uses one router ID. Router IDs of any two OSPF processes cannot be the same.


Area

OSPF supports multiple areas. An OSPF domain is divided into multiple areas to ease the computing pressure of a large-scale network.

An area is a logical group of routers, and each group is identified by an area ID. The border between areas is a router. A router may belong to one area or multiple areas. One network segment (link) can belong to only one area, or each OSPF-enabled interface must belong to a specified area.

Area 0 is the backbone area, and other areas are normal areas. Normal areas must be directly connected to the backbone area.

Figure 5-3 Division of the OSPF Areas



OSPF Router

The following types of routers are defined in OSPF, and assigned with different responsibilities:


- Internal router
- All interface of an interval router belong to the same OSPF area. As shown in Figure 5-3, A, C, F, G, I, M, J, K, and L are internal routers.
- Area border router (ABR)

- An ABR is used to connect the backbone area with a normal area. An ABR belongs to two or more areas, and one of the areas must be the backbone area. As shown in Figure 5-3, B, D, E, and H are ABRs.
- Backbone router
- A backbone router has at least one interface that belongs to the backbone area. All ABRs and all routers in area 0 are backbone routers. As shown in Figure 5-3, A, B, C, D, E, and H are backbone routers.
- AS boundary router (ASBR)
- An ASBR is used to exchange routing information with other ASs. An ASBR is not necessarily located on the border of an AS. It may be a router inside an area, or an ABR. As shown in Figure 5-3, A is an ASBR.

Virtual Link


OSPF supports virtual links. A virtual link is a logical link that belongs to the backbone area. It is used to resolve the problems such as a discontinuous backbone area or a failure to directly connect a normal area to the backbone area on the physical network. A virtual link supports traversal of only one normal area, and this area is called transit area. Routers on both ends of a virtual link are ABRs.

Figure 5-4 Discontinuous Backbone Area on the Physical Network



As shown in Figure 5-4, a virtual link is set up between A and B to connect two separated parts of Area 0. Area 1 is a transit area, and A and B are ABRs of Area 1.

Figure 5-5 Failure to Directly Connect a Normal Area to the Backbone Area on the Physical Network



As shown in Figure 5-5, a virtual link is set up between A and B to extend Area 0 to B so that Area 0 can be directly connected to Area 2 on B. Area 1 is a transit area, A is an ABR of Area 1, and B is an ABR of Area 0 and Area 2.

LSA

OSPF describes the routing information by means of Link State Advertisement (LSA).

LSA Type	Description
Router-LSA(Type1)	This LSA is originated by every router. It describes the link state and cost of the router, and is advertised only within the area where the originating router is located.
Network-LSA(Type2)	This LSA is originated by a designated router (DR). It describes the state of the current link, and is advertised only within the area where the DR is located.
Inter-Area-Prefix-LSA(Type3)	This LSA is originated by an ABR. It describes a route to another area, and is advertised to areas except totally stub areas or Not-So-Stubby Area (NSSA) areas.
Inter-Area-Router-LSA(Type4)	This LSA is originated by an ABR. It describes a route to an ASBR, and is advertised to areas except areas where the ASBR is located.
AS-external-LSA(Type5)	This LSA is originated by an ABR. It describes a route to a destination outside the AS, and is advertised to all areas except the stub and NSSA areas.
NSSA LSA(Type7)	This LSA is originated by an ABR. It describes a route to a destination outside the AS, and is advertised only within the NASSA areas.
Link-LSA(Type8)	This LSA is originated by every router. It describes the link-local address and IPv6 prefix address of each link, and provides the link option that will be set in the Network-LSA. It advertised only on the current link.
Intra-Area-Prefix-LSA(Type9)	Every router or DR generates one or more Intra-Area-Prefix-LSAs, which are advertised in the area to which the router or DR belongs. <ul style="list-style-type: none"> ▪ The Intra-Area-Prefix-LSA generated by a router describes the IPv6 prefix address associated with the Route-LSA.

	<ul style="list-style-type: none"> ▪ The Intra-Area-Prefix-LSA generated by a DR describes the IPv6 prefix address associated with the Network-LSA.
--	--

i Stub areas, NSSA areas, totally stub areas, and totally NSSA areas are special forms of normal areas and help reduce the load of routers and enhance reliability of OSPF routes.

OSPF Packet

The following table lists the protocol packets used by OSPF. These OSPF packets are encapsulated in IP packets and transmitted in multicast or unicast mode.

Packet Type	Description
Hello	Hello packets are sent periodically to discover and maintain OSPF neighbor relationships.
Database Description (DD)	DD packets carry brief information about the local Link-State Database (LSDB) and are used to synchronize the LSDBs between OSPF neighbors.
Link State Request (LSR)	LSR packets are used to request the required LSAs from neighbors. LSR packets are sent only after DD packets are exchanged successfully between OSPF neighbors.
Link State Update (LSU)	LSU packets are used to send the required LSAs to peers.
Link State Acknowledgment (LSAck)	LSAck packets are used to acknowledge the received LSAs.

Overview

Feature	Description
Link-State Routing Protocols	Run OSPF on the router to obtain routes to different destinations on the network.
OSPF Route Management	Properly plan or optimize OSPF routes through manual configuration to implement management of OSPF routes.
Enhanced Security and Reliability	Use functions such as authentication and BFD correlation to enhance security, stability, and reliability of OSPF.

5.3.1 Link-State Routing Protocols

OSPF is a type of link-state routing protocols. Its working process is as follows:

- Neighbor discovery → Bidirectional communication
- An OSPF neighbor relationship is set up between adjacent routers, and bidirectional communication is maintained.
- Database synchronization → Full adjacency
- A router uses LSAs to advertise all its link states. LSAs are exchanged between neighbors and the link state database (LSDB) is synchronized to achieve full adjacency.
- Shortest Path Tree (SPT) computation → Formation of a routing table
- The router computes the shortest path to each destination network based on the LSDB and forms an OSPF routing table.

Working Principle

Neighbor Discovery → Bidirectional Communication

Routers send Hello packets through all OSPF-enabled interfaces (or virtual links). If Hello packets can be exchanged between two routers, and parameters carried in the Hello packets can be successfully negotiated, the two routers become neighbors. Routers that are mutually neighbors find their own router IDs from Hello packets sent from neighbors, and bidirectional communication is set up.

A Hello packet includes, but is not limited to, the following information:

- Router ID of the originating router
- Area ID of the originating router interface (or virtual link)
- Instance ID of the originating router interface (or virtual link)
- Interface ID of the originating router interface (or virtual link)
- Priority of the originating router interface (used for DR/BDR election)
- Hello interval of the originating router interface (or virtual link)
- Neighbor dead interval of the originating router interface (or virtual link)
- IP addresses of the DR and Backup Designated Router (BDR)
- Router ID of the neighbor of the originating router

Database Synchronization → Full Adjacency

After bidirectional communication is set up between neighbor routers, the DD, LSR, LSU, and LSack packets are used to exchange LSAs and set up the adjacency. The brief process is as follows:

- A router generates an LSA to describe all link states on the router.
- The LSA is exchanged between neighbors. When a router receives the LSA from its neighbor, it copies the LSA and saves the copy in the local LSDB, and then advertises the LSA to other neighbors.
- When the router and its neighbors obtain the same LSDB, full adjacency is achieved.

i OSPF will be very quiet without changes in link costs or network addition or deletion. If any change takes place, the changed link states are advertised to quickly synchronize the LSDB.

SPT Computation → Formation of a Routing Table

After the complete LSDB is obtained from the router, the Dijkstra algorithm is run to generate an SPT from the local router to each destination network. The SPT records the destination networks, next-hop addresses, and costs. OSPF generates a routing table based on the SPT.

If changes in link costs or network addition or deletion take place, the LSDB will be updated. The router again runs the Dijkstra algorithm, generates a new SPT, and updates the routing table.

i The Dijkstra algorithm is used to find a shortest path from a vertex to other vertices in a weighted directed graph.

OSPF Network Types


A router does not necessarily need to exchange LSAs with every neighbor and set up an adjacency with every neighbor. To improve efficiency, OSPF classifies networks that use various link layer protocols into five types so that LSAs are exchanged in different ways to set up an adjacency:

- Broadcast
- Neighbors are discovered, and the DR and BDR are elected.
- The DR (or BDR) exchanges LSAs with all other routers to set up an adjacency. Except the DR and BDR, all other routers do not exchange LSAs with each other, and the adjacency is not set up.
- Ethernet and fiber distributed data interface (FDDI) belong to the broadcast network type by default.
- Non-broadcast multiple access (NBMA)
- Neighbors are manually configured, and the DR and BDR are elected.
- The DR (or BDR) exchanges LSAs with all other routers to set up an adjacency. Except the DR and BDR, all other routers do not exchange LSAs with each other, and the adjacency is not set up.
- X.25, frame relay, and ATM belong to NBMA networks by default.
- Point-to-point (P2P)
- Neighbors are automatically discovered, and the DR or BDR is not elected.
- LSAs are exchanged between routers at both ends of the link, and the adjacency is set up.
- PPP, HDLC, and LAPB belong to the P2P network type by default.
- Point-to-multipoint(P2MP)
- Neighbors are automatically discovered, and the DR or BDR is not elected.
- LSAs are exchanged between any two routers, and the adjacency is set up.

- Networks without any link layer protocol belong to the P2MP network type by default.
- P2MP broadcast
- Neighbors are manually configured, and the DR or BDR is not elected.
- LSAs are exchanged between any two routers, and the adjacency is set up.
- Networks without any link layer protocol belong to the P2MP network type by default.

OSPF Route Types

Figure 5-6



Display the OSPF routes (marked in red) in the routing table of Router C.

```
C#show ipv6 route ospf
```

IPv6 routing table name is Default(0) global scope - 7 entries

Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP

I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary

O - OSPF intra area, OI - OSPF inter area, OE1 - OSPF external type 1, OE2 - OSPF external type 2

ON1 - OSPF NSSA external type 1, ON2 - OSPF NSSA external type 2

[*] - NOT in hardware forwarding table

L ::1/128 via Loopback, local host

OI 3001::/64 [110/2] via FE80::21A:A9FF:FE15:4CB9, VLAN 200

C 3001:1::/64 via VLAN 200, directly connected

L 3001:1::2/128 via VLAN 200, local host

L FE80::/10 via ::1, Null0

C FE80::/64 via VLAN 200, directly connected

L FE80::21A:A9FF:FE01:FB1F/128 via VLAN 200, local host

A mark is displayed in front of each OSPF route to indicate the type of the route. There are six types of OSPF routes:

- O: Intra-area route
- This type of route describes how to arrive at a destination network in the local area. The cost of this type of route is equal to the cost of the route from the local router to the destination network.
- OI: Inter-area route
- This type of route describes how to arrive at a destination network in another area. The cost of this type of route is equal to the cost of the route from the local router to the destination network.
- OE1: Type 1 external route
- This type of route describes how to arrive at a destination network outside the AS. The cost of this type of route is equal to the cost of the route from the local router to the ASBR plus the cost of the route from the ASBR to the destination network. This type of route does not exist on routers in the stub or NSSA area.
- OE2: Type 2 external route
- This type of route describes how to arrive at a destination network outside the AS. The cost of this type of route is equal to the cost of the route from the ASBR to the destination network. This type of route does not exist on routers in the stub or NSSA area.
- ON1: Type 1 external route of the NSSA area
- This type of route describes how to arrive at a destination network outside the AS through the ASBR in the NSSA area. The cost of this type of route is equal to the cost of the route from the local router to the ASBR plus the cost of the route from the ASBR to the destination network. This type of route exists only on routers in the NSSA area.
- ON2: Type 2 external route of the NSSA area
- This type of route describes how to arrive at a destination network outside the AS through the ASBR in the NSSA area. The cost of this type of route is equal to the cost of the route from the ASBR to the destination network. This type of route exists only on routers in the NSSA area.

i Reliability of OE2 and ON2 routes is poor. OSPF believes that the cost of the route from the ASBR to a destination outside an AS is far greater than the cost of the route to the ASBR within the AS. Therefore, when the route cost is computed, only the cost of the route from the ASBR to a destination outside an AS is considered.

Related Configuration

Enabling OSPF

OSPF is disabled by default.

Run the **ipv6 router ospf 1** command to create an OSPF process on the router.

Run the **ipv6 ospf area** command to enable OSPF on an interface and specify the area ID.

Run the **area virtual-link** command to create a virtual link on the router. The virtual link can be treated as a logical interface.

Router ID

By default, the OSPF process elects the largest IPv4 address among the IPv4 addresses of all the loopback interfaces as the router ID. If the loopback interfaces configured with IPv4 addresses are not available, the OSPF process elects the largest IPv4 address among the IPv4 addresses of all the physical ports as the router ID.

Alternatively, you can run the **router-id** command to manually specify the router ID.

Protocol Control Parameters

Run the **ipv6 ospf hello-interval** command to modify the Hello interval on the interface. The default value is 10s (or 30s for NBMA networks).

Run the **ipv6 ospf dead-interval** command to modify the neighbor dead interval on the interface. The default value is four times the Hello interval.

Use the **poll-interval** parameter in the **ipv6 ospf neighbor** command to modify the neighbor polling interval on the NBMA interface. The default value is 120s.

Run the **ipv6 ospf transmit-delay** command to modify the LSU packet transmission delay on the interface. The default value is 1s.

Run the **ipv6 ospf retransmit-interval** command to modify the LSU packet retransmission interval on the interface. The default value is 5s.

Use the **hello-interval** parameter in the **area virtual-link** command to modify the Hello interval on the virtual link. The default value is 10s.

Use the **dead-interval** parameter in the **area virtual-link** command to modify the neighbor dead interval on the virtual link. The default value is four times the Hello interval.

Use the **transmit-delay** parameter in the **area virtual-link** command to modify the LSU packet transmission delay on the virtual link. The default value is 1s.

Use the **retransmit-interval** parameter in the **area virtual-link** command to modify the LSU packet retransmission interval on the virtual link. The default value is 5s.

Run the **timers throttle lsa all** command to modify parameters of the exponential backoff algorithm that generates LSAs. The default values of these parameters are 0 ms, 5000 ms, and 5000 ms.

Run the **timers pacing lsa-group** command to modify the LSA group update interval. The default value is 30s.

Run the **timers pacing lsa-transmit** command to modify the LS-UPD packet sending interval and the number of sent LS-UPD packets. The default values are 40 ms and 1.

Run the **timers lsa arrival** command to modify the delay after which the same LSA is received. The default value is 1000 ms.

Run the **timers throttle spf** command to modify the SPT computation delay, minimum interval between two SPT computations, and maximum interval between two SPT computations. The default values are 1000 ms, 5000 ms, and 10000 ms.

OSPF Network Types

By default, Ethernet and FDDI belong to the broadcast type, X.25, frame relay, and ATM belong to the NBMA type, and PPP, HDLC, and LAPB belong to the P2P type.

Run the **ipv6 ospf network** command to manually specify the network type of an interface.

Run the **ipv6 ospf neighbor** command to manually specify a neighbor. For the NBMA and P2MP non-broadcast types, you must manually specify neighbors.

Run the **ipv6 ospf priority** command to adjust the priorities of interfaces, which are used for DR/BDR election. The DR/BDR election is required for the broadcast and NBMA types. The router with the highest priority wins in the election, and the router with the priority of 0 does not participate in the election. The default value is 1.

5.3.2 OSPF Route Management

Properly plan or optimize OSPF routes through manual configuration to implement management of OSPF routes.

Working Principle

(Totally) Stub Area and (Totally) NSSA Area

The (totally) stub and (totally) NSSA areas help reduce the protocol interaction load and the size of the routing table.

- If an appropriate area is configured as a (totally) stub or NSSA area, advertisement of a large number of Type 5 and Type 3 LSAs can be avoided within the area.

Area	Type 1 and Type 2 LSAs	Type 3 LSA	Type 4 LSA	Type 5 LSA	Type 7 LSA
Non (totally) stub area and NSSA area	Allowed	Allowed	Allowed	Allowed	Not allowed

Stub area	Allowed	Allowed (containing one default route)	Not allowed	Not allowed	Not allowed
Totally stub area	Allowed	Only one default route is allowed.	Not allowed	Not allowed	Not allowed
NSSA area	Allowed	Allowed (containing one default route)	Allowed	Not allowed	Allowed
Totally NSSA area	Allowed	Only one default route is allowed.	Allowed	Not allowed	Allowed

- i The ABR uses Type 3 LSAs to advertise a default route to the (totally) stub or NSSA area.
- i The ABR converts Type 7 LSAs in the totally NSSA area to Type 5 LSAs, and advertise Type 5 LSAs to the backbone area.
- If an area is appropriately configured as a (totally) stub area or an NSSA area, a large number of OE1, OE2, and OI routes will not be added to the routing table of a router in the area.

Area	Routes Available in the Routing Table of a Router Inside the Area
Non (totally) stub area and NSSA area	O: a route to a destination network in the local area OI: a route to a destination network in another area OE1 or OE2: a route or default route to a destination network segment outside the AS (via any ASBR in the AS)
Stub area	O: a route to a destination network in the local area OI: a route or a default route to a destination network in another area
Totally stub area	O: a route to a destination network in the local area OI: a default route
NSSA area	O: a route to a destination network in the local area OI: a route or a default route to a destination network in another area ON1 or ON2: a route or default route to a destination network segment outside the AS (via an ASBR in the local area)
Totally NSSA area	O: a route to a destination network in the local area OI: a default route

	ON1 or ON2: a route or default route to a destination network segment outside the AS (via an ASBR in the local area)
--	--

Route Redistribution

Route redistribution refers to the process of introducing routes of other routing protocols, routes of other OSPF processes, static routes, and direct routes that exist on the device to an OSPF process so that these routes can be advertised to neighbors using Type 5 and Type 7 LSAs. A default route cannot be introduced during route redistribution.

Route redistribution is often used for interworking between ASs. You can configure route redistribution on an ASBR to advertise routes outside an AS to the interior of the AS, or routes inside an AS to the exterior of the AS.

Default Route Introduction

By configuring a command on an ASBR, you can introduce a default route to an OSPF process so that the route can be advertised to neighbors using Type 5 and Type 7 LSAs.

Default route introduction is often used for interworking between ASs. One default route is used to replace all the routes outside an AS.

Route Summarization

Route summarization is a process of summarizing routing information with the same prefix into one route, and advertising the summarized route (replacing a large number of individual routes) to neighbors. Route summarization helps reduce the protocol interaction load and the size of the routing table.

By default, the ABR advertises inter-area routing information by using Type3 LSAs within a network segment, and advertises redistributed routing information by using Type 5 and Type 7 LSAs. If continuous network segments exist, it is recommended that you configure route summarization.

Route Filtering

OSPF supports route filtering to ensure security and facilitate control when the routing information is being learned, exchanged, or used.

Using configuration commands, you can configure route filtering for the following items:

- Interface: The interface is prevented from sending routing information (any LSAs) or exchanging routing information (any LSAs) with neighbors.
- Routing information outside an AS: Only the routing information that meets the filtering conditions can be redistributed to the OSPF process (Type 5 and Type 7 LSAs).
- LSAs received by a router: In the OSPF routing table, only the routes that are computed based on the LSAs meeting the filtering conditions can be advertised.

Route Cost

If redundancy links or devices exist on the network, multiple paths may exist from the local device to the destination network. OSPF selects the path with the minimum total cost to form an OSPF route. The total cost of a path is equal to the sum of the costs of individual links along the path. The total cost of a path can be minimized by modifying the costs of individual links along the path. In this way, OSPF selects this path to form a route.

Using configuration commands, you can modify the following link costs:

- Cost from an interface to a directly connected network segment and cost from the interface to a neighbor
- Cost from an ABR to the default network segment
- Cost from an ASBR to an external network segment and cost from the ASBR to the default network segment

i Both the cost and the metric indicate the cost and are not differentiated from each other.

OSPF Administrative Distance

The administrative distance (AD) evaluates reliability of a route, and the value is an integer ranging from 0 to 255. A smaller AD value indicates that the route is more trustworthy. If multiples exist to the same destination, the route preferentially selects a route with a smaller AD value. The route with a greater AD value becomes a floating route, that is, a standby route of the optimum route.

By default, the route coming from one source corresponds to an AD value. The AD value is a local concept. Modifying the AD value affects route selection only on the current router.

Route Source	Directly-connected network	Static route	EBGP Route	OSPF Route	IS-IS Route	RIP Route	IBGP Route	Unreachable Route
Default AD	0	1	20	110	115	120	200	255

Related Configuration

Stub/NSSA Area

By default, no stub or nssa area is configured.

Run the **area stub** command to configure a specified area as a stub area.

Run the **area nssa** command to configure a specified area as a nssa area.

- i** A backbone area cannot be configured as a stub or nssa area.
- i** A transit area (with virtual links going through) cannot be configured as a stub or nssa area.
- i** An area containing an ASBR cannot be configured as a stub area.

Route Redistribution and Default Route Introduction

By default, routes are not redistributed and the default route is not introduced.

Run the **redistribute** command to configure route redistribution.

Run the **default-information originate** command to introduce a default route.

After configuring route redistribution and default route introduction, the router automatically becomes an ASBR.

Route Summarization

By default, routes are not summarized. If route summarization is configured, a discard route will be automatically added.

Run the **area range** command to summarize routes (Type 3 LSA) distributed between areas on the ABR.

Run the **summary-prefix** command to summarize redistributed routes (Type 5 and Type 7 LSAs) on the ASBR.

Route Filtering

By default, routes are not filtered.

Run the **passive-interface** command to configure a passive interface. Routing information (any LSAs) cannot be exchanged on a passive interface.

Use the **route-map** parameter in the **redistribute** command, or use the **distribute-list out** command to filter the external routing information of the AS on the ASBR. Only the routing information that meets the filtering conditions can be redistributed to the OSPF process (Type 5 LSAs).

Run the **distribute-list in** command to filter LSAs received by the router. In the OSPF routing table, only the routes that are computed based on the LSAs meeting the filtering conditions can be advertised.

Route Cost

- Cost from the interface to the directly-connected network segment (cost on the interface)
The default value is the auto cost. Auto cost = Reference bandwidth/Interface bandwidth
Run the **auto-cost reference-bandwidth** command to set the reference bandwidth of the auto cost.
The default value is 100 Mbps.
Run the **ipv6 ospf cost** command to manually set the cost of the interface. The configuration priority of this item is higher than that of the auto cost.
- Cost from the interface to a specified neighbor (that is, cost from the local device to a specified neighbor)
The default value is the auto cost.
Use the **cost** parameter in the **ipv6 ospf neighbor** command to modify the cost from the interface

to a specified neighbor. The configuration priority of this item is higher than that of the cost of the interface.

This configuration item is applicable only to P2MP-type interfaces.

- Cost from the ABR to the default network segment (that is, the cost of the default route that is automatically advertised by the ABR to the stub area)
The default value is 1.
Run the **area default-cost** command to modify the cost of the default route that the ABR automatically advertise to the stub or nssa areas.
- Cost from the ASBR to an external network segment (that is, the metric of an external route)
By default, the metric of a redistributed BGP route is 1, the metric of other types of redistributed routes is 20, and the route type is Type 2 External.
Run the **default-metric** command to modify the default metric of the external route.
Use the **metric**, **metric-type**, and **route-map** parameters in the **redistribute** command to modify the metric and route type of the external route.
- Cost from the ASBR to the default network segment (that is, the metric of the default route that is manually introduced)
By default, the metric is 1, and the route type is Type 2 External.
Use the **metric**, **metric-type**, and **route-map** parameters in the **default-information originate** command to modify the metric and route type of the default route that is manually introduced.
- Use the **metric** and **metric-type** parameters of **default-information originate** in the **area nssa** command to modify the metric and type of the default route that is manually introduced to the NSSA area.

OSPF Administrative Distance

By default, the OSPF AD is 110.

Run the **distance** command to set the AD of an OSPF route.

5.3.3 Enhanced Security and Reliability

Use functions such as authentication and BFD correlation to enhance security, stability, and reliability of OSPF.

Working Principle

Authentication

OSPFv3 uses the authentication mechanism, that is, IP authentication header (AH) and IP Encapsulating Security Payload (ESP), provided by IPv6 to prevent unauthorized routers that access the network and hosts that forge OSPF packets to participate in OSPF routing. OSPF packets received on the OSPF

interface (or at both ends of a virtual link) are authenticated. If authentication fails, the packets are discarded and the adjacency cannot be set up.

Enabling authentication can avoid learning unauthenticated or invalid routes, thus preventing advertising valid routes to unauthenticated devices. In the broadcast-type network, authentication also prevents unauthenticated devices from becoming designated devices, ensuring stability of the routing system and protecting the routing system against intrusions.

MTU Verification

On receiving a DD packet, OSPF checks whether the MTU of the neighbor interface is the same as the MTU of the local interface. If the MTU of the interface specified in the received DD packet is greater than the MTU of the interface that receives the packet, the adjacency cannot be set up. Disabling MTU verification can avoid this problem.

Two-Way Maintenance

OSPF routers periodically send Hello packets to each other to maintain the adjacency. On a large network, a lot of packets may be sent or received, occupying too much CPU and memory. As a result, some packets are delayed or discarded. If the processing time of Hello packets exceeds the dead interval, the adjacency will be destroyed.

If the two-way maintenance function is enabled, in addition to the Hello packets, the DD, LSU, LSR, and LSAck packets can also be used to maintain the bidirectional communication between neighbors, which makes the adjacency more stable.

Concurrent neighbor Interaction Restriction

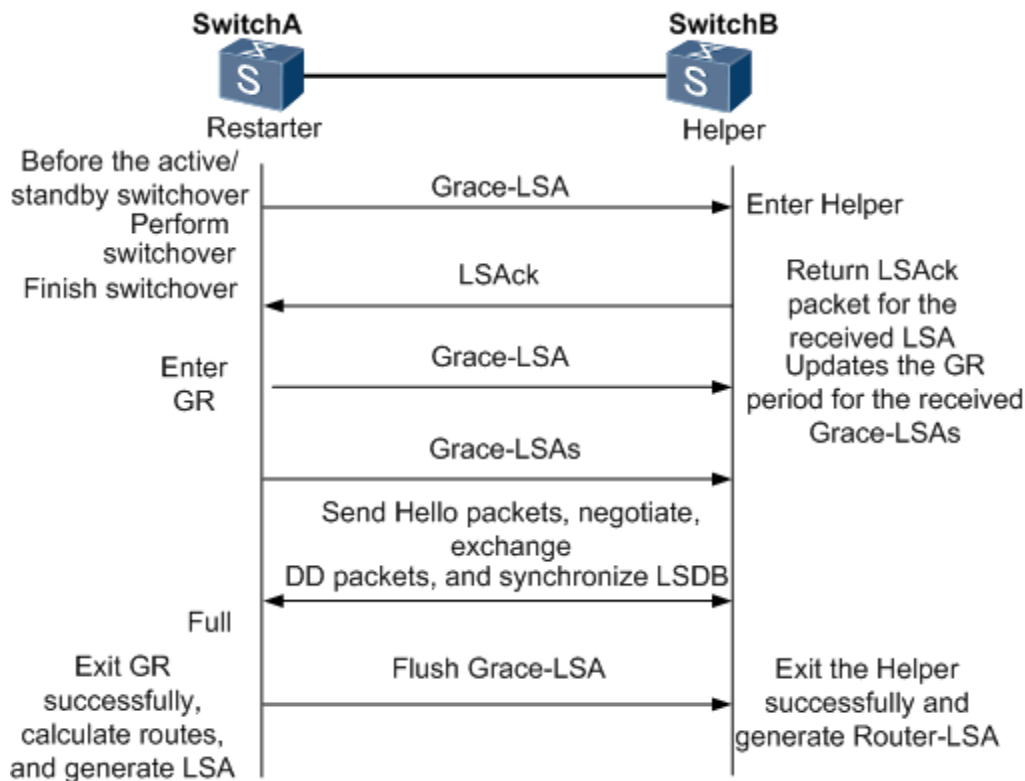
When a router simultaneously exchanges data with multiple neighbors, its performance may be affected. If the maximum number of neighbors that concurrently initiate or accept interaction with the OSPF process, the router can interact with neighbors by batches, which ensures data forwarding and other key services.

GR

The control and forwarding separated technology is widely used among routers. On a relatively stable network topology, when a GR-enabled router is restarted on the control plane, data forwarding can continue on the forwarding plane. In addition, actions (such as adjacency re-forming and route computation) performed on the control plane do not affect functions of the forwarding plane. In this way, service interruption caused by route flapping can be avoided, thus enhancing reliability of the entire network.

Currently, the GR function is used only during active/standby switchover and system upgrade.

Figure 5-7 Normal OSPF GR Process



- The GR process requires collaboration between the restarter and the helper. The restarter is the router where GR occurs. The helper is a neighbor of the restarter.
- When entering or exiting the GR process, the restarter sends a Grace-LSA to the neighbor, notifying the neighbor to enter or exit the helper state.
- When the adjacency between the restarter and the helper reaches the Full state, the router can exit the GR process successfully.

Fast Hello and BFD Correlation

After a link fault occurs, it takes a period of time (about 40s) before OSPF can sense the death of the neighbor. Then, OSPF advertises the information and re-computes the SPT. During this period, traffic is interrupted.

- After the fast Hello function is enabled (that is, the neighbor dead interval is set to 1s), OSPF can sense the death of a neighbor within 1s once a link is faulty. This greatly accelerates route convergence and prevents traffic interruption.
- BFD is used to test connectivity between devices. A link fault can be detected in as short as 150 ms. After OSPF is correlated with BFD, OSPF can sense the death of a neighbor in as short as 150 ms once a link is faulty. This greatly accelerates route convergence and prevents traffic interruption.

Related Configuration

OSPF Packet Authentication

By default, authentication is disabled.

- Run the **area authentication** command to enable authentication in the entire area so that the authentication function takes effect on all interfaces in this area. If authentication is enabled in area 0, the function also takes effect on the virtual link.
- Run the **area encryption** command to enable encryption and authentication in the entire area so that the encryption and authentication functions take effect on all interfaces in this area. If encryption and authentication are enabled in area 0, the functions also take effect on the virtual link.
- Run the **ipv6 ospf authentication** command to enable authentication on an interface. This configuration takes precedence over the area-based configuration.
- Run the **ipv6 ospf encryption** command to enable encryption and authentication on an interface. This configuration takes precedence over the area-based configuration.
- Use the **authentication** parameter in the **area virtual-link** command to enable authentication at both ends of a virtual link. This configuration takes precedence over the area-based configuration.
- Use the **encryption** parameter in the **area virtual-link** command to enable encryption and authentication at both ends of a virtual link. This configuration takes precedence over the area-based configuration.

MTU Verification

By default, MTU verification is disabled.

Run the **ipv6 ospf mtu-ignore** command to disable MTU verification on an interface.

Two-Way Maintenance

By default, bidirectional maintenance is enabled.

Run the **two-way-maintain** command to enable two-way maintenance.

Concurrent neighbor Interaction Restriction

Run the **max-concurrent-dd** command to modify the maximum number of neighbors that are concurrently interacting with the current OSPF process. The default value is 5.

Run the **ipv6 router ospf max-concurrent-dd** command to modify the maximum number of neighbors that are concurrently interacting with all OSPF processes on the router. The default value is 10.

GR

By default, the restarter function is disabled, and the helper function is enabled.

Run the **graceful-restart** command to configure the restarter function.

Run the **graceful-restart helper** command to configure the helper function.

Fast Hello

By default, the neighbor dead interval on the interface is 40s.

Run the **ipv6 ospf dead-interval minimal hello-multiplier** command to enable the Fast Hello function on an interface, that is, the neighbor dead interval is 1s.

Correlating OSPF with BFD

By default, OSPF is not correlated with BFD.

Run the **bfd interval min_rx multiplier** command to set the BFD parameters.

Run the **bfd all-interfaces** command to correlate OSPF with BFD on all interfaces.

Run the **ipv6 ospf bfd** command to correlate OSPF with BFD on the current interface.

5.3.4 Network Management Functions

Use functions such as the MIB and Syslog to facilitate OSPF management.

Working Principle

MIB

MIB is the device status information set maintained by a device. You can use the management program to view and set the MIB node.

Multiple OSPF processes can be simultaneously started on a router, but the OSPF MIB can be bound with only one OSPF process.

Trap

A trap message is a notification generated when the system detects a fault. This message contains the related fault information.

If the trap function is enabled, the router can proactively send the trap messages to the network management device.

Syslog

The Syslog records the operations (such as command configuration) performed by users on routers and specific events (such as network connection failures).

If the syslog is allowed to record the adjacency changes, the network administrator can view the logs to learn the entire process that the OSPF adjacency is set up and maintained.

Related Configuration

MIB

By default, the MIB is bound with the OSPF process with the smallest process ID.

Run the **enable mib-binding** command to bind the MIB with the current OSPF process.

Trap

By default, all traps functions are disabled, and the device is not allowed to send OSPF traps.

Run the **snmp-server enable traps ospf** command to allow the device to send OSPF traps.




Run the **enable traps** command to enable a specified trap function for an OSPF process.

Syslog

By default, the Syslog is allowed to record the adjacency changes.

Run the **log-adj-changes** command to allow the Syslog to record the adjacency changes.


5.4 Configuration

Configuration	Description and Command
Configuring OSPF Basic Functions	 (Mandatory) It is used to build an OSPF routing domain.
	ipv6routerospf Creates an OSPF process.
	router-id Configures a router ID.
	ipv6 ospfarea Enables OSPF on an interface and specifies an area ID.
	area virtual-link Creates a virtual link.
Setting the Network Type	 (Optional) The configurations are mandatory if the physical network is the X.25, frame relay, or ATM network.
	ipv6 ospf network Defines the network type.
	ipv6 ospf neighbor Specifies a neighbor.
	ipv6 ospf priority Configures the DR priority.
	 (Optional) The configurations are recommended if the OSPF routing domain is connected with an external network.

Configuring Route Redistribution and Default Route	redistribute	Configures route redistribution.
	default-information originate	Introduces a default route.
Configuring the Stub/NSSA Area	⚠ (Optional) It is used to reduce interaction of routing information and the size of routing table, and enhance stability of routes.	
	Area stub	Configures a stub area.
	Area nssa	Configures a nssa area.
Configuring Route Summarization	⚠ (Optional) It is used to reduce interaction of routing information and the size of routing table, and enhance stability of routes.	
	arearange	Summarizes routes that are advertised between areas.
	summary-prefix	Summarizes routes that are introduced through redistribution.
Configuring Route Filtering	⚠ (Optional) It is used to manually control interaction of routing information and filter available OSPF routes.	
	passive-interface	Configures a passive interface.
	distribute-list out	Filters routes that are introduced through redistribution.
	distribute-listin	Filters received LSAs.
Modifying the Route Cost and AD	⚠ (Optional) It is used to manually control the shortest route computed by OSPF and determine whether to select an OSPF route preferentially.	
	auto-costreference-bandwidth	Modifies the reference bandwidth of the auto cost.
	ipv6 ospf cost	Modifies the cost in the outbound direction of an interface.
	areadefault-cost	Modifies the cost of the default route in a stub/nssa area.

	default-metric	Modifies the default metric of a redistributed route.
	distance	Modifies the OSPF AD.
Enabling Authentication	⚠️ (Optional) It is used to prevent routers that illegally access the network and hosts that forge OSPF packets from participating in the OSPF protocol process.	
	areaauthentication	Enables authentication and sets the authentication mode in an area.
	areaencryption	Enables encryption and authentication and sets the authentication mode in an area.
	ipv6 ospf authentication	Enables authentication and sets the authentication mode on an interface.
	ipv6 ospf encryption	Enables encryption and authentication and sets the authentication mode on an interface.
Modifying the Maximum Number of Concurrent Neighbors	⚠️ (Optional) It is used to prevent the problem of performance deterioration caused by over-consumption of the CPU.	
	max-concurrent-dd	Modifies the maximum number of concurrent neighbors on the current OSPF process.
	ipv6 router ospf max-concurrent-dd	Modifies the maximum number of concurrent neighbors on all OSPF processes.
Disabling MTU Verification	⚠️ (Optional) It is used to prevent the problem that the adjacency cannot be set up due to MTU inconsistency on the neighbor interface.	
	ipv6 ospf mtu-ignore	Disables MTU verification on an interface.

Enabling Two-Way Maintenance	⚠ (Optional) It is used to prevent termination of the adjacency due to the delay or loss of Hello packets.	
	two-way-maintain	Enables two-way maintenance.
Enabling GR	⚠ (Optional) It is used to retain OSPF routing forwarding during restart or active/standby switchover of the OSPF processes to prevent traffic interruption.	
	graceful-restart	Enables the restarter function.
	graceful-restart helper	Enables the helper function.
Enabling Fast Hello	⚠ (Optional) It is used to quickly discover the death of a neighbor to prevent traffic interruption when a link is faulty.	
	ipv6 ospf dead-intervalminimal hello-multiplier	Enabling the Fast Hello function on an interface.
Enabling GR	⚠ (Optional) It is used to retain OSPF routing forwarding during restart or active/standby switchover of the OSPF processes to prevent traffic interruption.	
	graceful-restart	Enables the restarter function.
	graceful-restart helper	Enables the helper function.
Correlating OSPF with BFD	⚠ (Optional) It is used to quickly discover the death of a neighbor to prevent traffic interruption when a link is faulty.	
	bfd all-interfaces	Correlates OSPF with BFD on all interfaces.
	ipv6 ospf bfd	Correlates OSPF with BFD on the current interface.
Configuring Network	⚠ (Optional) The configurations enable users to use the SNMP network management software to manage OSPF.	

Management Functions	enable mib-binding	Bind MIB to the OSPF process.
	enable traps	Enables the trap function of the OSPF process.
	log-adj-changes	Allows the syslogs to record the changes in adjacency status.
Modifying Protocol Control Parameters	 (Optional) You are advised not to modify protocol control parameters unless necessary.	
	ipv6 ospf hello-interval	Modifies the Hello interval on an interface.
	ipv6 ospf dead-interval	Modifies the neighbor death interval on an interface.
	ipv6 ospf transmit-delay	Modifies the LSU packet transmission delay on an interface.
	ipv6 ospf retransmit-interval	Modifies the LSU packet retransmission interval on an interface.
	timers throttle lsa all	Modifies parameters of the exponential backoff algorithm that generates LSAs.
	timers pacing lsa-group	Modifies the LSA group update interval.
	timers pacing lsa-transmit	Modifies the LS-UPD packet sending interval.
	timers lsa arrival	Modifies the delay after which the same LSA is received.
	timers throttlespf	Modifies the SPT computation timer.

	timers throttle route inter-area	Modifies the inter-area route computation delay.
	timers throttle route ase	Modifies the inter-area route computation delay.

5.4.1 Configuring OSPF Basic Functions

Configuration

Effect

- Set up an OSPF routing domain on the network to provide IPv6 unicast routing service for users on the network.

Notes

- Ensure that the IPv6 routing function is enabled, that is, **ipv6 routing** is not disabled; otherwise, OSPF cannot be enabled.
- IPv6 must be enabled on the interface.
- It is strongly recommended that you manually configure the router ID.

Configuration

Steps

Creating an OSPF Process

- Mandatory.
- The configuration is mandatory for every router.

Configuring a Router ID

- (Optional) It is strongly recommended that you manually configure the router ID.
- If the router ID is not configured, OSPF selects an interface IP address. If the IP address is not configured for any interface, or the configured IP addresses have been used by other OSPF instances, you must manually configure the router ID.

Enabling OSPF on an Interface and Specifying an Area ID

- Mandatory.
- The configuration is mandatory for every router.

Verification

- Run the **show ipv6 route ospf** command to verify that the entries of the OSPF routing table are correctly loaded.
- Run the **ping** command to verify that the IPv6 unicast service is correctly configured.

Related Commands

Creating an OSPF Process

Command	ipv6 router ospf<i>process-id</i> [vrf<i>vrf-name</i>]
Parameter Description	<i>process-id</i> : Indicates the OSPFv3 process ID. If the process ID is not specified, process 1 is enabled. <i>vrf-name</i> : Specifies the VPN routing and forwarding (VRF) to which the OSPFv3 process belongs.
Command Mode	Global configuration mode
Usage Guide	After enabling the OSPFv3 process, the device enters the routing process configuration mode.

Configuring a Router ID

Command	router-id<i>router-id</i>
Parameter Description	<i>router-id</i> : Indicates the ID of the device, which is expressed in the IPv4 address.
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>Every device where OSPFv3 run must be identified by using a router ID. You can configure any IPv4 address as the router ID of the device, and ensure that the router ID is unique in an AS. If multiple OSPFv3 processes run on the same device, the router ID of each process must also be unique.</p> <p>After the router ID changes, OSPF performs a lot of internal processing. Therefore, you are advised not to change the router ID unless necessary. When an attempt is made to modify the router ID, a prompt is displayed, requesting you to confirm the modification. After the OSPFv3 process is enabled, you are advised to specify the router ID before configuring other parameters of the process.</p>

Enabling OSPF on an Interface and Specifying an Area ID

Command	ipv6 ospf<i>process-id</i> area<i>area-id</i> [instance<i>instance-id</i>]
----------------	---

Parameter Description	<p>process-id: Indicates the ID of an OSPFv3 process. The value ranges from 1 to 65,535.</p> <p>Areaarea-id: Indicates the ID of the OSPFv3 area in which the interface participates. It can be an integer or an IPv4 prefix.</p> <p>Instanceinstance-id: Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>Run this command in interface configuration mode to enable the interface to participate in OSPFv3, and then run the ipv6 router ospf command to configure the OSPFv3 process. After the OSPFv3 process is configured, the interface will automatically participate in the related process.</p> <p>Run the no ipv6 ospf area command so that the specified interface no longer participates in the OSPFv3 routing process.</p> <p>Run the no ipv6 router ospf command so that all interfaces no longer participate in the OSPFv3 routing process.</p> <p>The adjacency can be set up only between devices with the same <i>instance-id</i>.</p> <p>After this command is configured, all prefix information on the interface will participate in the OSPFv3 process.</p>


Creating a Virtual Link

Command	<pre>area <i>area-id</i> virtual-link <i>router-id</i> [hello-interval <i>seconds</i>] [dead-interval <i>seconds</i>] [retransmit-interval <i>seconds</i>] [transmit-delay <i>seconds</i>] [instance <i>instance-id</i>] [authentication ipsec sp <i>spi</i> [md5 sha1] [0 7] <i>key</i>] [encryption ipsec sp <i>spi</i> esp null [md5 sha1] [0 7] <i>key</i>]</pre>
Parameter Description	<p>area-id: Indicates the ID of the area where the virtual link is located. It can be an integer or an IPv4 prefix.</p> <p>router-id: Indicates the router ID of the neighbor connected to the virtual link.</p> <p>dead-interval seconds: Indicates the time that the local interface of the virtual link detects the failure of the neighbor. The unit is second. The value ranges from 1 to 65,535.</p> <p>hello-interval seconds: Indicates the time that the Hello packet is sent on the local interface of the virtual link. The unit is second. The value ranges from 1 to 65,535.</p> <p>retransmit-interval seconds: Indicates the interval at which the LSA is retransmitted on the local interface of the virtual link. The unit is second. The value ranges from 1 to 65,535.</p>

	<p>transmit-delay seconds: Indicates the delay after which the LSA is sent on the local interface of the virtual link. The unit is second. The value ranges from 1 to 65,535.</p> <p>instanceinstance-id: Indicates the ID of the instance corresponding to the virtual link. The value ranges from 0 to 255. A virtual link cannot be set up between devices with different instance IDs.</p> <p>spi: Indicates the security parameter index (SPI). The value ranges from 256 to 4,294,967,295.</p> <p>md5: Enables message digit 5 (MD5) authentication.</p> <p>sha1: Enables Secure Hash Algorithm 1 (SHA1) authentication.</p> <p>0: Indicates that the key is displayed in plain text.</p> <p>7: Indicates that the key is displayed in cipher text.</p> <p>key: Indicates the authentication key.</p> <p>null: Indicates that no encryption mode is used.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>In an OSPFv3 AS, all areas must be connected to the backbone area to properly learn the routing information of the entire OSPFv3 AS. If an area cannot be directly connected to the backbone area, the virtual link can be used to connect this area to the backbone area.</p> <p>The area where the virtual link is located cannot be a stub or nssa area.</p> <p>At both ends of neighbors between which the virtual link is set up, settings of hello-interval, dead-interval, and instance must be consistent; otherwise, the adjacency cannot be set up properly.</p>

Configuration Example

Scenario
Figure 5-8



	<table border="1"> <tr> <td data-bbox="342 264 480 527">Remark S</td> <td data-bbox="480 264 1463 527"> <p>The interface IP addresses are as follows:</p> <p>A: GE 0/1 2001:1::1/64 GE 0/2 2001:2::1/64</p> <p>B: GE 0/1 2001:1::2/64 GE 0/2 2001:3::1/64</p> <p>C: GE 0/3 2001:2::2/64</p> <p>D: GE 0/3 2001:3::2/64</p> </td> </tr> </table>	Remark S	<p>The interface IP addresses are as follows:</p> <p>A: GE 0/1 2001:1::1/64 GE 0/2 2001:2::1/64</p> <p>B: GE 0/1 2001:1::2/64 GE 0/2 2001:3::1/64</p> <p>C: GE 0/3 2001:2::2/64</p> <p>D: GE 0/3 2001:3::2/64</p>
Remark S	<p>The interface IP addresses are as follows:</p> <p>A: GE 0/1 2001:1::1/64 GE 0/2 2001:2::1/64</p> <p>B: GE 0/1 2001:1::2/64 GE 0/2 2001:3::1/64</p> <p>C: GE 0/3 2001:2::2/64</p> <p>D: GE 0/3 2001:3::2/64</p>		
Configuration Steps	<ul style="list-style-type: none"> ▪ Configure the interface IP addresses on all routers. ▪ Enable the IPv4 unicast routing function on all routers. (This function is enabled by default.) ▪ Configure the OSPF instances and router IDs on all routers. ▪ Enable OSPF on the interfaces configured on all routers. 		
A	<pre>A#configure terminal A(config)#interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#ipv6 enable A(config-if-GigabitEthernet 0/1)#ipv6 address 2001:1::1/64 A(config-if-GigabitEthernet 0/1)#ipv6 ospf 1 area 0 A(config-if-GigabitEthernet 0/1)#exit A(config)#interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)#ipv6 enable A(config-if-GigabitEthernet 0/2)#ipv6 address 2001:2::1/64 A(config-if-GigabitEthernet 0/2)#ipv6 ospf 1 area 1 A(config-if-GigabitEthernet 0/2)#exit A(config)#ipv6 router ospf 1 A(config-router)#router-id1.1.1.1</pre>		
B	<pre>B#configure terminal B(config)#interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)#ipv6 enable B(config-if-GigabitEthernet 0/1)#ipv6 address 2001:1::2/64 B(config-if-GigabitEthernet 0/1)#ipv6 ospf 1 area 0 B(config-if-GigabitEthernet 0/1)#exit B(config)#interface GigabitEthernet 0/2</pre>		

	<pre> B(config-if-GigabitEthernet 0/2)#ipv6 enable B(config-if-GigabitEthernet 0/2)#ipv6 address 2001:3::1/64 B(config-if-GigabitEthernet 0/2)#ipv6 ospf 1 area 2 B(config-if-GigabitEthernet 0/2)#exit B(config)#ipv6 router ospf 1 B(config-router)#router-id2.2.2.2 </pre>
C	<pre> C#configure terminal C(config)#interface GigabitEthernet 0/3 C(config-if-GigabitEthernet 0/3)#ipv6 enable C(config-if-GigabitEthernet 0/3)#ipv6 address 2001:2::2/64 C(config-if-GigabitEthernet 0/3)#ipv6 ospf 1 area 1 C(config-if-GigabitEthernet 0/3)#exit C(config)#ipv6 router ospf 1 C(config-router)#router-id3.3.3.3 </pre>
D	<pre> D#configure terminal D(config)#interface GigabitEthernet 0/3 D(config-if-GigabitEthernet 0/3)#ipv6 enable D(config-if-GigabitEthernet 0/3)#ipv6 address 2001:4::2/64 D(config-if-GigabitEthernet 0/3)#ipv6 ospf 1 area 2 D(config-if-GigabitEthernet 0/3)#exit D(config)#ipv6 router ospf 1 D(config-router)#router-id4.4.4.4 </pre>
Verification	<ul style="list-style-type: none"> ▪ Verify that the OSPF neighbors are correct on all routers. ▪ Verify that the routing table is correctly loaded on all routers. ▪ Verify that 2001:2::2/64 can be pinged successfully on Router D.

<p>A</p>	<pre> A#show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Instance ID Interface 2.2.2.2 1 Full/BDR 00:00:30 0 GigabitEthernet 0/1 3.3.3.31 Full/BDR 00:00:35 0 GigabitEthernet 0/2 A#show ipv6 route ospf IPv6 routing table name - Default - 0 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area O IA2001:3::/64 [110/20] via FE80::2D0:F8FF:FE22:4524, GigabitEthernet 0/1 </pre>
<p>B</p>	<pre> B# show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Instance ID Interface 1.1.1.11 Full/DR 00:00:30 0 GigabitEthernet 0/1 4.4.4.41 Full/BDR 00:00:35 0 GigabitEthernet 0/2 B#show ipv6 route ospf IPv6 routing table name - Default - 0 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area </pre>

```
O IA2001:2::/64 [110/20] via FE80::2D0:F8FF:FE22:4536, GigabitEthernet 0/1
```

C

```
C# show ipv6 ospf neighbor
```

```
OSPFv3 Process (1), 1 Neighbors, 1 is Full:
```

Neighbor ID	Pri	State	Dead Time	Instance ID	Interface
1.1.1.11	Full/DR	00:00:30	0		GigabitEthernet 0/3

```
C#show ipv6 route ospf
```

```
IPv6 routing table name - Default - 0 entries
```

```
Codes: C - Connected, L - Local, S - Static
```

```
R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
IA - Inter area
```

```
O IA2001:1::/64 [110/2] via FE80::2D0:F8FF:FE22:4537, GigabitEthernet 0/3
```

```
O IA2001:3::/64 [110/3] via FE80::2D0:F8FF:FE22:4537, GigabitEthernet 0/3
```

D

```
D# show ipv6 ospf neighbor
```

```
OSPFv3 Process (1), 1 Neighbors, 1 is Full:
```

Neighbor ID	Pri	State	Dead Time	Instance ID	Interface
2.2.2.2	1 Full/DR	00:00:30	0		GigabitEthernet 0/3

```
D#show ipv6 route ospf
```

```
IPv6 routing table name - Default - 0 entries
```

```
Codes: C - Connected, L - Local, S - Static
```

```
R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

IA - Inter area

```
O IA2001:1::/64 [110/2] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/3
```

```
O IA2001:2::/64 [110/3] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/3
```

```
D#
```

```
D#ping 2001:2::2
```

```
Sending 5, 100-byte ICMP Echoes to 2001:2::2, timeout is 2 seconds:
```

```
< press Ctrl+C to break >
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/9/14 ms.
```

Common

Errors

- IPv6 is disabled on the interface.
- OSPF cannot be enabled because the IPv6 unicast routing function is disabled.
- The area IDs enabled on adjacent interfaces are inconsistent.
- The same router ID is configured on multiple routers, resulting in a router ID conflict.

5.4.2 Setting the Network Type

Configuration

Effect

- If the physical network is X.25, Frame Relay, or ATM, OSPF can also run to provide the IPv6 unicast routing service.

Notes

- The OSPF basic functions must be configured.
- The broadcast network sends multicast OSPF packets, automatically discovers neighbors, and elects a DR and a BDR.
- The P2P network sends multicast OSPF packets and automatically discovers neighbors.
- The NBMA network sends unicast OSPF packets. Neighbors must be manually specified, and a DR and a BDR must be elected.
- The P2MP network (without carrying the **non-broadcast** parameter) sends multicast OSPF packets. Neighbors are automatically discovered.
- The P2MP network (carrying the **non-broadcast** parameter) sends unicast OSPF packets. Neighbors must be manually specified.

Configuration Steps

Configuring the Interface Network Type

- Optional.
- Perform this configuration on routers at both ends of the link.

Configuring a Neighbor

- (Optional)If the interface network type is set to NBMA or P2MP (carrying the **non-broadcast** parameter), neighbors must be configured.
- Neighbors are configured on routers at both ends of the NBMA or P2MP (carrying the **non-broadcast** parameter) network.

Configuring the Interface Priority

- (Optional)You must configure the interface priority if a router must be specified as a DR, or a router cannot be specified as a DR.
- Configure the interface priority on a router that must be specified as a DR, or cannot be specified as a DR.

Verification

- Run the **show ipv6 ospf interface** command to verify that the network type of each interface is correct.

Related Commands

Configuring the Interface Network Type

Command	<code>ipv6 ospf network {broadcast non-broadcast point-to-point point-to-multipoint[non-broadcast]}[instanceinstance-id]</code>
Parameter Description	<p>broadcast: Indicates the broadcast network type.</p> <p>non-broadcast: Indicates the non-broadcast network type.</p> <p>point-to-multipoint: Indicates the point-to-multipoint (P2MP) network type.</p> <p>point-to-multipoint non-broadcast: Indicates the P2MP non-broadcast network type.</p> <p>point-to-point: Indicates the point-to-point (P2P) network type.</p> <p>instanceinstance-id: Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.</p>

Command Mode	Interface configuration mode
Usage Guide	You can configure the network type of an interface based on the actual link type and topology.

Configuring a Neighbor

Command	ipv6 ospf neighbor <i>ipv6-address</i> { [<i>costcost</i>] [<i>poll-intervalseconds</i> <i>priorityvalue</i>] } [<i>instanceinstance-id</i>]
Parameter Description	<p><i>ip-address</i>: Indicates the link address of the neighbor interface.</p> <p>costcost: Indicates the cost required from the P2MP network to each neighbor. The cost is not defined by default. The cost configured on the interface is used. The value ranges from 1 to 65,535. Only a P2MP network supports this option.</p> <p>poll-interval seconds: Indicates the neighbor polling interval. The unit is second. The value ranges from 1 to 2,147,483,647. Only the non-broadcast (NBMA) network supports this option.</p> <p>priority value: Indicates the priority value of the non-broadcast network neighbor. The value ranges from 0 to 255. Only the non-broadcast network (NBMA) supports this option.</p> <p>instanceinstance-id: Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.</p>
Command Mode	Interface configuration mode
Usage Guide	You can configure neighbor parameters based on the actual network type.

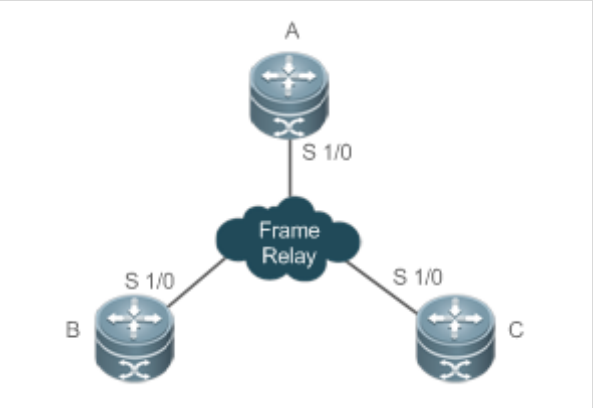
Configuring the Interface Priority

Command	ipv6 ospf priority <i>number-value</i> [<i>instanceinstance-id</i>]
Parameter Description	<p><i>number-value</i>: Indicates the priority of the interface. The value ranges from 0 to 255.</p> <p>instanceinstance-id: Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.</p>
Command Mode	Interface configuration mode

Usage Guide	<p>On a broadcast network, a DR or BDR must be elected. During the DR/BDR election, the device with a higher priority will be preferentially elected as a DR or BDR. If the priority is the same, the device with a larger router ID will be preferentially elected as a DR or BDR.</p> <p>A device with the priority 0 does not participate in the DR/BDR election.</p>
-------------	--

Configuration Example

Configuring the Interface Network Type

<p>Scenario</p> <p>Figure 5-9</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers. ▪ Configure the OSPF basic functions on all routers. ▪ Set the interface network type to P2MP on all routers.
<p>A</p>	<pre>A#configure terminal A(config)# interface Serial1/0 A(config-Serial1/0)# encapsulation frame-relay A(config-Serial1/0)# ipv6 ospf network point-to-multipoint</pre>
<p>B</p>	<pre>B#configure terminal B(config)# interface Serial1/0 B(config-Serial1/0)# encapsulation frame-relay B(config-Serial1/0)# ipv6 ospf network point-to-multipoint</pre>
<p>C</p>	<pre>C#configure terminal C(config)# interface Serial1/0</pre>

	<pre>C(config-Serial1/0)# encapsulation frame-relay C(config-Serial1/0)# ipv6 ospf network point-to-multipoint</pre>
Verification	<ul style="list-style-type: none"> Verify that the interface network type is P2MP.
A	<pre>A#show ipv6 ospf interface Serial1/0 Serial1/0 is up, line protocol is up Interface ID 2 IPv6 Prefixes fe80::2d0:f8ff:fe22:3346/64 (Link-Local Address) OSPFv3 Process (1), Area 0.0.0.1, Instance ID 0 Router ID 192.168.22.30, Network Type POINTOMULTIPOINT, Cost: 1 Transmit Delay is 1 sec, State Point-To-Point, Priority 1 Timer interval configured, Hello 30, Dead 120, Wait 40, Retransmit 10 Hello due in 00:00:06 Neighbor Count is 1, Adjacent neighbor count is 1 Hello received 40 sent 40, DD received 17 sent 9 LS-Req received 1 sent 3, LS-Upd received 6 sent 5 LS-Ack received 3 sent 4, Discarded 1</pre>

Common Errors

- The network types configured on interfaces at two ends are inconsistent, causing abnormal route learning.
- The network type is set to NBMA or P2MP (non-broadcast), but neighbors are not specified.

5.4.3 Configuring Route Redistribution and Default Route

Configuration Effect

- Introduce unicast routes for other AS domains to the OSPF domain to provide the unicast routing service to other AS domains for users in the OSPF domain.
- In the OSPF domain, inject a default route to another AS domain so that the unicast routing service to another AS domain can be provided for users in the OSPF domain.

Notes

- The OSPF basic functions must be configured.

Configuration Steps

Configuring External Route Redistribution

- (Optional) This configuration is mandatory if external routes of the OSPF domain should be introduced to the ASBR.
- Perform this configuration on an ASBR.

Generating a Default Route

- (Optional) Perform this configuration if the default route should be introduced to an ASBR so that other routers in the OSPF domain access other AS domains through this ASBR by default.
- Perform this configuration on an ASBR.

Verification

- On a router inside the OSPF domain, run the **show ipv6 route ospf** command to verify that the unicast routes to other AS domains are loaded.
- On a router inside the OSPF domain, run the **show ipv6 route ospf** command to verify that the default route to the ASBR is loaded.
- Run the **ping** command to verify that the IPv6 unicast service to other AS domains is correct.

Related Commands

Configuring Route Redistribution

Command	redistribute { bgp connected isis [<i>area-tag</i>] ospf <i>process-id</i> rip static }{ level-1 level-1-2 level-2 } match { internal external [1 2] nssa-external [1 2]} metric <i>metric-value</i> metric-type {1 2} route-map <i>route-map-name</i> / <i>tag-tag-value</i>]
Parameter Description	<p>bgp: Indicates redistribution from BGP.</p> <p>connected: Indicates redistribution from direct routes.</p> <p>isis [<i>area-tag</i>]: Indicates redistribution from IS-IS. area-tag specifies the IS-IS instance.</p> <p>ospf<i>process-id</i>: Indicates redistribution from OSPF. process-id specifies an OSPF instance. The value ranges from 1 to 65535. 1-65535</p> <p>rip: Indicates redistribution from RIP.</p> <p>static: Indicates redistribution from static routes.</p>

	<p>level-1 level-1-2 level-2: Used only when IS-IS routes are redistributed. Only the routes of the specified level are redistributed. By default, only level-2 IS-IS routes can be redistributed.</p> <p>match: Used only when OSPF routes are redistributed. Only the routes that match the specified criteria are redistributed. By default, all OSPF routes can be redistributed.</p> <p>metric<i>metric-value:</i> Indicates the metric of the OSPF external LSA. <i>metric-value</i> specifies the size of the metric. The value ranges from 0 to 16,777,214.</p> <p>metric-type {1 2}: Indicates the external route type, which can be E-1 or E-2.</p> <p>route-map<i>route-map-name:</i> Sets the redistribution filtering rules.</p> <p>tag<i>tag-value:</i> Specifies the tag value of the route that is redistributed into the OSPF routing domain. The value ranges from 0 to 4294967295.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>When the device supports multiple routing protocols, collaboration between protocols is very important. To run multiple routing protocols concurrently, the device must be able to redistribute routing information of a protocol to another protocol. This applies to all routing protocols.</p> <p>During redistribution of IS-IS routes, level-1, level-2, or level-1-2 can be configured to indicate that IS-IS routes of the specified level(s) will be redistributed. By default, IS-IS routes of level 2 are redistributed.</p> <p>During redistribution of OSPFv3 routes, match can be configured to indicate that OSPFv3 routes of the specified sub-type will be redistributed. By default, all types of OSPFv3 routes are redistributed.</p> <p>For the level parameter configured during redistribution of IS-IS routes and the match parameter configured during redistribution of OSPFv3 routes, the routes are matched against the route map only when the sub-type of the routes are correct.</p> <p>During configuration of route redistribution, the match rules configured in route map configuration mode are used based on the original information of routes. The priorities of tag, metric and metric-type in the route redistribution configuration are lower than the priority of the set rules configured in route map configuration mode.</p> <p>The set metric value of the associated routemap should fall into the range of 0 to 16,777,214. If the value exceeds this range, routes cannot be introduced.</p> <p>The configuration rules for the no form of the redistribute command are as follows:</p> <ol style="list-style-type: none"> 1. If some parameters are specified in the no form of the command, default values of these parameters will be restored.

	<p>2. If no parameter is specified in the no form of the command, the entire command will be deleted.</p> <p>For example, if redistribute isis 112 level-2 is configured, the no redistribute isis 112 level-2 command only restores the default value of level-2. As level-2 itself is the default value of the parameter, the configuration saved is still redistribute isis 112 level-2 after the preceding no form of the command is executed. To delete the entire command, you need to run the no redistribute isis 112 command.</p>
--	--

Introducing a Default Route

Command	default-information originate [always] [metric <i>metric</i>] [metric-type <i>type</i>] [route-map <i>map</i>]
Parameter Description	<p>always: Enables OSPF to generate a default route regardless of whether the local router has a default route.</p> <p>metric <i>metric</i>: Indicates the initial metric of the default route. The value ranges from 0 to 16, 777, 214. By default, the metric of the default route is 1.</p> <p>metric-type <i>type</i>: Indicates the type of the default route. OSPF external routes are classified into two types: Type 1: The metric varies with routers; Type 2: The metric is the same for all routers. Type 1 external routes are more trustworthy than Type 2 external routes.</p> <p>route-map <i>map-name</i>: Indicates the associated route-map name. By default, no route-map is associated.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>When the redistribute or default-information command is executed, the OSPFv3-enabled router automatically becomes an ASBR.</p> <p>The ASBR, however, does not automatically generate or advertise a default route to all routers in the OSPF routing domain. To have the ASBR generate a default route, configure the default-information originate command.</p> <p>If always is specified, the OSPFv3 process advertises an external default route to neighbors no matter whether a default route exists in the core routing table. This default route, however, is not displayed on the local router. To confirm whether the default route is generated, run the show ipv6 ospf database command to display the OSPFv3 link status database. On an OSPFv3 neighbor, you can run the show ipv6 route ospf command to see the default route.</p> <p>The metric of the external default route can only be defined in the default-information originate command, instead of the default-metric command.</p>

OSPFv3 has two types of external routes. The metric of the Type 1 external route changes, but the metric of the Type 2 external route is fixed. If two parallel paths to the same destination network have the same route metric, the priority of the Type 1 route is higher than that of the Type 2 route. Therefore, the **show ipv6 route ospf** command displays only the Type 1 route.

A router in a stub area cannot generate an external default route.

Configuration

Example

Configuring Route Redistribution

<p>Scenario Figure 5-10</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Enable IPv6 on interfaces of all routers. Configure the OSPF basic functions on all routers. Introduce an external static route to Router D.
<p>D</p>	<pre>D#configure terminal D(config)#ipv6 router ospf 1 D(config-router)# redistribute static</pre>
<p>Verification</p>	<ul style="list-style-type: none"> On Router D, run the show ipv6ospf database external brief command to verify that an LSA corresponding to an external route is generated. On Router C, run the show ipv6 route ospf command to verify that the external static route has been introduced.
<p>D</p>	<pre>D#show ipv6 ospf database external OSPFv3 Router with ID (4.4.4.4) (Process 1) AS-external-LSA</pre>

LS age: 7
 LS Type: AS-External-LSA
 Link State ID: 0.0.0.6
 Advertising Router: 4.4.4.4
 LS Seq Number: 0x80000001
 Checksum: 0x9C1F
 Length: 36
 Metric Type: 2 (Larger than any link state path)
 Metric: 20
 Prefix: 2001:10:10::/64
 Prefix Options: 0 (-|-|-|-)

C

C#show ipv6 route ospf

IPv6 routing table name - Default - 0 entries

Codes: C - Connected, L - Local, S - Static

R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

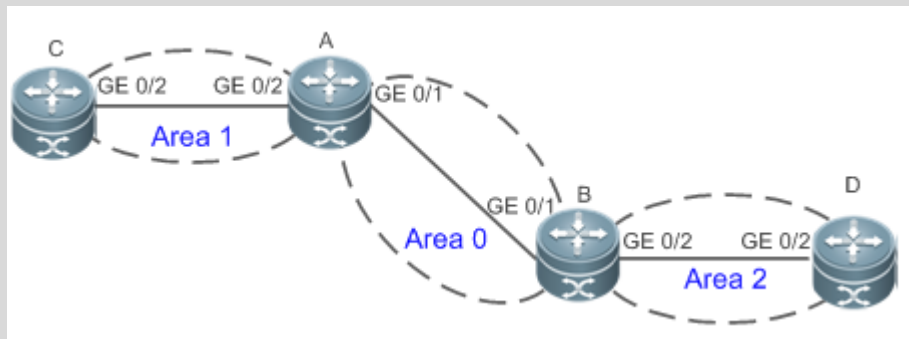
SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

IA - Inter area

O E2 2001:10:10::/64 [110/20] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/2

Configuring the Default Route

Scenario
Figure 5-11



Configuration
Steps

- Enable IPv6 on interfaces of all routers.

	<ul style="list-style-type: none"> Configure the OSPF basic functions on all routers. Configure the default route on Router D.
D	<pre>D#configure terminal D(config)#ipv6 router ospf 1 D(config-router)#default-information originate always</pre>
Verification	<ul style="list-style-type: none"> On Router D, run the show ipv6ospf database external brief command to verify that an LSA corresponding to the default route is generated. On Router C, run the show ipv6 route ospf command to verify that the OSPF default route exists.
D	<pre>D#show ipv6 ospf database external OSPFv3 Router with ID (4.4.4.4) (Process 1) AS-external-LSA LS age: 3 LS Type: AS-External-LSA Link State ID: 0.0.0.7 Advertising Router: 4.4.4.4 LS Seq Number: 0x80000001 Checksum: 0x1839 Length: 32 Metric Type: 2 (Larger than any link state path) Metric: 1 Prefix: ::/0 Prefix Options: 0 (- - -) External Route Tag: 1</pre>
C	<pre>C#show ipv6route ospf IPv6 routing table name - Default - 0 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2</pre>

SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

IA - Inter area

O E2::/0 [110/20] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/2

Common Errors

- A route loop is formed because the **default-information originate always** command is configured on multiple routers.
- Routes cannot be introduced because route redistribution is configured on a router in the stub area.

5.4.4 Configuring the Stub Area and NSSA Area

Configuration Effect

- Configure an area located on the stub as a stub area to reduce interaction of routing information and the size of routing table, and enhance stability of routes.

Notes

- The OSPF basic functions must be configured.
- A backbone or transit area cannot be configured as a stub or an NSSA area.
- A router in the stub area cannot introduce external routes, but a router in the NSSA area can introduce external routes.

Configuration Steps

Configuring a Stub Area

- (Optional) Perform this configuration if you wish to reduce the size of the routing table on routers in the area.
- Perform this configuration on all routers in the same area.

Configuring an NSSA Area

- (Optional) Perform this configuration if you wish to reduce the size of the routing table on routers in the area and introduce OSPF external routes to the area.
- The area must be configured as an NSSA area on all routers in this area.

Verification

Verifying the Stub Area

- On a router in the stub area, run the **show ipv6 route** command to verify that the router is not loaded with any external routes.

Verifying the NSSA Area

- On a router in the NSSA area, run the **show ipv6 ospf database** command to verify that the introduced external route generates Type 7 LSAs.
- On a router in the backbone area, run the **show ipv6 route** command to verify that the router is loaded with external routes introduced from the NSSA area.

Related Commands

Configuring a Stub Area

Command	area <i>area-id</i> stub [no-summary]
Parameter Description	<i>area-id</i> : Indicates the ID of the stub area. The value can be an integer or an IPv4 prefix. no-summary : This option is valid only on the ABR in a stub area. If this option is specified, the ABR only advertises one Type 3 LSA indicating the default route to the stub area, and does not advertise other Type 3 LSAs.
Command Mode	OSPF routing process configuration mode
Usage Guide	An area located on the stub of a network can be configured as a stub area. You must run the area stub command on all routers in a stub area. Devices in a stub area cannot learn the external routes (Type 5 LSAs) of the AS. In practice, external routes take up a large proportion of the link status database. Therefore, devices in a stub area can learn only a small amount of routing information, which reduces the amount of system resources required to run the OSPFv3 protocol. By default, an ABR in a stub area will generate a Type 3 LSA indicating the default fault, and advertise the LSA to the stub area. In this way, devices in the stub area can access devices outside the AS. To configure a totally stub area, add the no-summary keyword when running the area stub command on the ABR.

Configuring an NSSA Area

Command	area <i>area-id</i> nssa [no-redistribution] [default-information-originate [<i>metric-value</i>] [<i>metric-type type</i>]] [no-summary] [translator [<i>stability-interval seconds</i> always]] [translate type7 suppress-fa]
----------------	--

Parameter Description	<p><i>area-id</i>: Indicates the ID of the NSSA area.</p> <p>no-redistribution: Select this option if the router is an NSSA ABR and you want to use only the redistribute command to introduce the routing information into a common area instead of an NSSA area.</p> <p>default-information-originate: Indicates that a default Type 7 LSA is generated and introduced to the NSSA area. This option takes effect only on an NSSA ABR or ASBR.</p> <p>metricvalue: Specifies the metric of the generated default LSA. The value ranges from 0 to 16,777,214. The default value is 1.</p> <p>metric-type: Specifies the route type of the generated default LSA. The values include 1 and 2. 1 represents N-1, and 2 represents N-2. The default value is 2.</p> <p>no-summary: Prohibits the ABR in the NSSA area from sending summary LSAs (Type-3 LSA).</p> <p>translator: Indicates that the NSSA ABR is a translator.</p> <p>stability-intervalseconds: Indicates the stability interval after the NSSA ABR is changed from a translator to a non-translator. The unit is second. The default value is 40. The value ranges from 0 to 2,147,483,647.</p> <p>always: Indicates that the current NSSA ABR always acts as a translator. The default value is the standby translator.</p> <p>translate type7 suppress-fa: Indicates that the Type 5 LSA does not carry the forwarding address (FA) when the Type 7 LSA is converted to the Type 5 LSA.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The default-information-originate parameter is used to generate a default Type 7 LSA. This parameter has different functions on the ABR and the ASBR in the NSSA area. On the ABR, a Type 7 LSA default route is generated regardless of whether the default route exists in the routing table. On the ASBR (not an ABR), a Type 7 LSA default route is generated only when the default route exists in the routing table.</p> <p>If the no-redistribution parameter is configured on the ASBR, other external routes introduced by OSPF through the redistribute command cannot be advertised to the NSSA area. This parameter is generally used when a router in the NSSA area acts both as the ASBR and the ABR. It prevents external routing information from entering the NSSA area.</p> <p>To further reduce the number of LSAs sent to the NSSA area, you can configure the no-summary parameter on the ABR to prevent the ABR from sending the summary LSAs (Type 3 LSA) to the NSSA area.</p>

area default-cost is used on an ABR or ASBR connected to the NSSA area. This command configures the cost of the default route sent from the ABR/ASBR to the NSSA area. By default, the cost of the default route sent to the NSSA area is 1.

If an NSSA area has two or more ABRs, the ABR with the largest router ID is elected by default as the translator for converting Type 7 LSAs into Type 5 LSAs. If the current device is always the translator ABR for converting Type 7 LSAs into Type 5 LSAs, use the **translator always** parameter.

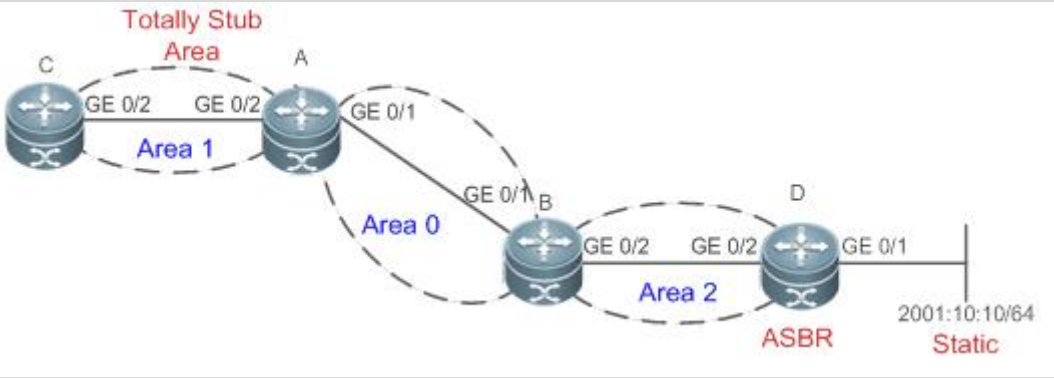
If the translator role of the current device is replaced by another ABR, the conversion capability is retained during the time specified by **stability-interval**. If the router does not become a translator again during **stability-interval**, LSAs that are converted from Type 7 to Type 5 will be deleted from the AS after **stability-interval** expires.

To prevent a routing loop, LSAs that are converted from Type 7 to Type 5 will be deleted from the AS immediately after the current device loses the translator role even if **stability-interval** does not expire.

In the same NSSA area, it is recommended that **translator always** be configured on only one ABR.

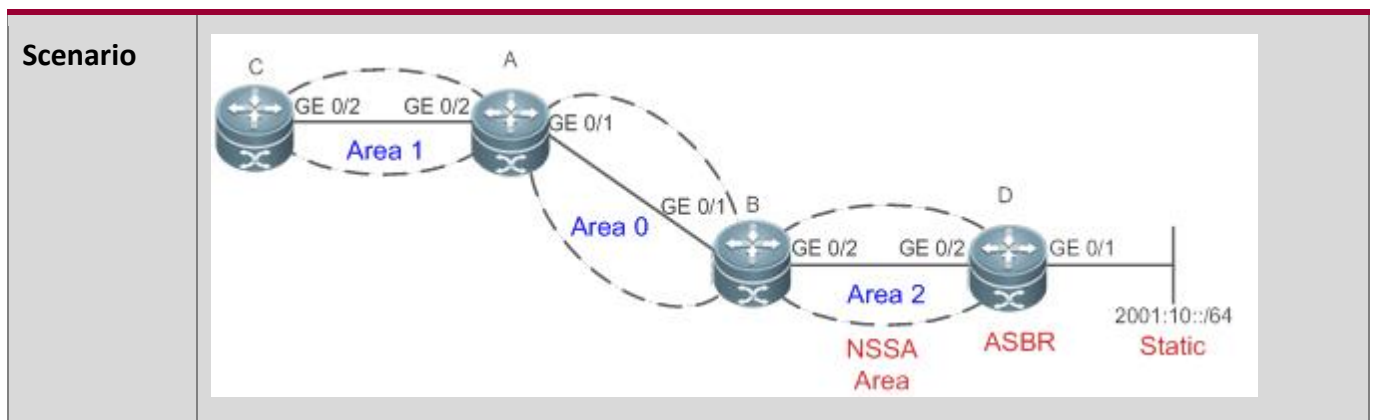
Configuration Example

Configuring a Stub Area

<p>Scenario</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Introduce an external static route to Router D. ▪ Configure area 1 as the stub area on Router A and Router C.
<p>D</p>	<pre>D#configure terminal D(config)#ipv6 router ospf 1 D(config-router)#redistribute static</pre>

A	<pre>A# configure terminal A(config)#ipv6 router ospf 1 A(config-router)#area 1 stub no-summary</pre>
C	<pre>C#configure terminal C(config)#ipv6 router ospf 1 C(config-router)#area 1 stub</pre>
Verification	<ul style="list-style-type: none"> On Router C, run the show ipv6 route ospf command to display the routing table. Verify that there is only one default inter-area route, and no external static route is introduced from Router D.
C	<pre>C#show ipv6 route ospf</pre> <p>IPv6 routing table name - Default - 0 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area</p> <pre>O IA::/0 [110/3] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/2</pre>

Configuring an NSSA Area



Configurati on Steps	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Introduce an external static route to Router D. ▪ Configure area 1 as the NSSA area on Router B and Router D.
D	<pre>D#configure terminal D(config)#ipv6 router ospf 1 D(config-router)#area 1 nssa D(config-router)#redistribute static</pre>
B	<pre>B#configure terminal B(config)#ipv6 router ospf 1 B(config-router)#area 1 nssa</pre>
Verification	<ul style="list-style-type: none"> ▪ On Router D, run the show ipv6 ospf database command to display the database information and verify that Type 7 LSAs are generated. ▪ On Router A, run the show ipv6 route ospf command to display the routing table and verify that an external static route is introduced by Router D.
D	<pre>D#show ipv6 ospf database nssa-external OSPFv3 Router with ID (1.1.1.1) (Process 1) NSSA-external-LSA (Area 0.0.0.1) LS age: 1196 LS Type: NSSA-external-LSA Link State ID: 0.0.0.3 Advertising Router: 1.1.1.1 LS Seq Number: 0x80000004 Checksum: 0x1F25 Length: 52 Metric Type: 2 (Larger than any link state path) Metric: 20 Prefix: 2001:10::/64 Prefix Options: 8 (P - - -) Forwarding Address: 4000::1</pre>
A	<pre>A#show ipv6 route ospf</pre>

IPv6 routing table name - Default - 0 entries

Codes: C - Connected, L - Local, S - Static

R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

IA - Inter area

O N2 2001:10::/64 [110/20] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/1

Common

Errors

- Configurations of the area type are inconsistent on routers in the same area.
- External routes cannot be introduced because route redistribution is configured on a router in the stub area.

5.4.5 Configuring Route Summarization

Configuration

Effect

- Summarize routes to reduce interaction of routing information and the size of routing table, and enhance stability of routes.
- Shield or filter routes.

Notes

- The OSPF basic functions must be configured.
- The address range of the summarize route may exceed the actual network range in the routing table. If data is sent to a network beyond the summarization range, a routing loop may be formed and the router processing load may increase. To prevent these problems, a discard route must be added to the routing table or shield or filter routes.

Configuration

Steps

Configuring Inter-Area Route Summarization

- (Optional) Perform this configuration when routes of the OSPF area need to be summarized.

- Unless otherwise required, perform this configuration on an ABR in the area where routes to be summarized are located.

Configuring External Route Summarization

- (Optional) Perform this configuration when routes external to the OSPF domain need to be summarized.
- Unless otherwise required, perform this configuration on an ASBR, to which routes that need to be summarized are introduced.

Verification

- Run the **show ipv6 route ospf** command to verify that individual routes do not exist and only the summarized route exists.

Related Commands

Configuring Inter-Area Route Summarization

Command	area <i>area-id</i> range <i>ipv6-prefix/prefix-length</i> [advertise not-advertise]
Parameter Description	<p><i>area-id</i>: Specifies the ID of the OSPF area to which the summarized route should be injected. The value can be an integer or an IPv4 prefix.</p> <p><i>ipv6-prefix/prefix-length</i>: Indicates the range of IP addresses to be summarized.</p> <p>advertise not-advertise: Specifies whether the summarized route should be advertised.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>This command takes effect only on an ABR, and is used to summarize multiple routes in an area into a route and advertise this route to other areas. Combination of the routing information occurs only on the boundary of an area. Routers inside the area can learn specific routing information, whereas routers in other areas can learn only one summarized route. In addition, you can set advertise or not-advertise to determine whether to advertise the summarized route to shield and filter routes. By default, the summarized route is advertised. You can use the cost parameter to set the metric of the summarized route.</p> <p>You can configure route summarization commands for multiple areas. This simplifies routes in the entire OSPF routing domain, and improves the network forwarding performance, especially for a large-sized network.</p>

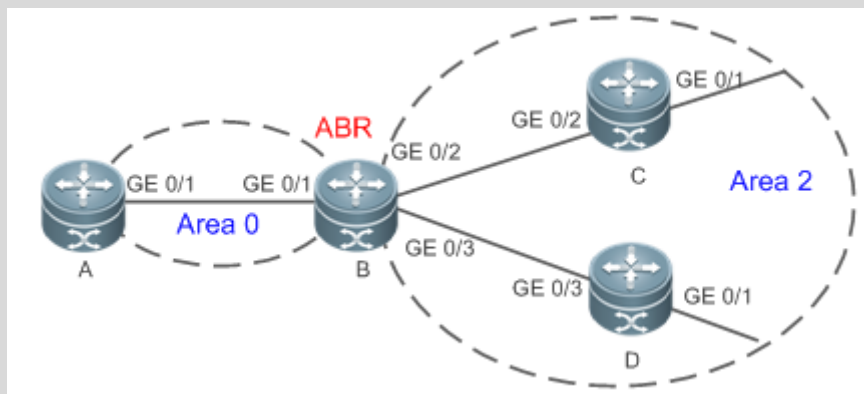
	When multiple route summarization commands are configured and have the inclusive relationship with each other, the area range to be summarized is determined based on the maximum match principle.
--	--

Configuring External Route Summarization

Command	summary-prefix <i>ipv6-prefix/prefix-length</i> [not-advertise tag number]
Parameter Description	<p><i>ipv6-prefix/prefix-length</i>: Indicates the range of IP addresses to be summarized.</p> <p>not-advertise: Indicates that the summarized route is not advertised. If this parameter is not specified, the summarized route is advertised.</p> <p>tagnumber: Specifies the tag value of the route that is redistributed into the OSPFv3 routing domain. The value ranges from 0 to 4,294,967,295.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>When routes are redistributed from other routing processes and injected to the OSPFv3 routing process, each route is advertised to the OSPFv3 routers using an external LSA. If the injected routes are a continuous address space, the ABR can advertise only one summarized route to significantly reduce the size of the routing table.</p> <p>area range summarizes the routes between OSPFv3 areas, whereas summary-prefix summarizes external routes of the OSPFv3 routing domain.</p> <p>When configured on the NSSA ABR translator, summary-prefix summarizes redistributed routes and routes obtained based on the LSAs that are converted from Type 7 to Type 5. When configured on the ASBR (not an NSSA ABR translator), summary-prefix summarizes only redistributed routes.</p>

Configuration Example

**Configurati
on Steps**
Figure 5-13



Remarks **The interface IPv6 addresses are as follows:**
B: GE0/2 2001:16:2::1/64 GE0/3 2001:16:3::1/64
C: GE0/2 2001:16:2::2/64 GE0/1 2001:16:4::2/64
D: GE0/3 2001:16:3::2/64 GE0/1 2001:16:5::1/64

**Configurati
on Steps**

- Enable IPv6 on interfaces of all routers.(Omitted)
- Configure the OSPF basic functions on all routers. (Omitted)
- Summarize routes of area 2 on Router B.

B

```
B#configure terminal
B(config)#ipv6 router ospf 1
B(config-router)#area 2 range 2001:16::/64
```

Verification

On Router A, check the routing table and verify that the entry 2001:16::/64 is generated and other routes do not exist.

A

```
A#show ipv6 route ospf
```

IPv6 routing table name - Default - 0 entries
Codes: C - Connected, L - Local, S - Static
R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
IA - Inter area

```
O IA 2001:16::/64 [110/2] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/1
```

Common

Errors

- Inter-area route summarization cannot be implemented because the **area range** command is configured on a non-ABR device.

5.4.6 Configuring Route Filtering

Configuration

Effect

- Routes that do not meet filtering conditions cannot be loaded to the routing table, or advertised to neighbors. Network users cannot access specified destination network.

Notes

- The OSPF basic functions must be configured.
- Filtering routes by using the **distribute-list in** command affects forwarding of local routes, but does not affect route computation based on LSAs. Therefore, if route filtering is configured on the ABR, Type 3 LSAs will still be generated and advertised to other areas because routes can still be computed based on LSAs. As a result, black-hole routes are generated. In this case, you can run the **area filter-list** or **area range** (containing the **not-advertise** parameter) command on the ABR to prevent generation of black-hole routes.

Configuration

Steps

Configuring Inter-Area Route Filtering

- (Optional) This configuration is recommended if users need to be restricted from accessing the network in a certain OSPF area.
- Unless otherwise required, perform this configuration on an ABR in the area where filtered routes are located.
- **Configuring Redistributed Route Filtering**
- (Optional) Perform this configuration if external routes introduced by the ASBR need to be filtered.
- Unless otherwise required, perform this configuration on an ASBR to which filtered routes are introduced.
- **Configuring Learned Route Filtering**
- (Optional) Perform this configuration if users need to be restricted from accessing a specified destination network.

- Unless otherwise required, perform this configuration on a router that requires route filtering.

Verification

- Run the **show ipv6 route** command to verify that the router is not loaded with routes that have been filtered out.
- Run the **ping** command to verify that the specified destination network cannot be accessed.

Related Commands

Configuring a Passive Interface

Command	passive-interface {default <i>interface-type</i> <i>interface-number</i> }
Parameter Description	<i>interface-type interface-number</i> : Indicates the interface that should be configured as a passive interface. default : Indicates that all interfaces will be configured as passive interfaces.
Command Mode	OSPF routing process configuration mode
Usage Guide	When an interface is configured as a passive interface, it no longer sends or receives Hello packets. This command takes effect only on an OSPFv3-enabled interface, and not on a virtual link.

Configuring Redistributed Route Filtering

Command	distribute-list { <i>name</i> prefix-list <i>prefix-list-name</i> }out[bgp connected isis[<i>area-tag</i>] ospf <i>process-id</i> rip static]
Parameter Description	<i>name</i> : Uses the ACL for filtering. prefix <i>prefix-list-name</i> : Uses the prefix list for filtering. bgp connected isis[area-tag] ospf process-id rip static : Indicates the source of routes to be filtered.
Command Mode	OSPF routing process configuration mode
Usage Guide	distribute-list out is similar to redistribute route-map , and is used to filter routes that are redistributed from other protocols to OSPFv3. The distribute-list out command itself

does not redistribute routes, and is generally used together with the **redistribute** command. The ACL and the prefix list filtering rules are mutually exclusive in the configuration. That is, if the ACL is used for filtering routes coming from a certain source, the prefix list cannot be configured to filter the same routes.


Configuring Learned Route Filtering

Command	distribute-list { <i>name</i> prefix-list <i>prefix-list-name</i> }in [<i>interface-type</i> <i>interface-number</i>]
Parameter Description	<p><i>name</i>: Uses the ACL for filtering.</p> <p>prefix<i>prefix-list-name</i>: Uses the prefix list for filtering.</p> <p><i>interface-type interface-number</i>: Specifies the interface for which LSA routes are filtered.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>Filter routes that are computed based on received LSAs. Only routes meeting the filtering conditions can be forwarded. The command does not affect the LSDB or the routing tables of neighbors. The ACL and the prefix list filtering rules are mutually exclusive in the configuration. That is, if the ACL is used for filtering routes on a specified interface, the prefix list cannot be configured to filter routes on the same interface.</p> <p>Filtering routes by using the distribute-list in command affects forwarding of local routes, but does not affect route computation based on LSAs. Therefore, if route filtering is configured on the ABR, Type 3 LSAs will still be generated and advertised to other areas because routes can still be computed based on LSAs. As a result, black-hole routes are generated. In this case, you can run the area range (containing the not-advertise parameter) command on the ABR to prevent generation of black-hole routes.</p>

Configuration

Example

Scenario
Figure 5-14



Remarks

The interface IPv6 addresses are as follows:
 B: GE0/2 2001:16:2::1/64 GE0/3 2001:16:3::1/64
 C: GE0/2 2001:16:2::2/64 GE0/1 2001:16:4::2/64
 D: GE0/3 2001:16:3::2/64 GE0/1 2001:16:5::1/64

Configuration Steps

- Enable IPv6 on interfaces of all routers. (Omitted)
- Configure the OSPF basic functions on all routers. (Omitted)
- On Router A, configure route filtering.

A

```
A#configure terminal
A(config)#ipv6 access-list test
A (config-ipv6-acl)#permit ipv6 2001:16:5::/64 any
A(config)#ipv6 router ospf 1
A(config-router)#distribute-list test in GigabitEthernet0/1
```

Verification

- On Router A, check the routing table. Verify that only the entry 2001:16:5::/64 is loaded.

A

```
A#show ipv6 route ospf

IPv6 routing table name - Default - 0 entries
Codes: C - Connected, L - Local, S - Static
       R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

IA - Inter area

O IA 2001:16:5::/64 [110/2] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/1

Common Errors

- Filtering routes by using the **distribute-list in** command affects forwarding of local routes, but does not affect route computation based on LSAs. Therefore, if route filtering is configured on the ABR, Type 3 LSAs will still be generated and advertised to other areas because routes can still be computed based on LSAs. As a result, black-hole routes are generated.

5.4.7 Modifying the Route Cost and AD

Configuration Effect

- Change the OSPF routes so that the traffic passes through specified nodes or bypasses specified nodes.
- Change the sequence that a router selects routes so as to change the priorities of OSPF routes.

Notes

- The OSPF basic functions must be configured.
- If you run the **ipv6 ospf cost** command to configure the cost of an interface, the configured cost will automatically overwrite the cost that is computed based on the auto cost.

Configuration Steps

Configuring the Reference Bandwidth

- Optional.
- A router is connected with lines with different bandwidths. This configuration is recommended if you wish to preferentially select the line with a larger bandwidth.

Configuring the Cost of an Interface

- Optional.
- A router is connected with multiple lines. This configuration is recommended if you wish to manually specify a preferential line.

Configuring the Default Metric for Redistribution

- Optional.

- This configuration is mandatory if the cost of external routes of the OSPF domain should be specified when external routes are introduced to an ASBR.

Configuring the Maximum Metric

- Optional.
- A router may be unstable during the restart process or a period of time after the router is restarted, and users do not want to forward data through this router. In this case, this configuration is recommended.

Configuring the AD

- Optional.
- Perform this configuration if you wish to change the priorities of OSPF routes on a router that runs multiple unicast routing protocols.

Verification

- Run the **show ipv6 ospf interface** command to verify that the costs of interfaces are correct.
- Run the **show ipv6 route** command to verify that the costs of external routes introduced by the ASBR are correct.
- Restart the router. Within a specified period of time, data is not forwarded through the restarted router.

Related Commands

Configuring the Reference Bandwidth

Command	auto-cost reference-bandwidth <i>ref-bw</i>
Parameter Description	<i>ref-bw</i> : Indicates the reference bandwidth. The unit is Mbps. The value ranges from 1 to 4, 294, 967.
Command Mode	OSPF routing process configuration mode
Usage Guide	You can run the ipv6 ospf cost command in interface configuration mode to specify the cost of the interface. The priority of this cost is higher than that of the metric computed based on the reference bandwidth.

Configuring the Cost of an Interface

Command	ipv6 ospf cost <i>cost</i>[<i>instanceinstance-id</i>]
Parameter Description	<i>cost</i> : Indicates the cost of an OSPF interface. The value ranges from 0 to 65,535. <i>instanceinstance-id</i> : Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.
Command Mode	Interface configuration mode
Usage Guide	By default, the cost of an OSPFv3 interface is equal to 100 Mbps/Bandwidth, where Bandwidth is the bandwidth of the interface and configured by the bandwidth command in interface configuration mode. The costs of OSPF interfaces on several typical lines are as follows: 64 Kbps serial line: The cost is 1562. E1 line: The cost is 48. 10M Ethernet: The cost is 10. 100M Ethernet: The cost is 1. If you run the ipv6 ospf cost command to configure the cost of an interface, the configured cost will automatically overwrite the cost that is computed based on the auto cost.

Configuring the Cost of the Default Route in a Stub/NSSA Area

Command	area <i>area-id</i> default-cost <i>cost</i>
Parameter Description	<i>area-id</i> : Indicates the ID of the stub/nssa area. <i>cost</i> : Indicates the cost of the default summarized route injected to the stub/nssa area. The value ranges from 0 to 16, 777, 215.
Command Mode	OSPF routing process configuration mode
Usage Guide	This command takes effect only on an ABR in a stub/nssa area.

Configuring the Default Metric for Redistribution

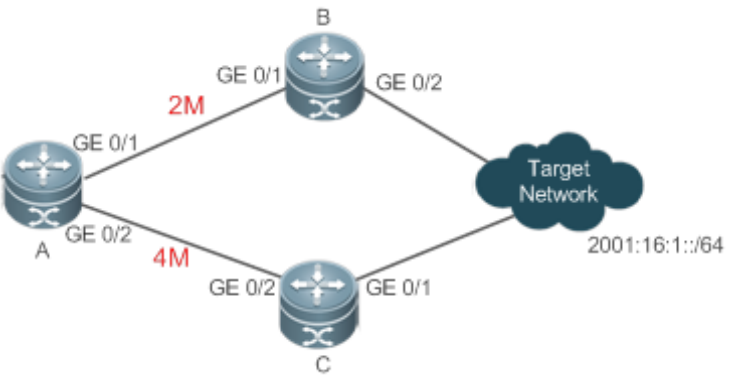
Command	default-metric <i>metric</i>
----------------	-------------------------------------

Parameter Description	<i>metric</i> : Indicates the default metric of the OSPF redistributed route. The value ranges from 1 to 16,777,214.
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The default-metric command must be used together with the redistribute command to modify the initial metrics of all redistributed routes. The default-metric command does not take effect on external routes that are injected to the OSPF routing domain by the default-information originate command.</p> <p>The default metric of a redistributed direct route is always 20.</p>

Configuring the AD

Command	distance { <i>distance</i> ospf{ [<i>intra-area distance</i>] [<i>inter-area distance</i>] [<i>external distance</i>]} }
Parameter Description	<p><i>distance</i>: Indicates the AD of a route. The value ranges from 1 to 255.</p> <p>intra-area<i>distance</i>: Indicates the AD of an intra-area route. The value ranges from 1 to 255.</p> <p>inter-area <i>distance</i>: Indicates the AD of an inter-area route. The value ranges from 1 to 255.</p> <p>external <i>distance</i>: Indicates the AD of an external route. The value ranges from 1 to 255.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>Use this command to specify different ADs for different types of OSPF routes.</p> <p>The AD allows different routing protocols to compare route priorities. A smaller AD indicates a higher route priority.</p> <p>The priorities of routes generated by different OSPFv3 processes must be compared based on ADs.</p> <p>If the AD of a route entry is set to 255, the route entry is not trustworthy and does not participate in packet forwarding.</p>

Configuration Example

0	
Configuration Steps	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ On Router A, configure the cost of each interface.
A	<pre>A#configure terminal A(config)#interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#ipv6 ospf cost 10 A(config)#interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)#ipv6 ospf cost 20</pre>
Verification	<ul style="list-style-type: none"> ▪ On Router A, check the routing table. The next hop of the optimum path to 2001:16:1::/64 is Router B.
A	<pre>A#show ipv6 route ospf</pre> <p>IPv6 routing table name - Default - 0 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area</p> <pre>O E2 2001:16:1::/64 [110/2] via FE80::2D0:F8FF:FE22:4547, GigabitEthernet 0/1</pre>

Common Errors

- If the cost of an interface is set to 0 in the **ipv6 ospf cost** command, a route computation error may occur. For example, a routing loop is obtained.

5.4.8 Enabling Authentication

Configuration

Effect

- All routers connected to the OSPF network must be authenticated to ensure stability of OSPF and protect OSPF against intrusions.

Notes

- The OSPF basic functions must be configured.
- If authentication is configured for an area, the configuration takes effect on all interfaces that belong to this area.
- If authentication is configured for both an interface and the area to which the interface belongs, the configuration for the interface takes effect preferentially.

Configuration

Steps

Configuring Authentication

- Optional.
- Perform this configuration if a router accesses a network that requires authentication.

Configuring Encryption

- Optional.
- Perform this configuration if a router accesses a network that requires encryption.

Configuring Virtual Link Authentication

- Optional.
- Perform this configuration if a router accesses a network that requires authentication.

Configuring Virtual Link Encryption

- Optional.
- Perform this configuration if a router accesses a network that requires encryption.

Verification

- If routers are configured with different authentication keys, run the **show ipv6 ospf neighbor** command to verify that there is no OSPF neighbor.
- If routers are configured with the same authentication key, run the **show ipv6 ospf neighbor** command to verify that there are OSPF neighbors.

Related Commands

Configuring Area-based Authentication

Command	area <i>area-id</i> authentication ipsec spi <i>spi</i> [md5 sha1] [0 7] key
Parameter Description	<p><i>area-id</i>: Indicates the area ID. The value can be an integer or an IPv4 prefix.</p> <p><i>spi</i>: Indicates the SPI. The value ranges from 256 to 4,294,967,295.</p> <p>md5: Enables MD5 authentication.</p> <p>sha1: Enables SHA1 authentication.</p> <p>0: Indicates that the key is displayed in plain text.</p> <p>7: Indicates that the key is displayed in cipher text.</p> <p><i>key</i>: Indicates the authentication key.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The RGOS supports three authentication types:</p> <ul style="list-style-type: none">▪ No authentication▪ MD5 authentication▪ SHA1 authentication <p>Configuration of area-based authentication for OSPFv3 takes effect on all interfaces (except virtual links) in the area, but the interface-based authentication configuration takes precedence over the area-based configuration.</p>

Configuring Area-based Encryption and Authentication

Command	area <i>area-id</i> encryption ipsec spi <i>spi</i> esp [null [des 3des] [0 7] <i>des-key</i>] [md5 sha1] [0 7] <i>key</i>
Parameter Description	<p><i>area-id</i>: Indicates the area ID. The value can be an integer or an IPv4 prefix.</p> <p><i>spi</i>: Indicates the SPI. The value ranges from 256 to 4,294,967,295.</p> <p>null: Indicates that no encryption mode is used.</p> <p>des: Indicates that the Data Encryption Standard (DES) mode is used.</p> <p>3des: Indicates that the Triple DES (3DES) mode is used.</p> <p><i>des-key</i>: Indicates the encryption key.</p> <p>md5: Enables MD5 authentication.</p>

	<p>sha1: Enables SHA1 authentication.</p> <p>0: Indicates that the key is displayed in plain text.</p> <p>7: Indicates that the key is displayed in cipher text.</p> <p>key: Indicates the authentication key.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The RGOS supports two encryption modes and two authentication modes.</p> <p>The two encryption modes are as follows:</p> <ul style="list-style-type: none"> ▪ DES ▪ 3DES <p>The two authentication modes are as follows:</p> <ul style="list-style-type: none"> ▪ MD5 ▪ SHA1 <p>Configuration of area-based encryption and authentication for OSPFv3 takes effect on all interfaces (except virtual links) in the area, but the interface-based encryption and authentication configuration takes precedence over the area-based configuration.</p>

Configuring Interface-based Authentication

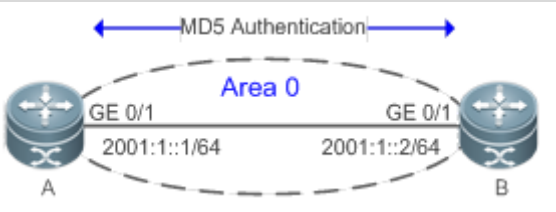
Command	ipv6 ospfauthentication[null ipsec spspi[md5 sha1] [0 7]key][instanceinstance-id]
Parameter Description	<p><i>area-id:</i> Indicates the area ID. The value can be an integer or an IPv4 prefix.</p> <p><i>spsi:</i> Indicates the SPI. The value ranges from 256 to 4,294,967,295.</p> <p>md5: Enables MD5 authentication.</p> <p>sha1: Enables SHA1 authentication.</p> <p>0: Indicates that the key is displayed in plain text.</p> <p>7: Indicates that the key is displayed in cipher text.</p> <p><i>key:</i> Indicates the authentication key.</p> <p>instance instance-id: Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.</p>
Command Mode	OSPF routing process configuration mode

Usage Guide	<p>The RGOS supports three authentication types:</p> <ul style="list-style-type: none"> ▪ No authentication ▪ MD5 authentication ▪ SHA1 authentication <p>OSPFv3 authentication parameters configured on interconnected interfaces must be consistent.</p>
--------------------	---

Configuring Interface-based Encryption and Authentication

Command	<pre>ipv6 ospf encryption ipsec spi <i>spi</i> esp [null [des 3des] [0 7] <i>des-key</i>] [md5 sha1] [0 7] <i>key</i> [<i>instance-id</i>]</pre>
Parameter Description	<p><i>spi</i>: Indicates the SPI. The value ranges from 256 to 4,294,967,295.</p> <p>null: Indicates that no encryption mode is used.</p> <p>des: Indicates that the DES mode is used.</p> <p>3des: Indicates that the 3DES mode is used.</p> <p><i>des-key</i>: Indicates the encryption key.</p> <p>md5: Enables MD5 authentication.</p> <p>sha1: Enables SHA1 authentication.</p> <p>0: Indicates that the key is displayed in plain text.</p> <p>7: Indicates that the key is displayed in cipher text.</p> <p><i>key</i>: Indicates the authentication key.</p> <p>instance <i>instance-id</i>: Indicates the ID of a specified OSPFv3 process of the interface. The value ranges from 0 to 255.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The RGOS supports two encryption modes and two authentication modes.</p> <p>The two encryption modes are as follows:</p> <ul style="list-style-type: none"> ▪ DES ▪ 3DES <p>The two authentication modes are as follows:</p> <ul style="list-style-type: none"> ▪ MD5 ▪ SHA1 <p>OSPFv3 encryption and authentication parameters configured on the local interface must be consistent with those configured on the interconnected interfaces.</p>

Configuration Example

Scenario Figure 5-16	
Configuration Steps	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Configure MD5 authentication for interfaces of all routers.
A	<pre>A#configure terminal A(config)#interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#ipv6 ospf authentication ipsec spi 256 md5 01234567890123456789012345678912</pre>
B	<pre>B# configure terminal B(config)#interface GigabitEthernet 0/3 B(config-if-GigabitEthernet 0/3)#ipv6 ospf authentication ipsec spi 256 md5 01234567890123456789012345678912</pre>
Verification	<ul style="list-style-type: none"> ▪ On Router A and Router B, verify that the OSPF neighbor status is correct.
A	<pre>A# show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Instance ID Interface 2.2.2.2 1 Full/DR 00:00:38 0 GigabitEthernet 0/1</pre>
B	<pre>B# show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Instance ID Interface 1.1.1.1 1 Full/BDR 00:00:38 0 GigabitEthernet 0/1</pre>

Common Errors

- The configured authentication modes are inconsistent.
- The configured authentication keys are inconsistent.

5.4.9 Modifying the Maximum Number of Concurrent Neighbors

Configuration

Effect

- Control the maximum number of concurrent neighbors on the OSPF process to ease the pressure on the device.

Notes

- The OSPF basic functions must be configured.

Configuration

Steps

Configuring the Maximum Number of Concurrent Neighbors on the OSPF Process

- (Optional) This configuration is recommended if you wish to set up the OSPF adjacency more quickly when a router is connected with a lot of other routers.
- Perform this configuration on a core router.

Verification

- Run the **show ipv6 ospf neighbor** command to display the number of neighbors that are concurrently interacting with the OSPF process.

Related

Commands

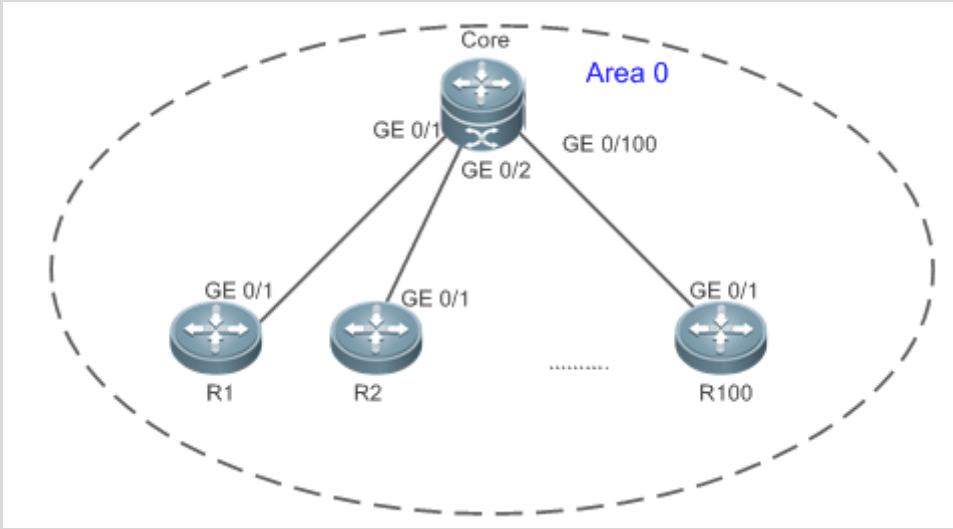
Configuring the Maximum Number of Concurrent Neighbors on the Current Process

Command	max-concurrent-ddnumber
Parameter Description	<i>number</i> : Specifies the maximum number of neighbors that are concurrently interacting with the OSPF process. The value ranges from 1 to 65,535.
Command Mode	OSPF routing process configuration mode
Usage Guide	When the performance of a router is affected because the router exchanges data with multiple neighbors, you can configure this command to restrict the maximum of neighbors with which each OSPF process can concurrently initiate or accept interaction.

Configuring the Maximum Number of Concurrent Neighbors on All Processes

Command	ipv6 router ospf max-concurrent-ddnumber
Parameter Description	<i>number</i> : Specifies the maximum number of neighbors that are concurrently interacting with the OSPF process. The value ranges from 1 to 65,535.
Command Mode	Global configuration mode
Usage Guide	When the performance of a router is affected because the router exchanges data with multiple neighbors, you can configure this command to restrict the maximum of neighbors with which all OSPF processes can concurrently initiate or accept interaction.

Configuration Example

Scenario Figure 5-17	
Configuration Steps	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ On the Router Core, set the maximum number of concurrent neighbors to 4.
Core	<pre>Core# configure terminal Core(config)# ipv6 router ospf max-concurrent-dd 4</pre>
Verification	<ul style="list-style-type: none"> ▪ On the Router Core, check the neighbor status and verify that at most eight neighbors concurrently interact with the OSPF process.

Common Errors

N/A

5.4.10 Disabling MTU Verification

Configuration Effect

- The unicast routing service can be provided even if the MTUs of interfaces on neighbor routers are different.

Notes

- The OSPF basic functions must be configured.

Configuration Steps

Disabling MTU Verification

- (Optional) MTU verification is disabled by default. You are advised to retain the default configuration.
- Perform this configuration on two routers with different interface MTUs.

Verification

- The adjacency can be set up between routers with different MTUs.

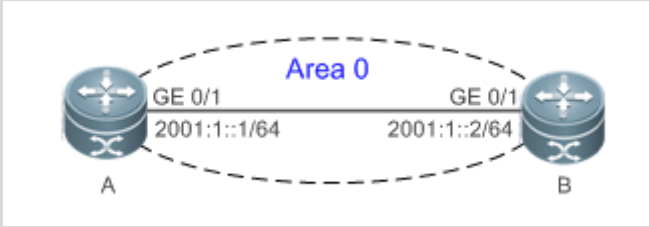
Related Commands

Disabling MTU Verification

Command	<code>ipv6 ospf mtu-ignore</code>
Parameter Description	N/A
Command Mode	Interface configuration mode
Usage Guide	On receiving the database description packet, OSPF checks whether the MTU of the interface on the neighbor is the same as the MTU of its own interface. If the interface

MTU specified in the received database description packet is greater than the MTU of the local interface, the adjacency cannot be set up. To resolve this problem, you can disable MTU verification.

Configuration Example

<p>Scenario Figure 5-18</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Configure different MTUs for interfaces on two routers. ▪ Disable MTU verification on all routers. (By default, the function of disabling MTU verification is enabled.)
<p>A</p>	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ipv6 mtu 1400 A(config-if-GigabitEthernet 0/1)#ipv6 ospf mtu-ignore</pre>
<p>B</p>	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ipv6 mtu 1600 B(config-if-GigabitEthernet 0/1)# ipv6 ospf mtu-ignore</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ On Router A, verify that the OSPF neighbor information is correct.
<p>A</p>	<pre>A# show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Instance ID Interface 2.2.2.2 1 Full/DR 00:00:38 0 GigabitEthernet 0/1</pre>

Common Errors

N/A

5.4.11 Enabling Two-Way Maintenance

Configuration

Effect

- Non-Hello packets can also be used to maintain the adjacency.

Notes

- The OSPF basic functions must be configured.

Configuration

Steps

Enabling Two-Way Maintenance

- (Optional) This function is enabled by default. You are advised to retain the default configuration.
- Perform this configuration on all routers.

Verification

- Non-Hello packets can also be used to maintain the adjacency.

Related

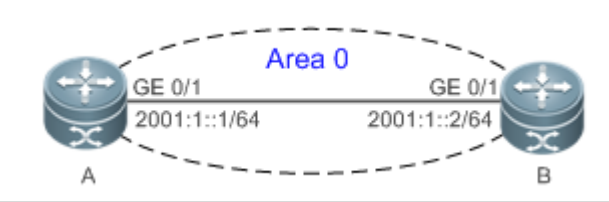
Commands

Enabling Two-Way Maintenance

Command	<code>two-way-maintain</code>
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	On a large network, a lot of packets may be sent or received, occupying too much CPU and memory. As a result, some packets are delayed or discarded. If the processing time of Hello packets exceeds the dead interval, the adjacency will be destroyed due to timeout. If the two-way maintenance function is enabled, in addition to the Hello packets, the DD, LSU, LSR, and LSAck packets can also be used to maintain the bidirectional communication between neighbors when a large number of packets exist

on the network. This prevents termination of the adjacency caused by delayed or discarded Hello packets.

Configuration Example

<p>Scenario Figure 5-19</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ On Router A, enable the two-way maintenance function. (This function is enabled by default.)
<p>A</p>	<pre>A# configure terminal A(config)# ipv6 routerospf 1 A(config-router)#two-way-maintain</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ When the adjacency is being set up, Router A checks the neighbor dead interval and updates the dead interval without waiting for Router B to send a Hello packet.
<p>A</p>	<pre>A# show ipv6 ospfneighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State Dead Time Instance ID Interface 2.2.2.2 1 Full/DR 00:00:38 0 GigabitEthernet 0/1</pre>

Common Errors

N/A

5.4.12 Correlating OSPF with BFD

Configuration Effect

- Once a link is faulty, OSPF can quickly detect the failure of the route. This configuration helps shorten the traffic interruption time.

Notes

- The OSPF basic functions must be configured.
- The BFD parameters must be configured for the interface in advance.
- If BFD is configured for both a process and an interface, the interface-based configuration takes effect preferentially.

Configuration

Steps

Correlating OSPF with BFD

- (Optional) Perform this configuration if you wish to accelerate OSPF network convergence.
- Perform this configuration on routers at both ends of the link.

Verification

- Run the **show bfd neighbor** command to verify that the BFD neighbors are normal.

Related

Commands

Correlating an OSPF Interface with BFD

Command	ipv6 ospf bfd [disable]
Parameter Description	disable: Disables BFD for link detection on a specified OSPF-enabled interface.
Command Mode	Interface configuration mode
Usage Guide	The interface-based configuration takes precedence over the bfd all-interfaces command used in process configuration mode. Based on the actual environment, you can run the ipv6 ospf bfd command to enable BFD on a specified interface for link detection, or run the bfd all-interfaces command in OSPF process configuration mode to enable BFD on all interface of the OSPF process, or run the ipv6 ospf bfd disable command to disable BFD on a specified interface.

Correlating an OSPF Process with BFD

Command	bfd all-interfaces
----------------	---------------------------

Parameter Description	N/A
Command Mode	OSPF process configuration mode
Usage Guide	<p>OSPF dynamically discovers neighbors through the Hello packets. After OSPF enables the BFD function, a BFD session will be set up to achieve the full adjacency, and use the BFD mechanism to detect the neighbor status. Once a neighbor failure is detected through BFD, OSPF performs network convergence immediately.</p> <p>You can also run the ipv6 ospf bfd [disable] command in interface configuration mode to enable or disable the BFD function on a specified interface, and this configuration takes precedence over the bfd all-interfaces command used in OSPF process configuration mode.</p>

Configuration Example

<p>Scenario</p> <p>Figure 5-20</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Configure the BFD parameters for interfaces of all routers. ▪ Correlate OSPF with BFD on all routers.
<p>A</p>	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet0/1)#bfd interval 200 min_rx 200 multiplier 5 A(config)# ipv6 router ospf 1 A(config-router)#bfd all-interfaces</pre>

B	<pre>B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 2/1)#bfd interval 200 min_rx 200 multiplier 5 B(config)# ipv6 router ospf 1 B(config-router)#bfd all-interfaces</pre>
Verification	<ul style="list-style-type: none"> On Router A and Router B, verify that the BFD state is Up. Disconnect Router B from the switch. On Router A, verify that a neighbor is found disconnected during BFD, and the corresponding OSPF route is deleted.
A	<pre>A# show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State BFD State Dead Time Instance ID Interface 2.2.2.2 1 Full/BDR Up 00:00:35 0 GigabitEthernet 0/1</pre>
B	<pre>B# show ipv6 ospf neighbor OSPFv3 Process (1), 1 Neighbors, 1 is Full: Neighbor ID Pri State BFD State Dead Time Instance ID Interface 1.1.1.1 1 Full/DR Up 00:00:35 0 GigabitEthernet 0/1</pre>

Common Errors

N/A

5.4.13 Enabling GR

Configuration Effect

- When a distributed route switches services from the active board to the standby board, traffic forwarding continues and is not interrupted.
- When the OSPF process is being restarted, data forwarding continues and is not interrupted.

Notes

- The OSPF basic functions must be configured.
- The neighbor router must support the GR helper function.
- The grace period cannot be shorter than the neighbor dead time of the neighbor router.

Configuration Steps

Configuring the OSPF GR Function

- (Optional) This function is enabled by default. You are advised to retain the default configuration.
- Perform this configuration on routers where hot standby switchover is triggered or the OSPF process is restarted.

Configuring the OSPF GR Helper Function

- (Optional) This function is enabled by default. You are advised to retain the default configuration.
- Perform this configuration on a router if hot standby switchover is triggered or the OSPF process is restarted on a neighbor of this router.

Verification

- When a distributed router switches services from the active board to the standby board, data forwarding continues and the traffic is not interrupted.
- When the OSPF process is being restarted, data forwarding continues and the traffic is not interrupted.

Related Commands

Configuring the OSPF GR Function

Command	<code>graceful-restart [grace-period <i>grace-period</i> inconsistent-lsa-checking]</code>
Parameter Description	<p>grace-period <i>grace-period</i>: Indicates the grace period, which is the maximum time from occurrence of an OSPF failure to completion of the OSPF GR. The value of the grace period varies from 1s to 1800s. The default value is 120s.</p> <p>inconsistent-lsa-checking: Enables topological change detection. If any topological change is detected, OSPF exits the GR process to complete convergence. After GR is enabled, topological change detection is enabled by default.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	The GR function is configured based on the OSPF process. You can configure different parameters for different OSPF processes based on the actual conditions. This command is used to configure the GR restarter capability of a device. The grace period is the maximum time of the entire GR process, during which link status is rebuilt so that the

original state of the OSPF process is restored. After the grace period expires, OSPF exits the GR state and performs common OSPF operations.

Run the **graceful-restart** command to set the grace period to 120s. The **graceful-restart grace-period** command allows you to modify the grace period explicitly.

The precondition for successful execution of GR and uninterrupted forwarding is that the topology remains stable. If the topology changes, OSPF quickly converges without waiting for further execution of GR, thus avoiding long-time forwarding black-hole.

- Disabling topology detection: If OSPF cannot converge in time when the topology changes during the hot standby process, forwarding black-hole may appear in a long time.
- Enabling topology detection: Forwarding may be interrupted when topology detection is enabled, but the interruption time is far shorter than that when topology detection is disabled.

In most cases, it is recommended that topology detection be enabled. In special scenarios, topology detection can be disabled if the topology changes after the hot standby process, but it can be ensured that the forwarding black-hole will not appear in a long time. This can minimize the forwarding interruption time during the hot standby process.

If the Fast Hello function is enabled, the GR function cannot be enabled.

Configuring the OSPF GR Helper Function

Command	<code>graceful-restart helper { disable strict-lsa-checking internal-lsa-checking }</code>
Parameter Description	<p>disable: Prohibits a device from acting as a GR helper for another device.</p> <p>strict-lsa-checking: Indicates that changes in Type 1 to Type 5 and Type 7 LSAs will be checked during the period that the device acts as a GR helper to determine whether the network changes. If the network changes, the device will stop acting as the GR helper.</p> <p>internal-lsa-checking: Indicates that changes in Type 1 to Type 3 LSAs will be checked during the period that the device acts as a GR helper to determine whether the network changes. If the network changes, the device will stop acting as the GR helper.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	This command is used to configure the GR helper capability of a router. When a neighbor router implements GR, it sends a Grace-LSA to notify all neighbor routers. If the GR helper function is enabled on the local router, the local router becomes the GR helper on receiving the Grace-LSA, and helps the neighbor to complete GR. The disable option indicates that GR helper is not provided for any device that implements GR.

After a device becomes the GR helper, the network changes are not detected by default. If any change takes place on the network, the network topology converges after GR is completed. If you wish that network changes can be quickly detected during the GR process, you can configure **strict-lsa-checking** to check Type 1 to 5 and Type 7 LSAs that indicate the network information or **internal-lsa-checking** to check Type 1 to 3 LSAs that indicate internal routes of the AS domain. When the network scale is large, it is recommended that you disable the LSA checking options (**strict-lsa-checking** and **internal-lsa-checking**) because regional network changes may trigger termination of GR and consequently reduce the convergence of the entire network.

Configuration Example

<p>Scenario Figure 5-21</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ On Router A, Router C, and Router D, enable the GR helper function. (This function is enabled by default.) ▪ On Router B, enable the GR function.
<p>B</p>	<pre>B# configure terminal B(config)# ipv6 router ospf1 B(config-router)# graceful-restart</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Trigger a hot standby switchover on Router B, and verify that the routing tables of destination Network 1 and Network 2 remain unchanged on Router A during the switchover. ▪ Trigger a hot standby switchover on Router B, ping destination Network 1 from Router A, and verify that traffic forwarding is not interrupted during the switchover.

Common Errors

- Traffic forwarding is interrupted during the GR process because the configured grace period is shorter than the neighbor dead time of the neighbor router.

5.4.14 Configuring Network Management Functions

Configuration Effect

- Use the network management software to manage OSPF parameters and monitor the OSPF running status.

Notes

- The OSPF basic functions must be configured.
- You must enable the MIB function of the SNMP server before enabling the OSPF MIB function.
- You must enable the trap function of the SNMP server before enabling the OSPF trap function.
- You must enable the logging function of the device before outputting the OSPF logs.

Configuration Steps

Binding the MIB with the OSPF Process

- (Optional) This configuration is required if you want to use the network management software to manage parameters of a specified OSPF process.
- Perform this configuration on all routers.

Enabling the Trap Function

- (Optional) This configuration is required if you want to use the network management software to monitor the OSPF running status.
- Perform this configuration on all routers.

Configuring the Logging Function

- (Optional) This function is enabled by default. You are advised to retain the default configuration. If you want to reduce the log output, disable this function.
- Perform this configuration on all routers.

Verification

- Use the network management software to manage the OSPF parameters.
- Use the network management software to monitor the OSPF running status.

Related Commands

Binding the MIB with the OSPF Process

Command	enable mib-binding
Parameter Description	N/A
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The OSPFv2 MIB does not have the OSPFv3 process information. Therefore, you can perform operations only on a single OSPFv2 process through SNMP. By default, the OSPFv3 MIB is bound with the OSPFv3 process with the smallest process ID, and all user operations take effect on this process.</p> <p>If you wish to perform operations on a specified OSPFv3 process through SNMP, run this command to bind the MIB with the process.</p>

Enabling the Trap Function

Command	enable traps[error [IfConfigError IfRxBadPacket VirtIfConfigError VirtIfRxBadPacket] state-change[IfStateChange NbrStateChange NssaTranslatorStatusChange VirtIfStateChange VirtNbrStateChange RestartStatusChange NbrRestartHelperStatusChange VirtNbrRestartHelperStatusChange]]
Parameter Description	<p>IfConfigError: Indicates that an interface parameter configuration error occurs.</p> <p>IfRxBadPacket: Indicates that the interface receives a bad packet.</p> <p>VirtIfConfigError: Indicates that a virtual interface parameter configuration error occurs.</p> <p>VirtIfRxBadPacket: Indicates that the virtual interface receives a bad packet.</p> <p>IfStateChange: Indicates that interface state changes.</p> <p>NbrStateChange: Indicates that the neighbor state changes.</p> <p>NssaTranslatorStatusChange: Indicates that the NSSA state changes.</p> <p>VirtIfStateChange: Indicates that the virtual interface state changes.</p> <p>VirtNbrStateChange: Indicates that the virtual neighbor state changes.</p> <p>RestartStatusChange: Indicates that the GR state of the local device changes.</p>

	<p>NbrRestartHelperStatusChange: Indicates that the state of the neighbor GR process changes.</p> <p>VirtNbrRestartHelperStatusChange: Indicates that the GR state of the virtual neighbor changes.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>The function configured by this command is restricted by the snmp-server command. You can configure snmp-server enable traps ospf and then enable traps command before the corresponding OSPF traps can be correctly sent out.</p> <p>This command is not restricted by the MIB bound with the process. The trap function can be enabled concurrently for different processes.</p>

Configuring the Logging Function

Command	log-adj-changes[detail]
Parameter Description	detail: Records all status change information.
Command Mode	OSPF routing process configuration mode
Usage Guide	N/A

Configuration Example

<p>Scenario</p> <p>Figure 5-22</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Enable IPv6 on interfaces of all routers.(Omitted) ▪ Configure the OSPF basic functions on all routers. (Omitted) ▪ Bind the MIB with the OSPF process on Router A. ▪ Enable the trap function on Router A.

A	<pre>A# configure terminal A(config)#snmp-server host 192.168.2.2 traps version 2c public A(config)#snmp-server community public rw A(config)#snmp-server enable traps A(config)# A(config)# ipv6 routerospf 10 A(config-router)# enable mib-binding A(config-router)# enable traps</pre>
Verification	<ul style="list-style-type: none"> Use the MIB tool to read and set the OSPF parameters and display the OSPF running status.

Common Errors

N/A

5.4.15 Modifying Protocol Control Parameters

Configuration Effect

- Modify protocol control parameters to change the protocol running status.

Notes

- The OSPF basic functions must be configured.
- The neighbor dead time cannot be shorter than the Hello interval.

Configuration Steps

Configuring the Hello Interval

- (Optional) You are advised to retain the default configuration.
- Perform this configuration on routers at both end of a link.

Configuring the Dead Interval

- (Optional) You are advised to retain the default configuration. This configuration can be adjusted if you wish to accelerate OSPF convergence when a link fails.
- Perform this configuration on routers at both end of a link.

Configuring the LSU Retransmission Interval

- (Optional) You are advised to adjust this configuration if a lot of routes exist in the user environment and network congestion is serious.

Configuring the LSA Generation Time

- (Optional) You are advised to retain the default configuration.

Configuring the LSA Group Refresh Time

- (Optional) You are advised to retain the default configuration. This configuration can be adjusted if a lot of routes exist in the user environment.
- Perform this configuration on an ASBR or ABR.

Configuring LSA Repeated Receiving Delay

- (Optional) You are advised to retain the default configuration.

Configuring the SPF Computation Delay

- (Optional) This configuration can be adjusted if network flapping frequently occurs.

Configuring the Inter-Area Route Computation Delay

- (Optional) You are advised to retain the default configuration.
- Perform this configuration on all routers.

Configuring the Inter-Area Route Computation Delay

- (Optional) You are advised to retain the default configuration.
- Perform this configuration on all routers.

Verification

- Run the **show ipv6 ospf** and **show ipv6 ospf neighbor** commands to display the protocol running parameters and status.

Related

Commands

Configuring the Hello Interval

Command	<code>ipv6ospf hello-intervalseconds</code>
Parameter Description	<i>seconds</i> : Indicates the interval at which OSPF sends the Hello packet. The unit is second. The value ranges from 1 to 65,535.

Command Mode	Interface configuration mode
Usage Guide	The Hello interval is contained in the Hello packet. A shorter Hello interval indicates that OSPF can detect topological changes more quickly, but the network traffic increases. The Hello interval must be the same on all routers in the same network segment. If you want to manually modify the neighbor dead interval, ensure that the neighbor dead interval is longer than the Hello interval.

Configuring the Dead Interval

Command	ipv6ospf dead-interval <i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the time that the neighbor is declared lost. The unit is second. The value ranges from 2 to 65535.
Command Mode	Interface configuration mode
Usage Guide	<p>The OSPF dead interval is contained in the Hello packet. If OSPF does not receive a Hello packet from a neighbor within the dead interval, it declares that the neighbor is invalid and deletes this neighbor record from the neighbor list. By default, the dead interval is four times the Hello interval. If the Hello interval is modified, the dead interval is modified automatically.</p> <p>When using this command to manually modify the dead interval, pay attention to the following issues:</p> <ol style="list-style-type: none"> 1. The dead interval cannot be shorter than the Hello interval. 2. The dead interval must be the same on all routers in the same network segment.

Configuring the LSU Transmission Delay

Command	ipv6ospf transmit-delay<i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the LSU transmission delay on the OSPF interface. The unit is second. The value ranges from 0 to 65, 535.
Command Mode	Interface configuration mode

Usage Guide	<p>Before an LSU packet is transmitted, the Age fields in all LSAs in this packet will increase based on the amount specified by the ip ospf transmit-delay command. Considering the transmission delay and line propagation delay on the interface, you need to set the LSU transmission delay to a greater value for a low-speed line or interface. The LSU transmission delay of a virtual link is defined by the transmit-delay parameter in the area virtual-link command.</p> <p>If the value of the Age field of an LSA reaches 3600, the packet will be retransmitted or a retransmission will be requested. If the LSA is not updated in time, the expired LSA will be deleted from the LSDB.</p>
-------------	---

Configuring the LSU Retransmission Interval

Command	ipv6ospf retransmit-intervalseconds
Parameter Description	<i>seconds</i> : Indicates the LSU retransmission interval. The unit is second. The value ranges from 0 to 65,535. This interval must be longer than the round-trip transmission delay of data packets between two neighbors.
Command Mode	Interface configuration mode
Usage Guide	<p>After a router finishes sending an LSU packet, this packet is still kept in the transmit buffer queue. If an acknowledgment from the neighbor is not received within the time defined by the ip ospf retransmit-interval command, the router retransmits the LSU packet.</p> <p>The retransmission delay can be set to a greater value on a serial line or virtual link to prevent unnecessary retransmission. The LSU retransmission delay of a virtual link is defined by the retransmit-interval parameter in the area virtual-link command.</p>

Configuring the LSA Generation Time

Command	timers throttle lsa all <i>delay-time hold-time max-wait-time</i>
Parameter Description	<p><i>delay-time</i>: Indicates the minimum delay for LSA generation. The first LSA in the database is always generated instantly. The value ranges from 0 to 600,000. The unit is ms.</p> <p><i>hold-time</i>: Indicates the minimum interval between the first LSA update and the second LSA update. The value ranges from 1 to 600,000. The unit is ms.</p> <p><i>max-wait-time</i>: Indicates the maximum interval between two LSA updates when the LSA is updated continuously. This interval is also used to determine whether the LSA is updated continuously. The value ranges from 1 to 600,000. The unit is ms.</p>

Command Mode	OSPF routing process configuration mode
Usage Guide	<p>If a high convergence requirement is raised when a link changes, you can set delay-time to a smaller value. You can also appropriately increase values of the preceding parameters to reduce the CPU usage.</p> <p>When configuring this command, the value of hold-time cannot be smaller than the value of delay-time, and the value of max-wait-time cannot be smaller than the value of hold-time.</p>

Configuring the LSA Group Refresh Time

Command	timers pacinglsa-group <i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the LSA group pacing interval. The value ranges from 10 to 1, 800. The unit is second.
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>Every LSA has a time to live (LSA age). When the LSA age reaches 1800s, a refreshment is needed to prevent LSAs from being cleared because their ages reaching the maximum. If LSA update and aging computation are performed for every LSA, the device will consume a lot of CPU resources. In order to use CPU resources effectively, you can refresh LSAs by group on the device. The interval of group refreshment is called group pacing interval. The group refreshment operation is to organize the LSAs generated within a group pacing interval into a group and refresh the group as a whole.</p> <p>If the total number of LSAs does not change, a larger group pacing interval indicates that more LSAs need to be processed after timeout. To maintain the CPU stability, the number of LSAs processes upon each timeout cannot be too large. If the number of LSAs is large, you are advised to reduce the group pacing interval. For example, if there are 1000 LSAs in the database, you can reduce the pacing interval; if there are 40 to 100 LSAs, you can set the pacing interval to 10-20 minutes.</p>

Configuring the LSA Group Refresh Interval

Command	timers pacing lsa-transmit <i>transmit-time transmit-count</i>
Parameter Description	<i>transmit-time</i> : Indicates the LSA group transmission interval. The value ranges from 10 to 600,000. The unit is ms.

	<i>transmit-count</i> : Indicates the number of LS-UPD packets in a group. The value ranges from 1 to 200.
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>If the number of LSAs is large and the device load is heavy in an environment, properly configuring transmit-time and transmit-count can limit the number of LS-UPD packets flooded on a network.</p> <p>If the CPU usage is not high and the network bandwidth load is not heavy, reducing the value of transmit-time and increasing the value of transmit-count can accelerate the environment convergence.</p>

Configuring LSA Repeated Receiving Delay

Command	timers lsa arrival <i>arrival-time</i>
Parameter Description	<i>arrival-time</i> : Indicates the delay after which the same LSA is received. The value ranges from 0 to 600,000. The unit is ms.
Command Mode	OSPF routing process configuration mode
Usage Guide	No processing is performed if the same LSA is received within the specified time.

Configuring the SPF Computation Delay

Command	timers throttle spf <i>spf-delay spf-holdtime spf-max-waittime</i>
Parameter Description	<p><i>spf-delay</i>: Indicates the SPF computation delay. The unit is ms. The value ranges from 1 to 600,000. When detecting a topological change, the OSPF routing process triggers the SPF computation at least after spf-delay elapses.</p> <p><i>spf-holdtime</i>: Indicates the minimum interval between two SPF computations. The unit is ms. The value ranges from 1 to 600,000.</p> <p><i>spf-max-waittime</i>: Indicates the maximum interval between two SPF computations. The unit is ms. The value ranges from 1 to 600,000.</p> <p><i>number</i>: Indicates the metric of the summarized route.</p>

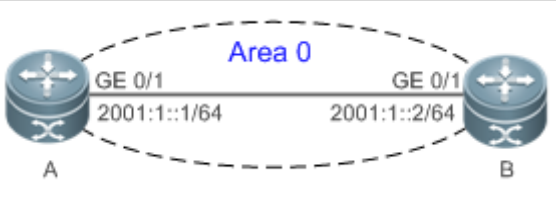
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>spf-delay indicates the minimum time between the occurrence of the topological change and the start of SPF computation. spf-holdtime indicates the minimum interval between the first SPF computation and the second SPF computation. After that, the interval between two SPF computations must be at least twice of the previous interval. When the interval reaches spf-max-waittime, the interval cannot increase again. If the interval between two SPF computations already exceeds the required minimum value, the interval is computed by starting from spf-holdtime.</p> <p>You can set spf-delay and spf-holdtime to smaller values to accelerate topology convergence, and set spf-max-waittime to a larger value to reduce SPF computation. Flexible settings can be used based on stability of the network topology.</p> <p>Compared with the timers spf command, this command supports more flexible settings to accelerate the convergence speed of SPF computation and further reduce the system resources consumed by SPF computation when the topology continuously changes. Therefore, you are advised to use the timers throttle spf command for configuration.</p> <ol style="list-style-type: none"> 1. The value of spf-holdtime cannot be smaller than the value of spf-delay; otherwise, spf-holdtime will be automatically set to the value of spf-delay. 2. The value of spf-max-waittime cannot be smaller than the value of spf-holdtime; otherwise, spf-max-waittime will be automatically set to the value of spf-holdtime. 3. The configurations of timers throttle spf and timers spf are mutually overwritten. 4. When both timers throttle spf and timers spf are not configured, the default values of timers throttle spf prevail.

Configuring the Computation Delays of Inter-Area Routes and External Routes

Command	timers throttle route{inter-area asease-delay}
Parameter Description	<p>inter-area asease-delay: Indicates the inter-area route computation delay. The unit is ms. The value ranges from 0 to 600,000.</p> <p>asease-delay: Indicates the external route computation delay. The unit is ms. The value ranges from 0 to 600,000.</p>
Command Mode	OSPF routing process configuration mode
Usage Guide	<p>If a strict requirement is raised for the network convergence time, use the default value.</p> <p>If a lot of inter-area or external routes exist on the network and the network is not stable, adjust the delays and optimize route computation to reduce the load on the device.</p>

Configuration Example

Configuring the Hello Interval and Dead Interval

<p>Scenario Figure 5-23</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Enable IPv6 on interfaces of all routers. (Omitted) Configure the OSPF basic functions on all routers. (Omitted) Configure the Hello interval and dead interval on all routers.
<p>A</p>	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ipv6 ospf hello-interval 15 A(config-if-GigabitEthernet 0/1)# ipv6 ospf dead-interval 50</pre>
<p>B</p>	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ipv6 ospf hello-interval 15 A(config-if-GigabitEthernet 0/1)# ipv6 ospf dead-interval 50</pre>
<p>Verification</p>	<ul style="list-style-type: none"> Check the interface parameters on Router A and Router B. Verify that the Hello interval is 10s and the dead interval is 50s. On Router A and Router B, verify that the OSPF neighbor information is correct.
<p>A</p>	<pre>A# show ipv6 ospf interface GigabitEthernet 0/1 is up, line protocol is up Interface ID 2 IPv6 Prefixes fe80::2d0:f8ff:fe22:3346/64 (Link-Local Address) 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 DR, Priority 1 Timer interval configured, Hello 15, Dead 50, Wait 40, Retransmit 10</pre>

```
Hello due in 00:00:06
Neighbor Count is 1, Adjacent neighbor count is 1
Hello received 40 sent 40, DD received 17 sent 9
LS-Req received 1 sent 3, LS-Upd received 6 sent 5
LS-Ack received 3 sent 4, Discarded 1
```

```
A# show ipv6 ospf neighbor
```

```
OSPFv3 Process (1), 1 Neighbors, 1 is Full:
```

Neighbor ID	Pri	State	Dead Time	Instance ID	Interface
2.2.2.21	Full/BDR	00:00:30	0		GigabitEthernet 0/1

B

```
B# show ipv6 ospf interface
```

```
GigabitEthernet 0/1 is up, line protocol is up
```

```
Interface ID 2
```

```
IPv6 Prefixes
```

```
fe80::2d0:f8ff:fe22:3446/64 (Link-Local Address)
```

```
OSPFv3 Process (1), Area 0.0.0.0, Instance ID 0
```

```
Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
```

```
Transmit Delay is 1 sec, State BDR, Priority 1
```

```
Timer interval configured, Hello 15, Dead 50, Wait 40, Retransmit 10
```

```
Hello due in 00:00:06
```

```
Neighbor Count is 1, Adjacent neighbor count is 1
```

```
Hello received 40 sent 40, DD received 17 sent 9
```

```
LS-Req received 1 sent 3, LS-Upd received 6 sent 5
```

```
LS-Ack received 3 sent 4, Discarded 1
```

```
B# show ipv6 ospf neighbor
```

```
OSPFv3 Process (1), 1 Neighbors, 1 is Full:
```

Neighbor ID	Pri	State	Dead Time	Instance ID	Interface
1.1.1.11	Full/DR	00:00:38	0		GigabitEthernet 0/1

Common Errors

- The configured neighbor dead time is shorter than the Hello interval.

5.5 Monitoring

Clearing

⚠ Running the **clear** commands may lose vital information and thus interrupt services.

Description	Command
Clears and resets an OSPF process.	clear ipv6 ospf [<i>process-id</i>] process

Displaying

Description	Command
Displays the OSPF process configurations.	show ipv6 ospf [<i>process-id</i>]
Displays information about the OSPF LSDB.	show ipv6 ospf[<i>process-id</i>] database[<i>lsa-type</i> [<i>adv-router</i><i>router-id</i>]]
Displays OSPF-enabled interfaces.	show ipv6 ospf [<i>process-id</i>] interface [<i>interface-type interface-number</i> brief]
Displays the OSPF neighbor list.	show ipv6 ospf[<i>process-id</i>] neighbor[<i>interface-type interface-number</i>[detail] <i>neighbor-id</i> detail]
Displays the OSPF routing table.	show ipv6 ospf [<i>process-id</i>] route[<i>count</i>]
Displays the summarized route of OSPF redistributed routes.	show ipv6 ospf[<i>process-id</i>] summary-prefix
Displays the OSPF network topology information.	show ipv6 ospf[<i>process-id</i>] topology [<i>area</i><i>area-id</i>]

Displays OSPF virtual links.	show ipv6 ospf [<i>process-id</i>] virtual-links
------------------------------	---

Debugging

⚠ System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Debugs OSPF events.	debug ipv6 ospf events [abr asbr os router vlink]
Debugs OSPF interfaces.	debug ipv6 ospf ifsm [events status timers]
Debugs OSPF neighbors.	debug ipv6 ospf n fsm [events status timers]
Debugs the OSPF NSM.	debug ipv6 ospf nsm [interface redistribute route]
Debugs OSPF LSAs.	debug ipv6 ospf lsa [flooding generate install maxage refresh]
Debugs OSPF packets.	debug ipv6 ospf packet [dd detail hello ls-ack ls-request ls-update recv send]
Debugs OSPF routes.	debug ipv6 ospf route [ase ia install spf time]

6 CONFIGURING IS-IS

6.1 Overview

Intermediate System to Intermediate System (IS-IS) is an extensible, robust, and easy-to-use Interior Gateway Protocol (IGP) for route selection and applicable to an IP-ISO CLNS dual environment network (ISO CLNS is short for International Organization for Standardization Connectionless Network Service).

IS-IS has the common characteristics of a link state protocol. It sends Hello packets to discover and maintain neighbor relationships, and sends Link State Protocol Data Units (LSPs) to neighbors to advertise its link state. IS-IS supports Level-1 routing and Level-2 routing. All devices at the same Level maintain the same Link State Database (LSDB), which stores the LSPs generated by the devices to notify each other of the Level's network topology. Each device uses the Dijkstra Shortest Path First (SPF) algorithm to perform best-route calculation, path selection, and fast convergence.

Protocols and

Standards

- RFC1142: OSI IS-IS Intra-domain Routing Protocol
- RFC1195: Use of OSI IS-IS for routing in TCP/IP and dual environments
- RFC3786: Extending the Number of Intermediate System to Intermediate System (IS-IS) Link State PDU (LSP) Fragments Beyond the 256 Limit
- RFC3373: Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
- RFC3358: Optional Checksums in Intermediate System to Intermediate System (ISIS)
- RFC3784: Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
- RFC2763: Dynamic Hostname Exchange Mechanism for IS-IS
- RFC6119(draft-ietf-isis-ipv6-te-00): IPv6 Traffic Engineering in IS-IS
- RFC 2966: Domain-wide Prefix Distribution with Two-Level IS-IS

6.2 Applications

Application	Description
-------------	-------------


Planar Topology	A planar topology is applicable to a small-scale network. At the initial stage of large-scale network construction, core devices are deployed to form an area based on a planar topology.
Hierarchical Topology	A hierarchical topology is applicable to a large-scale network with frequent link flapping.

6.2.1 Planar Topology

Scenario

A planar topology is formed by devices in the same area. See Figure 6-1.

Figure 6-1 Planar Topology



Deployment

- To facilitate future extension and reduce device burden, configure the devices in a planar topology as Level-2 devices.


6.2.2 Hierarchical Topology

Scenario

A hierarchical topology divides the network into the core layer, convergence layer, and access layer. See Figure 6-2.

- Route summarization at the convergence layer is facilitated by address planning.
- When primary and secondary routes exist, devices at the convergence layer leak Level-2 routes to Level-1 areas.

Figure 6-2 Hierarchical Topology



Remarks	Devices at the core layer must be connected consecutively.
---------	---

Deployment

- Design the network topology starting from the core layer.
- Configure devices at the core layer as Level-2 devices.
- Configure devices at the convergence layer as Level-1/Level-2 devices.
- Configure devices at the access layer as Level-1 devices.

6.3 Features

Basic Concepts

End System (ES)

An ES is a non-router device, for example, a host.

Intermediate System (IS)

An IS is a router, which is the basic unit used to transmit routing information and generate routes in IS-IS.

End System to Intermediate System Routing Exchange Protocol (ES-IS)

ES-IS is the protocol used for communication between ESs and ISs in Open System Interconnection (OSI) to dynamically discover Level-2 neighbor relationships.

Domain

A set of ISs in the same routing domain (RD) use the same routing protocol to exchange routing information.

Area

An RD can be divided into multiple areas.

Complete Serial Number PDU (CSNP)

CSNPs are sent by a Designated Intermediate System (DIS) every 10s to synchronize link states in a broadcast network.

Partial Sequence Number PDU (PSNP)

PSNPs are sent by a point-to-point (P2P) link to confirm LSPs, or request LSPs in a broadcast network.

Connectionless Network Protocol (CLNP)

CLNP is an OSI protocol used to transmit data and error messages at the network layer. It is similar to the IP protocol.

Connectionless Network Service (CLNS)

The CLNS is a type of unreliable connection and requires no circuit setup before data transmission.

Designated Intermediate System (DIS)

Similar to a DIS router (DR) in Open Shortest Path First (OSPF), a DIS propagates LSPs to other machines in a Local Area Network (LAN). Neighbor relationships are established not only between DISs and other machines but also between those machines. This characteristic is not possessed by OSPF.

Hello Packet

Hello packets are used to establish and maintain neighbor relationships.

LSP

LSPs describe link states, similar to link-state advertisement (LSA) in OSPF, but the former do not depend on TCP/IP information. LSPs are classified into Level-1 LSPs and Level-2 LSPs, depending on different route types.

Network Selector (NSEL)

An NSEL (sometimes referred to as SEL) specifies the target network-layer protocol service. It is similar to the TCP/UDP port for the Upper Layer Service in the IP protocol. In IS-IS, SEL is typically set to 00 to indicate a device.

Network Service Access Point (NSAP)

An NSAP is the CLNS complete address, including the OSI address and high-layer processes. It consists of an area ID, a system ID, and SEL. When SEL is set to 00, the NSAP is a Network Entity Title (NET), similar to an IP address plus a protocol number.

Sub-Network Point of Attachment (SNPA)

An SNPA provides physical connections and network-layer services. It is similar to a MAC address used in the IP protocol, a Data Link Connection Identifier (DLCI) used by frame relay (FR), or High-Level Data Link Control (HDLC) in a wide area network (WAN).

Level-1 Route

A Level-1 route is an intra-area route that only receives relevant information within the area. To reach other areas, you need to store in Level-1 a default route destined for the closest Level-2.

Level-2 Route

A Level-2 route is an inter-area backbone route. Level-1 and Level-2 cannot be connected directly.

Level-1/Level-2 Route

A Level-1/Level-2 route is a border route connecting a Level-1 route and a Level-2 route. It maintains two databases for the Level-1 and Level-2 routes respectively. It is similar to an area border router (ABR) in OSPF.

Pseudonode

A pseudonode identifies a broadcast subnet (LAN) and allows a broadcast medium to work as a virtual device, which has a route as its interface. The route-pseudonode relationship is managed by a DIS.

Network Entity Title (NET)

A NET is part of an OSI address and describes the area ID and system ID, but it does not define the NSEL, which is contained in the NSAP of the specified system.

Circuit

Circuit is an interface-related term used in IS-IS. Whereas NSAP and NET indicate whole devices, a circuit indicates an interface. The circuit ID of a P2P interface is one byte long. For example, the circuit ID of HDLC is 0x00. In a broadcast network (for example, a LAN), the circuit ID is seven bytes long, including the system ID, for example, 1921.6800.0001.01.

 For details about terms related to IS-IS, see ISO 10589 and RFC1195.

Overview

Feature	Description
IS-IS Network Hierarchy	An IS-IS network is divided into Level-1 and Level-2. The nodes on which devices exchange information in the same area form one Level (Level-1).
IS-IS Address Coding Mode	An IS-IS address is called a NET, which consists of an area ID, a system ID, and an NSAP identifier.

IS-IS Packet Types	There are three types of IS-IS packets: LSP, IS-IS Hello packet (IIH PDU), and serial number packet (SNP) classified into CSNP and PSNP.
DIS Election	A DIS simulates multiple access links as a pseudonode and generates LSPs for the pseudonode. The pseudonode sets up a relationship with each device in the local network and forbids direct communication between the devices.
IS-IS Supported TLV Types	IS-IS supports 21 types of Type-Length-Value (TLV).
LSP Fragment Extension	IS-IS floods LSPs to advertise link states. The size of an LSP is limited by the Maximum Transmission Unit (MTU) size of the link. When the content to be advertised exceeds one LSP, IS-IS will create LSP fragments to carry new link state information.
IS-IS VRF	VPN Routing and Forwarding (VRF) is mainly used for local routing and packet separation. It avoids route conflict caused by use of the same prefix by multiple VPNs.
IS-IS MTR	Multi-Topology Routing (MTR) is mainly used to calculate IPv4 and IPv6 unicast routes in IS-IS based on different topologies.

6.3.1 IS-IS Network Hierarchy


An IS-IS network is divided into Level-1 and Level-2. The nodes on which devices exchange information in the same area form one Level (Level-1).

Working Principle

All devices in an area know the area's network topology and exchange data within the area. A Level-1/Level-2 device is a border device that belongs to different areas and provides inter-area connections. Areas are connected by Level-2 devices. The border devices in various areas form a Level-2 backbone network for inter-area data exchange.

Level-1 devices are only interested in the local area's topology, including all nodes in the local area and the next-hop devices destined for the nodes. Level-1 devices access other areas through Level-2 devices and forward packets from a target network outside of the local area to the closest Level-2 device.

Figure 6-3 IS-IS Network Topology



Related Configuration

Setting the Circuit Type of an IS-IS Interface

By default, **circuit-type** is set to Level-1/Level-2.

Run the **isis circuit-type** command to change the Level of an interface.

If **circuit-type** is set to Level-1 or Level-2-only, IS-IS will only send PDUs of the corresponding Level.

Specifying the IS-IS Level

By default, **is-type** is set to Level-1/Level-2 if no IS-IS instance runs at Level-2 (including Level-1/Level-2). **is-type** is set to Level-1 if there are IS-IS instances running at Level-2 (including Level-1/Level-2).

Run the **is-type** command to specify the Level at which IS-IS will run.


Changing the **is-type** value will enable or disable the routes of a certain Level. A device can have only one instance running at Level-2 (including Level-1/Level-2).

6.3.2 IS-IS Address Coding Mode

An IS-IS address is called a NET, which consists of an area ID, a system ID, and an NSAP identifier, ranging from eight to 20 bytes.

Working Principle

Figure 6-4 NET Address Format



- The area ID identifies the RD length in an area and is fixed relative to the RD. It ranges from one to 13 bytes.
- The system ID is unique in an autonomous system (AS).
- The NSAP is a network selector and sometimes called SEL. In IS-IS, SEL is typically set to 00 to indicate a device.

[Related Configuration](#)

Configuring a NET Address in IS-IS

By default, no NET address is configured in IS-IS.

Run the **net** command to configure a NET address in IS-IS.

The command configures an area ID and a system ID in IS-IS. Different NET addresses must have the same system ID.

6.3.3 IS-IS Packet Types

There are three types of IS-IS packets:

- LSP
- IIH PDU
- SNP (classified into CSNP and PSNP)

[Working Principle](#)

There are three types of IS-IS packets:

- LSP

LSPs are used to transmit link state records within an area and are classified into Level-1 LSPs and Level-2 LSPs. LSPs are only flooded to the corresponding Level.

- IIH PDU

IIH PDUs are used to maintain neighbor relationships. They carry multicast MAC addresses used to determine whether other systems run IS-IS.

- SNP (classified into CSNP and PSNP)

CSNPs are used for LSDB synchronization. By default, a DIS sends a CSNP every 10s in a broadcast network. In a P2P network, a CSNP is sent only after a neighbor relationship is established.

PSNPs are also used for LSDB synchronization.

[Related Configuration](#)

Configuring the LSP Interval on an IS-IS Interface

By default, the LSP interval is 33 ms. If no Level is specified, the interval takes effect for Level-1 and Level-2 LSPs.

Run the **isis lsp-interval** command to configure the LSP interval on an IS-IS interface, in the unit of seconds.

The command changes the LSP interval.

Configuring the Maximum Number of LSP PDUs Sent by IS-IS Interface at a Time

The default value is 5. If no particular level is designated, the default level is set to **level-1-2**, that is, the configuration of the maximum number takes effect at Level-1 and Level-2 concurrently.

Run the **isis lsp-flood** command to set the number of LSP PDUs sent by the IS-IS interface at a time.

You can use this command to change the maximum number of LSP PDUs sent by the interface at a time.

Configuring the Interval for Receiving Duplicate LSP Packets

By default, this command is not configured and the function of setting the interval for receiving LSP packets is not supported. If no particular level is designated, the default level is set to **level-1-2**, that is, the interval configuration takes effect at Level-1 and Level-2 concurrently.

Run the **min-lsp-arrival** command to set the interval for receiving duplicate LSP packets.

During network flapping, LSP packets are frequently duplicated and the LSP packet sequence number constantly increases. If the device processes the duplicate LSP packets, more system resources are consumed. Therefore, setting the interval for receiving duplicate LSP packets can relieve device load.

Configuring the Hello Packet Interval on an IS-IS Interface

By default, the Hello packet interval is 10s for Level-1 and Level-2.

Run the **isis hello-interval** command to configure the Hello packet interval on an IS-IS interface, in the unit of seconds.

The command changes the Hello packet interval. A DIS sends Hello packets at a frequency three times that by non-DIS devices in a broadcast network. If an IS is elected as the DIS on the interface, by default, the interface sends a Hello packet every 3.3s.

Configuring the Minimum PSNP Interval

By default, the minimum PSNP interval is not configured, and the default interval 2s takes effect for Level-1 and Level-2 PSNPs.

Run the **isis psnp-interval** command to configure the minimum PSNP interval, in the unit of seconds.

PSNPs are mainly used to request LSPs that are absent locally or respond to received LSPs (in a P2P network). The PSNP interval should be minimized. If many LSPs exist and the device performance is low, you can increase the PSNP interval and LSP retransmission interval to reduce the device burden.

Configuring the CSNP Broadcast Interval on an IS-IS Interface

By default, CSNPs are sent at 10s intervals in a broadcast network. No CSNPs are sent in a P2P network. When you configure a new CSNP interval without Level-1 or Level-2 specified, the interval takes effect for Level-1 and Level-2 CSNPs.

Run the **isis csnp-interval** command to specify the CSNP broadcast interval on an IS-IS interface, in the unit of seconds.

The command changes the CSNP interval. By default, a DIS sends a CSNP every 10s in a broadcast network. In a P2P network, a CSNP is sent only after a neighbor relationship is established. An interface set to **mesh-groups** can be configured to periodically send CSNPs. No CSNPs are sent if the CSNP interval is set to 0.

6.3.4 DIS Election

A DIS is a designated device in a broadcast network and works like a DR in OSPF.

A pseudonode is generated by a DIS and sets up a relationship with each device in the local network.

Working Principle

A DIS simulates multiple access links as a pseudonode and generates LSPs for the pseudonode. The pseudonode sets up a relationship with each device in the local network and forbids direct communication between the devices. A broadcast subnet and a non-broadcast multiple access (NBMA) network are considered as pseudonodes externally. Non-DIS devices report their link states to the DIS in the same network, and the DIS maintains the link states reported by all ISs in the network. Like DR election in OSPF, a DIS is elected to reduce unnecessary neighbor relationships and route information exchanges.

DIS election in IS-IS is preemptive. The election result can be manually controlled through interface priority configuration. The device with a higher interface priority is more likely to be elected as the DIS.

Related Configuration

Configuring the Priority for DIS Election in a LAN

By default, Priority 64 takes effect for Level-1 and Level-2.

Run the **isis priority** command to configure the priority for DIS election in a LAN.

The command changes the priority carried in Hello packets in a LAN. The device with a lower priority is less likely to be elected as the DIS.

The command is invalid on a P2P network interface. The **no isis priority** command, with or without parameters, restores the priority to its default value. To change the configured priority, run the **isis priority** command with the priority specified to overwrite the existing configuration, or you can first restore the priority to its default value and then configure a new priority.

6.3.5 IS-IS Supported TLV Types

IS-IS supports 26 types of TLV.

Working Principle

The following table lists the IS-IS supported TLV types:

TLV Code	Description
Code = 1	Area ID
Code = 2	Priority of an IS neighbor
Code = 3	ES neighbor
Code = 6	MAC address of an IS neighbor
Code = 8	Filling field
Code = 9	LSP entity
Code = 10	Verification field
Code = 14	Size of the source LSP buffer

Code = 22	Extended IS reachability
Code = 128	IP internal reachability information
Code = 129	Supported protocol
Code = 130	IP external reachability information
Code = 131	Inter-domain routing protocol information
Code = 132	IP interface address
Code = 133	Verification information
Code = 135	Extended IP reachability TLV
Code = 137	Dynamic host name
Code = 211	Graceful Restart (GR)
Code = 222	Multi-Topology (MT) IS reachability
Code = 229	MT TLV
Code = 211	GR
Code=232	IPv6 interface
Code = 235	IPv4 MT IP reachability TLV
Code =236	IPv6 IP reachability TLV
Code = 237	IPv6 MT IP reachability TLV
Code = 240	P2P three-way handshake TLV

[Related Configuration](#)

Configuring the Neighbor Detection Protocol Carried in Hello Packets

By default, neighbor detection is enabled.

Run the **adjacency-check** command to configure the neighbor detection protocol carried in Hello packets.

6.3.6 LSP Fragment Extension

IS-IS floods LSPs to advertise link states. The size of an LSP is limited by the MTU size of the link. When the content to be advertised exceeds one LSP, IS-IS will create LSP fragments to carry new link state information. According to ISO standards, an LSP fragment is identified by a one-byte LSP number. An IS-IS device can generate up to 256 LSP fragments.

Working Principle

The 256 LSP fragments are insufficient in any of the following situations:

1. New applications (such as traffic engineering [TE]) extend new TLV or Sub-TLV.
2. The network is expanded continuously.
3. Routes with reduced granularity are advertised, or other routes are redistributed to IS-IS.

After LSP fragments are used up, new routing information and neighbor information will be discarded, causing network exceptions such as routing black holes or loops. LSP fragments must be extended to carry more link state information, thus ensuring normal network operation.

You can configure an additional system ID and enable fragment extension to allow IS-IS to advertise more link state information in extended LSP fragments. Each virtual system can be considered as a virtual device that establishes a neighbor relationship (with the path value being 0) with the originating system. Extended LSPs are published by the neighbor of the originating system, that is, the virtual system.

The following terms are related to fragment extension:

Normal System ID

The system ID defined by ISO is used to establish neighbor relationships and learn routes. It is further defined as the normal system ID in order to be distinguished from the additional system ID introduced to fragment extension.

Additional System ID

The additional system ID is configured by an administrator to generate extended LSPs. The additional system ID shares the usage rules of the normal system ID (for example, the additional system ID must be unique in the entire area), except that the additional system ID is not carried in Hello packets for neighbor relationship establishment.

Originating System

An originating system is an IS-IS-enabled device and maps a virtual system identified by the additional system ID.

Virtual System (Virtual IS)

A virtual system is identified by the additional system ID and used to generate extended LSPs. The virtual system concept is proposed by RFC for distinguishing from the originating system concept. Each virtual system can generate up to 256 LSP fragments. The administrator can configure multiple additional system IDs (virtual systems) to generate more LSP fragments.

Original LSP

An original LSP is the LSP whose system ID contained in the LSP ID is a normal system ID. Original LSPs are generated by an originating system.

Extended LSP

An extended LSP is the LSP whose system ID contained in the LSP ID is an additional system ID. Extended LSPs are generated by a virtual system.

Related Configuration

Enabling Fragment Extension

By default, fragment extension is disabled. If you do not specify a Level when enabling fragment extension, it will take effect for Level-1 and Level-2 LSPs.

Run the **lsp-fragments-extend** command to enable fragment extension.


There are up to 256 LSP fragments. When the fragments are used up, subsequent link state information, including neighbor information and IP route information, will be discarded, causing a network exception. To solve this problem, enable fragment extension at the specified Level and configure an additional system ID by using the **virtual-system** command.

6.3.7 IS-IS VRF

VRF is mainly used for local routing and packet separation. It avoids route conflict caused by use of the same prefix by multiple VPNs. IPv4 VPN and IPv6 VPN combine Multiprotocol Label Switching (MPLS) advantages in terms of Quality of Service (QoS) and security assurance, and are the primary solutions for interconnecting the geographically different office branches of an enterprise or industry user.

Working Principle

Figure 6-5 Separation of Different VPNs by VRF Tables Configured on Provider Edge (PE) Devices



In Figure 6-5, the following configuration requirements exist: Configure the two sites (CE1 and CE3) in VPN1 to access each other and the two sites (CE2 and CE4) in VPN2 to access each other, and forbid access between the sites in VPN1 and those in VPN2, because VPN1 and VPN2 belong to different customers or departments and may have identical IP addresses.

The customer edge (CE) devices connect the customer network to the PEs to exchange VPN routing information with the PEs, that is, advertise local routes to the PEs and learn remote routes from the PEs.

Each PE learns routes from directly connected CEs and exchanges the learned VPN routes with the other PE through the Border Gateway Protocol (BGP). The PEs provide access to the VPN service.

The Provider (P) device in the Service Provider (SP) network is not directly connected to the CEs. The P device only needs the MPLS forwarding capability and does not maintain VPN information.

The IS-IS protocol running between the PEs and CEs requires the VRF capability to separate routing information between VPN1 and VPN2. That is, IS-IS only learns routes through VRF.

Related Configuration

Binding an IS-IS Instance with a VRF Table

By default, an IS-IS instance is not bound with any VRF table.

Run the **VRF** command to bind an IS-IS instance with a VRF table.

Note the following constraints or conventions for the binding operation:

- The IS-IS instances bound with the same non-default VRF table must be configured with different system IDs. The IS-IS instances bound with different VRF tables can be configured with the same system ID.
- One IS-IS instance can be bound with only one VRF table, but one VRF table can be bound to multiple IS-IS instances.

- When the VRF table bound to an IS-IS instance is changed, all IS-IS interfaces associated with the instance will be deleted. That is, the **ip (or ipv6) router isis [tag]** interface configuration and the redistribution configuration in routing process mode will be deleted.

6.3.8 IS-IS MTR


IS-IS MTR is an extended feature used to separate IPv4 unicast route calculation and IPv6 unicast route calculation based on topologies. It complies with the specification of IS-IS MT extension defined in RFC 5120. New TLV types are introduced to IIH PDUs and LSPs to transmit IPv6 unicast topology information. One physical network can be divided into an IPv4 unicast logical topology and an IPv6 unicast logical topology. The two topologies perform SPF calculation separately and maintain independent IPv4 and IPv6 unicast routing tables respectively. In this way, IPv4 unicast service traffic and IPv6 unicast service traffic are forwarded by different paths. The IS-IS MTR technique helps users deploy IPv6 unicast networks without the constraint on consistency between IPv4 and IPv6 unicast topology information.

IS-IS MTR is derived from IS-IS MT, which is used to separate IPv4 and IPv6 unicast topologies, unicast and multicast topologies, and topologies using different protocol stacks (such as IPv4 and Pv6). IS-IS MTR separates IPv4 and IPv6 unicast topologies based on IS-IS MT.

Working Principle

Figure 6-6 shows a typical networking application. The following implementation requirements exist: Deploy an IPv6 unicast topology in incremental mode, and upgrade some devices to support IPv4 and IPv6 dual protocol stacks while keeping other IPv4-enabled devices unchanged.


Figure 6-6 Physical Topology for IPv4-IPv6 Hybrid Deployment



In Figure 6-6, each link is marked by a number indicating its metric. RB only supports the IPv4 protocol stack, whereas other devices support IPv4 and IPv6 dual protocol stacks.


The networking constraint on consistency between IPv4 and IPv6 unicast topologies must be canceled to retain the use of RB; otherwise, RB cannot establish a neighbor relationship with RA or RD, which will cause new problems.

Figure 6-7 IPv4-IPv6 Hybrid Topology



In Figure 6-7, without IS-IS MTR support, the SPF calculations performed by RA, RB, RC, and RD only take into account the single hybrid topology. The calculated shortest path is RA -> RB -> RD, with the overhead being 20. RB will discard IPv6 packets because it does not support IPv6.

Figure 6-8 Separation of IPv4 and IPv6 Topologies



In Figure 6-8, the IS-IS MTR technique is used to separate IPv4 and IPv6 unicast topologies. RA, RB, RC, and RD establish neighbor relationships based on the IPv4 unicast topology and IPv6 unicast topology respectively. The left part shows the IPv4 topology formed by IPv4-enabled routers. The calculated IPv4 shortest path is RA -> RB -> RC, which realizes IPv4 packet forwarding. The right part shows the IPv6

topology formed by IPv6-enabled routers. The calculated IPv6 shortest path is RA -> RC -> RD, which realizes IPv6 packet forwarding.

IS-IS MTR must be deployed to avoid routing black holes when some devices support only one protocol. IS-IS MTR is not required when all devices support IPv4 and IPv6 dual protocol stacks.

- Deployment of a new network: IS-IS MTR is not required when devices only support the IPv4 protocol stack. For devices that only support the IPv6 protocol stack or devices that support IPv4 and IPv6 dual protocol stacks, enable the MT mode of IS-IS MTR. You are advised not to enable Multi-Topology Transition (MTT); otherwise, loops may occur.
- Reconstruction of an existing network with devices supporting only one protocol stack: Enable the MTT mode of IS-IS MTR on devices that support IPv4 and IPv6 dual protocol stacks in sequence (starting from the device closest to a device supporting only one protocol stack in the network topology). After the MTT mode is enabled on all new devices, switch the MTT mode to the MT mode on these devices in sequence (starting from the device farthest from a device supporting only one protocol stack in the network topology).

Related Configuration

Enabling MTR for IS-IS Instances

By default, IS-IS instances are not enabled with MTR.

Run the **multi-topology** command to configure IS-IS to support IPv6 unicast topologies. After that, IPv4 and IPv6 unicast routes in IS-IS will be calculated based on different topologies.

Note the following constraints or conventions when you use the **multi-topology** command:

1. Set **metric-style** to **Wide** or **Transition** before you run the command.
2. The MTR feature will be disabled if **metric-style** is set to **Narrow** or only one Level is configured to support the Wide or Transition mode.

6.3.9 IS-IS Neighbor

The following conditions must be met for two routing devices to establish a neighbor relationship when IS-IS MTR is not configured:




- The interface addresses on both routing devices are in the same network segment.
- The interface Levels on both routing devices match.
- The routing devices are authenticated by each other.
- The routing devices support the same protocol.

The following conditions must be met for routing devices to establish a neighbor relationship when IS-IS MTR is configured:




- The interface addresses on both routing devices are in the same network segments.

- The interface Levels on both routing devices match.
- The routing devices are authenticated by each other.
- The routing devices have at least one consistent MT ID when P2P links are configured.
- There are no constraints on the MT IDs that the routing devices support when LAN links are configured.

6.4 Configuration

Configuration	Description and Command	
Enabling IS-IS	 (Mandatory) It is used to enable IS-IS on specified interfaces. You need to create an IS-IS routing process in advance.	
	router isis [tag]	Starts an IS-IS routing process. <i>tag</i> indicates the process name.
	net <i>areaAddress.SystemId.00</i>	Configures a NET address in IS-IS.
	ip router isis [tag]	Enables IS-IS on an interface. <i>tag</i> indicates the name of the IS-IS routing process.
Configuring IS-IS Hello Packets	 (Optional) It is used to configure the IS-IS Hello packet holdtime.	
	isis hello-interval { <i>interval</i> minimal } [level-1 level-2]	Configures the Hello packet interval on an interface. The value range is 1 to 65,535, in the unit of seconds.
	isis hello-multiplier <i>multiplier-number</i> [level-1 level-2]	Configures the Hello packet holdtime multiplier on an IS-IS interface. The value range is 2 to 100. The default value is 3.
Configuring IS-IS LSPs	 (Optional) It is used to perform time-related LSP configuration, determine whether to ignore LSP checksum errors, and enable/disable LSP fragment extension.	

	isis lsp-interval <i>interval</i> [level-1 level-2]	Configures the minimum LSP interval on an interface. The value range is 1 to 4,294,967,295, in the unit of milliseconds.
	isis lsp-flood <i>lsp-number</i> [level-1 level-2]	Indicates the maximum number of LSP packets sent by the IS-IS interface at a time. The value range is 1 to 1000 . The default value is 5 .
	isis retransmit-interval <i>interval</i> [level-1 level-2]	Configures the LSP retransmission interval by P2P links on an interface. The value range is 0 to 65,535, in the unit of seconds.
	lsp-gen-interval [level-1 level-2] <i>maximum-interval</i> [<i>initial-interval</i> <i>hold-interval</i>]	Configures LSP generation cycle. <i>maximum-interval</i> : Indicates the maximum interval for generating two consecutive LSP packets. The value range is 1 to 65535 (in seconds). The default value is 5 . <i>initial-interval</i> : Indicates the waiting time for generating an LSP packet for the first time. The value range is 0 to 60000 (in milliseconds). The default value is 50 . <i>hold-interval</i> : Indicates the minimum interval for generating an LSP packet for the second time. The value range is 10 to 60000 (in milliseconds). The default value is 200 .
	ignore-lsp-errors	Configures to ignore LSP checksum errors.
	lsp-fragments-extend [level-1 level-2] [compatible rfc3786]	Enables fragment extension.
	virtual-system <i>system-id</i>	Configures an additional system ID.

Configuring IS-IS SNPs	<p> (Optional) It is used to configure the CSNP broadcast interval.</p>	
	<p>isis csnp-interval <i>interval</i> [level-1 level-2]</p>	<p>Configures the CSNP interval on an interface. The value range is 0 to 65,535, in the unit of seconds. The default value is 10s.</p> <p>No CSNPs are sent if the CSNP interval is set to 0.</p>
Configuring the IS-IS Level Type	<p> (Optional) It is used to configure the system type or interface circuit type in IS-IS.</p>	
	<p>is-type { level-1 level-1-2 level-2-only }</p>	<p>Configures the system type.</p>
	<p>isis circuit-type { level-1 level-1-2 level-2-only [external] }</p>	<p>Configures the interface circuit type.</p>
Configuring IS-IS Authentication	<p> (Optional) It is used to configure interface authentication, area authentication, and RD authentication.</p>	
	<p>isis password [0 7] <i>password</i> [send-only] [level-1 level-2]</p>	<p>Configures the password for plaintext authentication of Hello packets on an interface.</p> <p>When send-only is included, the authentication password is only used to authenticate sent Hello packets. Received Hello packets are not authenticated.</p> <p>If no Level is specified, the configured authentication and password take effect for all Levels.</p> <p>This command does not take effect if the isis authentication mode command is executed. Both commands are used to configure IS-IS interface authentication, but the isis password command has a lower priority. Before you run the isis password command, delete the isis</p>

		authentication mode command configuration.
	isis authentication mode { text md5 } [level-1 level-2]	<p>Specifies authentication as plaintext or MD5.</p> <p>If no Level is specified, the authentication mode takes effect for all Levels.</p> <p>If you use this command after the isis password password [level-1 level-2] command is executed, the previous command configuration will be overwritten. Both commands are used to configure IS-IS interface authentication, but the isis authentication mode command has a higher priority.</p>
	isis authentication key-chain <i>name-of-chain</i> [level-1 level-2]	<p>Configures the password for interface authentication.</p> <p>If no Level is specified, the configured key chain takes effect for all Levels.</p> <p>This command must be used with the isis authentication mode command to configure IS-IS interface authentication.</p>
	isis authentication send-only [level-1 level-2]	<p>(Optional) Specifies that interface authentication is performed only on sent packets. Received packets are not authenticated.</p> <p>If no Level is specified, the send-only authentication mode takes effect for all Levels.</p> <p>This command is used to avoid network flapping caused by a temporary authentication failure when IS-IS authentication is configured. Before you deploy IS-IS authentication in the entire</p>

		<p>network, run the isis authentication mode { text md5 } [level-1 level-2] and isis authentication key-chain name-of-chain [level-1 level-2] commands on each device. After that, run the no isis authentication send-only command to restore the authentication of received packets. This realizes smooth authentication deployment and avoids network flapping.</p>
	<p>area-password [0 7] password [send-only]</p>	<p>Configures the password for area (Level-1) plaintext authentication.</p> <p>When send-only is included, the authentication password is only used to authenticate sent packets. Received packets are not authenticated.</p> <p>This command does not take effect if the authentication mode command is executed. Both commands are used to configure IS-IS area authentication, but the area-password command has a lower priority. Before you run the area-password command, delete the authentication mode command configuration.</p>
	<p>authentication mode { text md5 } level-1</p>	<p>Specifies the IS-IS area authentication mode.</p> <p>If you use this command after the area-password password command is executed, the previous command configuration will be overwritten. Both commands are used to configure IS-IS area authentication, but the authentication mode command has a higher priority.</p>

	authentication key-chain <i>name-of-chain</i> level-1	<p>Configures the key chain for IS-IS area authentication.</p> <p>This command must be used with the authentication mode command to configure IS-IS area authentication.</p>
	authentication send-only level-1	<p>(Optional) Specifies that IS-IS area authentication is performed only on sent packets. Received packets are not authenticated.</p> <p>This command is used to avoid network flapping caused by a temporary authentication failure when IS-IS authentication is configured. Before you deploy IS-IS authentication in the entire area, run the authentication mode { text md5 } level-1 and authentication key-chain <i>name-of-chain</i> level-1 commands on each device. After that, run the no authentication send-only command to restore the authentication of received packets. This realizes smooth authentication deployment and avoids network flapping.</p>
	domain-password [0 7] <i>password</i> [send-only]	<p>Configures the password for RD (Level-2) plaintext authentication.</p> <p>When send-only is included, the authentication password is only used to authenticate sent packets. Received packets are not authenticated.</p> <p>This command does not take effect if the authentication mode command is executed. Both commands are used to configure IS-IS RD authentication, but the domain-password command has a</p>

		<p>lower priority. Before you run the domain-password command, delete the authentication mode command configuration.</p>
	<p>authentication mode { text md5 } level-2</p>	<p>Specifies the IS-IS RD authentication mode.</p> <p>If you use this command after the domain-password <i>password</i> command is executed, the previous command configuration will be overwritten. Both commands are used to configure IS-IS RD authentication, but the authentication mode command has a higher priority.</p>
	<p>authentication key-chain name-of-chain level-2</p>	<p>Configures the password for IS-IS RD authentication.</p> <p>This command must be used with the authentication mode command to configure IS-IS RD authentication.</p>
	<p>authentication send-only level-2</p>	<p>(Optional) Specifies that IS-IS RD authentication is performed only on sent packets. Received packets are not authenticated.</p> <p>This command is used to avoid network flapping caused by a temporary authentication failure when IS-IS authentication is configured. Before you deploy IS-IS authentication in the entire RD, run the authentication mode { text md5 } level-2 and authentication key-chain name-of-chain level-2 commands on each device. After that, run the no authentication send-only command to restore the authentication of received packets.</p>

		This realizes smooth authentication deployment and avoids network flapping.
Configuring IS-IS GR	⚠ (Optional) It is used to enable IS-IS GR.	
	graceful-restart	Enables the GR Restart capability on the device that works as a Restarter. By default, the GR Restart capability is enabled.
	graceful-restart grace-period <i>seconds</i>	(Optional) Configures the IS-IS GR time on the device that works as a Restarter. The default value is 300s.
	no graceful-restart helper disable	Enables the IS-IS GR Help capability on the device that works as a Helper. By default, the GR Help capability is enabled.
Configuring IS-IS NSR	⚠ (Optional) It is used to enable IS-IS NSR.	
	nsr	Enables the IS-IS NSR. By default, it is disabled.
Configuring BFD Support for IS-IS	⚠ (Optional) It is used to enable BFD support for IS-IS.	
	bfd all-interfaces [anti-congestion]	Enables BFD support for IS-IS on all interfaces.
	isis bfd [disable anti-congestion]	Enables or disables BFD support for IS-IS on the current interface.
Setting the IS-IS Overload Bit	⚠ (Optional) It is used to set the overload bit in LSPs.	
	set-overload-bit [on-startup { <i>seconds</i> / wait-for-bgp [<i>bgpseconds</i>] }] [suppress { [interlevel] [external] }] [level-1 level-2]	Sets the overload bit.

Configuring IS-IS VRF	⚠ (Optional) It is used to bind an IS-IS instance with a VRF table.	
	vrf <i>vrf-name</i>	Binds an IS-IS instance with a VRF table.
Configuring IS-IS MTR	⚠ (Optional) It is used to calculate IPv4 and IPv6 unicast routes in IS-IS based on different topologies.	
	multi-topology [transition]	Configures IS-IS to support IPv6 unicast topologies.
Configuring Simple Network Management Protocol (SNMP) for IS-IS	⚠ (Optional) It is used to allow the SNMP software to perform Management Information Base (MIB) operations on IS-IS instances.	
	enable mib-binding	Performs MIB operations on the instance bound with Tag 1.
	configure terminal	Enters global configuration mode.
	snmp-server enable traps isis	Enables IS-IS trap globally.
	snmp-server host { <i>host-addr</i> ipv6 <i>ipv6-addr</i> } [vrf <i>vrfname</i>] [traps] [version { 1 2c 3 { auth noauth priv }] <i>community-string</i> [udp-port <i>port-num</i>]	Configures an SNMP host in global configuration mode to receive IS-IS trap messages.
	router isis	Enters IS-IS routing process configuration mode.
enable traps all	Allows the sending of all IS-IS trap messages to the host with the IP address 10.1.1.1.	
Configuring IS-IS Two-way Maintenance	⚠ Optional.	
	two-way-maintain	Enables IS-IS two-way maintenance.
Configuring Other IS-IS Parameters	⚠ Optional.	
	maximum-paths <i>maximum</i>	Configures the maximum number of IS-IS IPv4/IPv6 equal-cost paths.

isp-length <i>receive size</i>	Configures the maximum length allowed for received LSPs.
isp-length originate <i>size [level-1 level-2]</i>	Configures the maximum length allowed for sent LSPs.
passive-interface [default] { <i>interface-type interface-number</i> }	Configures a passive interface.
bandwidth-reference <i>value</i>	Configures the bandwidth reference for IS-IS to automatically calculate the link metric.
interfaces-protocol-compatible	Configures the TLV field of the IS-IS protocol to be padded based on protocols supported by an interface.
is-name <i>name</i>	Configures the instance name.
isis metric <i>metric [level-1 level-2]</i>	Configures the interface metric, which is valid only when metric-style is set to Narrow .
isis wide-metric <i>metric [level-1 level-2]</i>	Configures the interface wide-metric value, which is valid only when metric-style is set to Wide .
Isis passive	Configures the passive interface.
isis priority <i>value [level-1 level-2]</i>	Configures the priority for DIS election on an interface.
default-information originate [route-map <i>map-name</i>]	Generates a Level-2 default route, which will be advertised through LSPs. When the command includes the route-map option, a default route is generated only if the criteria in route-map are met.
summary-address <i>ip-address net-mask [level-1 level-2 level-1-2] [metric number]</i>	Configures an IPv4 summary route.
summary-prefix <i>ipv6-prefix/prefix-length [level-1 level-2 level-1-2]</i>	Configures an IPv6 summary route.

	ignore-lsp-errors	Configures to ignore LSP checksum errors.
	log-adjacency-changes	Activates logging of IS-IS neighbor relationship changes.
	redistribute	Configures route redistribution.

6.4.1 Enabling IS-IS

Configuration

Effect

- Before you run IS-IS, create an IS-IS routing process in global configuration mode. You can set the **tag** parameter after the **router isis** command to name the process. You can add different tags to configure different IS-IS routing processes. The setting of the **tag** parameter is optional.
- A system ID uniquely identifies an IS in a routing AS; therefore, the system ID must be unique across the AS. In IS-IS, each area may contain one or multiple area IDs. Normally, you only need to configure one area ID. You can configure multiple area IDs to realize area division. If an IS is configured with multiple area IDs, the system IDs must be the same.
- After an interface is added to the specified IS-IS routing process, the interface will establish a neighbor relationship.

Notes

- The Level-1 IS devices in an area must be configured with the same area ID.
- The core routing table does not distinguish the routing entries generated by different IS-IS routing processes.
- The IP addresses of interfaces connected between neighbors must be in the same network segment.
- If the two IP addresses are in different network segments, a neighbor relationship cannot be established.
- If you need to add an interface to the specified IS-IS routing process, set the **tag** parameter after the **ip router isis** command to indicate the process name.
- If you run the **no ip routing** command in global configuration mode, IS-IS will disable IPv4 routing on all interfaces. That is, the **no ip router isis [tag]** command is automatically executed on all interfaces. Other IS-IS settings remain unchanged.
- By default, CPU protection is enabled on devices. For packets mapped to the destination group addresses (AllISSystems, AllL1ISSystems, and AllL2ISSystems) in IS-IS, there is a default limit (for example, 400 pps) on the number of packets sent to the CPU. If a device has many neighbor

relationships or sends Hello packets at short intervals, the IS-IS packets that the device receives may exceed the default limit, causing frequent flapping of neighbor relationships. To solve the problem, you can use the CPU protection command in global configuration mode to increase the limit.

Configuration

Steps

Starting an IS-IS Routing Process

- Mandatory.
- Perform this configuration in global configuration mode on each device, unless otherwise specified.

Configuring a NET Address in IS-IS

- Mandatory.
- Perform this configuration in IS-IS routing process configuration mode on each device, unless otherwise specified.

Enabling IS-IS on an Interfaces

- Mandatory.
- Perform this configuration in interface configuration mode on each device, unless otherwise specified.

Verification

- Check whether devices send Hello packets.
- Check whether devices establish neighbor relationships.
- Check whether devices exchange LSPs.

Related

Commands

Starting an IS-IS Routing Process

Command	router isis [<i>tag</i>]
Parameter Description	<i>tag</i> : Indicates the name of an IS-IS instance.
Command Mode	Global configuration mode
Usage Guide	Use this command to initialize an IS-IS instance and enter IS-IS routing process configuration mode. An IS-IS instance will start running after a NET address is configured.

If you set the **tag** parameter when you start an IS-IS routing process, you need to add the **tag** parameter when closing the IS-IS routing process.

By default, CPU protection is enabled on devices. For packets mapped to the destination group addresses (AllISSystems, AllL1ISSystems, and AllL2ISSystems) in IS-IS, there is a default limit (for example, 400 pps) on the number of packets sent to the CPU. If a device has many neighbor relationships or sends Hello packets at short intervals, the IS-IS packets that the device receives may exceed the default limit, causing frequent flapping of neighbor relationships. To solve the problem, you can use the CPU protection command in global configuration mode to increase the limit.

Configuring a NET Address in IS-IS

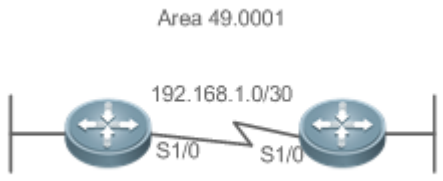
Command	net <i>net-address</i>
Parameter Description	<i>net-address</i> : The NET address is in the format of XX.XXXX.YYYY.YYYY.YYYY.00. XX.XXXX indicates the area ID, and YYYY.YYYY.YYYY indicates the system ID.
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to configure an area ID and a system ID in IS-IS. Different NET addresses must have the same system ID.

Enabling IS-IS on an Interface

Command	ip router isis [<i>tag</i>]
Parameter Description	<i>tag</i> : Indicates the name of an IS-IS instance.
Command Mode	Interface configuration mode
Usage Guide	Use this command to enable an interface to participate in IS-IS IPv4 routing. Use the no form of this command to disable the IS-IS routing process on the interface. If you run the no ip routing command in global configuration mode, IS-IS will disable IPv4 routing on all interfaces. That is, the no ip router isis [<i>tag</i>] command is automatically executed on all interfaces. Other IS-IS settings remain unchanged.

Configuration Example

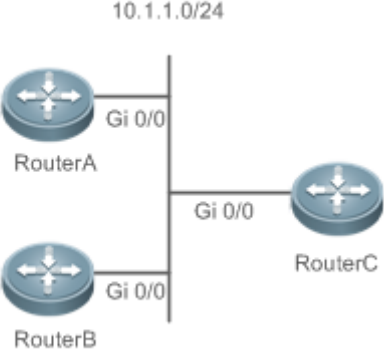
Establishing a Neighbor Relationship on an IS-IS P2P Link

Scenario	Router A and Router B are connected in P2P mode.
Figure 6-9 P2P Link Topology	
Configuration Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS. ▪ Configure Wide Area Network (WAN) interfaces.
A	<pre>A(config)# router isis A(config-router)# net 49.0001.0000.0000.0001.00 A(config)# interface Serial 1/0 A(config-if)# ip address 192.168.1.1 255.255.255.252 A(config-if)# ip router isis</pre>
B	<pre>B(config)# router isis B(config-router)# net 49.0001.0000.0000.0002.00 B(config)# interface Serial 1/0 B(config-if)# ip address 192.168.1.2 255.255.255.252 B(config-if)# ip router isis</pre>
Verification	<ul style="list-style-type: none"> ▪ Enable sending of Hello packets from the interface 192.168.1.1 on Router A to the interface 192.168.1.2 on Router B. ▪ Establish an IS-IS neighbor relationship between Router A and Router B, with the neighbor state being Up. ▪ Check the LSPs on Router A and Router B. The system IDs 0000.0000.0001 and 0000.0000.0002 should exist.
A	<pre>A# show isis neighbors A# show isis database detail</pre>

B

```
B# show isis neighbors
```

Establishing a Neighbor Relationship on an IS-IS Broadcast Link

Scenario	Router A, Router B, and Router C are interconnected through the Ethernet.
Figure 6-10 IS-IS Broadcast Link Topology	 <p>The diagram illustrates a network topology where three routers (Router A, Router B, and Router C) are interconnected through a shared Ethernet link. The link is labeled with the IP address range 10.1.1.0/24. Router A and Router B are positioned on the left side of the link, while Router C is on the right. Each router is connected to the Ethernet link via its GigabitEthernet 0/0 interface.</p>
Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS.▪ Configure Ethernet interfaces.
A	<pre>A(config)# router isis A(config-router)# net 49.0001.0000.0000.0001.00 A(config)# interface GigabitEthernet 0/0 A(config-if)# ip address 10.1.1.1 255.255.255.0 A(config-if)# ip router isis</pre>
B	<pre>B(config)# router isis B(config-router)# net 49.0001.0000.0000.0002.00 B(config)# interface GigabitEthernet 0/0 B(config-if)# ip address 10.1.1.2 255.255.255.0 B(config-if)# ip router isis</pre>
C	<pre>C(config)# router isis C(config-router)# net 49.0001.0000.0000.0003.00 C(config)# interface GigabitEthernet 0/0 C(config-if)# ip address 10.1.1.3 255.255.255.0 C(config-if)# ip router isis</pre>

Verification	<p>Enable sending of Hello packets from the interface 10.1.1.1 on Router A to the interface 10.1.1.2 on Router B and the interface 10.1.1.3 on Router C.</p> <ul style="list-style-type: none"> Establish IS-IS neighbor relationships between Router A and Router B and between Router A and Router C, with the neighbor state being Up. Check the LSPs on Router A, Router B, and Router C. The system IDs 0000.0000.0001, 0000.0000.0002, and 0000.0000.0003 should exist.
A	<pre>A# show isis neighbors A# show isis database detail</pre>
B	<pre>B# show isis neighbors</pre>
C	<pre>C# show isis neighbors</pre>

Performing Simple IS-ISv6 Configuration

Scenario	Router A and Router B are connected through the Ethernet.
Figure 6-11 IS-ISv6 Broadcast Link Topology	<p>The diagram shows two routers, Router A and Router B, connected by a line representing an Ethernet link. Both routers have their GigabitEthernet 0/0 interfaces connected to this link. Above the link, the IPv6 address range 1000::/112 is specified.</p>
Configuration Steps	<ul style="list-style-type: none"> Configure IS-IS. Configure Ethernet interfaces.
A	<pre>A(config)# router isis A(config-router)# net 49.0001.0000.0000.0001.00 A(config)# interface GigabitEthernet 0/0 A(config-if)# ipv6 address 1000 ::1/112 A(config-if)# ipv6 router isis</pre>
B	<pre>B(config)# router isis B(config-router)# net 49.0001.0000.0000.0002.00 B(config)# interface GigabitEthernet 0/0</pre>

	<pre>B(config-if)# ipv6 address 1000 ::2/112 B(config-if)# ipv6 router isis</pre>
Verification	<p>Enable sending of Hello packets from the interface 1000 ::1 on Router A to the interface 1000 ::2 on Router B.</p> <p>Establish an IS-IS neighbor relationship between Router A and Router B, with the neighbor state being Up.</p> <p>Check the LSPs on Router A and Router B. The system IDs 0000.0000.0001 and 0000.0000.0002 should exist.</p>
A	<pre>A# show isis neighbors A# show isis database detail</pre>
B	<pre>B# show isis neighbors</pre>

Common

Errors

- The IP addresses of the interfaces connected between neighbors are not in the same network segment.
- The **ip router isis** command is not executed on interfaces.
- No NET address is configured, or different NET addresses exist at Level-1.
- **max-area-addresses** is configured differently on both sides.
- **metric-style** is configured differently on both sides.
- The interface Levels on both sides are different. One side is Level-1, whereas the other side is Level-2.
- One side is configured with the P2P mode, whereas the other side is configured with the broadcast mode.
- One side is enabled with authentication, whereas the other side is not.

6.4.2 Configuring IS-IS Hello Packets

Configuration

Effect

- Configure the Hello packet interval on an interface. The value range is 1 to 65,535, in the unit of seconds.

- Configure the Hello packet holdtime multiplier on an IS-IS interface.

Notes

- You can change the Hello packet holdtime by using the **isis hello-multiplier** command or **isis hello-interval** command or both.
- By default, CPU protection is enabled on devices. For packets mapped to the destination group addresses (AllISSystems, AllL1ISSystems, and AllL2ISSystems) in IS-IS, there is a default limit (for example, 400 pps) on the number of packets sent to the CPU. If a device has many neighbor relationships or sends Hello packets at short interval, the IS-IS packets that the device receives may exceed the default limit, causing frequent flapping of neighbor relationships. To solve the problem, you can use the CPU protection command in global mode to increase the limit.

Configuration

Steps

Configuring the Hello Packet Interval on an Interface

- Perform this configuration based on requirements.
- Run the **isis hello-interval** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring the Hello Packet Holdtime Multiplier on an Interface

- Perform this configuration based on requirements.
- Run the **isis hello-multiplier** command in interface configuration mode on the desired device, unless otherwise specified.

Verification

- Enable Router A to send Hello packets to Router B and Router C, and capture packets to check the packet interval.
- Make Router B or Router C down. After the holdtime has elapsed, check whether the corresponding neighbor relationship on Router A is invalid.

Related

Commands

Configuring the Hello Packet Interval on an Interface

Command	isis hello-interval { <i>interval</i> / <i>minimal</i> } [<i>level-1</i> <i>level-2</i>]
Parameter Description	<p><i>interval</i>: Indicates the Hello packet interval. The value range is 1 to 65,535, in the unit of seconds. The default value is 10.</p> <p>minimal: Indicates the minimum value of the holdtime, which is 1.</p>

	<p>level-1: Applies the setting to Level-1 Hello packets.</p> <p>level-2: Applies the setting to Level-2 Hello packets.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>Use this command to change the Hello packet interval. The default interval is 10s. A DIS sends Hello packets at a frequency three times that by non-DIS devices in a broadcast network. If an IS is elected as the DIS on the interface, by default, the interface sends a Hello packet every 3.3s.</p> <p>If the keyword minimal is used, the Hello packet holdtime is set to 1. The Hello packet interval will be calculated based on the holdtime multiplier. If the holdtime multiplier is set to 4 and the isis hello-interval minimal command is executed, the Hello packet interval is equal to 1s divided by 4. The default Hello packet holdtime multiplier on an IS-IS interface is 3. The holdtime is equal to the holdtime multiplier multiplied by the packet interval. If the keyword minimal is used, the holdtime is set to 1. The packet interval is equal to 1 divided by the holdtime multiplier. If the holdtime multiplier is set to 4 and the isis hello-interval minimal command is executed, the packet interval is equal to 1 divided by 4s, which is 250 ms.</p>

Configuring Hello Packet Holdtime Multiplier on an Interface

Command	isis hello-multiplier <i>multiplier-number</i> [level-1 level-2]
Parameter Description	<i>multiplier-number</i> : Indicates the Hello packet holdtime multiplier. The value range is 2 to 100. The default value is 3.
Command Mode	Interface configuration mode
Usage Guide	The Hello packet holdtime is equal to the Hello packet interval multiplied by the holdtime multiplier.

Configuration

Example

Configuring the Hello Packet Interval and Holdtime on an IS-IS Interface

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure the Hello packet interval on an IS-IS interface. ▪ Configure the Hello packet holdtime multiplier on an IS-IS interface.
-----------------------------	--

	<pre>A(config)# interface GigabitEthernet 0/0 A(config-if)# isis hello-interval 5 A(config-if)# isis hello-multiplier 5</pre>
Verification	<p>Enable Router A to send Hello packets to Router B and Router C, and capture packets to check the packet interval.</p> <p>Make Router B or Router C down. After the holdtime has elapsed, check whether the corresponding neighbor relationship on Router A is invalid.</p>
	<pre>A# show isis neighbor</pre>

6.4.3 Configuring IS-IS LSPs

Configuration

Effect

- **isis lsp-interval:** Configures the LSP interval on an IS-IS interface.
- **isis lsp-flood:** Sets the number of LSP PDUs sent by the IS-IS interface at a time.
- **lsp-gen-interval:** Configures the minimum LSP generational interval. The LSP generational interval is the interval between the new-version LSP generation time and the old-version LSP generation time.
- **isis retransmit-interval:** After a device at one end of a P2P link sends an LSP packet, if the device receives no response within a period of time, it determines that the LSP packet is lost or dropped due to an error. The device will resend the LSP packet.
- **lsp-gen-interval:** Indicates the exponential backoff algorithm of LSP packet generation. Any update to related information forming the LSP packet leads to LSP packet generation. During network flapping, LSP packets are frequently generated, which increases system resource consumption. An appropriate value can be set by running the **lsp-gen-interval** command. In this way, LSP packets are generated and advertised in time when the network is stable. When the network becomes unstable, less LSP packets are generated as the flapping continues, reducing the device consumption.
- **lsp-refresh-interval:** All current LSPs are periodically retransmitted to enable each network node to maintain the latest LSPs. The retransmission period is called the LSP refresh interval, which aims to update and synchronize LSPs in the entire area.
- **max-lsp-lifetime:** An LSP contains a field to indicate its lifetime. When a device generates an LSP, the field is set to the maximum lifetime of the LSP. After the LSP is received by the peer device, its lifetime will decrease with time. The peer device will replace the old LSP with the newly received one. If the device receives no new LSP until the existing LSP's lifetime decreases to 0, the existing LSP is still maintained in the LSDB for another 60s. If the device still receives no new LSP during this

period, the existing LSP will be deleted from the LSDB. This mechanism updates and synchronizes LSPs in the entire area.

- **min-lsp-arrival:** During network flapping, LSP packets are frequently duplicated and the LSP packet sequence number constantly increases. As a result, the device keeps receiving duplicate LSP packets with different sequence numbers. If the device processes the duplicate LSP packets, more system resources are consumed. An appropriate value can be set by running the **min-lsp-arrival** command. In this way, LSPs are received and processed in time when the network is stable. When the network becomes unstable, less duplicate LSP packets are received and processed as the flapping continues, reducing the device consumption.
- **ignore-lsp-errors:** After receiving an LSP, the local IS-IS neighbor calculates its checksum and compares it with the checksum contained in the LSP. By default, if the two checksums are inconsistent, the LSP will be discarded. If you run the **ignore-lsp-errors** command to configure to ignore checksum errors, the LSP will be processed normally despite checksum inconsistency.
- **lsp-fragments-extend:** Enables LSP fragment extension, which is used to generate an extended LSP when the 256 fragments of the original LSP are used up.

Notes

- The LSP refresh interval must be smaller than the maximum LSP lifetime.
- The maximum LSP lifetime must be greater than the LSP refresh interval.
- The value of **initial-interval** cannot be greater than that of **maximum-interval**. Otherwise, the value of **initial-interval** will be used as the value of **maximum-interval**.
- The value of **hold-interval** cannot be greater than that of **maximum-interval**. Otherwise, the value of **hold-interval** will be used as the value of **maximum-interval**.
- The value of **initial-interval** cannot be greater than that of **hold-interval**. Otherwise, the value of **initial-interval** will be used as the value of **hold-interval**.

Configuration

Steps

Configuring the Minimum LSP Interval

- Perform this configuration based on requirements.
- Run the **isis lsp-interval** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring the Maximum Number of LSP Packets Sent at a Time

- Perform this configuration based on requirements.
- Run the **isis lsp-flood** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring the LSP Retransmission Interval

- Perform this configuration based on requirements.
-

- Run the **isis retransmit-interval** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring LSP Packet Generation Cycle

- Perform this configuration based on requirements.
- Run the **lsp-gen-interval** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring the LSP Refresh Interval

- Perform this configuration based on requirements.
- Run the **lsp-refresh-interval** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the LSP Lifetime

- Perform this configuration based on requirements.
- Run the **max-lsp-lifetime** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the Interval for Receiving Duplicate LSP Packets

- Perform this configuration based on requirements.
- Run the **min-lsp-arrival** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring to Ignore LSP Checksum Errors

- Perform this configuration based on requirements.
- Run the **ignore-lsp-errors** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring LSP Fragment Extension

- Perform this configuration based on requirements.
- Run the **lsp-fragments-extend** and **virtual-system** commands in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- Update LSPs continuously and capture LSPs to check the minimum LSP interval.
- Disable neighboring routes and capture LSPs to check the LSP retransmission interval.
- Capture LSPs to check the refresh interval.
- Check the LSP lifetime.
- Send an LSP with an incorrect checksum and check whether the LSP is discarded.
- Reduce the **lsp-length originate** command value, add routing information, and capture LSPs to check whether more than 256 LSP fragments are generated.

Related Commands

Configuring the Minimum LSP Interval

Command	isis lsp-interval <i>interval</i> [level-1 level-2]
Parameter Description	<i>milliseconds</i> : Indicates the LSP interval. The value range is 1 to 4,294,967,295, in the unit of milliseconds. level-1 : Applies the setting only to Level-1 LSPs. level-2 : Applies the setting only to Level-2 LSPs.
Command Mode	Interface configuration mode
Usage Guide	N/A

Configuring the Maximum Number of LSP Packets Sent at a Time

Command	isis lsp-flood <i>lsp-number</i> [level-1 level-2]
Parameter Description	<i>lsp-number</i> : Indicates the maximum number of LSP packets sent by the IS-IS interface at a time. The value range is 1 to 1000 . The default value is 5 . level-1 : Applies the configuration only to Level-1. level-2 : Applies the configuration only to Level-2.
Command Mode	Interface configuration mode
Usage Guide	N/A

Configuring LSP Packet Generation Cycle

Command	lsp-gen-interval [level-1 level-2] <i>maximum-interval</i> [<i>initial-interval hold-interval</i>]
Parameter Description	level-1 : Applies the configuration only to Level-1. level-2 : Applies the configuration only to Level-2.

	<p><i>maximum-interval</i>: Indicates the maximum interval for generating two consecutive LSP packets. The value range is 1 to 65535 (in seconds). The default value is 5.</p> <p><i>initial-interval</i>: Indicates the waiting time for generating an LSP packet for the first time. The value range is 0 to 60000 (in milliseconds). The default value is 50.</p> <p><i>hold-interval</i>: Indicates the minimum interval for generating an LSP packet for the second time. The value range is 10 to 60000 (in milliseconds). The default value is 200.</p>
Configuration Mode	IS-IS routing process configuration mode
Usage Guide	<p>The LSP packet generation interval refers to the interval for generating two different LSP packets. A smaller generation interval indicates faster network convergence, which, however, will be accompanied by frequent flooding on the network.</p> <p>The waiting time for generating an LSP packet for the first time is the initial interval. If the network becomes unstable, the LSP packet regeneration interval is changed to be less than the maximum interval, and the interval for generating an LSP packet for the second time becomes the hold interval. A corresponding penalty will be added to this interval: The next interval for regenerating a LSP packet doubles the previous interval for generating the same LSP packet, until the regeneration interval reaches the maximum interval. Subsequent LSP packets will be generated at the maximum interval. When the network becomes stable, the LSP packet regeneration interval becomes greater than the maximum interval, and the waiting time for LSP packet generation is restored to the initial interval.</p> <p>Link changes have high requirements for convergence. The initial interval can be set to a small value. The preceding parameters can also be adjusted to larger values to reduce CPU consumption.</p> <p>The value of initial-interval cannot be greater than that of maximum-interval. Otherwise, the value of initial-interval will be used as the value of maximum-interval.</p> <p>The value of hold-interval cannot be greater than that of maximum-interval. Otherwise, the value of hold-interval will be used as the value of maximum-interval.</p> <p>The value of initial-interval cannot be greater than that of hold-interval. Otherwise, the value of initial-interval will be used as the value of hold-interval.</p>

Configuring the LSP Retransmission Interval

Command	isis retransmit-interval <i>interval</i> [level-1 level-2]
Parameter Description	seconds : Indicates the LSP retransmission interval. The value range is 0 to 65,535, in the unit of seconds.

	<p>level-1: Applies the setting only to Level-1 LSPs.</p> <p>level-2: Applies the setting only to Level-2 LSPs.</p>
Command Mode	Interface configuration mode
Usage Guide	Use this command to configure the LSP retransmission interval. In a P2P network, after a device sends an LSP, if the device receives no PSNP response within the time specified by this command, it will resend the LSP. If the retransmission interval is set to 0, the LSP will not be resent,

Configuring the LSP Refresh Interval

Command	<code>lsp-refresh-interval interval</code>
Parameter Description	<i>interval</i> : Indicates the LSP refresh interval. The value range is 1 to 65,535, in the unit of seconds. The default value is 900.
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>After an LSP has remained stable for a period specified by this command, it will be refreshed and updated before being published.</p> <p>The LSP refresh interval must be smaller than the maximum LSP lifetime.</p>

Configuring the LSP Lifetime

Command	<code>max-lsp-lifetime value</code>
Parameter Description	<i>value</i> : Indicates the maximum time that LSPs keep alive. The value range is 1 to 65,535, in the unit of seconds. The default value is 1,200.
Command Mode	IS-IS routing process configuration mode
Usage Guide	The maximum LSP lifetime must be greater than LSP refresh interval.

Configuring the Interval for Receiving Duplicate LSP Packets

Command	<code>min-lsp-arrival [level-1 level-2] maximum-interval initial-interval hold-interval</code>
----------------	---

Parameter Description	<p>level-1: Applies the configuration only to Level-1.</p> <p>level-2: Applies the configuration only to Level-2.</p> <p><i>maximum-interval:</i> Indicates the maximum interval for receiving two duplicate LSP packets. The value range is 1 to 120 (in seconds).</p> <p><i>initial-interval:</i> Indicates the interval for receiving a duplicate LSP packet for the first time. The value range is 0 to 60000 (in milliseconds).</p> <p><i>hold-interval:</i> Indicates the minimum interval for receiving a duplicate LSP packet for the second time. The value range is 10 to 60000 (in milliseconds).</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>The interval for receiving a duplicate LSP packet for the first time is the initial interval. When the interval for receiving duplicate LSP packets is less than the maximum interval, the interval for receiving a duplicate LSP packet for the second time becomes the hold interval. In addition, a corresponding penalty is added to this interval: The next interval for receiving a duplicate LSP packet doubles the previous interval for receiving the same LSP packet, until this interval reaches the maximum interval. When the network becomes stable, the interval for receiving a duplicate LSP packet becomes greater than the maximum interval, and the interval for receiving a duplicate LSP packet is restored to the initial interval.</p> <p>Link changes have high requirements for convergence. The initial interval can be set to a small value. The preceding parameters can also be adjusted to larger values to reduce CPU consumption.</p> <p>The value of initial-interval cannot be greater than that of maximum-interval. Otherwise, the value of initial-interval will be used as the value of maximum-interval.</p> <p>The value of hold-interval cannot be greater than that of maximum-interval. Otherwise, the value of hold-interval will be used as the value of maximum-interval.</p> <p>The value of initial-interval cannot be greater than that of hold-interval. Otherwise, the value of initial-interval will be used as the value of hold-interval.</p>

Configuring to Ignore LSP Checksum Errors

Command	ignore-lsp-errors
Parameter Description	N/A

Command Mode	IS-IS routing process configuration mode
Usage Guide	After receiving an LSP, the local IS-IS neighbor calculates its checksum and compares it with the checksum contained in the LSP. By default, if the two checksums are inconsistent, the LSP will be discarded. If you run the ignore-lsp-errors command to configure to ignore checksum errors, the LSP will be processed normally despite checksum inconsistency.

Configuring LSP Fragment Extension

Command	lsp-fragments-extend [level-1 level-2] [compatible rfc3786]
Parameter Description	<p>level-1: Applies the setting only to Level-1 LSPs.</p> <p>level-2: Applies the setting only to Level-2 LSPs.</p> <p>compatible: Indicates compatibility with the RFC version of extended LSPs.</p> <p>rfc3786: Extends the LSP old version.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to enable LSP fragment extension.

Configuring an Additional System ID

Command	virtual-system <i>system-id</i>
Parameter Description	<i>system-id</i> : Indicates an additional system ID (6-byte).
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to configure the additional system ID of an IS-IS routing process, which is used by the extended LSP that is generated after the 256 fragments of the original LSP are used up. To enable fragment extension, run the lsp-fragments-extend command.

Configuration Example

Configuring the Minimum LSP Interval

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure the minimum LSP interval.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)# isis lsp-interval 100 level-2</pre>
Verification	Run the clear isis * command to update LSPs continuously and capture LSPs to check the minimum LSP interval.

Configuring the Maximum Number of LSP Packets Sent at a Time

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure the maximum number of LSP packets sent at a time.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)# isis lsp-flood 10 level-2</pre>
Verification	Run show isis interface command to check the maximum number of LSP packets sent at a time.

Configuring the LSP Retransmission Interval

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors in P2P mode. (Omitted)▪ Configure the LSP retransmission interval.
	<pre>A(config)# interface serial 0/1 A(config-if)# isis retransmit-interval 10 level-2</pre>
Verification	Disable neighboring routes and capture LSPs to check the LSP retransmission interval.

Configuring LSP Packet Generation Cycle

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure the LSP packet generation cycle.
	<pre>A(config)# router isis A(config-router)# lsp-gen-interval 5 50 100</pre>
Verification	Generate LSP packet frequently to see whether the exponential backoff algorithm is used.

Configuring the LSP Refresh Interval

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure the LSP refresh interval.
	<pre>A(config)# router isis A(config-router)# lsp-refresh-interval 600</pre>
Verification	Capture LSPs to check the refresh interval.

Configuring the LSP Lifetime

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure the LSP lifetime.
	<pre>A(config)# router isis A(config-router)# max-lsp-lifetime 1500</pre>
Verification	Check the LSP lifetime (LSP Holdtime field).
	<pre>A# show isis database</pre>

Configuring the Interval for Receiving Duplicate LSP Packets

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure the interval for receiving duplicate LSP packets.
	<pre>A(config)# router isis A(config-router)# min-lsp-arrival 10 100 200</pre>
Verification	Check LSP Seq Num field, that is, the receiving LSP status.
	<pre>A# show isis database</pre>

Configuring to Ignore LSP Checksum Errors

Configurati on Steps	<ul style="list-style-type: none">▪ Configure IS-IS neighbors. (Omitted)▪ Configure to ignore LSP checksum errors.
	<pre>A(config)# router isis A(config-router)# ignore-lsp-errors</pre>

Verification	Send an LSP with an incorrect checksum and check whether the LSP is discarded.
--------------	--

Configuring LSP Fragment Extension

Configuration Steps	<p>Configure IS-IS neighbors. (Omitted)</p> <ul style="list-style-type: none"> Configure LSP fragment extension. Configure the additional system ID of the IS-IS routing process.
	<pre>A(config)# router isis A(config-router)# lsp-fragments-extend A(config-router)# virtual-system 0000.0000.0034</pre>
Verification	Reduce the lsp-length originate command value, add routing information, and capture LSPs to check whether more than 256 LSP fragments are generated.

6.4.4 Configuring IS-IS SNPs

Configuration

Effect

- CSNPs are periodically broadcast by the DIS in a broadcast network for LSDB synchronization. In a P2P network, a CSNP is sent only after a neighbor relationship is established. An interface set to **mesh-groups** can be configured to periodically send CSNPs.
- When you need to set **mesh-group** on an IS-IS interface, run the **isis csnp-interval** command to configure the non-0 CSNP interval to ensure complete LSP synchronization between neighbors in the network. After that, CSNPs will be periodically sent to synchronize LSPs.

Configuration

Steps

- Perform this configuration based on requirements.
- Run the **isis csnp-interval interval [level-1 | level-2]** command in interface configuration mode on the desired device, unless otherwise specified.

Verification

Capture CSNPs in the broadcast network to check the CSNP interval.

Related

Commands

Configuring Source Registration Filter

Command	isis csnp-interval <i>interval</i> [level-1 level-2]
Parameter Description	<p><i>interval</i>: Indicates the CSNP interval. The value range is 0 to 65,535, in the unit of seconds.</p> <p>level-1: Applies the setting only to Level-1 CSNPs.</p> <p>level-2: Applies the setting only to Level-2 CSNPs.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>Use this command to change the CSNP interval. By default, a DIS sends a CSNP every 10s in a broadcast network.</p> <p>In a P2P network, a CSNP is sent only after a neighbor relationship is established. An interface set to mesh-groups can be configured to periodically send CSNPs.</p> <p>No CSNPs are sent if the CSNP interval is set to 0.</p>

Configuration Example

Configuring the CSNP Broadcast Interval

Configuration Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure the CSNP broadcast interval.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)# isis csnp-interval 20</pre>
Verification	Capture packets to check the CSNP interval.

6.4.5 Configuring the IS-IS Level Type

Configuration Effect

- IS-IS supports a two-Level system to realize routing management and extensible route selection in a large network. Each Level is only concerned about maintaining the topology of the corresponding area.
- You can run the **is-type** command in IS-IS routing process configuration mode to configure an IS-IS Level, or run the **isis circuit-type** command in interface configuration mode to configure the IS-IS

Level of an interface. The default Levels specified by the **is-type** and **isis circuit-type** commands are Level-1/Level-2. If you run both commands, the interface only sends the PDUs of the same Level specified by the two commands.

Notes

- If Level-1 or Level-2-only is configured using the **circuit-type** command, IS-IS will only send PDUs of the corresponding Level.
- If an interface is set to **external**, the interface will work as an external domain interface and IS-IS will not send PDUs of the corresponding Level.
- A device can have only one instance running at Level-2 (including Level-1/Level-2).

Configuration

Steps

Configuring the System Type

- Perform this configuration based on requirements.
- Run the **is-type** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the Interface Circuit Type

- Perform this configuration based on requirements.
- Run the **isis circuit-type** command in interface configuration mode on the desired device, unless otherwise specified.

Verification

- Check whether only the instances of the Level specified by the **is-type** command are processed, and neighbors of the corresponding Level are created.
- Check whether the interface only sends the PDUs of the same Level specified by the **is-type** and **circuit-type** commands.

Related

Commands

Configuring the System Type

Command	is-type { level-1 level-1-2 level-2-only }
Parameter Description	level-1: Indicates that IS-IS only runs at Level-1. level-1-2: Indicates that IS-IS runs at Level-1 and Level-2. level-2-only: Indicates that IS-IS only runs at Level-2.

Command Mode	IS-IS routing process configuration mode
Usage Guide	Changing the is-type value will enable or disable the routes of the corresponding level.

Configuring the Interface Circuit Type

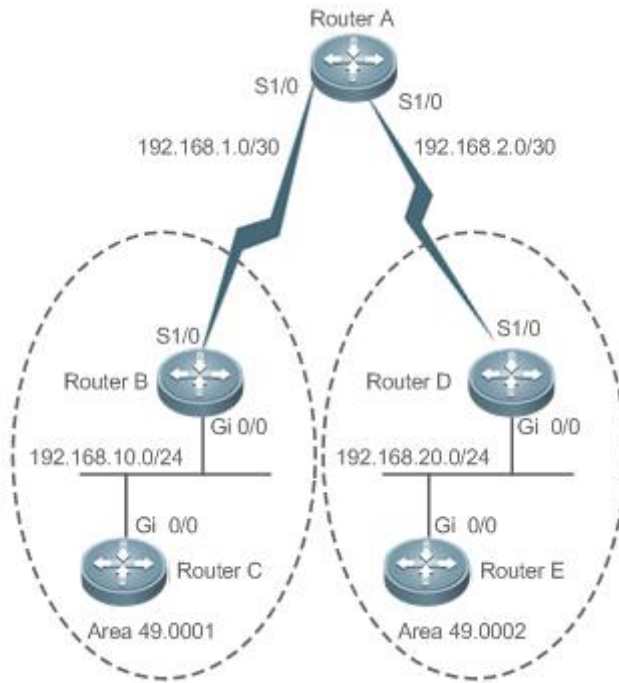
Command	isis circuit-type { level-1 level-1-2 level-2-only [external] }
Parameter Description	<p>level-1: Establishes a Level-1 neighbor relationship.</p> <p>level-2-only: Establishes a Level-2 neighbor relationship.</p> <p>level-1-2: Establishes a Level-1/Level-2 neighbor relationship.</p> <p>external: Uses the interface as an external domain interface.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>If the circuit type is set to Level-1 or Level-2-only, IS-IS will only send PDUs of the corresponding Level.</p> <p>If the system type is set to Level-1 or Level-2-only, IS-IS only processes the instances of the corresponding Level, and the interface only sends the PDUs of the same Level specified by the is-type and circuit-type commands.</p> <p>If the interface is set to external, the interface will work as an external domain interface and IS-IS will not send PDUs of the corresponding Level.</p>

Configuration

Example

Configuring IS-IS Levels

Configuration Requirements	Router A is connected to Router B and Router C by P2P serial links. Router B and Router C are connected by the Ethernet, and Router D and Router E are also connected by the Ethernet. On Router A, configure IS-IS area route summarization. Note that area route summarization can be configured only on border devices.
Figure 6-12 IS-IS Level Configuration	



Configuration Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS. ▪ Configure Ethernet interfaces. ▪ Configure the IS-IS Level structure.
A	<p>Configure IS-IS.</p> <pre>A(config)# router isis A(config-router)# net 50.0001.0000.0000.0001.00 A(config-router)# is-type level-2-only</pre>
	<p>Configure two serial link ports.</p> <pre>A(config)# interface Serial 1/0 A(config-if)# ip address 192.168.1.1 255.255.255.252 A(config-if)# ip router isis A(config)# interface Serial 1/1 A(config-if)# ip address 192.168.2.1 255.255.255.252 A(config-if)# ip router isis</pre>
B	<p>Configure IS-IS.</p>
	<pre>B(config)# router isis</pre>

	B(config-router)# net 49.0001.0000.0000.0002.00
	Configure an Ethernet interface.
	B(config)# interface GigabitEthernet 0/0 B(config-if)# ip address 192.168.10.1 255.255.255.0 B(config-if)# ip router isis
	Configure a serial link port.
	B(config)# interface Serial 1/0 B(config-if)# ip address 192.168.1.2 255.255.255.252 B(config-if)# ip router isis
C	Configure IS-IS.
	C(config)# router isis C(config-router)# net 49.0001.0000.0000.0003.00 C(config-router)# is-type level-1
	Configure an Ethernet interface.
	C(config)# interface GigabitEthernet 0/0 C(config-if)# ip address 192.168.10.2 255.255.255.0 C(config-if)# ip router isis
D	Configure IS-IS.
	D(config)# router isis D(config-router)# net 49.0002.0000.0000.0004.00
	Configure an Ethernet interface.
	D(config)# interface GigabitEthernet 0/0 D(config-if)# ip address 192.168.20.1 255.255.255.0 D(config-if)# ip router isis

	Configure a serial link port.
	D(config)# interface Serial 1/0 D(config-if)# ip address 192.168.2.2 255.255.255.252 D(config-if)# ip router isis
E	Configure IS-IS.
	E(config)# router isis E(config-router)# net 49.0002.0000.0000.0005.00 E(config-router)# is-type level-1
	Configure an Ethernet interface.
	E(config)# interface GigabitEthernet 0/0 E(config-if)# ip address 192.168.20.2 255.255.255.0 E(config-if)# ip router isis
Verification	<ul style="list-style-type: none"> ▪ Check whether neighbor relationships are established normally. ▪ Capture packets to check whether Router A only sends and receives Level-2 packets. ▪ Capture packets to check whether Router B and Router D only send and receive Level-1 and Level-2 packets. ▪ Capture packets to check whether Router C and Router E only send and receive Level-1 packets.
A	A# show isis neighbors A# show isis database detail
B	B# show isis neighbors B# show isis database detail
C	C# show isis neighbors C# show isis database detail
D	D# show isis neighbors D# show isis database detail

E

```
E# show isis neighbors
E# show isis database detail
```

6.4.6 Configuring IS-IS Authentication

Configuration

Effect

- Interface authentication is intended for establishing and maintaining neighbor relationships. A neighbor relationship cannot be established between two IS-IS devices with different interface authentication passwords. This prevents unauthorized or unauthenticated IS-IS devices from joining an IS-IS network that requires authentication. Interface authentication passwords are encapsulated in Hello packets before being sent.
- Area authentication and RD authentication in IS-IS are performed to verify LSPs, CSNPs, and PSNPs to prevent unauthorized or unauthenticated routing information from being injected into the LSDB. Authentication passwords are encapsulated in LSPs, CSNPs, and PSNPs before being sent.

Notes

- An interface authentication password is encapsulated in a Hello packet before being sent by an interface. When an interface receives a Hello packet, it checks the password in the packet against the existing one.
- Area authentication passwords are encapsulated in Level-1 LSPs, CSNPs, and PSNPs. When an interface receives an LSP, CSNP, or PSNP, it checks the password in the packet against the existing one.
- RD authentication passwords are encapsulated in Level-2 LSPs, CSNPs, and PSNPs. When an interface receives an LSP, CSNP, or PSNP, it checks the password in the packet against the existing one.

Configuration

Steps

Configuring Interface Authentication

- Perform this configuration based on requirements.
- Configure **isis password** in interface configuration mode on the desired device, unless otherwise specified.

Configuring Area Authentication

- Perform this configuration based on requirements.
- Run the **area-password** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring RD Authentication

- Perform this configuration based on requirements.
- Run the **domain-password** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- IS-IS plaintext authentication provides only limited security because the password transferred through a packet is visible.
- IS-IS MD5 authentication provides higher security because the password transferred through a packet is encrypted using the MD5 algorithm.

Related Commands

Configuring the Password for Plaintext Authentication of Hello Packets on an Interface

Command	<code>isis password [0 7] password [send-only] [level-1 level-2]</code>
Parameter Description	<p>0: Indicates that the key is displayed in plaintext.</p> <p>7: Indicates that the key is displayed in ciphertext.</p> <p>password-string: Indicates the password string for plaintext authentication. The string can contain up to 126 characters.</p> <p>send-only: Indicates that the plaintext authentication password is only used to authenticate sent packets. Received packets are not authenticated.</p> <p>level-1: Applies the setting to the Level-1 circuit type.</p> <p>level-2: Applies the setting to the Level-2 circuit type.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>Use this command to configure the password for Hello packet authentication on an interface. Use the no form of this command to clear the password.</p> <p>If no Level is specified, by default, the password takes effect for Level-1 and Level-2 circuit types.</p> <p>This command does not take effect if the isis authentication mode command is executed. You need to first delete the previous command configuration.</p>

If you include the **send-only** parameter when deleting the **isis authentication mode** command configuration, only the parameter setting is canceled.

Specifying Interface Authentication as Plaintext or MD5

Command	isis authentication mode { md5 text } [level-1 level-2]
Parameter Description	md5 : Uses MD5 authentication. text : Uses plaintext authentication. level-1 : Applies the setting to the Level-1 circuit type. level-2 : Applies the setting to the Level-2 circuit type.
Command Mode	Interface configuration mode
Usage Guide	Use this command to specify the authentication mode before you can make the key chain configured using the isis authentication key-chain command take effect. If no Level is specified, the authentication mode will take effect for Level-1 and Level-2 circuit types. If you use the isis authentication mode command after the isis password command is executed to configure plaintext authentication, the previous command configuration will be overwritten. The isis password command does not take effect if the isis authentication mode command is executed. To run the isis password command, delete the isis authentication mode command configuration first.

Configuring the Password for Interface Authentication

Command	isis authentication key-chain <i>name-of-chain</i> [level-1 level-2]
Parameter Description	<i>name-of-chain</i> : Indicates the name of a key chain. The maximum length is 255. level-1 : Indicates that the authentication key chain takes effect for Level-1. level-2 : Indicates that the authentication key chain takes effect for Level-2.
Command Mode	Interface configuration mode

Usage Guide	<p>Authentication is not performed if no key chain is configured using the key chain command. In addition to the key chain command, you also need to run the isis authentication mode command to make IS-IS key chain authentication take effect.</p> <p>The key chain is applicable to plaintext authentication and MD5 authentication. Which authentication mode to use can be determined using the isis authentication mode command.</p> <p>For plaintext authentication, the key-string in the key chain cannot exceed 80 characters; otherwise, the key chain will be invalid.</p> <p>Only one key chain can be used at a time. After you configure a new key chain, it will replace the original one.</p> <p>If no Level is specified, the key chain takes effect for Level-1 and Level-2.</p> <p>The key chain is applicable to Hello packets. IS-IS will send or receive passwords that belong to the key chain.</p> <p>A key chain may contain multiple passwords. A password with a smaller SN is preferentially used for sending a packet. When the packet arrives at the peer device, the device will receive the packet if the packet-carried password is consistent with a password in the key chain.</p> <p>The authentication commands (for example, authentication key-chain) executed in IS-IS routing process configuration mode are intended for LSPs and SNPs. They do not take effect for IS-IS interfaces.</p>
-------------	--

(Optional) Applying Interface Authentication Only to Sent Packets (Received Packets Are Not Authenticated)

Command	isis authentication send-only [level-1 level-2]
Parameter Description	<p>level-1: Sets send-only for Level-1 on an interface.</p> <p>level-2: Sets send-only for Level-2 on an interface.</p>
Command Mode	Interface configuration mode
Usage Guide	<p>Use this command to enable IS-IS to set an authentication password in the Hello packet sent by an interface. However, IS-IS does not authenticate the Hello packet received by the interface. You can use this command before you deploy IS-IS interface authentication on all devices in the network or before you change the authentication password or authentication mode. After you run the isis authentication send-only command, the devices will not authenticate received Hello packets to avoid network flapping when IS-IS interface authentication is deployed. After authentication is deployed in the entire</p>

	<p>network, run the no isis authentication send-only command to cancel the send-only setting.</p> <p>The isis authentication send-only command is applicable to plaintext authentication and MD5 authentication. You can run the isis authentication mode command to specify the authentication mode for an IS-IS interface.</p> <p>If no Level is specified, the authentication mode will take effect for Level-1 and Level-2 on the interface.</p>
--	--

Configuring the Password for Area (Level-1) Plaintext Authentication

Command	area-password [0 7] password [send-only]
Parameter Description	<p>0: Indicates that the key is displayed in plaintext.</p> <p>7: Indicates that the key is displayed in ciphertext.</p> <p>password-string: Indicates the password string for plaintext authentication. The string can contain up to 126 characters.</p> <p>send-only: Indicates that the plaintext authentication password is only used to authenticate sent Hello packets in Level-1 areas. Received Hello packets are not authenticate.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Run this command to enable authentication of received LSPs, CSNPs, and PSNPs in Level-1 areas and include authentication information in these packets before they are sent. All IS-IS devices in an area must be configured with the same password.</p> <p>This command does not take effect if the authentication mode command is executed. You need to first delete the previous command configuration.</p> <p>To delete the password, run the no area-password command. If you run the no area-password send-only command, only the send-only setting is canceled. If you run the area-password psw send-only and no area-password send-only commands in sequence, the configuration is changed to area-password psw.</p>

Configuring the Password for RD (Level-2) Plaintext Authentication

Command	domain-password [0 7] password [send-only]
----------------	---

Parameter Description	<p>0: Indicates that the key is displayed in plaintext.</p> <p>7: Indicates that the key is displayed in ciphertext.</p> <p>password-string: Indicates the password string for plaintext authentication. The string can contain up to 126 characters.</p> <p>send-only: Indicates that the plaintext authentication password is only used to authenticate sent Hello packets in Level-1 areas. Received Hello packets are not authenticated.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Run this command to enable authentication of received LSPs, CSNPs, and PSNPs in Level-2 domains and include authentication information in these packets before they are sent. All IS-IS devices in a Level-2 domain must be configured with the same password.</p> <p>This command does not take effect if the authentication mode command is executed. You need to first delete the previous command configuration.</p> <p>To delete the password, run the no domain-password command. If you run the no domain-password send-only command, only the send-only setting is canceled. If you run the domain-password psw send-only and no domain-password send-only commands in sequence, the configuration is changed to domain-password psw.</p>

Specifying the IS-IS RD Authentication Mode

Command	authentication mode { md5 text } [level-1 level-2]
Parameter Description	<p>md5: Uses MD5 authentication.</p> <p>text: Uses plaintext authentication.</p> <p>level-1: Indicates that the authentication mode takes effect for Level-1.</p> <p>level-2: Indicates that the authentication mode takes effect for Level-2.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Use this command to specify the authentication mode before you can make the key chain configured using the authentication key-chain command take effect.</p> <p>If no Level is specified, the authentication mode will take effect for Level-1 and Level-2.</p>

	<p>If you use the authentication mode command after the area-password or domain-password command is executed to configure plaintext authentication, the previous command configuration will be overwritten.</p> <p>The area-password or domain-password command does not take effect if the authentication mode command is executed. To run the area-password or domain-password command, delete the authentication mode command configuration first.</p>
--	--

Specifying the Key Chain for IS-IS Authentication

Command	authentication key-chain <i>name-of-chain</i> [level-1 level-2]
Parameter Description	<p><i>name-of-chain</i>: Indicates the name of a key chain. The maximum length is 255.</p> <p>level-1: Indicates that the authentication key chain takes effect for Level-1.</p> <p>level-2: Indicates that the authentication key chain takes effect for Level-2.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Authentication is not performed if no key chain is configured using the key chain command. In addition to the key chain command, you also need to run the authentication mode command to make IS-IS key chain authentication take effect.</p> <p>The key chain is applicable to plaintext authentication and MD5 authentication. Which authentication mode to use can be determined using the authentication mode command.</p> <p>For plaintext authentication, the key-string in the key chain cannot exceed 80 characters; otherwise, the key chain will be invalid.</p> <p>Only one key chain can be used at a time. After you configure a new key chain, it will replace the original one.</p> <p>If no Level is specified, the key chain takes effect for Level-1 and Level-2.</p> <p>The key chain is applicable to LSPs, CSNPs, and PSNPs. IS-IS will send or receive passwords that belong to the key chain.</p> <p>A key chain may contain multiple passwords. A password with a SN is preferentially used for sending a packet. When the packet arrives at the peer device, the device will receive the packet if the packet-carried password is consistent with a password in the key chain.</p>

Applying IS-IS Authentication Only to Sent Packets

Command	authentication send-only [level-1 level-2]
----------------	---

Parameter Description	<p>level-1: Applies the send-only setting to Level-1.</p> <p>level-2: Applies the send-only setting to Level-2.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Use this command to enable IS-IS to set an authentication password in the Hello packet to be sent. However, IS-IS does not authenticate received Hello packets. You can use this command before you deploy IS-IS authentication on all devices in the network or before you change the authentication password or authentication mode. After you run the authentication send-only command, the devices will not authenticate received packets to avoid network flapping when authentication passwords are deployed. After authentication is deployed in the entire network, run the no isis authentication send-only command to cancel the send-only setting.</p> <p>The authentication send-only command is applicable to plaintext authentication and MD5 authentication. You can run the authentication mode command to specify the authentication mode.</p> <p>If no Level is specified, the authentication mode will take effect for Level-1 and Level-2.</p>

Configuration Example

Configuring IS-IS Authentication

Configuration Requirements	<p>Router A, Router B, and Router C are connected through the Ethernet and run IS-IS. Router A is a Level-1 device, Router B is a Level-1/Level-2 device, and Router C is a Level-2 device. The following configuration requirements exist: Apply plaintext authentication to the Hello packets between Router A and Router B, as well as Level-1 LSPs and SNPs. Apply MD5 authentication to the Hello packets between Router B and Router C, as well as Level-2 LSPs and SNPs.</p>
Figure 6-13 IS-IS Authentication Topology	<p>The diagram illustrates a network topology with three routers: Router A (Level 1), Router B (Level 1/Level 2), and Router C (Level 2). Router A is connected to Router B via their G0/0 and G0/1 interfaces. Router B is connected to Router C via their G0/1 and G0/1 interfaces. The IP address for Router A is 192.168.2.0/24. The area ID for Router B is 49.0001, and for Router C is 49.0002. The IP address for Router C is 192.168.30.0/24.</p>

<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure IS-IS. ▪ Configure Ethernet interfaces. ▪ Configure the password for IS-IS authentication.
<p>A</p>	<p>Configure IS-IS.</p>
	<pre>A(config)# router isis A(config-router)# net 49.0001.0000.0000.0001.00 A(config-router)# is-type level-1 A(config-router)# area-password aa</pre>
	<p>Configure an Ethernet interface.</p> <pre>A(config)# interface GigabitEthernet 0/0 A(config-if)# ip address 192.168.20.1 255.255.255.0 A(config-if)# ip router isis A(config-if)# isis password cc</pre>
<p>B</p>	<p>Configure the password for IS-IS authentication.</p>
	<pre>B(config)# key chain kc1 B(config-keychain)# key 1 B(config-keychain-key)# key-string aa B(config)# key chain kc2 B(config-keychain)# key 1 B(config-keychain-key)# key-string bb B(config)# key chain kc3 B(config-keychain)# key 1 B(config-keychain-key)# key-string cc</pre>
	<p>Configure IS-IS.</p>
	<pre>B(config)# router isis B(config-router)# net 49.0001.0000.0000.0002.00 B(config-router)# authentication mode text level-1 B(config-router)# authentication key-chain kc1 B(config-router)# authentication mode md5 level-2</pre>

	B(config-router)# authentication key-chain kc2
	Configure two Ethernet interfaces.
	B(config)# interface GigabitEthernet 0/0 B(config-if)# ip address 192.168.20.2 255.255.255.0 B(config-if)# ip router isis B(config-if)# isis authentication mode text B(config-if)# isis authentication key-chain kc3 B(config)# interface GigabitEthernet 0/1 B(config-if)# ip address 192.168.30.2 255.255.255.0 B(config-if)# ip router isis B(config-if)# isis authentication mode md5 B(config-if)# isis authentication key-chain kc3
C	Configure the password for IS-IS authentication.
	C(config)# key chain kc2 C(config-keychain)# key 1 C(config-keychain-key)# key-string bb C(config)# key chain kc3 C(config-keychain)# key 1 C(config-keychain-key)# key-string cc
	Configure IS-IS.
	C(config)# router isis C(config-router)# net 49.0002.0000.0000.0002.00 C(config-router)# is-type level-2 C(config-router)# authentication mode md5 level-2 C(config-router)# authentication key-chain kc2 Configure an Ethernet interface.
	C(config)# interface GigabitEthernet 0/1 C(config-if)# ip address 192.168.30.3 255.255.255.0 C(config-if)# ip router isis

	<pre>C(config-if)# isis authentication mode md5 C(config-if)# isis authentication key-chain kc3</pre>
Verification	Check whether neighbor relationships are established normally.
A	<pre>A# show isis neighbors A# show isis database detail</pre>
B	<pre>B# show isis neighbors</pre>
C	<pre>C# show isis neighbors</pre>

Common

Errors

- Different authentication passwords are configured between neighbors.
- Different authentication modes are configured between neighbors.

6.4.7 Configuring IS-IS GR

Configuration

Effect

- IS-IS GR helps improve system reliability. On devices that separate the control plane from the forwarding plane, GR ensures that data forwarding is not interrupted during routing protocol restart.

IS-IS GR Working Mechanism

For GR to be successful, the following two conditions must be met: (1) The network topology is stable; (2) The device can ensure uninterrupted forwarding when it restarts IS-IS.

Two roles exist during the GR process: Restarter and Helper. Accordingly, IS-IS GR is divided into the IS-IS GR Restart capability and IS-IS GR Help capability. A device with the GR Restart capability can send a GR request and execute GR. A device with the GR Help capability can receive a GR request and help its neighbor with GR implementation. The GR process starts when the Restarter sends a GR request. After receiving the GR request, the neighboring device enters Help mode to help the Restarter reestablish its LSDB while maintaining the neighbor relationship with the Restarter. The main GR working mechanism is as follows:


When an IS-IS device needs to perform GR, it instructs its neighbor to maintain their neighbor relationship so that other devices in the network cannot sense the change in the topological relationship and the neighbor will not recalculate the route and update its forwarding table. The IS-IS

device synchronizes and restores the LSDB to its pre-GR state with the help of the neighbor to ensure that the route and forwarding table remain unchanged before and after GR implementation and data forwarding is not interrupted.

The Restarter performs the following operations during the GR process:

1. The GR Restarter notifies the GR Helpers that it will be restarted.


Figure 6-14 Restart Notification by the GR Restarter



Switch A is a GR Restarter, and Switch B and Switch C are the GR Helpers for Switch A. Switch A sends a GR request instructing all its neighbors not to delete the neighbor relationships with Switch A when it is restarted. After receiving the GR request, the neighbors send GR responses to the GR Restarter, and will maintain their neighbor relationships with the GR Restarter during the GR time (specified by **GR grace-period**) notified by the GR Restarter.

2. The GR Restarter is restarted.


Figure 6-15 Restart Performed by the GR Restarter



When the GR Restarter is restarted, its IS-IS interface goes from Down to Up. Because the GR Helpers know that the GR Restarter is in IS-IS restart state, they maintain their neighbor relationships with the GR Restarter during the GR time and retain the routes from the GR Restarter.

3. The GR Restarter synchronizes topology and routing information from the GR Helpers.


Figure 6-16 LSDB Synchronization



After IS-IS restart, the GR Restarter synchronizes topology or routing information from the GR Helpers and recalculates its routing table. During this process, any change in the routing table is not updated to the forwarding table.

4. GR is completed when the GR Restarter finishes LSDB synchronization. Then all devices enter IS-IS interaction state.

Figure 6-17 GR Completion



After the GR Restarter synchronizes all required data, all devices enter IS-IS interaction state. The GR Restarter's routing table is updated to the forwarding table and invalid entries are cleared. Because the GR Restarter is completely restored to the pre-restart state under stable network conditions, its routing table and forwarding table remain unchanged before and after GR.

Notes

- IS-IS GR is implemented based on RFC5306: Restart Signaling for IS-IS.

✔ All products support the IS-IS GR Helper capability.

Configuration

Steps

Enabling the IS-IS GR Restart Capability

- Perform this configuration based on requirements.
- Run the **graceful-restart** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the Maximum GR Time

- Perform this configuration based on requirements.
- Run the **graceful-restart grace-period** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Enabling the IS-IS GR Help Capability

- Perform this configuration based on requirements.
- Run the **graceful-restart helper** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- Check whether the routing table and forwarding table remain unchanged before and after GR.

Related

Commands

Enabling the IS-IS GR Restart Capability

Command	graceful-restart
Parameter Description	N/A
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to enable the IS-IS GR Restart capability. As long as the network conditions remain unchanged, IS-IS can be restarted and restored to the pre-restart state without impact on data forwarding.

Configuring the Maximum GR Time

Command	graceful-restart grace-period <i>seconds</i>
Parameter Description	<i>seconds</i> : Indicates the GR time. The value range is 1s to 65,535s. The default value is 300s.
Command Mode	IS-IS routing process configuration mode
Usage Guide	N/A

Enabling the IS-IS GR Help Capability

Command	graceful-restart helper disable
Parameter Description	N/A
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use the graceful-restart helper disable command to disable the IS-IS GR Help capability. The command enables IS-IS to ignore the GR request sent by the device to be restarted.

6.4.8 Configuring IS-IS NSR

Configuration

Effect

- IS-IS nonstop routing (IS-IS NSR) helps improve system reliability. During the active/standby switchover of devices in distributed or Virtual Switch Unit (VSU) mode, user data traffic keeps being forwarded without interruption.

Working Mechanism of IS-IS NSR

NSR backs up relevant IS-IS information from the master Supervisor Engine to the slave Supervisor Engine of the distributed device, or from the master device to the slave device in VSU mode, so that the device can automatically recover the link state and regenerate a route upon active/standby switchovers, without requiring help from neighbor devices during the recovery.

Notes

For the same IS-IS process, either NSP or GR is enabled, because they are exclusive.

The switchover of distributed devices and VSU devices takes a period of time. If the IS-IS neighbor keepalive duration is less than the switchover duration, IS-IS neighbor relationship with the neighbor device is removed, and the services are interrupted during the switchover. Therefore, it is recommended to set the IS-IS neighbor keepalive duration not less than the default value. When Fast Hello is enabled, the IS-IS neighbor keepalive duration is less than 1 second and the IS-IS neighbor relationship times out during the switchover, causing NSR failures. Therefore, it is recommended to disable Fast Hello when NSR is enabled.

Configuration Steps

Configuring IS-IS NSR

- Optional. IS-IS NSR is disabled by default. Perform this configuration as required.
- Configure NSR in IS-IS routing process configuration mode on the required devices unless otherwise specified.

Verification

- Verify that user data traffic keeps being forwarded without interruption during the hot-backup switchover of devices in distributed or VSU mode.

Related Commands

Enabling NSR

Command	<code>nsr</code>
Parameter Description	N/A
Configuration Mode	IS-IS routing process configuration mode
Usage Guide	<p>NSR backs up relevant IS-IS information from the master Supervisor Engine to the slave Supervisor Engine of the distributed device, or from the master device to the slave device in VSU mode, so that the device can automatically recover the link state and regenerate a route upon active/standby switchovers, without requiring help from neighbor devices during the recovery. Information such as the neighbor relationship and link status needs to be backed up.</p> <p>For the same IS-IS process, either NSP or GR is enabled, because they are exclusive.</p>
	The switchover of distributed devices and VSU devices takes a period of time. If the IS-IS neighbor keepalive duration

is less than the switchover duration, IS-IS neighbor relationship with the neighbor device is removed, and the services are interrupted during the switchover. Therefore, it is recommended to set the IS-IS neighbor keepalive duration not less than the default value. When Fast Hello is enabled, the IS-IS neighbor keepalive duration is less than 1 second and the IS-IS neighbor relationship times out during the switchover, causing NSR failures. Therefore, it is recommended to disable Fast Hello when NSR is enabled.

By default, if the DIS is a QTECH device, the neighbor relationship between the neighbor device and the QTECH device is maintained for 10 seconds at most (the DIS interface sends one HELLO packet every 3.3 seconds). However, the active/standby switchover of QTECH devices takes at least 7 seconds. If the switchover takes place when the interface is to send HELLO packets, the period for the neighbor device to receive HELLO packets again may be greater than 10 seconds, leading to neighbor down. To avoid neighbor down, when NSR is enabled, increase the time for maintaining the neighbor relationship with the neighbor device that is to send HELLO packets by 10 seconds.

6.4.9 Configuring BFD Support for IS-IS

Configuration

Effect

- IS-IS dynamically discovers neighbors through Hello packets. After IS-IS enables the BFD function, a BFD session will be set up with the neighbor in Up state. The BFD mechanism is used to detect the neighbor state. Once a neighbor failure is detected through BFD, IS-IS performs network convergence immediately. The convergence time can be reduced from 30s to less than 1s. By default, IS-IS Hello packets are sent at an interval of 10s in a P2P network, and the time required to detect a neighbor failure is three times the packet interval, that is 30s.

Notes

- You must set BFD session parameters before you enable BFD support for IS-IS.
- When you run the **bfd up-dampening** command on an interface with BFD support for IS-IS, you need to run the **bfd all-interfaces** command with the *[anti-congestion]* option selected.
- When you run the **bfd all-interfaces** command with the *[anti-congestion]* option selected, run the **bfd up-dampening** command on the interface.
- IP routing may cause a neighbor's interface for BFD session setup to be inconsistent with the interface for outgoing BFD packets. If this happens, the BFD session cannot be set up.

- If a neighbor's interface for BFD session setup is inconsistent with the interface for outgoing BFD packets, the BFD session cannot be set up.

Configuration

Steps

Enabling BFD Support for IS-IS on All Interfaces

- Perform this configuration based on requirements.
- Run the **bfd all-interfaces** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Enabling BFD Support for IS-IS on the Current Interface

- Perform this configuration based on requirements.
- Run the **isis bfd** command in interface configuration mode on the desired device, unless otherwise specified.

Verification

- Build a topology with two parallel lines. Typically, IS-IS selects one line as the master line and the other as the backup line. Enable BFD on the master line.
- Make the master line fail. Check whether IS-IS performs route convergence based on the BFD monitoring state and starts the backup line.

Related

Commands

Enabling BFD Support for IS-IS on the Current Interface

Command	bfd all-interfaces [<i>anti-congestion</i>]
Parameter Description	<i>anti-congestion</i> : Indicates the IS-IS BFD anti-congestion option.
Command Mode	IS-IS routing process configuration mode
Usage Guide	You can enable or disable BFD on an IS-IS interface by using any of the following two methods: Method 1: Run the bfd all-interfaces command in IS-IS routing process configuration mode to enable BFD on all IS-IS interfaces, and then run the no bfd all-interfaces command to disable BFD on all IS-IS interfaces.

Method 2: Run the **isis bfd [disable]** command in interface configuration mode to enable BFD on the specified IS-IS interface, and then run the **isis bfd disable** command to disable BFD on the interface.

Enabling BFD Support for IS-IS on the Current Interface

Command	isis bfd [<i>disable</i> <i>anti-congestion</i>]
Parameter Description	<i>disable</i> : Disables BFD support for IS-IS on the current interface. <i>anti-congestion</i> : Indicates the IS-IS BFD anti-congestion option.
Command Mode	Interface configuration mode
Usage Guide	<p>You can enable or disable BFD on an IS-IS interface by using any of the following two methods:</p> <p>Method 1: Run the [no] bfd all-interfaces [anti-congestion] command in IS-IS routing process configuration mode to enable or disable BFD on all IS-IS interfaces.</p> <p>Method 2: Run the isis bfd [disable anti-congestion] command in interface configuration mode to enable or disable BFD on the specified interface.</p> <p>Normally, BFD sends detection packets at millisecond intervals to detect the link state. When a link exception (such as a disconnected link) occurs, BFD can quickly detect it and instruct IS-IS to delete the neighbor relationship and the neighbor reachability information in LSPs. Then IS-IS recalculates and generates a new route to bypass the abnormal link, thus realizing fast convergence. With the introduction of new techniques such as the Multi-Service Transport Platform (MSTP), link congestion tends to occur during peak hours of data communication. BFD quickly detects the link exception and instructs IS-IS to delete the neighbor relationship and the neighbor reachability information in LSPs. Link switch is performed to bypass the congested link. A Hello packet for IS-IS neighbor detection is sent every 10s and its expiration time is 30s. The Hello packet can still be received normally when BFD detects an exception, and therefore an IS-IS neighbor relationship is reestablished quickly, causing the route to be restored to the congested link. Then BFD detects the abnormal link and link switch is performed again. This process is repeated, which makes the route be switched between the congested link and other links, causing repetitive flapping.</p> <p>The anti-congestion option is used to avoid routing flapping in case of link congestion. After the option is configured, the IS-IS neighbor state is still kept alive when link congestion occurs, but the neighbor reachability information in LSPs is deleted. The route is switched to a normal link. When the congested link is restored, the neighbor reachability information in LSPs is recovered and the route is switched back, which avoids route flapping.</p>

	When you run the bfd all-interfaces [<i>anti-congestion</i>] command, run the bfd up-dampening command on the interface. The two commands must be used together. If you run only one command, the route flap dampening feature may not take effect or other network exceptions may occur.
--	---

Configuration

Example

Enabling BFD Support for IS-IS on the Current Interface

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Set BFD session parameters. (Omitted) ▪ Enable BFD support for IS-IS on the current interface.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)# isis bfd</pre>
Verification	<p>Enable S1 (192.168.1.10) and S2 (192.168.2.10) to send packets to G1 (229.1.1.1) and G2 (229.1.2.1). Add User to the G1 and G2 groups.</p> <ul style="list-style-type: none"> ▪ Check the multicast packet that User receives. User should only receive the (S1, G1) packet. ▪ Check that the PIM-SM routing table does not have the (S1, G2), (S2, G1), and (S2, G2) entries.
	<pre>A# show bfd neighbors detail</pre>

Common

Errors

- BFD support for IS-IS is not enabled on neighbors.

6.4.10 Setting the IS-IS Overload Bit

Configuration

Effect

The overload bit is used in the following three situations:

- Device overload

The local IS-IS node has overload issues, such as insufficient memory or full CPU load; as a result, its routing table has incomplete routes or does not have resource forwarding data. You can set the overload bit in an LSP to instruct the neighbor not to use the local node as a forwarding device.

To set the overload bit, run the **set-overload-bit** command without the **on-startup** keyword. The overload bit can be configured or canceled manually. When the local IS-IS node is restored, manually cancel the command configuration; otherwise, the node is always in overload state.

- Instantaneous black hole

In the scenario described by RFC3277, the IS-IS convergence speed is faster than the BGP speed; as a result, after an IS-IS node is restarted, a route may be instantaneously unreachable, which is called an instantaneous black hole. You can set the overload bit in an LSP to instruct the neighbor not to use the local node as a forwarding device until the specified time has elapsed.

To set the overload bit, run the **set-overload-bit** command with the **on-startup** keyword. The overload bit can be configured or canceled automatically by the IS-IS node based on the configuration. If the **on-startup** keyword is selected, the IS-IS node automatically enters instantaneous black hole state after restart. When a neighbor relationship is established, the IS-IS node sends an LSP with the overload bit to notify the neighbor that the local node enters instantaneous black hole (or overload) state and instruct the neighbor not to use the local node as a forwarding device. After the specified time has elapsed, the IS-IS node immediately sends an LSP with the overload bit canceled to notify the neighbor that the local node has exited instantaneous black hole (or overload) state and can work as a forwarding device.

- Disabling real data forwarding on the local IS-IS node

If you only need to connect the local IS-IS node to a production network for testing or to meet other functional requirements, but does not require the node to forward real data in the network, you can set the overload bit in an LSP to instruct the neighbor not to use the local node as a forwarding device.

To set the overload bit, run the **set-overload-bit** command without the **on-startup** keyword. The overload bit can be configured or canceled manually. You can set the **suppress** keyword based on requirements to limit the routing information carried in an LSP in case of overload. For example, internal and external routes can be suppressed, and only the local direct route is advertised.

Notes

- At the same Level, the configuration with the **on-startup** keyword is mutually exclusive with the configuration without the **on-startup** keyword.

Configuration Steps

- Perform this configuration based on requirements.

- Run the **set-overload-bit** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- Capture packets and check that the neighbor does not forward LSPs from the local node.

Related Commands

Command	set-overload-bit [on-startup { seconds wait-for-bgp [bgpseconds] }] [suppress { [interlevel] [external] }] [level-1 level-2]
Parameter Description	<p>on-startup seconds: Indicates the duration when an IS-IS node remains in overload state after restart. The value range is 5s to 86,400s.</p> <p>wait-for-bgp bgpseconds: Indicates the duration in which an IS-IS node automatically enters and remains in the wait-for-BGP convergence state after restart. The value ranges from 5–86,400 in seconds. The default value is 600.</p> <p>suppress: Indicates not to advertise internal routes (intra-area and inter-area routes) or external routes to neighbors when the IS-IS node is in overload state.</p> <p>interlevel: Indicates not to advertise intra-area and inter-area routes to neighbors when the IS-IS node is in overload state. It is used with the suppress keyword.</p> <p>external: Indicates not to advertise external routes to neighbors when the IS-IS node is in overload state. It is used with the suppress keyword.</p> <p>level-1: Sends LSPs with the overload bit only to Level-1 neighbors.</p> <p>level-2: Sends LSPs with the overload bit only to Level-2 neighbors.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Use this command to force an IS-IS node to set the overload bit in a non-virtual LSP to instruct its IS-IS neighbors not to use the local node as a forwarding device.</p> <p>If you select the on-startup keyword, the IS-IS node automatically enters overload state after restart.</p> <p>If you do not select the on-startup keyword, the IS-IS node enters overload state immediately after restart.</p>

Configuration Example

Configuring the Overload Bit in Case of an Instantaneous Black Hole

Configurati on Steps	Configure IS-IS neighbors. (Omitted) <ul style="list-style-type: none">Verify that the IS-IS node enters instantaneous black hole state immediately after restart and remains in this state until the specified time (300s) has elapsed, and the IS-IS node only advertises local direct links to its neighbors during the specified time.
	<pre>A(config)# router isis A(config-router)#set-overload-bit on-startup 300 suppress interlevel external</pre>
Verification	<ul style="list-style-type: none">Capture packets to check LSPs.Verify that the IS-IS node automatically enters instantaneous black hole state after restart. Once a neighbor relationship is established, the IS-IS node sends an LSP with the overload bit.After the specified time has elapsed, the IS-IS node immediately sends an LSP with the overload bit canceled to notify its neighbors that the local node has exited instantaneous black hole (or overload) state.
	<pre>A# show isis neighbors</pre>

Disabling Real Data Forwarding on the Local IS-IS Node

Configurati on Steps	<ul style="list-style-type: none">Configure IS-IS neighbors. (Omitted)Connect the local IS-IS node as a test device to a production network. The node is not required to forward real data in the network to avoid impact on production.
	<pre>A(config)# router isis A(config-router)#set-overload-bit suppress interlevel external</pre>
Verification	Capture packets to check LSPs. Verify that the LSPs carry the overload bit and only advertise local direct routes.
	<pre>A# show isis neighbors</pre>

6.4.11 Configuring IS-IS VRF

Configuration Effect

- Each VRF table can be seen as a virtual device or a dedicated PE device.
- The virtual device contains the following elements: an independent routing table, as well as an independent address space; a set of interfaces that belong to the VRF table; a set of routing protocols applicable only to the VRF table.
- Each device can maintain one or more VRF tables and a public-network routing table (also called a global routing table). Multiple VRF instances are separated from each other.

Notes

- Note the following constraints or conventions when you bind IS-IS instances and VRF tables:
- The IS-IS instances bound with the same VRF table must be configured with different system IDs. The IS-IS instances bound with different VRF tables can be configured with the same system ID.
- One IS-IS instance can be bound with only one VRF table, but one VRF table can be bound to multiple IS-IS instances.
- When the VRF table bound to an IS-IS instance is changed, all IS-IS interfaces associated with the instance will be deleted. That is, the **ip router isis [tag]** interface configuration and the redistribution configuration in routing process configuration mode will be deleted.

Configuration

Steps

- Perform this configuration based on requirements.
- Run the **vrf** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- Check whether the local device establishes neighbor relationships with other devices specified in the VRF table.

Related

Commands

Configuring IS-IS VRF

Command	vrf vrf-name
Parameter Description	<i>vrf-name</i> : Indicates the name of an existing VRF table.

Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>Before you bind an IS-IS instance to a VRF table, ensure that the VRF table has been configured. If you need to establish an IS-ISv6 neighbor relationship, enable IPv6 and ensure that the table to be bound is a multiprotocol VRF table.</p> <p>Note the following constraints or conventions when you bind IS-IS instances and VRF tables:</p> <ul style="list-style-type: none"> ▪ The IS-IS instances bound with the same non-default VRF table must be configured with different system IDs. The IS-IS instances bound with different VRF tables can be configured with the same system ID. ▪ One IS-IS instance can be bound with only one VRF table, but one VRF table can be bound to multiple IS-IS instances. ▪ When the VRF table bound to an IS-IS instance is changed, all IS-IS interfaces associated with the instance will be deleted. That is, the ip (or ipv6) router isis [tag] interface configuration and the redistribution configuration in routing process configuration mode will be deleted.

Configuration

Example

Configuring IS-IS VRF

Configuration Steps	<p>Bind an IS-IS instance to a VRF table.</p> <p>Add interfaces to the VRF table and IS-IS instance. (Omitted)</p>
	<pre>A(config)#vrf definition vrf_1 A(config-vrf)#address-family ipv4 A(config-vrf-af)#exit-address-family A(config)# router isis A(config-router)# vrf vrf_1</pre>
Verification	<p>Check whether the local device establishes neighbor relationships with other devices specified in the VRF table.</p>
	<pre>A# show isis neighbors</pre>

Common

Errors

- Interfaces are not added to the VRF table.
- The IP addresses of the interfaces connected between neighbors are not in the same network segment.
- The **ip router isis** command is not executed on interfaces.
- No NET address is configured, or different NET addresses exist at Level-1.
- **max-area-addresses** is configured differently on both sides.
- **metric-style** is configured differently on both sides.
- The interface Levels on both sides are different. One side is Level-1, whereas the other side is Level-2.
- One side is configured with the P2P mode, whereas the other side is configured with the broadcast mode.
- One side is enabled with authentication, whereas the other side is not.

6.4.12 Configuring IS-IS MTR

Configuration

Effect

- If the **multi-topology** command is not executed, IPv4 and IPv6 share one IS-IS physical topology, also called the default topology. If the **multi-topology** command is executed without the **transition** parameter, routing devices run in MT mode. IS-ISv4 runs in the default topology, and IS-ISv6 runs in the IPv6 unicast topology. If the **multi-topology** command is executed with the **transition** parameter, routing devices run in MTT mode. IS-ISv6 runs in the default topology and IPv6 unicast topology. The three configurations are mutually exclusive. The routing devices in MTT mode can transfer the MT TLV or the default topology TLV. The MTT mode is applicable to incremental deployment to ensure smooth network migration. The MTT mode can cause route leaking between the default topology and IPv6 unicast topology. If the MTT mode is configured improperly, network failures such as routing black holes and loops may occur.

Notes

Note the following constraints or conventions when you configure the IS-IS MTR feature:

- Set **metric-style** to **Wide** or **Transition** before you run the **multi-topology** command.
- The MTR feature will be disabled if **metric-style** is set to **Narrow** or only one Level is configured to support the **Wide** or **Transition** mode.

Configuration

Steps

- Perform this configuration based on requirements.
- Configure the MTR feature in IS-IS address-family ipv6 configuration mode on the desired device, unless otherwise specified.

Verification

- Check whether the local device establishes neighbor relationships with other devices.

Related Commands

Configuring IS-IS MTR

Command	multi-topology [<i>transition</i>]
Parameter Description	<i>transition</i> : Configures the MTT mode, which supports smooth migration from an IPv4-IPv6 hybrid topology to separate IPv4 and IPv6 topologies.
Command Mode	IS-IS address-family ipv6 configuration mode
Usage Guide	<p>If the multi-topology command is not executed, IPv4 and IPv6 share one IS-IS physical topology, also called the default topology. If the multi-topology command is executed without the transition parameter, routing devices run in MT mode. IS-ISv4 runs in the default topology, and IS-ISv6 runs in the IPv6 unicast topology. If the multi-topology command is executed with the transition parameter, routing devices run in MTT mode. IS-ISv6 runs in the default topology and IPv6 unicast topology. The three configurations are mutually exclusive. The routing devices in MTT mode can transfer the MT TLV or the default topology TLV. The MTT mode is applicable to incremental deployment to ensure smooth network migration. The MTT mode can cause route leaking between the default topology and IPv6 unicast topology. If the MTT mode is configured improperly, network failures such as routing black holes and loops may occur.</p> <p>Set metric-style to Wide or Transition before you run the command. The MTR feature will be disabled if metric-style is set to Narrow or only one Level is configured to support the Wide or Transition mode.</p>

Configuration

Example

Configuring IS-IS MTR

Configuration Requirements	<p>The typical application scenario of MTR is to retain devices that only support IPv4 services in a network where IPv6 service extension will be performed.</p> <p>In Figure 1-20, Router 2 only supports the IPv4 protocol stack but does not support the MTR feature; therefore, it can only run IPv4 services. The network capacity needs to be scaled to support IPv6 services in order to meet service extension requirements. (Router 1, Router 3, and Router 4 that support the MTR feature will</p>
----------------------------	--


be added.) The device (Router 2) that supports only one protocol stack must be replaced to maintain the stability of the network running IPv4 and IPv6 dual protocol stacks; otherwise, IPv6 routing black holes may occur.

If you need to retain Router 2, you can configure the MTR feature on Router 1, Router 3, and Router 4. The MTR feature enables Router 2 to continue to run IPv4 services without interference on the IPv4 and IPv6 services on Router 1, Router 3, and Router 4. The MTR feature improves networking flexibility, indirectly prolongs the service life of old devices, and meets service extension requirements while maximizing the values of old devices.

The configuration requirements are as follows:

- Retain Router 2, which only supports IPv4 services.
- Add devices that support IPv4 and IPv6 dual topologies, and separate IPv4 route calculation and IPv6 route calculation based on different topologies.

Figure 6-18
IS-IS MTR
Topology



Router 1

Configure IS-IS and Ethernet interfaces.

Configure IS-IS:

```
QTECH(config)# router isis
QTECH(config-router)# net 49.0001.0000.0000.0001.00
QTECH(config-router)# is-type level-1
QTECH(config-router)# metric-style wide
QTECH(config-router)# address-family ipv6
QTECH(config-router-af)# multi-topology
```

Configure Ethernet interfaces:

```
QTECH(config)# interface gigabitEthernet 0/1
```

	<pre> QTECH(config-if-GigabitEthernet 0/1)# ipv6 enable QTECH(config-if-GigabitEthernet 0/1)# ipv6 address 1002: : 1/112 QTECH(config-if-GigabitEthernet 0/1)# ipv6 router isis QTECH(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0 QTECH(config-if-GigabitEthernet 0/1)# ip router isis QTECH(config-if-GigabitEthernet 0/1)# interface gigabitEthernet 0/2 QTECH(config-if-GigabitEthernet 0/2)# ipv6 enable QTECH(config-if-GigabitEthernet 0/2)# ipv6 address 1003: : 1/112 QTECH(config-if-GigabitEthernet 0/2)# ipv6 router isis QTECH(config-if-GigabitEthernet 0/2)# ip address 192.168.2.1 255.255.255.0 QTECH(config-if-GigabitEthernet 0/2)# ip router isis QTECH(config-if-GigabitEthernet 0/2)#isis wide-metric 11 </pre>
Router 2	Configure IS-IS and Ethernet interfaces.
	<pre> Configure IS-IS: QTECH(config)# router isis QTECH(config-router)# net 49.0001. 0000.0000.0002.00 QTECH(config-router)# is-type level-1 QTECH(config-router)# metric-style wide QTECH(config-router)#address-family ipv6 QTECH(config-router-af)#no adjacency-check Configure Ethernet interfaces: QTECH(config)# interface gigabitEthernet 0/1 QTECH(config-if-GigabitEthernet 0/1)# ip address 192.168.1.2 255.255.255.0 QTECH(config-if-GigabitEthernet 0/1)# ip router isis QTECH(config-if-GigabitEthernet 0/1)# interface gigabitEthernet 0/2 QTECH(config-if-GigabitEthernet 0/2)# ip address 192.168.3.2 255.255.255.0 QTECH(config-if-GigabitEthernet 0/2)# ip router isis </pre>
Router 3	Configure IS-IS and Ethernet interfaces.
	Configure IS-IS:

```

QTECH(config)# router isis
QTECH(config-router)# net 49.0001.0000.0000.0003.00
QTECH(config-router)# is-type level-1
QTECH(config-router)# metric-style wide
QTECH(config-router)# address-family ipv6
QTECH(config-router-af)# multi-topology

Configure Ethernet interfaces:

QTECH(config)# interface gigabitEthernet 0/1
QTECH(config-if-GigabitEthernet 0/1)# ipv6 enable
QTECH(config-if-GigabitEthernet 0/1)# ipv6 address 3001: : 1/112
QTECH(config-if-GigabitEthernet 0/1)# ipv6 router isis
QTECH(config-if-GigabitEthernet 0/1)# ip address 192.168.2.3 255.255.255.0
QTECH(config-if-GigabitEthernet 0/1)# ip router isis
QTECH(config-if-GigabitEthernet 0/1)#isis wide-metric 11
QTECH(config-if-GigabitEthernet 0/1)# interface gigabitEthernet 0/2
QTECH(config-if-GigabitEthernet 0/2)# ipv6 enable
QTECH(config-if-GigabitEthernet 0/2)# ipv6 address 3004: : 1/112
QTECH(config-if-GigabitEthernet 0/2)# ipv6 router isis
QTECH(config-if-GigabitEthernet 0/2)# ip address 192.168.4.3 255.255.255.0
QTECH(config-if-GigabitEthernet 0/2)# ip router isis
QTECH(config-if-GigabitEthernet 0/2)#isis wide-metric 12

```

Router 4

Configure IS-IS and Ethernet interfaces.

```

Configure IS-IS:
QTECH(config)# router isis
QTECH(config-router)# net 49.0001.0000.0000.0004.00
QTECH(config-router)# is-type level-1
QTECH(config-router)# metric-style wide
QTECH(config-router)# address-family ipv6
QTECH(config-router-af)# multi-topology

Configure Ethernet interfaces:

```

	<pre> QTECH(config)# interface gigabitEthernet 0/1 QTECH(config-if-GigabitEthernet 0/1)# ipv6 enable QTECH(config-if-GigabitEthernet 0/1)# ipv6 address 4002: : 1/112 QTECH(config-if-GigabitEthernet 0/1)# ipv6 router isis QTECH(config-if-GigabitEthernet 0/1)# ip address 192.168.3.4 255.255.255.0 QTECH(config-if-GigabitEthernet 0/1)# ip router isis QTECH(config-if-GigabitEthernet 0/1)# interface gigabitEthernet 0/2 QTECH(config-if-GigabitEthernet 0/2)# ipv6 enable QTECH(config-if-GigabitEthernet 0/2)# ipv6 address 4003: : 1/112 QTECH(config-if-GigabitEthernet 0/2)# ipv6 router isis QTECH(config-if-GigabitEthernet 0/2)# ip address 192.168.4.4 255.255.255.0 QTECH(config-if-GigabitEthernet 0/2)# ip router isis </pre>
Verification	<ul style="list-style-type: none"> ▪ Run the show command on Router 1 to check whether the next hop of the IPv4 route destined for Router 4 is Router 2. ▪ Run the show command on Router 1 to check whether the next hop of the IPv6 route destined for Router 4 is Router 3.
Checking the IPv4 route	<pre> QTECH#show ip route Codes: C - connected, S - static, R - RIP, B - BGP O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default Gateway of last resort is no set C 192.168.1.0/24 is directly connected, GigabitEthernet 0/1 C 192.168.1.1/32 is local host. C 192.168.2.0/24 is directly connected, GigabitEthernet 0/2 C 192.168.2.1/32 is local host. i L1 192.168.3.0/24 [115/20] via 192.168.1.2, 00:13:14, GigabitEthernet 0/1 i L1 192.168.4.0/24 [115/23] via 192.168.2.3, 00:02:40, GigabitEthernet 0/2 </pre>
Checking the IPv6 route	<pre> QTECH#show ipv6 route IPv6 routing table name is - Default - 16 entries </pre>


```

Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
      I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
      O - OSPF intra area, OI - OSPF inter area, OE1 - OSPF external type 1, OE2 - OSPF
external type 2
      ON1 - OSPF NSSA external type 1, ON2 - OSPF NSSA external type 2
L   ::1/128 via Loopback, local host
C   1002::/112 via GigabitEthernet 0/1, directly connected
L   1002::1/128 via GigabitEthernet 0/1, local host
C   1003::/112 via GigabitEthernet 0/2, directly connected
L   1003::1/128 via GigabitEthernet 0/2, local host
I1  3001::/112 [115/21] via FE80::C806:5FF:FEE8:38, GigabitEthernet 0/2
I1  3004::/112 [115/21] via FE80::C806:5FF:FEE8:38, GigabitEthernet 0/2
I1  4002::/112 [115/31] via FE80::C806:5FF:FEE8:38, GigabitEthernet 0/2
I1  4003::/112 [115/31] via FE80::C806:5FF:FEE8:38, GigabitEthernet 0/2
L   FE80::/10 via ::1, Null0
C   FE80::/64 via GigabitEthernet 0/2, directly connected
L   FE80::1614:4BFF:FE12:ADFC/128 via GigabitEthernet 0/2, local host
C   FE80::/64 via GigabitEthernet 0/1, directly connected
L   FE80::1614:4BFF:FE12:ADFD/128 via GigabitEthernet 0/1, local host
C   FE80::/64 via Local 0, directly connected
L   FE80::1614:4BFF:FE12:ADFC/128 via Local 0, local host

```

Common Errors

- metric-style is not set to Wide or Transition.
- The protocol types used by two neighbors do not match; therefore, a neighbor relationship cannot be established.
- The IP addresses of the interfaces connected between neighbors are not in the same network segment.
- The **ip router isis** command is not executed on interfaces.
- No NET address is configured, or different NET addresses exist at Level 1.
- **max-area-addresses** is configured differently on both sides.
- **metric-style** is configured differently on both sides.
- The interface Levels on both sides are different. One side is Level-1, whereas the other side is Level-2.

- One side is configured with the P2P mode, whereas the other side is configured with the broadcast mode.
- One side is enabled with authentication, whereas the other side is not.

6.4.13 Configuring SNMP for IS-IS

Configuration

Effect

- By default, the SNMP software can perform the MIB operation on the first IS-IS instance. To perform the MIB operation on other instances, you need to manually specify these instances.

Notes

- By default, the SNMP software can perform the MIB operation on the first displayed IS-IS instance.

Configuration

Steps

Binding the Instances on Which the IS-IS MIB Operation Will Be Performed

- Perform this configuration based on requirements.
- Run the **enable mib-binding** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Enabling IS-IS Trap Globally

- Perform this configuration based on requirements.
- Run the **snmp-server enable traps isis** command in global configuration mode on the desired device, unless otherwise specified.

Configuring an SNMP Host Globally

- Perform this configuration based on requirements.
- Run the **snmp-server host** command in global configuration mode on the desired device, unless otherwise specified.

Allowing the Sending of all IS-IS Trap Messages to the SNMP Host

- Perform this configuration based on requirements.
- Run the **enable traps all** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- Use the MIB tool to read and write IS-IS settings.

Related

Commands

Binding the Instances on Which the IS-IS MIB Operation Will Be Performed

Command	enable mib-binding
Parameter Description	N/A
Command Mode	IS-IS routing process configuration mode
Usage Guide	The latest standards stipulate that the MIB operation can be performed on a single instance. By default, the MIB operation is performed on the first displayed IS-IS instance. Because multiple IS-IS instances can be configured, the administrator can use this command to specify the instances on which the MIB operation will be performed.

Enabling IS-IS Trap Globally

Command	snmp-server enable traps [isis]
Parameter Description	<i>isis</i> : Enables IS-IS event trap.
Command Mode	Global configuration mode
Usage Guide	This command must be used with the snmp-server host command in global configuration mode so that trap messages can be sent.

Configuring an SNMP Host Globally

Command	snmp-server host { <i>host-addr</i> ipv6 <i>ipv6-addr</i> } [vrf <i>vrfname</i>] [traps] [version { 1 2c 3 { auth noauth priv } }] <i>community-string</i> [udp-port <i>port-num</i>] [<i>notification-type</i>]
Parameter Description	<i>host-addr</i> : Indicates the address of the SNMP host. <i>ipv6-addr</i> : Indicates the IPv6 address of the SNMP host. <i>vrfname</i> : Indicates the name of a VRF table. version : Indicates the SNMP version, which can be set to V1 , V2C , or V3 auth noauth priv : Indicates the security level of V3 users. <i>community-string</i> : Indicates the community string or user name (V3 version).

	<p><i>port-num</i>: Indicates the port number of the SNMP host.</p> <p><i>notification-type</i>: Indicates the type of trap messages that are actively sent, for example, snmp.</p>
Command Mode	Global configuration mode
Usage Guide	This command is used with the snmp-server enable traps command to actively send trap messages to a Network Management System (NMS). You can configure different SNMP hosts to receive trap messages. A host supports different trap types, ports, and VRF tables. For the same host (with the same port configuration and VRF configuration), the last configuration is combined with the previous configurations. That is, to send different trap messages to the same host, configure a type of trap messages each time. These configurations are finally combined.

Allowing the Sending of Trap Messages

Command	enable traps { all <i>traps set</i> }
Parameter Description	<p>all: Indicates all trap messages.</p> <p><i>traps set</i>: Indicates a trap message type in any set.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	IS-IS packets are classified into 18 types of trap messages, which are grouped into several sets, with each set containing several trap message types. To enable the sending of IS-IS trap messages, run the snmp-server enable traps isis command in global configuration mode and specify the recipient host and the type of trap messages that can be sent.

Configuration

Example

Configuring IS-IS SNMP

Configurati on Steps	<ul style="list-style-type: none"> ▪ Bind the instances on which the IS-IS MIB operation will be performed. ▪ Complete trap message-related settings.
	<pre>A(config)# router isis A(config-router)# enable mib-binding A# configure terminal A(config)#snmp-server enable traps isis A(config)#snmp-server host 10.1.1.1 traps version 2c public A(config)#router isis A(config-router)# enable traps all</pre>
Verification	Run the MIB tool to read and write IS-IS settings.
	<pre>A# show running-config</pre>

6.4.14 Configuring IS-IS Two-way Maintenance

Configuration

Effect

- Enable IS-IS two-way maintenance.

Notes

- The IS-IS basic functions must be configured.
- The neighbor relationship is successfully established.

Configuration

Steps

Configuring IS-IS Two-way Maintenance

- Configure the two-way maintenance function as required.
- Run the **two-way-maintain** command in IS-IS routing process configuration mode on the required devices unless otherwise specified.

Verification

- Run the **show isis neighbor** command to check the neighbor update time.
- Run the **show isis protocol** command to check whether the two-way maintenance function is enabled.

Related

Commands

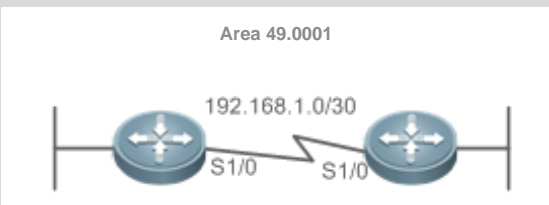
Configuring IS-IS Two-way Maintenance

Command	two-way-maintain
Parameter Description	N/A
Command Mode	IS-IS routing process configuration mode
Usage Guide	In a large-scale network, a large number of packets are sent and received, which occupies lots of CPU and memory resources, causing some IS-IS packets to be delayed or discarded. If the time required for processing HELLO packets exceeds the neighbor relationship maintenance time, the corresponding neighbor relationship times out and is removed. When the two-way maintenance function is enabled, if a large number of packets exist on the network, the LSP packets, CSNP packets, and PSNP packets from a neighbor in addition to HELLO packets can also be used to maintain the two-way relationship with the neighbor, preventing neighbor failure caused by delay or discard of HELLO packets.

Configuration Example

- i** The following example is implemented based on IS-IS basic functions. For details about the IS-IS basic functions, see preceding description

Configuring IS-IS Two-way Maintenance

Scenario	 <p>The diagram shows two routers connected via their S1/0 interfaces. They are in Area 49.0001 and share the IP address range 192.168.1.0/30.</p>
Configurati on Steps	<ul style="list-style-type: none"> Configure the ISIS basic functions on all devices. The neighbor relationship is successfully established.
Verification	<p>Verify that an ISIS interface neighbor is established on device A.</p> <p>Check the status of an ISIS instance on device A.</p>

Common Errors

6.4.15 Configuring Other IS-IS Parameters

Configuration

Effect

- **maximum-paths:** Configures the maximum number of IS-IS equal-cost paths to be installed to a routing table.
- **isp-length receive:** Configures the maximum length allowed for received LSPs.
- **isp-length originate:** Configures the maximum length allowed for sent LSPs.
- **passive-interface:** Prevents passive interfaces from receiving and sending IS-IS packets. That is, IS-IS neighbor relationships will not be established on passive interfaces. The IP addresses of passive interfaces are flooded through other interfaces.
- **bandwidth-reference:** Configures the bandwidth reference for IS-IS to automatically calculate the link metric. If a metric is configured, the metric prevails. If no metric is configured, the metric is automatically calculated based on the configured bandwidth reference.
- **interface-protocol-compatible:** Configures the TLV field of the IS-IS protocol to be padded based on protocols supported by an interface.
- **is-name:** Replaces the system ID of an instance with the configured name.
- **isis metric:** Stores the metric, which is used in SPF calculation, in the IP reachability information TLV. The greater the metric, the greater the routing consumption of the interface and the longer the path obtained by SPF calculation.
- **isis passive:** Prevents the interface from sending or receiving IS-IS packets, but the corresponding IP address will be flooded through other interfaces.
- **isis priority:** In a broadcast network, IS-IS needs to elect a DIS among all devices. The DIS will generate a pseudonode and related LSPs. The device with the highest priority is elected as the DIS. You can configure different priorities for different Levels.
- **default-information originate:** Generates a Level-2 default route, which will be advertised through LSPs.
- **spf-interval:** Configures the exponential backoff algorithm of SPF.
- **summary-address** and **summary-prefix:** Creates a summary route to represent a group of routes in a routing table. A summary route can include multiple routes of the specified Level. The interface metric of the summary route follows the smallest interface metric among all routes.
- **log-adjacency-changes:** Enables neighbor relationship event output to log IS-IS neighbor relationship changes.
- **redistribute:** Redistributes other routes to IS-IS; redistributes Level-1 routes to Level-2; redistributes Level-2 routes to Level-1.

Configuration

Steps

Configuring the Maximum Number of Equal-Cost Paths

- Perform this configuration based on requirements.
- Run the **maximum-paths** command in IS-IS routing process configuration mode or IS-IS address-family ipv6 configuration mode on the desired device, unless otherwise specified.

Configuring the Maximum Length Allowed for Received LSPs

- Perform this configuration based on requirements.
- Run the **lsp-length receive** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the Maximum Length Allowed for Sent LSPs

- Perform this configuration based on requirements.
- Run the **lsp-length originate** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring a Passive Interface

- Perform this configuration based on requirements.
- Run the **passive-interface** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the IS-IS Interface Bandwidth Reference

- Perform this configuration based on requirements.
- Run the **bandwidth-reference** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the TLV Field of the IS-IS Protocol to Be Padded Based on Protocols Supported by an Interface

- Perform this configuration based on requirements.
- Run the **interfaces-protocol-compatible** command in IS-IS routing process configuration mode on the desired devices, unless otherwise specified.

Configuring the IS-IS Instance Alias

- Perform this configuration based on requirements.
- Run the **is-name** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring the IS-IS Interface Metric

- Perform this configuration based on requirements.
- Run the **isis metric** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring an Interface as a Passive Interface

- Perform this configuration based on requirements.

- Run the **isis passive** command in interface configuration mode on the desired device, unless otherwise specified.

Configuring the Priority of the DIS

- Perform this configuration based on requirements.
- Run the **isis priority** command in interface configuration mode on the desired device, unless otherwise specified.

Generating a Default Route

- Perform this configuration based on requirements.
- Run the **default-information originate** command in IS-IS routing process configuration mode or IS-IS address-family ipv6 configuration mode on the desired device, unless otherwise specified.

Configuring the SPF Calculation Cycle

- Perform this configuration based on requirements.
- Run the **spf-interval** command in interface configuration mode on the desired device, unless otherwise specified.

Configure a Summary Route

- Perform this configuration based on requirements.
- Run the **summary-address** and **summary-prefix** commands in IS-IS routing process configuration mode or IS-IS address-family ipv6 configuration mode on the desired device, unless otherwise specified.

Enabling Neighbor Relationship Event Output

- Perform this configuration based on requirements.
- Run the **log-adjacency-changes** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Configuring Route Redistribution

- Perform this configuration based on requirements.
- Run the **redistribute** command in IS-IS routing process configuration mode on the desired device, unless otherwise specified.

Verification

- **maximum-paths:** Check whether the maximum number of equal-cost paths displayed by routing entries is the same as the configuration.
- **lsp-length receive:** Capture packets to check the length of LSPs.
- **lsp-length originate:** Capture packets to check the length of LSPs.
- **passive-interface:** Capture packets to check whether the interface receives and sends IS-IS packets.
- **bandwidth-reference:** Check the metric of the routing table.
- **interfaces-protocol-compatible:** Check the neighbor relationship establishment status.

- **is-name:** Check the database details of IS-IS
- **isis metric:** Check the database details of IS-IS.
- **isis priority:** Check whether the device with the changed priority setting is elected as the DIS.
- **default-information originate:** Check whether a default route is generated.
- **spf-interval:** Check whether the SPF calculation cycle works.
- **summary-address** and **summary-prefix:** Capture packets to check whether the summary route instead of detailed routes is advertised through LSPs.
- **log-adjacency-changes:** Change the neighbor state and verify that the change is recorded when debugging is disabled.
- **redistribute:** Check IS-IS routing entries.

Related

Commands

Configuring the Maximum Number of Equal-Cost Paths

Command	maximum-paths <i>maximum</i>
Parameter Description	<i>maximum</i> : Indicates the maximum number of IS-IS equal-cost routes to be installed to a routing table. The value range is 1 to 32.
Command Mode	IS-IS routing process configuration mode and IS-IS address-family ipv6 configuration mode
Usage Guide	This command is used by IS-IS to control the number of IS-IS equal-cost paths to be installed to a routing table. The routing table also has a command used to control the number of equal-cost paths. The number of effective equal-cost paths is determined by either of the two command values, whichever is smaller.

Configuring the Maximum Length Allowed for Received LSPs

Command	lsp-length receive <i>size</i>
Parameter Description	<i>size</i> : Indicates the maximum length allowed for received LSPs. According to RFC, the value range is 1,492 to 16,000, in the unit of bytes.
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to control the maximum length allowed for LSPs received by the local device. Intermediate nodes with sufficient memory are required to receive LSPs whose maximum length is equal to the interface MTU in order to avoid a route convergence

failure. From this perspective, the command is meaningless. The maximum length allowed for received LSPs cannot be smaller than that allowed for sent LSPs; otherwise, the former will be automatically adjusted to be equal to the latter.

Configuring the Maximum Length Allowed for Sent LSPs

Command	<code>lsp-length originate size [level-1 level-2]</code>
Parameter Description	<p><i>size</i>: Indicates the maximum length allowed for sent LSPs. The value range is 512 to 16,000, in the unit of bytes.</p> <p>level-1: Applies the setting only to Level-1 LSPs.</p> <p>level-2: Applies the setting only to Level-2 LSPs.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	In principle, the maximum length of LSPs and SNPs cannot be greater than the interface MTU; otherwise, the packets will be discarded when being sent.

Configuring a Passive Interface

Command	<code>passive-interface [default] { interface-type interface-number }</code>
Parameter Description	<p>default: Configures all IS-IS interfaces that are not enabled as passive interfaces.</p> <p><i>interface-type</i>: Indicates the interface type.</p> <p><i>interface-number</i>: Indicates the interface number.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>This command prevents the specified interface from receiving and sending IS-IS packets, but the IP address of the interface will be flooded by other interfaces.</p> <p>If the default option is selected and there are more than 255 IS-IS interfaces not enabled, only the first 255 interfaces will be configured as passive interfaces. The remaining interfaces are non-passive interfaces.</p>

Configuring the Bandwidth Reference

Command	<code>bandwidth-reference value</code>
---------	--

Parameter Description	<i>value</i> : Indicates the bandwidth value. The value ranges from 1 to 4294967 , in the unit of Mbps. The default value is 100 .
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to calculate the interface metric in an instance according to the bandwidth reference configured in the instance. If a metric is configured, the metric prevails.

Configuring the TLV Field of the IS-IS Protocol to Be Padded Based on Protocols Supported by an Interface

Command	interfaces-protocol-compatible
Parameter Description	N/A
Command Mode	IS-IS routing process configuration mode
Usage Guide	<p>When devices supporting IS-IS for IPv4 or IS-IS for IPv6 and those supporting both IS-IS for IPv4 and IS-IS for IPv6 on a network are configured to establish a neighbor relationship, the calculated routes may be unreachable, resulting in routing black holes. To prevent routing black holes, the TLV (#129) field of QTECH IS-IS protocol is padded based on protocols supported by an instance, that is, IS-IS instances that establish a neighbor relationship need to support the same protocol.</p> <p>The TLV (#129) field of the IS-IS protocol supported by devices of some vendors is padded based on protocols supported by an interface. In single topology mode, a QTECH device connects to a device of another vendor, a loopback interface is configured on both devices, both IS-IS for IPv4 and IS-IS for IPv6 are enabled, and only a single protocol stack (IS-IS for IPv4 or IS-IS for IPv6) is configured on the interface for interconnecting the two devices. The device of the other vendor sends TLV#129 (supporting only a single protocol stack) based on interfaces whereas QTECH device sends TLV#129 (supporting dual protocol stacks) based on instances. As a result, the displayed neighbor status on the device of the other vendor is "Init" but QTECH device cannot establish a neighbor relationship with the device of the other vendor.</p> <p>For this, the interfaces-protocol-compatible command can be configured to enable QTECH devices to establish a neighbor relationship with devices of other vendors.</p>

Configuring an IS Name

Command	is-name name
Parameter Description	<i>name</i> : Indicates the name for replacing the system ID. The name contains a maximum of 64 characters.
Command Mode	IS-IS routing process configuration mode
Usage Guide	Use this command to replace the system ID of an instance with the configured name. All system IDs displayed in the command output of show isis database and show isis neighbors are replaced with the configured name.

Configuring the IS-IS Interface Metric

Command	isis metric <i>metric</i> [level-1 level-2]
Parameter Description	<i>metric</i> : Indicates the metric value. The value range is 1 to 63. The default value is 10. level-1 : Applies the setting to Level-1 circuits. level-2 : Applies the setting to Level-2 circuits.
Command Mode	Interface configuration mode
Usage Guide	The metric, which is used in SPF calculation, is stored in the IP reachability information TLV. The greater the metric, the greater the routing consumption of the interface and the longer the path obtained by SPF calculation. The metric belongs to the narrow type and is valid only when metric-style is set to Narrow.

Configuring the Wide Metric of an Interface

Command	isis wide-metric <i>metric</i> [level-1 level-2]
Parameter Description	<i>metric</i> : Indicates the metric value. The value range is 1 to 16,777,214. The default value is 10. level-1 : Applies the setting to Level-1 circuits. level-2 : Applies the setting to Level-2 circuits.

Command Mode	Interface configuration mode
Usage Guide	The metric, which is used in SPF calculation, is stored in the IP reachability information TLV. The greater the metric, the greater the routing consumption of the interface and the longer the path obtained by SPF calculation. The metric is valid only when metric-style is set to Wide.

Configuring an Interface as a Passive Interface

Command	isis passive
Parameter Description	-
Command Mode	Interface configuration mode
Usage Guide	This command prevents the interface from sending or receiving ISIS packets, but the corresponding IP address will be flooded through other interfaces. It only takes effect on the generated ISIS interface.

Configuring the Priority of the DIS

Command	isis priority <i>value</i> [level-1 level-2]
Parameter Description	<i>value</i> : Indicates the priority. The value range is 0 to 127. The default value is 64. level-1: Applies the setting to Level-1 circuits. level-2: Applies the setting to Level-2 circuits.
Command Mode	Interface configuration mode
Usage Guide	Use this command to change the priority carried in Hello packets in a LAN. The device with a lower priority is less likely to be elected as the DIS. The command is invalid on a P2P network interface. The no isis priority command, with or without parameters, restores the priority to its default value. To change the configured priority, run the isis priority command with the

priority specified to overwrite the existing configuration, or you can first restore the priority to its default value and then configure a new priority.

Generating a Default Route

Command	default-information originate [route-map <i>map-name</i>]
Parameter Description	route-map <i>map-name</i> : Associates with a route map.
Command Mode	IS-IS routing process configuration mode and IS-IS address-family ipv6 configuration mode
Usage Guide	Because Level-2 domains do not generate any default route, use this command to allow a default route to enter a Level-2 domain.

Configuring SPF Calculation Cycle

Command	spf-interval [level-1 level-2] <i>maximum-interval</i> [<i>initial-interval hold-interval</i>]
Parameter Description	level-1 : Applies the configuration only to Level-1. level-2 : Applies the configuration only to Level-2. <i>maximum-interval</i> : Indicates the maximum interval for performing two consecutive SPF calculations. The value range is 1 to 120 (in seconds). The default value is 10 . <i>initial-interval</i> : Indicates the waiting time for performing the SPF calculation for the first time. The value range is 0 to 60000 (in milliseconds). The default value is 50 . <i>hold-interval</i> : Indicates the minimum interval for performing the SPF calculation for the second time. The value range is 10 to 60000 (in milliseconds). The default value is 200 .
Command Mode	IS-IS routing process configuration mode
Usage Guide	Increasing the maximum interval for performing SPF calculations can avoid frequent SPF calculations and waste of CPU resources. However, a larger minimum interval also leads to slower responses to route changes.
	The waiting time for performing the SPF calculation for the first time is the initial interval. If the network becomes unstable, the SPF calculation interval is less than the maximum interval, and the interval for performing the SPF calculation for the second time becomes the hold interval. A corresponding penalty is added to this interval: The next interval for

the SPF calculation doubles the previous interval for the same SPF calculation, until the SPF calculation interval reaches the maximum interval. Subsequent SPF calculations are performed at the maximum interval. When the network becomes stable, the interval for performing the SPF calculation becomes greater than the maximum interval, and the waiting time for performing the SPF calculation is restored to the initial interval.

Link changes have high requirements for convergence. The initial interval can be set to a small value. The preceding parameters can also be adjusted to larger values to reduce CPU consumption.

The value of **initial-interval** cannot be greater than that of **maximum-interval**. Otherwise, the value of **initial-interval** will be used as the value of **maximum-interval**.

The value of **hold-interval** cannot be greater than that of **maximum-interval**. Otherwise, the value of **hold-interval** will be used as the value of **maximum-interval**.

The value of **initial-interval** cannot be greater than that of **hold-interval**. Otherwise, the value of **initial-interval** will be used as the value of **hold-interval**.

Configuring an IPv4 Summary Route

Command	<code>summary-address ip-address net-mask [level-1 level-2 level-1-2] [metric number]</code>
Parameter Description	<p><i>ip-address</i>: Indicates the IP address of the summary route.</p> <p><i>net-mask</i>: Indicates the subnet mask of the summary route.</p> <p>level-1: Applies the setting only to Level-1.</p> <p>level-2: Applies the setting only to Level-2. By default, the setting takes effect for Level-2.</p> <p>level-1-2: Applies the setting to Level-1 and Level-2.</p> <p><i>number</i>: Indicates the metric of the summary route.</p>
Command Mode	IS-IS routing process configuration mode
Usage Guide	If the configured summary route contains routing information about a reachable address or network segment, the summary route, instead of detailed routes, is advertised externally.

Configuring an IPv6 Summary Route

Command	<code>summary-prefix ipv6-prefix/prefix-length [level-1 level-2 level-1-2]</code>
---------	---

Parameter Description	<p><i>ipv6-prefix/prefix-length</i>: Indicates the network address of the summary route and its IPv6 prefix length. The address format is X:X:X:X::X/<0-128>.</p> <p>level-1: Applies the setting only to Level-1.</p> <p>level-2: Applies the setting only to Level-2. By default, the setting takes effect for Level-2.</p> <p>level-1-2: Applies the setting to Level-1 and Level-2.</p>
Command Mode	IS-IS address-family ipv6 configuration mode
Usage Guide	If the configured summary route contains routing information about a reachable address or network segment, the summary route, instead of detailed routes, is advertised externally.

Enabling Neighbor Relationship Event Output

Command	log-adjacency-changes
Parameter Description	N/A
Command Mode	IS-IS routing process configuration mode
Usage Guide	You can also use the debug command to record IS neighbor state changes, but the command consumes many system resources.

Redistributing Other Routes to IS-IS

Command	redistribute { bgp ospf <i>process-id</i> [match { internal [external [1 2]] [nssa-external [1 2]] external [1 2] [internal] [nssa-external [1 2]] nssa-external [1 2] [internal] [external [1 2]] } } rip connected static } [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [route-map <i>map-tag</i>] [level-1 level-1-2 level-2]
Parameter Description	<p><i>process-id</i>: Indicates the OSPF process ID. The range is 1 to 65,535.</p> <p>match { internal external [1 2] nssa-external [1 2] }: When OSPF routes are redistributed, the routes are filtered by subtype. If the match option is not selected, routes of all OSPF types will be received. If match external is not followed by the number 1 or 2, OSPF routes specified by external 1 and external 2 will be redistributed. If match</p>

	<p>nssa-external is not followed by the number 1 or 2, OSPF routes specified by nssa-external 1 and nssa-external 2 will be redistributed.</p> <p>metric <i>metric-value</i>: Indicates the metric of redistributed routes. The value range is 0 to 4,261,412,864. The metric of external routes is used when the metric option is not specified.</p> <p>metric-type { internal external }: Indicates the metric type of redistributed routes. internal: Indicates that the metric belongs to the internal type. external: Indicates that the metric belongs to the external type. If metric-type is not specified, the metric belongs to the internal type.</p> <p>route-map <i>map-tag</i>: Indicates the route map used for external route redistribution. It is used to filter redistributed routes or configure the attributes of redistributed routes. The value of <i>map-tag</i> cannot exceed 32 characters. By default, route-map is not configured.</p> <p>level-1 level-1-2 level-2: Indicates the Level of redistributed routes received by IS-IS. If no Level is specified, routes are redistributed to Level-2. level-1: Redistributes routes to Level-1. level-1-2: Redistributes routes to Level-1 and Level-2. level-2: Redistributes routes to Level-2.</p>
Command Mode	IS-IS routing process configuration mode and IS-IS address-family ipv6 configuration mode
Usage Guide	<p>The no redistribute { bgp ospf process-id rip connected static } command is used to cancel the redistribution of routes mapped to the specified protocol. If no redistribute is followed by other parameters, the command will restore the default parameter settings, rather than cancel route redistribution. For example, no redistribute bgp cancels BGP route redistribution, whereas no redistribute bgp route-map aa cancels the route map named aa used for BGP route redistribution.</p> <p>When external routes are redistributed in IPv4 mode, the routing information is stored in LSPs' IP External Reachability Information TLV.</p> <p>When external routes are redistributed in IPv6 mode, the routing information is stored in LSPs' IPv6 Reachable TLV.</p> <p>In the old versions of some vendors, if metric-type is set to external, the metric of redistributed routes is added by 64 during route calculation and used to determine routing. This practice does not comply with the related protocol. In the actual application, external routes may be preferred over internal routes. If this happens during interworking with old versions of some vendors, you can modify the related setting (such as metric or metric-type) of each device to ensure that internal routes are preferred over external routes.</p>

Redistributing the Level-1 Reachable Routing Information of the Specified IS-IS Instance to Level-2 of the Current Instance

Command	redistribute isis [tag] level-1 into level-2 [route-map route-map-name / distribute-list access-list-name]
Parameter Description	<p><i>tag</i>: Indicates the name of the IS-IS instance whose routing information will be redistributed.</p> <p>route-map route-map-name: Indicates the route map used for rout redistribution. It is used to filter redistributed routes or configure the attributes of redistributed routes. The value of <i>route-map-name</i> cannot exceed 32 characters. By default, route-map is not configured.</p> <p>distribute-list access-list-name: Filters redistributed routes by using distribute-list. <i>access-list-name</i> indicates the associated prefix list, which can be a standard prefix list, an extended prefix list, or a name prefix list. It is in the format of {<1-99> <100-199> <1300-1999> <2000-2699> <i>acl-name</i> }. When the IS-IS address-family ipv6 configuration mode is applied, only the name prefix list can be used, in the format of <i>acl-name</i>.</p>
Command Mode	IS-IS routing process configuration mode and IS-IS address-family ipv6 configuration mode
Usage Guide	<p>You can use the route-map or distribute-list parameter to filter the specified instance's Level-1 routes to be redistributed. Only the routes that meet specific criteria can be redistributed to Level-2 of the current instance. The route-map and distribute-list parameters cannot be used at the same time.</p> <p>The no redistribute isis [tag] level-2 into level-1 command is used to cancel the redistribution of the specified instance's routes. If no redistribute is followed by other parameters, the command will restore the default parameter settings, rather than cancel route redistribution.</p> <p>For example, no redistribute isis tag1 level-1 into level-2 cancels the redistribution of the routes of the IS-IS instance name tag1. no redistribute isis tag1 level-1 into level-2 route-map aa cancels the use of the route map named aa to filter redistributed routes.</p>

Redistributing the Level-2 Reachable Routing Information of the Specified IS-IS Instance to Level-1 of the Current Instance

Command	redistribute isis [tag] level-2 into level-1 [route-map route-map-name / distribute-list access-list-name / (prefix ip-address net-mask / ipv6-prefix ipv6-address/length)]
Parameter Description	<p><i>tag</i>: Indicates the name of the IS-IS instance whose routing information will be redistributed.</p>

	<p>route-map <i>route-map-name</i>: Indicates the route map used for route redistribution. It is used to filter redistributed routes or configure the attributes of redistributed routes. The value of <i>route-map-name</i> cannot exceed 32 characters. By default, route-map is not configured.</p> <p>Distribute-list <i>access-list-name</i>: Filters redistributed routes by using distribute-list. <i>access-list-name</i> indicates the associated prefix list, which can be a standard prefix list, an extended prefix list, or a name prefix list. It is in the format of {<1-99> <100-199> <1300-1999> <2000-2699> <i>acl-name</i> }. When the IS-IS address-family ipv6 configuration mode is applied, only the name prefix list can be used, in the format of <i>acl-name</i>.</p> <p>prefix <i>ip-address net-mask</i>: Determines the routes to be redistributed by address and prefix length.</p> <p>ipv6-prefix <i>ipv6-address/length</i>: Determines the IPv6 routes to be redistributed by address and prefix length.</p>
Command Mode	IS-IS routing process configuration mode and IS-IS address-family ipv6 configuration mode
Usage Guide	<p>You can use the route-map, distribute-list, or prefix parameter to filter the specified instance's Level-2 routes to be redistributed. Only the routes that meet specific criteria can be redistributed to Level-1 of the current instance.</p> <p>The no redistribue isis [tag] level-2 into level-1 command is used to cancel the redistribution of the specified instance's routes. If no redistribute is followed by other parameters, the command will restore the default parameter settings, rather than cancel route redistribution.</p> <p>For example:</p> <p>no redistribue isis tag1 level-2 into level-1 cancels the redistribution of the routes of the IS-IS instance name tag1. no redistribue isis tag1 level-2 into level-1 route-map aa cancels the use of the route map named aa to filter redistributed routes.</p>

Configuration

Example

Configuring the Maximum Number of Equal-Cost Paths

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure the maximum number of equal-cost paths.
----------------------	---

	<pre>A(config)# router isis A(config-router)# maximum-paths 5</pre>
Verification	<ul style="list-style-type: none"> Check whether the maximum number of equal-cost paths displayed by routing entries is the same as the configuration.
	<pre>A# show ip route isis</pre>

Configuring the Maximum Length Allowed for Received LSPs

Configuration	<ul style="list-style-type: none"> Configure IS-IS neighbors. (Omitted) Configure the maximum length allowed for received LSPs.
	<pre>A(config)# router isis A(config-router)# lsp-length receive 512</pre>
Verification	Capture packets to check the length of received LSPs.

Configuring the Maximum Length Allowed for Sent LSPs

Configurations	<ul style="list-style-type: none"> Configure IS-IS neighbors. (Omitted) Configure the maximum length allowed for sent LSPs.
	<pre>A# configure terminal A(config)# router isis 1 A(config-router)# lsp-length originate 512 level-2</pre>
Verification	Capture packets to check the length of sent LSPs.

Configuring a Passive Interface

Configuration Steps	<ul style="list-style-type: none"> Configure IS-IS neighbors. (Omitted) Configure a passive interface.
	<pre>A# configure terminal A(config)# router isis 1 A(config-router)# passive-interface GigabitEthernet 0/0</pre>

Verification	Capture packets to check whether the interface receives and sends IS-IS packets.
---------------------	--

Configuring the Metric of an IS-IS Interface

Configuration Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure metric of the IS-IS interface.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)#isis metric 1</pre>
Verification	Check the database details of IS-IS.
	<pre>A# show isis database detail</pre>

Configuring an Interface as a Passive Interface

Configuration Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure the interface as a passive interface.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)# isis passive</pre>
Verification	Check the interface status of IS-IS.
	<pre>A# show isis interface GigabitEthernet 0/1</pre>

Configuring the Priority of the DIS

Configuration Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure the priority of the DIS.
	<pre>A(config)# interface GigabitEthernet 0/1 A(config-if)# isis priority 127 level-1</pre>
Verification	Check whether the device with the changed priority setting is elected as the DIS.
	<pre>A# show isis database detail</pre>

Generating a Default Route

Configuration Steps	<ul style="list-style-type: none"> Configure IS-IS neighbors. (Omitted) Generate a default route.
	<pre>A(config)# router isis A(config-router)# default-information originate</pre>
Verification	Capture packets to check whether the sent LSP contains a default route.

Configuring SPF Calculation Cycle

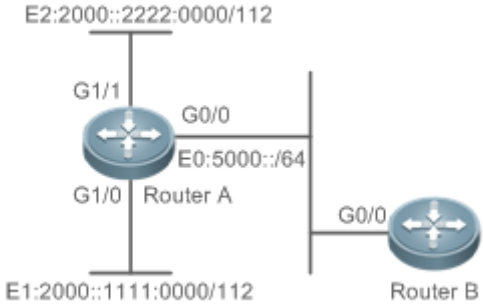
Configuration Steps	<ul style="list-style-type: none"> Configure IS-IS neighbors. (Omitted) Configuring SPF calculation cycle.
	<pre>A(config)# router isis A(config-router)# spf-interval 5 100 200</pre>
Verification	Check whether the SPF calculation cycle works.

Configuring an IS-IS Summary Route

Configuration Requirements	Router A and Router B are connected through the Ethernet and run IS-IS. Configure Router A to advertise only the 172.16.0.0/22 route instead of the 172.16.1.0/24 and 172.16.2.0/24 routes.
Figure 6-19 IS-IS Route Summary Topology	
Configuration Steps	<ul style="list-style-type: none"> Configure IS-IS. Configure Ethernet interfaces. Configure the password for IS-IS authentication.
A	Configure IS-IS.

	<pre>A(config)# router isis A(config-router)# net 49.0001.0000.0000.0001.00 A(config-router)# summary-address 172.16.0.0/16 level-1-2</pre>
	<p>Configure Ethernet interfaces.</p> <pre>A(config)# interface GigabitEthernet 0/0 A(config-if)# ip address 192.168.20.1 255.255.255.0 A(config-if)# ip router isis A(config)# interface GigabitEthernet 1/0 A(config-if)# ip address 172.16.1.1 255.255.255.0 A(config-if)# ip router isis A(config)# interface GigabitEthernet 1/1 A(config-if)# ip address 172.16.2.1 255.255.255.0 A(config-if)# ip router isis</pre>
B	Configure IS-IS.
	<pre>B(config)# router isis B(config-router)# net 49.0001.0000.0000.0002.00</pre>
	Configure an Ethernet interface.
	<pre>B(config)# interface GigabitEthernet 0/0 B(config-if)# ip address 192.168.20.2 255.255.255.0 B(config-if)# ip router isis</pre>
Verification	Run the show ip route command on Router B to check whether only one summary route exists.
B	<pre>B(config)# show ip route i L1 172.16.0.0/16 [115/20] via 192.168.20.1, FastEthernet0/0</pre>

Configuring an IS-ISv6 Summary Route

	<p>Router A and Router B are connected through the Ethernet and run IS-ISv6. Configure Router A to advertise only the 2000::/96 route instead of the 2000::1111:0/112 and 2000::2222::0/112 routes.</p>
<p>Figure 6-20 IS-ISv6 Route Summary Topology</p>	 <p>The diagram shows two routers, Router A and Router B, connected via their GigabitEthernet (G0/0) interfaces. Router A has three other interfaces: GigabitEthernet 1/1, GigabitEthernet 1/0, and Ethernet 0. Router B has only the GigabitEthernet 0/0 interface shown. IPv6 addresses are assigned to these interfaces: E2:2000::2222:0000/112 on G1/1 of Router A; E1:2000::1111:0000/112 on G1/0 of Router A; E0:5000::/64 on E0 of Router A; and G0/0 on G0/0 of Router B.</p>
<p>Configuratio n Steps</p>	<ul style="list-style-type: none"> ▪ Configure IS-IS. ▪ Configure Ethernet interfaces. ▪ Configure the password for IS-IS authentication.
<p>A</p>	<p>Configure IS-IS.</p>
	<pre>A(config)# ipv6 unicast-routing A(config)# router isis A(config-router)# net 49.0001.0000.0000.0001.00 A(config-router)# address-family ipv6 unicast A (config-router-af)# summary-prefix 2000::/96 level-1-2 A (config-router-af)# exit-address-family</pre>
	<p>Configure Ethernet interfaces.</p> <pre>A(config)# interface GigabitEthernet 0/0 A(config-if)# ipv6 address 5000::1/64 A(config-if)# ipv6 router isis A(config)# interface GigabitEthernet 1/0 A(config-if)# ipv6 address 2000::1111:0001/112 A(config-if)# ipv6 router isis A(config)# interface GigabitEthernet 1/1 A(config-if)# ipv6 address 2000::2222:0001/112 A(config-if)# ipv6 router isis</pre>

B	Configure IS-IS.
	<pre>B(config)# ipv6 unicast-routing B(config)# router isis B(config-router)# net 49.0001.0000.0000.0002.00</pre>
	Configure an Ethernet interface.
	<pre>B(config)# interface GigabitEthernet 0/0 B(config-if)# ipv6 address 5000::2/64 B(config-if)# ipv6 router isis</pre>
Verification	Run the show ipv6 route command on Router B to check whether only one summary route exists.
B	<pre>B(config)# show ipv6 route l1 2000::/96 [115/20] via FE80::C800:1BFF:FEF8:1C, FastEthernet1/0</pre>

Enabling Neighbor Relationship Event Output

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Enable neighbor relationship event output.
	<pre>A(config-router)# log-adjacency-changes</pre>
Verification	Change the neighbor state and verify that the change is recorded when debugging is disabled.

Configuring Route Redistribution

Configurati on Steps	<ul style="list-style-type: none"> ▪ Configure IS-IS neighbors. (Omitted) ▪ Configure OSPF routes. (Omitted) ▪ Configure route redistribution
	<pre>A(config)# router isis A(config-router)# redistribute ospf 1 metric 10 level-1</pre>
Verification	Check whether routing entries with redistributed routes exist.

```
A# show ip route isis
```

6.5 Monitoring

Clearing

! Running the **clear** commands may lose vital information and thus interrupt services.

Description	Command
Clears all IS-IS neighbor relationship tables.	clear clns neighbors
Clears all IS-IS data structures.	clear isis *
Clears all IS-IS counters.	clear isis [tag] counter

Displaying

Description	Command
Displays all IS neighbors and inter-device neighbor relationships.	show clns [tag] is-neighbors [interface-type interface-number] [detail]
Displays all IS neighbors and provides device information and information about the neighbor relationship with ESs.	show clns [tag] neighbors [interface-type interface-number] [detail]
Displays all IS-IS counters.	show isis [tag] counter
Displays the LSDB information.	show isis [tag] database [FLAGS] [LEVEL] [LSPID]
Displays the state information related to IS-IS GR.	show isis [tag] graceful-restart

Displays the relationship between the device name and system ID.	show isis [tag] hostname
Displays the details of an IS-IS interface.	show isis [tag] interface [interface-type interface-number] [counter]
Displays the mesh group configuration of all interfaces.	show isis [tag] mesh-groups
Displays IS-IS neighbor information.	show isis [tag] neighbors [detail]
Displays IS-IS NSR information.	show isis [tag] nsr
Displays the neighbor information of virtual systems in IS-IS.	show isis [tag] virtual-neighbors
Displays IS-IS information.	show isis [tag] protocol
Displays the topology of IS-IS device connection.	show isis [tag] topology [frr { self-originate WORD all }] [I1 I2 level-1 level-2]
Displays information of an IS-IS IPv6 unicast topology.	show isis [tag] ipv6 topology [I1 I2 level-1 level-2]

Debugging

⚠ System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Enables IS-IS debugging.	debug isis { all auth events frr gr ifsm lsp mtr n fsm nsm nsr pdu spf warn }

7 CONFIGURING BGP

7.1 Overview

The Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP) used for communication between routers in different autonomous systems (ASs). BGP is used to exchange network accessibility information between different ASs and eliminate routing loops by using its own mechanism.

BGP uses TCP as the transmission protocol. The reliable transmission mechanism of TCP is used to ensure the transmission reliability of BGP.

Routers running BGP are called BGP speakers. BGP speakers between which a BGP session is established are called BPG peers.

Two modes can be used to establish peers between BGP speakers: Internal BGP (IBGP) and External BGP (EBGP).

- IBGP refers to a BGP connection established within an AS and completes transition of routing information within the AS.
- EBGP refers to a BGP connection established between different ASs and completes exchange of routing information between different ASs.

Rules for BGP to select an optimum route:

1. Invalid routing table entries are not involved in optimum route selection.

i Invalid entries include entries of inaccessible next hops and flapping entries.

2. Otherwise, select a route with a large value of **LOCAL_PREF**.

3. Otherwise, select a route generated by a BGP speaker.

i Routes generated by a BGP speaker include routes generated by the **network**, **redistribute** and **aggregate** commands.

4. Otherwise, select a route with the shortest AS length.

5. Otherwise, select a route with a smaller value of **ORIGIN**.

6. Otherwise, select a route with the smallest value of **MED**.

7. Otherwise, EBGP routes have higher priorities than IBGP routes and routes in the AS alliance, and the IBGP routes have the same priorities as the routes in the AS alliance.

8. Otherwise, select a route with the smallest IGP metric value to the next hop.

9. Otherwise, select an EBGP route that is received first.

10. Otherwise, select a route advertised by a BGP speaker with a smaller router ID.

11. Otherwise, select a route with a large cluster length.
12. Otherwise, select a route with a large neighbor address.

i The preceding shows the route selection process under the default configurations. By using CLI commands, you can change the route selection process. For example, you can run the **bgp bestpath as-path ignore** command to make step 4 of the route selection process lose effect or run the **bgp bestpath compare-routerid** command to make step 9 lose effect.

Protocols and Standards

- RFC4271: A Border Gateway Protocol 4 (BGP-4)
- RFC4273: Definitions of Managed Objects for BGP-4
- RFC4360: Proposed Standard: BGP Extended Communities Attribute
- RFC4486: Proposed Standard: Subcodes for BGP Cease Notification Message
- RFC4724: Proposed Standard: Graceful Restart Mechanism for BGP
- RFC4760: Draft Standard: Multiprotocol Extensions for BGP-4
- RFC5492: Draft Standard: Capabilities Advertisement with BGP-4

7.2 Applications

Application	Description
Inter-AS Route Advertisement	Implement inter-AS route advertisement by using BGP.
Intra-AS Route Reflection	Set up a route reflection topology within an AS to reduce BGP connections.

7.2.1 Inter-AS Route Advertisement

Scenario

BGP implements route advertisement and maintenance across different ASs.

As shown in Figure 7-1, BGP transfers the route of AS 65536 to AS 65538 through AS 65537.

Figure 7-1



Remarks	<p>R1 is a device at the network edge of AS 65536.</p> <p>R2 and R3 are devices at the network edge of AS 65537.</p> <p>R4 is a device at the network edge of AS 65538.</p>
----------------	--

Deployment

- Establish the EBGP neighborhood between R1 and R2 to implement inter-AS route advertisement.
- Establish the IBGP neighborhood between R2 and R3 to implement intra-AS route advertisement.
- The Internet runs OSPF to ensure network accessibility between R2 and R3.
- Establish the EBGP neighborhood between R3 and R4 to implement inter-AS route advertisement.

7.2.2 Intra-AS Route Reflection

Scenario

According to the BGP route advertisement principles, routes learned by an IBGP neighbor will not be advertised to the next IBGP neighbor by default. Therefore within an AS, a device running BGP must implement full-mesh. When there are many BGP devices within the AS, implementing full-mesh may cause large difficulties for network deployment. In this case, route reflection can be used to solve this problem.

As shown in Figure 7-2, route reflection is deployed to implement BGP full-mesh among R1 to R4 and RR.

Figure 7-2



Remarks	RR is a route reflector. R1 to R4 are route reflection clients.
---------	--

Deployment

- Establish IBGP neighborships between R1 to R4 and RR respectively.
- Configure R1 to R4 as the route reflection clients of RR.

7.3 Features

Basic Concept

BGP Speaker and AS Number

A router enabled with BGP is called a BGP speaker.

After a route is enabled with BGP, a local AS number must be specified for the router. An AS number is a globally unique number allocated by IANA, ranging from 1 to 4294967295.

BGP Neighbor and Peer

Before a route is advertised between BGP speakers, a neighborship must be established in advance. You need to manually configure BGP neighbors on both BGP speakers. That is, configure the peer as a neighbor on the two BGP speakers respectively. Therefore, BGP neighbors are also called BGP peers.

Neighbor Type and Route Type

BGP neighborships are classified into the following types:

- IBGP neighborship: The neighborship between BGP speakers within an AS is called IBGP neighborship. Routes learned from IBGP neighbors are called IBGP routes.
- EBGP neighborship: The neighborship between BGP speakers in different ASs is called EBGP neighborship. Routes learned from EBGP neighbors are called EBGP routes.

BGP route attribute

When a BGP speaker advertises routes to its neighbors, the BGP speaker also advertises the attributes carried by the routes. Common BGP attributes are as follows:

- ORIGIN: Specifies the origin of a BGP route and can be set to **IGP**, **EGP**, or **INCOMPLETE**.
- AS-PATH: Lists the ASs passed by a route in a reverse order. The last AS is placed at the beginning of the list.
- NEXT-HOP: Specifies the IP address of the next hop to be reached by a BGP route.
- MULTI-EXIT-DISC: Distinguishes multiple output/input interfaces for reaching the same neighbor AS. A smaller value means a higher priority.
- LOCAL-PREF: Distinguishes the priorities of IBGP routes in an AS. A larger value means a higher priority.

Overview

Feature	Description
Creating a BGP Neighbor	Create a BGP neighbor.
Configuring a BGP Route Reflector	Set up a BGP route reflection topology to simplify network deployment for BGP neighbor full-mesh.
Configuring a BGP Alliance	Configure a BGP alliance to simplify network deployment for BGP neighbor full-mesh.
Re-distributing Local AS Network Information to BGP	Re-distribute routing information to BGP and advertise local routes through BGP.
Controlling Route Exchange	Configure the route exchange policy for a BGP peer and control routes to be received by and to be advertised to this peer.

<u>Between BGP Peers</u>	
<u>Obtaining Accessible Networks of Other ASs from BGP</u>	Re-distribute routing information in BGP into a core routing table or IGP.
<u>Configuring Synchronization Between BGP and IGP</u>	Configure BGP to check whether BGP routes are synchronized with IGP routes.
<u>Configuring BGP Soft Reset</u>	After a routing policy changes, use soft reset to apply a new policy.
<u>Configuring the Route Attributes of BGP</u>	Configure the route selection algorithms and routing policy control of BGP.
<u>Configuring BGP Route Aggregation</u>	Reduce routes by means of route aggregation.
<u>Configuring BGP Route Dampening</u>	Reduce the impacts of route flapping on a network topology.
<u>Configuring the Management Distance of BGP</u>	Change the priorities of BGP routes.
<u>Configuring Multi-path Load Balancing of BGP</u>	Configure multi-path load balancing for BGP to enhance the network reliability and increase the network bandwidth.
<u>Configuring BGP FRR</u>	Configure fast rerouting for BGP to enhance the network reliability.

<u>Configuring BGP Timers</u>	Modify the internal timer time of BGP.
<u>Configuring BGP Route Update Mechanisms</u>	Disable/Enable regular scanning for BGP routes and configure the route scanning interval.
<u>Configuring the Next-Hop Triggering Update Function of BGP</u>	Configure the next hop triggering update function of BGP.
<u>Configuring BGP LOCAL AS</u>	Configure the LOCAL AS for a BGP neighbor.
<u>Configuring BGP Capacity Protection</u>	Avoid non-predictable running status caused by consumption of device capacity.
<u>Configuring BGP GR</u>	Configure the BGP GR function to enhance the network reliability.
<u>Configuring 4-Byte AS Numbers of BGP</u>	Configure the display mode of a 4-byte AS number.
<u>Configuring a Regular Expression</u>	Use a regular expression to filter routing information.
<u>Configuring BGP Session Retention</u>	Configure BGP to ensure that after an address family with incorrect routing attributes is detected for a neighbor, other address family routes advertised by the neighbor will not be affected.
<u>Configuring BGP Delayed Advertisement upon System Restart</u>	Configure BGP to delay route advertisement to a neighbor within a period after the system is restarted.

Configuring BGP NSR	Enable the NSR function for BGP neighbor.
Configuring BGP Routes to Be Recursive Only to Host Routes	Enable IBGP or EBGP routes to be recursive only to host routes.
Configuring Outbound Loop Detection for a BGP Neighbor	Enable outbound loop detection for a BGP neighbor.
Shutting Down BGP Connection Gracefully	Shut down BGP connection gracefully. Make sure that the service traffic is not interrupted or is interrupted for a very short period of time.
Configuring Enhanced VPN Route Import	Configure the enhanced VPN route import function.
EVPN Route Attribute	Indicate EVPN route attribute.
Configuring BGP EVPN	Configure the BGP EVPN function.
Other Related Configurations	Configure extended BGP functions.

7.3.1 Creating a BGP Neighbor

A BGP neighbor is manually configured by a user. Two connection modes are supported: IBGP and EBGP. You can identify the connection mode between BGP speakers based on the AS where the BGP peer resides and the AS where the BGP speaker resides.

- i** Generally, BGP speakers between which an EBGP connection is established are directly connected whereas BGP speakers between which an IBGP connection is established can be at any location within an AS.

Working
Principle

A BGP speaker can initiate a TCP connection request to a BGP peer specified by a user. After the TCP connection is successfully created, the peers will exchange BGP packets to negotiate about connection parameters. The BGP neighborhood is successfully established after the negotiation succeeds.

Creating a TCP Connection

A BGP speaker initiates a TCP connection request to a neighbor. The destination IP address is the peer IP address specified by the user and the port number is fixed to 179.

The BGP speaker also listens on the port number 179 of the local TCP connection to receive connection requests from its peer.

Negotiating about Protocol Parameters

After the TCP connection is successfully created, the BGP speakers exchange OPEN packets to negotiate about BGP connection parameters. The parameters for negotiation include:

- Version: Indicates the BGP version number. At present, only version 4 is supported.
- Neighbor AS number: Determines whether the AS number of the neighbor is consistent with the local AS number. If not, the connection request will be denied.
- Hold Time: Negotiates about the timeout duration for the BGP connection. The default value is 180 seconds.
- Neighbor capability: Negotiates about various extended capabilities supported by the neighbor, including the address family, dynamic route update, and GR functions.

Maintaining Neighborhood

The Keepalive message is periodically sent between BGP speakers. If a new Keepalive packet is not received from the BGP neighbor after the Hold Time expires, the BGP speaker considers that the neighbor is not accessible, disconnects the TCP connection from the neighbor, and attempts to reconnect to it. The interval for a BGP speaker to send the Keepalive message is one third of the Hold Time determined through negotiation and is 60 seconds by default.

Related Configuration

Creating a BGP Neighbor

By default, a BGP speaker does not specify any neighbor. You can manually configure a BGP neighbor. You can run the **neighbor** { *peer-address* | *peer-group-name* } **remote-as** *as-number* command to manually create a BGP neighbor and specify the AS number of the neighbor.

Setting the Neighbor TTL

By default, The TTL field in a TCP packet sent by an IBGP neighbor is set to the maximum value (255). It is set to 1 by an EBGP neighbor.

You can run the **neighbor** { *peer-address* | *peer-group-name* } **ebgp-multihop** [*ttl*] command to set the TTL field of a TCP packet sent by a BGP neighbor.

A larger value of TTL means a longer distance between BGP neighbors. When TTL is 1, the BGP neighbor devices must be directly connected.

Setting the Source Address of TCP

By default, BGP automatically selects the source IP address of a TCP connection based on the IP address of the neighbor. Generally, the IP address of a local packet output interface is used.

You can run the **neighbor** { *peer-address* | *peer-group-name* } **update-source** { *interface-type interface-number* | *address* } command to adjust the source IP address of the neighbor's TCP connection.

Setting MD5 Encryption

By default, a BGP connection is not encrypted through MD5.

You can run the **neighbor** { *peer-address* | *peer-group-name* } **password** [*0* | *7*] *string* command to set encryption for a BGP neighbor's TCP connection.

Activating the Address Family Capability of a Neighbor

By default, a neighbor created in the BGP configuration mode activates only the IPv4 Unicast address family capability.

You can run the **address-family** command to enter a corresponding address family mode, and then run the **neighbor** { *peer-address* | *peer-group-name* } **activate** command to activate the address family capability for the BGP neighbor.

7.3.2 Configuring a BGP Route Reflector

According to the principle of BGP route advertisement, full mesh must be established for all BGP speakers within an AS (neighborships need to be established between each two BGP speakers). Too many BGP speakers within an AS will increase the resource overhead of the BGP speakers, increase the network administrator's workload and complexity of configuration and decrease the network expansion capability.

Using a route reflector is a method for reducing IBGP peer connections within an AS.

-
- i** The methods for reducing the IBGP peer connections within an AS include using a route reflector and using an AS alliance.
-

Working Principle

Configure a BGP speaker as a route reflector which classifies IBGP peers in an AS into two types: clients and non-clients.

The rules for implementing a route reflector within an AS are as follows:

- Configure a route reflector and specify clients for the route reflector. The route reflector and its clients form a cluster. The route reflector will connect to its clients.
- The clients of a route reflector in a cluster cannot connect to other BGP speakers out of the cluster.
- Within an AS, full mesh is established among IBGP peers of non-clients. The IBGP peers of non-clients include the following situations: Multiple route reflectors in a cluster; a route reflector in a cluster and BGP speakers (generally not supporting the route reflector function) not involved in the route reflector function out of the cluster; a route reflector in a cluster and route reflectors in other clusters.

The rules for processing a route received by a route reflector are as follows:

- A route update message received by an EBGP speaker will be sent to all clients and non-clients.
- A route update message received by a client will be sent to other clients and all non-clients.
- A route update message received by an IBGP speaker will be sent to all the other clients.

- i** Generally, only one route reflector is configured in a cluster. In this case, the Router ID of the route reflector can be used to identify this cluster. To increase the redundancy, you can set multiple route reflectors in a cluster. In this case, you must configure the cluster ID so that a route reflector can identify the route update messages from other route reflectors in the cluster.
- i** If multiple route reflectors are configured for a cluster, you must configure a cluster ID for the cluster.
- i** Generally, it is unnecessary to create connections between the clients of a route reflector in a cluster because the route reflector will reflect the routes between the clients. However, if full mesh has been established among all clients, you can cancel the client route reflection function of the route reflector.

Related Configuration

Configuring a BGP Route Reflector and Reflected Clients

By default, BGP is not configured with route reflection.

You can run the **neighbor peer-address route-reflector-client** command to configure a device as a route reflector and its neighbor devices as reflected clients.

Configuring BGP Client-Client Reflection

By default, BGP client-client route reflection is enabled, which means that routes received from a reflected client can be advertised to other clients.

You can run the **bgp client-to-client reflection** command to enable or disable (using the **no** form of this command) client-client reflection.

Configuring a BGP Reflection Cluster ID

By default, a BGP reflection cluster ID is the Router-ID of BGP. If multiple reflection clusters are deployed within an AS, different reflection cluster IDs must be configured for these reflection clusters.

You can run the **bgp cluster-id** *cluster-id* command to manually configure the cluster ID of a route reflector.

7.3.3 Configuring a BGP Alliance

An alliance is another method for reducing the IBGP peer connections within an AS.

Working Principle

Divide an AS into multiple sub ASs and configure a unified alliance ID (namely, the alliance AS NUMBER) for these sub ASs to form an alliance. Outside the alliance, the entire alliance is still considered as an AS and only the AS number of the alliance is visible. Inside the alliance, full mesh of IBGP peers can be established for BGP speakers within a sub AS, and EBGP connections can be established for BGP speakers in different sub ASs. Though EBGP connections are established between BGP speakers within a sub AS, when information is exchanged, NEXT_HOP, MED, LOCAL_PREF and other path attributes keep unchanged.

Related Configuration

Configuring a BGP Alliance ID

By default, no alliance ID is configured for a BGP speaker.

You can run the **bgp confederation identifier** *as-number* command to configure a BGP alliance ID. After the configuration is successful, the local AS (specified by the **router bgp** *as-number* command) of BGP becomes the private AS inside the alliance and is invisible to other ASs.

Configuring a BGP Alliance Neighbor

By default, no alliance neighbor is configured for BGP.

You can run the **bgp confederation peers** *as-number* [... *as-number*] command to configure a BGP alliance neighbor. After the configuration succeeds, the AS specified by this command and the local AS belong to the same alliance.

7.3.4 Re-distributing Local AS Network Information to BGP

BGP cannot automatically discover or learn accessible networks. The accessible network information of a local AS must be re-distributed to BGP. Then, BGP can advertise the information to neighbors.

Working Principle

Two methods can be used to re-distribute local AS network information to BGP:

- Manual static configuration: re-distribute the accessible network information within a specified range to BGP.
- Configuring route re-distribution: re-distribute accessible IGP network information to BGP.

i In addition, you can also re-distribute local AS network information to BGP routes by configuring route aggregation.

[Related Configuration](#)

Configuring a BGP Network

By default, no network is configured for BGP.

You can run the **network** *network-number* [**mask** *mask*] [**route-map** *map-tag*] [**backdoor**] command to configure a BGP network to re-distribute specified accessible network information to BGP. The prerequisite for successfully re-distributing routing information to BGP is that a route is available in the core routing table and this route can be an IGP, directly-connected or static route.

Configuring BGP Route Re-distribution

By default, BGP is not configured with route re-distribution.

Run the **redistribute** *protocol-type* command to import routes of other protocols to BGP, including OSPF, RIP, ISIS, static, and direct routes, host routes converted by ARP, and routes converted by ND.

Importing Routes with Multiple Paths or Next Hops to BGP

By default, routes imported to BGP have only one next hop.

Run the **bgp sourced-paths** *protocol-type* **all** command to import routes with multiple next hops of other protocols to BGP.

7.3.5 Controlling Route Exchange Between BGP Peers

BGP provides powerful route management functions. You can actively control the route exchange between BGP peers.

[Working Principle](#)

Configure the route exchange policy for a BGP peer and control routes to be received by and to be advertised to this peer.

[Related Configuration](#)

Configuring the Default Route to Be Advertised to a Peer

By default, BGP does not advertise the default route.

You can run the **neighbor** { *address* | *peer-group-name* } **default-originate** [**route-map** *map-tag*] command to advertise the default route to a peer (or a peer group).

Configuring Next-Hop-Self for a Peer

By default, BGP does not change the next hop of a route when it advertises the route to an IBGP neighbor and sets the next hop to the local BGP speaker when it advertises the route to an EBGP neighbor.

You can run the **neighbor** { *address* | *peer-group-name* } **next-hop-self** command to configure the next hop of a route to the local BGP speaker when distributing the route to a specified BGP peer (group).

Configuring Remove-Private-AS for a Peer

By default, BGP does not delete the private AS in the AS-PATH attribute when it advertises routing information to a peer.

You can run the **neighbor** { *address* | *peer-group-name* } **remove-private-as** command to require that the private AS number recorded in the AS path attribute should be deleted when routing information is distributed to an EBGP peer (group). This command does not apply to an IBGP neighbor.

Configuring Send-Community for a Peer

By default, BGP does not send the community attribute when it advertises routing information to a peer.

You can run the **neighbor** { *address* | *peer-group-name* } **send-community** command to specify that the community attribute can be sent to a specified BGP peer (group).

Configuring Maximum-Prefix for a Peer

By default, BGP does not restrict the records of routing information that can be received by a peer.

You can run the **neighbor** { *address* | *peer-group-name* } **maximum-prefix** *maximum* [**warning-only**] command to specify the records of routing information received from a specified peer (group).

Configuring Route Filtering for a BGP Neighbor

By default, a BGP neighbor is not enabled with any filtering policy and receives all legal routing information advertised by a neighbor.

BGP supports multiple methods of configuring the route filtering policies for a neighbor, including:

- **neighbor** { *peer-address* | *peer-group-name* } **distribute-list** { *access-list-number* | *access-list-name* } { **in** | **out** }

Use an ACL to filter routes in the input and output directions of the neighbor.

- **neighbor** { *peer-address* | *peer-group-name* } **filter-list** *access-list-number* { **in** | **out** }

Use an AS-PATH list to filter routes in the input and output directions of the neighbor.

- **neighbor** { *peer-address* | *peer-group-name* } **prefix-list** *prefix-list-name* { **in** | **out** }

Use a prefix-list to filter routes in the input and output directions of the neighbor.

- **neighbor** { *peer-address* | *peer-group-name* } **route-map** *map-tag* { **in** | **out** }

Use a route map to filter routes in the input and output directions of the neighbor.

- **neighbor** { *address* | *peer-group-name* } **unsuppress-map** *map-tag*
- Allow for advertising certain routing information previously suppressed by the **aggregate-address** command when distributing routing information to a specified peer.
- **bgp advertise-map** *map-tag*

Use command **bgp advertise-map** *map-tag* to advertise route policy to all neighbors

7.3.6 Obtaining Accessible Networks of Other ASs from BGP

Send routing information of other ASs exchanged by BGP to the routing table of a device so that the device can forward packets to other ASs.

Send routing information of other ASs exchanged by BGP to the routing table of a device so that the device can forward packets to other ASs.

Working Principle

BGP Sends Routing Information to a Core Routing Table

BGP controls routing information sent to the core routing table by using **table-map**. **table-map** can modify the attributes of routing information sent to the core routing table. If the route is matched, BGP modifies the attribute of the routing information and sends the route. If the route is not matched or route matching is denied, BGP does not modify the attribute of the routing information but sends the route.

Changes of **table-map** are not reflected in the core routing table immediately, but reflected a moment later. To update the application of **table-map** immediately, you can run the **clear ip bgp** [*vrf vrf-name*] **table-map** command to update the routing information in the core routing table immediately. This command does not clear the existing routes in the core routing table, but directly applies **table-map** to send the updated routing information, thereby not causing forwarding flapping.

Re-distributing BGP Routes to IGP

Re-distribute BGP routes on a BGP speaker to IGP to ensure that routers within an AS can obtain routes to other ASs.

Related Configuration

Configuring table-map

By default, BGP is not configured with a table-map and allows for sending all routes without modifying the attributes of the routes.

You can run the **table-map** *route-map-name* command to set a table-map and control the routing information to be sent to the core routing table. *route-map-name* specifies a route-map to be associated.

- i** Run the **table-map** command in the BGP configuration mode or in the IPv4 address family mode. The Match and Set statements supported in the table-map are as follows:
Match statements: as-path, community, ip address, ip next-hop, metric, origin and route-type
Set statements: metric, tag and next-hop
- i** You can run the **no table-map** command to delete the table-map configurations.

Configuring BGP Route Re-distribution by IGP

By default, IGP does not re-distribute BGP routes.

You can run the **redistribute bgp** [**route-map** *map-tag*] [**metric** *metric-value*] command to re-distribute BGP routes to IGP (RIP\OSPF\ISIS).

The **bgp redistribute-internal** command controls only whether to re-distribute routes learned from IBGP to IGP. By default, routes learned from IBGP can be re-distributed to IGP.

- i** You can run the **bgp redistribute-internal** command in the BGP configuration mode, IPv4/IPv6 address family mode or the IPv4 VRF address family mode.
- i** You can run the **no bgp redistribute-internal** command to delete the configuration.

7.3.7 Configuring Synchronization Between BGP and IGP

Generally, BGP speakers working as mutual IBGP neighbors are not directly connected. IGP devices between the BGP speakers may fail to learn routing information same as that learned by the BGP speakers. When a BGP speaker at the border of an AS forwards packets received from other domains to the next-hop IBGP neighbor, the packets pass an IGP device in the middle. In this case, the packets may be lost due to no routing information on the IGP device.

Working Principle

To keep synchronization between BGP and IGP, you must ensure that all routers within an AS can learn routing information to be sent to another AS before the routing information is advertised to this AS.

Synchronization between BGP and IGP is not required only in the following cases:

- Routing information passing through an AS is not available. For example, the AS is an end AS.
- All routers within an AS run BGP. Full mesh is established among all BGP speakers (neighborship is established between each two BGP speakers).

Related Configuration

Configuring BGP Route Synchronization

By default, synchronization between BGP and IBGP routes is disabled.

You can run the **synchronization** command to enable synchronization between BGP and IGP.

? You can run the **no synchronization** command to disable synchronization between BGP and IGP.

7.3.8 Configuring BGP Soft Reset

If routing policies (including **neighbor distribute-list**, **neighbor route-map**, **neighbor prefix-list** and **neighbor filter-list**) change, an effective method must be provided to implement new routing policies. A traditional method is to terminate a BGP connection and then create a new BGP connection. By configuring BGP Soft Reset, you can execute a new routing policy without terminating a BGP session connection.

Working Principle

i Routing policies that affect inbound routing information are called inbound routing policies (such as In-route-map and In-dist-list) and routing policies that affect outbound routing information are called outbound routing policies (such as Out-route-map and Out-dist-list).

When outbound routing policies change, BGP soft reset will re-advertise all routing information of a BGP speaker to its neighbors.

If inbound routing policies change, the operation is more complex than that when outbound routing policies change. This is because outbound routing policies are executed in the routing table of the local BGP speaker whereas inbound routing policies are executed for routing information received from the BGP peer. To reduce cost, the local BGP speaker does not store the original routing information received from the BGP peer.

If inbound routing policies change and a neighbor device supports route update, you can configure soft reset to send a route update request to the neighbor device. After receiving the request, the neighbor device re-advertises all routing information. You can also perform configuration to ensure that each BGP peer stores original routing information on the local BGP speaker and provides original routing information basis for modifying inbound routing policies subsequently.

i The "route update capability" allows for modifying and executing routing policies without storing original routing information. This product supports the route update capability. You can run the **show ip bgp neighbors** command to check whether a BGP peer supports route update. If yes, you do not need to run the **neighbor soft-reconfiguration inbound** command when inbound routing policies change.

Related Configuration

Configuring BGP Soft Reset

Run the **clear ip bgp { * | peer-address | peer-group peer-group-name | external } soft out** command to soft reset a BGP connection. You can activate execution of a routing policy without restarting the BGP session.

Saving Original Routing Information of Neighbors

By default, BGP does not save original routing information of neighbors.

Run the **neighbor { address | peer-group-name } soft-reconfiguration inbound** command to save unmodified routing information sent by a BGP peer (group).

7.3.9 Configuring the Route Attributes of BGP

BGP provides various control policies for route attributes. You can apply the policies based on actual conditions.

Working Principle

AS_PATH Attribute

BGP can control distribution of routing information in three modes:

- IP address. You can run the **neighbor distribute-list** and **neighbor prefix-list** commands for implementation.
- AS_PATH attribute. See the description in this section.
- COMMUNITY attribute. See the related configuration of the COMMUNITY attribute.

You can use an AS path-based access control list (ACL) to control the distribution of routing information. Where, the AS path-based ACL uses a regular expression to parse the AS path.

Based on the standard (RFC1771), BGP does not consider the AS path length when selecting the optimum path. Generally, a shorter AS path length means a higher path priority; therefore, QTECH considers the AS path length when selecting the optimum path. You can determine whether to consider the AS path length when selecting the optimum path based on the actual conditions.

-
- i** Within an AS, whether to consider the AS path should be consistent for all BGP speakers when the optimum path is selected; otherwise, the optimum paths selected by the BGP speakers may be different.
-

MULTI_EXIT_DISC Attribute

BGP uses the MED value as the basis for comparing priorities of paths learned from EBGP peers. A smaller MED value means a higher path priority.

- By default, the MED value is compared only for paths of peers from the same AS when the optimum path is selected.
- By default, the MED value is not compared for paths of peers from other sub ASs within an AS alliance.
- By default, if a path not configured with the MED attribute is received, it is considered that the MED value of this path is 0. Since a smaller MED value means a higher path priority, this path has the highest priority.
- By default, the MED value is not compared with paths from different ASs; instead, the sequence of receiving the paths is compared.

LOCAL_PREF Attribute

When sending routes received from EBGP peers to IBGP peers, a BGP speaker adds the LOCAL_PREF attribute. BGP uses the LOCAL_PREF attribute as the basis for comparing priorities of paths learned from IBGP peers. A larger value of LOCAL_PREF means a higher path priority.

You can also run the **set local-preference** command of a route map to modify the LOCAL_PREF attribute of the specified path.

COMMUNITY Attribute

The COMMUNITY attribute is another mode for controlling distribution of routing information.

A community is a set of destination addresses. The COMMUNITY attribute is intended to facilitate execution of a routing policy based on a community and thereby simplify the configuration of routing information distribution control on BGP speakers. Each destination address may belong to multiple communities. An AS administrator can define the communities to which a destination address belongs.

By default, all destination addresses belong to the Internet community and are carried in the community attribute of the path.

At present, four common community attribute values are pre-defined:

- Internet: Indicates the Internet community. All paths belong to this community.
- no-export: Indicates that the path is not advertised to EBGP peers.
- no-advertise: Indicates that the path is not advertised to any BGP peer.

- **local-as**: Indicates that a path is not advertised to other ASs. When an AS alliance is configured, the path is not advertised to other ASs or sub ASs.

By using the community attribute, you can control the receiving, prioritization and distribution of routing information. BGP speakers can set, add or modify the community attribute when learning, advertising or re-distributing routes. An aggregation path will contain the community attribute values of all aggregated paths.

-
- BGP supports up to 32 COMMUNITY attributes for each route and allows for up to 32 COMMUNITY attributes when match and set COMMUNITY of a route map are configured.
-

Others

During selection of the optimum path, if two paths with the same path attributes are received from different EBGP peers, the optimum path is selected based on the receiving sequence by default. You can disable comparison of the receiving sequence but use the path with a smaller router ID as the optimum path.

Related Configuration

Configuring AS_PATH Attribute

- **ip as-path access-list** *path-list-name* { **permit** | **deny** } *as-regular-expression*
- Defines an AS path list.
- **neighbor** { *address* | *peer-group-name* } **filter-list** *path-list-name* { **in** | **out** }
- By default, no filtering policy is configured for BGP peers.
- The configuration is the same as that for routing information receiving and sending for a specified BGP peer (group). Routing policies are executed based on the AS path list to advertise or receive only routes that match the policies.
- **neighbor** { *address* | *peer-group-name* } **route-map** *map-tag* { **in** | **out** }
- By default, no filtering policy is configured for BGP peers.
- The configuration is the same as when receiving and sending routing information for a specified BGP peer (group). Routing policies are executed based on a route map or the set rules in the route map are used to modify routing attributes.

-
- In the route-map configuration mode, you can run the **match as-path** command to modify AS path attributes by using an AS path list or directly run the **set as-path** command to modify AS attribute values.
-

- **bgp bestpath as-path ignore**

Allows BGP not to consider the AS path length when selecting the optimum path. The AS path length is compared by default.

By default, a smaller AS path length means a higher path priority.

Configuring MULTI_EXIT_DISC Attribute

- `bgp always-compare-med`
- Allows for comparing the MED values of paths from different ASs, which is disabled by default.
- `bgp bestpath med confed`
- Allows for comparing the MED values of paths of peers from other sub ASs in the same AS alliance, which is disabled by default.
- `bgp bestpath med missing-as-worst`
- Sets a path not configured with the MED attribute to the lowest priority, which is disabled by default.
- `bgp deterministic-med`
- Allows for comparing the paths of peers within the same AS, which is disabled by default.

Configuring LOCAL_PREF Attribute

- `bgp default local-preference value`
- Changes the default local preference value, ranging from 0 to 4,294,967,295. A larger value means a higher priority. The default value is 100.

Configuring COMMUNITY Attribute

- `ip community-list standard community-list-name { permit | deny } community-number`

Creates a community list. **community-list-name** indicates the name of the community list.

i `community-number`: Indicates a value (0 to 4,294,967,295) specified by a user or a known community attribute (internet, local-AS, no-advertise or no-export).

- `neighbor { address | peer-group-name } send-community`
- Allows for sending the community attribute to a specified BGP peer (group), which is not configured by default.
- `neighbor { address | peer-group-name } route-map map-tag { in | out }`
- The configuration is the same as that for routing information receiving and sending for a specified BGP peer (group). Routing policies are executed based on a route map. No filtering policy is configured for peers by default.

i In the route-map configuration mode, you can run the **match community-list [exact]** and **set community-list delete** commands to modify the community attribute by using a community list or directly run the **set community** command to modify the community value.

Others

- `bgp bestpath compare-routerid`
- Allows BGP to compare the router ID when selecting the optimum path, which is disabled by default.

7.3.10 Configuring BGP Route Aggregation

BGP-4 supports CIDR and therefore allows for creating aggregation entries to reduce the size of a BGP routing table. BGP aggregation entries can be added to a BGP routing table only when valid paths are available within the aggregation range.

Working Principle

Aggregate one or more detailed BGP routes into a BGP route with a shorter network mask.

- i** By default, BGP advertises all path information before and after aggregation. If you hope that only aggregated path information is advertised, you can run the **aggregate-address summary-only** command.
- i** When the **aggregate-address** command is used to configure an aggregated route, the aggregated route takes effect immediately as long as there are routes in the configured address range.

Related Configuration

Configuring BGP Route Aggregation

- `aggregate-address address mask`
- Configures BGP route aggregation. By default, BGP does not create any aggregated routing entry.
- `aggregate-address address mask as-set`
- Configures an aggregation address and stores the AS path information within the aggregation address range. By default, BGP does not store AS path information.
- `aggregate-address address mask summary-only`
- Configures an aggregation address and advertises only an aggregated path. By default, BGP advertises all path information within the aggregation range.
- `aggregate-address address mask as-set summary-only`
- Configures an aggregation address, stores the AS path information within the aggregation address range and advertises only aggregated paths.

7.3.11 Configuring BGP Route Dampening

If a route changes between being valid and invalid, route flapping occurs.

Route flapping often causes transmission of unstable routes in a network, and thereby causes network instability. BGP route dampening is a method for reducing route flapping. It reduces possible route flapping by monitoring routing information from EBGp peers.

Working Principle

Terms used in BGP route dampening are as follows:

- **Route Flap:** A route changes between being valid and invalid.
- **Penalty:** Once route flapping occurs, a BGP speaker enabled with route dampening adds a value to the penalty for this route. The penalty is accumulated until the Suppress Limit is reached.
- **Suppress Limit:** When the penalty of a route is greater than this value, the route will be suppressed.
- **Half-life-time:** The time used for the penalty to be halved.
- **Reuse Limit:** When the penalty value of a route is smaller than this value, route suppression will be canceled.
- **Max-suppress-time:** The longest time that a route can be suppressed.

A brief description of route dampening processing: BGP speaker punishes a route once (adds to the penalty) route flapping occurs. When the penalty reaches the Suppress Limit, the route will be suppressed. When the Half-life-time reaches, the penalty is halved. When the penalty is reduced to the Reuse Limit, the route is activated again. The Max-suppress-time indicates the longest time that the route can be suppressed.

[Related Configuration](#)

Configuring BGP Route Dampening

- `bgp dampening`
- Enables BGP dampening, which is disabled by default.
- `bgp dampening half-life-time reuse suppress max-suppress-time`

Configures the parameters of route dampening.

`half-life-time` (1~45minutes): The default value is 15 minutes. A larger value means a longer flapping suppression and dampening period.

`reuse` (1~10000): The default value is 750. A smaller value means longer time for continuous stabilization before a flapping route is enabled again.

`suppress` (1~20000): The default value is 2,000. A smaller value means more flapping times allowed before suppression.

`max-suppress-time` (1~255minutes): The default value is 4*half-life-time. A larger value means longer maximum suppression time.

Displaying BGP Route Dampening

- `show ip bgp dampening flap-statistics`

Displays the flapping statistics about all routes.

- `show ip bgp dampening dampened-paths`

Displays the statistics about suppressed routes.

Resetting BGP Route Dampening

- `clear ip bgp flap-statistics`

Clears the flapping statistics about all routes that are not suppressed.

- `clear ip bgp flap-statistics address mask`

Clears the flapping statistics about specified routes (excluding suppressed routes).

- `clear ip bgp dampening [address [mask]]`

Clears the flapping statistics about all routes, including routes whose suppression is cancelled.

7.3.12 Configuring the Management Distance of BGP

The management distance is used to evaluate the reliability of various route sources. A smaller management distance means a better route.

Working Principle

Management Distance of BGP

The management distance indicates the reliability of a route source, ranging from 1 to 255. A larger value means lower reliability. BGP sets different management distances for routing information learned from different sources, including External-distance, Internal-distance and Local-distance.

- External-distance: Indicates the management distance of routes learned from EBGp peers.
- Internal-distance: Indicates the management distance of routes learned from IBGP peers.
- Local-distance: Indicates the management distance for routes learned from peers but it is considered that better routes can be learned from IGP. Generally, these routes can be indicated by the **Network Backdoor** command.

i You are not advised to change the management distance of BGP. If you really need to change the management distance of BGP, please remember:

The external-distance should be shorter than the management distances of other IGP routing protocols (OSPF and RIP).

The internal-distance and local-distance should be longer than the management distances of other IGP routing protocols.

Backdoor Route

If you prefer an IGP route but do not use an EBGp route, you can set the EBGp route as the backdoor route. By default, the management distance for routes learned from a BGP speaker for which an EBGp connection is established is 20. You can run the **network backdoor** command to set the management distance of the network information to 200 so that the same network information learned from IGP has the highest priority. The networks learned from IGP are considered backdoor networks and are not advertised.

Related Configuration

Configuring the Management Distance of BGP

You can run the **distance bgp** *external-distance internal-distance local-distance* command to configure the management distance of BGP. The value ranges from 1 to 255.

The default value of *external-distance* is 20; the default value of *internal-distance* is 200; the default value of *local-distance* is 200.

A longer management distance means a lower route priority.

Configuring a Backdoor Route

Run the **network** *network-number* **mask** *network-mask* **backdoor** command to configure a backdoor route. By default, no backdoor route is configured.

7.3.13 Configuring Multi-path Load Balancing of BGP

Multi-path load balancing means that there are multiple paths to the same network and data packets are evenly forwarded by these paths. In a routing table, one route has multiple next hops.

According to the types of equivalent routes, multi-path load balancing of BGP is classified into the following types:

- EBGp load balancing: implement load balancing for routes learned from EBGp neighbors.
- IBGP load balancing: implement load balancing for routes learned from IBGP neighbors.

i Both the IPv4 and IPv6 protocol stacks support multi-path load balancing.

i Load balancing cannot be implemented between IBGP and EBGp routes (including EBGp routes in an alliance).

Working Principle

If a BGP routing table has multiple paths to the same network, BGP calculates the route with the highest priority by default. If there are optimum multiple routes with the same priorities, BGP still selects a unique route by using comparison rules, notifies the route to the forwarding plane and controls the forwarding of data streams. After multi-path load balancing is enabled, BGP calculates a unique optimum route and also lists paths with the same priorities as equivalent routes. Then, BGP notifies the optimum route and the equivalent routes to the forwarding plane to implement load balancing.

Equivalent routes have the same basic attributes and priorities. That is, according to the optimum path selection rules of BGP, the paths have the same priorities before router-IDs are compared.

AS_PATH Loose Comparison

By default, equivalent routes must have the same AS-PATH attributes. Under such strict conditions, load balancing cannot be implemented in certain environments. In this case, you are advertised to enable the AS-PATH loose comparison mode. In the AS-PATH loose comparison mode, when other

conditions for equivalent routes are met, as long as the AS-PATH lengths of routes and the AS-PATH lengths of routes from an alliance are the same respectively, it is considered that the conditions for equivalent routes are met.

⚠ When the next hops of multiple BGP equivalent paths recur to the same IGP output interface, load balancing cannot be implemented.

Related Configuration

Configuring Multi-path Load Balancing of BGP

- `maximum-paths { ebgp | ibgp } number`

Enables the multi-path load balancing function of BGP.

number indicates the number of equivalent next hops, ranging from 1 to device capacity. The default value is 1. A larger value means more equivalent next hops allowed.

Configuring AS_PATH Loose Comparison

- `bgp bestpath as-path multipath-relax`

Enables the BGP AS-PATH loose comparison mode.

7.3.14 Configuring BGP FRR

With high-speed development of IP technologies and application of various complex services, the requirements for network security and stability become increasingly higher. Especially, certain real-time services (audios and videos) are sensitive to network running status and may be largely affected by unstable networks. Therefore, more and more focus and importance are attached to network reliability. With these requirements, the IP FRR function comes into being. It is intended to use a backup link to maintain data forwarding during route platform convergence after a faulty link is detected, in order to achieve the ideal targets of "zero delay" and "zero loss" in packet forwarding.

BGP FRR is shorted for Fast Reroute.

Working Principle

If a BGP routing table has multiple paths to the same network, BGP calculates the route with the highest priority by default. After the BGP FRR function is used, BGP selects a backup route for each optimum route. After BFD FRR detects that the master link is faulty, it switches the data to the originally calculated backup link for forwarding. After route convergence is completed, data is switched to the optimum route re-calculated for forwarding. In this way, BGP FRR can avoid route disconnection due to a link fault before BGP route convergence is completed.

-
- i** Only one backup route can be generated and the next hop of the backup route cannot be the same as that of the preferred route.
 - i** A backup next hop cannot be generated for an Equal-Cost Multi-Path Routing (ECMP) route.
 - i** BGP FRR has a lower priority than VPN FRR. That is, if VPN FRR is enabled in the VRF mode, BGP FRR takes effect only when VPN FRR fails to calculate a backup route.
-

Related Configuration

Configuring BGP FRR

Run the **bgp fast-reroute** command to enable the BGP FRR function, which is disabled by default.

Configuring a BFD Session to a BGP Neighbor

Run the **neighbor** { *neighbor-address* | *peer-group-name* } **fall-over bfd** command to configure a BFD session to a BGP neighbor, which is not configured by default.

Manually Configuring a BGP BFD Session

If the BFD session to a BGP neighbor cannot be used to fast detect the failure of the master link, you can run the **bfd bind bgp peer-ip ip-address interface interface-type interface-index source-ip ip-address** command to configure a BGP BFD session, which is not configured by default.

7.3.15 Configuring Fast Withdrawal of Specified BGP Routes

With the rapid development of IP technologies and the application of various complex services, increasingly higher requirements are imposed on the network. Customers may need to implement fast convergence on specific routes. Fast route withdrawal can meet customers' requirements for preferentially withdrawal of specific routes.

Working Principle

By default, routes are withdrawn based on preferential calculation according to the sequence in the routing table. Fast route withdrawal adds the routes meeting conditions to a high-priority queue to implement preferential processing of specified routes, so that the specified routes can be withdrawn quickly.

-
- i** The L2VPN address family supports the route map mode only.
 - i** Only one of the access list and prefix list takes effect.
-

Related Configuration

Activating the Address Family Capability

By default, only the IPv4 unicast address family capability is activated for the neighbors created in BGP configuration mode.

Run the **address-family** command to enter the corresponding address family mode.

Configuring Fast Withdrawal of Specified BGP Routes

Run the **bgp fast-withdraw { access-list { *access-list-number* | *access-list-name* } | prefix-list *prefix-list-name* | route-map *map-tag* }** command to enable fast withdrawal of specified BGP routes in the address family. This function is disabled by default.

7.3.16 Configuring BGP Multi-path Bypass Protection

The original BGP FRR supports only 1:1 protection. After the ECMP function is enabled, the FRR becomes unavailable. The multi-path bypass protection enables BGP to support 1:N protection. Users can configure BGP multi-path bypass protection to enable the system to select a backup bypass path even if ECMP is configured. When all ECMP routes fail, the bypass route becomes the primary route for forwarding.

Working Principle

By default, when there are multiple paths reachable to the same network in the BGP routing table, BGP calculates one route with the highest priority. After the BGP multi-path bypass protection function is enabled, BGP preferentially selects a backup route when ECMP routes exist. When BFD detects a fault occurring on the master path, the system switches to the selected backup path for data forwarding. After route convergence is completed, the system switches to the optimum path that is re-calculated based on routes to forward data.

- i** Only one bypass route is generated and the next hop of the bypass route cannot be the same as next hops of ECMP routes.
- i** The BGP FRR function conflicts with the BGP multi-path bypass protection function. If either of them is configured when the other function is already available, a conflict prompt will be displayed.

Related Configuration

Configuring BGP Multi-path Bypass Protection

Run the **bgp install standby-path** command to enable BGP multi-path bypass protection, which is disabled by default.

Configuring a BFD Session to a BGP Neighbor

Run the **neighbor** { *neighbor-address* | *peer-group-name* } **fall-over bfd** command to configure a BFD session to a BGP neighbor, which is not configured by default.

Configuring a BGP BFD Session Manually

When faults occurring on the master path cannot be rapidly detected via the BFD session of a BGP neighbor, run the **bfd bind bgp peer-ip** *ip-address* **interface** *interface-type* *interface-index* **source-ip** *ip-address* command to configure a BGP BFD session. No BGP BFD session is configured by default.

7.3.17 Allowing the BGP Route Reflector to Modify Route Attributes

According to RFC4456, the BGP route reflector cannot modify route attributes when reflecting routes (the purpose is to prevent route loops). However, relevant route attributes need to be modified when network traffic paths need to be re-planned. This command is configured to allow the route reflector to modify route attributes (including **Route-map** or **nexthop-self**).

Working Principle

The route reflector is not allowed to modify route attributes by default. After this command is configured, the route reflector is allowed to modify route attributes.

Related Configuration

Configuring a BGP Route Reflector and Reflection Client

No BGP route reflector is configured by default.

Run the **neighbor** *neighbor-address* **route-reflector-client** command to configure the local device as the route reflector and the neighbor as the route reflection client.

Configuring BGP Client-Client Reflection

By default, client-client route reflection is automatically enabled on the BGP route reflector, and routes received from the reflection client are allowed to be advertised to other clients.

Run the **bgp client-to-client reflection** command to enable client-client reflection. Run the **no** form of this command to disable client-client reflection.

Configuring a BGP Reflection Cluster Identifier

By default, the BGP reflection cluster identifier directly uses the BGP Route-ID. When multiple reflection clusters are deployed in one AS domain, different reflection cluster identifiers need to be configured for them.

Run the **bgp cluster-id** *cluster-id* command to manually configure a cluster identifier for the route reflector.

Allowing the BGP Route Reflector to Modify Route Attributes

By default, BGP does not allow the route reflector to modify route attributes.

Run the **bgp route-reflector attribute-change** command to allow the route reflector to modify route attributes.

7.3.18 Configuring the BGP Device to Advertise Routes with the Lowest Priority upon Restart

By default, after the device is restarted and a neighborship is established, a BGP peer can advertise routes to its neighbors. However, in some cases, for example, there are a large number of neighbors or routes when the device is started but is slow in writing entries to hardware. The neighbors have learnt routes and start forwarding traffic, but the local device has not completed entry writing to the hardware, resulting in a traffic forwarding failure. Therefore, BGP needs to advertise routes with the lowest priority after device restart. In this way, other available routes instead of routes that pass through the local device are preferentially selected.

Working Principle

Configure the BGP device to advertise routes with the lowest priority upon restart. In this way, after the BGP device is restarted and neighborships are established, the BGP device advertises routes with the lowest priority to its neighbors. Adjust the **MED** value of routes to **4,294,967,295** for EBGP neighbors and adjust the value of **local-pref** to **0** for IBGP neighbors. After all ARP entries are delivered and all forwarding entries are written into hardware, run the **clear bgp advertise lowest-priority on-startup** command to restore the priority of advertised routes. If a user does not configure to restore the priority of advertised routes, the device automatically restores the priority of advertised routes after a period of time.

Related Configuration

Configuring the BGP Device to Advertise Routes with the Lowest Priority upon Restart

- `bgp advertise lowest-priority on-startup [recover-time]`

Configure BGP to adjust the priority of advertised routes to the lowest upon device restart. The priority is not adjusted by default. **recover-time** indicates the time (in seconds) of the timer for restoring the priority of advertised routes. The value range is from **1** to **65,535** and the default value is **600**.

7.3.19 Configuring BGP Timers

You can manually configure various timers within BGP to meet the neighbor keepalive and route management requirements in different network environments.

Working Principle

BGP Neighbor Keepalive Timer

BGP uses the Keepalive timer to maintain a valid connection with a peer and uses the Holdtime timer to identify whether a peer is valid. By default, the value of the Keepalive timer is 60 seconds and the value of the Holdtime timer is 180 seconds. When a BGP connection is established between two BGP speakers, the two BGP speakers negotiate about the Holdtime timer value and select a smaller value. 1/3 of the negotiated Holdtime timer value and the configured Keepalive timer value are compared and the smaller value is used as the Keepalive timer value.

Neighbor Reconnection Timer

To reduce the impacts of frequent BGP reconnection to a neighbor on the network bandwidth, after a BGP speaker detects failure of a neighbor connection, the BGP speaker attempts to reconnect the neighbor after the connect-retry timer expires. By default, the value of the connect-retry timer is 15s.

Route Advertisement Timer

To reduce the impacts of route update packets on the network bandwidth, after a BGP speaker detects a network topology change, the BGP speaker does not advertise the route update to its neighbors immediately. Instead, the BGP speaker uses a regular update mechanism to advertise all changed routing information to its neighbors.

Related Configuration

Configuring the BGP Neighbor Keepalive Timer

- **timers bgp** keepalive holdtime

Adjusts the BGP *keepalive* and *holdtime* values for all peers.

The *keepalive* value ranges from 0 to 65,535. The default value is 60 seconds.

The *holdtime* value ranges from 0 to 65,535. The default value is 180 seconds.

- **neighbor** { address | peer-group-name } **timers** keepalive holdtime

Configures the *keepalive* and *holdtime* values used for connecting to a specified BGP peer (group).

The *keepalive* value ranges from 0 to 65,535. The default value is 60 seconds.

The *holdtime* value ranges from 0 to 65,535. The default value is 180 seconds.

Configuring the Neighbor Re-connection Timer

- **neighbor** { address | peer-group-name } **timers connect** connect-retry

Configures the *connect-retry* value used for reconnecting to a specified BGP peer (group).

The value of *connect-retry* ranges from 1 to 65,535. The default value is 15 seconds.

Configuring the Route Advertisement Timer

- **neighbor** { *address* | *peer-group-name* } **advertisemet-interval** *seconds*

Configures the minimum interval for sending route updates to a specified BGP peer (group). The value of *advertisemet-interval* ranges from 0 to 600 seconds. The default value for IBGP peers is 0 seconds and the default value for EBGP peers is 30 seconds.

- **neighbor** { *address* | *peer-group-name* } **as-origination-interval** *seconds*

Configures the minimum interval for sending local initial route updates to a specified BGP peer (group). The value of *As-origination-interval* ranges from 1 to 65,535. The default value is 1 second.

7.3.20 Configuring BGP Route Update Mechanisms

Working

Principle

BGP provides two route update mechanisms: regular-scanning update and event-triggering update. Regular-scanning update indicates that BGP uses an internal timer to start scanning regularly and update the routing table. Event-triggering update indicates that BGP starts scanning and updates the routing table when the BGP configuration commands are changed due to user configuration or the next hop of a BGP route changes.

- ❗ This function is configured based on address families and can be configured in the IPv4, IPv6, IPv4 vrf and IPv6 VRF address family modes.
- ❗ If you set the BGP route update mechanism to event-triggering update (by running the **bgp scan-rib disable** command), you must disable synchronization (by running the **no synchronization** command) and enable the BGP next-hop triggering update function (by running the **bgp nexthop trigger enable** command). On the other hand, if you enable synchronization or disable the BGP next-hop triggering update function, the BGP routing table must be updated in the regular scanning mode.

Related

Configuration

Configuring Route Update Mechanisms

- **bgp scan-rib disable**
- Sets the BGP route update mechanism to event-triggering update. Regular-scanning update is used by default.
- **bgp scan-time** *scan-time*
- Configures the regular update interval of BGP. The value of *scan-time* ranges from 5 to 60 seconds. The default value is 60 seconds.

7.3.21 Configuring the Next-Hop Triggering Update Function of BGP

The next-hop triggering update function of BGP is a method for reducing the BGP convergence time. This function is used to optimize the method for monitoring the next hop of a route to ensure that BGP can increase the BGP route convergence speed when the network topology is stable.

Working Principle

When BGP connects to a neighbor, BGP automatically monitors the next hop of the BGP route learned from the neighbor. When the next hop changes in the core routing table, BGP receives an advertisement about the next hop change and updates the BGP routing table. This optimization measure improves the BGP route convergence performance by reducing the time for detecting next-hop changes.

If this function is disabled, BGP next hop update will be discovered through regular scanning specified by scan-timer.

- i This function is configured based on address families and can be configured in the IPv4, IPv6, and IPv4 vrf address family modes.
- i **bgp nexthop trigger delay** and **bgp scan-time** control the same timer. When bgp scan is enabled (it is enabled by default and can be disabled by the **bgp scan-rib disable** command), if the value of **bgp nexthop trigger delay** is larger than 60s, bgp scan does not take effect because the scan timer is always triggered before the delay.
- ⚠ If the network environment is unstable (with frequent next-hop changes), especially with many routes, this function performs unnecessary route calculations, which consumes more CPU resources. Therefore, you are advised to disable this function in this environment.

Related Configuration

Configuring the Next-Hop Triggering Update Function of BGP

- **bgp nexthop trigger enable**

Enables the BGP next-hop triggering function, which is enabled by default.

- **bgp nexthop trigger delay *delay-time***

Configures the delay of BGP next-hop triggering update. The value of *delay-time* ranges from 0 to 100 seconds. The default value is 5 seconds.

7.3.22 Configuring BGP LOCAL AS

The Local AS function of BGP is used to configure a local AS different from a router BGP AS for a specific peer. This is similar to deploying a new virtual AS between the peer devices. When the local router BGP

AS changes, you can establish a BGP connection without changing the BGP configurations on the peer device. This function is mainly used for AS migration and merging of large networks and ensures that the device configurations in other interconnected ASs are not affected.

Working Principle

In BGP, when a local device connects to a peer, the local device advertises the local AS number to the peer by using an Open message. The peer checks whether the BGP AS number advertised is the same as the local AS number. If the AS numbers are different, the peer will deny the BGP connection. By default, the local AS in the BGP connection is a route BGP AS. However, if a local AS is configured for the peer, the configured local AS will replace the route BGP AS when a BGP connection is established between the local device and the peer.

- i** The **neighbor peer-address local-as as-num** command for configuring the BGP Local AS function can be followed by more options. For details, see the Command Reference.
- i** The BGP Local AS function is applied only to EBGP peers, but is not applied to IBGP peers and alliance EBGP peers. In addition, the BGP Local AS function has the following restrictions:
 - 1) The configured local AS cannot be the same as the remote AS of a peer.
 - 2) The local AS cannot be configured independently for a member of a peer group.
 - 3) The configured local AS cannot be the same as the route BGP AS.
 - 4) If a device is a member of an AS alliance, the local AS cannot be the same as the AS alliance number.

Related Configuration

Configuring BGP LOCAL AS

- **neighbor { address | peer-group-name } local-as as-number**

Configures a local AS for a peer. By default, no local AS is configured for any peer. The local AS of a peer is the route BGP AS.

7.3.23 Configuring BGP Capacity Protection

There are often a large number of BGP routes, which may cause overload of a device, especially for a device with small memory. Protecting BGP capacity helps avoid non-predictable running status caused by consumption of device capacity.

Working Principle

Restricting the Number of BGP Routes

Restrict the number of BGP routes by setting the maximum number of routes in a BGP address family and the maximum number of routes that can be learned by a BGP neighbor.

Entering the OVERFLOW State in case of Insufficient Memory

If the memory is insufficient, BGP can enter the OVERFLOW state. In the OVERFLOW state, BGP generates a default route pointing to a NULL interface. If a newly learned route is not a refined route other than the default route in the current routing table, the route is discarded. In other words, general newly learned routes are discarded to ensure that the system memory is stable. The purpose of not discarding all routes is to avoid route loops in the entire network. Therefore, it is safe for BGP to enter the OVERFLOW state. BGP is allowed to enter the OVERFLOW state by default.

- i** By default, BGP enters the OVERFLOW state in case of insufficient memory. If you do not want to BGP to enter the OVERFLOW state, you can run the **no overflow memory-lack** command to disable this function.
- i** In the OVERFLOW state, BGP supports only the **clear bgp { addressfamily | all } *** command at present. You can also exit from the OVERFLOW state by disabling and enabling BGP again. When the memory becomes sufficient again, BGP can also automatically exit from the OVERFLOW state.

Related Configuration

Restricting the Number of BGP Routes

- **neighbor { neighbor-address | peer-group-name } maximum-prefix maximum [threshold] [warning-only]**

Restricts the maximum number of routes that can be learned from a BGP neighbor, which is not restricted by default.

- maximum-prefix *maximum*

Restricts the maximum number of routes in a BGP address family, which is not restricted by default..

- Run the **bgp maximum-prefix maximum [vrf vrf-name]** command to restrict the maximum number of routes in the BGP global or specified VRF. When a route advertisement in an address family causes the current number of BGP routes to exceed the maximum number, a prompt indicating route overflow in the global or specified VRF is displayed, and the BGP global or specified VRF is set to the overflow state. This function is disabled by default.

Configuring BGP Overflow

- overflow memory-lack

Enable BGP to enter the OVERFLOW state in case of insufficient memory, which is enabled by default.

7.3.24 Configuring BGP GR

Graceful Restart (GR) is intended to implement uninterrupted data forwarding during restart of BGP. During active/standby switching of the management boards, the GR function keeps the network topology stable, maintains the forwarding table and ensures that key services are not interrupted.

Working Principle

Comply with RFC4724: Graceful Restart Mechanism for BGP. [BGP GR] is used in the following description to indicate the RFC.

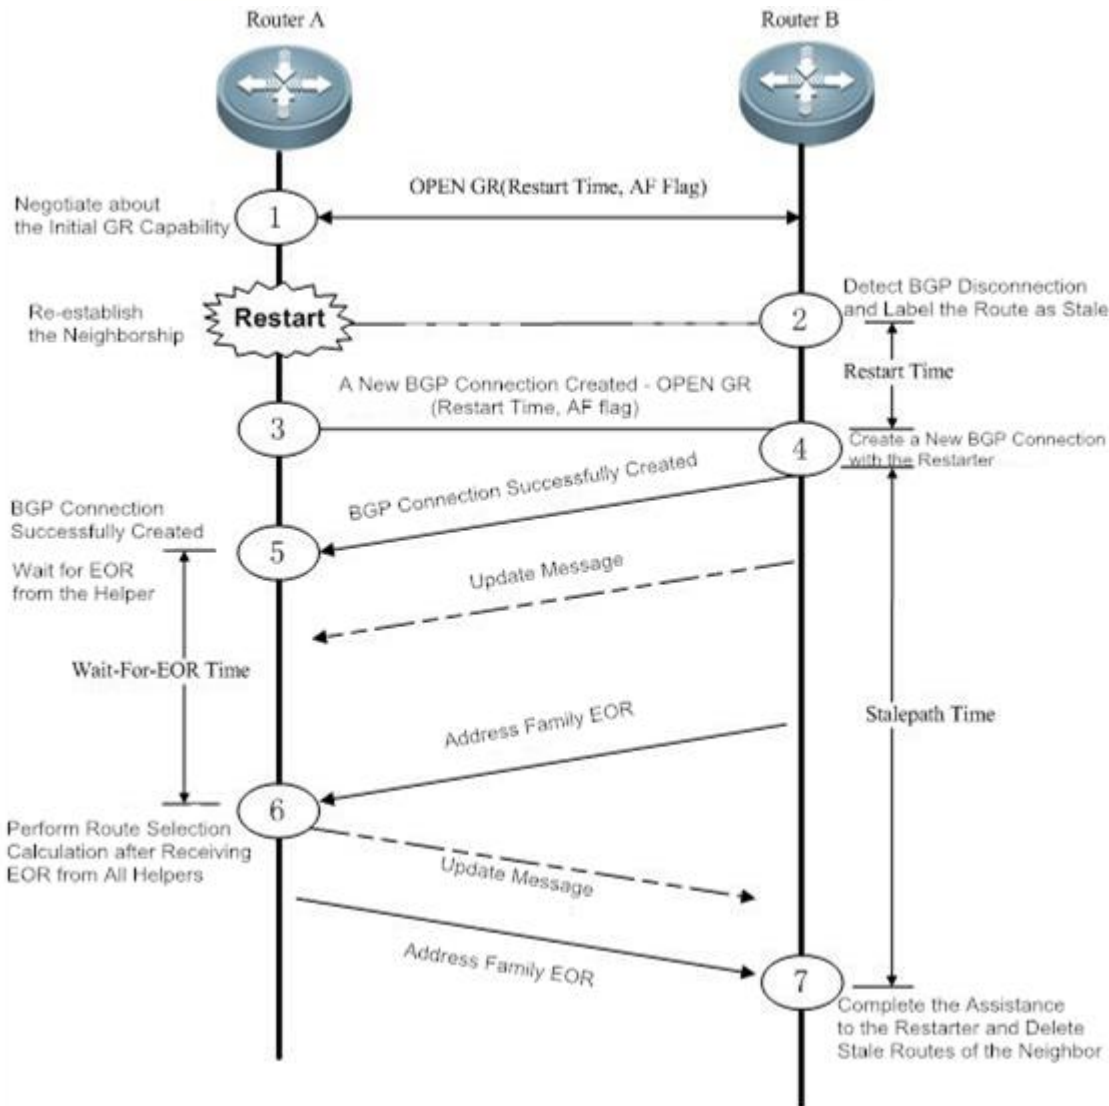
BGP GR is not an independent process, but is jointly completed by the Restarter and Helper.

- The Restarter performs restart and maintains the working capability of the route forwarding plane when the route control plane is faulty.
- The Helper is the BGP neighbor of the Restarter and helps the Restarter to complete GR.

A capability indicating GR is added to the OPEN message of BGP, which is called "Graceful Restart Capability". This capability is used by BGP to tell its neighbor it supports the graceful restart capability. During initialization of a BGP connection, two neighbors negotiate about the GR capability.

-
- The route update end flag (End-of-RIB, shorted as EOR) is added to the Update packet of BGP, which indicates that the routing information update to the neighbor is completed.
-

Figure 7-3 BGP GR Interaction Process



1. ① When BGP establishes neighborhood at the beginning, BGP uses the GR capability field in the OPEN message to negotiate about the GR capabilities of the two neighbors.
2. ② At a moment, the Restarter starts restart, and the BGP session is disconnected. The Helper detects the disconnection, keeps the route of the Restarter valid but adds the "Stale (aged but not updated)" flag to the route.
3. ③ and ④ The Restarter completes restart and connects to the Helper again.
4. ⑤ The Restarter waits for the route update message and EOR flag from the Helper.
5. ⑥ After receiving the EOR flag from all neighbors, the Restarter performs route calculation, update routing entries and then sends updated routes to the Helper.
6. ⑦ After receiving the updated routes, the Helper cancels the "Stale" flag of the routes. After receiving the EOR flag from the Restarter, the Helper deletes routes with the "Stale" flag (these routes are not updated), performs route calculation, and updates the routing entries. The entire GR process is completed.

BGP GR defines several extended and important timers:

- **Restart-Timer:** The GR Restarter advertises the time value to the GR Helper, which indicates the maximum waiting time that the GR Restarter hopes the Helper to wait before a new connection is established between them. You can run the **bgp graceful-restart restart-time** command to modify the time value.
- **Wait-For-EOR Timer:** Indicates the maximum time that the GR Restarter waits for the EOR flag from all GR Helpers. After receiving the EOR flag from all GR Helpers or after the Wait-For-EOR timer expires, the GR Restarter calculates the preferred route and updates the routing entries. You can run the **bgp update-delay** command to modify the time value.
- **StalePath Timer:** Indicates the maximum time that the GR-Helper waits for the EOR flag from the GR Restarter after a new connection is established between them. Within this period, the Helper keeps the original route of the Restarter valid. After receiving the EOR flag or after the StalePath timer expires, the Helper clears the routing entries still with the "Stale" tag. You can run the **bgp graceful-restart stalepath-time** command to modify the time value.

Related Configuration

Configuring BGP GR

- `bgp graceful-restart`

Enables the Restarter capability, which is enabled by default.

- `bgp graceful-restart restart-time time`

Sets the Restart Timer. The default value is 120 seconds.

- `bgp update-delay delay`

Sets the Wait-For-EOR Timer. The default value is 120 seconds.

- `bgp graceful-restart stalepath-time time`

Sets the StalePath Timer. The default value is 360 seconds.

- `bgp graceful-restart disable`

Disables the address family GR capability. The address family GR capability is enabled by default. After the global BGP GR is enabled, the GR capability is automatically enabled for all address families.

- ❗ When BGP GR is implemented, all BGP peers must enable the BGP GR capability. If certain peers do not support or enable GR, BGP GR may fail to be implemented. GR failure may cause a short route black-hole or route loop, which may affect the network. Therefore, you are advised to verify that all neighbors are enabled with the BGP GR capability. You can run the **show ip bgp neighbors** command to display the capabilities successfully negotiated between BGP peers and verify that

the GR capability negotiation is successful. In the BGP route configuration mode, run the **bgp graceful-restart** command to enable the BGP GR capability.

- i** The **bgp graceful-restart** command will not be applied to a successfully established BGP connection immediately. That is, when the BGP connection is in the Established state, the BGP peers will not re-negotiate about the GR capability immediately. To enable the BGP peers of the BGP connection to negotiate about the GR capability immediately, you need to forcibly restart the BGP peers to re-negotiate about the GR capability by running the **clear ip bgp 192.168.195.64** command (for example). To make GR enabling or disabling take effect immediately, you must restart the neighborship for capability negotiation, which may cause network flapping and affect normal use of users. Therefore, you can explicitly control whether to restart the neighborship.
- i** Supporting BGP GR does not mean that a device can be used as the Restarter to implement BGP GR. Whether to implement BGP GR also depends on the hardware capabilities of the device. QTECH devices must support the dual-engine hot backup when being used as the GR Restarter.
- i** The restart period configured by the **bgp graceful-restart restart-time** command should not be longer than the Hold Time of the BGP peers; otherwise, the Hold Time will be used as the restart time to be advertised to the BGP peers during GR capability negotiation.
- i** The **bgp graceful-restart disable** command is used to disable the GR capability in an address family in the address family configuration mode, which is not configured by default.

7.3.25 Configuring 4-Byte AS Numbers of BGP

A traditional AS number consists of 2 bytes, ranging from 1 to 65,535. A newly defined AS number consists of 4 bytes, ranging from 1 to 4,294,967,295. Newly defined AS numbers are used to cope with exhaustion of AS number resources.

Working Principle

4-byte AS numbers support two expression modes: the decimal mode and dot mode. The decimal mode is the same as the original expression mode, that is, expressing the 4 bytes of an AS number as decimal digits. The dot mode is expressed as ([higher 2 bytes.]lower 2 bytes). If the higher 2 bytes are 0, they will not be displayed.

For example, an AS number is 65534 in the decimal mode and is 65,534 in the dot mode (the 0 at the beginning is not displayed).

For example, an AS number is 65,536 in the decimal mode, and is 1.0 in the dot mode.

For example, an AS number is 65,538 in the decimal mode, and is 1.2 in the dot mode.

- i** Related protocols are as follows: RFC 4893 and RFC 5396.

Configuring the Display Mode of a 4-Byte AS Number

A 4-byte AS number is displayed in the decimal mode by default. You can manually set the display mode to the dot mode. After the setting, a regular expression will use the dot mode for matching 4-byte AS numbers.

Compatibility with Devices Supporting Only 2-Byte AS Numbers

With introduction of 4-byte AS numbers, BGP connections may be established between old BGP speakers supporting only 2-byte AS numbers and new BGP speakers supporting 4-byte AS numbers. If the AS where a new BGP speaker resides has a 4-byte AS number, when an old BGP speaker creates neighborhood with the new BGP speaker, the old BGP speaker uses the reserved AS number 23,456 to replace the 4-byte AS number of the new BGP speaker. In the OPEN packets sent by the new BGP speaker to the old BGP speaker, the 4-byte AS number in the **My Autonomous System** field will be replaced by 23,456. In addition, in UPDATE packets sent to the old BGP speaker, the 4-byte AS number in the AS-PATH and AGGREGATOR attributes will also be replaced by 23,456. In addition, new optional transfer attributes AS4-PATH and AS4- AGGREGATOR will be used to record the real 4-byte AS number so that the real AS-PATH and AGGREGATOR attributes can be restored when the route reaches a next new BGP speaker.

In other cases, the real AS number of the remote end is used to create neighborhood.

Related Configuration

Configuring the Display Mode of a 4-Byte AS Number

- `bgp asnotation dot`

Displays a 4-byte AS number in the dot mode. The decimal mode is used by default.

7.3.26 Configuring a Regular Expression

A regular expression is a formula that matches strings based on a template.

The formula is used to assess text data and return True or False to indicate whether the expression can correctly describe the data.

Working Principle

Regular expressions are used in BGP path attributes. The following table describes the usages of special characters in a regular expression.

Character	Symbol	Special Meaning
Period	.	Matches any single character.

Asterisk	*	Matches zero or any sequence in a string.
Plus sign	+	Matches one or any sequence in a string.
Question mark	?	Matches zero or one symbol in a string.
Caret	^	Matches the start of a string.
Dollar sign	\$	Matches the end of a string.
Underline	_	Matches the start, end and space of commas, brackets and strings.
Square brackets	[]	Matches a single character within a range.

Related Configuration

Using a Regular Expression in a show Command

- `show ip bgp regexp regexp`

Displays the BGP routing information in a specified regular expression matched by the AS-PATH attribute.

- `show ip bgp quote-regexp regexp`

Displays the BGP routing information in a regular expression within the specified double quotation marks matched by the AS-PATH attribute.

7.3.27 Configuring BGP Session Retention

By default, when an UPDATE packet is received from a neighbor, a BGP session will be disconnected if an error is detected on the multi-protocol routing attribute. This will cause flapping of the routes in all address families of this neighbor. That is, the routing error in an address family will affect the route stability in other address families.

Working Principle

After the BGP session retention function is enabled, if an error occurs in the routing attribute of an address family, only the routing information in this address family related to the neighbor is deleted. In addition, the BGP session and other address families are not affected, which enhances the stability of BGP.

recovery-time is used to configure the time for waiting for automatic route recovery, which requires that a neighbor should support the route-refresh capability. After the recovery-time, BGP sends the route-refresh message of the address family to the neighbor and re-advertises all routing information in the address family to this neighbor.

-
- i** In the session retention state, you can manually reset the neighbor to exit from the session retention state.
-

Related Configuration

Configuring BGP Session Retention

- `bgp mp-error-handle session-retain [recovery-time time]`

Enables the BGP session retention function, which is disabled by default.

recovery-time *time* configures the time for waiting for automatic route recovery, ranging from 10 to 65535 seconds. The default value is 120.

7.3.28 Configuring BGP Delayed Advertisement upon System Restart

By default, after the neighborship is established after system restart, a BGP peer can advertise route information to its neighbors. This is normal in most cases. However, in certain cases, for example, there are many neighbors or routes during startup but writing entries into the hardware is slow. In this case, the neighbors have learned the routes and started forwarding traffic, but the hardware has not completed writing of entries at the local end, which causes failure of traffic forwarding.

Working Principle

The BGP delayed advertisement upon system restart ensures that routes are not advertised to neighbors immediately after the neighborship is established upon system restart and that the routes are advertised after a period. This function has no effect on other behaviors such as route receiving performed by the neighbors. If part of the routes is not affected by the delay, configure prefix-list policy to match this part of routes so that route advertisement can be more flexible.

delay-time is used to configure the waiting time before routes are advertised to the neighbors. **startup-time** is used to configure the startup time. Within the startup-time, BGP sends routing information to the neighbors at the interval specified by **delay-time**.

-
- i** After the startup-time ends, the default route advertisement behavior recovers.
-

Related Configuration

Configuring BGP Delayed Advertisement upon System Restart

bgp initial-advertise-delay *delay-time* [*startup-time*] [*wait-for-controller*]

Enables BGP delayed advertisement upon system restart, which is disabled by default.

delay-time configures the delay time for advertising routes after the BGP neighborship is established upon system restart, ranging from 1 to 600 seconds. The default value is 1s.

startup-time configures the time range for system restart, ranging from 5 to 58,400 seconds. The delayed route advertisement mechanism is used within this range. The default value is 600s.

7.3.29 Configuring BGP Delayed Advertisement for First Routes

By default, after the neighborship is established, a BGP peer can advertise route information to its neighbors. However, neighbors with the neighborship newly established will send out the route information after a delayed period of time.

Working Principle

After BGP starts, BGP peers negotiate to establish the neighborship before sending route information (update packets). An optimum route is calculated on the local end and sent to the peer end. By default, routes are advertised directly. After a neighbor receives a better route, it updates the route information to the peer. As a result, extra route advertisement occurs. **update-delay** can be configured to shorten the route information update time.

In addition, after **update-delay** is configured on the local end, a specific neighbor sends route information to the local end, the local end selects an optimum route, and then sends out the route information immediately to other neighbors by default. In this case, extra advertisement also occurs. Therefore, the value of **update-delay** includes two parts: delay of route advertisement from the local end to a specific neighbor and delay of route advertisement from the local end to other neighbors after the specific neighbor advertises the route information to the local end.

- i** If BGP delayed advertisement upon system restart and BGP delayed advertisement for first routes are enabled at the same time, BGP delayed advertisement upon system restart takes precedence over BGP delayed advertisement for first routes.
- i** BGP GR is not affected by either BGP delayed advertisement upon system restart or BGP delayed advertisement for first routes, that is, the BGP GR route advertisement is not affected by the delay time.

Related Configuration

Creating a BGP Neighbor

By default, no neighbor is specified on a BGP speaker. You need to manually configure a BGP neighbor. Run the **neighbor** { *neighbor-address* | *peer-group-name* } **remote-as** *as-number* command to manually create a BGP neighbor and specify the AS number of the neighbor.

Configuring BGP Delayed Advertisement for First Routes

By default, BGP delayed advertisement for first routes is disabled for neighbors.

Run the **neighbor** { *neighbor-address* | *peer-group-name* } **update-delay** *delay-time* command to enable BGP delayed advertisement for first routes.

7.3.30 Configuring BGP NSR

None-Stop-Routing (NSR) is used to ensure uninterrupted routes during protocol restart upon a switchover between the active and standby management boards. During a switchover between the active and standby management boards, the NSR function ensures stable network topology, maintains neighbor status and forwarding tables, and guarantees uninterrupted key services.

Working

Principle

Different from BGP GR, BGP NSR is a function implemented individually by a BGP speaker and does not rely on neighbors.

- During a switchover between the active and standby management boards, connections between BGP neighbors are not interrupted.
- Neighbors do not perceive the restart behavior on the local end. The GR technology is used for restart restoration.

Neighbors with the BGP NSR function enabled do not need to negotiate the GR capability. In addition, neighbor assistance is not required and interworking problems do not exist as NSR is a reliability technology not perceived by neighbors.

Related

Configuration

Creating a BGP Neighbor

By default, no neighbor is specified on a BGP speaker. You need to manually configure a BGP neighbor. Run the **neighbor** { *neighbor-address* | *peer-group-name* } **remote-as** *as-number* command to manually create a BGP neighbor and specify the AS number of the neighbor.

Enabling the NSR Function for BGP Neighbors

By default, the NSR function is disabled for BGP neighbors.

Run the **neighbor** { *neighbor-address* | *peer-group-name* } **ha-mode nsr** command to enable the NSR function for BGP neighbors.

Enabling the NSR Function for BGP Neighbors Globally

By default, the NSR function is disabled.

Run the **bgp nsr** command to enable the NSR function for BGP neighbors globally.

- i** When the BGP NSR function is enabled for a neighbor, the connection to the neighbor will be reset, the tcp nss (none-stop-service) of the neighbor will be enabled, and the related neighbor and route information are backed up to the standby management board. In this case, if the NSR function is enabled after the neighborship is established, the neighborship will be re-established, causing neighbor down. It is recommended that the NSR function be enabled before the neighborship is established.
- i** The NSR function does not take effect on a neighbor that fails to establish the BGP connection, that is, if the BGP connection is not in the **Established** state, the neighbor will not execute NSR related backup operations. Only if the BGP connection is in the **Established** state, the neighbor performs NSR related backup operations.
- i** Not all devices that support the BGP NSR function can perform switchovers between the active and standby management boards. To perform switchovers between the active and standby management boards, the device hardware must meet specific requirements. NSR devices on QTECH networks must support the dual engine redundancy hot backup function.
- i** Currently, only the IPv4/IPv6 unicast address family and the EVPN address family support the NSR function. If a neighbor activates other address families except the unicast address family, NSR switchovers cannot be performed.
- !** If the BGP or TCP backup is not complete, the BGP NSR switchover may fail. When this status occurs, the BGP GR will reconnect the route. It is recommended to enable the BGP NSR and BGP GR at the same time.

7.3.31 Configuring BGP Routes to Be Recursive Only to Host Routes

BGP routes use optimal matching for route recursion by default. Therefore, BGP routes may be recursive to default routes or incorrect network segment routes, resulting in an egress or next hop error. After the function of making BGP routes recursive only to host routes is enabled, BGP routes are recursive only to 32-bit IPv4 host routes or 128-bit IPv6 host routes. This function is applied to routes learnt by neighbors of IBGP or multi-hop EBGP neighborships established using the Loopback interface.

Working Principle

After the function of making BGP routes recursive only to host routes is enabled, the function checks the validity of the next-hop addresses of IBGP or multi-hop EBGP routes via exact matching, and conducts route recursion on the IBGP or multi-hop EBGP routes via exact matching after they are written into the FIB. For example, the next-hop address of the IBGP route 192.168.2.0/24 is 1.1.1.1. The next-hop address of the IBGP route 192.168.2.0/24 is valid only when the route 1.1.1.1/32 exists, and the egress of the route written into the FIB is the egress of the route 1.1.1.1/32.

Related Configuration

Creating a BGP Neighbor

By default, no neighbor is specified for a BGP speaker. You need to manually configure a BGP neighbor.

Run the **neighbor** { *neighbor-address* | *peer-group-name* } **remote-as** *as-number* command to manually create a BGP neighbor and specify an AS number for the neighbor.

Enabling BGP Routes to Be Recursive Only to Host Routes

By default, recursion is performed on BGP routes via exact matching.

Run the **bgp recursion host** command to enable BGP routes to be recursive only to host routes.

7.3.32 Configuring Outbound Loop Detection for a BGP Neighbor

By default, BGP conducts loop detection on BGP routes when receiving the BGP routes from a neighbor. When the **AS Path** attribute carried in a BGP route contains the local AS number, BGP filters out the BGP route. The outbound loop detection function of a neighbor is to conduct loop detection on routes in advance when the routes are transmitted to a neighbor, so as to filter out loop routes.

Working Principle

When sending a route to an EBGP neighbor, the device judges whether the **AS Path** attribute carried in the BGP route contains the AS number of the neighbor. If yes, the route is looped and the device does not send the route to the EBGP neighbor.

Related Configuration

Creating a BGP Neighbor

By default, no neighbor is specified for a BGP speaker. You need to manually configure a BGP neighbor. Run the **neighbor** { *neighbor-address* | *peer-group-name* } **remote-as** *as-number* command to manually create a BGP neighbor and specify an AS number for the neighbor.

Enabling Outbound Loop Detection for a Neighbor

The outbound loop detection is disabled for a neighbor by default.

Run the **neighbor** { *neighbor-address* | *peer-group-name* } **as-loop-check out** command to enable the outbound loop detection for the BGP neighbor.

7.3.33 Shutting Down BGP Connections Gracefully

The RFC 6198 *Requirements for the Graceful Shutdown of BGP Sessions* defines a requirement for shutting down BGP connections gracefully, and requests that a BGP neighbor shut down connections under the precondition that service traffic is not interrupted or is interrupted for a very short period of time.

Working Principle

This feature adopts the "Make-Before-Break" mode to shut down BGP connections, to ensure that service traffic is not interrupted or is interrupted for a very short period of time in the process from the feature configuration to actual BGP connection shutdown. The steps are as follows:

- The feature enables the device to advertise a route with the lowest priority (the value of **local-preference** is **0** or the value of **med** is **4,294,967,295**) to a neighbor and carry the **gshut community** attribute in the route so that the neighbor updates routes and switch traffic to a backup link or other equivalent links in advance.
- The device shuts down the BGP connection with the neighbor after a period of time.

The delay period of time for BGP connection shutdown can be manually specified or automatically calculated. The automatically calculated delay time is 1 second for 1000 routes.

- i** If the **local-preference** or **med** attribute of the route with the lowest priority cannot enable the neighbor to switch traffic to backup links or other equivalent links in advance, configure an inbound policy on the neighbor to change the priority of routes carrying the **gshut community** or specified **community** attribute or to filter such routes.

Related Configuration

Shutting Down All Connections of a BGP Instance Gracefully

By default, all connections of BGP instances are not shut down.

Run the **bgp shutdown graceful [community value] [delay time]** command to shut down all connections of a BGP instance.

Shutting Down Connections of the BGP Neighbor Gracefully

By default, the connections of BGP neighbors are not shut down.

Run the **neighbor { neighbor-address | peer-group-name } shutdown graceful [community value] [delay time]** command to disable connections of a BGP neighbor.

7.3.34 Configuring Enhanced VPN Route Import

Working

Principle

During inter-VRF route import, import of L3VPN remote routes to VRF, or import of EVPN routes to the IP route table, only routes with preferred next hops are imported by default.

The enhanced VPN route import function is an extension of the inter-import of the preceding routes. It enables all routes with next hops or equivalent next hops to be imported.

Related

Configuration

Configuring the Route Import Policy

By default, only routes with preferred next hops are imported.

Run the **import path selection { all | bestpath | multipath }** command to import all routes with next hops, routes with preferred next hops, or routes with equivalent next hops.

7.3.35 BGP Route Update Group

The BGP route update group function is used to enhance the performance for advertising routes to neighbors.

Working

Principle

The BGP route update group function automatically classifies neighbors with the same outbound policy to the same update group. When routes are sent to neighbors, the update packet is encapsulated based on the update group and sent to all neighbors in the update group. In this case, the update packet is encapsulated for once and sent multiple times, improving the performance of route advertisement to neighbors.

7.3.36 Configuring BGP EVPN

As a layer-2 VPN technology, the Ethernet Virtual Private Network (EVPN) uses the MB-BGP on the control plane to distribute layer-2 and layer-3 routing information and uses the Virtual eXtensible Local Area Network (VXLAN) on the data plane to encapsulate service traffic.

Working Principle

The EVPN is a layer-2 VPN technology extended based on the MP-BGP. The MP-BGP expands the EVPN address family to transfer the required layer-2 and layer-3 routing information, and support multiple types of data planes. At present, only the VXLAN can serve as the data plane of the EVPN.

The MP-BGP adds not only the EVPN address family, but also the EVPN network layer reachability information, which includes the following:

1. Ethernet Auto-discovery Route: Automatically discovers Ethernet Segment (ES) information in a multi-homed network.
2. MAC/IP Advertisement Route: Advertises the MAC address and host ARP information.
3. Inclusive Multicast Ethernet Tag (IMET) Route: Automatically discovers VXLAN tunnel end points (VTEPs) and creates VXLAN tunnels.
4. Ethernet Segment Route: Advertises ES information and information about connected VTEPs.
5. IP Prefix Route: Advertises routing information of IP network segments.

At present, Ethernet Auto-discovery Route and Ethernet Segment Route are not supported.

The preceding EVPN routing information includes the RD (for EVPN routes of different VNIs) and route target (for importing EVPN routes between different devices).

Related Configuration

Configuring EVPN

By default, the EVPN mode is not activated.

Run the **evpn** command to enter the EVPN mode.

Creating an EVI Instance

By default, no EVI instance is configured.

Run the **vni vni-id** command to create an EVI instance and enter the EVI instance configuration mode.

Run the **vni range vni-id-list** command to create EVI instances in batches and enter the EVI instance batch configuration mode.

Configuring RD

By default, no RD is configured in evpn-vni mode.

Run the **rd** { **auto** | *rd_value* } command to add RD.

Configuring Route Target

By default, no route target is configured in evpn-vni mode.

Run the **route-target** { **import** | **export** | **both** } command to configure route target.

Configuring a Route Map for EVI Instances

By default, no route map is configured for EVI instances.

Run the **export map** *routemap-name* command to configure the route map for exporting the extended community attribute of EVPN routes from the local device to the remote device.

Run **import map** *routemap-name* command to configure the route map for importing the remote EVPN routes to the local VNI instance.

Configuring an EVPN Address Family

By default, no EVPN address family is activated.




Run the **address-family l2vpn evpn** command to activate an EVPN address family.

Configuring MAC Mobility Parameters


By default, if five MAC address moves are detected within 180s, it is considered that a MAC address conflict occurs.

Run the **bgp mac-mobility timer count** command to adjust the MAC mobility detection parameters.

7.3.37 Other Related Configurations

-  For configuration and application of BGP MCE, see section "VRF Configuration Guide".
-  For configuration and application of BGP L2VPN, see section "L2VPN Configuration Guide".
-  For configuration and application of the BGP MDT address family, see section "Multicast VPN (MD Configuration Guide)".

7.4 Configuration

Configuration	Description and Command
	 (Mandatory) It is used to create a BGP neighbor.

Configuring a BGP Peer (Group)	router bgp	Enables BGP.
	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>	Creates a BGP neighbor.
Configuring MD5 Authentication	⚠ (Optional) It is used to perform encrypted authentication for the BGP neighbor.	
	neighbor { <i>peer-address</i> <i>peer-group-name</i> } password [0 7] <i>string</i>	Configures the password for encryption.
Configuring a Route Reflector	⚠ (Optional) It is used to reduce the number of BGP neighbor connections.	
	neighbor { <i>peer-address</i> <i>peer-group-name</i> } route-reflector-client	Specifies a peer (group) as a reflector client.
Configuring an AS Alliance	⚠ (Optional) It is used to reduce the number of BGP neighbor connections.	
	bgp confederation identifier <i>as-number</i>	Configures the BGP alliance ID.
	bgp confederation peers <i>as-number</i> [... <i>as-number</i>]	Configures a BGP alliance neighbor.
Configuring Multi-path Load Balancing of BGP	⚠ (Optional) It is used to implement multi-path load balancing.	
	maximum-paths ibgp <i>number</i>	Configures IBGP load balancing.
	maximum-paths ebgp <i>number</i>	Configures EBGP load balancing.
	bgp bestpath as-path multipath-relax	Enables the BGP AS-PATH loose comparison mode.
Configuring EBG FRR	⚠ (Optional) It is used to increase the convergence speed when a network fault occurs.	
	bgp fast-reroute	Configures BGP FRR.
	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } fall-over bfd	Configures a BFD session to a BGP neighbor.
	⚠ (Optional) It is used to fast withdraw the specified BGP routes.	

Configuring Fast Withdrawal of Specified BGP Routes	bgp fast-withdraw { access-list { <i>access-list-number</i> <i>access-list-name</i> } prefix-list <i>prefix-list-name</i> route-map <i>map-tag</i> }	Configures fast withdrawal of specified BGP routes.
Configuring Local ASs	⚠ (Optional) It is used for transitional deployment during network migration.	
	neighbor { <i>peer-address</i> <i>peer-group-name</i> } local-as <i>as-number</i> [no-prepend [replace-as [dual-as]]]	Configures the local AS for a BGP neighbor.
Configuring BGP GR	⚠ (Recommended) It is used to improve the network reliability.	
	bgp graceful-restart	Enables the BGP GR capability.
	bgp graceful-restart restart-time <i>restart-time</i>	Configures the maximum time for BGP GR.
	bgp graceful-restart stalepath-time <i>time</i>	Configures the maximum retention time for BGP stable route.
Configuring a BGP IPv6 Address Family	⚠ (Optional) It is used to deploy an IPv6 network by using BGP.	
	address-family ipv6 unicast	Enters the BGP IPv6 unicast configuration mode.
	neighbor { <i>peer-address</i> <i>peer-group-name</i> } activate	Activates the address family capability of a BGP neighbor in the current configuration mode.
Configuring Interconnection with Devices Supporting Only 2-Byte AS Numbers	⚠ Optional. It is used for interconnecting with an old device that supports only 2-byte AS numbers.	
	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>	Creates a BGP neighbor.

7.4.1 Configuring a BGP Peer (Group)

Configuration

Effect

- Configure BGP and create IBGP and EBGP neighbors.

Notes

- If an IBGP neighbor is not directly connected, you need to configure IGP or a static routing protocol to implement interconnection.
- If an EBGP neighbor is not directly connected, you need to configure the **ebgp-multihop** parameter for the neighbor.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring a Source Interface for a BGP Neighbor

- (Optional) Perform this configuration in the BGP configuration mode. By default, BGP automatically selects a local interface that reaches the destination IP address of a peer as the source interface.

i For an IBGP neighbor, you are advised to use a Loopback interface as the source interface.

Verification

- Run the **show** command to display the neighbor status.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	<code>neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i></code>
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Creating a Source Interface for a BGP Neighbor


Command	<code>neighbor { <i>peer-address</i> <i>peer-group-name</i> } update-source { <i>interface-type interface-number</i> <i>address</i> }</code>
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>interface-type interface-number</i>: Indicates an interface name.</p> <p><i>address</i>: Directly specifies the network interface address used for creating a BGP connection.</p>
Command Mode	BGP configuration mode
Usage Guide	The source interface of a neighbor must be a local valid interface or address.

Configuration

Example

Configuring a BGP Peer (Group)

Scenario
Figure 7-4



**Configurati
on Steps**

- Enable BGP on all devices and set the AS numbers as shown in Figure 7-4.
- Configure a loopback interface on A, B, and C and create an IBGP neighbor based on the loopback interface.
- Create an EBGP neighborhood by using the directly connected interfaces on C and D.
- Create an IBGP peer group on C.

A

```
A# configure terminal
A(config)# interface loopback 0
A(config-if-Loopback 0)# ip address 10.1.1.1 255.255.255.255
A(config-if-Loopback 0)# exit
A(config)# interface GigabitEthernet 0/1
A(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0
A(config-if-GigabitEthernet 0/1)# exit
A(config)# router bgp 65536
A(config-router)# neighbor 10.1.1.3 remote-as 65536
A(config-router)# neighbor 10.1.1.3 update-source loopback 0
```

B

```
B# configure terminal
B(config)# interface loopback 0
B(config-if-Loopback 0)# ip address 10.1.1.2 255.255.255.255
B(config-if-Loopback 0)# exit
B(config)# interface GigabitEthernet 0/1
B(config-if-GigabitEthernet 0/1)# ip address 192.168.2.2 255.255.255.0
B(config-if-GigabitEthernet 0/1)# exit
```

	<pre> B(config)# router bgp 65536 B(config-router)# neighbor 10.1.1.3 remote-as 65536 B(config-router)# neighbor 10.1.1.3 update-source loopback 0 </pre>
C	<pre> C# configure terminal C(config)# interface loopback 0 C(config-if-Loopback 0)# ip address 10.1.1.3 255.255.255.255 C(config-if-Loopback 0)# exit C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.1.3 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# interface GigabitEthernet 0/2 C(config-if-GigabitEthernet 0/2)# ip address 192.168.2.3 255.255.255.0 C(config-if-GigabitEthernet 0/2)# exit C(config)# interface GigabitEthernet 0/3 C(config-if-GigabitEthernet 0/3)# ip address 192.168.3.3 255.255.255.0 C(config-if-GigabitEthernet 0/3)# exit C(config)# router bgp 65536 C(config-router)# neighbor ibgp-group peer-group C(config-router)# neighbor ibgp-group remote-as 65536 C(config-router)# neighbor ibgp-group update-source loopback 0 C(config-router)# neighbor 10.1.1.1 peer-group ibgp-group C(config-router)# neighbor 10.1.1.2 peer-group ibgp-group C(config-router)# neighbor 192.168.3.4 remote-as 65537 </pre>
D	<pre> D# configure terminal D(config)# interface GigabitEthernet 0/1 D(config-if-GigabitEthernet 0/1)# ip address 192.168.3.4 255.255.255.0 D(config-if-GigabitEthernet 0/1)# exit D(config)# router bgp 65537 D(config-router)# neighbor 192.168.3.3 remote-as 65536 </pre>
Verification	Run the show command to display the BGP neighbor status.

A

```
A# show ip bgp neighbor
BGP neighbor is 10.1.1.3, remote AS 65536, local AS 65536, internal link
BGP version 4, remote router ID 10.1.1.3
BGP state = Established, up for 00:00:05
Last read      , hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received (old and new)
  Four-octets ASN Capability: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 2 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:1
  refresh message:0 dynamic cap:0 notifications:0
Sent 2 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:1
  refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 0 seconds
Update source is Loopback 0

For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 0, Offset 0, Mask 0x1
0 accepted prefixes
0 announced prefixes

Connections established 1; dropped 0
Local host: 10.1.1.1, Local port: 1039
Foreign host: 10.1.1.3, Foreign port: 179
Nexthop: 10.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

Last Reset: , due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

B

```
B# show ip bgp neighbor
BGP neighbor is 10.1.1.3, remote AS 65536, local AS 65536, internal link
BGP version 4, remote router ID 10.1.1.3
BGP state = Established, up for 00:00:07
Last read , hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received (old and new)
  Four-octets ASN Capability: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 2 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:1
  refresh message:0 dynamic cap:0 notifications:0
Sent 2 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:1
  refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 0 seconds
Update source is Loopback 0

For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 0, Offset 0, Mask 0x1
0 accepted prefixes
0 announced prefixes

Connections established 1; dropped 0
Local host: 10.1.1.2, Local port: 1041
Foreign host: 10.1.1.3, Foreign port: 179
Nexthop: 10.1.1.2
Nexthop global: ::
```

	<p>Nexthop local: ::</p> <p>BGP connection: non shared network</p> <p>Last Reset: , due to BGP Notification received</p> <p>Notification Error Message: (Cease/Other Configuration Change.)</p>
C	<pre> C# show ip bgp neighbor BGP neighbor is 10.1.1.1, remote AS 65536, local AS 65536, internal link Member of peer-group ibgp-group for session parameters BGP version 4, remote router ID 10.1.1.1 BGP state = Established, up for 00:01:13 Last read , hold time is 180, keepalive interval is 60 seconds Neighbor capabilities: Route refresh: advertised and received (old and new) Four-octets ASN Capability: advertised and received Address family IPv4 Unicast: advertised and received Received 3 messages, 0 notifications, 0 in queue open message:1 update message:0 keepalive message:2 refresh message:0 dynamic cap:0 notifications:0 Sent 3 messages, 0 notifications, 0 in queue open message:1 update message:0 keepalive message:2 refresh message:0 dynamic cap:0 notifications:0 Route refresh request: received 0, sent 0 Minimum time between advertisement runs is 0 seconds Update source is Loopback 0 For address family: IPv4 Unicast BGP table version 1, neighbor version 1 Index 1, Offset 0, Mask 0x2 ibgp-group peer-group member 0 accepted prefixes 0 announced prefixes Connections established 1; dropped 0 </pre>

Local host: 10.1.1.3, Local port: 179
Foreign host: 10.1.1.1, Foreign port: 1039
Nexthop: 10.1.1.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 10.1.1.2, remote AS 65536, local AS 65536, internal link
Member of peer-group ibgp-group for session parameters
BGP version 4, remote router ID 10.1.1.2
BGP state = Established, up for 00:01:17
Last read , hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
Route refresh: advertised and received (old and new)
Four-octets ASN Capability: advertised and received
Address family IPv4 Unicast: advertised and received
Received 3 messages, 0 notifications, 0 in queue
open message:1 update message:0 keepalive message:2
refresh message:0 dynamic cap:0 notifications:0
Sent 3 messages, 0 notifications, 0 in queue
open message:1 update message:0 keepalive message:2
refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 0 seconds
Update source is Loopback 0

For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
ibgp-group peer-group member
0 accepted prefixes
0 announced prefixes

Connections established 1; dropped 0
Local host: 10.1.1.3, Local port: 179
Foreign host: 10.1.1.2, Foreign port: 1041
Nexthop: 10.1.1.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 192.168.3.4, remote AS 65536, local AS 65536, internal link
Member of peer-group ibgp-group for session parameters
BGP version 4, remote router ID 192.168.3.4
BGP state = Established, up for 00:01:01
Last read , hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
Route refresh: advertised and received (old and new)
Four-octets ASN Capability: advertised and received
Address family IPv4 Unicast: advertised and received
Received 3 messages, 0 notifications, 0 in queue
open message:1 update message:0 keepalive message:2
refresh message:0 dynamic cap:0 notifications:0
Sent 3 messages, 0 notifications, 0 in queue
open message:1 update message:0 keepalive message:2
refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 0 seconds
Update source is Loopback 0

For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
ibgp-group peer-group member
0 accepted prefixes
0 announced prefixes

Connections established 1; dropped 0
Local host: 192.168.3.3, Local port: 179
Foreign host: 192.168.3.4, Foreign port: 1018
Nexthop: 192.168.3.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

D

```
D# show ip bgp neighbor
BGP neighbor is 192.168.3.3, remote AS 65536, local AS 65536, internal link
Member of peer-group ibgp-group for session parameters
BGP version 4, remote router ID 10.1.1.3
BGP state = Established, up for 00:01:01
Last read      , hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received (old and new)
  Four-octets ASN Capability: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 3 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:2
  refresh message:0 dynamic cap:0 notifications:0
Sent 3 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:2
  refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 0 seconds
Update source is Loopback 0

For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
ibgp-group peer-group member
```

```
0 accepted prefixes
0 announced prefixes

Connections established 1; dropped 0
Local host: 192.168.3.4, Local port: 1018
Foreign host: 192.168.3.3, Foreign port: 179
Nexthop: 192.168.3.4
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

Common

Errors

- IGP is not enabled and the interconnection between the local loopback address and the loopback address on the IBGP neighbor fails, which causes that the neighbor fails to be created.
- `ebgp-multihop` is not configured when an EBGP is not directly connected, which causes that a TCP connection fails to be created.

7.4.2 Configuring MD5 Authentication

Configuration

Effect

- Configure MD5 for encrypted authentication between EBGP and IBGP neighbors.

Notes

- If an IBGP neighbor is not directly connected, you need to configure IGP or a static routing protocol to implement interconnection.
- If an EBGP neighbor is not directly connected, you need to configure the **`ebgp-multihop`** parameter for the neighbor.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Verification

- Run the **show** command to display the neighbor status.

Related Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>peer-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuring an MD5 Password for a BGP Neighbor

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } password [0 7] <i>string</i>
Parameter Description	<i>peer-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.

	<p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p>0: Displays a password not encrypted.</p> <p>7: Displays a password encrypted.</p> <p><i>string</i>: Indicates a password for TCP MD5 authentication, consisting of a maximum of 80 characters.</p>
Command Mode	BGP configuration mode
Usage Guide	The same passwords must be configured on the two ends of a BGP neighborship.

Configuration Example

Configuring BGP MD5 Authentication

<p>Scenario Figure 7-5</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Enable BGP on all devices and set the AS numbers as shown in Figure 7-5. Configure a loopback interface on B and C and create an IBGP neighbor based on the loopback interface. Create an EBGP neighborship by using the directly connected interfaces on A and B. Configure the passwords on A, B and C for their neighbors.
<p>A</p>	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0 A(config-if-GigabitEthernet 0/1)# exit A(config)# router bgp 65537</pre>

	<pre>A(config-router)# neighbor 192.168.1.2 remote-as 65536 A(config-router)# neighbor 192.168.1.2 password 7 ebgpneighbor</pre>
B	<pre>B# configure terminal B(config)# interface loopback 0 B(config-if-Loopback 0)# ip address 10.1.1.1 255.255.255.255 B(config-if-Loopback 0)# exit B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip address 192.168.1.2 255.255.255.0 B(config-if-GigabitEthernet 0/1)# exit B(config)# interface GigabitEthernet 0/2 B(config-if-GigabitEthernet 0/2)# ip address 192.168.2.2 255.255.255.0 B(config-if-GigabitEthernet 0/2)# exit B(config)# router bgp 65536 B(config-router)# neighbor 10.1.1.2 remote-as 65536 B(config-router)# neighbor 10.1.1.2 update-source loopback 0 B(config-router)# neighbor 10.1.1.2 password ibgpneighbor B(config-router)# neighbor 192.168.1.1 remote-as 65537 B(config-router)# neighbor 192.168.1.1 password 7 ebgpneighbor</pre>
C	<pre>C# configure terminal C(config)# interface loopback 0 C(config-if-Loopback 0)# ip address 10.1.1.2 255.255.255.255 C(config-if-Loopback 0)# exit C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.2.3 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# router bgp 65536 C(config-router)# neighbor 10.1.1.1 remote-as 65536 C(config-router)# neighbor 10.1.1.1 update-source loopback 0 C(config-router)# neighbor 10.1.1.1 password ibgpneighbor</pre>
Verification	Run the show command to display the BGP neighbor status.

A

```
A#show ip bgp neighbors
BGP neighbor is 192.168.1.2, remote AS 65536, local AS 65537, external link
  BGP version 4, remote router ID 10.1.1.1
  BGP state = Established, up for 00:04:54
  Last read      , hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 7 messages, 0 notifications, 0 in queue
    open message:1 update message:0 keepalive message:6
    refresh message:0 dynamic cap:0 notifications:0
  Sent 7 messages, 0 notifications, 0 in queue
    open message:1 update message:0 keepalive message:6
    refresh message:0 dynamic cap:0 notifications:0
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast
  BGP table version 1, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  0 accepted prefixes
  0 announced prefixes

Connections established 2; dropped 1
Local host: 192.168.1.1, Local port: 1026
Foreign host: 192.168.1.2, Foreign port: 179
Nexthop: 192.168.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:04:54, due to BGP Notification sent
```

Notification Error Message: (Cease/Administratively Reset.)

B

```
B# show ip bgp neighbors
BGP neighbor is 10.1.1.2, remote AS 65536, local AS 65536, internal link
BGP version 4, remote router ID 10.1.1.2
BGP state = Established, up for 00:04:01
Last read      , hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received (old and new)
  Four-octets ASN Capability: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 8 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:7
  refresh message:0 dynamic cap:0 notifications:0
Sent 8 messages, 0 notifications, 0 in queue
  open message:1 update message:0 keepalive message:7
  refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast
BGP table version 1, neighbor version 0
Index 1, Offset 0, Mask 0x2
0 accepted prefixes
0 announced prefixes

Connections established 2; dropped 1
Local host: 10.1.1.1, Local port: 179
Foreign host: 10.1.1.2, Foreign port: 1038
Nexthop: 10.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

Last Reset: 00:05:27, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)

BGP neighbor is 192.168.1.1, remote AS 65537, local AS 65536, external link

BGP version 4, remote router ID 192.168.1.1

BGP state = Established, up for 00:05:27

Last read , hold time is 180, keepalive interval is 60 seconds

Neighbor capabilities:

Route refresh: advertised and received (old and new)

Four-octets ASN Capability: advertised and received

Address family IPv4 Unicast: advertised and received

Received 8 messages, 0 notifications, 0 in queue

open message:1 update message:0 keepalive message:7

refresh message:0 dynamic cap:0 notifications:0

Sent 8 messages, 0 notifications, 0 in queue

open message:1 update message:0 keepalive message:7

refresh message:0 dynamic cap:0 notifications:0

Route refresh request: received 0, sent 0

Minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast

BGP table version 1, neighbor version 0

Index 1, Offset 0, Mask 0x2

0 accepted prefixes

0 announced prefixes

Connections established 2; dropped 1

Local host: 192.168.1.2, Local port: 179

Foreign host: 192.168.1.1, Foreign port: 1026

Nexthop: 192.168.1.2

Nexthop global: ::

Nexthop local: ::

BGP connection: non shared network

Last Reset: 00:05:27, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)

C

```
C# show ip bgp neighbors
BGP neighbor is 10.1.1.1, remote AS 65536, local AS 65536, internal link
  BGP version 4, remote router ID 10.1.1.1
  BGP state = Established, up for 00:04:01
  Last read      , hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 8 messages, 0 notifications, 0 in queue
    open message:1 update message:0 keepalive message:7
    refresh message:0 dynamic cap:0 notifications:0
  Sent 8 messages, 0 notifications, 0 in queue
    open message:1 update message:0 keepalive message:7
    refresh message:0 dynamic cap:0 notifications:0
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast
  BGP table version 1, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  0 accepted prefixes
  0 announced prefixes

Connections established 2; dropped 1
Local host: 10.1.1.2, Local port: 1038
Foreign host: 10.1.1.1, Foreign port: 179
Nexthop: 10.1.1.2
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network
Last Reset: 00:05:27, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)
```

Common Errors

- The passwords for MD5 encrypted authentication at the two ends of a BGP neighborhood are different.

7.4.3 Configuring a Route Reflector

Configuration Effect

- Configure a route reflector in the IBGP environment to reduce the number of BGP neighbor connections.

Notes

- If an IBGP neighbor is not directly connected, you need to configure IGP or a static routing protocol to implement interconnection.

Configuration Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Creating a BGP Reflector

- (Mandatory) Perform this configuration in the BGP configuration mode.

Verification

- Run the **show** command to display the neighbor status.

Related Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

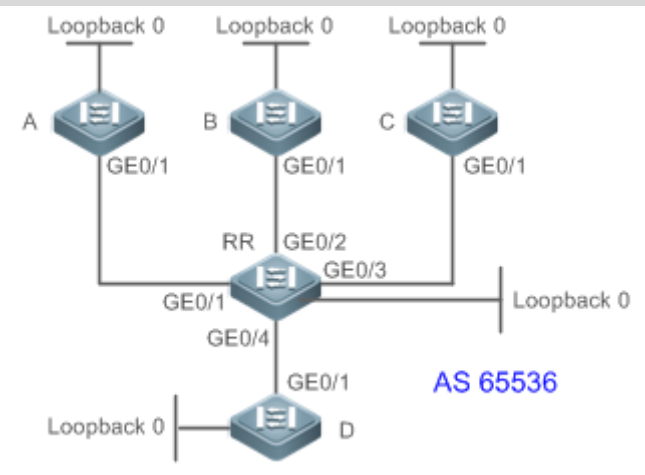
Creating a BGP Reflector

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } route-reflector-client
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p>
Command Mode	BGP configuration mode

Usage Guide	-
-------------	---

Configuration Example

Configuring a BGP Route Reflector

<p>Scenario Figure 7-6</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Enable BGP on all devices and set the AS numbers as shown in Figure 7-6. Configure a loopback interface on all devices and create an IBGP neighborhood by using the loopback interface according to the connection lines as shown in Figure 7-6. Configure route reflection on the device RR and specify A, B, C and D as reflector clients.
<p>A</p>	<pre>A# configure terminal A(config)# interface loopback 0 A(config-if-Loopback 0)# ip address 10.1.1.1 255.255.255.255 A(config-if-Loopback 0)# exit A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0 A(config-if-GigabitEthernet 0/1)# exit A(config)# router bgp 65536 A(config-router)# neighbor 10.1.1.5 remote-as 65536 A(config-router)# neighbor 10.1.1.5 update-source loopback 0 A(config-router)# network 192.168.1.0 mask 255.255.255.0</pre>

B	<pre>B# configure terminal B(config)# interface loopback 0 B(config-if-Loopback 0)# ip address 10.1.1.2 255.255.255.255 B(config-if-Loopback 0)# exit B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip address 192.168.2.2 255.255.255.0 B(config-if-GigabitEthernet 0/1)# exit B(config)# router bgp 65536 B(config-router)# neighbor 10.1.1.5 remote-as 65536 B(config-router)# neighbor 10.1.1.5 update-source loopback 0</pre>
C	<pre>C# configure terminal C(config)# interface loopback 0 C(config-if-Loopback 0)# ip address 10.1.1.3 255.255.255.255 C(config-if-Loopback 0)# exit C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.3.3 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# router bgp 65536 C(config-router)# neighbor 10.1.1.5 remote-as 65536 C(config-router)# neighbor 10.1.1.5 update-source loopback 0</pre>
D	<pre>C# configure terminal C(config)# interface loopback 0 C(config-if-Loopback 0)# ip address 10.1.1.4 255.255.255.255 C(config-if-Loopback 0)# exit C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.4.4 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# router bgp 65536 C(config-router)# neighbor 10.1.1.5 remote-as 65536 C(config-router)# neighbor 10.1.1.5 update-source loopback 0</pre>

RR	<pre> RR# configure terminal RR(config)# interface loopback 0 RR(config-if-Loopback 0)# ip address 10.1.1.5 255.255.255.255 RR(config-if-Loopback 0)# exit RR(config)# interface GigabitEthernet 0/1 RR(config-if-GigabitEthernet 0/1)# ip address 192.168.1.5 255.255.255.0 RR(config-if-GigabitEthernet 0/1)# exit RR(config)# interface GigabitEthernet 0/2 RR(config-if-GigabitEthernet 0/2)# ip address 192.168.2.5 255.255.255.0 RR(config-if-GigabitEthernet 0/2)# exit RR(config)# interface GigabitEthernet 0/3 RR(config-if-GigabitEthernet 0/3)# ip address 192.168.3.5 255.255.255.0 RR(config-if-GigabitEthernet 0/3)# exit RR(config)# interface GigabitEthernet 0/4 RR(config-if-GigabitEthernet 0/4)# ip address 192.168.4.5 255.255.255.0 RR(config-if-GigabitEthernet 0/4)# exit RR(config)# router bgp 65536 RR(config-router)# neighbor 10.1.1.1 remote-as 65536 RR(config-router)# neighbor 10.1.1.1 update-source loopback 0 RR(config-router)# neighbor 10.1.1.1 route-reflector-client RR(config-router)# neighbor 10.1.1.2 remote-as 65536 RR(config-router)# neighbor 10.1.1.2 update-source loopback 0 RR(config-router)# neighbor 10.1.1.2 route-reflector-client RR(config-router)# neighbor 10.1.1.3 remote-as 65536 RR(config-router)# neighbor 10.1.1.3 update-source loopback 0 RR(config-router)# neighbor 10.1.1.3 route-reflector-client RR(config-router)# neighbor 10.1.1.4 remote-as 65536 RR(config-router)# neighbor 10.1.1.4 update-source loopback 0 RR(config-router)# neighbor 10.1.1.4 route-reflector-client </pre>
Verification	Run the show command to display the BGP neighbor status.
RR	RR# show ip bgp summary

BGP router identifier 10.1.1.5, local AS number 65536

BGP table version is 1

0 BGP AS-PATH entries

0 BGP Community entries

1 BGP Prefix entries (Maximum-prefix:4294967295)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.1	4	65536	8	9	1	0	0	00:05:11	1
10.1.1.2	4	65536	9	9	1	0	0	00:05:24	0
10.1.1.3	4	65536	8	7	1	0	0	00:05:10	0
10.1.1.4	4	65536	9	8	1	0	0	00:05:14	0

RR# show ip bgp

BGP table version is 1, local router ID is 10.1.1.5

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

S Stale, b - backup entry

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i192.168.1.0	10.1.1.1	0	100	0	i

Total number of prefixes 1

D

D# show ip bgp summary

BGP router identifier 10.1.1.4, local AS number 65536

BGP table version is 1

0 BGP AS-PATH entries

0 BGP Community entries

1 BGP Prefix entries (Maximum-prefix:4294967295)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.5	4	65536	8	9	1	0	0	00:05:20	1

```

D# show ip bgp
BGP table version is 1, local router ID is 10.1.1.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
                S Stale, b - backup entry
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop         Metric   LocPrf   Weight Path
* i192.168.1.0    10.1.1.1         0       100      0   i

Total number of prefixes 1

```

7.4.4 Configuring an AS Alliance

Configuration

Effect

- Configure a BGP alliance to reduce the number of BGP neighbor connections.

Notes

- It is advised to use private AS numbers for sub ASs (also called member ASs) within an alliance. Private AS numbers range from 64,512 to 65,535.
- Within a sub AS of an alliance, full mesh must be established for all BGP speakers (route reflectors can be further configured within the sub AS).
- An EBGP neighborhood must be established between sub ASs of an alliance.
- All BGP speakers within an alliance must belong to a sub AS within the alliance.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Configuring a BGP Alliance ID

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring a BGP Alliance Member

- (Mandatory) Perform this configuration in the BGP configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring Multiple Hops for an EBGP Neighbor

- Perform this configuration in the BGP configuration mode. It is mandatory when an EBGP neighbor is not directly connected.

Configuring BGP Route Re-distribution to a Network

- (Optional) Perform this configuration in the BGP configuration mode. Perform this configuration when a local route needs to be advertised. You can also configure an alternative network by means of re-distribution.

Verification

- Run the **show** command to display the BGP neighbor status.
- Run the **show** command to display the BGP routing table information.

Related

Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Enabling a BGP Alliance ID

Command	bgp confederation identifier <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	BGP configuration mode
Usage Guide	-

Configuring a BGP Alliance Member

Command	bgp confederation peers <i>as-number</i> [... <i>as-number</i>]
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	BGP configuration mode
Usage Guide	All member ASs of a local EBGp alliance must be identified.

Creating a BGP Neighbor

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>peer-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuring Multiple Hops for an EBGp Neighbor

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } ebgp-multihop [<i>ttl</i>]
Parameter Description	<i>peer-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>ttl</i> : Indicates the maximum number of hops that are allowed, ranging from 1 to 255.

Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuring BGP Route Re-distribution to a Network

Command	network <i>network-number</i> [mask <i>mask</i>] [route-map <i>map-tag</i>] [backdoor]
Parameter Description	<p><i>network-number</i>: Indicates the network address.</p> <p><i>mask</i>: Indicates the subnet mask.</p> <p><i>map-tag</i>: Indicates the name of a route map, consisting of no more than 32 characters.</p> <p>backdoor: Indicates that the route is a backdoor route.</p>
Command Mode	BGP configuration mode
Usage Guide	The core routing table must contain same IGP (or static and directly connected) routes.

Configuration Example

Configuring a BGP Alliance

Scenario Figure 7-7	
Configuration Steps	<ul style="list-style-type: none"> Configure BGP on A and B, set the AS number to 65,536 and configure an IBGP neighborhood.

- Configure BGP on C and D, set the AS number to 65,537 and configure an IBGP neighborship.
- Configure BGP on ASBR1 and E, set the AS number to 65,538 and configure an IBGP neighborship.
- Configure an alliance ID 100 on A, B, C, D, E and ASBR1.
- Configure the alliance member 65,537 on A, configure C as an EBGP neighbor, and set the peer AS number to 65,537.
- Configure the alliance members 65,536 and 65,538 on C, configure A as an EBGP neighbor and set the peer AS number to 65,536, configure ASBR1 as an EBGP neighbor and set the peer AS number to 65,538.
- Configure the alliance members 65,537 on ASBR1, configure C as an EBGP neighbor and set the peer AS number to 65,537, configure ASBR2 as an EBGP neighbor and set the peer AS number to 200.
- Configure BGP on ASBR2 and set the AS number to 200; configure ASBR1 as an EBGP neighbor and set the peer AS number to 100.

A

```
A# configure terminal
A(config)# interface loopback 0
A(config-if-Loopback 0)# ip address 10.1.1.1 255.255.255.255
A(config-if-Loopback 0)# exit
A(config)# interface GigabitEthernet 0/1
A(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0
A(config-if-GigabitEthernet 0/1)# exit
A(config)# interface GigabitEthernet 0/2
A(config-if-GigabitEthernet 0/2)# ip address 192.168.2.1 255.255.255.0
A(config-if-GigabitEthernet 0/2)# exit
A(config)# router bgp 65536
A(config-router)# bgp confederation identifier 100
A(config-router)# bgp confederation peers 65537
A(config-router)# neighbor 10.1.1.2 remote-as 65536
A(config-router)# neighbor 10.1.1.2 update-source loopback 0
A(config-router)# neighbor 10.1.1.3 remote-as 65537
A(config-router)# neighbor 10.1.1.3 ebgp-multihop 2
A(config-router)# neighbor 10.1.1.3 update-source loopback 0
A(config-router)# network 192.168.1.0 mask 255.255.255.0
```

B

```
B# configure terminal
```

	<pre> B(config)# interface loopback 0 B(config-if-Loopback 0)# ip address 10.1.1.2 255.255.255.255 B(config-if-Loopback 0)# exit B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip address 192.168.1.2 255.255.255.0 B(config-if-GigabitEthernet 0/1)# exit B(config)# router bgp 65536 B(config-router)# neighbor 10.1.1.1 remote-as 65536 B(config-router)# neighbor 10.1.1.1 update-source loopback 0 </pre>
C	<pre> C# configure terminal C(config)# interface loopback 0 C(config-if-Loopback 0)# ip address 10.1.1.3 255.255.255.255 C(config-if-Loopback 0)# exit C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.3.3 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# interface GigabitEthernet 0/2 C(config-if-GigabitEthernet 0/2)# ip address 192.168.2.3 255.255.255.0 C(config-if-GigabitEthernet 0/2)# exit C(config)# interface GigabitEthernet 0/3 C(config-if-GigabitEthernet 0/3)# ip address 192.168.4.3 255.255.255.0 C(config-if-GigabitEthernet 0/3)# exit C(config)# router bgp 65537 C(config-router)# bgp confederation identifier 100 C(config-router)# bgp confederation peers 65536 65538 C(config-router)# neighbor 10.1.1.1 remote-as 65536 C(config-router)# neighbor 10.1.1.1 update-source loopback 0 C(config-router)# neighbor 10.1.1.1 ebgp-multihop 2 C(config-router)# neighbor 10.1.1.4 remote-as 65537 C(config-router)# neighbor 10.1.1.4 update-source loopback 0 C(config-router)# neighbor 10.1.1.5 remote-as 65538 C(config-router)# neighbor 10.1.1.5 update-source loopback 0 </pre>

	<pre>C(config-router)# neighbor 10.1.1.5 ebgp-multihop 2</pre>
D	<pre>D# configure terminal D(config)# interface loopback 0 D(config-if-Loopback 0)# ip address 10.1.1.4 255.255.255.255 D(config-if-Loopback 0)# exit D(config)# interface GigabitEthernet 0/1 D(config-if-GigabitEthernet 0/1)# ip address 192.168.3.4 255.255.255.0 D(config-if-GigabitEthernet 0/1)# exit D(config)# router bgp 65537 D(config-router)# neighbor 10.1.1.3 remote-as 65537 D(config-router)# neighbor 10.1.1.3 update-source loopback 0</pre>
E	<pre>E# configure terminal E(config)# interface loopback 0 E(config-if-Loopback 0)# ip address 10.1.1.6 255.255.255.255 E(config-if-Loopback 0)# exit E(config)# interface GigabitEthernet 0/1 E(config-if-GigabitEthernet 0/1)# ip address 192.168.5.6 255.255.255.0 E(config-if-GigabitEthernet 0/1)# exit E(config)# router bgp 65538 E(config-router)# neighbor 10.1.1.5 remote-as 65538 E(config-router)# neighbor 10.1.1.5 update-source loopback 0</pre>
ASBR1	<pre>ASBR1# configure terminal ASBR1(config)# interface loopback 0 ASBR1(config-if-Loopback 0)# ip address 10.1.1.5 255.255.255.255 ASBR1(config-if-Loopback 0)# exit ASBR1(config)# interface GigabitEthernet 0/1 ASBR1(config-if-GigabitEthernet 0/1)# ip address 192.168.5.5 255.255.255.0 ASBR1(config-if-GigabitEthernet 0/1)# exit ASBR1(config)# interface GigabitEthernet 0/2 ASBR1(config-if-GigabitEthernet 0/2)# ip address 192.168.6.5 255.255.255.0</pre>

	<pre> ASBR1(config-if-GigabitEthernet 0/2)# exit ASBR1(config)# interface GigabitEthernet 0/3 ASBR1(config-if-GigabitEthernet 0/3)# ip address 192.168.4.5 255.255.255.0 ASBR1(config-if-GigabitEthernet 0/3)# exit ASBR1(config)# router bgp 65538 ASBR1(config-router)# bgp confederation identifier 100 ASBR1(config-router)# bgp confederation peers 65537 ASBR1(config-router)# neighbor 10.1.1.3 remote-as 65537 ASBR1(config-router)# neighbor 10.1.1.3 update-source loopback 0 ASBR1(config-router)# neighbor 10.1.1.3 ebgp-multihop 2 ASBR1(config-router)# neighbor 10.1.1.6 remote-65538 ASBR1(config-router)# neighbor 10.1.1.6 update-source loopback 0 ASBR1(config-router)# neighbor 192.168.6.7 remote-as 200 </pre>
ASBR2	<pre> ASBR2# configure terminal ASBR2(config)# interface GigabitEthernet 0/1 ASBR2(config-if-GigabitEthernet 0/1)# ip address 192.168.6.7 255.255.255.0 ASBR2(config-if-GigabitEthernet 0/1)# exit ASBR2(config)# router bgp 200 ASBR2(config-router)# neighbor 192.168.6.5 remote-as 100 ASBR2(config-router)# network 192.168.6.0 mask 255.255.255.0 </pre>
Verification	Run the show command to display the information.
A	<pre> A# show ip bgp summary BGP router identifier 10.1.1.1, local AS number 65536 BGP table version is 1 1 BGP AS-PATH entries 0 BGP Community entries 1 BGP Prefix entries (Maximum-prefix:4294967295) Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 10.1.1.2 4 65536 3 3 1 0 0 00:00:05 0 10.1.1.3 4 65537 3 3 1 0 0 00:00:06 1 </pre>

Total number of neighbors 1

A# show ip bgp

BGP table version is 1, local router ID is 10.1.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S Stale, b - backup entry

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight Path
* 192.168.6.0	192.168.6.7	0	100	0 (65537 65538) 200 i

Total number of prefixes 1

ASBR1

A# show ip bgp summary

BGP router identifier 10.1.1.5, local AS number 200

BGP table version is 2

2 BGP AS-PATH entries

0 BGP Community entries

2 BGP Prefix entries (Maximum-prefix:4294967295)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.3	4	65537	3	3	2	0	0	00:00:10	1
10.1.1.6	4	65538	3	3	2	0	0	00:00:08	0
192.168.6.7	4	200	3	3	2	0	0	00:00:05	1

Total number of neighbors 1

A# show ip bgp

BGP table version is 1, local router ID is 10.1.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S Stale, b - backup entry

Origin codes: i - IGP, e - EGP, ? - incomplete

	<pre> Network Next Hop Metric LocPrf Weight Path * 192.168.1.0 10.1.1.1 0 100 0 (65537 65536) i *> 192.168.6.0 192.168.6.7 0 100 0 200 i Total number of prefixes 1 </pre>
ASBR2	<pre> A# show ip bgp summary BGP router identifier 192.168.6.7, local AS number 200 BGP table version is 1 1 BGP AS-PATH entries 0 BGP Community entries 1 BGP Prefix entries (Maximum-prefix:4294967295) Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 192.168.6.5 4 100 3 3 1 0 0 00:00:05 1 Total number of neighbors 1 A# show ip bgp BGP table version is 1, local router ID is 10.1.1.1 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S Stale, b - backup entry Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 192.168.1.0 192.168.6.5 0 100 0 (65537 65538) 200 i Total number of prefixes 1 </pre>

Common Errors

- No BGP alliance neighbor is configured.
- Full mesh is not established within sub-ASs of an alliance.

7.4.5 Configuring Multi-path Load Balancing of BGP

Configuration

Effect

- Implement multi-path load balancing for IBGP routes.
- Support AS-PATH loose comparison.

Notes

- Routes learned from an IBGP neighbor must have the same priority (the router-ID does not need to be compared).

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring BGP Load Balancing

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring AS-PATH Loose Comparison

- (Optional) Perform this configuration in the BGP configuration mode. Perform this configuration when load balancing needs to be implemented for routes learned from different ASs.

Verification

- Run the **show** command to display BGP routing information.
- Run the **show** command to display the core routing table information.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.

Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	<code>neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i></code>
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuring BGP Load Balancing

Command	<code>maximum-paths { <i>ebgp</i> <i>ibgp</i> } <i>number</i></code>
Parameter Description	<i>number</i> : Indicates the maximum number of equivalent paths, ranging from 1 to device capacity. If the value is 1, multi-path load balancing of IBGP will be disabled.
Command Mode	BGP configuration mode
Usage Guide	-

Configuring AS-PATH Loose Comparison

Command	<code>bgp bestpath as-path multipath-relax</code>
----------------	--

Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Configuration Example

Configuring Multi-path Load Balancing of IBGP

<p>Scenario Figure 7-8</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Enable BGP on all devices and set the AS numbers as shown in Figure 7-8. ▪ Establish IBGP neighborships between A and B and between A and C by using directly connected interfaces. ▪ Establish EBGP neighborships between B and D and between C and E by using directly connected interfaces. ▪ Re-distribute the same routes to D and E. ▪ Configure IBGP load balancing on A and enable the AS-PATH loose comparison mode.
<p>A</p>	<p>A# conf terminal</p>

	<pre> A(config)# interface fastEthernet 0/0 A(config-if-FastEthernet 0/0)# ip address 10.1.1.1 255.255.0.0 A(config-if-FastEthernet 0/0)# exit A(config)# interface fastEthernet 0/1 A(config-if-FastEthernet 0/1)# ip address 10.2.1.1 255.255.0.0 A(config-if-FastEthernet 0/1)# exit A(config)# ip route 10.3.0.0 255.255.0.0 10.1.1.2 A(config)# ip route 10.4.0.0 255.255.0.0 10.2.1.2 A(config)# router bgp 65530 A(config-router)# neighbor 10.1.1.2 remote-as 65530 A(config-router)# neighbor 10.2.1.2 remote-as 65530 A(config-router)# bgp maximum-paths ibgp 2 A(config-router)# bgp bestpath as-path multipath-relax </pre>
B	<pre> B# conf terminal B(config)# interface fastEthernet 0/0 B(config-if-FastEthernet 0/0)# ip address 10.1.1.2 255.255.0.0 B(config-if-FastEthernet 0/0)# exit B(config)# interface fastEthernet 0/1 B(config-if-FastEthernet 0/1)# ip address 10.3.1.2 255.255.0.0 B(config-if-FastEthernet 0/1)# exit B(config)# router bgp 65530 B(config-router)# neighbor 10.1.1.1 remote-as 65530 B(config-router)# neighbor 10.3.1.1 remote-as 65531 </pre>
C	<pre> C# conf terminal C(config)# interface fastEthernet 0/0 C(config-if-FastEthernet 0/0)# ip address 10.2.1.2 255.255.0.0 C(config-if-FastEthernet 0/0)# exit C(config)# interface fastEthernet 0/1 C(config-if-FastEthernet 0/1)# ip address 10.4.1.2 255.255.0.0 C(config-if-FastEthernet 0/1)# exit C(config)# router bgp 65530 </pre>

	<pre>C(config-router)# neighbor 10.2.1.1 remote-as 65530 C(config-router)# neighbor 10.4.1.1 remote-as 65532</pre>
D	<pre>D# conf terminal D(config)# interface fastEthernet 0/0 D(config-if-FastEthernet 0/0)# ip address 10.3.1.1 255.255.0.0 D(config-if-FastEthernet 0/0)# exit D(config)# interface loopback 1 D(config-if)#ip address 10.5.1.1 255.255.0.0 D(config-if-FastEthernet 0/1)# exit D(config)# router bgp 65531 D(config-router)# neighbor 10.3.1.2 remote-as 65530 D(config-router)# redistribute connected</pre>
E	<pre>E# conf terminal E(config)# interface fastEthernet 0/0 E(config-if-FastEthernet 0/0)# ip address 10.4.1.1 255.255.0.0 E(config-if-FastEthernet 0/0)# exit E(config)# interface loopback 1 E(config-if)#ip address 10.5.1.2 255.255.0.0 E(config-if-FastEthernet 0/1)# exit E(config)# router bgp 65532 E(config-router)# neighbor 10.4.1.2 remote-as 65530 E(config-router)# redistribute connected</pre>
Verification	Run the show command to display the information.
A	<pre>A# show ip bgp summary BGP router identifier 10.2.1.1, local AS number 65530 BGP table version is 9 2 BGP AS-PATH entries 0 BGP Community entries 3 BGP Prefix entries (Maximum-prefix:4294967295)</pre>

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down
State/PfxRcd
```

```
172.16.23.140 4 65530 29 25 8 0 0 00:18:48 2
```

```
172.16.23.141 4 65530 24 21 8 0 0 00:17:58 2
```

```
A# show ip bgp
```

```
BGP table version is 9, local router ID is 10.2.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i10.3.0.0/16	10.3.1.1	0	100	0	65531 ?
*>i10.4.0.0/16	10.4.1.1	0	100	0	65532 ?
* i10.5.0.0/16	10.3.1.1	0	100	0	65531 ?
*>i	10.4.1.1	0	100	0	65532 ?

```
Total number of prefixes 3
```

```
A# show ip bgp 10.5.0.0
```

```
BGP routing table entry for 10.5.0.0/16
```

```
Paths: (2 available, best #1, table Default-IP-Routing-Table)
```

```
Not advertised to any peer
```

```
65532
```

```
10.4.1.1 from 10.2.1.2 (172.16.24.1)
```

```
Origin incomplete, metric 0, localpref 100, valid, internal, multipath, best
```

```
Last update: Mon Mar 21 03:45:14 2011
```

```
65531
```

```
10.3.1.1 from 10.1.1.2 (172.16.25.1)
```

```
Origin incomplete, metric 0, localpref 100, valid, internal, multipath
```

```
Last update: Mon Mar 21 03:45:14 2011
```

```
A# show ip route
```

Codes: C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default

Gateway of last resort is no set

C 10.1.0.0/16 is directly connected, FastEthernet 0/0

C 10.1.1.1/32 is local host.

C 10.2.0.0/16 is directly connected, FastEthernet 0/1

C 10.2.1.1/32 is local host.

S 10.3.0.0/16 [1/0] via 10.1.1.2

S 10.4.0.0/16 [1/0] via 10.2.1.2

B 10.5.0.0/16 [200/0] via 10.3.1.1, 00:27:56

[200/0] via 10.4.1.1, 00:27:56

Common Errors

- The priorities of multi-hop BGP routes are different, which causes load balancing failure.

7.4.6 Configuring EBGP FRR

Configuration Effect

- Implement EBGP FRR.

Notes

- (Optional) Configure a neighbor BFD session to implement fast link fault detection.

Configuration Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring BGP FRR

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring a Neighbor BFD Session

- (Optional) Perform this configuration in the BGP configuration mode.

Verification

- Run the **show** command to display routing information.

Related

Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Configuring BGP FRR

Command	bgp fast-reroute
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Creating a BFD Session to a BGP Neighbor


Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } fall-over bfd
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters.
Command Mode	BGP configuration mode
Usage Guide	-

Configuration

Example

Configuring EBGP FRR

Scenario
Figure 7-9



**Configurati
on Steps**

- Enable BGP on all devices.
- Configure the addresses of the directly connected interfaces on A, B and C to establish EBGP neighborships.
- Configure a BFD session for the EBGP neighborship between B and C.
- Configure FRR on C.

A

```
A# conf terminal
A(config)# interface GigabitEthernet 0/1
A(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0
A(config-if-GigabitEthernet 0/1)# exit
A(config)# interface GigabitEthernet 0/2
A(config-if-GigabitEthernet 0/2)# ip address 192.168.2.1 255.255.255.0
A(config-if-GigabitEthernet 0/2)# exit
A(config)# router bgp 100
A(config-router)# neighbor 192.168.1.2 remote-as 300
A(config-router)# neighbor 192.168.2.2 remote-as 200
A(config-router)# redistribute connect
```

B

```
B# configure terminal
B(config)# interface GigabitEthernet 0/1
B(config-if-GigabitEthernet 0/1)# ip address 192.168.3.1 255.255.255.0
B(config-if-GigabitEthernet 0/1)# bfd interval 200 min_rx 200 multiplier 5
B(config-if-GigabitEthernet 0/1)# exit
B(config)# interface GigabitEthernet 0/2
B(config-if-GigabitEthernet 0/2)# ip address 192.168.2.2 255.255.255.0
B(config-if-GigabitEthernet 0/2)# exit
```

	<pre> B(config)# router bgp 200 B(config-router)# neighbor 192.168.3.2 remote-as 300 B(config-router)# neighbor 192.168.3.2 fall-over bfd B(config-router)# neighbor 192.168.2.1 remote-as 100 B(config-router)# redistribute connect </pre>
C	<pre> C# configure terminal C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.1.2 255.255.255.0 C(config-if-GigabitEthernet 0/1)# exit C(config)# interface fastEthernet 0/2 C(config-if-GigabitEthernet 0/2)# ip address 192.168.3.2 255.255.0.0 C(config-if-GigabitEthernet 0/2)# bfd interval 200 min_rx 200 multiplier 5 C(config-if-GigabitEthernet 0/2)# exit C(config)# router bgp 300 C(config-router)# neighbor 192.168.1.1 remote-as 100 C(config-router)# neighbor 192.168.3.1 remote-as 200 C(config-router)# neighbor 192.168.3.1 fall-over bfd C(config-router)# address-family ipv4 unicast C(config-router-af)# bgp fast-reroute C(config-router-af)# redistribute connect </pre>
Verification	Run the show command to display the information.
C	<pre> C# show ip bgp summary BGP router identifier 10.10.10.10, local AS number 300 BGP table version is 12 4 BGP AS-PATH entries 0 BGP Community entries 3 BGP Prefix entries (Maximum-prefix:4294967295) Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 192.168.1.1 4 100 76 77 12 12 0 00:59:27 3 192.168.3.1 4 200 30 30 12 12 0 00:19:03 3 </pre>

Total number of neighbors 2

C# show ip bgp

BGP table version is 12, local router ID is 10.10.10.10

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S Stale, b - backup entry

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
* 192.168.1.0	192.168.3.1	0		0 200	?
*	192.168.1.1	0	0	100	?
*>	0.0.0.0	0	32768		?
*> 192.168.2.0	192.168.3.1	0		0 200	?
*b	192.168.1.1	0	0	100	?
* 192.168.3.0	192.168.3.1	0		0 200	?
*	192.168.1.1	0	0	100 200	?
*>	0.0.0.0	0	32768		?

Total number of prefixes 3

C# show ip bgp 192.168.2.0

BGP routing table entry for 192.168.2.0/24

Paths: (2 available, best #1, table Default-IP-Routing-Table)

Advertised to non peer-group peers:

192.168.1.1

200

192.168.3.1 from 192.168.3.1 (3.3.3.3)

Origin incomplete, metric 0, localpref 100, valid, external, best

Last update: Tue Oct 5 00:26:52 1971

100

192.168.1.1 from 192.168.1.1 (44.44.44.44)

Origin incomplete, metric 0, localpref 100, valid, external, backup

Last update: Mon Oct 4 23:46:28 1971

C# show ip route

Codes: C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default

Gateway of last resort is no set

C 192.168.1.0/24 is directly connected, GigabitEthernet 1/9
C 192.168.1.2/32 is local host.
B 192.168.2.0/24 [20/0] via 192.168.3.1, 00:21:39
C 192.168.3.0/24 is directly connected, GigabitEthernet 1/11
C 192.168.3.2/32 is local host.

Common Errors

- No BFD session is configured for BGP neighbors.

7.4.7 Configuring Fast Withdrawal of Specified BGP Routes

Configuration Effect

- Implement fast withdrawal of specified BGP routes.

Notes

- Configure fast withdrawal of specified routes to preferentially advertise the withdrawal of specified routes.

Configuration Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Configuring Fast Withdrawal of Specified BGP Routes

- (Mandatory) Perform this configuration in BGP configuration mode.

Verification

- Run the **show** command to display the configurations.

Related Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number. The value ranges from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuring Fast Withdrawal of Specified BGP Routes

Command	bgp fast-withdraw { access-list { <i>access-list-number</i> <i>access-list-name</i> } prefix-list <i>prefix-list-name</i> route-map <i>map-tag</i> }
Parameter Description	<i>access-list-number</i> : Indicates the ACL number. The value ranges from 1 to 199 and 1300 to 2699. <i>access-list-name</i> : Indicates an ACL name. <i>prefix-list-name</i> : Indicates the name of a prefix list. <i>map-tag</i> : Indicates the name of a route map.
Command Mode	BGP configuration mode
Usage Guide	N/A

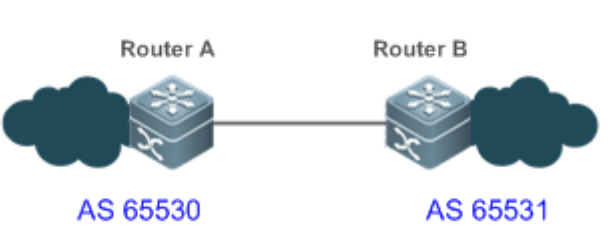
Creating a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
---------	---

Parameter Description	<p><i>neighbor-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, containing no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group). The value ranges from 1 to 4,294,967,295.</p>
Command Mode	BGP configuration mode
Usage Guide	The AS number specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuration Example

Configuring Fast Withdrawal of Specified BGP Routes

Scenario Figure 7-10	
Configuration Steps	<ul style="list-style-type: none"> Enable BGP on all devices. Establish a BGP neighbor relationship between Routers A and B. Configure ACLs on Router A. Configure fast withdrawal of specified routes on Router A.
A	A# conf terminal
Verification	Run the show command to display the configurations.
A	A#show run router bgp

7.4.8 Configuring Local ASs

Configuration

Effect

- Smoothly migrate the network configurations of router A from AS 23 to AS 3600.

Notes

N/A

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring the Local AS for a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Verification

- Run the **show** command to display the information.

Related

Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

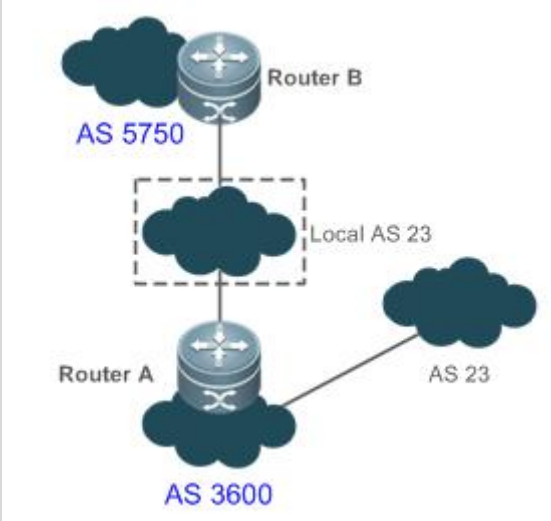
Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuring the Local AS for a BGP Neighbor

Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } local-as <i>as-number</i> [no-prepend [replace-as [dual-as]]]
Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>as-number</i>: Indicates a local AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.</p> <p>no-prepend: Does not add the local AS to the AS-PATH in the routing information received by a peer. This option is not available by default.</p> <p>replace-as: For the AS-PATH in the routing information sent by a peer, the local AS is used to replace the BGP AS. This option is not available by default.</p> <p>dual-as: Enables a peer to use the BGP AS or Local AS to establish a BGP connection with a device. This option is not available by default.</p>
Command Mode	BGP configuration mode
Usage Guide	-

Configuration Example

Configuring BGP Local-AS

<p>Scenario Figure 7-11</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Create an EBGP neighborhood with B on A and specify the Local-AS for the EBGP neighborhood. ▪ Create an EBGP neighborhood for connecting to A on B.
<p>A</p>	<pre>A# configure terminal A(config)# router bgp 3600 A(config-router)# neighbor 57.50.1.1 remote-as 5750 A(config-router)# neighbor 57.50.1.1 update-source loopback 0 A(config-router)# neighbor 57.50.1.1 ebgp-multihop 255 A(config-router)# neighbor 57.50.1.1 local-as 23 no-prepend replace-as dual-as</pre>
<p>B</p>	<pre>B# configure terminal B(config)# router bgp 5750 B(config-router)# neighbor 36.0.1.1 remote-as 23 B(config-router)# neighbor 36.0.1.1 update-source loopback 0 B(config-router)# neighbor 36.0.1.1 ebgp-multihop 255</pre>
<p>Verification</p>	<p>Run the show command to display the BGP neighbor status.</p>

A

```
A# show ip bgp neighbors 57.50.1.1
BGP neighbor is 57.50.1.1, remote AS 5750, local AS 23(using Peer's Local AS, no-
prepend, replace-as, dual-as), external link
BGP version 4, remote router ID 0.0.0.0
BGP state = Idle
Last read, hold time is 180, keepalive interval is 60 seconds
Received 0 messages, 0 notifications, 0 in queue
  open message:0 update message:0 keepalive message:0
  refresh message:0 dynamic cap:0 notifications:0
Sent 0 messages, 0 notifications, 0 in queue
```

7.4.9 Configuring BGP GR

Configuration

Effect

- Configure BGP GR to implement network deployment with high reliability.

Notes

- To successfully deploy the BGP GR function, you need to use a neighbor device as the GR Helper.
- In an BGP environment, you also need to configure IGP GR.
- After BGP GR is enabled, you need to reset a BGP neighbor connection to make it take effect.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Configuring BGP GR

- Perform this configuration in the BGP configuration mode, which is configured by default.

Configuring a BGP GR Timer

- (Optional) Perform this configuration in the BGP configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Verification

- Run the **show** command to display the neighbor status.

Related Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Configuring BGP GR

Command	<code>bgp graceful-restart</code>
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Configuring the BGP GR Restart Timer

Command	<code>bgp graceful-restart restart-time <i>restart-time</i></code>
Parameter Description	<i>restart-time</i> : Indicates the maximum waiting time that the GR Restarter hopes the GR Helper to wait before a new connection is created, ranging from 1 to 3600 seconds.
Command Mode	BGP configuration mode
Usage Guide	-

Configuring the BGP GR Route Stale Timer

Command	bgp graceful-restart stalepath-time <i>time</i>
Parameter Description	<i>time</i> : Indicates the maximum time that a stale route keeps valid after the connection with a neighbor GR device is recovered, ranging from 1 to 3600 seconds.
Command Mode	BGP configuration mode
Usage Guide	-

Creating a BGP Neighbor


Command	neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>peer-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuration

Example

Configuring BGP GR

Scenario
Figure 7-12



Configuration Steps

- Enable BGP on all devices and set the AS numbers as shown in Figure 7-12.
- Configure a loopback interface on R2, R3, and R4 and create an IBGP neighborhood based on the loopback interface.
- Create an EBGP neighborhood by using the directly connected interfaces on R1 and R2.
- Enable BGP GR on R1, R2, R3, and R4.

R1

```
R1# configure terminal
R1(config-router)# exit
R1(config)# router bgp 100
R1(config-router)# bgp graceful-restart
```

R2

```
R2# configure terminal
R2(config)# router ospf 1
R2(config-router)# graceful-restart
R2(config-router)# exit
R2(config)# router bgp 100
R2(config-router)# bgp graceful-restart
```

R3

```
R3# configure terminal
R3(config)# router ospf 1
R3(config-router)# graceful-restart
```

	<pre>R3(config-router)# exit R3(config)# router bgp 100 R3(config-router)# bgp graceful-restart</pre>
R4	<pre>R4# configure terminal R4(config)# router ospf 1 R4(config-router)# graceful-restart R4(config-router)# exit R4(config)# router bgp 100 R4(config-router)# bgp graceful-restart</pre>
Verification	Run the show command to display the BGP neighbor status.
R2	<pre>R2# show ip ospf Routing Process "ospf 1" with ID 10.0.0.2 Process uptime is 4 minutes Process bound to VRF default Conforms to RFC2328, and RFC1583Compatibility flag is enabled Supports only single TOS(TOS0) routes Supports opaque LSA This router is an ASBR (injecting external routing information) SPF schedule delay 5 secs, Hold time between two SPFs 10 secs LsaGroupPacing: 240 secs Number of incoming current DD exchange neighbors 0/5 Number of outgoing current DD exchange neighbors 0/5 Number of external LSA 4. Checksum 0x0278E0 Number of opaque AS LSA 0. Checksum 0x000000 Number of non-default external LSA 4 External LSA database is unlimited. Number of LSA originated 6 Number of LSA received 2 Log Neighbor Adjacency Changes : Enabled Graceful-restart enabled Graceful-restart helper support enabled</pre>


```
Number of areas attached to this router: 1
Area 0 (BACKBONE)
.....

R2# show ip bgp neighbors
BGP neighbor is 192.168.195.183, remote AS 200, local AS 100, external link
  BGP version 4, remote router ID 10.0.0.1
  BGP state = Established, up for 00:06:37
  Last read 00:06:37, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
  Address family IPv4 Unicast: advertised and received
  Graceful restart: advertised and received
  Remote Restart timer is 120 seconds
  Address families preserved by peer:
    None
.....
```

Common Errors

- GR is not enabled for IGP.
- GR is not enabled for a BGP neighbor device.

7.4.10 Configuring a BGP IPv6 Address Family

Configuration Effect

- Configure BGP IPv6 routes to implement IPv6 network access in different ASs.

Notes

- Generally, BGP uses IPv6 addresses to create neighborships and implement exchange of IPv6 routes.
- Configurations related to BGP IPv6 services must be configured in the BGP IPv6 address family mode.

Configuration Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring the BGP IPv4 Address Family Mode

- (Optional) Perform this configuration in the BGP configuration mode.

Disabling the IPv4 Address Family Capability for a BGP Neighbor

- (Optional) Perform this configuration in the BGP IPv6 configuration mode.

Configuring the BGP IPv6 Address Family Mode

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring the IPv6 Address Family Capability for a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP IPv6 configuration mode.

Configuring IPv6 Route Advertisement in BGP

- (Optional) Perform this configuration in the BGP IPv6 configuration mode.

Verification

- Run the **show** command to display the neighbor status.
- Run the **show** command to display the routing status.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>neighbor-address</i> : Indicates the address of a peer, which is usually an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Configuring the BGP IPv4 Address Family Mode

Command	address-family ipv4 unicast
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Disabling the IPv4 Address Family Capability for a BGP Neighbor

Command	no neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } activate
Parameter Description	<i>neighbor-address</i> : Indicates the address of a peer, which is usually an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters.
Command Mode	BGP IPv4 address family mode

Usage Guide	Neighbors with IPv6 addresses are used to exchange IPv6 routes. However, when a neighbor is configured in the BGP mode, BGP automatically activates the IPv4 unicast address family capability for the neighbor. Therefore, you are advised to manually disable the IPv4 unicast address family capability.
-------------	---

Configuring the BGP IPv6 Address Family Mode

Command	address-family ipv6 unicast
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Configuring the IPv6 Address Family Capability for a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } activate
Parameter Description	<i>neighbor-address</i> : Indicates the address of a peer, which is usually an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters.
Command Mode	BGP IPv6 address family mode
Usage Guide	-

Configuring IPv6 Route Advertisement in BGP

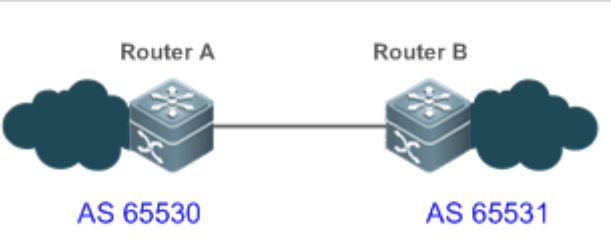
Command	network <i>network-number</i> [mask <i>mask</i>] [route-map <i>map-tag</i>] [backdoor]
---------	---

Parameter Description	<p><i>network-number</i>: Indicates the network number.</p> <p><i>mask</i>: Indicates the subnet mask.</p> <p><i>map-tag</i>: Indicates the name of a route map, consisting of no more than 32 characters.</p> <p>backdoor: Indicates that the route is a backdoor route.</p>
Command Mode	BGP IPv6 address family mode
Usage Guide	-

Configuration

Example

Configuring BGP to Implement IPv6 Route Exchange in Different ASs

<p>Scenario</p> <p>Figure 7-13</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Enable BGP on all devices and set the AS numbers as shown in Figure 7-13. Configure a BGP neighbor, disable the IPv4 address family capability for the neighbor and activate the IPv6 address family capability. Configure IPv6 route advertisement in BGP.
<p>A</p>	<pre> A# configure terminal A(config)# int loopback 0 A(config-if-Loopback)# ipv6 address 30::1/128 A(config-if-Loopback)# exit A(config)# router bgp 65530 A(config-router)# neighbor 100::1 remote-as 65531 A(config-router)# address-family ipv4 A(config-router-af)# no neighbor 100::1 activate A(config-router-af)# exit-address-family A(config-router)# address-family ipv6 </pre>

	<pre>A(config-router-af)# neighbor 100::1 activate A(config-router-af)# network 30::1/128</pre>
B	<pre>B# configure terminal B(config)# router bgp 65531 B(config-router)# neighbor 100::2 remote-as 65530 B(config-router)# address-family ipv4 B(config-router-af)# no neighbor 100::2 activate B(config-router-af)# exit-address-family B(config-router)# address-family ipv6 B(config-router-af)# neighbor 100::2 activate</pre>
Verification	Run the show command to display the BGP neighbor status.
A	<pre>A# show bgp ipv6 unicast summary BGP router identifier 1.1.1.1, local AS number 65530 BGP table version is 1 1 BGP AS-PATH entries 0 BGP Community entries 1 BGP Prefix entries (Maximum-prefix:4294967295) Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 100::1 4 65531 4 6 1 0 0 00:01:49 0 Total number of neighbors 1</pre>
B	<p>Run the show command to display BGP routing information.</p> <pre>B# show bgp ipv6 unicast BGP table version is 4, local router ID is 2.2.2.2 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S Stale, b - backup entry Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path</pre>

```
*> 30::1/128    100::2        0            0 65530 i
```

```
Total number of prefixes 1
```

Common

Errors

- The IPv6 address family capability is not activated for BGP neighbors.
- In non-6PE scenarios, IPv4 addresses are used to establish IPv6 routes for exchange between neighbors.

7.4.11 Configuring a BGP EVPN

Notes

- By default, the BGP routing mode is located in the IPv4 unicast address family and a BGP EVPN must be configured in L2VPN EVPN address family mode.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Activating the EVPN Capability for a BGP Neighbor

- (Mandatory) Perform this configuration in BGP L2VPN EVPN address family mode.

Creating an EVI Instance

- (Mandatory) Perform this configuration in EVPN configuration mode.

Creating EVI Instances in Batch

- (Optional) Perform this configuration in EVPN configuration mode.

Creating RD

- (Mandatory) Perform this configuration in EVPN-VNI configuration mode.

Creating Route Target

- (Mandatory) Perform this configuration in EVPN-VNI configuration mode.

Configuring Route Map for EVI Instance

- (Optional) Perform this configuration in EVPN-VNI address family mode.

Configuring MAC Mobility Parameters

- (Optional) Perform this configuration in BGP L2VPN EVPN configuration mode.

Verification

- Run the **show** command to display the information.

Related

Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Entering the L2VPN EVPN Address Family

Command	address-family l2vpn evpn
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Activating the EVPN Capability for a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } activate
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address must be an IPv4 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters.
Command Mode	BGP L2VPN EVPN address family mode
Usage Guide	-

Creating an EVI Instance

Command	vni <i>vni-id</i>
Parameter Description	<i>vni-id</i> : indicates the VNI ID. Ranges from 1 to 16777215.
Command Mode	EVPN configuration mode
Usage Guide	After VNI configuration is completed, related RD/RT needs to be configured.

Creating EVI Instances in Batches

Command	vni range <i>vni-id-list</i>
Parameter Description	<i>vni-id-list</i> : Specifies the name of a VNI ID list.
Command Mode	EVPN configuration mode
Usage Guide	<p>The configuration fails if the number of configured VNI instances exceeds the capacity or a VNI instance is being deleted.</p> <p>The RD values of different VNI instances must be different. Therefore, only the automatic RD configuration mode is available for configuring VNI instances in batches, and VNI instances cannot be configured manually.</p>

Configuring an RD

Command	rd { auto <i>rd_value</i> }
Parameter Description	<p>auto: Generates an RD value automatically.</p> <p><i>rd_value</i>: Specifies the value of an RD.</p> <p>The parameter <i>rd_value</i> has three forms:</p> <p>(1) rd_value = as_num:nn</p> <p>as_num indicates the 2-byte public AS number. nn is user-defined and ranges from 0 to 4294967295.</p> <p>(2) rd_value = ip_addr:nn</p> <p>ip_add refers to the global IP address. nn is user-defined and ranges from 0 to 65535.</p> <p>(3) rd_value = as4_num:nn</p> <p>as4_num indicates the 4-byte public AS number. nn is user-defined and ranges from 1 to 65535.</p> <p>The 4-byte AS number ranges from 1 to 4294967295, which is 1 to 65535.65535 in dot mode.</p>
Command Mode	EVPN VNI configuration mode
Usage Guide	If an EVI is defined and has an RD value configured, the RD value cannot be modified. To modify it, you need to delete the EVI first, and then define the EVI again and configure a new RD value for it.

	<p>One EVI has only one RD value.</p> <ul style="list-style-type: none"> i The RD value of the 4-byte AS is in the format of AS4:NN. AS4 can be expressed in decimal or dot mode. AS4 ranges from 1 to 4294967295, which is 1 to 65535.65535 in dot mode. NN ranges from 1 to 65535. i The AS number in the range of 1 to 65535 is displayed the same in both decimal mode and dot mode. Therefore, save the AS number in the range of 1 to 65535 as a 2-byte AS number. i The automatically generated RD value is in the format of ip_addr:nn. ip_addr indicates the BGP router ID, and nn indicates the VNI ID. Because the VNI ID occupies a maximum of three bytes but there is only two bytes available for nn, when the VNI ID occupies more than 2 bytes, the VNI ID is split before being filled in the position for nn. Once generated, the RD value remains unchanged.
--	--

Configuring a Route Target

Command	<code>route-target { import export both } { auto rt_value }</code>
Parameter Description	<p>import: Specifies the import RT value.</p> <p>export: Specifies the export RT value.</p> <p>both: Specifies the import and export RT values.</p> <p>auto: Generates an RT value automatically.</p> <p><i>rt_value</i>: Specifies the value of an RT.</p> <p>The parameter <i>rt_value</i> has three forms:</p> <p>(1) <code>rt_value = as_num:nn</code></p> <p>as_num indicates the 2-byte public AS number. nn is user-defined and ranges from 0 to 4294967295.</p> <p>(2) <code>rt_value = ip_addr:nn</code></p> <p>ip_addr refers to the global IP address. nn is user-defined and ranges from 0 to 65535.</p> <p>(3) <code>rt_value = as4_num:nn</code></p> <p>as4_num indicates the 4-byte public AS number. nn is user-defined and ranges from 1 to 65535.</p> <p>The 4-byte AS number ranges from 1 to 4294967295, which is 1 to 65535.65535 in dot mode.</p>
Command Mode	EVPN VNI configuration mode

Usage Guide	<p>One EVI supports multiple import and export RT values.</p> <ul style="list-style-type: none"> i The automatically generated RT value is in the format of AS2:nn. AS2 indicates the 2-byte AS number. If a 4-byte AS number is used, it is split, and a two 2-byte part is filled in the position for nn while the other 2-byte part is filled in the position for the RT value. nn indicates the VNI ID. There are four bytes of space provided for accommodating the 3-byte VNI ID. i If the BGP AS number changes, the automatically-generated RT value will also be changed. i If the manually configured RT value is consistent with the automatically generated one, both of them are displayed. If the automatically generated RT value is 100:1, it is directly deleted. Deleting this value does not cancel the auto configuration mode. To cancel the auto configuration mode, run the no form of the route-target auto command.
-------------	---

Configuring a Route Map for EVI Instances

Command	export map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Specifies the name of a route map.
Command Mode	EVPN VNI configuration mode
Usage Guide	<p>This command is used to modify the extended community attribute advertised by a type 5 route converted from a local EVPN route or IP route.</p> <p>This command supports only one route map, and the old configuration is overwritten by the new configuration.</p>

Command	import map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Specifies the name of a route map.
Command Mode	EVPN VNI configuration mode

Usage Guide	<p>This command is used to filter the remote EVPN routes to be imported to the local VNI instance, or modify the attribute of the remote EVPN routes imported to the local VNI instance.</p> <p>This command supports only one route map, and the old configuration is overwritten by the new configuration.</p>
-------------	--

Configuring MAC Mobility Parameters


Command	bgp mac-mobility <i>timer count</i>
Parameter Description	<p><i>timer</i>: Specifies the timer duration (in seconds) for MAC mobility detection. If a number of <i>count</i> MAC address moves are detected within the duration, it is considered that a MAC conflict occurs and syslogs are displayed to warn users. The value is 180 by default and ranges from 1 to 3600.</p> <p><i>count</i>: Specifies the detection times. The value is 5 by default and ranges from 1 to 360.</p>
Command Mode	BGP L2VPN EVPN address family mode
Usage Guide	<p>One MAC address may move for many times. When a MAC address move is detected by a PE through local learning, a timer of <i>timer</i> seconds (default value = 180s) will be enabled. If a number of <i>count</i> (default value = 5) MAC address moves are detected before the timer times out, it is considered that a MAC conflict occurs.</p>

Configuration

Example

Configuring a BGP EVPN Address Family

**Scenario
Figure 7-14**



**Configurati
on Steps**

- Enable BGP on all devices and set the AS numbers as shown in Figure 1-14.
- Configure BGP neighborships as shown in Figure 1-14.
- Activate the L2VPN EVPN address family capability for BGP neighbors.

A

```
A# config terminal
A(config)# interface gigabitEthernet 0/1
A(config-GigabitEthernet 0/1)# no shutdown
A(config-GigabitEthernet 0/1)# ip address 172.18.2.1 255.255.255.0
A(config-GigabitEthernet 0/1)# exit
A(config)# interface gigabitEthernet 0/2
A(config-GigabitEthernet 0/1)# no shutdown
A(config-GigabitEthernet 0/1)# ip address 172.18.1.1 255.255.255.0
A(config-GigabitEthernet 0/1)# exit
A(config)# router bgp 100
A(config-router)# neighbor 172.18.2.3 remote-as 100
A(config-router)# neighbor 172.18.1.2 remote-as 100
A(config-router)# address-family l2vpn evpn
A(config-router-af)# neighbor 172.18.2.3 activate
A(config-router-af)# neighbor 172.18.1.2 activate
A(config-router-af)# neighbor 172.18.2.3 route-reflector-client
```

	<pre>A(config-router-af)# neighbor 172.18.1.2 route-reflector-client A(config-router-af)# exit-address-family</pre>
B	<pre>B# config terminal B(config)# evpn B(config-evpn)# vni 100 B(config-evpn-vni)# rd auto B(config-evpn-vni)# route-target both auto B(config-evpn-vni)# exit B(config-evpn)# exit B(config)# interface gigabitEthernet 0/1 B(config-GigabitEthernet 0/1)# no shutdown B(config-GigabitEthernet 0/1)# ip address 172.18.2.3 255.255.255.0 B(config-GigabitEthernet 0/1)# exit B(config)# router bgp 100 B(config-router)# neighbor 172.18.2.1 remote-as 100 B(config-router)# address-family l2vpn evpn B(config-router-af)# neighbor 172.18.2.1 activate B(config-router-af)# exit-address-family</pre>
C	The same as that for B.
Verification	Run the show command to display the information.
A	<pre>A#sh bgp l2vpn evpn all summary For address family: L2VPN EVPN BGP router identifier 3.3.3.3, local AS number 100 BGP table version is 51 1 BGP AS-PATH entries 0 BGP Community entries 6 BGP Prefix entries (Maximum-prefix:4294967295) Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 172.18.2.3 4 100 1993 2033 49 0 0 1d04h40m 3</pre>

```
172.18.1.2 4 100 619 623 51 0 0 08:47:26 3
```

Total number of neighbors 2

```
A#sh bgp l2vpn evpn all
```

BGP table version is 51, local router ID is 3.3.3.3

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S Stale, b - backup entry

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1.1.1.1:100					
*>i0:6:0011.2233.2016:0:0.0.0.0/128					
1.1.1.1		0	100	0	i
*>i0:6:0011.2233.2016:32:100.1.1.2/128					
1.1.1.1		0	100	0	i
*>i0:32:1.1.1.1/72					
1.1.1.1		0	100	0	i

Total number of prefixes 3

Route Distinguisher: 2.2.2.2:100

*>i0:6:08c6.b322.33df:0:0.0.0.0/128					
2.2.2.2		0	100	0	i
*>i0:6:08c6.b322.33df:32:100.1.1.1/128					
2.2.2.2		0	100	0	i
*>i0:32:2.2.2.2/72					
2.2.2.2		0	100	0	i

Total number of prefixes 3

B

```
B#sh bgp l2vpn evpn all summary
```

For address family: L2VPN EVPN

BGP router identifier 1.1.1.1, local AS number 100

BGP table version is 16

1 BGP AS-PATH entries


```

0 BGP Community entries
6 BGP Prefix entries (Maximum-prefix:4294967295)

Neighbor    V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
172.18.2.1  4    100   639   625    16   0   0 08:55:10   3

Total number of neighbors 1

B#sh bgp l2vpn evpn all
BGP table version is 16, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
                S Stale, b - backup entry
Origin codes: i - IGP, e - EGP, ? - incomplete

Network      Next Hop      Metric  LocPrf  Weight Path
Route Distinguisher: 1.1.1.1:100 (Default for EVI 100)
*> 0:6:0011.2233.2016:0:0.0.0.0/128
           0.0.0.0                32768  i
*>i0:6:08c6.b322.33df:0:0.0.0.0/128
           2.2.2.2                0    100    0  i
*> 0:6:0011.2233.2016:32:100.1.1.2/128
           0.0.0.0                32768  i
*>i0:6:08c6.b322.33df:32:100.1.1.1/128
           2.2.2.2                0    100    0  i
*> 0:32:1.1.1.1/72  0.0.0.0                32768  i
*>i0:32:2.2.2.2/72  2.2.2.2                0    100    0  i

Total number of prefixes 6

```

Common Errors

- No L2VPN EVPN address family neighbor is configured.
- The EVPN address family capability is not activated for BGP neighbors.

7.4.12 Configuring Interconnection with Devices Supporting Only 2-Byte AS Numbers

Configuration

Effect

- Successfully interconnect devices supporting 4-byte AS numbers with devices supporting only 2-byte AS numbers.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring the Display Mode of a 4-Byte AS Number

- (Optional) Perform this configuration in the BGP configuration mode. By default, a 4-byte AS number is displayed as decimal digits.

Verification

- Run the show command to display the neighbor status.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	<code>neighbor { <i>peer-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i></code>
---------	---

Parameter Description	<p><i>peer-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, consisting of no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

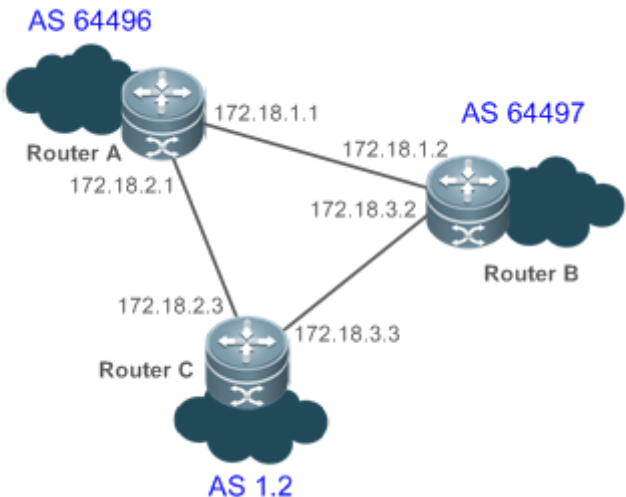
Configuring the Display Mode of a BGP 4-Byte AS Number

Command	bgp asnotation dot
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Configuration

Example

Configuring Compatibility Between BGP Devices Supporting 4-Byte AS Numbers and 2-Byte AS Numbers

<p>Scenario Figure 7-15</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Enable BGP on all devices and set the AS numbers as shown in Figure 7-15. ▪ Configure BGP neighborships.
<p>A</p>	<pre>A# configure terminal A(config)# router bgp 64496 A(config-router)# neighbor 172.18.1.2 remote-as 64497 A(config-router)# neighbor 172.18.2.3 remote-as 23456</pre>
<p>B</p>	<pre>B# configure terminal B(config)# router bgp 64497 B(config-router)# neighbor 172.18.1.1 remote-as 64496 B(config-router)# neighbor 172.18.3.3 remote-as 1.2 B(config-router)# bgp asnotation dot B(config-router)# end</pre>
<p>C</p>	<pre>C# configure terminal C(config)# router bgp 1.2 C(config-router)# neighbor 172.18.2.1 remote-as 64496 C(config-router)# neighbor 172.18.3.2 remote-as 64497</pre>
<p>Verification</p>	<p>Run the show command to display the BGP neighbor status.</p>
<p>A</p>	<pre>A# show ip bgp summary BGP router identifier 172.18.1.1, local AS number 64496 BGP table version is 1, main routing table version 1</pre>

	<pre>Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down Statd 172.18.1.2 4 64497 7 7 1 0 0 00:03:04 0 172.18.2.3 4 23456 4 4 1 0 0 00:00:15 0</pre>
B	<pre>B# show ip bgp summary BGP router identifier 172.18.3.2, local AS number 64497 BGP table version is 1, main routing table version 1 Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down Statd 172.18.1.1 4 64496 7 7 1 0 0 00:00:04 0 172.18.3.2 4 1.2 4 4 1 0 0 00:00:16 0</pre>

Common Errors

N/A

7.4.13 Using Local IPv6 Link Addresses to Establish BGP Neighborships

Configuration

Effect

- Use local IPv6 link addresses to establish BGP neighborships.

Notes

- Generally, global IPv4 addresses need to be used for establishing BGP neighborships.
- Local IPv6 link addresses can be used for establishing only single-hop BGP neighborships.
- When local IPv6 link addresses are used for establishing neighborships, using local IPv6 link addresses as information sources must be specified on the peer end.
- When local IPv6 link addresses are used for establishing neighborships, local IPv6 link addresses must be configured on both ends.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in the global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP configuration mode.

Specifying the Message Source for a BGP Neighbor

- Perform this configuration in the BGP configuration mode. If local IPv6 link addresses are configured for a neighbor, this configuration is mandatory; otherwise, this configuration is optional.

Configuring the BGP IPv4 Address Family Mode

- (Optional) Perform this configuration in the BGP configuration mode.

Disabling the IPv4 Address Family Capability for a BGP Neighbor

- (Optional) Perform this configuration in the BGP IPv6 configuration mode.

Configuring the BGP IPv6 Address Family Mode

- (Mandatory) Perform this configuration in the BGP configuration mode.

Configuring the IPv6 Address Family Capability for a BGP Neighbor

- (Mandatory) Perform this configuration in the BGP IPv6 configuration mode.

Configuring IPv6 Route Advertisement in BGP

- (Optional) Perform this configuration in the BGP IPv6 configuration mode.

Verification

- Run the show command to display the neighbor status.
- Run the show command to display the routing status.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in the dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	<code>neighbor <i>peer-address</i> remote-as <i>as-number</i></code>
----------------	---

Parameter Description	<i>peer-address</i> : Indicates the address of a peer, which is usually an IPv6 address. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS specified for a peer must be the same as the BGP AS number of a BGP speaker at the peer end.

Specifying the Message Source for a BGP Neighbor

Command	<i>neighbor peer-address update-source interface-type interface-number</i>
Parameter Description	<i>peer-address</i> : Indicates the address of a peer, which is usually an IPv6 address. <i>interface-type interface-number</i> : Indicates an interface name.
Command Mode	BGP configuration mode
Usage Guide	If the local IPv6 link address of a local interface is used when a BGP neighborhood is established with a neighbor device, this interface must be specified as the message source of the neighborhood when the BGP neighborhood is configured on the local device.

Configuring the BGP IPv4 Address Family Mode

Command	<i>address-family ipv4 unicast</i>
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Disabling the IPv4 Address Family Capability for a BGP Neighbor

Command	<i>no neighbor peer-address activate</i>
----------------	---

Parameter Description	<i>peer-address</i> : Indicates the address of a peer, which is usually an IPv6 address.
Command Mode	BGP IPv4 address family mode
Usage Guide	Neighbors with IPv6 addresses are used to exchange IPv6 routes. However, when a neighbor is configured in the BGP mode, BGP automatically activates the IPv4 unicast address family capability for the neighbor. Therefore, you are advised to manually disable the IPv4 unicast address family capability.

Configuring the BGP IPv6 Address Family Mode

Command	address-family ipv6 unicast
Parameter Description	-
Command Mode	BGP configuration mode
Usage Guide	-

Configuring the IPv6 Address Family Capability for a BGP Neighbor

Command	neighbor <i>peer-address</i> activate
Parameter Description	<i>peer-address</i> : Indicates the address of a peer, which is usually an IPv6 address.
Command Mode	BGP IPv6 address family mode
Usage Guide	-

Configuring IPv6 Route Advertisement in BGP

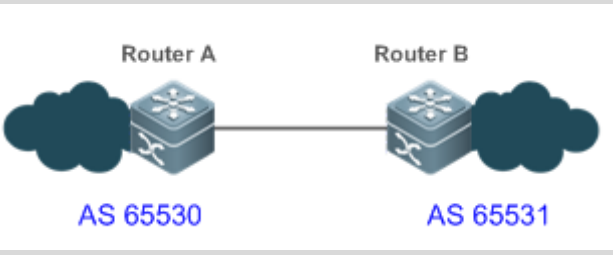
Command	network <i>network-number</i> [mask <i>mask</i>] [route-map <i>map-tag</i>] [backdoor]
---------	---

Parameter Description	<p><i>network-number</i>: Indicates the network number.</p> <p><i>mask</i>: Indicates the subnet mask.</p> <p><i>map-tag</i>: Indicates the name of a route map, consisting of no more than 32 characters.</p> <p>backdoor: Indicates that the route is a backdoor route.</p>
Command Mode	BGP IPv6 address family mode
Usage Guide	-

Configuration

Example

Using a Local IPv6 Link Address for Configuring a BGP Neighborhood to Implement IPv6 Route Exchange in Different ASs

Scenario Figure 7-16	
Configuration Steps	<ul style="list-style-type: none"> Enable BGP on all devices and set the AS numbers as shown in Figure 7-16. Configure a BGP neighbor, specify the update-source, disable the IPv4 address family capability for the neighbor and activate the IPv6 address family capability. Configure IPv6 route advertisement in BGP.
A	<pre> A# configure terminal A(config)# int loopback 0 A(config-if-Loopback)# ipv6 address 30::1/128 A(config-if-Loopback)# exit A(config)# int GigabitEthernet 0/1 A(config-if-Loopback)# ipv6 address fe80:100::2/64 A(config-if-Loopback)# exit A(config)# router bgp 65530 A(config-router)# neighbor fe80:100::1 remote-as 65531 </pre>

	<pre> A(config-router)# neighbor fe80:100::1 update-source GigabitEthernet 0/1 A(config-router)# address-family ipv4 A(config-router-af)# no neighbor fe80:100::1 activate A(config-router-af)# exit-address-family A(config-router)# address-family ipv6 A(config-router-af)# neighbor fe80:100::1 activate A(config-router-af)# network 30::1/128 </pre>
B	<pre> B# configure terminal A(config)# int GigabitEthernet 0/1 A(config-if-Loopback)# ipv6 address fe80:100::1/64 A(config-if-Loopback)# exit B(config)# router bgp 65531 B(config-router)# neighbor fe80:100::2 remote-as 65530 A(config-router)# neighbor fe80:100::2 update-source GigabitEthernet 0/1 B(config-router)# address-family ipv4 B(config-router-af)# no neighbor fe80:100::2 activate B(config-router-af)# exit-address-family B(config-router)# address-family ipv6 B(config-router-af)# neighbor fe80:100::2 activate </pre>
Verification	Run the show command to display the BGP neighbor status.
A	<pre> A# show bgp ipv6 unicast summary BGP router identifier 1.1.1.1, local AS number 65530 BGP table version is 1 1 BGP AS-PATH entries 0 BGP Community entries 1 BGP Prefix entries (Maximum-prefix:4294967295) Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd FE80:100::1 4 65531 4 6 1 0 0 00:01:49 0 </pre>

	Total number of neighbors 1												
B	<p>Run the show command to display BGP routing information.</p> <pre>B# show bgp ipv6 unicast BGP table version is 4, local router ID is 2.2.2.2 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S Stale, b - backup entry Origin codes: i - IGP, e - EGP, ? - incomplete</pre> <table border="1"> <thead> <tr> <th>Network</th> <th>Next Hop</th> <th>Metric</th> <th>LocPrf</th> <th>Weight</th> <th>Path</th> </tr> </thead> <tbody> <tr> <td>*> 30::1/128</td> <td>FE80:100::2</td> <td>0</td> <td></td> <td>0</td> <td>65530 i</td> </tr> </tbody> </table> <p>Total number of prefixes 1</p>	Network	Next Hop	Metric	LocPrf	Weight	Path	*> 30::1/128	FE80:100::2	0		0	65530 i
Network	Next Hop	Metric	LocPrf	Weight	Path								
*> 30::1/128	FE80:100::2	0		0	65530 i								

Common Errors

- When a neighborhood is configured, a local IPv6 link address is used to specify the neighborhood; however, no update source is specified as the interface for this local IPv6 link address during local configuration.
- Only one end uses a local IPv6 link address for establishing a neighborhood.

7.4.14 Configuring BGP NSR

Configuration Effect

- Configure the BGP NSR function to ensure highly reliable network deployment.

Notes

- To deploy the BGP NSR function successfully, the device must support the dual engine redundancy hot backup function.
- After the BGP NSR function is enabled, connections to the BGP neighbors will be reset for the function to take effect.

Configuration Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Configuring the NSR Function for BGP Neighbors

- (Mandatory) Perform this configuration in BGP configuration mode.

Verification

- Run the show command to query the BGP neighbor status.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number, ranging from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	-

Creating a BGP Neighbor

Command	<code>neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i></code>
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode

Usage Guide	The AS specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.
-------------	---


Configuring the NSR Function for BGP Neighbors

Command	<code>neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } ha-mode nsr</code>
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, consisting of no more than 32 characters.
Command Mode	BGP configuration mode
Usage Guide	-

Configuration

Example

Configuring BGP NSR

Scenario Figure 7-17	
Configuration Steps	<ul style="list-style-type: none"> Enable BGP on all devices and set the AS numbers as shown in Figure 7-17. Establish the EBGP neighborship between routers A and B via the directly connected port. Enable the BGP NSR function on A.
A	<pre>A# configure terminal A(config-router)# exit A(config)# router bgp 65530 A(config-router)# neighbor 192.168.195.182 ha-mode nsr</pre>
Verification	Run the show command to query the BGP neighbor status.
A	<pre>A# show ip bgp neighbors</pre>

	<p>BGP neighbor is 192.168.195.183, remote AS 65530, local AS 65531, external link Using BFD to detect fast fallover - BFD session state up BGP version 4, remote router ID 10.0.0.1 BGP state = Established, up for 00:06:37 Last read 00:06:37, hold time is 180, keepalive interval is 60 seconds Neighbor capabilities: Route refresh: advertised and received (old and new) Address family IPv4 Unicast: advertised and received Graceful restart: advertised and received Remote Restart timer is 120 seconds Address families preserved by peer: None </p>
B	<p>B# show ip bgp neighbors BGP neighbor is 192.168.195.183, remote AS 65530, local AS 65531, external link Using BFD to detect fast fallover - BFD session state up BGP version 4, remote router ID 10.0.0.1 BGP state = Established, up for 00:06:37 Last read 00:06:37, hold time is 180, keepalive interval is 60 seconds Neighbor capabilities: Route refresh: advertised and received (old and new) Address family IPv4 Unicast: advertised and received Graceful restart: advertised and received Remote Restart timer is 120 seconds Address families preserved by peer: None </p>

Common

Errors

- IGP is not configured.
- The NSR function is not enabled for BGP neighbors.

7.4.15 Configuring BGP Routes to Be Recursive Only to Host Routes

Configuration

Effect

- Configure BGP routes to be recursive only to host routes, to enable IBGP or multi-EBGP routes to be recursive only to host routes.

Notes

- This feature is effective only to IBGP or multi-hop EBGP routes. It is unavailable to direct EBGP routes or routes whose next-hop addresses are direct connection addresses.
- IBGP or multi-hop EBGP neighborships of the device need to be established using the address of the Loopback interface. Otherwise, this feature is unavailable.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Configuring BGP Routes to Be Recursive Only to Host Routes

- (Mandatory) Perform this configuration in BGP address family mode.

Verification

- Run the show running-config command to display the configuration.

Related

Commands

Enabling BGP

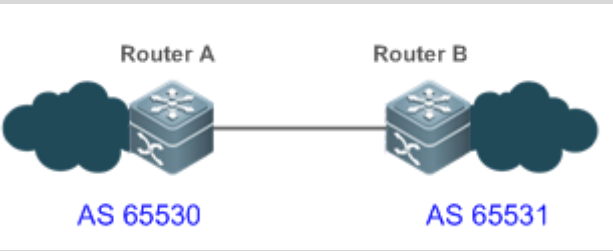
Command	<code>router bgp as-number</code>
Parameter Description	<i>as-number</i> : Indicates an AS number. The value range is from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuring BGP Routes to Be Recursive Only to Host Routes

Command	Bgp recursion host
Parameter Description	N/A
Command Mode	BGP address family mode
Usage Guide	N/A

Configuration Example

Configuring BGP Routes to Be Recursive Only to Host Routes

Scenario Figure 7-18	
Configuration Steps	<ul style="list-style-type: none">▪ Enable BGP on all devices and set the AS numbers as shown in the preceding figure.▪ Establish the multi-hop EBGP neighborships on Router A and Router B through the Loopback interface.▪ Enable the function of making BGP routes recursive only to host routes on Router A and Router B.
Verification	Run the show command to display the BGP neighbor status.

Common Errors

- If the feature is enabled on a device but IBGP or multi-hop EBGP neighborships are not established using the address of the Loopback interface, the device cannot find a host route that matches the next-hop address of a BGP route. As a result, the next hop of the BGP route is invalid and routes cannot be preferentially selected.

7.4.16 Configuring Outbound Loop Detection for a BGP Neighbor

Configuration

Effect

- Configure outbound loop detection for a BGP neighbor

Notes

- This feature is available only to EBGP neighbors.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Configuring Outbound Loop Detection for a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Verification

- Run the show command to display the neighbor status.

Related

Commands

Enabling BGP

Command	<code>router bgp <i>as-number</i></code>
Parameter Description	<i>as-number</i> : Indicates an AS number. The value range is from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	N/A

Creating a BGP Neighbor

Command	<code>neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i></code>
Parameter Description	<p><i>neighbor-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, containing no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS number specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

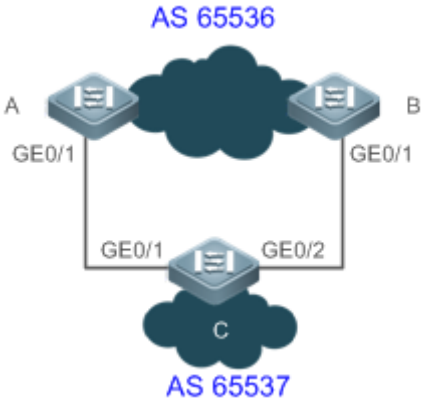
Configuring Outbound Loop Detection for a BGP Neighbor

Command	<code>neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } as-loop-check out</code>
Parameter Description	<p><i>neighbor-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, containing no more than 32 characters.</p>
Command Mode	BGP configuration mode
Usage Guide	N/A

Configuration

Example

Configuring Outbound Loop Detection for a BGP Neighbor

<p>Scenario Figure 7-19</p>	
<p>Configurati on Steps</p>	<ul style="list-style-type: none"> ▪ Enable BGP on all devices and set the AS numbers as shown in the preceding figure. ▪ Establish the EBGP neighborhood between Device A and Device C, and between Device B and Device C. ▪ On Device C, enable outbound loop detection for its neighbors: Device A and Device B.
<p>A</p>	<pre>A# configure terminal A(config)# router bgp 65536 A(config-router)# neighbor 10.1.1.1 remote-as 65537</pre>
<p>B</p>	<pre>B# configure terminal B(config)# router bgp 65536 B(config-router)# neighbor 20.1.1.1 remote-as 65537</pre>
<p>C</p>	<pre>C# configure terminal C(config)# router bgp 65537 C(config-router)# neighbor 10.1.1.2 remote-as 65536 C(config-router)# neighbor 10.1.1.2 as-loop-check out C(config-router)# neighbor 20.1.1.2 remote-as 65536 C(config-router)# neighbor 20.1.1.2 as-loop-check out</pre>
<p>Verification</p>	<p>Run the show command to display the BGP neighbor status.</p>
<p>C</p>	<pre>C# show ip bgp neighbors 10.1.1.2 BGP neighbor is 10.1.1.2, remote AS 65536, local AS 65537, external link Using as path loop detection in announcing route BGP version 4, remote router ID 10.0.0.1 BGP state = Established, up for 00:06:37 Last read 00:06:37, hold time is 180, keepalive interval is 60 seconds</pre>

Neighbor capabilities: Route refresh: advertised and received (old and new) Address family IPv4 Unicast: advertised and received Graceful restart: advertised and received Remote Restart timer is 120 seconds Address families preserved by peer: None ...
--

7.4.17 Shutting Down BGP Connections Gracefully

Configuration

Effect

- Shut down BGP connections gracefully. Service traffic is not interrupted or is interrupted for a very short period of time in the process from feature configuration to actual BGP connection shutdown.

Notes

N/A

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Shutting Down Connections of the BGP Neighbor Gracefully

- (Optional) Perform this configuration in BGP configuration mode.

Shutting Down All Connections of a BGP Instance Gracefully

- (Optional) Perform this configuration in BGP configuration mode.

Verification

- Run the **show** command to display the neighbor status.

Related

Commands

Enabling BGP

Command	router bgp <i>as-number</i>
Parameter Description	<i>as-number</i> : Indicates an AS number. The value range is from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	N/A

Creating a BGP Neighbor

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i>
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, containing no more than 32 characters. <i>as-number</i> : Indicates the AS number of a BGP peer (group).
Command Mode	BGP configuration mode
Usage Guide	The AS number specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Shutting Down Connections of the BGP Neighbor Gracefully

Command	neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } shutdown graceful [<i>community value</i>] [<i>delay time</i>]
Parameter Description	<i>neighbor-address</i> : Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address. <i>peer-group-name</i> : Specifies the name of a peer group, containing no more than 32 characters. <i>community value</i> : Indicates the community value carried in the route sent to the neighbor.

	delay <i>time</i> : Indicates the delay time (in seconds) for shutting down BGP connections. The value range is from 1 to 65,535.
Command Mode	BGP configuration mode
Usage Guide	N/A

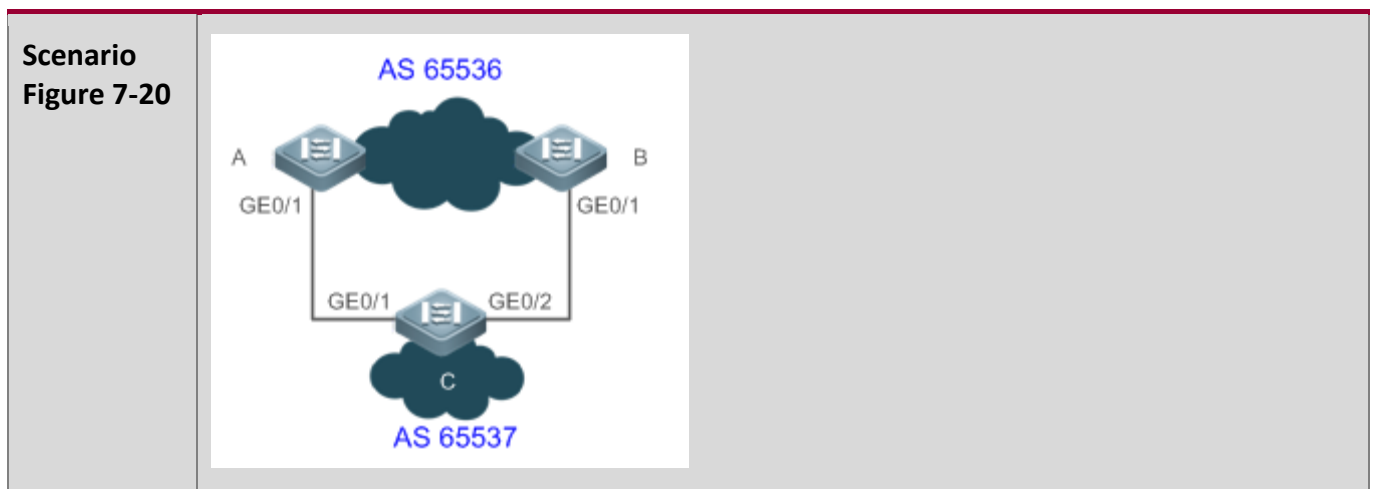
Shutting Down All Connections of a BGP Instance Gracefully

Command	bgp shutdown graceful [community <i>value</i>] [delay <i>time</i>]
Parameter Description	community <i>value</i> : Indicates the community value carried in the route sent to the neighbor. delay <i>time</i> : Indicates the delay time (in seconds) for shutting down BGP connections. The value range is from 1 to 65,535.
Command Mode	BGP configuration mode
Usage Guide	N/A

Configuration

Example

Shutting Down BGP Connections Gracefully



Configurati on Steps	<ul style="list-style-type: none"> ▪ Enable BGP on all devices and set the AS numbers as shown in the preceding figure. ▪ Establish the EBGP neighborship between Device A and Device C and between Device B and Device C. ▪ On Device A, shut down all connections of a BGP instance gracefully.
A	<pre>A# configure terminal A(config)# router bgp 65536 A(config-router)# neighbor 10.1.1.1 remote-as 65537 A(config-router)# bgp shutdown graceful med 4294967295</pre>
B	<pre>B# configure terminal B(config)# router bgp 65536 B(config-router)# neighbor 20.1.1.1 remote-as 65537</pre>
C	<pre>C# configure terminal C(config)# router bgp 65537 C(config-router)# neighbor 10.1.1.2 remote-as 65536 C(config-router)# neighbor 10.1.1.2 as-loop-check out C(config-router)# neighbor 20.1.1.2 remote-as 65536 C(config-router)# neighbor 20.1.1.2 as-loop-check out</pre>
Verification	Run the show command to display the BGP neighbor status.
A	<pre>A(config)#show ip bgp neighbor 10.1.1.2 For address family: IPv4 Unicast BGP neighbor is 10.10.10.2, remote AS 65537, local AS 65536, external link Administratively graceful shut down BGP version 4, remote router ID 0.0.0.0 BGP state = Idle Last read , hold time is 30, keepalive interval is 1 seconds Configured hold time is 30, keepalive interval is 1 seconds Received 0 messages, 0 notifications, 0 in queue open message:0 update message:0 keepalive message:0 refresh message:0 dynamic cap:0 notifications:0 Sent 0 messages, 0 notifications, 0 in queue</pre>

```
open message:0 update message:0 keepalive message:0
refresh message:0 dynamic cap:0 notifications:0
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds
Update source is Loopback 0

For address family: IPv4 Unicast
BGP table version 1, neighbor version 0
Index 1, Offset 0, Mask 0x2
0 accepted prefixes
0 announced prefixes

Connections established 0; dropped 0
BGP neighbor may be up to 255 hops away.
Sock_fd: -1
Last Reset:      , due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

7.4.18 Configuring EBGP Multi-path Bypass Protection

Configuration

Effect

- Implement EBGP multi-path bypass protection.

Notes

- BFD sessions can be configured for neighbors to implement fast link fault detection.

Configuration

Steps

Enabling BGP

- (Mandatory) Perform this configuration in global configuration mode.

Creating a BGP Neighbor

- (Mandatory) Perform this configuration in BGP configuration mode.

Configuring BGP ECMP

- (Mandatory) Perform this configuration in BGP configuration mode.

Configuring BGP Multi-path Bypass Protection

- (Mandatory) Perform this configuration in BGP configuration mode.

Configuring a BFD Session to a BGP Neighbor

- (Optional) Perform this configuration in BGP configuration mode.

Verification

- Run the **show** command to display route information.

Related

Commands

Enabling BGP

Command	<code>router bgp as-number</code>
Parameter Description	<i>as-number</i> : Indicates an AS number. The value range is from 1 to 4,294,967,295, which is 1 to 65535.65535 in dot mode.
Command Mode	Global configuration mode
Usage Guide	N/A

Configuring BGP Multi-path Bypass Protection

Command	<code>bgp install standby-path</code>
Parameter Description	N/A
Command Mode	BGP configuration mode
Usage Guide	N/A

Creating a BGP Neighbor

Command	<code>neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } remote-as <i>as-number</i></code>
Parameter Description	<p><i>neighbor-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, containing no more than 32 characters.</p> <p><i>as-number</i>: Indicates the AS number of a BGP peer (group).</p>
Command Mode	BGP configuration mode
Usage Guide	The AS number specified for a peer (group) must be the same as the BGP AS number of a BGP speaker at the peer end.

Creating a BFD Session to a BGP Neighbor

Command	<code>neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } fall-over bfd</code>
Parameter Description	<p><i>neighbor-address</i>: Specifies the address of a peer. This address may be an IPv4 address or an IPv6 address.</p> <p><i>peer-group-name</i>: Specifies the name of a peer group, containing no more than 32 characters.</p>
Command Mode	BGP configuration mode
Usage Guide	N/A


Configuration

Example

- i** The following configuration example describes only the configuration relevant to BGP multi-path bypass protection.

Configuring BGP Multi-path Bypass Protection

**Scenario
Figure 7-21**



**Configurati
on Steps**

- Enable BGP on all devices.
- Configure addresses for directly-connected ports on Device A, Device B, and Device C, and mutually establish the EBGP neighborship between every two of them.
- Configure a BFD session for the EBGP neighborship between Device A and Device C.
- Configure BGP multi-path bypass protection on Device C.
- Configure ECMP on Device C.

A

```
A# conf terminal
A(config)# interface GigabitEthernet 0/1
A(config-if-GigabitEthernet 0/1)# ip address 192.168.1.1 255.255.255.0
A(config-if-GigabitEthernet 0/1)# exit
A(config)# interface GigabitEthernet 0/2
A(config-if-GigabitEthernet 0/2)# ip address 192.168.2.1 255.255.255.0
A(config-if-GigabitEthernet 0/2)# exit
A(config)# interface GigabitEthernet 0/3
A(config-if-GigabitEthernet 0/3)# ip address 192.168.4.1 255.255.255.0
A(config-if-GigabitEthernet 0/3)# exit
A(config)# router bgp 100
A(config-router)# neighbor 192.168.1.2 remote-as 300
```

	<pre>A(config-router)# neighbor 192.168.4.2 remote-as 300 A(config-router)# neighbor 192.168.2.2 remote-as 200 A(config-router)# maximum-paths ebgp 2 A(config-router)# redistribute connect</pre>
B	<pre>B# configure terminal B(config)# interface GigabitEthernet 0/1 B(config-if-GigabitEthernet 0/1)# ip address 192.168.3.1 255.255.255.0 B(config-if-GigabitEthernet 0/1)# bfd interval 200 min_rx 200 multiplier 5 B(config-if-GigabitEthernet 0/1)# exit B(config)# interface GigabitEthernet 0/2 B(config-if-GigabitEthernet 0/2)# ip address 192.168.2.2 255.255.255.0 B(config-if-GigabitEthernet 0/2)# exit B(config)# router bgp 200 B(config-router)# neighbor 192.168.3.2 remote-as 300 B(config-router)# neighbor 192.168.3.2 fall-over bfd B(config-router)# neighbor 192.168.2.1 remote-as 100 B(config-router)# redistribute connect</pre>
C	<pre>C# configure terminal C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ip address 192.168.1.2 255.255.255.0 C(config-if-GigabitEthernet 0/1)# bfd interval 200 min_rx 200 multiplier 5 C(config-if-GigabitEthernet 0/1)# exit C(config)# interface fastEthernet 0/2 C(config-if-GigabitEthernet 0/2)# ip address 192.168.3.2 255.255.0.0 C(config-if-GigabitEthernet 0/2)# bfd interval 200 min_rx 200 multiplier 5 C(config-if-GigabitEthernet 0/2)# exit C(config)# interface fastEthernet 0/3 C(config-if-GigabitEthernet 0/3)# ip address 192.168.4.2 255.255.0.0 C(config-if-GigabitEthernet 0/3)# bfd interval 200 min_rx 200 multiplier 5 C(config-if-GigabitEthernet 0/3)# exit C(config)# router bgp 300</pre>

```

C(config-router)# neighbor 192.168.1.1 remote-as 100
C(config-router)# neighbor 192.168.4.1 remote-as 100
C(config-router)# neighbor 192.168.3.1 remote-as 200
C(config-router)# neighbor 192.168.1.1 fall-over bfd
C(config-router)# neighbor 192.168.4.1 fall-over bfd
C(config-router)# maximum-paths ebgp 2
C(config-router)# address-family ipv4 unicast
C(config-router-af)# bgp install standby-path
C(config-router-af)# redistribute connect

```

Verification Run the **show** command to display the configuration.

C

```

C# show ip bgp summary
BGP router identifier 10.10.10.10, local AS number 300
BGP table version is 12
4 BGP AS-PATH entries
0 BGP Community entries
4 BGP Prefix entries (Maximum-prefix:4294967295)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.1.1	4	100	76	77	12	12	0	00:59:27	3
192.168.3.1	4	200	30	30	12	12	0	00:19:03	3
192.168.4.1	4	100	76	77	12	12	0	00:59:01	3

Total number of neighbors 3

```

C# show ip bgp
BGP table version is 12, local router ID is 10.10.10.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
                S Stale, b - backup entry, m - multipath
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
---------	----------	--------	--------	--------	------

```

* 192.168.1.0 192.168.3.1 0 0 200 ?
* 192.168.4.1 0 0 100 ?
* 192.168.1.1 0 0 100 ?
*> 0.0.0.0 0 32768 ?
*> 192.168.2.0 192.168.4.1 0 0 200 ?
*m 192.168.1.1 0 0 200 ?
*b 192.168.3.1 0 0 100 ?
* 192.168.3.0 192.168.3.1 0 0 200 ?
* 192.168.4.1 0 0 100 200 ?
* 192.168.1.1 0 0 100 200 ?
*> 0.0.0.0 0 32768 ?
* 192.168.4.0 192.168.3.1 0 0 200 ?
* 192.168.4.1 0 0 100 ?
* 192.168.1.1 0 0 100 ?
*> 0.0.0.0 0 32768 ?

```

Total number of prefixes 4

C# show ip bgp 192.168.2.0

BGP routing table entry for 192.168.2.0/24

Paths: (3 available, best #3, table Default-IP-Routing-Table)

Advertised to non peer-group peers:

192.168.1.1

200

192.168.3.1 from 192.168.3.1 (3.3.3.3)

Origin incomplete, metric 0, localpref 100, valid, external, backup

Last update: Tue Oct 5 00:26:52 1971

100

192.168.1.1 from 192.168.1.1 (44.44.44.44)

Origin incomplete, metric 0, localpref 100, valid, external, multipath

Last update: Mon Oct 4 23:46:28 1971

100

```
192.168.4.1 from 192.168.4.1 (44.44.44.44)
Origin incomplete, metric 0, localpref 100, valid, external, multipath, best
Last update: Mon Oct 4 23:46:28 1971
```

```
C# show ip route
```

```
Codes: C - connected, S - static, R - RIP, B - BGP
```

```
O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
ia - IS-IS inter area, * - candidate default
```

```
Gateway of last resort is no set
```

```
C 192.168.1.0/24 is directly connected, GigabitEthernet 0/1
```

```
C 192.168.1.2/32 is local host.
```

```
B 192.168.2.0/24 [20/0] via 192.168.1.1, 00:21:39
```

```
[20/0] via 192.168.4.1, 00:21:39
```

```
C 192.168.3.0/24 is directly connected, GigabitEthernet 0/2
```

```
C 192.168.3.2/32 is local host.
```

```
C 192.168.4.0/24 is directly connected, GigabitEthernet 0/3
```

```
C 192.168.4.2/32 is local host.
```

7.4.19 Configuring Inter-VRF Multi-Path Route Import

Configuration

Effect

- Ensure inter-VRF route import and ECMP multi-path for inter-imported routes.

Notes

- N/A

Configuration

Steps

Configuring a BGP VRF Address Family

- (Mandatory) Perform this configuration in BGP configuration mode.

Importing Static Routes to BGP

- (Mandatory) Perform this configuration in BGP address family mode.

Importing Multi-Path Static Routes to BGP

- (Mandatory) Perform this configuration in BGP address family mode.

Configuring BGP ECMP

- (Mandatory) Perform this configuration in BGP address family mode.

Configuring Inter-VRF Import for All Path Routes

- (Mandatory) Perform this configuration in BGP address family mode.

Verification

- Run the **show** command to display route information.

Related

Commands

Configuring a BGP VRF Address Family

Command	address-family ipv4 vrf <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the VRF instance name.
Command Mode	BGP configuration mode
Usage Guide	N/A

Importing Static Routes to BGP

Command	redistribute <i>protocol-type</i> [route-map <i>map-tag</i>] [metric <i>metric-value</i>]
Parameter Description	<p><i>protocol-type</i>: Indicates the source protocol type of a redistributed route.</p> <p>route-map <i>map-tag</i>: Indicates the name of an associated route map.</p> <p><i>metric <i>metric-value</i></i>: Indicates the default metric value of a redistributed route. The value range is from 0 to 4,294,967,295.</p>

Command Mode	BGP address family mode
Usage Guide	N/A

Importing Multi-Path Static Routes to BGP

Command	bgp sourced-paths <i>protocol-type</i> all
Parameter Description	<i>protocol-type</i> : Indicates the source protocol type of a redistributed route.
Command Mode	BGP address family mode
Usage Guide	This command needs to be used together with the redistribution command to import routes with multiple next hops from other protocols to BGP.

Configuring BGP ECMP

Command	maximum-paths { ebgp ibgp } <i>number</i>
Parameter Description	<p>ebgp: Specifies the number of equivalent paths of the EBGp multipath load balancing function.</p> <p>ibgp: Specifies the number of equivalent paths of the IBGP multipath load balancing function.</p> <p><i>number</i>: Indicates the maximum number of equivalent paths. The minimum value is 1, and the maximum value depends on the device capability. If the value is 1, the EBGp multipath load balancing function is disabled.</p>
Command Mode	BGP address family mode
Usage Guide	<p>The maximum-paths ebgp command is also used to configure equivalence of confederation EBGp multiple paths and local inter-VRF import routes.</p> <p>IBGP and EBGp routes cannot form equivalent routes.</p>

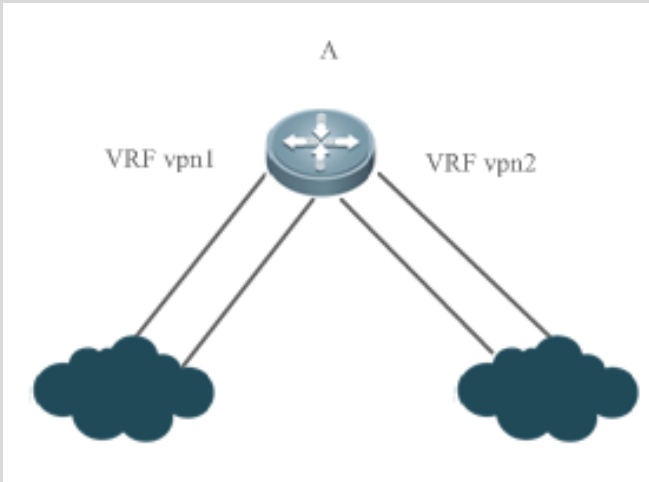
Configuring Inter-VRF Import for All Path Routes

Command	import path selection { all bestpath multipath }
Parameter Description	<p>all: Imports all routes with next hops.</p> <p>bestpath: Imports routes with preferred next hops. By default, only routes with preferred next hops are imported.</p> <p>multipath: Imports routes with preferred and equivalent next hops.</p>
Command Mode	BGP address family mode
Usage Guide	This command can be used to control inter-VRF route import, L3VPN remote routes import to VRF, and EVPN routes import to the IP route table.

Configuration

Example

Configuring BGP Multi-Path Bypass Protection

Scenario Figure 7-22	 <p>The diagram shows a central router labeled 'A' at the top. Below it are two cloud-like shapes representing networks. The left network is connected to the router via a line labeled 'VRF vpn1'. The right network is connected via a line labeled 'VRF vpn2'. The router 'A' has a circular icon with four arrows pointing outwards, indicating its central role in the network.</p>
	<p>Device A connects to two networks through VRF vpn1 and vpn2, and cross-VRF access is implemented through device A.</p>
Configuration Steps	<p>Configure VRF.</p> <p>Configure VRF static routes.</p> <p>Configure a VRF address family.</p> <p>Import VRF static routes to BGP.</p> <p>Enable multi-path static route import to BGP.</p>

	<p>Configure BGP ECMP.</p> <p>Configure inter-VRF import for all path routes.</p>
A	<pre>A# conf terminal A(config)# ip vrf vpn1 A(config-vrf)# rd 200:1 A(config-vrf)# route-target both 100:100 A(config-vrf)# exit A(config)# ip vrf vpn2 A(config-vrf)# rd 300:1 A(config-vrf)# route-target both 100:100 A(config-vrf)# exit A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# A(config-if-GigabitEthernet 0/1)# ip vrf forwarding vpn1 A(config-if-GigabitEthernet 0/1)# ip address 44.1.1.2 255.255.255.0 A(config-if-GigabitEthernet 0/1)# exit A(config)# interface GigabitEthernet 0/2 A(config-if-GigabitEthernet 0/2)# ip vrf forwarding vpn1 A(config-if-GigabitEthernet 0/2)# ip address 45.1.1.2 255.255.255.0 A(config-if-GigabitEthernet 0/2)# exit A(config)# interface GigabitEthernet 0/3 A(config-if-GigabitEthernet 0/3)# ip vrf forwarding vpn2 A(config-if-GigabitEthernet 0/3)# ip address 46.1.1.2 255.255.255.0 A(config-if-GigabitEthernet 0/3)# exit A(config)# interface GigabitEthernet 0/4 A(config-if-GigabitEthernet 0/4)# ip vrf forwarding vpn2 A(config-if-GigabitEthernet 0/4)# ip address 47.1.1.2 255.255.255.0 A(config-if-GigabitEthernet 0/4)# exit A(config)# ip route vrf vpn1 100.1.1.1 255.255.255.255 44.1.1.1</pre>

```

A(config)# ip route vrf vpn1 100.1.1.1 255.255.255.255
45.1.1.1
A(config)# ip route vrf vpn2 200.1.1.1 255.255.255.255
46.1.1.1
A(config)# ip route vrf vpn2 200.1.1.1 255.255.255.255
47.1.1.1
A(config)# router bgp 100
A(config-router)# address-family ipv4 vrf vpn1
A(config-router-af)# redistribute static
A(config-router-af)# maximum-paths ebgp 32
A(config-router-af)# bgp sourced-paths static all
A(config-router-af)# import path selection all
A(config-router-af)# exit-address-family
A(config-router)# address-family ipv4 vrf vpn2
A(config-router-af)# redistribute static
A(config-router-af)# maximum-paths ebgp 32
A(config-router-af)# bgp sourced-paths static all
A(config-router-af)# import path selection all
A(config-router-af)# exit-address-family

```

Verification

Run the show command to display the configurations.

A

```

A#show ip route vrf vpn1
Routing Table: vpn1

Codes: C - Connected, L - Local, S - Static
        R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow
route
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA
external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS
level-2
        IA - Inter area, EV - BGP EVPN, A - Arp to host
        * - candidate default

```

```
Gateway of last resort is no set
C    44.1.1.0/24 is directly connected, GigabitEthernet 0/1
C    44.1.1.2/32 is local host.
C    45.1.1.0/24 is directly connected, GigabitEthernet 0/2
C    45.1.1.2/32 is local host.
S    100.1.1.1/32 [1/0] via 44.1.1.1
      [1/0] via 45.1.1.1
B    200.1.1.1/32 [20/0] via 47.1.1.1, 02:32:01
      [20/0] via 46.1.1.1, 02:32:01
```

```
A#show ip route vrf vpn2
```

```
Routing Table: vpn2
```

```
Codes: C - Connected, L - Local, S - Static
        R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow
route
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA
external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS
level-2
        IA - Inter area, EV - BGP EVPN, A - Arp to host
        * - candidate default
```

```
Gateway of last resort is no set
C    46.1.1.0/24 is directly connected, GigabitEthernet 0/3
C    46.1.1.2/32 is local host.
C    47.1.1.0/24 is directly connected, GigabitEthernet 0/4
C    47.1.1.2/32 is local host.
B    100.1.1.1/32 [20/0] via 45.1.1.1, 03:27:07
      [20/0] via 44.1.1.1, 03:27:07
S    200.1.1.1/32 [1/0] via 46.1.1.1
      [1/0] via 47.1.1.1
```

7.5 Monitoring

Clearing

! Running the **clear** commands may lose vital information and thus interrupt services.

Description	Command
Clears BGP IPv4 unicast routes.	clear ip bgp [vrf <i>vrf-name</i>] { * <i>as-number</i> <i>peer-address</i> } [soft] [in out] clear bgp ipv4 unicast [vrf <i>vrf-name</i>] { * <i>as-number</i> <i>peer-address</i> } [soft] [in out]
Clears BGP IPv6 unicast routes.	clear bgp ipv6 unicast [vrf <i>vrf-name</i>] { * <i>as-number</i> <i>peer-address</i> } [soft] [in out]
Clears BGP L2VPN EVPN routes.	clear bgp l2vpn evpn { * <i>as-number</i> <i>neighbor-address</i> } [soft] [in out]
Restores the priority of advertised routes.	clear bgp advertise lowest-priority on-startup

Displaying

Description	Command
Displays BGP IPv4 unicast routes.	show ip bgp show bgp ipv4 unicast
Displays BGP IPv6 unicast routes.	show bgp ipv6 unicast
Displays BGP L2VPN EVPN routes.	show bgp l2vpn evpn all
Displays BGP statistics.	show bgp statistics [vrf <i>vrf-name</i>]

Debugging

⚠ System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Enables all BGP debugging.	debug ip bgp all
Debugs BGP route flapping.	debug ip bgp dampening
Debugs BGP event processing.	debug ip bgp event
Debugs BGP route filtering.	debug ip bgp filter
Debugs BGP status machine.	debug ip bgp fsm
Debugs BGP neighbor keepalive.	debug ip bgp keepalives
Debugs BGP core route processing.	debug ip bgp nsm
Debugs BGP UPDATE packets.	debug ip bgp update
Debugs BGP EVPN.	debug ip bgp evpn

8 CONFIGURING RIPNG

8.1 Overview

RIP next generation (RIPng) is a unicast routing protocol that applies to IPv6 networks. RIPng-enabled routers exchange routing information to obtain routes to remote networks.

As an Interior Gateway Protocol (IGP), RIPng can run only within the autonomous system (AS) and is applicable to small-sized networks with routes no more than 16 hops.

Protocols and Standards


- RFC2080: Defines the RIPng.

8.2 Application

RIPng is generally used on some small-sized networks, such as office networks of small companies.

As shown in the following figure, the company builds an IPv6 network, on which all routers support IPv6. The network is small in size, but the workload is still heavy if the network is maintained manually. In this case, RIPng can be configured to adapt to topological changes of the small-sized network, which reduces the workload.

Figure 8-1



8.3 Features

Basic Concepts

IGP and EGP

IGP runs within an AS. For example, RIPng is a type of IGP.

Exterior Gateway Protocol (EGP) runs between ASs. For example, BGP is a type of EGP.

Feature

Feature	Description
RIPng and RIP	RIPng is an extension of RIPv2 on the basis of IPv6. Both are similar in functions and configurations.
Exchanging Routing Information	By exchanging routing information, RIPng-enabled devices can automatically obtain routes to a remote network and update routes in real time.
Routing Algorithm	RIPng is a protocol based on the distance-vector algorithm. It uses the vector addition method to compute the routing information.
Avoiding Route Loops	RIPng uses functions, such as split horizon and poison reverse, to avoid route loops.

8.3.1 RIPng and RIP

RIP applies to IPv4 networks. Two RIP versions are available, including RIPv1 and RIPv2.

RIPng is an extension of RIPv2 on the basis of IPv6. Both are similar in functions and configurations.

Working Principle

RIPv2

RIPv2 packets are multicast. The multicast address is 224.0.0.9, and the UDP port ID is 520. RIPv2 can identify the subnet mask.

RIPng

RIPng packets are multicast. The multicast address is FF02::9, the source address is FE80::/10, and the UDP port ID is 521. RIPng can identify the subnet mask.

i This chapter describes functions and configurations of RIPng. For details about RIPv2, see "Configuring RIP".

Related Configuration

Enabling the RIPng Process

By default, the RIPng process is disabled.

Run the **ipv6 router rip** command to enable the RIPng process.

You must enable the RIPng process on a device; otherwise, all functions related to RIPng cannot take effect.

Running RIPng on an Interface

By default, RIPng does not run on an interface.

Run the **ipv6 rip enable** command to run RIPng on an interface.

After RIPng runs on an interface, RIPng packets can be exchanged on the interface and RIPng can learn routes to the network segments directly connected to the device.

Prohibiting an Interface from Sending or Receiving Packets

By default, a RIPng-enabled interface is allowed to send and receive RIPng packets.

Run the **passive-interface** command to prohibit an interface from sending RIPng packets.

8.3.2 Exchanging Routing Information

Compared with static routing, the dynamic routing protocol has a significant advantage, that is, by exchanging routing information, devices can automatically obtain routes to a remote network and update the routes in real time.

Working Principle

Initialization

After RIPng is enabled on a router, the router sends a request packet to its neighbor router, requesting for all routing information, that is, the routing table. After receiving the request message, the neighbor router returns a response packet containing the local routing table. After receiving the response packet, the router updates the local routing table, and sends an update packet to the neighbor router, informing the neighbor router of the route update information. After receiving the update packet, the

neighbor router updates the local routing table, and sends the update packet to other adjacent routers. After a series of updates, all routers can obtain and retain the latest routing information.

Periodical Update

By default, periodical update is enabled for RIPng. Adjacent routers exchange complete routing information with each other every 30s (update timer), that is, the entire routing table is sent to neighbor routers.

- i** For every non-local route, if the route is not updated within 180s (invalid timer), the metric of the route is changed to 16 (unreachable). If the route is still not updated in the next 120s (flush timer), the route is deleted from the routing table.

Default Route

In the routing table, a route to the destination network `::/0` is called default route.

The default route can be learned from a neighbor router, or sent to a neighbor router.

Route Redistribution

For RIPng, other types of routes (such as direct routes, static routes, and routes of other routing protocols) are called external routes.

External routes (excluding the default route) can be redistributed to RIPng and advertised to neighbors.

Route Filtering

Filtering conditions can be configured to limit the routing information exchanged between adjacent routers. Only the routing information that meets filtering conditions can be sent or received.

Related Configuration

RIPng Timers

By default, the update timer is 30s, the invalid timer is 180s, and the flush timer is 120s.

Run the **timers basic** command to modify durations of RIPng timers.

Increasing the duration of the flush timer can reduce the route flapping. Decreasing the duration of the flush timer helps accelerate route convergence.

The durations of RIPng timers must be consistent on adjacent routers. Unless otherwise required, you are advised not to modify the RIPng timers.

Default Route

Run the **ipv6 rip default-information** command to advertise the default route to neighbors on an interface.

Route Redistribution

Run the **redistribute** command to redistribute external routes (excluding the default route) to RIPng and advertise them to neighbors.

Route Filtering

Run the **distribute-list out** command to set filtering rules to limit the routing information sent by the device.

Run the **distribute-list in** command to set filtering rules to limit the routing information received by the device.

8.3.3 Routing Algorithm

RIPng is a protocol based on the distance-vector algorithm. It uses the vector addition method to compute the routing information.

Working Principle


Distance-Vector Algorithm

RIPng is a protocol based on the distance-vector algorithm. The distance-vector algorithm treats a route as a vector that consists of the destination network and distance (metric). The router obtains a route from its neighbor and adds the distance vector from itself to the neighbor to the route to form its own route.

RIPng uses the hop count to evaluate the distance (metric) to the destination network. By default, the hop count from a router to its directly connected network is 0, the hop count from a router to a network that can be reached through a router is 1, and so on. That is, the metric is equal to the number of routers from the local network to the destination network. To restrict the convergence time, RIPng stipulates that the metric must be an integer between 0 and 15. If the metric is equal to or greater than 16, the destination network or host is unreachable. For this reason, RIPng cannot be applied to a large-scale network.

As shown in the following figure 8-2, Router A is connected to the network 2::/64. Router B obtains the route (2::/64, 0) from Router A and adds the metric 1 to the route to obtain its own route (2::/64, 1), and the next hop points to Router A.

Figure 8-2




Selecting the Optimum Route

RIPng selects an optimum route based on the following principle: If multiple routes to the same destination network is available, a router preferentially selects the route with the smallest metric.

As shown in the following figure 8-3, Router A is connected to the network 2::/64. Router C obtains the route (2::/64, 0) from Router A and the route (2::/64, 1) from Router B. Router C will select the route that is obtained from Router A and add metric 1 to this route to form its own route (2::/64, 1), and the next hop points to Router A.

Figure 8-3



i When routes coming from different sources exist on a router, the route with the smaller distance is preferentially selected.

Route Source	Default Distance
Directly-connected network	0
Static route	1
OSPF route	110

IS-IS route	115
RIPng route	120
Unreachable route	255

Related Configuration

Modifying the Distance

By default, the distance of a RIPng route is 120.

Run the **distance** command to modify the distance of a RIPng route.

Modifying the Metric

For a RIPng route that is proactively discovered by a device, the default metric is equal to the number of hops from the local network to the destination network. The metric offset of the interface is 1.

For a RIPng router that is manually configured (default route or redistributed route), the default metric is 1.

Run the **ipv6 rip metric-offset** command to modify the metric offset of the interface.

Run the **default-metric** command to modify the default metric of an external route (redistributed route).

Run the **redistribute** command to modify the metric of an external route (redistributed route) when advertising this route.

Run the **ipv6 rip default-information** command to modify the metric of a default route when advertising the default route.

8.3.4 Avoiding Route Loops

RIPng uses functions, such as split horizon and poison reverse, to avoid route loops.

Working Principle


Route Loop

A RIPng route loop occurs due to inherent defects of the distance-vector algorithm.

As shown in the following figure 8-4, Router A is connected to the network 2::/64, and sends an update packet every 30s. Router B receives the route to 2::/64 from Router A every 30s. If Router A is

disconnected from 2::/64, the route to 2::/64 will be deleted from the routing table on Router A. Next time, the update packet sent by Router A no longer contains this route. As Router B does not receive an update packet related to 2::/64, Router B determines that the route to 2::/64 is valid within 180s and uses the update packet to send this route to Router A. As the route to 2::/64 does not exist on Router A, the route learned from Router B is added to the routing table. Router B determines that data can reach 2::/64 through Router A, and Router A determines that data can reach 2::/64 through Router B. In this way, a route loop is formed.

Figure 8-4




Split Horizon

Split horizon can prevent route loops. After split horizon is enabled, a route received on this interface will not be sent out from this interface.

As shown in the following figure 8-5, after split horizon is enabled on Router B, Router B will not send the route to 2::/64 back to Router A. Router B will learn 180s later that 2::/64 is not reachable.

Figure 8-5




Poison Reverse

Poison reverse can also prevent route loops. Compared with split horizon, poison reverse is more reliable, but brings more protocol packets, which makes network congestion more severe.

After poison reverse is enabled on an interface, a route received from this interface will be sent out from this interface again, but the metric of this router will be changed to 16 (unreachable).

As shown in the following figure, after poison reverse is enabled on Router A, if Router A detects a disconnection from 2::/64, Router A will not delete the route to 2::/64. Instead, Router A changes the number of hops to 16, and advertises the route through the update packet. On receiving the update packet, Router B learns that 2::/64 is not reachable.

Figure 8-6



Related Configuration

Split Horizon

By default, split horizon is enabled.


Run the **no split-horizon** command to disable split horizon.

Poison Reverse

By default, poison reverse is disabled.

Run the **split-horizon poisoned-reverse** command to enable poison reverse. (After poison reverse is enabled, split horizon is automatically disabled.)

8.4 Configuration

Configuration	Related Commands	
Configuring RIPng Basic Functions	 (Mandatory) It is used to build a RIPng routing domain.	
	ipv6 router rip	Enables a RIPng routing process and enters routing process configuration mode.
	ipv6 rip enable	Runs RIPng on an interface.

	split-horizon	Enables split horizon or poison reverse.
	passive-interface	Configures a passive interface.
Advertising the Default Route or External Routes	⚠ Optional.	
	ipv6 rip default-information	Advertise the default route to neighbors on an interface.
	redistribute	Redistributes routes and advertising external routes to neighbors.
Setting Route Filtering Rules	⚠ Optional.	
	distribute-list in	Filters the received RIPng routing information.
	distribute-list out	Filters the sent RIPng routing information.
Modifying Route Selection Parameters	⚠ Optional.	
	distance	Modifies the administrative distance of a RIPng route.
	ipv6 rip metric-offset	Modifies the metric offset on an interface.
	default-metric	Configure the default metric for route redistribution.
Modifying Timers	⚠ Optional.	
	timers	Modifies the update timer, invalid timer, and flush timer of RIPng.

8.4.1 Configuring RIPng Basic Functions

Configuration

Effect

- Build a RIPng routing domain on the network.
- Routers in the domain obtain routes to a remote network through RIPng.

Notes

- IPv6 addresses must be configured.
- IPv6 unicast routes must be enabled.

Configuration

Steps

Enabling a RIPng Routing Process

- Mandatory.
- Unless otherwise required, perform this configuration on every router in the RIPng routing domain.

Running RIPng on an Interface

- Mandatory.
- Unless otherwise required, perform this configuration on every interconnected interface of routers in the RIPng routing domain.

Enabling Split Horizon or Poison Reverse

- By default, split horizon is enabled and poison reverse is disabled.
- Unless otherwise required, enable split horizon on every interface connected to the broadcast network, such as the Ethernet. (Retain the default setting.)
- Unless otherwise required, enable split horizon on every interface connected to the point-to-point (P2P) network, such as the PPP and HDLC. (Retain the default setting.)
- It is recommended that split horizon and poison reverse be disabled on an interface connected to a non-broadcast multi-access network, such as FR and X.25; otherwise, some devices cannot learn the complete routing information.
- If the secondary IP address is configured for an interface connected to a non-broadcast, it is recommended that split horizon and poison reverse be disabled.

Configuring a Passive Interface

- This configuration is recommended.
- Use the passive interface to set the boundary of the RIPng routing domain. The network segment of the passive interface belongs to the RIPng routing domain, but RIPng packets cannot be sent over the passive interface.
- If RIPng routes need to be exchanged on an interface (such as the router interconnect interface) in the RIPng routing domain, this interface cannot be configured as a passive interface.

Verification

- Check the routing table on a router to verify that the route to a remote network can be obtained through RIPng.

Related Commands

Enabling a RIPng Routing Process

Command	ipv6 router rip
Parameter Description	N/A
Command Mode	Global configuration mode
Usage Guide	This command is used to create a RIPng routing process and enter routing process configuration mode.

Running RIPng on an Interface

Command	ipv6 rip enable
Parameter Description	N/A
Command Mode	Interface configuration mode
Usage Guide	The configuration for running the RIPng on an interface is different from that of RIPv2. In RIPv2, the network command is configured in routing process configuration mode to define an IP address range. If the IP address of an interface belongs to this IP address range, RIP automatically runs on this interface.

Enabling Split Horizon

Command	split-horizon [poisoned-reverse]
Parameter Description	poisoned-reverse: Indicates that the split horizon function contains the poison reverse function.

Command Mode	Routing process configuration mode
Usage Guide	Run the show ipv6 rip command to check whether split horizon is enabled. The configuration is different from that of RIPv2. In RIPv2, the split horizon function is configured in interface configuration mode.

Configuring a Passive Interface

Command	passive-interface { default <i>interface-type interface-num</i> }
Parameter Description	default: Indicates all interfaces. <i>interface-type interface-num</i>: Specifies an interface.
Command Mode	Routing process configuration mode
Usage Guide	First, run the passive-interface default command to configure all interfaces as passive interfaces. Then, run the no passive-interface <i>interface-type interface-num</i> command so that the interfaces used for interconnection between routers in the domain are not passive interface.

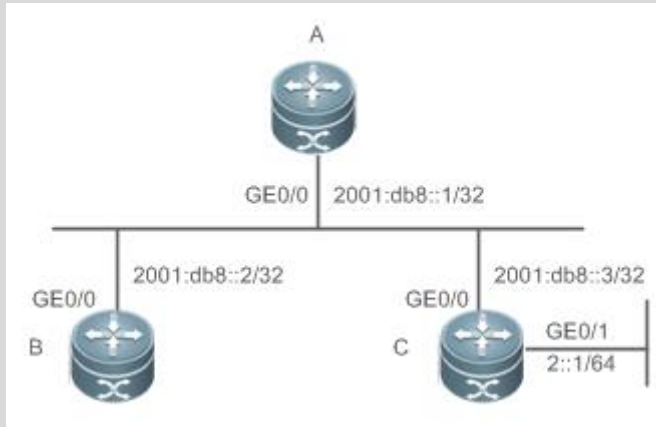
Displaying the IP Routing Table

Command	show ipv6 route
Parameter Description	N/A
Command Mode	Privileged EXEC mode or global configuration mode
Usage Guide	Check whether the routing table contains any route to a remote network that is learned through RIPng.

Configuration Example

Building a RIPng Routing Domain

Scenario
Figure 8-7



**Configurati
on Steps**

- Configure IPv6 addresses on all routers.
- Enable RIPng on all routers.

A

```
A# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
A(config)# ipv6 router rip
A(config-router)# exit
A(config)# interface GigabitEthernet 0/0
A(config-if-GigabitEthernet 0/0)# ipv6 address 2001:db8::1/32
A(config-if-GigabitEthernet 0/0)# ipv6 rip enable
```

B

```
B# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
B(config)# ipv6 router rip
B(config-router)# exit
B(config)# interface GigabitEthernet 0/0
B(config-if-GigabitEthernet 0/0)# ipv6 address 2001:db8::2/32
B(config-if-GigabitEthernet 0/0)# ipv6 rip enable
```

C

```
C# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
C(config)# ipv6 router rip
C(config-router)# exit
C(config)# interface GigabitEthernet 0/0
```

	<pre>C(config-if-GigabitEthernet 0/0)# C(config-if-GigabitEthernet 0/0)# ipv6 address 2001:db8::3/32 C(config-if-GigabitEthernet 0/0)# ipv6 rip enable C(config)# interface GigabitEthernet 0/1 C(config-if-GigabitEthernet 0/1)# ipv6 address 2::1/64 C(config-if-GigabitEthernet 0/1)# ipv6 rip enable</pre>
Verification	Check the routing tables on Router A, Router B, and Router C. The routing tables should contain routes to a remote network that are learned through RIPng.
A	<pre>A# show ipv6 route</pre> <p>IPv6 routing table name - Default - 6 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area</p> <pre>R 2::/64 [120/2] via FE80::2D0:F8FF:FEFB:D521, GigabitEthernet 0/0 C 2001:DB8::/32 via GigabitEthernet 0/0, directly connected L 2001:DB8::1/128 via GigabitEthernet 0/0, local host C FE80::/10 via ::1, Null0 C FE80::/64 via GigabitEthernet 0/0, directly connected L FE80::2D0:F8FF:FEFB:E7CE/128 via GigabitEthernet 0/0, local host</pre>
B	<pre>B# show ipv6 route</pre> <p>IPv6 routing table name - Default - 6 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2</p>

	<p>SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area</p> <p>R 2::/64 [120/2] via FE80::2D0:F8FF:FEFB:D521, GigabitEthernet 0/0 C 2001:DB8::/32 via GigabitEthernet 0/0, directly connected L 2001:DB8::2/128 via GigabitEthernet 0/0, local host C FE80::/64 via GigabitEthernet 0/0, directly connected L FE80::2D0:F8FF:FEFB:C9BA/128 via GigabitEthernet 0/0, local host</p>
C	<p>QTECH# show ipv6 route</p> <p>IPv6 routing table name - Default - 9 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area</p> <p>C 2::/64 via GigabitEthernet 0/1, directly connected L 2::2/128 via GigabitEthernet 0/1, local host C 2001:DB8::/32 via GigabitEthernet 0/0, directly connected L 2001:DB8::3/128 via GigabitEthernet 0/0, local host C FE80::/10 via ::1, Null0 C FE80::/64 via GigabitEthernet 0/0, directly connected L FE80::2D0:F8FF:FEFB:D521/128 via GigabitEthernet 0/0, local host C FE80::/64 via GigabitEthernet 0/1, directly connected L FE80::2D0:F8FF:FEFB:D521/128 via GigabitEthernet 0/1, local host</p>

Common Errors

- The IPv6 address is not configured on an interface.
- The interface used for interconnection between devices is configured as a passive interface.

8.4.2 Advertising the Default Route or External Routes

Configuration

Effect

- In the RIPng domain, introduce a unicast route of another AS so that the unicast routing service to this AS can be provided for users in the RIPng domain.
- In the RIPng domain, inject a default route to another AS so that the unicast routing service to this AS can be provided for users in the RIPng domain.

Notes

- The RIPng basic functions must be configured.

Configuration

Steps

Configuring External Route Redistribution

- Optional.
- Perform this configuration if external routes of the RIPng domain should be introduced to the AS border router (ASBR).

Generating a Default Route

- Optional.
- Perform this configuration if the default route should be introduced to an ASBR so that other routers in the RIPng domain access other AS domains through this ASBR by default.

Verification

- Run the **show ipv6 route rip** command on a non-ASBR to check whether the external routes of the domain and default route have been loaded.

Related

Commands

Advertising the Default Route to Neighbors on an Interface

Command	<code>ipv6 rip default-information { only originate } [metric <i>metric-value</i>]</code>
Parameter Description	only: Advertises only IPv6 default route. originate: Advertises the IPv6 default route and other routes. metric <i>metric-value</i>: Indicates the metric of the default route. The value ranges from 1 to 15. The default value is 1.

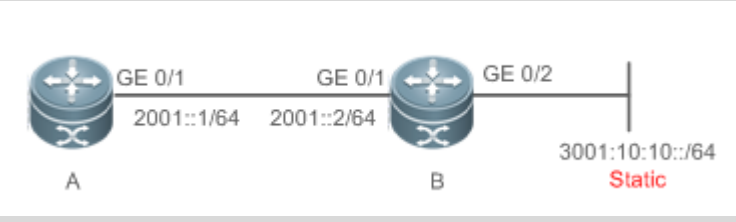
Command Mode	Interface configuration mode
Usage Guide	<p>After this command is configured on the interface, an IPv6 default route is advertised to the external devices through this interface, but the route itself is not added to the route forwarding table or the device and the RIPng route database.</p> <p>To prevent occurrence of a route loop, once this command is configured on an interface, RIPng refuses to receive the default route updates advertised by neighbors.</p>

Redistributing Routes and Advertising External Routes to Neighbors

Command	redistribute { bgp connected isis [<i>area-tag</i>] ospf <i>process-id</i> static } [metric <i>metric-value</i> route-map <i>route-map-name</i>]
Parameter Description	<p>bgp: Indicates redistribution from BGP.</p> <p>Connected: Indicates redistribution from direct routes.</p> <p>isis [<i>area-tag</i>]: Indicates redistribution from IS-IS. <i>area-tag</i> indicates the IS-IS process ID.</p> <p>ospf <i>process-id</i>: Indicates redistribution from OSPF. <i>process-id</i> indicates the OSPF process ID. The value ranges from 1 to 65535.</p> <p>static: Indicates redistribution from static routes.</p> <p>metric <i>metric-value</i>: Sets the metric of the route redistributed to the RIPng domain.</p> <p>route-map <i>route-map-name</i>: Sets the redistribution filtering rules.</p>
Command Mode	Routing process configuration mode
Usage Guide	<p>During route redistribution, it is not necessary to convert the metric of one routing protocol to the metric of another routing protocol because different routing protocols use completely different metric measurement methods. RIP measures the metric based on the hop count, and OSPF measures the metric based on the bandwidth. Therefore, the computed metrics cannot be compared with each other.</p>

Configuration

Example

Scenario Figure 8-8	
Configuration Steps	<ul style="list-style-type: none"> ▪ Configure the interface IPv6 addresses on all routers. (Omitted) ▪ Configure the RIPng basic functions on all routers. (Omitted) ▪ On Router B, configure redistribution of static routes. ▪ On the GE0/1 interface of Router A, configure advertisement of the default route.
A	<pre>A# configure terminal A(config)# interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)# ipv6 rip default-information originate</pre>
B	<pre>B# configure terminal B(config)# ipv6 router rip B(config-router)# redistribute static</pre>
Verification	<ul style="list-style-type: none"> ▪ Check the routing tables on Router A and Router B, and confirm that Router A can learn the route 3001:10:10::/64, and Router B can learn the default route ::/0.
A	<pre>A# show ipv6 route rip IPv6 routing table name - Default - 17 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area R 3001:10:10::/64 [120/2] via FE80::2D0:F8FF:FE22:334A, GigabitEthernet 0/1</pre>
B	<pre>B# show ipv6 route rip IPv6 routing table name - Default - 17 entries</pre>

```
Codes: C - Connected, L - Local, S - Static
R - RIP, O - OSPF, B - BGP, I - IS-IS
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
IA - Inter area
R  ::/0 [120/2] via FE80::21A:A9FF:FE41:5B06, GigabitEthernet 0/1
```

8.4.3 Setting Route Filtering Rules

Configuration

Effect

- Routes that do not meet filtering criteria cannot be loaded to the routing table, or advertised to neighbors. In this way, users within the network can be prevented from accessing specified destination networks.

Notes

- The RIPng basic functions must be configured.

Configuration

Steps

Filtering the Received RIP Routing Information

- To refuse receiving some specified routes, you can configure the route distribution control list to process all the received route update packets. If no interface is specified, route update packets received on all interfaces will be processed.

Filtering the Sent RIP Routing Information

- If this command does not contain any optional parameter, route update advertisement control takes effect on all interfaces. If the command contains the interface parameter, route update advertisement control takes effect only on the specified interface. If the command contains other routing process parameters, route update advertisement control takes effect only on the specified routing process.

Verification

- Run the **show ipv6 route rip** command to check that the routes that have been filtered out are not loaded to the routing table.

Related Commands

Command	distribute-list prefix-list <i>prefix-list-name</i> { in out } [<i>interface-type interface-name</i>]
Parameter Description	<p>prefix-list <i>prefix-list-name</i>: Indicates the name of the prefix list, which is used to filter routes.</p> <p>in out: Specifies update routes (received or sent routes) that are filtered.</p> <p>interface-type interface-name: Indicates that the distribution list is applied to the specified interface.</p>
Command Mode	Routing process configuration mode
Usage Guide	N/A

Configuration Example

Scenario Figure 8-9	
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IPv6 addresses on all routers. (Omitted) Configure the RIPng basic functions on all routers. (Omitted) On router A, configure route filtering.
A	<pre>A# configure terminal A(config)# ipv6 prefix-list hello permit 4001::/64 A(config)# ipv6 router rip A(config-router)# distribute-list prefix-list hello in</pre>
Verification	<ul style="list-style-type: none"> Check that Router A can learn only the route to 4001::/64.

A

```
A# show ipv6 route rip
IPv6 routing table name - Default - 17 entries
Codes: C - Connected, L - Local, S - Static
        R - RIP, O - OSPF, B - BGP, I - IS-IS
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
        IA - Inter area
R   4001::/64 [120/2] via FE80::2D0:F8FF:FE22:334A, GigabitEthernet 0/1
```

8.4.4 Modifying Route Selection Parameters

Configuration

Effect

- Change the RIPng routes to enable the traffic pass through specified nodes or avoid passing through specified nodes.
- Change the sequence that a router selects various types of routes so as to change the priorities of RIPng routes.

Notes

- The RIPng basic functions must be configured.

Configuration

Steps

Modifying the Administrative Distance of a RIPng Route

- Optional.
- Perform this configuration if you wish to change the priorities of RIPng routes on a router that runs multiple unicast routing protocols.

Modifying the Metric Offset on an Interface

- Optional.
- Unless otherwise required, perform this configuration on a router where the metrics of routes need to be adjusted.

Configuring the Default Metric of an External Route Redistributed to RIPng

- Optional.

- Unless otherwise required, perform this configuration on an ASBR to which external routes are introduced.

Verification

- Run the **show ipv6 rip** command to display the administrative distance of RIPng routes.
- Run the **show ipv6 rip data** command to display the metrics of external routes redistributed to RIPng.

Related

Commands

Modifying the Administrative Distance of a RIPng Route

Command	distance <i>distance</i>
Parameter Description	<i>distance</i> : Sets the administrative distance of a RIPng route. The value is an integer ranging from 1 to 254.
Command Mode	Routing process configuration mode
Usage Guide	Run this command to set the administrative distance of a RIPng route.

Modifying the Metric Offset on an Interface

Command	ipv6 rip metric-offset <i>value</i>
Parameter Description	<i>value</i> : Indicates the interface metric offset. The value ranges from 1 to 16.
Command Mode	Routing process configuration mode
Usage Guide	Before a route is added to the routing table, the metric of the route must be added with the metric offset set on the interface. You can control the use of a route by setting the interface metric offset.

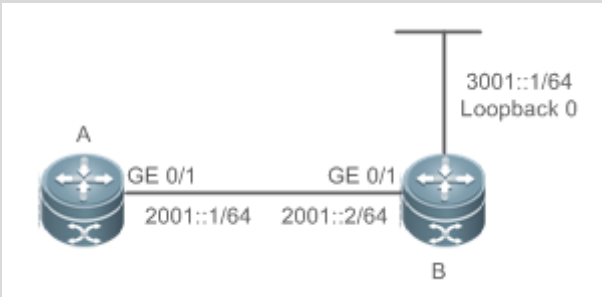
Configuring the Default Metric of an External Route Redistributed to RIPng

Command	default-metric <i>metric</i>
----------------	-------------------------------------

Parameter Description	<i>metric</i> : Indicates the default metric. The valid value ranges from 1 to 16. If the value is equal to or greater than 16, the RGOS determines that this route is unreachable.
Command Mode	Global configuration mode
Usage Guide	If the metric is not specified during redistribution of a routing protocol process, RIPng uses the metric defined by the default-metric command. If the metric is specified, the metric defined by the default-metric command is overwritten by the specified metric. If this command is not configured, the value of default-metric is 1.

Configuration Example

Modifying the Administrative Distance of a RIPng Route

Scenario Figure 8-10	
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IPv6 addresses on all routers. (Omitted) Configure the RIPng basic functions on all routers. (Omitted) On Router A, set the administrative distance of a RIPng route to 160.
	<pre>A# configure terminal A(config)# ipv6 router rip A(config-router)# distance 160</pre>
Verification	<ul style="list-style-type: none"> On Router A, check whether the administrative distance of a RIPng route is 160.
	<pre>A# show ipv6 route rip in 3001::/64 R 3001::/64 [160/2] via FE80::2D0:F8FF:FE22:334A, GigabitEthernet 0/1</pre>

8.4.5 Modifying Timers

Configuration

Effect

- Change the duration of RIPng timers to accelerate or slow down the change of the protocol state or occurrence of an event.

Notes

- The RIPng basic functions must be configured.
- Modifying the protocol control parameters may result in protocol running failures. Therefore, you are advised not to modify the timers.

Configuration

Steps

Modifying the Update Timer, Invalid Timer, and Flush Timer

- Mandatory.
- Unless otherwise required, perform this configuration on a router where RIPng timers need to be modified.

Verification

- Run the **show ipv6 rip** command to display settings of timers.

Related

Commands

Command	<code>timers update invalid flush</code>
Parameter Description	<p><i>Update:</i> Indicates the route update time in second. It defines the interval at which the device sends the route update packet. Each time an update packet is received, the invalid timer and flush timer are reset. By default, a route update packet is sent every 30s.</p> <p><i>Invalid:</i> Indicates the route invalid time in second, counted from the last time when a valid update packet is received. It defines the time after which the route in the routing list becomes invalid because the route is not updated. The duration of the invalid timer must be at least three times the duration of the update timer. If no update packet is received before the invalid timer expires, the corresponding route enters the invalid state. If the update packet is received before the invalid timer expires, the timer is reset. The default duration of the invalid timer is 180s.</p>

	<i>Flush</i> : Indicates the route flushing time in second, counted from the time when the RIPng route enters the invalid state. When the flush timer expires, the route in the invalid state will be deleted from the routing table. The default duration of the flush timer is 120s.
Command Mode	Routing process configuration mode
Usage Guide	By default, the update timer is 30s, the invalid timer is 180s, and the flush timer is 120s.

Configuration

Example

Scenario Figure 8-11	
Configuration Steps	<ul style="list-style-type: none"> Configure the interface IPv6 addresses on all routers. (Omitted) Configure the RIPng basic functions on all routers. (Omitted) On Router A, configure the update timer, invalid timer, and flush timer.
B	<pre>B# configure terminal B(config)# ipv6 router rip B(config-router)# timers 10 30 90</pre>
Verification	<ul style="list-style-type: none"> On Router B, check the settings of RIPng timers.
B	<pre>B# show ipv6 rip Routing Protocol is "RIPng" Sending updates every 10 seconds with +/-50%, next due in 12 seconds Timeout after 30 seconds, garbage collect after 90 seconds Outgoing update filter list for all interface is: not set Incoming update filter list for all interface is: not set</pre>

```
Default redistribution metric is 1
Default distance is 120
Redistribution:
  Redistributing protocol connected
Default version control: send version 1, receive version 1
Interface      Send  Recv
  GigabitEthernet 0/1  1    1
Routing Information Sources:
  Gateway: fe80::2d0:f8ff:fe22:334a Distance: 120
  Last Update: 00:00:02  Bad Packets: 0  Bad Routes: 0
```

Common Errors

- Settings of RIPng timers on devices connected to the same network are inconsistent. Consequently, routes cannot be learned properly.

8.5 Monitoring

Displaying

Description	Command
Displays information about the RIPng process.	show ipv6 rip
Displays the RIPng routing table.	show ipv6 rip database

Debugging

- ⚠ System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
-------------	---------

Debugs RIPng.

```
debug ipv6 rip [interface interface-type interface-num | nsm |  
restart ]
```

9 CONFIGURING PBR

9.1 Overview

Policy-based routing (PBR) is implemented by applying a route map including policies to interfaces and devices.

Similar to static routing, PBR is also manually configured and cannot automatically update with network changes. In addition, PBR is effective only for packets sent from local interfaces and devices. As compared with static and dynamic routing, PBR is more flexible. Static and dynamic routing can forward packets only based on destination addresses. PBR can forward packets based on source and destination addresses, packet length and input interface.

9.2 Applications


Application	Description
Selecting an ISP by Using PBR	Specify preferential output interfaces for packets from different subnets.
Implementing Traffic Classification by Using PBR	Specify QoS values for packets from different subnets.

9.2.1 Selecting an ISP by Using PBR

An existing user network often uses resources of multiple internet server providers (ISPs). PBR needs to be used since different bandwidths may be requested from different ISPs or the network resources for key users need to be protected. By controlling forwarding of certain data packets, you can make full use ISP resources as well as meet the requirements of flexible and diversified applications.

Scenario

Figure 9-1



A LAN has two output interfaces for connecting the Internet. PBR is configured on the layer-3 device DEV1 to enable the two output interfaces to implement load sharing and mutual backup.

The specific requirements are as follows:

- Data streams from subnet 1 are sent from GE 0/1.
- Data streams from subnet 2 are sent from GE 0/2.
- If the GE 0/1 link is disconnected, the data streams on GE 0/1 are switched to GE 0/2. Vice versa.

Deployment


- Configure two different ACLs on the layer-3 device DEV1:
ACL1: source addresses belong to subnet 1.
ACL2: source addresses belong to subnet 2.
- Configure two policies in the route map on the layer-3 device DEV1:
Policy 1: sets the next hops for packets matching ACL1 to GE0/1 and GE0/2 (Based on the configuration sequence, GE0/1 takes effect first and GE0/2 works in the backup mode).
Policy 2: sets the next hops for packets matching ACL2 to GE0/2 and GE0/1 (Based on the configuration sequence, GE0/2 takes effect first and GE0/1 works in the backup mode).
- Configure PBR on GE0/3 (by using a route map). Then, packets received on this interface are forwarded based on the policies.

9.2.2 Implementing Traffic Classification by Using PBR

Scenario

Networks of medium- and small-sized enterprises have simple structures. Different branch nodes are interconnected to the central nodes through carrier dedicated lines or the Internet VPN mode. Enterprise networks often need to implement three-in-one integration (of audio, video and data) to maximize the utilization of existing IP networks and save costs. Since all traffic is output from a single output interface, it is necessary to adjust the QoS policies for the output interface, in order to provide preferential communication quality for bandwidth- and delay-sensitive applications.

Figure 9-2



A LAN has an output interface for connecting the Internet. PBR is configured on the layer-3 device DEV1 to change the QoS values for packets from different networks.

The specific requirements are as follows:

For data streams from subnet 1, representing audio streams, set the DSCP value to 56.

For data streams from subnet 2, representing video streams, set the DSCP value to 40.

For data streams from subnet 3, representing data streams, set the DSCP value to 24.

Deployment

- Configure three different ACLs on the layer-3 device DEV1:

ACL1: source addresses belong to subnet 1.

ACL2: source addresses belong to subnet 2.

ACL3: source addresses belong to subnet 3.

- Configure three policies in the route map on the layer-3 device DEV1:

Policy 1: sets the DSCP value for packets matching ACL1 to 56.

Policy 2: sets the DSCP value for packets matching ACL2 to 40.

Policy 3: sets the DSCP value for packets matching ACL3 to 24.

- Configure PBR on GE0/3 (by using a route map). Then, the DSCP values for packets received on this interface are changed based on the policies.

9.3 Features

Feature	Description
Configuring a Policy	Before configuring PBR, configure policies in a route map.
Configuring PBR	Apply a route map including policies to interfaces and devices to implement PBR.

9.3.1 Configuring a Policy

A policy is a "match ..., set..." statement, which indicates that "if certain conditions are matched, perform certain processing actions".

i For detailed introduction to the policies, see the section "Route Map".

Executing Policies

In the global configuration mode, you can run the **route-map** *route-map-name* [**permit** | **deny**] [*sequence-number*] command to create a policy in a route map.

A route map may contain multiple policies. Each policy has a corresponding sequence number. A smaller sequence number means a higher priority. Policies are executed based on their sequence numbers. Once the matching condition of a policy is met, the processing action for this policy needs to be executed and the route map exits. If no matching condition of any policy is met, no processing action will be performed.

Policies have two working modes:

- **permit**: When the matching condition of a policy is met, perform the processing action for this policy and exit the route map.
- **deny**: When the matching condition of a policy is met, do not perform the processing action for this policy and exit the route map.



Matching conditions of policies





The matching conditions of a policy may contain 0, 1 or more matching rules.

- If 0 matching rule is contained, no packet will be matched.
- If one or more match rules are contained, all match rules must be matched at the same time to meet the matching conditions of the policy.

In the route map mode, run the **match** command to configure match rules. One **match** command is mapped to one match rule.

PBR supports the following **match** commands:

	Command	Description
	match ip address	The source IPv4 address (and the destination IPv4 address) is used as the matching condition.  Multiple match ip address commands can be configured in a policy.
IPv6 PBR	match ipv6 address	The source IPv6 address (and the destination IPv6 address) is used as the matching condition.  Only one match ipv6 policy command can be configured in a policy.

-  IPv4 PBR defines the source IP address (and destination IP address) ranges of packets by using the IP standard or extended ACLs. IPv6 PBR defines the source IPv6 address (and destination IPv6 address) ranges of packets by using the IPv6 extended ACLs.
-  Packet forwarding based on policies of IPv4 PBR interfaces supports expert-level and MAC name ACLs. Packet forwarding based on local policies does not support expert-level and MAC name ACLs.
-  When PBR uses an ACL that is unavailable, the route sub-map will not be matched and the next route sub-map will be matched instead. If no route sub-map is matched, a common route will be selected for forwarding. If only ACLs are configured but no ACE is configured, the PBR forwarding behavior is the same as that in a scenario where an ACL is unavailable
-  On a switch, if a route sub-map uses multiple ACLs in PBR, only the first ACL is matched.

Processing action for a policy













The processing action of a policy may contain 0, 1 or more set rules.

- If 0 set rule is contained, no processing action will be performed and the route map will directly exit.
- If one or more set rules are contained, all processing actions will be performed and the route map will exit.

 If set rules have different priorities, the set rule with the highest priority will take effect.

In the route map mode, run the **set** command to configure set rules. One **set** command is mapped to one set rule.

PBR supports the following **set** commands:

	Command	Description
IPv4 PBR	set ip tos	<p>Modifies the tos field of an IPv4 packet.</p> <p> Command set ip and dscp cannot work with each other.</p>
	set ip precedence	<p>Modifies the precedence field of an IPv4 packet.</p> <p> Command set ip and dscp cannot work with each other.</p>
	set ip dscp	<p>Modifies the dscp field of an IPv4 packet.</p> <p> Command set ip tos and precedence cannot work with each other.</p>
	 set ip next-hop	<ul style="list-style-type: none">  Configures the next hop of IPv4 packet forwarding. The next hop must be directly connected; otherwise, this command is invalid.  A packet matching the match rules will be forwarded to the next hop specified by set ip next-hop first, no matter whether the route selected for the packet in the routing table is consistent with the next hop specified by PBR.  On a switch, the output interfaces for next hops supported by PBR include the SVI, routing and layer-3 AP interfaces.
	 set ip next-hop recursive	<ul style="list-style-type: none">  Configures the recursive next hop of IPv4 packet forwarding. The next hop can be directly connected or not directly connected. A non-directly-connected next hop will recur to a static or dynamic route in the routing table.  This command supports recursion to multiple ECMP next hops of a static or dynamic route. A maximum of 32 next hops are supported. If a recursive route is a static route, only one next hop is supported for the static recursive route.  The redundant backup or load balancing mode of multiple recursive next hops is also determined by the ip policy { redundance load-balance } command.  A packet matching the match rules will be forwarded to the recursive next hop specified by set ip next-hop

		<p>recursive first, no matter whether the route selected for the packet in the routing table is consistent with the next hop specified by PBR.</p> <ul style="list-style-type: none"> Only when a static or dynamic route has an output interface and a next-hop IP address, the policy-based recursive next hop can take effect.
	<ul style="list-style-type: none"> set ip default next-hop 	<ul style="list-style-type: none"> Configures the default next hop of IPv4 packet forwarding. A packet matching the match rules will be forwarded to the default next hop specified by this command if a route fails to be selected or the default route is selected for this packet in the routing table.
<ul style="list-style-type: none"> IPv6 PBR 	<ul style="list-style-type: none"> set ipv6 precedence 	<ul style="list-style-type: none"> Modifies the precedence field of an IPv6 packet. IPv6 PBR does not support set ipv6 tos or set ipv6 dscp.
	<ul style="list-style-type: none"> set ipv6 next-hop 	<ul style="list-style-type: none"> Configures the next hop of IPv6 packet forwarding. An IPv6 packet matching the match rules will be forwarded to the next hop specified by set ipv6 next-hop first, no matter whether the route selected for the IPv6 packet in the routing table is consistent with the next hop specified by PBR. The next hop must be directly connected; otherwise, this command is invalid.
	<ul style="list-style-type: none"> set ipv6 default next-hop 	<ul style="list-style-type: none"> Configures the default next hop of IPv6 packet forwarding. An IPv6 packet matching the match rules will be forwarded to the default next hop specified by this command if a route fails to be selected or the default route is selected for this packet in the routing table. The next hop must be directly connected; otherwise, this command is invalid.

i The priority sequence is as follows: **set ip next-hop** > **set ip next-hop recursive** > common route > **set ip default next-hop** > default route. The preceding **set** commands can be configured at the same time but only the command with the highest priority takes effect.

- i** The priority sequence is as follows: **set ipv6 next-hop** > common route > **set ipv6 default next-hop** > default route. The preceding set commands can be configured at the same time but only the command with the highest priority takes effect.
- ✓** For switches, the **set ipv6 default next-hop** command does not take effect for IPv6 addresses whose mask length exceeds 64.

9.3.2 Configuring PBR

PBR

Apply a route map including policies to interfaces or devices to implement PBR.

- Apply a route map to an interface so that packets received by the interface are routed based on the policy.
The PBR is often used to control user packets received by a device. This command is effective only for forwarded packets, but not for locally initiated packets.
 - Apply a route map to a device so that packets locally initiated are routed based on the policy.
The PBR is often used to control protocol packets exchanged between devices (such as ping packets sent locally). This command is effective only for locally initiated packets, but not for forwarded packets.
-
- i** By default, PBR is not available on a device and packets are forwarded based on a routing table.
 - i** On a switch, the interfaces which support PBR are L3 Ethernet interface, SVI interface and L3 AP interface.

Redundant backup or load balancing

You can set multiple next hops in a policy. Either redundant backup or load balancing can be implemented among multiple next hops. Redundant backup is implemented by default.

- i** Redundant backup or load balancing is only effective for next hops configured in the **set ip next-hop**, **set ip next-hop recursive**, **set ip default next-hop**, **set ipv6 next-hop** and **set ipv6 default next-hop** commands, and only effective among multiple next hops in the same set rule.
- Redundant backup
Based on the configuration sequence, the first accessible next hop takes effect. When the currently effective next hop (R1) is faulty, the traffic automatically switches to the next accessible next hop (R2). When R1 becomes accessible again, the traffic automatically switches back to R1. A newly added next hop is arranged at the last of the sequence. Assume that the original sequence of multiple next hops is R1 > R2 > R3. After R1 is deleted and added again, the sequence changes to R2 > R3 > R1.
If no next hop is accessible, packets will be discarded.

Correlation with BFD

Correlation between PBR and BFD is effective only for next hops configured by the **set ip next-hop** or **set ipv6 next-hop** command.

The **set ip next-hop** and **set ipv6 next-hop** commands carry the **verify-availability** and **bfd [vrf vrf-name] interface-type interface-number gateway** parameters, which can establish correlation between PBR and a BFD session and monitor the accessibility of next hops.

Correlation between PBR and BFD helps enhance the PBR's perception about network environment changes. When BFD detects that the current next hop is not accessible, the BFD will immediately notify the PBR to switch the traffic to another accessible next hop (to implement redundant backup) or all the other accessible next hops (to implement load balancing).

- i** For the configuration and related commands for correlation between PBR and BFD, see the "BFD" section.

Correlation with Track

Correlation between PBR and Track is effective only for next hops configured by the **set ip next-hop** command.

The **set ip next-hop** command carries the **verify-availability** and **track track-obj-number** parameters, which can establish correlation between PBR and a Track session and monitor the accessibility of next hops.

Correlation between PBR and Track helps enhance the PBR's perception about network environment changes. When Track detects that the current next hop is not accessible, the Track will immediately notify the PBR to switch the traffic to another accessible next hop (to implement redundant backup) or all the other accessible next hops (to implement load balancing).

- i** Only IPv4 PBR supports correlation with Track.
- i** For the configuration and related commands for correlation between PBR and Track, see the "RNS" section.

Source-addressed-based PBR

Run the global configuration commands **ip policy-source in-interface** and **ipv6 policy-source in-interface** to directly generate source-address-based PBR. You do not need to configure a route map.

- If only source IPv4 or Ipv6 addresses need to be matched for packets forwarded by an interface, you can apply the command for source-address-based PBR. The procedure for configuring this command is simpler than that for interface-based PBR.
- If source-address-based PBR is applied to a specified interface, packets received on this interface will be routed based on policies according to the source addresses.

The PBR is often used to control user packets received by a device. This command is effective only for forwarded packets, but not for locally initiated packets.

- i** Source-address-based PBR has a higher priority than interface-based PBR. If source-address-based PBR and interface-based PBR are applied to the same interface, only interface-based PBR takes effect.
- i** By default, source-address-based routing is not available on devices and packets are forwarded based on the routing table.



9.4 Limitations



For QSW-6900 series products:

- If PBR refers ACL which does not exist, the route submap is not matched but match the next route submap. When all route submaps do not match, it is forwarded through common route. It is the same if only ACL is configured but the ACE is not.
- On a switch, if a route submap of PBR refers to multiple ACLs, it matches only the first ACL and does not match the subsequent ACL.

9.5 Configuration

Configuration	Description and Command
Configuring Basic Functions of PBR	 (Mandatory) It is used to apply PBR to forward packets.
	ip policy route-map Applies PBR for IPv4 packets received by an interface.
	ipv6 policy route-map Applies PBR for IPv6 packets received by an interface.
	ip local policy route-map Applies PBR for IPv4 packets locally initiated.
	ipv6 local policy route-map Applies PBR for IPv6 packets locally initiated.
Setting Redundant Backup or Load Balancing	 (Optional) It is used to set whether PBR implements redundant backup or load balancing among multiple next hops.
	ip policy { redundance load-balance } Sets whether IPv4 PBR implements redundant backup or

		load balancing among multiple next hops. The default setting is redundant backup.
	ipv6 policy { redundance load-balance }	Sets whether IPv6 PBR implements redundant backup or load balancing among multiple next hops. The default setting is redundant backup.
Configuring Source-Address-Based PBR	⚠ (Optional) It is used to apply source-address-based PBR to forward packets.	
	ip policy-source in-interface	Applies source-address-based PBR for IPv4 packets received by an interface.
	Ipv6 policy-source in-interface	Applies source-address-based PBR for IPv6 packets received by an interface.

9.5.1 Configuring Basic Functions of PBR

Configuration

Effect

Perform personalized routing management for user data streams by preparing flexible policies.

Perform personalized management for protocol interaction and network topologies by preparing flexible policies.

Notes

- A route map must be used when PBR is configured; therefore, you must configure a route map on a device.
- If an ACL is used when the route map is configured, you must configure the ACL on the device.

Configuration

Steps

Applying PBR for IPv4 packets received by an interface

- To perform personalized routing management for IPv4 user data streams passing a device, you should perform this configuration.

- Perform this configuration on the input interface for user data streams.
- Run the **ip policy route-map** command to apply a route map to an interface. Then, PBR is executed for IPv4 packets received on this interface.

Command	ip policy route-map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Defaults	By default, PBR is unavailable on a device and packets are forwarded based on a routing table.
Command Mode	Interface configuration mode
Usage Guide	<p>Only one ip policy route-map command can be configured for an interface. If multiple ip policy route-map commands are configured for an interface, only the last configuration takes effect.</p> <p>If the route map used in PBR is unavailable, the PBR does not take effect.</p> <p>It will not take effect, if the PBR is configured on DAILER, loopback or tunnel interface.</p>

Applying PBR for IPv6 packets received by an interface

- To perform personalized routing management for IPv6 user data streams passing a device, you should perform this configuration.
- Perform this configuration on the input interface for user data streams.
- Run the **ipv6 policy route-map** command to apply a route map to an interface. Then, PBR is executed for IPv6 packets received on this interface.

Command	ipv6 policy route-map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Defaults	By default, PBR is unavailable on a device and packets are forwarded based on a routing table.
Command Mode	Interface configuration mode

Usage Guide	<p>Only one ipv6 policy route-map command can be configured for an interface. If multiple ipv6 policy route-map commands are configured for an interface, only the last configuration takes effect.</p> <p>If the route map used in PBR is unavailable, the PBR does not take effect.</p> <p>It will not take effect, if the PBR is configured on DAILER, loopback or tunnel interface.</p>
-------------	---

Applying PBR for IPv4 packets locally initiated

- To perform personalized management for IPv4 protocol interaction and IPv4 network topologies, you should perform this configuration.
- Run the **ip local policy route-map** command to apply a route map to a device. Then, PBR is executed for IPv4 packets locally initiated.

Command	ip local policy route-map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Defaults	By default, PBR is unavailable on a device and packets are forwarded based on a routing table.
Command Mode	Global configuration mode
Usage Guide	<p>Only one ip local policy route-map command can be configured for a device.</p> <p>If the route map used in PBR is unavailable, the PBR does not take effect.</p>

Applying PBR for IPv6 packets locally initiated

- To perform personalized management for IPv6 protocol interaction and IPv6 network topologies, you should perform this configuration.
- Run the **ipv6 local policy route-map** command to apply a route map to a device. Then, PBR is executed for IPv6 packets locally initiated.

Command	ipv6 local policy route-map <i>route-map-name</i>
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Defaults	By default, PBR is unavailable on a device and packets are forwarded based on a routing table.

Command Mode	Global configuration mode
Usage Guide	Only one ipv6 local policy route-map command can be configured for a device. If the route map used in PBR is unavailable, the PBR does not take effect.

Verification

- Check the configurations of PBR.
- Check the configurations of the route map used by PBR.
- If an ACL is used when the route map is configured, you should check the configurations of the ACL.

Checking the configurations of IPv4 PBR

Command	show ip policy [<i>route-map-name</i>]
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<p>Check the interfaces configured with IPv4 PBR according to the output information and the name of the used route map.</p> <pre> QTECH# show ip policy Banlance mode: redundance Interface Route map local RM_for_PBR_1 GigabitEthernet 0/1 RM_for_PBR_2 </pre> <p>Local indicates applying policy-based routing for IPv4 packets locally initiated.</p>

Checking the configurations of IPv6 PBR

Command	show ipv6 policy [<i>route-map-name</i>]
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Command Mode	Privilege, global and interface configuration modes

Usage Guide	<p>Check the interfaces configured with IPv6 PBR according to the output information and the name of the used route map.</p> <pre>QTECH#show ipv6 policy Banalance mode: redundance Interface Route map local RM_for_PBR_1 VLAN 1 RM_for_PBR_2</pre> <p>Local indicates applying policy-based routing for IPv6 packets locally initiated.</p>
-------------	---

Checking the configurations of a route map

Command	show route-map [<i>route-map-name</i>]
Parameter Description	<i>route-map-name</i> : Indicates the name of a route map.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<p>Multiple route maps may be available on a device. Focus on the route map used in PBR and check its policy settings.</p> <pre>QTECH# show route-map route-map RM_FOR_PBR, permit, sequence 10 Match clauses: ip address acl1 Set clauses: ip next-hop 200.24.18.1 route-map RM_FOR_PBR, permit, sequence 20 Match clauses: ip address acl2 Set clauses: ip next-hop 200.24.19.1</pre>

Checking the configurations of an ACL

Command	show access-lists [<i>acl-id</i> <i>acl-name</i>]
---------	--

Parameter Description	<i>acl-id</i> : Indicates the ACL ID. <i>acl-name</i> : Indicates the ACL name.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	Multiple ACLs may be available on a device. Focus on the ACL used by a route map and check its configurations. <pre>QTECH# show access-lists 1 ip access-list standard 1 10 permit 200.24.16.0 0.0.0.255 ip access-list standard 2 10 permit 200.24.17.0 0.0.0.255</pre>

Checking the routing information of IPv4 PBR

Command	show ip pbr route [interface <i>if-name</i> local]
Parameter Description	<i>if-name</i> : Indicates an interface name. local : Indicates local.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	Specify a local interface or device and check the routing information of IPv4 PBR. <pre>QTECH# show ip pbr route PBR IPv4 Route Summay : 1 Interface : GigabitEthernet 0/1 Sequence : 10 Min Length : None Max Length : None VRF ID : 0 Route Flags : Route Type : PBR Direct : Permit Priority : High</pre>

```

Tos_Dscp   : None
Precedence : None
Tos_Dscp   : 0
Precedence : 0
Mode       : redundance
Nextthop Count : 1
Nextthop[0] : 192.168.8.100
Weight[0]   : 1
Ifindex[0]  : 2

```

Checking the routing information of IPv6 PBR

Command	show ipv6 pbr route [interface <i>if-name</i> local]
Parameter Description	<i>if-name</i> : Indicates an interface name. local : Indicates local.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	Specify a local interface or device and check the routing information of IPv6 PBR. <pre> QTECH# show ipv6 pbr route PBR IPv6 Route Summary : 1 Interface : GigabitEthernet 0/1 Sequence : 10 ACL[0] : 2900 ACL_CLS[0] : 5 Min Length : None Max Length : None VRF ID : 0 Route Flags : Route Type : PBR Direct : Permit Priority : High Tos_Dscp : None </pre>

```

Precedence : None
Tos_Dscp : 0
Precedence : 0
Mode : redundance
Nexthop Count : 1
Nexthop[0] : 10::2
Weight[0] : 1
Ifindex[0] : 2

```

Checking a route map used by IPv4 PBR

Command	show ip pbr route-map <i>rmap-name</i>
Parameter Description	<i>rmap-name</i> : Indicates the route map name.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<p>Specify a route map and check the route map used by IPv4 PBR.</p> <pre> QTECH# show ip pbr route-map rm PBR VRF: GLOBAL, ID: 0 Forward Mode: redundance Forwarding: On Route-map rm Route-map index: Sequence 10, permit Match rule: ACL ID : 2900, CLS: 1, Name: acl1 Set rule: IPv4 nexthop: 192.168.8.100, (VRF name: , ID: 0), Weight: 0 PBR state info ifx: 2, Connected: True, Track state: Up </pre>

Checking a route map used by IPv6 PBR

Command	show ipv6 pbr route-map <i>rmap-name</i>
----------------	---

Parameter Description	<i>rmap-name</i> : Indicates the route map name.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<p>Specify a route map and check the route map used by IPv6 PBR.</p> <pre> QTECH# show ipv6 pbr route-map rm6 PBR VRF: GLOBAL, ID: 0 Forward Mode: redundance Forwarding: On Route-map rm6 Route-map index: Sequence 10, permit Match rule: ACL ID : 2901, CLS: 5, Name: acl6 Set rule: IPv6 nexthop: 10::2, (VRF name: , ID: 0), Weight: 0 PBR state info ifx: 2, Connected: True, Track state: Up </pre>

Checking the statistics about packets forwarded by IPv4 PBR

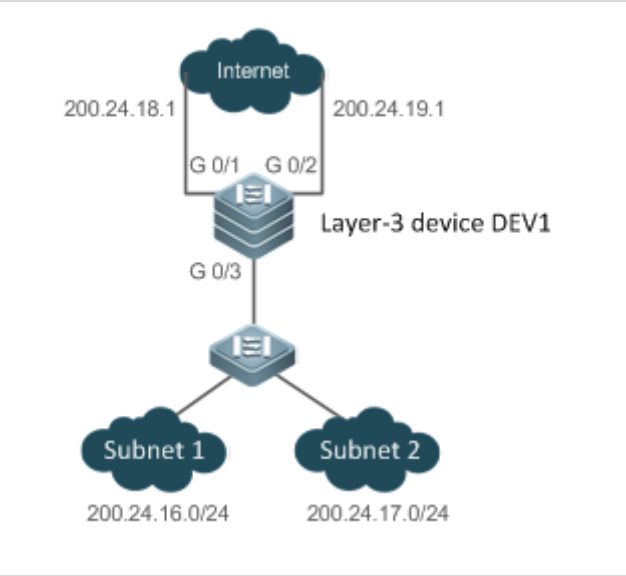
Command	show ip pbr statistics [interface <i>if-name</i> local]
Parameter Description	<p><i>if-name</i>: Indicates an interface name.</p> <p>local: Indicates local.</p>
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<pre> QTECH# show ip pbr statistics IPv4 Policy-based route statistic gigabitEthernet 0/1 statistics : 10 </pre>


Checking the statistics about packets forwarded by IPv6 PBR

Command	show ipv6 pbr statistics [interface <i>if-name</i> local]
Parameter Description	<i>if-name</i> : Indicates an interface name. local : Indicates local.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<pre> QTECH# show ipv6 pbr statistics IPv6 Policy-based route statistic gigabitEthernet 0/1 statistics : 20 </pre>

Configuration Example

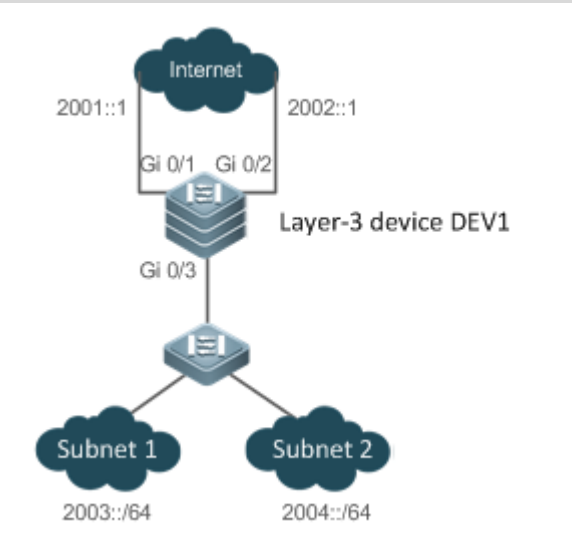
Configuring IPv4 PBR and selecting an output link based on source addresses of packets

<p>Scenario Figure 9-3</p>	 <p>The diagram illustrates a network topology. At the top, a cloud labeled 'Internet' is connected to a 'Layer-3 device DEV1' through two interfaces: G 0/1 with IP address 200.24.18.1 and G 0/2 with IP address 200.24.19.1. Below DEV1, another interface G 0/3 is connected to a second 'Layer-3 device'. This second device is then connected to two separate clouds representing 'Subnet 1' (200.24.16.0/24) and 'Subnet 2' (200.24.17.0/24).</p>
	<p>The layer-3 device DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 200.24.16.0/24 whereas the network segment where subnet 2 resides is 200.24.17.0/24.</p> <p>DEV 1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 200.24.18.1 and 200.24.19.1.</p>

	<p>This LAN has two output interfaces for connecting the Internet. The requirements are as follows:</p> <ul style="list-style-type: none"> ▪ Data streams from subnet 1 for accessing the Internet should pass GE 0/1. ▪ Data streams from subnet 2 for accessing the Internet should pass GE 0/2. ▪ If the GE 0/1 link is disconnected, the data streams on the GE 0/1 interface are switched to the GE 0/2 interface. Vice versa.
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure two ACLs to match packets from subnets 1 and 2 respectively. ▪ Set a policy to set the next hops for packets from subnet 1 to GE0/1 and GE0/2. (Pay attention to the configuration sequence.) ▪ Set a policy to set the next hops for packets from subnet 2 to GE0/2 and GE0/1. (Pay attention to the configuration sequence.) ▪ Apply the policy to GE 0/3. ▪ Set PBR to implement redundant backup among multiple next hops. (The default setting is redundant backup.) <hr/> <p> During redundant backup, based on the configuration sequence, the first next hop takes effect first.</p>
	<pre> DEV1(config)# access-list 1 permit 200.24.16.0 0.0.0.255 DEV1(config)# access-list 2 permit 200.24.17.0 0.0.0.255 DEV1(config)# route-map RM_FOR_PBR 10 DEV1(config-route-map)# match ip address 1 DEV1(config-route-map)# set ip next-hop 200.24.18.1 DEV1(config-route-map)# set ip next-hop 200.24.19.1 DEV1(config-route-map)# exit DEV1(config)# route-map RM_FOR_PBR 20 DEV1(config-route-map)# match ip address 2 DEV1(config-route-map)# set ip next-hop 200.24.19.1 DEV1(config-route-map)# set ip next-hop 200.24.18.1 DEV1(config-route-map)# exit DEV1(config)# interface GigabitEthernet 0/3 DEV1(config-if-GigabitEthernet 0/3)# ip policy route-map RM_FOR_PBR DEV1(config-if-GigabitEthernet 0/3)# exit DEV1(config)# ip policy redundancy </pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Check the configurations of IPv4 PBR. ▪ Check the configurations of the route map.

	<ul style="list-style-type: none"> Check the configurations of an ACL.
	<pre>DEV1# show ip policy Interface Route map GigabitEthernet 0/3 RM_FOR_PBR</pre>
	<pre>DEV1# show route-map route-map RM_FOR_PBR, permit, sequence 10 Match clauses: ip address 1 Set clauses: ip next-hop 200.24.18.1 200.24.19.1 route-map RM_FOR_PBR, permit, sequence 20 Match clauses: ip address 2 Set clauses: ip next-hop 200.24.19.1 200.24.18.1</pre>
	<pre>DEV1# show access-lists ip access-list standard 1 10 permit 200.24.16.0 0.0.0.255 ip access-list standard 2 10 permit 200.24.17.0 0.0.0.255</pre>

Configuring IPv6 PBR and selecting an output link based on source addresses of packets

<p>Scenario Figure 9-4</p>	
	<p>DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 2003::/64 whereas the network segment where subnet 2 resides is 2004::/64.</p> <p>DEV1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 2001::1/64 and 2002::1/64.</p>
	<p>This LAN has two output interfaces for connecting the Internet. The requirements are as follows:</p> <ul style="list-style-type: none"> ▪ Data streams from subnet 1 for accessing the Internet should pass GE 0/1. ▪ Data streams from subnet 2 for accessing the Internet should pass GE 0/2. ▪ If the GE 0/1 link is disconnected, the data streams on the GE 0/1 interface are switched to the GE 0/2 interface. Vice versa.
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure two ACLs to match packets from subnets 1 and 2 respectively. ▪ Set a policy to set the next hops for packets from subnet 1 to GE0/1 and GE0/2. (Pay attention to the configuration sequence.) ▪ Set a policy to set the next hops for packets from subnet 2 to GE0/2 and GE0/1. (Pay attention to the configuration sequence.) ▪ Apply the policy to GE 0/3. ▪ Set PBR to implement redundant backup among multiple next hops. <p>i During redundant backup, based on the configuration sequence, the first next hop takes effect first.</p>
	<pre>DEV1(config)# ipv6 access-list net1</pre>

	<pre> DEV1(config-ipv6-acl)# permit ipv6 2003::/64 any DEV1(config-ipv6-acl)# exit DEV1(config)# ipv6 access-list net2 DEV1(config-ipv6-acl)# permit ipv6 2004::/64 any DEV1(config-ipv6-acl)# exit DEV1(config)# route-map RM_FOR_PBR 30 DEV1(config-route-map)# match ipv6 address net1 DEV1(config-route-map)# set ipv6 next-hop 2001::1 DEV1(config-route-map)# set ipv6 next-hop 2002::1 DEV1(config-route-map)# exit DEV1(config)# route-map RM_FOR_PBR 40 DEV1(config-route-map)# match ipv6 address net2 DEV1(config-route-map)# set ipv6 next-hop 2002::1 DEV1(config-route-map)# set ipv6 next-hop 2001::1 DEV1(config-route-map)# exit DEV1(config)# interface GigabitEthernet 0/3 DEV1(config-if-GigabitEthernet 0/3)# ipv6 policy route-map RM_FOR_PBR DEV1(config-if-GigabitEthernet 0/3)# exit DEV1(config)# ipv6 policy redundance </pre>
Verification	<ul style="list-style-type: none"> ▪ Check the configurations of IPv6 PBR. ▪ Check the configurations of the route map. ▪ Check the configurations of an ACL.
	<pre> DEV1# show ipv6 policy Interface Route map GigabitEthernet 0/3 RM_FOR_PBR </pre>
	<pre> DEV1# show route-map route-map RM_FOR_PBR, permit, sequence 11 Match clauses: ipv6 address net1 Set clauses: </pre>

```

ipv6 next-hop 2001::1 2002::1
route-map RM_FOR_PBR, permit, sequence 21
Match clauses:
  ipv6 address net2
Set clauses:
  ipv6 next-hop 2002::1 2001::1

```


```

DEV1# show access-lists
ipv6 access-list net1
10 permit ipv6 2003::/64 any
(0 packets matched)
ipv6 access-list net2
10 permit ipv6 2004::/64 any
(0 packets matched)

```

Configuring correlation between IPv4 PBR and Track

Scenario Figure 9-5



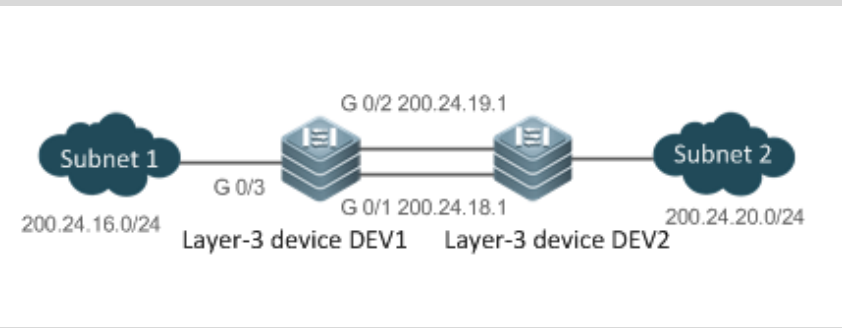
The layer-3 device DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 200.24.16.0/24 whereas the network segment where subnet 2 resides is 200.24.17.0/24.

DEV 1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 200.24.18.1 and 200.24.19.1.

	<ul style="list-style-type: none"> DEV1 can fast detect a faulty output link and switch to a backup link.
Configurati on Steps	<p>When configuring IPv4 PBR and selecting an output link based on source addresses of the packets, add or modify the following configurations (red fields):</p> <ul style="list-style-type: none"> Set two Track objects and track the accessibility of the next hops of the two output interfaces. When configuring a policy, set the correlation between the next hops and the Track objects.
DEV1	<pre> DEV1(config)# ip access-list extended 101 DEV1(config-ip-acl)# permit ip 200.24.16.0 0.0.0.255 any DEV1(config-ip-acl)# exit DEV1(config)# ip access-list extended 102 DEV1(config-ip-acl)# permit ip 200.24.17.0 0.0.0.255 any DEV1(config-ip-acl)# exit DEV1(config)#ip rns 1 DEV1(config-ip-rns)#icmp-echo 200.24.18.1 DEV1(config)#ip rns schedule 1 start-time now life forever DEV1(config)#track 1 rns 1 DEV1(config)#ip rns 2 DEV1(config-ip-rns)#icmp-echo 200.24.19.1 DEV1(config)#ip rns schedule 2 start-time now life forever DEV1(config)#track 2 rns 2 DEV1(config)# route-map RM_FOR_PBR 10 DEV1(config-route-map)# match ip address 101 DEV1(config-route-map)# set ip next-hop verify-availability 200.24.18.1 track 1 DEV1(config-route-map)# set ip next-hop verify-availability 200.24.19.1 track 2 DEV1(config-route-map)# exit DEV1(config)# route-map RM_FOR_PBR 20 DEV1(config-route-map)# match ip address 102 DEV1(config-route-map)# set ip next-hop verify-availability 200.24.19.1 track 2 DEV1(config-route-map)# set ip next-hop verify-availability 200.24.18.1 track 1 DEV1(config-route-map)# exit DEV1(config)# interface GigabitEthernet 0/3 </pre>

	<pre>DEV1(config-if-GigabitEthernet 0/3)# ip policy route-map RM_FOR_PBR DEV1(config-if-GigabitEthernet 0/3)# exit DEV1(config)# ip policy redundancy</pre>
Verification	<ul style="list-style-type: none"> Check whether the Track objects are up.
DEV1	<pre>QTECH#show track Track 1 Reliable Network Service 1 The state is Up 1 change, current state last: 120 secs Delay up 30 secs, down 50 secs Track 2 Reliable Network Service 2 The state is Up 1 change, current state last: 130 secs Delay up 30 secs, down 50 secs</pre>

Configuring IPv4 recursive PBR, selecting an output link based on source addresses of the packets, and recurring to the output link of a dynamic route

<p>Scenario Figure 9-6</p>	
	<p>The layer-3 device DEV 1 is connected to subnet 1 through GE0/3. The network segment where subnet 1 resides is 200.24.16.0/24.</p> <p>DEV 1 is connected to subnet 2 through GE0/1 and GE0/2 and their next hops are 200.24.18.1 and 200.24.19.1.</p>
	<p>Subnet 1 is connected to subnet 2 through two output interfaces of DEV1. The requirements are as follows:</p>

	<ul style="list-style-type: none"> Configure static or dynamic routes in advance to ensure that static or dynamic routes in the network segment 200.24.20.0 are available in the routing table of DEV1. Data streams from subnet 1 for accessing the Internet can recur to a dynamic route whose IP address is 200.24.20.1. If the GE 0/1 link is disconnected, the data streams on GE 0/1 are switched to GE 0/2. Vice versa.
Configurati on Steps	<ul style="list-style-type: none"> Configure an ACL to match packets from subnet 1. Set a policy to set the recursive next hop for packets from subnet 1 to 200.24.20.1. Apply the policy to GE 0/3. Set PBR to implement redundant backup among multiple next hops. (The default setting is redundant backup.) <p>i During redundant backup, the sequence for the next hops to take effect is related to the sequence for the static or dynamic routes to take effect.</p>
	<pre>DEV1(config)# access-list 1 permit 200.24.16.0 0.0.0.255 DEV1(config)# route-map RM_FOR_PBR 10 DEV1(config-route-map)# match ip address 1 DEV1(config-route-map)# set ip next-hop recursive 200.24.20.1 DEV1(config-route-map)# exit DEV1(config)# interface GigabitEthernet 0/3 DEV1(config-if-GigabitEthernet 0/3)# ip policy route-map RM_FOR_PBR DEV1(config-if-GigabitEthernet 0/3)# exit DEV1(config)# ip policy redundancy</pre>
Verification	<ul style="list-style-type: none"> Check the configurations of IPv4 PBR. Check the configurations of the route map. Check the configurations of the ACLs.
	<pre>DEV1# show ip policy Interface Route map GigabitEthernet 0/3 RM_FOR_PBR</pre>
	<pre>DEV1# show route-map route-map RM_FOR_PBR, permit, sequence 10 Match clauses: ip address 1</pre>

	<pre>Set clauses: ip next-hop recursive 200.24.20.1</pre>
	<pre>DEV1# show access-lists ip access-list standard 1 10 permit 200.24.16.0 0.0.0.255</pre>

Common Errors

- A route map is used when PBR is configured but the route map does not exist.
- An ACL is used when a route map is configured but the ACL does not exist.

9.5.2 Setting Redundant Backup or Load Balancing

Configuration Effect

- Using multiple next hops in the mutual backup mode can enhance the network reliability.
- Implementing load balancing among multiple next hops can expand the network bandwidth.

Notes

- The basic functions of PBR must be configured.
- Redundant backup and load balancing are effective only for the next hops set by the following **set** commands.

Command	Description
set ip next-hop	Configures the next hop of IPv4 packets. This command carries the <i>weight</i> parameter, which is used to set the weight of the WCMP. The default value is 1.
set ip default next-hop	Configures the default next hop of IPv4 packets. This command carries the <i>weight</i> parameter, which is used to set the weight of the WCMP. The default value is 1.
set ipv6 next-hop	Configures the next hop of IPv6 packets. This command carries the <i>weight</i> parameter, which is used to set the weight of the WCMP. The default value is 1.

set ipv6 default next-hop	Configures the default next hop of IPv6 packets. This command carries the <i>weight</i> parameter, which is used to set the weight of the WCMP. The default value is 1.
set ip next-hop recursive	Configures the recursive next hop of IPv4 packets. Only one command can be configured for a route map and packets can recur to multiple next hops (up to 32 next hops) of a static or dynamic ECMP route. The redundant backup or load balancing mode for recurring to multiple next hops is also determined by the ip policy { redundance load-balance } command.

i Up to eight next hops can be set for WCMP whereas up to 32 next hops can be set for ECMP.

Configuration Steps

Setting whether IPv4 PBR implements redundant backup or load balancing among multiple next hops

- If load balancing needs to be implemented among multiple next hops, this configuration needs to be performed.
- If load balancing is configured at present, you also need to perform this configuration to reset redundant backup.
- This configuration is effective for all PBRs configured on a device.

Command	ip policy { redundance load-balance }
Parameter Description	redundance: Indicates redundant backup. load-balance: Indicates load balancing.
Defaults	Redundant backup is configured by default.
Command Mode	Global configuration mode
Usage Guide	If redundant backup is selected, the first next hop takes effect based on the configuration sequence. If load balancing is selected, all next hops take effect at the same time and share traffic by weight.

Setting whether Ipv6 PBR implements redundant backup or load balancing among multiple next hops

- If load balancing needs to be implemented among multiple next hops, this configuration needs to be performed.

- If load balancing is configured at present, you also need to perform this configuration to reset redundant backup.
- This configuration is effective for all PBRs configured on a device.

Command	ipv6 policy { redundance load-balance }
Parameter Description	redundance: Indicates redundant backup. load-balance: Indicates load balancing.
Defaults	Redundant backup is configured by default.
Command Mode	Global configuration mode
Usage Guide	If redundant backup is selected, the first next hop takes effect based on the configuration sequence. If load balancing is selected, all next hops take effect at the same time and share traffic by weight.

Verification

- Check whether redundant backup or load balancing is implemented among multiple next hops.

Checking whether IPv4 PBR implements redundant backup or load balancing among multiple next hops

Command	show ip policy [route-map-name]
Parameter Description	route-map-name: Specifies a route map.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	See the following example and focus on the red field. <pre> QTECH# show ip policy Banlance mode: redundance Interface Route map local test GigabitEthernet 0/3 test </pre>

Checking whether IPv6 PBR implements redundant backup or load balancing among multiple next hops

Command	<code>show ipv6 policy [route-map-name]</code>
Parameter Description	<i>route-map-name</i> : Specifies a route map.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	See the following example and focus on the red field. <pre>PE#show ipv6 policy Banlance mode: redundance Interface Route map VLAN 1 RM_for_Vlan_1 VLAN 2 RM_for_Vlan_2</pre>

Configuration

Example

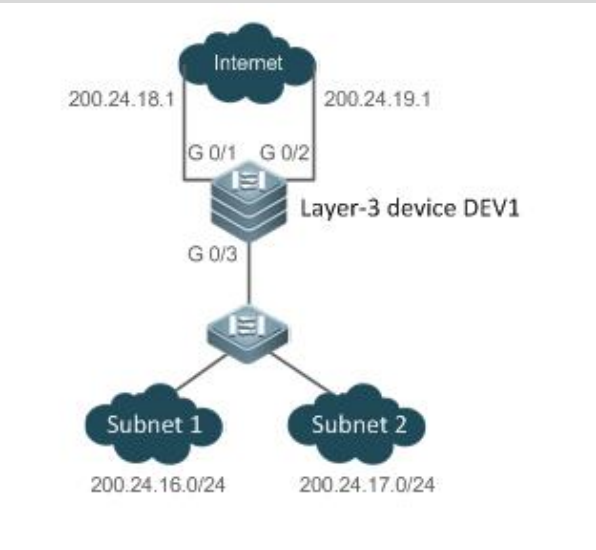
Configuring IPv4 PBR to implement redundant backup among multiple next hops

- See the preceding example: [Configuring IPv4 PBR and selecting an output link based on source addresses of packets](#)

Configuring IPv6 PBR to implement redundant backup among multiple next hops

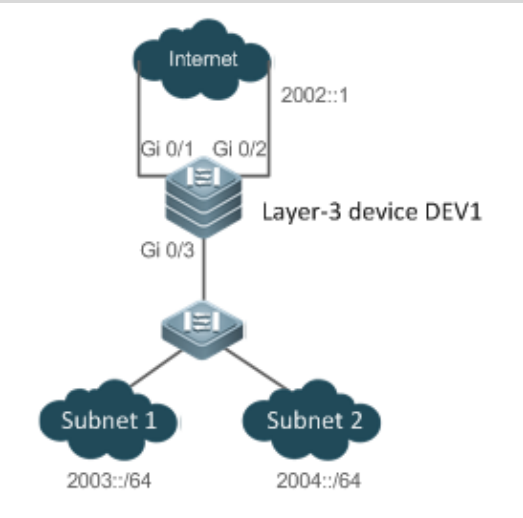
- See the preceding example: [Configuring IPv6 PBR and selecting an output link based on source addresses of packets](#)

Configuring IPv4 PBR to implement load balancing among multiple next hops

<p>Scenario</p> <p>Figure 9-7</p>	
	<p>The layer-3 device DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 200.24.16.0/24 whereas the network segment where subnet 2 resides is 200.24.17.0/24.</p> <p>DEV 1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 200.24.18.1 and 200.24.19.1.</p>
	<p>This LAN has two output interfaces for connecting the Internet. The requirements are as follows: The traffic is equally shared by GE0/1 and GE0/2.</p>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure basic functions of PBR. Specify multiple next hops. ▪ Set the load balancing mode.
	<pre> DEV1(config)# route-map RM_LOAD_PBR 10 DEV1(config-route-map)# set ip next-hop 200.24.18.1 DEV1(config-route-map)# set ip next-hop 200.24.19.1 DEV1(config-route-map)# exit DEV1(config)# interface GigabitEthernet 0/3 DEV1(config-if-GigabitEthernet 0/3)# ip policy route-map RM_LOAD_PBR DEV1(config-if-GigabitEthernet 0/3)# exit DEV1(config)# ip policy load-balance </pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Check the configurations of IPv4 PBR. ▪ Check the configurations of the route map.

	<pre>DEV1# show ip policy Balance mode: load-balance Interface Route map GigabitEthernet 0/3 RM_LOAD_PBR</pre>
	<pre>DEV1# show route-map route-map PBR-VRF-Selection, permit, sequence 10 Match clauses: Set clauses: ip next-hop 200.24.18.1 8 ip next-hop 200.24.19.1 8</pre>

Configuring IPv6 PBR to implement load balancing among multiple next hops

<p>Scenario Figure 9-8</p>	
	<p>DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 2003::/64 whereas the network segment where subnet 2 resides is 2004::/64.</p> <p>DEV1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 2001::1/64 and 2002::1/64.</p>
	<p>This LAN has two output interfaces for connecting the Internet. The requirements are as follows: The traffic is equally shared by GE0/1 and GE0/2.</p>
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Configure basic functions of PBR. Specify multiple next hops. ▪ Set the load balancing mode.
	<pre>DEV1(config)# route-map RM_LOAD_PBR 20 DEV1(config-route-map)# set ipv6 next-hop 2001::1 DEV1(config-route-map)# set ipv6 next-hop 2002::1 DEV1(config-route-map)# exit DEV1(config)# interface GigabitEthernet 0/3 DEV1(config-if-GigabitEthernet 0/3)# ipv6 policy route-map RM_LOAD_PBR DEV1(config-if-GigabitEthernet 0/3)# exit DEV1(config)# ipv6 policy load-balance</pre>
<p>Verification</p>	<ul style="list-style-type: none"> ▪ Check the configurations of IPv6 PBR. ▪ Check the configurations of the route map.
	<pre>DEV1# show ipv6 policy</pre>

```

Balance mode: load-balance
Interface          Route map
GigabitEthernet 0/3  RM_LOAD_PBR
DEV1# show route-map
route-map PBR-VRF-Selection, permit, sequence 10
  Match clauses:
  Set clauses:
  ipv6 next-hop 2001::1
  ipv6 next-hop 2002::1

```

9.5.3 Configuring Source-Address-Based PBR

Configuration

Effect

Perform personalized routing management for IPv4 or IPv6 addresses of user data streams by preparing flexible policies.

Notes

- Source-address-based PBR has a higher priority than interface-based PBR. When they are applied to an interface at the same time, interface-based PBR takes effect whereas source-address-based PBR does not take effect.

Configuration

Steps

Applying source-address-based PBR for IPv4 packets received by an interface

- To perform personalized routing management based on source IPv4 addresses for IPv4 user data streams passing a device, you should perform this configuration.
- The global configuration takes effect on the input interface of specified user data streams.
- Run the **ip policy-source in-interface** command to perform source-address-based PBR for IPv4 packets received by a specified interface.

Command	ip policy-source in-interface <i>interface-type</i> <i>sequence</i> {<i>source-address mask</i> <i>source-address/mask</i>} {[default] next-hop <i>ip-address</i> [<i>weight</i>] [default] interface <i>out-interface-type</i> vrf <i>vrf-name</i>}
Parameter Description	<i>interface-type</i> : Specifies the type of an interface to which source-address PBR is applied.

	<p><i>sequence</i>: Indicates the sequence number of a policy. A smaller sequence number means a higher priority.</p> <p><i>source-address</i>: Indicates the source IPv4 address.</p> <p><i>mask</i>: Indicates the mask of the source IPv4 address.</p> <p><i>ip-address</i>: Indicates the next-hop IPv4 address.</p> <p><i>weight</i>: Indicates the weight of a next hop.</p> <p><i>out-interface-type</i>: Indicates the type of the next hop output interface.</p> <p><i>vrf-name</i>: Indicates the name of a VRF instance.</p>
Defaults	By default, source-address-based PBR is unavailable on a device and packets are forwarded based on a routing table.
Command Mode	Global configuration mode
Usage Guide	<p>Configure multiple ip policy-source in-interface commands for the same interface. The sequence numbers of different source addresses are different. A smaller sequence number means a higher priority of the source-address-based PBR.</p> <p>When the sequence number is the same, the priorities of next hops are as follows: vrf vrf-name > next-hop ip-address > interface out-interface-type > default next-hop ip-address > default interface out-interface-type.</p> <p>Source-address-based PBR has a higher priority than interface-based PBR. When they are applied to an interface at the same time, interface-based PBR takes effect whereas source-address-based PBR does not take effect.</p>

Applying source-address-based PBR for IPv6 packets received by an interface

- To perform personalized routing management based on source IPv6 addresses for IPv6 user data streams passing a device, you should perform this configuration.
- The global configuration takes effect on the input interface of specified user data streams.
- Run the **ipv6 policy-source in-interface** command to perform source-address-based PBR for IPv6 packets received by a specified interface.

Command	ipv6 policy-source in-interface interface-type sequence {source-address/prefix-length} {[default] next-hop ipv6-address [weight] [default] interface out-interface-type vrf vrf-name}
---------	--

Parameter Description	<p><i>interface-type</i>: Specifies the type of an interface to which source-address PBR is applied.</p> <p><i>sequence</i>: Indicates the sequence number of a policy. A smaller sequence number means a higher priority.</p> <p><i>source-address</i>: Indicates the source IPv6 address.</p> <p><i>prefix-length</i>: Indicates the prefix length of a source IPv6 address.</p> <p><i>ipv6-address</i>: Indicates the next-hop IPv6 address.</p> <p><i>weight</i>: Indicates the weight of a next hop.</p> <p><i>out-interface-type</i>: Indicates the type of the next hop output interface.</p> <p><i>vrf-name</i>: Indicates the name of a VRF instance.</p>
Defaults	By default, source-address-based PBR is unavailable on a device and packets are forwarded based on a routing table.
Command Mode	Global configuration mode
Usage Guide	<p>Configure multiple ipv6 policy-source in-interface commands for the same interface. The sequence numbers of different source addresses are different. A smaller sequence number means a higher priority of the source-address-based PBR.</p> <p>When the sequence number is the same, the priorities of next hops are as follows: vrf vrf-name > next-hop ipv6-address > interface out-interface-type > default next-hop ipv6-address > default interface out-interface-type.</p> <p>Source-address-based PBR has a higher priority than interface-based PBR. When they are applied to an interface at the same time, interface-based PBR takes effect whereas source-address-based PBR does not take effect.</p>

Verification

Checking the routing information of source-address-based IPv4 PBR

Command	show ip pbr source-route [interface <i>if-name</i>]
Parameter Description	<i>if-name</i> : Indicates an interface name.
Command Mode	Privilege, global and interface configuration modes

Usage Guide	<p>Specify an interface and check the routing information of IPv4 source-address-based PBR.</p> <pre> QTECH# show ip pbr source-route PBR IPv4 Source Route Interface : GigabitEthernet 0/1 Sequence : 10 Source address : 10.1.1.1/24 VRF ID : 0 Route Flags : Route Type : PBR Direct : Permit Priority : High Match_ipaddr : Exist Mode : redundance Nextthop Count : 1 Nextthop[0] : 192.168.8.100 Weight[0] : 1 Ifindex[0] : 2 </pre>
-------------	--

Checking the routing information of IPv6 PBR


Command	show ipv6 pbr source-route [interface <i>if-name</i>]
Parameter Description	<i>if-name</i> : Indicates an interface name.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	<p>Specify an interface and check the routing information of IPv6 PBR.</p> <pre> QTECH# show ipv6 pbr source-route PBR IPv6 Source Route Interface : GigabitEthernet 0/1 Sequence : 10 Source address : 1000::1/64 </pre>

VRF ID : 0
Route Flags :
Route Type : PBR
Direct : Permit
Priority : High
Match_ipaddr : Exist
Mode : redundance
Nexthop Count : 1
Nexthop[0] : 1001::2
Weight[0] : 1
Ifindex[0] : 3

Configuration Example

Configuring IPv4 source-address-based PBR and selecting an output link based on source addresses of packets

Scenario Figure 9-9



The layer-3 device DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 200.24.16.0/24 whereas the network segment where subnet 2 resides is 200.24.17.0/24.

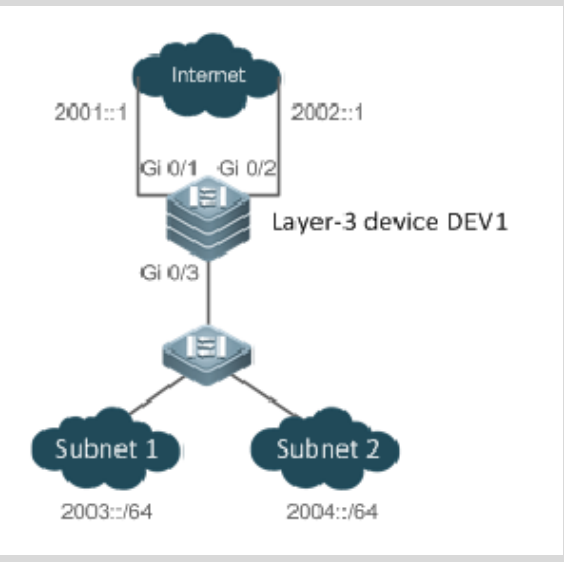
DEV 1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 200.24.18.1 and 200.24.19.1.

	<p>This LAN has two output interfaces for connecting the Internet. The requirements are as follows:</p> <ul style="list-style-type: none"> ▪ Data streams from subnet 1 for accessing the Internet should pass GE 0/1. ▪ Data streams from subnet 2 for accessing the Internet should pass GE 0/2. ▪ If the GE 0/1 link is disconnected, the data streams on GE 0/1 are switched to GE 0/2. Vice versa.
Configuration Steps	<ul style="list-style-type: none"> ▪ Set source-address-based PBR and set the next hops for packets from the GE0/3 subnet 1 to GE0/1 and GE0/2. (Pay attention to the configuration sequence.) ▪ Set source-address-based PBR and set the next hops for packets from the GE0/3 subnet 2 to GE0/2 and GE0/1. (Pay attention to the configuration sequence.) ▪ Set PBR to implement redundant backup among multiple next hops. (The default setting is redundant backup.) ❗ During redundant backup, based on the configuration sequence, the first next hop takes effect first.
	<pre>DEV1(config)# ip policy-source in-interface gigabitEthernet 0/3 1 200.24.16.0/24 next-hop 200.24.18.1 200.24.19.1 DEV1(config)# ip policy-source in-interface gigabitEthernet 0/3 2 200.24.17.0/24 next-hop 200.24.19.1 200.24.18.1 DEV1(config)# ip policy redundance</pre>
Verification	<ul style="list-style-type: none"> ▪ Check the routing information of source-address-based IPv4 PBR.
	<pre>DEV1# show ip pbr source-route PBR IPv4 Source Route Interface : GigabitEthernet 0/3 Sequence : 1 Source address : 200.24.16.0/24 VRF ID : 0 Route Flags : Route Type : PBR Direct : Permit Priority : High Match_ipaddr : Exist Mode : redundance</pre>

```
Nexthop Count : 2
Nexthop[0] : 200.24.18.1
Weight[0] : 1
Ifindex[0] : 1
Nexthop[1] : 200.24.19.1
Weight[1] : 1
Ifindex[1] : 2

Interface : GigabitEthernet 0/3
Sequence : 2
Source address : 200.24.17.0/24
VRF ID : 0
Route Flags :
Route Type : PBR
Direct : Permit
Priority : High
Match_ipaddr : Exist
Mode : redundance
Nexthop Count : 2
Nexthop[0] : 200.24.19.1
Weight[0] : 1
Ifindex[0] : 2
Nexthop[1] : 200.24.18.1
Weight[1] : 1
Ifindex[1] : 1
```

Configuring IPv6 source-address-based PBR and selecting an output link based on source addresses of packets

<p>Scenario Figure 9-10</p>	 <p>The diagram illustrates a network topology. At the top, a cloud labeled 'Internet' is connected to a 'Layer-3 device DEV1' through two interfaces: GE0/1 with IP address 2001::1 and GE0/2 with IP address 2002::1. Below DEV1, another interface GE0/3 is connected to a second Layer-3 device. This second device is connected to two subnets: 'Subnet 1' with IP address 2003::/64 and 'Subnet 2' with IP address 2004::/64.</p>
	<p>DEV 1 is connected to subnet 1 and subnet 2 through GE0/3. The network segment where subnet 1 resides is 2003::/64 whereas the network segment where subnet 2 resides is 2004::/64.</p> <p>DEV1 is connected to the Internet through GE0/1 and GE0/2 and their next hops are 2001::1/64 and 2002::1/64.</p>
	<p>This LAN has two output interfaces for connecting the Internet. The requirements are as follows:</p> <ul style="list-style-type: none"> ▪ Data streams from subnet 1 for accessing the Internet should pass GE 0/1. ▪ Data streams from subnet 2 for accessing the Internet should pass GE 0/2. ▪ If the GE 0/1 link is faulty, the data streams on GE 0/1 are switched to GE 0/2. Vice versa.
<p>Configuration Steps</p>	<ul style="list-style-type: none"> ▪ Set source-address-based PBR and set the next hops for packets from the GE0/3 subnet 1 to GE0/1 and GE0/2. (Pay attention to the configuration sequence.) ▪ Set source-address-based PBR and set the next hops for packets from the GE0/3 subnet 2 to GE0/2 and GE0/1. (Pay attention to the configuration sequence.) ▪ Set PBR to implement redundant backup among multiple next hops. <p>i During redundant backup, based on the configuration sequence, the first next hop takes effect first.</p>
	<pre>DEV1(config)# ipv6 policy-source in-interface gigabitEthernet 0/3 1 2003::/64 next-hop 2001::1 2002::1</pre>

	<pre>DEV1(config)# ip policy-source in-interface gigabitEthernet 0/3 2 2004::/64 next-hop 2002::1 2001::1 DEV1(config)# ipv6 policy redundance</pre>
Verification	<ul style="list-style-type: none"> Check the configuration of IPv6 source-address-based PBR.
	<pre>DEV1# show ipv6 pbr source-route PBR IPv6 Source Route Interface : GigabitEthernet 0/3 Sequence : 1 Source address : 2003::/64 VRF ID : 0 Route Flags : Route Type : PBR Direct : Permit Priority : High Match_ipaddr : Exist Mode : redundance Nexthop Count : 2 Nexthop[0] : 2001::1 Weight[0] : 1 Iindex[0] : 1 Nexthop[1] : 2002::1 Weight[1] : 1 Iindex[1] : 2 Interface : GigabitEthernet 0/3 Sequence : 2 Source address : 2004::/64 VRF ID : 0 Route Flags : Route Type : PBR Direct : Permit Priority : High</pre>

```

Match_ipaddr : Exist
Mode      : redundance
Nextthop Count : 2
Nextthop[0] : 2002::1
Weight[0]  : 1
Ifindex[0] : 2
Nextthop[1] : 2001::1
Weight[1]  : 1
Ifindex[1] : 1

```

9.6 Monitoring

Clearing

⚠ Running the **clear** commands may lose vital information and thus interrupt services.

Description	Command
Clears the statistics about packets forwarded by IPv4 PBR.	clear ip pbr statistics [interface <i>if-name</i> local]
Clears the statistics about packets forwarded by IPv6 PBR.	clear ipv6 pbr statistics [interface <i>if-name</i> local]

Displaying

	Command
Displays the configurations of IPv4 PBR.	show ip policy
Displays the configurations of IPv6 PBR.	show ipv6 policy
Displays the configurations of a route map.	show route-map [<i>name</i>]

Displays the configurations of an ACL.	show access-list
Displays the correlation between IPv4 PBR and BFD.	show ip pbr bfd
Displays the correlation between IPv6 PBR and BFD.	show ipv6 pbr bfd
Displays the routing information of IPv4 PBR.	show ip pbr route [interface <i>if-name</i> local]
Displays the routing information of IPv6 PBR.	show ipv6 pbr route [interface <i>if-name</i> local]
Displays a route map used by IPv4 PBR.	show ip pbr route-map <i>rmap-name</i>
Displays a route map used by IPv6 PBR.	show ipv6 pbr route-map <i>rmap-name</i>
Displays the routing information of IPv4 source-address-based PBR.	show ip pbr source-route [interface <i>if-name</i>]
Displays the routing information of IPv6 source-address-based PBR.	show ipv6 pbr source-route [interface <i>if-name</i>]
Displays the statistics about IPv4 PBR.	show ip pbr statistics [interface <i>if-name</i> local]
Displays the statistics about IPv6 PBR.	show ipv6 pbr statistics [interface <i>if-name</i> local]

Debugging

 System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
-------------	---------

Debugs PBR errors.	debug pbr error
Debugs PBR events.	debug pbr events
Debugs multiple service cards supported by PBR.	debug pbr ms
Debugs PBR message communication.	debug pbr msg
Debugs interaction between PBR and NSM.	debug pbr nsm
Debugs packet forwarding of PBR.	debug pbr packet
Debugs PBR GR.	debug pbr restart

10 MANAGING ROUTES

10.1 Overview

The network service module (NSM) manages the routing table, consolidates routes sent by various routing protocols, and selects and sends preferred routes to the routing table. Routes discovered by various routing protocols are stored in the routing table. These routes are generally classified by source into three types:

- Direct route: It is the route discovered by a link-layer protocol and is also called interface route.
- Static route: It is manually configured by the network administrator. A static route is easy to configure and less demanding on the system, and therefore applicable to a small-sized network that is stable and has a simple topology. However, when the network topology changes, the static route must be manually reconfigured and cannot automatically adapt to the topological changes.
- Dynamic route: It is the route discovered by a dynamic routing protocol.

10.2 Applications

Application	Description
Basic Functions of the Static Route	Manually configure a route.
Floating Static Route	Configure a standby route in the multipath scenario.
Load Balancing Static Route	Configure load balancing static routes in the multipath scenario.
Correlation of Static Routes with BFD	Use the Bidirectional Forwarding Detection (BFD) function to test whether the next hop of a static route is reachable.
Fast Reroute of Static Routes	Use the fast reroute function to improve the switching performance in the multipath scenario.
ARP-to-host Conversion	Use the ARP-to-host conversion to not transmit ARP packets sent by hosts across the whole network.

10.2.1 Basic Functions of the Static Route


Scenario

On a network with a simple topology, you can configure only static routes to implement network interworking. Appropriate configuration and use of static routes can improve the network performance and guarantee the bandwidth for important network applications.

As shown in Figure 10-1, to implement interworking between PC 1, PC 2, and PC 3, you can configure static routes on R 1, R 2, and R 3.

- On R 1, configure a route to the network segment of PC 2 through R 2, and a route to the network segment of PC 3 through R 3.
- On R 2, configure a route to the network segment of PC 1 through R 1, and a route to the network segment of PC 3 through R 3.
- On R 3, configure a route to the network segment of PC 1 through R 1, and a route to the network segment of PC 2 through R 2.

Figure 10-1



Deployment

- Configure the address and subnet mask of each interface.
- Configure static routes on R 1, R 2, and R 3.

10.2.2 Floating Static Route


Scenario

If no dynamic routing protocol is configured, you can configure floating static routes to implement dynamic switching of routes to prevent communication interruption caused by the network connection failures.

As shown in Figure 10-2, to prevent communication interruption caused by a line failure between R 1 and R 3, you can configure a floating static route respectively on R 1 and R 3. Normally, packets are forwarded on a path with a small administrative distance. If a link on this path is down, the route is automatically switched to the path with a large administrative distance.

- On R1, configure two routes to the network segment of PC 3, including a route through R 3 (default distance = 1) and a route through R 2 (default distance = 2).
- On R 3, configure two routes to the network segment of PC 1, including a route through R 1 (default distance = 1) and a route through R 2 (default distance = 2).

Figure 10-2



Deployment

- Configure the address and subnet mask of each interface.
- Configure static routes on R 1, R 2, and R 3.

10.2.3 Load Balancing Static Route

Scenario

If there are multiple paths to the same destination, you can configure load balancing routes. Unlike floating routes, the administrative distances of load balancing routes are the same. Packets are distributed among these routes based on the balanced forwarding policy.

As shown in Figure 10-3, load balancing routes are configured respectively on R 1 and R 3 so that packets sent to the network segment of PC 3 or PC 1 are balanced between two routes, including a route through R 2 and a route through R 4.

- On R 1, configure two routes to the network segment of PC 3, including a route through R 2 and a route through R 4.
- On R 3, configure two routes to the network segment of PC 1, including a route through R 2 and a route through R 4.

Figure 10-3



Remarks	On the switch, the load is balanced based on the destination IP address by default.
----------------	--

Deployment

- Configure the address and subnet mask of each interface.
- Configure static routes on R 1, R 2, R 3, and R 4.
- Configure the load balancing policy on R 1 and R 3.

10.2.4 Correlation of Static Routes with Track, BFD or ARP

Scenario

When the floating static routes or load balancing static routes are configured, the static routes may fail to sense the route failures if the line is faulty but the interface status is normal. To resolve this problem, the device needs to check whether the next hop of a static route is reachable. If the next hop is not reachable, the device can switch the traffic to the standby route.


You can use the Track, BFD or ARP function to check whether the next hop of a static route is reachable. The following scenario takes BFD as an example.

- ⚠ Three modes can all check whether the next hop of a static route is reachable. If more than one mode is configured, when one mode checks the next hop of a static route is unreachable, then the static route is invalid. Therefore, it is recommend to take only one mode.

As shown in Figure 10-4, to prevent communication interruption caused by a line failure between R 1 and R 3, you can configure a floating static route respectively on R 1 and R 3, and correlate static routes with BFD.

- On R 1, configure two routes to the network segment of PC 3, including a route through R 3 (default distance = 1) and a route through R 2 (default distance = 2). BFD is enabled on the first route to check whether the next hop 1.1.13.3 is reachable, and on the second route to check whether the next hop 1.1.12.2 is reachable.
- On R 3, configure two routes to the network segment of PC 1, including a route through R 1 (default distance = 1) and a route through R 2 (default distance = 2). BFD is enabled on the first route to check whether the next hop 1.1.13.1 is reachable, and on the second route to check whether the next hop 1.1.23.2 is reachable.

Figure 100-4



Deployment

- Configure the address and subnet mask of each interface.
- Configure the BFD parameters on each interface.
- Configure static routes and correlate these static routes with BFD on R 1, R 2, and R 3.

10.2.5 Fast Reroute of Static Routes


Scenario

To accelerate route switching and shorten the communication interruption time when no dynamic routing protocol is configured, you can either correlate static routes with Track, BFD or ARP to check whether the next hop is reachable. In addition, you can or configure fast reroute to further improve the convergence performance.

As shown in Figure 10-5, to prevent communication interruption caused by a line failure between R 1 and R 3, you can configure static fast reroute respectively on R 1 and R 3. Normally, packets are forwarded on the path between R 1 and R 3. When the link on this route is down, packets are automatically rerouted to R 2.

- On R 1, configure a route with the exit interface set to Gi0/3 and the next hop set to 1.1.13.3, and a standby route with the exit interface set to Gi0/2 and the next hop set to 1.1.12.2.
- On R 3, configure a route with the exit interface set to Gi0/1 and the next hop set to 1.1.13.1, and a standby route with the exit interface set to Gi0/2 and the next hop set to 1.1.23.2.

Figure 10-5



Deployment

- Configure the address and subnet mask of each interface.
- Configure static routes on R 1, R 2, and R 3.
- Configure static fast reroute on R 1, R 2, and R 3.

10.2.6 ARP-to-host Conversion


Scenario

In special requirement scenarios, communication is required between hosts, but ARP packets sent by hosts are not desired to transmit across the whole network. In this case, you can run the arp-to-host command on the devices connected to hosts to convert ARP entries to host routes, so that the remote devices learn the host routes through BGP redistribution to meet the scenario requirement.

As shown in the following figure, PC1 and PC2 are two hosts in the same network segment. PC1 and PC2 are required to communicate with each other, but ARP packets are not desired to transmit across the whole network. Run the arp-to-host command on R1 and R2 to generate ARP host routes, establish

a BGP neighbor relationship between R1 and R2, and enable ARP host route redistribution. In this case, both R1 and R2 can obtain routes to PC1 and PC2 to implement network interconnection.

Figure 10-6



Deployment

- Configure IP addresses and masks for interfaces.
- Enable ARP-to-host conversion on interfaces of R1 and R2.
- Establish a BGP neighbor relationship between R1 and R2, and redistribute ARP host routes.

10.3 Features

Feature	Description
Route Computation	Generate a valid route on a device.
Optimal Route Selection	Select an optimal route to forward packets.
Default Route	Forward all packets and help reduce the size of a routing table.
Route Reliability	Quickly detect a route failure and recover communication.

10.3.1 Route Computation

Routing Function

Routing functions are classified into IPv4 and IPv6 routing functions. If the routing functions are disabled, a device is equivalent to a host and cannot forward routes.

Dynamic Route

A dynamic routing protocol learns remote routes and dynamically updates routes by exchanging routes with neighbors. If a neighbor is the next hop of a route and this neighbor fails, the route fails as well.

Static Route

On a network with a simple topology, you can configure only static routes to implement network interworking. Appropriate configuration and use of static routes can improve the network performance and guarantee the bandwidth for important network applications.

Whether a static route is active is computed based on the status of the local interface. When the exit interface of a static route is located at layer 3 (L3) and is in Up status (the link status is Up and the IP address is configured), this route is active and can be used for packet forwarding.

A static route can go across VPN routing & forwarding (VRF) instances. The next hop or exit interface of a static route of VRF 1 can be configured on VRF 2.

ARP-to-host Conversion

Enable ARP-to-host conversion on a specified interface. If the interface is a Layer-3 interface in up status, ARP entries will be converted into 32-bit host routes on the interface.

10.3.2 Optimal Route Selection

Administrative Distance

When multiple routing protocols generate routes to the same destination, the priorities of these routes can be determined based on the administrative distance. A smaller administrative distance indicates a higher priority.

Equal-Cost Route

If multiple routes to the same destination have different next hops but the same administrative distance, these routes are mutually equal-cost routes. Packets are distributed among these routes to implement load balancing based on the balanced forwarding policy.

On a specific device, the total number of equal-cost routes is limited. Routes beyond the limit do not participate in packet forwarding.

Floating Route

If multiple routes to the same destination have different next hops and different administrative distances, these routes are mutually floating routes. The route with the smallest administrative distance will be first selected for packet forwarding. If this route fails, a route with a larger administrative

distance is further selected for forwarding, thus preventing communication interruption caused by a network line failure.

10.3.3 Default Route

In the forwarding routing table, the route with the destination network segment 0.0.0.0 and the subnet mask 0.0.0.0 is the default route. Packets that cannot be forwarded by other routes will be forwarded by the default route. The default route can be statically configured or generated by a dynamic routing protocol.

Static Default Route

On a L3 switch, a static route with the network segment 0.0.0.0 and the subnet mask 0.0.0.0 is configured to generate the default route.

Default Network

The default network is configured to generate a default route. If the **ip default-network** command is configured to specify a network (a classful network, such as a Class A, B, or C network), and this network exists in the routing table, the router will use this network as the default network and the next hop of this network is the default gateway. As the network specified by the **ip default-network** command is a classful one, if this command is used to identify a subnet in a classful network, the router automatically generates a static route of the classful network instead of any default route.

10.3.4 Route Reliability

When a device on a network is faulty, some routes become unreachable, resulting in traffic interruption. If connectivity of the next hop can be detected in real time, the route can be re-computed when a fault occurs, or traffic can be switched over to the standby route.

Correlation with Track

A track object is an abstract concept. It can be used to trace whether an IP address is reachable or whether an interface is up. If a dynamic routing protocol or a static route is correlated with the Track function, the dynamic routing protocol or the static route can quickly learn whether the next hop is reachable so as to respond quickly.

Correlation with BFD

The BFD protocol provides a light-load and fast method for detecting the connectivity of the forwarding path between two adjacent routers. If a dynamic routing protocol or a static route is correlated with

the BFD function, the dynamic routing protocol or the static route can quickly learn whether the next hop is reachable so as to respond quickly.

i The detection performance of BFD is better than that of Track.

Correlation with ARP

If a static route is correlated with the ARP function, when the ARP does not exist, the static route can quickly learn that the next hop is unreachable so as to respond quickly.

Fast Reroute

Fast reroute provides a standby route. When a dynamic routing protocol or a static route detects that the next hop is unreachable, it immediately switches traffic over to the standby route to recovery communication.

10.4 Limitations



QSW-6900 series products support 64 groups of ECMP, and each group supports a maximum of 32 entries. A group refers to multiple equivalent next hops of a prefix.

10.5 Configuration

Configuration Item	Description and Command	
Configuring a Static Route	! (Mandatory) It is used to configure a static route entry.	
	ip route	Configures an IPv4 static route.
	ipv6 route	Configures an IPv6 static route.
Configuring a Default Route	! (Optional) It is used to configure the default gateway.	
	ip route 0.0.0.0 0.0.0.0 gateway	Configures an IPv4 default gateway on a L3 device.
	ipv6 route ::/0 ipv6-gateway	Configures an IPv6 default gateway on a L3 device.

	ip default network	Configures an IPv4 default network on a L3 device.
Configuring Route Limitations	⚠ (Optional) It is used to limit the number of equal-cost routes and number of static routes, or disable routing.	
	maximum-paths	Configures the maximum number of equal-cost routes.
	ip static route-limit	Configures the maximum number of IPv4 static routes.
	ipv6 static route-limit	Configures the maximum number of IPv6 static routes.
	no ip routing	Disables IPv4 routing.
	noipv6 unicast-routing	Disables IPv6 routing.
	no ip route static inter-vrf	Prohibits static routing across VRFs.
Correlating a Static Route with BFD	⚠ (Optional) It is used to correlate a static route with BFD.	
	ip route static bfd	Correlates an IPv4 static route with BFD.
	ipv6 route static bfd	Correlates an IPv6 static route with BFD.
Configure Static Fast Reroute	⚠ (Optional) It is used to configure static fast reroute.	
	route-map	Configures a route map.
	set fast-reroute backup-next-hop	Configures the standby interface and standby next hop for fast reroute.
	ip fast-reroute	Configures static fast reroute.

Configure ARP-to-host Conversion	ip route arp-to-host interface <i>interface-name</i>	Configures ARP-to-host conversion on a specified interface.
	redistribute arp-host[route-map map-tag] [metricmetric-value]	Configures BGP redistribution of ARP host route.

10.5.1 Configuring a Static Route

Configuration

Effect

- Generate a static route in the routing table. Use the static route to forward packets to a remote network.

Notes

- If the **no ip routing** command is configured on a L3 switch, you cannot configure IPv4 static routes on this switch, and existing IPv4 static routes will also be deleted. Before the device is restarted, reconfiguring the **ip routing** command can recover the deleted IPv4 static routes. After the device is restarted, deleted IPv4 static routes cannot be recovered.
- If the **no ipv6 unicast- routing** command is configured on a L3 switch, you cannot configure IPv6 static routes on this switch, and existing IPv6 static routes will also be deleted. Before the device is restarted, reconfiguring the **ipv6 unicast- routing** command can recover the deleted IPv6 static routes. After the device is restarted, deleted IPv6 static routes cannot be recovered.
- To correlate a static route with the Track function, you must run the **track** command to configure a track object.

Configuration

Steps

Configuring a Static IPv4 Route

Configure the following command on an IPv4-enabled router.

Command	ip route [vrf vrf_name] networknet-mask {ip-address interface [ip-address]} [distance] [tag tag] [permanent track object-number/ arp] [weight number] [descriptiondescription-text] [disabled enabled] [global]	
Parameter Description	vrfvrf_name	(Optional) Indicates the routing VRF, which can be a single-protocol IPv4 VRF or a multi-protocol VRF of a configured IPv4 address family. The VRF is a global VRF by default.

	<i>network</i>	Indicates the address of the destination network.
	<i>net-mask</i>	Indicates the mask of the destination network.
	<i>ip-address</i>	(Optional) Indicates the next-hop address of the static route. You must specify at least one of <i>ip-address</i> and <i>interface</i> , or both of them. If <i>ip-address</i> is not specified, a static direct route is configured.
	<i>interface</i>	(Optional) Indicates the next-hop exit interface of the static route. You must specify at least one of <i>ip-address</i> and <i>interface</i> , or both of them. If <i>interface</i> is not specified, a recursive static direct route is configured. The exit interface is obtained by the next hop in the routing table.
	<i>distance</i>	(Optional) Indicates the administrative distance of the static route. The administrative distance is 1 by default.
	<i>tag</i>	(Optional) Indicates the tag of the static route. The tag is 0 by default.
	permanent	(Optional) Indicates the flag of the permanent route. The static route is not a permanent route by default.
	track <i>object-number</i>	(Optional) Indicates correlation with Track. <i>object-number</i> indicates the ID of the track object. By default, the static route is not correlated with the Track function.
	weight <i>number</i>	(Optional) Indicates the weight of the static route. The weight is 1 by default.
	description <i>description-text</i>	(Optional) Indicates the description of the static route. By default, no description is configured. <i>description-text</i> is a string of one to 60 characters.
	disabled/enabled	(Optional) Indicates the enable flag of the static route. The flag is enabled by default.
	global	(Optional) Indicates that the next hop belongs to a global VRF. By default, the VRF of the next hop is the same as the VRF specified by <i>vrf name</i> .
	arp	(Optional) Correlate the static route with ARP function. By default, it is not correlated.
Defaults	By default, no static route is configured.	

Command Mode	Global configuration mode
Usage Guide	<p>The simplest configuration of this command is ip route networknet-maskip-address.</p> <p>If the static route is correlated with Track and the down status of the trace object is detected, the static route is not active and does not participate in packet forwarding.</p> <p>If the static route is correlated with ARP, but no ARP information is detected, the static route is not active and does not participate in packet forwarding.</p>

Configuring an IPv6 Static Route

Configure the following command on an IPv6-enabled router.

Command	ipv6 route [vrfvrf-name] ipv6-prefix/prefix-length { ipv6-address [nexthop-vrf {vrf-name1/ default}] interface [ipv6-address [nexthop-vrf {vrf-name1/ default}]] } [distance] [weightnumber] [descriptiondescription-text]	
Parameter Description	vrfvrf-name	(Optional) Indicates the routing VRF, which must be a multi-protocol VRF of a configured IPv6 address family. The VRF is a global VRF by default.
	ipv6-prefix	Indicates the IPv6 prefix, which must comply with the address expression specified in RFC4291.
	prefix-length	Indicates the length of the IPv6 prefix. Note that a slash (/) must be added in front of the length.
	ipv6-address	(Optional) Indicates the next-hop address of the static route. You must specify at least one of <i>ipv6-address</i> and <i>interface</i> , or both of them. If <i>ipv6-address</i> is not specified, a static direct route is configured.
	interface	(Optional) Indicates the next-hop exit interface of the static route. You must specify at least one of <i>ipv6-address</i> and <i>interface</i> , or both of them. If <i>interface</i> is not specified, a recursive static direct route is configured. The exit interface is obtained by the next hop in the routing table.
	nexthop-vrf vrf-name1	(Optional) Indicates the routing VRF of the next hop, which must be a multi-protocol VRF of a configured IPv6 address family. By default, the VRF of the next hop is the same as the VRF specified by the VRF name. nexthop-vrf default indicates that the VRF of the next shop is a global VRF.
	distance	(Optional) Indicates the administrative distance of the static route. The administrative distance is 1 by default.

	weight number	(Optional) Indicates the weight of the static route, which must be specified when you configure equal-cost routes. The weight ranges from 1 to 8. When the weights of all equal-cost routes of a route are summed up, the sum cannot exceed the maximum number of equal-cost routes that can be configured for the route. Weighting of equal-cost routes of a route indicates the traffic ratio of these routes. The weight is 1 by default.
	description <i>description-text</i>	(Optional) Indicates the description of the static route. By default, no description is configured. <i>description-text</i> is a string of one to 60 characters.
Defaults	By default, no static route is configured.	
Command Mode	Global configuration mode	
Usage Guide	The simplest configuration of this command is ipv6 route <i>ipv6-prefix / prefix-length ipv6-address</i> .	

Verification


- Run the **show ip route** command to display the IPv4 routing table and check whether the configured IPv4 static route takes effect.
- Run the **show ipv6 route** command to display the IPv6 routing table and check whether the configured IPv6 static route takes effect.

Configuration

Example

Configuring Static Routes to Implement Interworking of the IPv4 Network

Scenario
Figure 10-7



**Configurati
on Steps**

- Configure interface addresses on each device.

R1

```
R1#configure terminal
R1(config)#interface gigabitEthernet 0/0
R1(config-if-GigabitEthernet 0/0)# ip address 1.1.1.1 255.255.255.0
R1(config-if-GigabitEthernet 0/0)# exit
R1(config)#interface gigabitEthernet 0/2
R1(config-if-GigabitEthernet 0/2)# ip address 1.1.12.1 255.255.255.0
R1(config-if-GigabitEthernet 0/2)# exit
R1(config)#interface gigabitEthernet 0/3
R1(config-if-GigabitEthernet 0/3)# ip address 1.1.13.1 255.255.255.0
```

R2

```
R2#configure terminal
R2(config)#interface gigabitEthernet 0/0
R2(config-if-GigabitEthernet 0/0)# ip address 1.1.2.1 255.255.255.0
R2(config-if-GigabitEthernet 0/0)# exit
R2(config)#interface gigabitEthernet 0/1
R2(config-if-GigabitEthernet 0/1)# ip address 1.1.12.2 255.255.255.0
R2(config-if-GigabitEthernet 0/1)# exit
R2(config)#interface gigabitEthernet 0/3
R2(config-if-GigabitEthernet 0/3)# ip address 1.1.23.2 255.255.255.0
```

R3	<pre> R3#configure terminal R3(config)#interface gigabitEthernet 0/0 R3(config-if-GigabitEthernet 0/0)# ip address 1.1.3.1 255.255.255.0 R3(config-if-GigabitEthernet 0/0)# exit R3(config)#interface gigabitEthernet 0/1 R3(config-if-GigabitEthernet 0/1)# ip address 1.1.13.3 255.255.255.0 R3(config-if-GigabitEthernet 0/0)# exit R3(config)#interface gigabitEthernet 0/2 R3(config-if-GigabitEthernet 0/2)# ip address 1.1.23.3 255.255.255.0 </pre>
	<ul style="list-style-type: none"> Configure static routes on each device.
R1	<pre> R1#configure terminal R1(config)#ip route 1.1.2.0 255.255.255.0 GigabitEthernet 0/2 1.1.12.2 R1(config)# ip route 1.1.3.0 255.255.255.0 GigabitEthernet 0/3 1.1.13.3 </pre>
R2	<pre> R2#configure terminal R2(config)#ip route 1.1.1.0 255.255.255.0 GigabitEthernet 0/1 1.1.12.1 R2(config)# ip route 1.1.3.0 255.255.255.0 GigabitEthernet 0/3 1.1.23.3 </pre>
R3	<pre> R3#configure terminal R3(config)#ip route 1.1.2.0 255.255.255.0 GigabitEthernet 0/2 1.1.23.2 R3(config)# ip route 1.1.1.0 255.255.255.0 GigabitEthernet 0/1 1.1.13.1 </pre>
Verification	<ul style="list-style-type: none"> Display the routing table.
R1	<pre> R1# show ip route Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area, * - candidate default </pre>


	<p>Gateway of last resort is no set</p> <p>C 1.1.1.0/24 is directly connected, GigabitEthernet 0/0</p> <p>C 1.1.1.1/32 is local host.</p> <p>S 1.1.2.0/24 [1/0] via 1.1.12.2, GigabitEthernet 0/2</p> <p>S 1.1.3.0/24 [1/0] via 1.1.13.3, GigabitEthernet 0/2</p> <p>C 1.1.12.0/24 is directly connected, GigabitEthernet 0/2</p> <p>C 1.1.12.1/32 is local host.</p> <p>C 1.1.13.0/24 is directly connected, GigabitEthernet 0/3</p> <p>C 1.1.13.1/32 is local host.</p>
R2	<p>R2# show ip route</p> <p>Codes: C - Connected, L - Local, S - Static</p> <p>R - RIP, O - OSPF, B - BGP, I - IS-IS</p> <p>N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2</p> <p>E1 - OSPF external type 1, E2 - OSPF external type 2</p> <p>SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2</p> <p>IA - Inter area, * - candidate default</p> <p>Gateway of last resort is no set</p> <p>S 1.1.1.0/24 [1/0] via 1.1.12.1, GigabitEthernet 0/0</p> <p>C 1.1.2.0/24 is directly connected, GigabitEthernet 0/0</p> <p>C 1.1.2.1/32 is local host.</p> <p>S 1.1.3.0/24 [1/0] via 1.1.23.3, GigabitEthernet 0/3</p> <p>C 1.1.12.0/24 is directly connected, GigabitEthernet 0/1</p> <p>C 1.1.12.2/32 is local host.</p> <p>C 1.1.23.0/24 is directly connected, GigabitEthernet 0/3</p> <p>C 1.1.23.2/32 is local host.</p>
R3	<p>R3# show ip route</p> <p>Codes: C - Connected, L - Local, S - Static</p> <p>R - RIP, O - OSPF, B - BGP, I - IS-IS</p> <p>N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2</p> <p>E1 - OSPF external type 1, E2 - OSPF external type 2</p>

SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
 IA - Inter area, * - candidate default

Gateway of last resort is no set

S 1.1.1.0/24 [1/0] via 1.1.13.1, GigabitEthernet 0/2
 S 1.1.2.0/24 [1/0] via 1.1.23.2, GigabitEthernet 0/2
 C 1.1.3.0/24 is directly connected, GigabitEthernet 0/0
 C 1.1.3.1/32 is local host.
 C 1.1.13.0/24 is directly connected, GigabitEthernet 0/1
 C 1.1.13.3/32 is local host.
 C 1.1.23.0/24 is directly connected, GigabitEthernet 0/2
 C 1.1.23.3/32 is local host.

Correlating IPv4 Static Routes with Track

<p>Scenario Figure 10-8</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure static routes on R 1 and R 2, and specify the exit interface or next hop as the interworking interface. Correlate static routes with Track on R 1 and R 2, and check the connectivity of the next hops of static routes.
<p>R1</p>	<pre>R1#configure terminal R1(config)#interface gigabitEthernet 0/1 R1(config-if-GigabitEthernet 0/1)# ip address 1.1.12.1 255.255.255.0 R1(config-if-GigabitEthernet 0/1)# exit R1(config)#track 2 interface gigabitEthernet 0/1 line-protocol R1(config)# ip route 1.1.2.0 255.0.0.0 gigabitEthernet 0/1 1.1.12.2 track 2</pre>
<p>R2</p>	<pre>R2#configure terminal R2(config)#interface gigabitEthernet 0/1 R2(config-if-GigabitEthernet 0/1)# ip address 1.1.12.2 255.255.255.0</pre>

	<pre>R1(config-if-GigabitEthernet 0/1)# exit R1(config)#track 2 interface gigabitEthernet 0/1 line-protocol R1(config)# ip route 1.1.1.0 255.0.0.0 gigabitEthernet 0/1 1.1.12.1 track 2</pre>
Verification	<ul style="list-style-type: none"> Display the Track status. Display the static routes correlated with Track. <pre>R1# show track 2 Track 2 Interface gigabitEthernet 0/1 The state is Up, delayed Down (5 secs remaining) 1 change, current state last: 300 secs Delay up 0 secs, down 0 secs R1#show ip route track-table ip route 1.1.2.0 255.0.0.0 GigabitEthernet 0/1 1.1.12.2 track 2 up</pre>

Configuring Static Routes to Implement Interworking of the IPv6 Network

Scenario Figure 10-9	
Configurati on Steps	<ul style="list-style-type: none"> Configure interface addresses on each device.
R1	<pre>R1#configure terminal R1(config)#interface gigabitEthernet 0/0 R1(config-if-GigabitEthernet 0/0)# ipv6 address 1111:1111::1/64 R1(config-if-GigabitEthernet 0/0)# exit R1(config)#interface gigabitEthernet 0/1 R1(config-if-GigabitEthernet 0/1)# ipv6 address 1111:1212::1/64</pre>
R2	<pre>R2#configure terminal</pre>

	<pre>R2(config)#interface gigabitEthernet 0/0 R2(config-if-GigabitEthernet 0/0)#ipv6 address 1111:2323::1/64 R2(config-if-GigabitEthernet 0/0)# exit R2(config)#interface gigabitEthernet 0/1 R2(config-if-GigabitEthernet 0/1)# ipv6 address 1111:1212::2/64</pre>
	<ul style="list-style-type: none"> ▪ Configure static routes on each device.
R1	<pre>R1#configure terminal R1(config)# ipv6 route 1111:2323::0/64 gigabitEthernet 0/1</pre>
R2	<pre>R2#configure terminal R2(config)#ipv6 route 1111:1111::0/64 gigabitEthernet 0/1</pre>
Verification	<ul style="list-style-type: none"> ▪ Display the routing table.
R1	<pre>R1# show ipv6 route IPv6 routing table name - Default - 10 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area C 1111:1111::/64 via GigabitEthernet 0/0, directly connected L 1111:1111::1/128 via GigabitEthernet 0/0, local host C 1111:1212::/64 via GigabitEthernet 0/1, directly connected L 1111:1212::1/128 via GigabitEthernet 0/1, local host S 1111:2323::/64 [1/0] via GigabitEthernet 0/1, directly connected C FE80::/10 via ::1, Null0 C FE80::/64 via GigabitEthernet 0/0, directly connected L FE80::2D0:F8FF:FEFB:C092/128 via GigabitEthernet 0/0, local host</pre>

	<pre> C FE80::/64 via GigabitEthernet 0/1, directly connected L FE80::2D0:F8FF:FEFB:C092/128 via GigabitEthernet 0/1, local host </pre>
R2	<pre> R2# show ipv6 route IPv6 routing table name - Default - 10 entries Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area C 1111:2323::/64 via GigabitEthernet 0/0, directly connected L 1111:2323::1/128 via GigabitEthernet 0/0, local host C 1111:1212::/64 via GigabitEthernet 0/1, directly connected L 1111:1212::1/128 via GigabitEthernet 0/1, local host S 1111:1111::/64 [1/0] via GigabitEthernet 0/1, directly connected C FE80::/10 via ::1, Null0 C FE80::/64 via GigabitEthernet 0/0, directly connected L FE80::2D0:F8FF:FEFB:C092/128 via GigabitEthernet 0/0, local host C FE80::/64 via GigabitEthernet 0/1, directly connected L FE80::2D0:F8FF:FEFB:C092/128 via GigabitEthernet 0/1, local host </pre>

Common Errors

- The link on the interface is not up.
- No IP address is configured for the interface.
- The static route is correlated with Track, but the track object is not configured.

10.5.2 Configuring a Default Route

Configuration Effect

- Generate a default route in the routing table. The default route is used to forward packets that cannot be forwarded by other routes.

Notes

- On a L3 switch, run the **ip route 0.0.0.0 0.0.0.0 gateway** or **ipv6 route ::/0 ipv6-gateway** command to configure the default gateway.
- If the **no ip routing** or **no ipv6 unicast-routing** command is configured on a L3 switch, you can run the **ip default gateway** or **ipv6 default gateway** command to configure the default gateway.

Configuration Steps

Configuring the IPv4 Gateway on a L2 Switch

Command	ip default-gateway gateway	
Parameter Description	<i>gateway</i>	indicates the IPv4 gateway address.
Defaults	By default, no static default route is configured.	
Command Mode	Global configuration mode	
Usage Guide	N/A	

Configuring the IPv6 Gateway on a L2 Switch

Command	ipv6 default-gateway gateway	
Parameter Description	<i>gateway</i>	indicates the IPv6 gateway address.
Defaults	By default, no static default route is configured.	
Command Mode	Global configuration mode	
Usage Guide	N/A	

Configuring the IPv4 Default Gateway on a L3 Switch

Command	ip route [<i>vrf vrf_name</i>] 0.0.0.00.0.0.0 { <i>ip-address</i> <i>interface</i> [<i>ip-address</i>]} [<i>distance</i>] [<i>tag tag</i>] [<i>permanent</i>] [<i>weight number</i>] [<i>descriptiondescription-text</i>] [<i>disabled</i> <i>enabled</i>] [<i>global</i>]	
Parameter Description	vrf <i>vrf_name</i>	(Optional) Indicates the routing VRF, which can be a single-protocol IPv4 VRF or a multi-protocol VRF of a configured IPv4 address family. The VRF is a global VRF by default.
	0.0.0.0	Indicates the address of the destination network.
	0.0.0.0	Indicates the mask of the destination network.
	<i>ip-address</i>	(Optional) Indicates the next-hop address of the static route. You must specify at least one of <i>ip-address</i> and <i>interface</i> , or both of them. If <i>ip-address</i> is not specified, a static direct route is configured.
	<i>interface</i>	(Optional) Indicates the next-hop exit interface of the static route. You must specify at least one of <i>ip-address</i> and <i>interface</i> , or both of them. If <i>interface</i> is not specified, a recursive static direct route is configured. The exit interface is obtained by the next hop in the routing table.
	<i>distance</i>	(Optional) Indicates the administrative distance of the static route. The administrative distance is 1 by default.
	<i>tag</i>	(Optional) Indicates the tag of the static route. The tag is 0 by default.
	permanent	(Optional) Indicates the flag of the permanent route. The static route is not a permanent route by default.
	weight <i>number</i>	(Optional) Indicates the weight of the static route. The weight is 1 by default.
	description <i>descr</i> <i>ption-text</i>	(Optional) Indicates the description of the static route. By default, no description is configured. <i>description-text</i> is a string of one to 60 characters.
	disabled /enabled	(Optional) Indicates the enable flag of the static route. The flag is enabled by default.
global	(Optional) Indicates that the next hop belongs to a global VRF. By default, the VRF of the next hop is the same as the VRF specified by <i>vrf name</i> .	
Defaults	By default, no static default route is configured.	

Command Mode	Global configuration mode
Usage Guide	The simplest configuration of this command is ip route0.0.0.0 0.0.0.0 ip-address .

Configuring the IPv6 Default Gateway on a L3 Switch

Command	ipv6 route [vrfvrf-name] ::/0 { ipv6-address [nexthop-vrf {vrf-name1/ default}] interface [ipv6-address [nexthop-vrf {vrf-name1/ default}]] } [distance] [weightnumber] [descriptiondescription-text]	
Parameter Description	vrfvrf-name	(Optional) Indicates the routing VRF, which must be a multi-protocol VRF of a configured IPv6 address family. The VRF is a global VRF by default.
	::	Indicates the IPv6 prefix, which must comply with the address expression specified in RFC4291.
	0	Indicates the length of the IPv6 prefix. Note that a slash (/) must be added in front of the length.
	ipv6-address	(Optional) Indicates the next-hop address of the static route. You must specify at least one of <i>ipv6-address</i> and <i>interface</i> , or both of them. If <i>ipv6-address</i> is not specified, a static direct route is configured.
	interface	(Optional) Indicates the next-hop exit interface of the static route. You must specify at least one of <i>ipv6-address</i> and <i>interface</i> , or both of them. If <i>interface</i> is not specified, a recursive static direct route is configured. The exit interface is obtained by the next hop in the routing table.
	nexthop-vrf vrf-name1	(Optional) Indicates the routing VRF of the next hop, which must be a multi-protocol VRF of a configured IPv6 address family. By default, the VRF of the next hop is the same as the VRF specified by <i>vrf name</i> . nexthop-vrf default indicates that the VRF of the next shop is a global VRF.
	distance	(Optional) Indicates the administrative distance of the static route. The administrative distance is 1 by default.
	weight number	(Optional) Indicates the weight of the static route, which must be specified when you configure equal-cost routes. The weight ranges

		from 1 to 8. When the weights of all equal-cost routes of a route are summed up, the sum cannot exceed the maximum number of equal-cost routes that can be configured for the route. Weighting of equal-cost routes of a route indicates the traffic ratio of these routes. The weight is 1 by default.
	description <i>description-text</i>	(Optional) Indicates the description of the static route. By default, no description is configured. <i>description-text</i> is a string of one to 60 characters.
Defaults	By default, no static default route is configured.	
Command Mode	Global configuration mode	
Usage Guide	The simplest configuration of this command is ipv6 route ::/0 ipv6-gateway .	

Configuring the IPv4 Default Network on a L3 Switch

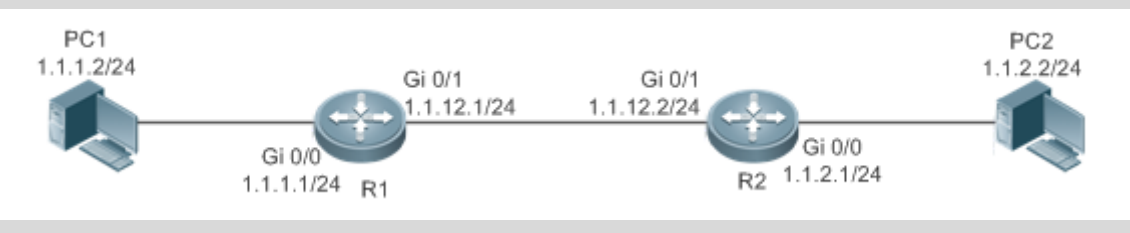
Command	ip default-network <i>network</i>	
Parameter Description	<i>network</i>	Indicates the address of the network. (The network must be a Class A, B, or C network.)
Defaults	By default, no default network is configured.	
Command Mode	Global configuration mode	
Usage Guide	If the network specified by the ip default-network command exists, a default route is generated and the next hop to this network is the default gateway. If the network specified by the ip default-network command does not exist, the default route is not generated.	

Verification

- On a L2 switch (or a L3 switch where routing is disabled), run the **show ip redirects** or **show ipv6 redirects** command to display the default gateway.
- On a L3 switch where routing is enabled, run the **show ip route** or **show ipv6 route** command to display the default route.

Configuration Example

Configuring IPv4 Default Routes on L3 Switches to Implement Network Interworking

<p>Scenario</p> <p>Figure 10-10</p>	
<p>Configuration Steps</p>	<ul style="list-style-type: none"> Configure IP addresses on L3 devices.
<p>R1</p>	<pre>R1#configure terminal R1(config)#interface gigabitEthernet 0/0 R1(config-if-GigabitEthernet 0/0)# ip address 1.1.1.1 255.255.255.0 R1(config-if-GigabitEthernet 0/0)# exit R1(config)#interface gigabitEthernet 0/1 R1(config-if-GigabitEthernet 0/1)# ip address 1.1.12.1 255.255.255.0 R1(config-if-GigabitEthernet 0/0)# exit</pre>
<p>R2</p>	<pre>R2#configure terminal R2(config)#interface gigabitEthernet 0/0 R2(config-if-GigabitEthernet 0/0)# ip address 1.1.2.1 255.255.255.0 R2(config-if-GigabitEthernet 0/0)# exit R2(config)#interface gigabitEthernet 0/1 R2(config-if-GigabitEthernet 0/1)# ip address 1.1.12.2 255.255.255.0 R2(config-if-GigabitEthernet 0/0)# exit</pre>
<p>R1</p>	<ul style="list-style-type: none"> Configure an IPv4 default gateway on R 1. <pre>R1#configure terminal R1(config)#ip route 0.0.0.0 0.0.0.0 GigabitEthernet 0/1 1.1.12.2</pre>
<p>R2</p>	<pre>R2#configure terminal R2(config)#ip route 0.0.0.0 0.0.0.0 GigabitEthernet 0/1 1.1.12.1</pre>

Verification	<ul style="list-style-type: none"> Display the routing table.
R1	<pre> R1# show ip route Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area, * - candidate default Gateway of last resort is 1.1.12.2 S* 0.0.0.0/0 [1/0] via 1.1.12.2, GigabitEthernet 0/1 C 1.1.1.0/24 is directly connected, GigabitEthernet 0/0 C 1.1.1.1/32 is local host. C 1.1.12.0/24 is directly connected, GigabitEthernet 0/1 C 1.1.12.1/32 is local host. </pre>

10.5.3 Configuring Route Limitations

Configuration

Effect

- Limit the number of equal-cost routes and number of static routes, or disable routing.

Notes

Route limitations cannot be configured on a L2 switch.

Configuration

Steps

Configuring the Maximum Number of Equal-Cost Routes

Command	<code>maximum-paths</code> <i>number</i>	
Parameter Description	<i>number</i>	Indicates the maximum number of equal-cost routes. The value ranges from 1 to 64.

Defaults	The default value varies from products.
Command Mode	Global configuration mode
Usage Guide	Run this command to configure the maximum number of next hops in the equal-cost route. In load balancing mode, the number of routes on which traffic is balanced does not exceed the configured number of equal-cost routes.

Configuring the Maximum Number of IPv4 Static Routes

Command	<code>ip static route-limitnumber</code>
Parameter Description	<i>number</i> Indicates the upper limit of routes. The value ranges from 1 to 10,000.
Defaults	By default, a maximum of 1, 024 IP static routes can be configured.
Command Mode	Global configuration mode
Usage Guide	Run this command to configure the maximum number of IPv4 static routes. If the maximum number of IPv4 static routes is reached, no more IPv4 static route can be configured.

Configuring the Maximum Number of IPv6 Static Routes

Command	<code>ipv6 static route-limitnumber</code>
Parameter Description	<i>number</i> Indicates the upper limit of routes. The value ranges from 1 to 10,000.
Defaults	By default, a maximum of 1,000 IPv6 static routes can be configured.
Command Mode	Global configuration mode
Usage Guide	Run this command to configure the maximum number of IPv6 static routes. If the maximum number of IPv6 static routes is reached, no more IPv6 static route can be configured.

Disabling IPv4 Routing

Command	no ip routing
Parameter Description	N/A
Defaults	By default, IPv4 routing is enabled.
Command Mode	Global configuration mode
Usage Guide	Run this command to disable IPv4 routing. If the device functions only as a bridge or a voice over IP (VoIP) gateway, the device does not need to use the IPv4 routing function of the RGOS software. In this case, you can disable the IPv4 routing function of the RGOS software.

Disabling IPv6 Routing

Command	no ipv6 unicast-routing
Parameter Description	N/A
Defaults	By default, IPv6 routing is enabled.
Command Mode	Global configuration mode
Usage Guide	Run this command to disable IPv6 routing. If the device functions only as a bridge or a VoIP gateway, the device does not need to use the IPv6 routing function of the RGOS software. In this case, you can disable the IPv6 routing function of the RGOS software.

Prohibiting Static Routing Across VRFs

Command	no ip route static inter-vrf
Parameter Description	N/A
Defaults	By default, static IP or IPv6 routing across VRFs is allowed.

Command Mode	Global configuration mode
Usage Guide	Run this command to prohibit static IP routing across VRFs. After this command is configured, the static IP route across VRFs is not active and cannot participate in packet forwarding.

Verification

Run the **show run** command to display the configuration file and verify that the preceding configuration commands exist.

Configuration

Example

Configuring at Most Two Static Routing Limitations

<p>Scenario Figure 10-6</p>	
<p>Configuration Steps</p>	<p>On R 1, configure the IP addresses, static routes, and maximum number of static routes.</p>
	<pre> R1#configure terminal R1(config)#interface gigabitEthernet 0/0 R1(config-if-GigabitEthernet 0/0)# ip address 1.1.1.1 255.255.255.0 R1(config-if-GigabitEthernet 0/0)# exit R1(config)#interface gigabitEthernet 0/2 R1(config-if-GigabitEthernet 0/2)# ip address 1.1.12.1 255.255.255.0 R1(config-if-GigabitEthernet 0/0)# exit </pre>

	<pre>R1(config)#interface gigabitEthernet 0/3 R1(config-if-GigabitEthernet 0/3)# ip address 1.1.13.1 255.255.255.0 R1(config-if-GigabitEthernet 0/3)# exit R1(config)#ip route 1.1.3.0 255.255.255.0 1.1.13.3 R1(config)#ip route 1.1.4.0 255.255.255.0 1.1.12.2 R1(config)#ip route 1.1.5.0 255.255.255.0 1.1.12.2 R1(config)#ip static route-limit 2 % Exceeding maximum static routes limit.</pre>
Verification	<ul style="list-style-type: none"> ▪ Check the static routes that really take effect in the routing table.
	<pre>R1(config)# show ip route Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area, * - candidate default Gateway of last resort is no set C 1.1.1.0/24 is directly connected, GigabitEthernet 0/0 C 1.1.1.1/32 is local host. S 1.1.3.0/24 [1/0] via 1.1.13.3 S 1.1.4.0/24 [1/0] via 1.1.12.2 C 1.1.12.0/24 is directly connected, GigabitEthernet 0/2 C 1.1.12.1/32 is local host. C 1.1.13.0/24 is directly connected, GigabitEthernet 0/3 C 1.1.13.1/32 is local host.</pre>

10.5.4 Correlating a Static Route with BFD

Configuration Effect

- A static route can quickly detect a route failure with the help of BFD.

Notes

- BFD correlation cannot be configured on a L2 switch.
- You must configure a static route.
- You must configure the BFD session parameters by running the **bfd interval xmin_rx xmultiplier x** command.

Configuration Steps

Correlating an IPv4 Static Route with BFD

Command	ip route static bfd [vrf <i>vrf-name</i>]<i>interface-type interface-number gateway</i>[<i>sourceip-address</i>]	
Parameter Description	vrf <i>vrf-name</i>	(Optional) Indicates the name of the VRF to which the static route belongs. The VRF is a global VRF by default.
	<i>interface-type</i>	Indicates the interface type.
	<i>interface-number</i>	Indicates the interface number.
	<i>gateway</i>	Indicates the IP address of the gateway, that is, the neighbor IP address of BFD. If the next hop of the static route is this neighbor, BFD is used to check the connectivity of the forwarding path.
	source <i>ip-address</i>	(Optional) Indicates the source IP address used for the BFD session. This parameter must be configured if the neighbor IP address involves multiple hops. By default, the source IP address is not specified.
Defaults	By default, a static route is not correlated with BFD.	
Command Mode	Global configuration mode	
Usage Guide	Run this command to correlate an IPv4 static route with BFD. If the down status of the BFD session is detected, the IPv4 static route is not active and does not participate in packet forwarding.	

Correlating an IPv6 Static Route with BFD

Command	ipv6 route static bfd [<i>vrfvrf-name</i>] <i>interface-type interface-number gateway</i> [<i>source ipv6-address</i>]
---------	---

Parameter Description	vrfvrf-name	(Optional) Indicates the name of the VRF to which the static route belongs. The VRF is a global VRF by default.
	interface-type	Indicates the interface type.
	interface-number	Indicates the interface number.
	gateway	Indicates the IP address of the gateway, that is, the neighbor IP address of BFD. If the next hop of the static route is this neighbor, BFD is used to check the connectivity of the forwarding path.
	sourceip-address	(Optional) Indicates the source IP address used for the BFD session. This parameter must be configured if the neighbor IP address involves multiple hops. By default, the neighbor IP address of the BFD session is a single hop, and the source IP address is not used.
Defaults	By default, a static route is not correlated with BFD.	
Command Mode	Global configuration mode	
Usage Guide	Run this command to correlate an IPv6 static route with BFD. If the down status of the BFD session is detected, the IPv6 static route is not active and does not participate in packet forwarding.	

Verification

- Run the `show bfd neighbors` command to display information about BFD neighbors.
- Run the `show ip route static bfd` or `show ipv6 route static bfd` command to display information about correlation of static routes with BFD.

Configuration

Example

Correlating an IPv4 Static Route with BFD

Scenario Figure 10-7	
Configuration Steps	<ul style="list-style-type: none"> Configure a BFD session on the interconnect interface between R 1 and R 2.

	<ul style="list-style-type: none"> Configure static routes on R 1 and R 2, and specify the exit interface or next hop as the interworking interface. Correlate static routes with BFD on R 1 and R 2, and check the connectivity of the next hops of static routes.
R1	<pre>R1#configure terminal R1(config)#interface gigabitEthernet 0/1 R1(config-if-GigabitEthernet 0/1)# no switchport R1(config-if-GigabitEthernet 0/1)# ip address 1.1.12.1 255.255.255.0 R1(config-if-GigabitEthernet 0/1)#bfd interval 50 min_rx 50 multiplier 3 R1(config-if-GigabitEthernet 0/1)# exit R1(config)# ip route 1.1.2.0 255.0.0.0 FastEthernet 0/1 1.1.12.2 R1(config)#ip route static bfd gigabitEthernet 0/1 1.1.12.2</pre>
R2	<pre>R2#configure terminal R1(config)#interface gigabitEthernet 0/1 R1(config-if-GigabitEthernet 0/1)# no switchport R1(config-if-GigabitEthernet 0/1)# ip address 1.1.12.2 255.255.255.0 R1(config-if-GigabitEthernet 0/1)#bfd interval 50 min_rx 50 multiplier 3 R1(config-if-GigabitEthernet 0/1)# exit R1(config)# ip route 1.1.1.0 255.0.0.0 FastEthernet 0/1 1.1.12.1 R1(config)#ip route static bfd gigabitEthernet 0/1 1.1.12.1</pre>
Verification	<ul style="list-style-type: none"> Display the status of BFD neighbors. Display the static routes correlated with BFD.
R1	<pre>R1#show bfd neighbors OurAddr NeighAddr LD/RD RH/RS Holdown(mult) State Int 1.1.12.1 1.1.12.2 8192/0 Up 0(3) Up GigabitEthernet 0/1 R1#show ip route static bfd S 1.1.2.0/24 via 1.1.12.2, GigabitEthernet 0/1, BFD state is Up</pre>

Common Errors

- The link on the interface is not up.

- No IP address is configured for the interface.
- No BFD session parameters are configured.
- No static route is configured.

10.5.5 Configure Static Fast Reroute

Configuration

Effect

- Configure and enable static fast reroute.

Notes

- Static fast reroute cannot be configured on a L2 switch.
- You must configure a static route.
- You must configure a route map.

Configuration

Steps

Defining a Standby Route in the Route Map

Command	set fast-reroute backup-nexthop <i>interface ip-address</i>	
Parameter Description	<i>interface</i>	Indicates the standby exit interface.
	<i>ip-address</i>	Indicates the standby next hop.
Defaults	N/A	
Command Mode	Global configuration mode	
Usage Guide	<p>Run the route-map <i>name</i> [permit deny] <i>sequence</i> command to create a road map.</p> <p>Run the match command to define matching conditions.</p> <p>Run the set fast-reroute backup-nexthop <i>interface ip-address</i> command to define the standby exit interface and standby next hop.</p> <p>If a route meets matching conditions, a standby route is generated for this route. If the match command is not configured, standby routes are generated for any static route with the exit interface and next hop.</p>	

Enabling Fast Reroute and Referencing the Route Map

Command	ip fast-reroute [vrf <i>vrf-name</i>] static route-map <i>route-map-name</i>	
Parameter Description	<i>vrf-name</i>	(Optional) Specifies a VRF. If the VRF is not specified, the command is executed on all VRFs.
	<i>route-map-name</i>	Indicates the name of the road map for the standby route.
Defaults	By default, static fast reroute is not configured.	
Command Mode	Global configuration mode	
Usage Guide	Run this command to enable fast reroute and reference the route map.	

Verification

Run the **show ip route fast-reroute** command to display the active and standby routes that take effect.

Configuration

Example

Configuring Fast Re-Routing

<p>Scenario Figure 10-8</p>	
<p>Configurati on Steps</p>	<p>On R 1, configure a static route to the network segment of PC 3, and the next hop of the exit interface is R 3.</p>

	<p>On R 1, configure static fast reroute. The next hop of the exit interface of the standby route is R2.</p>
	<pre> R1#configure terminal R1(config)#interface gigabitEthernet 0/0 R1(config-if-GigabitEthernet 0/0)# ip address 1.1.1.1 255.255.255.0 R1(config-if-GigabitEthernet 0/0)# exit R1(config)#interface gigabitEthernet 0/2 R1(config-if-GigabitEthernet 0/2)# ip address 1.1.12.1 255.255.255.0 R1(config-if-GigabitEthernet 0/0)# exit R1(config)#interface gigabitEthernet 0/3 R1(config-if-GigabitEthernet 0/3)# ip address 1.1.13.1 255.255.255.0 R1(config-if-GigabitEthernet 0/3)# exit R1(config)# ip route 1.1.3.0 255.255.255.0 GigabitEthernet 0/3 1.1.13.3 R1(config)#route-map fast-reroute R1(config-route-map)# set fast-reroute backup-interface GigabitEthernet 0/2 backup- nextHop 1.1.12.2 R1(config-route-map)# exit R1(config)#ip fast-reroute static route-map fast-reroute </pre>
<p>Verification</p>	<p>Display the active and standby routes on R 1.</p> <pre> R1#show ip route fast-reroute Codes: C - connected, S - static, R - RIP, B - BGP O - OSPF, IA - OSPF inter area N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area, * - candidate default Status codes: m - main entry, b - backup entry, a - active entry Gateway of last resort is no set S 1.1.3.0/24 [ma] via 1.1.13.3, GigabitEthernet 0/3 [b] via 1.1.12.2, GigabitEthernet 0/2 </pre>

Common Errors

- The link on the interface is not up.
- No static route is configured.
- The matching conditions are not configured or are not properly configured in the road map.

10.5.6 Configuring ARP-to-host Conversion

Configuration Effect

- Enable ARP-to-host conversion on a specified interface.

Notes

- This configuration cannot be performed on a Layer-2 switch.
- The interface must be a Layer-3 interface in up status.

Configuration Steps

Enabling ARP-to-host Conversion on a Specified Interface

Command	<code>ip route arp-to-host interface <i>interface-name</i></code>
Parameter Description	<i>interface-name</i> : Indicates an interface name.
Defaults	N/A
Command Mode	Global configuration mode
Usage Guide	This command is used to convert ARP entries into 32-bit host routes on a specified interface.

Enabling ARP Proxy

Command	<code>local-proxy-arp</code>
---------	------------------------------

Parameter Description	N/A
Defaults	N/A
Command Mode	Interface configuration mode
Usage Guide	After receiving an ARP request, the device uses a proxy to send an ARP response with its own Ethernet MAC address. In this case, hosts communicate with each other through Layer-3 routes.

Configuring BGP Redistribution of ARP Host Routes

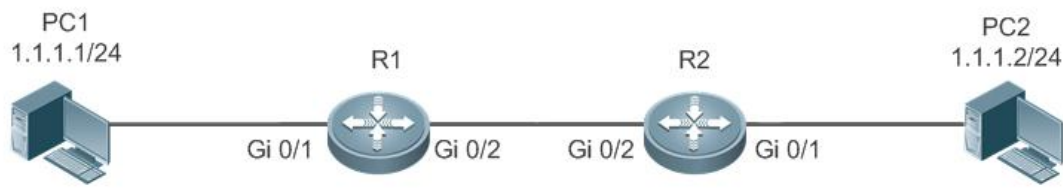
Command	redistribute arp-host [route-map <i>map-tag</i>] [metric <i>metric-value</i>]
Parameter Description	route-map <i>map-tag</i> : Indicates the name of an associated route map. No route map is associated by default. metric <i>metric-value</i> : Indicates the default metric value for redistributing routes, in the range from 0 to 4,294,967,295 . The value is not set by default.
Defaults	N/A
Command Mode	BGP configuration mode
Usage Guide	This command is used to configure BGP redistribution of ARP host routes.

Verification

Run the **show ip route arphost** command to display the generated host routes.

Configuration Example

Configuring ARP-to-host Conversion

<p>Scenario Figure 10-9</p>	
<p>Configurati on Steps</p>	<p>Configure interface addresses on R1 and R2. Enable ARP-to-host conversion on specified interfaces of R1 and R2, configure a BGP neighbor relationship, and redistribute ARP host routes.</p>
<p>R1</p>	<pre> R1#configure terminal R1(config)#interface vlan 40 R1(config-if-VLAN 40)#ip address 1.1.1.3 24 R1(config-if-VLAN 40)#local-proxy-arp R1(config-if-VLAN 40)#exit R1(config)#interface gigabitEthernet 0/1 R1(config-if-GigabitEthernet 0/1)#switchport R1(config-if-GigabitEthernet 0/1)#switchport mode trunk R1(config-if-GigabitEthernet 0/1)#switchport trunk allowed vlan only 40 R1(config-if-GigabitEthernet 0/1)#exit R1(config)#interface gigabitEthernet 0/2 R1(config-if-GigabitEthernet 0/2)# ip address 10.10.10.1 255.255.255.0 R1(config-if-GigabitEthernet 0/2)# exit R1(config)#router bgp 1 R1(config-router)#neighbor 10.10.10.2 remote-as 1 R1(config-router)#redistribute arp-host R1(config-router)#address-family ipv4 R1(onfig-router-af)#neighbor 10.10.10.2 next-hop-self R1(onfig-router-af)#exit R1(config-router)#exit R1(config)#ip route arp-to-host vlan 40 R1(config)#ip route arp-to-host delay-time 1 </pre>
<p>R2</p>	<pre> R2#configure terminal R2(config)#interface vlan 40 R2(config-if-VLAN 40)#ip address 1.1.1.4 24 R2(config-if-VLAN 40)#local-proxy-arp R2(config-if-VLAN 40)#exit R2(config)#interface gigabitEthernet 0/1 R2(config-if-GigabitEthernet 0/1)# switchport R2(config-if-GigabitEthernet 0/1)#switchport mode trunk R2(config-if-GigabitEthernet 0/1)# switchport trunk allowed vlan only 40 </pre>

	<pre> R2(config-if-GigabitEthernet 0/1)#exit R2(config)#interface gigabitEthernet 0/2 R2(config-if-GigabitEthernet 0/2)# ip address 10.10.10.2 255.255.255.0 R2(config-if-GigabitEthernet 0/2)# exit R2(config)#router bgp 1 R2(config-router)#neighbor 10.10.10.1 remote-as 1 R2(config-router)#redistribute arp-host R2(config-router)#address-family ipv4 R2(onfig-router-af)#neighbor 10.10.10.1 next-hop-self R2(onfig-router-af)#exit R2(config-router)#exit R2(config)#ip route arp-to-host vlan 40 R2(config)#ip route arp-to-host delay-time 1 </pre>
<p>Verification</p>	<p>Display the ARP host routes of R1 to check that PC1 can ping PC2.</p> <pre> R1(config)#show ip route arphost A 1.1.1.1/32 [1/0] via 1.1.1.1, gigabitEthernet 0/1 R1#show ip route Codes: C - Connected, L - Local, S - Static R - RIP, O - OSPF, B - BGP, I - IS-IS, V - Overflow route N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 E1 - OSPF external type 1, E2 - OSPF external type 2 SU - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 IA - Inter area, EV - BGP EVPN, A - Arp to host * - candidate default Gateway of last resort is 0.0.0.0 to network 0.0.0.0 S* 0.0.0.0/0 is directly connected, Null 0 C 10.10.10.0/24 is directly connected, GigabitEthernet 0/2 C 10.10.10.1/32 is local host. C 1.1.1.0/24 is directly connected, VLAN 40 A 1.1.1.1/32 [1/0] via 1.1.1.1, VLAN 40 C 1.1.1.3/32 is local host. B 1.1.1.2/32 [200/0] via 10.10.10.2, 00:57:12 </pre>

Common Errors

- The ARP proxy function is not enabled on specified interfaces.

10.6 Monitoring

Displaying

Description	Command
Displays the IPv4 routing table.	show ip route
Displays the IPv6 routing table.	show ipv6route

Debugging

! System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Debugs IPv4 route management.	debug nsm kernel ucast- v4
Debugs IPv6 route management.	debug nsm kernel ucast-v6
Debugs fast reroute management.	debug nsm kernel fr
Debugs default network management.	debug nsm kernel default-network
Debugs internal events of route management.	debug nsm events
Debugs sending of route management and routing protocol messages.	debug nsm packet send
Debugs receiving of route management and routing protocol messages.	debug nsm packet recv

11 CONFIGURING VRF

11.1 Overview

A Virtual Private Network (VPN) Routing and Forwarding (VRF) table is used for the forwarding of VPN packets. Each VPN corresponds to a VRF table.

A device that provides the VPN service has multiple routing tables, including a public network routing table and one or multiple VRF tables. The public-network routing table is used for the forwarding of public network packets, and the VRF tables are used for the forwarding of VPN packets. These routing tables are created to separate routes in the public network from those in VPNs and separate routes in different VPNs.

- i** A VPN is a private dedicated network built in the public network. "Virtual" means that the VPN is logically exclusive, instead of physically exclusive.

11.2 Applications

Application	Description
Local Inter-VPN Access	Provide the VPN service on a routing device and enable VPNs to access each other.
VRF only on Provider Edges (PEs)	Provide the VPN service in an IP network and connect one Customer Edge (CE) to one VPN.
VRF on CEs and PEs	Provide the VPN service in an IP network and connect one CE to multiple VPNs.

- i** CE: An edge device in a customer network
- i** PE: An edge device in a Service Provider (SP) network


11.2.1 Local Inter-VPN Access

Scenario

Provide the VPN service on a routing device and enable VPNs to access each other.

In Figure 11-1, Sub a runs the Routing Information Protocol (RIP), Sub b runs the Open Shortest Path First (OSPF) protocol, and Servers is a network segment directly connected to C. Provide the VPN service on C to Sub a, Sub b, and Servers, and enable Sub a and Sub b to access Servers.

Figure 11-1



Related Configuration

- On C, create a VRF table for Sub a, bind the interface directly connected to A, and associate the VRF table with A by using RIP.
- On C, create a VRF table for Sub b, bind the interface directly connected to B, and associate the VRF table with B by using OSPF.
- On C, create a VRF table for Servers and bind the interface directly connected to Servers.
- On C, configure the route targets (RTs) of the VRF tables for Suba, Subb, and Servers. Import the routes in the VRF tables for Sub a and Sub b to the VRF table for Servers, and import the routes in the VRF table for Servers to the VRF tables for Sub a and Sub b.
- Configure the Border Gateway Protocol (BGP) on C. Introduce the RIP routes to the VRF table for Sub a, introduce the OSPF routes to the VRF table for Sub b, and introduce the direct routes to the VRF table for Servers.

11.2.2 VRF only on PEs


Scenario

An Internet Service Provider (ISP) provides the VPN service in an IP backbone network.

In Figure 11-2, VPN1 runs RIP, and VPN2 runs OSPF.

- One CE is connected to one VPN, and all routes on the CE are exclusively used by the connected VPN. Therefore, no VRF table needs to be created to separate the routes.
- On each PE, VRF tables must be created to separate the routes in VPN1, those in VPN2, and those in the public network from each other.

Figure 11-2



Deployment

- On PE1, create a VRF table for VPN1 and bind the interface directly connected to CE1. On PE2, create a VRF table for VPN1 and bind the interface directly connected to CE3.
- On PE1, create a VRF table for VPN2 and bind the interface directly connected to CE2. On PE2, create a VRF table for VPN2 and bind the interface directly connected to CE4.
- On PE1, associate the VRF table for VPN1 with CE1 by using RIP. On PE2, associate the VRF table for VPN1 with CE3 by using RIP.
- On PE1, associate the VRF table for VPN2 with CE2 by using OSPF. On PE2, associate the VRF table for VPN2 with CE4 by using OSPF.
- In the VRF instance for VPN1 on PE1, redistribute RIP routes to BGP, and redistribute BGP routes to RIP. The configuration on PE2 is similar.
- In the VRF instance for VPN2 on PE1, redistribute OSPF routes to BGP, and redistribute BGP routes to OSPF. The configuration on PE2 is similar.

11.2.3 VRF on CEs and PEs (MCE Application)

Scenario

An ISP provides the VPN service in an IP backbone network.

In Figure 11-3, VPN a runs RIP, VPN b runs OSPF, and PE1 and PE2 are connected to BGP.

- One Multi-VPN-Instance CE (MCE) is connected to multiple VPNs. VRF tables must be created to separate the routes in VPN a from those in VPN b.
- On each PE, VRF tables must be created to separate the routes in VPN a, those in VPN b, and those in the public network from each other.

Figure 11-3



Deployment

- One MCE1, create VRF tables for VPN a and VPN b respectively, bind the interfaces directly connected to VPN a and VPN b, and bind the VLAN interface connected to PE1. The configuration on MCE2 is similar.
- On PE1, create VRF tables for VPN a and VPN b respectively, and bind the VLAN interface connected to MCE1. The configuration on PE2 is similar.
- On MCE1, associate the VRF table for VPN a with VPN a by using RIP. The configuration on MCE2 is similar.
- On MCE1, associate the VRF table for VPN b with VPN b by using OSPF. The configuration on MCE2 is similar.
- In the VRF instance for VPN a on MCE1, redistribute RIP routes to BGP, and redistribute BGP routes to RIP. The configuration on MCE2 is similar.
- In the VRF instance for VPN b on MCE1, redistribute OSPF routes to BGP, and redistribute BGP routes to OSPF. The configuration on MCE2 is similar.

11.3 Features

Overview

Feature	Description
VPN Instance	A VPN instance is used to provide the VPN service. It is typically represented by a VRF table.
VPN Route	A VPN route is used to forward VPN packets.
VPN Route Attribute	Route distinguisher (RD): Identifies the VPN to which a route belongs. RT: Indicates the route trade-off mode of VRF.

11.3.1 VPN Instance

A VPN instance is used to provide the VPN service. On a device that provides the VPN service, a VPN instance consists of the VRF table, interfaces, routing protocol processes, and configuration that belong to the same VPN. A VPN instance is typically represented by a VRF table.

Working Principle

A PE exchanges routes with a CE by using the related routing protocol in the corresponding VPN instance. A VRF table is bound to a specific interface to generate its interface set. Packets received on these interfaces will be associated with the VRF table and forwarded along corresponding routes.

Related Configuration

! Single-protocol VRF tables and multiprotocol VRF tables cannot be created at the same time. Single-protocol VRF tables only support IPv4, whereas multiprotocol VRF tables support IPv4 and IPv6.

Configuring a Single-Protocol VRF Table

By default, a device has no VRF table.

Run the **ip vrf** command to create a single-protocol VRF table.

Run the **ip vrf forwarding** command to bind an interface.

Currently, single-protocol VRF tables only support IPv4.

Configuring a Multiprotocol VRF Table

By default, a device has no VRF table.

Run the **vrf definition** command to create a multiprotocol VRF table.

Run the **address-family ipv4** command to enable the IPv4 address family.

Run the **address-family ipv6** command to enable the IPv6 address family.

Run the **vrf forwarding** command to bind an interface.

Multiprotocol VRF tables support IPv4 and IPv6.

11.3.2 VPN Route

A VPN route is only used to forward VPN packets. It comes from:

- Direct route and host route on the bound interface
- Direct route and host route on the configured import interface (not bound)
- Static and dynamic routes (RIP, RIPng, OSPFv2, OSPFv3, ISIS, and BGP) in the configured VPN instance

i For details about the static routes in a VPN instance, see "Configuration Guide > IP Route".


i For details about RIP in a VPN instance, see "Configuration Guide > IP Route > RIP".

i For details about RIPng in a VPN instance, see "Configuration Guide > IP Route > RIPng".

i For details about OSPFv2 in a VPN instance, see "Configuration Guide > IP Route > OSPFv2".

i For details about OSPFv3 in a VPN instance, see "Configuration Guide > IP Route > OSPFv3".

i For details about ISIS in a VPN instance, see "Configuration Guide > IP Route > ISIS".

 For details about BGP in a VPN instance, see "Configuration Guide > IP Route > BGP".

11.3.3 VPN Route Attribute

The BGP extended attributes include two attributes specific to VPN routes: RD and RT.

Working Principle

RD

Two routes with the same address but different RDs in two VRF tables can be advertised separately between PEs, because the routes are sent together with their RDs through multiprotocol BGP (MP-BGP).

RT

RT in essence indicates each VRF table's route trade-off and preferences. It is mainly used to control the advertising and installation policies for VPN routes. RT is divided into the import attribute and export attribute. The import attribute indicates the route of interest, and the export attribute indicates the advertised route. A PE advertises a route to other PEs based on the RT export rule in the corresponding VRF table. The peer PE checks all received routes against the RT import rule in each VRF table. If a route matches an RT export rule (the export rule contains the import rule), it will be added to the corresponding VRF table.

Related Configuration

RD

By default, no RD is configured in VRF mode.


Run the **rd** command to configure an RD.


RT

By default, no RT is configured in VRF mode or address family mode.

Run the **route-target { import | export | both }** command to configure an RT.

11.4 Configuration

Configuration	Description and Command
Configuring a Single-Protocol VRF Table	 Single-protocol VRF tables and multiprotocol VRF tables cannot be created at the same time. If IPv6 is supported, configure a multiprotocol VRF table; otherwise, you can configure a single-protocol VRF table or a

	multiprotocol VRF table. This configuration item creates a VRF table in an IPv4 network. IPv6 is not supported.	
	ip vrf <i>vrf-name</i>	Creates a VRF table.
	rd <i>rd_value</i>	Configures an RD.
	route-target { import export both } <i>rt_value</i>	Configures an RT.
	ip vrf forwarding <i>vrf-name</i>	Binds an interface and adds the direct route and host route on the interface to a VRF table.
	ip vrf receive <i>vrf_name</i>	Adds the direct route and host route on an interface to a VRF table without binding the interface.
Configuring a Multiprotocol VRF Table	 Single-protocol VRF tables and multiprotocol VRF tables cannot be created at the same time. If IPv6 is supported, configure a multiprotocol VRF table. otherwise, you can configure a single-protocol VRF table or a multiprotocol VRF table. This configuration item creates a VRF table in an IPv4 or IPv6 network.	
	vrf definition <i>vrf-name</i>	Creates a VRF table.
	description <i>string</i>	Configures a VRF descriptor.
	rd <i>rd_value</i>	Configures an RD.
	route-target { import export both } <i>rt_value</i>	Configures an RT.
	address-family ipv4	Enables the IPv4 address family.
	address-family ipv6	Enables the IPv6 address family.
	vrf forwarding <i>vrf-name</i>	Binds an interface and adds the direct route and host route on the interface to a VRF table.
	vrf receive <i>vrf-name</i>	Adds the direct route and host route on an interface to a VRF table without binding the interface.

11.4.1 Configuring a Single-Protocol VRF Table

Configuration

Effect

- Provide the VPN service on a device.
- With BGP assistance, flexibly control the separation and access between VPNs.
- With BGP assistance, provide the VPN service in an IP backbone network.
- Only IPv4 is supported.

Notes

- No VRF table needs to be created if the device only forwards packets from one VPN or from the public network.
- If the device needs to forward public network packets and VPN packets or forward packets from multiple VPNs, VRF tables must be created to separate routes.
- In many cases, static or dynamic routes (RIP, OSPF, ISIS, and BGP) need to be added to VRF tables.

Configuration

Steps

Creating a VRF Table

- Mandatory.
- Create a VRF table for each VPN.

Configuring an RD

- Optional.
- When routing information needs to be advertised through BGP in the backbone network, BGP may select the best route for advertising if overlapping network addresses exist in different VPNs, which will make some VPNs fail to obtain corresponding routing information. To solve this problem, you can configure RDs for routes to enable BGP to make routing decisions based on these RDs, thus ensuring that each VPN can obtain corresponding routing information.
- Run the **rd** command in single-protocol VRF mode.

Configuring an RT

- Optional.
- You can run the **route-target export** command to specify the attributes of the route to be advertised, and run the **route-target import** command to specify the attributes of the route to be received. You can also run the **route-target both** command to specify the export and import attributes.
- Run the **route-target** command in single-protocol VRF mode.

Binding an Interface and Adding the Direct Route and Host Route on the Interface to a VRF Table

- Mandatory.

- If the physical link for transmitting VPN packets is exclusively occupied by a VPN, bind the physical interface to the corresponding VRF table.
- If the physical link for transmitting VPN packets is shared by multiple VPNs, you need to create an independent logical link for each VPN, and bind the logical interface to the corresponding VRF table. A logical interface can be a subinterface or a VLAN interface.
- You must bind an interface to the corresponding single-protocol VRF table before you configure the IPv4 address of the interface. If you bind the interface after its IPv4 address is configured, the IPv4 address will be invalid (the IPv6 address of the interface is retained).
- If you bind an interface to the corresponding single-protocol VRF table and enable IPv6 on the interface, the device cannot forward the IPv6 packets received on the interface.

Adding the Direct Route and Host Route on an Interface to a VRF Table Without Binding the Interface

- Optional.
- If policy-based routing (PBR) is required for VRF table selection, run the **ip vrf receive** command on the interface to which PBR is applied, and import the direct route and host route on the interface to each VRF table available for choice.

Verification

- Check whether VRF tables are created correctly on the router.

Related Commands

Creating a VRF Table

Command	ip vrf <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the name of the VRF table to be created. It cannot exceed 31 characters.
Command Mode	Global configuration mode
Usage Guide	After you run the command, the system will enter VRF mode.

Configuring an RD

Command	rd <i>rd_value</i>
Parameter Description	<i>rd_value</i> has the following three different parameter forms: (1) <i>rd_value</i> = as_num: nn

	<p>as_num indicates the 2-byte number that identifies a public autonomous system (AS). nn is configurable in the range 0..4294967295.</p> <p>(2) <i>rd_value</i> = ip_addr: nn</p> <p>ip_addr must be a global IP address. nn is configurable in the range 0..65535.</p> <p>(3) <i>rd_value</i> = as4_num: nn</p> <p>as4_num indicates the 4-byte number that identifies a public AS. nn is configurable in the range 1..65535.</p>
Command Mode	VRF configuration mode
Usage Guide	<p>You cannot directly change the RD of an existing VRF table. You need to delete the VRF table first and then configure a new RD.</p> <p>A VRF table has only one RD. You cannot configure multiple RDs for one VRF table.</p>

Configuring an RT

Command	route-target { import export both } <i>rt_value</i>
Parameter Description	<p><i>rt_value</i> has the following three different parameter forms:</p> <p>(1) <i>rt_value</i> = as_num: nn</p> <p>as_num indicates the 2-byte number that identifies a public AS. nn is configurable in the range 0..4294967295.</p> <p>(2) <i>rt_value</i> = ip_addr: nn</p> <p>ip_addr must be a global IP address. nn is configurable in the range 0..65535.</p> <p>(3) <i>rt_value</i> = as4_num: nn</p> <p>as4_num indicates the 4-byte number that identifies a public AS. nn is configurable in the range 1..65535.</p>
Command Mode	VRF configuration mode
Usage Guide	A VRF table can be configured with multiple import and export RT attributes.

Binding an Interface

Command	ip vrf forwarding <i>vrf-name</i>
----------------	--

Parameter Description	<i>vrf-name</i> : Indicates the name of a VRF table.
Command Mode	Interface configuration mode
Usage Guide	<p>By default, an interface does not belong to any VRF table.</p> <p>After an interface is bound to the corresponding VRF table, the direct route and host route on the interface will be automatically added to the VRF table.</p> <p>You must bind an interface to the corresponding single-protocol VRF table before you configure the IPv4 address of the interface. If you bind the interface after its IPv4 address is configured, the IPv4 address will be invalid (the IPv6 address of the interface is retained).</p>

Adding the Direct Route and Host Route on an Interface to a VRF Table Without Binding the Interface

Command	ip vrf receive <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the name of a VRF table.
Command Mode	Interface configuration mode
Usage Guide	<p>This command is used to add the host route and direct route on an interface to a VRF table. If you need to add the host route and direct route on an interface to multiple VRF tables, run the command multiple times.</p> <p>Different from the ip vrf forwarding command, the ip vrf receive command does not bind an interface to the corresponding VRF table. The interface is still a global interface and does not belong to any VRF table.</p> <p>The ip vrf forwarding and ip vrf receive commands are mutually exclusive on the same interface.</p>

Displaying the VRF Information on a Device

Command	show ip vrf [<i>brief</i> <i>detail</i> <i>interfaces</i>] [<i>vrf-name</i>]
Parameter Description	<p>brief: Displays brief information.</p> <p>detail: Displays detailed information.</p>

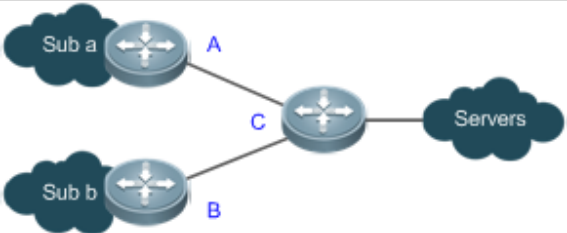
	interfaces: Displays the interface binding information. <i>vrf-name:</i> Indicates the name of a VRF table.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	This command is used to display the information of a specified VRF table to check whether the VRF table is bound with the correct interface.

Displaying the Routes in a VRF Table

Command	show ip route vrf <i>vrf-name</i>
Parameter Description	<i>vrf-name:</i> Indicates the name of a VRF table.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	This command is used to check whether a specified VRF table contains corresponding routes.

Configuration Example

Local Inter-VPN Access

Scenario Figure 11-4	
	Sub a, Sub b, and Servers are three VPNs that have separate address spaces. Sub a runs RIP, Sub b runs OSPF, and Servers is a network segment directly connected to C.
Configuration	Routes in Sub a are separated from those in Sub b, but both Sub a and Sub b can access Servers.

Requirements																									
Configuration Steps	<ul style="list-style-type: none"> ▪ On C, create a VRF table for Sub a, bind the interface directly connected to A, and associate the VRF table with A by using RIP. ▪ On C, create a VRF table for Sub b, bind the interface directly connected to B, and associate the VRF table with B by using OSPF. ▪ On C, create a VRF table for Servers and bind the interface directly connected to Servers. ▪ On C, configure the RTs of the VRF tables for Sub a, Sub b, and Servers. Import the routes in the VRF tables for Sub a and Sub b to the VRF table for Servers, and import the routes in the VRF table for Servers to the VRF tables for Sub a and Sub b. ▪ Configure the Border Gateway Protocol (BGP) on C. Introduce the RIP routes to the VRF table for Sub a, introduce the OSPF routes to the VRF table for Sub b (enabled with an address family), and introduce the direct routes to the VRF table for Servers (enabled with an address family). <p>i Planning of interfaces and addresses:</p> <table border="1" data-bbox="354 894 1490 1495"> <thead> <tr> <th>Interface Description</th> <th>Interface Name</th> <th>IP Address/Mask</th> <th>VRF Table</th> </tr> </thead> <tbody> <tr> <td>Interface on C connected to A</td> <td>GE0/1</td> <td>10.10.1.1/24</td> <td>VRF table for Sub a</td> </tr> <tr> <td>Interface on C connected to B</td> <td>GE0/2</td> <td>10.10.2.1/24</td> <td>VRF table for Sub b</td> </tr> <tr> <td>Interface on C connected to Servers</td> <td>GE0/3</td> <td>10.10.3.1/24</td> <td>VRF table for Servers</td> </tr> <tr> <td>Interface on A connected to C</td> <td>GE0/1</td> <td>10.10.1.2/24</td> <td>-</td> </tr> <tr> <td>Interface on B connected to C</td> <td>GE0/2</td> <td>10.10.2.2/24</td> <td>-</td> </tr> </tbody> </table>	Interface Description	Interface Name	IP Address/Mask	VRF Table	Interface on C connected to A	GE0/1	10.10.1.1/24	VRF table for Sub a	Interface on C connected to B	GE0/2	10.10.2.1/24	VRF table for Sub b	Interface on C connected to Servers	GE0/3	10.10.3.1/24	VRF table for Servers	Interface on A connected to C	GE0/1	10.10.1.2/24	-	Interface on B connected to C	GE0/2	10.10.2.2/24	-
Interface Description	Interface Name	IP Address/Mask	VRF Table																						
Interface on C connected to A	GE0/1	10.10.1.1/24	VRF table for Sub a																						
Interface on C connected to B	GE0/2	10.10.2.1/24	VRF table for Sub b																						
Interface on C connected to Servers	GE0/3	10.10.3.1/24	VRF table for Servers																						
Interface on A connected to C	GE0/1	10.10.1.2/24	-																						
Interface on B connected to C	GE0/2	10.10.2.2/24	-																						
A	<pre> A(config)#interface GigabitEthernet 0/1 A(config-if-GigabitEthernet 0/1)#no switchport port A(config-if-GigabitEthernet 0/1)#ip address 10.10.1.2 255.255.255.0 A(config-if-GigabitEthernet 0/1)#exit A(config)#router rip A(config-router)#version 2 A(config-router)#no auto-summary </pre>																								

	A(config-router)#network 10.10.1.0 0.0.0.255
B	<pre> B(config)#interface GigabitEthernet 0/2 B(config-if-GigabitEthernet 0/2)#no switchport port B(config-if-GigabitEthernet 0/2)#ip address 10.10.2.2 255.255.255.0 B(config-if-GigabitEthernet 0/2)#exit B(config)#router ospf 1 B(config-router)#network 10.10.2.0 0.0.0.255 area 0 </pre>
C	<pre> C(config)# ip vrf Suba C(config-vrf)# rd 100:1 C(config-vrf)# route-target import 100:3 C(config-vrf)# route-target export 100:1 C(config-vrf)# exit C(config)#interface GigabitEthernet 0/1 C(config-GigabitEthernet 0/1)#ip vrf forwarding Suba C(config-GigabitEthernet 0/1)#ip address 10.10.1.1 255.255.255.0 C(config-GigabitEthernet 0/1)# exit C(config)#router rip C(config-router)#address-family ipv4 vrf Suba C(config-router-af)# version 2 C(config-router-af)# no auto-summary C(config-router-af)#network 10.10.1.0 0.0.0.255 C(config-router-af)#exit </pre>
	<pre> C(config)# ip vrf Subb C(config-vrf)# rd 100:2 C(config-vrf)# route-target import 100:3 C(config-vrf)# route-target export 100:2 C(config-vrf)# exit C(config)#interface gigabitEthernet 0/2 C(config-GigabitEthernet 0/2)#ip vrf forwarding Subb C(config-GigabitEthernet 0/2)# ip address 10.10.2.1 255.255.255.0 </pre>

	<pre>C(config-GigabitEthernet 0/2)# exit C(config)# router ospf 2 vrf Subb C(config-router)# network 10.10.2.0 0.0.0.255 area 0 C(config-router)# exit</pre>
	<pre>C(config)# ip vrf Servers C(config-vrf)# rd 100:3 C(config-vrf)# route-target import 100:1 C(config-vrf)# route-target import 100:2 C(config-vrf)# route-target export 100:3 C(config-vrf)# exit C(config)# interface gigabitEthernet 0/3 C(config-GigabitEthernet 0/3)# ip vrf forwarding Servers C(config-GigabitEthernet 0/3)# ip address 10.10.3.1 255.255.255.0 C(config-GigabitEthernet 0/3)# exit</pre>
	<pre>C(config)# router bgp 200 C(config-router)# address-family ipv4 vrf vpna C(config-router-af)# redistribute rip C(config-router-af)# exit C(config-router)# address-family ipv4 vrf vpnb C(config-router-af)# redistribute ospf 1 C(config-router-af)# exit C(config-router)# address-family ipv4 vrf Servers C(config-router-af)# redistribute connected subnets C(config-router-af)# exit</pre>
Verification	<ul style="list-style-type: none"> ▪ Run the show ip vrf interface command on C to check the interface binding information. ▪ Run the show ip route vrf command on C to check whether two VRF tables are created to separate the routes in Sub a from those in Sub b and whether both VRF tables contain the routes in Servers.
C	<pre>C# show ip vrf interfaces Interface IP-Address VRF Protocol</pre>

```
GigabitEthernet 0/1 10.10.1.1 Suba up
GigabitEthernet 0/2 10.10.2.1 Subb up
GigabitEthernet 0/3 10.10.3.1 Servers up
```

```
C# show ip route vrf Subb
Routing Table: Subb

Codes: C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default

Gateway of last resort is no set
O 10.2.0.0/16 [20/0] via 0.0.0.0, 00:10:46, GigabitEthernet 0/2
O 10.10.2.0/24 [20/0] via 0.0.0.0, 00:10:46, GigabitEthernet 0/2
C 10.10.2.1/32 is local host.
C 10.10.3.0/24 is directly connected, GigabitEthernet 0/3
C 10.10.3.1/32 is local host.
```

Common Errors

- An interface is bound to a VRF table after the IP interface of the interface is configured.
- When a physical link is used to forward packets from multiple VPNs, the corresponding physical interface is bound to a VRF table.
- VPN routes are not introduced to BGP.

11.4.2 Configuring a Multiprotocol VRF Table

Configuration Effect

- Provide the VPN service on a device.
- With BGP assistance, flexibly control the separation and access between VPNs.
- With BGP assistance, provide the VPN service in an IP backbone network.
- Support IPv4 and IPv6 through address family configuration.

Notes

- No VRF table needs to be created if the device only forwards packets from one VPN or from the public network.
- If the device needs to forward public network packets and VPN packets or forward packets from multiple VPNs, VRF tables must be created to separate routes.
- In many cases, static or dynamic routes (RIP, OSPF, ISIS, and BGP) need to be added to VRF tables.

Configuration

Steps

Creating a VRF Table

- Mandatory.
- Create a VRF table for each VPN.

Configuring an Address Family

- Mandatory.
- Enable the corresponding address family for each created VRF table.

Configuring an RD

- Optional.
- When routing information needs to be advertised through BGP in the backbone network, BGP may select the best route for advertising if overlapping network addresses exist in different VPNs, which will make some VPNs fail to obtain corresponding routing information. To solve this problem, you can configure RDs for routes to enable BGP to make routing decisions based on these RDs, thus ensuring that each VPN can obtain corresponding routing information.

Configuring an RT

- Optional.
- You can run the **route-target export** command to specify the attributes of the route to be advertised, and run the **route-target import** command to specify the attributes of the route to be received. You can also run the **route-target both** command to specify the export and import attributes.
- Run the **route-target** command in multiprotocol VRF mode or multiprotocol VRF address family mode.

Binding an Interface and Adding the Direct Route and Host Route on the Interface to a VRF Table

- Mandatory.
- If the physical link for transmitting VPN packets is exclusively occupied by a VPN, bind the physical interface to the corresponding VRF table.
- If the physical link for transmitting VPN packets is shared by multiple VPNs, you need to create an independent logical link for each VPN, and bind the logical interface to the corresponding VRF table.

A logical interface can be a subinterface or a VLAN interface.

- Before you bind an interface to a multiprotocol VRF table, enable an address family for the table. If you do not enable the IPv4 address family in advance, you cannot configure the IPv4 address and VRRP IPv4 address of the bound interface. If you do not enable the IPv6 address family in advance, you cannot configure the IPv6 address and VRRP IPv6 address of the bound interface.
- You must bind an interface to the corresponding multiprotocol VRF table before you configure the IPv4 or IPv6 address of the interface. If you bind the interface after its IPv4 or IPv6 address is configured, the address will be invalid.

Adding the Direct Route and Host Route on an Interface to a VRF Table Without Binding the Interface

- Optional.
- If PBR is required for VRF table selection, run the **ip vrf receive** command on the interface to which PBR is applied, and import the direct route and host route on the interface to each VRF table available for choice.

Verification

- Check whether multiprotocol VRF tables are created correctly on the router and corresponding address families are enabled.

Related Commands

Creating a VRF Table

Command	vrf definition <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the name of the VRF table to be created. It cannot exceed 31 characters.
Command Mode	Global configuration mode
Usage Guide	After you run the command, the system will enter VRF mode.

Enabling the IPv4 Address Family

Command	address-family ipv4
Parameter Description	N/A

Command Mode	VRF mode
Usage Guide	After you run the command, the system will enter VRF IPv4 address family submode.

Enabling the IPv6 Address Family

Command	address-family ipv6
Parameter Description	N/A
Command Mode	VRF mode
Usage Guide	After you run the command, the system will enter VRF IPv6 address family submode.

Configuring an RD

Command	rd rd_value
Parameter Description	<p><i>rd_value</i> has the following three different parameter forms:</p> <p>(1) <i>rd_value</i> = as_num: nn as_num indicates the 2-byte number that identifies a public AS. nn is configurable in the range 0..4294967295.</p> <p>(2) <i>rd_value</i> = ip_addr: nn ip_addr must be a global IP address. nn is configurable in the range 0..65535.</p> <p>(3) <i>rd_value</i> = as4_num: nn as4_num indicates the 4-byte number that identifies a public AS. nn is configurable in the range 1..65535.</p>
Command Mode	VRF configuration mode
Usage Guide	<p>You cannot directly change the RD of an existing VRF table. You need to delete the VRF table first and then configure a new RD.</p> <p>A VRF table has only one RD. You cannot configure multiple RDs for one VRF table.</p>

Configuring an RT

Command	route-target { import export both } <i>rt_value</i>
Parameter Description	<p><i>rt_value</i> has the following three different parameter forms:</p> <p>(1) <i>rt_value</i> = as_num: nn</p> <p>as_num indicates the 2-byte number that identifies a public AS. nn is configurable in the range 0..4294967295.</p> <p>(2) <i>rt_value</i> = ip_addr: nn</p> <p>ip_addr must be a global IP address. nn is configurable in the range 0..65535.</p> <p>(3) <i>rt_value</i> = as4_num: nn</p> <p>as4_num indicates the 4-byte number that identifies a public AS. nn is configurable in the range 1..65535.</p>
Command Mode	VRF configuration mode or VRF address family submode
Usage Guide	One VRF table can be configured with multiple import and export RT attributes.

Binding an Interface

Command	vrf forwarding <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the name of a VRF table.
Command Mode	Interface configuration mode
Usage Guide	<p>By default, an interface does not belong to any VRF table.</p> <p>After an interface is bound to the corresponding VRF table, the direct route and host route on the interface will be automatically added to the VRF table.</p> <p>Before you bind an interface to a multiprotocol VRF table, enable an address family for the table. If you do not enable the IPv4 address family in advance, you cannot configure the IPv4 address and VRRP IPv4 address of the bound interface. If you do not enable the</p>

	<p>IPv6 address family in advance, you cannot configure the IPv6 address and VRRP IPv6 address of the bound interface.</p> <p>You must bind an interface to a multiprotocol VRF table before you configure the IPv4, IPv6, VRRP IPv4, and VRRP IPv6 addresses of the interface; otherwise, these addresses will be invalid and the IPv6 protocol on the interface will be disabled.</p> <p>If the IPv4 address family is deleted from the multiprotocol VRF table, the IPv4 and VRRP IPv4 addresses of all interfaces bound to the VRF table will be deleted, and the IPv4 static routes in the VRF table or next-hop routes are also deleted. If the IPv6 address family is deleted from the multiprotocol VRF table, the IPv6 and VRRP IPv6 addresses of all interfaces bound to the VRF table will be deleted, the IPv6 protocol on the interfaces will be disabled, and the IPv6 static routes in the VRF table or next-hop routes are also deleted.</p>
--	--

Adding the Direct Route and Host Route on an Interface to a VRF Table Without Binding the Interface

Command	vrf receive <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the name of a VRF table.
Command Mode	Interface configuration mode
Usage Guide	<p>This command is used to add the host route and direct route on an interface to a VRF table. If you need to add the host route and direct route on an interface to multiple VRF tables, run the command multiple times.</p> <p>Different from the vrf forwarding command, the vrf receive command does not bind an interface to the corresponding VRF table. The interface is still a global interface and does not belong to any VRF table.</p> <p>The vrf forwarding and vrf receive commands are mutually exclusive on the same interface.</p>

Displaying the VRF Information on a Device

Command	show vrf [<i>brief</i> <i>detail</i> <i>ipv4</i> <i>ipv6</i>] [<i>vrf-name</i>]
Parameter Description	<p>brief: Displays brief information.</p> <p>detail: Displays detailed information.</p> <p>ipv4: Displays the brief information of an IPv4 VRF table.</p> <p>ipv6: Displays the brief information of an IPv6 VRF table.</p>

	<i>vrf-name</i> : Indicates the name of a VRF table.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	This command is used to display the information of a specified VRF table to check whether the VRF table is bound with the correct interface.

Displaying the Routes in a VRF Table

Command	show ip route vrf <i>vrf-name</i>
Parameter Description	<i>vrf-name</i> : Indicates the name of a VRF table.
Command Mode	Privilege, global and interface configuration modes
Usage Guide	This command is used to check whether a specified VRF table contains corresponding routes.

Configuration

Example

- i** The following example only describes VRF-related configuration on A1, B1, MCE1, and PE1. The configuration on A2, B2, MCE2, and PE2 is similar.

VRF on CEs and PEs (MCE Application)

Scenario
Figure 11-5



VPN a and VPN b have independent address spaces.
VPN a runs RIP and VPN b runs OSPF.

Configuration Requirements

The routes in VPN a are separated from those in VPN b. A1 and A2 can access each other, and B1 and B2 can access each other.

Configuration Steps

- Connect MCE1 and A1 through RIP. Extend RIP routes on A1. On MCE1, create a VRF table for VPN a, bind the directly connected interface, and configure RIP routes.
- Connect MCE1 and B1 through OSPF. Extend OSPF routes on B1. On MCE1, create a VRF table for VPN b, bind the directly connected interface, and configure OSPF routes.
- Connect MCE1 and PE1 through BGP. On MCE1 and PE1, create a VRF table for each VPN, bind the VLAN interface, and configure BGP routes.
- Configure the physical link between MCE1 and PE1 in Trunk mode.
- In the VRF instance for VPN a on MCE1, redistribute the RIP routes to BGP, and redistribute the BGP routes to RIP.
- In the VRF instance for VPN b on MCE1, redistribute the OSPF routes to BGP, and redistribute the BGP routes to OSPF.

i Planning of interfaces and addresses:

Interface Description	Interface Name	IP Address/Mask	VRF Table
Physical interface on A1 connected to MCE1	GE0/1	10.10.1.2/24	-
Physical interface on B1 connected to MCE1	GE0/2	10.10.2.2/24	-
Physical interface on MCE1 connected to A1	GE0/1	10.10.1.1/24	VRF table for VPN a
Physical interface on MCE1 connected to B1	GE0/2	10.10.2.1/24	VRF table for VPN b
Logical interface on MCE1 connected to PE1	VLAN10	10.10.10.1/24	VRF table for VPN a
Logical interface on MCE1 connected to PE1	VLAN20	10.10.20.1/24	VRF table for VPN b

	Logical interface on PE1 connected to MCE1	VLAN10	10.10.10.2/24	VRF table for VPN a
	Logical interface on PE1 connected to MCE1	VLAN20	10.10.20.2/24	VRF table for VPN b
A1	<pre>A1(config)#interface GigabitEthernet 0/1 A1(config-if-GigabitEthernet 0/1)#no switchport port A1(config-if-GigabitEthernet 0/1)#ip address 10.10.1.2 255.255.255.0 A1(config-if-GigabitEthernet 0/1)#exit A1(config)#router rip A1(config-router)#version 2 A1(config-router)#no auto-summary A1(config-router)#network 10.10.1.0 0.0.0.255</pre>			
B1	<pre>B1(config)#interface GigabitEthernet 0/2 B1(config-if-GigabitEthernet 0/1)#no switchport port B1(config-if-GigabitEthernet 0/1)#ip address 10.10.2.2 255.255.255.0 B1(config-if-GigabitEthernet 0/1)#exit B1(config)#router ospf 1 B1(config-router)#network 10.10.2.0 0.0.0.255 area 0</pre>			
MCE1	<pre>#Create a VRF table for VPN a and a VRF table VPN b, and enable the IPv4 address family. MCE1(config)#vrf definition vpna MCE1(config-vrf)#address-family ipv4 MCE1(config-vrf-af)#exit MCE1(config-vrf)#exit MCE1(config)#vrf definition vpnb MCE1(config-vrf)#address-family ipv4 MCE1(config-vrf-af)#exit MCE1(config-vrf)#exit</pre>			
	<pre>#Bind interfaces to the VRF tables. MCE1(config)#interface GigabitEthernet 0/1</pre>			

```
MCE1(config-if-GigabitEthernet 0/1)#no switchport port
MCE1(config-if-GigabitEthernet 0/1)#vrf forwarding vpna
MCE1(config-if-GigabitEthernet 0/1)#ip address 10.10.1.1 255.255.255.0
MCE1(config-if-GigabitEthernet 0/1)#exit
MCE1(config)#interface GigabitEthernet 0/2
MCE1(config-if-GigabitEthernet 0/2)#no switchport port
MCE1(config-if-GigabitEthernet 0/2)#vrf forwarding vpb
MCE1(config-if-GigabitEthernet 0/2)#ip address 10.10.2.1 255.255.255.0
MCE1(config-if-GigabitEthernet 0/2)#exit
MCE1(config)#interface vlan 10
MCE1(config-if-VLAN 10)#vrf forwarding vpna
MCE1(config-if-VLAN 10)#ip address 10.10.10.1 255.255.255.0
MCE1(config-if-VLAN 10)#exit
MCE1(config)#interface vlan 20
MCE1(config-if-VLAN 20)#vrf forwarding vpb
MCE1(config-if-VLAN 20)#ip address 10.10.20.1 255.255.255.0
MCE1(config-if-VLAN 20)#exit
```

#Configure the interface connected to PE1 in Trunk mode.

```
MCE1(config)#interface GigabitEthernet 0/3
MCE1(config-if-GigabitEthernet 0/3)#switchport mode trunk
MCE1(config-if-GigabitEthernet 0/3)#exit
```

#Configure RIP and BGP routes in the VRF table for VPN a, and introduce routes in the two VRF tables to each other.

```
MCE1(config)#router rip
MCE1(config-router)#address-family ipv4 vrf vpna
MCE1(config-router-af)# version 2
MCE1(config-router-af)# no auto-summary
MCE1(config-router-af)#network 10.10.1.0 0.0.0.255
MCE1(config-router-af)#redistribute bgp subnets
MCE1(config-router-af)#exit
MCE1(config)# router bgp 100
MCE1(config-router)#address-family ipv4 vrf vpna
```

	<pre>MCE1(config-router-af)#neighbor 10.10.10.2 remote-as 200 MCE1(config-router-af)#redistribute rip MCE1(config-router-af)#exit</pre>
	<p>#Configure OSPF and BGP routes in the VRF table for VPN b, and introduce routes in the two VRF tables to each other.</p> <pre>MCE1(config)#router ospf 1 vrf vpnb MCE1(config-router)#network 10.10.2.0 0.0.0.255 area 0 MCE1(config-router)#redistribute bgp subnets MCE1(config-router)#exit MCE1(config)# router bgp 100 MCE1(config-router)#address-family ipv4 vrf vpnb MCE1(config-router-af)#neighbor 10.10.20.2 remote-as 200 MCE1(config-router-af)#redistribute ospf 1 MCE1(config-router-af)#exit</pre>
PE1	<p>#Create a VRF table for VPN a and a VRF table VPN b, and enable the IPv4 address family.</p> <pre>PE1(config)#vrf definition vpna PE1(config-vrf)#address-family ipv4 PE1(config-vrf-af)#exit PE1(config-vrf)#exit PE1(config)#vrf definition vpnb PE1(config-vrf)#address-family ipv4 PE1(config-vrf-af)#exit PE1(config-vrf)#exit</pre>
	<p>#Bind interfaces to the VRF tables.</p> <pre>PE1(config)#vlan 10 PE1(config-vlan)#exit PE1(config)#vlan 20 PE1(config-vlan)#exit PE1(config)#interface vlan 10 PE1(config-if-VLAN 10)#vrf forwarding vpna</pre>

	<pre>PE1(config-if-VLAN 10)#ip address 10.10.10.2 255.255.255.0 PE1(config-if-VLAN 10)#exit PE1(config)#interface vlan 20 PE1(config-if-VLAN 20)#vrf forwarding vpnb PE1(config-if-VLAN 20)#ip address 10.10.20.2 255.255.255.0 PE1(config-if-VLAN 20)#exit</pre>
	<pre>#Configure the interface on PE1 connected to MCE1 in Trunk mode. PE1(config)#interface GigabitEthernet 0/3 PE1(config-if-GigabitEthernet 0/3)#switchport mode trunk PE1(config-if-GigabitEthernet 0/3)#exit</pre>
	<pre>#Configure BGP routes in the VRF table for VPN a. PE1(config)# router bgp 200 PE1(config-router)#address-family ipv4 vrf vpna PE1(config-router-af)#neighbor 10.10.10.1 remote-as 100 PE1(config-router-af)#exit</pre>
	<pre>#Configure BGP routes in the VRF table for VPN b. PE1(config)# router bgp 200 PE1(config-router)#address-family ipv4 vrf vpnb PE1(config-router-af)#neighbor 10.10.20.1 remote-as 100 PE1(config-router-af)#exit</pre>
Verification	<ul style="list-style-type: none"> ▪ On A1, run the show ip route command to display the routes in VPN a. ▪ On B2, run the show ip route command to display the routes in VPN b. ▪ On MCE1, run the show ip route vrf vpna command to display the routes in VPN a, and run the show ip route vrf vpnb command to display the routes in VPN b. ▪ On PE1, run the show ip route vrf vpna command to display the routes in VPN a, and run the show ip route vrf vpnb command to display the routes in VPN b.
	-

Common Errors

- A multiprotocol VRF table is configured, but no address family is enabled.
- An interface is bound to a VRF table after the IP interface of the interface is configured.

- When a physical link is used to forward packets from multiple VPNs, the corresponding physical interface is bound to a VRF table.
- VPN routes are not introduced to BGP.

11.5 Monitoring

Clearing

 Running the **clear** commands may lose vital information and thus interrupt services.

Description	Command
Clears the routes in a specified VRF table.	clear ip route vrf <i>vrf-name</i>

Displaying

Description	Command
Displays the information of a single-protocol VRF table.	show ip vrf [brief detail interfaces] [<i>vrf-name</i>]
Displays the information of a multiprotocol VRF table.	show vrf [ipv4 ipv6 brief detail] [<i>vrf-name</i>]

Debugging

 System resources are occupied when debugging information is output. Therefore, disable debugging immediately after use.

Description	Command
Displays the debugging information during the processes where a VRF table is created, an address family is enabled, and an interface is bound to the VRF table.	debug vrf

Prints the information of interface-related VRF operation debugging.

debug vrf interface

12 CONFIGURING L3MAC

12.1 Overview

A layer-3 MAC address for forwarding can be configured for a device so that packets whose destination MAC addresses match the configured layer-3 MAC address are forwarded at layer 3. In data center scenarios, this MAC address can be configured in static ARP entries of neighbors so that routing can be triggered on any neighbor device.

12.2 Applications

Application	Description
Data Center Scenario	Layer-3 routing can be triggered on any neighbor without perceiving the layer-3 MAC address of the device.


12.2.1 Data Center Scenario

Scenario

A data center adopting the topology shown in the figure below is used as an example:

- A global layer-3 MAC address is configured on SW2 and SW3.
- A static ARP entry is configured on SW1.
- The MAC address in the static ARP entry configured on SW1 is the global layer-3 MAC address of its neighbors, that is, SW2 and SW3.

Figure 1-1



Deployment

- Packets from SW1 can be routed at layer 3 through its neighbors: SW2 and SW3.

12.3 Configuration

Configuration	Description and Command	
Adding a Layer-3 MAC Address	⚠ Mandatory.	
	I3intf-mac-add	Adds a layer-3 MAC address.

12.3.1 Adding a Layer-3 MAC Address

Configuration

Effect

- Add a global layer-3 MAC address so that layer-3 routing can be triggered when packets matching this layer-3 MAC address are received.

Notes

- Up to two layer-3 MAC addresses can be added.

Configuration

Steps

Adding a Layer-3 MAC Address

Command	I3intf-mac-add <i>mac</i>
Parameter Description	<i>mac</i> : Indicates the layer-3 MAC address.
Command Mode	Global configuration mode
Usage Guide	N/A

Deleting a Layer-3 MAC Address

Command	no I3intf-mac-add <i>mac</i>
----------------	-------------------------------------

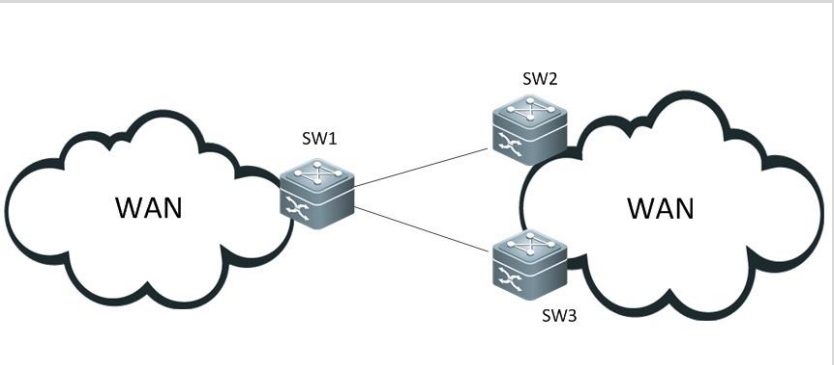
Parameter Description	<i>mac</i> : Indicates the layer-3 MAC address.
Command Mode	Global configuration mode
Usage Guide	N/A

Verification

1. Run the **show run** command to check whether the current configuration is correct.
2. Run the **show l3intf-mac-add** command to check whether the current configuration is correct.

Configuration Example

Adding a Layer-3 MAC Address

Scenario Figure 1-2	
Configuration Steps	Configure a static ARP entry on SW1, and a global layer-3 MAC address on SW2 and SW3. Other configurations are omitted.
SW2	SW2# configure terminal SW2#(config)#l3intf-mac-add 0000.0001.2233
SW3	SW3# configure terminal SW3#(config)#l3intf-mac-add 0000.0001.2233

SW1	SW1# configure terminal SW1(config)# arp 192.168.22.20 0000.0001.2233 arpa
Verification	Run the show command to display the global MAC address configured on SW2 and SW3. Pump traffic and verify that packets from SW1 are routed at layer 3 regardless of whether the next hop of SW1 is SW2 or SW3.
SW2	SW2(config)#show l3intf-mac-add l3intf mac address: 0000.0001.2235
SW3	SW3(config)#show l3intf-mac-add l3intf mac address: 0000.0001.2235

12.4 Monitoring

Displaying

Description	Command
Displays the configured global layer-3 MAC address of the device.	show l3intf-mac-add

13 CONFIGURING HARDWARE CAPACITY QUERY

13.1 Overview

Hardware capacity query provides the function of querying the capacity and usage condition of the Forwarding Information Base (FIB) hardware.

Protocol s and Standar ds

N/A

13.2 Features

Basic Concept s

Route hardware resources are classified into IPv4 route resources and IPv6 route resources by protocol, and host route resources and network route resources by route prefix.

The allocation of route resources varies with the chip hardware type.

- QSW-6900 series

Host route resources and network route resources are managed separately and their capacities are fixed. IPv4 and IPv6 host routes share the host route resources while IPv4 and IPv6 network routes share network route resources. One IPv6 route occupies twice resources of one IPv4 route. For example, if the total host route capacity is 200K, then 200K IPv4 host routes, 100K IPv6 host routes, or 100K IPv4 host routes and 50K IPv6 host routes are supported.

IPv6 network routes include network routes with the prefix length of 0-64 and network routes with the prefix length of 65-127. The two types of network route resources are managed separately and their capacities are fixed.

- i** When a hash conflict occurs during the configuration of host routes, network route resources are occupied and the actually used network route capacity may be greater than the displayed used capacity.

13.3 Monitoring

Displayi ng

Description	Command
-------------	---------

Displays the capacity and usage condition of the FIB hardware of the switch.

show route-res usage

Debugging

N/A

-  The monitoring command above is also applicable to chassis-type and box-type devices, standalone devices, and VSUs.