Advanced IGP Features

Foreword

- Both Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) are link-state Interior Gateway Protocols (IGPs). Routers that run them synchronize link state databases (LSDBs) and use the shortest path first (SPF) algorithm to calculate optimal routes.
- In response to network topology changes, OSPF and IS-IS support multiple fast convergence and protection mechanisms, which minimize traffic loss caused by network faults.
- To control the size of a routing table, OSPF and IS-IS support route selection and routing information control.
- This course describes the advanced features of OSPF and IS-IS, including fast convergence and route control.

Objectives

- On completion of this course, you will be able to:
  - Describe various fast convergence techniques of OSPF and IS-IS.
  - Configure OSPF and IS-IS equal-cost routes.
  - Configure OSPF and IS-IS to advertise default routes.
  - Describe the application scenarios of OSPF and IS-IS multi-process.
  - Describe the application scenarios of OSPF forwarding address (FA).
  - Describe the implementation of IS-IS LSP fragment extension.

Contents

Overview of OSPF Fast Convergence

- OSPF fast convergence is an extended feature of OSPF to speed up route convergence. It features partial route calculation (PRC) and the intelligent timer.
- In addition, OSPF supports fast convergence upon fault rectification. For example, OSPF IP fast reroute (FRR) can be used to implement fast traffic switchover to a backup link, and OSPF can also be associated with BFD to implement fast fault detection.

PRC

- PRC only calculates routes that have been changed on a network.
- PRC does not calculate nodes. Instead, it updates routes based on the shortest path tree (SPT) calculated using the SPF algorithm.

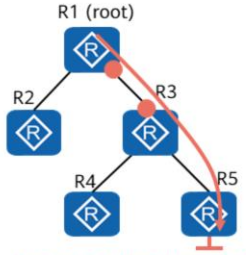

Loopback0: A directly connected network segment is added.

In route calculation, a node represents a router, and a leaf represents a route. PRC processes only the changed leaf information.

- Scenario description:
  - OSPF runs on a network. The figure shows the SPT with R1 as the root after network convergence. When R1 accesses R5, traffic is sent to the destination based on [outbound interface of R1's downlink, IP address of R3's uplink interface].
  - OSPF is enabled on Loopback0 of R5. This means that a new network segment is added to the OSPF network.
- PRC:
  - R5 floods a new LSA on the entire network.
  - After receiving the LSA, R1 creates a new route that inherits the original path and next hop used when R1 accesses R5. In this case, the SPT remains unchanged, and only a leaf is added to R5.
  - Therefore, when R1 accesses Loopback0 of R5, traffic is sent to the destination based on [outbound interface of R1's downlink, IP address of R3's uplink interface].
- Benefits:
  - PRC focuses only on routes that are changed due to the addition of network segments to an OSPF network, thereby speeding up route calculation.

- Note: On Huawei devices, OSPF PRC is enabled by default.

Intelligent Timer

- The intelligent timer is used for SPF calculation and LSA generation.
- It can quickly respond to a small number of external incidents and prevent excessive CPU consumption.

Controlling LSA Generation and Reception

- To prevent network connections or frequent route flapping from consuming excessive device resources, OSPF complies with the following rules:
  - After an LSA is generated, it cannot be generated again within 1s. The interval for updating LSAs is 5s.
  - The interval for receiving LSAs is 1s.
- On a stable network where routes need to be fast converged, the intelligent timer can be used to set the interval for updating LSAs to 0s in order to cancel this interval. In this manner, topology or route changes can be immediately advertised to the network through LSAs or be immediately detected, thereby speeding up route convergence on the network.

Controlling Route Calculation

- When a network changes, the OSPF LSDB changes, and the shortest path needs to be recalculated. If a network changes frequently, the shortest path is calculated accordingly, which results in excessive consumption of system resources and compromises device efficiency.
- You can configure the intelligent timer to set a proper interval for SPF calculation in order to prevent excessive consumption of a router's memory and bandwidth resources.

- If the interval for triggering route calculation is long, the network convergence speed is affected.
- The first timeout period of the intelligent timer is fixed. Before the intelligent timer expires, if an event that triggers the timer occurs, the next timeout period of the intelligent timer becomes longer.

Basic Intelligent Timer Configuration Commands (1)

1. Set an interval for updating OSPF LSAs.

   [Huawei-ospf-1] **lsa-originate-interval** { **0** | { **intelligent-timer** *max-interval start-interval hold-interval* | **other-type** *interval* } * }

   By default, the intelligent timer is enabled; the maximum interval, initial interval, and hold interval at which LSAs are updated are 5000 ms, 500 ms, and 1000 ms, respectively.

After the intelligent timer is used:

1. The initial interval for updating LSAs is specified by *start-interval*.
2. The interval at which LSAs are updated for the *n*th ($n \geq 2$) time equals *hold-interval* x $2^{(n-2)}$.
3. When the interval specified by *hold-interval* x $2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPF updates LSAs at the maximum interval for three consecutive times. Then, OSPF returns to the first step and updates LSAs at the initial interval specified by *start-interval*.

- Command: [Huawei-ospf] **lsa-originate-interval** { **0** | { **intelligent-timer** *max-interval start-interval hold-interval* | **other-type** *interval* } }
  - **0**: sets the interval for updating LSAs to 0s, that is, cancels the interval of 5s for updating LSAs.
  - **intelligent-timer**: uses the intelligent timer to set the update interval for router-LSAs and network-LSAs.
  - *max-interval*: specifies the maximum interval for updating OSPF LSAs. The value is an integer ranging from 1 to 120000, in milliseconds. The default value is 5000.
  - *start-interval*: specifies the initial interval for updating OSPF LSAs. The value is an integer ranging from 0 to 60000, in milliseconds. The default value is 500.
  - *hold-interval*: specifies the hold interval for updating OSPF LSAs. The value is an integer ranging from 1 to 60000, in milliseconds. The default value is 1000.
  - **other-type**: sets an update interval for OSPF LSAs except router-LSAs and network-LSAs.
  - *interval*: specifies the interval for updating LSAs. The value is an integer ranging from 0 to 10, in seconds. The default value is 5.

Basic Intelligent Timer Configuration Commands (2)

2. Set an interval for receiving OSPF LSAs.

[Huawei-ospf-1] **lsa-arrival-interval** { *interval* | **intelligent-timer** *max-interval start-interval hold-interval* }

By default, the intelligent timer is enabled; the maximum interval, initial interval, and hold interval at which LSAs are received are 1000 ms, 500 ms, and 500 ms, respectively.

After the intelligent timer is used:

1. The initial interval for receiving LSAs is specified by *start-interval*.
2. The interval at which LSAs are received for the *n*th ($n \geq 2$) time equals *hold-interval* x $2^{(n-2)}$.
3. When the interval specified by *hold-interval* x $2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPF receives LSAs at the maximum interval for three consecutive times. Then, OSPF returns to the first step and receives LSAs at the initial interval specified by *start-interval*.

- Command: [Huawei-ospf-1] **lsa-arrival-interval** { *interval* | **intelligent-timer** *max-interval start-interval hold-interval* }
  - *interval*: specifies the interval for receiving LSAs. The value is an integer ranging from 0 to 10000, in milliseconds.
  - **intelligent-timer**: uses the intelligent timer to set the receive interval for LSAs.
  - *max-interval*: specifies the maximum interval for receiving OSPF LSAs. The value is an integer ranging from 1 to 120000, in milliseconds. The default value is 1000.
  - *start-interval*: specifies the initial interval for receiving OSPF LSAs. The value is an integer ranging from 0 to 60000, in milliseconds. The default value is 500.
  - *hold-interval*: specifies the hold interval for receiving OSPF LSAs. The value is an integer ranging from 1 to 60000, in milliseconds. The default value is 500.

Basic Intelligent Timer Configuration Commands (3)

3. Set an interval for OSPF route calculation.

   [Huawei-ospf-1] **spf-schedule-interval** { *interval1* | **intelligent-timer** *max-interval start-interval hold-interval* | **millisecond** *interval2* }

   By default, the intelligent timer is enabled; the maximum interval, initial interval, and hold interval for SPF calculation are 10000 ms, 500 ms, and 1000 ms, respectively.

After the intelligent timer is used, the interval for SPF calculation is as follows:

1. The initial interval for SPF calculation is specified by *start-interval*.
2. The interval for SPF calculation for the *n*th ($n \geq 2$) time equals *hold-interval* x $2^{(n-2)}$.
3. When the interval specified by *hold-interval* x $2^{(n-2)}$ reaches the maximum interval specified by *max-interval*, OSPF performs SPF calculation at the maximum interval for three consecutive times. Then, OSPF returns to the first step and performs SPF calculation at the initial interval specified by *start-interval*.

- Command: [Huawei-ospf-1] **spf-schedule-interval** { *interval1* | **intelligent-timer** *max-interval start-interval hold-interval* | **millisecond** *interval2* }
  - *interval1*: specifies an interval for OSPF SPF calculation. The value is an integer ranging from 1 to 10, in seconds.
  - **intelligent-timer**: uses the intelligent timer to set the interval for OSPF SPF calculation.
  - *max-interval*: specifies the maximum interval for OSPF SPF calculation. The value is an integer ranging from 1 to 120000, in milliseconds. The default value is 10000.
  - *start-interval*: specifies the initial interval for OSPF SPF calculation. The value is an integer ranging from 1 to 60000, in milliseconds. The default value is 500.
  - *hold-interval*: specifies the hold interval for OSPF SPF calculation. The value is an integer ranging from 1 to 60000, in milliseconds. The default value is 1000.
  - **millisecond** *interval2*: specifies an interval for OSPF SPF calculation. The value is an integer ranging from 1 to 10000, in milliseconds.

OSPF IP FRR

- OSPF IP fast reroute (FRR) is a dynamic IP FRR technology that uses the loop-free alternate (LFA) algorithm to pre-calculate a backup path and saves it in the forwarding table. If the primary link fails, traffic is rapidly switched to the backup link, ensuring traffic continuity and achieving traffic protection. OSPF IP FRR can reduce the fault recovery time to less than 50 ms.
- The LFA algorithm calculates a backup link based on the following principles:
  - A device uses the SPF algorithm to calculate shortest paths to the destination, with each neighbor that provides a backup link as a root node. The device then uses the inequality to calculate a loop-free backup link with the minimum cost.

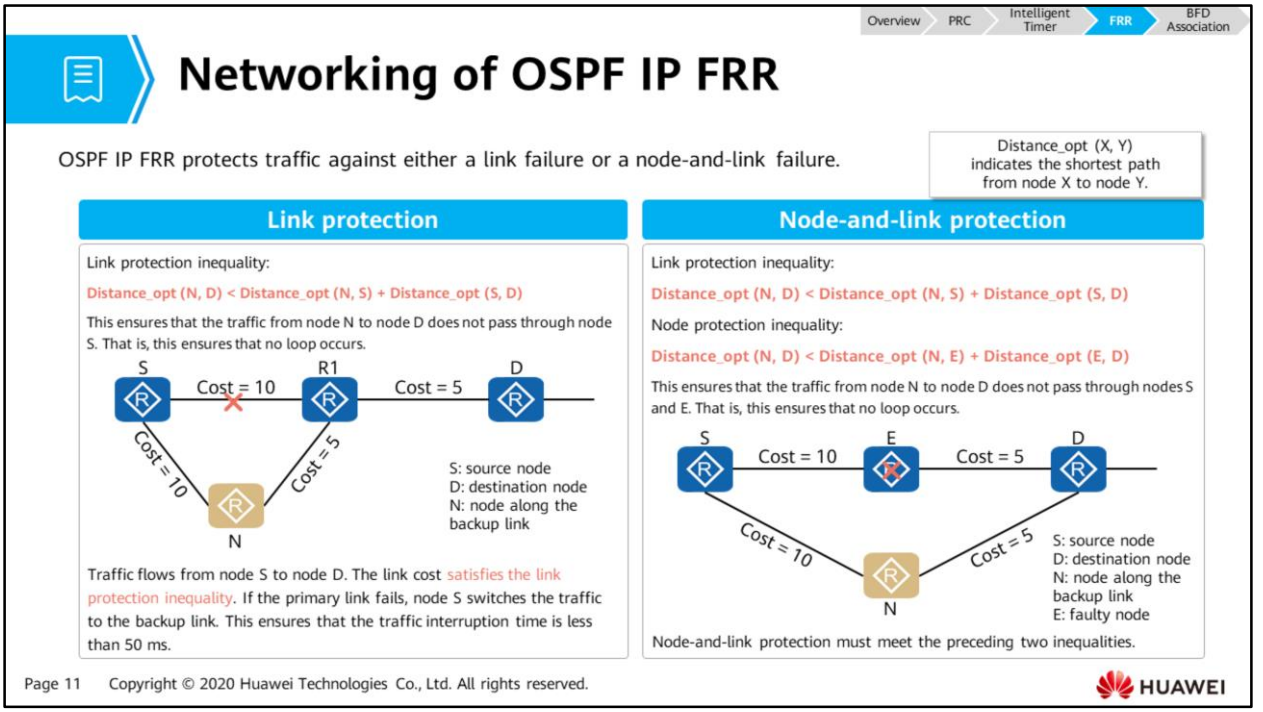

- Node-and-link protection:
  - As shown in the right figure, traffic flows from node S to node D. The link cost satisfies the node-and-link protection inequality. If the primary link fails, node S switches the traffic to the backup link. This ensures that the traffic interruption time is less than 50 ms.
- OSPF IP FRR protects traffic against either a link failure or a node-and-link failure.
  - Link protection takes effect when the traffic to be protected flows along a specified link.
  - Node-and-link protection takes effect when the traffic to be protected flows along a specified device. Node-and-link protection takes precedence over link protection.

Basic OSPF IP FRR Configuration Commands

1. Enable OSPF IP FRR.

```
[Huawei-ospf-1] frr
[Huawei-ospf-1-frr]
```

Create OSPF FRR and enter its view.

```
[Huawei-ospf-1-frr] loop-free-alternate
```

After OSPF IP FRR is enabled, the device uses the LFA algorithm to calculate the next hop and outbound interface for a backup link.

2. (Optional) Disable OSPF IP FRR on an interface.

```
[Huawei-GigabitEthernet0/0/1] ospf frr block
```

OSPF IP FRR can be disabled on an interface of a specific device that is running important services and resides on an FRR backup link. This setting prevents the device connected to this interface from being a part of a backup link and being burdened after FRR switches traffic to the backup link.

Example for Configuring OSPF IP FRR

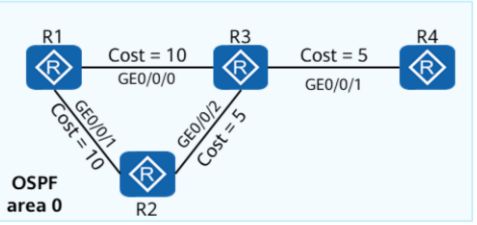

Device	Router ID	Interface	IP Address
R1	10.1.1.1	GE0/0/0	10.1.13.1/24
		GE0/0/1	10.1.12.1/24
R2	10.1.2.2	GE0/0/1	10.1.12.2/24
		GE0/0/2	10.1.23.2/24
R3	10.1.3.3	GE0/0/0	10.1.13.3/24
		GE0/0/1	10.1.34.3/24
		GE0/0/2	10.1.23.3/24
R4	10.1.4.4	GE0/0/1	10.1.34.4/24

If the link between R1 and R3 fails, traffic forwarded by R1 can be quickly switched to the backup link and forwarded by R2.

1. Assign an IP address to each interface and configure OSPF on each device. (The configuration details are not provided here.)
2. Configure an OSPF cost for each device. The following example uses the command output on R1.

```
[R1] interface GigabitEthernet 0/0/0
[R1-GigabitEthernet 0/0/0] ospf cost 10
[R1-GigabitEthernet 0/0/0] quit
[R1] interface GigabitEthernet 0/0/1
[R1-GigabitEthernet 0/0/1] ospf cost 10
[R1-GigabitEthernet 0/0/1] quit
```

3. Enable OSPF IP FRR on R1.

```
[R1] ospf
[R1-ospf-1] frr
[R1-ospf-1-frr] loop-free-alternate
[R1-ospf-1-frr] quit
[R1-ospf-1] quit
```

The cost configurations of R2, R3, and R4 are similar to the configuration of R1.

Checking the OSPF IP FRR Configuration

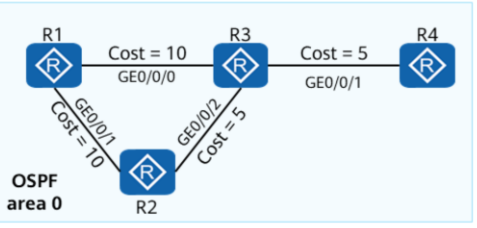

Device	Router ID	Interface	IP Address
R1	10.1.1.1	GE0/0/0	10.1.13.1/24
		GE0/0/1	10.1.12.1/24
R2	10.1.2.2	GE0/0/1	10.1.12.2/24
		GE0/0/2	10.1.23.2/24
R3	10.1.3.3	GE0/0/0	10.1.13.3/24
		GE0/0/1	10.1.34.3/24
		GE0/0/2	10.1.23.3/24
R4	10.1.4.4	GE0/0/1	10.1.34.4/24

If the link between R1 and R3 fails, traffic forwarded by R1 can be quickly switched to the backup link and forwarded by R2.

Check information about the route from R1 to GE0/0/1 of R4.

```
[R1]display ospf routing 10.1.34.4
          OSPF Process 1 with Router ID 10.1.1.1
 Destination :  10.1.34.0/24
 AdverRouter :  10.1.4.4      Area :      0.0.0.0
 Cost :         15            Type :      Transit
 NextHop :      10.1.13.3     Interface : GigabitEthernet0/0/0
 Priority :     Low           Age :       00h01m59s
 Backup Nexthop :  10.1.12.2  Backup Interface:GigabitEthernet0/0/1
 Backup Type :     LFA LINK
```

You can find that OSPF generates a backup link after OSPF IP FRR is enabled on R1.

BFD for OSPF

- A link fault or a topology change causes devices to recalculate routes. Fast and efficient routing protocol convergence is necessary to improve network availability.
- BFD for OSPF associates BFD with OSPF. If a fault occurs on the link between a device and its neighbor, BFD can rapidly detect the link fault to speed up OSPF's response to network topology changes.

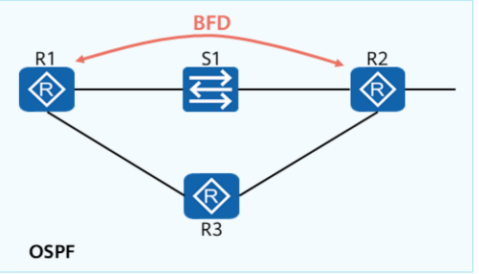

The working principle of BFD for OSPF is as follows:

- OSPF neighbor relationships are established between R1, R2, and R3. When the neighbor relationships enter the Full state, BFD is instructed to set up a BFD session.
- If a fault occurs on the link between R1 and R2, BFD detects the fault and notifies R1. R1 processes the BFD session down event and recalculates the route. The new path is R1-R3-R2.

  HUAWEI

- OSPF periodically sends Hello packets to neighbors to detect faults. It takes more than 1s to detect a fault. By default, when the OSPF Dead timer expires, the neighbor is considered invalid. The default value of the OSPF Dead timer is 40s. With the development of technologies, voice, video, and video on demand (VOD) services are widely used. These services are sensitive to the packet loss rate and delay. When the traffic rate reaches gigabit per second (Gbit/s), long-time fault detection causes a large number of packets to be lost. This cannot meet high reliability requirements of the carrier-class network.
- BFD for OSPF is introduced to resolve this problem. After BFD for OSPF is configured in a specified process or on a specified interface, the link status can be rapidly detected and fault detection can be completed in milliseconds. This speeds up OSPF convergence when the link status changes.

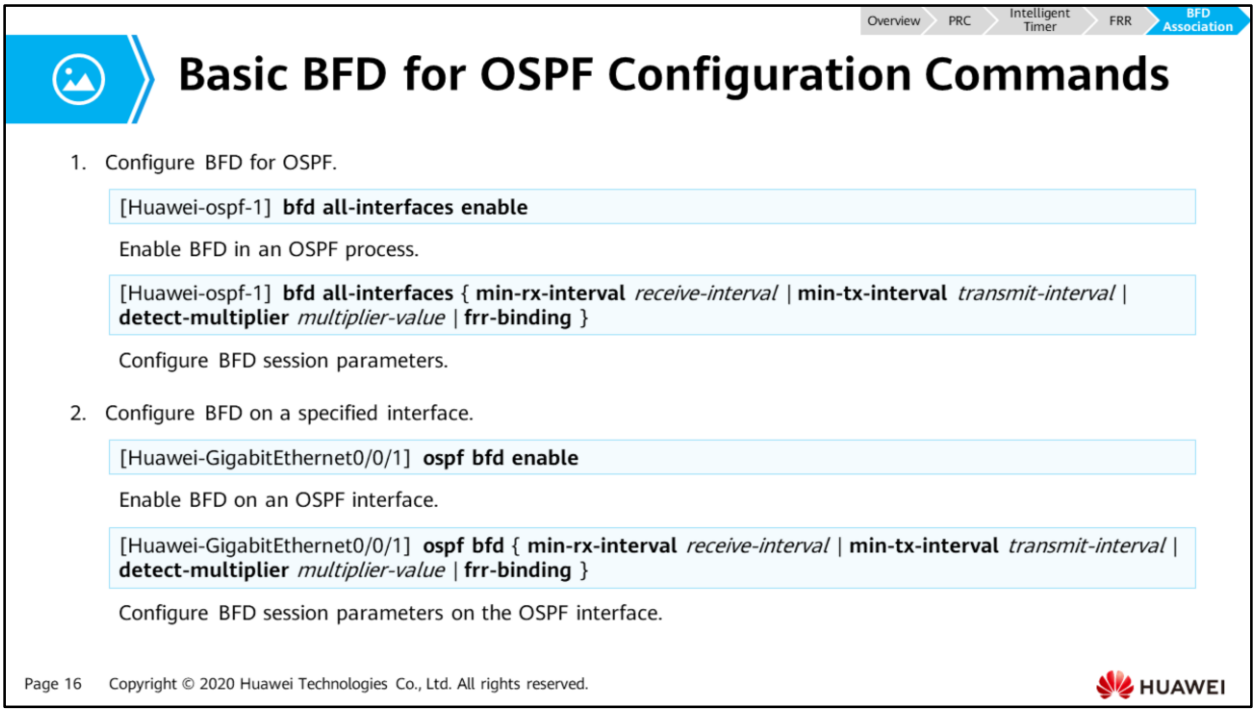

- Prerequisites:
  - Before using BFD to quickly detect link faults, run the **bfd** command in the system view to enable BFD globally.
- The BFD configuration on an interface takes precedence over that in a process. If BFD is enabled on an interface, the BFD parameters on the interface are used to establish BFD sessions.
- OSPF IP FRR can be associated with BFD.
  - During the OSPF IP FRR configuration, the underlying layer needs to fast respond to a link status change so that traffic can be switched to the backup link immediately.
  - OSPF IP FRR and BFD can be bound to rapidly detect link faults. This ensures that traffic is rapidly switched to the backup link in the case of link failures.
- Command: [Huawei-ospf-1] **bfd all-interfaces** { **min-rx-interval** *receive-interval* | **min-tx-interval** *transmit-interval* | **detect-multiplier** *multiplier-value* | **frr-binding** }
  - **min-rx-interval** *receive-interval*: specifies an expected minimum interval for receiving BFD packets from the peer. The value is an integer ranging from 10 to 2000, in milliseconds. The default value is 1000.
  - **min-tx-interval** *transmit-interval*: specifies a minimum interval for sending BFD packets to the peer. The value is an integer ranging from 10 to 2000, in milliseconds. The default value is 1000.
  - **detect-multiplier** *multiplier-value*: specifies a local detection multiplier. The value is an integer ranging from 3 to 50. The default value is 3.
  - **frr-binding**: binds the BFD session status to the link status of an interface. If a BFD session goes down, the physical link of the bound interface also goes down, triggering traffic to be switched to the backup link.

Contents

HUAWEI

Overview of OSPF Route Control

OSPF route control includes:

- Adjusting the OSPF interface cost
- Setting the maximum number of equal-cost routes for load balancing
- Importing external routes
- Configuring route summarization
- Configuring default route advertisement
- Configuring filter-policies
- Configuring OSPF to filter outgoing LSAs
- Configuring an ABR to filter Type3 LSAs
- Setting the maximum number of External LSAs in the LSDB

HUAWEI

- This course describes only equal-cost routes, default routes, and LSA filtering. For other information, see HCIP-Datacom-Core Technology.

Equal-Cost Route

- If the destinations and costs of the multiple routes discovered by one routing protocol are the same, these routes are equal-cost routes and can participate in load balancing.
- The device sends packets to the same destination address through multiple equal-cost routes in load balancing mode.

Set the maximum number of equal-cost routes for load balancing.

```
[Huawei-ospf-1] maximum load-balancing number
```


- Command: [Huawei-ospf-1] **maximum load-balancing** *number*
  - *number*: specifies the maximum number of equal-cost routes for load balancing. The value range varies according to the device model. For details, see the product documentation of the corresponding device.

Example for Configuring the Number of Equal-Cost Routes for Load Balancing

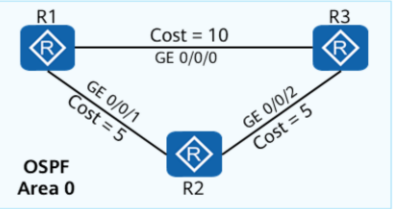

Device	Interface	IP Address
R1	GE 0/0/0	10.1.13.1/24
	GE 0/0/1	10.1.12.1/24
R2	GE 0/0/1	10.1.12.2/24
	GE 0/0/2	10.1.23.2/24
R3	Loopback0	10.1.3.3/32
	GE 0/0/0	10.1.13.3/24
	GE 0/0/2	10.1.23.3/24

It is required that R1 can access the loopback interface address of R3 through the path R1 -> R3 or the path R1 -> R2 -> R3.

1. Assign an IP address to each interface and configure OSPF on each device. (Details are not provided here.)

2. Configure the maximum number of OSPF equal-cost routes for load balancing on R1.

```
[R1] ospf
[R1-ospf-1] maximum load-balancing 2
```

3. Verify the configuration.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
10.1.3.3/32         OSPF    10   10          D   10.1.13.3       GigabitEthernet0/0/0
                    OSPF    10   10          D   10.1.12.2       GigabitEthernet0/0/1
```

Default Route

- On the area border and AS border of an OSPF network generally reside multiple routers for egress backup or traffic load balancing. In this case, a default route can be configured to reduce the number of routing entries in the routing table and ensure high availability of the network.
- OSPF default routes are generally applied to the following scenarios:
  - An ABR advertises Type 3 LSAs carrying the default route so that routers in an area forward inter-area packets accordingly.
  - An ASBR advertises Type 5 or Type 7 LSAs carrying the default route so that routers in an AS forward AS-external packets.

Area Type	Trigger Condition	Advertised by	LSA Type	Flooding Scope
Common area	The **default-route-advertise** command is run.	ASBR	Type 5 LSA	Common area
Stub area and totally stubby area	Automatically generated	ABR	Type 3 LSA	Stub area
NSSA	The **nssa [default-route-advertise]** command is run.	ASBR	Type 7 LSA	NSSA
Totally NSSA	Automatically generated	ABR	Type 3 LSA	NSSA

HUAWEI

- Default routes have all 0s as the destination address and mask. A device uses a default route to forward packets when no matching route is available. Hierarchical management of OSPF routes prioritizes the default route carried in Type 3 LSAs over the default route carried in Type 5 or Type 7 LSAs.
- Common area:
  - By default, routers in a common OSPF area do not generate default routes. To enable a router in a common OSPF area to advertise a default route to OSPF, run the **default-route-advertise** command on the router. After the command is run, the router generates a default ASE LSA (Type 5 LSA) and advertises it to the entire OSPF AS.
- Stub area:
  - Type 5 LSAs cannot be advertised within a stub area. All routers within a stub area can learn AS external routes only through an ABR.
  - The ABR in a stub area automatically generates a default Type 3 LSA and advertises it to the entire stub area. The ABR uses the default route to divert traffic destined for a destination outside the AS to itself and then forwards the traffic.

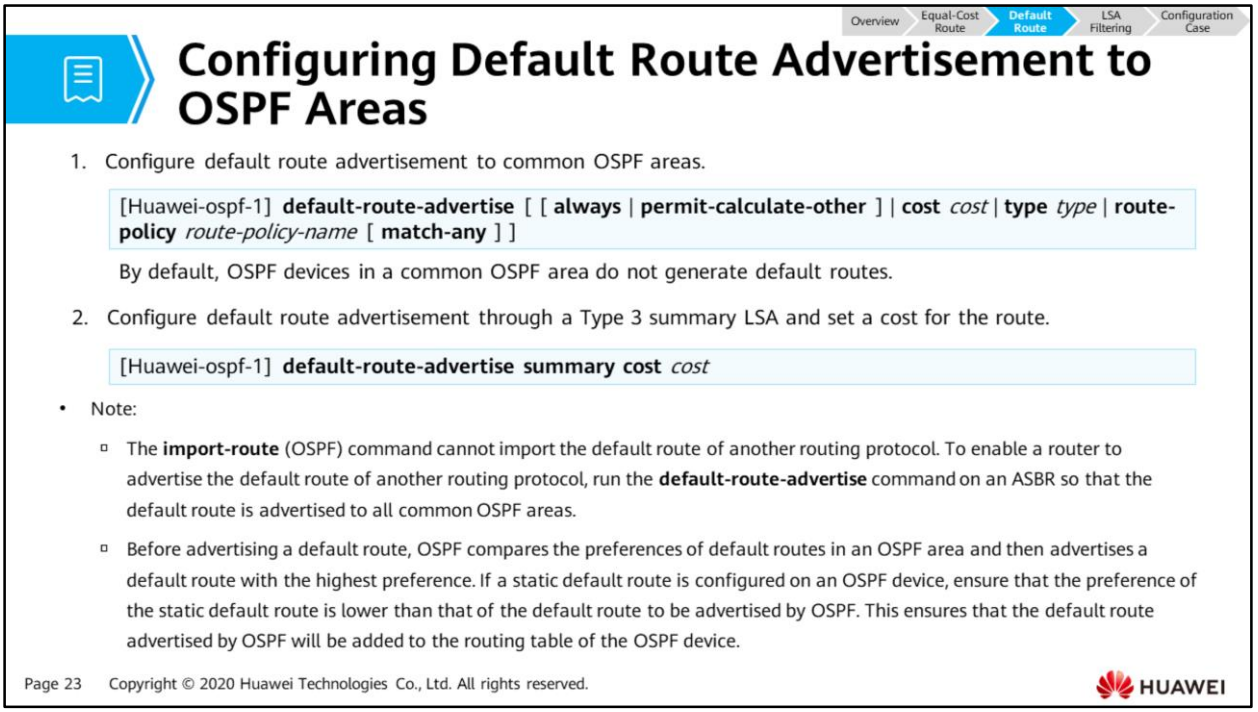

- Command: [Huawei-ospf-1] **default-route-advertise** [[**always | permit-calculate-other] | cost** *cost* | **type** *type* | **route-policy** *route-policy-name* [**match-any**]]
  - **always**: An LSA that describes the default route is generated and advertised regardless of whether the local device has an active default route that does not belong to the current OSPF process.
    - If **always** is configured, the device does not calculate the default routes from other devices.
    - If **always** is not configured, an LSA that describes the default route can be generated only if an active default route that does not belong to the current OSPF process exists in the routing table of the local device.
  - **permit-calculate-other**: An LSA that describes the default route is generated and advertised only if the device has an active default route that does not belong to the current OSPF process, and the device still calculates the default routes from other devices.
  - **type** *type*: specifies the type of an external route. The value is 1 or 2. The default value is 2.
    - 1: Type 1 external route
    - 2: Type 2 external route
  - **route-policy** *route-policy-name*: specifies the name of a route-policy. The device advertises default routes according to the configuration of the route-policy when the routing table of the device contains a default route that matches the route-policy but does not belong to the current OSPF process. The value is a string of 1 to 40 case-sensitive characters. If spaces are used, the string must start and end with double quotation marks (").

Configuring OSPF to Filter Outgoing LSAs

- When multiple links exist between two routers, you can enable the function to filter outgoing LSAs, preventing them from being sent to particular links. This function can help reduce unnecessary retransmission of LSAs and reduce bandwidth consumption.
- Filtering the outgoing LSAs on a specified OSPF interface can prevent unwanted LSAs from being sent to neighbors, thus reducing the LSDB sizes of the neighbors and speeding up network convergence.

To filter outgoing LSAs on an OSPF interface, run the following command:

[Huawei-GigabitEthernet0/0/1] **ospf filter-lsa-out** { **all** | { **summary** [**acl** { *acl-number* | *acl-name* }] | **ase** [**acl** { *acl-number* | *acl-name* }] | **nssa** [**acl** { *acl-number* | *acl-name* }] } }

The command configuration does not take effect for the LSAs that have been sent out before the command is run. The aging time of such an LSA is still 3600 seconds.

HUAWEI

- Command: [Huawei-GigabitEthernet0/0/1] **ospf filter-lsa-out** { **all** | { **summary** [**acl** { *acl-number* | *acl-name* }] | **ase** [**acl** { *acl-number* | *acl-name* }] | **nssa** [**acl** { *acl-number* | *acl-name* }] } }
  - **all**: filters all outgoing LSAs, except grace LSAs.
  - **summary**: filters outgoing network-summary-LSAs (Type 3).
  - **ase**: filters outgoing AS-external-LSAs (Type 5).
  - **nssa**: filters outgoing NSSA-LSAs (Type 7).
  - **acl** *acl-number*: specifies the number of a basic ACL. The value is an integer ranging from 2000 to 2999.
  - **acl** *acl-name*: specifies the name of an ACL. The value is a string of 1 to 32 case-sensitive characters. It cannot contain spaces and must start with a letter (a to z or A to Z).

Configuring OSPF to Filter ABR Type 3 LSAs

- Configure filtering policies for incoming and outgoing ABR Type 3 LSAs (network-summary-LSAs) in an area, allowing only those that pass the filtering to be sent and accepted.
- Filtering Type 3 LSAs in a specified OSPF area can prevent unwanted LSAs from being sent to neighbors, thus reducing the LSDB sizes and speeding up network convergence.

To filter outgoing Type3 LSAs in an OSPF area, run the following command:

[Huawei-ospf-1-area-0.0.0.1] **filter** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **export**

To filter incoming Type3 LSAs in an OSPF area, run the following command:

[Huawei-ospf-1-area-0.0.0.1] **filter** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **import**

- Command: [Huawei-ospf-1-area-0.0.0.1] **filter** { *acl-number* | **acl-name** *acl-name* | **ip-prefix** *ip-prefix-name* | **route-policy** *route-policy-name* } **export**
  - *acl-number*: specifies the number of a basic ACL. The value is an integer ranging from 2000 to 2999.

Overview of OSPF Database Overflow

- The LSDBs of OSPF devices in the same area are synchronized after routes are converged. However, achieving such a state can be difficult as the number of routes on a network continuously increases, causing some devices to be unable to carry excess routing information due to limited system resources. This is called an OSPF database overflow.
- One way to solve such an issue is to configure stub areas or NSSAs, which reduces the amount of routing information on devices. However, such an approach cannot prevent an OSPF database overflow caused by a sharp increase in dynamic routes. To resolve this issue, set the maximum number of non-default external LSAs allowed in the LSDB of a device to dynamically control the LSDB size.

  Set the maximum number of non-default external LSAs allowed in the LSDB.

  [Huawei-ospf-1] **lsdb-overflow-limit** *number*

HUAWEI

- When the number of external LSAs (Type 5 and Type 7) imported by OSPF exceeds the maximum number supported, excessive external LSAs cannot be processed properly and are discarded. To address this issue, you can set a proper upper limit for the number of non-default external LSAs in the LSDB, so as to adjust and optimize the OSPF network.
- Command: [Huawei-ospf-1] **lsdb-overflow-limit** *number*
  - *number*: specifies the maximum number of non-default external LSAs allowed in the LSDB. The value is an integer ranging from 1 to 1000000.

Preventive Measures for OSPF Database Overflows

- Setting the maximum number of non-default external routes on a router can prevent OSPF database overflows.
- Set the same maximum number for all routers on an OSPF network. When the number of non-default external routes on a router reaches the maximum number, the router enters the Overflow state and starts the Overflow state timer (default timeout period: 5s). After the timer expires, the router automatically exits the Overflow state.

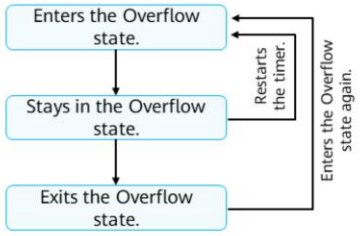

Overflow Phase	OSPF Processing
Enters the Overflow state.	• Deletes all non-default external routes generated by the router itself. • Starts the Overflow state timer.

Stays in the Overflow state.	• Does not generate non-default external routes.
	• Discards newly received non-default external routes and does not reply with LSAck packets.
	• Checks whether the number of external routes still exceeds the maximum number when the overflow state timer expires.
	  ▪ If not, the router exits the Overflow state.
	  ▪ If so, the router restarts the Overflow state timer.

Exits the Overflow state.	• Deletes the Overflow state timer. • Generates non-default external routes. • Accepts newly received non-default external routes and replies with LSAck packets. • Gets ready to enter the Overflow state again.

Example for Configuring OSPF Route Control (1)

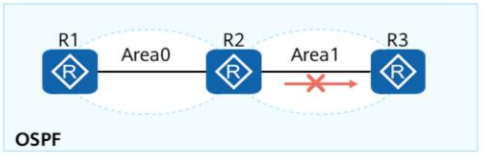

To reduce the number of LSAs on R3 and ensure that R3 can properly communicate with routers in other areas, check that the following requirements are met:

- R2 does not inject Type 3 LSAs into Area 1.
- R2 advertises default routes.

1. Assign an IP address to each interface and configure OSPF on each device. (omitted).
2. Configure R2 to filter Type 3 LSAs.

```
[R2] acl 2000
[R2-acl-basic-2000] rule deny
[R2-acl-basic-2000] quit
[R2] ospf
[R2-ospf-1] area 1
[R2-ospf-1-area-0.0.0.1] filter 2000 import
```

3. Configure R2 to advertise default routes.

```
[R2] ospf
[R2-ospf-1] default-route-advertise always
```

Example for Configuring OSPF Route Control (2)

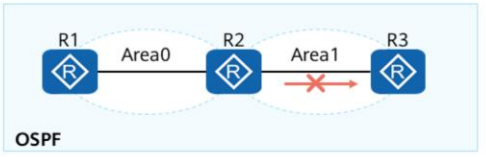

Type3 LSA

- The LSDB of R3 in Area1 does not contain Type 3 LSAs but contains a default Type 5 LSA.
- R3 can access devices in other areas through the default route.

1. Check the LSDB of R3.

```
[R3]display ospf lsdb
         OSPF Process 1 with Router ID 10.1.23.3
                     Link State Database
                   Area: 0.0.0.1
Type       LinkState ID    AdvRouter       Age  Len  Sequence   Metric
Router     10.1.23.3       10.1.23.3       731  48   80000004   1
Router     10.1.12.2       10.1.12.2       406  36   80000008   1
Network    10.1.23.2       10.1.12.2       730  32   80000002   0

                     AS External Database
Type       LinkState ID    AdvRouter       Age  Len  Sequence   Metric
External   0.0.0.0         10.1.12.2       406  36   80000001   1
```

2. Check the routing table of R3.

```
[R3]dis ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 9        Routes : 9
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
0.0.0.0/0           O_ASE   150  1         D     10.1.23.2       GigabitEthernet0/0/1
...
```

OSPF Route Control Case Analysis

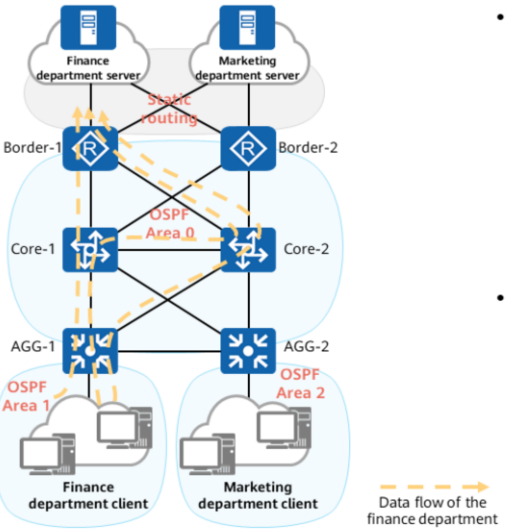

- Network deployment:
  - An enterprise network is divided into two networks, one for the finance department and the other for the marketing department.
  - The enterprise network uses OSPF to allow internal network connectivity. The backbone network is deployed in area 0, clients of the finance department's network are deployed in area 1, and clients of the marketing department's network are deployed in area 2.
  - Border devices access department servers through static routes, which are imported to the OSPF process.
- Requirements:
  - As long as the border-1 router and its uplink work properly, the data flows of the finance department are forwarded only through the border-1 router.
  - As long as the core-1 router and its uplink work properly, the data flows of the finance department are forwarded only through the core-1 router.
  - For details about the data forwarding requirements of the marketing department, see the comment of this slide.

  HUAWEI

- Data forwarding requirements of the marketing department:
  - As long as the border-2 router and its uplink work properly, the data flows of the marketing department are forwarded only through the border-2 router.
  - As long as the core-2 router and its uplink work properly, the data flows of the marketing department are forwarded only through the core-2 router.
- This case uses the data forwarding path of the finance department as an example. The data forwarding path of the marketing department is not described here.

Requirement Analysis

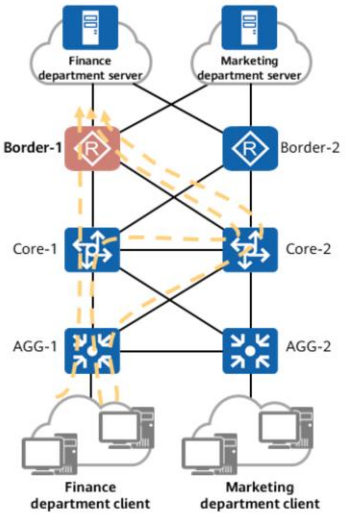

1. Controlling the network egress for data forwarding:
   - Data of the finance department is always forwarded through border-1.
   - Data of the marketing department is always forwarded through border-2.

To ensure that fixed ASBRs are used to forward data, the internal network changes must be ignored. That is, the internal route cost is not calculated.
—Use Type 2 external routes of OSPF.

2. Controlling the precise internal path of data flows:
   - Load-balancing path should not exist.

Data needs to be sent to a specific ASBR along the planned path on the network.
—Adjust the internal path cost.

- Type 2 external route:
  - Because a Type 2 external route offers low reliability, its cost is considered to be much greater than the cost of any internal route to an ASBR.
  - Cost of a Type 2 external route = Cost of the route from an ASBR to the destination

Traffic Egress Control

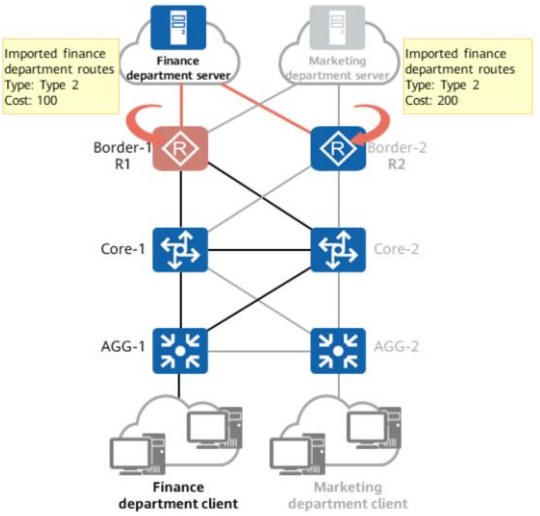

- Implementation:
  - Import static routes destined for the finance department server to the OSPF processes of R1 (border-1) and R2 (border-2) to implement egress backup through route-policies.
  - Set the type of the imported external route to Type 2.
  - On R1, set the cost of the external route to 100; on R2, set the cost of the external route to 200.
- Configuration result:
  - When there are two Type 2 external routes with different costs on the same network segment, the network device prefers the route with a smaller cost. In this case, each network device preferentially selects R1 as the egress.

HUAWEI

- The internal path cost to each ASBR is not considered during traffic egress control.

Controlling Internal Paths

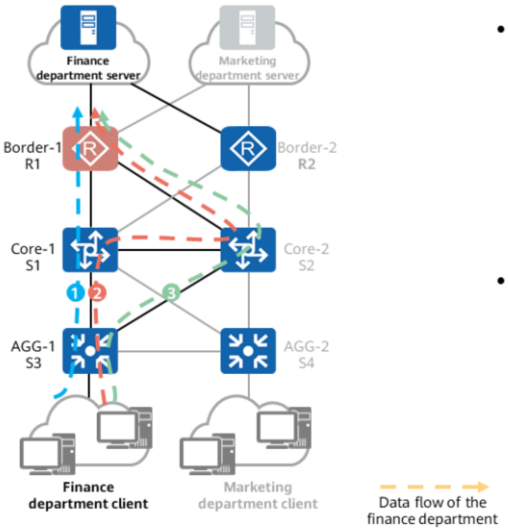

- Network requirement analysis:
  - If the network is running properly, S3 (AGG-1) selects path 1.
  - If the link between S1 (core-1) and R1 fails, S3 selects path 2.
  - If S1 fails, S3 selects path 3.
- Implementation:
  - Path 1-cost < Path 2-cost < Path 3-cost

Controlling Internal Paths: Adjusting the Cost Between Aggregation Devices and Core Devices

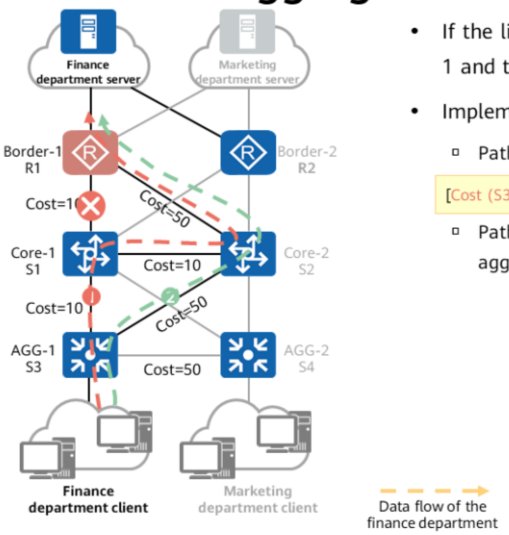

- If the link between S1 and R1 fails, S3 preferentially selects path 1 and then path 2 because S1 is working properly.
- Implementation:
  - Path 1-cost < Path 2-cost, that is:

[Cost (S3 - S1) + Cost (S1 - S2) + Cost (S2 - R1)] < [Cost (S3 - S2) + Cost (S2 - R1)]

  - Path 1 can be achieved by adjusting the path cost between aggregation devices and core devices.

Controlling the Internal Path: Adjusting the Cost Between Core Devices and Boundary Devices

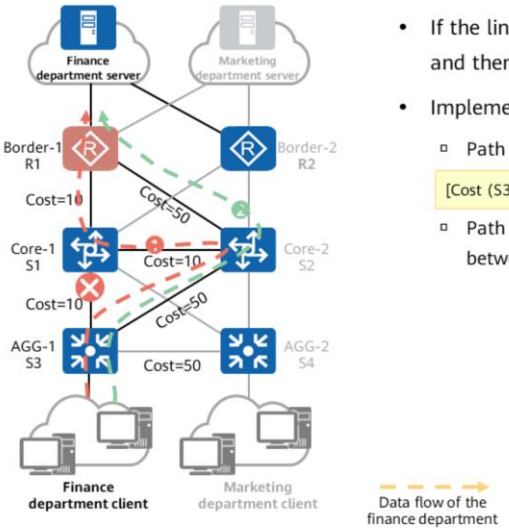

- If the link between S3 and S1 fails, S3 preferentially selects path 1 and then path 2 because S1 is working properly.
- Implementation:
  - Path 1-cost < Path 2-cost, that is:

[Cost (S3 - S2) + Cost (S2 - S1) + Cost (S1 - R1)] < [Cost (S3 - S2) + Cost (S2 - R1)]

  - Path 1 can be preferentially selected by adjusting the path cost between core devices and boundary devices.

Example for Configuring OSPF Route Control (1)

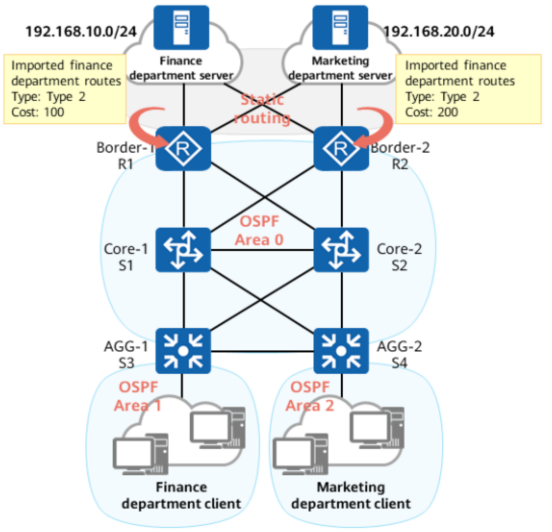

Configure a route-policy on R1 and set the cost to 100.

```
[R1] acl 2000
[R1-acl-basic-2000] rule permit source 192.168.10.0 0.0.0.255
[R1-acl-basic-2000] quit

[R1] route-policy static2ospf permit node 10
[R1-route-policy] if-match acl 2000
[R1-route-policy] apply cost 100
[R1-route-policy] quit
[R1] route-policy static2ospf permit node 20
[R1-route-policy] quit
```

Import static routes to the OSPF process on R1.

```
[R1] ospf
[R1-ospf-1] import-route static route-policy static2ospf type 2
```

Example for Configuring OSPF Route Control (2)

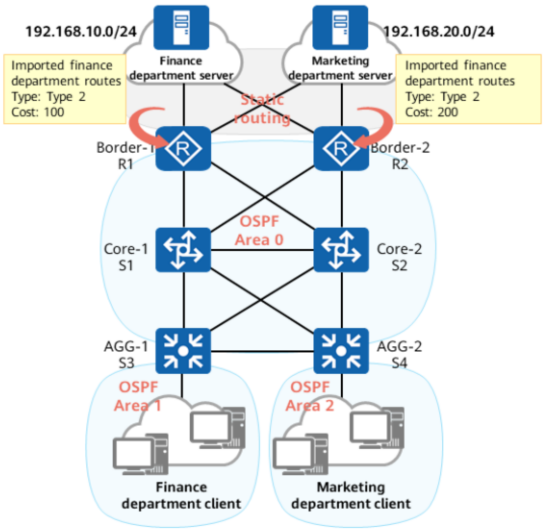

Configure a route-policy on R2 and set the cost to 200.

```
[R2] acl 2000
[R2-acl-basic-2000] rule permit source 192.168.10.0 0.0.0.255
[R2-acl-basic-2000] quit

[R2] route-policy static2ospf permit node 10
[R2-route-policy] if-match acl 2000
[R2-route-policy] apply cost 200
[R2-route-policy] quit
[R2] route-policy static2ospf permit node 20
[R2-route-policy] quit
```

Import static routes to the OSPF process on R2.

```
[R2] ospf
[R2-ospf-1] import-route static route-policy static2ospf type 2
```

Example for Configuring OSPF Route Control (3)

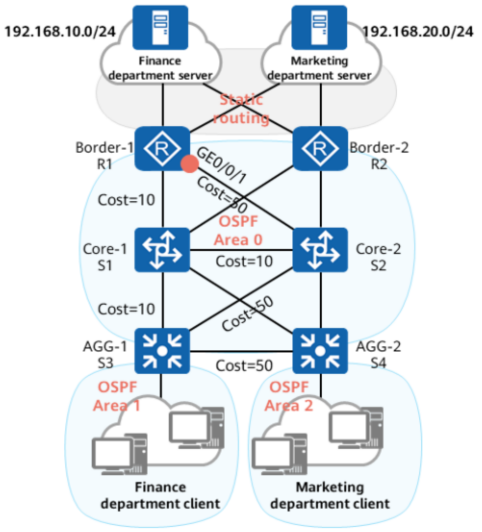

Set the OSPF cost of an interface (GE0/0/1 of R1 is used as an example).

```
[R1] interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1] ospf cost 50
```

Contents

OSPF Multi-Process

- OSPF supports multiple processes that can separately run on the same device and do not affect each other. Route exchange between different OSPF processes is similar to route exchange between different routing protocols.
- An interface on a router can belong to only one OSPF process.

- Usage scenario:
  - A typical application of OSPF multi-process is in the VPN scenario.
  - As shown in the figure, a PE connects to two different VPN customers, and OSPF is deployed between the PE and CEs. Therefore, multiple processes can be deployed on the PE to isolate the VPN customers.

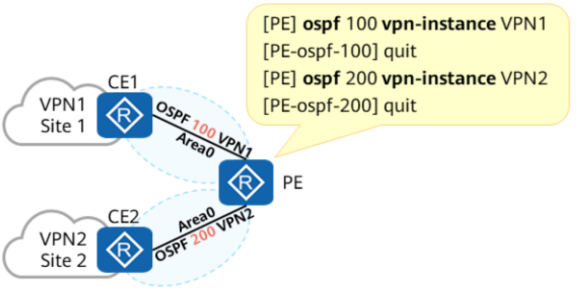

- VPN: virtual private network
- If a VPN instance is specified for an OSPF process that is to be created, the OSPF process belongs to this instance. Otherwise, the OSPF process belongs to the global instance.

Association Between OSPF and BGP (1)

When a new device is added or a device is restarted, network traffic may be lost during BGP convergence. This is because IGP route convergence is faster than BGP route convergence.

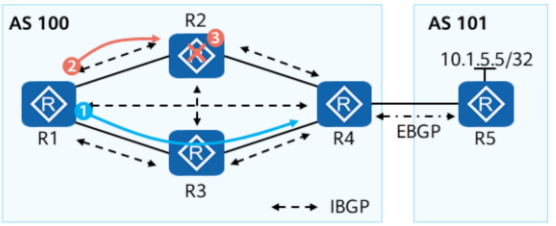

OSPF runs on R1, R2, R3, and R4 and full-mesh IBGP connections are established. R3 is the backup device of R2.

When the network is stable, the traffic from R1 to 10.1.5.5/32 passes through the path R1 -> R2 -> R4 -> R5.

1. If R2 fails, traffic is switched to the path R1 -> R3 -> R4 -> R5.
2. After R2 recovers, OSPF converges first because IGP route convergence is faster than BGP route convergence. If R1 needs to access 10.1.5.5/32, it searches for a BGP route with the next hop being R5. R2 then searches for the IGP route and sends traffic to R2 over the route.
3. After receiving the traffic, R2 searches for the BGP route. Because BGP route convergence is not complete, R2 does not find a route to 10.1.5.5/32 and therefore does not forward traffic. As a result, traffic is lost.

Association Between OSPF and BGP (2)

- OSPF-BGP synchronization can be enabled to prevent traffic loss.
- After OSPF-BGP synchronization is enabled on a device, the device remains as a stub router within the set synchronization period. That is, the link metric in the LSA advertised by the device is the maximum value 65535. Therefore, the device instructs other OSPF devices not to use it for data forwarding.

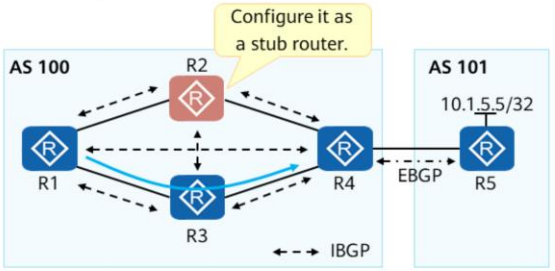

Enable BGP association on R2. In this way, R1 continues to forward traffic through R3, but does not forward traffic to R2 until BGP route convergence on R2 is complete.

Configure the stub router.

[Huawei-ospf-1] **stub-router** [**on-startup** [*interval*]]

- Configuring a stub router is a special route selection method. The path with a stub router configured is not preferred.
- The implementation is to set the metric to the maximum value (65535) to prevent data from being forwarded through the router. It is used to protect the link of the router and is usually used in maintenance scenarios, such as upgrade.

- Command: [Huawei-ospf-1] **stub-router** [**on-startup** [*interval*]]
  - **on-startup** [*interval*]: specifies the interval for a device to remain as a stub router when the device restarts or fails. The value is an integer ranging from 5 to 65535, in seconds. The default value is 500s.
    - If **on-startup** is not configured, the device is always a stub router. That is, the cost of all routes sent by this device is 65535.
    - If **on-startup** is specified, the device remains as a stub router only when it restarts or fails. The duration is determined by *interval.* If *interval* is not specified, the default value of 500s is used.

OSPF Forwarding Address

Forwarding address (FA):

- The FA is the address to which a data packet is forwarded before the packet reaches the advertised destination address. If the forwarding address is 0.0.0.0, the packet is forwarded to the originating ASBR.
- Type 5 AS-External LSAs and Type 7 NSSA LSAs:

LS Age	Options	LS Type
Link State ID		
Advertising Router		
LS Sequence Number		
LS Checksum	Length	
Network Mask		
E	0	Metric
Forwarding Address		
External Route Tag		
...		

OSPF Type 5 and Type 7 LSAs contain a special FA field. The introduction of FA enables OSPF to avoid the sub-optimal path problem in some special scenarios.

Problem Occurring When No FA Is Used

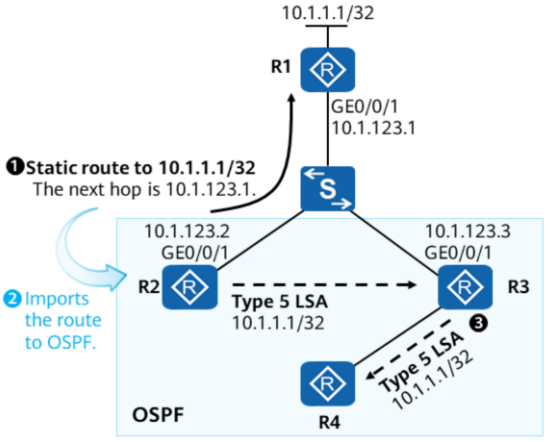

- OSPF runs on R2, R3, and R4 that are deployed in Area 0. OSPF is enabled on GE0/0/1 of R2 and GE0/0/1 of R3, and OSPF adjacencies are established between the two routers. However, no OSPF adjacencies are established between R1 and the two routers.
  1. Configure a static route destined for 10.1.1.1/32 on R2, with 10.1.123.1 as the next hop.
  2. R2 imports the static route to OSPF and generates a Type 5 LSA to be flooded in the area.
  3. After R3 receives the Type 5 LSA from R2, it calculates an external route to 10.1.1.1/32 and sets the next-hop address of the route to R2 (10.1.123.2).
- The path from a router (for example, R4) in the OSPF domain to 10.1.1.1/32 is R4 -> R3 -> R2 -> R1, which is the sub-optimal path.

Using FA to Solve the Suboptimal Path Problem

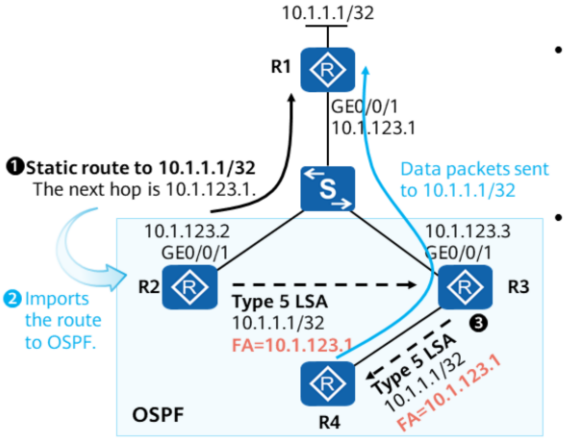

- When advertising an external route destined for 10.1.1.1/32 in the OSPF area, R2 sets the FA of the corresponding Type 5 LSA to 10.1.123.1, which is the next hop of the external route.
- When R3 receives the LSA, it calculates the route to 10.1.1.1/32 and finds that the FA is not 0. Therefore, R3 considers that the next hop to 10.1.1.1/32 is the address specified by the FA, that is, 10.1.123.1.

FA Values

- When an ASBR imports external routes, if the FA field in the Type 5 LSA is 0, the router considers that the data packets destined for the destination network segment should be sent to the ASBR. If the FA field in a Type 5 LSA is not 0, the router considers that the data packet destined for the destination network segment should be sent to the device identified by the FA.
- The FA field can be set to a non-zero value only when all the following conditions are met:
  - The ASBR activates OSPF on the interface (outbound interface of the external route) connected to the external network.
  - The preceding interface is not configured as a silent interface.
  - The OSPF network type of such an interface is broadcast or NBMA.
  - The IP address of the interface is within the network segment specified by the network command in the OSPF configuration.
- The route to the FA must be an OSPF intra-area route or an inter-area route. In this manner, the router that receives the external LSA can add the LSA into the routing table. The next hop of the route generated using the loaded external LSA is the same as the next hop to the FA.

  HUAWEI

- Type 7 LSAs in an NSSA are translated into Type 5 LSAs.
  - To advertise external routes imported by an NSSA to other areas, Type 7 LSAs must be translated into Type 5 LSAs. By default, the translator is the ABR with the largest router ID in an NSSA.
  - The propagate bit (P-bit) in the Options field of an LSA header is used to notify a translator whether the Type 7 LSA needs to be translated into a Type 5 LSA. A Type 7 LSA can be translated into a Type 5 LSA only when the P-bit is set to 1 and the FA is not 0.
  - The P-bit is not set for Type 7 LSAs generated by an ABR.
- Note: All OSPF LSAs have the same LSA header, and the P-bit is in the Options field of the LSA header.

Case: Typical FA Application in NSSA Scenarios

- When multiple ABRs exist in an NSSA, the system automatically selects an ABR as a translator to translate Type 7 LSAs into Type 5 LSAs. Other ABRs do not perform LSA translation.
- As shown in the figure, if the FA is not considered, R1 considers that the packet must pass through the ABR (R3) to reach the destination because the router ID of R3 is greater than that of R2. In this way, traffic is diverted to the low-bandwidth link R1 -> R3 -> R4 -> R5.

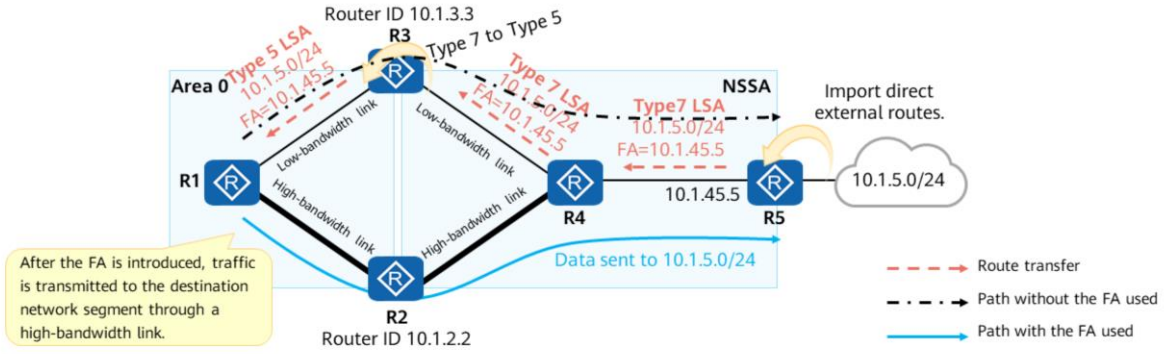

- As shown in the figure:
  - Configure R5 to import direct external routes and set the IP address of the FA to 10.1.45.5, which is used by R5 to access the destination network segment 10.1.5.0/24.
  - R3 translates Type 7 LSAs into Type 5 LSAs and the LSAs continue to carry the FA 10.1.45.5.
  - Upon receipt, R1 searches its OSPF routing table for a route to the FA and uses the next hop address of the route as the next hop address of the external route.
  - Therefore, R1 will finally access the destination network segment 10.1.5.0/24 through the path R1 -> R2 -> R4 -> R5.

Contents

HUAWEI

Overview of IS-IS Fast Convergence

- IS-IS fast convergence is an extended feature designed to speed up route convergence. It provides a series of functions, covering incremental SPF (I-SPF), PRC, intelligent timer, and LSP fast flooding.
- IS-IS supports fast convergence after a fault is rectified. For example, IS-IS auto FRR can be used to quickly switch traffic to a backup link, and IS-IS can be associated with BFD to quickly detect faults.

HUAWEI

- The functions, including PRC, intelligent timer, and FRR, of IS-IS are similar to those of OSPF and therefore not detailed here.

I-SPF

I-SPF implementation: When the network topology changes, I-SPF recalculates routes only for affected nodes instead of all nodes, speeding up route calculation.

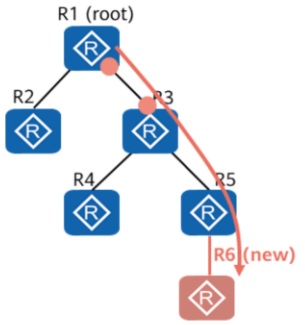

In route calculation, a node represents a router, and a leaf represents a route. I-SPF processes the information of only changed nodes.

- Scenario description:
  - On an IS-IS network, an SPT with R1 as the root is calculated after route convergence, as shown in the left figure. When R1 accesses R5, the traffic of R1 is forwarded to R5 based on [the outbound interface of R1's downlink and the IP address of the inbound interface of R3's uplink].
  - R6 is added as a downstream device of R5. IS-IS is enabled on R6, meaning that there is a new network node on the IS-IS network.
- I-SPF calculation:
  - R5 and R6 both flood LSPs carrying information about their neighbor relationship on the entire network.
  - After receiving such an LSP, R1 performs I-SPF calculation only for R5 and R6 to generate a new SPT.
  - Therefore, when R1 accesses R5 and R6, the traffic of R1 is forwarded to R5 and R6 based on [the outbound interface of R1's downlink and the IP address of the inbound interface of R3's uplink].

  HUAWEI

- SPF for route calculation: If a node on a network changes, SPF recalculates routes for all the nodes on the network, which takes a long time, consumes a large number of CPU resources, and consequently reduces the network-wide convergence speed.
- I-SPF is an improvement of SPF. Unlike SPF that calculates all nodes, I-SPF calculates only affected nodes. The SPT generated using I-SPF is the same as that generated using SPF. This significantly decreases CPU usage and speeds up network convergence.
- I-SPF and PRC are used together on an IS-IS network.
  - If the SPT calculated by I-SPF changes, PRC processes all the leaves (routes) of only the changed node.
  - If the SPT calculated by I-SPF does not change, PRC processes only the changed leaves (routes). For example, if IS-IS is newly enabled on an interface of a node, the SPT on the network remains unchanged. In this case, PRC updates only the routes of this interface, which consumes less CPU resources.

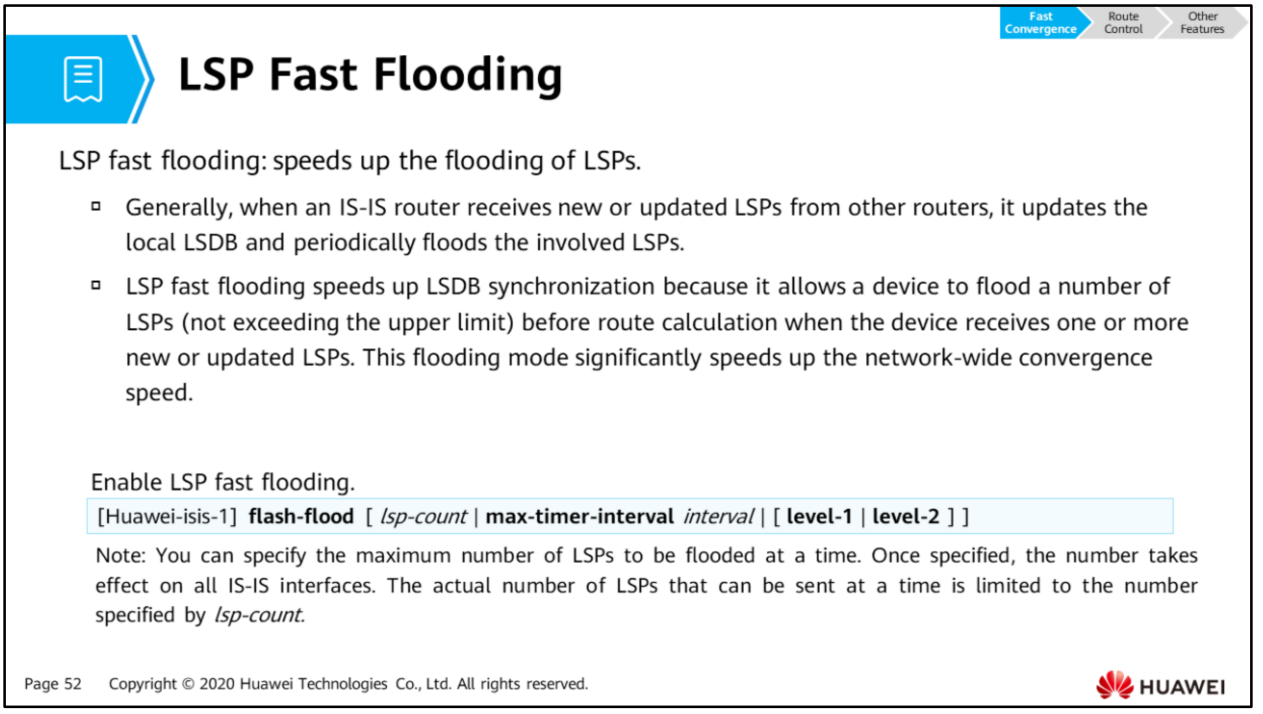

- Command: [Huawei-isis-1] **flash-flood** [*lsp-count* | **max-timer-interval** *interval* | [**level-1** | **level-2**]]
  - *lsp-count*: specifies the maximum number of LSPs that can be flooded on each interface at a time. The value is an integer ranging from 1 to 15. The default value is 5.
  - **max-timer-interval** *interval:* specifies the maximum interval at which LSPs are flooded. The value is an integer ranging from 10 to 50000, in milliseconds. The default value is 10.
  - **level-1**: enables the LSP flash-flood function in the Level-1 area. If no level is specified in the command, this function is enabled in both Level-1 and Level-2 areas.
  - **level-2**: enables the LSP flash-flood function in the Level-2 area. If no level is specified in the command, this function is enabled in both Level-1 and Level-2 areas.

Overview of IS-IS Route Control

In real-world applications, IS-IS routes calculated using SPF sometimes cannot meet network planning and traffic management requirements, which may lead to various problems, such as slow table lookup due to a large number of routing entries in routing tables and unbalanced link usage on a network. To optimize IS-IS networks and facilitate traffic management, more precise route control is required. You can use any of the following methods to implement the control:

- Adjusting the IS-IS preference
- Adjusting the IS-IS interface cost
- Configuring equal-cost routes
- Configuring IS-IS route leaking
- Configuring default route advertisement
- Importing external routes
- Configuring filter-policies

HUAWEI

- This course involves only equal-cost and default routes. For details about other control methods, see *HCIP-Datacom-Core Technology*.

Equal-Cost Route

If there are multiple redundant links on an IS-IS network, multiple equal-cost routes to the same destination may exist. In this case, you can use either of the following methods to configure equal-cost routes:

- Configure load balancing so that traffic is evenly distributed to relevant links.
  - This method improves link utilization and reduces the possibility of congestion caused by overloaded links. However, because traffic will be randomly forwarded, this method may make traffic management difficult.
- Configure a preference for each equal-cost route so that the route with the highest preference is preferentially selected and the others function as backups.
  - This method is used to specify the preferred route among multiple equal-cost routes, without the need to modify original configurations. It facilitates traffic management and improves network reliability.
  - Note: After preferences are configured for equal-cost routes, IS-IS devices forward traffic to the next hop with the highest preference, instead of forwarding traffic in load balancing mode.

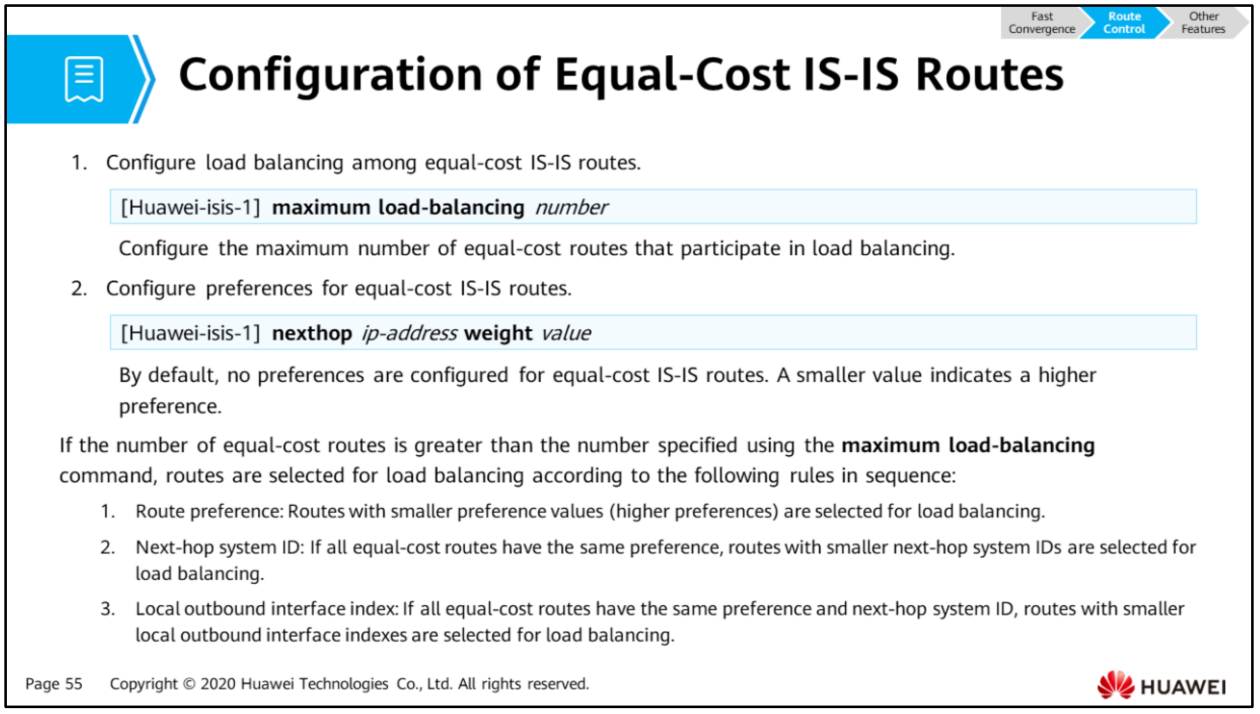

- Command: [Huawei-isis-1] **maximum load-balancing** *number*
  - *number*: specifies the maximum number of equal-cost routes that participate in load balancing. The value varies according to the device model.
- Command: [Huawei-isis-1] **nexthop** *ip-address* **weight** *value*
  - *ip-address*: specifies the IP address of the next hop. The value is in dotted decimal notation.
  - **weight** *value*: specifies the weight of the next hop. A smaller value indicates a higher preference. The value is an integer ranging from 1 to 254.

Example for Configuring Equal-Cost Routes (1)

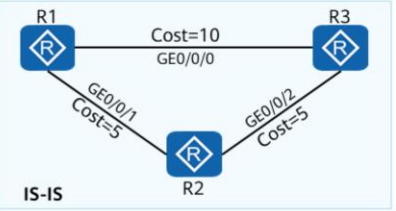

Device	Interface	IP Address
R1	GE0/0/0	10.1.13.1/24
	GE0/0/1	10.1.12.1/24
R2	GE0/0/1	10.1.12.2/24
	GE0/0/2	10.1.23.2/24
R3	Loopback0	10.1.3.3/32
	GE0/0/0	10.1.13.3/24
	GE0/0/2	10.1.23.3/24

On the network shown in the figure, IS-IS runs on R1, R2, and R3. It is required that R1 be able to access the loopback interface address of R3 through path R1 -> R3 or R1 -> R2 -> R3.

1. Assign IP addresses to device interfaces and configure IS-IS on each device. (Omitted)
2. Configure the number of IS-IS equal-cost routes for load balancing on R1.

```
[R1] isis
[R1-isis-1] maximum load-balancing 2
```

3. Verify the configuration.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
10.1.3.3/32         ISIS-L2 15   10          D   10.1.13.3       GigabitEthernet0/0/0
                    ISIS-L2 15   10          D   10.1.12.2       GigabitEthernet0/0/1
```

Example for Configuring Equal-Cost Routes (2)

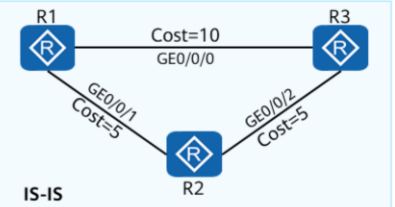

Device	Interface	IP Address
R1	GE0/0/0	10.1.13.1/24
	GE0/0/1	10.1.12.1/24
R2	GE0/0/1	10.1.12.2/24
	GE0/0/2	10.1.23.2/24
R3	Loopback0	10.1.3.3/32
	GE0/0/0	10.1.13.3/24
	GE0/0/2	10.1.23.3/24

On the network shown in the figure, IS-IS runs on R1, R2, and R3. It is required that R1 access the loopback interface address of R3 preferentially through path R1 -> R3.

1. Assign IP addresses to device interfaces and configure IS-IS on each device. (Omitted)
2. (Optional) Configure preference for equal-cost routes on R1.

```
[R1] isis
[R1-isis-1] nexthop 10.1.13.3 weight 1
[R1-isis-1] nexthop 10.1.12.2 weight 2
```

3. Verify the configuration.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
10.1.3.3/32         ISIS-L2 15   10          D   10.1.13.3       GigabitEthernet0/0/0
```

Default Route

- IS-IS allows you to control the generation and advertisement of default routes using the following methods:
  - On Level-1-2 devices, configure a rule for setting the attached (ATT) bit in Level-1 LSPs.
  - Configure Level-1 devices not to automatically generate default routes even if they receive Level-1 LSPs with the ATT bit set to 1.
  - Configure devices to advertise default routes to the IS-IS routing domain.
- LSP packet format:
  - The ATT bit is generated by a Level-1-2 router to identify whether the originating router connects to other areas. Note that the ATT field has four bits, and Huawei datacom products use only one of the four bits.

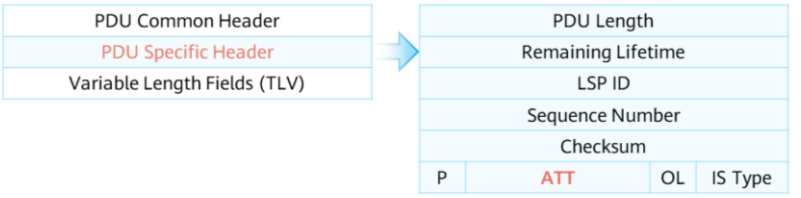

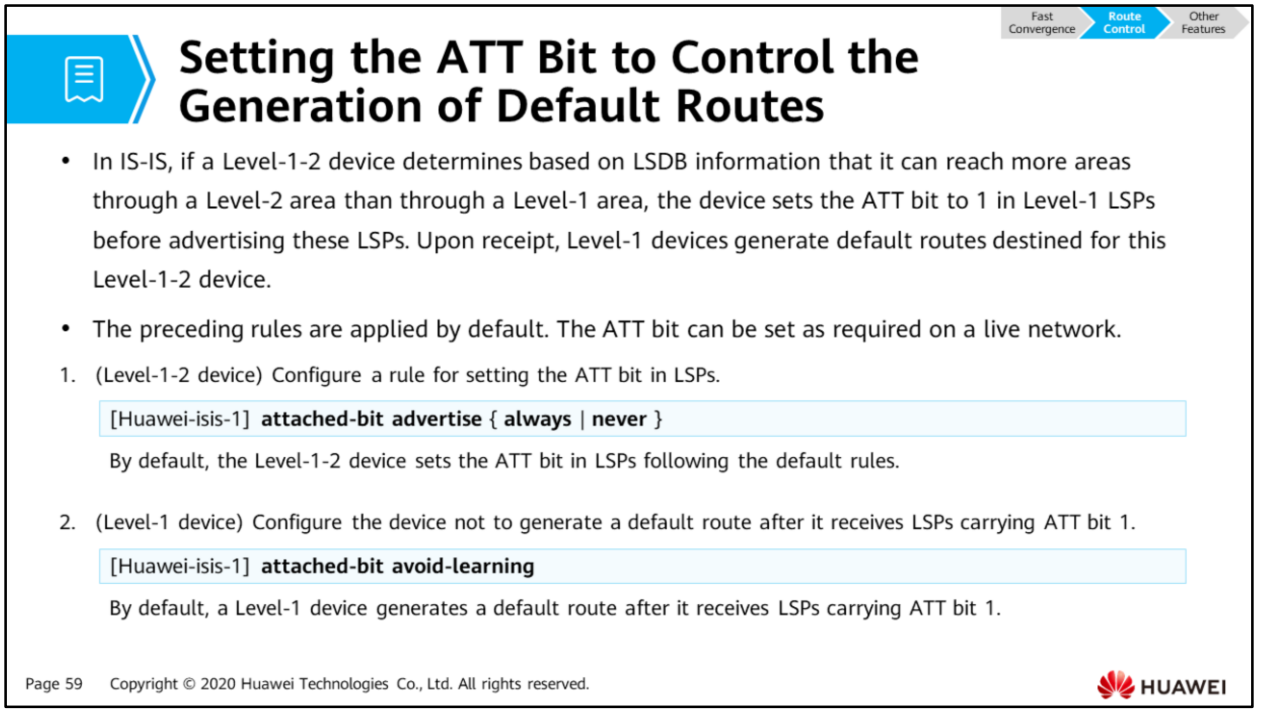

- Command: [Huawei-isis-1] **attached-bit advertise** { **always** | **never** }
  - **always**: indicates that the ATT bit is set to 1. After receiving an LSP with ATT bit 1, a Level-1 device generates a default route.
  - **never**: indicates that the ATT bit is set to 0. This prevents the Level-1 device from generating default routes and reduces the size of the routing table.
- Although the ATT bit is defined in both Level-1 and Level-2 LSPs, it is set to 1 only in Level-1 LSPs advertised by Level-1-2 devices. Therefore, this command takes effect only on Level-1-2 devices.
- To prevent Level-1 devices from advertising default routes to their routing tables, perform either of the following operations:
  - Run the **attached-bit advertise never** command on Level-1-2 devices to disable them from advertising LSPs with ATT bit 1.
  - Run the **attached-bit avoid-learning** command on Level-1 devices that connect to Level-1-2 devices.
- The difference between the preceding commands lies in that the **attached-bit avoid-learning** command applies to specified Level-1 devices.

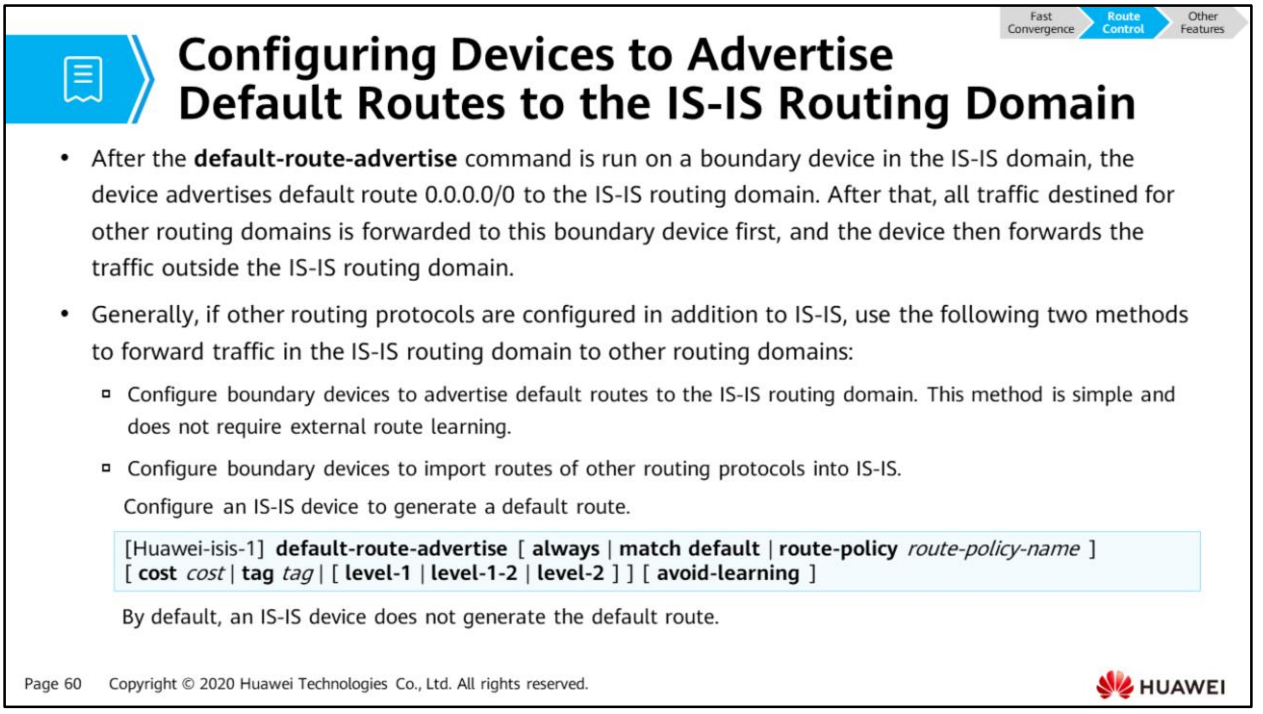

- Command: [Huawei-isis-1] **default-route-advertise** [**always** | **match default** | **route-policy** *route-policy-name*] [**cost** *cost* | **tag** *tag* | [**level-1** | **level-1-2** | **level-2**]] [**avoid-learning**]
  - **always**: configures an IS-IS device to unconditionally advertise the default route and set itself as the next hop in the route.
  - **match default**: advertises the default route generated by another routing protocol or IS-IS process through LSPs if such a route already exists in the routing table.
  - **route-policy** *route-policy-name*: specifies the name of the route-policy. A Level-1-2 device advertises the default route to the IS-IS routing domain only when external routes matching the route-policy exist in the routing table of the device. This prevents routing blackholes caused by the advertisement of the default route when link faults make some important external routes unavailable. This route-policy does not affect the import of external routes into IS-IS. The value is a string of 1 to 40 case-sensitive characters, spaces not supported. If spaces are used, the string must start and end with double quotation marks (").
  - **cost** *cost*: specifies the cost of the default route. The value is an integer. The value range depends on **cost-style**. When **cost-style** is **narrow**, **narrow-compatible**, or **compatible**, the value ranges from 0 to 63. When **cost-style** is **wide** or **wide-compatible**, the value ranges from 0 to 4261412864.

Example for Configuring Default Routes

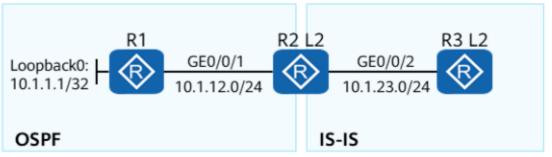

As shown in the figure, OSPF runs on GE0/0/1 and Loopback0 of R1 and GE0/0/1 of R2, and IS-IS runs on GE0/0/2 of R2 and R3.

To enable R3 to access the OSPF network, configure R2 to advertise the default route to the IS-IS routing domain.

Similarly, R2 also needs to advertise the default route to the OSPF routing domain.

1. Assign an IP address to each interface and configure OSPF/IS-IS on each device.
2. Configure R2 to advertise the default route to the IS-IS routing domain.

```
[R2] isis
[R2-isis-1] default-route-advertise always
```

3. Configure R2 to advertise the default route to the OSPF routing domain.

```
[R2] ospf
[R2-ospf-1] default-route-advertise always
```

4. Check routing entries on R3.

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 8        Routes : 8
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
0.0.0.0/0           ISIS-L2 15   10          D   10.1.23.2       GigabitEthernet0/0/2
...
```

IS-IS Multi-Instance and Multi-Process

- In IS-IS multi-instance, multiple IS-IS processes are created on the same router, with each process associated with one VPN instance.
- In IS-IS multi-process, multiple IS-IS processes are created in the same VPN instance (or in the same public network instance). These IS-IS processes are independent of each other. IS-IS processes function similarly to different routing protocols in route exchange.
- A network may carry different services, which need to be isolated for security. You can bind each IS-IS process to a different VPN instance.
- Application scenarios:
  - IS-IS multi-instance and multi-process are typically used in VPN scenarios.
  - As the figure shown, a PE is connected to two VPNs, each belonging to a different customer, and IS-IS is deployed between the PE and CEs. You can configure multiple IS-IS processes on the PE to isolate the VPN customers.

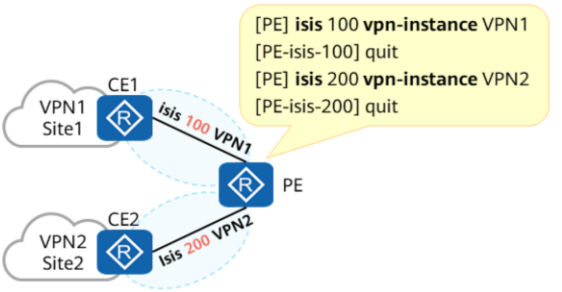

- IS-IS multi-process and multi-instance have the following characteristics:
  - In IS-IS multi-process, processes share the same global routing table. IS-IS multi-instance, however, uses the routing tables of VPNs, with each VPN having a separate routing table.
  - IS-IS multi-process allows a set of interfaces to be associated with a specified IS-IS process. This ensures that the protocol operations in the specified IS-IS process are confined only to this set of interfaces. In this way, multiple IS-IS processes can work on a single router, with each process responsible for a unique set of interfaces.
  - When creating an IS-IS process, you can bind it to a VPN instance. The IS-IS process then accepts and processes only the events related to the VPN instance. When the bound VPN instance is deleted, the IS-IS process is also deleted.
- Command: [Huawei] **isis** [*process-id*] [**vpn-instance** *vpn-instance-name*]
  - *process-id*: specifies the ID of an IS-IS process. If no IS-IS process is specified, IS-IS process 1 is started. The value is an integer ranging from 1 to 65535. The default value is 1.
  - **vpn-instance** *vpn-instance-name*: specifies the name of a VPN instance. If this parameter is not specified, no VPN instance is associated with the IS-IS process. The value is a string of 1 to 31 case-sensitive characters. If spaces are used, the string must start and end with double quotation marks (").

LSP Fragment

- When a PDU to be advertised by IS-IS contains too much information, an IS-IS router generates LSP fragments to carry the information.
- LSP packet format

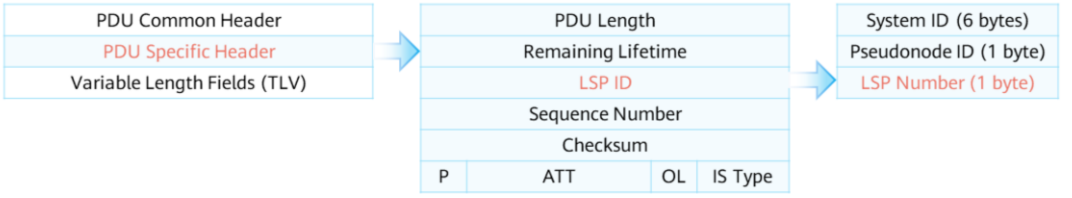

- An IS-IS LSP fragment is identified by the 1-byte LSP Number field in LSP ID. So, an IS-IS process can generate a maximum of 256 LSP fragments, which means only limited information can be carried.

Basic Concepts of LSP Fragment Extension

- You can configure a virtual system ID for IS-IS to generate virtual IS-IS LSPs to carry routing information.
  - Originating system: a router that actually runs IS-IS. A single IS-IS process can advertise LSPs like multiple virtual routers, and the originating system refers to the "actual" (not virtual) IS-IS process.
  - Normal system ID: system ID of the originating system.
  - Virtual system: system identified by an additional system ID to generate extended LSP fragments. These fragments carry additional system IDs in their LSP IDs.
  - Additional system ID: system ID of a virtual system. It is assigned by a network administrator to identify a virtual system. A maximum of 256 extended LSP fragments can be generated for each additional system ID.
  - TLV 24 (IS Alias ID TLV): carried in LSP fragments and describes the relationship between the originating and virtual systems.

Type=24
Length
Value: **Normal System-ID**

- After LSP fragment extension is configured, the system prompts you to restart the IS-IS process if information is lost because LSPs overflow. After the IS-IS process is restarted, the originating system loads as much routing information as possible to its LSPs. The information that cannot be loaded is placed in the LSPs of virtual systems for transmission. The originating system then notifies other routers of its relationship with the virtual systems through TLV 24.

HUAWEI

- The additional and normal system IDs must be unique throughout a routing domain.

Implementation of LSP Fragment Extension

- In IS-IS, each normal system ID identifies a system, which can generate a maximum of 256 LSP fragments. With additional system IDs, up to 50 virtual systems can be configured, and an IS-IS process can then generate a maximum of 13,056 LSP fragments.
- LSP fragment extension can work in two modes.

Mode-1

- Used when some routers on the network do not support LSP fragment extension.

- As the figure shown, R1 loads some routing information to the LSPs of R1-1 and R1-2 for transmission. When R2 receives the LSPs from R1, R1-1, and R1-2, it considers that there are three independent routers at the peer end and calculates routes as normal. The costs of the routes from R1 to R1-1 and from R1 to R1-2 are both 0, meaning that the routes from R2 to R1 and from R2 to R1-1/R1-2 share the same cost.

Mode-2

- Used when all the routers on the network support LSP fragment extension.

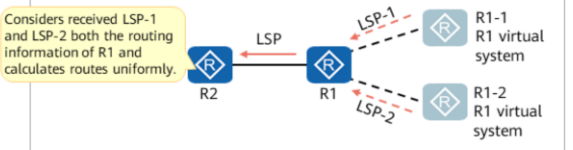

- As the figure shown, R1 loads some routing information to the LSPs of R1-1 and R1-2 for transmission. When R2 receives LSPs from R1-1 and R1-2, it knows that their originating system is R1 based on TLV 24. R2 then considers information advertised by R1-1 and R1-2 as information of R1.

- Mode-1 implementation:
  - Virtual systems participate in SPF calculation. The LSPs advertised by the originating system contain information about links to each virtual system. Similarly, the LSPs advertised by each virtual system contain information about links to the originating system. In this way, virtual systems function like physical routers that connect to the originating system.
  - Mode-1 is a transitional mode used to support earlier versions that are incapable of LSP fragment extension. In these earlier versions, IS-IS cannot identify TLV 24. As a result, the LSPs sent by a virtual system must look like LSPs sent by an originating system.
  - Precautions:
    - The LSPs sent by a virtual system must contain the same area address and overload bit as those in LSPs sent by an originating system. Other TLVs must also be the same.
    - The neighbor of a virtual system must point to an originating system, and the metric is the maximum value minus 1. The neighbor of the originating system must point to the virtual system, and the metric must be 0. This ensures that the virtual system is the downstream node of the originating system when other routers calculate routes.

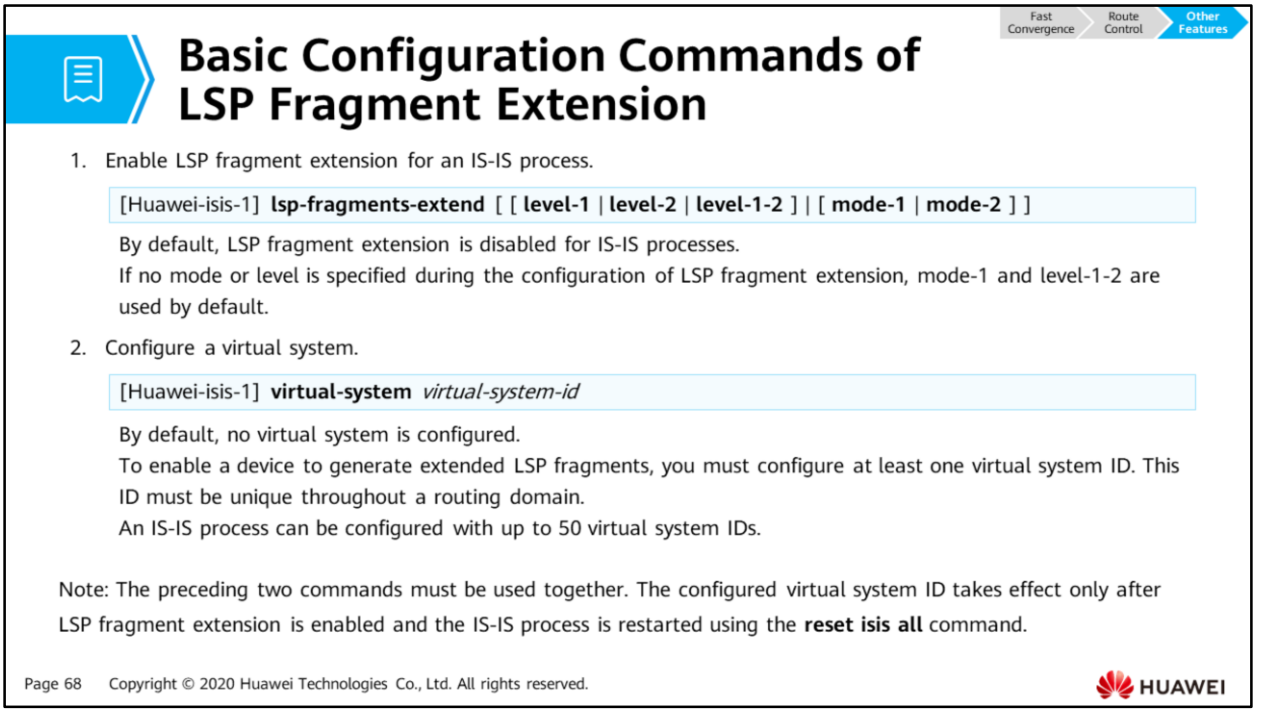

- Command: [Huawei-isis-1] **lsp-fragments-extend** [[**level-1** | **level-2** | **level-1-2**] | [**mode-1** | **mode-2**]]
  - **level-1**: enables LSP fragment extension in Level-1.
  - **level-2**: enables LSP fragment extension in Level-2.
  - **level-1-2**: enables LSP fragment extension in Level-1-2.
  - **mode-1**: allows routers to be compatible with other routers running earlier versions that are incapable of LSP fragment extension.
  - **mode-2**: requires all routers to support LSP fragment extension.
- Command: [Huawei-isis-1] **virtual-system** *virtual-system-id*
  - *virtual-system-id*: specifies a virtual system ID of an IS-IS process. The length is 6 bytes (48 bits), and the format is XXXX.XXXX.XXXX.

Quiz

1. (Multiple) Which of the following fast convergence mechanisms are supported by OSPF? ( )
   A. I-SPF
   B. LSP fast flooding
   C. Intelligent timer
   D. OSPF IP FRR
2. (TorF) The FA field in Type 5 LSAs of OSPF must be 0.0.0.0. ( )
   A. True
   B. False
3. (TorF) On an IS-IS network, if a device runs mode-2 LSP fragment extension, virtual systems do not participate in SPF calculation. All routers on the network know that the LSPs generated by the virtual systems actually belong to the originating system. ( )
   A. True
   B. False

1. CD
2. B
3. A

Summary

- To better adapt to network topology changes, OSPF and IS-IS support multiple fast convergence modes. I-SPF and PRC algorithm speed up route calculation; FRR enables fast traffic switching to the backup link; and intelligent timers allow you to control the speeds at which link state information is generated and routes are calculated.
- To control the size of routing tables and improve network performance, OSPF and IS-IS support route filtering, equal-cost routes, and default routes.
- To isolate protocol routing tables, OSPF and IS-IS support multi-process deployment on the same device. Protocol routing tables of different processes do not affect each other.
- To prevent a sub-optimal route from being selected when external routes are imported, OSPF uses the FA in Type 5 or Type 7 LSAs to guide data packet forwarding.
- To carry more routing information, IS-IS provides LSP fragment extension.
- Thanks to the preceding features, OSPF and IS-IS are widely and flexibly used on live networks.


Thank You
www.huawei.com