

Contents

Configuring IP forwarding basic settings	1
About FIB table	1
Saving the IP forwarding entries to a file.....	2
Enabling split horizon forwarding	2
Display and maintenance commands for FIB table	2
Configuring load sharing	3
About load sharing	3
Configuring load sharing mode	3
Enabling IPv4 bandwidth-based load sharing.....	3
Enabling local-first load sharing	4
Enabling symmetric load sharing	4
Display and maintenance commands for load sharing	5
Loading sharing configuration examples	5
Example: Configuring load sharing based on source and destination addresses.....	5

Configuring IP forwarding basic settings

About FIB table

A device uses the FIB table to make packet forwarding decisions.

A device selects optimal routes from the routing table, and puts them into the FIB table. Each FIB entry specifies the next hop IP address and output interface for packets destined for a specific subnet or host.

For more information about the routing table, see *Layer 3—IP Routing Configuration Guide*.

Use the `display fib` command to display the FIB table. The following example displays the entire FIB table.

```
<Sysname> display fib
```

```
Destination count: 8 FIB entry count: 8
```

```
Flag:
```

```
U:Usable   G:Gateway   H:Host     B:Blackhole D:Dynamic   S:Static  
R:Relay    F:FRR
```

Destination/Mask	Nexthop	Flag	OutInterface/Token	Label
0.0.0.0/32	127.0.0.1	UH	InLoop0	Null
127.0.0.0/8	127.0.0.1	U	InLoop0	Null
127.0.0.0/32	127.0.0.1	UH	InLoop0	Null
127.0.0.1/32	127.0.0.1	UH	InLoop0	Null
127.255.255.255/32	127.0.0.1	UH	InLoop0	Null
224.0.0.0/4	0.0.0.0	UB	NULL0	Null
224.0.0.0/24	0.0.0.0	UB	NULL0	Null
255.255.255.255/32	127.0.0.1	UH	InLoop0	Null

A FIB entry includes the following items:

- **Destination**—Destination IP address.
- **Mask**—Network mask. The mask and the destination address identify the destination network. A logical AND operation between the destination address and the network mask yields the address of the destination network. For example, if the destination address is 192.168.1.40 and the mask 255.255.255.0, the address of the destination network is 192.168.1.0. A network mask includes a certain number of consecutive 1s. It can be expressed in dotted decimal format or by the number of the 1s.
- **Nexthop**—IP address of the next hop.
- **Flag**—Route flag.
- **OutInterface**—Output interface.
- **Token**—MPLS Label Switched Path index number.
- **Label**—Inner label.

Saving the IP forwarding entries to a file

Restrictions and guidelines

The feature automatically creates the file if you specify a nonexistent file. If the file already exists, this feature overwrites the file content.

This feature triggers one-time saving of the IP forwarding entries.

To automatically save the IP forwarding entries periodically, configure a schedule for the device to automatically run the `ip forwarding-table save` command. For information about scheduling a task, see *Fundamentals Configuration Guide*.

Procedure

To save the IP forwarding entries to a file, execute the following command in any view:

```
ip forwarding-table save filename filename
```

Enabling split horizon forwarding

About split horizon forwarding

This feature prevents IPv4, IPv6, and MPLS packets from being forwarded out of the physical interface on which they were received, avoiding network loops.

Procedure

1. Enter system view.
`system-view`
2. Enable split horizon forwarding.
`forwarding split-horizon`
By default, split horizon forwarding is disabled.

Display and maintenance commands for FIB table

Execute `display` commands in any view.

Task	Command
Display FIB entries.	<code>display fib [vpn-instance <i>vpn-instance-name</i>] [<i>ip-address</i> [<i>mask</i> <i>mask-length</i>]]</code>

Configuring load sharing

About load sharing

If a routing protocol finds multiple equal-cost best routes to the same destination, the device forwards packets over the equal-cost routes to implement load sharing.

Configuring load sharing mode

About load sharing mode

In the per-flow load sharing mode, the device forwards flows over equal-cost routes. Packets of one flow travel along the same routes. You can configure the device to identify a flow based on the following criteria: source IP address, destination IP address, source port number, destination port number, IP protocol number, and ingress port.

In a complex network, when the criteria cannot distinguish flows, you can use the **algorithm** keyword to specify an algorithm to identify flows.

Procedure

1. Enter system view.

```
system-view
```

2. Configure load sharing.

```
ip load-sharing mode per-flow [ algorithm algorithm-number [ seed seed-number ] [ shift shift-number ] ] | [ dest-ip | dest-port | ingress-port | ip-pro | src-ip | src-port ] * | tunnel { inner | outer } ] { global | slot slot-number }
```

By default, the device performs per-flow load sharing based on the following criteria: source IP address, destination IP address, source port number, destination port number, IP protocol number, and ingress port.

3. Display the load sharing path selected for a flow.

```
display ip load-sharing path ingress-port interface-type interface-number packet-format { ipv4oe dest-ip ip-address [ src-ip ip-address ] | ipv6oe dest-ipv6 ipv6-address [ src-ipv6 ipv6-address ] } [ dest-port port-id | ip-pro protocol-id | src-port port-id ] *
```

The option settings in this command must match both the options displayed in the **display ip load-sharing mode** command and the field values in load shared packets. If the option settings do not meet the requirement, the path displayed by this command might be different from the real path for load sharing.

Enabling IPv4 bandwidth-based load sharing

About IPv4 bandwidth-based load sharing

This feature load shares flow traffic among multiple output interfaces based on their load percentages. The device calculates the load percentage for each output interface in terms of the interface expected bandwidth.

Devices that run load sharing protocols, such as Locator/ID Separation Protocol (LISP), implement load sharing based on the ratios defined by these protocols.

Restrictions and guidelines

This feature is mutually exclusive with enhanced ECMP mode, which is configured by using the `ecmp mode enhanced` command.

After you configure this feature, make sure the expected bandwidth of the interface does not exceed the physical bandwidth of the interface.

Procedure

1. Enter system view.
`system-view`
2. Enable IPv4 bandwidth-based load sharing.
`bandwidth-based-sharing`
By default, the IPv4 bandwidth-based load sharing is disabled.
3. (Optional.) Configure the expected bandwidth of the interface.
 - a. Enter interface view.
 - b. `interface interface-type interface-number`
 - c. Configure the expected bandwidth of the interface.
 - d. `bandwidth bandwidth`
 - e. By default, the expected bandwidth is the physical bandwidth of the interface.

Enabling local-first load sharing

About local-first load sharing

Local-first load sharing distributes traffic preferentially across the output interfaces on the receiving IRF member device if output interfaces for multiple equal-cost routes are on different members. This feature enhances packets forwarding efficiency.

Procedure

1. Enter system view.
`system-view`
2. Enable local-first load sharing.
`ip load-sharing local-first enable`
By default, local-first load sharing is enabled.

Enabling symmetric load sharing

About symmetric load sharing

Symmetric load sharing ensures that bidirectional traffic specific to a source and destination address pair flow along the same path.

Restrictions and guidelines

Before you enable this feature, first execute the `ip load-sharing mode` command to configure five tuple-based load sharing. The device will perform load sharing based on the source IP address, destination IP address, source port number, destination port number, and IP protocol number. If not doing so, the device will use the default load sharing mode. In this mode, the packet ingress port is included in load sharing calculation and symmetric load sharing does not take effect.

Procedure

1. Enter system view.

system-view

2. Enable symmetric load sharing.

ip load-sharing symmetric enable

By default, symmetric load sharing is disabled.

Display and maintenance commands for load sharing

Execute **display** commands in any view.

Task	Command
Display the load sharing mode in use.	display ip load-sharing mode slot slot-number
Display the load sharing path selected for a flow.	display ip load-sharing path ingress-port interface-type interface-number packet-format { ipv4oe dest-ip ip-address [src-ip ip-address] ipv6oe dest-ipv6 ipv6-address [src-ipv6 ipv6-address] } [dest-port port-id ip-pro protocol-id src-port port-id] *

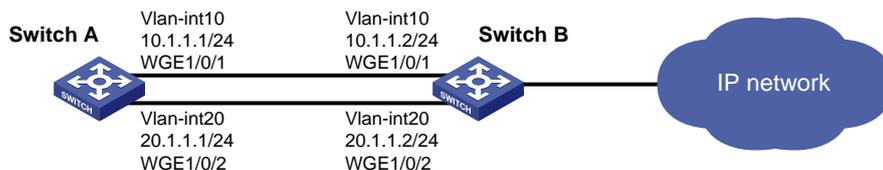
Loading sharing configuration examples

Example: Configuring load sharing based on source and destination addresses

Network configuration

As shown in [Figure 1](#), Switch A has two equal-cost routes to Switch B. Configure load sharing on Switch A to forward packets through Switch B to the destination IP address 1.2.3.4/24.

Figure 1 Network diagram



Procedure

On Switch A, assign Twenty-FiveGigE 1/0/1 to VLAN 10, and Twenty-FiveGigE 1/0/2 to VLAN 20.

```
<SwitchA> system-view
[SwitchA] vlan 10
[SwitchA-vlan10] port twenty-fivegige 1/0/1
[SwitchA-vlan10] quit
[SwitchA] vlan 20
[SwitchA-vlan20] port twenty-fivegige 1/0/2
[SwitchA-vlan20] quit
```

On Switch A, configure IP addresses for VLAN-interface 10 and VLAN-interface 20.

```
[SwitchA] interface vlan-interface 10
[SwitchA-Vlan-interface10] ip address 10.1.1.1 24
[SwitchA-Vlan-interface10] quit
[SwitchA] interface vlan-interface 20
[SwitchA-Vlan-interface20] ip address 20.1.1.1 24
[SwitchA-Vlan-interface20] quit
```

On Switch B, assign Twenty-FiveGigE 1/0/1 to VLAN 10, and Twenty-FiveGigE 1/0/2 to VLAN 20.

```
<SwitchB> system-view
[SwitchB] vlan 10
[SwitchB-vlan10] port twenty-fivegige 1/0/1
[SwitchB-vlan10] quit
[SwitchB] vlan 20
[SwitchB-vlan20] port twenty-fivegige 1/0/2
[SwitchB-vlan20] quit
```

On Switch B, configure IP addresses for VLAN-interface 10 and VLAN-interface 20.

```
[SwitchB] interface vlan-interface 10
[SwitchB-Vlan-interface10] ip address 10.1.1.2 24
[SwitchB-Vlan-interface10] quit
[SwitchB] interface vlan-interface 20
[SwitchB-Vlan-interface20] ip address 20.1.1.2 24
[SwitchB-Vlan-interface20] quit
```

On Switch A, configure two static routes to the destination IP address.

```
<SwitchA> system-view
[SwitchA] ip route-static 1.2.3.4 24 10.1.1.2
[SwitchA] ip route-static 1.2.3.4 24 20.1.1.2
[SwitchA] quit
```

On Switch A, display FIB entries matching the destination IP address 1.2.3.4.

```
<SwitchA> display fib 1.2.3.4
Destination count: 1 FIB entry count: 2
Flag:
  U:Usable  G:Gateway  H:Host  B:Blackhole  D:Dynamic  S:Static
  R:Relay   F:FRR
Destination/Mask  Nexthop          Flag      OutInterface/Token  Label
1.2.3.0/24       10.1.1.2         USGR      Vlan10              Null
1.2.3.0/24       20.1.1.2         USGR      Vlan20              Null
```

On Switch A, configure per-flow load sharing based on the source IP address and destination IP address.

```
<SwitchA> system-view
[SwitchA] ip load-sharing mode per-flow dest-ip src-ip global
```

Verifying the configuration

Verify that Switch A implements load sharing.

```
<SwitchA> display counters outbound interface GigabitEthernet
Interface          Total (pkts)  Broadcast (pkts)  Multicast (pkts)  Err (pkts)
WGE1/0/1           1045          0                  0                  0
WGE1/0/2           1044          0                  0                  0
```