Contents

1 Configuring IP REF

1.1 Introduction

1.1.1 Overview

On products that do not support hardware-based forwarding, IPv4/IPv6 packets are forwarded by software. To optimize the forwarding performance of software, the device introduces the IPv4/IPv6.

Orion maintains two tables: forwarding table and adjacency table. The forwarding table is used to store routing information. The adjacency table, equivalent to the Address Resolution Protocol (ARP) table and IPv6 neighbor table, is used to store link layer information of next hops.

1.1.2 Principles

1. Basic Concepts

- Forwarding table

  A forwarding table, also known as forwarding information base (FIB), is the lookup table where the device determines target devices for data exchange, and there are one-to-one mappings between entries in the FIB and entries in the IP routing table. Since the FIB contains all the necessary routing information, there is no need to maintain the route cache. When the network topology or route changes, the IP routing table is updated, and entries in the FIB change accordingly. During packet forwarding, IP REF selects packet transmission paths according to the forwarding table.

- Adjacency table

  An adjacency table contains output interface information of routed packets, for example, the next hop list, the next component used for packet processing, or the link layer output encapsulation. When a packet matches a certain adjacent node, the packet is directly encapsulated and then the send function of the node is called to forward the packet. For the sake of search and update, an adjacency table is often organized in the form of a hash table. To support routing load balancing, the next hop list of an adjacent node is organized in the form of a load balancing table. The next hop list of an adjacent node may not contain next hop information but may contain indexes of next components used for packet processing (such as other line cards and multi-service cards).

- Active resolution

  REF supports next hop resolution. If the MAC address of the next hop of an Ethernet interface is unknown, REF will actively resolve the next hop.

  IPv4 REF requests the ARP module to resolve the next hop. IPv6 REF requests the neighbor discovery (ND) module to resolve the next hops.

  .

- Packet forwarding path

  Packets are routed and forwarded based on their IPv4/IPv6 addresses. If the source and destination IPv4/IPv6 addresses of a packet are specified, the forwarding path of this packet is definite.

2. REF Load Balancing

REF Load balancing is used to distribute traffic load among channels of multiple network devices. In the REF model, a route is a multi-path route when the route prefix is associated with multiple next hops. In this case, the route will be associated with a load balancing table and load balancing will be implemented based on the route weights. When an IP packet matches the load balancing table based on the longest prefix match, REF performs hash calculation based on the IP address of the packet and selects a path to forward the packet.

3. ECMP Load Balancing

Equal cost multiple path (ECMP) load balancing works in such a way that the same packet flow is transmitted by one path and different packet flows are transmitted by multiple equal-cost paths according to the load balancing algorithms, to achieve load balancing.

ECMP differentiates packet flows by using one or several load balancing algorithms based on packet characteristics, such as the source IP address, destination IP address, L4 source port number, and L4 destination port number. It also distributes the same packet flow to the same path for transmission, and different packet flows to equal-cost paths evenly.

For example, in source IP address-based load balancing mode, packets are distributed to different paths based on the source IP addresses of the packets. Packets with different source IP addresses are evenly distributed to equal-cost paths, and packets with the same source IP address are forwarded by the same path.

4. Hash Load Balancing Control

Hash load balancing allows users to control load balancing flexibly in different scenarios. Hash load balancing functions are described as below:

- Elastic hash

  In normal cases, a load balancing cluster interworks with multiple top of rack (ToR) devices via the ECMP, and the ToR devices are required to distribute all packets of a session to the same scheduling server. However, when a server is faulty, all session streams will be reorganized and balanced according to the traditional load balancing mechanism of ToR devices. In this case, packets of the same session may be distributed to two servers. The elastic hash function ensures that, when one link is faulty, traffic on other links is not affected and only the traffic of the faulty link is distributed to other active links, without affecting the traffic of other links. In this way, traffic of a session will only be distributed to the same server.

- Hash algorithm mode setting

  When traffic models are different, an appropriate hash algorithm mode can be set to adapt to different traffic to achieve the optimal load balancing effect.

- Hash disturbance

  Device traffic is balanced using the hash algorithm. For the same stream from two devices of the same type, the same path will be calculated by the load balancing algorithm. When the ECMP is deployed, the same stream from the two devices may be distributed to the same destination device. This is called hash polarization. The hash disturbance factor is used to affect the load balancing algorithm of the device. Different disturbance factors can be configured on different devices so that different paths are calculated for the same stream.

- Hash synchronization

To ensure network security, a firewall cluster is deployed between the internal and external networks for traffic cleaning. This requires that both the uplink and downlink traffic of a session be transmitted to the same device in the firewall cluster for processing. The source and destination IP addresses contained in the uplink and downlink streams of a session are reversed. If a traditional hash algorithm is used, the uplink and downlink streams will be directed to different firewalls, while hash synchronization can ensure that the same path is calculated for the uplink and downlink streams.

- Hash factor acquisition

  For packets of the virtual extensible LAN (VXLAN) and Generic Routing Encapsulation (GRE) tunnels, there are inner and outer headers. The hash factor can be obtained from the inner or outer header to achieve better load balancing. For example, when tunnel packets have the same IP address in the outer header but different IP addresses in the inner header, the IP address in the inner header can be specified as a hash factor in order to achieve better load balancing.

1.1.3 Protocols and Standards

RFC 1812: Requirements for IP Version 4 Routers

1.2 Configuration Task Summary

All the configuration tasks below are optional. Perform the configuration tasks as required.

- Configuring REF Load Balancing
- Configuring ECMP to Conduct Load Balancing According to a Fixed Profile
- Configuring ECMP to Conduct Load Balancing According to an Enhanced Profile

1.3 Configuring REF Load Balancing

1.3.1 Overview

REF Load balancing is used to distribute traffic load among channels of multiple network devices. REF supports load balancing of packets. In the REF model, a route is a multi-path route when the route prefix is associated with multiple next hops. In this case, the route will be associated with a load balancing table and load balancing will be implemented based on the route weights. When an IP packet matches the load balancing table based on the longest prefix match, REF performs hash calculation based on the IP address of the packet and selects a path to forward the packet.

IP REF supports three load balancing modes, that is, three routing policies:

- The load is balanced based on the destination IP addresses of IP packets and hash calculation is performed on the destination IP addresses of the packets. A path with a greater weight is more likely to be selected. This policy is used by default.
- The load is balanced based on the source and destination IP addresses of IP packets and hash calculation is performed on the source and destination IP addresses of the packets. A path with a greater weight is more likely to be selected.
- The load is balanced based on the source IP addresses of IP packets and hash calculation is performed on the source IP addresses of the packets. A path with a greater weight is more likely to be selected.

1.3.2 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure an IP REF load balancing algorithm.

(IPv4 network)

ip ref load-sharing { **original | original-only** }

The default load balancing algorithm of IPv4 REF is load balancing based on the destination IP addresses.

(IPv6 network)

ipv6 ref load-sharing { **original | original-only** }

The default load balancing algorithm of IPv6 REF is load balancing based on the destination IPv6 addresses.

1.4 Configuring ECMP to Conduct Load Balancing According to a Fixed Profile

1.4.1 Overview

ECMP differentiates packet flows by using one or several load balancing algorithms based on packet characteristics, such as the source IP address, destination IP address, L4 source port number, and L4 destination port number. It also distributes the same packet flow to the same path for transmission, and different packet flows to equal-cost paths evenly.

1.4.2 Restrictions and Guidelines

- The command for configuring a load balancing policy in basic load balancing mode is not different to IPv4 and IPv6.
- When the **enhanced profile** *profile-name* parameter is configured in the **ip ref load-balance** command, load balancing can be performed according to the enhanced profile indicated by *profile-name*. For the enhanced profile configurations, see Configuring ECMP to Conduct Load Balancing According to an Enhanced Profile.
- When running the **ip ref load-balance** command, you can configure the following parameters to select the ECMP load balancing mode using a fixed profile.
  - **enhanced profile** *profile-name*: Performs load balancing based on the field configured in the enhanced profile.
  -
  - **src-dst-ip**: Performs load balancing based on the source and destination IP addresses of packets.
  - **src-dst-ip-src-dst-l4port**: Performs load balancing based on the source IP addresses, destination IP addresses, L4 source port numbers, and L4 destination port numbers of packets.
  - **src-ip**: Performs load balancing based on the source IP addresses of packets.

1.4.3 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure an ECMP load balancing policy.

ip ref load-balance { **enhanced profile** *profile-name* | **src-dst-ip** | **src-dst-ip-src-dst-l4port** | **src-ip** }The load is balanced according to an enhanced profile by default.

(4) (Optional) Enable the elastic hash of ECMP load balancing.

ip ref hash-elasticity enable

Elastic hash of ECMP load balancing is disabled by default.

1.5 Configuring ECMP to Conduct Load Balancing According to an Enhanced Profile

1.5.1 Overview

The enhanced load balancing mode allows users to define load balancing profiles and combine multiple fields of different types of packets for load balancing.

In enhanced load balancing mode, a device first determines the type of packets to be transmitted and performs load balancing based on the specified packet fields. For example, if IPv4 packets containing changing source IP addresses need to be output through an aggregate port, configure the device to perform load balancing based on the specified source IP address field in IPv4 packets.

Configure this function on aggregation or core switches that support enhanced load balancing, and specify parameters to achieve load balancing.

1.5.2 Restrictions and Guidelines

- **dst-ip**: Performs load balancing based on the destination IP addresses of incoming packets.
- **l4-dst-port**: Performs load balancing based on the L4 destination port numbers of incoming packets.
- **l4-src-port**: Performs load balancing based on the L4 source port numbers of incoming packets.
- **protocol**: Performs load balancing based on the protocol types of incoming packets.
- **src-ip**: Performs load balancing based on the source IP addresses of incoming packets.
- **src-port**: Performs load balancing based on the source port numbers of incoming packets.

1.5.3 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Rename the enhanced profile and enter the enhanced profile configuration mode of REF load balancing.

ip-ref-load-balance-profile *profile-name*

The default name of the enhanced profile is **default**.

(4) Configure the enhanced profile for load balancing. Perform at least one of the following configuration steps:

- Configure the load balancing mode for IP packets.

  (IPv4 network)

  ipv4 field [**dst-ip**] [**l4-dst-port**] [**l4-src-port**] [**protocol**] [**src-ip**] [**src-port**]

  The load is balanced based on the source IP addresses, destination IP addresses, L4 source port numbers, and L4 destination port numbers of packets by default.

  (IPv6 network)

  ipv6 field [**dst-ip**] [**l4-dst-port**] [**l4-src-port**] [**protocol**] [**src-ip**] [**src-port**]

  The load is balanced based on the source IP addresses, destination IP addresses, L4 source port numbers, and L4 destination port numbers of packets by default.

- Configure the hash disturbance factor.

  hash-disturb *string*

  The hash disturbance function of load balancing is disabled by default.

- Enable hash synchronization.

  hash-symmetrical { **ipv4** | **ipv6** }

  Hash synchronization is disabled by default.

1.6 Monitoring

Run the **show** command to check the running status of a configured function to verify the configuration effect.

Run the **clear** command to clear information.

Caution

Running the **clear** command may lose vital information and thus interrupt services.

Table 1-1Monitoring

Command	Purpose
show ip ref packet statistics	Displays current packet statistics of IPv4 REF.
clear ip ref packet statistics	Clears current packet statistics of IPv4 REF.
show ip ref adjacency [**glean** | **local** | *ip-address* | **interface** *interface-type interface-number* | **discard** | **statistics**]	Displays information about the specified glean adjacencies, local adjacencies, adjacency of a specified IP address, adjacencies associated with a specified interface, and all adjacent nodes of IPv4

Command	Purpose
	REF.
show ip ref resolve-list	Displays the next hop resolved by IPv4 REF.
show ip ref exact-route [**vrf** *vrf-name*] *source-ip-address destination-ip-address*	Displays the actual forwarding path of a specified packet.
show ip ref route [**vrf** *vrf-name*] [**default** | *ipv4-address mask* | **statistics**]	Displays routing information in the IPv4 REF table. The parameter **default** indicates that information about the default route will be displayed.
show ip ref route detail *ipv4-address* [**vrf** *vrf-name*]	Displays details about the current IPv4 route.
show ip ref load-balance	Displays the configuration of the ECMP load balancing using a fixed profile.
show ip ref-load-balance-profile [*profile-name*]	Displays the configuration of the ECMP load balancing using an enhanced profile.
ip ref synchronize all	Synchronizes IPv4 routes between the hardware and software REF.
show ipv6 ref packet statistics	Displays current packet statistics of IPv6 REF.
clear ipv6 ref packet statistics	Clears current packet statistics of IPv6 REF.
show ipv6 ref adjacency [**glean** | **local** | *ipv6-address* | **interface** *interface-type interface-number* | **discard** | **statistics**]	Displays information about the specified glean adjacencies, local adjacencies, adjacency of a specified IPv6 address, adjacencies associated with a specified interface, and all adjacent nodes of IPv6 REF.
show ipv6 ref resolve-list	Displays the next hop resolved by IPv6 REF.
show ipv6 ref exact-route [**vrf** *vrf-name*] *source-ipv6-address destination-ipv6-address*	Displays actual the forwarding path of a specified IPv6 packet.
show ipv6 ref route [**vrf** *vrf-name*] [**default** | **statistics** | *ipv6-address/prefix-length*]	Displays route information in the IPv6 REF table. The parameter **default** indicates that information about the default route will be displayed.

Command	Purpose
show ipv6 ref route detail *ipv6-address* [**vrf** *vrf-name*]	Displays details about the current IPv6 route.
ipv6 ref synchronize all	Synchronizes IPv6 routes between the hardware and software REF.

1.7 Configuration Examples

1.7.1 Configuring Load Balancing

1. Requirements

As shown in Figure 1-1, device A and device E run REF, and device B, device C and device D are forwarding devices.

Device A communicates with device E. Link 1, link 2, and link 3 are three equal-cost links. A specific load balancing algorithm is configured on device A so that traffic is balanced by these three equal-cost links.

2. Topology

Figure 1-1Topology of Configuring an ECMP Load Balancing Policy

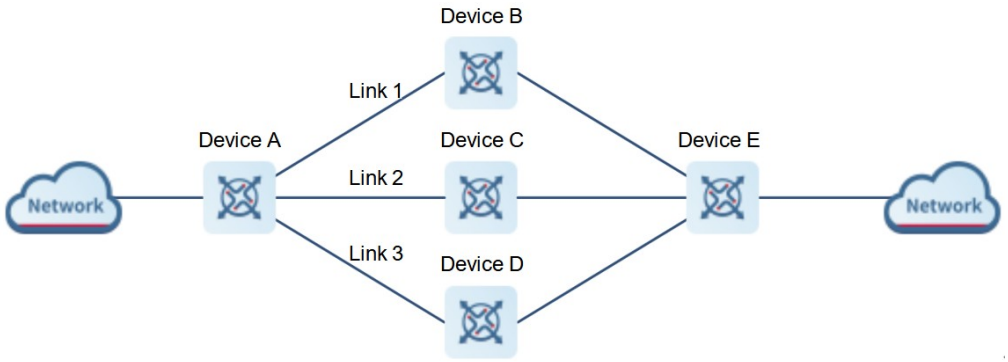

3. Notes

- Configure a policy of load balancing using a fixed profile on device A. Configure the load balancing mode, enable the elastic hash of ECMP load balancing, and configure the ECMP hash algorithm mode.
- Configure a policy of load balancing using an enhanced profile on device A: Rename the ECMP load balancing profile, configure the load balancing mode for IPv4 packets, configure the load balancing mode for IPv6 packets, configure the hash disturbance factor of load balancing, and disable the load balancing hash disturbance function.

4. Procedure

Configure a policy of load balancing using a fixed profile on device A.

```
DeviceA> enable
DeviceA# configure terminal
```

```
DeviceA(config)# ip ref load-balance src-ip
DeviceA(config)# ip ref hash-elasticity enable
```

Configure a policy of load balancing using an enhanced profile on device A.

```
DeviceA> enable
DeviceA# configure terminal
DeviceA(config)# ip-ref-load-balance-profile rpn
DeviceA(ref-ip-config-load-balance-profile)# ipv4 field src-ip dst-ip
DeviceA(ref-ip-config-load-balance-profile)# ipv6 field src-ip dst-ip protocol
DeviceA(ref-ip-config-load-balance-profile)# hash-disturb A
DeviceA(ref-ip-config-load-balance-profile)# no hash-symmetrical ipv6
```

5. Verification

Run the **show ip ref load-balance** command to display the configuration of load balancing using a fixed profile.

```
DeviceA# show ip ref load-balance
load-balance            : src-ip.
  hash-elasticity        : enable.
```

Run the **show ip ref load-balance-profile** command to display the configuration of load balancing using an enhanced profile.

```
DeviceA# show ip ref load-balance-profile
Load-balance-profile: rpn
Packet    Hash Field:
    IPV4: src-ip dst-ip
   IPV6: protocol src-ip dst-ip
```

6. Configuration Files

Device A configuration file

```
hostname DeviceA
ip-ref-load-balance-profile rpn
 ipv4 field src-ip dst-ip
 ipv6 field protocol src-ip dst-ip
 hash-disturb A
!
ip ref load-balance src-ip
ip ref hash-elasticity enable
```