Contents

Configuring IPv6 PBR

About IPv6 PBR

IPv6 policy-based routing (PBR) uses user-defined policies to route IPv6 packets. A policy can specify parameters for packets that match specific criteria such as ACLs. The parameters include the next hop and default next hop.

IPv6 packet forwarding process

A device forwards received IPv6 packets using the following process:

1. The device uses PBR to forward matching packets.
2. If one of the following events occurs, the device searches for a route (except the default route) in the routing table to forward packets:
   - The packets do not match the PBR policy.
   - The PBR-based forwarding fails.
3. If the forwarding fails, the device uses the default next hop defined in PBR to forward packets.
4. If the forwarding fails, the device uses the default route to forward packets.

IPv6 PBR types

IPv6 PBR includes the following types:

- **Local PBR**—Guides the forwarding of locally generated packets, such as the ICMP packets generated by using the **`ping`** command.
- **Interface PBR**—Guides the forwarding of packets received on an interface only.

Policy

An IPv6 policy includes match criteria and actions to be taken on the matching packets. A policy can have one or multiple nodes as follows:

- Each node is identified by a node number. A smaller node number has a higher priority.
- A node contains **`if-match`** and **`apply`** clauses. An **`if-match`** clause specifies a match criterion, and an **`apply`** clause specifies an action.
- A node has a match mode of **permit** or **deny**.

An IPv6 policy compares packets with nodes in priority order. If a packet matches the criteria on a node, it is processed by the action on the node. If the packet does not match any criteria on the node, it goes to the next node for a match. If the packet does not match the criteria on any node, the device performs a routing table lookup for the packet.

Relationship between if-match clauses

IPv6 PBR supports only the **`if-match acl`** clause to set an ACL match criterion. On a node, you can specify only one **`if-match`** clause.

Relationship between apply clauses

You can specify multiple **`apply`** clauses for a node, but some of them might not be executed. For more information about the relationship between the **`apply`** clauses, see "Configuring actions for an IPv6 node."

Relationship between the match mode and clauses on the node

Does a packet match all the if-match clauses on the node?	Match mode	
	In permit mode	**In deny mode**
Yes	• If the node contains **apply** clauses, IPv6 PBR executes the **apply** clauses on the node. If IPv6 PBR-based forwarding succeeds, IPv6 PBR does not compare the packet with the next node. • If the node does not contain **apply** clauses, the device performs a routing table lookup for the packet.	The device performs a routing table lookup for the packet.
No	IPv6 PBR compares the packet with the next node.	IPv6 PBR compares the packet with the next node.

NOTE:

A node that has no **if-match** clauses matches any packet.

IPv6 PBR and Track

IPv6 PBR can work with the Track feature to dynamically adapt the availability status of an **apply** clause to the link status of a tracked object. The tracked object can be a next hop or default next hop.

- When the track entry associated with an object changes to **Negative**, the **apply** clause is invalid.
- When the track entry changes to **Positive** or **NotReady**, the **apply** clause is valid.

For more information about Track and IPv6 PBR collaboration, see *High Availability Configuration Guide*.

Restrictions and guidelines: IPv6 PBR configuration

If a packet destined for the local device matches an IPv6 PBR policy, IPv6 PBR will execute the apply clauses in the policy, including the clause for forwarding. When you configure an IPv6 PBR policy, be careful to avoid this situation.

IPv6 PBR tasks at a glance

To configure IPv6 PBR, perform the following tasks:

1. Configuring an IPv6 policy
   a. Creating an IPv6 node
   b. Setting match criteria for an IPv6 node
   c. Configuring actions for an IPv6 node
2. Specifying a policy for IPv6 PBR

   Choose the following tasks as needed:

Configuring an IPv6 policy

Creating an IPv6 node

1. Enter system view.
   system-view
2. Create an IPv6 policy or policy node and enter its view.
   ipv6 policy-based-route *policy-name* [**deny** | **permit**] **node** *node-number*
3. (Optional.) Configure a description for the IPv6 policy node.
   description *text*
   By default, no description is configured for an IPv6 policy node.

Setting match criteria for an IPv6 node

1. Enter system view.
   system-view
2. Enter IPv6 policy node view.
   ipv6 policy-based-route *policy-name* [**deny** | **permit**] **node** *node-number*
3. Set match criteria.
   - Set an ACL match criterion.
     if-match acl { *ipv6-acl-number* | **name** *ipv6-acl-name* }
     By default, no ACL match criterion is set.
     The ACL match criterion cannot match Layer 2 information.
     When using the ACL to match packets, IPv6 PBR ignores the action (permit or deny) and time range settings in the ACL.
   - Set a local QoS ID match criterion.
     if-match qos-local-id *local-id-value* [**qppb-manipulation**]
     By default, no local QoS ID match criterion is set.
     If you specify the **qppb-manipulation** keyword, the command applies only to QPPB in which the device acts as a BGP receiver. The device determines that a received packet matches the policy node if the local QoS ID obtained from its matching route entry is identifical to the specified local QoS ID criterion. For more information about QPPB, see *ACL and QoS Configuration Guide*.
   - Set a service chain match criterion.
     if-match service-chain { **path-id** *service-path-id* [**path-index** *service-path-index*] }
     By default, no service chain match criterion is set.
     The S6820 switch series does not support this command.

Configuring actions for an IPv6 node

About apply clauses

The **apply** clauses allow you to specify actions to take on matching packets on a node.

The following **apply** clauses determine the packet forwarding paths in a descending order:

- **apply next-hop**
- **apply output-interface NULL0**
- **apply default-next-hop**

IPv6 PBR supports the **apply** clauses in Table 1.

Table 1 Apply clauses supported in IPv6 PBR

Clause	Meaning	Remarks
apply precedence	Sets an IP precedence.	This clause is always executed.
apply next-hop	Sets next hops.	N/A
apply output-interface NULL0	Sets the output interface to Null 0.	If both the **apply next-hop** and **apply output-interface NULL0** clauses are configured, only the **apply next-hop** clause is executed.
apply service-chain	Sets service chain information.	For this clause to take effect, you must also use the **apply next-hop** clause to specify a reachable next hop.
apply default-next-hop	Sets default next hops.	The clauses take effect only in the following cases: • No next hops are set or the next hops are invalid. • The IPv6 packet does not match any route in the routing table.
apply fail-action-drop next-hop	Sets the action that drops matching packets when all next hops become invalid.	If no next hops are specified on the IPv6 policy node, this clause drops all packets that match the IPv6 policy node. This clause is typically used in scenarios that require strict routing paths.

Restrictions and guidelines for action configuration

If you specify a next hop, IPv6 PBR periodically performs a lookup in the FIB table to determine its availability. Temporary service interruption might occur if IPv6 PBR does not update the route immediately after its availability status changes.

Setting an IP preference

1. Enter system view.
   system-view
2. Enter IPv6 policy node view.
   ipv6 policy-based-route *policy-name* [**deny** | **permit**] **node** *node-number*
3. Set an IP precedence.
   apply precedence { *type* | *value* }
   By default, no IP precedence is specified.

Configuring actions for a node

1. Enter system view.
   system-view
2. Enter IPv6 policy node view.

ipv6 policy-based-route *policy-name* [**deny** | **permit**] **node** *node-number*

3. Configure actions for a node.
   - Set next hops for permitted IPv6 packets.

     apply next-hop [**vpn-instance** *vpn-instance-name*] { *ipv6-address* [**direct**] [**track** *track-entry-number*] } &<1-2>

     By default, no next hops are specified.

     You can specify multiple next hops for backup in one command line or by executing this command multiple times. You can specify a maximum of two next hops for a node.
   - Set the output interface to Null 0.

     apply output-interface NULL 0

     By default, no action is configured to set the output interface to Null 0.
   - Set default next hops.

     apply default-next-hop [**vpn-instance** *vpn-instance-name*] { *ipv6-address* [**direct**] [**track** *track-entry-number*] }&<1-2>

     By default, no default next hops are specified.

     You can specify multiple default next hops for backup in one command line or by executing this command multiple times. You can specify a maximum of two default next hops for a node.
   - Set service chain information.

     apply service-chain path-id *service-path-id* [**path-index** *service-path-index*]

     By default, no service chain information is set.

     The S6820 switch series does not support this command.
   - Set the action that drops matching packets when all next hops specified on the IPv6 policy node are invalid.

     apply fail-action-drop next-hop

     By default, the drop action is not configured. The matching packets are forwarded based on the typical packet forwarding process as described in "IPv6 packet forwarding process."

     This command does not apply to software-forwarded packets.

Specifying a policy for IPv6 PBR

Specifying an IPv6 policy for IPv6 local PBR

About IPv6 local PBR

Perform this task to specify an IPv6 policy for IPv6 local PBR to guide the forwarding of locally generated packets.

Restrictions and guidelines

You can specify only one policy for IPv6 local PBR and must make sure the specified policy already exists. Before you apply a new policy, you must first remove the current policy.

IPv6 local PBR might affect local services, such as ping and Telnet. When you use IPv6 local PBR, make sure you fully understand its impact on local services of the device.

Procedure

1. Enter system view.

   system-view

2. Specify an IPv6 policy for IPv6 local PBR.
   ipv6 local policy-based-route *policy-name*
   By default, IPv6 local PBR is not enabled.

Specifying an IPv6 policy for IPv6 interface PBR

About interface PBR

Perform this task to apply an IPv6 policy to an interface to guide the forwarding of packets received on the interface only.

To apply an IPv6 policy to multiple VLAN interfaces at the same time, you can use the **ipv6 policy-based-route apply** command. Using this command simplifies configuration and saves device resources.

Restrictions and guidelines

You can apply only one policy to an interface and must make sure the specified policy already exists. Before you apply a new policy, you must first remove the current policy from the interface.

You can apply a policy to multiple interfaces.

Procedure

1. Enter system view.
   system-view
2. Enter interface view.
   interface *interface-type interface-number*
3. Specify an IPv6 policy for IPv6 interface PBR.
   ipv6 policy-based-route *policy-name*
   By default, no IPv6 policy is applied to the interface.

Display and maintenance commands for IPv6 PBR

Execute **display** commands in any view and **reset** commands in user view.

Task	Command
Display IPv6 PBR policy information.	**display ipv6 policy-based-route** [**policy** *policy-name*]

Task	Command
Display IPv6 interface PBR configuration and statistics.	**display ipv6 policy-based-route interface** *interface-type interface-number* [**slot** *slot-number*]
Display IPv6 local PBR configuration and statistics.	**display ipv6 policy-based-route local** [**slot** *slot-number*]
Display IPv6 PBR configuration.	**display ipv6 policy-based-route setup**
Clear IPv6 PBR statistics.	**reset ipv6 policy-based-route statistics** [**policy** *policy-name*]

IPv6 PBR configuration examples

Example: Configuring packet type-based IPv6 local PBR

Network configuration

As shown in Figure 1, Switch B and Switch C are connected through Switch A. Switch B and Switch C do not have a route to reach each other.

Configure IPv6 PBR on Switch A to forward all TCP packets to the next hop 1::2 (Switch B).

Figure 1 Network diagram

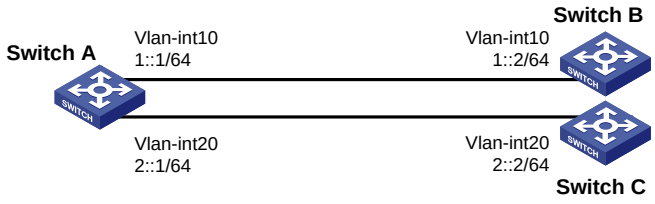

Procedure

1. Configure Switch A:

   # Create VLAN 10 and VLAN 20.

   ```
   <SwitchA> system-view
   [SwitchA] vlan 10
   ```

```
[SwitchA-vlan10] quit
[SwitchA] vlan 20
[SwitchA-vlan20] quit
```

Configure the IPv6 addresses of VLAN-interface 10 and VLAN-interface 20.

```
[SwitchA] interface vlan-interface 10
[SwitchA-Vlan-interface10] ipv6 address 1::1 64
[SwitchA-Vlan-interface10] quit
[SwitchA] interface vlan-interface 20
[SwitchA-Vlan-interface20] ipv6 address 2::1 64
[SwitchA-Vlan-interface20] quit
```

Configure ACL 3001 to match TCP packets.

```
[SwitchA] acl ipv6 advanced 3001
[SwitchA-acl-ipv6-adv-3001] rule permit tcp
[SwitchA-acl-ipv6-adv-3001] quit
```

Configure Node 5 for policy **aaa** to forward TCP packets to next hop 1::2.

```
[SwitchA] ipv6 policy-based-route aaa permit node 5
[SwitchA-pbr6-aaa-5] if-match acl 3001
[SwitchA-pbr6-aaa-5] apply next-hop 1::2
[SwitchA-pbr6-aaa-5] quit
```

Configure IPv6 local PBR by applying policy **aaa** to Switch A.

```
[SwitchA] ipv6 local policy-based-route aaa
```

2. Configure Switch B:

Create VLAN 10.

```
<SwitchB> system-view
[SwitchB] vlan 10
[SwitchB-vlan10] quit
```

Configure the IPv6 address of VLAN-interface 10.

```
[SwitchB] interface vlan-interface 10
[SwitchB-Vlan-interface10] ipv6 address 1::2 64
```

3. Configure Switch C:

Create VLAN 20.

```
<SwitchC> system-view
[SwitchC] vlan 20
[SwitchC-vlan20] quit
```

Configure the IPv6 address of VLAN-interface 20.

```
[SwitchC] interface vlan-interface 20
[SwitchC-Vlan-interface20] ipv6 address 2::2 64
```

Verifying the configuration

1. Perform telnet operations to verify that IPv6 local PBR on Switch A operates as configured to forward the matching TCP packets to the next hop 1::2 (Switch B), as follows:

Verify that you can telnet to Switch B from Switch A successfully. (Details not shown.)

Verify that you cannot telnet to Switch C from Switch A. (Details not shown.)

2. Verify that Switch A forwards packets other than TCP packets through VLAN-interface 20. For example, verify that you can ping Switch C from Switch A. (Details not shown.)

Example: Configuring packet type-based IPv6 interface PBR

Network configuration

As shown in Figure 2, Switch B and Switch C do not have a route to reach each other.

Configure IPv6 PBR on Switch A to forward all TCP packets received on VLAN-interface 11 to the next hop 1::2 (Switch B).

Figure 2 Network diagram

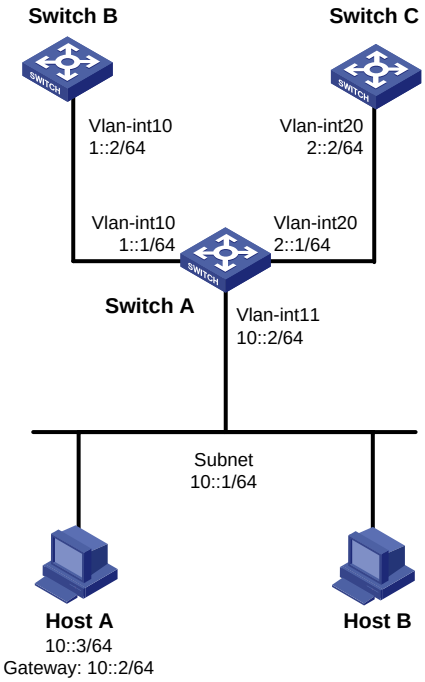

Procedure

1. Configure Switch A:

   # Create VLAN 10 and VLAN 20.

   ```
   <SwitchA> system-view
   [SwitchA] vlan 10
   [SwitchA-vlan10] quit
   [SwitchA] vlan 20
   [SwitchA-vlan20] quit
   ```

   # Configure RIPng.

   ```
   [SwitchA] ripng 1
   [SwitchA-ripng-1] quit
   [SwitchA] interface vlan-interface 10
   [SwitchA-Vlan-interface10] ipv6 address 1::1 64
   [SwitchA-Vlan-interface10] ripng 1 enable
   [SwitchA-Vlan-interface10] quit
   [SwitchA] interface vlan-interface 20
   [SwitchA-Vlan-interface20] ipv6 address 2::1 64
   [SwitchA-Vlan-interface20] ripng 1 enable
   [SwitchA-Vlan-interface20] quit
   ```

   # Configure ACL 3001 to match TCP packets.

```
[SwitchA] acl ipv6 advanced 3001
[SwitchA-acl-ipv6-adv-3001] rule permit tcp
[SwitchA-acl-ipv6-adv-3001] quit
```

Configure Node 5 for policy **aaa** to forward TCP packets to next hop 1::2.

```
[SwitchA] ipv6 policy-based-route aaa permit node 5
[SwitchA-pbr6-aaa-5] if-match acl 3001
[SwitchA-pbr6-aaa-5] apply next-hop 1::2
[SwitchA-pbr6-aaa-5] quit
```

Configure IPv6 interface PBR by applying policy **aaa** to VLAN-interface 11.

```
[SwitchA] interface vlan-interface 11
[SwitchA-Vlan-interface11] ipv6 address 10::2 64
[SwitchA-Vlan-interface11] undo ipv6 nd ra halt
[SwitchA-Vlan-interface11] ripng 1 enable
[SwitchA-Vlan-interface11] ipv6 policy-based-route aaa
```

2. Configure Switch B:

Create VLAN 10.

```
<SwitchB> system-view
[SwitchB] vlan 10
[SwitchB-vlan10] quit
```

Configure RIPng.

```
[SwitchB] ripng 1
[SwitchB-ripng-1] quit
[SwitchB] interface vlan-interface 10
[SwitchB-Vlan-interface10] ipv6 address 1::2 64
[SwitchB-Vlan-interface10] ripng 1 enable
[SwitchB-Vlan-interface10] quit
```

3. Configure Switch C:

Create VLAN 20.

```
<SwitchC> system-view
[SwitchC] vlan 20
[SwitchC-vlan20] quit
```

Configure RIPng.

```
[SwitchC] ripng 1
[SwitchC-ripng-1] quit
[SwitchC] interface vlan-interface 20
[SwitchC-Vlan-interface20] ipv6 address 2::2 64
[SwitchC-Vlan-interface20] ripng 1 enable
[SwitchC-Vlan-interface20] quit
```

Verifying the configuration

1. Enable IPv6 and configure the IPv6 address 10::3 for Host A.

```
C:\>ipv6 install
Installing...
Succeeded.
C:\>ipv6 adu 4/10::3
```

2. Perform telnet operations to verify that IPv6 interface PBR on Switch A operates as configured to forward the matching TCP packets to the next hop 1::2 (Switch B), as follows:

Verify that you can telnet to Switch B from Host A successfully. (Details not shown.)

Verify that you cannot telnet to Switch C from Host A. (Details not shown.)

3. Verify that Switch A forwards packets other than TCP packets through VLAN-interface 20. For example, verify that you can ping Switch C from Host A. (Details not shown.)