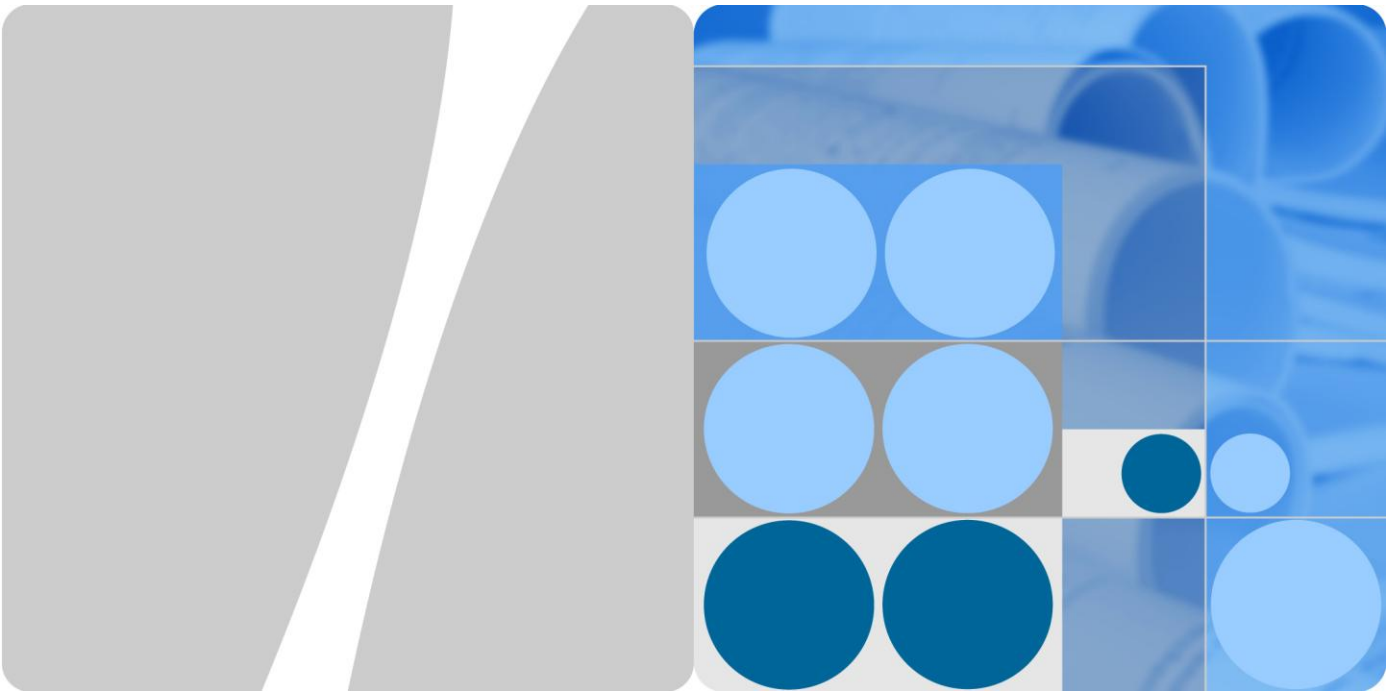

VRRP Technology White Paper

Issue **01**

Date **2012-08-31**

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

 and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://www.huawei.com

Email: support@huawei.com

1 VRRP

About This Chapter

1.1 Introduction to VRRP

Definition

The Virtual Router Redundancy Protocol (VRRP) groups multiple routing devices into a virtual router. One device functions as the master, and the others as the backup devices. When the next hop device of the master device fails, VRRP switches services to a backup device. This implementation ensures nonstop service transmission and reliability.

Purpose

As networks rapidly develop and applications become diversified, various value-added services, such as IPTV and video conferencing are widely used. Demands for network infrastructure reliability are increasing, especially for nonstop service transmission.

In practice, hosts use the default gateway to communicate with external networks. If the default gateway is faulty, hosts cannot access external networks. Configuring dynamic routing protocols such as RIP, OSPF, or ICMP can improve system reliability. However, dynamic routing protocols are difficult to configure and may not be supported by some hosts.

VRRP prevents communication failures in a better way than the preceding solution. VRRP virtualizes multiple routing devices into a virtual router without changing the networking. The

virtual router IP address is configured as the default gateway address. When the master device in the virtual router becomes faulty, VRRP uses a backup device to transmit service traffic.

Benefits

VRRP has the following benefits:

- Simplified network management: On a multicast or broadcast LAN such as an Ethernet network, a logical VRRP gateway ensures reliability for key links. VRRP configuration is simple and takes effect without modification in configurations of dynamic routing protocols and route discovery protocols.
- Strong adaptability: VRRP packets are encapsulated in IP packets, so VRRP supports multiple upper-layer protocols.
- Low cost: VRRP defines only VRRP protocol packets, which reduces burden on network devices.

1.2 References

The following table lists the references of this document.

Document	Description	Remarks
RFC 2281	Hot Standby Router Protocol (HSRP)	–
RFC 2338	Virtual Router Redundancy Protocol (version number One1998)	–
RFC 2787	Definitions of Managed Objects for the Virtual Router Redundancy Protocol	–
RFC 3768	Virtual Router Redundancy Protocol (version number Two 2004)	–
RFC 5798	Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6	–

1.3 Principles

1.3.1 Basic Concepts of VRRP

Figure 1-1 shows VRRP networking.

Figure 1-1 VRRP networking

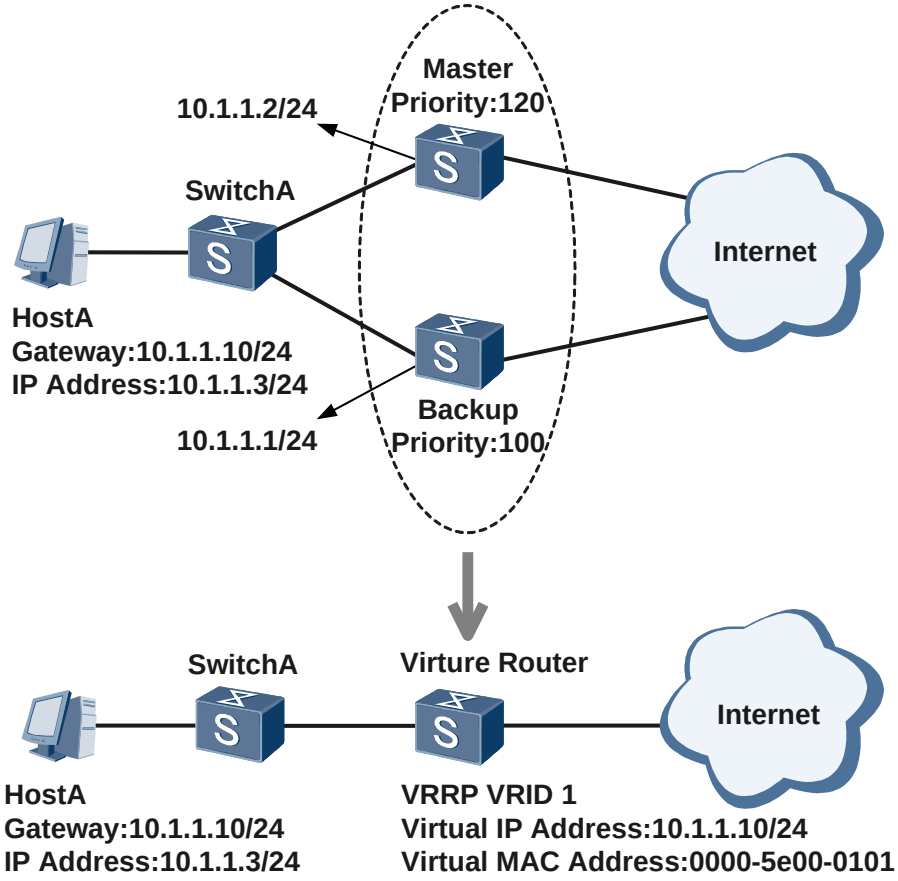

- VRRP router: device running VRRP. It may join one or more virtual routers.
- Virtual router: VRRP group. It consists of one master router and one or more backup routers. The VRRP group is used as the default gateway on a LAN.
- Virtual router master: VRRP router that forwards packets.
- Virtual router backup: a group of VRRP routers that do not forward packets. When the master router is faulty, a backup router preempts to be the new master router.
- VRID: virtual router ID.
- Virtual IP address: IP address of a virtual router. A virtual router can be assigned one or more virtual IP addresses. Virtual IP addresses are configurable.
- IP address owner: VRRP router that uses an IP address of a virtual router as the actual interface address. If an IP address owner is available, it usually functions as the virtual router master.
- Virtual MAC address: MAC address that is generated by the virtual router based on the virtual router ID. A virtual router has one virtual MAC address and is in the format of 00-00-5E-00-01-{VRID}(VRRP for IPv4) or 00-00-5E-00-02-{VRID}(VRRP for IPv6). The virtual router sends ARP Reply packets using the virtual MAC address instead of the interface MAC address.
- Primary IP address: is selected from one of actual IP addresses of interfaces. Usually, it is the first configured IP address. The primary IP address is often used as the source IP address for VRRP broadcast packets.

- Priority: priority of a VRRP router. The virtual router selects the master and backup routers based on the priority.
- Preemption mode: If the priority of a virtual router backup is higher than the priority of the current virtual router master, the virtual router backup automatically becomes the virtual router master.
- Non-preemption mode: As long as the virtual router master is working properly, the backup with a higher priority cannot become the virtual router master.

1.3.2 VRRP Packets

VRRP packets are sent to notify all backup routers in a VRRP group of the master router priority and status.

VRRP packets are encapsulated into IP packets and sent to the VRRP virtual IP address. In the IP packet header, the source address is the primary IP address of the interface that sends the packets, the destination address is 224.0.0.18, the TTL is 255, and the protocol number is 112. The primary IP address is not the virtual IP address.

VRRP has two versions: VRRPv2 and VRRPv3. VRRPv2 applies to the IPv4 network, and VRRPv3 applies to IPv4 and IPv6 networks.

VRRP is classified into VRRP for IPv4 and VRRP for IPv6 (VRRP6) by network type. VRRP for IPv4 supports VRRPv2 and VRRPv3, and VRRP for IPv6 supports only VRRPv3.

VRRPv2 and VRRPv3 have the following differences:

- Support different networks. VRRPv3 applies to IPv4 and IPv6 networks, whereas VRRPv2 applies to only the IPv4 network.
- Have different authentication functions. VRRPv3 does not support authentication, whereas VRRPv2 supports.
- Use different units for the interval at which VRRP Advertisement packets are sent. VRRPv3 uses the centiseconds, whereas VRRPv2 uses the seconds.

The following example describes how to configure the VRRP version number.

NOTE

S9300 V200R001C00 is used as an example.

Step 1 Run the **system-view** command to enter the system view.

```
<Quidway> system-view
```

Step 2 Run the **vrrp version { v2 | v3 }** command to configure a VRRP version number.

By default, VRRPv2 is used. Here, VRRPv3 is used.

```
[Quidway] vrrp version v3
```

Step 3 (Optional) Run the **vrrp version-3 send-packet-mode** { v2-only | v3-only | v2v3-both } command to set the mode in which VRRPv3 Advertisement packets are sent if VRRPv3 is used.

The default mode is **v3-only**. Here, **v2v3-both** is used.

```
[Quidway] vrrp version-3 send-packet-mode v2v3-both
```

1.3.3 VRRP Implementation

VRRP State Machine

VRRP defines three statuses: Initialize, Master, and Backup. Only the device in Master state can forward packets destined for the virtual IP address.

Figure 1-2 shows VRRP state transition.

Figure 1-2 VRRP state transition

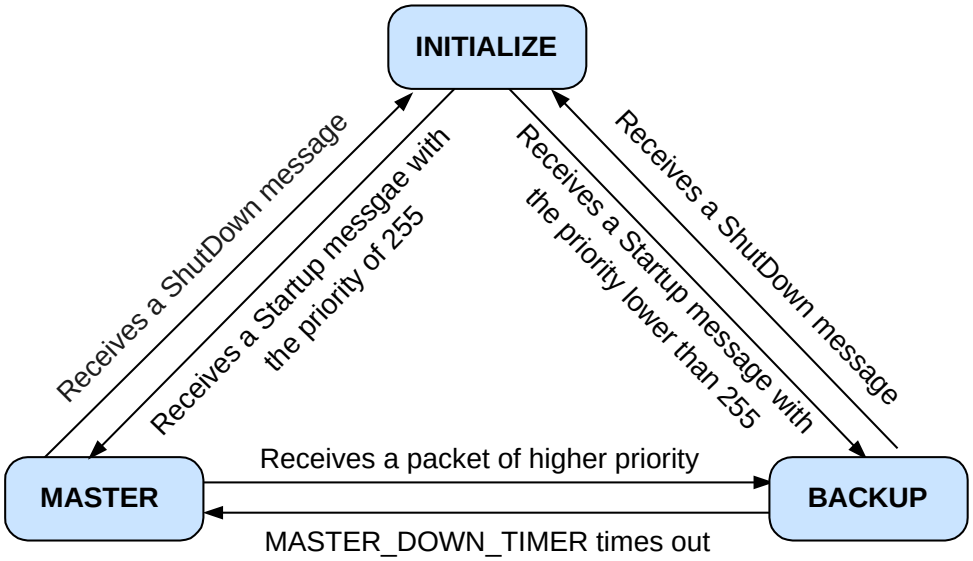

Table 1-1 describes VRRP states.

Table 1-1 VRRP states

State	Description
Initialize	VRRP is unavailable. The device in Initialize state cannot process VRRP packets. When a device starts or detects a fault, it enters the Initialize state.
Master	The VRRP device in Master state performs the following operations: • Sends VRRP Advertisement packets at intervals. • Uses the virtual MAC address to respond to ARP Request packets destined for the virtual IP address. • Forwards IP packets destined for the virtual MAC address. • Processes the IP packets destined for the virtual IP address if the device is an IP address owner. If the device is not the IP address owner, it discards the IP packets destined for the virtual IP address.
Backup	The VRRP device in Backup state performs the following

State	Description
	operations: • Receives VRRP Advertisement packets from the master and determines whether the master works properly. • Does not respond to ARP Request packets destined for the virtual IP address. • Discards IP packets destined for the virtual MAC address. • Discards IP packets destined for the virtual IP address. • Discards packets that carry a lower priority than the device and does not reset the Master_Down_Interval timer; resets the Master_Down_Interval timer and does not compare IP addresses if the received packet carries the same priority as the device. **NOTE:** Master_Down_Interval timer: If the backup does not receive Advertisement packets after the timer expires, the backup becomes the master. The calculation formula is as follows: Master_Down_Interval = 3xAdvertisement_Interval + Skew_time (offset time) Skew_Time = (256–Priority)/256

VRRP Working Process

The VRRP working process is as follows:

Devices in a VRRP group select the master based on device priorities. The master sends gratuitous ARP packets to notify the connected device or host of its virtual MAC address.

The master periodically sends VRRP Advertisement packets to all backups in the VRRP group to advertise its configuration and running status.

If the master becomes faulty, the backups in the group select a new master based on priorities.

When the VRRP group status changes, a new master is used. The new master sends gratuitous ARP packets carrying the virtual MAC address and virtual IP address of the virtual router to update the MAC address entry on the connected host or device. Then user traffic is switched to the new master. This process is transparent to users.

When the original master recovers and is the IP address owner, the original master directly switches to the Master state. If the original master is not the IP address owner, it first switches to the Backup state and its original priority is restored.

To ensure that the master and backup cooperate, VRRP must be able to:

- Select the master.
- Advertise the master status.

Selecting the Master

VRRP determines the device role in the virtual router based on device priorities. The device with a higher priority is more likely to become the master.

The VRRP-enabled device in the VRRP group first works in Initialize state. After receiving an interface Up message, the VRRP-enabled device with its priority less than 255 first

switches to the Backup state. After the Master_Down_Interval timer expires, the VRRP-enabled device switches to the Master state again. The device that first switches to the Master state obtains priorities of other devices in the group by exchanging VRRP Advertisement packets. Then the master is selected.

- If the master priority in VRRP packets is higher than or equal to the priority of the device, the backup retains in Backup state.
- If the master priority in VRRP packets is lower than the priority of the device, the backup in preemption mode switches to the Master state or the backup in non-preemption mode retains in Backup state.

NOTE

- If multiple devices in the group switch to the master, the devices with a lower priority switch to the Backup state and the device with the highest priority becomes the master after these devices exchange Advertisement packets. If multiple devices have the same priority, the device where the interface with the largest IP address resides is the master.
- If the device is the IP address owner, it switches to the Master state immediately after receiving an interface Up message.

Advertising the Master Status

The master periodically sends VRRP Advertisement packets to all backups in the VRRP group to advertise its configuration and running status. The backup determines whether the master works properly based on the received VRRP Advertisement packets.

- When the master does not retain the Master state, for example, the master leaves the group, it sends a VRRP Advertisement packet with priority 0. In this manner, a backup can switch to the master immediately without waiting for the Master_Down_Interval timer to expire. The switchover period is called Skew time, in seconds. The value is calculated using the following formula: Skew time = (256 - Backup priority)/256
- If the master cannot send VRRP Advertisement packets due to network faults, the backups cannot learn the running status of the master. The backups consider the master faulty only after the Master_Down_Interval timer expires. Then a backup switches to the Master state. Master_Down_Interval = 3 x Advertisement_Interval + Skew_time (in seconds)

NOTE

If congestion occurs on an unstable network, the backup may not receive VRRP Advertisement packets from the master within the period of Master_Down_Interval. A backup then switches to the Master state. If the VRRP Advertisement packet from the original master reaches the backup (new master), the new master switches to the Backup state. In this case, the VRRP group status changes frequently. To solve the problem, the preemption delay is used. When the Master_Down_Interval timer expires, the backup waits for the preemption delay. If the backup does not receive a VRRP Advertisement packet within the preemption delay, it switches to the Master state.

VRRP Authentication

Authentication modes and keys can be set based on network security requirements, and these settings are carried in the headers of VRRP Advertisement packets.

- On a highly secure network, you can use non-authentication. The device does not authenticate VRRP Advertisement packets to be sent. In addition, the device does not authenticate the received VRRP packets. It considers all the received packets as valid.
- On a vulnerable network, either simple or Message Digest 5 (MD5) authentication can be performed:

- Simple authentication: The device encapsulates the authentication mode and authentication key into an outgoing VRRP Advertisement packet. The device that receives the VRRP Advertisement packet compares the authentication mode and authentication key in the packet with those configured on itself. If the values are the same, the device considers the received VRRP Advertisement packet valid. If the values are different, the device considers the received VRRP Advertisement packet invalid and discards it.
- MD5 authentication: The device uses the MD5 algorithm to encrypt the authentication key and encapsulates the key in the Authentication Data field of an outgoing VRRP Advertisement packet. The device that receives the VRRP Advertisement packet matches the authentication mode with the decrypted authentication key in the packet.

NOTE

Only VRRPv2 supports authentication.

MD5 authentication provides higher security than simple authentication.

The following example describes how to configure an authentication mode of VRRP packets.

NOTE

S9300 V200R001C00 is used as an example.

Step 1 Run the **system-view** command to enter the system view.

```
<Quidway> system-view
```

Step 2 Run the **interface** *interface-type interface-number* command to enter the interface view.

Enter the view of VLANIF 100 enabled with VRRP.

```
[Quidway] interface vlanif 100
```

Step 3 Run the **vrrp vrid virtual-router-id authentication-mode** { **simple** { key | plain key | cipher cipher-key } | **md5** md5-key } command to configure an authentication mode of VRRP packets.

Configure MD5 authentication for VRRP group 1 and set the authentication key to hello.

```
[Quidway-Vlanif100] vrrp vrid 1 authentication-mode md5 hello
```

NOTE

Devices in a VRRP group must be configured with the same authentication mode and authentication key; otherwise, they cannot negotiate the Master and Backup status.

1.3.4 VRRP Active/Standby Mode

VRRP often uses the active/standby mode, as shown in Figure 1-3. In active/standby mode, a virtual router must be set up. The virtual router consists of a master router and multiple backup routers.

Figure 1-3 VRRP in active/standby mode

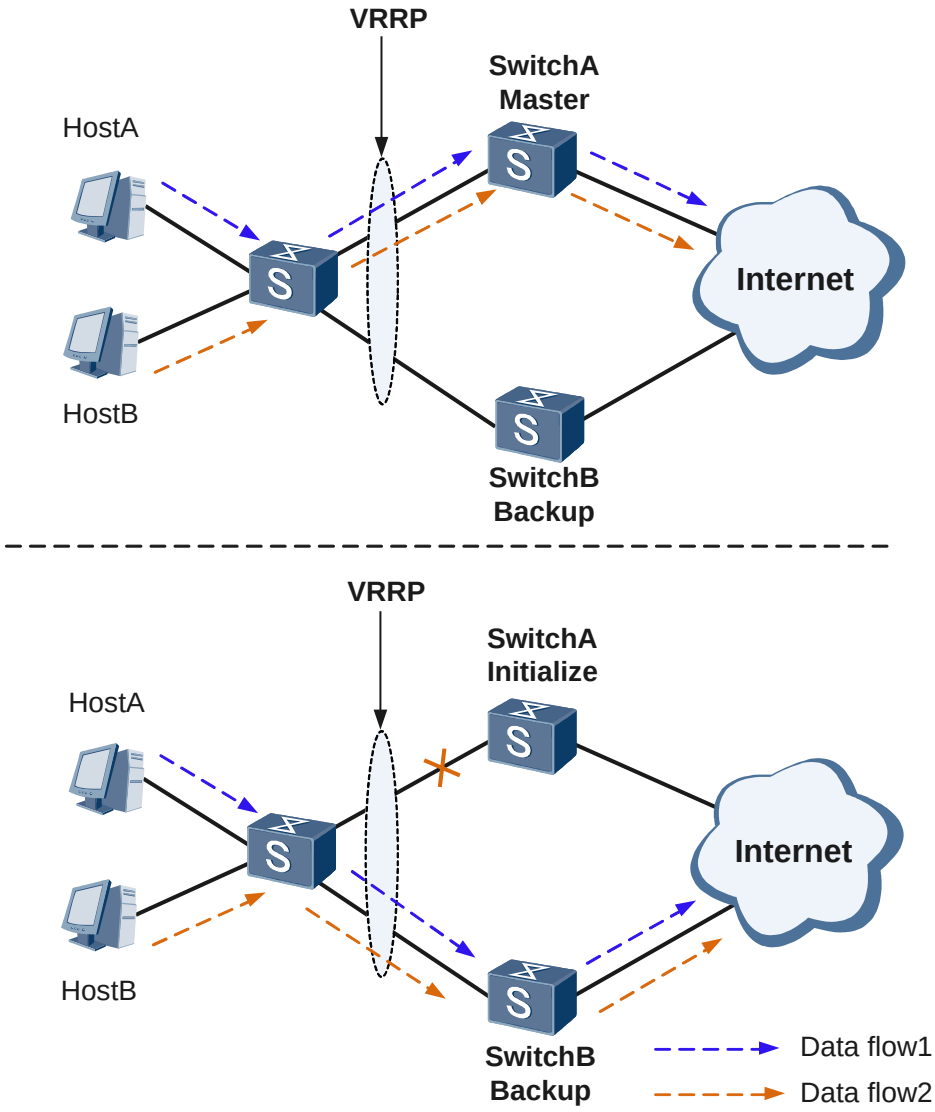

SwitchA is the master and forwards service packets. SwitchB is the backup device and does not forward services. SwitchA periodically sends VRRP Advertisement packets to SwitchB, notifying that SwitchA itself works properly. If SwitchA is faulty, SwitchB becomes the master to take over services.

After SwitchA recovers, it becomes the master in preemption mode. In non-preemption mode, SwitchA retains in Backup state.

The following example describes how to configure VRRP in active/standby mode.

NOTE

Only key procedures are provided here. S9300 V200R001C00 is used as an example.

Create VRRP group 1 on SwitchA and set the priority of SwitchA to 120 so that SwitchA functions as the master.

```
[SwitchA] interface vlanif 100
[SwitchA-Vlanif100] vrrp vrid 1 virtual-ip 10.1.1.111
```

```
[SwitchA-Vlanif100] vrrp vrid 1 priority 120
[SwitchA-Vlanif100] vrrp vrid 1 preempt-mode timer delay 20
[SwitchA-Vlanif100] quit
```

Configure VRRP group 1 on SwitchB. SwitchB uses default value 100. When the master fails, SwitchB becomes the master to take over services.

```
[SwitchB] interface vlanif 100
[SwitchB-Vlanif100] vrrp vrid 1 virtual-ip 10.1.1.111
[SwitchB-Vlanif100] quit
```

1.3.5 VRRP Load Balancing Mode

In load balancing mode, multiple VRRP groups transmit services simultaneously, as shown in Figure 1-4. The implementation and packet negotiation in load balancing mode are similar to those in active/standby mode. Each VRRP group has one master device and multiple backup devices. In load balancing mode, multiple VRRP groups need to be set up and use different master devices. A VRRP device can join multiple VRRP groups and has different priorities in these VRRP groups.

Figure 1-4 VRRP in load balancing mode

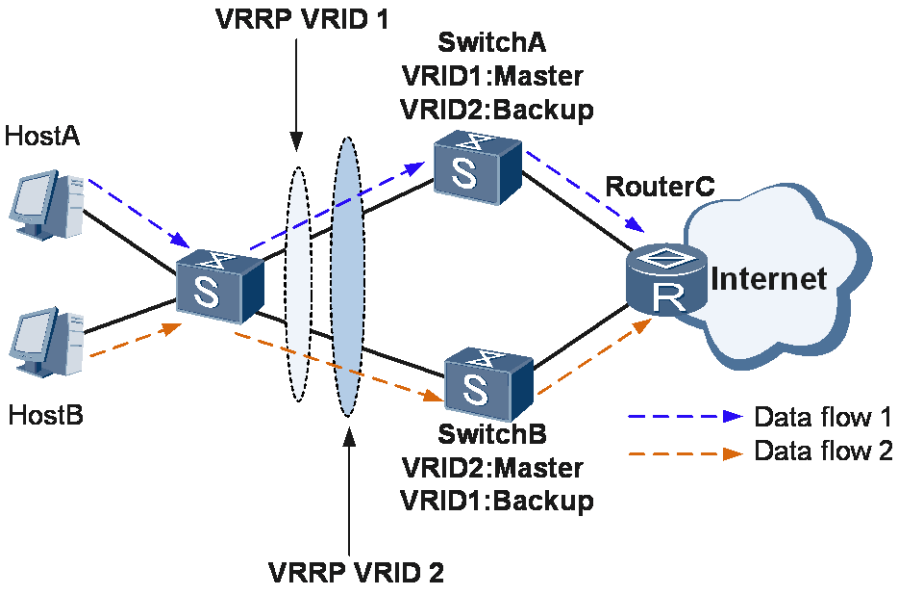

As shown in Figure 1-4, two VRRP groups are configured:

- VRRP group 1: SwitchA functions as the master and SwitchB as the backup.
- VRRP group 2: SwitchB functions as the master and SwitchA as the backup.
- Backup groups 1 and 2 are gateways for different hosts.

Multiple VRRP groups load balance traffic and back up each other.

The following example describes how to configure VRRP in load balancing mode.

NOTE

Only key procedures are provided here. S9300 V200R001C00 is used as an example.

Create VRRP 1 on SwitchA and SwitchB, set the priority of SwitchA to 120 and the preemption delay to 20s, and set the default priority for SwitchB.

```
[SwitchA] interface vlanif 100
[SwitchA-Vlanif100] vrrp vrid 1 virtual-ip 10.1.1.111
[SwitchA-Vlanif100] vrrp vrid 1 priority 120
[SwitchA-Vlanif100] vrrp vrid 1 preempt-mode timer delay 20
[SwitchA-Vlanif100] quit
[SwitchB] interface vlanif 100
[SwitchB-Vlanif100] vrrp vrid 1 virtual-ip 10.1.1.111
[SwitchB-Vlanif100] quit
```

Create VRRP 2 on SwitchA and SwitchB, set the priority of SwitchB to 120 and the preemption delay to 20s, and set the default priority for SwitchA.

```
[SwitchB] interface vlanif 100
[SwitchB-Vlanif100] vrrp vrid 2 virtual-ip 10.1.1.112
[SwitchB-Vlanif100] vrrp vrid 2 priority 120
[SwitchB-Vlanif100] vrrp vrid 2 preempt-mode timer delay 20
[SwitchB-Vlanif100] quit
[SwitchA] interface vlanif 100
[SwitchA-Vlanif100] vrrp vrid 2 virtual-ip 10.1.1.112
[SwitchA-Vlanif100] quit
```

1.3.6 Smooth VRRP Switching

When an active/standby switchover of the main control boards occurs on the master, the master cannot send VRRP protocol packets before a new master starts to work. After the Master_Down_Interval timer expires, the backup switches to the master if it does not receive VRRP Advertisement packets. In this situation, two master devices coexist. After the original master device completes the active/standby switchover, it detects that it has a higher priority than the other master device, and therefore retains the Master state in preemption mode. The other master device switches back to the Backup state. During this process, services are switched twice, causing unstable service transmission.

To prevent the impact of the active/standby switchover on service traffic, enable VRRP smooth switching on the master. During VRRP smooth switching, the master cooperates with the backup to ensure smooth service transmission.

- Before VRRP smooth switching, you must configure the backup to learn the interval at which VRRP packets are sent. After receiving an Advertisement packet from the master, the backup checks the interval in the packet. If the received interval is different from its interval, the backup learns the interval and adjusts its own interval to be the same as the learned interval.
- When starting an active/standby switchover, the master will save the current interval at which VRRP Advertisement packets are sent and set the smooth VRRP switching time to the new interval. During smooth VRRP switching, the master sends a VRRP Advertisement packet at the new interval.
- After receiving the packet, the backup learns the interval in the packet and adjusts its own interval to be the same as the learned interval.
- After the switchover is complete, the master restores its original interval and sends an Advertisement packet at the new interval. After receiving the packet, the backup learns the interval again.

NOTE

During VRRP smooth switching, the learning function takes precedence over the preemption function. When the interval carried in the received packet is different from the current interval and the priority carried in the received packet is lower than the configured priority, the learning function takes effect and the timer is reset.

VRRP smooth switching also depends on the system. If the system is busy since the switchover and cannot schedule the operation of the VRRP module, VRRP smooth switching cannot take effect.

The following example describes how to configure smooth VRRP switching.

NOTE

Only key procedures are provided here. S9300 V200R001C00 is used as an example.

Enable smooth VRRP switching and set the interval carried in VRRP Advertisement packets during smooth VRRP switching. Here, the interval carried in VRRP Advertisement packets during smooth VRRP switching is 80s.

```
<Quidway> system-view
[Quidway] vrrp smooth-switching timer 20
```

NOTE

By default, smooth VRRP switching is enabled, and the interval carried in VRRP Advertisement packets is 100 seconds.

1.3.7 mVRRP

A UPE is usually dual-homed to two NPEs to improve network reliability. Multiple VRRP groups can be configured on the two NPEs to transmit various types of services. Each VRRP group needs to maintain its own state machine; therefore, a large number of VRRP packets are transmitted between NPEs.

As shown in Figure 1-5, to decrease bandwidth and CPU resources occupied by protocol packets, configure a VRRP group as an mVRRP group and bind other service VRRP groups to the mVRRP group. The mVRRP group sends VRRP Advertisement packets to determine the master and backup status for its service VRRP groups.

Figure 1-5 mVRRP networking

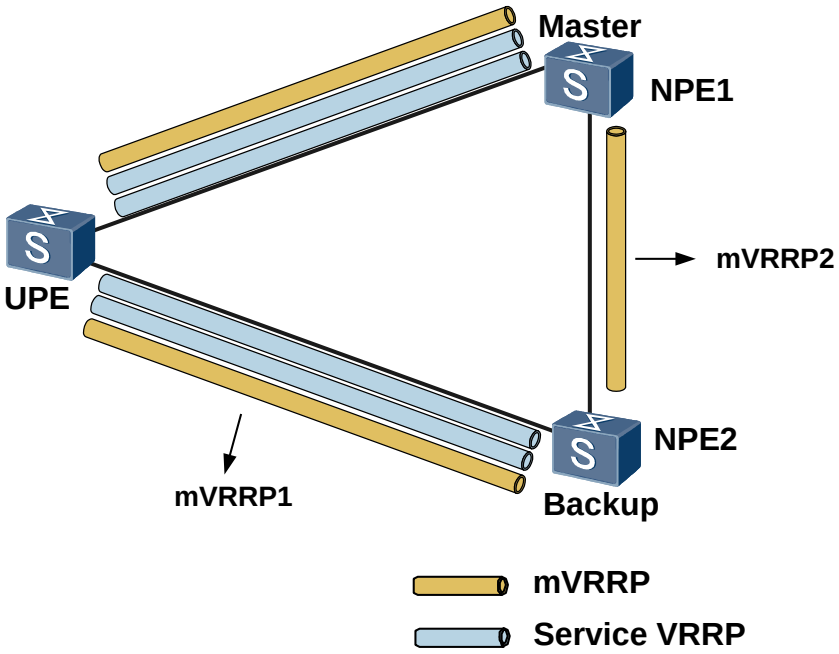

mVRRP is used in the following scenarios:

- When an mVRRP group functions as the gateway (mVRRP1 in Figure 1), the mVRRP group determines the Master and Backup status and forwards service traffic. You must create a VRRP group and configure a virtual IP address as the gateway address, and then configure this VRRP group as an mVRRP group
- When an mVRRP group does not function as the gateway (mVRRP2 in Figure 1), the mVRRP group only determines the master and backup status, and cannot forward service traffic. The mVRRP group does not require a virtual IP address, and you can directly create an mVRRP group on an interface. mVRRP simplifies maintenance.

The following example describes how to configure mVRRP.

NOTE

Only key procedures are provided here. S9300 V200R001C00 is used as an example.

Configure VRRP group 1 as an mVRRP group.

```
[Quidway-Vlanif10] admin-vrrp vrid 1
```

(Optional) Bind VRRP group 2 to mVRRP group 1.

```
[Quidway-Vlanif20] vrrp vrid 2 track admin-vrrp interface vlanif 20 vrid 1
unflowdown
```

1.4 Applications

1.4.1 Association Between a VRRP Group and the Interface Status to Monitor the Uplink

A VRRP group can only detect the status change of the interface on which the VRRP group is configured. VRRP cannot detect faults on the uplink interface or direct uplink of the master, so services are interrupted. You can associate a VRRP group with the interface status. When the uplink interface or direct uplink of the master fails, the priority of the master is adjusted. This triggers an active/standby switchover, ensuring proper traffic forwarding.

A VRRP group can monitor the interface status in Increased or Reduced mode:

- In Increased mode, when the monitored interface becomes Down, the priority of the device where the monitored interface resides increases.
- In Reduced mode, when the monitored interface becomes Down, the priority of the device where the monitored interface resides decreases.

Figure 1-6 Association between a VRRP group and the interface status

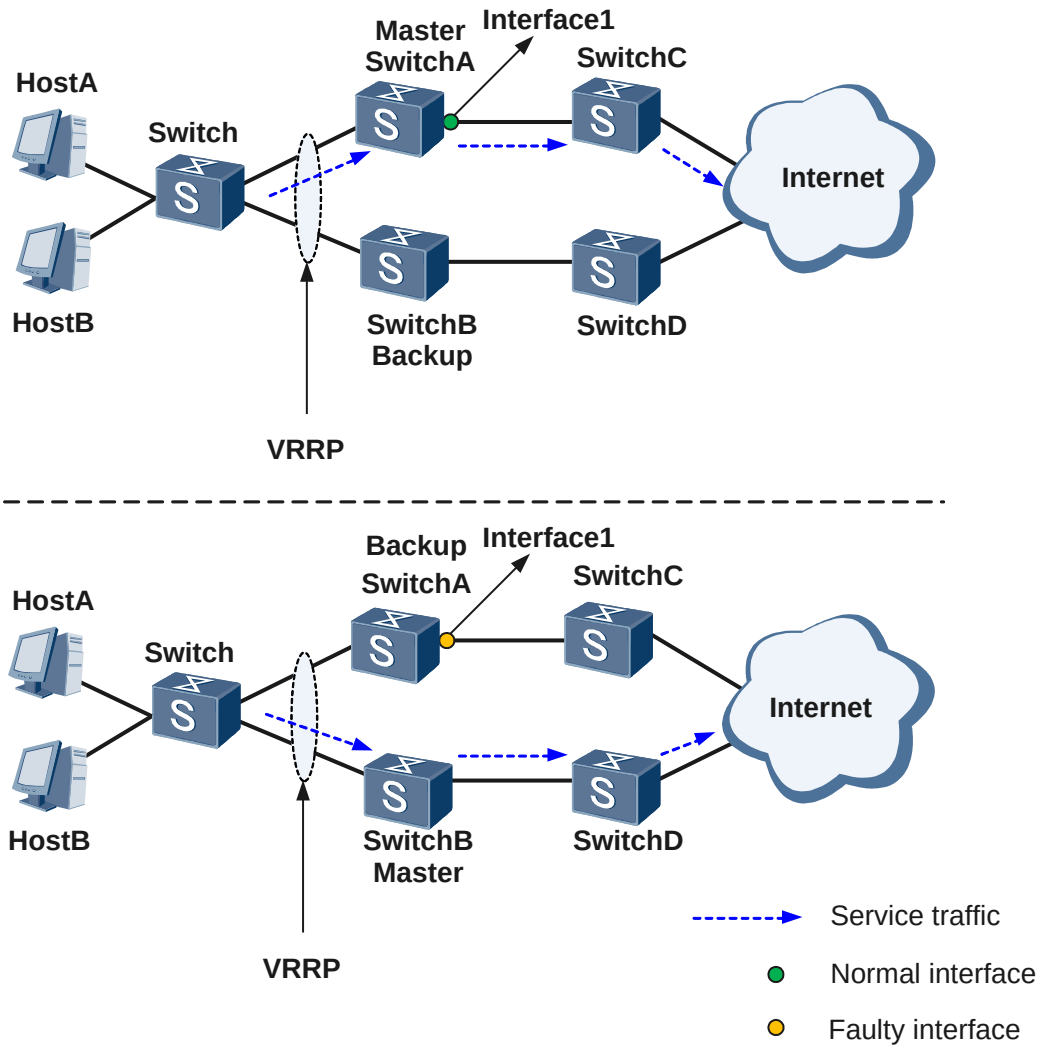

As shown in Figure 1-6, a VRRP group is configured between SwitchA and SwitchB. SwitchA is the master and SwitchB is the backup. SwitchA and SwitchB work in preemption mode. On SwitchA, the Reduced mode is used to monitor uplink interface Interface1. When Interface1 becomes faulty, the priority of SwitchA decreases. Then SwitchB becomes the master through negotiation so that user traffic is forwarded correctly.

1.4.2 Association Between VRRP and BFD/NQA/Routing to Monitor the Uplink

VRRP can detect only faults in VRRP groups. Association between a VRRP group and the interface status allows the Switch to detect faults on the uplink interface or or direct uplink of the master. When an indirect uplink of the master fails, VRRP cannot detect the fault, causing user traffic loss. You can configure association between VRRP and BFD/NQA/routing to solve this problem. Association between VRRP and BFD/NQA/routing allows the Router to detect faults on the uplink of the master. When the uplink of the master fails, BFD/NQA/routing rapidly detects the fault and notifies the master of adjusting its priority. This triggers an active/standby switchover, ensuring proper traffic forwarding.

Figure 1-7 Association between VRRP and BFD/NQA/routing to monitor the uplink

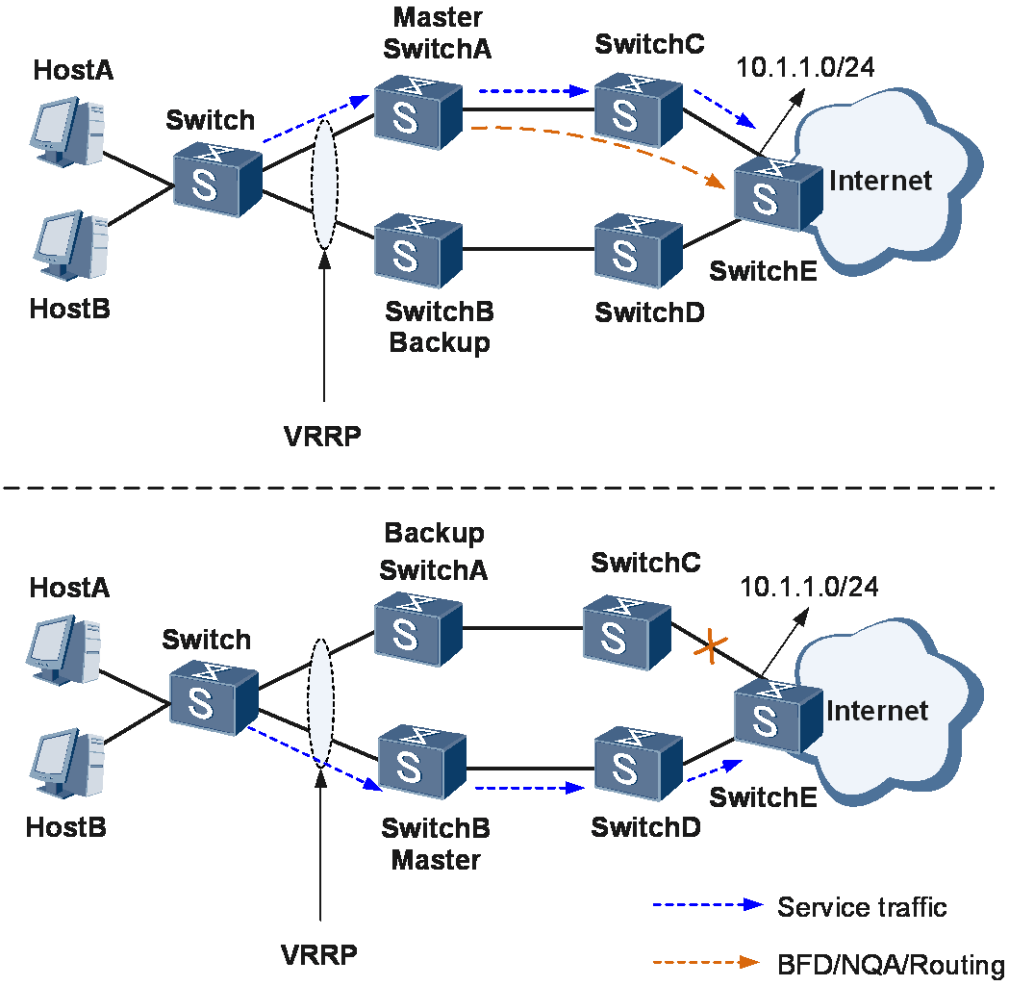

As shown in Figure 1-7, a VRRP group is configured between SwitchA and SwitchB. SwitchA is the master and SwitchB is the backup. SwitchA and SwitchB work in preemption mode. BFD/NQA/routing is configured to detect faults on the link from SwitchA to SwitchE, and association between VRRP and BFD/NQA/routing is configured on SwitchA. When BFD/NQA/routing detects the fault on the link from SwitchC to SwitchE, BFD/NQA/routing notifies the master of adjusting its priority. SwitchB becomes the master through negotiation so that user traffic is forwarded correctly.

1.4.3 Association Between VRRP and BFD to Implement a Rapid Active/Standby Switchover

A VRRP group sends and receives VRRP packets to determine the master and backup status. When a VRRP group is faulty, the backup detects the fault and switches to the master after the Master_Down_Interval timer expires. The switchover period is at least 3s. During the switchover period, service traffic is still sent to the original master, causing user traffic loss. Association between VRRP and BFD prevents traffic loss. Configure a BFD session between the master and the backup, and bind the BFD session to a VRRP group. The BFD session detects connectivity of the VRRP group. After detecting a fault, the BFD session notifies the VRRP group of an active/standby switchover. This mechanism implements millisecond-level switchover and reduces traffic loss.

Figure 1-8 Association between VRRP and BFD to implement a rapid active/standby switchover

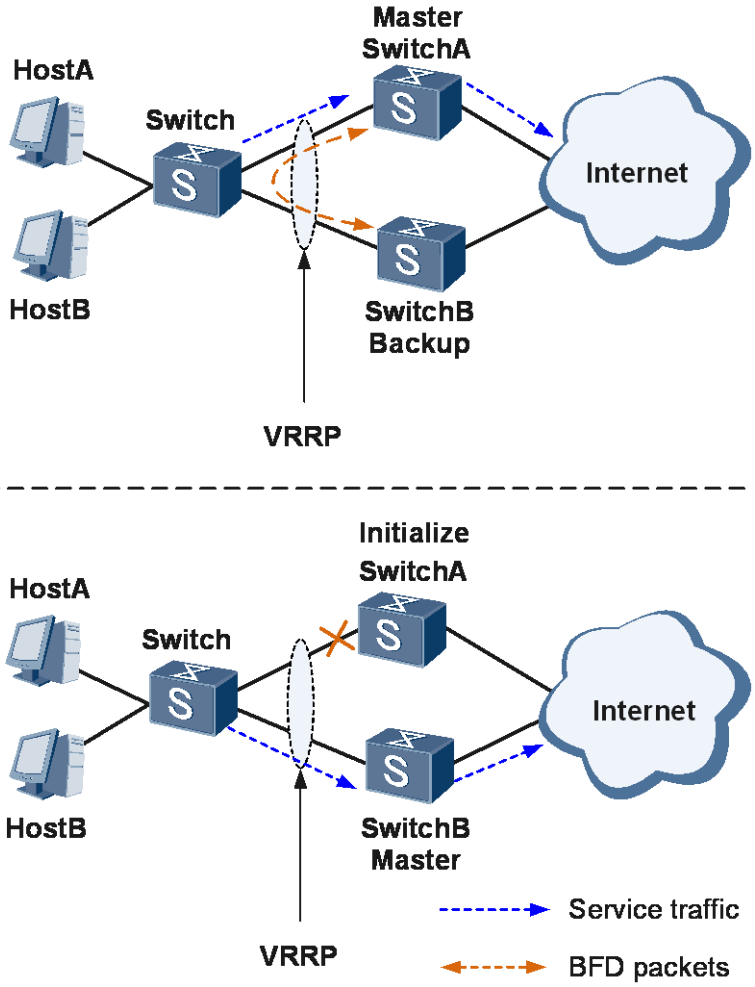

As shown in Figure 1-8, a VRRP group is configured between SwitchA and SwitchB. SwitchA is the master and SwitchB is the backup. User traffic is forwarded through SwitchA. Delayed preemption is configured on SwitchA and immediate preemption is configured on SwitchB. BFD sessions are configured on SwitchA and SwitchB and association between VRRP and BFD is configured on SwitchB.

When a fault occurs in the VRRP group, BFD rapidly detects the fault and notifies SwitchB of increasing the priority. In this case, SwitchB has a higher priority than SwitchA. SwitchB becomes the master and user traffic is forwarded through SwitchB. This implements a rapid active/standby switchover.

1.5 Troubleshooting Cases

1.5.1 Multiple Masters Coexist in One VRRP Group

Fault Description

Multiple masters exist in a VRRP group.

Procedure

Step 1 Ping masters to check network connectivity between masters.

If the ping operation fails, check whether the network connection is correct.

If the ping operation is successful, go to step 2.

Step 2 Run the **display vrrp protocol-information** command in any view to check the VRRP version on each master is compatible with the mode in which VRRP Advertisement packets are sent.

If the version is incompatible with the mode, run the **vrrp version** { **v2** | **v3** } command in the system view to change the version.

If the version is compatible with the mode, go to step 3.

NOTE

- A VRRPv2 group can only send and receive VRRPv2 Advertisement packets. It discards the received VRRPv3 Advertisement packets.
- A VRRPv3 group can send and receive both VRRPv2 and VRRPv3 Advertisement packets. You can configure the mode in which VRRPv3 Advertisement packets are sent. The mode can be **v2-only**, **v3-only**, or **v2v3-both**.

Step 3 Run the **display vrrp** *virtual-router-id* command in any view to check whether the master uses the same virtual IP address, interval at which VRRP Advertisement packets are sent, authentication mode, and authentication key.

- If the configured virtual IP addresses are different, run the **vrrp vrid** *virtual-router-id virtual-ip virtual-address* command to set the same virtual IP address.
- If the intervals are different, run the **vrrp vrid** *virtual-router-id* **timer advertise** *advertise-interval* command to set the same interval.

- If the authentication modes and authentication keys are different, run the **vrrp vrid** *virtual-router-id* **authentication-mode** { **simple** { **key** | *plain key* | **cipher** *cipher-key* } | **md5** *md5-key* } command to set the same authentication mode and authentication key.

1.5.2 VRRP Group Status Changes Frequently

Fault Description

The VRRP group status changes frequently.

Procedure

Step 1 Run the **display vrrp** *virtual-router-id* command in any view to check whether the VRRP group is associated with an interface, a BFD session, or an NQA test instance.

- If the VRRP group is associated with the interface, BFD session, or NQA test instance, flapping of the interface, BFD session, or NQA test instance causes VRRP group status flapping. Rectify the fault on the associated module.
- If association is not configured, go to step 2.

Step 2 Run the **display vrrp** *virtual-router-id* command in any view to check the preemption delay of the VRRP group.

- If the preemption delay is 0, run the **vrrp vrid** *virtual-router-id* **preempt-mode timer delay** *delay-value* command in the view of the interface where the VRRP group is configured to set the non-0 preemption delay.
- If the preemption is not 0, go to step 3.

Step 3 Run the **vrrp vrid** *virtual-router-id* **timer advertise** *advertise-interval* command in the view of the interface where the VRRP group is configured to set a larger interval at which VRRP Advertisement packets are sent, or run the **vrrp vrid** *virtual-router-id* **preempt-mode timer delay** *delay-value* command to set a larger preemption delay.

1.6 FAQs

1.6.1 Why Cannot the Virtual IP Address of a VRRP Group Be Pinged?

The ping to a virtual IP address may cause potential ICMP attacks. To prevent this problem, run the **undo vrrp virtual-ip ping enable** command to disable the ping function. The virtual IP address of the VRRP group cannot be pinged. To enable the ping to a virtual IP address, run the **vrrp virtual-ip ping enable** command in the system view.

1.6.2 How Can I Adjust the Interval Between Gratuitous ARP Packets Sent from a VRRP Group?

By default, the master sends gratuitous ARP packets every 2 minutes. You can use the **vrrp gratuitous-arp timeout** command to change the interval at which gratuitous ARP packets are sent.

1.6.3 When VRRP Packets Are Sent in a Super-VLAN, Why VRRP Heartbeat Packets Cannot Be Transmitted Normally Between VRRP Master and Backup Devices When the Link Between Them Does Not Allow the First Sub-VLAN of the Super-VLAN?

By default, VRRP Advertisement packets of a super-VLAN are sent only to the sub-VLAN that has the smallest VLAN ID among all the sub-VLANs in Up state. You can use the **vrrp advertise send-mode** command to cancel or configure the mode in which VRRP Advertisement packets of a super-VLAN are sent.

1.6.4 Can RRPP and VRRP Be Used Together on a Switch?

RRPP and VRRP can be configured simultaneously on a switch.

1.7 Terms and Abbreviations

Abbreviation	Full Name
VRRP	Virtual Router Redundancy Protocol
ARP	Address Resolution Protocol
BFD	Bidirectional Forwarding Detection
L2VPN	Layer 2 virtual private network
PW	Pseudo Wire
VSI	Virtual Switching Instance
QinQ	802.1Q in 802.1Q
ME	Metro Ethernet
mVRRP	Manage Virtual Router Redundancy Protocol
mVPLS	Manage Virtual Private LAN Service
mVSI	Manage Virtual Switching Instance