

Implementing VRRP

- Configuring VRRP, on page 1
- Enabling Multiple Group Optimization (MGO) for VRRP, on page 11
- Configuring SNMP Server Notifications for VRRP Events, on page 13

Configuring VRRP

The Virtual Router Redundancy Protocol (VRRP) feature allows for transparent failover at the first-hop IP router, enabling a group of routers to form a single virtual router. For more information on VRRP and related concepts, see Understanding VRRP, on page 1

Restrictions for Configuring VRRP

- ICMP redirects are not supported.
- The maximum VRRP supported is only 16, but this scale number may vary or reduce further based on your BFD, BVI V4 and V6 configuration. For example, if BFD is configured then the value becomes 16-1=15. If BVI v4 and BVI v6 are also configured along with BFD, then the value is 16-3=13 only.

Understanding VRRP

The Virtual Router Redundancy Protocol (VRRP) feature allows for transparent failover at the first-hop IP router, enabling a group of routers to form a single virtual router.

Note VRRP is supported over VRF.

VRRP Overview

A LAN client can use a dynamic process or static configuration to determine which router should be the first hop to a particular remote destination. The client examples of dynamic router discovery are as follows:

- Proxy ARP—The client uses Address Resolution Protocol (ARP) to get the destination it wants to reach, and a router responds to the ARP request with its own MAC address.

- Routing protocol—The client listens to dynamic routing protocol updates (for example, from Routing Information Protocol [RIP]) and forms its own routing table.
- IRDP (ICMP Router Discovery Protocol) client—The client runs an Internet Control Message Protocol (ICMP) router discovery client.

The drawback to dynamic discovery protocols is that they incur some configuration and processing overhead on the LAN client. Also, in the event of a router failure, the process of switching to another router can be slow.

An alternative to dynamic Cisco Discovery Protocols is to statically configure a default router on the client. This approach simplifies client configuration and processing, but creates a single point of failure. If the default gateway fails, the LAN client is limited to communicating only on the local IP network segment and is cut off from the rest of the network.

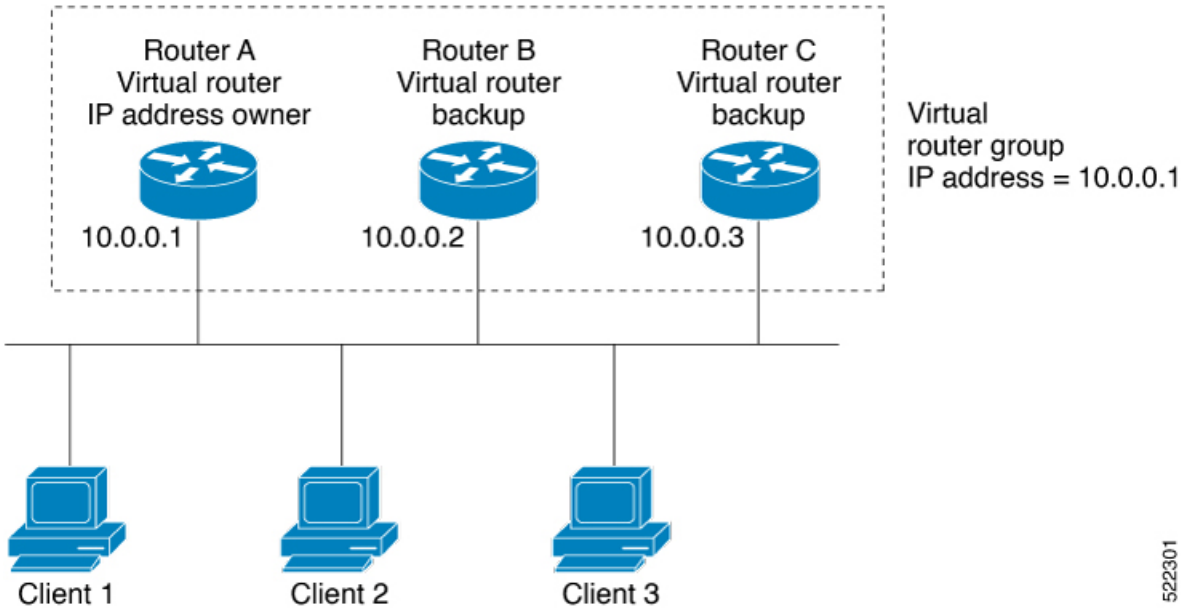
The Virtual Router Redundancy Protocol (VRRP) feature can solve the static configuration problem. VRRP is an IP routing redundancy protocol designed to allow for transparent failover at the first-hop IP router. VRRP enables a group of routers to form a single *virtual router* . The LAN clients can then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a *VRRP group*.

When the virtual router group IP address is the same as the IP address of the physical interface of any router in the VRRP group, then such router becomes the *IP address owner* and the VRRP group operates in the *Owner* mode. When a VRRP group operates in Owner mode, the IP address owner is responsible for forwarding packets that are sent to the VRRP group.

For operating in Owner mode in case of IPv6 VRRP sessions, the link-local address that is configured for the VRRP session must be the same as the link-local address of the physical interface in a router. The link-local address can be autoconfigured by the router or can be an address that is configured by the administrator.

For example, Figure 1: Basic VRRP Topology, on page 3 shows a LAN topology in which VRRP is configured. In this example, Routers A, B, and C are *VRRP routers* (routers running VRRP) that compose a virtual router. The IP address of the virtual router is the same as that configured for the interface of Router A (10.0.0.1).

Figure 1: Basic VRRP Topology



Because the virtual router uses the IP address of the physical interface of Router A, Router A assumes the role of the *IP address owner* and is responsible for forwarding packets that are sent to the VRRP group IP address. Clients 1 through 3 are configured with the default gateway IP address of 10.0.0.1.

Routers B and C function as *backup virtual routers*. If the router that is IP address owner fails, the router that is configured with the higher priority becomes the IP address owner and provides uninterrupted service for the LAN hosts. When Router A recovers, it becomes the IP address owner again.

Note We recommend that you disable Spanning Tree Protocol (STP) on switch ports to which the virtual routers are connected. Enable RSTP or rapid-PVST on the switch interfaces if the switch supports these protocols.

Multiple Virtual Router Support

You can configure up to 255 virtual routers on a router interface. The actual number of virtual routers that a router interface can support depends on the following factors:

- Router processing capability
- Router memory capability
- Router interface support of multiple MAC addresses

In a topology where multiple virtual routers are configured on a router interface, the interface can act as an IP address owner for one or more virtual routers and as a backup for one or more virtual routers.

VRRP Router Priority

An important aspect of the VRRP redundancy scheme is VRRP router priority. Priority determines the role that each VRRP router plays and what happens if the IP address owner virtual router fails.

If a VRRP router owns the IP address of the virtual router and the IP address of the physical interface, this router functions as a IP address owner virtual router.

If no VRRP router owns the IP address, the priority of a VRRP router, combined with the reempt settings, determines if a VRRP router functions as an IP address owner router or a backup virtual router. By default, the highest priority VRRP router functions as IP address owner router, and all the others function as backups. Priority also determines the order of ascendancy to becoming an IP address owner virtual router if the IP address owner virtual router fails. You can configure the priority of each backup virtual router with a value of 1 through 254, using the vrrp priority command.

For example, if Router A, the IP address owner virtual router in a LAN topology, fails, an election process takes place to determine if backup virtual Routers B or C should take over. If Routers B and C are configured with the priorities of 101 and 100, respectively, Router B is elected to become IP address owner virtual router because it has the higher priority. If Routers B and C are both configured with the priority of 100, the backup virtual router with the higher IP address is elected to become the IP address owner virtual router.

By default, a preemptive scheme is enabled whereby a higher-priority backup virtual router that becomes available takes over from the current IP address owner virtual router. You can disable this preemptive scheme using the vrrp preempt disable command. If preemption is disabled, the backup virtual router that is elected to become IP address owner router upon the failure of the original higher priority IP address owner router, remains the IP address owner router even if the original IP address owner virtual router recovers and becomes available again.

VRRP Advertisements

The IP address owner virtual router sends VRRP advertisements to other VRRP routers in the same group. The advertisements communicate the priority and state of the IP address owner virtual router. The VRRP advertisements are encapsulated in IP packets and sent to the IP Version 4 multicast address assigned to the VRRP group. The advertisements are sent every second by default; the interval is configurable.

Benefits of VRRP

The benefits of VRRP are as follows:

- Redundancy— VRRP enables you to configure multiple routers as the default gateway router, which reduces the possibility of a single point of failure in a network.
- Load Sharing—You can configure VRRP in such a way that traffic to and from LAN clients can be shared by multiple routers, thereby sharing the traffic load more equitably among available routers.
- Multiple Virtual Routers—VRRP supports up to 100 virtual routers (VRRP groups) on a router interface, subject to the platform supporting multiple MAC addresses. You can configure up to 256 virtual routers on a router interface. Multiple virtual router support enables you to implement redundancy and load sharing in your LAN topology.
- Multiple IP Addresses—The virtual router can manage multiple IP addresses, including secondary IP addresses. Therefore, if you have multiple subnets configured on an Ethernet interface, you can configure VRRP on each subnet.
- Preemption—The redundancy scheme of VRRP enables you to preempt a backup virtual router that has taken over for a failing IP address owner virtual router with a higher-priority backup virtual router that has become available.
- Text Authentication—You can ensure that VRRP messages received from VRRP routers that comprise a virtual router are authenticated by configuring a simple text password.

- Advertisement Protocol—VRRP uses a dedicated Internet Assigned Numbers Authority (IANA) standard multicast address (224.0.0.18) for VRRP advertisements. This addressing scheme minimizes the number of routers that must service the multicasts and allows test equipment to accurately identify VRRP packets on a segment. The IANA assigns VRRP the IP protocol number 112.

Hot Restartability for VRRP

In the event of failure of a VRRP process in one group, forced failovers in peer VRRP IP address owner router groups should be prevented. Hot restartability supports warm RP failover without incurring forced failovers to peer VRRP routers.

Customizing VRRP

Configuration Example

Customizing the behavior of VRRP is optional. Be aware that as soon as you enable a VRRP group, that group is operating. It is possible that if you first enable a VRRP group before customizing VRRP, the router could take over control of the group and become the master virtual router before you have finished customizing the feature. Therefore, if you plan to customize VRRP, it is a good idea to do so before enabling VRRP.

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2

router(config-vrrp)#delay minimum 2 reload 10
/* (Optional) Delays the startup of the state machine when an interface comes up. */

router(config-vrrp-if)#address-family ipv6
router(config-vrrp-address-family)#vrrp 3
/* The version keyword is available only if IPv4 address-family is selected. */



router(config-vrrp-virtual-router)#accept-mode disable
/* Disables the installation of routes for the VRRP virtual addresses. */

router(config-vrrp-virtual-router)#priority 254
/* (Optional) Sets the priority of the virtual router. */

router(config-vrrp-virtual-router)#preempt delay 15
/* (Optional) Controls which router becomes the master router. */

router(config-vrrp-virtual-router)#timer 4
/* (Optional) Configures the interval between successive advertisements by the master router
 in a VRRP virtual router. */

router(config-vrrp-virtual-router)#track interface TenGigE 0/0/0/2 30
/* (Optional) Configures the VRRP to track an interface. */

router(config-vrrp-virtual-router)#commit
```

Running Configuration

```
Router#show running-config router vrrp
router vrrp
interface TenGigE 0/0/0/2
delay minimum 2 reload 10
address-family ipv6
```

```
vrrp 3
text-authentication
accept-mode disable
priority 254
preempt delay 15
timer 4
track interface TenGigE 0/0/0/2 30
!
```

Verification

```
Router#show vrrp detail

TenGigE0/0/0/2 - IPv4 vrID 3
  State is Master, IP address owner
    1 state changes, last state change 00:01:00
    State change history:
    May 19 12:28:59.825 UTC  Init     -> Master   Virtual IP configured
  Last resign sent:     Never
  Last resign received: Never
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 0000.5E00.0103, state is active
  Master router is local
  Version is 2
  Advertise time 4 secs
    Master Down Timer 12.015 (3 x 4 + (1 x 4/256))
  Minimum delay 2 sec, reload delay 10 sec
  Current priority 255
    Configured priority 254, may preempt
      minimum delay 15 secs
  Authentication enabled, string "text1"
    Tracked items: 1/1 up: 30 decrement
      Object name                State     Decrement
      TenGigE0/0/0/2              Up            30
```

Enabling VRRP

Configuration Example

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv4
router(config-vrrp-address-family)#vrrp 3 version 3
/* The version keyword is available only if IPv4 address-family is selected. */

router(config-vrrp-virtual-router)#address 10.20.30.1
/* Enables VRRP on an interface and specifies the IP address of the virtual router. */

router(config-vrrp-virtual-router)#commit
```

Running Configuration

```
Router#show running-config router vrrp
router vrrp
interface TenGigE 0/0/0/2
address-family ipv4
vrrp 3 version 3
```

```
address 10.20.30.1
!
```

Verification

```
Router#show vrrp detail

TenGigE0/0/0/2 - IPv4 vrID 3
  State is Master, IP address owner
    1 state changes, last state change 00:01:00
    State change history:
    May 19 12:28:59.825 UTC  Init     -> Master   Virtual IP configured
  Last resign sent:     Never
  Last resign received: Never
  Virtual IP address is 10.20.30.1
  Virtual MAC address is 0000.5E00.0103, state is active
  Master router is local
  Version is 2
  Advertise time 4 secs
    Master Down Timer 12.015 (3 x 4 + (1 x 4/256))
  Current priority 255
```

Clearing VRRP Statistics

Clears all the software counters for the specified virtual router.

```
Router#clear vrrp statistics
/* If no interface is specified, statistics of all virtual routers are removed. */
```

Configuring a Global Virtual IPv6 Address

Configuration Example

Configures the global virtual IPv6 address for a virtual router.

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv6
router(config-vrrp-address-family)#vrrp 3
/* The version keyword is available only if IPv4 address-family is selected. */

router(config-vrrp-if-virtual-router)#address global 2001:db8::
router(config-vrrp-virtual-router)#commit
```

Running Configuration

```
Router#show running-config router vrrp
router vrrp
interface TenGigE 0/0/0/2
address-family ipv6
vrrp 3
address global 2001:db8::
!
```

Configuring the Primary and Secondary Virtual IPv4 Addresses

Configuration Example

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv4
router(config-vrrp-address-family)#vrrp 3 version 3
/* The version keyword is available only if IPv4 address-family is selected. */

router(config-vrrp-if-virtual-router)#address 10.20.30.1
/* Configures primary virtual IPv4 address for a virtual router. */

router(config-vrrp-if-virtual-router)#address 10.20.30.2 secondary
/* Configures secondary virtual IPv4 address for a virtual router. */

router(config-vrrp-virtual-router)#commit
```

Running Configuration

```
Router#show running-config router vrrp
router vrrp
interface TenGigE 0/0/0/2
address-family ipv4
vrrp 3 version 3
address 10.20.30.1
address 10.20.30.2 secondary
!
```

Verification

```
Router#show vrrp detail

TenGigE0/0/0/2 - IPv4 vrID 3
  State is Master, IP address owner
    1 state changes, last state change 00:01:00
    State change history:
    May 19 12:28:59.825 UTC  Init     -> Master   Virtual IP configured
  Last resign sent:     Never
  Last resign received: Never
  Virtual IP address is 10.20.30.1
  Virtual MAC address is 0000.5E00.0103, state is active
  Master router is local
  Virtual secondary IP address is 10.20.30.2
  Version is 2
  Advertise time 4 secs
    Master Down Timer 12.015 (3 x 4 + (1 x 4/256))
  Current priority 255
```

Configuring a Virtual Link-Local IPv6 Address

Configures either the virtual link-local IPv6 address for a virtual router or specifies that the virtual link-local IPv6 address be enabled and calculated automatically from the virtual router virtual MAC address.

The IPv6 address space is structured differently compared to IPv4. Link-local addresses are used to identify each interface on the local network. These addresses may either be configured or determined automatically in a standard way using the link-layer (hardware) address of the interface (MAC address for Ethernet interfaces). Link-local addresses have a standard format and are valid only on the local network (they cannot be routed to, from multiple hops away).

Global unicast IPv6 addresses occupy a disjoint subset of the IPv6 address space from link-local addresses. They can be routed to, from multiple hops away and have an associated prefix length (between 0 and 128 bits).

Each VRRP virtual router has an associated virtual link-local address. This may be configured or determined automatically from the virtual router's virtual MAC address. The virtual MAC address must be unique on the local network. The virtual link-local address is analogous to an IPv4 virtual router's primary virtual IPv4 address, except that its virtual IP (VIP) state is always considered to be up, since duplicate address detection is not required for addresses whose scope is local.

Configuration Example

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv6

/* Use one of the following address linklocal commands: */
router(config-vrrp-address-family)#vrrp 1 address linklocal FE80::260:3EFF:FE11:6770
/* Configures the virtual link-local IPv6 address for the virtual router. */

router(config-vrrp-address-family)#vrrp 1 address linklocal autoconfigure
/* Specifies that the virtual link-local IPv6 address should be enabled and calculated
automatically
from the virtual router virtual MAC address. */

router(config-vrrp-virtual-router)#commit
```

Running Configuration

```
Router#show running-config router vrrp
router vrrp
interface TenGigE 0/0/0/2
address-family ipv6
vrrp 1 address linklocal FE80::260:3EFF:FE11:6770
!
```

BFD for VRRP

Bidirectional Forwarding Detection (BFD) is a network protocol used to detect faults between two forwarding engines. BFD sessions operate in asynchronous mode. In asynchronous mode, both endpoints periodically send hello packets to each other. If a number of those packets are not received, the session is considered down.

Advantages of BFD

- BFD provides failure detection in less than one second.
- BFD supports all types of encapsulation.
- BFD is not tied to any particular routing protocol, supports almost all routing protocols.

BFD Process

VRRP uses BFD to detect a link failure and facilitate fast failover times without excessive control packet overhead.

The VRRP process creates BFD sessions as required. When a BFD session goes down, each backup group monitoring the session transitions to the active state.

After a transition to active state triggered by a BFD session going down, VRRP does not participate in any state elections for 10 seconds.



Note IPv4 only supports BFD for VRRP.

Configuring BFD

Enabling BFD

```
Router# configure
Router(config)#  router vrrp
Router(config-vrrp)# interface <type> <interface-path-id>
Router(config-vrrp-if)# address-family ipv4
Router(config-vrrp-ipv4)# vrrp[group number] version <version-no> bfd
fast-detect [peer ipv4 <ipv4-address> <interface-type> <interface-path-id>]
commit
```

Verifying BFD on VRRP

```
router vrrp
interface TenGigE0/0/0/3.1
bfd minimum-interval 4
bfd multiplier 3
address-family ipv4
vrrp 1
priority 200
address 41.41.1.3
bfd fast-detect peer ipv4 41.41.1.2
```

Modifying BFD timers (minimum interval)

Minimum interval determines the frequency of sending BFD packets to BFD peers (in milliseconds). The default minimum interval is 15ms.

```
Router# configure
```

```
Router(config)#  router vrrp
Router(config-vrrp)# interface <type> <interface-path-id>
Router(config-vrrp-if)# bfd minimum-interval <interval>
Router(config-vrrp-if)# bfd multiplier <multiplier>
router(config-vrrp-if)# address-family ipv4
commit
```

Modifying BFD timers (multiplier)

Multiplier is the number of consecutive BFD packets which must be missed from a BFD peer before declaring that peer unavailable. The default multiplier is 3.

```
Router# configure
Router(config)#  router vrrp
Router(config-vrrp)# interface <type> <interface-path-id>
Router(config-vrrp-if)#  bfd multiplier <multiplier>
router(config-vrrp-if)# address-family ipv4
commit
```

Disabling State Change Logging

Configuration Example

Disables the task of logging the VRRP state change events via syslog.

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#message state disable
router(config-vrrp)#commit
```

Enabling Multiple Group Optimization (MGO) for VRRP

Configuration Examples

Multiple Group Optimization for Virtual Router Redundancy Protocol (VRRP) provides a solution for reducing control traffic in a deployment consisting of many subinterfaces. By running the VRRP control traffic for just one session, the control traffic is reduced for the subinterfaces with identical redundancy requirements. All other sessions are subordinates of this primary session, and inherit their states from it.

Configuring VRRP Session Name

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv4
router(config-vrrp-address-family)#vrrp 1
/* Enables VRRP group configuration mode on a specific interface. */
```

```
router(config-vrrp-vritual-router)#name m1
/* Specifies the VRRP session name. */

router(config-vrrp-gp)#commit
```

Configuring the Subordinate Group to Inherit its State from a Specified Group

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv4

router(config-vrrp-address-family)#vrrp 2 slave
/* Enables VRRP slave configuration mode on a specific interface. */

router(config-vrrp-slave)#follow m1
/* Instructs the subordinate group to inherit its state from the specified group, m1 (MGO
session name). */

router(config-vrrp-slave)#address 10.2.3.2
/* Specifies the primary virtual IPv4 address for subordinate group. */

router(config-vrrp-slave)#address 10.2.3.3 secondary
/* Specifies the secondary virtual IPv4 address for subordinate group. */

router(config-vrrp-gp)#commit
```

Primary and Secondary Virtual IPv4 Addresses for the Subordinate Group

```
Router#configure
Router(config)#router vrrp
router(config-vrrp)#interface TenGigE 0/0/0/2
router(config-vrrp-if)#address-family ipv4

router(config-vrrp-address-family)#vrrp 2 slave
/* Enables VRRP slave configuration mode on a specific interface. */

router(config-vrrp-slave)#address 10.2.3.2
/* Specifies the primary virtual IPv4 address for subordinate group. */

router(config-vrrp-slave)#address 10.2.3.3 secondary
/* Specifies the secondary virtual IPv4 address for subordinate group. */

router(config-vrrp-slave)#commit
```

Running Configuration

```
Router#show running-config router vrrp 1
router vrrp
interface TenGigE 0/0/0/2
address-family ipv4
vrrp 1
name m1
!

/* Subordinate group */
Router#show running-config router vrrp 2
router vrrp
interface TenGigE 0/0/0/2
address-family ipv4
vrrp 2 slave
follow m1
address 10.2.3.2
address 10.2.3.3 secondary
```

!

Configuring SNMP Server Notifications for VRRP Events

MIB support for VRRP

VRRP enables one or more IP addresses to be assumed by a router when a failure occurs. For example, when IP traffic from a host reaches a failed router because the failed router is the default gateway, the traffic is transparently forwarded by the VRRP router that has assumed control. VRRP does not require configuration of dynamic routing or router discovery protocols on every end host. The VRRP router controlling the IP address(es) associated with a virtual router is called the IP address owner router, and forwards packets sent to these IP addresses. The election process provides dynamic fail over (standby) in the forwarding responsibility should the IP address owner router become unavailable. This allows any of the virtual router IP addresses on the LAN to be used as the default first hop router by end-hosts.

The advantage gained from using VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host. Simple Network Management Protocol (SNMP) traps provide information of the state changes, when the virtual routers (in standby) are moved to IP address owner router's state or if the standby router is made IP address owner router.

Configuration Example

Enables SNMP server notifications (traps) for VRRP.

```
Router#configure
Router(config)#snmp-server traps vrrp events
router(config)#commit
```

Use the **show snmp traps details** command to view details of SNMP server notifications.