Implementing VRRP

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

- If you configure the command hw-module vrrpscale enable on the router, upto 64 (IPv4 and IPv6 combined) VRRP groups are supported on Cisco NCS 5500 Series routers and upto 16 (IPv4 and IPv6 combined) VRRP groups are supported on Cisco NCS 540 Series routers. By default, 16 VRRP groups (IPv4 and IPv6 combined) are supported on Cisco NCS 5500 Series Routers and Cisco NCS 540 Series routers.

- If you do not configure the command hw-module vrrpscale enable on the router and if you configure BFD along with BVI and if all the BVIs are sharing the same Chassis (default) MAC, the VRRP scale is reduced to 13. You cannot use any custom BVI MAC in this mode until VRRP scale is reduced to 11.

- If you configure BFD along with BVI and all the BVIs are configured with one custom MAC or mix of one custom MAC and Chassis (default) MAC, the VRRP scale is reduced to 11.

- ICMP redirects are not supported.

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 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*.

For example, 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
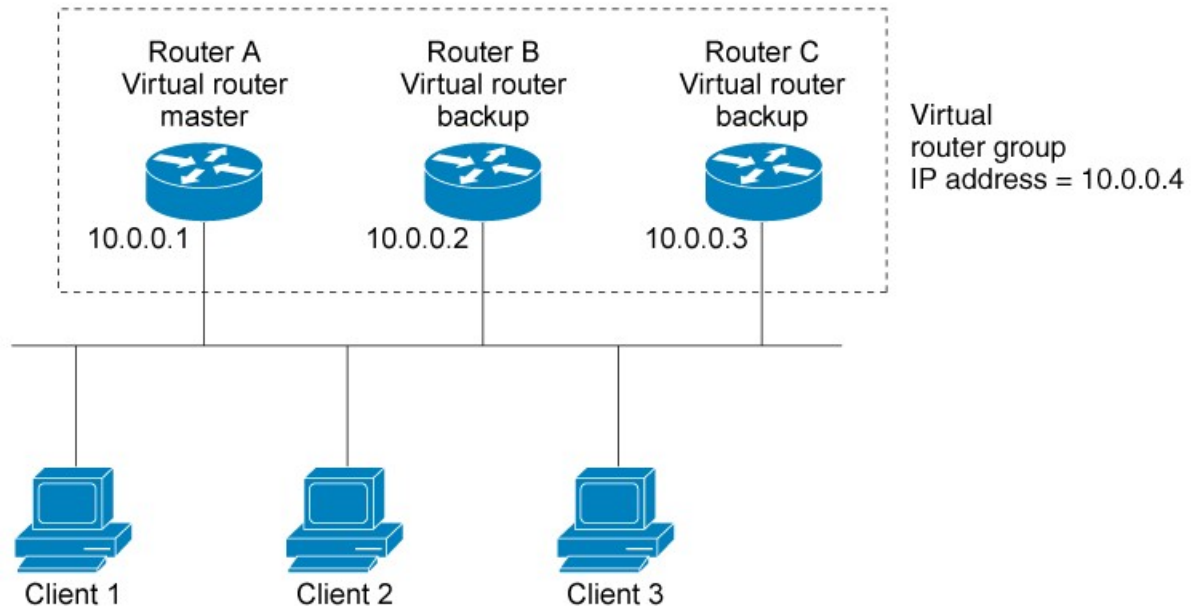
Because the virtual router uses the IP address of the physical interface of Router A, Router A assumes the role of the *master virtual router* and is also known as the *IP address owner.* As the master virtual router, Router A controls the IP address of the virtual router and is responsible for forwarding packets sent to this 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 master virtual router fails, the router configured with the higher priority becomes the master virtual router and provides uninterrupted service for the LAN hosts. When Router A recovers, it becomes the master virtual router 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 100 virtual routers on a router interface. You can configure up to 256 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 a master 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 master 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 master 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 a master or a backup virtual router. By default, the highest priority VRRP router functions as master, and all the others function as backups. Priority also determines the order of ascendancy to becoming a master virtual router if the master 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 master 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 master 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 master virtual router.

By default, a preemptive scheme is enabled whereby a higher-priority backup virtual router that becomes available takes over from the current master 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 master upon the failure of the original higher priority master, remains the master even if the original master virtual router recovers and becomes available again.

VRRP Advertisements

The master virtual router sends VRRP advertisements to other VRRP routers in the same group. The advertisements communicate the priority and state of the master 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 master 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 master router groups should be prevented. Hot restartability supports warm RP failover without incurring forced failovers to peer VRRP routers.

Understanding VRRP over BVI

The Virtual Router Redundancy Protocol (VRRP) protocol provides default gateway redundancy. It allows a group of routers to behave as a single virtual default gateway router in which one router acts as the Master router and others routers act as Backup routers.

BVI (Bridge-Group Virtual Interface) is a virtual interface which provides L3 or routed functionality to a Bridge Group. L2 functionality is applicable to the interfaces which are part of a Bridge Group and BVI is the routed interface for that Bridge Group.

Usually, VRRP sessions run on top of interfaces of the multiple routers which are in the same home network. However, you can configure VRRP session over BVI. Thereby, instead of physical interfaces, VRRP sessions can run between BVI interfaces of multiple routers.

Configuring VRRP for IPv4 Networks

This section describes the procedure for configuring and verifying VRRP for IPv4 networks.

Configuration

Use the following configuration for configuring VRRP for IPv4 networks.

Note Certain customizations (as mentioned) are recommended to control the behavior of the VRRP group on committing the VRRP configuration on the Router. If the following customizations are not configured, then the Router seizes control of the VRRP group, and immediately assumes the role of the master virtual Router.

```
/* Enter the interface configuration mode and configure an IPv4 address for the interface.
 */
Router(config)# interface gigabitEthernet 0/0/0/1
Router(config-if)# ipv4 address 10.10.10.1 255.255.255.0
Router(config-if)# no shut
Router(config-if)# commit
Fri Dec  8 13:49:24.142 IST
Router:Dec  8 13:49:24.285 : ifmgr[402]: %PKT_INFRA-LINK-3-UPDOWN : Interface
GigabitEthernet0/0/0/1, changed state to Down
Router:Dec  8 13:49:24.711 : ifmgr[402]: %PKT_INFRA-LINK-3-UPDOWN : Interface
GigabitEthernet0/0/0/1, changed state to Up
```

```
Router(config-if)# exit
Router(config)# do show ip int brief
Fri Dec  8 13:50:05.505 IST

Interface                      IP-Address      Status          Protocol Vrf-Name
GigabitEthernet0/0/0/0         unassigned      Shutdown        Down     default
GigabitEthernet0/0/0/1         10.10.10.1      Up              Up       default
GigabitEthernet0/0/0/2         unassigned      Shutdown        Down     default
GigabitEthernet0/0/0/3         unassigned      Shutdown        Down     default
GigabitEthernet0/0/0/4         unassigned      Shutdown        Down     default


/* Enter the VRRP configuration mode and add the configured interface. */
Router(config)# router vrrp
Router(config-vrrp)# interface GigabitEthernet 0/0/0/1

/* CUSTOMIZATION: Configure a delay for the startup of the state machine when the interface
 comes up. */
Router(config-vrrp)# delay minimum 2 reload 10 */

/* Configure VRRP version 3 for IPv4 */
Router(config-vrrp-if)# address-family ipv4 vrrp 100 version 3
Router(config-vrrp-virtual-router)# address 10.10.10.1

/* CUSTOMIZATION: Disable the installation of routes for the VRRP virtual addresses. */
Router(config-vrrp-virtual-Router)# accept-mode disable

/* CUSTOMIZATION: Set a priority for the virtual Router. */
Router(config-vrrp-virtual-Router)# priority 254

/* CUSTOMIZATION: Configure a preempt delay value that controls the selection of the master
 virtual Router. */
Router(config-vrrp-virtual-Router)# preempt delay 15

/* CUSTOMIZATION: Configure the interval between successive advertisements by the master
virtual Router. */
Router(config-vrrp-virtual-Router)#timer 4

/* CUSTOMIZATION: Configure VRRP to track an interface. */
Router(config-vrrp-virtual-Router)# track interface GigabitEthernet0/0/0/1 30

 /* Commit the configuration */
Router(config-vrrp-virtual-Router)# commit
```
_____

You have successfully configured VRRP for IPv4 networks.

Validation

Use the following commands to validate the configuration.

```
/* Validate the configuration */
Router(config-vrrp-virtual-router)# do show run interface GigabitEthernet 0/0/0/1
Fri Dec  8 15:04:38.140 IST
interface GigabitEthernet0/0/0/1
 ipv4 address 10.10.10.1 255.255.255.0
!
```
_____

```
Router(config)# show running-config router vrrp
Fri Dec  8 13:50:18.959 IST
router vrrp
 interface GigabitEthernet0/0/0/1
```

```
   delay minimum 2 reload 10
  address-family ipv4
   vrrp 100 version 3
    priority 254
    preempt delay 15
    timer 4
    track interface GigabitEthernet0/0/0/2 30
    address 10.10.10.1
    accept-mode disable
   !
  !
 !
```

```
Router(config-vrrp-virtual-router)# do show vrrp ipv4 interface gigabitEthernet 0/0/0/1
Fri Dec  8 15:02:56.952 IST
IPv4 Virtual Routers:
                   A indicates IP address owner
                   | P indicates configured to preempt
                   | |
Interface   vrID Prio A P State    Master addr    VRouter addr
Gi0/0/0/1    100  255 A P Master   local          10.10.10.1
```

```
Router(config-vrrp-virtual-router)# end
Router# show vrrp detail
Fri Dec  8 15:08:36.469 IST
GigabitEthernet0/0/0/1 - IPv4 vrID 100
  State is Master, IP address owner
    1 state changes, last state change 01:19:06
    State change history:
    Dec  8 13:49:30.147 IST  Init     -> Master   Delay timer expired
  Last resign sent:     Never
  Last resign received: Never
  Virtual IP address is 10.10.10.1
  Virtual MAC address is 0000.5E00.0164, state is active
  Master router is local
  Version is 3
  Advertise time 1 secs
    Master Down Timer 3.003 (3 x 1 + (1 x 1/256))
  Minimum delay 1 sec, reload delay 5 sec
  Current priority 255
    Configured priority 100, may preempt
      minimum delay 0 secs
```

You have successfully validated VRRP for IPv4 networks.

Configuring VRRP for IPv6 Networks

This section describes the procedure for configuring and verifying VRRP for IPv6 networks.

Configuration

The following sample includes the configuration and customization of VRRP for IPv6 networks.

Note Certain customizations (as mentioned) are recommended to control the behavior of the VRRP group on committing the VRRP configuration on the Router. If the following customizations are not configured, then the Router seizes control of the VRRP group, and immediately assumes the role of the master virtual Router.

```
/* Enter the interface configuration mode and configure an IPv6 address */
Router# interface GigabitEthernet 0/0/0/2
Router(config-if)# ipv6 address 10::1/64
Router(config-if)# no shut

/* Exit the interface configuration mode and enter the vrrp configuration mode */
Router(config-if)# exit
Router(config)# Router vrrp

/* Add the configured interface for VRRP */
Router(config-vrrp)# interface GigabitEthernet 0/0/0/2

/* CUSTOMIZATION: Configure a delay for the startup of the state machine when the interface
 comes up. */
Router(config-vrrp)# delay minimum 2 reload 10 */

/* Enable the IPv6 global and link local address family on the interface  */
Router(config-vrrp-if)# address-family ipv6 vrrp 50
Router(config-vrrp-virtual-Router)# address linklocal autoconfig

/* CUSTOMIZATION: Disable the installation of routes for the VRRP virtual addresses. */
Router(config-vrrp-virtual-Router)# accept-mode disable

/* CUSTOMIZATION: Set a priority for the virtual Router. */
Router(config-vrrp-virtual-Router)# priority 254

/* CUSTOMIZATION: Configure a preempt delay value that controls the selection of the master
 virtual Router. */
Router(config-vrrp-virtual-Router)# preempt delay 15

/* CUSTOMIZATION: Configure the interval between successive advertisements by the master
virtual Router. */
Router(config-vrrp-virtual-Router)#timer 4

/* CUSTOMIZATION: Configure VRRP to track an interface. */
Router(config-vrrp-virtual-Router)# track interface GigabitEthernet0/0/0/2 30

/* Commit the configuration */
Router(config-vrrp-virtual-Router)# commit
```

You have successfully configured VRRP for IPv6 networks.

Validation

Use the following commands to validate the configuration.

```
/* Validate the configuration */
Router(config-vrrp-virtual-router)# do show run interface GigabitEthernet 0/0/0/2
Fri Dec  8 14:55:48.378 IST
interface GigabitEthernet0/0/0/2
 ipv6 address 10::1/64
!
----------------------------------------------------------------------------------------------------------
Router(config-vrrp-virtual-router)# do show running-config router vrrp
...
router vrrp
 interface GigabitEthernet0/0/0/2
  delay minimum 2 reload 10
  address-family ipv6
   vrrp 50
    priority 254
```

```
      preempt delay 15
      timer 4
      track interface GigabitEthernet0/0/0/2 30
      address linklocal autoconfig
      accept-mode disable
     !
    !
   !
!
```

```
Router(config-vrrp-virtual-router)# do show vrrp ipv6 interface gigabitEthernet 0/0/0/2
Fri Dec  8 14:59:25.547 IST
IPv6 Virtual Routers:
                      A indicates IP address owner
                      | P indicates configured to preempt
                      | |
Interface    vrID Prio A P State    Master addr     VRouter addr
Gi0/0/0/2      50  254   P Master   local
                                              fe80::200:5eff:fe00:203
```

```
Router(config-vrrp-virtual-router)# end
Router# show vrrp detail
Fri Dec  8 15:08:36.469 IST
GigabitEthernet0/0/0/2 - IPv6 vrID 50
  State is Master
    2 state changes, last state change 00:18:01
    State change history:
    Dec  8 14:50:23.326 IST  Init     -> Backup  Virtual IP configured
    Dec  8 14:50:35.365 IST  Backup   -> Master  Master down timer expired
  Last resign sent:     Never
  Last resign received: Never
  Virtual IP address is fe80::200:5eff:fe00:203
  Virtual MAC address is 0000.5E00.0203, state is active
  Master router is local

  Advertise time 4 secs
    Master Down Timer 12.031 (3 x 4 + (2 x 4/256))
  Minimum delay 2 sec, reload delay 10 sec
  Current priority 254
    Configured priority 254, may preempt
      minimum delay 15 secs
    Tracked items: 1/1 up: 0 decrement
      Object name              State     Decrement
      GigabitEthernet0/0/0/2     Up          30
```

You have successfully validated VRRP for IPv6 networks.

Configure VRRP over BVI

To configure VRRP sessions over BVI, you must complete the following configurations:

1. Configure a set of interfaces as L2 interfaces.

2. Configure a bridge group.

3. Configure a BVI.

4. Configure VRRP over BVI.

Configuration Example

```
/* Enter the global configuration mode and Configure a set of interfaces as L2 interfaces
*/
Router# configure
Router(config)# interface HundredGigE0/0/0/0.1 l2transport
Router(config-subif)# exit
Router(config)# interface HundredGigE0/0/0/1.1 l2transport
Router(config-subif)# commit
Router(config-subif)# exit

/* Enter the Layer 2 VPN configuration mode and Configure a bridge group */
Router(config)# l2vpn
Router(config-l2vpn)# bridge group 5
Router(config-l2vpn-bg)# bridge-domain 5
Router(config-l2vpn-bg-bd)# interface HundredGigE 0/0/0/0.1
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# interface HundredGigE 0/0/0/1.1
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# routed interface BVI 10
Router(config-l2vpn-bg-bd-bvi)# commit
Router(config-l2vpn-bg-bd-bvi)# exit

/* Configure a BVI in the global configuration mode*/
Router(config-l2vpn-bg-bd)# interface BVI 10

Router(config-if)# ipv4 address 209.165.200.225 255.255.255.0
Router(config-if)# ipv6 address 2001:DB8:A:B::1/64
Router(config-if)# commit

/* Configure VRRP over BVI in the global configuration mode for IPv4 address*/
Router(config)# router VRRP
Router(config-vrrp)# interface BVI 10
Router(config-vrrp-if)# address-family ipv4
Router(config-vrrp-address-family)# VRRP 10
Router config-vrrp-virtual-router)# priority 101
Router config-vrrp-virtual-router)# 209.165.200.226
Router (config-vrrp-virtual-router)# commit


/* Configure VRRP over BVI in the global configuration mode for IPv6 address*/
Router(config)# router VRRP
Router(config-vrrp)# interface BVI 10
Router(config-vrrp-if)# address-family ipv6
Router(config-vrrp-address-family)# VRRP 11
Router config-vrrp-virtual-router)# address global 2001:DB8:A:B::2/64
Router config-vrrp-virtual-router)#  address linklocal autoconfig
Router (config-vrrp-virtual-router)# commit
```

Verification

Use the following command to verify the bridge domain details:

```
Router# show l2vpn bridge-domain detail

Legend: pp = Partially Programmed.
Bridge group: 10, bridge-domain: 10, id: 0, state: up, ShgId: 0, MSTi: 0
  Coupled state: disabled
  VINE state: BVI Resolved
  MAC learning: enabled
  MAC withdraw: enabled
    MAC withdraw for Access PW: enabled
```

```
      MAC withdraw sent on: bridge port up
      MAC withdraw relaying (access to access): disabled
   Flooding:
      Broadcast & Multicast: enabled
      Unknown unicast: enabled
   MAC aging time: 300 s, Type: inactivity
   MAC limit: 64000, Action: none, Notification: syslog
   MAC limit reached: no, threshold: 75%
   MAC port down flush: enabled
   MAC Secure: disabled, Logging: disabled
   Split Horizon Group: none
   Dynamic ARP Inspection: disabled, Logging: disabled
   IP Source Guard: disabled, Logging: disabled
   DHCPv4 Snooping: disabled
   DHCPv4 Snooping profile: none
   IGMP Snooping: disabled
   IGMP Snooping profile: none
   MLD Snooping profile: none
   Storm Control: disabled
   Bridge MTU: 1500
   MIB cvplsConfigIndex: 1
   Filter MAC addresses:
   P2MP PW: disabled
   Create time: 25/07/2018 19:04:41 (3w5d ago)
   No status change since creation
   ACs: 2 (2 up), VFIs: 0, PWs: 0 (0 up), PBBs: 0 (0 up), VNIs: 0 (0 up)
   List of ACs:
     AC: BVI10, state is up
       Type Routed-Interface
       MTU 1514; XC ID 0x80000002; interworking none
       BVI MAC address:
         008a.9651.e0e0
       Split Horizon Group: Access
     AC: Bundle-Ether10.1, state is up
       Type VLAN; Num Ranges: 1
       Rewrite Tags: []
       VLAN ranges: [10, 10]
       MTU 1500; XC ID 0xa0000001; interworking none
       MAC learning: enabled
       Flooding:
         Broadcast & Multicast: enabled
         Unknown unicast: enabled
       MAC aging time: 300 s, Type: inactivity
       MAC limit: 64000, Action: none, Notification: syslog
       MAC limit reached: no, threshold: 75%
       MAC port down flush: enabled
       MAC Secure: disabled, Logging: disabled
       Split Horizon Group: none
       Dynamic ARP Inspection: disabled, Logging: disabled
       IP Source Guard: disabled, Logging: disabled
       DHCPv4 Snooping: disabled
       DHCPv4 Snooping profile: none
       IGMP Snooping: disabled
       IGMP Snooping profile: none
       MLD Snooping profile: none
       Storm Control: bridge-domain policer
       Static MAC addresses:
       Statistics:
         packets: received 0 (multicast 0, broadcast 0, unknown unicast 0, unicast 0), sent
4429828
         bytes: received 0 (multicast 0, broadcast 0, unknown unicast 0, unicast 0), sent
344854904
         MAC move: 0
       Storm control drop counters:
```

```
        packets: broadcast 0, multicast 0, unknown unicast 0
        bytes: broadcast 0, multicast 0, unknown unicast 0
      Dynamic ARP inspection drop counters:
        packets: 0, bytes: 0
      IP source guard drop counters:
        packets: 0, bytes: 0
  List of Access PWs:
  List of VFIs:
  List of Access VFIs:
```

Use the following command to show the VRRP details:

```
Router# show vrrp ipv4 detail

BVI10 - IPv4 vrID 10
  State is Master
    63 state changes, last state change 5d20h
    State change history:
    Aug 15 20:30:47.986 UTC  Backup   -> Init     Interface Down update
    Aug 15 20:30:50.642 UTC  Init     -> Backup   Delay timer expired
    Aug 15 20:30:51.304 UTC  Backup   -> Init     Interface Down update
    Aug 15 20:31:44.957 UTC  Init     -> Backup   Delay timer expired
    Aug 15 20:31:48.562 UTC  Backup   -> Master   Master down timer expired
  Last resign sent:     Aug 15 20:30:20.700 UTC
  Last resign received: Jul 25 19:04:21.466 UTC
  Virtual IP address is 209.165.200.226
  Virtual MAC address is 0000.5E00.010a, state is active
  Master router is local
  Version is 2
  Advertise time 1 secs
    Master Down Timer 3.605 (3 x 1 + (155 x 1/256))
  Minimum delay 1 sec, reload delay 5 sec
  Current priority 101
    Configured priority 101, may preempt
      minimum delay 0 secs


Router# show vrrp ipv6 detail

BVI10 - IPv6 vrID 11
  State is Master
    63 state changes, last state change 5d20h
    State change history:
    Aug 15 20:30:48.032 UTC  Backup   -> Init     Interface Down update
    Aug 15 20:30:50.517 UTC  Init     -> Backup   Delay timer expired
    Aug 15 20:30:51.348 UTC  Backup   -> Init     Interface Down update
    Aug 15 20:31:44.996 UTC  Init     -> Backup   Delay timer expired
    Aug 15 20:31:48.605 UTC  Backup   -> Master   Master down timer expired
  Last resign sent:     Aug 15 20:30:20.702 UTC
  Last resign received: Never
  Virtual IP address is fe80::200:5eff:fe00:20b
    Secondary Virtual IP address is 2001:DB8:A:B::2/64
  Virtual MAC address is 0000.5E00.020b, state is active
  Master router is local
  Version is 3
  Advertise time 1 secs
    Master Down Timer 3.609 (3 x 1 + (156 x 1/256))
  Minimum delay 1 sec, reload delay 5 sec
  Current priority 100
    Configured priority 100, may preempt
      minimum delay 0 secs
```

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 slaves of this primary session, and inherit their states from it.

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 s1
/* Specifies the VRRP session name. */

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

Slave Follow

```
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
/* Configures a slave follow. Instructs the slave 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 slave group. */

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

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

Primary and Secondary Virtual IPv4 Addresses for the Slave 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 slave group. */

router(config-vrrp-slave)#address 10.2.3.3 secondary
/* Specifies the secondary virtual IPv4 address for slave 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 s1
!

/* Slave 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 master, and forwards packets sent to these IP addresses. The election process provides dynamic fail over (standby) in the forwarding responsibility should the master 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 master state or if the standby router is made master.

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.