



Implement VRRP

This module describes the concepts and tasks you will use to configure Virtual Router Redundancy Protocol (VRRP).

- [Implement VRRP, on page 1](#)
- [Generic Restrictions for VRRP Configuration, on page 2](#)
- [Understand VRRP, on page 2](#)
- [VRRP over Physical Interfaces and Bundle Interfaces, on page 5](#)
- [Unicast VRRP, on page 13](#)
- [VRRP over BVI, on page 14](#)
- [View VRRP statistics in the Router, on page 19](#)

Implement VRRP

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Virtual Router Redundancy Protocol	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100]) This feature is supported on: <ul style="list-style-type: none">• 8011-4G24Y4H-I
Virtual Router Redundancy Protocol	Release 24.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) You can now improve network reliability by enabling the Virtual Router Redundancy Protocol (VRRP) to automatically assign available IP routers to participating hosts, which enhances reliability. This improvement reduces downtime by allowing a backup router to take over if the primary router fails, ensuring continuous network availability. VRRP also supports load balancing, distributing traffic across multiple routers for better efficiency. It integrates seamlessly with existing network configurations, offering a streamlined setup process without compromising security or performance. This feature is especially beneficial for critical applications that require high availability. *This feature support is now extended to Cisco 8712-MOD-M routers.

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 [Understand VRRP, on page 2](#).

Generic Restrictions for VRRP Configuration

These are some restrictions to consider before you configure VRRP on supported interfaces on the Cisco 8000 platform.

- ICMP redirects are not supported.
- Protocol Independent Multicast (PIM) is not supported with VRRP.
- The Cisco 8000 Series Routers do not support the configuration of Virtual IP (VIP) addresses for VRRP on subinterfaces. As an alternative, Bridge-Group Virtual Interfaces (BVI) can be utilized for VRRP implementation on subinterfaces.

Understand VRRP

To implement VRRP on Cisco IOS XR software, you need to understand the following concepts:

- [VRRP Overview](#)
- [Multiple Virtual Router Support](#)
- [VRRP Router Priority](#)
- [VRRP Advertisements](#)
- [Benefits of VRRP](#)
- [Hot Restartability for VRRP](#)

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, [Basic VRRP Topology](#) 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 preempt 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.

VRRP over Physical Interfaces and Bundle Interfaces

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
--------------	---------------------	---------------------

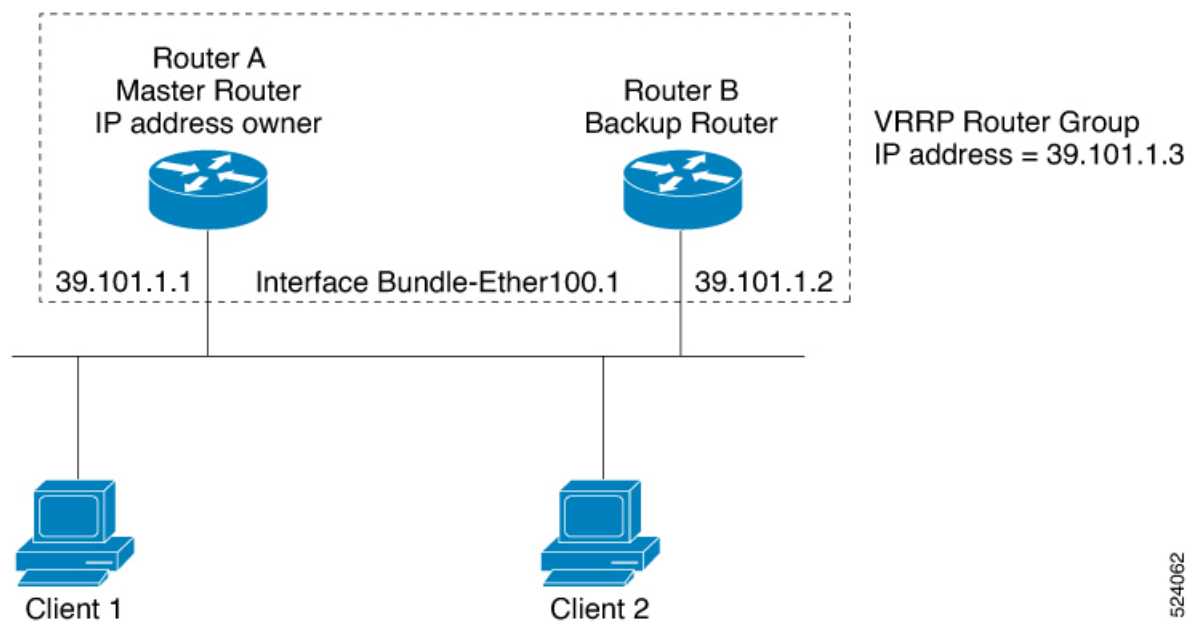
VRRP over Physical interfaces and Bundle interfaces	Release 24.2.11	<p>This feature ensures high availability of routing paths by mitigating any failure in the primary interfaces within a group of physical or bundle interfaces or sub-interfaces in a network with a failover to a backup physical or bundle interface.</p> <p>The feature achieves this failover through a backup router in the VRRP router group configured on the physical or bundle interfaces or sub-interfaces. The virtual IP addresses from the failed primary router are handed over to this backup router.</p>
---	-----------------	--

You can configure multiple routers from a VRRP router group and use one of these routers as the default gateway router to reduce the possibility of a single point of network failure. It also ensures that the traffic load received from and sent to LAN clients is equally shared between the available routers for optimal performance.

You can configure up to 510 IPv4 and IPv6 VRRP groups on the Cisco Silicon One Q200 and P100 ASIC-based systems.

Topology

This topology provides the VRRP implementation over a Bundle sub-interface on a network segment.



VRRP topology and underlying network elements:

- **Interface Bundle-Ether100.1:** This is the Bundle sub-interface which is configured on the Master router (IP Address: 39.101.1.1) and the Backup router (IP Address: 39.101.1.2).

- **VRRP Router Group:** The VRRP router group consists of two or more routers configured with VRRP. This group provides backup services to the backup router (Router B) in the event of the Master router (Router A) failure. The virtual IP address of the VRRP router group in this topology is 39.101.1.3.
- **Router A:** Router A is the Master router on which the Bundle sub-interface **Bundle-Ether100.1** is configured. This router receives and routes packets destined for the MAC address of the group. This router is known as the IP address owner of the VRRP router group.
- **Router B:** When the designated Master router fails, VRRP detects it and selects another router with the highest priority in the VRRP router group as the Backup router. This Backup router with the Bundle sub-interface **Bundle-Ether100.1** configuration takes over the job of Router A and assumes control of the MAC and IP Addresses of the VRRP router group.
- **Client 1 and Client 2:** These are source hosts from which traffic is forwarded through the Master router or Backup router to the destination hosts.

Usage Guidelines for VRRP Configuration

When setting up VRRP for both IPv4 and IPv6 groups with more than 510 sessions, or when using aggressive timer intervals below 1 second, it is crucial to adjust the LPTS flow parameters accordingly. Specifically, increase the LPTS policer rate to a minimum of 3000 pps. This ensures optimal performance and avoids any LPTS drops for VRRP traffic to maintain network stability and efficiency. For detailed instructions on configuring these settings, refer to the *Usage Guidelines* section within the [lpts pifib hardware police](#) command documentation.

Configure VRRP for IPv4 Networks on Physical Interfaces and Bundle Interfaces

This section describes the procedure for configuring and verifying VRRP for IPv4 networks on Physical interfaces, Physical sub-interfaces, Bundle interfaces, and Bundle sub-interfaces.

Configuration

Use the following configuration for configuring VRRP for IPv4 networks on Physical interfaces, Physical sub-interfaces, Bundle interfaces, and Bundle sub-interfaces.



Note It is recommended to apply customizations to control the behavior of the VRRP group on committing the VRRP configuration on the router. Without these customizations, the router seizes control of the VRRP group, and immediately assumes the role of the IP address owner virtual router.

Master Router Configuration

Perform the following configuration once you have configured the main bundle interface, *Bundle-Ether100* on the Master router and added the member links to *Bundle-Ether100*.

This example provides the VRRP configuration for IPv4 networks on the Master router.

```
/* Enter the interface configuration mode and configure an IPv4 address for the interface.
*/
Router(config)# interface Bundle-Ether100.1
Router(config-if)# ipv4 address 39.101.1.1/24
Router(config-if)# encapsulation dot1q 1
Router(config-if)# no shut
Router(config-if)# commit
```

```

Fri Feb 23 13:49:24.142 IST
Router:Feb 23 13:49:24.285 : ifmgr[402]: %PKT_INFRA-LINK-3-UPDOWN : Interface
Bundle-Ether100.1, changed state to Down
Router:Feb 23 13:49:24.711 : ifmgr[402]: %PKT_INFRA-LINK-3-UPDOWN : Interface
Bundle-Ether100.1, changed state to Up

Router(config-if)# exit
Router(config)# do show ip int brief
Fri Feb 23 13:50:05.505 IST

Interface                                IP-Address      Status          Protocol        Vrf-Name
Bundle-Ether100.1                       39.101.1.1      Up              Up              default
Bundle-Ether100.2                       unassigned      Shutdown        Down            default
Bundle-Ether100.3                       unassigned      Shutdown        Down            default
Bundle-Ether100.4                       unassigned      Shutdown        Down            default
Bundle-Ether100.5                       unassigned      Shutdown        Down            default

/* Enter the VRRP configuration mode and add the configured interface. */
Router(config)# router vrrp
Router(config-vrrp)# interface Bundle-Ether100.1

/* Configure VRRP for IPv4 */
Router(config-vrrp-if)# address-family ipv4 vrrp 1
Router(config-vrrp-virtual-router)# address 39.101.1.3

/* CUSTOMIZATION: Set a priority for the virtual Router. */
Router(config-vrrp-virtual-Router)# priority 100

/* CUSTOMIZATION: Configure a preempt delay value that controls the selection of the IP
address owner virtual Router. */
Router(config-vrrp-virtual-Router)# preempt delay 30

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

```

Backup Router Configuration

Perform the following configuration once you have configured the main bundle interface, *Bundle-Ether100* on a backup router and added the member links to *Bundle-Ether100*.

This example provides the VRRP configuration for IPv4 networks on a backup router.

```

/* Enter the interface configuration mode and configure an IPv4 address for the interface.
*/
Router(config)# interface Bundle-Ether100.1
Router(config-if)# ipv4 address 39.101.1.2/24
Router(config-if)# encapsulation dot1q 1
Router(config-if)# no shut
Router(config-if)# commit
Fri Feb 23 13:49:24.142 IST
Router:Feb 23 13:49:24.285 : ifmgr[402]: %PKT_INFRA-LINK-3-UPDOWN : Interface
Bundle-Ether100.1, changed state to Down
Router:Feb 23 13:49:24.711 : ifmgr[402]: %PKT_INFRA-LINK-3-UPDOWN : Interface
Bundle-Ether100.1, changed state to Up

Router(config-if)# exit
Router(config)# do show ip int brief
Fri Feb 23 13:50:05.505 IST

Interface                                IP-Address      Status          Protocol        Vrf-Name
Bundle-Ether100.1                       39.101.1.2      Up              Up              default
Bundle-Ether100.2                       unassigned      Shutdown        Down            default

```


Bundle-Ether100.3	unassigned	Shutdown	Down	default
Bundle-Ether100.4	unassigned	Shutdown	Down	default
Bundle-Ether100.5	unassigned	Shutdown	Down	default

```

/* Enter the VRRP configuration mode and add the configured interface. */
Router(config)# router vrrp
Router(config-vrrp)# interface Bundle-Ether100.1

/* Configure VRRP for IPv4 */
Router(config-vrrp-if)# address-family ipv4 vrrp 1
Router(config-vrrp-virtual-router)# address 39.101.1.3

/* CUSTOMIZATION: Set a priority for the virtual Router. */
Router(config-vrrp-virtual-Router)# priority 100

/* CUSTOMIZATION: Configure a preempt delay value that controls the selection of the IP
address owner virtual Router. */
Router(config-vrrp-virtual-Router)# preempt delay 30

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

```

You have successfully configured VRRP for IPv4 networks on the Master and backup routers.

Verification

Use the following commands to verify the VRRP configuration for IPv4 networks on the Master router.

```

/* Verify the configuration */
Router(config-vrrp-virtual-router)# do show run interface Bundle-Ether100.1
Fri Feb 23 15:04:38.140 IST
interface Bundle-Ether100.1
  ipv4 address 39.101.1.1/24
  encapsulation dot1q 1
!

Router(config)# show running-config router vrrp
Fri Feb 23 13:50:18.959 IST
router vrrp
  interface Bundle-Ether100.1
    address-family ipv4
      vrrp 1
        priority 100
        preempt delay 30
        address 39.101.1.3
  !
!

Router# show vrrp detail
Fri Feb 23 15:08:36.469 IST
Bundle-Ether100.1 - IPv4 vrID 1
  State is Master
    16 state changes, last state change 00:05:26
    State change history:
      Feb 23 05:09:11.811 UTC Backup -> Init Interface Down update
      Feb 23 05:13:25.018 UTC Init -> Backup Delay timer expired
      Feb 23 05:13:28.990 UTC Backup -> Master Master down timer expired
      Feb 23 05:13:59.911 UTC Master -> Backup Higher priority advert received
      Feb 23 05:28:20.344 UTC Backup -> Master Master down timer expired
  Last resign sent: Never

```

```

Last resign received: Feb 23 05:02:47.544 UTC
Virtual IP address is 39.101.1.3
Virtual MAC address is 0000.5E00.0101, state is active
Master router is local
Version is 2
Advertise time 1 secs
  Master Down Timer 3.960 (3 x 1 + (246 x 1/256))
Minimum delay 1 sec, reload delay 5 sec
Current priority 10
  Configured priority 10, may preempt
    minimum delay 30 secs

```

You have successfully verified the VRRP configuration for IPv4 networks.

Configure VRRP for IPv6 Networks on Physical Interfaces and Bundle Interfaces

This section describes the procedure for configuring and verifying VRRP for IPv6 networks on Physical interfaces, Physical sub-interfaces, Bundle interfaces, and Bundle sub-interfaces.

Configuration

The following sample includes the configuration and customization of VRRP for IPv6 networks on Physical interfaces, Physical sub-interfaces, Bundle interfaces, and Bundle sub-interfaces.



Note It is recommended to apply customizations to control the behavior of the VRRP group on committing the VRRP configuration on the router. Without these customizations, the router seizes control of the VRRP group, and immediately assumes the role of the IP address owner virtual router.

Master Router Configuration

Perform the following configuration once you have configured the main bundle interface, *Bundle-Ether100* on the Master router and added the member links to *Bundle-Ether100*.

This example provides the VRRP configuration for IPv6 networks on the Master router.

```

/* Enter the interface configuration mode and configure an IPv6 address */
Router# interface Bundle-Ether100.1
Router(config-if)# ipv6 address 39:100:0:1::1:1/64
Router(config-if)# encapsulation dot1q 1
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 Bundle-Ether100.1

/* Enable the IPv6 global and link local address family on the interface */
Router(config-vrrp-if)# address-family ipv6 vrrp 1
Router(config-vrrp-virtual-Router)#address global 39:100:0:1::1:3
Router(config-vrrp-virtual-Router)# address linklocal autoconfig

/* CUSTOMIZATION: Set a priority for the virtual Router. */
Router(config-vrrp-virtual-Router)# priority 100

```

```
/* CUSTOMIZATION: Configure a preempt delay value that controls the selection of the IP
address owner virtual Router. */
Router(config-vrrp-virtual-Router)# preempt delay 30

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

Backup Router Configuration

Perform the following configuration once you have configured the main bundle interface, *Bundle-Ether100* on a backup router and added the member links to *Bundle-Ether100*.

This example provides the VRRP configuration for IPv6 networks on a backup router.

```
/* Enter the interface configuration mode and configure an IPv6 address */
Router# interface Bundle-Ether100.1
Router(config-if)# ipv6 address 39:100:0:1::1:2/64
Router(config-if)# encapsulation dot1q 1
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 Bundle-Ether100.1

/* Enable the IPv6 global and link local address family on the interface */
Router(config-vrrp-if)# address-family ipv6 vrrp 1
Router(config-vrrp-virtual-Router)# address global 39:100:0:1::1:3
Router(config-vrrp-virtual-Router)# address linklocal autoconfig

/* CUSTOMIZATION: Set a priority for the virtual Router. */
Router(config-vrrp-virtual-Router)# priority 100

/* CUSTOMIZATION: Configure a preempt delay value that controls the selection of the IP
address owner virtual Router. */
Router(config-vrrp-virtual-Router)# preempt delay 30

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

You have successfully configured VRRP for IPv6 networks on the Master and backup routers.

Verification

Use the following commands to verify the VRRP configuration for IPv6 networks on the Master router.

```
/* Verify the configuration */
Router(config-vrrp-virtual-router)# do show run interface Bundle-Ether100.1
Fri Feb 23 14:55:48.378 IST
interface Bundle-Ether100.1
  ipv6 address 39:100:0:1::1:1/64
  encapsulation dot1q 1
!

Router(config-vrrp-virtual-router)# do show running-config router vrrp
...
router vrrp
  interface Bundle-Ether100.1
    address-family ipv6
```

```

vrrp 1
  priority 100
  preempt delay 30
  address global 39:100:0:1::1:3
  address linklocal autoconfig
!
!
!

```

Router# **show vrrp detail**

Fri Feb 23 15:08:36.469 IST

Bundle-Ether100.1 - IPv6 vrID 1

State is Master

18 state changes, last state change 00:05:26

State change history:

Feb 23 05:09:11.811 UTC	Backup	-> Init	Interface Down update
Feb 23 05:13:25.018 UTC	Init	-> Backup	Delay timer expired
Feb 23 05:13:28.990 UTC	Backup	-> Master	Master down timer expired
Feb 23 05:13:59.911 UTC	Master	-> Backup	Higher priority advert received
Feb 23 05:28:20.344 UTC	Backup	-> Master	Master down timer expired

Last resign sent: Never

Last resign received: Feb 22 16:57:09.690 UTC

Virtual IP address is fe80::200:5eff:fe00:201

Secondary Virtual IP address is 39:100:0:1::1:3

Virtual MAC address is 0000.5E00.0201, state is active

Master router is local

Version is 3

Advertise time 1 secs

Master Down Timer 3.960 (3 x 1 + (246 x 1/256))

Minimum delay 1 sec, reload delay 5 sec

Current priority 10

Configured priority 10, may preempt

minimum delay 30 secs

You have successfully verified the VRRP configuration for IPv6 networks.

Unicast VRRP

Table 3: Feature History Table

Feature Name	Release Name	Description
Unicast VRRP	Release 7.11.1	<p>We have now enabled Layer 3 unicast transport mode in VRRP, allowing it to enhance its capacity to send data to other networks, including cloud networks. Pairwise router redundancy enables high availability in cloud network scenarios. However, a virtual IP (VIP) address is required by the default route of the cloud native function because there is no pre-designated active member in paired routers. HSRP can provide a VIP, but cloud networks do not support Layer 2 multicast or broadcast transports. You can configure VRRP to support Layer 3 unicast transport to overcome the limitation of Layer 2 multicast and broadcast transports.</p> <p>The feature introduces these changes:</p> <p>New Command:</p> <p>CLI:</p> <ul style="list-style-type: none"> • unicast-peer <p>Modified Commands:</p> <ul style="list-style-type: none"> • show vrrp command is modified to support new fields: Mcast packet in Ucast mode, IPv4 Unicast Peer, and IPv4 Unicast Peer. <p>YANG Data Model:</p> <p>New Xpaths for:</p> <ul style="list-style-type: none"> • <code>Cisco-IOS-XR-ipv4-vrrp-cfg.yang</code> • <code>Cisco-IOS-XR-ipv4-vrrp-oper.yang</code> <p>(see GitHub, YANG Data Models Navigator)</p>

You can now configure VRRP to support Layer 3 unicast transport, allowing it to enhance its capacity to send data to cloud networks. Pairwise router redundancy enables high availability in cloud network scenarios. The default route of the cloud native function needs a virtual IP (VIP) address because the paired routers do not have a pre-designated active member. Though HSRP provides a VIP, the cloud networks do not support Layer 2 multicast or broadcast transports. To overcome the limitations of Layer 2 multicast and broadcast transports, configure VRRP in Layer 3 unicast mode to support Layer 3 unicast transport.

This feature also enables VRRP to communicate state transition notifications using event-driven telemetry.

VRRP over BVI

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
VRRP over BVI	Release 7.5.2	Virtual Router Redundancy Protocol (VRRP) runs on top of interfaces of multiple routers in the same home network that has both Cisco and other vendor routers. It allows a group of routers to behave as a single virtual default gateway router, thereby providing default gateway redundancy and minimizing traffic loss. VRRP now supports Bridge-Group Virtual Interface (BVI), which means that VRRP sessions can run between BVI interfaces of multiple routers.


The Virtual Router Redundancy Protocol (VRRP) protocol provides default gateway redundancy for a LAN. It allows a group of routers to behave as a single virtual default gateway router. In the group of routers, the router that has the highest VRRP priority acts as the primary router and the remaining routers act as backup routers.

Bridge-Group Virtual Interface (BVI) is a virtual interface which provides Layer 3 or routed functionality to a bridge group. Layer 2 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. You can configure VRRP session only over BVI.

Topology

This topology showcases how VRRP functions over BVI.



In this topology, PE1 and PE2 are paired in a redundant group. This group provides Layer 3 gateway service to CE1 and CE2. VRRP is configured over BVI interfaces on PE1 and PE2. VRRP ensures one BVI is the active gateway. The other is the standby gateway.

You can configure one of the BVIs to be active and the other BVI as standby by setting the VRRP priority value. The active BVI is programmed with the virtual MAC address chosen by VRRP. Hosts, CE1 and CE2 send the traffic to the virtual destination MAC address and the active BVI forwards the traffic.

During failover, the standby BVI becomes active and is programmed with the virtual MAC address. The traffic from the hosts is forwarded through this active BVI.

Supported Scale and Systems

VRRP over Bridge Virtual Interfaces (BVIs) is supported:

- You can configure upto 512 VRRP groups (IPv4 and IPv6 combined) over BVIs on:
 - Cisco Silicon One Q100 ASIC-based systems
 - Cisco Silicon One Q200 ASIC-based systems
 - Cisco Silicon One P100 ASIC-based systems
- Where the underlay IRB bridge domains consist of bridge members on L2 main or subinterfaces. Only physical and bundle interfaces are supported for L2 bridging in IRB.

- When both IPv4 and IPv6 are configured on a BVI interface, IPv4 and IPv6 each requires a session. A total of two sessions are consumed on a BVI interface.
- For IPv4 and IPv6 configurations, in both the default and VRF tables.
- On both the fixed and distributed systems.
- VRRP on BVI is supported on the 88-LC1-12TH24FH-E line card from Cisco IOS XR Release 25.1.1.

Restrictions

Consider these restrictions before you configure VRRP over BVIs.

- The minimum supported VRRP Hello timer is 100 ms. At the minimum timer, a total of 50 sessions are supported. Above 100 ms timers, the sessions scale goes up proportionately. A maximum of 255 VRRP groups and 510 sessions are supported.
- VRRP on BVI is not supported on the 88-LC1-12TH24FH-E line card in Cisco IOS XR Release 24.3.1.

Configure VRRP over BVI

To configure VRRP sessions over BVI, you must complete the following configurations on PE1 and PE2:

1. Configure a set of interfaces as Layer 2 interfaces and a set of VLAN sub-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 Layer 2 interfaces
   and a set of VLAN sub-interfaces */
Router# configure
Router(config)# interface HundredGigE0/0/1/0.1 l2transport
Router(config-subif)# encapsulation dot1q 1
Router(config-subif)# rewrite ingress tag pop 1 symmetric
Router(config-subif)# commit
Router(config-subif)# exit
Router(config)# interface HundredGigE0/0/1/1.1 l2transport
Router(config-subif)# encapsulation dot1q 1
Router(config-subif)# rewrite ingress tag pop 1 symmetric
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 HundredGigE0/0/1/0.1
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# interface HundredGigE0/0/1/1.1
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)# 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)# address 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
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: 5, bridge-domain: 5, id: 1, 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: 32768, 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: 2
Filter MAC addresses:

```

```

P2MP PW: disabled
Multicast Source: Not Set
Create time: 26/05/2020 17:08:54 (00:11:30 ago)
No status change since creation
ACs: 3 (3 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 0x80000001; interworking none
BVI MAC address:
c472.95a6.8b90
Virtual MAC addresses:
0000.5e00.010a
0000.5e00.020b
Split Horizon Group: Access
AC: HundredGigE0/0/1/0.1, state is up
Type VLAN; Num Ranges: 1
Rewrite Tags: []
VLAN ranges: [1, 1]
MTU 1500; XC ID 0x1; interworking none
MAC learning: enabled

```

Use the following command to show the VRRP details:

```
Router# show vrrp ipv4 detail
```

```

BVI10 - IPv4 vrID 10
State is Master
2 state changes, last state change 00:11:57
State change history:
May 26 17:08:59.470 UTC Init -> Backup Delay timer expired
May 26 17:09:03.075 UTC Backup -> Master Master down timer expired
Last resign sent: Never
Last resign received: Never
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
2 state changes, last state change 00:04:29
State change history:
May 26 17:16:43.476 UTC Init -> Backup Virtual IP configured
May 26 17:16:47.085 UTC Backup -> Master Master down timer expired
Last resign sent: Never
Last resign received: Never
Virtual IP address is fe80::200:5eff:fe00:20b
Secondary Virtual IP address is 2001:db8:a:b::2
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

Router# show vrrp interface BVI10 detail
BVI10 - IPv4 vrID 10
State is Master
2 state changes, last state change 00:12:35
State change history:
May 26 17:08:59.470 UTC Init -> Backup Delay timer expired
May 26 17:09:03.075 UTC Backup -> Master Master down timer expired
Last resign sent: Never
Last resign received: Never
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

BVI10 - IPv6 vrID 11
State is Master
2 state changes, last state change 00:04:51
State change history:
May 26 17:16:43.476 UTC Init -> Backup Virtual IP configured
May 26 17:16:47.085 UTC Backup -> Master Master down timer expired
Last resign sent: Never
Last resign received: Never
Virtual IP address is fe80::200:5eff:fe00:20b
Secondary Virtual IP address is 2001:db8:a:b::2
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

```

View VRRP statistics in the Router

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
View VRRP statistics in Router	Release 24.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature support is now extended to Cisco 8712-MOD-M routers.

Feature Name	Release Information	Feature Description
View VRRP statistics in Router	Release 7.9.1	<p>With this feature, you can view or clear statistics of one or all Virtual Router Redundancy Protocol (VRRP) groups or Virtual Router IDs (VRIDs). This information helps you monitor VRRP health in the routers. It's also helpful in debugging VRRP issues like packet exchange failures when all virtual routers in the VRRP topology function as backup virtual routers and there's no IP address owner.</p> <p>This feature introduces the following commands:</p> <ul style="list-style-type: none"> • CLI: <ul style="list-style-type: none"> • show vrrp statistics • clear vrrp statistics • YANG Data Model: <p>Cisco-IOS-XR-ipv4-vrrp-oper.yang Cisco native model (see GitHub, YANG Data Models Navigator)</p>

You can view statistics of one or all Virtual Router Redundancy Protocol (VRRP) virtual routers using the **show vrrp statistics** command as follows:

```
Router# show vrrp statistics
Invalid packets:
  Invalid checksum:                0
  Unknown/unsupported versions:    3
  Invalid vrID:                    1
  Too short:                       7
Protocol:
  Transitions to Master            4
Packets:
  Total received:                  54
  Adverts sent:                    0
  Bad TTL:                         0
  Short Packets:                   6
  Failed authentication:           0
  Unknown authentication:          2
  Conflicting authentication:       0
  Unknown Type field:              1
  Conflicting Advertise time:       0
  Conflicting Addresses:           0
  Received with zero priority:      9
  Sent with zero priority:          0
```

You can clear the VRRP statistics using the **clear vrrp statistics** as follows:

```
Router# clear vrrp statistics
Router# show vrrp statistics
Invalid packets:
  Invalid checksum: 0
  Unknown/unsupported versions: 0
  Invalid vrID: 0
  Too short: 0
Protocol:
  Transitions to Master 0
Packets:
  Total received: 0
  Adverts sent: 0
  Bad TTL: 0
  Short Packets: 0
  Failed authentication: 0
  Unknown authentication: 0
  Conflicting authentication: 0
  Unknown Type field: 0
  Conflicting Advertise time: 0
  Conflicting Addresses: 0
  Received with zero priority: 0
  Sent with zero priority: 0
```

