



Configuration Guide

VRRPv3 for AOS

This configuration guide will aid in setting up the Virtual Router Redundancy Protocol version 3 (VRRPv3) for ADTRAN Operating System (AOS) products using Internet Protocol version 4 (IPv4) or IP version 6 (IPv6). Included in the guide are step-by-step configuration instructions using the command line interface (CLI), configuration examples, and troubleshooting information.

This guide consists of the following sections:

- *VRRPv3 Overview on page 2*
- *Hardware and Software Requirements and Limitations on page 4*
- *VRRPv3 CLI Configuration on page 5*
- *Configuring Additional Virtual Router Settings (Optional) on page 8*
- *Example VRRPv3 Configurations on page 11*
- *VRRPv3 Configuration Command Summary on page 13*
- *Troubleshooting on page 15*

VRRPv3 Overview

VRRPv3 is a standard protocol that allows multiple physical routers to act as a single virtual router on an Ethernet network. The virtual router is comprised of two or more physical routers running VRRPv3, and acts as a single virtual router for hosts on a shared local area network (LAN). This functionality provides seamless redundancy to networked end-host devices, and allows the forwarding of host traffic if a router or interface port fails within the network. VRRPv3 is defined in RFC 5798, and functions similarly to IPv4 VRRPv2, with the primary difference being the support of both IPv4 and IPv6. IPv6 provides some default router failover support through Neighbor Unreachability Detection (NUD), but pairing VRRPv3 with IPv6 provides a quicker method to allow switchover to backup routers than is possible with regular IPv6 Neighbor Discovery (ND) processing.

In VRRPv3, one physical router acts as the master controller of the virtual router. The other VRRPv3 routers are designated as virtual router backups and will only assume forwarding responsibility for the virtual router should the current master fail or have its priority reduced. The master router is responsible for answering host requests for IPv6 ND address resolution requests or IPv4 Address Resolution Protocol (ARP) requests and forwarding packets sent to the virtual router. Hosts typically use the virtual router address as the default router. In IPv6, the virtual router also participates in ND router solicitation (RS) and router advertisement (RA) exchanges when these exchanges are associated with the virtual router address (as would be done if the virtual router were a standard IPv6 router). Through this process, hosts learn of the virtual router advertising itself as a default router. The virtual router transmits RAs with a non-zero router lifetime value. IPv6 hosts then store the link-local address found in the source address of the received RA as an entry in their default router list. By advertising virtual link-local addresses, multiple IPv6 routers can dynamically assign responsibility for one-way forwarding of packets sent by these IPv6 hosts, depending on router condition and priority. In IPv4, the router answers ARP requests for virtual IPv4 addresses.

In addition, the master controller sends periodic VRRPv3 advertisements indicating to any listening backup routers that it is the current master. Backup routers listen to the master advertisements, but do not normally send any advertisements of their own. If a backup router has a higher priority, based on received advertisements from the master, it will begin to send out packets and the current master transitions to a backup router. The typical IPv6 VRRPv3 network topology is illustrated in *Figure 1*. A typical IPv4 VRRPv3 network will be similar, except using IPv4 addresses.

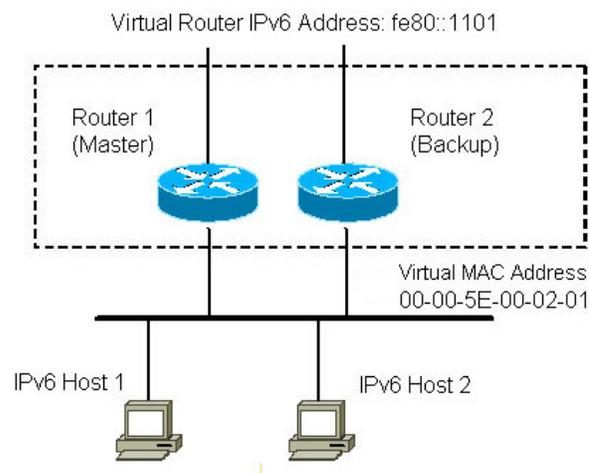


Figure 1. Typical IPv6 VRRPv3 Network Topology

VRRPv3 VRIDs

Virtual router IDs (VRIDs) are numbers used to associate a group of physical routers running VRRPv3. The virtual router master in the group will periodically send out VRRPv3 packets to all other routers within the group. These packets communicate the priority and state of the virtual router master. In addition, VRID values are advertised by VRRPv3, and are used to generate virtual router media access control (MAC) addresses. The VRID must be the same on all routers that are part of the virtual router group.

The VRRPv3 Virtual Master

The virtual router master responds to host requests for IP to Layer 2 address resolution (ND address requests (ARs) for IPv6 and ARP requests for IPv4) for the virtual router address. The master router, as part of a virtual router group, performs IPv4 ARP or IPv6 ND AR exchange functionality upon startup, or when transitioning from the backup to the master state. For each IPv6 address associated with the virtual router, unsolicited ND neighbor advertisement (NA) messages are sent. If a master virtual router sends an ND neighbor solicitation (NS) message for a host's IPv6 address, it must include the virtual MAC address for the virtual router if it sends a source link-layer address option in the NS message. All VRRPv3 packets are sent from the VRRPv3 unicast MAC address towards the VRRPv3 multicast MAC address to ensure MAC learning tables are populated with the appropriate port for each virtual router instance.

When a router is the master in a virtual router group, it also processes received RS and sends RA messages for the prefixes and options associated with the interface. To complete this task, the router must meet the following criteria:

- It must be a member of the solicited node multicast address for the IPv6 address associated with the virtual router.
- It must respond to ND NS messages for the IPv6 addresses associated with the virtual router.
- It must send ND RAs for the virtual router. Any backup routers must be configured by the user to send the same RA options as the current master. This includes options like prefixes, prefix option bits, maximum transmission unit (MTU) information, M and O bits, and any other RA information.
- A VRRPv3 router in the backup state does not log events that indicate a misconfiguration of virtual router addresses with respect to the current master. The backup router does not have the information available in its prefix tables for comparison.
- If RAs are suppressed on an interface, for example by using the **ipv6 nd ra suppress** command, then any RAs sent for a virtual router on that link will be suppressed when the router is in a master state.
- Any RA-related configuration on an interface also applies to any virtual router prefixes while the router is in the master state.

When the virtual router is a master, it responds to address resolution requests for the virtual router address, and forwards traffic from hosts using the virtual address as their default router. In addition, the virtual router can respond to packets addressed directly to the virtual address, for example, when using ping capabilities or other local access to the master device for debug purposes.

RA Messages and IPv6 VRRPv3

When using IPv6 VRRPv3, a router sends RAs to indicate that it is able to be a default router for listening hosts, to advertise prefixes and their options, to provide link configuration information such as MTU and to provide source link-layer address information. A host uses the link-local source address of the router in the RA as a default router candidate if the router lifetime is not zero. Global virtual router IPv6 addresses are not required by a host using the default router.

When VRRPv3 is enabled on an interface, the RA messages can change in the following ways:

- When a router is configured for VRRPv3, it stops sending RAs from the interface's default link-local address and begins sending out one or more RAs with a lifetime of zero (thus withdrawing its intention to be a default router using this address).
- If a virtual router transitions to the master state, it begins building and sending RAs from the virtual link-local address.
- A virtual router in the backup state continues to suppress sending RAs.
- If multiple virtual routers are configured on the AOS unit, each sends separate RAs if the virtual router is in the master state. Each virtual router sends its own source link-layer address option if it is included for address resolution.

ARP and IPv4 VRRPv3

When using IPv4 VRRPv3, the ARP functionality of VRRPv3 includes the following considerations:

- The IPv4 VRRPv3 router in the backup state will not respond to the virtual IP address. The master router will, however, respond to ARP requests for its own IPv4 address as well as the IPv4 virtual address.
- VRRPv3 advertisement packets generated from the router are sourced from the VRRPv3 IPv4 unicast MAC address (based off the VRID) and sent to the VRRPv3 IPv4 multicast MAC address. All other traffic forwarded to and through the unit is addressed to the virtual unicast MAC address.
- Because ARP responses differ when a router is in backup or master state, interface addresses are not allowed to match secondary addresses on the interface. This would cause problems with ARP responses to the secondary address where both VRRP units would claim they have the address.

Hardware and Software Requirements and Limitations

The IPv6 VRRPv3 feature is supported on AOS devices running AOS firmware R10.7.0 or later, as outlined in the *AOS Feature Matrix*, available online at <https://supportforums.adtran.com>. The IPv6 VRRPv3 syntax was modified in AOS firmware release R10.11.0 to require the address family designation. IPv4 VRRPv3 is supported on AOS devices running AOS firmware R10.11.0 or later, as outlined in the *AOS Feature Matrix*, available online at <https://supportforums.adtran.com>.

VRRPv3 can include support for either IPv4 or IPv6. Both versions can be enabled on an interface simultaneously.

VRRPv3 is disabled by default on all AOS products.

The maximum number of IPv6 VRRPv3 routers supported per interface is two, and the maximum number of IPv4 VRRPv3 routers supported per interface is two.

In VRRPv3, IPv6 duplicate address detection (DAD) is disabled for virtual addresses.

VRRPv3 CLI Configuration

To configure VRRPv3 using the CLI, you will need to access the CLI and complete the following:

- *Step 1: Enable IPv6 on the Interface (for IPv6 VRRPv3) on page 5*
- *Step 2: Create a Virtual Router Instance on the Interface on page 5*
- *Step 3: Specify the Priority Level of the Virtual Router Instance on page 7*

In addition, you can optionally choose to configure additional router settings. These settings are outlined in *Configuring Additional Virtual Router Settings (Optional) on page 8*.

Step 1: Enable IPv6 on the Interface (for IPv6 VRRPv3)

IPv6 must be enabled on the interface before any VRRPv3 configuration commands can be entered when using IPv6 VRRPv3. To enable IPv6 and create a link-local address on the interface, enter the **ipv6** command from the interface's configuration mode as follows:

```
(config)#interface ethernet 0/1  
(config-eth 0/1)#ipv6
```



You can also enable IPv6 on an interface and specify a unicast IPv6 address for that interface using the **ipv6 address** commands. For more information about these commands, or IPv6 functionality in AOS, refer to the configuration guide *Configuring IPv6 in AOS*, available online at <https://supportforums.adtran.com>.

Step 2: Create a Virtual Router Instance on the Interface

To create a virtual router instance on the interface, you must enable VRRPv3 and create a virtual router identifier (VID), as well as specify which IP VRRPv3 version you are using. IP VRRPv3 versions are specified using an address family; one address family for IPv4 and one address family for IPv6. The router instance is then configured from the VRRPv3 address family's configuration mode. Follow these steps to create a virtual router instance on the interface:

1. Enable VRRPv3, create a VRID, and specify an address family using the **vrrpv3 <vrid> address-family [ipv4 | ipv6]** command from the interface's configuration mode. This command creates a virtual router, identified by its VRID and address family, and enters the configuration mode for that particular virtual router. The **<vrid>** parameter specifies the VRID for the virtual router instance. This value is advertised by VRRPv3 and used to generate the virtual router MAC address. VRID values must be the same on all routers that are part of the virtual router group. Valid VRID range is **1** to **255**. VRID numbering is independent between VRRPv3 IPv4 and IPv6 address families. Address families are specified using the **ipv4** and **ipv6** keywords, where the keyword indicates the IP version used. Using the **no** form of this command removes the virtual router and its configuration from VRRPv3. For example, to enable IPv4 VRRPv3, and create an IPv4 virtual router instance, enter the command as follows:



For IPv4 VRRPv3, a primary IPv4 address must be configured on the interface before VRRPv3 configuration commands can be entered.

```
(config)#interface ethernet 0/1
(config-eth 0/1)#vrrpv3 15 address-family ipv4
(config-if-vrrpv3 15)#
```

To enable IPv6 VRRPv3, and create an IPv6 virtual router instance, enter the command as follows:

```
(config)#interface ethernet 0/1
(config-eth 0/1)#ipv6
(config-eth 0/1)#vrrpv3 15 address-family ipv6
(config-if-vrrpv3 15)#
```

2. Specify a primary VRRPv3 IP address for the virtual router by configuring an IPv4 or IPv6 address on the virtual router instance. A primary virtual address must be configured before the VRRPv3 router will become active, and this address must be configured to the same value on all VRRPv3 routers that are part of the same VRID group. This address is the address that the virtual router instance on this unit takes if it transitions from a backup to master device.

Configuring IPv6 Virtual Addresses

You assign virtual addresses to an IPv6 virtual router instance on the interface by using one of three commands: **address** *<ipv6 link-local address>* **primary**, **address** *<ipv6 global-address>* **secondary**, or **address generate primary**. The first command allows you to specify a link-local address for the virtual router instance, the second command allows optionally specifying virtual secondary IPv6 global addresses for the virtual router instance, and the third command assigns a default virtual link-local address to the virtual router based on a computation from the virtual interface ID and virtual MAC address associated with the VRID. A virtual link-local address must be configured before the VRRPv3 router becomes active. Only one virtual link-local address can be specified per virtual router. If one command is used to configure the link-local address, and then the command is entered a second time, the second command replaces the previous command in the configuration. The *<ipv6 link-local address>* parameter specifies the link-local address that will be used by the virtual router instance. A link-local IPv6 address is specified in the format **FE80::<hex-digits>**. The *<hex-digits>* are the lower 64 bits of the link-local IPv6 address, and since link-local addresses have no prefix, the bits entered form the entire IPv6 address. The **generate primary** parameter specifies that the virtual link-local address is computed from the VRID that is itself computed from the virtual MAC address associated with the VIRD. Using the **no** form of either command removes the link-local address and its configuration from the virtual router.

To create an IPv6 virtual router instance with a specified link-local address, enter the command from the virtual router instance's configuration mode as follows:

```
(config)#interface ethernet 0/1
(config-eth 0/1)#ipv6
(config-eth 0/1)#vrrpv3 15 address-family ipv6
(config-if-vrrpv3 15)#address fe80::220:8FF:fE54:F9D8 primary
```

To create an IPv6 virtual router instance with an assigned default link-local address, enter the command from the virtual router instance's configuration mode as follows:

```
(config)#interface ethernet 0/1
(config-eth 0/1)#ipv6
(config-eth 0/1)#vrrpv3 15 address-family ipv6
(config-if-vrrpv3 15)#address generate primary
```



*You can optionally specify a secondary virtual global IPv6 address to be used by the virtual router by entering the **address** <ipv6 global-address> **secondary** command. The address specified must be on the same subnet as a configured IPv6 address on the interface because it is needed for prefix lengths in RA messages. This address cannot match an IPv6 address on the interface.*

Configuring IPv4 Virtual Addresses

You can create an IPv4 virtual router instance address on the interface by using the **address** <ipv4 address> <**primary** | **secondary**> command from the virtual router instance's configuration mode. This command assigns an IPv4 address to the virtual router. The primary address must be configured before the virtual router will become active and before any secondary addresses can be added. The primary address is specified using the **primary** keyword, and any secondary addresses are specified using the **secondary** keyword. The primary address is the only address that can be an owner address (the same as the configured router IPv4 address for the interface). Entering a second primary address replaces any existing primary address. The specified address must match a prefix on this interface, and the address must not be the same as a virtual address in a different VRID group. The <ipv4 address> specifies the virtual address in decimal dotted notation, for example, **10.10.10.1**. The specified address must be on the same subnet as a configured IPv4 address on the interface. To create an IPv4 virtual router instance with an assigned address, enter the command from the virtual router instance's configuration mode as follows:

```
(config)#interface ethernet 0/1
(config-eth 0/1)#vrrpv3 15 address-family ipv4
(config-if-vrrpv3 15)#address 10.10.10.1 primary
```

Step 3: Specify the Priority Level of the Virtual Router Instance



It is possible for VRRPv3 to operate based on default priority level settings. However, it is important to understand and verify the election process to ensure the desired VRRPv3 router is ultimately selected as the master.

The virtual router master is selected using an election process based on the priority level value of each VRRPv3 router. There are two types of priority levels: configured and actual. The configured priority level is the numerical value originally assigned to the VRRPv3 router. For example, the default configured priority for a VRRPv3 router that is not the virtual link-local address owner is **100**. The actual priority level is the value that takes into account any adjustments resulting from a track state. For example, the actual priority level of a master router whose network path is down may be **90** versus the configured priority of **100**. The actual priority of a VRRPv3 router that is the virtual link-local address owner is always **255**. The VRRPv3 router with the highest actual priority level is the virtual router master. If there is a tie, the actual priority level will be used along with the interface link-local address to determine the master.

If the virtual link-local address of the virtual router is the same as the interface address on the VRRPv3 router (the link-local address owner), then the default priority level will be **255**. This is the highest possible priority level and means that when the link-local address owner is available, it will be the virtual router master.



Object tracking can be used to decrement the priority level in a VRRPv3 router. Refer to [Tracking Objects on page 9](#) for more details on object tracking.

To specify the virtual router instance priority, enter the **priority** *<level>* command from the virtual router instance's configuration mode. The *<level>* parameter specifies the configured priority level of the virtual router. Valid range is **1** to **254**. Using the **no** form of this command removes the configured priority from the virtual router and returns the router to the default setting. To configure the priority level for the virtual router, enter the command as follows:

```
(config-if-vrrpv3 15)#priority 200
```

Configuring Additional Virtual Router Settings (Optional)

After configuring the VRID, link-local address, and priority for the virtual router instance, you can optionally configure other parameters for the virtual router. These options include adding a description, specifying a preempt or start-up delay, configuring advertisement timers, tracking objects, and matching addresses. Each of these optional configurations are described in the following sections.

Text Description

An optional text description can be added to provide details about the VRRPv3 group using the **description** *<text>* command from the virtual router instance's configuration mode. The *<text>* parameter is a text string used to describe the virtual router instance. By default, no description is configured. Using the **no** form of this command removes the text description from the virtual router's configuration. To add a description, enter the command as follows:

```
(config-if-vrrpv3 15)#description RemoteRoutersA
```

Preempt

Whenever a VRRPv3 router with a higher actual priority level than the current master is added to a virtual router group, it attempts to take over or preempt the master router if preemption is enabled. This behavior may be desired, for example, when the new router is more powerful than the existing master router. To specify a virtual router can preempt the current master, enter the **preempt [delay minimum <delay>]** command from the virtual router instance's configuration mode. The optional **delay minimum <delay>** parameter specifies that the VRRPv3 router waits a specified amount of time before attempting to preempt the master. Delays are specified in seconds, with a valid range of **0** to **255** seconds. By default, a VRRPv3 router preempts with no additional delay. Using the **no** form of this command disables preemption, unless the master router is the link-local address owner (in which case it will always preempt, and this command is ignored). Enter the command as follows:

```
(config-if-vrrpv3 15)#preempt delay minimum 30
```



When VRRPv3 transitions occur between master and backup devices, security related configurations can be affected and security policies may need to be adjusted. For example, preemption might need to be disabled to prevent loss of state or session, particularly in use of long-term Transmission Control Protocol (TCP) connections, such as Telnet or Secure Socket Layer (SSL).

Startup Delay

You can optionally specify that the virtual router wait a specified amount of time before running the VRRPv3 state machine after initial activation. The delay can prevent inadvertent declarations of multiple VRRPv3 masters while waiting for potentially slower interfaces to begin passing advertisement traffic. By default, the startup delay is set to **35** seconds. To change the startup delay, enter the **startup-delay** *<delay>* command from the virtual router interface's configuration mode. The *<delay>* parameter is the time, in seconds, that the virtual router waits to start the virtual router. Valid range is **0** to **255** seconds. Using the **no** form of this command returns the startup delay to the default value. Enter the command as follows:

```
(config-if-vrrpv3 15)#startup-delay 100
```



If the virtual router receives an advertisement from an external master during the delay period, it assumes that the network path is open, cancels the timer, and transitions to the appropriate backup or master state.

Advertisement Interval

The master router sends VRRPv3 advertisements at a specified interval. By default, these advertisements are sent every **1** second. You can optionally specify a new advertisement interval using the **timers advertise** *<interval>* command in the virtual router instance's configuration mode. The *<interval>* parameter specifies the time between advertisements in seconds. Valid range is **1** to **40** seconds. Using the **no** form of this command returns the advertisement interval to the default value. To change the interval, enter the command as follows:

```
(config-if-vrrpv3 15)#timers advertise 5
```



All VRRPv3 routers in a virtual router group must use the same advertisement interval.

Tracking Objects

Object tracking can be used to change the priority level of a VRRPv3 router. One purpose of object tracking for VRRPv3 is to detect failure of an interface or main path connected to the master router. The object that is tracked is typically an interface towards corporate or the far end of the main path. Depending on the object type specified, as long as the line protocol for the interface is up the track will remain in a PASS state and nothing will change. However, if the line protocol for an interface goes down, the track will transition to a FAIL state. This transition will cause the priority level of the VRRPv3 router to be decremented by the numerical amount specified (or the default amount). Use the **track** *<name>* [**decrement** *<value>*] command to track objects. The *<name>* parameter is the name of the track. The

optional **decrement** *<value>* parameter specifies the numerical amount to decrement the priority level if the track transitions to a FAIL state. The decrement value range is **1** to **254**, and the default value is **10**. Using the **no** form of this command removes the track reference and its configuration from the virtual router.



*A track must be created before the **track** command can be issued. For detailed information on creating probes and track, refer to the configuration guide *Configuring Network Monitor in AOS*, available online at <https://supportforums.adtran.com>.*



*If a VRRPv3 router owns the virtual router link-local address, then the VRRPv3 router's priority level cannot be decremented as a result of the **track** command. Therefore, if object tracking is used to monitor paths and effectively decrement priority levels in case of interface or path failure, then it is important that no VRRPv3 router own the virtual router link-local address.*

To enable object tracking, enter the command as follows from the virtual router instance's configuration mode:

```
(config-if-vrrpv3 15)#track TRACK1 decrement 20
```

Shutdown

The **shutdown** command disables the VRRPv3 router and the **no shutdown** command enables the VRRPv3 router. If the VRRPv3 router is currently the master, entering the **shutdown** command causes the router to advertise a priority level of **0**, thus relinquishing its role as the master and allowing a backup router to take over. By default, VRRPv3 is enabled (**no shutdown**) when the virtual link-local address is configured. Enter the command from the virtual router instance's configuration mode as follows:

```
(config-if-vrrpv3 15)#shutdown
```

Secondary Address Matching

When using VRRPv3, you can optionally specify that VRRPv3 virtual routers will match secondary IP addresses when receiving advertisement packets from peer virtual routers. This matching activity is used to verify configuration consistency and is enabled to occur by default. Secondary address matching is enabled using the **match-address** command from the virtual router instance's configuration mode. Primary virtual address configuration consistency is always verified regardless of the value of this setting. By default, secondary address matching is enabled for VRRPv3. Using the **no** form of this command disables secondary address matching, but still forces primary address matching. To disable secondary address matching for VRRPv3, enter the command as follows:

```
(config-if-vrrpv3 15)#no match-address
```

Example VRRPv3 Configurations

The following example scenarios are designed to enhance understanding of VRRPv3 configuration on AOS products.

IPv6 VRRPv3 Configuration Example

In this example, a simple VRRPv3 network is configured where two routers provide redundant default gateway support to IPv6 hosts on an Ethernet LAN. The master router advertises the shared virtual router link-local address as a default router on the network. *Figure 2* illustrates the VRRPv3 network for this example.

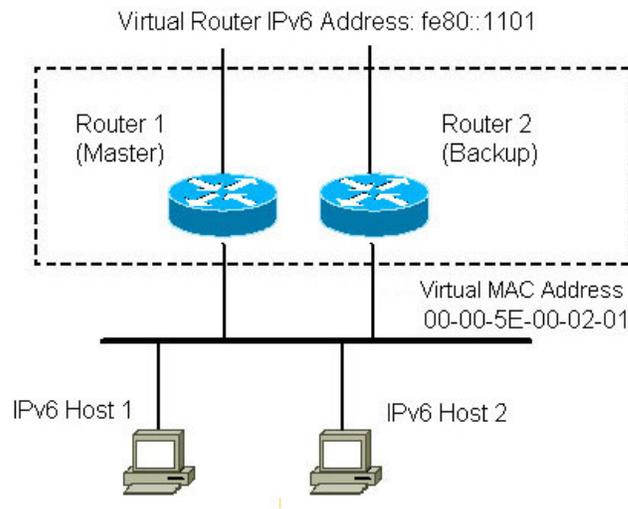


Figure 2. IPv6 VRRPv3 Network

The VRRPv3 configuration on each router is as follows:

Router 1 (Master):

```
interface ethernet 0/1
  ipv6
  vrrpv3 1 address-family ipv6
  address fe80::1101 primary
  priority 180
  exit
```

!

Router 2 (Backup):

```
interface ethernet 0/1
  ipv6
  vrrpv3 1 address-family ipv6
  address fe80::1101 primary
  exit
```

!

IPv4 and IPv6 VRRPv3 Configuration Example

In this example, a simple VRRPv3 network is configured where two routers provide redundant default gateway support to IPv4 and IPv6 hosts on an Ethernet LAN. The master router advertises the shared virtual router address as the default router on the network. *Figure 3* illustrates the VRRPv3 network for this example.

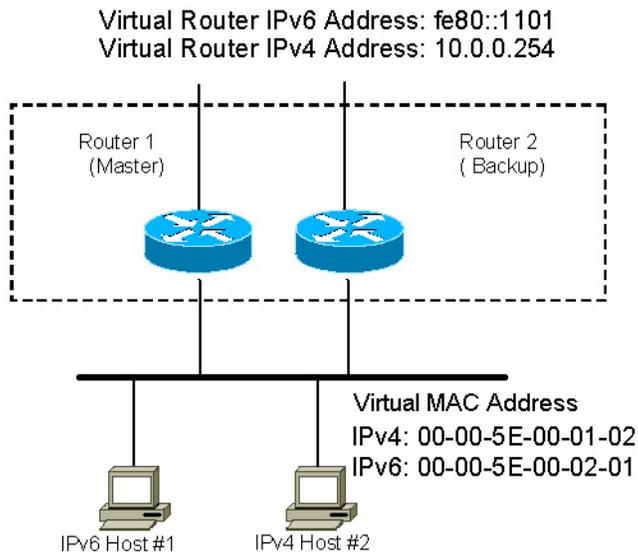


Figure 3. IPv4 and IPv6 VRRPv3 Network

The VRRPv3 configuration on each router is as follows:

Router 1 (master):

```
interface ethernet 0/1
  ip address 10.0.0.252 255.255.255.0
  ipv6
  vrrpv3 2 address-family ipv4
    address 10.0.0.254 primary
    address 10.0.0.1 secondary
    priority 180
  exit
  vrrpv3 1 address-family ipv6
    address fe80::1101 primary
    priority 180
  exit
```

!

Router 2 (Backup):

```
interface ethernet 0/1
  ip address 10.0.0.253 255.255.255.0
  ipv6
  vrrpv3 2 address-family ipv4
```

```

address 10.0.0.254 primary
address 10.0.0.1 secondary
priority 100
exit
vrrpv3 1 address-family ipv6
address fe80::1101 primary
priority 100
exit

```

!

VRRPv3 Configuration Command Summary

The following tables summarize the minimum steps and additional settings required to configure VRRPv3 on an AOS product.

Table 1. VRRPv3 Configuration Commands

Command	Description
(config-eth 0/1)# ipv6	Enables IPv6 on the interface and assigns the interface a link-local address. This command is for use with IPv6 VRRPv3.
(config-eth 0/1)# vrrpv3 <vrid> address-family ipv4 OR (config-eth 0/1)# vrrpv3 <vrid> address-family ipv6	Enables VRRPv3, specifies the virtual router ID, and specifies the IPv4 address family. VRID valid range is 1 to 255 . This command is for use with IPv4 VRRPv3. Enables VRRPv3, specifies the virtual router ID, and specifies the IPv6 address family. VRID valid range is 1 to 255 . This command is for use with IPv6 VRRPv3.
(config-if-vrrpv3 15)# address <ipv6 link-local address> primary OR (config-if-vrrpv3 15)# address generate primary	Assigns an IPv6 address to the virtual router. Specify IPv6 link-local addresses in colon hexadecimal notation (for example, fe80::1101). The primary keyword specifies the address is the primary address. The generate primary keywords specify that a default virtual link-local address is assigned to the virtual router. This command is for use with IPv6 VRRPv3.
(config-if-vrrpv3 15)# address <ipv6 global-address> secondary	Assigns a secondary global IPv6 address to the virtual router. Specify global IPv6 address in colon hexadecimal notation (for example, 2001:DB8::1). This command is for use with IPv6 VRRPv3.

Table 1. VRRPv3 Configuration Commands

Command	Description
(config-if-vrrpv3 15)# address <ipv4 address> < primary secondary >	Assigns an IPv4 address to the virtual router. Specify IPv4 addresses in decimal dotted notation (for example, 10.10.10.1). The primary keyword specifies the address is the primary address, and the secondary keyword specifies the address is a secondary address. This command is for use with IPv4 VRRPv3.
(config-if-vrrpv3 15)# priority <level>	Optional. Specifies the priority level of the VRRPv3 router. Level range is 1 to 254 , with a default value of 100 unless the router is the owner of the virtual IP address, in which case it is 255 . This command is for use with both IPv4 and IPv6 VRRPv3.

Table 2. Additional VRRPv3 Configuration Settings

Command	Description
(config-if-vrrpv3 15)# description <text>	Specifies a textual description for a virtual router group. This command is for use with both IPv4 and IPv6 VRRPv3.
(config-if-vrrpv3 15)# preempt [delay minimum <time>]	Allows a VRRPv3 router to preempt the current master router if its priority level is higher than the current master's. Preemption is enabled by default. The optional delay minimum keywords specify the amount of time (in seconds) the router will wait before attempting to preempt the current master router. Range is 0 to 255 seconds. This command is for use with both IPv4 and IPv6 VRRPv3.
(config-if-vrrpv3 15)# startup-delay <delay>	Configures the router with a hold down timer that prevents it from becoming active before the physical interface has had a chance to initialize (for example, when waiting on spanning-tree convergence for a connected switchport). Delay range is 0 to 255 seconds. By default, startup delay is set to 35 seconds. This command is for use with both IPv4 and IPv6 VRRPv3.
(config-if-vrrpv3 15)# timers advertise <interval>	Specifies the time (in seconds) between advertisements sent by the master router. Range is 1 to 255 seconds. This command is for use with both IPv4 and IPv6 VRRPv3.

Table 2. Additional VRRPv3 Configuration Settings (Continued)

Command	Description
(config-if-vrrpv3 15)# track <name> [decrement <value>]	Associates a track with the VRRPv3 group and, upon failure of this track, decrements the priority level of the router by the specified amount. Decrement value range is 1 to 254 . If the decrement value is not specified, the default decrement value of 10 is used. This command is for use with both IPv4 and IPv6 VRRPv3.
(config-if-vrrpv3 15)# shutdown	Disables the VRRPv3 router. This command is for use with both IPv4 and IPv6 VRRPv3.
(config-if-vrrpv3 15)# match-address	Specifies that the virtual router matches secondary IP addresses by default when receiving packets from nearby virtual routers. The primary address is always matched. Enabled by default. This command is for use with both IPv4 and IPv6 VRRPv3.

Troubleshooting

After configuring VRRPv3, several commands can be issued from Enable mode in the CLI to assist in troubleshooting. The following table contains the AOS troubleshooting commands that are implemented specifically for VRRPv3.

Table 3. Enable Mode Troubleshooting Commands

Command	Description
show vrrpv3 [brief statistics interface <interface>] [group <vrid>] [<ipv4 ipv6>]	Displays configuration and operating data for all VRRPv3 groups on all interfaces. The optional brief keyword limits the amount of data shown per group. The optional statistics keyword displays operating statistics associated with configured VRRPv3 interfaces. The optional interface keyword limits the output to VRRPv3 instances on the specified interface. The optional group parameter limits the output to a single group on a specified interface. The optional ipv4 and ipv6 parameter limits the output to a single address family.
clear vrrpv3 counters [interface <interface>] [group <vrid>] <ipv4 ipv6>]	Clears the VRRPv3 statistics counters for all VRRPv3 groups on all interfaces. Optionally clears the VRRPv3 statistics counters for a single interface (if specified) and a single group (if specified). Optionally clears the statistic counters for a single address family (if specified).

Table 3. Enable Mode Troubleshooting Commands (Continued)

Command	Description
debug vrrpv3 [error] [interface <interface> [group <vrid> <ipv4 ipv6>] [packet]	Displays all VRRPv3 debug messages for all groups on all interfaces. The optional error keyword limits the output to VRRPv3 errors. The optional interface keyword limits the output to VRRPv3 instances on the specified interface. The optional group keyword limits output to VRRPv3 instances of a specified group on the interface and optionally a specified address family. The optional packet keyword limits the output to a trace of VRRPv3 router packets.

Show Commands

The **show vrrpv3** commands are used to display configuration, operation, and statistical data for VRRPv3 routers. Output from this command can be limited to display a brief amount of data for all VRRPv3 groups (using the **brief** keyword), to show statistics for all VRRPv3 groups (using the **statistics** keyword), to show only information for a specified interface (using the **interface <interface>** keywords), or to only show information for a specific group on an interface (using the **group <vrid>** keywords in conjunction with the **interface <interface>** keywords). Interfaces are specified in the **<interface type> <slot/port | interface id>** format, for example, for an Ethernet interface enter **ethernet 0/1**. Groups are specified by VRID. Valid range is **1** to **255**. You can also optionally limit output for this command to a single address family using the optional **ipv4** and **ipv6** keywords.

To display information for all configured VRRPv3 instances, enter the command from the Enable mode as follows:

>enable

#show vrrpv3

eth 0/1

Group 1 - Address-Family IPv6

State: Master

Administrative state: UP

Description:

Configured Priority: 100, Actual Priority: 100

Number of Addresses: 1

Virtual Link-Local Address: FE80::7890

Virtual Global Address: 22::1

Virtual MAC Address: 00:00:5E:00:02:01

Accept-Mode is enabled

Advertisement interval: 1 second(s)

Preemption: Enabled - delay 0 second(s)

Last Transition: 0:00:00:02

Master Router Address: FE80::2A0:C8FF:FE23:21E0 (local) Priority: 100

Group 2 - Address-Family IPv4

State: Master

Administrative State: UP

Description:
Configured Priority: 100, Actual Priority: 100
Number of Addresses: 1
Virtual IP Address: 10.0.0.254
Virtual IP Address: 22.22.22.2
Virtual MAC Address: 00:00:5E:00:01:01
Accept-Mode is enabled
Advertisement interval: 1 second(s)
Preemption: Enabled - delay 0 second(s)
Last Transition: 0:00:00:02
Master Router Address; 10.0.0.254 (local) Priority: 100

To display statistical information for all configured VRRPv3 instances, enter the command from Enable mode as follows:

>enable

#show vrrpv3 statistics

Global statistics:

Unknown group vrids: 38
Invalid checksum: 0
Invalid version: 0
Invalid packet length: 0
Invalid header type: 0
Receive advertisement purges: 0

eth 0/1

Group 1 - Address-Family IPv6

Became Master:0
Priority Zero Packets Sent: 0
Priority Zero Packets Received: 0
Advertisements Sent: 0
Advertisements Received: 0
Advertisements Interval Errors: 0
Advertisements TTL Errors: 0
Advertisements Address List Errors: 0
Advertisements Packet Length Errors: 0

Group 2 - Address-Family IPv4

Became Master: 0
Priority Zero Packets Sent: 0
Priority Zero Packets Received: 0
Advertisements Sent: 0
Advertisements Received: 0
Advertisements Interval Errors: 0
Advertisements TTL Errors: 0
Advertisements Address List Errors: 0
Advertisements Packet Length Errors: 0

Clear Commands

The **clear vrrpv3 counters** command is used to clear all IPv6 VRRPv3 statistics on the router. You can limit cleared statistics to a single interface using the optional **interface** *<interface>* parameter, or you can limit cleared statistics to a single group on an interface using the optional **group** *<vrid>* parameter in conjunction with the **interface** *<interface>* parameter. Interfaces are specified in the *<interface type>* *<slot/port | interface id>* format, for example, for an Ethernet interface enter **ethernet 0/1**. Groups are specified by VRID. Valid range is **1** to **255**. Additionally you can clear statistics for a specific address family by specifying **ipv4** or **ipv6**.

To clear all VRRPv3 statistics, enter the command from Enable mode as follows:

```
>enable
#clear vrrpv3 counters
```

Debug Commands

The **debug vrrpv3** commands display debug information associated with all IPv6 VRRPv3 instances on the router. Debug information can be limited to a specific interface (using the **interface** *<interface>* keywords), to a specific group on an interface (using the **group** *<vrid>* keywords in conjunction with the **interface** *<interface>* keywords), to a specific group's address family (using the **ipv4** or **ipv6** keywords in conjunction with the group VRID), to VRRPv3 packet information (using the **packet** keyword), or to VRRPv3 error information (using the **error** keyword). Interfaces are specified in the *<interface type>* *<slot/port | interface id>* format, for example, for an Ethernet interface enter **ethernet 0/1**. Groups are specified by VRID. Valid range is **1** to **255**.

To enable debug information for all IPv6 VRRPv3 instances, enter the command from Enable mode as follows:

```
>enable
#debug vrrpv3
```



Turning on a large amount of debug information can adversely affect the performance of your unit.

Configuration Issues in VRRPv3

In certain cases, VRRPv3 can be configured incorrectly so that it does not function as expected. The following scenarios can be used to troubleshoot possible misconfigurations of VRRPv3, and to understand the effects of not configuring VRRPv3 properly.

Different Virtual IP Address with the Same VRID

If VRRPv3 is configured on two routers on the same LAN and with the same VRID but with different primary virtual addresses, then both routers will advertise different default routers when they are in the master state. Because the virtual routers do not agree on the primary virtual address, they would discard each other's advertisements and both assume the master state. Any IPv6 hosts using the virtual router as their default gateway can experience a temporary delay in forwarding packets after a VRRP transition due to waiting for NUD to declare the existing router address to be unreachable and purging it. In the case of

IPv4 virtual routers, IPv4 hosts will try to use whichever virtual default router address they were configured to use. In this scenario, if the **debug vrrpv3** command is entered, the router will generate a log message with each received advertisement packet. The message states that the received virtual IP address does not match the configuration. The router also increments an error statistic for this condition.

Same Virtual IP Address with Different VRID on Two Routers

If VRRPv3 is configured on two routers, on the same LAN and with the same virtual IP address, but with different VRID values, each router can assume the master state. If the **debug vrrpv3** command is entered, the router generates a log message with each received advertisement packet. The message states that the received VRRPv3 advertisement has an unknown VRID. The router also increments an error statistic for this condition.

Same Virtual IP Address with Different VRID on a Single Interface

If VRRPv3 is configured on two virtual routers on a single interface with the same virtual IP address, but with different VRID values, an error is created. The error message states that the IP address is already in use by another VRID. This message is not created, however, if the virtual addresses are link-local addresses on different interfaces. There is also no error generated when the virtual addresses are the same but in different virtual routing and forwarding (VRF) instances.

Mismatched Address Families

If a configured virtual router receives an IPv4 VRRPv3 packet, and it is only configured for IPv6 VRRPv3, it will ignore the IPv4 VRRPv3 packet.

Mismatched VRRP Versions

If a configured virtual router receives a VRRPv2 packet on an interface configured only for VRRPv3, then no action is taken and an error is not reported.