



High Availability Feature Guide for the NFX250 Network Services Platform

Release

15.1X53



Modified: 2016-04-06

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

High Availability Feature Guide for the NFX250 Network Services Platform
15.1X53
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	ix
	Documentation and Release Notes	ix
	Using the Examples in This Manual	ix
	Merging a Full Example	x
	Merging a Snippet	x
	Documentation Conventions	xi
	Documentation Feedback	xiii
	Requesting Technical Support	xiii
	Self-Help Online Tools and Resources	xiii
	Opening a Case with JTAC	xiv
Part 1	Virtual Router Redundancy Protocol	
Chapter 1	Configuring Virtual Router Redundancy Protocol	3
	Understanding VRRP	3
	Overview of VRRP	3
	Sample VRRP Topology	4
	Example: Configuring VRRP for Load Sharing	6
	Configuring Basic VRRP Support for QFX	11
	Configuring VRRP Authentication (IPv4 Only)	12
	Configuring the Startup Period for VRRP Operations	13
	Configuring the Advertisement Interval for the VRRP Master	13
	Modifying the Advertisement Interval in Seconds	14
	Modifying the Advertisement Interval in Milliseconds	14
	Configuring VRRP Preemption and Hold Time	15
	Configuring VRRP Preemption	15
	Configuring the Preemption Hold Time	15
	Overriding the Hold Time	15
	Configuring a Route to Be Tracked	16
	Configuring a Logical Interface to Be Tracked	17
	Configuring a Backup to Accept Packets Destined for the Virtual IP Address	18
	Configuring Passive ARP Learning for VRRP Backups	19
	Configuring the Silent Period	20
	Configuring Inheritance for a VRRP Group	20
	Troubleshooting VRRP	21
Part 2	Configuration Statements and Operational Commands	
Chapter 2	Configuration Statements (VRRP)	25
	accept-data	26
	advertise-interval	27

	asymmetric-hold-time	28
	authentication-key	29
	authentication-type	30
	bandwidth-threshold	31
	failover-delay	32
	fast-interval	33
	hold-time (VRRP)	34
	interface	35
	preempt (VRRP)	36
	priority (Protocols VRRP)	37
	priority-cost (VRRP)	38
	priority-hold-time	39
	route (Interfaces)	40
	startup-silent-period	41
	traceoptions	42
	track (VRRP)	44
	virtual-address	45
	vrrp-group	46
Chapter 3	Operational Mode Commands (VRRP)	49
	show vrrp	50

List of Figures

Part 1	Virtual Router Redundancy Protocol	
Chapter 1	Configuring Virtual Router Redundancy Protocol	3
	Figure 1: Basic VRRP Topology	5
	Figure 2: VRRP Load-Sharing Configuration	7

List of Tables

	About the Documentation	ix
	Table 1: Notice Icons	xi
	Table 2: Text and Syntax Conventions	xi
Part 1	Virtual Router Redundancy Protocol	
Chapter 1	Configuring Virtual Router Redundancy Protocol	3
	Table 3: Settings for VRRP Load-Sharing Example	7
	Table 4: Interface State and Priority Cost Usage	18
Part 2	Configuration Statements and Operational Commands	
Chapter 3	Operational Mode Commands (VRRP)	49
	Table 5: show vrrp Output Fields	50

About the Documentation

- Documentation and Release Notes on page ix
- Using the Examples in This Manual on page ix
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiii

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xsl; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xi defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>

- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Virtual Router Redundancy Protocol

- [Configuring Virtual Router Redundancy Protocol on page 3](#)

CHAPTER 1

Configuring Virtual Router Redundancy Protocol

- [Understanding VRRP on page 3](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 12](#)
- [Configuring the Startup Period for VRRP Operations on page 13](#)
- [Configuring the Advertisement Interval for the VRRP Master on page 13](#)
- [Configuring VRRP Preemption and Hold Time on page 15](#)
- [Configuring a Route to Be Tracked on page 16](#)
- [Configuring a Logical Interface to Be Tracked on page 17](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 18](#)
- [Configuring Passive ARP Learning for VRRP Backups on page 19](#)
- [Configuring the Silent Period on page 20](#)
- [Configuring Inheritance for a VRRP Group on page 20](#)
- [Troubleshooting VRRP on page 21](#)

Understanding VRRP

Supported Platforms [EX4600, OCX1100, QFabric System, QFX Series standalone switches](#)

Juniper Networks switches support the Virtual Router Redundancy Protocol (VRRP) and VRRPv3 (for IPv6). This topic covers:

- [Overview of VRRP on page 3](#)
- [Sample VRRP Topology on page 4](#)

Overview of VRRP

Configuring end hosts on your network with static default routes minimizes configuration effort and complexity and reduces processing overhead on the end hosts. When hosts are configured with static routes, the failure of the default gateway normally results in a catastrophic event, isolating all hosts that are unable to detect available alternate paths

to their gateway. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for end hosts if the primary gateway fails.

VRRP (defined in RFC 3768) provides dynamic failover of IP addresses from one router to another in the event of failure. You can implement VRRP to provide a highly available default path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.

Switches configured with VRRP share a virtual IP address, which is the address you configure as the default route on the hosts. At any time, one of the switches is the VRRP master, meaning that it owns the virtual IP address and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure (**1 through 255**). If the master fails, the backup switch with the highest priority becomes the master within a few seconds. This is done without any interaction with the hosts.

In VRRP operation, the master sends advertisements to the backup switches at regular intervals. The default interval is 1 second. If the backup switches do not receive an advertisement for a set period, the backup with the highest priority takes over as master within a few seconds and begins forwarding packets. This is done without any interaction with the hosts.



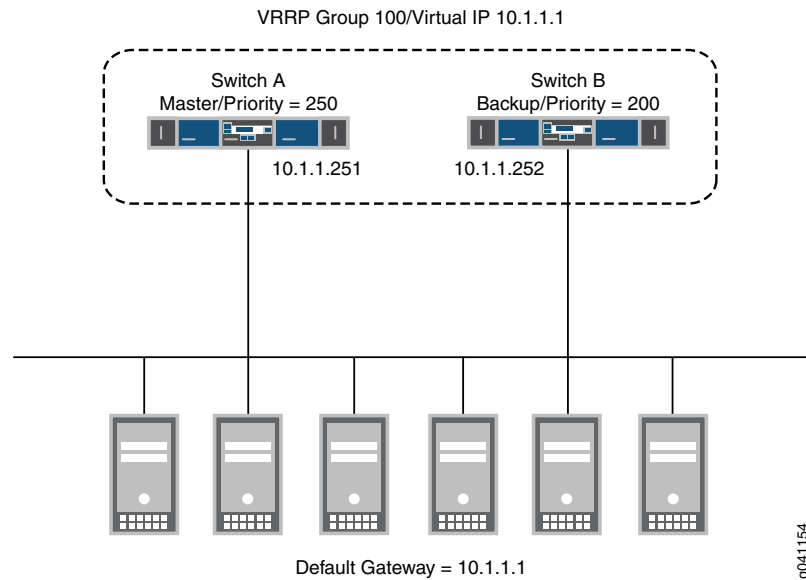
NOTE: Priority 255 cannot be set for routed VLAN interfaces (RVIs).

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. One benefit of this configuration is if you use VMware's vMotion, virtual machines can transition between hosts connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a host connected to a QFabric system in data center A can transition to a host connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address.

Sample VRRP Topology

[Figure 1 on page 5](#) illustrates a basic VRRP topology. In this example, switches A and B are running VRRP and share the virtual IP address 10.1.1.1. The default gateway for each of the clients is 10.1.1.1.

Figure 1: Basic VRRP Topology



The following illustrates basic VRRP behavior using [Figure 1 on page 5](#) for reference:

1. When any of the servers wants to send traffic out of the LAN, it sends the traffic to the default gateway address of 10.1.1.1. This is a virtual IP address (VIP) owned by VRRP group 100. Because switch A is the master of the group, the VIP is associated with the “real” address 10.1.1.251 on switch A, and traffic from the servers is actually sent to this address. (Switch A is the master because it has been configured with a higher priority value.)
2. If there is a failure on switch A that prevents it from forwarding traffic to or from the servers—for example, if the interface connected to the LAN fails—switch B becomes the master and assumes ownership of the VIP. The servers continue to send traffic to the VIP, but because the VIP is now associated with the “real” address 10.1.1.252 on switch B (because of change of master), the traffic is sent to switch B instead of switch A.
3. If the problem that caused the failure on switch A is corrected, switch A becomes the master again and reasserts ownership of the VIP. In this case, the servers resume sending traffic to switch A.

Notice that no configuration changes are required on the servers for them to switch between sending traffic to switch A and switch B. When the VIP moves between 10.1.1.251 and 10.1.1.252, the change is detected by normal TCP-IP behavior and no configuration or intervention is required on the servers.

Related Documentation

- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

Example: Configuring VRRP for Load Sharing

Supported Platforms [EX4600](#), [OCX1100](#), [QFabric System](#), [QFX Series standalone switches](#)

If you do not want to dedicate a switch to be a VRRP backup (and therefore leave it idle unless the master fails), you can create a load-sharing configuration in which each participating switch simultaneously acts as a master and a backup.

One reason to use a load-sharing (active-active) configuration is that you are more likely to actively monitor and maintain both switches and notice if a problem occurs on either of them. If you use a configuration in which one switch is only a backup (an active-backup configuration), you might be less likely to pay attention to the backup switch while it is idle. In the worst case, this could lead to the backup switch developing an undetected problem and not being able to perform adequately when a failover occurs.

- [Requirements on page 6](#)
- [Overview and Topology on page 6](#)
- [Configuring VRRP on Both Switches on page 7](#)
- [Verification on page 10](#)

Requirements

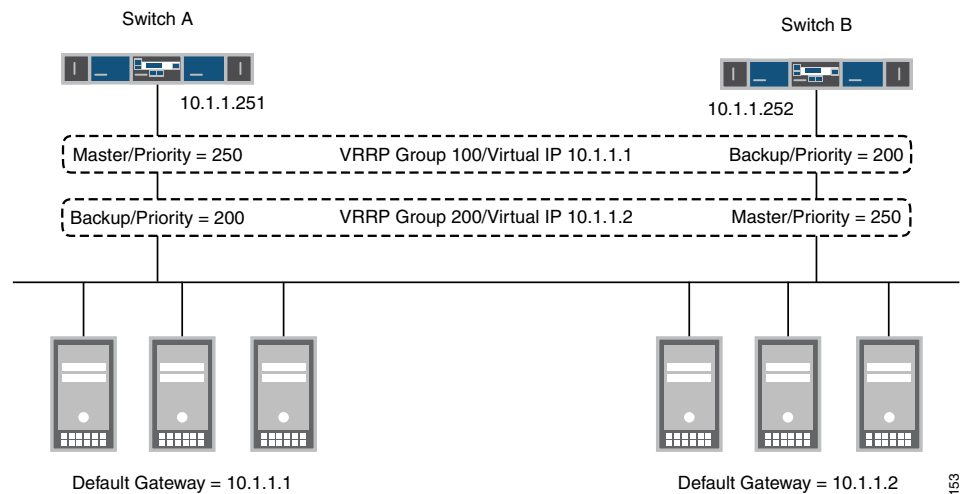
This example uses the following hardware and software components:

- Two switches
- Junos OS Release 11.3 or later
- Static routing or a dynamic routing protocol enabled on both switches.

Overview and Topology

This example uses two VRRP groups, each of which has its own virtual IP address. Devices on the LAN use one of these virtual IP addresses as their default gateway. If one of the switches fails, the other switch takes over for it. In the topology shown in [Figure 2 on page 7](#), for example, Switch A is the master for VRRP group 100. If Switch A fails, Switch B takes over and forwards traffic that the end devices send to the default gateway address 10.1.1.1.

Figure 2: VRRP Load-Sharing Configuration



9041153

This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 3 on page 7](#) lists VRRP settings for each switch.

Table 3: Settings for VRRP Load-Sharing Example

Switch A	Switch B
VRRP Group 100: <ul style="list-style-type: none"> Interface address: 10.1.1.251 VIP: 10.1.1.1 Priority: 250 	VRRP Group 100: <ul style="list-style-type: none"> Interface address: 10.1.1.252 VIP: 10.1.1.1 Priority: 200
VRRP Group 200: <ul style="list-style-type: none"> Interface address: 10.1.1.251 VIP: 10.1.1.2 Priority: 200 	VRRP Group 200: <ul style="list-style-type: none"> Interface address: 10.1.1.252 VIP: 10.1.1.2 Priority: 250

In addition to configuring the two switches as shown, you must configure your end devices so that some of them use one of the virtual IP addresses as their default gateway and the remaining end devices use the other virtual IP address as their default gateway.

Note that if a failover occurs, the remaining switch might be unable to handle all of the traffic, depending on the demand.

Configuring VRRP on Both Switches

CLI Quick Configuration

Enter the following on Switch A:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 priority 200
```

Enter the following on Switch B:

```
[edit]
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 priority 250
```

Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch A:

1. Create VRRP group 100 on Switch A and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 200
```

Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch B:

1. Create VRRP group 100 on Switch B and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
```

Switch A remains the master for group 100 because it has the highest priority for this group.

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
```

```
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
```

Switch B becomes the master for group 200 because it has the highest priority for this group.

Results Display the results of the configuration on Switch A:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.251 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 250
          }
        }
        vrrp-group 200 {
          virtual address 10.1.1.2
          priority 200
        }
      }
    }
  }
}
```

Display the results of the configuration on Switch B:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.252 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 200
          }
        }
        vrrp-group 200 {
          virtual address 10.1.1.2
          priority 250
        }
      }
    }
  }
}
```

Verification

- [Verifying that VRRP Is Working on Switch A on page 10](#)
- [Verifying that VRRP Is Working on Switch B on page 10](#)

Verifying that VRRP Is Working on Switch A

Purpose Verify that VRRP is active on Switch A and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch A and that the switch is master for group 100 and backup for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	master	A .0327 1c1 10.1.1.251 vip 10.1.1.1	
xe-0/0/0.0	up	200	backup	A .0327 1c1 10.1.1.251 vip 10.1.1.2	

Meaning The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 200 does not arrive before the timer expires, Switch A asserts itself as the master for this group.

Verifying that VRRP Is Working on Switch B

Purpose Verify that VRRP is active on Switch B and that the master and backup roles are correct.

Action Use the following command to verify that VRRP is active on Switch B and that the switch is backup for group 100 and master for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	backup	A .0327 1c1 10.1.1.252 vip 10.1.1.1	
xe-0/0/0.0	up	200	master	A .0327 1c1 10.1.1.252 vip 10.1.1.2	

Meaning The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 100 does not arrive before the timer expires, Switch B asserts itself as the master for this group.

- Related Documentation**
- [Understanding VRRP on page 3](#)
 - [Configuring Basic VRRP Support for QFX on page 11](#)

Configuring Basic VRRP Support for QFX

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

To configure basic VRRP support, configure VRRP groups on interfaces by including the **vrrp-group** statement:

```
vrrp-group group-id {  
    priority number;  
    virtual-address [ addresses ];  
}
```

An interface can be a member of multiple VRRP groups.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

For each interface, you must configure the following:

- Group identifier—Assign a value from 0 through 255. You must use the same identifier for each switch in the VRRP group.
- Priority—Assign a value from 1 through 255. The switch with the highest priority becomes the VRRP master. Assign different priorities to each switch in the VRRP group. If there are two or more switches with the same priority, the switch with the VRRP interface that has the highest IP address becomes the master.
- Virtual IP address—Normally, you configure only one address per group, but you can configure as many as eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
 - You must configure the same address on all the switches in the VRRP group.
 - If you configure a virtual IP address to be the same as a physical interface address, the switch with that interface becomes the master for the group. You must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
 - If the virtual IP address is not the same as the physical interface address, you must ensure that the address does not appear anywhere else in the switch configuration. For example, verify that you do not use this address for another interface (including an aggregated Ethernet interface) or for a static ARP entry.



NOTE: If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the `source-address-filter` statement at the `[edit interfaces interface-name]` hierarchy. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

- Related Documentation**
- [Understanding VRRP on page 3](#)
 - [Configuring the Startup Period for VRRP Operations on page 13](#)
 - [Configuring VRRP Authentication \(IPv4 Only\) on page 12](#)

Configuring VRRP Authentication (IPv4 Only)

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted switches participate in a VRRP group. By default, VRRP authentication is disabled. You can configure one of the following authentication methods for a group, and each switch in the same group must use the same method:

- Simple authentication—Uses a text password included in the transmitted packet. The receiving switch uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Adds an authentication header (AH) to the IP packet that encapsulates the VRRP packet. You create an authentication key that is used to create a hash of the packet, and the hash is stored in the AH. A receiving switch recalculates the hash on the incoming packet and compares the hashes. If they are identical, the packet is valid and is accepted. Otherwise the switch drops the incoming packet.

To enable authentication and specify an authentication method, include the **authentication-type** statement.

authentication-type *authentication;*

authentication can be **simple** or **md5**. The authentication type must be the same for all the switches in the VRRP group.

You can include this statement at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

authentication-key *key;*

key (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").



NOTE: The key must be the same for all switches in the VRRP group.

You can include this statement at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

Related Documentation

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)

Configuring the Startup Period for VRRP Operations

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

Configure the startup-silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets while an interface is coming online. The period starts when the state of a VRRP interface is changed from down to up. During this period, Master Down Events are ignored.

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```

Related Documentation

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

Configuring the Advertisement Interval for the VRRP Master

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

By default, the master switch sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master switch is still operational. If the master switch fails or becomes unreachable, the backup switch with the highest priority value becomes the new master switch.

You can modify the advertisement interval in seconds or in milliseconds; the interval must be the same for all routing platforms in the VRRP group.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 14](#)
- [Modifying the Advertisement Interval in Milliseconds on page 14](#)

Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

```
advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

Modifying the Advertisement Interval in Milliseconds

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the **fast-interval** statement:

```
fast-interval milliseconds;
```

The interval can be from 100 through 999 milliseconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]



NOTE: Junos OS sets the advertisement interval to 0 in VRRP packets. When you configure VRRP with other vendors' equipment, the **fast-interval** statement works correctly only when the other equipment also has the advertisement interval set to 0 in the VRRP packet. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

Related Documentation

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

Configuring VRRP Preemption and Hold Time

Supported Platforms EX4600, EX9200, M Series, QFabric System, QFX Series standalone switches, T Series

- [Configuring VRRP Preemption on page 15](#)
- [Configuring the Preemption Hold Time on page 15](#)
- [Overriding the Hold Time on page 15](#)

Configuring VRRP Preemption

By default, a higher-priority VRRP backup switch preempts a lower-priority master switch. To explicitly enable this behavior, include the following statement:

```
preempt;
```

To prohibit a higher-priority VRRP backup switch from preempting a lower-priority master switch, include the following statement on the lower-priority switch:

```
no-preempt;
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

Configuring the Preemption Hold Time

You can also configure a preemption hold time, which is the number of seconds a higher-priority backup router that has just started up waits before preempting the master router. You might want to configure a hold time so that routing protocols or other Junos OS components converge before preemption occurs.

The hold time is applied only on startup. By default, the hold-time value is 0 seconds, meaning that preemption can occur immediately after the backup router starts up.

To modify the preemption hold-time value, configure the following statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address vrrp-group *group-id*] preempt

Overriding the Hold Time

You can use the `asymmetric-hold-time` statement to configure a VRRP master to fail over to the backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.

When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.

You can include this statement at the following hierarchy level:

- `[edit protocols vrrp]`

**Related
Documentation**

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

Configuring a Route to Be Tracked

Supported Platforms EX4600, QFabric System, QFX Series standalone switches

A VRRP master can track a route and dynamically trigger a new master router election if the route becomes unreachable. To enable this behavior, you must configure a cost that will be subtracted from the priority of the master if the tracked route becomes unreachable. The new priority must be less than the priority of one of the backups so that the backup becomes the new master.

To configure a route to be tracked, include the following statements:

```
track {  
  priority-hold-time seconds;  
  route prefix/prefix-length routing-instance default priority-cost priority;  
}
```

You can include these statements at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

The **prefix** and **prefix-length** values specify the route to be tracked. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority of the master changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **priority-cost** option is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).

**Related
Documentation**

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

- [Configuring a Logical Interface to Be Tracked on page 17](#)

Configuring a Logical Interface to Be Tracked

Supported Platforms EX4600, QFabric System, QFX Series standalone switches

VRRP can track whether a logical interface is up, down, or not present, and can change the priority of the switch based on the state of the interface, which might trigger a new master election. VRRP can also track the operational speed of a logical interface and update the priority of the switch when the speed crosses a configured threshold. For each VRRP group, you can track as many as 10 logical interfaces.

When interface tracking is enabled on a switch, you cannot assign the switch a priority of 255 to make it the master for the group.

To configure a logical interface to be tracked, include the following statements:

```
track {
  interface interface-name {
    bandwidth-threshold bits-per-second priority-cost priority;
    priority-cost priority;
  }
  priority-hold-time seconds;
}
```

You can include these statements at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]`

The interface specified is the interface to be tracked for the VRRP group. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **bandwidth-threshold** statement specifies a threshold for the tracked interface. If the bandwidth of the tracked interface drops below the threshold value, the system subtracts the bandwidth threshold **priority-cost** value from the VRRP priority for the switch. You can create as many as five **bandwidth-threshold** statements for each tracked interface.

The interface **priority-cost** statement is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).



WARNING: On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface

belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

If you configure tracking for more than one interface, Junos OS subtracts the sum of the priority costs for the tracked interfaces from the VRRP priority if all the tracked interfaces fail. However, if you configure the interface **priority-cost** statement and the bandwidth threshold **priority-cost** statement, they are not added together. The switch uses only one priority cost for a tracked interface, as indicated in [Table 4 on page 18](#):

Table 4: Interface State and Priority Cost Usage

Tracked Interface State	Priority Cost Usage
Down	priority cost <i>priority</i>
Not down; media speed below one or more bandwidth thresholds	Priority-cost of the lowest applicable bandwidth threshold

You must configure an interface priority cost only if you do not configure any bandwidth thresholds. If you do not configure an interface **priority-cost** value and the interface fails, Junos OS subtracts the bandwidth threshold **priority-cost** value of the lowest bandwidth threshold from the priority of the switch.

Related Documentation

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)
- [Configuring a Route to Be Tracked on page 16](#)

Configuring a Backup to Accept Packets Destined for the Virtual IP Address

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

By default, a switch configured to be a VRRP backup but acting as the master does not process packets sent to the virtual IP address—that is, packets in which the destination address is the virtual IP address. To configure a backup switch to process packets sent to the virtual IP address while it is acting as the master, include the **accept-data** statement on the backup:

```
accept-data;
```

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group] *group-id***

To explicitly prohibit the backup from accepting packets destined for the virtual IP address while acting as master, include the **no-accept-data** statement:

```
no-accept-data;
```

If you include the **accept-data** statement, configure the connected hosts so that they:

- Process gratuitous ARP requests.
- Do not use packets other than ARP replies to update their ARP cache.

This statement is disabled by default. If you enable it, your configuration does not comply with RFC 3768.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.

Related Documentation

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

Configuring Passive ARP Learning for VRRP Backups

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

By default, VRRP backup switches drop ARP requests for the MAC address of the VRRP IP. This means that backups do not learn the ARP mappings (IP address to MAC address mappings) for the hosts sending the requests. If it becomes the master, the configured backup must learn all the entries that were present in the ARP cache of the original master. In environments with many directly attached hosts, the number of ARP entries to learn can be very large. This can cause a significant delay while the backup transitions to the master state, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup to learn approximately the same contents as the ARP cache in the master, thus preventing the problem of needing to learn many ARP entries quickly. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
passive-learning;
```

We recommend setting passive learning on both the backup and master VRRP switches. Doing so prevents the need to manually configure a master that fails and becomes a backup. While a switch operates as the master, the passive learning configuration has no impact. The configuration takes effect only when a switch operates as a backup.

- Related Documentation**
- [Understanding VRRP on page 3](#)
 - [Configuring Basic VRRP Support for QFX on page 11](#)
 - [Example: Configuring VRRP for Load Sharing on page 6](#)

Configuring the Silent Period

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

When the state of a VRRP interface changes from down to up, a silent period begins. During this period, any master down events are ignored. Configure the silent period interval to avoid problems that can be caused if incoming VRRP advertisement packets are delayed or interrupted while an interface starts up.

To configure the silent period, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]  
  startup-silent-period seconds;
```

- Related Documentation**
- [Understanding VRRP on page 3](#)
 - [Configuring Basic VRRP Support for QFX on page 11](#)
 - [Example: Configuring VRRP for Load Sharing on page 6](#)

Configuring Inheritance for a VRRP Group

Supported Platforms [EX4600, QFabric System, QFX Series standalone switches](#)

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. By configuring inheritance, you can prevent VRRP groups other than the active group from sending out VRRP advertisements. When the **vrrp-inherit-from** configuration statement is included in the configuration, only the active VRRP group from which the other VRRP groups are inheriting the state sends out VRRP advertisements; the groups inheriting the state do not send any VRRP advertisements, because the state is maintained only on the group from which the state is inherited.

If the **vrrp-inherit-from** statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]**:

```
[edit interfaces interface-name unit logical-unit-number family inet address address  
  vrrp-group group-id]  
  vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, note the following conditions:

- Both inheriting groups and active groups must be on the same physical interface and logical system. However, the groups need not necessarily be on the same VLAN or logical interface.
- Both inheriting groups and active groups must be on the same routing instances; however, this limitation does not apply for groups on the integrated routing and bridging (IRB) interfaces.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

Related Documentation

- [Understanding VRRP on page 3](#)
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)

Troubleshooting VRRP

Supported Platforms [EX4600, QFX Series](#)

Problem **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

Solution Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

Related Documentation • [failover-delay on page 32](#)

PART 2

Configuration Statements and Operational Commands

- [Configuration Statements \(VRRP\) on page 25](#)
- [Operational Mode Commands \(VRRP\) on page 49](#)

CHAPTER 2


Configuration Statements (VRRP)

- [accept-data on page 26](#)
- [advertise-interval on page 27](#)
- [asymmetric-hold-time on page 28](#)
- [authentication-key on page 29](#)
- [authentication-type on page 30](#)
- [bandwidth-threshold on page 31](#)
- [failover-delay on page 32](#)
- [fast-interval on page 33](#)
- [hold-time \(VRRP\) on page 34](#)
- [interface on page 35](#)
- [preempt \(VRRP\) on page 36](#)
- [priority \(Protocols VRRP\) on page 37](#)
- [priority-cost \(VRRP\) on page 38](#)
- [priority-hold-time on page 39](#)
- [route \(Interfaces\) on page 40](#)
- [startup-silent-period on page 41](#)
- [traceoptions on page 42](#)
- [track \(VRRP\) on page 44](#)
- [virtual-address on page 45](#)
- [vrrp-group on page 46](#)

accept-data

Supported Platforms	M Series, PTX Series, QFX Series, T Series
Syntax	(accept-data no-accept-data);
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not an interface accepts packets destined for the virtual IP address:</p> <ul style="list-style-type: none"> • accept-data—Enable the interface to accept packets destined for the virtual IP address. • no-accept-data—Prevent the interface from accepting packets destined for the virtual IP address.
Default	<p>If the accept-data statement is not configured, the master router responds to ARP requests only.</p> <p>The accept-data statement has the following restrictions and limitations:</p> <ul style="list-style-type: none"> • If the master router owns the virtual IP address or if the priority of the master router is set to 255, the accept-data statement becomes inapplicable. • If the master router owns the virtual IP address, the master router responds to Internet Control Message Protocol (ICMP) message requests. • If you want to restrict the incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets. • If you include the accept-data statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, <i>Virtual Router Redundancy Protocol (VRRP)</i>).
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group</i>

advertise-interval

Supported Platforms	EX Series, M Series, OCX1100, QFabric System, QFX Series standalone switches, T Series
Syntax	advertise-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets. All routers in the VRRP group must use the same advertisement interval.
<div>  <p>NOTE: When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div>	
Options	<i>seconds</i> —Interval between advertisement packets. Range: 1 through 255 seconds Default: 1 second
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Advertisement Interval for the VRRP Master Router</i> • fast-interval on page 33 • <i>inet6-advertise-interval</i> • <i>version-3</i>

asymmetric-hold-time

Supported Platforms	EX4600, QFX Series
Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	<p>Configure a VRRP master to fail over to a backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.</p> <p>When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.</p>
Default	asymmetric-hold-time is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP Preemption and Hold Time on page 15

authentication-key

Supported Platforms	EX Series, M Series, QFX Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650, T Series
Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS 12.3X48-D10 for the SRX Series.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement. All routers in the VRRP group must use the same authentication scheme and password.



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

Options	key —Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP Authentication (IPv4 Only) • Configuring VRRP Authentication (IPv4 Only) on page 12 • authentication-type on page 30 • version-3

authentication-type

Supported Platforms	EX Series, M Series, MX Series, OCX1100, QFabric System, QFX Series standalone switches, SRX Series, T Series
Syntax	authentication-type <i>authentication</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement. All routers in the VRRP group must use the same authentication scheme and password.



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

Options	<p>authentication—Authentication scheme:</p> <ul style="list-style-type: none"> simple—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure. md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. <p>Default: none (no authentication is performed).</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring VRRP Authentication (IPv4 Only) Configuring VRRP Authentication (IPv4 Only) on page 12 authentication-key on page 29 version-3


bandwidth-threshold

Supported Platforms	M Series, PTX Series, QFX Series, T Series
Syntax	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<p><i>bits-per-second</i>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p>Range: 1 through 10000000000000 bits per second</p> <p><i>priority-cost <i>priority</i></i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, and forces a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 17

failover-delay

Supported Platforms	EX4600, QFX Series
Syntax	failover-delay <i>milliseconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	<p>If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).</p> <p>If you configure a failover delay, the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.</p>
Options	<i>milliseconds</i> —Specify the failover delay time, in milliseconds. Range: 50 through 2000
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Troubleshooting VRRP on page 21• show vrrp on page 50

fast-interval

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	<code>fast-interval milliseconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS 12.3X48-D10 for the SRX Series.</p>
Description	<p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
Options	<p>milliseconds—Interval between advertisement packets.</p> <p>Range: 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>
<div>  <p>NOTE: When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for <code>fast-interval</code>. Commit check fails if a value less than 100 is configured.</p> </div>	
Default: 1 second	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Advertisement Interval for the VRRP Master Router • Configuring the Advertisement Interval for the VRRP Master on page 13 • advertise-interval on page 27 • advertise-interval on page 27 • inet6-advertise-interval • version-3

hold-time (VRRP)

Supported Platforms	EX Series, M Series, QFX Series, T Series
Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
Default	VRRP preemption is not timed.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 seconds Default: 0 seconds (VRRP preemption is not timed.)
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Backup Router to Preempt the VRRP Master Router• Configuring VRRP Preemption and Hold Time on page 15

interface

Supported Platforms	M Series, QFX Series, T Series
Syntax	<pre>interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>bandwidth-threshold statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p>interface-name—Interface to be tracked for this VRRP group.</p> <p>Range: 1 through 10 interfaces</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 17 • Junos OS Services Interfaces Library for Routing Devices

preempt (VRRP)

Supported Platforms EX9200, M Series, OCX1100, QFabric System, QFX Series standalone switches, T Series

Syntax (preempt | no-preempt) {
 hold-time seconds;
}

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*],
[edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router:

- **preempt**—Allow the master router to be preempted.



NOTE: By default, a higher-priority backup router can preempt a lower-priority master router.

- **no-preempt**—Prohibit the preemption of the master router. When **no-preempt** is configured, the backup router cannot preempt the master router even if the backup router has a higher priority.

The remaining statement is explained separately.

Default By default the **preempt** statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the **preempt** statement is not explicitly configured.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring a Backup Router to Preempt the VRRP Master Router](#)
- [Configuring VRRP Preemption and Hold Time on page 15](#)


priority (Protocols VRRP)

Supported Platforms	EX Series, M Series, OCX1100, QFabric System, QFX Series standalone switches, T Series
Syntax	<code>priority priority;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<p>priority—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p>Range: 1 through 255</p> <p>Default: 100 (for backup routers)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support • Configuring Basic VRRP Support for QFX on page 11

priority-cost (VRRP)

Supported Platforms	ACX2000 and ACX2100, EX Series, M120, MX240, OCX1100, QFabric System, QFX Series standalone switches, SRX210, SRX3400, T1600, T640
Syntax	<code>priority-cost priority;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <code>vrrp-group group-id track interface interface-name</code>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <code>vrrp-inet6-group group-id track interface interface-name</code>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <code>vrrp-group group-id track interface interface-name</code>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <code>vrrp-inet6-group group-id track interface interface-name</code>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	priority —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. Range: 1 through 254
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked for a VRRP Group• Configuring a Logical Interface to Be Tracked on page 17

priority-hold-time

Supported Platforms	EX Series, M Series, MX Series, OCX1100, QFabric System, QFX Series standalone switches, SRX Series, T Series
Syntax	priority-hold-time <i>seconds</i> ;
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
<div>  <p>NOTE: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div>	
Options	<p><i>seconds</i>—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0through 3600 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 17

route (Interfaces)

Supported Platforms	EX Series, M120, MX240, OCX1100, QFabric System, QFX Series standalone switches, SRX210, SRX3400, T1600, T640
Syntax	route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group group-id track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group group-id track]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS 11.3 for QFX Series. Statement introduced in Junos OS 12.1 for EX Series switches.
Description	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<i>prefix</i> —Route to be tracked for this VRRP group. <i>priority-cost priority</i> —The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. <i>routing-instance instance-name</i> —Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <i>instance-name</i> must be default .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Route to Be Tracked for a VRRP Group• Configuring a Route to Be Tracked on page 16

startup-silent-period

Supported Platforms	EX Series, M Series, MX Series, OCX1100, QFabric System, QFX Series standalone switches, SRX Series, T Series
Syntax	startup-silent-period <i>seconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
Options	seconds —Number of seconds for the startup period. Default: 4 seconds Range: 1 through 2000 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Startup Period for VRRP Operations</i>• Configuring the Startup Period for VRRP Operations on page 13

traceoptions

Supported Platforms EX4600, OCX1100, QFabric System, QFX Series standalone switches

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

To specify more than one tracing operation, include multiple **flag** statements.

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.



NOTE: The traceoptions statement is not supported on a QFabric system.

Default If you do not include this statement, no VRRP-specific tracing operations are performed.

Options **filename filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

files number—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

Range: 0 through 4,294,967,296 files

Default: 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations
- **database**—Database changes
- **general**—General events

- **interfaces**—Interface changes
- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

match *regex*—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—Restrict users from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your routing platform

Default: 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Tracing VRRP Operations</i>

track (VRRP)

Supported Platforms	EX Series, M Series, OCX1100, QFabric System, QFX Series standalone switches, T Series
Syntax	<pre>track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>priority-hold-time statement added in Junos OS Release 8.1.</p> <p>route statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.</p>
Description	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
Options	The remaining statements are described separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked for a VRRP Group • Configuring a Route to Be Tracked for a VRRP Group • Configuring a Logical Interface to Be Tracked on page 17 • Configuring a Route to Be Tracked on page 16

virtual-address

Supported Platforms	EX Series, M Series, MX Series, OCX1100, QFabric System, QFX Series standalone switches, SRX Series, T Series
Syntax	virtual-address [<i>addresses</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support • Configuring Basic VRRP Support for QFX on page 11

vrrp-group

Supported Platforms	EX Series, M Series, MX Series, OCX1100, QFabric System, QFX Series standalone switches, SRX Series, T Series
Syntax	<pre> vrrp-group <i>group-id</i> { (accept-data no-accept-data); advertise-interval <i>seconds</i>; advertisements-threshold <i>number</i>; authentication-key <i>key</i>; authentication-type <i>authentication</i>; fast-interval <i>milliseconds</i>; (preempt no-preempt) { hold-time <i>seconds</i>; } priority <i>number</i>; track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; } virtual-address [<i>addresses</i>]; vrrp-inherit-from <i>vrrp-group</i>; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group. As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the nonstop-routing statement at the [edit routing-options] or [edit logical system <i>logical-system-name</i> routing-options hierarchy level.
Options	<p>group-id—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Basic VRRP Support*
- *Configuring VRRP*
- [Configuring Basic VRRP Support for QFX on page 11](#)
- [Example: Configuring VRRP for Load Sharing on page 6](#)
- *vrrp-inet6-group*
- *nonstop-routing*

CHAPTER 3

Operational Mode Commands (VRRP)

- `show vrrp`

show vrrp

Supported Platforms [EX Series, OCX1100, QFabric System, QFX Series standalone switches](#)

Syntax `show vrrp`
`<brief | detail | extensive | summary>`
`<interface interface-name>`
`<track interfaces>`

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1x53-D20 for the OCX Series.

Description Display information and status about VRRP groups.

Options **none**—(Same as brief) Display brief status information about all VRRP interfaces.
brief | detail | extensive | summary—(Optional) Display the specified level of output.
interface *interface-name*—(Optional) Display information and status about the specified VRRP interface.
track interfaces—(Optional) Display information and status about VRRP track interfaces.

Required Privilege Level view

Related Documentation

- [Configuring VRRP for IPv6 \(CLI Procedure\)](#)

List of Sample Output [show vrrp on page 55](#)
[show vrrp brief on page 55](#)
[show vrrp detail \(IPv6\) on page 55](#)
[show vrrp detail \(Route Track\) on page 56](#)
[show vrrp extensive on page 56](#)
[show vrrp interface on page 57](#)
[show vrrp summary on page 58](#)
[show vrrp track detail on page 58](#)
[show vrrp track summary on page 59](#)

Output Fields [Table 5 on page 50](#) lists the output fields for the **show vrrp** command. Output fields are listed in the approximate order in which they appear.

Table 5: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	none, brief, extensive, summary
Interface index	Physical interface index number, which reflects its initialization sequence.	extensive

Table 5: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Groups	Total number of VRRP groups configured on the interface.	extensive
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	extensive
Interface VRRP PDU statistics	Nonerrored statistics for the logical interface: <ul style="list-style-type: none"> • Advertisement sent—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • Advertisement received—Number of VRRP advertisement PDUs received by the interface. • Packets received—Number of VRRP packets received for VRRP groups on the interface. • No group match received—Number of VRRP packets received for VRRP groups that do not exist on the interface. 	extensive
Interface VRRP PDU error statistics	Errored statistics for the logical interface: <ul style="list-style-type: none"> • Invalid IPAH next type received—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • Invalid VRRP ttl value received—Number of packets received whose IP time-to-live (TTL) value is not 255. • Invalid VRRP version received—Number of packets received whose VRRP version is not 2. • Invalid VRRP pdu type received—Number of packets received whose VRRP PDU type is not 1. • Invalid VRRP authentication type received—Number of packets received whose VRRP authentication is not none, simple, or md5. • Invalid VRRP IP count received—Number of packets received whose VRRP IP count exceeds 8. • Invalid VRRP checksum received—Number of packets received whose VRRP checksum does not match the calculated value. 	extensive
Physical interface	Name of the physical interface.	detail, extensive
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	none, brief, detail, extensive
Index	Physical interface index number, which reflects its initialization sequence.	detail, extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail, extensive
VRRP-Traps	Status of VRRP traps: Enabled or Disabled .	detail, extensive

Table 5: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type and Address	Identifier for the address and the address itself: <ul style="list-style-type: none"> • lcl—Configured local interface address. • mas—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. • vip—Configured virtual IP addresses. 	none, brief, summary
Interface state or Int state	State of the physical interface: <ul style="list-style-type: none"> • down—The device is present and the link is unavailable. • not present—The interface is configured, but no physical device is present. • unknown—The VRRP process has not had time to query the kernel about the state of the interface. • up—The device is present and the link is established. 	none, brief, extensive, summary
Group	VRRP group number.	none, brief, extensive, summary
State	VRRP state: <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • initializing—VRRP is initializing. • master—The interface is acting as the master router interface. • transition—The interface is changing between being the backup and being the master router. 	extensive
Priority	Configured VRRP priority for the interface.	detail, extensive
Advertisement interval	Configured VRRP advertisement interval.	detail, extensive
Authentication type	Configured VRRP authentication type: none , simple , or md5 .	detail, extensive
Preempt	Whether preemption is allowed on the interface: yes or no .	detail, extensive
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: yes or no .	detail, extensive
VIP count	Number of virtual IP addresses that have been configured on the interface.	detail, extensive
VIP	List of virtual IP addresses configured on the interface.	detail, extensive
Advertisement timer	Time until the advertisement timer expires.	detail, extensive

Table 5: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is N/A .	detail, extensive
Virtual router uptime	Time that the virtual router has been up.	detail, extensive
Master router uptime	Time that the master router has been up.	detail, extensive
Virtual MAC	MAC address associated with the virtual IP address.	detail, extensive
Tracking	Whether tracking is enabled or disabled .	detail, extensive
Current priority	Current operational priority for being the VRRP master.	detail, extensive
Configured priority	Configured base priority for being the VRRP master.	detail, extensive
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. Disabled indicates no minimum interval.	detail, extensive
Remaining-time	(track option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	detail extensive
Interface/Tracked interface	Name of the tracked interface.	detail extensive
Int state/Interface state	Current operational state of the tracked interface: up or down .	detail, extensive
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	detail, extensive
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	detail, extensive
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of down means that the corresponding priority cost is incurred when the interface is down.	detail, extensive
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	detail, extensive
Route count	The number of routes being tracked.	detail, extensive

Table 5: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Route	The IP address of the route being tracked.	detail, extensive
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	detail, extensive
Route state	The state of the route being tracked: up , down , or unknown .	detail, extensive
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	detail, extensive
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	detail, extensive
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	extensive
Group VRRP PDU error statistics	Errored statistics for the VRRP group: <ul style="list-style-type: none"> • Bad authentication type received—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be none, simple, or md5 and must be the same for all routers in the VRRP group. • Bad password received—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group. • Bad MD5 digest received—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. • Bad advertisement timer received—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. • Bad VIP count received—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • Bad VIPADDR received—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance. 	extensive
Group state transition statistics	State transition statistics for the VRRP group: <ul style="list-style-type: none"> • Idle to master transitions—Number of times that the VRRP instance transitioned from the idle state to the master state. • Idle to backup transitions—Number of times that the VRRP instance transitioned from the idle state to the backup state. • Backup to master transitions—Number of times that the VRRP instance transitioned from the backup state to the master state. • Master to backup transitions—Number of times that the VRRP instance transitioned from the master state to the backup state. 	extensive
Vlan-id	ID of Vlan	detail

Table 5: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
VR state	VRRP information: <ul style="list-style-type: none"> • backup—The interface is acting as the backup router interface. • bringup—VRRP is just starting, and the physical device is not yet present. • idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • initializing—VRRP is initializing. • master—The interface is acting as the master router interface. • transition—The interface is changing between being the backup and being the master router. 	none, brief
Timer	VRRP timer information: <ul style="list-style-type: none"> • A—Time, in seconds, until the advertisement timer expires. • D—Time, in seconds, until the Master is Dead timer expires. 	none, brief

Sample Output

show vrrp

```

user@host> show vrrp
Interface      State      Group  VR state  Timer  Type  Address
ge-0/0/0.121   up         1      master    A 1.052 1c1  gec0::12:1:1:1
                                     vip  ge80::12:1:1:99
                                     vip  gec0::12:1:1:99
ge-0/0/2.131   up         1      master    A 0.364 1c1  gec0::13:1:1:1
                                     vip  ge80::13:1:1:99
                                     vip  gec0::13:1:1:99

```

show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 55](#).

show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
 Interface state: up, Group: 1, State: master
 Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
 gec0::13:1:1:99
 Advertisement timer: 0.327s, Master router: ge80::13:1:1:1
 Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
 Virtual MAC: 00:00:5e:00:02:01
 Tracking: disabled

show vrrp detail (Route Track)

user@host> show vrrp detail

Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24
 Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
 Interface state: up, Group: 100, State: master
 Priority: 150, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
 Advertisement timer: 1.218s, Master router: 30.30.30.30
 Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
 Virtual MAC: 00:00:5e:00:01:64
 Tracking: enabled
 Current priority: 150, Configured priority: 150
 Priority hold-time: disabled
 Interface tracking: disabled
 Route tracking: enabled, Route count: 1

Route	VRF name	Route state	Priority cost
192.168.40.0/22	default	up	30

show vrrp extensive

user@host> show vrrp extensive

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

Interface VRRP PDU statistics

Advertisement sent	:	188
Advertisement received	:	0
Packets received	:	0
No group match received	:	0

Interface VRRP PDU error statistics

Invalid IPAH next type received	:	0
Invalid VRRP TTL value received	:	0
Invalid VRRP version received	:	0
Invalid VRRP PDU type received	:	0
Invalid VRRP authentication type received	:	0
Invalid VRRP IP count received	:	0
Invalid VRRP checksum received	:	0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
 Interface state: up, Group: 1, State: master
 Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
 gec0::12:1:1:99
 Advertisement timer: 1.034s, Master router: ge80::12:1:1:1
 Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
 Virtual MAC: 00:00:5e:00:02:01
 Tracking: disabled
 Group VRRP PDU statistics

```

    Advertisement sent          :          188
    Advertisement received      :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received          :           0
    Bad MD5 digest received         :           0
    Bad advertisement timer received:         0
    Bad VIP count received          :           0
    Bad VIPADDR received           :           0
Group state transition statistics
    Idle to master transitions       :           0
    Idle to backup transitions       :           1
    Backup to master transitions     :           1
    Master to backup transitions     :           0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
    Advertisement sent            :          186
    Advertisement received         :           0
    Packets received               :           0
    No group match received        :           0
Interface VRRP PDU error statistics
    Invalid IPAH next type received :           0
    Invalid VRRP TTL value received :           0
    Invalid VRRP version received   :           0
    Invalid VRRP PDU type received  :           0
    Invalid VRRP authentication type received:         0
    Invalid VRRP IP count received  :           0
    Invalid VRRP checksum received  :           0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
    Advertisement sent            :          186
    Advertisement received         :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received          :           0
    Bad MD5 digest received         :           0
    Bad advertisement timer received:         0
    Bad VIP count received          :           0
    Bad VIPADDR received           :           0
Group state transition statistics
    Idle to master transitions       :           0
    Idle to backup transitions       :           1
    Backup to master transitions     :           1
    Master to backup transitions     :           0

```

show vrrp interface

user@host> show vrrp interface

```

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
  Packets received             :           0
  No group match received      :           0
Interface VRRP PDU error statistics
  Invalid IPAH next type received :           0
  Invalid VRRP TTL value received :           0
  Invalid VRRP version received  :           0
  Invalid VRRP PDU type received :           0
  Invalid VRRP authentication type received:           0
  Invalid VRRP IP count received :           0
  Invalid VRRP checksum received :           0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent           :          205
  Advertisement received       :           0
Group VRRP PDU error statistics
  Bad authentication type received:           0
  Bad password received         :           0
  Bad MD5 digest received       :           0
  Bad advertisement timer received:           0
  Bad VIP count received        :           0
  Bad VIPADDR received         :           0
Group state transition statistics
  Idle to master transitions     :           0
  Idle to backup transitions     :           1
  Backup to master transitions   :           1
  Master to backup transitions   :           0

```

show vrrp summary

```

user@host> show vrrp summary

```

Interface	State	Group	VR state	Type	Address
ge-4/1/0.0	up	1	backup	lcl	10.57.0.2
				vip	10.57.0.100

show vrrp track detail

```

user@host> show vrrp track detail
Tracked interface: ae1.211
State: up, Speed: 400m
Incurred priority cost: 0

```

Threshold	Priority cost	Active
400m	10	
300m	60	
200m	110	
100m	160	
down	190	

```
Tracking VRRP interface: ae0.210, Group: 1
VR State: master
Current priority: 200, Configured priority: 200
Priority hold-time: disabled,    Remaining-time: 50.351
```

show vrrp track summary

```
user@host> show vrrp track summary
```

Track if	State	Speed	VRRP if	Group	VR State	Current priority
ae1.211	up	400m	ae0.210	1	master	200

