Configuring VRRP 9

This chapter describes how to configure the Virtual Router Redundancy Protocol (VRRP) in your system.

Overview

VRRP can prevent loss of network connectivity to end hosts if the static default IP gateway fails. By implementing VRRP, you can designate a number of routers as “backup” routers in the event that the default “master” router fails.

Note: *It is important to understand that the term “virtual router” as defined in the ERX System Basics Configuration Guide, Chapter 10, Configuring Virtual Routers, is different from what is implied by VRRP. In this chapter, the term “virtual router” always refers to a VRRP router; that is, a router that has enabled VRRP.*

In case of a failure, VRRP dynamically shifts the packet-forwarding responsibility to a backup router. A redundancy scheme is created by VRRP, which allows hosts to keep a single IP address for the default gateway, but maps the IP address to a well-known virtual MAC address. VRRP provides this redundancy without user intervention or additional configuration at the end hosts.

VRRP is supported on the FE-2 line module and the GE/FE line module. The ERX system currently supports 200 virtual router IDs (VRIDs) on the FE-2 line module, and supports 800 VRIDs on the GE/FE line module. VLANs are fully supported.

Terminology

Table 9-1 provides definitions for the basic VRRP terms used in this chapter.

Table 9-1 VRRP definitions

Term	Definition
VRRP router	A router that is running VRRP. It may participate in one or more VRIDs. An IP redundancy instance can: • Act as a master with associated addresses it owns at an IP interface • Act simultaneously as a backup for other routers with additional VRID mappings and priorities for those routers
Master router	The VRRP router that assumes the responsibility of forwarding packets sent to the IP address(es) associated with the virtual router, and that answers ARP requests for these IP addresses. If the IP address owner is available, then it always becomes the master.
Backup router	The VRRP router available to assume forwarding responsibility if the current master router fails.
IP address owner	The IP interface–VRID pair instance that has the associated IP address(es) as real interface address(es). This router, when up, responds to packets addressed to one of these IP addresses for ICMP pings or TCP connections. The IP address owner is the *primary router*.
Primary IP address	An IP address configured as primary from the set of real interface addresses. VRRP advertisements are always sent (by the master router) using the primary IP address as the source of the IP packet.

References

For more information about VRRP, see:

- RFC 2338 – Virtual Router Redundancy Protocol (April 1998)
- RFC 2787 – Definitions of Managed Objects for the Virtual Router Redundancy Protocol (March 2000)

Note: *It is important to have some background understanding of the Address Resolution Protocol (ARP) before configuring VRRP. See the section, Address Resolution Protocol (ARP), in Chapter 2, Configuring IP.*

How VRRP Works

The advantage of using VRRP is that you gain a higher availability for the default path without requiring configuration of dynamic routing or router discovery protocols on every end host.

VRRP routers viewed as a “redundancy group” share the responsibility for forwarding packets as if they “owned” the IP address corresponding to the default gateway configured on the hosts. At any time, one of the VRRP routers acts as the master and other VRRP routers act as backups. If the master router fails, a backup router becomes the new master. In this way, router redundancy is always provided, allowing traffic on the LAN to be routed without relying on a single router.

There is always a master for the shared IP address. If the master goes down, an elected VRRP router forwards packets on behalf of the owner by taking over the virtual MAC address used by the owner.

When implemented in your network, VRRP assumes that if any link to a subnet is active, the router has access to the entire subnet. VRRP leverages the broadcast capabilities of Ethernet. Provided that one of the routers in a VRRP configuration is running, ARP requests for the IP addresses assigned to the default gateway always receive replies. Additionally, end hosts can send packets outside their subnet without interruption.

Configuration Examples

Three VRRP configuration examples are described and illustrated in this section. They include:

- Basic VRRP configuration
- Common VRRP configuration
- VRRP configuration without the real address owner

Basic VRRP Configuration

As shown in Figure 9-1, the basic VRRP configuration uses a single VRID (VRID 1). Because R1 is the address owner, it serves as the master. Router R2 is the backup. The four end hosts on subnet 1 are configured to use 10.0.0.1/8 as the default router. IP address 10.0.0.1 is associated with VRID 1.

In this example, if R1 becomes unavailable, R2 takes over VRID 1 and its associated IP addresses. Packets sent to IP destinations outside the 10.x.x.x subnet using 10.0.0.1 as the router are then forwarded by R2. Even though R2 assumes R1's forwarding responsibilities, it never processes any packet with destination address (DA) 10.0.0.1. When R1 becomes active again, it takes over as the master and R2 reverts to backup.

The VRRP MAC address is always: 00-00-5e-00-01-*vrid*. The valid VRID range is 0x01–0xFF.

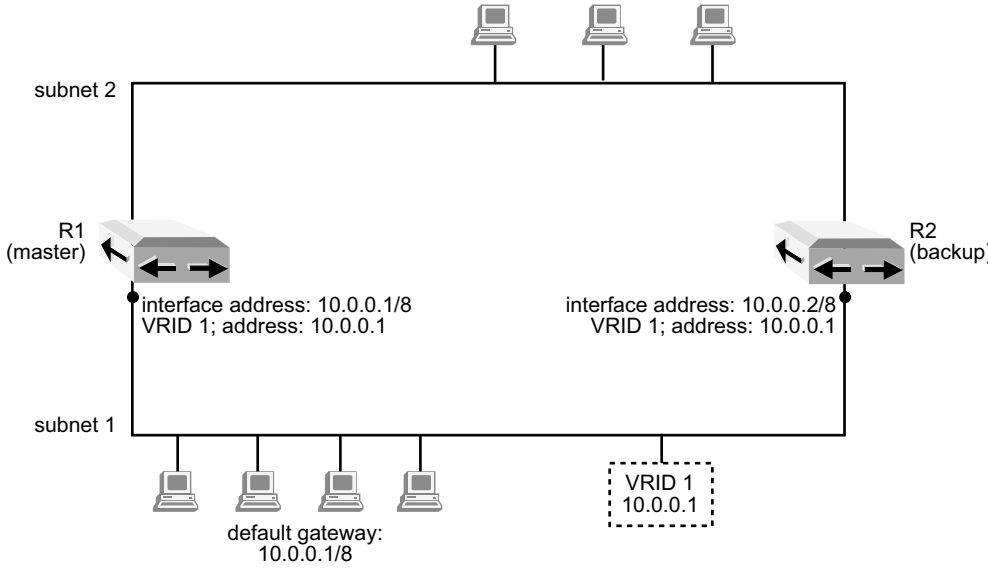

Figure 9-1 Basic VRRP configuration

Commonly Used VRRP Configuration

Figure 9-2 shows two physical routers backing each other up through VRRP. Routers R1 and R2 are both configured with VRID 1 and VRID 2. In this configuration, under normal circumstances, the routing load is distributed between the two routers.

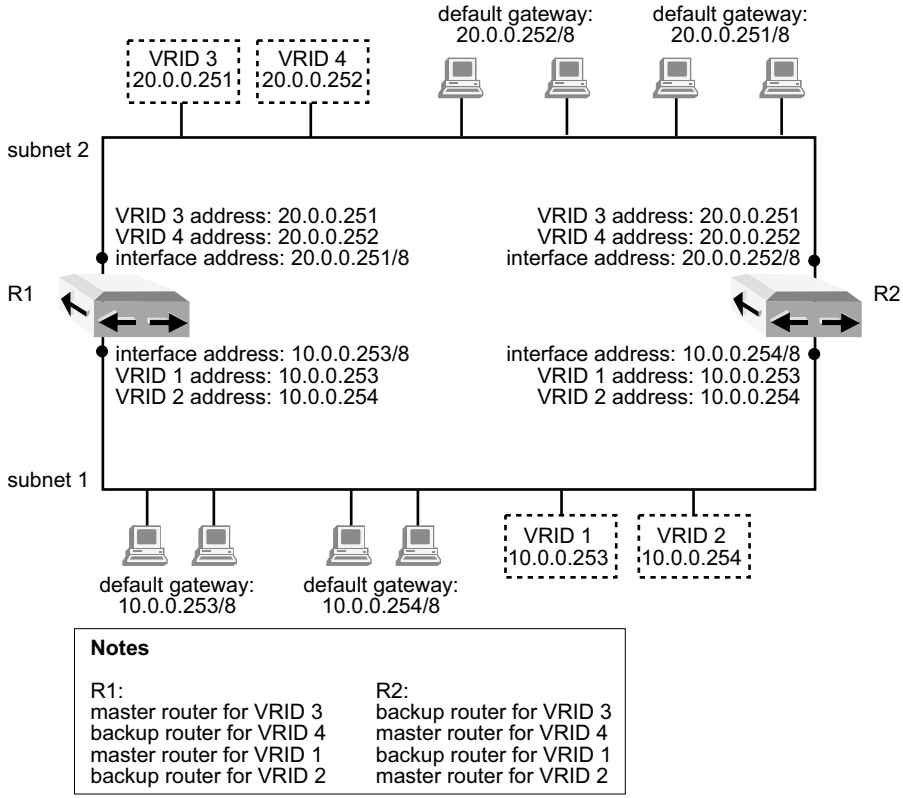

Figure 9-2 Commonly used VRRP configuration

VRRP Configuration Without the Real Address Owner

Figure 9-3 is noticeably similar to Figure 9-2 except that there is no real owner for the addresses configured by the VRIDs. Consequently, both routers R1 and R2 are configured as backup routers for VRID 1, VRID 2, VRID 3, and VRID 4.

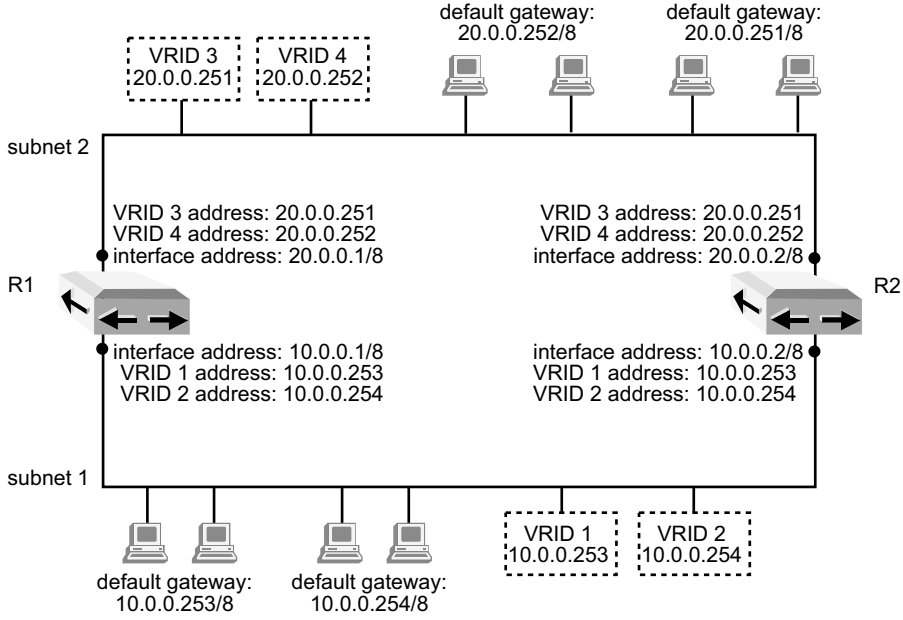

Figure 9-3 VRRP configuration without the real address owner

Assuming that preemption is enabled, the router that is configured with the highest priority for each VRID becomes the master. If priorities are the same, the router that has the highest primary address becomes the master.

This configuration shows how the address owner does not necessarily need to exist under VRRP. The configuration seems somewhat mysterious, since the PCs are never able to “ping” their gateway; however, they can talk to destinations outside their own subnet.

IP multicast packets are used by the election protocol specified in VRRP to provide the router with redundancy. Therefore, VRRP can operate over a variety of multiaccess LAN technologies supporting IP multicast. It is important to remember that there is always one master for an IP address shared by the redundancy group.

How VRRP Is Implemented in the ERX System

VRRP is implemented in the ERX system with the goal of meeting two requirements. The first requirement is to avoid the single point of failure inherent to hosts that have a single default gateway configured. The second requirement is to keep the complexity of redundancy away from the hosts themselves. These goals are in compliance with RFC 2338 and RFC 2787.

The association between VRIDs and IP addresses is coordinated among all participating VRRP routers. A scenario is provided to help you understand how VRRP is implemented in the ERX router.

1. The ERX system assigns common VRIDs to the group of routers that are going to share IP addresses.
2. The ERX system sends VRRP advertisements to well-known multicast addresses. The router that owns the addresses automatically becomes the master and sends periodic VRRP advertisement messages. A VRRP advertisement consists of the IP addresses that the master controls and the VRID.
3. If the master stops advertising for a predetermined period of time, the remaining routers using the same VRID enter an election process to determine which router takes over the master responsibilities.
4. If the master does not own the IP addresses for which it is responsible, it drops all packets that have DAs to these IP addresses.
5. If the elected master fails, backup routers start the election process again.
6. When the original master becomes operational again, it restarts broadcasting advertisements as long as preemption is enabled or the master is the address owner. Packet forwarding responsibility then shifts back to the original master router.

Router Election Rules

If the master router becomes unavailable, there are rules that govern election of the master router:

- The backup router assigned the highest priority for each VRID becomes the master.
- If two backup routers were assigned the same priority, the router that has the highest primary address becomes the master. For example, if several routers were all assigned the default priority of 100, then the IP addresses must be compared.

- Router election on a VRRP router can also be determined by whether or not the preemption option is enabled.

  When a backup router detects a master router with a lower priority than the backup has, the backup router may either choose to leave the current master alone or take over the current master and become the master itself.

  When preemption is enabled, a backup router always preempts or takes over the responsibility of the master router. When preemption is disabled, the lower-priority backup is left in the master state.

Note: *It is possible that ICMP redirect's source address may be overridden when VRRP is in use. When a backup VRID acts as a master on a given IP interface, its ICMP redirects must "fake" the source IP address of the IP address owner. The IP address must be faked because hosts accept only an ICMP redirect that was sent by the host's current gateway.*

Configuring VRRP

You must perform certain steps to configure VRRP; others are optional.

Configuring the IP Interface

Before you configure VRRP, you must configure an IP interface and assign a primary IP address and subnet mask. When the IP address belongs to the owner of the VRID, it must be associated with the VRID you create.

To configure an IP interface and IP address:

1 Configure an IP interface.

```
host1(config)#interface fastEthernet 4/0
```

2 Assign an IP address and subnet mask.

```
host1(config-if)#ip address 194.50.1.42 255.255.255.0
```

Note: *It is recommended that you complete all IP address configurations before you configure VRRP. If for any reason the IP address information is changed after you configure VRRP, you must revise the associated IP addresses configured on the related VRRP entries. If* ***auto*** *addresses are specified in the* ***ip vrrp virtual-address*** *command and they are used in conjunction with priority 255, you must disable and re-enable the VRRP entry for the association list to be updated.*

Creating VRIDs

A master or backup router running the VRRP protocol may participate in one or more VRID instances. There are several ways to create a VRID instance:

- It is recommended that you create and/or configure a VRID instance first, and enable it last. For example:

```
host1(config-if)#ip vrrp 198
host1(config-if)#ip vrrp 198 priority 255
```

- You can create and enable a VRID instance by using the **ip vrrp *vrid* enable** command, as shown below.

```
host1(config-if)#ip vrrp 25 enable
```

- You continue to configure the VRID by identifying the VRID each time a VRRP command is used. For example:

```
host1(config-if)#ip vrrp 175 authentication-type none
```

- Alternatively, you can create a new VRID when configuring any VRRP command, providing it is the first time the VRID instance is used. For example, if you execute the command **ip vrrp *vrid* preempt** and it is the first time that VRID is used, a new VRID is created.

```
host1(config-if)#ip vrrp 16 preempt
```

- Use the ip **vrrp *vrid* enable** command last. The new VRID is not enabled until you execute this command.

```
host1(config-if)#ip vrrp 198 enable
host1(config-if)#ip vrrp 16 enable
host1(config-if)#ip vrrp 175 enable
```

Configuration Steps

Before configuring VRRP, you may find it helpful to review the configuration examples in the earlier section *How VRRP Works*.

To configure VRRP parameters:

1 (Optional) Create a VRID instance.

```
host1(config-if)#ip vrrp 25
```

2 (Optional) Set a VRRP advertisement interval for the same VRID.

```
host1(config-if)#ip vrrp 25 advertise-interval 50
```

3 Set the VRRP router priority for owner and/or backup router(s).

This step is mandatory to configure priority for the owner VRID (255). This step is optional to configure priority for a backup VRID (1–254). The default is 100.

```
host1(config-if)#ip vrrp 25 priority 255
host1(config-if)#ip vrrp 22 priority 254
```

4 (Optional) Set the preemption option. In this example a new VRID is created.

```
host1(config-if)#ip vrrp 10 preempt
```

5 Associate an IP address with a VRID.

```
host1(config-if)#ip vrrp 25 virtual-address 194.2.1.63
  255.255.255.0
```

6 (Optional) Set the VRRP authentication type to either **text** or **none**.

```
host1(config-if)#ip vrrp 25 authentication-type none
```

7 (Optional) Configure the VRRP authentication key.

```
host1(config-if)#ip vrrp 25 authentication-key dublin
```

8 Enable the VRID instance.

```
host1(config-if)#ip vrrp 25 enable
```

ip vrrp

- Use to create a VRID instance.
- The VRID range is between 1–255.
- Example

```
host1(config-if)#ip vrrp 25
```

- Use the **no** version to remove a VRID instance.

ip vrrp advertise-interval

- Use to configure the advertisement interval time.
- The interval time can be configured for seconds or milliseconds.
- Use seconds to be in compliance with RFC 2338.
- If your VRID environment consists of only ERX systems, you can optionally use milliseconds.
- Example

```
host1(config-if)#ip vrrp 25 advertise-interval 50
```

- Use the **no** version to restore the default, 1 second.

ip vrrp authentication-key

- Use to specify the authentication key.
- Use the **key** keyword only when the authentication type is **text**. See **ip vrrp authentication-type**.
- Example

```
host1(config-if)#ip vrrp 25 authentication-key dublin
```

- Use **no ip vrrp authentication-key** to set the authentication key to its default, empty string.

ip vrrp authentication-type

- Use to specify the authentication type of either **text** or **none**.
- Example

```
host1(config-if)#ip vrrp 175 authentication-type none
```

- Use **no ip vrrp authentication-type** to set the authentication type to its default, **none**.

ip vrrp enable

- Use to enable an existing VRID instance.
- The VRID range is between 1–255.
- The default is disable.
- Example

```
host1(config-if)#ip vrrp 175 enable
```

- Use the **no** version to disable an existing VRID instance.

ip vrrp preempt

- Use to enable preemption.
- When preemption is enabled, a backup router always takes over the responsibility of the master router. When preemption is disabled, the lower-priority backup is left in the master state.
- Example

```
host1(config-if)#ip vrrp 10 preempt
```

- The default is enabled.
- Use the **no** version to disable preemption.

ip vrrp priority

- Use to configure the priority of VRRP routers.
- A value of 255 implies “master router” priority.
- A value of 1–254 implies “backup router” priority.
- Example

```
host1(config-if)#ip vrrp 25 priority 255
```

- Use the **no** version to set the priority to the default value, 100.

ip vrrp virtual-address

- Use to associate an IP address with a VRID.
- If the **auto** keyword is used, associated addresses are automatically learned or configured, depending on the priority attribute.
- There is no default.
- Example

```
host1(config-if)#ip vrrp 25 virtual-address 194.2.1.63
  255.255.255.0
```

- Use **no ip vrrp virtual-address** to remove an IP address association with a VRID. If **auto** addresses are used, the **no** version clears the **auto** flag.

Monitoring VRRP

You can use several VRRP show commands to monitor the details of your VRRP configuration.

baseline ip vrrp

- Sets the baseline on all VRRP statistics as the current value.
- Example

```
host1#baseline ip vrrp
```

show ip vrrp

- Use to display a detailed summary of all VRIDs configured.
- Use the **interface** keyword to specify a specific Fast Ethernet or Gigabit Ethernet interface.
- Field descriptions
  - Interface – Fast Ethernet or Gigabit Ethernet interface specifier and VRID
  - primary address – IP address used while in master state; not necessarily an associated address
  - operational state – state of the VRRP router: master, backup, or init; when the operational state is backup, the current master's ip address is given
  - admin state – administrative status: enabled or disabled
  - up time – number of seconds that the VRID has been enabled in non-init state
  - interval – VRRP advertisement interval: seconds/milliseconds
  - last error status – help text used to debug any error detected
  - priority – priority of VRRP router
  - auth type – the VRRP authentication type: none or text
  - preemption – VRRP router preemption status: enabled or disabled
  - auto assoc address(es) – IP addresses associated with VRID

- Example

```
Interface: fastEthernet5/0.0 vrrpVrid: 1
  primary address: 10.0.0.2
  operational state: backup (current master: 130.0.0.1)
  admin state: enabled
  up time: 145 seconds
  interval: 1 second
  last error status: no error
  priority: 101
  auth type: none
  preemption: enabled
  auto assoc address(es): 10.0.0.1, 100.0.0.1, 101.0.0.1
```

show ip vrrp brief

- Use to display a brief summary of all VRIDs configured.
- Use the **interface** keyword to specify a specific Fast Ethernet or Gigabit Ethernet interface.
- Field descriptions
  - Interface – Fast Ethernet or Gigabit Ethernet specifier
  - VRID – VRRP router instance configured on this interface
  - Primary Address – IP address used while in master state; not necessarily an associated address
  - State – operational state of VRRP router: master, backup, or init
  - Adv – advertisement interval in seconds
  - Pri – priority assigned to this router
  - Admin – administrative state of the VRID: enabled or disabled
- Example

```
Interface               VRID  Primary Address  State   Adv  Pri  Admin
----------------------  ----  ---------------  ------  ---  ---  --------
fastEthernet12/8.1.1     255  123.123.123.123    init    1  100  disabled
gigabitEthernet12/8.1.1    1          1.1.1.1  master    1  254  enabled
```

show ip vrrp neighbor

- Use to display neighbors known to the VRRP routers.
- Use the **interface** keyword to specify a specific Fast Ethernet or Gigabit Ethernet interface.
- Field descriptions
  - Interface – Fast Ethernet or Gigabit Ethernet interface specifier and VRID of neighbors known to VRRP router
  - time discovered – date and time neighbor was detected
  - primary address – primary IP address of neighbor
  - adv interval (sec) – VRRP advertisement interval in seconds
  - priority – priority status of VRRP router (255 = owner)
  - auth type – VRRP authentication type: none or text

- assoc address(es) – IP addresses associated with VRID that are advertised by neighbor

- Example

```
Interface: fastEthernet5/0.0 vrrpVrid: 1
  time discovered: 08/09/2001 07:44
  primary address: 10.0.0.1
  adv interval (sec): 1
  priority: 255 (owner)
  auth type: none
  assoc address(es): 10.0.0.1, 100.0.0.1, 101.0.0.1

Interface: fastEthernet5/0.1 vrrpVrid: 11
  time discovered: 08/09/2001 07:44
  primary address: 11.0.0.1
  adv interval (sec): 1
  priority: 255 (owner)
  auth type: none
  assoc address(es): 11.0.0.1, 110.0.0.1, 111.0.0.1
```

show ip vrrp statistics global

- Use to display the statistics of configured VRRP routers and each individual VRID.
- Use the **delta** keyword with the **show ip vrrp statistics** commands to view the baseline statistics.
- Field descriptions
  - checksumErrors – total number of VRRP packets received with an invalid VRRP checksum value
  - versionErrors – total number of VRRP packets received with an unknown or unsupported version number
  - vrIdErrors – total number of VRRP packets received with an invalid VRID for this virtual router
  - iccErrors – count of line card notifications that did not make it to the controller
  - txErrors – count of advertisements that did not get sent due to resource limitations
  - rxErrors – count of advertisements received that could not be parsed by VRRP applications
  - Example

```
host1#show ip vrrp statistics global
Globals:
  checksumErrors: 0
  versionErrors: 0
  vrIdErrors: 0
  iccErrors: 0
  txErrors: 0
  rxErrors: 0
```

show ip vrrp statistics

- Use to display statistics of configured VRRP routers and each individual VRID.
- Use the **delta** keyword with the **show ip vrrp statistics** command to view the baseline statistics.
- Use the **interface** keyword with the **show ip vrrp statistics** command to specify a specific Fast Ethernet or Gigabit Ethernet interface.
- Field descriptions
  - Interface – Fast Ethernet or Gigabit Ethernet interface specifier and VRID
  - becomeMaster – total number of times that this VRID state has transitioned to MASTER
  - advertiseRcvd – total number of VRRP advertisements received
  - advertiseIntervalErrors – total number of VRRP advertisement packets received for which the advertisement interval is different from the one configured for the VRID
  - authFailures – total number of VRRP packets received that do not pass the authentication check
  - ipTtlErrors – total number of VRRP packets received with IP TTL (time-to-live) not equal to 255
  - priorityZeroPktsRcvd – total number of VRRP packets received with a priority of 0
  - priorityZeroPktsSent – total number of VRRP packets sent with a priority of 0
  - invalidTypePktsRcvd – total number of VRRP packets received with an invalid value in the Type field
  - addressListErrors – total number of VRRP packets received for which the address list does not match the locally configured list for the VRID
  - invalidAuthType – total number of VRRP packets received with an unknown authentication type
  - authTypeMismatch – total number of VRRP packets received with Auth Type not equal to the locally configured authentication method
  - packetLengthErrors – total number of VRRP packets received with a packet length less than the length of the VRRP header
- Example of statistics per interface

```
host1#show ip vrrp statistics interface fastEthernet 4/0
Globals:
  checksumErrors: 0
  versionErrors: 0
  vrIdErrors: 1
  iccErrors: 0
  txErrors: 0
  rxErrors: 0

Interface: fastEthernet4/0 vrrpVrid: 1
  becomeMaster: 10
  advertiseRcvd: 0
  advertiseIntervalErrors: 0
  authFailures: 0
```

```
    ipTtlErrors: 0
    priorityZeroPktsRcvd: 0
    priorityZeroPktsSent: 9
    invalidTypePktsRcvd: 0
    addressListErrors: 0
    invalidAuthType: 0
    authTypeMismatch: 0
    packetLengthErrors: 0

  Interface: fastEthernet4/0 vrrpVrid: 50
    becomeMaster: 0
    advertiseRcvd: 1000
    advertiseIntervalErrors: 0
    authFailures: 0
    ipTtlErrors: 0
    priorityZeroPktsRcvd: 0
    priorityZeroPktsSent: 0
    invalidTypePktsRcvd: 0
    addressListErrors: 0
    invalidAuthType: 0
    authTypeMismatch: 0
    packetLengthErrors: 0
```

- Example of statistics per interface and VRID

```
host1#show ip vrrp statistics interface fastEthernet 4/0 1
Interface: fastEthernet4/0 vrrpVrid: 1
  becomeMaster: 0
  advertiseRcvd: 0
  advertiseIntervalErrors: 0
  authFailures: 0
  ipTtlErrors: 0
  priorityZeroPktsRcvd: 0
  priorityZeroPktsSent: 0
  invalidTypePktsRcvd: 0
  addressListErrors: 0
  invalidAuthType: 0
  authTypeMismatch: 0
  packetLengthErrors: 0
```

show ip vrrp summary

- Use to display a summary count on all VRIDs configured.
- Field descriptions
  - ip interfaces with vrrp – total number of VRIDs configured on ip interfaces
  - entries – total number of entries in all states
  - entries enabled – number of entries that were enabled
  - entries with owner priority – number of entries with owner priority
  - entries in init state – number of entries in the init state

  - entries in backup state – number of entries in backup state
  - entries in master state – number of entries in master state

- Example

```
host1#show ip vrrp summary
  ip interfaces with vrrp: 1
  entries: 10
  entries enabled: 10
  entries with owner priority: 1
  entries in init state: 0
  entries in backup state: 9
  entries in master state: 1
```