
In This Chapter

This chapter provides information about configuring Virtual Router Redundancy Protocol (VRRP) parameters. Topics in this chapter include:

- [VRRP Overview on page 282](#)
 - [Virtual Router on page 283](#)
 - [IP Address Owner on page 283](#)
 - [Primary and Secondary IP Addresses on page 284](#)
 - [Virtual Router Master on page 284](#)
 - [Virtual Router Backup on page 285](#)
 - [Owner and Non-Owner VRRP on page 285](#)
 - [Configurable Parameters on page 286](#)
- [VRRP Priority Control Policies on page 294](#)
 - [VRRP Virtual Router Policy Constraints on page 294](#)
 - [VRRP Virtual Router Instance Base Priority on page 294](#)
 - [VRRP Priority Control Policy Delta In-Use Priority Limit on page 295](#)
 - [VRRP Priority Control Policy Priority Events on page 296](#)
- [VRRP Non-Owner Accessibility on page 302](#)
 - [Non-Owner Access Ping Reply on page 302](#)
 - [Non-Owner Access Telnet on page 302](#)
 - [Non-Owner Access SSH on page 303](#)
 - [VRRP Advertisement Message IP Address List Verification on page 292](#)
- [VRRP Configuration Process Overview on page 304](#)
- [Configuration Notes on page 305](#)

VRRP Overview

The Virtual Router Redundancy Protocol (VRRP) for IPv4 is defined in the IETF RFC 3768, *Virtual Router Redundancy Protocol*. VRRP for IPv6 is specified in *draft-ietf-vrrp-unified-spec-02.txt*. VRRP describes a method of implementing a redundant IP interface shared between two or more routers on a common LAN segment, allowing a group of routers to function as one virtual router. When this IP interface is specified as a default gateway on hosts directly attached to this LAN, the routers sharing the IP interface prevent a single point of failure by limiting access to this gateway address. VRRP can be implemented on IES service interfaces and on core network IP interfaces.

If the master virtual router fails, the backup router configured with the highest acceptable priority becomes the master virtual router. The new master router assumes the normal packet forwarding for the local hosts.

Figure 13 displays an example of a VRRP configuration.

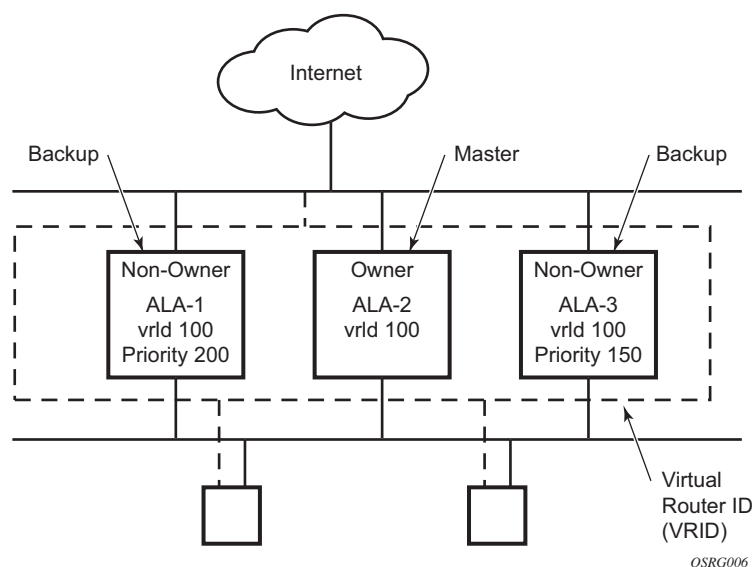


Figure 13: VRRP Configuration

VRRP Components

VRRP consists of the following components:

- [Virtual Router on page 283](#)
 - [IP Address Owner on page 283](#)
 - [Primary and Secondary IP Addresses on page 284](#)
 - [Virtual Router Master on page 284](#)
 - [Virtual Router Backup on page 285](#)
 - [Owner and Non-Owner VRRP on page 285](#)
-

Virtual Router

A virtual router is a logical entity managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier (VRID) and a set of associated IP addresses (or address) across a common LAN. A VRRP router can backup one or more virtual routers.

The purpose of supporting multiple IP addresses within a single virtual router is for multi-netting. This is a common mechanism that allows multiple local subnet attachment on a single routing interface. Up to four virtual routers are possible on a single Alcatel-Lucent IP interface. The virtual routers must be in the same subnet. Each virtual router has its own VRID, state machine and messaging instance.

IP Address Owner

VRRP can be configured in either an owner or non-owner mode. The owner is the VRRP router whose virtual router IP address is the same as the real interface IP address. This is the router that responds to packets addressed to one of the IP addresses for ICMP pings, TCP connections, etc. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

Alcatel-Lucent routers allow the virtual routers to be configured as non-owners of the IP address. VRRP on a router can be configured to allow non-owners to respond to ICMP echo requests when they become the virtual router master for the virtual router. Telnet and other connection-oriented protocols can also be configured for non-owner master response. However, the individual application conversations (connections) will not survive a VRRP failover. A non-owner VRRP

router operating as a backup will not respond to any packets addressed to any of the virtual router IP addresses.

Primary and Secondary IP Addresses

A primary address is an IP address selected from the set of real interface address. VRRP advertisements are always sent using the primary IP address as the source of the IP packet.

An IP interface must always have a primary IP address assigned for VRRP to be active on the interface. Alcatel-Lucent routers supports both primary and secondary IP addresses (multi-netting) on the IP interface. The virtual router's VRID primary IP address is always the primary address on the IP interface. VRRP uses the primary IP address as the IP address placed in the source IP address field of the IP header for all VRRP messages sent on that interface.

Virtual Router Master

The VRRP router which controls the IP address(es) associated with a virtual router is called the master. The master is responsible for forwarding packets sent to the VRRP IP addresses. An election process provides dynamic failover of the forwarding responsibility if the master becomes unavailable. This allows any of the virtual router IP addresses on the LAN to be used as the default first hop router by end hosts. This enables a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end host.

If the master is unavailable, each backup virtual router for the VRID compare the configured priority values to determine the master role. In case of a tie, the virtual router with the highest primary IP address becomes master.

The `preempt` parameter can be set to `false` to prevent a backup virtual router with a better priority value from becoming master when an existing non-owner virtual router is the current master. This is determined on a first-come, first-served basis.

While master, a virtual router routes and originates all IP packets into the LAN using the physical MAC address for the IP interface as the Layer 2 source MAC address, not the VRID MAC address. ARP packets also use the parent IP interface MAC address as the Layer 2 source MAC address while inserting the virtual router MAC address in the appropriate hardware address field. VRRP messages are the only packets transmitted using the virtual router MAC address as the Layer 2 source MAC.

Virtual Router Backup

A new virtual router master is selected from the set of VRRP routers available to assume forwarding responsibility for a virtual router should the current master fail.

Owner and Non-Owner VRRP

The owner controls the IP address of the virtual router and is responsible for forwarding packets sent to this IP address. The owner assumes the role of the master virtual router. Only one virtual router in the domain can be configured as owner. All other virtual router instances participating in this message domain must have the same VRID configured.

The most important parameter to be defined on a non-owner virtual router instance is the priority. The priority defines a virtual router's selection order in the master election process. The priority value and the preempt mode determine the virtual router with the highest priority to become the master virtual router.

The base priority is used to derive the in-use priority of the virtual router instance as modified by any optional VRRP priority control policy. VRRP priority control policies can be used to either override or adjust the base priority value depending on events or conditions within the chassis.

For information about non-owner access parameters, refer to [VRRP Non-Owner Accessibility on page 302](#).

Configurable Parameters

In addition to backup IP addresses, to facilitate configuration of a virtual router on Alcatel-Lucent routers, the following parameters can be defined in owner configurations:

- [Virtual Router ID \(VRID\) on page 286](#)
- [Message Interval and Master Inheritance on page 288](#)
- [VRRP Message Authentication on page 290](#)
- [Authentication Data on page 292](#)
- [Virtual MAC Address on page 292](#)

The following parameters can be defined in non-owner configurations:

- [Virtual Router ID \(VRID\) on page 286](#)
 - [Priority on page 286](#)
 - [Message Interval and Master Inheritance on page 288](#)
 - [Master Down Interval on page 289](#)
 - [Preempt Mode on page 289](#)
 - [VRRP Message Authentication on page 290](#)
 - [Authentication Data on page 292](#)
 - [Virtual MAC Address on page 292](#)
 - [Inherit Master VRRP Router's Advertisement Interval Timer on page 293](#)
 - [Policies on page 293](#)
-

Virtual Router ID (VRID)

The VRID must be configured with the same value on each virtual router associated with the redundant IP address (IP addresses). It is placed in all VRRP advertisement messages sent by each virtual router.

Priority

The priority value affects the interaction between this VRID and the same VRID of other virtual routers participating on the same LAN. A higher priority value defines a greater priority in becoming the virtual router master for the VRID. The priority value can only be configured when

the defined IP address on the IP interface is different than the virtual router IP address (non-owner mode).

When the IP address on the IP interface matches the virtual router IP address (owner mode), the priority value is fixed at 255, the highest value possible. This virtual router member is considered the owner of the virtual router IP address. There can only be one owner of the virtual router IP address for all virtual router members.

The priority value 0 is reserved for VRRP advertisement message purposes. It is used to tell other virtual routers in the same VRID that this virtual router is no longer acting as master, triggering a new election process. When this happens, each backup virtual router sets its master down timer equal to the skew time value. This shortens the time until one of the backup virtual routers becomes master.

The current master virtual router must transmit a VRRP advertisement message immediately upon receipt of a VRRP message with priority set to 0. This prevents another backup from becoming master for a short period of time.

Non-owner virtual routers may be configured with a priority of 254 through 1. The default value is 100. Multiple non-owners can share the same priority value. When multiple non-owner backup virtual routers are tied (transmit VRRP advertisement messages simultaneously) in the election process, both become master simultaneously, the one with the best priority will win the election. If the priority value in the message is equal to the master's local priority value, then the primary IP address of the local master and the message is evaluated as the tie breaker. The higher IP address becomes master. (The primary IP address is the source IP address of the VRRP advertisement message.)

The priority is also used to determine when to preempt the existing master. If the preempt mode value is true, VRRP advertisement messages from inferior (lower priority) masters are discarded, causing the master down timer to expire and the transition to master state.

The priority value also dictates the skew time added to the master timeout period.

IP Addresses

Each virtual router participating in the same VRID should be defined with the same set of IP addresses. These are the IP addresses being used by hosts on the LAN as gateway addresses. Multi-netting supports 16 IP addresses on the IP interface, up to 16 addresses can be assigned to a specific a virtual router instance.

Message Interval and Master Inheritance

Each virtual router is configured with a message interval per VRID within which it participates. This parameter must be the same for every virtual router on the VRID.

For IPv4, the default advertisement interval is 1 second and can be configured between 100 milliseconds and 255 seconds 900 milliseconds. For IPv6, the default advertisement interval is 1 second and can be configured between 100 milliseconds and 40 seconds 950 milliseconds.

As specified in the RFC, the advertisement interval field in every received VRRP advertisement message must match the locally configured advertisement interval. If a mismatch occurs, depending on the inherit configuration, the current master's advertisement interval setting can be used to operationally override the locally configured advertisement interval setting. If the current master changes, the new master setting is used. If the local virtual router becomes master, the locally configured advertisement interval is enforced.

If a VRRP advertisement message is received with an advertisement interval set to a value different than the local value and the inherit parameter is disabled, the message is discarded without processing.

The master virtual router on a VRID uses the advertisement interval to load the advertisement timer, specifying when to send the next VRRP advertisement message. Each backup virtual router on a VRID uses the advertisement interval (with the configured local priority) to derive the master down timer value.

VRRP advertisements messages that are fragmented, contain IP options (IPv4), or contain extension headers (IPv6) require a longer message interval to be configured.

Skew Time

The skew time is used to add a time period to the master down interval. This is not a configurable parameter. It is derived from the current local priority of the virtual router's VRID. To calculate the skew time, the virtual router evaluates the following formula:

For IPv4: $\text{Skew Time} = ((256 - \text{priority}) / 256) \text{ seconds}$

For IPv6: $\text{Skew Time} = (((256 - \text{priority}) * \text{Master_Adver_Interval}) / 256) \text{ centiseconds}$

The higher priority value, the smaller the skew time will be. This means that virtual routers with a lower priority will transition to master slower than virtual routers with higher priorities.

Master Down Interval

The master down interval is a calculated value used to load the master down timer. When the master down timer expires, the virtual router enters the master state. To calculate the master down interval, the virtual router evaluates the following formula:

$$\text{Master Down Interval} = (3 \times \text{Operational Advertisement Interval}) + \text{Skew Time}$$

The operational advertisement interval is dependent upon the state of the inherit parameter. When the inherit parameter is enabled, the operational advertisement interval is derived from the current master's advertisement interval field in the VRRP advertisement message. When inherit is disabled, the operational advertisement interval must be equal to the locally configured advertisement interval.

The master down timer is only operational when the local virtual router is operating in backup mode.

Preempt Mode

Preempt mode is a true or false configured value which controls whether a specific backup virtual router preempts a lower priority master. The IP address owner will always become master when available. Preempt mode cannot be set to false on the owner virtual router. The default value for preempt mode is true.

When preempt mode is true, the advertised priority from the incoming VRRP advertisement message from the current master is compared to the local configured priority. If the local priority is higher, the received VRRP advertisement message is discarded. This will result in the eventual expiration of the master down timer causing a transition to the master state. If the received priority is equal to the local priority, the message is not discarded and the current master will not be discarded. Note that when in the backup state, the received primary IP address is not part of the decision to preempt and is not used as a tie breaker when the received and local priorities are equal.

When `preempt` is enabled, the virtual router instance overrides any non-owner master with an in-use message priority value less than the virtual router instance in-use priority value. If `preempt` is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.

VRRP Message Authentication

The authentication type parameter defines the type of authentication used by the virtual router in VRRP advertisement message authentication. VRRP message authentication is applicable to IPv4 only. The current master uses the configured authentication type to indicate any egress message manipulation that must be performed in conjunction with any supporting authentication parameters before transmitting a VRRP advertisement message. The configured authentication type value is transmitted in the message authentication type field with the appropriate authentication data field filled in. Backup routers use the authentication type message field value in interpreting the contained authentication data field within received VRRP advertisement messages.

VRRP supports three message authentication methods which provide varying degrees of security. The supported authentication types are:

- 0 – No Authentication
- 1 – Simple Text Password
- 2 – IP Authentication Header

Authentication Type 0 – No Authentication

The use of type 0 indicates that VRRP advertisement messages are not authenticated (provides no authentication). The master transmitting VRRP advertisement messages will transmit the value 0 in the egress messages authentication type field and the authentication data field. Backup virtual routers receiving VRRP advertisement messages with the authentication type field equal to 0 will ignore the authentication data field in the message.

All compliant VRRP advertisement messages are accepted. The following fields within the received VRRP advertisement message are checked for compliance (the VRRP specification may require additional checks).

- IP header checks specific to VRRP
 - IP header destination IP address – Must be 224.0.0.18
 - IP header TTL field – Must be equal to 255, the packet must not have traversed any IP routed hops
 - IP header protocol field – must be 112 (decimal)

- VRRP message checks
 - Version field – Must be set to the value 2
 - Type field – Must be set to the value of 1 (advertisement)
 - Virtual router ID field – Must match one of the configured VRID on the ingress IP interface (All other fields are dependent on matching the virtual router ID field to one of the interfaces configured VRID parameters)
 - Priority field – Must be equal to or greater than the VRID in-use priority or be equal to 0 (Note, equal to the VRID in-use priority and 0 requires further processing regarding master/backup and senders IP address to determine validity of the message)
 - Authentication type field – Must be equal to 0
 - Advertisement interval field – Must be equal to the VRID configured advertisement interval
 - Checksum field – Must be valid
 - Authentication data fields – Must be ignored.

VRRP messages not meeting the criteria are silently dropped.

Authentication Type 1 – Simple Text Password

The use of type 1 indicates that VRRP advertisement messages are authenticated with a clear (simple) text password. All virtual routers participating in the virtual router instance must be configured with the same 8 octet password. Transmitting virtual routers place a value of 1 in the VRRP advertisement message authentication type field and put the configured simple text password into the message authentication data field. Receiving virtual routers compare the message authentication data field with the local configured simple text password based on the message authentication type field value of 1.

The same checks are performed for type 0 with the following exceptions (the VRRP specification may require additional checks):

- VRRP message checks
 - Authentication type field – Must be equal to 1
 - Authentication data fields – Must be equal to the VRID configured simple text password

Any VRRP message not meeting the type 0 verification checks with the exceptions above are silently discarded.

Authentication Failure

Any received VRRP advertisement message that fails authentication must be silently discarded with an invalid authentication counter incremented for the ingress virtual router instance.

Authentication Data

This feature is different than the VRRP advertisement message field with the same name. This is any required authentication information that is pertinent to the configured authentication type. The type of authentication data used for each authentication type is as follows:

<u>Authentication Type</u>	<u>Authentication Data</u>
0	None, authentication is not performed
1	Simple text password consisting of 8 octets

Virtual MAC Address

The MAC address can be used instead of an IP address in ARP responses when the virtual router instance is master. The MAC address configuration must be the same for all virtual routers participating as a virtual router or indeterminate connectivity by the attached IP hosts will result. All VRRP advertisement messages are transmitted with *ieee-mac-addr* as the source MAC.

VRRP Advertisement Message IP Address List Verification

VRRP advertisement messages contain an IP address count field that indicates the number of IP addresses listed in the sequential IP address fields at the end of the message.

The Alcatel-Lucent routers implementation always logs mismatching events. The decision on where and whether to forward the generated messages depends on the configuration of the event manager.

To facilitate the sending of mismatch log messages, each virtual router instance keeps the mismatch state associated with each source IP address in the VRRP master table. Whenever the state changes, a mismatch log message is generated indicating the source IP address within the message, the mismatch or match event and the time of the event.

With secondary IP address support, multiple IP addresses may be found in the list and it should match the IP address on the virtual router instance. Owner and non-owner virtual router instances have the supported IP addresses explicitly defined, making mismatched supported IP address within the interconnected virtual router instances a provisioning issue.

Inherit Master VRRP Router's Advertisement Interval Timer

The virtual router instance can inherit the master VRRP router's advertisement interval timer which is used by backup routers to calculate the master down timer.

The inheritance is only configurable in the non-owner nodal context. It is used to allow the current virtual router instance master to dictate the master down timer for all backup virtual routers.

IPv6 Virtual Router Instance Operationally Up

Once the IPv6 virtual router is properly configured with a minimum of one link-local backup address, the parent interface's router advertisement must be configured to use the virtual MAC address for the virtual router to be considered operationally up.

Policies

Policies can be configured to control VRRP priority with the virtual router instance. VRRP priority control policies can be used to override or adjust the base priority value depending on events or conditions within the chassis.

The policy can be associated with more than one virtual router instance. The priority events within the policy override or diminish the base priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority can eventually be restored to the base priority value.

Policies can only be configured in the non-owner VRRP context. For non-owner virtual router instances, if policies are not configured, then the base priority is used as the in-use priority.

VRRP Priority Control Policies

This implementation of VRRP supports control policies to manipulate virtual router participation in the VRRP master election process and master self-deprecation. The local priority value for the virtual router instance is used to control the election process and master state.

VRRP Virtual Router Policy Constraints

Priority control policies can only be applied to non-owner VRRP virtual router instances. Owner VRRP virtual routers cannot be controlled by a priority control policy because they are required to have a priority value of 255 that cannot be diminished. Only one VRRP priority control policy can be applied to a non-owner virtual router instance.

Multiple VRRP virtual router instances may be associated with the same IP interface, allowing multiple priority control policies to be associated with the IP interface.

An applied VRRP priority control policy only affects the in-use priority on the virtual router instance when the preempt mode has been enabled. A virtual router instance with preempt mode disabled will always use the base priority as the in-use priority, ignoring any configured priority control policy.

VRRP Virtual Router Instance Base Priority

Non-owner virtual router instances must have a base priority value between 1 and 254. The value 0 is reserved for master termination. The value 255 is reserved for owners. The default base priority for non-owner virtual router instances is the value 100.

The base priority is the starting priority for the VRRP instance. The actual in-use priority for the VRRP instance is derived from the base priority and an optional VRRP priority control policy.

VRRP Priority Control Policy Delta In-Use Priority Limit

A VRRP priority control policy enforces an overall minimum value that the policy can inflict on the VRRP virtual router instance base priority. This value provides a lower limit to the delta priority events manipulation of the base priority.

A delta priority event is a conditional event defined in the priority control policy that subtracts a given amount from the current, in-use priority for all VRRP virtual router instances to which the policy is applied. Multiple delta priority events can apply simultaneously, creating a dynamic priority value. The base priority for the instance, less the sum of the delta values derives the actual priority value in-use.

An explicit priority event is a conditional event defined in the priority control policy that explicitly defines the in-use priority for the virtual router instance. The explicitly defined values are not affected by the delta in-use priority limit. When multiple explicit priority events happen simultaneously, the lowest value is used for the in-use priority. The configured base priority is not a factor in explicit priority overrides of the in-use priority.

The allowed range of the Delta In-Use Priority Limit is 1 to 254. The default is 1, which prevents the delta priority events from operationally disabling the virtual router instance.

VRRP Priority Control Policy Priority Events

The main function of a VRRP priority control policy is to define conditions or events that impact the system's ability to communicate with outside hosts or portions of the network. When one or multiple of these events are true, the base priority on the virtual router instance is either overwritten with an explicit value, or a sum of delta priorities is subtracted from the base priority. The result is the in-use priority for the virtual router instance. Any priority event may be configured as an explicit event or a delta event.

Explicit events override all delta events. When multiple explicit events occur, the event with the lowest priority value is assigned to the in-use priority. As events clear, the in-use priority is reevaluated accordingly and adjusted dynamically.

Delta priority events also have priority values. When no explicit events have occurred within the policy, the sum of the occurring delta events priorities is subtracted from the base priority of each virtual router instance. If the result is lower than the delta in-use priority limit, the delta in-use priority limit is used as the in-use priority for the virtual router instance. Otherwise, the in-use priority is set to the base priority less the sum of the delta events.

Each event generates a VRRP priority event message indicating the policy-id, the event type, the priority type (delta or explicit) and the event priority value. Another log message is generated when the event is no longer true, indicating that it has been cleared.

Priority Event Hold-Set Timers

Hold-set timers are used to dampen the effect of a flapping event. A flapping event is where the event continually transitions between clear and set. The hold-set value is loaded into a hold set timer that prevents a set event from transitioning to the cleared state until it expires.

Each time an event transitions between cleared and set, the timer is loaded and begins to count down to zero. If the timer reaches zero, the event will be allowed to enter the cleared state once more. Entering the cleared state is always dependent on the object controlling the event conforming to the requirements defined in the event itself. It is possible, on some event types, to have a further set action reload the hold set timer. This extends the amount of time that must expire before entering the cleared state.

For an example of a hold-set timer setting, refer to [LAG Degrade Priority Event on page 297](#).

Port Down Priority Event

The port down priority event is tied to either a physical port or a SONET/SDH channel. The port or channel operational state is evaluated to determine a port down priority event or event clear.

When the port or channel operational state is up, the port down priority event is considered false or cleared. When the port or channel operational state is down, the port down priority event is considered true or set.

LAG Degrade Priority Event

The LAG degrade priority event is tied to an existing Link Aggregation Group (LAG). The LAG degrade priority event is conditional to percentage of available port bandwidth on the LAG. Multiple bandwidth percentage thresholds may be defined, each with its own priority value.

If the LAG transitions from one threshold to the next, the previous threshold priority value is subtracted from the total delta sum while the new threshold priority value is added to the sum. The new sum is then subtracted from the base priority and compared to the delta in-use priority limit to derive the new in-use priority on the virtual router instance.

The following example illustrates a LAG priority event and its interaction with the hold set timer in changing the in-use priority.

The following state and timer settings are used for the LAG events displayed in [Table 6](#):

- User-defined thresholds: 2 ports down 4 ports down 6 ports down
- LAG configured ports: 8 ports
- Hold set timer (hold-set): 5 seconds

Table 6: LAG Events

Time	LAG Port State	Parameter	State	Comments
0	All ports down	Event State	Set - 8 ports down	
		Event Threshold	6 ports down	
		Hold Set Timer	5 seconds	Set to hold-set parameter

Table 6: LAG Events (Continued)

Time	LAG Port State	Parameter	State	Comments
1	One port up	Event State	Set - 8 ports down	Cannot change until Hold Set Timer expires
		Event Threshold	6 ports down	
		Hold Set Timer	5 seconds	
2	All ports up	Event State	Set - 8 ports down	Still waiting for Hold Set Timer expires
		Event Threshold	6 ports down	
		Hold Set Timer	3 seconds	
5	All ports up	Event State	Cleared - All ports up	Event cleared
		Event Threshold	None	
		Hold Set Timer	Expired	
100	Five ports down	Event State	Set - 5 ports down	Set to hold-set parameter
		Event Threshold	4 ports down	
		Hold Set Timer	Expired	
102	Three ports down	Event State	Set - 5 ports down	
		Event Threshold	4 ports down	
		Hold Set Timer	3 seconds	
103	All ports up	Event State	Set - 5 ports down	
		Event Threshold	4 ports down	
		Hold Set Timer	2 second	
104	Two ports down	Event State	Set - 5 ports down	Current threshold is 5, so 2 down has no effect
		Event Threshold	4 ports down	
		Hold Set Timer	1 second	
105	Two ports down	Event State	Set - 2 ports down	
		Event Threshold	2 ports down	
		Hold Set Timer	Expired	
200	Four ports down	Event State	Set - 2 ports down	Set to hold-set parameter
		Event Threshold	4 ports down	
		Hold Set Timer	5 seconds	

Table 6: LAG Events (Continued)

Time	LAG Port State	Parameter	State	Comments
1	One port up	Event State	Set - 8 ports down	Cannot change until Hold Set Timer expires
		Event Threshold	6 ports down	
		Hold Set Timer	5 seconds	
2	All ports up	Event State	Set - 8 ports down	Still waiting for Hold Set Timer expires
		Event Threshold	6 ports down	
		Hold Set Timer	3 seconds	
5	All ports up	Event State	Cleared - All ports up	Event cleared
		Event Threshold	None	
		Hold Set Timer	Expired	
100	Five ports down	Event State	Set - 5 ports down	Set to hold-set parameter
		Event Threshold	4 ports down	
		Hold Set Timer	Expired	
102	Three ports down	Event State	Set - 5 ports down	
		Event Threshold	4 ports down	
		Hold Set Timer	3 seconds	
103	All ports up	Event State	Set - 5 ports down	
		Event Threshold	4 ports down	
		Hold Set Timer	2 second	
104	Two ports down	Event State	Set - 5 ports down	Current threshold is 5, so 2 down has no effect
		Event Threshold	4 ports down	
		Hold Set Timer	1 second	
105	Two ports down	Event State	Set - 2 ports down	
		Event Threshold	2 ports down	
		Hold Set Timer	Expired	
200	Four ports down	Event State	Set - 2 ports down	Set to hold-set parameter
		Event Threshold	4 ports down	
		Hold Set Timer	5 seconds	

Table 6: LAG Events (Continued)

Time	LAG Port State	Parameter	State	Comments
202	Seven ports down	Event State	Set - 7 ports down	Changed due to increase
		Event Threshold	6 ports down	
		Hold Set Timer	5 seconds	Set to hold-set due to threshold increase
206	All ports up	Event State	Set - 7 ports down	
		Event Threshold	6 ports down	
		Hold Set Timer	1 second	
207	All ports up	Event State	Cleared - All ports up	
		Event Threshold	None	Event cleared
		Hold Set Timer	Expired	

Host Unreachable Priority Event

The host unreachable priority event creates a continuous ping task that is used to test connectivity to a remote host. The path to the remote host and the remote host itself must be capable and configured to accept ICMP echo request and replies for the ping to be successful.

The ping task is controlled by interval and size parameters that define how often the ICMP request messages are transmitted and the size of each message. A historical missing reply parameter defines when the ping destination is considered unreachable.

When the host is unreachable, the host unreachable priority event is considered true or set. When the host is reachable, the host unreachable priority event is considered false or cleared.

Route Unknown Priority Event

The route unknown priority event defines a task that monitors the existence of a given route prefix in the system's routing table.

The route monitoring task can be constrained by a condition that allows a prefix that is less specific than the defined prefix to be considered as a match. The source protocol can be defined to indicate the protocol the installed route must be populated from. To further define match criteria when multiple instances of the route prefix exist, an optional next hop parameter can be defined.

When a route prefix exists within the active route table that matches the defined match criteria, the route unknown priority event is considered false or cleared. When a route prefix does not exist within the active route table matching the defined criteria, the route unknown priority event is considered true or set.

VRRP Non-Owner Accessibility

Although the RFC states that only VRRP owners can respond to ping and other management-oriented protocols directed to the VRID IP addresses, the routers allow an override of this restraint on a per VRRP virtual router instance basis.

Non-Owner Access Ping Reply

When non-owner access ping reply is enabled on a virtual router instance, ICMP echo request messages destined to the non-owner virtual router instance IP addresses are not discarded at the IP interface when operating in master mode. ICMP echo request messages are always discarded in backup mode.

When non-owner access ping reply is disabled on a virtual router instance, ICMP echo request messages destined to the non-owner virtual router instance IP addresses are silently discarded in both the master and backup modes.

Non-Owner Access Telnet

When non-owner access Telnet is enabled on a virtual router instance, authorized Telnet sessions may be established that are destined to the virtual router instance IP addresses when operating in master mode. Telnet sessions are always discarded at the IP interface when destined to a virtual router IP address operating in backup mode. Enabling non-owner access Telnet does not guarantee Telnet access, proper management and security features must be enabled to allow Telnet on this interface and possibly from the given source IP address.

When non-owner access Telnet is disabled on a virtual router instance, Telnet sessions destined to the non-owner virtual router instance IP addresses are silently discarded in both master and backup modes.

Non-Owner Access SSH

When non-owner access SSH is enabled on a virtual router instance, authorized SSH sessions may be established that are destined to the virtual router instance IP addresses when operating in master mode. SSH sessions are always discarded at the IP interface when destined to a virtual router IP address operating in backup mode. Enabling non-owner access SSH does not guarantee SSH access, proper management and security features must be enabled to allow SSH on this interface and possibly from the given source IP address. SSH is applicable to IPv4 VRRP only.

When non-owner access SSH is disabled on a virtual router instance, SSH sessions destined to the non-owner virtual router instance IP addresses are silently discarded in both master and backup modes.

VRRP Configuration Process Overview

Figure 14 displays the process to provision VRRP parameters.

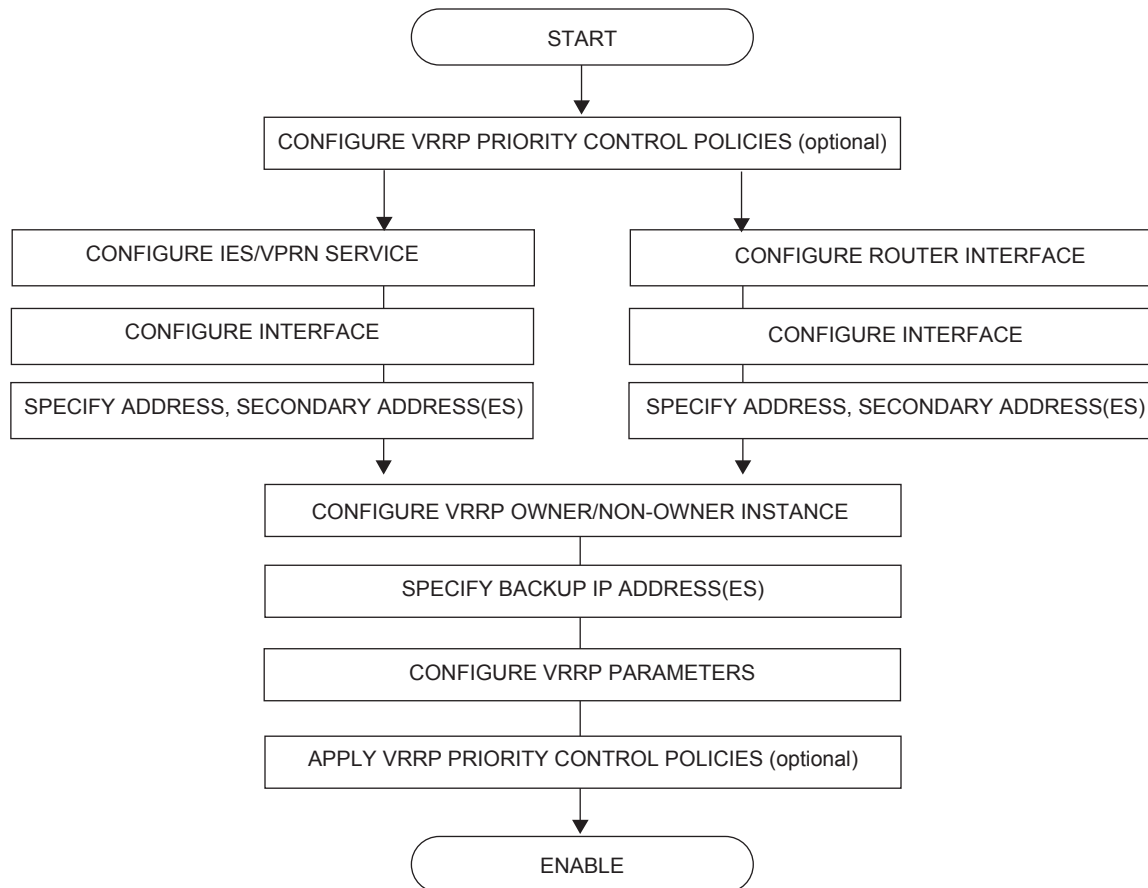


Figure 14: VRRP Configuration and Implementation Flow

Configuration Notes

This section describes VRRP configuration caveats.

General

- Creating and applying VRRP policies are optional.
- Backup command:
 - The backup IP address(es) must be on the same subnet. The backup addresses explicitly define which IP addresses are in the VRRP advertisement message IP address list.
 - In the owner mode, the backup IP address must be identical to one of the interface's IP addresses. The backup address explicitly defines which IP addresses are in the VRRP advertisement message IP address list.
 - For IPv6, one of the backup addresses configured must be the link-local address of the owner VRRP instance.

