
Configuration Commands

Interface Configuration Commands

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>if>vrrp
Description	<p>This command sets the simple text authentication key used to generate master VRRP advertisement messages and validates VRRP advertisements.</p> <p>If simple text password authentication is not required, the authentication-key command is not required.</p> <p>The command is configurable in both non-owner and owner vrrp nodal contexts.</p> <p>The <i>key</i> parameter identifies the simple text password to be used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses an eight octet long string that is inserted into all transmitted VRRP advertisement messages and is compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the <i>key</i>.</p> <p>The <i>key</i> string is case sensitive and is left justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field similarly holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with a 0 value in the corresponding octet.</p> <p>If the command is re-executed with a different password key defined, the new key is used ediatly.</p> <p>The authentication-key command can be executed at anytime.</p> <p>To change the current in-use password key on multiple virtual router instances:</p> <ol style="list-style-type: none"> 1. Identify the current master. 2. Shutdown the virtual router instance on all backups. 3. Execute the authentication-key command on the master to change the password key. 4. Execute the authentication-key command and no shutdown command on each backup. <p>The no form of the command reverts to the default value.</p>
Default	no authentication-key — The authentication key value is the null string.
Parameters	<i>authentication-key</i> — The authentication key. Allowed values are any string up to 8 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

hash-key — The hash key. The key can be any combination of ASCII characters up to 22 (*hash-key1*) or 121 (*hash-key2*) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup

Syntax	[no] backup <i>ip-address</i>
Context	config>router>if>vrrp
Description	<p>This command associates router IP addresses with the parental IP interface IP addresses.</p> <p>The backup command has two distinct functions when used in an owner or a non-owner context of the virtual router instance.</p> <p>Non-owner virtual router instances actually create a routable IP interface address that is operationally dependent on the virtual router instance mode (master or backup). The backup command in owner virtual router instances does not create a routable IP interface address; it simply defines the existing parental IP interface IP addresses that are advertised by the virtual router instance.</p> <p>For owner virtual router instances, the backup command defines the IP addresses that are advertised within VRRP advertisement messages. This communicates the IP addresses that the master is representing to backup virtual routers receiving the messages. Advertising a correct list is important. The specified <i>ip-addr</i> must be equal to one of the existing parental IP interface IP addresses (primary or secondary) or the backup command will fail.</p> <p>For non-owner virtual router instances, the backup command actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (ping-reply, telnet-reply, and ssh-reply). The specified <i>ip-addr</i> must be an IP address that is within one of the parental IP interface local subnets created with the address or secondary commands. If a local subnet does not exist that includes the specified <i>ip-addr</i> or if <i>ip-addr</i> is the same IP address as the parental IP interface IP address, the backup command will fail.</p> <p>The new interface IP address created with the backup command assumes the mask and parameters of the corresponding parent IP interface IP address. The <i>ip-addr</i> is only active when the virtual router instance is operating in the master state. When not operating as master, the virtual router instance acts as if it is operationally down. It will not respond to ARP requests to <i>ip-addr</i>, nor will it route packets received with its <i>vrid</i> derived source MAC address. A non-master virtual router instance always silently discards packets destined to <i>ip-addr</i>. A single virtual router instance may only have a single virtual router IP address from a given parental local subnet. Multiple virtual router instances can define a virtual router IP address from the same local subnet as long as each is a different IP address.</p>

In IPv4, up to sixteen **backup** *ip-addr* commands can be executed within the same virtual router instance. Executing **backup** multiple times with the same *ip-addr* results in no operation performed and no error generated. At least one successful **backup** *ip-addr* command must be executed before the virtual router instance can enter the operational state.

When operating as (non-owner) master, the default functionality associated with *ip-addr* is ARP response to ARP requests to *ip-addr*, routing of packets destined to the virtual router instance source MAC address and silently discarding packets destined to *ip-addr*. Enabling the non-owner-access parameters selectively allows ping, Telnet and SSH connectivity to *ip-addr* when the virtual router instance is operating as master.

The **no** form of the command removes the specified virtual router IP address from the virtual router instance. For non-owner virtual router instances, this causes all routing and local access associated with the *ip-addr* to cease. For **owner** virtual router instances, the **no backup** command only removes *ip-addr* from the list of advertised IP addresses. If the last *ip-addr* is removed from the virtual router instance, the virtual router instance will enter the operationally down state

Special Cases

Assigning the Virtual Router ID IP Address — Once the *vrid* is created on the parent IP interface, IP addresses need to be assigned to the virtual router instance. If the *vrid* was created with the keyword **owner**, the virtual router instance IP addresses must have one or more of the parent IP interface defined IP addresses (primary and secondary). For non-owner virtual router instances, the virtual router IP addresses each must be within one of the parental IP interface IP address defined local subnets. For both **owner** and non-owner virtual router instances, the virtual router IP addresses must be explicitly defined using the **backup** *ip-addr* command.

Virtual Router Instance IP Address Assignment Conditions — The RFC does not specify that the assigned IP addresses to the virtual router instance must be in the same subnet as the parent IP interface primary IP address or secondary IP addresses. The only requirement is that all virtual routers participating in the same virtual router instance have the same virtual router IP addresses assigned. To avoid confusion, the assigned virtual router IP addresses must be in a local subnet of one of the parent IP interfaces IP addresses. For **owner** virtual router instances the assigned virtual router IP address must be the same as one of the parental interface primary or secondary IP addresses.

The following rules apply when adding, changing, or removing parental and virtual router IP addresses:

Owner Virtual Router IP Address Parental Association — When an IP address is assigned to an **owner** virtual router instance, it must be associated with one of the parental IP interface-assigned IP addresses. The virtual router IP address must be equal to the primary or one of the secondary IP addresses within the parental IP interface.

Example - Owner Virtual Router Instance

Parent IP addresses:	10.10.10.10/24	
	11.11.11.11/24	
Virtual router IP addresses:	10.10.10.11	Invalid (not equal to parent IP address)
	10.10.10.10	Associated (same as parent IP address 10.10.10.10)
	10.10.11.11	Invalid (not equal to parent IP address)

11.11.11.254	Invalid (not equal to parent IP address)
11.11.11.255	Invalid (not equal to parent IP address)

Non-Owner Virtual Router IP Address Parental Association — When an IP address is assigned to a non-owner virtual router instance, it must be associated with one of the parental IP interface assigned IP addresses. The virtual router IP address must be a valid IP address within one of the parental IP interfaces local subnet. Local subnets are created by the primary or secondary IP addresses in conjunction with the IP addresses mask. If the defined virtual router IP address is equal to the associated subnet’s broadcast address, it is invalid. Virtual router IP addresses for non-owner virtual router instances that are equal to a parental IP interface IP address are also invalid.

The same virtual router IP address may not be assigned to two separate virtual router instances. If the virtual router IP address already exists on another virtual router instance, the virtual router IP address assignment will fail.

Example - Non-Owner Virtual Router Instance

Parent IP addresses:	10.10.10.10/24	
	11.11.11.11/24	
Virtual router IP addresses:	10.10.10.11	Associated with 10.10.10.10 (in subnet)
	10.10.10.10	Invalid (same as parent IP address)
	10.10.11.11	Invalid (outside of all Parent IP subnets)
	11.11.11.254	Associated with 11.11.11.11 (in subnet)
	11.11.11.255	Invalid (broadcast address of 11.11.11.11/24)

Virtual Router IP Address Assignment without Parent IP Address — When assigning an IP address to a virtual router instance, an associated IP address (see **Owner Virtual Router IP Address Parental Association** and **Non-Owner Virtual Router IP Address Parental Association**) on the parental IP interface must already exist. If an associated IP address on the parental IP interface is not configured, the virtual router IP address assignment fails.

Parent Primary IP Address Changed — When a virtual router IP address is set and the associated parent IP interface IP address is changed, the new parent IP interface IP address is evaluated to ensure it meets the association rules defined in **Owner Virtual Router IP Address Parental Association** or **Non-Owner Virtual Router IP Address Parental Association**. If the association check fails, the parental IP address change is not allowed. If the parental IP address change fails, the previously configured IP address definition remains in effect.

Only the primary parent IP address can be changed. Secondary addresses must be removed before the new IP address can be added. **Parent Primary or Secondary IP Address Removal** explains IP address removal conditions.

Parent Primary or Secondary IP Address Removal — When a virtual router IP address is successfully set, but removing the associated parent IP interface IP address is attempted and fails. All virtual router IP addresses associated with the parental IP interface IP address must be deleted prior

to removing the parental IP address. This includes virtual router IP address associations from multiple virtual router instances on the IP interface.

Default	no backup — No virtual router IP address is assigned.
Parameters	<i>ip-address</i> — The virtual router IP address expressed in dotted decimal notation. The IP virtual router IP address must be in the same subnet of the parental IP interface IP address or equal to one of the primary or secondary IP addresses for owner virtual router instances.
Values	1.0.0.1 - 223.255.255.254

backup

Syntax	config>router>if>ipv6>vrrp
Description	<p>This command associates router IPv6 addresses with the parental IP interface IP addresses.</p> <p>The backup command has two distinct functions when used in an owner or a non-owner context of the virtual router instance.</p> <p>Non-owner virtual router instances actually create a routable IP interface address that is operationally dependent on the virtual router instance mode (master or backup). The backup command in owner virtual router instances does not create a routable IP interface address; it simply defines the existing parental IP interface IP addresses that are advertised by the virtual router instance.</p> <p>For owner virtual router instances, the backup command defines the IP addresses that are advertised within VRRP advertisement messages. This communicates the IP addresses that the master is representing to backup virtual routers receiving the messages. Advertising a correct list is important. The specified <i>ipv6-addr</i> must be equal to one of the existing parental IP interface IP addresses (link-local or global) or the backup command will fail.</p> <p>For non-owner virtual router instances, the backup command actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (ping-reply, telnet-reply, and ssh-reply). The specified <i>ipv6-addr</i> must be an IP address that is within one of the parental IP interface local subnets created with the link-local-address or address commands. If a local subnet does not exist that includes the specified <i>ipv6-addr</i> or if <i>ipv6-addr</i> is the same IP address as the parental IP interface IP address, the backup command will fail.</p> <p>The new interface IP address created with the backup command assumes the mask and parameters of the corresponding parent IP interface IP address. The <i>ipv6-addr</i> is only active when the virtual router instance is operating in the master state. For IPv6 VRRP, the parental interface's IP address that is in the same subnet as the backup address must be manually-configured, non EUI-64 and configured to be in the preferred state.</p> <p>When not operating as master, the virtual router instance acts as if it is operationally down. It will not respond to ARP requests to <i>ipv6-addr</i>, nor will it route packets received with its <i>vrid</i> derived source MAC address. A non-master virtual router instance always silently discards packets destined to <i>ipv6-addr</i>. A single virtual router instance may only have a single virtual router IP address from a given parental local subnet. Multiple virtual router instances can define a virtual router IP address from the same local subnet as long as each is a different IP address.</p>

Executing **backup** multiple times with the same *ipv6-addr* results in no operation performed and no error generated. At least one successful **backup** *ipv6-addr* command must be executed before the virtual router instance can enter the operational state.

When operating as (non-owner) master, the default functionality associated with *ipv6-addr* is ARP response to ARP requests to *ip-addr*, routing of packets destined to the virtual router instance source MAC address and silently discarding packets destined to *ipv6-addr*. An IPv6 virtual router instance can enter the operational state only if one of the configured backup address is a link-local address and the router advertisement of the interface is configured to use the virtual MAC address. Enabling the non-owner-access parameters selectively allows ping, Telnet and traceroute connectivity to *ipv6-addr* when the virtual router instance is operating as master.

The **no** form of the command removes the specified virtual router IP address from the virtual router instance. For non-owner virtual router instances, this causes all routing and local access associated with the *ipv6-addr* to cease. For **owner** virtual router instances, the **no backup** command only removes *ipv6-addr* from the list of advertised IP addresses. If the last *ipv6-addr* or the link-local address is removed from the virtual router instance, the virtual router instance will enter the operationally down state

Special Cases

Assigning the Virtual Router ID Address — Once the *vrid* is created on the parent IP interface, IP addresses need to be assigned to the virtual router instance. If the *vrid* was created with the keyword **owner**, the virtual router instance IP addresses must have one or more of the parent IP interface defined IP addresses. For non-owner virtual router instances, the virtual router IP addresses each must be within one of the parental IP interface IP address defined local subnets. For both **owner** and non-owner virtual router instances, the virtual router IP addresses must be explicitly defined using the **backup** *ipv6-addr* command.

The following rules apply when adding, changing, or removing parental and virtual router IP addresses:

Owner Virtual Router IP Address Parental Association — When an IP address is assigned to an **owner** virtual router instance, it must be associated with one of the parental IP interface-assigned IP addresses.

Example - Owner Virtual Router Instance

Parent IP addresses:	10.10.10.10/24	
	11.11.11.11/24	
Virtual router IP addresses:	10.10.10.11	Invalid (not equal to parent IP address)
	10.10.10.10	Associated (same as parent IP address 10.10.10.10)
	10.10.11.11	Invalid (not equal to parent IP address)
	11.11.11.254	Invalid (not equal to parent IP address)
	11.11.11.255	Invalid (not equal to parent IP address)

Non-Owner Virtual Router IP Address Parental Association — When an IP address is assigned to a non-owner virtual router instance, it must be associated with one of the parental IP interface assigned IP addresses. The virtual router IP address must be a valid IP address within one of

the parental IP interfaces local subnet. Local subnets are created by the link-local or global IP addresses in conjunction with the IP addresses mask. If the defined virtual router IP address is equal to the associated subnet's broadcast address, it is invalid. Virtual router IP addresses for non-owner virtual router instances that are equal to a parental IP interface IP address are also invalid.

The same virtual router IP address may not be assigned to two separate virtual router instances. If the virtual router IP address already exists on another virtual router instance, the virtual router IP address assignment will fail.

One exception to this rule is for the IPv6 link-local address that is configured as a backup address. The same link-local address can be configured in all virtual routers that use the same vrid.

Example - Non-Owner Virtual Router Instance

Parent IP addresses:	10.10.10.10/24	
	11.11.11.11/24	
Virtual router IPv6 addresses:	10.10.10.11	Associated with 10.10.10.10 (in subnet)
	10.10.10.10	Invalid (same as parent IP address)
	10.10.11.11	Invalid (outside of all Parent IP subnets)
	11.11.11.254	Associated with 11.11.11.11 (in subnet)
	11.11.11.255	Invalid (broadcast address of 11.11.11.11/24)

Virtual Router IP Address Assignment without Parent IP Address — When assigning an IP address to a virtual router instance, an associated IP address (see **Owner Virtual Router IP Address Parental Association** and **Non-Owner Virtual Router IP Address Parental Association**) on the parental IP interface must already exist. If an associated IP address on the parental IP interface is not configured, the virtual router IP address assignment fails.

Virtual Router IPv6 Address Assignment — An IPv6 backup address requires that the parental IP address that is in the same subnet as the backup address must be manually configured, non-EUI-64 and configured to be in the preferred state.

Default	no backup — No virtual router IP address is assigned.	
Parameters	<i>ipv6-address</i> — The virtual router IP address expressed in dotted decimal notation. The IP virtual router IP address must be in the same subnet of the parental IP interface IP address or equal to one of the the parent interface addresses for owner virtual router instances.	
Values	ipv6-address	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x::d.d.d.d x: [0..FFFF]H d: [0..255]D

bfd-enable

Syntax	[no] bfd-enable [<i>service-id</i>] interface <i>interface-name</i> dst-ip <i>ip-address</i> [no] bfd-enable interface <i>interface-name</i> dst-ip <i>ip-address</i>						
Context	config>router>if>vrrp config>router>if>ipv6>vrrp						
Description	<p>This commands assigns a bi-directional forwarding detect (BFD) session to a given VRRP/SRRP instance. This BFD sessions provided a heartbeat mechanism that can be used to speed up the transition of the standby VRRP router to an active state. If the associated BFD session fails, the VRRP routers will immediately send a VRRP Advertisement message. In addition, the standby VRRP router(s) will transition to a Master state to speed convergence. The normal VRRP election process will then take place based on the Advertisement messages sent by all VRRP routers.</p> <p>There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.</p> <p>The parameters used for the BFD sessions are set by the BFD command under the IP interface.</p> <p>The no form of this command removes BFD from the configuration.</p>						
Default	none						
Parameters	<p><i>service-id</i> — Specifies the service ID of the interface running BFD.</p> <table> <tr> <td>Values</td> <td><i>service-id:</i></td> <td>1 — 2147483647</td> </tr> <tr> <td></td> <td><i>svc-name:</i></td> <td>64 characters maximum</td> </tr> </table> <p>interface <i>interface-name</i> — Specifies the name of the interface running BFD. The specified interface may not yet be configured with BFD. However, when it is, this virtual router will then initiate the BFD session.</p> <p>dst-ip <i>ip-address</i> — Specifies the destination address to be used for the BFD session.</p>	Values	<i>service-id:</i>	1 — 2147483647		<i>svc-name:</i>	64 characters maximum
Values	<i>service-id:</i>	1 — 2147483647					
	<i>svc-name:</i>	64 characters maximum					

init-delay

Syntax	init-delay <i>seconds</i> no init-delay		
Context	config>router>if>vrrp config>router>if>ipv6>vrrp		
Description	This command configures a VRRP initialization delay timer.		
Parameters	<p><i>seconds</i> — Specifies the initialization delay timer for VRRP, in seconds.</p> <table> <tr> <td>Values</td> <td>1 — 65535</td> </tr> </table>	Values	1 — 65535
Values	1 — 65535		

mac

Syntax	mac <i>mac-address</i> no mac
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>This command sets an explicit MAC address used by the virtual router instance overriding the VRRP default derived from the VRID.</p> <p>Changing the default MAC address is useful when an existing HSRP or other non-VRRP default MAC is in use by the IP hosts using the virtual router IP address. Many hosts do not monitor unessential ARPs and continue to use the cached non-VRRP MAC address after the virtual router becomes master of the host's gateway address.</p> <p>The mac command sets the MAC address used in ARP responses when the virtual router instance is master. Routing of IP packets with <i>mac-address</i> as the destination MAC is also enabled. The mac setting 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 <i>mac-address</i> as the source MAC.</p> <p>The command can be configured in both non-owner and owner vrrp nodal contexts.</p> <p>The mac command can be executed at any time and takes effect ediatly. When the virtual router MAC on a master virtual router instance changes, a gratuitous ARP is ediatly sent with a VRRP advertisement message. If the virtual router instance is disabled or operating as backup, the gratuitous ARP and VRRP advertisement message is not sent.</p> <p>The no form of the command restores the default VRRP MAC address to the virtual router instance.</p>
Default	no mac — The virtual router instance uses the default VRRP MAC address derived from the VRID.
Parameters	<i>mac-address</i> — The 48-bit MAC address for the virtual router instance in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

master-int-inherit

Syntax	[no] master-int-inherit
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>This command enables the virtual router instance to inherit the master VRRP router's advertisement interval timer which is used by backup routers to calculate the master down timer.</p> <p>The master-int-inherit command is only available in the non-owner nodal context and is used to allow the current virtual router instance master to dictate the master down timer for all backup virtual routers. The master-int-inherit command has no effect when the virtual router instance is operating as master.</p>

If **master-int-inherit** is not enabled, the locally configured **message-interval** must match the master's VRRP advertisement message advertisement interval field value or the message is discarded.

The **no** form of the command restores the default operating condition which requires the locally configured **message-interval** to match the received VRRP advertisement message advertisement interval field value.

Default no master-int-inherit — The virtual router instance does not inherit the master VRRP router's advertisement interval timer and uses the locally configured message interval.

message-interval

Syntax	message-interval {[seconds] [milliseconds milliseconds]} no message-interval
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>This command configures the administrative advertisement message timer used by the master virtual router instance to send VRRP advertisement messages and to derive the master down timer as backup.</p> <p>For an owner virtual router instance, the administrative advertisement timer directly sets the operational advertisement timer and indirectly sets the master down timer for the virtual router instance.</p> <p>Non-owner virtual router instances usage of the message-interval setting is dependent on the state of the virtual router (master or backup) and the state of the master-int-inherit parameter.</p> <ul style="list-style-type: none">• When a non-owner is operating as master for the virtual router, the configured message-interval is used as the operational advertisement timer similar to an owner virtual router instance. The master-int-inherit command has no effect when operating as master.• When a non-owner is in the backup state with master-int-inherit disabled, the configured message-interval value is used to match the incoming VRRP advertisement message advertisement interval field. If the locally configured message interval does not match the advertisement interval field, the VRRP advertisement is discarded.• When a non-owner is in the backup state with master-int-inherit enabled, the configured message-interval is ignored. The master down timer is indirectly derived from the incoming VRRP advertisement message advertisement interval field value. <p>VRRP advertisements messages that are fragmented, contain IP options (IPv4), or contain extension headers (IPv6) require a longer message interval to be configured.</p> <p>The in-use value of the message interval is used to derive the master down timer to be used when the virtual router is operating in backup mode based on the following formula:</p> $(3x \text{ (in-use message interval) } + \text{ skew time})$ <p>The skew time portion is used to slow down virtual routers with relatively low priority values when competing in the master election process.</p> <p>The command is available in both non-owner and owner vrrp nodal contexts.</p>

By default, a **message-interval** of 1 second is used.

The **no** form of the command reverts to the default value.

Default	1 — Advertisement timer set to 1 second
Parameters	<i>seconds</i> — The number of seconds that will transpire before the advertisement timer expires expressed as a decimal integer.
	Values IPv4: 1 — 255 IPv6: 1 — 40
	milliseconds <i>milliseconds</i> — Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on the 7450 ESS-1 chassis.
	Values 100 — 900 IPv6: 10 — 990

policy

Syntax	policy <i>policy-id</i> no policy
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>This command adds a VRRP priority control policy association with the virtual router instance.</p> <p>To further augment the virtual router instance base priority, VRRP priority control policies can be used to override or adjust the base priority value depending on events or conditions within the chassis.</p> <p>The policy can be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base priority set with the priority command 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.</p> <p>The policy command is only available in the non-owner vrrp nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the policy command is not executed, the base priority is used as the in-use priority.</p> <p>The no form of the command removes existing VRRP priority control policy associations from the virtual router instance. All associations must be removed prior to deleting the policy from the system.</p>
Default	no policy — No VRRP priority control policy is associated with the virtual router instance.
Parameters	<i>policy-id</i> — The policy ID of the VRRP priority control expressed as a decimal integer. The <i>vrrp-policy-id</i> must already exist for the command to function.
	Values 1 — 9999

preempt

Syntax	[no] preempt
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.</p> <p>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.</p> <p>The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.</p> <p>The default value for preempt mode is enabled.</p>
Default	preempt

priority

Syntax	priority <i>base-priority</i> no priority
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>This command configures the base router priority for the virtual router instance used in the master election process.</p> <p>The priority is the most important parameter set on a non-owner virtual router instance. The priority defines a virtual router's selection order in the master election process. Together, the priority value and the preempt mode allow the virtual router with the best priority to become the master virtual router.</p> <p>The <i>base-priority</i> 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.</p> <p>The priority command is only available in the non-owner vrrp nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed.</p> <p>For non-owner virtual router instances, the default base priority value is 100.</p> <p>The no form of the command reverts to the default value.</p>
Default	100

Parameters	<i>base-priority</i> — The base priority used by the virtual router instance expressed as a decimal integer. If no VRRP priority control policy is defined, the <i>base-priority</i> is the in-use priority for the virtual router instance.
	Values 1 — 254

ping-reply

Syntax	[no] ping-reply
Context	config>router>if>vrrp config>router>if>ipv6>vrrp
Description	<p>This command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses.</p> <p>Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses. Many network administrators find this limitation frustrating when troubleshooting VRRP connectivity issues.</p> <p>SR OS allows this access limitation to be selectively lifted for certain applications. Ping, Telnet and SSH can be individually enabled or disabled on a per-virtual-router-instance basis.</p> <p>The ping-reply command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses. The Ping request can be received on any routed interface. Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address).</p> <p>When ping-reply is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.</p> <p>Non-owner backup virtual routers never respond to ICMP echo requests regardless of the ping-reply setting.</p> <p>The ping-reply command is only available in non-owner vrrp nodal context.</p> <p>By default, ICMP echo requests to the virtual router instance IP addresses are silently discarded.</p> <p>The no form of the command configures discarding all ICMP echo request messages destined to the non-owner virtual router instance IP addresses.</p>
Default	no ping-reply — ICMP echo requests to the virtual router instance IP addresses are discarded.

shutdown

Syntax [no] shutdown

Context config>router>if>vrrp
config>router>if>ipv6>vrrp
config>vrrp>policy

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Special Cases **Non-Owner Virtual Router** — Non-owner virtual router instances can be administratively shutdown. This allows the termination of VRRP participation in the virtual router and stops all routing and other access capabilities with regards to the virtual router IP addresses. Shutting down the virtual router instance provides a mechanism to maintain the virtual routers without causing false backup/master state changes.

If the **shutdown** command is executed, no VRRP advertisement messages are generated and all received VRRP advertisement messages are silently discarded with no processing.

By default, virtual router instances are created in the **no shutdown** state.

Whenever the administrative state of a virtual router instance transitions, a log message is generated.

Whenever the operational state of a virtual router instance transitions, a log message is generated.

Owner Virtual Router — An owner virtual router context does not have a **shutdown** command. To administratively disable an owner virtual router instance, use the **shutdown** command within the parent IP interface node which administratively downs the IP interface.

ssh-reply

Syntax [no] ssh-reply

Context config>router>if>vrrp

Description This command enables the non-owner master to reply to SSH requests directed at the virtual router instance IP addresses. This command is only applicable to IPv4.

Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses.

This limitation can be disregarded for certain applications. Ping, Telnet and SSH can be individually enabled or disabled on a per-virtual-router-instance basis.

The **ssh-reply** command enables the non-owner master to reply to SSH requests directed at the virtual router instances IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or

based on the SSH source host address). Proper login and CLI command authentication is still enforced.

When **ssh-reply** is not enabled, SSH requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH requests regardless of the **ssh-reply** setting.

The **ssh-reply** command is only available in non-owner **vrrp** nodal context.

By default, SSH requests to the virtual router instance IP addresses are silently discarded.

The **no** form of the command discards all SSH request messages destined to the non-owner virtual router instance IP addresses.

Default **no ssh-reply** — SSH requests to the virtual router instance IP addresses are discarded.

standby-forwarding

Syntax **[no] standby-forwarding**

Context config>router>if>vrrp
config>router>if>ipv6>vrrp

Description This command specifies whether this VRRP instance allows forwarding packets to a standby router. When disabled, a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address. When enabled, a standby router should forward all traffic.

telnet-reply

Syntax **[no] telnet-reply**

Context config>router>if>vrrp
config>router>if>ipv6>vrrp

Description This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the virtual router instances' IP addresses.

Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses. Many network administrators find this limitation frustrating when troubleshooting VRRP connectivity issues.

This limitation can be disregarded for certain applications. Ping, SSH and Telnet can each be individually enabled or disabled on a per-virtual-router-instance basis.

The **telnet-reply** command enables the non-owner master to reply to Telnet requests directed at the virtual router instances' IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication is still enforced.

Interface Configuration Commands

When **telnet-reply** is not enabled, Telnet requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet requests regardless of the **telnet-reply** setting.

The **telnet-reply** command is only available in non-owner **vrrp** nodal context.

By default, Telnet requests to the virtual router instance IP addresses will be silently discarded.

The **no** form of the command configures discarding all Telnet request messages destined to the non-owner virtual router instance IP addresses.

Default **no telnet-reply** — Telnet requests to the virtual router instance IP addresses are discarded.

traceroute-reply

Syntax **[no] traceroute-reply**

Context config>router>if>vrrp
config>router>if>ipv6>vrrp

Description This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **traceroute-reply** status.

Default no traceroute-reply

vrrp

Syntax **vrrp vrid [owner]**
no vrrp vrid

Context config>router>interface *ip-int-name*
config>router>if>ipv6

Description This command creates the context to configure a VRRP virtual router instance. A virtual router is defined by its virtual router identifier (VRID) and a set of IP addresses.

The optional **owner** keyword indicates that 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.

All other virtual router instances participating in this message domain must have the same *vrid* configured and cannot be configured as **owner**. Once created, the **owner** keyword is optional when entering the *vrid* for configuration purposes.

A *vrid* is internally associated with the IP interface. This allows the *vrid* to be used on multiple IP interfaces while representing different virtual router instances.

For IPv4, up to four **vrrp vrid** nodes can be configured on a router interface. Each virtual router instance can manage up to 16 backup IP addresses. For IPv6, only one virtual router ID can be configured on a router interface.

The **no** form of the command removes the specified *vrid* from the IP interface. This terminates VRRP participation and deletes all references to the *vrid* in conjunction with the IP interface. The *vrid* does not need to be shutdown to remove the virtual router instance.

Special Cases

Virtual Router Instance Owner IP Address Conditions — It is possible for the virtual router instance **owner** to be created prior to assigning the parent IP interface primary or secondary IP addresses. When this is the case, the virtual router instance is not associated with an IP address. The operational state of the virtual router instance is down.

VRRP Owner Command Exclusions — By specifying the VRRP *vrid* as **owner**, The following commands are no longer available:

- **vrrp priority** — The virtual router instance **owner** is hard-coded with a **priority** value of 255 and cannot be changed.
- **vrrp master-int-inherit** — Owner virtual router instances do not accept VRRP advertisement messages; the advertisement interval field is not evaluated and cannot be inherited.
- **ping-reply, telnet-reply** and **ssh-reply** — The **owner** virtual router instance always allows Ping, Telnet and SSH if the management and security parameters are configured to accept them on the parent IP interface.
- **vrrp shutdown** — The **owner** virtual router instance cannot be shutdown in the **vrrp** node. If this was allowed, VRRP messages would not be sent, but the parent IP interface address would continue to respond to ARPs and forward IP packets. Another virtual router instance may detect the missing master due to the termination of VRRP advertisement messages and become master. This would cause two routers responding to ARP requests for the same IP addresses. To **shutdown** the **owner** virtual router instance, use the **shutdown** command in the parent IP interface context. This will prevent VRRP participation, IP ARP reply and IP forwarding. To continue parent IP interface ARP reply and forwarding without VRRP participation, remove the **vrrp vrid** instance.
- **traceroute-reply**

Default

no vrrp — No VRRP virtual router instance is associated with the IP interface.

Parameters

vrid — The virtual router ID for the IP interface expressed as a decimal integer.

Values 1 — 255

owner — Identifies this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. Once created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted and then recreated without the **owner** keyword to remove ownership.

Priority Policy Commands

delta-in-use-limit

Syntax	delta-in-use-limit <i>in-use-priority-limit</i> no delta-in-use-limit
Context	config>vrrp>policy <i>vrrp-policy-id</i>
Description	<p>This command sets a lower limit on the virtual router in-use priority that can be derived from the delta priority control events.</p> <p>Each <i>vrrp-priority-id</i> places limits on the delta priority control events to define the in-use priority of the virtual router instance. Setting this limit prevents the sum of the delta priority events from lowering the in-use priority value of the associated virtual router instances below the configured value.</p> <p>The limit has no effect on explicit priority control events. Explicit priority control events are controlled by setting the in-use priority to any value between 1 and 254.</p> <p>Only non-owner virtual router instances can be associated with VRRP priority control policies and their priority control events.</p> <p>Once the total sum of all delta events is calculated and subtracted from the base priority of the virtual router instance, the result is compared to the delta-in-use-limit value. If the result is less than the limit, the delta-in-use-limit value is used as the virtual router in-use priority value. If an explicit priority control event overrides the delta priority control events, the delta-in-use-limit has no effect.</p> <p>Setting the limit to a higher value than the default of 1 limits the effect of the delta priority control events on the virtual router instance base priority value. This allows for multiple priority control events while minimizing the overall effect on the in-use priority.</p> <p>Changing the <i>in-use-priority-limit</i> causes an ediate re-evaluation of the in-use priority values for all virtual router instances associated with this <i>vrrp-policy-id</i> based on the current sum of all active delta control policy events.</p> <p>The no form of the command reverts to the default value.</p>
Default	1 — The lower limit of 1 for the in-use priority, as modified, by delta priority control events.
Parameters	<p><i>in-use-priority-limit</i> — The lower limit of the in-use priority base, as modified by priority control policies. The <i>in-use-priority-limit</i> has the same range as the non-owner virtual router instance base-priority parameter. If the result of the total delta priority control events minus the virtual router instances base-priority, is less than the <i>in-use-priority-limit</i>, the <i>in-use-priority-limit</i> value is used as the virtual router instances in-use priority value.</p> <p>Setting the <i>in-use-priority-limit</i> to a value equal to or larger than the virtual router instance <i>base-priority</i> prevents the delta priority control events from having any effect on the virtual router instance in-use priority value.</p>
Values	1 — 254

description

Syntax	description <i>string</i> no description
Context	config>vrrp>policy <i>vrrp-policy-id</i>
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of the command removes the string from the configuration.
Default	none
Parameters	<i>string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

policy

Syntax	policy <i>policy-id</i> [context <i>service-id</i>] no policy <i>policy-id</i>
Context	config>vrrp
Description	This command creates the context to configure a VRRP priority control policy which is used to control the VRRP in-use priority based on priority control events. It is a parental node for the various VRRP priority control policy commands that define the policy parameters and priority event conditions. The virtual router instance priority command defines the initial or base value to be used by non-owner virtual routers. This value can be modified by assigning a VRRP priority control policy to the virtual router instance. The VRRP priority control policy can override or diminish the base priority setting to establish the actual in-use priority of the virtual router instance. The policy <i>policy-id</i> command must be created first, before it can be associated with a virtual router instance. Because VRRP priority control policies define conditions and events that must be maintained, they can be resource intensive. The number of policies is limited to 1000. The <i>policy-id</i> do not have to be consecutive integers. The range of available policy identifiers is from 1 to 9999. The no form of the command deletes the specific <i>policy-id</i> from the system. The <i>policy-id</i> must be removed first from all virtual router instances before the no policy command can be issued. If the <i>policy-id</i> is associated with a virtual router instance, the command will fail.
Default	none

Priority Policy Commands

- Parameters** *vrrp-policy-id* — The VRRP priority control ID expressed as a decimal integer that uniquely identifies this policy from any other VRRP priority control policy defined on the system. Up to 1000 policies can be defined.
- Values** 1 — 9999
- context** *service-id* — Specifies the service ID to which this policy applies. A value of zero (0) means that this policy does not apply to a service but applies to the base router instance.
- Values** 1 — 2147483647

priority-event

- Syntax** **[no]** **priority-event**
- Context** config>vrrp>policy *vrrp-priority-id*
- Description** This command creates the context to configure VRRP priority control events used to define criteria to modify the VRRP in-use priority.
- A priority control event specifies an object to monitor and the effect on the in-use priority level for an associated virtual router instance.
- Up to 32 priority control events can be configured within the **priority-event** node.
- The **no** form of the command clears any configured priority events.

Priority Policy Event Commands

hold-clear

Syntax	hold-clear <i>seconds</i> no hold-clear
Context	config>vrrp>policy>priority-event>port-down config>vrrp>policy>priority-event>lag-port-down config>vrrp>policy>priority-event>route-unknown
Description	<p>This command configures the hold clear time for the event. The <i>seconds</i> parameter specifies the hold-clear time, the amount of time in seconds by which the effect of a cleared event on the associated virtual router instance is delayed.</p> <p>The hold-clear time is used to prevent black hole conditions when a virtual router instance advertises itself as a master before other conditions associated with the cleared event have had a chance to enter a forwarding state.</p>
Default	no hold-clear
Parameters	<i>seconds</i> — Specifies the amount of time in seconds by which the effect of a cleared event on the associated virtual router instance is delayed.
Values	0 — 86400

hold-set

Syntax	hold-set <i>seconds</i> no hold-set
Context	config>vrrp>policy>priority-event>host-unreachable config>vrrp>policy>priority-event>lag-port-down config>vrrp>policy>priority-event>port-down config>vrrp>policy>priority-event>route-unknown
Description	<p>This command specifies the amount of time that must pass before the set state for a VRRP priority control event can transition to the cleared state to dampen flapping events. A flapping event continually transitions between clear and set.</p> <p>The hold-set command is used to dampen the effect of a flapping event. 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.</p> <p>Each time an event transitions between cleared and set, the timer is loaded and begins a countdown to zero. When the timer reaches zero, the event is allowed to enter the cleared state. Entering the cleared state is 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 another set action reload the hold-set timer. This extends the amount of time that must expire before entering the cleared state.</p>

Once the hold set timer expires and the event meets the cleared state requirements or is set to a lower threshold, the current set effect on the virtual router instances in-use priority can be removed. As with **lag-port-down** events, this may be a decrease in the set effect if the *clearing* amounts to a lower set threshold.

The **hold-set** command can be executed at anytime. If the hold-set timer value is configured larger than the new *seconds* setting, the timer is loaded with the new **hold-set** value.

The **no** form of the command reverts the default value.

Default 0 — The hold-set timer is disabled so event transitions are processed ediatly.

Parameters *seconds* — The number of seconds that the hold set timer waits after an event enters a set state or enters a higher threshold set state, depending on the event type.

The value of 0 disables the hold set timer, preventing any delay in processing lower set thresholds or cleared events.

Values 0 — 86400

priority

Syntax **priority** *priority-level* [{**delta** | **explicit**}]
no priority

Context config>vrrp>policy>priority-event>host-unreachable *ip-addr*
config>vrrp>policy>priority-event>lag-port-down *lag-id*>number-down *number-of-lag-ports-down*
config>vrrp>policy>priority-event>port-down *port-id*[*channel-id*]
config>vrrp>policy>priority-event>route-unknown *prefix/mask-length*

Description This command controls the effect the set event has on the virtual router instance in-use priority. When the event is set, the *priority-level* is either subtracted from the base priority of each virtual router instance or it defines the explicit in-use priority value of the virtual router instance depending on whether the **delta** or **explicit** keywords are specified.

Multiple set events in the same policy have interaction constraints:

- If any set events have an explicit **priority** value, all the delta **priority** values are ignored.
- The set event with the lowest explicit **priority** value defines the in-use priority that are used by all virtual router instances associated with the policy.
- If no set events have an explicit **priority** value, all the set events delta **priority** values are added and subtracted from the base priority value defined on each virtual router instance associated with the policy.
- If the delta priorities sum exceeds the **delta-in-use-limit** parameter, then the **delta-in-use-limit** parameter is used as the value subtracted from the base priority value defined on each virtual router instance associated with the policy.

If the **priority** command is not configured on the priority event, the *priority-value* defaults to 0 and the qualifier keyword defaults to **delta**, thus, there is no impact on the in-use priority.

The **no** form of the command reverts to the default values.

Default 0 delta — The set event will subtract 0 from the base priority (no effect).

Parameters *priority-level* — The priority level adjustment value expressed as a decimal integer.

Values 0 — 254

delta | explicit — Configures what effect the *priority-level* will have on the base priority value.

When **delta** is specified, the *priority-level* value is subtracted from the associated virtual router instance's base priority when the event is set and no explicit events are set. The sum of the priority event *priority-level* values on all set delta priority events are subtracted from the virtual router base priority to derive the virtual router instance in-use priority value. If the **delta** priority event is cleared, the *priority-level* is no longer used in the in-use priority calculation.

When **explicit** is specified, the *priority-level* value is used to override the base priority of the virtual router instance if the priority event is set and no other **explicit** priority event is set with a lower *priority-level*. The set **explicit** priority value with the lowest *priority-level* determines the actual in-use protocol value for all virtual router instances associated with the policy.

Default delta

Values delta, explicit

weight-down

Syntax **weight-down** *lag-ports-down-weight*
no weight-down

Context config>vrrp>policy>priority-event>lag-port-down

Description This command creates a context to configure an event set threshold within a lag-port-down priority control event. The weight-down command defines a sub-node within the lag-port-down event and is uniquely identified with the lag-ports-down-weight parameter. Each weight-down node within the same lag-port-down event node must have a unique lag-ports-down-weight value. Each weight-down node has its own priority command that takes effect whenever that node represents the current threshold. A single LAG can use either weight-threshold or port threshold. The command is required for proper operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.

The total number of sub-nodes (uniquely identified by the lag-ports-down-weight parameter) allowed in the system is 2048.

A **weight-down** node is not required for each possible number of ports that could be down. The active threshold is always the closest lower threshold.

The **no** form of the command deletes the event set threshold. The threshold may be removed at any time. If the removed threshold is the current active threshold, the event set thresholds must be re-evaluated after removal.

Default no weight-down

Parameters *lag-ports-down-weight* — The total weight of LAG ports down to create a set event threshold. This is the active threshold when the weight of down ports in the LAG equals or exceeds *lag-ports-down-weight*, but does not equal or exceed the next highest configured *lag-ports-down-weight*.

mc-ipsec-non-forwarding

Syntax	<code>[no] mc-ipsec-non-forwarding <i>tunnel-grp-id</i></code>
Context	<code>config>vrrp>policy>priority-event</code>
Description	This command configures an instance of a multi-chassis IPsec tunnel-group Priority Event used to override the base priority value of a VRRP virtual router instance depending on the operational state of the event.
Parameters	<i>tunnel-grp-id</i> — Identifies the multi-chassis IPsec tunnel group whose non-forwarding state is monitored by this priority control event.

Priority Policy Port Down Event Commands

port-down

Syntax	<code>[no] port-down port-id</code>
Context	<code>config>vrrp>policy>priority-event</code>
Description	<p>This command configures a port down priority control event that monitors the operational state of a port or SONET/SDH channel. When the port or channel enters the operational down state, the event is considered set. When the port or channel enters the operational up state, the event is considered cleared.</p> <p>Multiple unique port-down event nodes can be configured within the priority-event context up to the overall limit of 32 events. Up to 32 events can be defined in any combination of types.</p> <p>The port-down command can reference an arbitrary port or channel . The port or channel does not need to be pre-provisioned or populated within the system. The operational state of the port-down event is set as follows:</p> <ul style="list-style-type: none">• Set – non-provisioned• Set – not populated• Set – down• Cleared – up <p>When the port or channel is provisioned, populated, or enters the operationally up or down state, the event operational state is updated appropriately.</p> <p>When the event enters the operationally down, non-provisioned, or non-populated state, the event is considered to be set. When an event transitions from clear to set, the set is processed ediatly and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from cleared to set, a hold set timer is loaded with the value configured by the events hold-set command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold set timer expires, the timer is reset to the hold-set value, extending the time before another clear can take effect.</p> <p>When the event enters the operationally up state, the event is considered to be cleared. Once the events hold-set expires, the effects of the events priority value are ediatly removed from the in-use priority of all associated virtual router instances.</p> <p>The actual effect on the virtual router instance in-use priority value depends on the defined event priority and its delta or explicit nature.</p> <p>The no form of the command deletes the specific port or channel monitoring event. The event may be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances will be re-evaluated. The events hold-set timer has no effect on the removal procedure.</p>
Default	no port-down — No port down priority control events are defined.
Parameters	<i>port-id</i> — The port ID of the port monitored by the VRRP priority control event.

Priority Policy Port Down Event Commands

The *port-id* can only be monitored by a single event in this policy. The port can be monitored by multiple VRRP priority control policies. A port and a specific channel on the port are considered to be separate entities. A port and a channel on the port can be monitored by separate events in the same policy.

Values	port-id	<i>slot/mda/port[.channel]</i>
	aps-id	<i>aps-group-id[.channel]</i>
		aps keyword
		group-id 1 — 64
	bundle-type-slot/mda.<bundle-num>	
		bundle keyword
		type ima, ppp
		bundle-num 1 —256
	ccag-id	<i>ccag-id.path-id[cc-type]</i>
		ccag keyword
		id 1 — 8
		path-id a, b
		cc-type .sap-net, .net-sap

The POS channel on the port monitored by the VRRP priority control event. The *port-id.channel-id* can only be monitored by a single event in this policy. The channel can be monitored by multiple VRRP priority control policies. A port and a specific channel on the port are considered to be separate entities. A port and a channel on the port can be monitored by separate events in the same policy.

If the port is provisioned, but the *channel* does not exist or the port has not been populated, the appropriate event operational state is Set – non-populated.

If the port is not provisioned, the event operational state is Set – non-provisioned.

If the POS interface is configured as a clear-channel, the *channel-id* is 1 and the channel bandwidth is the full bandwidth of the port.

Priority Policy LAG Events Commands

lag-port-down

Syntax	<code>[no] lag-port-down lag-id</code>
Context	config>vrrp>policy>priority-event
Description	<p>This command creates the context to configure Link Aggregation Group (LAG) priority control events that monitor the operational state of the links in the LAG.</p> <p>The lag-port-down command configures a priority control event. The event monitors the operational state of each port in the specified LAG. When one or more of the ports enter the operational down state, the event is considered to be set. When all the ports enter the operational up state, the event is considered to be clear. As ports enter the operational up state, any previous set threshold that represents more down ports is considered cleared, while the event is considered to be set.</p> <p>Multiple unique lag-port-down event nodes can be configured within the priority-event node up to the maximum of 32 events.</p> <p>The lag-port-down command can reference an arbitrary LAG. The <i>lag-id</i> does have to already exist within the system. The operational state of the lag-port-down event will indicate:</p> <ul style="list-style-type: none">• Set – non-existent• Set – one port down• Set – two ports down• Set – three ports down• Set – four ports down• Set – five ports down• Set – six ports down• Set – seven ports down• Set – eight ports down• Cleared – all ports up <p>When the <i>lag-id</i> is created, or a port in <i>lag-id</i> becomes operationally up or down, the event operational state must be updated appropriately.</p> <p>When one or more of the LAG composite ports enters the operationally down state or the <i>lag-id</i> is deleted or does not exist, the event is considered to be set. When an event transitions from clear to set, the set is processed ediatly and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold set timer is loaded with the value configured by the events hold-set command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold set timer expires, the timer is reset to the hold-set value, extending the time before another clear can take effect.</p>

The **lag-port-down** event is considered to have a tiered event set state. While the priority impact per number of ports down is totally configurable, as more ports go down, the effect on the associated virtual router instances in-use priority is expected to increase (lowering the priority). When each configured threshold is crossed, any higher thresholds are considered further event sets and are processed ediatly with the hold set timer reset to the configured value of the **hold-set** command. As the thresholds are crossed in the opposite direction (fewer ports down then previously), the priority effect of the event is not processed until the hold set timer expires. If the number of ports down threshold again increases before the hold set timer expires, the timer is only reset to the **hold-set** value if the number of ports down is equal to or greater than the threshold that set the timer.

The event contains **number-down** nodes that define the priority delta or explicit value to be used based on the number of LAG composite ports that are in the operationally down state. These nodes represent the event set thresholds. Not all port down thresholds must be configured. As the number of down ports increase, the **number-down** *ports-down* node that expresses a value equal to or less than the number of down ports describes the delta or explicit priority value to be applied.

The **no** form of the command deletes the specific LAG monitoring event. The event can be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances must be reevaluated. The events **hold-set** timer has no effect on the removal procedure.

Default no lag-port-down — No LAG priority control events are created.

Parameters *lag-id* — The LAG ID that the specific event is to monitor expressed as a decimal integer. The *lag-id* can only be monitored by a single event in this policy. The LAG may be monitored by multiple VRRP priority control policies. A port within the LAG and the LAG ID itself are considered to be separate entities. A composite port may be monitored with the **port-down** event while the *lag-id* the port is in is monitored by a **lag-port-down** event in the same policy.

Values 1 — 800

number-down

Syntax [**no**] **number-down** *number-of-lag-ports-down*

Context config>vrrp>policy>priority-event>lag-port-down *lag-id*

Description This command creates a context to configure an event set threshold within a lag-port-down priority control event.

The **number-down** command defines a sub-node within the **lag-port-down** event and is uniquely identified with the *number-of-lag-ports-down* parameter. Each **number-down** node within the same **lag-port-down** event node must have a unique *number-of-lag-ports-down* value. Each **number-down** node has its own **priority** command that takes effect whenever that node represents the current threshold.

The total number of sub-nodes (uniquely identified by the *number-of-lag-ports-down* parameter) allowed in a single **lag-port-down** event is equal to the total number of possible physical ports allowed in a LAG.

A **number-down** node is not required for each possible number of ports that could be down. The active threshold is always the closest lower threshold. When the number of ports down equals a given threshold, that is the active threshold.

The **no** form of the command deletes the event set threshold. The threshold may be removed at any time. If the removed threshold is the current active threshold, the event set thresholds must be re-evaluated after removal.

Default no number-down — No threshold for the LAG priority event is created.

Parameters *number-of-lag-ports-down* — The number of LAG ports down to create a set event threshold. This is the active threshold when the number of down ports in the LAG equals or exceeds *number-of-lag-ports-down*, but does not equal or exceed the next highest configured *number-of-lag-ports-down*.

Values 1 — 64 (for 64-link LAG)
1 — 32 (for other LAGs)

Priority Policy Host Unreachable Event Commands

drop-count

Syntax	drop-count <i>consecutive-failures</i> no drop-count
Context	config>vrrp <i>vrrp-policy-id</i> >priority-event>host-unreachable <i>ip-addr</i>
Description	<p>This command configures the number of consecutively sent ICMP echo request messages that must fail before the host unreachable priority control event is set.</p> <p>The drop-count command is used to define the number of consecutive message send attempts that must fail for the host-unreachable priority event to enter the set state. Each unsuccessful attempt increments the event's consecutive message drop counter. With each successful attempt, the event's consecutive message drop counter resets to zero.</p> <p>If the event's consecutive message drop counter reaches the drop-count value, the host-unreachable priority event enters the set state.</p> <p>The event's hold-set value defines how long the event must stay in the set state even when a successful message attempt clears the consecutive drop counter. The event is not cleared until the consecutive drop counter is less than the drop-count value and the hold-set timer has a value of zero (expired).</p> <p>The no form of the command reverts to the default value.</p>
Default	3 — 3 consecutive ICMP echo request failures are required before the host unreachable priority control event is set.
Parameters	<p><i>consecutive-failures</i> — The number of ICMP echo request message attempts that must fail for the event to enter the set state. It also defines the threshold so a lower consecutive number of failures can clear the event state.</p> <p>Values 1 — 60</p>

host-unreachable

Syntax	[no] host-unreachable <i>ip-address</i> [no] host-unreachable <i>ipv6-address</i>
Context	config>vrrp>policy>priority-event
Description	<p>This command creates the context to configure a host unreachable priority control event to monitor the ability to receive ICMP echo reply packets from an IP host address.</p> <p>A host unreachable priority event creates a continuous ICMP echo request (ping) probe to the specified <i>ip-address</i>. If a ping fails, the event is considered to be set. If a ping is successful, the event is considered to be cleared.</p>

Multiple unique (different *ip-address*) **host-unreachable** event nodes can be configured within the **priority-event** node to a maximum of 32 events.

The **host-unreachable** command can reference any valid local or remote IP address. The ability to ARP a local IP address or find a remote IP address within a route prefix in the route table is considered part of the monitoring procedure. The **host-unreachable** priority event operational state tracks ARP or route table entries dynamically appearing and disappearing from the system. The operational state of the **host-unreachable** event can be one of the following:

Host Unreachable Operational State	Description
Set – no ARP	No ARP address found for <i>ip-addr</i> for drop-count consecutive attempts. Only applies when IP address is considered local.
Set – no route	No route exists for <i>ip-addr</i> for drop-count consecutive attempts. Only when IP address is considered remote.
Set – host unreachable	ICMP host unreachable message received for drop-count consecutive attempts.
Set – no reply	ICMP echo request timed out for drop-count consecutive attempts.
Set – reply received	Last ICMP echo request attempt received an echo reply but historically not able to clear the event.
Cleared – no ARP	No ARP address found for <i>ip-addr</i> - not enough failed attempts to set the event.
Cleared – no route	No route exists for <i>ip-addr</i> - not enough failed attempts to set the event.
Cleared – host unreachable	ICMP host unreachable message received - not enough failed attempts to set the event.
Cleared – no reply	ICMP echo request timed out - not enough failed attempts to set the event.
Cleared – reply received	Event is cleared - last ICMP echo request received an echo reply.

Unlike other priority event types, the **host-unreachable** priority event monitors a repetitive task. A historical evaluation is performed on the success rate of receiving ICMP echo reply messages. The operational state takes its cleared and set orientation from the historical success rate. The informational portion of the operational state is derived from the last attempt's result. It is possible for the previous attempt to fail while the operational state is still cleared due to an insufficient number of failures to cause it to become set. It is also possible for the state to be set while the previous attempt was successful.

When an event transitions from clear to set, the set is processed ediatly and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold set timer is loaded with the value configured by the events **hold-set** command. This timer

Priority Policy Host Unreachable Event Commands

prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The hold-set timer be expired and the historical success rate must be met prior to the event operational state becoming cleared.

The **no** form of the command deletes the specific IP host monitoring event. The event may be deleted at anytime. When the event is deleted, the in-use priority of all associated virtual router instances must be reevaluated. The event's **hold-set** timer has no effect on the removal procedure.

Default **no host-unreachable** — No host unreachable priority events are created.

Parameters *ip-addr* — The IP address of the host for which the specific event will monitor connectivity. The *ip-addr* can only be monitored by a single event in this policy. The IP address can be monitored by multiple VRRP priority control policies. The IP address can be used in one or multiple **ping** requests. Each VRRP priority control **host-unreachable** and **ping** destined to the same *ip-addr* is uniquely identified on a per message basis. Each session originates a unique identifier value for the ICMP echo request messages it generates. This allows received ICMP echo reply messages to be directed to the appropriate sending application.

Values

ipv4-address :	a.b.c.d
ipv6-address :	x:x:x:x:x:x:x[-interface]
x:	[0..FFFF]H
interface:	32 chars maximum, mandatory for link local addresses

Note that the link-local IPv6 address must have an interface name specified. The global IPv6 address must not have an interface name specified.

interval

Syntax **interval seconds**
no interval

Context config>vrrp>priority-event>host-unreachable

Description This command configures the number of seconds between host unreachable priority event ICMP echo request messages directed to the host IP address.

The **no** form of the command reverts to the default value.

Default 1

Parameters *seconds* — The number of seconds between the ICMP echo request messages sent to the host IP address for the host unreachable priority event.

Values 1 — 60

padding-size

Syntax	padding-size <i>size</i> no padding-size
Context	config>vrrp>priority-event>host-unreachable
Description	This command allows the operator to increase the size of IP packet by padding the PDU. The no form of the command reverts to the default.
Default	0
Parameters	<i>size</i> — Specifies amount of increase to to ICMP PDU.
Values	0 — 16384

timeout

Syntax	timeout <i>seconds</i> no timeout
Context	config>vrrp <i>vrrp-policy-id</i> >priority-event>host-unreachable <i>ip-addr</i>
Description	<p>This command defines the time, in seconds, that must pass before considering the far-end IP host unresponsive to an outstanding ICMP echo request message.</p> <p>The timeout value is not directly related to the configured interval parameter. The timeout value may be larger, equal, or smaller, relative to the interval value.</p> <p>If the timeout value is larger than the interval value, multiple ICMP echo request messages may be outstanding. Every ICMP echo request message transmitted to the far end host is tracked individually according to the message identifier and sequence number.</p> <p>With each consecutive attempt to send an ICMP echo request message, the timeout timer is loaded with the timeout value. The timer decrements until:</p> <ul style="list-style-type: none"> • An internal error occurs preventing message sending (request unsuccessful). • An internal error occurs preventing message reply receiving (request unsuccessful). • A required route table entry does not exist to reach the IP address (request unsuccessful). • A required ARP entry does not exist and ARP request timed out (request unsuccessful). • A valid reply is received (request successful). <p>Note that it is possible for a required ARP request to succeed or timeout after the message timeout timer expires. In this case, the message request is unsuccessful.</p> <p>If an ICMP echo reply message is not received prior to the timeout period for a given ICMP echo request, that request is considered to be dropped and increments the consecutive message drop counter for the priority event.</p> <p>If an ICMP echo reply message with the same sequence number as an outstanding ICMP echo request message is received prior to that message timing out, the request is considered successful. The consecutive message drop counter is cleared and the request message no longer is outstanding.</p>

Priority Policy Host Unreachable Event Commands

If an ICMP Echo Reply message with a sequence number equal to an ICMP echo request sequence number that had previously timed out is received, that reply is silently discarded while incrementing the priority event reply discard counter.

The **no** form of the command reverts to the default value.

Default 1

Parameters *seconds* — The number of seconds before an ICMP echo request message is timed out. Once a message is timed out, a reply with the same identifier and sequence number is discarded.

Values 1 — 60

Priority Policy Route Unknown Event Commands

less-specific

Syntax	[no] less-specific [allow-default]
Context	config>vrrp>policy>priority-event>route-unknown <i>prefix/mask-length</i>
Description	<p>This command allows a CIDR shortest match hit on a route prefix that contains the IP route prefix associated with the route unknown priority event.</p> <p>The less-specific command modifies the search parameters for the IP route prefix specified in the route-unknown priority event. Specifying less-specific allows a CIDR shortest match hit on a route prefix that contains the IP route prefix.</p> <p>The less-specific command eases the RTM lookup criteria when searching for the <i>prefix/mask-length</i>. When the route-unknown priority event sends the prefix to the RTM (as if it was a destination lookup), the result route table prefix (if a result is found) is checked to see if it is an exact match or a less specific match. The less-specific command enables a less specific route table prefix to match the configured prefix. When less-specific is not specified, a less specific route table prefix fails to match the configured prefix. The allow-default optional parameter extends the less-specific match to include the default route (0.0.0.0).</p> <p>The no form of the command prevents RTM lookup results that are less specific than the route prefix from matching.</p>
Default	no less-specific — The route unknown priority events requires an exact prefix/mask match.
Parameters	allow-default — When the allow-default parameter is specified with the less-specific command, an RTM return of 0.0.0.0 matches the IP prefix. If less-specific is entered without the allow-default parameter, a return of 0.0.0.0 will not match the IP prefix. To disable allow-default , but continue to allow less-specific match operation, only enter the less-specific command (without the allow-default parameter).

next-hop

Syntax	[no] next-hop <i>ip-address</i>
Context	config>vrrp>policy>priority-event>route-unknown <i>prefix/mask-length</i>
Description	<p>This command adds an allowed next hop IP address to match the IP route prefix for a route-unknown priority control event.</p> <p>If the next-hop IP address does not match one of the defined <i>ip-address</i>, the match is considered unsuccessful and the route-unknown event transitions to the set state.</p> <p>The next-hop command is optional. If no next-hop <i>ip-address</i> commands are configured, the comparison between the RTM prefix return and the route-unknown IP route prefix are not included in the next hop information.</p>

Priority Policy Route Unknown Event Commands

When more than one next hop IP addresses are eligible for matching, a **next-hop** command must be executed for each IP address. Defining the same IP address multiple times has no effect after the first instance.

The **no** form of the command removes the *ip-address* from the list of acceptable next hops when looking up the **route-unknown** prefix. If this *ip-address* is the last next hop defined on the **route-unknown** event, the returned next hop information is ignored when testing the match criteria. If the *ip-address* does not exist, the **no next-hop** command returns a warning error, but continues to execute if part of an **exec** script.

Default no next-hop — No next hop IP address for the route unknown priority control event is defined.

Parameters *ip-address* — The IP address for an acceptable next hop IP address for a returned route prefix from the RTM when looking up the **route-unknown** route prefix.

Values

ipv4-address :	a.b.c.d
ipv6-address :	x::x::x::x::x[-interface]
x:	[0..FFFF]H
interface:	32 chars maximum, mandatory for link local addresses

Note that the link-local IPv6 address must have an interface name specified. The global IPv6 address must not have an interface name specified.

protocol

Syntax **protocol** {**bgp** | **bgp-vpn** | **ospf** | **is-is** | **rip** | **static**}
no protocol

Context config>vrrp>policy>priority-event>route-unknown *prefix/mask-length*

Description This command adds one or more route sources to match the route unknown IP route prefix for a route unknown priority control event.

If the route source does not match one of the defined protocols, the match is considered unsuccessful and the **route-unknown** event transitions to the set state.

The **protocol** command is optional. If the **protocol** command is not executed, the comparison between the RTM prefix return and the **route-unknown** IP route prefix will not include the source of the prefix. The **protocol** command cannot be executed without at least one associated route source parameter. All parameters are reset each time the **protocol** command is executed and only the explicitly defined protocols are allowed to match.

The **no** form of the command removes protocol route source as a match criteria for returned RTM route prefixes.

To remove specific existing route source match criteria, execute the **protocol** command and include only the specific route source criteria. Any unspecified route source criteria is removed.

Default no protocol — No route source for the route unknown priority event is defined.

Parameters **bgp** — This parameter defines BGP as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **bgp** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **bgp** parameter,

a returned route prefix with a source of BGP will not be considered a match and will cause the event to enter the set state.

bgp-vpn — This parameter defines **bgp-vpn** as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **bgp-vpn** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **bgp-vpn** parameter, a returned route prefix with a source of **bgp-vpn** will not be considered a match and will cause the event to enter the set state.

ospf — This parameter defines OSPF as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **ospf** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **ospf** parameter, a returned route prefix with a source of OSPF will not be considered a match and will cause the event to enter the set state.

is-is — This parameter defines IS-IS as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **is-is** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **is-is** parameter, a returned route prefix with a source of IS-IS will not be considered a match and will cause the event to enter the set state.

rip — This parameter defines RIP as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **rip** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **rip** parameter, a returned route prefix with a source of RIP will not be considered a match and will cause the event to enter the set state.

static — This parameter defines a static route as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **static** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **static** parameter, a returned route prefix with a source of static route will not be considered a match and will cause the event to enter the set state.

route-unknown

Syntax	[no] route-unknown <i>prefix/mask-length</i>
Context	config>vrrp>policy>priority-event
Description	<p>This command creates a context to configure a route unknown priority control event that monitors the existence of a specific active IP route prefix within the routing table.</p> <p>The route-unknown command configures a priority control event that defines a link between the VRRP priority control policy and the Route Table Manager (RTM). The RTM registers the specified route prefix as monitored by the policy. If any change (add, delete, new next hop) occurs relative to the prefix, the policy is notified and takes proper action according to the priority event definition. If the route prefix exists and is active in the routing table according to the conditions defined, the event is in the cleared state. If the route prefix is removed, becomes inactive or fails to meet the event criteria, the event is in the set state.</p> <p>The command creates a route-unknown node identified by <i>prefix/mask-length</i> and containing event control commands.</p>

Multiple unique (different *prefix/mask-length*) **route-unknown** event nodes can be configured within the **priority-event** node up to the maximum limit of 32 events.

The **route-unknown** command can reference any valid IP address mask-length pair. The IP address and associated mask length define a unique IP router prefix. The dynamic monitoring of the route prefix results in one of the following event operational states:

route-unknown Operational State	Description
Set – non-existent	The route does not exist in the route table.
Set – inactive	The route exists in the route table but is not being used.
Set – wrong next hop	The route exists in the route table but does not meet the next-hop requirements.
Set – wrong protocol	The route exists in the route table but does not meet the protocol requirements.
Set – less specific found	The route exists in the route table but does not meet any less-specific requirements.
Set – default best match	The route exists in the route table as the default route but the default route is not allowed for route matching.
Cleared – less specific found	A less specific route exists in the route table and meets all criteria including the less-specific requirements.
Cleared – found	The route exists in the route table manager and meets all criteria.

An existing route prefix in the RTM must be active (used by the IP forwarding engine) to clear the event operational state. It may be less specific (the defined prefix may be contained in a larger prefix according to Classless Inter-Domain Routing (CIDR) techniques) if the event has the **less-specific** statement defined. The less specific route that incorporates the router prefix may be the default route (0.0.0.0) if the **less-specific allow-default** statement is defined. The matching prefix may be required to have a specific next hop IP address if defined by the event **next-hop** command. Finally, the source of the RTM prefix may be required to be one of the dynamic routing protocols or be statically defined if defined by the event **protocol** command. If an RTM prefix is not found that matches all the above criteria (if defined in the event control commands), the event is considered to be set. If a matching prefix is found in the RTM, the event is considered to be cleared.

When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The **no** form of the command is used to remove the specific *prefix/mask-length* monitoring event. The event can be removed at anytime. When the event is removed, the in-use priority of all associated

virtual router instances must be reevaluated. The events **hold-set** timer has no effect on the removal procedure.

Default **no route-unknown** — No route unknown priority control events are defined for the priority control event policy.

Parameters *prefix* — The IP prefix address to be monitored by the route unknown priority control event in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

mask-length — The subnet mask length expressed as a decimal integer associated with the IP *prefix* defining the route prefix to be monitored by the route unknown priority control event.

Values 0 — 32

ip-address — The IP address of the host for which the specific event will monitor connectivity. The *ip-addr* can only be monitored by a single event in this policy. The IP address can be monitored by multiple VRRP priority control policies. The IP address can be used in one or multiple **ping** requests. Each VRRP priority control **host-unreachable** and **ping** destined to the same *ip-addr* is uniquely identified on a per message basis. Each session originates a unique identifier value for the ICMP echo request messages it generates. This allows received ICMP echo reply messages to be directed to the appropriate sending application.

Values *ip-prefix/mask:* ip-prefix a.b.c.d (host bits must be 0)
mask 0 — 32

ipv6-address/prefix:
ipv6-address x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:d.d.d.d
x: [0..FFFF]H
prefix-length 1 — 128

