

---

## VPRN Service Configuration Commands

---

### Generic Commands

#### shutdown

<b>Syntax</b>	<code>[no] shutdown</code>
<b>Context</b>	<pre> config&gt;service&gt;vprn config&gt;service&gt;vprn&gt;dhcp6&gt;server&gt;failover config&gt;service&gt;vprn&gt;igmp-trk config&gt;service&gt;vprn&gt;red-if config&gt;service&gt;vprn&gt;router-advert&gt;if config&gt;service&gt;vprn&gt;gsmp config&gt;service&gt;vprn&gt;gsmp&gt;group config&gt;service&gt;vprn&gt;gsmp&gt;group&gt;neighbor config&gt;service&gt;vprn&gt;igmp config&gt;service&gt;vprn&gt;igmp&gt;if config&gt;service&gt;vprn&gt;igmp&gt;if&gt;mcac config&gt;service&gt;vprn&gt;igmp&gt;if&gt;mcac&gt;mc-constraints config&gt;service&gt;vprn&gt;if config&gt;service&gt;vprn&gt;if&gt;dhcp config&gt;service&gt;vprn&gt;if&gt;dhcp&gt;proxy config&gt;service&gt;vprn&gt;if&gt;vrrp config&gt;service&gt;vprn&gt;if&gt;ipv6&gt;vrrp config&gt;service&gt;vprn&gt;if&gt;sap config&gt;service&gt;vprn&gt;if&gt;sap&gt;static-host config&gt;service&gt;vprn&gt;bgp config&gt;service&gt;vprn&gt;bgp&gt;group config&gt;service&gt;vprn&gt;bgp&gt;group&gt;neighbor config&gt;service&gt;vprn&gt;isis config&gt;service&gt;vprn&gt;isis&gt;if config&gt;service&gt;vprn&gt;mvpn&gt;provider-tunnel&gt;inclusive&gt;pim config&gt;service&gt;vprn&gt;ospf config&gt;service&gt;vprn&gt;ospf&gt;area&gt;if config&gt;service&gt;vprn&gt;ospf3 config&gt;service&gt;vprn&gt;ospf3&gt;area&gt;if config&gt;service&gt;vprn&gt;ospf3&gt;area&gt;virtual-link config&gt;service&gt;vprn&gt;ospf&gt;area&gt;virtual-link config&gt;service&gt;vprn&gt;ospf&gt;area&gt;sham-link config&gt;service&gt;vprn&gt;red-if&gt;spoke-sdp config&gt;service&gt;vprn&gt;rip config&gt;service&gt;vprn&gt;rip&gt;group config&gt;service&gt;vprn&gt;rip&gt;group&gt;neighbor config&gt;service&gt;vprn&gt;pim config&gt;service&gt;vprn&gt;pim&gt;if </pre>

```
config>service>vprn>pim>rp>bsr-candidate
config>service>vprn>pim>rp>ipv6>bsr-candidate
config>service>vprn>pim>rp>ipv6>embedded-rp
config>service>vprn>pim>rp>ipv6>rp-candidate
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp>proxy-server
config>service>vprn>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>arp-host
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vprn>dhcp>server>failover
config>service>vprn>nw-if>dhcp
config>service>vprn>nw-if>eth-cfm>mep
config>service>vprn>radius-proxy>server>cache
config>service>vprn>radius-proxy>server
config>service>vprn>radius-server
config>service>vprn>ipsec-if
config>service>vprn>ipsec-if>sap>tunnel
config>service>vprn>log>log-id
```

### 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.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

If the AS number was previously changed, the BGP AS number inherits the new value.

### Special Cases:

**Service Admin State** — Bindings to an SDP within the service will be put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.

A service is regarded as operational providing that one IP Interface SAP and one SDP is operational.

**VPRN BGP and RIP** — This command disables the BGP or RIP instance on the given IP interface. Routes learned from a neighbor that is shutdown are immediately removed from the BGP or RIP database and RTM. If BGP or RIP is globally shutdown, then all RIP group and neighbor interfaces are shutdown operationally. If a BGP or RIP group is shutdown, all member neighbor interfaces are shutdown operationally. If a BGP or RIP neighbor is shutdown, just that neighbor interface is operationally shutdown.

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>service>vprn>if>dhcp config>service>vprn>bgp config>service>vprn>rip config>service>vprn config>service>vprn>l2tp config>service>vprn>red-if config>service>vprn>if config>service>vprn>if>sap config>service>vprn>if>dhcp config>service>vprn>if>dhcp5 config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor config>service>vprn>rip config>service>vprn>ripng config>service>vprn>rip>group config>service>vprn>ripng>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng>group>neighbor config>service>vprn>subscriber-interface config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp config>service>vprn>sub-if>grp-if>sap>atm config>service>vprn>dhcp config>service>vprn>dhcp>server>pool config>service>vprn>sub-if>grp-if>pppoe config>service>vprn>nw-if config>service>vprn>radius-proxy>server config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command creates a text description stored in the configuration file for a configuration context. The <b>description</b> command associates a text string with a configuration context to help identify the content in the configuration file. The <b>no</b> form of this command removes the string from the configuration.
<b>Default</b>	No description associated with the configuration context.
<b>Parameters</b>	<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.

---

## Global Commands

### vprn

<b>Syntax</b>	<b>vprn</b> <i>service-id</i> [ <b>customer</b> <i>customer-id</i> ] [ <b>create</b> ] <b>no vprn</b> <i>service-id</i>						
<b>Context</b>	config>service						
<b>Description</b>	<p>This command creates or edits a Virtual Private Routed Network (VPRN) service instance.</p> <p>If the <i>service-id</i> does not exist, a context for the service is created. If the <i>service-id</i> exists, the context for editing the service is entered.</p> <p>VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.</p> <p>IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.</p> <p>When a service is created, the <b>customer</b> keyword and <i>customer-id</i> must be specified and associates the service with a customer. The <i>customer-id</i> must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.</p> <p>When a service is created, the use of the <b>customer</b> <i>customer-id</i> is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect <i>customer-id</i> results in an error.</p> <p>Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.</p> <p>The <b>no</b> form of the command deletes the VPRN service instance with the specified <i>service-id</i>. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shutdown and deleted.</p>						
<b>Default</b>	None — No VPRN service instances exist until they are explicitly created.						
<b>Parameters</b>	<p><i>service-id</i> — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The <i>service-id</i> must be the same number used for every 7750 SR on which this service is defined.</p> <table border="0"> <tr> <td style="vertical-align: top;"><b>Values</b></td> <td><i>service-id</i>:</td> <td>1 — 2147483648</td> </tr> <tr> <td></td> <td><i>svc-name</i>:</td> <td>64 characters maximum</td> </tr> </table> <p><b>customer</b> <i>customer-id</i> — Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.</p>	<b>Values</b>	<i>service-id</i> :	1 — 2147483648		<i>svc-name</i> :	64 characters maximum
<b>Values</b>	<i>service-id</i> :	1 — 2147483648					
	<i>svc-name</i> :	64 characters maximum					

**Values** 1 — 2147483647

## aggregate

<b>Syntax</b>	<p><b>aggregate</b> <i>ip-prefix/ip-prefix-length</i> [<b>summary-only</b>] [<b>as-set</b>] [<b>aggregator</b> <i>as-number:ip-address</i>] [<b>black-hole</b>] [<b>community</b> <i>comm-id</i>] [<b>description</b> <i>description</i>]</p> <p><b>aggregate</b> <i>ip-prefix/ip-prefix-length</i> [<b>summary-only</b>] [<b>as-set</b>] [<b>aggregator</b> <i>as-number:ip-address</i>] [<b>community</b> <i>comm-id</i>] [<b>indirect</b> <i>ip-address</i>] [<b>description</b> <i>description</i>]</p> <p><b>no aggregate</b> <i>ip-prefix/ip-prefix-length</i></p>				
<b>Context</b>	config>service>vprn				
<b>Description</b>	<p>This command creates an aggregate route.</p> <p>Use this command to automatically install an aggregate in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more-specific match of the aggregate.</p> <p>The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.</p> <p>Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.</p> <p>A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.</p> <p>By default aggregate routes are not installed in the forwarding table, however there are configuration options that allow an aggregate route to be installed with a black-hole next hop or with an indirect IP address as next hop.</p> <p>The <b>no</b> form of the command removes the aggregate.</p>				
<b>Default</b>	No aggregate routes are defined.				
<b>Parameters</b>	<p><i>ip-prefix</i> — The destination address of the aggregate route in dotted decimal notation.</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="vertical-align: top;"><b>Values</b></td> <td> <p>ipv4-prefix                    a.b.c.d (host bits must be 0)</p> <p>ipv4-prefix-length            0 — 32</p> <p>ipv6-prefix                    x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p>                                  x:x:x:x:x:d.d.d.d</p> <p>                                  x:            [0 — FFFF]H</p> <p>                                  d:            [0 — 255]D</p> <p>ipv6-prefix-length            0 — 128</p> </td> </tr> </table> <p>The mask associated with the network address expressed as a mask length.</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="vertical-align: top;"><b>Values</b></td> <td>0 — 32</td> </tr> </table> <p><b>summary-only</b> — This optional parameter suppresses advertisement of more specific component routes for the aggregate.</p>	<b>Values</b>	<p>ipv4-prefix                    a.b.c.d (host bits must be 0)</p> <p>ipv4-prefix-length            0 — 32</p> <p>ipv6-prefix                    x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p>                                  x:x:x:x:x:d.d.d.d</p> <p>                                  x:            [0 — FFFF]H</p> <p>                                  d:            [0 — 255]D</p> <p>ipv6-prefix-length            0 — 128</p>	<b>Values</b>	0 — 32
<b>Values</b>	<p>ipv4-prefix                    a.b.c.d (host bits must be 0)</p> <p>ipv4-prefix-length            0 — 32</p> <p>ipv6-prefix                    x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p>                                  x:x:x:x:x:d.d.d.d</p> <p>                                  x:            [0 — FFFF]H</p> <p>                                  d:            [0 — 255]D</p> <p>ipv6-prefix-length            0 — 128</p>				
<b>Values</b>	0 — 32				

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

**as-set** — This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS\_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

**aggregator** *as-number:ip-address* — This optional parameter specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

**community** *comm-id* — This configuration option associates a BGP community with the aggregate route. The community can be matched in route policies and is automatically added to BGP routes exported from the aggregate route.

<b>Values</b>	comm-id	asn:comm-val   well-known-comm
	asn	0 — 65535
	comm-val	0 — 65535
	well-known-comm	no-advertise, no-export, no-export-subconfed

**black-hole** — This optional parameter installs the aggregate route, when activated, in the FIB with a black-hole next-hop; where packets matching this route are discarded.

**indirect** *ip-address* — This configuration option specifies that the aggregate route should be installed in the FIB with a next-hop taken from the route used to forward packets to ip-address.

<b>Values</b>	ipv4-prefix	a.b.c.d
	ipv6-prefix	x:x:x:x:x:x:x
		x:x:x:x:x:d.d.d.d
		x: [0 — FFFF]H
		d: [0 — 255]D

**description** *description-text* — Specifies a text description stored in the configuration file for a configuration context.

## allow-export-bgp-vprn

**Syntax** **allow-export-bgp-vprn**  
**no allow-export-bgp-vprn**

**Context** config>service>vprn

**Description** This command causes the vrf-export and vrf-target functions of the VPRN to include BGP-VPN routes installed in the VPRN route table. For split-horizon reasons, these routes are normally not re-advertisable as VPN-IP routes.

When a BGP-VPN route is re-exported, the route-distinguisher and label values are rewritten per the configuration of the re-exporting VPRN.

**Note** — Appropriate measures should be taken to ensure that routing updates do not loop back to the source when this command is used, otherwise the routes could become unstable.

**Default** **no allow-export-bgp-vprn**

## auto-bind-tunnel

<b>Syntax</b>	<b>auto-bind-tunnel</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command enables the context to configure automatic binding of a VPRN service using tunnels to MP-BGP peers.</p> <p>The <b>auto-bind-tunnel</b> node is simply a context to configure the binding of VPRN routes to tunnels. The user must configure the <b>resolution</b> option to enable auto-bind resolution to tunnels in TTM. If the <b>resolution</b> option is explicitly set to <b>disabled</b>, the auto-binding to tunnel is removed.</p> <p>If <b>resolution</b> is set to <b>any</b>, any supported tunnel type in VPRN context will be selected following TTM preference. If one or more explicit tunnel types are specified using the <b>resolution-filter</b> option, then only these tunnel types will be selected again following the TTM preference.</p> <p>The following tunnel types are supported in a VPRN context in order of preference: RSVP, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the <b>auto-bind-tunnel</b> context.</p> <p>The <b>ldp</b> value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.</p> <p>The <b>rsvp</b> value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.</p> <p>When the <b>sr-isis (sr-ospf)</b> value is enabled, a SR tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.</p> <p>The user must set <b>resolution</b> to <b>filter</b> to activate the list of tunnel-types configured under <b>resolution-filter</b>.</p> <p>When an explicit SDP to a BGP next-hop is configured in a VPRN service (<b>configure&gt;service&gt;vprn&gt;spoke-sdp</b>), it overrides the <b>auto-bind-tunnel</b> selection for that BGP next-hop only. There is no support for reverting automatically to the <b>auto-bind-tunnel</b> selection if the explicit SDP goes down. The user must delete the explicit spoke-sdp in the VPRN service context to resume using the <b>auto-bind-tunnel</b> selection for the BGP next-hop.</p>

## resolution

<b>Syntax</b>	<b>resolution {any   filter   disabled}</b>
<b>Context</b>	config>service>vprn>auto-bind-tunnel
<b>Description</b>	<p>This command configures the resolution mode in the automatic binding of a VPRN service to tunnels to MP-BGP peers.</p>

## Global Commands

- Parameters** **any** — Enables the binding to any supported tunnel type in VPRN context following TTM preference.
- filter** — Enables the binding to the subset of tunnel types configured under resolution-filter.
- disabled** — Disables the automatic binding of a VPRN service to tunnels to MP-BGP peers.

## resolution-filter

- Syntax** **resolution-filter**
- Context** config>service>vprn>auto-bind-tunnel
- Description** This command configures the subset of tunnel types which can be used in the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.
- The following tunnel types are supported in a VPRN context in order of preference: RSVP, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context.
- Parameters** **gre** — Selects the GRE tunnel type.
- ldp** — Selects the LDP tunnel type.
- rsvp** — Selects the RSVP-TE tunnel type.
- sr-isis** — Selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM.

## autonomous-system

- Syntax** **autonomous-system** *as-number*  
**no autonomous-system**
- Context** config>service>vprn
- Description** This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). This command defines the autonomous system to be used by this VPN routing.
- The **no** form of the command removes the defined AS from this VPRN context.
- Default** no autonomous-system
- Parameters** *as-number* — Specifies the AS number for the VPRN service.
- Values** 1 — 4294967295

## backup-path

- Syntax** [**no**] **backup-path** [**ipv4**] [**ipv6**]
- Context** config>router



```
config>service>vprn
```

**Description** This command enables the computation and use of a backup path for IPv4 and/or IPv6 BGP-learned prefixes belonging to the base router or a particular VPRN. Multiple paths must be received for a prefix in order to take advantage of this feature. When a prefix has a backup path and its primary path(s) fail the affected traffic is rapidly diverted to the backup path without waiting for control plane re-convergence to occur. When many prefixes share the same primary path(s), and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes. In some cases prefix independent convergence may require use of FP2 or later IOMs/IMMs/XMAs.

By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.

**Default** no backup-path

**Parameters** **ipv4** — Enables the use of a backup path for BGP-learned IPv4 prefixes

**ipv6** — Enables the use of a backup path for BGP-learned IPv6 prefixes

## carrier-carrier-vpn

<b>Syntax</b>	<b>[no] carrier-carrier-vpn</b>
<b>Context</b>	config>service>vpn
<b>Description</b>	<p>This command configures a VPRN service to support a Carrier Supporting Carrier model. It should be configured on a network provider's CSC-PE device.</p> <p>This command cannot be applied to a VPRN unless it has no SAP or spoke-SDP interfaces. Once this command has been entered one or more MPLS-capable CSC interfaces can be created in the VPRN.</p> <p>The <b>no</b> form of the command removes the Carrier Supporting Carrier capability from a VPRN.</p>
<b>Default</b>	no carrier-carrier-vpn

## confederation

<b>Syntax</b>	<b>confederation</b> <i>confed-as-num</i> <b>members</b> <i>as-number</i> [ <i>as-number...</i> (up to 15 max)] <b>no confederation</b> <i>confed-as-num</i> <b>members</b> <i>as-number</i> [ <i>as-number...</i> (up to 15 max)] <b>no confederation</b>
<b>Context</b>	config>service>vpn
<b>Description</b>	<p>This command configures the VPRN BGP instance to participate in a BGP confederation. BGP confederations can be used to reduce the number of IBGP sessions required within an AS.</p> <p>When a VPRN BGP instance is part of a confederation, it can form confederation-EBGP sessions with CE router peers in a different sub-autonomous systems of the same confederation as well as regular EBGP sessions with CE router peers outside the confederation. A VPRN BGP instance that is part of a confederation cannot import or export its routes to the base router instance (as VPN-IP routes).</p> <p>The <b>no</b> form of the command deletes the specified member AS from the confederation. When members are not specified in the no statement, the entire list is removed and confederations is disabled. When the last member of the list is removed, confederations is disabled.</p>
<b>Default</b>	No confederations are defined.
<b>Parameters</b>	<p><i>confed-as-num</i> — The confederation AS number defined as a decimal value.</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>members</b> <i>as-number</i> — The AS number(s) that are members of the confederation, each expressed as a decimal integer. Configure up to 15 members per confed-as-num.</p> <p><b>Values</b> 1 — 4294967295</p>

## dns

<b>Syntax</b>	<b>[no] dns</b>
<b>Context</b>	config>service>vpn

**Description** This command enables the context to configure domain name servers.  
The **no** form of the command disables DNS for this service.

## ipv4-source-address

**Syntax** **ipv4-source-address** *ipv4-address*  
**no ipv4-source-address**

**Context** config>service>vprn>dns

**Description** This command configures the IPv4 address of the default secondary DNS server for the subscribers using this interface. Subscribers that cannot obtain an IPv4 DNS server address by other means, can use this for DNS name resolution.

The *ipv4-address* value can only be set to a nonzero value if the value of VPRN type is set to **subscriber-split-horizon**.

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

**Default** none

**Parameters** *ipv4-address* — Specifies the IPv4 address of the default secondary DNS server.

**Values** *ipv4-address* - a.b.c.d

## ipv6-source-address

**Syntax** **ipv6-source-address** *ipv6-address*  
**no ipv6-source-address**

**Context** config>service>vprn>dns

**Description** This command configures the IPv6 address of the default secondary DNS server for the subscribers using this interface. Subscribers that cannot obtain an IPv6 DNS server address by other means, can use this for DNS name resolution.

The *ipv6-address* value can only be set to a nonzero value if the value of VPRN type is set to **subscriber-split-horizon**.

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

**Default** none

**Parameters** *ipv4-address* — Specifies the IPv6 address of the default secondary DNS server.

**Values** *ipv4-address* - a.b.c.d

## primary-dns

<b>Syntax</b>	<b>primary-dns</b> <i>ip-address</i> <b>no primary-dns</b>
<b>Context</b>	config>service>vprn>dns
<b>Description</b>	This command configures the primary DNS server used for DNS name resolution. DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files. The <b>no</b> form of the command removes the primary DNS server from the configuration.
<b>Default</b>	<b>no primary-dns</b> — No primary DNS server is configured.
<b>Parameters</b>	<i>ip-address</i> — The IP or IPv6 address of the primary DNS server.
<b>Values</b>	ipv4-address - a.b.c.d ipv6-address:    x::x::x::x::x::x[-interface] x::x::x::x::x::x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface - 32 chars max, for link local addresses.

## secondary-dns

<b>Syntax</b>	<b>secondary-dns</b> <i>ip-address</i> <b>no secondary-dns</b>
<b>Context</b>	config>service>vprn>dns
<b>Description</b>	This command configures the secondary DNS server for DNS name resolution. The secondary DNS server is used only if the primary DNS server does not respond. DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files. The <b>no</b> form of the command removes the secondary DNS server from the configuration.
<b>Default</b>	<b>no secondary-dns</b> — No secondary DNS server is configured.
<b>Parameters</b>	<i>ip-address</i> — The IP or IPv6 address of the secondary DNS server.
<b>Values</b>	ipv4-address - a.b.c.d ipv6-address:    x::x::x::x::x::x[-interface] x::x::x::x::x::x:d.d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface - 32 chars max, for link local addresses

## tertiary-dns

<b>Syntax</b>	<b>tertiary-dns</b> <i>ip-address</i> <b>no tertiary-dns</b>												
<b>Context</b>	config>service>vprn>dns												
<b>Description</b>	<p>This command configures the tertiary DNS server for DNS name resolution. The tertiary DNS server is used only if the primary DNS server and the secondary DNS server do not respond.</p> <p>DNS name resolution can be used when executing ping, traceroute, and service-ping, and also when defining file URLs. DNS name resolution is not supported when DNS names are embedded in configuration files.</p> <p>The <b>no</b> form of the command removes the tertiary DNS server from the configuration.</p>												
<b>Default</b>	<b>no tertiary-dns</b> — No tertiary DNS server is configured.												
<b>Parameters</b>	<i>ip-address</i> — The IP or IPv6 address of the tertiary DNS server.												
	<table> <tr> <td><b>Values</b></td> <td>ipv4-address - a.b.c.d</td> </tr> <tr> <td></td> <td>ipv6-address: x:x:x:x:x:x:x[-interface]</td> </tr> <tr> <td></td> <td>x:x:x:x:x:d.d.d.d[-interface]</td> </tr> <tr> <td></td> <td>x: [0..FFFF]H</td> </tr> <tr> <td></td> <td>d: [0..255]D</td> </tr> <tr> <td></td> <td>interface - 32 chars max, for link local addresses</td> </tr> </table>	<b>Values</b>	ipv4-address - a.b.c.d		ipv6-address: x:x:x:x:x:x:x[-interface]		x:x:x:x:x:d.d.d.d[-interface]		x: [0..FFFF]H		d: [0..255]D		interface - 32 chars max, for link local addresses
<b>Values</b>	ipv4-address - a.b.c.d												
	ipv6-address: x:x:x:x:x:x:x[-interface]												
	x:x:x:x:x:d.d.d.d[-interface]												
	x: [0..FFFF]H												
	d: [0..255]D												
	interface - 32 chars max, for link local addresses												

## ecmp

<b>Syntax</b>	<b>ecmp</b> <i>max-ecmp-routes</i> <b>no ecmp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command enables equal-cost multipath (ECMP) and configures the number of routes for path sharing. For example, the value of 2 means that 2 equal cost routes will be used for cost sharing.</p> <p>ECMP groups form when the system routes to the same destination with equal cost values. Routing table entries can be entered manually (as static routes), or they can be formed when neighbors are discovered and routing table information is exchanged by routing protocols. The system can balance traffic across the groups with equal costs.</p> <p>ECMP can only be used for routes learned with the same preference and same protocol. See the discussion on preferences in the <b>static-route</b> command.</p> <p>When more ECMP routes are available at the best preference than configured by the <i>max-ecmp-routes</i> parameter, then the lowest next-hop IP address algorithm is used to select the number of routes configured.</p> <p>The <b>no</b> form of the command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, the newly updated route is used.</p>
<b>Default</b>	no ecmp
<b>Parameters</b>	<i>max-ecmp-routes</i> — Specifies the maximum number of routes for path sharing.

**Values** 0—32

## enable-bgp-vpn-backup

<b>Syntax</b>	<b>enable-bgp-vpn-backup [ipv4] [ipv6]</b> <b>no enable-bgp-vpn-backup</b>
<b>Context</b>	config>service>vprn>bgp
<b>Description</b>	This command allows BGP-VPN routes imported into the VPRN to be used as backup paths for IPv4 and/or IPv6 BGP-learned prefixes.
<b>Parameters</b>	<b>ipv4</b> — Allows BGP-VPN routes to be used as backup paths for IPv4 prefixes. <b>ipv6</b> — Allows BGP-VPN routes to be used as backup paths for IPv6 prefixes.

## fib-priority

<b>Syntax</b>	<b>fib-priority {high   standard}</b>
<b>Context</b>	config>service>vpr
<b>Description</b>	This command specifies the FIB priority for VPRN.
<b>Parameters</b>	<b>high</b> — Specifies high FIB priority for VPRN. <b>standard</b> — Specifies standard FIB priority for VPRN.

## grt-lookup

<b>Syntax</b>	<b>grt-lookup</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command provides the context under which all Global Route Table (GRT) leaking commands are configured. If all the supporting commands in the context are removed, this command will also be removed.

## enable-grt

<b>Syntax</b>	<b>[no] enable-grt</b>
<b>Context</b>	config>service>vprn>grt-lookup
<b>Description</b>	This command enables the functions required for looking up routes in the Global Route Table (GRT) when the lookup in the local VRF fails. If this command is enabled without the use of a static-route option (as subcommand to this parent), a lookup in the local VRF is preferred over the GRT. When the local VRF returns no route table lookup matches, the result from the GRT is preferred.

The no form of this command disables the lookup in the GRT when the lookup in the local VRF fails.

**Default** no enable-grt

## export-grt

**Syntax** **export-grt** *policy-name* [*policy-name* ...(up to 5 max)]  
**no export-grt**

**Context** config>service>vprn>grt-lookup

**Description** This command uses route policy to determine which routes are exported from the VRF to the GRT along with all the forwarding information. These entries will be marked as BGP-VPN routes in the GRT. Routes must be in the GRT in order for proper routing to occur from the GRT to the VRF.

**Default** no export-grt

## export-limit

**Syntax** **export-limit** *num-routes*  
**no export-limit**

**Context** config>service>vprn>grt-lookup  
config>service>vprn>ospf  
config>service>vprn>ospf3  
config>service>vprn>rip

**Description** This command provides the ability to limit the total number of routes exported from the VRF to the GRT. The value zero (0) provides an override that disables the maximum limit. Setting this value to zero (0) will not limit the number of routes exported from the VRF to the GRT. Configuring a range of one (1) to 1000 will limit the number of routes to the specified value.

The **no** form of the command sets the export-limit to a default of five (5).

**Default** export-limit 5

**Parameters** *num-routes* — Specifies maximum number of routes that can be exported.

**Values** 0 — 1000

## export-v6-limit

**Syntax** **export-v6-limit** *num-routes*  
**no export-v6-limit**

**Context** config>service>vprn>grt-lookup

**Description** The export-limit range provides the ability to limit the total number of IPv6 routes exported from the VPRN to the GRT. The value “0” provides an override that disables the maximum limit. Setting this

value to “0” will not limit the the number of routes exported from the VPRN to the GRT. Configuring a range of 1-1000 will limit the number of routes to the specified value.

The **no** form of the command sets the export-limit to a default of 5.

<b>Default</b>	export-v6-limit 5
<b>Parameters</b>	<i>num-routes</i> — Specifies maximum number of routes that can be exported.
<b>Values</b>	0 — 1000

## allow-local-management

<b>Syntax</b>	<b>[no] allow-local-management</b>
<b>Context</b>	config>service>vprn>grt-lookup>enable-grt
<b>Description</b>	When enabled, both IPv4 and IPv6 base interfaces shall respond to leaked traffic from the VPRN.

## static-route

<b>Syntax</b>	<b>static-route</b> <i>{ip-prefix/prefix-length   ip-prefix netmask}</i> [ <b>preference preference</b> ] [ <b>metric metric</b> ] [ <b>enable disable</b> ] <b>grt</b> <b>no static-route</b>
<b>Context</b>	config>service>vprn>grt-lookup>enable-grt
<b>Description</b>	This command is a simplified version of the traditional static-route command pointing to the base routing instance. This instructs the route lookup function to look only in the GRT for a route matching destination static route and not look up the route in the local VPRN. The GRT keyword is a required parameter.  The no form casues the feature into the default mode of primary lookup for all routes in the local VPRN and failing a match in the local VPRN, the lookup result in the GRT will be used.
<b>Parameters</b>	<i>ip-prefix/prefix-length</i> — Specifies the IPv4 prefix and prefix length.  <b>Values</b> ip-prefix: a.b.c.d (host bits must be 0) ipv4-prefix-length: [0..32]  ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D ipv6-prefix-length 0 — 128  <i>netmask</i> — Specifies the netmask.  <b>Values</b> a.b.c.d (network bits all 1 and host bits all 0)  <i>preference</i> — Specifies the preference.



**Values** 1..255

*metric* — Specifies the metric.

**Values** 0..65535

**enable|disable** — Keyword; specifies the state of the static-route.

**grt** — Keyword; Global Route Table lookup.

## gsmp

<b>Syntax</b>	<b>gsmp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure GSMP connections maintained in this service.
<b>Default</b>	not enabled

## group

<b>Syntax</b>	<b>[no] group <i>name</i></b>
<b>Context</b>	config>service>vprn>gsmp
<b>Description</b>	This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.
<b>Parameters</b>	<i>name</i> — Specifies the group name up to 32 characters in length.

## ancp

<b>Syntax</b>	<b>ancp</b>
<b>Context</b>	config>service>vprn>gsmp>group
<b>Description</b>	This command configures ANCP parameters for this GSMP group.

## dynamic-topology-discover

<b>Syntax</b>	<b>[no] dynamic-topology-discover</b>
<b>Context</b>	config>service>vprn>gsmp>group>ancp
<b>Description</b>	This command enables the ANCP dynamic topology discovery capability. The <b>no</b> form of this command disables the feature.

## oam

<b>Syntax</b>	<b>[no] oam</b>
<b>Context</b>	config>service>vprn>gsmp>group>ancp
<b>Description</b>	This command specifies whether or not the GSMP ANCP OAM capability should be negotiated at startup of the GSMP connection. The <b>no</b> form of this command disables the feature.

## hold-multiplier

<b>Syntax</b>	<b>hold-multiplier multiplier</b> <b>no hold-multiplier</b>
<b>Context</b>	config>service>vprn>gsmp>group
<b>Description</b>	This command configures the hold-multiplier for the GSMP connections in this group.
<b>Parameters</b>	<i>multiplier</i> — Specifies the GSMP hold multiplier value. <b>Values</b> 1 — 100

## idle-filter

<b>Syntax</b>	<b>idle-filter</b> <b>no idle-filter</b>
<b>Context</b>	config>service>vpls>gsmp config>service>vprn>gsmp
<b>Description</b>	This command when applied will filter out new subscriber's ANCP messages from subscriber with "DSL-line-state" IDLE
<b>Default</b>	no idle-filter

## keepalive

<b>Syntax</b>	<b>keepalive seconds</b> <b>no keepalive</b>
<b>Context</b>	config>service>vprn>gsmp>group
<b>Description</b>	This command configures keepalive values for the GSMP connections in this group.
<b>Parameters</b>	<i>seconds</i> — Specifies the GSMP keepalive timer value in seconds. <b>Values</b> 1 — 25

## neighbor

<b>Syntax</b>	<b>[no] neighbor</b> <i>ip-address</i>
<b>Context</b>	config>service>vprn>gsmp>group
<b>Description</b>	This command adds or removes a neighbor in this group.
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address in dotted decimal notation.

## local-address

<b>Syntax</b>	<b>local-address</b> <i>ip-address</i> <b>no local-address</b>
<b>Context</b>	config>service>vprn>gsmp>group>neighbor
<b>Description</b>	This command configures the source ip-address used in the connection towards the neighbor.
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address in dotted decimal notation.

## priority-marking

<b>Syntax</b>	<b>priority-marking dscp</b> <i>dscp-name</i> <b>priority-marking prec</b> <i>ip-prec-value</i> <b>no priority-marking</b>
<b>Context</b>	config>service>vprn>gsmp>group>neighbor
<b>Description</b>	This command configures the type of priority marking to be used.
<b>Parameters</b>	<b>dscp</b> <i>dscp-name</i> — Specifies the DSCP code-point to be used.  <b>Values</b> be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63  <b>prec</b> <i>ip-prec-value</i> — Specifies the precedence value to be used.  <b>Values</b> 0 — 7

## persistency-database

<b>Syntax</b>	<b>persistency-database</b> <b>no persistency-database</b>
<b>Context</b>	config>service>vpls <service id>gsmp config>service>vprn<service id>gsmp
<b>Description</b>	This command enables the system to store DSL line information in memory. If the GSMP connection terminates, the DSL line information will remain in memory and accessible for Radius authentication and accounting.
<b>Default</b>	no persistency-database

---

## IGMP Commands

### igmp

<b>Syntax</b>	<b>[no] igmp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure IGMP parameters. The <b>no</b> form of the command disables IGMP.
<b>Default</b>	disabled

### group-interface

<b>Syntax</b>	<b>[no] group-interface</b> <i>ip-int-name</i> <b>[no] group-interface fwd-service</b> <i>service-id ip-int-name</i>
<b>Context</b>	config>service>vprn>igmp
<b>Description</b>	This command configures IGMP group interfaces. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. <b>fwd-service</b> <i>service-id</i> — Specifies the service ID. This is only configured in the retailer VRF. This construct references the wholesaler service under which the group-interface (and the subscriber) is actually defined.
<b>Values</b>	1 — 2147483650, svc-name up to 64 char maximum
<b>Default</b>	none

### disable-router-alert-check

<b>Syntax</b>	<b>[no] disable-router-alert-check</b>
<b>Context</b>	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
<b>Description</b>	This command enables the IGMP router alert check option. The <b>no</b> form of the command disables the router alert check.

## import

<b>Syntax</b>	<b>import</b> <i>policy-name</i> <b>no import</b>
<b>Context</b>	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
<b>Description</b>	This command specifies the policy that is to be applied on this interface.
<b>Parameters</b>	<i>policy-name</i> — Specify the policy to filter IGMP packets.

## max-groups

<b>Syntax</b>	<b>max-groups</b> <i>value</i> <b>no max-groups</b>
<b>Context</b>	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
<b>Description</b>	This command configures the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.  The <b>no</b> form of the command removes the value.
<b>Parameters</b>	<i>value</i> — Specifies the maximum number of groups for this interface.  <b>Values</b> 1 — 16000

## max-sources

<b>Syntax</b>	<b>max-sources</b> [1..1000] <b>no max-sources</b>
<b>Context</b>	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
<b>Description</b>	This command specifies the maximum number of sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of sources, the sources that are already accepted are not deleted. Only new sources will not be allowed.
<b>Parameters</b>	<i>sources</i> — Specifies the maximum number of sources for this interface.  <b>Values</b> 1 — 1000

## max-grp-sources

<b>Syntax</b>	<b>max-grp-sources</b> [1..32000] <b>no max-grp-sources</b>
<b>Context</b>	config>service>vprn>igmp>gr-if config>service>vprn>igmp>if
<b>Description</b>	This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed.  The no form of the command reverts to the default.
<b>Default</b>	0
<b>Parameters</b>	1 — 32000 — Specifies the maximum number of group source.
<b>Values</b>	1 — 32000

## mcac

<b>Syntax</b>	<b>mcac</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command enables the context to configure multicast CAC parameters.

## mc-constraints

<b>Syntax</b>	<b>mc-constraints</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command configures multicast CAC constraints.

## policy

<b>Syntax</b>	<b>policy</b> <i>policy-name</i> <b>no policy</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command references the global channel bandwidth definition policy that is used for (H)mcac and HQoS Adjust.  HQoS Adjustment is supported only with redirection enabled. In other words, the policy from the redirected interface is used for HQoS Adjustment.

Hierarchical mcac (Hmccac) is supported only with redirection enabled. In Hmccac, the subscriber is checked first against its bandwidth limits followed by the check on the redirected interface against the bandwidth limits defined under the redirected interface. In the Hmccac case the channel definition policy must be referenced under the redirected interface level.

<b>Parameters</b>	<i>policy-name</i> — Specifies the name of the global mcac channel definition policy defined under the hierarchy <b>configure&gt;router&gt;mcac&gt;policy</b> .
<b>Default</b>	No policy is referenced.

## unconstrained-bw

<b>Syntax</b>	<b>unconstrained-bw</b> <i>bandwidth</i> <b>mandatory-bw</b> <i>mandatory-bw</i> <b>no unconstrained-bw</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command configures unconstrained-bw for multicast cac policy on this interface. The <b>no</b> form of the command
<b>Parameters</b>	<p><i>bandwidth</i> — Specifies the bandwidth assigned for interface's multicast cac policy traffic in kilo-bits per second (kbps).</p> <p><b>Values</b> 0 — 2147483647</p> <p><b>mandatory-bw</b> <i>mandatory-bw</i> — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).</p> <p>If the <i>bandwidth</i> value is 0, no mandatory channels are allowed. If <i>bandwidth</i> is not configured, then all mandatory and optional channels are allowed.</p> <p>If the value of <i>mandatory-bw</i> is equal to the value of <i>bandwidth</i>, then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.</p> <p>The value of <i>mandatory-bw</i> should always be less than or equal to that of <i>bandwidth</i>. An attempt to set the value of <i>mandatory-bw</i> greater than that of <i>bandwidth</i>, will result in inconsistent value error.</p> <p><b>Values</b> 0 — 2147483647</p>

## query-src-ip

<b>Syntax</b>	<b>query-src-ip</b> <i>ip-address</i> <b>no query-src-ip</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command configures the query source IP address for the group interface. This IP address overrides the source IP address configured at the router level. The <b>no</b> form of the command removes the IP address.



<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Sets the source IPv4 address for all subscriber's IGMP queries.

## sub-hosts-only

<b>Syntax</b>	<b>[no] sub-hosts-only</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command enables the IGMP traffic from known hosts only. The <b>no</b> form of the command disable the IGMP traffic from known hosts only

## subnet-check

<b>Syntax</b>	<b>[no] subnet-check</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command enables local subnet checking for IGMP. The <b>no</b> form of the command disables local subnet checking for IGMP.

## version

<b>Syntax</b>	<b>version</b> <i>version</i> <b>no version</b>
<b>Context</b>	config>service>vprn>igmp>gr-if
<b>Description</b>	This command configures the version of IGMP. The <b>no</b> form of the command
<b>Parameters</b>	<i>version</i> — Specifies the IGMP version. <b>Values</b> 1, 2 or 3

## grp-if-query-src-ip

<b>Syntax</b>	<b>grp-if-query-src-ip</b> <i>ip-address</i> <b>no grp-if-query-src-ip</b>
<b>Context</b>	config>service>vprn>igmp
<b>Description</b>	This command configures the query source IP address for all group interfaces. The <b>no</b> form of the command removes the IP address.

## Global Commands

**Default** none

## interface

**Syntax** **interface** *ip-int-name*  
**no interface**

**Context** config>service>vprn>igmp

**Description** This command enables the context to configure IGMP interface parameters.

**Parameters** *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

**Values** 1 — 32 characters maximum

## import

**Syntax** **import** *policy-name*  
**no import**

**Context** config>service>vprn>igmp>if

**Description** This command imports a policy to filter IGMP packets. The **no** form of the command removes the policy association from the IGMP instance.

**Default** **no import** — No import policy specified.

**Parameters** *policy-name* — The import route policy name. Allowed values are any string up to 32 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.

The specified name(s) must already be defined.

## max-groups

**Syntax** **max-groups** *value*  
**no max-groups**

**Context** config>service>vprn>igmp>if

**Description** This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

**Default** 0, no limit to the number of groups.

**Parameters** *value* — Specifies the maximum number of groups for this interface.

**Values** 1 — 16000

## mcac

**Syntax** **mcac**

**Context** config>service>vprn>if  
config>service>vprn>pim>if

**Description** This command configures multicast CAC policy and constraints for this interface.

**Default** none

## mc-constraints

**Syntax** **mc-constraints**

**Context** config>service>vprn>igmp>if>mcac  
config>service>vprn>pim>if>mcac

**Description** This command enables the context to configure multicast CAC constraints.

**Default** none

## level

**Syntax** **level** *level-id* **bw** *bandwidth*  
**no level** *level-id*

**Context** config>service>vprn>igmp>if>mcac  
config>service>vprn>pim>if>mcac

**Description** This command configures interface levels and associated bandwidth for multicast CAC policy.

**Parameters** *level-id* — Specifies an entry for the multicast CAC policy constraint level configured on this system.

**Values** 1 — 8

*bandwidth* — Specifies the bandwidth in kilobits per second (kbps) for the level.

**Values** 1 — 2147483647

## number-down

**Syntax** `number-down number-lag-port-down level level-id`  
`no number-down`

**Context** `config>service>vprn>igmp>if>mcac>mc-constraints`  
`config>service>vprn>pim>if>mcac>mc-constraints`

**Description** This command configures the number of ports down and level for interface's multicast CAC policy.

**Default** not enabled

**Parameters** *number-lag-port-down* — If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface will be as per the levels configured in this context.

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

*level-id* — Specifies an entry for the multicast CAC policy constraint level configured on this system.

**Values** 1 — 8

## use-lag-port-weight

**Syntax** `[no] use-lag-port-weight`

**Context** `config>service>vprn>igmp>interface>mcac>mc-constraints`  
`config>service>vprn>mld>interface>mcac>mc-constraints`  
`config>service>vprn>pim>interface>mcac>mc-constraints`

**Description** This command enables port weight to be used when determining available bandwidth per level when LAG ports go down/come up. The command is required for proper operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.

**Default** `no use-lag-port-weight` — port number is used when determining available bandwidth per level when LAG ports go down/come up.

## policy

**Syntax** `policy policy-name`  
`no policy`

**Context** `config>service>vprn>igmp>if>mcac`  
`config>service>vprn>pim>if>mcac`

**Description** This command configures the mulitcast CAC policy name.

**Parameters** *policy-name* — The multicast CAC policy name. Allowed values are any string up to 32 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.

## unconstrained-bw

**Syntax** **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*  
**no unconstrained-bw**

**Context** config>service>vprn>igmp>if>mcac  
 config>service>vprn>pim>if>mcac

**Description** This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled (**no unconstrained-bw**) there will be no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the **unconstrained-bw** minus the **mandatory-bw** and the mandatory channels have to stay below the specified value for the **mandatory-bw**. After this interface check, the bundle checks are performed.

**Parameters** *bandwidth* — The bandwidth assigned for interface's MCAC policy traffic, in kilo-bits per second (kbps).

**Values** 0 — 2147483647

**mandatory-bw** *mandatory-bw* — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilo-bits per second (kbps).

If the *bandwidth* value is 0, no mandatory channels are allowed. If *bandwidth* is not configured, then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.

The value of *mandatory-bw* should always be less than or equal to that of *bandwidth*. An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

**Values** 0 — 2147483647

## static

**Syntax** **static**

**Context** config>service>vprn>igmp>if

**Description** This command tests forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

**Default** none

## group

**Syntax** [**no**] **group** *grp-ip-address*

**Context** config>service>vprn>igmp>if>static

## Global Commands

<b>Description</b>	<p>This command adds a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.</p> <p>When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>grp-ip-address</i> — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group. The address must be in dotted decimal notation

## source

<b>Syntax</b>	<b>source</b>
<b>Context</b>	config>service>vprn>igmp>if>static>group
<b>Description</b>	<p>This command specifies a IPv4 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group is to receive multicast traffic from, and from the source(s) that the traffic is expected.</p> <p>The <b>source</b> command is mutually exclusive with the specification of individual sources for the same group.</p> <p>The source command in combination with the group is used to create a specific (S,G) static group entry.</p> <p>Use the <b>no</b> form of the command to remove the source from the configuration.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IPv4 unicast address.

## starg

<b>Syntax</b>	<b>starg</b>
<b>Context</b>	config>service>vprn>igmp>if>static>group
<b>Description</b>	<p>This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.</p> <p>Use the <b>no</b> form of the command to remove the starg entry from the configuration.</p>
<b>Default</b>	none

## subnet-check

<b>Syntax</b>	<b>[no] subnet-check</b>
<b>Context</b>	config>service>vprn>igmp>if
<b>Description</b>	This command enables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.
<b>Default</b>	enabled

## version

<b>Syntax</b>	<b>version <i>version</i></b> <b>no version</b>
<b>Context</b>	config>service>vprn>igmp>if
<b>Description</b>	This command specifies the IGMP version. If routers run different versions of IGMP, they will negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP on that LAN.  For IGMPv3, note that a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.
<b>Default</b>	3
<b>Parameters</b>	<i>version</i> — Specifies the IGMP version number.
	<b>Values</b> 1, 2, 3

## query-interval

<b>Syntax</b>	<b>query-interval</b> <i>seconds</i> <b>no query-interval</b>
<b>Context</b>	config>service>vprn>igmp
<b>Description</b>	This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.
<b>Default</b>	125
<b>Parameters</b>	<i>seconds</i> — The time frequency, in seconds, that the router transmits general host-query messages. <b>Values</b> 2 — 1024

## query-last-member-interval

<b>Syntax</b>	<b>query-last-member-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vprn>igmp
<b>Description</b>	This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.
<b>Default</b>	1
<b>Parameters</b>	<i>seconds</i> — Specifies the frequency, in seconds, at which query messages are sent. <b>Values</b> 1 — 1024

## query-response-interval

<b>Syntax</b>	<b>query-response-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vprn>igmp
<b>Description</b>	This command specifies how long the querier router waits to receive a response to a host-query message from a host.
<b>Default</b>	10
<b>Parameters</b>	<i>seconds</i> — Specifies the the length of time to wait to receive a response to the host-query message from the host. <b>Values</b> 1 — 1023



## robust-count

<b>Syntax</b>	<b>robust-count</b> <i>robust-count</i> <b>no robust-count</b>
<b>Context</b>	config>service>vprn>igmp
<b>Description</b>	This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.
<b>Default</b>	2
<b>Parameters</b>	<i>robust-count</i> — Specifies the robust count value.
	<b>Values</b> 2 — 10

## ssm-translate

<b>Syntax</b>	<b>igmp</b>
<b>Context</b>	config>service>vprn>igmp config>service>vprn>igmp>if
<b>Description</b>	This command enables the context to configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the starg command is not enabled. An error message is generated if you try to configure the <b>source</b> command with <b>starg</b> command enabled.

## grp-range

<b>Syntax</b>	<b>[no] grp-range</b> <i>start end</i>
<b>Context</b>	config>service>vprn>igmp>ssm-translate
<b>Description</b>	This command is used to configure group ranges which are translated to SSM (S,G) entries.
<b>Parameters</b>	<i>start</i> — An IP address that specifies the start of the group range. <i>end</i> — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the <i>start</i> value.

## source

<b>Syntax</b>	<b>[no] source</b> <i>ip-address</i>
<b>Context</b>	config>service>vprn>igmp>ssm-translate>grp-range

## Global Commands

**Description** This command specifies the source IP address for the group range. Whenever a (\*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

**Parameters** *ip-address* — Specifies the IP address that will be sending data.

## igmp-host-tracking

**Syntax** **igmp-host-tracking**

**Context** config>service>vprn  
config>service>vprn>sap

**Description** This command enables the context to configure IGMP host tracking parameters.

## expiry-time

**Syntax** **expiry-time** *expiry-time*  
**no expiry-time**

**Context** config>service>vprn>igmp-trk  
config>service>vprn>sap>igmp-trk

**Description** This command configures the time that the system continues to track inactive hosts. The **no** form of the command removes the values from the configuration.

**Default** no expiry-time

**Parameters** *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host.

**Values** 1 — 65535

## import

**Syntax** **import** *policy-name*  
**no import**

**Context** config>service>vprn>sap>igmp-trk

**Description** This command associates an import policy to filter IGMP packets. The **no** form of the command removes the values from the configuration.

**Default** no import

**Parameters** *policy-name* — Specifies the import policy name.

## max-num-groups

<b>Syntax</b>	<b>max-num-groups</b> <i>max-num-groups</i> <b>no max-num-groups</b>
<b>Context</b>	config>service>vprn>sap>igmp-trk
<b>Description</b>	This command configures the maximum number of multicast groups allowed to be tracked. The <b>no</b> form of the command removes the values from the configuration.
<b>Default</b>	no max-num-groups
<b>Parameters</b>	<i>max-num-groups</i> — Specifies the maximum number of multicast groups allowed to be tracked. <b>Values</b> 1 — 196607

## max-num-sources

<b>Syntax</b>	<b>max-num-sources</b> <i>max-num-sources</i> <b>no max-num-sources</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if
<b>Description</b>	This command specifies the maximum number of multicast sources allowed to be tracked per group. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	no max-num-sources
<b>Parameters</b>	<i>max-num-sources</i> — Specifies the maximum number of multicast sources allowed to be tracked per group. <b>Values</b> 1 — 1000

## label-mode

<b>Syntax</b>	<b>label-mode</b> { <i>vrf</i>   <i>next-hop</i> } <b>no label-mode</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command controls the method by which service labels are allocated to routes exported by the VPRN as BGP-VPN routes. The <b>vrf</b> option selects service label per VRF mode while the <b>next-hop</b> option selects service label per next-hop mode. The <b>no</b> form of the command sets the mode to the default mode of service label per VRF.
<b>Default</b>	no label-mode
<b>Parameters</b>	<b>vrf</b> — Selects service label per VRF mode. <b>next-hop</b> — Selects service label per next-hop mode.

## maximum-ipv6-routes

<b>Syntax</b>	<b>maximum-ipv6-routes</b> <i>number</i> [ <b>log-only</b> ] [ <b>threshold</b> <i>percentage</i> ] <b>no maximum-ipv6-routes</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command specifies the maximum number of remote IPv6 routes that can be held within a VPN routing/ forwarding (VRF) context. Note that <b>local</b>, <b>host</b>, <b>static</b> and <b>aggregate</b> routes are not counted.</p> <p>Note that the VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.</p> <p>If the <b>log-only</b> parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.</p> <p>The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.</p> <p>The <b>no</b> form of the command disables any limit on the number of routes within a VRF context. Issue the <b>no</b> form of the command only when the VPRN instance is shutdown.</p>
<b>Default</b>	0 or disabled — The threshold will not be raised.
<b>Parameters</b>	<p><i>number</i> — An integer that specifies the maximum number of routes to be held in a VRF context.</p> <p><b>Values</b> 1 — 2147483647</p> <p><b>log-only</b> — This parameter specifies that if the maximum limit is reached, only log the event. <b>log-only</b> does not disable the learning of new routes.</p> <p><b>threshold</b> <i>percentage</i> — The percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ( <math>[\text{mid}+\text{max}] / 2</math> ).</p> <p><b>Values</b> 0 — 100</p>

## maximum-routes

<b>Syntax</b>	<b>maximum-routes</b> <i>number</i> [ <b>log-only</b> ] [ <b>threshold</b> <i>percentage</i> ] <b>no maximum-routes</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command specifies the maximum number of remote routes that can be held within a VPN routing/ forwarding (VRF) context. Note that <b>local</b>, <b>host</b>, <b>static</b> and <b>aggregate</b> routes are not counted.</p> <p>Note that the VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.</p> <p>If the <b>log-only</b> parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the</p>

VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of the command disables any limit on the number of routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

**Default** 0 or disabled — The threshold will not be raised.

**Parameters** *number* — An integer that specifies the maximum number of routes to be held in a VRF context.

**Values** 1 — 2147483647

**log-only** — This parameter specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

**threshold** *percentage* — The percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate (  $(\text{mid} + \text{max}) / 2$  ).

**Values** 0 — 100

## multicast-info-policy

**Syntax** **multicast-info-policy** *policy-name*  
**no multicast-info-policy**

**Context** config>service>vprn

**Description** This command configures multicast information policy.

**Parameters** *policy-name* — Specifies the policy name.

**Values** 32 chars max

## mc-maximum-routes

**Syntax** **mc-maximum-routes** *number* [**log-only**] [**threshold** *threshold*]

**Context** config>service>vprn

**Description** This command specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then no new joins will be processed.

The **no** form of the command disables the limit of multicast routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

**Default** no mc-maximum-routes

## Global Commands

<b>Parameters</b>	<i>number</i> — Specifies the maximum number of routes to be held in a VRF context.
<b>Values</b>	1 — 2147483647
	<b>log-only</b> — Specifies that if the maximum limit is reached, only log the event. <b>log-only</b> does not disable the learning of new routes.
	<b>threshold</b> <i>threshold</i> — The percentage at which a warning log message and SNMP trap should be sent.
<b>Values</b>	0 — 100
<b>Default</b>	10

## network

<b>Syntax</b>	<b>network</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure network parameters for the VPRN service.

## ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>service>vprn>network
<b>Description</b>	This command enables the context to configure network ingress parameters for the VPRN service.

## filter

<b>Syntax</b>	<b>filter ip</b> <i>ip-filter-id</i> <b>filter ipv6</b> <i>ipv6-filter-id</i> <b>no filter</b> [ <b>ip</b> <i>ip-filter-id</i> ] [ <b>ipv6</b> <i>ipv6-filter-id</i> ]
<b>Context</b>	config>service>vprn>network>ingress
<b>Description</b>	This command configures a network ingress filter for IPv4 or IPv6 traffic arriving over explicitly defined spokes or auto-bind network interfaces for the VPRN service. The <b>no</b> form of the command removes an IPv4, IPv6, or both filters.
<b>Default</b>	no filter
<b>Parameters</b>	<i>ip-filter-id/ipv6-filter-id</i> — Specifies an existing IP/IPv6 filter policy of a scope template.
<b>Values</b>	[1..65535]   <i>name</i> <i>name</i> : 64 characters maximum

## ptp

<b>Syntax</b>	<b>[no] ptp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure PTP parameters for the VPRN service.

## peer-limit

<b>Syntax</b>	<b>peer-limit <i>limit</i></b> <b>no peer-limit</b>
<b>Context</b>	configure>service>vprn>ptp
<b>Description</b>	<p>This command specifies an upper limit to the number of discovered peers permitted within the routing instance. This can be used to ensure that a routing instance does not consume all the possible discovered peers and blocking discovered peers in other routing instances.</p> <p>If it is desired to reserve a fixed number of discovered peers per router instance, then all router instances supporting PTP should have values specified with this command and the sum of all the peer-limit values must not exceed the maximum number of discovered peers supported by the system.</p> <p>If the user attempts to specify a peer-limit, and there are already more discovered peers in the routing instance than the new limit being specified, the configuration will not be accepted.</p>
<b>Default</b>	no limit
<b>Parameters</b>	<i>limit</i> — Specifies the maximum number of discovered peers allowed in the routing instance.
	<b>Values</b> 0 — 50
	<b>Default</b> 0 (The maximum number of discovered peers supported by the system.)

## peer

<b>Syntax</b>	<b>peer <i>a.b.c.d</i> [create]</b>
<b>Context</b>	config>system>ptp configure>service>vprn>ptp
<b>Description</b>	<p>This command configures a remote PTP peer. It provides the context to configure parameters for the remote PTP peer.</p> <p>Up to 20 remote PTP peers may be configured.</p> <p>The <b>no</b> form of the command deletes the specified peer.</p> <p>If the clock-type is ordinary slave or boundary, and PTP is no shutdown, the last peer cannot be deleted. This prevents the user from having PTP enabled without any peer configured and enabled.</p> <p>Peers are created within the routing instance associated with the context of this command. All configured PTP peers must use the same routing instance.</p>

## Global Commands

<b>Default</b>	none
<b>Parameters</b>	<i>a.b.c.d</i> — The IP address of the remote peer.
<b>Values</b>	ipv4-address <i>a.b.c.d</i>

## log-sync-interval

<b>Syntax</b>	<b>log-sync-interval</b> <i>log-interval</i>
<b>Context</b>	configure>service>vpn>ptp>peer
<b>Description</b>	<p>This command configures the message interval used for unicast event messages. It defines the message interval for both Sync and Delay_Resp messages that are requested during unicast negotiation to the specific peer. This controls the Sync and Delay_Resp message rate sent from remote peers to the local node. It does not affect the Sync or Delay_Resp packet rate that may be sent from the local node to remote peers. Remote peers may request a Sync or Delay_Resp packet rate anywhere within the acceptable grant range.</p> <p>The <b>log-sync-interval</b> cannot be changed unless the peer is shutdown.</p>
<b>Default</b>	-6 (64 packets per second) for 1588-2008 or -6 (64 packets per second) for g8265dot1-2010 or -4 (16 packets per second) for g8275dot1-2014
<b>Parameters</b>	<i>log-interval</i> — Specifies the sync message interval, in log form.
<b>Values</b>	[-6..0]

## local-priority

<b>Syntax</b>	<b>local-priority</b> <i>local-priority</i>
<b>Context</b>	configure>service>vpn>ptp>peer
<b>Description</b>	<p>This command configures the local priority used to choose between PTP masters in the best master clock algorithm (BMCA). This setting is relevant when the profile is set to either g8265dot1-2010 or g8275dot1-2014. The parameter is ignored when any other profile is selected.</p> <p>The value 1 is the highest priority and 255 is the lowest priority. The priority of a peer cannot be configured if the PTP profile is ieee1588-2008.</p> <p>For g8265dot1-2010, this parameter configures the priority used to choose between master clocks with the same quality (see G.8265.1 for more details).</p> <p>For g8275dot1-2014, this parameter sets the value of the <b>localPriority</b> associated with the Announce messages received from external clocks (<b>ptp&gt;peer</b> or <b>ptp&gt;port</b>), or the local clock (<b>ptp</b>). See G.8275.1 for more detailed information.</p>
<b>Default</b>	128
<b>Parameters</b>	<i>local-priority</i> — Specifies the value of the local priority.



**Values** 1 — 255

## reassemble-group

<b>Syntax</b>	<b>reassemble-group</b> <i>nat-group-id</i> <b>no reassemble-group</b>
<b>Context</b>	configure>router config>service>vpn
<b>Description</b>	This command associate reassemble-group consisting of multiple ISAs with the routing context in which the application requiring reassemble service resides.
<b>Default</b>	no route-distinguisher
<b>Parameters</b>	<i>nat-group-id</i> — Nat-group id. The nat-group contains up to 10 active ISAs. <i>asn:number</i> — The ASN is a 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value.

## route-distinguisher

<b>Syntax</b>	<b>route-distinguisher</b> [ <i>ip-address:number</i>   <i>asn:number</i> ] <b>route-distinguisher auto-rd</b> <b>no route-distinguisher</b>
<b>Context</b>	config>service>vpn
<b>Description</b>	This command sets the identifier attached to routes the VPN belongs to. Each routing instance must have a unique (within the carrier's domain) route distinguisher associated with it. A route distinguisher must be defined for a VPRN to be operationally active.  Alternatively, the <b>auto-rd</b> option allows the system to automatically generate a Route Distinguisher (RD) based on the <b>bgp-auto-rd-range</b> command configured at the service level.
<b>Default</b>	no route-distinguisher
<b>Parameters</b>	The route distinguisher is a 6-byte value that can be specified in one of the following formats: <i>ip-address:number</i> — Specifies the IP address in dotted decimal notation. The assigned number must not be greater than 65535. <i>asn:number</i> — The ASN is a 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value. <b>auto-rd</b> — The system will generate an RD for the service according to the IP address and range configured in the <b>bgp-auto-rd-range</b> command.

## router-id

<b>Syntax</b>	<b>router-id</b> <i>ip-address</i> <b>no router-id</b>
<b>Context</b>	config>service>vprn config>service>vprn>ospf config>service>vprn>bgp
<b>Description</b>	<p>This command sets the router ID for a specific VPRN context.</p> <p>When configuring the router ID in the base instance of OSPF it overrides the router ID configured in the <b>config&gt;router</b> context. The default value for the base instance is inherited from the configuration in the <b>config&gt;router</b> context. If the router ID in the <b>config&gt;router</b> context is not configured, the following applies:</p> <ul style="list-style-type: none"> <li>• The system uses the system interface address (which is also the loopback address).</li> <li>• If a system interface address is not configured, use the last 32 bits of the chassis MAC address.</li> </ul> <p>If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.</p> <p>This is a <b>required</b> command when configuring multiple instances and the instance being configured is not the base instance.</p> <p>When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.</p> <p>To force the new router ID to be used, issue the <b>shutdown</b> and <b>no shutdown</b> commands for the instance, or reboot the entire router.</p> <p>It is possible to configure an SR OS node to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.</p> <p>The <b>no</b> form of the command removes the router ID definition from the given VPRN context.</p>
<b>Default</b>	no router-id
<b>Parameters</b>	<i>ip-address</i> — The IP address must be given in dotted decimal notation.

## service-name

<b>Syntax</b>	<b>service-name</b> <i>service-name</i> <b>no service-name</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command configures an optional service name, up to 64 characters in length, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the and 7750 SR platforms.</p>

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

**Parameters** *service-name* — Specifies a unique service name to identify the service. Service names may not begin with an integer (0-9).

## sgt-qos

<b>Syntax</b>	<b>sgt-qos</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure DSCP/Dot1p re-marking for self-generated traffic.

## application

<b>Syntax</b>	<b>application</b> <i>dscp-app-name</i> <b>dscp</b> { <i>dscp-value</i>   <i>dscp-name</i> } <b>application</b> <i>dot1p-app-name</i> <b>dot1p</b> <i>dot1p-priority</i> <b>no application</b> { <i>dscp-app-name</i>   <i>dot1p-app-name</i> }
<b>Context</b>	config>service>vprn>sgt-qos
<b>Description</b>	This command configures DSCP/Dot1p re-marking for self-generated traffic. When an application is configured using this command, then the specified DSCP name/value is used for all packets generated by this application within the router instance it is configured.

Using the value configured in this command:

- Sets the DSCP bits in the IP packet.
- Maps to the FC. This value will be signaled from the CPM to the egress forwarding complex.
- Based on this signaled FC the egress forwarding complex QoS policy sets the IEEE802.1 dot1P and LSP EXP bits.
- The Dot1P and the LSP EXP bits are set by the egress complex for all packets based on the signaled FC. This includes ARP and IS-IS packets that, due to their nature, do not carry DSCP bits.
- The DSCP value in the egress IP header will be as configured in this command. The egress QoS policy will not overwrite this value.

Only one DSCP name/value can be configured per application, if multiple entries are configured then the subsequent entry overrides the previous configured entry.

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

<b>Parameters</b>	<i>dscp-app-name</i> — Specifies the DSCP application name.
	<b>Values</b> ldp, rsvp, bgp, rip, msdp, pim, ptp, ospf, igmp, mld, telnet, tftp, ftp, ssh, snmp, snmp-notification, syslog, icmp, traceroute, tacplus, dns, ntp, radius, cflowd, dhcp, bootp, ndis, vrrp, srrp
	<i>dscp-value</i> — Specifies a value when this packet egresses the respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or IEEE 802.1p (Dot1P) bits as appropriate otherwise the default mapping applies.
	<b>Values</b> 0 — 63
	<i>dscp-name</i> — Specifies the DSCP name.
	<b>Values</b> none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39,

cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

*dot1p-priority* — Specifies the Dot1P priority.

**Values** 0 — 7

*dot1p-app-name* — Specifies the Dot1P application name.

**Values** arp, isis

## dscp

<b>Syntax</b>	<b>dscp</b> <i>dscp-name</i> <b>fc</b> <i>fc-name</i> <b>no dscp</b> <i>dscp-name</i>
<b>Context</b>	config>service>vprn>sgt-qos
<b>Description</b>	<p>This command creates a mapping between the DiffServ Code Point (DSCP) of the self generated traffic and the forwarding class.</p> <p>Self generated traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all sixty-four DiffServ code points to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the default-action command.</p> <p>All DSCP names that defines a DSCP value must be explicitly defined.</p> <p>The <b>no</b> form of this command removes the DiffServ code point to forwarding class association. The default-action then applies to that code point value.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>dscp-name</i> — The name of the DiffServ code point to be associated with the forwarding class. DiffServ code point can only be specified by its name and only an existing DiffServ code point can be specified. The software provides names for the well known code points.</p> <p><b>Values</b> be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63</p> <p><b>fc</b> <i>fc-name</i> — Specifies the forwarding class name. All packets with DSCP value or MPLS EXP bits that is not defined will be placed in this forwarding class.</p> <p><b>Default</b> None, the fc name must be specified</p> <p><b>Values</b> be, l2, af, l1, h2, ef, h1, nc</p>

## single-sfm-overload

**Syntax** `single-sfm-overload [holdoff-time holdoff-time]`  
`no single-sfm-overload`

**Context** `config>service>vprn`

**Description** This command configures OSPF, OSPFv3 and IS-IS to set overload when the router has fewer than the full set of SFMs functioning, which reduces forwarding capacity. Setting overload enables a router to still participate in exchanging routing information, but routes all traffic away from it.

The conditions to set overload are as follows:

- 7750 SR-12/SR-7/SR-c12 and 7450 ESS-12/ESS-7/ESS-6 platforms: protocol sets overload if one of the SF/CPMs fails
- 7950 XRS and 7750 SR-12e platforms: protocol sets overload if two SFMs fail

The **no** form of this command configures the router to not set overload if an SFM fails.

**Default** `no single-sfm-overload`

**Parameters** *holdoff-time* — This parameter specifies the delay between detecting SFM failures and setting overload.

**Values** 1 — 600 seconds

**Default** 0 seconds

## snmp

**Syntax** `snmp`

**Context** `config>service>vprn`

**Description** This command enables the context to configure SNMP parameters for this VPRN.

## access

**Syntax** `[no] access`

**Context** `config>service>vprn>snmp`

**Description** This command enables/disables SNMP access on the VPRN interface. This command allows SNMP queries destined to the VPRN interface IP addresses for this VPRN (including VPRN interfaces that are bound to R-VPLS services) to be processed by the SNMP agent on the router. SNMP queries that arrive on VPRN interfaces but are destined to IP addresses in the Base routing context that can be accessed in the VPRN (for example, the router system address via grt leaking do not require snmp-access to be enabled but do require allow-local-management to be enabled.

Refer to the System Management Guide for detailed information about SNMP.

## community

**Syntax** **community** *community-name* [**hash** | **hash2**] [**access-permissions**] [**version** *SNMP-version*] [**src-access-list** *list-name*]  
**no community** [*community-name*]

**Context** config>service>vprn>snmp

**Description** This command sets the SNMP community name(s) to be used with the associated VPRN instance. These VPRN community names are used to associate SNMP v1/v2c requests with a particular vprn context and to return a reply that contains VPRN-specific data or limit SNMP access to data in a specific VPRN instance.

VPRN snmp communities configured with an access permission of 'r' are automatically associated with the default access group "snmp-vprn-ro" and the "vprn-view" view (read only). VPRN snmp communities configured with an access permission of 'rw' are automatically associated with the default access group "snmp-vprn" and the "vprn-view" view (read/write).

The community in an SNMP v1/v2 request determines the SNMP context (i.e., the vprn# for accessing SNMP tables) and not the VPRN of the incoming interface on which the request was received. When an SNMP request arrives on VPRN 5 interface "ringo" with a destination IP address equal to the "ringo" interface, but the community in the SNMP request is the community configured against VPRN 101, then the SNMP request will be processed using the VPRN 101 context. (the response will contain information about VPRN 101). It is recommended to avoid using a simple series of vprn snmp-community values that are similar to each other (for example, avoid my-vprncomm-1, my-vprn-comm-2, etc).

The **no** form of the command removes the SNMP community name from the given VPRN context.

**Default** None — The SNMP community must be explicitly specified.

**Parameters** *community-name* — Specifies the SNMP v1/v2c community name. This is a secret/confidential key used to access SNMP and specify a context (base vs vprn1 vs vprn2).  
**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.  
**hash2** — Specifies the key is entered in a more complex encrypted form.  
**version** *SNMP-version* — Specifies the SNMP version.

**Values** v1, v2c, both

*access-permissions* — Specifies the access rights to MIB objects.

**Values** **r** — Grants only read access to MIB objects. Creates an association of the community-name with the **snmp-vprn-ro** access group.  
**rw** — Grants read and write access to MIB objects. Creates an association of the community-name with the **snmp-vprn** access group.

*list-name* — Configures the **community** to reference a specific **src-access-list** (created under **configure system security snmp**), which will be used to validate the source IP address of all received SNMP requests that use this **community**. Multiple **community** (vprn or base router) and **usm-community** instances can reference the same **src-access-list**.

## source-address

<b>Syntax</b>	<b>source-address</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to specify the source address and application that should be used in all unsolicited packets.

## application

<b>Syntax</b>	<b>application</b> <i>app</i> [ <i>ip-int-name</i>   <i>ip-address</i> ] <b>no application</b> <i>app</i>
<b>Context</b>	config>service>vprn>source-address
<b>Description</b>	This command specifies the source address and application.
<b>Parameters</b>	<i>app</i> — Specify the application name.  <b>Values</b> cflowd, dns, ftp, ntp, ping, ptp, radius, snmptrap, sntp, ssh, syslog, tacplus, telnet, traceroute, mcreporter, icmp-error  <i>ip-int-name</i>   <i>ip-address</i> — Specifies the name of the IP interface or IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## application6

<b>Syntax</b>	<b>application6</b> <i>app</i> <i>ipv6-address</i>
<b>Context</b>	config>service>vprn>source-address
<b>Description</b>	This command specifies the IPv6 source address and application.
<b>Parameters</b>	<i>app</i> — Specify the application name.  <b>Values</b> cflowd, dns, ftp, ntp, ping, radius, snmptrap, syslog, tacplus, telnet, traceroute, icmp6-error  <i>ipv6-address</i> — Specifies the IPv6 address.

## static-route

<b>Syntax</b>	<b>[no] static-route</b> { <i>ip-prefix/prefix-length</i>   <i>ip-prefix netmask</i> } [ <b>preference</b> <i>preference</i> ] <b>[metric</b> <i>metric</i> ] [ <b>tag</b> <i>tag</i> ] [ <b>community</b> <i>comm-id</i> ] [ <b>enable</b>   <b>disable</b> ] { <b>next-hop</b> <i>ip-int-name</i>   <i>ip-address</i> [ <b>mcast-family</b> ]   <b>ipsec-tunnel</b> <i>ipsec-tunnel-name</i> } [ <b>bfd-enable</b>   { <b>cpe-check</b> <i>cpe-ip-address</i> [ <b>interval</b> <i>seconds</i> ] [ <b>drop-count</b> <i>count</i> ] [ <b>padding-size</b> <i>padding-size</i> ] [ <b>log</b> ]}} <b>[no] static-route</b> { <i>ip-prefix/prefix-length</i>   <i>ip-prefix netmask</i> } [ <b>preference</b> <i>preference</i> ] <b>[metric</b> <i>metric</i> ] [ <b>tag</b> <i>tag</i> ] [ <b>community</b> <i>comm-id</i> ] [ <b>enable</b>   <b>disable</b> ] <b>indirect</b> <i>ip-address</i> [ <b>cpe-</b>
---------------	--



**check** *cpe-ip-address* [**interval** *seconds*][**drop-count** *count*] [**padding-size** *padding-size*] [**log**]  
 [**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*]  
 [**metric** *metric*] [**tag** *tag*] [**community** *comm-id*] [**enable** | **disable**] **black-hole** [**mcast-family**]

**Context** config>service>vprn

**Description** This command creates a static route. A static route can have a directly-connected interface as a next-hop (specified using the IP interface name or an IP address of the interface), or an indirect IP address as a next-hop or a black-hole next-hop (specifying a discard action).

The no form of the command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered.

If a CPE connectivity check target address is already being used as the target address in a different static route, then cpe-check parameters must match. If they do not, the new configuration command will be rejected.

If a static-route command is issued with no cpe-check target but the destination prefix/netmask and next-hop matches a static route that did have an associated cpe-check, the cpe-check test will be removed from the associated static route.

**Default** No static routes are defined.

**Parameters** *ip-prefix* — The destination address of the aggregate route in dotted decimal notation.

<b>Values</b>	ipv4-prefix	a.b.c.d (host bits must be 0)
	ipv4-prefix-length	0 — 32
	ipv6-prefix	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
	ipv6-prefix-length	0 — 128

*netmask* — The subnet mask in dotted decimal notation.

**Values** 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

*ip-int-name* — The name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed with

*ip-address* — The IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

<b>Values</b>	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface]
		x: [0..FFFF]H
		d: [0..255]D
		interface: 32 characters maximum, mandatory for link local addresses

**community** *comm-id* — This configuration option associates a BGP community with the static route. The community can be matched in route policies and is automatically added to BGP routes exported from the static route.

<b>Values</b>	comm-id	asn:comm-val   well-known-comm
	asn	0 — 65535
	comm-val	0 — 65535
	well-known-comm	no-advertise, no-export, no-export-subconfed

**enable** — Static routes can be administratively enabled or disabled. Use the **enable** parameter to re-enable a disabled static route. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

**Default** enable

**disable** — Static routes can be administratively enabled or disabled. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

**Default** enable

**interval** *seconds* — This optional parameter specifies the interval between ICMP pings to the target IP address.

**Values** 1 —255 seconds

**Default** 1 seconds

**drop-count** *count* — This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to de-active the associated static route.

**Values** Value range: 1 —255

**Default** 3

**padding-size** *padding-size* — This optional parameter specifies the amount of padding to add to the ICMP packet in bytes. The parameter is only applicable when the **cpe-check** parameter is used with the **static-route** command.

**Values** 0 — 16384 bytes

**log** — This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events should be sent to the system log, syslog and SNMP traps.

**next-hop** [*ip-address* | *ip-int-name*] — Specifies the directly connected next hop IP address used to reach the destination. If the next hop is over an unnumbered interface, the *ip-int-name* of the unnumbered interface (on this node) can be configured.

The **next-hop** keyword and the **indirect** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **indirect** or **black-hole** parameters), then this static route will be replaced with the newly entered command, and unless specified, the respective defaults for preference and metric will be applied.

The *ip-addr* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

**ipsec-tunnel** *ipsec-tunnel-name* — specifies an IPsec tunnel name up to 32 characters in length.

**indirect** *ip-address* — Specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr* is not directly connected to a network configured on this node. The destination can be reachable via multiple paths. The static route remains valid as long as the address configured as the indirect address remains a valid entry in the routing table. Indirect static routes cannot use an *ip-prefix/mask* to another indirect static route.

The **indirect** keyword and the **next-hop** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **next-hop** or **black-hole** parameters), then this static route will be replaced with the newly entered command and unless specified the respective defaults for preference and metric will be applied.

The *ip-addr* configured can be either on the network or the access side and is normally at least one hop away from this node.

**black-hole** — Specifies a black hole route meaning that if the destination address on a packet matches this static route it will be silently discarded.

The **black-hole** keyword is mutually exclusive with either the **next-hop** or **indirect** keywords. If an identical command is entered, with exception of either the **next-hop** or **indirect** parameters, then the static route is replaced with the new command, and unless specified, the respective defaults for **preference** and **metric** are applied.

**preference** *preference* — The preference of this static route (as opposed to the routes from different sources such as BGP or OSPF), expressed as a decimal integer. When modifying the **preference** value of an existing static route, unless specified, the metric will not change.

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of which route to use is determined by the configuration of the ECMP command.

**Default** 5

**Values** 1 — 255

**metric** *metric* — The cost metric for the static route, expressed as a decimal integer. This value is used when importing this static route into other protocols such as OSPF. This value is also used to determine the static route to install in the forwarding table: When modifying the metrics of an existing static route, unless specified, the preference will not change.

If there are multiple static routes with the same preference but unequal metrics, the lower cost (metric) route is installed. If there are multiple static routes with equal preference and metrics then ECMP rules apply. If there are multiple routes with unequal preferences, then the lower preference route is installed.

**Default** 1

**Values** 0 — 65535

**tag** — Adds a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

**Values** 1— 4294967295

**bfd-enable** — Associates the state of the static route to a BFD session between the local system and the configured nexthop. This keyword cannot be configured if the nexthop is **indirect** or a **blackhole** keywords are specified.

**cpe-check target-ip-address** — This parameter specifies the IP address of the target CPE device. ICMP pings will be sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the associated static route. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.

**Default** no cpe-check enabled

**mcast-family** — Enables submission of the IPv4 static route into IPv4 multicast RTM.

**Values** mcast-ipv4

## ttl-propagate

**Syntax** **ttl-propagate**

**Context** config>service>vprn

**Description** This command enables the context to configure TTL propagation for transit and locally generated packets in a given VPRN routing context.

**Default** none

## local

**Syntax** **local [inherit | all | vc-only | none]**

**Context** config>service>vprn>ttl-propagate

**Description** This command overrides the global configuration of the TTL propagation for locally generated packets which are forwarded over a MPLS LSPs in a given VPRN service context.

The global configuration is performed under config>router>ttl-propagate>vprn-local.

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value

**Default** inherit

**Parameters** **inherit** — TheTTL propoagation behavior is inherited from the global configuration under config>router>ttl-propogate>vprn-local

**none** — TheTTL of the IP packet is not propagated into the VC label or labels in the transport label stack

**vc-only** — The TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.

**all** — The TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.

## transit

<b>Syntax</b>	<b>transit [inherit   all   vc-only   none]</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command overrides the global configuration of the TTL propagation for in transit packets which are forwarded over a MPLS LSPs in a given VPRN service context.</p> <p>The global configuration is performed under config&gt;router&gt;ttl-propagate&gt;vprn-transit.</p> <p>The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value.</p>
<b>Default</b>	inherit
<b>Parameters</b>	<p><b>inherit</b> — The TTL propagation behavior is inherited from the global configuration under config&gt;router&gt;ttl-propagate&gt;vprn-transit</p> <p><b>none</b> — The TTL of the IP packet is not propagated into the VC label or labels in the transport label stack</p> <p><b>vc-only</b> — The TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.</p> <p><b>all</b> — The TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.</p>

## type

<b>Syntax</b>	<b>type [hub   spoke   subscriber-split-horizon]</b> <b>no type</b>
<b>Context</b>	config>service>vprn>
<b>Description</b>	This command designates the type of VPRN instance being configured for hub and spoke topologies. Use the <b>no</b> form to reset to the default of a fully meshed VPRN.
<b>Default</b>	no type
<b>Parameters</b>	<p><b>hub</b> — Specifies a hub VPRN which allows all traffic from the hub SAPs to be routed to the destination directly, while all traffic from spoke VPRNs or network interfaces can only be routed to a hub SAP.</p> <p><b>spoke</b> — Specifies a spoke VPRN which allows traffic from associated SAPs or spoke terminations to only be forwarded through routes learned from separate VPRN, which should be configured as a type Hub VPRN.</p>

**subscriber-split-horizon** — Controls the flow of traffic for wholesale subscriber applications.

## vrf-export

<b>Syntax</b>	<b>vrf-export</b> <i>policy</i> [ <i>policy...</i> ] <b>no vrf-export</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command specifies the export policies to control routes exported from the local VPN routing/forwarding (VRF) to other VRFs on the same or remote PE routers (via MP-BGP). You can specify up to fifteen (15) policy names. The <b>no</b> form of the command removes all route policy names from the export list.
<b>Default</b>	None — No routes are exported from the VRF by default.
<b>Parameters</b>	<i>policy</i> — The route policy statement name.

## vrf-import

<b>Syntax</b>	<b>vrf-import</b> <i>policy</i> [ <i>policy...</i> ] <b>no vrf-import</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command sets the import policies to control routes imported to the local VPN routing/forwarding (VRF) from other VRFs on the same or remote PE routers (via MP-BGP). Up to fifteen (15) names may be specified. BGP-VPN routes imported with a vrf-import policy will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs on the same router, unless the preference is changed by the policy. The <b>no</b> form of the command removes all route policy names from the import list.
<b>Default</b>	None — No routes are accepted into the VRF by default.
<b>Parameters</b>	<i>policy</i> — The route policy statement name.

## vrf-target

<b>Syntax</b>	<b>vrf-target</b> { <b>ext-community</b>   <b>export</b> <i>ext-community</i>   <b>import</b> <i>ext-community</i> } <b>no vrf-target</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command facilitates a simplified method to configure the route target to be added to advertised routes or compared against received routes from other VRFs on the same or remote PE routers (via MP-BGP).

BGP-VPN routes imported with a `vrf-target` statement will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs in the same router.

Specified **vrf-import** or **vrf-export** policies override the **vrf-target** policy.

The no form of the command removes the `vrf-target`

**Default** no vrf-target

**Parameters** *ext-comm* — An extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

**Values** <ext-community> : target: {<ip-addr:comm-val>|<2byte-asnumber:ext-comm-val>|<4byte-asnumber:comm-val>}  
 ip-addr a.b.c.d  
 comm-val [0..65535]  
 2byte-asnumber [0..65535]  
 ext-comm-val [0..4294967295]  
 4byte-asnumber [0..4294967295]

**import** *ext-community* — Specify communities allowed to be accepted from remote PE neighbors.

**export** *ext-community* — Specify communities allowed to be sent to remote PE neighbors.

---

## Router L2TP Commands

### l2tp

<b>Syntax</b>	<b>l2tp</b> <b>no l2tp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure L2TP parameters. L2TP extends the PPP model by allowing Layer 2 and PPP endpoints to reside on different devices interconnected by a packet-switched network.

### avp-hiding

<b>Syntax</b>	<b>avp-hiding sensitive   always</b> <b>no avp-hiding</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as cleartext in an AVP. The <b>no</b> form of the command returns the value to <b>never</b> allow AVP hiding.
<b>Parameters</b>	<i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnels in this L2TP group. <b>Default</b> no avp-hiding <b>Values</b> sensitive — AVP hiding is used only for sensitive information (such as username/password). always — AVP hiding is always used.

### calling-number-format

<b>Syntax</b>	<b>calling-number-format ascii-spec</b> <b>no calling-number-format</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command what string to put in the Calling Number AVP, for L2TP control messages related to a session in this L2TP protocol instance.
<b>Parameters</b>	<i>ascii-spec</i> — Specifies the L2TP calling number AVP. <b>Values</b> ascii-spec char-specification ascii-spec char-specification ascii-char   char-origin



ascii-char	a printable ASCII character
char-origin	%origin
origin	S   c   r   s   l
S	- system name, the value of TIMETRA-CHASSIS-MIB::tmnxChassisName
c	- Agent Circuit Id
r	- Agent Remote Id
s	- SAP ID, formatted as a character string
l	- Logical Line ID

## challenge

<b>Syntax</b>	<b>challenge</b> <i>always</i> <b>no challenge</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures the use of challenge-response authentication. The <b>no</b> form of the command reverts to the default <b>never</b> value.
<b>Parameters</b>	<i>always</i> — Specifies that the challenge-response authentication is always used.
<b>Default</b>	no challenge
<b>Values</b>	always

## destruct-timeout

<b>Syntax</b>	<b>destruct-timeout</b> <i>destruct-timeout</i> <b>no destruct-timeout</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures the period of time that the data of a disconnected tunnel will persist before being removed. The <b>no</b> form of the command removes the value from the configuration.
<b>Default</b>	no destruct-timeout
<b>Parameters</b>	<i>destruct-timeout</i> — [Specifies the automatic removal of dynamic L2TP sessions, in seconds, that are no longer active.
<b>Default</b>	no destruct-timeout
<b>Values</b>	60 — 86400

## exclude-avps

<b>Syntax</b>	<b>exclude-avps calling-number</b> <b>no exclude-avps</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures the L2TP AVPs to exclude.

## ipcp-subnet-negotiation

<b>Syntax</b>	<b>[no] ipcp-subnet-negotiation</b>
<b>Context</b>	configure>router>l2tp>group>ppp configure>router>l2tp>group>tunnel>ppp configure>service>vprn>l2tp>group>ppp configure>service>vprn>l2tp>group>tunnel>ppp
<b>Description</b>	Enables IPCP negotiation for PPPoE hosts. If not enabled (default setting), the current behavior will apply even if subnet is allocated to the host. Enables IPCP negotiation for PPPoE hosts. If not enabled (default setting), the current behavior will apply even if subnet is allocated in the host.

## peer-address-change-policy

<b>Syntax</b>	<b>peer-address-change-policy {accept   ignore   reject}</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures the reaction to a change of tunnel peer address in this router.

## receive-window-size

<b>Syntax</b>	<b>receive-window-size [4..1024]</b> <b>no receive-window-size</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures the L2TP receive window size.

## session-limit

<b>Syntax</b>	<b>session-limit session-limit</b> <b>no session-limit</b>
<b>Context</b>	config>service>vprn>l2tp
<b>Description</b>	This command configures the L2TP session limit of this router.

**Parameters** *session-limit* — Specifies the session limit.

**Values** 1..131071

## group

**Syntax** **group** *tunnel-group-name* [**create**]  
**no group** *tunnel-group-name*

**Context** config>service>vprn>l2tp

**Description** This command configures an L2TP tunnel group.

**Parameters** *tunnel-group-name* — Specifies a name string to identify a L2TP group up to 63 characters in length.

**create** — This keyword is mandatory when creating a tunnel group name. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

## session-limit

**Syntax** **session-limit** *session-limit*  
**no session-limit**

**Context** config>service>vprn>l2tp

**Description** This command configures the L2TP session limit for the router. L2TP is connection-oriented. The L2TP Network Server (LNS) and LAC maintain state for each call that is initiated or answered by an LAC. An L2TP session is created between the LAC and LNS when an end-to-end PPP connection is established between a remote system and the LNS. Datagrams related to the PPP connection are sent over the tunnel between the LAC and LNS. There is a one to one relationship between established L2TP sessions and their associated calls.

**Parameters** *session-limit* — Specifies the number of sessions allowed.

**Default** no session-limit

**Values** 1 — 131071

## avp-hiding

**Syntax** **avp-hiding** *sensitive* | *always*  
**no avp-hiding**

**Context** config>service>vprn>l2tp>group

**Description** This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as cleartext in an AVP.

The **no** form of the command returns the value to **never** allow AVP hiding.

## Router L2TP Commands

<b>Parameters</b>	<i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnels in this L2TP group.
<b>Default</b>	no avp-hiding
<b>Values</b>	sensitive — AVP hiding is used only for sensitive information (such as username/password). always — AVP hiding is always used.

## challenge

<b>Syntax</b>	<b>challenge</b> <i>always</i> <b>no challenge</b>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command configures the use of challenge-response authentication. The <b>no</b> form of the command reverts to the default <b>never</b> value.
<b>Parameters</b>	<i>always</i> — Specifies when challenge-response is to be used for the authentication of the tunnels in this L2TP group.
<b>Default</b>	no challenge
<b>Values</b>	always

## destruct-timeout

<b>Syntax</b>	<b>destruct-timeout</b> <i>destruct-timeout</i> <b>no destruct-timeout</b>
<b>Context</b>	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command configures the period of time that the data of a disconnected tunnel will persist before being removed. The <b>no</b> form of the command removes the value from the configuration.
<b>Default</b>	no destruct-timeout
<b>Parameters</b>	<i>destruct-timeout</i> — [Specifies the automatic removal of dynamic L2TP sessions, in seconds, that are no longer active.
<b>Default</b>	no destruct-timeout
<b>Values</b>	60 — 86400

## hello-interval

<b>Syntax</b>	<b>hello-interval</b> <i>hello-interval</i> <b>no hello-interval</b>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command configures the time interval between two consecutive tunnel Hello messages. The Hello message is an L2TP control message sent by either peer of a LAC-LNS control connection. This control message is used as a keepalive for the tunnel.  The <b>no</b> form of the command removes the interval from the configuration.
<b>Default</b>	60
<b>Parameters</b>	<i>hello-interval</i> — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages.  <b>Default</b> no hello-interval <b>Values</b> 60 — 3600

## idle-timeout

<b>Syntax</b>	<b>idle-timeout</b> <i>idle-timeout</i> <b>no idle-timeout</b>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command configures the period of time that an established tunnel with no active sessions will persist before being disconnected.  Enter the <b>no</b> form of the command to maintain a persistent tunnel.  The <b>no</b> form of the command removes the idle timeout from the configuration.
<b>Default</b>	no idle-timeout
<b>Parameters</b>	<i>idle-timeout</i> — Specifies the idle timeout value, in seconds until the group is removed.  <b>Default</b> no idle-timeout <b>Values</b> 0 — 3600

## Ins-group

<b>Syntax</b>	<b>Ins-group</b> <i>Ins-group-id</i> <b>no Ins-group</b>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command configures the ISA LNS group.

## Router L2TP Commands

**Parameters** *lns-group-id* — Specifies the LNS group ID.

**Values** 1..4

### local-address

**Syntax** **local-address** *ip-address*  
**no local-address**

**Context** config>service>vprn>l2tp>group>tunnel

**Description** This command configures the local address.

**Parameters** *ip-address* — Specifies the IP address used during L2TP authentication.

### local-name

**Syntax** **local-name** *host-name*  
**no local-name**

**Context** config>service>vprn>l2tp>group  
config>service>vprn>l2tp>group>tunnel

**Description** This command creates the local host name used by this system for the tunnels in this L2TP group during the authentication phase of tunnel establishment. It can be used to distinguish tunnels.

The **no** form of the command removes the name from the configuration.

**Default** local-name

**Parameters** *host-name* — Specifies the host name, up to 64 characters in length, that the router will use to identify itself during L2TP authentication.

**Default** no local-name

### max-retries-estab

**Syntax** **max-retries-estab** *max-retries*  
**no max-retries-estab**

**Context** config>service>vprn>l2tp>group  
config>service>vprn>l2tp>group>tunnel

**Description** This command configures the number of retries allowed for this L2TP tunnel while it is established, before its control connection goes down.

The **no** form of the command removes the value from the configuration.

**Default** no max-retries-estab

**Parameters** *max-retries* — Specifies the maximum number of retries for an established tunnel.

**Default** no max-retries-estab  
**Values** 2 — 7

## max-retries-not-estab

**Syntax** **max-retries-not-estab** *max-retries*  
**no max-retries-not-estab**

**Context** config>service>vprn>l2tp>group  
config>service>vprn>l2tp>group>tunnel

**Description** This command configures the number of retries allowed for this L2TP tunnel while it is not established, before its control connection goes down.  
The **no** form of the command removes the value from the configuration.

**Default** no max-retries-not-estab

**Parameters** *max-retries* — Specifies the maximum number of retries for non-established tunnels.

**Default** no max-retries-not-estab  
**Values** 2 — 7

## password

**Syntax** **password** *password* [**hash** | **hash2**]  
**no password**

**Context** config>service>vprn>l2tp>group  
config>service>vprn>l2tp>group>tunnel

**Description** This command configures the password between L2TP LAC and LNS  
The no form of the command removes the password.

**Default** no password

**Parameters** *password* — Configures the password used for challenge/response calculation and AVP hiding. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the hash2 keyword is specified.

**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

**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.

**Default** no password

### ppp

<b>Syntax</b>	<b>ppp</b>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command configures PPP for the L2TP tunnel group.

### authentication

<b>Syntax</b>	<b>authentication {chap pap pref-chap}</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the PPP authentication protocol to negotiate.

### authentication-policy

<b>Syntax</b>	<b>authentication-policy <i>auth-policy-name</i></b> <b>no authentication-policy</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the authentication policy.
<b>Parameters</b>	<i>auth-policy-name</i> — Specifies the authentication policy name. <b>Values</b> 32 chars max

### default-group-interface

<b>Syntax</b>	<b>default-group-interface <i>ip-int-name</i> <i>service-id</i> <i>service-id</i></b> <b>no default-group-interface</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the default group interface.
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the interface name. <b>Values</b> 32 chars max <i>service-id</i> — Specifies the service ID. <b>Values</b> 1..2147483648 <i>svc-name</i> — Specifies the service name (instead of service ID). <b>Values</b> 64 chars max



## keepalive

<b>Syntax</b>	<b>keepalive</b> <i>seconds</i> [ <b>hold-up-multiplier</b> <i>multiplier</i> ] <b>no keepalive</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the PPP keepalive interval and multiplier.
<b>Parameters</b>	<i>seconds</i> — Specifies in seconds the interval. <b>Values</b> 10..300 <i>multiplier</i> — Specifies the multiplier. <b>Values</b> 1..5

## mtu

<b>Syntax</b>	<b>mtu</b> <i>mtu-bytes</i> <b>no mtu</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the maximum PPP MTU size.
<b>Parameters</b>	<i>mtu-bytes</i> — Specifies, in bytes, the maximum PPP MTU size. <b>Values</b> 512..9212

## proxy-authentication

<b>Syntax</b>	<b>[no] proxy-authentication</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the use of the authentication AVPs received from the LAC.

## proxy-lcp

<b>Syntax</b>	<b>[no] proxy-lcp</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the use of the proxy LCP AVPs received from the LAC.

## user-db

<b>Syntax</b>	<b>user-db</b> <i>local-user-db-name</i> <b>no user-db</b>
<b>Context</b>	config>service>vprn>l2tp>group>ppp
<b>Description</b>	This command configures the local user database to use for PPP PAP/CHAP authentication.
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the local user database name. <b>Values</b> 32 chars max

## session-assign-method

<b>Syntax</b>	<b>session-assign-method</b> <i>weighted</i> <b>no session-assign-method</b>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command specifies how new sessions are assigned to one of the set of suitable tunnels that are available or could be made available.
<b>Default</b>	no session-assign-method
<b>Parameters</b>	<i>weighted</i> — specifies that the sessions are shared between the available tunnels. If necessary, new tunnels are set up until the maximum number is reached. The distribution aims at an equal ratio of the actual number of sessions to the maximum number of sessions. <b>Default</b> no session-assign-method. All new sessions are placed by preference in existing tunnels. <b>Values</b> <i>weighted</i> — Enables weighted preference to tunnels in the group.

## session-limit

<b>Syntax</b>	<b>session-limit</b> <i>session-limit</i> <b>no session-limit</b>
<b>Context</b>	config>service>vprn>l2tp>group config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command configures the session limit. The value controls how many L2TP session will be allowed within a given context (system, group, tunnel). The no form of the command removes the value from the configuration.
<b>Default</b>	no session-limit
<b>Parameters</b>	<i>session-limit</i> — Specifies the allowed number of sessions within the given context. <b>Values</b> 1 — 131071

---

## Router L2TP Tunnel Commands

### tunnel

<b>Syntax</b>	<b>tunnel</b> <i>tunnel-name</i> [ <b>create</b> ] <b>no tunnel</b> <i>tunnel-name</i>
<b>Context</b>	config>service>vprn>l2tp>group
<b>Description</b>	This command configures an L2TP tunnel. A tunnel exists between a LAC-LNS pair and consists of a Control Connection and zero or more L2TP sessions. The tunnel carries encapsulated PPP datagrams and control messages between the LAC and the L2TP Network Server (LNS).
<b>Parameters</b>	<i>tunnel-name</i> — Specifies a valid string to identify a L2TP up to 32 characters in length. <b>create</b> — mandatory while creating a new tunnel

### auto-establish

<b>Syntax</b>	[ <b>no</b> ] <b>auto-establish</b>
<b>Context</b>	config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command specifies if this tunnel is to be automatically set up by the system.
<b>Default</b>	no auto-establish

### avp-hiding

<b>Syntax</b>	<b>avp-hiding</b> { <b>never</b>   <b>sensitive</b>   <b>always</b> } <b>no avp-hiding</b>
<b>Context</b>	config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as cleartext in an AVP.  Note that it is recommended that sensitive information not be sent in clear text.  The <b>no</b> form of the command removes the parameter of the configuration and indicates that the value on group level will be taken.
<b>Default</b>	no avp-hiding
<b>Parameters</b>	<i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnel.  <b>Values</b> never — AVP hiding is not used. sensitive — AVP hiding is used only for sensitive information (such as username/password). always — AVP hiding is always used.

## challenge

<b>Syntax</b>	<b>challenge</b> <i>challenge-mode</i> <b>no challenge</b>
<b>Context</b>	config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command configures the use of challenge-response authentication. The <b>no</b> form of the command removes the parameter from the configuration and indicates that the value on group level will be taken.
<b>Default</b>	no challenge
<b>Parameters</b>	<i>challenge-mode</i> — Specifies when challenge-response is to be used for the authentication of the tunnel. <b>Values</b> always — Always allows the use of challenge-response authentication. never — Never allows the use of challenge-response authentication.

## hello-interval

<b>Syntax</b>	<b>hello-interval</b> <i>hello-interval</i> <b>hello-interval infinite</b> <b>no hello-interval</b>
<b>Context</b>	config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command configures the number of seconds between sending Hellos for a L2TP tunnel. The no form removes the parameter from the configuration and indicates that the value on group level will be taken.
<b>Parameters</b>	<i>hello-interval</i> — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages. <b>Values</b> 60 — 3600 <b>infinite</b> — Specifies that no hello messages are sent.

## idle-timeout

<b>Syntax</b>	<b>idle-timeout</b> <i>idle-timeout</i> <b>idle-timeout infinite</b> <b>no idle-timeout</b>
<b>Context</b>	config>service>vprn>l2tp>group>tunnel
<b>Description</b>	This command configures the idle timeout to wait before being disconnect. The no form indicates that the parameter will be removed from the configuration and that the value specified on group level will be taken.

**Parameters** *idle-timeout* — Specifies the idle timeout, in seconds.

**Values** 0 — 3600

**infinite** — Specifies that the tunnel will not be closed when idle.

## peer

**Syntax** **peer** *ip-address*  
**no peer**

**Context** config>service>vprn>l2tp>group>tunnel

**Description** This command configures the peer address.  
The **no** form of the command removes the IP address from the tunnel configuration.

**Default** no peer

**Parameters** *ip-address* — Sets the LNS IP address for the tunnel.

## preference

**Syntax** **preference** *preference*  
**no preference**

**Context** config>service>vprn>l2tp>group>tunnel

**Description** This command configures a preference number that indicates the relative preference assigned to a tunnel when using a weighted session assignment.  
The **no** form of the command removes the preference value from the tunnel configuration.

**Default** no preference

**Parameters** *preference* — Specifies the tunnel preference number with its group. The value 0 corresponds to the highest preference.

**Values** 0 — 16777215

## remote-name

**Syntax** **remote-name** *host-name*  
**no remote-name**

**Context** config>service>vprn>l2tp>group>tunnel

**Description** This command configures a string to be compared to the host name used by the tunnel peer during the authentication phase of tunnel establishment.

**Parameters** *host-name* — Specifies a remote host name for the tunnel up to 64 characters in length.

---

## Router DHCP Configuration Commands

### dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure DHCP parameters.

### dhcp6

<b>Syntax</b>	<b>dhcp6</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure DHCP6 parameters.

### local-dhcp-server

<b>Syntax</b>	<b>local-dhcp-server</b> <i>server-name</i> [ <b>create</b> ] <b>no local-dhcp-server</b> <i>server-name</i>
<b>Context</b>	config>service>vprn>dhcp config>service>vprn>dhcp6 config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.
<b>Default</b>	none
<b>Parameters</b>	<i>server-name</i> — Specifies the name of local DHCP server. <b>create</b> — Keyword used to create the server name. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.

### failover

<b>Syntax</b>	<b>failover</b>
<b>Context</b>	config>service>vprn>dhcp
<b>Description</b>	This command enables the context to configure failover parameters.

## ignore-mclt-on-takeover

<b>Syntax</b>	<b>ignore-mclt-on-takeover</b> <b>no ignore-mclt-on-takeover</b>
<b>Context</b>	config>service>vprn>dhcp>server>failover configure>router>dhcp6>server>failover configure>router>dhcp6>server>pool configure>service>vprn>dhcp6>server>failover configure>service>vprn>dhcp6>server>pool
<b>Description</b>	<p>With this flag enabled, the ‘remote’ IP address/prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the MCLT to expire. Note that by setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, will still be reduced to the MCLT over time and all new lease times will be set to MCLT □ this behavior remain the same as originally intended for MCLT.</p> <p>Some deployments require that the ‘remote’ IP address/prefix range starts delegating new IP addresses/prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.</p> <p>This can be achieved by enabling the ignore-mclt-on-takeover flag and by configuring the partner-down-delay to 0.</p> <p>Enabling this functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (partner-down-delay and MCLT) would ensure that during their duration, the new IP addresses/prefixes are delegated only from one node – the one with local IP address-range/prefix. The drawback is of course that the new IP address delegation is delayed and thus service is impacted.</p> <p>But if one could ensure that the intercommunication link is always available, then the DHCP nodes would stay in sync and the two timers would not be needed. This is why it is of utmost importance that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses/prefixes.</p>
<b>Default</b>	no ignore-mclt-on-takeover

## maximum-client-lead-time

<b>Syntax</b>	<b>maximum-client-lead-time [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>]</b> <b>no maximum-client-lead-time</b>
<b>Context</b>	config>service>vprn>dhcp>server>failover config>service>vprn>dhcp>server>pool config>router>vprn>dhcp6>server>failover config>router>vprn>dhcp6>server>pool config>service>vprn>dhcp6>server>failover

```
config>service>vprn>dhcp6>server>pool
```

**Description** The maximum-client-lead-time (MCLT) is the maximum time that a DHCP server can extend client's lease time beyond the lease time currently known by the DHCP partner node. In dual-homed environment, the initial lease time for all DHCP clients is by default restricted to MCLT. Consecutive DHCP renewals are allowed to extend the lease time beyond the MCLT.

The MCLT is a safeguard against IP address/prefix duplication in cases of a lease synchronization failure when local-remote failover model is deployed

Once the intercommunication link failure between the redundant DHCP servers is detected, the DHCP IP address range configured as remote will not be allowed to start delegating new leases until the MCLT + partner-down-delay intervals expire. This is to ensure that the new lease that was delegated from the 'local' IP address-range/prefix on one node, but was never synchronized due to the intercommunication link failure, will expire before the same IP address/prefix is allocated from the remote IP address-range/prefix on the other node.

However, the already existing (and synchronized) lease times can be renewed from the remote IP address range at any time, regardless of the state of the intercommunication link (operational or failed).

Lease synchronization failure can be caused either by a node failure, or a failure of the link over which the DHCP leases are synchronized (intercommunication link). Synchronization failure detection can take up to 3 seconds.

During the failure, the DHCP lease time for the new clients will be restricted to MCLT while for the existing clients the lease time will over time (by consecutive DHCP renewals) be gradually reduced to the MCLT.

**Default** 10 minutes

**Parameters** **hrs** *hours* — Specifies the hour parameter of the MCLT.

**Values** 1 — 23

**min** *minutes* — Specifies the minute parameter of the MCLT.

**Values** 1 — 59

**sec** *seconds* — Specifies the seconds parameter of the MCLT.

**Values** 1 — 59

### partner-down-delay

**Syntax** **partner-down-delay** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]  
**no partner-down-delay**

**Context** config>service>vprn>dhcp>server>failover  
config>router>dhcp6>server>failover  
config>router>dhcp6>server>pool  
config>service>vprn>dhcp6>server>failover  
config>service>vprn>dhcp6>server>pool

**Description** Since the DHCP lease synchronization failure can be caused by the failure of the intercommunication link (and not necessary the entire node), there is a possibility the redundant DHCP servers become



isolated in the network. In other words, they can serve DHCP clients but they cannot synchronize the lease. This can lead to duplicate assignment of IP addresses, since the servers have configured overlapping IP address ranges but they are not aware of each other's leases.

The purpose of the `partner-down-delay` is to prevent the IP lease duplication during the intercommunication link failure by not allowing new IP addresses to be assigned from the remote IP address range. This timer is intended to provide the operator with enough time to remedy the failed situation and to avoid duplication of IP addresses/prefixes during the failure.

During the `partner-down-delay` time, the prefix designated as remote will be eligible only for renewals of the existing DHCP leases that have been synchronized by the peering node. Only after the sum of the `partner-down-delay` and the `maximum-client-lead-time` will the prefix designated as remote be eligible for delegation of the new DHCP leases. When this occurs, we say that the remote IP address range has been taken over.

It is possible to expedite the takeover of a remote IP address range so that the new IP leases can start being delegated from that range shortly after the intercommunication failure is detected. This can be achieved by configuring the `partner-down-delay` timer to 0 seconds, along with enabling the `ignore-mclt-on-takeover` CLI flag. Caution must be taken before enabling this functionality. It is safe to bypass safety timers (`partner-down-delay` + MCLT) only in cases where the operator is certain that the intercommunication between the nodes has failed due to the entire node failure and not due to the intercommunication (MCS) link failure. Failed intercommunication due to the nodal failure would ensure that only one node is present in the network for IP address delegation (as opposed to two isolated nodes with overlapping IP address ranges where address duplication can occur). For this reason, the operator MUST ensure that there are redundant paths between the nodes to ensure uninterrupted synchronization of DHCP leases.

In access-driven mode of operation, `partner-down-delay` has no effect.

**Default** 23 hours, 59minutes, and 59 seconds.

**Parameters** `hrs hours` — Specifies the hour parameter of the partner down delay feature.

**Values** 1 — 23

`min minutes` — Specifies the minute parameter of the partner down delay feature.

**Values** 1 — 59

`sec seconds` — Specifies the seconds parameter of the partner down delay feature.

**Values** 1 — 59

## peer

**Syntax** `peer ip-address tag sync-tag-name`  
`no peer ip-address`

**Context** `config>service>vprn>dhcp>server>failover`  
`config>router>dhcp6>server>failover`  
`config>router>dhcp6>server>pool`  
`config>service>vprn>dhcp6>server>failover`  
`config>service>vprn>dhcp6>server>pool`

## Router DHCP Configuration Commands

<b>Description</b>	<p>DHCP leases can be synchronized per DHCP server of DHCP pool. The pair of synchronizing servers or pools is identified by a tag. The synchronization information is carried over the Multi-Chassis Synchronization (MCS) link between the two peers. MCS link is a logical link (IP, or MPLS).</p> <p>MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5sec, MCS will send its own keepalive to the peer. If a reply is NOT received within 3sec, MCS will declare its operation state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this. It can take up to 3 seconds before the DHCP client realizes that the inter-chassis communication link has failed.</p> <p>Note that the inter-chassis communication link failure does not necessarily assume the same failed fate for the access links. In other words the two redundant nodes can become isolated from each other in the network. This would occur in cases where only the intercommunication (MCS) link fails. It is of utmost importance that this MCS link be highly redundant.</p>
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the IPv4 address of the peer</p> <p><b>sync-tag</b> <i>sync-tag</i> — Specifies a synchronization tag to be used while synchronizing DHCP server or pools.</p>

## startup-wait-time

<b>Syntax</b>	<code>[no] startup-wait-time [min <i>minutes</i>] [sec <i>seconds</i>]</code>
<b>Context</b>	<code>config&gt;service&gt;vprn&gt;dhcp6&gt;server&gt;failover</code> <code>config&gt;service&gt;vprn&gt;dhcp6&gt;server&gt;pool</code> <code>config&gt;router&gt;dhcp6&gt;server&gt;failover</code> <code>config&gt;router&gt;dhcp6&gt;server&gt;pool</code>
<b>Description</b>	This command enables startup-wait-time during which each peer waits after the initialization process before assuming the active role for the prefix designated as local or access-driven. This is to avoid transient issues during the initialization process.
<b>Default</b>	2 minutes
<b>Parameters</b>	<p><b>min</b> <i>minutes</i> — Specifies the minute parameter of the startup wait time feature.</p> <p><b>Values</b> 1 — 10</p> <p><b>sec</b> <i>seconds</i> — Specifies the seconds parameter of the startup wait time feature.</p> <p><b>Values</b> 1 — 59</p>

## ignore-rapid-commit

<b>Syntax</b>	<code>[no] ignore-rapid-commit</code>
<b>Context</b>	<code>config&gt;service&gt;vprn&gt;dhcp6&gt;server</code>
<b>Description</b>	This command specifies whether the Rapid Commit Option (RCO) sent by the DHCPv6 client is processed.

If enabled and the client has included an RCO in the solicit, the server ignores the option and processes the remainder of the message as if no RCO were present.

The **no** form of the command disables ignore-rapid-commit.

## lease-hold-time

<b>Syntax</b>	<b>lease-hold-time</b> [ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no lease-hold-time</b>								
<b>Context</b>	config>service>vprn>dhcp6>server								
<b>Description</b>	This command configures the time to remember this lease.								
<b>Parameters</b>	[ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] — S the lease hold time.								
<b>Values</b>	<table> <tr> <td>days:</td> <td>[0..3650]</td> </tr> <tr> <td>hours:</td> <td>[0..23]</td> </tr> <tr> <td>minutes:</td> <td>[0..59]</td> </tr> <tr> <td>seconds:</td> <td>[0..59]</td> </tr> </table>	days:	[0..3650]	hours:	[0..23]	minutes:	[0..59]	seconds:	[0..59]
days:	[0..3650]								
hours:	[0..23]								
minutes:	[0..59]								
seconds:	[0..59]								

## force-renews

<b>Syntax</b>	[ <b>no</b> ] <b>force-renews</b>
<b>Context</b>	config>service>vprn>dhcp>server
<b>Description</b>	This command enables the sending of sending forcerenew messages. The <b>no</b> form of the command disables the sending of forcerenew messages.
<b>Default</b>	no disable-force-renews

## pool

<b>Syntax</b>	<b>pool</b> <i>pool-name</i> [ <b>create</b> ] <b>no pool</b> <i>pool-name</i>
<b>Context</b>	config>service>vprn>dhcp>server
<b>Description</b>	This command configures a DHCP address pool on the router.
<b>Default</b>	none
<b>Parameters</b>	<p><i>pool name</i> — Specifies the name of this IP address pool. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters.</p> <p><b>create</b> — Keyword used to create the entity. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.</p>

## max-lease-time

<b>Syntax</b>	<b>max-lease-time</b> [days <i>days</i> ] [hrs <i>hours</i> ] [min <i>minutes</i> ] [sec <i>seconds</i> ] <b>no max-lease-time</b>								
<b>Context</b>	config>service>vprn>dhcp>server>pool								
<b>Description</b>	This command configures the maximum lease time. The <b>no</b> form of the command returns the value to the default.								
<b>Default</b>	10 days								
<b>Parameters</b>	<i>time</i> — Specifies the maximum lease time.								
<b>Values</b>	<table> <tr> <td>days :</td> <td>0 — 3650</td> </tr> <tr> <td>hours</td> <td>0 — 23</td> </tr> <tr> <td>minutes:</td> <td>0 — 59</td> </tr> <tr> <td>seconds</td> <td>0 — 59</td> </tr> </table>	days :	0 — 3650	hours	0 — 23	minutes:	0 — 59	seconds	0 — 59
days :	0 — 3650								
hours	0 — 23								
minutes:	0 — 59								
seconds	0 — 59								

## min-lease-time

<b>Syntax</b>	<b>min-lease-time</b> [days <i>days</i> ] [hrs <i>hours</i> ] [min <i>minutes</i> ] [sec <i>seconds</i> ] <b>no min-lease-time</b>								
<b>Context</b>	config>service>vprn>dhcp>server>pool								
<b>Description</b>	This command configures the minimum lease time. The <b>no</b> form of the command returns the value to the default.								
<b>Default</b>	10 minutes								
<b>Parameters</b>	<i>time</i> — Specifies the minimum lease time.								
<b>Values</b>	<table> <tr> <td>days :</td> <td>0 — 3650</td> </tr> <tr> <td>hours</td> <td>0 — 23</td> </tr> <tr> <td>minutes:</td> <td>0 — 59</td> </tr> <tr> <td>seconds</td> <td>0 — 59</td> </tr> </table>	days :	0 — 3650	hours	0 — 23	minutes:	0 — 59	seconds	0 — 59
days :	0 — 3650								
hours	0 — 23								
minutes:	0 — 59								
seconds	0 — 59								

## minimum-free

<b>Syntax</b>	<b>minimum-free</b> <i>minimum-free</i> [percent] [event-when-depleted] <b>no minimum-free</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool
<b>Description</b>	This command configures the minimum number of free addresses. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	1

**Parameters** *minimum-free* — Specifies the desired minimum number of free addresses in this pool. If the actual number of free addresses in this pool falls below this configured minimum, a notification is generated.

**Values** 0 — 255

**percent** — Indicates the value indicates a percentage.

**event-when-depleted** — This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

## offer-time

**Syntax** **offer-time** [**min** *minutes*] [**sec** *seconds*]  
**no offer-time**

**Context** config>service>vprn>dhcp>server>pool

**Description** This command configures the offer time.  
The **no** form of the command returns the value to the default.

**Default** 1 minute

**Parameters** *time* — Specifies the offer time.

**Values** minutes: 0 — 10  
seconds 0 — 59

## options

**Syntax** **options**

**Context** config>service>vprn>dhcp>server>pool

**Description** This command enables the context to configure pool options. The options defined here can be overruled by defining the same option in the local user database.

**Default** none

## custom-option

<b>Syntax</b>	<b>custom-option</b> <i>option-number</i> <b>address</b> <i>ip-address</i> [ <i>ip-address...</i> (up to 4 max)] (DHCP only) <b>custom-option</b> <i>option-number</i> <b>address</b> <i>ipv6-address</i> [ <i>ipv6-address...</i> (up to 4 max)] (DHCP6 only) <b>custom-option</b> <i>option-number</i> <b>domain</b> <i>domain-string</i> <b>custom-option</b> <i>option-number</i> <b>hex</b> <i>hex-string</i> <b>custom-option</b> <i>option-number</i> <b>string</b> <i>ascii-string</i> <b>no custom-option</b> <i>option-number</i>
<b>Context</b>	config>service>vprn>dhcp>server>pool>options config>service>vprn>dhcp>server>pool>subnet>options
<b>Description</b>	This command configures specific DHCP options. The options defined here can overrule options in the local user database.  The <b>no</b> form of the removes the option from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>option-number</i> — specifies the option number that the DHCP server uses to send the identification strings to the DHCP client. <b>Values</b> 1 — 254 <b>address</b> <i>ip-address</i> — Specifies the IP address of this host. <b>domain</b> <i>domain-string</i> — <b>hex</b> <i>hex-string</i> — Specifies the hex value of this option. <b>Values</b> 0x0..0xFFFFFFFF...(maximum 254 hex nibbles) <b>string</b> <i>ascii-string</i> — Specifies the value of this option. <b>Values</b> Up to 127 characters maximum.

## dns-server

<b>Syntax</b>	<b>dns-server</b> <i>ip-address</i> [ <i>ip-address...</i> (up to 4 max)](DHCP only) <b>dns-server</b> <i>ipv6-address</i> [ <i>ipv6-address...</i> (up to 4 max)] (DHCP6 only)
<b>Context</b>	config>service>vprn>dhcp>server>pool>options
<b>Description</b>	This command configures the IP address of the DNS server.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — The IP address of the DNS server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

## domain-name

<b>Syntax</b>	<b>domain-name</b> <i>domain-name</i> <b>no domain-name</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>options
<b>Description</b>	This command configures the default domain for a DHCP client that the router uses to complete unqualified hostnames (without a dotted-decimal domain name). The <b>no</b> form of the command removes the name from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>domain-name</i> — Specifies the domain name for the client. <b>Values</b> Up to 127 characters

## renew-timer

<b>Syntax</b>	<b>renew-timer</b> [ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no renew-timer</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix
<b>Description</b>	This command configures the renew-timer (T1), the time at which the client contacts the server from which the addresses in the IA_NA or IA_PD were obtained to extend the lifetimes of the addresses or prefixes assigned to the client.
<b>Default</b>	1800
<b>Parameters</b>	<i>seconds</i> — Specifies the time duration relative to the current time, expressed in units of seconds. A value of zero leaves the renew-time at the discretion of the client. <b>Values</b> 0-604,800

## rebind-timer

<b>Syntax</b>	<b>rebind-timer</b> [ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no rebind-timer</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix
<b>Description</b>	This command configures the rebind-timer (T2), the time at which the client contacts any available server to extend the lifetimes of the addresses or prefixes assigned to the client.
<b>Default</b>	2880
<b>Parameters</b>	<i>seconds</i> — T2 is a time duration relative to the current time. A value of zero leaves the rebind-time at the discretion of the client. <b>Values</b> 0-1,209,600

## Router DHCP Configuration Commands

**Parameters** [days *days*][hrs *hours*] [min *minutes*] [sec *seconds*] — Specifies the rebind timer.

<b>Values</b>	days:	[0..3650]
	hours:	[0..23]
	minutes:	[0..59]
	seconds:	[0..59]

## prefix

**Syntax** **prefix** *ipv6-address/prefix-length* [**failover** {**local** | **remote** | **access-driven**}] [**pd**] [**wan-host**] [**create**]  
**no prefix** *ipv6-address/prefix-length*

**Context** config>router>dhcp6>server>pool  
config>service>vprn>dhcp6>server>pool

**Description** This command allows a list of prefixes(using the prefix command multiple times) to be routed to hosts associated with this pool. Each prefix will be represented in the associated FIB with a reference to the pool. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).

**Default** Failover local.

**Parameters** *ipv6-address* — Specifies the 128-bit IPv6 address.

**Values** 128-bit hexadecimal IPv6 address in compressed form

*prefix-length* — Specifies the length of any associated aggregate prefix.

**Values** 32 — 63

**failover** — This command designates a IPv6 prefix as local, remote or access-driven. This is used when multi-chassis synchronization is enabled.

**local** — An IPv6 prefix designated as local is used for new lease grants or to renew the existing lease grants. Local prefix designation should be always paired with the remote designation of the same prefix on the peering node.

The IPv6 prefix configured as local on one node can only be configured as remote on the other node. No other combination is allowed between the two nodes for an IPv6 prefix that is configured as local.

The dhcpv6 relay could point to both IPv6 DHCP server addresses - the one hosting the local IPv6 prefix and the one hosting the corresponding remote IPv6 prefix. Under normal circumstances the new lease will always be allocated from the local IPv6 prefix while the leases can be renewed from either IPv6 prefix (local or remote). Under network failure, the remote IPv6 prefix can be taken over according to the intercommunication link state transitions and associated timers.

**remote** — An IPv6 prefix designated as remote is used only to renew the existing DHCP leases. The new leases will be delegated from it only after the maximum-client-lead-time + partner-down-delay time elapses. At that point we say that the remote IPv6 prefix has been taken over.

To ensure faster takeover, the partner-down-delay can be set to 0 and the MCLT time can be ignored. Extra caution should be exercised when enabling this mode of operation, as described in the configuration guides.



The IPv6 prefix configured as remote on one node can only be configured as local on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as remote.

**access-driven** — An IPv4 prefix designated as access-driven is used for new lease grants or to renew the existing lease grants regardless of the state of the intercommunication link (operational or failed). In this mode of operation the IPv6 prefix is actively shared between the two 7x50 DHCPv6 server nodes. This can be used on both DHCPv6 servers only in cases where the access protection mechanism (SRRP or MC-LAG) will ensure that there is only a single active path for DHCPv6 clients using the same IPv6 prefix available to one of the redundant 7x50 DHCPv6 nodes.

The IPv6 prefix configured as access-driven on one node can only be configured as access-driven on the other node. No other combination is allowed between the two nodes for an IPv6 prefix that is configured as access-driven.

There MUST be no crosslinks between the DHCPv6 servers that have IPv6 address ranges configured in access-driven failover mode. In other words, each node must have the dhcp-relay pointing to the IPv6 address of the local DHCPv6 server. This IPv6 address must be the same on both nodes. For example, both DHCPv6 servers should have a loopback address configured with the same IPv6 address (IPv4 or IPv6) and a DHCPv6 server associated with this loopback address. Those IPv6 addresses MUST not be advertised outside of each box. The DHCPv6 relay in each node would point to its local DHCPv6 server via this loopback IPv6 address.

**pd** — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

**wan-host** — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime</b> [ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no preferred-lifetime</b>								
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix								
<b>Description</b>	The preferred lifetime for the IPv6 prefix or address in the option, expressed in units of seconds. When the preferred lifetime expires, any derived addresses are deprecated.								
<b>Default</b>	3600								
<b>Parameters</b>	<i>time</i> — Specifies the preferred lifetime.								
<b>Values</b>	<table> <tr> <td>days:</td> <td>0 — 3650</td> </tr> <tr> <td>hours:</td> <td>0 — 23</td> </tr> <tr> <td>minutes:</td> <td>0 — 59</td> </tr> <tr> <td>seconds</td> <td>0 — 59</td> </tr> </table>	days:	0 — 3650	hours:	0 — 23	minutes:	0 — 59	seconds	0 — 59
days:	0 — 3650								
hours:	0 — 23								
minutes:	0 — 59								
seconds	0 — 59								

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime</b> [ <b>days</b> <i>days</i> ][ <b>hrs</b> <i>hours</i> ] [ <b>min</b> <i>minutes</i> ] [ <b>sec</b> <i>seconds</i> ] <b>no valid-lifetime</b>
<b>Context</b>	config>service>vprn>dhcp6>server>pool>prefix
<b>Description</b>	The valid lifetime for the IPv6 prefix or address in the option, expressed in units of seconds.
<b>Default</b>	86,400
<b>Parameters</b>	<i>time</i> — Specifies the valid lifetime.  <b>Values</b> days: 0 — 3650 hours: 0 — 23 minutes: 0 — 59 seconds 0 — 59

## use-link-address

<b>Syntax</b>	<b>use-link-address</b> [ <b>scope</b> <i>scope</i> ] <b>no use-link-address</b>
<b>Context</b>	config>service>vprn>dhcp6>server
<b>Description</b>	This command specifies whether the GI address selects a single subnet or a pool. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	subnet
<b>Parameters</b>	<b>scope</b> <i>scope</i> — Specifies the scope of the IP address selection.  <b>Values</b> subnet, pool

## user-ident

<b>Syntax</b>	<b>user-ident</b> <i>user-ident</i> <b>no user-ident</b>
<b>Context</b>	config>service>vprn>dhcp6>server
<b>Description</b>	This command specifies which method is used by the local DHCP server to uniquely identify a user. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	duid
<b>Parameters</b>	<i>user-ident</i> — Configures the user identification method.  <b>Values</b> duid, interface-id, interface-id-link-local

## lease-rebind-time

<b>Syntax</b>	<b>lease-rebind-time</b> [days <i>days</i> ] [hrs <i>hours</i> ] [min <i>minutes</i> ] [sec <i>seconds</i> ] <b>no lease-rebind-time</b>								
<b>Context</b>	config>service>vprn>dhcp>server>pool>options								
<b>Description</b>	This command configures the time the client transitions to a rebinding state. The <b>no</b> form of the command removes the time from the configuration.								
<b>Default</b>	none								
<b>Parameters</b>	<i>time</i> — Specifies the lease rebind time.								
<b>Values</b>	<table> <tr> <td>days:</td> <td>0 — 3650</td> </tr> <tr> <td>hours:</td> <td>0 — 23</td> </tr> <tr> <td>minutes:</td> <td>0 — 59</td> </tr> <tr> <td>seconds</td> <td>0 — 59</td> </tr> </table>	days:	0 — 3650	hours:	0 — 23	minutes:	0 — 59	seconds	0 — 59
days:	0 — 3650								
hours:	0 — 23								
minutes:	0 — 59								
seconds	0 — 59								

## lease-renew-time

<b>Syntax</b>	<b>lease-renew-time</b> [days <i>days</i> ] [hrs <i>hours</i> ] [min <i>minutes</i> ] [sec <i>seconds</i> ] <b>no lease-renew-time</b>								
<b>Context</b>	config>service>vprn>dhcp>server>pool>options								
<b>Description</b>	This command configures the time the client transitions to a renew state. The <b>no</b> form of the command removes the time from the configuration.								
<b>Default</b>	none								
<b>Parameters</b>	<i>time</i> — Specifies the lease renew time.								
<b>Values</b>	<table> <tr> <td>days:</td> <td>0 — 3650</td> </tr> <tr> <td>hours:</td> <td>0 — 23</td> </tr> <tr> <td>minutes:</td> <td>0 — 59</td> </tr> <tr> <td>seconds</td> <td>0 — 59</td> </tr> </table>	days:	0 — 3650	hours:	0 — 23	minutes:	0 — 59	seconds	0 — 59
days:	0 — 3650								
hours:	0 — 23								
minutes:	0 — 59								
seconds	0 — 59								

## lease-time

<b>Syntax</b>	<b>lease-time</b> [days <i>days</i> ] [hrs <i>hours</i> ] [min <i>minutes</i> ] [sec <i>seconds</i> ] <b>no lease-time</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>options
<b>Description</b>	This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address. The <b>no</b> form of the command removes the lease time parameters from the configuration.
<b>Default</b>	none

## Router DHCP Configuration Commands

<b>Parameters</b>	<i>time</i> — Specifies the lease time.
<b>Values</b>	days : 0 — 3650 hours 0 — 23 minutes: 0 — 59 seconds 0 — 59

### netbios-name-server

<b>Syntax</b>	<b>netbios-name-server ip-address</b> [ <i>ip-address</i> ...(up to 4 max)] <b>no netbios-name-server</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>options
<b>Description</b>	This command configures up to four Network Basic Input/Output System (NetBIOS) name server IP addresses.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — The IP address of the NetBIOS name server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

### netbios-node-type

<b>Syntax</b>	<b>netbios-node-type netbios-node-type</b> <b>no netbios-node-type</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>options
<b>Description</b>	This command configures the Network Basic Input/Output System (NetBIOS) node type.
<b>Default</b>	none
<b>Parameters</b>	<i>netbios-node-type</i> — Specifies the netbios node type.  <b>Values</b> B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name. P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name. M — Mixed node uses broadcasted queries to find a node, and if that fails, queries a known P-node name server for the address. H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.

## server

<b>Syntax</b>	<b>server</b> <i>server-name</i> <b>no server</b>
<b>Context</b>	configure>service>ies>sub-if>grp-if>local-address-assignment configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>grp-if>local-address-assignment configure>service>vprn>sub-if>local-address-assignment
<b>Description</b>	This command designates a local 7x50 DHCPv4 server for local pools management where IPv4 addresses for PPPoXv4 clients will be allocated without the need for the internal 7x50 DHCP relay-agent. Those addresses will be tied to PPPoX sessions and they will be de-allocated when the PPPoX session is terminated.
<b>Default</b>	none
<b>Parameters</b>	<i>server-name</i> — Specifies the name of the local 7x50 DHCP server.

## client-application

<b>Syntax</b>	<b>client-application</b> [ppp-v4] <b>no client-application</b>
<b>Context</b>	configure>service>ies>sub-if>grp-if>local-address-assignment configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>grp-if>local-address-assignment configure>service>vprn>sub-if>local-address-assignment
<b>Description</b>	This command enables local 7x50 DHCP Server pool management for PPPoXv4 clients. A pool of IP addresses can be shared between IPE clients that rely on DHCP protocol (lease renewal process) and PPPoX clients where address allocation is not dependent on DHCP messaging but instead an IP address allocation within the pool is tied to the PPPoX session.
<b>Default</b>	none

## default-pool

<b>Syntax</b>	<b>default-pool</b> <i>pool-name</i> <b>no default-pool</b>
<b>Context</b>	configure>service>ies>sub-if>grp-if>local-address-assignment configure>service>ies>sub-if>local-address-assignment configure>service>vprn>sub-if>grp-if>local-address-assignment configure>service>vprn>sub-if>local-address-assignment
<b>Description</b>	This command references a default DHCP address pool for local PPPoX pool management in case that the pool-name is not returned via Radius or LUDB.
<b>Default</b>	none

## Router DHCP Configuration Commands

**Parameters** *pool-name* — Specifies the name of the local 7x50 DHCP server pool.

### delayed-enable

**Syntax** **delayed-enable** *seconds* [**init-only**]  
**no delayed-enable**

**Context** configure>service>ies>sub-if>local-address-assignment  
configure>service>vprn>sub-if>local-address-assignment

**Description** This command will render the subscriber-interface non operation for the given amount of time once the node is rebooted or once the interface is enabled (no-shutdown). The purpose of this timer is to stall the operation of the subscriber-interface until the MCS database is synchronized.

A typical use case for this timer would be to prevent IP lease duplication for PPPoE clients using local PPPoXv4/v6 pools in redundant DHCPv4/v6 server configuration. Since there is no classical DHCP lease state maintained for local PPPoXv4/v6 pools, the IP addresses will not be synchronized via DHCP Server. Instead they will be synchronized via PPPoX clients whose state is maintained in 7x50. Once the PPPoX subscriber host is synchronized between the two 7x50 nodes, the respective IP address lease will be updated in the respective local pool.

One artifact of this behavior (IP address assignment in local DHCP pools is synchronized via PPPoX clients and not via DHCP server synchronization mechanism) is that during the node boot, the DHCP server must wait for the completion of PPPoX subscriber synchronization via MCS so that it learns which addresses/prefixes are already allocated on the peering node. Since the DHCP server can theoretically start assigning IP addresses before the PPPoX sync is completed, a duplicate address assignment may occur. For example an IP address lease can be granted via DHCP local pools while PPPoX sync is still in progress. Once the PPPoX sync is completed, the DHCP server may discover that the granted IP lease has already been allocated by the peering node. The most recent lease will be kept and the other will be removed from both systems. To prevent this scenario, a configurable timer is set to an arbitrary value that will render sub-if non-operational until the timer expires. The purpose of this timer is to allow the PPPoX sync to complete before subscribers under the sub-intf can be served.

**Default** none

**Parameters** *second* — Specifies in seconds.

**Values** 1 — 1200

### subnet

**Syntax** **subnet** {*ip-address/mask* | *ip-address netmask*} [**create**]  
**no subnet** {*ip-address/mask* | *ip-address netmask*}

**Context** config>service>vprn>dhcp>server>pool

**Description** This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created no IP addresses are made available until a range is defined.

<b>Default</b>	none
<b>Parameters</b>	<p><i>ip-address</i> — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>mask</i> — The subnet mask in dotted decimal notation. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.</p> <p><i>netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.</p> <p><b>create</b> — Keyword used to create the entity. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.</p>

## address-range

<b>Syntax</b>	<b>address-range</b> <i>start-ip-address end-ip-address</i> [ <b>failover</b> { <b>local</b>   <b>remote</b>   <b>access-driven</b> }] <b>no address-range</b> <i>start-ip-address end-ip-address</i>
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet configure>router>dhcp>server>pool>subnet
<b>Description</b>	<p>This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses will be included (other than specific excluded addresses).</p> <p>The only two valid failover combinations between the two redundant DHCP nodes are:</p> <ul style="list-style-type: none"> <li>• local - remote</li> <li>• access-driven - access-driven</li> </ul>
<b>Default</b>	Failover local
<b>Parameters</b>	<p><i>start-ip-address</i> — Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>end-ip-address</i> — Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><b>failover</b> — This command designates an address range as local, remote or access-driven. This is used when multi-chassis synchronization is enabled.</p> <p><b>local</b> — An IPv4 address-range designated as local is used for new lease grants or to renew the existing lease grants. Local address-range designation should be always paired with the remote designation of the same address-range on the peering node.</p> <p>The IP address range configured as local on one node can only be configured as remote on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as local.</p> <p>The dhcp relay could point to both IP DHCP server addresses - the one hosting the local IP address range and the one hosting the corresponding remote IP address range. Under normal circumstances the new lease will always be allocated from the local IP address range while the</p>

## Router DHCP Configuration Commands

leases can be renewed from either IP address range (local or remote). Under network failure, the remote IP address range can be taken over according to the intercommunication link state transitions and associated timers.

**remote** — An IPv4 address-range designated as remote is used only to renew the existing DHCP leases. The new leases will be delegated from it only after the maximum-client-lead-time + partner-down-delay time elapses. At that point we say that the remote IP address range has been taken over.

To ensure faster takeover, the partner-down-delay can be set to 0 and the MCLT time can be ignored. Extra caution should be exercised when enabling this mode of operation, as described in the configuration guides.

The IP address range configured as remote on one node can only be configured as local on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as remote.

**access-driven** — An IPv4 address-range designated as access-driven is used for new lease grants or to renew the existing lease grants regardless of the state of the intercommunication link (operational or failed). In this mode of operation the IP address-range is actively shared between the two 7x50 DHCP server nodes. This can be used on both DHCP servers only in cases where the access protection mechanism (SRRP or MC-LAG) will ensure that there is only a single active path for DHCP clients using the same IP address range available to one of the redundant 7x50 DHCP nodes.

The IP address range configured as access-driven on one node can only be configured as access-driven on the other node. No other combination is allowed between the two nodes for an IP address ranges that is configured as access-driven.

There MUST be no crosslinks between the DHCP servers that have IP address ranges configured in access-driven failover mode. In other words, each node must have the dhcp-relay pointing to the IP address of the local DHCP server. This IP address must be the same on both nodes. For example, both DHCP servers should have a loopback address configured with the same IP address (IPv4 or IPv6) and a DHCP server associated with this loopback address. Those IP addresses MUST not be advertised outside of each box. The DHCP relay in each node would point to its local DHCP server via this loopback IP address.

### drain

<b>Syntax</b>	[no] drain
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet
<b>Description</b>	This command subnet draining which means no new leases can be assigned from this subnet and existing leases are cleaned up upon renew/rebind. The <b>no</b> form of the command means the subnet is active and new leases can be assigned from it.

### exclude-addresses

<b>Syntax</b>	[no] <b>exclude-addresses</b> <i>start-ip-address</i> [ <i>end-ip-address</i> ]
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet



<b>Description</b>	This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.
<b>Default</b>	none
<b>Parameters</b>	<p><i>start-ip-address</i> — Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>end-ip-address</i> — Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p>

## maximum-declined

<b>Syntax</b>	<b>maximum-declined</b> <i>maximum-declined</i> <b>no maximum-declined</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet
<b>Description</b>	This command configures the maximum number of declined addresses allowed.
<b>Default</b>	64
<b>Parameters</b>	<p><i>maximum-declined</i> — Specifies the maximum number of declined addresses allowed.</p> <p><b>Values</b>     0 — 4294967295</p>

## minimum-free

<b>Syntax</b>	<b>minimum-free</b> <i>minimum-free</i> [ <b>percent</b> ] [ <b>event-when-depleted</b> ] <b>no minimum-free</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet
<b>Description</b>	This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.
<b>Default</b>	1
<b>Parameters</b>	<p><i>minimum-free</i> — Specifies the minimum number of free addresses in this subnet.</p> <p><b>Values</b>     0 — 255</p> <p><b>percent</b> — Indicates the value indicates a percentage.</p> <p><b>event-when-depleted</b> — This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.</p>

### default-router

<b>Syntax</b>	<b>default-router</b> <i>ip-address</i> [ <i>ip-address...</i> (up to 4 max)] <b>no default-router</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet
<b>Description</b>	This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.  The <b>no</b> form of the command removes the address(es) from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address of the default router. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

### subnet-mask

<b>Syntax</b>	<b>subnet-mask</b> <i>ip-address</i> <b>no subnet-mask</b>
<b>Context</b>	config>service>vprn>dhcp>server>pool>subnet
<b>Description</b>	This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.  The <b>no</b> form of the command removes the address from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

### use-gi-address

<b>Syntax</b>	<b>[no] use-gi-address</b>
<b>Context</b>	config>service>vprn>dhcp>server
<b>Description</b>	This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnets is not found. If the local-user-db-name is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address the server maps to a local user database and configures the user with no address.  A pool can include multiple subnets. Since the GI is shared by multiple subnets in a subscriber-interface the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.
<b>Default</b>	no use-gi-address

## use-pool-from-client

<b>Syntax</b>	<b>[no] use-pool-from-client</b>
<b>Context</b>	config>service>vprn>dhcp>server config>service>vprn>dhcp6>server
<b>Description</b>	<p>This command specifies if the IP address pool to be used by this server is the pool indicated by the vendor-specific sub-option 13 of the DHCP Option 82.</p> <p>When enabled, the pool indicated by the sub-option 13 is to be used.</p> <p>The <b>no</b> form of the command indicates that the pool selection is specified by the value of <b>use-gi-address</b> setting.</p>

## user-db

<b>Syntax</b>	<b>user-db</b> <i>local-user-db-name</i> <b>no user-db</b>
<b>Context</b>	config>service>vprn>dhcp>server
<b>Description</b>	This command configures a local user database for authentication.
<b>Default</b>	not enabled
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the name of a local user database.

---

## Log Commands

### log

<b>Syntax</b>	<b>log</b>
<b>Context</b>	configure>service>vprn configure>service>vprn>log-id
<b>Description</b>	This command enables the context to configure event stream logging.

### filter

<b>Syntax</b>	<b>[no] filter</b> <i>filter-id</i>
<b>Context</b>	config>service>vprn>log config>service>vprn>log>log-id
<b>Description</b>	<p>This command creates a context for an event filter. An event filter specifies whether to forward or drop an event or trap based on the match criteria.</p> <p>Filters are configured in the <b>filter</b> <i>filter-id</i> context and then applied to a log in the <b>log-id</b> <i>log-id</i> context. Only events for the configured log source streams destined to the log ID where the filter is applied are filtered.</p> <p>Any changes made to an existing filter, using any of the sub-commands, are immediately applied to the destinations where the filter is applied.</p> <p>The <b>no</b> form of the command removes the filter association from log IDs which causes those logs to forward all events.</p>
<b>Default</b>	No event filters are defined.
<b>Parameters</b>	<i>filter-id</i> — The filter ID uniquely identifies the filter.
<b>Values</b>	1 — 1000

### default-action

<b>Syntax</b>	<b>default-action {drop   forward}</b> <b>no default-action</b>
<b>Context</b>	config>service>vprn>log>filter
<b>Description</b>	<p>The default action specifies the action that is applied to events when no action is specified in the event filter entries or when an event does not match the specified criteria.</p> <p>When multiple <b>default-action</b> commands are entered, the last command overwrites the previous command.</p>

The **no** form of the command reverts the default action to the default value (forward).

**Default** **default-action forward** — The events which are not explicitly dropped by an event filter match are forwarded.

**Parameters** **drop** — The events which are not explicitly forwarded by an event filter match are dropped.

**forward** — The events which are not explicitly dropped by an event filter match are forwarded.

## entry

**Syntax** **[no] entry entry-id**

**Context** config>service>vprn>log>filter

**Description** This command is used to create or edit an event filter entry. Multiple entries may be created using unique *entry-id* numbers. The TiMOS implementation exits the filter on the first match found and executes the action in accordance with the action command.

Comparisons are performed in an ascending entry ID order. When entries are created, they should be arranged sequentially from the most explicit entry to the least explicit. Matching ceases when a packet matches an entry. The entry action is performed on the packet, either drop or forward. To be considered a match, the packet must meet all the conditions defined in the entry.

An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and are rendered inactive.

The **no** form of the command removes the specified entry from the event filter. Entries removed from the event filter are immediately removed from all log-id's where the filter is applied.

**Default** No event filter entries are defined. An entry must be explicitly configured.

**Parameters** *entry-id*. The entry ID uniquely identifies a set of match criteria corresponding action within a filter. Entry ID values should be configured in staggered increments so you can insert a new entry in an existing policy without renumbering the existing entries.

**Values** 1 — 999

## action

**Syntax** **action {drop | forward}**  
**no action**

**Context** config>service>vprn>log>filter>entry

**Description** This command specifies a drop or forward action associated with the filter entry. If neither **drop** nor **forward** is specified, the **default-action** will be used for traffic that conforms to the match criteria. This could be considered a No-Op filter entry used to explicitly exit a set of filter entries without modifying previous actions.

Multiple action statements entered will overwrite previous actions.

The **no** form of the command removes the specified **action** statement.

## Log Commands

**Default** Action specified by the **default-action** command will apply.

**Parameters** **drop** — Specifies packets matching the entry criteria will be dropped.  
**forward** — Specifies packets matching the entry criteria will be forwarded.

## match

**Syntax** **[no] match**

**Context** config>service>vprn>log>filter>entry

**Description** This command creates context to enter/edit match criteria for a filter entry. When the match criteria is satisfied, the action associated with the entry is executed.

If more than one match parameter (within one match statement) is specified, then all the criteria must be satisfied (AND functional) before the action associated with the match is executed.

Use the **application** command to display a list of the valid applications.

Match context can consist of multiple match parameters (application, event-number, severity, subject), but multiple **match** statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the *entry-id*.

**Default** No match context is defined.

## application

**Syntax** **application {eq | neq} application-id**  
**no application**

**Context** config>service>vprn>log>filter>entry>match

**Description** This command adds an OS application as an event filter match criterion.

An OS application is the software entity that reports the event. Applications include IP, MPLS, OSPF, CLI, SERVICES etc. Only one application can be specified. The latest **application** command overwrites the previous command.

The **no** form of the command removes the application as a match criterion.

**Default** **no application** — No application match criterion is specified.

**Parameters** **eq | neq** — The operator specifying the type of match. Valid operators are listed in the table below.

Operator	Notes
eq	equal to
neq	not equal to

*application-id* — The application name string.

**Values** port, ppp, rip, route\_policy, rsvp, security, snmp, stp, svcmgr, system, user, vrrp, vtr

## message

**Syntax** **message** {**eq** | **neq**} **pattern** *pattern* [**regex**]  
**no message**

**Context** config>service>vprn>log>filter>entry>match

**Description** This command adds system messages as a match criterion.  
The **no** form of the command removes messages as a match criterion.

**Parameters** **eq** — Determines if the matching criteria should be equal to the specified value.  
**neq** — Determines if the matching criteria should not be equal to the specified value.  
**pattern** *pattern* — Specifies a message up to 400 characters to be used in the match criteria.  
**regex** — Specifies the type of string comparison to use to determine if the log event matches the value of **message** command parameters. When the **regex** keyword is not specified, the default matching algorithm used is a basic substring match.

## number

**Syntax** **number** {**eq** | **neq** | **lt** | **lte** | **gt** | **gte**} *event-id*  
**no number**

**Context** config>service>vprn>log>filter>entry>match

**Description** This command adds an SR OS application event number as a match criterion.  
SR OS event numbers uniquely identify a specific logging event within an application.  
Only one **number** command can be entered per event filter entry. The latest **number** command overwrites the previous command.  
The **no** form of the command removes the event number as a match criterion.

**Default** **no event-number** — No event ID match criterion is specified.

**Parameters** **eq** | **neq** | **lt** | **lte** | **gt** | **gte** — This operator specifies the type of match. Valid operators are listed in the table below. Valid operators are:

Operator	Notes
eq	equal to
neq	not equal to
lt	less than

Operator	Notes
lte	less than or equal to
gt	greater than
gte	greater than or equal to

*event-id* — The event ID, expressed as a decimal integer.

**Values** 1 — 4294967295

## severity

<b>Syntax</b>	<b>severity</b> { <b>eq</b>   <b>neq</b>   <b>lt</b>   <b>lte</b>   <b>gt</b>   <b>gte</b> } <i>severity-level</i> <b>no severity</b>
<b>Context</b>	config>service>vprn>log>filter>entry>match
<b>Description</b>	This command adds an event severity level as a match criterion. Only one severity command can be entered per event filter entry. The latest severity command overwrites the previous command. The <b>no</b> form of the command removes the severity match criterion.
<b>Default</b>	<b>no severity</b> — No severity level match criterion is specified.
<b>Parameters</b>	<b>eq</b>   <b>neq</b>   <b>lt</b>   <b>lte</b>   <b>gt</b>   <b>gte</b> — This operator specifies the type of match. Valid operators are listed in the table below.

Operator	Notes
eq	equal to
neq	not equal to
lt	less than
lte	less than or equal to
gt	greater than
gte	greater than or equal to



*severity-name* — The ITU severity level name. The following table lists severity names and corresponding numbers per ITU standards M.3100 X.733 & X.21 severity levels.

Severity Number	Severity Name
1	cleared
2	indeterminate (info)
3	critical
4	major
5	minor
6	warning

**Values** cleared, intermediate, critical, major, minor, warning

## subject

**Syntax** **subject {eq|neq} subject [regexp]**  
**no subject**

**Context** config>service>vprn>log>filter>entry>match

**Description** This command adds an event subject as a match criterion.

The subject is the entity for which the event is reported, such as a port. In this case the port-id string would be the subject. Only one **subject** command can be entered per event filter entry. The latest **subject** command overwrites the previous command.

The **no** form of the command removes the subject match criterion.

**Default** **no subject** — No subject match criterion specified.

**Parameters** **eq | neq** — This operator specifies the type of match. Valid operators are listed in the following table:

Operator	Notes
eq	equal to
neg	not equal to

*subject* — A string used as the subject match criterion.

**regexp** — Specifies the type of string comparison to use to determine if the log event matches the value of **subject** command parameters. When the **regexp** keyword is specified, the string in the **subject** command is a regular expression string that will be matched against the subject string in the log event being filtered.

When **regexp** keyword is not specified, the **subject** command string is matched exactly by the event filter.

## Log Commands

### log-id

<b>Syntax</b>	<b>[no] log-id</b> <i>log-id</i>
<b>Context</b>	config>service>vprn>log
<b>Description</b>	<p>This command creates a context to configure destinations for event streams.</p> <p>The <b>log-id</b> context is used to direct events, alarms/traps, and debug information to respective destinations.</p> <p>A maximum of 10 logs can be configured.</p> <p>Before an event can be associated with this log-id, the <b>from</b> command identifying the source of the event must be configured.</p> <p>Only one destination can be specified for a <i>log-id</i>. The destination of an event stream can be an in-memory buffer, console, session, snmp-trap-group, syslog, or file.</p> <p>Use the <b>event-control</b> command to suppress the generation of events, alarms, and traps for all log destinations.</p> <p>An event filter policy can be applied in the log-id context to limit which events, alarms, and traps are sent to the specified log-id.</p> <p>Log-IDs 99 and 100 are created by the agent. Log-ID 99 captures all log messages. Log-ID 100 captures log messages with a severity level of major and above.</p> <p>Note that Log-ID 99 provides valuable information for the admin-tech file. Removing or changing the log configuration may hinder debugging capabilities. It is strongly recommended not to alter the configuration for Log-ID 99.</p> <p>The <b>no</b> form of the command deletes the log destination ID from the configuration.</p>
<b>Default</b>	No log destinations are defined.
<b>Parameters</b>	<i>log-id</i> — The log ID number, expressed as a decimal integer.
<b>Values</b>	1 — 100

### to snmp

<b>Syntax</b>	<b>to snmp</b> [ <i>size</i> ]
<b>Context</b>	config>service>vprn>log>log-id
<b>Description</b>	<p>This is one of the commands used to specify the log ID destination. This parameter is mandatory when configuring a log destination. This command instructs the alarms and traps to be directed to the <b>snmp-trap-group</b> associated with <i>log-id</i>.</p> <p>A local circular memory log is always maintained for SNMP notifications sent to the specified snmp-trap-group for the <i>log-id</i>.</p> <p>The source of the data stream must be specified in the <b>from</b> command prior to configuring the destination with the <b>to</b> command.</p> <p>The <b>to</b> command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.</p>

<b>Default</b>	none
<b>Parameters</b>	<i>size</i> — The <i>size</i> parameter defines the number of events stored in this memory log.
	<b>Default</b> 100
	<b>Values</b> 50 — 1024

## to syslog

<b>Syntax</b>	<b>to syslog</b> <i>syslog-id</i>
<b>Context</b>	config>service>vprn>log>log-id
<b>Description</b>	<p>This is one of the commands used to specify the log ID destination. This parameter is mandatory when configuring a log destination.</p> <p>This command instructs the alarms and traps to be directed to a specified syslog. To remain consistent with the standards governing syslog, messages to syslog are truncated to 1k bytes.</p> <p>The source of the data stream must be specified in the <b>from</b> command prior to configuring the destination with the <b>to</b> command.</p> <p>The <b>to</b> command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>syslog-id</i> — Instructs the events selected for the log ID to be directed to the <i>syslog-id</i> . The characteristics of the <i>syslog-id</i> referenced here must have been defined in the <b>config&gt;log&gt;syslog</b> <i>syslog-id</i> context.
	<b>Values</b> 1 — 10

## from

<b>Syntax</b>	<b>from</b> {[ <b>main</b> ] [ <b>change</b> ]} <b>no from</b>
<b>Context</b>	config>service>vprn>log>log-id
<b>Description</b>	<p>This command selects the source stream to be sent to a log destination.</p> <p>One or more source streams must be specified. The source of the data stream must be identified using the <b>from</b> command before you can configure the destination using the <b>to</b> command. The <b>from</b> command can identify multiple source streams in a single statement (for example: <b>from main change debug-trace</b>).</p> <p>Only one <b>from</b> command may be entered for a single <i>log-id</i>. If multiple <b>from</b> commands are configured, then the last command entered overwrites the previous <b>from</b> command.</p> <p>The <b>no</b> form of the command removes all previously configured source streams.</p>
<b>Default</b>	No source stream is configured.

## Log Commands

- Parameters**
- main** — Instructs all events in the main event stream to be sent to the destination defined in the **to** command for this destination *log-id*. The main event stream contains the events that are not explicitly directed to any other event stream. To limit the events forwarded to the destination, configure filters using the **filter** command.
  - change** — Instructs all events in the user activity stream to be sent to the destination configured in the **to** command for this destination *log-id*. The change event stream contains all events that directly affect the configuration or operation of this node. To limit the events forwarded to the change stream destination, configure filters using the **filter** command.

## time-format

- Syntax** **time-format** {**local** | **utc**}
- Context** config>service>vprn>log>log-id
- Description** This command specifies whether the time should be displayed in local or Coordinated Universal Time (UTC) format.
- Default** utc
- Parameters**
- local** — Specifies that timestamps are written in the system's local time.
  - utc** — Specifies that timestamps are written using the UTC value. This was formerly called Greenwich Mean Time (GMT) and Zulu time.

## syslog

- Syntax** [**no**] **syslog** *syslog-id*
- Context** config>service>vprn>log
- Description** This command creates the context to configure a syslog target host that is capable of receiving selected syslog messages from this network element.
- A valid *syslog-id* must have the target syslog host address configured.
- A maximum of 10 syslog-ids can be configured.
- No log events are sent to a syslog target address until the syslog-id has been configured as the log destination (**to**) in the log-id node.
- The syslog ID configured in the **configure/service/vprn** context has a local VPRN scope and only needs to be unique within the specific VPRN instance. The same ID can be reused under a different VPRN service or in the global log context under **config>log**.
- Default** No syslog IDs are defined.
- Parameters** *syslog-id* — The syslog ID number for the syslog destination, expressed as a decimal integer.
- Values** 1 — 10

## address

<b>Syntax</b>	<b>address</b> <i>ip-address</i> <b>no address</b>																						
<b>Context</b>	config>service>vprn>log>syslog																						
<b>Description</b>	<p>This command adds the syslog target host IP address to/from a syslog ID.</p> <p>This parameter is mandatory. If no <b>address</b> is configured, syslog data cannot be forwarded to the syslog target host.</p> <p>Only one address can be associated with a <i>syslog-id</i>. If multiple addresses are entered, the last address entered overwrites the previous address.</p> <p>The same syslog target host can be used by multiple log IDs.</p> <p>The <b>no</b> form of the command removes the syslog target host IP address.</p>																						
<b>Default</b>	<b>no address</b> — There is no syslog target host IP address defined for the syslog ID.																						
<b>Parameters</b>	<i>ip-address</i> — The IP address of the syslog target host in dotted decimal notation.																						
<b>Values</b>	<table> <tr> <td>ipv4-address</td> <td>a.b.c.d</td> </tr> <tr> <td>ipv6-address</td> <td>x:x:x:x:x:x:x:x[-interface]</td> </tr> <tr> <td></td> <td>x:x:x:x:x:d.d.d.d[-interface]</td> </tr> <tr> <td></td> <td>x: [0..FFFF]H</td> </tr> <tr> <td></td> <td>d: [0..255]D</td> </tr> <tr> <td></td> <td>interface: 32 characters maximum, mandatory for link local addresses</td> </tr> <tr> <td></td> <td>ipv6-address x:x:x:x:x:x:x:x[-interface]</td> </tr> <tr> <td></td> <td>x:x:x:x:x:d.d.d.d[-interface]</td> </tr> <tr> <td></td> <td>x: [0..FFFF]H</td> </tr> <tr> <td></td> <td>d: [0..255]D</td> </tr> <tr> <td></td> <td>interface: 32 characters maximum, mandatory for link local addresses</td> </tr> </table>	ipv4-address	a.b.c.d	ipv6-address	x:x:x:x:x:x:x:x[-interface]		x:x:x:x:x:d.d.d.d[-interface]		x: [0..FFFF]H		d: [0..255]D		interface: 32 characters maximum, mandatory for link local addresses		ipv6-address x:x:x:x:x:x:x:x[-interface]		x:x:x:x:x:d.d.d.d[-interface]		x: [0..FFFF]H		d: [0..255]D		interface: 32 characters maximum, mandatory for link local addresses
ipv4-address	a.b.c.d																						
ipv6-address	x:x:x:x:x:x:x:x[-interface]																						
	x:x:x:x:x:d.d.d.d[-interface]																						
	x: [0..FFFF]H																						
	d: [0..255]D																						
	interface: 32 characters maximum, mandatory for link local addresses																						
	ipv6-address x:x:x:x:x:x:x:x[-interface]																						
	x:x:x:x:x:d.d.d.d[-interface]																						
	x: [0..FFFF]H																						
	d: [0..255]D																						
	interface: 32 characters maximum, mandatory for link local addresses																						

## facility

<b>Syntax</b>	<b>facility</b> <i>syslog-facility</i> <b>no facility</b>
<b>Context</b>	cconfig>service>vprn>log>syslog
<b>Description</b>	<p>This command configures the facility code for messages sent to the syslog target host.</p> <p>Multiple syslog IDs can be created with the same target host but each syslog ID can only have one facility code. If multiple facility codes are entered, the last <i>facility-code</i> entered overwrites the previous facility-code.</p> <p>If multiple facilities need to be generated for a single syslog target host, then multiple <b>log-id</b> entries must be created, each with its own filter criteria to select the events to be sent to the syslog target host with a given facility code.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>

**Default** `local7` — syslog entries are sent with the local7 facility code.

**Parameters** `syslog-facility` — The syslog facility name represents a specific numeric facility code. The code should be entered in accordance with the syslog RFC. However, the software does not validate if the facility code configured is appropriate for the event type being sent to the syslog target host.

**Values** kernel, user, mail, systemd, auth, syslogd, printer, netnews, uucp, cron, authpriv, ftp, ntp, logaudit, logalert, cron2, local0, local1, local2, local3, local4, local5, local6, local7

Valid responses per RFC3164, *The BSD syslog Protocol*, are listed in the table below.

Numerical Code	Facility Code
0	kernel
1	user
2	mail
3	systemd
4	auth
5	syslogd
6	printer
7	net-news
8	uucp
9	cron
10	auth-priv
11	ftp
12	ntp
13	log-audit
14	log-alert
15	cron2
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

**Values** 0 — 23

## log-prefix

<b>Syntax</b>	<b>log-prefix</b> <i>log-prefix-string</i> <b>no log-prefix</b>
<b>Context</b>	config>service>vprn>log>syslog
<b>Description</b>	<p>This command adds the string prepended to every syslog message sent to the syslog host. RFC3164, <i>The BSD syslog Protocol</i>, allows an alphanumeric string (tag) to be prepended to the content of every log message sent to the syslog host. This alphanumeric string can, for example, be used to identify the node that generates the log entry. The software appends a colon (:) and a space to the string and it is inserted in the syslog message after the date stamp and before the syslog message content.</p> <p>Only one string can be entered. If multiple strings are entered, the last string overwrites the previous string. The alphanumeric string can contain lowercase (a-z), uppercase (A-Z) and numeric (0-9) characters.</p> <p>The <b>no</b> form of the command removes the log prefix string.</p>
<b>Default</b>	<b>no log-prefix</b> — no prepend log prefix string defined.
<b>Parameters</b>	<i>log-prefix-string</i> — An alphanumeric string of up to 32 characters. Spaces and colons (:) cannot be used in the string.

## level

<b>Syntax</b>	<b>level</b> <i>syslog-level</i> <b>no level</b>
<b>Context</b>	config>service>vprn>log>syslog
<b>Description</b>	<p>This command configures the syslog message severity level threshold. All messages with severity level equal to or higher than the threshold are sent to the syslog target host.</p> <p>Only a single threshold level can be specified. If multiple levels are entered, the last <b>level</b> entered will overwrite the previously entered commands.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Parameters</b>	<i>value</i> — The threshold severity level name.
<b>Values</b>	emergency, alert, critical, error, warning, notice, info, debug

Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
	0	emergency	system is unusable
3	1	alert	action must be taken immediately

Router severity level	Numerical Severity (highest to lowest)	Configured Severity	Definition
4	2	critical	critical condition
5	3	error	error condition
6	4	warning	warning condition
	5	notice	normal but significant condition
1 cleared 2 indeterminate	6	info	informational messages
	7	debug	debug-level messages

## port

<b>Syntax</b>	<b>port</b> <i>value</i> <b>no port</b>
<b>Context</b>	config>service>vprn>log>syslog
<b>Description</b>	<p>This command configures the UDP port that will be used to send syslog messages to the syslog target host.</p> <p>The port configuration is needed if the syslog target host uses a port other than the standard UDP syslog port 514.</p> <p>Only one port can be configured. If multiple <b>port</b> commands are entered, the last entered port overwrites the previously entered ports.</p> <p>The <b>no</b> form of the command reverts to default value.</p>
<b>Default</b>	<b>no port</b>
<b>Parameters</b>	<i>value</i> — The value is the configured UDP port number used when sending syslog messages.
	<b>Values</b> 1 — 65535

## snmp-trap-group

<b>Syntax</b>	<b>[no] snmp-trap-group</b> <i>log-id</i>
<b>Context</b>	config>service>vprn>log
<b>Description</b>	<p>This command creates the context to configure a group of SNMP trap receivers and their operational parameters for a given log-id.</p> <p>A group specifies the types of SNMP traps and specifies the log ID which will receive the group of SNMP traps. A trap group must be configured in order for SNMP traps to be sent.</p> <p>To suppress the generation of all alarms and traps see the <b>event-control</b> command. To suppress alarms and traps that are sent to this log-id, see the <b>filter</b> command. Once alarms and traps are</p>



generated they can be directed to one or more SNMP trap groups. Logger events that can be forwarded as SNMP traps are always defined on the main event source.

The **no** form of the command deletes the SNMP trap group.

**Default** There are no default SNMP trap groups.

**Parameters** *log-id* — The log ID value of a log configured in the **log-id** context. Alarms and traps cannot be sent to the trap receivers until a valid *log-id* exists.

**Values** 1 — 99

## trap-target

**Syntax** **trap-target** *name* [**address** *ip-address*] [**port** *port*] [**snmpv1** | **snmpv2c** | **snmpv3**] **notify-community** *communityName* | *snmpv3SecurityName* [**security-level** {**no-auth-no-privacy** | **auth-no-privacy** | **privacy**}] [**replay**]  
**no trap-target** *name*

**Context** config>service>vprn>log>snmp-trap-group

**Description** This command adds/modifies a trap receiver and configures the operational parameters for the trap receiver. A trap reports significant events that occur on a network device such as errors or failures.

Before an SNMP trap can be issued to a trap receiver, the **log-id**, **snmp-trap-group** and at least one **trap-target** must be configured.

The **trap-target** command is used to add/remove a trap receiver from an **snmp-trap-group**. The operational parameters specified in the command include:

- The IP address of the trap receiver
- The UDP port used to send the SNMP trap
- SNMP version
- SNMP community name for SNMPv1 and SNMPv2c receivers.
- Security name and level for SNMPv3 trap receivers.

A single **snmp-trap-group** *log-id* can have multiple trap-receivers. Each trap receiver can have different operational parameters.

An address can be configured as a trap receiver more than once as long as a different port is used for each instance.

To prevent resource limitations, only configure a maximum of 10 trap receivers.

Note that if the same **trap-target** *name* **port** *port* parameter value is specified in more than one SNMP trap group, each trap destination should be configured with a different *notify-community* value. This allows a trap receiving an application, such as NMS, to reconcile a separate event sequence number stream for each router event log when multiple event logs are directed to the same IP address and port destination.

The **no** form of the command removes the SNMP trap receiver from the SNMP trap group.

**Default** No SNMP trap targets are defined.

**Parameters** *name* — Specifies the name of the trap target up to 28 characters in length.

**address** *ip-address* — The IP address of the trap receiver in dotted decimal notation. Only one IP address destination can be specified per trap destination group.

**Values**

ipv4-address	a.b.c.d (host bits must be 0)
ipv6-address	x:x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d[-interface] x: [0..FFFF]H d: [0..255]D interface: 32 characters maximum, mandatory for link local addresses

**port** *port* — The destination UDP port used for sending traps to the destination, expressed as a decimal integer. Only one port can be specified per **trap-target** statement. If multiple traps need to be issued to the same address then multiple ports must be configured.

**Default** 162

**Values** 1 — 65535

*snmpv1* | *snmpv2c* | *snmpv3* — Specifies the SNMP version format to use for traps sent to the trap receiver.

The keyword **snmpv1** selects the SNMP version 1 format. When specifying **snmpv1**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv1**, then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv2c** selects the SNMP version 2c format. When specifying **snmpv2c**, the **notify-community** must be configured for the proper SNMP community string that the trap receiver expects to be present in alarms and traps messages. If the SNMP version is changed from **snmpv3** to **snmpv2c**, then the **notify-community** parameter must be changed to reflect the community string rather than the *security-name* that is used by **snmpv3**.

The keyword **snmpv3** selects the SNMP version 3 format. When specifying **snmpv3**, the **notify-community** must be configured for the SNMP *security-name*. If the SNMP version is changed from **snmpv1** or **snmpv2c** to **snmpv3**, then the **notify-community** parameter must be changed to reflect the *security-name* rather than the community string used by **snmpv1** or **snmpv2c**.

Pre-existing conditions are checked before the snmpv3SecurityName is accepted. These are:

- The user name must be configured.
- The v3 access group must be configured.
- The v3 notification view must be configured.

**Default** snmpv3

**Values** snmpv1, snmpv2c, snmpv3

**notify-community** *community* | *security-name* — Specifies the community string for **snmpv1** or **snmpv2c** or the **snmpv3** *security-name*. If no **notify-community** is configured, then no alarms nor traps will be issued for the trap destination. If the SNMP version is modified, the **notify-community** must be changed to the proper form for the SNMP version.

**community** — The community string as required by the **snmpv1** or **snmpv2c** trap receiver. The community string can be an ASCII string up to 31 characters in length.

*security-name* — The *security-name* as defined in the `config>system>security>user` context for SNMP v3. The *security-name* can be an ASCII string up to 31 characters in length.

**security-level** {*no-auth-no-privacy* | *auth-no-privacy* | *privacy*} — Specifies the required authentication and privacy levels required to access the views configured on this node when configuring an **snmpv3** trap receiver.

The keyword **no-auth-no-privacy** specifies no authentication and no privacy (encryption) are required.

The keyword **auth-no-privacy** specifies authentication is required but no privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication**.

The keyword **privacy** specifies both authentication and privacy (encryption) is required. When this option is configured the *security-name* must be configured for **authentication** and **privacy**.

**Default** no-auth-no-privacy. This parameter can only be configured if SNMPv3 is also configured.

**Values** no-auth-no-privacy, auth-no-privacy, privacy

**replay** — Enable replay of missed events to target. If replay is applied to an SNMP trap target address, the address is monitored for reachability. Reachability is determined by whether or not there is a route in the routing table by which the target address can be reached. Before sending a trap to a target address, the SNMP module asks the PIP module if there is either an in-band or out-of-band route to the target address. If there is no route to the SNMP target address, the SNMP module saves the sequence-id of the first event that will be missed by the trap target. When the routing table changes again so that there is now a route by which the SNMP target address can be reached, the SNMP module replays (for example, retransmits) all events generated to the SNMP notification log while the target address was removed from the route table. Note that because of route table change convergence time, it is possible that one or more events may be lost at the beginning or end of a replay sequence. The `cold-start-wait` and `route-recovery-wait` timers under `config>log>app-route-notifications` can help reduce the probability of lost events.

---

## Multicast VPN Commands

### mvpn

<b>Syntax</b>	<b>mvpn</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure MVPN-related parameters for the IP VPN.

### auto-discovery

<b>Syntax</b>	<b>[default   mdt-safi] [source-address <i>ip-address</i>]</b>
<b>Context</b>	config>service>vprn>mvpn
<b>Description</b>	<p>This command enables MVPN membership auto-discovery through BGP. When auto-discovery is enabled, PIM peering on the inclusive provider tunnel is disabled. Changing auto-discovery configuration requires shutdown of this VPRN instance.</p> <p>The <b>no</b> form of the command disables MVPN membership auto-discovery through BGP.</p>
<b>Default</b>	<p><b>default</b></p> <p><b>default</b> — Enable AD route exchange based on format defined in ng-MVPN (RFC6514).</p> <p><b>mdt-safi</b> — Enable AD route exchange based on mdt-safi format defined in <i>draft-rosen-vpn-mcast</i>.</p> <p>This command allows optionally to specify a <b>source-address</b> - an IP address to be used by Rosen M-VPN for core diversity non-default IGP instances (not using system IP). Two unique IP addresses for all MVPNs are supported. For instances using default System IP, source address configuration should not be specified to avoid consuming one of the addresses.</p> <p>Explicitly defined <b>source-address</b> allows GRE-encapsulated Rosen MVPN multicast traffic (Default and Data MDT) to originate from a configured IP address, so the source IP address of the GRE packets won't be the default system IP address.</p> <p>Value:</p> <p><i>ip-address</i> — an IPv4 address. To achieve the desired functionality the address should be a pre-configured non-default ISIS or OSPF loopback address for an IGP instance using loopback address different from the system IP loopback.</p>

### c-mcast-signaling

<b>Syntax</b>	<b>c-mcast-signaling {bgp   pim}</b> <b>no c-mcast-signaling</b>
<b>Context</b>	config>service>vprn>mvpn

<b>Description</b>	This command specifies BGP or PIM, for PE-to-PE signaling of CE multicast states. When this command is set to PIM and neighbor discovery by BGP is disabled, PIM peering will be enabled on the inclusive tree.  Changes may only be made to this command when the mvpn node is shutdown.  The <b>no</b> form of the command reverts it back to the default.
<b>Default</b>	mcast-signaling bgp
<b>Parameters</b>	<b>bgp</b> — Specifies to use BGP for PE-to-PE signaling of CE multicast states. Auto-discovery must be enabled.  <b>pim</b> — Specifies to use PIM for PE-to-PE signaling of CE multicast states.

## intersite-shared

<b>Syntax</b>	<b>intersite-shared [persistend-type5-adv] [kat-type5-adv-withdraw]</b> <b>no intersite-shared</b>
<b>Context</b>	config>service>vprn>mvpn
<b>Description</b>	This command specifies whether to use inter-site shared C-trees or not. Optional parameters allow enabling additional intersite-shared functionality as per below description. Not specifying an optional parameter when executing the command disables that parameter.
<b>Default</b>	intersite-shared
<b>Parameters</b>	<b>persistend-type5-adv</b> — when specified for intersite-shared trees enabled, this parameter ensures that Type 5 SA routes are generated for the multicast source even if no joins are present for that source. When the parameter is not specified, the Type 5 SA routes are withdrawn where the prune from the last receiver is received for the multicast source.  <b>kat-type5-adv-withdraw</b> — when specified for intersite-shared trees this parameter allows operators to enable KeepAlive Timers (KAT) on source PEs for ng-MVPN inter-site shared deployments . On a multicast source failure, a KAT expiry on source PEs will trigger a withdraw of Type-5 Source-Active (S-A) route and switch from (C-S,C-G) to (C-*,C-G). When receiver PEs process reflected Type-5 S-A route withdrawals, they will withdraw their Type-7 ng-MVPN routes to the failed multicast source. Note the following: <ul style="list-style-type: none"> <li>• KAT must only be enabled on source PEs.</li> <li>• Functionality is supported with mLDP and RSVP-TE in the P-instance.</li> <li>• Local receiver per (C-S, C-G) must be configured on source PEs running KAT.</li> </ul>

## mdt-type

<b>Syntax</b>	<b>mdt-type {sender-receiver   sender-only   receiver-only}</b>
<b>Context</b>	config>service>vprn>mvpn

## Log Commands

<b>Description</b>	<p>This command allows restricting MVPN instance per PE node to a specific role. By default, MVPN instance on a given PE node assumes the role of being a sender as well as receiver. This creates a mesh of MDT/PMSI across all PE nodes from this PE.</p> <p>This command provides an option to configure either a sender-only or receiver only mode per PE node. Restricting the role of a PE node avoids creating full mesh of MDT/PMSI across all PE nodes that are participating in MVPN instance</p> <p>The <b>no</b> version of this command restores the default (sender-receiver).</p>
<b>Default</b>	mdt-type sender-receiver
<b>Parameters</b>	<p><b>sender-receiver</b> — MVPN has both sender and receivers connected to PE node</p> <p><b>sender-only</b> — MVPN has only senders connected to PE node</p> <p><b>receiver-only</b> — MVPN has only receivers connected to PE node</p>

## red-source-list

<b>Syntax</b>	<b>red-source-list</b>
<b>Context</b>	config>service>vprn>mvpn
<b>Description</b>	This command enables context to configure list of redundant source prefixes for preferred source selection.

## src-prefix

<b>Syntax</b>	<b>src-prefix</b> <i>ip-address/mask</i> [ <i>ip-address/mask ...up to 8 maximum</i> ] <b>no src-prefix</b> <i>ip-address/mask</i>
<b>Context</b>	config>service>vprn>mvpn>red-source-list
<b>Description</b>	<p>This command configures up to 8 multicast source IPv4 prefixes for preferred source selection. Single or multi-line inputs are allowed.</p> <p>The <b>no</b> form of the command deletes specified prefix from the list.</p>
<b>Default</b>	no prefixes are specified
<b>Parameters</b>	<i>ip-address/mask</i> — IPv4 address prefix with mask.

## ipv6

<b>Syntax</b>	<b>ipv6</b>
<b>Context</b>	config>service>vprn>mvpn>red-source-list
<b>Description</b>	This command enables context to configure list of redundant IPv6 source prefixes for preferred source selection.

## src-prefix

<b>Syntax</b>	<b>src-prefix</b> <i>ipv6-ip-address/prefix-length</i> [ <i>ipv6-address/prefix-length ...up to 8 maximum</i> ] <b>no</b> <i>ipv6-ip-address/prefix-length</i>
<b>Context</b>	config>service>vprn>mvpn>red-source-list>ipv6
<b>Description</b>	This command configures up to 8 multicast source IPv6 prefixes for preferred source selection. Single or multi-line inputs are allowed.  The <b>no</b> form of the command deletes specified prefix from the list
<b>Default</b>	no prefixes are specified
<b>Parameters</b>	<i>ipv6-ip-address/mask</i> — IPv6 address prefix with prefix-length.

## rpf-select

<b>Syntax</b>	<b>rpf-select</b>
<b>Context</b>	config>service>vprn>mvpn
<b>Description</b>	This command enables context for VRF extranet mapping for C-instance receivers in this receiver MVPN instance to multicast streams in P-instance core MVPN instances.

## core-mvpn

<b>Syntax</b>	<b>[no] core-mvpn</b> <i>service-id</i>
<b>Context</b>	config>service>vprn>mvpn>rpf-select
<b>Description</b>	This command enables context for VRF extranet mapping for C-instance receivers in this receiver MVPN instance to multicast streams in the specified P-instance core MVPN instance.

## group-prefix

<b>Syntax</b>	<b>group-prefix</b> <i>ip-address/mask</i> [ <i>ip-address/mask...(upto 8 max)</i> ][ <b>starg</b> ] <b>no group-prefix</b> <i>ip-address/mask</i>
<b>Context</b>	config>service>vprn>mvpn>rpf-select>core-mvpn
<b>Description</b>	This command configures multicast group IPv4 prefixes for the MVPN with per-group mapping extranet functionality. Multiple lines are allowed. Duplicate prefixes are ignored.  When the starg option is specified, extranet functionality is enabled for PIM ASM as for the specified group. When the option is not specified (not recommended with PIM ASM), the PIM ASM join will be mapped and data plane will be established, but the control plane will not be updated on SPT switchover, unless the switchover is driven by a CPE router on a receiver side.

## Log Commands

The **no** form of the command deletes specified prefix from the list, or removes mapping of all prefixes if group-prefix any was specified.

**Default** none

**Parameters** *ip-address/mask* — Specifies the IPv4 multicast address prefix with mask.

## provider-tunnel

**Syntax** **provider-tunnel**

**Context** config>service>vprn>mvpn

**Description** This command enables context to configure tunnel parameters for the MVPN.

## inclusive

**Syntax** **inclusive**

**Context** config>service>vprn>mvpn>pt

**Description** This command enables the context for specifying inclusive provider tunnels

## bsr

**Syntax** **bsr {unicast | spmsi}**  
**no bsr**

**Context** config>service>vprn>mvpn>provider-tunnel>inclusive

**Description** This command configures the type of BSR signaling used.  
The no form of the command restores the default.

**Default** no bsr

**Parameters** **unicast** — BSR PDU's are sent/forwarded using unicast PDUs (default).  
**spmsi** — BSR PDU's are sent/forwarded using S-PMSI full mesh.

## mldp

**Syntax** **mldp**  
**no mldp**

**Context** config>service>vprn>mvpn>provider-tunnel>inclusive

**Description** This command enables use of mLDP LSP for the provider tunnel.



**Default** no mldp

## shutdown

**Syntax** **shutdown**  
**no shutdown**

**Context** config>service>vprn>mvpn>provider-tunnel>inclusive>mldp

**Description** This command administratively disables and enables use of mLDP LSP for the provider tunnel.

**Default** no shutdown

## pim

**Syntax** **pim {asm | ssm} grp-ip-address**  
**no pim**

**Context** config>service>vprn>mvpn>pt>inclusive

**Description** This command specifies the PIM mode to use, ASM or SSM, for PIM-based inclusive provider tunnels and the multicast group address to use. Also enables the context for specifying parameters for PIM peering on the inclusive provider tunnel.

Note that auto-discovery must be enabled in order for SSM to operate.

The **no** form of the command removes the pim context including the statements under the context.

**Default** no pim

**Parameters** **asm** — Specifies to use PIM ASM for inclusive provider tunnels.

**ssm** — Specifies to use PIM SSM for inclusive provider tunnels.

**group-address** — Specifies the multicast group address to use.

## hello-interval

**Syntax** **hello-interval hello-interval**  
**no hello-interval**

**Context** config>service>vprn>mvpn>provider-tunnel>inclusive>pim

**Description** This command specifies the interval at which PIM hello messages are transmitted on the PIM inclusive provider tunnel.

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

**Default** 30 seconds

**Parameters** *hello-interval* — Specifies the hello interval, in seconds. A 0 (zero) value disables the sending of hello messages.

**Values** 0 — 255

## hello-multiplier

<b>Syntax</b>	<b>hello-multiplier</b> <i>deci-units</i> <b>no hello-multiplier</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>pim
<b>Description</b>	This command specifies the hello multiplier. The hello-multiplier in conjunction with the hello-interval determines the hold time for a PIM neighbor. Hold time = (hello-interval * hello-multiplier) / 10. The <b>no</b> form of the command reverts the value to the default.
<b>Default</b>	35
<b>Parameters</b>	<i>deci-units</i> — Specifies the value, in multiples of 0.1, for the formula used to calculate the hold time
	<b>Values</b> 20 — 100

## improved-assert

<b>Syntax</b>	<b>[no] improved-assert</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>pim
<b>Description</b>	This command enables improved assert procedure on the PIM inclusive provider tunnel. The <b>no</b> form of the command disables improved assert procedure.
<b>Default</b>	enabled

## three-way-hello

<b>Syntax</b>	<b>[no] three-way-hello</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>pim
<b>Description</b>	This command enables PIM three-way hello on the inclusive provider tunnel. The no form of the command disables the PIM three-way hello.
<b>Default</b>	disabled

## tracking-support

<b>Syntax</b>	<b>[no] tracking-support</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>pim
<b>Description</b>	This command enables the setting of the T bit in the LAN Prune Delay option of the hello message. This indicates the router's capability to disable Join message suppression. The no form of the command disables the setting.
<b>Default</b>	disabled

## rsvp

<b>Syntax</b>	<b>rsvp</b> <b>no rsvp</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive
<b>Description</b>	This command enables the context for specifying RSVP P2MP LSP for the provider tunnel. The <b>no</b> form of the command removes the rsvp context including all the statements in the context.
<b>Default</b>	no rsvp

## enable-bfd-root

<b>Syntax</b>	<b>enable-bfd-root [transmit-interval] [multiplier <i>multiplier</i>]</b> <b>no enable-bfd-root</b>								
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>rsvp								
<b>Description</b>	This command enables unidirectional multi-point BFD session on a sender (Root) PE node for upstream fast failure detection over RSVP-TE P2MP LSP.								
<b>Parameters</b>	<b>transmit-interval</b> — Sets the transmit interval, in milliseconds. <table> <tr> <td><b>Default</b></td> <td>100</td> </tr> <tr> <td><b>Values</b></td> <td>10 — 100,000</td> </tr> </table> <b>multiplier <i>multiplier</i></b> — Sets the multiplier for the BFD session. <table> <tr> <td><b>Default</b></td> <td>3</td> </tr> <tr> <td><b>Values</b></td> <td>3 — 20</td> </tr> </table>	<b>Default</b>	100	<b>Values</b>	10 — 100,000	<b>Default</b>	3	<b>Values</b>	3 — 20
<b>Default</b>	100								
<b>Values</b>	10 — 100,000								
<b>Default</b>	3								
<b>Values</b>	3 — 20								

## enable-bfd-leaf

<b>Syntax</b>	<b>[no] enable-bfd-leaf</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>rsvp
<b>Description</b>	This command enables unidirectional multi-point BFD session on a receiver (leaf) PE node for upstream fast failure detection over RSVP-TE P2MP LSP.

## lsp-template

<b>Syntax</b>	<b>lsp-template</b> <b>no lsp-template</b>
<b>Context</b>	Context config>service>vprn>mvpn>provider-tunnel>inclusive>rsvp
<b>Description</b>	This command specifies the use of automatically created P2MP LSP as the provider tunnel. The P2MP LSP will be signaled using the parameters specified in the template, such as bandwidth constraints, etc.
<b>Default</b>	none

## shutdown

<b>Syntax</b>	<b>shutdown</b> <b>no shutdown</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>rsvp>lsp-template
<b>Description</b>	This command administratively disables and enables use of RSVP P2MP LSP for the provider tunnel.
<b>Default</b>	no shutdown

## wildcard-spmsi

<b>Syntax</b>	<b>wildcard-spmsi</b> <b>no wildcard-spmsi</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive
<b>Description</b>	This command enables RFC6625 (C-*, C-*) S-PMSI functionality for ng-MVPN. When enabled, (C-*, C-*) S-PMSI is used instead of I-PMSI for this MVPN. Wildcard S-PMSI uses I-PMSI LSP template.  The no form disables the (C-*, C-*) S-PMSI functionality.
<b>Default</b>	no wildcard-spmsi

## selective

<b>Syntax</b>	<b>selective</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel
<b>Description</b>	This command enables the context to specify selective provider tunnel parameters.
<b>Default</b>	none

## auto-discovery-disable

<b>Syntax</b>	<b>[no] auto-discovery-disable</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	This command disables C-trees to P-tunnel binding auto-discovery through BGP so it is signaled using PIM join TLVs. This command requires the <b>c-mcast-signaling</b> parameter to be set to PIM. The <b>no</b> form of the command enables multicast VPN membership auto-discovery through BGP.
<b>Default</b>	no auto-discovery-disable

## data-delay-interval

<b>Syntax</b>	<b>data-delay-interval</b> <i>value</i> <b>no data-delay-interval</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	This command specifies the interval, in seconds, before a PE router connected to the source switches traffic from the inclusive provider tunnel to the selective provider tunnel. The no form of the command reverts the value to the default.
<b>Default</b>	3 seconds
<b>Parameters</b>	<i>value</i> — Specifies the data delay interval, in seconds.
<b>Values</b>	3 — 180

## data-threshold

<b>Syntax</b>	<b>data-threshold</b> { <i>c-grp-ip-addr/mask</i>   <i>c-grp-ip-addr netmask</i> } <i>s-pmsi-threshold</i> <b>[pe-threshold-add</b> <i>pe-threshold-add</i> ] <b>[pe-threshold-delete</b> <i>pe-threshold-delete</i> ] <b>data-threshold</b> <i>c-grp-ipv6-addr/prefix-length</i> <i>s-pmsi-threshold</i> <b>[pe-threshold-add</b> <i>pe-</i> <i>threshold-add</i> ] <b>[pe-threshold-delete</b> <i>pe-threshold-delete</i> ] <b>no data-threshold</b> { <i>c-grp-ip-addr/mask</i>   <i>c-grp-ip-addr netmask</i> } <b>no data-threshold</b> <i>c-grp-ipv6-addr/prefix-length</i>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	<p>This command specifies the data rate threshold that triggers the switch from the inclusive provider tunnel to the selective provider tunnel for (C-S, C-G) within the group range. Optionally, PE thresholds for creating/deleting ng-MVPN S-PMSI may also be specified. Omitting the PE thresholds, preserves currently set value (or defaults if never set). Multiple statements (one per a unique group) are allowed in the configuration.</p> <p>The <b>no</b> form of the command removes the values from the configuration.</p>
<b>Default</b>	no data-threshold
<b>Parameters</b>	<p><i>group-address/mask</i> — Specifies a multicast group address and netmask length.</p> <p><i>c-grp-ip-addr/mask</i>   <i>c-grp-ip-addr netmask</i> — Specifies an IPv4 multicast group address and netmask length or network mask.</p> <p><i>c-grp-ipv6-addr/prefix-length</i> — Specifies an IPv6 multicast group address and prefix length.</p> <p><i>s-pmsi-threshold</i> — Specifies the rate, in kilobits per second (kbps). If the rate for a (C-S, C-G) within the specified group range exceeds the threshold, traffic for the (C-S, C-G) will be switched to the selective provider tunnel.</p> <p><i>s-pmsi-threshold-add</i> — Specifies the number of receiver PEs for creating S-PMSI. When the number of receiver PEs for a given multicast group configuration is non-zero and below the threshold and BW threshold is satisfied, S-PMSI is created.</p> <p><i>s-pmsi-threshold-delete</i> — Specifies the number of receiver PEs for deleting S-PMSI. When the number of receiver PEs for a given multicast group configuration is above the threshold, S-PMSI is deleted and the multicast group is moved to I-PMSI or a wildcard S-PMSI. It is recommended that the delete threshold be significantly larger than the add threshold, to avoid re-signaling of S-PMSI as the receiver PE count fluctuates.</p> <p><b>Values</b></p> <p><i>c-grp-ip-addr</i> : multicast group address a.b.c.d</p> <p><i>mask</i> [4..32]</p> <p><i>netmask</i> : a.b.c.d (network bits all 1 and host bits all 0)</p> <p><i>s-pmsi-threshold</i> : [1..4294967294](threshold in kbps)</p> <p><i>c-grp-ipv6-addr</i> : multicast ipv6-address x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p style="padding-left: 100px;">x:x:x:x:x:d.d.d.d</p> <p style="padding-left: 100px;">x [0..FFFF]H</p> <p style="padding-left: 100px;">d [0..255]D</p> <p style="padding-left: 100px;"><i>prefix-length</i> [1..128]</p> <p>pe-threshold-add: [1..65535], if never specified, 65535 is used (add threshold always met)</p> <p>pe-threshold-delete: [2..65535], if never specified, 65535 is used (delete threshold never met)</p>

## join-tlv-packing-disable

<b>Syntax</b>	<b>[no] join-tlv-packing-disable</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective>pim
<b>Description</b>	This command enables packing of MDT join TLVs into a single PDU to improve efficiency, if multiple Join TLVs are available at the time of transmission. The no form of the command disables packing of MDT join TLVs into a single PDU.
<b>Default</b>	no join-tlv-packing-disable

## pim-asm

<b>Syntax</b>	<b>[no] pim-asm {grp-ip-address/mask   grp-ip-address netmask}</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	This command specifies the range of PIM-ASM groups to use on the sender PE to setup ASM multicast tree for draft Rosen based Data MDT.

## rsvp

<b>Syntax</b>	<b>[no] rsvp</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	This command enables use of P2MP RSVP as inclusive or selective provider tunnel.
<b>Default</b>	no rsvp

## lsp-template

<b>Syntax</b>	<b>[no] lsp-template <i>lsp-template-name</i></b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive config>service>vprn>mvpn>provider-tunnel>selective>rsvp
<b>Description</b>	This command specifies the use of automatically created P2MP LSP as the inclusive or selective provider tunnel. The P2MP LSP will be signaled using the parameters specified in the template, such as bandwidth constraints, etc.
<b>Default</b>	no lsp-template

## mldp

<b>Syntax</b>	<b>[no] mldp</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	This command enables use of P2MP mLDP LSP as inclusive or selective PMSI tunnels.
<b>Default</b>	no mldp

## maximum-p2mp-spmsi

<b>Syntax</b>	<b>[no] maximum-p2mp-spmsi</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective
<b>Description</b>	This command specifies the maximum number of RSVP P2MP or LDP P2MP S-PMSI tunnels for the mVPN. Once the limit is reached, no more RSVP P2MP S-PMSI or LDP P2MP S-PMSI is created and traffic over the data-threshold will stay on I-PMSI.
<b>Default</b>	10
<b>Parameters</b>	<i>number</i> — specifies the maximum number of RSVP P2MP or LDP P2MP S-PMSI tunnel for the mVPN.
<b>Values</b>	1 — 4k
<b>Default</b>	10

## shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>inclusive>rsvp>lsp-template config>service>vprn>mvpn>provider-tunnel>inclusive>mldp config>service>vprn>mvpn>provider-tunnel>selective>rsvp>lsp-template config>service>vprn>mvpn>provider-tunnel>selective>mldp
<b>Description</b>	This command administratively disables/enables use of P2MP RSVP LSP template or mLDP LSP for inclusive or selective PMSI tunnels.
<b>Default</b>	no shutdown

## enable-asm-mdt

<b>Syntax</b>	<b>[no] enable-asm-mdt</b>
<b>Context</b>	config>service>vprn>mvpn>provider-tunnel>selective



**Description** This command enables Data MDT with PIM-ASM mode on the receiver PE node. PIM-ASM or PIM-SSM operation mode is derived based on the locally configured SSM range on the node.

If `asm-mode` is disabled using this command, then PIM-SSM mode is enabled for all groups, independent of the configured SSM range on the node.

## pim-ssm

**Syntax** `pim-ssm {grp-ip-address/mask | grp-ip-address netmask}`  
**no pim-ssm**

**Context** `config>service>vprn>mvpn>provider-tunnel>selective`

**Description** This command specifies the PIM SSM groups to use for the selective provider tunnel.

**Parameters** `group-address/mask` — Specifies a multicast group address and netmask length.

## umh-pe-backup

**Syntax** `umh-pe-backup`

**Context** `config>service>vprn>mvpn`

**Description** This command enables context to configure primary and standby upstream PE association for the MVPN.

## umh-pe

**Syntax** `umh-pe ip-address standby ip-address`  
**no umh-pe ip-address**

**Context** `config>service>vprn>mvpn>umh-pe-backup`

**Description** This command assigns a standby PE to each primary PE that must be selected as an alternative PE in case the UFD session on tunnel from primary PE is detected down. Standby for a PE cannot be modified without shutting down the MVPN instance.

If a primary PE is not assigned a standby PE then the UMH selection would fall back to the default method.

## umh-selection

**Syntax** `umh-selection {highest-ip|hash-based|tunnel-status|unicast-rt-pref}`  
**no umh-selection**

**Context** `config>service>vprn>mvpn`

## Log Commands

<b>Description</b>	This command specifies which UMH selection mechanism to use, highest IP address, hash based or provider tunnel status. The <b>no</b> form of the command resets it back to default.
<b>Default</b>	umh-selection highest-ip
<b>Parameters</b>	<b>highest-ip</b> — Specifies that the highest IP address is selected as UMH. <b>hash-based</b> — Specifies that the UMH selection is based on the hash based procedures. <b>tunnel-status</b> — Specifies that UMH selection is based on the state of the tunnel as well as the available unicast routes through the tunnel. <b>unicast-rt-pref</b> — When selected, best unicast route will decide which UMH is chosen. Note that, all PE routers shall prefer the same route to the UMH for the UMH selection criterion (for example BGP path selection criteria must not influence one PE to choose different UMH from another PE).

## vrf-export

<b>Syntax</b>	<b>vrf-export</b> { <b>unicast</b>   <i>policy-name</i> [ <i>policy-name...</i> (up to 16 max)]} <b>no vrf-export</b>
<b>Context</b>	config>service>vprn>mvpn
<b>Description</b>	This command specifies the export policy (up to 16) to control MVPN routes exported from the local VRF to other VRFs on the same or remote PE routers.
<b>Default</b>	vrf-export unicast
<b>Parameters</b>	<b>unicast</b> — Specifies to use unicast VRF export policy for the MVPN. <i>policy</i> — Specifies a route policy name.

## vrf-import

<b>Syntax</b>	<b>vrf-import</b> { <b>unicast</b>   <i>policy-name</i> [ <i>policy-name...</i> (up to 16 max)]} <b>no vrf-import</b>
<b>Context</b>	config>service>vprn>mvpn
<b>Description</b>	This command specifies the import policy (up to 16) to control MVPN routes imported to the local VRF from other VRFs on the same or remote PE routers.
<b>Default</b>	vrf-import unicast
<b>Parameters</b>	<b>unicast</b> — Specifies to use a unicast VRF import policy for the MVPN. <i>policy</i> — Specifies a route policy name.

## vrf-target

<b>Syntax</b>	<b>vrf-target</b> {unicast   <i>ext-community</i>   <b>export unicast</b>   <i>ext-community</i>   <b>import unicast</b>   <i>ext-community</i> } <b>no vrf-target</b>								
<b>Context</b>	config>service>vprn>mvpn								
<b>Description</b>	This command specifies the route target to be added to the advertised routes or compared against the received routes from other VRFs on the same or remote PE routers. vrf-import or vrf-export policies override the vrf-target policy.  The <b>no</b> form of the command removes the vrf-target.								
<b>Default</b>	no vrf-target								
<b>Parameters</b>	<b>unicast</b> — Specifies to use unicast vrf-target <i>ext-community</i> for the multicast VPN.  <i>ext-comm</i> — An extended BGP community in the <b>type:x:y</b> format. The value <b>x</b> can be an integer or IP address. The <b>type</b> can be the target or origin. <b>x</b> and <b>y</b> are 16-bit integers.  <b>Values</b> target: { <i>ip-address:comm-val</i>   <i>2byte-asnumber:ext-comm-val</i>   <i>4byte-asnumber:comm-val</i> } <table> <tr> <td><i>ip-address:</i></td> <td>a.b.c.d</td> </tr> <tr> <td><i>comm-val:</i></td> <td>0 — 65535</td> </tr> <tr> <td><i>2byte-asnumber:</i></td> <td>1 — 65535</td> </tr> <tr> <td><i>4byte-asnumber</i></td> <td>0 — 4294967295</td> </tr> </table> <b>import</b> <i>ext-community</i> — Specify communities allowed to be accepted from remote PE neighbors. <b>export</b> <i>ext-community</i> — Specify communities allowed to be sent to remote PE neighbors.	<i>ip-address:</i>	a.b.c.d	<i>comm-val:</i>	0 — 65535	<i>2byte-asnumber:</i>	1 — 65535	<i>4byte-asnumber</i>	0 — 4294967295
<i>ip-address:</i>	a.b.c.d								
<i>comm-val:</i>	0 — 65535								
<i>2byte-asnumber:</i>	1 — 65535								
<i>4byte-asnumber</i>	0 — 4294967295								

## export

<b>Syntax</b>	<b>export</b> {unicast   <i>ext-community</i> }								
<b>Context</b>	config>service>vprn>mvpn>vrf-target								
<b>Description</b>	This command specifies communities to be sent to peers.								
<b>Parameters</b>	<b>unicast</b> — Specifies to use unicast vrf-target <i>ext-community</i> for the multicast VPN.  <i>ext-comm</i> — An extended BGP community in the <b>type:x:y</b> format. The value <b>x</b> can be an integer or IP address. The <b>type</b> can be the target or origin. <b>x</b> and <b>y</b> are 16-bit integers.  <b>Values</b> target: { <i>ip-address:comm-val</i>   <i>2byte-asnumber:ext-comm-val</i>   <i>4byte-asnumber:comm-val</i> } <table> <tr> <td><i>ip-address:</i></td> <td>a.b.c.d</td> </tr> <tr> <td><i>comm-val:</i></td> <td>0 — 65535</td> </tr> <tr> <td><i>2byte-asnumber:</i></td> <td>1 — 65535</td> </tr> <tr> <td><i>4byte-asnumber</i></td> <td>0 — 4294967295</td> </tr> </table>	<i>ip-address:</i>	a.b.c.d	<i>comm-val:</i>	0 — 65535	<i>2byte-asnumber:</i>	1 — 65535	<i>4byte-asnumber</i>	0 — 4294967295
<i>ip-address:</i>	a.b.c.d								
<i>comm-val:</i>	0 — 65535								
<i>2byte-asnumber:</i>	1 — 65535								
<i>4byte-asnumber</i>	0 — 4294967295								

## import

<b>Syntax</b>	<b>import</b> {unicast   ext-community}
<b>Context</b>	config>service>vprn>mvpn>vrf-target
<b>Description</b>	This command specifies communities to be accepted from peers.
<b>Parameters</b>	<b>unicast</b> — Specifies to use unicast vrf-target ext-community for the multicast VPN. <i>ext-comm</i> — An extended BGP community in the <b>type:x:y</b> format. The value <b>x</b> can be an integer or IP address. The <b>type</b> can be the target or origin. <b>x</b> and <b>y</b> are 16-bit integers. <b>Values</b> target: { <i>ip-address:comm-val</i>   <i>2byte-asnumber:ext-comm-val</i>   <i>4byte-asnumber:comm-val</i> } <i>ip-address:</i> a.b.c.d <i>comm-val:</i> 0 — 65535 <i>2byte-asnumber:</i> 1 — 65535 <i>4byte-asnumber</i> 0 — 4294967295

---

## Network Time Protocol Commands

### ntp

<b>Syntax</b>	<b>[no] ntp</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure Network Time Protocol (NTP) and its operation. This protocol defines a method to accurately distribute and maintain time for network elements. Furthermore this capability allows for the synchronization of clocks between the various network elements. Use the no form of the command to stop the execution of NTP and remove its configuration.
<b>Default</b>	none

### authenticate

<b>Syntax</b>	<b>[no] authenticate</b>
<b>Context</b>	config>service>vprn>ntp
<b>Description</b>	This command enables authentication for the NTP server.

### authentication-check

<b>Syntax</b>	<b>[no] authentication-check</b>
<b>Context</b>	config>service>vprn>ntp
<b>Description</b>	<p>This command provides the option to skip the rejection of NTP PDUs that do not match the authentication key-id, type or key requirements. The default behavior when authentication is configured is to reject all NTP protocol PDUs that have a mismatch in either the authentication key-id, type or key.</p> <p>When <b>authentication-check</b> is enabled, NTP PDUs are authenticated on receipt. However, mismatches cause a counter to be increased, one counter for type and one for key-id, one for type, value mismatches. These counters are visible in a show command.</p> <p>The <b>no</b> form of this command allows authentication mismatches to be accepted; the counters however are maintained.</p>
<b>Default</b>	authentication-check — Rejects authentication mismatches.

## authentication-key

<b>Syntax</b>	<b>authentication-key</b> <i>key-id</i> { <b>key</b> <i>key</i> } [ <b>hash</b>   <b>hash2</b> ] <b>type</b> { <b>des</b>   <b>message-digest</b> } <b>no authentication-key</b> <i>key-id</i>
<b>Context</b>	config>service>vprn>ntp
<b>Description</b>	This command sets the authentication key-id, type and key used to authenticate NTP PDUs sent to or received by other network elements participating in the NTP protocol. For authentication to work, the authentication key-id, type and key value must match.  The <b>no</b> form of the command removes the authentication key.
<b>Default</b>	none
<b>Parameters</b>	<p><i>key-id</i> — Configure the authentication key-id that will be used by the node when transmitting or receiving Network Time Protocol packets.</p> <p>Entering the authentication-key command with a key-id value that matches an existing configuration key will result in overriding the existing entry.</p> <p>Recipients of the NTP packets must have the same authentication key-id, type, and key value in order to use the data transmitted by this node. This is an optional parameter.</p> <p><b>Default</b>     None</p> <p><b>Values</b>     1 — 255</p> <p><b>key</b> — The authentication key associated with the configured key-id, the value configured in this parameter is the actual value used by other network elements to authenticate the NTP packet.</p> <p>The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the hash or hash2 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 or hash2 parameter specified.</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, this means that hash2 encrypted variable can't be copied and pasted. If the hash or hash2 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 or hash2 parameter specified.</p> <p><b>type</b> — This parameter determines if DES or message-digest authentication is used.</p> <p>This is a required parameter; either DES or message-digest must be configured.</p> <p><b>Values</b>     <b>des</b> — Specifies that DES authentication is used for this key                        <b>message-digest</b> — Specifies that MD5 authentication in accordance with RFC 2104 is used for this key.</p>

**Note:** des is not permitted in FIPS-140-2 mode.

## broadcast

<b>Syntax</b>	<b>broadcast</b> { <b>interface</b> <i>ip-int-name</i> } [ <b>key-id</b> <i>key-id</i> ] [ <b>version</b> <i>version</i> ] [ <b>ttl</b> <i>ttl</i> ] <b>no broadcast</b> { <b>interface</b> <i>ip-int-name</i> }
<b>Context</b>	config>service>vprn>ntp
<b>Description</b>	This command configures the node to transmit NTP packets on a given interface. Broadcast and multicast messages can easily be spoofed, thus, authentication is strongly recommended. The <b>no</b> form of this command removes the address from the configuration.
<b>Parameters</b>	<p><i>ip-int-name</i> — Specifies the local interface on which to transmit NTP broadcast packets. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><b>Values</b> 32 character maximum</p> <p><b>key-id</b> <i>key-id</i> — Identifies the configured authentication key and authentication type used by this node to receive and transmit NTP packets to and from an NTP server and peers. If an NTP packet is received by this node both authentication key and authentication type must be valid otherwise the packet will be rejected and an event/trap generated.</p> <p><b>Values</b> 1 — 255</p> <p><b>Default</b> none</p> <p><b>version</b> <i>version</i> — Specifies the NTP version number that is generated by this node. This parameter does not need to be configured when in client mode in which case all versions will be accepted.</p> <p><b>Values</b> 1 — 4</p> <p><b>Default</b> 4</p> <p><b>tll</b> <i>ttl</i> — Specifies the IP Time To Live (TTL) value.</p> <p><b>Values</b> 1 — 255</p> <p><b>Default</b> none</p>

---

## Redundant Interface Commands

### redundant-interface

<b>Syntax</b>	<b>[no] redundant-interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command configures a redundant interface.
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### address

<b>Syntax</b>	<b>address</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>remote-ip</b> <i>ip-address</i> ] <b>no address</b>
<b>Context</b>	config>service>vprn>redundant-interface
<b>Description</b>	This command assigns an IP address mask or netmask and a remote IP address to the interface.
<b>Parameters</b>	<i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface. <i>ip-address netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains. Assigns an IP address netmask to the interface. <b>remote-ip ip-address</b> — Assigns a remote IP to the interface.



## SDP Commands

### spoke-sdp

**Syntax** [no] **spoke-sdp** *sdp-id*

**Context** config>service>vprn

**Description** This command binds a service to an existing Service Distribution Point (SDP). A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPRN service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

**Default** No *sdp-id* is bound to a service.

**Special Cases** **VPRN** — Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP binding is rejected.

**Parameters** *sdp-id* — The SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

*vc-id* — The virtual circuit identifier.

**Values** 1 — 4294967295

### control-channel-status

**Syntax** [no] **control-channel-status**

**Context** config>service>vprn>spoke-sdp

**Description** This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF.

A control-channel-status no shutdown is allowed only if all of the following are true:

## SDP Commands

- The system is using network chassis mode D
- SDP signaling is off.
- The control-word is enabled (the control-word is disabled by default)
- The service type is Epipe, Apipe, VPLS, Cpipe, or IES/VPRN
- Mate SDP signaling is off (in vc-switched services)
- The pw-path-id is configured for this spoke-SDP.

The **no** form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shut down.

**Default** no control-channel-status

## acknowledgment

**Syntax** [no] acknowledgment

**Context** config>service>vprn>spoke-sdp>control-channel-status

**Description** This command enables the acknowledgement of control channel status messages. By default, no acknowledgement packets are sent.

## refresh-timer

**Syntax** refresh-timer *value*  
no refresh-timer

**Context** config>service>vprn>spoke-sdp>control-channel-status

**Description** This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

**Default** no refresh-timer

**Parameters** *value* — Specifies the refresh timer value.

**Values** 10 — 65535 seconds

**Default** 0 (off)

## request-timer

**Syntax** request-timer *timer1* retry-timer *timer2* timeout-multiplier *multiplier*  
no request-timer

**Context** config>service>vprn>spoke-sdp>control-channel-status

**Description** This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel

status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

- Parameters**
- timer1* — Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the “request” bit set, are sent.
- Values** 10 — 65535 seconds
- retry-timer** *timer2* — specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.
- Values** 0, 3 — 60 seconds
- timeout-multiplier** *multiplier* — If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional.
- Values** 3 — 20 seconds

## control-word

- Syntax** [no] control-word
- Context** config>service>vprn>spoke-sdp
- Description**
- The control word command provides the option to add a control word as part of the packet encapsulation for pseudowire types for which the control word is optional. These are Ethernet pseudowires (Epipe). ATM N:1 cell mode pseudowires (apipe vc-types atm-vcc and atm-vpc) and VT pseudowire (apipe vc-type atm-cell).
- The configuration for the two directions of the pseudowire must match because the control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported. The C-bit in the pseudowire FEC sent in the label mapping message is set to 1 when the control word is enabled. Otherwise, it is set to 0.
- The service will only come up if the same C-bit value is signaled in both directions. If a spoke-sdp is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an “Illegal C-bit” status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer will withdraw its original label and will send a label mapping with the C-bit set to 1 and the VLL service will be up in both nodes. The control word must be enabled to allow MPLS-TP OAM to be used on a static spoke-sdp in a apipe, epipe and cpipe service.

## pw-path-id

- Syntax** [no] pw-path-id
- Context** config>service>vprn>spoke-sdp
- Description**
- This command enables the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID.

For an IES or VPRN spoke-sdp, the pw-path-id is only valid for ethernet spoke-sdps.

The **pw-path-id** is only configurable if all of the following is true:

- The system is using network chassis mode D
- SDP signaling is off
- control-word is enabled (control-word is disabled by default)
- the service type is epipe, vpls, cpipe, apipe, or IES/VPRN interface
- mate SDP signaling is off for vc-switched services

The **no** form of the command deletes the PW path ID.

**Default** no pw-path-id

## agi

**Syntax** **agi** *agi*  
**no agi**

**Context** config>service>vprn>spoke-sdp>pw-path-id

**Description** This command configures the attachment group identifier for an MPLS-TP PW.

**Parameters** *agi* — Specifies the attachment group identifier.

**Values** 0 — 4294967295

## saii-type2

**Syntax** **saii-type2** *global-id:node-id:ac-id*  
**no saii-type2**

**Context** config>service>vprn>spoke-sdp>pw-path-id

**Description** This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the taii-type2 of the mate spoke-sdp.

**Parameters** *global-id* — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

**Values** 0 — 4294967295

*node-id* — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

**Values** a.b.c.d or 0 — 4294967295

*ac-id* — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

**Values** 1 — 4294967295

## taii-type2

<b>Syntax</b>	<b>taii-type2</b> <i>global-id:node-id:ac-id</i> <b>no taii-type2</b>
<b>Context</b>	config>service>vprn>spoke-sdp>pw-path-id
<b>Description</b>	This command configures the target individual attachment identifier (TAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the saii-type2 of the mate spoke-sdp.
<b>Parameters</b>	<p><i>global-id</i> — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p><b>Values</b> 0 — 4294967295</p> <p><i>node-id</i> — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.</p> <p><b>Values</b> a.b.c.d or 0 — 4294967295</p> <p><i>ac-id</i> — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.</p> <p><b>Values</b> 1 — 4294967295</p>

## spoke-sdp

<b>Syntax</b>	<b>spoke-sdp</b> <i>sdp-id</i> [: <i>vc-id</i> ] <i>vc-type</i> { <b>ether</b>   <b>ipipe</b> } [ <b>create</b> ] <b>no spoke-sdp</b> <i>sdp-id</i> [: <i>vc-id</i> ] <i>vc-type</i> { <b>ether</b>   <b>ipipe</b> } [ <b>create</b> ]
<b>Context</b>	config>service>vprn>if
<b>Description</b>	<p>This command binds a service to an existing Service Distribution Point (SDP).</p> <p>A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.</p> <p>The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.</p> <p>The SDP must already be defined in the <b>config&gt;service&gt;sdp</b> context in order to associate an SDP with a service. If the <b>sdp</b> <i>sdp-id</i> is not already configured, an error message is generated. If the <i>sdp-id</i> does exist, a binding between that <i>sdp-id</i> and the service is created.</p> <p>SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SR devices can participate in the service.</p> <p>Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN services. All packets are forwarded over the default LSP.</p> <p>The <b>no</b> form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.</p>

## SDP Commands

<b>Default</b>	none
<b>Special Cases</b>	<b>VPRN</b> — Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7750 SR router. If two sdp-id bindings terminate on the same 7750 SR, an error occurs and the second SDP binding is rejected.  <i>sdp-id</i> — The SDP identifier. <b>Values</b> 1 — 17407  <i>vc-id</i> — The virtual circuit identifier. <b>Values</b> 1 — 4294967295  <i>vc-type</i> — The encapsulation and pseudowire type for the spoke SDP. <b>Values</b> ether—Ethernet pseudowire. ipipe—IP pseudowire. <b>Default</b> ether

## egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>vprn>if>spoke-sdp config>service>vprn>red-if>spoke-sdp
<b>Description</b>	This command configures an SDP context.

## hash-label

<b>Syntax</b>	<b>[no] hash-label</b>
<b>Context</b>	config>service>vprn config>service>vprn>spoke-sdp config>service>vprn>if>spoke-sdp
<b>Description</b>	<p>This command enables the use of the hash label on a VLL, VPLS, or VPRN service bound to LDP or RSVP SDP as well as to a VPRN service using the autobind mode with the <code>ldp</code>, <code>rsvp-te</code>, or <code>mpls</code> options. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the <code>gre</code> option..</p> <p>When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to 1 to indicate that.</p> <p>In order to allow for applications whereby the egress LER infers the presence of the Hash Label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static</p>

service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. Note however that for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets that are generated in CPM and forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The **no** form of this command disables the use of the hash label.

**Default** no hash-label

## ingress

**Syntax** **ingress**

**Context** config>service>vprn>if>spoke-sdp  
config>service>vprn>red-if>spoke-sdp

**Description** This command configures the SDP context.

## qos

**Syntax** **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*  
**no qos**

**Context** config>service>vprn>interface>spoke-sdp>ingress

**Description** This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers, which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC that needs to be redirected and optionally, for each traffic type (unicast, broadcast, unknown, or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface to which the pseudowire packets can be received. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step,

which means the same network QoS policy can redirect different pseudowires to different queue-group templates.

4. Apply this network QoS policy to the ingress context of a spoke-SDP inside a service, or to the ingress context of a pseudowire template, and specify the redirect queue-group name.
5. One or more spoke-SDPs can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-SDP to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:
  - a When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as *policer-output-queues*.
  - b When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
4. If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the MDA/FP.
5. If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
  - a the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP. This is the default behavior.
  - b a queue-group policer followed by the per-FP ingress shared queues referred to as *policer-output-queues* if the ingress context of the network IP interface from which the packet is received is redirected to a queue-group (*csc-policing*). The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from an R-VPLS spoke-SPD, which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet FC defined in the network-queue policy applied to the ingress of the MDA/FP is used.



When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless of whether an instance of the named policer queue-group exists on the ingress FP on which the pseudowire packet is received. The user can apply a QoS filter matching the dot1.p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload IP header if the user enabled the **ler-use-dscp** option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface on which the pseudowire packet is received.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

- Parameters** *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.
- Values** 1 — 65535
- fp-redirect-group** *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.
- ingress-instance** *instance-id* — Specifies the identification of a specific instance of the queue-group.
- Values** 1 — 16384

## vc-label

- Syntax** **vc-label** *egress-vc-label*  
**no vc-label** [*egress-vc-label*]
- Context** config>service>vprn>if>spoke-sdp>egress  
config>service>vprn>red-if>spoke-sdp>egress
- Description** This command configures the egress VC label.
- Parameters** *vc-label* — A VC egress value that indicates a specific connection.
- Values** 16 — 1048575

## vc-label

- Syntax** **vc-label** *ingress-vc-label*  
**no vc-label** [*ingress-vc-label*]
- Context** config>service>vprn>if>spoke-sdp>ingress  
config>service>vprn>red-if>spoke-sdp>ingress
- Description** This command configures the ingress VC label.

## SDP Commands

**Parameters** *vc-label* — A VC ingress value that indicates a specific connection.

**Values** 2048 — 18431

## egress

**Syntax** **egress**

**Context** config>service>vprn>network-interface

**Description** This command enables the context to configure egress network filter policies for the interface.

## filter

**Syntax** **filter ip** *ip-filter-id*  
**filter ipv6** *ipv6-filter-id*  
**no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

**Context** config>service>vprn>network-interface>egress  
config>service>vprn>if>spoke-sdp>egress  
config>service>vprn>if>spoke-sdp>ingress  
config>service>vprn>red-if>spoke-sdp>ingress  
config>service>vprn>red-if>spoke-sdp>egress  
config>service>vprn>nw-if>egress

**Description** This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. An IP filter policy can be associated with spoke SDPs. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria.

The filter command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress SAP. The *ip-filter-id* must already be defined before the filter command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use scope command within the filter definition to change the scope to local or global. The default scope of a filter is local.

**Parameters** **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

**Values** 1 — 65535

## qos

<b>Syntax</b>	<b>qos</b> <i>network-policy-id</i> <b>port-redirect-group</b> <i>queue-group-name</i> [ <b>instance</b> <i>instance-id</i> ] <b>no qos</b> [ <i>network-policy-id</i> ]
<b>Context</b>	configure>service>apipe>spoke-sdp>egress configure>service>cpipe>spoke-sdp>egress configure>service>epipe>spoke-sdp>egress configure>service>fpipe>spoke-sdp>egress configure>service>ipipe>spoke-sdp>egress config>service>vpls>spoke-sdp>egress config>service>vpls>mesh-sdp>egress config>service>pw-template>egress config>service>vprn>interface>spoke-sdp>egress config>service>ies>interface>spoke-sdp>egress
<b>Description</b>	<p>This command is used to redirect pseudowire packets to an egress port queue-group for the purpose of shaping.</p> <p>The egress pseudowire shaping provisioning model allows the mapping of one or more pseudowires to the same instance of queues, or policers and queues, which are defined in the queue-group template.</p> <p>Operationally, the provisioning model consists of the following steps:</p> <ol style="list-style-type: none"> <li>1. Create an egress queue-group template and configure queues only or policers and queues for each FC that needs to be redirected.</li> <li>2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface on which the pseudowire packets can be forwarded. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.</li> <li>3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different pseudowires to different queue-group templates.</li> <li>4. Apply this network QoS policy to the egress context of a spoke-SPD inside a service or to the egress context of a pseudowire template and specify the redirect queue-group name.</li> </ol> <p>One or more spoke-SPDs can have their FCs redirected to use queues only or queues and policers in the same queue-group instance.</p> <p>The following are the constraints and rules of this provisioning model:</p> <ol style="list-style-type: none"> <li>1. When a pseudowire FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface on which the pseudowire packet is forwarded. This queue can be a queue-group queue, or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a pseudowire packet.</li> </ol>

2. When a pseudowire FC is redirected to use a queue or a policer, and a queue in a queue-group and the queue-group name exists, but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SPD to the named queue-group. In such a case, the pseudowire packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the pseudowire packet is forwarded on.
3. When a pseudowire FC is redirected to use a queue, or a policer and a queue in a queue-group, and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports which have network IP interfaces. The handling of this is dealt with in the data path as follows:
  - a. When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
  - b. When a pseudowire packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the pseudowire packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
4. If a network QoS policy is applied to the egress context of a pseudowire, any pseudowire FC, which is not explicitly redirected in the network QoS policy, will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the pseudowire is redirected to exists and the redirection succeeds, the marking of the packet DEI/dot1.p/DSCP and the tunnel DEI/dot1.p/DSCP/EXP is performed; according to the relevant mappings of the (FC, profile) in the egress context of the network QoS policy applied to the pseudowire. This is true regardless, whether an instance of the queue-group exists or not on the egress port to which the pseudowire packet is forwarded. If the packet profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet DEI/dot1.p and the tunnel DEI/dot1.p/EXP, but the DSCP is not modified by the policer operation.

When the queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the marking of the packet DEI/dot1.p/DSCP and the tunnel DEI/dot1.p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface to which the pseudowire packet is forwarded.

The **no** version of this command removes the redirection of the pseudowire to the queue-group.

#### Parameters

*network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

**Values** 1 — 65535

**port-redirect-group** *queue-group-name* — This optional parameter specifies that the *queue-group-name* will be used for all egress forwarding class redirections within the network QoS policy ID. The specified *queue-group-name* must exist as a port egress queue group on the port associated with the IP interface.

**egress-instance** *instance-id* — Specifies the identification of a specific instance of the queue-group.

**Values** 1 — 16384

## Interface Commands

### interface

<b>Syntax</b>	<b>interface</b> <i>ip-int-name</i> <b>no interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.</p> <p>The <b>interface</b> command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The <b>interface</b> command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for internet access.</p> <p>Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for <b>config router interface</b> and <b>config service vprn interface</b>. Interface names must not be in the dotted decimal notation of an IP address. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.</p> <p>The available IP address space for local subnets and routes is controlled with the <b>config router service-prefix</b> command. The <b>service-prefix</b> command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into <b>config router</b> and <b>config service</b> domains.</p> <p>When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.</p> <p>By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.</p> <p>The <b>no</b> form of this command removes the interface and all the associated configuration. The interface must be administratively shutdown before issuing the <b>no interface</b> command.</p> <p>For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the <b>shutdown</b> command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.</p>
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for <b>config router interface</b> and <b>config service vprn interface</b> commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## Interface Commands

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

### active-cpm-protocols

<b>Syntax</b>	<b>[no] active-cpm-protocols</b>
<b>Context</b>	config>service>vprn>if
<b>Description</b>	This command enables CPM protocols on this interface.

### address

<b>Syntax</b>	<b>address</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>broadcast</b> <i>all-ones</i>   <b>host-ones</b> ] [ <b>track-srrp</b> <i>srrp-instance</i> ] <b>no address</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	<p>Assigns an IP address, IP subnet, and broadcast address format to a VPRN IP router interface. Only one IP address can be associated with an IP interface.</p> <p>An IP address must be assigned to each VPRN IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7750 SR.</p> <p>The local subnet that the <b>address</b> command defines must be part of the services address space within the routing context using the <b>config router service-prefix</b> command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the <b>config router interface</b> CLI context for network core connectivity with the <b>exclude</b> option in the <b>config router service-prefix</b> command.</p> <p>The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.</p> <p>By default, no IP address or subnet association exists on an IP interface until it is explicitly created.</p> <p>Use the <b>no</b> form of this command to remove the IP address assignment from the IP interface. When the <b>no address</b> command is entered, the interface becomes operationally down.</p>

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

*ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

*/* — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the “/” and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

*mask-length* — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (*/*) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 0 – 30. Note that a mask length of 32 is reserved for system IP addresses.

*mask* — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

**broadcast** — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface.

**Default**     host-ones

**all-ones** — The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

**host-ones** — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

**track-srrp** — Specifies the SRRP instance ID that this interface route needs to track.

## allow-directed-broadcasts

<b>Syntax</b>	<b>[no] allow-directed-broadcasts</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	<p>This command controls the forwarding of directed broadcasts out of the IP interface.</p> <p>A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The <b>allow-directed-broadcasts</b> command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.</p> <p>When enabled, a frame destined to the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.</p> <p>When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.</p> <p>By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.</p> <p>The <b>no</b> form of this command disables the forwarding of directed broadcasts out of the IP interface.</p>
<b>Default</b>	<b>no allow-directed-broadcasts</b> — Directed broadcasts are dropped.

## bfd

<b>Syntax</b>	<b>bfd transmit-interval [receive receive-interval] [multiplier multiplier] [echo-receive echo-interval] [type cpm-np]</b> <b>no bfd</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>if>ipv6 config>service>vprn>nw-if



<b>Description</b>	<p>This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.</p> <p>The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault.</p> <p>The <b>no</b> form of the command removes BFD from the associated IGP protocol adjacency.</p> <p><b>Important notes:</b> On the 7750 SR and , the <i>transmit-interval</i>, <i>receive receive-interval</i>, and <i>echo-receive echo-interval</i> values can only be modified to a value less than 100 when:</p> <ol style="list-style-type: none"> <li>1. The <b>type cpm-np option</b> is explicitly configured.</li> <li>2. The service is shut down (<b>shutdown</b>)</li> <li>3. The interval is specified 10 — 100000.</li> <li>4. The service is re-enabled (<b>no shutdown</b>)</li> </ol> <p>To remove the <b>type cpm-np</b> option, re-issue the <b>bfd</b> command without specifying the <b>type</b> parameter.</p>
<b>Default</b>	no bfd
<b>Parameters</b>	<p><i>transmit-interval</i> — Sets the transmit interval for the BFD session.</p> <p><b>Values</b> 10 — 100000 10 — 100000 (see Important Notes above)</p> <p><b>Default</b> 100</p> <p><i>receive receive-interval</i> — Sets the receive interval for the BFD session.</p> <p><b>Values</b> 10 — 100000 10 — 100000 (see Important Notes above)</p> <p><b>Default</b> 100</p> <p><i>multiplier multiplier</i> — Set the multiplier for the BFD session.</p> <p><b>Values</b> 3— 20</p> <p><b>Default</b> 3</p> <p><i>echo-receive echo-interval</i> — Sets the minimum echo receive interval, in milliseconds, for the BFD session.</p> <p><b>Values</b> 100 — 100000 10 — 100000 (see Important Notes above)</p> <p><b>Default</b> 100</p> <p><b>type cpm-np</b> — Specifies that BFD sessions associated with this interface will be created on the CPM network processor to allow for fast timers down to 10ms granularity.</p>

## cflowd-parameters

<b>Syntax</b>	<b>cflowd-parameters</b> <b>no cflowd-parameters</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	This command creates the configuration context to configure cflowd parameters for the associated IP interfaces.  <b>cflowd</b> is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement.  At a minimum, the <b>sampling</b> command must be configured within this context in order to enable cflowd sampling, otherwise traffic sampling will not occur.
<b>Default</b>	no cflowd-parameters

## sampling

<b>Syntax</b>	<b>sampling {unicast   multicast} type {acl   interface} [direction {ingress-only egress-only both}]</b> <b>no sampling {unicast   multicast}</b>
<b>Context</b>	config>service>vprn>if>cflowd-parameters config>service>vprn>nw-if>cflowd-parameters
<b>Description</b>	This command enables and configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis.  This command can be used to configure the sampling parameters for unicast and multicast traffic separately. If sampling is not configured for either unicast or multicast traffic, then that type of traffic will not be sampled.  If cflowd is enabled without either <b>egress-only</b> or <b>both</b> specified or with the <b>ingress-only</b> keyword specified, then only ingress sampling will be enabled on the associated IP interface.  The <b>no</b> form of the command disables the associated type of traffic sampling on the associated interface.
<b>Default</b>	no sampling
<b>Parameters</b>	<b>unicast</b> — Specifies that the <b>sampling</b> command will control the sampling of unicast traffic on the associated interface/SAP.  <b>multicast</b> — Specifies that the <b>sampling</b> command will control the sampling of multicast traffic on the associated interface/SAP.  <b>type</b> —  <b>Values</b> acl — Specifies that the sampled traffic is controlled via an IP traffic filter entry with the action “filter-sample” configured. interface — Specifies that all traffic entering or exiting the interface is subject to sampling.

**direction** — Specifies the direction to collect traffic flow samples.

**Values**      ingress-only — Enables ingress sampling only on the associated interface.  
                   egress-only — Enables egress sampling only on the associated interface.  
                   both — Enables both ingress and egress cflowd sampling.

## cpu-protection

<b>Syntax</b>	<b>cpu-protection</b> <i>policy-id</i> <b>no cpu-protection</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	<p>This command assigns an existing CPU protection policy to the associated service interface. For these interface types, the per-source rate limit is not applicable. The CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection&gt;policy</b> <i>cpu-protection-policy-id</i> context.</p> <p>If no CPU protection policy is assigned to a service interface, then a the default policy is used to limit the overall-rate.</p> <p>The <b>no</b> form of the command removes CPU protection policy association from the interface, resulting in no default rate limiting of control packets.</p>
<b>Default</b>	<p>cpu-protection 254 (for access interfaces)          cpu-protection 255 (for network interfaces)          none (for video-interfaces (where applicable), shown as <b>no cpu-protection</b> in CLI)</p> <p>The configuration of <b>no cpu-protection</b> returns the interface/SAP to the default policies as shown above.</p>
<b>Parameters</b>	<i>policy-id</i> — Specifies an existing CPU protection policy.
<b>Values</b>	1 — 255

## cpu-protection

<b>Syntax</b>	<b>cpu-protection</b> <i>policy-id</i> [ <b>mac-monitoring</b> ] <b>no cpu-protection</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>if>sap
<b>Description</b>	<p>This command assigns an existing CPU protection policy to the associated service group interface SAP, interface or MSAP policy. The CPU protection policies are configured in the <b>config&gt;sys&gt;security&gt;cpu-protection&gt;policy</b> <i>cpu-protection-policy-id</i> context.</p> <p>If no CPU protection policy is assigned to a service group interface SAP, then a the default policy is used to limit the overall-rate.</p>
<b>Default</b>	cpu-protection 254 (for access interfaces)

## Interface Commands

cpu-protection 255 (for network interfaces)

none (for video-interfaces (where applicable), shown as **no cpu-protection** in CLI)

The configuration of **no cpu-protection** returns the interface/SAP to the default policies as shown above.

**Parameters** *policy-id* — Specifies an existing CPU protection policy.

**Values** 1 — 255

**mac-monitoring** — When specified, the per MAC rate limiting should be performed, using the per-source-rate from the associated cpu-protection policy.

## dad-disable

**Syntax** [no] **dad-disable**

**Context** config>service>vprn>if>ipv6

**Description** This command disables duplicate address detection (DAD) on a per-interface basis. This prevents the router from performing a DAD check on the interface. All IPv6 addresses of an interface with DAD disabled, immediately enter a preferred state, without checking for uniqueness on the interface. This is useful for interfaces which enter a looped state during troubleshooting and operationally disable themselves when the loop is detected, requiring manual intervention to clear the DAD violation.

The **no** form of the command turns off **dad-disable** on the interface.

**Default** not enabled

## dist-cpu-protection

**Syntax** **dist-cpu-protection** *policy-name*  
**no dist-cpu-protection**

**Context** config>service>vprn>if>nw-if

**Description** This command assigns a Distributed CPU Protection (DCP) policy to the network interface. Only a valid created DCP policy can be assigned to a SAP or a network interface (note that this rule does not apply to templates such as an msap-policy)

**Default** no dist-cpu-protection

## delayed-enable

**Syntax** **delayed-enable** *seconds*  
**no delayed-enable**

**Context** config>service>vprn>if

```
config>service>vprn>if>nw-if
config>service>vprn>if>redundant-if
```

- Description** This command will cause a delay in the activation of an IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up. The **no** form of the command removes the command from the active configuration and removes the delay in activating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it expires.
- Parameters** *seconds* — Specifies a delay, in seconds, to make the interface operational.
- Values** 1 — 1200

## ingress

- Syntax** **ingress**
- Context** config>service>vprn>if
- Description** This command enters context to configure ingress parameters for network interfaces.

## policy-accounting

- Syntax** **policy-accounting** *<template-name>*  
**no policy-accounting**
- Context** config>service>vprn>if>ingress
- Description** This command configures the service vprn interface ingress policy accounting
- Parameters** *template-name* — Name of template (32 characters maximum)

## ip-mtu

- Syntax** **ip-mtu** *octets*  
**no ip-mtu**
- Context** config>service>vprn>if
- Description** This command configures the IP maximum transmit unit (packet) for this interface. The **no** form of the command returns the default value.
- Default** no ip-mtu

## Interface Commands

### ipcp

<b>Syntax</b>	<b>ipcp</b>
<b>Context</b>	config>service>vprn>if
<b>Description</b>	<p>This command creates allows access to the IPCP context within the interface configuration. Within this context, IPCP extensions can be configured to define such things as the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface.</p> <p>This command is only applicable if the associated SAP/port is a PPP/MLPPP interface.</p>
<b>Default</b>	none

### dns

<b>Syntax</b>	<b>dns ip-address [secondary ip-address]</b> <b>dns secondary ip-address</b> <b>no dns [ip-address] [secondary ip-address]</b>
<b>Context</b>	config>service>vprn>if>ipcp
<b>Description</b>	<p>This command defines the dns address(es) to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions.</p> <p>This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.</p> <p>The <b>no</b> form of the command deletes either the specified primary DNS address, secondary DNS address or both addresses from the IPCP extension peer-ip-address configuration.</p>
<b>Default</b>	no dns
<b>Parameters</b>	<p><i>ip-address</i> — This parameter specifies a unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.</p> <p><b>secondary ip-address</b> — This parameter specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions.</p>

### peer-ip-address

<b>Syntax</b>	<b>peer-ip-address ip-address</b> <b>no peer-ip-address</b>
<b>Context</b>	config>service>vprn>if>ipcp
<b>Description</b>	<p>This command defines the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions.</p> <p>This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.</p> <p>The interface must be shut down to modify the IPCP configuration.</p>

The **no** form of the command deletes the IPCP extension peer-ip-address configuration.

**Default** no peer-ip-address (0.0.0.0)

**Parameters** *ip-address* — Specifies a unicast IPv4 address to be signaled to the far-end of the associated PPP/MLPPP link by IPCP extensions.

## ipv6

**Syntax** [no] ipv6

**Context** config>service>vprn>if

**Description** This command configures an IPv6 interface.

## address

**Syntax** **address** *ipv6-address/mask* [**eui-64**] [**preferred**]  
**no address** *ipv6-address/prefix-length*

**Context** config>service>vprn>if>ipv6

**Description** This command assigns an IPv6 address to the VPRN interface.

**Parameters** *ipv6-address/prefix-length* — Specifies the IPv6 address on the interface.

**Values** *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  
d [0 — 255]D  
*prefix-length* 1 — 128

**eui-64** — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

**preferred** — specifies that the IPv6 address is the preferred IPv6 address for this interface. Preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. Preferred addresses maybe used as the source (or destination) address of packets sent from (or to) the interface. Preferred address doesn't go through the DAD process.

## dhcp6-relay

**Syntax** [no] dhcp6-relay

**Context** config>service>vprn>if>ipv6

**Description** This command configures DHCPv6 relay parameters for the VPRN interface.

## Interface Commands

### dhcp6-server

<b>Syntax</b>	<b>[no] dhcp6-server</b>
<b>Context</b>	config>service>vprn>if>ipv6
<b>Description</b>	This command configures DHCPv6 server parameters for the VPRN interface.

### icmp6

<b>Syntax</b>	<b>icmp6</b>
<b>Context</b>	config>service>vprn>if>ipv6
<b>Description</b>	This command configures ICMPv6 for the interface.

### link-local-address

<b>Syntax</b>	<b>link-local-address <i>ipv6-address</i> [preferred]</b> <b>no link-local-address</b>
<b>Context</b>	config>router>if>ipv6 config>service>ies>if>ipv6 config>service>vprn>if>ipv6
<b>Description</b>	This command configures the IPv6 link local address.  The no form of the command removes the configured link local address, and the router automatically generates a default link local address.  Note that removing a manually configured link local address may impact routing protocols or static routes that have a dependency on that address. It is not recommended to remove a link local address when there are active IPv6 subscriber hosts on an IES or VPRN interface.
<b>Parameters</b>	<b>preferred</b> — Disables duplicated address detection and sets the address to preferred, even if there is a duplicate address.

### local-proxy-nd

<b>Syntax</b>	<b>[no] local-proxy-nd</b>
<b>Context</b>	config>service>vprn>if>ipv6
<b>Description</b>	This command enables or disables neighbor discovery on the interface.



## neighbor

<b>Syntax</b>	<b>neighbor</b> <i>ipv6-address mac-address</i> <b>no neighbor</b> <i>ipv6-address</i>												
<b>Context</b>	config>service>vprn>if>ipv6												
<b>Description</b>	This command configures IPv6-to-MAC address mapping on the interface.												
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the IPv6 address on the interface. <table> <tr> <td><b>Values</b></td> <td><i>ipv6-address</i></td> <td>x:x:x:x:x:x:x (eight 16-bit pieces)</td> </tr> <tr> <td></td> <td></td> <td>x:x:x:x:x:d.d.d</td> </tr> <tr> <td></td> <td></td> <td>x [0 — FFFF]H</td> </tr> <tr> <td></td> <td></td> <td>d [0 — 255]D</td> </tr> </table> <i>mac-address</i> — Specifies the 48-bit MAC address for the static ARP 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.	<b>Values</b>	<i>ipv6-address</i>	x:x:x:x:x:x:x (eight 16-bit pieces)			x:x:x:x:x:d.d.d			x [0 — FFFF]H			d [0 — 255]D
<b>Values</b>	<i>ipv6-address</i>	x:x:x:x:x:x:x (eight 16-bit pieces)											
		x:x:x:x:x:d.d.d											
		x [0 — FFFF]H											
		d [0 — 255]D											

## neighbor-limit

<b>Syntax</b>	<b>neighbor-limit</b> <i>limit</i> [ <b>log-only</b> ] [ <b>threshold percent</b> ] <b>no neighbor-limit</b>				
<b>Context</b>	config>service>vprn>if>ipv6				
<b>Description</b>	This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface.  When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.  The <b>no</b> form of the command removes the <b>neighbor-limit</b> .				
<b>Default</b>	90 percent				
<b>Parameters</b>	<b>log-only</b> — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.  <i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent. <table> <tr> <td><b>Values</b></td> <td>0 — 100</td> </tr> </table> <i>limit</i> — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned. <table> <tr> <td><b>Values</b></td> <td>0 — 102400</td> </tr> </table>	<b>Values</b>	0 — 100	<b>Values</b>	0 — 102400
<b>Values</b>	0 — 100				
<b>Values</b>	0 — 102400				

## | proxy-nd-policy

<b>Syntax</b>	<b>proxy-nd-policy</b> <i>policy-name</i> [ <i>policy-name...</i> (up to 5 max)] <b>no proxy-nd-policy</b>
<b>Context</b>	config>service>vprn>if>ipv6
<b>Description</b>	This command configures a proxy neighbor discovery policy for the interface.
<b>Parameters</b>	<i>policy-name</i> — Specifies the existing policy name(s).

## load-balancing

<b>Syntax</b>	<b>load-balancing</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations.
<b>Default</b>	not applicable

## | egr-ip-load-balancing

<b>Syntax</b>	<b>egr-ip-load-balancing</b> { <b>source</b>   <b>destination</b>   <b>inner-ip</b> } <b>no egr-ip-load-balancing</b>
<b>Context</b>	config>service>vprn>if>load-balancing config>service>vprn>if>nw-if>load-balancing
<b>Description</b>	This command specifies whether to include the source address or destination address or both in the LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled, the command also applies to the inclusion of source/destination port in the hash inputs.  The <b>no</b> form of this command includes both source and destination parameters.
<b>Default</b>	no egr-ip-load-balancing
<b>Parameters</b>	<b>source</b> — Specifies using the source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port.  <b>destination</b> — Specifies using the destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.  <b>inner-ip</b> — Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

## lsr-load-balancing

<b>Syntax</b>	<b>lsr-load-balancing</b> <i>hashing-algorithm</i> <b>no lsr-load-balancing</b>
<b>Context</b>	config>service>vprn>nw-if>load-balancing
<b>Description</b>	This command specifies whether the IP header is used in the LAG and ECMP LSR hashing algorithm. This is the per interface setting.
<b>Default</b>	no lsr-load-balancing
<b>Parameters</b>	<p><b>lbl-only</b> — Only the label is used in the hashing algorithm.</p> <p><b>lbl-ip</b> — The IP header is included in the hashing algorithm.</p> <p><b>ip-only</b> — the IP header is used exclusively in the hashing algorithm</p> <p><b>eth-encap-ip</b> — The hash algorithm parses down the label stack (up to 3 labels supported) and once it hits the bottom, the stack assumes Ethernet II non-tagged header follows. At the expected Ethertype offset location, algorithm checks whether the value present is IPv4/v6 (0x0800 or 0x86DD). If the check passes, the hash algorithm checks the first nibble at the expected IP header location for IPv4/IPv6 (0x0100/0x0110). If the secondary check passes, the hash is performed using IP SA/DA fields in the expected IP header; otherwise (if any of the checks failed) label-stack hash is performed.</p>

## spi-load-balancing

<b>Syntax</b>	<b>[no] spi-load-balancing</b>
<b>Context</b>	config>service>vprn>if>load-balancing config>service>vprn>nw-if>load-balancing
<b>Description</b>	<p>This command enables use of the SPI in hashing for ESP/AH encrypted IPv4/v6 traffic. This is a per interface setting.</p> <p>The <b>no</b> form disables the SPI function.</p>
<b>Default</b>	disabled

## teid-load-balancing

<b>Syntax</b>	<b>[no] teid-load-balancing</b>
<b>Context</b>	config>service>vprn>interface>load-balancing config>service>vprn>nw-if>load-balancing
<b>Description</b>	This command enables inclusion of TEID in hashing for GTP-U/C encapsulates traffic for GTPv1/GTPv2. The <b>no</b> form of this command ignores TEID in hashing.
<b>Default</b>	disabled

## local-proxy-arp

<b>Syntax</b>	<b>[no] local-proxy-arp</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.
<b>Default</b>	no local-proxy-arp

## loopback

<b>Syntax</b>	<b>[no] loopback</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated interface cannot be bound to a SAP.  When using mtrace/mstat in a Layer 3 VPN context then the configuration for the VPRN should have a loopback address configured which has the same address as the core instance's system address (BGP next-hop).
<b>Default</b>	None

## mac

<b>Syntax</b>	<b>[no] mac <i>ieee-mac-address</i></b>
<b>Context</b>	config>service>vprn>if config>service>vprn>if>vrrp config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command assigns a specific MAC address to a VPRN IP interface.  The <b>no</b> form of this command returns the MAC address of the IP interface to the default value.
<b>Default</b>	The physical MAC address associated with the Ethernet interface that the SAP is configured on.
<b>Parameters</b>	<i>ieee-mac-address</i> — Specifies the 48-bit MAC address for the static ARP 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.

## ntp-broadcast

<b>Syntax</b>	<b>[no] ntp-broadcast</b>
<b>Context</b>	config>service>vprn>nw-if
<b>Description</b>	This command enables receiving of NTP/SNTP broadcasts on the interface/

## monitor-oper-group

<b>Syntax</b>	<b>monitor-oper-group</b> <i>name</i> <b>no monitor-oper-group</b>
<b>Context</b>	config>service>vprn>if
<b>Description</b>	This command specifies the operational group to be monitored by the object under which it is configured. The oper-group name must be already configured under the config>service context before its name is referenced in this command.  The <b>no</b> form of the command removes the association from the configuration.
<b>Default</b>	no monitor-oper-group
<b>Parameters</b>	<i>name</i> — Specifies a character string of maximum 32 ASCII characters identifying the group instance.

## proxy-arp

<b>Syntax</b>	<b>[no] proxy-arp</b>
<b>Context</b>	config>service>vprn>nw-if
<b>Description</b>	This command enables proxy ARP on the interface.
<b>Default</b>	no proxy-arp

## proxy-arp-policy

<b>Syntax</b>	<b>[no] proxy-arp-policy</b> <i>policy-name</i> [ <i>policy-name...</i> (up to 5 max)]
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command enables a proxy ARP policy for the interface.  The no form of this command disables the proxy ARP capability.
<b>Default</b>	no proxy-arp

## Interface Commands

**Parameters** *policy-name* — The export route policy name. Allowed values are any string up to 32 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.

The specified name(s) must already be defined.

### ptp-hw-assist

**Syntax** **[no] ptp-hw-assist**

**Context** config>service>vprn>if

**Description** This command configures the 1588 port based timestamping assist function for the interface. This capability is supported on a specific set of hardware. The command may be blocked if not all hardware has the required level of support.

If the SAP configuration of the interface is removed, the ptp-hw-assist configuration will be removed.

If the IPv4 address configuration of the interface is removed, the ptp-hw-assist configuration will be removed.

Only one interface per physical port can have ptp-hw-assist enabled.

**Default** no ptp-hw-assist

### qos-route-lookup

**Syntax** **qos-route-lookup [source | destination]**  
**no qos-route-lookup**

**Context** config>service>vprn>if  
config>service>vprn>if>ipv6  
config>service>vprn>sub-if>group-interface  
config>service>vprn>sub-if>grp-if>ipv6

**Description** This command enables QoS classification of the ingress IP packets on an interface based on the QoS information associated with routes in the forwarding table.

If the optional **destination** parameter is specified and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network QoS policy.

If the optional **source** parameter is specified and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network QoS policy.

If neither the optional **source** or **destination** parameter is present, then the default is **destination** address matching.

The functionality enabled by the qos-route-lookup command can be applied to IPv4 packets or IPv6 packets on an interface, depending on whether it is present at the interface context (applies to IPv4) or the interface>ipv6 context (applies to IPv6). The ability to specify source address based QoS lookup is not supported for IPv6.

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

<b>Default</b>	destination
<b>Parameters</b>	<p><b>source</b> — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.</p> <p><b>destination</b> — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.</p>

## redundant-interface

<b>Syntax</b>	<b>redundant-interface</b> <i>red-ip-int-name</i> <b>no redundant-interface</b>
<b>Context</b>	config>service>vprn config>service>vprn>sub-if>grp-if
<b>Description</b>	This command configures a redundant interface used for dual homing.
<b>Parameters</b>	<i>red-ip-int-name</i> — Specifies the redundant IP interface name.

## remote-proxy-arp

<b>Syntax</b>	<b>[no] remote-proxy-arp</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	<p>This command enables remote proxy ARP on the interface.</p> <p>Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.</p>
<b>Default</b>	no remote-proxy-arp

## secondary

<b>Syntax</b>	<b>secondary</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>broadcast all-ones</b>   <b>host-ones</b> ] [ <b>igp-inhibit</b> ] <b>no secondary</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> }
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if
<b>Description</b>	This command assigns an secondary IP address/IP subnet/broadcast address format to the interface.
<b>Default</b>	none
<b>Parameters</b>	<p><i>ip-address</i> — The IP address of the IP interface. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).</p> <p><i>mask</i> — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the <i>ip-address</i> from a traditional dotted decimal mask. The <i>mask</i> parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses.</p> <p><i>netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.</p> <p><b>broadcast</b> — The optional <b>broadcast</b> parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is <b>host-ones</b> which indicates a subnet broadcast address. Use this parameter to change the broadcast address to <b>all-ones</b> or revert back to a broadcast address of <b>host-ones</b>.</p> <p>The broadcast format on an IP interface can be specified when the IP address is assigned or changed. This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (<b>all-ones</b>) or the valid subnet broadcast address (<b>host-ones</b>) will be received by the IP interface. (<i>Default: host-ones</i>)</p> <p><b>all-ones</b> — The <b>all-ones</b> keyword following the <b>broadcast</b> parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.</p> <p><b>host-ones</b> — The <b>host-ones</b> keyword following the <b>broadcast</b> parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the <i>ip-address</i> and the <i>mask-length</i> or <i>mask</i> with all the host bits set to binary one. This is the default used by an IP interface. The <b>broadcast</b> parameter within the <b>address</b> command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the <b>broadcast</b> type to <b>host-ones</b> after being changed to <b>all-ones</b>, the <b>address</b> command must be executed with the <b>broadcast</b> parameter defined.</p> <p><b>igp-inhibit</b> — The optional <b>igp-inhibit</b> parameter signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the specified secondary IP interfaces will not be injected and used as passive interfaces and</p>



will not be advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces will not source RIP updates.

## static-arp

<b>Syntax</b>	<b>static-arp</b> <i>ieee-mac-address</i> <i>unnumbered</i> <b>no static-arp</b> <i>unnumbered</i>
<b>Context</b>	config>service>vpn>if config>service>vpn>nw-if
<b>Description</b>	This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface. If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.  The <b>no</b> form of this command removes a static ARP entry.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address for the static ARP in IP address dotted decimal notation.  <i>ieee-mac-address</i> — Specifies the 48-bit MAC address for the static ARP 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.  <i>unnumbered</i> — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

## static-tunnel-redundant-next-hop

<b>Syntax</b>	<b>static-tunnel-redundant-next-hop</b> <i>ip-address</i> <b>no static-tunnel-redundant-next-hop</b>
<b>Context</b>	config>service>vpn>if
<b>Description</b>	This command specifies redundant next-hop address on public or private IPsec interface (with public or private tunnel-sap) for static IPsec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them.  The next-hop address will be resolved in routing table of corresponding service.  The no form of the command removes the address from the interface configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the static ISA tunnel redundant next-hop address.

## strip-label

**Syntax** [no] strip-label

**Context** config>service>vprn>nw-if

**Description** This command forces packets to be stripped of all (max 5) MPLS labels before the packets are handed over for possible filter (PBR) processing.

If the packets do not have an IP header ediatly following the MPLS label stack after the strip, they are discarded. Only MPLS encapsulated IP, IGP shortcuts and VPRN over MPLS packets will be processed.

This command is only supported on:

- Optical ports
- IOM3-XP cards
- Null/Dot1q encaps
- Network ports
- IPv4

The **no** form removes the strip-label command.

In order to associate an interface that is configured with the strip-label parameter with a port, the port must be configured as single-fiber for the command to be valid.

**Default** no strip-label

## secure-nd

**Syntax** [no] secure-nd

**Context** config>service>vprn>if>ipv6

**Description** This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface.

The **no** form of the command reverts to the default and disabled SeND.

## allow-unsecured-msgs

**Syntax** [no] allow-unsecured-msgs

**Context** config>service>vprn>if>secure-nd

**Description** This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default.

The **no** form of the command disables accepting unsecured messages.

## link-local-modifier

**Syntax** **link-local-modifier** *modifier*  
**[no] link-local-modifier**

**Context** config>service>vprn>if>secure-nd

**Description** This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses.

**Parameters** *modifier* — Specifies the modifier in 32 hexadecimal nibbles.

**Values** 0x0–0xFFFFFFFF

## public-key-min-bits

**Syntax** **public-key-min-bits** *bits*  
**[no] public-key-min-bits**

**Context** config>service>vprn>if>secure-nd

**Description** This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA).

**Parameters** *bits* — Specifies the number of bits.

**Values** 512–1024

## security-parameter

**Syntax** **security-parameter** *sec*  
**[no] security-parameter**

**Context** config>service>vprn>if>secure-nd

**Description** This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA).

**Parameters** *sec* — Specifies the security parameter.

**Values** 0–1

## shutdown

**Syntax** **[no] shutdown**

**Context** config>service>vprn>if>secure-nd

**Description** This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

## stale-time

<b>Syntax</b>	<b>stale-time</b> <i>seconds</i> <b>no stale-time</b>
<b>Context</b>	config>service>vprn>ipv6 config>service>vprn>if>ipv6
<b>Description</b>	This command configures the time a neighbor discovery cache entry can remain stale before being removed.  The <b>no</b> form of the command removes the stale-time value.
<b>Default</b>	<b>no stale-time</b>
<b>Parameters</b>	<i>seconds</i> — The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.  <b>Values</b> 60 - 65535

## tcp-mss

<b>Syntax</b>	<b>tcp-mss</b> <i>mss-value</i> <b>no tcp-mss</b>
<b>Context</b>	service>vprn>if service>vprn>if>ipv6
<b>Description</b>	This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.  The <b>no</b> form of the command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value ( $tcp\_mss = ip\_mtu - 40$ ).
<b>Default</b>	<b>no tcp-mss</b>
<b>Parameters</b>	<i>mss-value</i> — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection. <b>Note:</b> $9158 = \max\text{-IP\_MTU} (9198) - 40$  <b>Values</b> 536 - 9158 (IPv4) 1220 - 9138 (IPv6)

## tos-marking-state

<b>Syntax</b>	<b>tos-marking-state</b> { <b>trusted</b>   <b>untrusted</b> } <b>no tos-marking-state</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if config>service>vprn>nw-if

<b>Description</b>	<p>This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted.</p> <p>When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions.</p> <p>Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.</p> <p>The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The <b>save config</b> command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.</p> <p>The <b>no tos-marking-state</b> command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.</p>
<b>Default</b>	trusted
<b>Parameters</b>	<p><b>trusted</b> — The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.</p> <p><b>untrusted</b> — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.</p>

## ipv6

<b>Syntax</b>	<b>[no] ipv6</b>
<b>Context</b>	config>service>vprn>sub-if config>service>vprn>sub-if>group-if
<b>Description</b>	This command configures IPv6 parameters.

## allow-unmatching-prefixes

<b>Syntax</b>	<b>[no] allow-unmatching-prefixes</b>
<b>Context</b>	config>service>vprn>sub-if
<b>Description</b>	This command allows address assignment to PPPoX hosts in cases where the assigned address falls outside the range of the configured subnets below the subscriber interface. Alternatively, if the interface is configured as unnumbered, this command cannot be enabled.
<b>Default</b>	no allow-unmatching-prefixes

## allow-unmatching-prefixes

<b>Syntax</b>	<b>[no] allow-unmatching-prefixes</b>
<b>Context</b>	configure>service>vprn>sub-if>ipv6> configure>service>ies>sub-if>ipv6>
<b>Description</b>	This command allows address assignment to IPv6 hosts in cases where the assigned address or the prefix falls outside of the range of the configured IPv6 subscriber-prefixes under the <b>configure&gt;service&gt;vprn/ies&gt;sub-if&gt;ipv6</b> hierarchy.  Unnumbered PPPoEv6 does not mean that the PPPoEv6 hosts do not have an IPv6 address or prefix assigned. It only means that the IPv6 address range (out of which the address or prefix is assigned to the host) does not have to be known in advance via configuration under the <b>subscriber-interface&gt;ipv6&gt;subscriber-prefixes</b> node
<b>Default</b>	no allow-unmatching-prefixes

## delegated-prefix-length

<b>Syntax</b>	<b>delegated-prefix-length <i>bits</i></b> <b>delegated-prefix-length <i>variable</i></b> <b>no delegated-prefix-length</b>
<b>Context</b>	configure>router>subscriber-interface>ipv6 configure>service>vprn>subscriber-interface>ipv6
<b>Description</b>	This command configures the subscriber-interface level setting for delegated prefix length. The delegated prefix length for a subscriber- interface can be either set to a fixed value that is explicitly configured under the subscriber-interface CLI hierarchy or a variable value that can be obtained from various sources. This command can be changed only when no IPv6 prefixes are configured under the subscriber-interface.
<b>Default</b>	no delegated-prefix-length This means that the delegated prefix length is 64.
<b>Parameters</b>	<i>bits</i> — The delegated prefix length in bits. This value will be applicable to the entire subscriber-interface. In case that the delegated prefix length is also supplied via other means (LUDB, Radius or DHCP Server), such supplied value must match the value configured under the subscriber-interface. Otherwise the prefix instantiation in 7x50 will fail.  <b>Values</b> 48 — 64  <b>variable</b> — The delegated prefix value can be of any length between 48..64. The value itself can vary between the prefixes and it will be provided at the time of prefix instantiation. The order of priority for the source of the delegated prefix length is: <ul style="list-style-type: none"> <li>• LUDB</li> <li>• Radius</li> <li>• DHCPv6 server</li> </ul>

## dhcp6

<b>Syntax</b>	<b>[no] dhcp6</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6
<b>Description</b>	This command allows access to the DHCP6 context within the group interface configuration. Within this context, DHCP6 parameters can be configured.
<b>Default</b>	no dhcp6

## option

<b>Syntax</b>	<b>[no] option</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6
<b>Description</b>	This command enables the context to configure DHCPv6 relay information options. The <b>no</b> form of the command disables DHCPv6 relay information options.

## interface-id

<b>Syntax</b>	<b>interface-id</b> <b>interface-id ascii-tuple</b> <b>interface-id ifindex</b> <b>interface-id sap-id</b> <b>interface-id string</b> <b>no interface-id</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>option
<b>Description</b>	This command enables the sending of interface ID options in the DHCPv6 relay packet. The <b>no</b> form of the command disables the sending of interface ID options in the DHCPv6 relay packet
<b>Parameters</b>	<b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “ ”. <b>ifindex</b> — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command <b>show&gt;router&gt;if&gt;detail</b> .) <b>sap-id</b> — Specifies that the SAP identifier will be used. <b>string</b> — Specifies a string of up to 32 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.

## remote-id

<b>Syntax</b>	<b>[no] remote-id</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>option
<b>Description</b>	This command enables the sending of remote ID option in the DHCPv6 relay packet. The client DHCP Unique Identifier (DUID) is used as the remote ID. The <b>no</b> form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

## proxy-server

<b>Syntax</b>	<b>[no] proxy-server</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6
<b>Description</b>	This command allows access to the DHCP6 proxy server context. Within this context, DHCP6 proxy server parameters of the group interface can be configured
<b>Default</b>	no proxy-server.

## renew-timer

<b>Syntax</b>	<b>renew-timer <i>seconds</i></b> <b>no renew-timer</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	This command configures the renew-timer (T1), the time at which the client contacts the server from which the addresses in the IA_NA or IA_PD were obtained to extend the lifetimes of the addresses or prefixes assigned to the client.
<b>Default</b>	1800
<b>Parameters</b>	<i>seconds</i> — Specifies the time duration relative to the current time, expressed in units of seconds. A value of zero leaves the renew-time at the discretion of the client. <b>Values</b> 0 — 604,800

## rebind-timer

<b>Syntax</b>	<b>rebind-timer <i>seconds</i></b> <b>no rebind-timer</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	This command configures the rebind-timer (T2), the time at which the client contacts any available server to extend the lifetimes of the addresses or prefixes assigned to the client.



<b>Default</b>	2880
<b>Parameters</b>	<i>seconds</i> — T2 is a time duration relative to the current time. A value of zero leaves the rebind-time at the discretion of the client.
<b>Values</b>	0 — 1,209,600

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime</b> [ <i>seconds</i>   <b>infinite</b> ] <b>no preferred-lifetime</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	The preferred lifetime for the IPv6 prefix or address in the option, expressed in units of seconds. When the preferred lifetime expires, any derived addresses are deprecated.
<b>Default</b>	3600
<b>Parameters</b>	<i>seconds</i> — Specifies a decimal time interval in seconds.
<b>Values</b>	600 — 424967295
	<b>infinite</b> — Specifies a 0xffffffff value, Dec = 4294967295

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime</b> [ <i>seconds</i>   <b>infinite</b> ] <b>no valid-lifetime</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	The valid lifetime for the IPv6 prefix or address in the option, expressed in units of seconds.
<b>Default</b>	86,400
<b>Parameters</b>	<i>seconds</i> — Specifies a decimal time interval in seconds.
<b>Values</b>	600 — 424967295
	<b>infinite</b> — Specifies a 0xffffffff value, Dec = 4294967295

## client-applications

<b>Syntax</b>	<b>client-applications</b> [ <b>dhcp</b> ] [ <b>pppoe</b> ] <b>no client-applications</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server
<b>Description</b>	This command configures the client host types to which the DHCP6 proxy server is allowed to assign addresses.

## Interface Commands

- Parameters**    **dhcp** — Specifies IP over Ethernet hosts.  
                  **pppoe** — Specifies PPP over Ethernet hosts.

### router-advertisements

- Syntax**        **[no] router-advertisements**
- Context**        config>service>vprn>sub-if>group-if>ipv6
- Description**    This command enables Router Advertisement transmission on this group interface.
- Default**        router-advertisements

### current-hop-limit

- Syntax**        **current-hop-limit** *hop-count*  
                  **no current-hop-limit**
- Context**        config>service>vprn>sub-if>group-if>ipv6>router-ad
- Description**    This command specifies the hop-limit advertised to hosts in router advertisements.
- Default**        64
- Parameters**    *hop-count* — Specifies the current hop limit (decimal) inserted into router advertisements.
- Values**        0 — 255

### managed-configuration

- Syntax**        **[no] managed-configuration**
- Context**        config>service>vprn>sub-if>group-if>ipv6>router-ad
- Description**    This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address auto-configured using stateless address auto-configuration. See RFC 3315, Dynamic Host Configuration Protocol (DHCP) for IPv6.
- Default**        no managed-configuration

### max-advertisement-interval

- Syntax**        **max-advertisement-interval** *seconds*  
                  **no max-advertisement-interval**
- Context**        config>service>vprn>sub-if>group-if>ipv6>router-ad
- Description**    This command configures the maximum interval between sending router advertisement messages.

<b>Default</b>	1800
<b>Parameters</b>	<i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages.
<b>Values</b>	900 — 1800

## min-advertisement-interval

<b>Syntax</b>	<b>min-advertisement-interval</b> <i>seconds</i> <b>no min-advertisement-interval</b>
<b>Context</b>	config>service>vprn>sub-if>group-if>ipv6>router-ad
<b>Description</b>	This command configures the minimum interval between sending router advertisement messages.
<b>Default</b>	900
<b>Parameters</b>	<i>seconds</i> — Specifies the minimum interval, in seconds, between sending router advertisement messages.
<b>Values</b>	900 — 1350

## mtu

<b>Syntax</b>	<b>mtu</b> <i>bytes</i> <b>no mtu</b>
<b>Context</b>	config>service>vprn>sub-if>group-if>ipv6>router-ad
<b>Description</b>	This command configures the MTU for the nodes to use to send packets on the link.
<b>Default</b>	no mtu
<b>Parameters</b>	<i>bytes</i> — Specifies the MTU for the nodes to use to send packets on the link.
<b>Values</b>	1280 — 9212

## other-stateful-configuration

<b>Syntax</b>	<b>[no] other-stateful-configuration</b>
<b>Context</b>	config>service>vprn>sub-if>group-if>ipv6>router-ad
<b>Description</b>	This command sets the “Other configuration” flag. This flag indicates that DHCPv6 is available for auto-configuration of other (non-address) information such as DNS-related information or information on other servers in the network. See RFC 3736, Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6.
<b>Default</b>	no other-stateful-configuration

## prefix-options

<b>Syntax</b>	<b>[no] prefix-options</b>
<b>Context</b>	config>service>vprn>sub-if>group-if>ipv6>router-ad
<b>Description</b>	This command configures router advertisement parameters for IPv6 prefixes returned via RADIUS Framed-IPv6-Prefix. All prefixes will inherit these configuration parameters.
<b>Default</b>	no prefix-options

## autonomous

<b>Syntax</b>	<b>[no] autonomous</b>
<b>Context</b>	config>services>vprn>sub-if>group-if>ipv6>router-ad>prefix-op
<b>Description</b>	This command specifies whether the prefix can be used for stateless address auto-configuration.
<b>Default</b>	no autonomous

## preferred-lifetime

<b>Syntax</b>	<b>preferred-lifetime</b> [ <i>seconds</i>   <b>infinite</b> ] <b>no preferred-lifetime</b>
<b>Context</b>	config>service>vprn>sub-if>group-if>ipv6>router-ad>prefix-op config>service>vprn>sub-if>group-if>ipv6>dhcp6>proxy-server
<b>Description</b>	This command configures the remaining length of time in seconds that this prefix will continue to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected.
<b>Default</b>	3600
<b>Parameters</b>	<i>seconds</i> — Specifies a decimal time interval in seconds. <b>Values</b> 0-4,294,967,295 <b>infinite</b> — Sepcifies a 0xffffffff value. Dec = 4,294,967,295.

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime</b> [ <i>seconds</i>   <b>infinite</b> ] <b>no valid-lifetime</b>
<b>Context</b>	config>service>vprn>sub-if>group-if>ipv6>router-ad>prefix-op config>service>vprn>sub-if>group-if>ipv6>dhcp6>proxy-server

<b>Description</b>	This command specifies the length of time, in seconds, that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity. The address generated from an invalidated prefix should not appear as the destination or source address of a packet.
<b>Default</b>	86400
<b>Parameters</b>	<i>seconds</i> — Specifies a decimal time interval in seconds.
	<b>Values</b> 0-4,294,967,295
	<b>infinite</b> — Specifies a 0xffffffff value. Dec = 4,294,967,295.

## reachable-time

<b>Syntax</b>	<b>reachable-time</b> <i>milliseconds</i> <b>no reachable-time</b>
<b>Context</b>	config>services>vprn>sub-if>group-if>ipv6>router-ad
<b>Description</b>	This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.
<b>Default</b>	no reachable-time
<b>Parameters</b>	<i>milliseconds</i> — The length of time the router should be considered reachable for default router selection.
	<b>Values</b> 0-3,600,000

## retransmit-time

<b>Syntax</b>	<b>retransmit-time</b> <i>milliseconds</i> <b>no retransmit-time</b>
<b>Context</b>	config>services>vprn>sub-if>group-if>ipv6>router-ad
<b>Description</b>	This command configures the retransmission frequency of neighbor solicitation messages.
<b>Default</b>	no retransmit-time
<b>Parameters</b>	<i>milliseconds</i> — Specifies how often the retransmission should occur.
	<b>Values</b> 0-1,800,000

## router-lifetime

<b>Syntax</b>	<b>router-lifetime</b> <i>seconds</i> <b>router-lifetime no-default-router</b> <b>no router-lifetime</b>
<b>Context</b>	config>services>vprn>sub-if>group-if>ipv6>router-ad

## Interface Commands

<b>Description</b>	This command sets the router lifetime. A value of zero indicates this router should not be used by hosts as a default router.
<b>Default</b>	4500
<b>Parameters</b>	<i>seconds</i> — Specifies how long this router is valid for default router selection. <b>Values</b> 2700-9000

## renew-timer

<b>Syntax</b>	<b>renew-timer</b> <i>seconds</i> <b>no renew-timer</b>
<b>Context</b>	config>services>vprn>sub-if>group-if>ipv6>dhcpv6
<b>Description</b>	This command configures the renew-timer (T1). The time at which the client contacts the server from the addresses in the IA_NA or IA_PD were obtained to extend the lifetimes of the addresses or prefixes assigned to the client.
<b>Default</b>	1800
<b>Parameters</b>	<i>seconds</i> — Time duration relative to the current time expressed in units of seconds. A value of zero (0) leaves the renew-time at the discretion of the client. <b>Values</b> 0-604,800

## rebind-timer

<b>Syntax</b>	<b>rebind-timer</b> <i>seconds</i> <b>no rebind-timer</b>
<b>Context</b>	config>services>vprn>sub-if>group-if>ipv6>dhcpv6
<b>Description</b>	This command configures the rebind-timer (T2), the time at which the client contacts any available server to extend the lifetimes of the addresses or prefixes assigned to the client.
<b>Default</b>	2880
<b>Parameters</b>	<i>seconds</i> — T2 is a time duration relative to the current time expressed in units of seconds. A value of zero (0) leaves the rebind-time at the discretion of the client. <b>Values</b> 0-1,209,600

## delegated-prefix-length

<b>Syntax</b>	[no] <b>delegated-prefix-length</b> <i>prefix-length</i>
<b>Context</b>	config>service>vprn>sub-if>ipv6

<b>Description</b>	This command defines the prefix-length used for all DHCPv6 prefix delegations on this subscriber interface.
<b>Parameters</b>	<i>prefix-length</i> — Specifies the prefix length in use on this subscriber interface for DHCPv6 IA_PD.
<b>Values</b>	48..64
<b>Default</b>	64

## subscriber-prefixes

<b>Syntax</b>	<b>subscriber-prefixes</b>
<b>Context</b>	config>service>vprn>sub-if>ipv6
<b>Description</b>	This command specifies aggregate off-link subscriber prefixes associated with this subscriber interface. Individual prefixes are specified under the prefix context list aggregate routes in which the next-hop is indirect via the subscriber interface.

## prefix

<b>Syntax</b>	<b>prefix</b> <i>ipv6-address/prefix-length</i> [ <b>pd</b> ] [ <b>wan-host</b> ] <b>no prefix</b> <i>ipv6-address/prefix-length</i>
<b>Context</b>	config>service>vprn>sub-if>ipv6>sub-prefixes
<b>Description</b>	This command allows a list of prefixes(using the prefix command multiple times) to be routed to hosts associated with this subscriber interface. Each prefix will be represented in the associated FIB with a reference to the subscriber interface. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).
<b>Parameters</b>	<i>ipv6-address</i> — Specifies the 128-bit IPv6 address. <b>Values</b> 128-bit hexadecimal IPv6 address in compressed form <i>prefix-length</i> — Specifies the length of any associated aggregate prefix. <b>Values</b> 32-63 <b>pd</b> — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation. <b>wan-host</b> — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

## private-retail-subnets

<b>Syntax</b>	<b>[no] private-retail-subnets</b>
<b>Context</b>	config>service>vprn>sub-if

## Interface Commands

**Description** This command controls the export of subnets to the forwarding service. When this attribute is configured, subnets defined on this retail subscriber interface will no longer be exported to the associated wholesale VPRN and will remain private to the retail VPRN. This is useful in a PPPoE business service context as it allows retail services to use overlapping IP address spaces even if these services are associated with the same wholesale service.

PPPoE sessions are actually terminated in the retail service although their traffic transits on a SAP belonging to the wholesale service. This configuration is incompatible, however, with IPoE host management (DHCP, static-host and ARP-host) as these host types require that the retail subnets are exported to the wholesale VPRN. Thus, if PPPoE sessions need to coexist with IPoE hosts, this attribute should not be configured on this retail interface.

This command will fail if the subscriber interface is not associated with a wholesale service.

If the retail VPRN is of the type **hub**, this attribute is mandatory. Then, it will be enabled by default and it will not be possible to deconfigure it.

## unnumbered

**Syntax** **unnumbered** [*ip-int-name* | *ip-address*]  
**no unnumbered**

**Context** config>service>vprn>if  
config>service>vprn>nw-if

**Description** This command configures the interface as an unnumbered interface. Unnumbered IP interface is supported on a Sonet/SDH access port with the PPP, ATM, or Frame Relay encapsulation. It is not supported on a TDM port or channel.

**Parameters** *ip-int-name* — Specifies the name of an IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.  
*ip-address* — Specifies an IP address.

## qos

**Syntax** **qos** *network-policy-id* **port-redirect-group** *queue-group-name* **egress-instance** *instance-id* **fp-redirect-group** *queue-group-name* **ingress-instance** *instance-id*  
**no qos**

**Context** config>service>vprn>nw-if

**Description** This command associates a network Quality of Service (QoS) policy with a network IP interface. Only one network QoS policy can be associated with an IP interface at one time. Attempts to associate a second QoS policy return an error.

Associating a network QoS policy with a network interface is useful for the following purposes:

- To apply classification rules for determining the forwarding-class and profile of ingress packets on the interface.



- To associate ingress packets on the interface with a queue-group instance applied to the ingress context of the interface's forwarding plane (FP). (This is only applicable to interfaces on IOM3 and later cards.) The referenced ingress queue-group instance may have policers defined in order to rate limit ingress traffic on a per-forwarding class (and forwarding type: unicast vs. multicast) basis.
- To perform 802.1p, DSCP, IP precedence and/or MPLS EXP re-marking of egress packets on the interface.
- To associate egress packets on the interface with a queue-group instance applied to the egress context of the interface's port. The referenced egress queue-group instance may have policers and/or queues defined in order to rate limit egress traffic on a per-forwarding class basis.

The **no** form of the command removes the network QoS policy association from the network IP interface, and the QoS policy reverts to the default.

**Default** **no qos**

**Parameters** *network-policy-id* — An existing network policy ID to associate with the IP interface.

**Values** 1 — 65535

**port-redirect-group** *queue-group-name* — This optional parameter specifies the egress queue-group used for all egress forwarding-class redirections specified within the network QoS policy ID. The specified *queue-group-name* must exist as an egress queue group applied to the egress context of the port associated with the IP interface.

**egress-instance** *instance-id* — Since multiple instances of the same egress queue-group can be applied to the same port this optional parameter is used to specify which particular instance to associate with this particular network IP interface.

**Values** 1 — 16384

**fp-redirect-group** *queue-group-name* — This optional parameter specifies the ingress queue-group used for all ingress forwarding-class redirections specified within the network QoS policy ID. The specified *queue-group-name* must exist as an ingress queue group applied to the ingress context of the forwarding plane associated with the IP interface.

**ingress-instance** *instance-id* — Since multiple instances of the same ingress queue-group can be applied to the same forwarding plane this parameter is required to specify which particular instance to associate with this particular network IP interface.

**Values** 1 — 16384

## urpf-check

**Syntax** **[no] urpf-check**

**Context** config>service>vprn>if  
config>service>vprn>nw-if  
config>service>vprn>interface>ipv6  
config>service>vprn>sub-if>grp-if

**Description** This command enables unicast RPF (uRPF) check on this interface.

The **no** form of the command disables unicast RPF (uRPF) Check on this interface.

## Interface Commands

**Default** disabled

### vas-if-type

**Syntax** **vas-if-type** {**to-from-access** | **to-from-network** | **to-from-both**}  
**no vas-if-type**

**Context** config>service>vprn>if

**Description** This command configures the type of a Value Added Service (VAS) facing interface.  
The **no** form of the command removes VAS interface type configuration.

**Default** **no vas-if-type**

**Parameters** **to-from-access** — Used when two separate (to-from-access and to-from-network) interfaces are used for Value Added Service (VAS) connectivity. For service chaining, traffic arriving from access interfaces (upstream) is redirected to a PBR target reachable over this interface for upstream VAS processing. Downstream traffic after VAS processing must arrive on this interface, so the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.

**to-from-network** — Used when two separate (to-from-access and to-from-network) interfaces are used for Value Added Service (VAS) connectivity. For service chaining, traffic arriving from network interfaces (downstream) is redirected to a PBR target reachable over this interface for downstream VAS processing. Upstream traffic after VAS processing must arrive on this interface, so regular routing can be applied.

**to-from-both** — Used when a single interface is used for Value Added Service (VAS) connectivity (no local-to-local traffic). For service chaining, both traffic arriving from access and from network is redirected to a PBR target reachable over this interface for upstream/downstream VAS processing. Traffic after VAS processing must arrive on this interface, so the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.

### mode

**Syntax** **mode** {**strict** | **loose** | **strict-no-ecmp**}  
**no mode**

**Context** config>service>vprn>if>urfp-check  
config>service>vprn>nw-if>urfp-check  
config>service>vprn>sub-if>grp-if>urpf-check

**Description** This command specifies the mode of unicast RPF check.  
The **no** form of the command reverts to the default (strict) mode.

**Default** strict

**Parameters** **strict** — When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

**loose** — In **loose** mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. This object is valid only when **urpf-check** is enabled.

**strict-no-ecmp** — When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF.

---

## DHCP Commands

### dhcp

<b>Syntax</b>	<b>dhcp</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>nw-if config>service>vprn>subscriber-interface config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables the context to configure DHCP parameters.

### client-applications

<b>Syntax</b>	<b>client-applications dhcp</b> <b>client-applications pppoe</b> <b>client-applications dhcp pppoe</b> <b>no client-applications</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command enables the clients that will try to contact the DHCP server(s). The <b>no</b> form of the command removes the server client type from the configuration.
<b>Parameters</b>	<b>dhcp</b> — Specifies that the DHCP relay will forward requests to the DHCP server(s). <b>pppoe</b> — Specifies that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s)ly assigned to PPPoE node.

### action

<b>Syntax</b>	<b>action {replace   drop   keep}</b> <b>no action</b>
<b>Context</b>	config>service>vprn>if>dhcp>option config>service>vprn>nw-if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
<b>Description</b>	This command configures the processing required when the SR-Series router receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet. The <b>no</b> form of this command returns the system to the default value.
<b>Default</b>	Per RFC 3046, <i>DHCP Relay Agent Information Option</i> , section 2.1.1, <i>Reforwarded DHCP requests</i> , the default is to keep the existing information intact. The exception to this is if the giaddr of the

received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

**Parameters** **replace** — In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

**drop** — The packet is dropped, and an error is logged.

**keep** — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

## circuit-id

<b>Syntax</b>	<b>circuit-id</b> [ascii-tuple   ifindex   sap-id   vlan-ascii-tuple] <b>no circuit-id</b>
<b>Context</b>	config>service>vprn>if>dhcp>option config>service>vprn>nw-if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
<b>Description</b>	When enabled, the router sends the interface index (If Index) in the <b>circuit-id</b> suboption of the DHCP packet. The If Index of a router interface can be displayed using the command <b>show&gt;router&gt;interface&gt;detail</b> . This option specifies data that must be unique to the router that is relaying the circuit.  If disabled, the <b>circuit-id</b> suboption of the DHCP packet will be left empty.  The <b>no</b> form of this command returns the system to the default.
<b>Default</b>	circuit-id
<b>Parameters</b>	<b>ascii-tuple</b> — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “[”].  <b>ifindex</b> — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command <b>show&gt;router&gt;interface&gt;detail</b> .  <b>sap-id</b> — Specifies that the SAP ID will be used.  <b>vlan-ascii-tuple</b> — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

## filter

<b>Syntax</b>	<b>filter</b> <i>filter-id</i> <b>no filter</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command configures the DHCP filter for this interface.
<b>Parameters</b>	<i>filter-id</i> — Specifies the filter policy. The filter ID must already exist within the created IP filters. <b>Values</b> 1 — 65535

## gi-address

<b>Syntax</b>	<b>gi-address</b> <i>ip-address</i> [ <i>src-ip-addr</i> ] <b>no gi-address</b>
<b>Context</b>	config>service>vprn>if>dhcp config>service>vprn>nw-if>dhcp config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.
<b>Default</b>	no gi-address
<b>Parameters</b>	<i>ip-address</i> — Specifies the host IP address to be used for DHCP relay packets. <i>src-ip-address</i> — Specifies the source IP address to be used for DHCP relay packets.

## lease-populate

<b>Syntax</b>	<b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>route-populate</b> [ <i>pd</i> ] <i>na</i> [ <i>ta</i> ] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>route-populate</b> <i>pd</i> [ <i>na</i> ] [ <i>ta</i> ] [ <b>exclude</b> ] <b>lease-populate</b> [ <i>nbr-of-leases</i> ] <b>route-populate</b> [ <i>pd</i> ] [ <i>na</i> ] <i>ta</i> <b>no lease-populate</b>
<b>Context</b>	config>service>vprn>if>ipv6 config>service>vprn>if>ipv6>dhcp-relay
<b>Description</b>	This command specifies the maximum number of DHCPv6 lease states allocated by the DHCPv6 relay function, allowed on this interface. Optionally, by specifying “route-populate” parameter, system could: <ul style="list-style-type: none"> <li>• Create routes based on the IA_PD/IA_NA/IA_TA prefix option in relay-reply message.</li> </ul>

- Create black hole routes based on `OPTION_PD_EXCLUDE` in `IA_PD` in relay-reply message.

These routes could be redistributed into IGP/BGP by using route-policy, following protocol types that could be used in “from protocol”:

- `dhcpv6-pd`
- `dhcpv6-na`
- `dhcpv6-ta`
- `dhcpv6-pd-excl`

**Parameters** *nbr-of-entries* — Defines the number lease state table entries allowed for this interface. If this parameter is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCPv6 ACK messages are discarded.

**Values** 1 — 8000

**route-populate** —

**Values** `pd/na/ta` — Create route based on specified option.

`exclude` — Create blackhole route based on `OPTION_PD_EXCLUDE`.

## neighbor-resolution

**Syntax** `[no] neighbor-resolution`

**Context** `config>service>vprn>if>ipv6>dhcp6-relay`

**Description** This command enables neighbor resolution with DHCPv6 relay. The **no** form of the command disables neighbor resolution.

## match-circuit-id

**Syntax** `[no] match-circuit-id`

**Context** `config>service>vprn>sub-if>grp-if>dhcp`

**Description** This command enables Option 82 circuit ID on relayed DHCP packet matching. For routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked.

When a response is received from the server the virtual router ID, transaction ID, and client hardware MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.

When the **match-circuit-id** command is enabled we use this as part of the key to guarantee correctness in our lookup. This is really only needed when we are dealing with an IP aware DSLAM that proxies the client hardware MAC address.

## Interface Commands

**Default** no match-circuit-id

### option

**Syntax** [no] option

**Context** config>service>vprn>if>dhcp  
config>service>vprn>nw-if>dhcp  
config>service>vprn>sub-if>dhcp  
config>service>vprn>sub-if>grp-if>dhcp

**Description** This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.  
The **no** form of this command returns the system to the default.

**Default** no option

### copy-82

**Syntax** [no] copy-82

**Context** config>service>vprn>nw-if>dhcp>option

**Description** This command enables the copy-82 option.  
The **no** form of the command disables the option.

### remote-id

**Syntax** remote-id [mac | string *string*]  
no remote-id

**Context** config>service>vprn>sub-if>grp-if>dhcp>option  
config>service>vprn>nw-if>dhcp>option

**Description** When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** suboption of the DHCP packet. This command identifies the host at the other end of the circuit. If disabled, the **remote-id** suboption of the DHCP packet will be left empty.  
The **no** form of this command returns the system to the default.

**Default** remote-id

**Parameters** **mac** — This keyword specifies the MAC address of the remote end is encoded in the suboption.  
**string *string*** — Specifies the remote-id.



## vendor-specific-option

<b>Syntax</b>	<b>[no] vendor-specific-option</b>
<b>Context</b>	config>service>vprn>if>dhcp>option config>service>vprn>nw-if>dhcp>option config>service>vprn>sub-if>grp-if>dhcp>option
<b>Description</b>	This command configures the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

## client-mac-address

<b>Syntax</b>	<b>[no] client-mac-address</b>
<b>Context</b>	config>service>vprn>if>dhcp>option config>service>vprn>nw-if>dhcp>option config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.  The <b>no</b> form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

## pool-name

<b>Syntax</b>	<b>[no] pool-name</b>
<b>Context</b>	config>service>vprn>if>dhcp>option
<b>Description</b>	This command enables the sending of the pool name in the Alcatel vendor-specific suboption of the DHCP relay packet.  The <b>no</b> form of the command disables the feature.

## if-name

<b>Syntax</b>	<b>[no] if-name</b>
<b>Context</b>	config>service>vprn>nw-if>dhcp>option
<b>Description</b>	This command enables the sending of the interface name in the Alcatel vendor specific suboption of the DHCP relay packet  The <b>no</b> form of the command disables the sending.

## port-id

<b>Syntax</b>	<b>[no] port-id</b>
<b>Context</b>	config>service>vprn>nw-if>dhcp>option
<b>Description</b>	This command enables sending of the port-id in the Alcatel vendor specific suboption of the DHCP relay packet The <b>no</b> form of the command disables the sending.

## sap-id

<b>Syntax</b>	<b>[no] sap-id</b>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet. The <b>no</b> form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

## service-id

<b>Syntax</b>	<b>[no] service-id</b>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command enables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet. The <b>no</b> form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

## string

<b>Syntax</b>	<b>[no] string <i>text</i></b>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command specifies the vendor specific suboption string of the DHCP relay packet. The <b>no</b> form of the command returns the default value.
<b>Parameters</b>	<i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“”).

## system-id

<b>Syntax</b>	<b>[no] system-id</b>
<b>Context</b>	config>service>vprn>if>dhcp>option>vendor config>service>vprn>nw-if>dhcp>option>vendor config>service>vprn>sub-if>grp-if>dhcp>option>vendor
<b>Description</b>	This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.
<b>Default</b>	None

## proxy-server

<b>Syntax</b>	<b>proxy-server</b>
<b>Context</b>	config>service>if>dhcp config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command configures the DHCP proxy server.

## emulated-server

<b>Syntax</b>	<b>emulated-server <i>ip-address</i></b> <b>no emulated-server</b>
<b>Context</b>	config>service>vprn>if>dhcp>proxy config>service>vprn>sub-if>grp-if>dhcp>proxy-server
<b>Description</b>	This command configures the IP address to be used as the DHCP server address in the context of this service. Typically, the configured address should be in the context of the subnet.  The <b>no</b> form of this command reverts to the default setting. The local proxy server will not become operational without a specified emulated server address.
<b>Parameters</b>	<i>ip-address</i> — Specifies the emulated server address.  <b>Default</b> Note that for a retail interface, the default is the local interface.

## lease-time

<b>Syntax</b>	<b>lease-time [days <i>days</i>] [hrs <i>hours</i>] [min <i>minutes</i>] [sec <i>seconds</i>] [radius-override]</b> <b>no lease-time</b>
<b>Context</b>	config>service>vprn>if>dhcp>proxy config>service>vprn>sub-if>grp-if>dhcp>proxy-server

## Interface Commands

<b>Description</b>	<p>This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.</p> <p>The no form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.</p>
<b>Default</b>	7 days 0 hours 0 seconds
<b>Parameters</b>	<p><b>radius-override</b> — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients.</p> <p><i>days</i> — Specifies the number of days that the given IP address is valid.</p> <p><b>Values</b> 0 — 3650</p> <p><i>hours</i> — Specifies the number of hours that the given IP address is valid.</p> <p><b>Values</b> 0 — 23</p> <p><i>minutes</i> — Specifies the number of minutes that the given IP address is valid.</p> <p><b>Values</b> 0 — 59</p> <p><i>seconds</i> — Specifies the number of seconds that the given IP address is valid.</p> <p><b>Values</b> 0 — 59</p>

## server

<b>Syntax</b>	<b>server</b> <i>server1</i> [ <i>server2</i> ...(up to 8 max)]
<b>Context</b>	config>service>vprn>if>dhcp config>service>vprn>nw-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	<p>This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. There can be a maximum of 8 DHCP servers configured.</p> <p>The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to "flood". This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream then can perform the full L3 DHCP relay function.</p>
<b>Default</b>	no server
<b>Parameters</b>	<i>server</i> — Specifies the DHCP server IP address.

## python-policy

<b>Syntax</b>	<b>python-policy</b> <i>name</i> <b>no python-policy</b>
<b>Context</b>	config>service>vprn>if>dhcp
<b>Description</b>	This command specifies a python policy to be used for DHCPv4. Python policies are configured in the <b>config&gt;python&gt; python-policy</b> <i>name</i> context.
<b>Parameters</b>	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

## python-policy

<b>Syntax</b>	<b>python-policy</b> <i>name</i> <b>no python-policy</b>
<b>Context</b>	config>service>vprn>if>dhcp6-relay
<b>Description</b>	This command specifies a python policy to be used for DHCPv6 relay. Python policies are configured in the <b>config&gt;python&gt; python-policy</b> <i>name</i> context.
<b>Parameters</b>	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

## relay-plain-bootp

<b>Syntax</b>	<b>[no] relay-plain-bootp</b>
<b>Context</b>	config>service>vprn>if>dhcp
<b>Description</b>	This command enables the relaying of plain BOOTP packets. The <b>no</b> form of the command disables the relaying of plain BOOTP packets.

## relay-unicast-msg

<b>Syntax</b>	<b>relay-unicast-msg</b> [ <b>release-update-src-ip</b> ] <b>no relay-unicast-msg</b>
<b>Context</b>	config>service>vprn>sub-if>dhcp config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	Relay unicast client DHCPv4 request (renew) messages. In the upstream direction: update the source-ip address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers). In the downstream direction: remove the gi-address and update the destination IP address to the value of the yiaddr (your IP address) field.

## Interface Commands

By default, unicast DHCPv4 release messages are forwarded transparently.

Additionally when the optional flag “relay-unicast-msg” is enabled, then the gi address and source IP address of relayed DHCPv4 messages can be configured to any local configured IP address in the same routing instance.

**Default** no relay-unicast-msg

**Parameters** **release-update-src-ip** — Updates the source IP address with the value used for relayed DHCPv4 messages

## snoop

**Syntax** [no] snoop

**Context** config>service>vprn>nw-if>dhcp

**Description** This command enables snooping of DHCP packets on this interface.  
The **no** form of the command disables snooping.

## trusted

**Syntax** [no] trusted

**Context** config>service>vprn>if>dhcp  
config>service>vprn>nw-if>dhcp  
config>service>vprn>sub-if>grp-if>dhcp

**Description** According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and which contains a Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit.

If trusted mode is enabled on an IP interface, the relay agent (the SR-Series) will modify the request's giaddr to be equal to the ingress interface and forward the request.

Note that this behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the relay agent (action = "replace"), the original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option.

The **no** form of this command returns the system to the default.

**Default** not enabled

## egress

**Syntax** egress

**Context** config>service>vprn>nw-if

**Description** This command enables the context to configure egress network filter policies for the interface.

## use-arp

**Syntax** `[no] use-arp`

**Context** `config>service>vprn>if>dhcp`

**Description** This command enables the use of ARP to determine the destination hardware address. The **no** form of the command disables the use of ARP to determine the destination hardware address

## user-db

**Syntax** `user-db local-user-db-name`  
`no user-db`

**Context** `config>service>vprn>sub-if>grp-if>dhcp`

**Description** This command configures the local user database to use for authentication. The **no** form of the command removes the value from the configuration.

**Default** `no user-db`

**Parameters** *local-user-db-name* — Specifies the local user database to use for authentication.

## dynamic-tunnel-redundant-next-hop

**Syntax** `dynamic-tunnel-redundant-next-hop ip-address`  
`no dynamic-tunnel-redundant-next-hop`

**Context** `config>service>vprn>if`

**Description** This command specifies redundant next-hop address on public or private IPsec interface (with public or private tunnel-sap) for dynamic IPsec tunnel. The specified next-hop address will be used by standby node to shunt traffic to master in case of it receives them.

The next-hop address will be resolved in routing table of corresponding service.

**Default** `none`

**Description** *ip-address* — Specifies the dynamic ISA tunnel redundant next-hop address.

## egr-ip-load-balancing

<b>Syntax</b>	<b>egr-ip-load-balancing</b> {source   destination   inner-ip} <b>no egr-ip-load-balancing</b>
<b>Context</b>	config>service>vprn>if>load-balancing config>service>vprn>if>nw-if>load-balancing
<b>Description</b>	This command specifies whether to include source address or destination address or both in LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled the command applies also to inclusion of source/destination port in the hash inputs.  The <b>no</b> form of this command includes both source and destination parameters.
<b>Default</b>	no egr-ip-load-balancing
<b>Parameters</b>	<b>source</b> — Specifies using source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port.  <b>destination</b> — Specifies using destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.  <b>inner-ip</b> — Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.

## enable-ingress-stats

<b>Syntax</b>	<b>[no] enable-ingress-stats</b>
<b>Context</b>	config>router>interface config>service>ies >interface config>service>vprn>interface config>service>ies>sub-if>grp-if config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.  If enabled, then the following statistics are collected:  IPv4 offered packets IPv4 offered octets IPv6 offered packets IPv6 offered octets  Note that octet statistics for IPv4 and IPv6 bytes at IP interfaces include the layer 2 frame overhead.
<b>Default</b>	no enable-ingress-stats



## enable-mac-accounting

<b>Syntax</b>	<b>[no] enable-mac-accounting</b>
<b>Context</b>	config>service>vprn>if
<b>Description</b>	This command enables MAC accounting functionality on this interface. The <b>no</b> form of the command disables MAC accounting functionality on this interface.

## host-connectivity-verify

<b>Syntax</b>	<b>host-connectivity-verify [source {vrrp   interface}] [interval <i>interval</i>] [action {remove   alarm}] [timeout <i>retry-timeout</i>] [retry-count <i>count</i>]</b> <b>host-connectivity-verify [interval <i>interval</i>] [action {remove   alarm}] [timeout <i>retry-timeout</i>] [retry-count <i>count</i>] [family <i>family</i>]</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>sub-if>grp-if>dhcp
<b>Description</b>	This command enables subscriber host connectivity verification on a given SAP within a service.  This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.
<b>Default</b>	no host-connectivity-verify
<b>Parameters</b>	<p><b>source {vrrp   interface}</b> — Specifies the source to be used for generation of subscriber host connectivity verification packets. The <b>interface</b> keyword forces the use of the interface mac and ip addresses. Note that there are up to 16 possible subnets on a given interface, therefore subscriber host connectivity verification tool will use always an address of the subnet to which the given host is pertaining. In case of group-interfaces. one of the parent subscriber-interface subnets (depending on host's address) will be used.</p> <p><b>interval <i>interval</i></b> — The interval, expressed in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.</p> <p><b>Values</b> 1— 6000 (Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.)</p> <p><b>action {remove   alarm}</b> — Defines the action taken on a subscriber host connectivity verification failure for a given host. The <b>remove</b> keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, etc.). DHCP-RELEASE will be signaled to corresponding DHCP server. Static hosts will never be removed. The <b>alarm</b> keyword raises an alarm indicating that the host is disconnected.</p> <p><b>timeout <i>retry-timeout</i></b> — Specifies the timeout in seconds between consecutive retries of subscriber host connectivity verification checks, in case the host does not respond.</p> <p><b>Values</b> 10— 60 seconds</p>

**retry-count** *count* — Specifies the number of retries that will be carried out before a subscriber host is considered to have failed the SHCV check.

**Values** 2— 29

**family** *family* — Indicates the IP address family for which subscriber host connectivity verification checks will be enabled. It can be set to ipv4 or ipv6 only, or both.

**Values** 2— 29

## Interface ICMP Commands

### icmp

<b>Syntax</b>	<b>icmp</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

### mask-reply

<b>Syntax</b>	<b>[no] mask-reply</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.  If a local node sends an ICMP mask request to the router interface, the <b>mask-reply</b> command configures the router interface to reply to the request.  By default, the router instance will reply to mask requests.  The <b>no</b> form of this command disables replies to ICMP mask requests on the router interface.
<b>Default</b>	<b>mask-reply</b> — Reply to ICMP mask requests.

### redirects

<b>Syntax</b>	<b>redirects</b> [ <i>number seconds</i> ] <b>no redirects</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.  When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.  The <b>redirects</b> command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters

## Interface Commands

by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

**Default** **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds.

**Parameters** *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

**Values** 10 — 1000

*seconds* — The time frame in seconds used to limit the *seconds* of ICMP redirect messages that can be issued.

**Values** 1 — 60

## tll-expired

**Syntax** **tll-expired** *number seconds*  
**no tll-expired**

**Context** config>service>vprn>if>icmp  
config>service>vprn>sub-if>grp-if>icmp  
config>service>vprn>nw-if>icmp

**Description** Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

**Default** tll-expired 100 10

**Parameters** *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

**Values** 10 — 1000

*seconds* — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

**Values** 1 — 60

## unreachables

<b>Syntax</b>	<b>unreachables</b> [ <i>number seconds</i> ] <b>no unreachable</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	<p>This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.</p> <p>The <b>unreachables</b> command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 10 second time interval.</p> <p>The <b>no</b> form of this command disables the generation of icmp destination unreachable messages on the router interface.</p>
<b>Default</b>	<b>unreachables 100 10</b>
<b>Parameters</b>	<p><i>number</i> — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 — 1000</p> <p><i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP unreachable messages that can be issued.</p> <p><b>Values</b> 1 — 60</p>

---

## Interface ICMP Commands

### icmp

<b>Syntax</b>	<b>icmp</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

### mask-reply

<b>Syntax</b>	<b>[no] mask-reply</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.  If a local node sends an ICMP mask request to the router interface, the <b>mask-reply</b> command configures the router interface to reply to the request.  By default, the router instance will reply to mask requests.  The <b>no</b> form of this command disables replies to ICMP mask requests on the router interface.
<b>Default</b>	<b>mask-reply</b> — Reply to ICMP mask requests.

### redirects

<b>Syntax</b>	<b>redirects</b> [ <i>number seconds</i> ] <b>no redirects</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.  When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.  The <b>redirects</b> command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters

by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

<b>Default</b>	<b>redirects 100 10</b> — Maximum of 100 redirect messages in 10 seconds.
<b>Parameters</b>	<i>number</i> — The maximum number of ICMP redirect messages to send. This parameter must be specified with the <i>seconds</i> parameter.
	<b>Values</b> 10 — 1000
	<i>seconds</i> — The time frame in seconds used to limit the <i>seconds</i> of ICMP redirect messages that can be issued.
	<b>Values</b> 1 — 60

## tll-expired

<b>Syntax</b>	<b>tll-expired</b> <i>number seconds</i> <b>no tll-expired</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.  By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.  The <b>no</b> form of this command disables the limiting the rate of TTL expired messages on the router interface.
<b>Default</b>	tll-expired 100 10
<b>Parameters</b>	<i>number</i> — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the <i>seconds</i> parameter.
	<b>Values</b> 10 — 1000
	<i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP TTL expired messages that can be issued, expressed as a decimal integer.
	<b>Values</b> 1 — 60

## unreachables

<b>Syntax</b>	<b>unreachables</b> [ <i>number seconds</i> ] <b>no unreachables</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	<p>This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.</p> <p>The <b>unreachables</b> command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 10 second time interval.</p> <p>The <b>no</b> form of this command disables the generation of icmp destination unreachable messages on the router interface.</p>
<b>Default</b>	<b>unreachables 100 10</b>
<b>Parameters</b>	<p><i>number</i> — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the <i>seconds</i> parameter.</p> <p><b>Values</b> 10 — 1000</p> <p><i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP unreachable messages that can be issued.</p> <p><b>Values</b> 1 — 60</p>





---

## Router Advertisement Commands

### router-advertisement

<b>Syntax</b>	<b>[no] router-advertisement</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.</p> <p>The <b>no</b> form of the command disables all IPv6 interface. However, the <b>no interface <i>interface-name</i></b> command disables a specific interface.</p>
<b>Default</b>	disabled

### dns-options

<b>Syntax</b>	<b>[no] dns-options</b>
<b>Context</b>	config>service>vprn>router-advertisement config>service>vprn>router-advertisement>interface
<b>Description</b>	<p>This command enables the context for configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.</p> <p>When specified at the router-advertisement level in the routing context, this command allows configuration of service-wide parameters. These can then be inherited at the interface level by specifying the <b>config&gt;service&gt;vprn&gt;router-advertisement&gt;interface&gt;dns-options&gt;include-dns</b> command.</p> <p>The <b>no</b> form of the command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.</p>
<b>Default</b>	disabled

### dns-servers

<b>Syntax</b>	<b>server <i>ipv6-address</i></b> <b>no server</b>
<b>Context</b>	config>service>vprn>router-advertisement>dns-options config>service>vprn>router-advertisement>interface>dns-options
<b>Description</b>	<p>This command specifies the IPv6 DNS servers to include in the RDNSS option in Router Advertisements. When specified at the router advertisement level this applies to all interfaces that have <b>include-dns</b> enabled, unless the interfaces have more specific <b>dns-options</b> configured.</p>
<b>Default</b>	none

**Parameters** *ipv6-address* — Specify the IPv6 address of the DNS server(s), up to 4 max. Specified as eight 16-bit hexadecimal pieces.

## include-dns

**Syntax** `[no] include-dns`

**Context** `config>service>vprn>router-advertisement>interface>dns-options`

**Description** This command enables the Recursive DNS Server (RDNSS) Option in router advertisements. This must be enabled for each interface on which the RDNSS option is required in router advertisement messages.

The **no** form of the command disables the RDNSS option in router advertisements.

**Default** disabled

## rdnss-lifetime

**Syntax** `rdnss-lifetime {seconds | infinite}`  
`no rdnss-lifetime`

**Context** `config>service>vprn>router-advertisement>dns-options`  
`config>service>vprn>router-advertisement>interface>dns-options`

**Description** This command specifies the maximum time that the RDNSS address may be used for name resolution by the client. The RDNSS Lifetime must be no more than twice MaxRtrAdvLifetime with a maximum of 3600 seconds.

**Default** infinite

**Parameters** **infinite** — specifies an infinite RDNSS lifetime.

*seconds* — Specifies the time in seconds.

**Values** 4— 3600

## interface

**Syntax** `[no] interface ip-int-name`

**Context** `config>service>vprn>router-advertisement`

**Description** This command configures router advertisement properties on a specific interface. The interface must already exist in the `config>router>interface` context.

**Default** No interfaces are configured by default.

**Parameters** *ip-int-name* — Specify the interface name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

### current-hop-limit

<b>Syntax</b>	<b>current-hop-limit</b> <i>number</i> <b>no current-hop-limit</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command configures the current-hop-limit in the router advertisement messages. It informs the nodes on the subnet about the hop-limit when originating IPv6 packets.
<b>Default</b>	64
<b>Parameters</b>	<i>number</i> — Specifies the hop limit. <b>Values</b> 0 — 255. A value of zero means there is an unspecified number of hops.

### managed-configuration

<b>Syntax</b>	<b>[no] managed-configuration</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address autoconfigured using stateless address autoconfiguration. See RFC 3315, <i>Dynamic Host Configuration Protocol (DHCP) for IPv6</i> .
<b>Default</b>	no managed-configuration

### max-advertisement-interval

<b>Syntax</b>	<b>[no] max-advertisement-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command configures the maximum interval between sending router advertisement messages.
<b>Default</b>	600
<b>Parameters</b>	<i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages. <b>Values</b> 4 — 1800

### min-advertisement-interval

<b>Syntax</b>	<b>[no] min-advertisement-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vprn>router-advert>if

<b>Description</b>	This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.
<b>Default</b>	200
<b>Parameters</b>	<i>seconds</i> — Specify the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages.
<b>Values</b>	3 — 1350

## mtu

<b>Syntax</b>	[no] mtu <i>mtu-bytes</i>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command configures the MTU for the nodes to use to send packets on the link.
<b>Default</b>	no mtu — The MTU option is not sent in the router advertisement messages.
<b>Parameters</b>	<i>mtu-bytes</i> — Specify the MTU for the nodes to use to send packets on the link.
<b>Values</b>	1280 — 9212

## other-stateful-configuration

<b>Syntax</b>	[no] other-stateful-configuration
<b>Description</b>	This command sets the "Other configuration" flag. This flag indicates that DHCPv6lite is available for autoconfiguration of other (non-address) information such as DNS-related information or information on other servers in the network. See RFC 3736, <i>Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6</i> .
<b>Default</b>	no other-stateful-configuration

## prefix

<b>Syntax</b>	[no] prefix [ <i>ipv6-prefix/prefix-length</i> ]				
<b>Context</b>	config>service>vprn>router-advert>if				
<b>Description</b>	This command configures an IPv6 prefix in the router advertisement messages. To support multiple IPv6 prefixes, use multiple prefix statements. No prefix is advertised until explicitly configured using prefix statements.				
<b>Default</b>	none				
<b>Parameters</b>	<i>ip-prefix</i> — The IP prefix for prefix list entry in dotted decimal notation.				
<b>Values</b>	<table> <tr> <td>ipv4-prefix</td> <td>a.b.c.d (host bits must be 0)</td> </tr> <tr> <td>ipv4-prefix-length</td> <td>0 — 32</td> </tr> </table>	ipv4-prefix	a.b.c.d (host bits must be 0)	ipv4-prefix-length	0 — 32
ipv4-prefix	a.b.c.d (host bits must be 0)				
ipv4-prefix-length	0 — 32				

## Router Advertisement Commands

ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d x: [0 — FFFF]H d: [0 — 255]D
ipv6-prefix-length	0 — 128

**prefix-length** — Specifies a route must match the most significant bits and have a prefix length.

**Values** 1 — 128

### autonomous

<b>Syntax</b>	<b>[no] autonomous</b>
<b>Context</b>	config>service>vprn>router-advert>if>prefix
<b>Description</b>	This command specifies whether the prefix can be used for stateless address autoconfiguration.
<b>Default</b>	enabled

### on-link

<b>Syntax</b>	<b>[no] on-link</b>
<b>Context</b>	config>service>vprn>router-advert>if>prefix
<b>Description</b>	This command specifies whether the prefix can be used for onlink determination.
<b>Default</b>	enabled

### preferred-lifetime

<b>Syntax</b>	<b>[no] preferred-lifetime {seconds   infinite}</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command configures the remaining length of time in seconds that this prefix will continue to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected.
<b>Default</b>	604800
<b>Parameters</b>	<i>seconds</i> — Specifies the remaining length of time in seconds that this prefix will continue to be preferred. <i>infinite</i> — Specifies that the prefix will always be preferred. A value of 4,294,967,295 represents infinity.

## valid-lifetime

<b>Syntax</b>	<b>valid-lifetime</b> { <i>seconds</i>   <b>infinite</b> }
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity.  The address generated from an invalidated prefix should not appear as the destination or source address of a packet.
<b>Default</b>	2592000
<b>Parameters</b>	<i>seconds</i> — Specifies the remaining length of time in seconds that this prefix will continue to be valid. <b>infinite</b> — Specifies that the prefix will always be valid. A value of 4,294,967,295 represents infinity.

## reachable-time

<b>Syntax</b>	<b>reachable-time</b> <i>milli-seconds</i> <b>no reachable-time</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.
<b>Default</b>	no reachable-time
<b>Parameters</b>	<i>milli-seconds</i> — Specifies the length of time the router should be considered reachable.  <b>Values</b> 0 — 3600000

## retransmit-time

<b>Syntax</b>	<b>retransmit-timer</b> <i>milli-seconds</i> <b>no retransmit-timer</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command configures the retransmission frequency of neighbor solicitation messages.
<b>Default</b>	no retransmit-time
<b>Parameters</b>	<i>milli-seconds</i> — Specifies how often the retransmission should occur.  <b>Values</b> 0 — 1800000

### router-lifetime

<b>Syntax</b>	<b>router-lifetime</b> <i>seconds</i> <b>no router-lifetime</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command sets the router lifetime.
<b>Default</b>	1800
<b>Parameters</b>	<i>seconds</i> — The length of time, in seconds, (relative to the time the packet is sent) that the prefix is valid for route determination. <b>Values</b> 0, 4 — 9000 seconds. 0 means that the router is not a default router on this link.

### use-virtual-mac

<b>Syntax</b>	<b>[no] use-virtual-mac</b>
<b>Context</b>	config>service>vprn>router-advert>if
<b>Description</b>	This command enables sending router advertisement messages using the VRRP virtual MAC address, provided that the virtual router is currently the master. If the virtual router is not the master, no router advertisement messages are sent. The <b>no</b> form of the command disables sending router advertisement messages.
<b>Default</b>	no use-virtual-mac



---

## NAT Commands

### nat

<b>Syntax</b>	<code>[no] nat</code>
<b>Context</b>	<code>config&gt;service&gt;vprn</code> <code>config&gt;router</code>
<b>Description</b>	This command configures, creates or deletes a NAT instance.

### inside

<b>Syntax</b>	<code>inside</code>
<b>Context</b>	<code>config&gt;service&gt;vprn&gt;nat</code> <code>config&gt;router&gt;nat</code>
<b>Description</b>	This command enters the “inside” context to configure the inside NAT instance.

### destination-prefix

<b>Syntax</b>	<code>[no] destination-prefix ip-prefix/length</code>
<b>Context</b>	<code>config&gt;service&gt;vprn&gt;nat&gt;inside</code> <code>config&gt;router&gt;nat&gt;inside</code>
<b>Description</b>	This command configures a destination prefix. An (internal) static route will be created for this prefix. All traffic that hits this route will be subject to NAT. The system will not allow a destination-prefix to be configured if the configured nat-policy refers to an IP pool that resides in the same service (as this would result in a routing loop).
<b>Parameters</b>	<p><i>ip-prefix</i> — Specifies the IP prefix; host bits must be zero (0).</p> <p><b>Values</b> a.b.c.d</p> <p><i>length</i> — Specifies the prefix length.</p> <p><b>Values</b> 0 — 32</p>

### dual-stack-lite

<b>Syntax</b>	<code>dual-stack-lite</code>
<b>Context</b>	<code>config&gt;service&gt;vprn&gt;nat</code> <code>config&gt;router&gt;nat&gt;inside</code>



**Values** 32 — 64, 128

**Default** 128

## l2-aware

**Syntax** **l2-aware**

**Context** config>services>vprn>nat>inside

**Description** This command enables the context to configure parameters specific to Layer 2-aware NAT.

## address

**Syntax** [**no**] **address** *ip-address/mask*

**Context** config>services>vprn>nat>inside>l2-aware

**Description** This command configures a Layer 2-aware NAT address. This address will act as a local address of the system. Hosts connected to the inside service will be able to ARP for this address. To verify connectivity, a host can also ping the address. This address is typically used as next hop of the default route of a Layer 2-aware host. The given mask defines a Layer 2-aware subnet. The (inside) IP address used by anLayer 2-aware host must match one of the subnets defined here or it will be rejected.

**Parameters** *ip-address* — Specifies the IP address in a.b.c.d format.

*mask* — Specifies the mask.

**Values** 16 — 32

## nat-policy

**Syntax** **nat-policy** *nat-policy-name*  
**no nat-policy**

**Context** config>services>vprn>nat>inside  
config>router>nat>inside

**Description** This command configures the NAT policy that will be used for large-scale NAT in this service.

**Parameters** *nat-policy-name* — Specifies the NAT policy name.

**Values** 32 chars max

### redundancy

<b>Syntax</b>	<b>redundancy</b>
<b>Context</b>	config>service>vprn>nat>inside config>service>vprn>nat>outside>pool
<b>Description</b>	This command enables the context to configure redundancy parameters.

### peer

<b>Syntax</b>	<b>peer</b> <i>ip-address</i> <b>no peer</b>
<b>Context</b>	config>service>vprn>nat>inside>redundancy
<b>Description</b>	This command configures the IP address of the NAT redundancy peer in the realm of this virtual router instance.

### steering-route

<b>Syntax</b>	<b>steering-route</b> <i>ip-prefix/length</i> <b>no steering-route</b>
<b>Context</b>	config>service>vprn>nat>inside>redundancy
<b>Description</b>	This command configures specifies the IP address and prefix length of the steering route. The steering route is used in the realm of this virtual router instance as an indirect next-hop for all the traffic that must be routed to the large scale NAT function.

### outside

<b>Syntax</b>	<b>outside</b>
<b>Context</b>	config>service>vprn>nat config>router>nat
<b>Description</b>	This command enters the “outside” context to configure the outside NAT instance.

### pool

<b>Syntax</b>	<b>pool</b> <i>nat-pool-name</i> [ <b>nat-group</b> <i>nat-group-id</i> <b>type</b> <i>pool-type</i> [ <b>no-allocate</b> ] [ <b>create</b> ]] <b>no pool</b> <i>nat-pool-name</i>
<b>Context</b>	config>service>vprn>nat>outside config>router>nat>outside

<b>Description</b>	This command configures a NAT pool.
<b>Parameters</b>	<i>nat-pool-name</i> — Specifies the NAT pool name. <b>Values</b> 32 chars max <i>nat-group-id</i> — Specifies the NAT group ID. <b>Values</b> 1 — 4 <b>create</b> — This parameter must be specified to create the instance. <i>pool-type</i> — Species the pool type, either large-scale or L2-aware.

## address-range

<b>Syntax</b>	<b>address-range</b> <i>start-ip-address end-ip-address</i> [ <b>create</b> ] <b>no address-range</b> <i>start-ip-address end-ip-address</i>
<b>Context</b>	config>service>vprn>nat>outside>pool config>router>nat>outside>pool
<b>Description</b>	This command configures a NAT address range.
<b>Parameters</b>	<i>start-ip-address</i> — Specifies the beginning IP address in a.b.c.d form. <i>end-ip-address</i> — Specifies the ending IP address in a.b.c.d. form. <b>create</b> — This parameter must be specified to create the instance.

## description

<b>Syntax</b>	<b>description</b> <i>description-string</i> <b>no description</b>
<b>Context</b>	config>service>vprn>nat>outside>pool>address-range config>service>vprn>nat>outside>pool config>router>nat>outside>pool>address-range config>router>nat>outside>pool
<b>Description</b>	This command configures the description for the NAT address range.
<b>Parameters</b>	<i>description-string</i> — Specifies the NAT address range description. <b>Values</b> 80 chars max

## drain

<b>Syntax</b>	[no] <b>drain</b>
<b>Context</b>	config>service>vprn>nat>outside>pool>address-range config>router>nat>outside>pool>address-range

## Router Advertisement Commands

**Description** This command starts or stops draining this NAT address range. When an address-range is being drained, it will not be used to serve new hosts. Existing hosts, however, will still be able to use the address that was assigned to them even if it is being drained. An address-range can only be deleted if the parent pool is shut down or if the range itself is effectively drained (no hosts are using the addresses anymore).

### mode

**Syntax** **mode** {**auto** | **napt**}  
**no mode**

**Context** config>service>vprn>nat>outside>pool

**Description** This command configures the mode of operation of this NAT address pool.  
The mode value is only relevant while the value of pool type is equal to largeScale; while the value of pool type is equal to l2Aware, the mode of operation is always NAPT.

### port-forwarding-range

**Syntax** **port-forwarding-range** *range-end*  
**no port-forwarding-range**

**Context** config>service>vprn>nat>outside>pool

**Description** This command configures the end of the port range available for port forwarding. The start of the range is always equal to one.  
The actual maximum value of the range end may be restricted to less than 65535 depending on the value of the objects port reservation type and port reservation value and on system specifications.

**Default** 1023

**Parameters** *range-end* — Specifies the mode of operation of this NAT pool

**Values** 1023 — 65535

### port-reservation

**Syntax** **port-reservation blocks** *num-blocks*  
**port-reservation ports** *num-ports*  
**no port-reservation**

**Context** config>service>vprn>nat>outside>pool  
config>router>nat>outside>pool

**Description** This command configures the size of the port-block that will be assigned to a host that is served by this pool. The number of ports configured here will be available to UDP, TCP and ICMP (as identifiers).

- Parameters** *num-blocks* — Specifies the number of port-blocks per IP address. Setting *num-blocks* to one (1) for large scale NAT will enable 1:1 NAT for IP addresses in this pool.
- Values** 1 — 64512
- num-ports* — Specifies the number of ports per block.
- Values** 1 — 32256

## export

- Syntax** **export** *ip-prefix/length*  
**no export**
- Context** config>service>vprn>nat>outside>pool>redundancy
- Description** This command installs the export route in the routing table for active NAT pools.
- Once the export route is in the routing table, it can be advertised in the network via a routing protocol. NAT pools in the standby or disabled state will not advertise the export route.
- A NAT pool will become active when it becomes operationally UP, AND there is no monitoring route (which is also the export route from the peer) present in the routing node (as received from the network). The pool will transition into standby state in case that the monitoring route (or export route from the peer) is already present in the routing table. In other words, the monitoring route is already advertised as an export route from the peering node with active NAT pool.
- The export route can be advertised only from :
- The active lead pool.
  - Active pool for which fate-sharing is disabled.
- Default** no export
- Parameters** *ip-prefix/length* — Specifies the IP prefix and length.
- Syntax:** ip-prefix/length : ip-prefix a.b.c.d  
ip-prefix-length 0 — 32
- Values** 0, 4, 16

## follow

- Syntax** **follow router** *router-instance pool name*  
**no follow**
- Context** configure>service>vprn>nat>outside>pool>redundancy  
configure>router> nat>outside>pool>redundancy
- Description** This command implicitly enables Pool Fate-Sharing Group (PFSG) which is required in case of multiple NAT policies per inside routing context. A NAT pool configured with this command will not advertise or monitor any route in order to change its (activity) state but instead it will directly follow

## Router Advertisement Commands

the state of the lead pool in the PFSG. Once the lead pool changes its (activity) state, all the remaining pools following the lead pool will change their state accordingly.

**Default** no follow

**Parameters** **router** *router-instance* — Specifies the routing instance where the lead pool resides.

**Values** <router-name>|<service-id>  
router-name - "Base"  
service-id - [1..2147483647]

**pool name** — The pool whose activity state is being shared up to 32 characters in length.

## monitor

**Syntax** **monitor** *ip-prefix/length*  
**no monitor**

**Context** configure>service>vprn>nat>outside>pool>redundancy  
configure>router> nat>outside>pool>redundancy

**Description** This command configures the monitoring route based on which the NAT multi-chassis switchover is triggered. Monitoring route of a NAT pool on the local node must match the export route of a corresponding NAT pool on the peering node. Presence of the monitoring route in the routing table is an indication that the peering NAT pool is active (since it is advertizing its export route). The disappearance of the monitoring route from the routing table is an indication that the peering pool has failed and consequently the nodal switchover is triggered, the local pool becomes active and its export route is consequently advertized. The export route can be advertized only from :

- The active lead pool.
- Active pool for which fate-sharing is disabled.

**Parameters** *ip-prefix/length* — Specifies the IP prefix and length.

**Syntax:** ip-prefix/length : ip-prefix a.b.c.d  
ip-prefix-length 0 — 32

## subscriber-limit

**Syntax** **subscriber-limit** [1..65535]  
**no subscriber-limit**

**Context** config>service>vprn>nat>outside>pool

**Description** This command configures the maximum number of subscribers per outside IP address.

If multiple port blocks per subscriber are used, the block size is typically small; all blocks assigned to a given subscriber belong to the same IP address; the subscriber limit guarantees that any subscriber can get a minimum number of ports.

**Parameters** *limit* — Specifies the maximum number of subscribers per outside IP address.



**Values** 1 — 65535

## watermarks

<b>Syntax</b>	<b>watermarks high</b> <i>percentage-high</i> <b>low</b> <i>percentage-low</i> <b>no watermarks</b>
<b>Context</b>	config>service>vprn>nat>outside>pool config>router>nat>outside>pool
<b>Description</b>	This command configures the watermarks for this NAT pool.
<b>Parameters</b>	<i>percentage-high</i> — Specifies the high percentage. <b>Values</b> 2 — 100 <i>percentage-low</i> — Specifies the low percentage. <b>Values</b> 1 — 99

---

## Subscriber Interface Commands

### subscriber-interface

<b>Syntax</b>	<b>subscriber-interface</b> <i>ip-int-name</i> [ <b>fwd-service</b> <i>service-id</i> <b>fwd-subscriber-interface</b> <i>ip-int-name</i> ] <b>no subscriber-interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.  Use the <b>no</b> form of the command to remove the subscriber interface.
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.  <b>fwd-service</b> <i>service-id</i> — Specifies the forwarding service ID for a subscriber interface in a retailer context.  <b>fwd-subscriber-interface</b> <i>ip-int-name</i> — Specifies the forwarding subscriber interface for a subscriber interface in a retailer context.

### address

<b>Syntax</b>	[ <b>no</b> ] <b>address</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } [ <b>gw-ip-address</b> <i>ip-address</i> ] [ <b>populate-host-routes</b> ] [ <b>track-srrp</b> <i>srrp-instance</i> [ <b>holdup-time</b> <i>msecs</i> ]]
<b>Context</b>	config>service>vprn>subscriber-interface
<b>Description</b>	This command configures the local subscriber subnets available on a subscriber IP interface. The configured <i>ip-address</i> and <i>mask</i> define the address space associated with the subscriber subnet. Up to 16 IP subnets can be created on a single subscriber IP interface. Each subnet supports a locally owned IP host address within the subnet that is not expected to appear on other routers that may be servicing the same subscriber subnet. For redundancy purposes, the keyword <b>gw-address</b> defines a separate IP address within the subnet for Subscriber Routed Redundancy Protocol (SRRP) routing. This IP address must be the same on the local and remote routers participating in a common SRRP instance.  In SRRP, a single SRRP instance is tied to a group IP interface. The group IP interface is contained directly within a subscriber IP interface context and thus directly associated with the subscriber subnets on the subscriber IP interface. The SRRP instance is also indirectly associated with any subscriber subnets tied to the subscriber interface through wholesale/retail VPRN configurations. With the directly-associated and the indirectly-associated subscriber interface subnets, a single SRRP instance can manage hundreds of SRRP gateway IP addresses. This automatic subnet association to the SRRP instance is different from VRRP where the redundant IP address is defined within the VRRP context.

Defining an SRRP gateway IP address on a subscriber subnet is not optional when the subnet is associated with a group IP interface with SRRP enabled. Enabling SRRP (**no shutdown**) will fail if one or more subscriber subnets do not have an SRRP gateway IP address defined. Creating a new subscriber subnet without an SRRP gateway IP address defined will fail when the subscriber subnet is associated with a group IP interface with an active SRRP instance. Once SRRP is enabled on a group interface, the SRRP instance will manage the ARP response and routing behavior for all subscriber hosts reachable through the group IP interface.

The `no` form of the command removes the address from a subscriber subnet. The `address` command for the specific subscriber subnet must be executed without the `gw-address` parameter. To succeed, all SRRP instances associated with the subscriber subnet must be removed or shutdown.

- Parameters**
- ip-address/mask* | *ip-address netmask* — Specifies the address space associated with the subscriber subnet
  - gw-ip-address** *ip-address* — Specifies a separate IP address within the subnet for SRRP routing purposes. This parameter must be followed by a valid IP interface that exists within the subscriber subnet created by the `address` command. The defined gateway IP address cannot currently exist as a subscriber host (static or dynamic). If the defined `ip-address` already exists as a subscriber host address, the `address` command will fail. The specified `ip-address` must be unique within the system.
 

The `gw-address` parameter may be specified at anytime. If the subscriber subnet was created previously, executing the `address` command with a `gw-address` parameter will simply add the SRRP gateway IP address to the existing subnet.

If the `address` command is executed without the `gw-address` parameter when the subscriber subnet is associated with an active SRRP instance, the `address` will fail. If the SRRP instance is inactive or removed, executing the `address` command without the `gw-address` parameter will remove the SRRP gateway IP address from the specified subscriber subnet.

If the `address` command is executed with a new `gw-address`, all SRRP instances currently associated with the specified subscriber subnet will be updated with the new SRRP gateway IP address.
  - populate-host-routes** — Specifies to populate subscriber-host routes in local FIB. Storing them in FIB benefits topologies only where the external router advertises more specific routes than the one corresponding to locally configured subscriber-interface subnets.

## allow-unmatching-subnets

- Syntax** `[no] allow-unmatching-subnets`
- Context** `config>service>vprn>subscriber-interface`
- Description** This command specifies whether subscriber hosts with a subnet that does not match any of the subnets configured on this interface, are allowed.

## group-interface

<b>Syntax</b>	<b>[no] group-interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn>subscriber-interface
<b>Description</b>	This command enables the context to configure a group interface. A group interface is an interface that may contain one or more SAPs. This interface is used in triple-play services where multiple SAPs are part of the same subnet.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-int-name</i> — Configures the interface group name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## arp-host

<b>Syntax</b>	<b>arp-host</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables the context to configure ARP host parameters.

## host-limit

<b>Syntax</b>	<b>host-limit</b> <i>max-num-hosts</i> <b>no host-limit</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>arp-host
<b>Description</b>	This command configures the maximum number of ARP hosts.
<b>Parameters</b>	<i>max-num-hosts</i> — Specifies the maximum number of ARP hosts. <b>Values</b> 1 — 32767

## min-auth-interval

<b>Syntax</b>	<b>min-auth-interval</b> <i>min-auth-interval</i> <b>no min-auth-interval</b>
<b>Context</b>	onfig>service>vprn>sub-if>grp-if>arp-host
<b>Description</b>	This command configures the minimum authentication interval.
<b>Parameters</b>	<i>min-auth-interval</i> — Specifies the minimum authentication interval. <b>Values</b> 1 — 6000

## sap-host-limit

<b>Syntax</b>	<b>sap-host-limit</b> <i>max-num-hosts-sap</i> <b>no sap-host-limit</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>arp-host
<b>Description</b>	This command configures the maximum number of ARP hosts per SAP.
<b>Parameters</b>	<i>max-num-hosts-sap</i> — Specifies the maximum number of ARP hosts per SAP allowed on this IES interface.
<b>Values</b>	1 — 32767

---

## PPPoE Commands

### pppoe

<b>Syntax</b>	<b>[no] pppoe</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables the context to configure PPPoE parameters.

### dhcp-client

<b>Syntax</b>	<b>dhcp-client</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>pppoe
<b>Description</b>	This command enables the context to configure the PPPoE-to-DHCP options.

### ccag-use-origin-sap

<b>Syntax</b>	<b>[no] ccag-use-origin-sap</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>pppoe>dhcp-client
<b>Description</b>	This command enables the original VPLS SAP to be included in the circuit-id option to send to the DHCP server (in case this interface is connected to a VPLS by a CCA MDA). The <b>no</b> form of the command disables the feature.
<b>Default</b>	no ccag-use-origin-sap

### pap-chap-user-db

<b>Syntax</b>	<b>pap-chap-user-db</b> <i>local-user-db-name</i> <b>no pap-chap-user-db</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>pppoe
<b>Description</b>	This command configures the local user database to use for PPP Challenge-Handshake Authentication Protocol/Password Authentication Protocol (PAP/CHAP) authentication. If an authentication policy is also configured, <b>pppoe-access-method</b> must be set to none in this authentication policy to use the local user database (in that case RADIUS authentication will not be used for PPPoE hosts).
<b>Parameters</b>	<i>local-user-db-name</i> — Specifies the local user database to use for authentication.

## pppoe-policy

<b>Syntax</b>	<b>pppoe-policy</b> <i>pppoe-policy-name</i> <b>no pppoe-policy</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>pppoe
<b>Description</b>	This command associates a PPPoE policy on this interface.
<b>Default</b>	default
<b>Parameters</b>	<i>pppoe-policy-name</i> — Specifies a a PPPoE policy up to 32 characters in length on this interface.

## sap-session-limit

<b>Syntax</b>	<b>sap-session-limit</b> <i>sap-session-limit</i> <b>no sap-session-limit</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>pppoe
<b>Description</b>	This command specifies the number of PPPoE hosts per SAP allowed for this group-interface.
<b>Default</b>	1
<b>Parameters</b>	<i>sap-session-limit</i> — Specifies the number of PPPoE hosts per SAP allowed. <b>Values</b> 1 — 20000

## session-limit

<b>Syntax</b>	<b>session-limit</b> <i>session-limit</i> <b>no session-limit</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>pppoe
<b>Description</b>	This command specifies the number of PPPoE hosts allowed for this group interface.
<b>Default</b>	1
<b>Parameters</b>	<i>session-limit</i> — Specifies the number of PPPoE hosts allowed <b>Values</b> 1 — 20000

---

## Interface ICMP Commands

### icmp

<b>Syntax</b>	<b>icmp</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

### mask-reply

<b>Syntax</b>	<b>[no] mask-reply</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	<p>This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.</p> <p>If a local node sends an ICMP mask request to the router interface, the <b>mask-reply</b> command configures the router interface to reply to the request.</p> <p>By default, the router instance will reply to mask requests.</p> <p>The <b>no</b> form of this command disables replies to ICMP mask requests on the router interface.</p>
<b>Default</b>	<b>mask-reply</b> — Reply to ICMP mask requests.

### redirects

<b>Syntax</b>	<b>redirects</b> [ <i>number seconds</i> ] <b>no redirects</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	<p>This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.</p> <p>When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.</p>



The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

<b>Default</b>	<b>redirects 100 10</b> — Maximum of 100 redirect messages in 10 seconds.
<b>Parameters</b>	<i>number</i> — The maximum number of ICMP redirect messages to send. This parameter must be specified with the <i>seconds</i> parameter.
	<b>Values</b> 10 — 1000
	<i>seconds</i> — The time frame in seconds used to limit the <i>seconds</i> of ICMP redirect messages that can be issued.
	<b>Values</b> 1 — 60

## tll-expired

<b>Syntax</b>	<b>tll-expired</b> <i>number seconds</i> <b>no tll-expired</b>
<b>Context</b>	config>service>vprn>if>icmp config>service>vprn>sub-if>grp-if>icmp config>service>vprn>nw-if>icmp
<b>Description</b>	Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.  By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.  The <b>no</b> form of this command disables the limiting the rate of TTL expired messages on the router interface.
<b>Default</b>	tll-expired 100 10
<b>Parameters</b>	<i>number</i> — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the <i>seconds</i> parameter.
	<b>Values</b> 10 — 1000
	<i>seconds</i> — The time frame in seconds used to limit the <i>number</i> of ICMP TTL expired messages that can be issued, expressed as a decimal integer.
	<b>Values</b> 1 — 60

### if-attribute

<b>Syntax</b>	<b>if-attribute</b>
<b>Context</b>	config>service>vprn>interface
<b>Description</b>	This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG).

### admin-group

<b>Syntax</b>	<b>admin-group</b> <i>group-name</i> [ <i>group-name...</i> (up to 5 max)] <b>no admin-group</b> <i>group-name</i> [ <i>group-name...</i> (up to 5 max)] <b>no admin-group</b>
<b>Context</b>	config>service>vprn>interface>if-attribute
<b>Description</b>	<p>This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface.</p> <p>Each single operation of the <b>admin-group</b> command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.</p> <p>The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.</p> <p>It should be noted that only the admin groups bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the <b>traffic-engineering</b> option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.</p> <p>The <b>no</b> form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.</p>
<b>Parameters</b>	<i>group-name</i> — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

### srlg-group

<b>Syntax</b>	<b>srlg-group</b> <i>group-name</i> [ <i>group-name...</i> (up to 5 max)] <b>no srlg-group</b> <i>group-name</i> [ <i>group-name...</i> (up to 5 max)] <b>no srlg-group</b>
<b>Context</b>	config>service>vprn>interface>if-attribute
<b>Description</b>	<p>This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.</p> <p>An interface can belong to up to 64 SRLG groups. However, each single operation of the <b>srlg-group</b> command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.</p>

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

It should be noted that only the SRLGs bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

**Parameters** *group-name* — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

## unreachables

**Syntax** **unreachables** [*number seconds*]  
**no unreachable**

**Context** config>service>vprn>if>icmp  
config>service>vprn>sub-if>grp-if>icmp  
config>service>vprn>nw-if>icmp

**Description** This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

**Default** **unreachables 100 10**

**Parameters** *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

**Values** 10 — 1000

*seconds* — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

**Values** 1 — 60

ip-mtu

<b>Syntax</b>	<b>ip-mtu</b> <i>octets</i> <b>no ip-mtu</b>
<b>Context</b>	config>service>vprn>nw-if
<b>Description</b>	This command configures the IP maximum transmit unit (packet) for the associated router IP interface.  The configured IP-MTU cannot be larger than the calculated IP MTU based on the port MTU configuration.  The MTU that is advertised from the IES size is:  MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))  The <b>no</b> form of the command returns the associated IP interfaces MTU to its default value, which is calculated, based on the port MTU setting. (For Ethernet ports this will typically be 1554.)
<b>Default</b>	<b>no ip-mtu</b>
<b>Parameters</b>	octets —  <b>Values</b> 512 — 9000

lag

<b>Syntax</b>	<b>lag</b> <i>lag-id[:encap-val]</i> <b>no lag</b>
<b>Context</b>	config>service>vprn>nw-if
<b>Description</b>	This command binds the interface to a Link Aggregation Group (LAG)  The <b>no</b> form of the command removes the LAG id from the configuration.
<b>Parameters</b>	<i>lag-ig[:encap-val]</i> — Specifies the LAG ID.  <b>Values</b> lag-id            1 — 800 encap-val        0        for null 0 — 4094 for dot1q

---

## Interface SAP Commands

### sap

<b>Syntax</b>	<b>sap</b> <i>sap-id</i> [ <b>create</b> ] <b>no sap</b> <i>sap-id</i>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if config>service>vprn>ipsec-if>sap
<b>Description</b>	<p>This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.</p> <p>All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.</p> <p>Enter an existing SAP without the <b>create</b> keyword to edit SAP parameters. The SAP is owned by the service in which it was created.</p> <p>A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the <b>config interface</b> <i>port-type port-id mode access</i> command. Channelized TDM ports are always access ports.</p> <p>If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.</p> <p>The <b>no</b> form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. The <b>no</b> form of the command causes the ptp-h-assist to be disabled.</p>
<b>Default</b>	No SAPs are defined.
<b>Special Cases</b>	<p><b>VPRN</b> — A VPRN SAP must be defined on an Ethernet interface.</p> <p><b>sap ipsec-id.private   public:tag</b> — This parameter associates an IPsec group SAP with this interface. This is the public side for an IPsec tunnel. Tunnels referencing this IPsec group in the private side may be created if their local IP is in the subnet of the interface subnet and the routing context specified matches with the one of the interface.</p> <p>This context will provide a SAP to the tunnel. The operator may associate an ingress and egress QoS policies as well as filters and virtual scheduling contexts. Internally this creates an Ethernet SAP that will be used to send and receive encrypted traffic to and from the MDA. Multiple tunnels can be associated with this SAP. The “tag” will be a dot1q value. The operator may see it as an identifier. The range is limited to 1 — 4094.</p>
<b>Parameters</b>	<p><i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.</p> <p><i>port-id</i> — Specifies the physical port ID in the <i>slot/mda/port</i> format.</p>

## Router Advertisement Commands

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot\_number/MDA\_number/port\_number format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels the port ID must include the channel ID. A period “.” separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

**create** — Keyword used to create a SAP instance.

**split-horizon-group** *group-name* — Specifies the name of the split horizon group to which the SAP belongs.

## aarp

<b>Syntax</b>	<b>aarp</b> <i>aarpId</i> <b>type</b> <i>type</i> <b>no aarp</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>if>spoke-sdp
<b>Description</b>	<p>This command associates an aarp instance to a multi-homed SAP or spoke-sdp. This instance is paired with the same aarp-id in the same node or in a peer node as part of a configuration to provide flow and packet asymmetry removal for traffic for a multi-homed SAP or spoke-sdp.</p> <p>The type specifies the role of this service point in the AARP: primary (dual-homed), secondary (dual-homed-secondary). The AA service attributes (app-profile, transit-policy) of the primary are inherited by the secondary endpoints. All endpoints within an aarp must be of the same type (sap or spoke), and all endpoints with an aarp must be within the same service.</p> <p>The <b>no</b> form of the command removes the association.</p>
<b>Default</b>	no aarp
<b>Parameters</b>	<p><i>aarpId</i> — Specifies the AARP instance associated with this SAP. If not configured, no AARP instance is associated with this SAP.</p> <p><b>Values</b> 1 —</p> <p><b>type</b> — Specifies the role of the SAP referenced by the AARP instance identified by AARP ID.</p> <p><b>Values</b> <b>dual-homed</b> — the primary dual homed aa-subscriber side service point of an aarp instance, only supported for IES and VPRN SAP and spoke-sdp <b>dual-homed-secondary</b> — One of the secondary dual homed aa-subscriber side service points of an aarp instance, only supported for IES and VPRN SAP and spoke-sdp.</p>

## tod-suite

<b>Syntax</b>	<b>tod-suite</b> <i>tod-suite-name</i> <b>no tod-suite</b>
<b>Context</b>	config>service>vprn>if>sap
<b>Description</b>	This command applies a time-based policy (filter or QoS policy) to the SAP. The suite name must already exist in the <b>config&gt;system&gt;cron</b> context.
<b>Default</b>	no tod-suite
<b>Parameters</b>	<i>tod-suite-name</i> — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP or a subscriber. The suite can be applied to more than one SAP.

## transit-policy

<b>Syntax</b>	<b>transit-policy</b> <i>ip-aasub-policy-id</i> <b>transit-policy prefix</b> <i>prefix-aasub-policy-id</i> <b>no transit-ip-policy</b>
<b>Context</b>	config>service>vprn>if>sap> config>service>vprn>if>spoke-sdp>
<b>Description</b>	This command associates a transit aa subscriber IP policy to the service. The transit IP policy must be defined prior to associating the policy with a SAP in the <b>config&gt;application assurance&gt;group&gt;policy&gt;transit-ip-policy</b> context.  Transit AA subscribers are managed by the system through the use of this policy assigned to services, which determines how transit subs are created and removed for that service.  The <b>no</b> form of the command removes the association of the policy to the service.
<b>Default</b>	no transit-ip-policy
	<i>ip-aasub-policy-id</i> — An integer that identifies a transit IP profile entry. <b>Values</b> 1 — 65535
	<i>prefix-aasub-policy-id</i> — An integer that identifies a prefix aasub-policy ID. <b>Values</b> 1 — 65535

## accounting-policy

<b>Syntax</b>	<b>accounting-policy</b> <i>acct-policy-id</i> <b>no accounting-policy</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>if>spoke-sdp
<b>Description</b>	This command creates the accounting policy context that can be applied to an interface SAP or interface SAP spoke SDP.

## Router Advertisement Commands

An accounting policy must be defined before it can be associated with a SAP.

If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SAP, and the accounting policy reverts to the default.

**Default** Default accounting policy.

**Parameters** *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

**Values** 1 — 99

## app-profile

**Syntax** **app-profile** *app-profile-name*  
**no app-profile**

**Context** config>service>vprn>if>spoke-sdp

**Description** This command configures the application profile name.

**Parameters** *app-profile-name* — Specifies the application profile name.

**Values** 32 chars max

## bfd-enable

**Syntax** **bfd-enable**  
**no bfd-enable**

**Context** config>service>vprn>interface>spoke-sdp

**Description** This command enables VCCV BFD on the PW associated with the VLL, BGP VPWS, or VPLS service. The parameters for the BFD session are derived from the named BFD template, which must have been first configured using the **bfd-template** command.

## bfd-template

**Syntax** **bfd-template** *name*  
**no bfd-template**

**Context** config>service>vprn>interface>spoke-sdp

**Description** This command configures a named BFD template to be used by VCCV BFD on PWs belonging to the VLL, BGP VPWS, or VPLS service. The template specifies parameters, such as the minimum



transmit and receive control packet timer intervals, to be used by the BFD session. Template parameters are configured under the **config>router>bfd** context.

**Default** no bfd-template

**Parameters** *name* — A text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

## collect-stats

**Syntax** [no] collect-stats

**Context** config>service>vprn>if>sap  
config>service>vprn>if>spoke-sdp

**Description** This command enables accounting and statistical data collection for either an interface SAP or interface SAP spoke SDP, or network port. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

**Default** no collect-stats

## cpu-protection

**Syntax** **cpu-protection** *policy-id* [mac-monitoring] | [eth-cfm-monitoring [aggregate][car]]  
**no cpu-protection**

**Context** config>service>vprn>sub-if>grp-if>sap

**Description** This command assigns an existing CPU protection policy to the associated group interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU-Protection policy is assigned to a group interface SAP, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.

The **no** form of the command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.

**Description** cpu-protection 254 (for access interfaces)  
cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

**Parameters** *policy-id* — Specifies an existing CPU protection policy.

## Router Advertisement Commands

**Values** 1 — 255

**mac-monitoring** — This keyword enables MAC monitoring.

**eth-cfm-monitoring** — This keyword enables Ethernet Connectivity Fault Management monitoring.

**aggregate** — This keyword applies the rate limit to the sum of the per peer packet rates.

**car** — (Committed Access Rate) This keyword causes Eth-CFM packets to be ignored when enforcing the overall-rate.

### dist-cpu-protection

<b>Syntax</b>	<b>dist-cpu-protection</b> <i>policy-name</i> <b>no dist-cpu-protection</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap config>service>>vprn>if>sap
<b>Description</b>	This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface (note that this rule does not apply to templates such as an msap-policy)
<b>Default</b>	no dist-cpu-protection

### default-host

<b>Syntax</b>	<b>default-host</b> <i>ip-address/mask next-hop next-hop-ip</i> <b>no default-host</b> <i>ip-address/mask</i>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	This command configures the default-host to be used. More than one default-host can be configured per SAP. The <b>no</b> form of the command removes the values from the configuration.
<b>Parameters</b>	<i>ip-address/mask</i> — Assigns an IP address/IP subnet format to the interface. <b>next-hop</b> <i>next-hop-ip</i> — Assigns the next hop IP address.

### source

<b>Syntax</b>	<b>source</b> <i>ip-address</i>
<b>Context</b>	config>service>ies>interface>sap>ip-tunnel
<b>Description</b>	This command configures the source IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The IPv4 or IPv6 address must belong to the one of the IP subnets associated with the public SAP interface of the tunnel-group. The <b>source</b> address, <b>remote-ip</b> address and <b>backup-remote-ip</b> address of a tunnel must all belong to the same

address family (IPv4 or IPv6). When the source address contains an IPv6 address it must be a global unicast address.

**Default** no source

**Parameters** *ip-address* — An IPv4 address or an IPv6 address.

## remote-ip

**Syntax** **remote-ip** *ip-address*  
**no remote-ip**

**Context** config>service>ies>interface>sap>ip-tunnel

**Description** This command configures the primary destination IPv4 or IPv6 address to use for an IP tunnel. This configuration applies to the outer IP header of the encapsulated packets. The **source** address, **remote-ip** address and **backup-remote-ip** address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the remote-ip address contains an IPv6 address it must be a global unicast address.

**Default** no remote-ip

**Parameters** *ip-address* — An IPv4 address or an IPv6 address.

## backup-remote-ip

**Syntax** **backup-remote-ip** *ip-address*  
**no remote-ip**

**Context** config>service>ies>interface>sap>ip-tunnel

**Description** This command configures the alternate destination IPv4 or IPv6 address to use for an IP tunnel. This destination address is used only if the primary destination configured with the **remote-ip** command is unreachable in the delivery service. The **source** address, **remote-ip** address and **backup-remote-ip** address of a tunnel must all belong to the same address family (IPv4 or IPv6). When the backup-remote-ip address contains an IPv6 address it must be a global unicast address.

**Default** no remote-ip

**Parameters** *ip-address* — An IPv4 address or an IPv6 address.

---

## Interface SAP ATM Commands

### atm

<b>Syntax</b>	<b>atm</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	<p>This command enables the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supporting ATM functionality such as:</p> <ul style="list-style-type: none"> <li>• Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality.</li> <li>• Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.</li> </ul> <p>If ATM functionality is not supported for a given context, the command returns an error.</p>

### egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
<b>Description</b>	This command configures egress ATM attributes for the SAP.

### encapsulation

<b>Syntax</b>	<b>encapsulation</b> <i>atm-encap-type</i>
<b>Context</b>	config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
<b>Description</b>	<p>This command configures RFC 2684, <i>Multiprotocol Encapsulation over ATM AAL5</i>, encapsulation for an ATM PVCC delimited SAP. This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition also references the ATM Forum LAN Emulation specification. Ingress traffic that does not match the configured encapsulation will be dropped.</p>
<b>Default</b>	The encapsulation is driven by the services for which the SAP is configured. For VPRN service SAPs, the default is <b>aal5snap-routed</b> .
<b>Parameters</b>	<i>atm-encap-type</i> — Specify the encapsulation type.

<b>Values</b>	<p><b>aal5snap-routed</b> — Routed encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684.</p> <p><b>aal5mux-ip</b> — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684.</p> <p><b>aal5snap-bridged</b> — Bridged encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684.</p> <p><b>aal5mux-bridged-eth-nofcs</b> — Bridged IP encapsulation for VC multiplexed circuit as defined in RFC 2684.</p>
---------------	---

## ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>service>vprn>if>sap>atm config>service>vprn>sub-if>grp-if>sap>atm
<b>Description</b>	This command configures ingress ATM attributes for the SAP.

## traffic-desc

<b>Syntax</b>	<b>traffic-desc</b> <i>traffic-desc-profile-id</i> <b>no traffic-desc</b>
<b>Context</b>	config>service>vprn>if>sap>atm>egress config>service>vprn>if>sap>atm>ingress config>service>vprn>sub-if>grp-if>sap>atm>egress config>service>vprn>sub-if>grp-if>sap>atm>ingress
<b>Description</b>	<p>This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP). When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.</p> <p>The <b>no</b> form of the command reverts the traffic descriptor to the default traffic descriptor profile.</p>
<b>Default</b>	The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.
<b>Parameters</b>	<i>traffic-desc-profile-id</i> — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

## oam

<b>Syntax</b>	<b>oam</b>
<b>Context</b>	config>service>vprn>if >sap>atm config>service>vprn>sub-if>grp-if>sap>atm
<b>Description</b>	This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

### alarm-cells

<b>Syntax</b>	<b>[no] alarm-cells</b>
<b>Context</b>	config>service>vprn>if>sap>atm>oam config>service>vprn>sub-if>grp-if>sap>atm>oam
<b>Description</b>	<p>This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCC's operational status.</p> <p>When alarm-cells functionality is enabled, a PVCC's operational status is affected when a PVCC goes into an AIS or RDI state because of an AIS/RDI processing (assuming nothing else affects PVCC's operational status, for example, if the PVCC goes DOWN, or enters a fault state and comes back UP, or exits that fault state). RDI cells are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI state, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).</p> <p>The <b>no</b> command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, a PVCC's operational status is no longer affected by a PVCC's OAM state changes due to AIS/RDI processing (note that when alarm-cells is disabled, a PVCC will change operational status to UP due to alarm-cell processing) and RDI cells are not generated as result of the PVCC going into AIS or RDI state. The PVCC's OAM status, however, will record OAM faults as described above.</p>
<b>Default</b>	Enabled for PVCCs delimiting VPRN SAPs

### periodic-loopback

<b>Syntax</b>	<b>[no] periodic-loopback</b>
<b>Context</b>	config>service>vprn>if >sap>atm>oam config>service>vprn>sub-if>grp-if>sap>atm
<b>Description</b>	<p>This command enables periodic OAM loopbacks on this SAP. This command is only configurable on VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the <code>config&gt;system&gt;atm&gt;oam&gt;loopback-period <i>period</i></code> context.</p> <p>If a response is not received and consecutive retry-down retries also result in failure, the endpoint will transition to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell will be transmitted every period as configured in the <code>loopback-period <i>period</i></code>. If a response is</p>

received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint will transition back to the up state.

The **no** form of the command sets the value back to the default.

**Default** no periodic-loopback

---

## Interface Anti-Spoofing Commands

### anti-spoof

<b>Syntax</b>	<b>anti-spoof</b> { <b>ip</b>   <b>mac</b>   <b>ip-mac</b>   <b>nh-mac</b> } <b>no anti-spoof-type</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	<p>This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the interface.</p> <p>The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (<b>ip</b>, <b>mac</b>, <b>ip-mac</b>, <b>nh-mac</b>) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.</p> <p>The <b>no</b> form of the command disables anti-spoof filtering on the SAP.</p>
<b>Default</b>	<p>Filter type default types:</p> <ul style="list-style-type: none"> <li>• Non-Ethernet encapsulation default anti-spoof filter type — When enabled on a non-Ethernet encapsulated SAP, the anti-spoof filter default type is <b>ip</b>.</li> <li>• Ethernet encapsulated default anti-spoof filter type — When enabled on an Ethernet encapsulated SAP, the anti-spoof default type is <b>ip-mac</b>.</li> <li>• Default anti-spoof filter state — Anti-spoof filtering is disabled by default on the SAP.</li> </ul>
<b>Parameters</b>	<p><b>ip</b> — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type <b>ip</b> command will fail.</p> <p><b>mac</b> — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. Setting the anti-spoof filter type to <b>mac</b> is not allowed on non-Ethernet encapsulated SAPs. If a static host exists on the SAP without a specified MAC address, the anti-spoof type <b>mac</b> command will fail. The anti-spoof type <b>mac</b> command will also fail if the SAP does not support Ethernet encapsulation.</p> <p><b>ip-mac</b> — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. If a static host exists on the SAP without both the IP address and MAC address specified, the anti-spoof type <b>ip-mac</b> command will fail. This is also true if the default anti-spoof filter type of the SAP is <b>ip-mac</b> and the default is not overridden. The anti-spoof type <b>ip-mac</b> command will also fail if the SAP does not support Ethernet encapsulation.</p> <p><b>nh-mac</b> — Indicates that the ingress anti-spoof is based on the source MAC address and the egress anti-spoof is based on the nh-ip-address.</p>



## app-profile

<b>Syntax</b>	<b>app-profile</b> <i>app-profile-name</i> <b>no app-profile</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	This command configures the application profile name.
<b>Parameters</b>	<i>app-profile-name</i> — Specifies an existing application profile name configured in the <b>config&gt;app-assure&gt;group&gt;policy</b> context.

## arp-limit

<b>Syntax</b>	<b>arp-limit</b> <i>limit</i> [ <b>log-only</b> ] [ <b>threshold</b> <i>percent</i> ] <b>no arp-limit</b>
<b>Context</b>	config>service>vprn>interface
<b>Description</b>	This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.  When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.  The <b>no</b> form of the command removes the <b>arp-limit</b> .
<b>Default</b>	90 percent
<b>Parameters</b>	<b>log-only</b> — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.  <i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.  <b>Values</b> 0 — 100  <i>limit</i> — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned.  <b>Values</b> 0 — 524288

## arp-populate

<b>Syntax</b>	<b>[no] arp-populate</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>subscriber-interface config>service>vprn>sub-if>grp-if
<b>Description</b>	This command enables populating static and dynamic hosts into the system ARP cache. When enabled, the host's IP address and MAC address are placed in the system ARP cache as a managed

entry. Static hosts must be defined on the interface using the **host** command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context. In the event that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed. Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and will be repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue.

The **arp-populate** command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.

Once **arp-populate** is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.

**arp-populate** can only be enabled on VPRN interfaces supporting Ethernet encapsulation.

Use the **no** form of the command to disable ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information in the systems ARP cache will be removed. Any existing static ARP entries previously inactive due to static or dynamic hosts will be populated in the system ARP cache.

When **arp-populate** is enabled, the system will not send out ARP Requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with arp-populate enabled.

**Default** not enabled

### arp-retry-timer

**Syntax** **arp-retry-timer** *timer-multiple*  
**no arp-retry-timer**

**Context** config>service>vprn>if  
config>service>vprn>network-interface

**Description** This command allows the arp retry timer to be configured to a specific value. The timer value is entered as a multiple of 100 ms. So a timer value of 1, means the ARP timer will be set to 100 ms. The **no** form of this command removes the command from the active configuration and returns the ARP retry timer to its default value of 5 seconds.

**Default** 5 seconds

**Parameters** *timer-multiple* — specifies the multiple of 100 ms that the ARP retry timer will be configured as.

**Values** 1 — 300 (equally a timer range of 100 ms to 30,000 ms)

## arp-timeout

<b>Syntax</b>	<b>arp-timeout</b> <i>seconds</i> <b>no arp-timeout</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if
<b>Description</b>	This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If <b>arp-timeout</b> is set to a value of zero seconds, ARP aging is disabled.  The <b>no</b> form of this command restores <b>arp-timeout</b> to the default value.
<b>Default</b>	14400 seconds
<b>Parameters</b>	<i>seconds</i> — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.  <b>Values</b> 0 — 65535

## authentication-policy

<b>Syntax</b>	<b>authentication-policy</b> <i>name</i> <b>no authentication-policy</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>sub-if>grp-if
<b>Description</b>	This command assigns an authentication policy to the interface.  The <b>no</b> form of this command removes the policy name from the group interface configuration.
<b>Default</b>	no authentication-policy
<b>Parameters</b>	<i>name</i> — Specifies the authentication policy name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## delayed-enable

<b>Syntax</b>	<b>delayed-enable</b> <i>seconds</i> [ <b>init-only</b> ] <b>no delayed-enable</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if config>service>vprn>nw-if
<b>Description</b>	This command delays making interface operational by the specified number of seconds.  In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has

## Router Advertisement Commands

time to be synchronized, the **delayed-enable** timer can be specified. The optional parameter **init-only** can be added to use this timer only after a reboot.

**Default** no delayed-enable

**Parameters** *seconds* — Specifies the number of seconds to delay before the interface is operational.

**Values** 1 — 1200

**init-only** — Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and/or to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

## calling-station-id

**Syntax** **calling-station-id** *calling-station-id*  
**no calling-station-id**

**Context** config>service>vprn>sub-if>grp-if>sap  
config>service>vprn>if>sap

**Description** This command enables the inclusion of the **calling-station-id** attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no value is set at the SAP level, an empty string is included.

**Default** This attribute is not sent by default.

## host

**Syntax** [**no**] **host** {[**ip** *ip-address* [**mac** *ieee-address*]} [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]  
**no host** {[**ip** *ip-address*] [**mac** *ieee-address*]}

**Context** config>service>vprn>if>sap

**Description** This command creates a static host for the SAP. Applications within the system that make use of static host entries include anti-spoof, and source MAC population into the VPLS forwarding database.

Multiple static hosts can be defined on the SAP. Each host is identified by a source IP address, a source MAC address, or both a source IP and source MAC address. When anti-spoof is enabled on the SAP, the host information will be populated into the SAP's anti-spoof table, allowing ingress packets matching the entry access to the SAP. When the MAC address exists in the host definition, the MAC address is populated into the VPLS forwarding database and associates it with the SAP. The static host definition overrides any static MAC entries using the same MAC and prevents dynamic learning of the MAC on another interface.

Defining a static host identical to an existing static host has no effect and will not generate a log or error message.

Every static host definition must have at least one address defined, IP or MAC.

Static hosts may exist on the SAP even with anti-spoof and **arp-populate** (VPRN) features disabled. When enabled, each feature has different requirements for static hosts.

**anti-spoof** — When enabled, this feature uses static and dynamic host information to populate entries into an anti-spoof filter table. The anti-spoof filter entries generated will be of the same type as specified in the anti-spoof type parameter. If the SAP anti-spoof filter is defined as *mac*, each static host definition must specify a MAC address. If the SAP anti-spoof filter is defined as *ip*, each static host definition must specify an IP address. If the SAP anti-spoof filter is defined as *ip-mac*, each static host definition must specify both an IP address and MAC address. If definition of a static host is attempted without the appropriate addresses specified for the enabled anti-spoof filter, the static host definition will fail.

**arp-populate** — When enabled, this feature uses static and dynamic host information to populate entries into the system's ARP cache. This is only available on the VPRN service SAPs. Both a MAC address and IP address are required to populate an ARP entry in the system. If definition of a static host is attempted without both a MAC and IP address specified when *arp-populate* is enabled, the static host definition will fail.

**fdb-populate** — This is an implicit feature that uses the static host definition as a static MAC in the VPLS forwarding database. It cannot be enabled or disabled and has no effect on the ability to create static hosts without a MAC address specified. When a MAC address is specified for a static host, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is created. The static host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning of the MAC on another interface. Existing static MAC entries with the same MAC address as a static host are marked as inactive but not deleted. If all static hosts are removed from the SAP, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a static host exists associated with the static MAC address.

The **no** form of the command removes a static entry from the system. The specified **ip address** and **mac address** must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the affect of its removal on the anti-spoof filter, ARP cache or the VPLS forwarding database is also evaluated.

#### Default

There are no default static entries.

#### Parameters

**ip ip-address** — Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip** and **anti-spoof ip-mac** commands. Only one static host can be configured on the SAP with a given IP address.  
The following rules apply to configure static hosts using an IP address:

- Only one static host can be defined using a specific IP address.
- Defining a static host with the same IP address as a previous static host overwrites the previous static host.
- If a static host has an IP address assigned, the MAC address for the host is optional (depending on the features enabled on the SAP).

**mac mac-address** — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac**, and **anti-spoof ip-mac**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address. The following rules apply to configuring static hosts using a MAC address:

- Multiple static hosts may share the same MAC address.
- Executing the host command with the same MAC address but a different IP address as an existing static host will create a new static host.

## Router Advertisement Commands

- If a static host has a MAC address assigned, the IP address for the host is optional (depending on the features enabled on the SAP).

**Values** 8k static and dynamic hosts per 10G forwarding complex. 64k8k per system.

**subscriber** *sub-ident-string* — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

- For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's *sub-ident-string* is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's split horizon group.

**sub-profile** *sub-profile-name* — Specify this optional parameter to specify an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

**sla-profile** *sla-profile-name* — Specify this optional parameter to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

## frame-relay

<b>Syntax</b>	<b>frame-relay</b>
<b>Context</b>	config>service>vprn>if>sap
<b>Description</b>	This command enables the context to configure Frame Relay parameters on the SAP.

## frf-12

<b>Syntax</b>	<b>[no] frf-12</b>
<b>Context</b>	config>service>vprn>if>sap

**Description** This command enables the use of FRF12 headers.  
The **no** form of the command disables the use of FRF12 headers.

## ete-fragment-threshold

**Syntax** **ete-fragment-threshold** *threshold*  
**no ete-fragment-threshold**

**Context** config>service>vprn>if>sap>frf-12

**Description** This command specifies the maximum length of a fragment to be transmitted.  
The **no** form of the command reverts to the default.

**Parameters** *threshold* — Specifies the maximum length of a fragment to be transmitted.

<b>Values</b>	128 — 512
<b>Default</b>	0

## interleave

**Syntax** [**no**] **interleave**

**Context** config>service>vprn>if>sap>frame-relay>frf.12

**Description** This command enables interleaving of high priority frames and low-priority frame fragments within a FR SAP using FRF.12 end-to-end fragmentation.

When this option is enabled, only frames of the FR SAP non expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI).

When this option is disabled, frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is however not included when the frame size is smaller than the user configured fragmentation size. In this mode, the SAP transmits all fragments of a frame before sending the next full or fragmented frame.

The receive direction of the FR SAP supports both modes of operation concurrently, with and without fragment interleaving.

The **no** form of this command restores the default mode of operation.

**Default** no interleave

## scheduling-class

**Syntax** **scheduling-class** *class-id*

**Context** config>service>vprn>if>sap

## Router Advertisement Commands

<b>Description</b>	This command specifies the scheduling class to use for this SAP.
<b>Parameters</b>	<i>class-id</i> — Specifies the scheduling class to use for this SAP.
<b>Values</b>	0 — 3
<b>Default</b>	0

## flowspec

<b>Syntax</b>	<b>flowspec</b> <b>no flowspec</b>
<b>Context</b>	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>sap>ingress config>service>ies>interface>spoke-sdp>ingress
<b>Description</b>	<p>This command enables IPv4 flowspec filtering on an access IP interface associated with a VPRN or IES service. Filtering is based on all of the IPv4 flowspec routes that have been received and accepted by the corresponding BGP instance. Ingress IPv4 traffic on an interface can be filtered by both a user-defined IPv4 filter and flowspec. Evaluation proceeds in this order:</p> <ol style="list-style-type: none"><li>1. user-defined IPv4 filter entries</li><li>2. flowspec-derived filter entries</li><li>3. user-defined IPv4 filter default-action</li></ol> <p>The <b>no</b> form of the command removes IPv4 flowspec filtering from an IP interface.</p>
<b>Default</b>	No access interfaces have IPv4 flowspec enabled.

## flowspec-ipv6

<b>Syntax</b>	<b>flowspec-ipv6</b> <b>no flowspec-ipv6</b>
<b>Context</b>	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>ies>interface>sap>ingress config>service>ies>interface>spoke-sdp>ingress
<b>Description</b>	<p>This command enables IPv6 flowspec filtering on an access IP interface associated with a VPRN or IES service. Filtering is based on all of the IPv6 flowspec routes that have been received and accepted by the corresponding BGP instance. Ingress IPv6 traffic on an interface can be filtered by both a user-defined IPv6 filter and flowspec. Evaluation proceeds in this order:</p> <ol style="list-style-type: none"><li>1. user-defined IPv6 filter entries</li><li>2. flowspec-derived filter entries</li><li>3. user-defined IPv6 filter default-action</li></ol> <p>The <b>no</b> form of the command removes IPv6 flowspec filtering from an IP interface.</p>



**Default** No access interfaces have IPv6 flowspec enabled.

## host-lockout-policy

**Syntax** **host-lockout-policy** *policy-name*  
**no host-lockout-policy**

**Context** config>service>vprn>if>sap

**Description** This command configures a host lockout policy.  
The no form of the command removes the policy name from the configuration.

## host-shutdown

**Syntax** [**no**] **host-shutdown**

**Context** config>service>vprn>if>sap

This command administratively enables host creation on this SAP.

## ip-tunnel

**Syntax** **ip-tunnel** *name* [**create**]  
**no ip-tunnel** *name*

**Context** config>service>vprn>if>sap

**Description** This command is used to configure an IP-GRE or IP-IP tunnel and associate it with a private tunnel SAP within an IES or VPRN service.  
The no form of the command deletes the specified IP/GRE or IP-IP tunnel from the configuration. The tunnel must be administratively shutdown before issuing the no ip-tunnel command.

**Default** No IP tunnels are defined.

**Parameters** *ip-tunnel-name* — Specifies the name of the IP tunnel. Tunnel names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## backup-remote-ip

**Syntax** **backup-remote-ip** *ip-address*  
**no backup-remote-ip**

**Context** config>service>interface>vprn>sap>ip-tunnel

## Router Advertisement Commands

- Description** This command sets the backup destination IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. If the primary destination address is not reachable in the delivery service (there is no route) or not defined then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The **no** form of the command deletes the backup-destination address from the GRE tunnel configuration.
- Parameters** *ip-address* — Specifies the destination IPv4 address of the GRE tunnel.
- Values** 1.0.0.0 — 223.255.255.255

## delivery-service

- Syntax** **delivery-service** {*service-id* | *svc-name*}  
**no delivery-service**
- Context** config>service>interface>vprn>sap>ip-tunnel
- Description** This command sets the delivery service for GRE encapsulated packets associated with a particular GRE tunnel. This is the IES or VPRN service where the GRE encapsulated packets are injected and terminated. The delivery service may be the same service that owns the private tunnel SAP associated with the GRE tunnel. The GRE tunnel does not come up until a valid delivery service is configured.
- The **no** form of the command deletes the delivery-service from the GRE tunnel configuration.
- Parameters** *service-id* — Identifies the service used to originate and terminate the GRE encapsulated packets belonging to the GRE tunnel.
- Values** 1—2147483648
- svc-name* — Identifies the service used to originate and terminate the GRE encapsulated packets belonging to the GRE tunnel.
- Values** 1—64 characters

## dscp

- Syntax** **dscp** *dscp-name*  
**no dscp**
- Context** config>service>interface>vprn>sap>ip-tunnel
- Description** This command sets the DSCP code-point in the outer IP header of GRE encapsulated packets associated with a particular GRE tunnel. The default, set using the no form of the command, is to copy the DSCP value from the inner IP header (after remarking by the private tunnel SAP egress qos policy) to the outer IP header.
- Default** no dscp
- Parameters** *dscp* — Specifies the DSCP code-point to be used.

**Values** be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

## source

<b>Syntax</b>	<b>source</b> <i>ip-address</i> <b>no source</b>
<b>Context</b>	config>service>interface>vprn>sap>ip-tunnel
<b>Description</b>	This command sets the source IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. It must be an address in the subnet of the associated public tunnel SAP interface. The GRE tunnel does not come up until a valid source address is configured.  The <b>no</b> form of the command deletes the source address from the GRE tunnel configuration. The tunnel must be administratively shutdown before issuing the <b>no source</b> command.
<b>Parameters</b>	<i>ip-address</i> — Specifies the source IPv4 address of the GRE tunnel.  <b>Values</b> 1.0.0.0 — 223.255.255.255

## remote-ip

<b>Syntax</b>	<b>remote-ip</b> <i>ip-address</i> <b>no remote-ip</b>
<b>Context</b>	config>service>interface>vprn>sap>ip-tunnel
<b>Description</b>	This command sets the primary destination IPv4 address of GRE encapsulated packets associated with a particular GRE tunnel. If this address is reachable in the delivery service (there is a route) then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.  The <b>no</b> form of the command deletes the destination address from the GRE tunnel configuration.
<b>Parameters</b>	<i>ip-address</i> — Specifies the destination IPv4 address of the GRE tunnel.  <b>Values</b> 1.0.0.0 — 223.255.255.255

---

## Interface SAP Filter and QoS Policy Commands

### egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap config>service>vprn>ipsec-if>sap
<b>Description</b>	This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.  If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

### ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap config>service>vprn>if>sap config>service>vprn>ipsec-if>sap
<b>Description</b>	This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.  If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

### agg-rate

<b>Syntax</b>	<b>[no] agg-rate</b>
<b>Context</b>	config>service>vprn>interface>sap>egress
<b>Description</b>	This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: <b>rate</b> , <b>limit-unused-bandwidth</b> , and <b>queue-frame-based-accounting</b> .

## rate

<b>Syntax</b>	<b>rate</b> {max   rate} <b>no rate</b>
<b>Context</b>	config>service>vprn>interface>sap>egress>agg-rate
<b>Description</b>	This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, VPORT etc.).

## limit-unused-bandwidth

<b>Syntax</b>	<b>[no] limit-unused-bandwidth</b>
<b>Context</b>	config>service>vprn>interface>sap>egress>agg-rate
<b>Description</b>	This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

## queue-frame-based-accounting

<b>Syntax</b>	<b>[no] queue-frame-based-accounting</b>
<b>Context</b>	config>service>vprn>interface>sap>egress>agg-rate
<b>Description</b>	This command is used to enabled (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

## agg-rate-limit

<b>Syntax</b>	<b>agg-rate-limit</b> <i>agg-rate</i> [ <b>queue-frame-based-accounting</b> ] <b>no agg-rate-limit</b>
<b>Context</b>	
<b>Description</b>	<p>This command defines a maximum total rate for all egress queues on a service SAP or multi-service site. The <b>agg-rate-limit</b> command is mutually exclusive with the egress scheduler policy. When an egress scheduler policy is defined, the <b>agg-rate-limit</b> command will fail. If the <b>agg-rate-limit</b> command is specified, an attempt to bind a <b>scheduler-policy</b> to the SAP or multi-service site will fail.</p> <p>A multi-service site must have a port scope defined that ensures all queues associated with the site are on the same port or channel. If the scope is not set to a port, the agg-rate-limit command will fail. Once an agg-rate-limit has been assigned to a multi-service site, the scope cannot be changed to card level.</p> <p>A port scheduler policy must be applied on the egress port or channel the SAP or multi-service site is bound to in order for the defined agg-rate-limit to take effect. The egress port scheduler enforces the</p>

## Router Advertisement Commands

aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.

If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.

The **no** form of the command removes the aggregate rate limit from the SAP or multi-service site.

**Parameters** *agg-rate* — Defines the rate, in kilobits-per-second, that the maximum aggregate rate that the queues on the SAP or multi-service site can operate.

**Values** 1 — 40000000, max

**queue-frame-based-accounting** — This keyword enables frame based accounting on all queues associated with the SAP or Multi-Service Site. If frame based accounting is required when an aggregate limit is not necessary, the max keyword should precede the queue-frame-based-accounting keyword. If frame based accounting must be disabled, execute *agg-rate-limit* without the queue-frame-based-accounting keyword present.

**Default** Frame based accounting is disabled by default

## filter

**Syntax** **filter ip** *ip-filter-id*  
**no filter**

**Context** config>service>vprn>if>sap>egress  
config>service>vprn>if>sap>ingress  
config>service>vprn>sub-if>grp-if>sap>egress  
config>service>vprn>ipsec-if>sap>egress  
config>service>vprn>ipsec-if>sap>ingress

**Description** This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. Filter policies control the forwarding and dropping of packets based on IP matching criteria.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress SAP. The *ip-filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

**Parameters** **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

**Values** 1 — 65535

## flowspec

**[no] flowspec**

<b>Context</b>	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress config>service>vprn>network-interface>ingress
<b>Description</b>	<p>This command enables flowspec filtering on an IP interface of a VPRN. Filtering is based on all of the flowspec routes that have been received and accepted by the VPRN. Ingress traffic on an IP interface can be filtered by both a user-defined ip filter and flowspec. Evaluation proceeds in this order:</p> <ol style="list-style-type: none"> <li>1. user-defined ip filter entries with entry numbers less than the configured insert-point</li> <li>2. flowspec-derived filter entries</li> <li>3. user-defined ip filter entries with entry numbers greater than or equal to the configured insert-point</li> <li>4. ip-filter default-action</li> </ol> <p>The <b>no</b> form of the command removes flowspec filtering from an IP interface.</p>
<b>Default</b>	No interfaces have flowspec enabled.

## flowspec-ipv6

<b>Syntax</b>	<b>[no] flowspec-ipv6</b>
<b>Context</b>	config>service>vprn>interface>sap>ingress config>service>vprn>interface>spoke-sdp>ingress
<b>Description</b>	<p>This command enables flowspec filtering on an IP interface of the base router. Filtering is based on all of the flowspec routes that have been received and accepted by the base router. Ingress traffic on an IP interface can be filtered by both a user-defined ip filter and flowspec. In this case, the user-defined ip filter entries are evaluated before the flowspec routes and the default action of the user-defined ip filter applies as the very last rule.</p> <p>The <b>no</b> form of the command removes flowspec filtering from an IP interface.</p>
<b>Default</b>	No interfaces have flowspec enabled.

## hsmdda-queue-override

<b>Syntax</b>	<b>[no] hsmdda-queue-override</b>
<b>Context</b>	config>service>vprn>if>sap>egress
<b>Description</b>	This command enables the context to configure HSMDDA queue overrides.

### queue

<b>Syntax</b>	<b>queue</b> <i>queue-id</i> [ <b>create</b> ] <b>no queue</b> <i>queue-id</i>
<b>Context</b>	config>service>vprn>if>sap>egress>hsmda-queue-overider
<b>Description</b>	This command configures overrides for a HSMDA queue. The actual valid values are those defined in the given SAP QoS policy.
<b>Parameters</b>	<i>queue-id</i> — Specifies the queue ID to override. <b>Values</b> 1 — 8 <b>create</b> — This keyword is mandatory while creating a new queue override.

### packet-byte-offset

<b>Syntax</b>	<b>packet-byte-offset</b> { <b>add</b> <i>add-bytes</i>   <b>subtract</b> <i>sub-bytes</i> } <b>no packet-byte-offset</b>
<b>Context</b>	config>service>vprn>if>sap>egress>hsmda-queue-overider
<b>Description</b>	<p>This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4 byte CRC (everything except the preamble and inter-frame gap). As an example, the packet-byte-offset command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.</p> <p>The accounting functions affected include:</p> <ul style="list-style-type: none"><li>• Offered High Priority / In-Profile Octet Counter</li><li>• Offered Low Priority / Out-of-Profile Octet Counter</li><li>• Discarded High Priority / In-Profile Octet Counter</li><li>• Discarded Low Priority / Out-of-Profile Octet Counter</li><li>• Forwarded In-Profile Octet Counter</li><li>• Forwarded Out-of-Profile Octet Counter</li><li>• Peak Information Rate (PIR) Leaky Bucket Updates</li><li>• Committed Information Rate (CIR) Leaky Bucket Updates</li><li>• Queue Group Aggregate Rate Limit Leaky Bucket Updates</li></ul> <p>The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.</p> <p>The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. Up to 31 bytes may be added to the packet and up to 32 bytes may be removed from the packet. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation,</p>



the command `packet-byte-offset subtract 14` would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the `packet-byte-offset` command is targeted at the queue and queue group level. The `packet-byte-offset`, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber's packets on an Ethernet aggregation network.

The `packet-byte-offset` value may be overridden at the queue-group level.

<b>Parameters</b>	<p><b>add</b> <i>add-bytes</i> — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the <b>add</b> or <b>subtract</b> keyword must be specified. The corresponding byte value must be specified when executing the <code>packet-byte-offset</code> command. The <b>add</b> keyword is mutually exclusive with the <b>subtract</b> keyword.</p> <p><b>Values</b>      0 — 31</p> <p><b>subtract</b> <i>sub-bytes</i> — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The <b>subtract</b> keyword is mutually exclusive with the <b>add</b> keyword. Either the <b>add</b> or <b>subtract</b> keyword must be specified. The corresponding byte value must be specified when executing the <code>packet-byte-offset</code> command.</p> <p><b>Values</b>      1 — 64</p>
-------------------	---

## slope-policy

<b>Syntax</b>	<b>slope-policy</b> <i>hsmda-slope-policy-name</i> <b>no slope-policy</b>
<b>Context</b>	config>service>vprn>if>sap>egress>hsmda-queue-overider
<b>Description</b>	This command specifies an existing slope policy name.

## wrr-weight

<b>Syntax</b>	<b>wrr-weight</b> <i>value</i> <b>no wrr-weight</b>
<b>Context</b>	config>service>vprn>if>sap>egress>hsmda-queue-overider>queue
<b>Description</b>	This command assigns the weight value to the HSMDA queue.  The <b>no</b> form of the command returns the weight value for the queue to the default value.
<b>Parameters</b>	<i>percentage</i> — Specifies the weight for the HSMDA queue.  <b>Values</b> 1 — 32

## wrr-policy

<b>Syntax</b>	<b>wrr-policy</b> <i>hsmda-wrr-policy-name</i> <b>no wrr-policy</b>
<b>Context</b>	config>service>vprn>if>sap>egress>hsmda-queue-overider
<b>Description</b>	This command associates an existing HSM DA weighted-round-robin (WRR) scheduling loop policy to the HSM DA queue.
<b>Parameters</b>	<i>hsmda-wrr-policy-name</i> — Specifies the existing HSM DA WRR policy name to associate to the queue.

## secondary-shaper

	<b>secondary-shaper</b> <i>secondary-shaper-name</i> <b>no secondary-shaper</b>
<b>Context</b>	config>service>vprn>if>sap>egress>hsmda-queue-overider
<b>Description</b>	This command configures an HSM DA secondary shaper. Note that an shaper override can only be configured on an HSM DA SAP.
<b>Parameters</b>	<i>secondary-shaper-name</i> — Specifies a secondary shaper name up to 32 characters in length.

## match-qinq-dot1p

<b>Syntax</b>	<b>match-qinq-dot1p</b> { <b>top</b>   <b>bottom</b> } <b>no match-qinq-dot1p</b>
<b>Context</b>	config>service>vprn>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>ingress config>service>vprn>ipsec-if>sap>ingress
<b>Description</b>	<p>This command specifies which Dot1Q tag position Dot1P bits in a QinQ encapsulated packet should be used to evaluate Dot1P QoS classification.</p> <p>The <b>match-qinq-dot1p</b> command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy’s Dot1P entries. The <b>top</b> and <b>bottom</b> keywords specify which position should be evaluated for QinQ encapsulated packets.</p> <p>The <b>no</b> form of the command restores the default dot1p evaluation behavior for the SAP.</p> <p>By default, the bottom most service delineating Dot1Q tags Dot1P bits are used. The following table defines the default behavior for Dot1P evaluation when the <b>match-qinq-dot1p</b> command is not executed.</p>

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

**Default** no match-qinq-dot1p - No filtering based on p-bits.  
top or bottom must be specified to override the default QinQ dot1p behavior.

**Parameters** **top** — The top parameter is mutually exclusive to the bottom parameter. When the top parameter is specified, the top most PBits are used (if existing) to match any dot1p dot1p-value entries. The following table defines the dot1p evaluation behavior when the top parameter is specified.

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	TopQ PBits

**bottom** — The bottom parameter is mutually exclusive to the top parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-

value entries. The following tables define the bottom position QinQ and TopQ SAP dot1p evaluation and the default dot1p explicit marking actions.

**Table 14: Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

<b>Port / SAP Type</b>	<b>Existing Packet Tags</b>	<b>PBits Used for Match</b>
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	BottomQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ / TopQ	TopQ	TopQ PBits
QinQ / TopQ	TopQ BottomQ	BottomQ PBits
QinQ / QinQ	TopQ BottomQ	BottomQ PBits

<b>Egress SAP Type</b>	<b>Ingress Packet Preserved Dot1P State</b>	<b>Marked (or Remarked) PBits</b>
Null	No preserved Dot1P bits	None
Null	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
Dot1Q	No preserved Dot1P bits	New PBits marked using dot1p-value
Dot1Q	Preserved Dot1P bits	Preserved tag PBits remarked using dot1p-value
TopQ	No preserved Dot1P bits	TopQ PBits marked using dot1p-value
TopQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits marked using dot1p-value, BottomQ PBits preserved
QinQ	No preserved Dot1P bits	TopQ PBits and BottomQ PBits marked using dot1p-value
QinQ	Preserved Dot1P bits (used as TopQ and BottomQ PBits)	TopQ PBits and BottomQ PBits marked using dot1p-value

The `dot1p dot1p-value` command must be configured without the `qinq-mark-top-only` parameter to remove the TopQ PBits only marking restriction.

## qinq-mark-top-only

<b>Syntax</b>	<code>[no] qinq-mark-top-only</code>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;if&gt;sap&gt;egress config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;sap&gt;engress config&gt;service&gt;vprn&gt;ipsec-if&gt;sap&gt;egress</pre>
<b>Description</b>	When enabled (the encapsulation type of the access port where this SAP is defined as qinq), the <b>qinq-mark-top-only</b> command specifies which P-bits/DEI bit to mark during packet egress. When disabled, both set of P-bits/DEI bit are marked. When the enabled, only the P-bits/DEI bit in the top Q-tag are marked.
<b>Default</b>	no qinq-mark-top-only

## qos

<b>Syntax</b>	<code>qos policy-id [port-redirect-group queue-group-name instance instance-id]</code> <code>no qos</code>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;if&gt;sap&gt;egress config&gt;service&gt;vprn&gt;sub-if&gt;grp-if&gt;sap&gt;engress config&gt;service&gt;vprn&gt;ipsec-if&gt;sap&gt;egress</pre>
<b>Description</b>	<p>This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP).</p> <p>QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the policy- id does not exist, an error will be returned.</p> <p>The <code>qos</code> command is used to associate both ingress and egress QoS policies. The <code>qos</code> command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.</p> <p>Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.</p> <p>By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p> <p>The <b>no</b> form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.</p>

## Router Advertisement Commands

<b>Default</b>	none				
<b>Parameters</b>	<p><i>policy-id</i> — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.</p> <p>1 — 65535</p> <p><b>port-redirect-group</b> — This keyword associates a SAP egress with an instance of a named queue group template on the egress port of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command.</p> <p><i>queue-group-name</i> — Specifies the name of the egress port queue group of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid egress queue group, created under config&gt;port&gt;ethernet&gt;access&gt;egress.</p> <p><b>instance</b> <i>instance-id</i> — Specifies the instance of the named egress port queue group on the IOM/IMM/XMA.</p> <table><tr><td><b>Values</b></td><td>1 — 40960</td></tr><tr><td><b>Default</b></td><td>1</td></tr></table>	<b>Values</b>	1 — 40960	<b>Default</b>	1
<b>Values</b>	1 — 40960				
<b>Default</b>	1				

## qos

<b>Syntax</b>	<b>qos</b> <i>policy-id</i> [ <b>shared-queuing</b>   <b>multipoint-shared</b> ] <b>fp-redirect-group</b> <i>queue-group-name</i> <b>instance</b> <i>instance-id</i> <b>no qos</b>
<b>Context</b>	config>service>vprn>if>sap>ingress config>service>vprn>sub-if>grp-if>sap>ingress config>service>vprn>ipsec-if>sap>ingress
<b>Description</b>	<p>This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).</p> <p>QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy- id does not exist, an error will be returned.</p> <p>The qos command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.</p> <p>Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.</p> <p>By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.</p> <p>The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.</p>
<b>Default</b>	none

**Parameters** *policy-id* — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

**Values** 1 — 65535

**shared-queuing** — Specifies the ingress shared queue policy used by this SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues instead of the shared ones.

**multipoint-shared** — This keyword specifies that this *queue-id* is for multipoint forwarded traffic only. This *queue-id* can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.

A queue must be created as **multipoint**. The **multipoint** designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the **multipoint** keyword, an error is generated and the command will not execute.

The **multipoint** keyword can be entered in the command line on a pre-existing multipoint queue to edit *queue-id* parameters.

**Default** Present (the queue is created as non-multipoint).

**Values** **Multipoint** or not present.

**fp-redirect-group** — This keyword creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM/XMA. The queue-group-name and instance instance-id are mandatory parameters when executing the command. The named queue group template can contain only policers. If it contains queues, then the command will fail.

*queue-group-name* — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters in length. The queue-group-name must correspond to a valid ingress queue group template name, configured under *config>qos>queue-group-templates*.

*instance-id* — Specifies the instance of the named queue group to be created on the IOM/IMM/XMA ingress forwarding plane.

## scheduler-policy

**Syntax** **scheduler-policy** *scheduler-policy-name*  
no scheduler-policy

**Context** config>service>vprn>if>sap>ingress  
config>service>vprn>if>sap>egress  
config>service>vprn>sub-if>grp-if>sap>engress  
config>service>vprn>sub-if>grp-if>sap>ingress  
config>service>vprn>ipsec-if>sap>egress  
config>service>vprn>ipsec-if>sap>ingress

**Description** This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port,

## Router Advertisement Commands

channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues associated with the customer site. Queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

*scheduler-policy-name*: — The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.

**Values** Any existing valid scheduler policy name.

### ipsec-gw

<b>Syntax</b>	<b>ipsec-gw</b> <i>name</i> <b>no ipsec-gw</b>
<b>Context</b>	config>service>vprn>if>sap
<b>Description</b>	This command configures the IPSec gateway.
<b>Parameters</b>	<i>name</i> — Specifies the IPSec gateway name up to 32 characters in length.

### cert

<b>Syntax</b>	<b>cert</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw
<b>Description</b>	This command enables the context to configure certificate parameters.

### cert

<b>Syntax</b>	<b>cert</b> <i>filename</i> <b>no cert</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw>cert
<b>Description</b>	This command configures the cert with local file URL.
<b>Default</b>	none



**Parameters** *filename* — Specifies the local file URL of the certificate to be used with this SAP IPsec tunnel.

## cert-profile

**Syntax** **cert-profile** *profile-name*  
**no cert-profile**

**Context** config>service>vprn>if>sap>ipsec-gw>cert  
config>service>vprn>if>sap>ipsec-tun>dyn>cert

**Description** This command specifies the cert-profile for the ipsec-tunnel or ipsec-gw. This command will override “cert” and “key” configuration under the ipsec-tunnel or ipsec-gw.

**Default** none

**Parameters** *profile-name* — Specifies the name of cert-profile.

## key

**Syntax** **key** *filename*  
**no key**

**Context** config>service>vprn>if>sap>ipsec-gw>cert

**Description** This command configures the key-pair file to be used for X.509 certificate authentication with this SAP IPsec tunnel.

**Default** none

**Parameters** *filename* — Specifies a key with the CA profile.

## trust-anchor

**Syntax** **trust-anchor** *ca-profile-name*  
**no trust-anchor**

**Context** config>service>vprn>if>sap>ipsec-gw>cert

**Description** This command configures the Certificate-Authority Profile name associated with this SAP IPsec tunnel certificate.

**Default** none

**Parameters** *ca-profile-name* — Specify a CA profile name up to 32 characters in length.

## default-secure-service

<b>Syntax</b>	<b>default-secure-service</b> <i>service-id ipsec-interface ip-int-name</i> <b>no default-secure-service</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw
<b>Description</b>	This command specifies a service ID or service name of the default security service used by this SAP IPsec gateway.
<b>Parameters</b>	<i>service-id</i> — Specifies a default secure service.
	<b>Values</b> <i>service-id</i> : 1 — 2147483648 <i>svc-name</i> : Specifies an existing service name up to 64 characters in length.

## default-tunnel-template

<b>Syntax</b>	<b>default-tunnel-template</b> <i>ipsec template identifier</i> <b>no default-tunnel-template</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw
<b>Description</b>	This command configures the default tunnel policy template for the gateway.
<b>Parameters</b>	<i>ipsec template id*</i> — [1..2048]

## ike-policy

<b>Syntax</b>	<b>ike-policy</b> <i>ike-policy-id</i> <b>no ike-policy</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw config>service>vprn>ipsec-if>sap>tunnel>dynamic-keying
<b>Description</b>	This command configures the IKE policy for the gateway.
<b>Parameters</b>	<i>ike-policy-id</i> — Specifies the IKE policy ID.
	<b>Values</b> 1 — 2048

## local-gateway-address

<b>Syntax</b>	<b>local-gateway-address</b> <i>ip-address</i> <b>no local-gateway-address</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw
<b>Description</b>	This command configures the ipsec-gateway local address.
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP unicast address.

## local-id

<b>Syntax</b>	<b>local-id type</b> {ipv4 fqdn} [value [value]] <b>no local-id</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw
<b>Description</b>	This command specifies the local ID of the router used for IDi or IDr for IKEv2 tunnels. The local-id can only be changed or removed when tunnel or gateway is shutdown.  Default: Depends on local-auth-method such as: <ul style="list-style-type: none"> <li>• Psk:local tunnel ip address</li> <li>• Cert-auth: subject of the local certificate</li> </ul>
<b>Parameters</b>	<b>type</b> — Specifies the type of local ID payload, it could be ipv4 address/FQDN domain name/distinguish name of subject in X.509 certificate.  <b>Values</b> <ul style="list-style-type: none"> <li>ipv4 — Use ipv4 as the local ID type, the default value is the local tunnel end-point address.</li> <li>fqdn — Use FQDN as the local ID type, the value must be configured.</li> <li>dn — Use the subject of the certificate configured for the tunnel or gateway.</li> </ul>

## pre-shared-key

<b>Syntax</b>	<b>pre-shared-key key</b> <b>no pre-shared-key</b>
<b>Context</b>	config>service>vprn>if>sap>ipsec-gw config>service>vprn>ipsec-if>sap>tunnel>dynamic-keying
<b>Description</b>	This command specifies the shared secret between the two peers forming the tunnel.
<b>Parameters</b>	<i>key</i> — Specifies a pre-shared-key for dynamic-keying.

## radius-accounting-policy

<b>Syntax</b>	<b>radius-accounting-policy policy-name</b> <b>no radius-accounting-policy</b>
<b>Context</b>	
<b>Description</b>	

## radius-authentication-policy

<b>Syntax</b>	<b>radius-authentication-policy</b> <i>name</i> <b>no radius-authentication-policy</b>
<b>Context</b>	
<b>Description</b>	

## lag-link-map-profile

<b>Syntax</b>	<b>lag-link-map-profile</b> <i>lnk-map-profile-id</i> <b>no lag-link-map-profile</b>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn> sub-if>grp-if >sap
<b>Description</b>	<p>This command assigns a pre-configured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/de-assigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration.</p> <p>The <b>no</b> form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.</p>
<b>Default</b>	<b>no lag-link-map-profile</b>
<b>Parameters</b>	<i>link-map-profile-id</i> — An integer from 1 to 64 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

## multi-service-site

<b>Syntax</b>	<b>multi-service-site</b> <i>customer-site-name</i> <b>no multi-service-site</b> <i>customer-site-name</i>
<b>Context</b>	config>service>vprn>if>sap config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	<p>This command creates a new customer site or edits an existing customer site with the <i>customer-site-name</i> parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).</p> <p>The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.</p>

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

**Default** None — Each customer site must be explicitly created.

**Parameters** *customer-site-name*: — Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

- The maximum number of customer sites defined for the chassis slot has not been met.
- The *customer-site-name* is valid.
- The **create** keyword is included in the command line syntax (if the system requires it).

When the maximum number of customer sites has been exceeded a configuration error occurs; the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs; the command will not execute and the CLI context will not change.

**Values** Valid names consist of any string up to 32 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.

## static-host

**Syntax** **static-host ip** *ip/did-address* [**mac** *ieee-address*] [**create**]  
**static-host mac** *ieee-address* [**create**]  
**no static-host** [**ip** *ip-address*>] **mac** *ieee-address*>  
**no static-host all** [**force**]  
**no static-host ip** *ip-address*

**Context** config>service>vprn>if>sap  
config>service>vprn>sub-if>grp-if>sap

**Description** This command configures a static host on this SAP.

**Syntax** **ip** *ip-address* — Specifies the IPv4 unicast address.

**mac** *ieee-address* — Specify this optional parameter when defining a static host. Every static host definition must have at least one address defined, IP or MAC.

**force** — Specifies the forced removal of the static host addresses.

*sla-profile sla-profile-name*

## Router Advertisement Commands

This optional parameter is used to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

### ancp-string

<b>Syntax</b>	<b>ancp-string</b> <i>ancp-string</i> <b>no ancp-string</b>
<b>Context</b>	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
<b>Description</b>	This command specifies the ANCP string associated to this SAP host.
<b>Parameters</b>	<i>ancp-string</i> — Specifies the ANCP string up to 63 characters in length.

### app-profile

<b>Syntax</b>	<b>app-profile</b> <i>app-profile-name</i> <b>no app-profile</b>
<b>Context</b>	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
<b>Description</b>	This command specifies an application profile name.
<b>Parameters</b>	<i>app-profile-name</i> — Specifies the application profile name up to 32 characters in length.

### inter-dest-id

<b>Syntax</b>	<b>inter-dest-id</b> <i>intermediate-destination-id</i> <b>no inter-dest-id</b>
<b>Context</b>	config>service>vprn>if>sap>static-host config>service>vprn>sub-if>grp-if>sap>static-host
<b>Description</b>	This command specifies to which intermediate destination (for example a DSLAM) this host belongs.
<b>Parameters</b>	<i>intermediate-destination-id</i> — Specifies the intermediate destination ID.

### managed-routes

<b>Syntax</b>	<b>managed-routes</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes
<b>Description</b>	This command configures managed routes.

## route

<b>Syntax</b>	<b>route</b> <i>{ip-prefix/length   ip-prefix netmask}</i> [ <b>create</b> ] <b>no route</b> <i>{ip-prefix/length   ip-prefix netmask}</i>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes
<b>Description</b>	This command assigns managed-route to a given subscriber-host. As a consequence, a static-route pointing subscriber-host ip address as a next hop will be installed in FIB. Up to 16 managed routes per subscriber-host can be configured.  The <b>no</b> form of the command removes the respective route. Per default, there are no managed-routes configured.

## sla-profile

<b>Syntax</b>	<b>sla-profile</b> <i>sla-profile-name</i> <b>no sla-profile</b>
<b>Context</b>	config>service>vprn>if>sap>static-host
<b>Description</b>	This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context.
<b>Parameters</b>	<i>sla-profile-name</i> — Specifies the SLA profile name.

## sub-profile

<b>Syntax</b>	<b>sub-profile</b> <i>sub-profile-name</i> <b>no sub-profile</b>
<b>Context</b>	config>service>vprn>if>sap>static-host
<b>Description</b>	This command specifies an existing subscriber profile name to be associated with the static subscriber host.
<b>Parameters</b>	<i>sub-profile-name</i> — Specifies the sub-profile name.

## subscriber

<b>Syntax</b>	<b>subscriber</b> <i>sub-ident</i> <b>no subscriber</b>
<b>Context</b>	config>service>vprn>if>sap>static-host
<b>Description</b>	This command specifies an existing subscriber identification profile to be associated with the static subscriber host.
<b>Parameters</b>	<i>sub-ident</i> — Specifies the subscriber identification/

## subscriber-sap-id

<b>Syntax</b>	<b>[no] subscriber-sap-id</b>
<b>Context</b>	config>service>vprn>if>sap>static-host
<b>Description</b>	This command enables using the SAP ID as subscriber id.
<b>Parameters</b>	<b>subscriber-sap-id</b> — Specifies to use the sap-id as the subscriber-id.

## queue-override

<b>Syntax</b>	<b>[no] queue-override</b>
<b>Context</b>	config>service>vprn>if>sap>egress config>service>vprn>if>sap>ingress config>service>vprn>ipsec-if>sap>egress config>service>vprn>ipsec-if>sap>ingress
<b>Description</b>	This command enables the context to configure override values for the specified SAP egress or ingress QoS queue. These values override the corresponding ones specified in the associated SAP egress or ingress QoS policy.

## queue

<b>Syntax</b>	<b>[no] queue <i>queue-id</i></b>
<b>Context</b>	config>service>vprn>if>sap>egress>queue-override config>service>vprn>if>sap>ingress>queue-override config>service>vprn>ipsec-if>sap>egress>queue-override config>service>vprn>ipsec-if>sap>ingress>queue-override
<b>Description</b>	This command specifies the ID of the queue whose parameters are to be overridden.
<b>Parameters</b>	<i>queue-id</i> — The queue ID whose parameters are to be overridden. <b>Values</b> 1 — 32

## adaptation-rule

<b>Syntax</b>	<b>adaptation-rule [pir <i>adaptation-rule</i>] [cir <i>adaptation-rule</i>] no adaptation-rule</b>
<b>Context</b>	config>service>vprn>if>sap>egress>queue-override>queue config>service>vprn>if>sap>ingress>queue-override>queue config>service>vprn>ipsec-if>sap>egress>queue-override>queue config>service>vprn>ipsec-if>sap>ingress>queue-override>queue



<b>Description</b>	<p>This command can be used to override specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.</p> <p>The <b>no</b> form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific <b>adaptation-rule</b> is removed, the default constraints for <b>rate</b> and <b>cir</b> apply.</p>
<b>Default</b>	no adaptation-rule
<b>Parameters</b>	<p><b>pir</b> — The <b>pir</b> parameter defines the constraints enforced when adapting the PIR rate defined within the <b>queue queue-id rate</b> command. The <b>pir</b> parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the <b>rate</b> command is not specified, the default applies.</p> <p><b>cir</b> — The <b>cir</b> parameter defines the constraints enforced when adapting the CIR rate defined within the <b>queue queue-id rate</b> command. The <b>cir</b> parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the <b>cir</b> parameter is not specified, the default constraint applies.</p> <p><i>adaptation-rule</i> — Specifies the criteria to use to compute the operational CIR and PIR values for this queue, while maintaining a minimum offset.</p> <p><b>Values</b></p> <p><b>max</b> — The <b>max</b> (maximum) keyword is mutually exclusive with the <b>min</b> and <b>closest</b> options. When <b>max</b> is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the <b>rate</b> command.</p> <p><b>min</b> — The <b>min</b> (minimum) keyword is mutually exclusive with the <b>max</b> and <b>closest</b> options. When <b>min</b> is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the <b>rate</b> command.</p> <p><b>closest</b> — The <b>closest</b> parameter is mutually exclusive with the <b>min</b> and <b>max</b> parameter. When <b>closest</b> is defined, the operational PIR for the queue will be the rate closest to the rate specified using the <b>rate</b> command.</p>

## avg-frame-overhead

<b>Syntax</b>	<pre>avg-frame-overhead percent no avg-frame-overhead</pre>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;if&gt;sap&gt;egress&gt;queue-override&gt;queue config&gt;service&gt;vprn&gt;ipsec-if&gt;sap&gt;egress&gt;queue-override&gt;queue</pre>
<b>Description</b>	<p>This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).</p> <p>When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:</p>

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be  $10000 \times 0.1$  or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be  $50 \times 20$  or 1000 octets.

- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
  - Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queue's offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be  $1000 / 10000$  or 0.1. When in use, the avg-frame-overhead will be the same as the packet to frame factor making this calculation unnecessary.
  - Frame based CIR — The frame based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be  $500 \times 1.1$  or 550 octets.
  - Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).
- As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.
- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be  $7500 \times 1.1$  or 8250 octets.
  - Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-

cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

**Default** 0

**Parameters** *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

**Values** 0 — 100

## cbs

**Syntax** **cbs** *size-in-kbytes*  
**no cbs**

**Context** config>service>vprn>if>sap>egress>queue-override>queue  
config>service>vprn>if>sap>ingress>queue-override>queue  
config>service>vprn>ipsec-if>sap>egress>queue-override>queue  
config>service>vprn>ipsec-if>sap>ingress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's CBS parameters. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS setting into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

If the CBS value is larger than the MBS value, an error will occur, preventing the CBS change.

The **no** form of this command returns the CBS size to the default value.

**Default** no cbs

## Router Advertisement Commands

**Parameters** *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

**Values** 0 — 131072 or default

## high-prio-only

**Syntax** **high-prio-only** *percent*  
**no high-prio-only**

**Context** config>service>vprn>if>sap>egress>queue-override>queue  
config>service>vprn>if>sap>ingress>queue-override>queue  
config>service>vprn>ipsec-if>sap>egress>queue-override>queue  
config>service>vprn>ipsec-if>sap>ingress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's high-prio-only parameters. The **high-prio-only** command configures the percentage of buffer space for the queue, used exclusively by high priority packets.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The defined **high-prio-only** value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the **high-prio-only** value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command restores the default high priority reserved size.

**Parameters** *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue will be reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

**Values** 0 — 100 | default

## mbs

**Syntax** **mbs** {*size-in-kbytes* | **default**}  
**no mbs**

**Context** config>service>vprn>if>sap>egress>queue-override>queue  
config>service>vprn>if>sap>egress>hsmda-queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is

controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size assigned to the queue.

**Default** default

**Parameters** *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

For **sap>egress>queue-override>queue**:

**Values** 0 — 1073741824 or **default** in **bytes** or **kilobites**.

For **sap>egress>hsmda-queue-override>queue**:

**Values** [0 — 2625][**kilobytes**] | [0 — 2688000]**bytes** | **default**

## mbs

**Syntax** **mbs** {*size-in-kbytes* | **default**}  
**no mbs**

**Context** config>service>vprn>if>sap>ingress>queue-override>queue  
config>service>vprn>ipsec-if>sap>ingress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The defined high-prio-only value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the high-prio-only value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns the MBS size assigned to the queue to the value.

**Default** default

## Router Advertisement Commands

**Parameters** *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

**Values** 0 — 131072 or default

## rate

**Syntax** **rate** *pir-rate* [**cir** *cir-rate*]  
**no rate**

**Context** config>service>vprn>if>sap>egress>queue-override>queue  
config>service>vprn>if>sap>ingress>queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

**Default** **rate max cir 0** — The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

**Parameters** *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values** 1 — 2000000000, **max**

**Default** max

*cir-rate* — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

**Values** 0 — 2000000000, **max**

**Default** 0

## rate

**Syntax** **rate** *pir-rate*  
**no rate**

**Context** config>service>vprn>if>sap>egress>hsmda-queue-override>queue

**Description** This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR). The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The **rate** command can be executed at any time, altering the PIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR parameters (**max**, 0).

**Default** *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

**Values** 1 — 100000000

**Default** max

## scheduler-override

**Syntax** [**no**] **scheduler-override**

**Context** config>service>vprn>if>sap>egress  
config>service>vprn>if>sap>ingress  
config>service>vprn>ipsec-if>sap>egress  
config>service>vprn>ipsec-if>sap>ingress

**Description** This command specifies the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

## scheduler

<b>Syntax</b>	<b>scheduler</b> <i>scheduler-name</i> <b>no scheduler</b> <i>scheduler-name</i>
<b>Context</b>	config>service>vprn>if>sap>egress>sched-override config>service>vprn>if>sap>ingress>sched-override config>service>vprn>ipsec-if>sap>egress>sched-override config>service>vprn>ipsec-if>sap>ingress>sched-override
<b>Description</b>	<p>This command can be used to override specific attributes of the specified scheduler name.</p> <p>A scheduler defines a bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier, except for schedulers created in tier 1). A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.</p> <p>Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If <i>scheduler-name</i> already exists within the policy tier level (regardless of the inclusion of the keyword <b>create</b>), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).</p> <p>If the <i>scheduler-name</i> exists within the policy on a different tier (regardless of the inclusion of the keyword <b>create</b>), an error occurs and the current CLI context will not change.</p> <p>If the <i>scheduler-name</i> does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:</p> <ol style="list-style-type: none"> <li>1. The maximum number of schedulers has not been configured.</li> <li>2. The provided <i>scheduler-name</i> is valid.</li> <li>3. The <b>create</b> keyword is entered with the command if the system is configured to require it (enabled in the <b>environment create</b> command).</li> </ol> <p>When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.</p> <p>If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.</p>
<b>Parameters</b>	<p><i>scheduler-name</i> — The name of the scheduler.</p> <p><b>Values</b> Valid names consist of any string up to 32 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.</p> <p><b>Default</b> None. Each scheduler must be explicitly created.</p> <p><i>create</i> — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given <i>scheduler-name</i>. If the <b>create</b> keyword is omitted, <b>scheduler-name</b> is not created when the</p>



system environment variable `create` is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword `has` has no effect when the object already exists.

## parent

<b>Syntax</b>	<b>parent</b> [ <b>weight</b> <i>weight</i> ] [ <b>cir-weight</b> <i>cir-weight</i> ] <b>no parent</b>
<b>Context</b>	config>service>vprn>if>sap>ingress>sched-override>scheduler config>service>vprn>if>sap>egress>sched-override>scheduler
<b>Description</b>	<p>This command can be used to override the scheduler's parent weight and cir-weight information. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.</p> <p>The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy – this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non default weightings for fostered schedulers.</p> <p>The no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.</p>
<b>Default</b>	no parent
<b>Parameters</b>	<p><b>weight</b> <i>weight</i> — <b>Weight</b> defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict <b>level</b> defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.</p> <p>A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.</p> <p><b>Values</b>      0 to 100</p> <p><b>Default</b>      1</p> <p><b>cir-weight</b> <i>cir-weight</i> — The <b>cir-weight</b> keyword defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same <i>cir-level</i> defined by the <b>cir-level</b> parameter in the applied scheduler policy. Within the strict <b>cir-level</b>, all <b>cir-weight</b> values from active children at that level are summed and the ratio of each active child's <b>cir-weight</b> to the total is used to distribute the available bandwidth at that level. A <b>cir-weight</b> is considered to be active when the queue or scheduler that the <b>cir-weight</b> pertains to has not reached the CIR and still has packets to transmit.</p> <p>A 0 (zero) <b>cir-weight</b> value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.</p>

**Values** 0 — 100  
**Default** 1

## rate

<b>Syntax</b>	<b>rate</b> <i>pir-rate</i> [ <i>cir cir-rate</i> ] <b>no rate</b>
<b>Context</b>	config>service>vprn>if>sap>egress>sched-override>scheduler config>service>vprn>ipsec-if>sap>egress>sched-override config>service>vprn>ipsec-if>sap>ingress>sched-override
<b>Description</b>	<p>This command can be used to override specific attributes of the specified scheduler rate. The <b>rate</b> command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.</p> <p>The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.</p> <p>When a scheduler is defined without specifying a rate, the default rate is <b>max</b>. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.</p> <p>The <b>no</b> form of this command returns all queues created with this <i>queue-id</i> by association with the QoS policy to the default PIR and CIR parameters.</p>
<b>Parameters</b>	<p><i>pir-rate</i> — The <b>pir</b> parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue will operate. A value of 0 to 100000000 or the keyword <b>max</b> or <b>sum</b> is accepted. Any other value will result in an error without modifying the current PIR rate.</p> <p>To calculate the actual PIR rate, the rate described by the queue's <b>rate</b> is multiplied by the <i>pir-rate</i>.</p> <p>The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default <b>pir</b> and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue will be allowed to forward packets in a given second, thus shaping the queue's output.</p> <p>The PIR parameter for SAP ingress queues do not have a negate (<b>no</b>) function. To return the queue's PIR rate to the default value, that value must be specified as the PIR value.</p>
	<b>Values</b> 1 — 100000000, <b>max</b>
	<b>Default</b> max

*cir cir-rate* — The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue will operate. A value of 0 to 250 or the keyword **max** is accepted. Any other value will result in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir** *pir-rate* is multiplied by the *cir cir-rate*. If the **cir** is set to **max**, then the CIR rate is set to infinity.

The SAP ingress context for CIR is dependent on the defined forwarding class (*fc*) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

**Values**      0 — 10000000, **max**, **sum**

**Default**     sum

---

## Routed VPLS Commands

### vpls

**Syntax** `vpls service-name`

**Context** `config>service`  
`config>service>vprn>if`

**Description** The **vpls** command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (**no shutdown**). Once the IP interface is administratively up, the system will scan the available VPLS services that have the **allow-ip-int-bind** flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

If a VPLS service is found associated with the name and with the **allow-ip-int-bind** flag set, the IP interface will be attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the **allow-ip-int-bind** flag set or a non-VPLS service associated with the name will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service **allow-ip-int-bind** flag is set at the time the name is applied, the VPLS service will be automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the **allow-ip-int-bind** flag set, the system will not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the **allow-ip-int-bind** flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the **allow-ip-int-bind** flag cannot be removed until the attached IP interface is unbound from the service name.

Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the **sap** or **spoke-sdp** commands on the interface.

**IES CHASSIS MODE DEPENDENCY**

An IES IP interface cannot be bound to a service name unless the system is configured in chassis mode D. Once an IES interface is bound to a service name, the chassis mode of the system cannot be changed to B or C.

**VPRN HARDWARE DEPENDENCY**

When a service name is bound to a VPRN IP interface, all SAPs associated with the VPRN service must be on hardware based on the FlexPath forwarding plane. Currently, these include the IOM3-XP, the various IMM modules and the SR7710c12. If any SAPs are associated with the wrong hardware type, the service name binding to the VPRN IP interface will fail. Once an IP interface within the VPRN service is bound to a service name, attempting to create a SAP on excluded hardware will fail.

**ROUTE EXPORT AND IMPORT BETWEEN ROUTING CONTEXTS**

The IES chassis mode dependency and the VPRN hardware dependency each are designed to prevent a condition where an ingress routing decision on hardware that does not support the mixed Layer 2 and Layer 3 behavior of routed VPLS is asked to route to a VPLS based next-hop.

Even with these restrictions, it is still possible using route leaking or import/export routing policies to create a condition where a FlexPath forwarding plane resolves a route to a VPLS next-hop. In this case, the forwarding plane handles the resolved next-hop as if it points to a null IP interface. Packets associated with a null next-hop egress IP interface will be discarded and an ICPM unreachable message will be generated when enabled.

**IP INTERFACE MTU AND FRAGMENTATION**

A VPLS service is affected by two MTU values; port MTUs and the VPLS service MTU. The MTU on each physical port defines the largest Layer 2 packet (including all DLC headers and CRC) that may be transmitted out a port. The VPLS itself has a service level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (QinQ, Dot1Q, TopQ or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, the virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational virtual port is ensured to support the largest packet traversing the VPLS service. The service delineation overhead on each Layer 2 packet is removed before forwarding into a VPLS service. VPLS services do not support fragmentation and must discard any Layer 2 packet larger than the service MTU after the service delineation overhead is removed.

IP interfaces have a configurable up MTU that defines the largest packet that may egress the IP interface without being fragmented. This MTU encompasses the IP portion of the packet and does not include any of the egress DLC header or CRC. This MTU does not affect the size of the largest ingress packet on the IP interface. If the egress IP portion of the packet is larger than the IP interface MTU and the IP header do not fragment flag is not set, the packet is fragmented into smaller packets that will not exceed the configured MTU size. If the do not fragment bit is set, the packet is silently discarded at egress when it exceeds the IP MTU.

When the IP interface is bound to a VPLS service, the IP MTU must be at least 18 bytes less than the VPLS service MTU. This allows for the addition of the minimal Ethernet encapsulation overhead; 6 bytes for the DA, 6 bytes for the SA, 2 bytes for the Etype and 4 bytes for the trailing CRC. Any remaining egress virtual port overhead (Dot1P, Dot1Q, QinQ, TopQ or SDP) required above the minimum is known to be less than the egress ports MTU since the virtual port would not be operational otherwise.

## Router Advertisement Commands

If the IP interface IP MTU value is too large based on the VPLS service MTU, the IP interface will enter the operationally down state until either the IP MTU is adequately lowered or the VPLS service MTU is sufficiently increased.

The **no** form of the command on the IP interface is used to remove the service name binding from the IP interface. If the service name has been resolved to a VPLS service context and the IP interface has been attached to the VPLS service, the IP interface will also be detached from the VPLS service.

**Default** none

**Parameters** service-name

The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an allow-ip-int-bind enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

## ingress

**Syntax** ingress

**Context** config>service>vprn>if>vpls

**Description** The ingress node in this context under the vpls binding is used to define the routed IPv4 and IPv6 optional filter overrides.

## v4-routed-override-filter

**Syntax** v4-routed-override-filter *ipv4-filter-id*  
no v4-routed-override-filter

**Context** config>service>vprn>if>vpls>ingress

**Description** The v4-routed-override-filter command is used to specify an IPv4 filter ID that will be applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IPv4 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IPv4 routed packets will use the any existing ingress IPv4 filter on the VPLS virtual port.

The **no** form of the command is used to remove the IPv4 routed override filter from the ingress IP interface. When removed, the IPv4 ingress routed packets within a VPLS service attached to the IP interface will use the IPv4 ingress filter applied to the packets virtual port when defined.

**Default** none

**Parameters** *ipv4-filter-id* — The ipv4-filter-id parameter is required when executing the v4-routed-override-filter command. The specified filter ID must exist as an IPv4 filter within the system or the override command will fail.

## v6-routed-override-filter

<b>Syntax</b>	<b>v6-routed-override-filter</b> <i>ipv6-filter-id</i> <b>no v6-routed-override-filter</b>
<b>Context</b>	config>service>vprn>if>vpls>ingress
<b>Description</b>	<p>The v6-routed-override-filter command is used to specify an IPv6 filter ID that will be applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IPv6 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IPv6 routed packets will use the any existing ingress IPv6 filter on the VPLS virtual port.</p> <p>The no v6-routed-override-filter command is used to remove the IPv6 routed override filter from the ingress IP interface. When removed, the IPv6 ingress routed packets within a VPLS service attached to the IP interface will use the IPv6 ingress filter applied to the packets virtual port when defined.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>ipv6-filter-id</i> — The ipv6-filter-id parameter is required when executing the v6-routed-override-filter command. The specified filter ID must exist as an IPv6 filter within the system or the override command will fail.

## egress

<b>Syntax</b>	<b>egress</b>
<b>Context</b>	config>service>vprn>if>vpls
<b>Description</b>	The egress node under the vpls binding is used to define the optional sap-egress QoS policy that will be used for reclassifying the egress forwarding class or profile for routed packets associated with the IP interface on the attached VPLS service context.

## reclassify-using-qos

<b>Syntax</b>	<b>reclassify-using-qos</b> <i>sap-egress-qos-id</i> <b>no reclassify-using-qos</b>
<b>Context</b>	config>service>vprn>if>vpls>egress
<b>Description</b>	<p>The reclassify-using-qos command is used to specify a sap-egress QoS policy that will be used to reclassify the forwarding class and profile of egress routed packets on the VPLS service. When routed packets associated with the IP interface egress a VPLS SAP, the reclassification rules within the sap-egress QoS policy applied to the SAP are always ignored (even when reclassify-using-qos is not defined).</p> <p>Any queues or policers defined within the specified QoS policy are ignored and are not created on the VPLS egress SAPs. Instead, the routed packets continue to use the forwarding class mappings, queues and policers from the sap-egress QoS policy applied to the egress VPLS SAP.</p>

While the specified sap-egress policy ID is applied to an IP interface it cannot be deleted from the system.

The **no** form of the command removes the sap-egress QoS policy used for reclassification from the egress IP interface. When removed, IP routed packets will not be reclassified on the egress SAPs of the VPLS service attached to the IP interface.

**Parameters** *sap-egress-qos-id* — The sap-egress-qos-id parameter is required when executing the reclassify-using-qos command. The specified SAP egress QoS ID must exist within the system or the command will fail.

### allow-ip-int-bind

**Syntax** **[no] allow-ip-int-bind**

**Context** config>service>vpls

**Description** This command sets a flag on the VPLS service that enables the ability to attach an IES or VPRN IP interface to the VPLS service in order to make the VPLS service routable. When the allow-ip-int-bind command is not enabled, the VPLS service cannot be attached to an IP interface.

#### VPLS Configuration Constraints for Enabling allow-ip-int-bind

When attempting to set the allow-ip-int-bind VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. In Release 8.0 the following VPLS features must be disabled or not configured for the allow-ip-int-bind flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- SDPs used in spoke or mesh SDP bindings cannot be configured as GRE
- The VPLS service type cannot be B-VPLS or M-VPLS and it cannot be an I-VPLS service bound to a B-VPLS context
- MVR from Routed VPLS and to another SAP is not supported
- Enhanced and Basic Subscriber Management features
- Network Domain on SDP bindings

Once the VPLS allow-ip-int-bind flag is set on a VPLS service, the above features cannot be enabled on the VPLS service.

#### NETWORK PORT HARDWARE CONSTRAINTS

The system also checks to ensure that all ports configured in network mode are associated with FlexPath forwarding planes. If a port is currently in network mode and the port is associated with a FlexPath forwarding plane, the allow-ip-int-bind command will fail. Once the allow-ip-int-bind flag is set on any VPLS service, attempting to enable network mode on a port associated with a FlexPath forwarding plane will fail.

#### VPLS SAP HARDWARE CONSTRAINTS

Besides VPLS configuration and network port hardware association, the system also checks to that all SAPs within the VPLS are created on Ethernet ports and the ports are associated with FlexPath forwarding planes. Certain Ethernet ports and virtual Ethernet ports are not supported which include HSMDA ports and CCAG virtual ports (VSM based). If a SAP in the VPLS exists on an unsupported port type or is associated with a FlexPath forwarding plane, the allow-ip-int-bind command will fail.



Once the `allow-ip-int-bind` flag is set on the VPLS service, attempting to create a VPLS SAP on the wrong port type or associated with a FlexPath forwarding plane will fail.

### **VPLS SERVICE NAME BOUND TO IP INTERFACE WITHOUT ALLOW-IP-INT-BIND FLAG SET**

In the event that a service name is applied to a VPLS service and that service name is also bound to an IP interface but the `allow-ip-int-bind` flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface will fail. After the `allow-ip-int-bind` flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied or the IP interface must be re-initialized using the `shutdown / no shutdown` commands. This will cause the system to reattempt the name resolution process between the IP interface and the VPLS service.

The `no` form of the command resets the `allow-ip-int-bind` flag on the VPLS service. If the VPLS service currently has an IP interface from an IES or VPRN service attached, the `no allow-ip-int-bind` command will fail. Once the `allow-ip-int-bind` flag is reset on the VPLS service, the configuration and hardware restrictions associated with setting the flag are removed. The port network mode hardware restrictions are also removed.

---

## ETH-CFM Service Commands

### eth-cfm

<b>Syntax</b>	<b>eth-cfm</b>
<b>Context</b>	config>service>vprn config>service>vprn>if>sap config>service>vprn>if>spoke-sdp config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	This command enables the context to configure ETH-CFM parameters.

### collect-lmm-stats

<b>Syntax</b>	<b>collect-lmm-stats</b> <b>no collect-lmm-stats</b>
<b>Context</b>	config>service>vprn>interface>sap>eth-cfm config>service>vprn>interface>spoke-sdp>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
<b>Description</b>	This command enables the collection of statistics on the SAP or MPLS SDP binding on which the ETH- LMM test is configured. The collection of LMM statistics must be enabled if a MEP is launching or responding to ETH-LMM packets. If LMM statistics collection is not enabled, the counters in the LMM and LMR PDU do not represent accurate measurements and all measurements should be ignored. The <b>show sap-using eth-cfm collect-lmm-stats</b> command and the <b>show sdp-using eth-cfm collect-lmm-stats command</b> can be used to display which entities are collecting stats. The <b>no</b> form of the command disables and deletes the counters for this SAP or MPLS SDP binding.
<b>Default</b>	<b>no collect-lmm-stats</b>

### mep

<b>Syntax</b>	<b>mep</b> <i>mep-id</i> <b>domain</b> <i>md-index</i> <b>association</b> <i>ma-index</i> [ <b>direction</b> {up   down}] <b>no mep</b> <i>mep-id</i> <b>domain</b> <i>md-index</i> <b>association</b> <i>ma-index</i>
<b>Context</b>	config>service>vprn>if>sap>eth-cfm config>service>vprn>if>spoke-sdp>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
<b>Description</b>	This command configures the ETH-CFM maintenance endpoint (MEP).
<b>Parameters</b>	<i>mep-id</i> — Specifies the maintenance association end point identifier. <b>Values</b> 1 — 8191

*md-index* — Specifies the maintenance domain (MD) index value.

**Values** 1 — 4294967295

*ma-index* — Specifies the MA index value.

**Values** 1 — 4294967295

**direction up|down** — Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction UP is not supported on VPRN MEPs.

down — Sends continuity check messages away from the MAC relay entity.

up — Sends continuity check messages towards the MAC relay entity.

## ais-enable

<b>Syntax</b>	<b>[no] ais-enable</b>
<b>Context</b>	config>service>vprn>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm
<b>Description</b>	This command configures the reception of Alarm Indication Signal (AIS) message.

## interface-support-enable

<b>Syntax</b>	<b>[no] interface-support-enable</b>
<b>Context</b>	config>service>vprn>sap>eth-cfm>mep>ais-enable config>service>vprn>spoke-sdp>eth-cfm>mep>ais-enable
<b>Description</b>	This command enables the AIS function to consider the operational state of the entity on which it is configured. With this command, ETH-AIS on DOWN MEPs will be triggered and cleared based on the operational status of the entity on which it is configured. If CCM is also enabled then transmission of the AIS PDU will be based on either the non operational state of the entity or on ANY CCM defect condition. AIS generation will cease if BOTH operational state is UP and CCM has no defect conditions. If the MEP is not CCM enabled then the operational state of the entity is the only consideration assuming this command is present for the MEP.
<b>Default</b>	<b>no interface-support-enabled</b> (AIS will not be generated or stopped based on the state of the entity on) which the DOWN MEP is configured.

## ccm-enable

<b>Syntax</b>	<b>[no] ccm-enable</b>
<b>Context</b>	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
<b>Description</b>	This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

### ccm-ltm-priority

<b>Syntax</b>	<b>ccm-ltm-priority</b> <i>priority</i> <b>no ccm-ltm-priority</b>
<b>Context</b>	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
<b>Description</b>	This command specifies the priority value for CCMs and LTMs transmitted by the MEP. The <b>no</b> form of the command removes the priority value from the configuration.
<b>Default</b>	The highest priority on the bridge-port.
<b>Parameters</b>	<i>priority</i> — Specifies the priority of CCM and LTM messages.
	<b>Values</b> 0 — 7

### ccm-padding-size

<b>Syntax</b>	<b>[no] ccm-padding-size</b> <i>ccm-padding</i>
<b>Context</b>	config>service>vprn>interface>sap>eth-cfm>mep config>service>vprn>interface>spoke-sdp>eth-cfm>mep config>service>vprn>subscriber-interface>group-interface>sap>eth-cfm>mep
<b>Description</b>	This command sets the byte size of the optional Data TLV to be included in the ETH-CC PDU. This will increase the size of the ETH-CC PDU by the configured value. The base size of the ETH-CC PDU, including the Interface Status TLV and Port Status TLV, is 83 bytes not including the Layer Two encapsulation. CCM padding is not supported when the CCM-Interval is less than one second.
<b>Default</b>	ccm-padding-size
<b>Parameters</b>	<i>ccm-padding</i> — specifies the byte size of the Optional Data TLV
	<b>Values</b> 3 — 1500

### csf-enable

<b>Syntax</b>	<b>[no] csf-enable</b>
<b>Context</b>	config>service>vprn>interface>sap>eth-cfm>mep config>service>vprn>interface>spoke-sdp>eth-cfm>mep config>service>vprn>subscriber-interface>group-interface>sap>eth-cfm>mep
<b>Description</b>	This command enables the reception and local processing of ETH-CSF frames.

## multiplier

<b>Syntax</b>	<b>multiplier</b> <i>multiplier-value</i> <b>no multiplier</b>
<b>Context</b>	config>service>vprn>interface>sap>eth-cfm>mep>cfs-enable config>service>vprn>interface>spoke-sdp>eth-cfm>mep>cfs-enable config>service>vprn>subscriber-interface>group-interface>sap>eth-cfm>mep>cfs-enable
<b>Description</b>	This command enables the multiplication factor applied to the receive time used to clear the CSF condition in increments of .5.
<b>Default</b>	3.5
<b>Parameters</b>	<i>multiplier-value</i> — Specifies the multiplier used for timing out CSF. <b>Values</b> 0.0, 2.0 .. 30.0

## eth-test-enable

<b>Syntax</b>	<b>[no] eth-test-enable</b>
<b>Context</b>	config>service>vprn>if>sap>eth-cfm>mep config>service>vprn>if>spoke-sdp>eth-cfm>mep config>service>vprn>sub-if>grp-if>sap>eth-cfm
<b>Description</b>	This command enables eth-test functionality on MEP. For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:  oam eth-cfm eth-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [ <i>priority priority</i> ] [ <i>data-length data-length</i> ]  A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

## test-pattern

<b>Syntax</b>	<b>test-pattern {all-zeros   all-ones} [crc-enable]</b> <b>no test-pattern</b>
<b>Context</b>	config>service>vprn>if>sap>eth-cfm>mep>eth-test-enable config>service>vprn>if>spoke-sdp>eth-cfm>mep>eth-test-enable config>service>vprn>sub-if>grp-if>sap>eth-cfm>eth-test-enable
<b>Description</b>	This command configures the test pattern for eth-test frames.  The <b>no</b> form of the command removes the values from the configuration.
<b>Parameters</b>	<b>all-zeros</b> — Specifies to use all zeros in the test pattern.

## ETH-CFM Service Commands

**all-ones** — Specifies to use all ones in the test pattern.

**crc-enable** — Generates a CRC checksum.

**Default** all-zeros

### bit-error-threshold

**Syntax** **bit-error-threshold** *bit-errors*

**Context** config>service>vprn>if>sap>eth-cfm>mep  
config>service>vprn>if>spoke-sdp>eth-cfm>mep  
config>service>vprn>sub-if>grp-if>sap>eth-cfm

**Description** This command specifies the lowest priority defect that is allowed to generate a fault alarm.

**Default** 1

**Parameters** *bit-errors* — Specifies the lowest priority defect.

**Values** 0 — 11840

### one-way-delay-threshold

**Syntax** **one-way-delay-threshold** *time*

**Context** config>service>vprn>if>sap>eth-cfm  
config>service>vprn>if>spoke-sdp>eth-cfm  
config>service>vprn>sub-if>grp-if>sap>eth-cfm

**Description** This command enables one way delay threshold time limit.

**Default** 3 seconds

**Parameters** *priority* — Specifies the value for the threshold.

**Values** 0 — 600

### snelch-ingress-levels

**Syntax** **snelch-ingress-levels** [*md-level* [*md-level*...]]  
**no snelch-ingress-levels**

**Context** config>service>vprn>if>sap>eth-cfm  
config>service>vprn>if>spoke-sdp>eth-cfm  
config>service>vprn>sub-if>grp-if>sap>eth-cfm

**Description** This command defines the levels of the ETH-CFM PDUs that will silently be discarded on ingress into the SAP or SDP Binding from the wire. All ETH-CFM PDUs inbound to the SAP or SDP binding will be dropped that match the configured levels without regard for any other ETH-CFM

criteria. No statistical information or drop count will be available for any ETH-PDU that is silently discarded by this option. The operator must configure a complete contiguous list of md-levels up to the highest level that will be dropped. The command must be retyped in complete form to modify a previous configuration, if the operator does not want to delete it first.

The **no** form of the command removes the silent discarding of previously matching ETH-CFM PDUs.

**Default** no squelch-ingress-levels

**Parameters** *md-level* — Identifies the level.

**Values** [0..7]

## tunnel-fault

**Syntax** tunnel-fault {accept | ignore}

**Context** config>service>vprn>eth-cfm  
config>service>vprn>if>sap>eth-cfm  
config>service>vprn>sub-if>grp-if>sap>eth-cfm

**Description** Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the command ais-enable under epipe>sap>eth-cfm>ais-enable for more details. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is “ignore” and the sap level command is “accept”. This means simply changing the service level command to “accept” will enable the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.

**Parameters** **accept** — Share fate with the facility tunnel MEP  
**ignore** — Do not share fate with the facility tunnel MEP

**Default** ignore (Service Level)  
accept (SAP Level for Epipe and VPLS)

## fault-propagation-enable

**Syntax** fault-propagation-enable {use-if-tlv | suspend-ccm}  
no fault-propagation-enable

**Context** config>service>vprn>if>sap>eth-cfm>mep  
config>service>vprn>if>spoke-sdp>eth-cfm>mep  
config>service>vprn>sub-if>grp-if>sap>eth-cfm

**Description** This command configures the fault propagation for the MEP.

## ETH-CFM Service Commands

- Parameters**
- use-if-tlv** — Specifies to use the interface TLV.
  - suspend-ccm** — Specifies to suspend the continuity check messages.

### low-priority-defect

- Syntax** `low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}`
- Context**  
config>service>vprn>if>sap>eth-cfm>mep  
config>service>vprn>if>spoke-sdp>eth-cfm>mep  
config>service>vprn>sub-if>grp-if>sap>eth-cfm
- Description** This command specifies the lowest priority defect that is allowed to generate a fault alarm.
- Default** macRemErrXcon
- | Values        |  |
|---------------|--|
| allDef        | DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM |
| macRemErrXcon | Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM       |
| remErrXcon    | Only DefRemoteCCM, DefErrorCCM, and DefXconCCM                     |
| errXcon       | Only DefErrorCCM and DefXconCCM                                    |
| xcon          | Only DefXconCCM; or  |
| noXcon        | No defects DefXcon or lower are to be reported                     |



## SAP Subscriber Management Commands

### sub-sla-mgmt

<b>Syntax</b>	<b>[no] sub-sla-mgmt</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap
<b>Description</b>	This command enables the context to configure subscriber management parameters for this SAP.
<b>Default</b>	no sub-sla-mgmt

### def-sla-profile

<b>Syntax</b>	<b>def-sla-profile</b> <i>default-sla-profile-name</i> <b>no def-sla-profile</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
<b>Description</b>	<p>This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sla-profile</b> context.</p> <p>An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.</p> <p>The <b>no</b> form of the command removes the default SLA profile from the SAP configuration.</p>
<b>Default</b>	no def-sla-profile
<b>Parameters</b>	<i>default-sla-profile-name</i> — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sla-profile</b> context.

### def-sub-profile

<b>Syntax</b>	<b>def-sub-profile</b> <i>default-subscriber-profile-name</i>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
<b>Description</b>	This command specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-profile</b> context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of the command removes the default SLA profile from the SAP configuration.

**Parameters** *default-sub-profile* — Specifies a default subscriber profile for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscriber-mgmt>sub-profile** context.

## multi-sub-sap

**Syntax** **multi-sub-sap** [*number-of-sub*]  
**no multi-sub-sap**

**Context** config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

**Description** This command configures the maximum number of subscribers for this SAP. It is used in conjunction with the **profiled-traffic-only** command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of this command returns the default value.

**Default** 1

**Parameters** *number-of-sub* — Specifies the maximum number of subscribers for this SAP.

**Values** 2 — 8000

## single-sub-parameters

**Syntax** **single-sub-parameters**

**Context** config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

**Description** This command enables the context to configure single subscriber parameters for this SAP.

## non-sub-traffic

**Syntax** **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*]  
**no non-sub-traffic**

**Context** config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

**Description** This command configures non-subscriber traffic profiles. It is used in conjunction with the **profiled-traffic-only** command on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of the command removes the profiles and disables the feature.

<b>Parameters</b>	<p><b>sub-profile</b> <i>sub-profile-name</i> — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the <b>config&gt;subscr-mgmt&gt;sub-profile</b> context.</p> <p><b>sla-profile</b> <i>sla-profile-name</i> — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the <b>config&gt;subscr-mgmt&gt;sla-profile</b> context.</p> <p><b>subscriber</b> <i>sub-ident-string</i> — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the <b>config&gt;subscr-mgmt&gt;sub-ident-policy</b> context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.</p> <ul style="list-style-type: none"> <li>For VPRN SAPs with <b>arp-reply-agent</b> enabled with the optional <i>sub-ident</i> parameter, the static subscriber host's <i>sub-ident-string</i> is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.</li> </ul> <p>If the static subscriber host's <i>sub-ident</i> string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.</p> <p>If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.</p> <p>If <i>sub-ident</i> is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.</p> <p>ARP requests are never forwarded back to the same SAP or within the receiving SAP's split horizon group.</p>
-------------------	---

## profiled-traffic-only

<b>Syntax</b>	<b>[no] profiled-traffic-only</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
<b>Description</b>	<p>This command enables profiled traffic only for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the <i>sla-profile</i> and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).</p> <p>The <b>no</b> form of the command disables the command.</p>

## sub-ident-policy

<b>Syntax</b>	<b>sub-ident-policy</b> <i>sub-ident-policy-name</i>
<b>Context</b>	config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
<b>Description</b>	<p>This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-ident-policy</b> context.</p> <p>Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.</p> <p>For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.</p> <p>When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.</p> <p>The <b>no</b> form of the command removes the default subscriber identification policy from the SAP configuration.</p>
<b>Default</b>	no sub-ident-policy
<b>Parameters</b>	<i>sub-ident-policy-name</i> — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the <b>config&gt;subscriber-mgmt&gt;sub-ident-policy</b> context.

## srrp

<b>Syntax</b>	<b>[no] srrp</b> <i>srrp-id</i>
<b>Context</b>	config>service>vprn>sub-if>grp-if
<b>Description</b>	<p>This command creates an SRRP instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface.</p> <p>The <b>no</b> form of the command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using the redundant IP interface associated with the group IP interface and will stop routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC will also be silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).</p>
<b>Default</b>	no srrp
<b>Parameters</b>	<i>srrp-id</i> — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context.

SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets.

**Values** 1 — 4294967295

## gw-mac

<b>Syntax</b>	<b>gw-mac</b> <i>mac-address</i> <b>no gw-mac</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>srrp
<b>Description</b>	<p>This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context.</p> <p>One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address.</p> <p>The <b>no</b> form of the command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.</p>
<b>Parameters</b>	<p><i>mac-address</i> — Specifies a MAC address that is used to override the default SRRP base MAC address</p> <p><b>Values</b> Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.</p> <p>If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.</p>

## keep-alive-interval

<b>Syntax</b>	<b>keep-alive-interval</b> <i>interval</i> <b>no keep-alive-interval</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>srrp
<b>Description</b>	<p>This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better than the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master.</p>

When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval.

The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.

The **no** form of the command restores the default interval.

**Parameters** *interval* — Specifies the interval, in milliseconds, between SRRP advertisement messages sent when operating in the master state.

**Values** 1 — 100

**Default** 10 milliseconds

## message-path

**Syntax** **message-path** *sap-id*  
**no message-path**

**Context** config>service>vprn>sub-if>grp-if>srrp

**Description** This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.

The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.

Changing the current SRRP message SAP on an active pair of routers should be done in the following manner:

1. Shutdown the backup SRRP instance.
2. Change the message SAP on the shutdown node.
3. Change the message SAP on the active master node.
4. Re-activate the shutdown SRRP instance.

Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs.

If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP will be sent from each member.

The **no** form of the command can only be executed when the SRRP instance is shutdown. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.

**Parameters** *sap-id* — Specifies the physical port identifier portion of the SAP definition.

## policy

<b>Syntax</b>	<b>[no] policy</b> <i>vrrp-policy-id</i>
<b>Context</b>	config>service>vprn>sub-if>grp-if>srrp
<b>Description</b>	<p>This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach Layer 2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.</p> <p>More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance.</p> <p>VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance.</p> <p>The <b>no</b> form of the command removes the association with <i>vrrp-policy-id</i> from the SRRP instance.</p>
<b>Parameters</b>	<p><i>vrrp-policy-id</i> — Specifies one or more VRRP policies with the SRRP instance.</p> <p><b>Values</b>      1 — 9999</p>

## priority

<b>Syntax</b>	<b>priority</b> <i>priority</i> <b>no priority</b>
<b>Context</b>	config>service>vprn>sub-if>grp-if>srrp
<b>Description</b>	<p>This command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state.</p> <p>The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the <i>becoming backup</i> state.</p> <p>When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.</p> <p>The <b>no</b> form of the command restores the default base priority to the SRRP instance. If a VRRP policy is associated with the SRRP instance, it will use the default base priority as the basis for any modifications to the SRRP instances in-use priority.</p>

**Parameters** *priority* — Specifies a base priority for the SRRP instance to override the default.

**Values** 1 — 254

**Default** 100

## send-fib-population-packets

**Syntax** **send-fib-population-packets (all | outer-tag-only)**  
**no send-fib-population-packets**

**Context** config>service>vprn>sub-if>grp-if>srrp

**Description** This command sends FIB population packets.  
The **no** form of the command disables sending FIB population packets.

**Default** all

**Parameters** **all** — Sends FIB population packets to all VLANs.  
**outer-tag-only** — Sends FIB population packets to only outer VLAN tags.

## generate-garp-on-outer-vlan

**Syntax** **send-fib-population-packets (all | outer-tag-only)**  
**no send-fib-population-packets**

**Context** config>service>vprn>sub-if>grp-if>srrp

**Description** This command sends GARP packets to outer VLANs only.  
The **no** form of the command disables sending GARP packets to outer VLANs only.

**Default** no send-fib-population-packets



## Interface VRRP Commands

### vrrp

<b>Syntax</b>	<b>vrrp</b> <i>virtual-router-id</i> [ <b>owner</b> ] <b>no vrrp</b> <i>virtual-router-id</i>
<b>Context</b>	config>service>vprn>if
<b>Description</b>	<p>This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.</p> <p>Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of <code>vrrp virtual-router-id</code> is used to define the configuration parameters for the VRID.</p> <p>The <b>no</b> form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown in order to remove the virtual router instance.</p>
<b>Default</b>	No default
<b>Parameters</b>	<p><i>virtual-router-id</i> — The <i>virtual-router-id</i> parameter specifies a new virtual router ID or one that can be modified on the IP interface.</p> <p><b>Values</b>      1 — 255</p>

### authentication-key

<b>Syntax</b>	<b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ] <b>no authentication-key</b>
<b>Context</b>	config>service>vprn>if>vrrp
<b>Description</b>	<p>The <b>authentication-key</b> command, within the <code>vrrp virtual-router-id</code> context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.</p> <p>The <b>authentication-key</b> command is one of the few commands not affected by the presence of the <b>owner</b> keyword. If simple text password authentication is not required, this command is not required. If the command is re-executed with a different password key defined, the new key will be used immediately. If a <b>no authentication-key</b> command is executed, the password authentication key is restored to the default value. The <b>authentication-key</b> command may be executed at any time.</p> <p>To change the current in-use password key on multiple virtual router instances:</p> <ul style="list-style-type: none"> <li>• Identify the current master</li> <li>• Shutdown the virtual router instance on all backups</li> <li>• Execute the authentication-key command on the master to change the password key</li> <li>• Execute the authentication-key command and no shutdown command on each backup key</li> </ul>

The **no** form of this command restores the default null string to the value of key.

**Default** No default. The authentication data field contains the value 0 in all 16 octets.

**Parameters** *authentication-key* — The *key* parameter identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string eight octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key.

The *key* parameter is expressed as a string consisting of up to eight alpha-numeric characters. Spaces must be contained in quotation marks ( “ ” ). The quotation marks are not considered part of the string.

The 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 holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

**Values** Any 7-bit printable ASCII character.

Exceptions:	Double quote (")	ASCII 34
	Carriage Return	ASCII 13
	Line Feed	ASCII 10
	Tab	ASCII 9
	Backspace	ASCII 8

*hash-key* — The hash key. The key can be any combination of ASCII characters up to 22 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 ip-address**

**Context** config>service>vprn>if>vrrp  
config>service>vprn>if>ipv6>vrrp

**Description** This command configures virtual router IP addresses for the interface.

## bfd-enable

<b>Syntax</b>	<b>bfd-enable interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i> <b>bfd-enable service-id interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i> <b>no bfd-enable interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i> <b>no bfd-enable service-id interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i>		
<b>Context</b>	config>service>vprn>if>vrrp config>service>vprn>sub-if>grp-if>srrp config>service>vprn>if>ipv6>vrrp		
<b>Description</b>	This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. 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. If the interface used is configured with centralized BFD, the BFD transmit and receive intervals need to be set to at least 300ms.  BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface. The specified interface may not be configured with BFD; when it is, the virtual router will then initiate the BFD session.  The <b>no</b> form of this command removes BFD from the configuration.		
<b>Default</b>	none		
<b>Parameters</b>	<i>service-id</i> — Specifies the service ID of the interface running BFD. <table> <tr> <td><b>Values</b></td> <td><i>service-id</i>: 1 — 2147483648  <i>svc-name</i>: Specifies an existing service name up to 64 characters in length.  No service ID indicates a network interface.</td> </tr> </table> <i>interface interface-name</i> — Specifies the name of the interface running BFD. <i>dst-ip ip-address</i> — Specifies the destination address to be used for the BFD session.	<b>Values</b>	<i>service-id</i> : 1 — 2147483648 <i>svc-name</i> : Specifies an existing service name up to 64 characters in length. No service ID indicates a network interface.
<b>Values</b>	<i>service-id</i> : 1 — 2147483648 <i>svc-name</i> : Specifies an existing service name up to 64 characters in length. No service ID indicates a network interface.		

## init-delay

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

mac

<b>Syntax</b>	<b>[no] mac</b> <i>ieee-mac-address</i>
<b>Context</b>	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
<b>Description</b>	This command assigns a specific MAC address to an IP interface. The <b>no</b> form of this command returns the MAC address of the IP interface to the default value.
<b>Default</b>	The physical MAC address associated with the Ethernet interface that the SAP is configured on.
<b>Parameters</b>	<i>ieee-mac-address</i> — Specifies the 48-bit MAC address for the static ARP 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

<b>Syntax</b>	<b>[no] master-int-inherit</b>
<b>Context</b>	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
<b>Description</b>	This command allows the master instance to dictate the master down timer (non-owner context only).
<b>Default</b>	no master-int-inherit

message-interval

<b>Syntax</b>	<b>message-interval</b> {[ <i>seconds</i> ] [ <b>milliseconds</b> <i>milliseconds</i> ]} <b>no message-interval</b>
<b>Context</b>	config>service>vprn>if config>service>vprn>if>ipv6>vrrp
<b>Description</b>	This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different than the virtual router instance configured message-interval value will be silently discarded.  The message-interval command is available in both non-owner and owner <b>vrrp</b> <i>virtual-router-id</i> nodal contexts. If the message-interval command is not executed, the default message interval of 1 second will be used.  The <b>no</b> form of this command restores the default message interval value of 1 second to the virtual router instance.
<b>Parameters</b>	<i>seconds</i> — The number of seconds that will transpire before the advertisement timer expires.
	<b>Values</b> 1 — 255
	<b>Default</b> 1

**milliseconds** *milliseconds* — Specifies the milliseconds time interval between sending advertisement messages. This parameter is not supported on single-slot chassis.

**Values** 100 — 900

## ping-reply

<b>Syntax</b>	<b>[no] ping-reply</b>
<b>Context</b>	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
<b>Description</b>	<p>This 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.</p> <p>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). 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 setting of ping-reply configuration.</p> <p>The ping-reply command is only available in non-owner <b>vrrp</b> <i>virtual-router-id</i> nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses will be silently discarded.</p> <p>The <b>no</b> form of this command restores the default operation of discarding all ICMP Echo Request messages destined to the non-owner virtual router instance IP addresses.</p>
<b>Default</b>	no ping-reply

## policy

<b>Syntax</b>	<b>policy vrrp-policy-id</b> <b>no policy</b>
<b>Context</b>	config>service>vprn>if>vrrp config>service>vprn>if>ipv6>vrrp
<b>Description</b>	This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).
<b>Parameters</b>	<i>vrrp-policy-id</i> — Specifies a VRRP priority control policy.
<b>Values</b>	1 — 9999

## preempt

<b>Syntax</b>	<b>[no] preempt</b>
<b>Context</b>	config>service>vprn>if

config>service>vprn>if>ipv6>vrrp

**Description** The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.

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.

The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.

The default value for preempt mode is enabled.

**Default** preempt

### priority

**Syntax** **priority** *priority*  
**no priority**

**Context** config>service>vprn>if>vrrp  
config>service>vprn>if>ipv6>vrrp

**Description** The priority command provides the ability to configure a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The priority command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base-priority will be set to 100.

The **no** form of this command restores the default value of 100 to base-priority.

**Parameters** *base-priority* — The base-priority parameter configures the base priority used by the virtual router instance. If a VRRP priority control policy is not also defined, the base-priority will be the in-use priority for the virtual router instance.

**Values** 1 — 254

**Default** 100

### ssh-reply

**Syntax** [**no**] **ssh-reply**

**Context** config>service>vprn>if>vrrp

**Description** This command enables the non-owner master to reply to SSH Requests directed at the virtual router instance's 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 packets to non-owner master virtual IP addresses are silently discarded. Non-owner backup virtual routers never respond to SSH regardless of the `ssh-reply` configuration.

The `ssh-reply` command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the `ssh-reply` command is not executed, SSH packets to the virtual router instance IP addresses will be silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

**Default** no `ssh-reply`

## standby-forwarding

**Syntax** **[no] standby-forwarding**

**Context** `config>service>vprn>if>vrrp`  
`config>service>vprn>if>ipv6>vrrp`

**Description** This command allows the forwarding of packets by a standby router.  
 The **no** form of the command specifies that 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.

**Default** no `standby-forwarding`

## telnet-reply

**Syntax** **[no] telnet-reply**

**Context** `config>service>vprn>if>vrrp`  
`config>service>vprn>if>ipv6>vrrp`

**Description** This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance's 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.

When `telnet-reply` is not enabled, TCP port 23 Telnet packets 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` configuration.

The `telnet-reply` command is only available in non-owner **VRRP** nodal context. If the `telnet-reply` command is not executed, Telnet packets to the virtual router instance IP addresses will be silently discarded.

## ETH-CFM Service Commands

The **no** form of this command restores the default operation of discarding all Telnet packets destined to the non-owner virtual router instance IP addresses.

**Default** no telnet-reply

## traceroute-reply

**Syntax** [no] traceroute-reply

**Context** config>service>vprn>if>vrrp  
config>service>vprn>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



---

## PIM Commands

### pim

<b>Syntax</b>	<b>[no] pim</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command configures a Protocol Independent Multicast (PIM) instance in the VPRN service. When an PIM instance is created, the protocol is enabled. PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The router supports PIM sparse mode (PIM-SM).</p> <p>The <b>no</b> form of the command deletes the PIM protocol instance removing all associated configuration parameters.</p>
<b>Default</b>	none

### apply-to

<b>Syntax</b>	<b>apply-to {all   none}</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	<p>This command creates a PIM interface with default parameters.</p> <p>If a manually created interface or modified interface is deleted, the interface will be recreated when the <b>apply-to</b> command is executed. If PIM is not required on a specific interface, then execute a <b>shutdown</b> command.</p> <p>The <b>apply-to</b> command is saved first in the PIM configuration structure, all subsequent commands either create new structures or modify the defaults as created by the <b>apply-to</b> command.</p>
<b>Default</b>	none (keyword)
<b>Parameters</b>	<p><b>all</b> — Specifies that all VPRN and non-VPRN interfaces are automatically applied in PIM.</p> <p><b>none</b> — No interfaces are automatically applied in PIM. PIM interfaces must be manually configured.</p>

### grt-extranet

<b>Syntax</b>	<b>[no] grt-extranet</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command enables the context to configure GRT/VRF extranet for this MVPN instance.

## group-prefix

**Syntax** **group-prefix** *ip-address/mask* [*ip-address/mask...(upto 8 max)*][**starg**]  
**group-prefix any**  
**no group-prefix** *ip-address/mask*  
**no group-prefix any**

**Context** config>service>vprn>pim>rpf-select>grt-extranet

**Description** This command configures multicast group IPv4 prefixes for the multicast GRT/VRF with per group mapping extranet functionality. Multiple lines are allowed. Duplicate prefixes are ignored. Operator can either configure specific groups for extranet or specify all groups by using key-word any. The two options are mutually exclusive in configuration.

When the starg option is specified, extranet functionality is enabled for PIM ASM as for the specified group. When the option is not specified (not recommended with PIM ASM), the PIM ASM join will be mapped and data plane will be established, but the control plane will not be updated on SPT switchover, unless the switchover is driven by a CPE router on a receiver side.

The no form of the command deletes specified prefix from the list, or removes mapping of all prefixes if group-prefix any was specified.

**Default** none

**Parameters** *ip-address/mask* — Specifies the IPv4 multicast address prefix with mask.

## import

**Syntax** **import** {**join-policy** | **register-policy**} [*policy-name* [*.. policy-name*] *policy-name*]  
**no import** {**join-policy** | **register-policy**}

**Context** config>service>vprn>pim

**Description** This command specifies the import route policy to be used for determining which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context. When an import policy is not specified, BGP routes are accepted by default.

The **no** form of the command removes the policy association from the IGMP instance.

**Default** no import join-policy  
no import register-policy

**Parameters** **join-policy** — Use this command to filter PIM join messages which prevents unwanted multicast streams from traversing the network.

**register-policy** — This keyword filters register messages. PIM register filters prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message.

*policy-name* — The route policy name. Allowed values are any string up to 32 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. Route policies are configured in the **config>router>policy-options** context.

## interface

<b>Syntax</b>	<b>[no] interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	<p>This command enables PIM on an interface and enables the context to configure interface-specific parameters. By default interfaces are activated in PIM based on the <b>apply-to</b> command, and do not have to be configured on an individual basis unless the default values must be changed.</p> <p>The <b>no</b> form of the command deletes the PIM interface configuration for this interface. If the <b>apply-to</b> command parameter is configured, then the <b>no interface</b> form must be saved in the configuration to avoid automatic (re)creation after the next <b>apply-to</b> is executed as part of a reboot.</p> <p>The <b>shutdown</b> command can be used to disable an interface without removing the configuration for the interface.</p>
<b>Default</b>	Interfaces are activated in PIM based on the apply-to command.
<b>Parameters</b>	<i>ip-int-name</i> — Specify the interface name. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## assert-period

<b>Syntax</b>	<b>assert-period</b> <i>assert-period</i> <b>no assert-period</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	<p>This command configures the period in seconds for periodic refreshes of PIM Assert messages on an interface.</p> <p>The <b>no</b> form of the command reverts to the default.</p>
<b>Default</b>	60
	<i>assert-period</i> — Specifies the period, in seconds, for periodic refreshes of PIM Assert messages on an interface.
	<b>Values</b> 1 — 300

## bfd-enable

<b>Syntax</b>	<b>[no] bfd-enable [ipv4 ipv6]</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	<p>This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.</p> <p>The <b>no</b> form of this command removes BFD from the associated IGP protocol adjacency.</p>

## PIM Commands

**Default** no bfd-enable

### bsm-check-rtr-alert

**Syntax** [no] **bsm-check-rtr-alert**

**Context** config>service>vprn>pim>if

**Description** This command enables the checking of router alert option in the bootstrap messages received on this interface.

**Default** no bsm-check-rtr-alert

### hello-interval

**Syntax** **hello-interval** *hello-interval*  
**no hello-interval**

**Context** config>service>vprn>pim>if

**Description** This command configures the frequency at which PIM Hello messages are transmitted on this interface.

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

**Default** 30

**Parameters** *hello-interval* — Specifies the hello interval in seconds. A 0 (zero) value disables the sending of hello messages.

**Values** 0 — 255 seconds

### hello-multiplier

**Syntax** **hello-multiplier** *deci-units*  
**no hello-multiplier**

**Context** config>service>vprn>pim>if

**Description** This command configures the multiplier to determine the holdtime for a PIM neighbor.

The **hello-multiplier** in conjunction with the **hello-interval** determines the holdtime for a PIM neighbor.

**Parameters** *deci-units* — Specify the value, specified in multiples of 0.1, for the formula used to calculate the hello-holdtime based on the hello-multiplier:

$(\text{hello-interval} * \text{hello-multiplier}) / 10$

This allows the PIMv2 default timeout of 3.5 seconds to be supported.

**Values** 20 — 100

**Default** 35

## improved-assert

**Syntax** `[no] improved-assert`

**Context** `config>service>vprn>pim>if`

**Description** This command enables improved assert processing on this interface. The PIM assert process establishes a forwarder for a LAN and requires interaction between the control and forwarding planes.

The assert process is started when data is received on an outgoing interface. This could impact performance if data is continuously received on an outgoing interface.

When enabled, the PIM assert process is done entirely on the control-plane with no interaction between the control and forwarding plane.

**Default** enabled

## instant-prune-echo

`[no] instant-prune-echo`

## max-groups

**Syntax** `max-groups value`  
`no max-groups`

**Context** `config>service>vprn>pim>if`

**Description** This command configures the maximum number of groups for which PIM can have downstream state based on received PIM Joins on this interface. This does not include IGMP local receivers on the interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. When this object has a value of 0, there is no limit to the number of groups.

**Parameters** *value* — Specifies the maximum number of groups for this interface.

**Values** 1 — 16000

## multicast-senders

<b>Syntax</b>	<b>multicast-senders {auto   always   never}</b> <b>no multicast-senders</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	This command configures the way subnet matching is done for incoming data packets on this interface. An IP multicast sender is an user entity to be authenticated in a receiving host.
<b>Parameters</b>	<b>auto</b> — Subnet matching is automatically performed for incoming data packets on this interface. <b>always</b> — Subnet matching is always performed for incoming data packets on this interface. <b>never</b> — Subnet matching is never performed for incoming data packets on this interface.

## p2mp-ldp-tree-join

<b>Syntax</b>	<b>p2mp-ldp-tree-join</b> <b>p2mp-ldp-tree-join ipv4</b> <b>p2mp-ldp-tree-join ipv6</b> <b>p2mp-ldp-tree-join ipv4 ipv6</b> <b>no p2mp-ldp-tree-join [ipv4] [ipv6]</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	This command configures the option to join P2MP LDP tree towards the multicast source for the VPRN service. If p2mp-ldp-tree-join is enabled, a PIM multicast join received on an interface is processed to join P2MP LDP LSP using the in-band signaled P2MP tree for the same multicast flow. LDP P2MP tree is setup towards the multicast source. Route to source of the multicast node is looked up from the RTM. The next-hop address for the route to source is set as the root of LDP P2MP tree.  The <b>no</b> form of command disables joining P2MP LDP tree for IPv4 or IPv6 or both (if both or none is specified).
<b>Parameters</b>	<b>ipv4</b> — Enables dynamic mLDP in-band signaling for IPv4 PIM joins. IPv4 multicast must be enabled; see <b>ipv4-multicast-disable</b> . For backward compatibility <b>p2mp-ldp-tree-join</b> is equivalent to <b>p2mp-ldp-tree-join ipv4</b> .  <b>ipv6</b> — Enables dynamic mLDP in-band signaling for IPv6 PIM joins. IPv6 multicast must be enabled; see <b>ipv6-multicast-disable</b> .
<b>Default</b>	no p2mp-ldp-tree-join

## priority

<b>Syntax</b>	<b>priority <i>dr-priority</i></b> <b>no priority</b>
<b>Context</b>	config>service>vprn>pim>if

<b>Description</b>	This command sets the priority value to become the rendezvous point (RP) that is included in bootstrap messages sent by the router. The RP is sometimes called the bootstrap router. The <b>priority</b> command indicates whether the router is eligible to be a bootstrap router. The <b>no</b> form of the command disqualifies the router to participate in the bootstrap election.
<b>Default</b>	1 (The router is the least likely to become the designated router.)
<b>Parameters</b>	<i>dr-priority</i> — Specifies the priority to become the designated router. The higher the value, the higher the priority.
<b>Values</b>	1 — 4294967295

## sticky-dr

<b>Syntax</b>	<b>sticky-dr [priority <i>dr-priority</i>]</b> <b>no sticky-dr</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	This command enables sticky-dr operation on this interface. When enabled, the priority in PIM hellos sent on this interface when elected as the designated router (DR) will be modified to the value configured in <i>dr-priority</i> . This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.  By enabling <b>sticky-dr</b> on this interface, it will continue to act as the DR for the LAN even after the old DR comes back up.  The <b>no</b> form of the command disables sticky-dr operation on this interface.
<b>Default</b>	disabled
<b>Parameters</b>	<b>priority</b> <i>dr-priority</i> — Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled.
<b>Values</b>	1 — 4294967295

## three-way-hello

<b>Syntax</b>	<b>three-way-hello [compatibility-mode]</b> <b>no three-way-hello</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	This command configures the compatibility mode for enabling the three way hello.
<b>Parameters</b>	<b>compatibility-mode</b> — Specifies to enable the three way hello.

## tracking-support

<b>Syntax</b>	<b>[no] tracking-support</b>
<b>Context</b>	config>service>vprn>pim>if
<b>Description</b>	This command sets the the T bit in the LAN Prune Delay option of the Hello Message. This indicates the router's capability to disable Join message suppression.
<b>Default</b>	no tracking-support

## ipv4-multicast-disable

<b>Syntax</b>	<b>[no] ipv4-multicast-disable</b>
<b>Context</b>	config>service>vprn>pim config>service>vprn>pim>interface
<b>Description</b>	This command administratively disables/enables PIM operation for IPv4.
<b>Default</b>	no ipv4-multicast-disable

## ipv6-multicast-disable

<b>Syntax</b>	<b>ipv6-multicast-disable</b>
<b>Context</b>	config>service>vprn>pim config>service>vprn>pim>interface
<b>Description</b>	This command administratively disables/enables PIM operation for IPv6.
<b>Default</b>	ipv6-multicast-disable

## mc-ecmp-balance

<b>Syntax</b>	<b>[no] mc-ecmp-balance</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command enables multicast balancing of traffic over ECMP links. When enabled, each multicast stream that needs to be forwarded over an ECMP link will be re-evaluated for the total multicast bandwidth utilization. Re-evaluation occurs on the ECMP interface in question.  The <b>no</b> form of the command disables the multicast balancing.



## mc-ecmp-balance-hold

<b>Syntax</b>	<b>mc-ecmp-balance-hold</b> <i>minutes</i> <b>no mc-ecmp-balance-hold</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command configures the hold time for multicast balancing over ECMP links.
<b>Parameters</b>	<i>minutes</i> — Specifies the hold time, in minutes, that applies after an interface has been added to the ECMP link.

## mc-ecmp-hashing-enabled

<b>Syntax</b>	<b>[no] mc-ecmp-hashing-enabled</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	<p>This command distributes PIM joins over the multiple ECMP paths based on a hash of S and G. When a link in the ECMP set is removed, the multicast streams that were using that link are re-distributed over the remaining ECMP links using the same hash algorithm. When a link is added to the ECMP set, new joins may be allocated to the new link based on the hash algorithm, but existing multicast streams using the other ECMP links stay on those links until they are pruned.</p> <p>The <b>no mc-ecmp-hashing-enabled</b> form of the command means that the use of multiple ECMP paths if enabled at the <b>config&gt;router</b> or <b>config&gt;service&gt;vprn</b> context is controlled by the existing implementation and CLI commands <b>mc-ecmp-balance</b>.</p>
<b>Default</b>	no mc-ecmp-hashing-enabled

## non-dr-attract-traffic

<b>Syntax</b>	<b>[no] non-dr-attract-traffic</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	<p>This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designater router.</p> <p>An operator can configure an interface (router or IES or VPRN interfaces) to IGMP and PIM. The interface IGMP state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to attract traffic to both routers a flag non-dr-attract-traffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. Note that while using this flag the router may not send the stream down to the DSLAM while not DR.</p> <p>When enabled, the designated router state is ignored. When disabled, <b>no non-dr-attract-traffic</b>, the designated router value is honored.</p>
<b>Default</b>	no non-dr-attract-traffic

## PIM Commands

### rp

<b>Syntax</b>	<b>rp</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	<p>This command enables access to the context to configure the rendezvous point (RP) ) of a PIM protocol instance.</p> <p>An Alcatel-Lucent PIM router acting as an RP must respond to a PIM register message specifying an SSM multicast group address by sending to the first hop router stop register message(s). It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be either from the SSM default range of 232/8 or from a multicast group address range that was explicitly configured for SSM.</p>
<b>Default</b>	rp enabled when PIM is enabled.

### anycast

<b>Syntax</b>	<b>[no] anycast</b> <i>rp-ip-address</i>
<b>Context</b>	config>service>vprn>pim>rp
<b>Description</b>	<p>This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.</p> <p>The <b>no</b> form of the command removes the anycast instance from the configuration.</p>
<b>Default</b>	none
<b>Parameters</b>	<p><i>rp-ip-address</i> — Configure the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no ip-address is entered then the command is simply used to enter the anycast CLI level.</p> <p><b>Values</b> Any valid loopback address configured on the node.</p>

### rp-set-peer

<b>Syntax</b>	<b>[no] rp-set-peer</b> <i>ip-address</i>
<b>Context</b>	config>service>vprn>pim>rp>anycast
<b>Description</b>	<p>This command configures a peer in the anycast rp-set. The address identifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.</p> <p>This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this rp-set.</p>

Although there is no set maximum of addresses that can be configured in an rp-set, up to 15 multicast addresses is recommended.

The **no** form of the command removes an entry from the list.

**Default** None

**Parameters** *ip-address* — Specifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.

## auto-rp-discovery

**Syntax** **[no] auto-rp-discovery**

**Context** config>service>vprn>pim>rp

**Description** This command enables Auto-RP protocol in discovery mode. In discovery mode, RP-mapping and RP-candidate messages are received and forwarded to downstream nodes. RP-mapping messages are received locally to learn about availability of RP nodes present in the network.

Note that chassis mode D or higher must be enabled for **auto-rp-discovery**.

The **no** form of the command disables auto RP.

**Default** disabled

## bootstrap-export

**Syntax** **bootstrap-export** *policy-name* [*policy-name*... up to five]  
**no bootstrap-export**

**Context** config>service>vprn>pim>rp

**Description** This command exports policies to control the flow of bootstrap messages from the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

**Default** none

**Parameters** *policy-name* — Specify the policy name. The policy statement must already be configured in the config>router>policy-options context.

## bootstrap-import

**Syntax** **bootstrap-import** *policy-name* [*policy-name*... up to five]  
**no bootstrap-import** *policy-name* [*policy-name*... up to five]

**Context** config>service>vprn>pim>rp

## PIM Commands

<b>Description</b>	This command imports policies to control the flow of bootstrap messages into the RP. Up to five policies can be defined. The <b>no</b> form of this command removes the specified policy names from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>policy-name</i> — Specify the policy name. The policy statement must already be configured in the config>router>policy-options context.

## bsr-candidate

<b>Syntax</b>	<b>bsr-candidate</b>
<b>Context</b>	config>service>vprn>pim>rp config>service>vprn>pim>rp>ipv6
<b>Description</b>	This command enables the context to configure a local rendezvous point (RP) of a PIM protocol instance.
<b>Default</b>	Enabled when PIM is enabled.

## address

<b>Syntax</b>	<b>[no] address</b> <i>ip-address</i>
<b>Context</b>	config>service>vprn>pim>rp>bsr-candidate config>service>vprn>pim>rp>rp-candidate
<b>Description</b>	This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router. Use the <b>no</b> form of this command to remove the static RP from the configuration.
<b>Default</b>	No IP address is specified.
<b>Parameters</b>	<i>ip-address</i> — The static IP address of the RP. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.
<b>Values</b>	1.0.0.0 – 223.255.255.255

## address

<b>Syntax</b>	<b>[no] address</b> <i>ipv6-address</i>
<b>Context</b>	config>service>vprn>pim>rp>ipv6>bsr-candidate config>service>vprn>pim>rp>ipv6>rp-candidate
<b>Description</b>	This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router.

Use the **no** form of this command to remove the static RP from the configuration.

<b>Default</b>	No IP address is specified.
<b>Parameters</b>	<i>ipv6-address</i> — The static IP address of the RP. The <i>ip-address</i> portion of the <b>address</b> command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.
<b>Values</b>	<pre> ipv6-address      : x:x:x:x:x:x:x (eight 16-bit pieces)                   : x:x:x:x:x:d.d.d                   : x [0..FFFF]H                   : d [0..255]D </pre>

## hash-mask-len

<b>Syntax</b>	<b>hash-mask-len</b> <i>hash-mask-length</i> <b>no hash-mask-len</b>
<b>Context</b>	config>service>vprn>pim>rp>bsr-candidate
<b>Description</b>	This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.
<b>Default</b>	30
<b>Parameters</b>	<i>hash-mask-length</i> — The hash mask length.
<b>Values</b>	0 — 32

## hash-mask-length

<b>Syntax</b>	<b>hash-mask-length</b> <i>hash-mask-length</i> <b>no hash-mask-length</b>
<b>Context</b>	config>service>vprn>pim>rp>ipv6>bsr-candidate
<b>Description</b>	This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.
<b>Default</b>	126
<b>Parameters</b>	<i>hash-mask-length</i> — The hash mask length.
<b>Values</b>	0 — 128

## priority

<b>Syntax</b>	<b>priority</b> <i>bootstrap-priority</i>
<b>Context</b>	config>service>vprn>pim>rp>bsr-candidate config>service>vprn>pim>rp>ipv6>bsr-candidate
<b>Description</b>	This command defines the priority used to become the rendezvous point (RP) . The higher the priority value the more likely that this router becomes the RP. If there is a tie, the router with the highest IP address is elected.
<b>Parameters</b>	<i>bootstrap-priority</i> — The priority to become the bootstrap router.
	<b>Values</b> 0 — 255
	<b>Default</b> 0 (the router is not eligible to be the bootstrap router)

## ipv6

<b>Syntax</b>	<b>ipv6</b>
<b>Context</b>	config>service>vprn>pim>rp
<b>Description</b>	This command enables access to the context to configure the rendezvous point (RP) of a PIM IPv6 protocol instance.  An Alcatel-Lucent IPv6 PIM router acting as an RP must respond to an IPv6 PIM register message specifying an SSM multicast group address by sending to the first hop router stop register message(s). It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be either from the SSM default range or from a multicast group address range that was explicitly configured for SSM.
<b>Default</b>	ipv6 RP enabled when IPv6 PIM is enabled.

## anycast

<b>Syntax</b>	<b>anycast</b> <i>ipv6-address</i> <b>no anycast</b> <i>ipv6-address</i>
<b>Context</b>	config>service>vprn>pim>rp>ipv6
<b>Description</b>	This command configures an IPv6 PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.  The <b>no</b> form of the command removes the anycast instance from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ipv6-address</i> — Configures the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If another anycast command is entered with an address then the old address will be replaced with the new address. If no address is



## group-range

<b>Syntax</b>	<b>[no] group-range</b> <i>{ipv6-address/prefix-length}</i>
<b>Context</b>	config>service>vprn>pim>rp>ipv6>embedded-rp config>service>vprn>pim>rp>ipv6>rp-candidate
<b>Description</b>	This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).  Use the no form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ipv6-address</i> — Specify the addresses or address ranges that this router can be an RP. <i>prefix-length</i> — Specify the address prefix length.
<b>Values</b>	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  <i>prefix-length</i> [8..128] //for embedded-rp <i>prefix-length</i> [16..128] //for rp-candidate

## group-prefix

<b>Syntax</b>	<b>[no] group-prefix grp-ipv6-address/prefix-length</b>
<b>Context</b>	config>service>vprn>pim>rp>ipv6>static
<b>Description</b>	The group-prefix for a static-rp defines a range of multicast-ip-addresses for which this static RP is applicable.  The no form of the command removes the criterion.
<b>Default</b>	none
<b>Parameters</b>	<i>grp-ipv6-address</i> — Specifies the multicast IPv6 address. <i>prefix-length</i> — Specifies the address prefix length.
<b>Values</b>	grp-ipv6-address 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  <i>prefix-length</i> [8..128]



## rp-candidate

<b>Syntax</b>	<b>rp-candidate</b>
<b>Context</b>	config>service>vprn>pim>rp config>service>vprn>pim>rp>ipv6
<b>Description</b>	This command enables the context to configure the candidate rendezvous point (RP) parameters.
<b>Default</b>	Enabled when PIM is enabled.

## group-range

<b>Syntax</b>	<b>[no] group-range</b> { <i>ip-prefix/mask</i>   <i>ip-prefix netmask</i> }																
<b>Context</b>	config>service>vprn>pim>rp>rp-candidate config>service>vprn>pim>ssm																
<b>Description</b>	This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).  Use the <b>no</b> form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.																
<b>Default</b>	none																
<b>Parameters</b>	<i>ip-prefix</i> — Specify the addresses or address ranges that this router can be an RP.  <table> <tr> <td><b>Values</b></td> <td>ipv4-prefix - a.b.c.d</td> </tr> <tr> <td></td> <td>ipv4-prefix-le - [0..32]</td> </tr> <tr> <td></td> <td>ipv6-prefix - x:x:x:x:x:x:x (eight 16-bitpieces)</td> </tr> <tr> <td></td> <td>x:x:x:x:x:d.d.d.d</td> </tr> <tr> <td></td> <td>x - [0..FFFF]H</td> </tr> <tr> <td></td> <td>d - [0..255]D</td> </tr> <tr> <td></td> <td>ipv6-prefix-le - [0..128]</td> </tr> </table> <i>mask</i> — Specify the address mask with the address to define a range of addresses. <i>netmask</i> — Specify the subnet mask in dotted decimal notation. <table> <tr> <td><b>Values</b></td> <td>:a.b.c.d (network bits all 1 and host bits all 0)</td> </tr> </table>	<b>Values</b>	ipv4-prefix - a.b.c.d		ipv4-prefix-le - [0..32]		ipv6-prefix - x:x:x:x:x:x:x (eight 16-bitpieces)		x:x:x:x:x:d.d.d.d		x - [0..FFFF]H		d - [0..255]D		ipv6-prefix-le - [0..128]	<b>Values</b>	:a.b.c.d (network bits all 1 and host bits all 0)
<b>Values</b>	ipv4-prefix - a.b.c.d																
	ipv4-prefix-le - [0..32]																
	ipv6-prefix - x:x:x:x:x:x:x (eight 16-bitpieces)																
	x:x:x:x:x:d.d.d.d																
	x - [0..FFFF]H																
	d - [0..255]D																
	ipv6-prefix-le - [0..128]																
<b>Values</b>	:a.b.c.d (network bits all 1 and host bits all 0)																

## holdtime

<b>Syntax</b>	<b>holdtime</b> <i>holdtime</i> <b>no holdtime</b> <i>holdtime</i>
<b>Context</b>	config>service>vprn>pim>rp>rp-candidate config>service>vprn>pim>rp>ipv6>rp-candidate
<b>Description</b>	Use this command to define the length of time neighboring router consider this router to be up. Use the <b>no</b> form of this command to revert to the default value.

## PIM Commands

**Default** 150

**Parameters** *holdtime* — Specify the length of time, in seconds, that neighbor should consider the sending router to be operational.

**Values** 0 — 255

## priority

**Syntax** **priority** *priority*  
**no priority** *priority*

**Context** config>router>pim>rp>local  
config>service>vprn>pim>rp>rp-candidate

**Description** This command defines the priority used to become the rendezvous point (RP). The higher the priority value, the more likely that this router will become the RP.

Use the **no** form of this command to revert to the default value.

**Default** 1

**Parameters** *priority* — Specify the priority to become the designated router. The higher the value the more likely the router will become the RP.

**Values** 0 — 255

## static

**Syntax** **static**

**Context** config>service>vprn>pim>rp

**Description** This command enables access to the context to configure a static rendezvous point (RP) of a PIM-SM protocol instance.

**Default** none

## address

**Syntax** [**no**] **address** *ip-address*

**Context** config>service>vprn>pim>rp>static

**Description** This command configures the static rendezvous point (RP) address.

The **no** form of this command removes the static RP entry from the configuration.

**Default** none

## group-prefix

<b>Syntax</b>	<b>[no] group-prefix</b> { <i>grp-ip-address/mask</i>   <i>grp-ip-address netmask</i> }
<b>Context</b>	config>service>vprn>pim>rp>static
<b>Context</b>	The <b>group-prefix</b> for a static-rp defines a range of multicast-ip-addresses for which a certain RP is applicable. The <b>no</b> form of the command removes the criterion.
<b>Default</b>	none
<b>Parameters</b>	<i>grp-ip-address</i> — Specify the multicast IP address. <i>mask</i> — Defines the mask of the multicast-ip-address. <b>Values</b> 4 — 32 <i>netmask</i> — Enter the subnet mask in dotted decimal notation. <b>Values</b> 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

## override

<b>Syntax</b>	<b>[no] override</b>
<b>Context</b>	config>service>vprn>pim>rp>static
<b>Description</b>	This command changes the precedence of static RP over dynamically learned Rendezvous Point (RP). When enabled, the static group-to-RP mappings take precedence over the dynamically learned mappings.
<b>Default</b>	no override

## rpf-table

<b>Syntax</b>	<b>[no] rpf-table</b> { <i>rtable-m</i>   <i>rtable-u</i>   <b>both</b> }
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command configures the sequence of route tables used to find a Reverse Path Forwarding (RPF) interface for a particular multicast route. By default, only the unicast route table is looked up to calculate RPF interface towards the source/ rendezvous point. However, the operator can specify the following: <ul style="list-style-type: none"> <li>a) Use unicast route table only</li> <li>b) Use multicast route table only or</li> <li>c) Use both the route tables.</li> </ul>

## PIM Commands

<b>Default</b>	rpf-table rtable-u
<b>Parameters</b>	<b>rtable-m</b> — pecified that only the multicast route table is to be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by static routes and OSPF. <b>rtable-u</b> — Specifies that only the unicast route table is to be used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by all the unicast routing protocols. <b>both</b> — Specifies that PIM always lookup first in the multicast route table, and if there is a route, PIM use it. If PIM does not find a route in the first lookup, it will try to find it in the unicast route table. rtable-m is checked before rtable-u.

## spt-switchover-threshold

<b>Syntax</b>	<b>spt-switchover-threshold</b> { <i>grp-ip-address/mask</i>   <i>grp-ip-address netmask</i> } <i>spt-threshold</i> <b>no spt-switchover-threshold</b> { <i>grp-ip-address/mask</i>   <i>grp-ip-address netmask</i> }
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command configures a shortest path tree (SPT tree) switchover threshold for a group prefix.
<b>Parameters</b>	<i>grp-ip-address</i> — Specify the multicast group address. <i>mask</i> — Defines the mask of the multicast-ip-address. <b>Values</b> 4 — 32 <i>netmask</i> — Enter the subnet mask in dotted decimal notation. <b>Values</b> 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0) <i>spt-threshold</i> — Specifies the configured threshold in kilo-bits per second(kbps) for the group to which this (S,G) belongs. For a group G configured with a threshold, switchover to SPT for an (S,G) is attempted only if the (S,G)'s rate exceeds this configured threshold.

## ssm-assert-compatible-mode

<b>Syntax</b>	<b>ssm-assert-compatible-mode</b> [enable disable]
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command specifies whether SSM assert is enabled in compatibility mode for this PIM protocol instance. When enabled, for SSM groups, PIM will consider the SPT bit to be implicitly set to compute the value of CouldAssert (S,G,I) as defined in RFC 4601, <i>Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)</i> . When disabled, for SSM groups, PIM will not assume the SPT bit to be set. The SPT bit will be set by Update_SPTbit(S,G,iif) macro defined in RFC 4601.
<b>Default</b>	disable
<b>Parameters</b>	<b>enable</b> — Enables SSM assert in compatibility mode for this PIM protocol instance. <b>disable</b> — Disabled SSM assert in compatibility mode for this PIM protocol instance.

## ssm-default-range-disable

<b>Syntax</b>	<b>ssm-default-range-disable ipv4</b>
<b>Context</b>	config>service>vprn>pim
<b>Description</b>	This command specifies whether to disable the use of default range (232/8) for SSM so that it can be used by ASM to process (*,G). When enabled, the use of default range is disabled for SSM and it can be used by ASM. When disabled, the SSM default range is enabled.
<b>Default</b>	disable

## ssm-groups

<b>Syntax</b>	<b>[no] ssm-groups</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables access to the context to enable a source-specific multicast (SSM) configuration instance.
<b>Default</b>	none

---

## C-MLDP Commands

### mld

<b>Syntax</b>	<b>[no] mld</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure Multicast Listener Discovery (MLD) parameters. The <b>no</b> form of the command disables MLD.
<b>Default</b>	no mld

### interface

<b>Syntax</b>	<b>[no] interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn>mld
<b>Description</b>	This command enables the context to configure an Multicast Listener Discovery (MLD) interface. The interface is a local identifier of the network interface on which reception of the specified multicast address is to be enabled or disabled.  The <b>no</b> form of the command deletes the MLD interface. The <b>shutdown</b> command in the <b>config&gt;router&gt;mld&gt;interface</b> context can be used to disable an interface without removing the configuration for the interface.
<b>Default</b>	<b>no interface</b> — No interfaces are defined.
<b>Parameters</b>	<i>ip-int-name</i> — The IP interface name. Interface names must be unique within the group of defined IP interfaces for <b>config router interface</b> and <b>config service ies interface</b> commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 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.  If the IP interface name does not exist or does not have an IP address configured an error message will be returned.  If the IP interface exists in a different area it will be moved to this area.

### disable-router-alert-check

<b>Syntax</b>	<b>[no] disable-router-alert-check</b>
<b>Context</b>	config>service>vprn>mld>interface
<b>Description</b>	This command enables router alert checking for MLD messages received on this interface.

The no form of the command disables the router alert checking.

**Default** none

## import

**Syntax** **import** *policy-name*  
**no import**

**Context** config>service>vprn>mld>interface

**Description** This command specifies the import route policy to be used for determining which membership reports are accepted by the router. Route policies are configured in the **config>router>policy-options** context.

When an import policy is not specified, all the MLD reports are accepted.

The **no** form of the command removes the policy association from the MLD instance.

**Default** **no import** — No import policy specified.

**Parameters** *policy-name* — The route policy name. Allowed values are any string up to 32 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. Route policies are configured in the **config>router>policy-options** context.

## max-groups

**Syntax** **max-groups** *value*  
**no max-groups**

**Context** config>service>vprn>mld>interface

**Description** This command specifies the maximum number of groups for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

**Default** 0, no limit to the number of groups.

**Parameters** *value* — Specifies the maximum number of groups for this interface.

**Values** 1 — 16000

## query-interval

<b>Syntax</b>	<b>query-interval</b> <i>seconds</i> <b>no query-interval</b>
<b>Context</b>	config>service>vprn>mld>interface
<b>Description</b>	This command specifies the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.
<b>Default</b>	125
<b>Parameters</b>	<i>seconds</i> — The time frequency, in seconds, that the router transmits general host-query messages. <b>Values</b> 2 — 1024

## query-last-member-interval

<b>Syntax</b>	<b>query-last-member-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vprn>mld>interface
<b>Description</b>	This command configures the frequency at which the querier sends group-specific query messages including messages sent in response to leave-group messages. The lower the interval, the faster the detection of the loss of the last member of a group.
<b>Default</b>	1
<b>Parameters</b>	<i>seconds</i> — Specifies the frequency, in seconds, at which query messages are sent. <b>Values</b> 1 — 1024

## query-response-interval

<b>Syntax</b>	<b>query-response-interval</b> <i>seconds</i>
<b>Context</b>	config>service>vprn>mld>interface
<b>Description</b>	This command specifies how long the querier router waits to receive a response to a host-query message from a host.
<b>Default</b>	10
<b>Parameters</b>	<i>seconds</i> — Specifies the the length of time to wait to receive a response to the host-query message from the host. <b>Values</b> 1 — 1023



## static

<b>Syntax</b>	<b>static</b>
<b>Context</b>	config>service>vprn>mld>interface
<b>Description</b>	This command tests multicast forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.
<b>Default</b>	none

## group

<b>Syntax</b>	<b>[no] group</b> <i>ipv6-address</i>
<b>Context</b>	config>service>vprn>mld>interface>static
<b>Description</b>	<p>This command enables the context to add a static multicast group either as a (*,G) or one or more (S,G) records. Use MLD static group memberships to test multicast forwarding without a receiver host. When MLD static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.</p> <p>When static MLD group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static MLD group entries do not generate join messages toward the RP.</p> <p>The <b>no</b> form of the command removes the IPv6 address from the configuration.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>ipv6-address</i> — Specifies an MLD multicast group address that receives data on an interface. The IP address must be unique for each static group.

## source

<b>Syntax</b>	<b>[no] source</b> <i>ipv6-address</i>
<b>Context</b>	config>service>vprn>mld>interface>static>group
<b>Description</b>	<p>This command specifies an IPv6 unicast address that sends data on an interface. This enables a multicast receiver host to signal a router the group to receive multicast traffic from, and from the source(s) that the traffic is expected.</p> <p>The <b>source</b> command is mutually exclusive with the specification of individual sources for the same group.</p> <p>The source command, in combination with the group, is used to create a specific (S,G) static group entry.</p> <p>The <b>no</b> form of the command removes the source from the configuration.</p>
<b>Default</b>	none
<b>Parameters</b>	<i>ip-address</i> — Specifies the IPv6 unicast address.

## starg

<b>Syntax</b>	<b>[no] starg</b>
<b>Context</b>	config>service>vprn>mld>interface>static>group
<b>Description</b>	This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.  Use the <b>no</b> form of the command to remove the starg entry from the configuration.
<b>Default</b>	none

## version

<b>Syntax</b>	<b>version <i>version</i></b> <b>no version</b>
<b>Context</b>	config>service>vprn>mld>interface
<b>Description</b>	This command specifies the MLD version. If routers run different versions, they will negotiate the lowest common version of MLD that is supported by hosts on their subnet and operate in that version. For MLD to function correctly, all routers on a LAN should be configured to run the same version of MLD on that LAN.
<b>Default</b>	1
<b>Parameters</b>	<i>version</i> — Specifies the MLD version number.  <b>Values</b> 1, 2

## robust-count

<b>Syntax</b>	<b>robust-count <i>robust-count</i></b> <b>no robust-count</b>
<b>Context</b>	config>service>vprn>mld
<b>Description</b>	This command configures the robust count. The robust-count variable allows tuning for the expected packet loss on a subnet. If a subnet anticipates losses, the robust-count variable can be increased.
<b>Default</b>	2
<b>Parameters</b>	<i>robust-count</i> — Specify the robust count value.  <b>Values</b> 2 — 10

## ssm-translate

<b>Syntax</b>	<b>ssm-translate</b>
<b>Context</b>	config>service>vprn>mld
<b>Description</b>	This command enables the context to configure group ranges which are translated to SSM (S,G) entries. If the static entry needs to be created, it has to be translated from a IGMPv1 IGMPv2 request to a Source Specific Multicast (SSM) join. An SSM translate source can only be added if the <b>starg</b> command is not enabled. An error message is generated if you try to configure the <b>source</b> command with <b>starg</b> command enabled.

## grp-range

<b>Syntax</b>	<b>[no] grp-range start end</b>
<b>Context</b>	config>service>vprn>mld>ssm-translate
<b>Description</b>	This command is used to configure group ranges which are translated to SSM (S,G) entries.
<b>Parameters</b>	<p><i>start</i> — An IP address that specifies the start of the group range.</p> <p><i>end</i> — An IP address that specifies the end of the group range. This value should always be greater than or equal to the value of the <i>start</i> value.</p>

## source

<b>Syntax</b>	<b>[no] source ip-address</b>
<b>Context</b>	config>service>vprn>mld>ssm-translate>grp-range
<b>Description</b>	This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by <b>grp-range start</b> and <b>end</b> parameters, it is translated to an (S,G) report with the value of this object as the source address.
<b>Parameters</b>	<i>ip-address</i> — Specifies the IP address that will be sending data.

## Network Interface Commands

### network-interface

- Syntax** `network-interface interface-name [create]`  
`no network-interface interface-name`
- Context** config>service>vprn
- Description** This command configures a network interface in a VPRN that acts as a CSC interface to a CSC-CE in a Carrier Supporting Carrier IP VPN deployment model.

---

## Network Ingress Commands

### network

<b>Syntax</b>	<b>network</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure network parameters for the VPRN service.

### ingress

<b>Syntax</b>	<b>ingress</b>
<b>Context</b>	config>service>vprn>network
<b>Description</b>	This command enables the context to configure network ingress parameters for the VPRN service.

### qos

<b>Syntax</b>	<b>qos network-policy-id fp-redirect-group queue-group-name instance instance-id</b> <b>no qos</b>
<b>Context</b>	config>service>vprn>network>ingress
<b>Description</b>	<p>This command is used to redirect unicast packets arriving on an automatically (using the <b>auto-bind</b> command) or manually configured (using a <b>spoke-sdp</b> command, but not the <b>spoke-sdp</b> command under the VPRN IP interface) binding in a VPRN to a policer in an ingress forwarding plane queue-group for the purpose of rate-limiting.</p> <p>For the policer to be used, the following must be true:</p> <ol style="list-style-type: none"> <li>1. The configured queue group template name must be applied to the forwarding plane on which the ingress traffic arrives using the instance id specified.</li> <li>2. The policer referenced in the FC-to-policer mappings in the ingress context of a network QoS policy must be present in the specified queue group template.</li> </ol> <p>The command will fail if the queue group template name does not exist or if the policer specified in the network QoS policy does not exist in the queue group template. If the queue group template name with the specified instance is not applied to the forwarding plane on which the VPRN binding unicast traffic arrives then this traffic will use the ingress network queues related to the network interface, however, the ingress classification is still based on the applied network QoS policy.</p> <p>The unicast traffic can be redirected to a policer under the forwarding class <b>fp-redirect-group</b> command in the ingress section of a network QoS policy; any <b>fp-redirect-group multicast-policer</b>, <b>broadcast-policer</b> or <b>unknown-policer</b> commands are ignored for this traffic. Note that multicast traffic would use the ingress network queues or queue group related to the network interface.</p>

## Network Ingress Commands

Ingress classification is based on the configuration of the ingress section of the specified network QoS policy, noting that the dot1p and exp classification is based on the outer Ethernet header and MPLS label whereas the DSCP applies to the outer IP header if the tunnel encapsulation is GRE, or the DSCP in the first IP header in the payload if **ler-use-dscp** is enabled in the ingress section of the referenced network QoS policy.

When this command is applied, it overrides the QoS applied to the related network interfaces for unicast traffic arriving on bindings in that VPRN.

The **no** version of this command removes the redirection of VPRN binding traffic to the queue-group policers.

**Parameters** *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

**Values** 1—65535

**fp-redirect-group** *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

**instance** *instance-id* — Specifies the identification of a specific instance of the queue-group.

**Values** 1—65535

## BGP Commands

### bgp

<b>Syntax</b>	<b>[no] bgp</b>
<b>Context</b>	service>vprn
<b>Description</b>	This command enables the BGP protocol with the VPRN service. The <b>no</b> form of the command disables the BGP protocol from the given VPRN service.
<b>Default</b>	no bgp

### bgp-shared-queue

<b>Syntax</b>	<b>bgp-shared-queue [cir rate] [pir rate]</b> <b>no bgp-shared-queue</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables all BGP peers within a VPRN instance to share a single CPM queue. This command takes affect on new BGP connections established; already established BGP peers continue to use their own CPM queue. Any changes to PIR/CIR of the shared queue takes effect only after BGP connections are re-established.
<b>Parameters</b>	<b>cir rate</b> — Specifies the CIR rate for the shared queue. <b>pir rate</b> — Specifies the PIR rate for the shared queue.

### advertise-inactive

<b>Syntax</b>	<b>[no] advertise-inactive</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command enables or disables the advertising of inactive BGP routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.
<b>Default</b>	no advertise-inactive

## aggregator-id-zero

<b>Syntax</b>	<b>[no] aggregator-id-zero</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.</p> <p>When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.</p> <p>When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.</p> <p>The <b>no</b> form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.</p> <p>The <b>no</b> form of the command used at the group level reverts to the value defined at the group level.</p> <p>The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

## always-compare-med

<b>Syntax</b>	<b>always-compare-med {zero   infinity}</b> <b>no always-compare-med strict-as {zero   infinity}</b> <b>no always-compare-med</b>
<b>Context</b>	config>router>bgp>best-path-selection
<b>Description</b>	<p>This command configures the comparison of BGP routes based on the MED attribute. The default behavior of SR-OS (equivalent to the <b>no</b> form of the command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The <b>always-compare-med</b> command without the <b>strict-as</b> keyword allows MED to be compared even if the paths have a different neighbor AS; in this case, if neither <b>zero</b> or <b>infinity</b> is specified, the <b>zero</b> option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the <b>strict-as</b> keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, <b>zero</b> or <b>infinity</b> is mandatory and tells BGP how to interpret paths without a MED attribute.</p>
<b>Default</b>	<b>no always-compare-med</b>
<b>Parameters</b>	<p><b>zero</b> — Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.</p> <p><b>infinity</b> — Specifies for routes learned without a MED attribute that a value of infinity (<math>2^{32}-1</math>) is used in the MED comparison. This in effect makes these routes the least desirable.</p>



**strike-as** — Specifies BGP paths to be compared even with different neighbor AS.

## as-path-ignore

<b>Syntax</b>	<b>[no] as-path-ignore</b>
<b>Context</b>	config>service>vprn>bgp
<b>Description</b>	This command determines whether the AS path is used to determine the best BGP route. If this option is present, the AS paths of incoming routes are not used in the route selection process. The <b>no</b> form of the command removes the parameter from the configuration.
<b>Default</b>	no as-path-ignore

## deterministic-med

<b>Syntax</b>	<b>[no] deterministic-med</b>
<b>Context</b>	config>service>vprn>bgp>best-path-selection
<b>Description</b>	This command controls how the BGP decision process compares routes on the basis of MED. When <b>deterministic-med</b> is configured, BGP groups paths that are equal up to the MED comparison step based on neighbor AS, and then compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without <b>deterministic-med</b> , the overall best path selection is sometimes dependent on the order of the route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.
<b>Default</b>	no deterministic-med

## as-override

<b>Syntax</b>	<b>[no] as-override</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.  This command breaks BGP's loop detection mechanism. It should be used carefully.
<b>Default</b>	as-override is not enabled by default.

## authentication-key

<b>Syntax</b>	<b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ] <b>no authentication-key</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command configures the BGP authentication key.  Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.  The no form of the command removes the authentication password from the configuration and effectively disables authentication.
<b>Default</b>	Authentication is disabled and the authentication password is empty.
<b>Parameters</b>	<i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).  <i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 342 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.  <b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> 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 <b>hash</b> parameter specified.  <b>hash2</b> — Specifies the key is entered in a more complex encrypted form. If the <b>hash2</b> parameter is not used, the less encrypted <b>hash</b> form is assumed.

## auth-keychain

<b>Syntax</b>	<b>auth-keychain</b> <i>name</i>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command configures the BGP authentication key for all peers.  The keychain allows the rollover of authentication keys during the lifetime of a session.
<b>Default</b>	<b>no auth-keychain</b>
<b>Parameters</b>	<i>name</i> — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

## backup-path

<b>Syntax</b>	<b>[no] backup-path [ipv4] [ipv6]</b>
<b>Context</b>	config>router config>router>bgp config>service>vprn>bgp
<b>Description</b>	This command enables the computation and use of a backup path for IPv4 and/or IPv6 BGP-learned prefixes belonging to the base router or a particular VPRN. Multiple paths must be received for a prefix in order to take advantage of this feature. When a prefix has a backup path and its primary path(s) fail the affected traffic is rapidly diverted to the backup path without waiting for control plane re-convergence to occur. When many prefixes share the same primary path(s), and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes. In some cases prefix independent convergence may require use of FP2 or later IOMs/IMMs/XMAs.  By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.
<b>Default</b>	<b>no backup-path</b>
<b>Parameters</b>	<b>ipv4</b> — enable the use of a backup path for BGP-learned IPv4 prefixes. <b>ipv6</b> — enable the use of a backup path for BGP-learned IPv6 prefixes.

## best-path-selection

<b>Syntax</b>	<b>best-path-selection</b>
<b>Context</b>	config>service>vprn>bgp
<b>Description</b>	This command enables path selection configuration.

## ignore-nh-metric

<b>Syntax</b>	<b>ignore-nh-metric</b> <b>no ignore-nh-metric</b>
<b>Context</b>	config>router>bgp>best-path-selection config>service>vprn config>service>vprn>bgp>best-path-selection
<b>Description</b>	This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>best-path-selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

## BGP Commands

The **no** form of the command (**no ignore-nh-metric**) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

**Default** **no ignore-nh-metric**

## ignore-router-id

**Syntax** **ignore-router-id**  
**no ignore-router-id**

**Context** config>router>bgp>best-path-selection  
config>service>vprn>bgp>best-path-selection

**Description** When the **ignore-router-id** command is present and the current best path to a destination was learned from EBGp peer X with BGP identifier x and a new path is received from EBGp peer Y with BGP identifier y the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x. The **no** form of the command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

**Default** **no ignore-router-id**

## bfd-enable

**Syntax** [**no**] **bfd-enable**

**Context** config>router>bgp  
config>router>bgp>group  
config>router>bgp>group>neighbor

**Description** This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated BGP protocol peering.

**Default** **no bfd-enable**

## cluster

**Syntax** **cluster** *cluster-id*  
**no cluster**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

Confederations can also be used to remove the full IBGP mesh requirement within an AS.

The **no** form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.

**Default** no cluster — No cluster ID is defined.

**Parameters** *cluster-id* — The route reflector cluster ID is expressed in dot decimal notation.

**Values** Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

## connect-retry

**Syntax** **connect-retry** *seconds*  
**no connect-retry**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

**Default** 120 seconds

**Parameters** *seconds* — The BGP Connect Retry timer value in seconds, expressed as a decimal integer.

**Values** 1 — 65535

## damp-peer-oscillations

<b>Syntax</b>	<b>damp-peer-oscillations</b> [ <i>idle-hold-time initial-wait second-wait max-wait</i> ] [ <b>error-interval minutes</b> ]
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command controls how long a BGP peer session remains in the idle-state after some type of error causes the session to reset. In the idle state, BGP does not initiate or respond to attempts to establish a new session. Repeated errors that occur a short while after each session reset cause longer and longer hold times in the idle state. This command supports the DampPeerOscillations FSM behavior described in section 8.1 of RFC 4271, <i>A Border Gateway Protocol 4 (BGP-4)</i>.</p> <p>The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.</p>
<b>Default</b>	<i>no damp-peer-oscillations</i>
<b>Parameters</b>	<p><i>initial-wait</i> — The amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while.</p> <p><b>Values</b> 0 — 2048</p> <p><b>Default</b> 0</p> <p><i>second-wait</i> — A period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time.</p> <p><b>Values</b> 0 — 2048</p> <p><b>Default</b> 5</p> <p><i>max-wait</i> — The maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability.</p> <p><b>Values</b> 0 — 2048</p> <p><b>Default</b> 60</p> <p><i>minutes</i> — The interval of time, in minutes after a session reset, during which the session must be error-free in order to reset the penalty counter and return to idle-hold-time to initial-wait.</p> <p><b>Values</b> 0 — 2048</p> <p><b>Default</b> 30</p>

## damping

<b>Syntax</b>	[no] <b>damping</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor

<b>Description</b>	<p>This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.</p> <p>The <b>no</b> form of the command used at the global level disables route damping.</p> <p>The <b>no</b> form of the command used at the group level reverts to the value defined at the global level.</p> <p>The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p> <p>When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:</p> <pre>Half-life:          15 minutes Max-suppress:     60 minutes Suppress-threshold:3000 Reuse-threshold  750</pre>
<b>Default</b>	no damping — Learned route damping is disabled.

## disable-4byte-asn

<b>Syntax</b>	<b>[no] disable-4byte-asn</b>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;bgp config&gt;service&gt;vprn&gt;bgp&gt;group config&gt;service&gt;vprn&gt;bgp&gt;group&gt;neighbor</pre>
<b>Description</b>	<p>This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.</p> <p>If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s).</p> <p>The <b>no</b> form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.</p>

## disable-capability-negotiation

<b>Syntax</b>	<b>[no] disable-capability-negotiation</b>
<b>Context</b>	<pre>config&gt;service&gt;vprn&gt;bgp&gt;group config&gt;service&gt;vprn&gt;bgp&gt;group&gt;neighbor</pre>
<b>Description</b>	<p>This command disables the exchange of capabilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and will strictly support IPv4 routing exchanges with that peer.</p> <p>The <b>no</b> form of the command removes this command from the configuration and restores the normal behavior.</p>
<b>Default</b>	no disable-capability-negotiation

## disable-client-reflect

<b>Syntax</b>	<b>[no] disable-client-reflect</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.  The <b>no</b> form re-enables client reflection of routes.
<b>Default</b>	no disable-client-reflect — Client routes are reflected to all client peers.

## disable-communities

<b>Syntax</b>	<b>disable-communities [standard] [extended]</b> <b>no disable-communities</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command configures BGP to disable sending communities.
<b>Parameters</b>	<b>standard</b> — Specifies standard communities that existed before VPRNs or 2547. <b>extended</b> — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

## disable-fast-external-failover

<b>Syntax</b>	<b>[no] disable-fast-external-failover</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command configures BGP fast external failover.

## eibgp-loadbalance

<b>Syntax</b>	<b>[no] eibgp-loadbalance</b>
<b>Context</b>	config>service>vprn>bgp



- Description** This command enables eiBGP load sharing so routes with both MP-BGP and IPv4 next-hops can be used simultaneously.
- In order for this command to be effective, the **ecmp** and **multipath** commands for the associated VPRN instance must also be configured to allow for multiple routes to the same destination.
- The **no** form of the command used at the global level reverts to default values.
- Default** no eibgp-loadbalance — Multipath disabled.

## enable-bgp-vpn-backup

- Syntax** **enable-bgp-vpn-backup [ipv4] [ipv6]**  
**no enable-bgp-vpn-backup**
- Context** config>service>vprn
- Description** This command enables BGP-VPN routes imported into the VPRN to have backup paths calculated for them (when they are the best path) and to be considered as backup path candidates (for other VPN-IP routes and VPRN BGP routes learned from CEs).
- Default** no enable-bgp-vpn-backup

## ebgp-link-bandwidth

- Syntax** **ebgp-link-bandwidth family [family ... (up to 2 max)]**  
**no ebgp-link-bandwidth**
- Context** config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor
- Description** When the **ebgp-link-bandwidth** command is configured, BGP automatically adds a link-bandwidth extended community to every route (of the selected types) received from directly connected (single-hop) EBGP peers within the scope of the command.
- The link-bandwidth extended community added by this command encodes the local-AS number of receiving BGP instance and the bandwidth of the interface to the directly connected EBGP peer.
- Default** no ebgp-link-bandwidth — No link bandwidth extended community is automatically added to received BGP routes.
- Parameters** *family* — The BGP address family.
- |               |             |  |
|---------------|-------------|--|
| <b>Values</b> | <i>ipv4</i> | The command applies to IPv4 and label-IPv4 routes. |
|               | <i>ipv6</i> | The command applies to IPv6 and 6PE routes.        |

## enable-peer-tracking

- Syntax** **[no] enable-peer-tracking**
- Context** config>service>vprn>bgp

## BGP Commands

```
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor
```

**Description** This command enables BGP peer tracking.

**Default** no enable-peer-tracking

## graceful-restart

**Syntax** [no] **graceful-restart**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command enables or disables graceful-restart for all VPRN BGP peers.

## enable-notification

**Syntax** **enable-notification**  
**no enable-notification**

**Context** config>service>vprn>bgp>graceful-restart  
config>service>vprn>bgp>group>graceful-restart  
config>service>vprn>bgp>group>neighbor>graceful-restart

**Description** When this command is present, the graceful restart capability sent by this router indicates support for NOTIFICATION messages. If the peer also supports this capability then the session can be restarted gracefully (while preserving forwarding) if either peer needs to send a NOTIFICATION message due to some type of event or error.

**Default** no enable-notification

## restart-time

**Syntax** **restart-time** *seconds*  
**no restart-time**

**Context** config>service>vprn>bgp>graceful-restart  
config>service>vprn>bgp>group>graceful-restart  
config>service>vprn>bgp>group>neighbor>graceful-restart

**Description** This command sets the value of the restart-time that is advertised in the router's graceful-restart capability. If this command is not configured, the default is 300.

**Default** no restart time

<b>Parameters</b>	<i>seconds</i> — The restart-time that is advertised in the router's graceful-restart capability.
<b>Values</b>	0 — 4095 seconds
<b>Default</b>	300

## stale-routes-time

<b>Syntax</b>	<b>[no] stale-routes-time</b> <i>time</i>
<b>Context</b>	config>service>vprn>bgp>graceful-restart config>service>vprn>bgp>group>graceful-restart config>service>vprn>bgp>group>neighbor>graceful-restart
<b>Description</b>	This command configures the time period to keep stale routes before the END-OF-RIB message is received from the restarting router.
<b>Parameters</b>	<i>time</i> — [1..3600 seconds]
<b>Default</b>	360 seconds

## error-handling

<b>Syntax</b>	<b>error-handling</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command specifies whether the error handling mechanism for optional transitive path attributes is enabled for this peer group.

## update-fault-tolerance

<b>Syntax</b>	<b>[no] update-fault-tolerance</b>
<b>Context</b>	config>service>vprn>bgp>error-handling config>service>vprn>bgp>group>error-handling config>service>vprn>bgp>group>neighbor>error-handling
<b>Description</b>	This command enables <b>treat-as-withdraw</b> and other similarly non-disruptive approaches for handling a wide range of UPDATE message errors, as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed.
<b>Default</b>	no fault-tolerance

## BGP Commands

### export

<b>Syntax</b>	<b>export</b> <i>policy</i> [ <i>policy...</i> ] <b>no export</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor
<b>Description</b>	<p>This command specifies the export policies to be used to control routes advertised to BGP neighbors. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.</p> <p>Note that if a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used.</p> <p>The <b>no</b> form of this command removes all route policy names from the export list.</p>
<b>Default</b>	no export — BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy.
<b>Parameters</b>	<i>policy</i> — A route policy statement name.

### family

<b>Syntax</b>	<b>family</b> [ <i>ipv4</i> ] [ <i>ipv6</i> ] [ <i>mcast-ipv4</i> ] [ <i>flow-ipv6</i> ] [ <i>flow-ipv4</i> ] <b>no family</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command specifies the address families to be negotiated with one or more multi-protocol BGP peers of the VPRN.</p> <p>The <b>no</b> form of the command removes the specified address family from the associated BGP sessions.</p>
<b>Default</b>	ipv4
<b>Parameters</b>	<i>ipv4</i> — Provisions IPv4 support. <i>ipv6</i> — Provisions IPv6 support. <b>mcast-ipv4</b> — Provisions Multicast IPv4 support. <b>[flow-ipv6]</b> — Exchanges IPv4 flowspec routes belonging to AFI 2 and SAFI 133. <b>[flow-ipv4]</b> — Specifies to use an address of variable size consisting of 1 or 2-byte NLRI(Network Layer Reachability Information) length followed by a variable length NLRI value.

## flowspec-validate

<b>Syntax</b>	<b>flowspec-validate</b> <b>no flowspec-validate</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command enables/disables validation of received flowspec routes. A flow route with a destination prefix subcomponent received from a particular peer is considered valid if and only if that peer also advertised the best unicast route to the destination prefix and any of its more-specific components. If validation is enabled and a flowspec route is not valid, it is not eligible for import into the RIB, it is not used for filtering, and it is not propagated to other flowspec peers.  The <b>no</b> form of the command disables the validation procedure.
<b>Default</b>	no flowspec-validate

## group

<b>Syntax</b>	<b>group</b> <i>name</i> [ <b>esm-dynamic-peer</b> ] <b>no group</b>
<b>Context</b>	config>service>vprn>bgp
<b>Description</b>	This command creates a context to configure a BGP peer group.  The <b>no</b> form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be shutdown before it can be deleted.
<b>Default</b>	None — No peer groups are defined.
<b>Parameters</b>	<i>name</i> — The peer group name. Allowed values is a string up to 32 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.  <b>esm-dynamic-peer</b> — This flag designates that the given BGP group will be used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any given VPRN. No BGP neighbors can be manually configured in a BGP group with this flag set.
	<b>Default</b> disabled

## neighbor

<b>Syntax</b>	[ <b>no</b> ] <b>neighbor</b> <i>ip-address</i>
<b>Context</b>	config>service>vprn>bgp>group
<b>Description</b>	This command creates a BGP peer/neighbor instance within the context of the BGP group.

## BGP Commands

This command can be issued repeatedly to create multiple peers and their associated configuration.

The **no** form of the command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

**Default** none — No neighbors are defined.

**Parameters** *ip-address* — The IP address of the BGP peer router in dotted decimal notation.

<b>Values</b>	ipv4-address :	a.b.c.d
	ipv6-address :	x:x:x:x:x:x:x[-interface]
		x:x:x:x:x:d.d.d.d[-interface]
	x:	[0..FFFF]H
	d:	[0..255]D
	interface:	32 chars maximum, mandatory for link local addresses

## family

**Syntax** **family [ipv4] [ipv6] [mcast-ipv4]**  
**no family**

**Context** config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the **family** command adds the specified address family to the list.

The **no** form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, then reset the supported address family back to the default.

**Default** ipv4

**Parameters** **ipv4** — Provisions support for IPv4 routing information.  
**ipv6** — Exchange IPv6 routing information.  
**mcast-ipv4** — Provisions Multicast IPv4 support.

## hold-time

**Syntax** **hold-time seconds [min seconds2]**  
**no hold-time**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either **keepalive** or **update**) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the router OS implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances:

1. If the specified hold-time is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
2. If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

**Default** 90 seconds

**Parameters** *seconds* — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

**Values** 0, 3 — 65535

*seconds2* — The minimum hold-time that will be accepted for the session. If the peer proposes a hold-time lower than this value the session attempt will be rejected.

## ibgp-multipath

**Syntax** **[no] ibgp-multipath**

**Context** config>service>vprn>bgp

**Description** This command defines the type of IBGP multipath to use when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple nexthops.

The **no** form of the command disables the IBGP multipath load balancing feature.

**Default** none

## import

**Syntax** **import** *policy* [*policy...*]  
**no import**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. When multiple policy

## BGP Commands

names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

The **no** form of this command removes all route policy names from the import list.

**Default** no import — BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

**Parameters** *policy* — A route policy statement name.

## keepalive

**Syntax** **keepalive** *seconds*  
**no keepalive**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The *seconds* parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the OS implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value:

If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored, and the **keepalive** is set to one third of the current **hold-time** value.

If the specified **hold-time** interval is less than the configured **keepalive** value, then the **keepalive** value is reset to one third of the specified **hold-time** interval.

If the **hold-time** interval is set to zero, then the configured value of the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

**Default** 30 seconds

**Parameters** *seconds* — The keepalive timer in seconds, expressed as a decimal integer.

**Values** 0 — 21845



## local-address

<b>Syntax</b>	<b>local-address</b> <i>ip-address</i> <b>no local-address</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>Configures the local IP address used by the group or neighbor when communicating with BGP peers. Outgoing connections use the <b>local-address</b> as the source of the TCP connection when initiating connections with a peer.</p> <p>When a local address is not specified, the SR OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.</p> <p>The <b>no</b> form of the command removes the configured local-address for BGP. The <b>no</b> form of the command used at the group level reverts to the value defined at the global level. The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	<p><b>no local-address</b> — The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.</p> <p><i>ip-address</i> — The local address expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.</p>

## local-as

<b>Syntax</b>	<b>local-as</b> <i>as-number</i> [ <b>private</b> ] [ <b>no-prepend-global-as</b> ] <b>no local-as</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command configures a BGP virtual autonomous system (AS) number.</p> <p>In addition to the global AS number configured for BGP in the config&gt;router&gt;autonomous-system context, a virtual (local) AS number can be configured to support various AS number migration scenarios. The local AS number is added to the beginning of the as-path attribute ahead of the router's AS number.</p> <p>This configuration parameter can be set at three levels: global level (applies to all EBGP peers), group level (applies to all EBGP peers in peer-group) or neighbor level (only applies to EBGP specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate local-as per EBGP session. The local-as command is not supported for IBGP sessions. When the optional <b>private</b> keyword is specified in the command the local-as number is not added to inbound routes from the EBGP peer that has <b>local-as</b> in effect.</p> <p>When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The <b>private</b> attribute can be added or removed dynamically by reissuing the command.</p>

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of the command used at the global level will remove any virtual AS number configured. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

<b>Default</b>	no local-as
<b>Parameters</b>	<i>as-number</i> — The virtual autonomous system number, expressed as a decimal integer.
	<b>Values</b> 1 — 65535
	<b>private</b> — Specifies the local-as is hidden in paths learned from the peering.
	<b>no-prepend-global-as</b> — Specifies that the global-as is hidden in paths announced to the EBGp peer.

## local-preference

<b>Syntax</b>	<b>local-preference</b> <i>local-preference</i> <b>no local-preference</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the <b>local-preference</b> integer set.</p> <p>The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.</p> <p>The <b>no</b> form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.</p> <p>The <b>no</b> form of the command used at the group level reverts to the value defined at the global level.</p> <p>The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	<b>no local-preference</b> — Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

**Parameters** *local-preference* — The local preference value to be used as the override value, expressed as a decimal integer.

**Values** 0 — 4294967295

## loop-detect

**Syntax** **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop**| **off**}  
**no loop-detect**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command configures how the BGP peer session handles loop detection in the AS path. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Note that dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

**Default** loop-detect ignore-loop

**Parameters** **drop-peer** — Sends a notification to the remote peer and drops the session.  
**discard-route** — Discards routes received with loops in the AS path.  
**ignore-loop** — Ignores routes with loops in the AS path but maintains peering.  
**off** — Disables loop detection.

## med-out

**Syntax** **med-out** {**number** | **igp-cost**}  
**no med-out**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

## BGP Commands

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default where the MED is not advertised.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

<b>Default</b>	no med-out
<b>Parameters</b>	<i>number</i> — The MED path attribute value, expressed as a decimal integer.
<b>Values</b>	0 — 4294967295
	<b>igp-cost</b> — The MED is set to the IGP cost of the given IP prefix.

## min-route-advertisement

<b>Syntax</b>	<b>min-route-advertisement</b> <i>seconds</i> <b>no min-route-advertisement</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.  This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.  The <b>no</b> form of the command reverts to default values.
<b>Default</b>	30 seconds
<b>Parameters</b>	<i>seconds</i> — The minimum route advertising interval, in seconds, expressed as a decimal integer.
<b>Values</b>	1— 255

## multihop

<b>Syntax</b>	<b>multihop</b> <i>tll-value</i> <b>no multihop</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

This parameter is meaningful only when configuring EBGp peers. It is ignored if set for an IBGP peer.

The **no** form of the command is used to convey to the BGP instance that the EBGp peers are directly connected.

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

**Default** 1 — EBGp peers are directly connected.

64 — IBGP

**Parameters** *tll-value* — The TTL value, expressed as a decimal integer.

**Values** 1 — 255

## multipath

**Syntax** **multipath** *max-paths* [**eibgp**]  
**no multipath**

**Context** config>service>vprn>bgp

**Description** This command enables BGP multipath.

When multipath is enabled BGP load shares traffic across multiple links. Multipath can be configured to load share traffic across a maximum of 16 routes. If the equal cost routes available are more than the configured value, then routes with the lowest next-hop IP address value are chosen.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Multipath is effectively disabled if the value is set to one. When multipath is disabled, and multiple equal cost routes are available, the route with the lowest next-hop IP address will be used.

The **no** form of the command used at the global level reverts to default values.

**Default** **no multipath** — Multipath disabled.

**Parameters** *integer* — The number of equal cost routes to use for multipath routing. If more equal cost routes exist than the configured value, routes with the lowest next-hop value are chosen. Setting this value to 1 disables multipath.

**Values** 1 — 16

**eibgp** — Enables EIBGP load balancing so that routes with both MP-BGP and IPv4 next-hops can be used simultaneously. Enabling this option will disable the nexthop type (MP-BGP or IPv4 and also the next-hop metric comparison).

## next-hop-resolution

<b>Syntax</b>	<b>next-hop-resolution</b>
<b>Context</b>	config>service>vprn>bgp
<b>Description</b>	This command enables the context to configure next-hop resolution parameters.

## next-hop-self

<b>Syntax</b>	<b>[no] next-hop-self</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.</p> <p>This is primarily used to avoid third-party route advertisements when connected to a multi-access network.</p> <p>The <b>no</b> form of the command used at the group level allows third-party route advertisements in a multi-access network.</p> <p>The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	<b>no next-hop-self</b> — Third-party route advertisements are allowed.

## passive

<b>Syntax</b>	<b>[no] passive</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command enables passive mode for the BGP group or neighbor.</p> <p>When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.</p> <p>The <b>no</b> form of the command used at the group level disables passive mode where BGP actively attempts to connect to its peers.</p> <p>The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	<b>no passive</b> — BGP will actively try to connect to all the configured peers.

## peer-as

<b>Syntax</b>	<b>peer-as <i>as-number</i></b>
<b>Context</b>	config>service>vprn>bgp>group

```
config>service>vprn>bgp>group>neighbor
```

**Description** This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router.

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

**Default** No AS numbers are defined.

**Parameters** *as-number* — The autonomous system number, expressed as a decimal integer.

**Values** 1 — 65535

## policy

**Syntax** **policy** *policy-name*  
**no policy**

**Context** config>service>vprn>bgp>next-hop-res

**Description** This command specifies the name of a policy statement to use with the BGP next-hop resolution process. The policy controls which IP routes in RTM are eligible to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next-hops to MPLS tunnels. If a BGP next-hop of an IPv4 or IPv6 route R is resolved in RTM and the longest matching route for the next-hop address is an IP route N that is rejected by the policy then route R is unresolved; if the route N is accepted by the policy then it becomes the resolving route for R.

The default next-hop resolution policy (when the **no policy** command is configured) is to use the longest matching active route in RTM that is not a BGP route (unless **use-bgp-routes** is configured), an aggregate route or a subscriber management route.

**Default** no policy

**Parameters** *policy-name* — The route policy name. Allowed values are any string up to 32 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. Route policies are configured in the **config>router>policy-options** context.

## peer-tracking-policy

<b>Syntax</b>	<b>peer-tracking-policy policy-name</b> <b>no peer-tracking-policy</b>
<b>Context</b>	config>router>bgp config>service>vprn>bgp
<b>Description</b>	<p>This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where this is enabled. The policy controls which IP routes in RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in RTM for a BGP neighbor address is an IP route that is rejected by the policy, or it is a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.</p> <p>The default peer-tracking policy (when the no peer-tracking-policy command is configured) is to use the longest matching active route in RTM that is not an LDP shortcut route or an aggregate route.</p>
<b>Default</b>	no peer-tracking-policy
<b>Parameters</b>	<i>policy-name</i> — The route policy name. Allowed values are any string up to 32 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. Route policies are configured in the <b>config&gt;router&gt;policy-options</b> context.

## preference

<b>Syntax</b>	<b>[no] preference preference</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group
<b>Description</b>	<p>This command configures the route preference for routes learned from the configured peer(s).</p> <p>This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.</p> <p>The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.</p> <p>The <b>no</b> form of the command used at the global level reverts to default value.</p> <p>The <b>no</b> form of the command used at the group level reverts to the value defined at the global level.</p> <p>The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	170
<b>Parameters</b>	<i>preference</i> — The route preference, expressed as a decimal integer.
<b>Values</b>	1 — 255



## prefix-limit

<b>Syntax</b>	<b>prefix-limit</b> <i>family limit</i> [ <b>log-only</b> ] [ <b>threshold percentage</b> ] [ <b>idle-timeout</b> { <i>minutes</i>   <b>forever</b> }] [ <b>post-import</b> ] <b>no prefix-limit</b> <i>family</i>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command configures the maximum number of BGP routes that can be received from a peer before some administrative action is taken. The administrative action can be the generation of a log event or taking down the session. If a session is taken down, then it can be brought back up automatically after an idle-timeout period, or else it can be configured to stay down ('forever') until the operator performs a reset.</p> <p>The <b>prefix-limit</b> command allows each address family to have its own limit; a set of address family limits can be applied to one neighbor or to all neighbors in a group.</p> <p>The <b>no</b> form of the command removes the <b>prefix-limit</b>.</p>
<b>Default</b>	No prefix limits for any address family.
<b>Parameters</b>	<p><b>log-only</b> — Enables the the warning message to be sent at the specified threshold percentage, and also when the limit is reached. However, the BGP session is not taken down.</p> <p><i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p><b>Values</b> 1 — 100</p> <p><i>family</i> — The address family to which the limit applies.</p> <p><b>Values</b> ipv4 vpn-ipv4 ipv6 vpn-ipv6 mcast-ipv4 l2-vpn mvpn-ipv4 mdt-safi ms-pw flow-ipv4 route-target mcast-vpn-ipv4 mvpn-ipv6 flow-ipv6 evpn mcast-ipv6</p> <p><i>limit</i> — The number of routes that can be learned from a peer expressed as a decimal integer.</p> <p><b>Values</b> 1 — 4294967295</p> <p><i>minutes</i> — Specifies duration in minutes before automatically re-establishing a session.</p> <p><b>Values</b> 1 — 1024</p> <p><b>forever</b> — Specifies that the session is reestablished only after <b>clear router bgp</b> command is executed.</p> <p><b>post-import</b> — Specifies that the limit should be applied only to the number of routes that are accepted by import policies.</p>

## rapid-withdrawal

<b>Syntax</b>	[no] <b>rapid-withdrawal</b>
<b>Context</b>	config>service>vprn>bgp

## BGP Commands

<b>Description</b>	This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates.  The <b>no</b> form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.
<b>Default</b>	no rapid-withdrawal

## remove-private

<b>Syntax</b>	<b>remove-private</b> [ <b>limited</b> ] [ <b>skip-peer-as</b> ] <b>no remove-private</b>
<b>Context</b>	config>service>vprn>bgp config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.  When the <b>remove-private</b> parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.  The OS software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.  The <b>no</b> form of the command used at the global level reverts to default value. The <b>no</b> form of the command used at the group level reverts to the value defined at the global level. The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.
<b>Default</b>	<b>no remove-private</b> — Private AS numbers will be included in the AS path attribute.
<b>Parameters</b>	<b>limited</b> — This optional keyword removes private ASNs up to the first public ASN encountered. It then stops removing private ASNs.  <b>skip-peer-as</b> — This optional keyword causes this command to not remove a private ASN from the AS-Path if that ASN is the same as the BGP peer AS number.

## leak-import

<b>Syntax</b>	<b>leak-import</b> <i>policy-name</i> [ <i>policy-name .. (up to 15 max)</i> ] <b>no leak-import</b>
<b>Context</b>	config>service>vprn>bgp>rib-management>ipv4 config>service>vprn>bgp>rib-management>ipv6
<b>Description</b>	This command specifies a chain of policy-statements (up to 15) that are evaluated against all the leakable IPv4 (or IPv6) BGP routes in all routing instances (VPRN or base). If a leak-import policy of routing instance T (the target routing instance) accepts a BGP route originally learned and marked as

leakable by another routing instance S (the source routing instance), and the BGP route is valid in S (a best or non-best path) then a copy of the route is added to the BGP RIB-IN of instance T.

**Default** No default

**Parameters** *policy-name* — The name of a route policy statement (up to 32 characters)

## split-horizon

**Syntax** **split-horizon**  
**no split-horizon**

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** This command enables the use of split-horizon. When applied globally, to a group, or a specific peer, split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer; confed-EBGP, EBGP and IBGP.

The configuration default is **no split-horizon**, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.

**NOTE:** Use of the **split-horizon** command may have a detrimental impact on peer and route scaling and therefore operators are encouraged to use it only when absolutely needed.

The **no** form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.

**Default** **no split-horizon**

## third-party-nexthop

**Syntax** third-party-nexthop  
no third-party-nexthop

**Context** config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor

**Description** Use this command to enable the router to send third-party next-hop to EBGP peers in the same subnet as the source peer, as described in RFC 4271. If enabled when an IPv4 or IPv6 route is received from one EBGP peer and advertised to another EBGP peer in the same IP subnet, the BGP next-hop is left unchanged. Third-party next-hop is not done if the address family of the transport does not match the address family of the route.

The **no** form of the command prevents BGP from performing any third party next-hop processing toward any single-hop EBGP peers within the scope of the command. No third-party next-hop means the next-hop will always carry the IP address of the interface used to establish the TCP connection to the peer.

**Default** no third-party-nexthop

## type

<b>Syntax</b>	<b>[no] type {internal   external}</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command designates the BGP peer as type internal or external.</p> <p>The type of <b>internal</b> indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.</p> <p>By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered <b>internal</b>. If the local AS is different, then the peer is considered <b>external</b>.</p> <p>The <b>no</b> form of the command used at the group level reverts to the default value. The <b>no</b> form of the command used at the neighbor level reverts to the value defined at the group level.</p>
<b>Default</b>	<b>no type</b> — Type of neighbor is derived on the local AS specified.
<b>Parameters</b>	<b>internal</b> — Configures the peer as internal. <b>external</b> — Configures the peer as external.

## updated-error-handling

<b>Syntax</b>	<b>[no] updated-error-handling</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	<p>This command controls whether SROS utilizes the new neighbor-complete bit when processing optional transitive path attributes and advertising them to the associated BGP neighbor.</p> <p>This command also control if SROS utilizes the error handling mechanism for optional-transitive path attributes.</p>
<b>Default</b>	no updated-error-handling

## ttl-security

<b>Syntax</b>	<b>ttl-security min-ttl-value</b> <b>no ttl-security</b>
<b>Context</b>	config>service>vprn>bgp>group config>service>vprn>bgp>group>neighbor
<b>Description</b>	Configure TTL security parameters for incoming packets.
<b>Parameters</b>	<i>min-ttl-value</i> — Specify the minimum TTL value for an incoming BGP packet.

**Values** 1 — 255

**Default** 1

---

## ISIS Commands

### isis

<b>Syntax</b>	<b>[no] isis</b> <i>isis-instance</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the ISIS protocol instance with the VPRN service. The <b>no</b> form of the command disables the ISIS protocol instance from the given VPRN service.
<b>Default</b>	no ISIS

### advertise-passive-only

<b>Syntax</b>	<b>[no] advertise-passive-only</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command enables and disables IS-IS for the VPRN instance to advertise only prefixes that belong to passive interfaces.

### advertise-router-capability

<b>Syntax</b>	<b>advertise-router-capability</b> { <i>area</i>   <i>as</i> } <b>no advertise-router-capability</b>
<b>Context</b>	config>service>vprn>isis config>service>vprn>isis>level
<b>Description</b>	This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A new TLV as defined in RFC 4971 advertises the TE Node Capability Descriptor capability.  The parameters ( <i>area</i> & <i>as</i> ) control the scope of the capabilities advertisements.  The no form of this command, disables this capability.
<b>Default</b>	<b>no advertise-router-capability</b>
<b>Parameters</b>	<b>area</b> — Capabilities are only advertised within the area of origin. <b>as</b> — Capabilities are only advertised throughout the entire autonomous system

## all-l1isis

<b>Syntax</b>	<b>all-l1isis</b> <i>ieee-address</i> <b>no all-l1isis</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command specifies the MAC address to use for the VPRN instance of the L1 IS-IS routers. The MAC address should be a multicast address. You should shut/no shut the IS-IS instance to make the change operational.
<b>Default</b>	all-l1isis 01-80-C2-00-01-00
<b>Parameters</b>	<i>ieee-address</i> — Specifies the destination MAC address for all L1 I-IS neighbors on the link for this ISIS instance.

## all-l2isis

<b>Syntax</b>	<b>all-l2isis</b> <i>ieee-address</i> <b>no all-l2isis</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command specifies the MAC address to use for L2 IS-IS routers for the VPRN instance. The MAC address should be a multicast address. You should shut/no shut the IS-IS instance to make the change operational.
<b>Default</b>	all-l2isis 01-80-C2-00-02-11
<b>Parameters</b>	<i>ieee-address</i> — Specifies the destination MAC address for all L2 ISIS neighbors on the link for this ISIS instance.

## area-id

<b>Syntax</b>	<b>[no] area-id</b> <i>area-address</i>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command configures the area ID portion of NSAP addresses for the VPRN instance. This identifies a point of connection to the network, such as a router interface, and is called a Network Service Access Point (NSAP). Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses.</p> <p>A maximum of 3 area addresses can be configured for the VPRN instance.</p> <p>NSAP addresses are divided into three parts. Only the area ID portion is configurable.</p> <ul style="list-style-type: none"> <li>• Area ID — A variable length field between 1 and 13 bytes long. This includes the Authority and Format Identifier (AFI) as the most significant byte and the area ID.</li> <li>• System ID — A six-byte system identification. This value is not configurable. The system ID is derived from the system or router ID.</li> </ul>

- Selector ID — A one-byte selector identification that must contain zeros when configuring a NET. This value is not configurable. The selector ID is always 00.

The NET is constructed like an NSAP but the selector byte contains a 00 value. NET addresses are exchanged in hello and LSP PDUs. All net addresses configured on the node are advertised to its neighbors.

For Level 1 interfaces, neighbors can have different area IDs, but, they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For Level 2 (only) interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only Level 2 neighbors and Level 2 LSPs are exchanged.

For Level 1 and Level 2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging Level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the first area address.

The **no** form of the command removes the area address.

## auth-keychain

<b>Syntax</b>	<b>auth-keychain</b> <i>name</i>
<b>Context</b>	config>service>vprn>isis< config>service>vprn>isis>level
<b>Description</b>	This command configures an authentication keychain to use for the protocol interface for the VPRN instance. The keychain allows the rollover of authentication keys during the lifetime of a session.
<b>Default</b>	no auth-keychain
<b>Parameters</b>	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

## authentication-check

<b>Syntax</b>	<b>[no] authentication-check</b>
<b>Context</b>	confi>service>vprn>isis
<b>Description</b>	This command sets an authentication check to reject PDUs that do not match the type or key requirements for the VPRN instance.  The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.  When <b>no authentication-check</b> is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, mismatches cause an event to be generated and will not be rejected.



The **no** form of this command allows authentication mismatches to be accepted and generate a log event.

**Default** **authentication-check** — Rejects authentication mismatches.

## authentication-key

**Syntax** **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]  
**no authentication-key**

**Context** config>service>vprn>isis  
config>service>vprn>isis>level

**Description** This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface for the VPRN instance.

Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication *key* and the authentication *type* on a segment must match. The **OSPF Commands** statement must also be included.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated including the hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, hello PDUs are not authenticated.

The **no** form of the command removes the authentication key.

**Default** **no authentication-key** — No authentication key is configured.

**Parameters** *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

*hash-key* — The hash key. The key can be any combination of ASCII characters up to 342 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.

## authentication-type

<b>Syntax</b>	<b>authentication-type</b> { <b>password</b>   <b>message-digest</b> } <b>no authentication</b>
<b>Context</b>	config>service>vprn>isis config>service>vprn>isis>level
<b>Description</b>	<p>This command enables either simple password or message digest authentication or must go in either the global IS-IS or IS-IS level context.</p> <p>Both the authentication key and the authentication type on a segment must match. The <b>authentication-key</b> statement must also be included.</p> <p>Configure the authentication type on the global level in the <b>config&gt;router&gt;isis</b> context.</p> <p>Configure or override the global setting by configuring the authentication type in the <b>config&gt;router&gt;isis&gt;level</b> context.</p> <p>The <b>no</b> form of the command disables authentication.</p>
<b>Default</b>	<b>no authentication-type</b> — No authentication type is configured and authentication is disabled.
<b>Parameters</b>	<p><b>password</b> — Specifies that simple password (plain text) authentication is required.</p> <p><b>message-digest</b> — Specifies that MD5 authentication in accordance with RFC2104 is required.</p>

## csnp-authentication

<b>Syntax</b>	<b>[no] csnp-authentication</b>
<b>Context</b>	config>service>vprn>isis config>service>vprn>isis>level
<b>Description</b>	This command enables authentication of individual ISIS packets of complete sequence number PDUs (CSNP) type for the VPRN instance.

## default-route-tag

<b>Syntax</b>	<b>default-route-tag</b> <i>tag</i> <b>no default-route-tag</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command configures the route tag for default route for the router or VPRN service.
<b>Parameters</b>	<p><i>tag</i> — Assigns a default tag.</p> <p><b>Values</b>      1 — 4294967295</p>

## export

<b>Syntax</b>	<b>[no] export</b> <i>policy-name</i> [ <i>policy-name</i> ...up to 5 max]
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command configures export routing policies that determine the routes exported from the routing table to IS-IS.</p> <p>If no export policy is defined, non IS-IS routes are not exported from the routing table manager to IS-IS.</p> <p>If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.</p> <p>If an <b>aggregate</b> command is also configured in the <b>config&gt;router</b> context, then the aggregation is applied before the export policy is applied.</p> <p>Routing policies are created in the <b>config&gt;router&gt;policy-options</b> context.</p> <p>The <b>no</b> form of the command removes the specified <i>policy-name</i> or all policies from the configuration if no <i>policy-name</i> is specified.</p>
<b>Default</b>	<b>no export</b> — No export policy name is specified.
<b>Parameters</b>	<i>policy-name</i> — The export policy name. Up to five <i>policy-name</i> arguments can be specified.

## export-limit

<b>Syntax</b>	<b>export-limit</b> <i>number</i> [ <b>log</b> <i>percentage</i> ] <b>no export-limit</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command configures the maximum number of routes (prefixes) that can be exported into IS-IS from the route table for the VPRN instance.</p> <p>The <b>no</b> form of the command removes the parameters from the configuration.</p>
<b>Default</b>	no export-limit, the export limit for routes or prefixes is disabled..
<b>Parameters</b>	<p><i>number</i> — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.</p> <p><b>Values</b> 1 — 4294967295</p> <p><b>log percentage</b> — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.</p> <p><b>Values</b> 1 — 100</p>

## graceful-restart

<b>Syntax</b>	<b>[no] graceful-restart</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command enables graceful-restart helper support for ISIS for the VPRN instance. The router will act as a helper to neighbors who are graceful-restart-capable and are restarting.</p> <p>When the control plane of a graceful-restart-capable router fails, the neighboring routers (graceful-restart helpers) temporarily preserve adjacency information so packets continue to be forwarded through the failed graceful-restart router using the last known routes. If the control plane of the graceful-restart router comes back up within the timer limits, then the routing protocols re-converge to minimize service interruption.</p> <p>The <b>no</b> form of the command disables graceful restart and removes all graceful restart configurations in the ISIS instance.</p>
<b>Default</b>	<b>disabled</b>

## helper-disable

<b>Syntax</b>	<b>[no] helper-disable</b>
<b>Context</b>	config>service>vprn>isis>graceful-restart
<b>Description</b>	<p>This command disables the helper support for graceful restart for the VPRN instance.</p> <p>When <b>graceful-restart</b> is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The router supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the router will not help the neighbors to restart).</p> <p>The <b>no helper-disable</b> command enables helper support and is the default when graceful-restart is enabled.</p>
<b>Default</b>	disabled

## hello-authentication

<b>Syntax</b>	<b>[no] hello-authentication</b>
<b>Context</b>	config>service>vprn>isis config>service>vprn>isis>interface config>service>vprn>isis>level
<b>Description</b>	<p>This command enables authentication of individual ISIS packets of HELLO type for the VPRN instance.</p> <p>The <b>no</b> form of the command suppresses authentication of HELLO packets.</p>

## hello-padding

<b>Syntax</b>	<b>[no] hello-padding {adaptive loose strict}</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command enables the padding of IS-IS hello messages for the VPRN instance, such that the message will be padded to a length of either 1492 or up to the maximum MTU value.</p> <p>The <b>no</b> form of the command disables IS-IS hello padding.</p>
<b>Default</b>	<b>no hello-padding</b> — hello padding is not configured
<b>Parameters</b>	<p><b>adaptive</b> — Specifies the adaptive padding option; this option is able to detect MTU asymmetry from one side of the connection but uses more overhead than loose padding.</p> <ol style="list-style-type: none"> <li>1. point-to-point interface—Hello PDUs are padded until the sender declares an adjacency on the link to be in state up. If the implementation supports RFC 3373/5303, "Three-Way Handshake for IS-IS Point-to-Point Adjacencies" then this is when the three-way state is up. If the implementation use the "classic" algorithm described in ISO 10589, this is when adjacency state is up. If the neighbor does not support the adjacency state TLV, then padding continues.</li> <li>2. broadcast interface—Padding starts until at least one adjacency is up on the interface.</li> </ol> <p><b>loose</b> — Specifies the loose padding option; the loose padding may not be able to detect certain situations such as asymmetrical MTUs between the routing devices.</p> <ol style="list-style-type: none"> <li>1. point-to-point interface—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the INIT state.</li> <li>2. broadcast interface—Padding starts until at least one adjacency (bcast only has up/down) is up on the interface.</li> </ol> <p><b>strict</b> — Specifies the strict padding option.</p> <ol style="list-style-type: none"> <li>1. point-to-point interface—Padding is done for all adjacency states, and is continuous. Strict padding has the most overhead but detects MTU issues on both sides of a link.</li> <li>2. broadcast interface—Padding is done for all adjacency states, and is continuous. Strict padding has the most overhead but detects MTU issues on both sides of a link.</li> </ol>

## ignore-lsp-errors

<b>Syntax</b>	<b>[no] ignore-lsp-errors</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command specifies that for this VPRN instance, ISIS will ignore LSP packets with errors. When enabled, IS-IS LSP errors will be ignored and the associated record will not be purged.</p> <p>This command enables ISIS to ignore the ATT bit and therefore suppress the installation of default routes.</p> <p>The <b>no</b> form of the command specifies that ISIS will not ignore LSP errors.</p>

## iid-tlv-enable

<b>Syntax</b>	<b>[no] iid-tlv-enable</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command enables or disables the Instance Identifier (IID) TLV for the ISIS instance. When enabled, each I-IS instance marks its packets with the IID TLV containing its unique 16-bit IID for the routing domain. You should shut/no shut the isis instance to make the change operational.
<b>Default</b>	no iid-tlv-enable

## interface

<b>Syntax</b>	<b>[no] interface ip-int-name</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command creates the context to configure an IS-IS interface. When an area is defined, the interfaces belong to that area. Interfaces cannot belong to separate areas. When the interface is a POS channel, the OSINCP is enabled when the interface is created and removed when the interface is deleted. The <b>no</b> form of the command removes IS-IS from the interface. The <b>shutdown</b> command in the <b>config&gt;router&gt;isis&gt;interface</b> context administratively disables IS-IS on the interface without affecting the IS-IS configuration.
<b>Default</b>	<b>no interface</b> — No IS-IS interfaces are defined.
<b>Parameters</b>	<i>ip-int-name</i> — Identify the IP interface name created in the <b>config&gt;router&gt;interface</b> context. The IP interface name must already exist.

## bfd-enable

<b>Syntax</b>	<b>[no] bfd-enable {ipv4   ipv6} [include-bfd-tlv]</b>
<b>Context</b>	config>service>vprn>interface
<b>Description</b>	This command enables the use of bi-directional forwarding (BFD) to control IPv4 adjacencies. By enabling BFD on an IPv4 or IPv6 protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set by the BFD command under the IP interface. This command must be given separately to enable/disable BFD for both IPv4 and IPv6. The <b>no</b> form of this command removes BFD from the associated adjacency.
<b>Default</b>	no bfd-enable ipv4

## csnp-interval

<b>Syntax</b>	<b>csnp-interval</b> <i>seconds</i> <b>no csnp-interval</b>
<b>Context</b>	config>service>vpern>isis>interface
<b>Description</b>	This command configures the time interval, in seconds, to send complete sequence number (CSN) PDUs from the interface. IS-IS must send CSN PDUs periodically. The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	<b>csnp-interval 10</b> — CSN PDUs are sent every 10 seconds for LAN interfaces. <b>csnp-interval 5</b> — CSN PDUs are sent every 5 seconds for point-to-point interfaces.
<b>Parameters</b>	<i>seconds</i> — The time interval, in seconds between successive CSN PDUs sent from this interface expressed as a decimal integer. 1 — 65535

## hello-auth-keychain

<b>Syntax</b>	<b>hello-auth-keychain</b> <i>name</i>
<b>Context</b>	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
<b>Description</b>	This command configures an authentication keychain to use for the protocol interface. The keychain allows the rollover of authentication keys during the lifetime of a session.
<b>Default</b>	no hello-auth-keychain
<b>Parameters</b>	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

## hello-authentication-key

<b>Syntax</b>	<b>hello-authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ] <b>no hello-authentication-key</b>
<b>Context</b>	config>service>vprn>isis>interface config>service>vprn>isis>interface> level
<b>Description</b>	This command configures the authentication key (password) for hello PDUs. Neighboring routers use the password to verify the authenticity of hello PDUs sent from this interface. Both the hello authentication key and the hello authentication type on a segment must match. The <b>hello-authentication-type</b> must be specified. To configure the hello authentication key in the interface context use the <b>hello-authentication-key</b> in the <b>config&gt;router&gt;isis&gt;interface</b> context.

To configure or override the hello authentication key for a specific level, configure the **hello-authentication-key** in the **config>router>isis>interface>level** context.

If both IS-IS and hello-authentication are configured, hello messages are validated using hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS (including hello) protocol PDUs.

When the hello authentication key is configured in the **config>router>isis>interface** context, it applies to all levels configured for the interface.

The **no** form of the command removes the authentication-key from the configuration.

- Default**     **no hello-authentication-key** — No hello authentication key is configured.
- Parameters**
  - authentication-key* — The hello authentication key (password). The key can be any combination of ASCII characters up to 254 characters in length (un-encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).
  - hash-key* — The hash key. The key can be any combination of ASCII characters up to 342 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.

## hello-authentication-type

- Syntax**     **hello-authentication-type {password | message-digest}**  
**no hello-authentication-type**
- Context**     config>service>vprn>isis>interface  
config>service>vprn>isis>interface>level
- Description**

This command enables hello authentication at either the interface or level context. Both the hello authentication key and the hello authentication type on a segment must match. The hello **authentication-key** statement must also be included.

To configure the hello authentication type at the interface context, use **hello-authentication-type** in the **config>router>isis>interface** context.

To configure or override the hello authentication setting for a given level, configure the **hello-authentication-type** in the **config>router>isis>interface>level** context.

The **no** form of the command disables hello authentication.
- Default**     **no hello-authentication-type** — Hello authentication is disabled.
- Parameters**
  - password** — Specifies simple password (plain text) authentication is required.



**message-digest** — Specifies MD5 authentication in accordance with RFC2104 (HMAC: Keyed-Hashing for Message Authentication) is required.

## interface-type

<b>Syntax</b>	<b>interface-type {broadcast   point-to-point}</b> <b>no interface-type</b>
<b>Context</b>	config>service>vprn>isis>interface
<b>Description</b>	<p>This command configures the IS-IS interface type as either broadcast or point-to-point.</p> <p>Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the designated IS-IS overhead if the link is used as a point-to-point.</p> <p>If the interface type is not known at the time the interface is added to IS-IS and subsequently the IP interface is bound (or moved) to a different interface type, then this command must be entered manually.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Special Cases</b>	<p><b>SONET</b> — Interfaces on SONET channels default to the point-to-point type.</p> <p><b>Ethernet or Unknown</b> — Physical interfaces that are Ethernet or unknown default to the broadcast type.</p>
<b>Default</b>	<p><b>point-to-point</b> — For IP interfaces on SONET channels.</p> <p><b>broadcast</b> — For IP interfaces on Ethernet or unknown type physical interfaces.</p>
<b>Parameters</b>	<p><b>broadcast</b> — Configures the interface to maintain this link as a broadcast network.</p> <p><b>point-to-point</b> — Configures the interface to maintain this link as a point-to-point link.</p>

## ipv4-multicast-disable

<b>Syntax</b>	<b>[no] ipv4-multicast-disable</b>
<b>Context</b>	config>service>vprn>isis>interface
<b>Description</b>	This command administratively disables/enables ISIS operation for IPv4.
<b>Default</b>	no ipv4-multicast-disable

## ipv6-unicast-disable

<b>Syntax</b>	<b>[no] ipv6-unicast-disable</b>
<b>Context</b>	config>router>isis>if config>service>vprn>isis>interface
<b>Description</b>	This command disables IS-IS IPv6 unicast routing for the interface.

By default IPv6 unicast on all interfaces is enabled. However, IPv6 unicast routing on IS-IS is in effect when the **config>router>isis>ipv6-routing mt** command is configured.

The **no** form of the command enables IS-IS IPv6 unicast routing for the interface.

## hello-interval

<b>Syntax</b>	<b>hello-interval</b> <i>seconds</i> <b>no hello-interval</b>
<b>Context</b>	config>router>isis>if>level <i>level-number</i> config>service>vprn>isis>interface>level
<b>Description</b>	This command configures the interval in seconds between hello messages issued on this interface at this level.  The <b>no</b> form of the command to reverts to the default value.
<b>Default</b>	<b>3</b> — Hello interval default for the designated intersystem. <b>9</b> — Hello interval default for non-designated intersystems.
<b>Parameters</b>	<i>seconds</i> — The hello interval in seconds expressed as a decimal integer.  <b>Values</b> 1 — 20000

## hello-multiplier

<b>Syntax</b>	<b>hello-multiplier</b> <i>multiplier</i> <b>no hello-multiplier</b>
<b>Context</b>	config>router>isis>if>level <i>level-number</i> config>service>vprn>isis>interface>level
<b>Description</b>	This command configures the number of missing hello PDUs from a neighbor after the router declares the adjacency down.  The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	<b>3</b> — The router can miss up to 3 hello messages before declaring the adjacency down.
<b>Parameters</b>	<i>multiplier</i> — The multiplier for the hello interval expressed as a decimal integer.  <b>Values</b> 2 — 100

## ipv4-multicast-metric

<b>Syntax</b>	<b>ipv4-multicast-metric</b> <i>metric</i> <b>no ipv4-multicast-metric</b>
<b>Context</b>	config>service>vprn>isis>interface>level
<b>Description</b>	This command configures IS-IS interface metric for IPv4 multicast for the VPRN instance. The <b>no</b> form of this command removes the metric from the configuration.
<b>Parameters</b>	<i>metric</i> — Specifies the IS-IS interface metric for IPv4 multicast.
<b>Values</b>	1 — 16777215

## ipv6-unicast-metric

<b>Syntax</b>	<b>ipv6-unicast-metric</b> <i>metric</i> <b>no ipv6-unicast-metric</b>
<b>Context</b>	config>service>vprn>isis>interface>level
<b>Description</b>	This command configures IS-IS interface metric for IPv6 unicast. The <b>no</b> form of this command removes the metric from the configuration.
<b>Parameters</b>	<i>metric</i> — Specifies the IS-IS interface metric for IPv6 unicast.
<b>Values</b>	1 — 16777215

## metric

<b>Syntax</b>	<b>metric</b> <i>metric</i> <b>no metric</b>
<b>Context</b>	config>service>vprn>isis>interface>level
<b>Description</b>	This command configures the metric used for the level on the interface. In order to calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different. If the metric is not configured, the default of 10 is used unless reference bandwidth is configured. The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	<b>10</b> — A metric of 10 for the level on the interface is used.
<b>Parameters</b>	<i>metric</i> — The metric assigned for this level on this interface.
<b>Values</b>	1 — 16777215

passive

<b>Syntax</b>	<b>[no] passive</b>
<b>Context</b>	config>service>vprn>isis>interface config>service>vprn>isis>interface>level
<b>Description</b>	<p>This command adds the passive attribute which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.</p> <p>When the passive mode is enabled, the interface or the interface at the level ignores ingress IS-IS protocol PDUs and will not transmit IS-IS protocol PDUs.</p> <p>The <b>no</b> form of the command removes the passive attribute.</p>
<b>Special Cases</b>	<p><b>Service Interfaces</b> — Service interfaces (defined using the service-prefix command in <b>config&gt;router</b>) are passive by default.</p> <p><b>All other Interfaces</b> — All other interfaces are not passive by default.</p>
<b>Default</b>	<p><b>passive</b> — Service interfaces are passive.</p> <p><b>no passive</b> — All other interfaces are not passive.</p>

priority

<b>Syntax</b>	<b>priority <i>number</i></b> <b>no priority</b>
<b>Context</b>	config>service>vprn>isis>interface>level
<b>Description</b>	<p>This command configures the priority of the IS-IS router interface for designated router election on a multi-access network.</p> <p>This priority is included in hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority is the preferred designated router. The designated router is responsible for sending LSPs with regard to this network and the routers that are attached to it.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	<b>64</b>
<b>Parameters</b>	<p><i>number</i> — The priority for this interface at this level.</p> <p><b>Values</b>      0 — 127</p>

## sd-offset

<b>Syntax</b>	<b>sd-offset</b> <i>offset-value</i> <b>no sd-offset</b>
<b>Context</b>	config>service>vprn>isis>interface>level
<b>Description</b>	If the pre-FEC error rate of the associated DWDM port crosses the configured <b>sd-threshold</b> , this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the <b>sd-threshold</b> value is configured under that port.  The <b>no</b> form of the command reverts the offset value to 0.
<b>Default</b>	no sd-offset
<b>Parameters</b>	<i>offset-value</i> — Specifies the amount the interface metric is increased by if the <b>sd-threshold</b> is crossed.  <b>Values</b> 0 — 16777215

## sf-offset

<b>Syntax</b>	<b>sf-offset</b> <i>offset-value</i> <b>no sf-offset</b>
<b>Context</b>	config>service>vprn>isis>interface>level
<b>Description</b>	If the pre-FEC error rate of the associated DWDM port crosses the configured <b>sf-threshold</b> , this offset-value is added to the IS-IS interface metric. This parameter is only effective if the interface is associated with a DWDM port and the <b>sf-threshold</b> value is configured under that port.  The <b>no</b> form of the command reverts the offset value to 0.
<b>Default</b>	no sf-offset
<b>Parameters</b>	<i>offset-value</i> — Specifies the amount the interface metric is increased by if the <b>sf-threshold</b> is crossed.  <b>Values</b> 0 — 16777215

## lfa-policy-map

<b>Syntax</b>	<b>lfa-policy-map route-nh-template</b> <i>template-name</i> <b>no lfa-policy-map</b>
<b>Context</b>	config>service>vprn>isis>interface
<b>Description</b>	This command applies a route next-hop policy template to the IS-IS interface for the VPRN instance.  When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the

interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

**Parameters** *template-name* — Specifies the name of the template, up to 32 characters.

## loopfree-alternate-exclude

**Syntax** **[no] loopfree-alternate**

**Context** config>service>vprn>isis>interface  
config>service>vprn>isis>level

**Description** This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.

When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

The **no** form of this command re-instates the default value for this command.

**Default** no loopfree-alternate-exclude

## lsp-pacing-interval

**Syntax** **lsp-pacing-interval** *milliseconds*  
**no lsp-pacing-interval**

**Context** config>service>vprn>isis>interface

**Description** This command configures the interval between LSP PDUs sent from this interface. To avoid bombarding adjacent neighbors with excessive data, pace the Link State Protocol Data Units (LSP's). If a value of zero is configured, no LSP's are sent from the interface.

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

**Default** **100** — LSPs are sent in 100 millisecond intervals.

**Parameters** *milliseconds* — The interval in milliseconds that IS-IS LSP's can be sent from the interface expressed as a decimal integer.

**Values** 0 — 65535

## mesh-group

**Syntax** **mesh-group** {**value** | **blocked**}  
**no mesh-group**

**Context** config>service>vprn>isis>interface

**Description** This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than a copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified. Configure mesh groups carefully. It is easy to create isolated islands that do not receive updates as (other) links fail.

The **no** form of the command removes the interface from the mesh group.

**Default** **no mesh-group** — The interface does not belong to a mesh group.

**Parameters** **value** — The unique decimal integer value distinguishes this mesh group from other mesh groups on this or any other router that is part of this mesh group.

**Values** 1 — 2000000000

**blocked** — Prevents an interface from flooding LSPs.

## retransmit-interval

**Syntax** **retransmit-interval** *seconds*  
**no retransmit-interval**

**Context** config>service>vprn>isis>interface

**Description** This command configures the minimum time between LSP PDU retransmissions on a point-to-point interface.

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

**Default** 100

**Parameters** *seconds* — The interval in seconds that IS-IS LSPs can be sent on the interface.

1 — 65535

## tag

<b>Syntax</b>	<b>tag</b> <i>tag</i> <b>no tag</b>
<b>Context</b>	config>service>vprn>isis>interface
<b>Description</b>	This command configures a route tag to the specified IP address of an interface.
<b>Parameters</b>	<i>tag</i> — [1 — 4294967295]

## ipv4-multicast-routing

<b>Syntax</b>	<b>ipv4-multicast-routing</b> { <b>native</b>   <b>mt</b> } <b>[no] ipv4-multicast-routing</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	The multicast RTM is used for Reverse Path Forwarding checks. This command controls which IS-IS topology is used to populate the IPv4 multicast RTM.  The <b>no</b> ipv4-multicast-routing form of the command results in none of the IS-IS routes being populated in the IPv4 multicast RTM and would be used if multicast is configured to use the unicast RTM for the RPF check.
<b>Default</b>	ipv4-multicast-routing native
<b>Parameters</b>	<b>native</b> — Causes IPv4 routes from the MT0 topology to be added to the multicast RTM for RPF checks.  <b>mt</b> — Causes IPv4 routes from the MT3 topology to be added to the multicast RTM for RPF checks.

## ipv4-routing

<b>Syntax</b>	<b>[no] ipv4-routing</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command specifies whether this IS-IS instance supports IPv4. The <b>no</b> form of the command disables IPv4 on the IS-IS instance.
<b>Default</b>	ipv4-routing

## ipv6-routing

<b>Syntax</b>	<b>[no] ipv6-routing</b> { <b>native</b>   <b>mt</b> }
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command enables IPv6 routing.



The **no** form of the command disables support for IS-IS IPv6 TLVs for IPv6 routing.

**Default** disabled

**Parameters** **native** — Enables IS-IS IPv6 TLVs for IPv6 routing and enables support for native IPv6 TLVs.  
**mt** — Enables IS-IS multi-topology TLVs for IPv6 routing. When this parameter is specified, the support for native IPv6 TLVs is disabled.

## level

**Syntax** **level** *level-number*

**Context** config>service>vprn>isis>  
 config>service>vprn>isis>interface  
 config>service>vprn>isis>link-group

**Description** This command creates the context to configure IS-IS Level 1 or Level 2 area attributes.

A router can be configured as a Level 1, Level 2, or Level 1-2 system. A Level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A Level 2 adjacency cannot be established over this interface.

Level 1/2 adjacency is created if the neighbor is also configured as Level 1/2 router and has at least one area address in common. A Level 2 adjacency is established if there are no common area IDs.

A Level 2 adjacency is established if another router is configured as Level 2 or a Level 1/2 router with interfaces configured as Level 1/2 or Level 2. Level 1 adjacencies will not be established over this interface.

To reset global and/or interface level parameters to the default, the following commands must be entered independently:

```
level> no hello-authentication-key
level> no hello-authentication-type
level> no hello-interval
level> no hello-multiplier
level> no metric
level> no passive
level> no priority
```

**Special Cases** **Global IS-IS Level** — The **config>router>isis** context configures default global parameters for both Level 1 and Level 2 interfaces.

**IS-IS Interface Level** — The **config>router>isis>interface** context configures IS-IS operational characteristics of the interface at Level 1 and/or Level 2. A logical interface can be configured on one Level 1 and one Level 2. In this case, each level can be configured independently and parameters must be removed independently.

By default an interface operates in both Level 1 and Level 2 modes.

**Default** level 1 or level 2

**Parameters** *level-number* — The IS-IS level number.

**Values** 1, 2

## default-ipv4-multicast-metric

<b>Syntax</b>	<b>default-ipv4-multicast-metric</b> <i>metric</i> <b>no default-ipv4-multicast-metric</b>
<b>Context</b>	config>service>vprn>isis>level
<b>Description</b>	This command configures the default metric to be used for the IS-IS interface in the IPv4 multicast topology (MT3).  The <b>no</b> form of this command deletes the specified default metric and reverts to using the system default of 10.
<b>Default</b>	10
<b>Parameters</b>	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT3)
	<b>Values</b> 1 — 16777215

## default-ipv6-multicast-metric

<b>Syntax</b>	<b>default-ipv6-multicast-metric</b> <i>metric</i> <b>no default-ipv6-multicast-metric</b>
<b>Context</b>	config>service>vprn>isis>level
<b>Description</b>	This command configures the default metric to be used for the IS-IS interface in the IPv6 multicast topology (MT4).  The <b>no</b> form of this command deletes the specified default metric and reverts to using the system default of 10.
<b>Default</b>	10
<b>Parameters</b>	<i>metric</i> — Specifies the default metric for interfaces in the IPv4 multicast topology (MT4).
	1 — 16777215

## default-ipv6-unicast-metric

<b>Syntax</b>	<b>default-ipv6-unicast-metric</b> <i>ipv6 metric</i> <b>no default-ipv6-unicast-metric</b>
<b>Context</b>	config>service>vprn>isis>level
<b>Description</b>	This command specifies the default metric for IPv6 unicast.
<b>Default</b>	no default-ipv6-unicast-metric
<b>Parameters</b>	<i>ipv6-metric</i> — Specifies the default metric for IPv6 unicast.
	<b>Values</b> 1 — 16777215

## default-metric

<b>Syntax</b>	<b>default-metric</b> <i>ipv4 metric</i> <b>no default-metric</b>
<b>Context</b>	config>service>vprn>isis>level
<b>Description</b>	This command specifies the configurable default metric used for all IS-IS interfaces on this level. This value is not used if a metric is configured for an interface.
<b>Default</b>	10
	<i>ipv4 metric</i> — Specifies the default metric for IPv4 unicast.
<b>Values</b>	1 — 16777215

## external-preference

<b>Syntax</b>	<b>external-preference</b> <i>preference</i> <b>no external-preference</b>
<b>Context</b>	config>service>vprn>isis>level
<b>Description</b>	<p>This command configures the external route preference for the IS-IS level.</p> <p>The <b>external-preference</b> command configures the preference level of either IS-IS level 1 or IS-IS level 2 external routes. By default, the preferences are as listed in the table below.</p> <p>A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference decides the route to use.</p> <p>Different protocols should not be configured with the same preference, if this occurs the tiebreaker is dependent on the default preference table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of the route to use is determined by the configuration of the <b>ecmp</b> in the <b>config&gt;router</b> context.</p>
<b>Default</b>	Default preferences are listed in the following table:

Route Type	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
MPLS	7	—
OSPF internal routes	10	No
IS-IS Level 1 internal	15	Yes <sup>a</sup>
IS-IS Level 2 internal	18	Yes*

Route Type	Preference	Configurable
OSPF external	150	Yes
IS-IS Level 1 external	160	Yes
IS-IS Level 2 external	165	Yes
TMS	167	No
BGP	170	Yes
BGP	170	Yes

- a. Internal preferences are changed using the **preference** command in the `config>router>isis>level level-number` context

**Parameters** *preference* — The preference for external routes at this level as expressed.

**Values** 1 — 255

## preference

**Syntax** **preference** *preference*  
**no preference**

**Context** `cconfig>service>vprn>isis>level`

**Description** This command configures the preference level of either IS-IS Level 1 or IS-IS Level 2 internal routes. By default, the preferences are listed in the table below.

A route can be learned by the router by different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the table below. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision what route to use is determined by the configuration of the **ecmp** in the `config>router` context.

**Default** Default preferences are listed in the following table:

Route Type	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
MPLS	7	—
OSPF internal routes	10	No
IS-IS level 1 internal	15	Yes

Route Type	Preference	Configurable
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes <sup>a</sup>
IS-IS level 2 external	165	Yes <sup>a</sup>
BGP	170	Yes

a. External preferences are changed using the **external-preference** command in the `config>router>isis>level level-number` context.

**Parameters** *preference* — The preference for external routes at this level expressed as a decimal integer.

**Values** 1 — 255

## wide-metrics-only

**Syntax** **[no] wide-metrics-only**

**Context** `config>service>vprn>isis>level`

**Description** This command enables the exclusive use of wide metrics in the LSPs for the level number.. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added, again, one for the adjacency and one for the IP prefix.

By default, both sets of TLVs are generated. When `wide-metrics-only` is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

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

## level-capability

**Syntax** **level-capability {level-1 | level-2 | level-1/2}**  
**no level-capability**

**Context** `config>service>vprn>isis`  
`config>service>vprn>isis>interface`

**Description** This command configures the routing level for an instance of the IS-IS routing process. An IS-IS router and an IS-IS interface can operate at Level 1, Level 2 or both Level 1 *and* 2.

Table 15 displays configuration combinations and the potential adjacencies that can be formed.

**Table 15: Potential Adjacency Capabilities**

Global Level	Interface Level	Potential Adjacency
L 1/2	L 1/2	Level 1 and/or Level 2
L 1/2	L 1	Level 1 only
L 1/2	L 2	Level 2 only
L 2	L 1/2	Level 2 only
L 2	L 2	Level 2 only
L 2	L 1	none
L 1	L 1/2	Level 1 only
L 1	L 2	none
L 1	L 1	Level 1 only

The **no** form of the command removes the level capability from the configuration.

**Special Cases** **IS-IS Router** — In the **config>router>isis** context, changing the **level-capability** performs a restart on the IS-IS protocol instance.

**IS-IS Interface** — In the **config>router>isis>interface** context, changing the **level-capability** performs a restart of IS-IS on the interface.

**Default** level-1/2

**Parameters** **level-1** — Specifies the router/interface can operate at Level 1 only.  
**level-2** — Specifies the router/interface can operate at Level 2 only.  
**level-1/2** — Specifies the router/interface can operate at both Level 1 and Level 2.

## link-group

**Syntax** **[no] link-group link-group-name**

**Context** config>service>vprn>isis

**Description** This command configures a link-group for the router or VPRN instance.  
 The **no** form of the command removes the specified link-group.

**Parameters** *link-group-name* — Name of the link-group to be added or removed from the router or VPRN service.

## description

<b>Syntax</b>	<b>description</b> <i>string</i> <b>no description</b>
<b>Context</b>	config>service>vprn>isis>link-group
<b>Description</b>	This command adds a description string to the associated link-group. The string can be up to 256 characters long and can only contain printable characters. If the command is issued in the context of a link-group that already contains a description then the previous description string is replaced. The <b>no</b> form of the command removes the description from the associated link-group.
<b>Parameters</b>	<i>string</i> — Character string to be associated with the associated link-group.
<b>Default</b>	revert-members <i>oper-members</i>

## ipv4-multicast-metric-offset

<b>Syntax</b>	<b>ipv4-multicast-metric-offset</b> <i>offset-value</i> <b>no ipv4-multicast-metric-offset</b>
<b>Context</b>	config>service>vprn>isis>link-group
<b>Description</b>	This command sets the offset value for the IPv4 multicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric for the IPv4 multicast topology. The <b>no</b> form of the command reverts the offset value to 0.
<b>Default</b>	no ipv4-multicast-metric-offset
<b>Parameters</b>	<i>offset-value</i> — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold
<b>Values</b>	0 — 6777215

## ipv4-unicast-metric-offset

<b>Syntax</b>	<b>ipv4-unicast-metric-offset</b> <i>offset-value</i> <b>no ipv4-unicast-metric-offset</b>
<b>Context</b>	config>service>vprn>isis>link-group
<b>Description</b>	This command sets the offset value for the IPv4 unicast address family. If the number of operational links drops below the oper-members threshold, the configured offset is applied to the interface metric. The <b>no</b> form of the command reverts the offset value to 0.
<b>Default</b>	no ipv4-unicast-metric-offset

**Parameters** *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

**Values** 0 — 6777215

## ipv6-unicast-metric-offset

**Syntax** **ipv6-unicast-metric-offset** *offset-value*  
**no ipv6-unicast-metric-offset**

**Context** config>service>vprn>isis>link-group

**Description** This command sets the offset value for the IPv6 unicast address family. If the number of operational links drops below the **oper-members** threshold, the configured offset is applied to the interface metric for the IPv6 topology.

The **no** form of the command reverts the offset value to 0.

**Default** no ipv6-unicast-metric-offset

**Parameters** *offset-value* — Specifies the amount the interface metric for the associated address family is to be increased if the number of operational members in the associated link-group drops below the oper-members threshold.

**Values** 0 — 6777215

## member

**Syntax** **[no] member** *interface-name*

**Context** config>service>vprn>isis>link-group>level

**Description** This command adds or removes a links to the associated link-group. The interface name should already exist before it is added to a link-group.

The **no** form of the command removes the specified interface from the associated link-group.

**Parameters** *interface-name* — Name of the interface to be added or removed from the associated link-group.

## oper-members

**Syntax** **oper-members** **[0-8]**  
**no oper-members**

**Context** config>service>vprn>isis>link-group>level

**Description** This command sets the threshold for the minimum number of operational links for the associated link-group. If the number of operational links drops below this threshold, the configured offsets are



applied. For example, oper-members=3. The metric of the member interfaces is increased when the number of interfaces is lower than 3.

The **no** form of the command reverts the oper-members limit to 1.

**Default** oper-members 0

## revert-members

**Syntax** **revert-members [0-8]**  
**no revert-members**

**Context** config>router>isis>link-group  
config>service>vprn>isis>link-group>level

**Description** This command sets the threshold for the minimum number of operational links to return the associated link-group to its normal operating state and remove the associated offsets to the IS-IS metrics. If the number of operational links is equal to or greater than the configured revert-member threshold then the configured offsets are removed.

The **no** form of the command reverts the revert-members threshold back to the default which is equal to the oper-member threshold value.

## loopfree-alternate

**Syntax** [**no**] **loopfree-alternate**

**Context** config>service>vprn>isis

**Description** This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level or under the OSPF routing protocol instance level.

When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.

The **no** form of this command disables the LFA computation by IGP SPF.

**Default** no loopfree-alternate

## loopfree-alternate-exclude

**Syntax** **loopfree-alternate-exclude prefix-policy prefix-policy [prefix-policy... up to 5]**  
**no loopfree-alternate-exclude**

**Context** config>service>vprn>isis

**Description** This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.

The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.

If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF. Note that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.

The default action of the **loopfree-alternate-exclude** command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.

The **no** form deletes the exclude prefix policy.

**Parameters** **prefix-policy** *prefix-policy* — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

## lsp-lifetime

**Syntax** **lsp-lifetime** *seconds*  
**no lsp-lifetime**

**Context** config>service>vprn>isis

**Description** This command sets the time, in seconds, the router wants the LSPs it originates to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the **lsp-lifetime** expires unless the originating router refreshes the LSP. By default, each router refreshes its LSP’s every 20 minutes (1200 seconds) so other routers will not age out the LSP.

The LSP refresh timer is derived from this formula:  $\text{lsp-lifetime}/2$

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

**Default** **1200** — LSPs originated by the router should be valid for 1200 seconds (20 minutes).

**Parameters** *seconds* — The time, in seconds, that the router wants the LSPs it originates to be considered valid by other routers in the domain.

**Values** 350 — 65535

## lsp-mtu-size

**Syntax** **lsp-mtu-size** *size*  
**no lsp-mtu-size**

**Context** config>service>vprn>isis

**Description** This command configures the LSP MTU size. If the *size* value is changed from the default using CLI or SNMP, then ISIS must be restarted in order for the change to take effect. This can be done by performing a **shutdown** command and then a **no shutdown** command in the **config>router>isis** context.

Note: Using the **exec** command to execute a configuration file to change the LSP MTU-size from its default value will automatically bounce IS-IS for the change to take effect.

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

<b>Default</b>	1492
<b>Parameters</b>	<i>size</i> — Specifies the LSP MTU size.
<b>Values</b>	490 — 9190

## lsp-refresh-interval

<b>Syntax</b>	<b>lsp-refresh-interval</b> <i>seconds</i> <b>no lsp-refresh-interval</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command configures the IS-IS LSP refresh timer interval for the VPRN instance. When configuring the LSP refresh interval, the value that is specified for <b>lsp-lifetime</b> must also be considered. The LSP refresh interval cannot be greater than 90% of the LSP lifetime.  The no form of the command reverts to the default (600 seconds), unless this value is greater than 90% of the LSP lifetime. For example, if the LSP lifetime is 400, then the no lsp-refresh-interval command will be rejected.
<b>Default</b>	600
<b>Parameters</b>	<i>seconds</i> — Specifies the refresh interval.
<b>Values</b>	150— 65535

## lsp-wait

<b>Syntax</b>	<b>lsp-wait</b> <i>lsp-wait</i> [ <i>lsp-initial-wait</i> [ <i>lsp-second-wait</i> ]]
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command is used to customize the throttling of IS-IS LSP-generation. Timers that determine when to generate the first, second and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the second <b>lsp-wait</b> timer until a maximum value is reached.
<b>Parameters</b>	<i>lsp-max-wait</i> — Specifies the maximum interval in seconds between two consecutive occurrences of an LSP being generated.
<b>Values</b>	1 — 120
<b>Default</b>	5
	<i>lsp-initial-wait</i> — Specifies the initial LSP generation delay in seconds.
<b>Values</b>	0 — 100
<b>Default</b>	0

## ISIS Commands

*lsp-second-wait* — Specifies the hold time in seconds between the first and second LSP generation.

**Values** 1 — 100

**Default** 1

## multi-topology

**Syntax** [no] multi-topology

**Context** config>service>vprn>isis

**Description** This command enables IS-IS multi-topology support.

**Default** disabled

## ipv4-multicast

**Syntax** [no] ipv4-multicast

**Context** config>service>vprn>isis>multi-topology

**Description** This command enables support for the IPv4 topology (MT3) within the associate IS-IS instance. The **no** form of this command disables support for the IPv4 topology (MT3) within the associated IS-IS instance.

**Default** no ipv4-multicast

## ipv6-unicast

**Syntax** [no] ipv6-unicast

**Context** config>service>vprn>isis>multi-topology

**Description** This command enables multi-topology TLVs. The no form of the command disables multi-topology TLVs.

## multicast-import

**Syntax** [no] multicast-import

**Context** config>service>vprn>isis

**Description** This command enables ISIS to submit routes into the multicast Route Table Manager (RTM). The **no** form of the command disables the submission of routes into the multicast RTM.

**Default** no multicast-import

## overload

<b>Syntax</b>	<b>overload</b> [timeout <i>seconds</i> ] <b>no overload</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.</p> <p>During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not be used for other transit traffic.</p> <p>If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.</p> <p>The <b>overload</b> command can be useful in circumstances where the router is overloaded or used prior to executing a <b>shutdown</b> command to divert traffic around the router.</p> <p>The <b>no</b> form of the command causes the router to exit the overload state.</p>
<b>Default</b>	<b>no overload</b>
<b>Parameters</b>	<p><i>seconds</i> — The time, in seconds, that this router must operate in overload state.</p> <p><b>Default</b>      infinity (overload state maintained indefinitely)</p> <p><b>Values</b>        60 — 1800</p>

## overload-on-boot

<b>Syntax</b>	<b>overload-on-boot</b> [timeout <i>seconds</i> ] <b>no overload-on-boot</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	<p>When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:</p> <ol style="list-style-type: none"> <li>1. The timeout timer expires.</li> <li>2. A manual override of the current overload state is entered with the <b>config&gt;router&gt;isis&gt;no overload</b> command.</li> </ol> <p>The <b>no overload</b> command does not affect the <b>overload-on-boot</b> function.</p> <p>If no timeout is specified, IS-IS will go into overload indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:</p> <p>L1 LSDB Overload : Manual on boot (Indefinitely in overload)</p> <p>L2 LSDB Overload : Manual on boot (Indefinitely in overload)</p> <p>This state can be cleared with the <b>config&gt;router&gt;isis&gt;no overload</b> command.</p>

When specifying a timeout value, IS-IS will go into overload for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time the system stays in overload:

L1 LSDB Overload : Manual on boot (Overload Time Left : 17)

L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **config>router>isis>no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

**Default** no overload-on-boot

Use the **show router isis status** command to display the administrative and operational state as well as all timers.

**Parameters** **timeout seconds** — Configure the timeout timer for overload-on-boot in seconds.

**Values** 60 — 1800

## poi-tlv-enable

**Syntax** **poi-tlv-enable**  
**no poi-tlv-enable**

**Context** config>service>vprn>isis

**Description** Enable use of Purge Originator Identification (POI) TLV for this IS-IS instance. The POI is added to purges and contains the system ID of the router that generated the purge, which simplifies troubleshooting and determining what caused the purge.

The **no** form of the command removes the POI functionality from the configuration.

**Default** no poi-tlv-enable

## psnp-authentication

**Syntax** [**no**] **psnp-authentication**

**Context** config>service>vprn>isis  
config>service>vprn>isis>level

**Description** This command enables authentication of individual ISIS packets of partial sequence number PDU (PSNP) type.

The **no** form of the command suppresses authentication of PSNP packets.

## reference-bandwidth

<b>Syntax</b>	<b>reference-bandwidth</b> <i>bandwidth-in-kbps</i> <b>reference-bandwidth</b> [ <b>tbps</b> <i>Tera-bps</i> ] [ <b>gbps</b> <i>Giga-bps</i> ] [ <b>mbps</b> <i>Mega-bps</i> ] [ <b>kbps</b> <i>Kilo-bps</i> ] <b>no reference-bandwidth</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command configures the reference bandwidth that provides the basis of bandwidth relative costing.  In order to calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, then the cost is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$$

If the reference bandwidth is configured as 10 Gigabits (10,000,000,000), a 100 M/bps interface has a default metric of 100. In order for metrics in excess of 63 to be configured, wide metrics must be deployed. (See `wide-metrics-only` in the `config>router>isis` context.)

If the reference bandwidth is not configured, then all interfaces have a default metric of 10.

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

<b>Default</b>	<b>no reference-bandwidth</b> — No reference bandwidth is defined. All interfaces have a metric of 10.
<b>Parameters</b>	<i>bandwidth-in-kbps</i> — The reference bandwidth in kilobits per second expressed as a decimal integer. <b>Values</b> 1 — 1000000000 <i>tbps Tera-bps</i> — The reference bandwidth in terabits per second expressed as a decimal integer. <b>Values</b> 1 — 4 <i>gbps Giga-bps</i> — The reference bandwidth in gigabits per second expressed as a decimal integer. <b>Values</b> 1 — 999 <i>mbps Mega-bps</i> — The reference bandwidth in megabits per second expressed as a decimal integer. <b>Values</b> 1 — 999 <i>kbps Kilo-bps</i> — reference bandwidth in kilobits per second expressed as a decimal integer. <b>Values</b> 1 — 999

## rib-priority

**Syntax** **rib-priority** {**high**} *prefix-list-name* | **tag** *tag-value*  
**no rib-priority**

**Context** config>service>vprn>isis

**Description** This command enabled RIB prioritization for the IS-IS protocol and specifies the prefix list or IS-IS tag value that will be used to select the specific routes that should be processed through the IS-IS route calculation process at a higher priority.

The no rib-priority form of command disables RIB prioritization.

**Default** no rib-priority

**Parameters** *prefix-list-name* — specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.  
**tag** *tag-value* — specifies the tag value that is used to match IS-IS routes that are to be processed at a higher priority through the route calculation process.

**Values** 1 – 4294967295

## router-id

**Syntax** **router-id** *ip-address*  
**no router-id**

**Context** config>service>vprn>isis

**Description** This command sets the router ID for a specific VPRN context.  
 If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.

The **no** form of the command removes the router ID definition from the given VPRN context.

**Default** no router-id

**Parameters** *ip-address* — The IP address must be given in dotted decimal notation.

## rsvp-shortcut

**Syntax** [**no**] **rsvp-shortcut**

**Context** config>service>vprn>isis

**Description** This command enables the use of an RSVP-TE shortcut for resolving IGP routes by IS-IS or OSPF routing protocols.

This command instructs IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. If the user enabled the **relative-metric** option for this LSP, IGP will apply the



shortest IGP cost between the endpoints of the LSP plus the value of the offset, instead of the LSP operational metric, when computing the cost of a prefix which is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When `rsvp-shortcut` is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in `configure>router>mpls>lsp>to`, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by entering the `configure>router>mpls>lsp>no igp-shortcut` command.

The SPF in OSPF or IS-IS will only use RSVP LSPs as forwarding adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application, which takes precedence over the LDP-over-RSVP application.

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
  - Ⓜ select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
  - Ⓜ exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
  - Ⓜ prefer tunnel next-hop over ip next-hop
  - Ⓜ within tunnel next-hops:
    - i. select lowest endpoint to destination cost
    - ii. if same endpoint to destination cost, select lowest endpoint node router-id
    - iii. if same router-id, select lowest tunnel-index
  - Ⓜ within ip next-hops:
    - i. select lowest downstream router-id
    - ii. if same downstream router-id, select lowest interface-index
- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel endpoint lies in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

The ingress IOM will spray the packets for this prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

This feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the `rsvp-shortcut` and the `multicast-import` options are enabled in IGP. The unicast RTM can still make use of the tunnel next-hop for the same prefix. This change is made possible with the enhancement by which SPF keeps track of both the direct first hop and the tunneled first hop of a node which is added to the Dijkstra tree.

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.  
 The **no** form of this command disables the resolution of IGP routes using RSVP shortcuts.

**Default** **no rsvp-shortcut**  
**Values**

## spf-wait

**Syntax** **[no] spf-wait** *spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

**Context** config>service>vprn>isis

**Description** This command defines the maximum interval between two consecutive SPF calculations in seconds. Timers that determine when to initiate the first, second and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the *spf-wait* value. The SPF interval will stay at *spf-wait* value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

**Default** no spf-wait

**Parameters** *spf-wait* — Specifies the maximum interval in seconds between two consecutive spf calculations.

**Values** 1 — 120

**Default** 10

*spf-initial-wait* — Specifies the initial SPF calculation delay in milliseconds after a topology change.

**Values** 10 — 100000

**Default** 1000

*spf-second-wait* — Specifies the hold time in milliseconds between the first and second SPF calculation.

**Values** 1 — 100000

**Default** 1000

## strict-adjacency-check

**Syntax** **[no] strict-adjacency-check**

**Context** config>service>vprn>isis

**Description** This command enables strict checking of address families (IPv4 and IPv6) for IS-IS adjacencies. When enabled, adjacencies will not come up unless both routers have exactly the same address families configured. If there is an existing adjacency with unmatched address families, it will be torn down. This command is used to prevent black-holing traffic when IPv4 and IPv6 topologies are dif-

ferent. When disabled (no strict-adjacency-check) a BFD session failure for either IPv4 or Ipv6 will cause the routes for the other address family to be removed as well.

When disabled (**no strict-adjacency-check**), both routers only need to have one common address family to establish the adjacency.

**Default** no strict-adjacency-check

## summary-address

<b>Syntax</b>	<b>summary-address</b> { <i>ip-prefix/mask</i>   <i>ip-prefix</i> [ <i>netmask</i> ]} <i>level</i> [ <b>tag</b> <i>tag</i> ] <b>no summary-address</b> { <i>ip-prefix/mask</i>   <i>ip-prefix</i> [ <i>netmask</i> ]}	
<b>Context</b>	config>service>vprn>isis	
<b>Description</b>	This command creates summary-addresses for the specified router or VPRN instance.	
<b>Default</b>	none	
<b>Parameters</b>	<i>ip-prefix/mask</i> — Specifies information for the specified IP prefix and mask length.	
	<b>Values</b>	ipv4-prefix: a.b.c.d (host bits must be 0) ipv4-prefix-length: 0 — 32 ipv6-prefix: 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 ipv6-prefix-length: [0 — 128]
	<i>netmask</i> — The subnet mask in dotted decimal notation.	
	<b>Values</b>	0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)
	<i>level</i> — Specifies IS-IS level area attributes.	
	<b>Values</b>	level-1, level-2, level-1/2
	<b>tag</b> <i>tag</i> — Assigns a route tag to the summary address.	
	<b>Values</b>	1 — 4294967295

## ignore-attached-bit

<b>Syntax</b>	<b>ignore-attached-bit</b> <b>no ignore-attached-bit</b>	
<b>Context</b>	config>service>vprn>isis	
<b>Description</b>	This command configures IS-IS to ignore the attached bit on received Level 1 LSPs to disable installation of default routes.	

## suppress-attached-bit

<b>Syntax</b>	<b>suppress-attached-bit</b> <b>no suppress-attached-bit</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command configures IS-IS to suppress setting the attached bit on originated Level 1 LSPs to prevent all L1 routers in the area from installing a default route to it.

## unicast-import-disable

<b>Syntax</b>	<b>[no] unicast-import-disable</b>
<b>Context</b>	config>service>vprn>isis
<b>Description</b>	This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM. Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured.
<b>Default</b>	disabled

## OSPF Commands

### ospf

<b>Syntax</b>	<b>[no] ospf</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command enables access to the context to enable an OSPF protocol instance.</p> <p>When an OSPF instance is created, the protocol is enabled. To start or suspend execution of the OSPF protocol without affecting the configuration, use the <b>no shutdown</b> command.</p> <p>The <b>no</b> form of the command deletes the OSPF protocol instance removing all associated configuration parameters.</p>
<b>Default</b>	<b>no ospf</b> — The OSPF protocol is not enabled.

### ospf3

<b>Syntax</b>	<b>ospf3</b> [ <i>instance-id</i> ] [ <i>router-id</i> ] <b>[no] ospf3</b> <i>instance-id</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	<p>This command creates an OSPFv3 routing instance and then enters the associated context to configure associated protocol parameters.</p> <p>When an OSPFv3 instance is created, the protocol is enabled. To start or suspend execution of the OSPF.</p> <p>The <b>no</b> form of the command deletes the OSPFv3 protocol instance, removing all associated configuration parameters.</p>
<b>Default</b>	<b>no default</b>
<b>Parameters</b>	<p><i>instance-id</i> — Specify the instance ID for the OSPFv3 instance being created or modified. The instance ID must match the specified range based on the address family. For ipv6-unicast, the instance id must be between 0 and 31. For ipv4-unicast the instance id must be between 64-95.</p> <p><b>Values</b>     0 — 31: IPV6 unicast</p> <p><b>Values</b>     64—95: IPV4 unicast</p>

## advertise-router-capability

<b>Syntax</b>	<b>advertise-router-capability</b> { <b>link</b>   <b>area</b>   <b>as</b> } <b>no advertise-router-capability</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command enables advertisement of a router's capabilities to its neighbors for informational and troubleshooting purposes. A Router Information (RI) LSA as defined in RFC 4970 advertises the following capabilities:</p> <ul style="list-style-type: none"> <li>• OSPF graceful restart capable: no</li> <li>• OSPF graceful restart helper: yes, when enabled</li> <li>• OSPF Stub Router support: yes</li> <li>• OSPF Traffic Engineering support: yes, when enabled</li> <li>• OSPF point-to-point over LAN: yes</li> <li>• OSPF Experimental TE: no</li> </ul> <p>The parameters (link, area &amp; as) control the scope of the capabilities advertisements. The no form of this command, disables this capability.</p>
<b>Default</b>	<b>no advertise-router-capability</b>
<b>Parameters</b>	<p><b>link</b> — Capabilities are only advertised over local link and not flooded beyond</p> <p><b>area</b> — Capabilities are only advertised within the area of origin</p> <p><b>as</b> — Capabilities are only advertised throughout the entire autonomous system</p>

## area

<b>Syntax</b>	<b>[no] area</b> <i>area-id</i>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.</p> <p>The <b>no</b> form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, sham-links, and address-ranges etc., that are currently assigned to this area.</p>
<b>Default</b>	<b>no area</b> — No OSPF areas are defined.
<b>Parameters</b>	<p><i>area-id</i> — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.</p> <p><b>Values</b>     0.0.0.0 — 255.255.255.255 (dotted decimal)                  0 — 4294967295 (decimal integer)</p>

## area-range

<b>Syntax</b>	<b>area-range</b> <i>ip-prefix/prefix-length</i> [ <b>advertise</b>   <b>not-advertise</b> ] <b>no area-range</b> <i>ip-prefix/mask</i> <b>area-range</b> <i>ipv6-prefix/prefix-length</i> [ <b>advertise</b>   <b>not-advertise</b> ] <b>no area-range</b> <i>ip-prefix/mask</i>
<b>Context</b>	config>service>vprn>ospf>area ospf>service>vprn>nssa config>service>vprn>ospf3>area
<b>Description</b>	<p>This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.</p> <p>ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.</p> <p>The <b>no</b> form of the command deletes the range (non) advertisement.</p>
<b>Default</b>	<b>no area-range</b> — No range of addresses are defined.
<b>Special Cases</b>	<b>NSSA Context</b> — In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs. <b>Area Context</b> — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.
<b>Parameters</b>	<i>ipv6-prefix/prefix-length</i> — The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area. <b>Values</b> <i>ipv6-prefix</i> - 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 <i>prefix-length</i> - [0..128] <i>mask</i> — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation. <b>Values</b> 0 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal) <b>advertise</b>   <b>not-advertise</b> — Specifies whether or not to advertise the summarized range of addresses into other areas. The <b>advertise</b> keyword indicates the range will be advertised, and the keyword <b>not-advertise</b> indicates the range will not be advertised. The default is <b>advertise</b> .

## blackhole-aggregate

<b>Syntax</b>	[no] <b>blackhole-aggregate</b>
<b>Context</b>	config>service>vprn>ospf>area

config>service>vprn>ospf3>area

- Description** This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.
- It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem configure the blackhole aggregate option.
- The **no** form of this command removes this option.
- Default** blackhole-aggregate

## interface

- Syntax** **[no] interface** *ip-int-name* [**secondary**]
- Context** config>service>vprn>ospf>area  
config>service>vprn>ospf3>area
- Description** This command creates a context to configure an OSPF interface.
- By default interfaces are not activated in any interior gateway protocol such as OSPF unless explicitly configured.
- The **no** form of the command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context can be used to disable an interface without removing the configuration for the interface.
- Default** **no interface** — No OSPF interfaces are defined.
- Parameters** *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 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.
- If the IP interface name does not exist or does not have an IP address configured an error message will be returned.
- If the IP interface exists in a different area it will be moved to this area.
- secondary** — Allows multiple secondary adjacencies to be established over a single IP interface.

## sham-link

- Syntax** **sham-link** *ip-int-name ip-address*
- Context** config>service>vprn>ospf>area
- Description** This command is similar to a virtual link with the exception that metric must be included in order to distinguish the cost between the MPLS-VPRN link and the backdoor.



- Parameters**
- ip-int-name* — The local interface name used for the sham-link. This is a mandatory parameter and interface names must be unique within the group of defined IP interfaces for **config>router>interface**, **config>service>ies>interface** and **config>service>vprn>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters, the entire string must be enclosed within double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.
- ip-address* — The IP address of the SHAM-link neighbor in IP address dotted decimal notation. This parameter is the remote peer of the sham link's IP address used to set up the SHAM link. This is a mandatory parameter and must be a valid IP address.

## advertise-subnet

- Syntax** **[no] advertise-subnet**
- Context** config>service>vprn>ospf>area>if
- Description** This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.
- Note that this command is not supported in the OSPF3 context.
- The **no** form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.
- Default** **advertise-subnet** — Advertises point-to-point interfaces as subnet routes.

## authentication

- Syntax** **authentication bidirectional sa-name**  
**authentication inbound sa-name outbound sa-name**  
**no authentication**
- Context** config>service>vprn>ospf3>area>if
- Description** This command configures OPSFv3 confidentiality authentication.
- The no form of the command removes the SA name from the configuration.
- Parameters**
- bidirectional sa-name** — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated.
- inbound sa-name** — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated.
- outbound sa-name** — Specifies the IPSec security association name in case the OSPFv3 traffic on the interface has to be authenticated.

## authentication-key

<b>Syntax</b>	<b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ] <b>no authentication-key</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf>area>sham-link
<b>Description</b>	<p>This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.</p> <p>Note that this command is not valid in the OSPF3 context.</p> <p>All neighboring routers must use the same type of authentication and password for proper protocol communication. If the <b>authentication-type</b> is configured as password, then this key must be configured.</p> <p>By default, no authentication key is configured.</p> <p>Note that this command is not supported in the OSPF context.</p> <p>The <b>no</b> form of the command removes the authentication key.</p>
<b>Default</b>	<b>no authentication-key</b> — No authentication key is defined.
<b>Parameters</b>	<p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“”).</p> <p><i>hash-key</i> — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“”).</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> 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 <b>hash</b> parameter specified.</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form. If the <b>hash2</b> parameter is not used, the less encrypted <b>hash</b> form is assumed.</p>

## authentication-type

<b>Syntax</b>	<b>authentication-type</b> { <b>password</b>   <b>message-digest</b> } <b>no authentication-type</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link
<b>Description</b>	<p>This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.</p> <p>Note that this command is not valid in the OSPF3 context.</p>

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

The **no** form of the command disables authentication on the interface.

Note that this command is not supported in the OSPF context.

**Default** **no authentication** — No authentication is enabled on an interface.

**Parameters** **password** — This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

**message-digest** — This keyword enables message digest MD5 authentication in accordance with RFC1321. If this option is configured, then at least one message-digest-key must be configured.

## bfd-enable

**Syntax** **bfd-enable [remain-down-on-failure]**  
**no bfd-enable**

**Context** config>service>vprn>ospf>interface>if  
config>service>vprn>ospf3>area>if

### Description

**Description** This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

**Default** **no bfd-enable**

**Parameters** **remain-down-on-failure** — Forces adjacency down on BFD failure.

## dead-interval

**Syntax** **dead-interval** *seconds*  
**no dead-interval**

**Context** config>service>vprn>ospf>area>if  
config>service>vprn>ospf>area>virtual-link  
config>service>vprn>ospf3>area>if  
config>service>vprn>ospf3>area>virtual-link  
config>service>vprn>ospf>area>sham-link

**Description** This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

## OSPF Commands

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

**Default** 40

**Special Cases** **OSPF Interface** — If the **dead-interval** configured applies to an interface, then all nodes on the subnet must have the same dead interval.

**Virtual Link** — If the **dead-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.

**Sham-link** — If the **dead-interval** configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same dead interval.

**Parameters** *seconds* — The dead interval expressed as a decimal integer.

**Values** 2 — 2147483647 seconds

## graceful-restart

**Syntax** **[no] graceful-restart**

**Context** config>service>vprn>ospf

**Description** This command enables or disables graceful-restart for VPRN OSPF.  
This command is not available for OSPF3.

## helper-disable

**Syntax** **helper-disable**

**Context** config>service>vprn>ospf>graceful-restart

**Description** This command disables the helper support for graceful restart.  
When **graceful-restart** is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The SR OS supports only helper mode. This facilitates the graceful restart of neighbors but will not act as a restarting router (meaning that the SR OS will not help the neighbors to restart).

This command is not available for OSPF3.

The **no helper-disable** command enables helper support and is the default when graceful-restart is enabled.

**Default** disabled

## ignore-dn-bit

**Syntax** **[no] ignore-dn-bit**

**Context** config>service>vprn>ospf

```
config>service>vprn>ospf3
```

**Description** This command specifies whether to suppress the setting of the DN bit for OSPF or OSPF3 LSA packets generated by this instance of OSPF or OSPF3 on the router.

The **no** form of the command enables the OSPF or OSPF3 router to follow the normal procedure to determine whether to set the DN bit.

**Default** no ignore-dn-bit

## import

**Syntax** **import** *policy-name* [*policy-name...*(up to 5 max)]  
**no import**

**Context** config>service>vprn>ospf  
config>service>vprn>ospf3

**Description** This command applies one or more (up to 5) route policies as OSPF import policies. When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination. The flooding of LSAs is unaffected by OSPF import policy actions.

**Default** If an OSPF route has the lowest preference value among all routes to a destination it is installed in the routing table.

**Parameters** *policy-name* — The import route policy name. Allowed values are any string up to 32 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.

The specified name(s) must already be defined.

## hello-interval

**Syntax** **hello-interval** *seconds*  
**no hello-interval**

**Context** config>service>vprn>ospf>area>if  
config>service>vprn>ospf3>area>if  
config>service>vprn>ospf>area>virtual-link  
config>service>vprn>ospf3>area>virtual-link  
config>service>vprn>ospf>area>sham-link

**Description** This command configures the interval between OSPF hellos issued on the interface, virtual link, or sham-link.

The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link and/or router failures at the cost of higher processing costs.

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

<b>Default</b>	<b>hello-interval 10</b> — A 10-second hello interval.
<b>Special Cases</b>	<p><b>OSPF Interface</b> — If the <b>hello-interval</b> configured applies to an interface, then all nodes on the subnet must have the same hello interval.</p> <p><b>Virtual Link</b> — If the <b>hello-interval</b> configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval.</p> <p><b>Sham Link</b> — If the <b>hello-interval</b> configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same hello interval</p>
<b>Parameters</b>	<p><i>seconds</i> — The hello interval in seconds expressed as a decimal integer.</p> <p><b>Values</b> 1 — 65535</p>

## interface-type

<b>Syntax</b>	<b>interface-type {broadcast   point-to-point}</b> <b>no interface-type</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if
<b>Description</b>	<p>This command configures the interface type to be either broadcast or point-to-point.</p> <p>Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided the link is used as a point-to-point.</p> <p>If the interface type is not known at the time the interface is added to OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.</p> <p>The <b>no</b> form of the command reverts to the default value.</p>
<b>Default</b>	<p><b>point-to-point</b> — If the physical interface is SONET.</p> <p><b>broadcast</b> — If the physical interface is Ethernet or unknown.</p>
<b>Special Cases</b>	<b>Virtual-Link</b> — A virtual link is always regarded as a point-to-point interface and not configurable.
<b>Parameters</b>	<p><b>broadcast</b> — Configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.</p> <p><b>point-to-point</b> — Configures the interface to maintain this link as a point-to-point link.</p>

## loopfree-alternate-exclude

<b>Syntax</b>	<b>[no] loopfree-alternate-exclude</b>
<b>Context</b>	configure>service>vpn>ospf>area configure>service>vpn>ospf3>area configure>service>vpn>ospf>area>interface configure>service>vpn>ospf3>area>interface
<b>Description</b>	<p>This command instructs IGP to not include a specific interface or all interfaces participating in a specific IS-IS level or OSPF area in the SPF LFA computation. This provides a way of reducing the LFA SPF calculation where it is not needed.</p> <p>When an interface is excluded from the LFA SPF in IS-IS, it is excluded in both level 1 and level 2. When it is excluded from the LFA SPF in OSPF, it is excluded in all areas. However, the above OSPF command can only be executed under the area in which the specified interface is primary and once enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.</p> <p>The <b>no</b> form of this command re-instates the default value for this command.</p>
<b>Default</b>	no loopfree-alternate-exclude.

## lsa-filter-out

<b>Syntax</b>	<b>lsa-filter-out [all   except-own-rtrlsa   except-own-rtrlsa-and-defaults] no lsa-filter-out</b>
<b>Context</b>	config>router>ospf>area>interface config>router>ospf3>area>interface config>service>vpn>ospf>area>interface config>service>vpn>ospf3>area>interface
<b>Description</b>	<p>This command enables filtering of outgoing OSPF LSAs on the selected OSPFv2 or OSPFv3 interface. Three filtering options are provided:</p> <ul style="list-style-type: none"> <li>• Do not flood any LSAs out the interface. This option is suitable if the neighbor is simply-connected and has a statically configured default route with the address of this interface as next-hop.</li> <li>• Flood the router's own router-LSA out the interface and suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and has a statically configured default route with a loopback or system interface address (contained in the router-LSA) as next-hop.</li> <li>• Flood the router's own router-LSA and all self-generated type-3, type-5 and type-7 LSAs advertising a default route (0/0) out the interface; suppress all other flooded LSAs. This option is suitable if the neighbor is simply-connected and does not have a statically configured default route.</li> </ul> <p>The <b>no</b> form of this command disables OSPF LSA filtering (normal operation).</p>
<b>Default</b>	no lsa-filter-out

## multicast-import

<b>Syntax</b>	<b>[no] multicast-import</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	This command enables the submission of routes into the multicast Route Table Manager (RTM) by OSPF.  The <b>no</b> form of the command disables the submission of routes into the multicast RTM.
<b>Default</b>	<b>no multicast-import</b>

## message-digest-key

<b>Syntax</b>	<b>message-digest-key</b> <i>keyid</i> <b>md5</b> [ <i>key</i>   <i>hash-key</i> ] [ <b>hash</b> ] <b>no message-digest-key</b> <i>keyid</i>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf>area>virtual-link config>service>vprn>ospf>area>sham-link
<b>Description</b>	This command configures a message digest key when MD5 authentication is enabled on the interface, virtual-link or sham-link. Multiple message digest keys can be configured.  Note that this command is not valid in the OSPF3 context.  The <b>no</b> form of the command removes the message digest key identified by the <i>key-id</i> .
<b>Default</b>	No message digest keys are defined.
<b>Parameters</b>	<b>keyid</b> — The <i>keyid</i> is expressed as a decimal integer.  <b>Values</b> 1 — 255  <b>md5</b> <i>key</i> — The MD5 key. The <i>key</i> can be any alphanumeric string up to 16 characters in length.  <b>md5</b> <i>hash-key</i> — The MD5 hash key. The key can be any combination of ASCII characters up to 32 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.  <b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> 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 <b>hash</b> parameter specified.



## metric

<b>Syntax</b>	<b>metric</b> <i>metric</i> <b>no metric</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if config>service>vprn>ospf>area>sham-link
<b>Description</b>	This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.  The <b>no</b> form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the <b>reference-bandwidth</b> command setting and the speed of the underlying link.
<b>Default</b>	<b>no metric</b> — The metric is based on <b>reference-bandwidth</b> setting and the link speed.
<b>Parameters</b>	<i>metric</i> — The metric to be applied to the interface expressed as a decimal integer.  <b>Values</b> 1 — 65535

## mtu

<b>Syntax</b>	<b>mtu</b> <i>bytes</i> <b>no mtu</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if
<b>Description</b>	This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:  <pre> config&gt;port&gt;ethernet config&gt;port&gt;sonet-sdh&gt;path config&gt;port&gt;tdm&gt;t3-e3 config&gt;port&gt;tdm&gt;t1-e1&gt;channel-group </pre> If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.  To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.  Use the <b>no</b> form of this command to revert to default.
<b>Default</b>	<b>no mtu</b> — Uses the value derived from the MTU configured in the <b>config&gt;port</b> context.
<b>Parameters</b>	<i>bytes</i> — The MTU to be used by OSPF for this logical interface in bytes.  <b>Values</b> 512 — 9198 (9212-14) (Depends on the physical media)

## passive

<b>Syntax</b>	<b>[no] passive</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if
<b>Description</b>	<p>This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.</p> <p>By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The <b>passive</b> parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.</p> <p>While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.</p> <p>The <b>no</b> form of the command removes the passive property from the OSPF interface.</p>
<b>Default</b>	Service interfaces defined in <b>config&gt;router&gt;service-prefix</b> are passive. All other interfaces are not passive.

## priority

<b>Syntax</b>	<b>priority</b> <i>number</i> <b>no priority</b>
<b>Context</b>	config>service>vprn>ospf>area>if config>service>vprn>ospf3>area>if
<b>Description</b>	<p>This command configures the priority of the OSPF interface that is used an election of the designated router on on the subnet.</p> <p>This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.</p> <p>The <b>no</b> form of the command reverts the interface priority to the default value.</p>
<b>Default</b>	<b>priority 1</b>
<b>Parameters</b>	<i>number</i> — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router of Backup Designated Router on the interface subnet.
	<b>Values</b> 0 — 255

## retransmit-interval

<b>Syntax</b>	<b>retransmit-interval</b> <i>seconds</i> <b>no retransmit-interval</b>
<b>Context</b>	config>service>vprn>ospf>area>if

```
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf3>area>if
config>service>vprn>ospf3>area>virtual-link
config>service>vprn>ospf>area>sham-link
```

**Description** This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit-interval expires and no acknowledgement has been received, the LSA will be retransmitted.

The **no** form of this command reverts to the default interval.

**Default** retransmit-interval 5

**Parameters** *seconds* — The retransmit interval in seconds expressed as a decimal integer.

**Values** 1 — 3600

## rib-priority

**Syntax** **rib-priority** {**high**} *prefix-list-name*  
**no rib-priority**

**Context** config>service>vprn>ospf>area>interface  
config>service>vprn>ospf3>area>interface

**Description** This command enables RIB prioritization for the OSPF/OSPFv3 protocol. When enabled at the OSPF interface level, all routes learned through the associated OSPF interface will be processed through the OSPF route calculation process at a higher priority.

The **no** form of **rib-priority** command disables RIB prioritization at the associated level.

**Default** no rib-priority

**Parameters** *prefix-list-name* — specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

## transit-delay

**Syntax** **transit-delay** *seconds*  
**no transit-delay**

**Context** config>service>vprn>ospf>area>if  
config>service>vprn>ospf3>area>if  
config>service>vprn>ospf>area>virtual-link  
config>service>vprn>ospf3>area>virtual-link  
config>service>vprn>ospf>area>sham-link

**Description** This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link or sham-link.

## OSPF Commands

The **no** form of this command reverts to the default delay time.

<b>Default</b>	<b>transit-delay 1</b>
<b>Parameters</b>	<i>seconds</i> — The transit delay in seconds expressed as a decimal integer.
<b>Values</b>	0 — 3600

## key-rollover-interval

<b>Syntax</b>	<b>key-rollover-interval</b> <i>key-rollover-interval</i>
<b>Context</b>	config>service>vprn>ospf3>area
<b>Description</b>	This command configures the key rollover interval. The <b>no</b> form of the command reverts to the default.
<b>Default</b>	10
<b>Parameters</b>	<i>key-rollover-interval</i> — Specifies the time, in seconds, after which a key rollover will start.
<b>Values</b>	10 — 300

## loopfree-alternate-exclude

<b>Syntax</b>	<b>[no] loopfree-alternate-exclude</b>
<b>Context</b>	config>service>vprn>ospf3>area
<b>Description</b>	This command specifies whether or not the OSPF area should be excluded during LFA calculations. When enabled, the OSPF area is excluded from LFA calculations. When disabled (the default), the OSPF area is included in LFA calculations. The <b>no</b> form of the command includes the OSPF area in LFA calculations.
<b>Default</b>	disabled

## nssa

<b>Syntax</b>	<b>[no] nssa</b>
<b>Context</b>	config>service>vprn>ospf>area config>service>vprn>ospf3>area
<b>Description</b>	This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area. NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain.

Existing virtual links of a non-stub or NSSA area will be removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of the command removes the NSSA designation and configuration context from the area.

**Default** **no nssa** — The OSPF area is not an NSSA.

## originate-default-route

**Syntax** **originate-default-route [type-7]**  
**no originate-default-route**

**Context** config>service>vprn>ospf>area>nssa  
config>service>vprn>ospf3>area>nssa

**Description** This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR)

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of the command disables origination of a default route.

**Default** **no originate-default-route** — A default route is not originated.

**Parameters** **type-7** — Specifies a type 7 LSA should be used for the default route.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

**Default** Type 3 LSA for the default route.

## redistribute-external

**Syntax** **[no] redistribute-external**

**Context** config>service>vprn>ospf>area>nssa  
config>service>vprn>ospf3>area>nssa

**Description** This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF domain.

## OSPF Commands

The **no** form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

**Default** **redistribute-external** — External routes are redistributed into the NSSA.

## summaries

**Syntax** **[no] summaries**

**Context** config>service>vprn>ospf>area>nssa  
config>service>vprn>ospf>area>stub  
config>service>vprn>ospf3>area>nssa

**Description** This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR). This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or nssa area. By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

**Default** **summaries** — Summary routes are advertised by the ABR into the stub area or NSSA.

## stub

**Syntax** **[no] stub**

**Context** config>service>vprn>ospf>area  
config>service>vprn>ospf3>area

**Description** This command enables access to the context to configure an OSPF stub area and adds/removes the stub designation from the area. External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF area cannot be both an NSSA and a stub area. Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB.

By default, an area is not a stub area.

The **no** form of the command removes the stub designation and configuration context from the area.

**Default** **no stub** — The area is not configured as a stub area.

## default-metric

**Syntax** **default-metric** *metric*  
**no default-metric**

**Context** config>service>vprn>ospf>area>stub  
config>service>vprn>ospf3>area>stub

<b>Description</b>	This command configures the metric used by the area border router (ABR) for the default route into a stub area. The default metric should only be configured on an ABR of a stub area. An ABR generates a default route if the area is a <b>stub</b> area.  The <b>no</b> form of the command reverts to the default value.
<b>Default</b>	default-metric 1
<b>Parameters</b>	<i>metric</i> — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.  <b>Values</b> 1 — 16777215

## virtual-link

<b>Syntax</b>	<b>[no] virtual-link</b> <i>router-id</i> <b>transit-area</b> <i>area-id</i>
<b>Context</b>	config>service>vprn>ospf>area config>service>vprn>ospf3>area
<b>Description</b>	This command configures a virtual link to connect area border routers to the backbone via a virtual link. The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone (see area 0.0.0.2 in the picture below) then the area border routers (routers 1 and 2 in the picture below) must be connected via a virtual link. The two area border routers will form a point-to-point like adjacency across the transit area (area 0.0.0.1 in the picture below). A virtual link can only be configured while in the area 0.0.0.0 context.  The <i>router-id</i> specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).  The <b>no</b> form of the command deletes the virtual link.
<b>Default</b>	No virtual link is defined.
<b>Parameters</b>	<i>router-id</i> — The router ID of the virtual neighbor in IP address dotted decimal notation.  <b>transit-area</b> <i>area-id</i> — The area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.  The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see Area 0.0.0.5 in <a href="#">Figure 30</a> ) then the area border routers (such as routers Y and Z) must be connected via a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area (see Area 0.0.0.4).

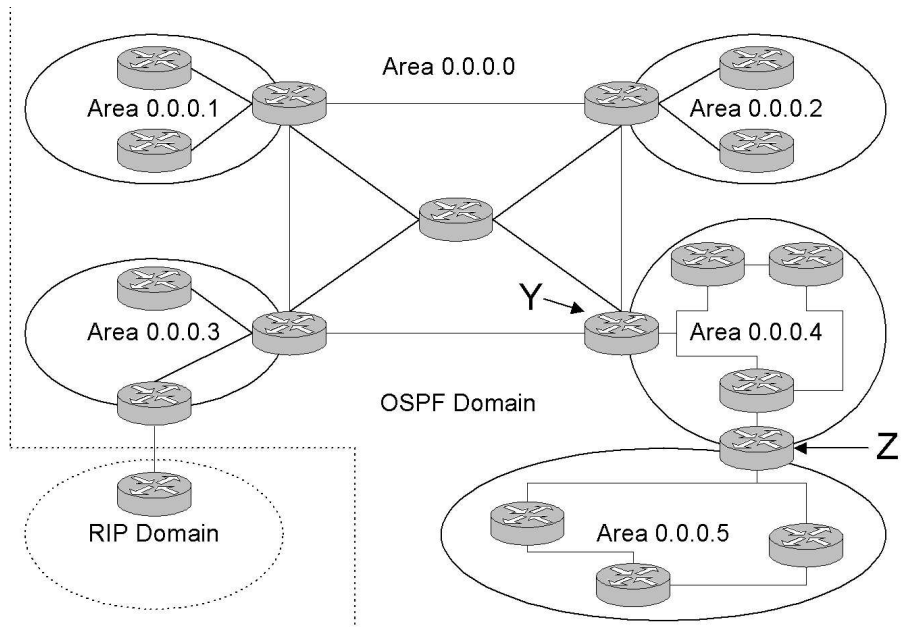


Figure 30: OSPF Areas

## compatible-rfc1583

<b>Syntax</b>	<b>[no] compatible-rfc1583</b>
<b>Context</b>	config>service>vprn>ospf
<b>Description</b>	<p>This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.</p> <p>RFC1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.</p> <p>Although it would be favorable to require all routers to run a more current compliancy level, this command allows the router to use obsolete methods of calculation.</p> <p>This command is not supported in OSPF3.</p> <p>The <b>no</b> form of the command enables the post-RFC1583 method of summary and external route calculation.</p>
<b>Default</b>	<b>compatible-rfc1583</b> — RFC1583 compliance is enabled.

## export

<b>Syntax</b>	<b>export</b> <i>policy-name</i> [ <i>policy-name...</i> ] <b>no export</b>
<b>Context</b>	config>service>vprn>ospf



```
config>service>vprn>ospf3
```

- Description** This command associates export route policies to determine which routes are exported from the routing table to OSPF. Export policies are only in effect if OSPF is configured as an ASBR.
- If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.
- If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.
- The **no** form of the command removes all policies from the configuration.
- Default** **no export** — No export route policies specified.
- Parameters** *policy-name* — The export route policy name. Allowed values are any string up to 32 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.
- The specified name(s) must already be defined.

## external-db-overflow

- Syntax** **external-db-overflow** *limit interval*  
**no external-db-overflow**
- Context** config>service>vprn>ospf  
config>service>vprn>ospf3
- Description** This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.
- The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.
- The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.
- The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.
- The **no** form of the command disables limiting the number of non-default AS-external-LSA entries.
- Default** **no external-db-overflow** — No limit on non-default AS-external-LSA entries.
- Parameters** *limit* — The maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.
- Values** -1 — 2147483647

*interval* — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

**Values** 0 — 2147483647

## external-preference

- Syntax** **external-preference** *preference*  
**no external-preference**
- Context** config>service>vprn>ospf  
 config>service>vprn>ospf3
- Description** This command configures the preference for OSPF external routes.  
 A route can be learned by the router from different protocols in which case the costs are not comparable; when this occurs the preference is used to decide which route will be used.  
 Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.  
 If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.  
 The **no** form of the command reverts to the default value.
- Default** **external-preference 150** — OSPF external routes have a default preference of 150.
- Parameters** *preference* — The preference for external routes expressed as a decimal integer.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes <sup>a</sup>
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

a. Preference for OSPF internal routes is configured with the **preference** command.

**Values** 1 — 255

## ignore-dn-bit

<b>Syntax</b>	<b>[no] ignore-dn-bit</b>
<b>Context</b>	config>service>vprn>ospf
<b>Description</b>	This command specifies whether to ignore the DN bit for OSPF LSA packets for this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets will be ignored. When disabled, the DN bit will not be ignored for OSPF LSA packets.

## loopfree-alternate

<b>Syntax</b>	<b>[no] loopfree-alternate</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command enables Loop-Free Alternate (LFA) computation by SPF under the IS-IS routing protocol level, or under the OSPF routing protocol instance level.</p> <p>When this command is enabled, it instructs the IGP SPF to attempt to pre-compute both a primary next-hop and an LFA next-hop for every learned prefix. IS-IS computes the primary SPF first and then computes the LFA SPF. The LFA backup next-hop is only available after the LFA SPF is completed. When found, the LFA next-hop is populated into the routing table along with the primary next-hop for the prefix.</p> <p>The <b>no</b> form of this command disables the LFA computation by IGP SPF.</p>
<b>Default</b>	no loopfree-alternate

## overload

<b>Syntax</b>	<b>overload [timeout seconds]</b> <b>no overload</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continue to reach the router.</p> <p>To put the IGP in an overload state enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a <b>no overload</b> command is executed.</p> <p>If the <b>overload</b> command is encountered during the execution of an <b>overload-on-boot</b> command then this command takes precedence. This could occur as a result of a saved configuration file where</p>

## OSPF Commands

both parameters are saved. When the file is saved by the system the **overload-on-boot** command is saved after the **overload** command.

Use the **no** form of this command to return to the default. When the **no overload** command is executed, the overload state is terminated regardless the reason the protocol entered overload state.

<b>Default</b>	no overload
<b>Parameters</b>	<b>timeout</b> <i>seconds</i> — Specifies the number of seconds to reset overloading.
<b>Values</b>	60 —1800
<b>Default</b>	60

## if-attribute

<b>Syntax</b>	<b>if-attribute</b>
<b>Context</b>	config>router config>router>interface config>service>ies>interface config>service>vprn>interface
<b>Description</b>	This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG).

## admin-group

<b>Syntax</b>	<b>admin-group</b> <i>group-name</i> [ <i>group-name...</i> (up to 5 max)] <b>no admin-group</b> <i>group-name</i> [ <i>group-name...</i> (up to 5 max)] <b>no admin-group</b>
<b>Context</b>	config>router>interface>if-attribute config>service>ies>interface>if-attribute config>service>vprn>interface>if-attribute config>router>mpls>interface
<b>Description</b>	<p>This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface.</p> <p>Each single operation of the <b>admin-group</b> command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.</p> <p>The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.</p> <p>It should be noted that only the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the <b>traffic-engineering</b> option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.</p>

The **no** form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

**Parameters** *group-name* — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

## srlg-group

**Syntax** **srlg-group** *group-name* [*group-name...*(up to 5 max)]  
**no srlg-group** *group-name* [*group-name...*(up to 5 max)]  
**no srlg-group**

**Context** config>router>interface>if-attribute  
 config>service>ies>interface>if-attribute  
 config>service>vprn>interface>if-attribute  
 config>router>mpls>interface

**Description** This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.

An interface can belong to up to 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

It should be noted that only the SRLGs bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

**Parameters** *group-name* — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

## lfa-policy-map

**Syntax** **lfa-policy-map** *route-nh-template* *template-name*  
**no lfa-policy-map**

**Context** config>router>ospf>area>interface  
 config>router>ospf3>area>interface  
 config>router>isis>interface  
 config>service>vprn>ospf>area>interface  
 config>service>vprn>ospf3>area>interface

**Description** This command applies a route next-hop policy template to an OSPF or IS-IS interface.

When a route next-hop policy template is applied to an interface in IS-IS, it is applied in both level 1 and level 2. When a route next-hop policy template is applied to an interface in OSPF, it is applied in

all areas. However, the command in an OSPF interface context can only be executed under the area in which the specified interface is primary and then applied in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

If the user excluded the interface from LFA using the command **loopfree-alternate-exclude**, the LFA policy, if applied to the interface, has no effect.

Finally, if the user applied a route next-hop policy template to a loopback interface or to the system interface, the command will not be rejected, but it will result in no action being taken.

The **no** form deletes the mapping of a route next-hop policy template to an OSPF or IS-IS interface.

**Parameters** *template-name* — Specifies the name of the template, up to 32 characters.

## loopfree-alternate-exclude

<b>Syntax</b>	<b>loopfree-alternate-exclude prefix-policy <i>prefix-policy</i> [<i>prefix-policy</i>... up to 5]</b> <b>no loopfree-alternate-exclude</b>
<b>Context</b>	config>router>ospf config>router>ospf3 config>router>isis config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command excludes from LFA SPF calculation prefixes that match a prefix entry or a tag entry in a prefix policy.</p> <p>The implementation already allows the user to exclude an interface in IS-IS or OSPF, an OSPF area, or an IS-IS level from the LFA SPF.</p> <p>If a prefix is excluded from LFA, then it will not be included in LFA calculation regardless of its priority. The prefix tag will, however, be used in the main SPF. Note that prefix tags are defined for the IS-IS protocol but not for the OSPF protocol.</p> <p>The default action of the <b>loopfree-alternate-exclude</b> command, when not explicitly specified by the user in the prefix policy, is a “reject”. Thus, regardless if the user did or did not explicitly add the statement “default-action reject” to the prefix policy, a prefix that did not match any entry in the policy will be accepted into LFA SPF.</p> <p>The <b>no</b> form deletes the exclude prefix policy.</p>
<b>Parameters</b>	<b>prefix-policy <i>prefix-policy</i></b> — Specifies the name of the prefix policy, up to 32 characters. The specified name must have been already defined.

## overload-include-ext-2

<b>Syntax</b>	<b>[no] overload-include-ext-2</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	This command is used to control if external type-2 routes should be re-advertised with a maximum metric value when the system goes into overload state for any reason. When this command is enabled and the router is in overload, all external type-2 routes will be advertised with the maximum metric.
<b>Default</b>	<b>no overload-include-ext-2</b>

## overload-include-stub

<b>Syntax</b>	<b>[no] overload-include-stub</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	This command is used to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.
<b>Default</b>	no overload-include-stub

## overload-on-boot

<b>Syntax</b>	<b>overload-on-boot [timeout seconds]</b> <b>no overload</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur: <ul style="list-style-type: none"> <li>• The timeout timer expires.</li> <li>• A manual override of the current overload state is entered with the <b>no overload</b> command.</li> </ul> <p>The <b>no overload</b> command does not affect the <b>overload-on-boot</b> function.</p> <p>The <b>no</b> form of the command removes the overload-on-boot functionality from the configuration.</p>
<b>Default</b>	no overload-on-boot
<b>Parameters</b>	<b>timeout seconds</b> — Specifies the number of seconds to reset overloading.

**Values** 60 —1800  
**Default** 60

preference

**Syntax** **preference** *preference*  
**no preference**

**Context** config>service>vprn>ospf  
 config>service>vprn>ospf3

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols in which case the costs are not comparable, when this occurs the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

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

**Default** **preference 10** — OSPF internal routes have a preference of 10.

**Parameters** *preference* — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in the following table.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes <sup>a</sup>
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

a. Preference for OSPF internal routes is configured with the **preference** command.



**Values** 1 — 255

## reference-bandwidth

<b>Syntax</b>	<b>reference-bandwidth</b> <i>reference-bandwidth</i> <b>no reference-bandwidth</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.</p> <p>The default interface cost is calculated as follows:</p> $\text{cost} = \text{reference-bandwidth} \div \text{bandwidth}$ <p>The default <i>reference-bandwidth</i> is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as follows:</p> <ul style="list-style-type: none"> <li>• 10 Mbs link default cost of 10000</li> <li>• 100 Mbs link default cost of 1000</li> <li>• 1 Gbps link default cost of 100</li> <li>• 10 Gbps link default cost of 10</li> </ul> <p>The <b>reference-bandwidth</b> command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the <b>metric</b> <i>metric</i> command in the <b>config&gt;router&gt;ospf&gt;area&gt;interface</b> <i>ip-int-name</i> context.</p> <p>The <b>no</b> form of the command reverts the reference-bandwidth to the default value.</p>
<b>Default</b>	<b>reference-bandwidth 100000000</b> — Reference bandwidth of 100 Gbps.
<b>Parameters</b>	<i>reference-bandwidth</i> — The reference bandwidth in kilobits per second expressed as a decimal integer.
	<b>Values</b> 1 — 1000000000

## rib-priority

<b>Syntax</b>	<b>rib-priority</b> { <b>high</b> } <i>prefix-list-name</i> <b>no rib-priority</b>
<b>Context</b>	config>service>vprn>ospf config>service>vprn>ospf3
<b>Description</b>	<p>This command enabled RIB prioritization for the OSPF protocol and specifies the prefix list that will be used to select the specific routes that should be processed through the OSPF route calculation process at a higher priority.</p> <p>The <b>no</b> form of <b>rib-priority</b> command disables RIB prioritization at the associated level.</p>

## OSPF Commands

**Default** no rib-priority

**Parameters** *prefix-list-name* — specifies the prefix list which is used to select the routes that are processed at a higher priority through the route calculation process.

## super-backbone

**Syntax** [no] super-backbone

**Context** config>service>vprn>ospf

**Description** This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

**Default** no super-backbone

## suppress-dn-bit

**Syntax** [no] suppress-dn-bit

**Context** config>service>vprn>ospf  
config>service>vprn>ospf3

**Description** This command specifies whether to suppress the setting of the DN bit for OSPF LSA packets generated by this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets generated by this instance of the OSPF router will not be set. When disabled, this instance of the OSPF router will follow the normal procedure to determine whether to set the DN bit.

**Default** no suppress-dn-bit

## timers

**Syntax** timers

**Context** config>service>vprn>ospf  
config>service>vprn>ospf3

**Description** This command enables the context that allows for the configuration of OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affect CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

**Default** none

## spf-wait

<b>Syntax</b>	<b>spf-wait</b> <i>max-spf-wait</i> [ <i>spf-initial-wait</i> [ <i>spf-second-wait</i> ]] <b>no spf-wait</b>
<b>Context</b>	config>service>vprn>ospf>timers config>service>vprn>ospf3>timers
<b>Description</b>	<p>This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the <i>spf-second-wait</i> interval. For example, if the <i>spf-second-wait</i> interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the <b>spf-wait</b> value. The SPF interval will stay at the <b>spf-wait</b> value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to <i>spf-initial-wait</i>.</p> <p>The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.</p> <p>Use the <b>no</b> form of this command to return to the default.</p>
<b>Default</b>	no spf-wait
<b>Parameters</b>	<p><i>max-spf-wait</i> — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.</p> <p><b>Values</b> 10 — 120000</p> <p><b>Default</b> 10000</p> <p><i>spf-initial-wait</i> — Specifies the initial SPF calculation delay in milliseconds after a topology change.</p> <p><b>Values</b> 10 — 100000</p> <p><b>Default</b> 1000</p> <p><i>spf-second-wait</i> — Specifies the hold time in milliseconds between the first and second SPF calculation.</p> <p><b>Values</b> 10 — 100000</p> <p><b>Default</b> 1000</p>

## unicast-import-disable

<b>Syntax</b>	<b>[no] unicast-import-disable</b>
<b>Context</b>	config>service>vprn>ospf
<b>Description</b>	<p>This command allows one IGP to import its routes into RPF RTM while another IGP imports routes only into the unicast RTM.</p> <p>Import policies can redistribute routes from an IGP protocol into the RPF RTM (the multicast routing table). By default, the IGP routes will not be imported into RPF RTM as such an import policy must be explicitly configured</p>

## OSPF Commands

**Default** no unicast-import-disable

### vpn-domain

**Syntax** **vpn-domain** [*type* {0005 | 0105 | 0205 | 8005}] **id** *id*  
**no vpn-domain**

**Context** config>service>vprn>ospf

**Description** This command specifies type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance. The parameters are mandatory and can be entered in either order. This command is not applicable in the **config>service>vprn>ospf3** context.

This command is not supported in OSPF3.

**Default** no vpn-domain

**Parameters** *id* — Specifies the OSPF VPN domain in the “xxxx.xxxx.xxxx” format. This is exchanged using BGP in the extended community attribute associated with a prefix. This object applies to VPRN instances of OSPF only.

*type* — Specifies the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID.

**Values** 0005, 0105, 0205, 8005

### vpn-tag

**Syntax** **vpn-tag** *vpn-tag*  
**no vpn-tag**

**Context** config>service>vprn>ospf

**Description** This command specifies the route tag for an OSPF VPN on a PE router. This field is set in the tag field of the OSPF external LSAs generated by the PE. This is mainly used to prevent routing loops. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance.

This command is not supported in OSPF3.

**Default** vpn-tag 0

### lsa-arrival

**Syntax** **lsa-arrival** *lsa-arrival-time*  
**no lsa-arrival**

**Context** config>service>vprn>ospf>timers

```
config>service>vprn>ospf3>timers
```

<b>Description</b>	This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors. It is recommended that the neighbors configured ( <b>lsa-generate</b> ) <i>lsa-second-wait</i> interval is equal or greater than the <b>lsa-arrival</b> timer configured here. Use the <b>no</b> form of this command to return to the default.
<b>Default</b>	no lsa-arrival
<b>Parameters</b>	<i>lsa-arrival-time</i> — Specifies the timer in milliseconds. Values entered that do not match this requirement will be rejected.
<b>Values</b>	0 — 600000

## lsa-generate

<b>Syntax</b>	<b>lsa-generate</b> <i>max-lsa-wait</i> [ <i>lsa-initial-wait</i> [ <i>lsa-second-wait</i> ]] <b>no lsa-generate-interval</b>
<b>Context</b>	config>service>vprn>ospf>timers config>service>vprn>ospf3>timers
<b>Description</b>	This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the <i>lsa-second-wait</i> timer until a maximum value is reached. Configuring the <b>lsa-arrival</b> interval to equal or less than the <i>lsa-second-wait</i> interval configured in the <b>lsa-generate</b> command is recommended. Use the <b>no</b> form of this command to return to the default.
<b>Default</b>	no lsa-generate
<b>Parameters</b>	<i>max-lsa-wait</i> — Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated. The timer must be entered as either 1 or in millisecond increments. Values entered that do not match this requirement will be rejected.
<b>Values</b>	1 — 600000

---

## RIP Commands

### rip

<b>Syntax</b>	<b>[no] rip</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the RIP protocol on the given VPRN IP interface. The <b>no</b> form of the command disables the RIP protocol from the given VPRN IP interface.
<b>Default</b>	no rip

### ripng

<b>Syntax</b>	<b>[no] ripng</b>
<b>Context</b>	config>router
<b>Description</b>	This command creates the context to configure the RIPng protocol instance. When a RIPng instance is created, the protocol is enabled by default. To start or suspend execution of the RIP protocol without affecting the configuration, use the <b>[no] shutdown</b> command. The <b>no</b> form of the command deletes the RIP protocol instance removing all associated configuration parameters.
<b>Default</b>	<b>no ripng</b> — No RIPng protocol instance defined.

### authentication-key

<b>Syntax</b>	<b>authentication-key</b> <i>[authentication-key   hash-key]</i> <b>[hash   hash2]</b> <b>no authentication-key</b>
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor
<b>Description</b>	This command sets the authentication password to be passed between RIP neighbors. The authentication type and authentication key must match exactly for the RIP message to be considered authentic and processed. The <b>no</b> form of the command removes the authentication password from the configuration and disables authentication.
<b>Default</b>	<b>no authentication-key</b> — Authentication is disabled and the authentication password is empty.

- Parameters**
- authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“”).
- hash-key* — The hash key. The key can be any combination of ASCII characters up to 33 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.

## authentication-type

- Syntax** **authentication-type {none | password | message-digest}**  
**no authentication-type**
- Context** config>service>vprn>rip  
config>service>vprn>rip>group  
config>service>vprn>rip>group>neighbor
- Description** This command defines the type of authentication to be used between RIP neighbors. The type and password must match exactly for the RIP message to be considered authentic and processed.
- The **no** form of the command removes the authentication type from the configuration and effectively disables authentication.
- Default** no authentication-type
- Parameters** *none* — No authentication is used.
- simple* — A simple clear-text password is sent.
- md5* — MD5 authentication is used.

## check-zero

- Syntax** **check-zero {enable | disable}**  
**no check-zero**
- Context** config>service>vprn>rip  
config>service>vprn>rip>group  
config>service>vprn>rip>group>neighbor  
config>service>vprn>ripng  
config>service>vprn>ripng>group  
config>service>vprn>ripng>group>neighbor

## RIP Commands

<b>Description</b>	<p>This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.</p> <p>The <b>no</b> form of the command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.</p>
<b>Default</b>	no check-zero
<b>Parameters</b>	<p><b>enable</b> — Enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting non-compliant RIP messages.</p> <p><b>disable</b> — Disables the checking and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.</p>

## split-horizon

<b>Syntax</b>	<b>split-horizon {enable   disable}</b> <b>no split-horizon</b>
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
<b>Description</b>	<p>This command enables the use of split-horizon. RIP uses split-horizon with poison-reverse to protect from such problems as “counting to infinity”. Split-horizon with poison reverse means that routes learned from a neighbor through a given interface are advertised in updates out of the same interface but with a metric of 16 (infinity).</p> <p>The <b>split-horizon disable</b> command enables split horizon without poison reverse. This allows the routes to be re-advertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.</p> <p>This configuration parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group) or neighbor level (only applies to the specified neighbor interface). The most specific value is used. In particular if no value is set (<b>no split-horizon</b>), the setting from the less specific level is inherited by the lower level.</p> <p>The <b>no</b> form of the command disables split horizon command which allows the lower level to inherit the setting from an upper level.</p>
<b>Default</b>	enabled

## export

<b>Syntax</b>	<b>export <i>policy</i> [<i>policy</i>...]</b> <b>no export</b>
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group



```
config>service>vprn>rip>group>neighbor
config>service>vprn>ripng
config>service>vprn>ripng>group
config>service>vprn>ripng>group>neighbor
```

- Description** This command specifies the export policies to be used to control routes advertised to RIP neighbors. By default, RIP advertises routes from other RIP routes but does not advertise any routes from other protocols unless directed by an export policy.
- The **no** form of the command removes all route policy names from the export list.
- Default** no export
- Parameters** *policy* — A route policy statement name.

## export-limit

- Syntax** **export-limit** *number* [**log** *percentage*]  
**no export-limit**
- Context** config>service>vprn>rip  
config>service>vprn>ripng
- Description** This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.
- The **no** form of the command removes the parameters from the configuration.
- Default** no export-limit, the export limit for routes or prefixes is disabled..
- Parameters** *number* — Specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table.
- Values** 1 — 4294967295
- log percentage** — Specifies the percentage of the export-limit, at which a warning log message and SNMP notification would be sent.
- Values** 1 — 100

## import

- Syntax** **import** *policy* [*policy*...]  
**no import**
- Context** config>service>vprn>rip  
config>service>vprn>rip>group  
config>service>vprn>rip>group>neighbor  
config>service>vprn>ripng  
config>service>vprn>ripng>group  
config>service>vprn>ripng>group>neighbor

## RIP Commands

<b>Description</b>	This command specifies the import policies to be used to control routes advertised from RIP neighbors. By default, RIP accepts all routes from configured RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.  The <b>no</b> form of the command removes all route policy names from the import list.
<b>Default</b>	no import
<b>Parameters</b>	<i>policy</i> — A route policy statement name.

## message-size

<b>Syntax</b>	<b>message-size</b> <i>max-num-of-routes</i> <b>no message-size</b>				
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor				
	This command sets the maximum number of routes per RIP update message. The <b>no</b> form of the command resets the maximum number of routes back to the default of 25.				
<b>Default</b>	no message-size				
<b>Parameters</b>	<i>size</i> — Integer. <table><tr><td><b>Default</b></td><td>25</td></tr><tr><td><b>Values</b></td><td>25 — 255</td></tr></table>	<b>Default</b>	25	<b>Values</b>	25 — 255
<b>Default</b>	25				
<b>Values</b>	25 — 255				

## metric-in

<b>Syntax</b>	<b>metric-in</b> <i>metric</i> <b>no metric-in</b>
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
<b>Description</b>	This command sets the metric added to routes that were received from a RIP neighbor. The <b>no</b> form of the command reverts the <i>metric</i> value back to the default.
<b>Default</b>	no metric-in

**Parameters** *metric* — The value added to the metric of routes received from a RIP neighbor, expressed as a decimal integer.

**Values** 1 — 16

## metric-out

**Syntax** **metric-out** *metric*  
**no metric-out**

**Context** config>service>vprn>rip  
config>service>vprn>rip>group  
config>service>vprn>rip>group>neighbor  
config>service>vprn>ripng  
config>service>vprn>ripng>group  
config>service>vprn>ripng>group>neighbor

**Description** This command sets the metric added to routes that were exported into RIP and advertised to RIP neighbors.

The **no** form of the command removes the command from the config and resets the metric-in value back to the default.

**Default** no metric-out

**Parameters** *metric* — The value added to the metric for routes exported into RIP and advertised to RIP neighbors, expressed as a decimal integer.

**Values** 1 — 16

## preference

**Syntax** **preference** *preference*  
**no preference**

**Context** config>service>vprn>rip  
config>service>vprn>rip>group  
config>service>vprn>rip>group>neighbor  
config>service>vprn>ripng  
config>service>vprn>ripng>group  
config>service>vprn>ripng>group>neighbor

**Description** This command sets the route preference assigned to RIP routes. This value can be overridden by route policies.

The **no** form of the command resets the *preference* to the default.

**Default** no preference

**Parameters** *preference* — An integer.

**Values** 1 — 255

**Default** 100

## propagate-metric

**Syntax** **[no] propagate-metric**

**Context** config>service>vprn>rip  
config>service>vprn>ripng

**Description** This command allows the RIP metric to be used to set the MP-BGP MED attribute when RIP is used as the CE-PE routing protocols for VPRNs. This is similar to the way the OSPF metric can be used to set the MP-BGP metric when OSPF is used as the CE-PE protocol.

MP-BGP will use the RIP metric to set the MED attribute, this attribute gets flooded through out the MP-BGP peers and will then be used to set the RIP metric at the other end and re-advertise the RIP metric to the far-end RIP neighbors.

## receive

**Syntax** **receive {both | none | version-1 | version-2}**  
**no receive**

**Context** config>service>vprn>rip  
config>service>vprn>rip>group  
config>service>vprn>rip>group>neighbor  
config>service>vprn>ripng  
config>service>vprn>ripng>group  
config>service>vprn>ripng>group>neighbor

**Description** This command configures the type(s) of RIP updates that will be accepted and processed. If **both** or **version-2** is specified, the RIP instance listens for and accepts packets sent to the broadcast and multicast (224.0.0.9) addresses.

If **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

This control can be issued at the global, group or interface level. The default behavior accepts and processes both RIPv1 and RIPv2 messages.

The **no** form of the command resets the type of messages accepted to both.

**Default** **no receive** — Accepts both formats.

**Parameters** **both** — Receive RIP updates in either Version 1 or Version 2 format.

**none** — Do not accept and RIP updates.

**version-1** — Router should only accept RIP updates in Version 1 format.

**version-2** — Router should only accept RIP updates in Version 2 format.

## send

<b>Syntax</b>	<b>send {broadcast   multicast   none   version-1   both}</b> <b>no send</b>
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
<b>Description</b>	This command specifies the type of RIP messages sent to RIP neighbors. This control can be issued at the global, group or interface level. The default behavior sends RIPv2 messages with the multicast (224.0.0.9) destination address.  If <b>version-1</b> is specified, the router only listens for and accepts packets sent to the broadcast address. The <b>no</b> form of this command resets the type of messages sent back to the default value.
<b>Default</b>	<b>no send</b> — Sends RIPv2 to the broadcast address.
<b>Parameters</b>	<b>broadcast</b> — Send RIPv2 formatted messages to the broadcast address. <b>multicast</b> — Send RIPv2 formatted messages to the multicast address. <b>none</b> — Do not send any RIP messages (i.e. silent listener). <b>version-1</b> — Send RIPv1 formatted messages to the broadcast address. <b>both</b> — Send both RIP v1 & RIP v2 updates to the broadcast address.

## timers

<b>Syntax</b>	<b>timers update timeout flush</b> <b>no timers</b>
<b>Context</b>	config>service>vprn>rip config>service>vprn>rip>group config>service>vprn>rip>group>neighbor config>service>vprn>ripng config>service>vprn>ripng>group config>service>vprn>ripng>group>neighbor
<b>Description</b>	This command sets the values for the update, timeout, and flush timers. <ul style="list-style-type: none"> <li>• Update timer — Determines how often RIP updates are sent.</li> <li>• Timeout timer — If a router is not updated by the time the timer expires, the route is declared invalid, but maintained in the RIP database.</li> <li>• Flush timer — Determines how long a route is maintained in the RIP database, after it has been declared invalid. Once this timer expires it is flushed from the RIP database completely.</li> </ul> The <b>no</b> form of the command resets all timers to their default values of 30, 180, and 120 seconds respectively.

## RIP Commands

<b>Default</b>	no timers
<b>Parameters</b>	<i>update</i> — The RIP update timer value in seconds. <b>Values</b> 1 — 600 <b>Default</b> 30 <i>timeout</i> — The RIP timeout timer value in seconds. <b>Values</b> 1 — 1200 <b>Default</b> 180 <i>flush</i> — The RIP flush timer value in seconds. <b>Values</b> 1 — 1200 <b>Default</b> 120

## unicast-address

<b>Syntax</b>	<b>[no] unicast-address</b> <i>ipv6-address</i>
<b>Context</b>	config>service>vprn>ripng>group>neighbor
<b>Description</b>	This command configures the unicast IPv6 address, RIPng updates messages will be sent to if the RIPng <b>send</b> command is set to <b>send unicast</b> . Multiple unicast-address entries can be configured, in which case unicast messages will be sent to each configured unicast IPv6 address. The <b>no</b> form of the command deletes the specified IPv6 unicast address from the configuration.
<b>Default</b>	<i>ipv6-address</i> — IPv6 unicast address to which unicast RIPng updates should be sent.

## group

<b>Syntax</b>	<b>[no] group</b> <i>group-name</i>
<b>Context</b>	config>service>vprn>rip config>service>vprn>ripng
<b>Description</b>	This command creates a context for configuring a RIP group of neighbors. RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces. The <b>no</b> form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.
<b>Default</b>	<b>no group</b> — No group of RIP neighbor interfaces defined
<b>Parameters</b>	<i>group-name</i> — The RIP group name. Allowed values are any string up to 32 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.

## neighbor

<b>Syntax</b>	<b>[no] neighbor</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn>rip>group config>service>vprn>ripng>group
<b>Description</b>	<p>This command creates a context for configuring a RIP neighbor interface. By default, interfaces are not activated in any interior gateway protocol such as RIP unless explicitly configured.</p> <p>The <b>no</b> form of the command deletes the RIP interface configuration for this interface. The <b>shutdown</b> command in the <b>config&gt;router&gt;rip&gt;group group-name&gt;neighbor ip-int-name</b> context can be used to disable an interface without removing the configuration for the interface.</p>
<b>Default</b>	<b>no neighbor</b> — No RIP interfaces defined
<b>Parameters</b>	<p><i>ip-int-name</i> — The IP interface name. Interface names must be unique within the group of defined IP interfaces for <b>config router interface</b> and <b>config service vprn interface</b> commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 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.</p> <p>If the IP interface name does not exist or does not have an IP address configured an error message will be returned.</p>

---

## IPSec Configuration Commands

### ipsec

<b>Syntax</b>	<b>ipsec</b>
<b>Context</b>	config>service>vpn>ipsec
<b>Description</b>	This command enables the context to configure IPSec policies.
<b>Default</b>	none

### security-policy

	<b>security-policy</b> <i>security-policy-id</i> [ <b>create</b> ] <b>no security-policy</b> <i>security-policy-id</i>
<b>Context</b>	config>service>vpn>ipsec
<b>Description</b>	This command configures a security policy to use for an IPSec tunnel.
<b>Default</b>	none
<b>Parameters</b>	<i>security-policy-id</i> — specifies a value to be assigned to a security policy. <b>Values</b> 1 — 8192 <b>create</b> — Keyword used to create the security policy instance. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.

### entry

<b>Syntax</b>	<b>entry</b> <i>entry-id</i> [ <b>create</b> ] <b>no entry</b> <i>entry-id</i>
<b>Context</b>	config>service>vpn>ipsec>sec-plcy
<b>Description</b>	This command configures an IPSec security policy entry.
<b>Parameters</b>	<i>entry-id</i> — Specifies the IPSec security policy entry. <b>Values</b> 1 — 16 <b>create</b> — Keyword used to create the security policy entry instance. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.



## local-ip

<b>Syntax</b>	<b>local-ip</b> { <i>ip-prefix/prefix-length</i>   <i>ip-prefix netmask</i>   <b>any</b> }				
<b>Context</b>	config>service>vpn>ipsec>sec-plcy>entry				
<b>Description</b>	This command configures the local (from the VPN ) IP prefix/mask for the policy parameter entry. Only one entry is necessary to describe a potential flow. The <b>local-ip</b> and <b>remote-ip</b> commands can be defined only once. The system will evaluate the local IP as the source IP when traffic is examined in the direction of VPN to the tunnel and as the destination IP when traffic flows from the tunnel to the VPN. The remote IP will be evaluated as the source IP when traffic flows from the tunnel to the VPN when traffic flows from the VPN to the tunnel.				
<b>Parameters</b>	<i>ip-prefix</i> — The destination address of the aggregate route in dotted decimal notation. <table> <tr> <td><b>Values</b></td> <td>a.b.c.d (host bits must be 0)</td> </tr> <tr> <td>prefix-length</td> <td>1 — 32</td> </tr> </table> <i>netmask</i> — The subnet mask in dotted decimal notation. <b>any</b> — keyword to specify that it can be any address.	<b>Values</b>	a.b.c.d (host bits must be 0)	prefix-length	1 — 32
<b>Values</b>	a.b.c.d (host bits must be 0)				
prefix-length	1 — 32				

## remote-ip

<b>Syntax</b>	<b>remote-ip</b> <i>ip-prefix/prefix-length</i>   <i>ip-prefix netmask</i>   <b>any</b> }				
<b>Context</b>	config>service>vpn>ipsec>sec-plcy>entry				
<b>Description</b>	This command configures the remote (from the tunnel) IP prefix/mask for the policy parameter entry. Only one entry is necessary to describe a potential flow. The <b>local-ip</b> and <b>remote-ip</b> commands can be defined only once. The system will evaluate the local IP as the source IP when traffic is examined in the direction of VPN to the tunnel and as the destination IP when traffic flows from the tunnel to the VPN. The remote IP will be evaluated as the source IP when traffic flows from the tunnel to the VPN when traffic flows from the VPN to the tunnel.				
<b>Parameters</b>	<i>ip-prefix</i> — The destination address of the aggregate route in dotted decimal notation. <table> <tr> <td><b>Values</b></td> <td>a.b.c.d (host bits must be 0)</td> </tr> <tr> <td>prefix-length</td> <td>1 — 32</td> </tr> </table> <i>netmask</i> — The subnet mask in dotted decimal notation. <b>any</b> — keyword to specify that it can be any address.	<b>Values</b>	a.b.c.d (host bits must be 0)	prefix-length	1 — 32
<b>Values</b>	a.b.c.d (host bits must be 0)				
prefix-length	1 — 32				

## ipsec-interface

<b>Syntax</b>	<b>ipsec-interface</b> <i>ip-int-name</i> [create] <b>no ipsec-interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vpn

## IPSec Configuration Commands

<b>Description</b>	This command configures an IPSec interface.
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. <b>Values</b> 1 — 32 characters maximum <b>create</b> — Keyword used to create the IPSec interface instance. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.

### address

<b>Syntax</b>	<b>[no] address</b> { <i>ip-address/mask</i>   <i>ip-address netmasks</i> }
<b>Context</b>	config>service>vprn>ipsec-if
<b>Description</b>	This command assigns an IP address/IP subnet to the interface
<b>Parameters</b>	<i>ip-address</i> — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets). <i>mask</i> — The subnet mask in dotted decimal notation. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252. Note that a mask of 255.255.255.255 is reserved for system IP addresses. <i>netmask</i> — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

### ip-mtu

<b>Syntax</b>	<b>ip-mtu</b> <i>octets</i> <b>no ip-mtu</b>
<b>Context</b>	config>service>vprn>ipsec-if
<b>Description</b>	This command configures the IP maximum transmit unit (packet) for this interface. The <b>no</b> form of the command returns the default value.
<b>Default</b>	no ip-mtu
<b>Parameters</b>	<i>octets</i> — Specifies the MTU size for this interface. <b>Values</b> 512 — 9000

## tunnel

<b>Syntax</b>	<b>tunnel</b> <i>ipsec-tunnel-name</i> [ <b>create</b> ] <b>no tunnel</b> <i>ipsec-tunnel-name</i>
<b>Context</b>	config>service>vprn>ipsec-if>sap
<b>Description</b>	This command specifies an IPsec tunnel name. An IPsec client sets up the encrypted tunnel across public network. The 7750 SR IPsec MDA acts as a concentrator gathering, and terminating these IPsec tunnels into an IES or VPRN service. This mechanism allows as service provider to offer a global VPRN service even if node of the VPRN are on an uncontrolled or insecure portion of the network.
<b>Default</b>	none
<b>Parameters</b>	<i>ipsec-tunnel-name</i> — Specifies an IPsec tunnel name up to 32 characters in length. <b>create</b> — Keyword used to create the IPsec tunnel instance. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context.

## bfd-designate

<b>Syntax</b>	[ <b>no</b> ] <b>bfd-designate</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command specifies whether this IPsec tunnel is the BFD designated tunnel.
<b>Default</b>	none

## bfd-enable

<b>Syntax</b>	[ <b>no</b> ] <b>bfd-enable service</b> <i>service-id</i> <b>interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command assign a BFD session provide heart-beat mechanism for given IPsec tunnel. There can be only one BFD session assigned to any given IPsec tunnel, but there can be multiple IPsec tunnels using same BFD session. BFD control the state of the associated tunnel, if BFD session goes down, system will also bring down the associated non-designated IPsec tunnel.
<b>Default</b>	none
<b>Parameters</b>	<b>service</b> <i>service-id</i> — Specifies where the service-id that the BFD session resides. <b>interface</b> <i>interface-name</i> — Specifies the name of the interface used by the BFD session. <b>dst-ip</b> <i>ip-address</i> — Specifies the destination address to be used for the BFD session.

### clear-df-bit

<b>Syntax</b>	<b>[no] clear-df-bit</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel config>service>interface>vprn>sap>ip-tunnel
<b>Description</b>	This command specifies whether to clear the Do not Fragment (DF) bit in the outgoing packets in this tunnel.

### dynamic-keying

<b>Syntax</b>	<b>[no] dynamic-keying</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command enables dynamic keying for the IPsec tunnel.
<b>Default</b>	none

### auto-establish

<b>Syntax</b>	<b>[no] auto-establish</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command specifies whether to attempt to establish a phase 1 exchange automatically. The <b>no</b> form of the command disables the automatic attempts to establish a phase 1 exchange.
<b>Default</b>	no auto-establish

### local-id

<b>Syntax</b>	<b>[no] local-id type {ipv4 &lt;v4address&gt;   fqdn &lt;fqdn-value&gt;}</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel>dynamic-keying
<b>Description</b>	This command specifies the local id of the 7750 SR used for IDi or IDr for IKEv2 tunnels. The local-id command can only be changed or removed when tunnel or gw is shutdown. The default value depends on the local-auth-method such as: <ul style="list-style-type: none"><li>- Psk:local tunnel ip address</li><li>- Cert-auth: subject of the local certificate</li></ul>
<b>Default</b>	no local-id
<b>Parameters</b>	<b>type</b> — Specifies the type of local ID payload, it could be ipv4 address. <b>ipv4</b> — Specifies IPv4 as the local ID type. The default value is the local tunnel end-point address

*v4address* — Specifies an IPv4 address. A value must be configured.

**fqdn** — Specifies FQDN as the local ID type. A value must be configured.

*fqdn-value* — Specifies a FQDN vaue. A value must be configured.

## transform

<b>Syntax</b>	<b>transform</b> <i>transform-id</i> [ <i>transform-id</i> ...(up to 4 max)] <b>no transform</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel>dynamic-keying
<b>Description</b>	This command associates the IPSec transform sets allowed for this tunnel. A maximum of four transforms can be specified. The transforms are listed in decreasing order of preference (the first one specified is the most preferred).
<b>Default</b>	none
<b>Parameters</b>	<i>transform-id</i> — Specifies the value used for transforms for dynamic keying. <b>Values</b> 1 — 2048

## local-gateway-address

<b>Syntax</b>	<b>local-gateway-address</b> <i>ip-address</i> <b>peer</b> <i>ip-address</i> <b>delivery-service</b> <i>service-id</i> <b>no local-gateway-address</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command specifies the local gateway address used for the tunnel and the address of the remote security gateway at the other end of the tunnelremote peer IP address to use.
<b>Default</b>	The base routing context is used if the delivery-router option is not specified.
<b>Parameters</b>	<i>ip-address</i> — IP address of the local end of the tunnel. <b>delivery-service</b> <i>service-id</i> — The ID of the IES or VPRN (front-door) delivery service of this tunnel. Use this service-id to find the VPRN used for delivery. <b>Values</b> <i>service-id</i> : 1 — 2147483648 <i>svc-name</i> : Specifies an existing service name up to 64 characters in length.

## manual-keying

<b>Syntax</b>	<b>[no] manual-keying</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel

## IPSec Configuration Commands

<b>Description</b>	This command configures Security Association (SA) for manual keying. When enabled, the command specifies whether this SA entry is created manually by the user or dynamically by the IPsec sub-system.
<b>Default</b>	none

### security-association

<b>Syntax</b>	<b>security-association</b> <i>security-entry-id</i> <b>authentication-key</b> <i>authentication-key</i> <b>encryption-key</b> <i>encryption-key</i> <b>spi</b> <i>spi</i> <b>transform</b> <i>transform-id</i> <b>direction</b> {inbound   outbound} <b>no security-association</b> <i>security-entry-id</i> <b>direction</b> {inbound   outbound}
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel>manual-keying
<b>Description</b>	This command configures the information required for manual keying SA creation.
<b>Default</b>	none
<b>Parameters</b>	<i>security-entry-id</i> — Specifies the ID of an SA entry. <b>Values</b> 1 — 16 <b>encryption-key</b> <i>encryption-key</i> — specifies the key used for the encryption algorithm. <b>Values</b> none or 0x0..0xFFFFFFFF...(max 64 hex nibbles) <b>authentication-key</b> <i>authentication-key</i> — <b>Values</b> none or 0x0..0xFFFFFFFF...(max 40 hex nibbles) <b>spi</b> <i>spi</i> — Specifies the SPI (Security Parameter Index) used to look up the instruction to verify and decrypt the incoming IPSec packets when the direction is inbound. When the direction is outbound, the SPI that will be used in the encoding of the outgoing packets. The remote node can use this SPI to lookup the instruction to verify and decrypt the packet. <b>Values</b> 256 — 16383 <b>transform</b> <i>transform-id</i> — specifies the transform entry that will be used by this SA entry. This object should be specified for all the entries created which are manual SAs. If the value is dynamic, then this value is irrelevant and will be zero. <b>Values</b> 1 — 2048 <b>direction</b> {inbound   outbound} — Specifies the direction of an IPSec tunnel.

### replay-window

<b>Syntax</b>	<b>replay-window</b> {32   64   128   256   512} <b>no replay-window</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel

<b>Description</b>	This command specifies the size of the anti-replay window. The anti-replay window protocol secures IP against an entity that can inject messages in a message stream from a source to a destination computer on the Internet.
<b>Default</b>	none
<b>Parameters</b>	{32   64   128   256   512} — Specifies the size of the SA anti-replay window.

## security-policy

<b>Syntax</b>	<b>security-policy</b> <i>security-policy-id</i> <b>no security-policy</b>
<b>Context</b>	config>service>vprn>ipsec-if>sap>tunnel
<b>Description</b>	This command configures an IPSec security policy. The policy may then be associated with tunnels defined in the same context.
<b>Default</b>	none
<b>Parameters</b>	<i>security-policy-id</i> — Specifies the IPSec security policy entry that the tunnel will use. <b>Values</b> 1 — 8192

---

## Threat Management Service Interface Commands

### tms-interface

<b>Syntax</b>	<b>tms-interface</b> <i>interface-name</i> [ <b>create</b> ] [ <b>off-ramp-vprn</b> <i>off-ramp-svc</i> ] [ <b>mgmt-vprn</b> <i>mgmt-svc</i> ] <b>no tms-interface</b> <i>interface-name</i>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command configure a Threat Management Service interface. The <b>no</b> form of the command removes the interface name from the configuration.
<b>Parameters</b>	<i>interface-name</i> — Specifies the interface name up to 22 characters in length. <b>create</b> — Keyword used to create the interface name. The <b>create</b> keyword requirement can be enabled/disabled in the <b>environment&gt;create</b> context. <i>off-ramp-svc</i> — Specifies the service ID or service name for the off-ramp VPRN. <b>Values</b> <i>service-id</i> : 1 — 2148007978 <i>svc-name</i> : An existing service name up to 64 characters in length. <i>mgmt-svc</i> — Specifies the service ID or service name for management VPRN. <i>service-id</i> : 1 — 2148007978 <i>svc-name</i> : An existing service name up to 64 characters in length.

### address

<b>Syntax</b>	<b>address</b> { <i>ip-address/mask</i>   <i>ip-address netmask</i> } <b>no address</b>
<b>Context</b>	config>service>vprn>tms-if
<b>Description</b>	This command assigns an IP address/IP subnet/broadcast address to the TMS instance for communications between Arbor CP collectors/managers and the TMS instance operating within the Service Router. The <b>no</b> form of the command removes the IP address information from the interface configuration.
<b>Parameters</b>	<i>ip-address/mask ip-address netmask</i> — Specifies IP address information. <b>Values</b> < <i>ip-address</i> [/ <i>mask</i> ]> <i>ip-address</i> a.b.c.d <i>mask</i> 32 < <i>netmask</i> >           a.b.c.d (all 1 bits)



## description

<b>Syntax</b>	<b>description</b> <i>long-description-string</i> <b>no description</b>
<b>Context</b>	config>service>vprn>tms-if
<b>Description</b>	This command configures a description for the interface. The <b>no</b> form of the command removes the description from the interface configuration.

## ipv6

<b>Syntax</b>	<b>[no] ipv6</b>
<b>Context</b>	config>service>vprn>tms-if
<b>Description</b>	This command configures IPv6 for a threat-management service interface. The <b>no</b> form of the command removes the IP address information from the interface configuration.

## password

<b>Syntax</b>	<b>password</b> [ <i>password</i> ] <b>no password</b>
<b>Context</b>	config>service>vprn>tms-if
<b>Description</b>	This command configures a password for the user. The <b>no</b> form of the command removes the password.
<b>Parameters</b>	<i>password</i> — Specifies the password for the TMS configuration.
<b>Values</b>	<password>key1<delim>value1 key2<delim>value2 ... <delim> is one of the following: '=' value is unencrypted and remain unencrypted '!' value is unencrypted and to be encrypted '%' value is encrypted and remain encrypted

## port

<b>Syntax</b>	<b>port</b> <i>mda-id</i> <b>no port</b>
<b>Context</b>	config>service>vprn>tms-if
<b>Description</b>	This command specifies a chassis slot and MDA to bind the interface to a physical port. The no form of the command removes the MDA ID from the interface configuration.

## IPSec Configuration Commands

**Parameters**    *mda-id* — Specifies the chassis slot and MDA.

<b>Values</b>	<slot>/<mda>	slot	[1..10]
		mda	[1..2]

---

## RADIUS Proxy Commands

### radius-proxy

<b>Syntax</b>	<b>radius-proxy</b>
<b>Context</b>	config>service>vprn
<b>Description</b>	This command enables the context to configure RADIUS proxy commands.

### server

<b>Syntax</b>	<b>server</b> <i>server-name</i> [ <b>create</b> ] [ <b>purpose</b> {[ <b>accounting</b> ][ <b>authentication</b> ]}] <b>no server</b> <i>server-name</i>
<b>Context</b>	config>service>vprn>radius-proxy
<b>Description</b>	This command configures the name of this RADIUS proxy server.
<b>Parameters</b>	<b>purpose accounting</b> — Specifies that this RADIUS proxy server will be used for accounting purposes. <b>purpose authentication</b> — Specifies that this RADIUS proxy server will be used for authentication purposes.

### cache

<b>Syntax</b>	<b>cache</b>
<b>Context</b>	config>service>vprn>radius-proxy>server
<b>Description</b>	This command enables the context to configure caching parameters. The <b>no</b> form of the command disables caching.

### key

<b>Syntax</b>	<b>key</b> <i>packet-type</i> { <b>accept</b>   <b>request</b> } <b>attribute-type</b> <i>attribute-type</i> [ <b>vendor-id</b> <i>vendor-id</i> ] <b>no key</b>
<b>Context</b>	config>service>vprn>radius-proxy>server>cache
<b>Description</b>	This command configures cache key parameters. The <b>no</b> form of the command removes the parameters from the configuration.
<b>Default</b>	no key

**Parameters** *packet-type* — Specifies the packet type of the RADIUS messages to use to generate the key for the cache of this RADIUS proxy server.

In order to generate the key associated with a RADIUS Access-Accept message, the system uses the attribute of the type specified by the value of `tmnxRadProxSrvCacheKeyAttrType`, within the associated RADIUS message of the type specified by the value of `tmnxRadProxSrvCacheKeyPktType`.

**Values** accept, request

**attribute-type** *attribute-type* — Specifies the RADIUS attribute type to cache for this RADIUS Proxy server.

In order to generate the key associated with a RADIUS Access-Accept message, the system uses the attribute of the type specified by the value of `tmnxRadProxSrvCacheKeyAttrType`, within the associated RADIUS message of the type specified by the value of `tmnxRadProxSrvCacheKeyPktType`.

**Values** 1 — 255

**vendor-id** *vendor-id* — Specifies the RADIUS Vendor-Id.

If the value of `tmnxRadProxSrvCacheKeyVendorId` is equal to zero, the attribute type specified by `tmnxRadProxSrvCacheKeyAttrType` must be used if it appears outside of a Vendor-Specific attribute.

If the value of `tmnxRadProxSrvCacheKeyVendorId` is not equal to zero, the attribute type specified by `tmnxRadProxSrvCacheKeyAttrType` must be used if it appears as a sub-attribute within a Vendor-Specific attribute with Vendor-Id equal to the value of `tmnxRadProxSrvCacheKeyVendorId`.

**Values** 1 — 16777215, alu

## timeout

**Syntax** `timeout [hrs hours] [min minutes] [sec seconds]`  
**no timeout**

**Context** `config>service>vprn>radius-proxy>server>cache`

**Description** This command configures the timeout, in seconds, after which an entry in the cache will expire. The **no** form of the command reverts to the default.

**Default** 300

**Parameters** *timeout* — Configures the timeout.

<b>Values</b>	hours	1..1
	minutes	1 — 59
	seconds	1 — 59

## track-accounting

**Syntax** **track-accounting** [start] [stop] [interim-update]  
**no track-accounting**

**Context** config>service>vprn>radius-proxy>server>cache

**Description** This command specifies which RADIUS accounting packets have impact on the cache of this RADIUS proxy server. Use it to configure what RADIUS accounting packets have impact on the cache.

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

**Default** no track-accounting

**Default** none

## default-accounting-server-policy

**Syntax** **default-accounting-server-policy** *policy-name*  
**no default-accounting-server-policy**

**Context** config>service>vprn>radius-proxy>server

**Description** This command configures the name of the default RADIUS server policy associated with this RADIUS proxy server for accounting purposes.

This default policy is used if no policy can be derived from the user name.

The **no** form of the command removes the policy from the configuration.

**Parameters** *policy-name* — Specifies the default accounting RADIUS server policy up to 32 characters in length.

## default-authentication-server-policy

**Syntax** **default-authentication-server-policy** *policy-name*  
**no default-authentication-server-policy**

**Context** config>service>vprn>radius-proxy>server

**Description** This command configures the name of the default RADIUS server policy associated with this RADIUS proxy server for authentication purposes.

This default policy is used if no policy can be derived from the user name.

The **no** form of the command removes the policy from the configuration.

**Parameters** *policy-name* — Specifies the default authentication RADIUS server policy up to 32 characters in length.

## interface

<b>Syntax</b>	<b>[no] interface</b> <i>ip-int-name</i>
<b>Context</b>	config>service>vprn>radius-proxy>server
<b>Description</b>	This command associates an interface to the proxy server. The no form of the command removes the interface name from the proxy server configuration.
<b>Default</b>	none
<b>Parameters</b>	<i>ip-int-name</i> — Specifies the name of the IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## load-balance-key

<b>Syntax</b>	<b>load-balance-key vendor</b> <i>vendor-id</i> [ <i>vendor-id...</i> (up to 5 max)] <b>attribute-type</b> <i>attribute-type</i> [ <i>attribute-type...</i> (up to 5 max)] <b>load-balance-key source-ip-udp</b> <b>no load-balance-key</b>
<b>Context</b>	config>service>vprn>radius-proxy>server
<b>Description</b>	This command configures how to construct the key for load-balancing RADIUS messages between RADIUS servers.
<b>Default</b>	load-balance-key
<b>Parameters</b>	<b>vendor</b> <i>vendor-id</i> — Specifies the RADIUS Vendor-Id. <b>Values</b> 0 — 16777215 <b>attribute-type</b> <i>attribute-type</i> — Specifies a RADIUS attribute that must be used to construct the key for load-balancing RADIUS messages between RADIUS servers. <b>Values</b> 1 — 255

## python-policy

<b>Syntax</b>	<b>python-policy</b> <i>name</i> <b>no python-policy</b>
<b>Context</b>	config>service>vprn>radius-proxy>server config>service>vprn>radius-server>server
<b>Description</b>	This command specifies a python policy. Python policies are configured in the <b>config&gt;python&gt;python-policy</b> <i>name</i> context.
<b>Parameters</b>	<i>name</i> — Specifies the name of an existing python script up to 32 characters in length.

## secret

<b>Syntax</b>	<b>secret</b> <i>secret</i> [ <b>hash</b>   <b>hash2</b> ] <b>no secret</b>
<b>Context</b>	config>service>vprn>radius-proxy>server
<b>Description</b>	This command configures the secret key associated with the RADIUS server. The no form of the command removes the key from the configuration.
<b>Parameters</b>	<i>secret</i> — The secret key (password) to access the RADIUS server. This secret key must match the password on the RADIUS server. <b>Values</b> Up to 32 characters in length. <b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> 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 <b>hash</b> parameter specified. <b>hash2</b> — Specifies the key is entered in a more complex encrypted form. If the <b>hash2</b> parameter is not used, the less encrypted <b>hash</b> form is assumed.

## send-accounting-response

<b>Syntax</b>	<b>[no] send-accounting-response</b>
<b>Context</b>	config>service>vprn>radius-proxy>server
<b>Description</b>	This command specifies if this RADIUS proxy server itself responds with an Accounting-Response message to each received Accounting-Request instead of proxying them to a configured RADIUS server. The <b>no</b> form of the command disables the accounting response messages.
<b>Default</b>	disabled

## username

<b>Syntax</b>	<b>username</b> [1..32] <b>prefix-string</b> [128 chars max] [ <b>accounting-server-policy</b> <i>policy-name</i> ] [ <b>authentication-server-policy</b> <i>policy-name</i> ] <b>no username</b> [1..32]
<b>Context</b>	config>service>vprn>radius-proxy>server
<b>Description</b>	This command configures username-to-RADIUS-server-policy associations. The <b>no</b> form of the command removes the associations from the configuration.
<b>Default</b>	none
<b>Parameters</b>	<b>username</b> — Specifies the user name. <b>Values</b> 1 — 32

**prefix-string** — Specifies the prefix-string for the association.

**Values** 128 characters, maximum

**accounting-server-policy** *policy-name* — Specifies the default accounting RADIUS server policy up to 32 characters in length.

**authentication-server-policy** *policy-name* — Specifies the default authentication RADIUS server policy up to 32 characters in length.

### radius-server

**Syntax** **radius-server**

**Context** config>service>vprn

**Description** This command enables the context to configure RADIUS server parameters.

### server

**Syntax** **server** *server-name* [**address** *ip-address*] [**secret** *key*] [**hash|hash2**] [**port** *port*] [**create**]  
**no server** *server-name*

**Context** config>service>vprn>radius-server

**Description** This command configures RADIUS server parameters.  
The **no** form of the command removes the parameters from the configuration.

**Parameters** *server-name* — Specifies the name of this RADIUS server.  
**address** *ip-address* — Specifies the IP address of the RADIUS server.  
**secret** *key* — Specifies the secret key (password) to access the RADIUS server. This secret key must match the password on the RADIUS server.  
**Values** Up to 32 characters in length.  
**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.  
**port** *port* — Specifies the UDP port number on which to contact the RADIUS server.

### accept-coa

**Syntax** [**no**] **accept-coa**

**Context** config>service>vprn>radius-server



**Description** This command specifies if this RADIUS server is allowed to process Change of Authorization messages.

## coa-script-policy

**Syntax** **coa-script-policy** *script-policy-name*  
**no coa-script-policy**

**Context** config>service>vprn>radius-server

**Description** This command specifies the RADIUS script policy used to change the RADIUS attributes of the Change-of-Authorization messages.  
 The **no** form of the command removes the script policy from the configuration.

**Default** none

**Parameters** *script-policy-name* — Specifies a Python script policy to modify Change-of-Authorization messages.

## pending-requests-limit

**Syntax** **pending-requests-limit** *limit*  
**no pending-requests-limit**

**Context** config>router  
 config>service>vprn>radius-server>server

**Description** This command specifies the limit of the number of pending RADIUS authentication requests.

**Default** 4096

**Parameters** *limit* — Configure the limit of the number of pending RADIUS requests.

**Values** 1 — 4096

## wpp

**Syntax** [**no**] **wpp**

**Context** config>router  
 config>service>vprn

**Description** This command enters the configuration context of web portal protocol (WPP) under router or vprn.  
 The **no** form of this command removes configuration under wpp.

**Default** no

### portals

<b>Syntax</b>	<b>portals</b>
<b>Context</b>	config>router>wpp config>service>vprn>wpp
<b>Description</b>	This command enters the configuration context of web portal server.

### portal

<b>Syntax</b>	<b>portal</b> <i>name</i> <b>address</b> <i>ip-address</i> [ <b>create</b> ] <b>portal</b> <i>name</i> <b>no portal</b> <i>name</i>
<b>Context</b>	config>router>wpp>portals config>service>vprn>wpp>portals
<b>Description</b>	This command either creates a new web portal server or enters an existing web portal server.
<b>Default</b>	no
<b>Parameters</b>	<i>name</i> — Specifies the name of the web portal server. <i>ip-address</i> — Specifies IPv4 address of the web portal server.

### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>router>wpp>portals>portal config>service>vprn>wpp>portals>portal
<b>Description</b>	This command cause system stops receiving web portal protocol packet from the web portal server.
<b>Default</b>	shutdown

### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>router>wpp config>service>vprn>wpp
<b>Description</b>	This command cause system stops receiving web portal protocol packet from all web portal servers defined in the routing instance
<b>Default</b>	shutdown

## wpp

<b>Syntax</b>	<b>[no] wpp</b>
<b>Context</b>	config>service>ies>sub-if>grp-if> config>service>vprn>sub-if>grp-if>
<b>Description</b>	This command enters the configuration context of web portal protocol (WPP) under group-interface. The <b>no</b> form of this command removes configuration under WPP.
<b>Default</b>	no

## initial-app-profile

<b>Syntax</b>	<b>initial-app-profile</b> <i>profile-name</i> <b>no initial-app-profile</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command specifies the initial app-profile for the hosts created on the group-interface. This initial app-profile is replaced after hosts pass the web portal authentication.
<b>Default</b>	no
<b>Parameters</b>	<i>profile-name</i> — Specifies the name of app-profile.

## initial-sla-profile

<b>Syntax</b>	<b>initial-sla-profile</b> <i>profile-name</i> <b>no initial-sla-profile</b>
<b>Context</b>	config>router>wpp config>service>vprn>wpp
<b>Description</b>	This command specifies the initial sla-profile for the hosts created on the group-interface. This initial sla-profile is replaced after hosts pass the web portal authentication.
<b>Default</b>	no
<b>Parameters</b>	<i>profile-name</i> — Specifies the name of sla-profile.

## initial-sub-profile

<b>Syntax</b>	<b>initial-sub-profile</b> <i>profile-name</i> <b>no initial-sub-profile</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp

## IPSec Configuration Commands

<b>Description</b>	This command specifies the initial sub-profile for the hosts created on the group-interface. This initial sub-profile will be replaced after hosts pass web portal authentication.
<b>Default</b>	no
<b>Parameters</b>	<i>profile-name</i> — Specifies the name of sub-profile.

### portal

<b>Syntax</b>	<b>portal router <i>router-instance</i> name <i>wpp-portal-name</i></b> <b>no portal</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command specifies the web portal server that system talks to for the hosts on the group-interface.
<b>Default</b>	no
<b>Parameters</b>	<i>router-instance</i> — Specifies the routing-instance that web portal server is defined. <i>profile-name</i> — Specifies the name of the web portal server.

### restore-disconnected

<b>Syntax</b>	<b>[no] restore-disconnected</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command enable the behavior that system will restore the initial-sla-profile/initial-sub-profile/initial-aa-prfofile when hosts disconnects instead of removing them.
<b>Default</b>	restore-disconnected

### shutdown

<b>Syntax</b>	<b>[no] shutdown</b>
<b>Context</b>	config>service>ies>sub-if>grp-if>wpp config>service>vprn>sub-if>grp-if>wpp
<b>Description</b>	This command disables web port protocol for the group-interface.
<b>Default</b>	shutdown

|