Configuration Commands

Generic Commands

shutdown

Syntax **[no] shutdown**

Context config>router>interface

Description The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. 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.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

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

Default no shutdown

description

Syntax **description** *description-string*
no description

Context config>router>if
config>router>if>dhcp
config>router>if>vrrp
config>router>l2tp>group
config>router>l2tp>group>tunnel

Description This command creates a text description stored in the configuration file for a configuration context.

The **no** form of the command removes the description string from the context.

Default No description is associated with the configuration context.

Parameters *description-string* — 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.

Router Global Commands

router

Syntax **router** *router-name*

Context config

Description This command enables the context to configure router parameters, and interfaces, route policies, and protocols.

Parameters *router-name* — Specify the router-name.

Values router-name: Base, management

Default Base

aggregate

Syntax **aggregate** *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*] [**black-hole**] [**community** *comm-id*] [**description** *description*]

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*] [**community** *comm-id*] [**indirect** *ip-address*] [**description** *description*]

no aggregate *ip-prefix/ip-prefix-length*

Context config>router

Description This command creates an aggregate route.

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.

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.

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.

A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.

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.

The **no** form of the command removes the aggregate.

Default No aggregate routes are defined.

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

Values			
	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: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	

The mask associated with the network address expressed as a mask length.

Values 0 — 32

summary-only — This optional parameter suppresses advertisement of more specific component routes for the aggregate.

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.

Values		
	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.

Values		
	ipv4-prefix	a.b.c.d
	ipv6-prefix	x:x: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.

autonomous-system

Syntax **autonomous-system** *autonomous-system*
no autonomous-system

Context config>router

Description This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number with an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself.

If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling/enabling (**shutdown**/ **no shutdown**) the BGP instance or rebooting the system with the new configuration.

Default No autonomous system number is defined.

Parameters *autonomous-system* — The autonomous system number expressed as a decimal integer.

Values 1 — 4294967295

confederation

Syntax **confederation** *confed-as-num* **members** *as-number* [*as-number...*up to 15 max]
no confederation [*confed-as-num* **members** *as-number...*up to 15 max]

Context config>router

Description This command creates confederation autonomous systems within an AS.

This technique is used to reduce the number of IBGP sessions required within an AS. Route reflection is another technique that is commonly deployed to reduce the number of IBGP sessions.

The **no** form of the command deletes the specified member AS from the confederation.

When no members are specified in the **no** statement, the entire list is removed and **confederation** is disabled.

When the last member of the list is removed, **confederation** is disabled.

Default no confederation - no confederations are defined.

Parameters *confed-as-num* — The confederation AS number expressed as a decimal integer.

Values 1 — 65535

members *member-as-num* **—** The AS number(s) of members that are part of the confederation, expressed as a decimal integer. Up to 15 members per *confed-as-num* can be configured.

Values 1 — 65535

ecmp

Syntax **ecmp** *max-ecmp-routes*
no ecmp

Context config>router

Description This command enables ECMP and configures the number of routes for path sharing; for example, the value 2 means two equal cost routes will be used for cost sharing.

ECMP can only be used for routes learned with the same preference and same protocol. See the discussion on preferences in the **static-route** command.

When more ECMP routes are available at the best preference than configured in *max-ecmp-routes,* then the lowest next-hop IP address algorithm is used to select the number of routes configured in *max-ecmp-routes*.

The **no** form of the command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, then the route with the lowest next-hop IP address is used.

Default no ecmp

Parameters *max-ecmp-routes* — The maximum number of equal cost routes allowed on this routing table instance, expressed as a decimal integer. Setting ECMP *max-ecmp-routes* to 1 yields the same result as entering **no ecmp**.

Values 0 — 32

weighted-ecmp

Syntax **weighted-ecmp**
no ecmp

Context config>router

Description This command enables the weighted load-balancing, or weighted ECMP, over MPLS LSP.

When this command is enabled, packets of IGP, BGP, and static route prefixes resolved to a set of ECMP tunnel next-hops are sprayed proportionally to the weights configured for each MPLS LSP in the ECMP set.

Weighted load-balancing over MPLS LSP is supported in the following forwarding contexts:

- IGP prefix resolved to IGP shortcuts in RTM (**rsvp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance).
- BGP prefix with the BGP next-hop resolved to IGP shortcuts in RTM (**rsvp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance).
- Static route prefix resolved to an indirect next-hop which itself is resolved to a set of equal-metric MPLS LSPs in TTM. The user can allow automatic selection or specify the names of the equal-metric MPLS LSPs in TTM to be used in the ECMP set.
- Static route prefix resolved to an indirect next-hop which itself is resolved to IGP shortcuts in RTM.

- BGP prefix with a BGP next-hop resolved to a static route which itself resolves to set of tunnel next-hops towards an indirect next-hop in RTM or TTM.
- BGP prefix resolving to another BGP prefix which next-hop is resolved to set of ECMP tunnel next-hops with a static route in RTM or TTM or to IGP shortcuts in RTM.

IGP computes the normalized weight for each prefix tunnel next-hop. IGP updates the route in RTM with the set of tunnel next-hops and normalized weights. RTM downloads the information to IOM for inclusion in the FIB.

If one or more LSPs in the ECMP set of a prefix do not have a weight configured, the regular ECMP spraying for the prefix will be performed.

The weight assigned to an LSP impacts only the forwarding decision, not the routing decision. In other words, it does not change the selection of the set of ECMP tunnel next-hops of a prefix when more next-hops exist than the value of the router **ecmp** option. Once the set of tunnel next-hops is selected, the LSP weight is used to modulate the amount of packets forwarded over each next-hop. It also does not change the hash routine, but only the spraying of the flows over the tunnel next-hops is modified to reflect the normalized weight of each tunnel next-hop.

The no version of the command resumes regular ECMP spraying of packets of IGP, BGP, and static route prefixes over MPLS LSP.

fib-priority

Syntax **fib-priority** {**high** | **standard**}

Context config>router

Description This command specifies the FIB priority for VPRN.

icmp-tunneling

Syntax **icmp-tunneling**
no icmp-tunneling

Context config>router

Description This command enables the tunneling of ICMP reply packets over MPLS LSP at a LSR node as per RFC 3032.

The LSR part of this feature consists of crafting the reply ICMP packet of type=11- 'time exceeded', with a source address set to a local address of the LSR node, and appending the IP header and leading payload octets of the original datagram. The system skips the lookup of the source address of the sender of the label TTL expiry packet, which becomes the destination address of the ICMP reply packet. Instead, CPM injects the ICMP reply packet in the forward direction of the MPLS LSP the label TTL expiry packet was received from. The TTL of pushed labels should be set to 255.

The source address of the ICMP reply packet is determined as follows. The LSR uses the address of the outgoing interface for the MPLS LSP. Note that with LDP LSP or BGP LSP multiple ECMP next-hops can exist and in such a case the first outgoing interface is selected. If that interface does not have an address of the same family (IPv4 or IPv6) as the ICMP packet, then the system address of the same family is selected. If one is not configured, the packet is dropped.

When the packet is received by the egress LER, it performs a regular user packet lookup in the data path in the GRT context for BGP shortcut, 6PE, and BGP label route prefixes, or in VPRN context for VPRN and 6VPE prefixes. It then forwards it to the destination, which is the sender of the original packet which TTL expired at the LSR.

If the egress LER does not have a route to the destination of the ICMP packet, it drops the packets.

The rate of the tunneled ICMP replies at the LSR can be directly or indirectly controlled by the existing IOM level and CPM levels mechanisms. Specifically, the rate of the incoming UDP traceroute packets received with a label stack can be controlled at ingress IOM using the distributed CPU protection feature. The rate of the ICMP replies by CPM can also be directly controlled by configuring a system wide rate limit for packets ICMP replies to MPLS expired packets which are successfully forwarded to CPM using the command 'configure system security vprn-network-exceptions'. Note that while this command's name refers to VPRN service, this feature rate limits ICMP replies for packets received with any label stack, including VPRN and shortcuts.

The 7x50 implementation supports appending to the ICMP reply of type Time Exceeded the MPLS label stack object defined in RFC 4950. It does not include it in the ICMP reply type of Destination unreachable.

The new MPLS Label Stack object permits an LSR to include label stack information including label value, EXP, and TTL field values, from the encapsulation header of the packet that expired at the LSR node. The ICMP message continues to include the IP header and leading payload octets of the original datagram.

In order to include the MPLS Label Stack object, the SROS implementation adds support of RFC 4884 which defines extensions for a multi-part ICMPv4/v6 message of type Time Exceeded.

The **no** form of command disables the tunneling of ICMP reply packets over MPLS LSP at a LSR node.

Default no icmp-tunneling

ignore-icmp-redirect

Syntax [no] **ignore-icmp-redirect**

Context config>router

Description This command drops ICMP redirects received on the management interface.

The no form of the command accepts ICMP redirects received on the management interface.

ip-fast-reroute

Syntax [no] **ip-fast-reroute**

Context config>router

Description This command enables IP Fast-Reroute (FRR) feature on the system.

This feature provides for the use of a Loop-Free Alternate (LFA) backup next-hop for forwarding in-transit and CPM generated IP packets when the primary next-hop is not available. IP FRR is supported on IPv4 and IPv6 OSPF/IS-IS prefixes forwarded in the base router instance to a network IP interface or to an IES SAP interface or spoke interface. It is also supported for VPRN VPN-IPv4 OSPF prefixes and VPN-IPv6 OSPF prefixes forwarded to a VPRN SAP interface or spoke interface.

IP FRR also provides a LFA backup next-hop for the destination prefix of a GRE tunnel used in an SDP or in VPRN auto-bind.

When any of the following events occurs, IGP instructs in the fast path on the IOMs to enable the LFA backup next-hop:

a. OSPF/IS-IS interface goes operationally down: physical or local admin shutdown.

b. Timeout of a BFD session to a next-hop when BFD is enabled on the OSPF/IS-IS interface

When the SPF computation determines there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Thus, the IP prefix will resolve to the multiple equal-cost primary next-hops that provide the required protection.

The **no** form of this command disables the IP FRR feature on the system

Default no ip-fast-reroute

mc-maximum-routes

Syntax **mc-maximum-routes** *number* [**log-only**] [**threshold** *threshold*]
no mc-maximum-routes

Context config>router

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

Parameters *number* — Specifies the maximum number of routes to be held in a VRF context.

Values 1 — 2147483647

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

threshold *threshold* **—** The percentage at which a warning log message and SNMP trap should be sent.

Values 0 — 100

Default 10

mpls-labels

Syntax **mpls-labels**

Context config>router

Description This command creates a context for the configuration of glocal parameters related to MPLS labels.

static-label-range

Syntax **static-label-range** *static-range*
no static-label-range

Context config>router>mpls-labels

Description This command configures the range of MPLS static label values shared among static LSP, MPLS-TP LSP, and static service VC label. Once this range is configured, it is reserved and cannot be used by other protocols such as RSVP, LDP, BGP, or Segment Routing to assign a label dynamically.

Parameters *static-range* — Size of the static label range in number of labels. The minimum label value in the range is 32. The maximum label value is thus computed as {32+ static-range-1}.

Values 0 — 131040 for chassis mode C

Values 0 — 262112 for chassis mode D

Default 18400

sr-labels

Syntax **sr-labels start** *start-value* **end** *end-value*
no sr-labels

Context config>router>mpls-labels

Description This command configures the range of the Segment Routing Global Block (SRGB). It is a label block which is used for assigning labels to segment routing prefix SIDs originated by this router. This range is carved from the system dynamic label range and is not instantiated by default.

This is a reserved label and once configured it cannot be used by other protocols such as RSVP, LDP, and BGP to assign a label dynamically.

Parameters **start** *start-value* **—** start label value in the SRGB

Values 18432 — 524287

Default none

end *end-value* **—** end label value in the SRGB

Values 18432 — 524287

Default none

multicast-info

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

Context configure>router

Description This command configures multicast information policy.

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

Values 32 chars max

network-domains

Syntax **network-domains**

Context config>router

Description This command opens context for defining network-domains. This command is applicable only in the base routing context.

description

Syntax [**no**] **description** *string*

Context config>router>network-domains>network-domain

Description This command creates a text description stored in the configuration file for a configuration context.

The **no** form of the command removes the description string from the context.

Default no description

Parameters *string* — 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 character (#, $, space, etc.), the entire string must be enclosed within double quotes.

network-domain

Syntax **network-domain** *network-domain-name* [**create**]
no network-domain *network-domain-name*

Context config>router>network-domains

Description This command creates network-domains that can be associated with individual interfaces and SDPs.

Default **network-domain** "default"

Parameters *network-domain-name* — Network domain name character string.

rpki-session

Syntax **rpki-session** *ip-address*
no rpki-session *ip-address*

Context config>router>origin-validation

Description This command configures a session with an RPKI local cache server by using the RPKI-Router protocol. It is over these sessions that the router learns dynamic VRP entries expressing valid origin AS and prefix associations. SR-OS supports the RPKI-Router protocol over TCP/IPv4 or TCP/IPv6 transport. A 7x50 router can setup an RPKI-Router session using the base routing table or the management router.

Default **no rpki-session**

Parameters *ip-address* — An IPv4 address or an IPv6 address. If the IPv6 address is link-local then the interface name must be appended to the IPv6 address after a hyphen (-).

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>router>origin-validation>rpki-session

Description This command configures the time in seconds to wait between one TCP connection attempt that fails and the next attempt. The default (with no connect-retry) is 120 seconds.

Default **no connect-retry**

Parameters *seconds* — Specifies time in seconds.

Values 1-65535

description

Syntax **description** *description-string*
no description

Context config>router>origin-validation>rpki-session

Description This command configures a description for an RPKI-Router session.

Default **no description**

Parameters *description-string* — Specifies a text string up to 80 characters in length.

local-address

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

Context config>router>origin-validation>rpki-session

Description This command configures the local address to use for setting up the TCP connection used by an RPKI-Router session. The default local-address is the outgoing interface IPv4 or IPv6 address. The local-address cannot be changed without first shutting down the session.

Default **no local-address**

Parameters *ip-address* — Specifies an IPv4 address or an IPv6 address.

port

Syntax **port** *port-id*
no port

Context config>router>origin-validation>rpki-session

Description This command configures the destination port number to use when contacting the cache server. The default port number is 323. The port cannot be changed without first shutting down the session.

Default **no port**

Parameters *port-id* — Specifies a port-id.

Values 0-65535

refresh-time

Syntax **refresh-time** *seconds1* **hold-time** *seconds2*
no refresh-time

Context config>router>origin-validation>rpki-session

Description This command is used to configure the **refresh-time** and **hold-time** intervals that are used for liveness detection of the RPKI-Router session. The **refresh-time** defaults to 300 seconds and is reset whenever a Reset Query PDU or Serial Query PDU is sent to the cache server. When the timer expires, a new Serial Query PDU is sent with the last known serial number.

The **hold-time** specifies the length of time in seconds that the session is to be considered UP without any indication that the cache server is alive and reachable. The timer defaults to 600 seconds and must be at least 2x the refresh-time (otherwise the CLI command is not accepted). Reception of any PDU from the cache server resets the hold timer. When the **hold-time** expires, the session is considered to be DOWN and the stale timer is started.

Default **no referesh-time**

Parameters *seconds1* — Specifies a time in seconds.

Values 30-32767

seconds2 — Specifies a time in seconds.

Values 60-65535

shutdown

Syntax **shutdown**
no shutdown

Context config>router>origin-validation>rpki-session

Description This command administratively disables an RPKI-Router session. The no form of the command enables the RPKI-Router session.

Default **no shutdown**

stale-time

Syntax **stale-time** *seconds*
no stale-time

Context config>router>origin-validation>rpki-session

Description This command configures the maximum length of time that prefix origin validation records learned from the cache server remain useable after the RPKI-Router session goes down. The default stale-time is 3600 seconds (1 hour). When the timer expires all remaining stale entries associated with the session are deleted.

Default **no stale-time**

Parameters *seconds* — Specifies a time in seconds.

Values 60-3600

static-entry

Syntax **static-entry** *ip-prefix/ip-prefix-length* **upto** *prefix-length2* **origin-as** *as-number* [**valid** | **invalid**]
no static-entry *ip-prefix/ip-prefix-length* **upto** *prefix-length2* **origin-as** *as-number*

Context config>router>origin-validation

Description This command configures a static VRP entry indicating that a particular origin AS is either valid or invalid for a particular IP prefix range. Static VRP entries are stored along with dynamic VRP entries (learned from local cache servers using the RPKI-Router protocol) in the origin validation database of the router. This database is used for determining the **origin-validation** state of IPv4 and/or IPv6 BGP routes received over sessions with the **enable-origin-validatio**n command configured.

Note that static entries can only be configured under the **config>router>origin-validation** context of the base router.

Default **no static entries**

Parameters *ip-prefix/ip-prefix-length* — Specifies an IPv4 or IPv6 address with a minimum prefix length value.

Values 60-3600

prefix-length2 — Specifies the maximum prefix length.

as-number — Specifies as-number.

Values 0-4294967295

valid **—** Specifies a keyword meaning the static entry expresses a valid combination of origin AS and prefix range.

invalid **—** Specifies a keyword meaning the static entry expresses an invalid combination of origin AS and prefix range.

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>router

Description This command configures the router ID for the router instance.

The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager. IS-IS uses the router ID value as its system ID.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs.

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.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

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

Default The system uses the system interface address (which is also the loopback address).
If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

Parameters *router-id* — The 32 bit router ID expressed in dotted decimal notation or as a decimal value.

service-prefix

Syntax **service-prefix** *ip-prefix/mask* **|** *ip-prefix netmask* [**exclusive**]
no service-prefix *ip-prefix/mask* | *ip-prefix netmask*

Context config>router

Description This command creates an IP address range reserved for IES or VPLS services.

The purpose of reserving IP addresses using **service-prefix** is to provide a mechanism to reserve one or more address ranges for services.

When services are defined, the address must be in the range specified as a service prefix. If a service prefix is defined, then IP addresses assigned for services must be within one of the ranges defined in the **service-prefix** command. If the **service-prefix** command is not configured, then no limitations exist.

Addresses in the range of a service prefix can be allocated to a network port unless the exclusive parameter is used. Then, the address range is exclusively reserved for services.

When a range that is a superset of a previously defined service prefix is defined, the subset is replaced with the superset definition; for example, if a service prefix exists for 10.10.10.0/24, and a service prefix is configured as 10.10.0.0/16, then 10.10.10.0/24 is replaced by the new 10.10.0.0/16 configuration.

When a range that is a subset of a previously defined service prefix is defined, the subset replaces the existing superset, providing addresses used by services are not affected; for example, if a service prefix exists for 10.10.0.0/16, and a service prefix is configured as 10.10.10.0/24, then the 10.10.0.0/16 entry is removed as long as no services are configured that use 10.10.x.x addresses other than 10.10.10.x.

The **no** form of the command removes all address reservations. A service prefix cannot be removed while one or more service uses an address or addresses in the range.

Default no service-prefix - no IP addresses are reserved for services.

Parameters *ip-prefix/mask* — The IP address prefix to include in the service prefix allocation in dotted decimal notation.

Values		
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: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	

Values exclusive

When this option is specified, the addresses configured are exclusively used for services and cannot be assigned to network ports.

sgt-qos

Syntax **sgt-qos**

Context config>router

Description This command configures DSCP/Dot1p re-marking for self-generated traffic.

application

Syntax

application *dscp-app-name* **dscp** {*dscp-value* |*dscp-name*}
application *dot1p-app-name* **dot1p** *dot1p-priority*
no application {*dscp-app-name*|*dot1p-app-name*}

Context config>router>sgt-qos

Description This command configures DSCP/Dot1p re-marking for applications.

Parameters *dscp-app-name* — Specifies the DSCP application name.

Values bgp, cflowd, dhcp, dns, ftp, icmp, igmp, igmp-reporter, l2tp, ldp, mld, msdp, ndis, ntp, ospf, pim, ptp, radius, rip, rsvp, snmp, snmp-notification, srrp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp

dscp-value — Specifies the DSCP value

Values 0 — 63

dscp-name — Specifies the DSCP name.

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 none, 0 — 7

dot1p-app-name — Specifies the Dot1p application name.

Values arp, isis, pppoe

dscp

Syntax

dscp *dscp-name* **fc** *fc-name*
no dscp *dscp-name*

Context config>router>sgt-qos

Description This command configures DSCP name to FC mapping.

Parameters *dscp-name* — Specifies the DSCP name.

Values 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

fc-name — Specifies the forward class name.

Values be, l2, af, l1, h2, ef, h1, nc

bfd-template

Syntax **bfd-template** *name* [**create**]
no bfd-template *name*

Context config>router>bfd

Description This command creates or edits a BFD template. A BFD template defines the set of configurable parameters used by a BFD session. These include the transmit and receive timers used for BFD CC packets, the transmit timer interval used when the session is providing a CV function, the multiplier value, the echo-receive interval, and whether ther BFD session terminates in the CPM network processor.

Default no bfd-template

Parameters *name* — Specifies a text string name for the template up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

transmit-interval

Syntax **transmit-interval** *transmit-interval*
no transmit-interval

Context config>router>bfd>bfd-template

Description This command specifies the transmit timer used for BFD packets. If the template is used for a BFD session on an MPLS-TP LSP, then this timer is used for CC packets.

Default no transmit-interval

Parameters *transmit-interval* — Specifies the transmit interval. Note that the minimum interval that can be configured is hardware dependent.

Values 10 ms — 100,000 ms in 1 ms intervals

Default 10 ms for CPM3 or higher; 1 second for other hardware

receive-interval

Syntax **receive-interval** *receive-interval*
no receive-interval

Context config>router>bfd>bfd-template

Description This command specifies the receive timer used for BFD packets. If the template is used for a BFD session on an MPLS-TP LSP, then this timer is used for CC packets.

Default no receive-interval

Parameters *receive-interval* — Specifies the receive interval. Note that the minimum interval that can be configured is hardware dependent.

Values 10 ms — 100,000 ms in 1 ms intervals

Default 10 ms for CPM3 or higher; 1 second for other hardware

cv-tx

Syntax **cv-tx** *transmit-interval*
no cv-tx

Context config>router>bfd>bfd-template

Description This command specifies the transmit interval used by BFD packets used for MPLS-TP proactive CV.

Default no cv-tx

Parameters *transmit-interval* — Specifies the transmit interval. This parameter is only used if a BFD session is enabled with CV on an MPLS-TP LSP.

Values 1 sec to 30 sec in 1 second increments

Default 1 second

echo-receive

Syntax **echo-receive** *echo-interval*
no echo-receive

Context config>router>bfd>bfd-template

Description This command sets the minimum echo receive interval, in milliseconds, for a session. This is not used by a BFD session for MPLS-TP.

Default no echo-receive

Parameters *echo-interval* — Specifies the echo receive interval.

Values 100 ms — 100,000 ms in 1 ms increments

Default 100

multiplier

Syntax **multiplier** *multiplier*
no multiplier

Context config>router>bfd>bfd-template

Description This command specifies the detect multiplier used for a BFD session. If a BFD control packet is not received for a period of *multiplier* x *receive-interval*, then the session is declared down.

Default 3

Parameters *multiplier* — Specifies the multiplier.

Values 3 — 20, integers

Default 3

type

Syntax [**no**] **type cpm-np**

Context config>router>bfd>bfd-template

Description This command selects the CPM network processor as the local termination point for the BFD session. This is enabled by default.

Default type cpm-np

single-sfm-overload

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

Context config>router

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

static-route

Syntax [**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**community** *comm-id*] [**enable** | **disable**] **next-hop** *ip-int-name* | *ip-address* [*mcast-family*] [**bfd-enable** |{**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**padding-size** *padding-size*] [**log**]} {**prefix-list** *prefix-list-name* [**all** | **none**]} |{**fc** *fc-name* [**priority** {**low** | **high**}] [**ldp-sync**] [**validate-next-hop]**

[**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**community** *comm-id*] [**enable** | **disable**] **indirect** *ip-address* [**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**padding-size** *padding-size*] [**log**]] {**prefix-list** *prefix-list-name* [**all** | **none**]} |{**fc** *fc-name* [**priority** {**low** | **high**}]}

[**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*] {**prefix-list** *prefix-list-name* [**all** | **none**]}

Context config>router

Description This command creates static route entries for both the network and access routes.
When configuring a static route, either **next-hop**, **indirect** or **black-hole** must be configured.
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/prefix-length* — The destination address of the static route.

Values			
	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: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	

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.

Values		
	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x: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 characters maximum, mandatory for link local addresses

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)

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.

Values			
	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

ldp-sync — Extends the LDP synchronization feature to a static route. When an interface comes back up, it is possible that a preferred static route using the interface as next-hop for a given prefix is enabled before the LDP adjacency to the peer LSR comes up on this interface. In this case, traffic on an SDP that uses the static route for the far-end address would be black-holed until the LDP session comes up and the FECs exchanged.

This option when enabled delays the activation of the static route until the LDP session comes up over the interface and the ldp-sync-timer configured on that interface has expired.

preference *preference* **—** The preference of this static route versus the routes from different sources such as BGP or OSPF, expressed as a decimal integer. When modifing the preference of an existing static route, the metric will not be changed unless specified.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is according to the default preference table defined in Table 6 on page 135.

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 route to use is determined by the configuration of the **ecmp** command

prefix-list *prefix-list-name* [**all | none**] **—** Specifies the prefix-list to be considered.

metric *metric* **—** The cost metric for the static route, expressed as a decimal integer. This value is used when importing the static route into other protocols such as OSPF. When the metric is configured as 0 then the metric configured in OSPF, default-import-metric, applies. When modifying the metric of an existing static route, the preference will not change unless specified. This value is also used to determine which static route to install in the forwarding table:

- If there are multiple static routes with the same preference but different metrics then the lower cost (metric) route will be installed.
- If there are multiple static routes with equal preferences and metrics then ECMP rules apply .
- If there are multiple routes with different preferences then the lower preference route will be installed.

Default	1
Values	0 — 65535

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 or a point-to-point interface, the ip-int-name of the unnumbered or point-to-point 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-address* 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.

Values	ip-int-name	32 chars max
	ipv4-address	a.b.c.d
	ipv6-address	x:x: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 characters maximum, mandatory for link local addresses

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 indirect address can only resolved from dynamic routing protocol. Another static route cannot be used to resolve the indirect address.

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 the route is a black hole route. If the destination address on a packet matches this static route, it will be silently discarded.

The **black-hole** keyword and the **next-hop** or **indirect** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **next-hop** or **indirect** 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.

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.

validate-next-hop — This configuration option tracks the state of the next-hop in the IPv4 ARP cache or IPv6 Neighbor Cache. When the next-hop is not reachable and is removed from the ARP or Neighbor Cache, the next-hop will no longer be considered valid. When the next-hop is again reachable and present in the ARP/Neighbor Cache, the static route will be considered valid. **Note:** This feature is supported for directly connected next-hops only, and is exclusive with indirect routes.

Table 6: Default Route Preferences

Label	Preference	Configurable
Direct attached	0	No
Static-route	5	Yes
OSPF Internal routes	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Default 5

Values 1 — 255

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

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 **indirect** or **blackhole** keywords are specified. The remote end of the BFD session must also be configured to originate or accept the BFD session controlling the static-route state.

mcast-family — Enables submission of the IPv4 or IPv6 static route into IPv4 or IPv6 multicast RTM.

Values **mcast-ipv4**, **mcast-ipv6**

cpe-check *target-ip-address* **—** This parameter specifies the IP address of the target CPE device. This option initiates a background ICMP ping test to the configured target IP address. This address can either be an IPv4 address for IPv4 static routes or an IPv6 address for IPv6 static routes. 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

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 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.

Sample Output

```
*B:Dut-C# configure router "management"
*B:Dut-C>config>router# info
----------------------------------------------
       static-route 1.1.1.0/24 next-hop 172.31.117.1
        static-route 1::/96 next-hop 3000::AC1F:7567
----------------------------------------------
*B:Dut-C>config>router#


*B:Dut-C>config>router# show router "management" route-table
===============================================================================
Route Table (Router: management)
===============================================================================
Dest Prefix                                   Type    Proto    Age         Pref
       Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
1.1.1.0/24                                    Remote  Static   00h01m29s   0
       172.31.117.1                                                  1
138.203.0.0/16                                Remote  Static   05h01m11s   0
       172.31.117.1                                                  1
172.31.117.0/24                               Local   Local    05h04m10s   0
       management                                                    0
-------------------------------------------------------------------------------
No. of Routes: 3
===============================================================================
*B:Dut-C>config>router#


*B:Dut-C>config>router# show router "management" route-table ipv6
===============================================================================
IPv6 Route Table (Router: management)
===============================================================================
Dest Prefix                                   Type    Proto    Age         Pref
       Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
1::/96                                        Remote  Static   00h01m09s   5
       3000::AC1F:7567                                               1
3000::/96                                     Local   Local    05h04m12s   5
       management                                                    0
3FFE::/96                                     Remote  Static   00h00m11s   5
       3000::AC1F:7567                                               0
-------------------------------------------------------------------------------
No. of Routes: 3
===============================================================================
*B:Dut-C>config>router#
```

Note that the help info output (?) is inherited from the basic router context and does not reflect the specific syntax for the management context.

```
Only next-hop is allowed with any extra parameters.

*B:Dut-C>config>router# show router "management" static-?
static-arp      static-route


*B:Dut-C>config>router# show router "management" static-route
===============================================================================
Static Route Table (Router: management)  Family: IPv4
```

```
===============================================================================
Prefix                                        Tag         Met    Pref Type Act
   Next Hop                                    Interface
-------------------------------------------------------------------------------
1.1.1.0/24                                    0           1      5    NH   Y
   172.31.117.1                                n/a
-------------------------------------------------------------------------------
No. of Static Routes: 1
===============================================================================
*B:Dut-C>config>router#


*B:Dut-C>config>router# show router "management" static-route ipv6
===============================================================================
Static Route Table (Router: management)  Family: IPv6
===============================================================================
Prefix                                        Tag         Met    Pref Type Act
Next Hop                                    Interface
-------------------------------------------------------------------------------
1::/96                                        0           1      5    NH   Y
   3000::AC1F:7567                             management
-------------------------------------------------------------------------------
No. of Static Routes: 1
===============================================================================
*B:Dut-C>config>router#
```

static-route-entry

Syntax **static-route-entry** {*ip-prefix*/*prefix-length*} [**mcast**] **indirect** {*ip-address*}

Context config>router

Description This command enables the resolution of a static route prefix to an indirect tunnel next-hop.

Default No static routes are defined.

Parameters *ip-prefix/prefix-length* — The destination address of the static route.

Values			
	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: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	

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 indirect address can only resolved from dynamic routing protocol. Another static route cannot be used to resolve the indirect address.

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.

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.

Values	ipv4-address	a.b.c.d (host bits must be 0)
	ipv6-address	x:x: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 characters maximum, mandatory for link local addresses

tunnel-next-hop

Syntax **tunnel-next-hop**

Context config>router>static-route-entry

Description This command enables the context to configure the resolution of a static route prefix to an indirect tunnel next-hop.

The existing **static-route command** is still supported with all other options, including the **indirect** option which can be used to resolve the indirect next-hops in RTM.

The new command is an add-on to configure the resolution to tunnel next-hops in TTM. As such, the user must first configure the prefix with the existing command and the **indirect** option and then enter the new command with the **indirect** option and with the new **static-route-entry** command.

If **tunnel-next-hop** context is configured and **resolution** is set to **disabled**, the binding to tunnel is removed and resolution resumes in RTM to IP next-hops.

If **resolution** is set to **any**, any supported tunnel type in static route context will be selected following TTM preference.

The following tunnel types are supported in a static route context: RSVP and LDP.

The **ldp** value instructs the code to search for an LDP LSP with a FEC prefix corresponding to the address of the indirect next-hop.

The **rsvp** value instructs the code to search for the best metric RSVP LSP to the address of the indirect next-hop. This address can correspond to the system interface or to another loopback used 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, the code selects the LSP with the lowest tunnel-id.

If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference. In the case of RSVP-TE tunnel type, the user can further restrict the selection by providing a list of LSP names.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

If **disallow-igp** is enabled, the static-route will not be activated using IGP next-hops in RTM if no tunnel next-hops are found in TTM.

disallow-igp

Syntax **disallow-igp**
no disallow-igp

Context config>router>static-route-entry>tunnel-next-hop

Description This command is for indirect static routes using tunnel next-hops. When enabled, the static route will not be activated using IGP next-hops in RTM if no tunnel next-hops are found in TTM.

resolution

Syntax **resolution** {**any**|**filter**|**disabled**}

Context config>router>static-route-entry>tunnel-next-hop

Description This command configures the resolution mode in the resolution of a static route using tunnels to an indirect next-hop.

Parameters **any** — enables the binding to any supported tunnel type in a static route context following TTM preference.

filter — enables the binding to the subset of tunnel types configured under **resolution-filter**.

disabled — disables the resolution of a static route using tunnels to an indirect next-hop.

resolution-filter

Syntax **resolution-filter** [**ldp**] [**rsvp-te** [**lsp** *lsp name*]...[**lsp** *lsp name*]]

Context config>router>static-route-entry>tunnel-next-hop

Description This command configures the susbset of tunnel types which can be used in the resolution of a static route using tunnels to an indirect next-hop.
The following tunnel types are supported in a static route context RSVP and LDP. In the case of RSVP-TE tunnel type, the user can further restrict the selection by providing a list of LSP names.

Parameters **ldp —** selects the LDP tunnel type.

rsvp-te [**lsp** *lsp-name*]...[**lsp** *lsp-name*] **—** selects the RSVP-TE tunnel type or a set of specific RSVP LSP names.

triggered-policy

Syntax **triggered-policy**
no triggered-policy

Context config>router

Description This command triggers route policy re-evaluation.

By default, when a change is made to a policy in the **config router policy options** context and then committed, the change is effective immediately. There may be circumstances when the changes should or must be delayed; for example, if a policy change is implemented that would affect every BGP peer on a router, the consequences could be dramatic. It would be more effective to control changes on a peer-by-peer basis.

If the **triggered-policy** command is enabled, and a given peer is established, and you want the peer to remain up, in order for a change to a route policy to take effect, a **clear** command with the *soft* or *soft inbound* option must be used; for example, **clear router bgp neighbor x.x.x.x soft**. This keeps the peer up, and the change made to a route policy is applied only to that peer or group of peers.

ttl-propagate

Syntax **ttl-propagate**

Context config>router

Description This command enables the context to configure TTL propagation for transit and locally generated packets in the Global Routing Table (GRT) and VPRN routing contexts

Default none

label-route-local

Syntax **label-route-local [all | none]**

Context config>router>ttl-propagate

Description This command configures the TTL propagation for locally generated packets which are forwarded over a BGP label route in the Global Routing Table (GRT) context.

For IPv4 and IPv6 packets forwarded using a RFC 3107 label route in the global routing instance, including 6PE, the all value of the command enables TTL propagation from the IP header into all labels in the transport label stack. The none value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack. This command does not have a no version.

Note that the TTL of the IP packet is always propagated into the RFC 3107 label itself, and this command only controls the propagation into the transport labels, for example, labels of the RSVP or LDP LSP to which the BGP label route resolves and which are pushed on top of the BGP label.

Note that if the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves:

RSVP LSP shortcut:

- configure router mpls shortcut-local-ttl-propagate

LDP LSP shortcut:

- configure router ldp shortcut-local-ttl-propagate

This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for RSVP or LDP LSP shortcut listed.

Default none

Parameters **none —** The TTL of the IP packet is not propagated into the transport label stack.

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

label-route-transit

Syntax **label-route-transit [all | none]**

Context cconfig>router>ttl-propagate

Description This command configures the TTL propagation for transit packets which are forwarded over a BGP label route in the Global Routing Table (GRT) context.

For IPv4 and IPv6 packets forwarded using a RFC 3107 label route in the global routing instance, including 6PE, the all value of the command enables TTL propagation from the IP header into all labels in the transport label stack. The none value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack. This command does not have a no version.

Note that the TTL of the IP packet is always propagated into the RFC 3107 label itself, and this command only controls the propagation into the transport labels, for example, labels of the RSVP or LDP LSP to which the BGP label route resolves and which are pushed on top of the BGP label.

Note that if the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves.

RSVP LSP shortcut:

- configure router mpls shortcut-transit-ttl-propagate

LDP LSP shortcut:

- configure router ldp shortcut-transit-ttl-propagate

This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for the listed RSVP or LDP LSP shortcut.

Default none

Parameters **none —** The TTL of the IP packet is not propagated into the transport label stack.

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

lsr-label-route

Syntax **ttl-propagate [all | none]**

Context config>router>ttl-propagate

Description This command configures the TTL propagation for transit packets at a router acting as an LSR for a BGP label route.

When an LSR swaps the BGP label for a ipv4 prefix packet, thus acting as a ABR, ASBR, or data-path Route-Reflector (RR) in the base routing instance, or swaps the BGP label for a vpn-ipv4 or vpn-ipv6 prefix packet, thus acting as an inter-AS Option B VPRN ASBR or VPRN data path Route-Reflector (RR), the all value of this command enables TTL propagation of the decremented TTL of the swapped BGP label into all outgoing LDP or RSVP transport labels.

Note that when an LSR swaps a label or stitches a label, it always writes the decremented TTL value into the outgoing swapped or stitched label. What this feature controls is whether this decremented TTL value is also propagated to the transport label stack pushed on top of the swapped or stitched label.

The none value reverts to the default mode which disables TTL propagation. Note this changes the existing default behavior which propagates the TTL to the transport label stack. When a customer upgrades, the new default becomes in effect. This command does not have a no version.

This feature also controls the TTL propagation at an LDP-BGP stitching LSR in the LDP to BGP stitching direction. It also controls the TTL propagation in Carrier Supporting Carrier (CsC) VPRN at both the CsC CE and CsC PE.

Note that SROS does not support ASBR or data path RR functionality for labeled IPv6 routes in the global routing instance (6PE). As such the CLI command of this feature has no impact on prefix packets forwarded in this context.

Default none

Parameters **none —** The TTL of the swapped label is not propagated into the transport label stack.

all — The TTL of the swapped label is propagated into all labels of the transport label stack.

vprn-local

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

Context config>router>ttl-propagate

Description This command configures the TTL propagation for locally generated packets which are forwarded over a MPLS LSPs in all VPRN service contexts.

For vpn-ipv4 and vpn-ipv6 packets forwarded in the context of all VPRN services in the system, including 6VPE packets, the all value of the command enables TTL propagation from the IP header into all labels in the stack:

The user can enable the TTL propagation behavior separately for locally generated packets by CPM (vprn-local) and for user and control packets in transit at the node (vprn-transit).

The vc-only value reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. The user can explicitly set the default behavior by configuring the vc-only value. This command does not have a no version.

The value none allows the user to disable the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP trace-route in VPRN inter-AS option B such that the ingress and egress ASBR nodes are not traced.

The user can override the global configuration within each VPRN instance using the following commands:

- config service vprn ttl-propagate local [inherit | none | vc-only | all]
- config service vprn ttl-propagate transit [inherit | none | vc-only | all]

Note however 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.

When a packet is received in a VPRN context but is looked up in the Global Routing Table (GRT), for example, leaking to GRT is enabled, the behavior of the TTL propagation is governed by the RSVP or LDP shortcut configuration when the matching routing is a LSP shortcut route. It is governed by the BGP label route configuration when the matching route is a RFC 3107 label route or a 6PE route.

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance.

Default vc-only

Parameters **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.

vprn-transit

Syntax **vprn-transit [all | vc-only | none**]

Context config>router>ttl-propagate

Description This command configures the TTL propagation for in transit packets which are forwarded over a MPLS LSPs in all VPRN service contexts. For vpn-ipv4 and vpn-ipv6 packets forwarded in the context of all VPRN services in the system, including 6VPE packets, the all value of the command enables TTL propagation from the IP header into all labels in the stack:

The user can enable the TTL propagation behavior separately for locally generated packets by CPM (vprn-local) and for user and control packets in transit at the node (vprn-transit).

The vc-only value reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. The user can explicitly set the default behavior by configuring the vc-only value. This command does not have a no version.

The value none allows the user to disable the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP trace-route in VPRN inter-AS option B such that the ingress and egress ASBR nodes are not traced.

The user can override the global configuration within each VPRN service instance using the following commands:

- config service vprn ttl-propagate local [inherit | none | vc-only | all]
- config service vprn ttl-propagate transit [inherit | none | vc-only | all]

Note 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.

When a packet is received in a VPRN context but is looked up in the Global Routing Table (GRT), for example, leaking to GRT is enabled, the behavior of the TTL propagation is governed by the RSVP or LDP shortcut configuration when the matching routing is a LSP shortcut route. It is governed by the BGP label route configuration when the matching route is a RFC 3107 label route or a 6PE route.

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance

Default vc-only

Parameters **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.

Router L2TP Commands

l2tp

Syntax **l2tp**

Context config>router

Description 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.

calling-number-format

Syntax **calling-number-format** *ascii-spec*
no calling-number-format

Context config>router>l2tp

Description This command what string to put in the Calling Number AVP, for L2TP control messages related to a session in this L2TP protocol instance.

Parameters *ascii-spec* — Specifies the L2TP calling number AVP.

Values 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

exclude-avps

Syntax **exclude-avps** *calling-number*
no exclude-avps

Context config>router>l2tp

Description This command configures the L2TP AVPs to exclude.

next-attempt

Syntax **next-attempt** {**same-preference-level** | **next-preference-level**}
no next-attempt

Context configure>router>l2tp
configure>service>vprn>l2tp

Description This command enables tunnel selection algorithm based on the tunnel preference level.

Parameters **same-preference-level —** In case that the tunnel-spec selection algorithm evaluates into a tunnel that is currently unavailable (for example tunnel in a blacklist) then the next elected tunnel, if available, will be chosen within the same preference-level as the last attempted tunnel. Only when all tunnels within the same preference level are exhausted, the tunnel selection algorithm will move to the next preference level.

In case that a new session setup request is received while all tunnels on the same preference level are blacklisted, the L2TP session will try to be established on blacklisted tunnels before the tunnel selection moves to the next preference level.

next-preference-level — In case that the tunnel-spec selection algorithm evaluates into a tunnel that is currently unavailable (for example tunnel in a blacklist) then the selection algorithm will try to select the tunnel from the next preference level, even though the tunnels on the same preference level might be available for selection.

Default next-preference-level

replace-result-code

Syntax **replace-result-code** *code* [code...(upto 3 max)]
no replace-result-code

Context configure>router>l2tp
configure>service>vprn>l2tp

Description This command will replace CDN Result-Code 4, 5 and 6 on LNS with the Result Code 2. This is needed for interoperability with some implementation of LAC which only take action based on CDN Result-Code 2, while ignore CDN Result-Code 4, 5 and 6.

Default no replace-result-code

Parameters *code* — Specifies the L2TP Result codes that need to be replaced.

Values cdn-tmp-no-facilities — CDN Result-Code 4 on LNS will be replaced with the result code 2 before it is sent to LAC.
cdn-prem-no-facilities — CDN Result-Code 5 on LNS will be replaced with the result code 2 before it is sent to LAC.
cdn-inv-dest — CDN Result-Code 6 on LNS will be replaced with the result code 2 before it is sent to LAC.

tunnel-selection-blacklist

Syntax **tunnel-selection-blacklist**

Context config>router>l2tp

Description This command enables the context to configure L2TP Tunnel Selection Blacklist parameters.

add-tunnel

Syntax **add-tunnel never**
add-tunnel on *reason* [*reason*...(upto 8 max)]
no add-tunnel

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command will force the tunnel to the blacklist and render it unavailable for new sessions for the duration of pre-configured time. Peers are always forced to the black list in case that they time out (failure to receive response to control packets). In addition to time outs, certain events can be used to trigger placement of the tunnel on the black list.

Parameters *reason* — Specifies the return codes or events that determine which tunnels are added to the blacklist

Values **cdn-err-code** — A tunnel will be forced to the blacklist in case that CDN message with the Result Code 2 (Call disconnected for the reasons indicated in error code) is received.
cdn-inv-dest — A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 6 (Invalid destination) is received.
cdn-tmp-no-facilities — A tunnel will be forced to the blacklist in case that CDN message with the Result Code 4 is received (Call failed due to lack of appropriate facilities being available - temporary condition) is received.
cdn-perm-no-facilities — A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 5 (Call failed due to lack of appropriate facilities being available - permanent condition) is received.
tx-cdn-not-established-in-time — A tunnel will be forced to the blacklist in case that CDN message with the Result Code 10 (Call was not established within time allotted by LAC) is sent from the LAC to the LNS.
stop-ccn-err-code — A tunnel will be forced to the blacklist in case that StopCCN message with the Result Code 2 (General error – Error Code indicates the problem) is sent or received.
stop-ccn-other — A tunnel will be forced to the blacklist in case that StopCCN message with the following Result Codes is received:

(1) General request to clear control connection
(4) Requestor is not authorized to establish a control channel
(5) Protocol version not supported
(6) Requestor is being shutdown
Or in the case that the StopCCN with the following result codes is transmitted:

(4) Requestor is not authorized to establish a control channel.
(5) Protocol version not supported

The receipt of the following Result Codes will NEVER blacklist a tunnel:
(0) Reserved
(3) Control channel already exist
(7) Finite state machine error
(8) Undefined

Transmission of the following Result Codes will NEVER blacklist a tunnel:

(1) General request to clear control connection
(3) Control channel already exist
(6) Requestor is being shutdown
(7) Finite state machine error

addr-change-timeout — A timed-out tunnel for which the peer IP address has changed mid-session (from the one that is provided initially during configuration) will be forced to the blacklist. In absence of this configuration option, only the configured peer for the tunnel will be blacklisted, but not the tunnel itself which now has a different peer address than the one initially configured.

never — When specified, no tunnels will be placed on blacklist under any circumstance. This parameter will available to preserve backward compatibility.

max-list-length

Syntax **max-list-length unlimited**
max-list-length *count*
no max-list-length

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command configured the maximum length of the peer/tunnel blacklist.

This command specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist. If a tunnel or peer needs to be added to the tunnel-selection-blacklist and the tunnel-selection-blacklist is full, the system will remove the item (tunnel or peer) from the blacklist that was in this blacklist forthe longest time.

Default unlimited

Parameters **unlimited —** Specifies there is no limit.

count — Specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist.

Values 1..65635

max-time

Syntax **max-time** *minutes*
no max-time

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command configures time for which an entity (peer or a tunnel) are kept in the blacklist.

Default 5 minutes

Parameters *minutes* — Specifies the maximum time a tunnel or peer may remain in the blacklist

Values 1..60

timeout-action

Syntax **timeout-action** *action*
no timeout-action

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command defines an action that will be executed on the entity (peer/tunnel) in the blacklist once the entity becomes eligible for selection again.

Default remove-from-blacklist

Parameters *action* — Specifies the Action to be taken when a tunnel or peer has been in the blacklist for the max-period of time.

Values remove-from-blacklist — The peer or tunnel in the blacklist will be removed completely from the blacklist and made eligible for the selection process once the max-time expires. In this mode of operation, multiple new sessions can be mapped into the same, newly released tunnel from the blacklist. The first such session will try to setup the tunnel, while the other will be buffered until the tunnel establishment process is completed. In case that the tunnel remains unavailable, it will be placed in the blacklist again. Consequently all new sessions will have be re-negotiated over an alternate tunnel.
try-one-session — Once the max-time expired, the peer or tunnel in the blacklist is made available for selection only to a single new session request. Only upon successful tunnel establishment will the incoming new sessions be eligible to be mapped into this tunnel. This behavior will avoid session establishment delays in case that the tunnel just removed from the blacklist is still unavailable.

peer-address-change-policy

Syntax **peer-address-change-policy** {**accept** | **ignore** | **reject**}

Context config>router>l2tp

Description This command specifies what to do in case the system receives a L2TP responsefrom another address than the one the request was sent to.

Parameters **accept —** Specifies that this system accepts any source IP address change of received L2TP control messages related to a locally originated tunnel in the state waitReply and rejectsany peer address change for other tunnels; in case the new peer IPaddress is accepted, it is learned and used as destination addressin subsequent L2TP messages.

ignore — Specifiesthat this system ignores any source IP address change of received L2TP control messages, does not learn anynew peer IP address and does not change the destination address insubsequent L2TP messages.

reject — Specifies that this system rejects any source IP address change of received L2TP control messages and drops those messages.

receive-window-size

Syntax **receive-window-size** [4..1024]
no receive-window-size

Context config>router>l2tp

Description This command configures the L2TP receive window size.

session-limit

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

Context config>router>l2tp

Description 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>router>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>router>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>router>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.

Parameters *avp-hiding* — Specifies the method to be used for the authentication of the tunnels in this L2TP group.

Default no avp-hiding

Values sensitive — AVP hiding is used only for sensitive information (such as username/ password).
always — AVP hiding is always used.

challenge

Syntax **challenge** *always*
no challenge

Context config>router>l2tp>group

Description This command configures the use of challenge-response authentication.

The **no** form of the command reverts to the default **never** value.

Parameters *always* — Specifies that the challenge-response authentication is always used.

Default no challenge

Values always

df-bit-lac

Syntax **df-bit-lac** {**always**|**never**}
no df-bit-lac

Context config>router>l2tp
config>service>vprn>l2tp

Description By default, the LAC df-bit-lac is always set and sends all L2TP packets with the DF bit set to 1. The DF bit is configurable to allow downstream routers to fragment the L2TP packets. The LAC itself will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped.

Default df-bit-lac always

Parameters **always —** Specifies that the LAC will send all L2TP packets with the DF bit set to 1.

never — Specifies that the LAC will send all L2TP packets with the DF bit set to 0.

df-bit-lac

Syntax **df-bit-lac** {**always**|**never**|**default**}
no df-bit-lac

Context config>router/service>vprn>l2tp>group
config>router/service>vprn>l2tp>group>tunnel

Description By default, the LAC df-bit-lac is set to default and sends all L2TP packets with the DF bit set to 1. The DF bit is configurable to allow downstream routers to fragment the L2TP packets. The LAC itself will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped. The configuration of the df-bit can be overridden at different levels: l2tp, tunnel, and group. The configuration at the tunnel level overrides the configuration on both group and l2tp. The configuration at the group level overrides the configuration on l2tp.

Default df-bit-lac default

Parameters **always —** Specifies that the LAC will send all L2TP packets with the DF bit set to 1.

never — Specifies that the LAC will send all L2TP packets with the DF bit set to 0.

default — Follows the DF-bit configuration specified on upper levels.

destruct-timeout

Syntax **destruct-timeout** *destruct-timeout*
no destruct-timeout

Context config>router>l2tp>group
config>router>l2tp>group>tunnel

Description This command configures the period of time that the data of a disconnected tunnel will persist before being removed.

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

Default no destruct-timeout

Parameters *destruct-timeout* — [Specifies the automatic removal of dynamic L2TP sessions, in seconds, that are no longer active.

Default no destruct-timeout

Values 60 — 86400

hello-interval

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

Context config>router>l2tp>group

Description 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 **no** form of the command removes the interval from the configuration.

Default 60

Parameters *hello-interval* — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages.

Default no hello-interval

Values 60 — 3600

idle-timeout

Syntax **idle-timeout** *idle-timeout*
no idle-timeout

Context config>router>l2tp>group

Description This command configures the period of time that an established tunnel with no active sessions will persist before being disconnected.

Enter the **no** form of the command to maintain a persistent tunnel.

The **no** form of the command removes the idle timeout from the configuration.

Default no idle-timeout

Parameters *idle-timeout* — Specifies the idle timeout value, in seconds until the group is removed.

Default no idle-timeout

Values 0 — 3600

lns-group

Syntax **lns-group** *lns-group-id*
no lns-group

Context config>router>l2tp>group

Description This command configures the ISA LNS group.

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

Values 1 — 4

load-balance-method

Syntax **load-balance-method** {**per-session|per-tunnel**}
no load-balance-method

Context config>router>l2tp>group
config>router>l2tp>group>tunnel

Description This command describes how new sessions are assigned to an L2TP ISA MDA.

Parameters **per-session —** Specifies that the lowest granularity for load-balancing is a session; each session can be assigned to a different

ISA MDA.

per-tunnel — Specifies that the lowest granularity for load-balancing is a tunnel; all sessions associated with the same tunnel are assigned to the same ISA MDA; this may be useful or required in certain cases, for example:

- MLPPP with multiple links per bundle;
- HPol intermediate destination arbiters where the intermediate destination is an L2TP tunnel.

local-address

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

Context config>router>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>router>l2tp>group
config>router>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 thename 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>router>l2tp>group
config>router>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>router>l2tp>group
config>router>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>router>l2tp>group
config>router>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

Syntax **ppp**

Context config>router>l2tp>group

Description This command configures PPP for the L2TP tunnel group.

authentication

Syntax **authentication** {**chap**|**pap**|**pref-chap**}

Context config>router>l2tp>group>ppp

Description This command configures the PPP authentication protocol to negotiate.

authentication-policy

Syntax **authentication-policy** *auth-policy-name*
no authentication-policy

Context config>router>l2tp>group>ppp

Description This command configures the authentication policy.

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

Values 32 chars max

default-group-interface

Syntax **default-group-interface** *ip-int-name* **service-id** *service-id*
no default-group-interface

Context config>router>l2tp>group>ppp

Description This command configures the default group interface.

Parameters *ip-int-name* — Specifies the interface name.

Values 32 chars max

service-id — Specifies the service ID.

Values 1..2147483648

svc-name — Specifies the service name (instead of service ID).

Values 64 chars max

keepalive

Syntax **keepalive** *seconds* [**hold-up-multiplier** *multiplier*]
no keepalive

Context config>router>l2tp>group>ppp

Description This command configures the PPP keepalive interval and multiplier.

Parameters *seconds* — Specifies in seconds the interval.

Values 10 — 300

multiplier — Specifies the multiplier.

Values 1 — 5

mtu

Syntax **mtu** *mtu-bytes*
no mtu

Context config>router>l2tp>group>ppp

Description This command configures the maximum PPP MTU size.

Parameters *mtu-bytes* — Specifies, in bytes, the maximum PPP MTU size.

Values 512 — 9212

proxy-authentication

Syntax [**no**] **proxy-authentication**

Context config>router>l2tp>group>ppp

Description This command configures the use of the authentication AVPs received from the LAC.

proxy-lcp

Syntax [**no**] **proxy-lcp**

Context config>router>l2tp>group>ppp

Description This command configures the use of the proxy LCP AVPs received from the LAC.

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>router>l2tp>group>ppp

Description This command configures the local user database to use for PPP PAP/CHAP authentication.

Parameters *local-user-db-name* — Specifies the local user database name.

Values 32 chars max

session-assign-method

Syntax **session-assign-method** *weighted*
no session-assign-method

Context config>router>l2tp>group

Description This command specifies how new sessions are assigned to one of the set of suitable tunnels that are available or could be made available.

Default no session-assign-method

Parameters *weighted* — 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.

Default no session-assign-method. All new sessions are placed by preference in existing tunnels.

Values weighted — Enables weighted preference to tunnels in the group.

session-limit

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

Context config>router>l2tp>group
config>router>l2tp>group>tunnel

Description 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.

Default no session-limit

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

Values 1 — 131071

Router L2TP Tunnel Commands

tunnel

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

Context config>router>l2tp>group

Description 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).

Parameters *tunnel-name* — Specifies a valid string to identify a L2TP up to 32 characters in length.

create — mandatory while creating a new tunnel

auto-establish

Syntax [**no**] **auto-establish**

Context config>router>l2tp>group>tunnel

Description This command specifies if this tunnel is to be automatically set up by the system.

no auto-establish

avp-hiding

Syntax **avp-hiding** {**never** | **sensitive** | **always**}
no avp-hiding

Context config>router>l2tp>group>tunnel

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.

Note that it is recommended that sensitive information not be sent in clear text.

The **no** form of the command removes the parameter of the configuration and indicates that the value on group level will be taken.

Default no avp-hiding

Parameters *avp-hiding* — Specifies the method to be used for the authentication of the tunnel.

Values 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

Syntax **challenge** *challenge-mode*
no challenge

Context config>router>l2tp>group>tunnel

Description This command configures the use of challenge-response authentication.

The **no** form of the command removes the parameter from the configuration and indicates that the value on group level will be taken.

Default no challenge

Parameters *challenge-mode* — Specifies when challenge-response is to be used for the authentication of the tunnel.

Values always — Always allows the use of challenge-response authentication.
never — Never allows the use of challenge-response authentication.

hello-interval

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

Context config>router>l2tp>group>tunnel

Description 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.

Parameters *hello-interval* — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages.

Values 60 — 3600

infinite — Specifies that no hello messages are sent.

idle-timeout

Syntax **idle-timeout** *idle-timeout*
idle-timeout infinite
no idle-timeout

Context config>router>l2tp>group>tunnel

Description 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>router>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>router>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>router>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.

tunnel-selection-blacklist

Syntax **tunnel-selection-blacklist**

Context config>router>l2tp

Description This command enables the context to configure L2TP Tunnel Selection Blacklist parameters.

add-tunnel

Syntax **add-tunnel never**
add-tunnel on *reason* [*reason*...(upto 8 max)]
no add-tunnel

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command will force the tunnel to the blacklist and render it unavailable for new sessions for the duration of pre-configured time. Peers are always forced to the black list in case that they time out (failure to receive response to control packets). In addition to time outs, certain events can be used to trigger placement of the tunnel on the black list.

Parameters *reason* — Specifies the return codes or events that determine which tunnels are added to the blacklist

Values **cdn-err-code** — A tunnel will be forced to the blacklist in case that CDN message with the Result Code 2 (Call disconnected for the reasons indicated in error code) is received.
cdn-inv-dest — A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 6 (Invalid destination) is received.
cdn-tmp-no-facilities — A tunnel will be forced to the blacklist in case that CDN message with the Result Code 4 is received (Call failed due to lack of appropriate facilities being available - temporary condition) is received.
cdn-perm-no-facilities — A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 5 (Call failed due to lack of appropriate facilities being available - permanent condition) is received.
tx-cdn-not-established-in-time — A tunnel will be forced to the blacklist in case that CDN message with the Result Code 10 (Call was not established within time allotted by LAC) is sent from the LAC to the LNS.
stop-ccn-err-code — A tunnel will be forced to the blacklist in case that StopCCN message with the Result Code 2 (General error – Error Code indicates the problem) is sent or received.
stop-ccn-other — A tunnel will be forced to the blacklist in case that StopCCN message with the following Result Codes is received:

(1) General request to clear control connection
(4) Requestor is not authorized to establish a control channel
(5) Protocol version not supported
(6) Requestor is being shutdown
Or in the case that the StopCCN with the following result codes is transmitted:
(4) Requestor is not authorized to establish a control channel.

(5) Protocol version not supported
The receipt of the following Result Codes will NEVER blacklist a tunnel:
(0) Reserved
(3) Control channel already exist
(7) Finite state machine error
(8) Undefined
Transmission of the following Result Codes will NEVER blacklist a tunnel:
(1) General request to clear control connection
(3) Control channel already exist
(6) Requestor is being shutdown
(7) Finite state machine error

addr-change-timeout — A timed-out tunnel for which the peer IP address has changed mid-session (from the one that is provided initially during configuration) will be forced to the blacklist. In absence of this configuration option, only the configured peer for the tunnel will be blacklisted, but not the tunnel itself which now has a different peer address than the one initially configured.

never — When specified, no tunnels will be placed on blacklist under any circumstance. This parameter will available to preserve backward compatibility.

max-list-length

Syntax **max-list-length unlimited**
max-list-length *count*
no max-list-length

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command configured the maximum length of the peer/tunnel blacklist.

This command specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist. If a tunnel or peer needs to be added to the tunnel-selection-blacklist and the tunnel-selection-blacklist is full, the system will remove the item (tunnel or peer) from the blacklist that was in this blacklist forthe longest time.

Default unlimited

Parameters **unlimited —** Specifies there is no limit.

count — Specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist.

Values 1..65635

max-time

Syntax **max-time** *minutes*
no max-time

Context configure>router>l2tp>tunnel-selection-blacklist configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command configures time for which an entity (peer or a tunnel) are kept in the blacklist.

Default 5 minutes

Parameters *minutes* — Specifies the maximum time a tunnel or peer may remain in the blacklist

Values 1..60

timeout-action

Syntax **timeout-action** *action*
no timeout-action

Context configure>router>l2tp>tunnel-selection-blacklist
configure>service>vprn>l2tp>tunnel-selection-blacklist

Description This command defines an action that will be executed on the entity (peer/tunnel) in the blacklist once the entity becomes eligible for selection again.

Default remove-from-blacklist

Parameters *action* — Specifies the Action to be taken when a tunnel or peer has been in the blacklist for the max-period of time.

Values remove-from-blacklist — The peer or tunnel in the blacklist will be removed completely from the blacklist and made eligible for the selection process once the max-time expires. In this mode of operation, multiple new sessions can be mapped into the same, newly released tunnel from the blacklist. The first such session will try to setup the tunnel, while the other will be buffered until the tunnel establishment process is completed. In case that the tunnel remains unavailable, it will be placed in the blacklist again. Consequently all new sessions will have be re-negotiated over an alternate tunnel.
try-one-session — Once the max-time expired, the peer or tunnel in the blacklist is made available for selection only to a single new session request. Only upon successful tunnel establishment will the incoming new sessions be eligible to be mapped into this tunnel. This behavior will avoid session establishment delays in case that the tunnel just removed from the blacklist is still unavailable.

Router Interface Commands

interface

Syntax [**no**] **interface** *ip-int-name* [**unnumbered-mpls-tp**]
[**no**] **interface** *ip-int-name* **gmpls-loopback**

Context config>router

Description This command creates a logical IP routing or unnumbered MPLS-TP interface. Once created, attributes like IP address, port, or system can be associated with the IP interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for **config router interface** and **config service ies interface**. 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 the interface names or the IP addresses. Ambiguity can exist if an IP address is used as an IP address and an interface name. Duplicate interface names can exist in different router instances, although this is not recommended because it is confusing.

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.

Although not a keyword, the ip-int-name "**system**" is associated with the network entity (such as a specific 7750 SR), not a specific interface. The system interface is also referred to as the loopback address.

An unnumbered MPLS-TP interface is a special type of interface that is only intended for MPLS-TP LSPs. IP routing protocols are blocked on interfaces of this type. If an interface is configured as unnumbered-mpls-tp, then it can only be associated with an Ethernet port or VLAN, using the port command, then either a unicast, multicast, or broadcast remote MAC address may be configured. Only static ARP is supported.

A GMPLS loopback interface is a special type of loopback interface that is used as the IP interface for a GMPLS IP Control Channel (IPCC). RSVP and LMP packets associated with GMPLS are associated with this loopback interface. All other IP protocols are blocked on this interface. One **gmpls-loopback** interface is required for each GMPLS peer node.

The **no** form of the command removes the IP interface and all the associated configurations. The interface must be administratively shut down before issuing the **no interface** command.

Default No interfaces or names are defined within the system.

Parameters *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 within double quotes.

Values 1 — 32 alphanumeric characters.

If the *ip-int-name* already exists, the context is 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 the context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and the context is changed to that interface for further command processing.

unnumbered-mpls-tp — Specifies that an interface is of type Unnumbered MPLS-TP. An unnumbered MPLS-TP interface is a special type of interface that is only intended for MPLS-TP LSPs. IP routing protocols are blocked on interfaces of this type. If an interface is configured as **unnumbered-mpls-tp**, then it can only be associated with an Ethernet port or VLAN, using the **port** command. Either a unicast, multicast or broadcast remote MAC address may be configured using the **static-arp** command. Only static ARP is supported.

gmpls-loopback — Specifies that the interface is a loopback interface for GMPLS control plane packets.

address

Syntax **address** {*ip-address*/*mask*|*ip-address netmask*} [**broadcast** *all-ones* | **host-ones**] [**track-srrp** *srrp-instance*]
no address

Context config>router>interface

Description This command assigns an IP address, IP subnet, and broadcast address format to an IP interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IP interface. An IP address and a mask combine 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 router.

The local subnet that the **address** command defines must not be part of the services address space within the routing context by use of the **config router service-prefix** command. Once a portion of the address space is allocated as a service prefix, that portion is not available to IP interfaces for network core connectivity.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. **Show** commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The **no** form of the command removes the IP address assignment from the IP interface. Interface specificconfigurations for MPLS/RSVP are also removed. This will operationally stop any MPLS LSPs that explicitly reference that IP address. When a new IP address is defined, interface specific configurations for MPLS/RSVP will need to be re-added. If the **no** form of the command is executed then **ptp-hw-assist** is disabled. If a new address is entered while another address is still active, the new address will be rejected.

Default No IP address is assigned to the IP interface.

Parameters *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.

Values 1.0.0.0 — 223.255.255.255

/ — The forward slash is a parameter delimiter that separates the *ip-addr* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-addr*, the "/" and the *mask-length* parameter. If a forward slash does not ediately follow the *ip-addr*, 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-addr* 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 1— 32. Note that a mask length of 32 is reserved for system IP addresses.

Values 1 — 32

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-addr* 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. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 — 255.255.255.255

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)

broadcast {**all-ones** | **host-ones**} **—** 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 indictates 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 **all-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

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-addr* and the *mask-length* or *mask* with all the host bits set to binary 1. 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.

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

Values **all-ones**, **host-ones**

track-srrp — Specifies the SRRP instance ID that this interface route needs to track.

allow-directed-broadcasts

Syntax [**no**] **allow-directed-broadcasts**

Context config>router>interface

Description This command enables the forwarding of directed broadcasts out of the IP interface.

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address of another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface is sent as a subnet broadcast out this interface. **NOTE**: Allowing directed broadcasts is a well-known mechanism used for denial-of-service attacks.

By default, directed broadcasts are not allowed and are discarded at this egress IP interface.

The **no** form of the command disables directed broadcasts forwarding out of the IP interface.

Default no allow-directed-broadcasts — Directed broadcasts are dropped.

arp-limit

Syntax **arp-limit** *limit* **[log-only] [threshold** *percent***]**
no arp-limit

Context config>router>interface

Description 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 **no** form of the command removes the **arp-limit.**

Default 90 percent

Parameters **log-only —** 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.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 — 100

limit — 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.

Values 0 — 524288

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>router>interface

Description This command configures the minimum time, in seconds, an ARP entry learned on the IP interface is 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 the **arp-timeout** value is set to 0 seconds, ARP aging is disabled.

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

Default 14400 seconds (4 hours)

Parameters *seconds* — The minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of 0 specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*] [**type cpm-np**]
no bfd

Context config>router>interface
config>router>interface>ipv6

Description This command specifies the bi-directional forwarding detection (BFD) parameters for the associated IP interface. If no parameters are defined the default values are used.

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.

The **no** form of the command removes BFD from the router interface regardless of the IGP/RSVP.

Important notes: On the 7750-SR, the *transmit-interval* and **receive *receive-interval*** values can only be modified to a value less than 100 ms when:

1. The **type cpm-np option** is explicitly configured.
2. The service is shut down (**shutdown**)

3. The interval is specified 10 — 100000.
4. The service is re-enabled (**no shutdown**)

To remove the **type cpm-np** option, re-issue the **bfd** command without specifying the **type** parameter.

Default no bfd

Parameters *transmit-interval* — Sets the transmit interval, in milliseconds, for the BFD session.

Values 10 — 100000
10 — 100000 (see Important Notes above)

Default 100

receive *receive-interval* — Sets the receive interval, in milliseconds, for the BFD session.

Values 10 — 100000
10 — 100000 (see Important Notes above)

Default 100

multiplier *multiplier* — Set the multiplier for the BFD session.

Values 3— 20

Default 3

echo-receive *echo-interval* **—** Sets the minimum echo receive interval, in milliseconds, for the session.

Values 100 — 100000

Default 0

type cpm-np — Selects the CPM network processor as the local termination point for the BFD session. See Important Notes, above.

cflowd-parameters

Syntax **cflowd-parameters**
no cflowd-parameters

Context config>router>interface

Description This command creates the configuration context to configure cflowd parameters for the associated IP interfaces.

cflowd 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 **sampling** command must be configured within this context in order to enable cflowd sampling, otherwise traffic sampling will not occur.

Default no cflowd-parameters

sampling

Syntax **sampling {unicast | multicast} type {acl | interface} [direction {ingress-only|egress-only|both}]**
no sampling {unicast | multicast}

Context config>router>interface>cflowd-parameters

Description 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 **egress-only** or **both** specified or with the **ingress-only** keyword specified, then only ingress sampling will be enabled on the associated IP interface.

The **no** form of the command disables the associated type of traffic sampling on the associated interface.

Default no sampling

Parameters **unicast —** Specifies that the **sampling** command will control the sampling of unicast traffic on the associated interface/SAP.

mulitcast — Specifies that the **sampling** command will control the sampling of multicast traffic on the associated interface/SAP.

type —

Values acl — Specifies that the sampled traffic is controlled via an IP traffic filter entry with the action “filter-sample” configured.
interface — Specfies 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

Syntax **cpu-protection** *policy-id*
no cpu-protection

Context config>router>interface

Description This command assigns an existing CPU protection policy for the interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

Parameters *policy-id* — Specifies an existing CPU protection policy.

Values 1 — 255

delayed-enable

Syntax **delayed-enable** *seconds*
no delayed-enable

Context config>router>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

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>router>if

Description This command assigns a Distributed CPU protection policy for the interface.

enable-ingress-stats

Syntax [**no**] **enable-ingress-stats**

Context 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

Description This command enables the collection of ingress interface IP stats. This command is only appliable 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.

Default no enable-ingress-stats

enable-mac-accounting

Syntax [**no**] **enable-mac-accounting**

Context config>router>interface

Description This command enables MAC Accounting functionality for the interface.

if-attribute

Syntax **if-attribute**

Context config>router>interface

Description This command adds and removes interface attributes.

if-admin-group

Syntax [**no**] **if-admin-group** *group-name* [*group-name*...(upto 5 max)]

Context config>router>interface

Description This command configures interface Admin Group memberships for this interface.

if-srlg-group

Syntax [**no**] **if-srlg-group** *group-name* [*group-name*...(upto 5 max)]

Context config>router>interface

Description This command configures interface SRLG Group memberships for this interface

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>router>interface

Description This command enables local proxy ARP on the interface.

Default no local-proxy-arp

ip-mtu

Syntax **ip-mtu octets**
no ip-mtu

Context config>router>if

Description This command configures the IP maximum transmit unit (packet) for the associated router IP interface.

The configured IP-MTU cannot be larger then the calculated IP MTU based on the port MTU configuration.

The MTU that will be used is:

MINIMUM((Port_MTU - EtherHeaderSize), (Configured ip-mtu))

The **no** 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.)

Default **no ip-mtu**

Parameters *octets* — 5

Values 12 – 9000

lag-link-map-profile

Syntax **lag-link-map-profile** *link-map-profile-id*
no lag-link-map-profile

Context config>router>if

Description 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.

The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

Default **no lag-link-map-profile**

Parameters *link-map-profile-id* — An integer from 1 to 32 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

ldp-shortcut

Syntax [**no**] **ldp-shortcut**

Context config>router

Description This command enables the resolution of IGP routes using LDP LSP across all network interfaces participating in the IS-IS and OSPF routing protocol in the system.

When LDP shortcut is enabled, LDP populates the routing table with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in the system routing table. One route corresponds to the LDP shortcut next-hop and has an owner of LDP. The other route is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.

All user and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP.

When an IPv4 packet is received on an ingress network interface, a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress IOM will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.

If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabelled.

When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the ingress IOM will spray the packets for this route based on hashing routine currently supported for IPv4 packets. When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix..

The no form of this command disables the resolution of IGP routes using LDP shortcuts.

Default no ldp-shortcut

ldp-sync-timer

Syntax **ldp-sync-timer** *seconds*
no ldp-sync-timer

Context config>router>interface

Description This command enables synchronization of IGP and LDP. When a link is restored after a failure, IGP sets the link cost to infinity and advertises it. The actual value advertised in OSPF is 0xFFFF (65535). The actual value advertised in IS-IS regular metric is 0x3F (63) and in IS-IS wide-metric is 0xFFFFFE (16777214). This feature is not supported on RIP interfaces.

Note that if an interface belongs to both IS-IS and OSPF, a physical failure will cause both IGPs to advertise infinite metric and to follow the IGP-LDP synchronization procedures. If only one IGP bounced on this interface or on the system, then only the affected IGP advertises the infinite metric and follow the IGP-LDP synchronization procedures.

Next LDP hello adjacency is brought up with the neighbour. The LDP synchronization timer is started by IGP from the time the LDP session to the neighbor is UP over the interface. This is to allow time for the label-FEC bindings to be exchanged.

When the LDP synchronization timer expires, the link cost is restored and is re-advertised. IGP will announce a new best next-hop and LDP will use it if the label binding for the neighbor’s FEC is available.

If the user changes the cost of an interface, the new value is advertised at the next flooding of link attributes by IGP. However, if the LDP synchronization timer is still running, the new cost value will only be advertised after the timer expired. Also, the new cost value will be advertised after the user executes any of the following commands if the currently advertised cost is different:

- tools>perform>router>isis>ldp-sync-exit
- tools>perform>router>ospf>ldp-sync-exit
- config>router>interface>no ldp-sync-timer
- config>router>ospf>disable-ldp-sync
- router>isis>disable-ldp-sync

If the user changes the value of the LDP synchronization timer parameter, the new value will take effect at the next synchronization event. In other words, if the timer is still running, it will continue using the previous value.

If parallel links exist to the same neighbor, then the bindings and services should remain UP as long as there is one interface that is UP. However, the user configured LDP synchronization timer still applies on the failed then restored interface. In this case, the router will only consider this interface for forwarding after IGP re-advertized its actual cost value.

Note that the LDP Sync Timer State is not always synched across to the standby CPM, so after an activity switch the timer state might not be same as it was on the previous active CPM.

The **no** form of this command disables IGP/LDP synchronization and deletes the configuration

Default no ldp-sync-timer

Parameters *seconds* — Specifies the time interval for the IGP-LDP synchronization timer in seconds.

Values 1 – 1800

load-balancing

Syntax **load-balancing**

Context config>router>if

Description 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.

Default not applicable

egr-ip-load-balancing

Syntax **egr-ip-load-balancing** {**source** | **destination** | **inner-ip**}
no egr-ip-load-balancing

Context config>router>interface>load-balancing

Description 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 **no** form of this command includes both source and destination parameters.

Default no egr-ip-load-balancing

Parameters **source —** Specifies using source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port.

destination — Specifies using destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.

inner-ip — 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

Syntax **lsr-load-balancing** *hashing-algorithm*
no lsr-load-balancing

Context config>router>if>load-balancing

Description This command specifies whether the IP header is used in the LAG and ECMP LSR hashing algorithm. This is the per interface setting.

Default no lsr-load-balancing

Parameters **lbl-only —** Only the label is used in the hashing algorithm.

lbl-ip — The IP header is included in the hashing algorithm.

ip-only — the IP header is used exclusively in the hashing algorithm

eth-encap-ip — 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 or0x86DD). 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 (any of the check failed) label-stack hash is performed.

spi-load-balancing

Syntax [**no**] **spi-load-balancing**

Context config>router>if>load-balancing

Description This command enables use of the SPI in hashing for ESP/AH encrypted IPv4/v6 traffic. This is a per interface setting.

The **no** form disables the SPI function.

Default disabled

teid-load-balancing

Syntax [**no**] **teid-load-balancing**

Context config>router>interface>load-balancing

Description This command enables inclusion of TEID in hashing for GTP-U/C encapsulates traffic for GTPv1/GTPv2. The **no** form of this command ignores TEID in hashing.

Default disabled

loopback

Syntax [**no**] **loopback**

Context config>router>interface

Description This command configures the interface as a loopback interface.

Default Not enabled

mac

Syntax **mac** *ieee-mac-addr*
no mac

Context config>router>interface

Description This command assigns a specific MAC address to an IP interface. Only one MAC address can be assigned to an IP interface. When multiple **mac** commands are entered, the last command overwrites the previous command.

The **no** form of the command returns the MAC address of the IP interface to the default value.

Default IP interface has a system-assigned MAC address.

Parameters *ieee-mac-addr* — Specifies the 48-bit MAC address for the IP interface in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa*-*bb*-*cc*-*dd*-*ee*-*ff,* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

multihoming

Syntax [**no**] **multihoming primary|secondary** [**hold-time** *holdover-time*]

Context config>router>interface

Description This command sets the associated loopback interface to be an anycast address used in multi-homing resiliency, as either the primary or a secondary (a primary address on the alternate router). The optional hold-time parameter is only applicable for the secondary context and specifies how long label information learned about the secondary anycast address should be kept after that peer is declared down. This timer should be set to a value large enough for the remainder of the network to detect the failure and complete the reconvergence process.

The no form of the command disables this setting.

Default no multihoming

Parameters *holdover-time* — Specifies the number of seconds the router should hold label information learned from the alternate router in its secondary table. This is to allow the reset of the network to reconverge after a router failure before the anycase based label assignments are flushed from the forwarding plane.

Values 0 - 65535

Default 90

network-domain

Syntax **network-domain** *network-domain-name*
no network-domain

Context config>router>interface

Description This command assigns a given interface to a given network-domain. The network-domain is then taken into account during sap-ingress queue allocation for VPLS SAP.

The network-domain association can only be done in a base-routing context. Associating a network domain with an loop-back or system interface will be rejected. Associating a network-domain with an interface that has no physical port specified will be accepted, but will have no effect as long as a corresponding port, or LAG, is defined..

Single interfaces can be associated with multiple network-domains.

Default per default “default” network domain is assigned

ntp-broadcast

Syntax [**no**] **ntp-broadcast**

Context config>router>interface

Description This command enables SNTP broadcasts received on the IP interface. This parameter is only valid when the SNTP **broadcast-client** global parameter is configured.

The **no** form of the command disables SNTP broadcast received on the IP interface.

Default no ntp-broadcast

port

Syntax **port** *port-name*
no port

Context config>router>interface

Description This command creates an association with a logical IP interface and a physical port.

An interface can also be associated with the system (loopback address).

The command returns an error if the interface is already associated with another port or the system. In this case, the association must be deleted before the command is re-attempted. The *port-id* can be in one of the following forms:

- Ethernet interfaces

  If the card in the slot has MDAs, *port-id* is in the slot_number/MDA_number/port_number format; for example, **1/1/3** specifies port 3 of the MDA installed in MDA slot 1 on the card installed in chassis slot 1.

- SONET/SDH interfaces

  When the *port-id* represents a POS interface, the *port-id* must include the *channel-id.* The POS interface must be configured as a **network** port.

The **no** form of the command deletes the association with the port. The **no** form of this command can only be performed when the interface is administratively down.

Default No port is associated with the IP interface.

Parameters *port-name* — The physical port identifier to associate with the IP interface.

Values	*port-name*	*port-id*[:*encap-val*]		
		encap-val	0	for null
			0..4094	for dot1q
			0..4094.*	for qinq
	port-id	*slot/mda/port*[.*channel*]		
		bundle-id	- bundle-*type-slot/mda.bundle-num*	
			bundle	keyword
			type	ima, fr, ppp
			bundle-num	1..336
		bpgrp-id	bpgrp-*type-bpgrp-num*	
			bpgrp	keyword
			type	ima, ppp
			bpgrp-num	1..2000
		aps-id	aps-*group-id*[.*channel*]	
			aps	keyword
			group-id	1..64
		ccag-id	ccag-*id.path-id*[*cc-type*]	
			ccag	keyword
			id	1..8
			path-id	a, b
			cc-type	.sap-*net*, .net-*sap*

		lag-id	lag-*id*	
			lag	keyword
			id	1..800

Values	port-id	slot/mda/port[.channel]		
		bundle-id	bundle-type-slot/mda.bundle-num	
			bundle	keyword
			type	ima, ppp
			bundle-num	1 — 336
		bpgrp-id	bpgrp-type-bpgrp-num	
			bpgrp	keyword
			type	ima, ppp
			bpgrp-num	1 — 256
		aps-id	aps-group-id[.channel]	
			aps	keyword
			group-id	1 — 16
		lag-id	lag-id	
			lag	keyword
			id	1 — 64

Values	port-id	slot/mda/port[.channel]		
		ccag-id	ccag-id.path-id[cc-type]	
			ccag	keyword
			id	1 — 8
			path-id	a, b
			cc-type	.sap-net, .net-sap
		lag-id	lag-id	
			lag	keyword
			id	1 — 200
		gtg-id	gmpls-tun-grp-id	
			gmpls-tun-grp	keyword
			id	1 – 1024

proxy-arp-policy

Syntax [**no**] **proxy-arp-policy** *policy-name* [*policy-name*...(up to 5 max)]

Context config>router>interface

Description This command enables and configure proxy ARP on the interface and specifies an existing policy-statement to analyze match and action criteria that controls the flow of routing information to and from a given protocol, set of protocols, or a particular neighbor. The policy-name is configured in the **config>router>policy-options** context.

Use proxy ARP so the router responds to ARP requests on behalf of another device. Static ARP is used when a 7750 SR needs to know about a device on an interface that cannot or does not respond to ARP requests. Thus, the SR OS configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address.

Default no proxy-arp-policy

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 policy name(s) must already be defined.

ptp-hw-assist

Syntax [**no**] **ptp-hw-assist**

Context config>router>interface

Description This command configures the 1588 port based timestamping assist function for the interface. Various checks are performed to ensure that this feature can be enabled. If a check fails:

- The command is blocked/rejected with an appropriate error message.
- 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.

Note: The port will validate the destination IP address on received 1588 messages. If the 1588 messages are sent to a loopback address within the node rather than the address of the interface, then the loopback address must be configured in the **configure**>**system**>**security**>**source-address application ptp** context.

Default no ptp-hw-assist

qos-route-lookup

Syntax **qos-route-lookup** [**source** | **destination**]
no qos-route-lookup

Context config>router>interface
config>router>interface>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). Subscriber management group interfaces also do not support the source QPPB option.

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

Default destination

Parameters **source —** Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.

destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.

qos

Syntax **qos** *network-policy-id* [**egress-port-redirect-group** *queue-group-name*] [**egress-instance** *instance-id*]] [**ingress-fp- redirect-group** *queue-group-name* **ingress-instance** *instance-id*]
no qos

Context config>router>interface

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

egress-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

ingress-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

remote-proxy-arp

Context config>router>interface

Description This command enables remote proxy ARP on the interface.

Default no remote-proxy-arp

secondary

Syntax **secondary** {[*ip-address/mask* **|** *ip-address netmask*]} [**broadcast** {**all-ones | host-ones**}] [**igp-inhibit**]
no secondary *ip-addr*

Context config>router>interface

Description Use this command to assign up to 16 secondary IP addresses to the interface. Each address can be configured in an IP address, IP subnet or broadcast address format.

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.

Values 1.0.0.0 — 223.255.255.255

/ — The forward slash is a parameter delimiter that 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-addr*, the "/" and the *mask-length* parameter. If a forward slash does not ediately follow the *ip-addr*, 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 1— 32. Note that a mask length of 32 is reserved for system IP addresses.

Values 1 — 32

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-addr* 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. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 — 255.255.255.255

broadcast {**all-ones** | **host-ones**} **—** 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 **all-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

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-addr* and the *mask-length* or *mask* with all the host bits set to binary 1. 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.

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.

igp-inhibit — The secondary IP address should not be recognized as a local interface by the running IGP.

static-arp

Syntax **static-arp** *ip-addr ieee-mac-addr unnumbered*
no static-arp *unnumbered*

Context config>router>interface

Description This command configures a static Address Resolution Protocol (ARP) entry associating an IP address with a MAC address for the core router instance. This static ARP appears 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 is replaced by the new MAC address.
The number of static-arp entries that can be configured on a single node is limited to 1000.
Static ARP is used when a 7750 SR needs to know about a device on an interface that cannot or does not respond to ARP requests. Thus, the SR OS configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address. Use proxy ARP so the 7750 SR responds to ARP requests on behalf of another device.

The **no** form of the command removes a static ARP entry.

Default No static ARPs are defined.

Parameters *unnumbered* — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.

ieee-mac-addr — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa*-*bb*-*cc*-*dd*-*ee*-*ff*, where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

strip-label

Syntax [**no**] **strip-label**

Context config>router>interface

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 immediately 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 of the command 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

tos-marking-state

Syntax **tos-marking-state** {**trusted** | **untrusted**}
no tos-marking-state

Context config>router>interface

Description This command is used on a network IP interface 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 IES and network IP interface as untrusted.
When the ingress network IP 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.
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.
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 **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** form of the command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default trusted

Parameters **trusted —** 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

untrusted — 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.

unnumbered

Syntax **unnumbered** [*ip-address* | *ip-int-name*]
no unnumbered

Context config>router>interface

Description This command sets an IP interface as an unnumbered interface and specifies the IP address to be used for the interface.

To conserve IP addresses, unnumbered interfaces can be configured. The address used when generating packets on this interface is the *ip-addr* parameter configured.
An error message will be generated if an **unnumbered** interface is configured, and an IP address already exists on this interface.

The **no** form of the command removes the IP address from the interface, effectively removing the unnumbered property. The interface must be **shutdown** before **no unnumbered** is issued to delete the IP address from the interface, or an error message will be generated.

Parameters *ip-addr* | *ip-int-name* — Optional. The IP address or IP interface name to associate with the unnumbered IP interface in dotted decimal notation. The configured IP address must exist on this node. It is recommended to use the system IP address as it is not associated with a particular interface and is therefore always reachable. The system IP address is the default if no *ip-addr* or *ip-int-name* is configured.

Default no unnumbered

qos-route-lookup

Syntax **qos-route-lookup** [**source** | **destination**]
no qos-route-lookup

Context config>router>if
config>router>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. Subscriber management group interfaces also do not support the source QPPB option.

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

Default destination

Parameters **source —** Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.

destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.

secure-nd

Syntax [**no**] **secure-nd**

Context config>router>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>router>if>ipv6>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>router>if>ipv6>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>router>if>ipv6>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>router>if>ipv6>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>router>if>ipv6>secure-nd

Description This command enables or disables Secure Neighbor Discovery (SeND) on the interface.

stale-time

Syntax **stale-time** *seconds*
no stale-time

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

Description This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of the command removes the stale-time value.

Default **no stale-time**

Parameters *seconds* — The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

Values 60–65535

tcp-mss

Syntax **tcp-mss** *mss-value*
no tcp-mss

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

Description This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.

The **no** 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).

Default **no tcp-mss**

Parameters *mss-value* — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.
Note: 9158 = max-IP_MTU (9198)-40

Values 536 - 9158 (IPv4)
1220 - 9138 (IPv6)

urpf-check

Syntax [**no**] **urpf-check**

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

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.

Default disabled

vas-if-type

Syntax **vas-if-type** {**to-from-access** | **to-from-network** | **to-from-both**}
no vas-if-type

Context config>router>interface

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>router>if>urpf-check
config>router>if>>ipv6>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.

mh-primary-interface

Syntax [**no**] **mh-primary-interface**

Context config>router

Description This command creates a loopback interface for use in multihoming resiliency. Once active, this interface can be used to advertise reachability information to the rest of the network using the primary address, which is backed up by the secondary.

The reachability for this address is advertised via IGPs and LDP protocols to allow the resolution of BGP routes advertised with this address.

The no form of the command disables this setting.

Default no multihoming

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
no address

Context config>router>mh-primary-interface
config>router>mh-secondary-interface

Description This command assigns an IP address, IP subnet and broadcast address format to an IP interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each IP interface for the interface to be active. An IP address and a mask combine 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 interface in the same routing context within the router.

The local subnet that the address command defines must not be part of the services address space within the routing context by use of the config>router>service-prefix command. Once a portion of the address space is allocated as a service prefix, that portion is not available to IP interfaces for network core connectivity. The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. Show commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The no form of the command removes the IP address assignment from the IP interface. Interface specific configurations for IGP protocols like OSPF are also removed. The no form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface will operationally stop any protocol interfaces or MPLS LSPs that explicitly reference that IP address. When a new IP address is defined, the IP interface can be administratively enabled (no shutdown), which reinitializes the protocol interfaces and MPLS LSPs associated with that IP interface.

If a new address is entered while another address is still active, the new address wil be rejected.

Parameters *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.

Values 1.0.0.0 - 223.255.255.255

/ — The forward slash is a parameter delimiter that separates the ipp-addr portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the ip-addr, the "/" and the mask-length parameter. If a forward slash does not immediately follow the ip-addr, 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-addr 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 1-32. Note that a mask length of 32 is reserved for system IP addresses.

Values 1-32

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the ip-addr from a traditional dotted decimal mask. The mask parameters indicates the complete mask that will be used ina logical 'AND' function to derive the local subnet of the IP address. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 - 255.255.255.255

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 - 255.255.255.255 (nework bits all 1 and host bits all 0).

description

Syntax **description** *description-string*
no description

Context config>router>mh-primary-interface
config>router>mh-secondary-interface

Description This command creates a text description stored in the configuration file for a configuration context.

The no form of the command removes the description string from the context.

Default no description

Parameters *description-string* — 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 character (#, $, space, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>router>mh-primary-interface
config>router>mh-secondary-interface

Description The shutdown command administratively disables an entity. 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.

Unlike other commands and parameters where the default state is not indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.

The no form of the command puts an entity into the administratively enabled state.

Default no shutdown

if-attribute

Syntax **if-attribute**

Context config>router
config>router>interface
config>service>ies>interface
config>service>vprn>interface

Description 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

Syntax **admin-group** *group-name* **value** *group-value*
no admin-group *group-name*

Context config>router>if-attribute

Description This command defines an administrative group (admin-group) that can be associated with an IP or MPLS interface.

Admin groups, also known as affinity, are used to tag IP and MPLS interfaces that share a specific characteristic with the same identifier. For example, an admin group identifier can represent all links that connect to core routers, or all links that have a bandwidth higher than 10G, or all links that are dedicated to a specific service.

The user first configures locally on each router the name and identifier of each admin group. A maximum of 32 admin groups can be configured per system.

The user then configures the admin group membership of an interface. The user can apply admin groups to a IES, VPRN, network IP, or MPLS interface.

When applied to MPLS interfaces, the interfaces can be included or excluded in the LSP path definition by inferring the admin-group name. CSPF will compute a path that satisfies the admin-group include and exclude constraints.

When applied to IES, VPRN, or network IP interfaces, the interfaces can be included or excluded in the route next-hop selection by inferring the admin-group name in a route next-hop policy template applied to an interface or a set of prefixes.

The following provisioning rules are applied to admin group configuration. The system will reject the creation of an admin-group if it re-uses the same name but with a different group value than an existing group. The system will also reject the creation of an admin-group if it re-uses the same group value but with a different name than an existing group.

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 **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

Parameters *group-name* — Specifies the name of the group with up to 32 characters. The association of group name and value hsould be unique within an IP/MPLS domain.

value *group-value* **—** Specifies the integer value associated with the group. The association of group name and value should be unique within an IP/MPLS domain.

Values 0 — 31

admin-group

Syntax
admin-group *group-name* [*group-name*...**(up to 5 max)**]
no admin-group *group-name* [*group-name*...**(up to 5 max)**]
no admin-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 admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface.

Each single operation of the **admin-group** 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.

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.

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 **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 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* **value** *group-value* [**penalty-weight** *penalty-weight*]
no srlg-group *group-name*

Context config>router>if-attribute

Description This command defines a Shared Risk Link Group (SRLG) which can be associated with an IP or MPLS interface.

SRLG is used to tag IP or MPLS interfaces which share a specific fate with the same identifier. For example, an SRLG group identifier could represent all links which use separate fibers but are carried in the same fiber conduit. If the conduit is accidentally cut, all the fiber links are cut which means all interfaces using these fiber links will fail.

The user first configures locally on each router the name and identifier of each SRLG group. A maximum of 1024 SRLGs can be configured per system.

The user then configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface. A maximum of 64 SRLGs can be applied to a given interface.

When SRLGs are applied to MPLS interfaces, CSPF at an LER will exclude the SRLGs of interfaces used by the LSP primary path when computing the path of the secondary path. CSPF at an LER or LSR will also exclude the SRLGs of the outgoing interface of the primary LSP path in the computation of the path of the FRR backup LSP. This provides path disjointness between the primary path and the secondary path or FRR backup path of an LSP.

When SRLGs applied to IES, VPRN, or network IP interfaces, they are evaluated in the route next-hop selection by adding the **srlg-enable** option in a route next-hop policy template applied to an interface or a set of prefixes. For instance, the user can enable the SRLG constraint to select a LFA next-hop for a prefix which avoids all interfaces that share fate with the primary next-hop.

The following provisioning rules are applied to SRLG configuration. The system will reject the creation of a SRLG if it re-uses the same name but with a different group value than an existing group. The system will also reject the creation of an SRLG if it re-uses the same group value but with a different name than an existing group.

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.

A user may specify a penalty weight (**penalty-weight**) associated with an SRLG. This controls the likelihood of paths with links sharing SRLG values with a primary path being used by a bypass or detour LSP. The higher the penalty weight, the less desirable it is to use the link with a given SRLG.

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.

value *group-value* — Specifies the integer value associated with the group. The association of group name and value should be unique within an IP/MPLS domain.

Values 0 — 4294967295

penalty-weight *penalty-weight* — Specifies the integer value of the penalty weight that is assigned to the SRLG group.

Values 0 — 65535

Default 0

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 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.

route-next-hop-policy

Syntax **route-next-hop-policy**

Context config>router

Description This command creates the context to configure route next-hop policies.

template

Syntax [**no**] **template** *template-name*

Context config>router>route-next-hop-policy

Description This command creates a template to configure the attributes of a Loop-Free Alternate (LFA) Shortest Path First (SPF) policy. An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of an LFA backup next-hop for a subset of prefixes that resolve to a specific primary next-hop.

The user first creates a route next-hop policy template under the global router context and then applies it to a specific OSPF or IS-IS interface in the global routing instance or in a VPRN instance.

A policy template can be used in both IS-IS and OSPF to apply the specific criteria to prefixes protected by LFA. Each instance of IS-IS or OSPF can apply the same policy template to one or more interface.

The commands within the route next-hop policy template use the **begin-commit-abort** model. The following are the steps to create and modify the template:

1. To create a template, the user enters the name of the new template directly under the route-next-hop-policy context.
2. To delete a template that is not in use, the user enters the **no** form for the template name under the route-next-hop-policy context.
3. The user enters the editing mode by executing the begin command under the route-next-hop-policy context. The user can then edit and change any number of route next-hop policy templates. However, the parameter value will still be stored temporarily in the template module until the commit is executed under the route-next-hop-policy context. Any temporary parameter changes will be lost if the user enters the abort command before the commit command.
4. The user is allowed to create or delete a template instantly once in the editing mode without the need to enter the commit command. Furthermore, the abort command, if entered, will have no effect on the prior deletion or creation of a template.

Once the commit command is issued, IS-IS or OSPF will re-evaluate the templates and if there are any net changes, it will schedule a new LFA SPF to re-compute the LFA next-hop for the prefixes associated with these templates.

Parameters *template-name* — Specifies the name of the template, up to 32 characters.

include-group

Syntax **include-group** *group-name* [**pref** *pref*]
no include-group *group-name*

Context config>router>route-next-hop-policy>template

Description This command configures the admin group constraint into the route next-hop policy template.

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links which belong to one or more of the specified admin groups. A link which does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in a include-group statement but also belongs to other groups which are not part of any include-group statement in the route next-hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select a LFA backup next-hop which is a member of the corresponding admin group. If none is found, then the admin group with the next higher preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred, i.e., numerically the highest preference value.

When evaluating multiple **include-group** statements within the same preference, any link which belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both include and exclude statements, the exclude statement will win. It other words, the exclude statement can be viewed as having an implicit preference value of 0.

Note the admin-group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

Parameters *group-name* — Specifies the name of the group, up to 32 characters.

pref *pref* **—** An integer specifying the relative preference of a group.

Values 1 — 255

Default 255

exclude-group

Syntax **exclude-group** *group-name*
no exclude-group *group-name*

Context config>router>route-next-hop-policy>template

Description This command configures the admin group constraint into the route next-hop policy template.

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links that belong to one or more of the specified admin groups. A link that does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in an include-group statement but also belongs to other groups that are not part of any include-group statement in the route next-hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select an LFA backup next-hop that is a member of the corresponding admin group. If none is found, then the admin group with the next highest preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred (i.e., numerically the highest preference value).

When evaluating multiple **include-group** statements within the same preference, any link that belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both include and exclude statements, the exclude statement will win. It other words, the exclude statement can be viewed as having an implicit preference value of zero (0).

Note that the admin-group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

Parameters *group-name* — Specifies the name of the group, up to 32 characters.

srlg-enable

Syntax [**no**] **srlg-enable**

Context config>router>route-next-hop-policy>template

Description This command configures the SRLG constraint into the route next-hop policy template.

When this command is applied to a prefix, the LFA SPF will attempt to select an LFA next-hop, among the computed ones, which uses an outgoing interface that does not participate in any of the SLRGs of the outgoing interface used by the primary next-hop.

Note that the SRLG criterion is applied before running the LFA next-hop selection algorithm.

The **no** form deletes the SRLG constraint from the route next-hop policy template.

protection-type

Syntax **protection-type** {**link** | **node**}
no protection-type

Context config>router>route-next-hop-policy>template

Description This command configures the protection type constraint into the route next-hop policy template.

The user can select if link protection or node protection is preferred in the selection of an LFA next-hop for all IP prefixes and LDP FEC prefixes to which a route next-hop policy template is applied. The default in SR OS implementation is node protection. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.

When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop will follow the protection type preference specified in the template.

The **no** form deletes the protection type constraint from the route next-hop policy template.

Parameters {**link** | **node**} — Specifies the two possible values for the protection type.

Default node

nh-type

Syntax **nh-type** {**ip** | **tunnel**}
no nh-type

Context config>router>route-next-hop-policy>template

Description This command configures the next-hop type constraint into the route next-hop policy template.

The user can select if tunnel backup next-hop or IP backup next-hop is preferred. The default in SROS implementation is to prefer IP next-hop over tunnel next-hop. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.

When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop will follow the next-hop type preference specified in the template.

The **no** form deletes the next-hop type constraint from the route next-hop policy template.

Parameters {**ip** | **tunnel**} — Specifies the two possible values for the next-hop type.

Default ip

mh-secondary-interface

Syntax [**no**] **mh-secondary-interface**

Context config>router

Description This command creates a loopback interface for use in multihoming resiliency. This address is considered the secondary multihoming address and is only used to resolve routes advertised by the primary router in the event that router becomes unavailable. For this purpose, the reachability for this address is advertised via IGPs and LDP protocols to allow the resolution of BGP routes advertised with this address by the primary multihoming router.

The no form of the command disables this setting.

Default no mh-secondary-interface

hold-time

Syntax **hold-time** *holdover-time*
no hold-time

Context config>router>mh-secondary-interface

Description The optional hold-time parameter is only applicable for the secondary context and specifies how long label information leraned about the secondary anycast address should be kept after that peer is declared down. This timer should be set to a value large enough for the remainder of the network to detect the failure and complete the reconvergence process.

The no form of the command resets the hold-time back to the default value.

Default no hold-time

Parameters *holdover-time* — Specifies the number of seconds the router should hold label information learned from the alternate router in its secondary label table. This is to allow the reset of the network to reconverge after a router failure before the anycast based label assignments are flushed from the forwarding plane.

Values 0-65535

Default 90

Router Interface Filter Commands

egress

Syntax **egress**

Context config>router>interface

Description This command enables access to the context to configure egress network filter policies for the IP interface. If an egress filter is not defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>router>interface

Description This command enables access to the context to configure ingress network filter policies for the IP interface. If an ingress filter is not defined, no filtering is performed.

flowspec

Syntax [**no**] **flowspec**

Context config>router>interface>ingress

Description This command enables IPv4 flowspec filtering on a network IP interface. Filtering is based on all of the IPv4 flowspec routes that have been received and accepted by the base router 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:

1.user-defined IPv4 filter entries

2.flowspec-derived filter entries

3.user-defined IPv4 filter default-action

The **no** form of the command removes IPv4 flowspec filtering from the network IP interface.

Default No network interfaces have IPv4 flowspec enabled.

flowspec-ipv6

Syntax [**no**] **flowspec**

Context config>router>interface>ingress

Description This command enables IPv6 flowspec filtering on a network IP interface. Filtering is based on all of the IPv6 flowspec routes that have been received and accepted by the base router BGP instance. Ingress IPv6 traffic on an interface can be filtered by both a user-defined IPv4 filter and flowspec. Evaluation proceeds in this order:

1.user-defined IPv6 filter entries

2.flowspec-derived filter entries

3.user-defined IPv6 filter default-action

The **no** form of the command removes IPv6 flowspec filtering from the network IP interface.

Default No network interfaces have IPv6 flowspec enabled.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
no filter [**ip** *ip-filter-ip*] [**ipv6** *ipv6-filter-id*]

Context config>router>if>ingress
config>router>if>egress

Description This command associates an IP filter policy with an IP interface.

Filter policies control packet forwarding and dropping based on IP match criteria.

The *ip-filter-id* must have been pre-configured before this **filter** command is executed. If the filter ID does not exist, an error occurs.

Only one filter ID can be specified.

The **no** form of the command removes the filter policy association with the IP interface.

Default No filter is specified.

Parameters **ip** *ip-filter-id* — The filter name acts as the ID for the IP filter policy expressed as a decimal integer. The filter policy must already exist within the **config>filter>ip** context.

Values 1 — 16384

ipv6 *ipv6-filter-id* **—** The filter name acts as the ID for the IPv6 filter policy expressed as a decimal integer. The filter policy must already exist within the **config>filter>ipv6** context.

Values 1— 65535

Router Interface ICMP Commands

icmp

Syntax **icmp**

Context config>router>interface

Description This command enables access to the context to configure Internet Control Message Protocol (ICMP) parameters on a network IP interface. ICMP is a message control and error reporting protocol that also provides information relevant to IP packet processing.

mask-reply

Syntax [**no**] **mask-reply**

Context config>router>if>icmp

Description This command enables responses to ICMP mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

The **no** form of the command disables replies to ICMP mask requests on the router interface.

Default mask-reply — Replies to ICMP mask requests.

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>router>if>icmp

Description This command enables and configures the rate for 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 **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects are issued can be controlled with the optional *number* and *time* 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 the 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, expressed as a decimal integer. This parameter must be specified with the *time* parameter.

Values 10 — 1000

seconds — The time frame, in seconds, used to limit the *number* of ICMP redirect messages that can be issued,expressed as a decimal integer.

Values 1 — 60

ttl-expired

Syntax **ttl-expired** [*number seconds*]
no ttl-expired

Context config>router>if>icmp

Description This command configures the rate that Internet Control Message Protocol (ICMP) Time To Live (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 the command disables the generation of TTL expired messages.

Default ttl-expired 100 10 — Maximum of 100 TTL expired message in 10 seconds.

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

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

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>router>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 that can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of the command disables the generation of ICMP destination unreachables on the router interface.

Default unreachables 100 10 — Maximum of 100 unreachable messages in 10 seconds.

Parameters *number* — The maximum number of ICMP unreachable messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

Values 10 — 1000

seconds — The time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued, expressed as a decimal integer.

Router Interface IPv6 Commands

ipv6

Syntax [**no**] **ipv6**

Context config>router>interface

Description This command configures IPv6 for a router interface.

The **no** form of the command disables IPv6 on the interface.

Default not enabled

address

Syntax **address** {*ipv6-address/prefix-length*} [**eui-64**]
no address {*ipv6-address/prefix-length*}

Context config>router>if>ipv6

Description This command assigns an IPv6 address to the interface.

Default none

Parameters *ipv6-address/prefix-length* — Specify the IPv6 address on the interface.

Values	ipv6-address/prefix: ipv6-address	x: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
	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 POS interfaces, the Base MAC address of the chassis should be used.

dad-disable

Syntax [**no**] **dad-disable**

Context config>router>interface>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

icmp6

Syntax **icmp6**

Context config>router>if>ipv6

Description This command enables the context to configure ICMPv6 parameters for the interface.

packet-too-big

Syntax **packet-too-big** [*number seconds*]
no packet-too-big

Context config>router>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 packet-too-big messages.

Parameters *number* — Limits the number of packet-too-big messages issued per the time frame specifed in the *seconds* parameter.

Values 10 — 1000

seconds — Determines the time frame, in seconds, that is used to limit the number of packet-too-big messages issued per time frame.

Values 1 — 60

param-problem

Syntax **param-problem** [*number seconds*]
no param-problem

Context config>router>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 param-problem messages.

Parameters *number* — Limits the number of param-problem messages issued per the time frame specifed in the *seconds* parameter.

Values 10 — 1000

seconds — Determines the time frame, in seconds, that is used to limit the number of param-problem messages issued per time frame.

Values 1 — 60

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>router>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 redirect messages. When configured, ICMPv6 redirects are generated when routes are not optimal on the router and another router on the same subnetwork has a better route to alert that node that a better route is available.

The **no** form of the command disables ICMPv6 redirects.

Default 100 10 (when IPv6 is enabled on the interface)

Parameters *number* — Limits the number of redirects issued per the time frame specifed in *seconds* parameter.

Values 10 — 1000

seconds — Determines the time frame, in seconds, that is used to limit the number of redirects issued per time frame.

Values 1 — 60

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>router>if>ipv6>icmp6

Description This command configures rate for ICMPv6 time-exceeded messages.

Parameters *number* — Limits the number of time-exceeded messages issued per the time frame specifed in *seconds* parameter.

Values 10 — 1000

seconds — Determines the time frame, in seconds, that is used to limit the number of time-exceeded messages issued per time frame.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>router>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 unreachable messages. When enabled, ICMPv6 host and network unreachable messages are generated by this interface.

The **no** form of the command disables the generation of ICMPv6 host and network unreachable messages by this interface.

Default 100 10 (when IPv6 is enabled on the interface)

Parameters *number* — Determines the number destination unreachable ICMPv6 messages to issue in the time frame specified in *seconds* parameter.

Values 10 — 1000

seconds — Sets the time frame, in seconds, to limit the number of destination unreachable ICMPv6 messages issued per time frame.

Values 1 — 60

link-local-address

Syntax **link-local-address** *ipv6-address* [**preferred**]

Context config>router>if>ipv6
config>service>ies>if>ipv6
config>service>vprn>if>ipv6

Description 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.

Parameters **preferred —** Disables duplicated address detection and sets the address to preferred, even if there is a duplicate address.

local-proxy-nd

Syntax [**no**] **local-proxy-nd**

Context config>router>if>ipv6

Description This command enables local proxy neighbor discovery on the interface.

The **no** form of the command disables local proxy neighbor discovery.

neighbor

Syntax **neighbor** [*ipv6-address*] [*mac-address*]
no neighbor [*ipv6-address*]

Context config>router>if>ipv6

Description This command configures an IPv6-to-MAC address mapping on the interface. Use this command if a directly attached IPv6 node does not support ICMPv6 neighbor discovery, or for some reason, a static address must be used. This command can only be used on Ethernet media.

The *ipv6-address* must be on the subnet that was configured from the IPv6 **address** command or a link-local address.

Parameters *ipv6-address* — The IPv6 address assigned to a router interface.

Values	ipv6-address:	x: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

mac-address — Specifies the MAC address for the neighbor in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

neighbor-limit

Syntax **neigbor-limit** *limit* **[log-only] [threshold** *percent***]**
no neighbor-limit

Context config>router>if>ipv6

Description 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 **no** form of the command removes the neighbor-limit.

Default 90 percent

Parameters **log-only —** 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.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 — 100

limit — 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.

Values 0 — 102400

proxy-nd-policy

Syntax **proxy-nd-policy** *policy-name* [*policy-name*...(up to 5 max)]
no proxy-nd-policy

Context config>router>if>ipv6

Description This command configure a proxy neighbor discovery policy for the interface.

Parameters *policy-name* — The neighbor discovery 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 policy name(s) must already be defined.

Router Interface DHCP Commands

dhcp

Syntax **dhcp**

Context config>router>if

Description This command enables the context to configure DHCP parameters.

gi-address

Syntax **gi-address** *ip-address* [*src-ip-addr*]
no gi-address

Context config>router>if>dhcp

Description This command configures the gateway interface address for the DHCP relay. 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.

Default no gi-address

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies the source IP address to be used for DHCP relay packets.

option

Syntax [**no**] **option**

Context config>router>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

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

Context config>router>if>dhcp>option

Description 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 **no** form of this command returns the system to the default value.

Default Per RFC 3046, *DHCP Relay Agent Information Option* , section 2.1.1, *Reforwarded DHCP requests*, 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

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>router>if>dhcp>option

Description When enabled, the router sends the interface index (If Index) in the **circuit-id** suboption of the DHCP packet. The If Index of a router interface can be displayed using the command **show>router>interface>detail**. This option specifies data that must be unique to the router that is relaying the circuit.

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default circuit-id

Parameters **ascii-tuple —** Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “|”.

ifindex — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command **show>router>interface>detail**.

sap-id — Specifies that the SAP ID will be used.

vlan-ascii-tuple — 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.

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
no remote-id

Context config>router>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

Syntax [**no**] **vendor-specific-option**

Context config>router>if>dhcp>option

Description This command configures the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>router>if>dhcp>option

Description This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** 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

Syntax [**no**] **pool-name**

Context config>router>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the pool name in the Alcatel vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the feature.

port-id

Syntax [**no**] **port-id**

Context config>router>if>dhcp>option>vendor-specific-option

Description This command enables sending of the port-id in the Alcatel vendor specific suboption of the DHCP relay packet

The **no** form of the command disables the sending.

service-id

Syntax [**no**] **service-id**

Context config>router>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>router>if>dhcp>option>vendor-specific-option

Description This command specifies the vendor specific suboption string of the DHCP relay packet.

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

Parameters *text* — 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

Syntax [**no**] **system-id**

Context config>router>if>dhcp>option>vendor-specific-option

Description This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

Default None

relay-plain-bootp

Syntax [**no**] **relay-plain-bootp**

Context config>router>if>dhcp

Description This command enables the relaying of plain BOOTP packets.

The **no** form of the command disables the relaying of plain BOOTP packets.

server

Syntax **server** *server1* [*server2*...(up to 8 max)]

Context config>router>if>dhcp

Description 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.

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.

Default no server

Parameters *server* — Specifies the DHCP server IP address.

trusted

Syntax [**no**] **trusted**

Context config>router>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

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>router>if>dhcp

Description This comman specifies a python policy. Python policies are configured in the **config>python> python-policy** *name* context.

Parameters *name* — Specifies the name of an existing python script up to 32 characters in length.

Router Advertisement Commands

router-advertisement

Syntax [**no**] **router-advertisement**

Context config>router

Description This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.

The **no** form of the command disables all IPv6 interface. However, the **no interface** *interface-name* command disables a specific interface.

Default disabled

dns-options

Syntax [**no**] **dns-options**

Context config>router>router-advertisement
config>router>router-advertisement>interface

Description This command enables the context for configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

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 **config>router>router-advertisement>interface>dns-options>include-dns** command.

The **no** form of the command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.

Default disabled

dns-servers

Syntax **server** *ipv6-address*
no server

Context config>router>router-advertisement>dns-options
config>router>router-advertisement>interface>dns-options

Description 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 **include-dns** enabled, unless the interfaces have more specific **dns-options** configured.

Default 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>router>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>router>router-advertisement>dns-options
config>router>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>router>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

Syntax **current-hop-limit** *number*
no current-hop-limit

Context config>router>router-advert>if

Description 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.

Default 64

Parameters *number* — Specifies the hop limit.

Values 0 — 255. A value of zero means there is an unspecified number of hops.

managed-configuration

Syntax [**no**] **managed-configuration**

Context config>router>router-advert>if

Description 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, *Dynamic Host Configuration Protocol (DHCP) for IPv6*.

Default no managed-configuration

max-advertisement-interval

Syntax [**no**] **max-advertisement-interval** *seconds*

Context config>router>router-advert>if

Description This command configures the maximum interval between sending router advertisement messages.

Default 600

Parameters *seconds* — Specifies the maximum interval in seconds between sending router advertisement messages.

Values 4 — 1800

min-advertisement-interval

Syntax [**no**] **min-advertisement-interval** *seconds*

Context config>router>router-advert>if

Description This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.

Default 200

Parameters *seconds* — Specify the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages.

Values 3 — 1350

mtu

Syntax [**no**] **mtu** *mtu-bytes*

Context config>router>router-advert>if

Description This command configures the MTU for the nodes to use to send packets on the link.

Default no mtu — The MTU option is not sent in the router advertisement messages.

Parameters *mtu-bytes* — Specify the MTU for the nodes to use to send packets on the link.

Values 1280 — 9212

other-stateful-configuration

Syntax [**no**] **other-stateful-configuration**

Description 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, *Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6*

Default no other-stateful-configuration

prefix

Syntax [**no**] **prefix** [*ipv6-prefix*/*prefix-length*]

Context config>router>router-advert>if

Description 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.

Default none

Parameters *ip-prefix —* The IP prefix for prefix list entry in dotted decimal notation.

Values

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: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	

prefix-length — Specifies a route must match the most significant bits and have a prefix length.

Values 1 — 128

autonomous

Syntax [**no**] **autonomous**

Context config>router>router-advert>if>prefix

Description This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default enabled

on-link

Syntax [**no**] **on-link**

Context config>router>router-advert>if>prefix

Description This command specifies whether the prefix can be used for onlink determination.

Default enabled

preferred-lifetime

Syntax [**no**] **preferred-lifetime** {*seconds* **|** **infinite**}

Context config>router>router-advert>if

Description 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.

Default 604800

Parameters *seconds* — Specifies the remaining length of time in seconds that this prefix will continue to be preferred.

infinite — Specifies that the prefix will always be preferred. A value of 4,294,967,295 represents infinity.

valid-lifetime

Syntax **valid-lifetime** {*seconds* | **infinite**}

Context config>router>router-advert>if

Description 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.

Default 2592000

Parameters *seconds* — Specifies the remaining length of time in seconds that this prefix will continue to be valid.

infinite — Specifies that the prefix will always be valid. A value of 4,294,967,295 represents infinity.

reachable-time

Syntax **reachable-time** *milli-seconds*
no reachable-time

Context config>router>router-advert>if

Description This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

Default no reachable-time

Parameters *milli-seconds* — Specifies the length of time the router should be considered reachable.

Values 0 — 3600000

retransmit-time

Syntax **retransmit-timer** *milli-seconds*
no retransmit-timer

Context config>router>router-advert>if

Description This command configures the retransmission frequency of neighbor solicitation messages.

Default no retransmit-time

Parameters *milli-seconds* — Specifies how often the retransmission should occur.

Values 0 — 1800000

router-lifetime

Syntax **router-lifetime** *seconds*
no router-lifetime

Context config>router>router-advert>if

Description This command sets the router lifetime.

Default 1800

Parameters *seconds* — The length of time, in seconds, (relative to the time the packet is sent) that the prefix is valid for route determination.

Values 0, 4 — 9000 seconds. 0 means that the router is not a default router on this link.

use-virtual-mac

Syntax [**no**] **use-virtual-mac**

Context config>router>router-advert>if

Description 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 **no** form of the command disables sending router advertisement messages.

Default no use-virtual-mac