MPLS Configuration Commands

Generic Commands

shutdown

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

Context config>router>mpls
config>router>mpls>interface
config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

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

MPLS is not enabled by default and must be explicitely enabled (**no shutdown**).

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

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

Default **no shutdown**

MPLS Commands

mpls

Syntax [**no**] **mpls**

Context config>router

Description This command enables the context to configure MPLS parameters. MPLS is not enabled by default and must be explicitely enabled (**no shutdown**). The **shutdown** command administratively disables MPLS.

The **no** form of this command deletes this MPLS protocol instance; this will remove all configuration parameters for this MPLS instance.

MPLS must be **shutdown** and all SDP bindings to LSPs removed before the MPLS instance can be deleted. If MPLS is not shutdown, when the **no mpls** command is executed, a warning message on the console displays indicating that MPLS is still administratively up.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>router>mpls>ingr-stats
config>router>mpls>lsp>egr-stats
config>router>mpls>lsp-template>egr-stats

Description This command associates an accounting policy to the MPLS instance.

An accounting policy must be defined before it can be associated else an error message is generated.

The **no** form of this command removes the accounting policy association.

Default none

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>router>mpls>ingr-stats
config>router>mpls>lsp>egr-stats

config>router>mpls>lsp-template>egr-stats

Description This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the forwarding engine. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

max-stats

Syntax [**no**] **max-stats**

Context config>router>mpls>ingr-stats
config>router>mpls>lsp>egr-stats
config>router>mpls>lsp-template>egr-stats

Description This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no max-stats** command is issued the statistics are still accumulated by the forwarding engine. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **max-stats** command is issued then the counters written to the billing file include all the traffic while the **no max-stats** command was in effect.

Default max-stats

dynamic-bypass

Syntax **dynamic-bypass** [**enable** | **disable**]
no dynamic-bypass

Context config>router>mpls

Description This command disables the creation of dynamic bypass LSPs in FRR. One or more manual bypass LSPs must be configured to protect the primary LSP path at the PLR nodes.

Default enable

egress-statistics

Syntax [**no**] **egress-statistics**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command configures statistics in the egress data path of an originating LSP at a head-end node. The user must execute the no shutdown for this command to effectively enable statistics.

The same set of counters is updated for packets forwarded over any path of the LSP and over the lifetime of the LSP. In steady state, the counters are updated for packets forwarded over the active path of the LSP. The active path can be the primary path, one of the secondary paths, the FRR detour path, or the FRR bypass path when the head-end node is also the PLR.

LSP statistics are not collected on a dynamic or a static bypass tunnel itself. LSP egress statistics are also not collected if the head-end node is also the Penultimate-Popping Hop (PHP) node for a single-hop LSP using an implicit null label.

When a hierarchy of LSPs is in use, statistics collection on the outermost label corresponding to the tunneling LSP and on the inner labels, corresponding to the tunneled LSPs, are mutually exclusive. A consequence of this is that when the user enables statistics collection on an RSVP LSP which is also used for tunneling LDP FECs with the LDP over RSVP feature, then statistics will be collected on the RSVP LSP only. There will be no statistics collected from an LDP FEC tunneled over this RSVP LSP regardless if the user enabled statistics collection on this FEC. When, the user disables statistics collection on the RSVP LSP, then statistics collection, if enabled, will be performed on a tunneled LDP FEC.

The **no** form of this command disables the statistics in the egress data path and removes the accounting policy association from the RSVP LSP.

Default no egress-statistics

exponential-backoff-retry

Syntax **exponential-backoff-retry**
no exponential-backoff-retry

Context configure>router>mpls

Description This command enables the use of an exponential back-off timer when re-trying an LSP. When an LSP path establishment attempt fails, the path is put into retry procedures and a new attempt will be performed at the expiry of the user-configurable retry timer (config>router>mpls>lsp>retry-timer). By default, the retry time is constant for every attempt. The exponential back-off timer procedures will double the value of the user configured retry timer value at every failure of the attempt to adjust to the potential network congestion that caused the failure. An LSP establishment fails if no Resv message was received and the Path message retry timer expired or a PathErr message was received before the timer expired.

admin-group-frr

Syntax [**no**] **admin-group-frr**

Context config>router>mpls

Description This command enables the use of the admin-group constraints in the association of a manual or dynamic bypass LSP with the primary LSP path at a Point-of-Local Repair (PLR) node.

When this command is enabled, each PLR node reads the admin-group constraints in the FAST_REROUTE object in the Path message of the LSP primary path. If the FAST_REROUTE object is not included in the Path message, then the PLR will read the admin-group constraints from the Session Attribute object in the Path message.

If the PLR is also the ingress LER for the LSP primary path, then it just uses the admin-group constraint from the LSP and/or path level configurations.

The PLR node then uses the admin-group constraints along with other constraints, such as hop-limit and SRLG, to select a manual or dynamic bypass among those that are already in use.

If none of the manual or dynamic bypass LSP satisfies the admin-group constraints, and/or the other constraints, the PLR node will request CSPF for a path that merges the closest to the protected link or node and that includes or excludes the specified admin-group IDs.

If the user changes the configuration of the above command, it will not have any effect on existing bypass associations. The change will only apply to new attempts to find a valid bypass.

The **no** form of this command disables the use of administrative group constraints on a FRR backup LSP at a PLR node.

Default no frr-admin-group

frr-object

Syntax [**no**] **frr-object**

Context config>router>mpls

Description This command specifies whether fast reroute for LSPs using the **facility** bypass method is signalled with or without the fast reroute object using the **one-to-one** keyword. The value is ignored if fast reroute is disabled for the LSP or if the LSP is using one-to-one Backup.

Default frr-object — The value is by default inherited by all LSPs.

hold-timer

Syntax **hold-timer** *seconds*
no hold-timer

Context config>router>mpls

Description This command specifies the amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module. This occurs anytime the ingress node brings up an LSP path or switches traffic from a working path to another working path of the same LSP.

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

Parameters *seconds* — Specifies the time, in seconds, for which the ingress node holds before programming its data plane and declaring the LSP up to the service module.

Values 0 — 10

Default 1 second

ingress-statistics

Syntax **ingress-statistics**

Context config>router>mpls

Description This command provides the context for the user to enter the LSP names for the purpose of enabling ingress data path statistics at the terminating node of the LSP, for example, egress LER.

Default none

least-fill-min-thd

Syntax **least-fill-min-thd** *percent*
no least-fill-min-thd

Context config>router>mpls

Description This parameter is used in the least-fill path selection process. When comparing the percentage of least available link bandwidth across the sorted paths, whenever two percentages differ by less than the value configured as the least-fill-min-thresh, CSPF will consider them equal and will apply a random number generator to select the path among these paths

The **no** form of the command resets this parameter to its default value.

Default 5

Parameters *percentage* — Specifies the least fill minimum threshold value as a percentage.

Values 1 — 100%

least-fill-reoptim-thd

Syntax **least-fill-reoptim-thd** *percent*
no least-fill-reoptim-thd

Context config>router>mpls

Description This parameter is used in the least-fill path selection method. During a timer-based re-signaling of an LSP path which has the least-fill option enabled, CSPF will first update the least-available bandwidth figure for the current path of this LSP. It then applies the least-fill path selection method to select a new path for this LSP. If the new computed path has the same cost as the current path, it will compare the least-available bandwidth figures of the two paths and if the difference exceeds the user configured optimization threshold, MPLS will generate a trap to indicate that a better least-fill path is available for this LSP. This trap can be used by an external SNMP based device to trigger a manual re-signaling of the LSP path since the timer-based re-signaling will not re-signal the path in this case. MPLS will generate a path update trap at the first MBB event which results in the re-signaling of the LSP path. This should clear the eligibility status of the path at the SNMP device.

The **no** form of this command resets this parameter to its default value.

Default 10

Parameters *percentage* — Specifies the least fill reoptimization threshold value as a percentage.

Values 1 — 100%

lsp

Syntax [**no**] **lsp** *lsp-name* **sender** *sender-address*

Context config>router>mpls>ingress-statistics

Description This command configures statistics in the ingress data path of a terminating RSVP LSP at an egress LER. The LSP name must correspond to the name configured by the operator at the ingress LER. It must not contain the special character ":" which is used as a field separator by the ingress LER for encoding the LSP and path names into the RSVP session name field in the session_attribute object. The operator must execute the **no shutdown** for this command to effectively enable statistics.

The same set of counters is updated for packets received over any path of this LSP and over the lifetime of the LSP. In steady-state, the counters are updated for packets received over the active path of the LSP. The active path can be the primary path, one of the secondary paths, the FRR detour path, or the FRR bypass path when the tail-end node is also the MP.

When a hierarchy of LSPs is in use, statistics collection on the outermost label corresponding to the tunneling LSP and on the inner labels, corresponding to the tunneled LSPs are mutually exclusive. A consequence of this is that when the operator enables statistics collection on an RSVP LSP which is also used for tunneling LDP FECs with the LDP over RSVP feature, then statistics will be collected on the RSVP LSP only. There will be no statistics collected for an LDP FEC tunneled over this RSVP LSP and also egressing on the same node regardless if the operator enabled statistics collection on this FEC. When, the operator disables statistics collection on the RSVP LSP, then statistics collection, if enabled, will be performed on a tunneled LDP FEC.

The operator can enable statistics collection on a manual bypass terminating on the egress LER. However all LSPs which primary path is protected by the manual bypass will not collect statistics when they activate forwarding over the manual bypass. When, the operator disables statistics collection on the manual bypass LSP, then statistics collection on the protected LSP, if enabled, will continue when the bypass LSP is activated.

The **no** form of this command disables statistics for this RSVP LSP in the ingress data path and removes the accounting policy association from the LSP.

Default none

Parameters **sender-address** *ip-address* **—** A string of 15 characters representing the IP address of the ingress LER for the LSP.

lsp-name — A string of up to 32 characters identifying the LSP name as configured at the ingress LER.

p2p-template-lsp

Syntax [**no**] **p2p-template-lsp rsvp-session-name** *SessionNameString* **sender** *sender-address*

Context config>router>mpls>ingress-stats

Description This command configures an ingress statistics context matching on the RSVP session name for a RSVP P2P auto-LSP at the egress LER.

When the ingress LER signals the path of a template based **one-hop-p2p** or **mesh-p2p auto-lsp**, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: *lsp-name::path-name*, where *lsp-name* component is encoded as follows:

1. P2MP LSP via user configuration for L3 multicast in global routing instance: “LspNameFrom-Config”
1. P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: templateName-SvcId-mTTmIndex
1. P2MP LSP as I-PMSI in VPLS/B-VPLS: templateName-SvcId-mTTmIndex.

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2P auto-LSP or on a context that matches on the template name and the destination of the LSP at the ingress LER.

When the matching is performed on a context, the user must enter the RSVP session name string in the format “templateName-svcId” to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration. In this case, the user is provided with CLI parameter max-stats to limit the maximum number of stat indices which can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected.

Note the following rules when configuring an ingress statistics context based on template matching:

- **max-stats**, once allocated, can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
- In order to delete ingress statistics context matching a template name, a shutdown is required.

- An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shut down.
- After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a **no shut** is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2P LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same steam crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

The **no** form deletes the ingress statistics context matching on the RSVP session name.

Parameters **rsvp-session-name** *SessionNameString* **—** Specifies the name of the RSVP P2MP session in the format of “templateName-svcId”. This field can hold up to 64 characters.

sender *sender-address* **—** Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

p2mp-template-lsp

Syntax [**no**] **p2mp-template-lsp rsvp-session-name** *SessionNameString* **sender** *sender-address*

Context config>router>mpls>ingress-stats

Description This command configures an ingress statistics context matching on the RSVP session name for a RSVP P2MP LSP at the egress LER.

When the ingress LER signals the path of the S2L sub-LSP, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: <lsp-name::path-name>, where lsp-name component is encoded as follows:

- P2MP LSP via user configuration for L3 multicast in global routing instance: “LspNameFromConfig”
- P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: templateName-SvcId-mTTmIndex
- P2MP LSP as I-PMSI in VPLS/B-VPLS: templateName-SvcId-mTTmIndex

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2MP LSP as configured at the ingress LER or on a context that matches on the template name and the service-id as configured at the ingress LER.

When the matching is performed on a context, the user must enter the RSVP session name string in the format “templateName-svcId” to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration and the user is provided with CLI parameter max-stats to limit the maximum number of stat indices that can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected. A background tasks monitors the ingress statistics templates which have one or more matching LSP instances with unassigned stat index and assigns one to them as they are freed.

Note the following rules when configuring an ingress statistics context based on template matching:

- max-stats, once allocated, can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
- In order to delete ingress statistics context matching a template name, a shutdown is required.
- An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shut down.
- After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a "no shut" is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2MP LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same steam crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

The **no** form deletes the ingress statistics context matching on the RSVP session name.

Parameters **rsvp-session-name** *SessionNameString* — Specifies the name of the RSVP P2MP session in the format of "templateName-svcId". This field can hold up to 64 characters.

sender *sender-address* — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

logger-event-bundling

Syntax [**no**] **logger-event-bundling**

Context configure>router>mpls

Description This feature merges two of the most commonly generated MPLS traps, vRtrMplsXCCreate and vRtrMplsXCDelete, which can be generated at both LER and LSR into a new specific trap vRtrMplsSessionsModified. In addition, this feature will perform bundling of traps of multiple RSVP sessions, i.e., LSPs, into this new specific trap.

The intent is to provide a tool for the user to minimize trap generation in an MPLS network. Note that the MPLS trap throttling will not be applied to this new trap.

The **no** version of this command disables the merging and bundling of the above MPLS traps.

lsp-template

Syntax **lsp-template** template-name [**p2mp | one-hop-p2p | mesh-p2p**]
no lsp-template *template-name*

Context config>router>mpls

Description This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. The LSP template type p**2mp, one-hop-p2p**, or **mesh-p2p** is mandatory.

The **no** form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.

Parameters *lsp-template-name* — Specifies the name of the LSP template. Any LSP template name and LSP name must not be the same.

p2mp | one-hop-p2p | mesh-p2p — Identifies the t ype of the LSP this template will signal.

lsp-init-retry-timeout

Syntax **lsp-init-retry-timeout** *seconds*
no lsp-init-retry-timeout

Context config>router>mpls

Description This command configures the initial LSP path retry-timer.

The new LSP path initial retry-timer is used instead of the retry-timer to abort the retry cycle when no RESV is received. The retry-timer will govern exclusively the time between two retry cycles and to handle retrying of an LSP path in a failure case with PATH errors or RESVTear.

The intent is that the user can now control how many refreshes of the pending PATH state can be performed before starting a new retry-cycle with a new LSP-id. This is all done without affecting the ability to react faster to failures of the LSP path, which will continue to be governed by the retry-timer.

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

Parameters *seconds* — Specifies the value, in seconds, used as the fast retry timer for a secondary path.

Values 10—600

Default 30

lsp-template

Syntax **lsp-template** *template-name* [**p2mp | one-hop-p2p | mesh-p2p**]
no lsp-template *template-name*

Context config>router>mpls

Description This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. The LSP template type **p2mp, one-hop-p2p**, or **mesh-p2p** is mandatory.

The **no** form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.

Parameters *lsp-template-name* — Specifies the name to identify LSP template. ANy LSP template name and LSP name must not be the same.

p2mp | one-hop-p2p | mesh-p2p — Identifies the type of the LSP this template will signal.

propagate-admin-group

Syntax [**no**] **propagate-admin-group**

Context config>router>mpls>lsp>fast-reroute
config>router>mpls>lsp-template>fast-reroute

Description The command enables the signaling of the primary LSP path admin-group constraints in the FRR object at the ingress.

When this command is executed, the admin-group constraints configured in the context of the P2P LSP primary path, or the ones configured in the context of the LSP and inherited by the primary path, are copied into the FAST_REROUTE object. The admin-group constraints are copied into the ‘include-any’ or ‘exclude-any’ fields.

The ingress LER thus propagates these constraints to the downstream nodes during the signaling of the LSP to allow them to include the admin-group constraints in the selection of the FRR backup LSP for protecting the LSP primary path.

The ingress LER will insert the FAST_REROUTE object by default in a primary LSP path message. If the user disables the object using the following command, the admin-group constraints will not be propagated: **configure>router>mpls>no frr-object** .

Note that the same admin-group constraints can be copied into the Session Attribute object. They are intended for the use of an LSR, typically an ABR, to expand the ERO of an inter-area LSP path. They are also used by any LSR node in the path of a CSPF or non-CSPF LSP to check the admin-group constraints against the ERO regardless if the hop is strict or loose. These are governed strictly by the command:

configure>router>mpls>lsp>propagate-admin-group

In other words, the user may decide to copy the primary path admin-group constraints into the FAST_REROUTE object only, or into the Session Attribute object only, or into both. Note, however, that the PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.

This feature is supported with the following LSP types and in both intra-area and inter-area TE where applicable:

- Primary path of a RSVP P2P LSP.
- S2L path of an RSVP P2MP LSP instance
- LSP template for an S2L path of an RSVP P2MP LSP instance.

The **no** form of this command disables the signaling of administrative group constraints in the FRR object.

Default no propagate-admin-group

max-bypass-associations

Syntax **max-bypass-associations** *integer*
no max-bypass-associations

Context config>router>mpls

Description This command allows the user to set a maximum number of LSP primary path associations with each manual or dynamic bypass LSP that is created in the system.

By default, a Point of Local Repair (PLR) node will associate a maximum of 1000 primary LSP paths with a given bypass before using the next available manual bypass or signaling a new dynamic bypass.

Note that a new bypass LSP may need to be signaled if the constraint of a given primary LSP path is not met by an existing bypass LSP even if the max-bypass-associations for this bypass LSP has not been reached.

The **no** form of the command re-instates the default value of this parameter.

Default no max-bypass-associations

Values 1 — 131,072

mbb-prefer-current-hops

Syntax [**no**] **mbb-prefer-current-hops**

Context config>router>mpls

Description This command implements a new option in the CSPF path computation during a Make-Before-Break (MBB) procedure of an RSVP LSP.

When MPLS performs an MBB for the primary or secondary path of a P2P LSP, or the S2L path of a P2MP LSP, and the new **mbb-prefer-current-hops** option is enabled in MPLS context, CSPF will select a path, among equal-cost candidate paths, with the most overlapping links with the current path. Normally, CSPF selects the path randomly.

The procedures of the new MBB CSPF path selection apply to LSP without the least-fill option enabled. If the least-fill rule results in a different path, the LSP path will be moved though. Users can still favor stability over least-fill condition by applying a larger value to the parameter **least-fill-min-**

thd under the MPLS context such that a path will only be moved when the difference of the least-available bandwidth becomes significant enough between the most used links in the equal cost paths. If that difference is not significant enough, CSPF will select the path with the most overlapping links instead of selecting a path randomly.

The procedures when the new **mbb-prefer-current-hops** option is enabled apply to all MBB types. Thus, it applies to the auto-bandwidth MBB, the configuration change MBB, the soft pre-emption MBB, the TE graceful shutdown MBB, the delayed retry MBB (for SRLG secondary LSP path), the path change MBB, the timer resignal MBB, and the manual resignal MBB.

During the FRR global revertive MBB, CSPF selects a random link among the ones available between the PLR node and the Merge Point node, including the failed link if it has restored in the meantime. These links cannot be checked for overlap with the current path.

The TE graceful shutdown MBB will still avoid the link or node that is in maintenance and the soft pre-emption MBB will still avoid the link that is overbooked.

For an inter-area LSP, this feature applies to the subset of the path from the ingress LER to the exit ABR.

The procedures of this feature are not applied to a zero bandwidth CSFP LSP, including an auto-bandwidth CSPF LSP while its operational bandwidth is zero, and to a non-CSPF LSP.

resignal-timer

Syntax **resignal-timer** *minutes*
no resignal-timer

Context config>router>mpls

Description This command specifies the value for the LSP resignal timer. The resignal timer is the time, in minutes, the software waits before attempting to resignal the LSPs.

When the resignal timer expires, if the new computed path for an LSP has a better metric than the current recorded hop list, an attempt is made to resignal that LSP using the make-before-break mechanism. If the attempt to resignal an LSP fails, the LSP will continue to use the existing path and a resignal will be attempted the next time the timer expires.

The **no** form of the command disables timer-based LSP resignalling.

Default no resignal-timer

Parameters *minutes* — The time the software waits before attempting to resignal the LSPs.

Values 30 — 10080

retry-on-igp-overload

Syntax [**no**] **retry-on-igp-overload**

Context config>router>mpls

Description This command enables tearing down LSPs when IGP is in overload state.

secondary-fast-retry-timer

Syntax **secondary-fast-retry-timer** *seconds*
no secondary-fast-retry-timer

Context config>router>mpls

Description This command specifies the value used as the fast retry timer for a secondary path. If the first attempt to set up a secondary path fails due to a path error, the fast retry timer will be started for the secondary path so that the path can be retried sooner. If the next attempt also fails, further retries for the path will use the configured value for LSP retry timer.

If retry-timer for the LSP is configured to be less than the MPLS secondary-fast-retry-timer, all retries for the secondary path will use the LSP retry-timer.

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

Default no secondary-fast-retry-timer

Parameters *seconds* — specifies the value, in seconds, used as the fast retry timer for a secondary path

Values 1 — 10

srlg-frr

Syntax **srlg-frr** [**strict**]
no srlg-frr

Context config>router>mpls

Description This command enables the use of the Shared Risk Loss Group (SRLG) constraint in the computation of FRR bypass or detour to be associated with any primary LSP path on this system.

When this option is enabled, CSPF includes the SRLG constraint in the computation of a FRR detour or bypass for protecting the primary LSP path.

CSPF prunes all links with interfaces which belong to the same SRLG as the interface which is being protected, i.e., the outgoing interface at the PLR the primary path is using. If one or more paths are found, the MPLS/RSVP task will select one based on best cost and will signal the bypass/detour. If not and the user included the strict option, the bypass/detour is not setup and the MPLS/RSVP task will keep retrying the request to CSPF. Otherwise, if a path exists which meets the other TE constraints, other than the SRLG one, the bypass/detour is setup.

A bypass or a detour LSP path is not guaranteed to be SRLG disjoint from the primary path. This is because only the SRLG constraint of the outgoing interface at the PLR the primary path is using is checked.

When the MPLS/RSVP task is searching for a SRLG bypass tunnel to associate with the primary path of the protected LSP, it will first check if any configured manual bypass LSP with CSPF enabled satisfies the SLRG constraints. The MPLS/RSVP skips any non-CSPF bypass LSP in the search as there is no ERO returned to check the SLRG constraint. If no path is found, it will check if an existing dynamic bypass LSP satisfies the SLRG and other primary path constraints. If not, then it will make a request to CSPF.

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG group membership of an interface the primary path is using would not be considered by the MPLS/RSVP task at the PLR for bypass/detour association until the next opportunity the primary path is re-signaled. The path may be re-signaled due to a failure or to a make-before break operation. Make-before break occurs as a result of a global revertive operation, a timer based or manual re-optimization of the LSP path, or a user change to any of the path constraints.

Once the bypass or detour path is setup and is operationally UP, any subsequent changes to the SRLG group membership of an interface the bypass/detour path is using would not be considered by the MPLS/RSVP task at the PLR until the next opportunity the association with the primary LSP path is re-checked. The association is re-checked if the bypass path is re-optimized. Detour paths are not re-optimized and are re-signaled if the primary path is down.

Enabling or disabling srlg-frr only takes effect at the next opportunity the LSP paths are resignaled. The user can wait for the resignal timer to expire or can cause the paths to be resignaled immediately by executing at the ingress LER the **tools perform router mpls resignal** command. Note that in order to force the dynamic bypass LSP to be resignaled using the SRLG constraint of the primary paths it is associated with, it is recommend to first disable dynamic bypass LSPs on the system using the “configure router mpls dynamicbypass” command, then manually resignal the LSP paths using the above tools perform command finally re-enable dynamic bypass LSPs on the system. Before performing this procedure, the user must ensure that no dynamic bypass LSP on the node is active to avoid causing the primary LSP path to go down.

An RSVP interface can belong to a maximum of 64 SRLG groups. The user configures the SRLG groups using the command **config>router>mpls>srlg-group**. The user configures the SRLG groups an RSVP interface belongs to using the **srlg-group** command in the **config>router>mpls>interface** context.

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

Default no srlg-frr

Parameters **strict** — Specifies the name of the SRLG group within a virtual router instance.

Values no slr-frr (default)
srlg-frr (non-strict)
srlg-frr **strict** (strict)

srlg-group

Syntax [**no**] **srlg-group** *group-name* [*group-name*...(**up to 5 max**)]
no srlg-group

Context config>router>interface>if-attribute
config>service>ies>interface>if-attribute
config>service>vprn>interface>if-attribute
config>router>mpls>interface

Description This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.

An interface can belong to up to 64 SRLG groups. However, each single operation of the srlg-group command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

It should be noted that only the SRLGs bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The no form of this command deletes one or more of the SRLG memberships of an interface.

The user can also delete all memberships of an interface by not specifying a group name.

Default no srlg-group

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.

user-srlg-db

Syntax **user-srlg-db** [**enable** | **disable**]

Context config>router>mpls

Description This command enables the use of CSPF by the user SRLG database. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF will query the local SRLG and compute a path after pruning links that are members of the SRLG IDs of the associated primary path. When MPLS makes a request to CSPF for an FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links that are members of the SRLG IDs of the PLR outgoing interface.

If an interface was not entered into the user SRLG database, it is assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information.

The disable keyword disables the use of the user SRLG database. CSPF will then resume queries into the TE database for SRLG membership information. The user SRLG database is maintained.

Default user-srlg-db disable

srlg-database

Syntax [**no**] **srlg-database**

Context config>router>mpls

Description This command provides the context for the user to enter manually the link members of SRLG groups for the entire network at any node that needs to signal LSP paths (for example, a head-end node).

The **no** form of the command deletes the entire SRLG database. CSPF will assume all interfaces have no SRLG membership association if the database was not disabled with the command **config>router>mpls>user-srlg-db disable**.

router-id

Syntax [**no**] **router-id** *ip*

Context config>router>mpls>srlg-database

Description This command provides the context for the user to manually enter the link members of SRLG groups for a specific router in the network. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. Use by CSPF of all interface SRLG membership information of a specific router ID may be temporarily disabled by shutting down the node. If this occurs, CSPF will assume these interfaces have no SRLG membership association.

The **no** form of this command will delete all interface entries under the router ID.

Parameters *ip-address* — 9.Specifies the router ID for this system. This must be the router ID configured under the base router instance, the base OSPF instance or the base IS-IS instance.

interface

Syntax **interface** *ip-address* **srlg-group** *group-name* [*group-name*...(up to 5 max)]
no interface *ip-address* [**srlg-group** *group-name*...(up to 5 max)]

Context config>router>mpls>srlg-database>router-id

Description This command allows the operator to manually enter the SRLG membership information for any link in the network, including links on this node, into the user SRLG database.

An interface can be associated with up to 5 SRLG groups for each execution of this command. The operator can associate an interface with up to 64 SRLG groups by executing the command multiple times.

CSPF will not use entered SRLG membership if an interface is not validated as part of a router ID in the routing table.

The **no** form of the command deletes a specific interface entry in this user SRLG database. The **group-name** must already exist in the **config>router>mpls>srlg-group** context.

Default none

Parameters *ip-int-name* — The name of the network IP interface. An interface name cannot be in the form of an IP address.

srlg-group *group-name* — Specifies the SRLG group name. Up to 1024 group names can be defined in the **config>router>mpls** context. The SRLG group names must be identical across all routers in a single domain.

load-balancing-weight

Syntax **load-balancing-weight** *integer*
no load-balancing-weight

Context config>router>mpls>lsp

Description This command assigns a weight to an MPLS LSP for use in the weighted load-balancing, or weighted ECMP, over MPLS feature.

Parameters *value* — 32-bit integer representing the weight of the LSP.

Values 0 — 4294967295

Default none

MPLS Interface Commands

interface

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

Context config>router>mpls

Description This command specifies MPLS protocol support on an IP interface. No MPLS commands are executed on an IP interface where MPLS is not enabled. An MPLS interface must be explicitly enabled (**no shutdown**).

The **no** form of this command deletes all MPLS commands such as **label-map** which are defined under the interface. The MPLS interface must be shutdown first in order to delete the interface definition. If the interface is not shutdown, the **no interface** *ip-int-name* command does nothing except issue a warning message on the console indicating that the interface is administratively up.

Default **shutdown**

Parameters *ip-int-name* — The name of the network IP interface. 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 to 32 alphanumeric characters.

admin-group

Syntax [**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 in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. IES andVPRN 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.

Default no admin-group

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.

auto-bandwidth-multipliers

Syntax **auto-bandwidth-multipliers sample-multiplier** *number1* **adjust-multiplier** *number2*
no auto-bandwidth-multipliers

Context config>router>mpls

Description This command specifies the number of collection intervals in the adjust interval.

Parameters **sample-multiplier** *number1* **—** Specifies the mulitplier for collection intervals in a sample interval.

Values 1 — 511

Default 1

adjust-multiplier *number2* **—** Specifies the number of collection intervals in the adjust interval.

Values 1 — 16383

Default 288

auto-lsp

Syntax **auto-lsp lsp-template** *template-name* {**policy** *peer-prefix-policy* [**peer-prefix-policy**...(upto 5 max)] | **one-hop**}
no auto-lsp lsp-template *template-name*

Context config>router>mpls

Description This command enables the automatic creation of an RSVP point-to-point LSP to a destination node whose router-id matches a prefix in the specified peer prefix policy. This LSP type is referred to as auto-LSP of type mesh.

The user can associate multiple templates with same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result in the instantiation of a single CSPF computed LSP primary path using the LSP template parameters as long as the prefix corresponds to a router-id for a node in the TE database. This feature does not support the automatic signaling of a secondary path for an LSP. If the user requires the signaling of multiple LSPs to the same destination node, s/he must apply a separate LSP template to the same or different prefix list that contains the same destination node. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID. This feature also does not support the signaling of a non-CSPF LSP. The selection of the **no cspf** option in the LSP template is thus blocked.

Up to five (5) peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or if a new LSP needs to be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed.

The user must perform a **no shutdown** of the template before it takes effect. Once a template is in use, the user must shutdown the template before effecting any changes to the parameters except for those LSP parameters for which the change can be handled with the Make-Before-Break (MBB) procedures. These parameters are **bandwidth** and enabling **fast-reroute** with or without the **hop-limit** or node-protect options. For all other parameters, the user shuts down the template and once a it is added, removed or modified, the existing instances of the LSP using this template are torn down and re-signaled.

The trigger to signal the LSP is when the router with a router-id the matching a prefix in the prefix list appears in the Traffic Engineering database. The signaled LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. It is, however, not available to be used as a provisioned SDP for explicit binding or auto-binding by services.

Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer based re-signaling of the LSP, for TE Graceful Shutdown and for soft pre-emption are supported.

The **one-to-one** option under **fast-reroute**, the LSP Diff-Serv **class-type** and **backup-class-type** parameters are not supported. If **diffserv-te** is enabled under RSVP, the auto-created LSP will still be signaled but with the default LSP class type.

If the **one-hop** option is specified instead of a prefix list, this command enables the automatic signaling of one-hop point-to-point LSPs using the specified template to all directly connected neighbors. This LSP type is referred to as auto-LSP of type one-hop. Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link that comes up to a directly connected IGP neighbor which router-id is discovered. It then instructs MPLS to signals an LSP with a destination address matching the router-id of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. Thus, the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the neighboring router.

An auto-created mesh or one-hop LSP can have egress statistics collected at the ingress LER by adding the **egress-statistics** node configuration into the LSP template. The user can also have **ingress statistics** collected at the egress LER using the same ingress-statistics node in CLI used with a provisioned LSP. The user must specify the full LSP name as signaled by the ingress LER in the RSVP session name field of the Session Attribute object in the received Path message.

The **no** form of this command deletes all LSP signaled using the specified template and prefix policy. When the **one-hop** option is used, it deletes all one-hop LSPs signaled using the specified template to all directly connected neighbors.

Parameters **lsp-template** *template-name* — Specifies an LSP template name up to 32 characters in length.

policy *peer-prefix-policy* — Specifies an peer prefix policy name up to 32 characters in length.

bypass-resignal-timer

Syntax **bypass-resignal-timer** *minutes*
no bypass-resignal-timer

Context config>router>mpls

Description This command triggers the periodic global re-optimization of all dynamic bypass LSP paths associated with RSVP P2P LSP. The operation is performed at each expiry of the user configurable bypass LSP re-signal timer.

When this command is enabled, MPLS makes a request to CSPF for the best path for each dynamic bypass LSP originated on this node. The constraints of the first associated LSP primary path and which originally triggered the signaling of the bypass LSP must be satisfied. In order to do this, MPLS saves the original Path State Block (PSB) of that LSP primary path even if the latter is torn down.

If CSPF returns no path or returns a new path that is equal in terms of cost to the current path, the PSB associations are not updated. If CSPF returns a new path with a different cost from the current one, MPLS will signal it.

Once the new path is successfully signaled, MPLS will evaluate each PSB of each PLR (i.e., each unique avoid-node or avoid-link constraint) associated with the older bypass LSP path to check if the corresponding LSP primary path constraints are still satisfied by the new bypass LSP path. If so, the PSB association is moved to the new bypass LSP.

Each PSB whose constraints are not satisfied remains associated with the older bypass LSP and will be checked at the next background PSB re-evaluation, or at the next timer or manual bypass re-optimization. Furthermore, if the older bypass LSP is SRLG disjoint with a primary path that has the non-strict SRLG constraint while the new bypass LSP is not SRLG disjoint, the PSB association is not moved.

If a specific PLR associated with a bypass LSP is active, the corresponding PSBs remain associated with the older bypass LSP until the Global Revertive Make-Before-Break (MBB) tears down all corresponding primary paths, which will also cause the older bypass LSP to be torn down.

This feature also implements a background PSB re-evaluation task which audits in the background each RSVP session and determines if an existing manual or dynamic bypass is more optimal for that session. If so, it moves the PSB association to this bypass. If the PLR for this session is active, no action is taken and the PSB will be re-examined at the next re-evaluation.

The periodic bypass re-optimization feature evaluates only the PSBs of the PLRs associated with that bypass LSP and only against the new bypass LSP path. The background re-evaluation task will, however, audit all PSBs on the system against all existing manual and dynamic bypass LSPs.

Furthermore, PSBs that have not been moved by the dynamic or manual re-optimization of a bypass LSP, due to the PSB constraints not being met by the new signaled bypass LSP path, will be re-evaluated by the background task against all existing manual and dynamic bypass LSPs.

Finally, the background re-evaluation task will check for PSBs that have requested node-protect bypass LSP but are currently associated with a link-protect bypass LSP, as well as PSBs that requested FRR protection and that have no association. This is in addition to the attempt made at the receipt of a Resv on the protected LSP path such that the association is speed up.

This feature is not supported with inter-area dynamic bypass LSP and bypass LSP protecting S2L paths of a P2MP LSP.

The **no** form of this command disables the periodic global re-optimization of dynamic bypass LSP paths.

Default no bypass-resignal timer. The periodic global re-optimization of dynamic bypass LSP paths is disabled.

Parameters *minutes* — Specifies the time, in minutes, MPLS waits before attempting to re-signal dynamic bypass LSP paths originated on the system.

Values 30 — 10080

cspf-on-loose-hop

Syntax [**no**] **cspf-on-loose-hop**

Context config>router>mpls

Description This command enables the option to do CSPF calculations until the next loose hop or the final destination of LSP on LSR. On receiving a PATH message on LSR and processing of all local hops in the received ERO, if the next hop is loose, then the LSR node will first do a CSPF calculation until the next loose hop. On successful completion of CSPF calculation, ERO in PATH message is modified to include newly calculated intermediate hops and propagate it forward to the next hop. This allows setting up inter-area LSPs based on ERO expansion method.

NOTE: The LSP may fail to set up if this option is enabled on an LSR that is not an area border router and receives a PATH message without proper next loose hop in ERO. The 'cspf-on-loose-hop' configuration is allowed to change dynamically and applied to new LSP setup after change.

Default no cspf-on-loose-hop

srlg-group

Syntax [**no**] **srlg-group** *group-name* [*group-name*...(up to 5 max)]

Context config>router>mpls>interface

Description This command defines the association of RSVP interface to an SRLG group. An interface can belong to up to 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of 5 groups to be specified at a time.

The **no** form of this command deletes the association of the interface to the SRLG group.

Default none

Parameters *group-name* — Specifies the name of the SRLG group within a virtual router instance up to 32 characters.

te-metric

Syntax **te-metric** *value*
no te-metric

Context config>router>mpls>interface

Description This command configures the traffic engineering metric used on the interface. This metric is in addition to the interface metric used by IGP for the shortest path computation.

This metric is flooded as part of the TE parameters for the interface using an opaque LSA or an LSP. The IS-IS TE metric is encoded as sub-TLV 18 as part of the extended IS reachability TLV. The metric value is encoded as a 24-bit unsigned integer. The OSPF TE metric is encoded as a sub-TLV Type 5 in the Link TLV. The metric value is encoded as a 32-bit unsigned integer.

When the use of the TE metric is enabled for an LSP, CSPF will first prune all links in the network topology which do not meet the constraints specified for the LSP path. Such constraints include bandwidth, admin-groups, and hop limit. Then, CSPF will run an SPF on the remaining links. The shortest path among the all SPF paths will be selected based on the TE metric instead of the IGP metric which is used by default.

The TE metric in CSPF LSP path computation can be configured by entering the command **config>router>mpls>lsp>cspf>use-te-metric**.

Note that the TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

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

Default no te-metric

The value of the IGP metric is advertised in the TE metric sub-TLV by IS-IS and OSPF.

Parameters *value* — Specifies the metric value.

Values 1 — 16777215

node-id-in-rro

Syntax [**no**] **node-id-in-rro** <include | exclude>

Context config>router>rsvp>

Description This command enables the option to include node-id sub-object in RRO. Node-ID sub-object propagation is required to provide fast reroute protection for LSP that spans across multiple area domains.

If this option is disabled, then node-id is not included in RRO object.

Default node-id-in-rro exclude

p2p-merge-point-abort-timer

Syntax **p2p-merge-point-abort-timer** [*1.. 65535*] *seconds*
no p2p-merge-point-abort-timer

Context config>router>rsvp

Description This command configures a timer to abort Merge-Point (MP) node procedures for a P2P LSP path. When a value higher than zero is configured for this timer, it will enter into effect anytime this node activates Merge-Point procedures for one or more P2P LSP paths. As soon an ingress interface goes operationally down, the Merge-Point node starts the abort timer. Upon expiry of the timer, MPLS will clean up all P2P LSP paths which ILM is on the failed interface and which have not already received a Path refresh over the bypass LSP.

Default 0 (disabled)

p2mp-merge-point-abort-timer

Syntax **p2mp-merge-point-abort-timer** [*1.. 65535*] *seconds*
no p2mp-merge-point-abort-timer

Context config>router>rsvp

Description This command specifies a configurable timer to abort Merge-Point (MP) node procedures for an S2L of a P2MP LSP instance. When a value higher than zero is configured for this timer, it will enter into effect anytime this node activates Merge-Point procedures for one or more P2MP LSP S2L paths. As soon an ingress interface goes operationally down, the Merge-Point node starts the abort timer. Upon expiry of the timer, MPLS will clean up all P2MP LSP S2L paths which ILM is on the failed interface and which have not already received a Path refresh over the bypass LSP.

Default 0 (disabled)

p2p-active-path-fast-retry

Syntax **p2p-active-path-fast-retry** *seconds* [*1..10*] *seconds*
no p2p-active-path-fast-retry

Context config>router>mpls

Description This command configures a global parameter to allow the user to apply a shorter retry timer for the first try after an active LSP path went down due to a local failure or the receipt of a RESVTear. This timer is used only in the first try. Subsequent retries will continue to be governed by the existing LSP level retry-timer.

Default 0 (disabled)

p2mp-s2-fast-retry

Syntax **p2mp-s2-fast-retry** *seconds* [*1..10*] *seconds*
no p2mp-s2-fast-retry

Context config>router>rsvp

Description This command configures a global parameter to allow the user to apply a shorter retry timer for the first try after an active LSP path went down due to a local failure or the receipt of a RESVTear. This timer is used only in the first try. Subsequent retries will continue to be governed by the existing LSP level retry-timer.

Default 0 (disabled)

preemption-timer

Syntax **preemption-timer** *seconds*
no preemption-timer

Context config>router>rsvp

Description This parameter configures the time in seconds a node holds to a reservation for which it triggered the soft pre-emption procedure.

The pre-empting node starts a separate preemption timer for each pre-empted LSP path. While this timer is on, the node should continue to refresh the Path and Resv for the pre-empted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.

A value of zero means the LSP should be pre-empted immediately; hard pre-empted.

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

Default 300

Parameters *seconds* — Specifies the time, in seconds, of the preemption timer.

Values 0 — 1800 seconds

label-map

Syntax [**no**] **label-map** *in-label*

Context config>router>mpls>interface

Description This command is used on transit routers when a static LSP is defined. The static LSP on the ingress router is initiated using the **config router mpls static-lsp** *lsp-name* command. An *in-label* can be associated with either a **pop** or a **swap** action, but not both. If both actions are specified, the last action specified takes effect.

The **no** form of this command deletes the static LSP configuration associated with the *in-label.*

Parameters *in-label* — Specifies the incoming MPLS label on which to match.

Values 32 — 1023

pop

Syntax [**no**] **pop**

Context config>router>mpls>if>label-map

Description This command specifies that the incoming label must be popped (removed). No label stacking is supported for a static LSP. The service header follows the top label. Once the label is popped, the packet is forwarded based on the service header.

The **no** form of this command removes the **pop** action for the *in-label*.

Default none

shutdown

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

Context config>router>mpls>if>label-map

Description This command disables the label map definition. This drops all packets that match the specified *in-label* specified in the **label-map** *in-label* command.

The **no** form of this command administratively enables the defined label map action.

Default **no shutdown**

swap

Syntax **swap** {*out-label* **| implicit-null-label**} **nexthop** *ip-address*
no swap {*out-label* **| implicit-null-label**}

Context config>router>mpls>interface>label-map

Description This command swaps the incoming label and specifies the outgoing label and next hop IP address on an LSR for a static LSP.

The **no** form of this command removes the swap action associated with the *in-label*.

Default none

Parameters **implicit-null-label** — Specifies the use of the implicit label value for the outgoing label of the swap operation.

out-label — Specifies the label value to be swapped with the in-label. Label values 16 through 1,048,575 are defined as follows:

Label values 16 through 31 are reserved.

Label values 32 through 1,023 are available for static assignment.

Label values 1,024 through 2,047 are reserved for future use.

Label values 2,048 through 18,431 are statically assigned for services.

Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services.

Label values 131,072 through 1,048,575 are reserved for future use.

Values 16 — 1048575

nexthop *ip-address* **—** The IP address to forward to. If an ARP entry for the next hop exists, then the static LSP will be marked operational. If ARP entry does not exist, software will set the operational status of the static LSP to down and continue to ARP for the configured nexthop. Software will continuously try to ARP for the configured nexthop at a fixed interval.

mpls-tp-mep

Syntax [**no**] **mpls-tp-mep**

Context config>router>mpls>interface

Description This command enables the context for a section layer MEP for MPLS-TP on an MPLS interface.

Default none

if-num

Syntax **if-num** *if-num*
no if-num

Context config>router>mpls>interface>mpls-tp-mep

Description This command configures the MPLS-TP interface number for the MPLS interface. This is a 32-bit unsigned integer that is node-wide unique.

Parameters *if-num* — This is a 32-bit value that is unique to the node.

Values 1 — 4,294,967,295

if-num-validation

Syntax **if-num-validation** {**enable|disable**}
no if-num-validation

Context config>router>mpls>interface>mpls-tp-mep

Description The if-num-validation command is used to enable or disable validation of the if-num in LSP Trace packet against the locally configured if-num for the interface over which the LSP Trace packet was received at the egress LER. This is because some 3rd-party implementations may not perform interface validation for unnumbered MPLS-TP interfaces and instead set the if-num in the dsmap TLV to 0. If the value is 'enabled', the node performs the validation of the ingress and egress if-nums received in the LSP echo request messages that ingress on this MPLS-interface. It validates that the message arrives on the interface as identified by the ingress if-num, and is forwarded on the interface as identified by the egress if-num.

If the value is 'disabled', no validation is performed for the ingress and egress if-nums received in the LSP echo request messages that ingress on this MPLS-interface."

Default enable

Parameters **enable —** Enables interface number validation.

disable — Disables interface number validation.

MPLS-TP Commands

mpls-tp

Syntax [**no**] **mpls-tp**

Context config>router>mpls

Description Generic MPLS-TP parameters and MPLS-TP trabsit paths are configured under this context. If a user configures **no mpls**, normally the entire mpls configuration is deleted. However, in the case of mpls-tp, a check is made that there is no other mpls-tp configuration (e.g., services or LSPs using mpls-tp on the node). The mpls-tp context cannot be deleted if MPLS-TP LSPs or SDPs exist on the system.

A **shutdown** of mpls-tp will bring down all MPLS-TP LSPs on the system.

Default no mpls-tp

tp-tunnel-id-range

Syntax **tp-tunnel-id-range** *start-id end-id*
no tp-tunnel-id-range

Context config>router>mpls>mpls-tp

Description This command configures the range of MPLS tunnel IDs reserved for MPLS-TP LSPs. The maximum difference between the start-id and end-id is 4K.

The tunnel ID referred to here is the RSVP-TE tunnel ID. This maps to the MPLS-TP Tunnel Number. There are some cases where the dynamic LSPs may have caused fragmentation to the number space such that contiguous range [*end-id – start-id*] is not available. In these cases, the command will fail.

There are no default values for the *start-id* and *end-id* of the tunnel id range, and they must be configured to enable MPLS-TP.

Default no tunnel-id-range

Parameters *start-id* — Specifies the start ID.

Values 1 — 61440

end-id — Specifies the end ID.

Values 1 — 61440

oam-template

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

Context config>router>mpls>mpls-tp

Description This command creates or edits an OAM template Generally applicable proactive OAM parameters are configured using templates. The top-level template is the OAM template.

Generic MPLS-TP OAM and fault management parameters are configured in the OAM Template.

Proactive CC/CV uses BFD and parameters such as Tx/Rx timer intervals, multiplier and other session/fault management parameters specific to BFD are configured using a BFD Template, which is referenced from the OAM template.

Default no oam-template

Parameters *name* — Specifies a text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes. Named OAM templates are referenced from the MPLS-TP path MEP configuration.

hold-time-down

Syntax **hold-time-down** *timer*
no hold-time-down

Context config>router>mpls>mpls-tp>oam-template

Description This command configures the hold-down dampening timer. It is equivalent to a hold-off timer.

Default no hold-time-down

Parameters *interval* — Specifies the hold-down dampening timer interval.

Values 0 — 5000 deciseconds in 10 ms increments

hold-time-up

Syntax **hold-time-up** *timer*
no hold-time-up

Context config>router>mpls>mpls-tp>oam-template

Description This command configures the hold-up dampening timer. This can be used to provide additional dampening to the state of proactive CC BFD sessions.

Default no hold-time-up

Parameters *interval* — Specifies the hold-up dampening timer interval.

Values 0 — 500 deciseconds, in 100 ms increments

Default 2 seconds

bfd-template

Syntax **bfd-template** *name*
no bfd-template

Context config>router>mpls>mpls-tp>oam-template

Description This command configures a named BFD template to be referenced by an OAM template.

Default no bfd-template

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

Values

protection-template

Syntax **protection-template** *name*
no protection-template

Context config>router>mpls>mpls-tp

Description Protection templates are used to define generally applicable protection parameters for MPLS-TP tunnels. Only linear protection is supported, and so the application of a named template to an MPLS-TP LSP implies that linear protection is used. A protection template is applied under the MEP context of the protect-path of an MPLS-TP LSP.

The protection-template command creates or edits a named protection template.

Default no protection-template

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

revertive

Syntax [**no**] **revertive**

Context config>router>mpls>mpls-tp>protection-template

Description This command configured revertive behavior for MPLS-TP linear protection. The protect-tp-path MEP must be in the shutdown state for of the MPLS-TP LSPs referencing this protection template in order to change the revertve parameter.

Default revertive

wait-to-restore

Syntax **wait-to-restore** *interval*
no wait-to-restore

Context config>router>mpls>mpls-tp>protection-template

Description This command configures the WTR timer. It determines how long to wait until the active path of an MPLS-TP LSP is restored to the working path following the clearing of a defect on the working path. It is appliable to revertive mode, only.

Default no wait-to-restore

Parameters *interval* — Specifies the WTR timer interval.

Values 0 — 720 seconds in 1 second increments

rapid-psc-timer

Syntax **rapid-psc-timer** *interval*
no rapid-psc-timer

Context config>router>mpls>mpls-tp>protection-template

Description This command configures the rapid timer value to be used for protection switching coordination (PSC) packets for MPLS-TP linear protection (RFC 6378).

Default no rapid-psc-timer

Parameters *interval* — Specifies the rapid timer interval.

Values [10, 100, 1000 ms]

Default 10 ms

slow-psc-timer

Syntax **slow-psc-timer** *interval*
no slow-psc-timer

Context config>router>mpls>mpls-tp>protection-template

Description This command configures the slow timer value to be used for protection switching coordination (PSC) packets for MPLS-TP linear protection (RFC 6378).

Default no rapid-psc-timer

Parameters *interval* — Specifies the slow timer interval.

Values [10, 100, 1000 ms]

global-id

Syntax **global-id** *global-id*
no global-id

Context config>router>mpls>mpls-tp

Description This command configures the MPLS-TP Global ID for the node. This is used as the 'from' Global ID used by MPLS-TP LSPs originating at this node. If a value is not entered, the Global ID is taken to be Zero. This is used if the global-id is not configured. If an operator expects that inter domain LSPs will be configured, then it is recommended that the global ID should be set to the local ASN of the node, as configured under config>system. If two-byte ASNs are used, then the most significant two bytes of the global-id are padded with zeros.

In order to change the value of the global-id, config>router>mpls>mpls-tp must be in the shutdown state. This will bring down all of the MPLS-TP LSPs on the node. New values a propagated to the system when a no shutdown is performed.

Default no global-id

Parameters *global-id* — Specifies the global ID for the node.

Values 0 — 4294967295

node-id

Syntax **node-id** *node-id*
no node-id

Context config>router>mpls>mpls-tp

Description This command configures the MPLS-TP Node ID for the node. This is used as the 'from' Node ID used by MPLS-TP LSPs originating at this node. The default value of the node-id is the system interface IPv4 address. The Node ID may be entered in 4-octed IPv4 address format, <a.b.c.d>, or as an unsigned 32 bit integer. Note that it is not treated as a routable IP address from the perspective of IP routing, and is not advertised in any IP routing protocols.

The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value is configured.

Default no node-id

Parameters *node-id* — Specifies the MPLS-TP node ID for the node.

Values <a.b.c.d> or [1— 4294967295]

Default System interface IPv4 address

transit-path

Syntax **transit-path** *path-name*
no transit-path

Context config>router>mpls>mpls-tp

Description This command enables the configuration or editing of an MPLS-TP transit path at an LSR.

Default no transit-path

Parameters *path-name* — Specifies the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

path-id

Syntax **path-id** {**lsp-num l***sp-num* **| working-path | protect-path** [**src-global-id** *src-global-id*] **src-node-id** *src-node-id* **src-tunnel-num** *src-tunnel-num* [**dest-global-id** *dest-global-id*] **dest-node-id** *dest-node-id* [**dest-tunnel-num** *dest-tunnel-num*]}
no path-id

Context config>router>mpls>mpls-tp>transit-path

Description This command configures path ID for an MPLS-TP transit path at an LSR. The path ID is equivalent to the MPLS-TP LSP ID and is used to generate the maintenance entity group intermediate point (MIP) identifier for the LSP at the LSR. A path-id must be configured for on-demand OAM to verify an LSP at the LSR.

The path-id must contain at least the following parameters: **lsp-num, src-node-id, src-global-id, src-tunnel-num, dest-node-id**.

The path-id must be unique on a node. It is recommended that his is also configured to be a globally unique value.

The **no** form of the command removes the path ID from the configuration.

Default no path-id

Parameters *lsp-num* — Specifues the LSP number.

Values 1 — 65535, or **working path**, or **protect-path**. A **working-path** is equivalent to a lsp-num of 1, and a **protect-path** is an lsp-num of 2.

src-global-id — Specifies the source global ID.

Values 0 — 4294967295

src-node-id — Specifies the source node ID.

Values a.b.c.d or 1 — 4294967295

src-tunnel-num — Specifies the source tunnel number.

Values 1 — 61440

dest-global-id — Specifies the destination global ID. If the destination global ID is not entered, then it is set to the same value as the source global ID.

Values 0 — 4294967295

dest-node-id — Specifies the destination node ID.

Values a.b.c.d or 1 — 4294967295

dest-tunnel-num — Specifies the destination tunnel number. If the destination tunnel number is not entered, then it is set to the same value as the source tunnel number.

Values 1 — 61440

forward-path

Syntax [**no**] **forward-path**

Context config>router>mpls>mpls-tp>transit-path

Description This command enables the forward path of an MPLS-TP transit path to be created or edited.

The forward path must be created before the reverse path.

The **no** form of this command removes the forward path. The forward path cannot be removed if a reverse exists.

Default no forward-path

reverse-path

Syntax [**no**] **reverse-path**

Context config>router>mpls>mpls-tp>transit-path

Description This command enables the reverse path of an MPLS-TP reverse path to be created or edited.

The reverse path must be created after the forward path.

The **no** form of this command removes the reverse path. The reverse path must be removed before the forward path.

Default no reverse-path

in-label

Syntax **in-label** *in-label* **out-label** *out-label* **out-link** *if-name* [**next-hop** *next-hop*]
no in-label

Context config>router>mpls>mpls-tp>transit-path>forward-path
config>router>mpls>mpls-tp>transit-path>reverse-path

Description This command configures the label mapping associated with a forward path or reverse path of an MPLS-TP transit path to be configured.

The incoming label, outgoing label and outgoing interface must be configured, using the **in-label**, **out-label** and **out-link** parameters. If the out-link refers to a numbered IP interface, the user may optionally configure the **next-hop** parameter and the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the *out-link* corresponds to the link returned by the system. If they do not correspond, then the path will not come up.

Default no in-label

Parameters *in-label* — Specifies the in label.

Values 32 — 16415

out-label — Specifies the out label.

Values 32 — 16415

if-name — Specifies the name of the outgoing interface use for the path.

next-hop — Specifies the next-hop.

Values a.b.c.d

shutdown

Syntax **[no] shutdown**

Context config>router>mpls>mpls-tp>transit-path

Description This command administratively enables or disables an MPLS-TP transit path.

Default no shutdown

LSP Commands

lsp

Syntax [**no**] **lsp** *lsp-name* [**bypass-only | p2mp-lsp | mpls-tp** *src-tunnel-num*]

Context config>router>mpls

Description This command creates an LSP that is either signaled dynamically by the router, or a statically provisioned MPLS-TP LSP.

When the LSP is created, the egress router must be specified using the **to** command and at least one **primary** or **secondary** path must be specified for signaled LSPs, or at least one working path for MPLS-TP LSPs. All other statements under the LSP hierarchy are optional.

LSPs are created in the administratively down (**shutdown**) state.

The **no** form of this command deletes the LSP. All configuration information associated with this LSP is lost. The LSP must be administratively shutdown and unbound from all SDPs before it can be deleted.

Default none

Parameters *lsp-name* — Name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

bypass-only — Defines an LSP as a manual bypass LSP exclusively. When a path message for a new LSP requests bypass protection, the PLR first checks if a manual bypass tunnel satisfying the path constraints exists. If one if found, the router selects it. If no manual bypass tunnel is found, therouter dynamically signals a bypass LSP in the default behavior. The CLI for this feature includes a knob that provides the user with the option to disable dynamic bypass creation on a per node basis.

p2mp-lsp — Defines an LSP as a point-to-multipoint LSP. The following parameters can be used with a P2MP LSP: adaptive, adspec, cspf, exclude, fast-reroute, from, hop-limit, include, metric, retry-limit, retry-timer, resignal-timer. The following parameters cannot be used with a P2MP LSP: primary, secondary, to, dest-global-id, dest-tunnel-number, working-tp-path, protect-tp-path.

mpls-tp *src-tunnel-num* **—** Defines an LSP as an MPLS-TP LSP. The *src-tunnel-num* is a mandatory create time parameter for mpls-tp LSPs, and has to be assigned by the user based on the configured range of tunnel IDs. The following parameters can only be used with an MPLS-TP LSP: to, dest-global-id, dest-tunnel-number, working-tp-path, protect-tp-path. Other parameters defined for the above LSP types cannot be used.

adaptive

Syntax [**no**] **adaptive**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables the make-before-break functionality for an LSP or LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP.

Default adaptive

adspec

Syntax [**no**] **adspec**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description When enabled, the ADSPEC object will be included in RSVP messages for this LSP. The ADSPEC object is used by the ingress LER to discover the minimum value of the MTU for links in the path of the LSP. By default, the ingress LER derives the LSP MTU from that of the outgoing interface of the LSP path.

Note that a bypass LSP always signals the ADSPEC object since it protects both primary paths which signal the ADSPEC object and primary paths which do not. This means that MTU of LSP at ingress LER may change to a different value from that derived from the outgoing interface even if the primary path has ADSPEC disabled.

Default **no adspec —** No ADSPEC objects are included in RSVP messages.

auto-bandwidth

Syntax [**no**] **auto-bandwidth**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables (and the no form disables) automatic adjustments of LSP bandwidth.

Auto-bandwidth at the LSP level cannot be executed unless **adaptive** is configured in the **config**>**router**>**mpls**>**lsp** context.

Default **no auto-bandwidth**

adjust-down

Syntax **adjust-down** *percent* [**bw** *mbps*]
no adjust-down

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the minimum threshold for decreasing the bandwidth of an LSP based on active measurement of LSP bandwidth.

The **no** form of this command is equivalent to adjust-down 5.

Default **no adjust-down**

Parameters *percent* — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as a percentage of the current bandwidth, for decreasing the bandwidth of the LSP.

Values 1 — 100

Default 5

mbps — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as an absolute bandwidth (mbps), for decreasing the bandwidth of the LSP.

Values 0 — 100000

Default 0

adjust-up

Syntax **adjust-up** *percent* [**bw** *mbps*]
no adjust-up

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the minimum threshold for increasing the bandwidth of an LSP based on active measurement of LSP bandwidth.

The **no** form of this command is equivalent to adjust-up 5.

Default **no adjust-up**

Parameters *percent* — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as a percentage of the current bandwidth, for increasing the bandwidth of the LSP.

Values 1-100

Default 5

mbps — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as an absolute bandwidth (mbps), for increasing the bandwidth of the LSP

Values 0 — 100000

Default 0

fc

Syntax **fc** *fc-name* **sampling-weight** *sampling-weight*
no fc

Context config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the sampling weight.

max-bandwidth

Syntax **max-bandwidth** *mbps*
no max-bandwidth

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the maximum bandwidth that auto-bandwidth allocation is allowed to request for an LSP.

The LSP maximum applies whether the bandwidth adjustment is triggered by normal adjust-interval expiry, the overflow limit having been reached, or manual request.

The **no** form of the command means max-bandwidth is 100 Gbps.

The max-bandwidth must be greater than the min-bandwidth.

Default no max-bandwidth

Parameters *mbps* — Specifies the maximum bandwidth in mbps.

Values 0 — 100000

Default 0

min-bandwidth

Syntax **min-bandwidth** *mbps*
no min-bandwidth

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the minimum bandwidth that auto-bandwidth allocation is allowed to request for an LSP.

The LSP minimum applies whether the bandwidth adjustment is triggered by normal adjust-timer expiry or manual request.

The **no** form of the command means min-bandwidth is zero.

Default **no min-bandwidth**

Parameters *mbps* — Specifies the minimum bandwidth in mbps.

Values 0 — 100000

Default 0

monitor-bandwidth

Syntax [**no**] **monitor-bandwidth**

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command enables the collection and display of auto-bandwidth measurements, but prevents any automatic bandwidth adjustments from taking place.

This command is mutually exclusive with the overflow-limit command.

The **no** form of the command the collection and display of auto-bandwidth measurements.

multipliers

Syntax **multipliers sample-multiplier** *num1* **adjust-multiplier** *num2*
no multipliers

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the sample-multiplier and adjust-multiplier applicable to one particular LSP.

The sample-multiplier configures the number of collection intervals between measurements of the number of bytes that have been transmitted on the LSP. The byte counts include the layer 2 encapsulation of MPLS packets and represent traffic of all forwarding classes and priorities (in-profile vs, out-of-profile) belonging to the LSP. The router calculates the average data rate in each sample interval. The maximum of this average data rate over multiple sample intervals is the measured bandwidth input to the auto-bandwidth adjustment algorithms.

The adjust-multiplier is the number of collection intervals between periodic evaluations by the ingress LER about whether to adjust the LSP bandwidth. The router keeps track of the maximum average data rate of each LSP since the last reset of the adjust-count.

The adjust-multiplier is not allowed to be set to a value less than the sample-multiplier. It is recommended that the adjust-multiplier be a multiple of the sample-multiplier.

The **no** form of this command instructs the system to take the value from the auto-bandwidth-defaults command.

Default **no multipliers**

Parameters *number1* — The number of collection intervals in a sample interval.

Values 1 — 511

Default inherited

number2 — The number of collection intervals in an adjust interval.

Values 1 — 16383

Default inherited

overflow-limit

Syntax

overflow-limit *number* **threshold** *percent* [**bw** *mbps*]
no overflow-limit

Context

config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures overflow-triggered auto-bandwidth adjustment. It sets the threshold at which bandwidth adjustment is initiated due to the configured number of overflow samples having been reached, regardless of how much time remains until the adjust interval ends.

A sample interval is counted as an overflow if the average data rate during the sample interval is higher than the currently reserved bandwidth by at least the thresholds configured as part of this command.

If overflow-triggered auto-bandwidth adjustment is successful the overflow count, maximum average data rate and adjust count are reset. If overflow-triggered auto-bandwidth adjustment fails then the overflow count is reset but the maximum average data rate and adjust count maintain current values.

This command is mutually exclusive with the monitor-bandwidth command.

The **no** form of this command disables overflow-triggered automatic bandwidth adjustment.

Default **no overflow-limit**

Parameters *number* — The number of overflow samples that triggers an overflow auto-bandwidth adjustment attempt.

Values 1 — 10

Default 0 (disabled)

percent — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as a percentage of the current bandwidth, for counting an overflow sample.

Values 1 — 100

Default 0 (disabled)

mbps — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as an absolute bandwidth (Mbps) relative to the current bandwidth, for counting an overflow sample.

Values 1— 100000

Default 0 (disabled)

underflow-limit

Syntax **underflow-limit** *number* **threshold** *percent* [**bw** *mbps*]
no underflow-limit

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures underflow-triggered auto-bandwidth adjustment. An underflow auto-bandwidth adjustment can occur any time during the adjust-interval; it is triggered when the number of consecutive underflow samples reaches the threashold N configured as part of this command. The new bandwidth of the LSP after a successful underflow adjustment is the maximum data rate observed in the last N consecutive underflow samples.

A sample interval is counted as an underflow if the average data rate during the sample interval is lower than the currently reserved bandwidth by at least the thresholds configured as part of this command.

This command is mutually exclusive with the **monitor-bandiwdth** command.

The **no** form of this command disables underflow-triggered automatic bandwidth adjustment.

Default no underflow-limit

Parameters *number* — The number of consecutive underflow samples that triggers an underflow auto-bandwidth adjustment attempt.

Values 0 — 10

Default 0 (disabled)

percent — The minimum difference between the current bandwidth of the LSP and the sampeld data rate, expressed as a percentage of the current bandwidth, for counting an underflow sample.

Values 0 —100

Default 0 (disabled)

mbps — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as an absolute bandwidth (Mbps) relative to the current bandwidth, for counting an underflow sample.

Values 0 —100,000

Default 0 (disabled)

bfd

Syntax bfd

Context config>router>mpls>lsp
config>router>mpls>lsp>primary
config>router>mpls>lsp-template

Description The bfd command creates a context for the configuration of LSP BFD commands.

bfd-enable

Syntax **bfd-enable**
no bfd-enable

Context config>router>mpls>lsp>bfd
config>router>mpls>lsp>primary>bfd
config>router>mpls>lsp-template>bfd

Description This command enables LSP BFD on the LSP. Lsp-bfd must also be configured under **config>router** to enable LSP BFD. The parameters for the BFD session are derived from the named BFD Template, which must have been configured prior to the **bfd-enable** command and associated with the service using the **bfd-template** command.

Default **no bfd-enable**

bfd-template

Syntax **bfd-template** *name*
no bfd-template

Context config>router>mpls>lsp>bfd
config>router>mpls>lsp>primary>bfd
config>router>mpls>lsp-template>bfd

Description This command references a named BFD template to be used by LSP BFD. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, to be used by the BFD session. Templates are configured under the **config>router>bfd** context.

Default **no bfd-template**

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

Default none

lsp-ping-interval

Syntax **lsp-ping-interval** *seconds*
no lsp-ping-interval

Context config>router>mpls>lsp>bfd
config>router>mpls>lsp>primary>bfd
config>router>mpls>lsp-template>bfd

Description This command configures the interval for the periodic LSP ping for LSPs on which **bfd-enable** has been configured. This is used to bootstrap and maintain the LSP BFD session. The default interval is 60 seconds, with a maximum of 300 seconds. A value of 0 disables periodic LSP Ping, such that an LSP Ping containing a bootstrap TLV is only sent when the BFD session is first initialized.

In scaled environments, LSP BFD sessions should use longer timers to reduce the chance of congestion and loading of common resources. Unless required, the **lsp-ping-interva**l should not be set lower than 300 seconds.

Default no lsp-ping interval This sets the periodic LSP Ping interval to a default of 60 seconds.

Parameters *seconds* — Sets the periodic LSP Ping interval.

Values 0, 60 – 300 seconds

Default 60 seconds

bgp-shortcut

Syntax [**no**] **bgp-shortcut**

Context config>router>mpls>lsp

Description This command enables the use of RSVP LSP for IPv4 BGP routes.

bgp-transport-tunnel

Syntax **bgp-transport-tunnel** *include* | *exclude*

Context config>router>mpls>lsp

Description This command allows or blocks RSVP-TE LSP to be used as a transport LSP for BGP tunnel routes.

Default bgp-transport-tunnel include

Parameters *include* — Allows RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop eBGP peers with ASBR to ASBR adjacency.

exclude — Blocks RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop eBGP peers with ASBR to ASBR adjacency.

class-forwarding

Syntax [**no**] **class-forwarding**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables the context to configure class based forwarding parameters for a given LSP or LSP-template.

The **no** form removes any Class-Based Forwarding configuration associated to that LSP or LSP-template.

Note that a change in the Class-Based Forwarding configuration may result in a change of forwarding behavior.

Default **no class-forwarding**

fc

Syntax **fc {be|l2|af|l1|h2|ef|h1|nc}**
no fc [{be|l2|af|l1|h2|ef|h1|nc}]

Context config>router>mpls>lsp>class-forwarding
config>router>mpls>lsp-template>class-forwarding

Description This command assigns a forwarding class to a given LSP or LSP-template. This command can only be passed with a single forwarding class but by passing the command multiple times it is possible to assign multiple forwarding classes (up to 8) to the same LSP or LSP-template.

The **no** form of this command removes the assignment of the forwarding classes from the LSP or LSP-template. It can only be passed with either a single or no forwarding class. If no forwarding class is specified, all the assignments are removed. In the other case, only the assignment of the specified forwarding class is removed.

Note that a change in the Class-Based Forwarding configuration may result in a change of forwarding behavior.

Default **no fc**

default-lsp

Syntax [**no**] **default-lsp**

Context config>router>mpls>lsp>class-forwarding
config>router>mpls>lsp-template>class-forwarding

Description This command assigns the **default-lsp** configuration to a given LSP or LSP-template. The Default LSP is the LSP on which will be forwarded any packet associated to a given class but for which no LSP with the corresponding class explicitly assigned exists.

The **no** form of this command removes the **default-lsp** assignment from the LSP or LSP-template.

Note that a change in the Class-Based Forwarding configuration may result in a change of forwarding behavior.

Default **no default-lsp**

class-type

Syntax **class-type** *ct-number*
no class-type

Context config>router>mpls>lsp
config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

Description This command configures the Diff-Serv Class Type (CT) for an LSP, the LSP primary path, or the LSP secondary path. The path level configuration overrides the LSP level configuration. However, only one CT per LSP path will be allowed as per RFC 4124.

The signaled CT of a dynamic bypass is always be CT0 regardless of the CT of the primary LSP path. The setup and hold priorities must be set to default values, i.e., 7 and 0 respectively. This assumes that the operator configured a couple of TE classes, one which combines CT0 and a priority of 7 and the other which combines CTO and a priority of 0. If not, the bypass LSP will not be signaled and will go into the down state.

The operator cannot configure the CT, setup priority, and hold priority of a manual bypass. They are always signaled with CT0 and the default setup and holding priorities.

The signaled CT and setup priority of a detour LSP must match those of the primary LSP path it is associated with.

If the operator changes the CT of an LSP or of an LSP path, or changes the setup and holding priorities of an LSP path, the path will be torn down and retried.

An LSP which does not have the CT explicitly configured will behave like a CT0 LSP when Diff-Serv is enabled.

If the operator configured a combination of a CT and a setup priority and/or a combination of a CT and a holding priroty for an LSP path that are not supported by the user-defined TE classes, the LSP path will be kept in a down state and an error code will be displayed in the show command output for the LSP path.

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

Default **no class-type**

Parameters *ct-number* — The Diff-Serv Class Type number.

Values 0 – 7

Default 0

bandwidth

Syntax **bandwidth** *rate-in-mbps*

Context config>router>mpls>lsp>primary-p2mp-instance
config>router>mpls>lsp-template

Description This command specifies the amount of bandwidth to be reserved for the P2MP instance.

Parameters *rate-in-mbps* — specifies the bandwidth, in Mbps.

Values 0 — 100000

cspf

Syntax [**no**] **cspf** [*use-te-metric*]

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables Constrained Shortest Path First (CSPF) computation for constrained-path LSPs. Constrained-path LSPs are the ones that take configuration constraints into account. CSPF is also used to calculate the detour routes when fast-reroute is enabled.

Explicitly configured LSPs where each hop from ingress to egress is specified do not use CSPF. The LSP will be set up using RSVP signaling from ingress to egress.

If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.

Default no cspf

Parameters *use-te-metric* — Specifies to use the use of the TE metric for the purpose of the LSP path computation by CSPF.

dest-global-id

Syntax **dest-global-id** *dest-global-id*
no dest-global-id

Context config>router>mpls>lsp

Description This optional command configures the MPLS-TP Global ID of the far end node of the MPLS-TP LSP. This command is only allowed for MPLS-TP LSPs. Global ID values of 0 indicate that the local node's configured global ID is used. If the local global-id is 0, then the dest-global-id must also be 0. The dest-global-id cannot be changed if an LSP is in use by an SDP.

Default 0

Parameters *dest-global-id* — Specifies the destination global ID.

Values 0 — 4294967295

Default 0

dest-tunnel-number

Syntax **dest-tunnel-number** *dest-tunnel-number*
no dest-tunnel-number

Context config>router>mpls>lsp

Description This optional command configures the MPLS-TP tunnel number of the LSP at the far end node of the MPLS-TP LSP. This command is only allowed for MPLS-TP LSPs. If it is not entered, then the system will take the dest-tunnel-number to be the same as the src-tunnel-num for the LSP.

Default The default value is the configured *src-tunnel-num*.

Parameters *dest-tunnel-number* — Specifies the destination tunnel number.

Values 1 — 61440

Default *src-tunnel-number*

working-tp-path

Syntax [**no**] **working-tp-path**

Context config>router>mpls>lsp

Description This command creates or edits the working path for an MPLS-TP LSP. At least one working path (but not more than one working path) must be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, then this is the path that is used as the default working path for the LSP, and it must be created prior to the protect path. The working-tp-path can only be deleted if no protect-tp-path exists for the LSP.

The following commands are applicable to the working-tp-path: **lsp-num, in-label, out-label, mep, shutdown**.

Default no working-tp-path

protect-tp-path

Syntax [**no**] **protect-tp-path**

Context config>router>mpls>lsp

Description This command creates or edits the protect path for an MPLS-TP LSP. At least one working path must exist before a protect path can be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, then this is the path that is used as the default protect path for the LSP. The protect path must be deleted before the wokring path. Only one protect path can be created for each MPLS-TP LSP.

The following commands are applicable to the working-tp-path: **lsp-num, in-label, out-label, mep, shutdown**.

lsp-num

Syntax **lsp-num** *lsp-num*
no lsp-num

Context config>mpls>lsp>working-tp-path
config>mpls>lsp>protect-tp-path

Description This command configures the MPLS-TP LSP Number for the working TP path or the Protect TP Path.

Default no lsp-num

Parameters *lsp-num* — Specifies the LSP number.

Values 1 — 65535

Default 1 for a working path, 2 for a protect path

in-label

Syntax **in-label** *in-label*
no in-label

Context config>mpls>lsp>working-tp-path
config>mpls>lsp>protect-tp-path

Description This command configures the incoming label for the reverse path or the working path or the protect path of an MPLS-TP LSP. MPLS-TP LSPs are bidirectional, and so an incoming label value must be specified for each path.

Default no in-label

Parameters *in-label* — Specifies the in label.

Values 32 — 16415

out-label

Syntax **out-label** *out-label* **out-link** *if-name* [**next-hop** *ip-address*]
no out-label

Context config>mpls>lsp>working-tp-path
config>mpls>lsp>protect-tp-path

Description This command configureds the outgoing label value to use for an MPLS-TP working or protect path. The out-link is the outgoing interface on the node that this path will use, and must be specified. If the out-link refers to a numbered IP interface, the user may optionally configure the **next-hop** parameter and the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the *out-link* corresponds to the link returned by the system. If they do not correspond, then the path will not come up.

Default no out-label

Parameters *out-label* — Specifies the out label.

Values 32 — 16415

if-name — Specifies the interface name.

ip-address — Specifies the IPv4 address in a.b.c.d

mep

Syntax [**no**] **mep**

Context config>mpls>lsp>working-tp-path
config>mpls>lsp>protect-tp-path

Description This command creates or edits an MPLS-TP maintenance entity group (MEG) endpoint (MEP) on and MPLS-TP path. MEPs reporesent the termination point for OAM flowing on the path, as well as linear protection for the LSP. Only one MEP can be configured at each end of the path.

The following commands are applicable to a MEP on an MPLS-TP working or protect path: oam-template, bfd-enable, and shutdown. In addition, a protection-template may be configured on a protect path.

The **no** form of the command removes a MEP from an MPLS-TP path.

mip

Syntax [**no**] **mip**

Context config>router>mpls>lsp>transit-path>forward-path
config>router>mpls>lsp>transit-path>reverse-path

Description This command creates a context for maintenence entity group intermediate point (MIP) parameters for the forward path and the reverse path of an MPLS-TP LSP at an LSR.

Default none

dsmap

Syntax **dsmap** *if-num*
no dsmap

Context config>router>mpls>lsp>working-tp-path>mep
config>router>mpls>lsp>protect-tp-path>mep
config>router>mpls>lsp>transit-path>forward-path>mip
config>router>mpls>lsp>transit-path>reverse-path>mip

Description This command is used to configure the values to use in the DSMAP TLV sent by a node in an LSP Trace echo request for a static MPLS-TP LSP. A node sending a DSMAP TLV will include the in-if-num and out-if-num values. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.

Parameters *if-num* — This is a 32-bit value corresponding to the expected ingress interface if-num used by an MPLS-TP LSP for the next hop downstream. A value of zero means that no interface validation will be performed.

Values 0 — 4,294,967,295

Default 0

oam-template

Syntax **oam-template** *name*
no oam-template

Context config>mpls>lsp>working-tp-path
config>mpls>lsp>protect-tp-path

Description This command applies a OAM template to an MPLS-TP working or protect path. It contains configuration paraeters for proactive OAM mechanisms that can be enabled on the path e.g. BFD. Configuration of an OAM template is optional.

The **no** form of the command removes the OAM template from the path.

Default no oam-template

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

bfd-enable

Syntax **bfd-enable** [**cc** | **cc_cv**]
no bfd-enable

Context config>mpls>lsp>working-tp-path
config>mpls>lsp>protect-tp-path

Description The command associates the operational state of an MPLS-TP path with a BFD session whose control packets flow on the path. The BFD packets are encapsulated in a generic associated channel

(G-ACh) on the path. The timer parameters of the BFD session are taken from the the OAM template of the MEP.

A value of cc means that the BFD session is only used for continuity check of the the MPLS-TP path. In this case, the cc timer parameters of the OAM template apply. A value of cv means that the BFD session is used for both continuity checking and connectivity verification, and the cv timers of the OAM template apply.

This form of the bfd-enable command is only applicable when it is configured under a MEP used on an MPLS-TP working or protect path.

Default no bfd-enable

Parameters **cc | cc_cv —** cc indicates that BFD runs in CC only mode. This mode uses GACh channel type 0x07. cc_cv indicates that BFD runs in combined CC and CV mode. This mode uses channel type 0x22 for MPLS-TP CC packets, and 0x23 for MPLS-TP CV packets.

protection-template

Syntax **protection-template** *name*
no protection-template

Context config>mpls>lsp>protect-tp-path

Description This command applies a protection template name to an MPLS-TP LSP that the protect path is configured under. If the template is applied, then MPLS-TP 1:1 linear protection is enabled on the LSP, using the parameters specified in the named template.

A named protection template can only be applied to the protect path context of an MPLS-TP LSP.

The no form of the command removes the template and thus disables mpls-tp linear protection on the LSP.

Default no protection-template

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

exclude

Syntax [**no**] **exclude** *group-name* [*group-name*...(up to 5 max)]

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command specifies the admin groups to be excluded when an LSP is set up in the primary or secondary contexts. Each single operation of the exclude command allows a maximum of 5 groups to be specified at a time. However, a maximum of 32 groups can be specified per LSP through multiple operations. The admin groups are defined in the **config>router>if-attribute>admin-group** context.

The exclude statement instructs the CSPF algorithm to avoid TE links which belong to any of the specified admin groups. A link which belongs one or more of the specified admin groups is excluded and thus pruned from the TE database before the CSPF computation.

Use the **no** form of the command to remove the exclude command.

Default no exclude

Parameters *group-name* — Specify the existing group-name to be excluded when an LSP is set up.

exclude-node

Syntax [**no**] **exclude-node** *ip-address*

Context config>router>mpls>lsp

Description This command enables the option to include XRO object in the bypass LSP PATH message object. The exclude-node option is required for manual bypass LSP with XRO to FRR protect ABR node in a multi-vendor network depolyment. This command must be configured on the PLR node that protects the ABR node. The ABR node IP address must be configured as exclude-node.

Default no exclude-node

fast-reroute

Syntax **fast-reroute** *frr-method*
no fast-reroute

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command creates a pre-computed detour LSP from each node in the path of the LSP. In case of failure of a link or LSP between two nodes, traffic is immediately rerouted on the pre-computed detour LSP, thus avoiding packet-loss.

When **fast-reroute** is enabled, each node along the path of the LSP tries to establish a detour LSP as follows:

- Each upstream node sets up a detour LSP that avoids only the immediate downstream node, and merges back on to the actual path of the LSP as soon as possible.

  If it is not possible to set up a detour LSP that avoids the immediate downstream node, a detour can be set up to the downstream node on a different interface.
- The detour LSP may take one or more hops (see **hop-limit**) before merging back on to the main LSP path.
- When the upstream node detects a downstream link or node failure, the ingress router switches traffic to a standby path if one was set up for the LSP.

Fast reroute is available only for the primary path. No configuration is required on the transit hops of the LSP. The ingress router will signal all intermediate routers using RSVP to set up their detours. TE must be enabled for fast-reroute to work.

If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.

The **no** form of the **fast-reroute** command removes the detour LSP from each node on the primary path. This command will also remove configuration information about the hop-limit and the bandwidth for the detour routes.

The **no** form of **fast-reroute hop-limit** command reverts to the default value.

Default **no fast-reroute** — When fast-reroute is specified, the default fast-reroute method is one-to-one.

Parameters **one-to-one** — In the one-to-one technique, a label switched path is established which intersects the original LSP somewhere downstream of the point of link or node failure. For each LSP which is backed up, a separate backup LSP is **facility** — This option, sometimes called **many-to-one**, takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created which serves to backup up a set of LSPs. This LSP tunnel is called a bypass tunnel.

The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair (PLR). Naturally, this constrains the set of LSPs being backed-up via that bypass tunnel to those that pass through a common downstream node. All LSPs which pass through the PLR and through this common node which do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

bandwidth

Syntax **bandwidth** *rate-in-mbps*
no bandwidth

Context config>router>mpls>lsp>fast-reroute
config>router>mpls>lsp-template>fast-reroute

Description This command is used to request reserved bandwidth on the detour path. When configuring an LSP, specify the traffic rate associated with the LSP.

When configuring fast reroute, allocate bandwidth for the rerouted path. The bandwidth rate does not need to be the same as the bandwidth allocated for the LSP.

Default no bandwidth — Bandwidth is not reserved for a rerouted path.

Parameters *rate-in-mbps* — Specifies the amount of bandwidth in Mbps to be reserved for the LSP path.

hop-limit

Syntax **hop-limit** *limit*
no hop-limit

Context config>router>mpls>lsp>fast-reroute
config>router>mpls>lsp-template>fast-reroute

Description For fast reroute, how many more routers a detour is allowed to traverse compared to the LSP itself. For example, if an LSP traverses four routers, any detour for the LSP can be no more than ten router hops, including the ingress and egress routers.

Default 16

Parameters *limit* — Specify the maximum number of hops.

Values 0 — 255

node-protect

Syntax [**no**] **node-protect**

Context config>router>mpls>lsp>fast-reroute
config>router>mpls>lsp-template>fast-reroute

Description This command enables or disables node and link protection on the specified LSP. Node protection ensures that traffic from an LSP traversing a neighboring router will reach its destination even if the neighboring router fails.

Default node-protect

from

Syntax **from** *ip-address*

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This optional command specifies the IP address of the ingress router for the LSP. When this command is not specified, the system IP address is used. IP addresses that are not defined in the system are allowed. If an invalid IP address is entered, LSP bring-up fails and an error is logged.

If an interface IP address is specified as the **from** address, and the egress interface of the nexthop IP address is a different interface, the LSP is not signaled. As the egress interface changes due to changes in the routing topology, an LSP recovers if the **from** IP address is the system IP address and not a specific interface IP address.

Only one **from** address can be configured.

Default The system IP address

Parameters *ip-address* — This is the IP address of the ingress router. This can be either the interface or the system IP address. If the IP address is local, the LSP must egress through that local interface which ensures local strictness.

Default System IP address

Values System IP or network interface IP addresses

hop-limit

Syntax **hop-limit** *number*
no hop-limit

Context config>router>mpls>lsp
config>router>mpls>lsp>fast-reroute
config>router>mpls>lsp-template>fast-reroute

Description This command specifies the maximum number of hops that an LSP can traverse, including the ingress and egress routers. An LSP is not set up if the hop limit is exceeded. This value can be changed dynamically for an LSP that is already set up with the following implications:

> If the new value is less than the current number of hops of the established LSP, the LSP is brought down. Software then tries to re-establish the LSP within the new **hop-limit** number. If the new value is equal to or greater than the current number hops of the established LSP, then the LSP is not affected.

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

Default **255**

Parameters *number* — The number of hops the LSP can traverse, expressed as an integer.

Values 2 — 255

Values 0 — 255

igp-shortcut

Syntax **igp-shortcut** [**lfa-protect** | **lfa-only**] [**relative-metric** [*offset*]]
[**no**] **igp-shortcut**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or as a forwarding adjacency for resolving IGP routes.

When the **rsvp-shortcut** or the advertise-tunnel-link option is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router-id of a remote node.

The **lfa-protect** option allows an LSP to be included in both the main SPF and the Loop-Free Alternate (LFA) SPF. For a given prefix, the LSP can be used either as a primary next-hop or as an LFA next-hop, but not both. If the main SPF computation selected a tunneled primary next-hop for a prefix, the LFA SPF will not select an LFA next-hop for this prefix and the protection of this prefix will rely on the RSVP LSP FRR protection. If the main SPF computation selected a direct primary next-hop, then the LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

The **lfa-only** option allows an LSP to be included in the LFA SPF only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next-hop. The LFA SPF will select a an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

When the **relative-metric** option is enabled, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset (instead of the LSP operational metric) when computing the cost of a prefix which is resolved to the LSP. The offset value is optional and it defaults to zero. The minimum net cost for a prefix is one (1) after applying the offset. Note that the TTM continues the show the LSP operational metric as provided by MPLS. In other words, applications such as LDP-over-RSVP (when IGP shortcut is disabled) and BGP and static route shortcuts will continue to use the LSP operational metric.

The **relative-metric** option is mutually exclusive with the **lfa-protect** or the **lfa-only** options. In other words, an LSP with the **relative-metric** option enabled cannot be included in the LFA SPF and vice-versa when the **rsvp-shortcut** option is enabled in the IGP.

Finally, the **relative-metric** option is ignored when forwarding adjacency is enabled in IS-IS or OSPF. In this case, IGP advertises the LSP as a point-to-point unnumbered link along with the LSP operational metric as returned by MPLS and capped to maximum link metric allowed in that IGP. Both the main SPF and the LFA SPFs will use the local IGP database to resolve the routes.

The **no** form of this command disables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or a forwarding adjacency for resolving IGP routes.

Default igp-shortcut. All RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP corresponds to a router-id of a remote node.

Parameters **lfa-protect —** An LSP is included in both the main SPF and the LFA SPF.

lfa-only — An LSP is included in the LFA SPF only.

relative-metric [*offset*] **—** The shortest IGP cost between the endpoints of the LSP plus the configured offset, instead of the LSP operational metric returned by MPLS, is used when calculating the cost of prefix resolved to this LSP. The offset parameter is an integer and is optional. An offset value of zero is used when the relative-metric option is enabled without specifying the offset parameter value.

Values [-10, +10]

least-fill

Syntax [**no**] **least-fill**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables the use of the least-fill path selection method for the computation of the path of this LSP.

When MPLS requests the computation of a path for this LSP, CSPF will find all equal cost shortest paths which satisfy the constraints of this path. Then, CSPF identifies the single link in each of these paths which has the least available bandwidth as a percentage of its maximum reservable bandwidth. It then selects the path which has the largest value of this percentage least available bandwidth figure. CSPF identifies the least available bandwidth link in each equal cost path after it has accounted for the bandwidth of the new requested path of this LSP.

CSPF applies the least-fill path selection method to all requests for a path, primary and secondary, of an LSP for which this option is enabled. The bandwidth of the path can be any value, including zero.

CSPF applies the least-fill criterion separately to each pre-emption priority in the base TE. A higher setup priority path can pre-empt lower holding priority paths.

CSPF also applies the least-fill criterion separately to each Diff-Serv TE class if Diff-Serv TE is enabled on this node. A higher setup priority path can pre-empt lower holding priority paths within a Class Type.

MPLS will re-signal and move the LSP to the new path in the following cases:

- Initial LSP path signaling.
- Re-try of an LSP path after failure.
- Make-before-break (MBB) due to pending soft preemption of the LSP path.
- MBB due to LSP path configuration change, i.e., a user change to bandwidth parameter of primary or secondary path, or a user enabling of fast-reroute option for the LSP.
- MBB of secondary path due to an update to primary path SRLG.
- MBB due to FRR Global Revertive procedures on the primary path.
- Manual re-signaling of an LSP path or of all LSP paths by the user.

During a manual re-signaling of an LSP path, MPLS will always re-signal the path regardless of whether the new path is exactly the same or different than the current path and regardless or whether the metric of the new path is different or not from that of the current path.

During a timer-based re-signaling of an LSP path which has the least-fill option enabled, MPLS will only re-signal the path if the metric of the new path is different than the one of the current path.

The user deletes a specific node entry in this database by executing the no form of this command.

Default no least-fill. The path of an LSP is randomly chosen among a set of equal cost paths.

ldp-over-rsvp

Syntax [**no**] **ldp-over-rsvp** [**include** | **exclude**]

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command configures an LSP so that it can be used by the IGP to calculate its SPF tree.

When the **ldp-over-rsvp** option is also enabled in ISIS or OSPF, the IGP provides LDP with all ECMP IP next-hops and tunnel endpoints that it considers to be the lowest cost path to its destination.

IGP provides only the endpoints which are the closest to the destination in terms of IGP cost for each IP next-hop of a prefix. If this results in more endpoints than the ECMP value configured on the router, it will further prune the endpoints based on the lowest router-id and for the same router-id, it will select lowest interface-index first.

LDP then looks up the tunnel table to select the actual tunnels to the endpoint provided by IGP and further limits the endpoint selection to the ones which are the closest to destination across all the IP next-hops provided by IGP for a prefix. For each remaining endpoint, LDP selects a tunnel in a round-robin fashion until the router ECMP value is reached. Note that for each endpoint, only tunnels with the same lowest metric are candidate and if more than one tunnel qualify, the selection begins with the lowest tunnel-id.

The no form of the command reverts to default operation.

Default ldporsvp include

include

Syntax [**no**] **include** *group-name* [*group-name*...(up to 5max)]

Context config>router>mpls>lsp
config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary
config>router>mpls>lsp-template

Description This command specifies the admin groups to be included when an LSP is set up. Up to 5 groups per operation can be specified, up to 32 maximum. The **include** statement instructs the CSPF algorithm to pick TE links among the links which belong to one or more of the specified admin groups. A link that does not belong to at least one of the specified admin groups is excluded and thus pruned from the TE database before the CSPF computation. However, a link can still be selected if it belongs to one of the groups in a **include** statement but also belongs to other groups which are not part of any **include** statement in the LSP or primary/secondary path configuration. In other words, the **include** statements implements the "include-any" behavior.

The **no** form of the command deletes the specified groups in the specified context.

Default no include

Parameters *group-name* — Specifies admin groups to be included when an LSP is set up.

priority

Syntax **priority setup-priority** *hold-priority*
no priority

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

Description This command enables the soft pre-emption procedures for this LSP path. The operator enables the soft pre-emption mechanism on a specific LSP name by explicitly configuring the setup and holding priorities for the primary path at the 7x50 head-end node. The operator can similarly configure priority values for a secondary path for this LSP name. Different values could be used for the primary and for any of the secondary paths. In the absence of explicit user configuration, the setup priority is internally set to the default value of 7 and the holding priority is set to the default value of 0. Note however that valid user-entered values for these two parameters require that the holding priority be numerically lower than or equal to the setup priority, otherwise pre-emption loops can occur.

Pre-emption is effected when a 7x50 pre-empting node processes a new RSVP session reservation and there is not enough available bandwidth on the RSVP interface, or the Class Type (CT) when Diff-Serv is enabled, to satisfy the bandwidth in the Flowspec object while there exist other session reservations for LSP paths with a strictly lower holding priority (numerically higher holding priority value) than the setup priority of the new LSP reservation. If enough available bandwidth is freed on the link or CT to accommodate the new reservation by pre-empting one or more lower priority LSP paths, the pre-empting node allows temporary overbooking of the RSVP interface and honors the new reservation.

The 7x50 pre-empting node will immediately set the ‘Preemption pending’ flag (0x10) in the IPv4 Sub-Object in the RRO object in the Resv refresh for each of the pre-empted LSP paths. The IPv4 Sub-Object corresponds to the outgoing interface being used by the pre-empting and pre-empted LSP paths. Note however that the bandwidth value in the Flowspec object is not changed. The Resv flag must also be set if the pre-empting node is a merge point for the primary LSP path and the backup bypass LSP or detour LSP and the backup LSP is activated.

When evaluating if enough available bandwidth will be freed, the 7x50 pre-empting node considers the reservations in order from the lowest holding priority (numerically higher holding priority value) to the holding priority just below the setup priority of the new reservation. A new reservation cannot pre-empt a reservation which has a value of the holding priority equal to the new reservation setup priority.

When Diff-Serv is enabled on the pre-empting node and the MAM bandwidth allocation model is used, a new reservation can only pre-empt a reservation in the same Class Type (CT).

LSP paths which were not flagged at the head-end for soft pre-emption will be hard pre-empted. LSP paths with the default holding priority of 0 cannot be pre-empted. LSP paths with zero bandwidth do not pre-empt other LSP paths regardless of the values of the path setup priority and the path holding priority. They can also not be pre-empted.

When evaluating if enough available bandwidth will be freed, the 7x50 pre-empting node considers the reservations in order from the lowest holding priority (numerically higher holding priority) to the holding priority just below the setup priority of the new reservation. There is no specific order in which the reservations in the same holding priority are considered.

The 7x50 pre-empting node starts a preemption timer for each of the pre-empted LSP paths. While this timer is on, the node should continue to refresh the Path and Resv for the pre-empted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.

A 7x50 head-end node upon receipt of the Resv refresh message with the ‘Preemption pending’ flag must immediately perform a make-before-break on the affected adaptive CSPF LSP. Both IGP metric and TE metric based CSPF LSPs are included. If an alternative path that excludes the flagged interface is not found, then the LSP is put on a retry in a similar way to the Global Revertive procedure at a 7x50 head-end node. However, the number of retries and the retry timer are governed by the values of the retry-limit and retry-timer parameters: config>router>mpls>lsp>retry-limit; config>router>mpls>lsp>retry-timer.

Note that MPLS will keep the address list of flagged interfaces for a maximum of 60 seconds (not user-configurable) from the time the first Resv message with the ‘Preemption pending’ flag is received. This actually means that MPLS will request CSPF to find a path that excludes the flagged interfaces in the first few retries until success or until 60 seconds have elapsed. Subsequent retries after the 60 seconds will not exclude the flagged interfaces as it is assumed IGP has converged by then and the Unreserved Bandwidth sub-TLV for that priority, or TE Class, in the TE database will show the updated value taking into account the pre-empting LSP path reservation or a value of zero if overbooked.

If the LSP has a configured secondary standby which is operationally UP, the 7x50 will switch the path of the LSP to it and then start the MBB. If no standby path is available and a secondary non-standby is configured, the 7x50 will start the MBB and signal the path of the secondary. The LSP path will be switched to either the secondary or the new primary, whichever comes up first.

The no form of the command reverts the LSP path priority to the default values and results in setting the setup priority to 7, in setting the holding priority to 0, and in clearing the ‘soft preemption desired’ flag in the RRO in the Resv refresh message.

Default no priority.

Parameters *setup-priority* — The priority of the reservation for this session at setup time.

Values 0 — 7 (0 is the highest priority and 7 is the lowest priority.)

Default 7 — This session does not pre-empt any other session.

holding-priority — The priority of the reservation for this session at pre-emption action.

Values 0 — 7 (0 is the highest priority and 7 is the lowest priority.)

Default 0 — This session does not get pre-empted by any other session.

main-ct-retry-limit

Syntax **main-ct-retry-limit number**
no main-ct-retry-limit

Context config>router>mpls>lsp

Description This command configures the maximum number of retries the LSP primary path should be retried with the LSP Diff-Serv main Class Type (CT).

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new main-ct-retry-limit parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a “shut/no-shut” on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

If the user entered a value of the main-ct-retry-limit parameter that is greater than the value of the LSP retry-limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry-limit. In other words, the meaning of the LSP retry-limit parameter is not changed and always represents the upper bound on the number of retries. The unmapped LSP primary path behavior applies to both CSPF and non-CSPF LSPs.

The **no** form of this command sets the parameter to the default value of zero (0) which means the LSP primary path will retry forever.

Default no main-ct-retry-limit

Parameters *number* — The number of times MPLS will attempt to re-establish the LSP primary path using the Diff-Serv main CT. Allowed values are integers in the range of zero (0) to 10,000, where zero indicates to retry infinitely.

Values 0 — 1000, integer

metric

Syntax [**no**] **metric** *metric*

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command allows the user to override the LSP operational metric with a constant administrative value that will not change regardless of the actual path the LSP is using over its lifetime.

The LSP operational metric will match the metric the active path of this LSP is using at any given time. For a CSPF LSP, this metric represents the cumulative IGP metric of all the links the active path is using. If CSPF for this LSP is configured to use the TE metric, the LSP operational metric is set to the maximum value. For a non-CSPF LSP, the operational metric is the shortest IGP cost to the destination of the LSP.

The LSP operational metric is used by some applications to select an LSP among a set of LSPs that are destined to the same egress router. The LSP with the lowest operational metric will be selected. If more than one LSP with the same lowest LSP metric exists, the LSP with the lowest tunnel index will be selected. The configuration of a constant metric by the user will make sure the LSP always maintains its preference in this selection regardless of the path it is using at any given time. Applications that use the LSP operational metric include LDP-over-RSVP, VPRN auto-bind, and IGP, BGP and static route shortcuts.

The **no** form of this command disables the administrative LSP metric and reverts to the default setting in which the metric value will represent the LSP metric returned by MPLS. The same behavior is obtained if the user entered a metric of value zero (0).

Default `no metric`. The LSP operational metric defaults to the metric retuned by MPLS.

Parameters *metric* — Specifies the integer value which specifies the value of the LSP administrative metric. A value of zero command reverts to the default setting and disables the administrative LSP metric.

Values 0— 16777215

to

Syntax **to** [*ip-address* | **node-id** *[a.b.c.d* | *1...4,294,967,295*]]

Context config>router>mpls>lsp

Description This command specifies the system IP address or MPLS-TP node-id of the egress router for the LSP. This command is mandatory to create an LSP.

An IP address for which a route does not exist is allowed in the configuration. If the LSP signaling fails because the destination is not reachable, an error is logged and the LSP operational status is set to down.

For a non MPLS-TP LSP, the **to** *ip-address* ***must*** be the system IP address of the egress router. If the **to** address does not match the SDP address, the LSP is not included in the SDP definition.

For an MPLS-TP LSP, the **to node-id** may be either in 4-octet IPv4 address format, or a 32bit unsigned integer. This command is mandatory to create an MPLS-TP LSP. Note tha a value of zero is invalid. This to address is used in the MPLS-TP LSP ID, and the MPLS-TP MEP ID for the LSP.

Default No default

Parameters *ip-address* — The system IP address of the egress router.

node-id *a.b.c.d. | 1...4,294,967,295* — 4-octet IPv4 formatted or unsigned 32-bit integer MPLS-TP node-id of the egress router.

propagate-admin-group

Syntax [**no**] propagate-admin-group

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables propagation of session attribute object with resource affinity (C-type 1) in PATH message. If a session attribute with resource affinity is received at an LSR, then it will check the compatibility of admin-groups received in PATH message against configured admin-groups on the egress interface of LSP.

To support admin-group for inter-area LSP, the ingress node must configure propagating admin-groups within the session attribute object. If a PATH message is received by an LSR node that has the **cspf-on-loose** option enabled and the message includes admin-groups, then the ERO expansion by CSPF to calculate the path to the next loose hop will include the admin-group constraints received from ingress node.

If this option is disabled, then the session attribute object without resource affinity (C-Type 7) is propagated in PATH message and CSPF at the LSR node will not include admin-group constraints.

This admin group propagation is supported with a P2P LSP, a P2MP LSP instance, and an LSP template.

The user can change the value of the **propagate-admin-group** option on the fly. A RSVP P2P LSP will perform a Make-Before-Break (MBB) on changing the configuration. A S2L path of an RSVP P2MP LSP will perform a Break-Before-Make on changing the configuration.

Default no propagate-admin-group

vprn-auto-bind

Syntax **vprn-auto-bind** [**include** | **exclude**]

Context config>router>mpls>lsp

config>router>mpls>lsp-template

Description This command determines whether the associated names LSP can be used or no as part of the auto-bind feature for VPRN services. By default a names LSP is available for inclusion to used for the auto-bind feature.

By configuring the command vprn-auto-bind exclude, the associated LSP will not be used by the auto-bind feature within VPRN services.

The **no** form of the command resets the flag backto the default value.

Default include

Parameters **include —** Allows an associated LSPto be used by auto-bin for vprn services

exclude — Disables the use of the associated LSP to be used with the auto-bind feature for VPRN services.

retry-limit

Syntax **retry-limit** *number*
no retry-limit

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This optional command specifies the number of attempts software should make to re-establish the LSP after it has failed LSP. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the LSP path is put into the **shutdown** state.

Use the config router **mpls lsp** *lsp-name* **no shutdown** command to bring up the path after the retry-limit is exceeded.

For P2MP LSP created based on LSP template, all S2Ls must attempt to retry-limit before client application is informed of failure.

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

Default 0 (no limit, retries forever)

Parameters *number* — The number of times software will attempt to re-establish the LSP after it has failed. Allowed values are integers in the range of 0 to 10000 where 0 indicates to retry forever.

Values 0 — 10000

retry-timer

Syntax **retry-timer** *seconds*
no retry-timer

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command configures the time, in seconds, for LSP re-establishment attempts after it has failed. The retry time is jittered to +/- 25% of its nominal value.

For P2MP LSP created based on LSP template, all S2Ls must attempt to retry-limit before client application is informed of failure.

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

Default **30**

Parameters *seconds* — The amount of time, in seconds, between attempts to re-establish the LSP after it has failed. Allowed values are integers in the range of 1 to 600.

Values 1 — 600

revert-timer

Syntax **revert-timer** *timer-value*
no revert-timer

Context config>router>mpls>lsp

Description This command configures a revert timer on an LSP. The timer starts whn the LSP prmary path recovers from a failure. The LSP reverts from a secondary path to the primary path when the timer expires, or when the secondary path fails.

The **no** form of this command cancels any currently outstanding revert timer. If the LSP is up when a no revert-timer is issued, the LSP will revert to the primary path. Otherwise the LSP reverts when the primary path is restored.

Default **no revert-timer**

timer-value — The amount of time, in one minute increments, between attempts to re-establish the LSP after it has failed. Allowed values are integers in the range of 0-4320.

Values 0— 4320

rsvp-resv-style

Syntax **rsvp-resv-style** [*se* | *ff*]

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command specifies the RSVP reservation style, shared explicit (se) or fixed filter (ff). A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration.

Default **se**

Parameters *ff* — Fixed filter is single reservation with an explicit scope. This reservation style specifies an explicit list of senders and a distinct reservation for each of them. A specific reservation request is created for data packets from a particular sender. The reservation scope is determined by an explicit list of senders.

se — Shared explicit is shared reservation with a limited scope. This reservation style specifies a shared reservation environment with an explicit reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders. Because each sender is explicitly listed in the RESV message, different labels can be assigned to different sender-receiver pairs, thereby creating separate LSPs.

shutdown

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

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command disables the existing LSP including the primary and any standby secondary paths.

To shutdown only the primary enter the **config router mpls lsp** *lsp-name* **primary** *path-name* **shutdown** command.

To shutdown a specific standby secondary enter the **config router mpls lsp** *lsp-name* **secondary** *path-name* **shutdown** command. The existing configuration of the LSP is preserved.

Use the **no** form of this command to restart the LSP. LSPs are created in a shutdown state. Use this command to administratively bring up the LSP.

Default **shutdown**

lsp-template

Syntax **[no] lsp-template** *lsp-template-name* **p2mp-lsp**

Context config>router>mpls

Description This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. ‘p2mp-lsp’ keyword is mandatory.

The **no** form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.

Default none

Parameters *lsp-template-name* — Name to identify LSP template. Any LSP template name and LSP name must not be same.

default-path

Syntax **[no] default-path** *path-name*

Context config>router>mpls>lsp-template

Description A default path binding must be provided before LSP template can be used for signaling LSP. LSP template must be shutdown to modify default-path binding.

The **no** form of command should delete path binding.

Default **none**

Parameters *path-name*

lsp-bfd

Syntax **lsp-bfd**
no lsp-bfd

Context config>router

Description This command creates a context for the configuration of LSP BFD parameters.

Default no lsp-bfd

bfd-sessions

Syntax **bfd-sessions** *max-limit*
no bfd-sessions

Context config>router>lsp-bfd

Description This command enables or disables LSP BFD at the tail end of LSPs on the system. It is also used to limit the maximum number of LSP BFD sessions that may be established at the tail-end of LSPs on a node to *max-limit*. It has no impact on the number of LSP BFD sessions that may be configured at the head end.

Default no bfd-sessions : The establishment of LSP BFD sessions by the node at the tail end of LSPs is disabled.

Parameters *max-limit* — The maximum number of LSP BFD sessions at the tail end of LSPs that can be established on a system. The maximum value that can be entered is constrained by the system wide limit for centralized BDF sessions.

Values 1- max, where max is the platform specific limit on centralized BFD sessions.

Primary and Secondary Path Commands

primary

Syntax **primary** *path-name*
no primary

Context config>router>mpls>lsp

Description This command specifies a preferred path for the LSP. This command is optional only if the **secondary** *path-name* is included in the LSP definition. Only one primary path can be defined for an LSP.

Some of the attributes of the LSP such as the bandwidth, and hop-limit can be optionally specified as the attributes of the primary path. The attributes specified in the **primary path** *path-name* command, override the LSP attributes.

The **no** form of this command deletes the association of this *path-name* from the LSP *lsp-name*. All configurations specific to this primary path, such as record, bandwidth, and hop limit, are deleted. The primary path must be shutdown first in order to delete it. The **no primary** command will not result in any action except a warning message on the console indicating that the primary path is administratively up.

Default none

Parameters *path-name* — The case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

secondary

Syntax [**no**] **secondary** *path-name*

Context config>router>mpls>lsp

Description This command specifies an alternative path that the LSP uses if the primary path is not available. This command is optional and is not required if the **config router mpls lsp** *lsp-name* **primary** *path-name* command is specified. After the switch over from the primary to the secondary, the software continuously tries to revert to the primary path. The switch back to the primary path is based on the **retry-timer** interval.

Up to eight secondary paths can be specified. All the secondary paths are considered equal and the first available path is used. The software will not switch back among secondary paths.

Software starts the signaling of all non-standby secondary paths at the same time. Retry counters are maintained for each unsuccessful attempt. Once the retry limit is reached on a path, software will not attempt to signal the path and administratively shuts down the path. The first successfully established path is made the active path for the LSP.

The **no** form of this command removes the association between this *path-name* and *lsp-name*. All specific configurations for this association are deleted. The secondary path must be shutdown first in order to delete it. The **no secondary** *path-name* command will not result in any action except a warning message on the console indicating that the secondary path is administratively up.

Default none

Parameters *path-name* — The case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

adaptive

Syntax [**no**] **adaptive**

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

Description This command enables the make-before-break functionality for an LSP or a primary or secondary LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP.

Default adaptive

backup-class-type

Syntax **backup-class-type ct-number**
no backup-class-type

Context config>router>mpls>lsp>primary

Description This command enables the use of the Diff-Serv backup Class-Type (CT), instead of the Diff-Serv main CT, to signal the LSP primary path when it fails and goes into retry. The Diff-Serv main CT is configured at the LSP level or at the primary path level using the following commands:

config>router>mpls>lsp>class-type *ct-number*

config>router>mpls>lsp>primary>class-type *ct-numbe*r

When a LSP primary path retries due a failure, for example, it fails after being in the UP state, or undergoes any type of Make-Before-Break (MBB), MPLS will retry a new path for the LSP using the main CT. If the first attempt failed, the head-end node performs subsequent retries using the backup CT. This procedure must be followed regardless if the currently used CT by this path is the main or backup CT. This applies to both CSPF and non-CSPF LSPs.

The triggers for using the backup CT after the first retry attempt are:

1. A local interface failure or a control plane failure (hello timeout etc.).
2. Receipt of a PathErr message with a notification of a FRR protection becoming active downstream and/or Receipt of a Resv message with a ‘Local-Protection-In-Use’ flag set. This invokes the FRR Global Revertive MBB.
3. Receipt of a PathErr message with error code=25 (“Notify”) and sub-code=7 (“Local link maintenance required”) or a sub-code=8 (“Local node maintenance required”). This invokes the TE Graceful Shutdown MBB.

4. Receipt of a Resv refresh message with the 'Preemption pending' flag set or a PathErr message with error code=34 ("Reroute") and a value=1 ("Reroute request soft preemption"). This invokes the soft preemption MBB.
5. Receipt of a ResvTear message.
6. A configuration change MBB.
7. The user executing the clear>router>mpls>lsp command.

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new **main-ct-retry-limit** parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a 'shut/no-shut' on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

When the re-signal timer expires, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP even if the new path found by CSPF is identical to the existing one since the idea is to restore the main CT for the primary path. A path with main CT is not found, the LSP remains on its current primary path using the backup CT.

When the user performs a manual re-signal of the primary path, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP as in current implementation.

The **no** form of this command disables the use of the Diff-Serv backup CT.

Default no backup-class-type

Parameters *ct-number* — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0-7, integer

bandwidth

Syntax **bandwidth** *rate-in-mbps*
no bandwidth

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary
config>router>mpls>lsp-template>fast-reroute

Description This command specifies the amount of bandwidth to be reserved for the LSP path.

The **no** form of this command resets bandwidth parameters (no bandwidth is reserved).

Default **no bandwidth** (bandwidth setting in the global LSP configuration)

Parameters *rate-in-mbps* — The amount of bandwidth reserved for the LSP path in Mbps. Allowed values are integers in the range of 1 to 100000.

Values 0 — 100000

exclude

Syntax [**no**] **exclude** *group-name* [*group-name*...(up to 5 max)]

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

Description This command specifies the admin groups to be excluded when an LSP is set up. Up to 5 groups per operation can be specified, up to 32 maximum. The admin groups are defined in the **config>router>if-attribute>admin-group** context.

Use the **no** form of the command to remove the exclude command.

Default no exclude

Parameters *group-name* — Specifies the existing group-name to be excluded when an LSP is set up.

hop-limit

Syntax **hop-limit** *number*
no hop-limit

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary

Description This optional command overrides the **config router mpls lsp** *lsp-name* **hop-limit** command. This command specifies the total number of hops that an LSP traverses, including the ingress and egress routers.

This value can be changed dynamically for an LSP that is already set up with the following implications:

If the new value is less than the current hops of the established LSP, the LSP is brought down. MPLS then tries to re-establish the LSP within the new hop-limit number. If the new value is equal or more than the current hops of the established LSP then the LSP will be unaffected.

The **no** form of this command reverts the values defined under the LSP definition using the **config router mpls lsp** *lsp-name* **hop-limit** command.

Default **no hop-limit**

Parameters *number* — The number of hops the LSP can traverse, expressed as an integer.

Values 2 — 255

record

Syntax [**no**] **record**

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary
config>router>mpls>lsp-template

Description This command enables recording of all the hops that an LSP path traverses. Enabling **record** increases the size of the PATH and RESV refresh messages for the LSP since this information is carried end-to-end along the path of the LSP. The increase in control traffic per LSP may impact scalability.

The **no** form of this command disables the recording of all the hops for the given LSP. There are no restrictions as to when the **no** command can be used. The **no** form of this command also disables the **record-label** command.

Default **record**

record-label

Syntax [**no**] **record-label**

Context config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary
config>router>mpls>lsp-template

Description This command enables recording of all the labels at each node that an LSP path traverses. Enabling the **record-label** command will also enable the **record** command if it is not already enabled.

The **no** form of this command disables the recording of the hops that an LSP path traverses.

Default **record-label**

srlg

Syntax [**no**] **srlg**

Context config>router>mpls>lsp>secondary

Description This command enables the use of the SRLG constraint in the computation of a secondary path for an LSP at the head-end LER.

When this feature is enabled, CSPF includes the SRLG constraint in the computation of the secondary LSP path. This requires that the primary LSP already be established and is up since the head-end LER needs the most current ERO computed by CSPF for the primary path. CSPF would return the list of SRLG groups along with the ERO during primary path CSPF computation. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS/RSVP task will query again CSPF providing the list of SLRG group numbers to be avoided. CSPF prunes all links with interfaces which belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds a path, the secondary is setup. If not, MPLS/RSVP will keep retrying the requests to CSPF.

If CSPF is not enabled on the LSP name, then a secondary path of that LSP which has the SRLG constraint included will be shut down and a specific failure code will indicate the exact reason for the failure in **show>router>mpls>lsp>path>detail** output.

At initial primary LSP path establishment, if primary does not come up or primary is not configured, SRLG secondary will not be signaled and will put to down state. A specific failure code will indicate the exact reason for the failure in **show>router>mpls>lsp>path>detail** output. However, if a non-SRLG secondary path was configured, such as a secondary path with the SRLG option disabled, MPLS/RSVP task will signal it and the LSP use it.

As soon as the primary path is configured and successfully established, MPLS/RSVP moves the LSP to the primary and signals all SRLG secondary paths.

Any time the primary path is re-optimized, has undergone MBB, or has come back up after being down, MPLS/RSVP task checks with CSPF if the SRLG secondary should be re-signaled. If MPLS/RSVP finds that current secondary path is no longer SRLG disjoint, for example, it became ineligible, it puts it on a delayed MBB immediately after the expiry of the retry timer. If MBB fails at the first try, the secondary path is torn down and the path is put on retry.

At the next opportunity the primary goes down, the LSP will use the path of an eligible SRLG secondary if it is UP. If all secondary eligible SLRG paths are Down, MPLS/RSVP will use a non SRLG secondary if configured and UP. If while the LSP is using a non SRLG secondary, an eligible SRLG secondary came back up, MPLS/RSVP will not switch the path of the LSP to it. As soon as primary is re-signaled and comes up with a new SLRG list, MPLS/RSVP will re-signal the secondary using the new SRLG list.

A secondary path which becomes ineligible as a result of an update to the SRLG membership list of the primary path will have the ineligibility status removed on any of the following events:

1. A successful MBB of the standby SRLG path which makes it eligible again.
1. The standby path goes down. MPLS/RSVP puts the standby on retry at the expiry of the retry timer. If successful, it becomes eligible. If not successful after the retry-timer expired or the number of retries reached the number configured under the retry-limit parameter, it is left down.
1. The primary path goes down. In this case, the ineligible secondary path is immediately torn down and will only be re-signaled when the primary comes back up with a new SRLG list.

Once primary path of the LSP is setup and is operationally up, any subsequent changes to the SRLG group membership of an interface the primary path is using would not be considered until the next opportunity the primary path is re-signaled. The primary path may be re-signaled due to a failure or to a make-before-break operation. Make-before-break occurs as a result of a global revertive operation, a timer based or manual re-optimization of the LSP path, or an operator change to any of the path constraints.

One an SRLG secondary path is setup and is operationally UP, any subsequent changes to the SRLG group membership of an interface the secondary path is using would not be considered until the next opportunity secondary path is re-signaled. The secondary path is re-signaled due to a failure, to a re-signaling of the primary path, or to a make before break operation. Make-before break occurs as a result of a timer based or manual re-optimization of the secondary path, or an operator change to any of the path constraints of the secondary path, including enabling or disabling the SRLG constraint itself.

Also, the user-configured include/exclude admin group statements for this secondary path are also checked together with the SRLG constraints by CSPF. Finally, note that enabling SRPG on a secondary standby path that is in the up state will case the path to be torn down and re-signaled using the SRLG constraint.

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

Default no srlg

standby

Syntax [**no**] **standby**

Context config>router>mpls>lsp>secondary

Description The secondary path LSP is normally signaled once the primary path LSP fails. The **standby** keyword ensures that the secondary path LSP is signaled and maintained indefinitely in a hot-standby state. When the primary path is re-established then the traffic is switched back to the primary path LSP.

The **no** form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.

Default none

path-preference

Syntax [**no**] **path-preference** *value*

Context config>router>mpls>lsp>secondary

Description This command enables use of path preference among configured standby secondary paths per LSP. If all standby secondary paths have a default path-preference value then a non-standby secondary path will remain the active path while a standby secondary is available. A standby secondary path configured with highest priority (lowest path-preference value) must be made the active path when the primary is not in use. Path preference can be configured on standby secondary path.

The **no** form of this command resets the path-preference to the default value.

Default 255

Parameters *value* — Specifies an alternate path for the LSP if the primary path is not available,

1–255

LSP Path Commands

hop

Syntax **hop** *hop-index ip-address* {**strict | loose**}
no hop *hop-index*

Context config>router>mpls>path

Description This command specifies the IP address of the hops that the LSP should traverse on its way to the egress router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified then the LSP can choose the best available interface.

Optionally, the LSP ingress and egress IP address can be included as the first and the last hop. A hop list can include the ingress interface IP address, the system IP address, and the egress IP address of any of the hops being specified.

The **no** form of this command deletes hop list entries for the path. All the LSPs currently using this path are affected. Additionally, all services actively using these LSPs are affected. The path must be shutdown first in order to delete the hop from the hop list. The **no hop** *hop-index* command will not result in any action except a warning message on the console indicating that the path is administratively up.

Default none

Parameters *hop-index* — The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 — 1024

ip-address — The system or network interface IP address of the transit router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified then the LSP can choose the best available interface. A hop list can also include the ingress interface IP address, the system IP address, and the egress IP address of any of the specified hops.

loose — This keyword specifies that the route taken by the LSP from the previous hop to this hop can traverse through other routers. Multiple hop entries with the same IP address are flagged as errors. Either the **loose** *or* **strict** keyword must be specified.

strict — This keyword specifies that the LSP must take a direct path from the previous hop router to this router. No transit routers between the previous router and this router are allowed. If the IP address specified is the interface address, then that is the interface the LSP must use. If there are direct parallel links between the previous router and this router and if system IP address is specified, then any one of the available interfaces can be used by the LSP. The user must ensure that the previous router and this router have a direct link. Multiple hop entries with the same IP address are flagged as errors. Either the **loose** *or* **strict** keyword must be specified.

path

Syntax [**no**] **path** *path-name*

Context config>router>mpls

Description This command creates the path to be used for an LSP. A path can be used by multiple LSPs. A path can specify some or all hops from ingress to egress and they can be either **strict** or **loose**. A path can also be empty (no *path-name* specified) in which case the LSP is set up based on IGP (best effort) calculated shortest path to the egress router. Paths are created in a **shutdown** state. A path must be shutdown before making any changes (adding or deleting hops) to the path. When a path is shutdown, any LSP using the path becomes operationally down.

To create a strict path from the ingress to the egress router, the ingress and the egress routers must be included in the path statement.

The **no** form of this command deletes the path and all its associated configuration information. All the LSPs that are currently using this path will be affected. Additionally all the services that are actively using these LSPs will be affected. A path must be **shutdown** and unbound from all LSPs using the path before it can be deleted. The **no path** *path-name* command will not result in any action except a warning message on the console indicating that the path may be in use.

Parameters *path-name* — Specify a unique case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

shutdown

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

Context config>router>mpls>path

Description This command disables the existing LSPs using this path. All services using these LSPs are affected. Binding information, however, is retained in those LSPs. Paths are created in the **shutdown** state.

The **no** form of this command administratively enables the path. All LSPs, where this path is defined as primary or defined as standby secondary, are (re)established.

Default **shutdown**

Static LSP Commands

static-lsp

Syntax [**no**] **static-lsp** *lsp-name*

Context config>router>mpls

Description This command is used to configure a static LSP on the ingress router. The static LSP is a manually set up LSP where the nexthop IP address and the outgoing label (push) must be specified.

The **no** form of this command deletes this static LSP and associated information.

The LSP must be shutdown first in order to delete it. If the LSP is not shut down, the **no static-lsp** *lsp-name* command does nothing except generate a warning message on the console indicating that the LSP is administratively up.

Parameters *lsp-name* — Name that identifies the LSP.

Values Up to 32 alphanumeric characters.

static-lsp-fast-retry

Syntax **static-lsp-fast-retry** *seconds*
no static-lsp-fast-retry

Context config>router>mpls

Description This command specifies the value used as the fast retry timer for a static LSP.

When a static LSP is trying to come up, the MPLS request for the ARP entry of the LSP next-hop may fail when it is made while the next-hop is still down or unavailable. In that case, MPLS starts a retry timer before making the next request. This enhancement allows the user to configure the retry timer, so that the LSP comes up as soon as the next-hop is up.

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

Default no static-fast-retry-timer

Parameters *seconds* — specifies the value, in seconds, used as the fast retry timer for a static LSP.

Values 1-30

push

Syntax **push** {*label* | **implicit-null-label**} **nexthop** *ip-address*
no push {*out-label* | **implicit-null-label**}

Context config>router>mpls>static-lsp

Description This command specifies the label to be pushed on the label stack and the next hop IP address for the static LSP.

The **no** form of this command removes the association of the label to push for the static LSP.

Parameters **implicit-null-label** — Specifies the use of the implicit label value for the push operation.

label — The label to push on the label stack. Label values 16 through 1,048,575 are defined as follows:

Label values 16 through 31 are reserved.

Label values 32 through 1,023 are available for static assignment.

Label values 1,024 through 2,047 are reserved for future use.

Label values 2,048 through 18,431 are statically assigned for services.

Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services.

Label values 131,072 through 1,048,575 are reserved for future use.

Values 16 — 1048575

nexthop *ip-address* — This command specifies the IP address of the next hop towards the LSP egress router. If an ARP entry for the next hop exists, then the static LSP is marked operational. If ARP entry does not exist, software sets the operational status of the static LSP to down and continues to ARP for the configured nexthop. Software continuously tries to ARP for the configured nexthop at a fixed interval.

shutdown

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

Context config>router>mpls>static-lsp

Description This command is used to administratively disable the static LSP.

The **no** form of this command administratively enables the static LSP.

Default **shutdown**

to

Syntax **to** *ip-address*

Context config>router>mpls>static-lsp

Description This command specifies the system IP address of the egress router for the static LSP. When creating an LSP this command is required. For LSPs that are used as transport tunnels for services, the **to** IP address ***must*** be the system IP address. If the **to** address does not match the SDP address, the LSP is not included in the SDP definition.

Parameters *ip-address* — The system IP address of the egress router.

Default none

Point-to-Multipoint MPLS (P2MP) Commands

p2mp-id

Syntax **p2mp-id** *id*

Context config>router>mpls>lsp

Description This command configures the identifier of an RSVP P2MP LSP. An RSVP P2MP LSP is fully identified by the combination of: <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the p2mp sender_template object.

The **p2mp-id** is a 32-bit identifier used in the session object that remains constant over the life of the P2MP tunnel. It is unique within the scope of the ingress LER.

The **no** form restores the default value of this parameter.

Default 0

Parameters *id* — Specifies a P2MP identifier.

Values 0 — 65535

primary-p2mp-instance

Syntax [**no**] **primary-p2mp-instance** *instance-name*

Context config>router>mpls>lsp

Description This command creates the primary instance of a P2MP LSP. The primary instance of a P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSP's. The root, for example a head-end node triggers signaling using one path message per S2L path. The leaf sub-LSP paths are merged at branching points.

Default none

Parameters *instance-name* — Specifies a name that identifies the P2MP LSP instance. The instance name can be up to 32 characters long and must be unique.

s2l-path

Syntax [**no**] **s2l-path** *path-name* **to** *ip-address*

Context config>router>mpls>lsp>primary-inst

Description This command creates a root-to-leaf (S2L) sub-LSP path for the primary instance of a P2MP LSP. The primary instance of a P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSPs. The root, for example, head-end node, triggers signaling using one path message per S2L path. The leaf sub-LSP paths are merged at branching points.

Each S2L sub-LSP is signaled in a separate path message. Each leaf node will respond with its own RESV message. A branch LSR node will forward the path message of each S2L sub-LSP to the downstream LSR without replicating it. It will also forward the RESV message of each S2L sub-LSP to the upstream LSR without merging it with the RESV messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and RESV states.

The S2L paths can be empty paths or can specify a list of explicit hops. The path name must exist and must have been defined using the **config>router>mpls>path** command. The same path name can be re-used by more than one S2L of the primary P2MP instance. However, the **to** keyword must have a unique argument per S2L as it corresponds to the address of the egress LER node.

Default none

Parameters *path-name* — Specifies the name of the path which consists of up to 32 alphanumeric characters.

to *ip-address* **—** Specifies the system IP address of the egress router.

p2mp-resignal-timer

Syntax **p2mp-resignal-timer** *minutes*
no p2mp-resignal-timer

Context config>router>mpls

Description This command configures the re-signal timer for a P2MP LSP instance. MPLS will request CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the P2MP re-signal timer expires. The P2MP re-signal timer is configured separately from the P2P LSP parameter. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful, regardless of the cost of the new S2L path.

The **no** form of this command disables the timer-based re-signaling of P2MP LSPs on this system.

Parameters *minutes* — Specifies the time MPLS waits before attempting to re-signal the P2MP LSP instance.

Values 60 — 10080

RSVP Configuration Commands

Generic Commands

shutdown

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

Context config>router>rsvp
config>router>rsvp>interface

Description This command disables the RSVP protocol instance or the RSVP-related functions for the interface. The RSVP configuration information associated with this interface is retained. When RSVP is administratively disabled, all the RSVP sessions are torn down. The existing configuration is retained.

The **no** form of this command administratively enables RSVP on the interface.

Default **shutdown**

RSVP Commands

rsvp

Syntax [**no**] **rsvp**

Context config>router

Description This command enables the context to configure RSVP protocol parameters. RSVP is not enabled by default and must be explicitely enabled (**no shutdown**).

RSVP is used to set up LSPs. RSVP should be enabled on all router interfaces that participate in signaled LSPs.

The **no** form of this command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance. To suspend the execution and maintain the existing configuration, use the **shutdown** command. RSVP must be shutdown before the RSVP instance can be deleted. If RSVP is not shutdown, the **no rsvp** command does nothing except issue a warning message on the console indicating that RSVP is still administratively enabled.

Default no shutdown

diffserv-te

Syntax **diffserv-te** [**mam** | **rdm**]
no diffserv-te

Context config>router>rsvp

Description This command enabled Diff-Serv Traffic Engineering on the node.

When this command is enabled, IS-IS and OSPF will start advertising available bandwidth for each TE class configured under the diffserv-te node. This command will only have effect if the operator has already enabled traffic engineering at the IS-IS and/or OSPF routing protocol levels:

config>router>isis>traffic-engineering

and/or:

config>router>ospf>traffic-engineering

IGP will advertize for each RSVP interface in the system the available bandwidth in each TE class in the unreserved bandwidth TE parameter for that class. In addition, IGP will continue to advertize the existing Maximum Reservable Link Bandwidth TE parameter to mean the maximum bandwidth that can be booked on a given interface by all classes. The value advertized is adjusted with the link **subscription** *percentage* factor configured in the **config>router>rsvp>interface** context.

The user configures the following parameters for the operation of Diff-Serv:

- Definition of TE classes, TE Class = {Class Type (CT), LSP priority}.
- Mapping of the system forwarding classes to the Diff-Serv Class Type (CT).

- Configuration of the percentage of RSVP interface bandwidth each CT shares, i.e., the Bandwidth Constraint (BC).

When Diff-Serv TE is enabled, the system will automatically enable the Max Allocation Model (MAM) Admission Control Policy. MAM represents the bandwidth constraint model for the admission control of an LSP reservation to a link. This is the only Admission Control Policy supported in this release.

Each CT shares a percentage of the Maximum Reservable Link Bandwidth via the user configured Bandwidth Constraint (BC) for this CT. The Maximum Reservable Link Bandwidth is the link bandwidth multiplied by the RSVP interface subscription factor.

The sum of all BC values across all CTs will not exceed the Maximum Reservable Link Bandwidth. In other words, the following rule is enforced:

SUM (BCc) =< Max-Reservable-Bandwidth, 0 <= c <= 7

An LSP of class-type CTc, setup priority p, holding priority h (h=<p), and bandwidth B is admitted into a link if the following condition is satisfied:

B <= Unreserved Bandwidth for TE-Class[i]

where TE-Class [i] maps to < CTc , p > in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in TE-class [j] = <CTc, h>. Thus, the reserved bandwidth for CTc and the unreserved bandwidth for the TE classes using CTc are updated as follows:

Reserved(CTc) = Reserved(CTc) + B

Unreserved TE-Class [j] = BCc - SUM (Reserved(CTc,q)) for 0<= q <= h

Unreserved TE-Class [i] = BCc - SUM (Reserved(CTc,q)) for 0<= q <= p

The same is done to update the unreserved bandwidth for any other TE class making use of the same CTc. These new values are advertised to the rest of the network at the next IGP-TE flooding.

The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types. It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs.

The RDM model is defined using the following equations:

SUM (Reserved (CTc)) <= BCb,

where the SUM is across all values of c in the range b <= c <= (MaxCT - 1), and BCb is the bandwidth constraint of CTb.

BC0= Max-Reservable-Bandwidth, so that

SUM (Reserved(CTc)) <= Max-Reservable-Bandwidth,

where the SUM is across all values of c in the range 0 <= c <= (MaxCT - 1).

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight pre-emption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

The enabling or disabling of Diff-Serv TE on the system requires the RSVP and MPLS protocol be shutdown.

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

Default no diffserv-te

Parameters **mam** — Defines the default admission control policy for Diff-Serv LSPs.

rdm — Defines Russian doll model for the admission control policy of Diff-Serv LSPs.

class-type-bw

Syntax **class-type-bw ct0** *%-link-bandwidth* **ct1***%-link-bandwidth* **ct2***%-link-bandwidth* **ct3***%-link-bandwidth* **ct4***%-link-bandwidth* **ct5***%-link-bandwidth* **ct6***%-link-bandwidth* **ct7***%-link-bandwidth*
no class-type-bw

Context config>router>rsvp>diffserv-te
config>router>rsvp>interface

Description This command configures the percentage of RSVP interface bandwidth each CT shares, for example, the Bandwidth Constraint (BC).

The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the Maximum Reservable Link Bandwidth TE parameter, for example, the link bandwidth multiplied by the RSVP interface **subscription** *percentage* parameter.

Note this configuration also exists at RSVP interface level and the interface specific configured value overrides the global configured value. The BC value can be changed at any time.

The RSVP interface **subscription** *percentage* parameter is configured in the **config>router>rsvp>interface** context.

The operator can specify the Bandwidth Constraint (BC) for a CT which is not used in any of the TE class definition but that does not get used by any LSP originating or transiting this node.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight pre-emption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

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

Parameters **ct0** (**ct1**/**ct2**/ — **ct7**) % *link-bandwidth* — The Diff-Serv Class Type number. One or more system forwading classes can be mapped to a CT.

Values 0 — 100 %

Default 0

fc

Syntax **fc** *fc-name* **class-type** *ct-number*
no fc fc-name

Context config>router>rsvp>diffserv-te

Description This command maps one or more system forwarding classes to a Diff-Serv Class Type (CT).

The default mapping is shown in the following table.

FC ID	FC Name	FC Designation	Class Type (CT)
7	Network Control	NC	7
6	High-1	H1	6
5	Expedited	EF	5
4	High-2	H2	4
3	Low-1	L1	3
2	Assured	AF	2
1	Low-2	L2	1
0	Best Effort	BE	0

The **no** form of this command reverts to the default mapping for the forwarding class name.

Parameters **class-type** *ct-number* — The Diff-Serv Class Type number. One or more system forwading classes can be mapped to a CT.

Values 0 — 7

te-class

Syntax **te-class** *te-class-number* **class-type** *ct-number* **priority** *priority*
no te-class te-class-number

Context config>router>rsvp>diffserv-te

Description This command configures a traffic engineering class. A TE class is defined as:

TE Class = {Class Type (CT), LSP priority}

Eight TE classes are supported. There is no default TE class once Diff-Serv is enabled. The user has to explicitly define each TE class.

When when Diff-Serv is disabled there will be an internal use of the default CT (CT0) and eight preemption priorities as shown in the following table.

Class Type (CT internal)	LSP Priority
0	7
0	6
0	5
0	4
0	3
0	2
0	1
0	0

The **no** form of this command deletes the TE class.

Parameters **te-class** *te-class-number* — The traffic engineering class number.

Values 0 — 7

class-type *ct-number* — The Diff-Serv Class Type number. One or more system forwading classes can be mapped to a CT.

Values 0 — 7

priority *priority* — The LSP priority.

Values 0 — 7

gr-helper

Syntax **gr-helper** [**enable** | **disable**]

Context config>router>rsvp>if

Description This command enables the RSVP Graceful Restart Helper feature.

The RSVP-TE Graceful Restart helper mode allows the SR OS based system (the helper node) to provide another router that has requested it (the restarting node) a grace period, during which the system will continue to use RSVP sessions to neighbors requesting the grace period. This is typically used when another router is rebooting its control plane but its forwarding plane is expected to continue to forward traffic based on the previously available Path and Resv states.

The user can enable Graceful Restart helper on each RSVP interface separately. When the GR helper feature is enabled on an RSVP interface, the node starts inserting a new Restart_Cap Object in the Hello packets to its neighbor. The restarting node does the same and indicates to the helper node the desired Restart Time and Recovery Time.

The GR Restart helper consists of a couple of phases. Once it loses Hello communication with its neighbor, the helper node enters the Restart phase. During this phase, it preserves the state of all RSVP sessions to its neighbor and waits for a new Hello message.

Once the Hello message is received indicating the restarting node preserved state, the helper node enters the recovery phase in which it starts refreshing all the sessions that were preserved. The restarting node will activate all the stale sessions that are refreshed by the helper node. Any Path state which did not get a Resv message from the restarting node once the Recovery Phase time is over is considered to have expired and is deleted by the helper node causing the proper Path Tear generation downstream.

The duration of the restart phase (recovery phase) is equal to the minimum of the neighbor's advertised Restart Time (Recovery Time) in its last Hello message and the locally configured value of the max-restart (max-recovery) parameter.

When GR helper is enabled on an RSVP interface, its procedures apply to the state of both P2P and P2MP RSVP LSP to a neighbor over this interface.

Default disable

graceful-shutdown

Syntax [**no**] **graceful-shutdown**

Context config>router>rsvp
config>router>rsvp>interface

Description This command initiates a graceful shutdown of the specified RSVP interface or all RSVP interfaces on the node if applied at the RSVP level. These are referred to as maintenance interface and maintenance node, respectively.

To initiate a graceful shutdown the maintenance node generates a PathErr message with a specific error sub-code of Local Maintenance on TE Link required for each LSP that is exiting the maintenance interface.

The node performs a single make-before-break attempt for all adaptive CSPF LSPs it originates and LSP paths using the maintenance interfaces. If an alternative path for an affected LSP is not found, then the LSP is maintained on its current path. The maintenance node also tears down and re-signals any detour LSP path using listed maintenance interfaces as soon as they are not active.

The maintenance node floods an IGP TE LSA/LSP containing Link TLV for the links under graceful shutdown with Traffic Engineering metric set to 0xffffffff and Unreserved Bandwidth parameter set to zero (0).

A head-end LER node, upon receipt of the PathErr message performs a single make-before-break attempt on the affected adaptive CSPF LSP. If an alternative path is not found, then the LSP is maintained on its current path.

A node does not take any action on the paths of the following originating LSPs after receiving the PathErr message:

a. An adaptive CSPF LSP for which the PathErr indicates a node address in the address list and the node corresponds to the destination of the LSP. In this case, there are no alternative paths which can be found.

b. An adaptive CSPF LSP whose path has explicit hops defined using the listed maintenance interface(s)/node(s).

c. A CSPF LSP with the adaptive option disabled and which current path is over the listed maintenance interfaces in the PathErr message. These are not subject to make-before-break.

d. A non CSPF LSP which current path is over the listed maintenance interfaces in the PathErr message.

The head-end LER node upon receipt of the updates IPG TE LSA/LSP for the maintenance interfaces updates the TE database. This information will be used at the next scheduled CSPF computation for any LSP which path may traverse any of the maintenance interfaces.

The **no** form of the command disables the graceful shutdown operation at the RSVP interface level or at the RSVP level. The configured TE parameters of the maintenance links are restored and the maintenance node floods the links.

Default none

gr-helper-time

Syntax **gr-helper-time max-recovery** *recovery-interval* [*1..1800*] *seconds* **max-restart** *restart-interval*
no gr-helper-time

Context config>router>rsvp

Description This command configures the local values for the max-recovery and the max-restart intervals used in the RSVP Graceful Restart Helper feature.

The values are configured globally in RSVP but separate instances of the timers are applied to each RSVP interface that has the RSVP Graceful Restart Helper enabled.

The **no** version of this command re-instates the default value for the delay timer.

Parameters *recovery-interval* — Specifies the max recovery interval value in seconds.

Values 1 — 1800

Default 300

restart-interval — Specifies the max restart interval value in seconds.

Values 1 — 300

Default 120

implicit-null-label

Syntax [**no**] **implicit-null-label**
implicit-null-label

Context config>router>rsvp

Description This command enables the use of the implicit null label.

Signalling the IMPLICIT NULL label value for all RSVP LSPs can be enabled for which this node is the egress LER. RSVP must be shutdown before being able to change this configuration option.

The egress LER does not signal the implicit null label value on P2MP RSVP LSPs. However, the Penultimate Hop Popping (PHP) node can honor a resv message with the label value set to the implicit null.

The **no** form of this command disables the signaling of the implicit null label.

Default no implicit-null-label

keep-multiplier

Syntax [**no**] **keep-multiplier** *number*
no keep-multiplier

Context config>router>rsvp

Description The **keep-multiplier** *number* is an integer used by RSVP to declare that a reservation is down or the neighbor is down.

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

Default 3

Parameters *number* — The **keep-multiplier** value.

Values 1 — 255

refresh-reduction-over-bypass

Syntax **refresh-reduction-over-bypass** [**enable** | **disable**]

Context config>router>rsvp

Description This command enables the refresh reduction capabilities over all bypass tunnels originating on this PLR node or terminating on this Merge Point (MP) node.

By default, this is disabled. Since a bypass tunnel may merge with the primary LSP path in a node downstream of the next-hop, there is no direct interface between the PLR and the MP node and it is possible the latter will not accept summary refresh messages received over the bypass.

When disabled, the node as a PLR or MP will not set the “Refresh-Reduction-Capable” bit on RSVP messages pertaining to LSP paths tunneled over the bypass. It will also not send Message-ID in RSVP messages. This effectively disables summary refresh.

Default disable

rapid-retransmit-time

Syntax **rapid-retransmit-time** *hundred-milliseconds*
no rapid-retransmit-time

Context config>router>rsvp

Description This command defines the value of the Rapid Retransmission Interval. It is used in the re-transmission mechanism to handle unacknowledged message_id objects and is based on an exponential back-off timer.

Re-transmission interval of a RSVP message with the same message_id = 2 * rapid-retransmit-time interval of time.

The node stops re-transmission of unacknowledged RSVP messages:

- If the updated back-off interval exceeds the value of the regular refresh interval.
- If the number of re-transmissions reaches the value of the **rapid-retry-limit** parameter, whichever comes first.

The Rapid Retransmission Interval must be smaller than the regular refresh interval configured in **config>router>rsvp>refresh-time**.

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

Default 5

Parameters *hundred-milliseconds* — Specifies the rapid retransmission interval.

Values 1 – 100, in units of 100 msec.

rapid-retry-limit

Syntax **rapid-retry-limit** *number*
no rapid-retry-limit

Context config>router>rsvp

Description This command is used to define the value of the Rapid Retry Limit. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first.

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

Default 3

Parameters *number* — Specifies the value of the Rapid Retry Limit.

Values 1 – 6, integer values

refresh-time

Syntax **refresh-time** *seconds*
no refresh-time

Context config>router>rsvp

Description The **refresh-time** controls the interval, in seconds, between the successive Path and Resv refresh messages. RSVP declares the session down after it misses **keep-multiplier** *number* consecutive refresh messages.

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

Default 30 seconds

Parameters *seconds* — The refresh time in seconds.

Values 1 — 65535

te-threshold-update

Syntax [**no**] **te-threshold-update**

Context config>router>rsvp

Description This command is used to control threshold-based IGP TE updates. The **te-threshold-update** command must enable IGP TE update based only on bandwidth reservation thresholds per interface and must block IGP TE update on bandwidth changes for each reservation. Threshold levels can be defined using the **te-up-threshold** and **te-down-threshold** commands at the global RSVP or per-interface level.

The **no** form of this command should reset te-threshold-update to the default value and disable threshold based update.

Default no te-threshold-update

on-cac-failure

Syntax [**no**] **on-cac-failure**

Context config>router>rsvp>te-threshold-update

Description This command is used to enable a CAC failure-triggered IGP update.

The **no** form of this command should reset on-cac-failure to the default value and disable the CAC failure-triggered IGP update.

Default no on-cac-failure

update-timer

Syntax **update-timer** *seconds*
no update-timer

Context config>router>rsvp>te-threshold-update

Description This command is to control timer-based IGP TE updates. Timer-based IGP updates can be enabled by specifying a non-zero time value. Default value of update-timer is 0.

The **no** form of this command should reset update-timer to the default value and disable timer-based IGP update.

Default no update-timer (time - 0 seconds)

Parameters *seconds* — The time in seconds.

Values 0-300

te-up-threshold

Syntax **te-up-threshold** *threshold-level* [*threshold-level*...(up to 16 max)]
no te-up-threshold

Context config>router>rsvp
config>router>rsvp>interface

Description This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The **te-threshold-update** command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported.

Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.

The **no** form of this command resets te-up-threshold to its default value.

Default 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100

Parameters *threshold-level* — Integer value

Values 0 — 100

te-down-threshold

Syntax **te-down-threshold** *threshold-level* [*threshold-level*...(up to 16 max)]
no te-down-threshold

Context config>router>rsvp
config>router>rsvp>interface

Description This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The **te-threshold-update** command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported.

Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.

The **no** form of this command resets te-down-threshold to its default value.

Default 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0

Parameters *threshold-level* — Integer value

Values 0 — 100

Interface Commands

interface

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

Context config>router>rsvp

Description This command enables RSVP protocol support on an IP interface. No RSVP commands are executed on an IP interface where RSVP is not enabled.

The **no** form of this command deletes all RSVP commands such as **hello-interval** and **subscription**, which are defined for the interface. The RSVP interface must be **shutdown** it can be deleted. If the interface is not shut down, the **no interface** *ip-int-name* command does nothing except issue a warning message on the console indicating that the interface is administratively up.

Default **shutdown**

Parameters *ip-int-name* — The name of the network IP interface. 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.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash | hash2**]
no authentication-key

Context config>router>rsvp>interface

Description This command specifies the authentication key to be used between RSVP neighbors to authenticate RSVP messages. Authentication uses the MD-5 message-based digest.

When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface.

A node maintains a security association using one authentication key for each interface to a neighbor. The following items are stored in the context of this security association:

- The HMAC-MD5 authentication algorithm.
- Key used with the authentication algorithm.
- Lifetime of the key. The user-entered key is valid until the user deletes it from the interface.
- Source Address of the sending system.
- Latest sending sequence number used with this key identifier.

A router RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an integrity object which also contains a flags field, a key identifier field, and a sequence number field. The

RSVP sender complies to the procedures for RSVP message generation in RFC 2747, *RSVP Cryptographic Authentication*.

A RSVP receiver uses the key together with the authentication algorithm to process received RSVP messages.

When a PLR node switches the path of the LSP to a bypass LSP, it does not send the Integrity object in the RSVP messages sent over the bypass tunnel. If the PLR receives an RSVP message with an Integrity object, it will perform the digest verification for the key of the interface over which the packet was received. If this fails, the packet is dropped. If the received RSVP message is a RESV message and does not have an Integrity object, then the PLR node will accept it only if it originated from the MP node.

An MP node will accept RSVP messages received over the bypass tunnel with and without the Integrity object. If an Integrity object is present, the proper digest verification for the key of the interface over which the packet was received is performed. If this fails, the packet is dropped.

The MD5 implementation does not support the authentication challenge procedures in RFC 2747.

The **no** form of this command disables authentication.

Default **no authentication-key** - The authentication key value is the null string.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

hash-key — The hash key. The key can be any combination of up 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (“ ”)

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

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

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

auth-keychain

Syntax **auth-keychain** *name*

Context config>router>rsvp>interface

Description This command configures an authentication keychain to use for authentication of protocol messages sent and received over the associated interface. The keychain must include a valid entry to properly authenticate protocol messages, including a key, specification of a supported authentication algorithm, and beginning time. Each entry may also include additional options to control the overall lifetime of each entry to allow for the seamless rollover of without affecting the protocol adjacencies.

The **no** form of the auth-keychain command removes the association between the routing protocol and any keychain currently used.

Default no auth-keychain

Parameters name — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>router>rsvp>interface

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated RSVP interface. This causes RSVP to register the interface with the BFD session on that interface.

The user configures the BFD session parameters, such as, **transmit-interval**, **receive-interval**, and **multiplier**, under the IP interface in the **config>router> interface>bfd** context.

Note that it is possible that the BFD session on the interface was started because of a prior registration with another protocol, for example, OSPF or IS-IS.

The registration of an RSVP interface with BFD is performed at the time of neighbor gets its first session. This means when this node sends or receives a new Path message over the interface. If however the session did not come up, due to not receiving a Resv for a new path message sent after the maximum number of re-tries, the LSP is shutdown and the node de-registers with BFD. In general, the registration of RSVP with BFD is removed as soon as the last RSVP session is cleared.

The registration of an RSVP interface with BFD is performed independent of whether RSVP hello is enabled on the interface or not. However, hello timeout will clear all sessions towards the neighbor and RSVP de-registers with BFD at clearing of the last session.

Note that an RSVP session is associated with a neighbor based on the interface address the path message is sent to. If multiple interfaces exist to the same node, then each interface is treated as a separate RSVP neighbor. The user will have to enable BFD on each interface and RSVP will register with the BFD session running with each of those neighbors independently

Similarly the disabling of BFD on the interface results in removing registration of the interface with BFD.

When a BFD session transitions to DOWN state, the following actions are triggered. For RSVP signaled LSPs, this triggers activation of FRR bypass/detour backup (PLR role), global revertive (head-end role), and switchover to secondary if any (head-end role) for affected LSPs with FRR enabled. It triggers switchover to secondary if any and scheduling of re-tries for signaling the primary path of the non-FRR affected LSPs (head-end role).

The **no** form of this command removes BFD from the associated RSVP protocol adjacency.

Default no bfd-enable

hello-interval

Syntax **hello-interval** *milli-seconds*
no hello-interval

Context config>router>rsvp>interface

Description This command configures the time interval between RSVP hello messages.

RSVP hello packets are used to detect loss of RSVP connectivity with the neighboring node. Hello packets detect the loss of neighbor far quicker than it would take for the RSVP session to time out based on the refresh interval. After the loss of the of number keep-multiplier consecutive hello packets, the neighbor is declared to be in a down state.

The **no** form of this command reverts to the default value of the hello-interval. To disable sending hello messages, set the value to zero.

Default **3000** milliseconds

Parameters *milli-seconds* — Specifies the RSVP hello interval in milliseconds, in multiples of 1000. A 0 (zero) value disables the sending of RSVP hello messages.

Values 0 — 60000 milliseconds (in multiples of 1000)

implicit-null-label

Syntax **implicit-null-label** [**enable** | **disable**]
no implicit-null-label

Context config>router>rsvp>interface

Description This command enables the use of the implicit null label over a specific RSVP interface.

All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, then the incoming interface for the path refresh message over the bypass dictates if the packet will use the implicit null label or not. The same for a 1-to-1 detour LSP.

The user must shutdown the RSVP interface before being able to change the implicit null configuration option.

The **no** form of this command returns the RSVP interface to use the RSVP level configuration value.

Default disable

Parameters **enable —** This parameter enables the implicit null label.

disable — This parameter disables the implicit null label.

refresh-reduction

Syntax [**no**] **refresh-reduction**

Context config>router>rsvp>interface

Description This command enables the use of the RSVP overhead refresh reduction capabilities on this RSVP interface.

When this option is enabled, a node will enable support for three capabilities. It will accept bundles RSVP messages from its peer over this interface, it will attempt to perform reliable RSVP message delivery to its peer, and will use summary refresh messages to refresh path and resv states. The reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled. The other two capabilities are enabled immediately.

A bundle message is intended to reduce overall message handling load. A bundle message consists of a bundle header followed by one or more bundle sub-messages. A sub-message can be any regular RSVP message except another bundle message. A node will only process received bundled RSVP messages but will not generate them.

When reliable message delivery is supported by both the node and its peer over the RSVP interface, an RSVP message is sent with a message_id object. A message_id object can be added to any RSVP message when sent individually or as a sub-message of a bundled message.

if the sender sets the ack_desired flag in the message_id object, the receiver acknowledges the receipt of the RSVP message by piggy-backing a message_ack object to the next RSVP message it sends to its peer. Alternatively, an ACK message can also be used to send the message_ack object. In both cases, one or many message_ack objects could be included in the same message.

The router supportsthe sending of separate ACK messages only but is capable of processing received message_ack objects piggy-backed to hop-by-hop RSVP messages, such as path and resv.

The router sets the ack_desired flag only in non refresh RSVP messages and in refresh messages which contain new state information.

A retransmission mechanism based on an exponential backoff timer is supported in order to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The rapid-retransmit-time is referred to as the rapid retransmission interval as it must be smaller than the regular refresh interval configured in the **config>router>rsvp>refresh-time** context. There is also a maximum number of retransmissions of an unacknowledged RSVP message rapid-retry-limit. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first. These two parameters are configurable globally on a system in the **config>router>rsvp** context.

Refresh summary consists of sending a summary refresh message containing a message_id list object. The fields of this object are populated each with the value of the message_identifier field in the message_id object of a previously sent individual path or resv message. The summary refresh message is sent every refresh regular interval as configured by the user using the refresh-time command in the **config>router>rsvp** context. The receiver checks each message_id object against the saved path and resv states. If a match is found, the state is updated as if a regular path or resv refresh message was received from the peer. If a specific message_identifier field does not match, then the node sends a message_id_nack object to the originator of the message.

The above capabilities are referred to collectively as “refresh overhead reduction extensions”. When the refresh-reduction is enabled on an RSVP interface, the node indicates this to its peer by setting a “refresh-reduction-capable” bit in the flags field of the common RSVP header. If both peers of an RSVP interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP messages from the peer on the interface. As soon as this bit is cleared, the router stops sending summary refresh messages. If a peer did not set the “refresh-reduction-capable” bit, a node does not attempt to send summary refresh messages.

However, if the peer did not set the “refresh-reduction-capable” bit, a node, with refresh reduction enabled and reliable message delivery enabled, will still attempt to perform reliable message delivery with this peer. If the peer does not support the message_id object, it returns an error message “unknown object class”. In this case, the node retransmits the RSVP message without the message_id object and reverts to using this method for future messages destined to this peer. The RSVP Overhead Refresh Reduction is supported with both RSVP P2P LSP path and the S2L path of an RSVP P2MP LSP instance over the same RSVP instance.

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

Default no refresh-reduction

reliable-delivery

Syntax [**no**] **reliable-delivery**

Context config>router>rsvp>interface>refresh-reduction

Description This command enables reliable delivery of RSVP messages over the RSVP interface. When refresh-reduction is enabled on an interface and reliable-delivery is disabled, then the router will send a message_id and not set ACK desired in the RSVP messages over the interface. Thus 7750 does not expect an ACK and but will accept it if received. The node will also accept message ID and reply with an ACK when requested. In this case, if the neighbor set the “refresh-reduction-capable” bit in the flags field of the common RSVP header, the node will enter summary refresh for a specific message_id it sent regardless if it received an ACK or not to this message from the neighbor.

Finally, when ‘reliable-delivery’ option is enabled on any interface, RSVP message pacing is disabled on all RSVP interfaces of the system, for example, the user cannot enable the msg-pacing option in the **config>router>rsvp** context, and error message is returned in CLI. Conversely, when the msg-pacing option is enabled, the user cannot enable the reliable delivery option on any interface on this system. An error message will also generated in CLI after such an attempt.

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

Default no reliable-delivery

subscription

Syntax **subscription** *percentage*
no subscription

Context config>router>rsvp>interface

Description This command configures the percentage of the link bandwidth that RSVP can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface.

When the **subscription** is set to zero, no new sessions are permitted on this interface. If the *percentage* is exceeded, the reservation is rejected and a log message is generated.

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

Default **100**

Parameters *percentage* — The percentage of the interface's bandwidth that RSVP allows to be used for reservations.

Values 0 — 1000

te-up-threshold

Syntax **te-up-threshold** *threshold-level* [*threshold-level*...(up to 16 max)]
no te-up-threshold

Context config>router>rsvp
config>router>rsvp>interface

Description This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The **te-threshold-update** command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported.

Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured.

The **no** form of this command resets the default value.

Default 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100

Parameters *threshold-level* — Integer value

Values 0 — 100

te-down-threshold

Syntax **te-down-threshold** *threshold-level* [*threshold-level*...(up to 16 max)]
no te-down-threshold

Context config>router>rsvp
config>router>rsvp>interface

Description This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The **te-threshold-update** command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates is supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported.

Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface must be used to trigger IGP updates if non-default threshold levels are configured.

The **no** form of this command resets the default value.

Default 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0

Parameters *threshold-level* — Integer value

Values 0 — 100

Message Pacing Commands

msg-pacing

Syntax [**no**] **msg-pacing**

Context config>router>rsvp

Description This command enables RSVP message pacing in which the specified number of RSVP messages, specified in the **max-burst** command, are sent in a configured interval, specified in the **period** command. A count is kept of the messages that were dropped because the output queue for the interface used for message pacing was full.

Default **no msg-pacing**

max-burst

Syntax **max-burst** *number*
no max-burst

Context config>router>rsvp>msg-pacing

Description This command specifies the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.

Default 650

Parameters *number —*

Values 100 — 1000 in increments of 10

period

Syntax **period** *milli-seconds*
no period

Context config>router>rsvp>msg-pacing

Description This command specifies the time interval, in milliseconds, when the router can send the specified number of RSVP messages which is specified in the **max-burst** command.

Default 100

Parameters *milli-seconds —*

Values 100 — 1000 milliseconds in increments of 10 milliseconds