

Chapter 6

Differentiated-Services-Aware Traffic Engineering Configuration Guidelines

Differentiated-Services-aware traffic engineering provides a way to guarantee a specified level of service over a Multiprotocol Label Switching (MPLS) network. The routers providing Differentiated-Services-aware traffic engineering are part of a differentiated services network domain. All routers participating in a differentiated services domain must have Differentiated-Services-aware traffic engineering enabled.

To help ensure that the specified service level is provided, it is necessary to ensure that no more than the amount of traffic specified is sent over the differentiated services domain. You can accomplish this goal by configuring a policer to police or rate limit the volume of traffic transiting the differentiated service domain. For more information about how to configure policers for label-switched paths (LSPs), see “Configuring Policers for LSPs” on page 171.

This feature can help to improve the quality of Internet services such as voice over IP (VoIP). It also makes it possible to better emulate an Asynchronous Transfer Mode (ATM) circuit over an MPLS network.

This chapter describes how to configure Differentiated-Services-aware traffic engineering for LSPs and multiclass LSPs:

- Differentiated-Services-Aware Traffic Engineering Standards on page 126
- Differentiated-Services-Aware Traffic Engineering Terminology on page 126
- Differentiated-Services-Aware Traffic Engineering Overview on page 127
- Configuring Differentiated-Services-Aware Traffic Engineering on page 133
- Bandwidth Oversubscription Overview on page 137
- Configuring the Bandwidth Subscription Percentage for LSPs on page 142
- Configuring Differentiated-Services-Aware Traffic Engineering for LSPs on page 144
- Configuring Multiclass LSPs on page 147

Differentiated-Services-Aware Traffic Engineering Standards

The following RFCs and Internet drafts provide information on Differentiated-Services-aware traffic engineering and multiclass LSPs:

- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*

These RFCs and Internet drafts are available on the IETF Web site at <http://www.ietf.org/>.

Differentiated-Services-Aware Traffic Engineering Terminology

The following terminology applies to Differentiated-Services-aware traffic engineering:

- Bandwidth model—The bandwidth model determines the values of the available bandwidth advertised by the interior gateway protocols (IGPs).
- CAC—Call admission control (CAC) checks to ensure there is adequate bandwidth on the path before the LSP is established. If the bandwidth is insufficient, the LSP is not established and an error is reported.
- Class type—A collection of traffic flows that is treated equivalently in a differentiated services domain. A class type maps to a queue and is much like a class-of-service (CoS) forwarding class in concept. It is also known as a traffic class.
- Differentiated Services—Differentiated Services make it possible to give different treatment to traffic based on the experimental (EXP) bits in the MPLS header. Traffic must be marked appropriately and CoS must be configured.
- Differentiated Services domain—The routers in a network that have Differentiated Services enabled.
- Differentiated-Services-aware traffic engineering—A type of constraint-based routing. It can enforce different bandwidth constraints for different classes of traffic. It can also do CAC on each traffic engineering class when an LSP is established.

- Multiclass LSP—A multiclass LSP functions like a standard LSP, but it also allows you to reserve bandwidth from multiple class types. The EXP bits of the MPLS header are used to distinguish between class types.
- MAM—The maximum allocation bandwidth constraint model divides the available bandwidth between the different classes. Sharing of bandwidth between the class types is not allowed.
- RDM—The Russian dolls bandwidth constraint model makes efficient use of bandwidth by allowing the class types to share bandwidth.
- Traffic engineering class—A paired class type and priority.
- Traffic engineering class map—A map between the class types, priorities, and traffic engineering classes. The traffic engineering class mapping must be consistent across the Differentiated Services domain.

Differentiated-Services-Aware Traffic Engineering Overview

Differentiated Services give different treatment to traffic based on the EXP bits in the MPLS label header and allow you to provide multiple classes of service.

The following sections describe Differentiated-Services-aware traffic engineering:

- Differentiated-Services-Aware Traffic Engineering Features on page 127
- Differentiated-Services-Aware Traffic Engineered LSPs on page 128
- Multiclass LSPs on page 131

Differentiated-Services-Aware Traffic Engineering Features

Differentiated-Services-aware traffic engineering provides the following features:

- Traffic engineering at a per-class level rather than at an aggregate level
- Different bandwidth constraints for different class types (traffic classes)
- Different queuing behaviors per class, allowing the router to forward traffic based on the class type

In comparison, standard traffic engineering does not consider CoS, and it completes its work on an aggregate basis across all Differentiated Service classes.

Differentiated-Services-aware traffic engineering provides the following advantages:

- Traffic engineering can be performed on a specific class type instead of at the aggregate level.
- Bandwidth constraints can be enforced on each specific class type.
- It forwards traffic based on the EXP bits.

This makes it possible to guarantee service and bandwidth across an MPLS network. With Differentiated-Services-aware traffic engineering, among other services, you can provide ATM circuit emulation, VoIP, and a guaranteed bandwidth service.

The following describes how the IGP, Constrained Shortest Path First (CSPF), and Resource Reservation Protocol (RSVP) participate in Differentiated-Services-aware traffic engineering:

- The IGP can advertise the unreserved bandwidth for each traffic engineering class to the other members of the differentiated services domain. The traffic engineering database (TED) stores this information.
- A CSPF calculation is performed considering the bandwidth constraints for each class type. If all the constraints are met, the CSPF calculation is considered successful.
- When RSVP signals an LSP, it requests bandwidth for specified class types.

Differentiated-Services-Aware Traffic Engineered LSPs

A Differentiated-Services-aware traffic engineered LSP is an LSP configured to reserve bandwidth for one of the supported class types and to carry traffic for that class type. The following sections discuss this type of LSPs:

- Differentiated-Services-Aware Traffic Engineered LSPs Overview on page 129
- Differentiated-Services-Aware Traffic Engineered LSPs Operation on page 130

Differentiated-Services-Aware Traffic Engineered LSPs Overview

A Differentiated-Services-aware traffic engineered LSP is an LSP configured with a bandwidth reservation for a specific class type. This LSP can carry traffic for a single class type. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

The class type must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network. You can unambiguously map a class type to a queue. On each node router, the class-of-service queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

For more information about topics related to LSPs and Differentiated-Services-aware traffic engineering, see the following:

- For forwarding classes and class of service, see the *JUNOS Class of Service Configuration Guide*.
- For EXP bits, see “Label Allocation” on page 28.
- For differentiated services, see RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*.
- For information about how the IGPs and RSVP have been modified to support Differentiated Services-aware MPLS traffic engineering, see Internet draft draft-ietf-tewg-diff-te-proto-07.txt.

Differentiated-Services-Aware Traffic Engineered LSPs Operation

When configuring a Differentiated-Services-aware traffic engineered LSP, you specify the class type and the bandwidth associated with it. The following occurs when an LSP is established with bandwidth reservation from a specific class type:

1. The IGP advertises how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for an LSP, CSPF is used to ensure that the bandwidth constraints are met for the class type carried by the LSP at the specified priority level.

CSPF also checks to ensure that the bandwidth model is configured consistently on each router participating in the LSP. If the bandwidth model is inconsistent, CSPF does not compute the path (except for LSPs from class type `ct0`).

3. Once a path is found, RSVP signals the LSP using the Classtype object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up.

An LSP that requires bandwidth from a particular class (except class type `ct0`) cannot be established through routers that do not understand the Classtype object. Preventing the use of routers that do not understand the Classtype object helps to ensure consistency throughout the Differentiated Services domain by preventing the LSP from using a router that cannot support Differentiated Services.

By default, LSPs are signaled with setup priority 7 and holding priority 0. An LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both LSPs configured for Differentiated-Services-aware traffic engineering and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings (either by remarking the EXP settings or by assuming that the traffic arrived with the correct EXP settings from the upstream router).

Multiclass LSPs

Multiclass LSPs function like standard LSPs, but they also allow you to configure multiple class types with guaranteed bandwidth. The EXP bits of the MPLS header are used to distinguish between class types. Multiclass LSPs can be configured for a variety of purposes. For example, you can configure a multiclass LSP to emulate the behavior of an ATM circuit. An ATM circuit can provide service-level guarantees to a class type. A multiclass LSP can provide a similar guaranteed level of service.

The following sections discuss multiclass LSPs:

- Multiclass LSP Overview on page 131
- Establishing a Multiclass LSP on the Differentiated Services Domain on page 132

Multiclass LSP Overview

A multiclass LSP is an LSP that can carry several class types. One multiclass LSP can be used to support up to four class types. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

Once a multiclass LSP is configured, traffic from all of the class types can:

- Follow the same path
- Be rerouted along the same path
- Be taken down at the same time

Class types must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network.

You can unambiguously map a class type to a queue. On each node router, the CoS queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

The combination of a class type and a priority level forms a traffic engineering class. The IGP can advertise up to eight traffic engineering classes for each link.

For more information about the EXP bits, see “Label Allocation” on page 28.

For more information about forwarding classes, see the *JUNOS Class of Service Configuration Guide*.

Establishing a Multiclass LSP on the Differentiated Services Domain

The following occurs when a multiclass LSP is established on the differentiated services domain:

1. The IGPs advertise how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for a multiclass LSP, CSPF is used to ensure that the constraints are met for all the class types carried by the multiclass LSP (a set of constraints instead of a single constraint).
3. Once a path is found, RSVP signals the LSP using an RSVP object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up. The RSVP object is a hop-by-hop object. Multiclass LSPs cannot be established through routers that do not understand this object. Preventing routers that do not understand the RSVP object from carrying traffic helps to ensure consistency throughout the differentiated services domain by preventing the multiclass LSP from using a router that is incapable of supporting differentiated services.

By default, multiclass LSPs are signaled with setup priority 7 and holding priority 0. A multiclass LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both multiclass LSPs and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings.

Configuring Differentiated-Services-Aware Traffic Engineering

To configure Differentiated-Services-aware traffic engineering, include the `diffserv-te` statement:

```
diffserv-te {
  bandwidth-model {
    extended-mam;
    mam;
    rdm;
  }
  te-class-matrix {
    traffic-class {
      tnumber {
        priority priority;
        traffic-class ctnumber priority priority;
      }
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-routers *logical-router-name* protocols mpls]

You must configure the `diffserv-te` statement on all routers participating in the Differentiated Services domain. However, you are not required to configure the traffic engineering class matrix (configured with the `te-class-matrix` statement).

To configure Differentiated-Services-aware traffic engineering, complete the procedures in the following sections:

- Configuring the Bandwidth Model on page 134
- Configuring Traffic Engineering Classes on page 135

Configuring the Bandwidth Model

You must configure a bandwidth model on all routers participating in the Differentiated Services domain. The bandwidth models available are MAM, extended MAM, and RDM:

- Maximum allocation bandwidth constraints model (MAM)—Defined in Internet draft draft-ietf-tewg-diff-te-mam-03.txt.
- Extended MAM—A proprietary bandwidth model that behaves much like standard MAM. If you configure multiclass LSPs, you must configure the extended MAM bandwidth model.
- Russian-dolls bandwidth allocation model (RDM)—Makes efficient use of bandwidth by allowing the class types to share bandwidth. RDM is defined in Internet draft draft-ietf-tewg-diff-te-russian-05.txt.

To configure a bandwidth model, include the `bandwidth-model` statement and specify one of the bandwidth model options:

```
bandwidth-model {
  extended-mam;
  mam;
  rdm;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls diffserv-te]
- [edit logical-routers *logical-router-name* protocols mpls diffserv-te]



NOTE: If you change the bandwidth model on an ingress router, all the LSPs enabled on the router are taken down and resignaled.

Configuring Traffic Engineering Classes

You are not required to configure the traffic engineering classes. Table 6 shows the default values for everything in the traffic engineering class matrix. The default mapping is expressed in terms of the default forwarding classes defined in the CoS configuration.

Table 6: Default Values for the Traffic Engineering Class Matrix

Traffic Engineering Class	Class Type	Queue	Priority
te0	ct0	0	7
te1	ct1	1	7
te2	ct2	2	7
te3	ct3	3	7
te4	ct0	0	0
te5	ct1	1	0
te6	ct2	2	0
te7	ct3	3	0

You can override the default mapping by configuring other values for the traffic engineering classes. You can configure traffic engineering classes 0 through 7. For each traffic engineering class, you configure a class type (or queue) from 0 through 3. For each class type, you configure a priority from 0 through 7.

To configure traffic engineering classes explicitly, include the `te-class-matrix` statement:

```
te-class-matrix {
  tnumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls diffserv-te]
- [edit logical-routers *logical-router-name* protocols mpls diffserv-te]

The following example shows how to configure traffic engineering class `te0` with class type `ct1` and a priority of 4:

```
[edit protocols mpls diffserv-te]
te-class-matrix {
  te0 traffic-class ct1 priority 4;
}
```



NOTE: If you explicitly configure a value for one of the traffic engineering classes, all the default values in the traffic engineering class matrix are dropped.

When you explicitly configure traffic engineering classes, you must also configure a bandwidth model; otherwise, the configuration commit fails. See “Configuring the Bandwidth Model” on page 134.

Requirements and Limitations for the Traffic Engineering Class Matrix

When you configure a traffic engineering class matrix, be aware of the following requirements and limitations:

- A mapping configuration is local and affects only the router on which it is configured. It does not affect other systems participating in the differentiated services domain. However, for a Differentiated Services domain to function properly, you need to configure the same traffic engineering class matrix on all the routers participating in the same domain.
- When explicitly configuring traffic engineering classes, you must configure the classes in sequence (`te0`, `te1`, `te2`, `te3`, and so on); otherwise, the configuration commit fails.

The first traffic engineering class you configure must be `te0`; otherwise, the configuration commit fails.

Bandwidth Oversubscription Overview

LSPs are established with bandwidth reservations configured for the maximum amount of traffic you expect to traverse the LSP. Not all LSPs carry the maximum amount of traffic over their links at all times. For example, even if the bandwidth for link A has been completely reserved, actual bandwidth might still be available but not currently in use. This excess bandwidth can be used by allowing other LSPs to also use link A, oversubscribing the link. You can oversubscribe the bandwidth configured for individual class types or specify a single value for all of the class types using an interface.

You can use oversubscription to take advantage of the statistical nature of traffic patterns and to permit higher utilization of links. The following examples describe how you might use bandwidth oversubscription and undersubscription:

- Use oversubscription on class types where peak periods of traffic do not coincide in time.
- Use oversubscription of class types carrying best-effort traffic. You take the risk of temporarily delaying or dropping traffic in exchange for making better utilization of network resources.
- Give different degrees of oversubscription or undersubscription of traffic for the different class types. For instance, you configure the subscription for classes of traffic as follows:
 - Best effort—ct0 1000
 - Voice—ct3 1

When you undersubscribe a class type for a multiclass LSP, the total demand of all RSVP sessions is always less than the actual capacity of the class type. You can use undersubscription to limit the utilization of a class type.

The bandwidth oversubscription calculation occurs on the local router only. Because no signaling or other interaction is required from other routers in the network, the feature can be enabled on individual routers without being enabled or available on other routers which might not support this feature. Neighboring routers do not need to know about the oversubscription calculation, they rely on the IGP.

The following sections describe the types of bandwidth oversubscription available in the JUNOS software:

- LSP Size Oversubscription on page 138
- Link Size Oversubscription on page 138
- Class Type Oversubscription and Local Oversubscription Multipliers on page 138

LSP Size Oversubscription

For LSP size oversubscription, you simply configure less bandwidth than the peak rate expected for the LSP. You also might need to adjust the configuration for automatic policers. Automatic policers manage the traffic assigned to an LSP, ensuring that it does not exceed the configured bandwidth values. LSP size oversubscription requires that the LSP can exceed its configured bandwidth allocation.

Policing is still possible. However, the policer must be manually configured to account for the maximum bandwidth planned for the LSP, rather than for the configured value.

Link Size Oversubscription

You can increase the maximum reservable bandwidth on the link and use the inflated values for bandwidth accounting. Use the **subscription** statement to oversubscribe the link. The configured value is applied to all class type bandwidth allocations on the link. For more information about link size oversubscription, see “Configuring the Bandwidth Subscription Percentage for LSPs” on page 142.

Class Type Oversubscription and Local Oversubscription Multipliers

Local oversubscription multipliers (LOMs) allow different oversubscription values for different class types. LOMs are useful for networks where the oversubscription ratio needs to be configured differently on different links and where oversubscription values are required for different classes. You might use this feature to oversubscribe class types handling best-effort traffic, but use no oversubscription for class types handling voice traffic. An LOM is calculated locally on the router. No information related to an LOM is signaled to other routers in the network.

An LOM is configurable on each link and for each class type. The per-class type LOM allows you to increase or decrease the oversubscription ratio. The per-class-type LOM is factored into all local bandwidth accounting for admission control and IGP advertisement of unreserved bandwidths.

The LOM calculation is tied to the bandwidth model (MAM, extended MAM, and russian dolls) used, because the effect of oversubscription across class types must be accounted for accurately.

The formulas related to the oversubscription of class types are described in the following sections:

- Class Type Bandwidth and the LOM on page 139
- LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 139
- LOM Calculation for the Russian Dolls Bandwidth Model on page 140
- Example: LOM Calculation on page 140



NOTE: All LOM calculations are performed by the JUNOS software and require no user intervention.

Class Type Bandwidth and the LOM

The following formula expresses the relationship between the bandwidth of the class type and the LOM. The normalized bandwidth of the class type (N_B) is equal to the reserved bandwidth of the class type (R_B) divided by the LOM of the class type (L_C):

$$N_B = R_B / L_C$$

When calculating available bandwidth, you need to subtract the normalized bandwidth from the relevant bandwidth constraint.



NOTE: When using an LOM, values advertised for the available bandwidth might be larger than the bandwidth constraint values. However, the values advertised in the maximum link bandwidth advertisement are not affected by local oversubscription.

LOM Calculation for the MAM and Extended MAM Bandwidth Models

The following formulas show how the LOM is calculated for the MAM and extended MAM bandwidth models.

$$\text{Unreserved TE-Class}(i) = LOM_c \times [BC_c - \text{SUM} (\text{Normalized} (CT_c, q))] \text{ for } q \leq p$$

Or

$$\text{Unreserved TE-Class}(i) = (LOM_c \times BC_c) - \text{SUM} (\text{Reserved} (CT_c, q)) \text{ for } q \leq p$$

where:

- LOM_c —LOM for class type c .
- BC_c —Bandwidth constraint for class type c .
- CT_c —Class type c .
- $\text{TE-Class}(i) \leftrightarrow (CT_c, \text{preemption } p)$ in the configured TE-Class mapping.

LOM Calculation for the Russian Dolls Bandwidth Model

The following formulas show how the LOM is calculated for the russian dolls bandwidth model:

$$\begin{aligned} \text{Unreserved TE-Class (i)} &= \text{LOM}_c \times \text{MIN} [\\ &[\text{BC}_c - \text{SUM} (\text{Normalized} (\text{CT}_b, q))] \text{ for } q \leq p \text{ and } c \leq b \leq 7, \\ &\dots \\ &[\text{BC}_0 - \text{SUM} (\text{Normalized} (\text{CT}_b, q))] \text{ for } q \leq p \text{ and } 0 \leq b \leq 7, \\ &] \end{aligned}$$

where:

- LOM_c —LOM for class type c .
- BC_c —Bandwidth constraint for class type c .
- $\text{TE-Class}(i) \leftrightarrow (\text{CT}_c, \text{preemption } p)$ in the configured TE-Class mapping.

Note that the impact of an LSP on the unreserved bandwidth of a class type does not depend only on the LOM for that class type—it also depends on the LOM for the class type of the LSP.

Example: LOM Calculation

The following example illustrates how an LOM calculation is made for four classes of traffic: ct_0 , ct_1 , ct_2 , and ct_3 .

The class types have been assigned the following values:

$$\begin{aligned} \text{ct}_0 &= 40 \\ \text{ct}_1 &= 30 \\ \text{ct}_2 &= 20 \\ \text{ct}_3 &= 10 \end{aligned}$$

These class type values yield the following bandwidth constraints:

$$\begin{aligned} \text{BC}_0 &= (\text{ct}_3 + \text{ct}_2 + \text{ct}_1 + \text{ct}_0) = 100 \\ \text{BC}_1 &= (\text{ct}_3 + \text{ct}_2 + \text{ct}_1) = 60 \\ \text{BC}_2 &= (\text{ct}_3 + \text{ct}_2) = 30 \\ \text{BC}_3 &= (\text{ct}_3) = 10 \end{aligned}$$

LSPs from class type ct_0 can take up to 100 percent of bandwidth on the link. LSPs from class type ct_1 can take up to 60 percent of the bandwidth on the link, and so on.

If you assume for this example that the class types have the following LOM values:

$$\begin{aligned} \text{LOM}(\text{ct}_0) &= 8 \\ \text{LOM}(\text{ct}_1) &= 4 \\ \text{LOM}(\text{ct}_2) &= 2 \\ \text{LOM}(\text{ct}_3) &= 1 \end{aligned}$$

In the absence of any other reservation, LSPs from class type **ct0** can take up to 800 percent of the available bandwidth ($8 \times 100 = 800$). In the absence of any other reservation, LSPs from class type **ct1** can take up to 240 percent of the available bandwidth ($4 \times 60 = 240$). etc.

The maximum amount of bandwidth that can be reserved is:

$$\begin{aligned} \text{ct0} &= \text{LOM}(\text{ct0}) \times \text{BC0} = 800 \\ \text{ct1} &= \text{LOM}(\text{ct1}) \times \text{BC1} = 240 \\ \text{ct2} &= \text{LOM}(\text{ct2}) \times \text{BC2} = 60 \\ \text{ct3} &= \text{LOM}(\text{ct3}) \times \text{BC3} = 10 \end{aligned}$$

For the undersubscribed class type **ct3**, the maximum reservable bandwidth is the same as the bandwidth constraint. For the overbooked class types, these values are not the values of the bandwidth constraint-taking into account the oversubscription for each class type separately. The oversubscription per class type in the sum is not taken into account because ultimately the entire bandwidth constraint can be filled with the bandwidth reservation of just one class type, so you have to account for that class type's bandwidth oversubscription only.

When calculating the available bandwidth for **CTc**, you need to express reservations from other classes as if they were from **CTc**. The reservation from class **ctx** is normalized with the LOM of **ctx**, but it is then multiplied by the LOM of **CTc**.

For the previous example, assume that **LSP1** has class type **ct3** configured with bandwidth of 10 and a priority of 0.

The values for the reservable bandwidth will be:

$$\begin{aligned} \text{ct0} &= 8 \times (100 - 10) = 720 \\ \text{ct1} &= 4 \times \min((100-10), (60-10)) = 200 \\ \text{ct2} &= 2 \times \min((100-10), (60-10), (30-10)) = 40 \\ \text{ct3} &= 1 \times \min((100-10), (60-10), (30-10), (10-10)) = 0 \end{aligned}$$

These numbers can be rationalized as follows: the normalized reservation is 10 percent. If this bandwidth came from class type **ct0**, it would be equivalent to an overbooked reservation of 80 percent. You can see that 720 percent ($800 - 80 = 720$) of the bandwidth remains available for other LSPs.

Configuring the Bandwidth Subscription Percentage for LSPs

By default, RSVP allows all of a class type's bandwidth (100 percent) to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

If you want to oversubscribe or undersubscribe all of the class types on an interface using the same percentage bandwidth, configure the percentage using the `subscription` statement:

```
subscription percentage;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

To undersubscribe or oversubscribe the bandwidth for each class type, configure a percentage for each class type (`ct0`, `ct1`, `ct2`, and `ct3`) option for the `subscription` statement. When you oversubscribe a class type, an LOM is applied to calculate the actual bandwidth reserved. See "Class Type Oversubscription and Local Oversubscription Multipliers" on page 138 for more information.

```
subscription {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

percentage is the percentage of class type bandwidth that RSVP allows to be used for reservations. It can be a value from 0 through 65,000 percent. If you specify a value greater than 100, you are oversubscribing the interface or class type.

The value you configure when you oversubscribe a class type is a percentage of the class type bandwidth that can actually be used. The default subscription value is 100 percent.

You can use the `subscription` statement to disable new RSVP sessions for one or more class types. If you configure a percentage of 0, no new sessions (including those with zero bandwidth requirements) are permitted for the class type.

Existing RSVP sessions are not affected by changing the subscription factor. To clear an existing session, issue the `clear rsvp session` command. For more information on the `clear rsvp session` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Bandwidth Subscription Troubleshooting

Be aware of the following issues when configuring bandwidth subscription:

- If you configure the `subscription` statement at the `[edit protocols rsvp interface all]` hierarchy level, and then configure a different value at the `[edit protocols rsvp interface interface-name]` hierarchy level, the value configured at the `[edit protocols rsvp interface interface-name]` hierarchy level overrides the other value.
- Per-class-type subscription can only be configured if a bandwidth model is also configured. If no bandwidth model is configured, the configuration fails with the following error message:

```
user@host# commit check
[edit protocols rsvp interface all]
  'subscription'
```

```
RSVP: Must have a diffserv-te bandwidth model configured when configuring
subscription per traffic class.
error: configuration check-out failed
```

- If you configure the `subscription` statement for a specific class type, you cannot configure the `subscription` statement for the whole interface. The configuration commits fails with the following error message:

```
user@host# commit check
[edit protocols rsvp interface all]
  'subscription'
```

```
RSVP: Cannot configure both link subscription and per traffic class
subscription.
error: configuration check-out failed
```

Configuring Differentiated-Services-Aware Traffic Engineering for LSPs

You must configure the Differentiated Services domain (see “Configuring Differentiated-Services-Aware Traffic Engineering” on page 133) before you can enable Differentiated-Services-aware traffic engineering for LSPs. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in the LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the LSP to function properly.



NOTE: You must configure either MAM or RDM as the bandwidth model when you configure Differentiated-Services-aware traffic engineering for LSPs. See “Configuring the Bandwidth Model” on page 134.

The actual data transmitted over this Differentiated Services domain is carried by an LSP. Each LSP relies on the EXP bits of the MPLS packets to enable Differentiated-Services-aware traffic engineering. Each LSP can carry traffic for a single class type.

All the routers participating in the LSP must be Juniper Networks routing platforms running JUNOS Release 6.3 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the JUNOS software. However, the Differentiated-Services-aware traffic engineering LSP cannot traverse these routers.



NOTE: You cannot simultaneously configure multiclass LSPs and Differentiated-Services-aware traffic engineering LSPs on the same router.

To enable Differentiated-Services-aware traffic engineering for LSPs, you need to configure the following:

- Configuring Class of Service for the Interfaces on page 145
- Configuring IGP on page 145
- Configuring a Traffic Engineered LSP on page 145
- Configuring Policing for LSPs on page 146
- Configuring Fast Reroute for Traffic Engineered LSPs on page 146

Configuring Class of Service for the Interfaces

The existing class-of-service infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to accomplish this are configured using the existing JUNOS software CoS features.



NOTE: ATM interfaces are not supported.

For information on how to configure CoS, see the *JUNOS Class of Service Configuration Guide*.

Configuring IGP

You can configure either IS-IS or OSPF as the IGP. The IS-IS and OSPF configurations for routers supporting LSPs are standard. For information on how to configure these protocols, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring a Traffic Engineered LSP

You configure an LSP by using the standard LSP configuration statements and procedures. To configure Differentiated-Services-aware traffic engineering for the LSP, specify a class type bandwidth constraint by including the **bandwidth** statement:

```
label-switched-path lsp-name {
  bandwidth {
    ctnumber bps;
  }
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for this statement.

If you do not specify a bandwidth for a class type, **ct0** is automatically specified as the queue for the LSP. You can configure only one class type for each LSP, unlike multiclass LSPs.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1
- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

You can configure setup and holding priorities for an LSP, but the following restrictions apply:

- The combination of class and priority must be one of the configured traffic engineering classes. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit to fail.
- Automatic bandwidth allocation is not supported. If you configure automatic bandwidth allocation, the commit fails.
- LSPs configured with the `bandwidth` statement but without specifying a class type use the default class type `ct0`.
- For migration issues, see Internet draft `draft-ietf-tewg-diff-te-07.txt`.

Configuring Policing for LSPs

Policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each LSP.

For information on how to configure a policer for an LSP, see “Configuring Policers for LSPs” on page 171.

Configuring Fast Reroute for Traffic Engineered LSPs

You can configure fast reroute for traffic engineered LSPs (LSPs carrying a single class of traffic). It is also possible to reserve bandwidth on the detour path for the class of traffic when fast reroute is enabled. The same class type number is used for both the traffic engineered LSP and its detour.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling Differentiated-Services-aware traffic engineering and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

You can configure the amount of bandwidth to reserve for detours using either the `bandwidth` statement or the `bandwidth-percent` statement. You can only configure one these statements at a time. If you do not configure either the `bandwidth` statement or the `bandwidth-percent` statement, the default setting is to not reserve bandwidth for the detour path (the bandwidth guarantee will be lost if traffic is switched to the detour).

When you configure the `bandwidth` statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. For information, see “Configuring Fast Reroute” on page 77.

When you configure the `bandwidth-percent` statement, the detour path bandwidth is computed by multiplying it to the bandwidth configured for the main traffic engineered LSP. For information on how to configure the bandwidth for a traffic engineered LSP, see “Configuring a Traffic Engineered LSP” on page 145.

To configure the percent of bandwidth used by the detour path based on the bandwidth of the protected path, include the `bandwidth-percent` statement:

```
bandwidth-percent percent;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-path-name* fast-reroute]
- [edit logical-routers *logical-router-name* protocols mpls label-switched-path *lsp-path-name* fast-reroute]

Configuring Multiclass LSPs

A multiclass LSP is an LSP configured to reserve bandwidth for multiple class types and also carries the traffic for these class types. The differentiated service behavior is determined by the EXP bits.

You must configure the Differentiated Services domain (see “Configuring Differentiated-Services-Aware Traffic Engineering” on page 133) before you can enable a multiclass LSP. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in a multiclass LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the multiclass LSP to function properly.



NOTE: You must configure extended MAM as the bandwidth model when you configure multiclass LSPs. See “Configuring the Bandwidth Model” on page 134.

All the routers participating in a multiclass LSP must be Juniper Networks routing platforms running JUNOS Release 6.2 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the JUNOS software. However, the multiclass LSP cannot traverse these routers.

To enable multiclass LSPs, you need to configure the following:

- Configuring Class of Service for the Interfaces on page 148
- Configuring the IGP on page 148
- Configuring a Multiclass LSP on page 148
- Configuring Policing for Multiclass LSPs on page 149
- Configuring Fast Reroute for Multiclass LSPs on page 150

Configuring Class of Service for the Interfaces

The existing class-of-service infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to consistently mark traffic are configured with the existing JUNOS software CoS features.



NOTE: ATM interfaces are not supported.

For information on how to configure CoS, see the *JUNOS Class of Service Configuration Guide*.

Configuring the IGP

You can configure either IS-IS or OSPF. The IS-IS and OSPF configurations for routers supporting multiclass LSPs are standard. For information about how to configure these protocols, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring a Multiclass LSP

You configure a multiclass LSP by using the standard LSP configuration statements and procedures. To configure an LSP as a multiclass LSP, specify the class type bandwidth constraints by including the **bandwidth** statement:

```
bandwidth {
  ct0 bps;
  ct1 bps;
  ct2 bps;
  ct3 bps;
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for these statements.

The class type statements specify bandwidth (in bits per second) for the following classes:

- ct0—Bandwidth reserved for class 0
- ct1—Bandwidth reserved for class 1
- ct2—Bandwidth reserved for class 2
- ct3—Bandwidth reserved for class 3

For example, to configure 50 megabytes of bandwidth for class type 1 and 30 megabytes of bandwidth for class type 2, include the `bandwidth` statement as follows:

```
[edit protocols mpls]
label-switched-path traffic-class {
  bandwidth {
    ct1 50M;
    ct2 30M;
  }
}
```

You cannot configure a bandwidth for a class type and also configure a bandwidth at the `[label-switched-path lsp-path-name bandwidth]` hierarchy level. For example, the following configuration cannot be committed:

```
[edit protocols mpls]
label-switched-path traffic-class {
  bandwidth {
    20M;
    ct1 10M;
  }
}
```

You can configure setup and holding priorities for a multiclass LSP, but the following restrictions apply:

- The setup and holding priorities apply to all classes for which bandwidth is requested.
- The combination of class and priority must be one of the configured traffic engineering classes. The default traffic engineering class configuration results in multiclass LSPs that cannot preempt and cannot be preempted. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit to fail.
- Automatic bandwidth allocation is not supported for multiclass LSPs. If you configure automatic bandwidth allocation, the commit fails.
- LSPs configured with the `bandwidth` statement but without specifying a class type use the default class type `ct0`.

Configuring Policing for Multiclass LSPs

Policing allows you to control the amount of traffic forwarded through a particular multiclass LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each multiclass LSP. You can also enable automatic policing for multiclass LSPs.

For information on how to configure a policer for a multiclass LSP, see “Configuring Policers for LSPs” on page 171 and “Configuring Automatic Policers” on page 174.

Configuring Fast Reroute for Multiclass LSPs

You can enable fast reroute for multiclass LSPs. The bandwidth guarantees for the class types can be carried over to the detour path in case the primary path of the multiclass LSP fails. The same traffic class types configured for the primary multiclass LSP are also signaled for the detour LSP.

The bandwidth guarantee for the detour path is a percentage of the bandwidth configured for the class types of the primary path. For example, you configure a value of 50 percent for the detour path and the protected LSP carries traffic for class types CT0 through CT3. The detour path is signaled with the same class types (CT0 through CT3) but with 50 percent of the bandwidth configured for the protected LSP.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling Differentiated-Services-aware traffic engineering, that all of the traffic class types needed are available, and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

The bandwidth percentage for fast reroute is signaled from the ingress router to the egress router. All of the intermediate devices must complete their own CSPF computations and signaling.

When you configure the `bandwidth-percent` statement, the detour path bandwidth is computed by multiplying it to the bandwidth configured for the primary multiclass LSP. For information on how to configure the bandwidth for the multiclass LSP, see “Configuring a Traffic Engineered LSP” on page 145.

To configure the percent of bandwidth used by the detour path based on the bandwidth of the protected path, include the `bandwidth-percent` statement:

```
bandwidth-percent percent;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-path-name* fast-reroute]
- [edit logical-routers *logical-router-name* protocols mpls label-switched-path *lsp-path-name* fast-reroute]