Chapter 6

Configure Differentiated-Services-Aware Traffic Engineering

Differentiated-services-aware traffic engineering provides a way to guarantee a specified level of service over an 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 “Configure Policers for LSPs” on page 160.

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 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

The following IETF 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*

Internet draft draft-ietf-tewg-diff-te-proto-*version*.txt, *Protocol Extensions for Support of Differentiated Services-aware MPLS Traffic Engineering*

Internet draft draft-ietf-tewg-diff-te-mam-*version*.txt, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-Aware MPLS Traffic Engineering*

Internet draft draft-ietf-tewg-diff-te-russian-*version*.txt, *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 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.

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

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 experimental (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 Network Interfaces and Class of Service Configuration Guide*.
- For EXP bits, see "Label Allocation" on page 26.
- For differentiated services, see Internet RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*.
- For how the IGPs and RSVP have been modified to support differentiated services-aware MPLS traffic engineering, see the Internet draft draft-ietf-tewg-diff-te-proto-05.txt, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*.

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 IGPs advertise 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 is incapable of supporting 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).

Differentiated-Services-Aware Traffic Engineered LSPs Limitations

The following limitations apply to differentiated-services-aware traffic engineered LSPs:

- ATM interfaces are not supported.

- You can configure fast reroute on differentiated-services-aware traffic engineering LSPs. However, if the primary path fails and the fast reroute detour path is used, the service guarantees are lost.

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

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 IGPs can advertise up to eight traffic engineering classes for each link.

For more information about the EXP bits, see “Label Allocation” on page 26.

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

Establish 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.

Multiclass LSP Limitations

The following limitations apply to multiclass LSPs:

- ATM interfaces are not supported.

- You can configure fast reroute on a multiclass LSP. However, if the primary path fails and the fast reroute detour path is used, the service guarantees are lost.

Configure 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 {
            tenumber {
                priority priority;
                traffic-class ctnumber priority priority;
            }
        }
    }
}
```

You can include the diffserv-te statement at the following hierarchy levels:

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

[edit 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:

Configure the Bandwidth Model on page 130

Configure Traffic Engineering Classes on page 131

Configure 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 the Russian dolls bandwidth model (RDM):

MAM—Defined in the IETF Internet draft draft-ietf-tewg-diff-te-mam-version.txt, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-Aware MPLS Traffic Engineering*.

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.

RDM—Makes efficient use of bandwidth by allowing the class types to share bandwidth. RDM is defined in the Internet draft draft-ietf-tewg-diff-te-russian-03.txt, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*.

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 the bandwidth-model statement at the following hierarchy levels:

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

[edit 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.

Configure Traffic Engineering Classes

You are not required to configure the traffic engineering classes. Table 3 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 3: 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 {
    tenumber traffic-class ctnumber priority priority;
}
```

You can configure this statement at the following hierarchy levels:

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

[edit protocols mpls diffserv-te]

The following example shows how you would 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 “Configure the Bandwidth Model” on page 130.

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.

Configure Differentiated-Services-Aware Traffic Engineering for LSPs

You must configure the differentiated services domain (see “Configure Differentiated-Services-Aware Traffic Engineering” on page 129) 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 “Configure the Bandwidth Model” on page 130.

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 software 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:

Configure Class of Service for the Interfaces on page 133

Configure IGP on page 133

Configure a Traffic Engineered LSP on page 133

Configure Policing for LSPs on page 134

Configure 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 CoS features.

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

Configure 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*.

Configure 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 bandwidth;
    }
}
```

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 the Internet draft draft-ietf-tewg-diff-te-proto-*version*.txt, *Protocol Extensions f or Support of Differentiated Services-a ware MPLS Traffic Engineering* .

Configure 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 “Configure Policers for LSPs” on page 160.

Configure 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 “Configure Differentiated-Services-Aware Traffic Engineering” on page 129) 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 “Configure the Bandwidth Model” on page 130.

All the routers participating in a multiclass LSP must be Juniper Networks routing platforms running JUNOS software 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:

Configure 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 CoS features.

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

Configure 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*.

Configure 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:

```
label-switched-path lsp-name {
    bandwidth {
        ct0 bandwidth;
        ct1 bandwidth;
        ct2 bandwidth;
        ct3 bandwidth;
    }
}
```

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.

Configure 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 “Configure Policers for LSPs” on page 160 and “Configure Automatic Policers” on page 163.