

Francois Le Faucheur
Thomas D. Nadeau
Cisco Systems, Inc.

Angela Chiu
Celion Networks

William Townsend
Tenor Networks

Darek Skalecki
Nortel Networks

IETF Internet Draft

Expires: August, 2001

Document: [draft-ietf-mpls-diff-te-ext-01.txt](#)

February 2001

**Extensions to RSVP-TE and CR-LDP
for support of Diff-Serv-aware MPLS Traffic Engineering**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are Working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at

<http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#), reference [[BCP14](#)].

Abstract

A companion document [[DIFF-TE-REQTS](#)] defines the requirements for support of Diff-Serv-aware MPLS Traffic Engineering on a per-Class-

Extensions for Diff-Serv Traffic Engineering February 2001

Type basis, as discussed in the Traffic Engineering Working Group Framework document [[TEWG-FW](#)].

This document proposes corresponding extensions to RSVP-TE and CR-LDP for support of Traffic Engineering on a per-Class-Type basis.

Two companion documents [[DIFF-TE-OSPF](#)] [[DIFF-TE-ISIS](#)] propose corresponding extensions to OSPF and ISIS for support of Traffic Engineering on a per-Class-Type basis.

1. Introduction

As Diff-Serv becomes prominent in providing scalable multi-class of services in IP networks, performing traffic engineering at a per-class level instead of an aggregated level is needed in networks where fine optimisation of resources is sought in order to further enhance performance and efficiency. By mapping a traffic trunk in a given class on a separate LSP, it allows the traffic trunk to utilize resources available on both shortest path(s) and non-shortest path(s) and follow paths that meet constraints which are specific to the given class. It also allows each class to select the proper protection/restoration mechanism(s) that satisfy its survivability requirements in a cost-effective manner.

Besides the set of parameters defined for the general aggregate TE [[TE-REQ](#)], a new set of per-class parameters needs to be provided at each LSR interface and propagated via extensions to the IGP (ISIS/OSPF) [[TEWG-FW](#)]. Furthermore, the per-class parameters can be aggregated into per-Class-Type parameters. The main motivation for grouping a set of classes into a Class-Type is to improve the scalability of the IGP link state advertisements by propagating information on a per-Class-Type basis instead of on a per-class basis. This approach also has the benefit of allowing better bandwidth sharing between classes in the same Class-Type.

A Class-Type [[TEWG-FW](#)] is defined as a set of classes that satisfy the following two conditions:

- 1) Classes in the same Class-Type possess common aggregate maximum and minimum bandwidth requirements to guarantee the required performance level.
- 2) There is no maximum or minimum bandwidth requirement to be enforced at the level of an individual class within the Class-Type. One can still implement some "priority" policies for classes within the same Class-Type in terms of accessing the Class-Type bandwidth (e.g. via the use of preemption

priorities).

An example of Class-Type comprising multiple Diff-Serv classes is a low-loss Class-Type that includes both AF1-based and AF2-based Ordering Aggregates.

Note that with per Class-Type TE, Constraint-Based Routing is performed with bandwidth constraints on a per Class-Type basis but LSPs may carry a single Diff-Serv class (Ordered Aggregate) with Diff-Serv scheduling (i.e. PHB) performed separately for each class.

In this document, we will only discuss "per Class-Type TE" because "per Class TE" can be viewed as a special case of per Class-Type TE (where each Class-Type is degenerated into a single Diff-Serv class).

This document focuses on intra-domain operations. Inter-domain operations is for further study.

A companion document [[DIFF-TE-REQTS](#)] defines the requirements for support of MPLS Traffic Engineering on a per-Class-Type basis. The following sections propose detailed extensions to RSVP and CR-LDP that meet those requirements. Two companion documents [[DIFF-TE-OSPF](#)] and [[DIFF-TE-ISIS](#)] propose extensions to OSPF and ISIS that also meet those requirements.

2. RSVP Extensions

In this section we describe extensions to RSVP for support of Diff-Serv Traffic Engineering on a per-Class-Type basis which meet the requirements defined in [[DIFF-TE-REQTS](#)]. These extensions are in addition to the extensions to RSVP defined in [[RSVP-TE](#)] for support of (aggregate) MPLS Traffic Engineering and to the extensions to RSVP defined in [[DIFF-MPLS](#)] for support of Diff-Serv over MPLS.

2.1. Diff-Serv related RSVP Messages Format

One new RSVP Object is defined in this document: the CLASSTYPE Object. Detailed description of this Object is provided below. This new Object is applicable to Path messages. This specification only defines the use of the CLASSTYPE Object in Path messages used to establish LSP Tunnels in accordance with [[RSVP-TE](#)] and thus containing a Session Object with a C-Type equal to LSP_TUNNEL_IPv4 and containing a LABEL_REQUEST object.

Restrictions defined in [[RSVP-TE](#)] for support of establishment of LSP Tunnels via RSVP are also applicable to the establishment of LSP Tunnels supporting Diff-Serv Traffic Engineering. For instance, only

CT : 2 bits

Indicates the Class-Type. Values currently allowed are 1, 2 and 3.

2.3. Handling CLASSTYPE Object

Le Faucheur et. al

4

Extensions for Diff-Serv Traffic Engineering February 2001

To establish an LSP tunnel with RSVP, the sender LSR creates a Path message with a session type of LSP_Tunnel_IPv4 and with a LABEL_REQUEST object as per [[RSVP-TE](#)]. The sender LSR may also include the DIFFSERV object as per [[DIFF-MPLS](#)].

If the LSP is associated with Class-Type 0, the sender LSR must not include the CLASSTYPE object in the Path message.

If the LSP is associated with Class-Type N (N=1,2,3), the sender LSR must include the CLASSTYPE object in the Path message with the Class-Type (CT) field set to N.

If a path message contains multiple CLASSTYPE objects, only the first one is meaningful; subsequent CLASSTYPE object(s) must be ignored and must not be forwarded.

Each LSR along the path records the CLASSTYPE object, when present, in its path state block.

If the CLASSTYPE object is not present in the Path message, the LSR must associate the Class-Type 0 to the LSP.

The destination LSR responds to the Path message by sending a Resv message without a CLASSTYPE object (whether the Path message contained a CLASSTYPE object or not).

During establishment of an LSP corresponding to the Class-Type N, the LSR performs admission control over the bandwidth available for that particular Class-Type, which is computed using the smallest of:

- the Class-Type N bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Bandwidth for Class-Type N and the bandwidth reserved by existing Class-Type N LSPs).
- the aggregate bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Aggregate Bandwidth and the bandwidth reserved by existing LSPs of all Class-Types).

[Editor's Note: the admission control algorithm described in the previous paragraph depends on the Bandwidth Reservation Scheme discussed in section 2.1 of [[DIFF-TE-REQTS](#)]]

In order to accurately apportion the resources associated with a Class-Type among the classes comprised in this Class-Type, the LSR may automatically adjust Diff-Serv scheduling parameters associated with a class within a Class-Type based on the bandwidth currently reserved by LSPs currently established in that class.

An LSR that recognizes the CLASSTYPE object and that receives a path message which contains the CLASSTYPE object but which does not contain a LABEL_REQUEST object or which does not have a session type of LSP_Tunnel_IPv4, must send a PathErr towards the sender with the

error code 'Diff-Serv TE Error' and an error value of 'Unexpected CLASSTYPE object'. Those are defined below in [section 4.5](#).

An LSR receiving a Path message with the CLASSTYPE object, which recognizes the CLASSTYPE object but does not support the particular Class-Type, must send a PathErr towards the sender with the error code 'Diff-Serv TE Error' and an error value of 'Unsupported Class-Type'. Those are defined below in [section 4.5](#).

An LSR receiving a Path message with the CLASSTYPE object, which recognizes the CLASSTYPE object but determines that the Class-Type value is not valid (i.e. Class-Type value 0), must send a PathErr towards the sender with the error code 'Diff-Serv TE Error' and an error value of 'Invalid Class-Type value'. Those are defined below in [section 4.5](#).

An LSR MUST handle the situations where the LSP can not be accepted for other reasons than those already discussed in this section, in accordance with [[RSVP-TE](#)] and [[DIFF-MPLS](#)] (e.g. a reservation is rejected by admission control, a label can not be associated).

2.4. Non-support of the CLASSTYPE Object

An LSR that does not recognize the CLASSTYPE object Class-Num must behave in accordance with the procedures specified in [RSVP] for an unknown Class-Num whose format is 0bbbbbbb (i.e. it must send a PathErr with the error code 'Unknown object class' toward the sender).

An LSR that recognizes the CLASSTYPE object Class-Num but does not recognize the CLASSTYPE object C-Type, must behave in accordance with the procedures specified in [RSVP] for an unknown C-type (i.e. it must send a PathErr with the error code 'Unknown object C-Type' toward the sender).

In both situations, this causes the path set-up to fail. The sender should notify management that a LSP cannot be established and possibly might take action to retry reservation establishment

without the CLASSTYPE object.

2.5. Error Codes For Diff-Serv TE

In the procedures described above, certain errors must be reported as a 'Diff-Serv TE Error'. The value of the 'Diff-Serv TE Error' error code is (TBD).

The following defines error values for the Diff-Serv TE Error:

Value	Error
1	Unexpected CLASSTYPE object
2	Unsupported Class-Type

Le Faucheur et. al

6

Extensions for Diff-Serv Traffic EngineeringFebruary 2001

3	Invalid Class-Type value
---	--------------------------

3. CR-LDP Extensions

CR-LDP, defined in [CR-LDP], is an extension to LDP, defined in [LDP], for support of (aggregate) MPLS Traffic Engineering. In this section we describe extensions to CR-LDP for support of Diff-Serv Traffic Engineering on a per-Class-Type basis which meet the requirements defined in [DIFF-TE-REQTS]. These extensions are in addition to the extensions to LDP defined in [DIFF-MPLS] for support of Diff-Serv over MPLS. They closely resemble the extensions to RSVP defined in the previous section.

Note that extensions of this section for support of Diff-Serv Traffic Engineering are not applicable to LDP due to the fact that LDP does not support MPLS Traffic Engineering and bandwidth reservation in particular.

3.1. Diff-Serv related CR-LDP Messages Encoding

One new CR-LDP TLV is defined in this document: the Class Type TLV. Detailed description of this TLV is provided below. This new TLV is applicable to Label Request messages.

Restrictions defined in [CR-LDP] for support of establishment of LSPs via CR-LDP are also applicable to the establishment of LSPs supporting Diff-Serv Traffic Engineering: for instance, only unicast LSPs are supported and multicast LSPs are for further study.

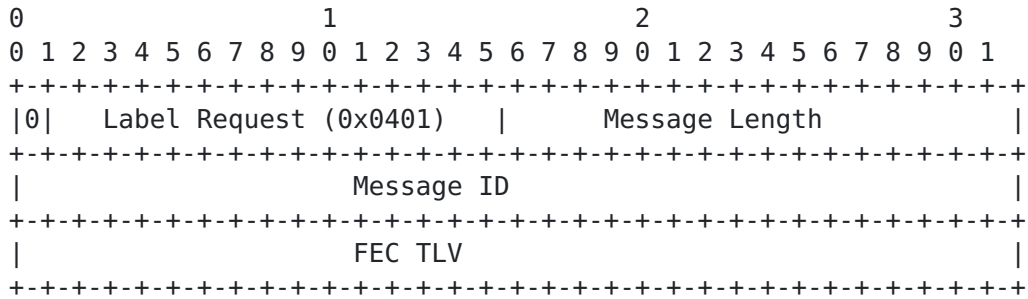
This new Class Type TLV is optional with respect to CR-LDP so that general CR-LDP implementations not concerned with per-Class-Type Diff-Serv Traffic Engineering are not required to support this TLV.

An LSR supporting Diff-Serv Traffic Engineering on a per-Class-Type

basis in compliance with this specification MUST support the Class Type TLV. It MUST support Class-Type value 1, and MAY support other Class-Type values.

3.1.1. Label Request Message Encoding


The encoding for the CR-LDP Label Request message is extended as follows, to optionally include the Class Type TLV:



Le Faucheur et. al

7


Extensions for Diff-Serv Traffic Engineering February 2001



The extension is based on a related LDP extension, defined in [DIFF-MPLS], for support of Diff-Serv TLV but further extended for CR-LDP with CR-LDP TLVs.

3.2. Class Type TLV

The Class Type TLV has the following form:



Reserved : 30 bits

This field is reserved. It must be set to zero on transmission and must be ignored on receipt.

CT : 2 bits

Indicates the Class-Type. Values currently allowed are 1, 2 and 3.

3.3. Handling Class Type TLV

To establish an LSP using CR-LDP, an ingress LSR generates a Label Request message as per [CR-LDP]. This Label Request may optionally include the Diff-Serv TLV as defined in [DIFF-MPLS] for LDP but extended to CR-LDP.

If the LSP is associated with Class-Type 0, the ingress LSR must not include the Class Type TLV in the Label Request message.

If the LSP is associated with Class-Type N (N=1,2,3), the ingress LSR must include the Class Type TLV in the Label Request message with the Class-Type (CT) field set to N.

If a Label Request message contains multiple Class Type TLVs, only the first one is meaningful; subsequent Class Type TLV(s) must be ignored and not forwarded.

If the Class Type TLV is not present in the Label Request message, an LSR must associate the Class-Type 0 to the LSP.

Le Faucheur et. al

8

Extensions for Diff-Serv Traffic Engineering February 2001

A downstream LSR sending a Label Mapping message in response to a Label Request message must not include the Class-Type TLV (whether the Class-Type TLV was included in the Label Request message or not).

During establishment of an LSP corresponding to the Class-Type N, an LSR performs admission control over the bandwidth available for that particular Class-Type, which is computed using the smallest of:

- the Class-Type N bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Bandwidth for Class-Type N and the bandwidth reserved by existing Class-Type N LSPs).
- the aggregate bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Aggregate Bandwidth and the bandwidth reserved by existing LSPs of all Class-Types).

[Editor's Note: the admission control algorithm described in the previous paragraph depends on the Bandwidth Reservation Scheme discussed in section 2.1 of [DIFF-TE-REQTS]]

In order to accurately apportion the resources associated with a Class-Type among the classes comprised in this Class-Type, an LSR may automatically adjust Diff-Serv scheduling parameters associated with a class within a Class-Type based on the bandwidth currently reserved by LSPs currently established in that class.

An LSR that recognizes the Class Type TLV and receives a Label Request message which contains the Class Type TLV but which does not

contain any of the CR-LDP TLVs, must reject the label request by sending upstream a Notification message which includes the Status TLV with a Status Code of 'Unexpected Class-Type TLV'. This is defined below in [section 5.4](#). This error can only occur when an LDP LSP as opposed to CR-LDP LSP is being established. As was already mentioned, Class Type TLV extension for Diff-Serv Traffic Engineering is not applicable to LDP.

An LSR receiving a Label Request message with the Class Type TLV, which recognizes the Class Type TLV but does not support the particular Class-Type, must reject the label request by sending upstream a Notification message which includes the Status TLV with a Status Code of 'Unsupported Class-Type'. This is defined below in [section 5.4](#).

An LSR receiving a Label Request message with the Class Type TLV, which recognizes the Class Type TLV but determines that the Class-Type value is not valid (i.e. Class-Type value 0), must reject the label request by sending upstream a Notification message which includes the Status TLV with a Status Code of 'Invalid Class-Type value'. This is defined below in [section 5.4](#).

An LSR MUST handle the situations where the LSP can not be accepted for other reasons than those already discussed in this section, in accordance with [[CR-LDP](#)], [[LDP](#)] and [[DIFF-MPLS](#)] (e.g. reservation rejected by admission control, a label can not be associated).

[3.4.](#) Status Code Values for Diff-Serv TE

In the procedures described above, certain errors must be reported. The following values are defined for the Status Code field of the Status TLV:

Status Code	E	Status Data
Unexpected Class Type TLV	0	TBD
Unsupported Class-Type	0	TBD
Invalid Class-Type value	0	TBD

[4.](#) Security Considerations

This document raises no new security issues for IS-IS, OSPF, RSVP or CR-LDP. The security mechanisms already proposed for these technologies may be used.

[5.](#) Acknowledgments

This document has benefited from discussions with Carol Iturralde and Rob Goguen.

References

[TE-REQ] Awduche et al, Requirements for Traffic Engineering over MPLS, [RFC2702](#), September 1999.

[TEWG-FW] Awduche et al, A Framework for Internet Traffic Engineering, [draft-ietf-tewg-framework-02.txt](#), July 2000.

[DIFF-TE-REQTS] Le Faucheur et al, Requirements for support of Diff-Serv-aware MPLS Traffic Engineering, [draft-ietf-tewg-diff-te-reqts-00.txt](#), February 2001.

[DIFF-TE-OSPF] Le Faucheur et al, Extension to OSPF for support of Diff-Serv-aware MPLS Traffic Engineering, [draft-ietf-ospf-diff-te-00.txt](#), February 2001.

[DIFF-TE-ISIS] Le Faucheur et al, Extension to ISIS for support of Diff-Serv-aware MPLS Traffic Engineering, [draft-ietf-isis-diff-te-00.txt](#), February 2001.

Le Faucheur et. al

10

Extensions for Diff-Serv Traffic Engineering February 2001

[OSPF-TE] Katz, Yeung, Traffic Engineering Extensions to OSPF, [draft-katz-yeung-ospf-traffic-03.txt](#), September 2000.

[ISIS-TE] Smit, Li, IS-IS extensions for Traffic Engineering, [draft-ietf-isis-traffic-02.txt](#), September 2000.

[RSVP-TE] Awduche et al, "Extensions to RSVP for LSP Tunnels", [draft-ietf-mpls-rsvp-lsp-tunnel-07.txt](#), August 2000.

[DIFF-MPLS] Le Faucheur et al, "MPLS Support of Diff-Serv", [draft-ietf-mpls-diff-ext-08.txt](#), February 2001

[LDP] Andersson et al., "LDP Specification", [draft-ietf-mpls-ldp-011.txt](#), August 2000

[CR-LDP] Jamoussi et al., "Constraint-Based LSP Setup using LDP", [draft-ietf-mpls-cr-ldp-04.txt](#), July 2000

[BCP14] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

Authors' Address:

Francois Le Faucheur
Cisco Systems, Inc.
Petra B - Les Lucioles - 291, rue Albert Caquot - 06560 Valbonne -
France
Phone: +33 4 92 96 75 64
Email: flefauch@cisco.com

Angela Chiu
Celion Networks
1 Sheila Drive, Suite 2
Tinton Falls, NJ 07724
Phone: +1-732 747 9987
Email: angela.chiu@celion.com

William Townsend
Tenor Networks
100 Nagog Park
Acton, MA 01720
Phone: +1-978-264-4900
Email: btownsend@tenornetworks.com

Thomas D. Nadeau
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
Phone: +1-978-244-3051
Email: tnadeau@cisco.com

Le Faucheur et. al

11

Extensions for Diff-Serv Traffic Engineering February 2001

Darek Skalecki
Nortel Networks
3500 Carling Ave,
Nepean K2H 8E9
Phone: +1-613-765-2252
Email: dareks@nortelnetworks.com

