

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-ospf-diff-te-00.txt](#)

February, 2001

**Extensions to OSPF
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/1id-abstracts.txt>.

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

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

Abstract

A companion document [[DIFF-TE-REQTS](#)] defines the requirements for support of Diff-Serv-aware MPLS Traffic Engineering on a per-Class-Type basis, as discussed in the Traffic Engineering Working Group Framework document [[TEWG-FW](#)].

This document proposes corresponding extensions to OSPF for support of Traffic Engineering on a per-Class-Type basis.

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

1. Introduction

As Diffserv 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 paths 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.

Le Faucheur et. al

2

OSPF Extensions for Diff-Serv Traffic Engineering Feb 2001

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 OSPF that meet those requirements.

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

2. OSPF Extensions

In this section we propose extensions to OSPF 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 already defined for support of (aggregate) MPLS Traffic Engineering in [[OSPF-TE](#)].

2.1. Existing TE Sub-TLVs

[OSPF-TE] defines a new LSA for support of (aggregate) Traffic Engineering, which is referred to as the Traffic Engineering LSA. This LSA contains a Link TLV (Type 2) comprising a number of sub-TLVs.

In this document we refer to the sub-TLV 7 (maximum reservable bandwidth) of the Link TLV (as defined in [[OSPF-TE](#)]) as the "Maximum Reservable Aggregate Bandwidth".

We also refer to the sub-TLV 8 (unreserved bandwidth) of the Link TLV (as defined in [[OSPF-TE](#)]) as the "Unreserved Bandwidth for Class-Type 0".

2.2. New Sub-TLVs

The following additional sub-TLVs are defined for the Link TLV of the Traffic Engineering LSA (sub-TLV numbers to be allocated)

TBD1 - Unreserved Bandwidth for Class-Type 1 (32 octets)

TBD2 - Unreserved Bandwidth for Class-Type 2 (32 octets)

TBD3 - Unreserved Bandwidth for Class-Type 3 (32 octets)

Each sub-TLV may occur only once. Unrecognized types are ignored.

Le Faucheur et. al

3

OSPF Extensions for Diff-Serv Traffic Engineering Feb 2001

Unlike the sub-TLVs defined for the Link TLV in [[OSPF-TE](#)], the additional sub-TLVs defined above are optional.

The Link TLV may include the sub-TLVs for any subset of the three additional Class-Types. In other words, the Link TLV may contain none of the three sub-TLVs defined above, any one of those, any two of those, or the three sub-TLVs.

As discussed in [[DIFF-TE-REQTS](#)], where a Class-Type is not effectively used in a network, it is recommended that the corresponding sub-TLV is not included in the Link TLV. Therefore, the Class-Types to be advertised in OSPF should be configurable. For instance, a Network Administrator may elect to use Diff-Serv Traffic Engineering in order to compute separate routes for data traffic and voice traffic (and apply different bandwidth constraints to the route computation for those). In that case, the IGP would only advertise the sub-TLV for one additional Class-Type (i.e. the Link TLV would contain sub-TLV 7 for the Maximum Reservable Aggregate Bandwidth, sub-TLV 8 for the Unreserved Bandwidth for Class-Type 0 and sub-TLV TBD1 for Unreserved Bandwidth for Class-Type 1).

An LSR which supports Class-Type N and which receives a Link TLV without the sub-TLV corresponding to Class-Type N, interprets this as meaning that the corresponding link does not support Class-Type N. For Constraint Based Routing purposes, the LSR may consider this equivalent to the case where the Link TLV contains an Unreserved Bandwidth for Class-Type N sub-TLV set to zero.

An LSR which does not support Class-Type N and which receives a Link TLV containing the sub-TLV corresponding to Class-Type N, must ignore this sub-TLV. However, the Link TLV must be flooded transparently, so that the sub-TLV for Class-Type N is kept in the Link TLV when reflooded by this LSR.

2.3. Sub-TLV Details

The Unreserved Bandwidth for Class-Type N ($N = 1, 2, 3$) sub-TLV specifies the amount of bandwidth not yet reserved at each of the eight preemption priority levels for Class-Type N. Each value will be less than or equal to the Maximum Reservable Bandwidth for Class-Type N.

When the bandwidth value for preemption Z ($Z > 0$) is identical to the bandwidth value for preemption Z-1, the bandwidth value for preemption Z is not explicitly repeated in the sub-TLV. Rather, the fact that it is identical to the value of preemption Z-1, is encoded in a "repetition octet".

Thus, the sub-TLV comprises:

- P ($1 \leq P \leq 8$) bandwidth values. These values correspond to the bandwidth that can be reserved with a holding priority of 0 through 7, arranged in increasing order with priority 0 occurring at the start of the sub-TLV, and priority 7 towards the end of the sub-TLV, but omitting all repeated values. The units are bytes per second and the values are encoded in IEEE floating point format.

- a "repetition octet" where each bit is referred to as bitZ , $0 \leq Z < 8$, and is defined to have the following meaning:

- * if bitZ = 0 then "Unreserved Bandwidth" for preemption level Z is explicitly included in the sub-TLV,
- * if bitZ = 1 then "Unreserved Bandwidth" for preemption level Z is not explicitly included in the sub-TLV but is defined to be equal to "Unreserved Bandwidth" for preemption level Z-1.

Note that the highest preemption level (level 0) is always advertised and the first bit (Bit0) in the "repetition octet" is always set to 0.

[Editor's note: should the "repetition octet" be moved before the bandwidth values?]

The Unreserved Bandwidth for Class-Type N sub-TLV is TLV type (TBDN). Its length is $(P * 4 + 1)$, where $1 \leq P \leq 8$ and where P is the number of non-equal bandwidth values across all preemption levels for that Class-Type.

For example, when a link supports LSPs of preemption levels 2 and 4 only (for a particular Class-Type) with "Unreserved Bandwidth" (for the particular Class-Type) on that link for preemption levels 0, 2,

and 4 currently of 10Mb/s, 5Mb/s and 3Mb/s, respectively, then "Unreserved Bandwidth" (for the particular Class-Type) for preemption levels 0, 2, and 4 of 10Mb/s, 5Mb/s and 3Mb/s, respectively, are explicitly advertised for that link as well as "repetition octet" of 01010111 in binary form. The sub-TLV length is 13.

3. Security Considerations

This document raises no new security issues for OSPF. The security mechanisms already proposed for OSPF may be used.

4. Acknowledgments

This document has benefited from discussions with Carol Iturralde.

Le Faucheur et. al

5

OSPF Extensions for Diff-Serv Traffic Engineering Feb 2001

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-EXT] Le Faucheur et al, Extension to RSVP and CR-LDP for support of Diff-Serv-aware MPLS Traffic Engineering, [draft-ietf-mpls-diff-te-ext-01.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-te00.txt](#), 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.

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

Le Faucheur et. al

6

OSPF Extensions for Diff-Serv Traffic Engineering Feb 2001

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

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

