

Network Working Group
Internet Draft
Category: Standards Track
Expires: March 2003
[draft-katz-yeung-ospf-traffic-08.txt](#)

D. Katz
Juniper Networks
D. Yeung
Procket Networks
K. Kompella
Juniper Networks
September 2002

Traffic Engineering Extensions to OSPF Version 2
***** Draft *****

Status

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>.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

This document describes extensions to the OSPF protocol version 2 to support intra-area Traffic Engineering, using Opaque Link State Advertisements.

Changes

(This section to be removed before publication).

Changes -06 version to -07 version

Comments from the ADs incorporated, as well as minor editing changes to conform to [draft-rfc-editor-rfc2223bis-02.txt](#).

Clean up front page headers.

(First page)

Clarify that this memo is for *intra-area* TE.

(Title, Abstract, [section 1.2](#))

Add a "Conventions" section ([rfc 2119](#)).

([Section 1.3](#))

Clarify what should be done with Reserved field in [section 2.2](#).

([Section 2.2](#))

Add an IANA Considerations section.

([Section 2.3.2](#), 2.4.2 and 8)

Clarify "IEEE Floating Point Format", and add reference.

([Section 2.4.2](#))

Clarify text for Resource Class/Color (match IS-IS TE text).

([Section 2.5.9](#))

Add text on originating TE LSAs.

([Section 3](#))

Broke up references into Normative and Informative (for now, IS-IS TE is Informative, pending reply from Routing ADs).

Add IPR Notices and Full Copyright Notice, as per [rfc 2026](#).

Changes from -07 to -08

Removed the Reserved field from the LSA ID: Instance is now 24 bits ([section 2.2](#)).

Added wording that if there is a BGP router ID, it should be used as the Router Address (as in ISIS TE draft) ([section 2.4.1](#)).

Rewrote the Security Considerations; added reference 11.

1. Introduction

This document specifies a method of adding traffic engineering capabilities to OSPF Version 2 [[1](#)]. The architecture of traffic engineering is described in [[2](#)]. The semantic content of the extensions is essentially identical to the corresponding extensions to IS-IS [[3](#)]. It is expected that the traffic engineering extensions to OSPF will continue to mirror those in IS-IS.

The extensions provide a way of describing the traffic engineering topology (including bandwidth and administrative constraints) and distributing this information within a given OSPF area. This topology does not necessarily match the regular routed topology, though this proposal depends on Network LSAs to describe multiaccess links.

1.1. Applicability

Many of the extensions specified in this document are in response to the requirements stated in [[2](#)], and thus are referred to as "traffic engineering extensions", and are also commonly associated with MPLS Traffic Engineering. A more accurate (albeit bland) designation is "extended link attributes", as what is proposed is simply to add more attributes to links in OSPF advertisements.

The information made available by these extensions can be used to build an extended link state database just as router LSAs are used to build a "regular" link state database; the difference is that the extended link state database (referred to below as the traffic engineering database) has additional link attributes. Uses of the traffic engineering database include:

- o monitoring the extended link attributes;
- o local constraint-based source routing; and
- o global traffic engineering.

For example, an OSPF-speaking device can participate in an OSPF area, build a traffic engineering database, and thereby report on the reservation state of links in that area.

In "local constraint-based source routing", a router R can compute a path from a source node A to a destination node B; typically, A is R itself, and B is specified by a "router address" (see below). This path may be subject to various constraints on the attributes of the links and nodes that the path traverses, e.g., use green links that have unreserved bandwidth of at least 10Mbps. This path could then be used to carry some subset of the traffic from A to B, forming a simple but effective means of traffic engineering. How the subset of

traffic is determined, and how the path is instantiated is beyond the scope of this document; suffice it to say that one means of defining the subset of traffic is "those packets whose IP destinations were learned from B", and one means of instantiating paths is using MPLS tunnels. As an aside, note that constraint-based routing can be NP-hard, or even unsolvable, depending on the nature of the attributes and constraints and thus many implementations will use heuristics. Consequently, we don't attempt to sketch an algorithm here.

Finally, for "global traffic engineering", a device can build a traffic engineering database, input a traffic matrix and an optimization function, crunch on the information, and thus compute optimal or near-optimal routing for the entire network. The device can subsequently monitor the traffic engineering topology and react to changes by recomputing the optimal routes.

1.2. Limitations

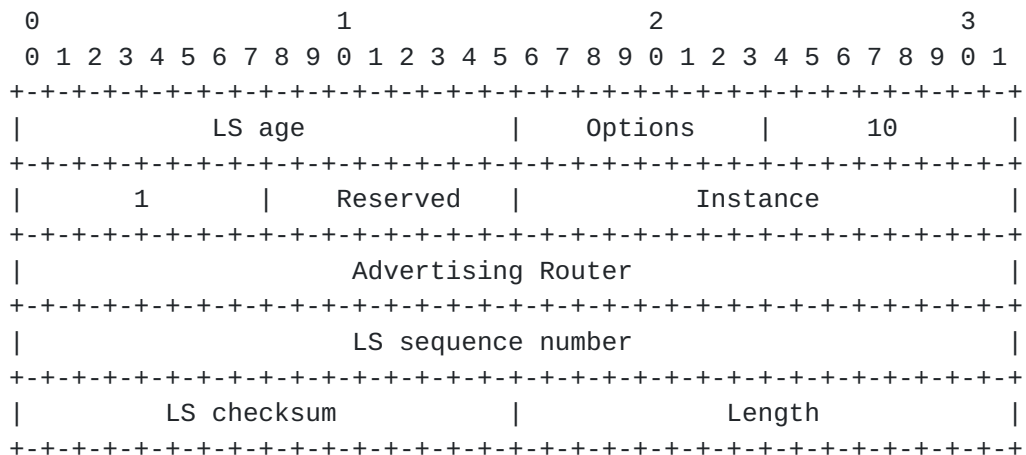
As mentioned above, this document specifies extensions and procedures for intra-area distribution of Traffic Engineering information. Methods for inter-area and inter-AS (Autonomous System) are not discussed here.

The extensions specified in this document capture the reservation state of point-to-point links. The reservation state of multiaccess links is not accurately reflected, except in the special case that there are only two devices in the multiaccess subnetwork.

This document also does not support unnumbered links. This deficiency is addressed in [\[4\]](#); see also [\[5\]](#) and [\[6\]](#).


1.3. Conventions

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](#) [\[7\]](#).



2.3.2. TLV Header

The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets for extensibility. The format of each TLV is:



The Length field defines the length of the value portion in octets (thus a TLV with no value portion would have a length of zero). The TLV is padded to four-octet alignment; padding is not included in the length field (so a three octet value would have a length of three, but the total size of the TLV would be eight octets). Nested TLVs are also 32-bit aligned. Unrecognized types are ignored.

This memo defines Types 1 and 2. See the IANA Considerations section for allocation of new Types.

2.4. LSA payload details

An LSA contains one top-level TLV.

There are two top-level TLVs defined:

- 1 - Router Address
- 2 - Link

2.4.1. Router Address TLV

The Router Address TLV specifies a stable IP address of the advertising router that is always reachable if there is any connectivity to it. This is typically implemented as a "loopback address"; the key attribute is that the address does not become unusable if an interface is down. In other protocols this is known as the "router ID," but for obvious reasons this nomenclature is avoided here. If a router advertises BGP routes with the BGP next hop attribute set to the BGP router ID, then the Router Address SHOULD be the same as the BGP router ID.

If IS-IS is also active in the domain, this address can also be used to compute the mapping between the OSPF and IS-IS topologies. For example, suppose a router R is advertising both IS-IS and OSPF Traffic Engineering LSAs, and suppose further that some router S is building a single Traffic Engineering Database (TED) based on both IS-IS and OSPF TE information. R may then appear as two separate nodes in S's TED; however, if both the IS-IS and OSPF LSAs generated by R contain the same Router Address, then S can determine that the IS-IS TE LSA and the OSPF TE LSA from R are indeed from a single router.

The router address TLV is type 1, and has a length of 4, and the value is the four octet IP address. It must appear in exactly one Traffic Engineering LSA originated by a router.

2.4.2. Link TLV

The Link TLV describes a single link. It is constructed of a set of sub-TLVs. There are no ordering requirements for the sub-TLVs.

Only one Link TLV shall be carried in each LSA, allowing for fine granularity changes in topology.

The Link TLV is type 2, and the length is variable.

The following sub-TLVs are defined:

The Link Type sub-TLV is TLV type 1, and is one octet in length.

2.5.2. Link ID

The Link ID sub-TLV identifies the other end of the link. For point-to-point links, this is the Router ID of the neighbor. For multiaccess links, this is the interface address of the designated router. The Link ID is identical to the contents of the Link ID field in the Router LSA for these link types.

The Link ID sub-TLV is TLV type 2, and is four octets in length.

2.5.3. Local Interface IP Address

The Local Interface IP Address sub-TLV specifies the IP address(es) of the interface corresponding to this link. If there are multiple local addresses on the link, they are all listed in this sub-TLV.

The Local Interface IP Address sub-TLV is TLV type 3, and is 4N octets in length, where N is the number of local addresses.

2.5.4. Remote Interface IP Address

The Remote Interface IP Address sub-TLV specifies the IP address(es) of the neighbor's interface corresponding to this link. This and the local address are used to discern multiple parallel links between systems. If the Link Type of the link is Multiaccess, the Remote Interface IP Address is set to 0.0.0.0 .

The Remote Interface IP Address sub-TLV is TLV type 4, and is 4N octets in length, where N is the number of neighbor addresses.

2.5.5. Traffic Engineering Metric

The Traffic Engineering Metric sub-TLV specifies the link metric for traffic engineering purposes. This metric may be different than the standard OSPF link metric. Typically, this metric is assigned by a network administrator.

The Traffic Engineering Metric sub-TLV is TLV type 5, and is four octets in length.

2.5.6. Maximum Bandwidth

The Maximum Bandwidth sub-TLV specifies the maximum bandwidth that can be used on this link in this direction (from the system originating the LSA to its neighbor), in IEEE floating point format. This is the true link capacity. The units are bytes per second.

The Maximum Bandwidth sub-TLV is TLV type 6, and is four octets in

length.

2.5.7. Maximum Reservable Bandwidth

The Maximum Reservable Bandwidth sub-TLV specifies the maximum bandwidth that may be reserved on this link in this direction, in IEEE floating point format. Note that this may be greater than the maximum bandwidth (in which case the link may be oversubscribed). This SHOULD be user-configurable; the default value should be the Maximum Bandwidth. The units are bytes per second.

The Maximum Reservable Bandwidth sub-TLV is TLV type 7, and is four octets in length.

2.5.8. Unreserved Bandwidth

The Unreserved Bandwidth sub-TLV specifies the amount of bandwidth not yet reserved at each of the eight priority levels, in IEEE floating point format. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7, arranged in increasing order with priority 0 occurring at the start of the sub-TLV, and priority 7 at the end of the sub-TLV. The initial values (before any bandwidth is reserved) are all set to the Maximum Reservable Bandwidth. Each value will be less than or equal to the Maximum Reservable Bandwidth. The units are bytes per second.

The Unreserved Bandwidth sub-TLV is TLV type 8, and is 32 octets in length.

2.5.9. Administrative Group

The Administrative Group sub-TLV contains a 4-octet bit mask assigned by the network administrator. Each set bit corresponds to one administrative group assigned to the interface. A link may belong to multiple groups.

By convention the least significant bit is referred to as 'group 0', and the most significant bit is referred to as 'group 31'.

The Administrative Group is also called Resource Class/Color [2].

The Administrative Group sub-TLV is TLV type 9, and is four octets in length.

3. Elements of Procedure

Routers shall originate Traffic Engineering LSAs whenever the LSA contents change, and whenever otherwise required by OSPF (an LSA refresh, for example). Note that this does not mean that every change must be flooded immediately; an implementation MAY set thresholds (for example, a bandwidth change threshold) that trigger immediate flooding, and initiate flooding of other changes after a short time interval. In any case, the origination of Traffic Engineering LSAs SHOULD be rate-limited to at most one every MinLSInterval [1].

Upon receipt of a changed Traffic Engineering LSA or Network LSA (since these are used in traffic engineering calculations), the router should update its traffic engineering database. No SPF or other route calculations are necessary.

4. Compatibility Issues

There should be no interoperability issues with routers that do not implement these extensions, as the Opaque LSAs will be silently ignored.

The result of having routers that do not implement these extensions is that the traffic engineering topology will be missing pieces; however, if the topology is connected, TE paths can still be calculated and ought to work.

5. Normative References

- [1] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [4] Kompella, K., Rekhter, Y., et al, "OSPF Extensions in Support of Generalized MPLS," work in progress.
- [6] Kompella, K., and Y. Rekhter, "Signalling Unnumbered Links in RSVP-TE," work in progress.
- [7] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [8] Coltun, R., "The OSPF Opaque LSA Option," [RFC 2370](#), July 1998.
- [9] IEEE, "IEEE Standard for Binary Floating-Point Arithmetic", Standard 754-1985, 1985 (ISBN 1-5593-7653-8).

6. Informative References

- [2] Awduche, D., et al, "Requirements for Traffic Engineering Over MPLS," [RFC 2702](#), September 1999.
- [3] Smit, H. and T. Li, "ISIS Extensions for Traffic Engineering," work in progress.
- [5] Kompella, K., Rekhter, Y., and A. Kullberg, "Signalling Unnumbered Links in CR-LDP," work in progress.
- [10] Narten, T., and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 2434](#), [BCP 26](#), October 1998.
- [11] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", [RFC 2154](#), June 1997.

7. Security Considerations

This document specifies the contents of Opaque LSAs in OSPFv2. As Opaque LSAs are not used for SPF computation or normal routing, the extensions specified here have no affect on IP routing. Tampering with TE LSAs may have an effect on traffic engineering computations, however, and it is suggested that whatever mechanisms are used for securing the transmission of normal OSPF LSAs be applied equally to all Opaque LSAs, including the TE LSAs specified here.

Note that the mechanisms in [[1](#)] and [[11](#)] apply to Opaque LSAs. It is suggested that future mechanisms proposed to secure/authenticate OSPFv2 LSA exchanges be made general enough to be used with Opaque LSAs.

8. IANA Considerations

The top level Types in a TE LSA as well as Types for sub-TLVs in a TE Link TLV are to be registered with IANA.

Following the guidelines set in [[10](#)], top level Types in TE LSAs from 3 through 32767 are to be assigned by Expert Review (the said Expert to be decided by the IESG). Types from 32768 through 65535 are reserved for Private Use. In all cases, assigned values Types MUST be registered with IANA.

Also, sub-Types of a TE Link TLV from 10 to 32767 are to be assigned by Expert Review; values from 32768 through 32772 are reserved for Private Use; and values from 32773 through 65535 are to be assigned

First Come First Served. In all cases, assigned values are to be registered with IANA.

9. Authors' Addresses

Dave Katz
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089 USA

Phone: +1 408 745 2000
Email: dkatz@juniper.net

Derek M. Yeung
Procket Networks, Inc.
1100 Cadillac Court
Milpitas, CA 95035 USA

Phone: +1 408 635-7900
Email: myeung@procket.com

Kireeti Kompella
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089 USA

Phone: +1 408 745 2000
Email: kireeti@juniper.net

10. IPR Notices

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any

copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

11. Full Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

