

Network Working Group
Internet Draft
Expiration Date: February 2001

Dave Katz
Juniper Networks, Inc.
Derek Yeung
Procket Networks, Inc.

Traffic Engineering Extensions to OSPF

[draft-katz-yeung-ospf-traffic-02.txt](#)

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>.

Abstract

This document describes extensions to the OSPF protocol to support Traffic Engineering, using opaque LSAs.

[1](#). Introduction

This document specifies a method of adding traffic engineering capabilities to OSPF [[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.

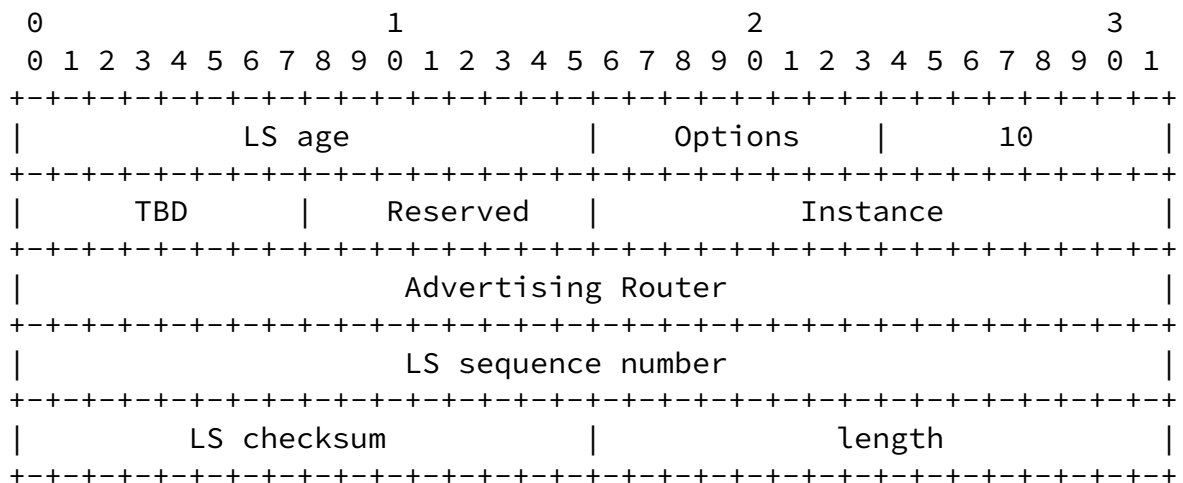
Internet Draft draft-katz-yeung-ospf-traffic-02.txt

August 2000

[2.3](#) LSA Format Overview

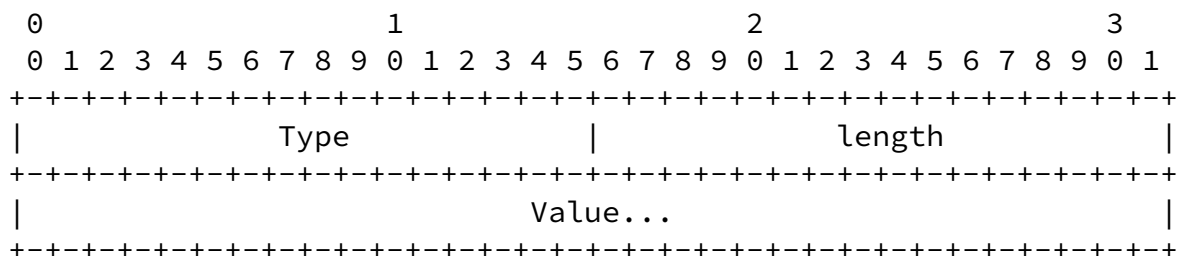
[2.3.1](#) LSA Header

The Traffic Engineering LSA starts with the standard LSA header:



[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 bytes (thus a TLV with no value portion would have a length of zero). The

TLV is padded to four-byte alignment; padding is not included in the length field (so a three byte value would have a length of three, but the total size of the TLV would be eight bytes). Nested TLVs are also 32-bit aligned. Unrecognized types are ignored. All types between 32768 and 65535 are reserved for vendor-specific extensions. All other undefined type codes are reserved for future assignment by IANA.

[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.

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:

- 1 - Link type (1 octet)
- 2 - Link ID (4 octets)
- 3 - Local interface IP address (4 octets)
- 4 - Remote interface IP address (4 octets)
- 5 - Traffic engineering metric (4 octets)
- 6 - Maximum bandwidth (4 octets)
- 7 - Maximum reservable bandwidth (4 octets)
- 8 - Unreserved bandwidth (32 octets)
- 9 - Resource class/color (4 octets)

32768-32772 - Reserved for Cisco-specific extensions

Each sub-TLV may occur only once. Unrecognized types are ignored. All of the defined sub-TLVs are mandatory (though future sub-TLVs may not necessarily be mandatory.)

[2.5](#) Sub-TLV Details

[2.5.1](#) Link Type

The Link Type sub-TLV defines the type of the link:

- 1 - Point-to-point
- 2 - Multiaccess

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 of the interface corresponding to this link. If the link is unnumbered, the Local Interface IP Address is set to the Router Address of the advertising router.

The Local Interface IP Address sub-TLV is TLV type 3, and is four

octets in length.

[2.5.4](#) Remote Interface IP Address

The Remote Interface IP Address sub-TLV specifies the IP address 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 is unnumbered, the first octet of the Remote Interface IP address is set to zero, and the last three octets are set to an interface index. The choice of the interface index is a local matter and is up to the implementation; the only requirements are that each unnumbered link on the advertising router has a unique interface

index; and that the OSPF module and the MPLS signalling module agree on the index. The SNMP IfIndex is a reasonable candidate.

The Remote Interface IP Address sub-TLV is TLV type 4, and is four octets in length.

[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.

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). 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. 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](#) Resource Class/Color

The Resource Class/Color sub-TLV specifies administrative group membership for this link, in terms of a bit mask. A link that is a member of multiple groups will have multiple bits set.

The Resource Class/Color 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).

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. Security Considerations

This document raises no new security issues for OSPF.

6. References

- [1] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.
- [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," [draft-ietf-isis-traffic-01.txt](#), work in progress.
- [4] Coltun, R., "The OSPF Opaque LSA Option," [RFC 2370](#), July 1998.

7. Authors' Addresses

Dave Katz
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94086 USA

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

Derek M. Yeung
Procket Networks, Inc.
3850 North First Street
San Jose, CA 95134 USA

Phone: +1 408 954-7900
Fax: +1 408 987-6166