

Internet Engineering Task Force
Internet Draft
Expiration Date: May 2001

Don Fedyk
(Nortel Networks)

Anoop Ghanwani
Rajesh Balay
(Cosine Communications)

Jerry Ash
(AT&T)

Alain Vedrenne
(SITA Equant)

November 2000

Multiple Metrics for Traffic Engineering with IS-IS and OSPF

[draft-fedyk-isis-ospf-te-metrics-01.txt](#)

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 obsolete 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 draft describes optional sub-TLVs that can be used to extend IGPs to support multiple metrics for use with traffic engineering. These optional extensions are an addendum to the existing work on traffic engineering extensions to OSPF and ISIS. While the current specifications allow advertising only one metric, the ability to

advertise multiple metrics has proven to be very useful in similar systems. The encoding for the proposed optional sub-TLVs is also specified. This draft is intended to complement and extend the existing work in [[OSPF](#)], [[ISIS](#)] and [[GMPLSR](#)].

1. Introduction

Present-day IP networks do not have many of the features required for supporting traffic engineering. To address this issue, several IETF Working Groups are working together to define the methodologies that are required for providing traffic engineering features [[TEREQ](#)]. Basically, there has been activity in the following areas:

- Extensions to IGP routing protocols (OSPF and IS-IS) for carrying an additional metric, and bandwidth and resource class information about links [[OSPF](#), [ISIS](#)]; and
- Signaling support for traffic engineering via MPLS [CRLDP, RSVP-TE].

The extensions for traffic engineering currently allow for only one routing metric. However, it is advantageous to allow the advertisement of multiple metrics because they can be used very effectively in a network running MPLS. In IP networks in the past, multiple routing metrics, even if defined, were not widely used. This is because traditional IP networks are based on a connectionless routing paradigm. For such networks, it has proven difficult to provide a scalable solution that uses multiple metrics. Since the routes must be computed in advance before any data arrives at the router, it places a significant burden on the router if it needs to compute separate routing tables for each metric. Furthermore, there is the issue of how to select which route a particular packet should use. See [[QOSR](#)] for a comprehensive study of the QoS routing problem and issues with multiple metrics and how these can be used.

The scalability issue described above is not a concern for networks running MPLS, because the routes for traffic engineering label switched paths (LSPs) may be computed on demand, and therefore a different routing metric may be used for each LSP. Existing connection-oriented systems that are non-IP based already have multiple metrics. For example, it may be desirable to minimize the economic cost of the path for one LSP, while for another LSP the goal may be minimizing the end-to-end delay. In other words, we are not recommending that the multiple advertised metrics be used simultaneously during path selection; that problem is known to be NP-complete. Instead, we would like to be able to use a different metric for path selection depending on the requirements of the LSP. Of course, clever approximation algorithms may be able to make use of more than one metric for path selection for a given LSP.

Therefore in this draft, a set of new sub-TLVs is proposed that will allow OSPF and ISIS to advertise multiple metrics for traffic

engineering. These sub-TLVs are optional; a particular implementation may optionally choose to support one or more of them.

The remainder of this draft is organized as follows. [Section 2](#) discusses the extensions to IS-IS and OSPF for advertising multiple metrics. [Section 3](#) discusses reasons for different routing metrics.

[2. Optional traffic engineering metrics for IS-IS and OSPF](#)

As a part of the recent IGP extensions for traffic engineering, a new routing metric called the Traffic Engineering Default Metric was introduced. The introduction of only a single metric is limiting. For example, it may be desirable to minimize hops, minimize delay or distance, minimize cost, or minimize resource usage, in the same network for different LSPs.

Three optional metrics are defined by this document for use in traffic engineering. These are in addition to the Default TE Metric [ISIS, OSPF]. The most pressing need is for a delay-based metric, but defining additional metrics has added advantages.

For example, having multiple metrics allows easy migration from one metric to another. A network operator may start out by using the first metric for delay expressed as milliseconds, and then later migrate to using delays expressed as hundreds of microseconds by using one of the `_free_` metrics in the network. Also, it makes the introduction of new metrics easier eliminating the need for vendor proprietary extensions, even if the metric is only for experimental use.

The actual metric type that is being advertised is specific to a service provider and does not need to be standardized. Consistency of the metric type is only required within the domain on which the path selection algorithm is being run.

[2.2 New sub-TLVs for IS-IS](#)

The optional metrics are sub-TLVs carried within the Extended IS Reachability TLV, whose TLV type is 22. The new sub-TLVs are:

- Traffic Engineering Optional Metric 1 (sub-TLV type TBD, length 3 octets);
- Traffic Engineering Optional Metric 2 (sub-TLV type TBD, length 3 octets);
- Traffic Engineering Optional Metric 3 (sub-TLV type TBD, length 3 octets).

[2.3 New sub-TLVs for OSPF](#)

The optional metrics are sub-TLVs carried within the Link TLV, whose TLV type is 2. The new sub-TLVs are:

- Traffic Engineering Optional Metric 1 (sub-TLV type TBD, length 4 octets);
- Traffic Engineering Optional Metric 2 (sub-TLV type TBD, length 4 octets);
- Traffic Engineering Optional Metric 3 (sub-TLV type TBD, length 4 octets).

3. Routing metrics background

This section is provided for background only. It is not intended to be an exhaustive list of routing metrics. Some attributes may be more appropriately represented as Resource Class constraints or colors. This is highlighted out where applicable.

A metric is a weight that is assigned to links in the network to indicate the relative preference of a link during path computation. Usually, metrics are selected so that they allow the computation of paths that meet or minimize or maximize (e.g., available link bandwidth on a per hop basis) certain constraints. Metrics for path selection can be classified using the following two criteria:

- Additive/non-additive: This refers to whether or not the path metric is obtained by adding the metric value for all links in the path. Examples of additive constraints include delay and hop count.
- Static/dynamic: A static metric doesn't change with time. This includes metrics such as hop count and propagation delay of a link. A dynamic metric changes with traffic conditions in the network. An example is the available bandwidth and queuing delays on a link.

Parameters such as bandwidth, delay, delay-jitter, and loss or error probability are used for driving path selection when setting up connections with QoS guarantees. Other metrics that can be considered include hop count, administrative cost, and economic cost.

A brief description of each of these metric types follows.

3.1 Bandwidth

For a connection-oriented network the unreserved bandwidth or bandwidth available allows path selection to be constrained to links that can support a request. The bandwidth available at each of the eight holding priorities is now advertised as a part of the IGP

Fedyk, et. al.	Expires September 2000	4
Internet Draft	draft-fedyk-mpls-te-metrics .00.txt	March, 2000

extensions for traffic engineering [[OSPF](#), [ISIS](#), [GMPLSR](#)]. LSPs that specify traffic parameters with bandwidth allocations can use this usage-based metric instead of the bandwidth metric to ensure sufficient bandwidth allocation.

3.2 Delay

Static delay-based metrics may be used for representing the propagation and the transmission delay of a link. This metric is used to minimize the end-to-end delay experienced by a packet. The propagation delay is a fixed quantity that is related to the transmission medium and the physical distance over which the signal travels. The transmission delay is the time that it takes to transmit a packet of a given size at a given link speed, and is thus a function of the bandwidth of the link. The combination of transmission delay for a given packet size and the propagation delay yields a value that represents the minimum delay that a packet would experience when transmitted over a given link.

The delay metric may represent the real value of the delay in milliseconds or microseconds. It may or may not include the transmission delay, especially since with higher link speeds, the transmission delay becomes negligible and propagation delay dominates.

3.3 Delay-jitter

Delay-jitter is caused due to queuing delays at various points in a network (usually the links). The exact value depends on the instantaneous traffic load. One option for advertising delay-jitter could be to advertise a static value that represents the average delay-jitter for a given link. With the advent of multi-queue schedulers such as static priority and fair queuing, delay-jitter can vary for different traffic types on the same link. A resource class bit could be used to indicate links that support such a queuing discipline, negating the need for a delay-jitter metric.

3.4 Loss Probability

The loss probability depends on the type of link used (optical fiber, satellite, etc.). Another way to express this would be bit error rate. This can be advertised as a static value. Alternatively, some bits in the resource class vector can be used to indicate the loss characteristics of the link again removing the requirement for a loss probability metric.

3.5 Hop Count

A Hop count metric is used to minimize the number of intermediate switches traversed, thereby minimizing the resources used in the network. Typically, minimizing the hop count does not require an explicit advertisement since it can be deduced while doing path

Fedyk, et. al.	Expires September 2000	5
Internet Draft	draft-fedyk-mpls-te-metrics .00.txt	March, 2000

computation. In some cases however, advertising the hop count may be desirable. For instance, when tunnels are advertised by the IGP, it may be useful to know the number of hops that the tunnel traverses.

3.6 Administrative weight

This metric is assigned based on a combination of factors. Sometimes, the administrative weight attempts to meld together several of the

other metric attributes into a single metric. It is important to realize that doing this is a significant compromise for traffic engineering since melding the various metrics together reduces the flexibility for path selection.

3.7 Economic cost or expense-based metric types

Cost-based metrics are assigned based on the economic cost for using a given link. These may be different from the administrative cost or weight of a link. With the advent of usage-based billing for services such as Frame Relay and ATM, and the use of tunnels over these technologies, the true cost of sending a certain amount and type of traffic can be calculated.

4. Security Considerations

This document raises no new security issues for IS-IS or OSPF.

5. Acknowledgements

The authors would like to thank Bilel Jamoussi, Peter Ashwood-Smith and Darek Skalecki for their helpful comments on the document.

6. References

[TEREQ] D. Awduche et al., "Requirements for Traffic Engineering Over MPLS," [RFC 2702](#), September 1999

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

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

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

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

Fedyk, et. al.	Expires September 2000	6
Internet Draft	draft-fedyk-mpls-te-metrics .00.txt	March, 2000

[QOSR] E. Crawley et al., "A Framework for QoS-based routing in the Internet," Request for Comments 2386, August 1998.

[GMPLSR] K. Kompella et al., "IS-IS Extensions in Support of Generalized MPLS," Internet Draft, [draft-ietf-isis-gmpls-extensions-00.txt](#), September 2000.

7. Authors' Addresses

Don Fedyk
Nortel Networks Corp.
[600](#) Technology Park Drive
Billerica, MA 01821
USA
Phone: +1-978-288-4506
Email: dwfedyk@nortelnetworks.com

Anoop Ghanwani
Cosine Communications
1200 Bridge Parkway
Redwood City, CA 94065
USA
Phone: +1-650-628-4225
Email: anoop@cosinecom.com

Rajesh Balay
Cosine Communications
[1200](#) Bridge Parkway
Redwood City, CA 94065
Phone: +1-650-628-4893
Email: rbalay@cosinecom.com

Gerald R. (Jerry) Ash
AT&T Labs
Room MT E3-3C37
200 Laurel Avenue
Middletown, NJ 07748
USA
Email: gash@att.com

Alain Vedrenne
SITA EQUANT
Strategic Technology Planning
[400](#) Galleria Parkway,
Atlanta, Georgia 30339
USA
Phone: +1 (678) 346-3466
Email: Alain.Vedrenne@sitaint

Full Copyright Statement

"Copyright (C) The Internet Society (March 2000). 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

Fedyk, et. al.	Expires September 2000	7
Internet Draft	draft-fedyk-mpls-te-metrics .00.txt	March, 2000

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.

