

Traffic Engineering Working Group
Internet Draft
Expiration Date: April 2000

Don Fedyk
Anoop Ghanwani

Nortel Networks Corp.

October 1999

Metrics and Resource Classes for Traffic Engineering

[draft-fedyk-mpls-te-metrics-00.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/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

There has recently been a lot of activity in defining the components for MPLS-based traffic engineering. This includes the recent extensions of the Interior Gateway Protocols (IGPs) to carry metrics and resource class information about links. This memo explores additional metrics that are useful for traffic engineering. We also describe some of the subtleties associated with the use of the resource class (or color) vector that has been defined in the IGP extensions.

Fedyk, et. al.
Internet Draft [draft-fedyk-mpls-te-metrics](#).00.txt September, 1999

1

[1](#). Introduction

Present-day IP networks do not have much support for traffic

engineering features. To address this issue, many of the 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:

- Extension of the IGP routing protocols (OSPF and IS-IS) for carrying additional metrics and resource class information about links [[OSPF](#),[ISIS](#)]; and
- Providing signaling support for traffic engineering via MPLS [[RSVP](#),[CRLDP](#)].

This draft provides a brief overview of different metrics that are useful for traffic engineering. We argue that it is better to allow the advertisement of multiple metrics because they can in fact be used when performing connection-oriented routing as would be the case in a network running MPLS. In the past, there have been attempts to define multiple metric types. However, IP is based on a connectionless routing paradigm and therefore it is difficult to provide a scalable solution that is able to use multiple metrics. One of the reasons for this is that every router in the network would need to maintain separate shortest-path spanning trees, one for each metric type. This is not an issue with a network running MPLS-based traffic engineering. The routes for traffic engineering label switched paths (LSPs) can be computed on demand, and therefore different metrics may be used for each LSP. 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. Note that we are not recommending that the multiple advertised metrics be used simultaneously; the multiple constraint problem is known to be NP-complete. It is simply desirable to use a different metric for path selection depending on the requirements of the LSP.

The remainder of this draft is organized as follows. [Section 2](#) discusses different routing metrics. [Section 3](#) discusses the use of metrics in traffic engineering and provides a recommendation for the metrics that should be advertised by routing protocols. [Section 4](#) discusses the use of the resource class vector, a new field that is being advertised by the IGP routing protocols, but about which little has been said so far.

[2](#). Routing Metrics

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 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. Bandwidth is an example of a non-additive constraint. On the other hand, it is also possible to look at bandwidth as an additive metric by using link costs that are in inverse proportion to available bandwidth; in such cases, the shortest path corresponds to the path with the most bandwidth.
- 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 bandwidth and queuing delays on a link.

Metrics are typically assigned positive numbers. Additive metrics can have a fairly large range of values starting from a small positive number and going all the way up to infinity, which in reality means there is no connection. As long as the additive metric is a positive number the routing systems can converge on loop free paths. In all metric schemes, there is a maximum metric, that means no connection, but there may be other limits due to the dynamic range. It is often challenging to find good ways to efficiently represent the complete dynamic range of metric values. One result of this is that many of the standards-based metrics are simple small integer values with a limited dynamic range.

The four factors of bandwidth, delay, delay-jitter and loss or error probability are often driving forces in the determination of metrics. The commonly used metrics include but are not limited to:

- Hop count
- Bandwidth
- Delay
- Administrative cost
- Economic cost

A brief description of each of these metric types follows.

[2.1](#) Hop-based metrics

A Hop based metric is used to minimize the number of intermediate switches traversed, thereby minimizing the resources used in the network. It also, therefore maximizes the number of equal cost paths in a network. The hop-based metric is also the degenerative form of many other metric types when the link attributes are considered equal. This type can be used to maximize the number of equal cost paths to a destination for load sharing schemes.

[2.2](#) Bandwidth-based metrics

A Bandwidth based metric is used to find the fewest large pipes to the destination. This metric can be assigned in a relative fashion or it can be algorithmically determined. Algorithmic determination has proven to be a challenge since the dynamic range of links has changed from 10^2 to 10^{12} over the years, and is expected to go beyond this in future. This complete range is not usually used at a single time in a domain a range of 10^5 between the highest and lowest link speeds are not uncommon. The formula for determination is usually in the form of the reciprocal of bandwidth. A 16 to 32 bit number is usually required to represent this effectively in an integer format.

Bandwidth metrics that are not algorithmically determined use a coarse representation of bandwidth rather than a true value. A bandwidth-based metric is a true minimizing metric since traffic that uses the higher speed links is more likely to get through.

2.2.1 Bandwidth Usage

An absolute bandwidth available at holding priority has been added for Traffic Engineering [[OSPF](#),[ISIS](#)]. LSPs that specify traffic parameters with bandwidth allocations can use this usage-based metric instead of the bandwidth metric to ensure sufficient bandwidth allocation. In a sense, this is a dynamic metric.

[2.3](#) Delay-based metrics

Static delay based metrics have been used for representing the propagation and the emission 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 medium and the physical distance that the signal travels. The emission delay is related to the time that it takes to transmit a packet from the first bit sent to the last bit sent. This is a function of the bandwidth of the link. The combination of emission and propagation delay for a given packet size yields a value that represents the minimum delay that a packet would experience when sent from one node to the other. This is a fixed number.

The delay metric represents a number of milliseconds or microseconds. Again the dynamic range of links has some bearing on the emission delay but the propagation delay tends to be stable in the millisecond range. The link speed is a gauge of how reliable the delay metric is since, as link speed increases the emission component becomes negligible.

[2.4](#) Economic cost or expense-based metrics

Cost based metrics are assigned based on attributes of links that represent a usage cost for a link. These are not to be confused with administrative weight. With the advent of usage charging for services such as Frame Relay and ATM, and the use of tunnels over

Fedyk, et. al.

April 2000

4

Internet Draft [draft-fedyk-mpls-te-metrics.00.txt](#) September, 1999

these technologies, a true cost per traffic sent can be calculated. The value of this type of metric is that traffic that does not want to incur charges for using a link or set of links.

[2.5](#) Administrative weight

This metric is assigned based on a combination of factors. The administrated weight attempts to meld several of the other metric attributes into a single metric. The weight can be adjusted. This type of metric is commonly used to differentiate links that differ from one another by some characteristic. It is important to realize that this is a significant compromise for traffic engineering since the melding of metric types reduces the ability to differentiate routes.

[3.](#) Traffic engineering trends

Recently, in the MPLS traffic engineering space, the addition of bandwidth reservation and allocation statistics represents dynamic information that can be used to engineer LSPs that have traffic parameters.

[3.1](#) Multiple traffic engineering (TE) metrics

Included with the recent work is the introduction of a new single metric for traffic engineering.

However, the introduction of a single metric is limiting since traffic engineering can have the goal of maximizing throughput, minimizing hops, minimizing delay or distance, minimizing cost or minimizing resource usage.

In such systems the metric type chosen is a guide to generating paths and will attempt to minimize the metric subject to the other constraints. There is little overhead for advertising multiple metrics for traffic engineering since only the static metrics need to be propagated.

[3.2](#) Bandwidth allocation for MPLS TE

Bandwidth allocation of MPLS TE LSPs with traffic parameters represents a usage-based constraint that is dynamic. As bandwidth is

consumed on a link, new LSPs must look for other links to satisfy their bandwidth requirements.

Interestingly enough, the value of maximizing throughput is not required for paths that specify minimum traffic profiles and allocate bandwidth.

[3.3](#) Metric recommendation

Fedyk, et. al.

April 2000

5

Internet Draft [draft-fedyk-mpls-te-metrics.00.txt](#) September, 1999

It is recommended that at least four additional metrics be defined for traffic engineering, in addition to the single metric that is used for connectionless IP routing. This would yield the following 5 metrics:

- Connectionless Metric (for use with connectionless routing)
- Traffic Engineering Metric 1 (32 bits)
- Traffic Engineering Metric 2 (32 bits)
- Traffic Engineering Metric 3 (32 bits)
- Reserved (32 bits)

The connectionless metric may be used to route LSPs along routes that would be selected by the IGP for connectionless routing.

It is network administrative decision on whether a metric is supported. The dynamic range of 32 is more than sufficient for the metrics.

A possible definition of the TE metrics could be:

- TE Metric 1: Delay Based Metric (Minimize delay)
- TE Metric 2: Cost/Administrative Based Metric (Minimize cost)
- TE Metric 3: Hop based Metric (Minimize resource usage)

[4.](#) Resource classes

The IGPs have recently been extended to carry a resource class or color vector for every link. This vector is 32 bits long and represents certain characteristics that the link satisfies. Little has been said about the expected use of this field. This section describes some of the ways it may be used and the subtleties associated with it.

4.1 Using the resource class or color vector

The resource class vector is one of the attributes flooded by the IGP, and is used by routers for computing a constraint-based path that uses only those links that satisfy the required criteria. The flooding scope of this information is a single area in OSPF and

level-1 in IS-IS. When computing a constraint-based route that traverses links in a single area, the router computing the path uses this information to determine which links satisfy the requested criteria. In this case, it is not necessary to carry the color vector of the request in the setup message. However, when setting up a hierarchical constraint-based route (i.e. one that traverses multiple areas), it becomes necessary to carry the resource class vector that is desired for the label switched path (LSP). When a router receives a request message during path setup, it would be expected to perform the following logical operation:

Fedyk, et. al.

April 2000

6

Internet Draft [draft-fedyk-mpls-te-metrics.00.txt](#) September, 1999

```
((color vector of request) & (color vector of link)) ==  
(color vector of request)
```

The request gets forwarded only if the above expression evaluates to TRUE. If a suitable link cannot be found, the request is rejected.

Using the above operation, there are two basic definitions of bits in the color vector that can be supported:

The link attribute is taken from a set of attributes. A given link may match one or more of these attributes requiring a bit to be set in that position.

A link satisfies a certain attribute up to a numeric value (e.g. a link has a security level of 5). In that case, all of the bits that correspond to that numeric value or lower must be set high at the time of configuration. Note that in the second case, there is some interdependence between the bits that must be taken into account while creating the color vector for the link. For example, in a system where links are rated with one of 8 levels of a certain attribute, the links at each level would have the following colors:

```
Level 0: 10000000  
Level 1: 11000000  
Level 2: 11100000  
..  
Level 8: 11111111
```

There are two drawbacks to using the simple operation defined above:

First, the usage of bits is sub-optimal; n levels of a particular characteristic require n bits in the color vector. While n levels of a certain characteristic can be represented with just $\lg(n)$ bits, it would require the capability to perform operations such as `_less than_` and `_greater than_`.

Second, the above operation cannot support the case where there are

three or more attributes in a set, of which some unordered subset of those is acceptable. For instance, if we have `_link type_` as a characteristic and we classify links as `terrestrial`, `low-orbit satellite` and `high-orbit satellite` links, and if we wish to request either `terrestrial` or `low-orbit satellite` links, there is no way to do that. However, if the same set was ordered by lower delay, (e.g. `terrestrial`, `low-orbit satellite`, `high orbit satellite`) then as an ordered subset this would still work.

4.2 Link characteristics

A few examples of link characteristics that might be desirable to have in the color vector include:

Fedyk, et. al.

April 2000

7

Internet Draft [draft-fedyk-mpls-blue-metrics](#).00.txt September, 1999

Satellite/terrestrial: A link can be either a satellite link or a terrestrial link. When making a request, the user can specify the use of only terrestrial links, or satellite links, or either.

Security level: A link can be rated as being at one of eight security levels ranging from completely insecure to highest security with levels in between. A typical request would be to use only links with a security of 5 or higher for the connection.

Link reliability: A link can be rated as being very reliable (possibly a link with automatic protection switching at the link layer) to one which is highly unreliable (i.e. loss characteristics of the link may be subject to weather or solar activity). This factor really is a function of loss rate and availability of a link. Some forms of transmission such as radio-based wireless and microwave are susceptible to weather conditions. These mediums are suitable for certain types of traffic.

4.3 Standardizing resource classes

It is worthwhile to explore the demand for standardized resource masks. This will allow requests to propagate seamlessly across areas since the semantics of the bits will have a universal value. One way to do this could be to have global and local subsets of the color mask. The bit semantics of the global portion would be standardized, while the bit semantics for the local portion would be left for definition by the developers/users of the router. The present color vector can be split equally into global and local portions as follows:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Globally significant										Local use																													

[illegible]

- Bit 0: satellite
- Bit 1: terrestrial
- Bit 2: highest reliability
- Bit 3: high reliability
- Bit 4: medium reliability
- Bit 5: low reliability
- Bits 6-15: reserved

6. Security considerations

Fedyk, et. al. April 2000
Internet Draft [draft-fedyk-mpls-te-metrics.00.txt](#) September, 1999

5. Acknowledgements

6. References

[ISIS] H. Smit, T. Lee, "IS-IS extensions for traffic engineering," Internet Draft, [draft-ietf-isis-traffic-01.txt](#), May 1999.

[CRLDP] B. Jamoussi et al., "Constraint-based LSP setup using LDP," Internet Draft, [draft-ietf-mpls-cr-ldp-03.txt](#), September 1999.

7. Authors' Addresses

Anoop Ghanwani
Nortel Networks Corp.
600 Technology Park Drive
Billerica, MA 01821
USA

Phone: +1-978-288-4506
dwfedyk@nortelnetworks.com

Phone: +1-978-288-4514
aghanwan@nortelnetworks.com

Full Copyright Statement

"Copyright (C) The Internet Society (1999). 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

Fedyk, et. al. April 2000 9
Internet Draft [draft-fedyk-mpls-te-metrics](#).00.txt September, 1999

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.

Fedyk, et. al. April 2000 10