

Network Working Group
Request for Comments: 3785
BCP: 87
Category: Best Current Practice

F. Le Faucheur
R. Uppili
Cisco Systems, Inc.
A. Vedrenne
P. Merckx
Equant
T. Telkamp
Global Crossing
May 2004

Use of Interior Gateway Protocol (IGP) Metric
as a second MPLS Traffic Engineering (TE) Metric

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This document describes a common practice on how the existing metric of Interior Gateway Protocols (IGP) can be used as an alternative metric to the Traffic Engineering (TE) metric for Constraint Based Routing of MultiProtocol Label Switching (MPLS) Traffic Engineering tunnels. This effectively results in the ability to perform Constraint Based Routing with optimization of one metric (e.g., link bandwidth) for some Traffic Engineering tunnels (e.g., Data Trunks) while optimizing another metric (e.g., propagation delay) for some other tunnels with different requirements (e.g., Voice Trunks). No protocol extensions or modifications are required. This text documents current router implementations and deployment practices.

1. Introduction

Interior Gateway Protocol (IGP) routing protocols (OSPF and IS-IS) as well as MultiProtocol Label Switching (MPLS) signaling protocols (RSVP-TE and CR-LDP) have been extended (as specified in [ISIS-TE], [OSPF-TE], [RSVP-TE] and [CR-LDP]) in order to support the Traffic Engineering (TE) functionality as defined in [TE-REQ].

These IGP routing protocol extensions currently include advertisement of a single additional MPLS TE metric to be used for Constraint Based Routing of TE tunnels.

However, the objective of traffic engineering is to optimize the use and the performance of the network. So it seems relevant that TE tunnel placement may be optimized according to different optimization criteria. For example, some Service Providers want to perform traffic engineering of different classes of service separately so that each class of Service is transported on a different TE tunnel. One example motivation for doing so is to apply different fast restoration policies to the different classes of service. Another example motivation is to take advantage of separate Constraint Based Routing in order to meet the different Quality of Service (QoS) objectives of each Class of Service. Depending on QoS objectives one may require either (a) enforcement by Constraint Based Routing of different bandwidth constraints for the different classes of service as defined in [DS-TE], or (b) optimizing on a different metric during Constraint Based Routing or (c) both. This document discusses how optimizing on a different metric can be achieved during Constraint Based Routing.

The most common scenario for a different metric calls for optimization of a metric reflecting delay (mainly propagation delay) when Constraint Based Routing TE Label Switched Paths (LSPs) that will be transporting voice, while optimizing a more usual metric (e.g., reflecting link bandwidth) when Constraint Based Routing TE LSPs that will be transporting data.

Additional IGP protocol extensions could be defined so that multiple TE metrics could be advertised in the IGP (as proposed for example in [METRICS]) and would thus be available to Constraint Based Routing in order to optimize on a different metric. However this document describes how optimizing on a different metric can be achieved today by existing implementations and deployments, without any additional IGP extensions beyond [ISIS-TE] and [OSPF-TE], by effectively using the IGP metric as a "second" TE metric.

2. Common Practice

In current MPLS TE deployments, network administrators often want Constraint Based Routing of TE LSPs carrying data traffic to be based on the same metric as the metric used for Shortest Path Routing. Where this is the case, this practice allows the Constraint Based Routing algorithm running on the Head-End LSR to use the IGP metric advertised in the IGP to compute paths for data TE LSPs instead of the advertised TE metric. The TE metric can then be used to convey

another metric (e.g., a delay-based metric) which can be used by the Constraint Based Routing algorithm on the Head-End LSR to compute path for the TE LSPs with different requirements (e.g., Voice TE LSP).

In some networks, network administrators configure the IGP metric to a value factoring the link propagation delay. In that case, this practice allows the Constraint Based Routing algorithm running on the Head-End LSR to use the IGP metric advertised in the IGP to compute paths for delay-sensitive TE LSPs (e.g., Voice TE LSPs) instead of the advertised TE metric. The TE metric can then be used to convey another metric (e.g., bandwidth based metric) which can be used by the Constraint Based Routing algorithm to compute paths for the data TE LSPs.

More generally, the TE metric can be used to carry any arbitrary metric that may be useful for Constraint Based Routing of the set of LSPs which need optimization on another metric than the IGP metric.

2.1. Head-End LSR Implementation Practice

A Head-End LSR implements the current practice by:

- (i) Allowing configuration, for each TE LSP to be routed, of whether the IGP metric or the TE metric is to be used by the Constraint Based Routing algorithm.
- (ii) Enabling the Constraint Based Routing algorithm to make use of either the TE metric or the IGP metric, depending on the above configuration for the considered TE-LSP

2.2. Network Deployment Practice

A Service Provider deploys this practice by:

- (i) Configuring, on every relevant link, the TE metric to reflect whatever metric is appropriate (e.g., delay-based metric) for Constraint Based Routing of some LSPs as an alternative metric to the IGP metric
- (ii) Configuring, for every TE LSP, whether this LSP is to be constraint based routed according to the TE metric or IGP metric

2.3. Constraints

The practice described in this document has the following constraints:

- (i) it only allows TE tunnels to be routed on either of two metrics (i.e., it cannot allow TE tunnels to be routed on one of three, or more, metrics). Extensions (for example such as those proposed in [METRICS]) could be defined in the future if necessary to relax this constraints, but this is outside the scope of this document.
- (ii) it can only be used where the IGP metric is appropriate as one of the two metrics to be used for constraint based routing (i.e., it cannot allow TE tunnels to be routed on either of two metrics while allowing IGP SPF to be based on a third metric). Extensions (for example such as those proposed in [METRICS]) could be defined in the future if necessary to relax this constraints, but this is outside the scope of this document.
- (iii) it can only be used on links which support an IGP adjacency so that an IGP metric is indeed advertised for the link. For example, this practice can not be used on Forwarding Adjacencies (see [LSP-HIER]).

Note that, as with [METRICS], this practice does not recommend that the TE metric and the IGP metric be used simultaneously during path computation for a given LSP. This is known to be an NP-complete problem.

2.4. Interoperability

Where path computation is entirely performed by the Head-End (e.g., intra-area operations with path computation on Head-end), this practice does not raise any interoperability issue among LSRs since the use of one metric or the other is a matter purely local to the Head-End LSR.

Where path computation involves another component than the Head-End (e.g., with inter-area operations where path computation is shared between the Head-End and Area Boundary Routers or a Path Computation Server), this practice requires that which metric to optimize on, be signaled along with the other constraints (bandwidth, affinity) for the LSP. See [PATH-COMP] for an example proposal on how to signal which metric to optimize, to another component involved in path computation when RSVP-TE is used as the protocol to signal path computation information.

3. Migration Considerations

Service Providers need to consider how to migrate from the current implementation to the new one supporting this practice.

Although the head-end routers act independently from each other, some migration scenarios may require that all head-end routers be upgraded to the new implementation to avoid any disruption on existing TE-LSPs before two metrics can effectively be used by TE. The reason is that routers with current implementation are expected to always use the TE metric for Constraint Based Routing of all tunnels; so when the TE metric is reconfigured to reflect the "second metric" (say to a delay-based metric) on links in the network, then all TE-LSPs would get routed based on the "second metric" metric, while the intent may be that only the TE-LSPs explicitly configured so should be routed based on the "second metric".

A possible migration scenario would look like this:

- 1) upgrade software on all head-end routers in the network to support this practice.
- 2) change the TE-LSPs configuration on the head-end routers to use the IGP metric (e.g., bandwidth-based) for Constraint Based Routing rather than the TE metric.
- 3) configure TE metric on the links to reflect the "second metric" (e.g., delay-based).
- 4) modify the LSP configuration of the subset of TE-LSPs which need to be Constraint Based routed using the "second metric" (e.g., delay-based), and/or create new TE-LSPs with such a configuration.

It is desirable that step 2 is non-disruptive (i.e., the routing of a LSP will not be affected in any way, and the data transmission will not be interrupted) by the change of LSP configuration to use "IGP metric" as long as the actual value of the "IGP metric" and "TE metric" are equal on every link at the time of LSP reconfiguration (as would be the case at step 2 in migration scenario above which assumed that TE metric was initially equal to IGP metric).

4. Security Considerations

The practice described in this document does not raise specific security issues beyond those of existing TE. Those are discussed in the respective security sections of [TE-REQ], [RSVP-TE] and [CR-LDP].

5. Acknowledgment

This document has benefited from discussion with Jean-Philippe Vasseur.

6. References

6.1. Normative References

- [TE-REQ] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M. and J. McManus, Requirements for Traffic Engineering over MPLS, RFC 2702, September 1999.

- [OSPF-TE] Katz, D., Kompella, K. and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

- [ISIS-TE] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, May 2004.

- [RSVP-TE] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V. and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

- [CR-LDP] Jamoussi, B., Andersson, L., Callon, R., Dantu, R., Wu, L., Doolan, P., Worster, T., Feldman, N., Fredette, A., Girish, M., Gray, E., Heinanen, J., Kilty, T. and A. Malis, "Constraint-Based LSP Setup using LDP", RFC 3212, January 2002.

6.1. Informative References

- [METRICS] Fedyk, et al., "Multiple Metrics for Traffic Engineering with IS-IS and OSPF", Work in Progress, November 2000.

- [DIFF-TE] Le Faucheur, F. and W. Lai, "Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering", RFC 3564, July 2003.

- [PATH-COMP] Vasseur, et al., "RSVP Path computation request and reply messages", Work in Progress, June 2002.

- [LSP-HIER] Kompella, et al., "LSP Hierarchy with Generalized MPLS TE", Work in Progress, September 2002.

7. Authors' Addresses

Francois Le Faucheur
Cisco Systems, Inc.
Village d'Entreprise Green Side - Batiment T3
400, Avenue de Roumanille
06410 Biot-Sophia Antipolis
France

Phone: +33 4 97 23 26 19
EMail: flefauch@cisco.com

Ramesh Uppili
Cisco Systems,
2000 Innovation Drive
Kanata,
ONTARIO,
Canada - K2K 3E8

Phone: 01-613-254 4578
Email: ruppili@cisco.com

Alain Vedrenne
Equant
Heraklion, 1041 route des Dolines, BP347
06906 Sophia Antipolis Cedex
FRANCE

Phone: +33 4 92 96 57 22
EMail: alain.vedrenne@equant.com

Pierre Merckx
Equant
1041 route des Dolines - BP 347
06906 SOPHIA ANTIPOLIS Cedex
FRANCE

Phone: +33 (0)492 96 6454
EMail: pierre.merckx@equant.com

Thomas Telkamp
Global Crossing, Ltd.
Croeselaan 148
NL-3521CG Utrecht
The Netherlands

Phone: +31 30 238 1250
EMail: telkamp@gblx.net

8. Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.