

Internet Engineering Task Force (IETF)
Request for Comments: 7471
Category: Standards Track
ISSN: 2070-1721

S. Giacalone
Unaffiliated
D. Ward
Cisco Systems
J. Drake
A. Atlas
Juniper Networks
S. Previdi
Cisco Systems
March 2015

OSPF Traffic Engineering (TE) Metric Extensions

Abstract

In certain networks, such as, but not limited to, financial information networks (e.g., stock market data providers), network performance information (e.g., link propagation delay) is becoming critical to data path selection.

This document describes common extensions to RFC 3630 "Traffic Engineering (TE) Extensions to OSPF Version 2" and RFC 5329 "Traffic Engineering Extensions to OSPF Version 3" to enable network performance information to be distributed in a scalable fashion. The information distributed using OSPF TE Metric Extensions can then be used to make path selection decisions based on network performance.

Note that this document only covers the mechanisms by which network performance information is distributed. The mechanisms for measuring network performance information or using that information, once distributed, are outside the scope of this document.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7471>.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions Used in This Document	4
3. TE Metric Extensions to OSPF TE	4
4. Sub-TLV Details	6
4.1. Unidirectional Link Delay Sub-TLV	6
4.1.1. Type	6
4.1.2. Length	6
4.1.3. Anomalous (A) Bit	7
4.1.4. Reserved	7
4.1.5. Delay Value	7
4.2. Min/Max Unidirectional Link Delay Sub-TLV	7
4.2.1. Type	7
4.2.2. Length	7
4.2.3. Anomalous (A) Bit	8
4.2.4. Reserved	8
4.2.5. Min Delay	8
4.2.6. Reserved	8
4.2.7. Max Delay	8
4.3. Unidirectional Delay Variation Sub-TLV	9
4.3.1. Type	9
4.3.2. Length	9
4.3.3. Reserved	9
4.3.4. Delay Variation	9
4.4. Unidirectional Link Loss Sub-TLV	9
4.4.1. Type	10
4.4.2. Length	10
4.4.3. Anomalous (A) Bit	10
4.4.4. Reserved	10
4.4.5. Link Loss	10

4.5. Unidirectional Residual Bandwidth Sub-TLV	10
4.5.1. Type	11
4.5.2. Length	11
4.5.3. Residual Bandwidth	11
4.6. Unidirectional Available Bandwidth Sub-TLV	11
4.6.1. Type	12
4.6.2. Length	12
4.6.3. Available Bandwidth	12
4.7. Unidirectional Utilized Bandwidth Sub-TLV	12
4.7.1. Type	12
4.7.2. Length	13
4.7.3. Utilized Bandwidth	13
5. Announcement Thresholds and Filters	13
6. Announcement Suppression	14
7. Network Stability and Announcement Periodicity	14
8. Enabling and Disabling Sub-TLVs	15
9. Static Metric Override	15
10. Compatibility	15
11. Security Considerations	15
12. IANA Considerations	16
13. References	16
13.1. Normative References	16
13.2. Informative References	17
Acknowledgments	18
Authors' Addresses	19

1. Introduction

In certain networks, such as, but not limited to, financial information networks (e.g., stock market data providers), network performance information (e.g., link propagation delay) is becoming as critical to data path selection as other metrics.

Because of this, using metrics such as hop count or cost as routing metrics is becoming only tangentially important. Rather, it would be beneficial to be able to make path selection decisions based on network performance information (such as link propagation delay) in a cost-effective and scalable way.

This document describes extensions to OSPFv2 and OSPFv3 TE (hereafter called "OSPF TE Metric Extensions"), that can be used to distribute network performance information (viz link propagation delay, delay variation, link loss, residual bandwidth, available bandwidth, and utilized bandwidth).

The data distributed by OSPF TE Metric Extensions is meant to be used as part of the operation of the routing protocol (e.g., by replacing cost with link propagation delay or considering bandwidth as well as

cost), by enhancing Constrained Shortest Path First (CSPF), or for use by a PCE [RFC4655] or an Application-Layer Traffic Optimization (ALTO) server [RFC7285]. With respect to CSPF, the data distributed by OSPF TE Metric Extensions can be used to set up, fail over, and fail back data paths using protocols such as RSVP-TE [RFC3209].

Note that the mechanisms described in this document only distribute network performance information. The methods for measuring that information or acting on it once it is distributed are outside the scope of this document. A method for measuring loss and delay in an MPLS network is described in [RFC6374].

While this document does not specify the method for measuring network performance information, any measurement of link propagation delay SHOULD NOT vary significantly based upon the offered traffic load and, hence, SHOULD NOT include queuing delays. For a forwarding adjacency (FA) [RFC4206], care must be taken that measurement of the link propagation delay avoids significant queuing delay; this can be accomplished in a variety of ways, e.g., measuring with a traffic class that experiences minimal queuing or summing the measured link propagation delay of the links on the FA's path.

2. Conventions Used in This Document

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 [RFC2119].

In this document, these words should convey that interpretation only when in ALL CAPS. Lowercase uses of these words are not to be interpreted as carrying this significance.

3. TE Metric Extensions to OSPF TE

This document defines new OSPF TE sub-TLVs that are used to distribute network performance information. The extensions in this document build on the ones provided in OSPFv2 TE [RFC3630] and OSPFv3 TE [RFC5329].

OSPFv2 TE Link State Advertisements (LSAs) [RFC3630] are opaque LSAs [RFC5250] with area flooding scope while OSPFv3 Intra-Area-TE-LSAs have their own LSA type, also with area flooding scope; both consist of a single TLV with one or more nested sub-TLVs. The Link TLV is common to both and describes the characteristics of a link between OSPF neighbors.

This document defines several additional sub-TLVs for the Link TLV:

Type	Length	Value
27	4	Unidirectional Link Delay
28	8	Min/Max Unidirectional Link Delay
29	4	Unidirectional Delay Variation
30	4	Unidirectional Link Loss
31	4	Unidirectional Residual Bandwidth
32	4	Unidirectional Available Bandwidth
33	4	Unidirectional Utilized Bandwidth

As can be seen in the list above, the sub-TLVs described in this document carry different types of network performance information. Many (but not all) of the sub-TLVs include a bit called the Anomalous (or A) bit. When the A bit is clear (or when the sub-TLV does not include an A bit), the sub-TLV describes steady state link performance. This information could conceivably be used to construct a steady state performance topology for initial tunnel path computation, or to verify alternative failover paths.

When network performance violates configurable link-local thresholds a sub-TLV with the A bit set is advertised. These sub-TLVs could be used by the receiving node to determine whether to move traffic to a backup path or whether to calculate an entirely new path. From an MPLS perspective, the intent of the A bit is to permit LSP ingress nodes to:

- A) Determine whether the link referenced in the sub-TLV affects any of the LSPs for which it is ingress. If there are, then:
- B) The node determines whether those LSPs still meet end-to-end performance objectives. If not, then:
- C) The node could then conceivably move affected traffic to a pre-established protection LSP or establish a new LSP and place the traffic in it.

The A bit was intentionally omitted from some sub-TLVs to help mitigate oscillations. See Section 7.1 for more information.

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								
+										+										+										+									
										27										4																			
+										+										+										+									
A										RESERVED										Delay																			
+										+										+										+									

[Page 6]

4.1.3. Anomalous (A) Bit

This field represents the Anomalous (A) bit. The A bit is set when the measured value of this parameter exceeds its configured maximum threshold. The A bit is cleared when the measured value falls below its configured reuse threshold. If the A bit is clear, the sub-TLV represents steady state link performance.

4.1.4. Reserved


This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

4.1.5. Delay Value

This 24-bit field carries the average link delay over a configurable interval in microseconds, encoded as an integer value. When set to the maximum value 16,777,215 (16.777215 sec), then the delay is at least that value, and it may be larger.

4.2. Min/Max Unidirectional Link Delay Sub-TLV

This sub-TLV advertises the minimum and maximum delay values between two directly connected OSPF neighbors. The delay advertised by this sub-TLV MUST be the delay from the advertising node to its neighbor (i.e., the forward path delay). The format of this sub-TLV is shown in the following diagram:



4.2.1. Type

This sub-TLV has a type of 28.

4.2.2. Length

The length is 8.

4.2.3. Anomalous (A) Bit

This field represents the Anomalous (A) bit. The A bit is set when one or more measured values exceed a configured maximum threshold. The A bit is cleared when the measured value falls below its configured reuse threshold. If the A bit is clear, the sub-TLV represents steady state link performance.

4.2.4. Reserved

This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

4.2.5. Min Delay

This 24-bit field carries minimum measured link delay value (in microseconds) over a configurable interval, encoded as an integer value.

Implementations MAY also permit the configuration of an offset value (in microseconds) to be added to the measured delay value to advertise operator specific delay constraints.

When set to the maximum value 16,777,215 (16.777215 sec), then the delay is at least that value, and it may be larger.

4.2.6. Reserved

This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

4.2.7. Max Delay

This 24-bit field carries the maximum measured link delay value (in microseconds) over a configurable interval, encoded as an integer value.


Implementations may also permit the configuration of an offset value (in microseconds) to be added to the measured delay value to advertise operator specific delay constraints.

It is possible for min delay and max delay to be the same value.

When the delay value is set to the maximum value 16,777,215 (16.777215 sec), then the delay is at least that value, and it may be larger.

4.3. Unidirectional Delay Variation Sub-TLV

This sub-TLV advertises the average link delay variation between two directly connected OSPF neighbors. The delay variation advertised by this sub-TLV MUST be the delay from the advertising node to its neighbor (i.e., the forward path delay variation). The format of this sub-TLV is shown in the following diagram:



4.3.1. Type

This sub-TLV has a type of 29.

4.3.2. Length

The length is 4.

4.3.3. Reserved


This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

4.3.4. Delay Variation

This 24-bit field carries the average link delay variation over a configurable interval in microseconds, encoded as an integer value. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), then the delay is at least that value, and it may be larger.

4.4. Unidirectional Link Loss Sub-TLV

This sub-TLV advertises the loss (as a packet percentage) between two directly connected OSPF neighbors. The link loss advertised by this sub-TLV MUST be the packet loss from the advertising node to its neighbor (i.e., the forward path loss). The format of this sub-TLV is shown in the following diagram:



4.4.1. Type

This sub-TLV has a type of 30

4.4.2. Length

The length is 4.

4.4.3. Anomalous (A) Bit

This field represents the Anomalous (A) bit. The A bit is set when the measured value of this parameter exceeds its configured maximum threshold. The A bit is cleared when the measured value falls below its configured reuse threshold. If the A bit is clear, the sub-TLV represents steady state link performance.

4.4.4. Reserved

This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.


4.4.5. Link Loss

This 24-bit field carries link packet loss as a percentage of the total traffic sent over a configurable interval. The basic unit is 0.000003%, where $(2^{24} - 2)$ is 50.331642%. This value is the highest packet loss percentage that can be expressed (the assumption being that precision is more important on high speed links than the ability to advertise loss rates greater than this, and that high speed links with over 50% loss are unusable). Therefore, measured values that are larger than the field maximum SHOULD be encoded as the maximum value.

4.5. Unidirectional Residual Bandwidth Sub-TLV

This sub-TLV advertises the residual bandwidth between two directly connected OSPF neighbors. The residual bandwidth advertised by this sub-TLV MUST be the residual bandwidth from the advertising node to its neighbor.

The format of this sub-TLV is shown in the following diagram:



4.5.1. Type

This sub-TLV has a type of 31.

4.5.2. Length

The length is 4.


4.5.3. Residual Bandwidth

This field carries the residual bandwidth on a link, forwarding adjacency [RFC4206], or bundled link in IEEE floating point format with units of bytes per second. For a link or forwarding adjacency, residual bandwidth is defined to be Maximum Bandwidth [RFC3630] minus the bandwidth currently allocated to RSVP-TE LSPs. For a bundled link, residual bandwidth is defined to be the sum of the component link residual bandwidths.

The calculation of Residual Bandwidth is different than that of Unreserved Bandwidth [RFC3630]. Residual Bandwidth subtracts tunnel reservations from Maximum Bandwidth (i.e., the link capacity) [RFC3630] and provides an aggregated remainder across priorities. Unreserved Bandwidth, on the other hand, is subtracted from the Maximum Reservable Bandwidth (the bandwidth that can theoretically be reserved) and provides per priority remainders. Residual Bandwidth and Unreserved Bandwidth [RFC3630] can be used concurrently, and each has a separate use case (e.g., the former can be used for applications like Weighted ECMP while the latter can be used for call admission control).

4.6. Unidirectional Available Bandwidth Sub-TLV

This sub-TLV advertises the available bandwidth between two directly connected OSPF neighbors. The available bandwidth advertised by this sub-TLV MUST be the available bandwidth from the advertising node to its neighbor. The format of this sub-TLV is shown in the following diagram:



4.6.1. Type

This sub-TLV has a type of 32.

4.6.2. Length


The length is 4.

4.6.3. Available Bandwidth

This field carries the available bandwidth on a link, forwarding adjacency, or bundled link in IEEE floating point format with units of bytes per second. For a link or forwarding adjacency, available bandwidth is defined to be residual bandwidth (see Section 4.5) minus the measured bandwidth used for the actual forwarding of non-RSVP-TE LSP packets. For a bundled link, available bandwidth is defined to be the sum of the component link available bandwidths.

4.7. Unidirectional Utilized Bandwidth Sub-TLV

This Sub-TLV advertises the bandwidth utilization between two directly connected OSPF neighbors. The bandwidth utilization advertised by this sub-TLV MUST be the bandwidth from the advertising node to its neighbor. The format of this Sub-TLV is shown in the following diagram:



4.7.1. Type

This sub-TLV has a type of 33.

4.7.2. Length

The length is 4.

4.7.3. Utilized Bandwidth

This field carries the bandwidth utilization on a link, forwarding adjacency, or bundled link in IEEE floating-point format with units of bytes per second. For a link or forwarding adjacency, bandwidth utilization represents the actual utilization of the link (i.e., as measured by the advertising node). For a bundled link, bandwidth utilization is defined to be the sum of the component link bandwidth utilizations.

5. Announcement Thresholds and Filters

The values advertised in all sub-TLVs (except min/max delay and residual bandwidth) MUST represent an average over a period or be obtained by a filter that is reasonably representative of an average. For example, a rolling average is one such filter.

Min and max delay MAY be the lowest and/or highest measured value over a measurement interval or MAY make use of a filter, or other technique, to obtain a reasonable representation of a min and max value representative of the interval with compensation for outliers.

The measurement interval, any filter coefficients, and any advertisement intervals MUST be configurable for each sub-TLV.

In addition to the measurement intervals governing re-advertisement, implementations SHOULD provide for each sub-TLV configurable accelerated advertisement thresholds, such that:

1. If the measured parameter falls outside a configured upper bound for all but the min delay metric (or lower bound for min delay metric only) and the advertised sub-TLV is not already outside that bound, or
2. If the difference between the last advertised value and current measured value exceed a configured threshold, then
3. The advertisement is made immediately.

4. For sub-TLVs, which include an A bit (except min/max delay), an additional threshold SHOULD be included corresponding to the threshold for which the performance is considered anomalous (and sub-TLVs with the A bit are sent). The A bit is cleared when the sub-TLV's performance has been below (or re-crosses) this threshold for an advertisement interval(s) to permit fail back.

To prevent oscillations, only the high threshold or the low threshold (but not both) may be used to trigger any given sub-TLV that supports both.

Additionally, once outside of the bounds of the threshold, any re-advertisement of a measurement within the bounds would remain governed solely by the measurement interval for that sub-TLV.

6. Announcement Suppression

When link performance values change by small amounts that fall under thresholds that would cause the announcement of a sub-TLV, implementations SHOULD suppress sub-TLV re-advertisement and/or lengthen the period within which they are refreshed.

Only the accelerated advertisement threshold mechanism described in Section 5 may shorten the re-advertisement interval.

All suppression and re-advertisement interval back-off timer features SHOULD be configurable.

7. Network Stability and Announcement Periodicity

Sections 5 and 6 provide configurable mechanisms to bound the number of re-advertisements. Instability might occur in very large networks if measurement intervals are set low enough to overwhelm the processing of flooded information at some of the routers in the topology. Therefore, care should be taken in setting these values.

Additionally, the default measurement interval for all sub-TLVs should be 30 seconds.

Announcements must also be able to be throttled using configurable inter-update throttle timers. The minimum announcement periodicity is 1 announcement per second. The default value should be set to 120 seconds.

Implementations should not permit the inter-update timer to be lower than the measurement interval.

Furthermore, it is recommended that any underlying performance measurement mechanisms not include any significant buffer delay, any significant buffer induced delay variation, or any significant loss due to buffer overflow or due to active queue management.

8. Enabling and Disabling Sub-TLVs

Implementations **MUST** make it possible to individually enable or disable the advertisement of each sub-TLV.

9. Static Metric Override

Implementations **SHOULD** permit the static configuration and/or manual override of dynamic measurements for each sub-TLV in order to simplify migration and to mitigate scenarios where dynamic measurements are not possible.

10. Compatibility

As per [RFC3630], an unrecognized TLV should be silently ignored. That is, it should not be processed but it should be included in LSAs sent to OSPF neighbors.

11. Security Considerations

This document does not introduce security issues beyond those discussed in [RFC3630]. OSPFv2 HMAC-SHA [RFC5709] provides additional protection for OSPFv2. OSPFv3 IPsec [RFC4552] and OSPFv3 Authentication Trailer [RFC7166] provide additional protection for OSPFv3.

OSPF Keying and Authentication for Routing Protocols (KARP) [RFC6863] provides an analysis of OSPFv2 and OSPFv3 routing security, and OSPFv2 Security Extensions [OSPFSEC] provides extensions designed to address the identified gaps in OSPFv2.

12. IANA Considerations

IANA maintains the registry for the Link TLV sub-TLVs. For OSPF TE Metric Extensions, one new type code for each sub-TLV defined in this document has been registered, as follows:

Value	Sub-TLV
27	Unidirectional Link Delay
28	Min/Max Unidirectional Link Delay
29	Unidirectional Delay Variation
30	Unidirectional Link Loss
31	Unidirectional Residual Bandwidth
32	Unidirectional Available Bandwidth
33	Unidirectional Utilized Bandwidth

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, September 2008, <<http://www.rfc-editor.org/info/rfc5329>>.
- [IEEE754] Institute of Electrical and Electronics Engineers, "Standard for Floating-Point Arithmetic", IEEE Standard 754, August 2008.

13.2. Informative References

- [RFC3209] 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, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005, <<http://www.rfc-editor.org/info/rfc4206>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, June 2006, <<http://www.rfc-editor.org/info/rfc4552>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC5709] Bhatia, M., Manral, V., Fanto, M., White, R., Barnes, M., Li, T., and R. Atkinson, "OSPFv2 HMAC-SHA Cryptographic Authentication", RFC 5709, October 2009, <<http://www.rfc-editor.org/info/rfc5709>>.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011, <<http://www.rfc-editor.org/info/rfc6374>>.
- [RFC6863] Hartman, S. and D. Zhang, "Analysis of OSPF Security According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6863, March 2013, <<http://www.rfc-editor.org/info/rfc6863>>.
- [RFC7166] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 7166, March 2014, <<http://www.rfc-editor.org/info/rfc7166>>.
- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, September 2014, <<http://www.rfc-editor.org/info/rfc7285>>.

[OSPFSEC] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed.,
"Security Extension for OSPFv2 when Using Manual Key
Management", Work in Progress, draft-ietf-ospf-security-
extension-manual-keying, November 2014.

Acknowledgments

The authors would like to recognize Nabil Bitar, Edward Crabbe, Don Fedyk, Acee Lindem, David McDysan, and Ayman Soliman for their contributions to this document.

The authors would also like to acknowledge Curtis Villamizar for his significant comments and direct content collaboration.

Authors' Addresses

Spencer Giacalone
Unaffiliated

EMail: spencer.giacalone@gmail.com

Dave Ward
Cisco Systems
170 West Tasman Dr.
San Jose, CA 95134
United States

EMail: dward@cisco.com

John Drake
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
United States

EMail: jdrake@juniper.net

Alia Atlas
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
United States

EMail: akatlas@juniper.net

Stefano Previdi
Cisco Systems
Via Del Serafico 200
00142 Rome
Italy

EMail: sprevidi@cisco.com

