

Internet Engineering Task Force (IETF)
Request for Comments: 8330
Category: Standards Track
ISSN: 2070-1721

H. Long
M. Ye
Huawei Technologies Co., Ltd.
G. Mirsky
ZTE
A. D'Alessandro
Telecom Italia S.p.A.
H. Shah
Ciena
February 2018

OSPF Traffic Engineering (OSPF-TE) Link Availability Extension
for Links with Variable Discrete Bandwidth

Abstract

A network may contain links with variable discrete bandwidth, e.g., microwave and copper. The bandwidth of such links may change discretely in response to a changing external environment. The word "availability" is typically used to describe such links during network planning. This document defines a new type of Generalized Switching Capability-Specific Information (SCSI) TLV to extend the Generalized Multiprotocol Label Switching (GMPLS) Open Shortest Path First (OSPF) routing protocol. The extension can be used for route computation in a network that contains links with variable discrete bandwidth. Note that this document only covers the mechanisms by which the availability information is distributed. The mechanisms by which availability information of a link is determined and the use of the distributed information for route computation are outside the scope of this document. It is intended that technology-specific documents will reference this document to describe specific uses.

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 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8330>.

Copyright Notice

Copyright (c) 2018 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 (<https://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
1.1. Conventions Used in This Document	3
2. Abbreviations	4
3. Overview	4
4. TE Metric Extension to OSPF-TE	5
4.1. Availability SCSI-TLV	5
4.2. Processing Procedures	6
5. Security Considerations	6
6. IANA Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Acknowledgments	10
Authors' Addresses	10

1. Introduction

Some data-plane technologies, e.g., microwave and copper, allow seamless changes of maximum physical bandwidth through a set of known discrete values. The parameter "availability", as described in [G.827], [F.1703], and [P.530], is often used to describe the link capacity. The availability is a time scale, representing a proportion of the operating time that the requested bandwidth is ensured. To set up a Label Switched Path (LSP) across these links, availability information is required by the nodes to verify the bandwidth before making a bandwidth reservation. Assigning different availability classes over such links provides for more efficient planning of link capacity to support different types of services. The link availability information will be determined by the operator and is statically configured. It will usually be determined from the availability requirements of the services expected to be carried on the LSP. For example, voice service usually needs "five nines" availability, while non-real-time services may adequately perform at four or three nines availability. For the route computation, both the availability information and the bandwidth resource information are needed. Since different service types may need different availability guarantees, multiple <availability, bandwidth> pairs may be required to be associated with a link.

In this document, a new type of Generalized SCSI-TLV, the Availability SCSI-TLV, is defined. It is intended that technology-specific documents will reference this document to describe specific uses. The signaling extension to support links with variable discrete bandwidth is defined in [RSVP-TE-Availability].

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Abbreviations

The following abbreviations are used in this document:

GMPLS	Generalized Multiprotocol Label Switching
ISCD	Interface Switching Capability Descriptor
LSA	Link State Advertisement
LSP	Label Switched Path
OSPF	Open Shortest Path First
SCSI	Switching Capability-Specific Information
SPF	Shortest Path First
TE	Traffic Engineering
TLV	Type-Length-Value

3. Overview

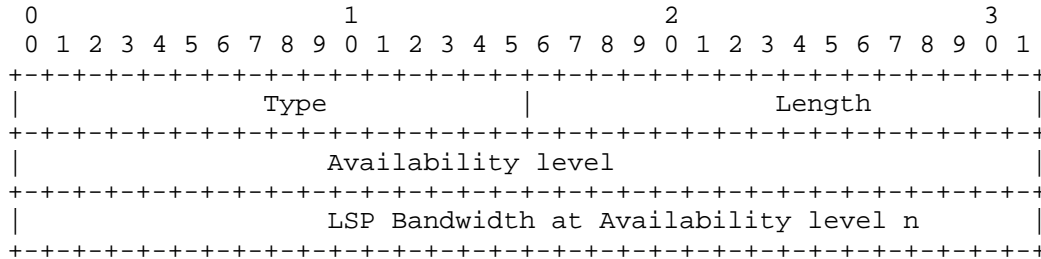
A node that has link(s) with variable discrete bandwidth attached should include an <availability, bandwidth> information list in its OSPF-TE LSA messages. The list provides the mapping between the link nominal bandwidth and its availability level. This information is used for path calculation by the node(s). The setup of an LSP requires this information to be flooded in the network and used by the nodes or the PCE for the path computation. In this document, a new type of Generalized SCSI-TLV, the Availability SCSI-TLV, is defined. The computed path can then be provisioned via the signaling protocol [RSVP-TE-Availability].

Note: The mechanisms described in this document only distribute availability information. The methods for measuring the information or using the information for route computation are outside the scope of this document.

4. TE Metric Extension to OSPF-TE

4.1. Availability SCSI-TLV

The Generalized SCSI is defined in [RFC8258]. This document defines a new type of Generalized SCSI-TLV called the Availability SCSI-TLV. The Availability SCSI-TLV can be included one or more times. It has the following format:



Type: 0x000A, 16 bits

Length: 2 octets (16 bits)

Availability level: 32 bits

This field is a binary32-format floating-point number as defined by [IEEE754-2008]. The bytes are transmitted in network order; that is, the byte containing the sign bit is transmitted first. This field describes the decimal value of the availability guarantee of the Switching Capability in the Interface Switching Capability Descriptor object [RFC4202]. The value MUST be less than 1. The Availability level field is usually expressed as the value 0.99/0.999/0.9999/0.99999.

LSP Bandwidth at Availability level n: 32 bits

This field is a 32-bit IEEE floating-point number as defined by [IEEE754-2008]. The bytes are transmitted in network order; that is, the byte containing the sign bit is transmitted first. This field describes the LSP bandwidth for the availability level represented in the Availability level field. The units are bytes per second.

4.2. Processing Procedures

The ISCD allows routing protocols such as OSPF to carry technology-specific information in the "Switching Capability-specific information" field; see [RFC4203]. A node advertising an interface with a Switching Capability that supports variable discrete bandwidth attached SHOULD contain one or more Availability SCSI-TLVs in its OSPF-TE LSA messages. Each Availability SCSI-TLV provides information about how much bandwidth a link can support for a specified availability. This information may be used for path calculation by the node(s).

The Availability SCSI-TLV MUST NOT be sent in ISCDs with Switching Capability field values that have not been defined to support the Availability SCSI-TLV. Non-supporting nodes would see such an ISCD/LSA as malformed.

The absence of the Availability SCSI-TLV in an ISCD containing Switching Capability field values that have been defined to support the Availability SCSI-TLV SHALL be interpreted as representing the fixed-bandwidth link with the highest availability value.

Only one Availability SCSI-TLV for the specific availability level SHOULD be sent. If multiple TLVs are present, the Availability SCSI-TLV with the lowest bandwidth value SHALL be processed. If an Availability SCSI-TLV with an invalid value (e.g., larger than 1) is received, the Availability SCSI-TLV will be ignored.

5. Security Considerations

This document specifies the contents of Opaque LSAs in OSPFv2. Tampering with GMPLS-TE LSAs may have an effect on TE computations. [RFC3630] suggests such mechanisms as the mechanism described in [RFC2154] to protect the transmission of this information, and those or other mechanisms should be used to secure and/or authenticate the information carried in the Opaque LSAs. An analysis of the security of OSPF is provided in [RFC6863] and applies to the OSPF extension defined in this document. Any new mechanisms developed to protect the transmission of information carried in Opaque LSAs will also automatically protect the extension defined in this document.

Please refer to [RFC5920] for details on security threats; defensive techniques; monitoring, detection, and reporting of security attacks; and requirements.

6. IANA Considerations

This document introduces a new type of Generalized SCSI-TLV (Availability) that is carried in the OSPF-TE LSA messages. Technology-specific documents will reference this document to describe the specific use of this Availability SCSI-TLV.

IANA created a registry called the "Generalized SCSI (Switching Capability Specific Information) TLV Types" registry [RFC8258]. The registry has been updated to include the following Availability SCSI-TLV:

Type	Description	Switching Type	Reference
0x000A	Availability	5, 52	RFC 8330

New switching types are required in order to use the Availability SCSI-TLV. IANA has registered the following in the "Switching Types" registry:

Value	Name	Reference
5	PSC with GSCSI support	RFC 8330
52	L2SC with GSCSI support	RFC 8330

7. References

7.1. Normative References

- [IEEE754-2008] IEEE, "IEEE Standard for Floating-Point Arithmetic", IEEE 754-2008, DOI 10.1109/IEEESTD.2008.4610935.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, DOI 10.17487/RFC4202, October 2005, <<https://www.rfc-editor.org/info/rfc4202>>.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<https://www.rfc-editor.org/info/rfc4203>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8258] Ceccarelli, D. and L. Berger, "Generalized SCSI: A Generic Structure for Interface Switching Capability Descriptor (ISCD) Switching Capability Specific Information (SCSI)", RFC 8258, DOI 10.17487/RFC8258, October 2017, <<https://www.rfc-editor.org/info/rfc8258>>.

7.2. Informative References

- [F.1703] International Telecommunication Union, "Availability objectives for real digital fixed wireless links used in 27 500 km hypothetical reference paths and connections", ITU-R Recommendation F.1703-0, January 2005, <<https://www.itu.int/rec/R-REC-F.1703-0-200501-I/en>>.
- [G.827] International Telecommunication Union, "Availability performance parameters and objectives for end-to-end international constant bit-rate digital paths", ITU-T Recommendation G.827, September 2003, <<https://www.itu.int/rec/T-REC-G.827/en>>.
- [P.530] International Telecommunication Union, "Propagation data and prediction methods required for the design of terrestrial line-of-sight systems", ITU-R Recommendation P.530-17, December 2017, <<https://www.itu.int/rec/R-REC-P.530/en>>.
- [RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, DOI 10.17487/RFC2154, June 1997, <<https://www.rfc-editor.org/info/rfc2154>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.

[RFC6863] Hartman, S. and D. Zhang, "Analysis of OSPF Security According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6863, DOI 10.17487/RFC6863, March 2013, <<https://www.rfc-editor.org/info/rfc6863>>.

[RSVP-TE-Availability]

Long, H., Ye, M., Mirsky, G., D'Alessandro, A., and H. Shah, "Ethernet Traffic Parameters with Availability Information", Work in Progress, draft-ietf-ccamp-rsvp-te-bandwidth-availability-08, January 2018.

Acknowledgments

The authors would like to thank Acee Lindem, Daniele Ceccarelli, and Lou Berger for their comments on the document.

Authors' Addresses

Hao Long
Huawei Technologies Co., Ltd.
No. 1899, Xiyuan Avenue, Hi-tech Western District
Chengdu 611731
China

Phone: +86-18615778750
Email: longhao@huawei.com

Min Ye
Huawei Technologies Co., Ltd.
No. 1899, Xiyuan Avenue, Hi-tech Western District
Chengdu 611731
China

Email: amy.yemin@huawei.com

Greg Mirsky
ZTE

Email: gregimirsky@gmail.com

Alessandro D'Alessandro
Telecom Italia S.p.A.

Email: alessandro.dalessandro@telecomitalia.it

Himanshu Shah
Ciena Corp.
3939 North First Street
San Jose, CA 95134
United States of America

Email: hshah@ciena.com