

Internet Engineering Task Force (IETF)
Request for Comments: 7688
Category: Standards Track
ISSN: 2070-1721

Y. Lee, Ed.
Huawei
G. Bernstein, Ed.
Grotto Networking
November 2015

GMPLS OSPF Enhancement for Signal and Network Element Compatibility
for Wavelength Switched Optical Networks

Abstract

This document provides Generalized Multiprotocol Label Switching (GMPLS) Open Shortest Path First (OSPF) routing enhancements to support signal compatibility constraints associated with Wavelength Switched Optical Network (WSON) elements. These routing enhancements are applicable in common optical or hybrid electro-optical networks where not all the optical signals in the network are compatible with all network elements participating in the network.

This compatibility constraint model is applicable to common optical or hybrid electro-optical systems such as optical-electronic-optical (OEO) switches, regenerators, and wavelength converters, since such systems can be limited to processing only certain types of WSON signals.

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/rfc7688>.

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
1.1. Conventions Used in This Document	3
2. The Optical Node Property TLV	3
2.1. Resource Block Information	4
2.2. Resource Accessibility	5
2.3. Resource Wavelength Constraints	5
2.4. Resource Block Pool State	5
2.5. Resource Block Shared Access Wavelength Availability	5
3. Interface Switching Capability Descriptor (ISCD) Format Extensions	5
3.1. Switching Capability Specific Information (SCSI)	6
4. WSON-Specific Scalability and Timeliness	7
5. Security Considerations	8
6. IANA Considerations	8
6.1. Optical Node Property TLV	8
6.1.1. Optical Node Property Sub-TLV	8
6.2. WSON-LSC Switching Type TLV	9
6.2.1. WSON-LSC SCSI Sub-TLVs	9
7. References	10
7.1. Normative References	10
7.2. Informative References	10
Authors' Addresses	12

1. Introduction

The documents [RFC6163], [RFC7446], and [RFC7581] explain how to extend the Wavelength Switched Optical Network (WSON) control plane to support both multiple WSON signal types and common hybrid electro-optical systems as well hybrid systems containing optical switching and electro-optical resources. In WSON, not all the optical signals in the network are compatible with all network elements participating in the network. Therefore, signal compatibility is an important constraint in path computation in a WSON.

This document provides GMPLS OSPF routing enhancements to support signal compatibility constraints associated with general WSON network elements. These routing enhancements are applicable in common optical or hybrid electro-optical networks where not all optical signals in the network are compatible with all network elements participating in the network.

This compatibility constraint model is applicable to common optical or hybrid electro-optical systems such as OEO switches, regenerators, and wavelength converters, since such systems can be limited to processing only certain types of WSON signals.

Related to this document is [RFC7580], which provides GMPLS OSPF routing enhancements to support the generic routing and label assignment process that can be applicable to a wider range of technologies beyond WSON.

1.1. 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 [RFC2119].

2. The Optical Node Property TLV

[RFC3630] defines OSPF Traffic Engineering (TE) Link State Advertisement (LSA) using an opaque LSA. This document adds a new top-level TLV for use in the OSPF TE LSA: the Optical Node Property TLV. The Optical Node Property TLV describes a single node. It is comprised of a set of optional sub-TLVs. There are no ordering requirements for the sub-TLVs.

When using the extensions defined in this document, at least one Optical Node Property TLV MUST be advertised in each LSA. To allow for fine-grained changes in topology, more than one Optical Node Property TLV MAY be advertised in a single LSA. Implementations MUST support receiving multiple Optical Node Property TLVs in an LSA.

The Optical Node Property TLV contains all WSON-specific node properties and signal compatibility constraints. The detailed encodings of these properties are defined in [RFC7581].

The following sub-TLVs of the Optical Node Property TLV are defined:

Value	Length	Sub-TLV Type
1	variable	Resource Block Information
2	variable	Resource Accessibility
3	variable	Resource Wavelength Constraints
4	variable	Resource Block Pool State
5	variable	Resource Block Shared Access Wavelength Availability

The detailed encodings of these sub-TLVs are found in [RFC7581] as indicated in the table below.

Sub-TLV Type	Section from [RFC7581]
Resource Block Information	4
Resource Accessibility	3.1
Resource Wavelength Constraints	3.2
Resource Block Pool State	3.3
Resource Block Shared Access Wavelength Availability	3.4

All sub-TLVs defined here may occur at most once in any given Optical Node TLV under one TE LSA. If more than one copy of the sub-TLV is received in the same LSA, the redundant sub-TLV SHOULD be ignored. If the same sub-TLV is advertised in a different TE LSA (which would only occur if there was a packaging error), then the sub-TLV with the largest LSA ID (Section 2.2 of RFC 3630) SHOULD be picked. These restrictions need not apply to future sub-TLVs. Unrecognized sub-TLVs are ignored.

Among the sub-TLVs defined above, the Resource Block Pool State sub-TLV and Resource Block Shared Access Wavelength Availability are dynamic in nature, while the rest are static. As such, they can be separated out from the rest and be advertised with multiple TE LSAs per OSPF router, as described in [RFC3630] and [RFC5250].

2.1. Resource Block Information

As defined in [RFC7446], this sub-TLV is used to represent resource signal constraints and processing capabilities of a node.

2.2. Resource Accessibility

This sub-TLV describes the structure of the resource pool in relation to the switching device. In particular, it indicates the ability of an ingress port to reach a resource block and of a resource block to reach a particular egress port.

2.3. Resource Wavelength Constraints

Resources, such as wavelength converters, etc., may have limited input or output wavelength ranges. Additionally, due to the structure of the optical system, not all wavelengths can necessarily reach or leave all the resources. The Resource Wavelength Constraints sub-TLV describes these properties.

2.4. Resource Block Pool State

This sub-TLV describes the usage state of a resource that can be encoded as either a list of integer values or a bitmap indicating whether a single resource is available or in use. This information can be relatively dynamic, i.e., can change when a connection is established or torn down.

2.5. Resource Block Shared Access Wavelength Availability

Resource blocks may be accessed via a shared fiber. If this is the case, then wavelength availability on these shared fibers is needed to understand resource availability.

3. Interface Switching Capability Descriptor (ISCD) Format Extensions

The ISCD describes the switching capability of an interface [RFC4202]. This document defines a new Switching Capability value for WSON as follows:

Value	Type
-----	-----
151	WSON-LSC

Switching Capability and Encoding values MUST be used as follows:

Switching Capability = WSON-LSC

Encoding Type = Lambda (as defined in [RFC3471])

When Switching Capability and Encoding fields are set to values as stated above, the Interface Switching Capability Descriptor MUST be interpreted as in [RFC4203] with the optional inclusion of one or more Switching Capability Specific Information sub-TLVs.

3.1. Switching Capability Specific Information (SCSI)

The technology-specific part of the WSON ISCD may include a variable number of sub-TLVs called Bandwidth sub-TLVs. Two types of Bandwidth sub-TLV are defined:

- Type 1: Available Labels
- Type 2: Shared Backup Labels

A SCSI may contain multiple Available Label sub-TLVs and multiple Shared Backup Label sub-TLVs. The following figure shows the format for a SCSI that contains these sub-TLVs, where the Available Label Sub-TLV and Shared Backup Label sub-TLV are as defined in [RFC7579]. The order of the sub-TLVs in the SCSI is arbitrary.

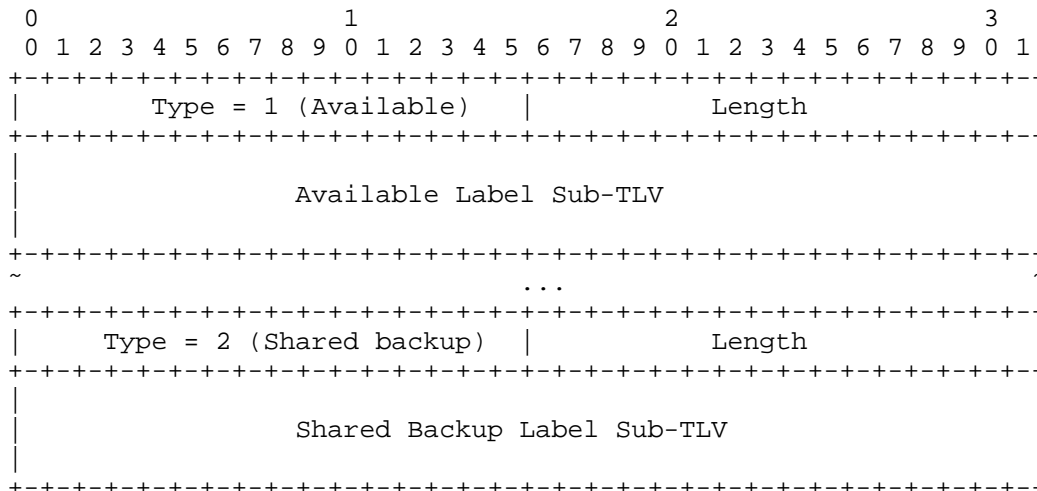


Figure 1: SCSI Format

If duplicated sub-TLVs are advertised, the router/node will ignore the duplicated labels that are identified by the Label format defined in [RFC6205].

The label format defined in [RFC6205] MUST be used when advertising interfaces with a WSON-LSC type Switching Capability.

4. WSON-Specific Scalability and Timeliness

This document has defined five sub-TLVs specific to WSON. The examples given in [RFC7581] show that very large systems, in terms of channel count, ports, or resources, can be very efficiently encoded.

There has been concern expressed that some possible systems may produce LSAs that exceed the IP Maximum Transmission Unit (MTU). In a typical node configuration, the Optical Node Property TLV will not exceed the IP MTU. A typical node configuration refers to a system with several hundreds of channels with an OEO element in the node. This would give the Optical Node Property TLV less than 350 bytes. In addition, [RFC7581] provides mechanisms to compactly encode required information elements. In a rare case where the TLV exceeds the IP MTU, IP fragmentation/reassembly can be used, which is an acceptable method. For IPv6, a node may use the IPv6 Fragment header to fragment the packet at the source and have it reassembled at the destination(s).

If the size of this LSA is greater than the MTU, then these sub-TLVs can be packed into separate LSAs. From the point of view of path computation, the presence of the Resource Block Information sub-TLV indicates that resources exist in the system and may have signal compatibility or other constraints. The other four sub-TLVs indicate constraints on access to and availability of those resources.

Hence, the "synchronization" procedure is quite simple from the point of view of path computation. Until a Resource Block Information sub-TLV is received for a system, path computation cannot make use of the other four sub-TLVs since it does not know the nature of the resources, e.g., whether the resources are wavelength converters, regenerators, or something else. Once this sub-TLV is received, path computation can proceed with whatever sub-TLVs it may have received (their use is dependent upon the system type).

If path computation proceeds with out-of-date or missing information from these sub-TLVs, then there is the possibility of either (a) path computation yielding a path that does not exist in the network, (b) path computation failing to find a path through the network that actually exists. Both situations are currently encountered with GMPLS, i.e., out-of-date information on constraints or resource availability.

If the new sub-TLVs or their attendant encodings are malformed, a proper implementation SHOULD log the problem and MUST stop sending the LSA that contains malformed TLVs or sub-TLVs.

Errors of this nature SHOULD be logged for the local operator. Implementations MUST provide a rate limit on such logs, and that rate limit SHOULD be configurable.

Note that the connection establishment mechanism (signaling or management) is ultimately responsible for the establishment of the connection, and this implies that such mechanisms must ensure signal compatibility.

5. Security Considerations

This document does not introduce security issues other than those discussed in [RFC3630] and [RFC4203].

As with [RFC4203], it specifies the contents of Opaque LSAs in OSPFv2. As Opaque LSAs are not used for Shortest Path First (SPF) computation or normal routing, the extensions specified here have no direct effect on IP routing. Tampering with GMPLS TE LSAs may have an effect on the underlying transport. [RFC3630] notes that the security mechanisms described in [RFC2328] apply to Opaque LSAs carried in OSPFv2.

For general security aspects relevant to GMPLS-controlled networks, please refer to [RFC5920].

6. IANA Considerations

6.1. Optical Node Property TLV

This document introduces a new Top-Level Node TLV (Optical Node Property TLV) under the OSPF TE LSA defined in [RFC3630]. IANA has registered a new TLV for "Optical Node Property". The new TLV is in the "Top Level Types in TE LSAs" registry in "Open Shortest Path First (OSPF) Traffic Engineering TLVs" located at <http://www.iana.org/assignments/ospf-traffic-eng-tlvs>, and is as follows:

Value	TLV Type	Reference
6	Optical Node Property	RFC 7688

6.1.1. Optical Node Property Sub-TLV

Additionally, a new IANA registry has been created named "Types for sub-TLVs of Optical Node Property (Value 6)" in the "Open Shortest Path First (OSPF) Traffic Engineering TLVs" registry located at <http://www.iana.org/assignments/ospf-traffic-eng-tlvs>. New sub-TLVs and their values have been assigned as follows:

Value	Length	Sub-TLV	Reference
0		Reserved	
1	variable	Resource Block Information	RFC 7688
2	variable	Resource Accessibility	RFC 7688
3	variable	Resource Wavelength Constraints	RFC 7688
4	variable	Resource Block Pool State	RFC 7688
5	variable	Resource Block Shared Access Wavelength Availability	RFC 7688
6-65535		Unassigned	

The registration procedure for this registry is Standards Action as defined in [RFC5226].

6.2. WSON-LSC Switching Type TLV

IANA has registered a new switching type in the "Switching Types" registry in "GMPLS Signaling Parameters", located at <http://www.iana.org/assignments/gmpls-sig-parameters>, as follows:

Value	Description	Reference
151	WSON-LSC	RFC 7688

Also, IANA has added the following entry to the IANAGmplsSwitchingTypeTC MIB:

```
wsonlsc(151), -- WSON-LSC
```

6.2.1. WSON-LSC SCSI Sub-TLVs

Additionally, a new IANA registry has been created for sub-TLVs of the WSON-LSC SCSI sub-TLV. It is named "Types for sub-TLVs of WSON-LSC SCSI (Switching Capability Specific Information)" and is in the "Open Shortest Path First (OSPF) Traffic Engineering TLVs" registry. It contains the following sub-TLVs:

Value	Sub-TLV	Reference
0	Reserved	
1	Available Labels	RFC 7688
2	Shared Backup Labels	RFC 7688
3-65535	Unassigned	

The registration procedure for this registry is Standards Action as defined in [RFC5226].

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, 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, DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [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, <<http://www.rfc-editor.org/info/rfc4203>>.
- [RFC6205] Otani, T., Ed., and D. Li, Ed., "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, DOI 10.17487/RFC6205, March 2011, <<http://www.rfc-editor.org/info/rfc6205>>.
- [RFC7579] Bernstein, G., Ed., Lee, Y., Ed., Li, D., Imajuku, W., and J. Han, "General Network Element Constraint Encoding for GMPLS-Controlled Networks", RFC 7579, DOI 10.17487/RFC7579, June 2015, <<http://www.rfc-editor.org/info/rfc7579>>.
- [RFC7580] Zhang, F., Lee, Y., Han, J., Bernstein, G., and Y. Xu, "OSPF-TE Extensions for General Network Element Constraints", RFC 7580, DOI 10.17487/RFC7580, June 2015, <<http://www.rfc-editor.org/info/rfc7580>>.
- [RFC7581] Bernstein, G., Ed., Lee, Y., Ed., Li, D., Imajuku, W., and J. Han, "Routing and Wavelength Assignment Information Encoding for Wavelength Switched Optical Networks", RFC 7581, DOI 10.17487/RFC7581, June 2015, <<http://www.rfc-editor.org/info/rfc7581>>.

7.2. Informative References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.

- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, DOI 10.17487/RFC3471, January 2003, <<http://www.rfc-editor.org/info/rfc3471>>.
- [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, <<http://www.rfc-editor.org/info/rfc4202>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.
- [RFC6163] Lee, Y., Ed., Bernstein, G., Ed., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, DOI 10.17487/RFC6163, April 2011, <<http://www.rfc-editor.org/info/rfc6163>>.
- [RFC7446] Lee, Y., Ed., Bernstein, G., Ed., Li, D., and W. Imajuku, "Routing and Wavelength Assignment Information Model for Wavelength Switched Optical Networks", RFC 7446, DOI 10.17487/RFC7446, February 2015, <<http://www.rfc-editor.org/info/rfc7446>>.

Authors' Addresses

Young Lee (editor)
Huawei Technologies
5340 Legacy Drive, Building 3
Plano, TX 75024
United States

Phone: (469) 277-5838
Email: leeyoung@huawei.com

Greg M. Bernstein (editor)
Grotto Networking
Fremont, CA
United States

Phone: (510) 573-2237
Email: gregb@grotto-networking.com