

CCAMP Working Group
Ed.
Internet-Draft
Ericsson
Intended status: Standards Track
Zhang
Expires: April 24, 2014
Technologies

D. Ceccarelli,

F.

Huawei

S.

Alcatel-

R.

Infinera

J.

Belotti

Lucent

Rao

Corporation

Drake

Juniper

October 21,

2013

**Traffic Engineering Extensions to OSPF for Generalized MPLS (GMPLS)
Control of Evolving G.709 OTN Networks
draft-ietf-ccamp-gmpls-ospf-g709v3-10**

Abstract

This document describes Open Shortest Path First - Traffic Engineering (OSPF-TE) routing protocol extensions to support Generalized MPLS (GMPLS) control of Optical Transport Networks (OTN) specified in ITU-T Recommendation G.709 as published in 2012. It extends mechanisms defined in [RFC4203](#).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months

and may be updated, replaced, or obsoleted 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."

This Internet-Draft will expire on April 24, 2014.

Copyright Notice

Copyright (c) 2013 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

Ceccarelli, et al.
1]

Expires April 24, 2014

[Page

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 1. Introduction
- 3 1.1. Terminology
- 3 2. OSPF-TE Extensions
- 3 3. TE-Link Representation
- 5 4. ISCD format extensions
- 6 4.1. Switching Capability Specific Information
- 7 4.1.1. Switching Capability Specific Information for
fixed containers
- 8 4.1.2. Switching Capability Specific Information for
variable containers
- 8 4.1.3. Switching Capability Specific Information - Field
values and explanation
- 9 5. Examples
- 12 5.1. MAX LSP Bandwidth fields in the ISCD
- 12 5.2. Example of T,S and TS granularity utilization
- 14 5.2.1. Example of different TS Granularities
- 15 5.3. Example of ODUflex advertisement
- 18 5.4. Example of single stage muxing
- 20 5.5. Example of multi stage muxing - Unbundled link
- 22 5.6. Example of multi stage muxing - Bundled links
- 24 5.7. Example of component links with non-homogeneous
hierarchies

25	6.	OSPFv2 scalability
28	7.	Compatibility
29	8.	Security Considerations
29	9.	IANA Considerations
29	9.1.	Switching types
29	9.2.	New sub-TLVs
30	10.	Contributors
30	11.	Acknowledgements
33	12.	References
33	12.1.	Normative References
33	12.2.	Informative References
33		Authors' Addresses
34		

1. Introduction

G.709 Optical Transport Network (OTN) [G.709-2012] includes new fixed and flexible ODU (Optical channel Data Unit) containers, two types of Tributary Slots (i.e., 1.25Gbps and 2.5Gbps), and supports various multiplexing relationships (e.g., ODU_j multiplexed into ODU_k (j<k)), two different tributary slots for ODU_k (K=1, 2, 3) and ODUflex service type. In order to present this information in routing, this document provides OTN technology specific encoding for use in GMPLS OSPF-TE as defined in [[RFC4203](#)].

For a short overview of OTN evolution and implications of OTN requirements on GMPLS routing please refer to [[OTN-FWK](#)]. The information model and an evaluation against the current solution are provided in [[OTN-INFO](#)].

Routing information for Optical Channel Layer (OCh) (i.e., wavelength) is beyond the scope of this document. Please refer to [[RFC6163](#)] and [[RFC6566](#)] for further information.

The reader is supposed to be familiar with OTN framework [[OTN-FWK](#)] and GMPLS evaluation against OTN [[OTN-INFO](#)].

1.1. Terminology

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. OSPF-TE Extensions

In terms of GMPLS based OTN networks, each OTU_k can be viewed as a component link, and each component link can carry one or more types of ODU_j (j<k).

Each TE Link State Advertisement (LSA) can carry a top-level link Type Length Value (TLV) with several nested sub-TLVs to describe different attributes of a TE link. Two top-level TLVs are defined in [[RFC3630](#)]. (1) The Router Address TLV (referred to as the Node TLV) and (2) the TE link TLV. One or more sub-TLVs can be nested into two top-level TLVs. The sub-TLV set for the two top-level TLVs are also defined in [[RFC3630](#)] and [[RFC4203](#)].

As discussed in [[OTN-FWK](#)] and [[OTN-INFO](#)], OSPF-TE must be extended to be able to advertise the termination and switching capabilities of each different ODU_j and ODU_k/OTU_k (Optical Transport Unit) and the

advertisement of related multiplexing capabilities. These

Ceccarelli, et al.
3]

Expires April 24, 2014


[Page

capabilities are carried in the Interface Switching Capability Descriptor (ISCD) Switching Capability-specific information field using formats defined in this document. As discussed in [SWCAP-UPDT], the use of a technology specific Switching Capability-specific information field necessitates the definition of a new Switching Capability value and associated new Switching Capability.

In the following, we will use ODUj to indicate a service type that is multiplexed into a higher order ODU, ODUk to indicate a higher order ODU including an ODUj and ODUk/OTUk to indicate the layer mapped into the OTUk. Moreover, ODUj(S) and ODUk(S) are used to indicate ODUj and ODUk supporting switching capability only, and the ODUj->ODUk format is used to indicate the ODUj into ODUk multiplexing capability.

This notation can be repeated as needed depending on the number of multiplexing levels. In the following, the term "multiplexing tree" is used to identify a multiplexing hierarchy where the root is always a server ODUk/OTUk and any other supported multiplexed container is represented with increasing granularity until reaching the leaf of the tree. The tree can be structured with more than one branch if the server ODUk/OTUk supports more than one hierarchy.

For example, if a multiplexing or example a multiplexing hierarchy like the following one is considered:



The ODU4 is the root of the muxing tree, ODU3 and ODU2 are containers directly multiplexed into the server and then ODU2, ODU0 are the leaves of the ODU3 branch, while ODUflex and ODU0 are the leaves of the ODU2 one. This means that it is possible to have the following multiplexing capabilities:

ODU2->ODU3->ODU4

ODU0->ODU3->ODU4

Ceccarelli, et al.
4]

Expires April 24, 2014

[Page

ODUflex->ODU2->ODU4
 ODU0->ODU2->ODU4

3. TE-Link Representation

G.709 ODUk/OTUk Links are represented as TE-Links in GMPLS Traffic Engineering Topology for supporting ODUj layer switching. These TE-Links can be modeled in multiple ways.

OTUk physical Link(s) can be modeled as a TE-Link(s). Figure 1 below provides an illustration of one hop OTUk TE-links.




Figure 1: OTUk TE-Links

It is possible to create TE-Links that span more than one hop by creating FAs between non-adjacent nodes (see Figure 2). As in the one hop case, multiple hop TE-links advertise ODU switching capacity.

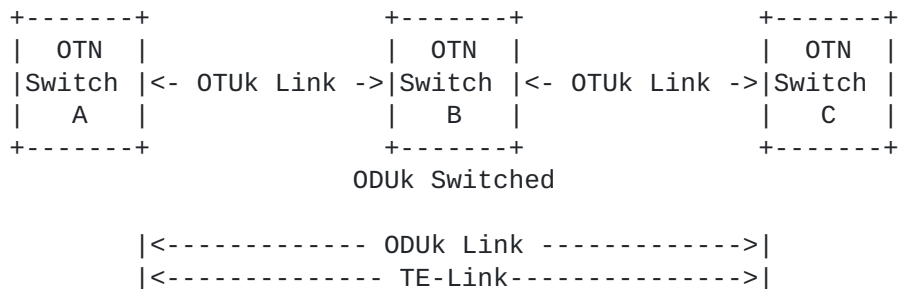


Figure 2: Multiple hop TE-Link

4. ISCD format extensions

The ISCD describes the switching capability of an interface and is defined in [\[RFC4203\]](#). This document defines a new Switching Capability value for OTN [G.709-2012] as follows:

Value	Type
-----	----
110 (TBA by IANA)	OTN-TDM capable (OTN-TDM)

When supporting the extensions defined in this document, for both fixed and flexible ODUs, the Switching Capability and Encoding values

MUST be used as follows:

- Switching Capability = OTN-TDM
- Encoding Type = G.709 ODUk (Digital Path) as defined in [\[RFC4328\]](#)

The same switching type and encoding values must be used for both fixed and flexible ODUs. When Switching Capability and Encoding fields are set to values as stated above, the Interface Switching Capability Descriptor MUST be interpreted as defined in [\[RFC4203\]](#).

Maximum LSP Bandwidth

The MAX LSP Bandwidth field is used according to [\[RFC4203\]](#): i.e., $0 \leq \text{MAX LSP Bandwidth} \leq \text{ODUK/OTUK}$, and intermediate values are those on the branch of OTN switching hierarchy supported by the interface. For example, in the OTU4 link it could be possible to have ODU4 as MAX LSP Bandwidth for some priorities, ODU3 for others, ODU2 for some

others, etc. The bandwidth unit is in bytes per second and the encoding MUST be in Institute of Electrical and Electronic Engineers (IEEE) floating point format. The discrete values for various ODUs is shown in the table below (please note that there are 1000 bits in a kbit according to normal practices in telecommunications).

ODU Type	ODU nominal bit rate	Value in Byte/ (floating p. val)
ODU0	1,244,160 kbit/s	0x4D1450C0
ODU1	239/238 x 2,488,320 kbit/s	0x4D94F048
ODU2	239/237 x 9,953,280 kbit/s	0x4E959129
ODU3	239/236 x 39,813,120 kbit/s	0x4F963367
ODU4	239/227 x 99,532,800 kbit/s	0x504331E3
ODU2e	239/237 x 10,312,500 kbit/s	0x4E9AF70A
ODUflex for CBR Client signals	239/238 x client signal bit rate	MAX LSP BANDWIDTH
ODUflex for GFP-F Mapped client signal	Configured bit rate	MAX LSP BANDWIDTH
ODU flex resizable	Configured bit rate	MAX LSP BANDWIDTH

A single ISCD MAY be used for the advertisement of unbundled or bundled links supporting homogeneous multiplexing hierarchies and the same TS (Tributary Slot) granularity. A different ISCD MUST be used for each different muxing hierarchy (muxing tree in the following examples) and different TS granularity supported within the TE Link.

In case a received LSA is not formatted accordingly to the

requirements indicated in this document, the problem SHOULD be logged and the wrongly formatted LSA, TLV or Sub-TLV MUST NOT be used for the path computation until a newer version correctly formatted is received.

4.1. Switching Capability Specific Information

The technology specific part of the OTN-TDM ISCD may include a variable number of sub-TLVs called Bandwidth sub-TLVs. Each sub-TLV is encoded with the sub-TLV header as defined in [[RFC3630](#)] [section 2.3.2](#). The muxing hierarchy tree MUST be encoded as an order independent list. Two types of Bandwidth sub-TLV are defined (TBA by IANA). Note that type values are defined in this document and not in [[RFC3630](#)].

- Type 1 - Unreserved Bandwidth for fixed containers

- Type 2 - Unreserved/MAX LSP Bandwidth for flexible containers

The Switching Capability-Specific Information (SCSI) MUST include one Type 1 sub-TLV for each fixed container and one Type 2 sub-TLV for each variable container. Each container type is identified by a Signal Type. Signal Type values are defined in [OTN-SIG].

With respect to ODUflex, three different signal types are allowed:

- ODUflex Constant Bit Rate (CBR), 21 - ODUflex Generic Framing Procedure-Frame mapped (GFP-F) resizable and 22 - ODUflex (GFP-F) non-resizable. Each MUST always be advertised in separate Type 2 sub-TLVs as each uses different adaptation functions [G.805]. In the case that both GFP-F resizable and non-resizable (i.e., 21 and 22) are supported, only Signal Type 21 SHALL be advertised as this type also implies support for type 22 adaptation.

4.1.1. Switching Capability Specific Information for fixed containers

The format of the Bandwidth sub-TLV for fixed containers is depicted in the following figure:

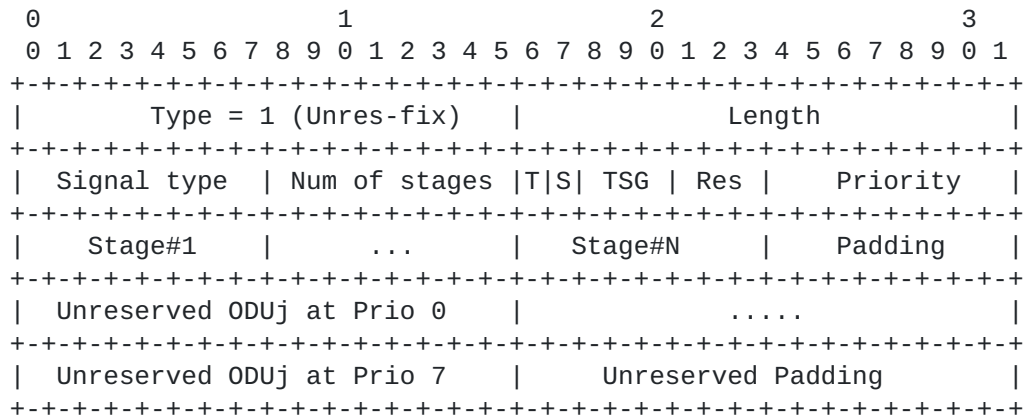


Figure 3: Bandwidth sub-TLV - Type 1 -

The values of the fields shown in figure 4 are explained in [section 4.1.3](#).

4.1.2. Switching Capability Specific Information for variable containers

The format of the Bandwidth sub-TLV for variable containers is depicted in the following figure:




Figure 4: Bandwidth sub-TLV - Type 2 -

The values of the fields shown in figure 4 are explained in [section 4.1.3](#).

4.1.3. Switching Capability Specific Information - Field values and explanation

The fields in the Bandwidth sub-TLV MUST be filled as follows:

- Signal Type (8 bits): Indicates the ODU type being advertised. Values are defined in [\[OTN-SIG\]](#).

- Number of stages (8 bits): This field indicates the number of multiplexing stages used to transport the indicated signal type. It MUST be set to the number of stages represented in the sub-

TLV.

- Flags (8 bits):

- T Flag (bit 17): Indicates whether the advertised bandwidth can be terminated. When the signal type can be terminated T MUST be set, while when the signal type cannot be terminated T MUST be cleared.

- S Flag (bit 18): Indicates whether the advertised bandwidth can be switched. When the signal type can be switched S MUST be set, while when the signal type cannot be switched S MUST

be

cleared.

The value 0 in both T and S bits MUST NOT be used.

- TS Granularity: Tributary Slot Granularity (3 bits): Used for the advertisement of the supported Tributary Slot granularity. The following values MUST be used:

- 0 - Ignored
- 1 - 1.25Gbps/2.5Gbps
- 2 - 2.5Gbps only
- 3 - 1.25Gbps only
- 4-7 - Reserved

A value of 1 MUST be used on interfaces which are configured to support the fall back procedures defined in [G.798-a2]. A value of 2 MUST be used on interfaces that only support 2.5Gbps time slots, such as [[RFC4328](#)] interfaces. A value of 3 MUST be used

on

interfaces that are configured to only support 1.25Gbps time slots. A value of 0 MUST be used for non-multiplexed signal

types

(i.e., a non-OTN client).

- Res (3 bits): reserved bits. MUST be set to 0 and ignored on receipt.

- Priority (8 bits): A bitmap used to indicate which priorities are being advertised. The bitmap is in ascending order, with the leftmost bit representing priority level 0 (i.e., the highest)

and

the rightmost bit representing priority level 7 (i.e., the lowest). A bit MUST be set (1) corresponding to each priority represented in the sub-TLV, and MUST NOT be set (0) when the corresponding priority is not represented. At least one priority level MUST be advertised that, unless overridden by local policy, SHALL be at priority level 0.

- Stage (8 bits): Each Stage field indicates a signal type in the multiplexing hierarchy used to transport the signal indicated in the Signal Type field. The number of Stage fields included in a sub-TLV MUST equal the value of the Number of Stages field. The Stage fields MUST be ordered to match the data plane in ascending order (from the lowest order ODU to the highest order ODU). The

values of the Stage field are the same as those defined for the Signal Type field. When the Number of stage field carries a 0, then the Stage and Padding fields MUST be omitted.

- Padding (variable): The Padding field is used to ensure the 32 bit alignment of stage fields. The length of the Padding field is always a multiple of 8 bits (1 byte). Its length can be calculated, in bytes, as: $4 - (\text{"value of Number of Stages field"} \% 4)$. The Padding field MUST be set to a zero (0) value on transmission and MUST be ignored on receipt.

- Unreserved ODUj (16 bits): This field indicates the Unreserved Bandwidth at a particular priority level. This field MUST be set to the number of ODUs at the indicated the Signal Type for a particular priority level. One field MUST be present for each bit set in the Priority field, and is ordered to match the Priority field. Fields MUST NOT be present for priority levels that are not indicated in the Priority field.

- Unreserved Padding (16 bits): The Padding field is used to ensure the 32 bit alignment of Unreserved ODUj fields. When present the Unreserved Padding field is 16 bits (2 byte) long. When the number of priorities is odd, the Unreserved Padding field MUST be included. When the number of priorities is even, the Unreserved Padding MUST be omitted.

- Unreserved Bandwidth (32 bits): This field indicates the Unreserved Bandwidth at a particular priority level. This field MUST be set to the bandwidth, in Bytes/sec in IEEE floating point format, available at the indicated Signal Type for a particular priority level. One field MUST be present for each bit set in the Priority field, and is ordered to match the Priority field. Fields MUST NOT be present for priority levels that are not indicated in the Priority field.

- Maximum LSP Bandwidth (32 bit): This field indicates the maximum bandwidth that can be allocated for a single LSP at a particular priority level. This field MUST be set to the maximum bandwidth, in Bytes/sec in IEEE floating point format, available to a single LSP at the indicated Signal Type for a particular priority level. One field MUST be present for each bit set in the Priority field, and is ordered to match the Priority field. Fields MUST NOT be present for priority levels that are not indicated in the Priority field. The advertisement of the MAX LSP Bandwidth MUST take into account HO OPUk bit rate tolerance and be calculated according to the following formula:

$$\text{Max LSP BW} = (\# \text{ available TSS}) * (\text{ODTUK.ts nominal bit rate}) * (1 - \text{HO OPUk bit rate tolerance})$$

5. Examples

The examples in the following pages are not normative and are not intended to imply or mandate any specific implementation.

5.1. MAX LSP Bandwidth fields in the ISCD

This example shows how the MAX LSP Bandwidth fields of the ISCD are filled accordingly to the evolving of the TE-link bandwidth occupancy. In the example an OTU4 link is considered, with supported priorities 0,2,4,7 and muxing hierarchy ODU1->ODU2->ODU3->ODU4.

At time T0, with the link completely free, the advertisement would be:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| SwCap=OTN_TDM | Encoding = 12 |   Reserved (all zeros)   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 0 = 100Gbps       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 1 = 0             |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 2 = 100Gbps       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 3 = 0             |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 4 = 100Gbps       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 5 = 0             |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 6 = 0             |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           MAX LSP Bandwidth at priority 7 = 100Gbps       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Switching Capability Specific Information       |
|           (variable length)                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 5: Example 1 - MAX LSP Bandwidth fields in the ISCD at T0

At time T1, an ODU3 at priority 2 is set-up, so for priority 0 the MAX LSP Bandwidth is still equal to the ODU4 bandwidth, while for priorities from 2 to 7 (excluding the non-supported ones) the MAX

LSP

Bandwidth is equal to ODU3, as no more ODU4s are available and the next supported ODUj in the hierarchy is ODU3. The advertisement is updated as follows:

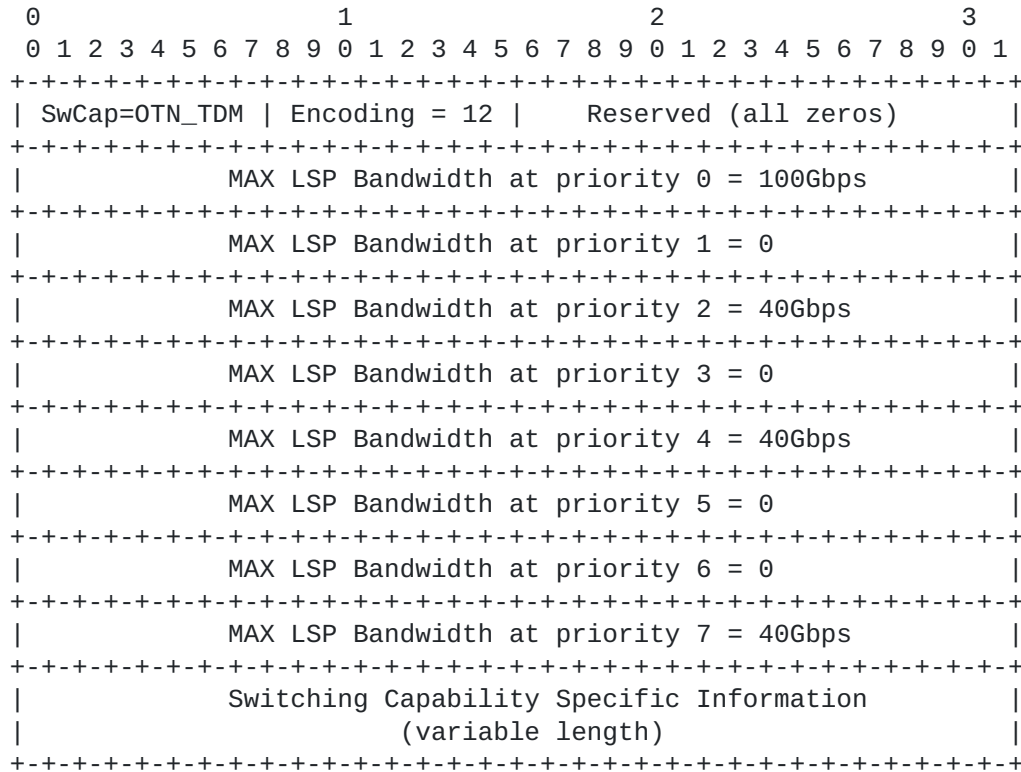


Figure 6: Example 1 - MAX LSP Bandwidth fields in the ISCD at T1

At time T2, an ODU2 at priority 4 is set-up. The first ODU3 is no longer available since T1, as it was kept by the ODU3 LSP, while the second is no more available and just 3 ODU2 are left in it. ODU2 is now the MAX LSP Bandwidth for priorities higher than 4. The advertisement is updated as follows:

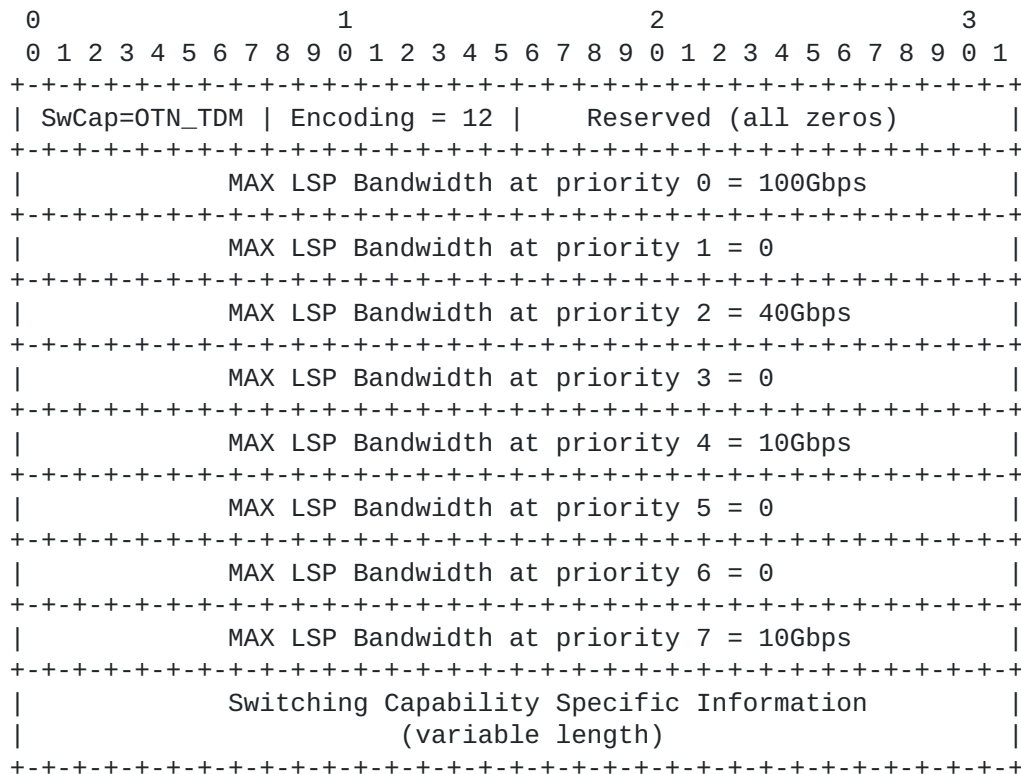


Figure 7: Example 1 - MAX LSP Bandwidth fields in the ISCD at T2

5.2. Example of T,S and TS granularity utilization

In this example, an interface with Tributary Slot Type 1.25Gbps and fallback procedure enabled is considered (TS granularity=1). It supports the simple ODU1->ODU2->ODU3 hierarchy and priorities 0 and 3. Suppose that in this interface the ODU3 signal type can be both switched or terminated, the ODU2 can only be terminated, and the ODU1 switched only. Please note that since the ODU1 is not being advertised to support ODU0, the value of is "ignored" (TS granularity=0). For the advertisement of the capabilities of such interface, a single ISCD is used and its format is as follows:

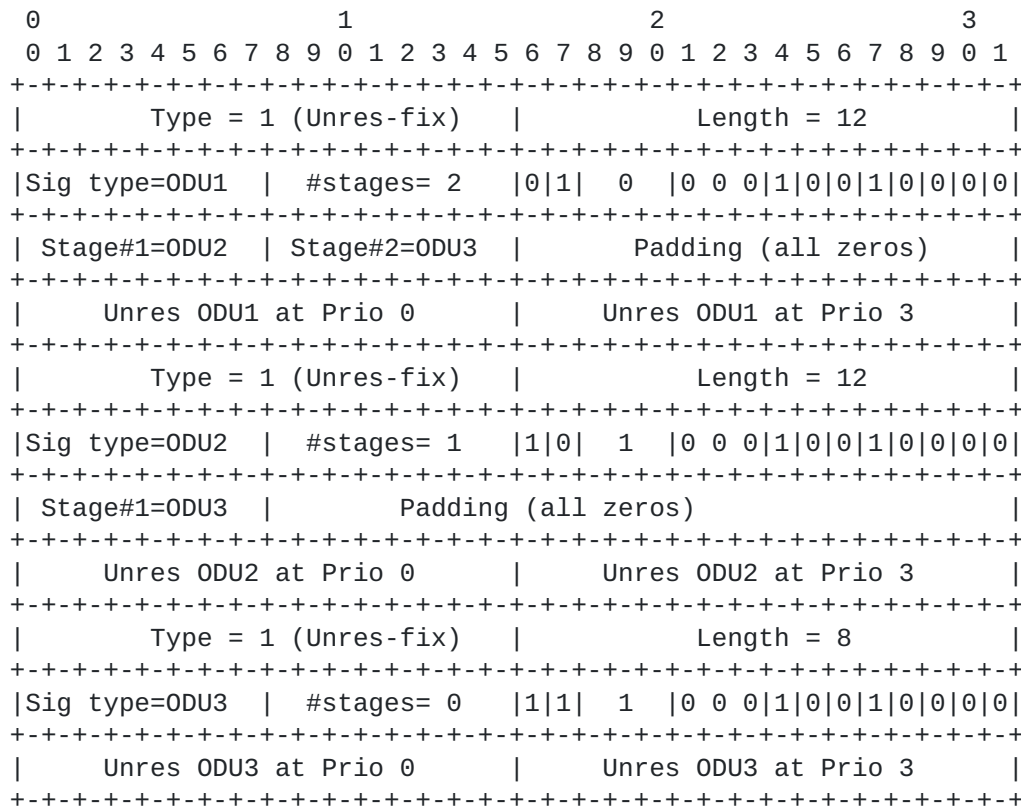


Figure 8: Example 2 - TS granularity, T and S utilization

5.2.1. Example of different TS Granularities

In this example, two interfaces with homogeneous hierarchies but different Tributary Slot Types are considered. The first one supports a [RFC4328](#) interface (TS granularity=2) while the second one supports G.709-2012 interface with fallback procedure disabled (TS granularity=3). Both of them support ODU1->ODU2->ODU3 hierarchy and priorities 0 and 3. T and S bits values are not relevant to this example. For the advertisement of the capabilities of such interfaces, two different ISCDs are used and the format of their SCSIs is as follows:

SCSI of ISCD 1 - TS granularity=2

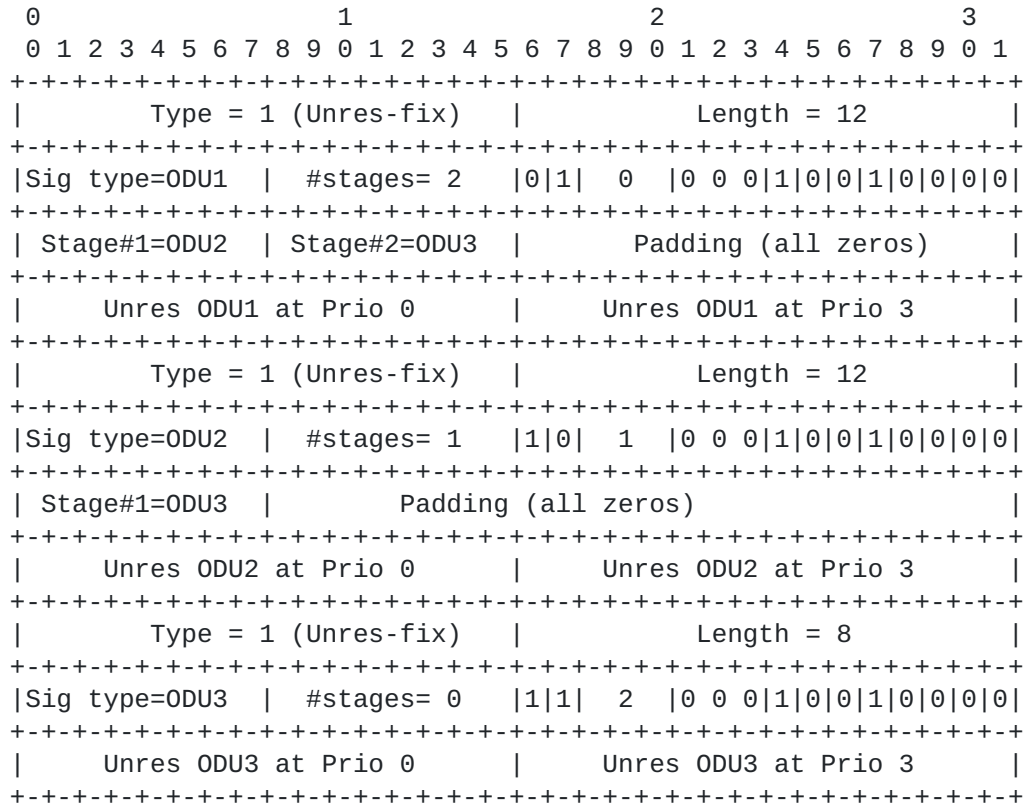


Figure 9: Example 2.1 - Different TS Granularities utilization - ISCD




Figure 10: Example 2.1 - Different TS Granularities utilization - ISCD 2

A particular case in which hierarchies with the same muxing tree but with different exported TS granularity MUST be considered as non-homogenous hierarchies. This is the case in which an H-LPS and the client LSP are terminated on the same egress node. What can happen is that a loose Explicit Route Object (ERO) is used at the hop where the signaled LSP is nested into the Hierarchical-LSP (H-LSP) (penultimate hop of the LSP).

In the following figure, node C receives from A a loose ERO towards node E and must choose between the ODU2 H-LSP on if1 or the one on if2. In this case, the H-LSP on if1 exports a TS=1.25Gbps, and if2

a

TS=2.5Gbps, the service LSP being signaled needs a 1.25Gbps

tributary

slot, only the H-LSP on if1 can be used to reach node E. For further details, please see [section 4.1](#) of the [\[OTN-INFO\]](#).

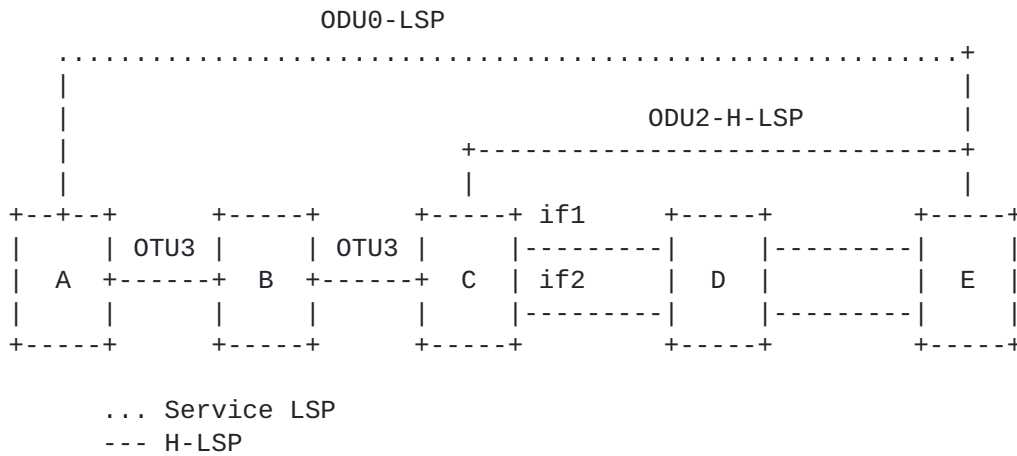


Figure 11: Example - Service LSP and H-LSP terminating on the same node

5.3. Example of ODUFlex advertisement

In this example, the advertisement of an ODUFlex->ODU3 hierarchy is shown. In case of ODUFlex advertisement, the MAX LSP Bandwidth needs to be advertised and, in some cases, information about the Unreserved bandwidth could also be useful. The amount of Unreserved bandwidth does not give a clear indication of how many ODUFlex LSP can be set up either at the MAX LSP Bandwidth or at different rates, as it gives no information about the spatial allocation of the free TSS.

An indication of the amount of Unreserved bandwidth could be useful during the path computation process, as shown in the following example. Supposing there are two TE-links (A and B) with MAX LSP Bandwidth equal to 10 Gbps each. In the case where 50Gbps of Unreserved Bandwidth are available on Link A, 10Gbps on Link B, and 3 ODUFlex LSPs of 10 GBps each have to be restored, for sure only one can be restored along Link B and it is probable, but not certain, that two of them can be restored along Link A. T, S and TS granularity fields are not relevant to this example.

In the case of ODUFlex advertisement, the Type 2 Bandwidth sub-TLV is used.



```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type = 2 (Unres/MAX-var)  |          Length = 72          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|S. type=ODUflex| #stages= 1  |X|X|X X X|0 0 0| Priority(8)  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Stage#1=ODU3 |          Padding (all zeros)          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 0          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 2          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 3          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 4          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 5          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 6          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 7          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 0          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 2          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 3          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 4          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 5          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 6          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 7          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Figure 12: Example 3 - ODUflex advertisement

5.4. Example of single stage muxing

Supposing there is 1 OTU4 component link supporting single stage muxing of ODU1, ODU2, ODU3 and ODUflex, the supported hierarchy can be summarized in a tree as in the following figure. For sake of simplicity, we also assume that only priorities 0 and 3 are supported. T, S and TS granularity fields are not relevant to this example.




and the related SCSIs as follows:

Figure 13: Example 4 - Single stage muxing

5.5. Example of multi stage muxing - Unbundled link

Supposing there is 1 OTU4 component link with muxing capabilities as shown in the following figure:



and supported priorities 0 and 3, the advertisement is composed by the following Bandwidth sub-TLVs (T and S fields are not relevant to this example):


```


+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type = 2 (Unres/MAX-var)  |          Length = 24          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|S.type=ODUflex | #stages= 2  |X|X|  0  |0 0 0|1|0|0|1|0|0|0|0|0|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Stage#1=ODU2  | Stage#2=ODU4 |  Padding (all zeros)      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 0 =100Gbps      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Unreserved Bandwidth at priority 3 =100Gbps      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 0 =10Gbps          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          MAX LSP Bandwidth at priority 3 =10Gbps          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 14: Example 5 - Multi stage muxing - Unbundled link

5.6. Example of multi stage muxing - Bundled links

In this example, 2 OTU4 component links with the same supported TS granularity and homogeneous muxing hierarchies are considered. The following muxing capabilities trees are supported:



Considering only supported priorities 0 and 3, the advertisement is as follows (T, S and TS granularity fields are not relevant to this example):

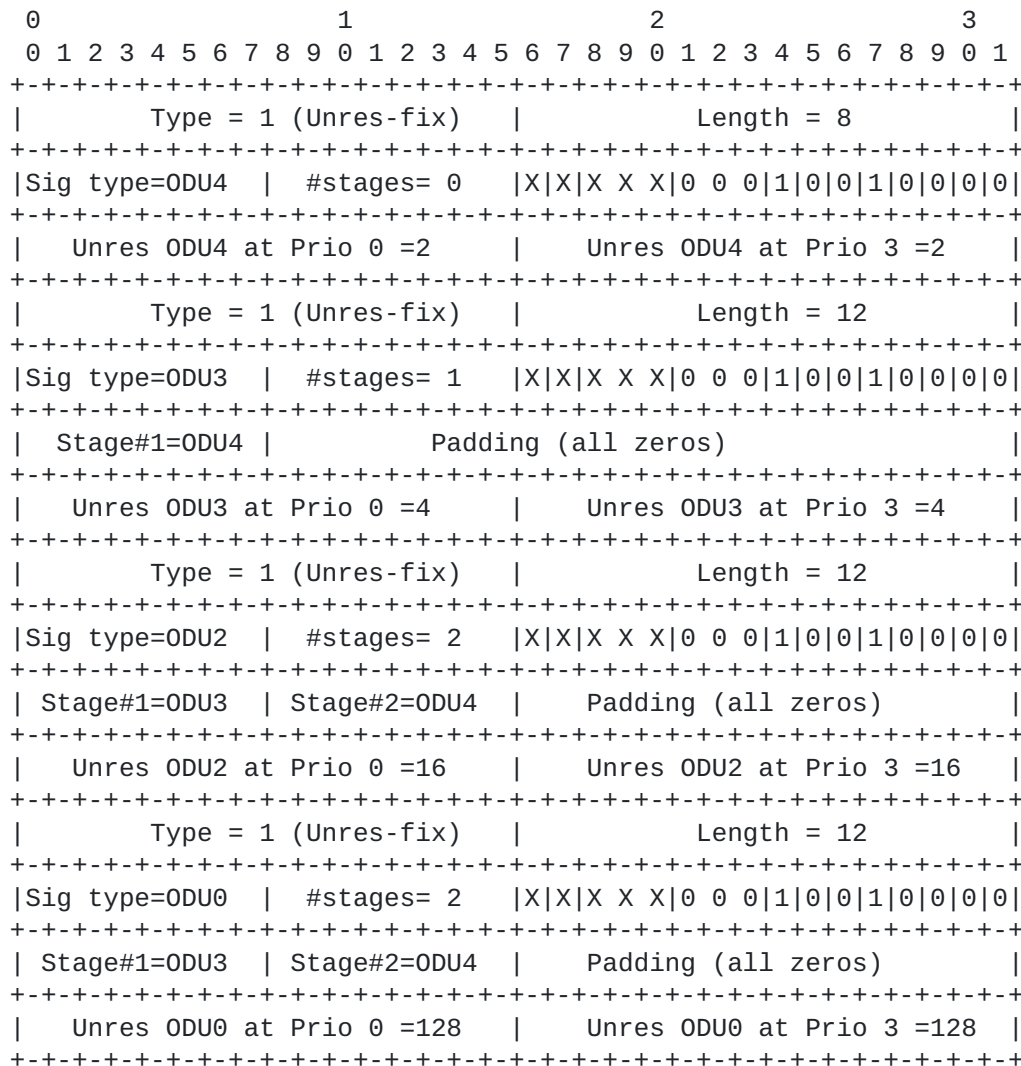



Figure 15: Example 6 - Multi stage muxing - Bundled links

5.7. Example of component links with non-homogeneous hierarchies

In this example, 2 OTU4 component links with the same supported TS granularity and non-homogeneous muxing hierarchies are considered. The following muxing capabilities trees are supported:



Considering only supported priorities 0 and 3, the advertisement uses two different ISCDs, one for each hierarchy (T, S and TS granularity fields are not relevant to this example). In the following figure, the SCSI of each ISCD is shown:

SCSI of ISCD 1 - Component Link#1

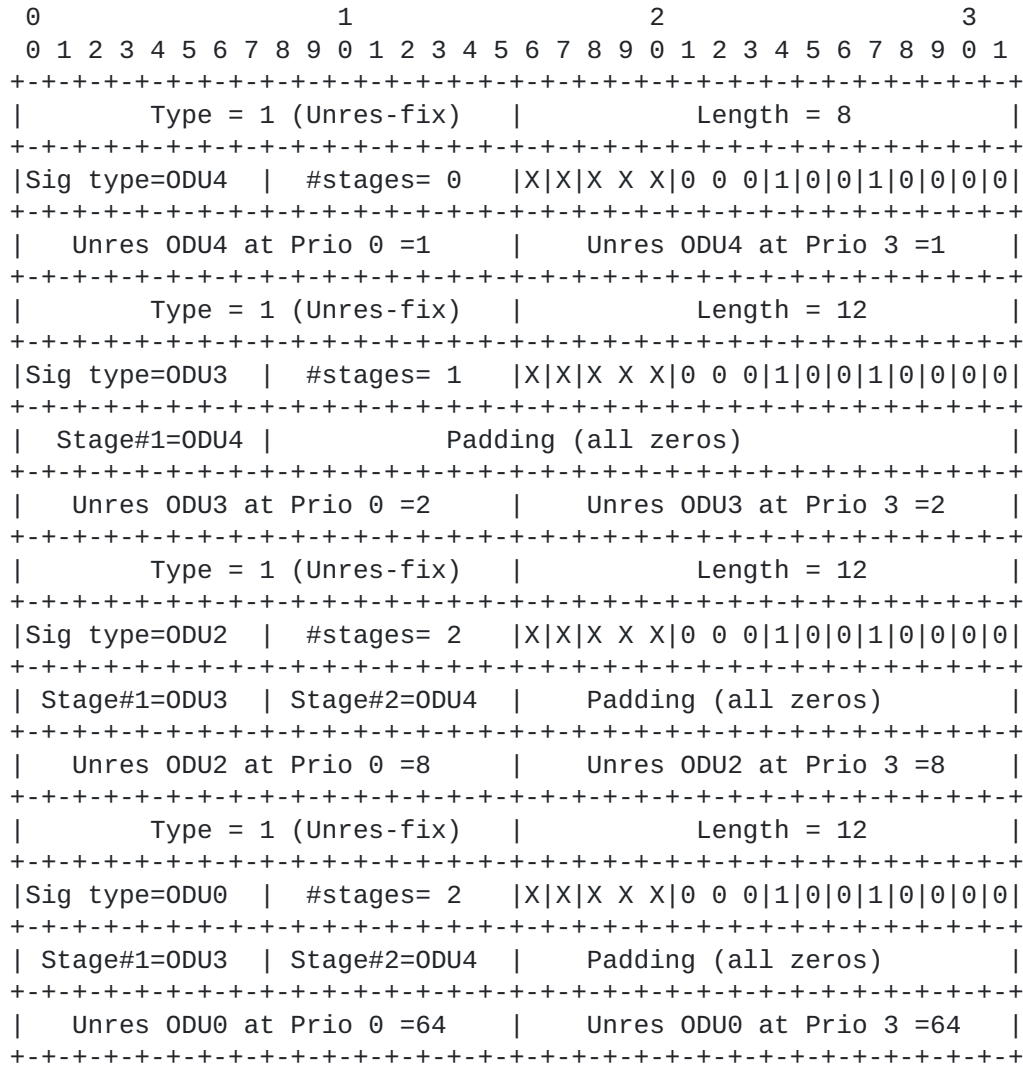


Figure 16: Example 7 - Multi stage muxing - Non-homogeneous hierarchies - ISCD 1

SCSI of ISCD 2 - Component Link#2

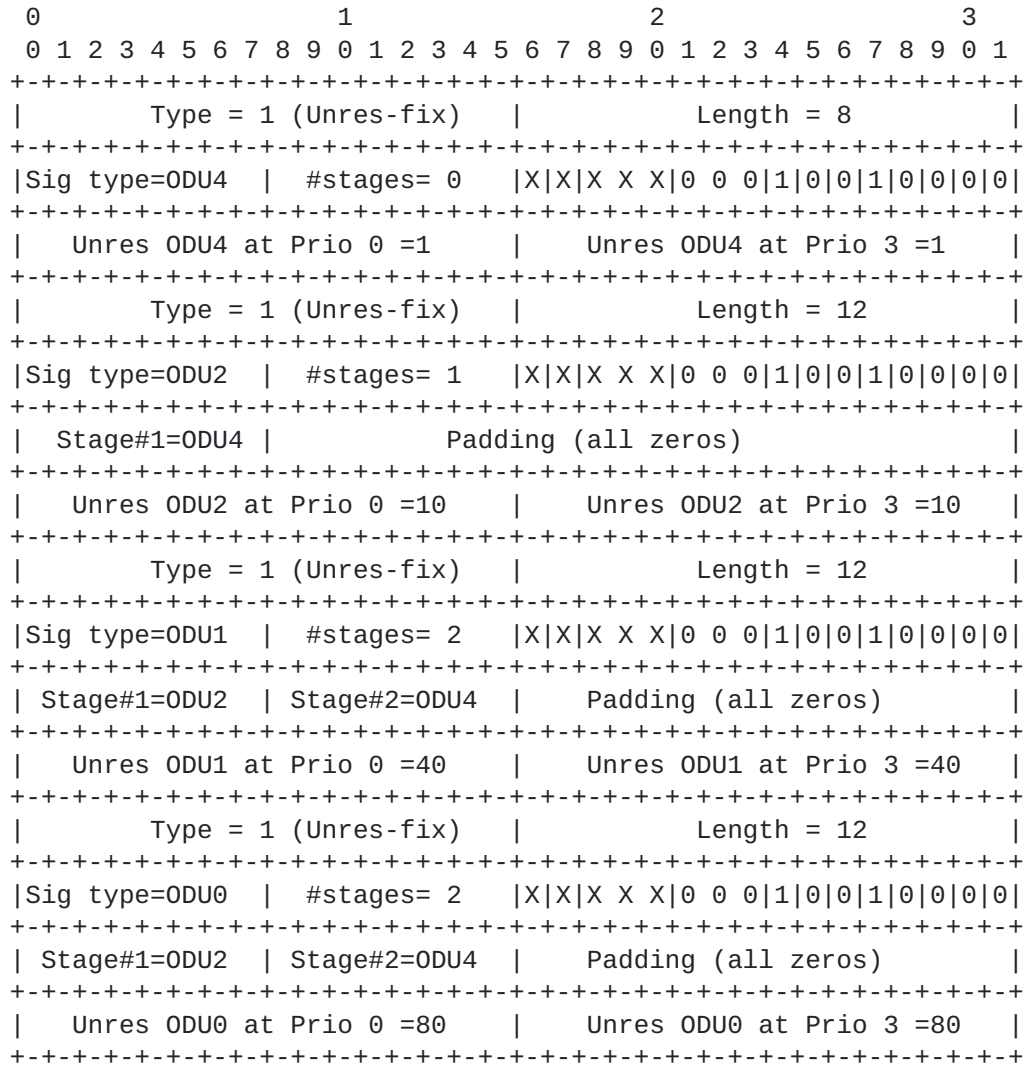


Figure 17: Example 7 - Multi stage muxing - Non-homogeneous hierarchies - ISCD 2

6. OSPFv2 scalability

This document does not introduce OSPF scalability issues with respect

to existing GMPLS encoding and does not require any modification to flooding frequency. Moreover the design of the encoding has been carried out taking into account bandwidth optimization, and in particular:

- Only unreserved and MAX LSP Bandwidth related to supported priorities is advertised
- With respect of fixed containers, only the number of available containers is advertised instead of available bandwidth so to use only 16 bits per container instead of 32 (as per former GMPLS encoding)

In order to further reduce the amount of data advertised it is RECOMMENDED to bundle component links with homogeneous hierarchies as described in [[RFC4201](#)] and illustrated in [Section 5.6](#).

7. Compatibility

All implementations of this document MAY also support advertisement as defined in [[RFC4328](#)]. When nodes support both advertisement methods, implementations MUST support the configuration of which advertisement method is followed. The choice of which is used is based on policy and beyond the scope of this document. This enables nodes following each method to identify similar supporting nodes and compute paths using only the appropriate nodes.

8. Security Considerations

This document extends [[RFC4203](#)]. As with [[RFC4203](#)], it specifies the contents of Opaque LSAs in OSPFv2. As Opaque LSAs are not used for 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 (optical and/or SONET-SDH)

network. [[RFC3630](#)] suggests mechanisms such as [[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.

For security threats, defensive techniques, monitoring/detection/reporting of security attacks and requirements please refer to [[RFC5920](#)].

9. IANA Considerations

9.1. Switching types

Upon approval of this document, IANA will make the assignment in the "Switching Types" section of the "GMPLS Signaling Parameters" registry located at

<http://www.iana.org/assignments/gmpls-sig-parameters>:

Value	Name	Reference
110 (*)	OTN-TDM capable (OTN-TDM)	[This.I-D]

(*) Suggested value

Same type of modification needs to applied to the IANA-GMPLS-TC-MIB at <https://www.iana.org/assignments/ianagmplstc-mib/ianagmplstc-mib>

9.2. New sub-TLVs

This document defines 2 new sub-TLVs that are carried in Interface Switching Capability Descriptors [[RFC4203](#)] with Signal Type OTN-TDM. Each sub-TLV includes a 16-bit type identifier (the T-field). The same T-field values are applicable to the new sub-TLV.

Upon approval of this document, IANA will create and maintain a new sub-registry, the "Types for sub-TLVs of OTN-TDM SCSI (Switch Capability-Specific Information)" registry under the "Open Shortest Path First (OSPF) Traffic Engineering TLVs" registry, see <http://www.iana.org/assignments/ospf-traffic-eng-tlvs/ospf-traffic-eng-tlvs.xml>, with the sub-TLV types as follows:

This document defines new sub-TLV types as follows:

Value	Sub-TLV	Reference
0	Reserved	[This.I-D]
1	Unreserved Bandwidth for fixed containers	[This.I-D]
2	Unreserved/MAX Bandwidth for flexible containers	[This.I-D]
3-65535	Unassigned	

Types are to be assigned via Standards Action as defined in [[RFC5226](#)].

10. Contributors

Diego Caviglia, Ericsson

Internet-Draft
2013

OSPF-TE extensions for OTN support

October

Via E.Melen, 77 - Genova - Italy

Email: diego.caviglia@ericsson.com

Dan Li, Huawei Technologies

Bantian, Longgang District - Shenzhen 518129 P.R.China

Email: danli@huawei.com

Pietro Vittorio Grandi, Alcatel-Lucent

Via Trento, 30 - Vimercate - Italy

Email: pietro_vittorio.grandi@alcatel-lucent.com

Khuzema Pithewan, Infinera Corporation

140 Caspian CT., Sunnyvale - CA - USA

Email: kpithewan@infinera.com

Xiaobing Zi, Huawei Technologies

Email: zixiaobing@huawei.com

Francesco Fondelli, Ericsson

Email: francesco.fondelli@ericsson.com

Marco Corsi

EMail: corsi.marco@gmail.com

Internet-Draft
2013

OSPF-TE extensions for OTN support

October

Eve Varma, Alcatel-Lucent

EMail: eve.varma@alcatel-lucent.com

Jonathan Sadler, Tellabs

EMail: jonathan.sadler@tellabs.com

Lyndon Ong, Ciena

EMail: lyong@ciena.com

Ashok Kunjidhpatham

akunjidhpatham@infinera.com

Snigdho Bardalai

sbardalai@infinera.com

Steve Balls

Steve.Balls@metaswitch.com

Jonathan Hardwick

Jonathan.Hardwick@metaswitch.com

Xihua Fu

fu.xihua@zte.com.cn

Cyril Margaria

cyril.margaria@nsn.com

Malcolm Betts

Malcolm.betts@zte.com.cn

11. Acknowledgements

The authors would like to thank Fred Gruman and Lou Berger for the precious comments and suggestions.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", [RFC 4201](#), October 2005.
- [RFC4203] Kompella, K. and Y. Rekhter, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4203](#), October 2005.
- [RFC4328] Papadimitriou, D., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control", [RFC 4328](#), January 2006.

12.2. Informative References

- [OTN-FWK] F.Zhang, D.Li, H.Li, S.Belotti, D.Ceccarelli, "Framework for GMPLS and PCE Control of G.709 Optical Transport networks, work in progress [draft-ietf-ccamp-gmpls-g709-framework-13](#)", June 2013.
- [OTN-INFO] S.Belotti, P.Grandi, D.Ceccarelli, D.Caviglia, F.Zhang,

- D.Li, "Information model for G.709 Optical Transport Networks (OTN), work in progress
[draft-ietf-ccamp-otn-g709-info-model-09](#)", June 2013.
- [OTN-SIG] F.Zhang, G.Zhang, S.Belotti, D.Ceccarelli, K.Pithewan, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for the evolving G.709 Optical Transport Networks Control, work in progress
[draft-ietf-ccamp-gmpls-signaling-g709v3-11](#)", June 2013.
- [RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", [RFC 2154](#), June 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), July 2010.
- [RFC6163] Lee, Y., Bernstein, G., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSOs)", [RFC 6163](#), April 2011.
- [RFC6566] Lee, Y., Bernstein, G., Li, D., and G. Martinelli, "A Framework for the Control of Wavelength Switched Optical Networks (WSOs) with Impairments", [RFC 6566](#), March 2012.
- [SWCAP-UPDT]
F.Zhang, D.Li, H.Li, S.Belotti, D.Ceccarelli, "Framework for GMPLS and PCE Control of G.709 Optical Transport networks, work in progress
[draft-ietf-ccamp-gmpls-g709-framework-13](#)", June 2013.

Authors' Addresses

Daniele Ceccarelli (editor)
Ericsson
Via E.Melen 77
Genova - Erzelli
Italy

Email: daniele.ceccarelli@ericsson.com

Internet-Draft
2013

OSPF-TE extensions for OTN support

October

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Shenzhen 518129 P.R.China Bantian, Longgang District
Phone: +86-755-28972912

Email: zhangfatai@huawei.com

Sergio Belotti
Alcatel-Lucent
Via Trento, 30
Vimercate
Italy

Email: sergio.belotti@alcatel-lucent.com

Rajan Rao
Infinera Corporation
140, Caspian CT.
Sunnyvale, CA-94089
USA

Email: rrao@infinera.com

John E Drake
Juniper

Email: jdrake@juniper.net

