

Internet Engineering Task Force (IETF)
Request for Comments: 7138
Category: Standards Track
ISSN: 2070-1721

D. Ceccarelli, Ed.
Ericsson
F. Zhang
Huawei Technologies
S. Belotti
Alcatel-Lucent
R. Rao
Infinera Corporation
J. Drake
Juniper
March 2014

Traffic Engineering Extensions to OSPF
for GMPLS Control of Evolving G.709 Optical Transport Networks

Abstract

This document describes Open Shortest Path First - Traffic Engineering (OSPF-TE) routing protocol extensions to support GMPLS control of Optical Transport Networks (OTNs) specified in ITU-T Recommendation G.709 as published in 2012. It extends mechanisms defined in RFC 4203.

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/rfc7138>.

Copyright Notice

Copyright (c) 2014 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	4
1.1. Terminology	4
2. OSPF-TE Extensions	4
3. TE-Link Representation	6
4. ISCD Format Extensions	6
4.1. Switching Capability Specific Information	8
4.1.1. Switching Capability Specific Information for Fixed Containers	9
4.1.2. Switching Capability Specific Information for Variable Containers	10
4.1.3. Switching Capability Specific Information -- Field Values and Explanation	10
5. Examples	13
5.1. MAX LSP Bandwidth Fields in the ISCD	13
5.2. Example of T, S, and TS Granularity Utilization	17
5.2.1. Example of Different TS Granularities	18
5.3. Example of ODUFlex Advertisement	20
5.4. Example of Single-Stage Muxing	22
5.5. Example of Multi-Stage Muxing -- Unbundled Link	23
5.6. Example of Multi-Stage Muxing -- Bundled Links	25
5.7. Example of Component Links with Non-Homogeneous Hierarchies	27
6. OSPFv2 Scalability	29
7. Compatibility	30
8. Security Considerations	30
9. IANA Considerations	31
9.1. Switching Types	31
9.2. New Sub-TLVs	31
10. Contributors	32
11. Acknowledgements	33
12. References	33
12.1. Normative References	33
12.2. Informative References	34

1. Introduction

G.709 ("Interfaces for the Optical Transport Network (OTN)") [G.709-2012] includes new fixed and flexible ODU (Optical channel Data Unit) containers, includes two types of tributary slots (i.e., 1.25 Gbps and 2.5 Gbps), and supports various multiplexing relationships (e.g., ODUj multiplexed into ODUk (j<k)), two different tributary slots for ODUk (K=1, 2, 3), and the ODUFlex service type. In order to advertise this information in routing, this document provides encoding specific to OTN technology 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 [RFC7062]. The information model and an evaluation against the current solution are provided in [RFC7096]. The reader is supposed to be familiar with both of these documents.

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

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 Optical channel Transport Unit-k (OTUk) can be viewed as a component link, and each component link can carry one or more types of ODUj (j<k).

Each TE-Link State Advertisement (LSA) can carry a top-level link 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 the two top-level TLVs. The sub-TLV set for the two top-level TLVs are also defined in [RFC3630] and [RFC4203].


As discussed in [RFC7062] and [RFC7096], OSPF-TE must be extended to be able to advertise the termination and Switching Capabilities of each different ODUj and ODUk/OTUk (Optical Transport Unit) and the advertisement of related multiplexing capabilities. These capabilities are carried in the Switching Capability specific information field of the Interface Switching Capability Descriptor

(ISCD) using formats defined in this document. As discussed in [RFC7062], 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 ODU_j to indicate a service type that is multiplexed into a higher-order (HO) ODU, ODU_k to indicate a higher-order ODU including an ODU_j, and ODU_k/OTU_k to indicate the layer mapped into the OTU_k. Moreover, ODU_j(S) and ODU_k(S) are used to indicate the ODU_j and ODU_k supporting Switching Capability only, and the ODU_j->ODU_k format is used to indicate the ODU_j-into-ODU_k 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 ODU_k/OTU_k 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 ODU_k/OTU_k supports more than one hierarchy.

For example, if 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 ODU₂ and 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
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 forwarding adjacencies (FAs) between non-adjacent nodes (see Figure 2). As in the one-hop case, multiple-hop TE-Links advertise the ODU Switching Capability.

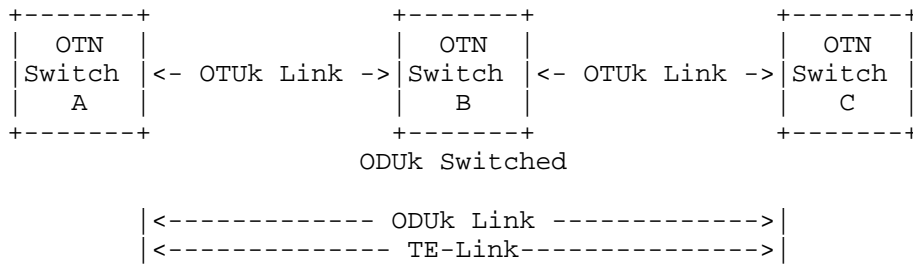


Figure 2: Multiple-Hop TE-Link

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	OTN-TDM capable

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:

- o Switching Capability = OTN-TDM
- o 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].

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 the 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/second and the encoding MUST be in IEEE floating point format. The discrete values for various ODUs are shown in the table below (please note that there are 1000 bits in a kilobit according to normal practices in telecommunications).

ODU Type	ODU nominal bit rate	Value in Byte/Sec (floating p. val)
ODU0	1,244,160 kbps	0x4D1450C0
ODU1	239/238 x 2,488,320 kbps	0x4D94F048
ODU2	239/237 x 9,953,280 kbps	0x4E959129
ODU3	239/236 x 39,813,120 kbps	0x4F963367
ODU4	239/227 x 99,532,800 kbps	0x504331E3
ODU2e	239/237 x 10,312,500 kbps	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
ODUflex 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.

When a received LSA includes a sub-TLV not formatted accordingly to the precise specifications in this document, the problem SHOULD be logged and the wrongly formatted sub-TLV MUST NOT be used for path computation.

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-TLVs are defined (TBA by IANA). Note that type values are defined in this document and not in [RFC3630].

- o Type 1 - Unreserved Bandwidth for fixed containers
- o 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 [RFC7139].

With respect to ODUflex, three different Signal Types are allowed:

- o 20 - ODUflex(CBR) (i.e., 1.25*N Gbps)
- o 21 - ODUflex(GFP-F), resizable (i.e., 1.25*N Gbps)
- o 22 - ODUflex(GFP-F), non-resizable (i.e., 1.25*N Gbps)

where CBR stands for Constant Bit Rate, and GFP-F stands for Generic Framing Procedure - Framed.

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.

4.1.1.1. Switching Capability Specific Information for Fixed Containers

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

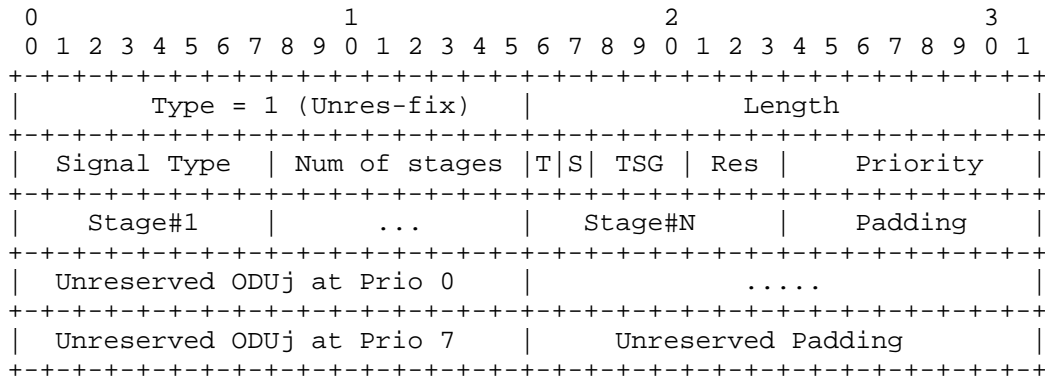


Figure 3: Bandwidth Sub-TLV -- Type 1

The values of the fields shown in Figure 3 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:

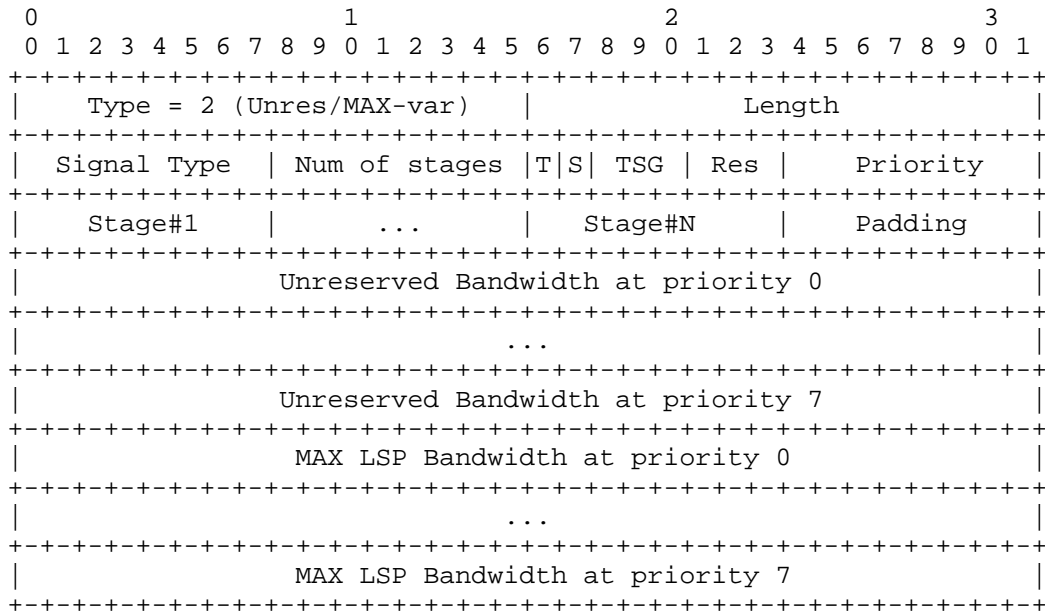


Figure 4: Bandwidth Sub-TLV -- Type 2

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:

- o Signal Type (8 bits): Indicates the ODU type being advertised. Values are defined in [RFC7139].
- o Num 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.

- o 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; when the Signal Type cannot be switched, S MUST be cleared.
 - * The value 0 in both the T bit and S bit MUST NOT be used.
- o TSG (3 bits): Tributary Slot Granularity. Used for the advertisement of the supported tributary slot granularity. The following values MUST be used:
 - * 0 - Ignored
 - * 1 - 1.25 Gbps / 2.5 Gbps
 - * 2 - 2.5 Gbps only
 - * 3 - 1.25 Gbps only
 - * 4-7 - Reserved

A value of 1 MUST be used on interfaces that are configured to support the fallback procedures defined in [G.798]. A value of 2 MUST be used on interfaces that only support 2.5 Gbps tributary slots, such as [RFC4328] interfaces. A value of 3 MUST be used on interfaces that are configured to only support 1.25 Gbps tributary slots. A value of 0 MUST be used for non-multiplexed Signal Types (i.e., a non-OTN client).

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

- o 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 Num 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 Num of stages field carries a 0, then the Stage and Padding fields MUST be omitted.
 - * Example: For the ODU1->ODU2->OD3 hierarchy, the Signal Type field is set to ODU1 and two Stage fields are present, the first indicating ODU2 and the second ODU3 (server layer).
- o 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: $4 - (\text{"value of Num of stages field"} \% 4)$. The Padding field MUST be set to a zero (0) value on transmission and MUST be ignored on receipt.
- o 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 the fields are ordered to match the Priority field. Fields MUST NOT be present for priority levels that are not indicated in the Priority field.
- o Unreserved Padding (16 bits): The Padding field is used to ensure the 32-bit alignment of the Unreserved ODUj fields. When present, the Unreserved Padding field is 16 bits (2 bytes) 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.
- o 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/second 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 the fields are ordered to match the Priority field. Fields MUST NOT be present for priority levels that are not indicated in the Priority field.

- o Maximum LSP Bandwidth (32 bits): 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/second 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 the fields are 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:

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

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 according to the evolving of the TE-Link bandwidth occupancy. In this 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:

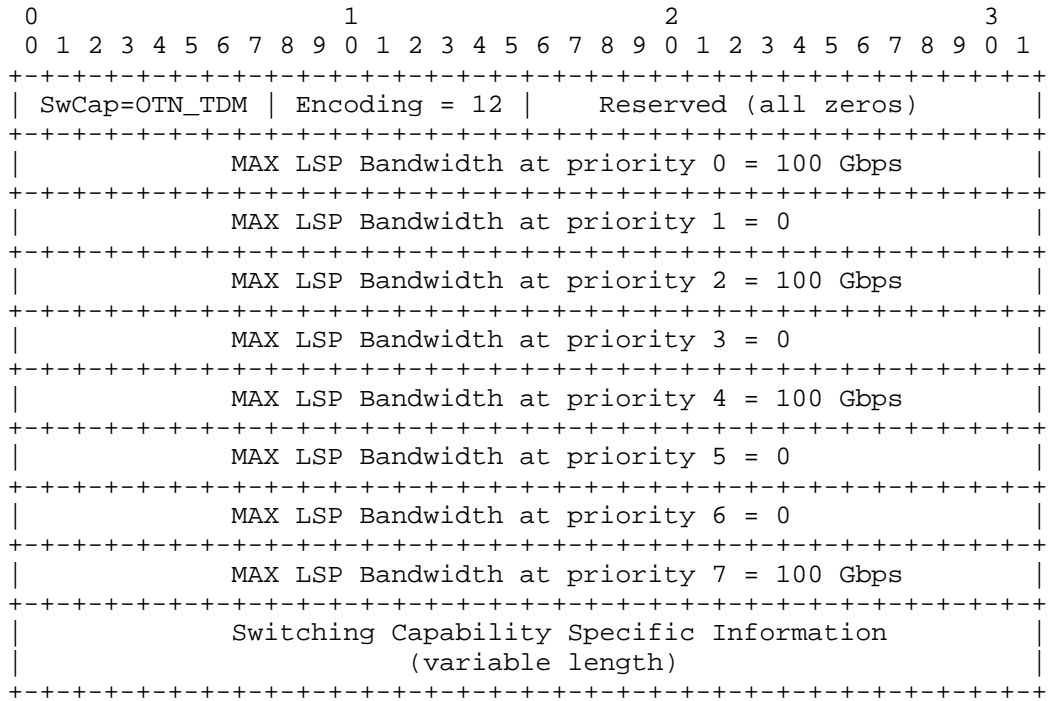


Figure 5: MAX LSP Bandwidth Fields in the ISCD at T0

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:

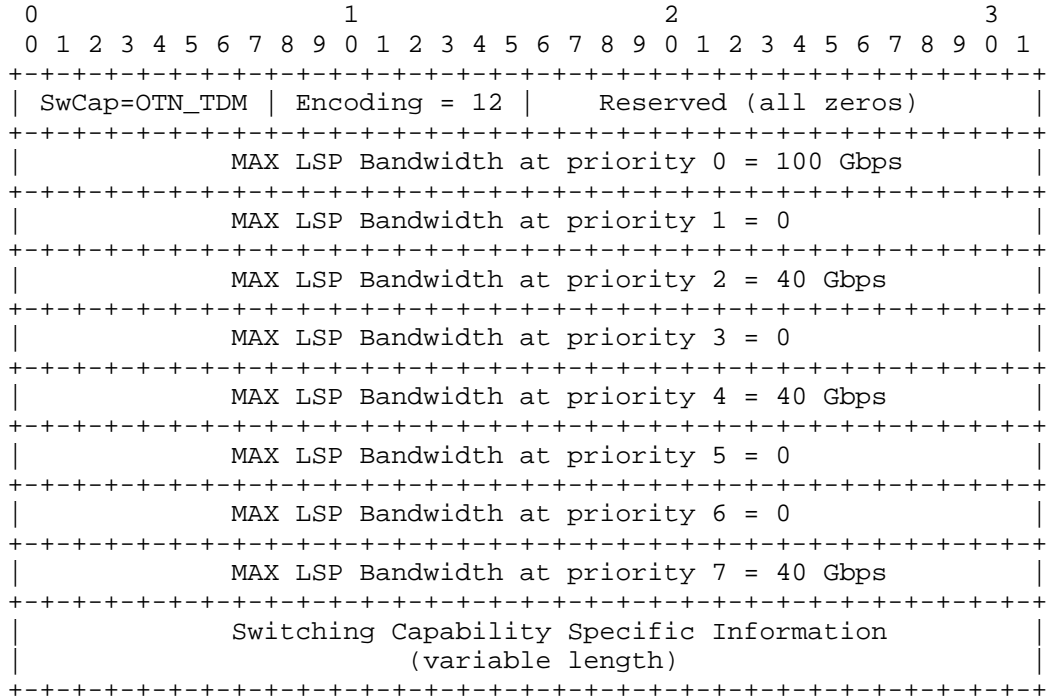


Figure 6: MAX LSP Bandwidth Fields in the ISCD at T1

At time T2, an ODU2 at priority 4 is set up. The first ODU3 has not been available since T1 as it was kept by the ODU3 LSP, while the second is no longer available and just 3 ODU2s are left in it. ODU2 is now the MAX LSP Bandwidth for priorities higher than 4. The advertisement is updated as follows:

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 = 100 Gbps																																															
MAX LSP Bandwidth at priority 1 = 0																																															
MAX LSP Bandwidth at priority 2 = 40 Gbps																																															
MAX LSP Bandwidth at priority 3 = 0																																															
MAX LSP Bandwidth at priority 4 = 10 Gbps																																															
MAX LSP Bandwidth at priority 5 = 0																																															
MAX LSP Bandwidth at priority 6 = 0																																															
MAX LSP Bandwidth at priority 7 = 10 Gbps																																															
Switching Capability Specific Information																																															
(variable length)																																															

Figure 7: 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 1.25 Gbps 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 can only be switched. Please note that since the ODU1 is not being advertised to support ODU0, the value of its TSG field is "ignored" (TS granularity=0). For the advertisement of the capabilities of such an interface, a single ISCD is used. Its format is as follows:

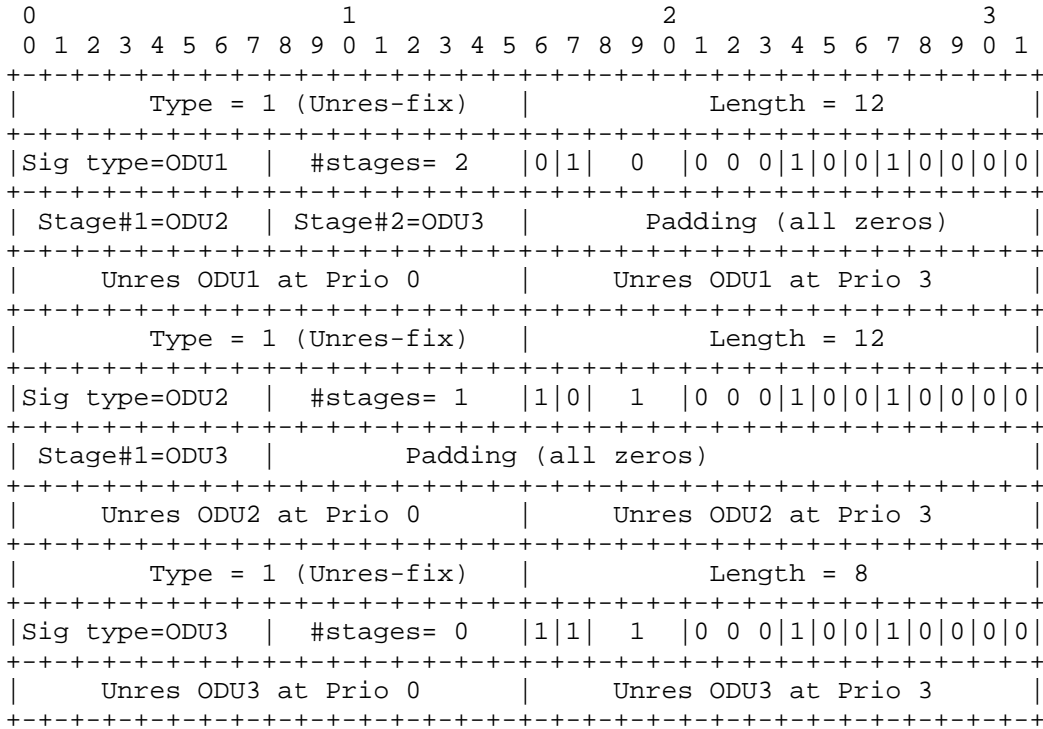


Figure 8: T, S, and TS Granularity Utilization

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 an [RFC4328] interface (TS granularity=2) while the second one supports a G.709-2012 interface with fallback procedure disabled (TS granularity=3). Both support the 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 can only be switched. For the advertisement of the capabilities of such interfaces, two different ISCDs are used. The format of their SCSIs is as follows:

SCSI of ISCD 1 -- TS granularity=2

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	2	3	4	5	6	7	8	9
Type = 1 (Unres-fix)										Length = 12																													
Sig type=ODU1										#stages= 2										0 1 0										0 0 0 1 0 0 1 0 0 0 0									
Stage#1=ODU2										Stage#2=ODU3																				Padding (all zeros)									
Unres ODU1 at Prio 0																				Unres ODU1 at Prio 3																			
Type = 1 (Unres-fix)										Length = 12																													
Sig type=ODU2										#stages= 1										1 0 1										0 0 0 1 0 0 1 0 0 0 0									
Stage#1=ODU3																														Padding (all zeros)									
Unres ODU2 at Prio 0																				Unres ODU2 at Prio 3																			
Type = 1 (Unres-fix)										Length = 8																													
Sig type=ODU3										#stages= 0										1 1 2										0 0 0 1 0 0 1 0 0 0 0									
Unres ODU3 at Prio 0																														Unres ODU3 at Prio 3									

Figure 9: Utilization of Different TS Granularities -- ISCD 1

SCSI of ISCD 2 -- TS granularity=3




Figure 10: Utilization of Different TS Granularities -- ISCD 2

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-LSP 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 a loose ERO from A; the ERO goes towards node E, and node C 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=1.25 Gbps, and the H-LSP on if2 exports a TS=2.5 Gbps; because the service LSP being signaled needs a 1.25 Gbps tributary slot, only the H-LSP on if1 can be used to reach node E. For further details, please see Section 3.2 of [RFC7096].

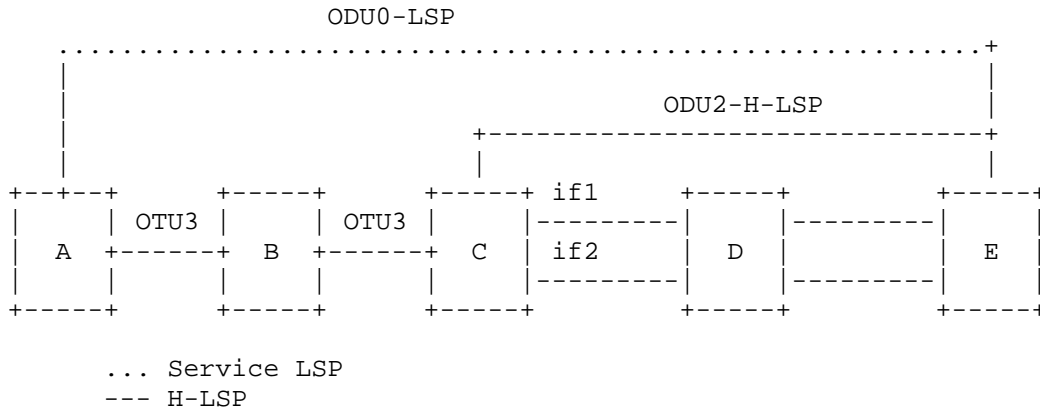


Figure 11: Example of Service LSP and H-LSP Terminating on the Same Node

5.3. Example of ODUflex Advertisement

In this example, the advertisement of an ODUflex->ODU3 hierarchy is shown. In the 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 LSPs 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. Suppose there are two TE-Links (A and B) with MAX LSP Bandwidth equal to 10 Gbps each. In the case where 50 Gbps of Unreserved Bandwidth are available on Link A, 10 Gbps 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. The T, S, and TSG fields are not relevant to this example (filled with Xs).

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

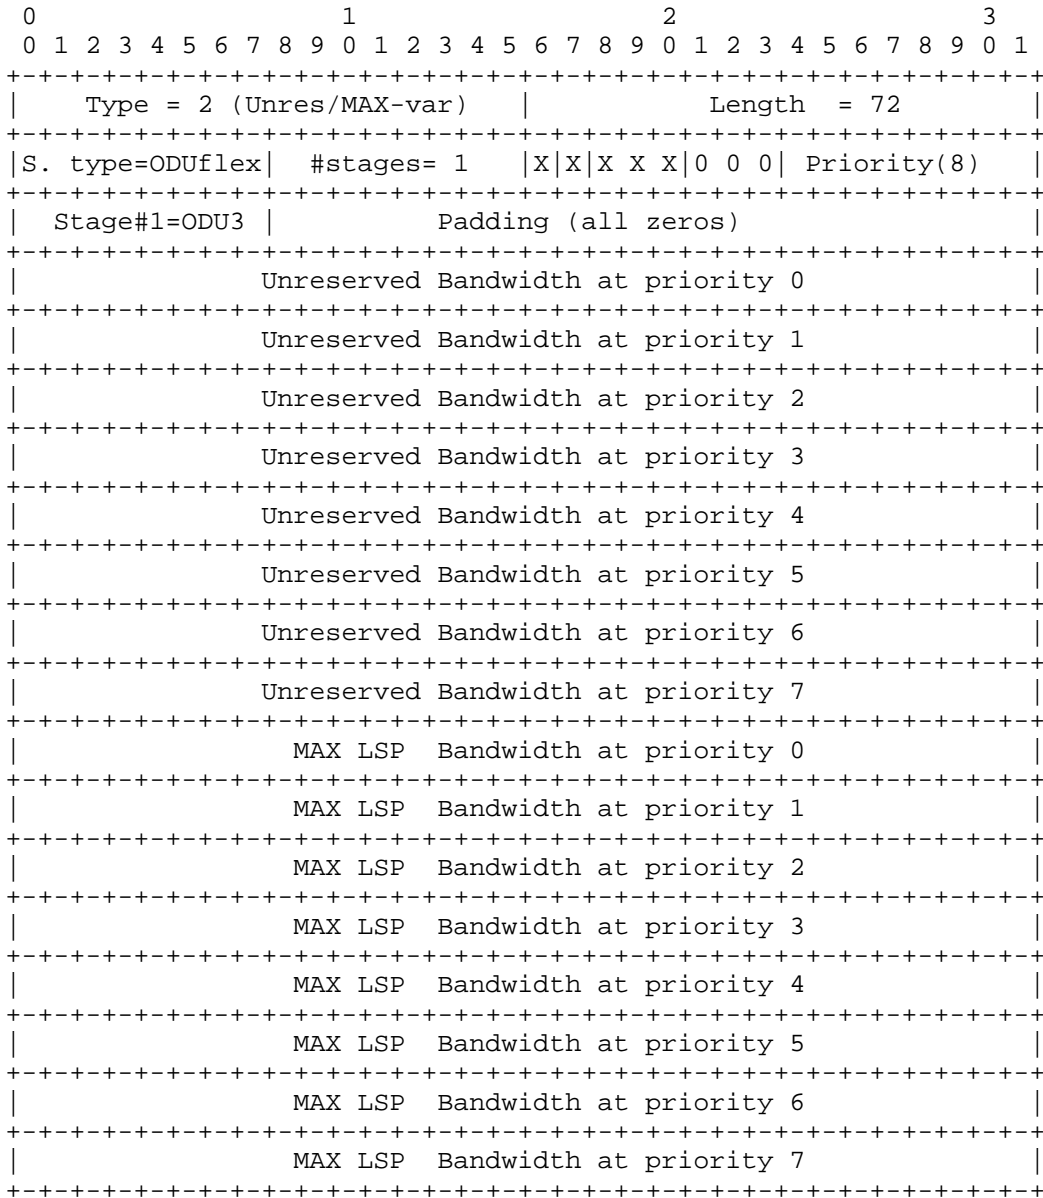


Figure 12: ODUflex Advertisement

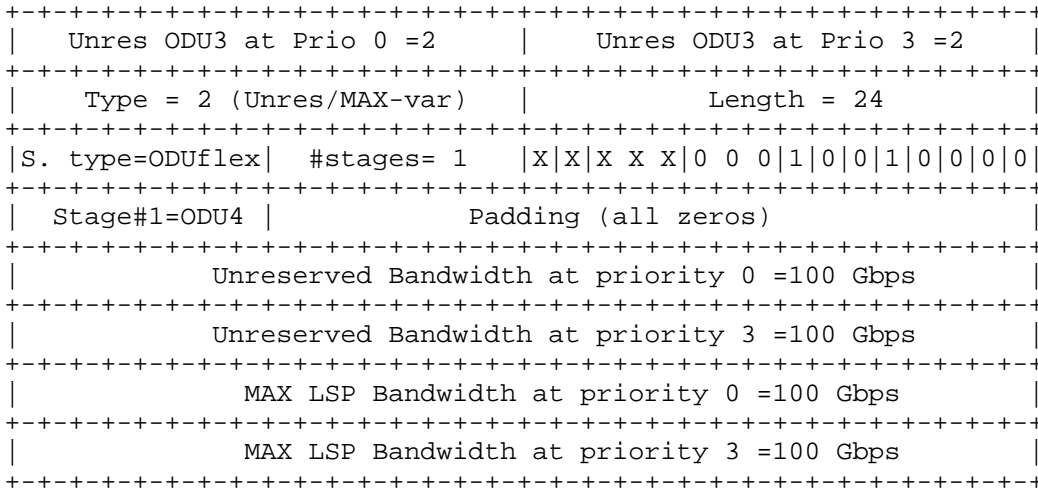



Figure 13: Single-Stage Muxing

5.5. Example of Multi-Stage Muxing -- Unbundled Link

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



Considering only 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 and filled with Xs):

```

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 = 1 (Unres-fix)          |          Length = 8          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Sig type=ODU4 | #stages= 0 |X|X| 1 |0 0 0|1|0|0|1|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Unres ODU4 at Prio 0 =1          |  Unres ODU4 at Prio 3 =1          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Type = 1 (Unres-fix)          |          Length = 12          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Sig type=ODU3 | #stages= 1 |X|X| 1 |0 0 0|1|0|0|1|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
| Stage#1=ODU4 |          Padding (all zeros)          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Unres ODU3 at Prio 0 =2          |  Unres ODU3 at Prio 3 =2          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Type = 1 (Unres-fix)          |          Length = 12          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Sig type=ODU2 | #stages= 1 |X|X| 1 |0 0 0|1|0|0|1|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
| Stage#1=ODU4 |          Padding (all zeros)          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Unres ODU2 at Prio 0 =10          |  Unres ODU2 at Prio 3 =10          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Type = 1 (Unres-fix)          |          Length = 12          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Sig type=ODU2 | #stages= 2 |X|X| 0 |0 0 0|1|0|0|1|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
| Stage#1=ODU3 | Stage#2=ODU4 |          Padding (all zeros)          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Unres ODU2 at Prio 0 =8          |  Unres ODU2 at Prio 3 =8          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Type = 1 (Unres-fix)          |          Length = 12          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Sig type=ODU0 | #stages= 2 |X|X| 0 |0 0 0|1|0|0|1|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
| Stage#1=ODU3 | Stage#2=ODU4 |          Padding (all zeros)          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Unres ODU0 at Prio 0 =64          |  Unres ODU0 at Prio 3 =64          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Type = 1 (Unres-fix)          |          Length = 12          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Sig type=ODU0 | #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)          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Unres ODU0 at Prio 0 =80          |  Unres ODU0 at Prio 3 =80          |

```



```


+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|   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 =100 Gbps           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           Unreserved Bandwidth at priority 3 =100 Gbps           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           MAX LSP Bandwidth at priority 0 =10 Gbps               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           MAX LSP Bandwidth at priority 3 =10 Gbps               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 14: 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 (the T, S, and TSG fields are not relevant to this example and filled with Xs):

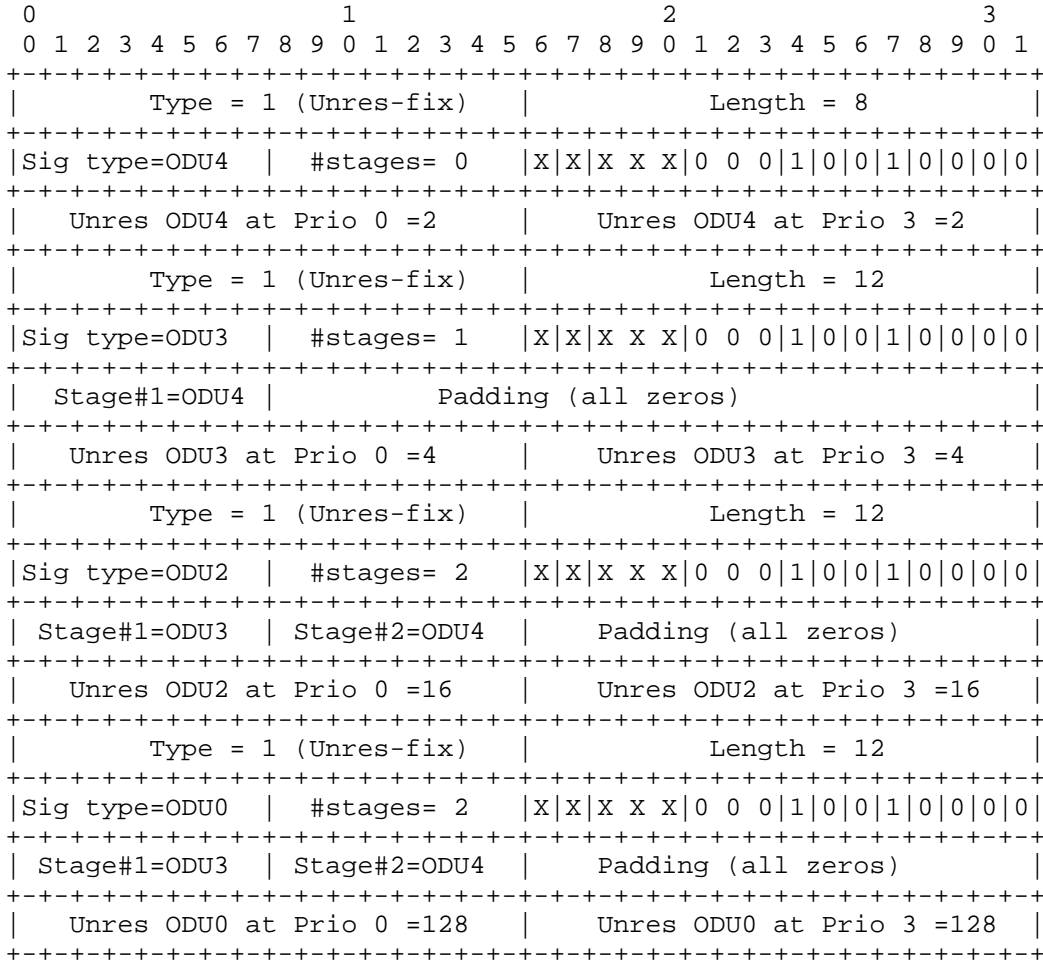



Figure 15: Multi-Stage Muxing -- Bundled Links

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 (the T, S, and TSG fields are not relevant to this example and filled with Xs). In the following figure, the SCSI of each ISCD is shown:

SCSI of ISCD 1 -- Component Link#1

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 = 1 (Unres-fix)								Length = 8																							
Sig type=ODU4								#stages= 0								X X X X X 0 0 0 1 0 0 1 0 0 0 0															
Unres ODU4 at Prio 0 =1								Unres ODU4 at Prio 3 =1																							
Type = 1 (Unres-fix)								Length = 12																							
Sig type=ODU3								#stages= 1								X X X X X 0 0 0 1 0 0 1 0 0 0 0															
Stage#1=ODU4								Padding (all zeros)																							
Unres ODU3 at Prio 0 =2								Unres ODU3 at Prio 3 =2																							
Type = 1 (Unres-fix)								Length = 12																							
Sig type=ODU2								#stages= 2								X X X X X 0 0 0 1 0 0 1 0 0 0 0															
Stage#1=ODU3								Stage#2=ODU4								Padding (all zeros)															
Unres ODU2 at Prio 0 =8								Unres ODU2 at Prio 3 =8																							
Type = 1 (Unres-fix)								Length = 12																							
Sig type=ODU0								#stages= 2								X X X X X 0 0 0 1 0 0 1 0 0 0 0															
Stage#1=ODU3								Stage#2=ODU4								Padding (all zeros)															
Unres ODU0 at Prio 0 =64								Unres ODU0 at Prio 3 =64																							

Figure 16: Multi-Stage Muxing -- Non-Homogeneous Hierarchies -- ISCD 1

SCSI of ISCD 2 -- Component Link#2

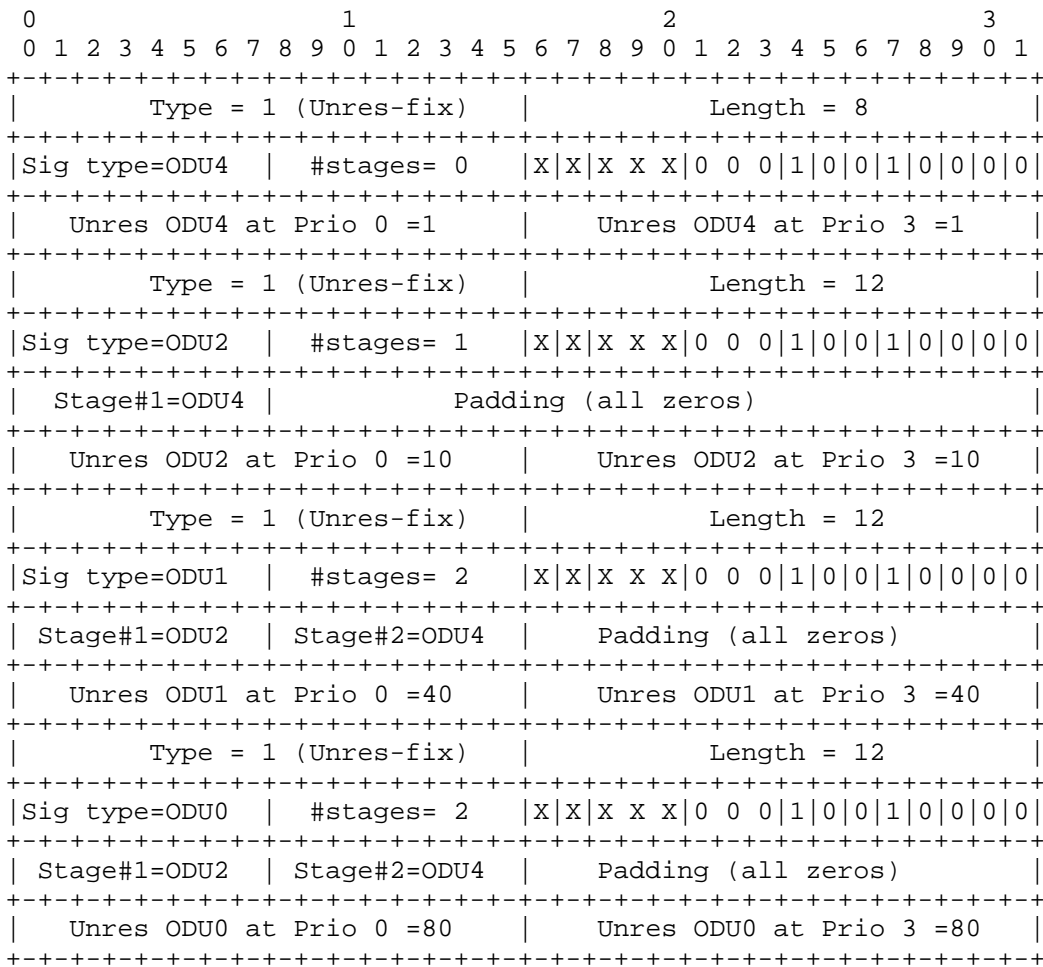


Figure 17: Multi-Stage Muxing -- Non-Homogeneous Hierarchies -- ISCD 2

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, in particular:

- o Only unreserved and MAX LSP Bandwidth related to supported priorities are advertised.
- o For fixed containers, only the number of available containers is advertised instead of the available bandwidth in order 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 [RFC4203]. When nodes support both the advertisement method in [RFC4203] and the one in this document, 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 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 (optical and/or Synchronous Optical Network - Synchronous Digital Hierarchy (SONET-SDH) network. [RFC3630] notes that the security mechanisms described in [RFC2328] apply to Opaque LSAs carried in OSPFv2. An analysis of the security of OSPF is provided in [RFC6863] and applies to the extensions to OSPF as described in this document. Any new mechanisms developed to protect the transmission of information carried in Opaque LSAs will also automatically protect the extensions defined in this document.

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

9. IANA Considerations

9.1. Switching Types

IANA has made the following assignment in the "Switching Types" section of the "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" registry located at <http://www.iana.org/assignments/gmpls-sig-parameters>:

Value	Name	Reference
110	OTN-TDM capable	[RFC7138]

The same type of modification has been applied to the IANA-GMPLS-TC-MIB at <https://www.iana.org/assignments/ianagmplstc-mib>, where the value:

OTN-TDM (110), -- Time-Division-Multiplex OTN-TDM capable

has been added to the IANAGmplsSwitchingTypeTC ::= TEXTUAL-CONVENTION syntax list.

9.2. New Sub-TLVs

This document defines 2 new sub-TLVs that are carried in Interface Switching Capability Descriptors [RFC4203] with the 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.

IANA has created and will maintain a new sub-registry, the "Types for sub-TLVs of OTN-TDM SCSI (Switching 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>, with the sub-TLV types as follows:

Value	Sub-TLV	Reference
0	Reserved	[RFC7138]
1	Unreserved Bandwidth for fixed containers	[RFC7138]
2	Unreserved/MAX Bandwidth for flexible containers	[RFC7138]
3-65535	Unassigned	

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

10. Contributors

Diego Caviglia
Ericsson
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

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 Kunjidhapatham
EMail: akunjidhapatham@infinera.com

Snigdho Bardalai
EMail: sbardalai@infinera.com

Steve Balls
EMail: Steve.Balls@metaswitch.com

Jonathan Hardwick
EMail: Jonathan.Hardwick@metaswitch.com

Xihua Fu
EMail: fu.xihua@zte.com.cn

Cyril Margaria
EMail: cyril.margaria@nsn.com

Malcolm Betts
EMail: Malcolm.betts@zte.com.cn

11. Acknowledgements

The authors would like to thank Fred Gruman and Lou Berger for their valuable comments and suggestions.

12. References

12.1. Normative References

- [G.709-2012] ITU-T, "Interface for the optical transport network", Recommendation G.709/Y.1331, February 2012.
- [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

- [G.798] ITU-T, "Characteristics of optical transport network hierarchy equipment functional blocks", Recommendation G.798, December 2012.
- [G.805] ITU-T, "Generic functional architecture of transport networks", Recommendation G.805, March 2000.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [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 (WSONs)", RFC 6163, April 2011.
- [RFC6566] Lee, Y., Bernstein, G., Li, D., and G. Martinelli, "A Framework for the Control of Wavelength Switched Optical Networks (WSONs) with Impairments", RFC 6566, March 2012.
- [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.
- [RFC7062] Zhang, F., Li, D., Li, H., Belotti, S., and D. Ceccarelli, "Framework for GMPLS and PCE Control of G.709 Optical Transport Networks", RFC 7062, November 2013.

- [RFC7096] Belotti, S., Grandi, P., Ceccarelli, D., Ed., Caviglia, D., and F. Zhang, "Evaluation of Existing GMPLS Encoding against G.709v3 Optical Transport Networks (OTNs)", RFC 7096, January 2014.
- [RFC7139] Zhang, F., Ed., Zhang, G., Belotti, S., Ceccarelli, D., and K. Pithewan, "GMPLS Signaling Extensions for Control of Evolving G.709 Optical Transport Networks", RFC 7139, March 2014.

Authors' Addresses

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

E-Mail: daniele.ceccarelli@ericsson.com

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129
P.R. China

Phone: +86-755-28972912
E-Mail: zhangfatai@huawei.com

Sergio Belotti
Alcatel-Lucent
Via Trento, 30
Vimercate
Italy

E-Mail: sergio.belotti@alcatel-lucent.com

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

E-Mail: rrao@infinera.com

John E. Drake
Juniper

E-Mail: jdrake@juniper.net

