CCAMP Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 9, 2012

D. Ceccarelli, Ed.
D. Caviglia
Ericsson
F. Zhang
D. Li
Huawei Technologies
S. Belotti
P. Grandi
Alcatel-Lucent
R. Rao
K. Pithewan
Infinera Corporation
J. Drake
Juniper
September 6, 2011

Traffic Engineering Extensions to OSPF for Generalized MPLS (GMPLS) Control of Evolving G.709 OTN Networks
draft-ceccarelli-ccamp-gmpls-ospf-g709-07

Abstract

The recent revision of ITU-T Recommendation G.709 [G709-V3] has introduced new fixed and flexible ODU containers, enabling optimized support for an increasingly abundant service mix.

This document describes OSPF routing protocol extensions to support Generalized MPLS (GMPLS) control of all currently defined ODU containers, in support of both sub-lambda and lambda level routing granularity.

Table of Contents

1. Introduction

G.709 OTN [G709-V3] includes new fixed and flexible ODU containers, two types of Tributary Slots (i.e., 1.25Gbps and 2.5Gbps), and supports various multiplexing relationships (e.g., ODUj multiplexed into ODUk (j<k)), two different tributary slots for ODUk (K=1, 2, 3) and ODUflex service type, which is being standardized in ITU-T. In order to present this information in the routing process, this document provides OTN technology specific encoding for OSPF-TE.

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].

The routing information for Optical Channel Layer (OCh) (i.e., wavelength) is out of the scope of this document. Please refer to [WSON-Frame] 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 OTUk can be viewed as a component link, and each component link can carry one or more types of ODUj (j<k).

Each TE 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 [RFC 3630]. (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 [RFC 3630] and [RFC 4203].

As discussed in [OTN-FWK] and [OTN-INFO], the OSPF-TE must be extended so to be able to advertise the termination and switching capabilites related to each different ODUj and ODUk/OTUk and the advertisement of related multiplexing capabilities. This leads to the need to define a new Switching Capability value and associated new Switching Capability for the ISCD.

In the following we will use ODUj to indicate a service type that is multiplexed into an higher order ODU, ODUk an 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 iterated 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.

If for example a multiplexing hierarchy like the following one is considered:

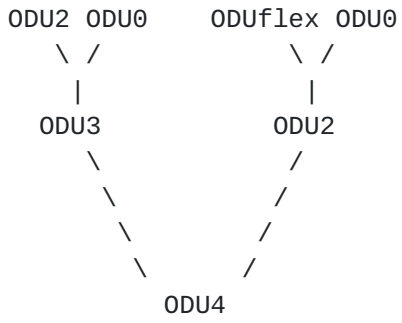

The ODU4 is the root of the muxing tree, ODU3 and ODU2 are containers directly multiplexed into the server and then ODU2, ODU0 are the leaves of the ODU3 branch, while ODUflex and ODU0 are the leaves of the ODU2 one. This means that on this traffic card it is possible to have the following multiplexing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

representations are captured below.

OTUk physical Link(s) can be modeled as a TE-Link(s). The TE-Link is termed as OTUk-TE-Link. The OTUk-TE-Link advertises ODUj switching capacity. The advertised capacity could include ODUk switching capacity. Figure-1 below provides an illustration of one hop ODUk TE-links.

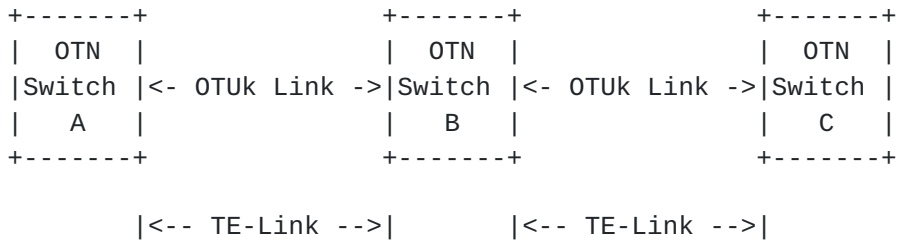

Figure 1: ODUk TE-Links

It is possible to create TE-Links that span more than one hop by creating FA between non-adjacent nodes. Such TE-Links are also termed ODUk-TE-Links. As in the one hop case, these types of ODUk-TE-Links also advertise ODUj switching capacity. The advertised capacity could include ODUk switching capacity.

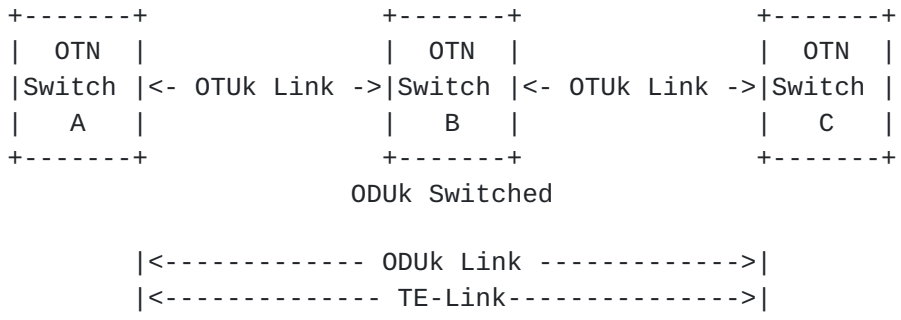

Figure 2: Multiple hop TE-Link

4. ISCD format extensions

The Interface Switching Capability Descriptor describes switching capability of an interface [RFC 4202]. This document defines a new Switching Capability value for OTN [G.709-v3] as follows:

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

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]

Both fixed and flexible ODUs use the same switching type and encoding values. When Switching Capability and Encoding fields are set to values as stated above, the Interface Switching Capability Descriptor should be interpreted as follows:

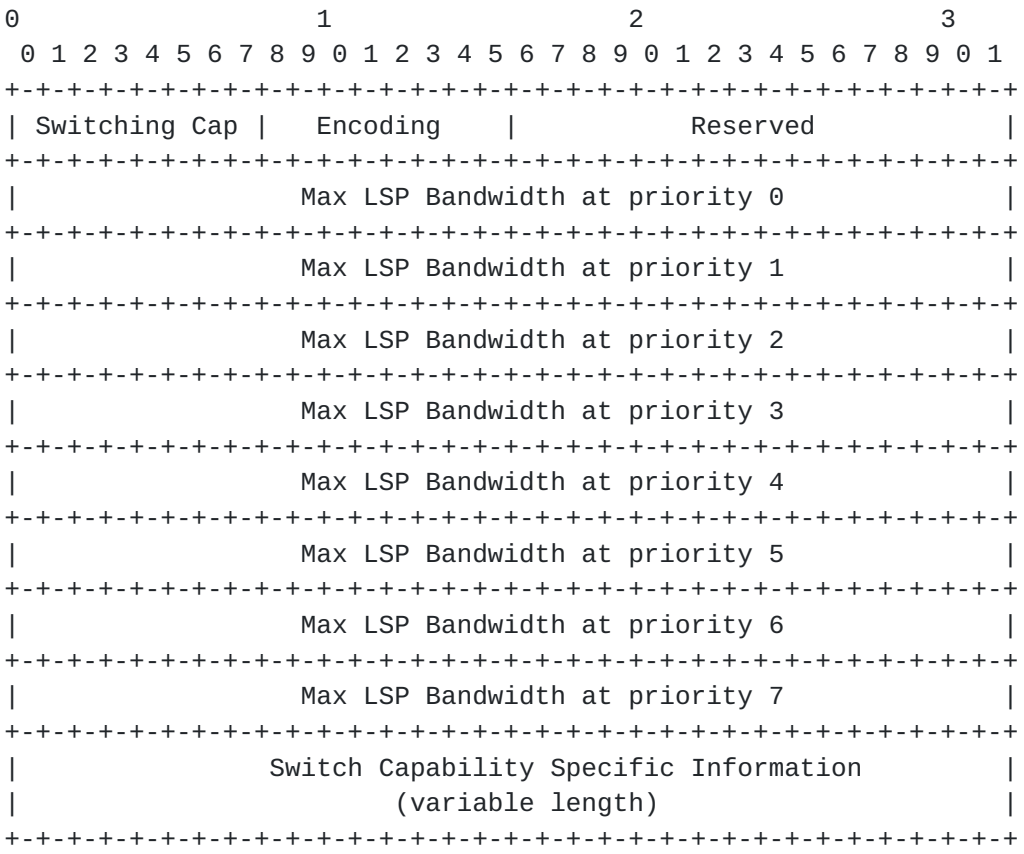

Maximum LSP Bandwidth

The MAX LSP bandwidth field is used according to [RFC4203]: i.e. 0 <= Max LSP Bandwidth <= ODUk/OTUk and intermediate values are those on

the branch of OTN switching hierarchy supported by the interface. E.g. 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 is in IEEE floating point format. The discrete values for various ODUs is shown in the table below.

ODU Type	ODU nominal bit rate	Value in Byte/Sec
ODU0	1 244 160 kbits/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 Tributary Slot Granularity (TSG). A different ISCD MUST be used for each different muxing hierarchy (muxing tree in the following examples) and different TSG supported within the TE Link, if it includes component links with differing characteristics.

4.1. Switch Capability Specific Information

The technology specific part of the OTN ISCD can include a variable number of sub-TLVs called Bandwidth sub-TLVs. The muxing hierarchy tree is encoded as an order independent list of them. Three types of Bandwidth TLV are defined (TBA by IANA):

- Type 1 - Unreserved Bandwidth for fixed containers

- Type 2 - Unreserved Bandwidth for flexible containers

- Type 3 - Max LSP Bandwidth for flexible containers

The format of the SCSI is depicted in the following figure:

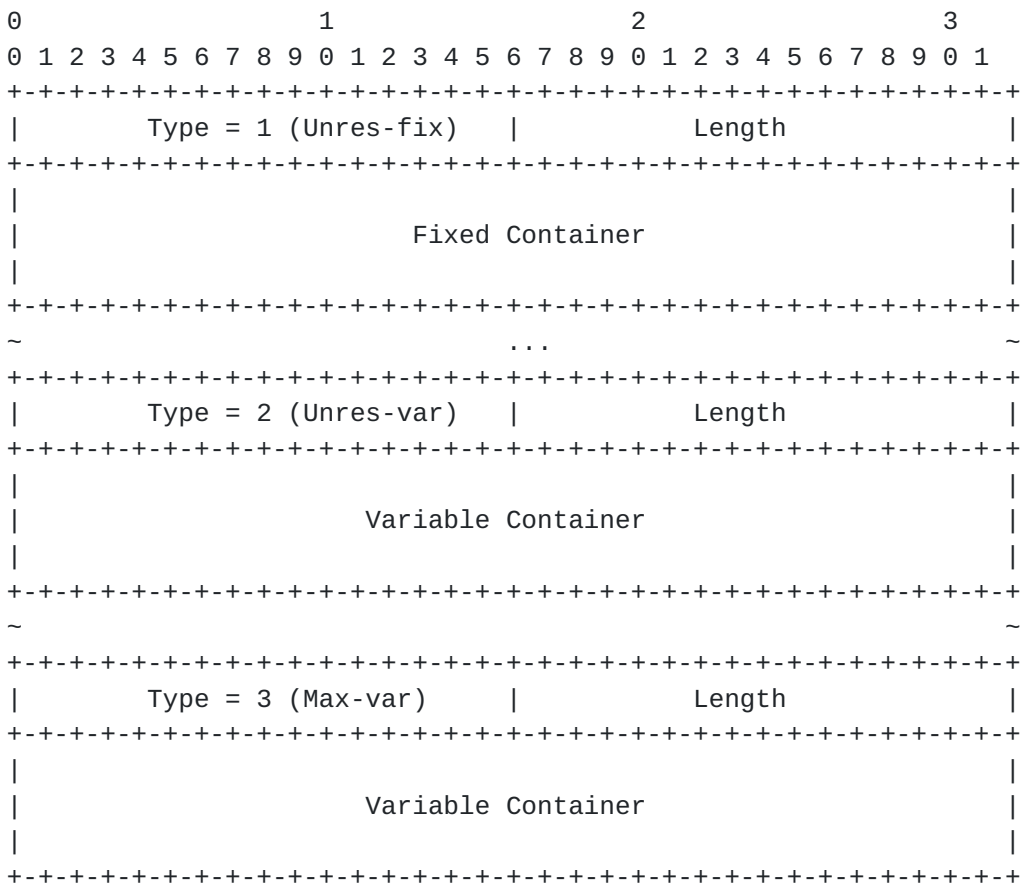

Figure 3: SCSI format

The format of the three different types of Bandwidth TLV are depicted in the following figures:

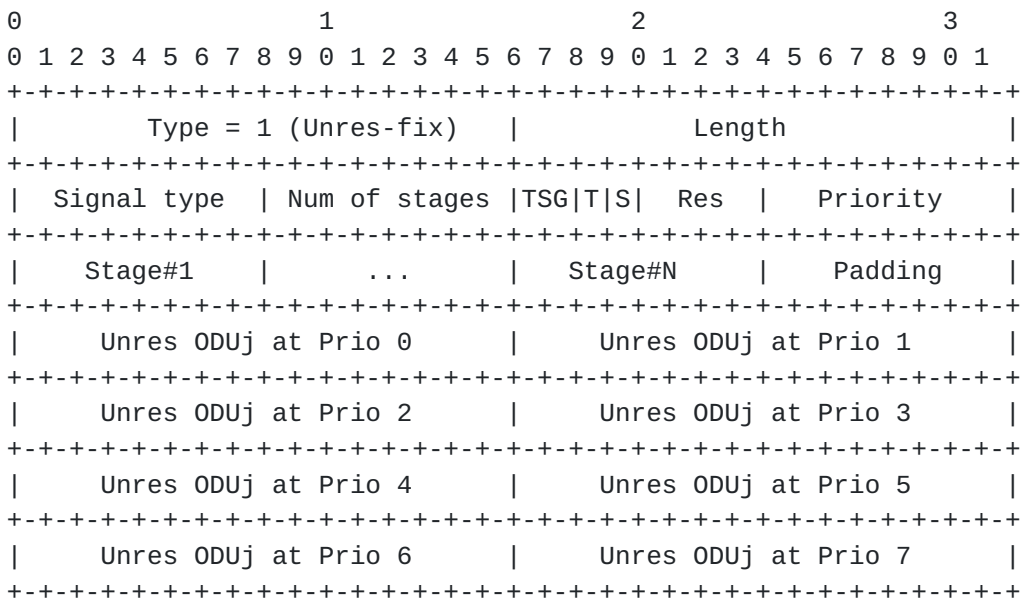

Figure 4: Bandwidth TLV - Type 1 -

The values of the fields shown in figure 4 are explained after figure 6.

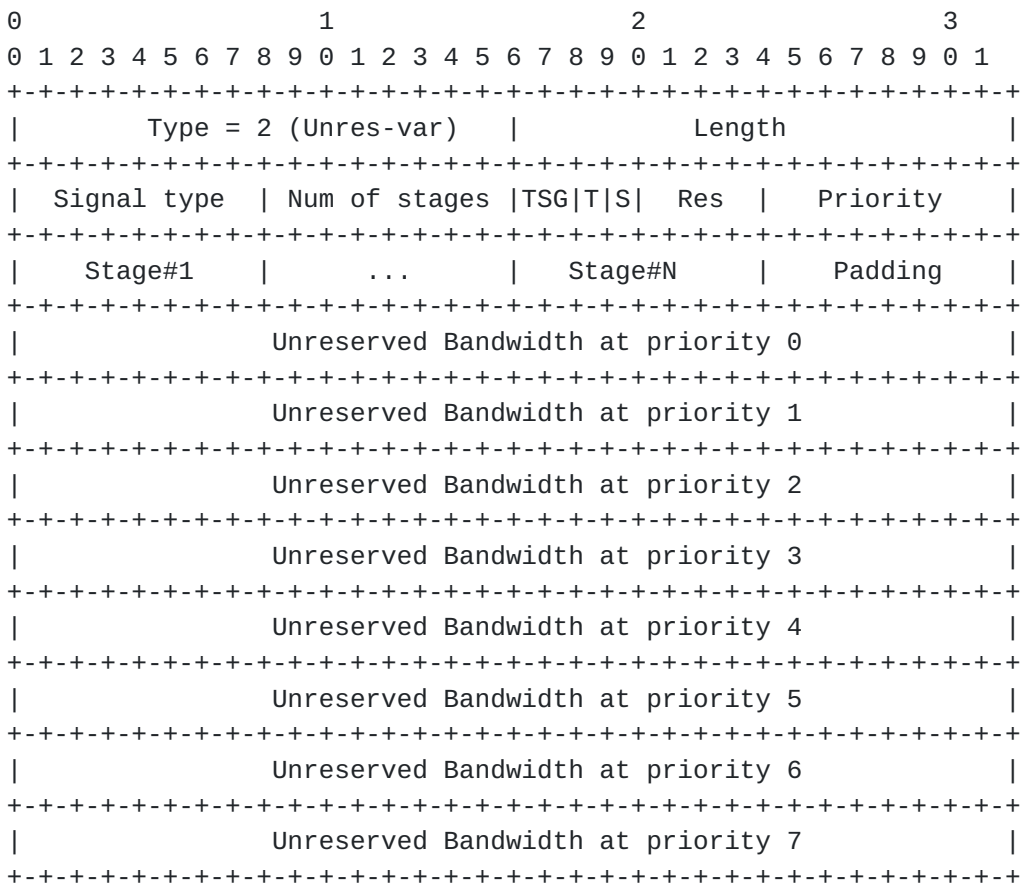

Figure 5: Bandwidth TLV - Type 2 -

While Bandwidth TLV Type 3 is as follows:

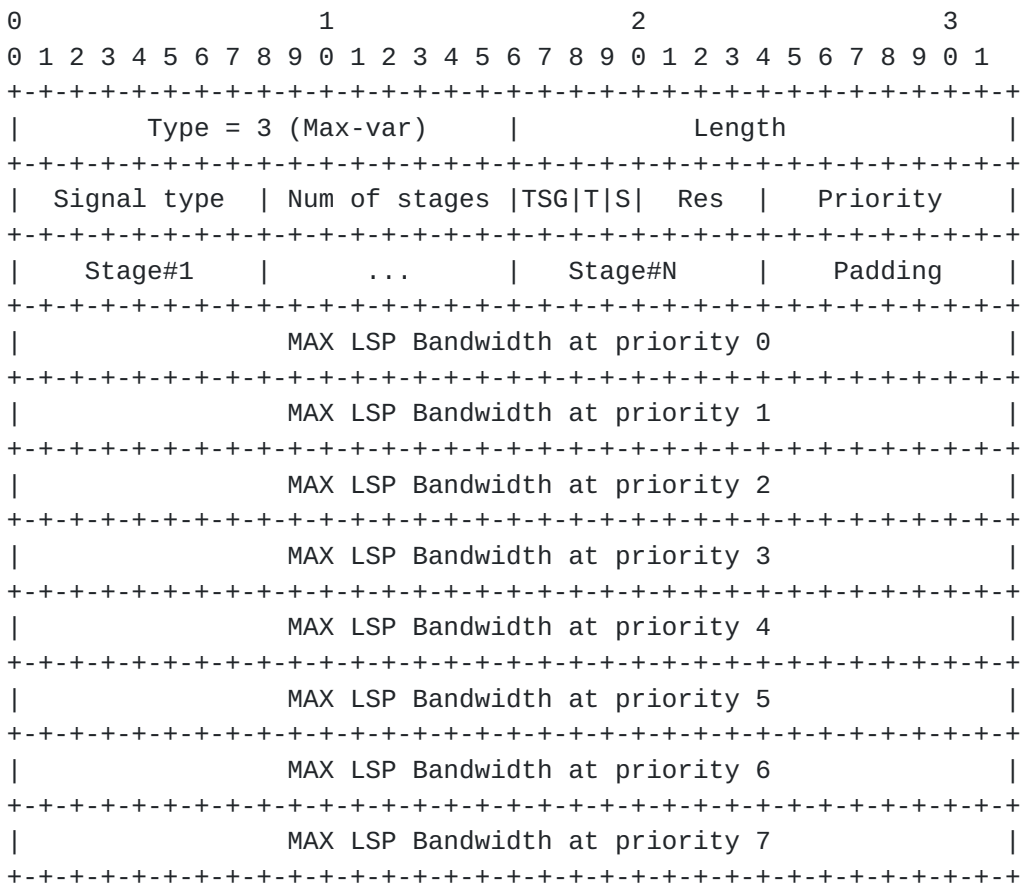

Figure 6: Bandwidth TLV - Type 3 -

- Signal Type: Indicates the ODU type being advertised

Value	Type
1	ODU1
2	ODU2
3	ODU3
4	ODU4
10	ODU0
11	ODU2e
20	ODUflex CBR
21	ODUflex GFP-F resizable
22	ODUflex GFP-F non resizable
60000-65535	Experimental

- Number of stages: Indicates the number of multiplexing stages level. It is equal to 0 when a server layer is being advertised, 1 in case of single stage muxing, 2 in case of dual stage muxing, etc.

- TSG: Tributary Slot Granularity (2bit): Used for the advertisement of the supported Tributary Slot granularity

  - 00 - Reserved

  - 01 - 1.25 Gbps

  - 10 - 2.5 Gbps

  - 11 - Reserved

- Flags:

  - T Flag (bit 17): Indicates whether the advertised bandwidth can be terminated. When T=1, the signal type can be terminated, when T=0, the signal type cannot be terminated.

  - S Flag (bit 18): Indicates whether the advertised bandwidth can be switched. When S=1, the signal type can be switched, when S=0, the signal type cannot be switched.

  The value 00 in both T and S bits is not permitted.

- Priority :8 bits field with 1 flag for each priority. Bit set indicates priority supported, bit cleared means priority not supported. The priority 0 is related to the most significant bit. When no priority is supported, priority 0 MUST be advertised.

- Stage#1 ... Stage#N : These fields are 8 bits long. Their number is variable and a field is present for each stage of the muxing hierarchy. The last one is always indicating the server ODU container (ODUk/OTUk). The values of the Stage fields are the same ones defined for the Signal Type field. If the number of stages is 0, then no Stage fields are included.

- Padding: Given that the number of Stages is variable, padding to 32 bits field is used as needed.

- Unreserved Bandwidth/Max LSP BW : In case of fixed containers (Type=1) the Unreserved Bandwidth field MUST be 16 bits long and indicates the Unreserved Bandwidth in number of available containers. Only Unreserved/MAX LSP BW fields for supported priorities are included, in order of increasing prioritiy (0 to

7). In case the number of supported priorities is odd, a 16 bits all zeros padding field is added. On the other hand, in case of variable containers (Type 2-3) the Unreserved/MAX LSP Bandwidth fields MUST be 32 bits long and expressed in IEEE floating point format. Only Unreserved/MAX LSP bandwidth for supported priorities MUST be advertised.

5. Examples

The examples in the following pages are not normative and are not intended to infer 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:

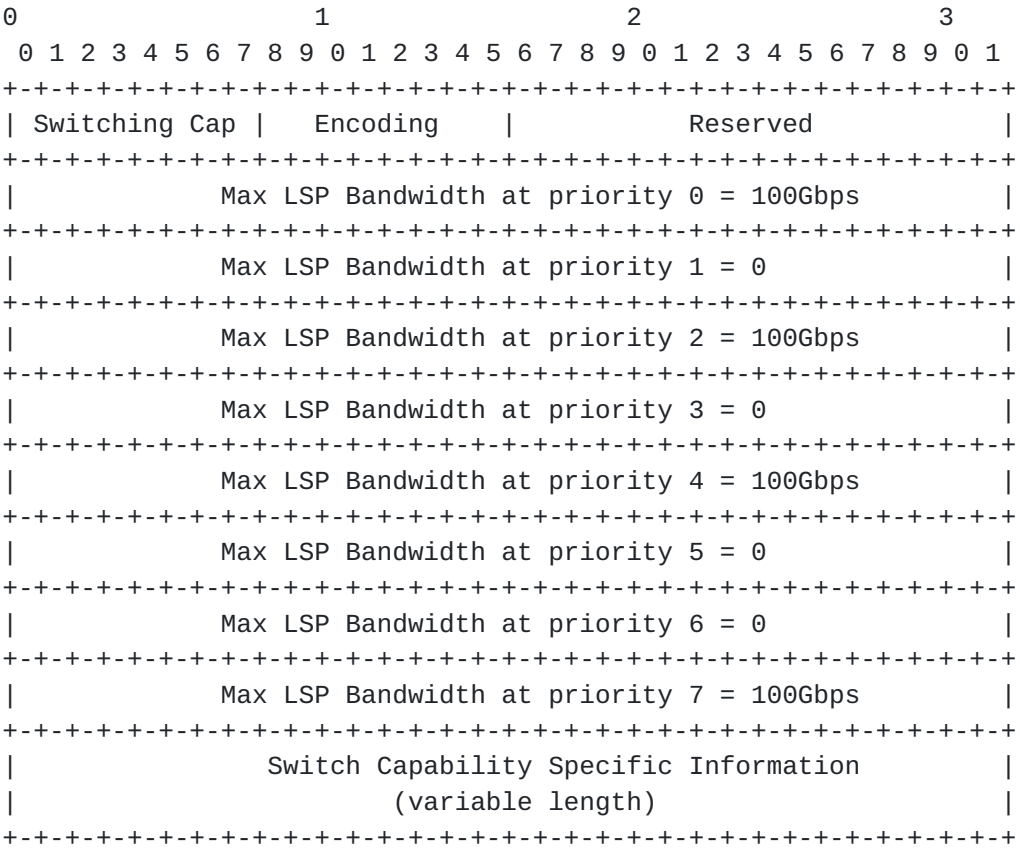

Figure 7: Example 1 - MAX LSP Bandwidth fields in the ISCD @T0

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:

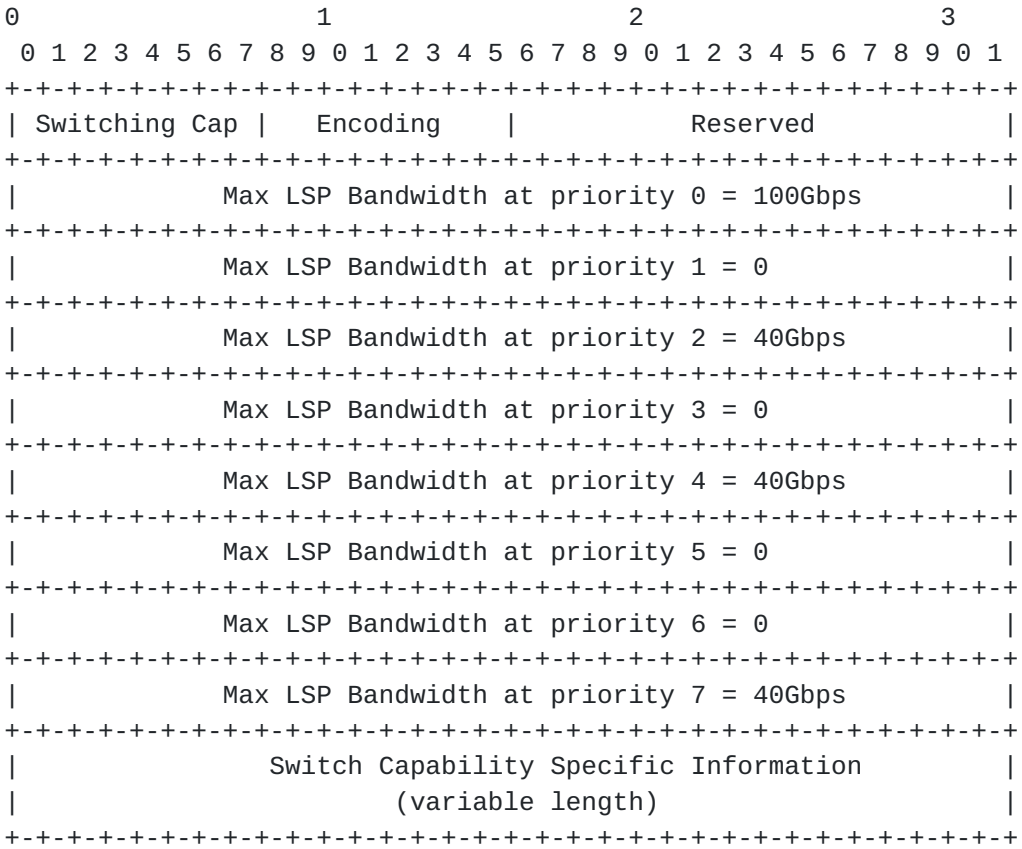

Figure 8: Example 1 - MAX LSP Bandwidth fields in the ISCD @T1

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:

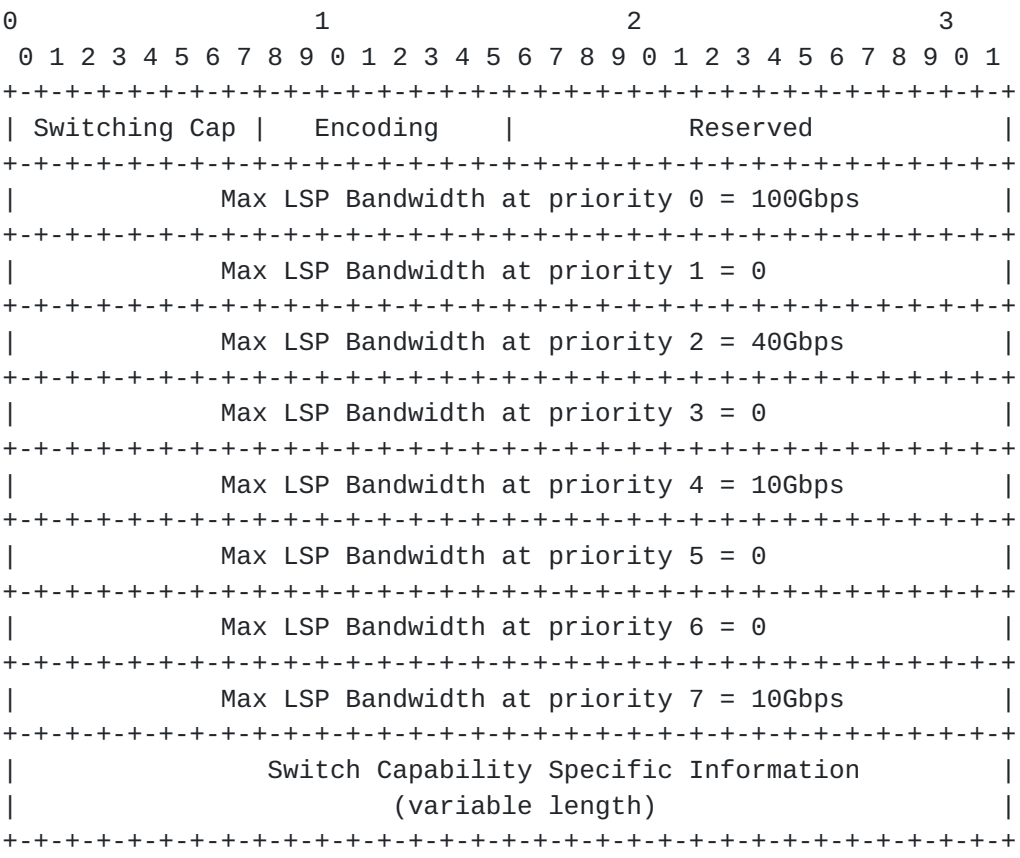

Figure 9: Example 1 - MAX LSP Bandwidth fields in the ISCD @T2

5.2. Example of TSG, T and S utilization

In this example an interface with Tributary Slot Type 1.25 Gbps is considered. 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. For the advertisement of the capabilities of such interface a single ISCD is used and its format is as follows:

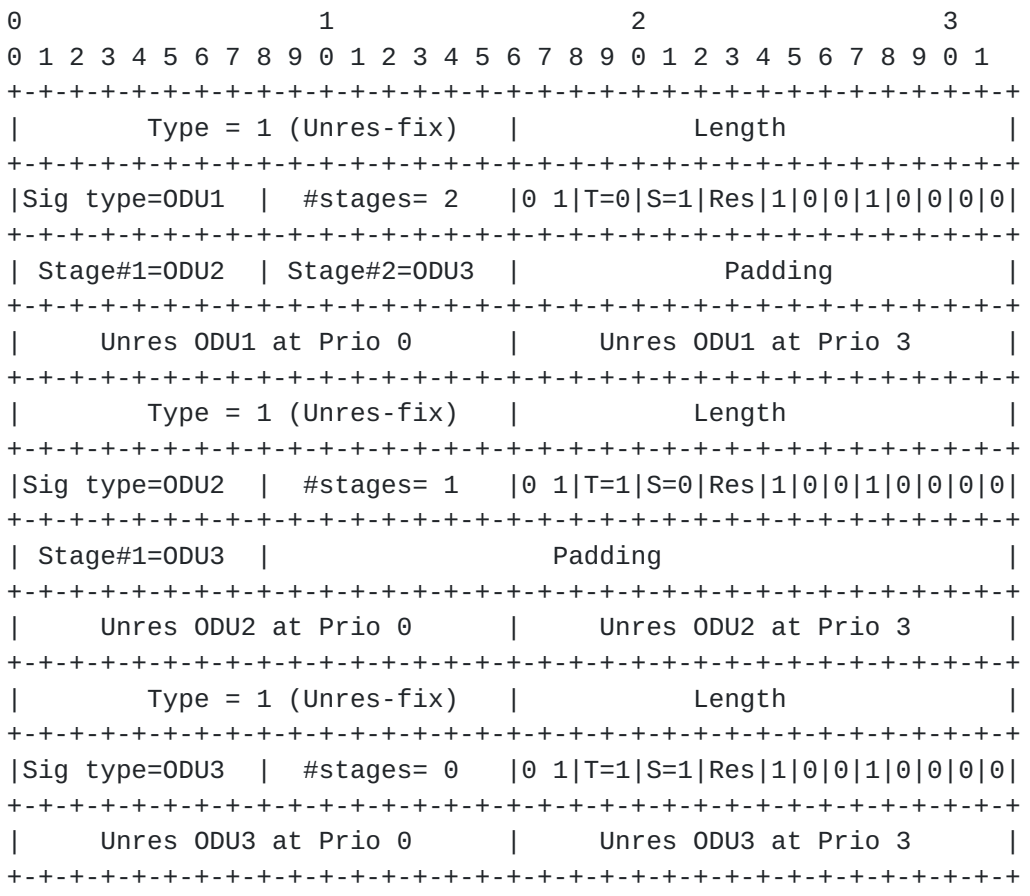

Figure 10: Example 2 - TSG, T and S utilization

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 but in some cases also information about the Unreserved bandwidth could 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 case 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 sure) that two of them can be restored along Link A.

In the case of ODUflex advertisement both the Type 2 and Type 3 Bandwidth TLVs are used.

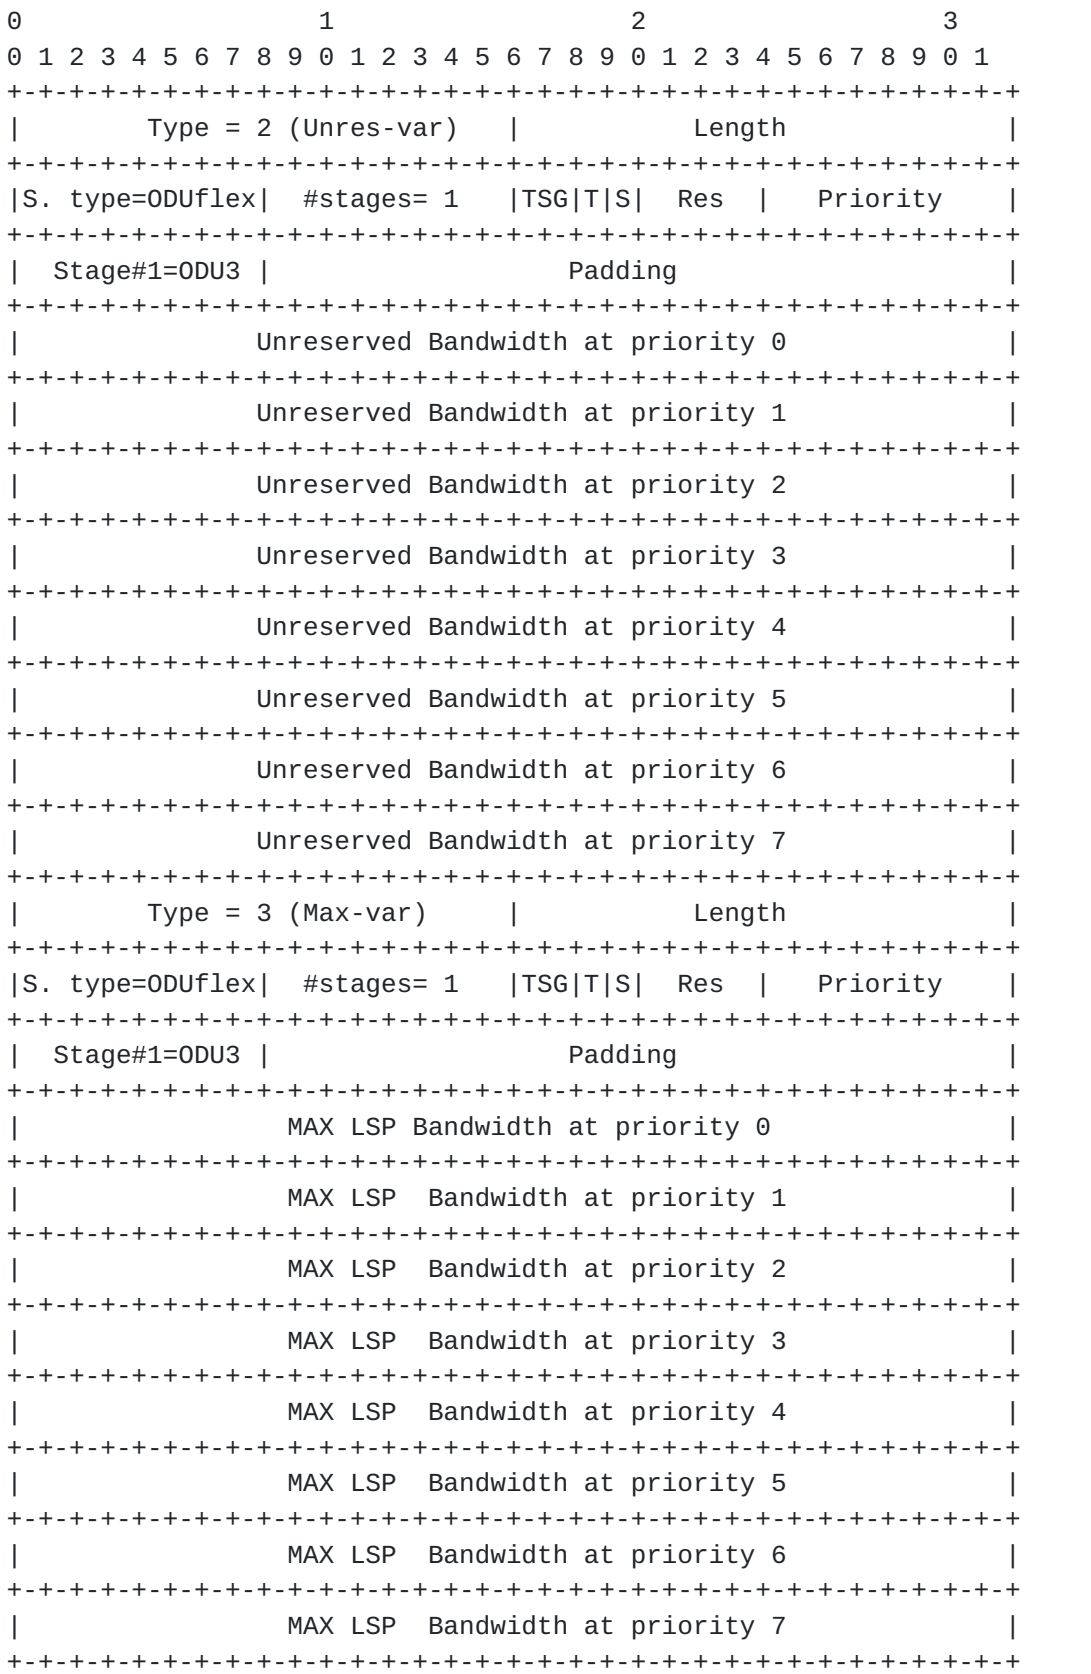
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Type = 2 (Unres-var)
Length
S. type=ODUflex
#stages= 1
TSG
T
S
Res
Priority
Stage#1=ODU3
Padding
Unreserved Bandwidth at priority 0
Unreserved Bandwidth at priority 1
Unreserved Bandwidth at priority 2
Unreserved Bandwidth at priority 3
Unreserved Bandwidth at priority 4
Unreserved Bandwidth at priority 5
Unreserved Bandwidth at priority 6
Unreserved Bandwidth at priority 7
Type = 3 (Max-var)
Length
S. type=ODUflex
#stages= 1
TSG
T
S
Res
Priority
Stage#1=ODU3
Padding
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 11: 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 assume that also in this case only priorities 0 and 3 are supported.

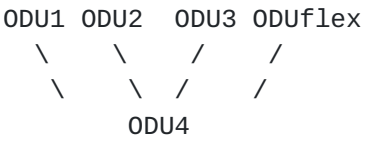

and the related SCSIs as follows:

```
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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Type = 1 (Unres-fix)   |           Length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Sig type=ODU4  |  #stages= 0   |TSG|T|S|  Res  |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              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Sig type=ODU1  |  #stages= 1   |TSG|T|S|  Res  |1|0|0|1|0|0|0|0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Stage#1=ODU4  |                  Padding                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Unres ODU1 at Prio 0 =40    |    Unres ODU1 at Prio 3 =40   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Type = 1 (Unres-fix)   |           Length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Sig type=ODU2  |  #stages= 1   |TSG|T|S|  Res  |1|0|0|1|0|0|0|0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Stage#1=ODU4  |                  Padding                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Unres ODU2 at Prio 0 =10    |    Unres ODU2 at Prio 3 =10   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Type = 1 (Unres-fix)   |           Length              |
```

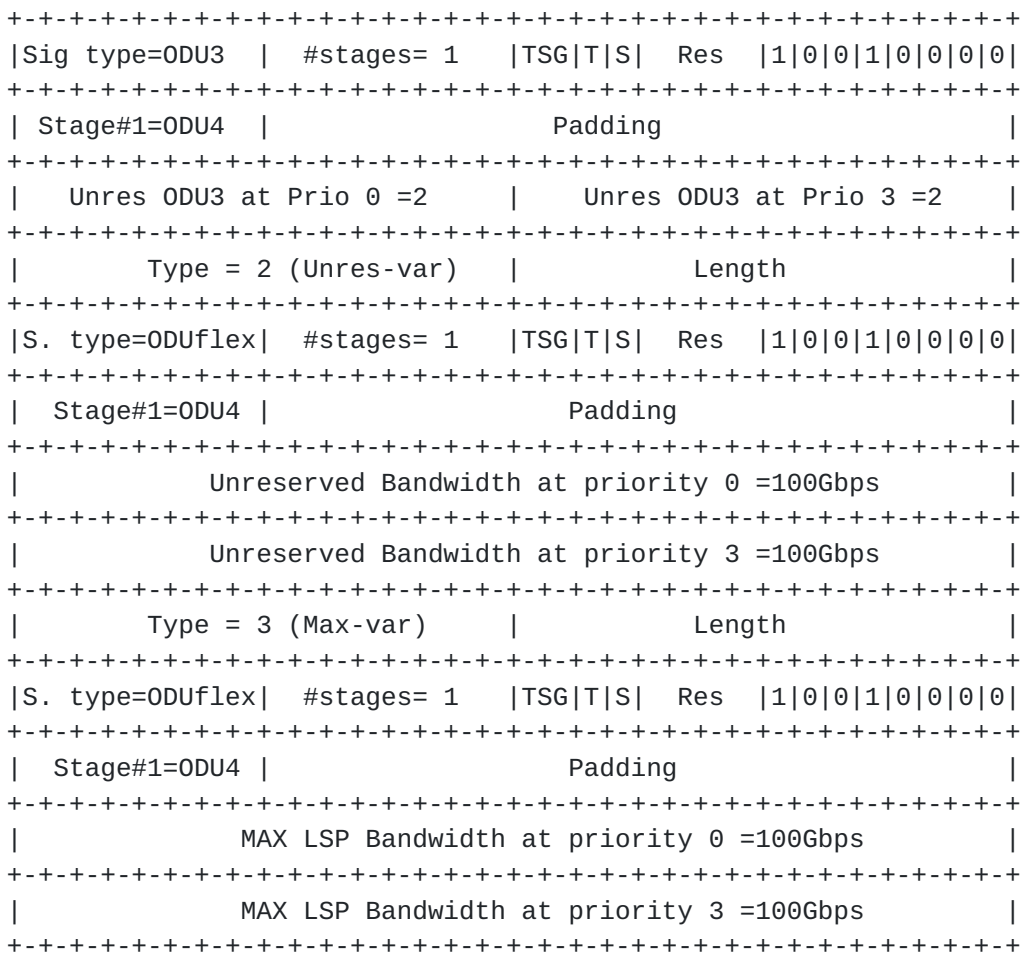

Figure 12: Example 4 - Single stage muxing

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:

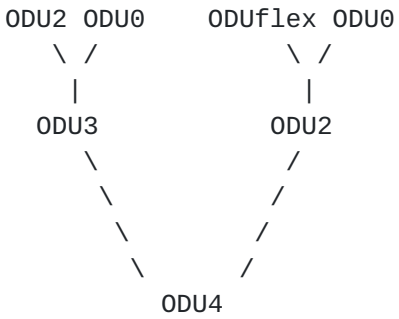

and supported pririties 0 and 3, the advertisement is composed by the following Bandwidth TLVs:

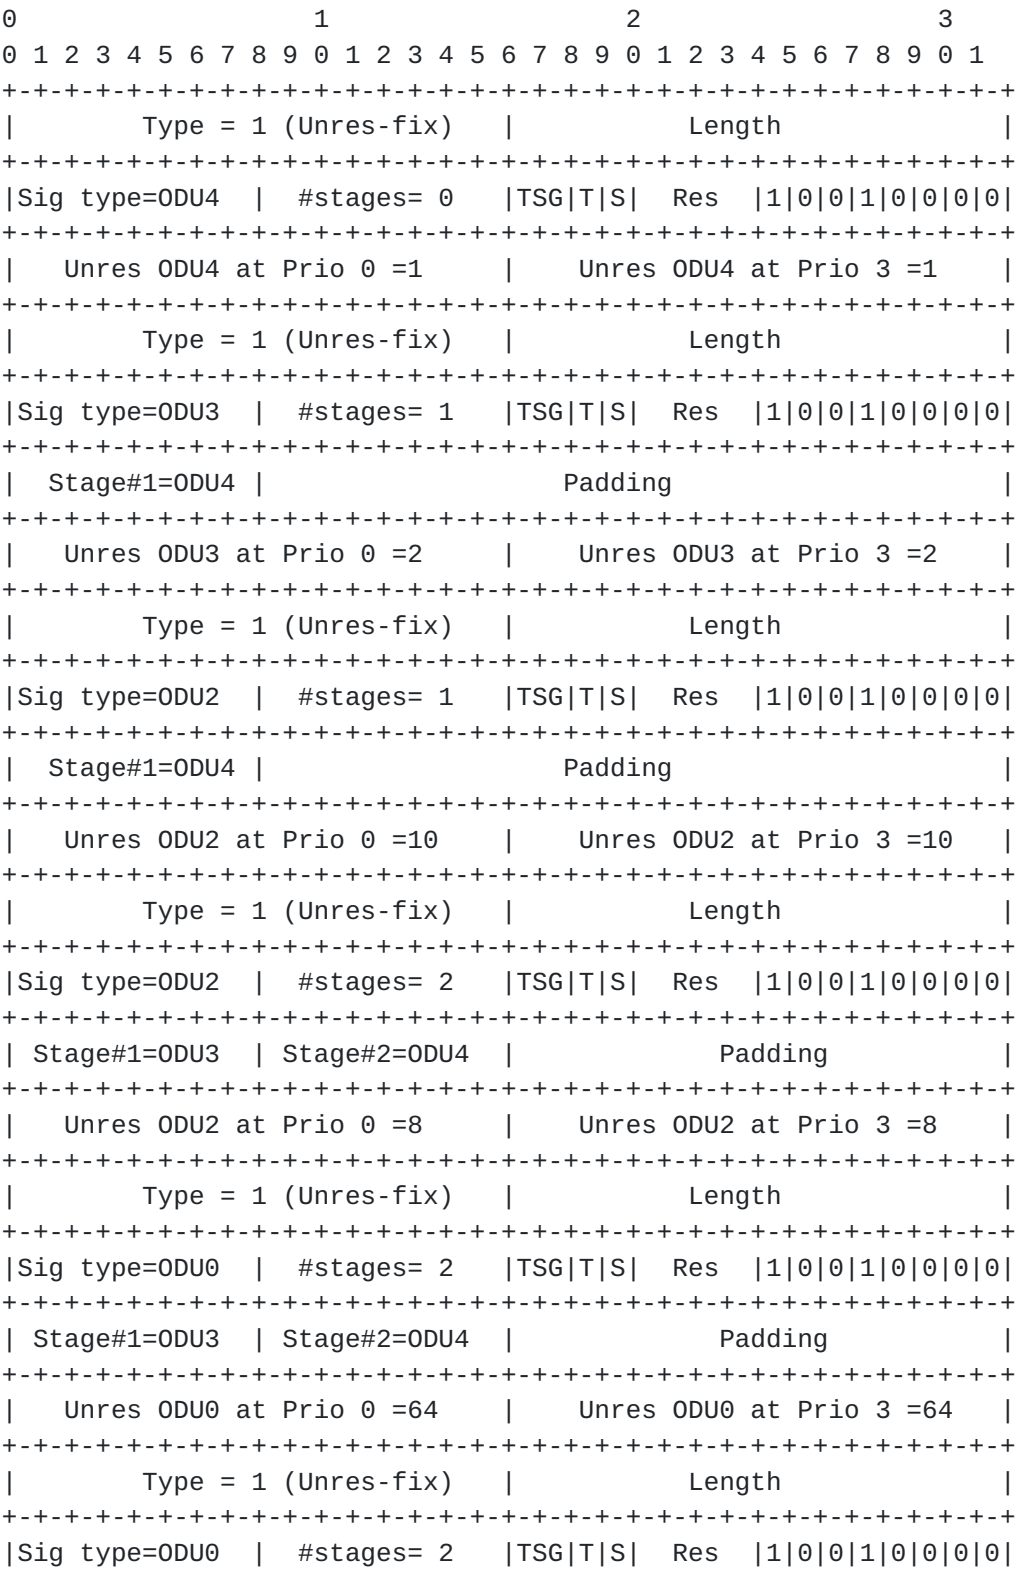

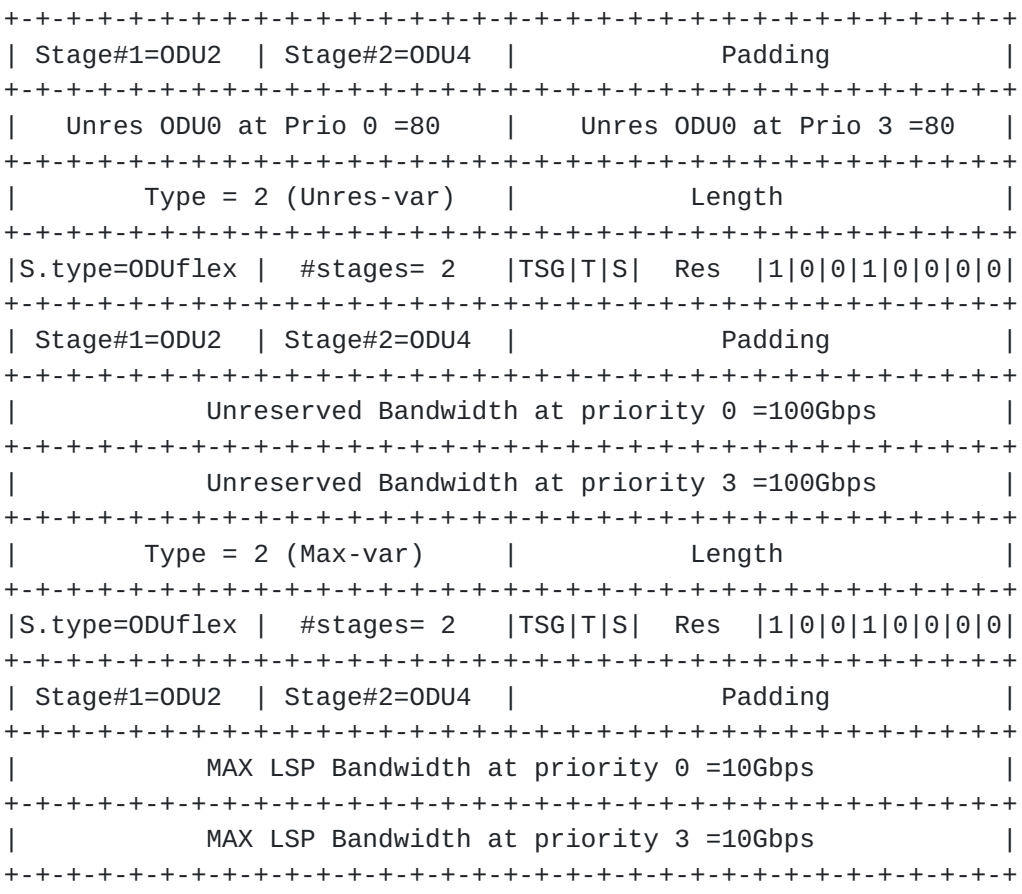

Figure 13: Example 5 - Multi stage muxing - Unbundled link

5.6. Example of multi stage muxing - Bundled links

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

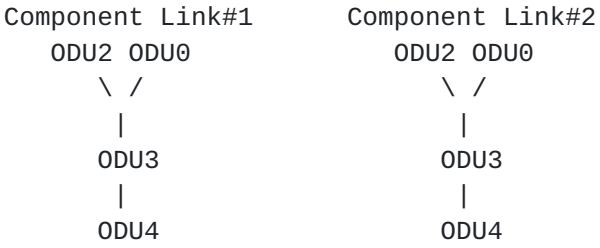

Considering only supported priorities 0 and 3, the advertisement is as follows:

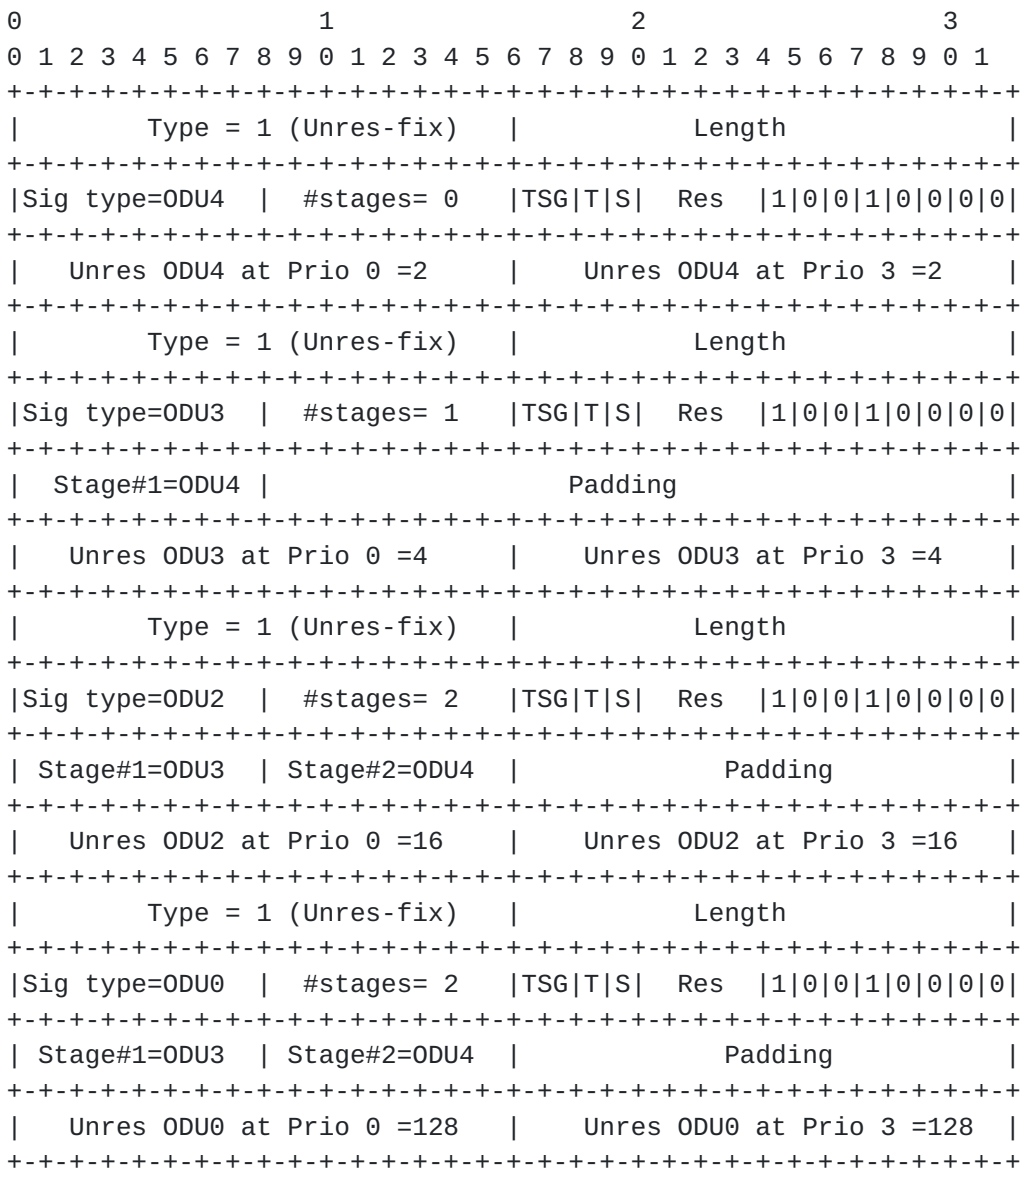

Figure 14: Example 6 - Multi stage muxing - Bundled lilnks

6. Compatibility

Backwards compatibility with implementations based on [RFC4328] can be achieved advertising the [RFC4328] based ISCDs in addition to the ISCD defined in this document.

7. Security Considerations

This document 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.

8. IANA Considerations

TBD

9. Contributors

Xiaobing Zi, Huawei Technologies

Email: zixiaobing@huawei.com

Francesco Fondelli, Ericsson

Email: francesco.fondelli@ericsson.com

Marco Corsi, Altran Italia

EMail: marco.corsi@altran.it

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

akunjidhapatham@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

10. Acknowledgements

11. References

11.1. Normative 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-04", March 2011.

[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-00", April 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC2370] Coltun, R., "The OSPF Opaque LSA Option", RFC 2370, July 1998.

[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.

[RFC4202] Kompella, K. and Y. Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.

[RFC4203] Kompella, K. and Y. Rekhter, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF

for IPv6", RFC 5340, July 2008.

[RFC6001] Papadimitriou, D., Vigoureux, M., Shiomoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, October 2010.

11.2. Informative References

[G.709] ITU-T, "Interface for the Optical Transport Network (OTN)", G.709 Recommendation (and Amendment 1), February 2001.

[G.709-v3] ITU-T, "Draft revised G.709, version 3", consented by ITU-T on Oct 2009.

[Gsup43] ITU-T, "Proposed revision of G.sup43 (for agreement)", December 2008.

Authors' Addresses

Daniele Ceccarelli (editor)
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: daniele.ceccarelli@ericsson.com

Diego Caviglia
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: diego.caviglia@ericsson.com

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

Dan Li
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Shenzhen 518129 P.R.China Bantian, Longgang District
Phone: +86-755-28973237

Email: danli@huawei.com

Sergio Belotti
Alcatel-Lucent
Via Trento, 30
Vimercate
Italy

Email: sergio.belotti@alcatel-lucent.com

Pietro Vittorio Grandi
Alcatel-Lucent
Via Trento, 30
Vimercate
Italy

Email: pietro_vittorio.grandi@alcatel-lucent.com

Rajan Rao
Infinera Corporation
169, Java Drive
Sunnyvale, CA-94089
USA

Email: rrao@infinera.com

Khuzema Pithewan
Infinera Corporation
169, Java Drive
Sunnyvale, CA-94089
USA

Email: kpithewan@infinera.com

John E Drake
Juniper

Email: jdrake@juniper.net