CCAMP Working Group
Internet-Draft
Intended status: Standards Track

Z. Fan
Huawei Technologies
R. Valiveti
I. Hussain
Infinera
Q. Wang
ZTE
Z. Ali
Cisco

Expires: April 30, 2018 October 30, 2017

OSPF Extensions for the GMPLS Control of OTN B100G Network

draft-merge-ccamp-otn-b100g-routing-ext-00

Abstract

ODUCn signal is recently introduced to OTN to support B100G feature. This document provides the OSPF extensions to control the OTN B100G Network.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on April XX, 2018.

Copyright Notice

Copyright (c) 2017 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

G.709 edition 5 [G709-2016] introduces ODUCn signal to support beyond 100G data rate. ODUCn signal, as a HO ODU, can carry OTN signals such as ODUk and ODUflex. The tributary slot granularity of ODUCn is 5 Gbps. The OSPF-TE extensions defined in [RFC7138] cannot support the OTN B100G features.

B100G framework document [I-D.merge-ccamp-otn-b100g-fwk] provides the requirements of protocol extensions to support the GMPLS control of OTN B100G. This document provides OSPF-TE extensions to support the control of ODUCn.

Note: This document considers routing information for OTN electrical layer only. Routing information for OTN optical layer (i.e., OCh, OTSiA, and FlexO interfaces) is beyond the scope of this document.

1.1. Requirements Language

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 [RFC2119].

2. Terminology

OPUCn: Optical Payload Unit-Cn

ODUCn: Optical Data Unit-Cn

OTUCn: completely standardized Optical Transport Unit-Cn

OTUCn-M: Optical Transport Unit-Cn with n OxUC overhead instances and M 5G tributary slots

TS: Tributary Slot

TSG: Tributary Slot Granularity

3. Overview of OSPF-TE Extensions for Support ODUCn

As described in [I-D.merge-ccamp-otn-b100g-fwk], OSPF-TE should be extended to advertise the 5G tributary slot granularity, the multiplexing capabilities of ODUCn, and the available bandwidth information of ODUCn.

The advertisement of ODUCn information is used to synchronize the two end nodes of an ODUCn link. If the two ends have different tributary slot granularities, this ODUCn link should not be setup. If the two ends have different multiplexing hierarchies for ODUCn, the supported ODUk multiplexing should be the ODUk supported by both ends. If the two ends mark different tributary slots as unavailable, each end node should calculate the actual available TS (i.e., the intersection of available TS from two ends), and convert the actual available bandwidth to equivalent available ODUk bandwidth.

4. ISCD Format Extensions

As defined in [RFC4203], ISCD is used to describe the switching capability. Although ODUCn is not switchable, as discussed in Section 3, we still need advertise some capabilities to the other end of the ODUCn link. We re-use the OTN-TDM switching capability

defined in [RFC7138]. A new LSP encoding type is defined for ODUCn in [I-D.merge-ccamp-b100g-signaling].

4.1. Switching Capability Specific Information

Besides ODUCn signal, [G709-2016] also introduces ODUflex for FlexE-aware signal and ODUflex with IMP. Three new signal type need to be defined:

- TBA1 - ODUCn
- TBA2 - ODUflex (IMP)
- TBA3 - ODUflex (FlexE-aware)

The Bandwidth sub-TLV defined in [RFC7138] contains two types. As ODUCn is a HO ODU, the multiplexing hierarchy is affected to have more stages. Type 1 Bandwidth sub-TLV need to be modified, and a new type Bandwidth sub-TLV is needed for ODUCn.

4.1.1. Modification of Type 1 Container

The multiplexing hierarchy is represented by stages in [RFC7138]. As ODUk can be multiplexed into ODUCn, one more multiplexing stage can be introduced in both type 1 (fixed container) and type 2 (flexible container) Bandwidth sub-TLV. The extreme case for type 1 is that ODU0->ODU1->ODU2->ODU3->ODU4->ODUCn, which contains 5 stages. The original one-row space for stage field could be insufficient. Therefore, the Stage field needs to be modified to support multiplexing to ODUCn. The modified format of type 1 Bandwidth sub-TLV is depicted in the following figure:

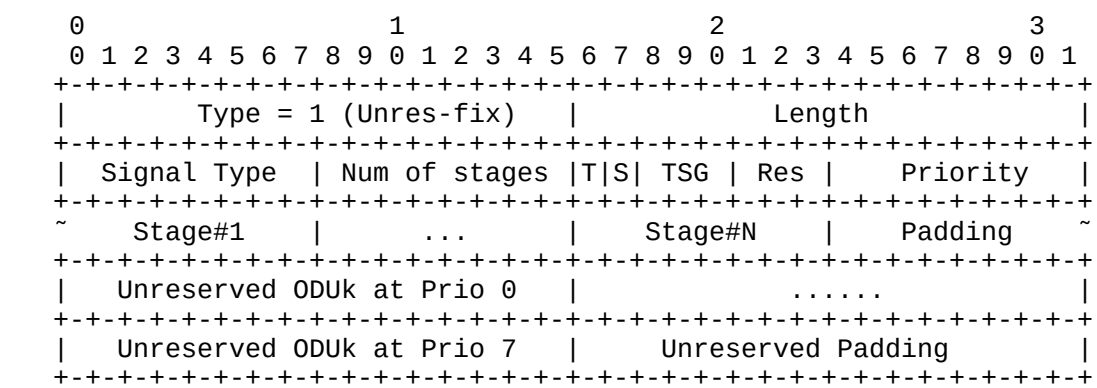

Figure 1: Modified Bandwidth sub-TLV for Type 1 containers

4.1.2. Type 3 Container for advertising Unreserved ODUCn

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

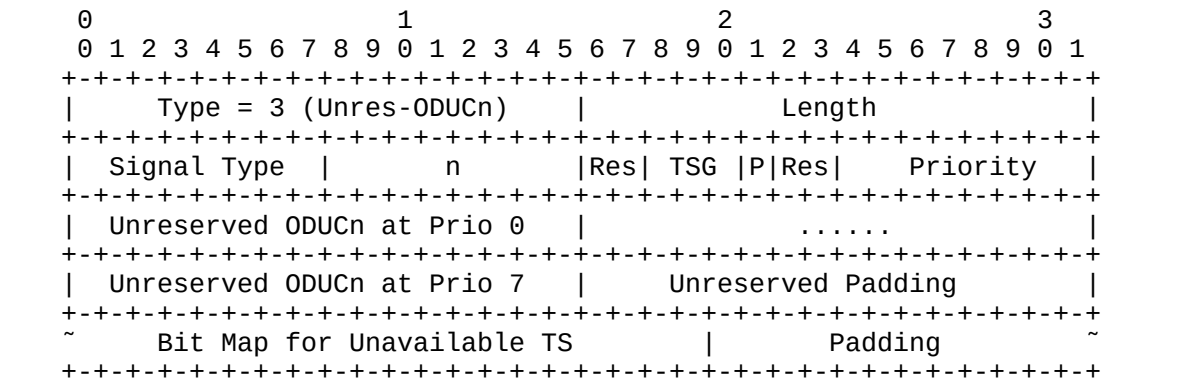

Figure 2: Extended Bandwidth sub-TLV for Type 3 containers

o Signal Type (8 bits): Same as the definition in [RFC7138]. The value can only be ODUCn signal.

o n (8 bits): Indicates the number of ODUC instance in an ODUCn signal.

o Flags (8 bits):

  * P Flag (bit 22): Indicates whether the advertised ODUCn link is mapped to sub-rate OTUCn-M, which means some TS in this link are marked as unavailable. When ODUCn contains unavailable TS, P MUST be set, while when ODUCn does not contain unavailable TS, P MUST be cleard.

o TSG (3 bits): Inherits the definition in [RFC7138] by adding a new value indicating the 5 Gbps TSG:

  * 4 - 5 Gbps only

Priority (8 bits): Same as the definition in [RFC7138].

Unreserved ODUCn (16 bits): Indicates the Unreserved Bandwidth at a particular priority level. This field MUST be set to the number of the specific ODUCn, which is identified by the Signal Type field, the n field, and the Bit Map for Unavailable TS field, 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.

Unreserved Padding (16 bits): Same as the definition in [RFC7138].

Bit Map for Unavailable TS (variable): Indicates which tributary slots are marked as unavailable due to the bandwidth limitation from lower layer connection, which is different from occupied/allocated TS. The total number of unavailable TS can be calculated by summing this field. The length of this field is derived from the n field (the length is 20 x n). The sequence of this field follows the joint sequence of the tributary slots in the ODUCn and the order of ODUC instances. The first 20 bits are respectively for ODUC#1, the second 20 bits are respectively for ODUC#2, and so on. Each bit in the bit map represents the corresponding tributary slot in the ODUCn with a value of 1 or 0 indicating whether the tributary slot is marked as unavailable or not. When P bit is cleared, the Bit Map field is not required and MUST NOT be included.

Padding (variable): Are added after the Bit Map field to make the whole label a multiple of four bytes if necessary. Padding bits MUST be set to 0 and MUST be ignored on receipt.

5. Examples

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

5.1. Multiplexing ODUk over ODUCn

This example shows the advertisement of the ISCD for ODUCn. An OTUC2 link is considered with supported priorities 0,3 and multiplexing hierarchy ODU4->ODUC2.

The format of the advertised ISCD is depicted by the following figure:

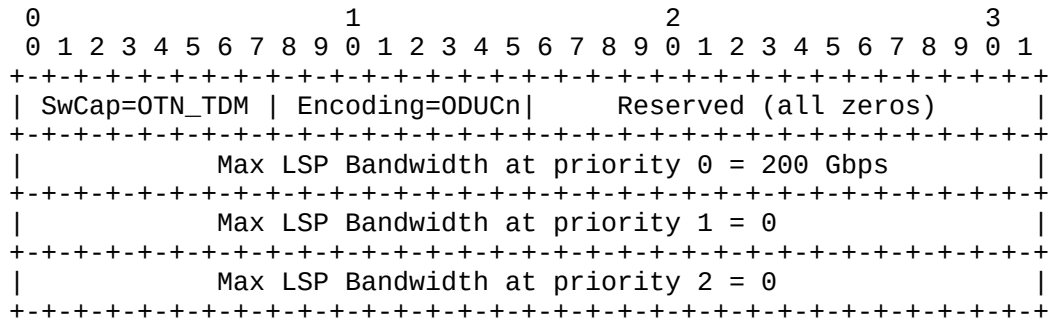

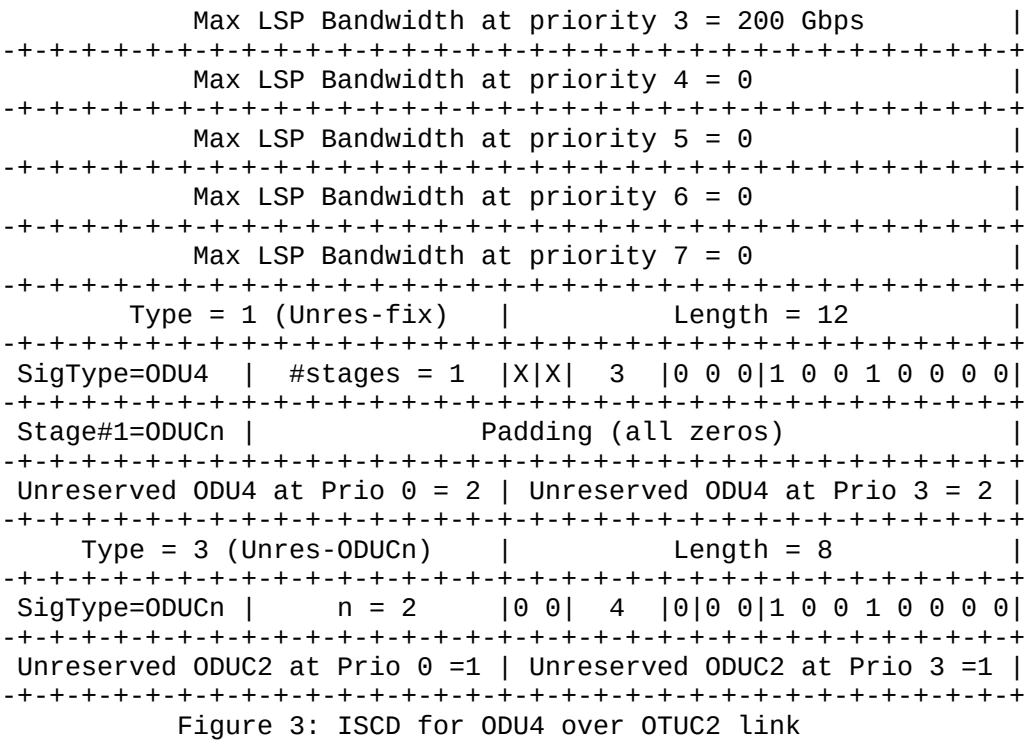

Figure 3: ISCD for ODU4 over OTUC2 link

The Max LSP Bandwidth is filled with the bandwidth of ODUC2 (i.e., 200 Gbps).

According to the multiplexing hierarchy, the advertised ODU4 has one stage to ODUCn. The number of unreserved ODU4 is 2 in this example.

The advertised ODUC2 has signal type as ODUCn, n as 2, and P bit cleared. The TSG value is 4, which means 5 Gbps granularity. The number of unreserved ODUC2 is 1 in this example.

5.2. Advertising Unavailable TS Information of ODUCn

This example shows the advertisement of unavailable TS information. An OTUC2-30 link is considered with supported priorities 0,3 and multiplexing hierarchy ODU4->ODUC2.

The format of the advertised ISCD is depicted by the following figure:

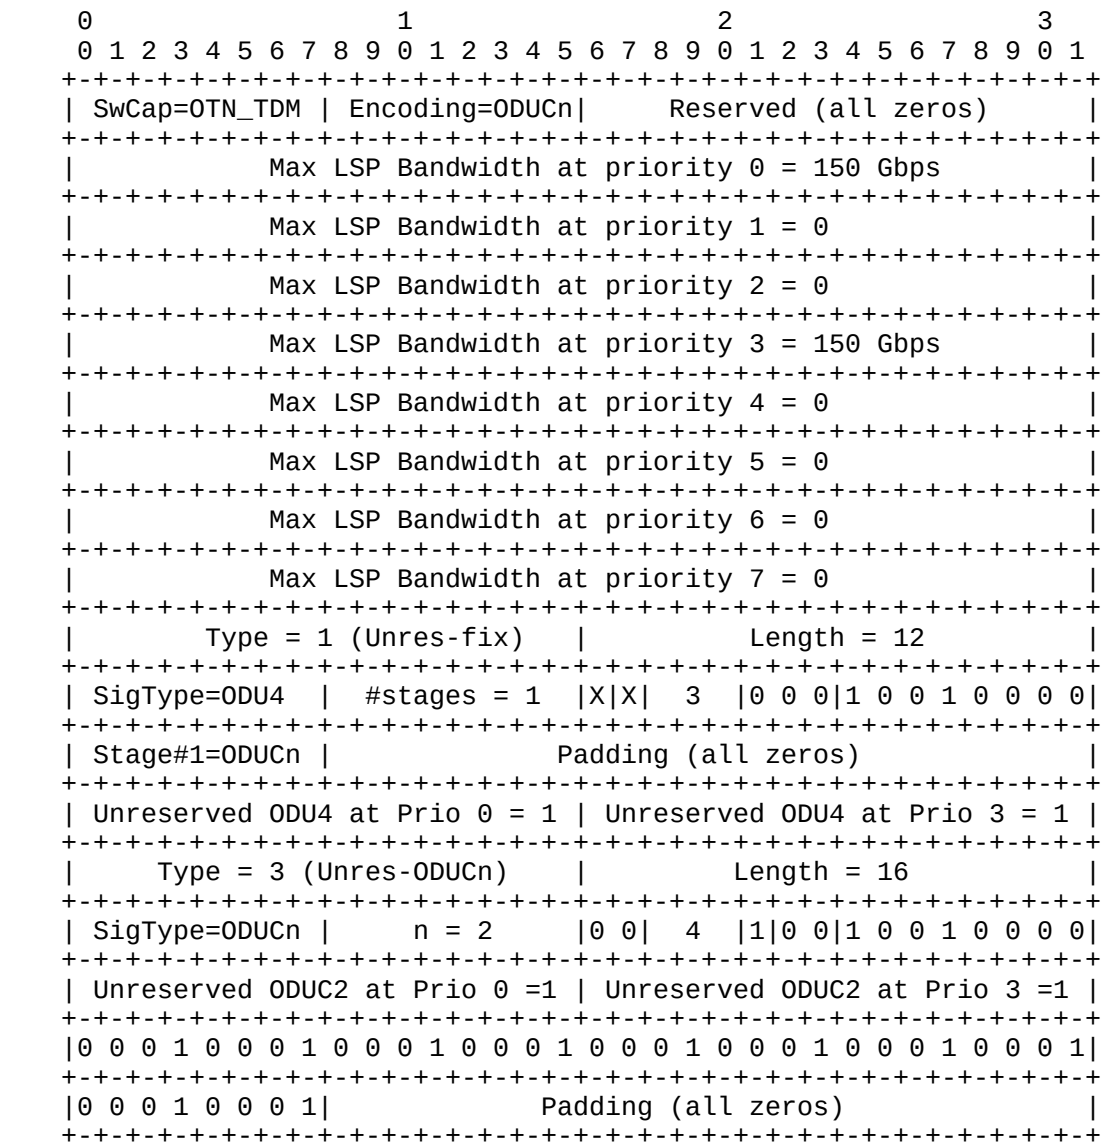

Figure 4: ISCD for ODU4 over OTUC2-30 link

The Max LSP Bandwidth is filled with 150 Gbps, as ODUC2 has 10 unavailable tributary slots.

As the bandwidth of ODUC2 is reduced, the number of unreserved ODU4 is 1 in this example.

The advertised ODUC2 has signal type as ODUCn, n as 2, and P bit set. The TSG value is 4, which means 5 Gbps granularity. The number of unreserved ODUC2 is 1 in this example. The Bit Map field indicates which tributary slot is marked as unavailable, where the marking policy is vendor specific. In this example, bit-4, bit-8, bit-12, bit-16, bit-20, bit-24, bit-28, bit-32, bit-36, and bit-40 are set, which means the corresponding tributary slots are marked as unavailable.

6. Security Considerations

TBD.

7. IANA considerations

TBD.

8. Contributors' Addresses

Haomian Zheng
Huawei Technologies

Email: zhenghaomian@huawei.com

Sergio Belotti
Nokia

Email: sergio.belotti@nokia.com

Yunbin Xu
CAICT

Email: xuyunbin@ritt.cn

Rajan Rao
Infinera

Email: rrao@infinera.com

Huub van Helvoort
Hai Gaoming B.V

Email: huubatwork@gmail.com

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.

[G709-2016] ITU-T, "Interface for the Optical Transport Network (OTN)", G.709/Y.1331 Recommendation, June 2016.

[RFC7138] Ceccarelli D., Zhang, F., Belotti, S., Rao, R., and J. Drake, "Traffic Engineering Extensions to OSPF for GMPLS Control of Evolving G.709 Optical Transport Networks", RFC7138, March 2014.

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

9.2. Informative References

[I-D.merge-ccamp-otn-b100g-fwk] Wang, Q., Ed., Valiveti, R., Ed., Zheng, H., Ed., Helvoort, H., and S. Belotti, "GMPLS Routing and Signaling Framework for B100G", draft-merge-ccamp-otn-b100g-fwk-02 (work in process), July 2017.

[I-D.merge-ccamp-b100g-signaling] Wang, Q., Ed., Zheng, H., Valiveti, R., Helvoort, H., and Z. Ali, " GMPLS Signalling Extensions for control of B100G OTUCn/ODUCn Network ", draft-merge-ccamp-100g-signalling-00 (work in process), October 2017.

Authors’ Addresses

Zheyu Fan
Huawei Technologies

Email: fanzheyu2@huawei.com

Radhakrishna Valiveti
Infinera

Email: rvaliveti@infinera.com

Iftekhar Hussain
Infinera

Email: IHussain@infinera.com

Qilei Wang
ZTE

Email: wang.qilei@zte.com.cn

Zafar Ali
Cisco

Email: zali@cisco.com