Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: May 3, 2018

Q. Wang, Ed.
ZTE
R. Valiveti, Ed.
Infinera Corp
H. Zheng, Ed.
Huawei
H. Helvoort
Hai Gaoming B.V
S. Belotti
Nokia

October 30, 2017

GMPLS Routing and Signaling Framework for B100G
draft-merge-ccamp-otn-b100g-fwk-02

Abstract

The 2016 revision of G.709 introduces support for OTU links with rates larger than 100G. This document provides a framework to address the GMPLS routing and signalling extensions that enable GMPLS to setup paths through network that contain these newly introduced OTUCn links.

Table of Contents

1. Introduction

The current GMPLS routing [RFC7138] and signaling extensions [RFC7139] includes coverage for all the OTN capabilities that were defined in the 2012 version of G.709 [ITU-T_G709_2012].

The 2016 version of G.709 [ITU-T_G709_2012] introduces support for higher rate OTU signals, termed OTUCn (which have a nominal rate of n x 100 Gbps). The newly introduced OTUCn represent a very powerful extension to the OTN capabilities, and one which naturally scales to transport any newer clients with bit rates in excess of 100G, as they are introduced.

This document presents an overview of the changes introduced in [ITU-T_G709_2016] and analyzes them to identify the extensions that would be required in GMPLS routing and signaling to enable the new OTN capabilities.

1.1. Scope

For the purposes of the B100G control plane discussion, the OTN should be considered as a combination of ODU and OTSi layers. Note that [ITU-T_G709_2016] is deprecating the use of the term "Och" for B100G entities, and leaving it intact only for maintaining continuity in the description of the signals with bandwidth upto 100G. This document focuses on only the control of the ODU layer. The control of the OTSi layer is out of scope of this document.

2. Terminology

2.1. 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 RFC 2119 [RFC2119].

2.2. OTN terminology used in this document

a. OPUCn: Optical Payload Unit -Cn.

b. ODUCn: Optical Data Unit - Cn.

c. OTUCn: Fully standardized Optical Transport Unit - Cn.

d. OTUCn-M: This signal is an extension of the OTUCn signal introduced above. This signal contains the same amount of overhead as the OTUCn signal, but contains a reduced amount of payload area. Specifically the payload area consists of M 5G tributary slots (where M is strictly less than 20*n).

e. PSI: OPU Payload structure Indicator. This is a multi-frame message and describes the composition of the OPU signal. This field is a concatenation of the Payload type (PT) and the Multiplex Structrure Indicator (MSI) defined below.

f. MSI: Multiplex Structure Indicator. This structure indicates the grouping of the tributary slots in an OPU payload area to realize a client signal that is multiplexed into an OPU. The individual clients multiplexed into the OPU payload area are distinguished by the Tributary Port number (TPN).

g. GMP: Generic Mapping Procedure.

Detailed description of these terms can be found in [ITU-T_G709_2016].

3. Overview of B100G in G.709

This section provides an overview of new features in [ITU-T_G709_2016].

3.1. OTUCn

In G.709 [ITU-T_G709_2012], the standard mechanism for transporting a client signal is to first map it into an ODU signal (of the appropriate rate), and then switch the resulting ODU signal through the OTN network. In the course of its traversal through the OTN network, the ODU signal generated by the mapper is either (a) multiplexed into higher-order ODU, and then encapsulated to form an OTU or (b) directly encapsulated into an OTU signal that defines the section layer. The option (b), i.e. direct encapsulation into an OTU was possible only for ODU1/ODU2/ODU3/ODU4; ODU signals with other rates (e.g. ODUflex) would first have to be processed per option (a) above. The term "client signal" is generic in the sense that it encompasses both Constant Bit rate (CBR) clients (e.g. 10GBASE-R, SONET OC-768), or packet traffic -- where the goal is to transfer the payload from end-to-end (without regard for bit transparency at the PCS layer). Given that OTU4 was the highest rate section layer signal supported in [ITU-T_G709_2012], the client signal rates were limited to be less than 100G (if ODU-VCAT was not used).

In order to carry client signals with rates greater than 100Gbps, [ITU-T_G709_2016] takes a general and scalable approach that decouples the rates of OTU signals from the client rate evolution.

The new OTU signal is called OTUCn; this signal is defined to have a rate of (approximately) n*100G. The following are the key characteristics of the OTUCn signal:

a. The OTUCn signal contains one ODUCn, which in turn contains one OPUCn signal. The OTUCn and ODUCn signals perform digital section roles only (see [ITU-T_G709_2016]:Section 6.1.1). The OTUCn and ODUCn can be seen as being analogous to the regenerator section, and multiplex section in SDH respectively.

b. The OTUCn signals can be viewed as being formed by interleaving n OTUC signals (where are labeled 1, 2, ..., n), each of which has the format of a standard OTUk signal without the FEC columns (per [ITU-T_G709_2016]:Figure 7-1). The ODUCn, and OPUCn have a similar structure, i.e. they can be seen as being formed by interleaving n instances of ODUC, OPUC signals (respectively) The OTUC signal contains the ODUC, and OPUC signals, just as in the case of fixed rate OTUs defined in G.709 [ITU-T_G709_2016].

c. Each of the OTUC "slices" have the same overhead (OH) as the standard OTUk signal in G.709 [ITU-T_G709_2016]. The combined signal OTUCn has n instances of OTUC OH, ODUC OH, and OPUC OH.

d. The OTUC signal has a slightly higher rate compared to the OTU4 signal (without FEC); this is to ensure that the OPUC payload area can carry an ODU4 signal.

3.1.1. Carrying OTUCn between 3R points

As explained above, within G.709 [ITU-T_G709_2016], the OTUCn, ODUCn and OPUCn signal structures are presented in a (physical) interface independent manner, by means of n OTUC, ODUC and OPUC instances that are marked #1 to #n. Specifically, the definition of the OTUCn signal does not cover aspects such as FEC, modulation formats, etc. These details are defined as part of the adaptation of the OTUCn layer to the optical layer(s). The specific interleaving of OTUC/ODUC/OPUC signals onto the optical signals is interface specific and specified for OTN interfaces with standardized application codes in the interface specific recommendations (G.709.x).

The following scenarios of OTUCn transport need to be considered (see Figure 1):

a. inter-domain interfaces: These types of interfaces are used for connecting OTN edge nodes to (a) client equipment (e.g. routers) or (b) hand-off points from other OTN networks. ITU-T has standardized the Flexible OTN (FlexO) interfaces to support these functions. Recommendation [ITU-T_G709.1] specifies a flexible interoperable short-reach OTN interface over which an OTUCn (n >=1) is transferred, using bonded FlexO interfaces which belong to a FlexO group. The FlexO group supports physical interface bonding, management of the group members, overhead for communication between FlexO peers etc. (these overheads are separate from the GCC0 channel defined over OTUCn). In its current form, Recommendation [ITU-T_G709.1] is limited to the case of transporting OTUCn signals using n 100G Ethernet PHY(s). The mechanisms for transporting the OTUCn signals over 100G optical interfaces are specified in [ITU-T_G709.1] and are not repeated here. When the PHY(s) for the emerging set of Ethernet signals, e.g. 200GbE and 400GbE, become available, new recommendations can define the required adaptations.

b. intra-domain interfaces: In these cases, the OTUCn is transported using a proprietary (vendor specific) encapsulation, FEC etc. In future, it may be possible to transport OTUCn for intra-domain links using future variants of FlexO.

```
==================================================================


      +---------------------------------------------------------+
      |                  OTUCn signal                           |
      +---------------------------------------------------------+
      |  Inter+Domain     |  Intra+Domain     |  Intra+Domain    |
      |  Interface (IrDI)|  Interface (IaDI)|  Interface        |
      |  FlexO (G.709.1) |  FlexO (G.709.x) |  Proprietary      |
      |                   |  (Future)         |  Encap, FEC etc. |
      +---------------------------------------------------------+


==================================================================
```

Figure 1: OTUCn transport possibilities

It is possible for an OTUCn signal to be transported via multiple hops of lower-layer adaptation (see Figure 2). In this scenario, the OTUCn spans multiple optical paths joined by a FlexO segment. An end-to-end OTUCn LSP needs to be setup after the optical circuits are established. The information about the FlexO interfaces (and group) are configured at the FlexO endpoints, and there is no dynamic setup.

```
  ==================================================================

         Optical Segment                    Optical Segment
      <---------------->                 <---------------->

   +------+         +------+          +------+         +------+
   |      |         |      +-----------+      |         |      |
   |  A   +---------+  B   |-----------|  C   +---------+  D   |
   |      |         |      +-----------+      |         |      |
   +------+         +------+     ^    +------+         +------+
                                 |
                                 +
                             FlexO Group
      <----------------------------------------------------->  OTUCn

           +-------+          +-------+           +-------+
           | OTUCn |          | OTUCn |           | OTUCn |
           +-------+          +-------+           +-------+
           | OTSiG |          | FlexO |           | OTSiG |
           +-------+          +-------+           +-------+

  ==================================================================
```

Figure 2: OTUCn transport - Multihop

This document views FlexO (even if there are some digital sub-layers involved in the adaptation) and other OTUCn transport mechanisms as "lower layers", and are therefore considered out-of-scope. The OTUCn layer operates independent of the method used to transport the signal.

3.2. ODUCn

The ODUCn signal [ITU-T_G709_2016] can be viewed as being formed by the appropriate interleaving of content from n ODUC signal instances. The ODUC frames have the same structure as a standard ODU -- in the sense that it has the same Overhead (OH) area, and the payload area -- but has a higher rate since its payload area can embed an ODU4 signal. The ODUCn signal can be formed in one of the following ways:

By multiplexing lower-rate (i.e. both low-order and high-order) ODUk signals.

Each of the n instances of ODUC can carry the NULL signal (as specified in [ITU-T_G709_2016]: Section 17.5.1)

Each of the n instances of ODUC can carry the PN-11 PRBS test sequence (as specified in [ITU-T_G709_2016]: Section 17.5.2)

It is concievable that vendors might implement proprietary mappings (Payload Type values of 0x80-x8F)of non-OTN client signals. An interoperable control plane cannot make use of these proprietary ODUCn signals, and hence this case isn't considered in this document.

The ODUCn signals have a rate that is captured in Table 1.

ODU Type	ODU Bit Rate
ODUCn	n x 239/226 x 99,532,800 kbit/s = n x 105,258,138.053 kbit/s

Table 1: ODUCn rates

The ODUCn is a multiplex section ODU signal, and is mapped into an OTUCn signal which provides the regenerator section layer. In some scenarios, the ODUCn, and OTUCn signals will be co-terminous, i.e. they will have identical source/sink locations. [ITU-T_G709_2016] and [ITU-T_G872] allow for the ODUCn signal to pass through a digital regenerator node which will terminate the OTUCn layer, but will pass the regenerated (but otherwise untouched) ODUCn towards a different OTUCn interface where a fresh OTUCn layer will be initiated (see Figure 3). In this case, an ODUCn LSP needs to be set up to traverse the 3 OTUCn segments.

Specifically, the OPUCn signal flows through these regenerators unchanged. That is, the set of client signals, their TPNs, trib-slot allocation remains unchanged. Note however that the ODUCn Overhead (OH) might be modified if TCM sub-layers are instantiated in order to monitor the performance of the repeater hops. In this sense, the ODUCn should not be seen as a general ODU which can be switched via an ODUk cross-connect.

```
   ==================================================================

                ^ FlexO group                                  ^ FlexO group
               |                                              |
+--------+     |     +--------+                +--------+     |    +--------+
|        +-----------+        |                |        +----------+        |
| OTN    |-----------| OTN    |                | OTN    |----------| OTN    |
| DXC    +-----------+ WXC    +----------------+ WXC    +----------+ DXC    |
|        |           | 3R     |                | 3R     |          |        |
+--------+           +--------+                +--------+          +--------+

     <------------------>  <----------------------> <------------------>
           OTUCn/FlexO            OTUCn/OTSiG            OTUCn/FlexO

     <-----------------------------------------------------------------> ODUCn

   ==================================================================
```

Figure 3: Multi-hop ODUCn signal

3.3. OTUCn-M

The standard OTUCn signal has the same rate as that of the ODUCn signal as captured in Table 1. This implies that the OTUCn signal can only be transported over wavelength groups which have a total capacity of multiples of (approximately) 100G. Modern DSPs support a variety of bit rates per wavelength, depending on the reach requirements for the optical link. With this in mind, ITU-T supports the notion of a reduced rate OTUCn signal, termed the OTUCn-M. The OTUCn-M signal is derived from the OTUCn signal by retaining all the n instances of overhead (one per OTUC slice) and crunching the OPUC tributary slots marked as "unavailable".

3.4. OPUCn Time Slot Granularity

[ITU-T_G709_2012] introduced the support for 1.25G granular tributary slots in OPU2, OPU3, and OPU4 signals. With the introduction of higher rate signals such as the OPUCn, it is no longer practical for the optical networks (and the datapath hardware) to support a very large number of flows at such a fine granularity. ITU-T has defined the OPUC with a tributary slot granularity of 5G. This means that the ODUCn signal has 20*n tributary slots (of 5Gbps capacity).

3.5. Structure of OPUCn MSI with Payload type 0x22

As mentioned above, the OPUCn signal has 20*n 5G tributary slots. The OPUCn contains n PSI structures, one per OPUC instance. The PSI structure consists of the Payload Type (of 0x22), followed by a Reserved Field (1 byte), followed by the MSI. The OPUCn MSI field has a fixed length of 40*n bytes and indicates the ODTU content of each TS of an OPUCn. Two bytes are used for each of the 20*n tributary slots, and each such information structure has the following format ([ITU-T_G709_2016] G.709:Section 20.4.1):

a. The TS availability bit 1 indicates if the tributary slot is available or unavailable

b. The TS occupation bit 9 indicates if the tributary slot is allocated or unallocated

3.6. Client Signal Mappings

Note that [ITU-T_G709_2016] introduces support for OTUCn signals with rates of n*100G and also introduces support for client signals with rates larger than 100G (e.g. the future 400GBASE-R client being standardized by IEEE, higher packet streams from NPUs). The approach taken by the ITU-T to map non-OTN client signals to the appropriate ODU containers is as follows:

a. All client signals with rates less than 100G are mapped as specified in [ITU-T_G709_2016]:Clause 17. These mappings are identical to those specified in the earlier revision of G.709 [ITU-T_G709_2012]. Thus, for example, the 1000BASE-X/10GBASE-R signals are mapped to ODU0/ODU2e respectively (see Table 2 -- based on Table 7-2 in [ITU-T_G709_2016])

b. Always map the new and emerging client signals to ODUflex signals of the appropriate rates (see Table 2 -- based on Table 7-2 in [ITU-T_G709_2016])

c. Drop support for ODU Virtual Concatenation. This simplifies the network, and the supporting hardware since multiple different mappings for the same client are no longer necessary. Note that legacy implementations that transported sub-100G clients using ODU VCAT shall continue to be supported.

d. ODUflex signals are low-order signals only. If the ODUflex entities have rates of 100G or less, they can be transported using either an ODUk (k=1..4) or an ODUCn server layer. On the other hand, ODUflex connections with rates greater than 100G will require the server layer to be ODUCn. The ODUCn signals must be adapted to an OTUCn signal. Figure 4 illstrates the hierarchy of the digital signals defined in [ITU-T_G709_2016].

ODU Type	ODU Bit Rate
ODU0	1,244,160 Kbps
ODU1	239/238 x 2,488,320 Kbps
ODU2	239/237 x 9,953,280 Kbps
ODU2e	239/237 x 10,312,500 Kbps
ODU3	239/236 x 39,813,120 Kbps
ODU4	239/227 x 99,532,800 Kbps
ODUflex for CBR client signals	239/238 x Client signal Bit rate
ODUflex for GFP-F mapped packet traffic	Configured bit rate
ODUflex for IMP mapped packet traffic	s x 239/238 x 5 156 250 kbit/s: s=2,8,5*n, n >= 1
ODUflex for FlexE aware transport	103 125 000 x 240/238 x n/20 kbit/s, where n is total number of available tributary slots among all PHYs which have been crunched and combined.

Note that this table doesn't include ODUCn -- since it cannot be generated by mapping a non-OTN signal. An ODUCn is always formed by multiplexing multiple LO-ODUs.

Table 2: Types and rates of ODUs usable for client mappings

```
    ==================================================================


                     Clients (e.g. SONET/SDH, Ethernet)
                          +       +      +
                          |       |      |
       +------------------+-------+------+-----------------------+
       |                     OPUk                                |
       +---------------------------------------------------------+
       |                     ODUk                                |
       +-----------------------+---------------------------+------+
       | OTUk, OTUk.V, OTUkV   |          OPUk             |      |
       +----------+--------------------------------------+       |
       | OTLk.n   |            |          ODUk             |      |
       +----------+            +---------------------+-----+      |
                               | OTUk, OTUk.V, OTUkV |   OPUCn    |
                               +----------+-----------------------+
                               | OTLk.n   |          |   ODUCn    |
                               +----------+          +------------+
                                                     |   OTUCn    |
                                                     +------------+


    ==================================================================
```

Figure 4: Digital Structure of OTN interfaces (from G.709:Figure 6-1)

4. Usecases

This section introduces various usecases that provide the rationale for the requirements that any solution must satisfy. At a later point in time, it is possible to consolidate these usecases so that all the multiplexing (and demultiplexing) variants are encountered along the path of an end-to-end ODU circuit.

Note-1: These usecases present scenarios in which OTUCn links are depicted. These illustrations do not highlight how the OTUCn is transported between the 3R points. That is, these usecases do not cover cases in which a standard FlexO interface (e.g. as defined in [ITU-T_G709.1]) is used, or whether a vendor specific mapping of OTUCn to OTSiG (as defined in [ITU-T_G872]) is used. In other words, multiple variants of these usecases based on FlexO usage (or not) are not included in this document.

4.1. 100GE Client Service with a homogeneous chain of OTUC1 links

In the scenario illustrated in Figure 5 a 100GBASE-R client is mapped into an ODU4 at NE1. The resulting ODU4 signal is multiplexed into the ODUC1 server layer (using GMP) and further encapsulated to form the OTUC1 signal. The links NE1-NE2, and NE2-NE3 are both OTUC1 links -- and they can carry one 100GE client mapped into an ODU4 server layer. Actions performed at NE2 are: (a) terminate OTUC1, and ODUC1 towards NE1 (b) demultiplex the ODU4 signal from ODUC1 (c) map the ODU4 signal onto a different ODUC1/OTUC1 towards NE3. NE3 performs the inverse sequence of steps performed at NE1, and recovers the 100GBASE-R client from the ODU4 signal. Note that the ODU4 and ODUC1 signals are not "interoperable" and that the ODUC1 is a server layer to the ODU4 signal.

This illustration is also applicable to the usecase in which members of a FlexE group are transported in a flexe-unaware mode in the transport network. Although this illustration included only OTUC1 signals, any higher rate OTUCn signal can be substituted for these signals. In this particular scenario, there are two adjacent ODUC1 hops, and the NE2 demultiplexs (and multiplexes) the ODU4 onto the ODUC1. It is possible to construct an alternative scenario in the case when NE2 acts as a regenerator, and doesn’t terminate the ODUC1 signals in the two hops, and instead repeats the ODUC1 signal; this scenario is specifically discussed in Section 4.4.

```
 ==================================================================


       +----------+                                +----------+
       |  100GE   |                                |  100GE   |
       +----------+        +---------------+       +----------+
       |  ODU4    |        |      ODU4     |       |  ODU4    |
       +----------+        +---------------+       +----------+
       |  ODUC1   |        | ODUC1 | ODUC1 |       |  ODUC1   |
       +----------+        +---------------+       +----------+
       |  OTUC1   +--------+ OTUC1 | OTUC1 +-------+  OTUC1   |
       +----------+        +---------------+       +----------+
          NE1                    NE2                   NE3

                +------------->         +------------->
                   Scope of                Scope of
                   OTUC1, ODUC1            OTUC1, ODUC1


 ==================================================================
```

Figure 5: 100GE Client service

4.2. 100GE Client Service with a mix of ODU4, and ODUC1 connections

In the scenario illustrated in Figure 6 a 100GBASE-R client is mapped into an ODU4 at NE1. The resulting ODU4 signal is encapsulated with an OTU layer to form the OTU4 signal. Actions performed at NE2 are: (a) terminate OTU4 layer, and extract the ODU4 signal (b) map the ODU4 signal onto a different ODUC1/OTUC1 towards NE3. NE3 performs the same set of actions that were performed by NE3 in Figure 5. This usecase illustrates a scenario in which an ODU4 signal can span between network elements regardless of whether they support the OTUCn interfaces or not.

```
==================================================================


    +----------+                                  +----------+
    |  100GE   |                                  |  100GE   |
    +----------+       +----------------+         +----------+
    |  ODU4    |       |      ODU4      |         |  ODU4    |
    +----------+       +-------+--------+         +----------+
    |  OTU4    +-------+ OTU4  | ODUC1  |         |  ODUC1   |
    +----------+       +----------------+         +----------+
                               | OTUC1  +---------+  OTUC1   |
                               ---------+         +----------+
         NE1                 NE2                     NE3

            +-------------->          +---------------->
                Scope of                     Scope of
                OTU4 layer                   OTUC1, ODUC1


==================================================================
```

Figure 6: 100GE Client Service with a mix of OTU4, and OTUC1 links

4.3. Transport of non-OTN Client Signal over ODUCn connection

Editor Note: this section may not be needed, as this section mainly describes the setup of client signal over ODUflex, then over ODUCn. Setup of ODUk/ODUflex can reuse mechanisms defined in RFC7139.

4.3.1. 400GE Client Service with a mix of OTUCn links

In the scenario illustrated in Figure 7 a 400GBASE-R client is mapped into an ODUflex at NE1. The resulting ODUflex signal is multiplexed into an ODUC4 (using GMP), and then transformed into an OTUC4 signal. The links between NE1-NE2, and NE2-NE3 are OTUC4 and OTUC6 (respectively). Actions performed at NE2 are: (a) terminate OTUC4, and ODUC4 towards NE1 (b) demultiplex the ODUflex signal from ODUC4 (c) map the ODUflex signal onto ODUC6/OTUC6 towards NE3. NE3 performs the inverse sequence of steps performed at NE1, and recovers the 400GBASE-R client from the ODUflex signal.

Although not specifically illustrated in this figure, the 200G of spare capacity in the NE2-NE3 links can be used to carry other client signals.. Although the scenario illustrated in Figure 7 is specific to 400GE, the treatment for packet clients at other rates (e.g. 25G, 50G, 200G) follows a very similar processing sequence. In the case of 25GBASE-R clients , the 25GE client signal will be mapped to an ODUflex, and can be multiplexed into an ODU4 signal, or an ODUCn signal as illustrated here.

```
==================================================================


    +----------+                                +----------+
    |  400GE   |                                |  400GE   |
    +----------+        +---------------+       +----------+
    |  ODUflex |        |    ODUflex    |       |  ODUflex |
    +----------+        +-------+-------+       +----------+
    |  ODUC4   |        | ODUC4 | ODUC6 |       |  ODUC6   |
    +----------+        +---------------+       +----------+
    |  OTUC4   +--------+ OTUC4 | OTUC6 +-------+  OTUC6   |
    +----------+        +-------+-------+       +----------+
       NE1                   NE2                     NE3

             <------------->           <------------->
                Scope of                   Scope of
                OTUC4, ODUC4               OTUC6, ODUC6



==================================================================
```

Figure 7: 400GE transport over OTUCn links

4.3.2. FlexE aware transport over OTUCn links

In the scenario illustrated in Figure 8 NE1 interfaces to a client equipment which includes the FlexE SHIM functions which originate/ terminate a FlexE group. The transport network edge node NE2 is FlexE aware -- but doesn't terminate the FlexE group. NE1 may (as defined in the FlexE draft [I-D.izh-ccamp-flexe-fwk]), crunch the unavailable tributary slots in the FlexE PHY signals, and map the resultant stream to one or more ODUflex signals. The links between NE1-NE2, and NE2-NE3 are OTUC4 and OTUC6 (respectively). Actions performed at NE2 are: (a) terminate OTUC4, and ODUC4 towards NE1 (b) demultiplex the ODUflex signal from ODUC4 (c) map the ODUflex signal onto ODUC6/OTUC6 towards NE3. NE3 recovers the Crunched and combined PHY(s) from the ODUflex signal, re-adds the unavailable calendar slots, and outputs the resulting stream towards the FlexE PHY(s).

In the scenario illustrated in Figure 8 the lowest rate OTUCn link is the OTUC4 link between NE1-NE2. This means that the size of the FlexE group is at most 4. FlexE groups with greater sizes can be handled by utilizing appropriate OTUCn links. Note that at most 400G of the capacity of OTUC6 (or 600G) NE2-NE3 link is occupied by the ODUflex signal; the remaining bandwidth can be allocated to other client signals.

```
 ==================================================================


FlexE    +----------+                      +----------+       FlexE
group    | Crunched |                      | Crunche  |       Group
  +------> and      |                      | and      +-------->
         | Combined |                      | Combined | Add
         | PHY(s)   |                      | PHY(s)   | unavailable
         +----------+   +---------------+  +----------+ Calendar
         |  ODUflex |   |    ODUflex    |  |  ODUflex | slots
         +----------+   +-------+-------+  +----------+
         |  ODUC4   |   | ODUC4 | ODUC6 |  |  ODUC6   |
         +----------+   +---------------+  +----------+
         |  OTUC4   +---+ OTUC4 | OTUC6 +---+  OTUC6   |
         +----------+   +-------+-------+  +----------+
            NE1            NE2               NE3

                  <--------->        <----------->
                  Scope of            Scope of
                  OTUC4, ODUC4        OTUC6, ODUC6


 ==================================================================
```

Figure 8: FlexE aware transport over OTUCn links

4.3.3. FlexE Client transport over OTUCn links

This use case (see Figure 9) concerns the scenario in which a FlexE group is terminated at the transport network edge node (via the FlexE SHIM function), and the FlexE clients are demultiplexed, and independently transported through the OTN network. In the scenario illustrated in Figure 9 the lowest rate OTUCn link is the OTUC4 link between NE1-NE2. This means that the maximum bit rate of the FlexE client is at most 400G. FlexE clients with greater sizes can be handled by utilizing appropriate OTUCn links. This figure illustrates the case in which one FlexE client is transported between NE1 and NE3. Other FlexE clients recovered at NE1 can routed independently to NE3, or to other network elements.

```
 ==================================================================


      +------------------+                       +----------------+
      |     FlexE        |                       |    FlexE       |
      |     Client       |                       |    Client      |
      +------------------+                       +----------------+
      | FlexE |          |   +----------------+  |       | Flex   |
      | SHIM  | ODUflex  |   |    ODUflex     |  |ODUflex| SHIM   |
      +------------------+   +----------------+  +----------------+
      | FlexE | ODUC4    |   | ODUC4 | ODUC6  |  | ODUC6 | FlexE |
 +----+ PHY(s +----------+   +----------------+  +-------+ PHY(s)+---->
FlexE |       | OTUC4    +---+ OTUC4 | OTUC6 +---+ OTUC6 |       | FlexE
Group +------------------+   +----------------+  +----------------+ Group
                 NE1            NE2                  NE3

                 <------------>      <----------->
                    Scope of           Scope of
                    OTUC4, ODUC4       OTUC6, ODUC6


 ==================================================================
```

Figure 9: FlexE client transport over OTUCn links

4.4. Multihop ODUCn link

As mentioned in the introductory section, the ODUCn is not a switchable entity. The ODUCn layer is a server layer, which more-or-less occupies the position of a section layer in OTN networks. As such, the ODUCn signal must be terminated and the contained low-order ODU flows can be switched independently to other OTN interfaces. G.709 and G.872 however allow for digital regenerators to terminate the OTUCn layer, and reinject the ODUCn layer towards another interface (where a new OTUCn section layer is started). This scenario is illustrated in Figure 10. In this figure, NE3 is the regenerator. The ODUC2 signal is terminated at NE2, and NE4. At the regeneration points, all the clients embedded inside the ODUCn signal are not touched (i.e. no TS changes can occur). More specifically, the OPUC2 signal is not modified in any way. However, the ODUC2 OH may be modified if intrusive TCM monitoring points are applied to the ODUC2 signal at NE3. It is for this reason that the ODUC2 entity must be visible at NE3.

In scenarios involving multi-hop ODUCn links, GMPLS signalling will be required to setup multiple ODUCn LSPs, each covering a regenerator section (since an end-to-end ODUCn LSP is not possible except in very simple configurations). A LO-ODU can then be switched across multiple ODUCn LSPs (possibly with different rates).

```
   ==================================================================


 +----------+                                                +----------+
 |  100GE   |                                                |  100GE   |
 +----------+     +---------------+                          +----------+
 |  ODU4    |     |      ODU4     |                          |  ODU4    |
 +----------+     +-------+-------+       +----------+       +----------+
 |  ODUC1   |     | ODUC1 | ODUC2 |       |  ODUC2   |       |  ODUC2   |
 +----------+     +---------------+       +----------+       +----------+
 |  OTUC1   +-----+ OTUC1 | OTUC2 +-------+  OTUC2   +-------+  OTUC2   |
 +----------+     +-------+-------+       +----------+       +----------+
    NE1                NE2                     NE3              NE4

         <------------->     <------------->    <------------->
            Scope of              OTUC2                OTUC2
            OTUC1, ODUC1

                             <--------------------------------->
                                          ODUC2


   ==================================================================
```

Figure 10: Multihop ODUCn link

4.5. Use of OTUCn-M links

The scenario illustrated in Figure 11 is a variant of the basic usecase presented in Figure 5. The only difference is that the second hop of the ODU4 connection makes use of a OTUC2-30 link which has a capacity of 150G.

```
    =================================================================

          +----------+                              +-----------+
          |  100GE   |                              |  100GE    |
          +----------+       +------------------+   +-----------+
          |  ODU4    |       |       ODU4       |   |  ODU4     |
          +----------+       +-------+----------+   +-----------+
          |  ODUC1   |       | ODUC1 | ODUC2    |   |  ODUC2    |
          +----------+       +------------------+   +-----------+
          |  OTUC1   +-------+ OTUC1 | OTUC2-30 +------+  OTUC2-30 |
          +----------+       +-------+----------+   +-----------+
             NE1                 NE2                    NE3

                   +------------->           +------------->
                      Scope of                   Scope of
                      OTUC1, ODUC1               OTUC2-30
                                                 ODUC2


    =================================================================
```

Figure 11: 100GE Client service over OTUCn-M links

4.6. Intermediate State of ODU mux

The ODUCn links have a tributary slot granularity of 5G -- and this makes it a bit inefficient if a small number of ODU0 flows have to be switched across an ODUCn links. In these cases, it is conceivable that the intermediate nodes may offer the convenience of a intermediate-stage multiplexing, whereby multiple ODU0 flows are first multiplexed into a higher rate container (e.g. ODU2), and then multiplexed into an ODUCn signal. This however assumes that all these ODU0 flows are co-routed in the network. If this assumption cannot be made, the only solution is to multiplex these ODU0 flows into higher rate flows, from the source of the traffic. This usecase isn’t elaborated in this document. We can add details if required.

5. GMPLS Implications

5.1. OTN ODUCn/OTUCn hierarchy

As described in [ITU-T_G872], the digital layers of the OTN are divided into the OTU layer and a hierarchy of one or more ODU layers. As an ODUCn cannot be used to support non-OTN client signals, the OTN client signals (e.g. ODU0, ODU1, ODU2, ODU2e, ODU3, ODU4, ODUflex) are first multiplexed into an ODUCn container, then the ODUCn container is then mapped into OTUCn (see Figure 1). The signal hierarchy supported by the ODUCn and OTUCn needs to be taken into consideration in control plane Routing and Signaling.

ODUCn based connection management is concerned with controlling the connectivity of ODUCn paths. According to [ITU-T_G872], the intermediate nodes with ODUCn do not support the switching of ODUCn tributary slot. Intermediate ODUCn points are only considered as a forwarding node. Once an ODUCn path is used to transport client signal, the TS occupied will not change across the ODUCn network.

5.2. OTUCn/OTUCn-M/ODUCn LSP

OTUCn/OTUCn-M Link is different from tranditional OTUk link. The OTUk link is already configured once two matched OTU interfaces are connected. But for setup of OTUCn link, the first thing that needs to do is to bond several different OTUC instances together as one group, which is seen as one OTUCn link. Control plane mechanisms are needed to finish the bonding of these instances.

For transportation of client signal over ODUCn signal, an ODUCn LSP is also needed to be configured with control plane mechanisms in advance.

Once an ODUCn LSP is set up, the signaling mechanism defined in [RFC7139] can be reused to set up ODUk LSP over ODUCn link. Setup of ODUk LSP over ODUCn LSP is out of the scope of this document.

5.3. Implications for GMPLS Signaling

[RFC7139] extends the base RSVP-TE signaling specification [RFC4328] to define RSVP-TE signaling extensions that can used to control OTN networks built in accordance with [ITU-T_G709_2012]. [ITU-T_G709_2016] introduced some new containers, such as OPUCn, ODUCn, and OTUCn. The mechanisms defined in [RFC7139] do not support these new OTN features. Therefore, GMPLS signaling protocols MUST be extended to support this new functionality. The following summarizes key aspects that should be considered for GMPLS signaling extensions:

a. Per the description in clause 7 of [ITU-T_G872], "the digital layers of the OTN are divided into the OTU layer and a hierarchy of one or more ODU layers". In B100G links, the ODUCn layer is the bottom of the ODU hierarchy, and an ODUCn (induced) LSP needs to be established before the LO-ODUs can flow across this link. The traffic parameters in a signaling message should be extended to support the new signal type(s) for the ODUCn signals. This approach keeps the treatment for ODUCn signals consistent with that of other ODU(s).

b. Support the new TS granularity: the signaling protocol should be able to identify the TS granularity (i.e., the new 5 Gbps TS granularity) to be used for establishing a Hierarchical LSP that will be used to carry service LSP(s) requiring a specific TS granularity.

c. A new label format MUST carry the information about one or more OTUC/ODUC instances to be bonded together.

d. Support for LSP setup of OTUCn sub rates (OTUCn-M) path: based on previous extensions, there should be new signal mechanism to declare the OTUCn-m information. The GMPLS signalling protocol SHALL support the setup of OTUCn sub rates (OTUCn-M) LSP, which includes the negotiation of unavaliable slots number, slots postion and allocation of slot resources.

e. The GMPLS signalling protocol should be able to specify the new ODUCn/OTUCn signal types and related traffic information. The traffic parameters should be extended in a signalling message to support the new ODUCn/OTUCn signal types

5.4. Implications for GMPLS Routing

The path computation process needs to select a suitable route for an ODUCn/OTUCn/OTUCn-M connection request. In order to perform the path computation, it needs to evaluate the available bandwidth/slots available on one or more candidate links. The routing protocol SHOULD be extended to carry sufficient information to represent ODU Traffic Engineering (TE) topology.

The Interface Switching Capability Descriptors defined in [RFC4203] present a new constraint for LSP path computation. [RFC4203] defines the Switching Capability, related Maximum LSP Bandwidth, and Switching Capability specific information. [RFC7138] updates the ISCD to support ODU4, ODU2e and ODUflex. The new Switching Capability specific information provided in [RFC7138] have to be adapted to support new features contained in [G709-2016]. The following requirements should be considered:

a. Support for carrying the link multiplexing capability: As discussed in Section 3.1.2, many different types of low-order ODU(s) (e.g. ODUflex, ODU4) can be multiplexed into the ODUCn. An ODUCn path may support one or more types of ODUk signals. The routing protocol should be capable of carrying this multiplexing capability.

b. Support for advertising 5G Tributary Slot Granularity introduced [ITU-T_G709_2016].

c. Support for advertisement of available bandwidth in an ODUCn path.

6. Acknowledgements

7. Authors (Full List)

Qilei Wang (editor)

ZTE

Nanjing, China

Email: wang.qilei@zte.com.cn

Radha Valiveti (editor)

Infinera Corp

Sunnyvale, CA, USA

Email: rvaliveti@infinera.com

Haomian Zheng (editor)

Huawei

CN

EMail: zhenghaomian@huawei.com

Huub van Helvoort

Hai Gaoming B.V

EMail: huubatwork@gmail.com

Sergio Belotti

Nokia

EMail: sergio.belotti@nokia.com

Iftekhar Hussain

Infinera Corp

Sunnyvale, CA, USA

Email: IHussain@infinera.com

Daniele Ceccarelli

Ericsson

Email: daniele.ceccarelli@ericsson.com

8. Contributors

Rajan Rao, Infinera Corp, Sunnyvale, USA, rrao@infinera.com

Fatai Zhang, Huawei,zhangfatai@huawei.com

Italo Busi, Huawei,italo.busi@huawei.com

Zheyu Fan, Huawei, fanzheyu2@huawei.com

Yuanbin Zhang, ZTE, Beiing, zhang.yuanbin@zte.com.cn

Zafar Ali, Cisco Systems, zali@cisco.com

Daniel King, d.king@lancaster.ac.uk

Manoj Kumar, Cisco Systems, manojk2@cisco.com

Antonello Bonfanti, Cisco Systems, abonfant@cisco.com

Akshaya Nadahalli, Cisco Systems, anadahal@cisco.com

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

None.

11. References

11.1. Normative References

[ITU-T_G709.1]
ITU-T, "ITU-T G.709.1: Flexible OTN short-reach interface; 2016", , 2016.

[ITU-T_G709_2012]
ITU-T, "ITU-T G.709: Optical Transport Network Interfaces; 02/2012", http://www.itu.int/rec/T-REC-G..709-201202-S/en, February 2012.

[ITU-T_G709_2016]
ITU-T, "ITU-T G.709: Optical Transport Network Interfaces; 07/2016", http://www.itu.int/rec/T-REC-G..709-201606-P/en, July 2016.

[ITU-T_G872]
ITU-T, "ITU-T G.872: The Architecture of Optical Transport Networks; 2017", http://www.itu.int/rec/T-REC-G.872/en, January 2017.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4328] Papadimitriou, D., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control", RFC 4328, DOI 10.17487/RFC4328, January 2006, <https://www.rfc-editor.org/info/rfc4328>.

[RFC7138] Ceccarelli, D., Ed., Zhang, F., Belotti, S., Rao, R., and J. Drake, "Traffic Engineering Extensions to OSPF for GMPLS Control of Evolving G.709 Optical Transport Networks", RFC 7138, DOI 10.17487/RFC7138, March 2014, <https://www.rfc-editor.org/info/rfc7138>.

[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, DOI 10.17487/RFC7139, March 2014, <https://www.rfc-editor.org/info/rfc7139>.

11.2. Informative References

[I-D.izh-ccamp-flexe-fwk]
Hussain, I., Valiveti, R., Pithewan, K., Wang, Q., Andersson, L., Zhang, F., Chen, M., Dong, J., Du, Z., zhenghaomian@huawei.com, z., Zhang, X., Huang, J., and Q. Zhong, "GMPLS Routing and Signaling Framework for Flexible Ethernet (FlexE)", draft-izh-ccamp-flexe-fwk-00 (work in progress), October 2016.

Authors' Addresses

Qilei Wang (editor)
ZTE
Nanjing
CN

Email: wang.qilei@zte.com.cn

Radha Valiveti (editor)
Infinera Corp
Sunnyvale
USA

Email: rvaliveti@infinera.com

Haomian Zheng (editor)
Huawei
CN

Email: zhenghaomian@huawei.com

Huub van Helvoort
Hai Gaoming B.V

Email: huubatwork@gmail.com

Sergio Belotti
Nokia

Email: sergio.belotti@nokia.com