IA OIF-ENNI-OTNv3-AM-01.0

OTNv3 Amendment to OIF UNI 2.0 and E-NNI 2.0/2.1 Common Part

OIF-ENNI-OTNv3-AM-01.0

May 19, 2014

Implementation Agreement created and approved
by the Optical Internetworking Forum
www.oiforum.com

The OIF is an international non profit organization with over 90 member companies, including the world's leading carriers and vendors. Being an industry group uniting representatives of the data and optical worlds, OIF's purpose is to accelerate the deployment of interoperable, cost-effective and robust optical internetworks and their associated technologies. Optical internetworks are data networks composed of routers and data switches interconnected by optical networking elements.

With the goal of promoting worldwide compatibility of optical internetworking products, the OIF actively supports and extends the work of national and international standards bodies. Working relationships or formal liaisons have been established with IEEE 802.1, IEEE 802.3ba, IETF, IP-MPLS Forum, IPv6 Forum, ITU-T SG13, ITU-T SG15, MEF, ATIS-OPTXS,ATIS-TMOC, TMF and the XFP MSA Group.

For additional information contact:
The Optical Internetworking Forum, 48377 Fremont Blvd.,
Suite 117, Fremont, CA 94538
510-492-4040 ☞ info@oiforum.com

www.oiforum.com

Working Group: **Networking & Operations**

TITLE: **OTNv3 Amendment to OIF UNI 2.0 and E-NNI 2.0/2.1**

SOURCE: **TECHNICAL EDITORS**

Fred Gruman
Fujitsu Network Communications, Inc.
2801 Telecom Parkway
Richardson TX 75082
Phone: +1 972 479 2477
fred.gruman@us.fujitsu.com

WORKING GROUP CHAIR

Remi Theillaud
Marben Products
remi.theillaud@marben-products.com

ABSTRACT: This implementation agreement specifies the extension to UNI Signaling 2.0, E-NNI Routing 2.1 and E-NNI Signaling 2.1 to support OTNv3 (ITU-T G.709 Edition 4, 2009-12).

Table of Contents

List of Figures

List of Tables

1 Introduction

This document specifies the architecture and protocol extensions to support the electrical ODU switching layer defined in ITU-T G.709 Edition 4 [G.709Ed4], herein referred to as OTNv3, and the procedures for hitless adjustment of ODUflex as defined in ITU-T G.7044 [G.7044]. It is an amendment to the OIF UNI Signaling 2.0 [OIF-UNI-02.0], E-NNI Signaling 2.0 [OIF-E-NNI-Sig-02.0] and E-NNI Routing 2.0 [OIF-E-NNI-OSPF-02.0] Implementation Agreements and reuses the terms and concepts described in the base OIF IAs.

1.1 Problem Statement

The advent of the automatic switched optical network has necessitated the development of interoperable procedures for requesting and establishing dynamic connection services across heterogeneous networks. The development of such procedures requires the definition of:

- Control domains and associated reference points (E-NNI, I-NNI, UNI)
- Services offered by the optical transport network across control domains
- Signaling protocols used to invoke the services across E-NNI interfaces
- Mechanisms used to transport signaling messages
- Routing protocols used to distribute topology across E-NNI interfaces
- Mechanisms for auto-discovery of link adjacency and services across E-NNI interfaces

The first phases of specifying the user-provider call control signaling interface and External Network Node Interface (E-NNI) signaling protocols has been completed in [OIF-UNI-01.0-R2] and [OIF-ENNI-SIG-01.0]. The second phase has been completed in [OIF-UNI-02.0] and [OIF-E-NNI-Sig-02.0]. The first phase of specifying OSPF-TE based routing has been completed in [OIF-ENNI-OSPF-01.0] and the second phase has been completed in [OIF-ENNI-OSPF-02.0].

The third phase has been the development of amendments to [OIF-E-NNI-Sig-02.0] and [OIF-ENNI-OSPF-02.0] to address multilayer routing and signaling [OIF-ENNI-ML-AM-01.0].

This document is an additional amendment to add support for OTNv3 to [OIF-UNI-02.0], [OIF-E-NNI-Sig-02.0], and [OIF-ENNI-OSPF-02.0]. It also addresses OTN-specific updates to the OIF E-NNI Multilayer Amendment [OIF-ENNI-ML-AM-01.0].

1.2 Scope

The scope of this Implementation Agreement is the specification of OTNv3-specific routing, signaling, and signaling-based discovery abstract messages, attributes, and flows enabling end-to-end dynamic establishment of transport connections across multiple control domains supporting multiple layer networks. The Implementation Agreement also provides a concrete realization of the abstract messages and attributes based on the RSVP-TE and OSPF-TE protocols.

This document addresses the electrical ODU switching layer. Specifically, this document will address the following OTNv3 features:

- OTU containers: OTU1, OTU2, OTU3 and OTU4
- ODU containers (fixed): ODU0, ODU1, ODU2, ODU2e, ODU3 and ODU4
- ODU containers (flexible): ODUflex(CBR), ODUflex(GFP) non-resizeable and ODUflex(GFP) resizeable
- 2.5G and 1.25G tributary slot granularity
- Single-stage and multi-stage multiplexing hierarchy
- Hitless resize of ODUflex(GFP)

The virtual concatenation of ODU is not addressed in this document. Additionally, signaling support for the configuration of OTN Operations, Administration and Maintenance (OAM) is not addressed.

The ITU-T G.709 optical OCh switching layer is outside the scope of this document.

OIF UNI Signaling 2.0, E-NNI Signaling 2.0, and E-NNI Routing 2.0 provide support for the OTN (ODU) layer as defined in ITU-T G.709 Edition 2 [G.709Ed2], herein referred to as OTNv1. This includes support for the OTU1, OTU2 and OTU3 line rates and the ODU1, ODU2 and ODU3 signal types. This amendment will address backwards compatibility issues with interoperating with implementations that support OTNv1.

Within this document OTNv3 refers to ITU-T G.709 Edition 4 [G.709Ed4] while OTNv1 refers to ITU-T G.709 Edition 2 [G.709Ed2]. When OTN is used without reference to a specific edition of ITU-T G.709, it refers to OTNv3. References to ITU-T G.709 Edition 2 will always be explicitly identified as OTNv1.

1.3 Relationship to Other Standards Bodies

This Implementation Agreement, to the greatest extent possible, uses available global standards documents.

This document covers the OTN technology layers as specified by ITU-T in [G.709Ed4] and the hitless adjustment of ODUflex as specified in [G.7044].

Signaling and routing extensions for the ODU switching layer is based on the work in IETF as described in [RFC7062], [RFC7096], [RFC7138] and [RFC7139].

1.4 Merits to OIF

The OTNv3 Amendment is a key stride towards the implementation of an open inter-domain routing and signaling protocol that enables dynamic setup and release of various services across OTN (ODU) networks. This activity supports the overall mission of the OIF.

1.5 Working Groups

Networking and Operations Working Group

Carrier Working Group

Interoperability Working Group

1.6 Document Organization

This document is organized as follows:

- Section 1: Introduction
- Section 2: Terminology and Abbreviations
- Section 3: OTNv3 Feature Set
- Section 4: OTNv3 Architecture
- Section 5: Abstract Messages and Flows
- Section 6: Security and Logging for OTNv3
- Section 7: Compatibility with UNI and E-NNI
- Section 8: References
- Appendix A: OTN Overview

1.7 Keywords

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 Terminology and Abbreviations

ASON Automatically Switched Optical Network
BEI Backward Error Indication
BIP Bit Interleaved Parity
CBR Constant Bit Rate
E-NNI External Network-Network Interface
G-PID Generalized Payload Identifier
GbE Gigabit Ethernet
GCC General Communication Channel
GFP Generic Framing Procedure
GMPLS Generalized Multi-Protocol Label Switching
HO-ODU Higher-Order ODU
HO-ODUk Higher-Order ODU of order k
HO-OPU Higher-Order OPU
HO-OPUk Higher-Order OPU of order k
IA Implementation Agreement
IEEE Institute of Electrical and Electronics Engineers
IETF Internet Engineering Task Force
I-NNI Internal Network-Network Interface
ITU-T International Telecommunications Union – Telecommunications Standardization Sector
LC Link Connection
LO-ODU Lower-Order ODU
LO-ODUk Lower-Order ODU of order k
LO-OPU Lower-Order OPU
LO-OPUk Lower-Order OPU of order k
LSP Label Switched Path
MEP: Maintenance End Point
MIP: Maintenance Intermediate Point
MPLS Multi-Protocol Label Switching
MT Multiplier
NMC Number of Multiplexed Components
NNI Network Node Interface
NE Network Element
NVC Number of Virtual Components
OAM: Operations, Administration and Maintenance
OCh Optical Channel
ODUflex Optical channel Data Unit flexible
ODTU Optical channel Data Tributary Unit
ODTUjk Optical channel Data Tributary Unit j into k
ODTUk.ts Optical channel Data Tributary Unit k with ts tributary slots
ODTUG Optical channel Data Tributary Unit Group
ODTUGk Optical channel Data Tributary Unit Group of order k
ODU Optical channel Data Unit
ODUk Optical channel Data Unit of order k
OIF Optical Internetworking Forum
OMS Optical Multiplex Section
OPS Optical Physical Section
OPU Optical channel Payload Unit
OPUk Optical channel Payload Unit of order k
OSPF Open Shortest Path First
OSPF-TE OSPF with Traffic Engineering extensions
OTM Optical Transport Module

OTN Optical Transport Network
OTNv1 Optical Transport Network as defined in [G.709Ed2]
OTNv3 Optical Transport Network as defined in [G.709Ed4]
OTS Optical Transmission Section
OTU Optical channel Transport Unit
OTUk Optical channel Transport Unit of order k
PM Path Monitoring
RSVP Resource reSerVation Protocol
RSVP-TE RSVP with Traffic Engineering extensions
SC Switched Connection
SCN Signaling Communications Network
SM Section Monitoring
SNC Subnetwork Connection
SPC Soft permanent connection
ST Signal Type
TCM Tandem Connection Monitoring
TDM Time Division Multiplexing
TLV Type-Length-Value encoding
TTI Trail Trace Identifier
TSG Tributary Slot Granularity
UNI User Network Interface
UNI-N UNI Signaling Agent – Network
UNI-C UNI Signaling Agent – Client

3 OTNv3 Architecture

The OTNv3 architecture is defined in ITU-T G.709 Edition 4 [G.709Ed4]. This document addresses the ODU switching layer including the associated OPU and OTU layers. This document does not address the optical OCh switching layer and associated layers (e.g., OMS, OTS, OPS, OTM)

3.1 OTN Hierarchy

The OTNv3 architecture supports a number of data rates *k* at the OPU, ODU and OTU layers. The client payload (which may be a lower-order ODU) is mapped into the OPU with OPU-specific overhead. The OPU is mapped into an ODU with ODU-specific overhead. The ODU is either mapped into an OTU with OTU-specific overhead or mapped into an ODTU. The ODTU is multiplexed into an ODTU Group (ODTU) which is then mapped into a higher-order OPU. The OTN mapping and multiplexing hierarchy is shown in Figure 1 and Figure 2.

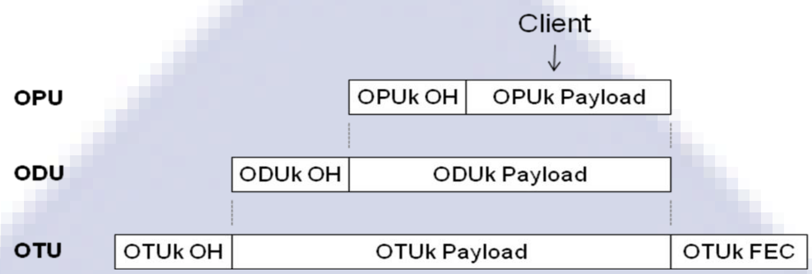

Figure 1: OTN Mapping Hierarchy, Direct Mapping of ODUk into OTUk

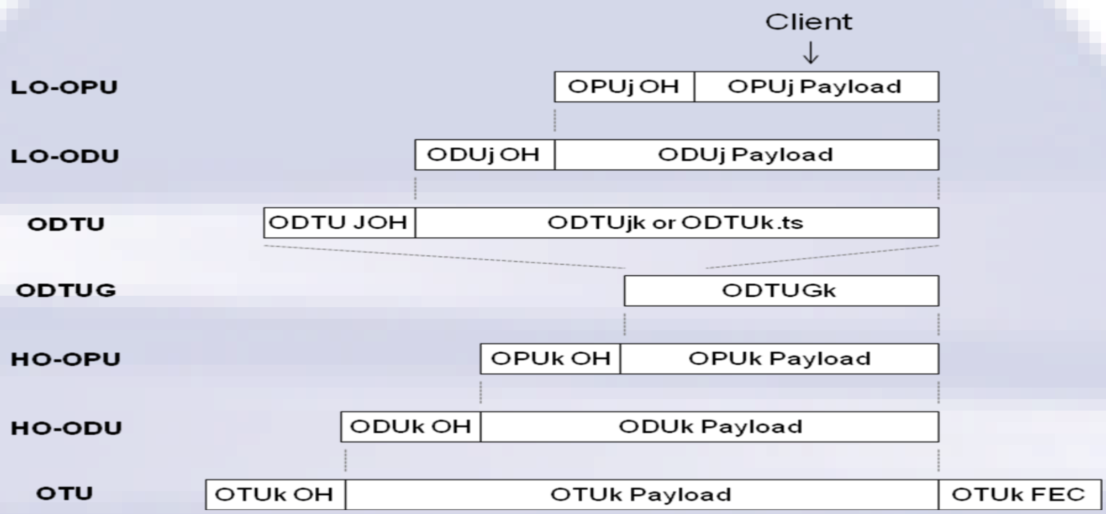

Figure 2: OTN Mapping and Multiplexing Hierarchy, Single-Level Multiplexing

3.2 OTNv3 Support for OPU, ODU and OTU Types

The OTNv3 architecture supports the OPU, ODU and OTU types shown in Table 1.

Table 1: OTNv3 Supported OPU, ODU and OTU Types

Layer	Supported Types
OPU	OPU0, OPU1, OPU2, OPU2e, OPU3, OPU4, OPUflex(CBR), OPUflex(GFP-F)
ODU	ODU0, ODU1, ODU2, ODU2e, ODU3, ODU4, ODUflex(CBR), ODUflex(GFP-F)
OTU	OTU1, OTU2, OTU3, OTU4

3.3 Tributary Slot Granularity (TSG) and Payload Type (PT)

The OTNv3 architecture supports two TSGs: 2.5G and 1.25G. The 2.5G tributary slot granularity was defined in OTNv1 and was used to map an ODU1 into an ODU2 or ODU3 or to map an ODU2 into an ODU3. A payload type of 20 is used for these mappings.

The 1.25G tributary slot granularity is defined in OTNv3. OTNv3 supports the mapping of all legal ODUj into ODUk using the 1.25G TSG, including the ability to map ODU1 into an ODU2 or ODU3 or an ODU2 into an ODU3. A payload type of 21 is used for these mappings with the exception of ODU0 into ODU1 which uses a payload type of 20.

For OTUk links supporting multiplexing of ODUj into ODUk (j<k), the ODUk tributary slot granularity must match at both ends. Implementations that support 1.25G TSG MUST support the ability to fallback to using 2.5G TSG if the other end does not support 1.25G TSG. The 1.25G TSG should be used if both ends support 1.25G TSG; however, implementations may choose to allow the user to force a link to use the 2.5G TSG.

The TSG capability of the neighbor across the UNI or E-NNI can be automatically discovered or manually configured. The automatic learning of neighbor capabilities is outside the scope of this Implementation Agreement.

3.4 Tributary Slot (TS)

When multiplexing an ODUj into an ODUk (j<k), the ODUj is mapped into a number of tributary slots (TS) in the OPUk payload. The number of TS varies depending on the rates of the ODUj and ODUk.

Table 2 and Table 3 show the number of tributary slots required for mapping a fixed ODUj into an ODUk for 1.25G and 2.5G TSG, respectively.

Table 2: Tributary Slots for 1.25G TSG

LO ODU	OPU1	OPU2	OPU3	OPU4
ODU0	1	1	1	1
ODU1	--	2	2	2
ODU2	--	--	8	8
ODU2e	--	--	9	8
ODU3	--	--	--	31

Table 3: Tributary Slots for 2.5G TSG

LO ODU	OPU2	OPU3
ODU1	1	1
ODU2	--	4

The mapping of ODUflex(CBR) into an ODUk uses a variable number of TS depending upon the rate and tolerance of the ODUflex(CBR) and the rate and tolerance of the higher-order OPUk/ODUk. The number of tributary slots is determined by the following formula:

$$\text{Number of TS} = \left\lceil \frac{ODUflex(CBR)\ nominal\ bit\ rate \times (1 + ODUflex\ bit\ rate\ tolerance)}{T \times ODTUk.ts\ nominal\ bit\ rate \times (1 - HO\ OPUk\ bit\ rate\ tolerance)} \right\rceil$$

where:

T = Transcoding Factor

The ODTUk.ts nominal bit rate is shown in Table 4. The HO OPUk bit rate tolerance is +/- 20 ppm for k=2, 3 and 4.

Table 4: ODTUk.ts Bit Rates

ODTU Type	Nominal bit-rate (kbps)
ODTU2.ts	1 249 409.620
ODTU3.ts	1 254 703.729
ODTU4.ts	1 301 709.251

The ODUflex(CBR) bit rate tolerance is +/- 100 ppm.

Note that the number of tributary slots may not be constant end-to-end. For example,ODUflex(CBR) carrying the InfiniBand SDR client (2.5G) maps into 3 TS when carried on an ODU3 or 2 TS when carried on an ODU4.

The mapping of ODUflex(GFP) into an ODUk also uses a variable number of TS. However, as opposed to ODUflex(CBR), the number of TS is constant end-to-end and is dependent only on the bit rate of the ODUflex(GFP). Also, instead of supporting a continuous range of bit rates, ODUflex(GFP) supports a discrete set of bit rates depending on the number of TS.

ITU chose a discrete set of bit rates partially to ensure that different vendors supported the same bit rates, but more importantly to ensure that the same number of tributary slots were used end-to-end regardless of the TS rate of the HO-ODUk used to carry the ODUflex(GFP), as an ODUflex could be carried by a mixture of HO-ODUk types (e.g., ODU2, ODU3 or ODU4). To meet this objective, the TS size was chosen based on the smallest ODUk that could carry that particular ODUflex(GFP). An ODUflex(GFP) that uses 1-8 TSes would be based on the TS rate of an ODU2 (N x ODU2.ts rate), an ODUflex(GFP) that uses 9-32 TSes would be based on the TS rate of an ODU3 (N x ODU3.ts rate) and an ODUflex(GFP) that uses 33-80 TSes would be based the TS rate of an ODU4 (N x ODU4.ts rate). Even if a particular ODUflex(GFP) was not be routed over the smallest ODU type, it was believed to be safest to still use the TS size of the smallest HO-ODU over which it might be carried, as perhaps it could get later be rerouted over the smallest possible HO-ODUk.

Table 5 defines the set of discrete bit rates supported by ODUflex(GFP) as a function of the number of tributary slots (N) and the ODUk.ts rate. Table 6 defines the bit rate for the ODUk.ts.

Table 5: ODUflex(GFP) Bit Rates

# TS (N)	Nominal bit-rate
1-8	N x ODU2.ts
9-32	N x ODU3.ts
33-80	N x ODU4.ts

Table 6: ODUk.ts Bit Rates

ODU Type	Nominal bit-rate (kbps)
ODU2.ts	1 249 177.230
ODU3.ts	1 254 470.354
ODU4.ts	1 301 467.133

The explicit list of discrete ODUflex(GFP) rates and encodings are given in Appendix B.

3.5 Tributary Port Number (TPN)

The tributary port number (TPN) must be assigned and signaled for each ODUj multiplexed into an ODUk (j<k). The TPN identifies the tributary slot or the set of tributary slots used to multiplex an ODUj into an ODUk.

It is not always possible to define a static relationship between the tributary slot # and the TPN. For instance, the tributary slots assigned to a resizable ODUflex may change over the lifetime of the ODUflex as tributary slots are added and removed, but the TPN would remain unchanged.

However, there are cases where a fixed relationship between TPN and Tributary Slot # is possible. For example, when an ODU0 uses tributary slot #1 of a HO-ODU1, then the TPN=1. When the ODU0 uses tributary slot #2 of a HO-ODU1, then the TPN=2.

The TPN is used in conjunction with an ODTU type. Within an ODTU type, the TPN must be unique for each ODUk supported by that ODTU type. The following ODTU types are supported:

- ODTU01: mapping an ODU0 into an ODU1 with a 1.25G TSG.
- ODTU12: mapping an ODU1 into an ODU2 with either a 1.25G or 2.5G TSG
- ODTU13: mapping an ODU1 into an ODU3 with either a 1.25G or 2.5G TSG
- ODTU23: mapping an ODU2 into an ODU3 with either a 1.25G or 2.5G TSG
- ODTU2.ts: mapping an ODU0 or ODUflex into an ODU2 with a 1.25G TSG
- ODTU3.ts: mapping an ODU0, ODU2e or ODUflex into an ODU3 with a 1.25G TSG
- ODTU4.ts: mapping an ODU0, ODU1, ODU2, ODU2e, ODU3 or ODUflex into an ODU4 with a 1.25G TSG

Table 7 and Table 8 specify the rules for TPN assignment for 1.25G and 2.5G TSG, respectively.

Fixed relationship means that the TPN assignment is based on the tributary slot in use. I.e., if the LO-ODU maps into the first tributary slot then TPN=1, if it maps into the second tributary slot then TPN=2, etc. The mapping between tributary slot and TPN cannot be changed when the relationship is fixed.

Flexible mapping means that an implementation may choose the TPN within the valid range (as long as that TPN is not in use by the specific mapping) independent of the assigned tributary slots.

Table 7: TPN Assignment for 1.25G TSG

HO OPU	LO ODU	Mapping	TPN Range	Relationship to Tributary Slot #
OPU1	ODU0	ODTU01	1, 2	Fixed
OPU2	ODU1	ODTU12	1-4	Flexible
	ODU0 ODUflex	ODTU2.ts	1-8	Flexible
OPU3	ODU1	ODTU13	1-16	Flexible
	ODU2	ODTU23	1-4	Flexible
	ODU0 ODU2e ODUflex	ODTU3.ts	1-32	Flexible
OPU4	ODU0 ODU1 ODU2 ODU2e ODU3 ODUflex	ODTU4.ts	1-80	Flexible

Table 8: TPN Assignment for 2.5G TSG

HO OPU	LO ODU	Mapping	TPN Range	Relationship to Tributary Slot #
OPU2	ODU1	ODTU12	1-4	Fixed
OPU3	ODU1	ODTU13	1-16	Fixed
	ODU2	ODTU23	1-4	Flexible

3.6 OTN Communication Channels for the Signaling Communication Network (SCN)

The OTN architecture supports the following communication channels within the OTN overhead:

- General Communication Channel 0 (GCC0) within the OTU overhead
- General Communication Channel 1 (GCC1) within the ODU overhead
- General Communication Channel 2 (GCC2) within the ODU overhead
- Combined General Communication Channels 1 and 2 (GCC1+2) within the ODU overhead.

The GCC nominal bit rates are shown in Table 9.

Table 9: GCC Rates for ODU/OTU Types

ODU/OTU Type	GCC0 Nominal Bit Rate (kbit/s)	GCC1 Nominal Bit Rate (kbit/s)	GCC2 Nominal Bit Rate (kbit/s)	GCC1+2 Nominal Bit Rate (kbit/s)
ODU1/OTU1	326.723	326.723	326.723	653.445
ODU2/OTU2	1 312.405	1 312.405	1 312.405	2 624.810
ODU3/OTU3	5 271.864	5 271.864	5 271.864	10 543.729
ODU4/OTU4	13 702.203	13 702.203	13 702.203	27 404.405

4 OTNv3 Feature Set

4.1 Multiplexing Hierarchy

The OTN architecture supports the ability to map an ODU into an OTU at the same *k*-order without multiplexing. It also supports the ability to multiplex a lower-order ODU into a higher-order ODU. The OTN architecture supports a single stage of multiplexing as well as multiple stages of multiplexing.

The direct mapping of an ODUk into an OTUk is shown in Figure 3. In this case, there is no multiplexing hierarchy.

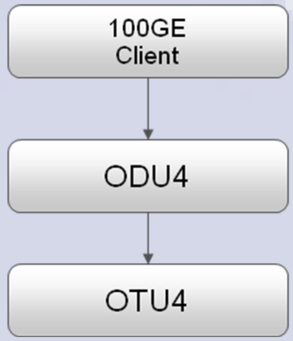

Figure 3: Example of Direct ODU Mapping

Single-stage multiplexing is when a non-ODU payload is mapped into an OPUj/ODUj that is then multiplexed into an OPUk/ODUk that is then mapped into an OTUk. Figure 4 shows an example where there is a single level of the multiplexing hierarchy.

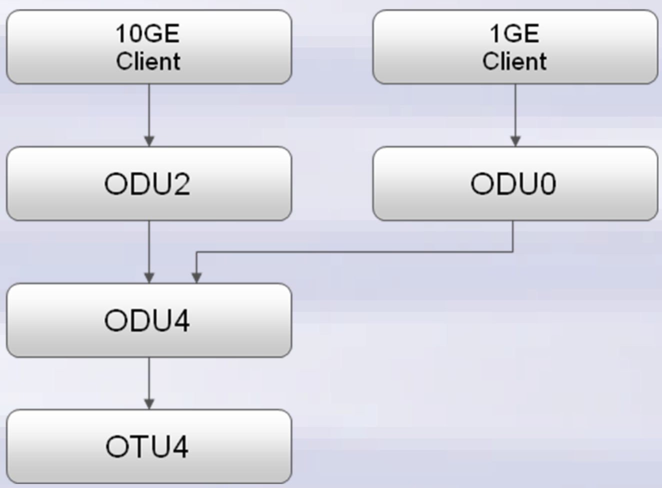

Figure 4: Example of Single-Level Multiplexing Hierarchy

Multi-stage multiplexing is where there is one or more intermediate ODU levels. For instance, a non-ODU payload is mapped into an OPUi/ODUi. The ODUi is multiplexed into an intermediate OPUj/ODUj. The intermediate ODUj is multiplexed into an OPUk/ODUk that is transported as an OTUk (Figure 5)

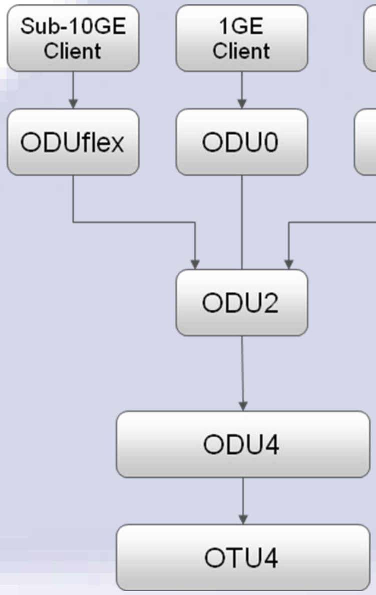

Figure 5: Example of Multi-Level Multiplexing Hierarchy

Many of the ODU services can be realized using a combination of a direct mapping of the ODU into an OTU or with single-stage multiplexing of the ODU into a higher-order ODU. However, there are some cases where there would be a need to support ODU services over multiple stages of the OTN multiplexing hierarchy.

One reason for multi-stage multiplexing is when there exists a portion of an OTN network that supports switching only for a subset of the ODU rates. For example, there may be a subnetwork of legacy OTNv1 equipment that only supports switching at the ODU1, ODU2 and ODU3 rates. In this example, it would be necessary to tunnel non-supported ODU rates, such as ODU0 and ODUflex through an ODU1, ODU2 or ODU3 tunnel. Tunneling leads to multi-stage multiplexing when the ODU tunnel is multiplexed into a higher-order OTU link (e.g., ODU0 → ODU2 tunnel → ODU3/OTU3).

Figure 6 shows an example of multistage multiplexing (MSM) where an ODUflex service is carried in an OTU2 tunnel through the legacy OTNv1 Network 2. In this figure, Nodes D and I support the capability of performing multistage multiplexing, i.e., putting the ODUflex into the ODU2 tunnel for transport over the OTU3 link.

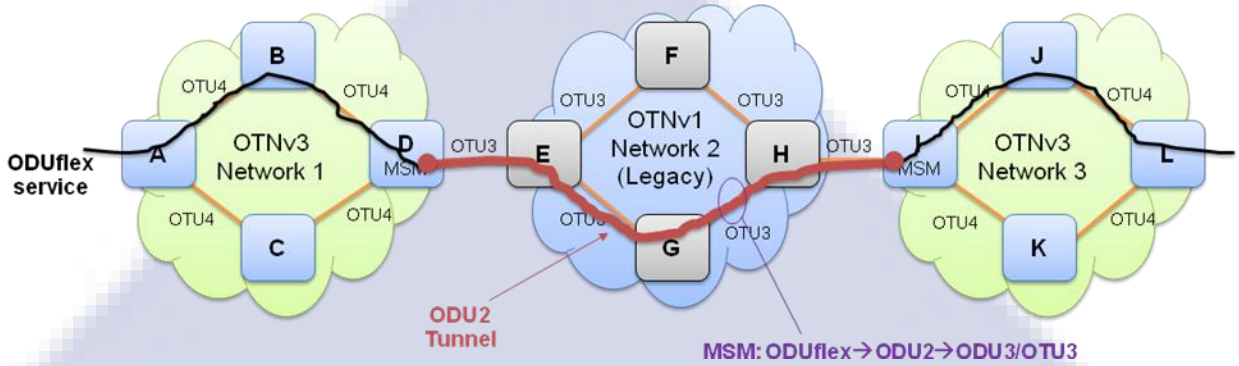

Figure 6: Example of Multistage Multiplexing (MSM) – ODUflex through Legacy Network

It is also possible that ODU tunnels may be established due to service provider preference. A service provider may choose to tunnel lower-order ODU services through higher-order ODU tunnels across part of his OTN network for scalability reasons in order to reduce the number of ODU services to manage. Or the service provider may be offering a transparent ODU tunneling service.

4.1.1 Routing for OTN Multiplexing Hierarchy

The support for single-stage and multi-stage multiplexing within an OTN network is advertised by explicitly identifying the multiplexing hierarchy supported on each link. From this information, a path through the OTN network can be determined including the need for single-stage or multi-stage multiplexing.

Figure 7 shows the link advertisements from the network example given in Figure 6. The nodes at the edges of the OTNv3 subnetwork (i.e., Nodes D and I) support the ability to perform multi-stage multiplexing. This is shown in the bandwidth advertisement of the OTU3 link from Node D to Node E and in the OTU3 link from Node I to Node H.

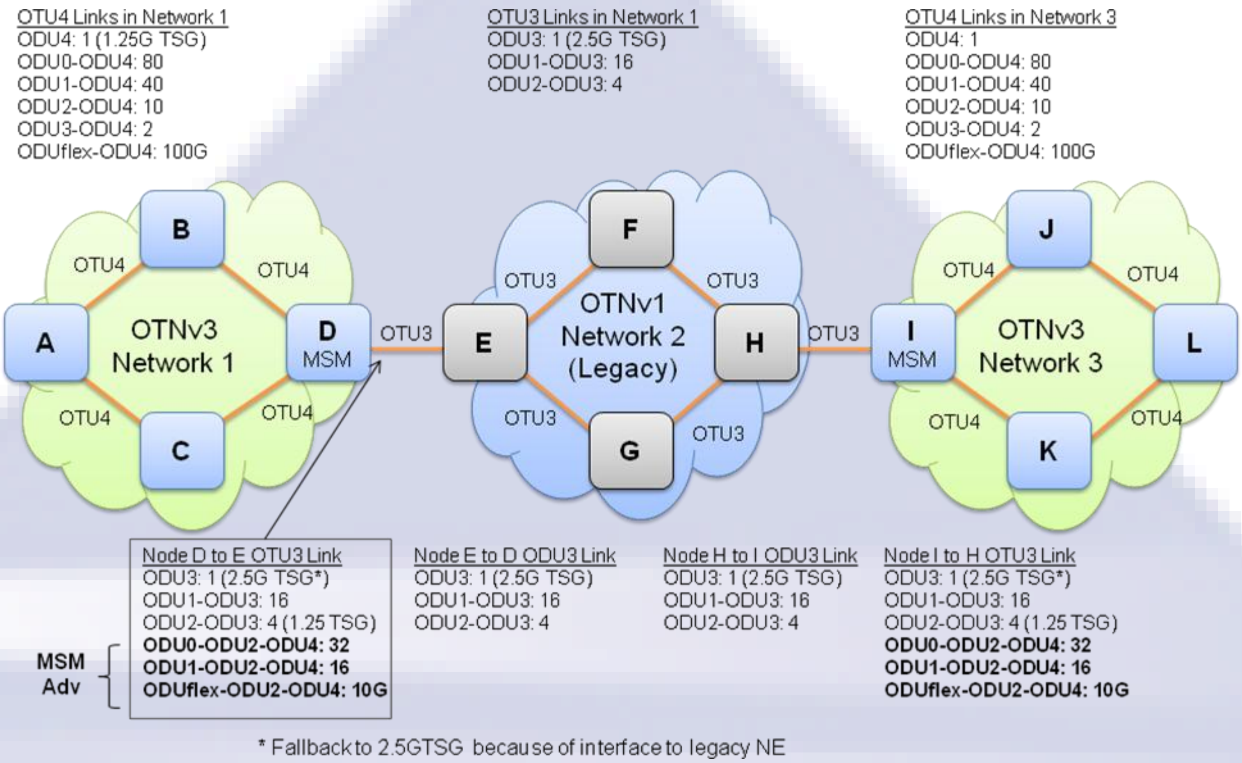

Figure 7: Example of Multistage Multiplexing Bandwidth Advertisement

To route an ODUflex service between Nodes A and L in Figure 7, it is necessary to use multi-stage multiplexing at Nodes D and I because the nodes in the legacy subnetwork (Nodes E-H) cannot switch at the ODUflex rate. The routing could plan an ODU2 tunnel through the legacy network, starting and terminating at Nodes D and I. Then the ODUflex service could be routed through the OTNv3 Networks 1 and 3 and carried through Network 2 via the ODU2 tunnel.

Additionally, the ODU2 tunnel can be advertised in the OTN routing instance as a link that connects Nodes D and I. The bandwidth that would be advertised would be the ODU2 bandwidth minus the bandwidth used by the ODUflex service. This would allow other service requests to be routed over the ODU2 tunnel.

4.1.2 Signaling for OTN Multiplexing Hierarchy

The signaling of multi-stage multiplexing requires the establishment of a separate signaling session for each intermediate ODUj layer. For example, for the two-stage multiplexing case of an ODUflex over an intermediate ODU2 there would be a signaling session for the ODUflex and a separate signaling session for the ODU2.

The example described in the previous section is shown in Figure 1 from signaling point of view. The ODU2 tunnel would be signaled between Nodes D and I with intermediate signaling nodes at E, F and H within Network 2. To allow other ODU services (bandwidth less than ODU2), a direct signaling adjacency is formed between Nodes D and I to support services mapped to this tunnel.

The ODUflex would be signaled from A, B, D, I, J and K. At Node D, the ODUflex is mapped into the ODU2 tunnel so the signaling would proceed directly to Node I. The ODUflex signaling would not be visible to the nodes in Network 2.

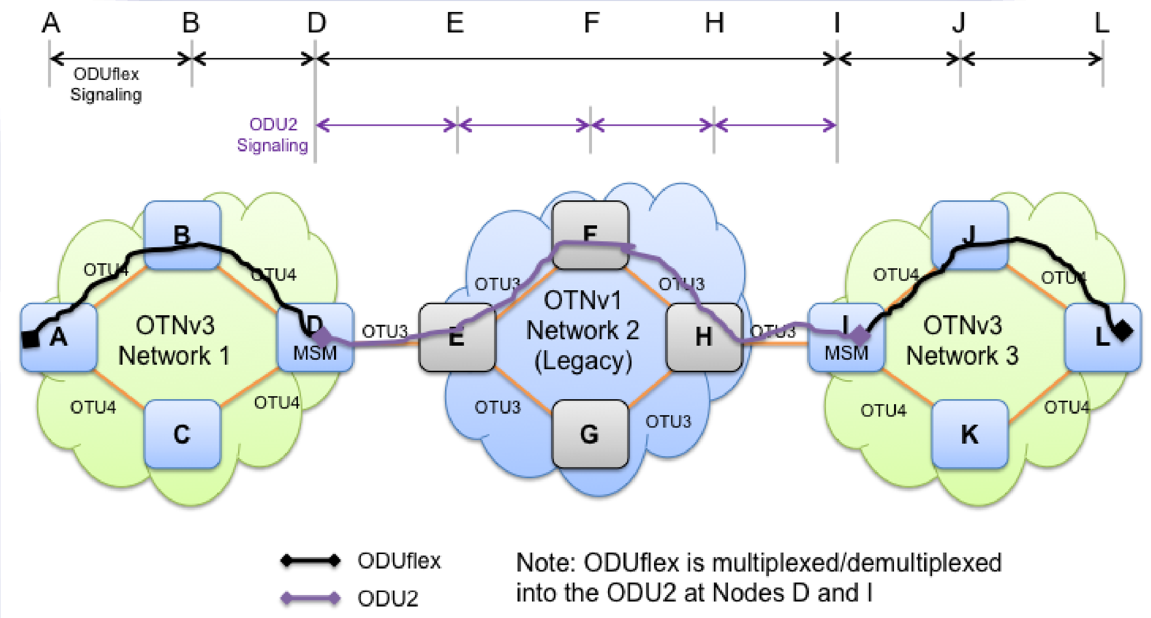

Figure 8: Multi-Stage Multiplexing Signaling Example

The signaling session for the intermediate ODUj layer in a multi-stage multiplexing scenario could be established prior to the signaling of the LO-ODU service that would use the ODUj as a tunnel. Or the signaling session for the ODUj could be automatically triggered by the LO-ODU service signaling session.

When signaling an ODU that will be used to multiplex lower-order ODU services, information about the desired OTN adaptation may be necessary. The TSG may need to be explicitly specified to coordinate the TSG provisioning at the head and tail ends. In certain cases it may also be necessary to request the specific LO-ODU rates that must be supported by the ODU tunnel. This information may also be use by the penultimate hop to select the correct outgoing link that the tail-end node can demultiplex to the required LO-ODUs.

The procedure for signaling the ODU service layer and intermediate ODU tunnel layer(s) follows the procedures described in the OIF Multilayer Amendment [OIF-ENNI-ML-AM-01.0]. Further details, including abstract messages and message flows, are provided in Section 5.1.

4.2 Hitless Adjustment of ODUflex

The transport plane procedure for the hitless adjustment of ODUflex is described in [G.7044]. ODUflex resizing only applies to ODUflex(GFP) services. Procedures are defined for the hitless increase and decrease of ODUflex bandwidth. In order to support ODUflex resizing all nodes along the path must support the procedures described in G.7044. This is different than VCAT resizing where it was sufficient for only the end nodes to support the resize procedure.

The ODUflex resize procedure operates in both directions at the same time. There are two main parts to the ODUflex resizing procedure: Link Connection Resize (LCR) and Bandwidth Resize (BWR). LCR operates at the link level (i.e., the link connection) between the upstream and downstream nodes. This part adjusts each link connection to use the newly added (in case of bandwidth increase) or to no longer use the deleted (in case of bandwidth decrease) tributary slots. Although the number of available tributary slots will change during LCR, LCR does not change the bandwidth rate.

BWR operates end-to-end to increase or decrease the bandwidth rate. For bandwidth increase operations, LCR completes first to ensure that the new tributary slots are available before BWR increases the bandwidth rate. For bandwidth decrease, BWR decreases the bandwidth prior to LCR deleting the tributary slots no longer in use.

The ODUflex resizing procedure requires that at least one tributary slot change end-to-end. In the case of bandwidth increase, at least one tributary slot must be added. For bandwidth decrease, at least one tributary slot must be deleted. Note that for ODUflex(GFP), the number of tributary slots is constant end-to-end.

In this implementation agreement, ODUflex resizing will be triggered by the control plane during bandwidth adjustment signaling. The control plane is also responsible for identifying the added or deleted tributary slots and providing this information to the transport plane. The control plane must ensure that the added or deleted tributary slots are consistent at both ends of a link.

For ODUflex bandwidth decrease, the highest number tributary slot cannot be removed. This tributary slot is used to transmit the GMP mapping information for multiplexing the ODUflex into a HO-ODU. The highest number tributary slot can only be removed when the ODUflex service is deleted.

For example, an ODUflex services uses 4 tributary slots: TS1-4. The bandwidth is reduced to 1 tributary slot. In this example, the ODUflex would retain the highest tributary slot, TS4, and remove tributary slots 1-3.

The ODUflex bandwidth adjustment procedure is rate limited to prevent transport plane buffer overflow. The rate limit is set to 512000 kbit/s^2. Implementations that support the ODUflex bandwidth resizing procedure must implement timers sufficient to account for this delay.

Upon detecting an error condition, the ODUflex bandwidth adjustment procedure will abort the increase or decrease operation and revert back to the original state. A notification of the ODUflex resizing error is sent to the management system. Additionally, the management plane or control plane may also initiate an abort to cancel an in-progress ODUflex bandwidth adjustment procedure.

If the transport plane detects an error during ODUflex resizing or if there is a management/control plane abort request, the local signaling controller shall initiate a graceful teardown of the signaling session associated witht the new bandwidth request. This teardown may be initiated by any node along the ODUflex path and normal graceful teardown procedure shall be used.

A node that receive a teardown request for the new bandwidth session while the original session is intact shall interpret this as an abort request. That node shall propagate the teardown request and should issue an abort to the local ODUflex bandwidth adjustment component (see Figure 9).

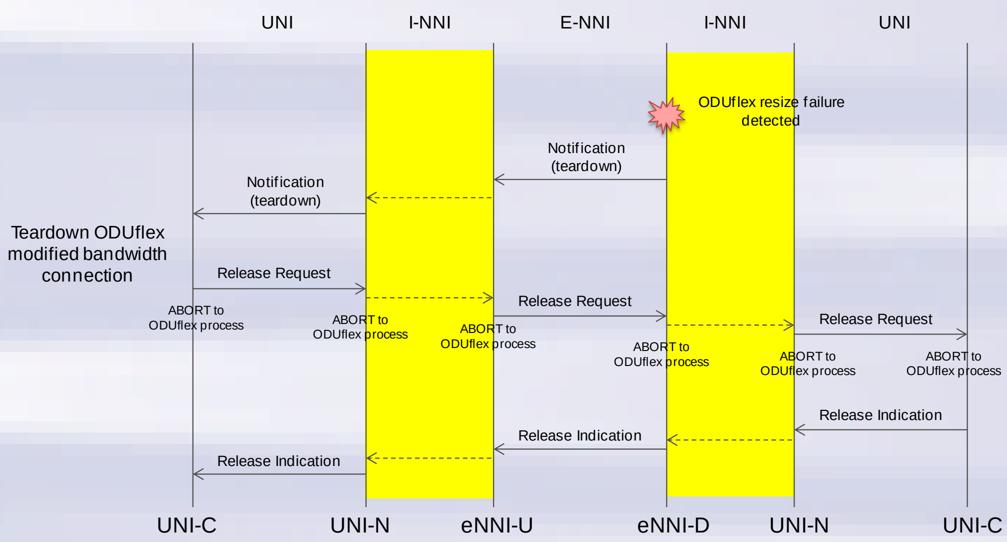

Figure 9: ODUflex Error Handling

4.3 OTNv1 Transport Plane

The OTNv3 extensions can be used to support transport plane links that are based on OTNv1 [G.709Ed2]. OTNv1 allowed the TPN to be explicitly configured or automatically detected through the Auto-MSI function. The Auto-MSI feature could also be used to establish a Tributary Slot Group. In this application, a customer orders a set of tributary slots (i.e., the Tributary Slot Group) and manages the specific structure and configuration of those tributary slots without involvement of the service provider.

The Auto-MSI function has been deprecated in [G.798Ed3]. Because of this, the OTNv3 signaling and routing extensions do not support the Auto-MSI function. OTNv3 requires that the TPN be explicitly configured, regardless if the transport plane supports OTNv1 or OTNv3. Additionally, the use of Auto-MSI for defining a Tributary Slot Group is not supported in the OTNv3 amendment. (Note that functionality similar to Tributary Slot Group can be obtained through the use of hierarchical ODU multiplexing under OTNv3.)

If Auto-MSI behavior is required for backwards compatibility, then the control plane should be configured to support the UNI/E-NNI 2.0 Signaling and Routing extensions (without the OTNv3 extensions) for those interfaces.

4.4 Multilayer Support for OTNv3

OTNv3 supports the multiplexing of LO-ODUs into a HO-ODU (see Section 4.1). This multiplexing of ODUs is a special case of multilayer where the client and server are both from the same switching technology and is referred to as intra-layer in this section. The intralayer impact for ODU multiplexing is described in this section.

4.4.1 Resource Relationships

The resource relationship for intralayer ODU multiplexing is typically m:1. Multiple client ODUs (i.e., LO-ODUs) are multiplexed into a single server ODU (i.e., HO-ODU). For example, there can be an 3xODU0, ODU1, and ODUflex(3) multiplexed into an ODU2.

The resource relationship can also be 1:1 under the case of an ODUflex that uses the full bandwidth of the HO-ODU.

Although the OTN transport plane supports the 1:n relationship via ODU VCAT, support for the 1:n relationship (and ODU VCAT) is out-of-scope of this amendment.

4.4.2 Intra-layer Routing

The ODU multiplexing hierarchy supported on an interface is described using the local client adaptation attribute and potentially supplemented by a transitional link.

The local client adpatation attribute contains the following information for intra-layer ODU multiplexing:

- Multiplexing hierarchy as a hierarchical list of ODU signal types
- Link capacity (i.e., bandwdith) per ODU signal type at each level of the multiplexing hierarchy
- Local connection type per ODU signal type at each level of the multiplexing hierarchy

Transitional links may be used to provide optional supplemental information such as weight, resource class, availability, etc. per ODU signal type at each level of the multiplexing hierarchy.

Abstract node, abstract link or pseudo-node representations are not supported as mechanisms to describe the intra-layer ODU multiplexing hierarchy across the E-NNI. However, this does not prohibit the use of these mechanisms within an I-NNI if they are wholly contained within the I-NNI domain.

4.4.3 Identifiers and Name Resolutions

TNAs are used to represent the set of access points that may be used to support the LO-to-HO ODU intra-layer adaptation. It is recommended that the same TNA value be used to represent the set of ODU signal types and multiplexing hierarchy supported on an OTU interface.

4.4.4 Intra-layer Processing Abstract Model

The procedures for ODU intra-layer setup is the same as the interlayer setup procedures described in Section 6.6 of [OIF-ENNI-ML-AM-01.0]. Additional detail for ODU intra-layer setup is provided in Sections 4.1 and 5.1.

4.4.5 Call Modifications

Call modification is supported to change the bandwidth of an ODUflex(GFP) service.

OTNv3 does not support intra-layer call modifications as call modification is only supported on the ODUflex(GFP) signal type. This signal type is always at the top of the ODU multiplexing hierarchy (i.e., it is always a LO-ODU and never serves as a HO-ODU).

Interlayer call modification is supported for ODUflex(GFP). In this case, the ODUflex call modification would be triggered by a cleint layer (e.g., Ethernet) call modification.

Additional details of ODUflex call modification is provided in Sections 4.2 and 5.2.

4.4.6 Interlayer Signaling Interface Abstract Attributes

Table 10 shows the update to the Interlayer Signaling Interface Abstract Attributes, as given in Table 11 of [OIF-ENNI-ML-AM-01.0].

Table 10: Update to [OIF-ENNI-ML-AM-01.0] Table 11, Interlayer Signaling Abstract Attributes

Abstract Attributes		Reference	Cross Ref
Service Attributes	Encoding Type	[OIF-ENNI-ML-AM-01.0]	[OIF-ENNI-Sig-02.0] Section 10.2.4.2

		Section 6.6.3.2.2 Section 4.4.6.1.1	[RFC7139]
	Switching Type	[OIF-ENNI-ML-AM-01.0] Section 6.6.3.2.3 Section 4.4.6.1.2	[OIF-ENNI-Sig-02.0] Section 10.2.4.3 [RFC7139]
	Technology specific traffic parameters	[OIF-ENNI-ML-AM-01.0] Section 6.6.3.2.4 Section 4.4.6.1.3	[OIF-ENNI-Sig-02.0] Section 10.2.4.4 [RFC7139]

4.4.6.1 Service-Related Attributes

4.4.6.1.1 Encoding Type

The Encoding Type specifies the encoding format of the signal to be transported. The encoding options are specified in [OIF-ENNI-Sig-02.0] and [RFC7139].

4.4.6.1.2 Switching Type

This indicates the type of switching that should be performed on a particular link. The switching type options are specified in [OIF-ENNI-Sig-02.0] and [RFC7139].

4.4.6.1.3 Traffic Parameters

Traffic parameters are specified in [OIF-ENNI-Sig-02.0] and [RFC7139].

5 Abstract Messages and Flows

5.1 Abstract Messages and Flows for Multi-stage Multiplexing

OTN multi-stage multiplexing uses the multilayer procedures described in [OIF-ENNI-ML-AM-01.0]. In this section, it is assumed that there is an ODUj that is multiplexed into an ODUk (j<k). The relationship between the ODUj and ODUk can be considered a client-server relationship. The ODUk is the server layer and the ODUj is the client layer. Note that more than one layer of multiplexing is possible (e.g., an ODU0 into an ODU2 into an ODU4).

Each level of the multiplexing hierarchy is considered a separate ODU trail and signaled independently. The exception is a HO-ODUk that maps directly into an OTUk, traverses a single hop and terminates at both ends of the OTUk. In this case the HO-ODUk trail can be implicitly setup without requiring an explicit control plane session. If an ODU trail traverses multiple hops, then the ODU trail must be explicitly signaled.

Just like the multilayer scenarios, the server ODU trail (if it must be explicitly signaled) can be established before the signaling of the client ODU trail. Or the client ODU trail can trigger the setup of the server ODU trail.

An example of the latter case is shown in Figure 10 where an ODU0 session is established and will trigger an ODU2 tunnel between Nodes B and D. The steps in this figure are as follows:

1) An ODU0 call is initiated at Node A. Node A begins the signaling for the ODU0 connection.

2) Node B receives the ODU0 connection setup request and initiates the setup of the ODU2 server connection. The decision to trigger the creation of an ODU2 could be based on a local policy decision or could be a result of information in the ODU0 signaling (e.g., multilayer ERO)

   The signaling-based discovery mechanism is used to establish a link in the OTN topology to represent the ODU2 tunnel. Node B inserts its local discovery information such as the ODU2 link's local identifier and Node B's SC PC IP address and SC PC ID. (Additional discovery information may also be supplied.)

3) Node C processes and forwards the ODU2 connection setup request.

4) Node D processes and terminates the ODU2 connection setup request and responds with an ODU2 connection setup indication.

   The connection setup indication will contain Node D's information to support the signaling-based discovery of the ODU2 link. This information includes Node D's local identifier for the ODU2 link and Node D's SC PC IP address and SC PC ID. Node D has sufficient information to establish the ODU2 link. Node D may also advertise the ODU2 link in routing.
   Node C processes and forwards the ODU2 connection setup indication. The ODU2 matrix connection is established.

5) Node B terminates the ODU2 connection setup indication and the ODU2 connection is established. Node B may respond with an ODU2 connection setup confirmation (not shown in the figure).

   Node B establishes the ODU2 link based on the discovery information received in the ODU2 signaling session. The ODU2 tunnel is now available for use by ODU clients. Node B establishes a signaling adjacency with Node D to support the ODU2 link. Node B may also advertise the ODU2 link in routing.

   The ODU0 is routed over the ODU2 link. Node B continues the ODU0 signaling directly to Node D (bypassing Node C).

6) Node D processes and forwards the ODU0 connection setup request.

7) Node E terminates the ODU0 connection setup request and responds with an ODU0 connection setup indication.

8) Node D processes and forwards the ODU0 connection setup indication. It is at this point that the ODU0 is multiplexed into the ODU2 server connection.

9) Node B receives the connection setup indication from Node D and also binds the multiplexing of the ODU0 into the ODU2 server connection. Node B sends the ODU0 connection setup indication to Node A.

10) Finally Node A receives the ODU0 connection setup indication completing the process. The ODU0 connection is established. Node A may also send an ODU0 connection setup confirmation (not shown in the figure).

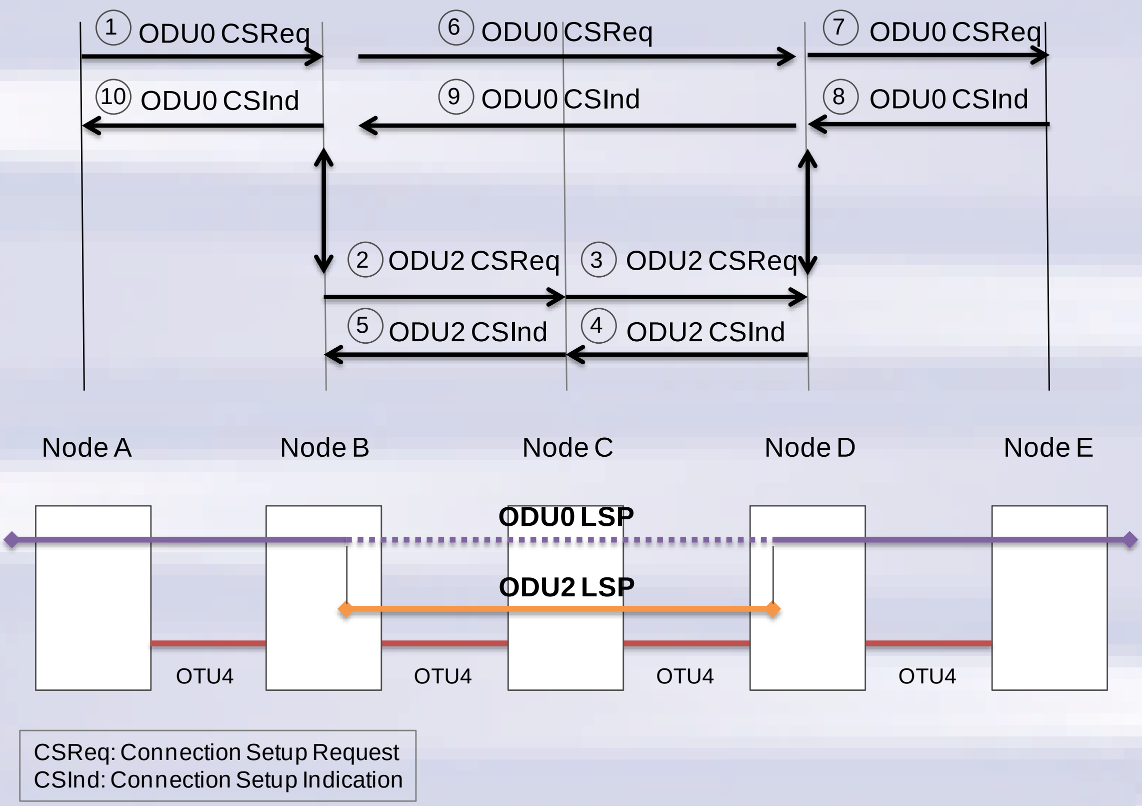

Figure 10: Multi-stage Multiplexing Signaling Flow Example

5.2 Abstract Messages and Flows for Hitless Adjustment of ODUflex

The hitless adjustment of ODUflex uses the non-disruptive call modification procedure defined in [OIF-UNI-02.0] and [OIF-E-NNI-Sig-02.0]. Specifically, the call and connection modification procedure is used to modify the bandwidth of the existing ODUflex connection. The abstract message flow for the hitless adjustment of ODUflex is shown in Figure 11.

When modifying the bandwidth of ODUflex, at least one tributary slot is added or deleted. On a given connection, only one operation, either a bandwidth increase or bandwidth decrease, can be in effect at a time. The simultaneous addition of tributary slots while deleting other tributary slots is not supported.

The connection modification procedure must not complete before the transport plane bandwidth adjustment has completed. Specifically, the connection modification confirmation message and completion notification to the management systems should not occur until after the transport plane operation is completed. Also for multilayer calls, the interlayer call modification indication should also wait until the transport plane operation completes. It is up to implementations to determine when the transport plane bandwidth adjustment operation is complete.

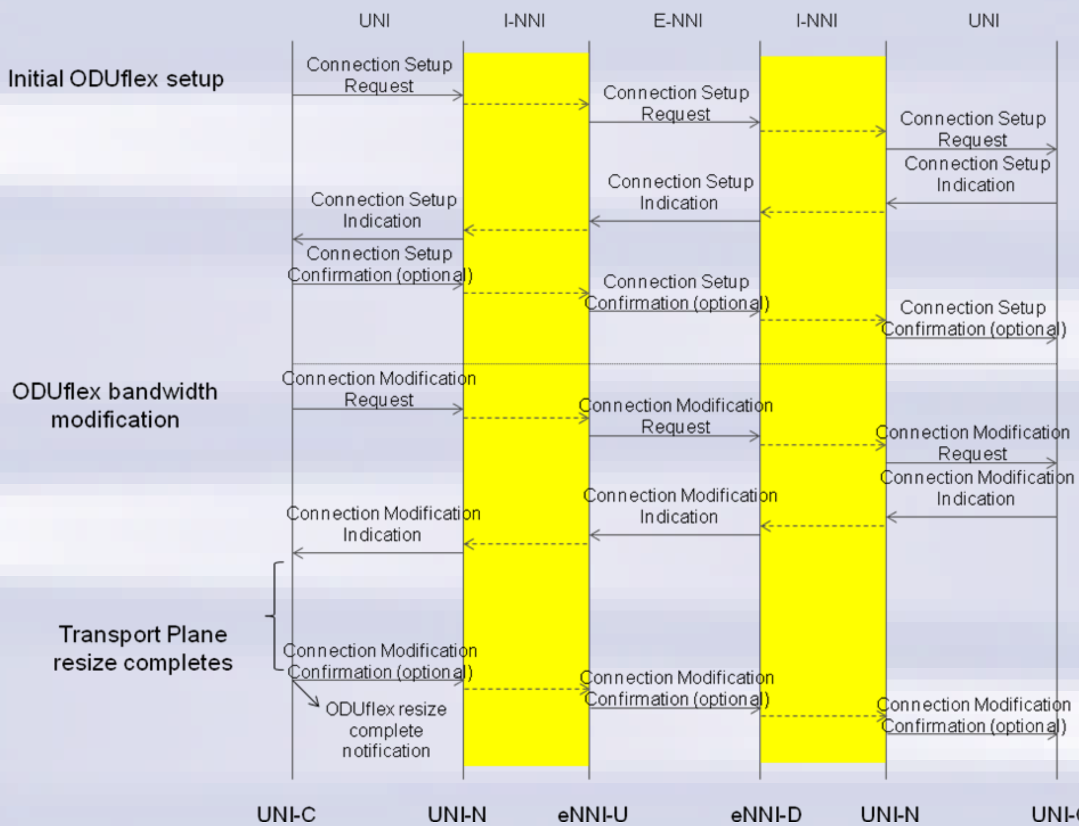

Figure 11: ODUflex Resize Abstract Message Flow

6 Security and Logging for OTNv3

6.1 Security

This amendment does not introduce new control plane protocols. Methods for securing control plane protocols are provided in [OIF-UNI-02.0], [OIF-E-NNI-Sig-02.0] and [OIF-E-NNI-OSPF-02.0]. Security extensions for the UNI and E-NNI 2.0 are provided in [OIF-

SECEXT-03.0]. Security for management interfaces to the network element is described in [OIF-SMI-01.0] and [OIF-SMIADD-02.1].

6.2 Logging

If control plane logging with syslog as specified in [OIF-SYSLOG-01.1] is implemented, then the protocol errors listed in this implementation agreement SHOULD be logged with a Severity higher (i.e., lower-numbered) than Informational. Also, to verify the correct operation of OTNv3 routing and signaling, the PROT@26041 logging capability can be turned on at one or more network elements implementing the OTNv3 capability to generate a secure, time-stamped trace of control plane traffic.

Further requirements for logging protocol messages are given in [OIF-E-NNI-Sig-02.0], [OIF-E-NNI-OSPF-02.0] and [OIF-ENNI-ML-AM-01.0].

7 Compatibility with UNI and E-NNI

7.1 Compatibility with UNI Signaling 2.0

UNI Signaling 2.0 [OIF-UNI-02.0] supports G.709 Edition 2 (OTNv1) [G.709Ed2]. UNI 2.0 supports switching only at the ODU1, ODU2 and ODU3 rates. Additionally, UNI 2.0 supports the label format, OTN traffic parameter and OTN switching capability defined in [RFC4328]. These OTN encodings are different that what is supported in this amendment.

An implementation MAY support backward compatibility with UNI Signaling 2.0 OTNv1[1]. Backward compatibility is achieved by learning the OIF UNI release and features supported by the neighbor across the UNI. The neighbor UNI protocol release and features can be automatically learned or manually configured. The automatic learning of neighbor capabilities is outside the scope of this Implementation Agreement

For a network element that supports UNI 2.0 and the OTNv3 extensions and that supports backward compatibility with UNI 2.0 OTNv1, then:

- If it detects that a neighbor supports UNI 2.0 but not the OTNv3 extensions, then the network element MUST revert to UNI 2.0 without OTNv3 extensions to that neighbor. Only OTNv1 ODU rates are supported.
- Otherwise, the OTNv3 extensions in this document SHOULD be used. In this case, OTNv3 ODU rates are supported.

[1] It is expected that most OTN UNI implementations will support OTNv3. Therefore, backward compatibility with UNI 2.0 OTN is not mandatory.

Operational deployment procedure must ensure that a network element that supports OTNv3 UNI signaling without UNI 2.0 OTNv1 backward compatibility will interface only to other OTNv3-capable network elements across OTN-capable UNIs.

Implementations that support the OTNv3 extensions may support the ability to disable OTNv3 on some UNI interfaces. In this case, the network element should revert to the OTNv1 support as specified in UNI 2.0 for those interfaces.

UNI 2.0 features other than OTN are not affected by this amendment and are fully compatible with this amendment.

7.2 Compatibility with E-NNI Signaling 2.0

Similar to UNI Signaling 2.0, E-NNI Signaling 2.0 [OIF-E-NNI-Sig-02.0] supports G.709 Edition 2 (OTNv1) [G.709Ed2]. The OTN encodings for OTNv1 are different than the OTNv3 encodings supported in this amendment.

An implementation MAY support backward compatibility with E-NNI Signaling 2.0 OTNv1[2]. Backwards compatibility is achieved by learning the OIF E-NNI release and features supported by the neighbor across the E-NNI. The neighbor E-NNI protocol release and features can be automatically learned or manually configured. The automatic learning of neighbor capabilities is outside the scope of this Implementation Agreement.

For a network element that supports E-NNI Signaling 2.0 and the OTNv3 extensions and that supports backward compatibility with E-NNI 2.0 OTNv1, then:

- If it detects that a neighbor supports E-NNI 2.0 but not the OTNv3 extensions, then the network element MUST revert to E-NNI 2.0 without OTNv3 extensions to that neighbor. Only OTNv1 ODU rates are supported
- Otherwise, the OTNv3 extensions in this document SHOULD be used. In this case, OTNv3 ODU rates are supported.

Operational deployment procedures must ensure that a network element that supports OTNv3 E-NNI signaling without E-NNI 2.0 OTNv1 backward compatibility will interface only to other OTNv3-capable network elements across OTN-capable E-NNIs.

Implementations that support the OTNv3 extensions may support the ability to disable OTNv3 on some E-NNI interfaces. In this case, the network element should revert to the OTNv1 support as specified in E-NNI 2.0 for those interfaces.

E-NNI Signaling 2.0 features other than OTN are not affected by this amendment and are fully compatible with this amendment.

[2] It is expected that most OTN E-NNI implementations will support OTNv3. Therefore, backward compatibility with E-NNI 2.0 OTN is not mandatory.

7.3 Compatibility with E-NNI Routing 2.0

E-NNI Routing 2.0 [OIF-E-NNI-OSPF-02.0] supports G.709 Edition 2 (OTNv1) [G.709Ed2]. E-NNI Routing 2.0 supports the advertisements of the ODU1, ODU2 and ODU3 signal types. E-NNI Routing 2.0 uses the Switching Capability 100 (TDM).

If an OSPF area contains one or more routing controllers that support only E-NNI Routing 2.0 without the OTNv3 routing extensions, then all OTNv3-capable routing controllers in that same area MAY advertise both the E-NNI 2.0 OTN ISCD and the OTNv3 OTN ISCD. These ISCDs are distinguished by the Switching Capability. The E-NNI 2.0 OTN ISCD uses the Switching Capability 100 (TDM). The OTNv3 ISCD uses the Switching Capability 110 (TDM-OTN).

An E-NNI Routing 2.0 routing controller would ignore the OTNv3 advertisements as they would use an switching capability (110) unknown to the routing controller.

An implementation MAY support backward compatibility with E-NNI Routing 2.0 OTNv1[3]. Backward compatibility is achieved by the advertisement of both E-NNI 2.0 OTN ISCD and OTNv3 ISCD. An implementation may provide a flag to enable or disable advertisement of the E-NNI 2.0 OTN ISCD (Switching Capability =100).

If a routing controller has disabled the advertisement of the E-NNI 2.0 OTN ISCD but receives only E-NNI 2.0 OTN ISCD advertisements from another router, then the routing controller may use the E-NNI 2.0 OTN ISCD advertisement when establishing ODU1, ODU2 or ODU3 services.

Implementations that support the OTNv3 extensions may support the ability to disable OTNv3 on some E-NNI interfaces. In this case, the routing controller should revert to advertising the E-NNI 2.0 OTN ISCD (Switching Capability = 100) for these E-NNI interfaces. It is recommended that the routing controller not advertise the OTNv3 OTN ISCD, but if the OTNv3 OTN ISCD is advertised, then it must only advertise support for ODU1, ODU2 or ODU3 signal types (because the signaling controller will use the E-NNI 2.0 OTNv1 extensions and other ODU types are not supported).

When both E-NNI 2.0 ISCD and OTNv3 ISCD are advertised by an OSPF router, the content of the two must be consistent for the common rates and hierarchy supported by OTNv1 and OTNv3 (i.e., single-stage ODU1, ODU2 and ODU3).

When both E-NNI 2.0 ISCD and OTNv3 ISCD are received from an OSPF router, the content of the two are expected to be the same for the common ODU rates and hierarchy supported by OTNv1 and OTNv3 (i.e., single-stage ODU1, ODU2 and ODU3). Therefore, an E-NNI 2.0 OTNv3 OSPF router SHOULD ignore the E-NNI 2.0 ISCD information.

[3] It is expected that most OTN E-NNI implementations will support OTNv3. Therefore backward compatibility with E-NNI 2.0 OTN is not mandatory.

7.4 Compatibility with E-NNI Multilayer 1.0

E-NNI Multilayer 1.0 features are fully compatible with this amendment.

The OTNv3 amendment builds upon the E-NNI Multilayer 1.0 extensions, especially in the signaling of ODU hierarchy. The OTNv3 extensions are fully compatible with the E-NNI Multilayer 1.0 extensions. Where needed, the Multilayer objects have been updated to support OTNv3-specific attributes such as layer and adaptation definitions.

7.5 Compatibility with UNI Signaling 1.0R2 and E-NNI Signaling 1.0

There are no backwards compatibility issues with UNI Signaling 1.0R2 [OIF-UNI-01.0-R2] and E-NNI Signaling 1.0 [OIF-ENNI-SIG-01] as these interfaces do not support the OTN (ODU) switching layer.

If the UNI or E-NNI neighbor supports the OIF 1.0 signaling protocol, then the local network element MUST not attempt to establish OTN (ODU) services over the interface.

Neighbor UNI and E-NNI protocol release and features can be automatically learned or manually configured. The automatic learning of neighbor capabilities is outside the scope of this Implementation Agreement.

7.6 Compatibility with E-NNI Routing 1.0

There are no backwards compatibility issues with E-NNI Routing 1.0 [OIF-E-NNI-OSPF-01.0] as E-NNI Routing 1.0 does not support the OTN (ODU) switching layer.

Network elements supporting E-NNI Routing 1.0 do not advertise OTN links nor support path computation for ODU services. These network elements should ignore OTN advertisements received from network elements that support E-NNI Routing 2.0 (with or without the OTNv3 extensions).

8 References

8.1 Normative references

8.1.1 IETF

[RFC2119] Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels“, IETF RFC

[RFC4328] Papadimitriou, D., Ed., Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control, RFC 4328, January 2006.

[RFC7062] Zhang, F., Ed., Li, D., Li, H., Belotti, S., and Ceccarelli, D., "Framework for GMPLS and PCE Control of G.709 Optical Transport Networks," RFC 7062, November 2013.

[RFC7096] Belotti, S., Ed., Grandi, P., Ceccarelli, D., Caviglia, D., Zhang, F., Li, D., "Evaluation of existing GMPLS encoding against G.709v3 Optical Transport Networks (OTNs)," RFC 7096, January 2014.

[RFC7138] Ceccarelli, D., Ed., Caviglia, D., Zhang, F., Li, D., Beloti, S., Grandi, P., Rao, R., Pithewan, K. Drake, J., "Traffic Engineering Extensions to OSPF for GMPLS Control of Evolving G.709 Optical Transport Networks", RFC 7138, March 2014.

[RFC7139] Zhang, F., Ed., Zhang, G., Belotti, S., Ceccarelli, D., and Pithewan, K., "GMPLS Signaling Extensions for Control of Evolving G.709 Optical Transport Networks," RFC 7139, March 2014.

8.1.2 ITU-T

[G.709Ed4] ITU-T Rec. G.709 (02/2012), Interfaces for the Optical Transport Network (OTN).

[G.798Ed4] ITU-T Rec. G.798 (10/2010), Characteristics of Optical Transport Network Hierarchy Equipment Functional Blocks.

[G.7044] ITU-T Rec G.7044 (10/2011), Hitless Adjustment of ODUflex(GFP) (HAO).

[G.7044Amd1] ITU-T Rec G.7044 Amendment 1 (02/2012), Hitless Adjustment of ODUflex(GFP) (HOA) Amendment 1.

8.1.3 OIF

[OIF-UNI-01.0-R2] OIF Implementation Agreement OIF-UNI-01.0-R2, "User Network Interface (UNI) 1.0 Signaling Specification, Release 2: Common Part", February 2004

[OIF-ENNI-SIG-01] OIF Implementation Agreement OIF-E-NNI-Sig-01.0, "Intra-Carrier E-NNI Signaling Specification", February 2004

[OIF-E-NNI-OSPF-01.0] OIF Implementation Agreement, "External Network-Network Interface (E-NNI) OSPF-based Routing - 1.0 (Intra-Carrier) Implementation Agreement", OIF-ENNI-OSPF-01.0, January 2007

[OIF-UNI-02.0] OIF Implementation Agreement, "User Network Interface (UNI) 2.0 Signaling Specification Common Part," OIF-UNI-02.0-Common, February 2008.

[OIF-E-NNI-Sig-02.0] OIF Implementation Agreement, "OIF E-NNI Signaling Specification", OIF-E-NNI-Sig-02.0, April 2009

[OIF-E-NNI-OSPF-02.0] OIF Implementation Agreement, “External Network-Network Interface (E-NNI) OSPF-based Routing - 2.0 (Intra-Carrier) Implementation Agreement”, OIF-ENNI-OSPF-02.0, July 2011.

[OIF-SECEXT-03.0] OIF Implementation Agreement, “Security Extension for UNI and E-NNI 2.0,” OIF-SEP-03.0, November 2010.

[OIF-SMI-01.0] OIF IA, “Security for Management Interfaces to Network Elements,” OIF-SMI-01.0, September 2003.

[OIF-SMIADD-02.1] OIF IA, “Addendum to the Security for Management Interfaces to Network Elements,” OIF-SMI-02.1, March 2006.

[OIF-SYSLOG-01.1] OIF IA, “OIF Control Plane Logging and Auditing with Syslog version 1.1,” OIF-SLG-01.1, July 2010.

[OIF-ENNI-ML-AM-01.0] OIF Implementation Agreement, “Multi-Layer Amendment to E-NNI 2.0 – Common Part,” OIF-ENNI-ML-AM-01.0, April 2013.

8.2 Informative references

8.2.1 ITU-T

[G.709Ed2] ITU-T Rec. G.709 (03/2003), Interfaces for the Optical Transport Network (OTN)

[G.709Ed3] ITU-T Rec. G.709 (12/2009), Interfaces for the Optical Transport Network (OTN).

[G.709Amend1] ITU-T Rec G.709 Amendment 1 (07/2010), Interfaces for the Optical Transport Network (OTN), Amendment 1.

[G.709Amend2] ITU-T Rec G.709 Amendment 2 (04/2011), Interfaces for the Optical Transport Network (OTN), Amendment 2.

[G.798Ed3] ITU-T Rec. G.798 (12/2006), Characteristics of Optical Transport Network Hierarchy Equipment Functional Blocks.

9 Appendix A: OTN Overview

9.1 OTN Types and Bit Rates

Table 11, Table 12 and Table 13 specifies the supported types, bit rates and bit rate tolerances at the OTU, ODU and OPU layers, respectively.

Table 11 OTU Types and Bit Rate

OTU Type	OTU nominal bit rate (kbits/s) Exact (Approximate)	OTU bit-rate tolerance (ppm)
OTU1	255/238 x 2 488 320 (2 666 057.143)	±20
OTU2	255/237 x 9 953 280 (10 709 225.316)	±20
OTU3	255/236 x 39 813 120 (43 018 413.559)	±20
OTU4	255/227 x 99 532 800 (111 809 973.568)	±20

Table 12: ODU Types and Bit Rates

ODU Type	ODU nominal bit rate (kbits/s) Exact (Approximate)	ODU bit-rate tolerance (ppm)
ODU0	1 244 160 (1 244 160.000)	±20
ODU1	239/238 x 2 488 320 (2 498 775.126)	±20
ODU2	239/237 x 9 953 280 (10 037 273.924)	±20
ODU3	239/236 x 39 813 120 (40 319 218.983)	±20
ODU4	239/227 x 99 532 800 (104 794 445.815)	±20
ODU2e	239/237 10 312 500 (10 399 525.316)	±100
ODUflex(CBR)	239/238 x Client signal bit rate	Client signal bit rate tolerance, max ±100
ODUflex(GFP)	Configured bit rate	±100

Table 13: OPU Types and Bit Rate

OPU Type	OPU nominal bit rate (kbits/s) Exact (Aprroximate)	OPU bit-rate tolerance (ppm)
OPU0	238/239 x 1 244 160 (1 238 954.310)	±20
OPU1	2 488 320 (2 488 320.000)	±20
OPU2	238/237 x 9 953 280 (9 995 276.962)	±20
OPU3	238/236 x 39 813 120 (40 150 519.322)	±20
OPU4	238/227 x 99 532 800 (104 355 975.330)	±20
OPU2e	238/237 10 312 500 (10 356 012.658)	±100
OPUflex(CBR)	Client signal bit rate	Client signal bit rate tolerance, max ±100
OPUflex(GFP)	238/239 x ODUflex signal rate	±20

9.2 ODUflex(GFP) Bit Rates

ODUflex(GFP) supports a discrete set of bit rates as shown in Table 14.

Table 14: ODUflex(GFP) Bit Rates

# TS	Rate (kbps)	# TS	Rate (kbps)	# TS	Rate (kbps)	# TS	Rate (kbps)
1	1249177.230	21	26343877.434	41	53360152.453	61	79389495.113
2	2498354.460	22	27598347.788	42	54661619.586	62	80690962.246
3	3747531.690	23	28852818.142	43	55963086.719	63	81992429.379
4	4996708.920	24	30107288.496	44	57264553.852	64	83293896.512
5	6245886.150	25	31361758.850	45	58566020.985	65	84595363.645
6	7495063.380	26	32616229.204	46	59867488.118	66	85896830.778
7	8744240.610	27	33870699.558	47	61168955.251	67	87198297.911
8	9993417.840	28	35125169.912	48	62470422.384	68	88499765.044
9	11290233.186	29	36379640.266	49	63771889.517	69	89801232.177
10	12544703.540	30	37634110.620	50	65073356.650	70	91102699.310
11	13799173.894	31	38888580.974	51	66374823.783	71	92404166.443
12	15053644.248	32	40143051.328	52	67676290.916	72	93705633.576

13	16308114.602	33	42948415.389	53	68977758.049	73	95007100.709
14	17562584.956	34	44249882.522	54	70279225.182	74	96308567.842
15	18817055.310	35	45551349.655	55	71580692.315	75	97610034.975
16	20071525.664	36	46852816.788	56	72882159.448	76	98911502.108
17	21325996.018	37	48154283.921	57	74183626.581	77	100212969.241
18	22580466.372	38	49455751.054	58	75485093.714	78	101514436.374
19	23834936.726	39	50757218.187	59	76786560.847	79	102815903.507
20	25089407.080	40	52058685.320	60	78088027.980	80	104117370.640

10 Appendix B: List of companies belonging to OIF when document is approved

Acacia Communications
ADVA Optical Networking
Agilent Technologies & Co.
Alcatel-Lucent
Altera
AMCC
Amphenol Corp.
Anritsu
Applied Communication Sciences
AT&T
Avago Technologies Inc.
Broadcom
Brocade
Centellax, Inc.
China Telecom
Ciena Corporation
Cisco Systems
ClariPhy Communications
Coriant
Cortina Systems
CPqD
Department of Defense
Deutsche Telekom
Emcore
Ericsson
FCI USA LLC
Fiberhome Technologies Group
Finisar Corporation
Fujikura
Fujitsu
Furukawa Electric Japan
Google
Hewlett Packard
Hitachi
Hittite Microwave Corp
Huawei Technologies
IBM Corporation

Infinera
Inphi
Intel
JDSU
Juniper Networks
Kaiam
Kandou
KDDI R&D Laboratories
Kikaua, LLC
LeCroy
LSI Corporation
Luxtera
M/A-COM Technology Solutions, Inc.
Marben Products
Marvell Technology
Mellanox Technologies
Metaswitch
Microsoft Corporation
Mindspeed
Mitsubishi Electric Corporation
Molex
MoSys, Inc.
MultiPhy Ltd
NEC
NeoPhotonics
NTT Corporation
Oclaro
Optelian
Orange
PETRA
PMC Sierra
QLogic Corporation
Ranovus
Semtech
Skorpios Technologies
Sumitomo Electric Industries
Sumitomo Osaka Cement
TE Connectivity
Tektronix
TELUS Communications, Inc.

TeraXion
Texas Instruments
Time Warner Cable
TriQuint Semiconductor
US Conec
Verizon
Xilinx
Xtera Communications
Yamaichi Electronics Ltd.
ZTE Corporation