


The New Ethernet-Driven OTN



This paper details how emerging Ethernet transport requirements have shaped the evolution of new OTN capabilities from the 2001 OTN V1 to the 2010 OTN V3, and how these capabilities support existing and new applications.

OTN V1 and V2 – History and Background

ITU SG 15 initially standardized the G.709 Optical Transport Network (OTN) in February 2001. The goal was to define the rates, formats, hierarchy and client mapping for a new circuit switched multiwavelength optical network. Rates were based on SDH/SONET. Initially, there were three OPUk/ODUk containers, OPU1 (2.488G STM-16 rate), OPU2 (238 /237 X 9.953G STM-64 rate) and OPU3 (238/236 x 39.813G STM-256 rate). The basic signal structure consists of three electrical levels, OPUk/ODUk/OTUk and a frame format of 3824 byte-wide columns by 4 rows. The OPUk/ODUk/OTUk structure is carried by a CWDM or DWDM wavelength known as an optical channel. The first defined client mapping was for STM-16/STM-64/STM-256 and GFP-F mapped packet frames.

G.709 Amendment 1, in November 2001, added ODUk multiplexing and virtual concatenation. For ODUk multiplexing, the OPUk payload area was partitioned into 2.488G tributary slots, and a new multiframe ODTUjk structure was defined ($j = \text{LO ODU}$ and $k = \text{HO ODU}$). A client is either asynchronously (AMP), bit-synchronously (BMP), or GFP mapped into a low-order OPUj / ODUj. The LO OPUj/ODUj is mapped to an ODTUjk multi-frame structure which is bit-synchronously multiplexed into N high order (HO) OPUk/ODUk tributary slots. The HO OPU overhead includes a Multiplex Structure Identifier (MSI) field, which provides the mapping between an HO OPU tributary slot and an ODTUjk/LO ODU port.

G.709 was updated in March 2003 to V2 with the addition of an ODUk protection switching communication channel. Between 2003 and 2007, G.709 V2 was updated with editorial changes which culminated in Amendment 2, dated November 2007. Since its introduction in 2001, the G.709 OPUk/ODUk/OTUk frame structure, which supports FEC, OAM, and protection capability, has been widely deployed as the electrical overhead for DWDM transport. The deployment of OTN switching was very limited. As Ethernet applications proliferated the requirement to interconnect 10 GbE switch ports by OTN brought to light in 2004 some issues with the SDH/SONET rate based OTN V1 of 2001. This started a series of Ethernet-driven events, which have culminated in the Ethernet optimized OTN V3 of 2010.

New Requirement – 10GBaseR (LAN PHY) Transparent Transport

When IEEE 802.3 standardized 10 GbE in 2002, the 802.3ae standard included both a SDH/SONET rate-based 10GBaseW WAN PHY and a 10 Gbps 10GBaseR LAN PHY. For the 10 GbE WAN PHY option, the 64B/66B coded MAC data can be mapped into a SDH VC-4-64c frame which can be mapped to an ODU2. This provided transparent transport, but there were two problems; the 10 GbE WAN PHY did not provide the full 10G MAC data rate, and its subsequent deployment in routers and Ethernet switches was very limited.

It is possible to GFP-F frame-map 10 GbE LAN PHY into a standard-rate ODU2 without the IPG, preamble, and SFD. The issue with this mapping is that it does not provide sufficient transport transparency to meet service provider requirements. By 2004, service providers were contributing 10 GbE LAN PHY/OTN transparent transport requirements to ITU SG 15. This foreshadowed two important paradigm shifts for OTN transport:

- Bit-transparent client transport
- Growing dominance of Ethernet as an OTN client

The problem was that the 2001 OTN V1 rates were based on SDH/SONET rates, and the 10 GbE client mapping was not transparent. Drivers behind the move to transparent client transport included:

- Multi-administrative domain applications
- 10 GbE transport for large customers
- Client OAM transparency, which simplifies fault propagation

Between 2004 and 2005, several proprietary solutions emerged and were proposed to ITU SG 15 for 10 GbE LAN PHY transparent OTN transport. These included an overspeed ODU2e and a semi-transparent ODU2 mapping, which GFP-F frame multiplexes 10 GbE PCS layer ordered sets in separate GFP-F frames. Initially the overspeed ODU2e was rejected due to its non-standard rate, and the semi-transparent approach was rejected because it was not transparent enough. Eventually a compromise was reached by which all 10 GbE over OTN approaches would be documented in ITU Supplement G.sup43, which was published in November 2006.

Impact of 40/100 GbE on OTN

Driven by applications and trends such as data centers, cloud computing and switch port capacity growth, IEEE 802.3 started work on a 100 GbE standard in 2006. During the initial study project, 40 GbE was added. Applying lessons learned from 10 GbE, ITU SG 15 immediately initiated work on PCS sub-layer bit-transparent transport of 40/100 GbE over OTN. Four issues were identified:

- Setting a new OPU4 rate to be an integral multiple of the OPU3 rate to simplify HO ODU4 multiplexing would result in a photonic transmission penalty
- The 2001 SDH based 40.151G OPU3 rate is insufficient for a bit transparent mapping of 40 GbE
- The OPU3 rate is insufficient for multiplexing 4 X 10 GbE LAN PHY.
- For the OPU4, a related requirement was the ability to multiplex ten overspeed ODU2e for 10 GbE LAN PHY.

If the OPU4 were a three-times multiple of the OPU3 rate, the OTU4 optical channel rate would be approximately 130G. If the OPU4 were sized to be the minimum rate for 100 GbE transparent and 10xODU2e, then the OTU4 optical channel rate would be approximately 112 Gb. The higher 130 Gb optical channel rate would gain a multiplexing simplification at the expense of reduced optical reach and optical impairment tolerance. The ITU SG 15 OTN WG eventually chose an OTU4 rate of 111.81G.

For 40 GbE PCS sub-layer mapped to an OPU3, block-code transcoding provided a solution. The transcoding process preserves the block code control codewords while reducing the aggregate bit rate. For 40 GbE mapped to an OPU3, it was decided to use a two-stage transcoding process. The first stage applies a 512B/513B block code to transcode the 64B/66B coded 40 GbE client, and the second stage applies a 1027B code, which provides a framing mechanism to locate the start of 513B blocks and an increase in Ethernet mean time to false packet acceptance. The transcoded 40 GbE PCS bit rate is reduced to a value that will fit in a standard rate OPU3.

Transcoding could not provide a solution for multiplexing four 10 GE LAN PHY circuits into an ODU3. Two possible solutions emerged:

- Use the G.sup43 semi-transparent 10 GbE mapping into a standard rate LO ODU2 which is multiplexed into a HO ODU3
- Multiplex overspeed ODU2e into an overspeed ODU3e. The ITU SG 15 agreed to document the overspeed ODU3e in an update to G.sup43.

Sub-ODU1 Clients and OTN V2 Amendment 3

In 2007, a new requirement emerged for OTN transport of clients with a bit rate less than the 2.488G OPU1. Paramount among these clients is GbE. Other clients that fall in this category are STM-1/OC-3, STM-4/OC-12, 1G/2G Fiber Channel, and DVB-ASI video (270M). This precipitated the need for:

- A smaller OPUk/ODUk container and OTN tributary slot size
- Enhanced mapping procedure for the case where there is a large disparity between the client rate and the LO ODU rate, or between the LO ODU rate and the rate of N HO ODU tributary slots.

The choice of the ODU4 rate also mandated development of an enhanced mapping and multiplexing procedure.

Driven by the importance of GbE, in 2008 ITU SG 15 agreed to define a new 1.25G ODU0. The preferred mapping was GbE PCS sub-layer bit transparent, which required GFP-T 8B/10B to 64B/66B transcoding, to fit the 1.25G GbE into the ODU0. For the GbE/ODU0, ITU SG 15 also developed an enhanced mapping procedure, which uses a sigma-delta-based justification distribution, allowing a larger timing differential between the GbE client and the OPU0 than OTN V1 AMP does. In April 2009, ITU SG 15 approved Amendment 3 to G.709 V2, which added the GbE ODU0, and the ODU2e and semi-transparent ODU2 for 10GbE, which were moved from G.sup43. The overspeed ODU3e, which supports four ODU2e, was added to G.sup43.

OTN V3

The introduction of the 1.25G ODU0 necessitated defining ODU2, ODU3 and ODU4 with 1.25G tributary slots to support ODU0 to HO ODUk multiplexing. This created a new ODU multiplex structure with 1.25G tributary slots. The ODTUjk was extended to a new ODTUk.ts structure (k = HO ODU, ts = number of required HO ODU tributary slots). These two OTN multiplex structures are distinguished by different OPU payload types

(PT 20 = ODTUjk, PT 21 = ODTUk.ts & ODTUjk). The concepts that were developed for enhanced mapping for GbE/ODU0 were extended to become a new GMP. GMP is superior to OTN V1 AMP in that it allows for a greater differential between client and LO OPU bit rates and between the LO ODUj bit rate and the ODTUk.ts bit rate, which is locked to the HO ODUk bit rate. For LO ODUj mapped to ODTUjk or ODTUk.ts, OTN V1/V2 (PT 20) only supports AMP, and OTN V3 (PT 20 and PT 21) supports both AMP and GMP. OTN V3 uses GMP for all new applications (such as LO ODUj multiplexed into ODU4), and only supports AMP for OTN V1/V2 legacy applications (ODU1 multiplexed into ODU2/ODU3 and ODU2 multiplexed into ODU3) and for OTN V2 Amendment 3 ODU0 multiplexed into ODU1.

The new 1.25G tributary slot OTN hierarchy, plus the addition of GMP, have facilitated support of a new flexible LO ODU structure, known as the ODUflex, along with a broad range of new clients. There are two types of ODUflex:

- CBR ODUflex for constant bit rate clients
- GFP-F ODUflex for GFP-F frame mapped clients

The nominal bit rate of the CBR ODUflex is 239/238 times the client bit rate. The client is bit-synchronously mapped to the CBR ODUflex, and GMP is used to map the CBR ODUflex into an integral number of HO ODUk tributary slots. The GFP-F ODUflex can be any bit rate, but the recommended bit rate is close to a multiple of the tributary slot size of the HO ODUk. For the GFP-F ODUflex, the packet client is GFP-F frame-mapped to the ODUflex, which is GMP-mapped into an integral number of HO ODUk tributary slots. GFP-F ODUflex can support all GFP-F packet clients, including Ethernet and IP/MPLS. The Ethernet client rate can be a sub-rate of the Ethernet port rate, and VLAN sets can be mapped to different GFP-F ODUflex.

ITU SG 15 approved G.709 V3 for the AAP comment process in October 2009, and final approval was achieved in January 2010. OTN V3 includes the above new functionality plus the ODU4, the mapping of 40/100 GbE to ODU3/ODU4, and ODU2e multiplexed into ODU3 and ODU4. Figure 1 illustrates the OTN V3 hierarchy and client support.

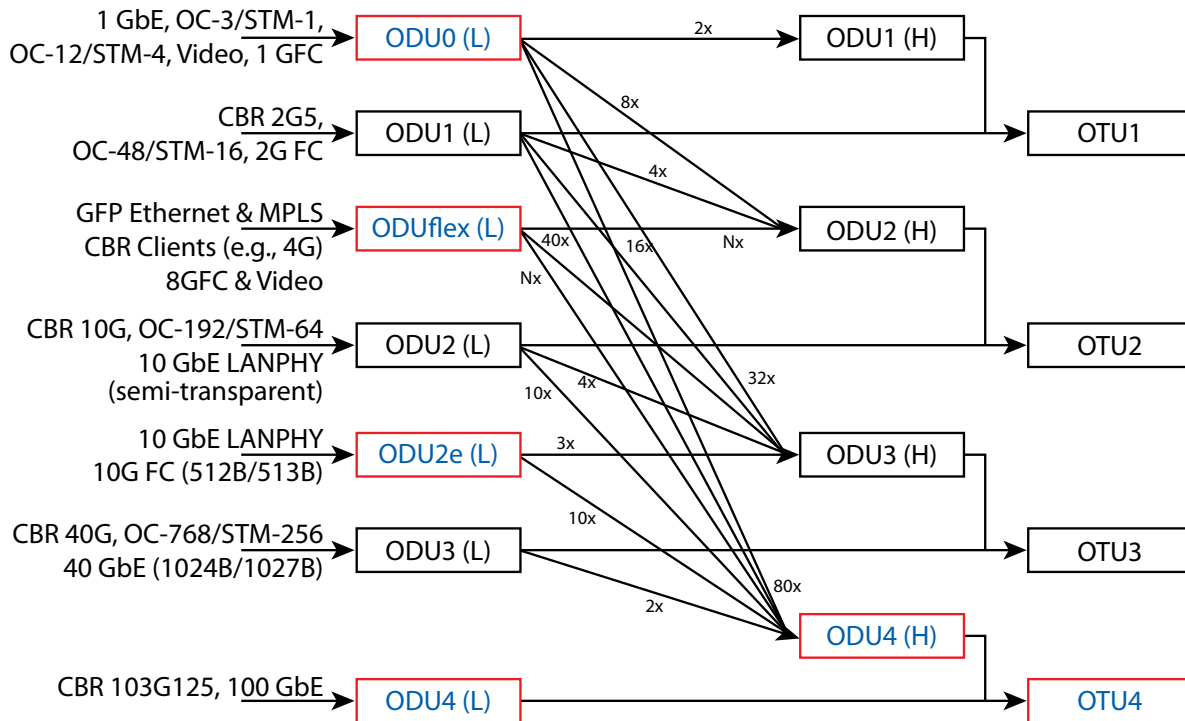


Figure 1: OTN V3 Multiplexing Hierarchy and Client Support

New Applications Enabled by OTN V3

The 2001 OTN V1 proposal was optimized for SDH/SONET transport and 100M/1Gb/10Gb Ethernet frame transport over SDH/SONET mapped to OTN. The new 2010 V3 OTN has been extended to 100G, and now supports 1G/10G/40G/100G Ethernet bit-transparent transport directly over OTN. OTN V3 also supports bit-transparent transport of storage (1G/2G/4G/8G/10G Fiber Channel), SDH/SONET (STM-1/OC-3/STM-4/OC-12 where SDH/SONET overhead is not terminated), and video (such as DVB-ASI). The OTN V3 ODUflex enables a new range of frame -mapped Ethernet and MPLS over OTN applications such as mapping VLANs or MPLS tunnels to an ODUflex.

OTN V3 will enable service providers to better address the expanding need for bit transparent Ethernet port transport (including client fault messages). A key advantage of OTN transport is that it provides virtually error-free transport with very low latency. OTN V3 also includes in service delay measurement. The attributes and new features of OTN V3 will enable service providers to better address existing and emerging applications.

Existing applications include:

- Carrier's Carrier and multiadministrative transport with very high QoS
- Transparent transport of high-capacity Ethernet ports for large customers
- OTN switching and grooming to ensure that wavelengths are fully packed
- Robust in-service OAM that goes beyond SONET/SDH to add OTN segment monitoring and protection

New/Emerging applications include:

- Data center interconnect for which OTN provides error free low latency transport for storage, high-performance computing and mission-critical cloud computing applications
- Optical bypass for backbone router interconnect
- High-performance MEF E-LINE and E-TREE (via OTN drop and continue) Ethernet service

7.0 References

[1] G.709/Y.1731, "Interfaces for the optical transport network (OTN)," 02/2001, ITU-T

[2] TD95R2 (PLEN/15), G.709 Issue 3.0 "Interfaces for the optical transport network (OTN)," 10/2009, ITU-T SG 15

Term	Definition
AAP	Alternative Approval Process
AMP	Asynchronous Mapping Procedure
BMP	Bit-synchronous Mapping Procedure
CBR	Constant Bit Rate
DWDM	Dense Wavelength Division Multiplexing
FEC	Forward Error Correction
GbE	Gigabit Ethernet
GMP	Generic Mapping Procedure
HO	High Order
IEEE	Institute of Electrical and Electronic Engineers
IP	Internet Protocol
ITU	International Telecommunications Union
LO	Low Order
MPLS	Multiprotocol Label Switching
MSI	Multiplex Structure Identifier

Term	Definition
OAM	Operations, Administration and Maintenance
ODU	Optical Data Unit
OPU	Optical Payload Unit
OSNR	Optical Signal-to-Noise Ratio
OTN	Optical Transport Network
OTU	Optical Transport Unit
PCS	Personal Communications Service
PHY	Open Systems Interconnection reference model-Physical Layer
QoS	Quality of Service
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
STM	Synchronous Transport Module
WAN	Wide-Area Network
WDM	Wavelength Division Multiplexing

© Copyright 2010 Fujitsu Network Communications Inc.
 FUJITSU (and design)™ are trademarks of Fujitsu Limited.
 All Rights Reserved. All other trademarks are the property of their respective owners.
 Configuration requirements for certain uses are described in the product documentation.

