

Next-Generation Optical Transport Networks

by Reza Vaez-Ghaemi, Ph.D.

Increasing deployments of broadband services demand optical transport networks that can cost-effectively scale. And, the emergence of high-speed Ethernet technologies is driving significant change in optical networks. Advances in transport technology, optical components, and optical network design are important enablers for this disruptive network change.

This paper introduces Ethernet-services transport over circuit-switched networks and then describes recent progress in high-speed Ethernet and optical transport network (OTN) technologies.

Introduction

The ever-increasing demand for broadband services has significantly contributed to the rising average return per user (ARPU). Services with the greatest impact include broadcast TV, video on demand (VoD), and mobile Internet. To be profitable, providing these services requires a transport technology with significantly lower cost-per-bit than traditional circuit-switched technology. This has driven the growing demand for Ethernet- and OTN/wavelength division multiplexing (WDM)-based solutions in metro and core/long-haul networks.

OTN reduces transport costs and delivers enhanced network and performance management functions. Forward error correction (FEC) algorithms improve the reach of transmission links, helping to reduce regenerators and/or optimize spectral efficiency. Additionally, an OTN “digital wrapper” includes many layers and components known from synchronous optical network/synchronous digital hierarchy (SONET/SDH) but at enhanced performance levels.

Optical Transport Networks

OTNs reduce operations and capital expenses and increase the scalability of WDM systems by offering the following features:

- simpler than SONET/SDH
- optimized for carrier WDM networks
- scalable for higher rates
- cost-effective for transporting wide and storage area network (WAN/SAN) protocols
- extended reach between difficult nodes
- transparent delivery of client signals.

An OTN introduces three or more sublayers: an optical transmission section (OTS), an optical multiplexing section (OMS), an optical channel (OCH), plus G.709 sublayers (Figure 1). A client-server model is applied between sublayers of the OTN. Each sublayer offers well-defined services to its client layers and each sublayer has its own layer-management functions including fault, performance, and configuration management.

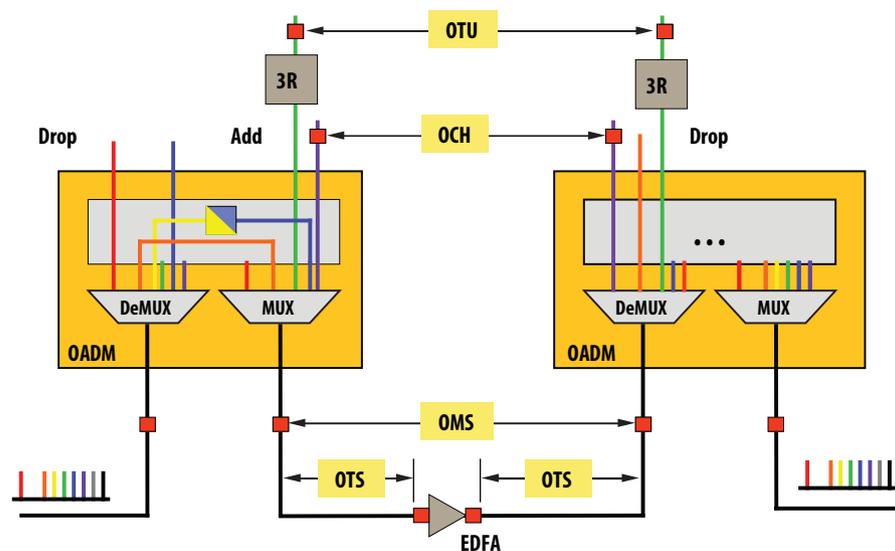


Figure 1: Optical transport network

An OTN provides a much simpler transport protocol than SONET/SDH that is optimized for transport applications and is not burdened with the SONET/SDH provisioning-intensive switching-layer functionality down to 1.5 or 2 Mbps client signals. OTN optical services offer many of the same protections and management features of SONET/SDH networks; however, an OTN is free of the complexity and cost associated with them. OTN services are also better for the transparent mapping and transport of native client traffic through metro and long-haul networks. This point is extremely important for client traffic where the preservation of clock and management information is necessary for sustaining end-to-end path link communications without degradation in performance.

High-Speed Ethernet

The widespread use of 10 gigabit Ethernet (GigE) links in data centers and video applications has driven the need for higher-speed Ethernet interfaces in data centers and metro/core networks (Figure 2). WDM technologies can further increase bandwidth in data and transport applications. However, the scalability of WDM-based systems is limited. To improve scalability and reduce transport cost-per-bit, higher-speed (>10 G) Ethernet and transmission technologies are developed for data and transport applications. They help:

- reduce the number of wavelengths leading to reduced network complexity
- improve fiber and wavelength utilization
- reduce network cost by increasing statistical-multiplexing efficiency
- future-proof systems, keeping them scalable to manage expected bandwidth demands.

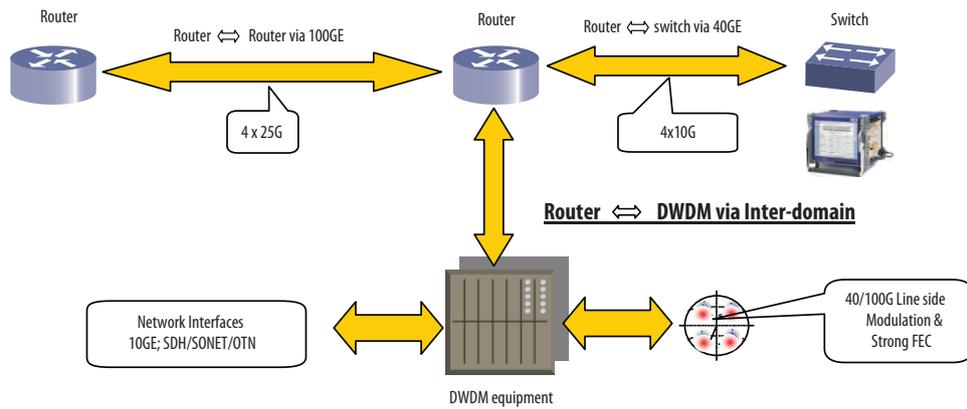


Figure 2: High-speed Ethernet applications

The Institute of Electrical and Electronic Engineers (IEEE) 40 Gbps and 100 Gbps Ethernet task force defined a number of standards identified in Figure 3. These standards define a number of 40 GigE and 100 GigE physical media dependent (PMD) sublayers with various attributes in terms of achievable reach and necessary types of cable and optical components. For inter-office applications, four 10 G fibers can be bundled in a ribbon fiber yielding a total capacity of 40 GigE (40GBASE-SR4) or 100 GigE (100GBASE-SR10). For short-reach metro applications up to 10 km, four 10/25 G wavelengths can be bundled inside a single-mode (SM) fiber operating between 1295 and 1310 nm (100GBASE-LR4). Delivery of 100 GigE can also be delivered up to 40 km with four 25 G wavelengths can be achieved using extended-reach optics (100GBASE-ER4).

IEEE 802.3ba

MMF < 100 m	40GBASE-SR4 Ribbon Fiber 850 nm 4 x 10 G	100GBASE-SR10 Ribbon Fiber 850 nm 10 x 10 G
SMF < 10 km	40GBASE-LR4 CWDM 1310 nm 4 x 10 G	100GBASE-SR4 LAN-WDM 1310 nm 4 x 25 G
SMF < 40 km		100GBASE-ER4 LAN-WDM 1310 nm 4 x 25 G

IEEE 802.3bg

SMF < 2 km	40GBASE-FR Serial TX: 1550 nm
------------	----------------------------------

Figure 3: High-speed Ethernet technologies

Next-Generation Optical Transport Networks

The initial adoption of OTNs was relatively slow due to the large installed base of SONET/SDH- and WDM-based systems. However, a significant rise in bandwidth demand associated with video services and an increasing use of local area networking (LAN) and SAN services drive the need for OTN as a transport technology for point-to-point transmission applications in core/long haul networks and for aggregation/switching applications in metro networks. OTN networks complement packet-based transport (Ethernet, MPLS) and existing SONET/SDH-based networks that serve as clients for OTN networks in metro/core networks (Figure 4).

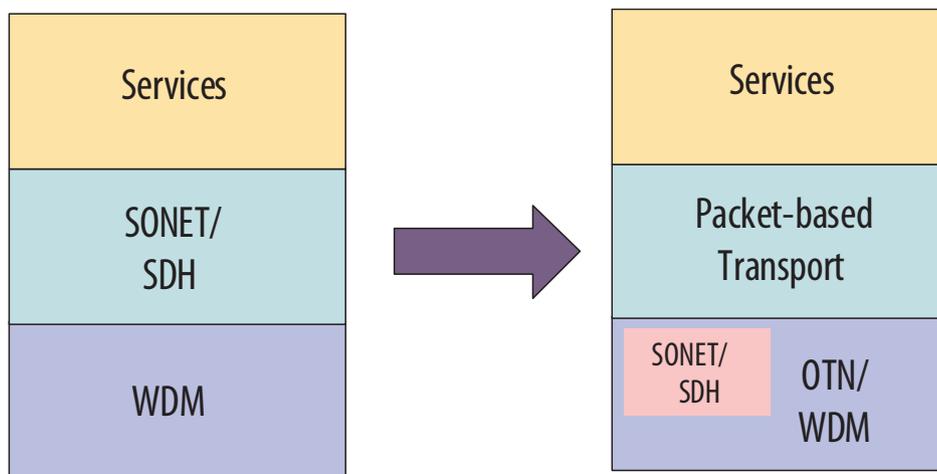


Figure 4: Migration to packet and optical transport networks

To enable the migration of OTN-based systems in metro/core networks, several enhancements are needed in order to enable:

- switching of OTN services
- transmission of higher-speed Ethernet interfaces developed by IEEE
- efficient transporting of 10 G LAN and any-rate SAN services
- direct mapping of lower rate (<2.5 G) clients into an OTN.

OTN Multiplexing and Mapping

The increased use of OTN as a client and as an inter-carrier handoff interface drives the need for OTN multiplexing and switching without the need to de-wrap and use SONET/SDH- or Ethernet-based multiplexing and switching. The ITU G.709 standard indicates that for OTN multiplexing, the FEC need not be considered as it will be recalculated after the multiplexing process. The optical data unit (ODU) portion of the OTN signal can be aggregated with other ODU signals to form a higher order ODU signal. Both lower- and higher-order (LO/HO) ODUs have the same frame format but different line rates. The LO ODU signal corresponds to the client signal; the HO ODU relates to the line interface (Figure 5).

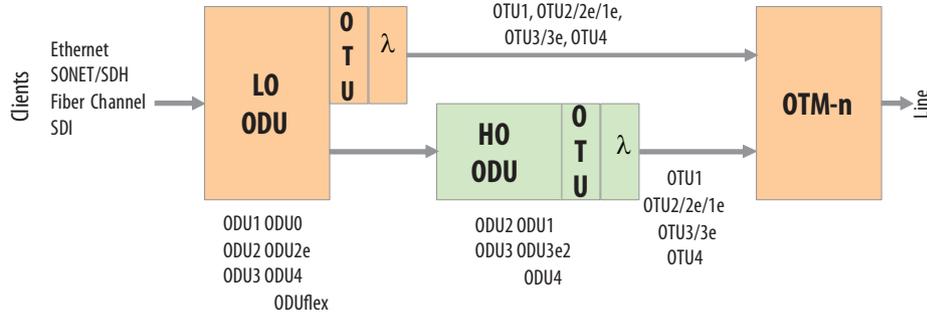


Figure 5: OTN multiplexing
 Note: this OTN diagram includes references to normative ITU documents and G.Sup43. The latter document is not normative.

New mapping procedures further increase the flexibility of OTN for multiplexing and aggregation applications. The original OTN standard mapped constant bit rate (CBR) signals using an asynchronous mapping procedure (AMP). Examples for CBR include 2.5 G or 10 G SONET/SDH (CBR2G5 or CBR10 G). In this scheme, the clock is created locally. The justification control (JC) bytes and their associated negative justification opportunity (NJO) and positive justification opportunity (PJO) bytes in OTN overhead adapt the optical channel payload unit (OPUk) payload rate and the client signal rate.

There are two principal methods for mapping 10 G Ethernet clients as illustrated in Figure 6. Primarily used for 10 GigE LAN clients, bit synchronous mapping (BMP) uses the client clock for the OTN signal. Because the OPU is frequency- and phase-locked to the client signal, there is no need for frequency justification. The JC bytes contain fixed values, the NJO contains a justification byte, and the PJO contains a data byte.

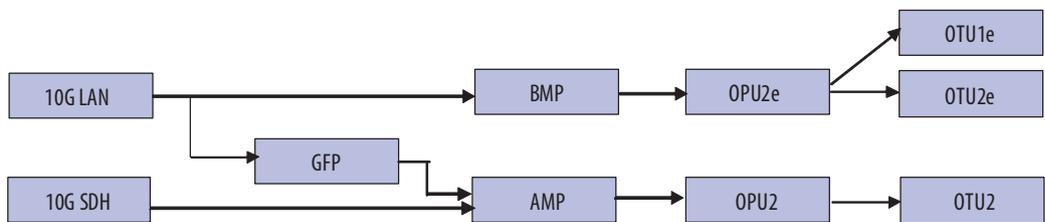


Figure 6: Mapping procedures

Alternatively, Ethernet clients can map into OTN with a combination of generic framing procedure (GFP) and AMP. The GFP maps the Ethernet frames into the OPU payload area. GFP idle frames that are sent in absence of client frames perform rate adaptation. The next step in mapping the OPU signal into the ODU is similar to the one used for CBR streams involving AMP.

A new mapping process proposed in ITU, the generic mapping procedure (GMP), promises to deliver a flexible mapping for clients of any rates. It is based on dividing the OPU segment into 1.25 G slots and filling them with client signals using a sigma-delta distribution method. Unlike other mapping methods, GMP uses stuffing bytes for rate adaptation between the client and ODU signal.

The new GMP process complements the traditional AMP supported in today's OTU2 and OTU3 line ports, enabling support of ODU0, ODU2e, and ODUflex signals by those line ports. GMP also maps new sub-2.5 G CBR clients (for example, 1 GigE, STM-1/4, FC-100/200) into ODU0/ODU1, and maps 40 and 100 GigE into ODU3 and ODU4).

10 G Transport

Transporting 10 GigE client signals has become one of the largest opportunities for OTN and one of its most contentious problems. Ten GigE LAN is a Layer 1 signal that can be used in data centers and metro networks. It is very cost-efficient, but has limitations in reach and jitter/wander performance. Ten GigE WAN is based on a thin SONET/SDH Layer 1 mechanism and has slightly lower throughput capability.

OTN can also carry 10 GigE. Figure 7 depicts several mechanisms for 10 GigE transport over OTN. Ten GigE can map with GFP into OTU2 (10.7 G) signals. The challenge consists in rate adjustment between 10 GigE LAN (10.3 G) and ODU (10.03 G) restricting the OTN signal from carrying the full 10 GigE signal. To adjust the bandwidth in this method, the GFP process removes the Ethernet preamble and inter-frame gap (IFG) characters. While this approach may be acceptable to some carriers, some other carriers demand a full transparent mapping including the IFG and preamble signals that may be used by some vendors for proprietary administrative processes and for new applications such as the transparent transport of synchronization information.

To deliver a fully transparent link, overclocked OTN signals are introduced. They add an OTN wrapper/FEC to the original Ethernet frames but exceed the OTU2 (10.7 G) rate. OTU1e works at 11.05 G while OTU2e operates at 11.1 G. Exceeding the OTU2 rate (10.7 G) is a limiting factor for lack of compatibility with OTU3 rates (43 G) preventing the aggregation of four 11.05/11.1 G bases OTN signals into 43 G. Both overclocked methods inherit the 10 GigE LAN 100 ppm clock tolerance and jitter/wander performance. Unfortunately, rate compatibility with OTU3 and clock/jitter performance conflict as a full transparency rate is reached. This conflict has been one of the major obstacles in agreeing on a single 10 G OTN rate for carriers.

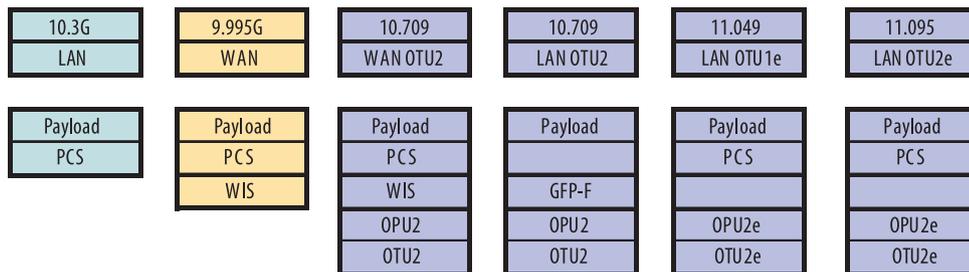


Figure 7: 10 G clients and mappings

High-Speed Transport (>10 G)

The original G.709 standard defined the OTU3 signal (43 G) for highest-speed applications. The emergence of higher-speed Ethernet interfaces demanded significant enhancements to original OTN signal specification. Several applications drove the need for OTN enhancements (Figure 8):

- multiplexing of four 10 GigE LAN clients
- mapping of 40 GigE parallel clients
- mapping of 40 GigE serial interface (to be standardized).

The original OTU3 signal allows the mapping of four 10 GigE LAN clients, if GFP-based, G.709 rate-compliant (10.7 G) mapping is used. For multiplexing of overclocked signals, new mechanisms are defined: OTU3e1 and OTU3e2.

The ODU3e1 is designed to carry four ODU2e signals as its only client. The frame format and multiplexing techniques for the ODU3e1 are similar to the ODU3 with the following exception. In order to accommodate the ± 100 ppm clock tolerance of the ODU2e clients, the ODU3e1 justification mechanism has been extended with an additional PJO and NJO byte and a corresponding JC byte format to use them. The ODU3e1 bit rate is $(\text{ODU2e rate})(4)(239/238) = 41.774364407 \text{ Gbps} \pm 20 \text{ ppm}$. The carriers that prefer ODU3e1 are those that have made extensive use of Ethernet as a network technology and use 10 GigE over ODU2e in much of their transport networks. Since ODU2e is their primary 10 Gbps OTN signal, they need an efficient means to carry four ODU2e signals over their 40 Gbps OTN links and are not as concerned about also carrying ODU2 signals. NTT has already deployed ODU3e1 in its network.

The OTU3e2 allows an ODU3 to carry up to four ODU2e signals and a mixture of ODU2 and ODU2e clients and is favored by several carriers with substantial OTN deployments. These carriers prefer using non-overclocked approaches as their typical method for 10 GigE client transport and see applications requiring full bit transparent transport as sufficiently rare so that the mapping inefficiency is not important. The mapping must extend the OPU3 justification capability to accommodate the ODU2e rate and a ± 100 ppm clock range. The current agreement is that the GMP currently under study by SG15 will be used to provide this extended justification capability.

For mapping of 40 and 100 GigE parallel interfaces, ITU specifies respective parallel OTU3 and OTU4 interfaces. The parallel OTU3/OTU4 interfaces benefit from 10 GigE cost advantage. These parallel interfaces use an inverse-multiplexing method across physical lanes aligned with the OTU3/OTU4 frames. The frames are distributed in round-robin manner between the lanes. Three additional OTLx.y rates are defined. An x value of 3 (4) corresponds to an aggregated signal rate of 40 G (100 G). The y stands for the number of lanes and can take the values of 4 or 10.

OTU3 43G	OTU3e1 44.57G	OTU3e2 44.58G	OTL3.4 4x10.7G	OTU4 111G	OTL4.4 4x28G	OTL4.10 10x11G
-------------	------------------	------------------	-------------------	--------------	-----------------	-------------------

Figure 8: 40/100G OTN rates

Finally, IEEE is developing a new 40 G serial interface with a nominal bit rate of 41.25 G. The carriers demand a full transparent mapping method that maintains all characters including administrative and synchronization information without introducing an overclocked rate as with 10 G. To maintain the original 43 G OTU3 line rate, a new method is proposed that transcodes 64 B/66 B blocks into 1024 B/1027 B blocks resulting in a client signal rate of 40.117088 Gbps. This client signal maps into OPU3 using GMP.

ODU0

One of the main drawbacks of the original G.709 was the lack of a multiplexing scheme for lower-speed (sub-2.5 G) client signals. One important application for sub-2.5 G multiplexing is for connections to customer-located equipment (CLE) for enterprise customers. Today, these optical connections commonly use SONET/SDH which requires full SONET/SDH functionality in the CLE and/or the CPE.

Deploying an OTN interface to the customer can eliminate the need for the CLE and/or CPE to support SONET/SDH. For example, if the enterprise customer uses a SONET/SDH UNI, the OTN CLE can be transparent to the SONET/SDH signal. Also, if a GigE or SAN UNI is used, the OTN CLE can likewise be transparent to this signal and not require SONET/SDH with virtual concatenation (VCat) to transport it.

Another important application for sub-2.5 G multiplexing is to enable efficient Layer 1 aggregation in the access and metro networks. For the next several years, there will be SONET/SDH signals coming from both legacy carrier access equipment and enterprise network interfaces. Rather than maintaining an entire SONET/SDH access network, considerable network simplification can be achieved if these signals are carried transparently through an OTN. Aggregation and multiplexing is important for bandwidth efficiency, and Layer 1 aggregation is much less expensive than aggregation at higher layers. Layer 2 or 3 aggregation is better performed deeper into the network where greater statistical multiplexing gains can be achieved. As GigE signals begin replacing SONET/SDH for enterprise customer connections and interfaces to broadband access systems and wireless base stations, it is important that a sub-ODU1 multiplexing method also be able to efficiently support GigE.

The OTN networks for Layer 1 access network and enterprise customer aggregation will typically use point-to-point or tree structures in which no add/drop capability is required. This is an important consideration, since providing any switching capability, including add/drop multiplexing, requires a path-type overhead that remains with the client signal end to end. Network management and equipment are significantly simpler with a non-switched point-to-point or tree-multiplexed network.

Multiple schemes are being considered and proposed to support lower-speed signals. While some are relatively early in the standardization process, others such as ODU0 are becoming ready for approval. ODU0 allows the mapping of 1 G Ethernet clients into ODU signals (Figure 9). For example, two ODU0 signals can be mapped inside an ODU1 signal. A timing transparent transcoding (TTT) method synchronously maps a 1 GigE signal (8 B/10 B coded, nominal bit rate of 1.25 G, ± 100 ppm) into a CBR stream using GFP. GMP maps the CBR stream into OPU0.

The main benefits of the ODU0 are that it is a switchable entity, it is consistent with the other ODUk signals, and it is optimized for efficient transport of the most important sub-ODU1 client (GigE). The main drawback of the ODU0 is that in the short term, no switch fabrics or network management systems support it. Consequently, there will be a transition period for several years for ODU0 deployment. The first stage of this transition will use just ODU0 multiplexing into ODU1 with the switching performed at the ODU1 level. This approach is adequate for many applications, since the end nodes for the ODU1 can handle some grooming of ODU0s into the ODU1s. Eventually, full ODU0 support will be available in the network.

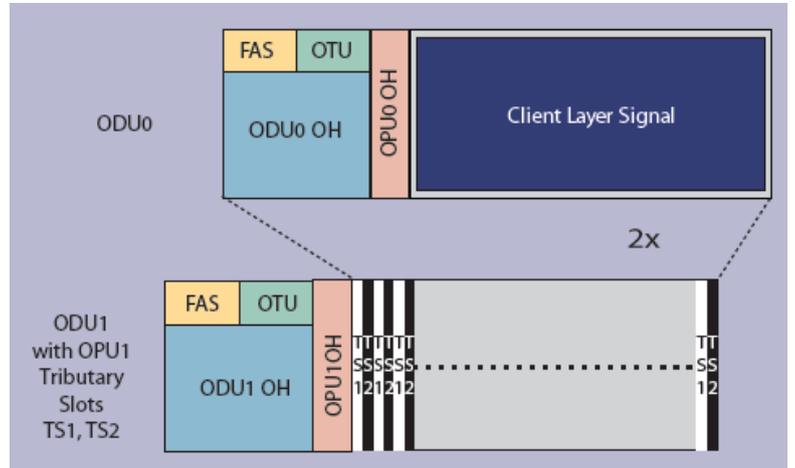
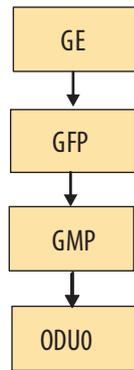


Figure 9: 1 GigE and ODU0 mapping

OPU Virtual Concatenation OPUk-Xv

Analog to SONET/SDH VCat, G.709 defined an OPU VCat concept for efficient transport of signals whose bandwidth do not match with existing ODU line rates. One example is 4 or 8 G Fibre Channel for which 4 G can be transported more efficiently if two OPU1 signals are concatenated rather than by using an OPU2 signal.

The OPU members of the VCat group are identified in Figure 10 using VCat overhead (VCOH) and a payload structure identifier (PSI). Each member can take different routes through the network that results in further advantages for VCat: better utilization of bandwidth and timeslots on existing links and improved resiliency in case of outage on one of the links. To let different members take different routes through the network, the OPU VCat mechanisms allows for the characterization and compensation of delays (differential delays) between various members going through different routes. Finally, a link capacity adjustment scheme (LCAS) provides a mechanism for hitless adaptation of bandwidth.

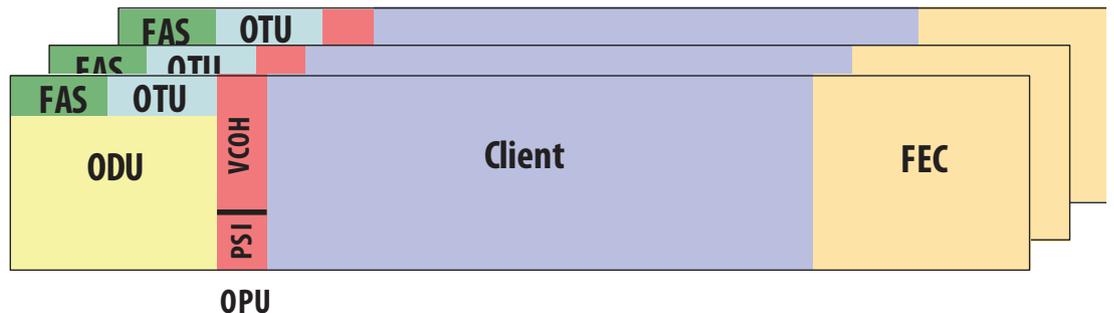


Figure 10: OPU Virtual concatenation (OPUk-Xv)

ODUFlex

While VCat delivers superior flexibility and resiliency, it is a relatively complex technology to build and manage. ODUFlex is proposed as an alternative mapping method without the complexity of VCat. ODUFlex applies to signals with rates higher than OPU1. They are bit-synchronously mapped into an ODUFlex frame that is located into a flexible collection of 1.25 G timeslots using GMP (Figure 11).

For CBR clients, ODUFlex carries the clock tolerance of the client signal. For packet clients, the client maps into the ODUFlex at a rate slightly above the service rate (by 239/238) allowing for a clock tolerance ± 100 ppm.

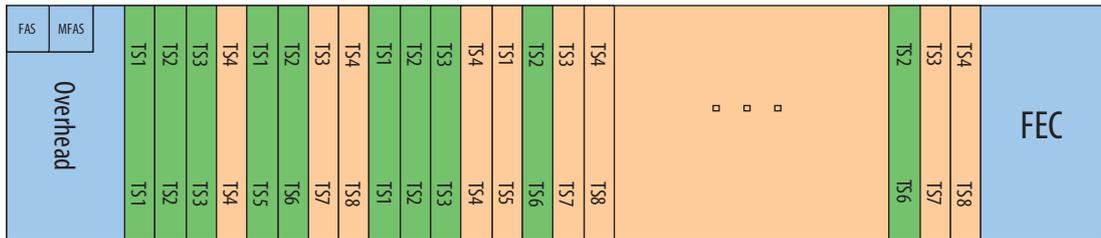


Figure 11: ODUFlex Multiplexing

Note that the new OTU4 and OTU3e2 line ports treat ODU_k (k=0,1,2,e,3) signals as specific bit-rate versions of the generic ODUFlex signals (for example, all those ODU_k/ODUFlex signals are mapped into the HO ODU4 and HO ODU3e2 by means of the GMP). One only needs to configure the bit rate and tolerance of the LO ODU signal and the tributary slots in which it is transported in the line port to have the GMP process perform its task to map/demap this LO ODU signal into/from the HO ODU tributary slots. This programmability also enables the hitless resizing of the ODUFlex connections when such an ODUFlex connection carries GFP-encapsulated packet streams (for example, Ethernet, IP, MPLS).

Other new CBR clients map into an ODUFlex using the very simple BMP which wraps the ODUFlex around the CBR client.

Conclusion

OTN delivers a cost-effective transport technology that will complement packet-based technologies (Ethernet and MPLS). Several enhancements are being introduced for OTN to efficiently work with packet-based technologies and migrate from point-to-point long-haul applications into metro/core networks. They allow efficient and transparent transport of emerging 40/100 G clients and packet-based clients and enable aggregation and switching at the OTN layer. Major enhancements include the introduction of low- and high-order ODU schemes, GMP, OTU3e1/OTU3e2, OTL3.4 & OTL 4.4, ODU0, and ODUFlex technologies.

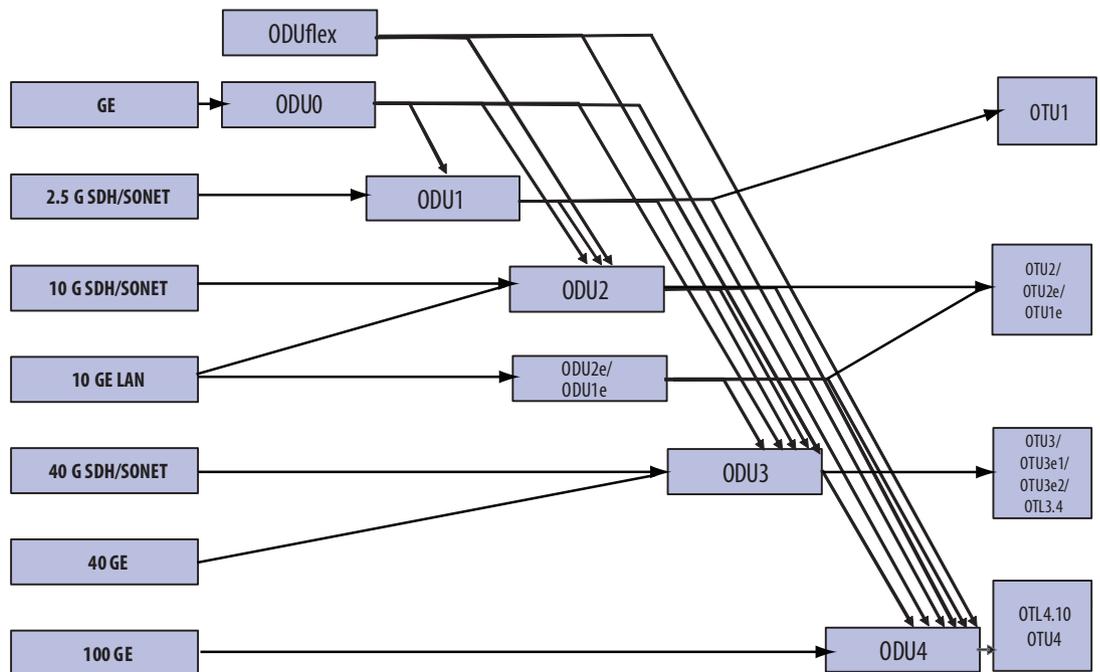


Figure 12: Enhanced OTN hierarchy

References

ITU-T Recommendation G.709 (2009), *Interfaces for the Optical Transport Network (OTN)*

ITU-T Recommendation G.798 (2004), *Characteristics of optical network hierarchy equipment functional blocks*

ITU-T Recommendation G.872 (2001), *Architecture of optical transport networks*

ITU-T Recommendation G.692 (1998), *Optical interfaces for multichannel systems with optical amplifiers*

ITU-T Recommendation G.707 (1996), *Synchronous Digital Hierarchy Bit Rates.*

ITU-T Recommendation G.7042/Y.1305 (2001), *Link capacity adjustment scheme for virtual concatenated signals*

ITU-T Recommendation G.694.1 (2002), *Spectral grids for WDM applications: DWDM frequency grid*

ITU-T Recommendation G.694.2 (2003), *Spectral grids for WDM applications: CWDM frequency grid*

ITU-T Recommendation G.975.1 (2004), *Forward error correction for high bit rate DWDM submarine systems*

ITU-T Supplement G.Sup43 (2008), *Transport of IEEE 10 G Base-R in Optical Transport Networks (OTN)*

ITU-T Recommendation G.7041 (2008), *Generic Framing Procedure (GFP)*

ntz (2009), *OTN als Transportmedium der Zukunft und Anforderungen an die Messtechnik* Peter Winterling

ITU-T Recommendation G.959.1 (2009), *Optical transport network physical layer interfaces*

IEEE P802.3ba (2010), *Draft Amendment: Media Access Control Parameters, Physical Layers and Management Parameters for 40 Gbps and 100 Gbps Operation*

Test & Measurement Regional Sales

NORTH AMERICA TEL: 1 866 228 3762 FAX: +1 301 353 9216	LATIN AMERICA TEL: +1 954 688 5660 FAX: +1 954 345 4668	ASIA PACIFIC TEL: +852 2892 0990 FAX: +852 2892 0770	EMEA TEL: +49 7121 86 2222 FAX: +49 7121 86 1222	WEBSITE: www.jdsu.com/test
---	--	---	---	---