

# Future-proofing the Optical Transport Network

By Geoff Bennett, Infinera Corp.

Technology is at its most useful when it's flexible enough to adapt to changing needs. Take WiFi as an example; imagine how complex it would be if there were half a dozen different WiFi standards out there. But wait...there actually are at least nine different 802.11 WiFi variants— but part of the WiFi standards process is to ensure that different implementations are flexible enough to negotiate the highest data rate for most combinations of client and base station.

This kind of flexibility could be useful in other areas too. In particular, the Optical Transport Network (OTN) containers that carry services in a reliable and manageable way across the internet have to be better able to respond to the appearance of new service data rates without waiting for several years for the standards committee to create a new container size. Let's take a look at how this technology is likely to evolve beyond the current 100-Gbps data rates.

## OTN transport containers

The OTN is defined in a family of ITU-T standards (G.709, G.798, G.872) that were designed to take over where SONET/SDH left off. So the first OTN container was defined as Optical Data Unit (ODU1), with a payload size of about 2.5 Gbps and compatible with OC-48/STM-16. Each higher data rate OTN container initially multiplied the capacity by a factor of 4 (as both SONET and SDH had tended to do previously). Thus ODU2 has a capacity of about 10 Gbps and ODU3 a capacity of 40 Gbps.

These ODU containers can be multiplexed together – so, for example, four ODU1 containers can be combined in an ODU2. At some point in the multiplexing process the service provider will map the ODU container onto a DWDM wavelength for long-haul transmission. If this was an ODU2 then the DWDM equipment would add the appropriate management fields, plus a forward error correction (FEC) field. The resulting structure is called an Optical Transport Unit 2 (OTU2), which is now ready to be transmitted on the fiber.

Figure 1 shows a simplified view of this process, where a client signal is mapped into an Optical Payload Unit, or OPU, at level “ $k$ ” in the hierarchy (where  $k = 1$  for ODU1,  $k=2$  for ODU2, etc.). This is then mapped into an Optical Data Unit (ODU $k$ ), to which a “digital wrapper” consisting of management overhead plus FEC is added to become an Optical Transport Unit (OTU), which is then transmitted over an optical fiber as an Optical Channel (OCh).

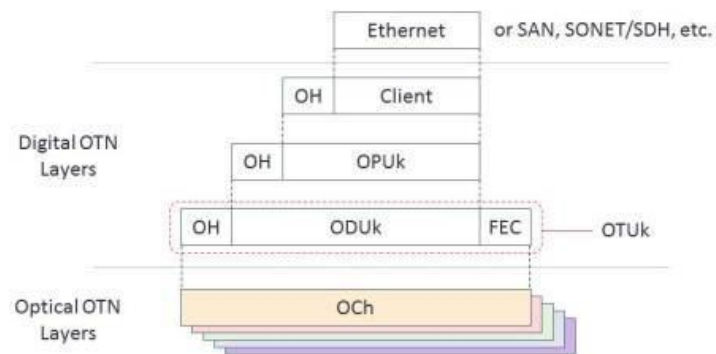


Figure 1. The fundamentals of OTN encapsulation.

While this process may appear convoluted, the OTN encapsulation and multiplexing standards were relatively straightforward -- until Ethernet started to be accepted as a carrier service. Suddenly the existing OTN containers seemed to be all the wrong sizes. For example the ODU2 container, despite being a nominal 10 Gbps in size, wasn't quite big enough to fit a full 10 Gbps Ethernet Physical Code Sublayer (PCS) stream; and the ODU1 container at a nominal 2.5 Gbps capacity was far too big to carry a single 1.25-Gbps Gigabit Ethernet PCS stream efficiently.

The result of this mismatch in client signal rates and ODU containers was that OTN vendors began to improvise, which meant that for several years Ethernet services were carried in proprietary containers until the OTN standards eventually caught up. Fortunately 40GbE could just about be carried in an ODU3 with some clever transcoding of the signal, and the new ODU4 container was (by design) actually the right size for a 100-Gbps PCS stream.

Most people would agree that OTN containers, while perfectly well designed for backwards compatibility with SONET/SDH, needed to be far more flexible to keep up with more modern services such as Ethernet, storage-area networks (SANs), and native digital video services. In particular, there had to be a way for services to be carried in standardized containers without having to wait for the ITU working group to create those standards.

Chronologically the first new development was to carry Gigabit Ethernet services more efficiently in a standard OTN container. Thus "ODU0," a 1.244-Gbps container, was created. The next step was to use the 1.25G tributary slots as the building blocks for a flexible container type -- dubbed ODUflex. Using this technique, client services can be efficiently mapped into transport capacity without having to wait for a new standard to be ratified and new multiplexing hardware to be developed. More importantly, ODUflex works within the OTN multiplexing hierarchy, and so a service that is mapped into an ODUflex container can itself be multiplexed as a client signal into a larger capacity "fixed" OTN container, such as ODU2, ODU3, or ODU4.

### **The advent of superchannels**

A DWDM superchannel is made up of multiple coherent optical carriers that are implemented on the same line card so that they can be brought into service in a single operational cycle. In mid-2012 the first commercial superchannel deployments began to roll out, using 500-Gbps polarization-multiplexed quadrature phase-shift keying (PM-QPSK) superchannels based on large-scale photonic integrated circuits (PICs). This type of coherent modulation technique is capable of around 3,000 km of optical reach, and with very good spectral efficiency. For even greater reach the modulation can be switched to, for example, PM-binary phase-shift keying (PM-BPSK). This delivers trans-Pacific distances, but as a consequence the capacity of the superchannel is reduced to 250 Gbps.

Some vendors are also showing demonstrations of non-PIC-based 400-Gbps superchannels using PM 16QAM modulation, and these may become commercially available in 2013. But 16QAM has a relatively short optical reach -- typically in the metro range. These "400 Gbps" superchannels will actually become 200-Gbps superchannels when switched to PM-QPSK mode (for long haul reach), and 100-Gbps superchannels in BPSK mode (for ultra long-haul and submarine operations).

In other words, superchannels will vary in their capacity depending on the modulation used and the reach required. This flexible coherent capability is a significant advantage for superchannel deployments because it allows the network designer to “trade capacity for reach.” But it does beg the question of what the associated OTN container should be for these transport technologies. If, for example, ITU-T were to define a fixed-bandwidth, 400-Gbps OTN container, it would work well with today’s PIC-based 500G superchannels (the extra 100 Gbps of capacity could be used by, for example, an ODU4 container), but it would be too large a container for the long-haul 200-Gbps PM-QPSK superchannels that are currently planned to be shipping next year.

### ITU-T SG15 “beyond 100G”

Partly inspired by the success of ODUflex, ITU-T Study Group 15 Question 11 is investigating flexible OTN technologies beyond the currently defined OTU4 and ODU4 containers at the 100-Gbps rate. Vendors and carriers such as Infinera, Verizon, Huawei, and Deutsche Telekom have co-authored or submitted like-minded proposals to this group, and a significant consensus is now beginning to form.

As with ODUflex, the first step is to determine the size of the building block for the flexible container. Since superchannels will be “above 100 Gbps” in data rate, it makes sense that the building block should be in the “tens of Gbps” range, and the possible choices for candidates are 25 Gbps, 50 Gbps, and 100 Gbps. Different vendors will use different superchannel implementations and different coherent modulation techniques, so the granularity chosen should be as inclusive as possible, while still offering efficient service mapping and multiplexing. At least one of the proposals is dubbed “OTUadapt,” although the official name may be different.

OTUadapt allows the transport infrastructure to adapt to the flexible capacity of the underlying coherent superchannels, but it also allows the “future proofing” of the OTN hierarchy in the face of emerging service types. For example, the IEEE Higher Speed Study Group (HSSG) is now debating the speed of the next generation of Ethernet standards beyond 100GbE. The electronic component vendors are pointing out that anything above a 400-Gbps data rate might cause several years of delay in getting the standard to market.

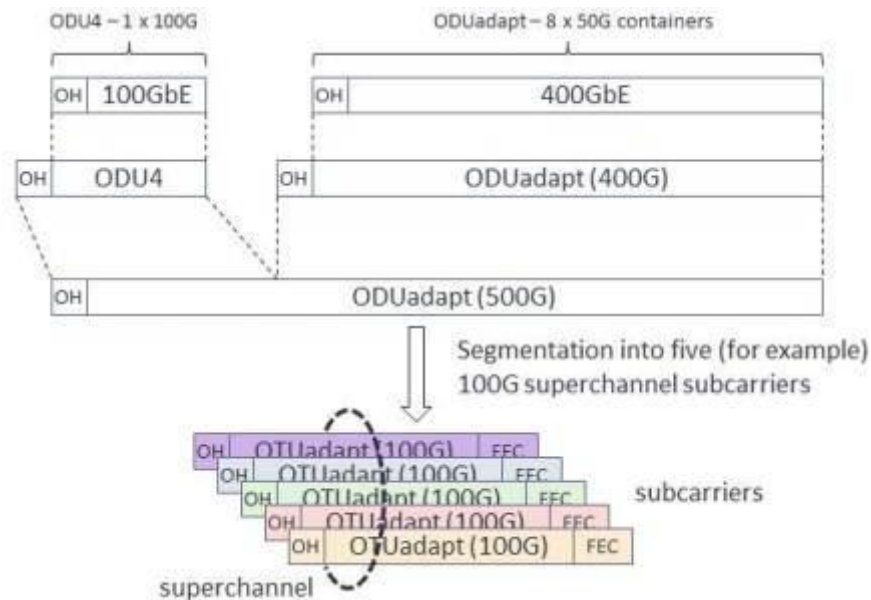
Why is ODUflex defined at the data unit level, while OTUadapt is defined at the transport unit level?

ODUflex was designed to build “odd sized,” but sub-lambda transport containers for client services, and then to be multiplexed with other ODUs into a larger capacity OTU, which would then be mapped onto a single wavelength.

In contrast, OTUadapt is a “super-lambda” sized OTU container that spans multiple wavelengths of a super-channel, and multiple lower-order (rate) ODUs can be multiplexed into the higher-order ODU signal (ODUadapt) supported by an OTUadapt signal.

On the other hand, some representatives of the end-user community would prefer to set a “stretch goal” of Terabit Ethernet, pointing out that moving from one Ethernet speed to another is a significant development and it should not be taken for a small increment in performance. Some observers have pointed to the confusion created by the dual-speed IEEE 802.3ba standard (40GbE and 100GbE), and feel that this served to introduce delays in getting the cost of these higher-speed Ethernet technologies down to an economic level.

In other words, there is likely to be some debate about the next Ethernet data rate, so a flexible OTN container would be ideal to avoid adding to any confusion. For example, an OTUadapt approach would ipso facto be compatible with 400GbE, 800GbE, and even Terabit Ethernet.



**Figure 2. A look at OTUadapt.**

Figure 2 shows an example of how OTUadapt might work. Two high data rate services, a single 400GbE and a single 100GbE, have been multiplexed into a 500G ODUadapt container. The upper section of Figure 2 shows how a future 400GbE service could be carried over a 500-Gbps PM-QPSK long-haul superchannel, along with an additional 100GbE service, using a 400-Gbps OTUadapt container and a 100-Gbps OTU4. The OTUadapt and OTU4 containers would coexist in the same superchannel. Note that 400GbE has not yet been determined as the next Ethernet data rate, and this rate was chosen purely as an example of a “greater than 100GbE” service. The ODUadapt container is mapped to a 500G superchannel by segmentation into OTUadapt units that are sized to fit a subcarrier, in this case a 100-Gbps subcarrier.

A more likely scenario would be for large numbers of 1GbE and 10GbE services to be multiplexed into an appropriate ODU (for example, ODU4s), which can then be combined to fill a 200G, 400G, 500G, or even a 1T superchannel. Since OTUadapt containers can be flexibly sized, this enables the network designer to choose the most efficient container sizes for multiplexing these services and for matching them to the underlying superchannel rates.

The reason that the second example is more likely is that, according to analyst forecasts, services such as 1GbE and 10GbE will massively dominate the service mix in the foreseeable future.