www.pipelinepub.com
Volume 21, Issue 12

Unlocking New Possibilities for Network Designs with Thin Transponders

By: Fady Masoud, M.Eng

Major improvements in increasing fiber capacity have come from the evolution of optical engines – from direct-detect to coherent – and from increased wavelength bit rate enabled by the higher order modulation schemes. However, optical engines are approaching the Shannon limit. As a result, fiber capacity has begun to plateau as spectral efficiency becomes constrained.

This has triggered service providers and network operators to begin looking at new strategies for increasing the capacity of their networks, including the use of multi-fiber strategies to connect locations. As network operators embrace these multi-fiber strategies, it can cause a significant shift in their priorities when selecting network solutions. When networks are designed using a single fiber pair between locations, capacity per fiber is typically the most important factor in selecting an optical engine. However, when multi-fiber strategies are used, capacity per fiber becomes significantly less important, enabling network operators to focus on other factors – such as power and space efficiency.

This shift in priorities is aligning perfectly with the evolution of coherent pluggables. Advances in DSP (digital signal processor) and CMOS (complementary metal-oxide semiconductor) geometries – from 28nm to 7nm to 3nm – have led to significant enhancements in optical performance, increasing capacity-reach in compact form factors such as QSFP-DD (Quad Small Form-factor Pluggable Double Density) and OSFP (Octal Small Form-factor Pluggable). And while these coherent pluggable optical engines provide less fiber capacity compared to their embedded optical engine counterparts, they provide a significant reduction in power and footprint per bit.

With multiple coherent optical solutions and deployment models available, network operators need to understand the pros and cons of each approach and determine which is right for their network and applications. Embedded optical engines with their high bitrate (1.2+ Tb/s) per wavelength capacity-reach and spectral efficiency are the optical transmission workhorses for terrestrial and subsea networks where fiber is scarce or expensive, or both. However, coherent pluggables, such as the emerging 800G ZR/ZR+, deliver greater performance than ever before while offering significant reductions in space and power vs. their embedded counterparts. While IPoDWDM, with coherent pluggables deployed directly into IP routers, is one way to deploy coherent pluggables, it is not the only one. By understanding the benefits and trade-offs for each network application and deployment model, network operators can make the best choice for today and tomorrow.

Every application will have its own set of drivers. Embedded engines continue to be the technology of choice when fibers are scarce and the cost of adding incremental fibers is high. However, the compelling economics of coherent pluggables, coupled with the shift in priorities away from maximizing fiber capacity and toward cost, space, and power efficiency, are significantly expanding the applications' scope. Network operators can secure a significant portion of the benefits of coherent pluggables while avoiding increased operational complexity or loss of optical functionality by selecting a third deployment model.

But what are the two current deployment models?

Current deployment models

There are currently two main deployment models of coherent optical engines being promoted by the industry:

1. High-performance optical engines embedded in a transponder card in an optical transport platform
2. Coherent pluggables hosted in IP/Ethernet platforms, such as switches and routers

Embedded optical engine-based transponders are designed and built from the ground up to maximize fiber capacity with best-in-class optical performance (capacity-reach). They aggregate traffic from various types of client interfaces onto one or two high-bit-rate wavelengths, such as 1.2 Tb/s. Designed for deployment in fully integrated optical transport platforms, they typically support multiple client ports for aggregation, in addition to a full suite of integrated optical functions designed for traffic grooming, aggregation, hair pinning, protection, alarm correlation, and other purposes.

In addition to the performance and host of traffic management features, another advantage of embedded transponders is the operational domain separation, acting as a demarcation point between the IP (Internet Protocol) and the optical domain. This is useful in enforcing service level agreements (SLAs) through setting clear boundariesand business demarcation. However, embedded transponders come with trade-offs such as high-power consumption and a large footprint.

The second model that the industry is currently focusing on is the deployment of coherent pluggables directly into routers, commonly referred to as IP over Dense Wavelength Division Multiplexing (IPoDWDM). This model has the advantages of low capital expenditure (CAPEX), low power consumption, and reduced footprint by eliminating the need for an optical transport platform. Moreover, coherent pluggables benefit from a complete ecosystem of Multi-source Agreements (MSA) and interoperability forums for a seamless line interworking and service provisioning.

However, IPoDWDM presents operational challenges related to management and compatibility across different host devices and operating systems. Furthermore, this model does not come to par with fully-fledged transponder-based platforms due to the lack of complete traffic aggregation capabilities and operational functionality.

Additionally, the IPoDWDM model has a direct one-to-one mapping between the router port speed and the pluggable speed – for example, where 400G pluggables are deployed in 400G router ports and 800G pluggables in 800G router ports. Nonetheless, sometimes the coherent pluggables must be dialed down to operate at a lower bit rate, such as 600 Gb/s or 400 Gb/s, to close a specific link, which results in a wasted router port capacity. Moreover, pluggables cannot fully replicate the optical functionality of embedded transponders.

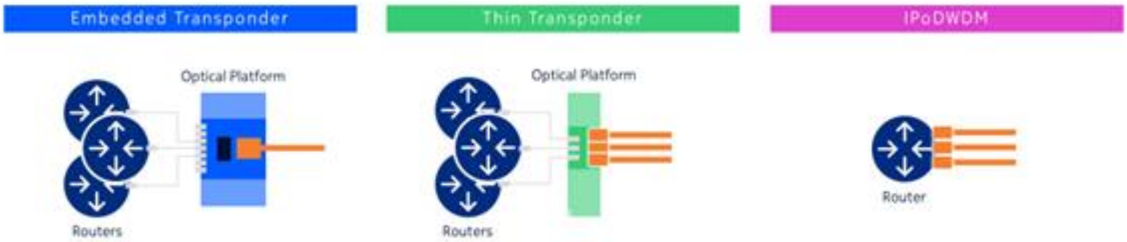

Figure 1: Expanding deployment models with thin transponders
click to enlarge

Expanding deployment modules with thin transponders

Thin transponders introduce a third deployment model that combines some of the strengths from the other two models without their associated disadvantages. Thin transponders are a set of small modules or "sleds" that are optimized for coherent pluggables. These are normally equipped in compact modular optical platforms where multiple sleds can be equipped in a single rack unit. Thin transponders offer multiple client ports (100G, 200G, 400G, or 800G) for grey optics to carry traffic from other platforms, such as routers, and multiple line ports for high-capacity coherent pluggables, such as 400G ZR, 400G ZR+, 800G ZR, or ZR+. Similar to IPoDWDM solutions, thin transponders offer lower CAPEX, lower power consumption, and smaller footprints, but without the operational challenges mentioned earlier. They also combine some of the advantages of embedded transponders, such as multiple client-side aggregation, operational domain separation, and some of the optical capabilities of fully fledged embedded transponders. Figure 1 (above) depicts each deployment model.

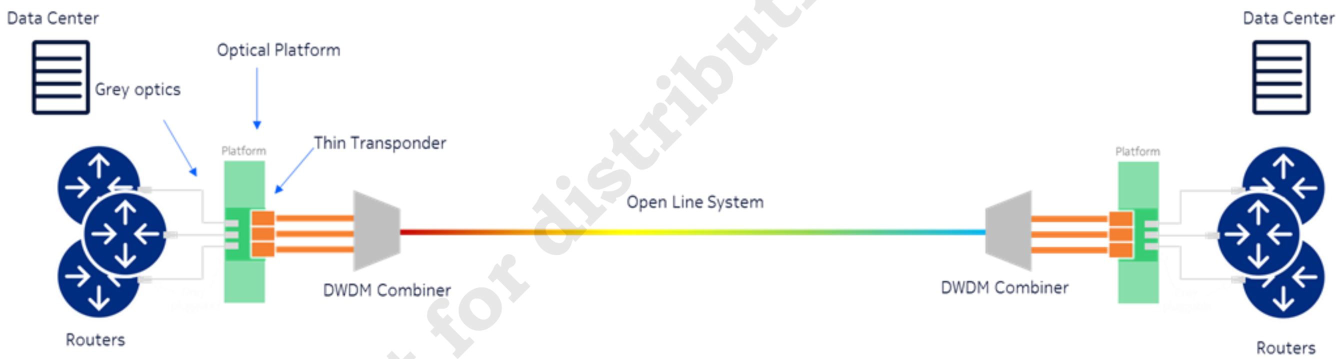

Figure 2: Typical Thin Transponder Application
click to enlarge

Thin transponders also enable a technological lifecycle separation between the long-lasting refresh cycle of the photonic layer and the shorter cycle of the IP layer. This allows network operators to benefit from the latest generation of coherent pluggables, such as ICE-X 800G ZR/ZR+, in existing 400G routers. This leads to a maximized ROI and operational flexibility by avoiding the network-wide upgrade of all routers to the latest 800G-capable generation.

Use cases of thin transponders

As noted, thin transponders combine key attributes of embedded transponders and coherent routing, providing cost-effective, flexible, and highly reliable optical transport leveraging the latest generation of coherent pluggables. Thin transponders are the solution of choice when space and power are limited, and a low variety of client services is required. Another advantage is the reduction in sparing costs, as the same 800G coherent pluggables can be used in IPoDWDM applications and in thin transponder modules. Figure 2 (on previous page) depicts a typical thin transponder application, such as connecting data centers.

One key application is backhauling data center traffic at Submarine Landing Terminal Equipment sites (SLTE). Embedded transponders are used to maximize fiber capacity from the wet plant. In contrast, thin transponders are used to reduce power consumption and footprint for traffic in the dry plant, connecting in-land data centers as depicted in Figure 3.

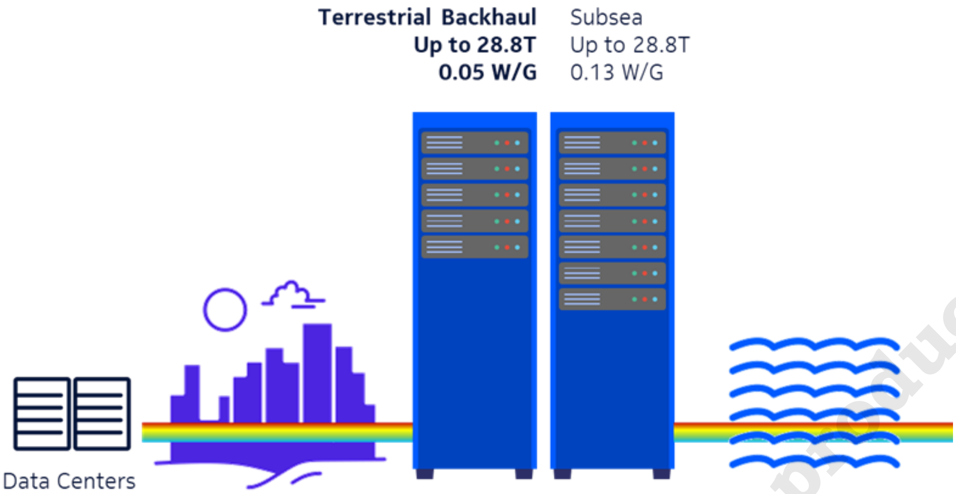

Figure 3: DC Backhaul using Thin Transponders
click to enlarge

Case study

To quantify the benefits of thin transponders, a network analysis was performed on a full-filled (C-Band) 1,000 km long optical link where all three deployment options were compared side-by-side. Compared to embedded transponders, thin transponders offer compelling reduction in CAPEX ($/G), power consumption (W/G), and footprint (RU). However, they offer slightly less (20%) capacity per fiber as depicted in Figure 4.

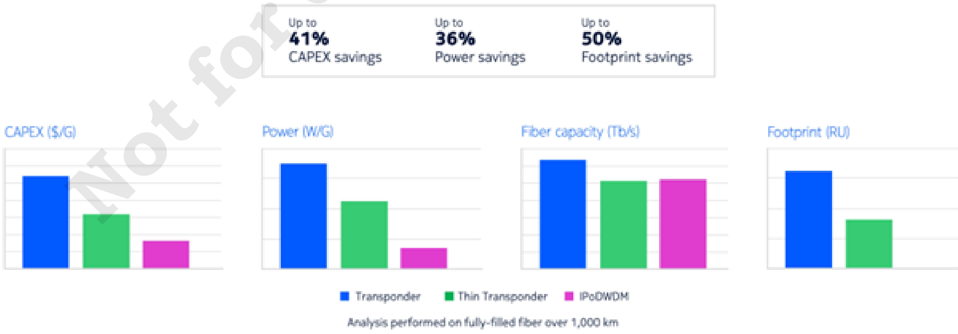

*Reflects different volumes and scale
Figure 4: Quantifying the benefits of thin transponders
click to enlarge

Creating flexibility and efficiency for optical networks

Thin transponders combine the key advantages and benefits of embedded transponders and IPoDWDM to provide deployment flexibility, while offering compelling economics through reduced CAPEX, power consumption, and footprint.