PennWell®

LIGHTWAVE®

EDITORIAL GUIDE

Data Center Interconnect Trends

The emerging data center interconnect application promises to drive new requirements as well as new technology – not only for metro access and core as well as regional networks, but for long-haul links as well. This editorial guide reviews the requirements for data center interconnect as well as both current technology offerings and likely future innovations. It also examines the optical technologies inside the data center that will meet growing capacity requirements.

Originally published June 13, 2016

A Snapshot of Today's Data Center Interconnect

By **JIM THEODORAS**

NOW THAT "DCI" has become a "thing," a well-known acronym with a wide range of connotations, "data center interconnect" has gone from being unknown and somewhat misunderstood to overused and overhyped. But what exactly the term DCI constitutes is changing all the time. What it meant last month might not be what it means in a few weeks' time. With that in mind, let's try to take a snapshot of this rapidly evolving beast. Let's attempt to pin down what DCI currently is and where it's headed.

Power Still Dominates

With rapidly advancing technology, one might be forgiven for thinking that the power situation in data centers is improving. Unfortunately, the growth in the cloud is outpacing advances in power efficiency, keeping power at the top of the list of concerns for data center operators (DCOs). And the industry has responded with a two-pronged attack on power: architectural efficiency and technology brevity.

Architectural efficiency: It sounds somewhat obvious to say, "Just get more from less," but it can actually be a very powerful concept. Reduce the number of network layers and the amount of processing needed to accomplish a task, and work-per-watt goes up. For example, while some have criticized the proliferation of Ethernet data rates, the truth is that these new intermediate rates are desperately needed. Having better granularity at exact multiples of each other greatly improves processing efficiency and dramatically lowers the overall power needed to perform identical amounts of switching and processing.

Throwing it all away: One of the most heavily debated data center mythologies is that everything gets discarded every 18 months. While admittedly there is some hyperbole to this, there is also a lot of truth. Data centers are not on a rapid update cycle just for bragging rights; simple math drives the process. If you have 50 MW available at a site and cloud growth forces you to double capacity, your options are to either double the size of the site (and hence power to 100 MW) or replace everything with new hardware that has twice the processing power at the same electrical power. The replaced hardware works its way down the value chain of processing jobs until it's obsoleted. Only then, by the way, is it destroyed and recycled.

The same logic is driving the push for higher-order modulation schemes in DCI. Each generation of coherent technology is dramatically more power efficient than the last. Just look at the rapid progression of 100-Gbps coherent technology from rack-sized behemoths to svelte CFP2 modules. Cloud growth and power demands are forcing DCOs to be on the bleeding edge of optical transport technology and will continue to do so for the foreseeable future.

Disaggregation — All the Rage

Optical transport systems are a complex mix of transponders, channel multiplexers, optical amplifiers, and more. These systems are traditionally centrally managed with a single controller, with each geographic location representing a single network node. In an attempt to make optical transport products better fit the DCI application, the industry has looked at disaggregating all of the optical transport functions into discretely managed network elements (see Figure 1). However, after several false starts and lab trials, it quickly became evident that fully disaggregating everything would overwhelm the orchestration layer.

Today, DCI disaggregation efforts appear to be converging toward dividing an optical transport network into two functional groups: transponding and optical line system (OLS). The transponding functions typically reside close to the switches, routers, and servers and depreciate on a similarly fast timeline, as short as the fabled 18 months. The optical amplifiers, multiplexers, and supervisory channel functions of the OLS are sometimes separately located from the transponders, closer to the fiber entry and exit of building. The OLS is also on a much slower technology cycle, with depreciations that can last decades, though 5-7 years is more common. Disaggregating the two seems to make a lot of sense.

4 KEY FACTS ABOUT DATA CENTERS AND FIBER

1 Exciting times for a dynamic sector

Networks and data centers worldwide are growing to accommodate the digital economy.

2 Fiber-optic technology is a key driver both inside and between data centers.

10G, 40G, 100G and 400G, and even higher rates soon.

3 Fiber technology is constantly evolving

- Wideband multimode fiber (WBMMF) is under development.
- The Encircled Flux metric is well implemented and contributes to the repeatability of multimode measurements.
- Singlemode fiber is available to reach longer distances.
- Multifiber push-on (MPO) connectors are increasingly popular, and provide the density and economy required by data centers.

4 Transceiver technology is rapidly changing

- Transceivers that function at 10G or 25G line rates are currently the norm.
- The standards bodies are now working on 100G, 200G and even 400G, with upcoming rollouts.
- **Future-proofing** your fiber-optic network infrastructure is now more critical than ever.

KEEPING UP WITH TECHNOLOGIES IS ALWAYS A CHALLENGE. WHEN IT COMES TO FIBER AND HIGH-SPEED NETWORKS, TRUST THE EXPERT.

CONSTRUCTION AND OPERATIONS

- Fiber connector inspection (FIP)
- Fiber certification (OLTS)
- Fiber characterizaton and fault finding (OTDR/iOLM)
- On-site reporting

SERVICE VALIDATION

- Ethernet testing from 10M to 100G+
- Multiservice testing (Fibre Channel/SONET/SDH/OTN)
- iOptics transceivers validation

DATA CENTER INTERCONNECT (DCI) TESTING AND MONITORING

- Advanced fiber characterization (CD/PMD/OSA)
- 100G+ performance validation
- Infrastructure, network and services assurance
- Virtualization, cloud and user experience

Explore our test solutions for data centers:
EXFO.com/Data-Center

Join our LinkedIn group:
Data Center Fiber

EXFO

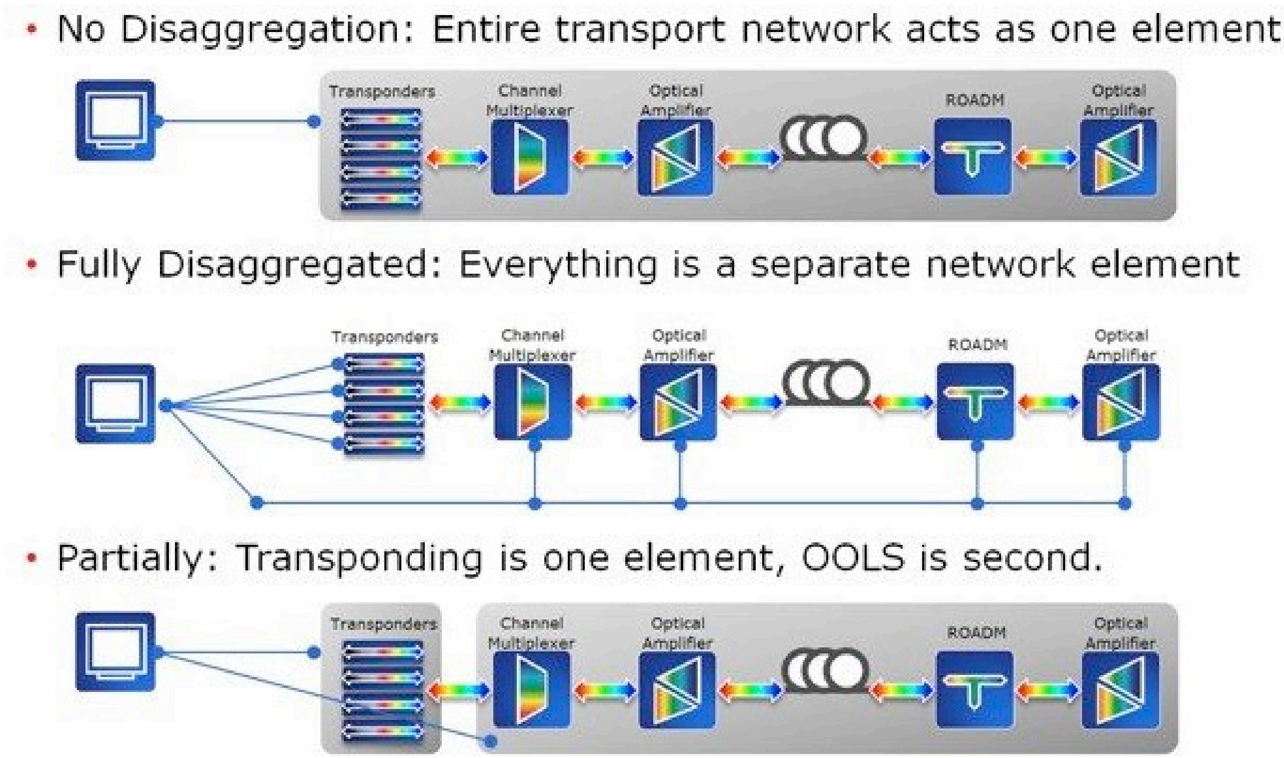

FIGURE 1. *Different optical transport disaggregation scenarios.*

At the risk of oversimplifying DCI, it involves merely the interconnection of the routers/switches within data centers over distances longer than client optics can achieve and at spectral densities greater than a single grey client per fiber pair (see Figure 2). This is traditionally achieved by connecting said client with a WDM system that transponds, multiplexes, and amplifies that signal. However, at data rates of 100 Gbps and above, there is a performance/price gap that continues to grow as data rates climb.

The Definition of "Overkill"

The 100 Gbps and above club is currently served by coherent optics that can achieve distances over 4,000 km with over 256 signals per fiber pair. The best a 100G client can achieve is 40 km, with a signal count of one. The huge performance gap between the two is forcing DCOs to use coherent optics for any and all connectivity. Using 4,000-km transport gear to interconnect adjacent data center properties is a serious case of overkill (like hunting squirrels with a bazooka). And when a need arises, the market will respond. In this case, with direct-detect alternatives to coherent options.

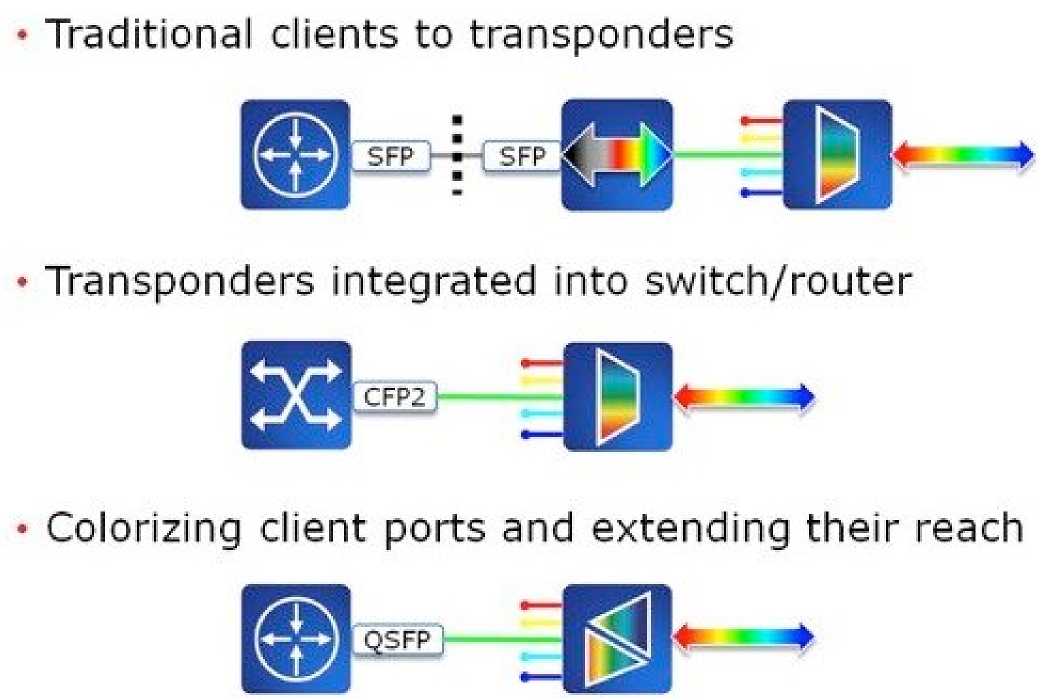

FIGURE 2. *Three ways of connecting routers.*

Direct-detect transmission is not easy at 100-Gbps and higher data rates, as all the fiber impairments that are cancelled out in the digital signal processing chain in coherent detection rear their ugly heads again. But fiber impairments scale with distance in the fiber, so as long as links are kept short enough, the gain is worth the pain.

The most recent example of a DCI direct-detect approach is a PAM4 QSFP28 module that is good for 40-80 km and 40 signals per fiber pair. Besides the great capacity this solution offers over shorter distances, it has the additional benefit of being able to plug into switch/router client ports, instantly transforming them into the WDM terminal portion of a disaggregated architecture. And with all the excitement this new solution has generated, more vendors will soon be jumping into the fray.

Tapping into the Remaining Two-Thirds

While cloud traffic is still growing almost asymptotically, most of this traffic is machine-to-machine; the growth in the total number of internet users is actually slowing down. One-third of the planet is now online (depending upon your definition), and now the onus is on the industry to figure out a way to bring online the remaining two-thirds. For that to happen, costs have to drop dramatically. Not subscription costs – networking costs. Even if the internet were totally free (or at least freemium), the costs of delivering connectivity would be prohibitive. A new paradigm is needed in networking, and that paradigm is "openness."

There is a serious misperception of open systems in the optical industry. Openness is not about eroding margins or destroying the value-add of WDM system vendors. It's about tapping into the full manufacturing capacity of our industry.

The current model of vendor lock-in cannot possibly scale to the levels of connectivity needed to reach the remaining two-thirds of the global population. That's why the Telecom Infra Project (TIP) has been launched with the aim to bring together operators, infrastructure providers, system integrators, and other technology companies. The TIP community hopes to reimagine the traditional approach to building and deploying network infrastructure. Only by working together as an industry can we hope to bring the benefits of the digital economy to all.

A Vision of an Open DCI Future

Fiber is forever — an asset planted in the ground with the unique property that its value increases as technology advances. Open optical line systems mated to that fiber create a communication highway that anyone can jump onto and ride. Open terminal systems from any and all vendors are like vehicles riding the highway – varied. And it's that variety, from a wide supply base, that will drive optical communications' continuing growth.

JIM THEODORAS is vice president of global business development at ADVA Optical Networking.

Originally published July 21, 2016

Segmenting the Data Center Interconnect Market

By **STU ELBY**

IN 2014, WHEN the first purpose-built data center interconnect (DCI) platform debuted, the DCI market was growing at 16% per year, only slightly faster than overall data center infrastructure networking market, according to Ovum's March 2016 "DCI Market Share Report." The new system addressed pent-up demand for a server-like optical platform designed specifically for DCI. Since then, sales of purpose-built, small form factor DCI products have taken off, growing at a 50-100% clip, depending on analysts' forecasts, and projected to outpace the overall data center growth rate for the next several years.

What drives this surge in spending? Multiple factors are at work, but two stand out:

- Data center to data center traffic (server-to-server) is growing much faster than data center to user due to distributed application design and replication. In addition, for performance and cost reasons, large cloud and content providers need to locate data centers closer to the end customers, which has resulted in the rapid development and growth of metro data centers. Unlike rural mega data centers, cloud providers typically build multiple smaller data centers in heavily populated metro areas due to the cost and availability of real estate and power. The result is a large number of metro data centers that require a high degree of interconnectivity.
- IT organizations have finally begun to adopt the cloud in an aggressive fashion. Whether the approach leverages software as a service (SaaS) applications such as Salesforce.com or Microsoft 365, or migration of some or all back office infrastructure into third-party data centers running infrastructure as a service (IaaS), the "chasm" has been crossed.

With this growth, the market has begun to segment, which means one size platform does not fit all (see Figure 1). To evaluate the technologies

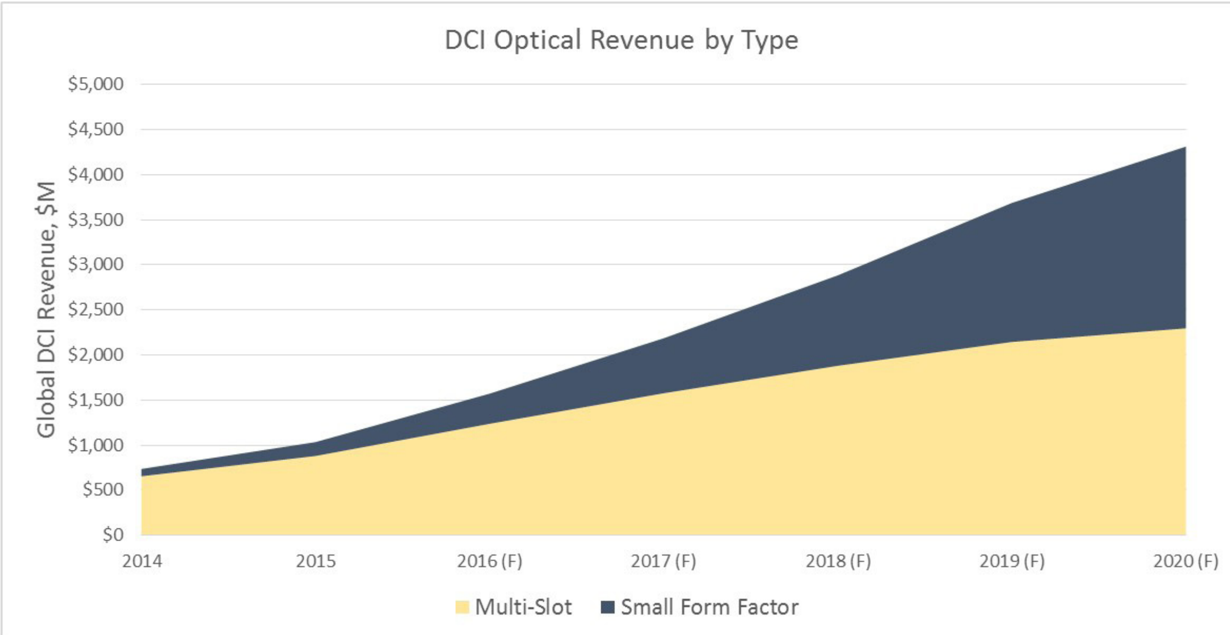

FIGURE 1. *DCI optical systems forecast by type. (Source: ACG Research, "2H-2015 Worldwide Optical Data Center Interconnect Forecast," April 2016)*

and platforms that play a role in DCI, an understanding of the customer segmentation is necessary.

DCI Market Segmentation

As Figure 2 illustrates, a small number of internet content providers (ICPs) compose the largest group of spenders on DCI. Also part of this DCI ecosystem, but to a much lesser degree, are carrier-neutral providers (CNPs). Together these two account for almost half of the total DCI spend worldwide, and over half in North America.

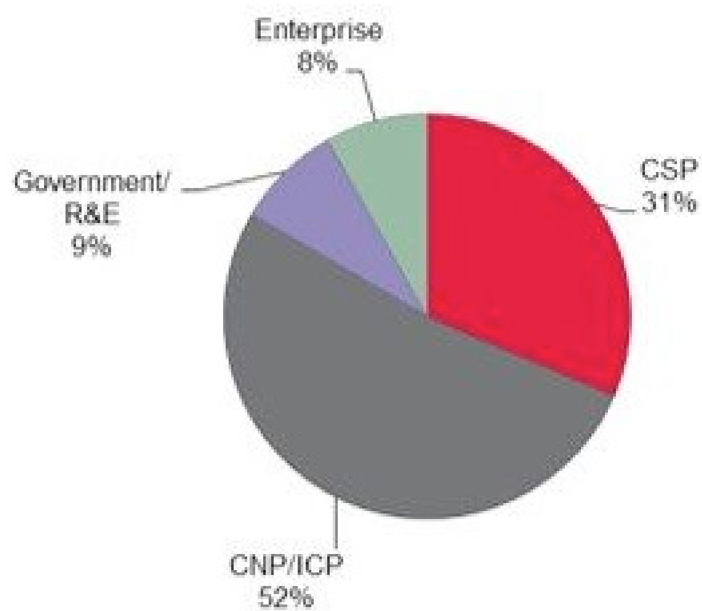

FIGURE 2. *North American 2015 DCI market segmentation by customer type. (Source: Ovum, "Market Share Report: 4Q15 and 2015 Data Center Interconnect (DCI)," March 2016)*

The vast volume of DCI traffic the largest ICPs generate has made it economically favorable for them to build their own DCI networks in many locations rather than buy transport services from service providers. We have seen a large increase in ICP-owned/operated terrestrial networks; Google, Microsoft, and Facebook have opted to build their own subsea networks as well.

Another sub-segment of this DCI market is the smaller ICPs and CNPs, e.g., AOL and Equinix. CNPs are of increasing importance in the cloud ecosystem because they provide the meet-me points where ICPs gain direct connections to the multiplicity of enterprises as well as local internet service providers (ISPs) that provide last-mile broadband access for consumers and businesses.

Communication service providers (CSPs) form another large segment of the DCI market. This group comprises a number of different provider types, such as Tier 1s, Tier 2s, cable, wholesale carriers, and even some data center operators. In particular, wholesale carriers such as Level3 and Telia Carrier are seeing a tremendous growth in traffic from ICP customers. In geographies or on routes where an ICP cannot access fiber or cannot cost-effectively build and operate its own DCI network, the ICP will purchase DCI from wholesale enterprise carriers. In these cases, ICPs still often want significant control over the operations and technologies used to carry their traffic. This is leading to interesting new business models between ICPs and CSPs:

- Managed fiber – where the CSP manages and operates a dedicated transport system (terminal gear, fiber, amplifiers, ROADMs) for the ICP. Either the CSP or the ICP can own the terminal equipment.
- Spectrum – the CSP provides a contiguous amount of spectrum (e.g., 250 GHz) as defined by filter tunings through a shared line system. The ICP can then light that spectrum using whatever optical technology it chooses.
- Wavelengths – the CSP provides a single wavelength or set of wavelengths for an ICP on a shared line system either through an add/drop multiplexer or a muxponder.

Data center operators (DCOs) include CSP subsidiaries that directly provide colocation and IaaS, such as Verizon (including assets from its Terremark acquisition) and CenturyLink (Savvis) as well as pure-play DCOs such as CyrusOne. These DCOs do not have the large volume of inter-data center traffic common in the higher tiers, but do have significant volume between their data

One-Stop Shopping

VeEX® products enable operators to deliver quality network services quickly and reliably. From DSL, Transport (DSn, SONET, OTN), and Ethernet to Fiber, VeEX offers a complete range of products for all your testing requirements. Designed with an intuitive user interface consistent across all platforms, VeEX products will reduce training costs and improve performance throughout the entire organization. Benefit from our VeSion R300 server Cloud-based management software to store your valuable data, control your assets, and keep your testers current with the latest features. No matter what network problem keeps you up at night, we have a one-stop shop solution for you.

Contact VeEX for all your testing needs.

VeEX
The Verification Experts
Website: www.veexinc.com
Email: info@veexinc.com
Tel: +1.510.651.0500
Fax: +1.510.651.0505

Fiber/CaTV solutions

centers and their customers' private data centers. Most often this connectivity is provided by a CSP via MEF Carrier Ethernet services. We also see some internet protocol (IP)/multi-protocol label switching (MPLS) virtual private network (VPN) services used for this application as well.

Another segment of the DCI market comprises enterprise, government, and research and education (R&E) organizations. Although more of this group's traffic will move to third-party clouds over time, at least some of these organizations continue to grow their own data center infrastructure. This growth occurs within a single metro area – such as an extensive university campus and its satellite offices – or across a large government agency such as the United States Postal Service or Internal Revenue Service that spans many locations across the country.

Because the traffic volume between these data centers is significantly smaller than that of ICPs (perhaps with the exception of certain government agencies), the economics of interconnection often favor purchasing CSP services rather than building dedicated DCI networks. But not always: the largest of these organizations can realize a strong return on investment by building their own DCI networks, and others may choose to do so based on other considerations such as security and control.

Differing Requirements

The market segmentation is critical to understand, as each segment has a different set of DCI requirements. Therefore, beyond cost, several features and capabilities may be prioritized differently across segments:

- Fiber capacity – the maximum bandwidth per fiber pair
- Power efficiency – measured as watts per gigabit per second
- Density – measured in gigabits per second per rack unit (RU).

Simplicity and ease of use – a true appliance (single box) approach without the need for external devices (such as multiplexers, amplifiers, or dispersion compensation modules), and with simple plug-and-play configuration and operation

Programmability and automation – the ability for the customer to use open application programming interfaces (APIs) to directly program the platforms and easily integrate them into its operations environment

Multi-protocol clients and multiple network services – the ability to support a range of interface protocols (e.g., Ethernet, Fibre Channel, TDM, optical data unit) and/or networking features (e.g., VPNs or MEF Carrier Ethernet services).

Table 1 illustrates the relative importance of each attribute in platform selection for each customer segment.

	Large ICPs	Smaller ICPs/CNPs	CSPs	Government, Enterprise
Fiber Capacity	HIGH	MEDIUM	HIGH	LOW
Power Efficiency	HIGHEST	HIGH	HIGH	HIGH
Density	HIGH	HIGH	HIGH	LOW
Simplicity and Ease of Use	HIGH	HIGH	MEDIUM	HIGH
Programmability and Automation	HIGHEST	HIGH	MEDIUM	MEDIUM
Multi-protocol, Multi-service	LOW	HIGH	HIGH	HIGH

Table 1: Highest Priorities for Data Center Interconnect Systems, by Customer Segment

The largest operational expense for an ICP is power, so power efficiency is typically the most critical DCI attribute following cost. Programmability is another critical attribute for the top ICPs as they have painstakingly optimized and automated their abilities to manage their infrastructures and grow capacity through internally developed operations systems. Any new platforms entering their environments must be able to quickly (read: weeks, not years) and easily (read: a few developers, not a 100-person team) be integrated into their systems.

Fiber capacity and density are important to ICPs, as they drive space and dark fiber costs, but are not as critical as power efficiency and programmability. Services support is the least critical factor for ICP optical transport DCI platforms. All DCI client-side interfaces have migrated or are migrating to 100 Gigabit

Ethernet. There are no additional services required at the transport layer, and all packet protocols and traffic management are handled at the application layers within the data centers.

CNPs are concerned with power and space, but unlike the large ICPs, CNPs require the support of a range of services that may place additional requirements on the DCI transport equipment. For example, some CNPs provide Ethernet services (e.g. Ethernet Virtual Private Line and multiplexed virtual LAN-based services) to enable their customers to interconnect colocation sites between metro data centers, interconnect to cloud and "x as a service" providers via cloud exchange services, and peer with Border Gateway Protocol (BGP)/IP via internet exchange services. Programmability is also important, as CNPs are striving for zero-touch provisioning and operations automation.

Fiber capacity is less important, as it is rare that the inter-data center traffic of the CNPs exceeds several terabits per span. Remember that much of the traffic in and out of a CNP data center is carried by the carriers/CSPs colocated in that data center.

CSPs often prefer a chassis-based system versus the small form factor platform ICPs prefer. However, this may change as CSPs adopt more of a scale-out ICP infrastructure approach. CSPs are concerned with space, power, and density but they absolutely need a rich set of protocols and interfaces to handle legacy services. CSPs are interested in ease of use but they typically have large, skilled optical operations teams to handle configuration and operation and can deal with more complex products. For this same reason, while they desire programmability and automation they are willing to adopt it more slowly and with greater patience than ICPs.

Enterprises have a similar concern for power efficiency as CNPs, but in this case it is because they must pay a monthly fee to the CNPs where they colocate based on the maximum power per rack they can draw. It behooves the enterprise customer to use the most power-efficient equipment on the market. Service support also is key to the enterprise as they are likely to still require native Fibre Channel support and potentially other variants like TDM or Infiniband.

Programmability is less important, as the scale of an enterprise cloud does not typically drive the level of operations automation required by the larger cloud providers. Fiber capacity is also not typically relevant, as the overall capacity per spans is measured in a few hundred gigabits per second.

However, over the last year several newer features have bubbled up to become highly desirable and even must-have for an increasing number of customers across all market segments:

- In-flight data encryption at both Layer 1 and Layer 2 (MacSEC)
- Link Layer Discovery Protocol (LLDP) for the automated detection of routers and switches
- Configuration/provisioning APIs based on YANG data models for automation and programmability.

The Need for Platform Diversity

Net-net, these differences in market drivers make it difficult to develop a single DCI platform that can serve all markets optimally.

One approach, which we refer to as "telco" or "general purpose," is a chassis-based system in which a customer can mix and match client interface cards and line cards. This approach serves the enterprise and CSP market well, but the added cost, power, and size of a chassis system is not aligned with what the large ICPs' need. Likewise, a pure Ethernet-over-DWDM pizza box can meet the requirements of the ICPs but may fail to serve the enterprise market in terms of service types and bandwidth granularity.

A developer of DCI equipment will either need a full portfolio of DCI products or will need to pick which segment of the market to address.

STU ELBY is senior vice president, Data Center Business Group, at Infinera.

Originally published July 22, 2016

Data Center Interconnects: The Road to 400G and Beyond

By **JUSTIN ABBOTT** and **DAVID HORN**

THE RAPID GLOBAL growth in cloud computing is driving massive demand for high-capacity data center infrastructure. While fixed and mobile broadband services primarily drove recent network bandwidth upgrades, the latest wave derives from a rapid increase in machine-to-machine traffic due to an expansion in server virtualization and software defined networks (SDNs). The result is a dramatically different growth rate, illustrated in Figure 1, that is outpacing consumer traffic expansion.

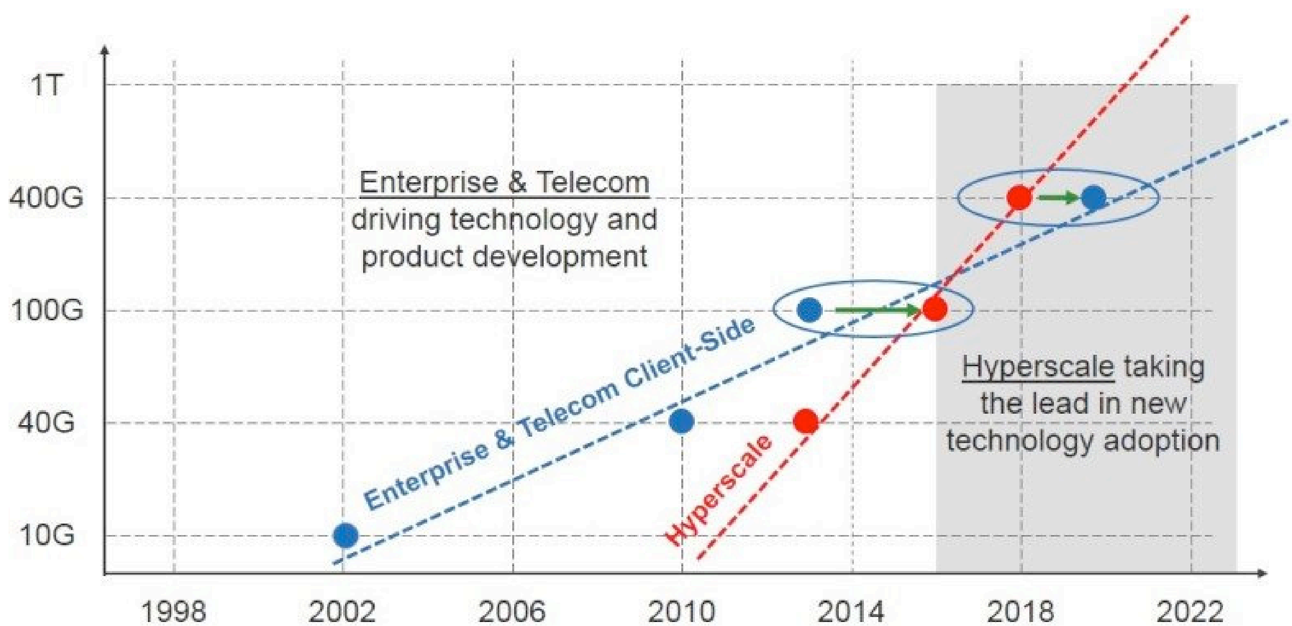

FIGURE 1. *Telecom and datacom adoption timelines.*

Similarly, the telecom/enterprise market traditionally has led early adoption and hence influenced the development of new module technology. However, the new hyperscale growth of data center interconnects has changed the dynamic of module innovation and deployment as well.

Hyperscale data centers are changing the game

Telecom/enterprise applications first adopted 100G technology in the form of CFP modules. Data centers generally did not adopt 100G interfaces until the technology matured and evolved towards denser, lower power interfaces, particularly in the form of QSFP28 modules.

However, as the hyperscale data center market scales to keep pace with machine-to-machine communications needs, data center operators have become the first to demand transmission modules for data rates of 400G and beyond. The drive to meet this requirement will lead to a couple of fundamental shifts in the market. First, as the data centers replace telecom/enterprise as the early adopters of new technology, that emerging technology will reflect their requirements. Second, those requirements are radically different from those typically associated with telecom. The telecom market generally required specific environmental, reliability, and interoperability/standardization features. But 400G optical modules will be focused on the needs of the data center and its associated challenges: high faceplate density, low power consumption, ever lower cost per bit, and reliable large-scale manufacturing capabilities.

Smaller, faster, lower cost per bit optical modules

Meeting these different requirements will take a variety of strategies. To enable maximum density, data center modules are ideally sized such that the number of modules in one rack unit (1RU) is aligned with the throughput of the switch ASIC inside that unit. For example, in existing 100G data center switches, 32 QSFP28 modules fit in 1RU, which aligns with a 3.2-Tbps switch ASIC. This trend will continue as switch ASICs scale to approximately 12-14 Tbps interfaces on a single 1RU faceplate.

The concept of multiple lanes to increase bandwidth has been widely used for 20 years in line-side transmission networks since the advent of DWDM. Adoption in shorter reach data center applications has been more recent, starting with 40G and now 100G modules using either multiple fibers or WDM. Future modules will need to strike a balance between number of lanes, bandwidth (baud rate), and bits per symbol to fit within the allowed size and power envelope while achieving the lowest possible cost. Generally, the fewer lasers and receivers required, the lower the cost will be.

There are several core technologies that will enable 400-Gbps interfaces, as well as higher data rates, within these balanced design parameters. One is the use of higher order modulation. Today's data center interconnect transceivers are based on non-return-to-zero (NRZ) modulation, which transmits one bit per symbol. The next level of modulation provides four-level pulse amplitude modulation (PAM4), which transmits two bits per symbol, therefore doubling the data rate without doubling the required overall bandwidth over conventional NRZ (see Figure 2).

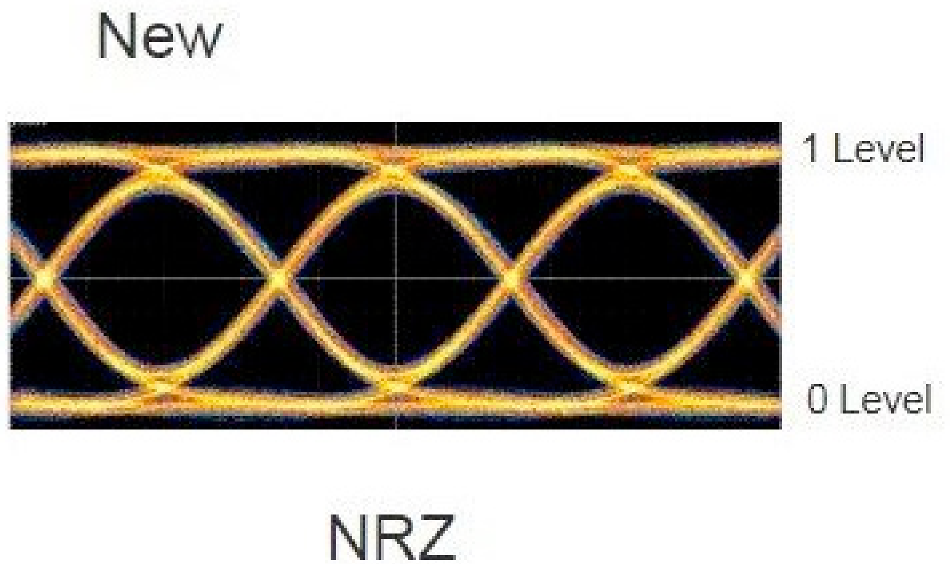

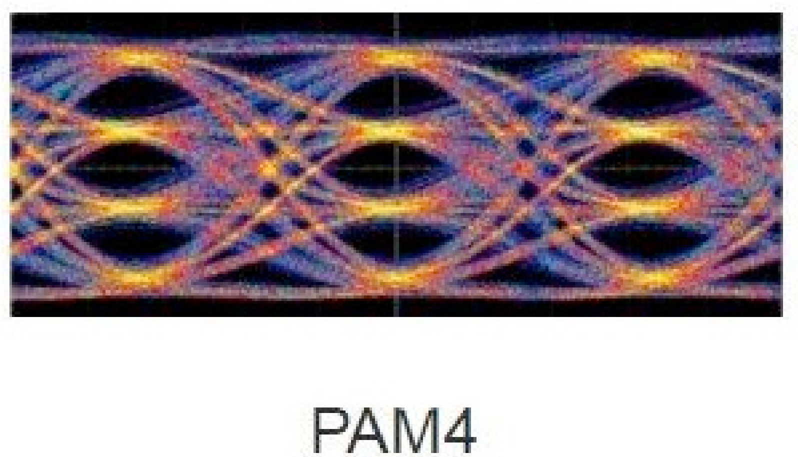

FIGURE 2. *NRZ and PAM4 modulation. (Source: http://globaltek.us.com/awg6010/)*

The consequence for this move to additional levels per symbol is a requirement for higher signal-to-noise ratio (SNR). To achieve higher SNR, PAM4 components require greater linearity in the modulated laser source and detector, higher extinction ratio, low relative intensity noise (RIN), greater output power, and higher responsivity. These characteristics will drive the development and choice of components to provide the necessary performance while still achieving the cost and density needed, even as hyperscale applications require enhanced loss budgets that extend beyond the standard 3-5 dB.

Components and packaging

For example, one of the key components that drive performance is the laser and its modulation. In existing <2 km 100G NRZ applications, both directly modulated lasers (DMLs) and externally modulated lasers (EMLs) provide acceptable performance. In the former, the drive current to the DML is varied to provide different levels of light amplitude representing the 1s and 0s. EMLs have a fixed laser driver current; the light is modulated by varying the voltage to the monolithically integrated electro-absorption modulator or Mach-Zehnder modulator, which then modulates the laser light.

In moving from NRZ to PAM4, the variation of the DML drive current intensity leads to a lack of linearity, additional noise, and low extinction ratio. There are emerging DML technologies that promise improved performance through a photon-photon resonance (PPR) effect between the distributed Bragg reflector (DBR) and the distributed feedback (DFB) laser structures. The PPR effect extends the modulation bandwidth beyond that of the DFB alone. However, the additional complexity of the grating and electrode structure will increase the component cost and reduce yield, while the reliability risks will remain a concern in these aluminum-containing structures. Therefore, it is too soon to determine if these devices can compete with the maturity, performance, and lower cost that EML technology delivers.

In addition to high performance, 400G transceivers also must support low cost per bit and the ability to reliably and efficiently scale to high-volume manufacturing. This has driven innovation in transceiver component design with a focus on leveraging the advantages of integration, automation, wafer-scale production, and non-hermetic packaging.

Such innovation is apparent in the receive optical subassembly (ROSA) and the transmit optical sub-assembly (TOSA). The cost of TOSAs and ROSAs is driven by the assembly of discrete components, alignment, burn in, and high cost of yield loss at the subassembly/module level. To address this, new TOSA and ROSA designs are emerging that leverage the use of wafer-level integration in assembly, packaging, and testing, based on both silicon photonics and complementary techniques within indium phosphide (InP).

One Test Platform - RF to Optics

Think RXT

Broadband Service Providers are challenged to expand the capacity of their Hybrid Fiber-Coax (HFC) networks. Evolutionary modifications are now well underway to ensure the HFC access network continues to enjoy a long lifetime and is ready to deliver the insatiable demands of CaTV subscribers.

Contact VeEX to discover how the RXT is truly the industry's only All-in-One RF to Optics test platform and how this tester can help you deploy DOCSIS 3.1 and other emerging technologies.

Silicon photonics offers the use of mature, large-scale processes compatible with complementary metal–oxide–semiconductor (CMOS) technology, today's standard for building integrated circuits, to precisely generate thousands of optical components on a monolithic substrate in a fully automated manufacturing environment. Enabled elements include optical waveguides, splitters, couplers, modulators, detectors, multiplexers, and demultiplexers.

In practice, silicon photonics components are defined through CMOS processes that involve lithographic patterning, material deposition, removal, and modification. These processes are repeated to build up the desired structure of the optical components and circuitry. Once complete, the wafer containing a patterned grid of devices can be burned in and tested before singulation. Testing at this comparatively earlier, lower cost point in the manufacturing process improves yield versus conventional photonic device manufacture.

Unfortunately, silicon has two photonic design challenges. A weak electro-optic effect makes silicon active elements like modulators relatively complex. Silicon's indirect bandgap structure also has not produced effective gain structures for lasers or optical amplifiers.

To resolve these issues, materials such as InP with higher electro-optic effect and well characterized optical lasing and amplification properties must be integrated into the monolithic silicon component. There are several ways to do this, including optical coupling between InP silicon waveguide edges, fiber coupling of the laser/modulator source to the silicon, and flip-chip integration of the laser to the silicon. The wafer-level integration of InP performance with the low-cost aspect of silicon photonics provides a platform that is well positioned to address the cost and density needs of the hyperscale market.

Silicon photonics offers additional benefits as well. For example, historically TOSA/ROSAs have been hermetically sealed to protect materials and free-space optics from environmental contamination that could reduce performance and reliability. Sealing is a time consuming and expensive process. Silicon is "self-hermetic" and therefore does not require hermetic packaging. This attribute greatly reduces the constraints on the design, materials used, and fabrication complexity required to build optical subassembly (OSA) packages.

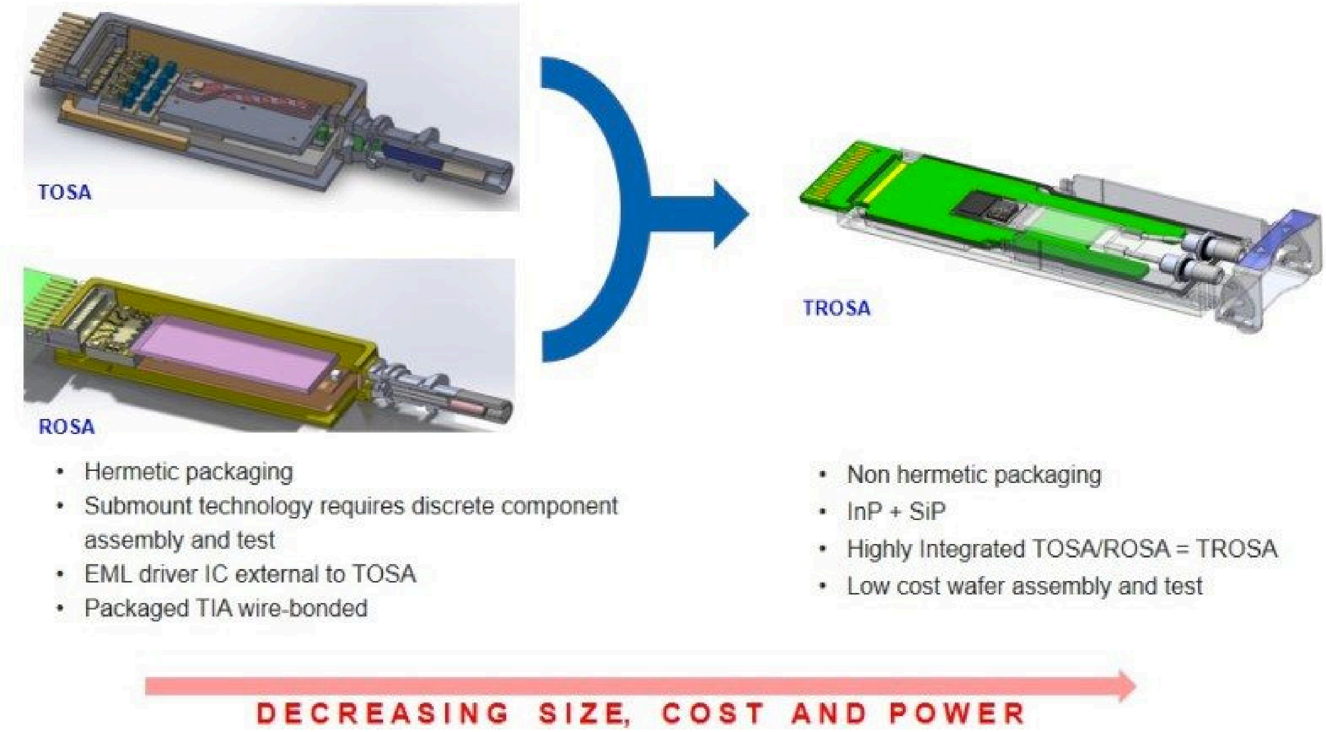

FIGURE 3. *Evolution of module optics.*

Some materials, including InP, will still need hermetic protection. But there are several ways to achieve this cost-effectively on a chip or wafer level that preserves wafer-level packaging and test. As a result, such OSAs use less material, require fewer process and test steps, and produce higher yields through final assembly and burn-in, all of which results in lower cost.

Silicon photonic OSAs also can be made small enough to be assembled into smaller transceiver form factors that can increase faceplate density. Figure 3 illustrates an example of the evolution of hermetically sealed TOSAs and ROSAs to InP and silicon photonics based optics.

Building the foundation for future data center transceivers

In summary, data centers are leading the migration to 400G interconnects, given the high growth in machine-to-machine capacity demand. The reduced size, low cost, and low power requirements of data center applications necessitates performance-oriented components that support the use of higher-order modulation, simultaneously with highly integrated and wafer-level packaging techniques to achieve critical low device costs.

The use of silicon photonics transforms discrete components and InP elements into monolithic structures with small, less complex packages. Component developers have begun to merge industry-proven InP laser/modulator sources with wafer-level silicon photonics integration and packaging techniques to provide the required performance and low cost, power, and size. Combine such designs with high-volume manufacturing and testing experience, and the new generation of components creates a technology foundation that will enable low-cost transceivers for the emerging data center market.

JUSTIN ABBOTT is senior product line manager for next generation products and **DAVID HORN** is director, datacom product and technology strategy, at Lumentum.

Anritsu

envision : ensure

Company Description:

Anritsu Company (www.anritsu.com) is the United States subsidiary of Anritsu Corporation, a global provider of innovative communications test and measurement solutions for more than 110 years. Anritsu provides solutions for existing and next-generation wired and wireless communication systems and operators. Anritsu products include wireless, optical, microwave/RF, and digital instruments as well as operations support systems for R&D, manufacturing, installation, and maintenance. Anritsu also provides precision microwave/RF components, optical devices, and high-speed electrical devices for communication products and systems

LINKS:

- NEW PRODUCT: MT1000A Network Master Pro
- Portable PIM Test Analyzer
- MT9083x2 Series with Fiber Visualizer
- Network Master MT9090A
- LMR Master Land Mobile Radio Modulation Analyzer S412E

Company Description:

EXFO provides communications service providers (CSPs) with test orchestration and performance intelligence solutions to ensure the smooth deployment, maintenance and management of next-generation, physical, virtual, fixed and mobile networks. The company has also forged strong relationships with network equipment manufacturers (NEMs) to develop deep expertise that migrates from the lab to the field and beyond. EXFO is no. 1 in portable optical testing and boasts the largest active service assurance deployment worldwide.

LINKS:

- Data Center Fiber LinkedIn Group
- EXFO Data Center Resource Center
- Next-generation, 100G, advanced multiservice test solution
- Fiber Certifier OLTS
- Intelligent Optical Link Mapper (iOLM)

Company Description:

Lumentum is a market-leading manufacturer of innovative optical and photonic products enabling optical networking and commercial laser customers worldwide. Lumentum's optical components and subsystems are part of virtually every type of telecom, enterprise, and data center network. Lumentum's commercial lasers enable advanced manufacturing techniques and diverse applications including next-generation 3D sensing capabilities. Lumentum is headquartered in Milpitas, California with R&D, manufacturing, and sales offices worldwide.

LINKS:

- https://www.lumentum.com/en/optical-communications/products/optical-transceivers
- https://www.lumentum.com/en/products/qsfp28-cwdm4-optical-transceiver-100ge
- https://www.lumentum.com/en/optical-communications/products/transmission-components-and-modules/trueflex-transport-modules/sdn-whiteboxes
- https://www.lumentum.com/en/media-room/videos/ofc-2016-100g-transceiver-portfolio-100g-live-data-demo
- https://www.lumentum.com/en/media-room/videos/ofc-2016-optical-whitebox-platform

Company Description:

VeEX Inc. develops innovative test solutions for next generation communication networks and services. With a blend of advanced technologies and vast technical expertise, VeEXÆs products address all stages of network deployment, maintenance, field service turn-up, and service verification features across DSL, Fiber Optics, CATV/DOCSIS, Mobile, next-generation Transport Network, Fibre Channel, Synchronous and Carrier Ethernet technologies. For more information visit www.veexinc.com.

LINKS:

- Link to our main website
- Link to our RXT-4111 DWDM OTDR Module Video
- Link to BTR Special Report, "Outside Plant Must Evolve to Meet Bandwidth Growth"
- Link to "Fiberizer Cloud-Based Fiber Optic Data Analysis Video"