

Factors driving PSM4 and silicon photonics for data center architectures

January 28, 2014

By Brent Hatfield, Molex Inc.

Multiple forces and trends are driving the implementation of parallel single-mode quad (PSM4) and [silicon photonics](#) (SiPh) in next-generation data center designs and architectures. PSM4 and SiPh fill an important gap in data center interconnect options – the gap between the reach of 25G multimode options and that of long range (LR) optical approaches. Using pigtail designs with these technologies can ease implementation via easy mating to structured cabling. From a “future proofing” perspective, PSM4 is an approach that can accommodate both current and future bandwidth upgrades.

The cloud

While most consumers have used the cloud in one form or another for many years (Flicker, iTunes, etc.), one recent major trend has been the aggressive proliferation of cloud use by private enterprises. A 2011 study by IBM found that 70% of mid-sized businesses were using cloud-based analytics and 66% had either already deployed or were planning to deploy cloud-based technologies.

Rather than increase IT capacity internally, moving items to the cloud has proven to be both cost-effective and timely. Microsoft recently reported that its Skydrive service has 17 million customers that store 10 petabytes of data. The trend [continues to accelerate](#).

Entertainment on-demand and entertainment mobility

From Hulu to Apple TV to Amazon Prime, Netflix, and others, entertainment is transitioning from the set-time model to an on-demand model. Furthermore, entertainment is now viewed as a product that is expected to be available not just when you want it but also where you want it.

Multiple viewing/experience platforms and options exist. Internet-connected televisions, tablets, PCs, and smartphones enable viewers to access any level of information or entertainment at the place and time of their choosing. This, in turn, requires large-scale infrastructure to store and support programming available on-demand by consumers.

Big Data

[According to Roberto V. Zicari of Object Database Management Systems](#), “Every day, 2.5 quintillion bytes of data are created. This data comes from digital pictures, videos, posts to social media sites, intelligent sensors, purchase transaction records, cell phone GPS signals to name a few. This is Big Data.”

Previously most of the data was released and deleted, viewed as having little value. Or, at a minimum, it was viewed as too cumbersome to store and use. However, recent new technologies in database management (most notably the development of Hadoop) enable large-scale data-based relational information, even where the data is in disparate types, formats, and across different systems. Essentially, all data has become both manageable and potentially useful. As a

result, all data has begun to be stored indefinitely and mined for information and relationships.

New data center architectures

As a result of these multiple factors, data centers are becoming larger and more distributed. Virtualization means that system architectures that once took additional hardware can be built with less equipment, but with more interconnections for ensuring resources are fully used and interactive. Next-generation data center designs typically require two main attributes: scalability and uniformity of performance coupled with low latency. Scalability is required to handle the expansion of services, addition of customers, and the increase in data. Uniformity of performance is required to provide a smooth flow of data between nodes.

The traditional three-tier data center architecture of core, aggregation, and access was well suited for traditional use of email, webpages, and traffic. However, new requirements for video and content delivery, virtual machines, cloud access, and social networking content assembly required low latency across the data center. Non-blocking bandwidth became both a requirement and a technical challenge that demanded a new data center architecture.

“Spine and Leaf” architectures resolve these two issues. A central spine handles high-bandwidth data between leafs, whereas a leaf controls the traffic flow between clusters of servers. The performance is thus balanced, while the structure supports the addition of leafs as the system scales. Traffic between nodes is balanced and accessible with low-latency east-west traffic flow.

This type of architecture has many attributes; one physical one is the need for longer links. And that could mean a significant change in cabling.

Multimode fiber and the distance/data rate challenge

In the last decade, copper high-speed interconnects have faced a primary challenge. As data rates increased, the distances that copper interconnects could accommodate decreased to the point that alternative approaches had to be found. To handle increased data rates, new and exotic ways were developed to expand the reach of copper: higher awg cable, new dielectric materials, and new equalization schemes, CDRs, and other active signal integrity devices. Each method came with increasing costs and implementation challenges. Ultimately, the crossover point occurred as [multimode fiber](#) became the only viable solution for exceeding a certain combination of data rate and distance.

Over time, however, the same challenges have arisen with multimode optics, propelling demand for more exotic fiber and connector constructions:

- 1GBase-SR with OM2 - fiber reach is 550 m
- 40GBase-SR4 with OM3 – fiber reach is 150 m
- 100GBase-SR4 with OM4 - fiber reach is 100 m.

Each new upgrade requires significant data center down time, including additional price premiums for structured cabling. And the time between these upgrades is shrinking. 10 Gigabit Ethernet (GbE) is deployed, and 40GbE systems are being installed now. We’ll soon see 100GbE begin deployment, and a consortium of suppliers announced an MSA for 400GbE last year.

Terabit Ethernet is on the horizon. The traditional economic model of replacing an outdated infrastructure and expecting pay off within a few short years is increasingly under pressure.

Bandwidth and technical challenges have been aptly referred to by some as a “coming data tsunami.” The other significant challenge is rising cost – and consumer expectations. Rightly or wrongly, consumers have come to expect noticeable performance increases with little or no cost increase. From photo transfer, to video, to HD video, to “everything available on demand through every device,” customers want more but do not want to pay more. This places increased pressure on providers to lower costs.

At a time of increasing cost pressures, constant multimode cabling upgrades run counter to industry needs.

A solution: SiPh PSM4 for 4-km reach

A SiPh-based PSM4 approach helps resolve the multiple challenges of needing higher bandwidth, longer distances, low power, and future proofing. SiPh-based long-reach PSM4 products offer seven primary advantages:

1. Distance: With transmission distances up to 4 km, they can accommodate most, if not all, new data center requirements.
2. Power: SiPh-based active optical cables (AOCs) use basically the same power consumption as VCSEL-based products: under 1 W for a 10G QSFP+ AOC and 1.5 W at 25 Gbps.
3. Cost: Cost is roughly the same as VCSEL-based AOCs, but as they are singlemode they use much less expensive singlemode cable as a transmission medium. As VCSEL-based product speeds increase, they require ever more expensive types of fiber to transmit effectively. After most networks have upgraded to OM3, another upgrade to OM4 has serious cost implications. And what comes after OM4, and how soon will it be needed? With singlemode SiPh AOCs the fiber stays inexpensive singlemode – and consistent as data rates increase.
4. Future-proof structured cabling: This is becoming an increasingly important factor in driving the adoption of SiPh in the market, particularly in new installations. Upgrading structured cabling is less economically viable as speeds continue to increase. PSM4 can accommodate:
 - 10GbE – 4x10 Gbps
 - 40GbE – 4x10 Gbps aggregated
 - QDR Infiniband – 4x10 Gbps
 - FDR Infiniband – 4x14 Gbps
 - EDR Infiniband – 4x25 Gbps
 - 100GbE – 4x25 Gbps aggregated.
5. Pigtail plug-in options for ease of installation and upgrades: PSM4 AOCs are available in pigtail versions. With a pigtail, one end is a traditional QSFP+ active interface, and the other either MPOs or LCs. This enables a quick connection to structured cabling (so the AOC side doesn't have to be pulled long distances) and the ability to quickly upgrade to new products as needed. So when the 4x28G zQSFP product is introduced, it can be immediately connected to structured cabling already in place. Upgrades from

10GbE/40GbE to 100GbE can be done immediately with no costly structured cabling upgrade and associated data center downtime.

6. SiPh actives becoming broadly available: An ecosystem of proven, reliable, cost-effective SiPh devices exists to accommodate singlemode data center links. SiPh is widely deployed today in PSM4 AOCs, and several new entrants into the market are offering a variety of options.
7. SiPh devices offer a clear technology path beyond 25G: VCSELs face increasing design challenges as speeds increase. On the other hand, with SiPh-based systems most modulating schemes have a clear and well-understood technology path to 50G, 100G, and beyond.

The future is now

Until recently, it was generally believed that bandwidth/distance design challenges were limited to high speed copper interconnects; most thought the transition to fiber optics in the nineties resolved these issues. Few predicted that traditional VCSEL-based optics would begin to experience some of the same challenges as copper interconnects did over a decade ago. But they have.

Singlemode SiPh bridges the gaps caused by the dual requirements of longer distances and higher data transmission speeds. For many customers this technology provides a lower-cost, lower-power option for what can be referred to as medium-reach distances that is future-proof for the next generations of data transmission speeds. With a broad base of products both available and being developed by multiple suppliers, PSM4 and SiPh AOCs can be deployed today to enable new data center architectures, with the assurance that they'll provide a path to cost-effectively meet future generation upgrade requirements.

Brent Hatfield is product manager, Fiber Optics Division, at [Molex Inc.](#)