Introduction

A key eco-system component impacting the Data Center is the emergence of Silicon Photonics technology. This new technology can change how data center systems exchange data and make rack equipment leaner by reducing total cost of ownership and removing networking bottlenecks that can result in stranded compute capacity. Specifically, Data Center architectures can be made significantly more efficient and scalable by bringing silicon photonics to switching nodes.

Silicon Photonics is widely accepted as a key technology in next-generation systems and data interconnects. It brings the advantages of both semiconductor integration and photonics—high data densities and transmission over longer distances—in a platform where high levels of integration can be achieved with low manufacturing costs using conventional silicon integrated circuit infrastructure.

Developments are underway to move switching functions into the optical domain to relax the digital functionality in CMOS and vastly increase Spine and Leaf switch node throughput. High-density integration of silicon photonics with Spine and Leaf switching electronics and use of additional degrees of freedom available only in the optical domain enable a dramatic increase in the bandwidth per fiber and scalability of the data processed in the switch fabric. Data Center engineers and architects are planning for large-scale deployment of silicon photonics-based switching appliances in upcoming hardware replacement cycles. The general market assumptions are that silicon photonic interconnects costs will drop to $1/Gbps due to high-density integration and WDM. Therefore, Silicon Photonics is changing the rules of Data Center IT equipment power consumption and associated ancillary equipment requirements.

Photonics changes the rules for Bandwidth-per-Watt

PHOTONICS:

- Modulate/receive ultra-high bandwidth data stream once per communication event.
- Broadband switch routes entire multi-wavelength stream.
- Off-chip Bandwidth = on-chip Bandwidth for nearly same Power(P).

ELECTRONICS:

- Buffer, receive and re-transmit at every device.
- Each bus lane routed independently. ($P \propto N_{LANES}$)
- Off-chip Bandwidth requires much more power than on-chip Bandwidth.

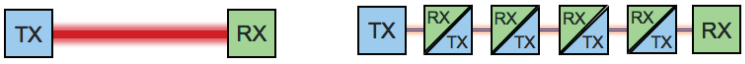

As silicon photonics becomes ubiquitous throughout the data center switching fabrics, organizations that once considered 1-Gigabit Ethernet (1GE) or 10GE sufficient bandwidths will be driven to deploy 40GE, 100GE, 200GE and beyond. The use of fiber optic cable throughout data centers is quickly becoming a reality.

The Yang – Managing the Fiber Interconnects

Adoption of Data Center architectures such as Spine and Leaf are driving the scale of fiber cabling infrastructure, requiring greater fiber counts. As this architecture is deployed in Data Centers, fiber counts multiply very quickly when compared with traditional three-layer distribution architectures. Fiber optics cabling is the best medium to support these always increasing speeds of Silicon Phonics equipment, so as not to require an infrastructure refresh before full depreciation of the Layer-0 infrastructure. Fiber also provides longer runs than copper with lighter cable weights, can resist electromagnetic interference (EMI), and deters hacking.

Optical Trends

- Optical software revolution – SDN/NFV
  - Programmable optical networking
- Faster adoption of 100G than expected driven by long haul and Data Center Interconnectivity
- Expected rapid adoption to 200G/400G and beyond
  - More coherent networks
  - Broader range of modulation formats
- New optical networks are being built
  - Mobile (Fronthaul/Cloud-RAN)
  - DCI networks (cloud operators)
- Changing dynamic in the optical industry
  - Web 2.0, mobile…

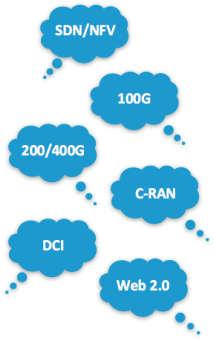

The only sure way to manage the exponential growth in fiber counts is by using an automated fiber cross-connect, inserted as an eco-system component of a Data Center POD within the deployment architecture, creating a "Flexi-Pod" deployment approach. As part of the eco-system to help simplify fiber deployments, advanced Data Center products such as Robotic Optical Cross-Connects, support either single-mode or multi-mode fiber plants.

While the ability to keep cabling channels in check will reduce cooling power consumption, intelligent Layer-0 infrastructures using Robotic Automated Optical Cross-Connects will also reduce power needs of active network equipment. When designed with a central robotic automated optical cross-connect zone, an intelligent Layer-0 infrastructure architecture can help ensure that all Silicon Photonics based switch ports are utilized — decreasing the power needs for electronics by keeping unused ports to a minimum.

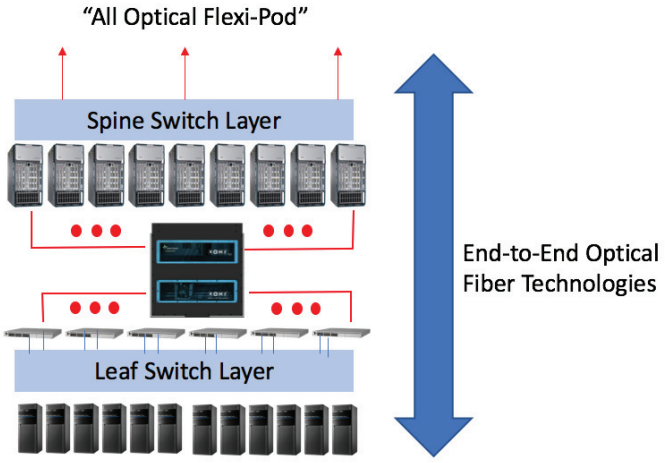

Incorporating an automated optical cross-connect with ultra-low latency enables the instant delivery and provisioning of any-to-any physical fiber optic connections, and automation provides the ability to virtualize these connections via software. An automated connectivity fabric makes Data Center Layer-0 networks dynamically reconfigurable, with security, speed and accuracy that traditional manual processes cannot match. This on demand connectivity of Silicon Photonics based switching products opens the door to a broad range of new applications within the Data Center ecosystem. By inserting Wave2Wave's ROME automated Robotic Automated optical cross-connects products into next generation Silicon Photonics equipment based Data Center architectures, operators can simplify Layer-0 infrastructures, reduce operational costs, and improve redundancy and resiliency.

