Optical Networks with ODIN in Smarter Data Centers: 2015 and Beyond

Dr. Casimer DeCusatis
Distinguished Engineer & CTO, System Networking Strategic Alliances
IBM Corporation, Poughkeepsie, NY

Overview

- Motivation & Framework for Enterprise Optical Data Center Networking
  - Integrated, Automated, Optimized: Solving Problems in Traditional Data Center Networks
  - Optical/Copper Link Applications: Cost and Power Tradeoffs
  - Virtualization & Software-Defined Networking
- High Performance Computing Applications
  - The Road to Exascale Computing
- Research Topics
  - Silicon Photonics
  - Optical PCBs
  - Advanced Packaging Technologies
- Conclusions

The Need for Progress is Clear

30 percent

Energy costs alone represent about 30% of an office building's total operating costs

42 percent

Worldwide, buildings consume 42% of all electricity – up to 50% of which is wasted

85%

In distributed computing 85% of computing capacity sits idle

18%

Anticipated annual increase in energy costs

20x

Growth in density of technology during this decade. Energy costs higher than capital outlay

50%+

More than half of our clients have plans in place to build a new data center/network facilities as they are out of power, cooling and/or space

Landscape of Interconnect Technology

From A. Benner, “Optical interconnect opportunities in supercomputers and high end computing”, OFC 2012

PHYSICAL Link Types

Distinguished by Length & Packaging

PHYSICAL Link Types	MAN & WAN	Cables – Long	Cables – Short	Backplane / Card-to-Card	Intra-Card	Intra-Module	Intra-chip
Length	Multi-km	10, - 300 m	1 m - 10 m	0.3 m - 1 m	0.1 m - 0.3 m	5 mm - 100 mm	0 mm - 20 mm
Typical # lanes per link	1	1 - 10s	1 - 10s	1 - 100s	1 - 100s	1 - 100s	1 - 100s
Use of optics	Since 80s	Since 90s	Since late 00's	Since 2010-2011	2012-2015	After 2015	Later

LOGICAL Link Types

LOGICAL Link Types	Internet	Local Area Network	Cluster / Data Center	Storage Area Network	Direct Attach Storage	I/O	Mezzanine Bus	SMP Coherency Bus	Memory Bus
Distinguished by Function & Link Protocol	Traffic: IP Stds: Ethernet, ATM, SONET,	Traffic: HTML pages to laptops,.. Stds: 1G Ethernet, WiFi	Traffic: Intra-application, or intra-distributed-application Stds: InfiniBand, 1G Ethernet, 10/40/100Enet	Traffic: Read/Write to disk, shared Std: Fibre Channel	Traffic: Read/Write to disk, unshared Stds: SAS, SATA	Traffic: Load/store to I/O adapters Stds: PCI/PCIe	Traffic: Load/store to Hubs & bridges Stds: Hyper-Transport	Traffic: Load/store coherency ops to other CPUs' caches Stds: Hyper-transport	Traffic: Load/Store to DRAM or Memory Fanout chip Stds: DDR3/2/.
Key Characteristic	Inter-operability with “Everybody”	100-300m over RJ-45 / CAT5 cabling, or wireless	BW & latency to <60 meters	Dominated by FC	Shared tech. between servers & desktops	Shared tech between servers & desktops	Reliability	Reliability, massive BW, reliability	Reliability & cost vs. DRAM
Use of optics	Since 80s	Maybe Never? (Wireless, Building re-wiring, BW demand)	Since 2000s	Since 90s	Not yet	Scattered	Not yet	Coming	Coming later

Link Technology: Single-mode Optics | Mixed multi-mode optics & copper | Copper

Data Center Interconnect

PHYSICAL Link Types

Distinguished by Length & Packaging

	MAN & WAN	Cables – Long	Cables – Short	Backplane / Card-to-Card	Intra-Card	Intra-Module	Intra-chip
Length	Multi-km	10, - 300 m	1 m - 10 m	0.3 m - 1 m	0.1 m - 0.3 m	5 mm - 100 mm	0 mm - 20 mm
Typical # lanes per link	1	1 - 10s	1 - 10s	1 - 100s	1 - 100s	1 - 100s	1 - 100s
Use of optics	Since 80s	Since 90s	Since late 00's	Since 2010-2011	2012-2015	After 2015	Later

LOGICAL Link Types

Distinguished by Function & Link Protocol

	Internet	Local Area Network	Cluster / Data Center	Storage Area Network	Direct Attach Storage	I/O	Mezzanine Bus	SMP Coherency Bus	Memory Bus
	Traffic: IP Stds: Ethernet, ATM, SONET,	Traffic: HTML pages to laptops,.. Stds: 1G Ethernet, WiFi	Traffic: Intra-application, or intra-distributed-application Stds: InfiniBand, 1G Ethernet, 10/40/100Enet	Traffic: Read/Write to disk, shared Std: Fibre Channel	Traffic: Read/Write to disk, unshared Stds: SAS, SATA	Traffic: Load/store to I/O adapters Stds: PCI/PCIe	Traffic: Load/store to Hubs & bridges Stds: Hyper-Transport	Traffic: Load/store coherency ops to other CPUs' caches Stds: Hyper-transport	Traffic: Load/Store to DRAM or Memory Fanout chip Stds: DDR3/2/.
Key Characteristic	Inter-operability with "Everybody"	100-300m over RJ-45 / CAT5 cabling, or wireless	BW & latency to <60 meters	Dominated by FC	Shared tech. between servers & desktops	Shared tech between servers & desktops	Reliability	Reliability, massive BW, reliability	Reliability & cost vs. DRAM
Use of optics	Since 80s	Maybe Never? (Wireless, Building re-wiring, BW demand)	Since 2000s	Since 90s	Not yet	Scattered	Not yet	Coming	Coming later

Link Technology: Single-mode Optics | Mixed multi-mode optics & copper | Copper

Transformations in the Modern Data Center

Forecasted evolution of Ethernet (IEEE)

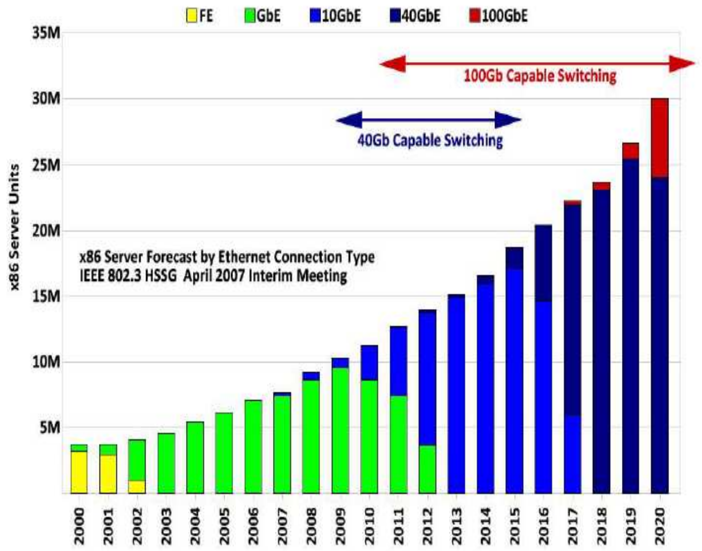

FC and FCoE SAN Switch and Adapter Port Shipments

Source: Dell'Oro Ethernet Switch Report
Five Year Forecast 2011 - 2015

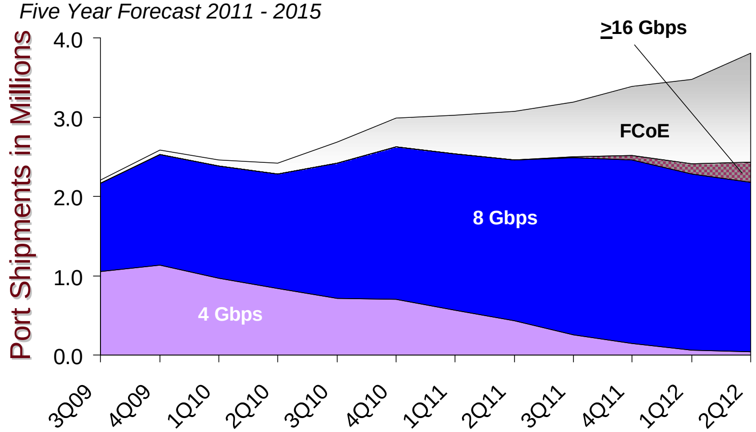

Forecasted evolution of FibreChannel (Infornetics)

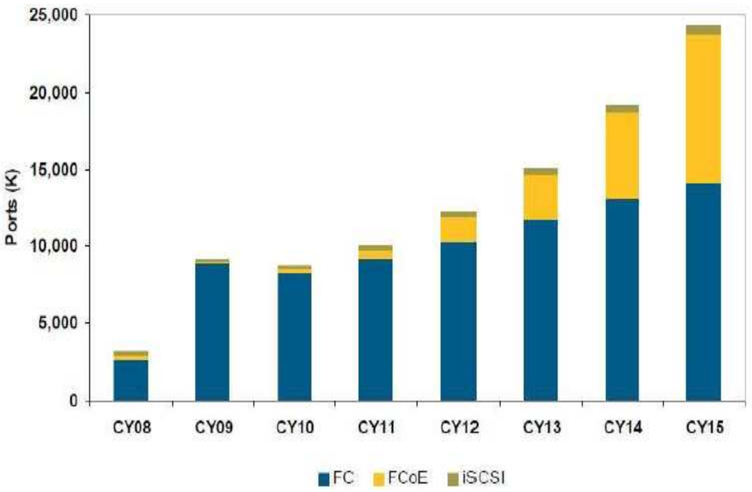

Total FCoE Port Shipments and Revenue

Source: Dell'Oro Ethernet Switch Report
Five Year Forecast 2011 - 2015

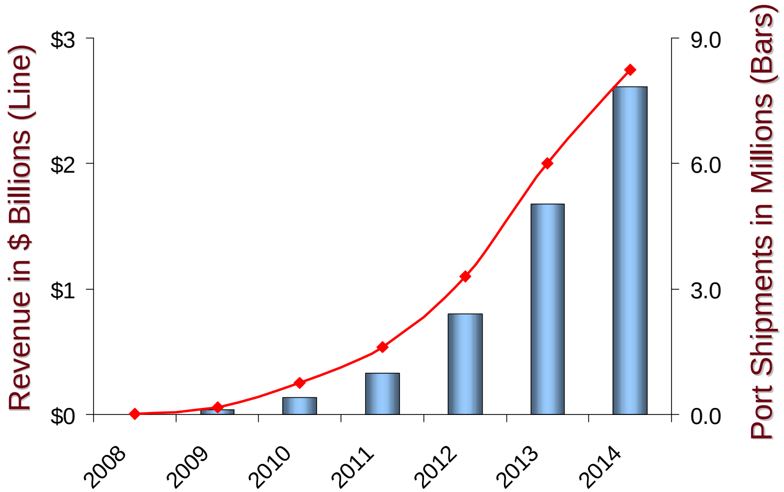

Transformations in the Modern Data Center

Forecasted evolution of Ethernet (IEEE)

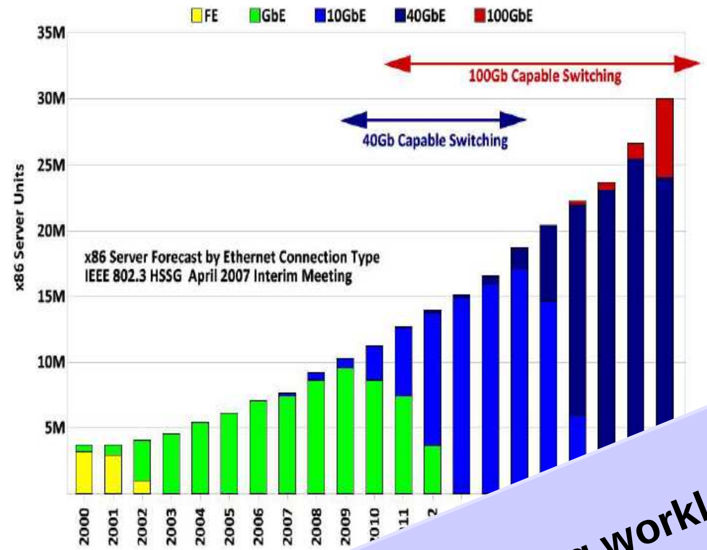

FC and FCoE SAN Switch and Adapter Port Shipments

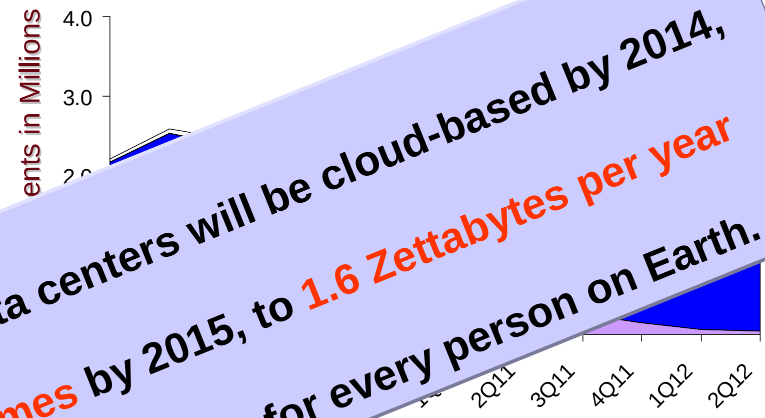

More than 50 percent of computing workloads in data centers will be cloud-based by 2014, and global cloud traffic will grow more than 12 times by 2015, to 1.6 Zettabytes per year — the equivalent of over four days of business-class video for every person on Earth.

Forecasted evo…

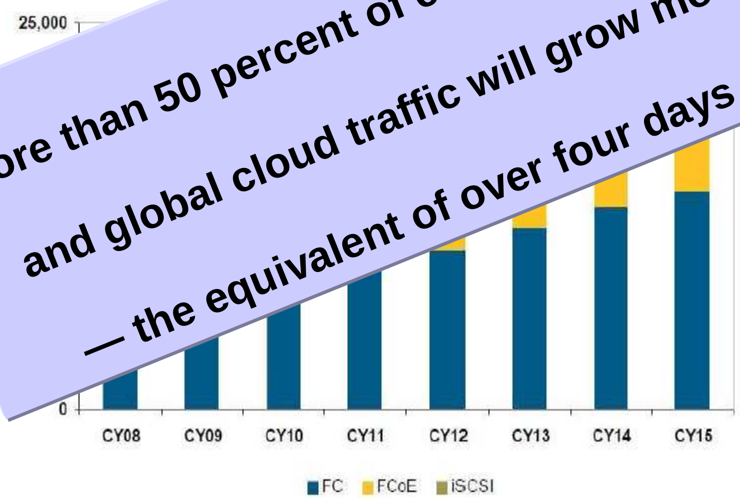

…t Shipments and Revenue

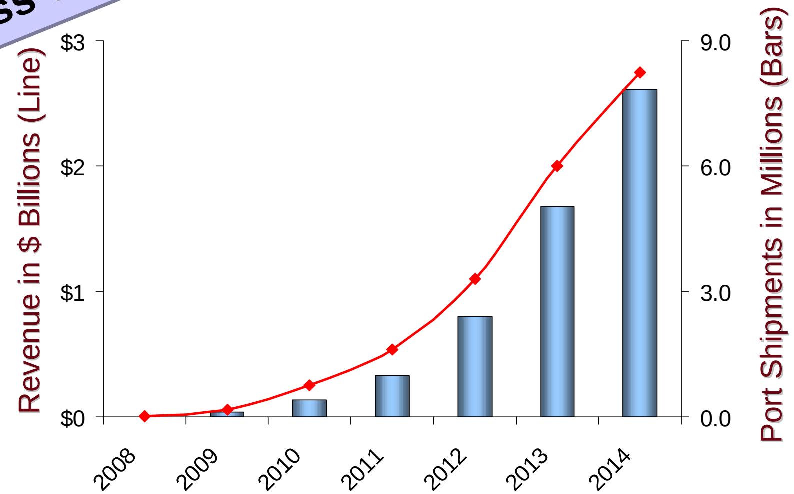

According to the Cisco Cloud Index, January 2012

The Open Datacenter Interoperable Network (ODIN)

- **Standards and best practices for data center networking**
  - Announced May 8 as part of InterOp 2012
    Five technical briefs (8-10 pages each), 2 page white paper, Q&A
    http://www-03.ibm.com/systems/networking/solutions/odin.html
  - Standards-based approach to data center network design, including descriptions of the standards that IBM and our partners agree upon

- **IBM System Networking will publish additional marketing assets describing how our products support the ODIN recommendations**
  - Technical white papers and conference presentations describing how IBM products can be used in these reference architectures
  - See IBM's Data Center Networking blog: https://www-304.ibm.com/connections/blogs/DCN/entry/odin_sets_the_standard_for_open_networking21?lang=en_us
  - And Twitter feed: https://twitter.com/#!/IBMCasimer

Traditional Closed, Mostly Proprietary Data Center Network

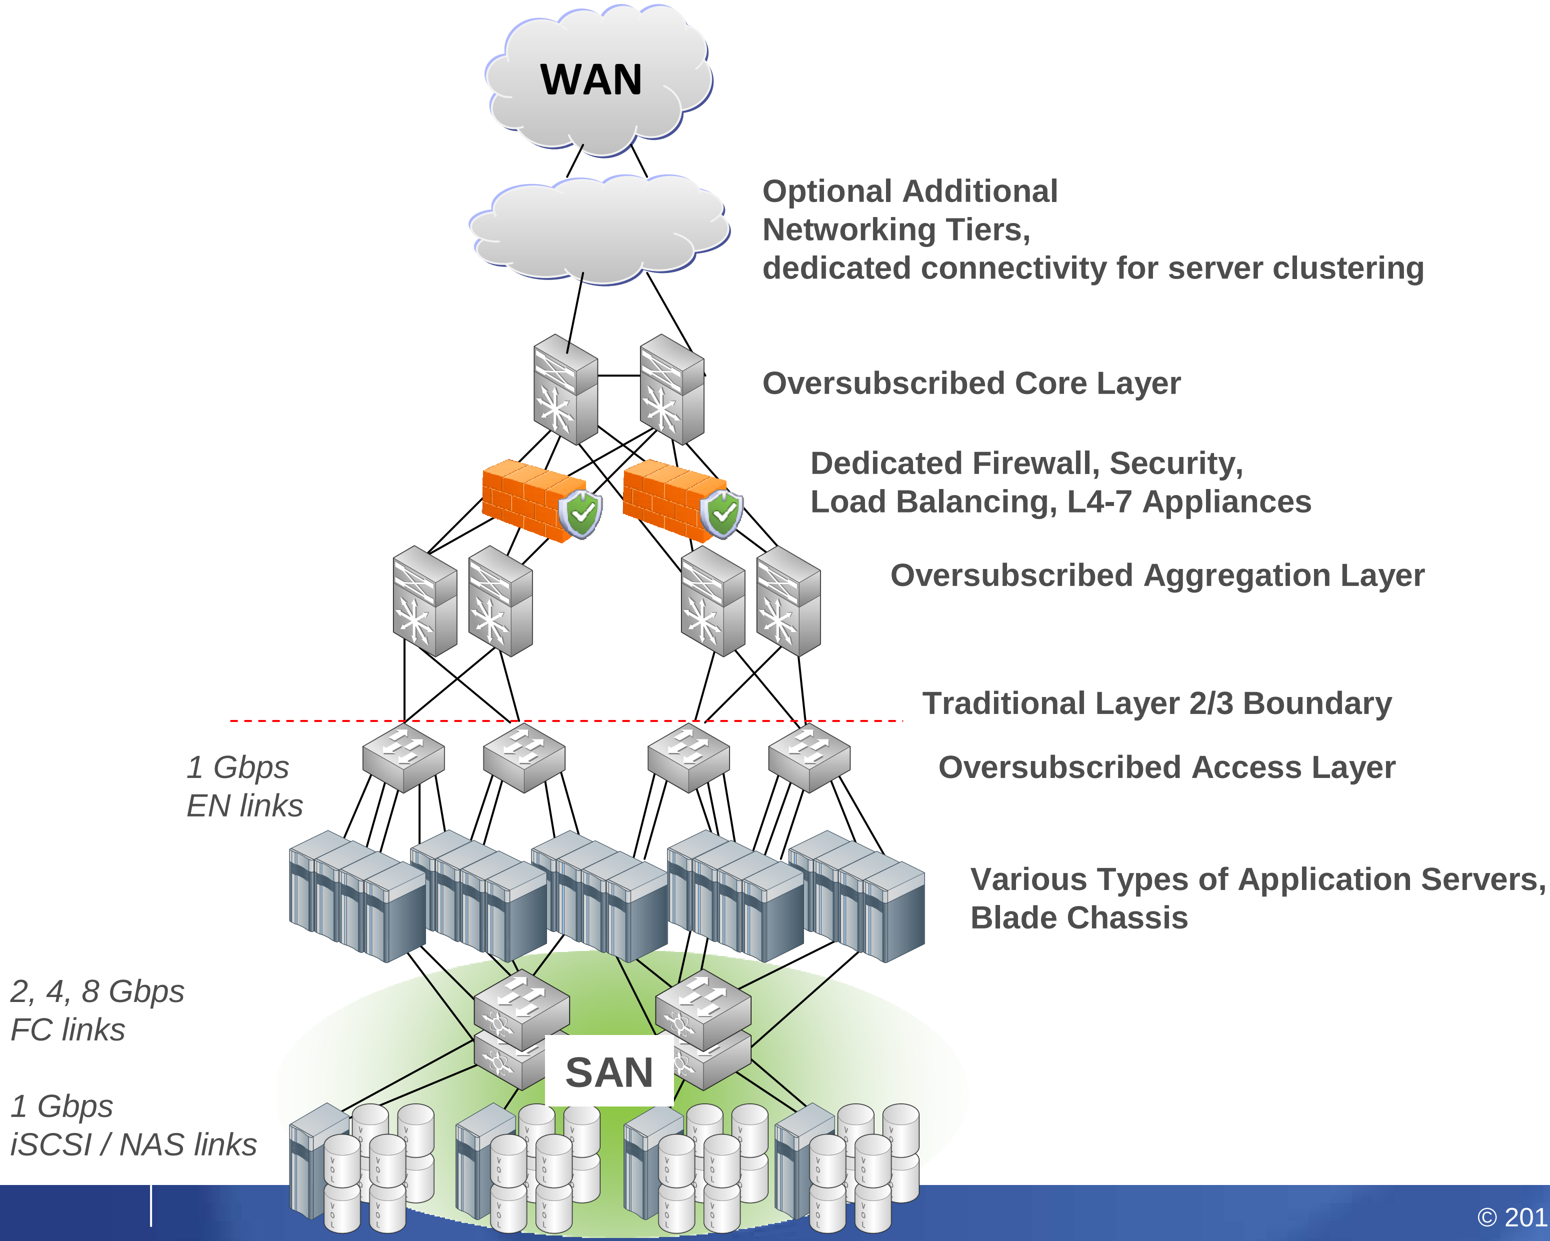

Traditional Data Center Networks: B.O. (Before ODIN)

- **Historically, Ethernet was used to interconnect "stations" (dumb terminals), first through repeaters and hubs, eventually through switched topologies**
  - Not knowing better, we designed our data centers the same way
- **The Ethernet campus network evolved into a structured network characterized by access, aggregation, services, and core layers, which could have 3, 4, or more tiers**
- **These networks are characterized by:**
  - Mostly north-south traffic patterns
  - Oversubscription at all tiers
  - Low virtualization, static network state
  - Use of spanning tree protocol (STP) to prevent loops
  - Layer 2 and 3 functions separated at the access layer
  - Services (firewalls, load balancers, etc) dedicated to each application in a silo structure
  - Network management centered in the switch operating system
  - Complex, often proprietary features and functions

Problems with Traditional Networks

- **Too many tiers**
  - Each tier adds latency (10-20 us or more); cumulative effect degrades performance
  - Oversubscription (in an effort to reduce tiers) can result in lost packets
- **Does not scale in a cost effective or performance effective manner**
  - Scaling requires adding more tiers, more physical switches, and more physical service appliances
  - Management functions do not scale well
  - STP restricts topologies and prevents full utilization of available bandwidth
  - Physical network must be rewired to handle changes in application workload
  - Manually configured SLAs and security prone to errors
  - Potential shortages of IP Addresses
- **Not optimized for new functions**
  - Most modern data center traffic is east-west
  - Oversubscription / lossy network requires separate storage infrastructure
  - Increasing use of virtualization means significantly more servers which can be dynamically created, modified, or destroyed
  - Desire to migrate VMs for high availability and better utilization
  - Multi-tenancy for cloud computing and other applications
- **High Operating and Capital Expense**
  - Too many protocol specific network types
  - Too many network, service, and storage managers
  - Too many discrete components lowers reliability, poorly integrated
  - Too much energy consumption / high cooling costs
  - Sprawl of lightly utilized servers and storage
  - Redundant networks required to insure disjoint multi-pathing for high availability
  - Moving VMs to increase utilization limited by Layer 2 domain boundaries, low bandwidth links, & manual management issues
  - Significant expense just to maintain current network, without deploying new resources

Open Datacenter with an Interoperable Network (ODIN)

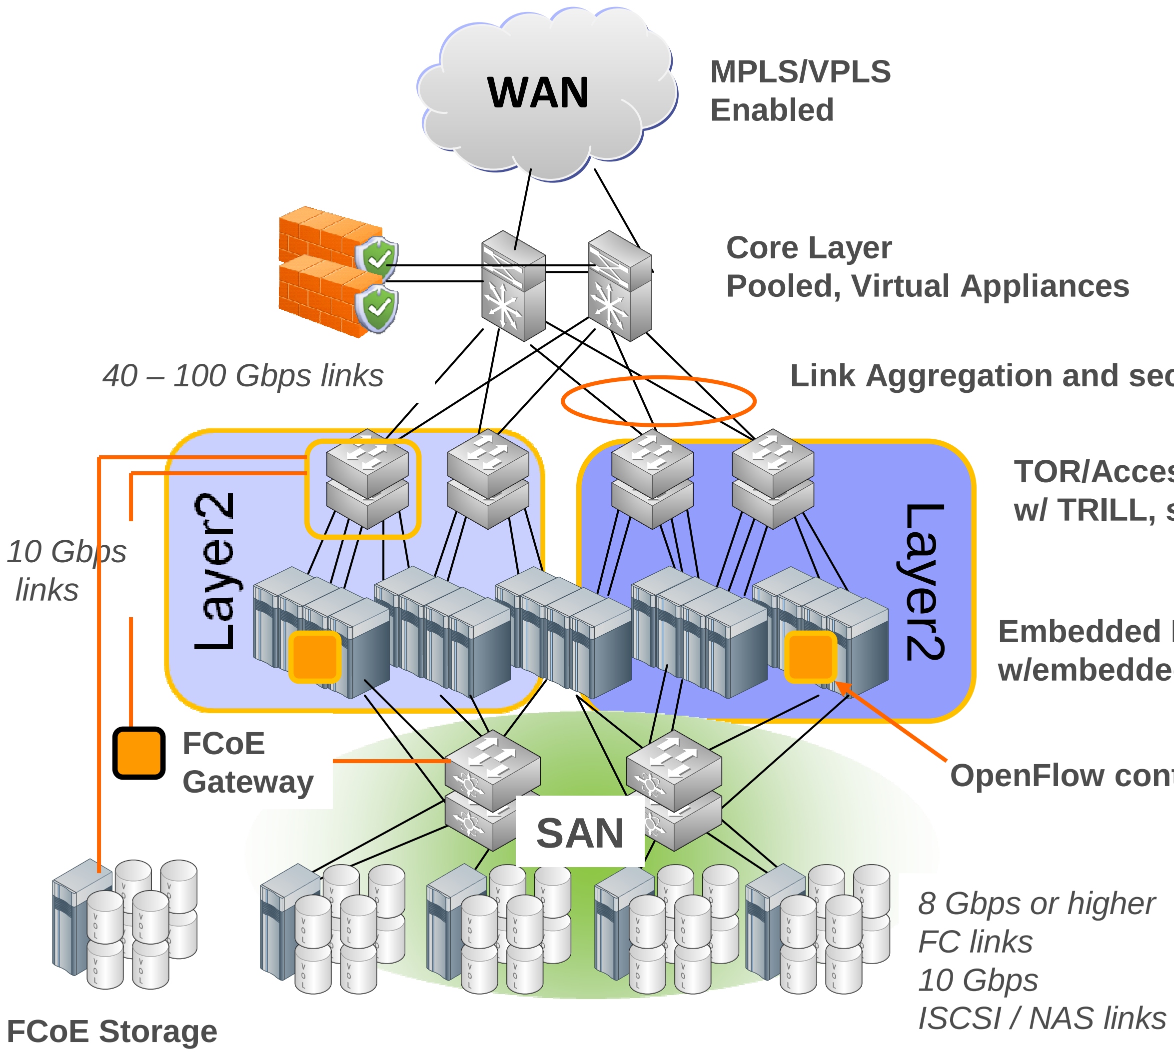

Modern Data Center Networks: A.O. (After ODIN)

- **Modern data centers are characterized by:**
  - 2 tier designs (with embedded Blade switches and virtual switches within the servers)
    - Lower latency and better performance
    - Cost effective scale-out to 1000s of physical ports, 10,000 VMs (with lower TCO)
    - Scaling without massive oversubscription
    - Less moving parts → higher availability and lower energy costs
    - Simplified cabling within and between racks
    - Enabled as an on-ramp for cloud computing, integrated PoDs, and end-to-end solutions
  - Optimized for east-west traffic flow with efficient traffic forwarding
  - Large Layer 2 domains and networks enable VM mobility across different physical servers
    - "VM Aware" fabric; network state resides in vSwitch, automated configuration & migration of port profiles
    - Options to move VMs either through hypervisor Vswitch or external switch
  - "Wire once" topologies with virtual, software-defined overlay networks
  - Pools of service appliances shared across multi-tenant environments
  - Arbitrary topologies (not constrained by STP) with numerous redundant paths, higher bandwidth utilization, switch stacking, and link aggregation
  - Options to converge SAN (and other RDMA networks) into a common fabric with gateways to existing SAN, multi-hop FCoE, disjoint fabric paths, and other features
  - Management functions are centralized, moving into the server, and require fewer instances with less manual intervention and more automation
    - Less opportuity for human error in security and other configurations
  - Based on open industry standards from the IEEE, IETF, ONF, and other groups, which are implemented by multiple vendors (lower TCO per Gartner Group report)
- **ODIN Provides a Data Center Network Reference Design based on Open Standards**

Broad Ecosystem of Support for ODIN

JUNIPER NETWORKS

"In order to contain both capital and operating expense, this network transformation should be based on open industry standards."

BROCADE

"ODIN…facilitates the deployment of new technologies"

HUAWEI

InterOp Webinar: "How to Prepare Your Infrastructure for the Cloud Using Open Standards"

big switch networks

"ODIN is a great example of how we need to maintain openness and interoperability"

NEC Empowered by Innovation

preferred approach to solving Big Data and network bottleneck issues

MARIST COLLEGE ORARE ET LABORARE 1929 POUGHKEEPSIE, N.Y.

National Science Foundation interop lab & Wall St. client engagement

CIENA Simply Smarter Light

"…one of the fundamental "change agents" in the networking industry…associated with encouraging creativity… a nearly ideal approach…is on its way to becoming industry best-practice for transforming data-centers"

ADVA Optical Networking

"…the missing piece in the cloud computing puzzle"

IBM

extreme networks

"We are proud to work with industry leaders like IBM"

How is the data center evolving ?

Integrated

8. **Self-contained** expandable infrastructure
9. Platform Manager & Software Defined Networking Stack

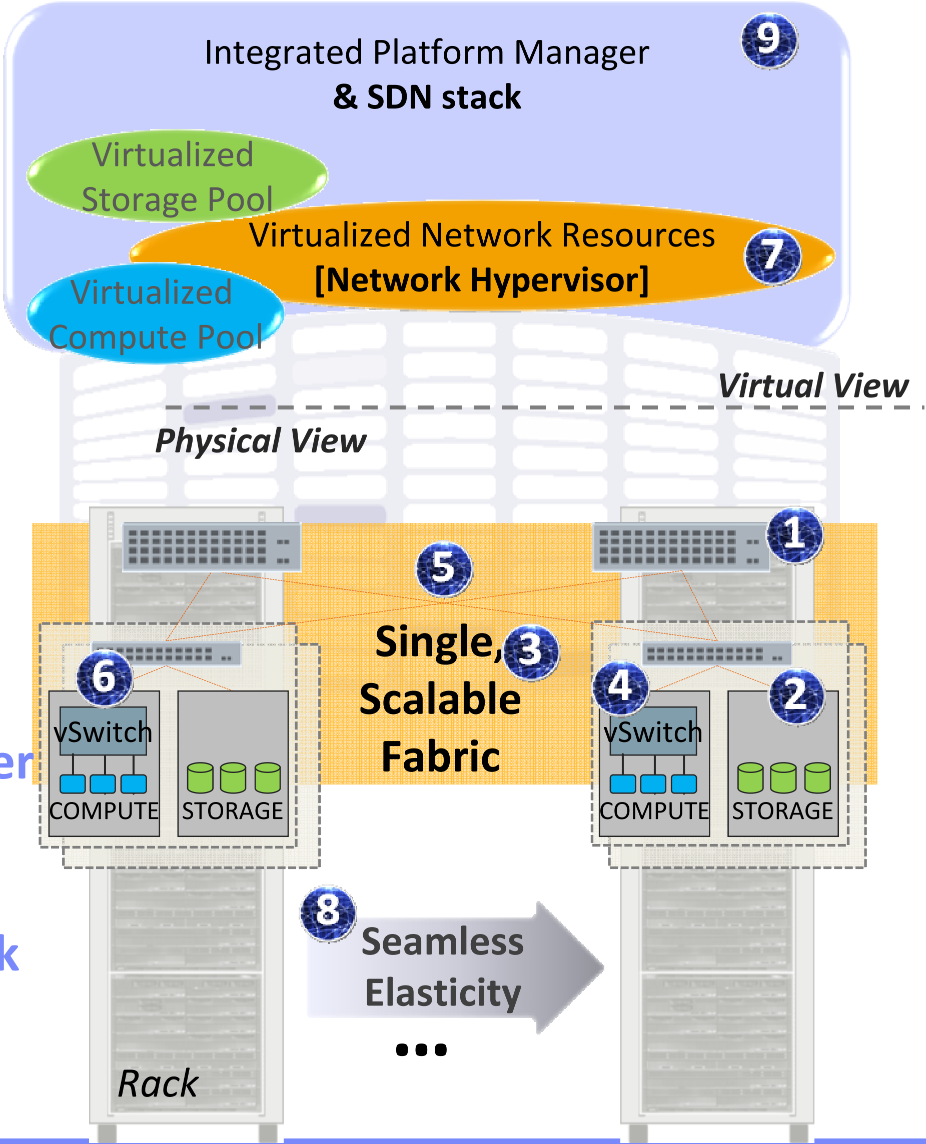

Automated

6. **Virtual Machine** network state automation
7. **Multi-tenant aware** Network Hypervisor

Optimized

1. **Fabric managed as a** single switch
2. Converged **fabric**
3. Scalable **fabric**
4. Flexible Bandwidth
5. Optimized Traffic

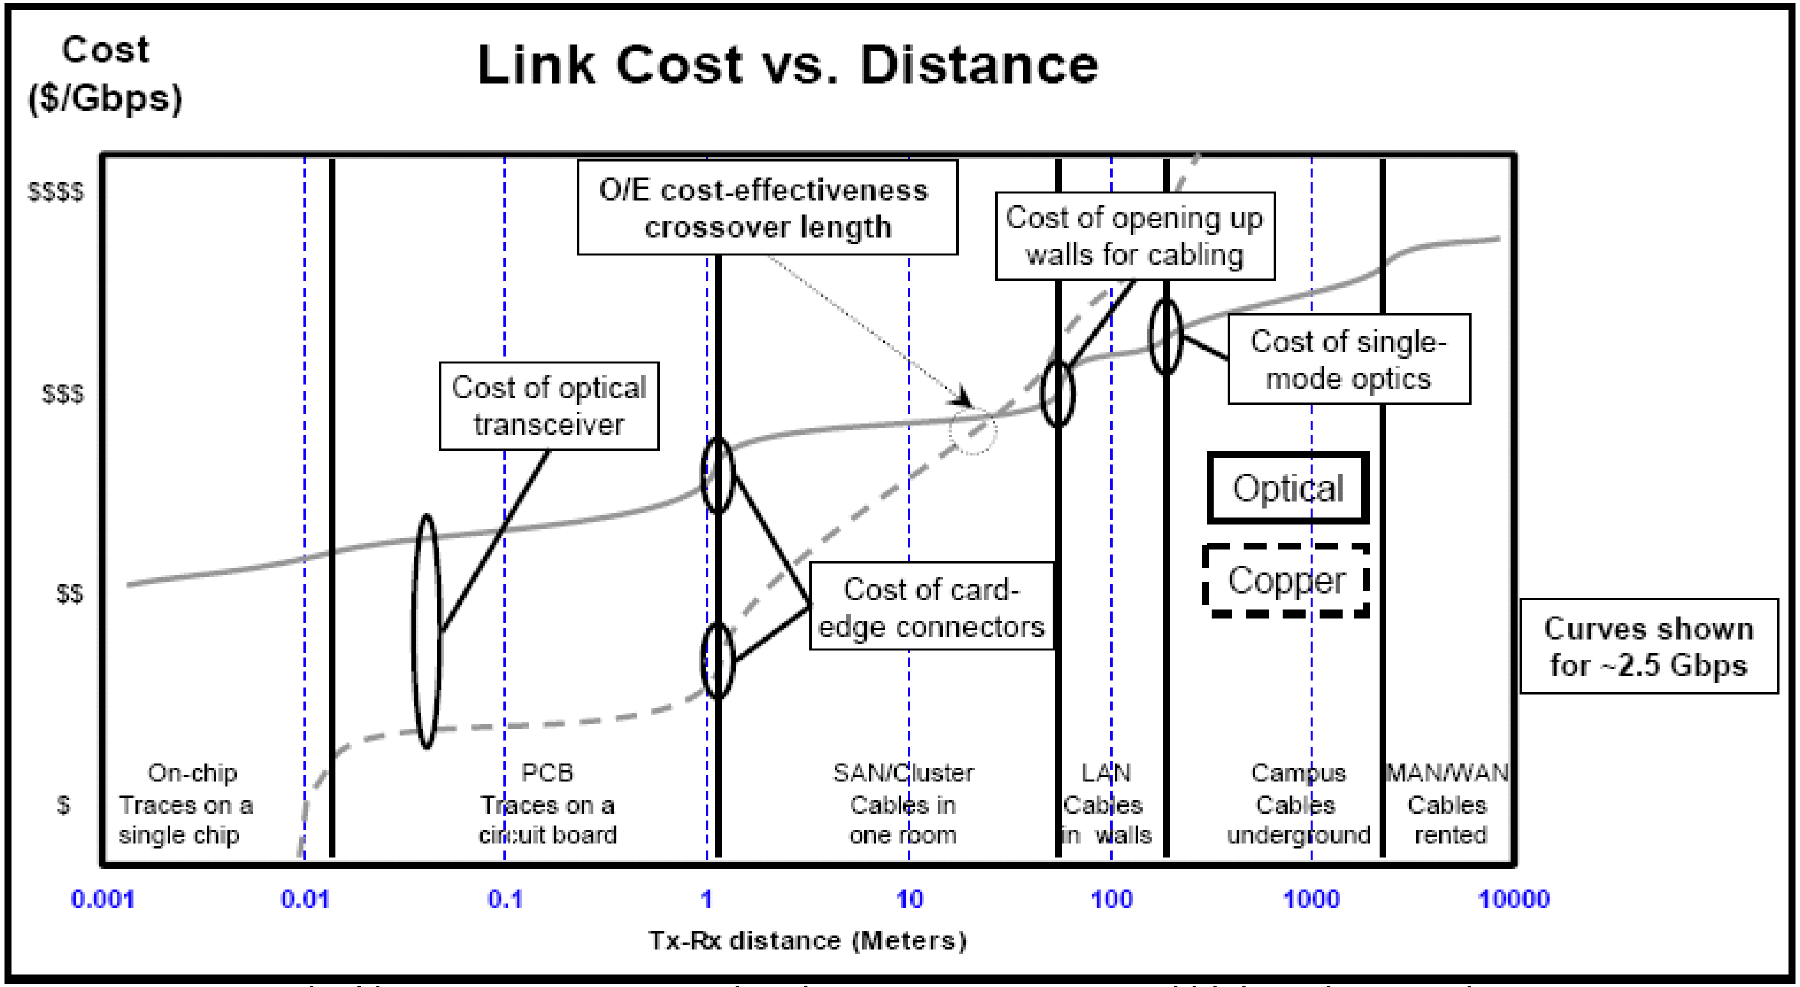

From A. Benner, "Optical interconnect opportunites in supercomputers and high end computing", OFC 2012

Cost curve for copper links is divided into several distinctly shaped regions

Cost curve for optical links is essentially flat with a higher startup cost

At short distances, copper is less expensive; at long distances, optics is less expensive

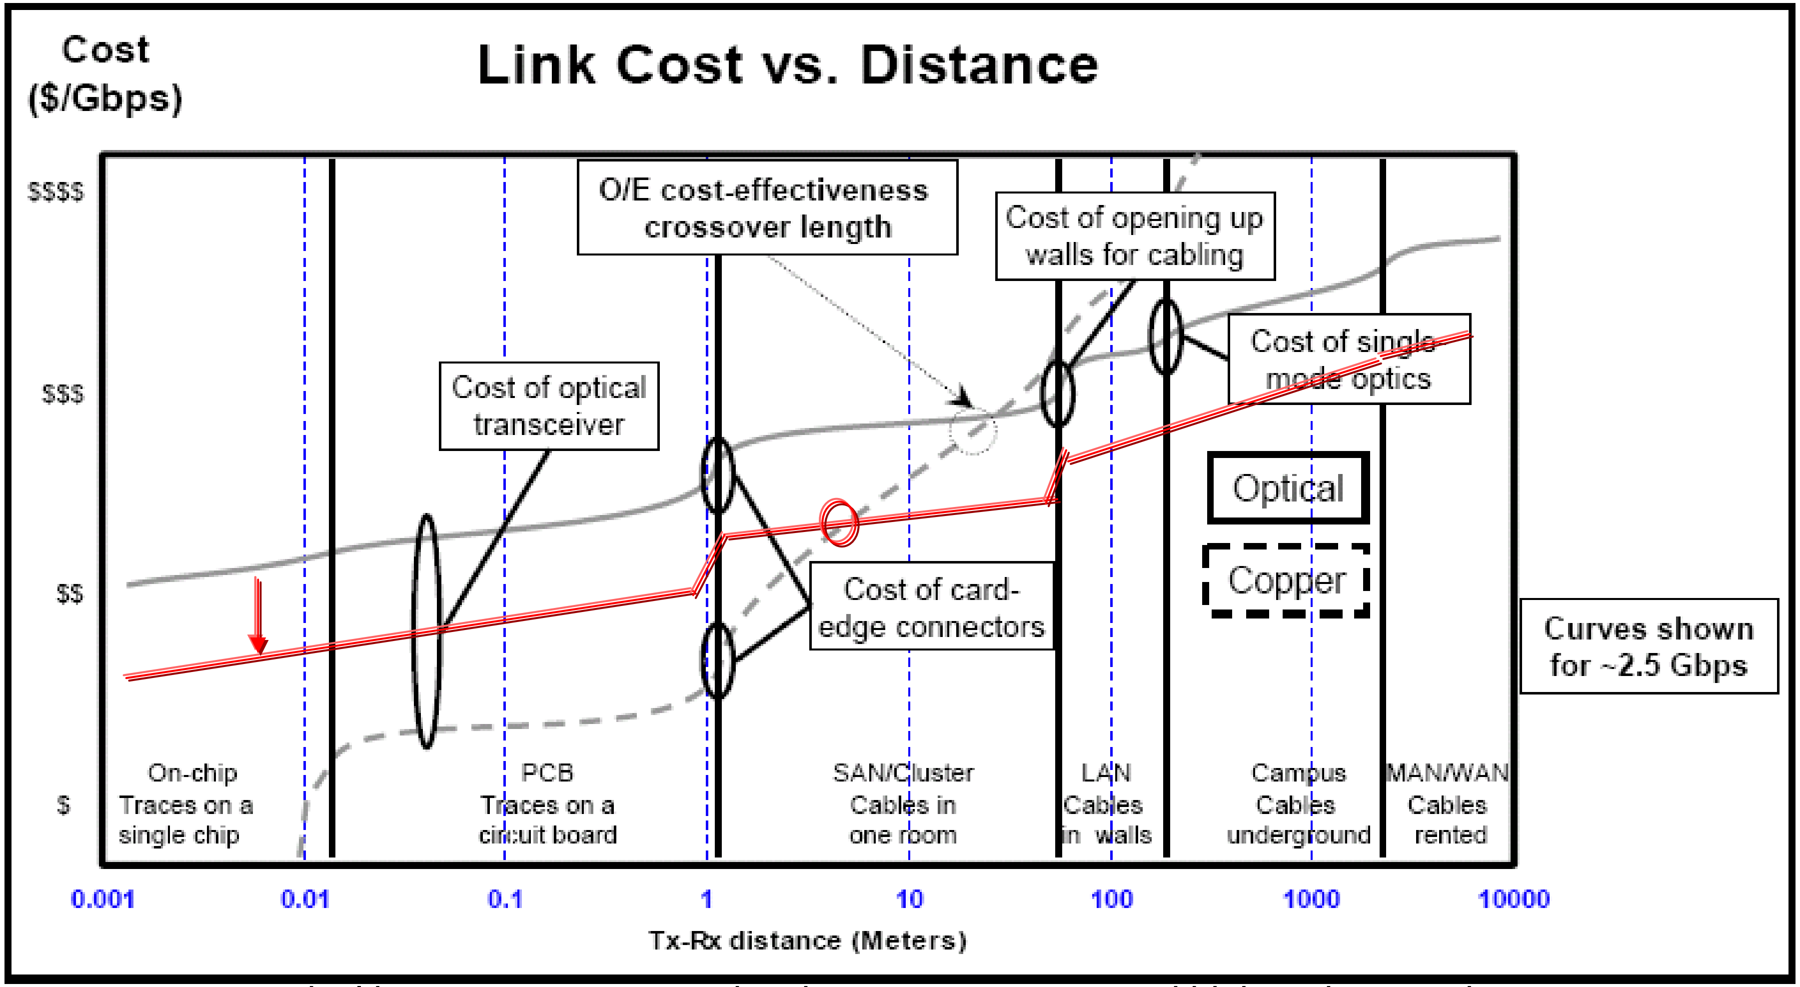

From A. Benner, "Optical interconnect opportunites in supercomputers and high end computing", OFC 2012

10G LOM tends to lower the entry point cost of optics
(and thus moves the crossover length to shorter distances)

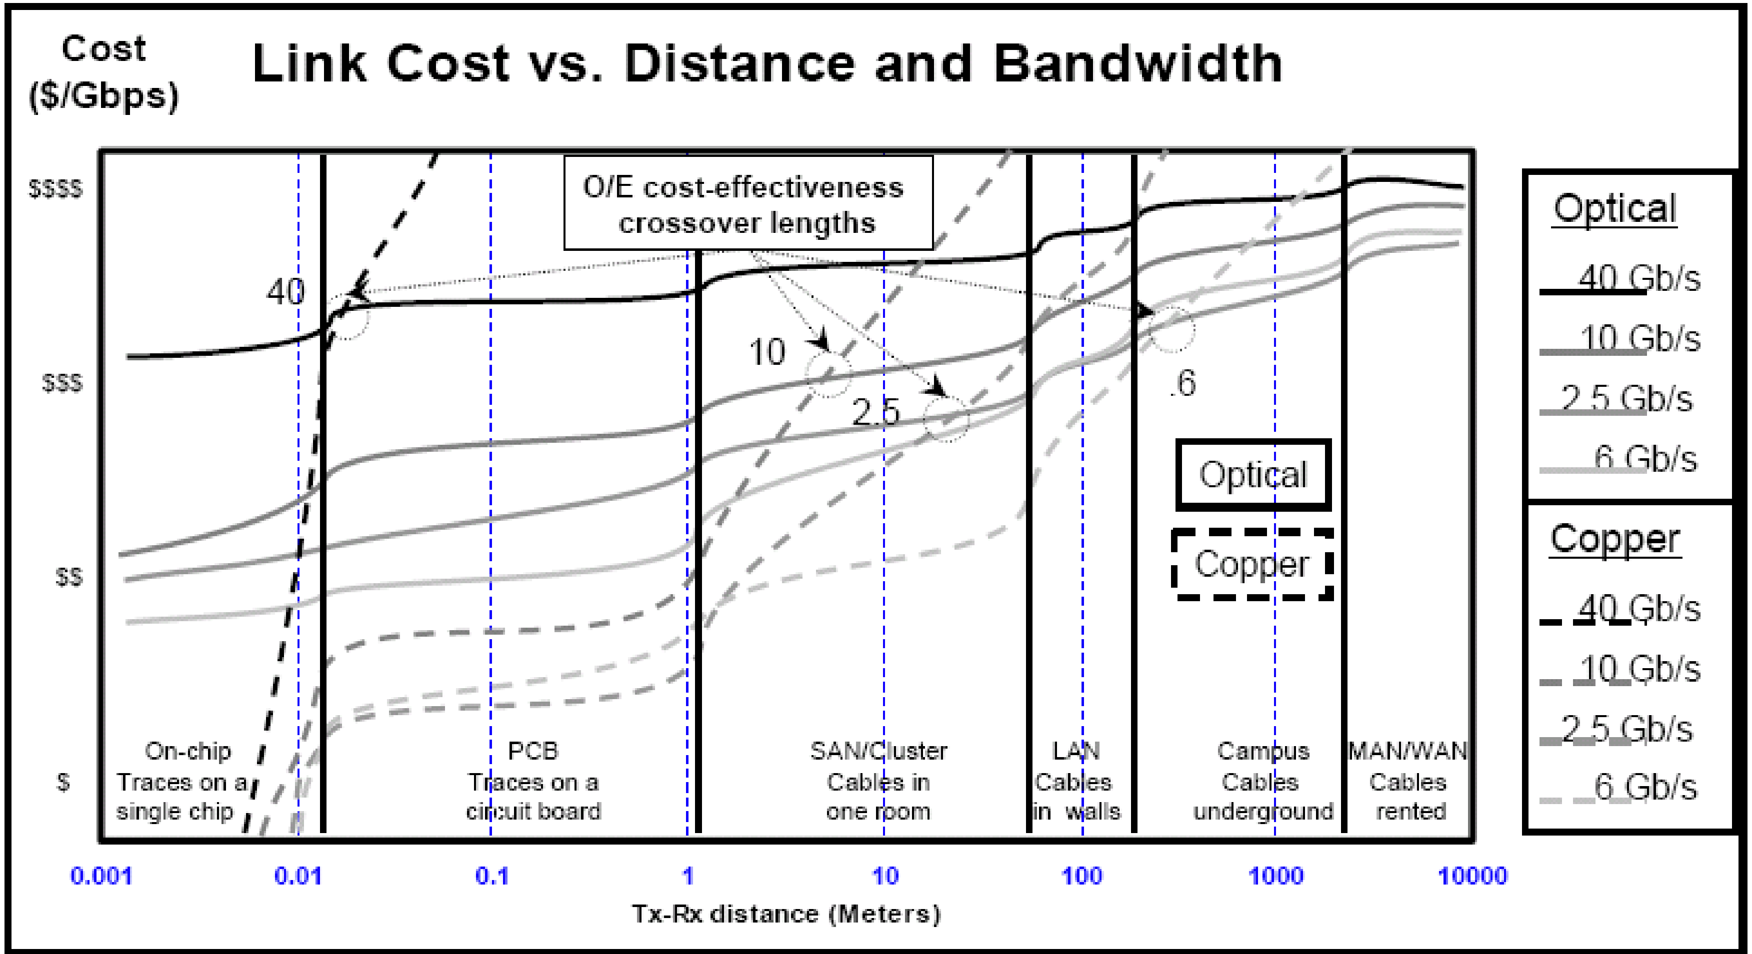

From A. Benner, “Optical interconnect opportunites in supercomputers and high end computing”, OFC 2012

copper & optics costs decline at about the same rate (crossover length at a given bit rate remains constant)

At higher bit rates, crossover point moves to shorter distances

At 2.5 Gbit/s, copper works up to 2-3 meters (within a rack, Blade chassis, or PoD) - 75% of the traffic

A conventional data center may contain thousands of servers,
each capable of hosting tens of virtual machines using current technology;
the number of network endpoints can easily reach tens to hundreds of thousands

Data centers are growing in scale…

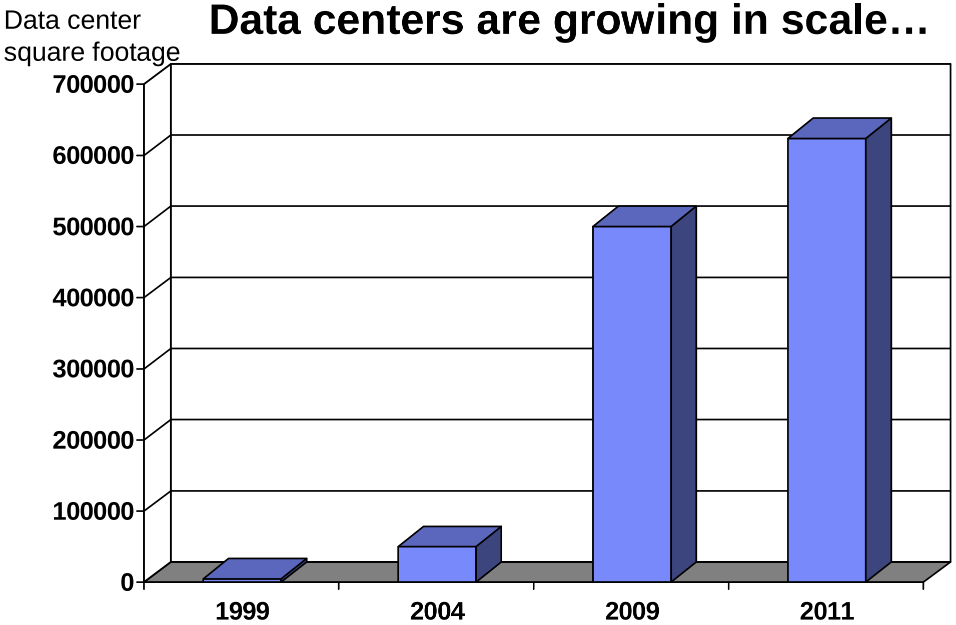

….and Power !

In 2001 there were 1 or 2 data centers in the world which required 10 MW of power

in 2011 there are dozens, and new data centers in the 60-65 MW range are under development

There are unique problems in a highly virtualized environment:

- not just massive numbers of network endpoints
- dynamically create, modify, and destroy VMs at any time
- Multi-tenancy & large number of isolated, independent sub-networks

High performance networks become very important

- High bandwidth links enable flexible placement of workloads & data for higher server utilization; this is a direct benefit to commercial enterprises

IBM Raleigh Leadership Data Center

Reduces energy costs by $1.8 M USD/year (over 50%)

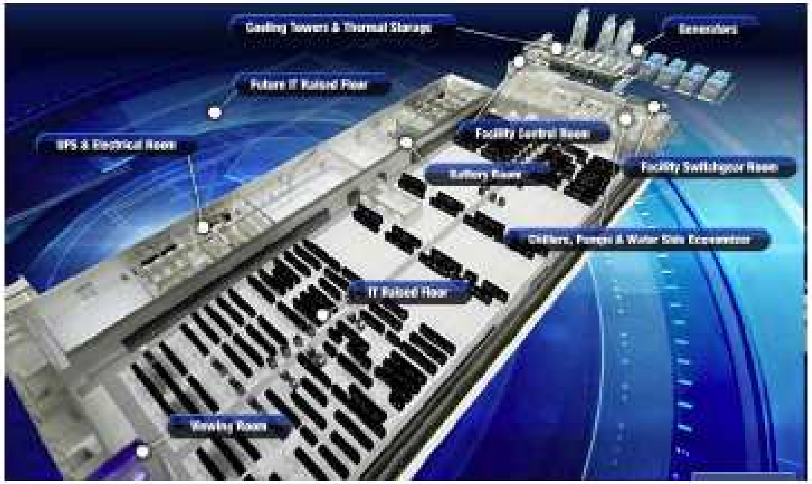


Two-cell 1300 ton cooling tower with variable speed fans


Six 2.5MW diesel generators for emergency power

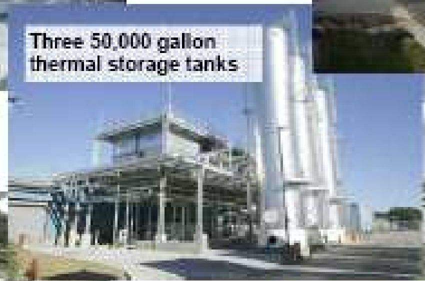
Three 50,000 gallon thermal storage tanks

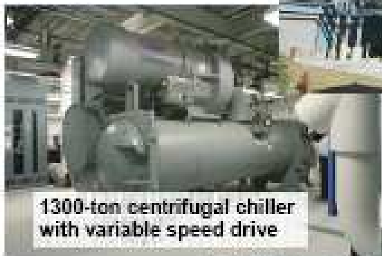
1300-ton centrifugal chiller with variable speed drive


Water side economizer for 3900 hrs/yr of free cooling


60,000 sq. ft. of IT raised floor space


3 MW electtrical switchgear from two independent sources

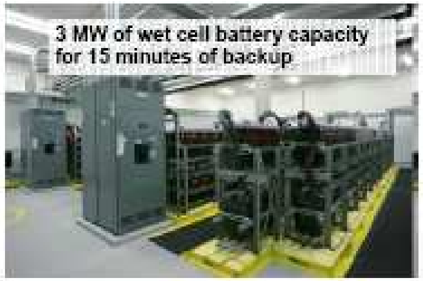
3 MW of wet cell battery capacity for 15 minutes of backup


Piping to support water cooling

Modern data centers require industrial-scale equipment

The benefits of server virtualization & consolidation become clear

http://www-935.ibm.com/services/us/cio/smarterdc/rtp_popup.html

Energy for Data Transport is a major issue for Exascale Systems

- Energy involved in data transport dominates energy used in compute cycles
  - Energy needed per floating point operation: 0.1 pJ/bit
  - Data transport on card (3-10 inches) is 200X higher
  - Data transport across a large system is 2000x higher !

Optics is a more power efficient technology:

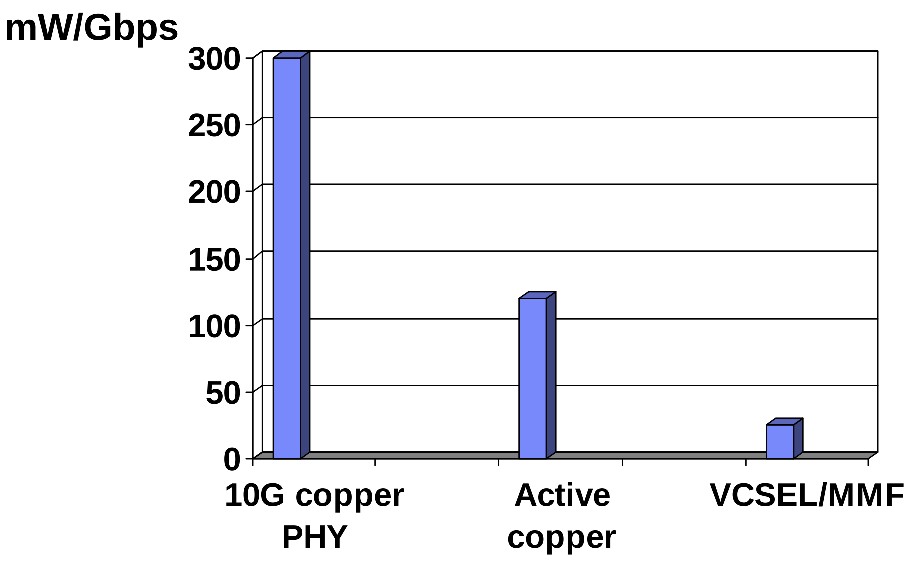

Assuming $1M per MWatt per year, total lifetime cost of ownership for optical solutions is significantly less expensive

OpenFlow and Software-Defined Networking

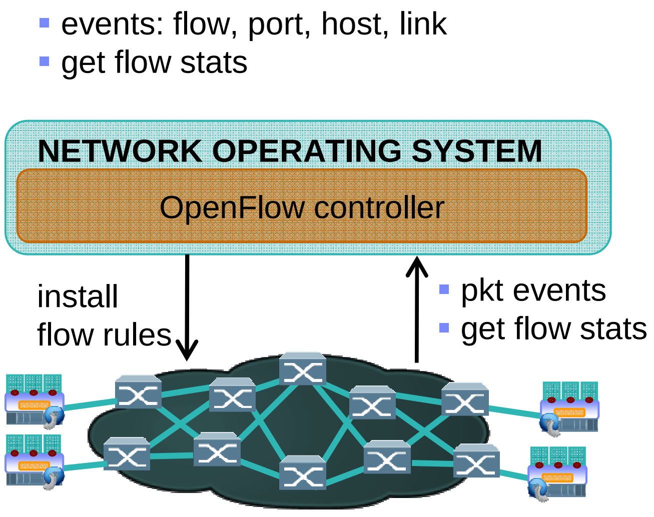

- OpenFlow provides an industry-standard API and protocol to program forwarding tables in switches
  - usually done by the co-resident control processor and software
  - alternatives exist like e/c-DFP and other vendor proprietary control protocols (e.g., Arista EoS)
- OpenFlow provides a "base layer" mechanism for the network operating system
- Much more is needed beyond OpenFlow to complete the NOS
  - state distribution and dissemination
  - other types of configuration beyond ACL state (std. forwarding, VLANs, port profiles, etc.)

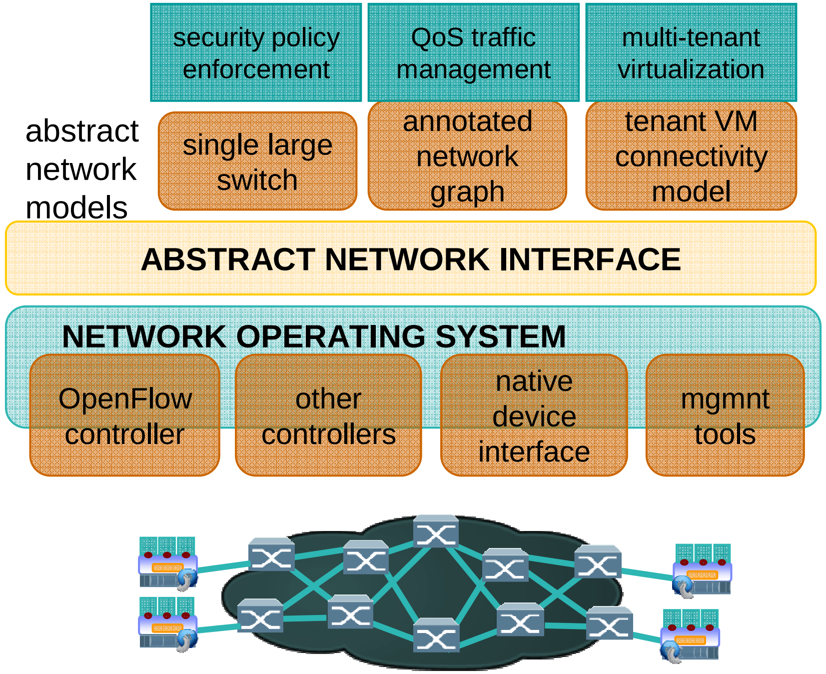

- NOS layer supports more functions
  - create global network view
  - distribute and manage network state
  - configure the physical network, including vSwitches
- ANI: map control and configuration goals on abstract view to the global network view
- models provide simplified view of the network that makes it easy to specify goals via control programs

Overlapping Address Spaces

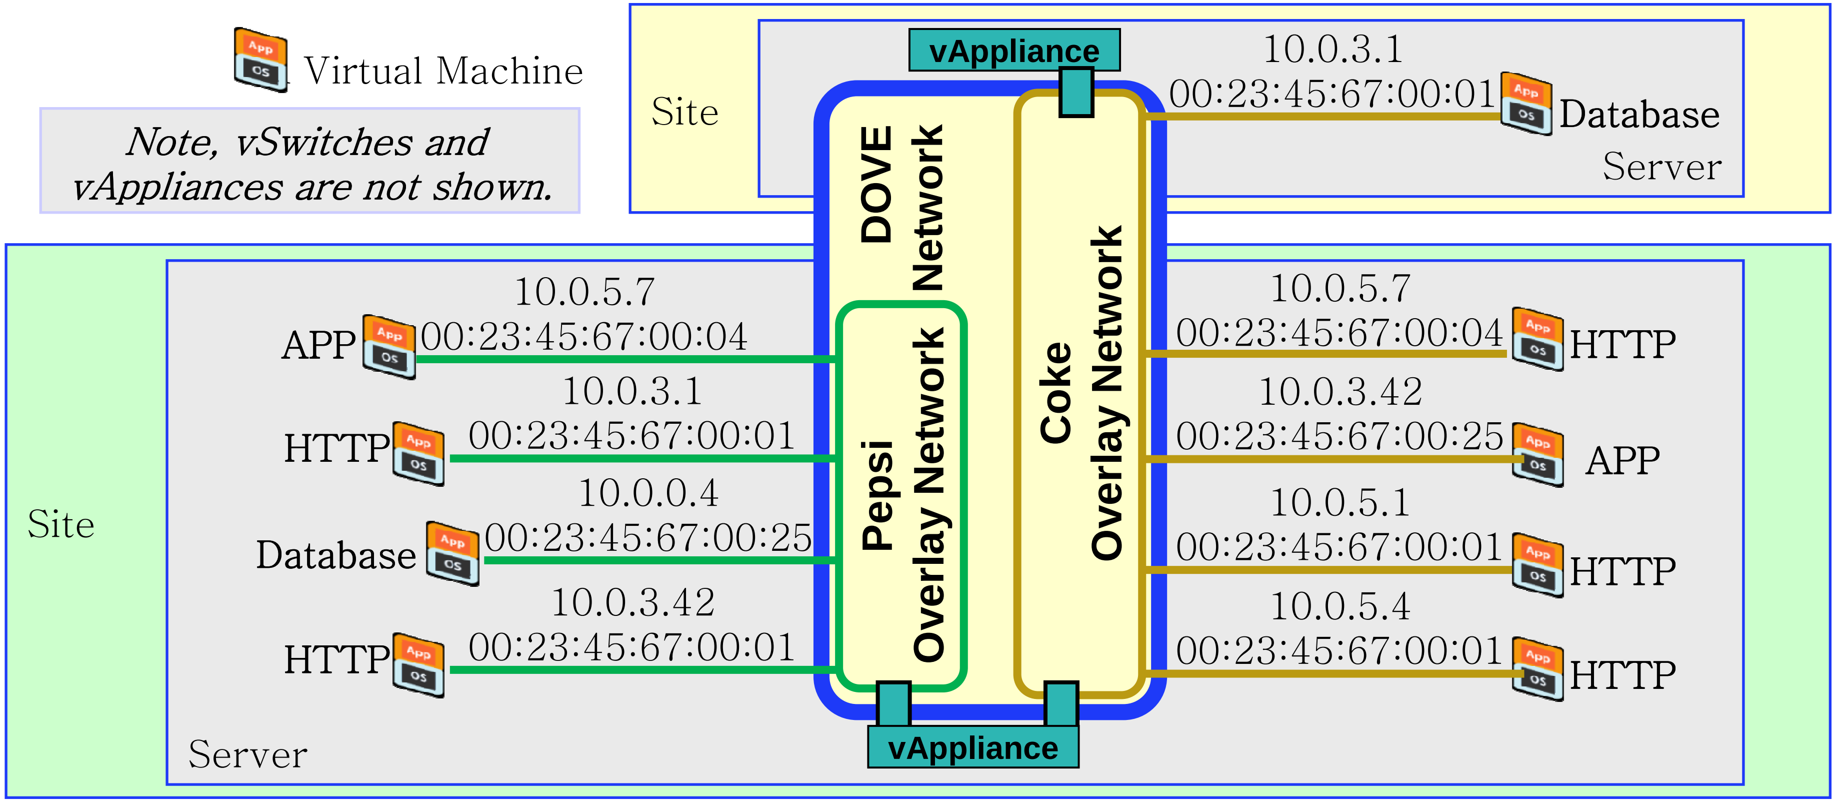

- Multi-tenant, Cloud environments require multiple IP address spaces within same server, within a Data Center and across Data Centers (see above).
  - Layer-3 Distributed Overlay Virtual Ethernet (DOVE) switches enable multi-tenancy all the way into the Server/Hypervisor, with overlapping IP Address spaces for the Virtual Machines.

Data Center Interconnect

HPC

PHYSICAL Link Types

Distinguished by Length & Packaging

PHYSICAL Link Types	MAN & WAN	Cables – Long	Cables – Short	Backplane / Card-to-Card	Intra-Card	Intra-Module	Intra-chip
Length	Multi-km	10, - 300 m	1 m - 10 m	0.3 m - 1 m	0.1 m - 0.3 m	5 mm - 100 mm	0 mm - 20 mm
Typical # lanes per link	1	1 - 10s	1 - 10s	1 - 100s	1 - 100s	1 - 100s	1 - 100s
Use of optics	Since 80s	Since 90s	Since late 00's	Since 2010-2011	2012-2015	After 2015	Later

LOGICAL Link Types

LOGICAL Link Types	Internet	Local Area Network	Cluster / Data Center	Storage Area Network	Direct Attach Storage	I/O	Mezzanine Bus	SMP Coherency Bus	Memory Bus
Distinguished by Function & Link Protocol	Traffic: IP Stds: Ethernet, ATM, SONET,	Traffic: HTML pages to laptops,.. Stds: 1G Ethernet, WiFi	Traffic: Intra-application, or intra-distributed-application Stds: InfiniBand, 1G Ethernet, 10/40/100Enet	Traffic: Read/Write to disk, shared Std: Fibre Channel	Traffic: Read/Write to disk, unshared Stds: SAS, SATA	Traffic: Load/store to I/O adapters Stds: PCI/PCIe	Traffic: Load/store to Hubs & bridges Stds: Hyper-Transport	Traffic: Load/store coherency ops to other CPUs' caches Stds: Hyper-transport	Traffic: Load/Store to DRAM or Memory Fanout chip Stds: DDR3/2/.
Key Characteristic	Inter-operability with "Everybody"	100-300m over RJ-45 / CAT5 cabling, or wireless	BW & latency to <60 meters	Dominated by FC	Shared tech. between servers & desktops	Shared tech between servers & desktops	Reliability	Reliability, massive BW, reliability	Reliability & cost vs. DRAM
Use of optics	Since 80s	Maybe Never? (Wireless, Building re-wiring, BW demand)	Since 2000s	Since 90s	Not yet	Scattered	Not yet	Coming	Coming later

Link Technology | Single-mode Optics | Mixed multi-mode optics & copper | Copper

Exascale Computing Grand Challenge

1 PetaFlop
72 BG/P Racks
2.5 MW
$150 M

10 PetaFlop
100 BG/P Racks
5 MW
$225 M

Overall Performance = 1000X
Performance / watt = 135X
Performance / $ = 1000X
Footprint = <2%

Innovation in technology, architecture, software and algorithms

ExaFlop System:
1 PetaFlop = 1/3 rack
20 MW
$500 M

1 PetaFlop = 1/3 rack
2019

Historical trend (top500.org and green500.org, also see A. Gana & J. Kash, OFC 2011)
10X performance, 4 years later, costs 1.5X more dollars, consumes 2X more power

Future System Interface Requirements

	2012	2018 - 2020
Peak Performance (PFLop)	10	1000 (1 ExaFlop)
Memory Bandwidth (TB/s)	0.05	3 – 15 (local 3D packaging)
Bidirectional IO Bandwidth (TB/s)	0.02	5 (potential optical solution)
Number of IO Pins (at 8 Gbps)	80	20,000 (potential optical solution)
Number of 10G Optical Channels	2×10^6	16×10^8 (higher bit rate/channel may reduce this somewhat)
Optical power consumption	25 mW/Gbps	1 mW/Gbps
Optical cost per Tbps	$1000	$25

Over time, increasing compute performance & bandwidth requirements will cause optical boundary to move closer to processor

IBM Research has active programs in a variety of optical interconnect technologies:

Silicon Photonics
Optical PCBs
Advanced Packaging, Transceivers, and Optical Vias

Data Center Interconnect | HPC | Research

PHYSICAL Link Types	MAN & WAN	Cables – Long	Cables – Short	Backplane / Card-to-Card	Intra-Card	Intra-Module	Intra-chip
Distinguished by Length & Packaging							
Length	Multi-km	10, - 300 m	1 m - 10 m	0.3 m - 1 m	0.1 m - 0.3 m	5 mm - 100 mm	0 mm - 20 mm
Typical # lanes per link	1	1 - 10s	1 - 10s	1 - 100s	1 - 100s	1 - 100s	1 - 100s
Use of optics	Since 80s	Since 90s	Since late 00's	Since 2010-2011	2012-2015	After 2015	Later

LOGICAL Link Types	Internet	Local Area Network	Cluster / Data Center	Storage Area Network	Direct Attach Storage	I/O	Mezzanine Bus	SMP Coherency Bus	Memory Bus
Distinguished by Function & Link Protocol	Traffic: IP Stds: Ethernet, ATM, SONET,	Traffic: HTML pages to laptops,.. Stds: 1G Ethernet, WiFi	Traffic: Intra-application, or intra-distributed-application Stds: InfiniBand, 1G Ethernet, 10/40/100Enet	Traffic: Read/Write to disk, shared Std: Fibre Channel	Traffic: Read/Write to disk, unshared Stds: SAS, SATA	Traffic: Load/store to I/O adapters Stds: PCI/PCIe	Traffic: Load/store to Hubs & bridges Stds: Hyper-Transport	Traffic: Load/store coherency ops to other CPUs' caches Stds: Hyper-transport	Traffic: Load/Store to DRAM or Memory Fanout chip Stds: DDR3/2/.
Key Characteristic	Inter-operability with "Everybody"	100-300m over RJ-45 / CAT5 cabling, or wireless	BW & latency to <60 meters	Dominated by FC	Shared tech. between servers & desktops	Shared tech between servers & desktops	Reliability	Reliability, massive BW, reliability	Reliability & cost vs. DRAM
Use of optics	Since 80s	Maybe Never? (Wireless, Building re-wiring, BW demand)	Since 2000s	Since 90s	Not yet	Scattered	Not yet	Coming	Coming later

Link Technology: Single-mode Optics | Mixed multi-mode optics & copper | Copper

Silicon Photonics-Related: Coupling to on-chip waveguides

- **Edge-coupling of optical waveguides in silicon photonics chip matches well with standard IC packaging practice & power/cooling requirements.**
- **Key problem: low-loss coupling to standard optical fiber**

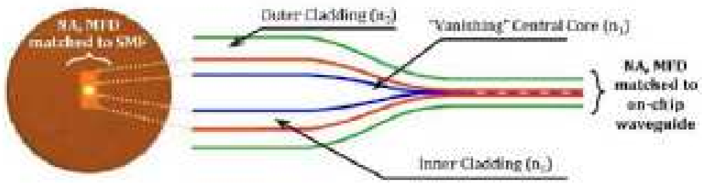

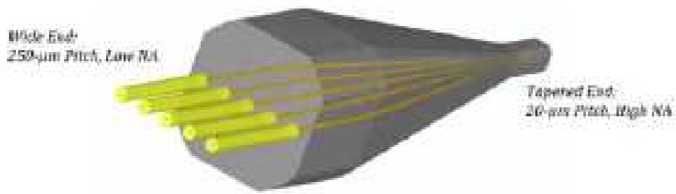

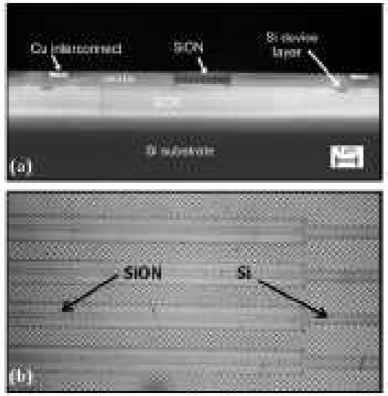

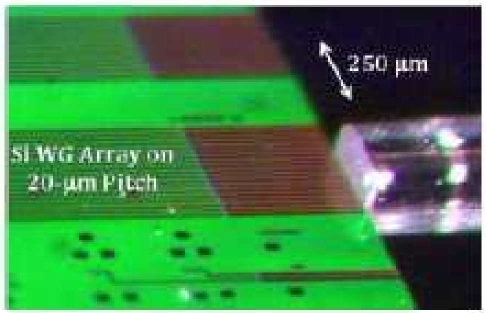

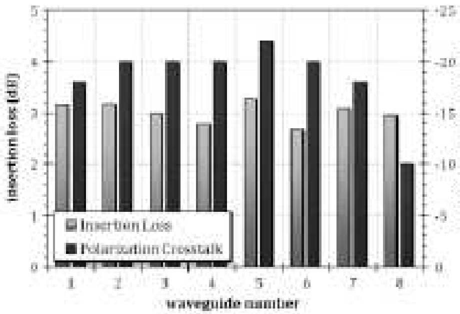

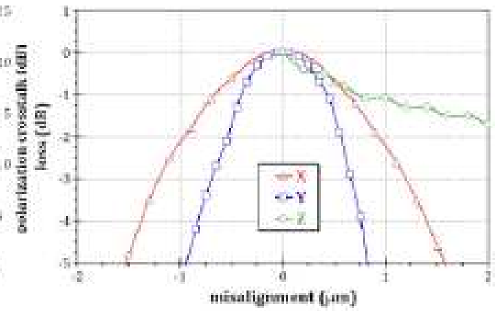

Y. Vlasov et al, multiple papers can be downloaded at: http://www.ibm.research.com/photonics

- F. E. Doany et al., "Multichannel High-Bandwidth Coupling of Ultradense Silicon Photonic Waveguide Array to Standard-Pitch Fiber Array", JLT, Vol. 29, No. 4, Feb.2011

P7-IH – Cable Density

- Many many optical fibers
  - Each of these cables is a 24-fiber multimode cable, carrying (10+10) GBytes/sec of traffic

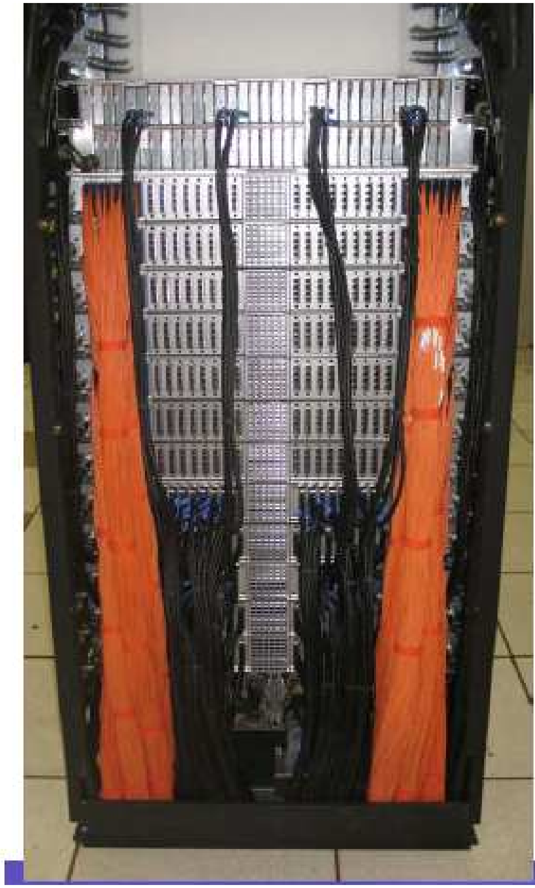

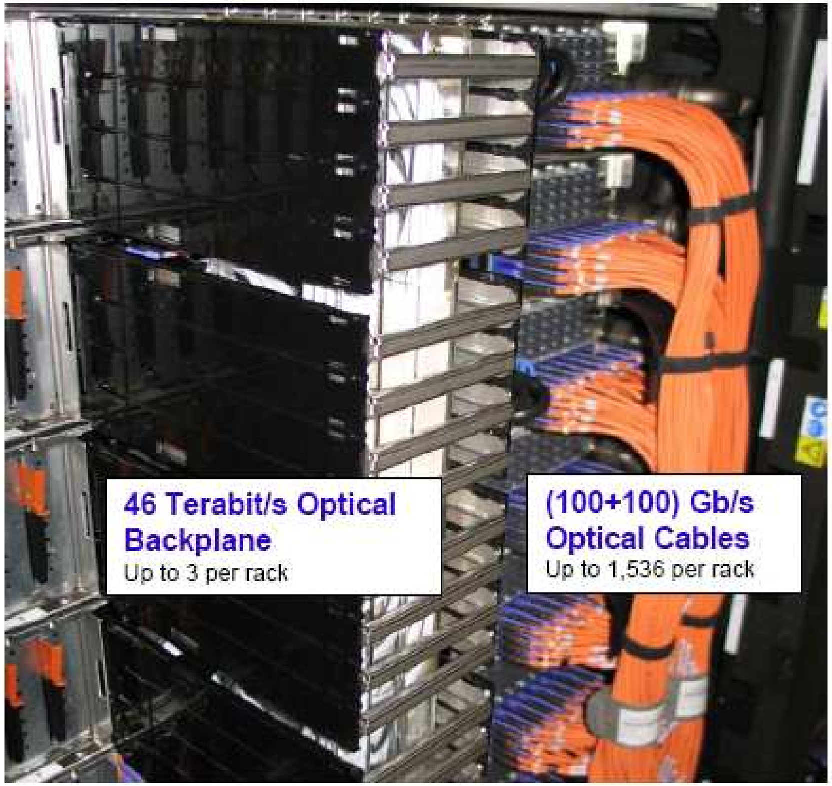

I/O Bandwidth for various levels of packaging is critical

Optical printed circuit boards

In collaboration with Varioprint

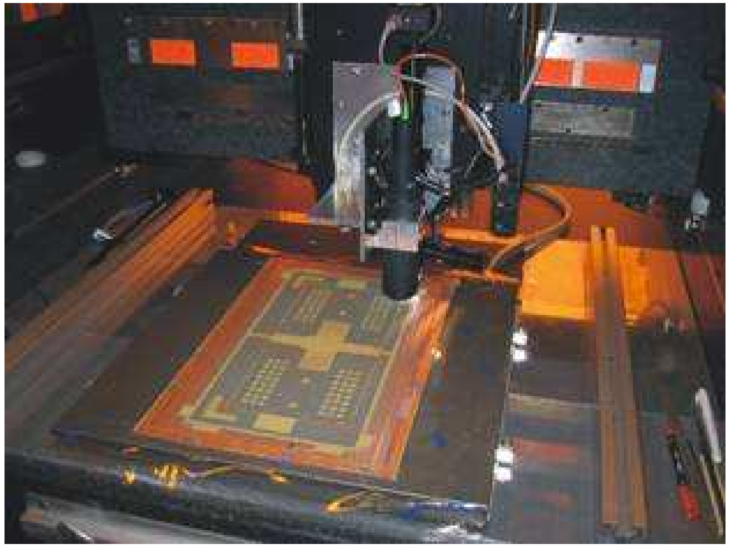

Waveguide processing on large panels,
305 mm x 460 mm

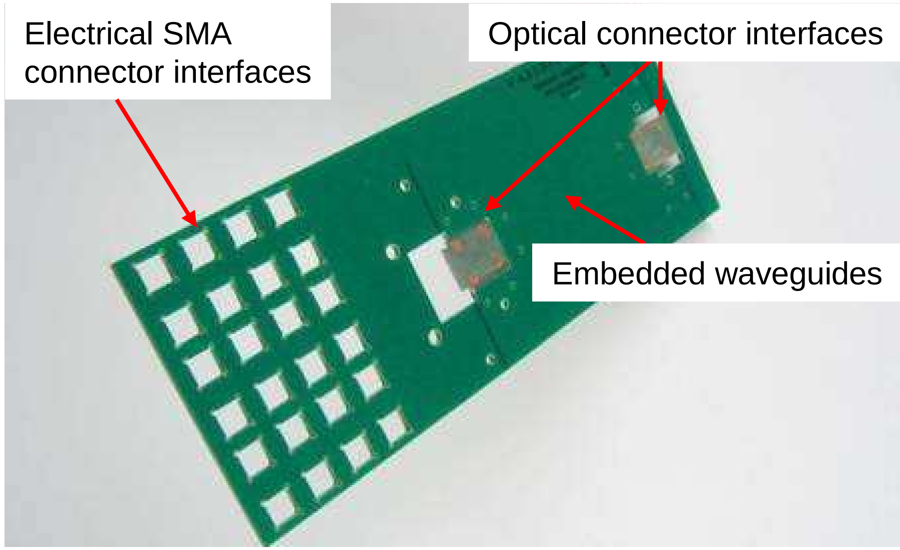

Finished optical board with optical and mechanical interfaces

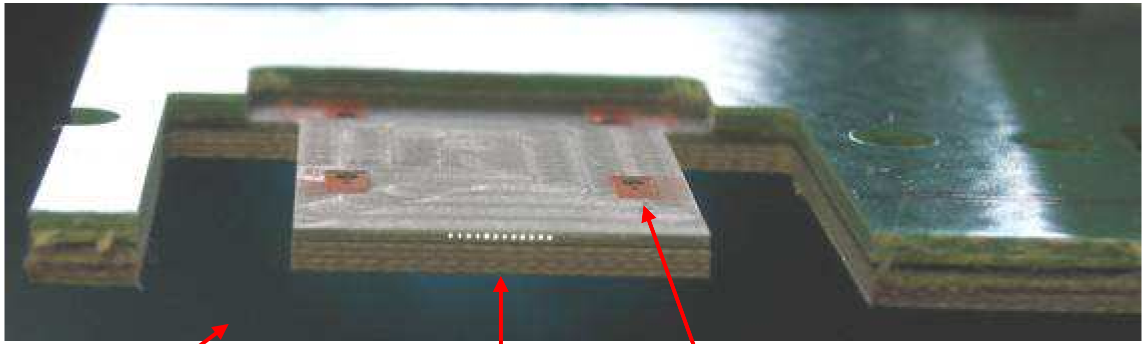

F. E. Doany et al, ECTC 2010, pp. 58–65
D. Jubin et al, Proc. SPIE, Vol. 7607, 76070K (2010)
F. D. Doany et al, IEEE Trans. Adv. Packag., Vol. 32, No. 2, pp. 345-359, May 2009.

Conclusions

- Accelerating change in enterprise data center networks
  - Rising energy costs, under-utilized servers, limited scalability, dynamic workload management
  - Need to automate, integrate, and optimize data center networks

- Over time, the boundary for optical interconnect is moving closer to the server
  - Higher data rates accelerate this trend
  - Virtualization, integrated blade switches, and SDN increase the need for flexible, high bandwidth links

- HPC performance is growing at an exponential rate, driving to exascale computing by 2018 – 2020 timeframe
  - Increasing aggregate system performance will demand more optical links in the future; bandwidth is steadily increasing (higher channel rates, parallel channels)
  - State of the art electronic packaging are engineering marvels but can be complex, difficult to test, and expensive.
  - Integration of high bandwidth technologies into first and second level packaging will be key; density requirements for connectors become increasingly important as the number of links per system grows
  - Single card systems may have integrated on-board optics within 5 years

- There are many enabling technologies on the horizon
  - Including silicon photonics, high speed VCSELs and transceivers, optical vias and PCBs, and polymer waveguides (to name just a few)

Thank You !