Optical Networking for Data Center Interconnects Across Wide Area Networks

Joe Berthold

Hot Interconnects 2009

August 27, 2009


MONTY PYTHON'S
And Now For Something
COMPLETELY DIFFERENT
DVD

ciena

Overview

Data center interconnect applications

Optical transmission and switching –current status

Optical networking economics in the WAN

100G – the next step in optical networking

Data Centers Across the Globe

an extreme example

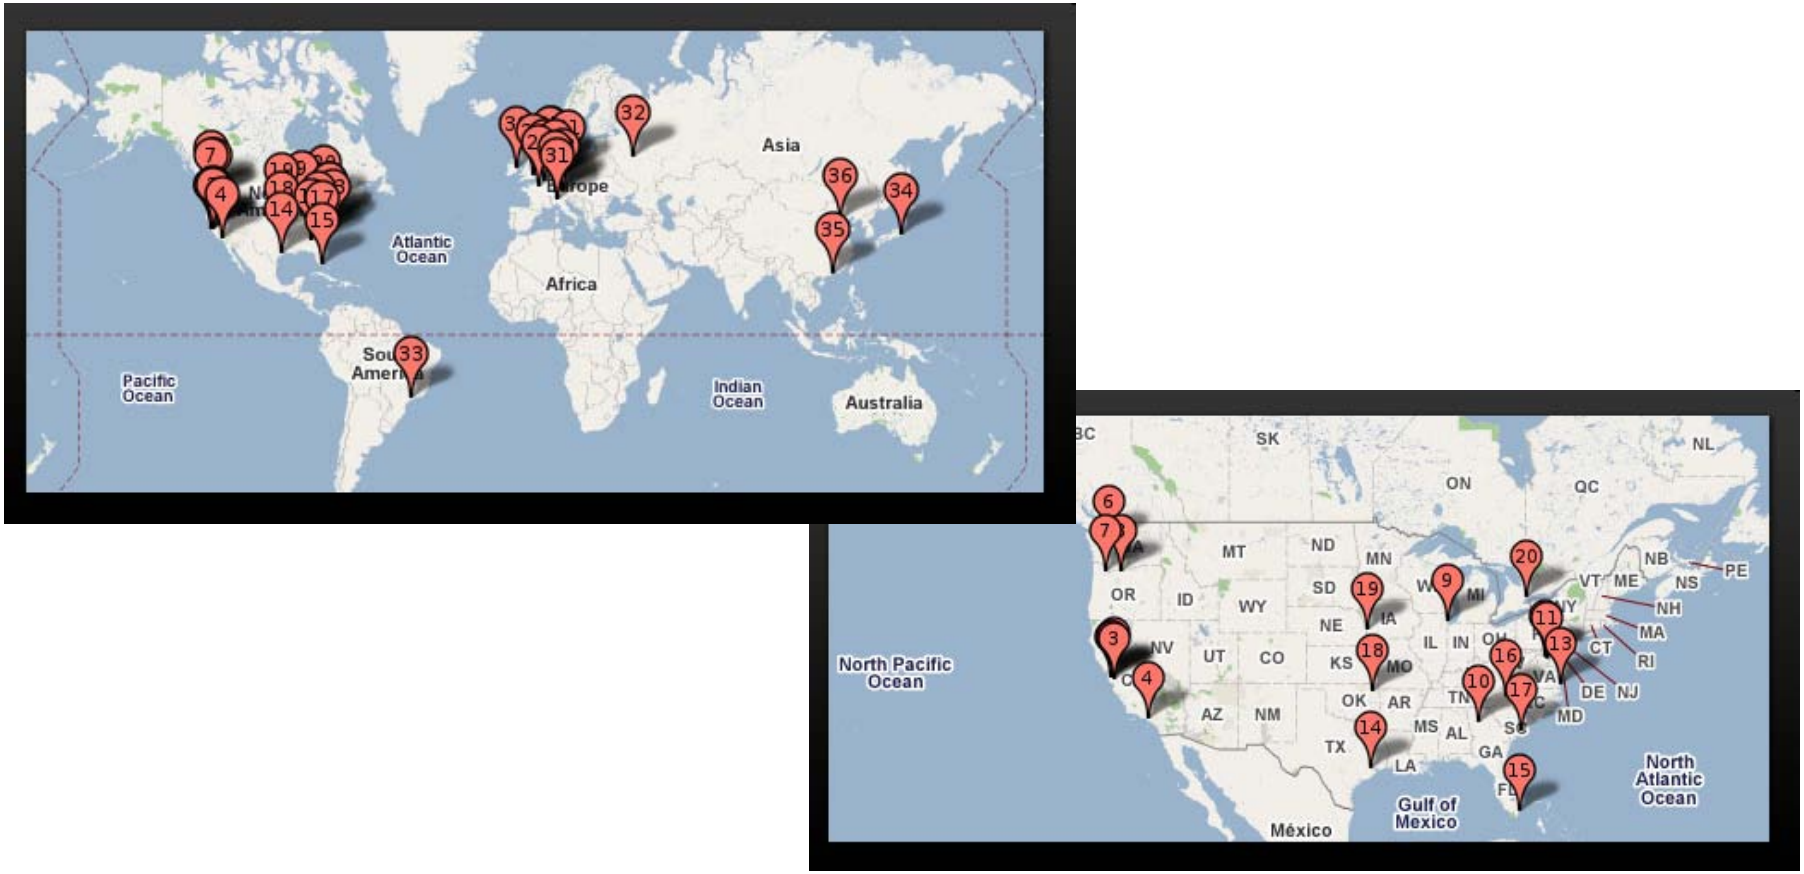

Google data centers, as reported by Royal Pingdom
http://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/

ciena

Mission Critical Data Center Interconnection

Interconnect applications

- → Data replication
  - → scheduled or synchronous
- → Survivability
  - → Business continuity/disaster recovery
- → Data center consolidation
- → Economical space and power
- → Virtualization

Key Requirements

- → Deterministic performance
- → Low and bounded latency
- → Lossless connectivity
- → High availability
- → Security

ciena

The Virtualized Data Center

Move to private, internal clouds is now

→ move to *federated internal* and *external* clouds is next

Virtualization will place new demands on the WAN

→ E.g., long-distance VMotion to migrate VMs outside the internal data centre

→Needing low-latency, zero-loss, high-performance WANs

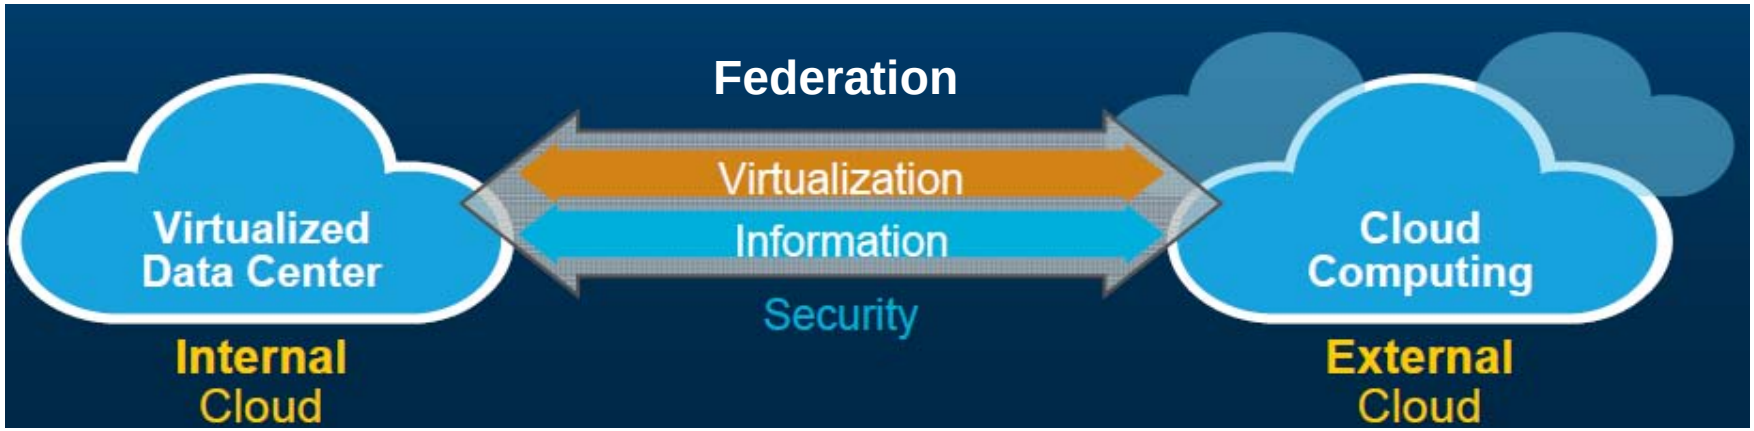

Source: Tucci Keynote EMC World 2009

ciena

Moving Data at on 100G Waves

Supercomputing'08 Demo

> 8Tb/s on fiber

CALTECH

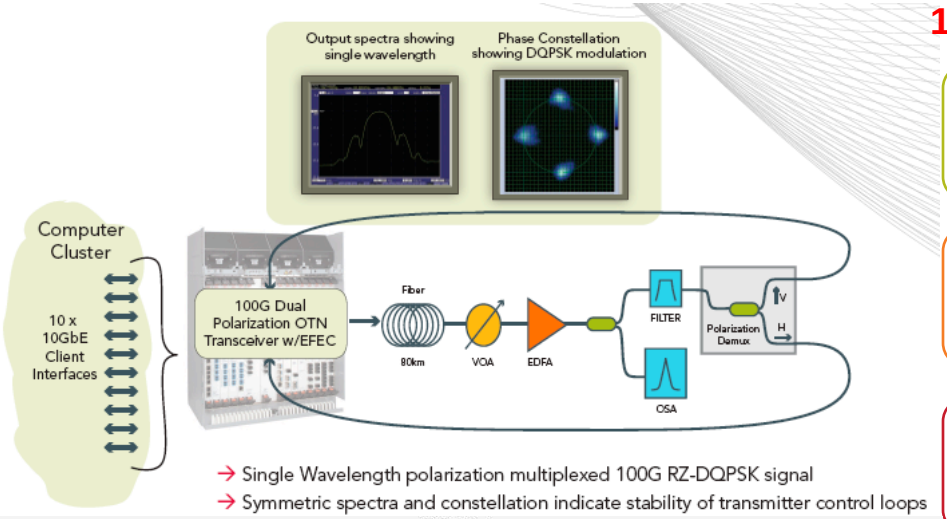

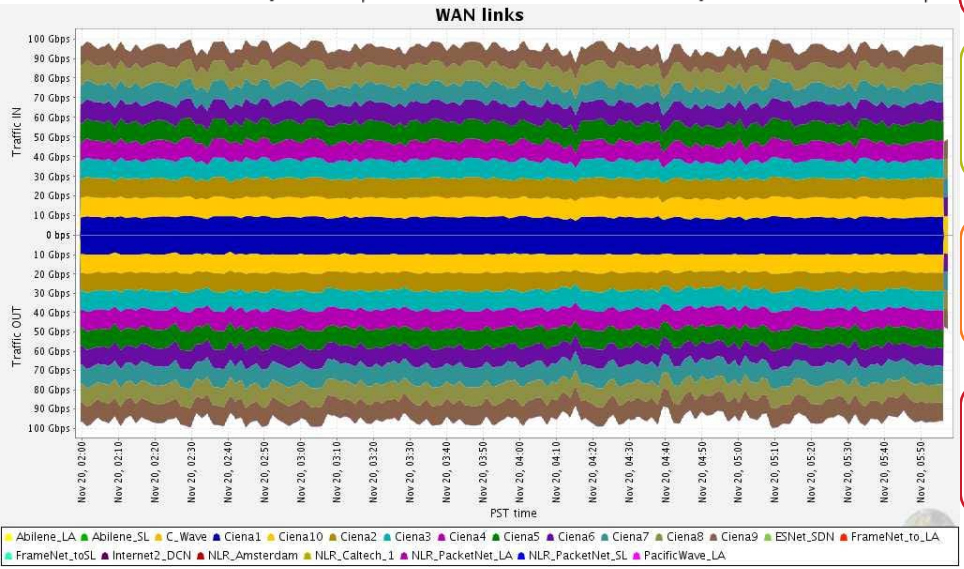

1 Petabyte of data transfer in 12hours on 1 wavelength

Single Wavelength 100G @ 50Ghz Spacing

Full C-Band Tunable, ROADM Compliant

Asynchronous Multiplexing 10x10G

Any 10G Client on Any Port

100G 8dB+ Enhanced FEC

Polarization Multiplexed

RZ-DQPSK Modulation Format

Low Latency

Commercial Form Factor

ciena

Infiniband Over SONET – Joint with ORNL/NLCF

collaborators: Makia Minich, Steven Carter

- Infiniband is effective data transport protocol for storage networks (few meters):
- TCP is not easily extended or not optimal for such data transfers

Question: Is IB effective over wide-area? - Yes

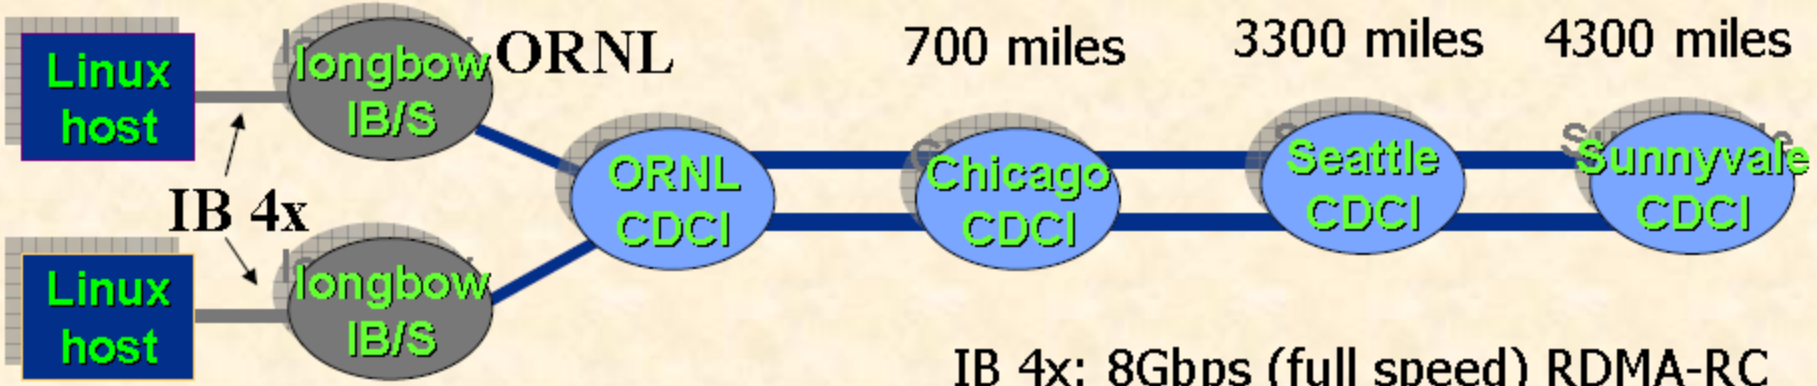

IB 4x: 8Gbps (full speed) RDMA-RC
Host-to-host local switch:7.5Gbps

ORNL loop -0.2 mile: **7.5**Gbps

ORNL-Chicago loop – 1400 miles: **7.46**Gbps

ORNL- Chicago - Seattle loop – 6600 miles: **7.23**Gbps

ORNL – Chicago – Seattle - Sunnyvale loop – 8600 miles: **7.20**Gbps

"Applications and networks, How to bridge the performance gap", OFC2007 Nagi Rao, Qishi Wu, Steven Carter and Bill Wing, ORNL

Internet2 Dynamic Circuit Network

Intelligent Optical Switching Nodes

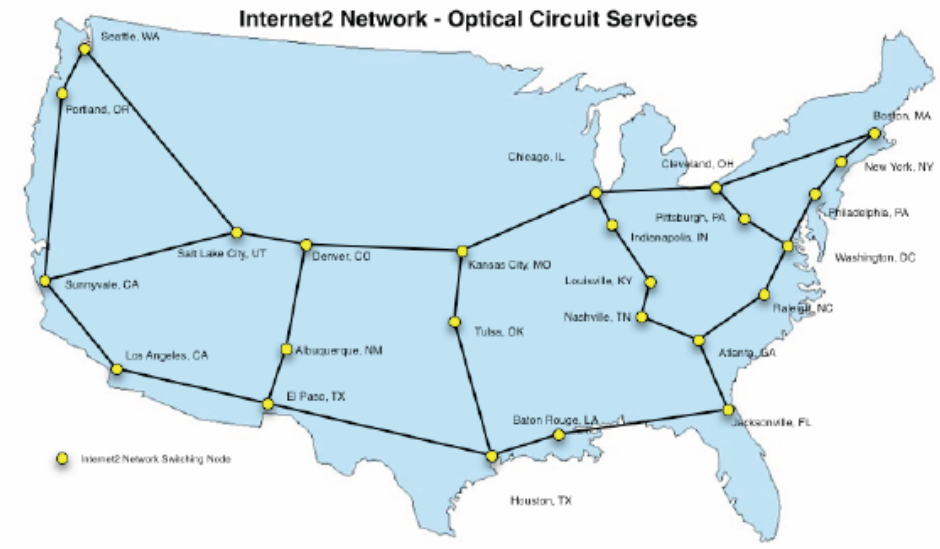

Optical Networking Cost Improvement Curve

Backbone DWDM per-bit, per-km cost improvements over time

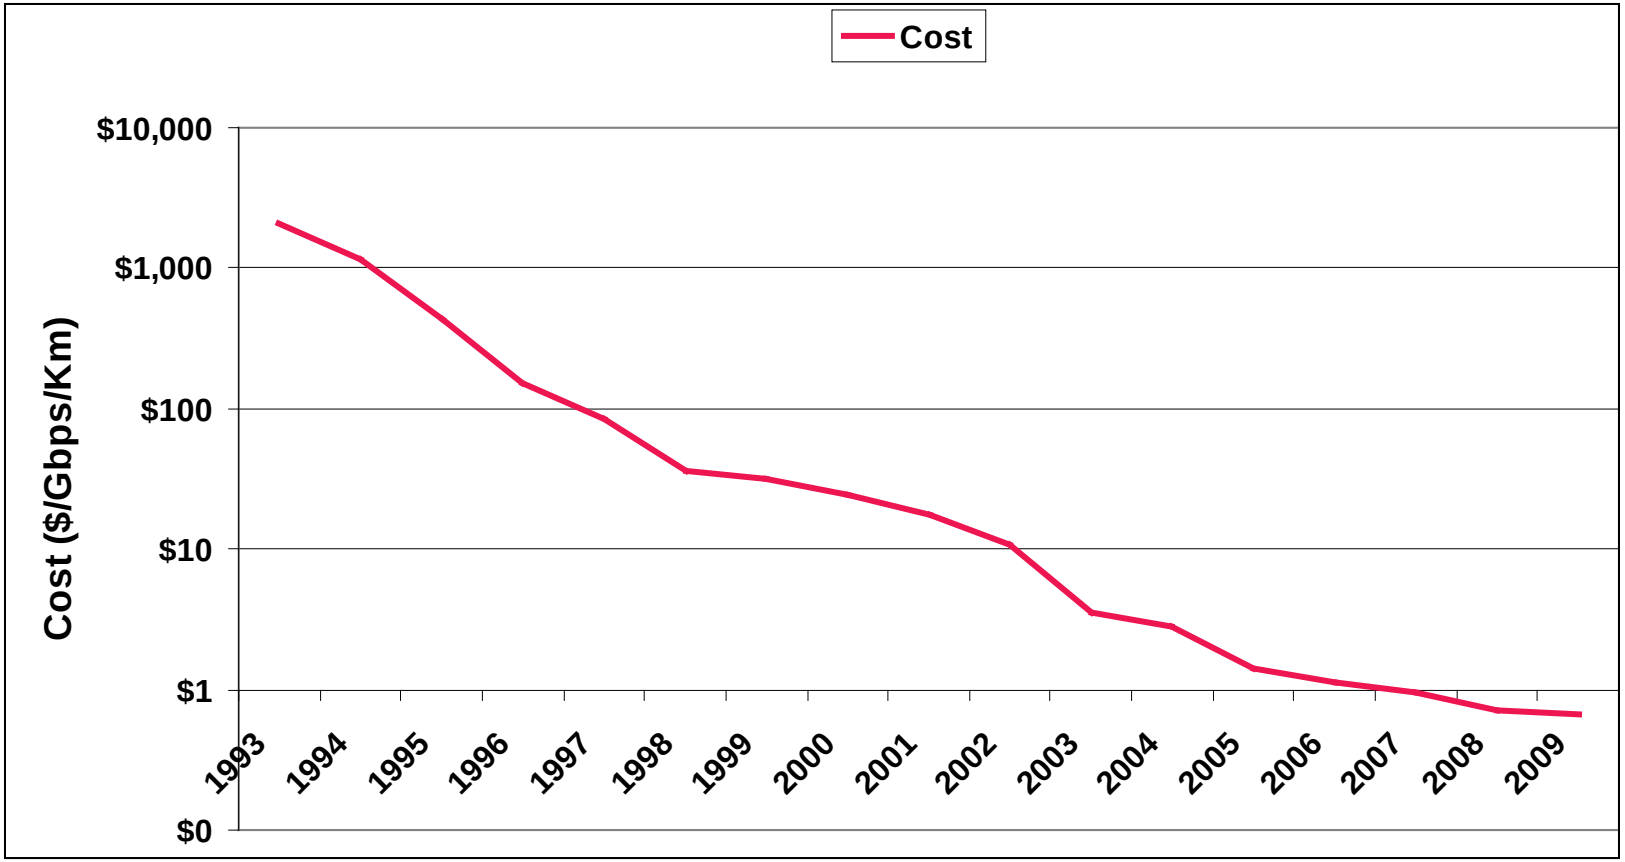

Source: Ovum

ciena

Economic Metrics for the Optical WAN

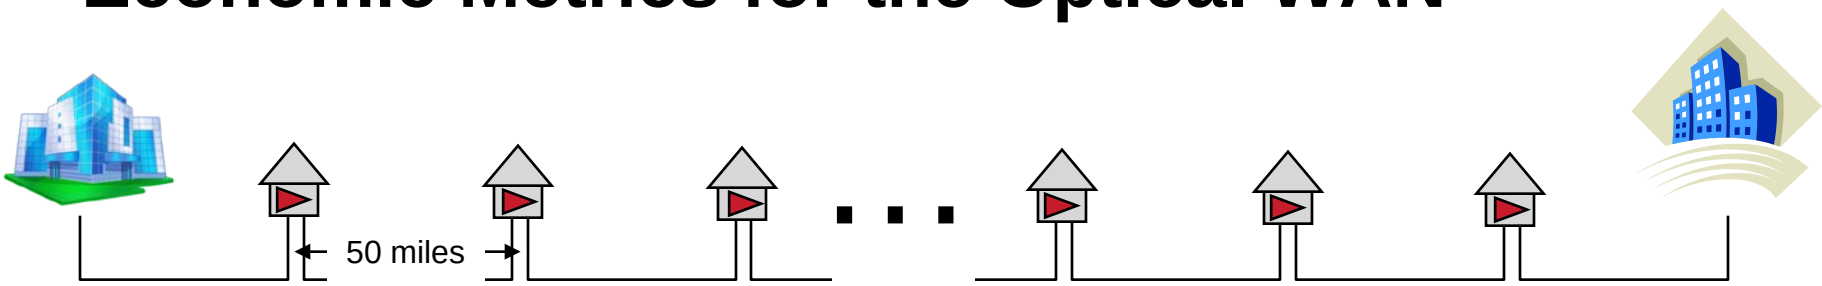

Optical Line System - transmission medium

- → **Optical fiber cost - $/mile**
- → **Optical amplifiers - $/unit**

Cost Relationships for Optical Line System

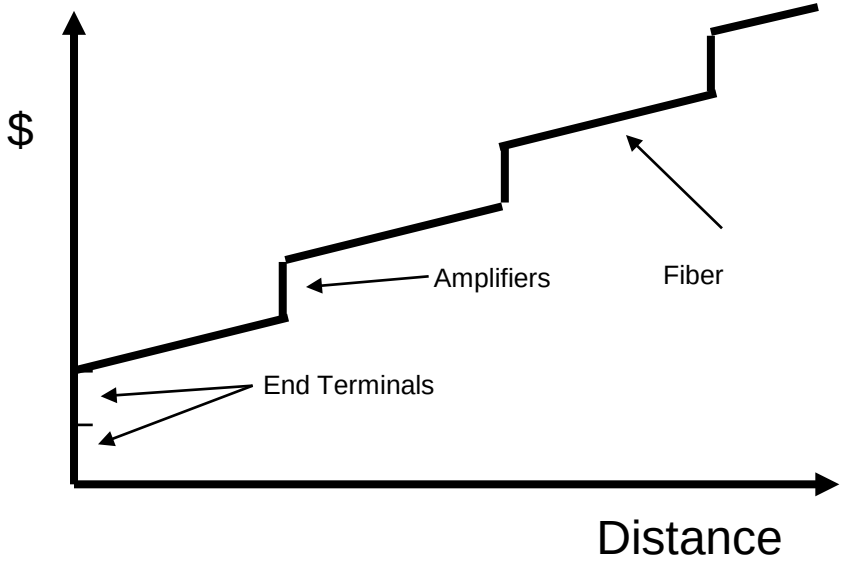

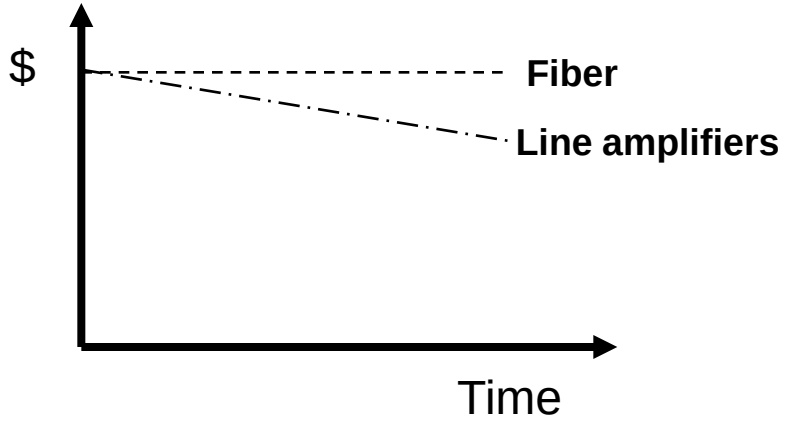

Cost ~ independent of capacity

ciena

DWDM Transceiver Cost Trends

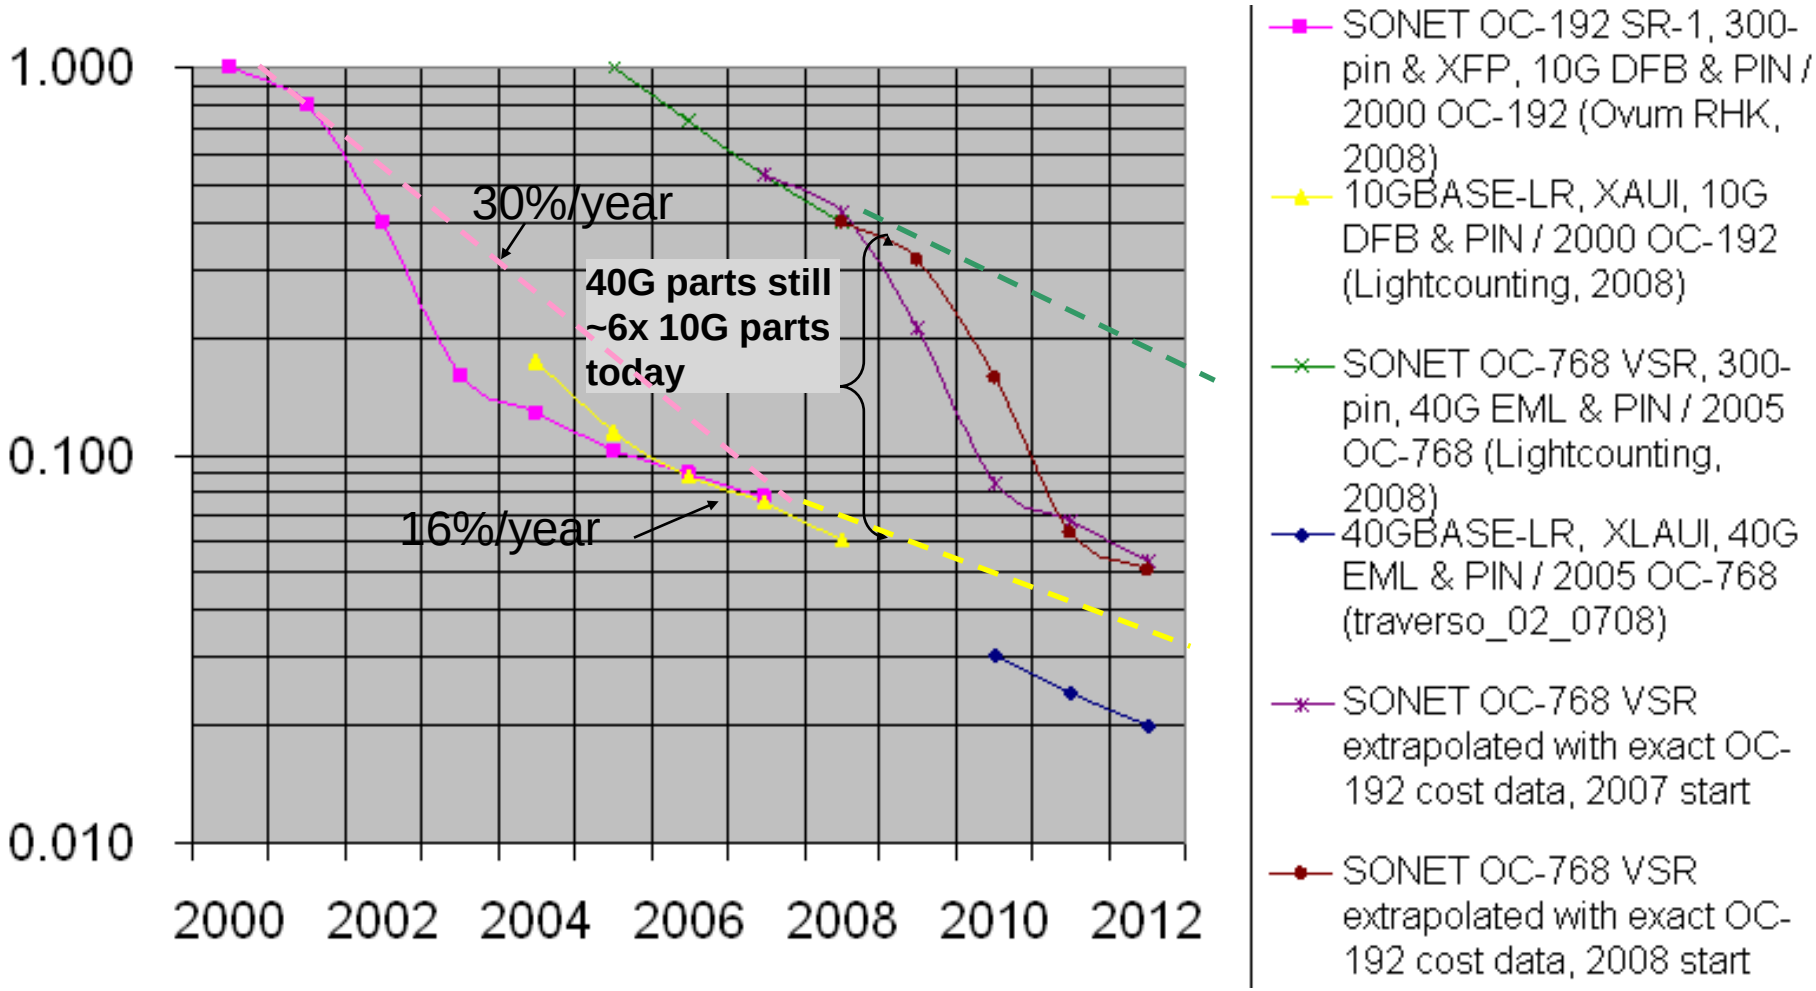

From IEEE 802.3ba Task Force

ciena

Economic Metrics for the Optical WAN

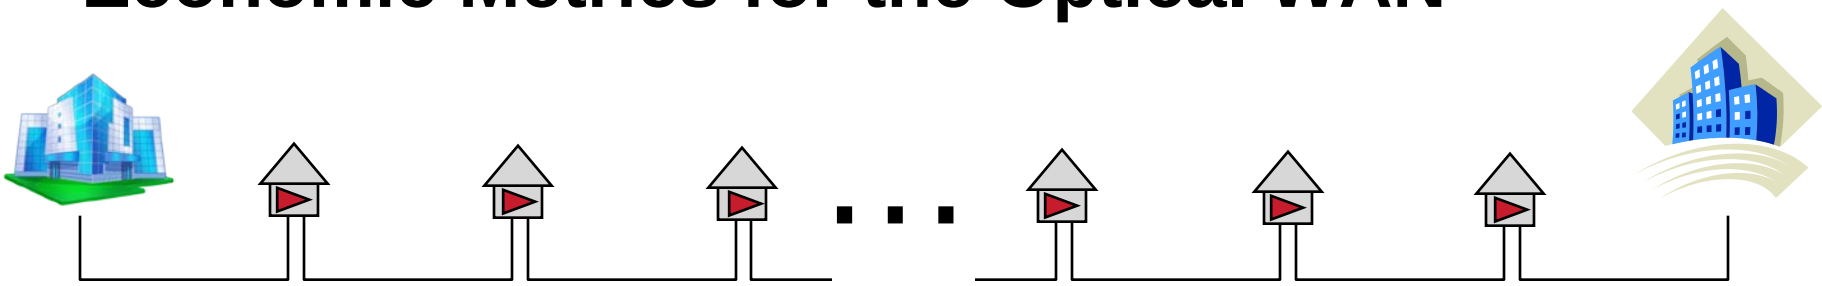

Optical Line System - transmission medium

- → **Optical fiber cost - $/mile**
- → **Optical amplifiers - $/unit**

DWDM transceivers

- → **Cost per bit per second**
- → **System capacity - Spectral efficiency**
  - → **bits per second per Hertz**
- → **Maximum unregenerated reach – miles**

Key Metrics

$/bits per second — **Low bandwidth needs** — $/bits per second per mile

$/bits per second per Hertz — **High bandwidth needs** — $/bits per second per Hertz per mile

← Metro — Regional — Long Haul →

ciena

Optical WAN
Cost Elements

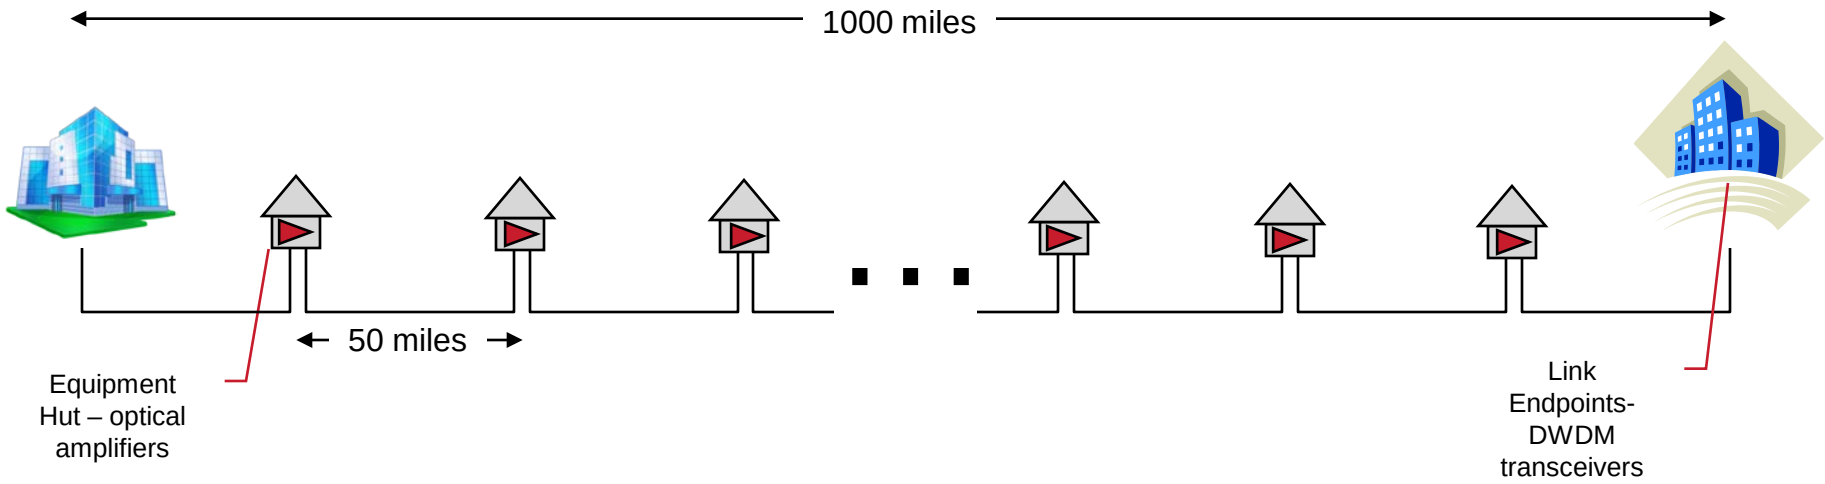

Capital Cost Assumptions

→ Optical Fiber	$1,000/mile
→ Optical amplifiers	$50,000/unit
→ DWDM transceivers	$1000/Gbps

Operational Cost Assumptions

→ Fiber route maintenance	$100/mile/year
→ Equipment space/power	$10,000/site/year
→ Equipment maintenance	10% of capital cost

Line costs: Fiber/amps

→ 200 mile example	$ 450,000
→ 1000 mile example	$2,050,000

Transceiver costs – fully populated DWDM link

→ 10G transceivers	0.8Tbps	$ 1,600,000
→ 40G transceivers	3.2Tbps	$ 6,400,000
→ 100G transceivers	8Tbps	$16,000,000

ciena

Key Observations on LD Transmission Costs

Fibers and amplifiers represent a fixed cost that is only slowly improving

→ **Increasing fiber bandwidth spreads this cost**

At 10G costs for transceivers and fiber paths are similar for LD links

→ ***Even at constant cost/bit*** **for 100G transceivers, total cost/bit/mile declines**

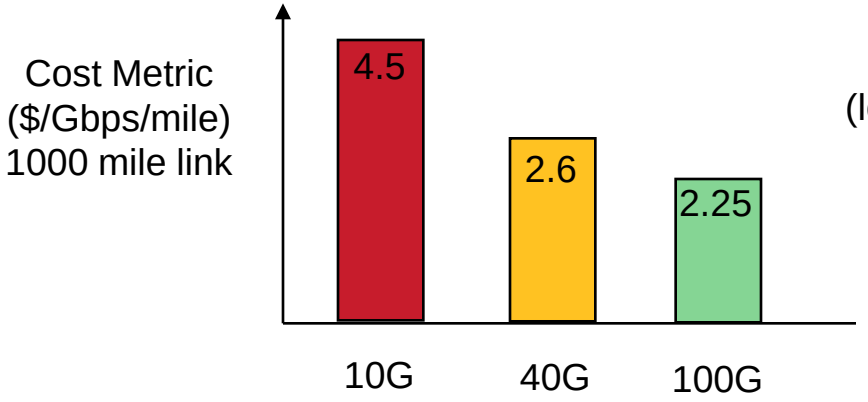

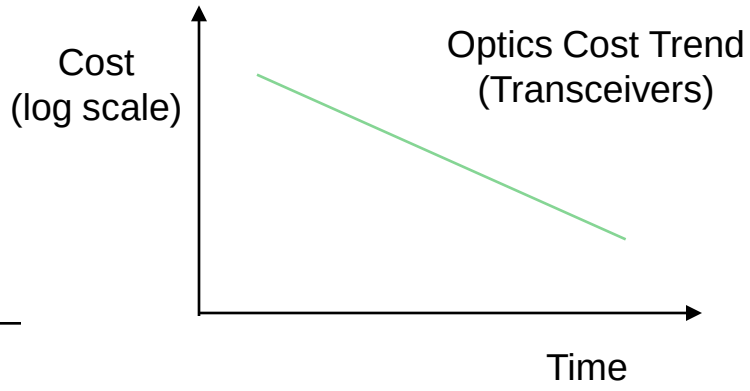

100G DWDM Transceivers – Today's Technology

10 X 10G multiplexed to 100 G on a single wavelength

Muxceiver **PolTracker**

10GbE Client

56Gbps Line

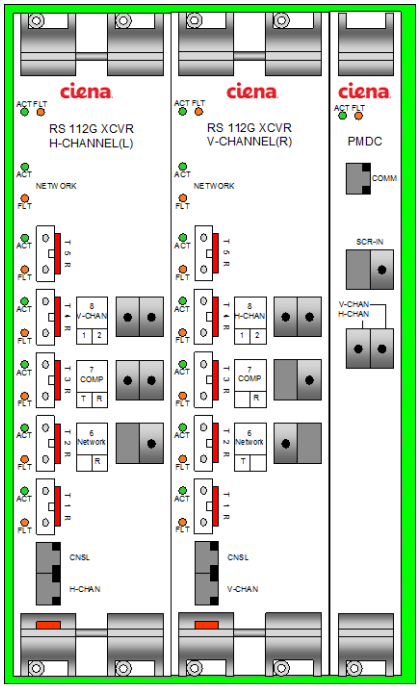

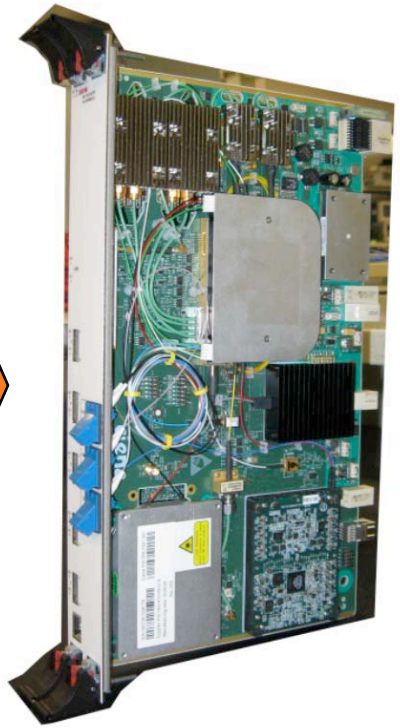

100G Technology Direction

Integrated photonic components for modulator and demodulator

Today
Separate Components

Ciena Direct Detection Transceiver

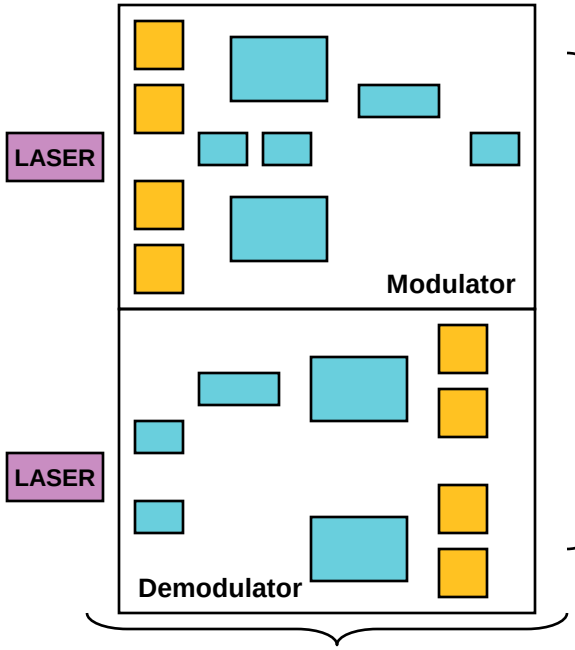

5 card slots
(two for clients, two for DWDM)

Near Future
Module Integration

Current phase of OIF 100G project

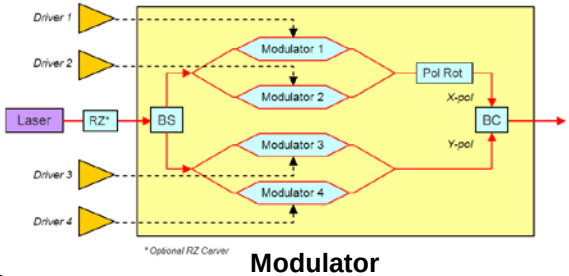

Modulator

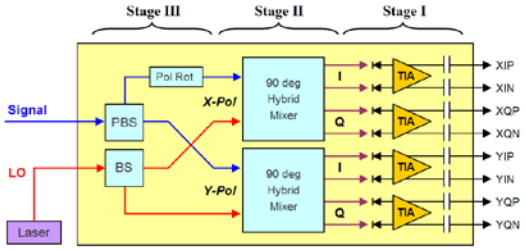

Demodulator

2 card slots
(one for clients, one for DWDM)

Long Term Target
Monolithic Integration

InP or Silicon Photonics

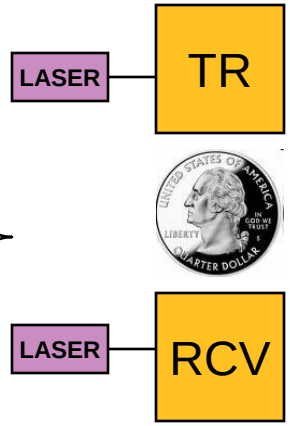

Future
Multiple 100Gs
Per line module

ciena

Summary

100G Transmission is here today!

Optical transmission systems capacity will scale 10X with 100G

Initial high cost of 10G will improve with new technologies

High capacity optical transmission and switching can enable rapid data movement among distributed data centers

ciena