



The Ethernet EcoSystem - Broad Market Potential for 100 Gb Ethernet

Joel Goergen Force10
Mark Nowell Cisco

Supporters

- Pete Anslow – Nortel Networks
- Med Belhadj – Cortina Systems
- Mike Bennett – LBNL
- Adam Bechtel - Yahoo
- Joe Berthold – Ciena
- Frank Chang – Vitesse
- Chris Cole – Finisar
- Ed Cornejo – OpNext
- Dan Dove - ProCurve Networking by HP
- Peter Harrison - Netflix
- Ken Jackson – Emcore
- John Jaeger – Infinera
- Wenbin Jiang – JDSU
- Jack Jewell – Picolight
- Alan Judge – Amazon
- Mark Kortekaas – CBS
- Mike Krause – HP
- Joe Lawrence – Level3 Communication
- Andy Moorwood - Extreme Networks
- Terry Morris – HP
- Gary Nicholl – Cisco
- Tom Palkert – Xilinx
- Shashi Patel – Foundry Networks
- Neil Peers - Adva
- Petar Pepeljugoski – IBM
- Drew Perkins – Infinera
- Liam Quinn – Dell
- Bill Ryan – Foundry Networks

Supporters (cont'd)


- Vik Saxena - Comcast
- Peter Schoenmaker – NTT
- Ted Seely - Sprint
- Hank Steenman – AMS IX
- Takejiro Takabayashi – JPIX
- Sashi Thiagarajan – Ciena
- Matt Traverso – OpNext
- Bill Trubey – Time Warner Cable
- Brad Turner – Juniper Networks
- Jason Weil – Cox Communications
- Doug Wilson - Microsoft
- Robert Winter – Dell
- Masato Yamanishi – Softbank BB
- Vish Yelsangikar – Netflix

Broad Market Potential

- IEEE 802 LMSC Criteria
 - Broad sets of applications
 - Multiple vendors and numerous users
 - Balanced cost (LAN versus attached stations)

The Ethernet Ecosystem Looking in the Core

Broadband Access

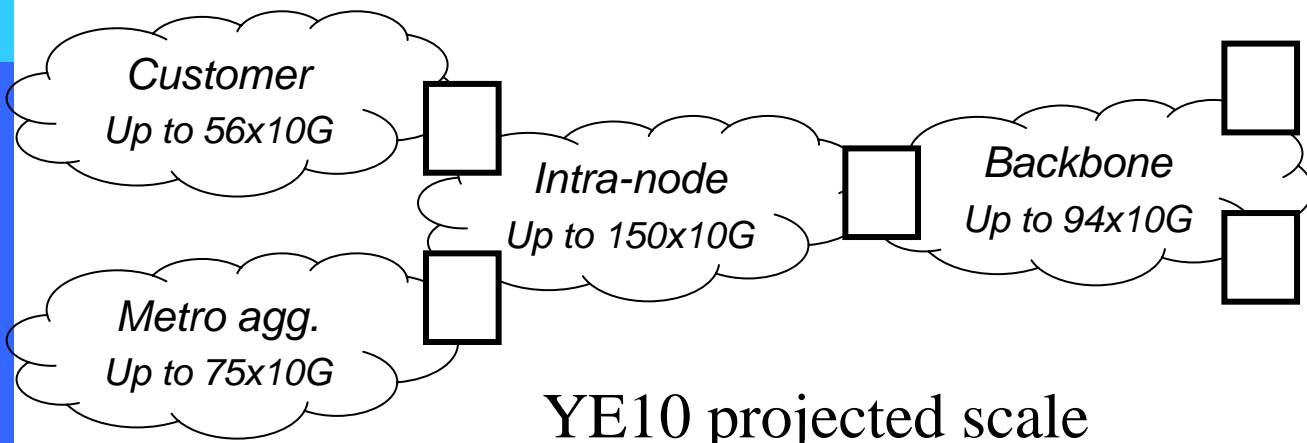
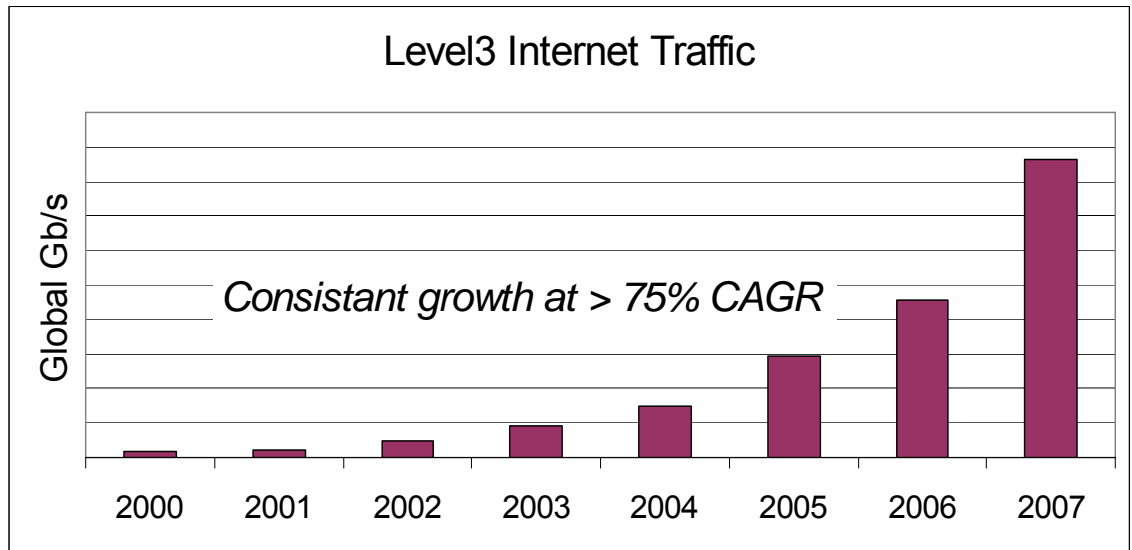



Data Centers and Enterprise

Research, Education and Government Facilities

Level 3

Ongoing annual growth of 75% drives a 9.4x increase in four years

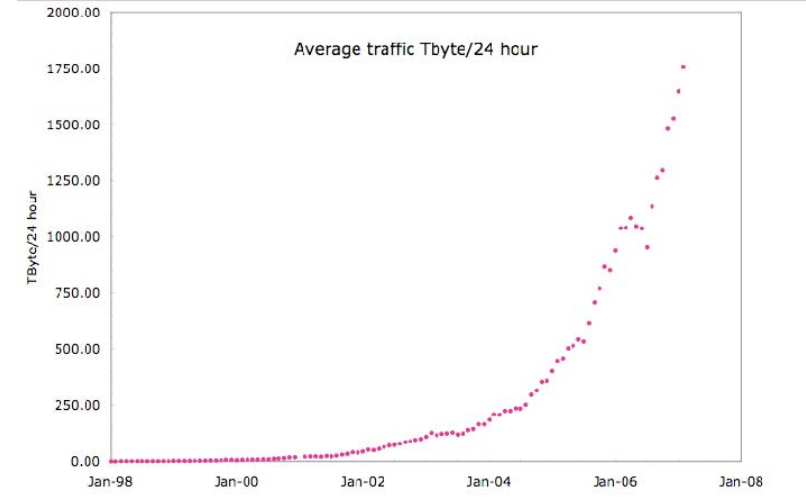



Unlikely that 10GE alone can meet this scale


YE10 projected scale

AMS-IX Daily Traffic

Feb 2006



Feb 2007



Feb 2007


- 71 10GbE
- 10 Queued Requests
- 8 (2 x10GbE)
- 2 (3x10GbE)

End of 2007

- Estimate 128 10GbE Ports

The Ethernet Ecosystem: Looking Outside the Core

Broadband Access



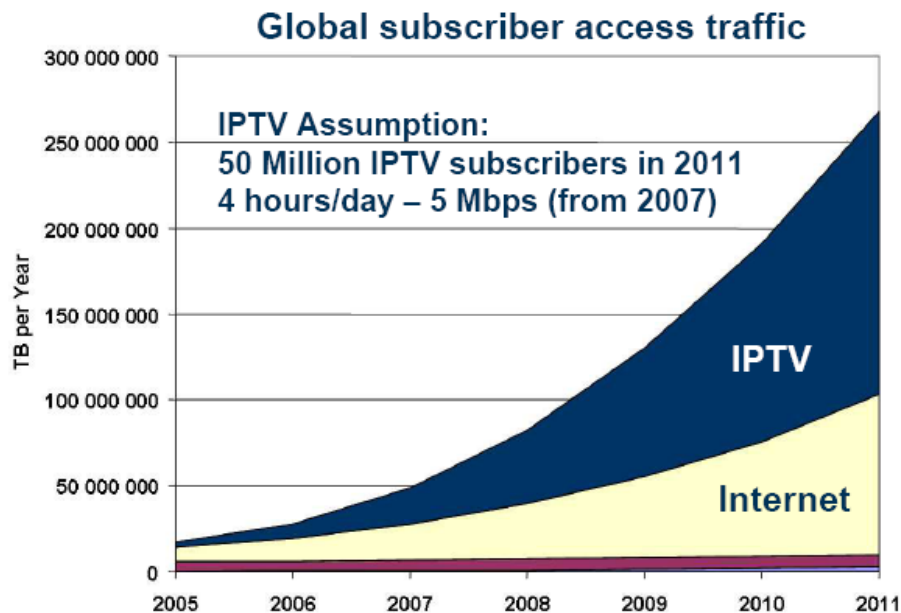
Data Centers and Enterprise

Research, Education and Government Facilities

Traffic Change & Growth

Broadband access

Fixed access traffic growth



Access traffic growth through increased:

- broadband penetration
- bandwidth demanding services

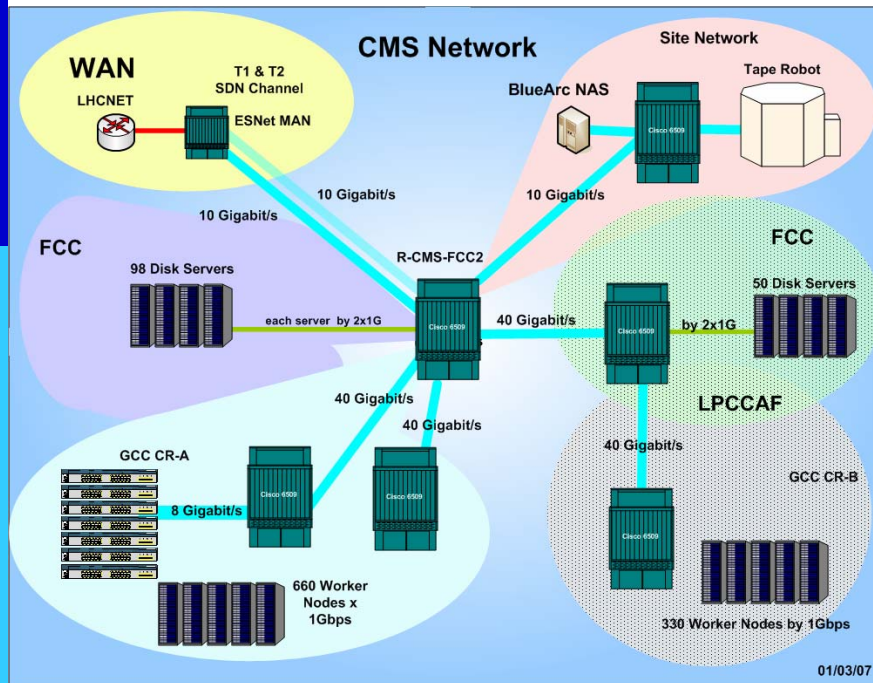
Service	Bandwidth
VoIP	100 kbps
IMS/Video conf	0.7 - 1.5 Mbps
Internet	0.2 - 5.0 Mbps
Gaming	0.2 - 0.5 Mbps
SDTV (MPEG-2)	6 Mbps
SDTV (MPEG-4)	3 Mbps
HDTV (MPEG-2)	20 Mbps
HDTV (MPEG-4)	10 Mbps

Typical service portfolio (2010):

1 HDTV, 2 SDTV, gaming, voice, high-speed internet → 25 – 30 Mbps (MPEG-4)

Reference: Alping_01_1106.pdf

R&D Example




CMS Experiment, Based in CERN Switzerland

- Data Distribution Model is based on Tiers
- Tier 0: 1 at CERN (raw data collected)
- Tier 1: 7 located throughout the World
 - data storage – some % of total data
 - data distribution
 - data reprocessing

Local Connectivity Model

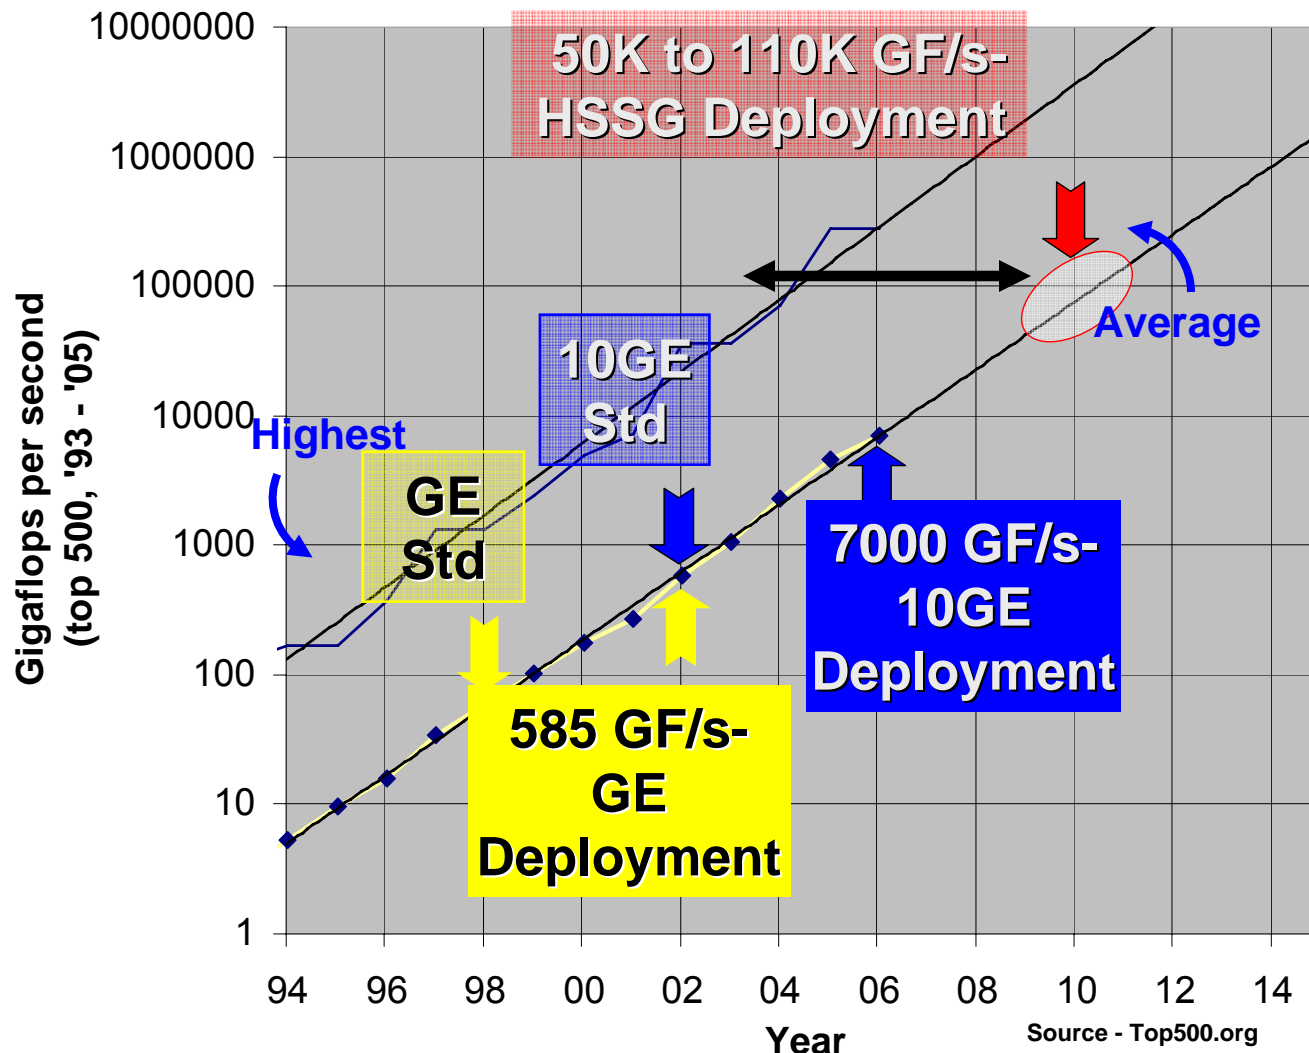
- Connectivity via large dense switch fabrics where possible
- More computer rooms are in progress
- High bandwidth interconnectivity between switches
 - Currently 40Gb aggregated
 - Minimal oversubscription
 - Scaling Issues are a concern

Network Utilization Between CMS Tier-1 Computing Facility Core Switches



Data courtesy of Donna Lamore, Fermilab

High Performance Computing



Historical Observations

- 4 years between standard and deployment
- GE captures 50% of market in 3 years (2002 - 2005)
- 10GE enters top500 in 2006 (included in Gigabit Ethernet #'s)
- 12x increase in Avg Gigaflops: 10x increase in Ethernet Interconnect

Forecast

- 10x increase needed at 84,000 Gigaflops (2010)

Per pepeljugoski_01_0107 - Future systems may have 10,000 processors & tens of thousands of fibers for each machine.

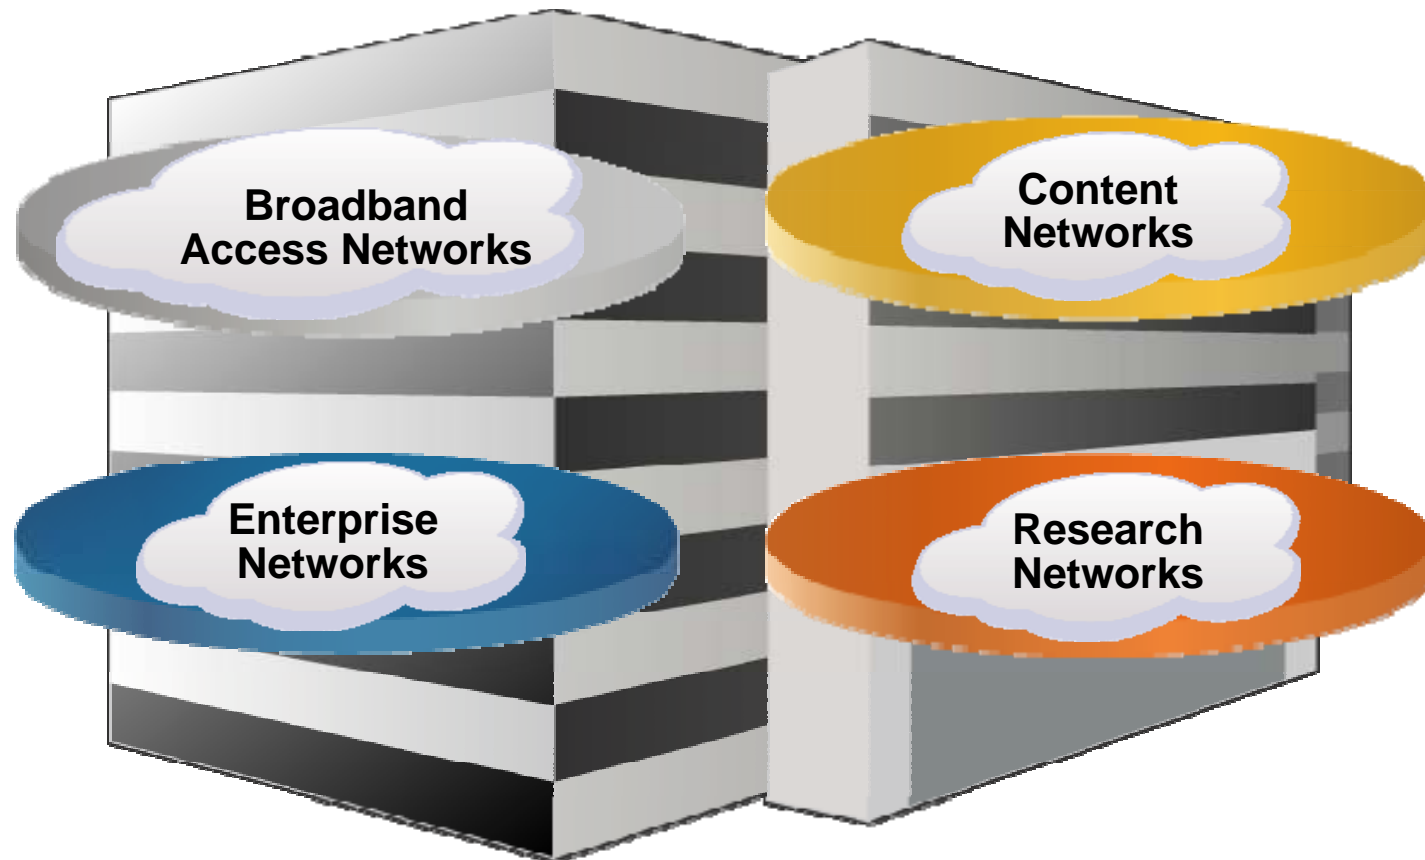
Enterprise Networks and Data Centers

- “Need for 100 GbE is apparent”
- Lagged (nx10) links are already being used internally
- “Carriers are turning down 10G services today”
 - Impacts their networks and growth
 - “Don’t judge our trunk capacities by what goes out of our facility”
 - Data may or may not leave facility
- Content
 - Unicast video
 - Cable operators
 - Carriers
 - Retailers
 - Rental
 - Beyond video
 - Amazon Web Services (<http://www.amazonaws.com/>)
 - Amazon Elastic Compute Cloud (ECS)
 - Amazon Simple Storage Service (Amazon S3)
- “Virtualization” driving demand for 100 GbE
- “Backup driving demand for 100 GbE”

Server Aggregation

- Consider rack mount servers
 - Gigabit Ethernet I/O
 - Horizontal server architectures can justify 100G links for aggregation
 - 7 years after 1000 BASE-T
 - 10 Gigabit Ethernet I/O
 - Transition over the next 3 to 5 years
 - 10GBASE-T completion, implementation, deployments
 - Offload, bus, and virtualization technologies will mature to fully use bandwidth.
- Consider blade servers & backplanes
 - Per IDC \approx 2.8M blade servers 2008 – 2009, estimate 500k chassis
 - Rule of Thumb
 - Create uplinks to support $\frac{1}{2}$ the total aggregate bandwidth
 - 1000BASE-KX
 - $\frac{1}{2} * 14 \text{ slots} * \text{GbE} = 7\text{G Uplink}$
 - $\frac{1}{2} * 14 \text{ slots} * 4 \text{ GbE} = 28\text{G Uplink}$
 - XAUI / 10GBASE-KX4 / 1 Port 10GBASE-KR-
 - $\frac{1}{2} * 14 * 10\text{GbE} = 70\text{G Uplink}$
 - **500K Chassis * 2 Ports = 1M 100G Ports Opportunity**
 - 4 Ports 10GBASE-KR
 - $\frac{1}{2} * 14 * 4 * 10 \text{ GbE} = 280\text{G Uplink}$
- **Existing network infrastructure is modeled on x10 aggregation scheme.**
- **Importance of cost for deployment**
 - **Today's cost**
 - **Long term ROI**


Single Site Networks



Network consolidation – supporting multiple faces of the Eco-System

Access

Broadband Access



 SoftBank BB

Data Contributed by Softbank BB Corp.

Provider – Softbank BB

Service Area – Japan

Subscribers – 5M (Jan 07)

Access Types – ADSL, ADSL2,
ADSL2+, GE-PON

Rates:

ADSL – 8Mb/s to 50 Mb/s

GE-PON – 100 Mb/s

Traffic Projection (to 2010) – 30%
Annual Growth

10 GbE (2010) – 1.5K to 2K Ports

100 GbE (2010) – 50 to 100 Ports

Summary

- Demand is clear to see for:
 - Carriers
 - Internet exchanges
- Demand not isolated to carriers and Internet exchanges:
 - R&D
 - HPC
 - Enterprise
 - Access
 - Server Aggregation
- Nature of traffic changing – Content is King!
- Carriers turning down 10GbE services
 - Impacts 10GbE deployment
 - Impacts customers trunk requirements
- 100GE intersects the market need. 10x is the right increment step for multiple applications.