Next Generation BASE-T Call For Interest

July 2012,
San Diego, CA

CFI Panel Members

- Presenters:
  - Bill Woodruff – Broadcom
  - Dave Chalupsky – Intel
- Supporters and experts for the Question & Answer session:
  - Dan Dove – APM
  - Brad Booth – Dell
  - David Koenen – HP

Supporters (58 Individuals from 39 Companies)

Rick Rabinovich	Alcatel-Lucent	Thananya Baldwin	Ixia
Dan Dove	APM	Jerry Pepper	Ixia
Venkatesh Nagapudi	APM	Anthony Ng	JDSU
Kamal Dalmia	Aquantia	Alan Flatman	LAN Technologies
Jeff Hirschman	Arista	Michael Bennett	LBNL
Yakov Belopolsky	Bel Stewart	Keith Kosanovich	Leviton
Will Bliss	Broadcom	David Dwelley	Linear Tech
Wael Diab	Broadcom	LK Bhupathi	Marvell
Michael Grimwood	Broadcom	Sudhakar Gundubogula	Marvell
Heng-Hsin Liao	Celestica	Yair Darshan	Microsemi
Peter Anslow	Ciena	Harry Forbes	Nexans
Hugh Barrass	Cisco	Paul Vanderlaan	Nexans
Mark Nowell	Cisco	Sterling A. Vaden	OCC
George Zimmerman	CME/Commscope	Shimon Muller	Oracle
Mabud Choudhury	Commscope	Ronald Nordin	Panduit
Paul Kolesar	Commscope	Rick Pimpinella	Panduit
Wayne Larsen	Commscope	Ronald Cates	PLX
Richard Mei	Commscope	Stephen Bates	PMC-Sierra
Brad Booth	Dell	Harshang Pandya	Psiber Data
John Dambrosia	Dell	Zhu Xing	Psiber Data
Rich Hernandez	Dell	Thuyen Dinh	Pulse
Steve O'Hara	Fluke Networks	Joseph Chou	Realtek
Steve Carlson	High-Speed Design	Valerie Maguire	Siemon
David Koenen	HP	Bruce Tolley	Solarflare
Robert Grow	Independent	Allan Nielsen	TE Connectivity
Siddharth Sheth	Inphi	David Estes	UNH-IOL
Andre Szczepanek	Inphi	Jeff Lapak	UNH-IOL
Ilango S. Ganga	Intel	Mandeep Chadha	Vitesse
Kent Lusted	Intel		
Richard Melitz	Intel		

Objectives for the meeting

- To measure the interest in starting a study group for Next Generation BASE-T.

- At this time, we **don't need to**
  - Fully explore the problem
  - Debate strengths and weaknesses of solutions
  - Choose any one solution
  - Create PAR or five criteria
  - Create a standard or specification

Agenda

- Market Opportunity
- Technical Viability
- Q&A
- Straw Polls

Market Opportunity

Defining the Market Opportunity

- What applications will benefit from a Next Generation (higher speed) BASE-T?
- Why will a twisted pair solution be desired?
- Projection of volumes

Server Ethernet Connections

- Server LAN ports presented here as a primary driver for a Next Gen. BASE-T
- Next Gen. BASE-T Study Group may define objectives based on the requirements of this application
  - Example: reach need much less than 100m…
    - Determine reach based upon application need while balancing power, cost and complexity
  - The needs of other applications may also be considered by the Study Group

Twisted Pair Enables…

Data center switching with flat 2-layer topology

- Server connections to access switch
  - Top of Rack (ToR) switch topology
  - Middle of Row / End of Row topologies
- Access link benefits from Twisted Pair distances
  - Eliminates stranded ports inherent in 42U rack height
  - Not all servers are 1RU, not all racks have 40 servers
  - Not all access switches are at the top of rack
    - Flexibility in physical topology
- Multiple speed generations on compatible infrastructure
  - Allows incremental upgrades, mixed environments
  - Use the ‘big pipes” where you need them
  - Structured cabling for flexible moves/adds/changes

Twisted Pair on Servers

- New LAN technology starts as add-in card, moves to Motherboard
  - 10/100/1000BASE-T LOM is the incumbent
  - 10GbE has been mostly add-in, 10GBASE-T ramping as LOM now
    - Support for 1000BASE-T & 10GBASE-T on same port is a key enabler for LOM
- Add-in Card Port Density
  - Quad port 1000BASE-T NICs very common, shipping more ports than dual port NICs.
  - Quad port 10GbE cards already in the market
  - Reasonable to assume that quad 40GbE will be needed
- Path to quad port 40GbE adapters
  - PCIe Low Profile card the most common add-in form factor
  - Four RJ45's fit on a PCIe Low Profile card
  - Four QSFP+ too large for the PCIe form factor

RJ45 and four-pair cabling provide density and compatibility

Twisted Pair on Switches

48 ports + Uplinks in a 1RU form factor

- 640Gb/s as 48x10Gb/s + 4x40Gb/s is available today
- 2.4Tb/s as 48x40Gb/s + 4x100Gb/s likely available in 2014
- RJ45-size connectors enable 48 ports in 1RU
- Twisted pair enables
  - End of Row, Middle of Row switch
  - Full utilization; no stranded ports
  - Use of structured cabling

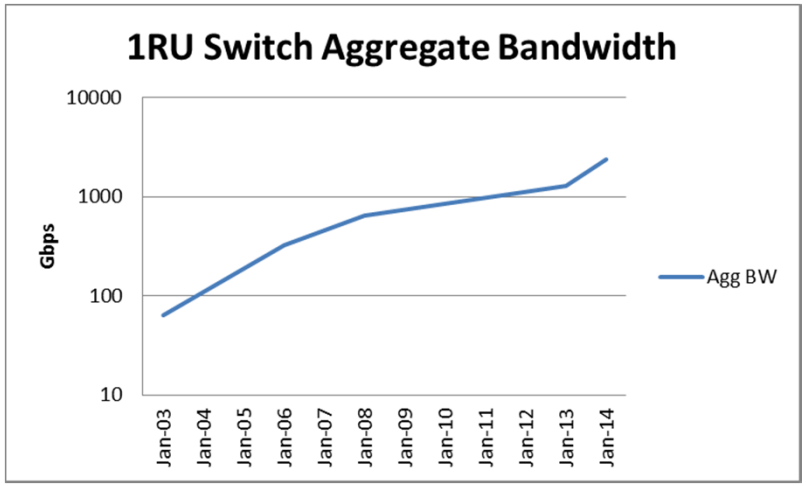

Source: Dell'oro

Server Market Perspective

- BASE-T Family is still the highest volume Ethernet port type today
- New technology introduced on add-in cards…
  - But LAN On Motherboard (LOM) drives the highest port volume
- Platform transitions drive new networking requirements
  - 2 to 3 year design cycle for Volume Servers
  - New platforms are the opportunity to design in new LOM technology
- Committing to new LOM technology requires:
  - Low Cost / Reasonable Power
  - High adoption / utilization rate
  - Compatibility with legacy speeds & infrastructure
- 10GbE LOM is growing
  - Led by blade servers with backplane Ethernet
  - Underway now with 10GBASE-T for rack servers
  - LOM-replacement daughter cards providing options during the transition

Server Market Trends

- A rapidly growing & changing environment
- System innovation on all fronts to serve different use models
  - Drive to Exascale this decade for high performance computing
    - 1000x increase over today's fastest computers
  - Density Optimized Servers
    - For Internet Portal Datacenter (IPDC), Cloud Computing, Social Networking
  - MicroServers
    - Applying many, many small processors to the problem.
    - Keeps 1GbE around longer, move to 10GbE drives need for 40GbE uplinks
- Variety of workloads and applications need different balance of compute vs. I/O capability
  - The variety of processor grades is increasing
  - 1GbE through 100GbE ports will co-exist in the market

This is not a homogeneous market!

Simplifying that Complex Story

Adoption can be summarized in three categories:

- **Fringe**: Users who demand the most possible bandwidth. The servers that would need this bandwidth would typically be the high end 4 or 8 socket versions.
- **Performance**: Users who demand more I/O performance due to virtualization or, in some cases, the desire to converge the SAN and LAN networks within the rack
- **Free**: A large portion of server buyers will only implement what is offered as the base configuration. These buyers would choose the “Free” option

    - Kimball Brown, LightCounting, 2011.
    - http://grouper.ieee.org/groups/802/3/ad_hoc/bwa/public/jul11/brown_01a_0711.pdf

- Brown did not apply percentages… reasonable estimate:
- Fringe <5%, Performance <20%, Free ~75% of the market

Server Port Forecasts

Past & Present

- x86 Server Ethernet port speed forecasts used in prior CFIs and Study Groups to:
  - Indicate broad market potential
  - Determine project timing
    - Note
      - x86 Servers are only part of the server market
      - A lot of “server-like” platforms are in service as network & storage appliances, not counted as servers.
      - They also need a switch connection
      - Result: Overall market opportunity several times higher than just x86 Server

Let’s look at prior Server NIC forecasts given in IEEE 802.3 and compare to today…

x86 Server Ethernet Connection Speeds with 40GbE & 100GbE

Hays_01_0407.pdf

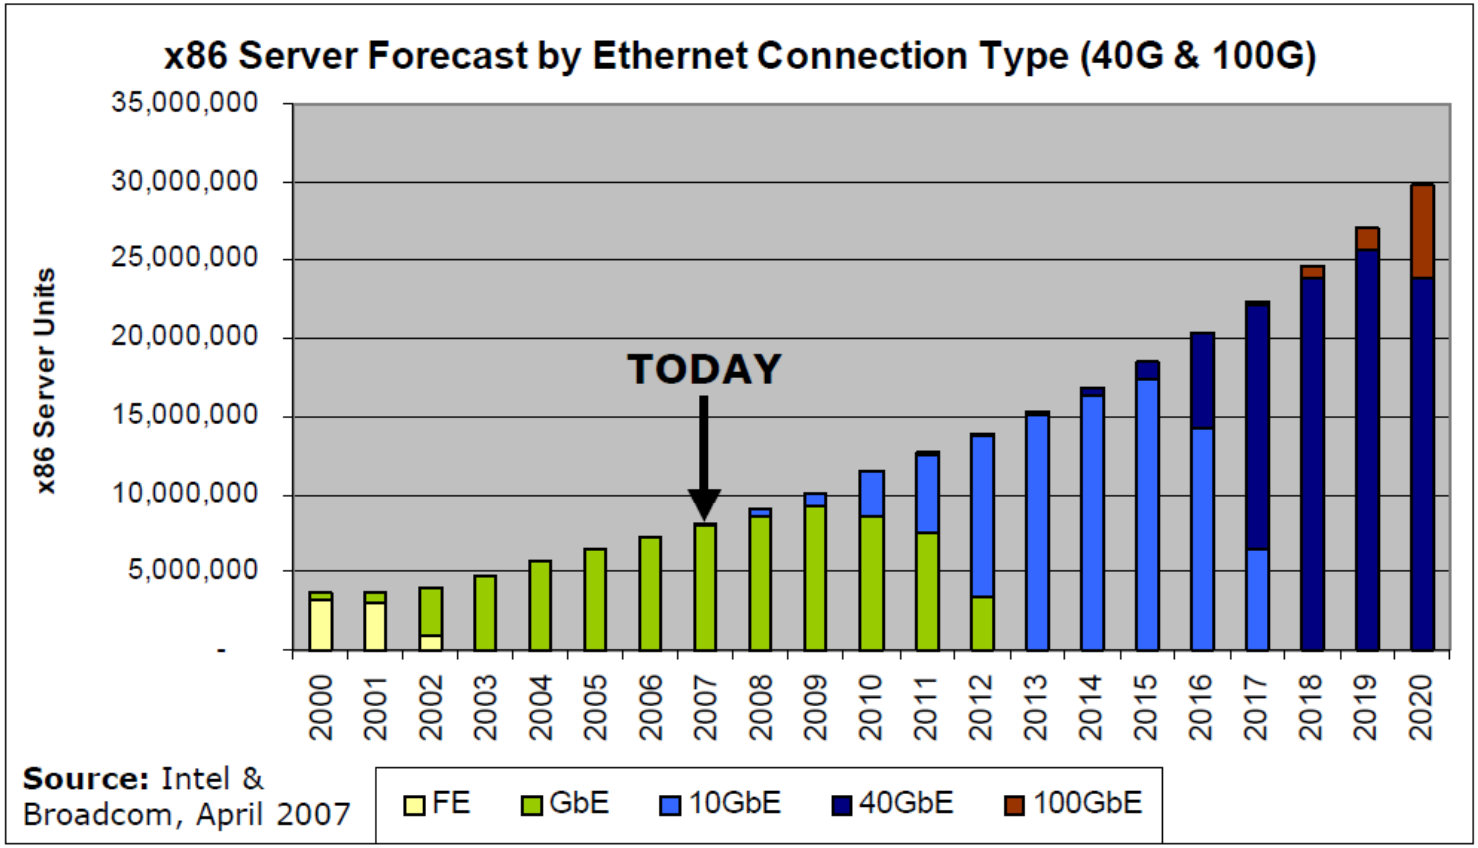

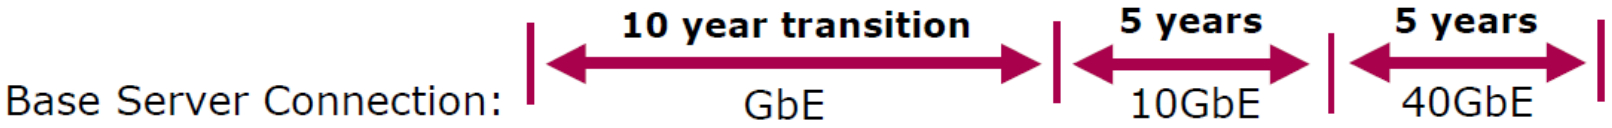

IEEE 802.3 HSSG April 2007 Interim Meeting

x86 servers by Ethernet connection speed (2010 forecast)

From the 100G Cu Backplane & Twinax CFI

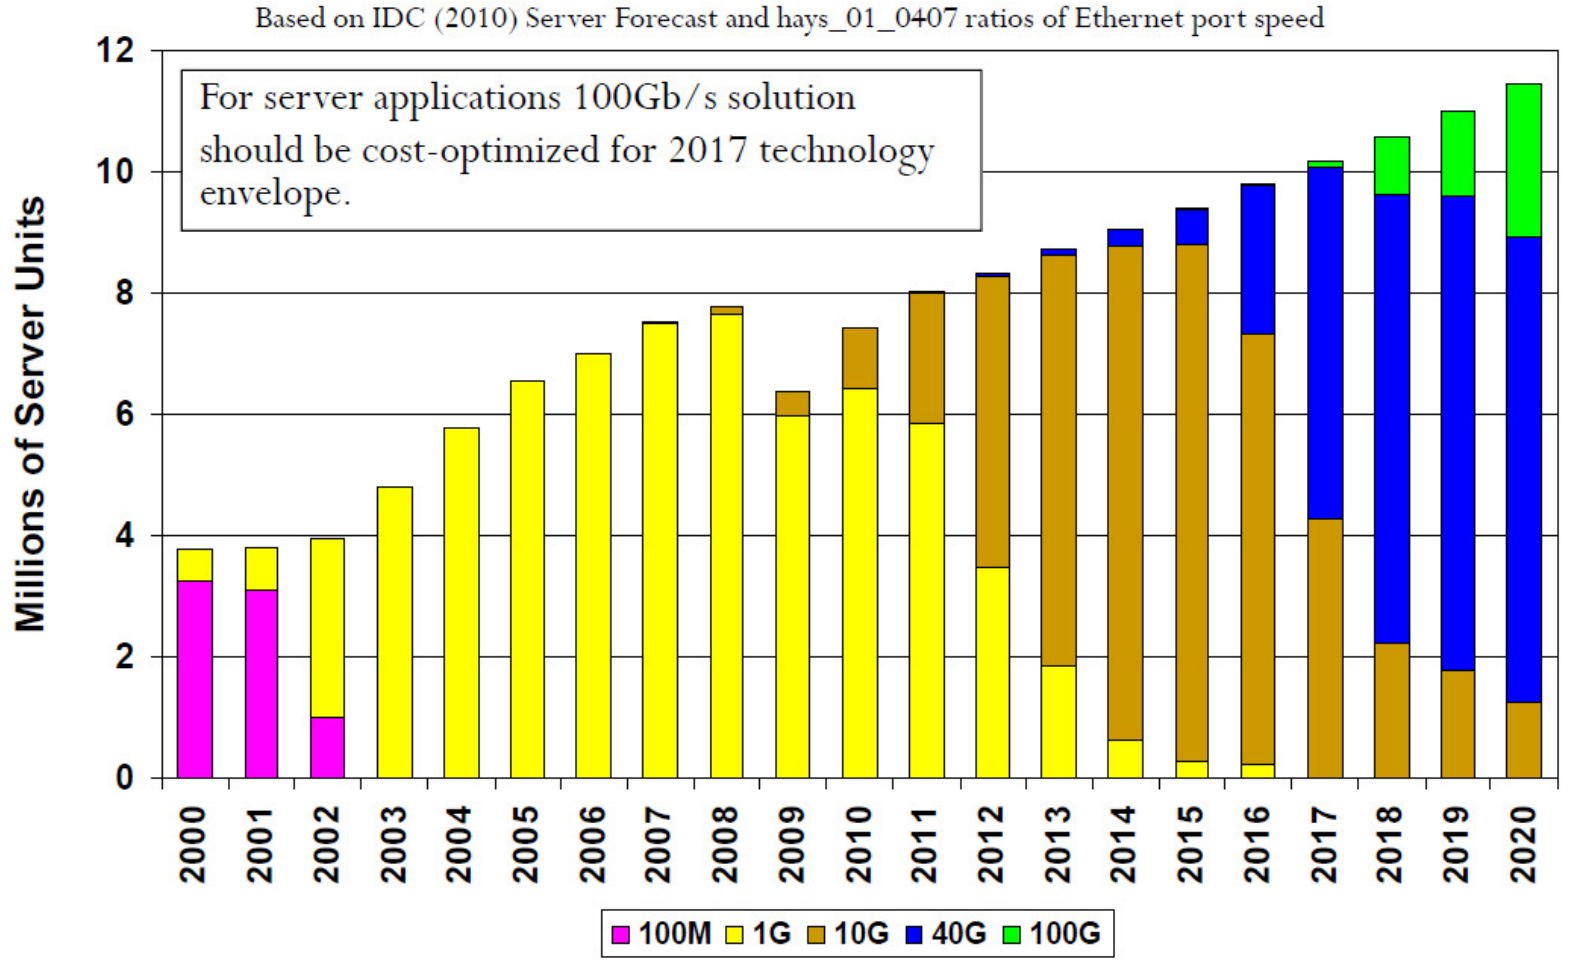

x86 Servers by Ethernet Connection Speed (2012 Forecast)

Based on IDC, Dell Oro, Crehan Research and Intel data from 2H'11 – 1Q'12

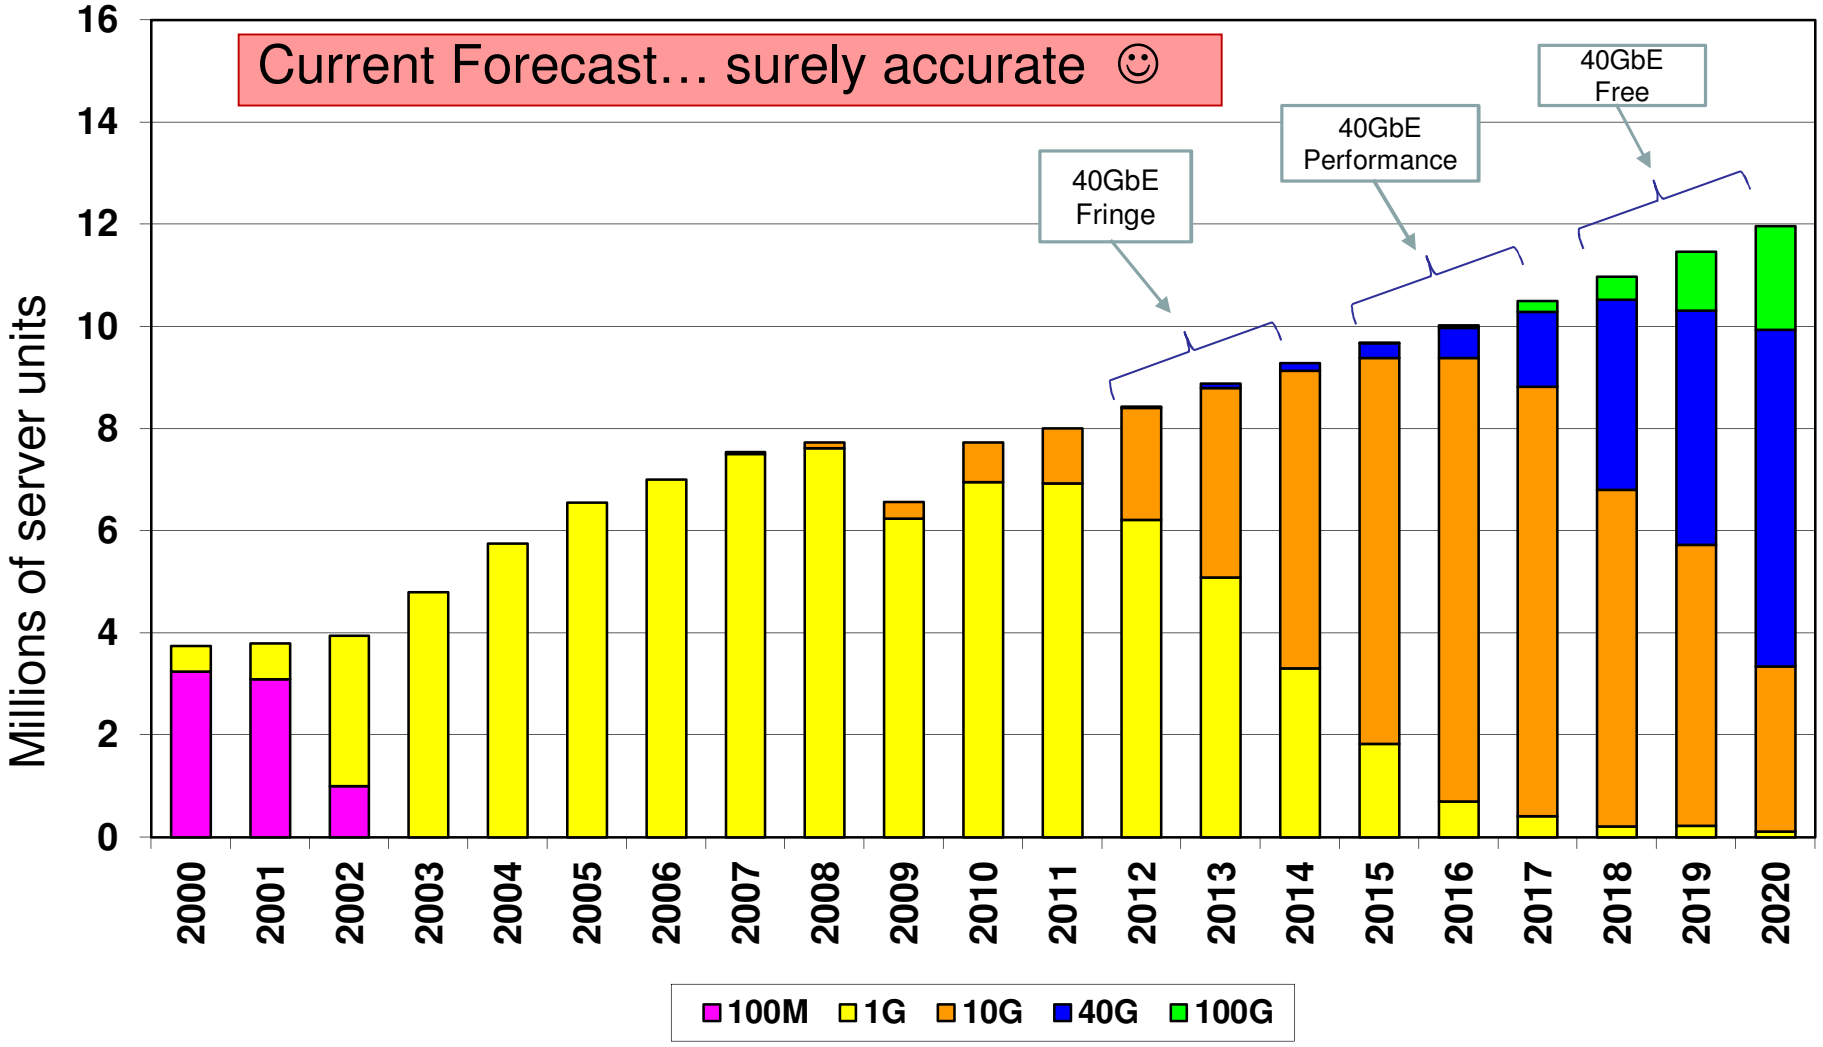

Observations

Facts

- Server units growth slowed from 2007 forecast (10% -> 5%)
  - 2009 recession hit, then ramping again
- Ethernet port speed transitions slower by percentage
  - 10GbE ports grew, but so did 1GbE ports

Conjecture on why

- Recession slowed development, qualification of new technology, and capital expenditure
- Server platform launch delays
  - New technology goes with the new platform
- The volume market did not get the cheap 10GBASE-T it was promised
  - We made it too hard

More Observations

- Total Server BW shipped increased from 2010 forecast
  - 10GbE grew, but
  - 1000BASE-T shipments far exceeded earlier forecast

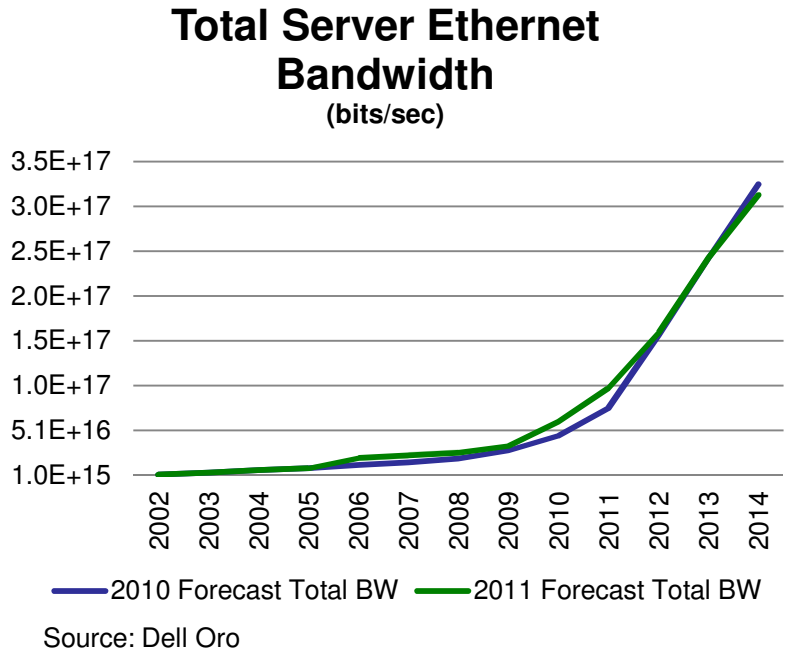

Source: Dell Oro

- Average Ports/Server grew to >4 in 2010 and expected to maintain

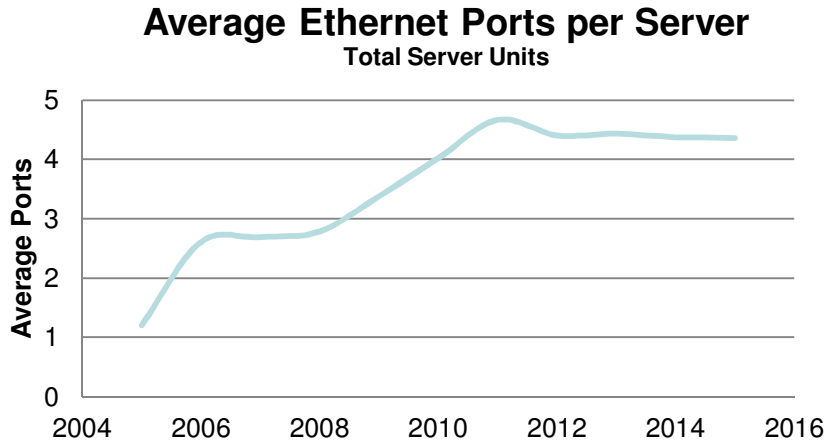

Server bandwidth trends

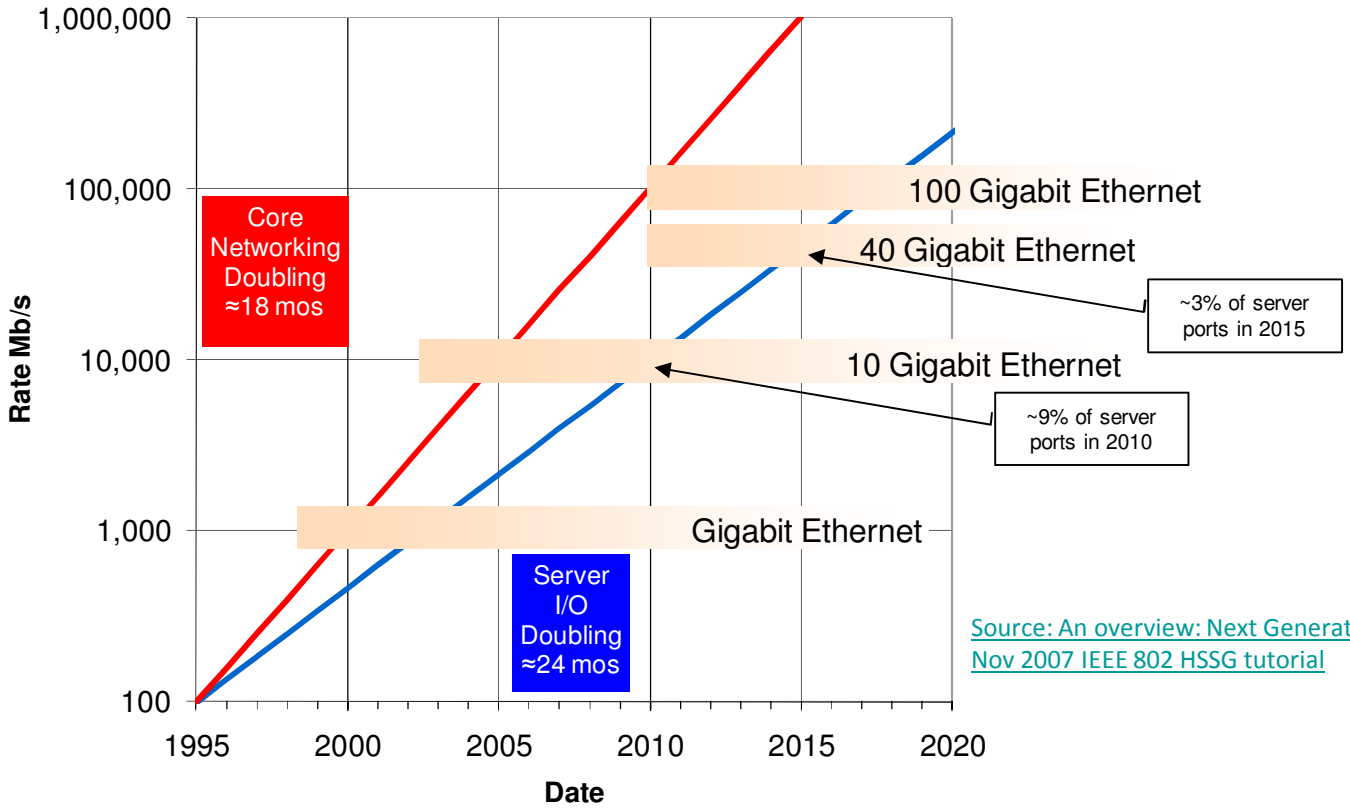

Source: An overview: Next Generation of Ethernet – Nov 2007 IEEE 802 HSSG tutorial

Typical Development cycles

- Standards: 2-3+ years
- System and silicon: 2-3+ years
- Typically ~5 years from start of 802.3 project to first products
  - Varies widely based on many factors

Server market has a wide variety of data rate needs.

- **Technology can persist for many years after first introduction.**

Market Need Summary

Status

- Total Server bandwidth capability and deployment continues to increase
- 10GBASE-T ramping with LOM
- “Fringe” deployment of 40GbE starting now

Need

- 10GBASE-T market needs an upgrade path
- Higher speed BASE-T Study Group now will help the Server market move from Fringe to wider adoption

Technical Viability

The Channel

- ISO/IEC and TIA are actively studying the problem
- Both define that sufficient capacity exists for 40GbE over four twisted pairs
- Both await IEEE 802.3's input to define length and to finalize impairment levels
  - IEEE Study group will need to balance objectives related to complexity, power and cost

Body of Work

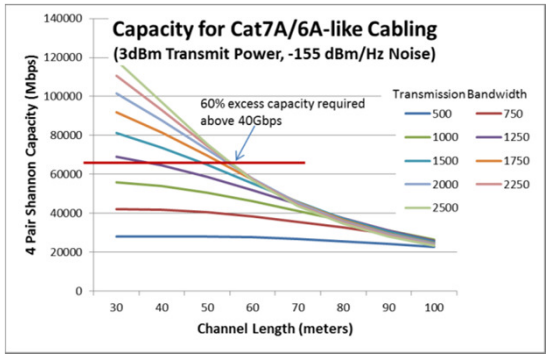

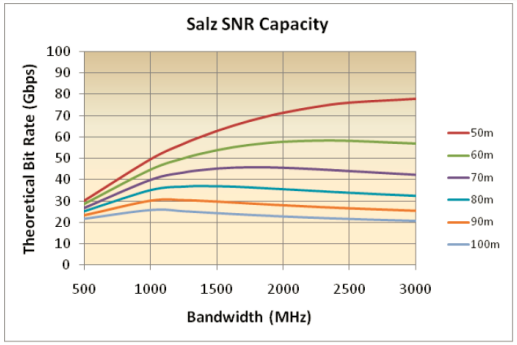

Figure 6. Salz SNR Capacity of Category 40G Channels versus Bandwidth

Class I fmax_system = 1.6 CN = 40 CF = 40 CR = 60 Tp = 3 BNp(1) = 63

L := 50

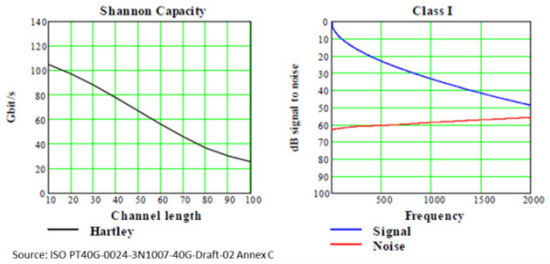

Source: ISO PT40G-0024-3N1007-40G-Draft-02 Annex C

Single-Carrier Data-Rate vs. Frequency

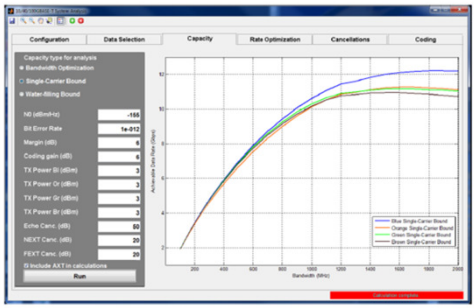

Source: Nexans ISO Contribution PT40G-018A – Performance of Cat7A

10

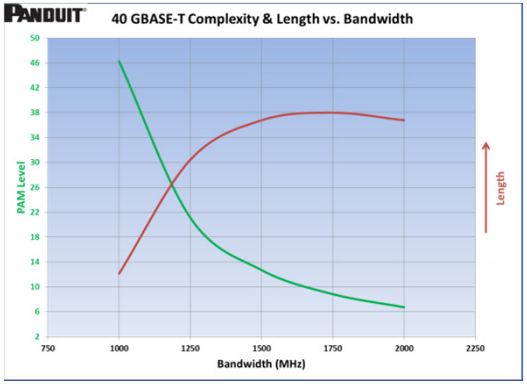

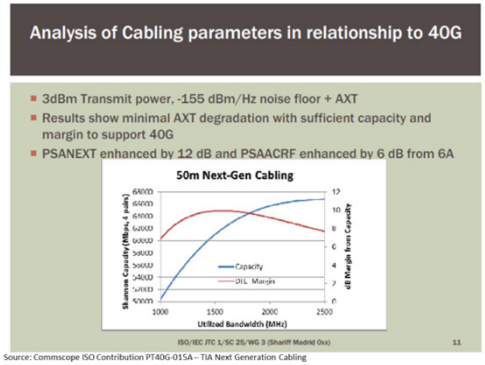

Source: Commscope ISO Contribution PT40G-015A – TIA Next Generation Cabling

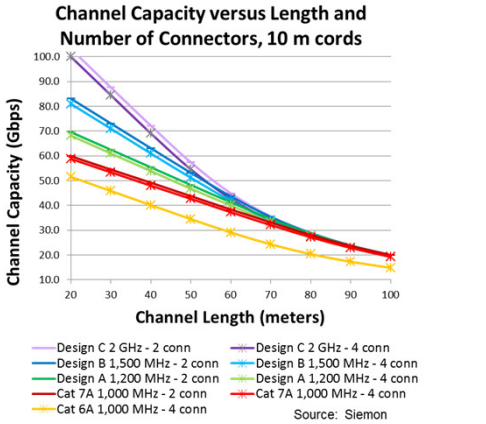

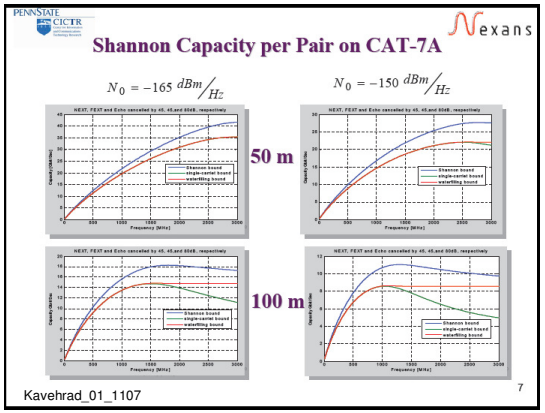

"Sources: Various contributions & published whitepapers: IEEE 802.3 HSSG, TIA TR42.7, ISO/IEC 11801 and whitepapers (2007-2012)"

The PHY

- One option is to maintain baseband modulation (like 10GBASE-T)
  - A combination of increased bandwidth and/or increased SNR is required
- Individuals from multiple semiconductor companies concur that NGBASE-T baseband circuits for 40GbE are viable in the time frame of this standard

A Study Group is needed to investigate!

Historic Power Trends

- Power will be a critical factor for 40GBASE-T
- History shows that BASE-T solutions have been Moore's Law friendly, showing significant power reduction with advancing silicon process.
- No reason to expect 40GBASE-T will not follow that trend.

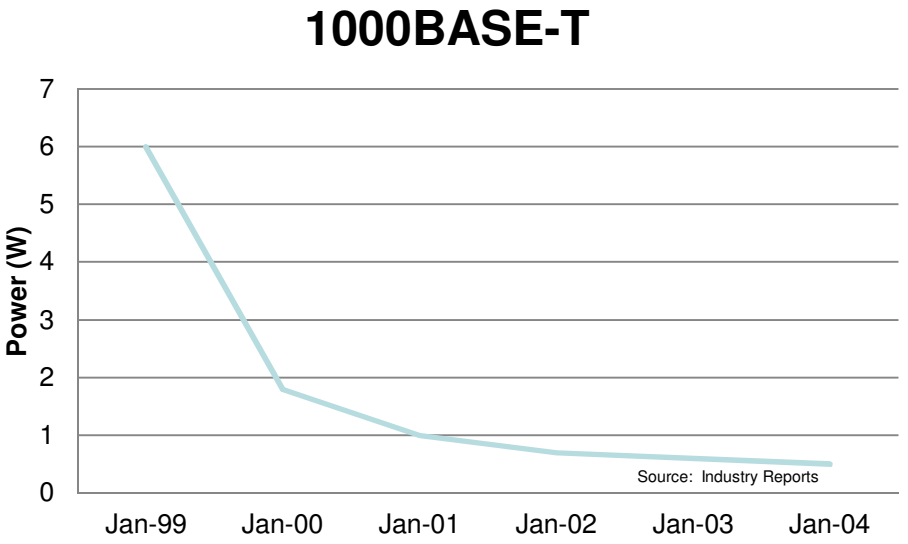

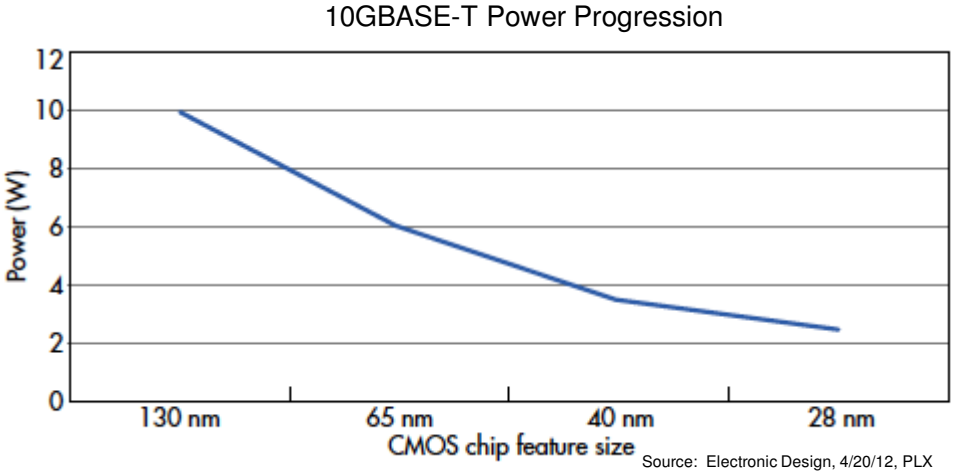

Attributes

Next Generation BASE-T provides capabilities not otherwise available in the portfolio

	Next Gen BASE-T	**40GBASE-CR4**
Backwards Compatibility / Autonegotiation	Autoneg will be integral, compatible with 10GBASE-T & 1000BASE-T	Breakout can interface with SFP+ or SFP, no Autonegotiation
Latency	Non-zero, for Study Group	Low
Density in 1RU	2Tb+, 48 ports plus uplinks (assuming RJ45 size)	1.28Tb+, 32 ports plus uplinks (assuming QSFP)
Cabling	4 twisted pair	8 pair, twinax
Cost factors	No active elements, assume twisted pair, widely sourced	QSFP end point with finite cost and assembly complexity, some vendors restrict sourcing
Reach	Can do End of Row, 30m or more, defined by Study Group	7m, then can use active cable
Integration	LoM'able, compatible with silicon integration	CR4 well suited for integration into CMOS ASICs

Data Center Topologies

- Next Gen BASE-T well suited to cover Server to Switch connections within the row

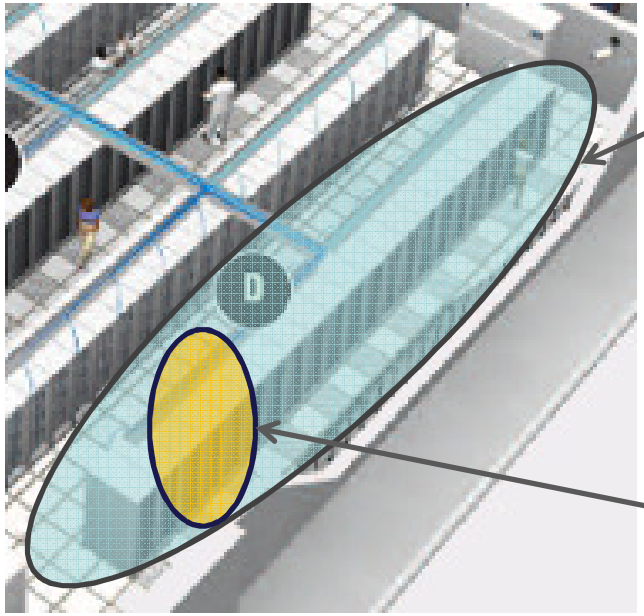

Distance served by NGBASE-T
- Within the rack
- Neighboring racks, stranded ports
- End of row

Distance served by CR4
- Within the rack
- Neighboring racks

Study Group Topics

- Cable reach
  - Understanding reach vs. power trade-off
- Channel objectives
  - Cable, connectors, magnetics
- Data Rate
- Coding / modulation
- EEE

Questions and Discussion

Next Gen BASE-T CFI Straw Poll

- Should an 802.3 Study Group be formed for

Next Generation BASE-T?

People in the Room

Y: _____

N: _____

A: _____

Dot 3 Voters Only

Y: _____

N: _____

A: _____

Straw Polls

- __ Number of people in the room
- __ Individuals who would attend and contribute to a

  Next Generation BASE-T Study Group
- __ Companies that support participation in a

  Next Generation BASE-T Study Group

Thank you!