

Ethernet Alliance



Ethernet 103: Introduction to 25 Gb/s Ethernet

Scott Kipp, Brocade
Mark Nowell, Cisco
John D'Ambrosia, Dell



Disclaimer

- Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance.



THE ETHERNET ALLIANCE

**A global community of end users, system vendors,
component suppliers and academia**

- **Our Mission**

- To promote industry awareness, acceptance and advancement of technology and products based on, or dependent upon, both existing and emerging IEEE 802 Ethernet standards and their management.
- To accelerate industry adoption and remove barriers to market entry by providing a cohesive, market responsive, industry voice.
- Provide resources to establish and demonstrate multi-vendor interoperability.

- **Activities**

- Promote marketing and education awareness
- Interoperability testing and demonstration
- Industry consensus building
- Technology and standards incubation



ETHERNET ALLIANCE STRATEGY

Expand Ethernet Ecosystem

- Facilitate interop testing
- Expand the market
- Go global

Support Ethernet Development

- Support consensus building
- Host Technology Exploration Forums (TEFs)
- Team with other organizations

Promote Ethernet

Marketing

Education



UNIVERSITY OF ETHERNET WEBINAR SERIES

Ethernet 101: Introduction to Ethernet

Physical Layer

Ethernet 102:
The Physical
Layer Of Ethernet

This Webinar



Ethernet 202:
10GBASE-T
Revamped

Ethernet 103:
Introduction to 25
Gb/s Ethernet

Ethernet 301:
40/100GbE Fiber
Cabling and
Migration Practices

Ethernet 203:
40G & 100G
Ethernet

Protocols

Ethernet 211:
Data Center
Convergence

Ethernet 311:
Multi-tenancy

Ethernet Topics

Ethernet 151:
Ethernet
Alliance Plugfests

All Webinars archived at: <http://www.ethernetalliance.org/library/university-of-ethernet-webinar-series/>

Today's Speakers



Scott G. Kipp
Principal Technologist
Brocade



Mark Nowell
Sr. Director, Engineering
Cisco
Chair, IEEE 802.3 25GbE
Project



John D'Ambrosia
Chief Ethernet Evangelist
CTO Office, Networking
Dell

Agenda

- Introduction
- 25 GbE Technology
- Deployment Architectures
- 25 GbE Standards
- Market Forecast
- In Closing
- Q&A

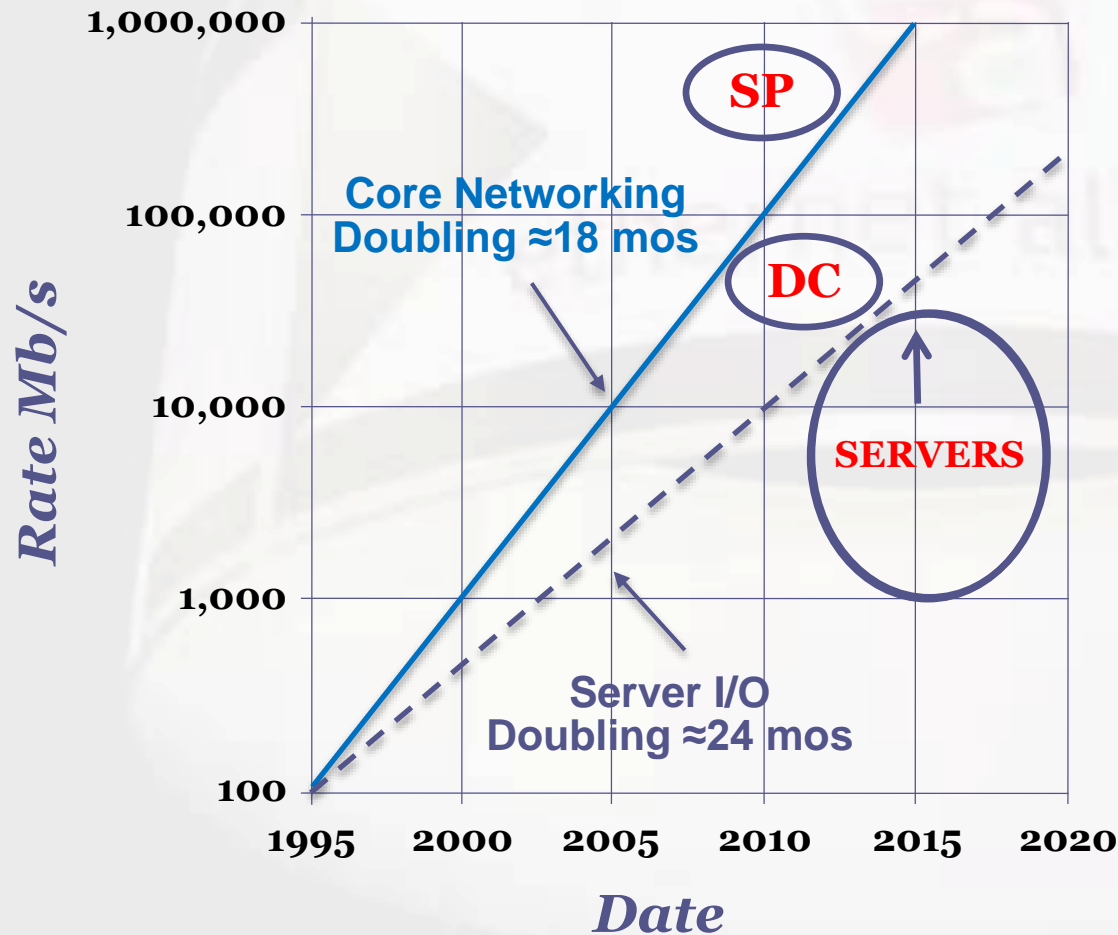




Introduction



Bandwidth Diversity Trends

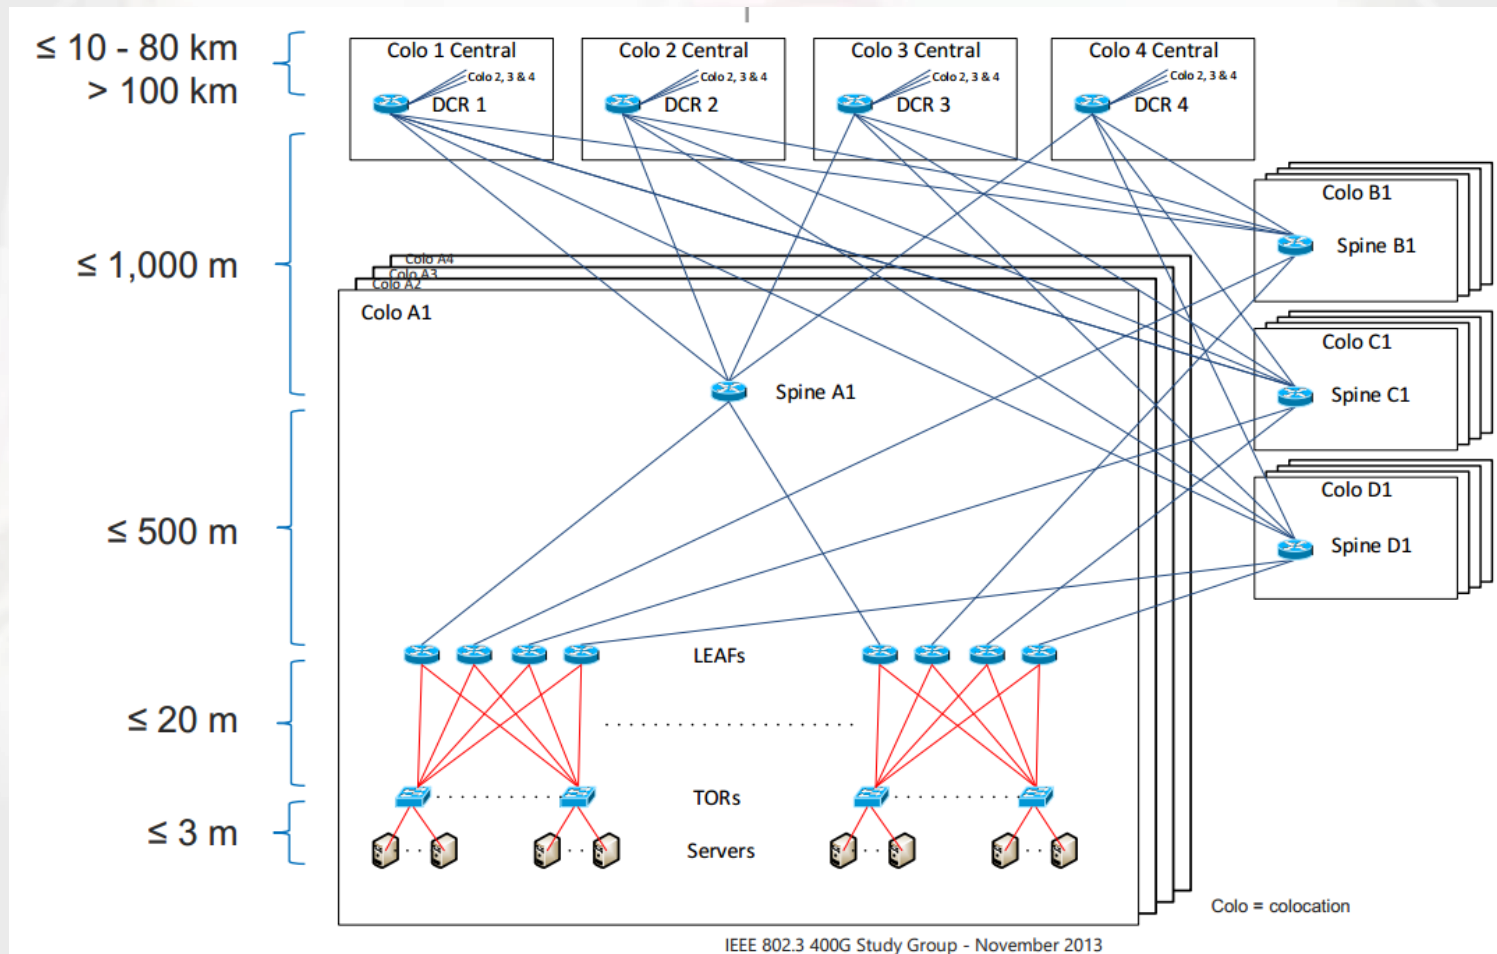


- In 2007 it was projected that servers would require 40GbE by today
- In 2014, majority of servers shipping based on GbE
- In 2014, cloud scale data centers looking beyond 10GbE

Based on IEEE 802.3 2007 HSSG Tutorial
http://www.ieee802.org/3/hssg/public/nov07/HSSG_Tutorial_1107.zip



Microsoft Data Center Architecture



Source: http://www.ieee802.org/3/400GSG/public/13_11/booth_400_01a_1113.pdf

Link Volumes for Data Center

Interconnection Volume

- Four sections per colo & multiple colos (≥ 4) per data center
- Volumes below are per section (except DCR to Metro)

A End	Z End	Volume	Reach (max)	Medium	Cost Sensitivity	Market Space
Server ‡	TOR	10k – 100k	3 m	Copper	Extreme	LAN
TOR	LEAF	1k – 10k	20 m	Fiber (AOC)	High	
LEAF	SPINE	1k – 10k	400 m	SMF	High	
SPINE	DCR	100 – 1000	1,000 m	SMF	Medium	Campus
DCR	Metro	100 – 300	10 - 80 km	SMF	Low	WAN

‡ Server-TOR links may be served by breakout cables

IEEE 802.3 400G Study Group - November 2013

Source: http://www.ieee802.org/3/400GSG/public/13_11/booth_400_01a_1113.pdf



Benefits of 25GbE for Server IO

- Maximize Switch Throughput & Efficiency
- Minimize CAPEX
 - Minimize # of ToR switches
 - Minimize # of Cables
- Minimize OPEX
 - Less Power
 - Less Cooling
- Minimize Cost per Bit
- Technology Exists today!

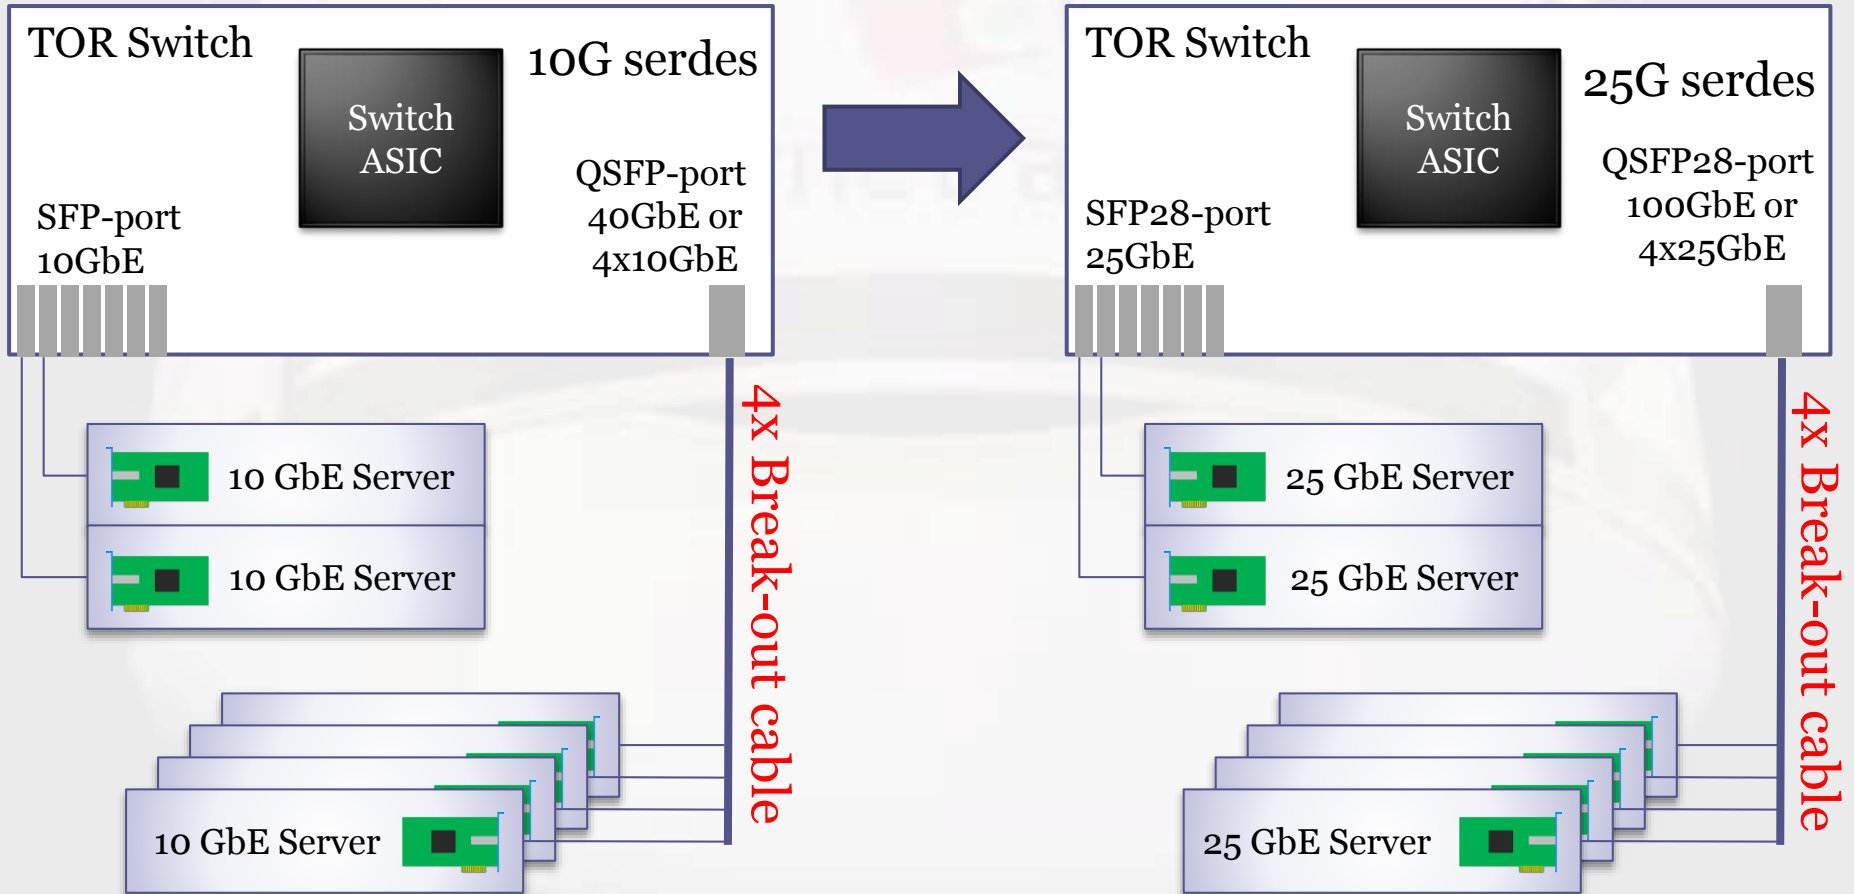




25GbE Technology



From 10GbE to 25GbE server IO



25Gb/s Related IEEE 802.3 Standards

Technology	Nomenclature	Description
Chip-to-Chip	CAUI-4	4 x 25 Gb/s
Chip-to-Module	CAUI-4	4 x 25 Gb/s
Backplanes	100GBASE-KR4	4 x 25 Gb/s (NRZ)
	100GBASE-KP4	4 x 25 Gb/s (PAM-4)
Cu Twin-Axial	100GBASE-CR4	4 x 25 Gb/s
MMF Optics	100GBASE-SR4	4 x 25 Gb/s
SMF Optics (10km)	100GBASE-LR4	4 x 25 Gb/s
SMF Optics (40km)	100GBASE-ER4	4 x 25 Gb/s

25Gb/s Capable Form Factors





Image courtesy of TE Connectivity


Switch ASIC I/O are Valuable!

- Today's I/O's can run from 10Gb/s to 25 Gb/s
- Switch ASIC connectivity limited by I/O pins
- Single Lane port maximizes server connectivity
- Single lane 25Gb/s I/O maximizes both port count and total bandwidth



Port Speed	Lane Speed (Gb/s)	Lanes /Port	Usable Ports	Total BW (Gbps)
10 GbE	10	1	128	1280
25 GbE	25	1	128	3200
40 GbE	10	4	32	1280
100 GbE	25	4	32	3200

Building a Cloud Scale Data Center



3.2 Tb/s Switch Based on 32 QSFP Ports					# TORs for a 100K Server Data Center
Server I/O	Servers	100 GbE Uplinks	Throughput (Tb/s) per ToR Switch	Utilization (%)	
40GbE (4x10G)	28	4	1.52	47.5	3572
25GbE (Single Lane)	96	8	3.2	100	1042




Deployment Architectures



Typical Architectures

Top of Rack
(ToR)




Intra-rack – 3m
Adjacent-rack – 5m

Middle of Row
(MoR)



3 to 15 m


End of Row
(EoR)



3 to 50 m

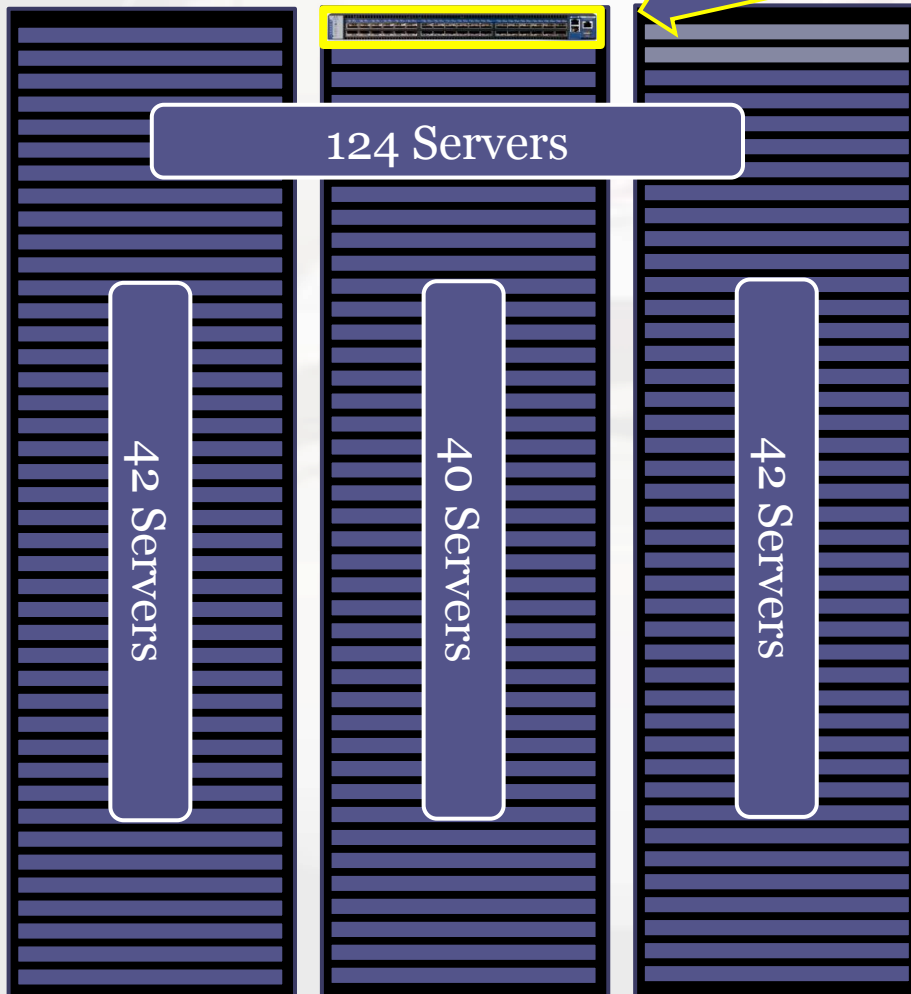
Typical interconnect distances

Source of graphics: jimenez_3bq_01_0711



TOR/Adjacent Rack Example

42 RU Racks



E.g.. 36 QSFP28 = 144 25GbE

124 Servers

5 QSFP28 ports for
500Gb/s of uplinks

One ToR switch can
support multiple
racks of 1U servers

1U Server EoR/MoR Design

2,400 Servers
in 57 racks

42 Servers

42 Servers

42 Servers


42 Servers

42 Servers

42 Servers

42 Servers

42 RU Rack



Modular switches
can have hundreds
of 25GbE Ports

This example has
864 25GbE
Ports/Switch

This rack of 3
modular switches
could support 2,400
servers with 192
100GbE ports
(19.2Tb/s) of
uplinks

What Servers Need from 25GbE?



Top of
Rack

Adjacent
Rack

Middle
of Row

End of
Row

3m Cu

5m Cu

Active Cables or MMF (100m)

25GBASE-T (30m)

- No one solution for “server to switch” interconnect
 - Backplane
 - Cu Cables
 - Active Cables
 - Multi-mode fiber
 - Future - 25GBASE-T
- “Right” media dictated by customers architectural choices



25GbE Standards



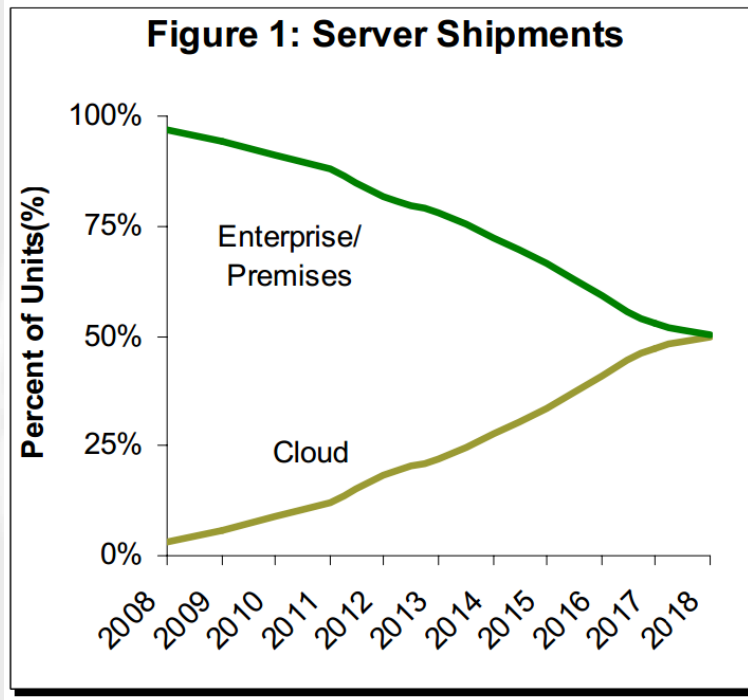
IEEE 802.3 25GbE Project

- IEEE 802.3 25GbE Study Group
 - <http://www.ieee802.org/3/25GSG/index.html>
- Reach objectives include:
 - Backplanes consistent with 802.3bj (100GbE) backplanes
 - Twinax support to 3 meters
 - Twinax support to 5 meters
 - MMF consistent with 802.3bm (100GBASE-SR4)
- IEEE P802.3by Task Force approved Dec 10.



IEEE 802.3 25GBASE-T Project

- Nov 14 - Successful Call-for-Interest
 - http://www.ieee802.org/3/cfi/1114_2/CFI_02_1114.pdf



Source: Dell'Oro Controller and Adapter Report, July 2014

- Targeting Enterprise-class data centers
- Targeting 0 to 30m
- Support ToR, MoR, and EoR architectures

Schedule for 25GbE

- **IEEE P802.3by Task Force**
 - 1st Task Force meeting Jan 2015
 - Estimated Ratification - 2016
 - Pre-standard product availability anticipated
- **IEEE 25GBASE-T Study Group**
 - 1st Study Group meeting Jan 2015
 - Similarities to 40GBASE-T anticipated
 - Estimated Ratification - Q2, 2016
 - (40GBASE-T schedule)







Market Forecast



Speed Migration for Cloud Servers



Speed Migration for Enterprise Servers





In Closing



The Case for 25GbE

- **Cost optimized servers beyond 10G needed**
- **25 Gb/s Ethernet**
 - **Leverages single-lane 25Gb/s signaling technology developed to support 100GbE**
 - **Maximize efficiency of server to access switch interconnect**
- **Web-scale data centers and cloud based services are leading applications**



Conclusions

- Compared to 40GbE today, 25GbE will deliver:
 - Lower cost
 - Higher density
 - Lower power architectures/bit
- No architectural changes necessary to move to 25GbE
- Rapid development and pre-standard 25GbE products coming soon!







Q&A



Thank You!

- If you have any questions or comments, please email admin@ethernetalliance.org
- Ethernet Alliance: visit www.ethernetalliance.org;
 -  Follow @EthernetAllianc on Twitter
 -  Join the Ethernet Alliance [LinkedIn group](#)
 -  Visit the Ethernet Alliance [Facebook page](#)

