

NGBASE-T Considerations: Reach and Migration Path

Brad Booth, Dell

David Koenen, HP

Supporters

- Robert Winter, Dell
- Rich Hernandez, Dell
- Jeff Lynch, IBM
- Mike Bennett, LBL
- Jacky Chang, HP
- Dave Chalupsky, Intel
- Dan Dove, AppliedMicro
- Kent Lusted, Intel
- Pravin Patel, IBM
- George Zimmerman, CME Consulting (Commscope)
- Bruce Tolley, Solarflare
- Mark Nowell, Cisco
- Steve Carlson, High Speed Design
- John D'Ambrosia, Dell
- Wael Diab, Broadcom
- Nathan Tracy, TE Connectivity
- Bernie Hammond, TE Connectivity
- Alan Flatman, LAN Technologies

Legacy BASE-T

- Migration
 - In the 90's, 10M to 10/100 devices
 - In the 00's, 10/100 to 10/100/1000 devices
 - In the 10's, 10/100/1000 to 100/1G/10G
 - 10M dropped
- Reach of 100 m
 - Based on horizontal cabling specifications
 - Cat 3 to 5 to 5e to 6A, majority using RJ45*

IEEE 802.3 NGBASE-T Interim Meeting -
Geneva, CH - Sept. 2012
* Patch panels and non-DTE cabling may use other connectors.

What's Changed

- During P802.3an, new topologies were occurring
 - Flatter networks, de-centralized switching
 - Virtualization, convergence
- Server-to-switch links
 - ToR deployment growing
 - Doesn't require a 100m cabling infrastructure
 - High density servers demand lowest-cost and lowest-power interconnect

Open Compute Example



- Reach to ToR switch is $< 5\text{m}$
 - Zero connector channel (i.e. point-to-point)

What's Changing (cont.)

- Pre-Assembled PODs
 - 10s of racks with 100s of servers pre-wired to end-of-row (EoR) switches in shipping container (POD)
 - MTO/factory terminated cable lengths up to 15m*
 - Mix of 1G, 10G & NGBASE-T rates
- Scale Out DC Deployments
 - Rapid, low cost, & easy deployment of multiple racks
 - Servers to EoR chassis switch cable topologies
 - Up to 3 cable segments with reach of 25m

Pre-Assembled POD Examples



Analyst Data

- Dell'Oro July 2012
 - “we do not anticipate 10G Base-T to out ship SFP+ direct attach copper until the end of 2015...”

Figure 9: 10 GE EN SW Port Shipments (in 000's)

	2010	2011	2012	2013	2014	2015	2016
Copper	1,620	3,936	9,130	16,336	25,284	34,154	43,694
Fiber	3,495	5,142	5,977	6,770	7,853	9,279	10,668
Total	5,115	9,078	15,106	23,106	33,137	43,432	54,362
% Copper	32%	43%	60%	71%	76%	79%	80%

- “For 2011... the majority of copper port shipments are SFP+ direct attach copper... remaining port shipments are mainly KR ports...”

Next Gen BASE-T

- Server-to-switch
 - Connectivity is typically less than 25m*
 - Should not bear the burden of cost & power for 100m
 - Needs a migration path
 - 1G → 10G → 40G for example; multi-speed PHYs
 - Wake-on-LAN supported on copper links
- Beyond server-to-switch
 - Switch uplinks using optical interconnect
 - Use of copper not what it once was

* Equal to 90% of all server-to-switch link lengths as shown in flatman_01_0311 as presented to IEEE P802.3bj Task Force - Sept. 2012

Recommendations

- Learn the lessons from 10GBASE-T
 - Customers care less about reach than they do about cost and power
 - But only reach goes into the standard
 - Network topologies have changed → embrace it!
- Migration path
 - Task Force needs to take this into consideration
 - Cabling is important part of equation

Considerations/Discussion Points

- Is there an inflection point on reach vs. power?
 - In-the-rack vs. in-the-row reach spec
 - PHY with two operating modes
- Enhance 10GBASE-T spec?
 - Service to humanity
 - Better migration strategy
- EEE in a virtualized/converged network?