

50 Gb/s Ethernet CFI considerations

9/16/15

Mark Nowell - Cisco

So far...

Interest in exploring Pre-CFI for 50Gb/s Ethernet first declared in spring 2015.

First presentation @ ad hoc meeting held at May 2015 IEEE Interim to gauge interest.

- Well attended.
- Email distribution list established (currently >140 members)

Hawaii Plenary was next opportunity to gather input. Multiple presentations.

- John D'Ambrosia – Dell – 200Gb/s Ethernet
- Yong Kim – Broadcom - 50G Project considerations
- Kohichi Tamura - Oclaro – Standard for 25G SMF
- Brad Booth – Microsoft – Market needs
- Mark Nowell – Cisco - Discussion

To summarize the key takeaway's presented and expressed

- The scope and scale of next IEEE 802.3 project(s) was the dominant discussion
- Interest in 25GE SMF but feedback that it should be standalone project
- CFI to include 50Gb/s Ethernet well supported
- Strong interest in establishing an effort with Nx50G variants (100G & 200G)
- Considerable discussion on how to merge or partition the work to effectively align with industry needs

Work areas of interest (as well as what's already there)

		Existing Rates									New Rates	
Media		10G	25G	40G	40G (G2)	100G (G1)	100G (G2)	100G (Gn)	400G (G1)	400G (G2)	50G	200G
PCB Traces		1x10G	1X25G	4X10G		10x10G	4X25G	2x50G 1x100G ?	16X25G 8x50G		1x50G	4x50G
BP		1x10G	1X25G	4x10G			4X25G				1x50G	4x50G
Cu Cable		1x10G	1X25G	4x10G		10x10G	4X25G				1x50G	4x50G
MMF		1x10G	1X25G	4x10G	1X40G	10x10G	4X25G	2X50G 1x100G	16X25G		1x50G	4x50G
SMF	500m		1X25G		1X40G			2X50G 1x100G	4x100G (PAM4)		1x50G	4x50G 2x100G
	2km		1X25G	1x40G	1X40G			2X50G 1x100G	8x50G WDM (PAM4)		1x50G	4x50G 2x100G
	10km	1x10G	1X25G	4x10G WDM			4x25G WDM		8x50G WDM (PAM4)			
	40km	1x10G		4x10G WDM			4x25G WDM					

50 GbE: Application areas

Media		Existing Rates								New Rates		
		10G	25G	40G	40G (G2)	100G (G1)	100G (G2)	100G (Gn)	400G (G1)	400G (G2)	50G	200G
PCB Traces		1x10G	1X25G	4X10G		10x10G			5X25G 1x50G		1x50G	
		1x10G	1X25G	4x10G				4X25G			1x50G	
Cu Cable		1x10G	1X25G	4x10G		10x10G					1x50G	
MMF		1x10G	1X25G	4x10G		10x10G		4X25G	16X25G		1x50G	
SMF	500m								100G (PAM4)		1x50G	
	2km			1x40G					TBD		1x50G	
	10km	1x10G		4x10G WDM			4x25G WDM		8x50G WDM (PAM4)			
	40km	1x10G		4x10G WDM			4x25G WDM					

C2C C2M AUI

Ethernet server IO

Ethernet Switch IO



IEEE 802.3 Ethernet Technology Overview: 50 GbE

Media		Existing Rates									New Rates	
		10G	25G	40G	40G (G2)	100G (G1)	100G (G2)	100G (Gn)	400G (G1)	400G (G2)	50G	200G
PCB Traces		1x10G	1X25G	4X10G		10x10G	4X25G	2x50G 1x100G ?	16X25G 8x50G		1x50G	
BP		1x10G	1X25G	4x10G		<div style="background-color: green; color: white; padding: 10px; border-radius: 15px; display: inline-block;">New work</div>					1x50G	
Cu Cable		1x10G	1X25G	4x10G							1x50G	
MMF		1x10G	1X25G	4x10G		10x10G	4X25G		16X25G		1x50G	
SMF	500m										1x50G	
	2km			1x40G							1x50G	
	10km	1x10G		4x10G WDM			4x25G WDM			8x50G WDM (PAM4)		
	40km	1x10G		4x10G WDM			4x25G WDM					

New work

Strong leverage of 802.3bs work



Other application areas discussed

		Existing Rates								New Rates		
Media		10G	25G	40G	40G (G2)	100G (G1)	100G (G2)	100G (Gn)	400G (G1)	400G (G2)	50G	200G
PCB Traces		1x10G	1X25G	4X10G		10x10G	4X25G				1x50G	4x50G
BP		1x10G	1X25G	4x10G			4X25G				1x50G	4x50G
Cu Cable		1x10G	1X25G	4x10G		10x10G	4X25G				1x50G	4x50G
MMF		1x10G	1X25G	4x10G	1X40G	10x10G	4X25G	2X50G 1x100G			1x50G	4x50G
Enterprise Campus	500m		1X25G		1X40G	Low-cost Ethernet PMDs		2X50G 1x100G			1x50G	4x50G 2x100G
			1X25G	1x40G	1X40G			2X50G 1x100G			1x50G	1x50G 2x100G
	10km	1x10G	1X25G	4x10G WDM		4x25G WDM			8x50G WDM (PAM4)			
	40km	1x10G		4x10G WDM		4x25G WDM						

C2C C2M AUI

NG BP upgrade

Ethernet server

Ethernet ToR uplink

Low-cost Ethernet PMDs

Enterprise Campus

The Case for 50 Gb/s Ethernet Single Lane

Single lane architecture proven to provide lowest cost solution when homogeneous throughout system

- 50 Gb/s ASIC IO technology under development (and standards specification under development in IEEE and OIF)
- Aligns with ongoing Network bandwidth demand growth, cost pressures, etc...

Tier 1 Datacenters looking to drive server adoption to 50 Gb/s IO. Early adopters already planning to deploy very soon (based on 2x25Gb/s)

Close technology overlap with 40 Gb/s Ethernet single lane

- Work previously done on assessing market need/interest
- Mid-2014 series of mtgs <http://www.ethernetalliance.org/library/presentations-2/>
- Cost of Quad lane PMDs @ 40Gb/s >> 4x single lane PMDs @ 10G
- Increased Broad Market potential for project
 - Server IO, Backplane, DC Switch Fabric IO

IEEE target of 2018 completion to align with industry deployments

The challenges ahead for a 50 Gb/s Ethernet project

Strong ability to leverage ongoing work in IEEE – 802.3bs

- SMF, CDAUI...

New work in PMDs

- 50G MMF, 50G Twinax, 50G Backplane

Architecture and interop open issues

- FEC choice? Continue on single lane approach from 802.3by (KR4 based) vs. leverage of 802.3bs (KP4 based with pre-interleaving)
- Interop with existing 2x25G MAC architectures that will exist by 2018
- 40Gb/s & FEC

Predict some strong areas of immediate consensus, some areas of necessary consensus building. 2018 completion will be a challenge.

The Case for Nx50 Gb/s Ethernet

Quad implementation has strong track record of market adoption

- 50 Gb/s ASIC IO under development (and standards specification in IEEE and OIF)

Switch Uplink IO to support 50Gb/s server IO

- Usually MMF, AOC or SMF

Backplane upgrade from existing 4x25G to 4x50G

CPUs rising to 200 Gb/s throughput

- Twinax cable requirement

The challenges ahead for a Nx50 Gb/s Ethernet project

Strong ability to leverage ongoing work in IEEE – 802.3bs

- SMF, CDAUI...

New work in PMDs – 100G (2x50G) & 200G (4x50G)

- Twinax, Backplane, MMF, SMF
- Argument that SMF not required is inconsistent with project justification

Architecture and interop open issues

- FEC & PCS architecture. 802.3bs was 1st project with multi-lane PCS and FEC.
- Need to define for both 100G & 200G. PCS/FEC architecture, as always, needs to consider fwd looking PMDs to ensure no obvious issues
- Interop with existing 100G MAC rate

New issues

- Optimized single MAC rates vs. scalable or flexible MAC

Workload

This is a lot of work!!

More of a reflection of the accelerating market applications than anything else.

Single vs. multiple projects will not reduce the work that needs to be done.

A single project has some efficiencies - mostly for the chair(s) - but this doesn't outweigh the considerations around scope and schedule impacts.

Work areas of interest (collapsed view)

		Existing Rates			New Rates	
Media		25G	40G (G2)	100G (Gn)	50G	200G
PCB Traces		1X25G		2x50G 1x100G ?	1x50G	4x50G
BP		1X25G			1x50G	4x50G
Cu Cable		1X25G			1x50G	4x50G
MMF		1X25G	1X40G	2X50G 1x100G	1x50G	4x50G
SMF	500m	1X25G	1X40G	2X50G 1x100G	1x50G	4x50G 2x100G
	2km	1X25G	1X40G	2X50G 1x100G	1x50G	1x50G 2x100G
	10km	1X25G				
	40km					

Option 1: One fully inclusive project.

- Considerable challenges too many aspects
- Schedule concerns with regards to addressing industry relevance
- Significant amount of work to be accomplished. All must end together.
- Dropping work items out of project **needs** to accelerate if possible but not at risk of weakening justification

Work areas of interest (collapsed view)

		Existing Rates			New Rates	
Media		25G	40G (G2)	100G (Gn)	50G	200G
PCB Traces		1X25G		2x50G 1x100G ?	1x50G	4x50G
BP		1X25G			1x50G	4x50G
Cu Cable		1X25G			1x50G	4x50G
MMF		1X25G	1X40G	2X50G 1x100G	1x50G	4x50G
SMF	500m	1X25G	1X40G	2X50G 1x100G	1x50G	4x50G 2x100G
	2km	1X25G	1X40G	2X50G 1x100G	1x50G	1x50G 2x100G
	10km	1X25G				
	40km					

Option 2: Two parallel projects

- Each project has mostly independent challenges
- Both leverage the 50Gb/s technology work of 802.3bs
- Independent schedule enables faster completion to address industry needs
- Some chair-level work coordination needed between projects.

Work areas of interest (collapsed view)

		Existing Rates			New Rates	
Media		25G	40G (G2)	100G (Gn)	50G	200G
PCB Traces		1X25G		2x50G 1x100G ?	1x50G	4x50G
	BP	1X25G			1x50G	4x50G
Cu Cable		1X25G			1x50G	4x50G
MMF		1X25G	1X40G	2X50G 1x100G	1x50G	4x50G
SMF	500m	1X25G	1X40G	2X50G 1x100G	1x50G	4x50G 2x100G
	2km	1X25G	1X40G	2X50G 1x100G	1x50G	1x50G 2x100G
	10km	1X25G				
	40km					

Option 3: One project – “minor” 200G inclusion

- Brings most of the PCS/FEC challenges of 200G into project with associated schedule impacts.
- Leverages the 50Gb/s technology work of 802.3bs
- IEEE no longer defining optical PMDs? Leave this to MSAs
- Not clear how option 3 doesn't become option 1 through SG.

Moving forward

802.3by's progress clears way for Nov 2015 CFI for Single lane 50G based project

Ensuring nx50G (200G) is standardized is necessary.

- Option 1 – Large scope project risks adding 1+ year to project
- Option 2 – Definitely achievable. Needs identification of who/how/when
 - 1 or 2 CFIs → 1 or 2 SGs → 1 or 2 TFs
 - Could be run independently unless some paranoia exists
 - 1 CFI → 1 SG → 1 TF is the traditional model
- Option 3 – High likelihood this becomes Option 1 through SG discussions



Summary

Believe that starting a Single lane 50Gb/s Ethernet CFI later than Nov 2015 is really too late. (Note – no planned scope limitations to just server IO as was done in 802.3by)

- Goal to not be late (again) on next single lane project.

Clear interest in Nx50G implementations being defined for 100G and 200G applications

- Concerned about the under-estimation of project complexity and impact of just adding it to 50 Gb/s project
- Parallel project initiation seems to make sense

ICAID activity (approval needed!) enables us to really ensure this and other topics gets figured out to ensure IEEE continues to serve industry needs