

IEEE 802.3 10Mb/s Backplane Ethernet Call For Interest Consensus Presentation

Amrik Bains/Dylan Walker, Cisco

David Tremblay, HPE

George Zimmerman, CME Consulting

Jon Lewis, Dell EMC

Mandeep Chadha, Microsemi

Dallas, TX

Nov 9th-12th, 2017

Introductions for today's presentation

Presenters and Expert Panel:

Amrik Bains/Dylan Walker

Cisco

David Tremblay

HPE

George Zimmerman

CME Consulting

Jon Lewis

Dell EMC

Mandeep Chadha

Microsemi

CFI Objectives

- To gauge the interest in studying 10Mb/s Backplane Ethernet
- We do not need to:
 - Fully explore the problem
 - Debate strengths and weaknesses of solutions
 - Choose a solution
 - Create a PAR or 5 Criteria
 - Create a standard
- Anyone in the room may vote or speak

Overview: About 802.3cg/10SPE

- Single pair connection
- Low cost/data rate suitable for sensors
- Ability to support power and data
- Multiple reaches (e.g. 15m, 1000m, etc.) for different use cases
- Targeting installed base cable in Industrial and Building Automation

• **It's Ethernet!!**

Overview: Intra-system Management Interface

- Many systems (e.g. server, switch) use a large number of internal control networks
- 10SPE could replace a number of other internal control networks interfaces within systems (e.g., I2C/SMBus, MDIO).
- Many systems that ship in the order of millions of systems per year, contain high 10s to low 100s of these legacy links.
- These legacy links contribute to a relatively high percentage of design, debug and support issues.
- Many endpoint devices are implemented using FPGAs or micro controllers, and could support a “fast & richer” interface than I2C etc.

What Are We Talking About?

- *What needs to be done to enable this transition?*
- *New link segment?*
- *Protocol definitions (out of scope)*
- *?????*

Agenda

- Overview Discussion Jon Lewis, Dell EMC
- Presentations
 - 10BP in Servers Jon Lewis, Dell EMC
 - 10BP in Fixed Switches Amrik Bains/Dylan Walker, Cisco
 - 10BP in Modular Switches David Tremblay, HPE
 - 10BP Technical Feasibility Mandeep Chadha, Microsemi
 - 10BP Why Now? George Zimmerman, CME
- Q&A
- Straw Polls

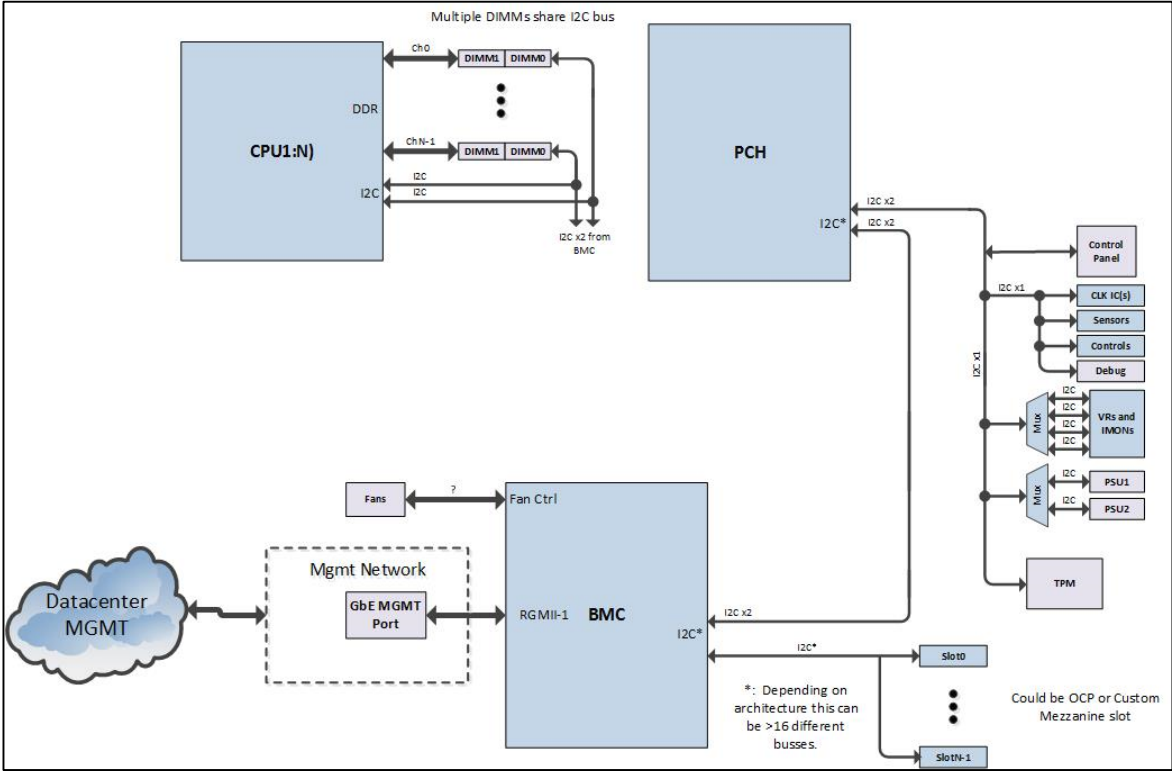
10BP in Servers

Jon Lewis, Dell EMC

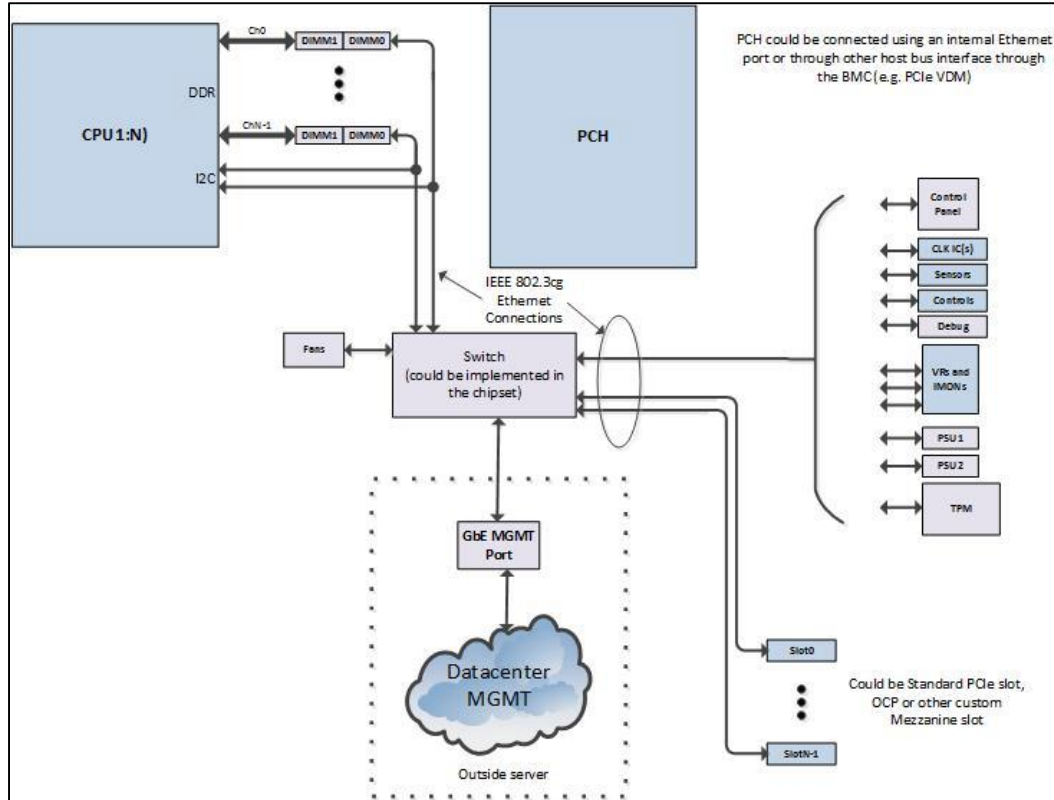
Server market size info

- *Server Market is currently ~10 million units per year*
- *Average number of devices in standard servers - >20 endpoints per server*
- *Total endpoint device >200 million parts per year*
- *Total switch device >10 million per year*

Current Architecture



Desired Architecture



Why 10BP?

Same number of pins as SMBus which would maintain current PCIe standard pin count while adding network functionality.

- Using Ethernet allows for discovery of devices using a common BMC driver.
 - Going from >10 custom drivers to a standard Ethernet driver reduces complexity on the BMC coding and will greatly reduce validation time required for all custom implementations.
- As the compute node and networking “converge” there is a fine line between what is in a traditional server and what is in the networking “Ether”
 - External PCIe Expanders
 - Chassis servers
 - Modular server implementations. Ethernet provides a standard ubiquitous management communication path
- Easier for automated alerts than multi-master SMBus.
 - Using SMBus multiplexers makes multi-master difficult ☹ impossible. Thus scanning the bus continuously is required.
 - Using Ethernet the endpoint could transmit the alert at any time without a master/slave relationship.

10BP in Fixed Switches

Amrik Bains/Dylan Walker, Cisco

Current Architecture

- Internal control plane is used to perform configuration/monitoring of components in the switch
- Many different components with various control plane interfaces, e.g.

➤ I2C/SMB Bus

- Optical Module, AC/DC Power Supplies, FAN Control, DC-DC Converters, Temp monitors, EPROM etc....
- 2 wire – Clock + Shared Data (Tx/Rx)
- Clock speed 100KHz, data BW (25 to 30Kb/s) limited due half-duplex and protocol overhead

➤ UART: Universal Asynchronous Rx/Tx

- Micro-controllers/CPU
- 2 wire – Rx/Tx
- 9.6Kb/s

➤ MDIO (IEEE 802.3 Clause 22/45): Copper PHYs and Fiber PHYs

- 2 wire
- Tx/Rx shared (half-duplex)
- Max specified MDC clock of 2.5MHz (avg. BW 1Mb/s)

➤ SPI :

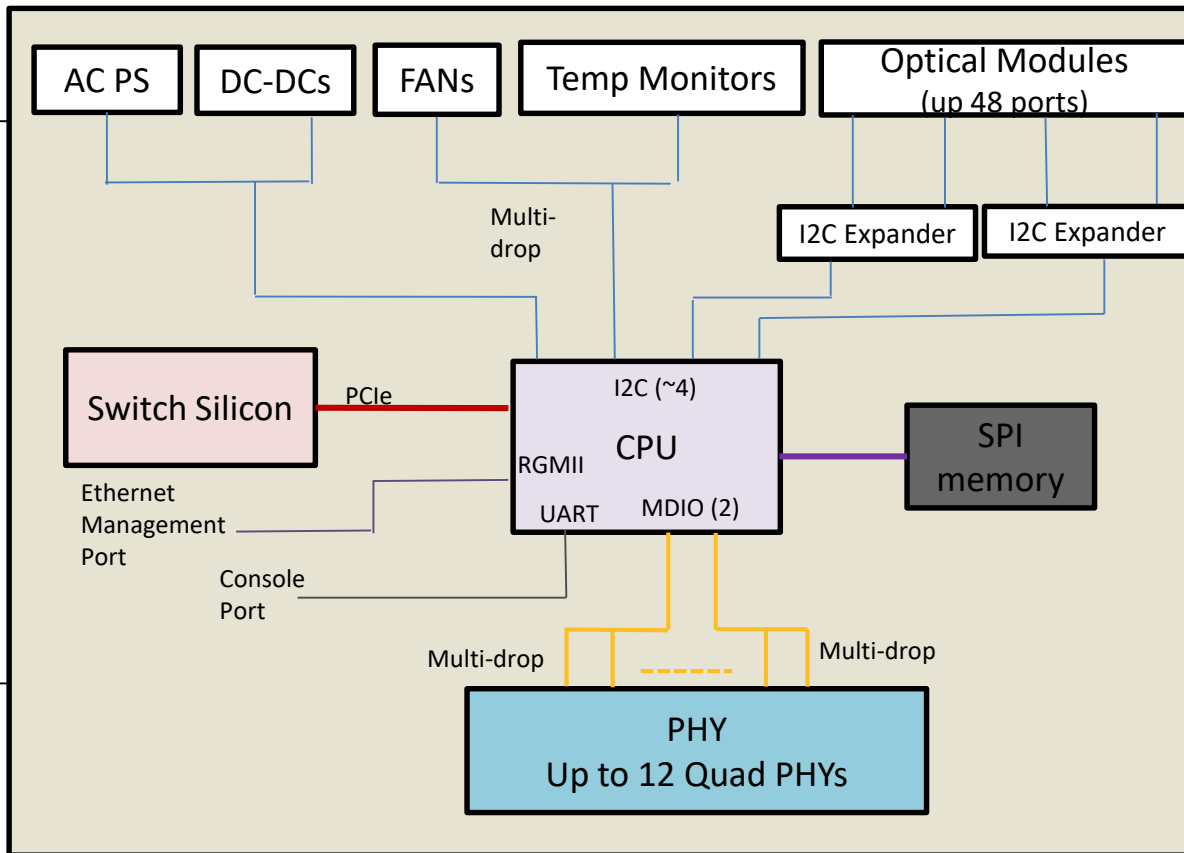
- SD card, Sensors, eMMC
- Minimum of 4 wire – Clock, Rx, Tx and Slave Select (incase of multiple devices connected to same data pins)
- Typically 12 to 25Mb/s
- Allows more data wires to used for higher bandwidth

Current Architecture

- Current management interfaces were specified in isolation, to solve particular issues, not ideal for overall system management
- Management Interfaces have not kept up with BW requirements
 - Devices require more configuration and status monitoring
 - AC Power Supplies: Temp, voltage, current, fan monitoring
 - DC-DC converters : Temp, voltage, current
 - Copper PHY and Copper PHY => Firmware image increases with complexity
 - Memories: Size increases with every node, requires higher BW for writes and reads
- As complexity increases, management interface needs to be provide reliable and faster access
 - e.g., I2C/SMBUS has reliability issues when used as multi-drop
- Each interface requires different “Software Drivers”
 - Creates development and maintenance issues

Current Architecture

- Too many interface types
- I/O types limited on CPU – require external devices to expand I/O
- Multi-drop – subject multiple failure conditions
- Limited BW
- Multiple Software Driver requirements

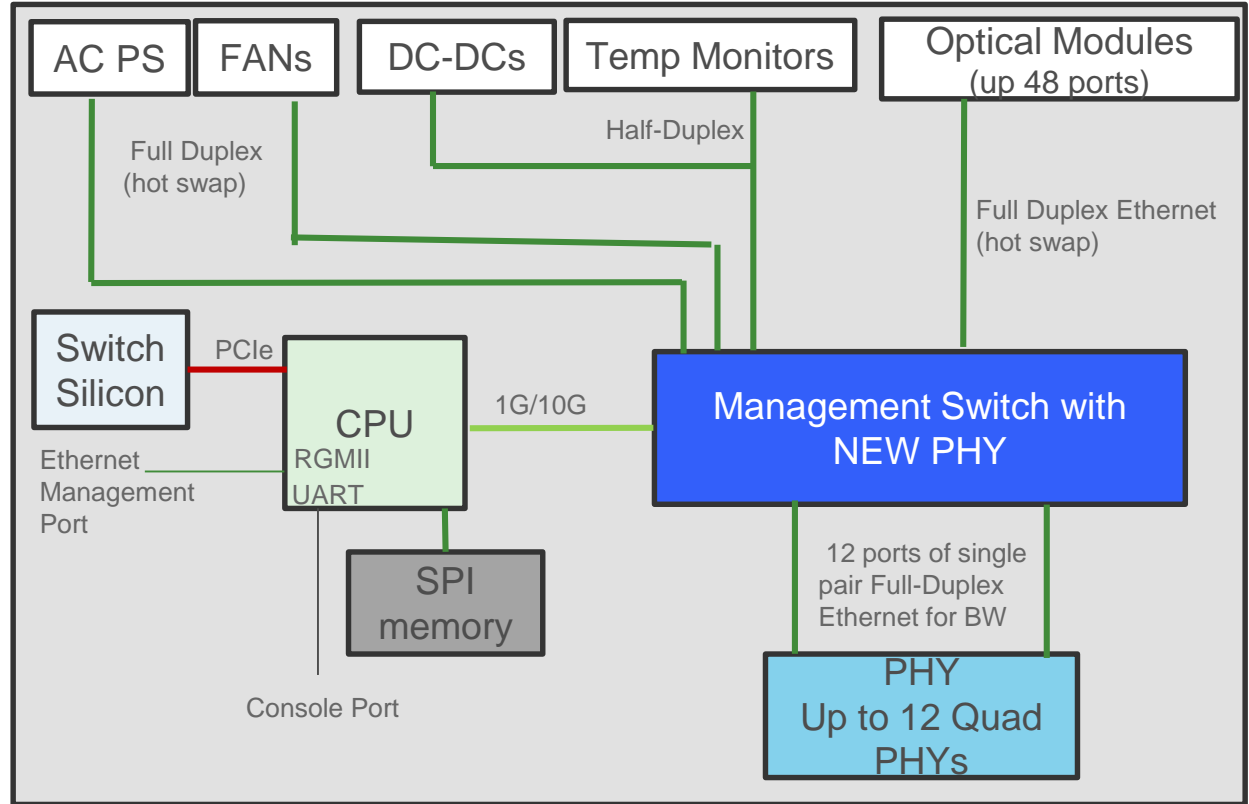


Why 10BP?

- 2-wire: Reduce inter-connect
- 10Mb/s is higher usable bandwidth than we have today
 - 100Mb/s may be interesting in future
- Ethernet based – one common driver
- Switched Point-to-Point
 - for BW or hot-puggable devices
- Half-Duplex for low bandwidth and fixed devices
 - e.g., DC-DC, Temp monitors etc.

Desired Architecture

- Number of different I/O types reduced by using
- CPU I/O limitation for I2C/MDO removed
- Point-Point Ethernet for BW and or Hot-pluggable interfaces
- Ethernet Driver

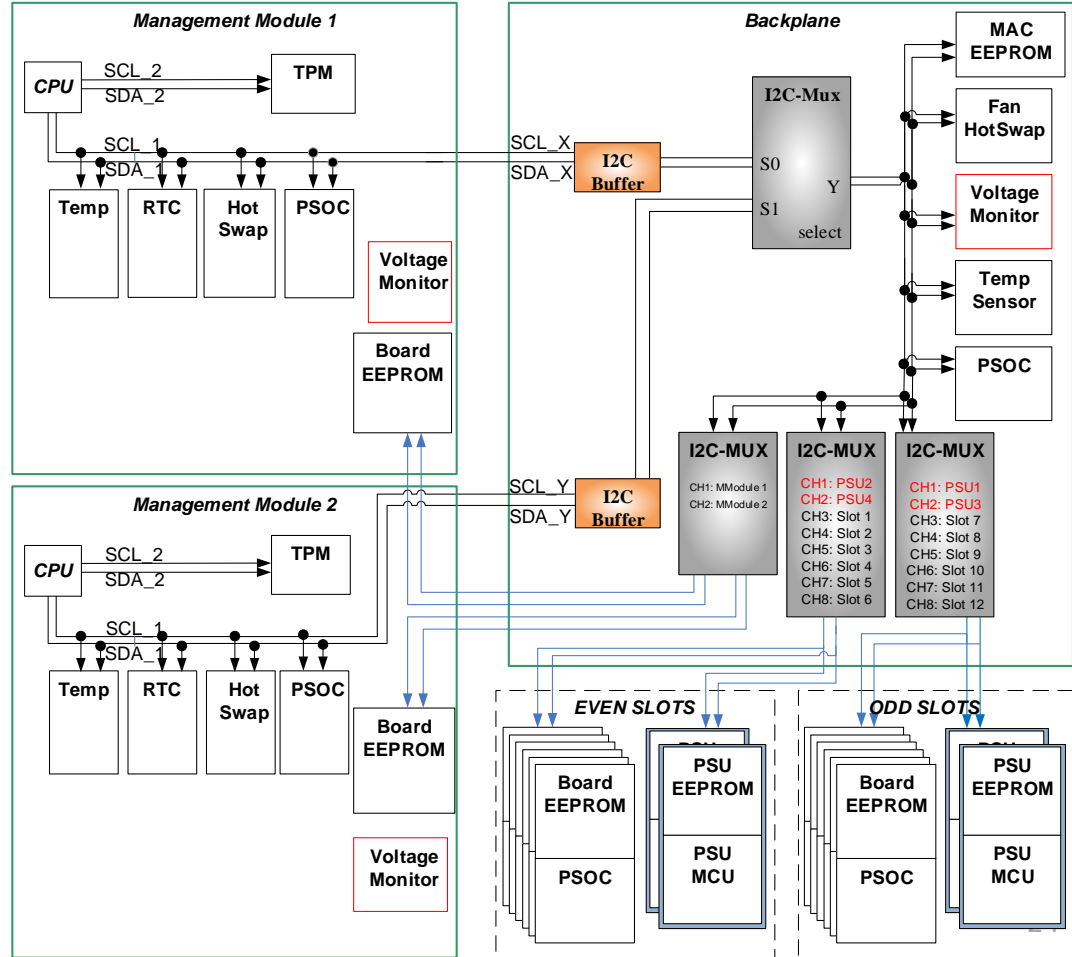


10BP in Modular Switches

David Tremblay, HPE

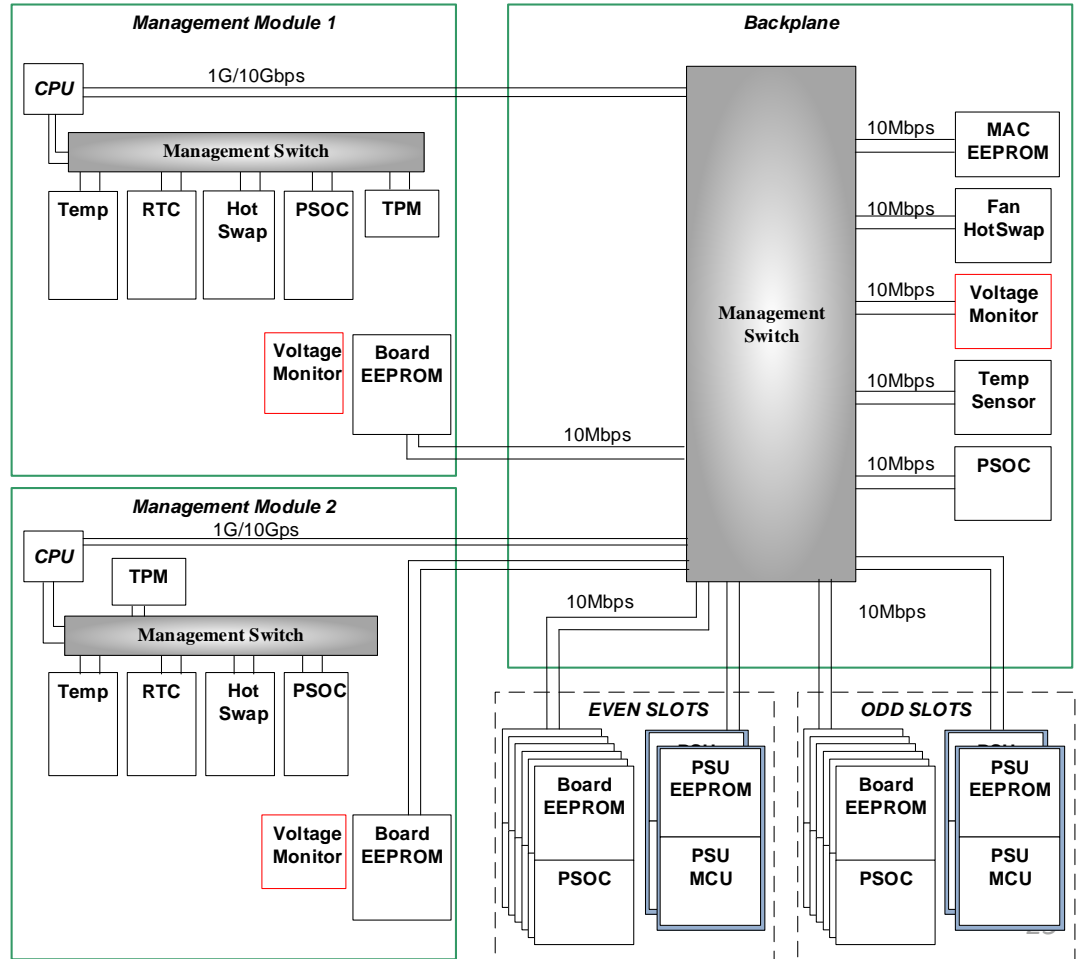
Current Modular Architecture

- CPU has limited I/O – requires external devices to expand I/O
- Multi-drop – subject to failure
- SW Driver Variation – requires excessive tuning and validation
- Limited Bandwidth



Desired Modular Architecture

- RBOM Reduction
- Point to Point Connectivity
- Single SW Driver
- Full Bandwidth Support



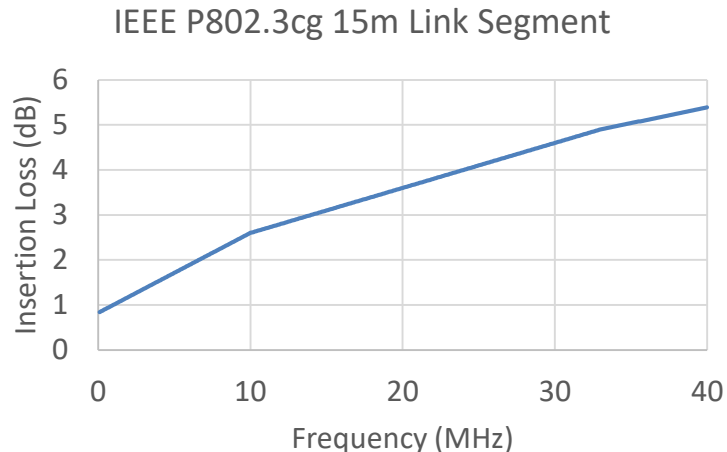
10BP Technology Feasibility

Mandeep Chadha, Microsemi

Lengths/link segment

- IEEE 802.3cg has a short reach link segment adopted
 - Based on cabling and connector measurements
- Consistent with needs of backplane channels (Jon Lewis is working on this)

- Adopt the equations on slide 18 of http://www.ieee802.org/3/cg/public/Sept2017/DiBiao_Bergner_01c_0917.pdf as a baseline for the 10SPE short reach link segment.



IL <	$1 + 1.6 (f-1)/9$ dB	f=0.3.... 10MHz
	$2.6 + 2.3 (f-10)/23$ dB	f=10 33MHz
	$4.9 + 2.3 (f-33)/33$ dB	f=3340MHz
RL >	14 dB	f=0.3...10MHz
	$14 - 10 * \text{LOG}_{10}(f/10)$ dB	f=10...40MHz
MC >	30 dB	f=0.3...20MHz
	$30 - 20 * \text{LOG}_{10}(f/20)$ dB	f=20...200MHz

10Mb/s Ethernet in Micros/FPGAs

- TBD
- Low insertion loss of channel gives lots of choices
- Many encoding types available
 - Differential Manchester, NRZ, etc.
 - All are low-complexity and have been implemented in FPGA format
- Duplexing method is new for micros/FPGAs
 - E.g., echo-cancelled full-duplex or half-duplex on medium
 - Already under study in 802.3cg for short-reach channel
- [some kind of complexity estimate based on either NRZ or DME half-duplex could go on next slide]

Summary of 10MbE Backplane Technical Feasibility & Application Potential

- Prospective silicon area or number of ports on a switch chip?

10BP - Why Now?

George Zimmerman, CME

Why do it at all?

- Interest from Network Equipment and Computer OEMs has created the potential for large volume short-reach interconnects
 - Should be useful in industrial applications too
 - Needs are very consistent with short-reach PHY
 - Should be manageable without significant delay, and little new work
- The Ethernet Ecosystem has been very successful
 - Open and common specifications
 - Ensured Interoperability
 - Security of development investment
 - 802.3cg is evidence Ethernet is filling more control-plane applications
 - Every time I turn around, there's a new application being proposed

Why Now?

- Leverage investment in standardization of 10Mbps single-pair technology
 - Let's be honest – 10Mbps Ethernet systems are old hat – but there aren't any single-pair versions standardized
 - Relevant PHY experts are gathered in IEEE 802.3cg
 - Likely commonality of requirements with 802.3cg short-reach
- BUT: Existing 802.3cg project documentation is specific to “single balanced twisted pair copper cabling”
 - Could be a new project, but that would cause more interference with 802.3cg

The Rub: 802.3cg PAR Scope, very specific

5.2.b. Scope of the project: Specify additions to and appropriate modifications of IEEE Std 802.3 to add 10 Mb/s Physical Layer (PHY) specifications and management parameters for operation, and associated optional provision of power, on single balanced twisted-pair copper cabling.

- Limitations:
 - Rate = 10Mbps
 - The project will define PHY specifications (allows more than 1 PHY)
 - The project will define optional powering as long as it is associated with 10Mbps PHYs
 - The medium is single balanced twisted-pair copper cabling
 - 2-pair, 4-pair are out of scope, shared sheath specifications using multi-pair cabling may be
 - Backplane, parallel pairs, or “open wire” pairs are unallowed (even though they usually work)
- Addressing backplane applications is a natural
 - Including PCB “pairs”: The electrons don’t care about twisting the wires or read PAR documents...
- If we don’t do this now, the 802.3cg solutions will likely be drafted, perhaps with proprietary extensions

Process

- Call for Interest on 10Mbps Backplane Ethernet to be held in November
- If successful, Study group will meet
 - Expect most work will be done on ad hoc calls before January
 - PAR and CSD modifications to 802.3cg to remove the “twisted pair” limitations
 - If on plan, working group will vote to forward PAR to IEEE-SA in January or March at latest
- Meanwhile, 802.3cg continues its work on short-reach and long reach PHYs
 - Contributors work offline to prepare and build consensus for any necessary text
 - Folds in anything necessary for backplane in March
- This process has been followed before, most recently adding 25GBASE-T to IEEE P802.3bq (originally just 40GBASE-T)
 - Actually, this should be MUCH easier than 802.3bq, because there are none of the differences associated with adding another speed (MAC interface, registers, or frequency translations)

Proposed Path Forward

- Resources to develop 10BP reside in P802.3cg
- Propose extending the work of P802.3cg to include 10BP
 - Modest incremental effort, anticipate common technology
 - P802.3cg motion showed unanimous support for that path : [P802.3cg Sept'17 motion #15](#)
- Proposed Next Steps
 - 802.3 approve formation of 10BP Study Group (this plenary)
 - SG determines objectives for 10BP
 - SG prepares updated PAR scope Objectives & CSD to modify the scope of P802.3cg
 - 802.3 approval of modified P802.3cg project documentation
 - By March 2018
- Enough time before 802.3cg adopted WG ballot milestone in May

Q&A

Presenters

Amrik Bains/Dylan Walker

David Tremblay

George Zimmerman

Jon Lewis

Mandeep Chadha

Cisco

HPE

CME Consulting

Dell EMC

Microsemi

Expert Panel

<TBD>

Straw Polls

Call-for-Interest Consensus

- Should a study group be formed for “10Mb/s Backplane Ethernet”?
- Y: N: A:
- Room count:

Participation

- I would participate in a “10Mb/s Backplane Ethernet” study group in IEEE 802.3
 - Tally:

- My company would support participation in a “10Mb/s Backplane Ethernet” study group
 - Tally:

Future Work

- Ask 802.3 at Thursday's closing meeting to form study group
- If approved:
 - 802 EC votes on Friday to approve the formation of the study group
 - First study group meeting would be during the January 2018 802.3 interim meeting (in Geneva)

End