IEEE 802.3
Beyond 400 Gb/s Ethernet Study Group

28 October 2021

Draft

Project Overview –

IEEE P802.3df:

200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet

Contributors

- John D’Ambrosia, Futurewei, U.S. Subsidiary of Huawei
- Kent Lusted, Intel
- Gary Nicholl, Cisco
- Dave Ofelt, Juniper
- Mark Nowell, Cisco
- Matt Brown, Huawei

Agenda

- Project Overview
- Defining the Architecture – A Holistic Approach
- Addressing Multiple Physical Layer Specifications
- Managing IEEE P802.3df
- Discussion

Introduction

- Goals of the presentation:
  - Provide a broad overview of the proposed IEEE P802.3df PAR, its objectives, and the opportunities for technology reuse
  - Communicate the importance of utilizing a “holistic” approach to initially address all of the IEEE P802.3df objectives simultaneously
  - Discuss the potential future timeline and management of the project

- What do we mean by holistic? According to Merriam-Webster:

  Holistic: relating to or concerned with wholes or with complete systems rather than with the analysis of, treatment of, or dissection into parts

Project Information

- IEEE 802.3 Beyond 400 Gb/s Ethernet Study Group
  - Website: https://www.ieee802.org/3/B400G/index.html
  - Charter
    - The IEEE 802 LMSC Executive Committee has chartered a Study Group under the IEEE 802.3 Ethernet Working Group to develop a Project Authorization Request (PAR) and Criteria for Standards Development (CSD) responses for:
      1. Beyond 400 Gb/s Ethernet
      2. Physical Layer specifications for existing Ethernet rates based on Physical Layer specifications for beyond 400 Gb/s Ethernet.

- Draft IEEE P802.3df Project
  - Title
    - Media Access Control Parameters, Physical Layers and Management Parameters for 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Operation
  - Scope
    - Define Ethernet Media Access Control (MAC) parameters, physical layer specifications; and management parameters for the transfer of Ethernet format frames at 800 Gb/s and 1.6 Tb/s over copper, multi-mode fiber, and single-mode fiber, and use this work to define derivative physical layer specifications and management parameters for the transfer of Ethernet format frames at 200 Gb/s and 400 Gb/s.
  - Draft Project Documentation
    - PAR: https://mentor.ieee.org/802-ec/dcn/21/ec-21-0224-01-00EC-par-ieee-p802-3df.pdf
    - CSD: https://mentor.ieee.org/802-ec/dcn/21/ec-21-0225-00-00EC-csd-ieee-p802-3df.pdf
    - Objectives: https://www.ieee802.org/3/B400G/proj_doc/objectives_b400g_210826.pdf

The Starting Point for the Study Group

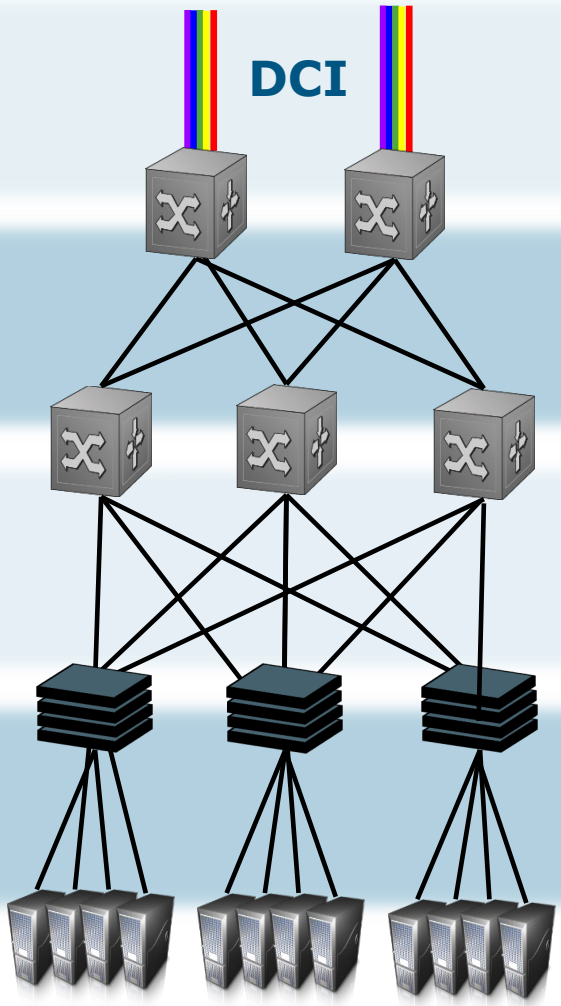

Example Network Topology

	Next Speed	Reach	Medium
Router	> 400 GbE	80 km	DWDM SMF
Leaf/Spine	> 400 GbE	500 m to 40 km	Duplex SMF
TOR / MOR / Leaf	> 400 GbE	50 m to 500 m	MMF / SMF (duplex or parallel)
Server	> 100 GbE	< 50 m	MMF Twin Axial

Urgency for an 800 Gb/s Ethernet Standard

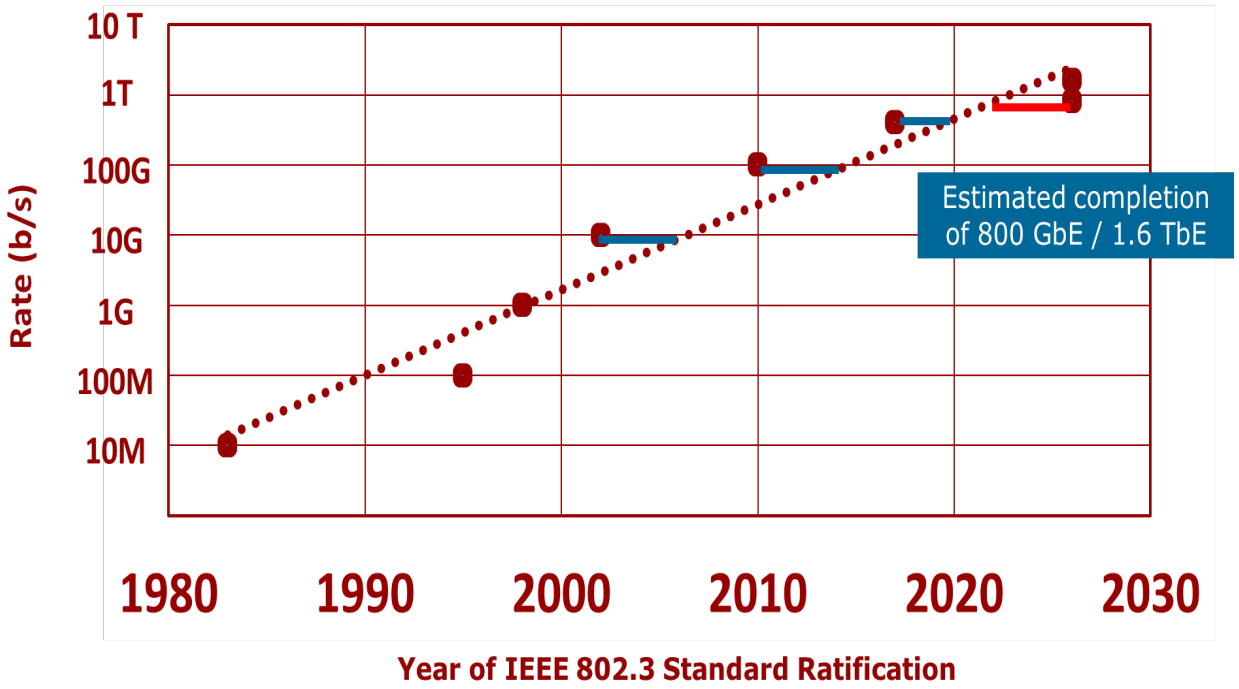

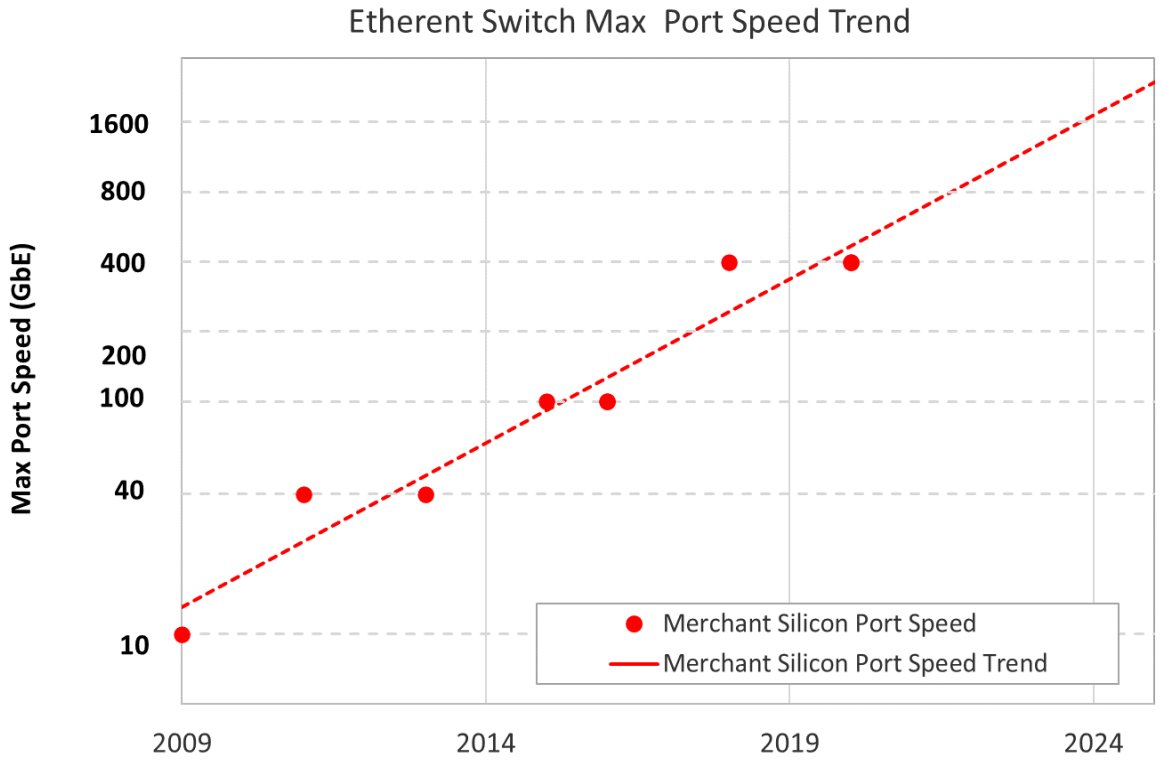

Data provided by Rob Stone, Facebook

Ethernet Technology Consortium publicly released 800 GbE specification
(https://ethernettechnologyconsortium.org/wp-content/uploads/2020/03/800G-Specification_r1.0.pdf)

- Draft 1.0 dated March 10, 2020
- Based on 8x100G (references to various IEEE 802.3 clauses)

The Relationship Between Ethernet & Signaling Rates

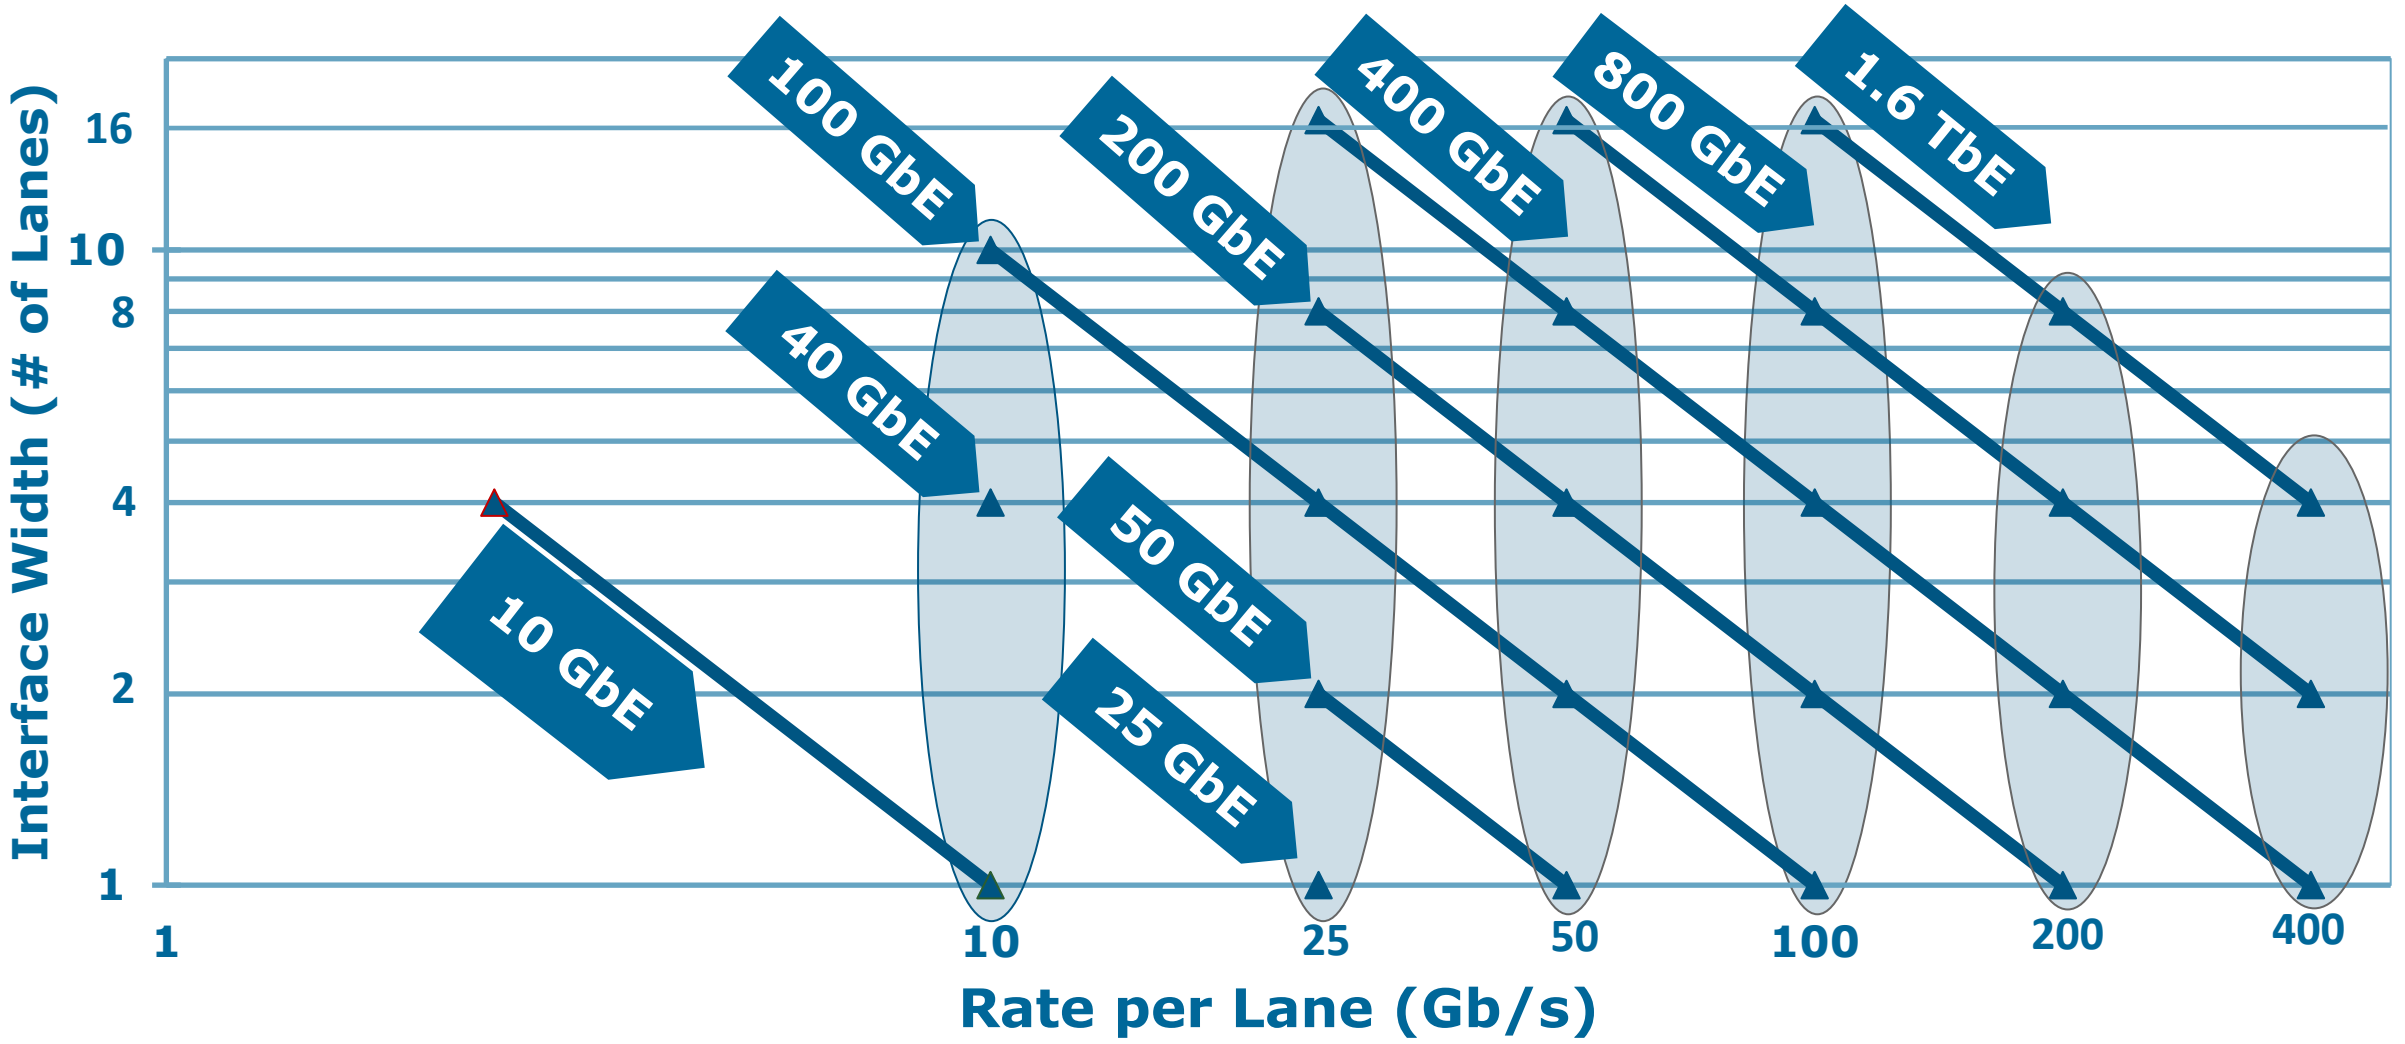

Growing Acceptance of x8 Solutions

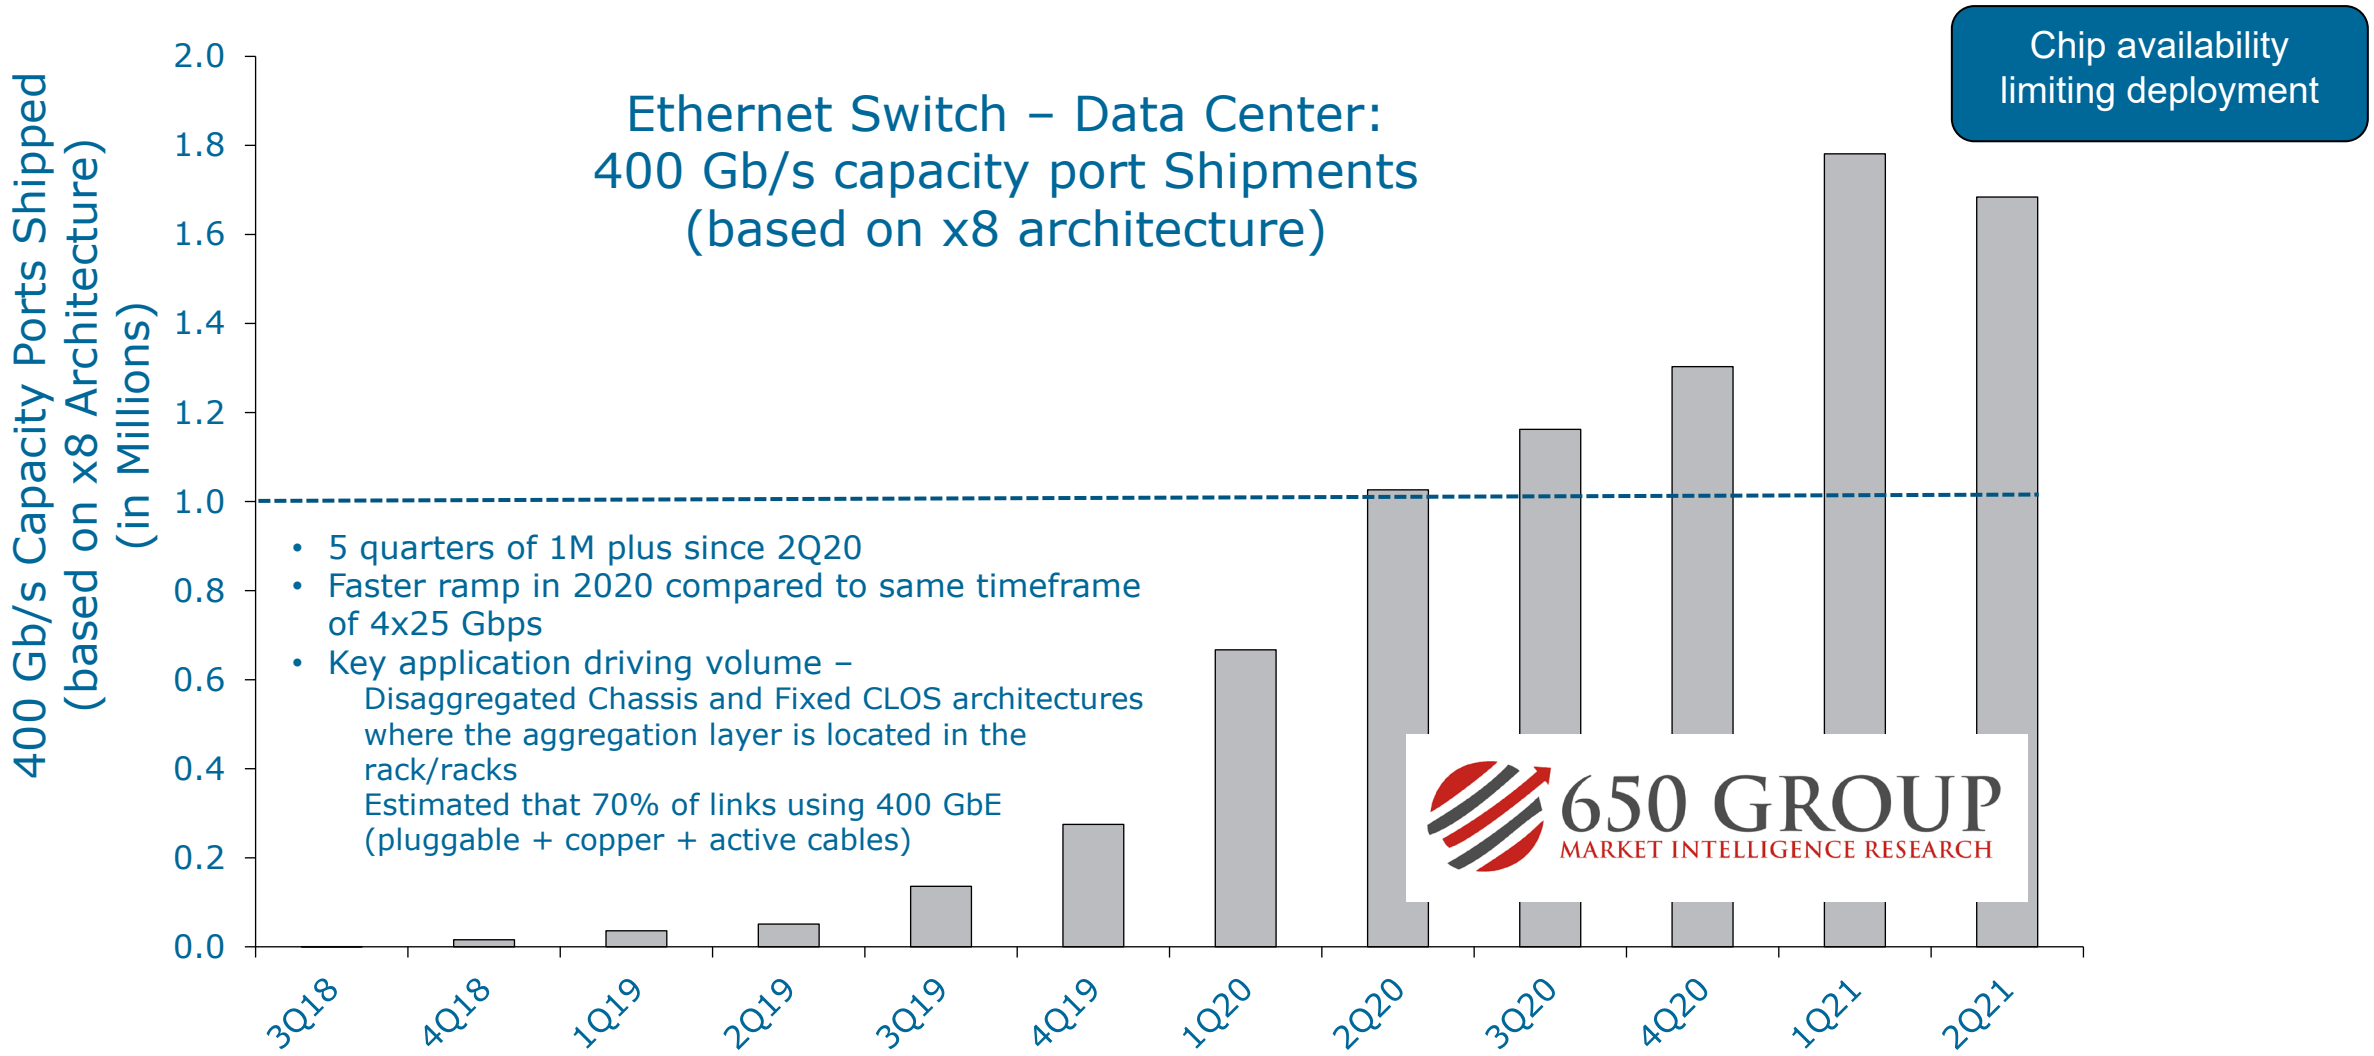

Data used with permission, Alan Weckel, 650 Group

Form Factor Adoption Comparison

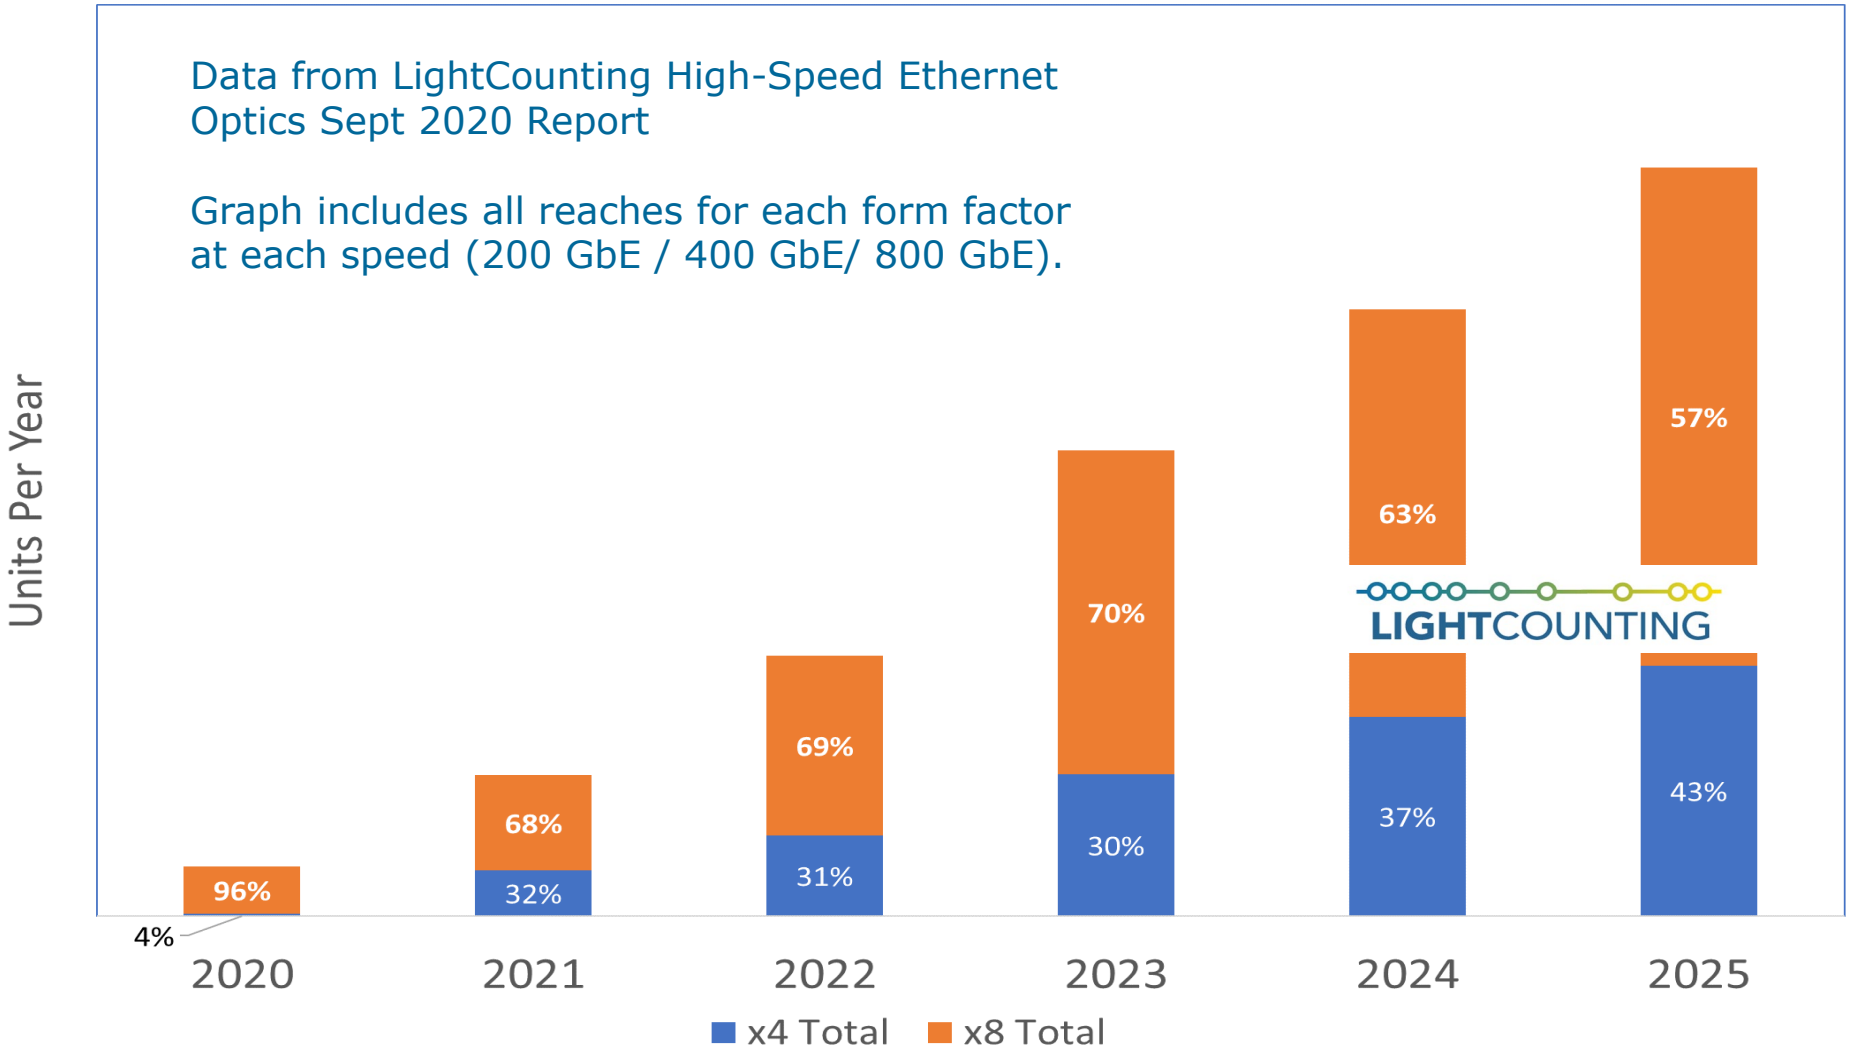

Leveraging Signaling Technologies Across Ethernet Rates

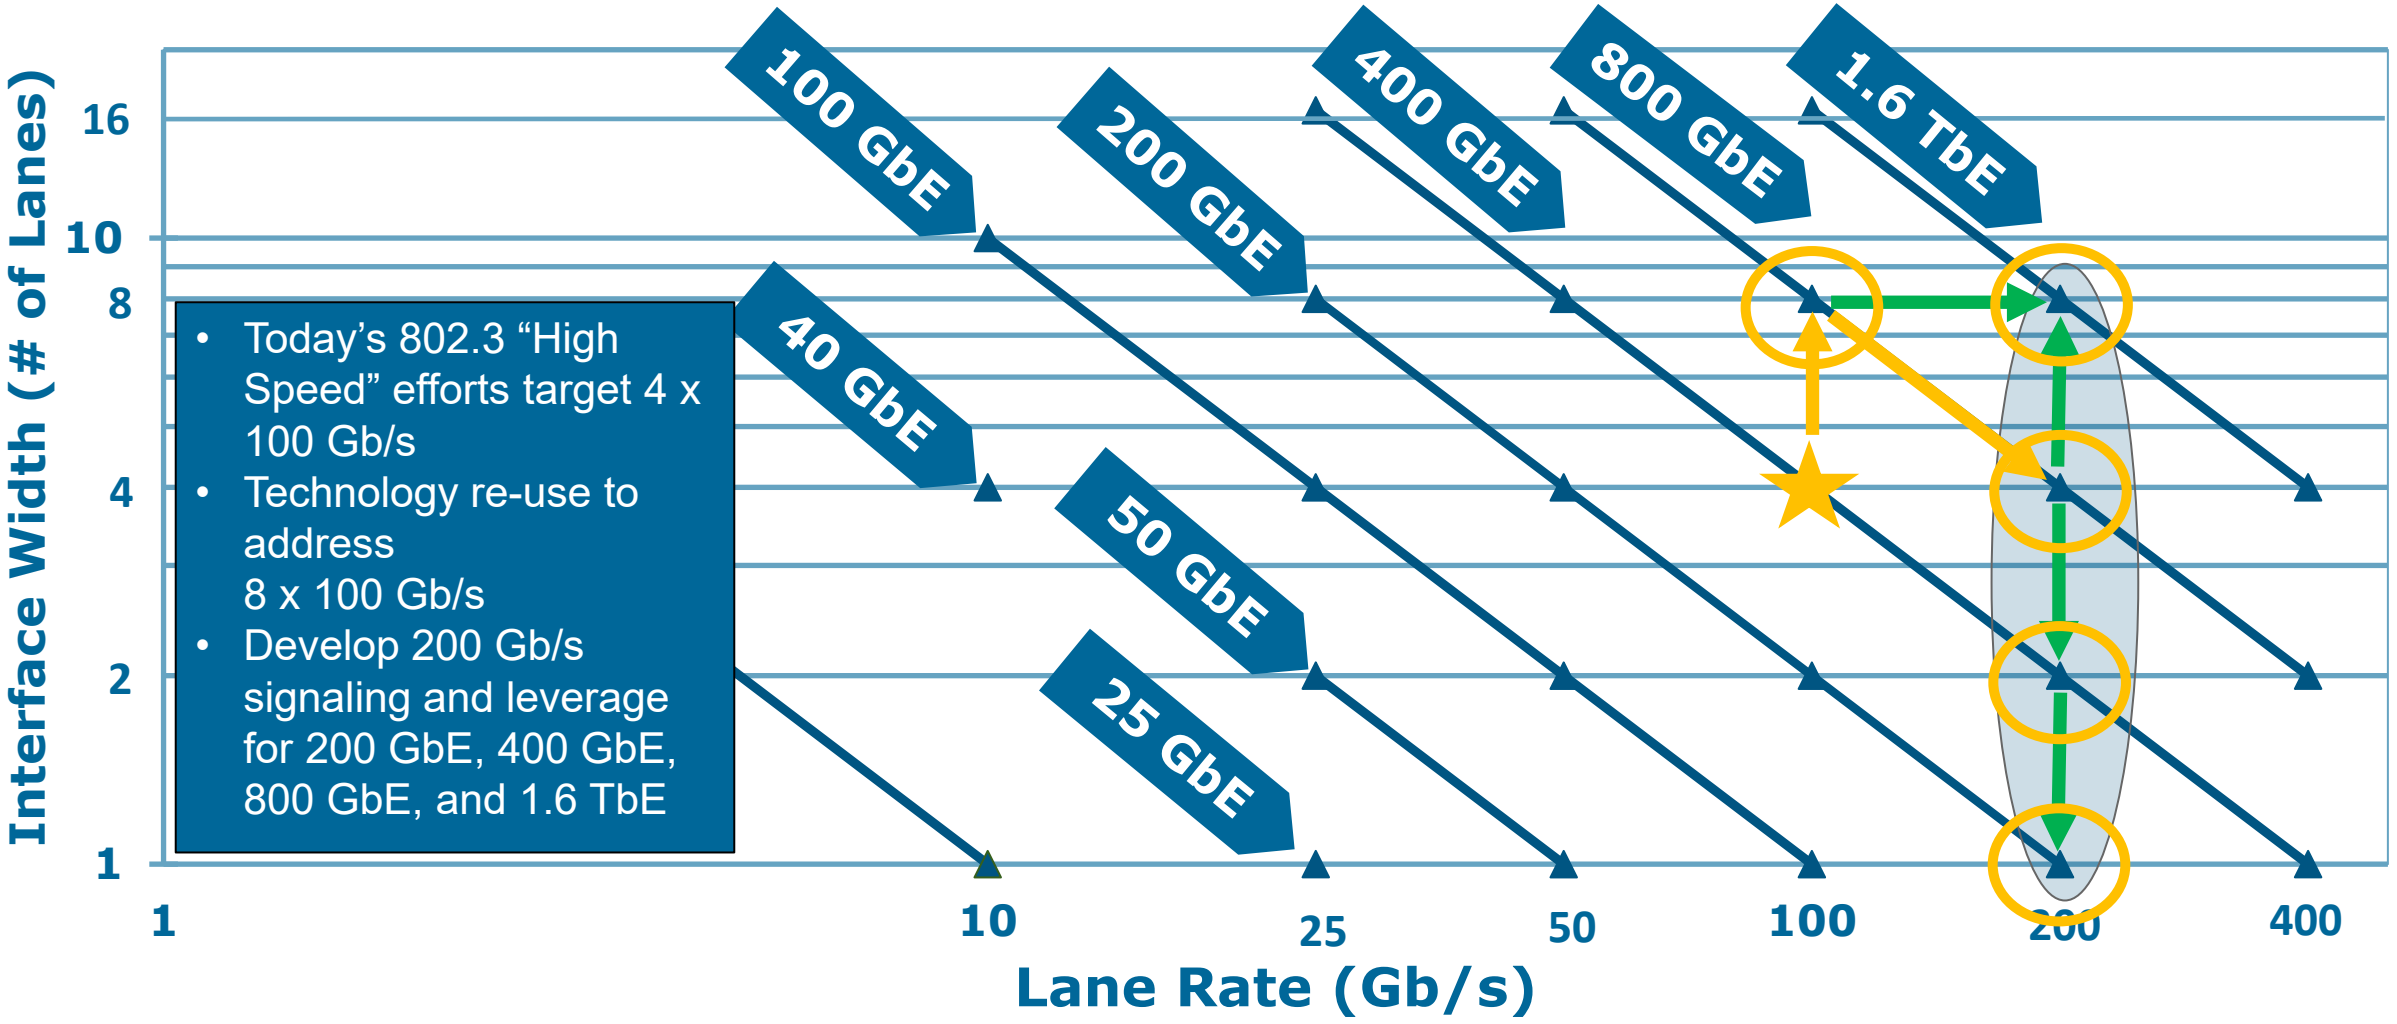

- Today's 802.3 "High Speed" efforts target 4 x 100 Gb/s
- Technology re-use to address 8 x 100 Gb/s
- Develop 200 Gb/s signaling and leverage for 200 GbE, 400 GbE, 800 GbE, and 1.6 TbE

Adopted Physical Layer Objectives

Ethernet Rate	Assumed Signaling Rate	AUI	BP	Cu Cable	MMF 50m	MMF 100m	SMF 500m	SMF 2km	SMF 10km	SMF 40km
200 Gb/s	200 Gb/s	Over 1 lane		Over 1 pair			Over 1 Pair	Over 1 Pair		
400 Gb/s	200 Gb/s	Over 2 lanes		Over 2 pairs			Over 2 Pair			
800 Gb/s	100 Gb/s	Over 8 lanes	Over 8 lanes	Over 8 pairs	Over 8 pairs	Over 8 pairs	Over 8 pairs	Over 8 pairs		
	200 Gb/s	Over 4 lanes		Over 4 pairs			Over 4 pairs	1) Over 4 pairs 2) Over 4 λ's		
	TBD								Over single SMF in each direction	Over single SMF in each direction
1.6 Tb/s	100 Gb/s	Over 16 lanes								
	200 Gb/s	Over 8 lanes		Over 8 pairs			Over 8 pairs	Over 8 pairs		

https://www.ieee802.org/3/B400G/proj_doc/objectives_b400g_210826.pdf

DEFINING THE ARCHITECTURE: A HOLISTIC APPROACH

Lessons Learned From Past Projects and Deployments

- Today's "Ethernet port" expected to support
  - Multiple Ethernet rates
    - ASICs support multiple MACs, PCS, and multiple SERDES signaling speeds
      e.g. (a port could support 200 GbE based on 4 x 50 Gb/s or 400 GbE based on 4 x 100 Gb/s)
    - "Breakout" of port capacity. Commonly used with physical layer specifications based on multiple conductors or fibers.
  - Multiple media and reaches
  - Multiple topologies
- Today's expectations were enabled by
  - Standards leveraging technology reuse of signaling rates per lane
  - Focus by IEEE 802.3 on backwards / forward-looking compatibility
- Independent development of standards can lead to long term system port usability issues

Example Issue

	FEC	PHY
802.3ba-2010	No FEC for 25 Gb/s per lane optics specified	100GBASE-LR4 100GBASE-ER4
802.3bm-2015	RS (528, 514) KR4 FEC	100GBASE-SR4

100GBASE-SR4 optics not supported by 802.3ba based ports

Learning from the Past as We Go Forward

An initial holistic approach is valuable to define an architecture:

- Maximize technology reuse while enabling support for:
  - Different Ethernet rates
  - Different signaling rates per lane
  - Different FEC codes / schemes
  - Different media / different reaches
- Benefits
  - Reduce time to develop multiple Ethernet specifications
  - Reduce cost / time to develop products
  - Improve potential for long term port usability of equipment
    - Beneficial to component and system suppliers ROI
    - Beneficial to end-users for long term usage of equipment

Architecture Inputs & Constraints

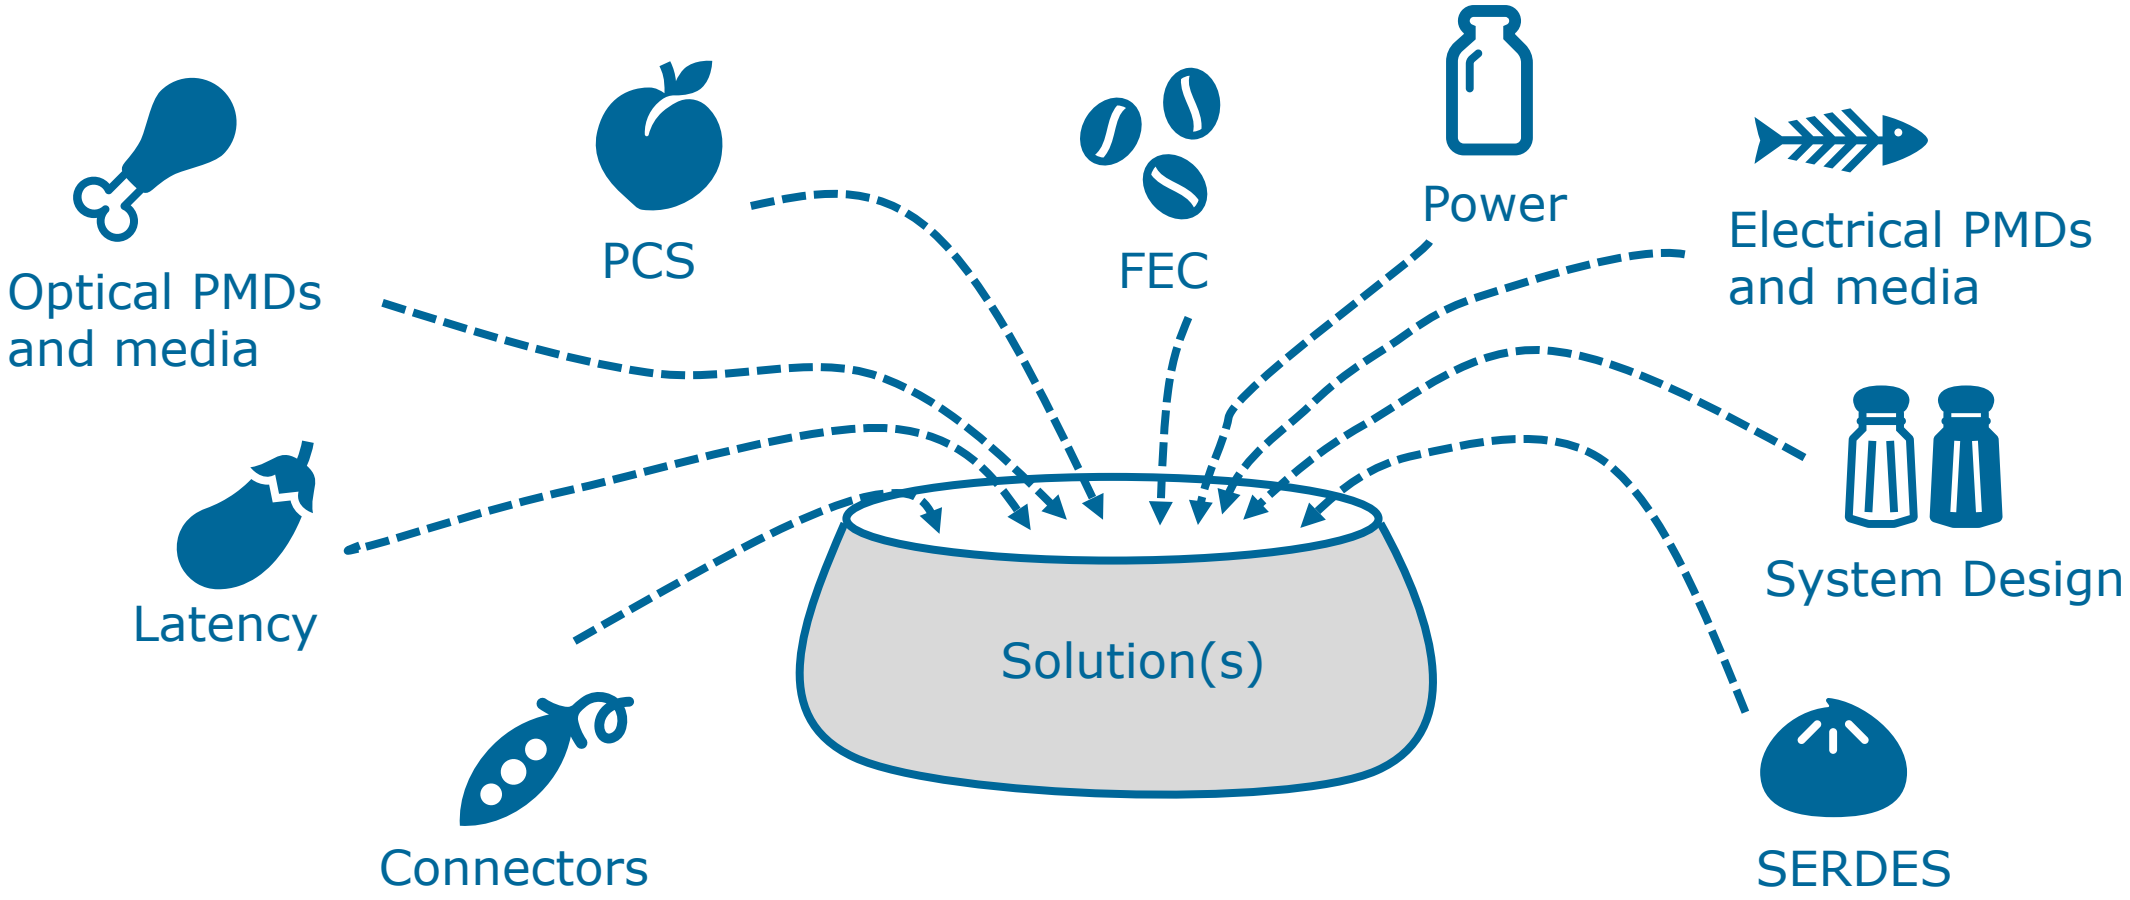

Look at all elements together to find the best solution(s)

Universal Multi-rate / Multi-medium Ports

PAST

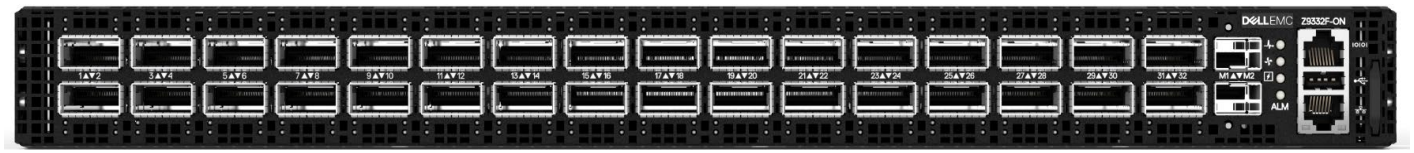

Image courtesy of David Piehler, Dell Technologies

- 32 x 400 Gb/s Capacity Ports
- Supports 32 x 400 GbE Ports or 128 x 100 GbE Ports
- FEC plays role in supporting mediums @ 100 GbE: -CR, -SR, -DR, -FR, -LR, -ER, -ZR

DESIRED FUTURE

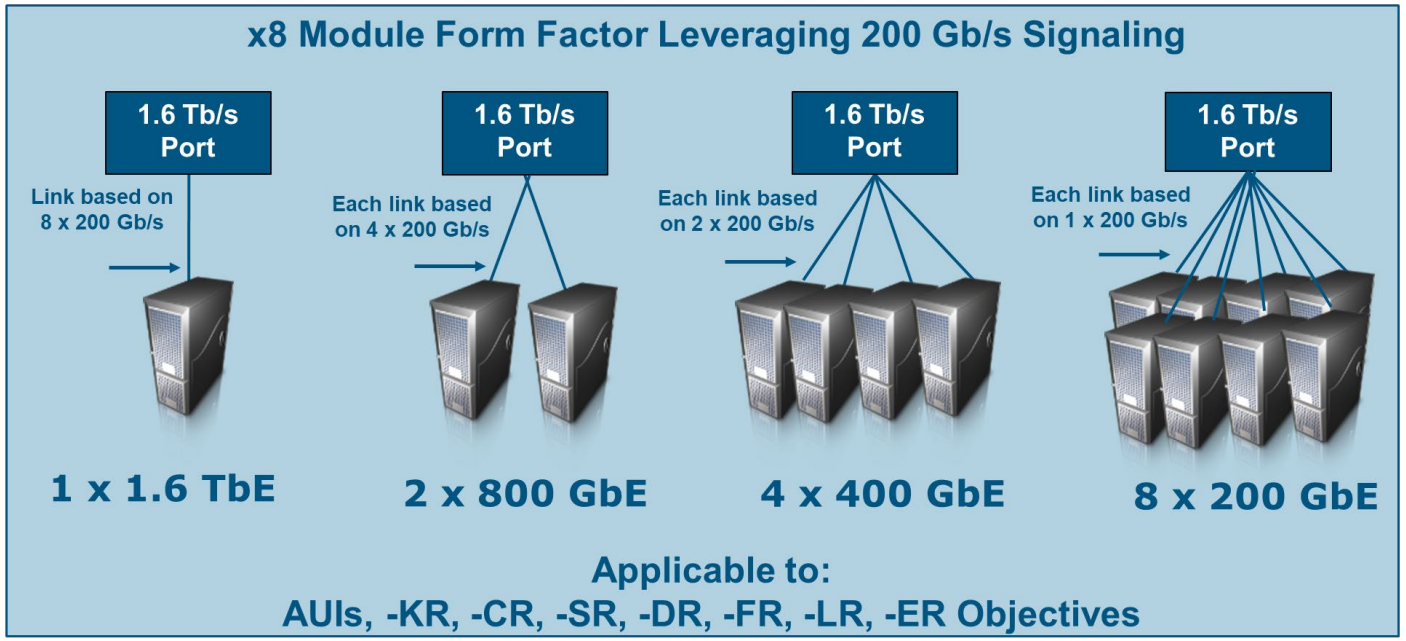

- 32 x 1.6 Tb/s Capacity Ports
- Supports 32 x 1.6 TbE Ports, 64 x 800 GbE Ports, 128 x 100 GbE Ports, 256 x 200 GbE Ports
- Optimize FEC to support all mediums - CR, -SR, - DR, -FR, - LR, -ER, -ZR

Consider 800 GbE

AUI 4 x 200 Gb/s 8 x 100 Gb/s	SMF Cabling (10 km / 40 km) – 1 Pair SMF Cabling (500 m / 2 km) – 4 Pairs SMF Cabling (500 m / 2 km) – 8 Pairs MMF Cabling (50 m / 100 m) – 8 Pairs	**AUI** 4 x 200 Gb/s 8 x 100 Gb/s
-CR 4 x 200 Gb/s 8 x 100 Gb/s	Twin-Axial Cabling – 4 Pairs Twin-Axial Cabling – 8 Pairs	**-CR** 4 x 200 Gb/s 8 x 100 Gb/s
-KR 8 x 100 Gb/s	Backplane – 8 Pairs	**-KR** 8 x 100 Gb/s

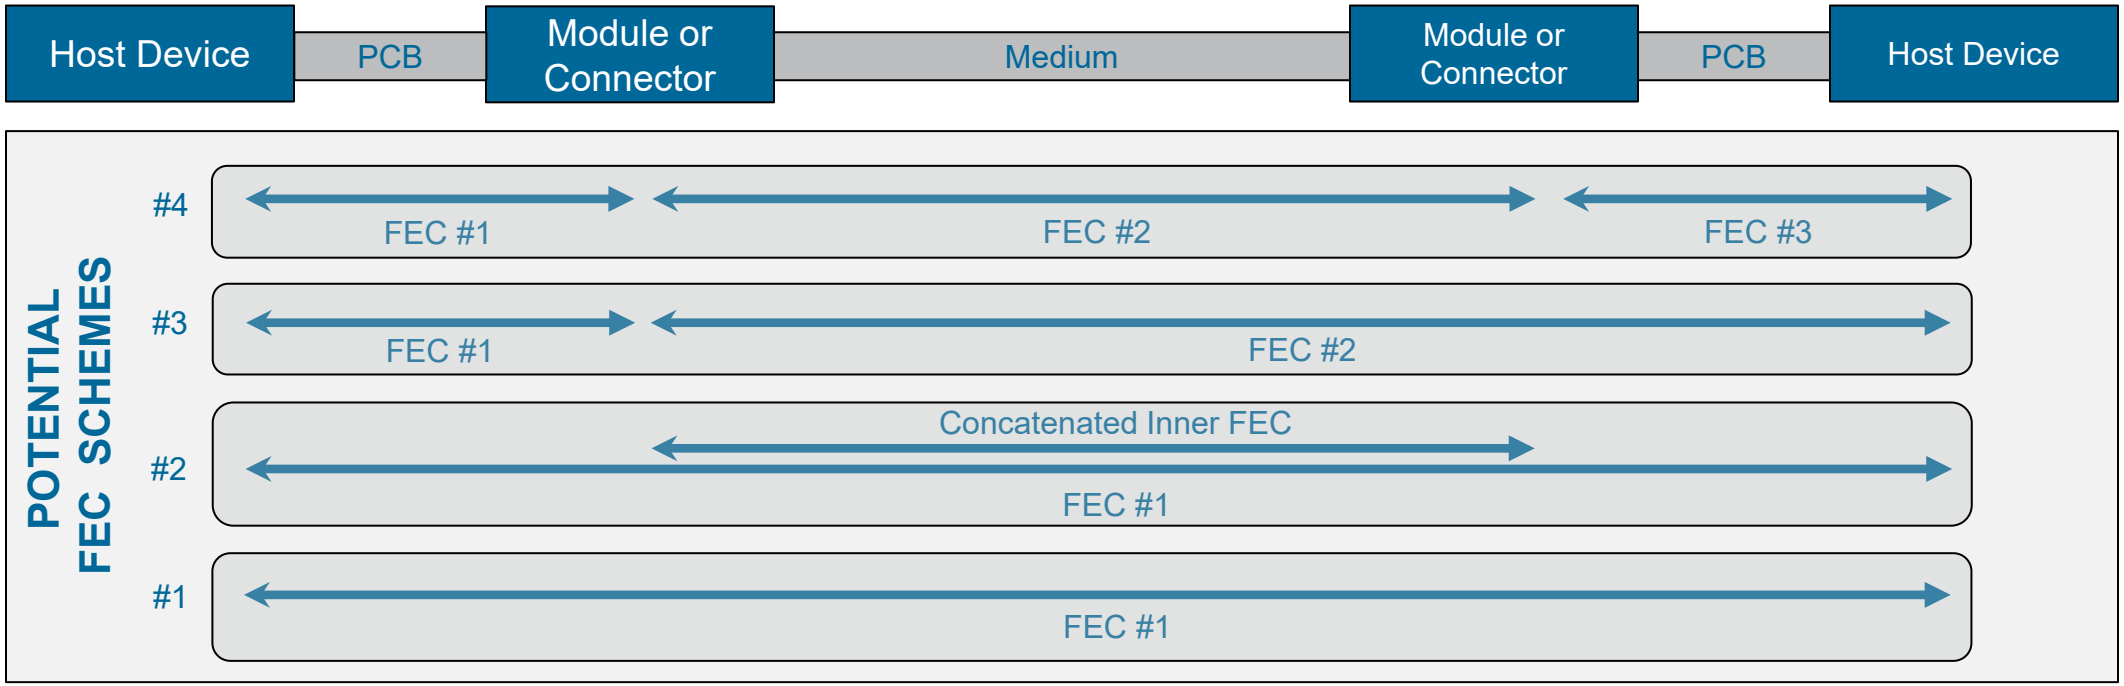

ARCHITECTURE

may need to enable one or more FEC schemes

Architecture Observations

- Taking an initial holistic approach to defining an architecture that will be capable of supporting one or more potential FEC schemes is key to minimizing risk of stranded ports in the future
- Architecture will be based on AUI's using different signaling rates
- The architecture needs to be flexible to address all physical layer specifications:
  - Leverage technology reuse!
- Variables to consider:
  - Power
  - Backwards / forward interoperability
  - FEC Coding Gain needed per physical layer specification
  - Latency

ADDRESSING MULTIPLE PHYSICAL LAYER SPECIFICATIONS

Adopted Physical Layer Objectives

Ethernet Rate	Assumed Signaling Rate	AUI	BP	Cu Cable	MMF 50m	MMF 100m	SMF 500m	SMF 2km	SMF 10km	SMF 40km
200 Gb/s	200 Gb/s	Over 1 lane		Over 1 pair			Over 1 Pair	Over 1 Pair		
400 Gb/s	200 Gb/s	Over 2 lanes		Over 2 pairs			Over 2 Pair			
800 Gb/s	100 Gb/s	Over 8 lanes	Over 8 lanes	Over 8 pairs	Over 8 pairs	Over 8 pairs	Over 8 pairs	Over 8 pairs		
	200 Gb/s	Over 4 lanes		Over 4 pairs			Over 4 pairs	1) Over 4 pairs 2) Over 4 λ's		
	TBD								Over single SMF in each direction	Over single SMF in each direction
1.6 Tb/s	100 Gb/s	Over 16 lanes								
	200 Gb/s	Over 8 lanes		Over 8 pairs			Over 8 pairs	Over 8 pairs		

Making it all work together

Technology Reuse

Leverage existing or work-in-progress 100 Gb/s per lane (e.g. 3cu, 3ck, 3db) to higher lane counts

Develop 200 Gb/s per lane electrical signaling for 1/2/4/8 lane variants of AUIs and electrical PMDs

Develop 200 Gb/s per optical fiber for 1/2/4/8 fiber based optical PMDs and 4 lambda WDM optical PMD

Potential for either direct detect and / or coherent signaling technology

MANAGING IEEE P802.3df

Looking Ahead…

- IEEE P802.3df is a big project…
- The adopted objectives have obvious groupings of objectives where **technology reuse is possible**.
- **Start together with an initial holistic approach** to considering the development of an architecture to support all of the project objectives is desirable.
- There is a potential that different objectives **may progress on different timelines.**

Looking closer at this potential scenario

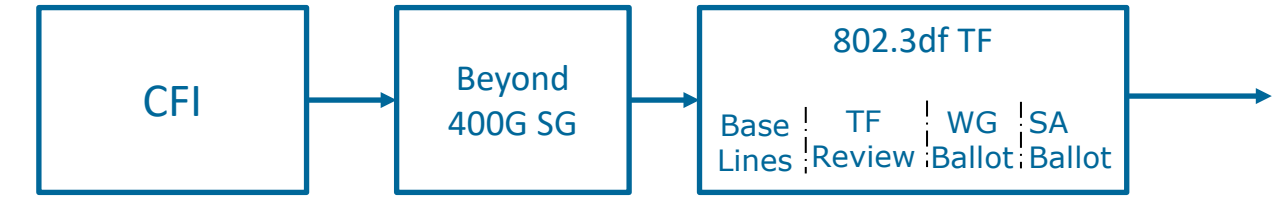

PAR & CSD Approval

Need to specify common foundational aspects before adopting baselines for all objectives to ensure consistency and interop (e.g. FEC architecture)

Will require significant analysis and technical decisions by the Task Force

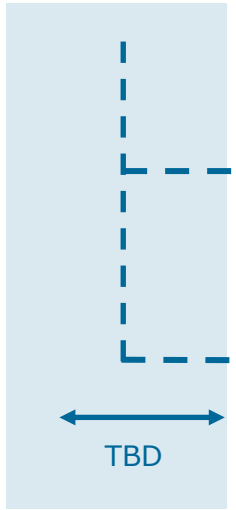

If some specifications need to take longer to specify, they are candidates to be handled by a subsequent Task Force(s)

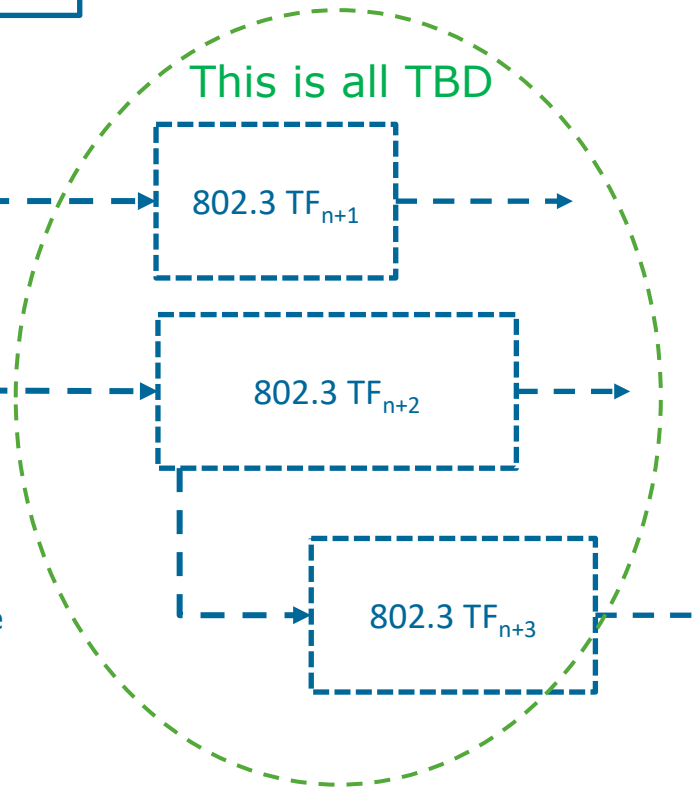

Possible objectives that might migrate include:

- 200 Gb/s based optical PMDs
- 200 Gb/s based copper PMD
- Coherent PMDs

Up to the 802.3df Task Force (or other future Task Force) and Working Group to decide

"Splitting the PAR" is a known process in 802

- PER IEEE 802 Operations Manual, Section 9.2 IEEE 802 LMSC approval

  At the discretion of the IEEE 802 LMSC Chair, PARs for ordinary items (e.g., Maintenance PARs) and PAR changes essential to the orderly conduct of business (e.g., division of existing work items or name changes to harmonize with equivalent ISO JTC-1 work items) may be placed on the IEEE 802 LMSC agenda if delivered to IEEE 802 LMSC members 48 hours in advance

- Most recent example: IEEE Beyond 10km Study Group
  - Initial PAR – 802.3cn (Mar 2018)
    - 50GBASE-ER, 200GBASE-ER4, 400GBASE-ER8, 100GBASE-ZR, 400GBASE-ZR
  - 802.3cn PAR split (Feb 2019)
    - 802.3cn (PAR Modification) - 50GBASE-ER, 200GBASE-ER4, 400GBASE-ER8
    - 802.3ct (New PAR) - 100GBASE-ZR, 400GBASE-ZR
  - 802.3ct PAR split (Feb 2020)
    - 802.3ct (PAR Modification) - 100GBASE-ZR
    - 802.3cw (New PAR) - 400GBASE-ZR

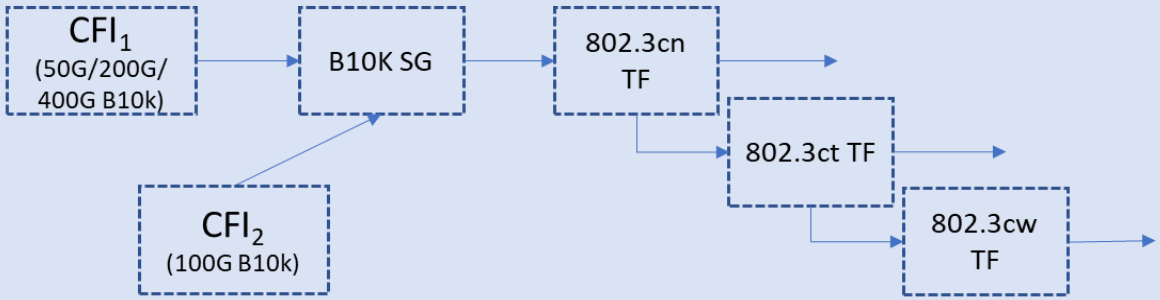

- 2nd fast-follower related CFI.
- Chartering motion merged into common SG
- Objectives agreed but concerns raised about schedule harmony
  - New TF's spun out of original to allows subsets of specs to complete while further work of more challenging objectives
- One(+) CFI → 3 TF's

Summary

- The proposed P802.3df project spans multiple Ethernet rates, signaling rates per lane, lane widths and media types
- **An initial holistic approach to start** is recommended to develop an architecture that would support one or more FEC schemes
  - This requires significant analysis and technical decisions by the Task Force
- The adopted Physical Layer objectives have obvious groupings where technology reuse is possible.
  - Many objectives are simply 1x, 2x, 4x, 8x lane variants
  - **If physical layer specifications evolve into different timelines, they can be managed**
- Time to get started!

Thank You!

QUESTIONS?