TEF 2025
Ethernet for AI

December 2-3, 2025
Hyatt Centric Mountain View, CA, USA

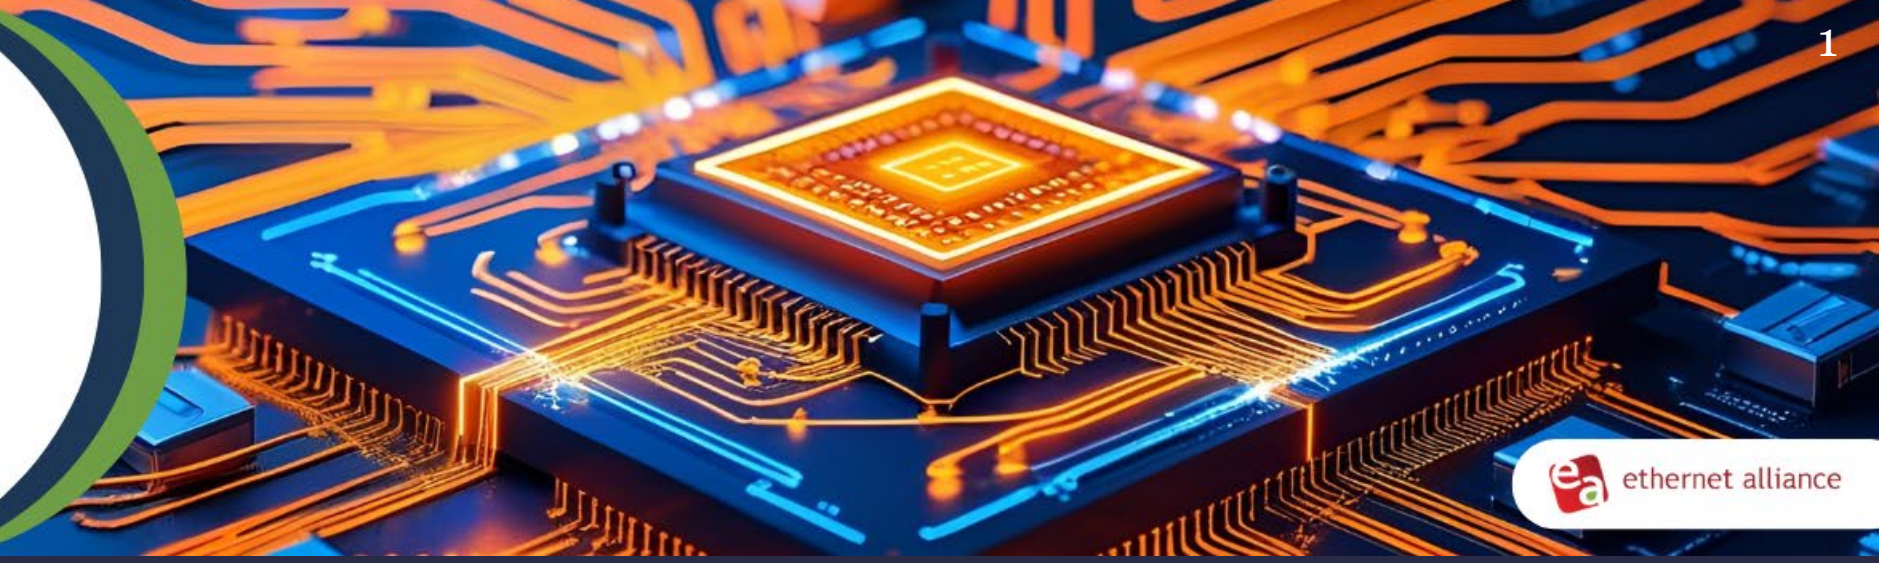

ethernet alliance

Ethernet for AI Isn't Just Evolving – Its Exploding!

John D'Ambrosia
Chair, IEEE 802.3 NEA "Ethernet for AI" Assessment,
Futurewei Technologies

December 2-3, 2025

This presentation has been developed within the Ethernet Alliance and is intended to educate and promote the exchange of information. Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance

Regarding the Views Expressed

My Industry Involvement

- Chair, IEEE 802.3 NEA “Ethernet for AI” Assessment Activity
- Ethernet Evangelist, High-speed Ethernet, Ethernet Alliance
- Chair, IEEE P802.3dj 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet Task Force
- Technical Vice-President, Futurewei Technologies, U.S. Subsidiary of Huawei
- Email – jdambrosia@ieee.org
- Follow me on LinkedIn

Per IEEE-SA Standards Board Bylaws

“At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE.”

Introduction

- TEF 2024 Launched industry wide discussions on 400 Gb/s signaling
- IEEE 802.3 launched the IEEE 802.3 New Ethernet Applications "Ethernet for AI" (E4AI) Assessment in January 2025
  - https://www.ieee802.org/3/ad_hoc/E4AI/index.html

Exploring Technical Feasibility

- Addressing “Ethernet for AI” will require
  - Understanding application / performance requirements and potential trade-offs
  - Technical work
    - Signaling rate
    - Modulation
    - FEC
    - Channels
- Any future effort in IEEE 802.3 will need to explore technical feasibility
- Exploration of technical feasibility may uncover items related to “broad market potential” or “economic feasibility” that need to be addressed

AI and IEEE 802.3

ARTIFICAL INTELLIGENCE & COMPUTE

- First Era (Before 2012)
  - Moore's Law – 2-year doubling
  - Uncommon to use GPUs for machine learning
- Modern Era (2012 and later)
  - 2012 – 2014: most results used 1-8 GPUs rated at 1-2 TFLOPS
  - 2014 – 2016: large-scale results used 10-100 GPUs rated at 5-10 TFLOPS
  - 2016 – 2017: greater algorithmic parallelism (huge batch sizes, architecture search, expert iteration), specialized hardware (TPUs), faster interconnects

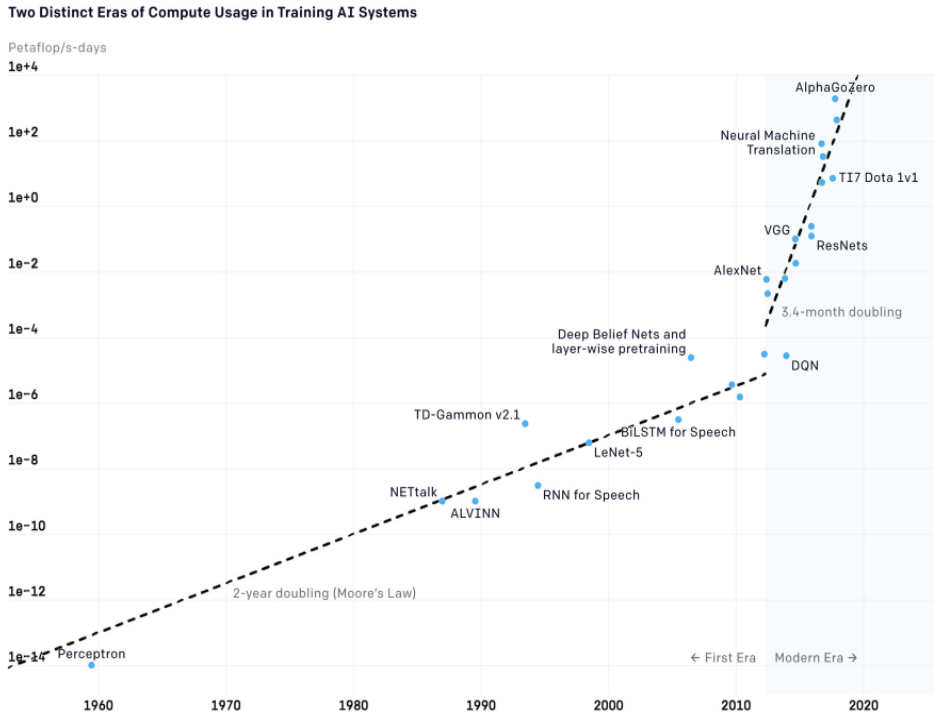

Source – OpenAI blog post 'AI and Compute' addendum 'Compute used in older headline results' posted 7th November 2019 by Girish Sastry, Jack Clark, Greg Brockman and Ilya Sutskever <https://openai.com/blog/ai-and-compute/>.

- AI was cited as a factor influencing bandwidth growth in the "Beyond 400GbE" CFI

Source: https://www.ieee802.org/3/ad_hoc/ngrates/public/calls/20_1029/CFI_Beyond400GbE_Rev7_201029.pdf

AI Application Requirements

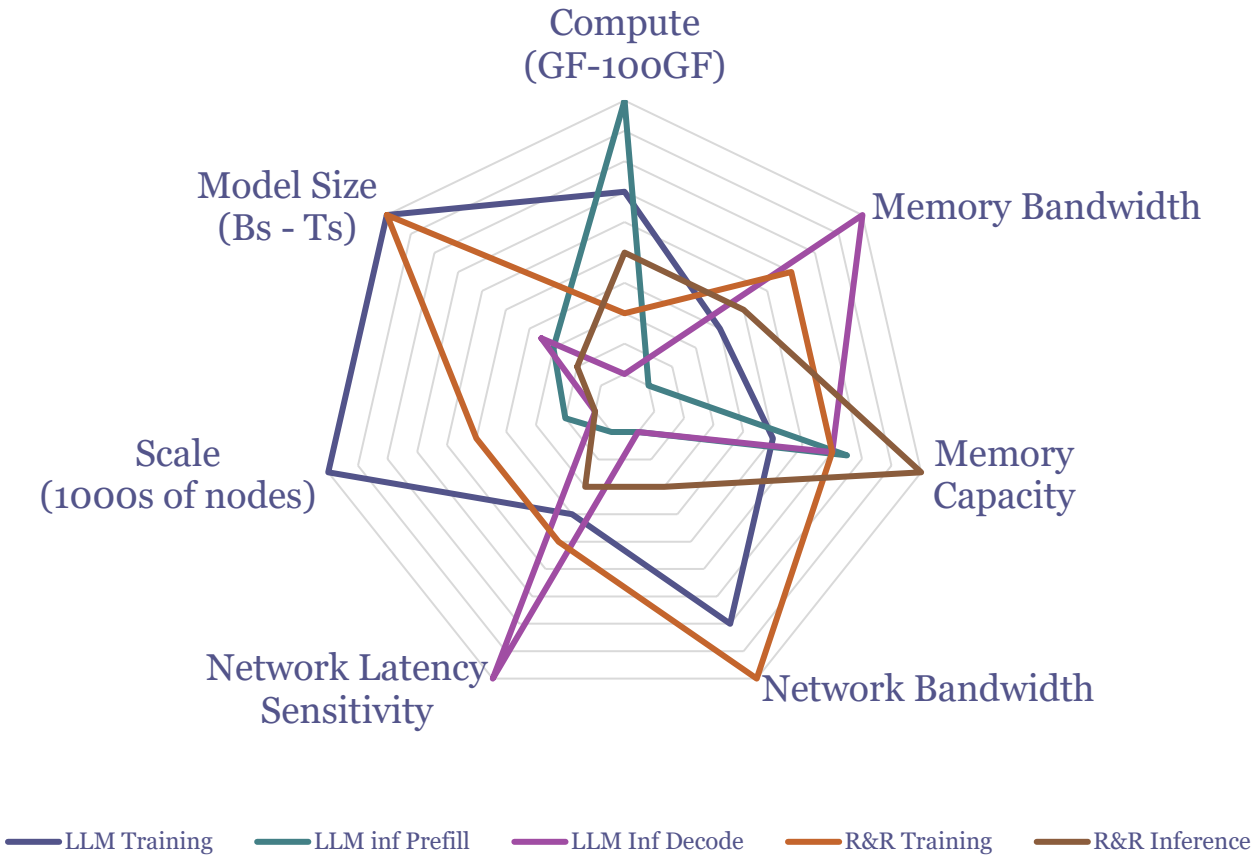

- Ethernet Priorities (2 – 5 years)
  - Resilience
  - Reach
  - Beachfront (I/O BW Limitation)
  - Power
- Latency can be traded off for resilience, reach and power

Source: Ethernet Alliance TEF 2024 Keynote – Ethernet – The foundation of AI @ Meta, Nic Viljoen

AI Datacenter Network Hierarchy

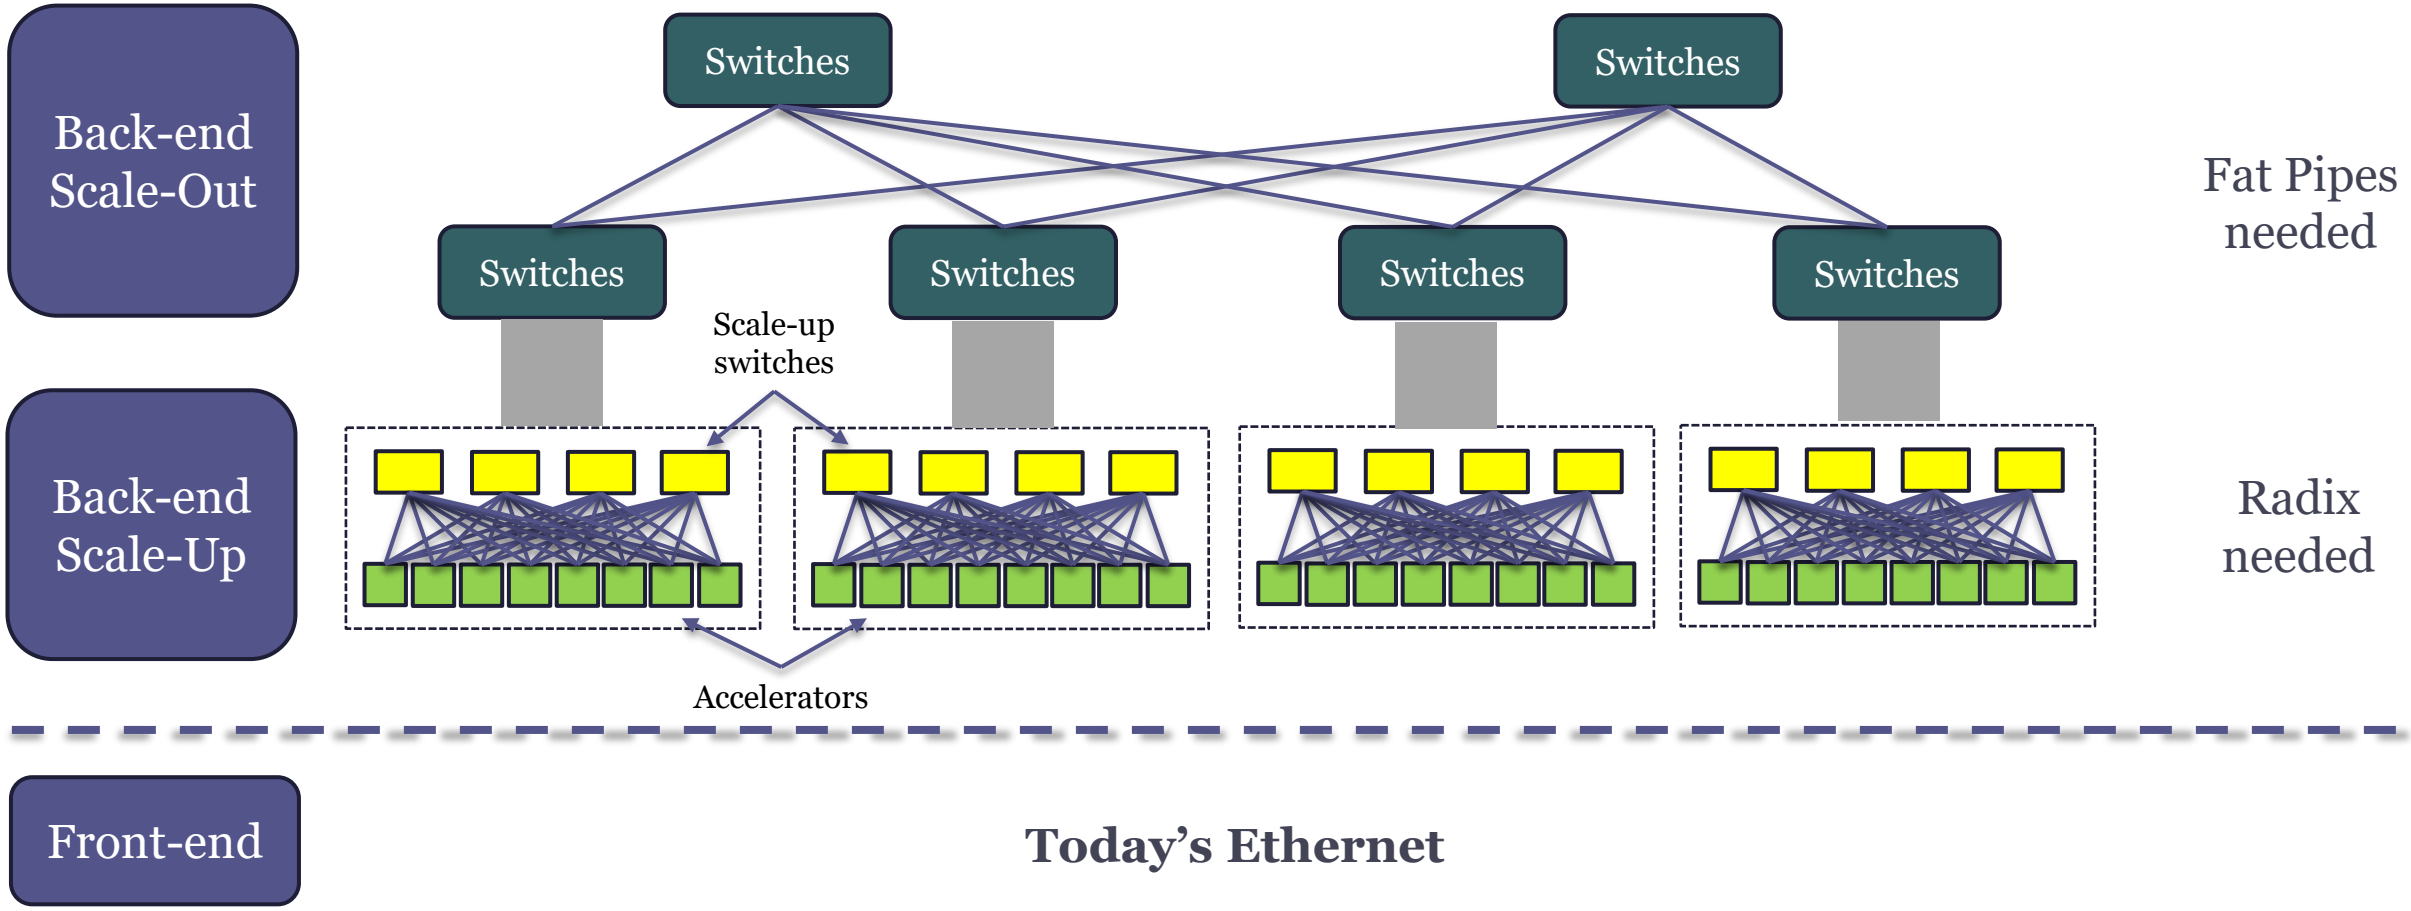

Based on AI Datacenters and their Diverse Network Requirements, Ram Huggahalli (Microsoft), Ethernet Alliance TEF 2024, Oct 2024.

Bandwidth Trends

DATA CENTERS CONTINUE AS A PRIMARY DRIVER

DC Traffic Continues to Grow Rapidly (Regular Servers)

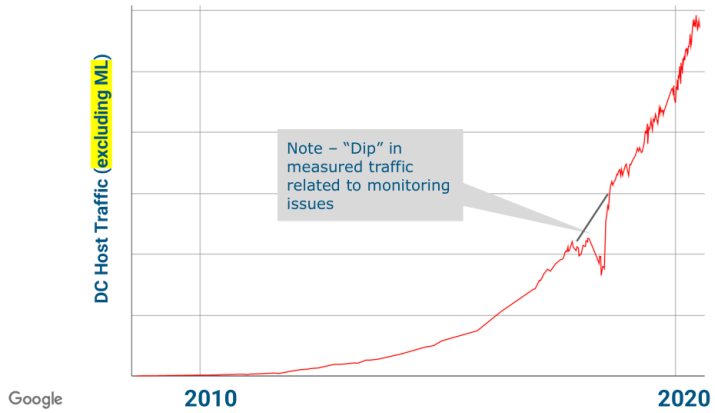

- Highlights growth in network fabric
- Excludes traffic due to machine learning

Courtesy - Cedric Lam, Google

Oct 2020 *

Introduction

- Future bandwidth requirements [1]
  - Front-end - 1x
  - Backend Scale-out - 10x
  - Backend Scale-up - 100x

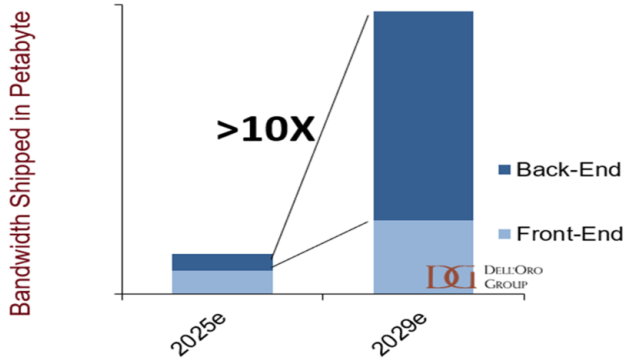

1. Alan Weckel, 650 Group, "AI Networking: What do scaleup and scaleout really mean for networking demand," 27 Mar 2025, https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0327/weckel_e4ai_01_250327.pdf.

- Source: Sameh Boujelbene, Dell'Oro AI Networks for AI Workloads Report, 2025
- Back-end includes scale-out only, excluding scale-up
- Back-end includes both Ethernet and InfiniBand

Aug 2025 **

TEF 2025 Ethernet for AI

* IEEE 802.3 CFI Consensus Presentation, D'Ambrosia, et. all, https://www.ieee802.org/3/ad_hoc/ngrates/public/calls/20_1029/CFI_Beyond400GbE_Rev7_201029.pdf, 29 Oct 2020.

** Aug 2025 Progress Update – Moving Towards Multiple CFIs, D'Ambrosia, Lusted, Nowell, https://www.ieee802.org/3/ad_hoc/E4AI/public/25_0819/dambrosia_e4ai_01_250819.pdf, 19 Aug 2025.

Lane Rate versus Ethernet Rate

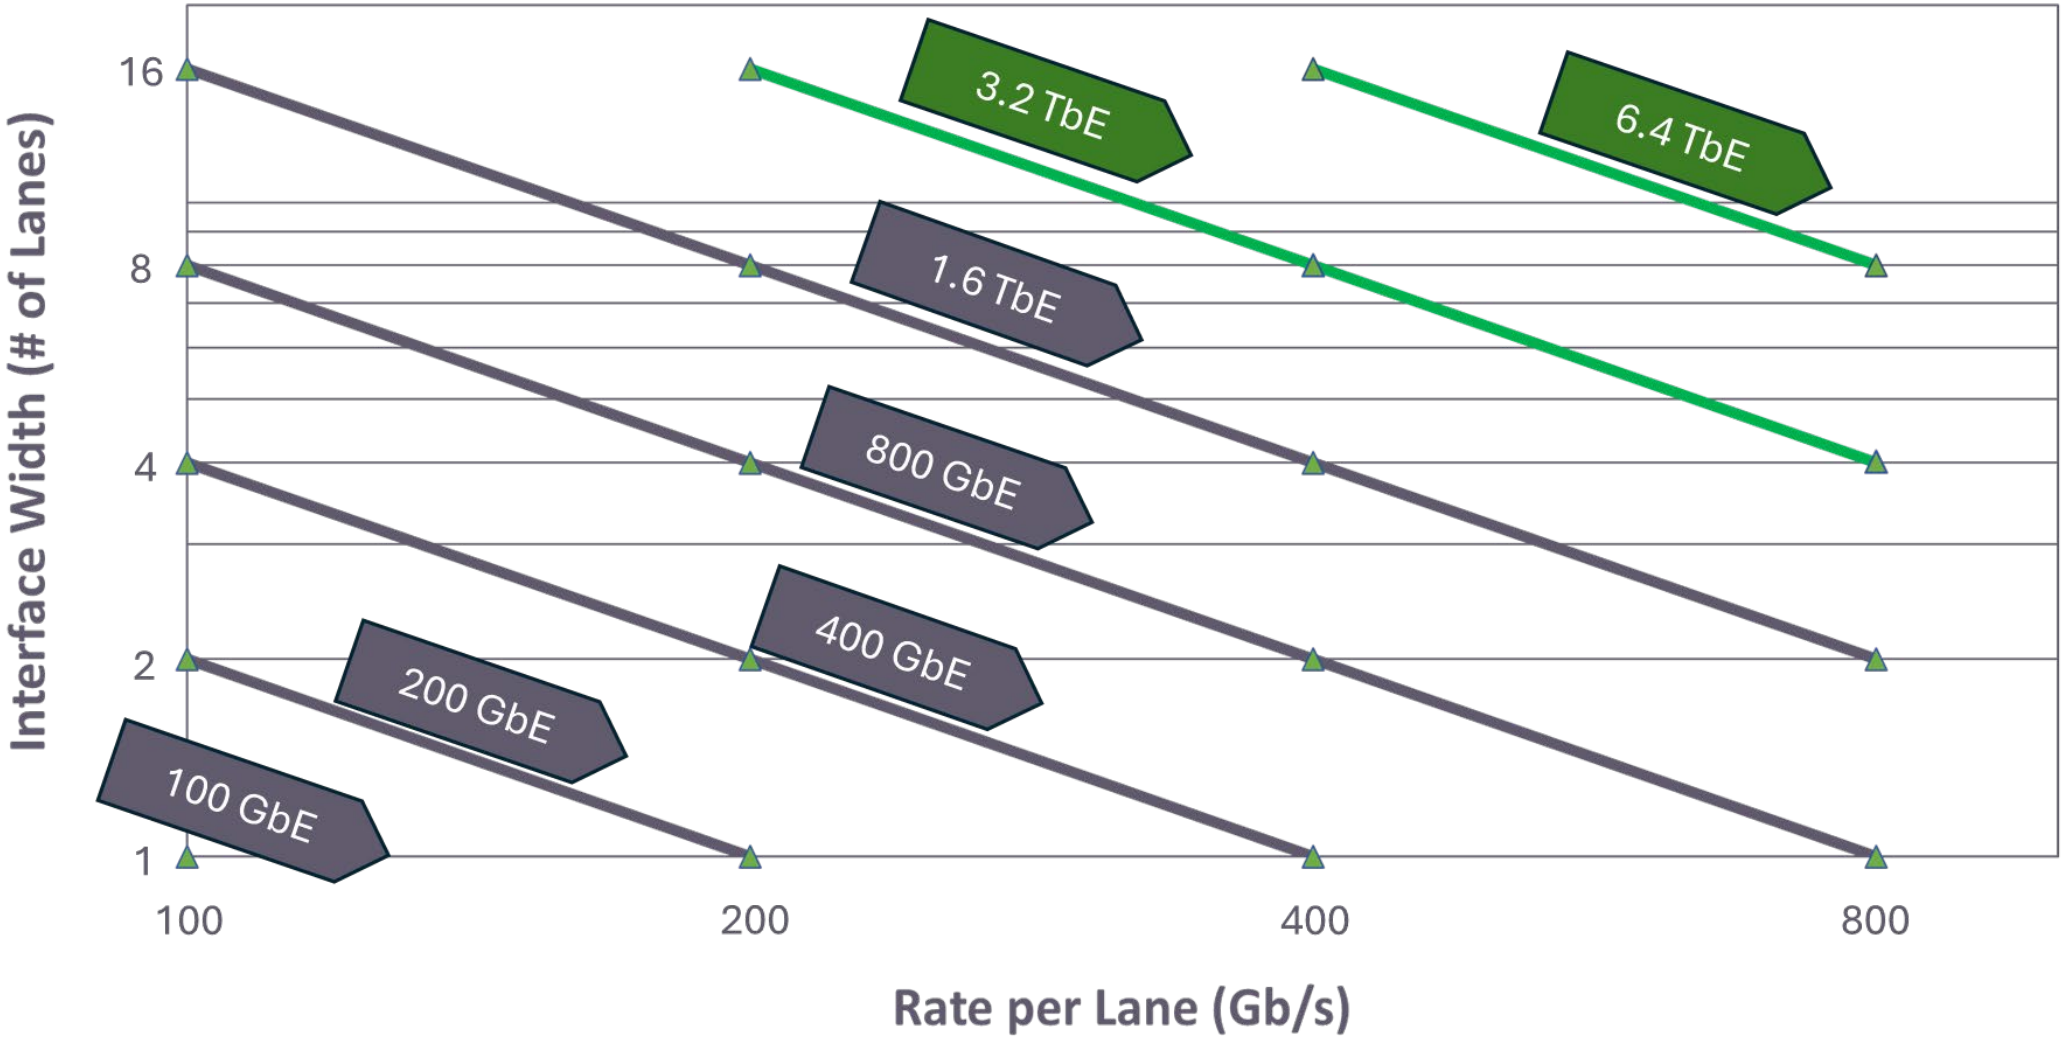

IEEE 802.3 and the Physical Layer

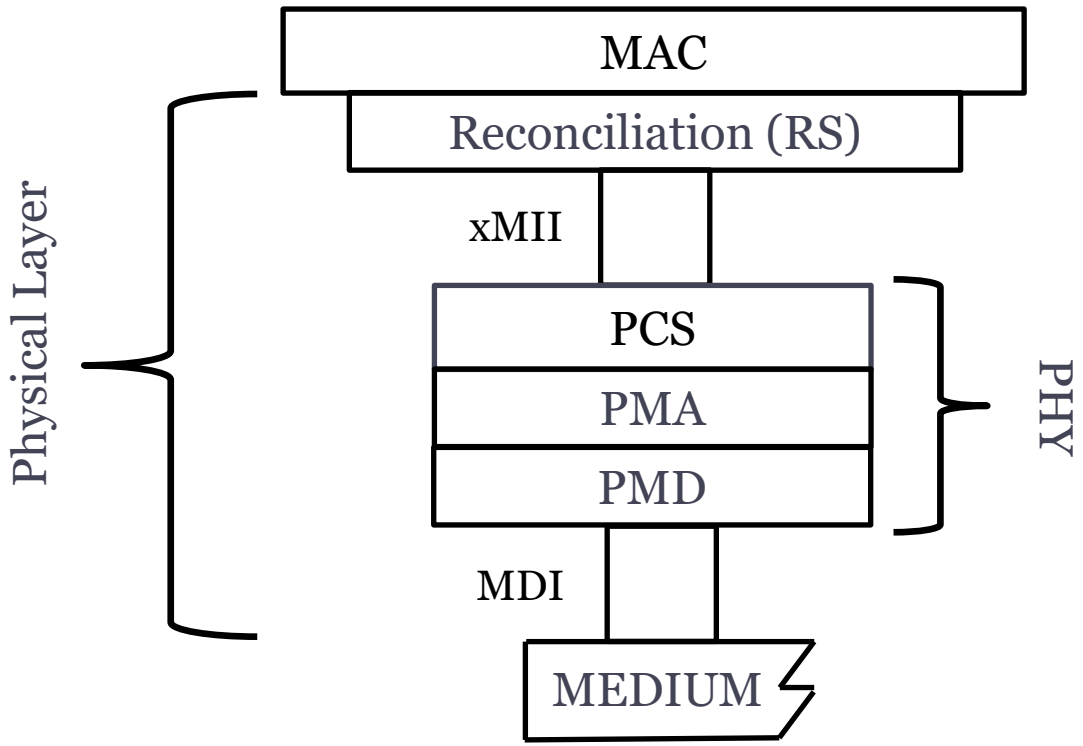

- IEEE 802.3 focuses on the Physical Layer, including
  - Modulation
  - Forward Error Correction (FEC)
  - Logic
- IEEE 802.3 E4AI Assessment
  - Explored modulation & FEC for electrical and optical signaling

Leveraging the IEEE P802.3dj Architecture

The IEEE P802.3dj architecture can support all of these FEC schemes:

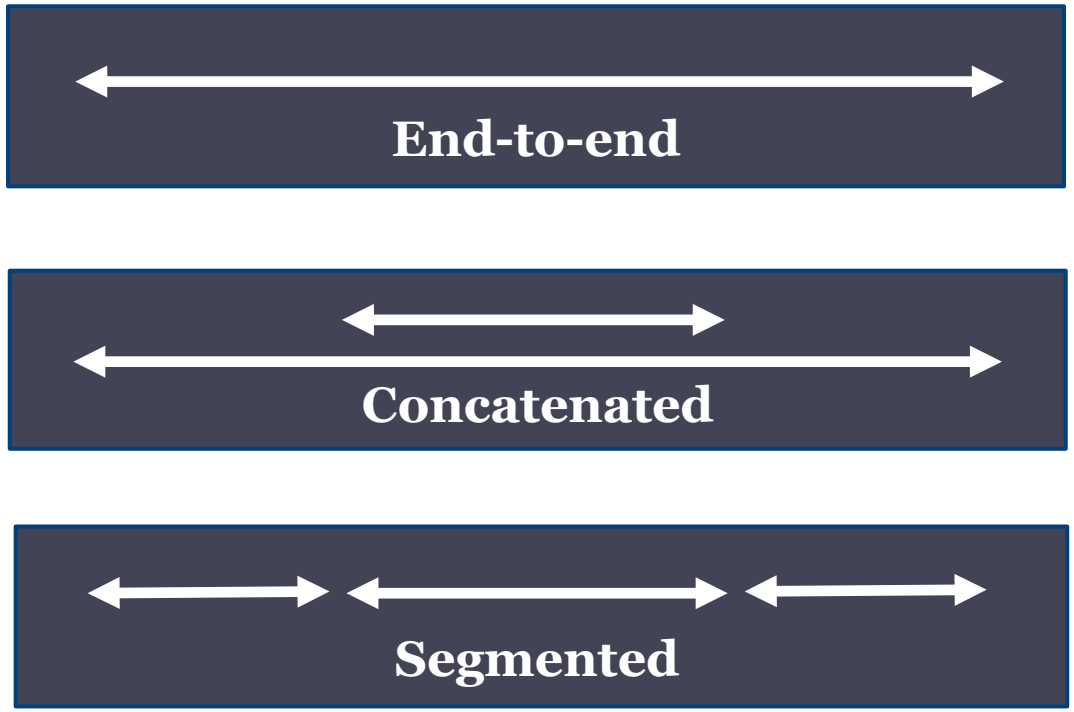

- **End-to-end FEC**
  - Not optimal for all targeted solutions
  - Lowest latency
- **Concatenated FEC**
  - Common “outer” FEC
  - Inner FEC – can be tailored on per solution basis
  - Reduced latency compared to “Segmented” approach
- **Segmented FEC**
  - Common “outer” FEC
  - Stronger internal FEC

Looking Upward

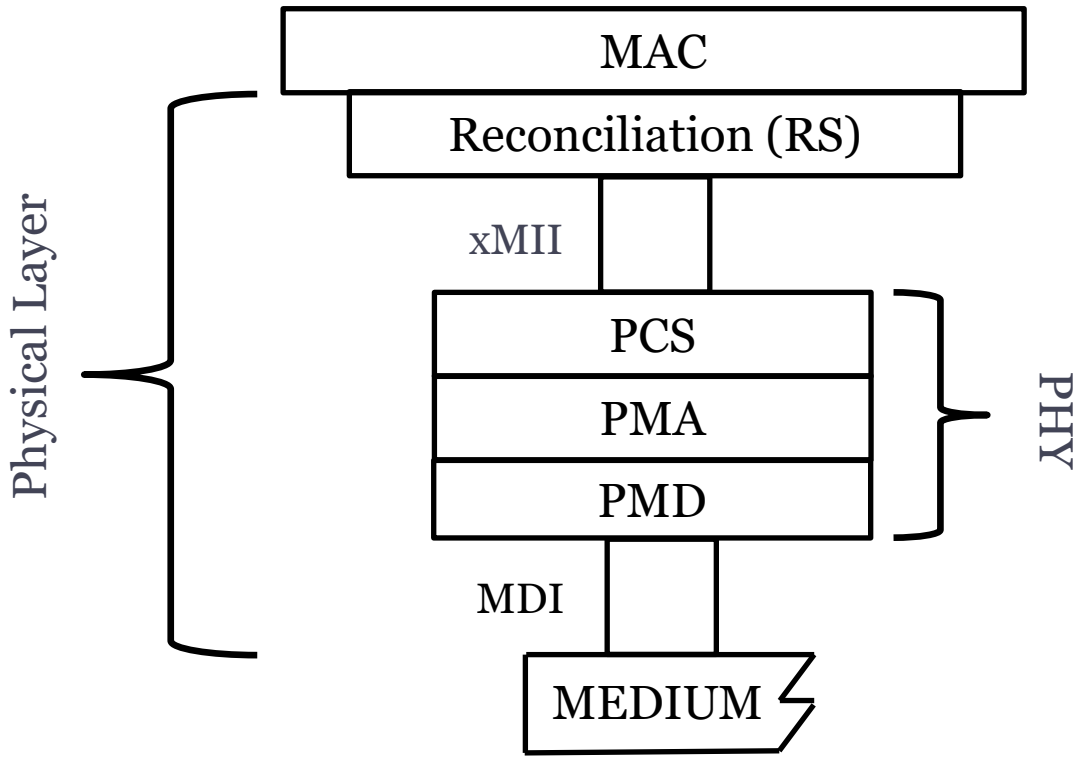

- IEEE 802.3 E4AI Assessment
  - Incubated effort to explore metadata service extensions for Ethernet
    - Support for features such as UEC's Link Layer Retry (LLR) and Credit-Based Flow Control (CBFC)
  - See IEEE EMS Study Group (IEEE P802.3dt)

Looking Down

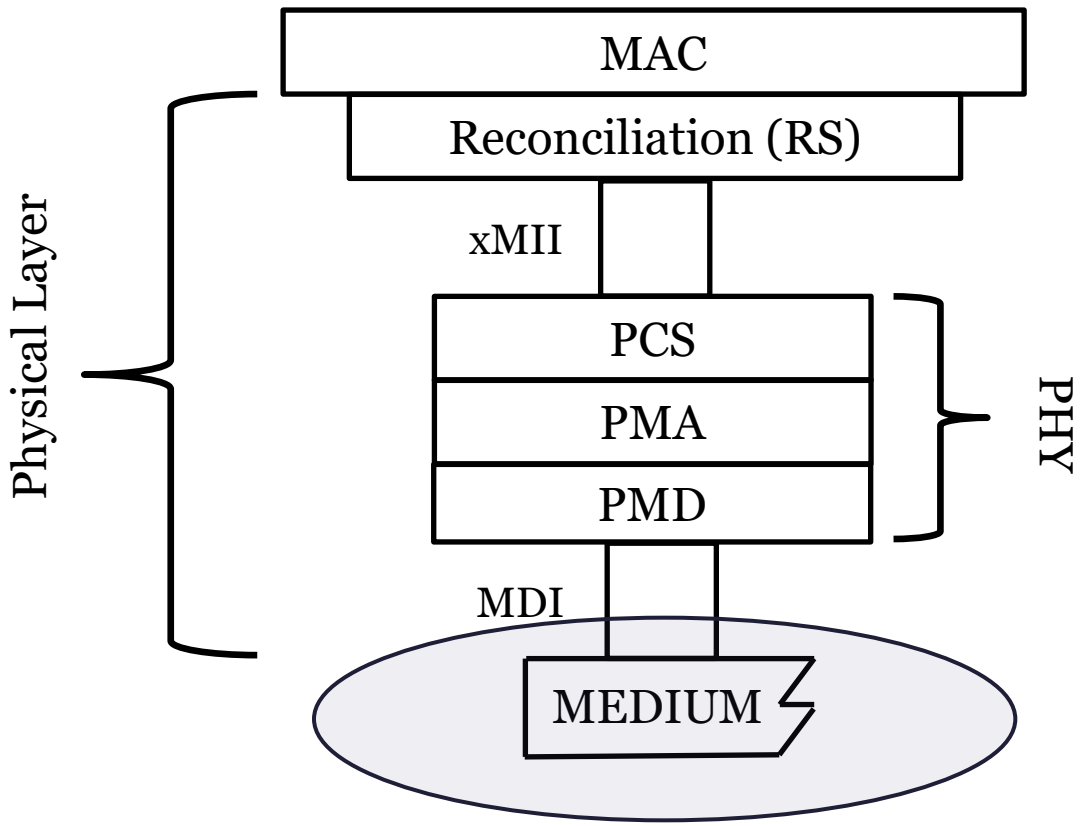

- IEEE 802.3 E4AI Assessment
  - Explored copper “channels”
    - Channel Library
  - Hosting “Fiber for AI” Workshop Teleconference (24 Feb 2026)
    - Future SMF / MMF
    - Multi-core
    - Hollow core
    - Etc

Consider this Trend….

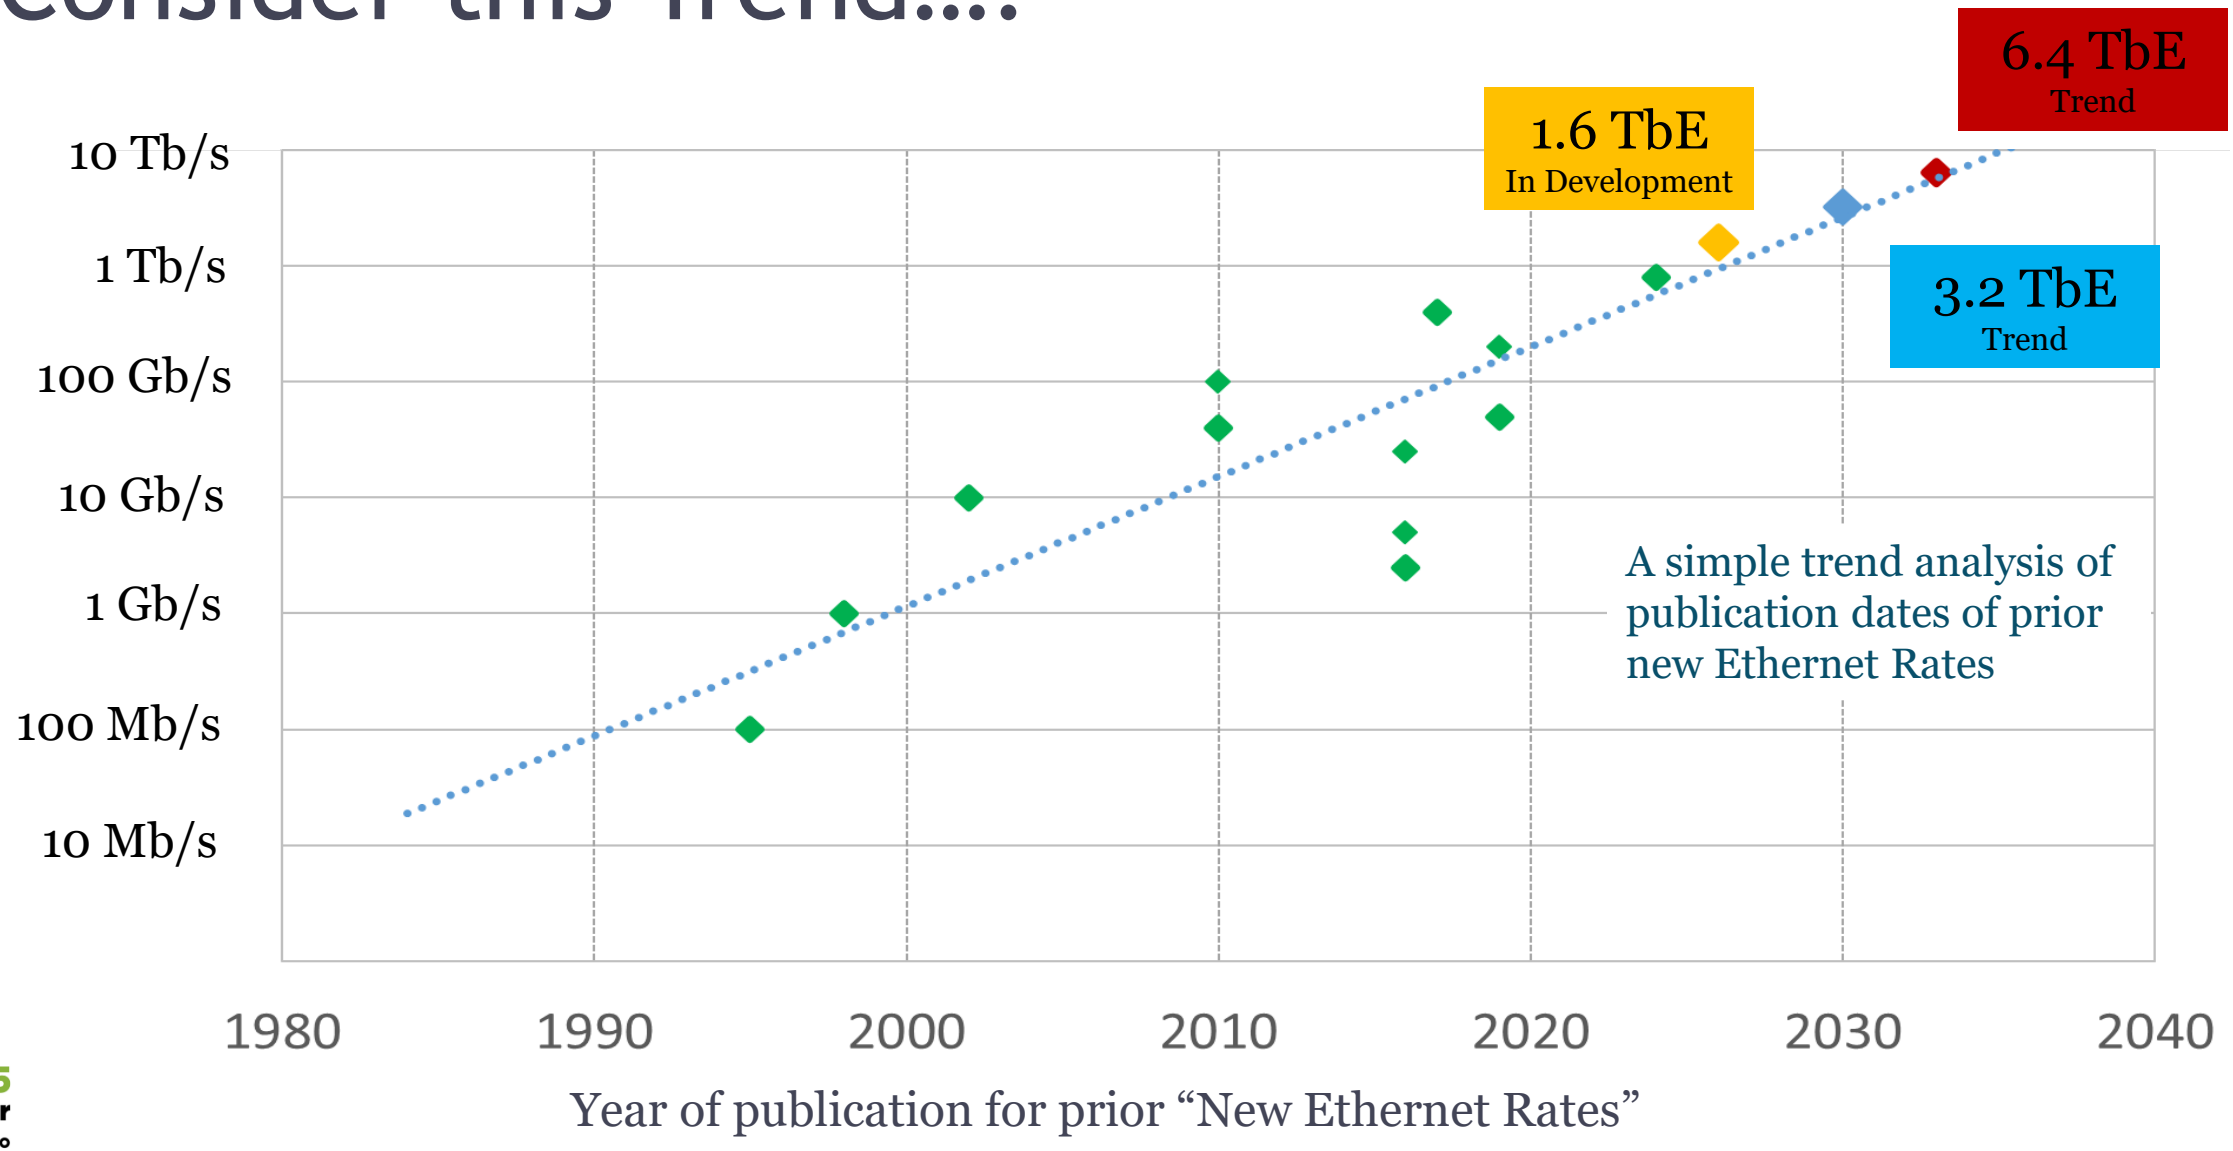

Front-end (FE) Vs Back-end (Scale-out) (BE)*

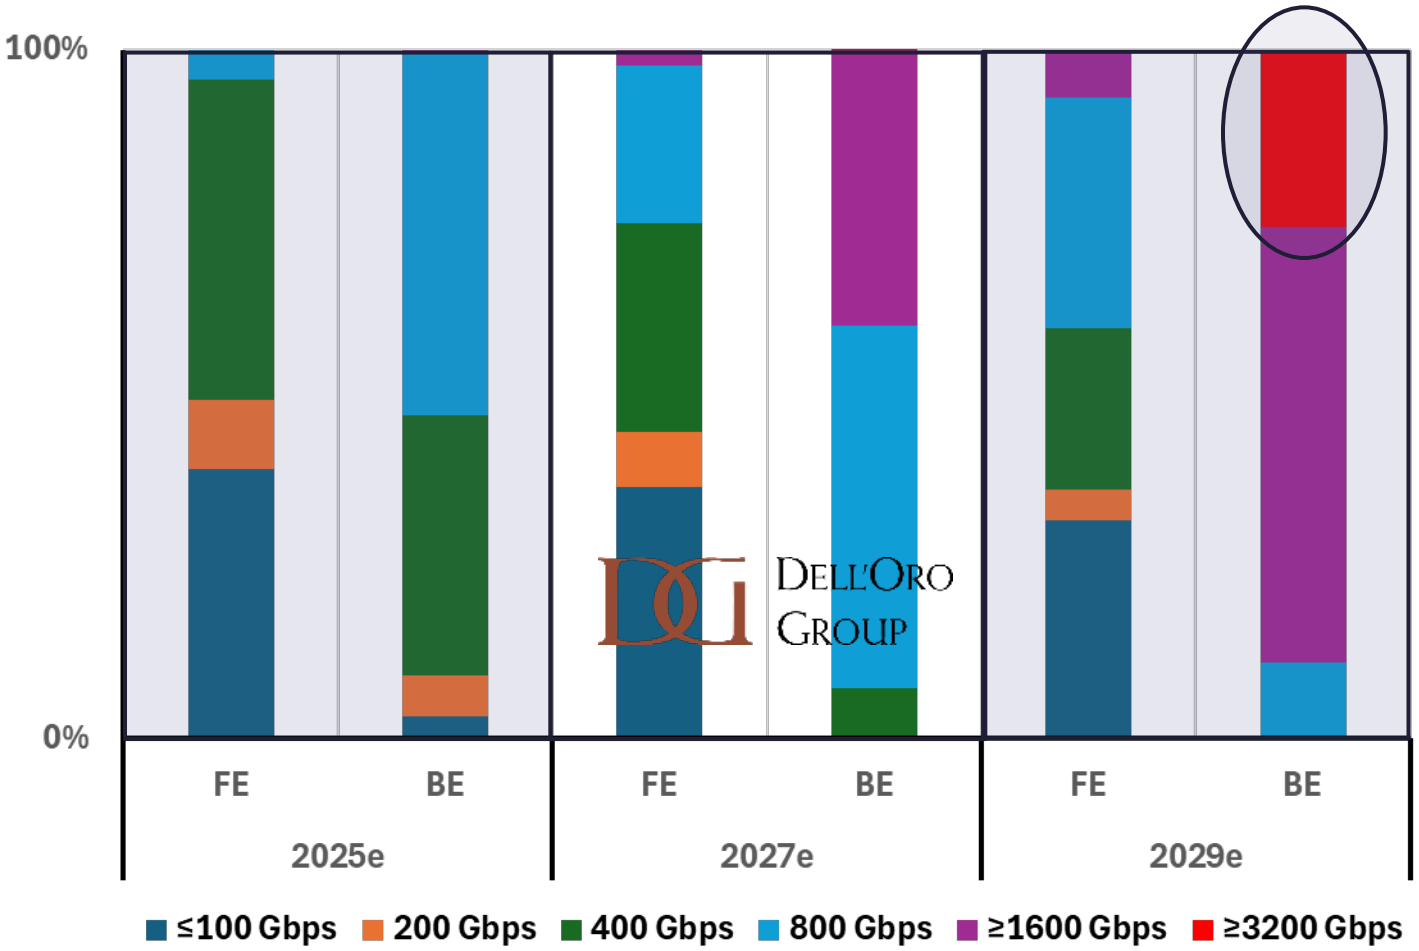

- Back-end scale-out networks will use higher capacity ports sooner and larger volumes than Front- End networks
- 2025 – Back-end networks : >50% using 800 Gbps
- 2027 – Back-end networks > 50% using 800 Gbps, significant ≥ 1600 Gbps penetration
- 2029 – Back-end networks >50% 1600 Tbps, ≥ 3200 Tbps penetration

Data provided by Sameh Boujelbene, Dell'Oro, Sept 2025.

* Back-end (Scale-out) include Ethernet and Infiniband

Dealing with Power Constraints

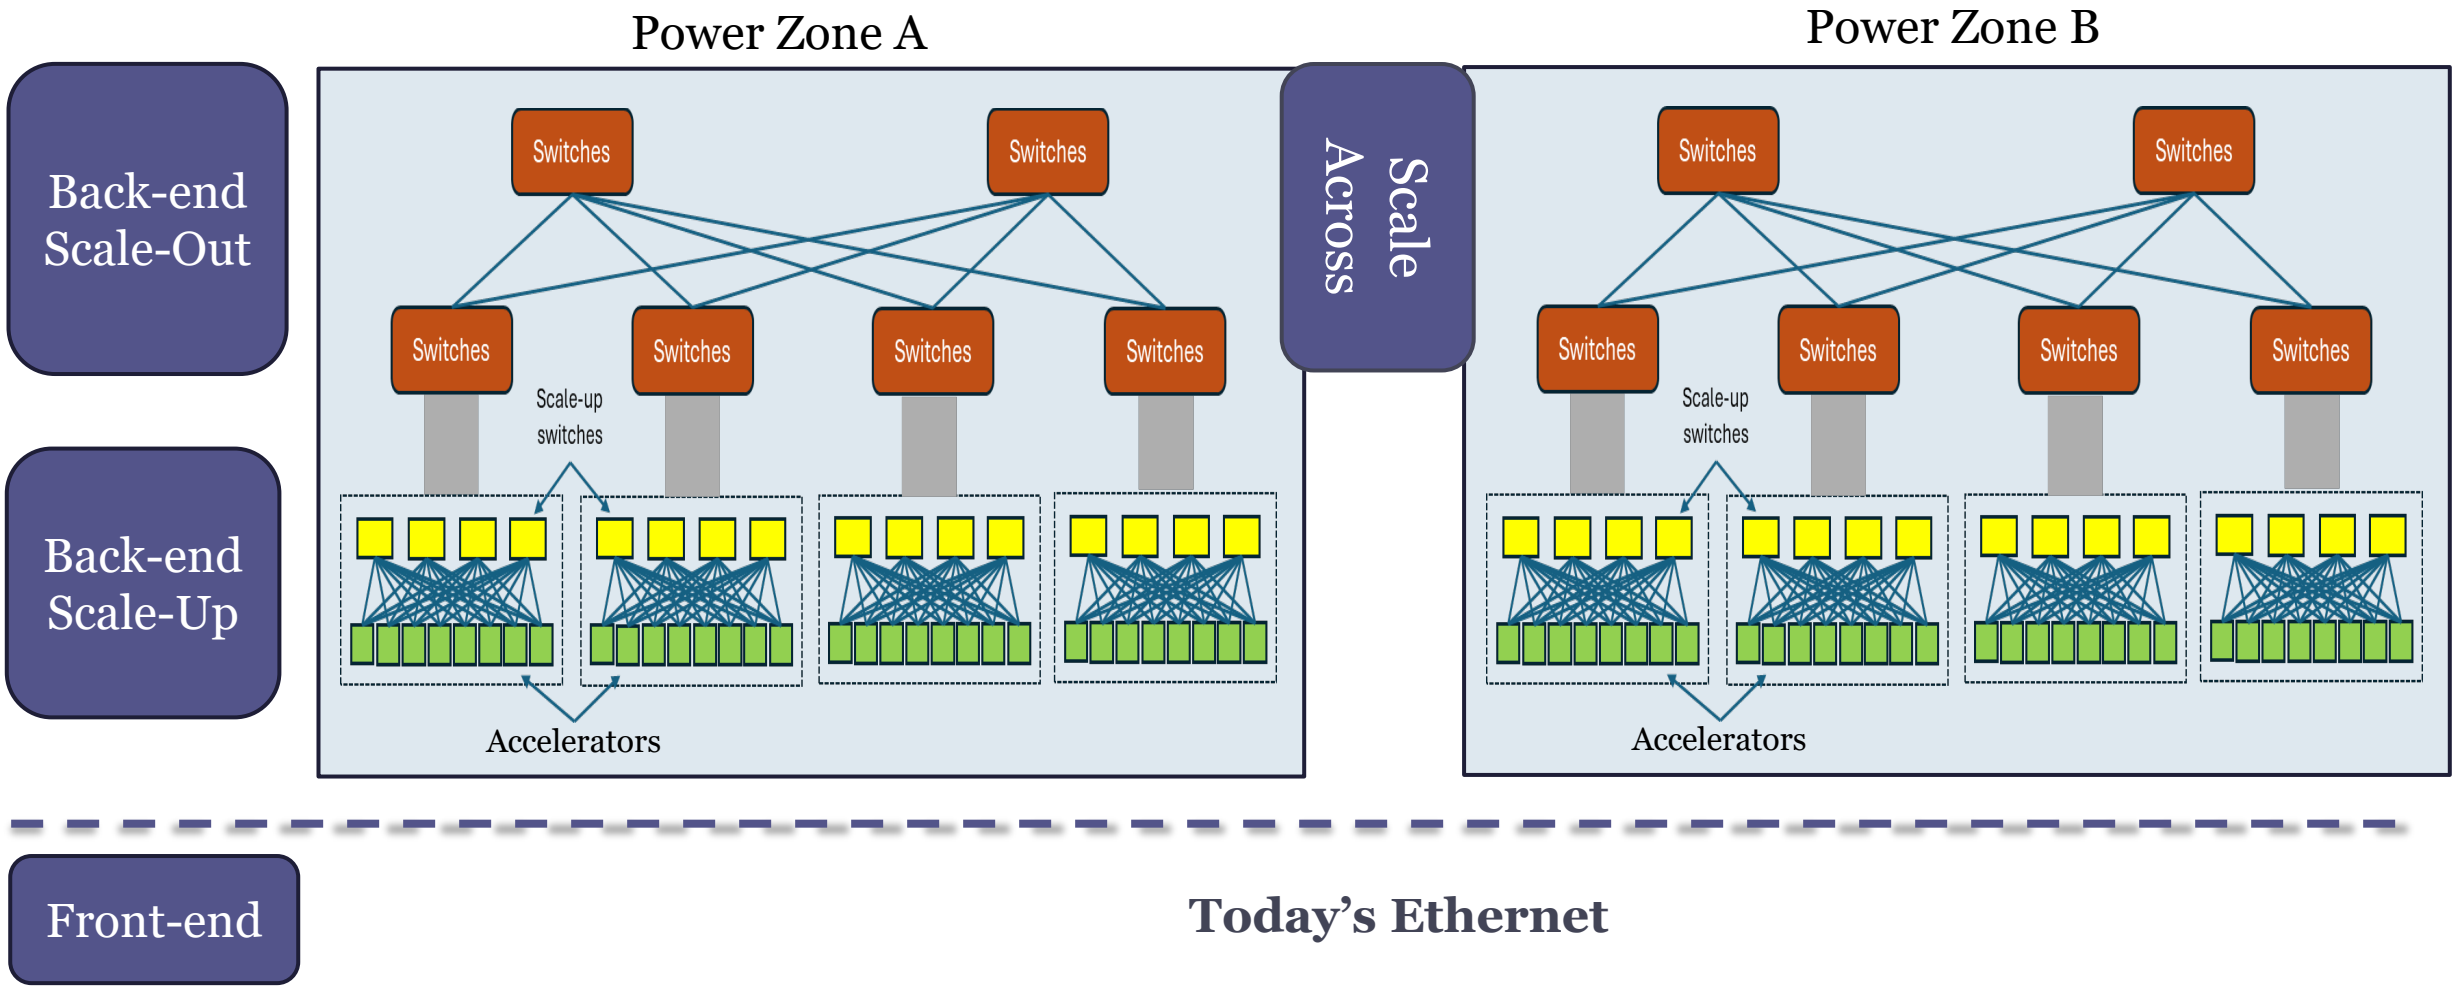

Time is of the Essence
Or How to Develop IEEE 802.3 Standards Quickly

- New “Higher Speed” Ethernet data rate projects
  - Longer overall duration
    - Study Groups have lasted a year or more
  - Targets more PHYs
  - Projects also developing new signaling technologies took **twice** the time in TF Draft development

- In general, PHY projects have taken significantly less time than new “Higher Speed” Ethernet data rate projects
  - Study Groups have lasted 4 months
  - Limited set of PHYs

- To enable rapid progress –
  - Focus on consensus building
  - Prioritize!
    - Develop the building blocks! (Signaling technologies)
    - Leverage those tools to then develop new “Higher Speed” Ethernet data rate solutions faster!

Moving Forward in IEEE 802.3

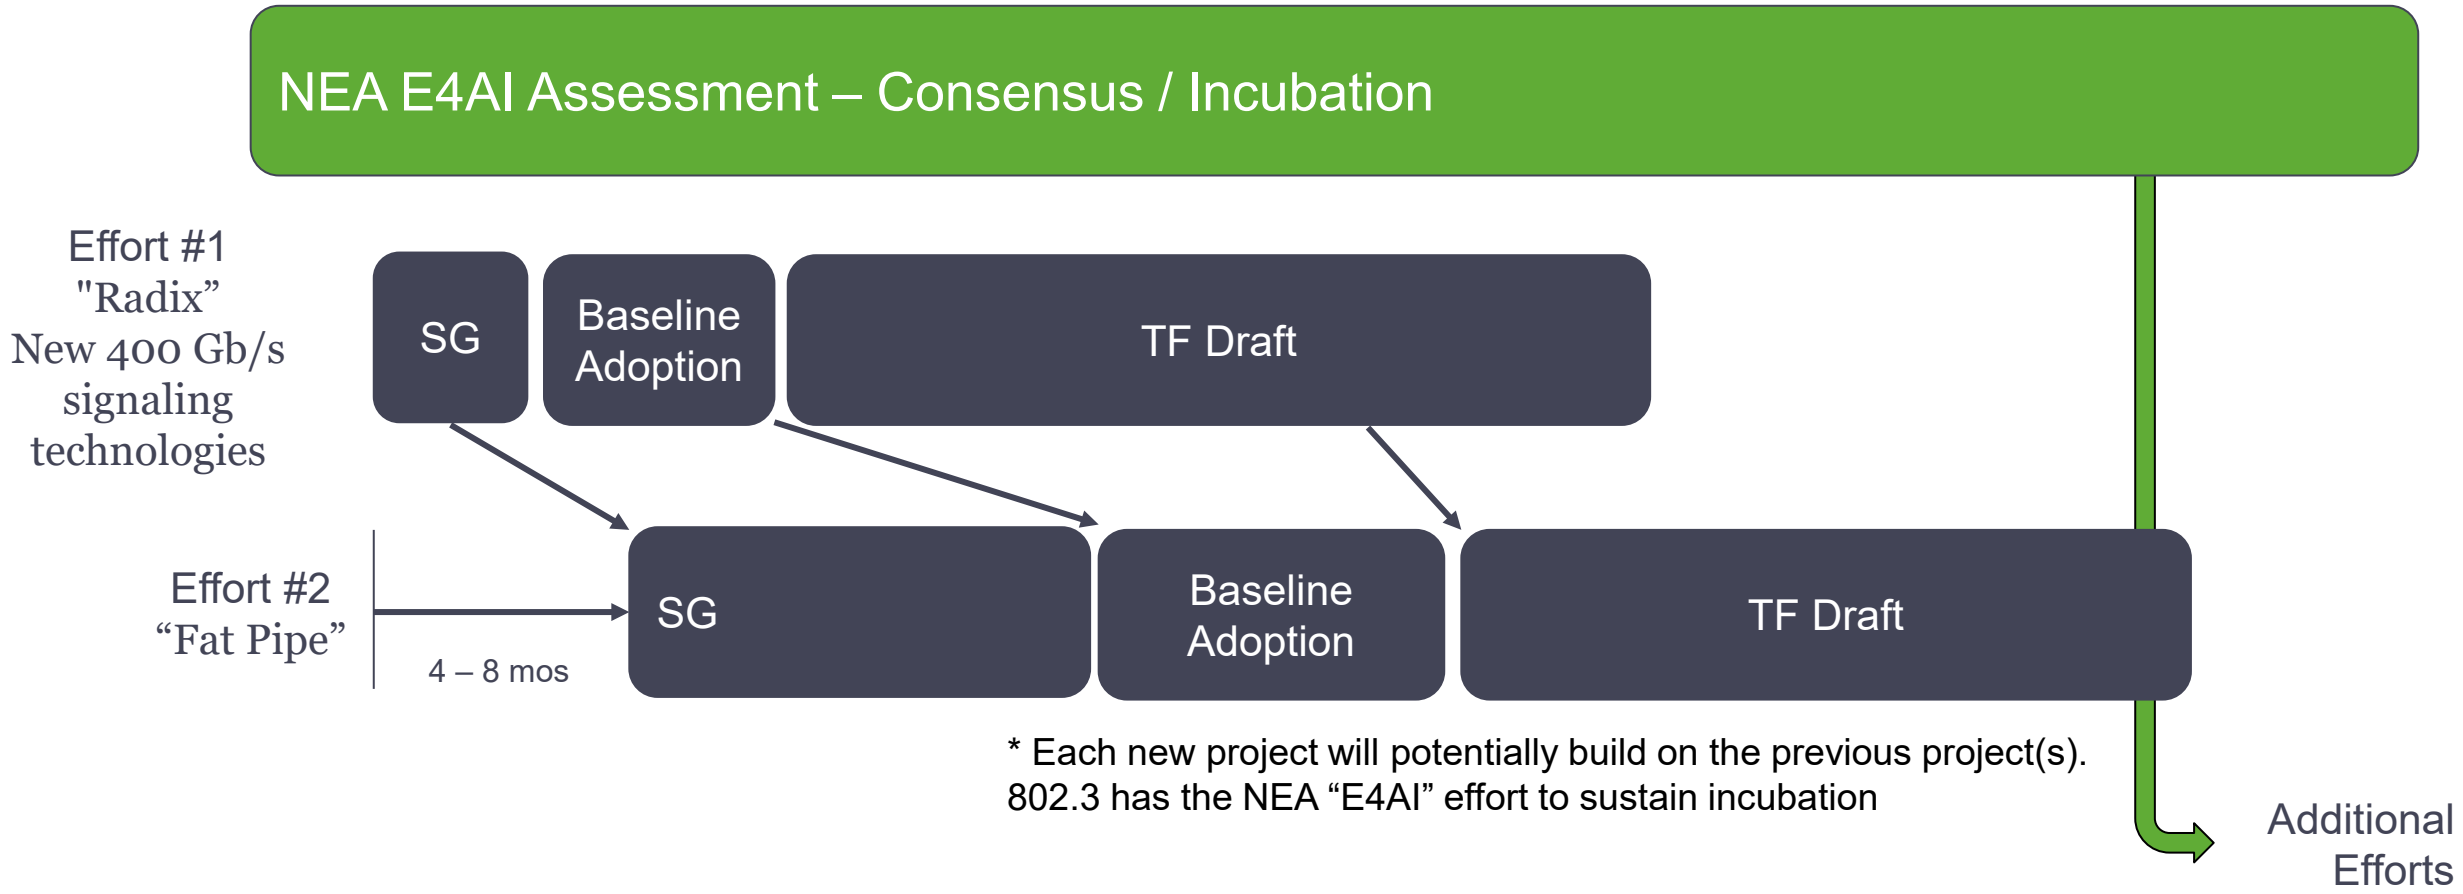

Summary

- Ethernet for AI represents a large growth in the Ethernet Ecosystem and is leading in terms of bandwidth, need, and volume over the traditional front-end
- Leaders in the IEEE 802.3 community are leveraging their experience to propose a faster overall approach to satisfy these needs
- The initiation of a new effort to develop 400 Gb/s electrical and optical signaling in IEEE 802.3 has been requested
- The industry needs to focus on building consensus!!!

Thank You!

Please see the following for further information:

- IEEE 802.3 400 Gb/s CFI Request
  https://www.ieee802.org/3/cfi/request_0326_1.html
- Ethernet Allliance TEF 2024: Ethernet in the Age of AI
  https://ethernetalliance.org/tef-2024-ethernet-in-the-age-of-ai/
- IEEE 802.3 NEA “Ethernet for AI” Assessment
  https://www.ieee802.org/3/ad_hoc/E4AI/index.html
- IEEE P802.3dj 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet Task Force
  https://www.ieee802.org/3/dj/index.html

QUESTIONS?