The 2015 Ethernet Roadmap for Networked Storage

June 30, 2015

Presented by The Ethernet Alliance and SNIA

Webcast Presenters

Scott Kipp, President, Ethernet Alliance - Brocade

David Fair, SNIA –ESF Chair – Intel

David Chalupsky, Chair, IEEE P802.3bq/bz TFs, Ethernet Alliance BASE-T Subcommittee - Intel

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Why The Storage Industry Needs To Pay Attention To The Ethernet Roadmap

The Good Old Days

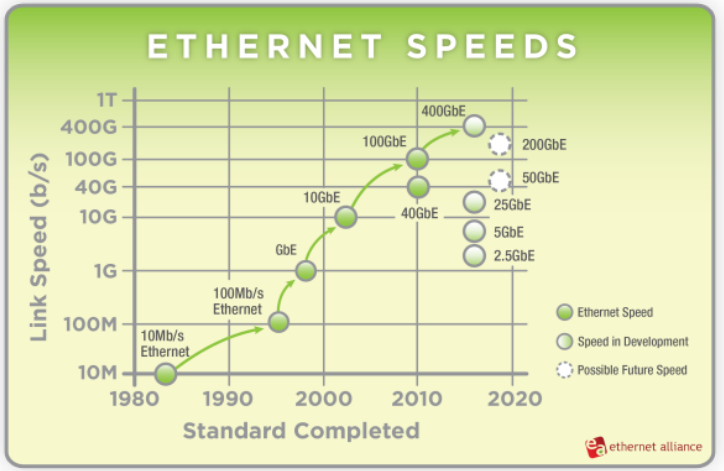

- The Ethernet community used to bump the speed by a Power of 10
  - "Build it and they will come"
  - Whenever the storage community discovered they needed faster
  - Ethernet network technology, Ethernet was waiting
  - Made a lot of sense given the cost of upgrading network infrastructure
- Storage media speed increases weren't happening that fast anyway
- Life was simple and good (relatively speaking)

Flash – Plan For The Disruption

SNIA Ethernet Storage Forum

- SNIA-ESF warned you in our 2012 webcast, *Flash – Plan For The Disruption*
  - Flash is ~100 times faster, ~40 times better in $ / IOPS, and ~600 times better in power / IOPS (2012)
- "A little dab will do you"
  - You don't need an all-flash array to see the benefit
  - Adding less than 1% capacity in the form of tiered flash can increase IOPS by more than 25%
- 3D flash will become a secondary disruptive wave
  - "The advent of 3D NAND has become a game-changer for the storage industry by increasing SSD capacity and dropping SSD prices…. Once parity [to HDDs] is achieved [in 2016], the transition to SSDs will become a tsunami. " - Jim O'Reilly, *Network Computing*, June 19, 2015

For Example…

- A single enterprise SSD available today can do sequential reads at 2.8 GB/s
- 22.4 Gbps is faster than what even a 10GbE adapter can support by more than a factor of two
- Network infrastructure isn't keeping up

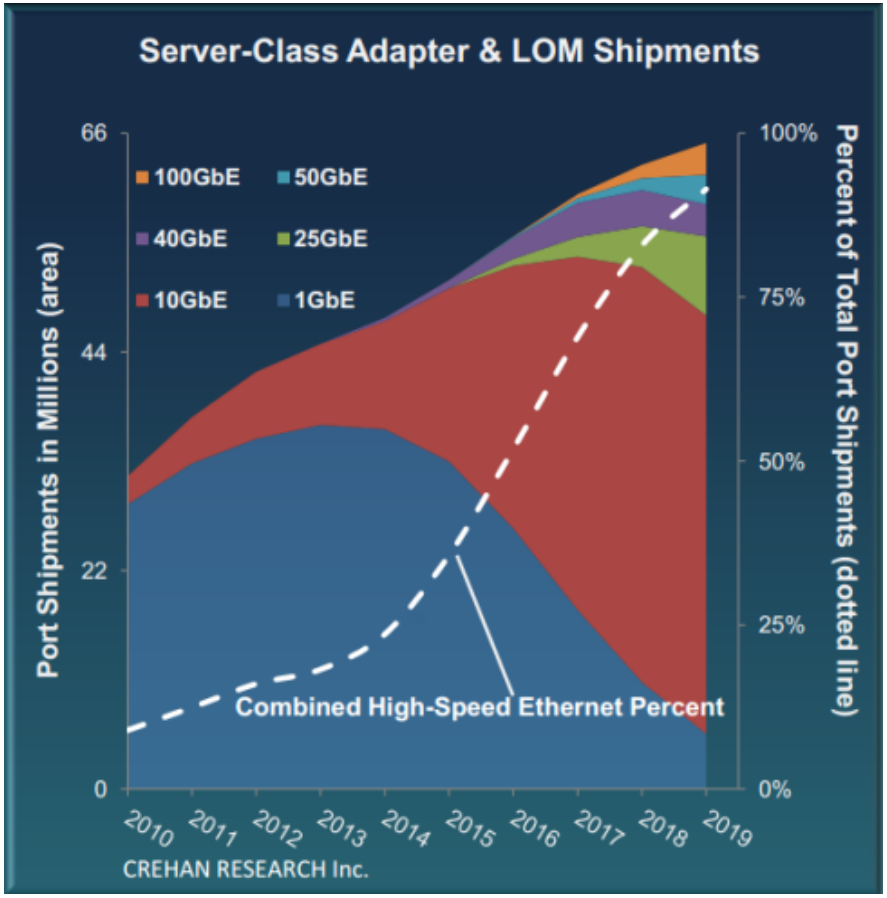

But It's Worse Than That: Data Is Exploding

SNIA Ethernet Storage Forum

- Expect the amount of data that needs to be stored and networked to increase by ~50% per year

- Not only do you need a faster network, you need more of it

The Future Of Storage Networking Depends On Ethernet

SNIA Ethernet Storage Forum

- The time has come for the storage community to attend to the Ethernet Roadmap and hopefully participate in influencing it
- These trends we see in storage are only a part of what is shaking up the Ethernet community

- SNIA-ESF is pleased to share with you the Ethernet Alliance's "2015 Ethernet Roadmap"
  - Scott Kipp, President, Ethernet Alliance
  - Dave Chalupsky, Chair, IEEE P802.3bq/bz TFs, Ethernet Alliance BASE-T Subcommittee

THE 2015 ETHERNET ROADMAP

Scott Kipp and Dave Chalupsky

June 30, 2015

ethernet alliance

The 2015 Ethernet Roadmap

SNIA™ Ethernet Storage Forum

What is it?

- 5,000 printed maps
- 500 T-shirts
- Many articles
- Many videos
- Pdf of map
- Whitepaper
- Graphics – use all the graphics in this presentation… From our website

Where is it?

- Tradeshows like OFC, Ethernet Technology Summit, Supercomputing, ECOC
- www.ethernetalliance.org/roadmap/

2015 ETHERNET ROADMAP

Front Side of Map

As shown on the long and winding road, Ethernet could have 12 speeds before 2020 with 6 new speeds introduced in the next 5 years. The progression of speeds is not in chronological order because 40GbE and 100GbE were primarily based on multiple lanes of 10Gb/s technology that was available before 25Gb/s serial technology enabled 25GbE. Lanes running at 25Gb/s are becoming practical in 2015 and will be used in 25GbE SFP+ and 4x25Gb/s 100GbE QSFP28. The next serial lane speed is expected to be 50Gb/s and enable 50GbE SFP28, 200GbE QSFP28 (4X50G) and 400GbE CFP2 (8X50G).

Beyond 400GbE, the map shows the unknown distant future that will become clearer as we approach 2020. Terabit links are expected when single lanes can be modulated at 100Gb/s and grouped into 10 or 16 lanes to form TbE or 1.6TbE. Significant investments in technology are needed before 100Gb/s lanes are economically feasible.

Low cost 100Gb/s lane technology that can fit in an SFP+ is not expected to be available until after 2020. The Ethernet Alliance will award the first company that produces a 100GbE SFP+ with the Holy Grail of the 100GbE SFP+.

The twisted pair or BASE-T roadmap in the lower right corner of the map shows how 10GBASE-T technology is being used in 4 new speeds — 2.5, 5, 25 and 40Gb/s. All four of these speeds are expected to be standardized in 2016 but they are targeting different cabling infrastructure. 2.5 and 5GBASE-T are being designed for Cat 5e cabling up to 100 meters while 25 and 40GBASE-T are being designed for 30 meters of Cat 8 cabling.

10 Tb/s
1.6 Tb/s >2020
6.4 Tb/s >>2020
1 Tb/s >2020
800 Gb/s >2020
400 Gb/s 2017 (est)
200 Gb/s 2018-2020 (est)
100 Gb/s 2010
50 Gb/s 2018-2020 (est)
40 Gb/s 2010
25 Gb/s 2016 (est)
10 Gb/s 2002
5 Gb/s 2016 (est)
2.5 Gb/s 2016 (est)
1 Gb/s 1998
100 Mb/s 1995
10 Mb/s 1983

Ethernet Speed
Speed in Development
Possible Future Speed

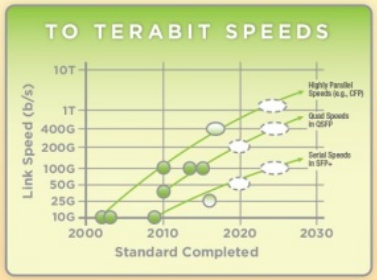

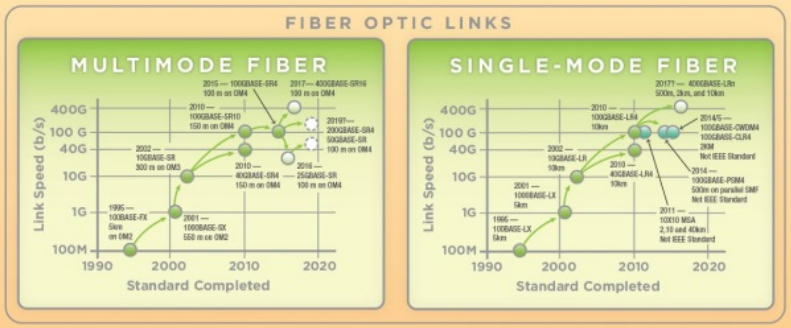

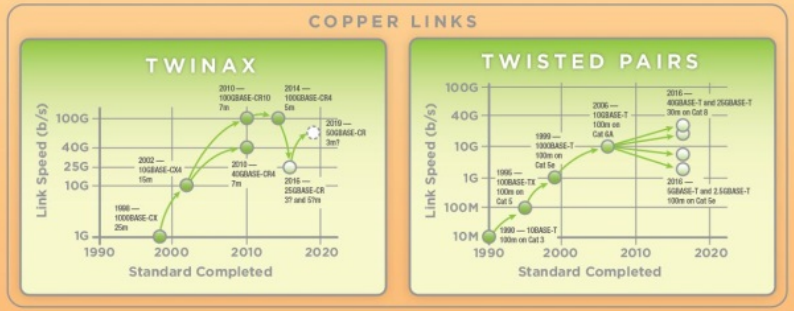

To get a pdf version of the roadmap and to find out more about the roadmap, please go to: **www.ethernetalliance.org/roadmap/**

MEDIA AND MODULES

Ethernet is wired technology and supports a variety of media including backplanes, twisted pair, twinax, multimode fiber and single-mode fiber. Most people know Ethernet by the twisted pair or Cat "x" cabling with RJ45 connectors because close to a billion ports a year are sold. Cat 8 is the latest generation of twisted pair cabling that will be used in 25GBASE-T and 40GBASE-T.

Another popular copper interface is Twinax copper cables that are also known as direct attach cables (DAC)s. DACs may be passive or active and provide very low cost connectivity to servers. Passive DACs are limited to 25 meters or less while active optical cables can go hundreds of meters.

For links longer than 100 meters, fiber optics are required and the graphic below shows three of many module types. The SFP family is the most popular module and supports a single channel or lane in each direction and duplex fibers. The QSFP family supports 4 channels while the CFP2 supports up to 10 channels and duplex or parallel fibers. For 40GbE and beyond, the electrical interface to the module is being defined in IEEE and supports a variety of optical interfaces from IEEE and other sources.

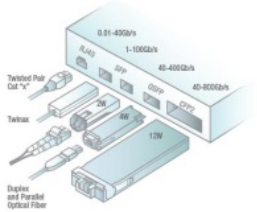

2015
ETHERNET
ROADMAP

THE PAST, PRESENT AND FUTURE OF ETHERNET

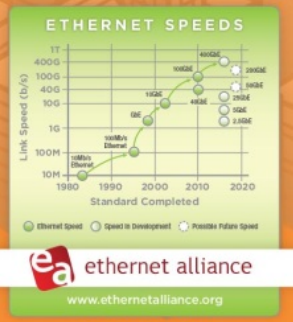

ethernet alliance
www.ethernetalliance.org

Designed by Scott Kipp

$9.95

Back Side of Map

ETHERNET ECOSYSTEM

Represented as a city, the Ethernet Ecosystem is divided into four quadrants that are interconnected by multiple MANs that are typically not Ethernet. While each quadrant has overlapping technologies and requirements, this map organizes the environments with a broad brush. Specific implementations may vary considerably.

The top half of the map represents applications where cost and connectivity are driving concerns. In the home, small office and car, link distances are less than 100 meters and speeds are typically under 10Gb/s, so copper cabling and wireless are ideal. As enterprises scale in size and requirements, they shift towards fiber and 10Gb/s speeds and beyond.

The lower half of the map captures applications that consistently push the bounds of Ethernet and require higher speeds and massive scalability. For example, service providers and hyperscale data centers will be the early adopters of 400GbE. These users may deploy hundreds of thousands of servers in data centers that span multiple football fields and consume hundreds of megawatts of power.

ENTERPRISE AND CAMPUS

Enterprises consume more Ethernet ports than the other environments by connecting desktop computers, devices and Voice over IP (VoIP) phones. The wired Ethernet networks are supplemented with wireless access points (WAPs) that are connected to Ethernet cables. 802.11ac WAPs are driving the need for 2.5 and 5GBASE-T and eventually 10GBASE-T. Most enterprise data centers are less than 10,000 sq ft and use Cat "x" cabling to connect to servers.

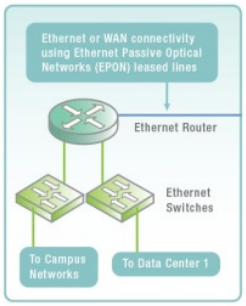

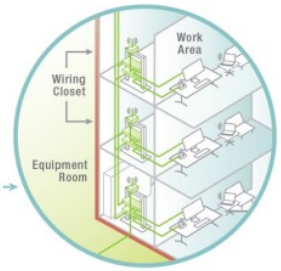

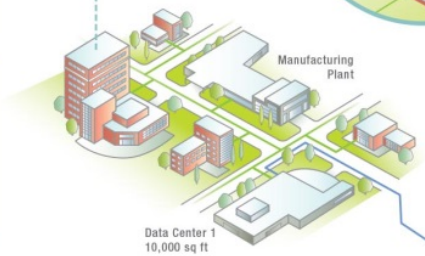

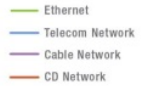

BACKBONE TO OTHER CITIES

BACKBONE TO OTHER CITIES

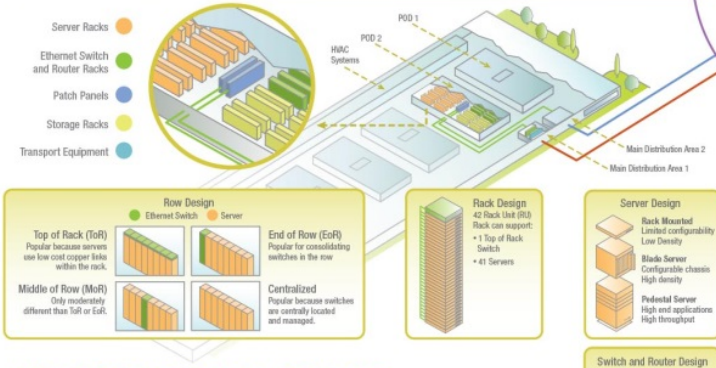

HYPERSCALE DATA CENTER

Hyperscale data centers, also known as warehouse scale computing and mega data centers, are known by their massive size and scalability. Cloud service providers, large enterprises and service providers pack over 100,000 servers that are often divided into several pods. Thousands of 25GbE servers and eventually 50GbE servers in these data centers drive the need for 400GbE to the MAN and WAN.

MANs

Metropolitan Area Networks (MANs) come in many varieties and deliver services to a variety of enterprises, organizations and consumers. Some MANs are based on Ethernet, but the largest MANs are based on Optical Transport Networks (OTN) technologies.

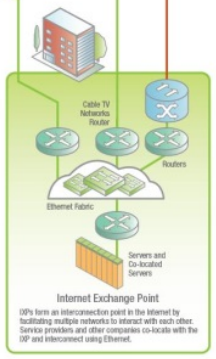

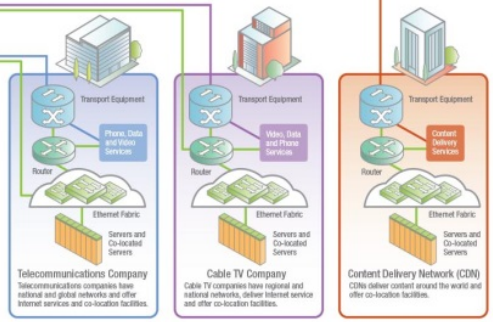

SERVICE PROVIDERS

Service providers deploy MANs and Wide Area Networks (WANs) to deliver a variety of services including Carrier Ethernet. Service providers may use Ethernet Passive Optical Networks (EPON) and cable companies send EPON Protocol over Coax (EPOC). Service providers use routers to interconnect various networks.

The Long and Winding Road

Some Nomenclature

Ethernet Speed

Speed in Development

Possible Future Speed

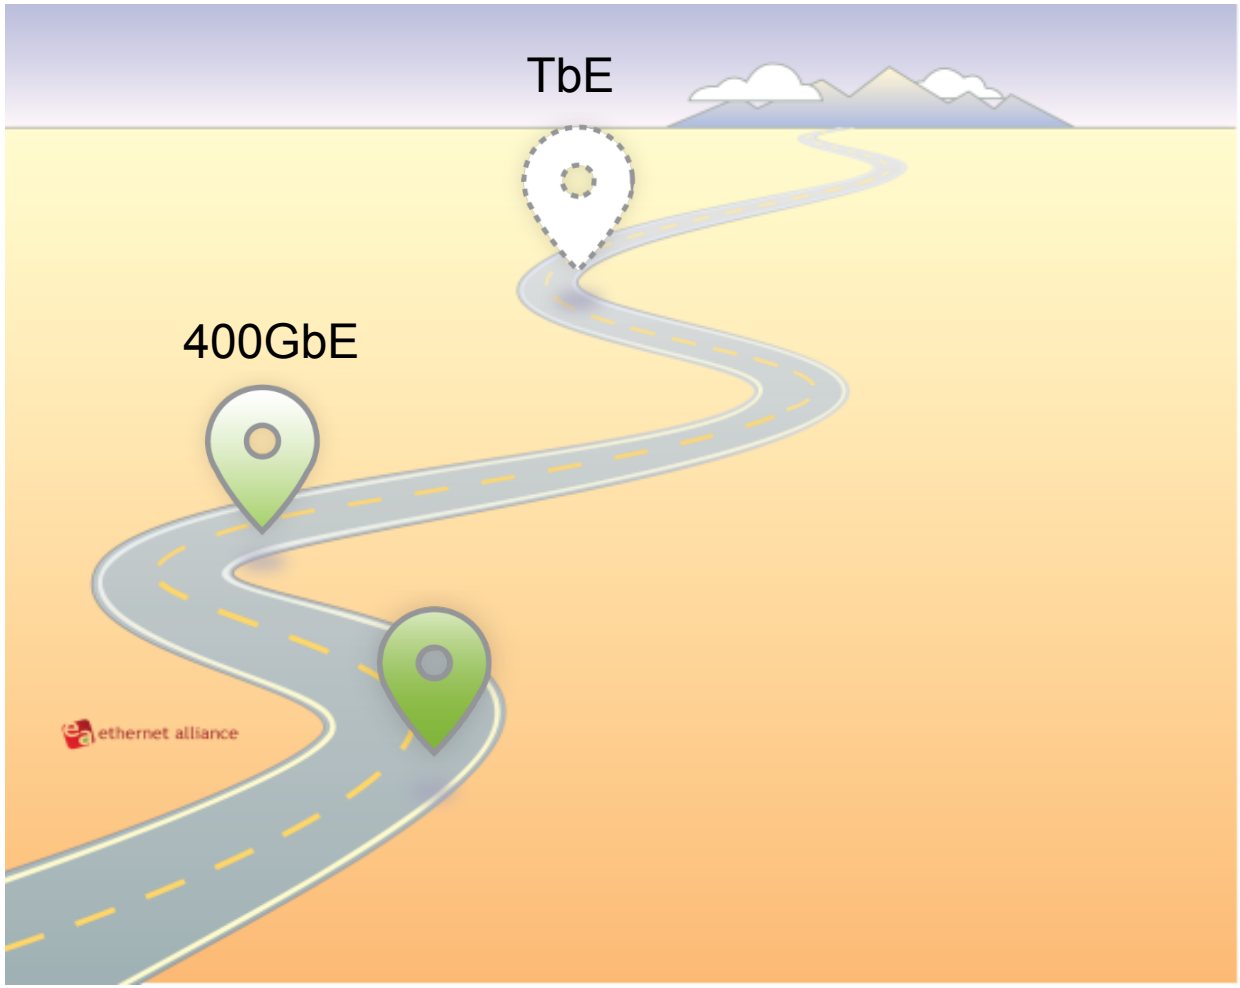

A Dozen Ethernet Speeds

SNIA™
Ethernet Storage Forum

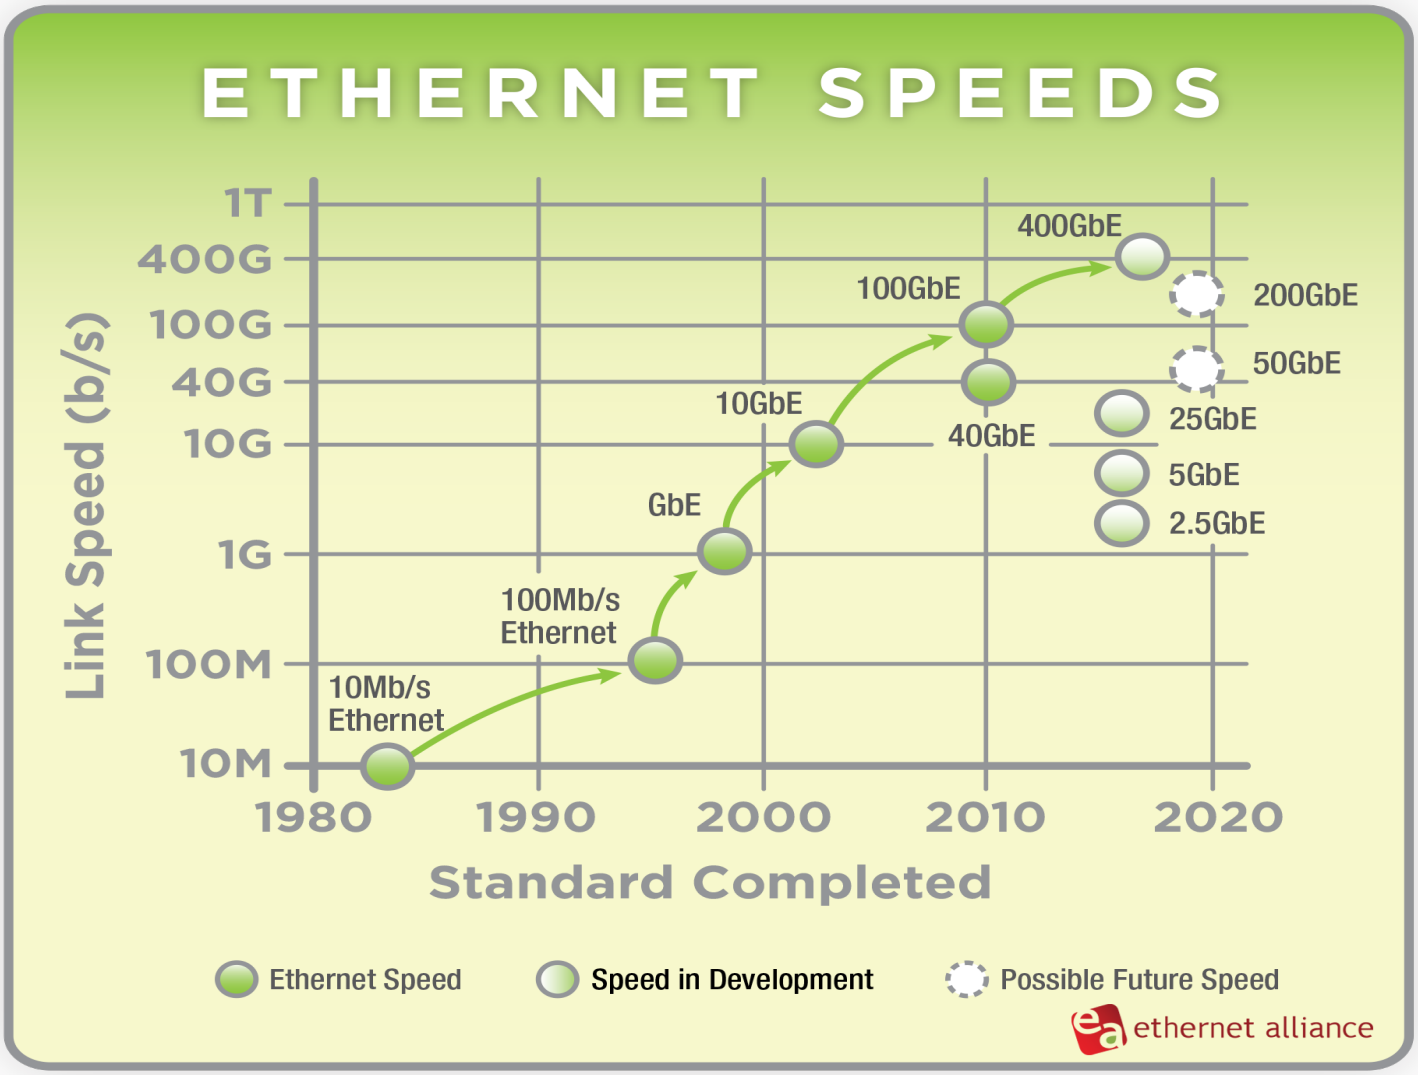

Optical Fiber Roadmaps

SNIA
Ethernet Storage Forum

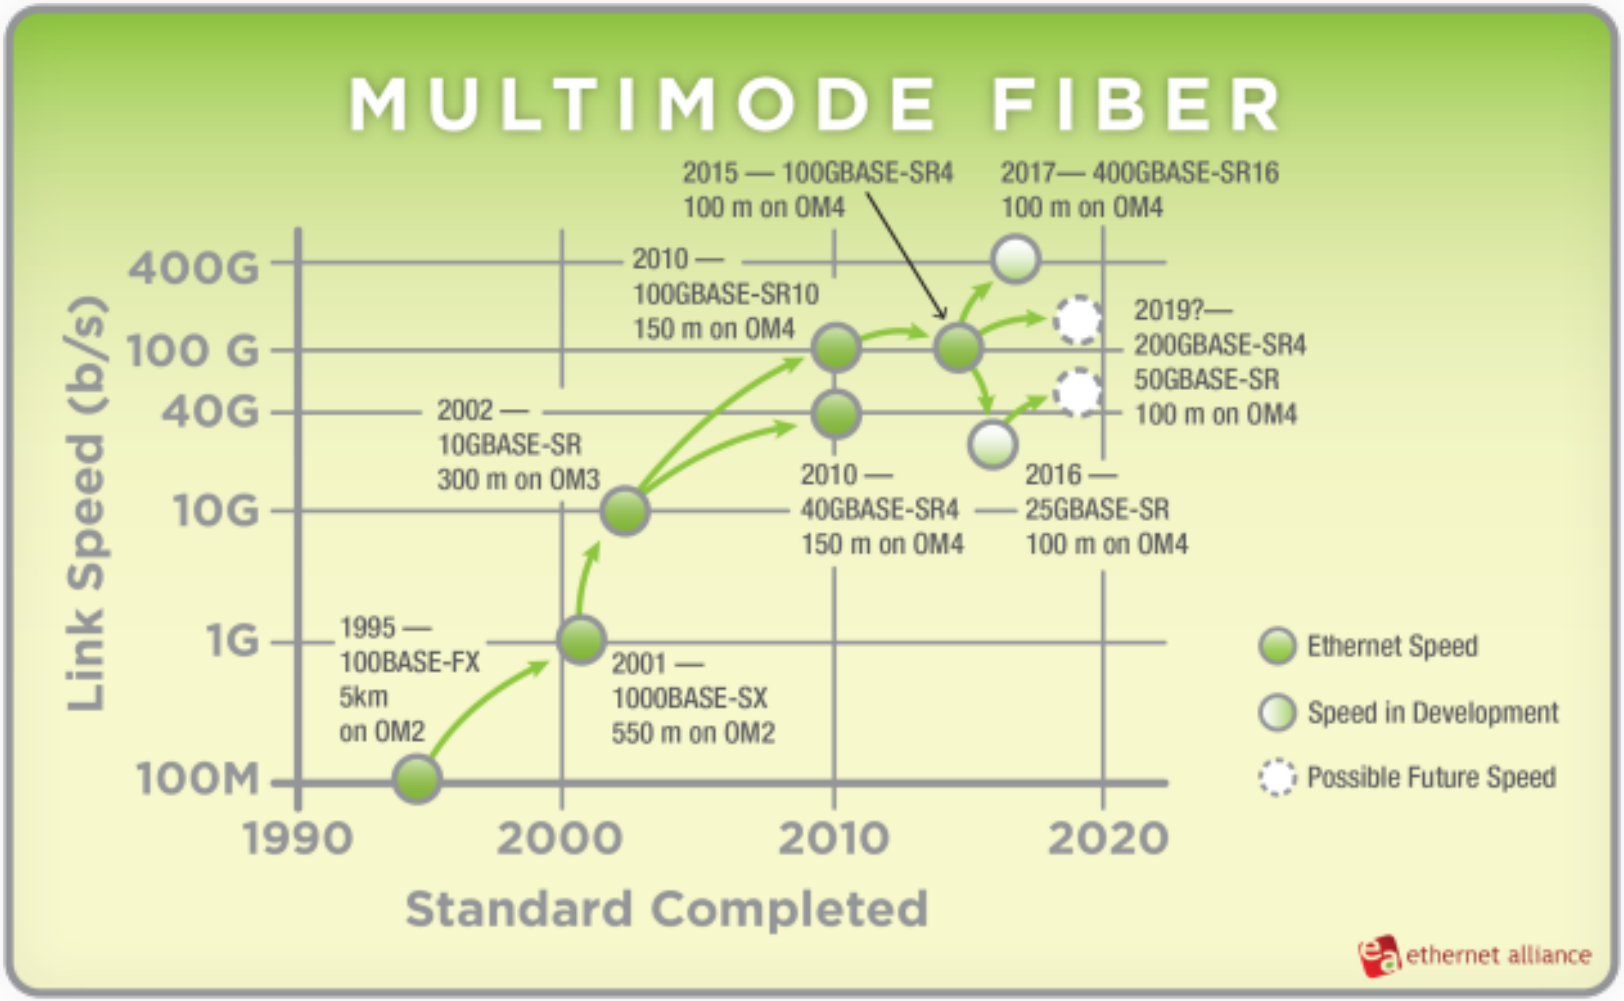

Optical Fiber Roadmaps

SNIA
Ethernet Storage Forum

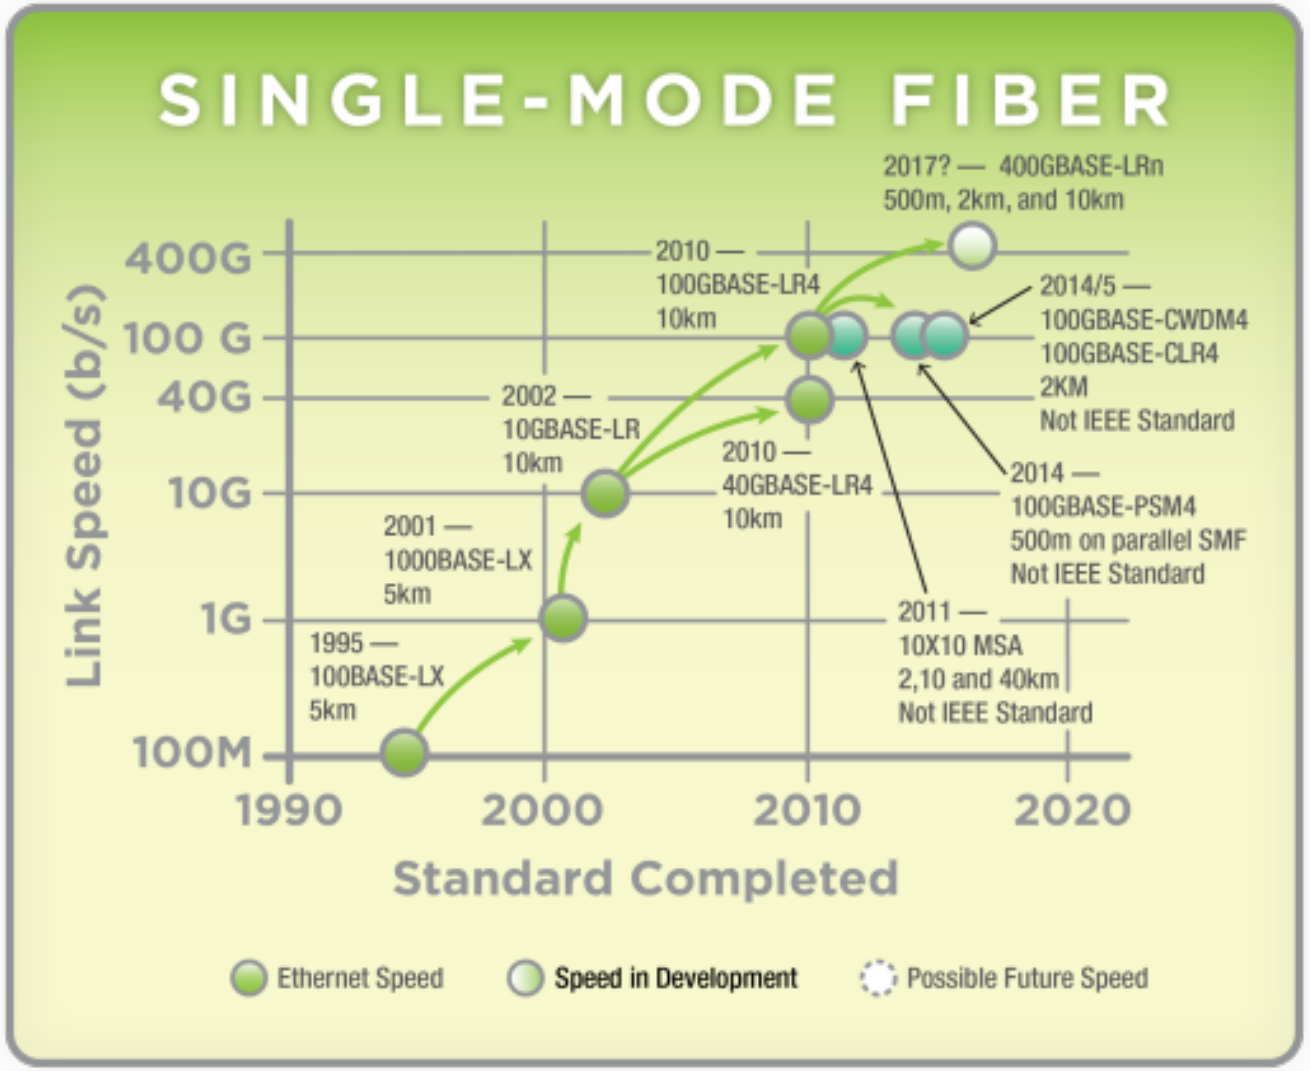

What's Really Going on?

SNIA™
Ethernet Storage Forum

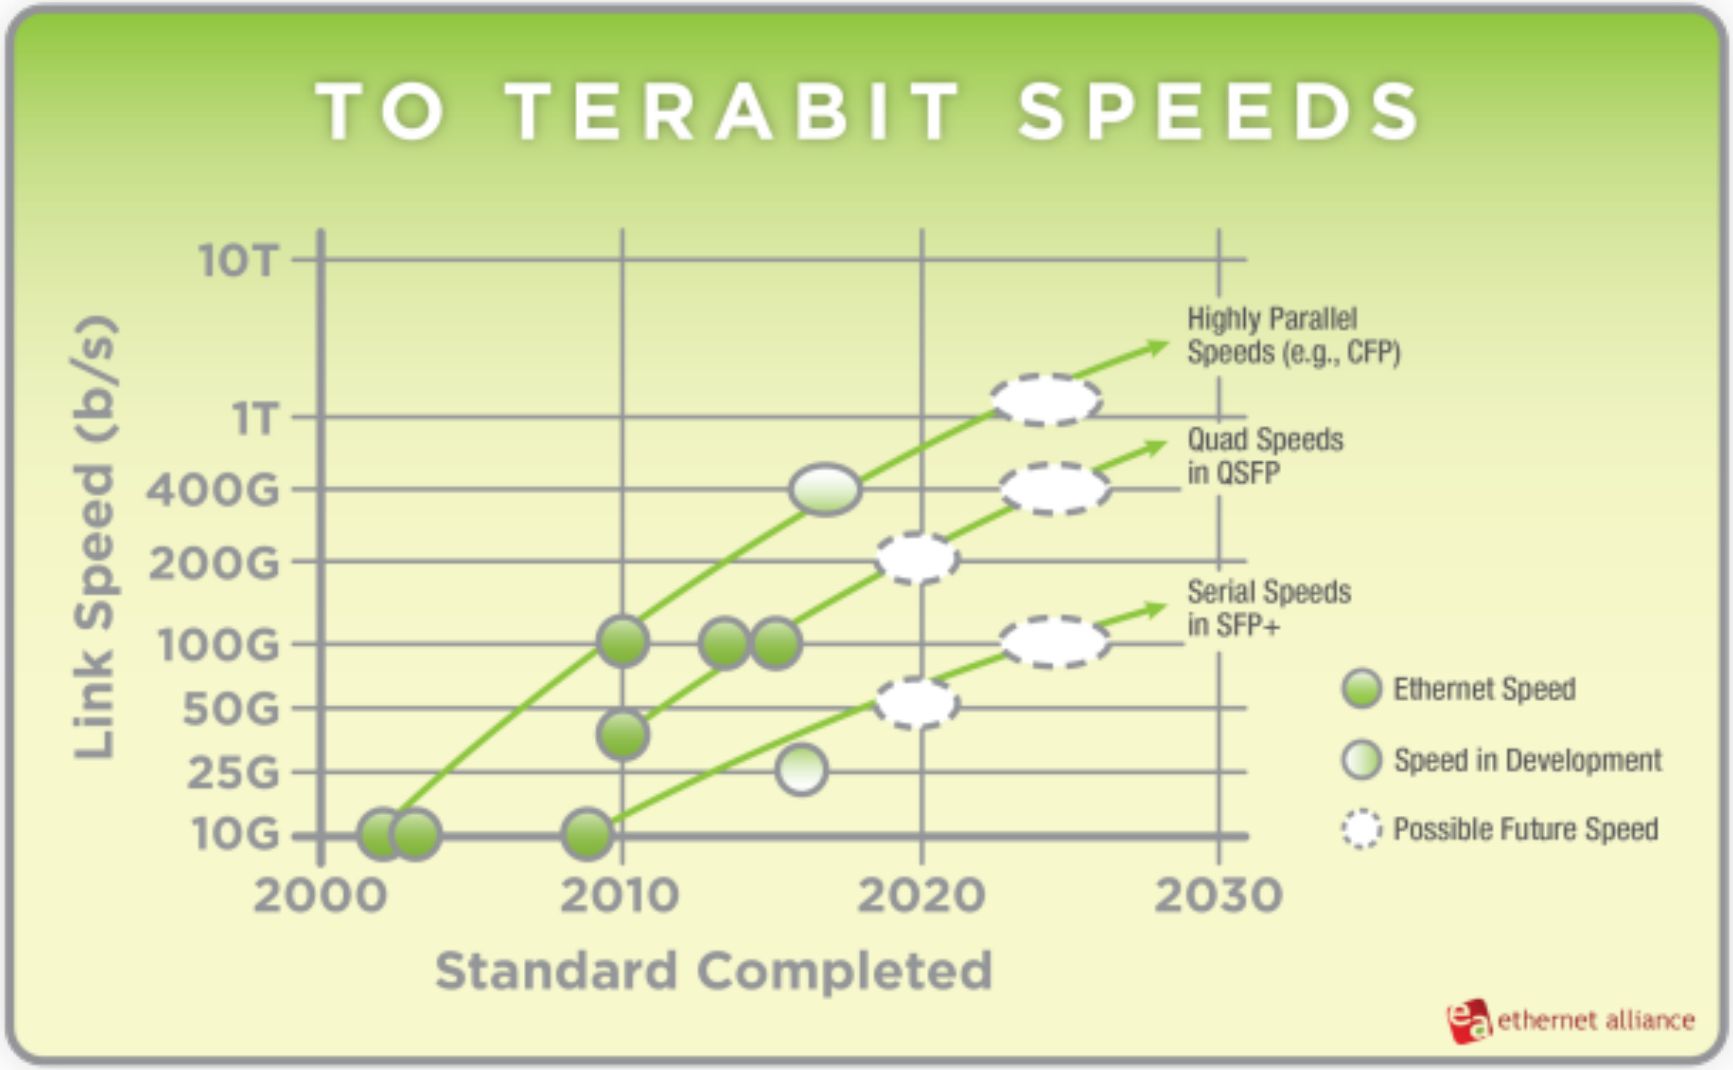

Media and Modules

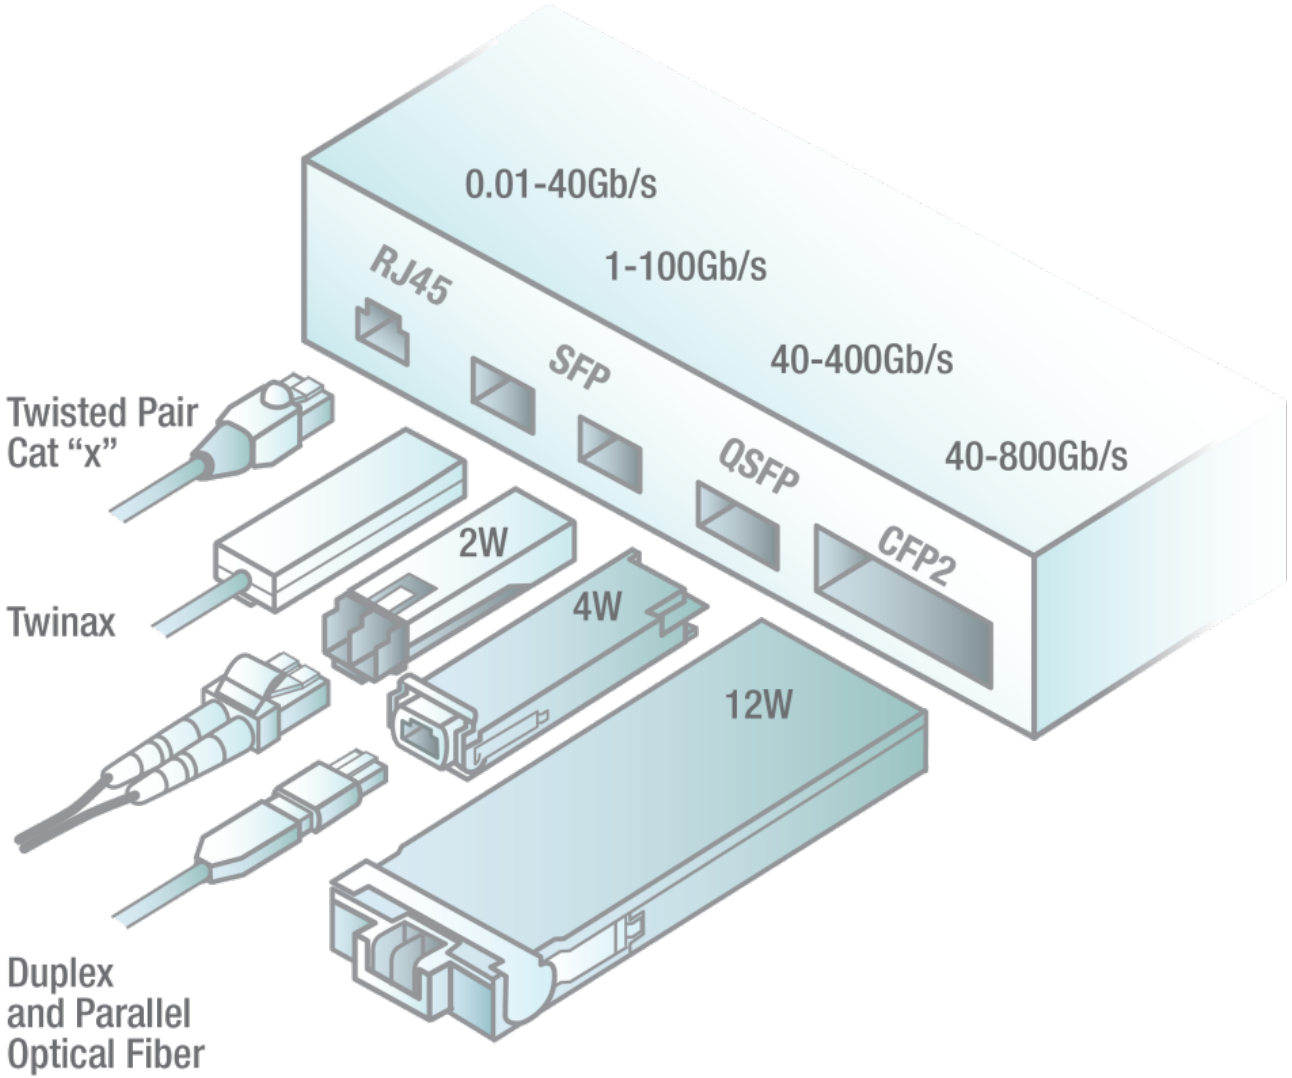

Twisted Pairs Roadmap

SNIA™
Ethernet Storage Forum

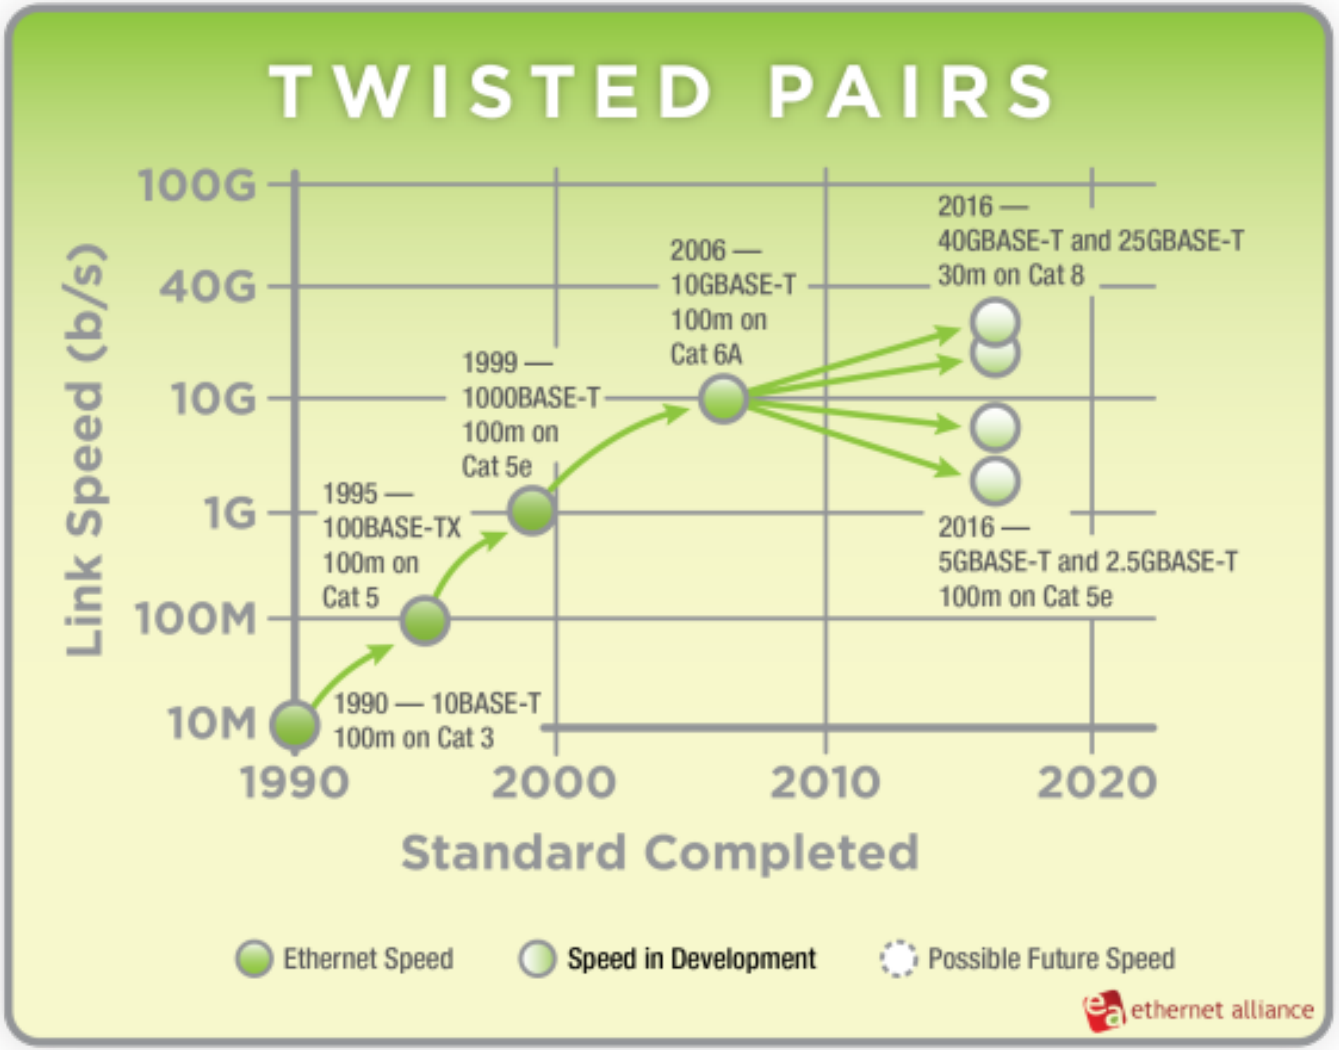

Twinax Roadmap

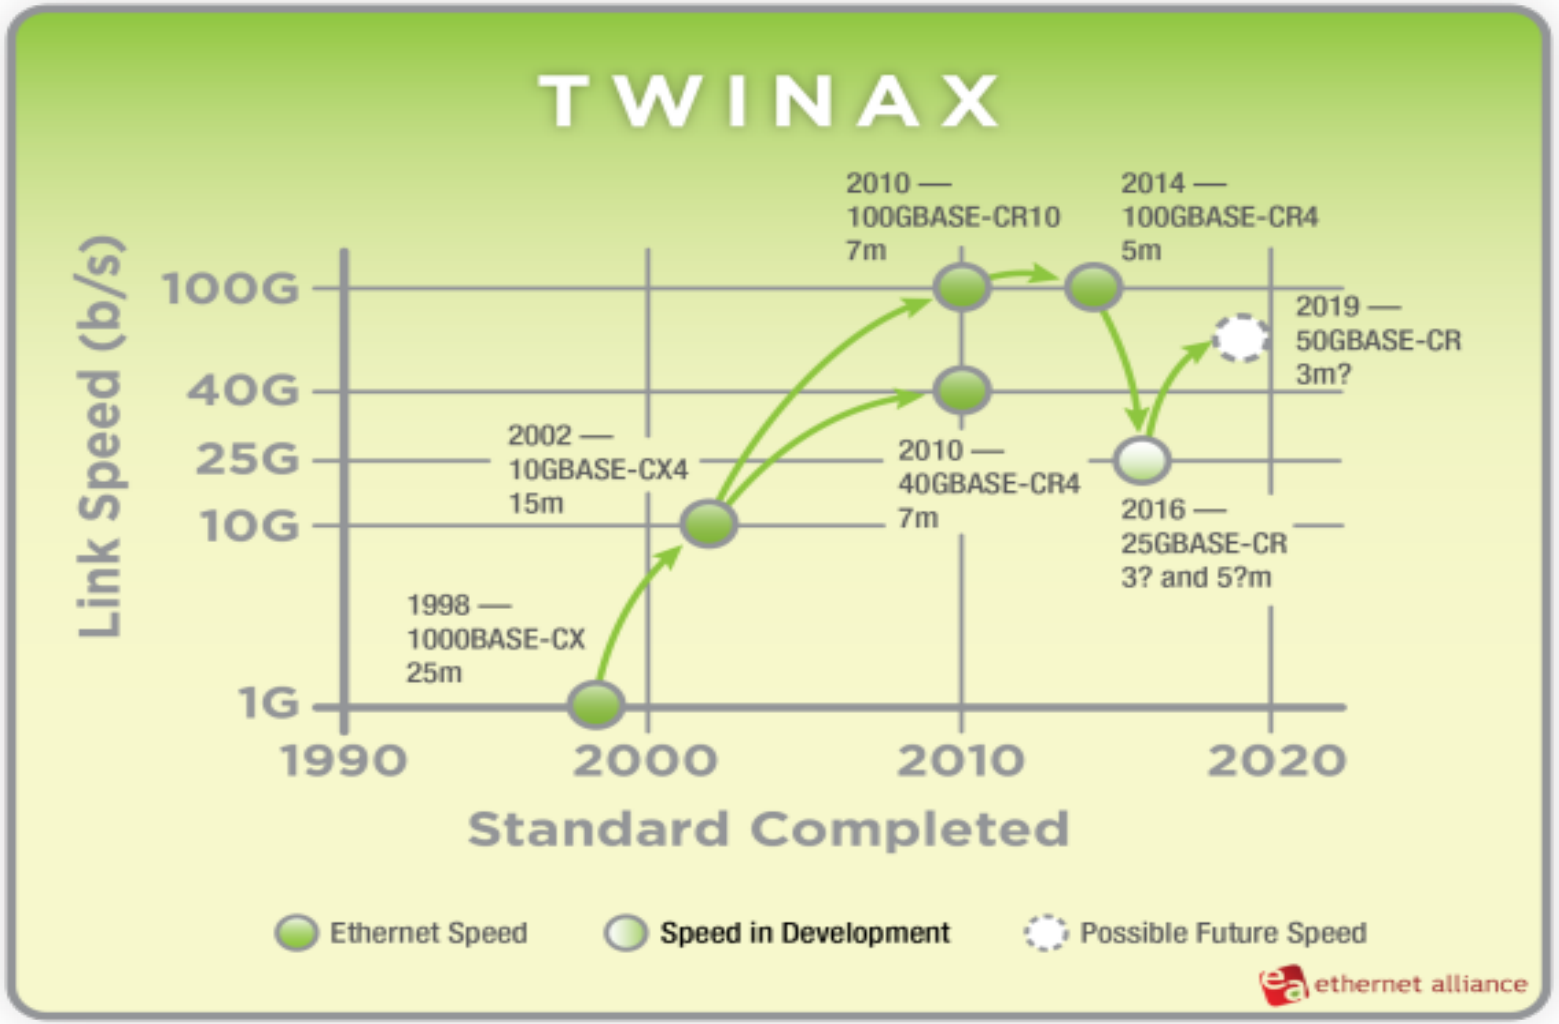

Switch Design

SNIA
Ethernet Storage Forum

Switch and Router Design

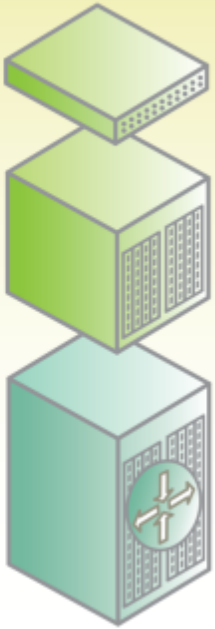

Fixed Port Switch
Limited configurability
No backplane

Modular Switch
Configurable chassis
Backplane for high bandwidth

Modular Router
Layer 3 Routing
High throughput

ethernet alliance

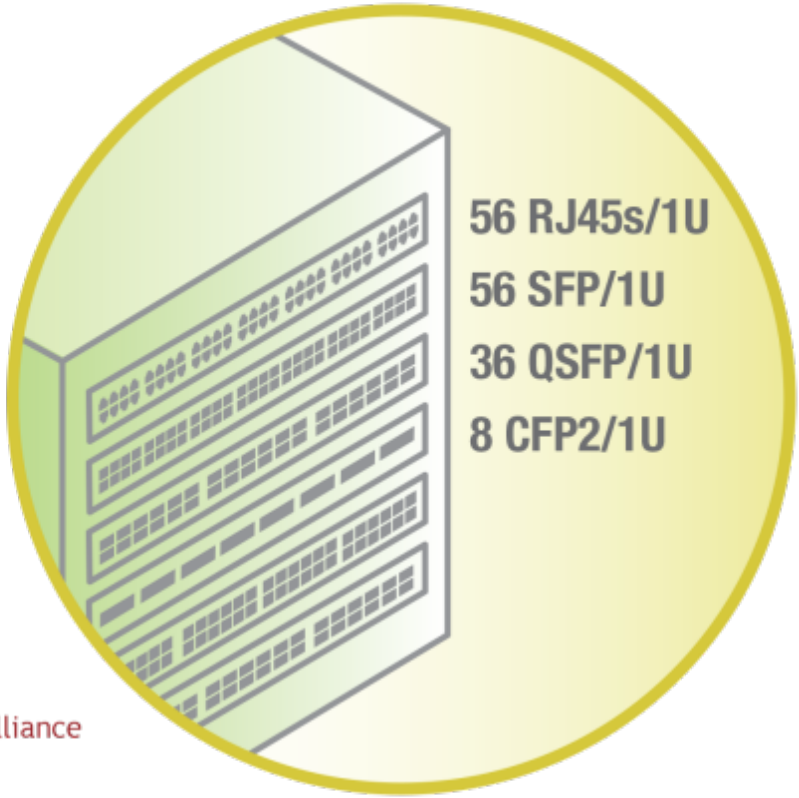

Hyperscale Data Centers

SNIA™
Ethernet Storage Forum

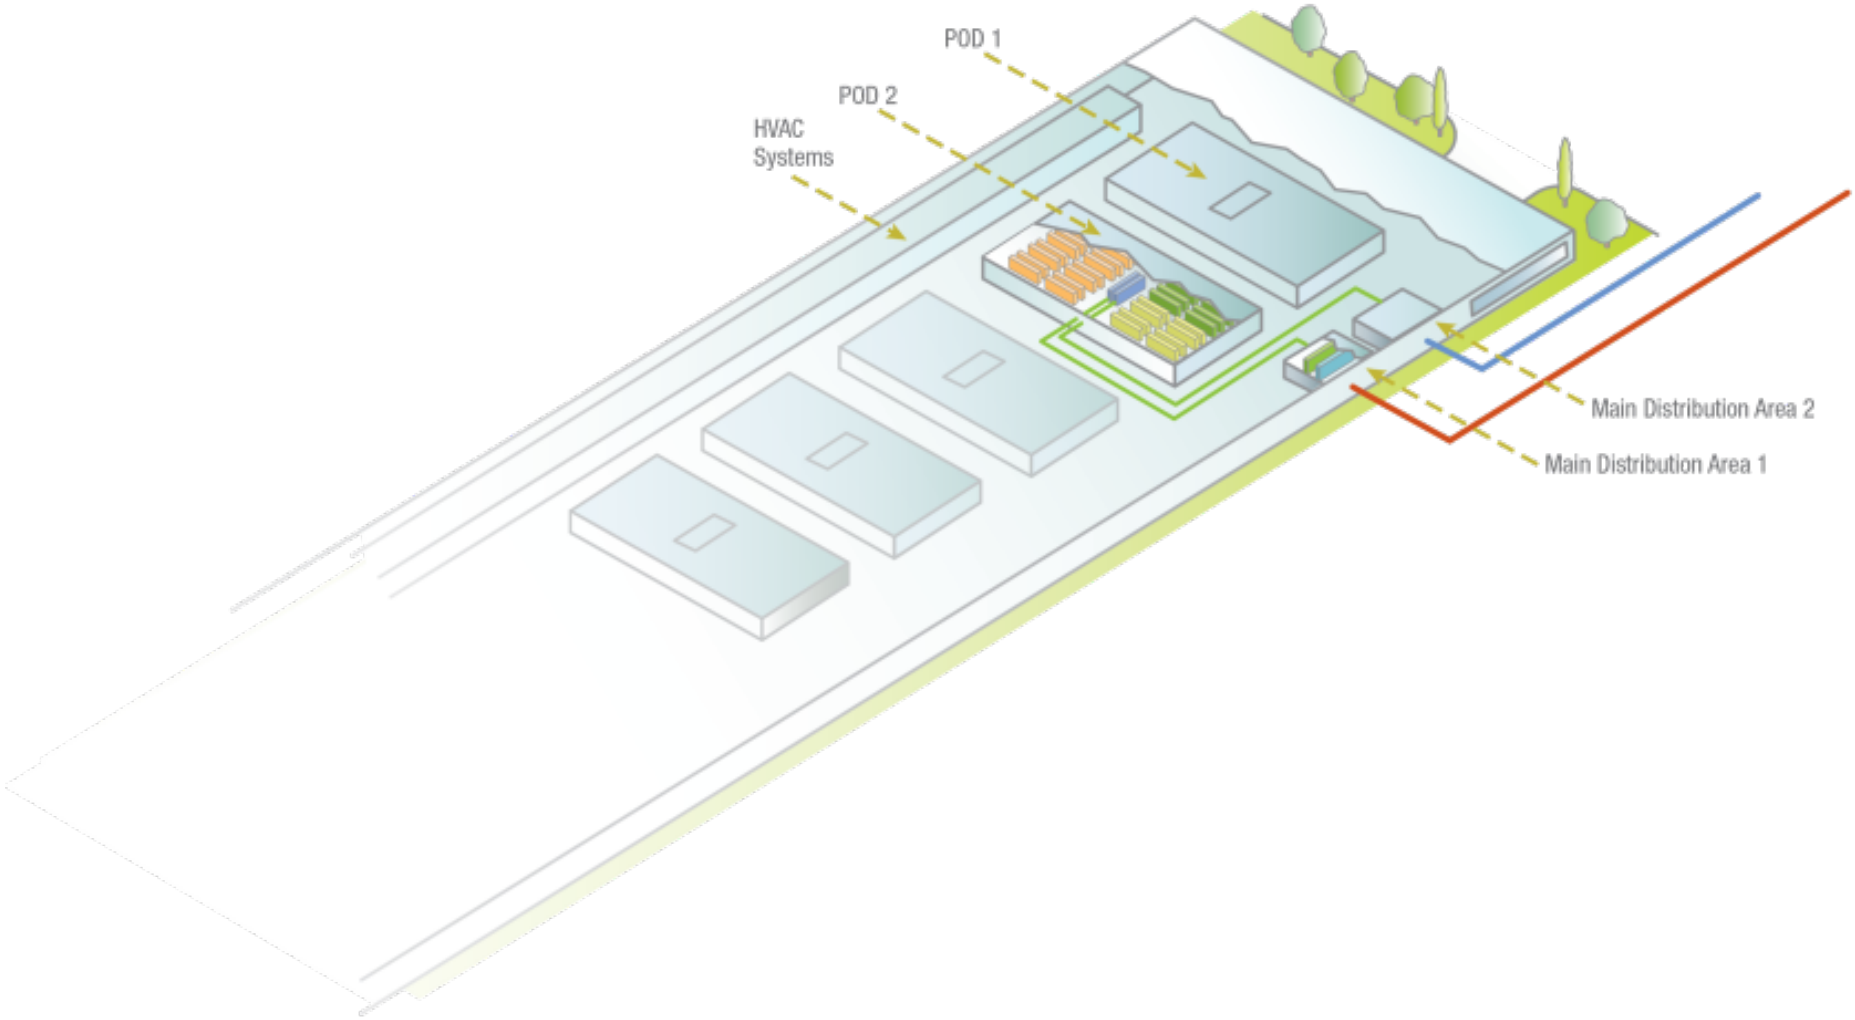

Racks and Racks in a POD

SNIA™
Ethernet Storage Forum

Server Racks

Ethernet Switch and Router Racks

Patch Panels

Storage Racks

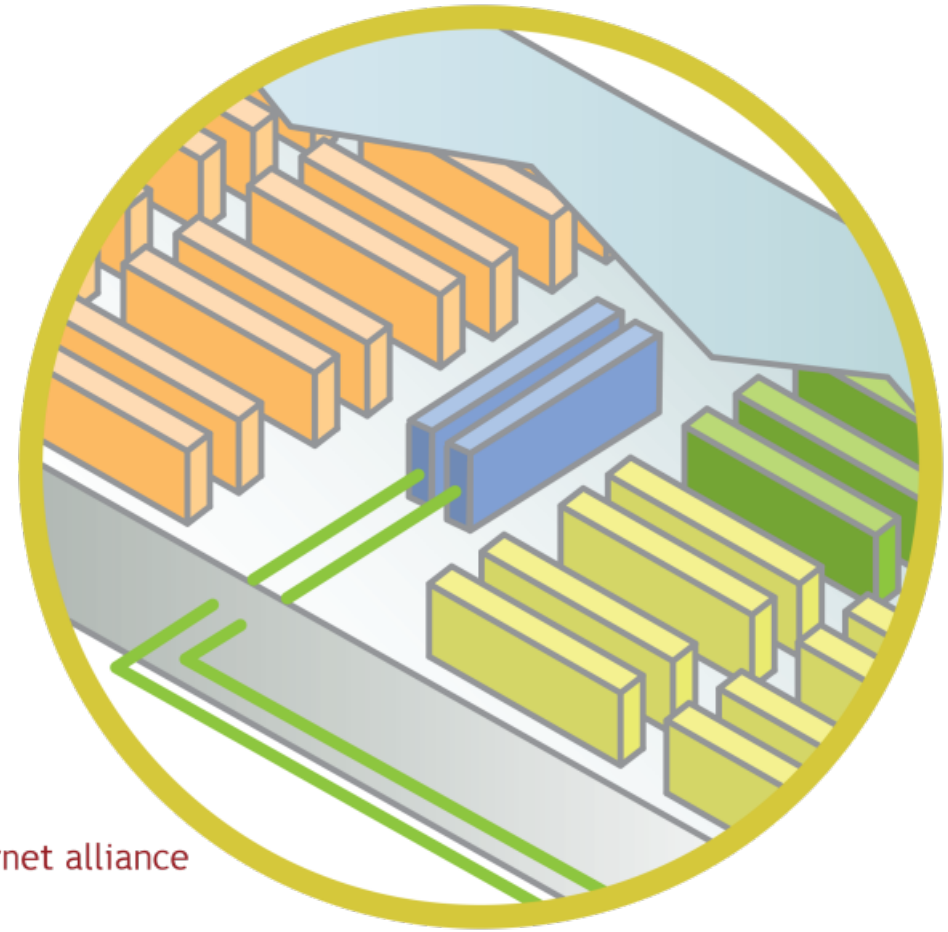

ethernet alliance

Row Design

SNIA
Ethernet Storage Forum

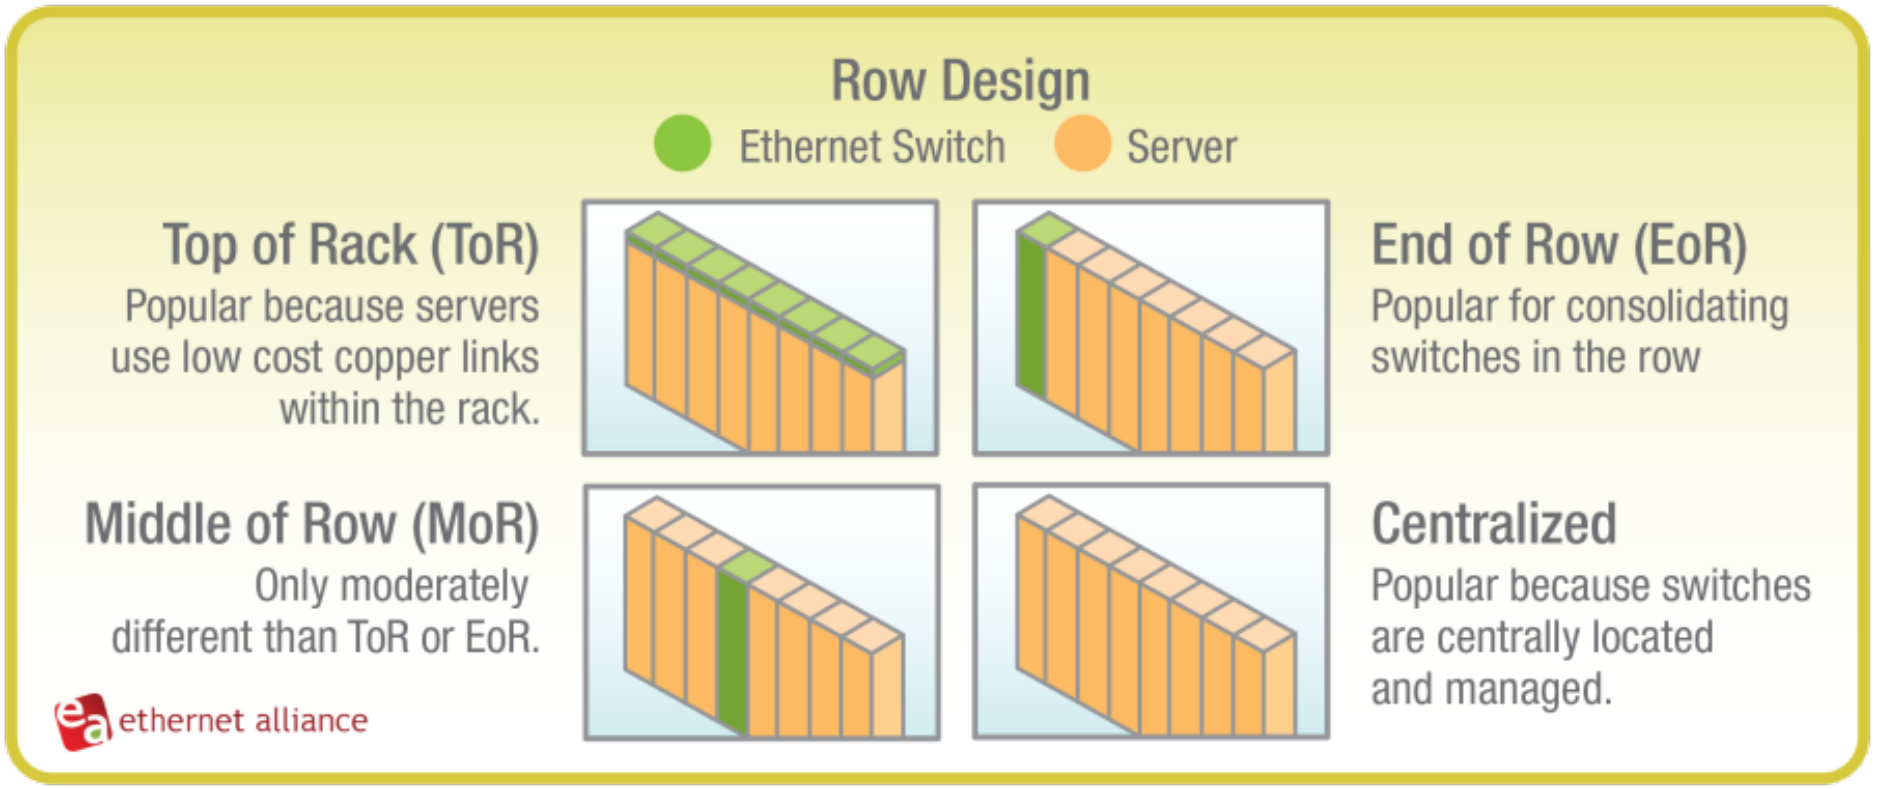

Rack and Server Design

SNIA™
Ethernet Storage Forum

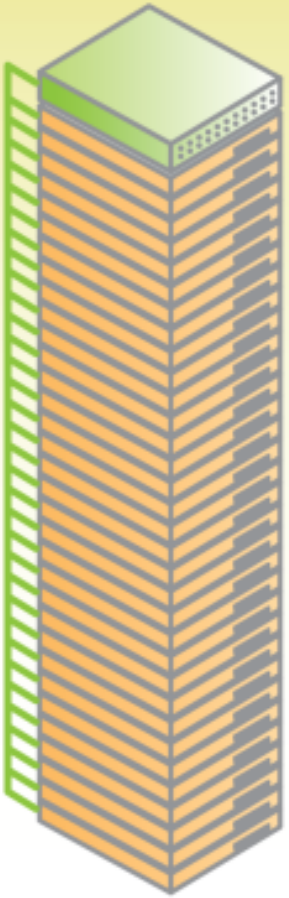

Rack Design

42 Rack Unit (RU)
Rack can support:

- 1 Top of Rack Switch
- 41 Servers

ethernet alliance

Server Design

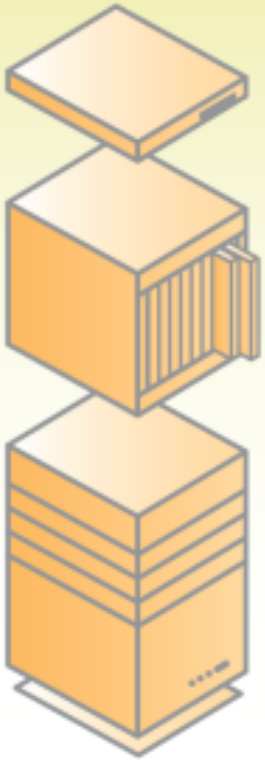

Rack Mounted
Limited configurability
Low Density

Blade Server
Configurable chassis
High density

Pedestal Server
High end applications
High throughput

ethernet alliance

Service Providers

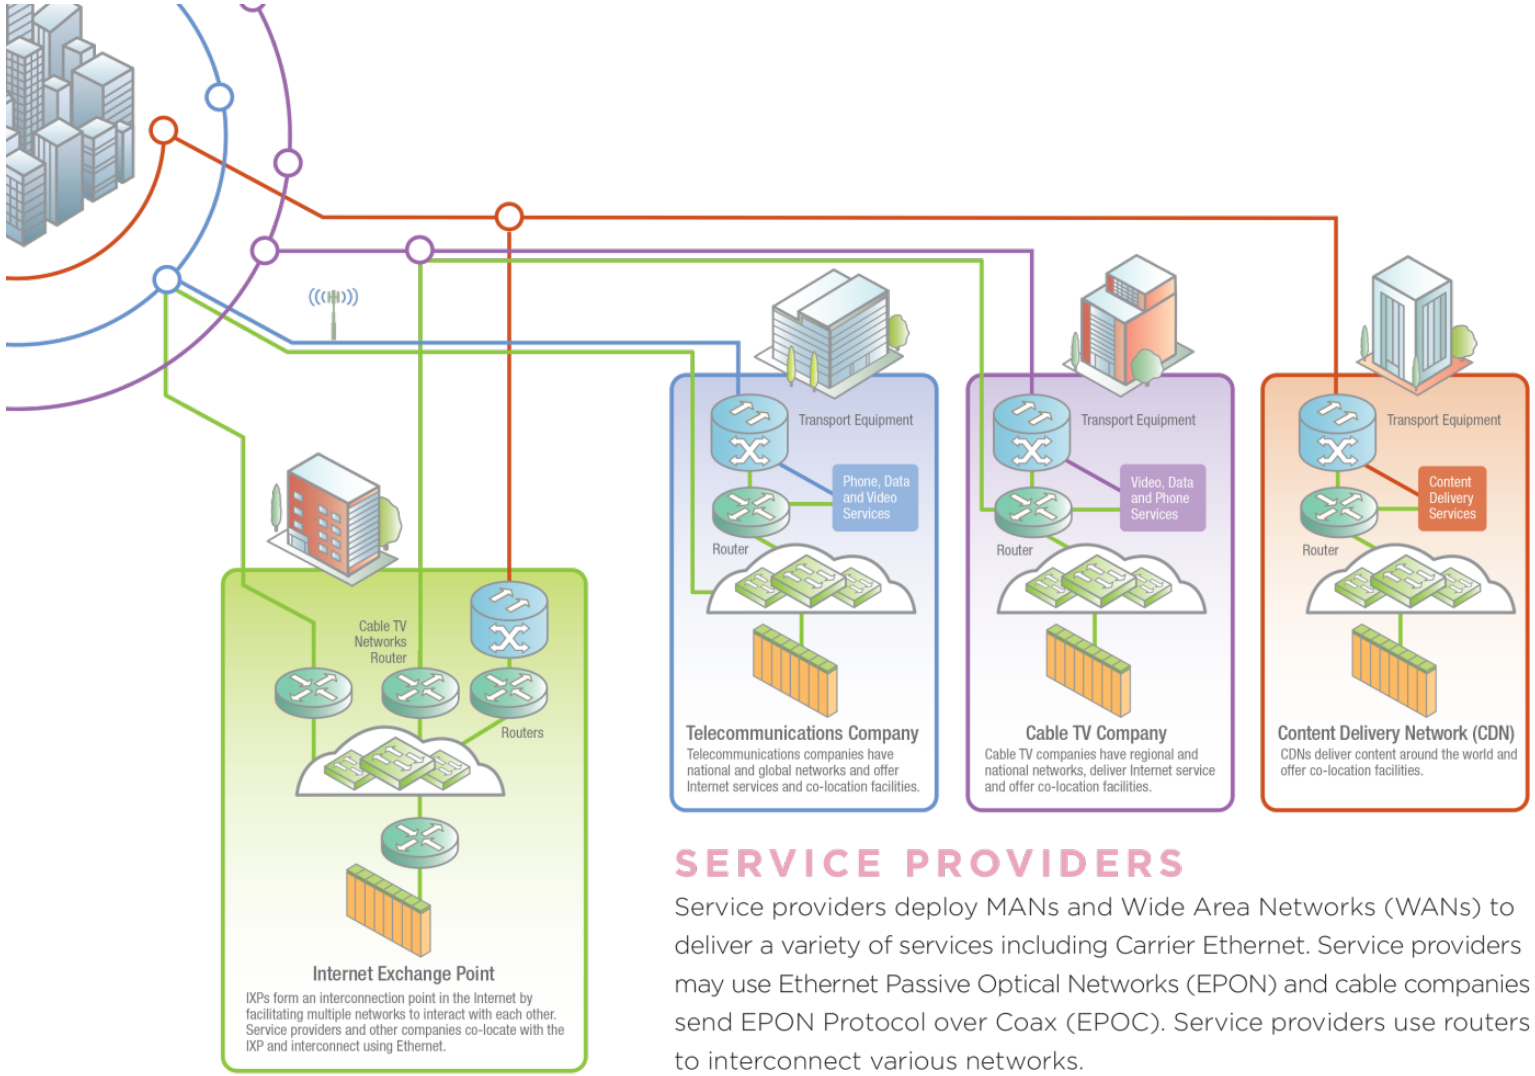

SERVICE PROVIDERS

Service providers deploy MANs and Wide Area Networks (WANs) to deliver a variety of services including Carrier Ethernet. Service providers may use Ethernet Passive Optical Networks (EPON) and cable companies send EPON Protocol over Coax (EPOC). Service providers use routers to interconnect various networks.

Our First Infographic

SNIA
Ethernet Storage Forum

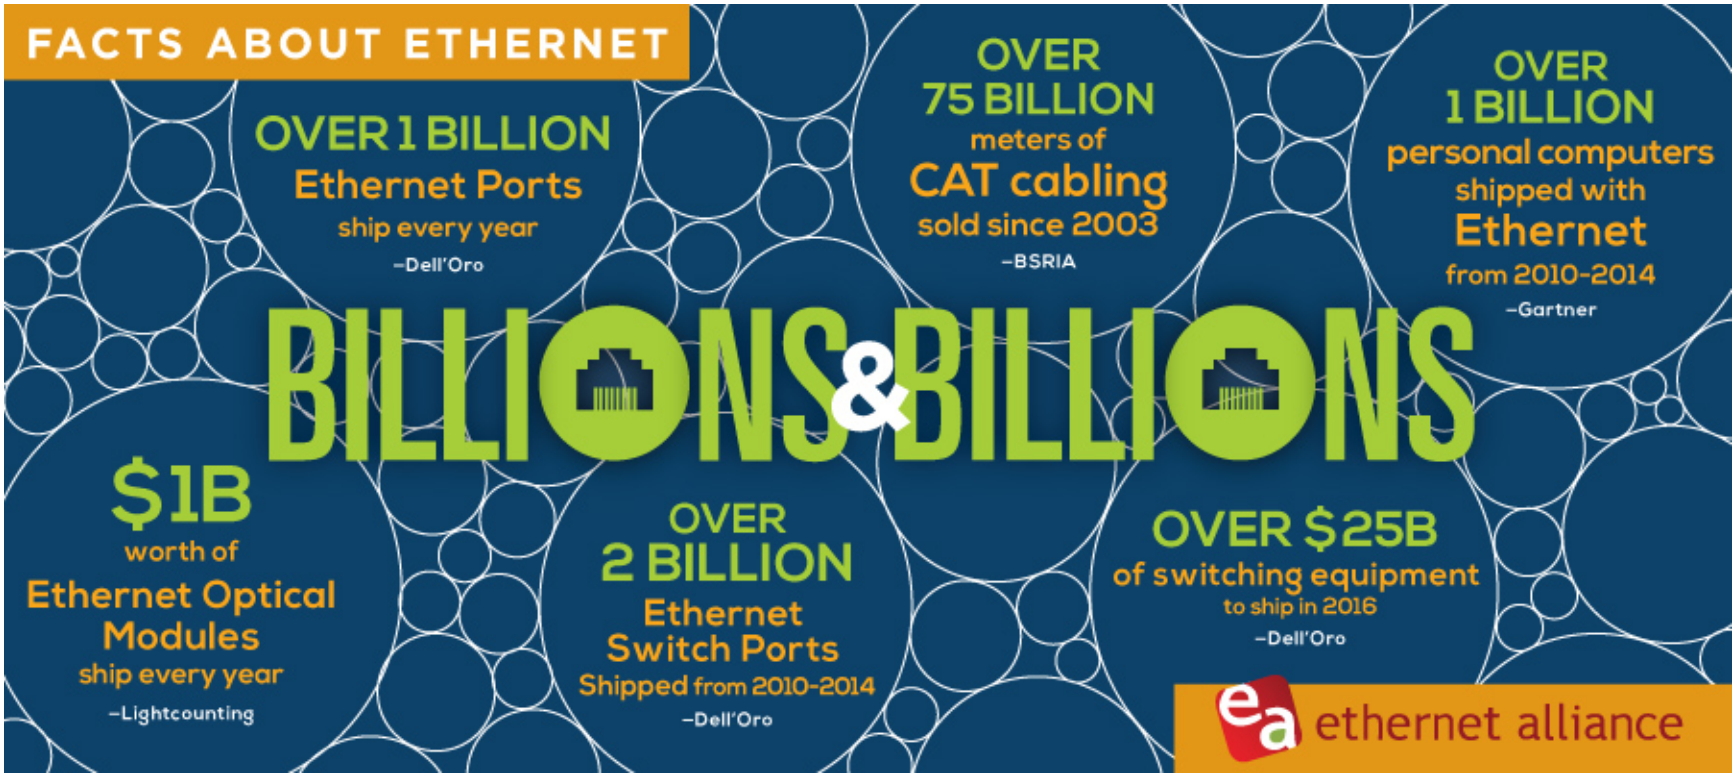

More Roadmap Information

SNIA
Ethernet Storage Forum

- Free downloads at www.ethernetalliance.org/roadmap/
  - PDF of map
  - White paper
  - Presentation with graphics for your use
- See us at SC15 – November 2015

After This Webcast

SNIA Ethernet Storage Forum

- This webcast and a PDF of the slides will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand
  - http://www.snia.org/forums/esf/knowledge/webcasts
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog
  - http://sniaesfblog.org/
- Follow us on Twitter@SNIAESF

Thank You

SNIA™
Ethernet Storage Forum