

“This diagram was hand drawn by Robert M. Metcalfe and photographed by Dave R. Boggs in 1976 to produce a 35 mm slide used to present Ethernet to the National Computer Conference in June of that year.”

Evolution of Ethernet Speeds: What's New and What's Next

Greg Hankins <greg.hankins@alcatel-lucent.com>
DENOG7 - Darmstadt, Germany

Image source: http://www.ieee802.org/3/ethernet_diag.html


Agenda

- 1. Ethernet Speed Evolution**
2. What's Next: 2.5 GE and 5 GE
3. What's Next: 25 GE
4. What's New: 40 GE
5. What's New: 100 GE
6. What's Next: 400 GE

Ethernet Speed Evolution Over 40+ years

New Speeds Driven by Diverse Market Requirements

- Market requirements for Ethernet are changing for different applications
 - Speed
 - Distance
 - Cost
- Different new speeds are needed for
 - Wireless access points: 2.5 GE and 5 GE
 - Servers: 25 GE
 - Core networks: 400 GE
- New Ethernet speeds under development will address these different requirements

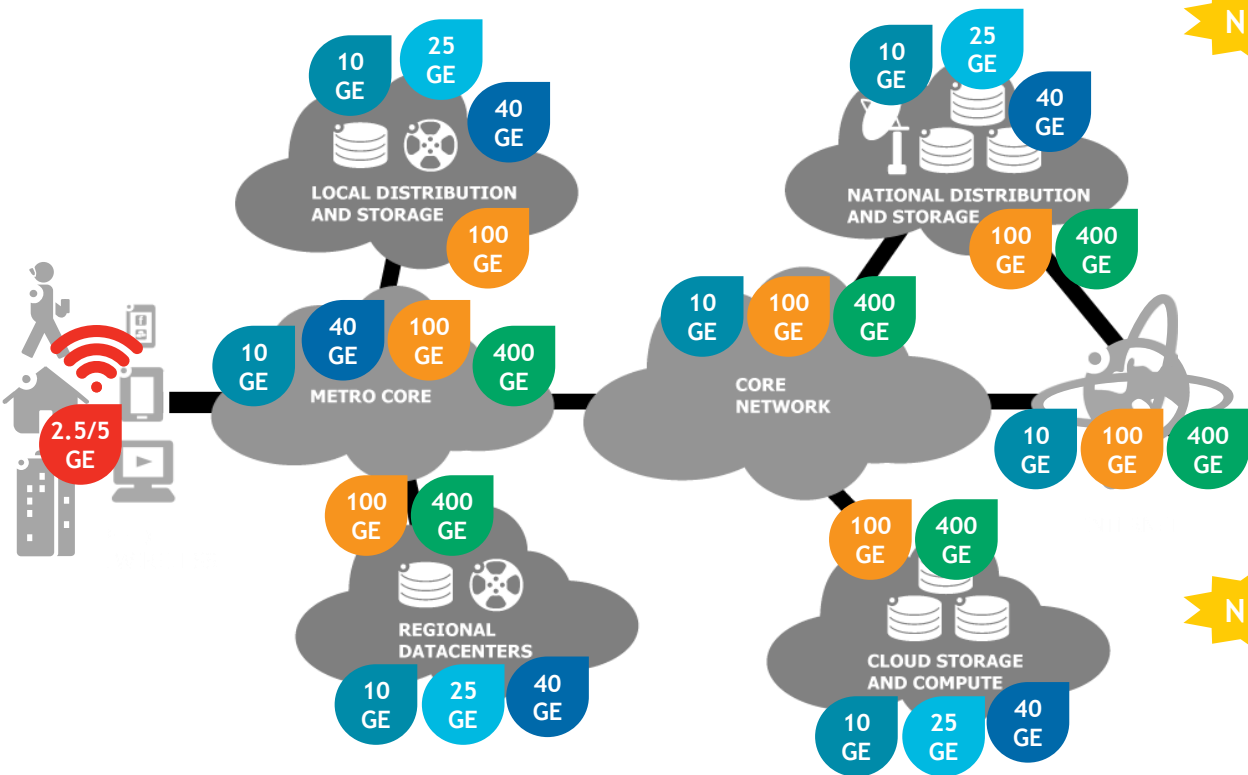


Six New Ethernet Speeds May be Coming Soon - Same Amount as in the Past 30 Years

Roadmap courtesy of the Ethernet Alliance: <http://www.ethernetalliance.org/roadmap/>

Higher Speed Ethernet Target Applications

Key Application Drivers



NEW 2.5/5 GE Applications (~2016)

- Higher Speed Wireless
- Large Cat 5e/6 Installed Base

NEW 25 GE Applications (~2016)

- Data Center Access
- Server NICs

40 GE Applications MORE

- Data Center Aggregation and Core
- Data Center Access
- Server NICs
- Metro Core

100 GE Applications MORE

- Service Provider Aggregation and Core
- Data Center Core
- Metro Core

NEW 400 GE Applications (~2017)

- Service Provider Core
- Large Data Center Core
- Large Metro Core

Agenda

1. Ethernet Speed Evolution
- 2. What's Next: 2.5 GE and 5 GE**
3. What's Next: 25 GE
4. What's New: 40 GE
5. What's New: 100 GE
6. What's Next: 400 GE

Market Drivers for 2.5 GE and 5 GE

Higher Speed Wireless

- Wireless access speeds are exceeding the wired connection speed from access points (APs) to the switching infrastructure
- APs are now capable of transmitting multiple Gb/s of aggregate capacity
 - 802.11n (2007+): 600 Mb/s
 - 802.11ac (2013+): 3.47 Gb/s
 - 802.11ac (2017+): 6.93 Gb/s
 - 802.11ax (2019+): 4x faster
- AP's wired connection speed should be at least 75% of maximum wireless speed to avoid throughput bottlenecks
 - Requires wired speeds of 2.5 GE and 5 GE over Cat 5e or Cat 6
 - PoE is also required to power the access points

802.11ac
Up to 6.93 Gb/s
Aggregate Capacity



1 Gb/s
on Cat 5e/6


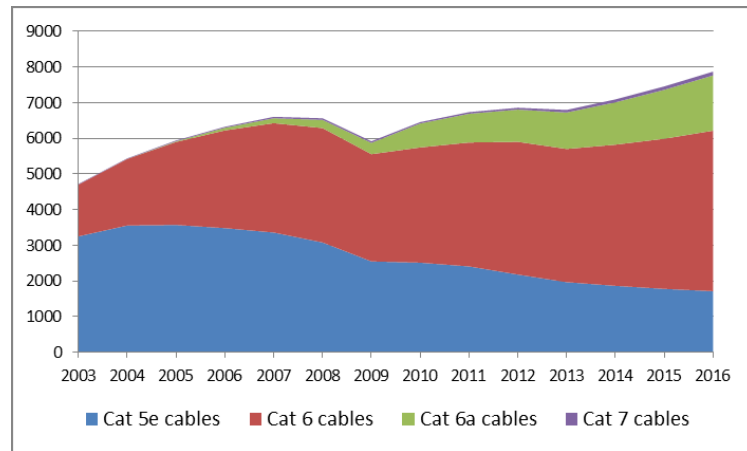


Cat
Videos

Market Drivers for 2.5 GE and 5 GE

Large Cabling Installed Base

- Would like to continue to use large installed based of existing Cat 5e/6 at higher speeds than 1 GE
 - 10GBASE-T requires Cat 6A to reach 100 m, or Cat 6 for up to 55 m depending on the installation
 - 25GBASE-T and 40GBASE-T require Cat 8 and are limited to 30 m
- Cat 5e/6 is widely deployed around the world in every type of building
 - BSRIA cabling report from 2003 - 2016 used in IEEE CFI
 - 58 B meters sold world-wide in 2014
 - 1282 M (>90%) of installed outlets
 - Data centers is only ~4%
- Installed base is not going away anytime soon, and is still growing
- Applications for wireless, desktops, small cell, security, etc




Data and diagram source: http://www.ieee802.org/3/NGEBASET/public/jan15/jones_ngeabt_04c_0115.pdf

2.5 GE and 5 GE Developments

- IEEE P802.3bz 2.5/5GBASE-T Task Force started in March 2015
 - 2.5GBASE-T: 4 x 625 Mb/s over 100 m Cat 5e (Class D) or Cat 6 (Class E) unshielded twisted-pair copper cabling
 - 5GBASE-T: 4 x 1.250 Gb/s over 100 m Cat 5e (Class D) or Cat 6 (Class E) unshielded twisted-pair copper cabling
 - MultiGBASE-T auto-negotiation between 2.5GBASE-T, 5GBASE-T, 10GBASE-T, 25GBASE-T, 40GBASE-T
 - Automatic MDI/MDI-X configuration
 - PoE support including IEEE 802.3bt amendment (power over 4 pairs)
 - Optional Energy Efficient Ethernet (EEE) support
- Generated Draft 1.1 for Task Force review
- Standard expected in December 2016
- Interfaces expected on the market in 2016+
- Task Force web page <http://www.ieee802.org/3/bz/>
- IEEE Study Group will be formed for backplane and short reach twinax copper cable interfaces
 - July 2015 CFI http://www.ieee802.org/3/cfi/0715_2/CFI_02_0715.pdf

2.5 GE and 5 GE Industry Groups

| | MGBASE-T Alliance | NBASE-T Alliance |
|------------------|---|---|
| Founded | June 2014 | October 2014 |
| Supporters | Component Vendors, System Vendors, Broadcom | Component Vendors, System Vendors, Cisco, !Broadcom |
| More Information | <p>MGBASE-T Alliance</p> <p>http://www.mgbasetalliance.org/</p> | <p></p> <p>http://www.nbaset.org/</p> |

Agenda

1. Ethernet Speed Evolution
2. What's Next: 2.5 GE and 5 GE
- 3. What's Next: 25 GE**
4. What's New: 40 GE
5. What's New: 100 GE
6. What's Next: 400 GE

Market Drivers for 25 GE

- Provide a server connection speed faster than 10 GE that is optimized for cost, throughput and efficiency
- Maximize efficiency of server connections to access switches in data centers
- Use a single 25 Gb/s signaling lane based on existing 25 Gb/s technology
 - 100 GE for backplanes and copper cables
 - CAUI-4 signaling
 - SFP28, QSFP28, CFP2, CFP4 media modules
 - 100 GE QSFP28 to 4 x 25 GE SFP28 breakout
- But what about 40 GE?
 - Inefficient 4 x 10 Gb/s signaling
 - Higher cost and size of QSFP+ compared to SFP28
 - Market requirements vary, multiple speeds are needed for different applications



25 Gb/s Maximizes Bandwidth and Efficiency

Overall Lower CapEx and OpEx

| 3.2 Tb/s Switch | Servers | 100 GE Uplinks | Capacity Utilization (Tb/s) | Capacity Utilization (%) | ToR Switches for 100K Servers |
|---------------------|---------|----------------|-----------------------------|--------------------------|-------------------------------|
| 25 GE (1 x 25 Gb/s) | 96 | 8 | 3.2 | 100 | 1042 |
| 40 GE (4 x 10 Gb/s) | 28 | 4 | 1.52 | 47.5 | 3572 |

| Port Speed (Gb/s) | Lane Speed (Gb/s) | Lanes per Port | Usable Ports | Total Capacity (Tb/s) |
|-------------------|-------------------|----------------|--------------|-----------------------|
| 10 | 10 | 1 | 128 | 1.28 |
| 25 | 25 | 1 | 128 | 3.2 |
| 40 | 10 | 4 | 32 | 1.28 |
| 100 | 25 | 4 | 32 | 3.2 |

- Connections to switch ASICs is limited by SERDES count and bandwidth
- Single higher speed 25 Gb/s lanes maximize bandwidth and switch fabric utilization vs. 4 x 10 Gb/s lanes
- A single lane per physical port maximizes the number of connected servers or uplinks per switch
- Overall higher port count, utilization and total server interconnect bandwidth vs. 40 GE

Sources: http://www.ieee802.org/3/cfi/0714_1/CFI_01_0714.pdf and <http://www.brighttalk.com/webcast/6205/135253/>

25 GE Developments

BASE-T

- IEEE 25GBASE-T Study Group combined meetings with IEEE P802.3bq 40GBASE-T Task Force
 - 25GBASE-T: 1 x 25 Gb/s over 30 m Cat 8 4-pair twisted-pair copper cabling (ISO/IEC JTC1 SC25 WG3 and TIA TR-42.7)
- 25GBASE-T added to P802.3bq standard in January 2015
- No change in standard schedule for either speed, expected March 2016

25 GE Developments

Copper and Fiber Optic Cables

- IEEE P802.3by 25 Gb/s Ethernet Task Force started in December 2014
 - 25GBASE-KR-S: 1 x 25 Gb/s NRZ 25 GBd over 1 m Megtron 6 backplane (no FEC, BASE-R FEC)
 - 25GBASE-KR: 1 x 25 Gb/s NRZ 25 GBd over 1 m Megtron 6 backplane (no FEC, BASE-R FEC, RS FEC)
 - 25GBASE-CR-S: 1 x 25 Gb/s over 3 m copper twinax cable (no FEC, BASE-R FEC, RS FEC)
 - 25GBASE-CR: 1 x 25 Gb/s over 5 m copper twinax cable (RS FEC)
 - 25GBASE-SR: 1 x 25 Gb/s over 70 m OM3 and 100 m OM4 duplex MMF
 - Auto-negotiation between copper interface types and FECs
 - Passive direct attach cable (DAC) types CA-N (2 m; no FEC, BASE-R FEC, RS FEC), CA-S (3 m; no FEC, BASE-R FEC, RS FEC), and CA-L (5 m; RS FEC)
 - Optional Energy Efficient Ethernet (EEE) support
- Generated Draft 2.2 for Working Group recirculation ballot
- Standard expected in September 2016
- Interfaces expected on the market in 2016+
- Task Force web page <http://www.ieee802.org/3/by/>

| FEC | Latency (ns) | BER |
|--------------------|--------------|------------|
| None | 0 | 10^{-12} |
| BASE-R (Clause 74) | 82 | 10^{-8} |
| RS (Clause 91) | 250 | 10^{-5} |


25 Gigabit Ethernet Consortium

- Founded in July 2014 by Arista, Broadcom, Google, Mellanox and Microsoft after the first 25 GE CFI failed in the IEEE in March 2014
- Developing 25 GE and 50 GE standards outside of IEEE
 - 25 GE: 1 x 25 Gb/s
 - 50 GE: 2 x 25 Gb/s
 - Based on 100GBASE-KR4 and 100GBASE-CR4
- Optional FEC modes
- Optional auto-negotiation
- Specifications only for backplane and twinax copper cable, but does not address or preclude active optical cable or fiber interfaces
- Full draft 1.4 specification only available to members
- More information at <http://25gethernet.org/>



25 GE Technology Reference

| Physical Layer Reach | 1 m Backplane | 3 m Copper Cable | 5 m Copper Cable | 30 m Twisted-Pair | 70 m OM3 / 100 m OM4 |
|-----------------------------|--------------------------------|--------------------------------|--------------------------------|----------------------------|--------------------------------|
| Name | 25GBASE-KR-S 25GBASE-KR | 25GBASE-CR-S | 25GBASE-CR | 25GBASE-T | 25GBASE-SR |
| Standard | September 2016 IEEE 802.3by | September 2016 IEEE 802.3by | September 2016 IEEE 802.3by | March 2016 IEEE 802.3bq | September 2016 IEEE 802.3by |
| Electrical Signaling (Gb/s) | 1 x 25 | 1 x 25 | 1 x 25 | 1 x 25 | 1 x 25 |
| Media Signaling (Gb/s) | 1 x 25 | 1 x 25 | 1 x 25 | 1 x 25 | 1 x 25 |
| Media Type | Backplane | Twinax Copper | Twinax Copper | Cat 8 | Duplex MMF |
| Module Type | Backplane | SFP28 | SFP28 | RJ45 | SFP28 |
| Market Availability | 2016+ | 2016+ | 2016+ | 2016+ | 2016+ |



SFP28 Pluggable Module
(Same Size as SFP and SFP+)

Image courtesy of the SFF Committee.

Agenda

1. Ethernet Speed Evolution
2. What's Next: 2.5 GE and 5 GE
3. What's Next: 25 GE
- 4. What's New: 40 GE**
5. What's New: 100 GE
6. What's Next: 400 GE

40 Gb/s QSFP+ Modules Overview

Quad Small Form-factor Pluggable+

- Created for high density interfaces primarily short reach interfaces for data center applications
 - Small compact form factor enables low power consumption and high density
 - Also used for longer reach 40 GE
- Used for a variety of Ethernet, Fibre Channel and InfiniBand applications
 - 40 GE uses 4 x 10 Gb/s bidirectional channels
- Supports a variety of copper and fiber 40 GE interfaces
 - Breakout from 40 GE to 4 x 10 GE
- Same faceplate size as an XFP but slightly shorter















40GBASE-SR4



40GBASE-LR4

Images courtesy of Finisar.

40 GE QSFP+ Pluggable Modules

| | Data Center Server and Access 40 GE to 4 x 10 GE Breakout | | | Aggregation and Core Native 40 GE | | |
|----------------------|---|---|--|--|---|---|
| Physical Layer Reach | 10 m Passive Copper Cable | 100 m OM3/OM4 | 7 m Passive Copper Cable | 100 m OM3/ 150 m OM4 | 10 km SMF | 40 km SMF |
| Pluggable Module |  10GSFP+Cu |  10GBASE-SR |  40GBASE-CR4 |  40GBASE-SR4 |  40GBASE-LR4 |  40GBASE-ER4 |
| Media |  Integrated Twinax (QSFP+ to 4 x SFP+) |  Parallel MMF (MPO to 4 x Duplex LC) |  Integrated Twinax (QSFP+ to QSFP+) |  Parallel MMF (MPO12) |  Duplex SMF (LC) |  Duplex SMF (LC) |
| Standard | July 2009 SFF-8431 | June 2002 IEEE 802.3ae | June 2010 IEEE 802.3ba | June 2010 IEEE 802.3ba | June 2010 IEEE 802.3ba | February 2015 IEEE 802.3bm |

40GBASE-SR4, 40GBASE-LR4, and 40GBASE-ER4 QSFP+ images courtesy of Finisar.

40 GE Developments

- IEEE 802.3bm 40 Gb/s and 100 Gb/s Operation Over Fiber Optic Cables Task Force started in September 2012
 - 40GBASE-ER4: 4 x 10 Gb/s over 40 km SMF
 - Optional EEE operation for 40 GE and 100 GE fiber interfaces
 - IEEE Std 802.3bm-2015 approved on February 17, 2015
- IEEE P802.3bq 40GBASE-T Task Force started in May 2013
 - 40GBASE-T: 4 x 10 Gb/s over 30 m Cat 8 4-pair twisted-pair copper cabling (ISO/IEC JTC1 SC25 WG3 and TIA TR-42.7)
 - 25GBASE-T: 1 x 25 Gb/s (added January 2015)
 - Optional Energy Efficient Ethernet (EEE) support
 - Generated Draft 2.2 for Working Group recirculation ballot
 - Standard expected in March 2016
 - Task Force web page <http://www.ieee802.org/3/bq/>

40 GE Technology Reference

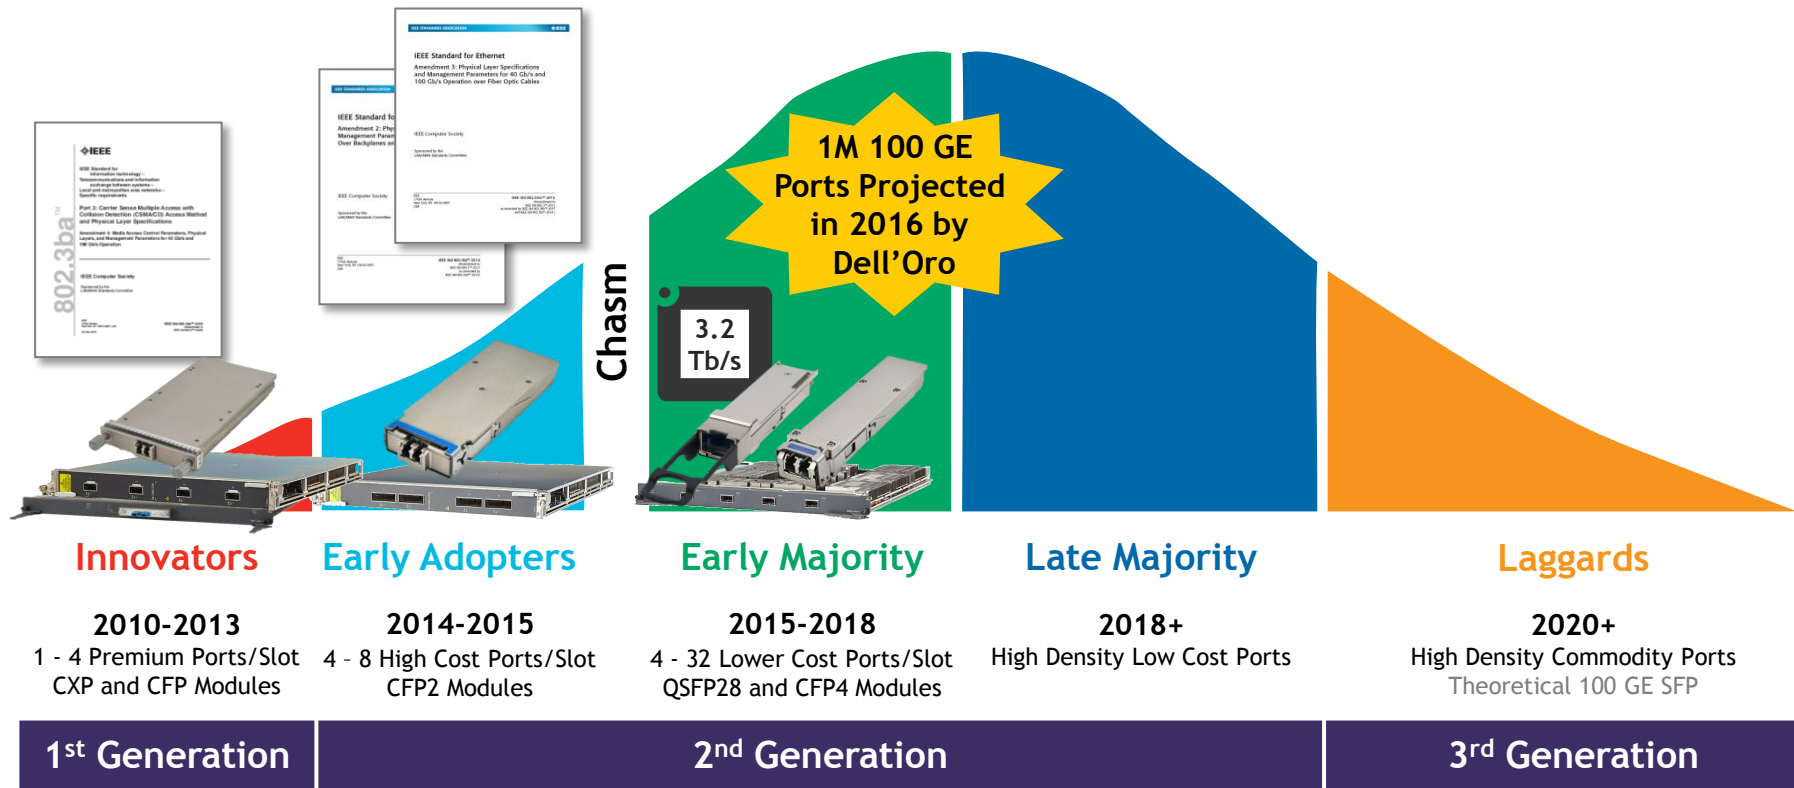
| Physical Layer Reach | 1 m Backplane | 7 m Copper Cable | 30 m Twisted-Pair | 100 m OM3 / 150 m OM4 | 2 km SMF | 10 km SMF | 40 km SMF |
|-----------------------------|------------------------|------------------------|-------------------------|------------------------------|--|-------------------------------|-------------------------------|
| Name | 40GBASE-KR4 | 40GBASE-CR4 | 40GBASE-T | 40GBASE-SR4 | 40GBASE-FR | 40GBASE-LR4 | 40GBASE-ER4 |
| Standard | June 2010 IEEE 802.3ba | June 2010 IEEE 802.3ba | March 2016 IEEE 802.3bq | June 2010 IEEE 802.3ba | March 2011 IEEE 802.3bg | June 2010 IEEE 802.3ba | February 2015 IEEE 802.3bm |
| Electrical Signaling (Gb/s) | 4 x 10 | 4 x 10 | 4 x 10 | 4 x 10 | 4 x 10 | 4 x 10 | 4 x 10 |
| Media Signaling (Gb/s) | 4 x 10 | 4 x 10 | 4 x 10 | 4 x 10 850 nm λ s | 1 x 40 1310 nm λ (RX) 1550 nm λ s (TX, RX) | 4 x 10 1310 nm λ s | 4 x 10 1310 nm λ s |
| Media Type | Backplane | Twinax | Cat 8 | Parallel MMF (MPO12) | Duplex SMF | Duplex SMF | Duplex SMF |
| Module Type | Backplane | QSFP+ | RJ45 | CFP, QSFP+ | CFP | CFP, QSFP+ | QSFP+ |
| Market Availability | None Known | 2010 | 2016+ | 2010 | 2012 | CFP 2010 QSFP+ 2011 | 2015 |

Agenda

1. Ethernet Speed Evolution
2. What's Next: 2.5 GE and 5 GE
3. What's Next: 25 GE
4. What's New: 40 GE
- 5. What's New: 100 GE**
6. What's Next: 400 GE

100 GE Technology Adoption Lifecycle

Crossing the Chasm With 2nd Generation 100 GE



Pluggable module images courtesy of Finisar. Adoption lifecycle curve based on “Crossing the Chasm” by Geoffrey Moore.

1st Generation vs. 2nd Generation 100 GE

- Fundamental 1st generation technology constraints limit higher 100 GE density and lower cost
- Electrical signaling to the CFP
 - 100 Gb/s Attachment Unit Interface (CAUI) uses 10 x 10 Gb/s lanes (CAUI-10)
- Optical signaling on the media
 - 100GBASE-SR10: 10 x 10 Gb/s parallel
 - 10x10 MSA: 10 x 10 Gb/s λ s
 - 100GBASE-LR4 and 100GBASE-ER4: 4 x 25 Gb/s λ s
- CFP module size, complexity and power consumption
- 2nd generation modules based on 4 x 25 Gb/s electrical signaling are available now





Image courtesy of Finisar.

100 GE Developments












Fiber Optic Cables

- IEEE P802.3bm 40 Gb/s and 100 Gb/s Operation Over Fiber Optic Cables Task Force started in September 2012
 - 40GBASE-ER4: 4 x 10 Gb/s over 40 km SMF
 - 100GBASE-SR4: 4 x 25 Gb/s over 70 m OM3 and 100 m OM4 parallel MMF
 - ~~4 x 25 Gb/s over 20 m MMF~~
 - Removed because there is not enough economic or technical advantage vs. existing MMF alternatives
 - ~~4 x 25 Gb/s over 500 m SMF~~
 - Removed due to lack of industry consensus that any of the proposals (CWDM, DMT, PAM-n, PSM4) provided sufficient size, cost and power reduction vs. existing SMF alternatives
 - CAUI-4 electrical signaling to the CFP2, CFP4 and QSFP28
 - Optional Energy Efficient Ethernet (EEE) support for 40 GE and 100 GE fiber interfaces
- IEEE Std 802.3bm-2015 approved on February 17, 2015



100 GE Pluggable Module Evolution

Each Module Increases Density, While Reducing Cost and Power

| | 1 st Generation | | | 2 nd Generation | |
|--|---|---|---|---|--|
| Market Availability | 2010 | 2010 | 2014 | 2015 | 2015 |
| Approximate Module Dimensions (Length x Width to Scale) |  |  |  |  |  |
| Front Panel Density (1 RU) | 4 Ports | 12 Ports | 8 Ports | 22/44 Ports | 16/32 Ports |
| Electrical Signaling (Gb/s) | 10 x 10 CAUI-10 | 10 x 10 CPPI | 10 x 10 CAUI-10 4 x 25 CAUI-4 | 4 x 25 CAUI-4 | 4 x 25 CAUI-4 |
| Media | MMF, SMF | Twinax, MMF | MMF, SMF | MMF, SMF | MMF, SMF |
| Power Consumption (W) | < 24 W (100GBASE-LR4) < 20 W (2 nd Generation CFP) | < 6 W (100GBASE-SR10) | < 12 W (100GBASE-LR4) | 3.5 W | 9 W |
| Industry Standard Modules |  CFP (82 mm Wide) |  CXP (27 mm Wide) |  CFP2 (41.5 mm Wide) |  QSFP28 (18.35 mm Wide) |  CFP4 (21.7 mm Wide) |
| Cisco Proprietary Module | | |  CPAK (34.84 mm Wide) | | |

CXP, CFP and QSFP28 images courtesy of Finisar.

100 GE Pluggable Module Evolution

Graphical View of Module Form Factors






CFP vs. iPhone 6



CFP2 vs. iPhone 6



CFP4 vs. iPhone 6



CFP diagrams courtesy of the CFP MSA. QSFP28 diagram courtesy of the SFF Committee.

27 *Not quite to scale.











COPYRIGHT © 2015 ALCATEL-LUCENT. ALL RIGHTS RESERVED.

Every success has its network

Alcatel-Lucent 

100 GE MSAs

This space is a little crowded...

| | 10x10 MSA | 100G PSM4 MSA | CWDM4 MSA Group | 100G CLR4 Alliance | OpenOptics MSA |
|--------------------------|---|---|---|---|---|
| Pluggable Module | CFP | CFP4, QSFP28 | CFP2, CFP4, QSFP28 | QSFP28 | QSFP28 |
| Media |  Duplex SMF (SC, LC) |  Parallel SMF (MPO12) |  Duplex SMF (SC) |  Duplex SMF (SC) |  Duplex SMF (SC) |
| Physical Layer Reach | 2 km, 10 km, 40 km | 500 m | 2 km | 2 km | > 2 km |
| Optical Signaling (Gb/s) | 10 x 10 1550 nm λ s | 4 x 25 1310 nm λ s | 4 x 25 1310 nm λ s | 4 x 25 1310 nm λ s | 4 x 25 1550 nm λ s |
| Founded | December 2010 | January 2014 | March 2014 | March 2014 | March 2014 |
| Main Supporters | Component Vendors, System Vendors, Network Operators | Component Vendors, System Vendors, Microsoft | Component Vendors, System Vendors | Component Vendors, System Vendors, Network Operators | Ciena, Mellanox, Oracle, Ranovus |
| More Information |  www.10x10msa.org |  www.psm4.org |  www.cwdm4-msa.org |  www.clr4-alliance.org |  www.openopticsmsa.org |

“MSA” means Multisource Agreement.

100 GE Technology Reference

| Physical Layer Reach | 1 m Backplane | 5 m Copper Cable | 7 m Copper Cable | 70 m OM3 / 100 m OM4 | 100 m OM3 / 150 m OM4 | 10 km SMF | 40 km SMF |
|-----------------------------|------------------------------|---------------------------|---------------------------|-------------------------------|------------------------------------|-------------------------------|-------------------------------|
| Name | 100GBASE-KP4 100GBASE-KR4 | 100GBASE-CR4 | 100GBASE-CR10 | 100GBASE-SR4 | 100GBASE-SR10 | 100GBASE-LR4 | 100GBASE-ER4* |
| Standard | June 2014 IEEE 802.3bj | June 2014 IEEE 802.3bj | June 2010 IEEE 802.3ba | February 2015 IEEE 802.3bm | June 2010 IEEE 802.3ba | June 2010 IEEE 802.3ba | June 2010 IEEE 802.3ba |
| Electrical Signaling (Gb/s) | 4 x 25 | 4 x 25 | 10 x 10 | 4 x 25 | 10 x 10 | 10 x 10 | 10 x 10 |
| Media Signaling (Gb/s) | 4 x 25 NRZ and PAM-4 | 4 x 25 | 10 x 10 | 4 x 25 850 nm λ s | 10 x 10 850 nm λ s | 4 x 25 1550 nm λ s | 4 x 25 1550 nm λ s |
| Media Type | Backplane | Twinax Copper | Twinax Copper | Parallel MMF (MPO12) | Parallel MMF (MPO24) | Duplex SMF (SC, LC) | Duplex SMF (SC, LC) |
| Module Type | Backplane | CFP2, CFP4, QSFP28 | CXP, CFP2, CFP4, QSFP28 | CFP2, CFP4, CPAK, QSFP28 | CFP, CFP2, CFP4, CPAK, CXP, QSFP28 | CFP, CFP2, CFP4, CPAK, QSFP28 | CFP, CFP2, CFP4 |
| Market Availability | None Known | 2014 | 2010 | 2015 | 2012 | 2010 | 2012 |

*100GBASE-ER4-lite (ITU-T application code 4L1-9D1F) has different optical specifications and uses the KR4 FEC.


Agenda

1. Ethernet Speed Evolution
2. What's Next: 2.5 GE and 5 GE
3. What's Next: 25 GE
4. What's New: 40 GE
5. What's New: 100 GE
- 6. What's Next: 400 GE**

Industry Challenges for 400 GE and Beyond

Solutions are Good, Fast, or Cheap - Pick Any Two

Ethernet at Terabit Speeds is
Technically and Economically
Impractical Today



Economics Dictate the Solution

IEEE Provides an Open Industry Forum to Make Decisions

IEEE 802.3 BWA Ad Hoc Report: http://www.ieee802.org/3/ad_hoc/bwa/BWA_Report.pdf

400 GE Developments

- IEEE P802.3bs 400 Gb/s Ethernet Task Force started in March 2014
 - 400GBASE-SR16: 16 x 25 Gb/s over 100 m parallel MMF (based on 100GBASE-SR4)
 - 400GBASE-DR4: 4 x 100 Gb/s over 500 m parallel SMF
 - 400GBASE-FR8: 8 x 50 Gb/s over 2 km duplex SMF
 - 400GBASE-LR8: 8 x 50 Gb/s over 10 km duplex SMF
 - Electrical interfaces: 25 Gb/s (NRZ) and 50 Gb/s (PAM-4)
- Strong desire to support 400 GE to 4 x 100 GE breakout functionality based on 40 GE to 4 x 10 GE success
- Generated Draft 1.0 for Task Force review
- Standard expected in December 2017 (was March 2017)
- First interfaces expected to be available in 2017+
- Task Force web page <http://www.ieee802.org/3/bs/>

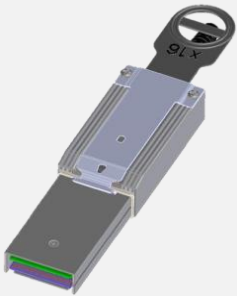
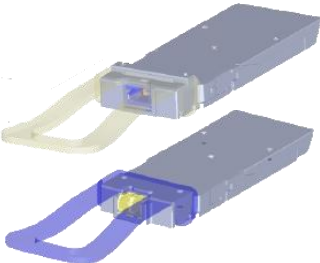
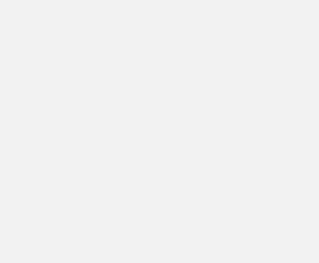
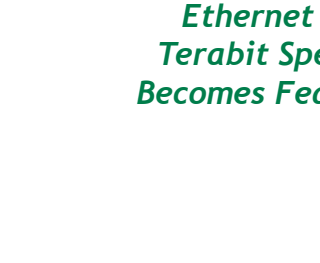
Many Technology Options to Consider

| Physical Layer Reach | 100 m MMF | 500 m SMF | 2 km SMF | 10 km SMF |
|----------------------|--------------------------|-----------------------------|------------------------|------------------------|
| 25 Gb/s NRZ | 16 λ x 16 MMF | | | |
| 50 Gb/s NRZ | | BiDi λ s x 8 SMF | 2 λ x 4 SMF | 8 λ x 1 SMF |
| 50 Gb/s PAM-4 | | | 8 λ x 1 SMF | 8 λ x 1 SMF |
| 100 Gb/s PAM-4 | | 1 λ x 4 SMF | 4 λ x 1 SMF | |
| 100 Gb/s DMT | | | | 4 λ x 1 SMF |

Four Ways to Go Faster: Signaling Speed, Modulation, Number of λ s, Number of Fibers

400 GE Pluggable Module Evolution Estimates

Each Module Increases Density, While Reducing Cost and Power

| | 1 st Generation | | 2 nd Generation | 3 rd Generation |
|------------------------------|---|---|--|---|
| Market Availability Estimate | ~2016 | ~2016 | ~2017+ | ~2020+ |
| Electrical Signaling | 16 x 25 Gb/s | 16 x 25 Gb/s 8 x 50 Gb/s | 8 x 50 Gb/s | 4 x 100 Gb/s |
| Module |  CDFP Style 3 |  CFP8 (~CFP2 Size, 40 mm Wide) |  CFP16 (~CFP4 Size) and/or QSFP-DD |  CFP16 and/or QSFP112 <i>Ethernet at Terabit Speeds Becomes Feasible</i> |

CDFP image courtesy of the CDFP MSA. CFP images courtesy of the CFP MSA.

What's Next in the IEEE?

- Three potential new Study Groups could be formed based on CFIs (Calls for Interest) for new higher speed projects at the November 2015 plenary meeting
- Single lane 50 Gb/s, next generation 100 Gb/s and 200 Gb/s Ethernet
 - 50 GE single lane
 - n x 50 Gb/s for 2 x 50 Gb/s 100 GE, and 4 x 50 Gb/s 200 GE
 - http://ieee802.org/3/cfi/request_1115_1.html
- 25 Gb/s single mode fibre
 - 2 km, 10 km, 40 km interfaces
 - http://ieee802.org/3/cfi/request_1115_2.html

Ethernet Speed Evolution Futures

- Ethernet continues to evolve to meet new and diverse market requirements
- Different new speeds are needed for different new applications
- Old 10x performance for 3x cost model doesn't work anymore as we get to higher speeds
 - 10 ME \Rightarrow 100 ME \Rightarrow 1 GE \Rightarrow 10 GE \Rightarrow 100 GE
- Current best technical and economic solutions are 4x to 8x the highest lane rate
 - 4 x 10 Gb/s for 40 GE
 - 4 x 25 Gb/s for 100 GE
 - 8 x 50 Gb/s for 400 GE
- New technology based on 50 Gb/s could be the basis of a new generation of speeds
 - 1 x 50 Gb/s for 50 GE
 - 2 x 50 Gb/s for 100 GE
 - 4 x 50 Gb/s for 200 GE

Ethernet Speed Evolution Summary

- 2.5 GE and 5 GE is coming soon for higher speed Cat 5e/6 applications
- 10 GE is being widely deployed in every part of the network
- 25 GE is coming soon for server and ToR applications
- 40 GE is increasingly deployed in data center networks
 - Popular for 40 GE and 4 x 10 GE breakout
- 100 GE has transitioned to 2nd generation technology with CFP2, CFP4 and QSFP28
 - Still at least a generation away from 100 Gb/s serial signaling
- 400 GE development is well under way and will leverage 100 GE technology
- Ethernet at Terabit speeds is still unfeasible in the near future, but we'll get there eventually (2020+)

More Information

- IEEE P802.3bz 2.5/5GBASE-T Task Force
 - <http://www.ieee802.org/3/bz/>
- IEEE P802.3.by 25 Gb/s Ethernet Task Force
 - <http://www.ieee802.org/3/by/>
- IEEE 802.3 25GBASE-T PHY Study Group
 - <http://www.ieee802.org/3/25GBASET/>
- IEEE P802.3bq 40GBASE-T Task Force
 - <http://www.ieee802.org/3/bq/>
- IEEE P802.3bs 400 Gb/s Ethernet Task Force
 - <http://www.ieee802.org/3/bs/>
- CFP MSA
 - <http://www.cfp-msa.org/>
- SFF Committee
 - <http://www.sffcommittee.com/>
- CDFP MSA
 - <http://www.cdfp-msa.com/>
- µQSFP MSA
 - <http://www.microqsfp.com/>
- COBO
 - <http://onboardoptics.org/>

Questions?

Acknowledgements

Chris Cole, Finisar

John D'Ambrosia, Dell Networking

Scott Kipp, Brocade

Steve Trowbridge, Alcatel-Lucent

Lots of reference slides are next...

Key Industry Developments for the Next Couple of Years

Making Ethernet Faster and Cheaper

Ethernet Standards



- P802.3bz 2.5/5GBASE-T Task Force
- P802.3by 25 Gb/s Ethernet Task Force, 25GBASE-T Study Group
- P802.3bq 40GBASE-T Task Force
- P802.3bs 400 Gb/s Ethernet Task Force

Ethernet MSAs and Consortia



- 100 GE 2 km, 10 km, 40 km
- 100 GE 500 m
- 100 GE 2 km
- 100 GE 2 km



MGBASE-T Alliance



- 100 GE 2 km
- 2.5 GE and 5 GE
- 25 GE and 50 GE

Pluggable Module MSAs



Next Generation Form Factors

- CFP4 (100 GE), CFP8, CFP16 (400 GE)
- QSFP28 (100 GE)
- μ QSFP (25 GE, 50 GE, 100 GE)



Component Standards



Electrical Signaling from ASICs to Pluggable Modules

- 50 Gb/s (CEI-56G-VSR-PAM and CEI-56G-MR-PAM)

“MSA” means Multisource Agreement.


MPO Cable Assemblies

High Density Ribbon Fiber Cabling

- 40 GE and 100 GE short reach pluggable modules use a Multifiber Push-On (MPO) cable assembly to interconnect network devices
 - Also called MTP by US Conec
- Widely available in a variety of high density multimode fiber (MMF) and single-mode fiber (SMF) cabling options for data centers
 - MPO to MPO
 - MPO cassette for patch panels with into LC, SC, etc
 - Keyed connectors maintain correct transmit/receive orientation
- 40GBASE-SR4 uses a 12-fiber OM3 or OM4 MMF MPO12 cable
 - 8 fibers used, left 4 for transmit and right 4 for receive
 - 4 middle fibers are unused
- 100GBASE-SR10 uses a 24-fiber OM3 or OM4 MMF MPO24 cable
 - 20 fibers used, top middle 10 for receive and bottom middle 10 for transmit
 - 2 fibers on each end are unused



MPO12 Cable Connector




100GBASE-SR10 CXP

40 GE Transmission

Multimode and Single-mode Fiber

- Multimode ribbon fiber


- Used for distances of 100 m on OM3 and 150 m on OM4 MMF
- Data is sent using multiple 850 nm lasers transmitting over multiple parallel fibers
- MPO cables provide multiple separate transmit and receive strands of multimode fiber in a ribbon cable assembly



4 x 10 Gb/s Over Parallel MMF
40GBASE-SR4

- Single-mode duplex fiber

- Used for distances of 2 km, 10 km and 40 km on standard duplex SMF
- WDM component in the pluggable module multiplexes four transmit λ s over one strand of fiber and four receive λ s over the other strand of fiber in the 1310 nm CWDM band for 40GBASE-LR4 and 40GBASE-ER4
- 40 Gb/s serial transmit over one strand of fiber and receive over the other strand of fiber on one 1550 nm λ is used for 40GBASE-FR



4 x 10 Gb/s Over Duplex SMF
40GBASE-LR4 and 40GBASE-ER4

100 GE Transmission


Multimode and Single-mode Fiber

- Multimode ribbon fiber


- Used for distances of 100 m on OM3 and 150 m on OM4 MMF
- Data is sent using multiple 850 nm lasers transmitting over multiple parallel fibers
- MPO cables provide multiple separate transmit and receive strands of multimode fiber in a ribbon cable assembly

- Single-mode duplex fiber


- Used for distances of 2 km, 10 km and 40 km on standard duplex SMF
- WDM component in the pluggable module multiplexes all transmit λ s over one strand of fiber and all receive λ s over the other strand of fiber
 - 10x10 MSA standards use 10 x 10 Gb/s λ s in the 1550 nm DWDM band
 - IEEE standards use 4 x 25 Gb/s λ s in the 1310 nm CWDM band



10 x 10 Gb/s Over Parallel MMF
100GBASE-SR10



10 x 10 Gb/s Over Duplex SMF
10x10-2km, 10x10-10km and 10x10-40km

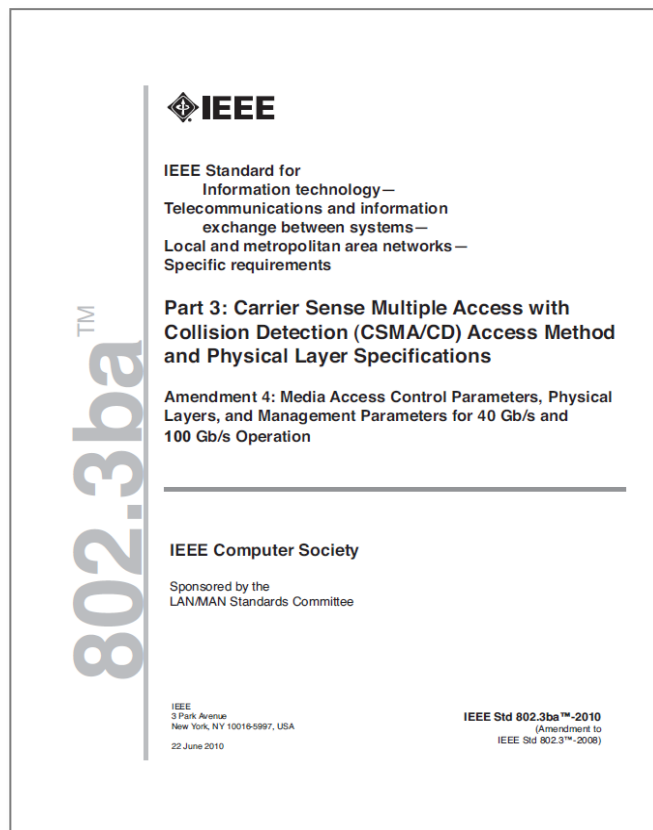


4 x 25 Gb/s Over Duplex SMF
100GBASE-LR4 and 100GBASE-ER4

Ethernet Standards Development Summary

Continuing Technology Evolution

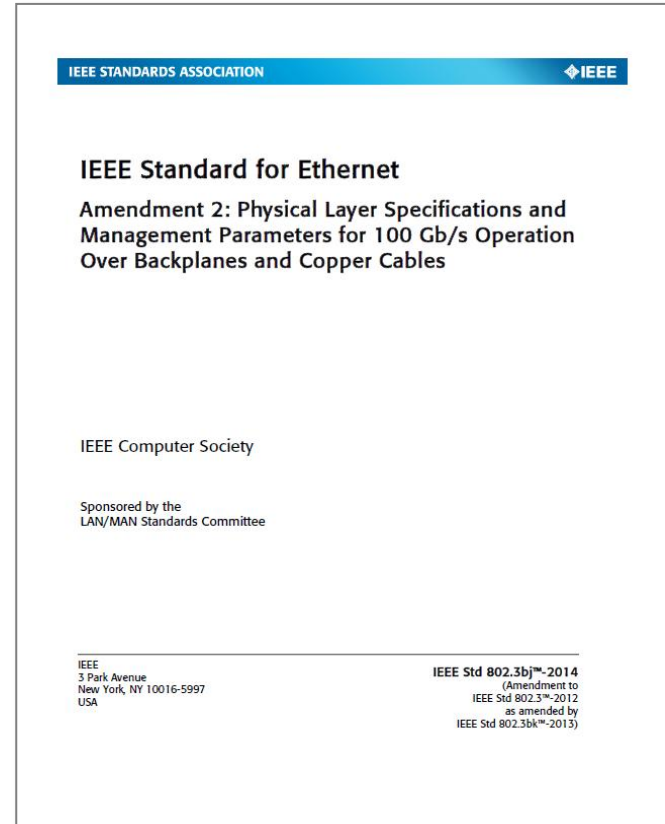
- IEEE 802.3ba-2010 standard for 40 GE and 100 GE approved June 17, 2010
 - 340 pages added to IEEE Std 802.3-2012
- Shipping 1st generation media, test equipment, router interfaces, and optical transport gear in 2011/2012
 - Mature, interoperable technology with broad vendor support
- 2nd generation technology is finished and available on the market now
- 400 GE under development as the next Ethernet speed
 - Expected on the market in 2017+



100 GE Developments

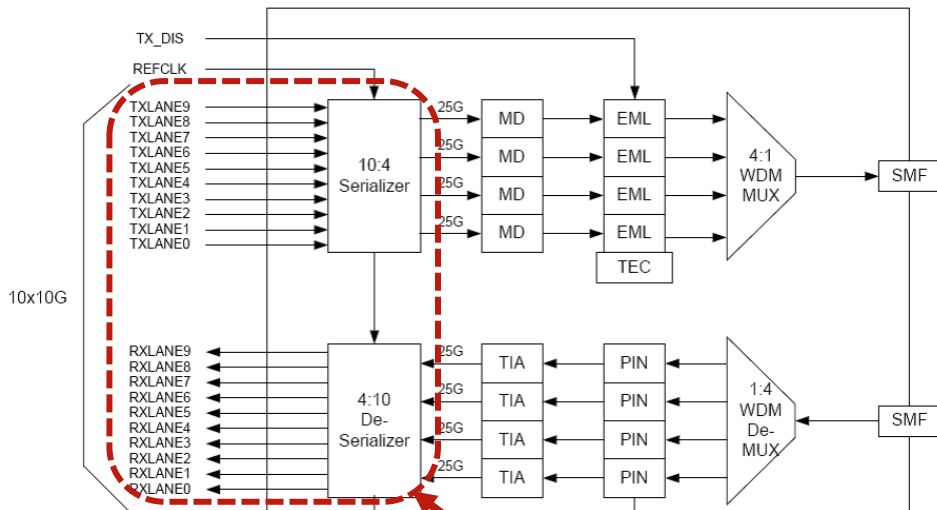
Backplane and Copper Cable

- IEEE P802.3bj 100 Gb/s Backplane and Copper Cable Task Force started in September 2011
 - 100GBASE-KR4: 4 x 25 Gb/s NRZ 25 GBd over 1 m Megtron 6 backplane
 - 100GBASE-KP4: 4 x 25 Gb/s PAM-4 12.5 GBd over 1 m enhanced FR4 backplane
 - 100GBASE-CR4: 4 x 25 Gb/s over 5 m copper twinax cable
 - Optional Energy Efficient Ethernet (EEE) operation for 40 GE and 100 GE backplane links and copper cable interfaces
- IEEE Std 802.3bj-2014 approved on June 12, 2014



1st Generation vs 2nd Generation 100 GE Signaling

1st Generation 100 GE 10 x 10 Gb/s Electrical and 4 x 25 Gb/s Optical



10 Gb/s Electrical Signaling and 10:4 Gearbox Adds Complexity, Cost, Space, and Consumes Power

2nd Generation 100 GE 4 x 25 Gb/s Electrical and Optical

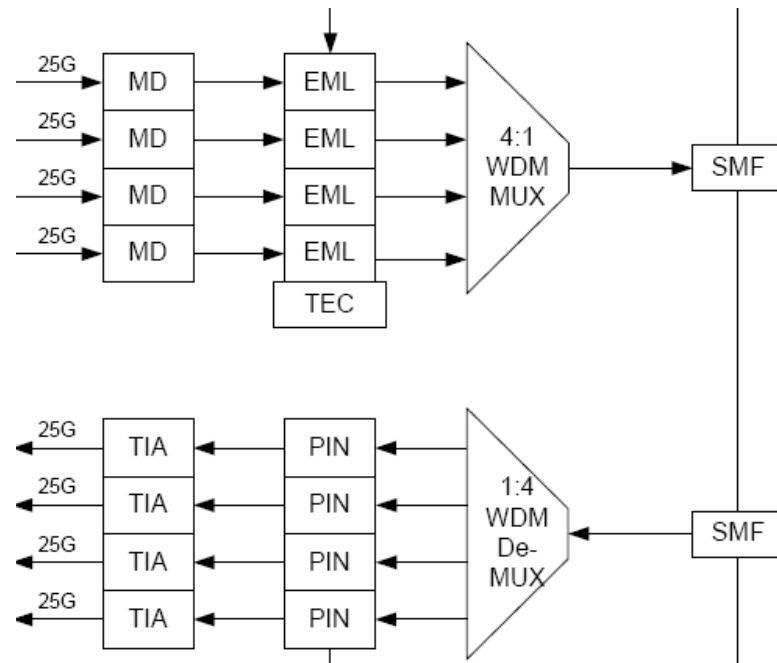



Diagram source: http://grouper.ieee.org/groups/802/3/ba/public/jul08/cole_03_0708.pdf


CFP2 Module Overview

- CFP2 supports electrical lanes that can run at multiple speeds
 - 10 x 10 Gb/s lanes (CAUI-10) for 100 GE
 - 8 x 25 Gb/s lanes (CAUI-4) for 100 GE
- Optional external universal gearbox can convert electrical signaling so that all current IEEE 100 GE standards can be supported in the CFP2 module
- Smaller size and lower cost, complexity and power consumption than the CFP




CFP4 Module Overview

- CFP4 uses 4 x 25 Gb/s electrical lanes (CAUI-4) for 100 GE
- Supports LC, MPO12, and MPO24 connectors for SMF and MMF
- Smaller size and lower cost, complexity and power consumption than the CFP2
- Slightly larger than the QSFP28, but with two key differences
 - The CFP4 can support up to 9 W power and the QSFP28 is limited to 3.5 W, so we may not see longer reach interfaces in the QSFP28 (100GBASE-ER4?)
 - The management interface for the CFP4 is MDIO but is a simpler I²C for the QSFP28, so the QSFP28 may not be able to support other applications such as coherent optics (ACO)




CFP Module Evolution for 100 GE and 400 GE



Higher Density Cages and Front Panel Density



CFP2
8 Ports Per Card
Front Panel Density



CFP4
16 Ports Per Card
Front Panel Density



Diagrams courtesy of the CFP MSA.

IEEE Bandwidth Requirement Projections






Diagram source: http://www.ieee802.org/3/hssg/public/nov07/HSSG_Tutorial_1107.zip

400 Gb/s CDFP Module Overview

Designed to Support 4 Tb/s Per Slot



- CD = 400 in Roman numerals, C = 100 and D = 500
- Optimized for short reaches and targeted for distances up to 500 m
- Supports copper cables, active optical cables and transceivers

Images courtesy of TE Connectivity and the CDFP MSA.

OPEN template

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

The information presented is subject to change without notice.
Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

This slide must be kept when distributed externally.