ethernet alliance

ETHERNET 202: 10GBASE-T REVAMPED

Ron Cates
PLX Technology

Valerie Maguire
Siemon

June 26, 2012

ethernet alliance

THE PRESENTERS

Ron Cates

Vice President of Marketing, Networking Products, at PLX Technology.

Prior to PLX, he was Sr. Vice President and General Manager of Wide Area Networking Products at Mindspeed Technologies. Mr. Cates has over 30 years of experience in the semiconductor industry and holds BSEE and MSEE degrees from the University of California at Los Angeles and an MBA from San Diego State University.

Valerie Maguire

- Vice Chair of the TIA TR-42 Telecommunications Cabling Committee
- Vice Chair of the TIA TR-42.7 Copper Cabling Subcommittee
- TIA Liaison to IEEE 802.3 Working Group
- Treasurer of IEEE 802.3 Working Group

ethernet alliance

DISCLAIMER

The views WE ARE expressing in this presentation are our own personal views and should not be considered the views or positions of the Ethernet Alliance®.

AGENDA

- Introduction to Ethernet Alliance
- An introduction to 10GBASE-T technology
- Appropriate Cabling for 10GBASE-T
- EMI Mitigation and Power Saving Techniques for 10GBASE-T

ethernet alliance

Ethernet Alliance Mission

- To promote industry awareness, acceptance and advancement of technology and products based on, or dependent upon, both existing and emerging IEEE 802 Ethernet standards and their management.
- To accelerate industry adoption and remove barriers to market entry by providing a cohesive, market responsive, industry voice.
- Provide resources to establish and demonstrate multi-vendor interoperability.

ethernet alliance

Expand Ethernet Ecosystem

- Facilitate interop testing
- Expand the market
- Go global

Support Ethernet Development

- Support consensus building
- Host Technology Exploration Forums (TEFs)
- Team with other orgs

Promote Ethernet

Marketing

Education

ethernet alliance

UNIVERSITY OF ETHERNET CURRICULUM

- Completed and available online
- Planned
- Concept

Ethernet 101: Introduction to Ethernet

Physical Layer x00 Series	Protocols x10 Series	Applications x20 Series	Products x30 Series
Ethernet 102: The Physical Layer Of Ethernet	**Ethernet 111:** 802.1:Protocols Of Ethernet	**Ethernet 121:** The Applications Of Ethernet	**Ethernet 131:** Ethernet Products
Ethernet 202: 10GBASE-T Revamped	**Ethernet 211:** Data Center Convergence	**Ethernet 221:** Data Center Applications	**Ethernet 231:** Ethernet Switches
Ethernet 301: 40/100GbE Fiber Cabling and Migration Practices	**Ethernet 311:** Congestion Notification	**Ethernet 321:** Industrial Applications	**Ethernet 331:** Ethernet Server Adapters

ethernet alliance

Cloud computing

Unified data/storage connectivity

Server virtualization

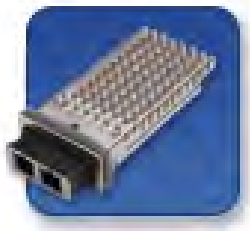

ethernet alliance

IEEE Standard Based:

10GBASE-SR optical; 850 nm lasers; multi-mode fiber; up to 300m.

10GBASE-LR optical; 1310 nm lasers; single-mode fiber; up to 10Km.

10GBASE-LRM optical; 1310 nm lasers; multi-mode fiber; up to 260m.

10GBASE-ER optical; 1550 nm lasers; single mode fiber; up to 40Km.

10GBASE-KX4 four copper backplane lanes with distance up to 1 meter.

10GBASE-KR single backplane lane with distance up to 1 meter.

10GBASE-T twisted pair copper cabling with distance up to 100 meters.

SFF Committee Based:

SFP+ Direct Attach

Connectorized twin-ax cable assembly which connects directly into a SFP+ module housing; up to 7m.

SFF is the Small Form Factor committee, an electronics industry group formed to develop specifications

- Ratified in June 2006 as IEEE 802.3an
- Hallmarks:
  - Uses twisted-pair cable (4-pairs)
    - Categories 6 up to 55m; 6A, 7, and 7A up to 100m
  - RJ-45 connectors
  - Backwards compatible with prior BASE-T generations
  - Full duplex transmission (2.5Gbps per twisted-pair)
  - 128 DSQ line code limits bandwidth to 400MHz
  - DSP based line equalization, cross-talk and echo cancellation
  - LDPC FEC (Low-density parity-check forward error correction)

ethernet alliance

SUPPORTED 10GBASE-T LINK SEGMENT DISTANCES

	Distance	Cabling References
Category 6	55m – 100m[1]	ISO/IEC TR 24750 / TIA TSB-155
Category 6: unscreened	55m	ISO/IEC TR 24750 / TIA TSB-155
Category 6: screened[2]	100m	ISO/IEC TR 24750 / TIA TSB-155
Category 6A	100m	Amendment 1 to ISO/IEC 11801:2002 / ANSI/TIA-568-B.2-10
Category 7	100m	ISO/IEC TR 24750

[1] Link segments must meet specified alien crosstalk to insertion loss requirements
[2] "screened" and "shielded" are terms that are used interchangeably in some geographies
Note: TIA recommends category 6A for greenfield installations

Source: IEEE 802.3-2008

TIA-568-C.2 100-M, 4-CONN CHANNEL MODEL

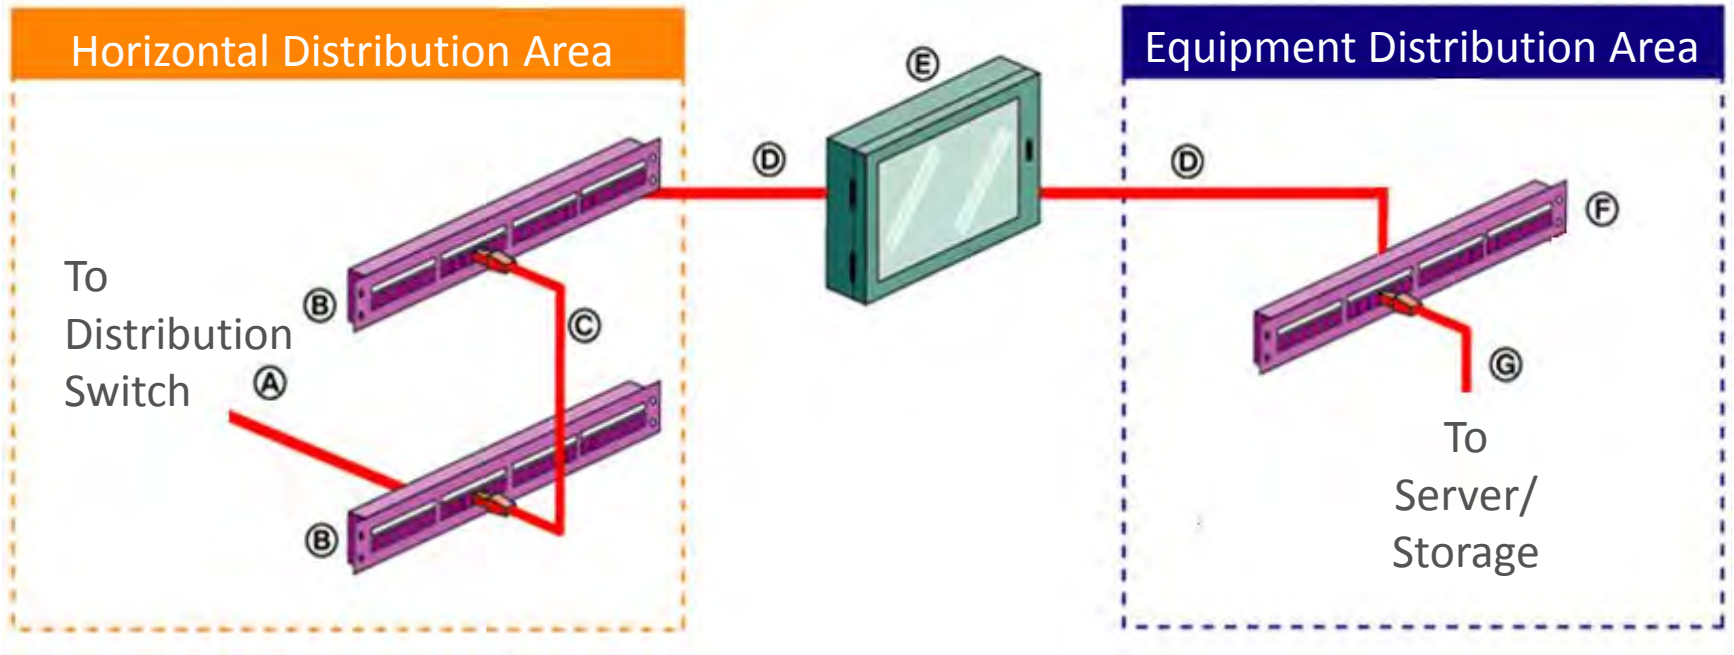

A. Equipment Cord
B. Full Cross-Connect
C. Patch Cord
D. Horizontal cable 90m (295 ft.) max.
E. Consolidation Point (optional)
F. Equipment Outlet
G. Patch Cord

End-of-Row Structured Cabling Topology

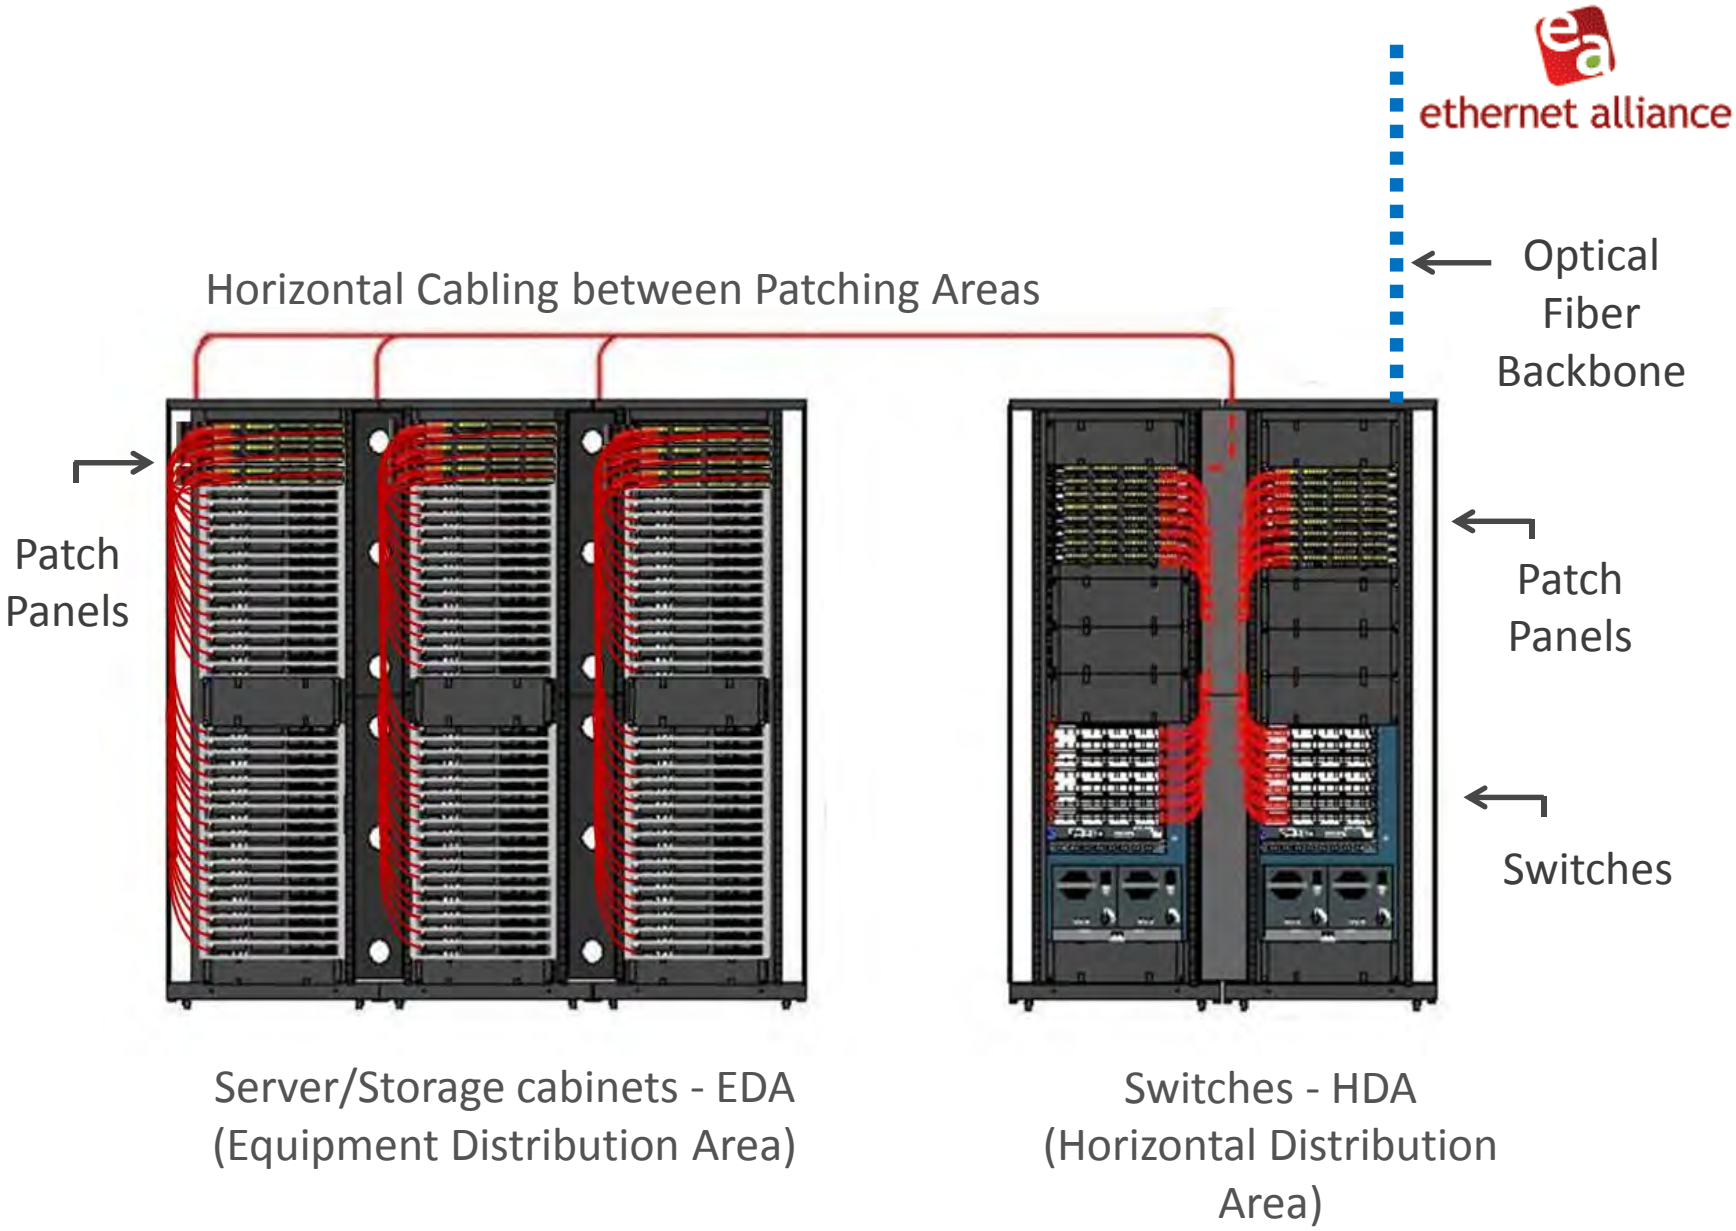

ethernet alliance

- Backwards Compatibility with older BASE-T standards
  - Allows incremental upgrading
  - "Future Proof" data center infrastructure
- 100m Structured Wiring Reach
  - Allows for data center organization flexibility
    - End of Row switching
    - "Fat Tree" topology with central switching pods
- Use of Inexpensive Cables and Connectors
  - Category 6 or 6A and RJ-45 vs. Optical Fiber or Twin-Ax with SFP+ connectors
- Uniform Transmission Media
  - One type of cable/connector for all situations in data center

- Electro Magnetic Interference (EMI)
- Caused by RF emitters (such as radio transmitters) in the vicinity of the network
- An important subject to discuss in any copper cabling based network
- Early implementations of 10GBASE-T systems were susceptible to EMI
- EMI mitigation is addressed in more modern implementations

- Two distinct EMI immunity tests in the industry from kHz to GHz
  - Telcordia GR1089: 8.5V/m (margined up to 10V/m)
  - CISPR24: 3V/m
- Immunity performance can be classified into in-band (up to ~500MHz) and out-of-band (>~500MHz) performance
  - Out-of-band interferers are rejected using the EMI filter (included on-chip)
  - In-band interferers require active cancellation using fast and accurate signal processing

- Modern 10GBASE-T products have three distinct methods of EMI cancellation built-in:
- Differential cancellation
  - Transient response to first learn interferer is microseconds.
  - No interoperability required.
  - Robustness is good with a Ps=99%
- Common Mode sense cancellation (Requires 5 channel magnetics – extra RBOM cost)
  - Similar to differential cancellation, but uses & requires 5 channel magnetics
  - Transient response to first learn interferer is microseconds.
  - No interoperability required.
  - Robustness is better with a Ps=99.9%

- Fast re-train method
  - Transient response to first learn interferer is 12ms
  - Requires interoperability between PHYs
  - Robustness is excellent with a Ps=99.99%
  - Is a backup mechanism to differential and common mode cancellation

ethernet alliance

- Industry advances in chip technology have rapidly brought down per port power consumption:
  - 3.5 watts @100m and 2W @10m today
- New technology allows integration of MAC/PHY to facilitate LAN-on-Motherboard (LOM)

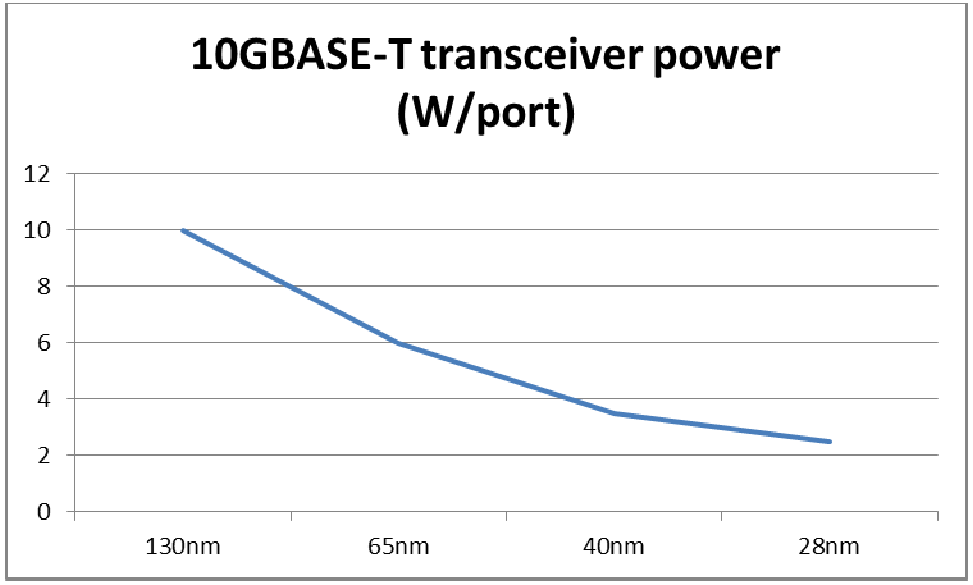

PHY volume production data

- 10GBASE-T transceivers can use these three strategies:
  - Power tuning vs. cable distance
  - Wake on LAN (W-O-L)
  - Energy Efficient Ethernet (EEE)

- Automatic measurement of cable length
- Power reductions when less than 100m cable is detected
  - Reduction of transmit power
  - Reduced number of taps in signal cancellation filters
  - Reduced dynamic range in A/D conversion
  - Lower power mode FEC
- Example in 40nm:

100m	**3.5W**
30m	**2.5W**
<10m	**2.0W**

- System-wide power saving by setting controller and PHY into sleep mode when network traffic on either direction is idle
- In sleep mode, PHY links at 100Mbps line-rate
- PHY or MAC waits for a “magic” packet from link partner to initiate wake-up

- Based on IEEE 802.3az Clause 78
- PHY-level and/or System-level power saving mode when network traffic is idle
  - up to 80%; typically 50% power savings
- Can operate at any supported line-rate (100M/1G/10G)
  - Maintains link at same line-rate
- When system needs to transmit data, packets can be buffered at MAC/controller level while the system wakes up and resume traffic.

ethernet alliance

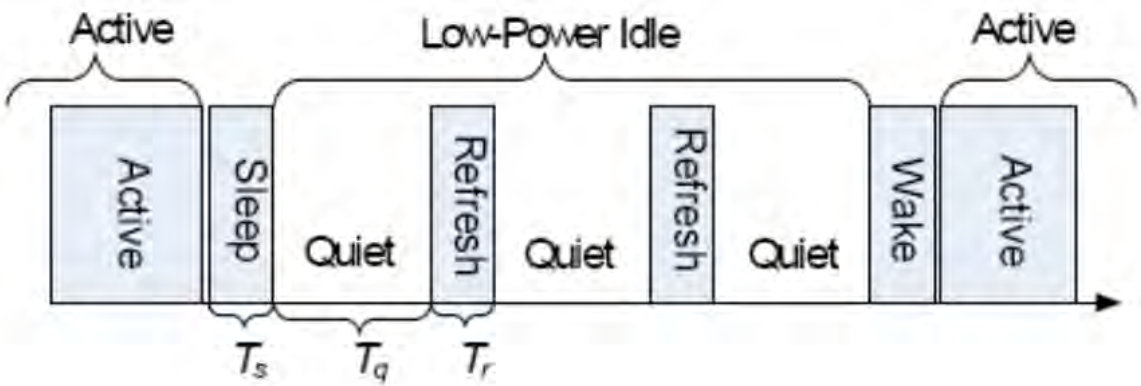

- The LPI mode can scale easily from ~0% utilization to 100%
- The Refresh is required to keep receiver parameters current
  - Timing lock, equalizer coefficients, canceller coefficients, etc.
  - This is critical to enable fast transitions from LPI to active

ethernet alliance

COMPARING EEE AND WOL

	EEE	W-O-L
Low Power Mode	LPI	Sleep
LP Mode Line-Rate	Any	100Mb/s (100Base-TX)
Wait time before entering LP Mode	Configurable down to a few Idle frames	Typically in seconds
Power Consumption in LP mode	~1W (40nm)	< 500mW
Wake-Up Time	Micro-seconds	Seconds
Transparent to Upper Layer Protocols	Yes	No

- Where WOL can be better:
  - Bursty traffic pattern with long idle time in between bursts
- Where EEE can be better:
  - Random traffic pattern with short idle time in between bursts
  - Applications where normal traffic flow is more in one direction and less in the other
  - When security association is required or when the link connection cannot be broken

- ToR switches
- Modular switches
- Servers
- NIC cards
- Storage arrays
- Network appliances
- More coming out soon

- 10GBASE-T migration underway!
- Fundamental advantages:
  - Backwards compatibility with slower BASE-T applications
  - 100m reach
  - Use of inexpensive cables and connectors
  - Uniform transmission media for all data center needs

- Be part of the Voice of Ethernet!
  - Network with Ethernet Thought Leaders
  - Participate in the Debate of Ethernet Futures
  - Contribute to Ethernet Alliance Social Media
- Visibility Through Participation
  - Global Exposure
  - Broad Market Exposure
- Prove Your Interoperability
  - Plugfest
  - Live Demonstrations
- Education

DISCUSSION AND Q&A

ethernet alliance

THANK YOU