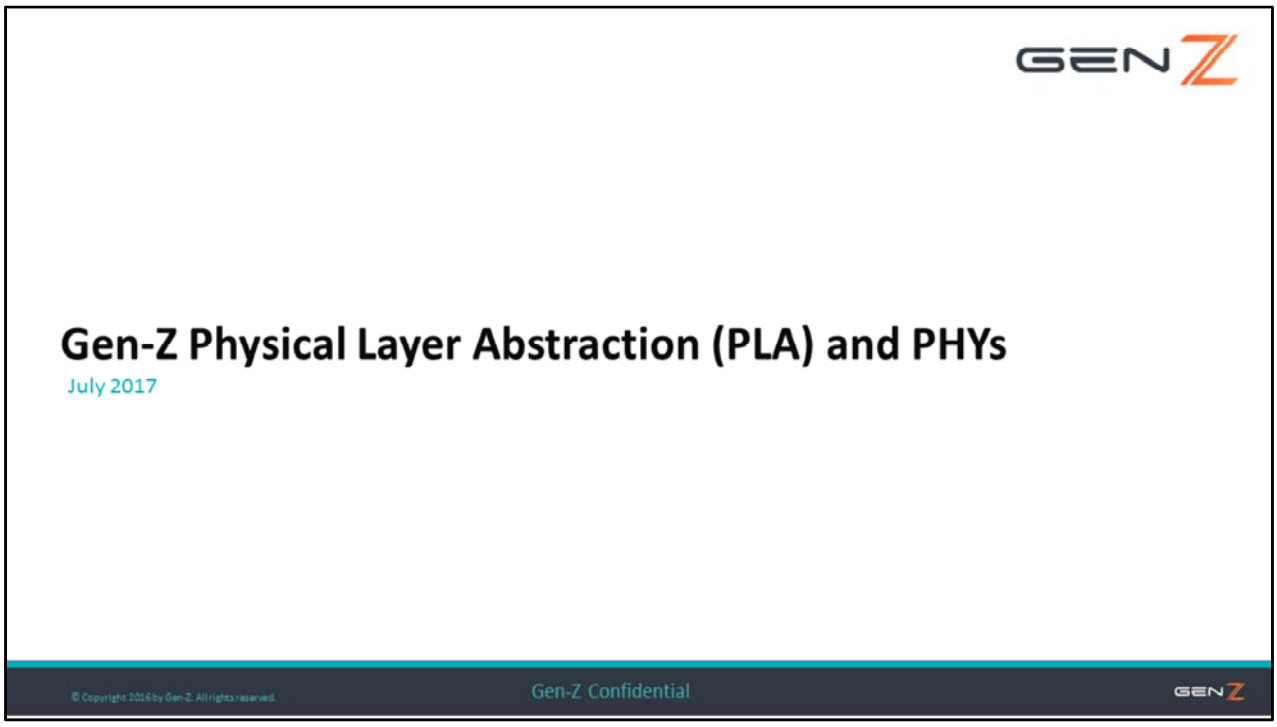

This presentation covers Gen-Z's Physical Layer Abstraction (PLA). The PLA enables an implementation to operate across multiple physical layers without requiring redesign of the higher-level logic.

Disclaimer

This document is provided 'as is' with no warranties whatsoever, including any warranty of merchantability, noninfringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. Gen-Z Consortium disclaims all liability for infringement of proprietary rights, relating to use of information in this document. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Gen-Z is a trademark or registered trademark of the Gen-Z Consortium.

All other product names are trademarks, registered trademarks, or servicemarks of their respective owners.

All material is subject to change at any time at the discretion of the Gen-Z Consortium

http://genzconsortium.org/

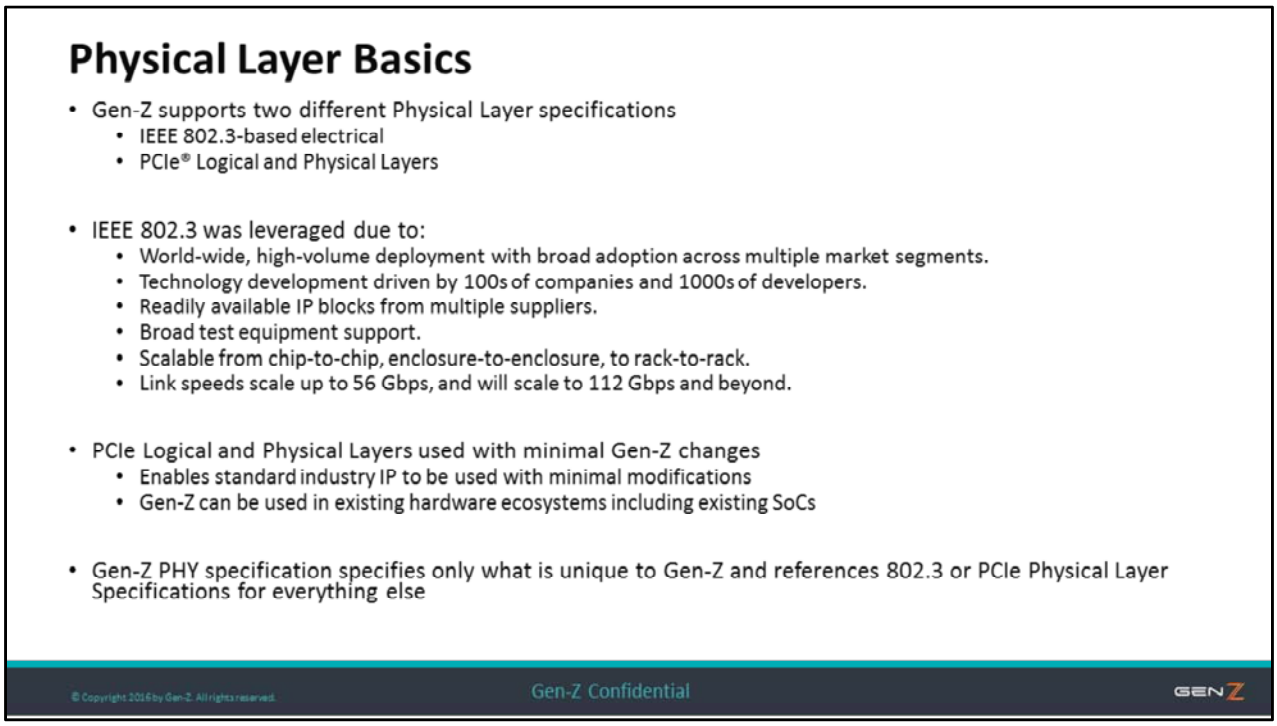

Gen-Z IP developers may choose to implement the IEEE 802.3-based PHY, the PCIe PHY, or both.

802.3 PHY is based on the 25G-KR and 50G-KR PMD specs with enhancements to enable Gen-Z features link asymmetric links widths, lane reversal, etc. Whereas, PCIe Logical and Physical Layers are used with very minimal changes.

IEEE 802.3 was leveraged due to:

IEEE 802.3 physical layer is one of the highest-volume technologies, with broad adoption across multiple market segments and deployed throughout the world.
IEEE 802.3 technology development is driven by 100s of companies and 1000s of developers.
IEEE 802.3 physical layer IP blocks are readily available from multiple suppliers. Further, these IP blocks are constantly evolving to deliver new capabilities and performance levels.
IEEE 802.3 physical layer is supported by many test equipment providers
IEEE 802.3 physical layer can scale from chip-to-chip, enclosure-to-enclosure, to rack-to-rack.
IEEE 802.3 physical layer presently scales up to 56 Gbps, and will scale to 112 Gbps and beyond.

Physical Layer Optimizations

- Gen-Z optimizes the physical layer as the follows:
  - Supports multiple loss budgets tailored to solution needs

PHY Type	Bit Rate	Interface	Reach	Max Insertion Loss	Bit-Error Rate
Gen-Z-E-NRZ-Gen4-PCIe	16 Gbps	Electrical	Medium/ Long	20dB	10^{-12} or better
Gen-Z-E-NRZ-25G-Local	25 Gbps	Electrical	Short/Very Short	10dB	10^{-15} or better
Gen-Z-E-NRZ-25G-Fabric	25 Gbps	Electrical	Long/Very Long	30dB	10^{-12} or better
Gen-Z-E-PAM4-56G-Local	~56 Gbps	Electrical	Short/Very Short	10dB	10^{-15} or better
Gen-Z-E-PAM4-56G-Fabric	~56 Gbps	Electrical	Medium/ Long	20dB	10^{-12} or better

  - No Auto-Negotiation
    - Target link rate is programmed by firmware prior to the start of link training
    - Link width reduction is automatically determined during full rate training
  - Simplified package and board design with automatic lane polarity inversion and link reversal detection

Physical Layer Architecture

- The Physical Layer is divided into two sub-layers
  - Physical Layer Logical
    - Contains link training state machine
    - Performs scrambling, encoding, data striping, serialization, and alignment
    - Connects to the Gen-Z core via the Physical Layer Abstraction Interface
  - Physical Layer Electrical
    - Contains circuits required to drive and recover the signal on the channel
    - Transmitter and receivers must comply with channel-dependent spec requirements
- Physical Layer Power Management controls power state transitions
- Physical Layer Management Control and Status block provide an interface with management firmware

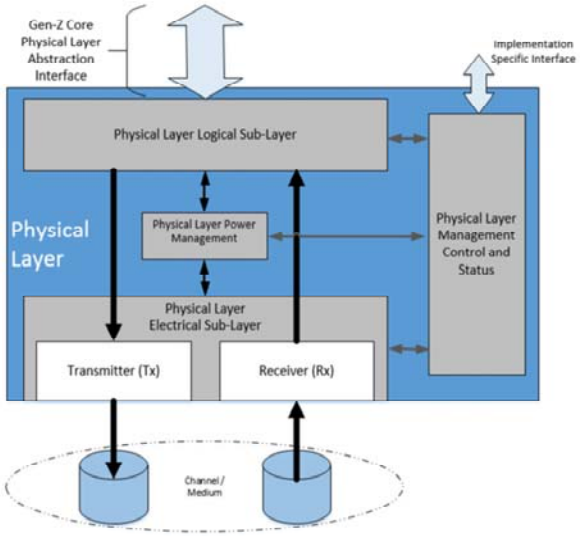

Physical Layer Training State Machine

- Required States:
  - PHY Idle – Transmitters and Receivers are off
  - Signal Detect – Link Endpoint is discovered via electrical idle exit through squelch detection
  - Align – Bit and symbol lock are achieved, lanes are deskewed, unused or broken RX lanes are dropped, link reversal and lane polarity are determined
  - Config – Training sequences are exchanged which notify the transmitter of which lanes are unused
  - PHY Up – Link training is complete and Gen-Z link and protocol packets are send and received
- Optional States:
  - Channel Optimize: Transmitter and receiver settings are dynamically selected based on channel characteristics
  - IBIST: Test patterns are exchanged to determine link quality
  - PHY Low Power – The link enters a low power state such as L1 or burst

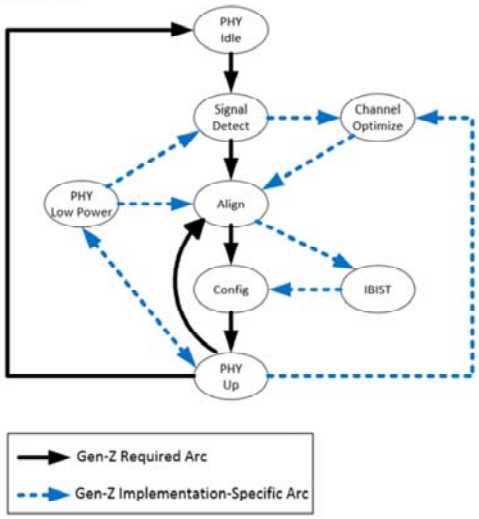

Polarity, Reversal, and Asymmetry

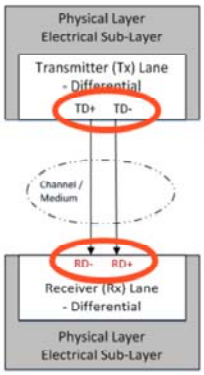

Lane Polarity is supported. TD+ can connect to RD- and TD- can connect to RD+

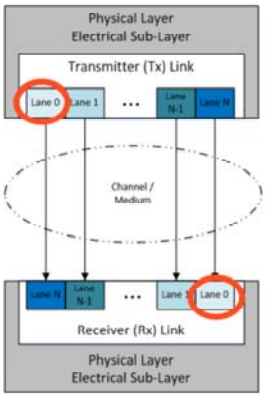

Link Reversal is supported. Tx Lane 0 can connect to Rx Lane N

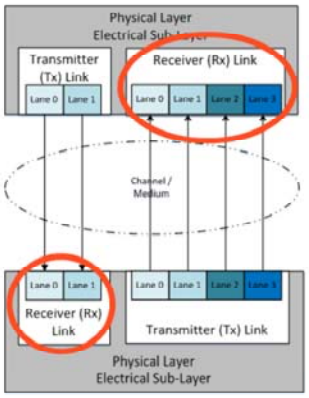

The number of Tx and Rx lanes can be different, which enables asymmetric link bandwidth.

Physical Layer Abstraction (PLA)

- Physical Layer types and implementations are isolated from Gen-Z Core by the Physical Layer Abstraction (PLA)
- PLA provides a common interface between the Physical Layer and the Gen-Z Core regardless of the physical transmission media (electrical or optical), signal transmission type (single-ended or differential), signal modulation (e.g., NRZ or PAM-4) and Physical Layer implementation that is selected
- PLA provides a mechanism for the Link Layer to change the PHY states
- PLA matches the bandwidth of the link with the core. The Core may run faster than the link to provide clock margin or if the link's bandwidth is reduced due to link width reduction

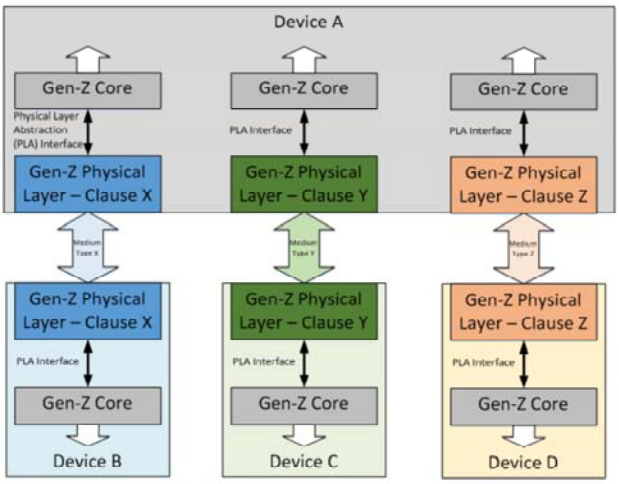

Physical Layer Abstraction Interface

- Control and Status
  - Gen-Z Core requests link state transitions with CORE_STATEREQ
  - The Physical Layer reports the current state with PHY_STATE
  - A Physical Layer state transition is complete when CORE_STATEREQ is the same value as PHY_STATE
- Transmit
  - The Physical Layer asserts PHY_TXDATAREQ when it is ready to receive transmit data and deasserts PHY_TXDATAREQ to backpressure the Core
  - Gen-Z Core sends transmit data on CORE_TXDATA by asserting CORE_TXVALID
- Receive
  - The Physical Layer sends data received from the link to the Core over PHY_RXDATA when PHY_RXVALID is asserted
  - The Physical Layer notifies the Core of an error by asserting one of the PHY_RXERROR* signals

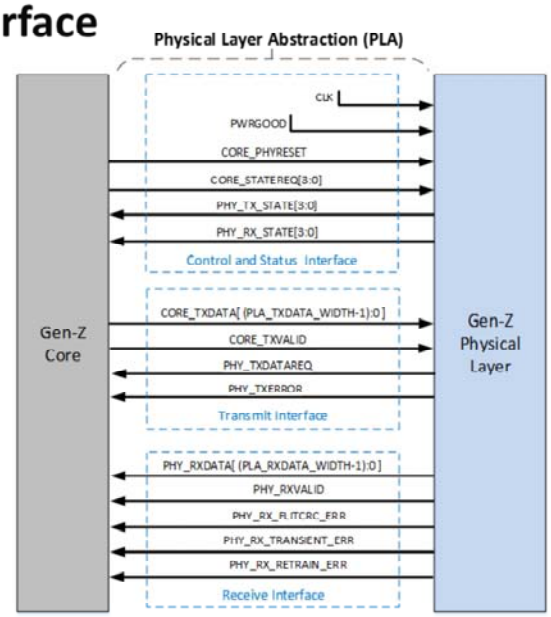

Physical Layer Abstraction Link States

- The Gen-Z Core controls link states with the CORE_STATEREQ signal
- PHY-Down corresponds to the Physical Layer Idle Link Training State
- A PHY-Down->PHY-Up transition of the CORE_STATEREQ signal initiates the Physical Layer Link Training State machine to start link training
- Once the Physical Layer Link Training State machine trains the link, the Physical Layer reports PHY-Up on the PHY_TX_STATE and PHY_RX_STATE signals
- The PHY-Down-Retrain state is used to initiate a retraining which is faster than going all the way back to PHY-Down

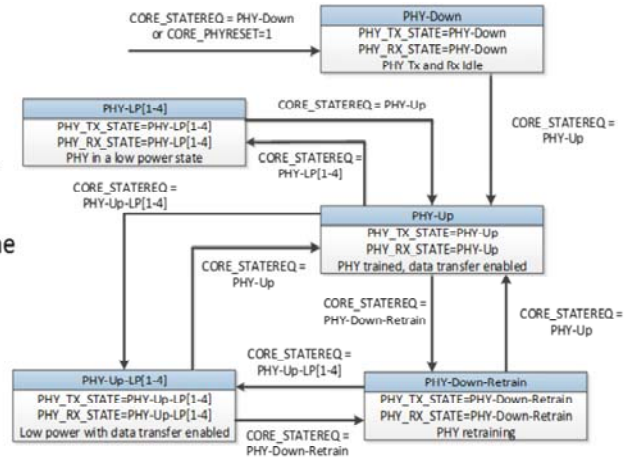

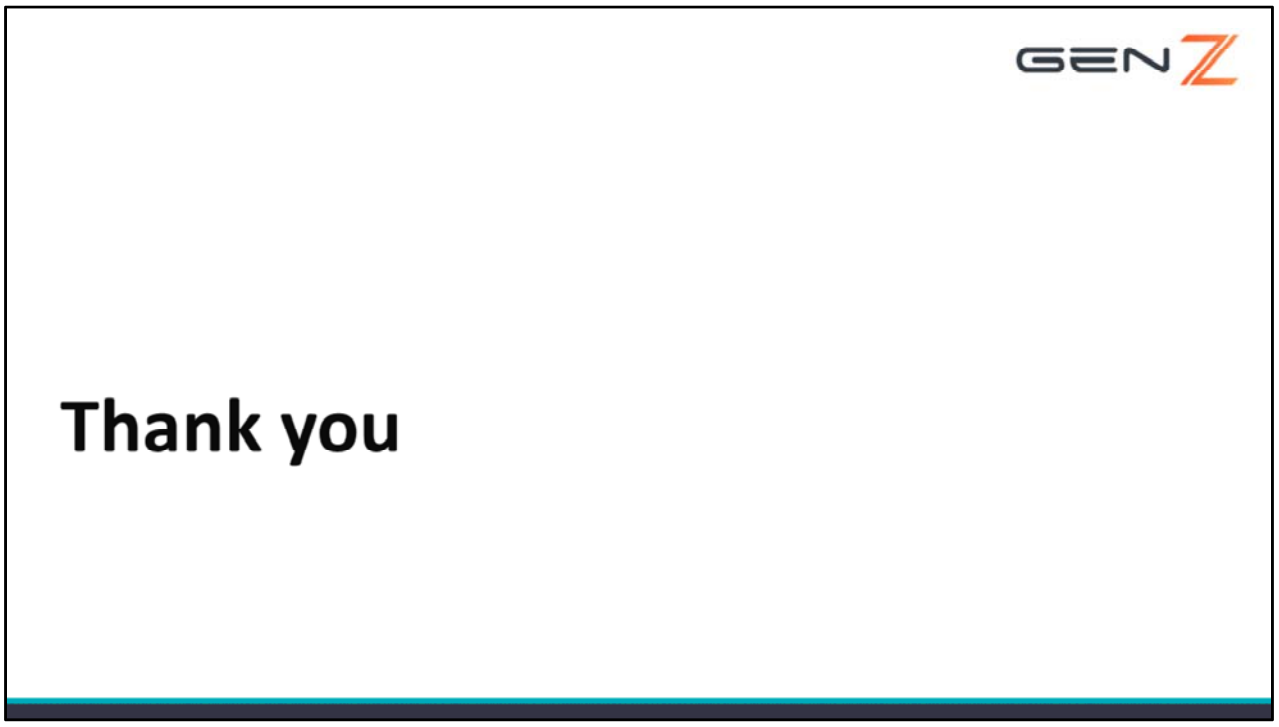

This concludes this presentation. Thank you.