TELEDYNE LECROY
Everywhereyoulook™

800GE vs 1.6TE

WHITE PAPER

Upgrading from 112G to 224G SerDes

Table of Contents

Introduction . . . 3
The path to 1.6Tbps . . . 3
- SerDes speed increase . . . 3
- Lane parallelization . . . 4
- Modulation format . . . 4
- Forward Error Correction (FEC) . . . 5
- Equalization . . . 5
- DSPs . . . 6

Overview of changes in 802.3dj . . . 8
- PCS lanes . . . 8
- PSC FEC . . . 8
- Concatenated FEC . . . 10

Transceivers and cables . . . 10
Implementations . . . 11
- CMIS-LT . . . 12
- Inter Sublayer communication . . . 12

Typical test scenarios . . . 13
- Silicon chip testing . . . 13
- Active cable tests . . . 14
- Network equipment testing . . . 14
- Data center testing . . . 15

Teledyne LeCroy 1.6TE test equipment . . . 15

Introduction

The rapid adoption of Artificial Intelligence (AI) and Machine Learning (ML) models has created an urgent demand for increasingly faster interconnects in data centers. These applications require thousands of processing units to be interconnected and to exchange data during both the training and inference phases. The interconnects typically include in-rack cabling, backplane connections, and chip-to-chip interfaces on processor boards.

Ethernet is the foundational networking technology for data center networks, and the industry is currently advancing toward the next leap in link speed: 1.6Tbps. This doubles the speed of the 800Gbps standard ratified by the IEEE in 2024. 1.6T Ethernet is enabled by using eight lanes of 224Gbps SerDes, compared to 112Gbps per lane used in 800G Ethernet. 200G, 400G and 800G will also be possible using 1, 2, and 4 lanes of 224 Gbps, respectively.

At these ultra-high transmission speeds, signal integrity becomes the most critical Layer 1 concern. As frequencies increase, frequency-dependent losses—including conductor and dielectric losses—rise sharply, degrading signal quality. Discontinuities in the transmission path, such as vias and connectors, introduce impedance mismatches that can cause reflections and radiated emissions, contributing to electromagnetic interference (EMI). Additionally, intersymbol interference (ISI) becomes more pronounced as limited bandwidth causes signal overlap between adjacent bits, making it harder for receivers to distinguish individual symbols. These challenges demand advanced modulation, equalization, error correction, and careful physical design to ensure reliable data transmission.

The 1.6T Ethernet standard is currently being developed under IEEE amendment 802.3dj and will inherit several technical features from earlier 800G and 400G implementations, such as the PAM4 modulation format and other foundational design elements. However, beyond the increase in speed, the new standard introduces significant changes—particularly in areas like Forward Error Correction (FEC), Physical Coding Sublayer (PCS) lane configurations, and equalization schemes. The Optical Internetworking Forum (OIF) is actively working on defining the 224G SerDes in the OIF-CEI-224G standard, which can be used with multiple protocols. Ethernet will use this standard.

In this White Paper, we present a technical overview of the similarities and differences between 800G and 1.6T Ethernet. We will explore the architectural foundations shared by both standards, as well as the key innovations introduced in 1.6T Ethernet. Additionally, we provide practical guidelines for testing and validating the performance of next-generation semiconductors, network elements, and active optical or copper cables designed for 1.6T Ethernet. This includes considerations for signal integrity, compliance testing, and interoperability to ensure robust and reliable system performance.

The path to 1.6Tbps

Transmitting 1.6Tbps means sending a bit every 0.625×10^{-12} second. Electronic circuitry is not capable of working this fast, so several tricks are used to make this happen. Overall, this includes speed increase, parallelization, PAM4 modulation, Forward Error Correction (FEC), and equalization as will be described in the following.

SerDes speed increase

The most straightforward way to increase data throughput is by raising the line transmission speed and, for 1.6Tbps Ethernet the line rate is being raised to 200Gbps compared to the 100Gbps in the previous 802.3df standard.

Although the raw Ethernet data rate at the MAC level is 200Gbps, the actual line rate on the wire is 212.50Gbps. The additional 6.25 % overhead is required for features such as Forward Error Correction (FEC) and 256/257 block coding.

The SerDes used for 200Gbps per lane Ethernet will be based on the OIF-CEI-224G standard which enables line rates up to 224Gbps to support various protocols like Interlaken and fiber channel.

Lane parallelization

A common technique in high-speed communication is parallelization, where a serial data stream is de-multiplexed into multiple parallel lanes. To achieve 1.6Tbps Ethernet, eight lanes of 200Gbps each are used, as illustrated in Figure 1. This represents a doubling of the per lane speed from 100Gbps in the previous 802.3df standard to 200Gbps in 802.3dj. A cable carrying 1.6Tbps will therefore physically contain eight separate connections in each direction.

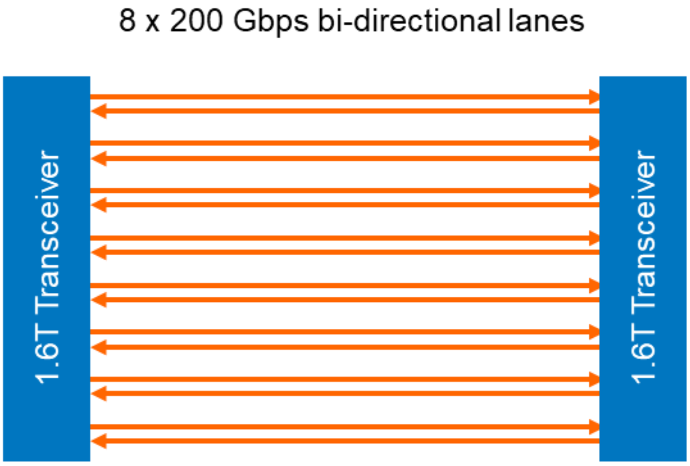

Figure 1: A 1.6 Tbps link is realized using 8 lanes of 200Gbps.

Modulation format

The modulation format for single lane 1Gbps to 25Gbps Ethernet over fiber, twin axial cable and backplanes is Non-Return-to-Zero (NRZ) where each symbol period carries information about a single bit – i.e either a “1” or a “0” as illustrated on Figure 2.

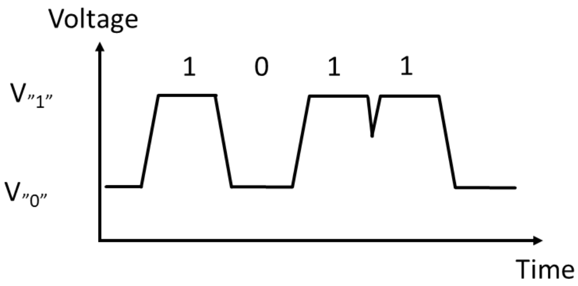

Figure 2: NRZ bit sequence with two voltage levels: V"0" represents a “0” and V"1" represents a “1”.

As the transmission rate increases, it becomes increasingly difficult to transmit NRZ modulation with acceptable signal quality. Therefore, for 50Gbps, 100Gbps and 200Gbps transmission 4-level Pulse Amplitude Modulation (PAM4) is used. The PAM4 coding scheme is illustrated on Figure 3 and uses four voltage levels, where each level represents a two-bit number. In this way the number of bits transmitted per symbol period is doubled compared to NRZ modulation.

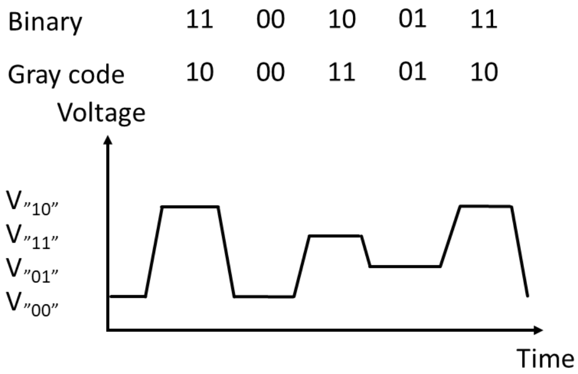

Figure 3: PAM4 Gray-coded bit sequence with four voltage levels that each represent a two-bit value.

The same Gray coding and pre-coding used for 112G is also applied to 224G. As shown in Figure 3, the highest voltage level represents the value “10” rather than “11,” which might seem counterintuitive. This type of coding, known as Gray coding, is used to reduce bit errors. Gray coding ensures that a transition between two adjacent voltage levels results in only one bit changing in the symbol. This means that such a voltage shift can cause at most a single-bit error.

Forward Error Correction (FEC)

Forward Error Correction (FEC) is a coding technique used to detect and correct errors in a bitstream by adding redundant bits and error-checking codes at the transmitter. At the receiver, the FEC decoder uses these additional bits to identify and correct any errors.

The (544, 514, 15) RS-FEC consists of codewords made up of 544 FEC symbols, with each symbol being 10 bits long. This FEC scheme can correct up to 15 symbol errors per codeword, regardless of how many bits are incorrect within each symbol.

The 224G standard uses the same (544, 514, 15) Reed-Solomon FEC (RS-FEC) as 112G. However, the method of distributing bits across the lanes of a 1.6T link has changed. Also, for optical connections, an additional inner FEC is added. These changes will be described in more detail in a subsequent section.

Equalization

When a signal is transmitted from one port to another over a cable, PCB trace or backplane, the signal is degraded by several factors. Electrical cables exhibit frequency-dependent loss at high frequencies due to a combination of skin effect, dielectric losses, and the distributed capacitance and inductance of the cable, which together form a low-pass filter that increasingly attenuates higher-frequency components of the signal. The signal level will also be reduced due to the inherent resistance of the wires. Limitations in the bandwidth will lead to Inter-Symbol Interference (ISI) and inductive coupling between electrical lanes as well as connectors lead to cross talk. Impedance mismatches cause reflections (Return Loss). At both the transmitter and receiver jitter can occur and finally the signal will be degraded by thermal noise.

To address the signal integrity issues mentioned above, equalizers are used at both the transmitter and receiver, as shown in Figure 4.

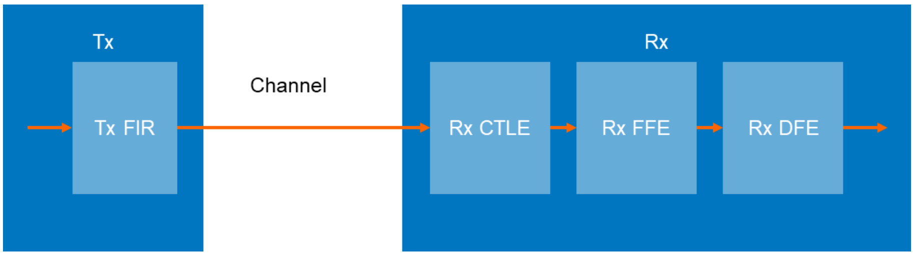

Figure 4: Types of equalizers found in transmitter (Tx) and Receiver (Rx) of 224 G SerDes.

The types of equalizers in 224G SerDes are fundamentally like those in 112G SerDes, but they are generally more complex, featuring additional stages (taps). In the transmitter, a finite impulse response (FIR) equalizer of the Feed-Forward Equalizer (FFE) type is used, primarily to compensate for ISI. FFEs can correct both pre-cursor and post-cursor ISI. In the receiver, the continuous-time linear equalizer (CTLE) compensates for frequency-dependent loss up to the Nyquist frequency, which, for 224Gbps PAM4, is 56 GHz. The receiver's Decision Feedback Equalizer (DFE) and FFE both contribute to ISI mitigation. In a subsequent section, we will examine how these equalizers are commonly implemented at such high speeds.

DSPs

Reaching 224Gbps per lane pushes nearly every component from Analogue to Digital Converters (ADCs) over clocks, equalizers to drivers in the high-speed electronics chain to its physical limits.

Analog-to-Digital Converters (ADCs) face some of the most visible challenges in high-speed data systems. To process 224Gbps signals using PAM4 modulation, ADCs must sample at rates exceeding 100 GSamples/s. This far surpasses the maximum clock speeds achievable by current Silicon technologies.

To overcome this limitation, designers typically employ time-interleaved, parallel ADC architectures, as illustrated in Figure 5. Using this approach allows the ADCs to operate at a reduced clock frequency. For example, a four-way interleaved ADC operating at 25 GSamples/s per channel can collectively achieve an effective sampling rate of 100 GSamples/s.

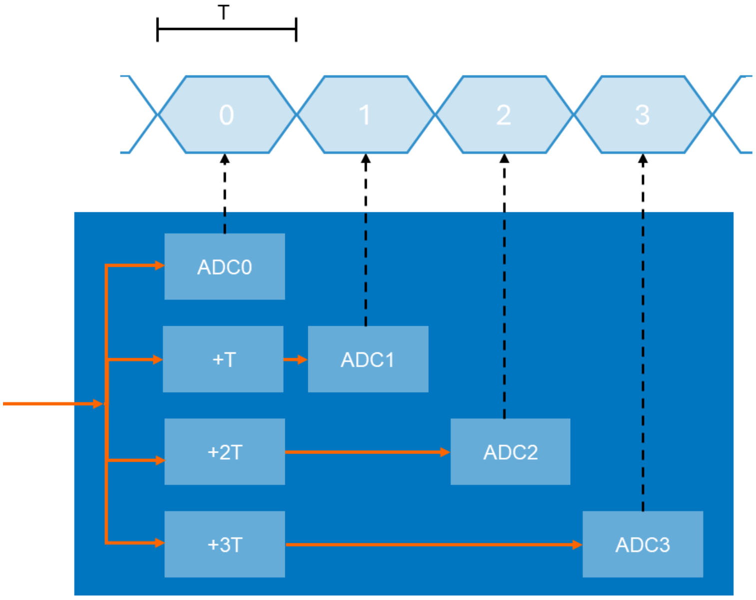

Figure 5: 1:4 Time-interleaved ADC sampling.

Clock generation and recovery also present significant hurdles. Phase-Locked Loops (PLLs) must generate reference clocks in the 50-100 GHz range with ultra-low jitter. Meanwhile, Clock and Data Recovery (CDR) circuits must reliably lock onto increasingly narrow eye openings, even in the presence of noise and jitter. Many systems now rely on DSP-based CDR techniques with adaptive equalization to maintain signal fidelity.

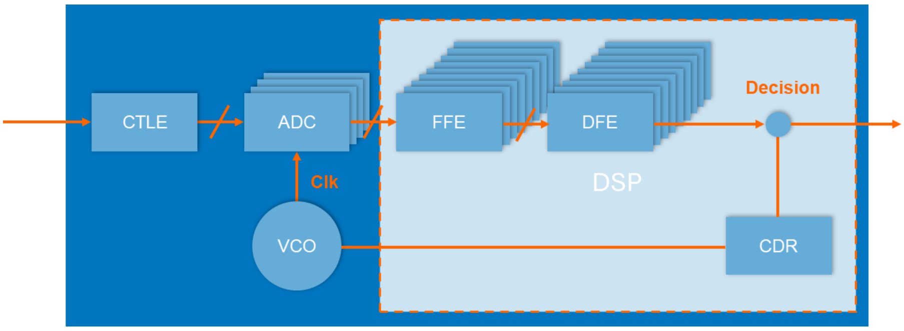

Figure 6: High level block diagram of 224Gbps PAM SerDes Rx.

As described earlier, equalization is essential for mitigating signal integrity issues. However, implementing digital equalizers is constrained by clock limitations and thus parallelization is also used for the FFE and DFEs in the receiver.

Finally, transmit drivers must deliver clean, high-speed signals with controlled swing and pre-/de-emphasis, all while maintaining power efficiency. This is where PHY developers must innovate to balance performance and energy consumption.

Overview of changes in 802.3dj

The IEEE802.3dj standard introduces significant enhancements to support 1.6Tbps Ethernet, with the most substantial changes occurring at Layer 1. Key updates include advances in the three main sub-layers – Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) layer and Physical Medium Dependent (PMD) layer.

PCS lanes

The PCS is responsible for data encoding and decoding, scrambling and descrambling, alignment marker insertion and removal, block and symbol redistribution, FEC encoding and decoding, and lane block synchronization and de-skew.

The 802.3dj standard introduces a significant change to the PCS for 1.6T Ethernet compared to previous high-speed Ethernet specifications. While the overall structure re-uses the two flows with multiple virtual lanes as introduced in 802.3df for 800 Gbps Ethernet, the number and speed of these lanes have changed, as illustrated in Figure 7. For 800G Ethernet based on 112G SerDes (e.g., 802.3df), there are 32 virtual lanes operating at 25 Gbps. In contrast, 1.6T Ethernet uses 16 lanes running at 100 Gbps.

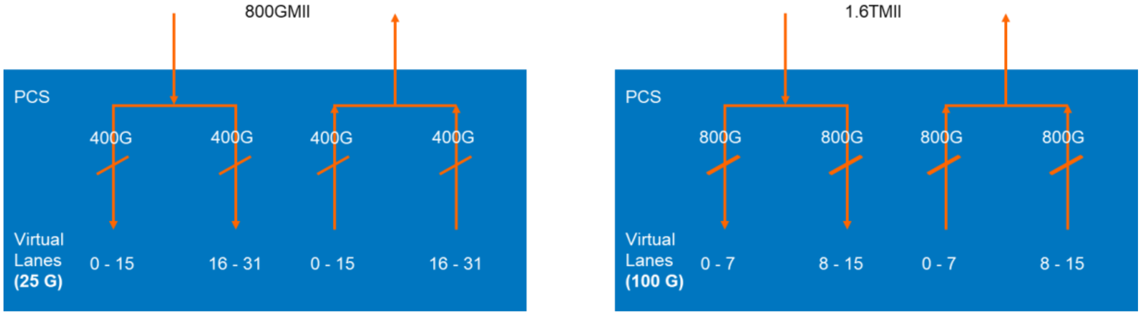

a) IEEE802.3df: 32 PCS lanes at 25 Gbps to support 800 Gbps Ethernet

b) IEEE802.3dj: 16 PCS lanes at 100 Gbps to support 1.6 Tbps Ethernet

Figure 7: PCS virtual lanes: a) 32 lanes of 25Gbps for 800Gbps Ethernet based on 112G SerDes and b) 16 lanes of 100Gbps for 1.6Tbps Ethernet based on 224G SerDes.

PSC FEC

The FEC in Ethernet is part of the PCS layer and, as mentioned above, the 224G standard uses the same RS-FEC as the 112G standard. For 800Gbps Ethernet, each group of 4 PCS lanes is bit-multiplexed into a 100Gbps PMA lane as illustrated in Figure 8. However, this scheme has the unfortunate effect of distributing burst errors occurring on a single lane across multiple symbol periods. Since the FEC can correct only up to 15 symbols per codeword, there is a high risk that it will not be able to correct all errors. At 200Gbps per lane this problem becomes even worse since more burst errors can be expected.

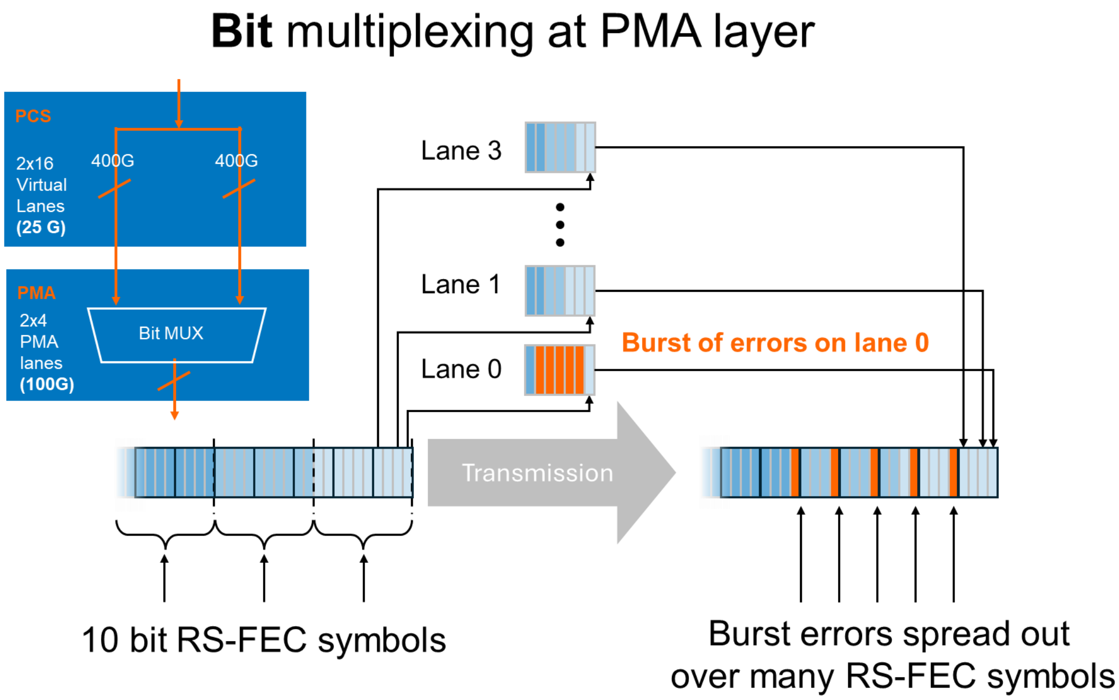

Figure 8: Bit multiplexing 10-bit FEC symbols spreads burst errors out over many symbol periods.

Therefore, symbol multiplexing has been chosen for 224G to improve the performance. As illustrated in Figure 9, this approach ensures that burst errors on a single lane are contained within as few symbol periods as possible, increasing the likelihood that the FEC will correct all errors.

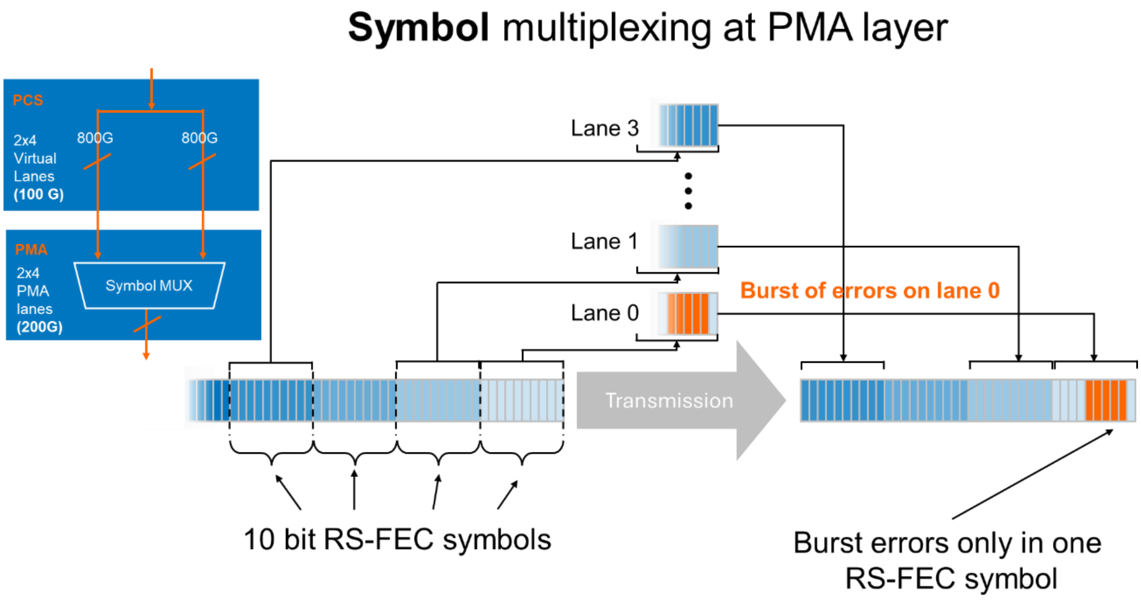

Figure 9: Symbol multiplexing 10-bit FEC symbols contains burst errors into a few symbol periods.

Concatenated FEC

Another new feature introduced with 802.3dj is the use of concatenated FECs for optical interconnects (BASE-FR). The just described PSC FEC (Outer FEC) is supplemented by a Hamming (128, 120) inner code (Inner FEC) at the optical PMD sublayer as illustrated on Figure 10. The Inner FEC will correct errors at the optical PMD layer while the Outer FEC will correct the remaining errors occurring on the electrical interface between the PMD and PCS layer.

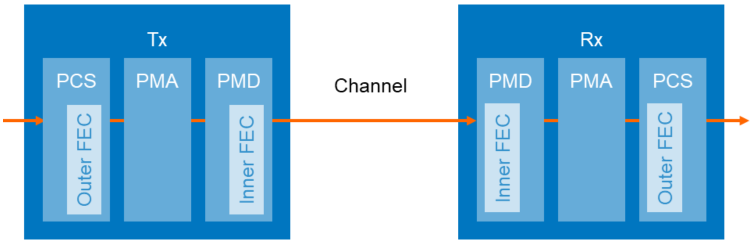

Figure 10: Concatenated FECs with Inner and Outer FEC.

Transceivers and cables

There are several transceiver module form factor options for 1.6T Ethernet. The OSFP (Octal Small Form Factor Pluggable) is well suited as a transceiver module for 1.6T Ethernet due to its inherent support for eight high-speed lanes. The OSFP Multi-Source Agreement (MSA) has already defined an OSFP1600 module supporting eight lanes at 224Gbps each. While the standard OSFP module includes an integrated heat sink, a version without a built-in heat sink also exists, the so-called SFP-RHS, where RHS stands for Riding Heat Sink. In the OSFP-RHS design, the heat sink is mounted on the cage of the network port and makes thermal contact with OSFP-RHS module when it is latched into place.

Similarly, the Double-Density (DD) version of the Quad Small Form Factor Pluggable (QSFP-DD) has also been specified to support eight lanes at 224Gbps.

A third option, also defined by the OSFP MSA, is the OSFP-XD (eXtra Dense), which integrates 16 lanes at 112Gbps within a single module, thereby achieving 1.6T Ethernet. This module targets the broad industry interest in 1.6Tbps interfaces based on 112 Gbps-per-lane technology and, with 224Gbps per lane the OSFP-XD can enable 3.2Tbps connectivity. Table below lists an overview of transceiver options for 1.6T Ethernet.

Transceiver options for 1.6T Ethernet

	OSFP	OSFP-RHS	QSFP-DD	OSFP-XD
Number of lanes	8	8	8	16
Per lane speed	224G	224G	224G	112G

In-rack connections in data centers are typically made by electrical cables rather than optical fiber cables due to cost and power consumption concerns. At 224Gbps the reach is expected to be 1 – 2 meters for passive cables (Direct Attach Cables – DACs). This reach can be extended by using Active Electrical Cables (AECs) that include equalizers and optionally also drivers and re-timers in the cable.

High-speed connectivity between racks and spine switches in data centers typically requires optical cables. To address the power consumption issues of optical interfaces, Linear Pluggable Optics (LPO) modules have been introduced by several vendors. These modules omit retiming, equalization, and error correction, instead relying on similar functionality built into the Switch or NIC ASIC, thereby reducing power consumption. LPOs typically fit into the OSFP and QSFP-DD format.

To further reduce the power consumption of optics, Co-Packaged Optics (CPO) is emerging as a promising solution. In CPO technology, the optics and the ASIC are integrated on the same substrate, which reduces losses, latency, and power consumption even further. However, since the ASIC and optics are now integrated, they do not fit into a traditional transceiver cage like an OSFP.

Implementations

Figure 11 illustrates a common implementation of a high-speed Ethernet switch network card with the switch ASIC connected via short Chip-to-Module (C2M) links to 224G PHY modules. The PHYs may include a 2:1 gearbox, which supports 112G toward the host switch ASIC and 224G toward the line interfaces. This gearbox enables switches with 112G ports to support 224G connectivity.

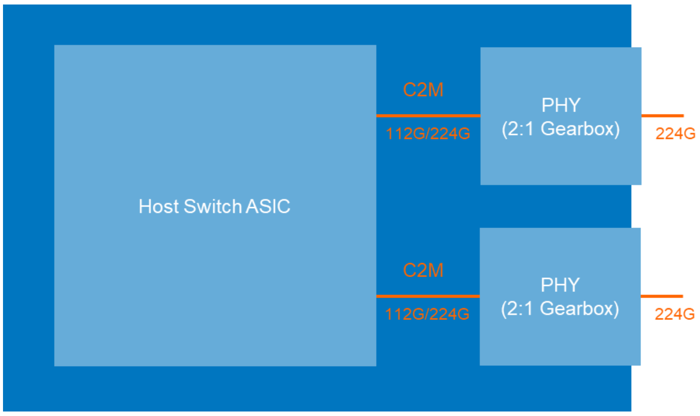

Figure 11: Typical implementation with host ASIC and PHY modules on a switch network card.

The use of an external PHY segments the Layer 1 sublayers between the host and PHY module, as shown in Figure 12. The PCS and PMA resides in the host whereas the PHY module contains its own PMA and the PMD. The interface between the two PMA sublayers is the Attachment Unit Interface (AUI).

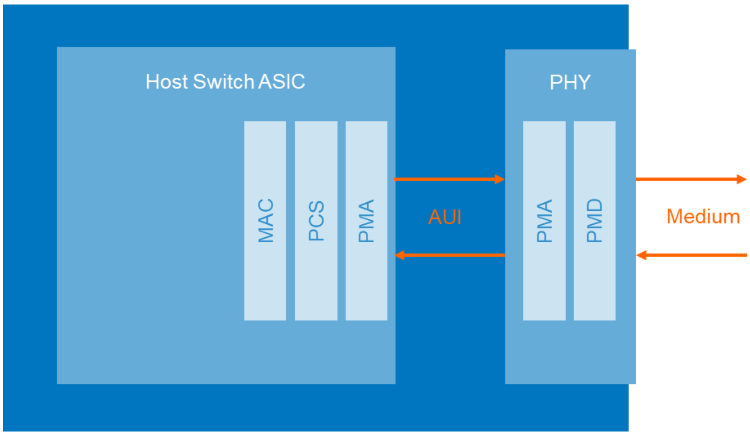

Figure 12: Segmentation of Layer 1 between host and PHY module.

CMIS-LT

At 112G / 224G speeds, even the short C2M links between the ASIC and the PHY modules require signal pre- and post-conditioning (equalization) to ensure adequate signal integrity. Due to variations in physical layout and temperature, there is rarely an equalizer setting that works in all situations. To address this, the OIF is defining CMIS-LT for Very Short Reach (VSR) and Extremely Short Reach (XRS) applications. CMIS-LT allows two modules on the same PCB to tune each other's transmit equalizer for optimal signal integrity. CMIS-LT uses out-of-band messaging based on the Common Management Interface Specification (CMIS).

CMIS-LT differs from the in-band 802.3 Ethernet Link Training (PMD control) which is intended for backplane and cable connections within a rack. The OIF CMIS-LT is a general-purpose solution and is not limited to Ethernet connectivity, although it uses a similar messaging structure as the IEEE 802.3 LT.

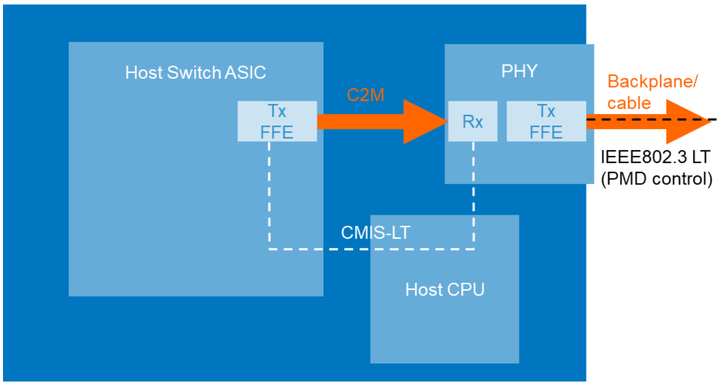

Figure 13: Illustration of CMIS-LT on the C2M connection and IEEE02.3 LT on the backplan/cable.

CMIS-LT is also intended for die-to-optical-module connections in applications such as CPOs where the optics is integrated with the switch ASIC on the same substrate using XSR or XSR+ connections.

Inter Sublayer communication

As described, the link between the ports on two host switches is segmented into several high-speed links which are trained segment-by segment. For the backplane or cable segment, both Auto

Negotiation (AN) and Link Training must complete successfully before the link is ready. For the AUI links, the CMIS-LT must also complete successfully. All segments must be ready for transmitting data before the end-to-end data transmission can begin, which necessitates an inter-sublayer service interface to support basic status communication between the sublayers.

Typical test scenarios

As semiconductor devices and network equipment are developed for 1.6Tbps Ethernet, there is naturally a need to test these new products. Furthermore, as this network gear is deployed in data centers and service provider networks, operators must also carry out a range of performance, interoperability, and maintenance tests. In the following we will show some common examples of test cases at the various levels of the 1.6Tbps Ethernet ecosystem.

Silicon chip testing

Testing of a new semiconductor device for terabit Ethernet starts at Layer 1. Once the basic bring-up tests are passed it is time for signal integrity tests. The silicon is tested on a test PCB which not necessarily has a full transceiver mounted but rather use SMPX-connectors.

Using PRBS, it is possible to verify the overall BER performance of a device. However, to verify the functionality of each sub-module, like FEC, PCS lane distribution and alignment, it necessary to test these individually. This can be done by deliberately inserting specific errors, such as FEC symbol errors, and check that the device handles these situations correctly.

Figure 14 how an example where the Z1600 Edun is used with an OSFP-to-SMPX connector for testing of a high-speed chip.

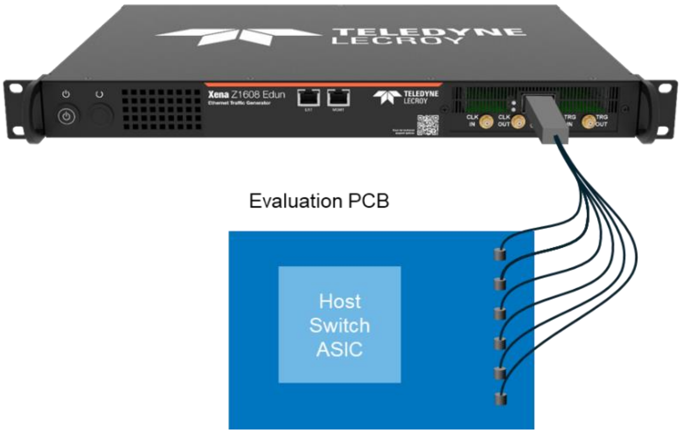

Figure 14: Functional testing of switch ASIC with Teledyne LeCroy Z1600 Edun and OSFP-to-SMPX connector.

When the chip is mounted with a PHY it becomes important to be able to test and tune the C2M equalizers and the CMIS-LT.

On top of the Layer 1 testing, it is off course crucial to verify the switch’s Layer 2 performance such as frame forwarding, MAC address learning, latency, jitter, VLANs, and QoS. The IEC RFC2544 and RFC2889 are good starting points for benchmarking and are available as test suites for all Teledyne LeCroy traffic generators.

Active cable tests

As mentioned earlier, active cables can extend the reach of passive DACs. However, active cables include equalizers which must be tuned to optimize the BER at various speeds and lengths. Manual equalization tuning can be slow and tedious, involving hours of testing to find the right tap setting.

At Teledyne LeCroy we have worked intensively to develop automated test sequences that via the CMIS interface run equalizer tests of active cables very efficiently and in a short time frame. Figure 15 shows a typical example of such a test using two Z1600 Edun traffic generators. The test is typically done in one direction at a time using one traffic generator to adjust the equalizer and generate PRBS and the other traffic generator to calculate the BER.

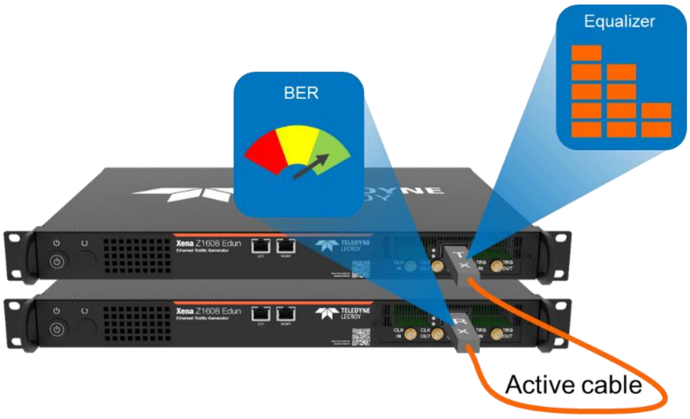

Figure 15: Testing of active cable using two Teledyne LeCroy z1600 Edun traffic generators.

Network equipment testing

Just as rigorous testing is essential for silicon manufacturers, verifying the Layer 1 - 3 functionality of switches and NICs is critical for Network Equipment Manufacturers (NEMs) developing 1.6Tbps-capable systems. Realistic traffic patterns and stress conditions must be emulated at wire speed without building massive test networks.

Traffic generators and analyzers like the Z1600 Edun are indispensable tools for validating functionality, performance and multi-vendor interoperability with minimum network building. A common test is the so-called snake test, where all ports on a switch is tested using two traffic generators as shown in Figure 16.

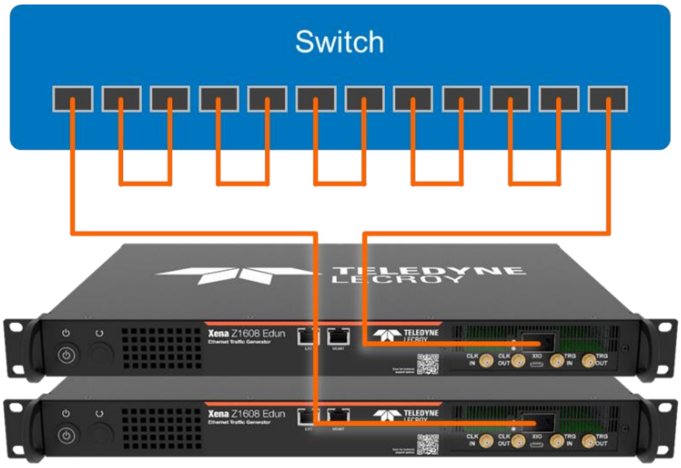

Figure 16: Snake test of all switch ports using two traffic generators.

Data center testing

All the previously discussed components come together in large data centers, where the increase in per-port speed plays a key role in boosting overall network capacity and efficiency. As operators gradually upgrade their infrastructure to 1.6Tbps Ethernet, interoperability with existing equipment becomes essential. Thorough testing before deployment allows potential performance issues to be identified and resolved in a controlled environment, reducing the risk of disruptions in live networks. Troubleshooting can be complex, with possible root causes including faulty cables, transceivers, network elements, or configuration errors.

Xena traffic generators are designed to address these challenges. They enable comprehensive testing of high-speed links, from verifying Layer 1 signal integrity and link training to validating Layer 2/3 protocol functionality. This allows data center operators to assess interoperability, ensure standards compliance, and fine-tune system performance—before going live. Whether for development, qualification, or troubleshooting, Teledyne LeCroy's Xena solutions provide the visibility and precision needed to confidently deploy next-generation Ethernet technologies.

Teledyne LeCroy 1.6TE test equipment

Teledyne LeCroy offers a range of Ethernet Traffic Generator and Analyzers covering all Ethernet speeds from 10Mbps to 1.6Tbps. All modules are conveniently managed through the same XenaManager 3 (XM3) software. XM3 is the new version of XenaManager loaded with improvements to the user interface to make testing of high-speed Ethernet even easier.

Our XOA open-source Python environment lets you integrate Ethernet testing with any of our traffic generators and analyzers into your own test environment. This makes it easy to migrate your existing test cases to new, higher-speed test modules.

Our Teledyne LeCroy Xena Z1608 Edun module supports 100Gbps, 200Gbps, 400Gbps, 800Gbps and 1.6Tbps using either 112G or 224G SerDes and a single OSFP port.

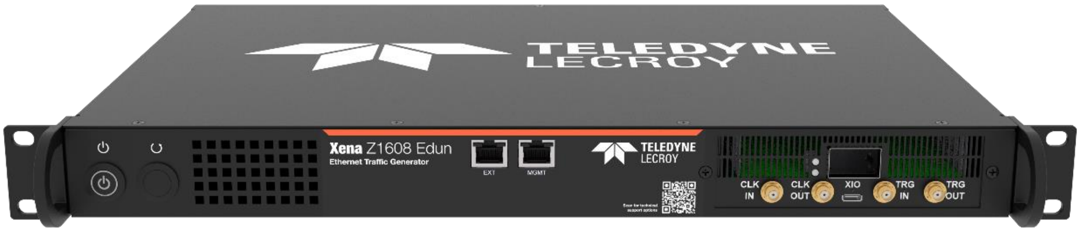

Figure 17: Xena Z1608 Edun 1.6T Ethernet Traffic Generator

The Z1608 Edun uses the same software as all our other traffic generators, ensuring a consistent user experience regardless of hardware configuration. XM3 is an easy-to-use tool for configuring ports, setting up traffic flows and analyzing traffic properties. Furthermore, the XOA open-source Python environment integrates with your own testing environment for automation and scripting.

The high-level feature set of Z1600 Edun is as follows:

- 5-speeds: 1.6TE, 800GE, 400GE, 200GE & 100GE
- 1 x OSFP cage
- Supports 224G SerDes (PAM4 224G)
- Supports 112G SerDes (PAM4 112G)
- Test with optics and DACs
- Extensive L1, L2 and L3 test features

For more information, please visit:

224G SerDes **Z1608 Edun** **Terabit Ethernet**