



OIF OPTICAL
INTERNETWORKING
FORUM

**White Paper:
Multi-Vendor Interoperability
Testing of CFP2 and QSFP28
with CEI-28G-VSR Interface
During OFC 2013 Exhibition**

March 19, 2013

White Paper created by the
Optical Internetworking Forum (OIF)
www.oiforum.com

Editors:

Ali Ghiasi, Mike Li, and Tom Palkert

Interoperability Chair:

Edward Frlan

Demo Participants and Contributors:

Agilent Technologies: Bob Hasenick and Joachim Vobis

Altera: Lux Joshi and Mike Li

Broadcom: Hong Aha, David Chung and Ali Ghiasi

Finisar: Chris Cole and Craig Schulz

Fujitsu Optical Components: Hidekazu Takeda

Inphi: Subramanian Seshadri and Richard Ward

Luxtera: Mike Musciano and Tom Palkert

Molex: Mehrdad Saberi, Scott Sommers, and Joe Dambach

Oclaro: Kiyo Hiramoto and Jon Anderson

Semtech: Ian Kalas and Edward Frlan

Tektronix: Christopher Loberg

Xilinx: Martin Gilpatric, Richard Cosgrave and Mark Gustlin

Yamaichi: Takeshi Nishimura

1.0 Executive Summary

The Optical Interworking Forum (OIF) is developing a set of Interoperability Agreements focused on new common electrical interfaces (CEI) applicable to higher speed optical systems requiring interconnect baud rates of 19.60 Gbaud to 28.10 Gbaud using NRZ coding. The OIF membership, consisting of semiconductor, connector, optical module suppliers, system suppliers, test equipment providers, allows a unique perspective for developing industry requirements, and a comprehensive understanding of the technology trade-offs necessary to enable the development and support for these requirements.

A multi-vendor interoperability event to demonstrate applications of CEI-28G-VSR interfaces for two popular module form factors, namely CFP2 and QSFP28, was successfully demonstrated during OFC 2013 Exhibition by several members of the OIF Physical and Link Layer (PLL) Working Group. This new demonstration follows the successful OIF interoperability events that took place at the 2012 OFC and ECOC 2012 Exhibitions. The OFC 2013 PLL interoperability builds on the previous interoperability events and consists of three individual demonstrations with 13 participating vendors. The three demonstrations are:

- Demo 1: Altera FPGA with 100 GbE MAC is driving CEI-28G-VSR host card with Finisar CFP2 module plugged into it and at the far side Oclaro CFP2 module is plugged into CEI-28G-VSR host card with Inphi 100 GbE gearbox looping the traffic back. Yamaichi CFP2 connectors are used at CFP2 module mating interface.
- Demo 2: Xilinx FPGA with 100 GbE MAC is driving CEI-28G-VSR host card with Fujitsu Optical Component CFP2 module plugged into it and at the far side Finisar CFP2 module is plugged into CEI-28G-VSR host card with Broadcom 100 GbE gearbox looping the traffic back. Yamaichi CFP2 connectors are used at CFP2 module mating interface.
- Demo 3: Inphi Gearbox is driving CEI-28G-VSR channel and followed by Semtech CDR that drives one of the Molex QSFP28 Active Optical Cable (AOC) based on Luxtera Silicon Photonics, on the far side of the AOC the signal is retimed by the Semtech CDR then the signal passes through the CEI-28G-VSR channel and at the far side it is terminated by the Tektronix BERTScope BSA286C. Molex zQSFP+ connectors are used at QSFP28 connector module mating interface.

1.1 OIF PLL Multi-Vendor Interoperability Testing Objectives

Interoperability is one of the keys to the success of any standard or implementation agreement. In order to promote the acceptance and demonstrate the viability of CEI-28G-VSR, the OIF sponsored a private, closed door interoperability Plugfest in February 2013. This Plugfest proved that CEI-28G-VSR is widely supported by various companies across the industry with participants including five semiconductor manufacturers (Altera, Broadcom, Inphi, Semtech, and Xilinx), two connector vendors (Molex and Yamaichi), four optical module vendors (Finisar, Fujitsu Optical Component, Molex, and Oclaro), one optical component vendor (Luxtera) and two test equipment manufacturers (Agilent and Tektronix). Cooperation between semiconductor, connector, optical module, optical component and test equipment suppliers is crucial to enable implementation and integration of high-speed signaling by system and operator vendors. OFC 2013 will showcase working demonstrations of the CEI-28G-VSR electrical interface implemented on a CFP2 100Gbase-LR4 optical module and a QSFP28 AOC, using test scenarios similar to the private interoperability tests.

2.0 CEI-28G-VSR Background Information and Applications

The communications and networking industries' data rates have been increasing along with the demands for higher levels of traffic aggregation. Additionally, the industry desires interoperable interfaces that support these next generation data rates. The OIF has been at the forefront of this demand, leading the industry with the development of CEI implementation agreements (IA) that support up to 28.10 GBd in chip-to-module interfacing applications. This development work is unique and timely in enabling the industry to re-use conventional chassis and line-cards as they develop equipment that is able to meet the evolving and challenging bandwidth demands of the communications industry. The CEI-28G-VSR Clause is in final development phase and implementation agreement is expected to be published by the mid of 2013.

The use of pluggable optical transceivers and direct attach copper cables is a common practice in equipment developed for the communications market. Developing interoperable pluggable solutions that keeps up with the demanding bandwidth needs of industry is critical to enabling next generation equipment that supports the communications service providers.

The CEI-28G-VSR application reference diagram is shown in Figure 1. It consists of 100 Ω differential PCB traces, vias, one connector and AC coupling capacitors. The CEI-28G-VSR IA is intended to be used for "Very Short Reach" channels, with length up to 150 mm, and loss up to 10.0 dB loss at Nyquist rate. The module that is being initially targeted for the next generation of optical modules are retimed interfaces operating at 25.78 – 28.10 GBd. The overall CEI-

28G-VSR link needs to operate with a bit error ratio (BER) of $1e-15$ or less. VSR interface supports multiple lanes and is hot pluggable.




Fig. 1: CEI-28G-VSR Reference Model

2.1 CEI-28G-VSR Specification Channel Requirement

A VSR link supports a maximum loss of 10 dB at Nyquist for operation from 19.60 to 28.10 GBd with NRZ modulation. The link loss budget is divided into a maximum host PCB loss of 7.3 dB, a connector loss of 1.2 dB and a module PCB loss of 1.5 dB, Figure 2 shows a diagram of the loss breakdown. TP1a and TP4a are respectively the host electrical output and input measurement points measured with Host Compliance Board (HCB) with PCB loss of 2.0 dB. TP1 and TP4 are respectively the module electrical input and output measurement points measured with module compliance board (MCB) with PCB loss of 1.25 dB.

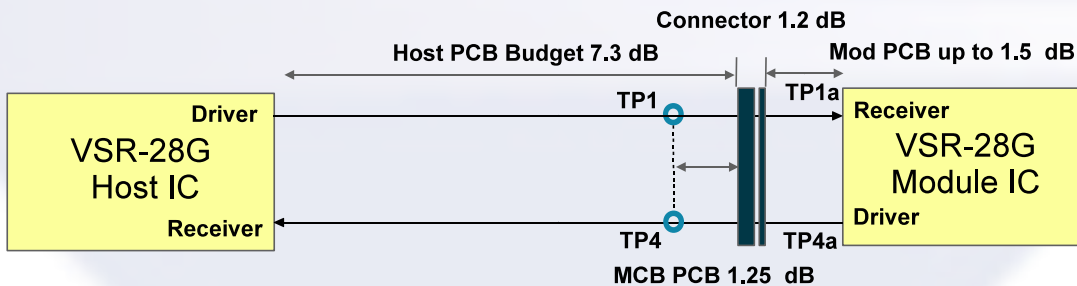


Fig. 2: CEI-28G-VSR Reference Model

Figure 3, shows the CEI 28G-VSR channel response scaled for operation at 28 GBd with loss of 10 dB at Nyquist frequency of 14 GHz.




Fig. 3: CEI-28G-VSR Channel Response

3.0 Demo 1: 100 GbE CFP2 with Live Ethernet Traffic

Participating companies in Demo 1 were Altera provided FPGA with 100 GbE MAC and Inphi provided 100 GbE 10:4 Gearbox. Finisar and Oclaro provided 100Gbase-LR4 CFP2 modules. Yamaichi provided CFP2 connector and cage.

3.1 Demo 1 Link Component Overview

Altera Stratix® V GT FPGAs are designed and fabricated on a 28 nm CMOS high-performance, low-power process, integrated with 4x28.1 GBd, and 32x12.5 GBd transceivers, up to 4x72 bit DDR3 memory interfaces. It has more than quarter million adaptive logic modules (ALMs), up to 622,000 equivalent logic elements (LEs), and up to 50 Mb on chip SRAM memory. Stratix® V GT FPGA transceivers are designed to interface with CFP2 modules meeting/exceeding the CEI-28G-VSR specification requirements. Finally Stratix® V GT FPGA MAC, PCS, and PMA are optimized to support high bandwidth applications such as 40G/100G/400G Ethernet and OTN, with low latency and low power.

Finisar CFP2 100GBase-LR4 modules use a quad 28G DFB DML PIC (Photonic Integrated Circuit) TOSA and quad 28G PIN PIC ROSA, with four 2x28G Semtech CDR. This results in 7 W total module power dissipation, and a significantly reduced cost structure compared to Gen1 28G EML TOSA technology. The CFP2 LR4 modules are now carrying live traffic in field trials and will be in full production this year.

Oclaro's first CFP2 design is roughly 50% the size CFP. The smaller size of the CFP2 transceiver is achieved by reducing the number of electrical lanes from 10 to 4 by applying a 25 GBd electrical interface per channel instead of 10 GBd and hence removing the Gearbox from the module. In order to achieve the smaller size, more integrated in-house optical components are used based on Oclaro's leading-edge optical device technologies. Oclaro's CFP2 transceiver is fully interoperable with existing CFP transceivers based on 100Gbase-LR4, which allows system OEMs to quickly and cost-effectively transition to CFP2 designs as market requirements evolve.

The Inphi IN112510-LD is a single-chip, low-power tri-rate CMOS PHY for 10:4 gearbox/10:10 retimer applications in 10/40G/100G Ethernet and OTU4 line cards and optical modules. Optimized for low-latency operation, it delivers high system performance and ultra-low jitter for demanding applications. It is designed to significantly exceed CEI-28G-VSR requirements and allows real-time monitoring of link margin and protocol checking as well as dynamic high resolution eye scan diagnostics.

Yamaichi CFP2 connector CN121 series complies with CFP-MSA specification and achieved better connector signal loss performance than CEI-28G-VSR requirement. Dual slot and single slot mechanical parts are already available, and it's also capable belly to belly mount as well.

3.2 Demo 1 Description

This demonstration consists of the following blocks: Altera FPGA with 100 GbE MAC, Finisar 100Gbase-LR4 CFP2 module, Oclaro 100GBase-LR4 CFP2 module, and Inphi 10:4 Gearbox. The link operated at 25.78 GBd full duplex error free over all 4 channels with PRBS31 as well as live 100 GbE Ethernet traffic. Demo 1 CFP2 100Gbase-LR4 modules support reach up to 10 km on duplex SMF.

Altera's FPGA host implementation was constructed of approximately 140 mm long 100 Ω microstrip traces with two vias constructed on Megtron-6 material that connects to a Yamaichi CFP2 connector. Finisar's CFP2 optical module was plugged into the Altera FPGA host board and a duplex fiber connects to the remote CFP2 module from Oclaro.

Inphi gearbox host implementation was constructed of 50 mm long fine stripline traces with two blind vias constructed on Megtron 6 material that connects to Yamaichi CFP2 connector. Oclaro CFP2 was plugged into Inphi gearbox host card then a duplex fiber connects to the remote CFP2 module from Finisar.

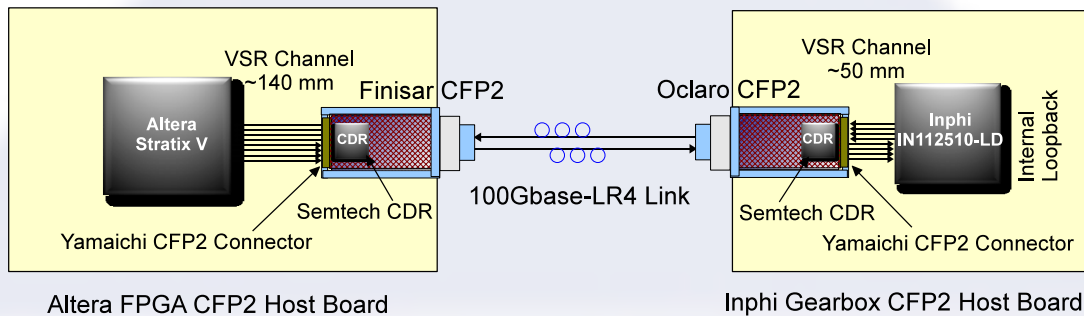


Fig. 4: Demo No. 1 Block Diagram

4.0 Demo 2: 100 GbE CFP2 with Live Ethernet Traffic

Participating companies in Demo 2 were Broadcom providing 100 GbE 10:4/10:10 Gearbox and Xilinx providing an FPGA with 100 GbE MAC. Finisar and Fujitsu Optical Component provided 100Gbase-LR4 CFP2 modules. Yamaichi provided CFP2 connector, cage, and CFP2 HCB (Host Compliant Board).

4.1 Link Component Overview

The Broadcom BCM84793 is a low-power single-chip dual mode 10:4 gearbox/10:10 retimer in 40 nm CMOS. BCM84793 supports 10 GbE, 40 GbE, 100 GbE, and OUT4 line card applications. The dual mode Gearbox/Retimer enables next generation 100 GbE applications as well as high-density 10/40 GbE applications. The device is optimized for low-latency operation and flexibility. The device delivers significantly lower jitter generation, improved jitter tolerance, and improved sensitivity compared to the CEI-28G-VSR specification. BCM84793 has advanced diagnostics, pattern generation and checking, and with ability for real time link margin and eye scan diagnostics.

Finisar CFP2 100GBase-LR4 modules use a quad 28G DFB DML PIC (Photonic Integrated Circuit) TOSA and quad 28G PIN PIC ROSA, with four 2x28G Semtech CDR. This results in 7 W total module power dissipation, and a significantly reduced cost structure compared to Gen1 28G EML TOSA technology. The CFP2 LR4 modules are now carrying live traffic in field trials and will be in full production this year.

Fujitsu Optical Components CFP2 LR4 transceiver is designed to target OIF 28G-VSR specification, which is the high speed electrical interface to support next generation 100G technology. The introduction of this new interface, together with the high density integration technology of optical and electrical parts, enables the compact sized (about half width of the current CFP) and low power consumption transceivers. This CFP2 module, which includes an integrated CDR, connects to the Xilinx gearbox via VSR channel boards directly.

The Xilinx 7 Series HT family of FPGAs utilize heterogeneous stacked 3D IC device integrating both FPGA slices and high performance 25-28.05 GBd transceivers via a silicon interposer. The FPGA slices, implemented on a 28nm high-performance/low-power (HPL) process, contains 580,480 logic cells, 1,680 DSP slices, 33,840Kbits of block RAM (BRAM), 48x 13.1 GBd GTH transceivers, and 8x 28.05 GBd GTZ transceivers. The GTZ transceivers are designed to exceed the OIF-28G-VSR electrical specification and are targeted specifically to interoperate with CFP2 optical modules and are arranged to directly connect to CFP2 connectors. For the demonstration, the device has been configured with multiple 100GbE MAC IP where the device is simultaneously transmitting and receiving packetized IP data and checking the IP data for errors.

Yamaichi CFP2 connector CN121 series complies with CFP-MSA specification and achieved better connector signal loss performance than CEI-28G-VSR requirement. Dual slot and single slot mechanical parts are already available, and it's also capable of belly to belly mount.

4.2 Demo 2 Test Equipment

Test equipment supplied to the PLL Interop demonstration by Agilent Technologies included the following:

- 86100D DCA-X Oscilloscope: Modular platform that accommodates up to 4 measurement modules and 16 measurement channels. Includes Option 200 Enhanced Jitter Analysis and Option 201 Advanced Waveform Analysis software.
- 86108B Wide Bandwidth Module: Has integrated dual receiver to 50 GHz, clock recovery to 32 Gb/s and precision timebase with typical intrinsic jitter of less than 50 fs.
- N1012A OIF CEI 3.0 Compliance and Debug Application: Measures the ~125 parameters for CEI 3.0 and the draft VSR implementation agreements.

4.3 Demo 2 Description

This demonstration system consists of the following blocks: Broadcom host gearbox IC, Finisar 100Gbase-LR4 CFP2 module, Fujitsu Optical Component 100GBase-LR4 CFP2 module, and Xilinx FPGA with 100 GbE MAC. The link operated at 25.78 GBd full duplex error free over all 4 channels with PRBS31 as well as live 100 GbE Ethernet traffic. Demo 2 CFP2 100Gbase-LR4 modules support reach up to 10 km on duplex SMF.

Broadcom gearbox host implementation was constructed of about 125 mm long fine stripline traces with two back drilled vias constructed on Megtron 6 material that connects to Yamaichi CFP2 connector. Finisar CFP2 was plugged into Broadcom gearbox host card then a duplex fiber connects to the remote CFP2 module (Fujitsu Optical Component).

Xilinx FPGA host implementation with dual CFP2 ports was constructed of about 50 mm long fine stripline traces with two vias constructed on Megtron 6 material that connects to Yamaichi far side CFP2 connector. Fujitsu Optical Component CFP2 was plugged into Xilinx FPGA host then a duplex fiber connects to the far side CFP2 module (Finisar).

Xilinx 2nd host ports is used to observe VSR transmitter compliance at TP1a. Yamaichi CFP2 HCB was inserted into the 2nd port to measure TP1a on Agilent Digital Sampling Scope with built-in CDR and VSR reference software CTLE for post equalized eye measurement and compliance to CEI-28G-VSR specification.

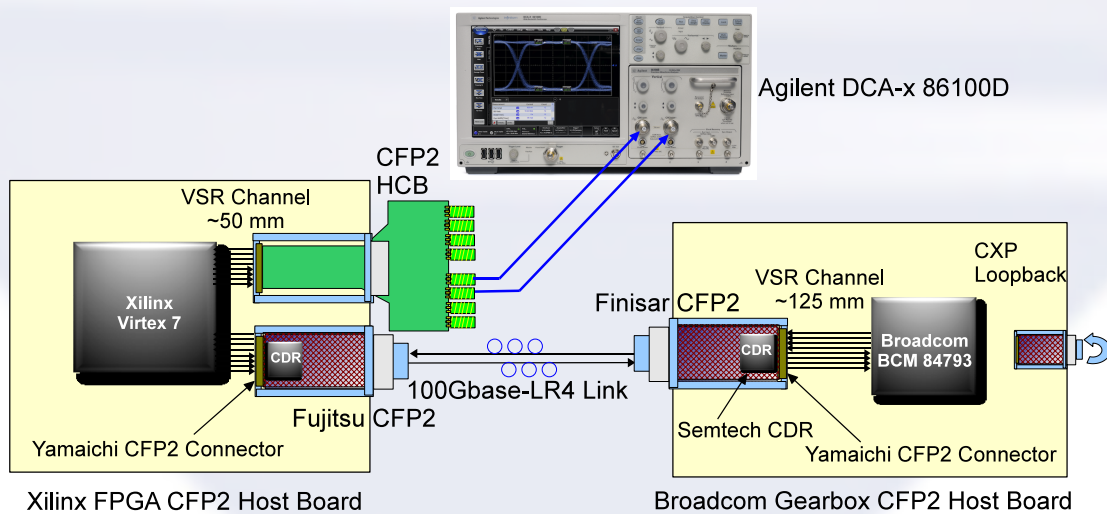


Fig. 5: Demo No. 2 Block Diagram

5.0 Demo 3: CEI-28G-VSR QSFP28 Application

Participating companies in Demo 3 were Inphi providing 10:4 gearbox, Semtech providing Quad CDRs, Molex and Luxtera proving the QSFP28 AOC, and Tektronix providing digital sampling scope and BERT.

5.1 Link Component Overview

The Inphi IN112510-LD is a single-chip, low-power tri-rate CMOS PHY for 10:4 gearbox/10:10 retimer applications in 10/40G/100G Ethernet and OTU4 line cards and optical modules. Optimized for low-latency operation, it delivers high system performance and ultra-low jitter for demanding applications. It is designed to significantly exceed CEI-28G-VSR requirements and allows real-time monitoring of link margin and protocol checking as well as dynamic high resolution eye scan diagnostics.

Luxtera's single chip Silicon Photonics optical transceiver includes four integrated transmit and receive channels operating from 1 to 28.1 Gbps. The transmit optical modulators are powered from a single laser source delivering superior optical performance when mated with standard single mode ribbon fiber cables. Luxtera's Silicon Photonics technology coupled with single mode fiber permits greater reach compared to multimode fiber solutions. The built-in optical receivers retrieve the optical signal and produce electrical signaling for downstream processing. Luxtera's optical transceiver is compatible with CEI-28G-VSR applications as illustrated in the demo 3. The device is targeted for 100 GbE, OTN and InfiniBand. Utilizing Silicon Photonics technology, the Luxtera optical transceiver IC enables low latency, low bit error rate, long reach, high reliability in a small footprint.

Molex zQSFP+™¹ connectors compatible with QSFP28 are used on the MCB boards. The zQSFP+ system from Molex supports next-generation CEI-28G-VSR, 100 Gbps Ethernet and InfiniBand Enhanced Data Rate (EDR) applications with excellent cooling, improved signal integrity (SI), superior electro magnetic interference (EMI) protection and low power consumption. Luxtera's Silicon Photonics (SIP) technology is heart of the Molex PSM4 (parallel single mode quad) AOC, which can transmit 100 GbE signal at up to 4 km reach and enabling compact QSFP28 form factor.

¹ zQSFP is trade mark of Molex Inc and is compatible with QSFP28 as defined by SFF-8661.

Semtech's GN2425/26 are low-power retimers, leveraging a generation of innovation in 10G module CDRs, optimized for reference-free 25-28.1 GBd operation in next-generation optical modules and active cables. By resetting the jitter budgets within the module in both directions, Semtech's CDRs are optimized for challenging new applications such as 100GBASE-LR4/ER4 and OTU4 CFP2 optical modules. The GN2425/26 feature superior receive sensitivity, input jitter tolerance and output jitter, allowing for reliable signal recovery and clean, wide-open transmit eyes.

5.2 Demo 3 Test Equipment

Test equipment supplied to the PLL Interop demonstration by Tektronix includes the following:

- DSA8300 Sampling Oscilloscope: The DSA8300 Series is a modular oscilloscope system that allows for electrical and optical sampling modules. For the OIF Interop, Tektronix provides the 80E09B Electrical Module. When used with the DSA8300 Series, this new module provides sub-100 femtosecond intrinsic jitter, low vertical noise, >45GHz bandwidth performance to enable high fidelity measurements of OIF 28G-VSR system.
- BSA286C BERTScope generates and checks the data at the far end of the CEI-28G-VSR link. The BSA286C provides support for up to 28.6 GBd PRBS31 measurements of Clock and Data with the ability to do advanced analysis in addition to the BER checking. Analysis tools include precise error location analysis, BER Contour and jitter mapping capability.

5.3 Demo 3 Description

This demonstration system consists of the following blocks: Inphi 10:4 Gearbox generating and checking data traffic, a VSR channel that included QSFP28 connector bolted to the Inphi Gearbox, the Semtech dual CDRs were added to the channel to retime the data in both direction, a QSFP28 MCB bolted to Semtech CDRs, Molex/Luxtera AOC was then inserted into the near end MCB, on the far end Molex/Luxtera AOC is plugged into the 2nd QSFP28 MCB, Semtech dual CDRs bolted to the far end MCB to retime the data in both directions, a VSR channel was then bolted to the Semtech dual CDRs, then one lane of the VSR channel is terminated to Tektronix BERTScope for independent data generation and checking. Semtech CDR on the far end is providing the recovered clock to the Tektronix BERTScope.

The local and far side of the link each had a Semtech VSR host channel trace PCB (contributing approximately 9dB of insertion loss), the total loss with addition of the QSFP28 MCB is approximately 13dB.

Demo 3 operated error free on one lane full duplex with PRBS31. One of the Semtech CDR outputs that recovered the signal through the CEI-28G-VSR channel is displayed on the Tek DSA8300 sample Oscilloscope, then the scope applies S/W CTLE function per OIF CEI-28G-VSR specifications.




Fig. 6: Demo No. 3 Block Diagram

6.0 Conclusions

This interoperability demonstration successfully brought semiconductor, connector, optical module and test equipment vendors together to show OIF CEI-28G-VSR instantiation of CFP2 as well as QSFP28 AOC. OIF CEI-28G-VSR is widely accepted and now the IEEE 802.3bm CAUI-4 specification is being developed largely based on OIF CEI-28G-VSR specification. OIF CEI-28G-VSR is the key building block for next generation 100 GbE, OTN, 32 GFC, and IB EDR retimed optical modules.

This demonstration for the 1st time showed CFP2 host and CFP2 module implementations of the 100Gbase-LR4 (100 GbE). Building on the QSFP28 ECOC 2012 demo, a full duplex end to end QSFP28 AOC link was demonstrated for the 1st time.