OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.

OCP NIC 3.0 Card Design and Interoperability Testing

Samit Ashdhir, Facebook
Jon Lewis, Dell EMC
Rick Eads, Keysight

Agenda

- OCP 3.0 NIC Design
- Interoperability Testing
  - Thermal
  - Electrical
- PCIe Conformance Testing for OCP platforms
- Future Work

Methodology

- Extend the success of OCP 2.0 by clearly specifying characteristics that allow ease of use across multiple systems.
- Thermal:
  - Standardize thermal requirements testing and details system requirements for airflow.
- Form Factor:
  - Clearly specify design limits and tolerances.
  - Industry standard SFF-TA-1002 (Rev 1.1)
- Systems Management:
  - Defines the minimum set of commands necessary to integrate any OCP 3.0 card into a compatible system.

Mechanical comparison to other Form Factors

Usable Space:
SFF to HHHL = 20% less
SFF to OCP2 = equivalent
LFF to FHFL = 50% less
LFF to FHHL = 13% less

Key difference for cooling is the topside keepout:
PCIe = 14.47 mm
OCP3 = 11.5 mm

Key difference for component placement Bottom side Keepout:
PCIe = 2.67mm
OCP = ~1.75mm

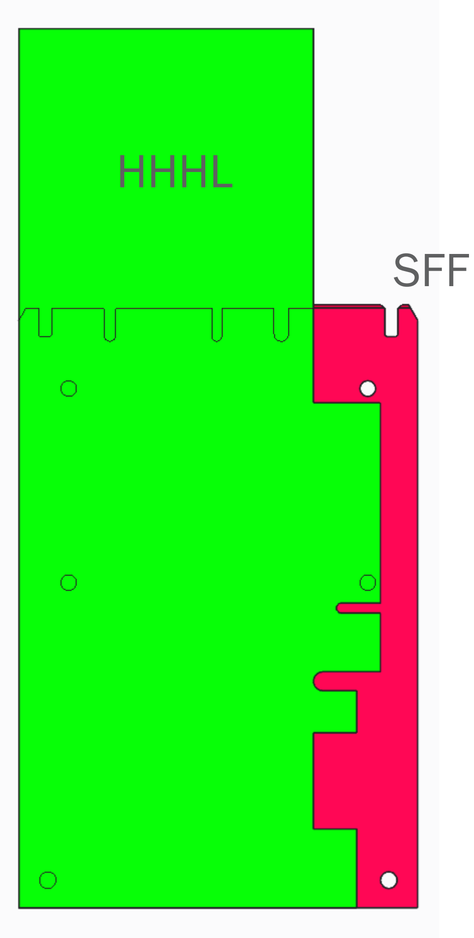

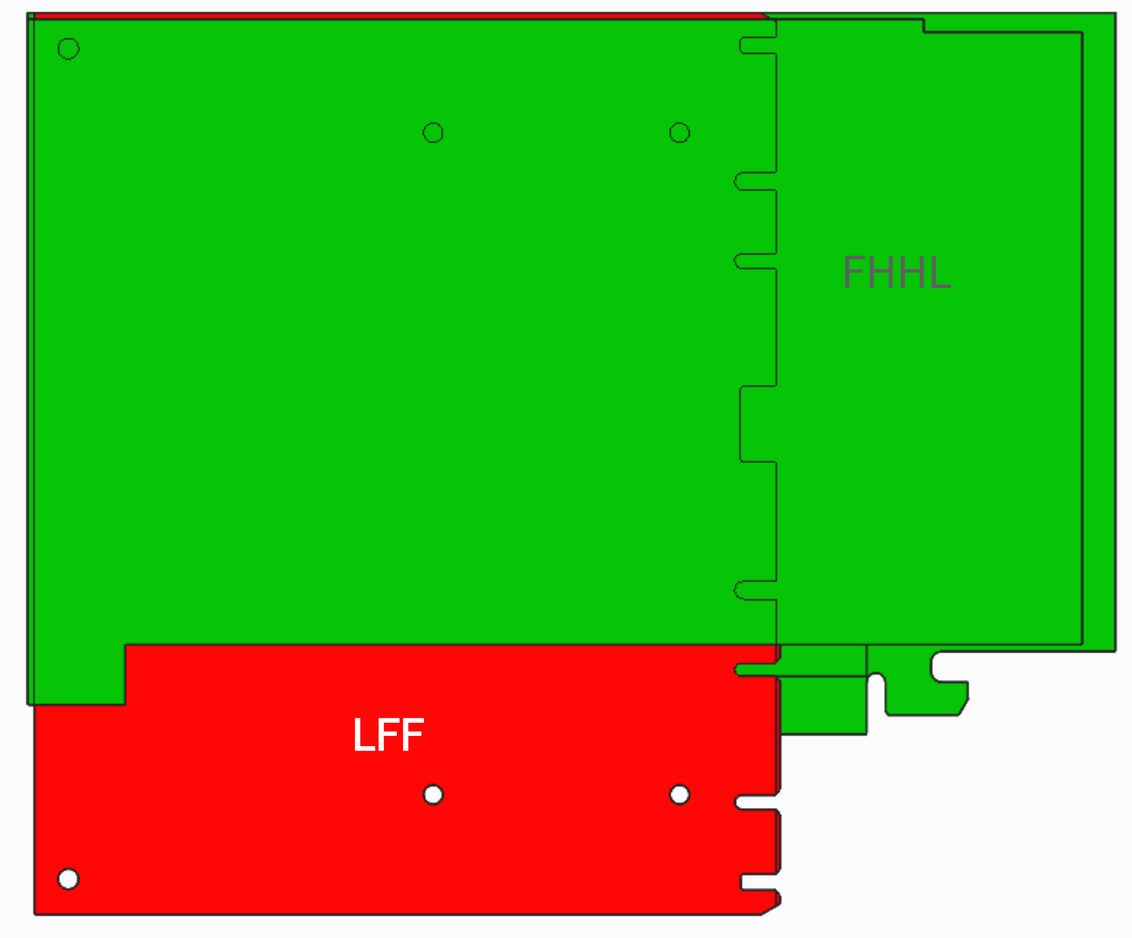

Systems Management Overview

- Power
  - AUX mode and MAIN mode power requirements contained in the FRU. Allows system to determine if sufficient power is available prior to moving from ID mode to AUX mode.
- Thermal
  - FRU contains worst-case airflow requirements. Allows system to determine if sufficient airflow is possible prior to enabling the adapter.
  - Real-time temperature monitoring of ASIC and Optics (if installed) allow the system to adjust airflow to properly cool the adapter.

Thermal comparison to PCIe and OCP2.0

- OCP 3.0 has significantly less cooling capacity due to topside keepout restrictions.
  - 11.5 mm for OCP 3.0
  - 14.67 mm for PCIe CEM
- Tiers have been added to PCIe SIG to allow for ease of design-in and a placeholder table is in the draft OCP 3.0 specification pending further analysis.

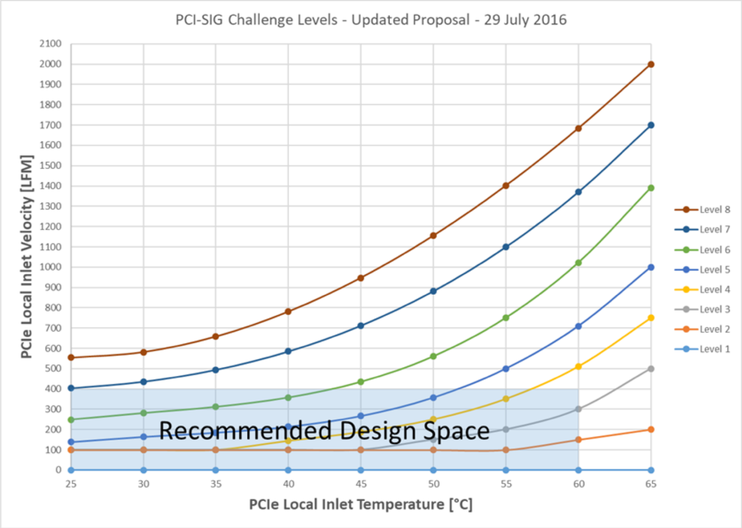

Table 59: Hot Aisle Card Cooling Tier Definitions (LFM)

	Target Operating Region				Server Airflow High Fan Speed		Non-Typical Server Airflow - Subject to System Capability					
OCP NIC 3.0 Local Inlet Temperature [°C]	Tier 1	Tier 2	Tier 3	Tier 4	Tier 5	Tier 6	Tier 7	Tier 8	Tier 9	Tier 10	Tier 11	Tier 12
5												
10												
15												
20												
25												
30												
35												
40												
45												
50												
55	50	100	150	200	250	300	350	400	450	500	750	1000
60												
65												

Work in Progress

SFF Hot Aisle Thermal restrictions

- Hot aisle simulations show 85C optics required with 400 LFM for QSFP
- SFP simulations only require 250 LFM

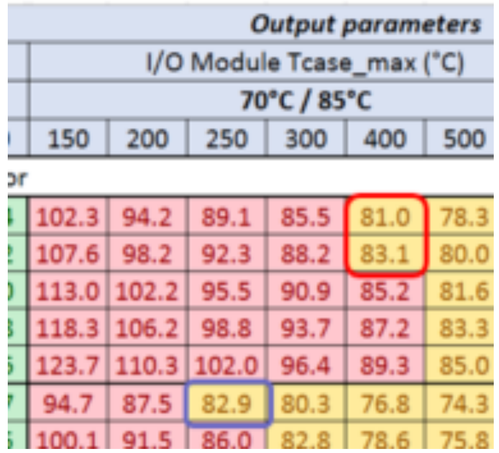

Output parameters						
I/O Module Tcase_max (°C)						
70°C / 85°C						
150	200	250	300	400	500	
or						
102.3	94.2	89.1	85.5	81.0	78.3	
107.6	98.2	92.3	88.2	83.1	80.0	
113.0	102.2	95.5	90.9	85.2	81.6	
118.3	106.2	98.8	93.7	87.2	83.3	
123.7	110.3	102.0	96.4	89.3	85.0	
94.7	87.5	82.9	80.3	76.8	74.3	
100.1	91.5	86.0	82.8	78.6	75.8	

Case #	Input parameters						Output pa																			
	I/O	ASIC 1 Power (W)	ASIC 1 Location	PCB Board Size	DRAM Location	Airflow Direction or I/O Location	ASIC Tcase_max (°C)							I/O Module Tcase_m							I/O Module approach Tair							
							105°C							70°C / 85°C													
							150	200	250	300	400	500	750	150	200	250	300	400	500	750	150	200	250	300	400	500	750	LFM
Hot Aisle - Straddle Mount 76x115mm Small Form Factor																											
A6 - with QSFP HS	2x QSFP	15	Top	14-small	N/A	Hot-aisle	84.0	79.9	77.3	75.6	73.3	71.8	69.4	102.3	94.2	89.1	85.5	81.0	78.3	74.4	74.1	69.4	66.5	64.6	62.3	60.9	59.0	
A7 - with QSFP HS	2x QSFP	20	Top	14-small	N/A	Hot-aisle	93.4	88.0	84.7	82.4	79.4	77.4	74.2	107.6	98.2	92.3	88.2	83.1	80.0	75.5	80.2	74.0	70.3	67.8	64.7	62.8	60.3	
A8 - with QSFP HS	2x QSFP	25	Top	14-small	N/A	Hot-aisle	102.7	96.1	92.0	89.1	85.4	82.9	79.0	113.0	102.2	95.5	90.9	85.2	81.6	76.7	86.4	78.7	74.0	70.9	67.1	64.7	61.6	
A9 - with QSFP HS	2x QSFP	30	Top	14-small	N/A	Hot-aisle	112.1	104.2	99.3	95.9	91.4	88.5	83.8	118.3	106.2	98.8	93.7	87.2	83.3	77.8	92.5	83.3	77.8	74.1	69.5	66.7	62.9	
A10 - with QSFP HS	2x QSFP	35	Top	14-small	N/A	Hot-aisle	121.5	112.3	106.6	102.7	97.5	94.0	88.6	123.7	110.3	102.0	96.4	89.3	85.0	78.9	98.6	88.0	81.6	77.2	71.9	68.6	64.2	
A6 - with QSFP HS	4x SFP	15	Top	14-small	N/A	Hot-aisle	83.7	79.9	77.4	75.7	73.5	72.0	69.7	94.7	87.5	82.9	80.3	76.8	74.3	70.7	77.7	72.8	69.8	67.7	65.1	63.4	61.0	
A7 - with QSFP HS	4x SFP	20	Top	14-small	N/A	Hot-aisle	93.1	88.0	84.7	82.5	79.5	77.5	74.5	100.1	91.5	86.0	82.8	78.6	75.8	71.8	84.5	78.2	74.3	71.6	68.1	66.0	62.9	
A8 - with QSFP HS	4x SFP	25	Top	14-small	N/A	Hot-aisle	102.4	96.0	92.1	89.3	85.6	83.1	79.4	105.5	95.5	89.7	85.6	81.0	77.7	72.9	91.3	83.5	78.8	75.5	71.3	68.6	64.7	
A9 - with QSFP HS	4x SFP	30	Top	14-small	N/A	Hot-aisle	111.7	104.2	99.4	96.0	91.7	88.7	84.3	110.9	99.2	92.9	88.4	83.3	79.5	74.3	98.1	88.9	83.3	79.4	74.5	71.2	66.7	
A10 - with QSFP HS	4x SFP	35	Top	14-small	N/A	Hot-aisle	121.1	112.3	106.8	102.8	97.7	94.3	89.1	116.3	103.3	96.1	91.1	85.0	81.3	75.3	104.8	94.3	87.8	83.3	77.5	73.8	68.5	

Validation: Thermal

Goal: Standardize Thermal Validation for OCP

- Deterministic and standardize methodology across compute and storage platforms
  - Enabling both OCP Mezzanine v2.0 and v3.0
- Enable interoperability across Mezzanine add-in cards (NIC) and platforms
- Open sourced to OCP Community to enable independent qualification
  - System vendors
  - 3rd party labs
  - Mezzanine add-in card vendors.
- Design can be customized to adapt to different thermal test configurations.

Thermal Validation Test Fixture

- 3D Models released on Wiki
- Initial PCIe Gen4 capable schematic/layout/gerber package is posted for community feedback
- Allows validation of the OCP 3.0 card with/without the presence of a PCIe adapter immediately above the card.
- Correlation at the system level will allow a system designer to support 3rd party cards using FRU information.

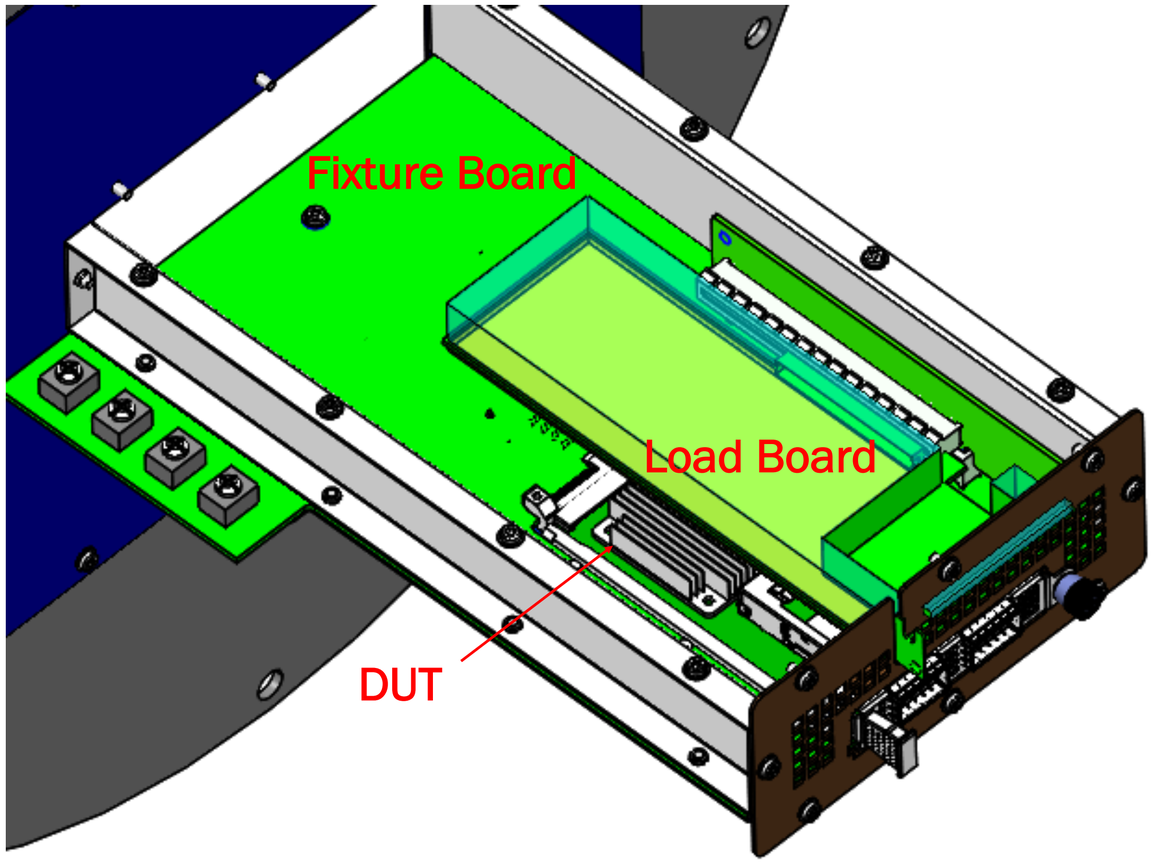

Validation: Electrical: PCIe

Goal: Standardize PCIe Conformance for OCP

- Deterministic and standardize methodology across compute and storage platforms
- Enable interoperability across Mezzanine add-in cards (NIC) and platforms
- PCIe Gen3+ interoperability and conformance
- Open sourced to OCP Community to enable independent qualification
  - System vendors
  - 3rd party labs
  - Mezzanine add-in card vendors.
- Full integration with standard industry equipment

Methodology

- Road to formalize conformance guidelines for OCP 3.0
  - Use OCP 2.0 systems as enablement platform
  - Incorporate SI guidelines and requirements into OCP 3.0 specifications
    - Target v0.9
- Cross community collaboration
  - Platform developers
  - NIC Developers
  - Test equipment vendors
- Integrate test solution with test equipment for automation

OCP vs PCI-CEM (Gen3+) Conformance Testing

- OCP 2.0 vs PCI-CEM

Attribute	*OCP*	*PCI-SIG*
Topology	1-3 Connectors	1-2 Connectors
Connector Z-height	3x versions (5mm, 8mm, 12mm)	1
System Loss	~-14.5dB	~-12dB
AIC Form-Factor	Mezzanine	CEM

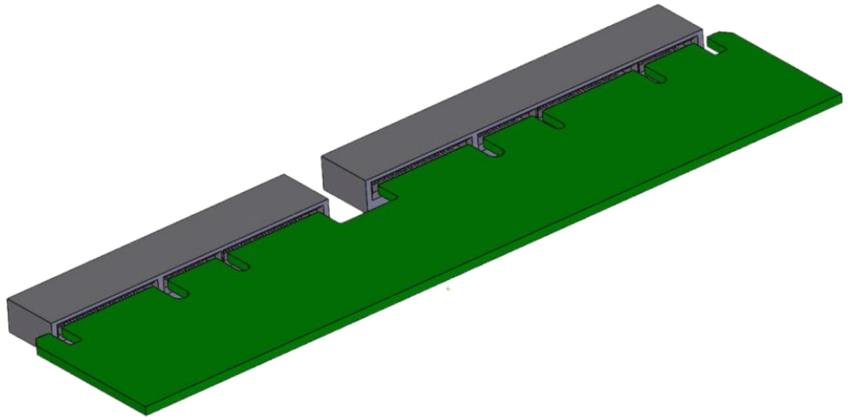

SFF-TA-1002 (Rev 1.1) 4C+ for OCP3.0

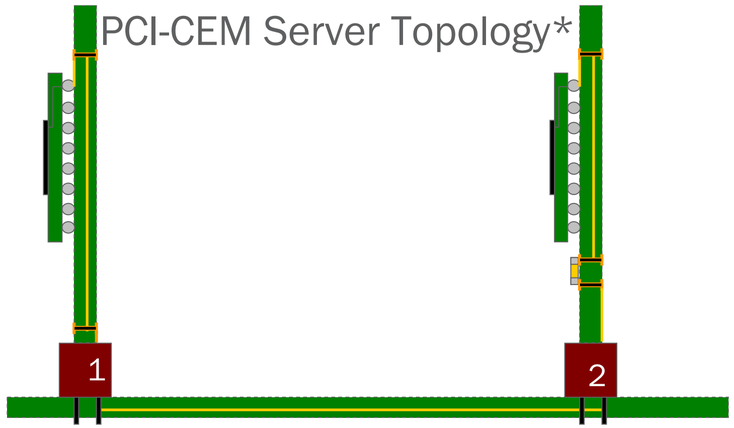

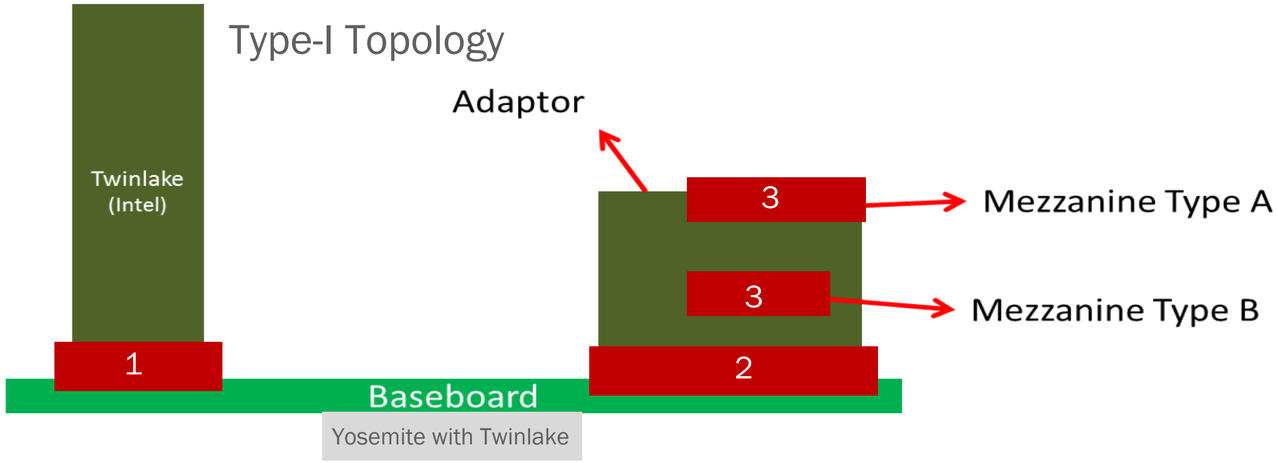

Yosemite with Twinlake

Source: Keysight

OCP Conformance Guideline

- Key conformance parameters
  - Tx SQ at far-end (Rx pins)
  - Rx sensitivity and jitter tolerance
- Acceptance guideline
  - Meet PCI-SIG Rx requirements at far-end
  - BER with Stressed Jitter Eye

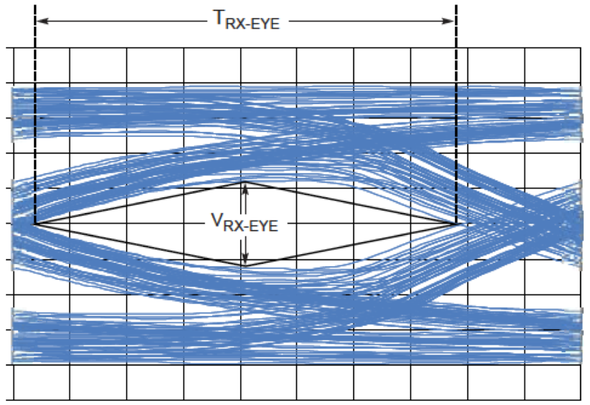

Table 4-23: Stressed Jitter Eye Parameters

Symbol	Parameter	Limits at 8.0 GT/s	Units	Comments
$V_{RX-LAUNCH-8G}$	Generator launch voltage	800 (nominal)	mVPP	Measured at TP1, see Figure 4-65. See Note 1.
$T_{RX-UI-8G}$	Unit Interval	125.00	ps	Nominal value is sufficient for Rx tolerancing. Value does not account for SSC.
$V_{RX-ST-8G}$	Eye height at TP2P	25 (min) 35 (max)	mVPP	At BER=10^{-12}. See Note 2.
$T_{RX-ST-8G}$	Eye width at TP2P	0.30	UI	At BER=10^{-12}. See Note 2.
$T_{RX-ST-SJ-8G}$	Sinusoidal Jitter	0.1 – 1.0	UI PP	See Figure 4-74 Measured at TP1. See Note 3.
$T_{RX-ST--RJ-8G}$	Random Jitter	3.0	ps RMS	Rj spectrally flat before filtering. Measured at TP1. See Note 4.

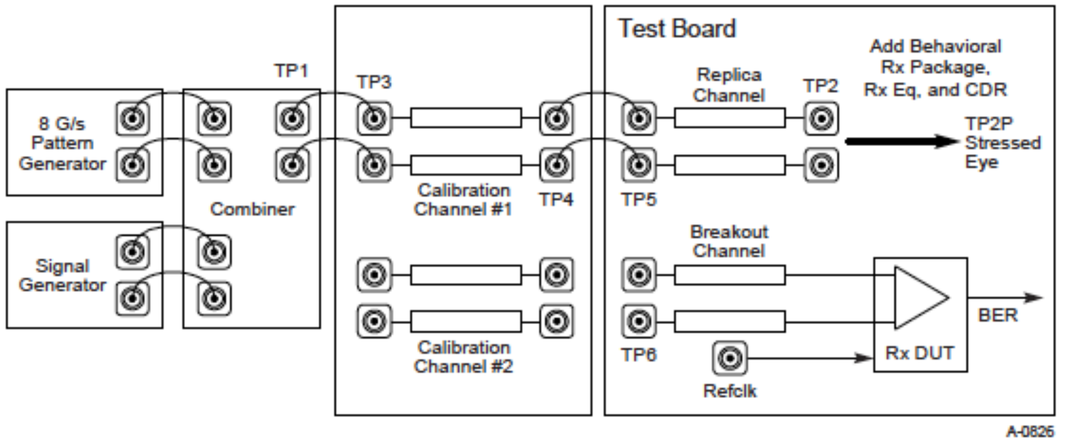

Figure 4-65: Rx Testboard Topology

*Source: PCI Express® Base Specification Revision 3.0

OCP Conformance Testing

- Test fixtures
  - Base board
  - Load board
  - Calibration board
- Conformance Measurement setup
  - Calibration
  - Test
  - SigTest Processing for Tx
    - Embed loss (pkg + worst channel)

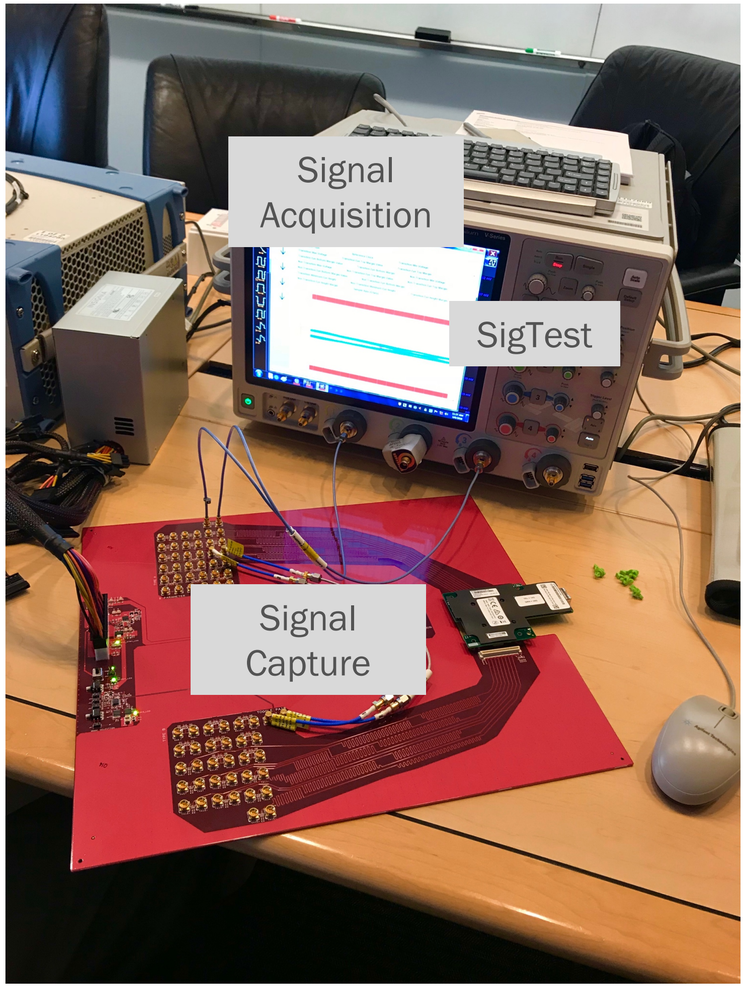

OCP 2.0 PCIe Tx Test Setup

OCP PCIe Test Fixtures

- OCP 2.0
  - Compute and Storage
  - 3x sets
    - Load Board
      - Type A
      - Type B
    - Base Board
  - Calibration

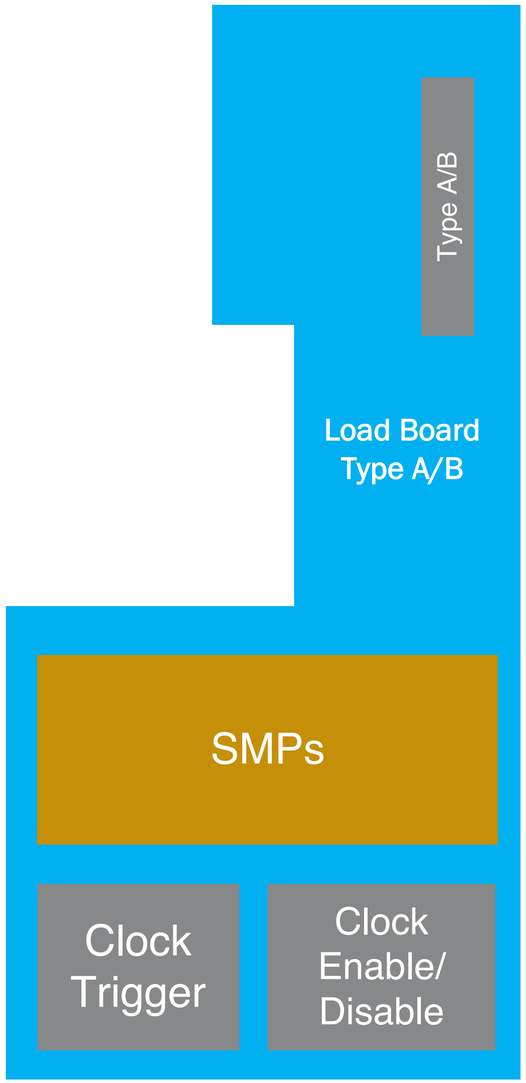

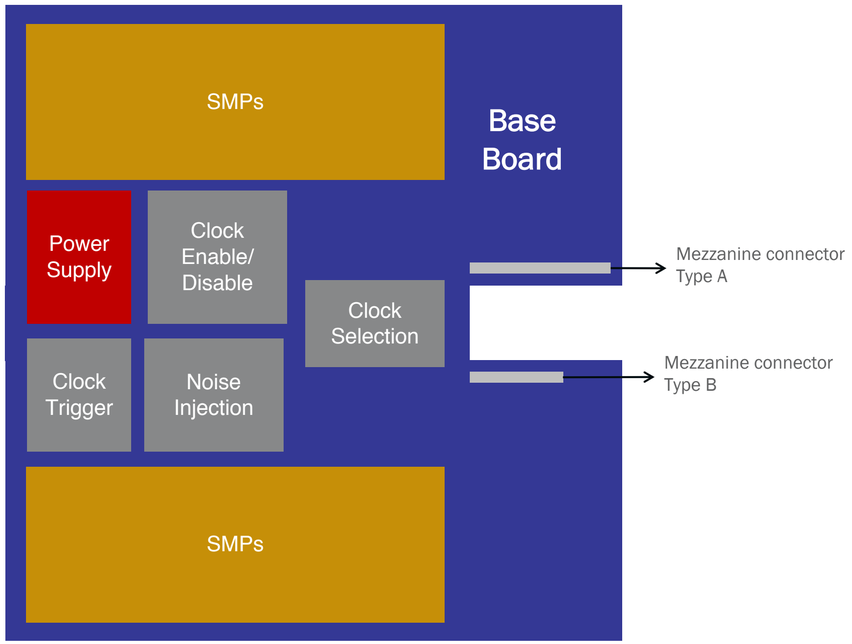

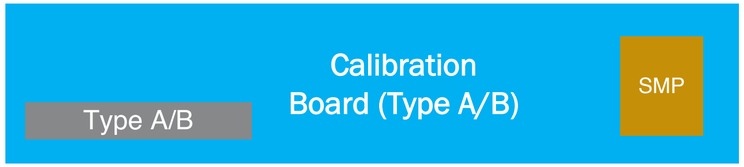

OCP3.0: Test Fixtures

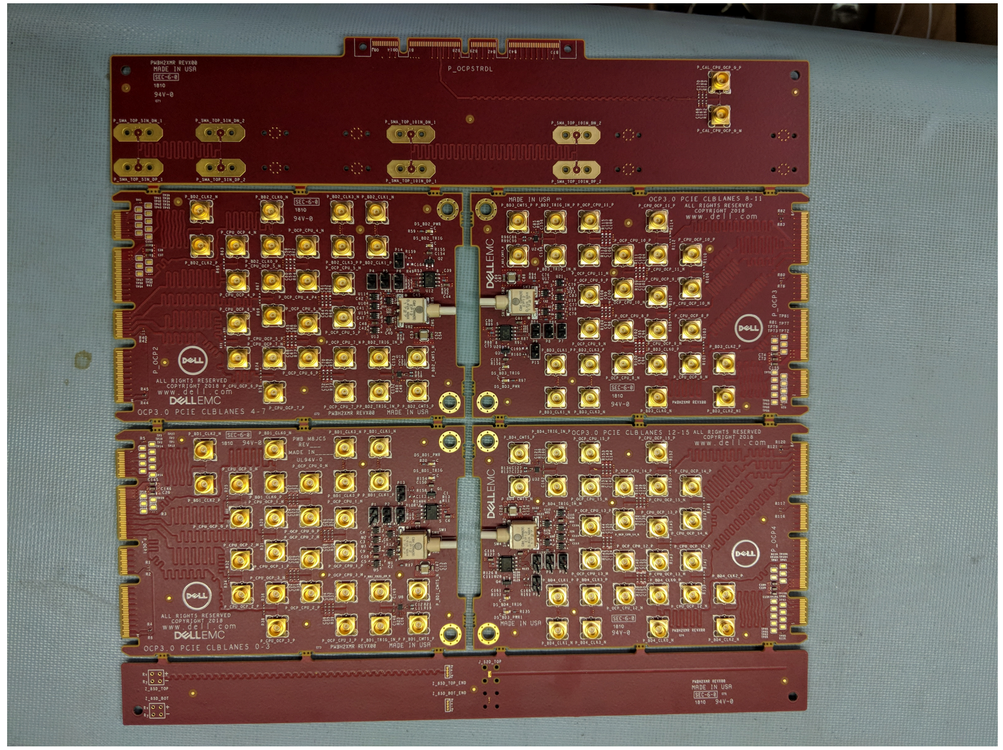

Load Board (1 of 4)

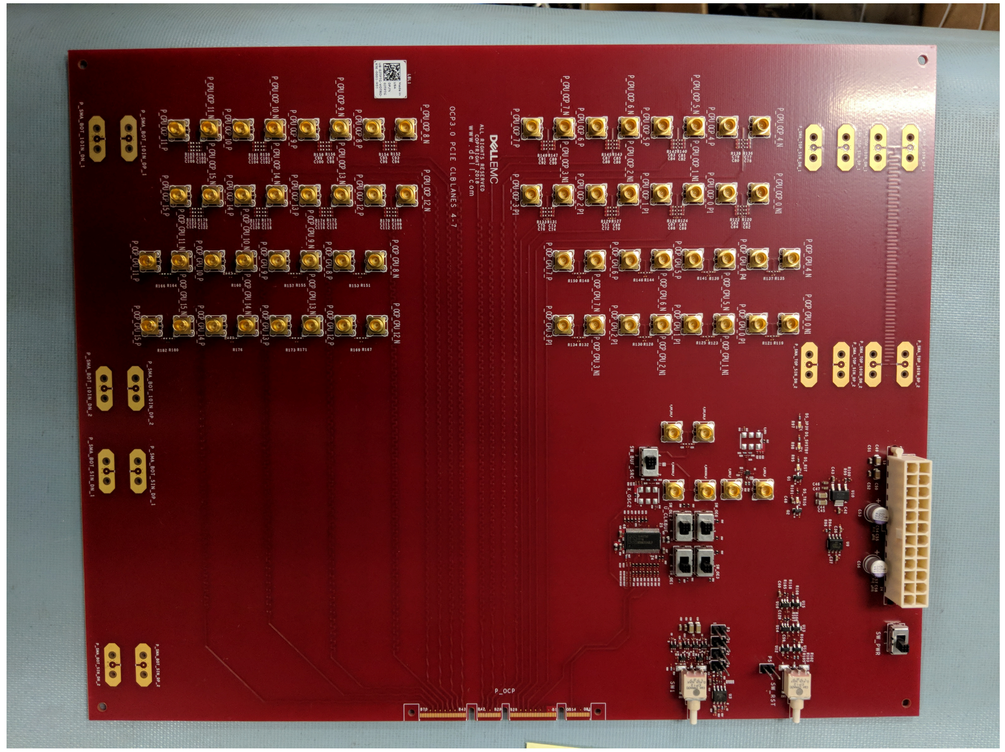

Base Board

**Courtesy: Dell-EMC*

Fixture Design

- Design
  - Material
    - OCP2.0
      - TU862 for BB and TU883 for LB
    - OCP3.0
      - TU863 for BB and TU883 for LB
  - Equal length channels
    - All Tx/Rx channels, including calibration
  - Clock trigger/power delivery similar to PCI-SIG CLB/CBB fixtures

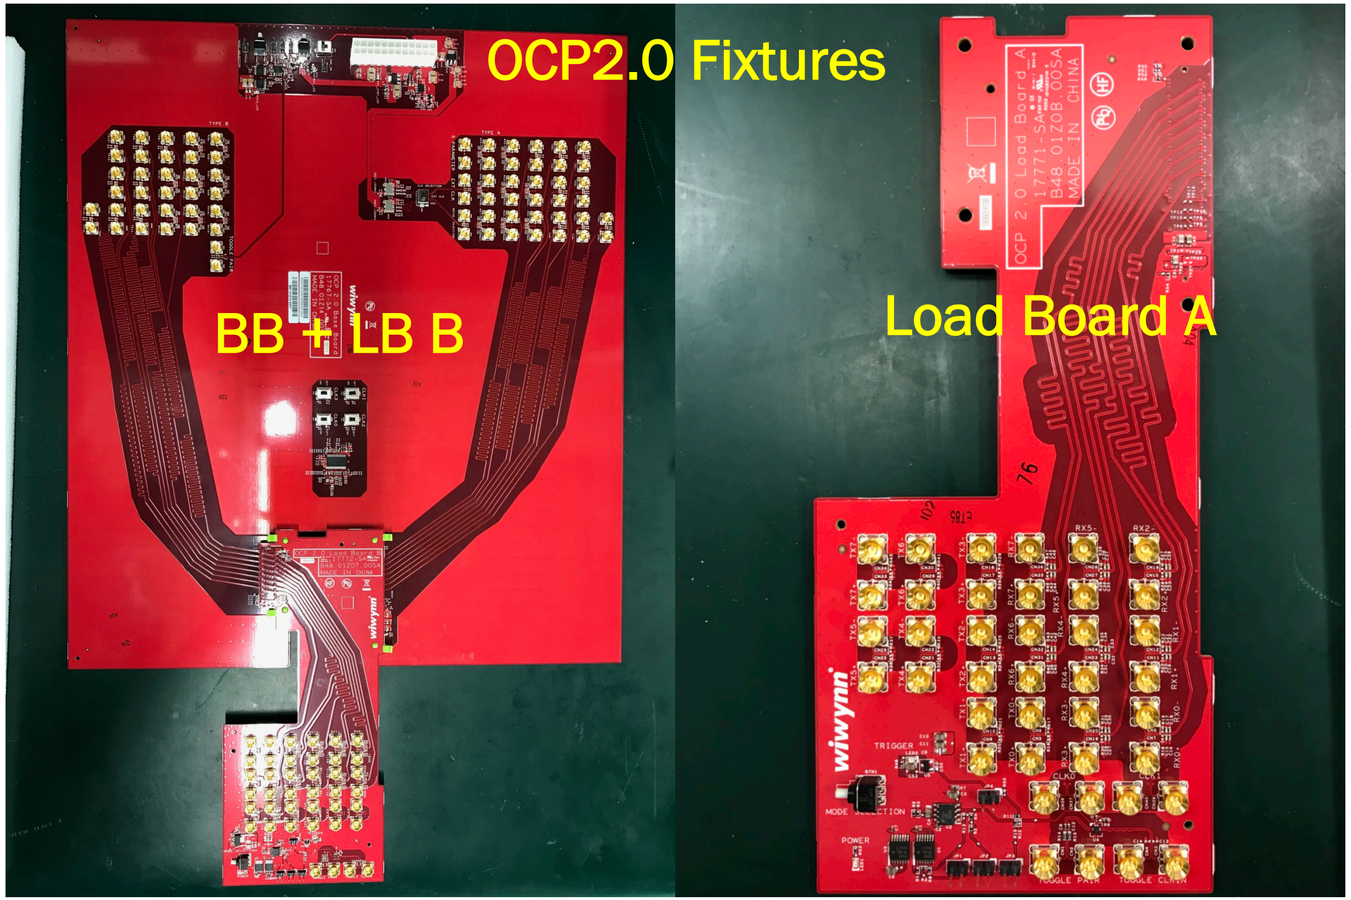

OCP2.0 PCIe Test Fixtures: Loss Characteristics

Base Board (bare)

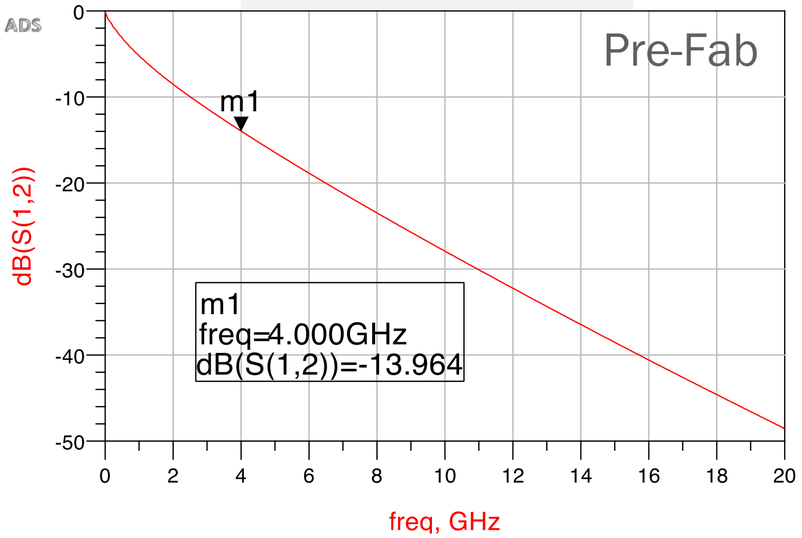

Load Board (bare)

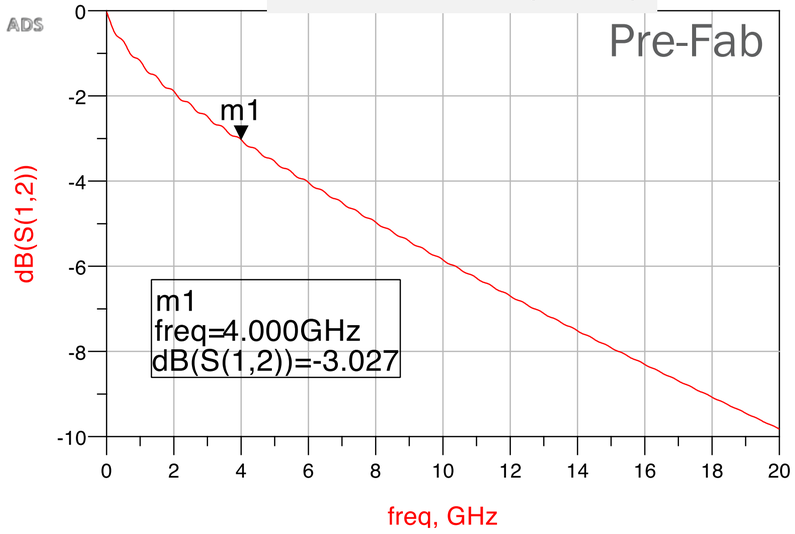

Test Fixture	***Design Target*** *(with Connectors)*	***Measured **** *(with Connectors)*
Baseboard	-14.75dB	-14.5 dB
Load Board Type A	-3.45dB	-3.52 dB
Load Board Type B	-3.45dB	-3.24 dB

** Sample average*

OCP2.0: Assessing Channel Margin

- Channel Margin
  - Per Simulation
    - AIC: -4.43dB
    - System: -14.93dB
  - Per Measurement
    - *Under work*

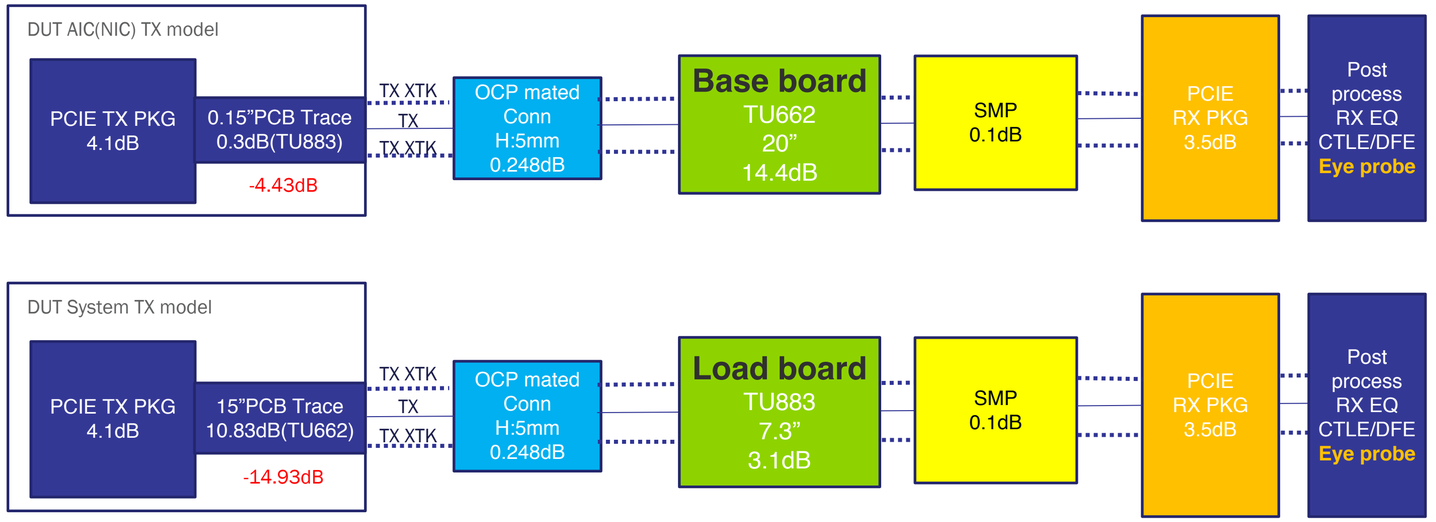

Topology	EH(mV) BER@e-12 (CEM:34mv)	EW(ps) BER@e-12 (CEM:41.25ps)	Jitter PP/RMS
DUT NIC + Mezz Connector + Baseboard + SMP + Rx pkg	48mv (TXEQ:P7)	66.88ps	46.88/7.37 (ps)
Yosemite V2 + Mezz Conn + LoadBoard + SMP + Rx pkg	39mv (TXEQ:P8)	61.87ps	46.88/ 7.63(ps)

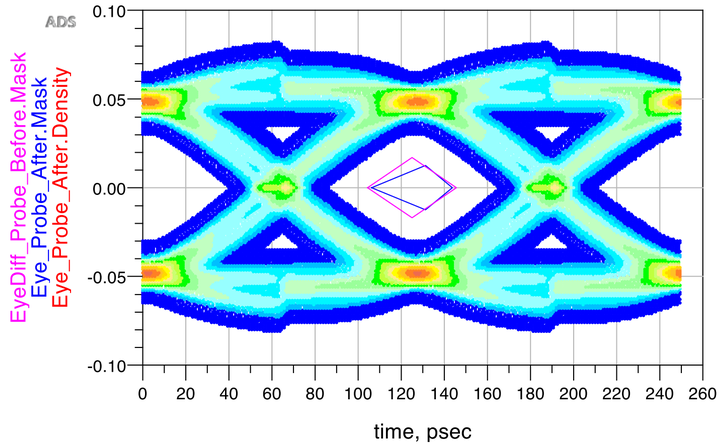

Calibration Board

Characterize channel loss of test boards to enable AFR (Automatic Fixture Removal)

- Design replica channel of test fixture
  - 2-L philosophy
  - Board built on same panel as test board
  - Serialization with test board
  - Reference Channel
- Advantage
  - Characterize individual test board
  - Accurate loss embedding

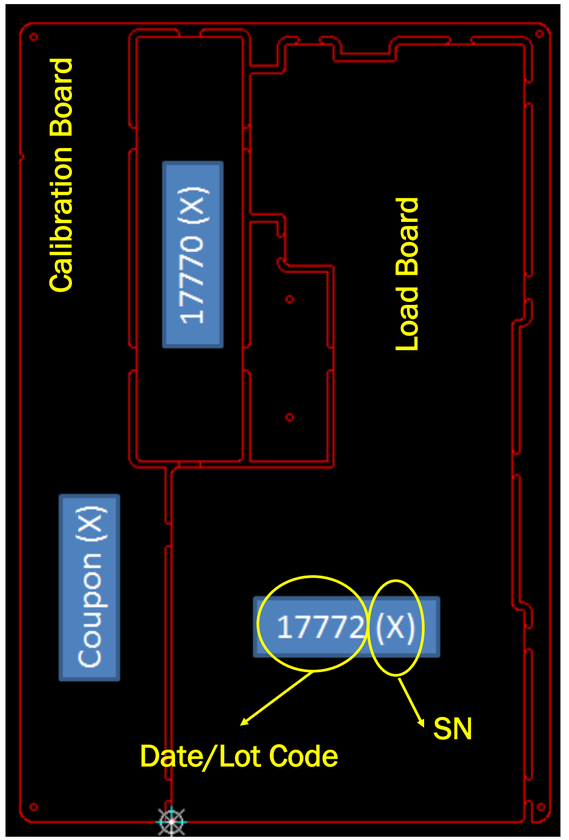

Serialization of Test Board, Calibration and Coupons

Calibration Design and Configuration

- Replica Channel
  - 2-L measured insertion-loss
  - Divide by 2 to obtain loss for individual board
- Single Differential channel
  - All Tx/Rx channels equal length within < 1%
- 1:1 association with test board
  - Same panel and serialized

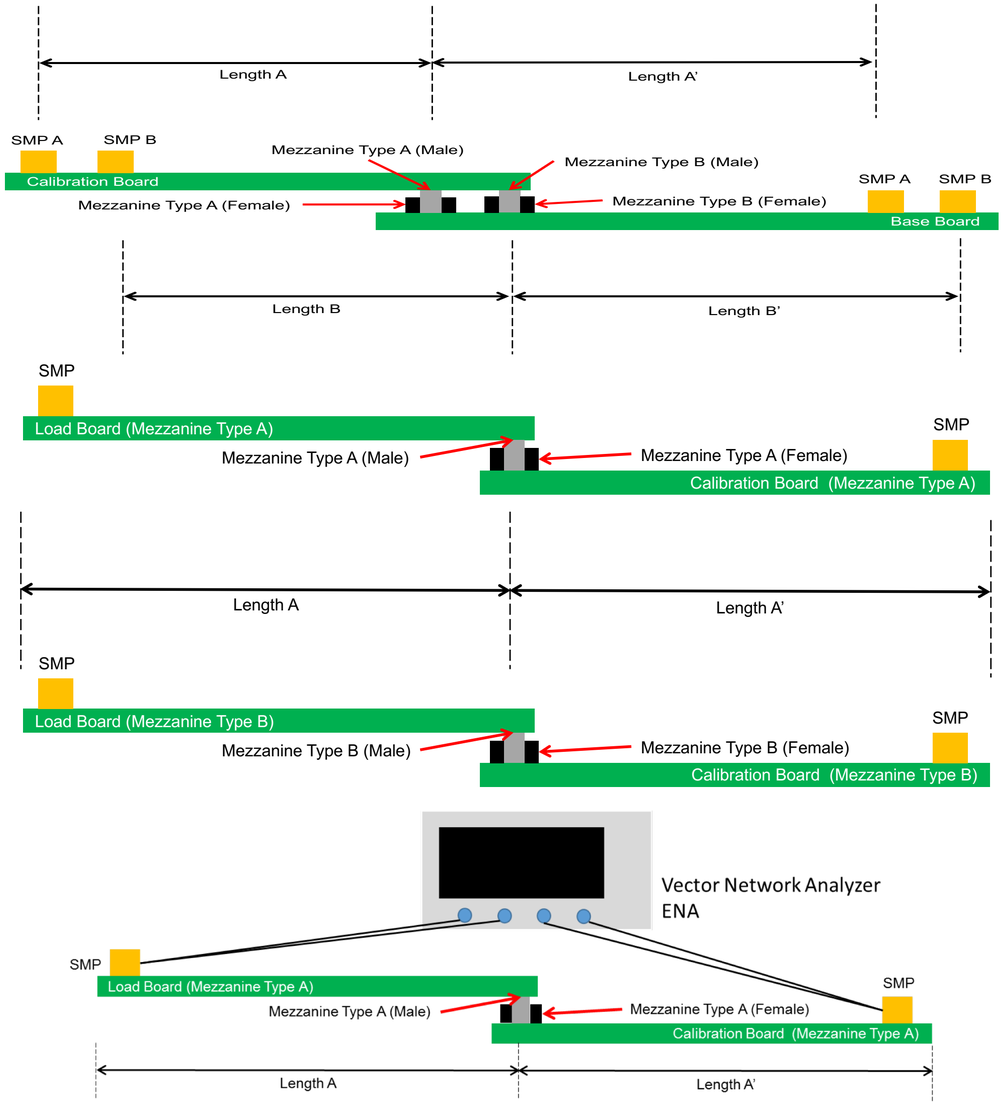

OCP2.0: Fixtures and Platforms

Bryce Canyon with LB A

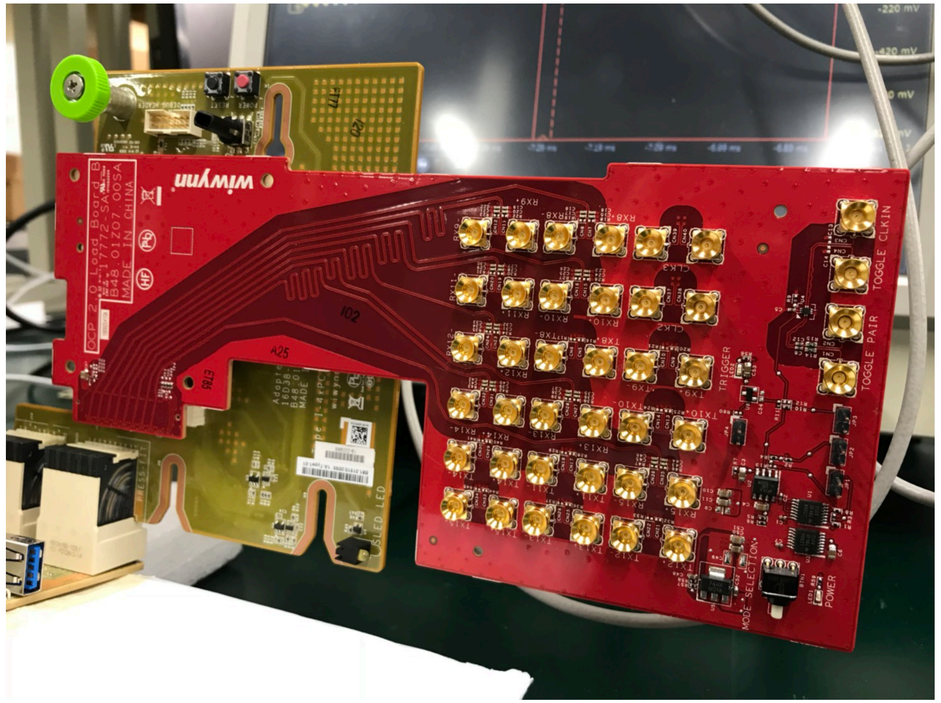

Yosemite V2 with LB B

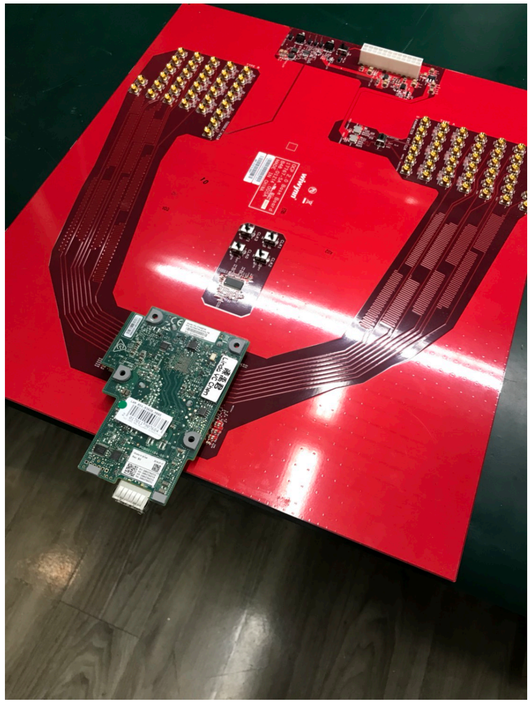

Mezzanine NIC with BB

PCIe Conformance Testing for OCP platforms: Rick Eads, Keysight

KEYSIGHT TECHNOLOGIES

OPEN. FOR BUSINESS.

OCP SUMMIT

PCIe Physical Layer Compliance

- Focus is on the interoperability predictors
  - Signal quality
  - Receiver testing
- The PCIe Base Specification architecture
  - Transmitter (TX)
  - Receiver (RX)
  - Channel
- Three main PCIe Spec Documents
  - PCIe BASE Specification (BASE Spec)
  - PCIe Card Electromechanical Specification (CEM Spec)
  - PCIe Test Specification (Test Spec)

KEYSIGHT TECHNOLOGIES

PCIe Channels Drive Spec Definition

Target Client Topology

10”, 1 connector
4-layer microstrip

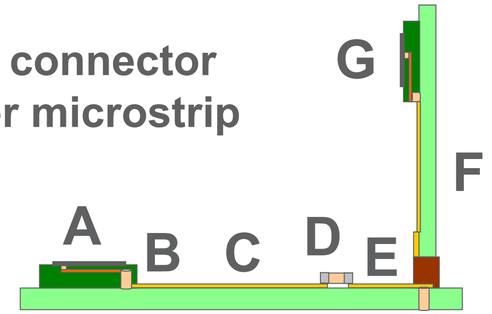

Seg	Description
A	RC PKG (transmitter)
B	RC break-out ~1”
C	MB Main 3-7”
D	MB coupling cap
E	Add in card break in 3”
F	EP break-out
G	EP PKG (receiver)

Target Server Topology

Stripline with via stubs
6- 8 layers, ”, 1 or 2 connectors

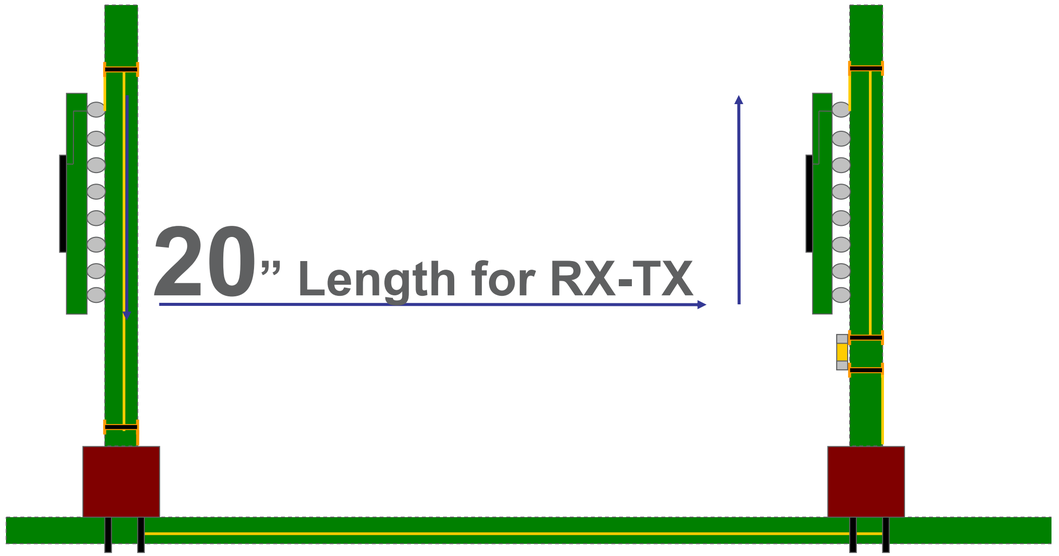

KEYSIGHT TECHNOLOGIES

PCIe Transmitter Test

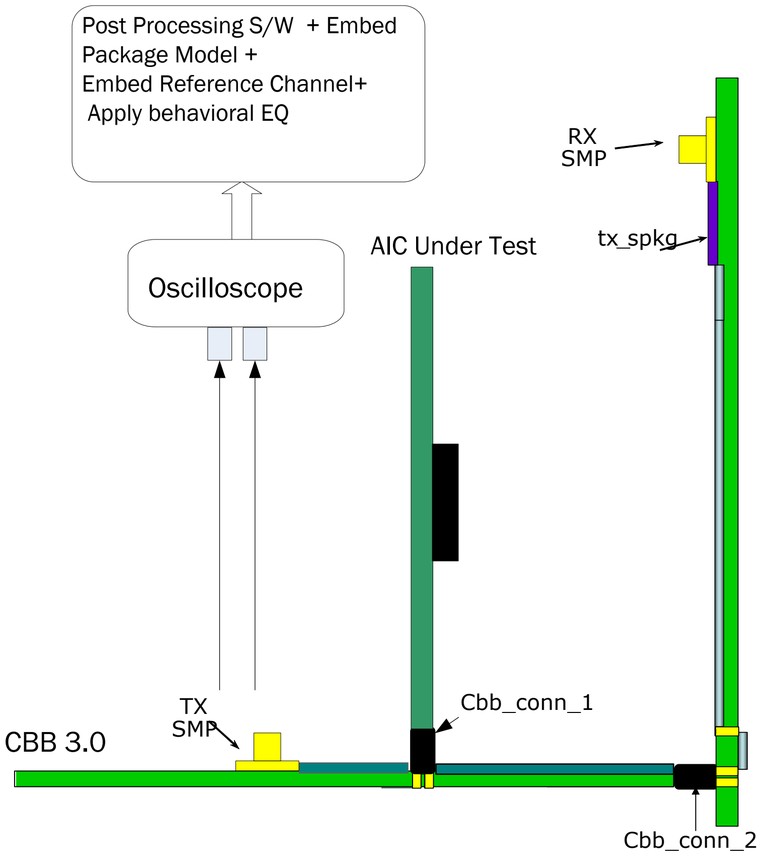

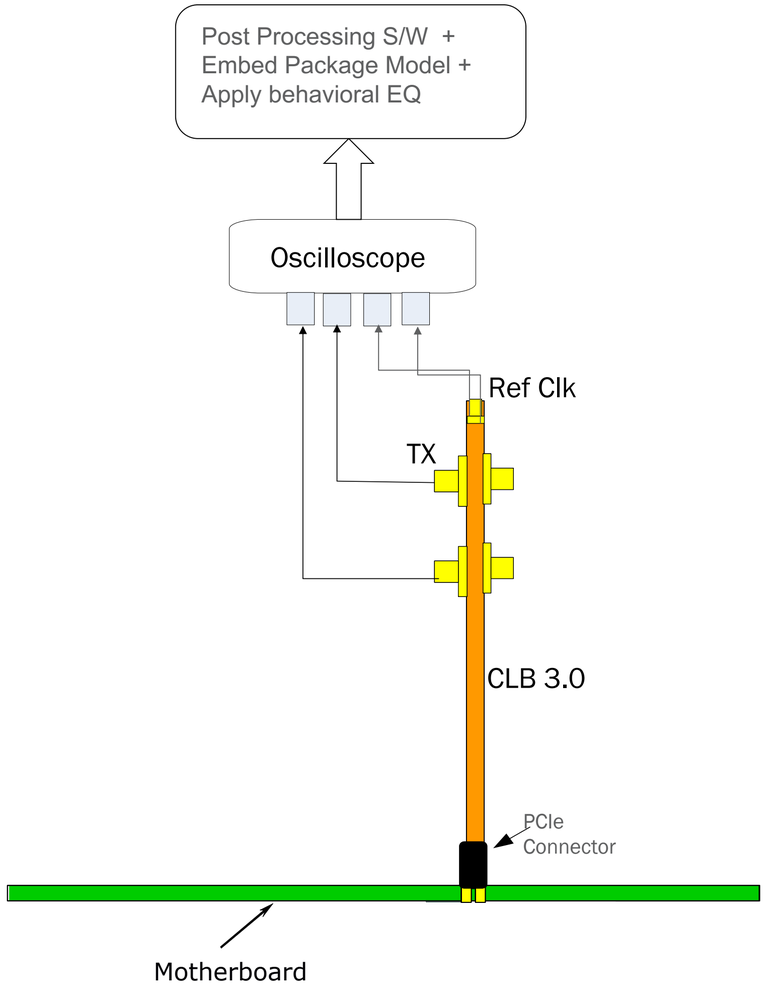

Conventional TX Test Setup

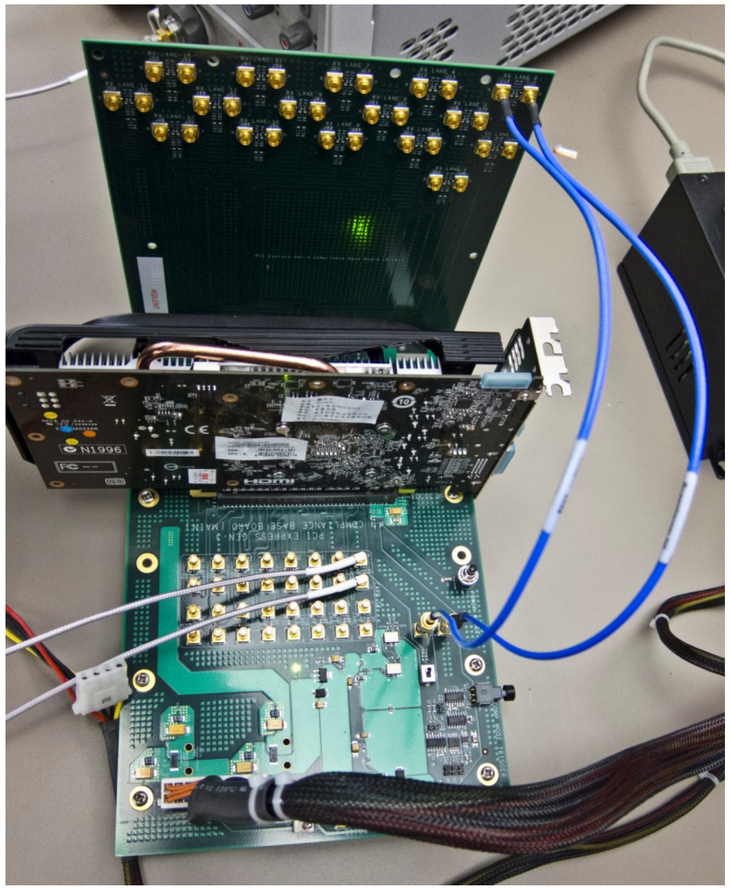

Add-in Card Test Example

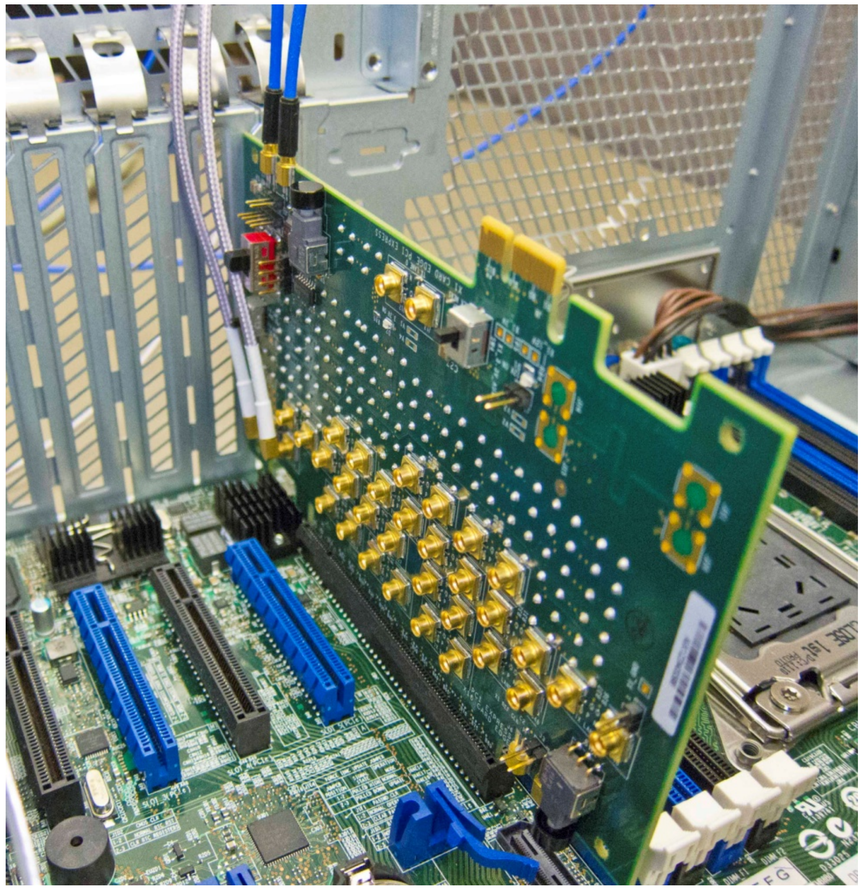

Motherboard Test Example

PCIe 3.0 Test Procedure

1. Setup DUT connection to Instrument
   - Calibrate and de-skew cables/connectors
2. Confirm Compliance Pattern at 8GT/s
3. Capture waveform on instrument
   - Differential TX signal for AIC test.
   - Differential TX and CLK signal for Motherboard test.
4. Post Process Data
   - Load data waveform in Sigtest (v3.2.0)
5. Change physical connection to next lane and repeat

KEYSIGHT TECHNOLOGIES

PCIe Compliance Patterns

2.5GT/s Compliance Pattern

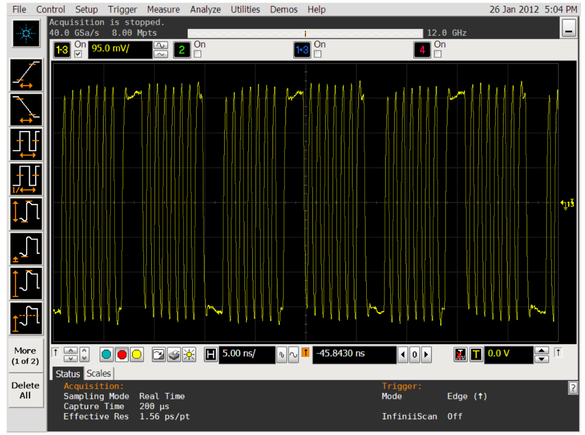

5GT/s Compliance Pattern

8GT/s Compliance Pattern

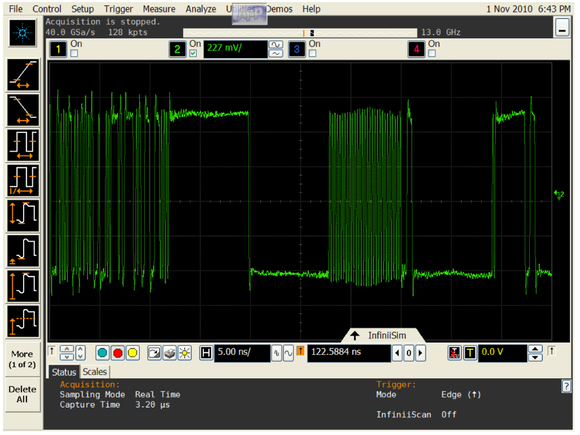

2.5GT/s Transition Eye

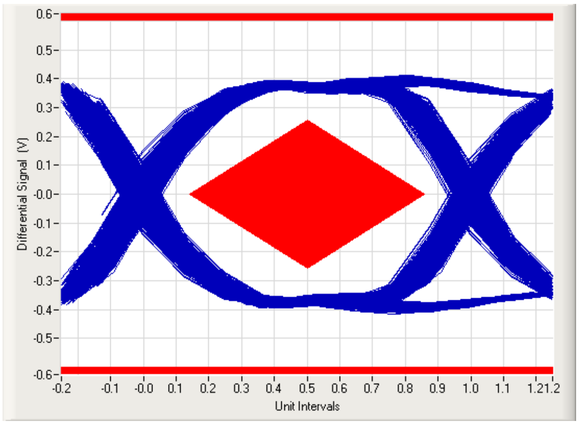

5GT/s Transition Eye

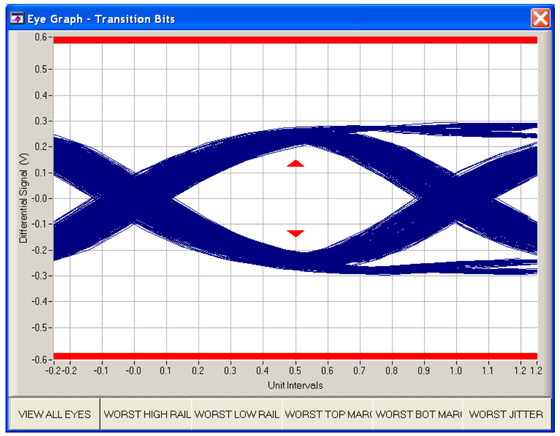

8GT/s Transition Eye

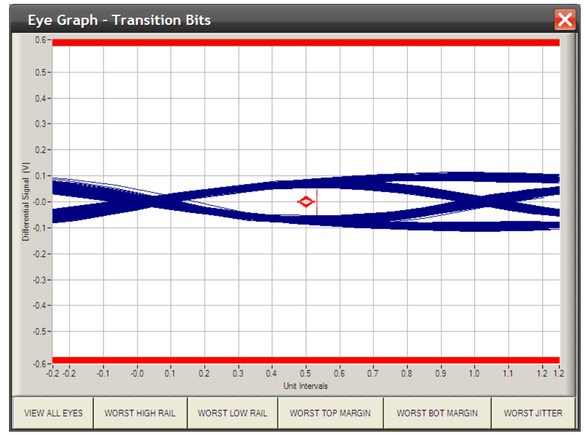

KEYSIGHT TECHNOLOGIES

AIC Rx Test Setup

- RX Testing required for 8GT/s
- Measurement Method
- BERT Based stressed jitter test

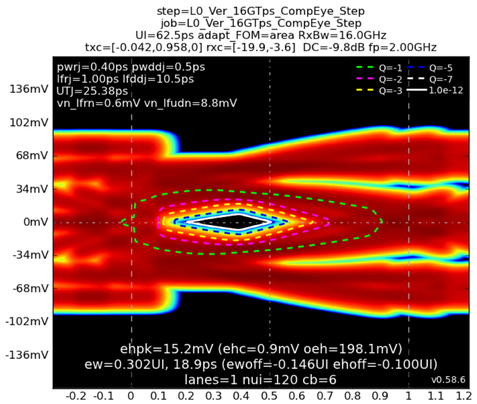

Stressed Eye: 8G: 25mV, 16G 15mV EH

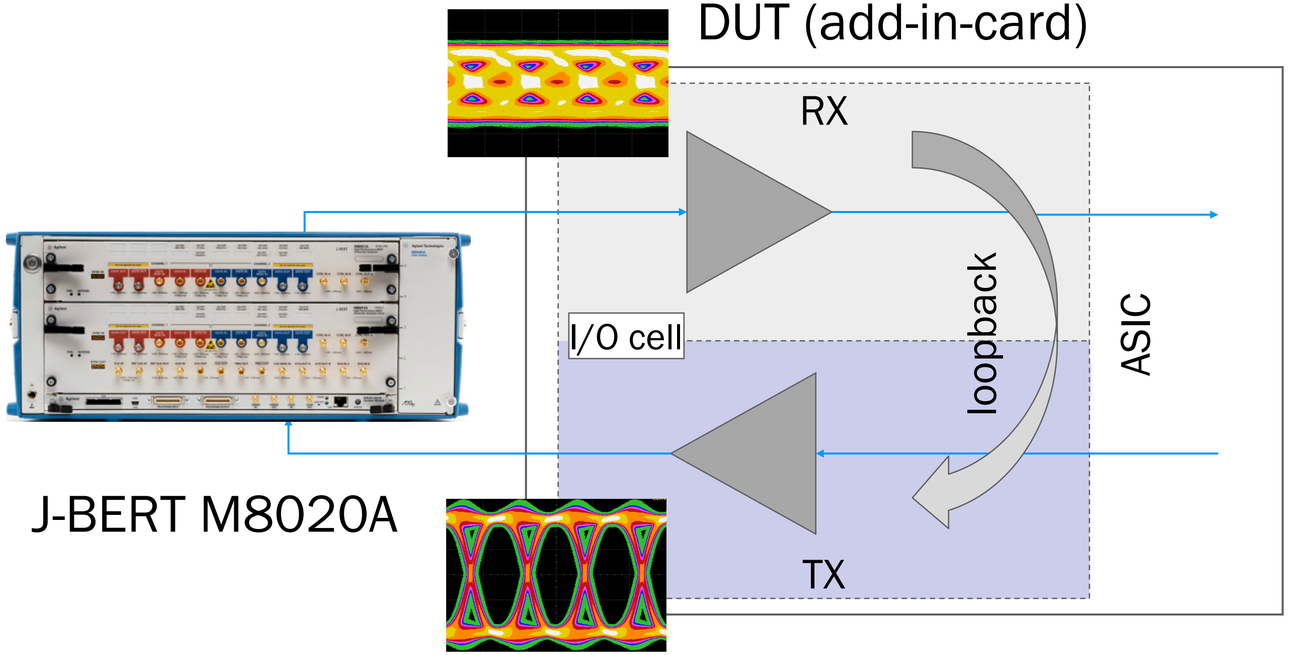

PCIe 3.0 Jitter Transfer Measurements

- Shows sensitivity of DUT TX output to Jitter on Reference Clock
- Requires the ability to drive a precision modulated, external reference clock

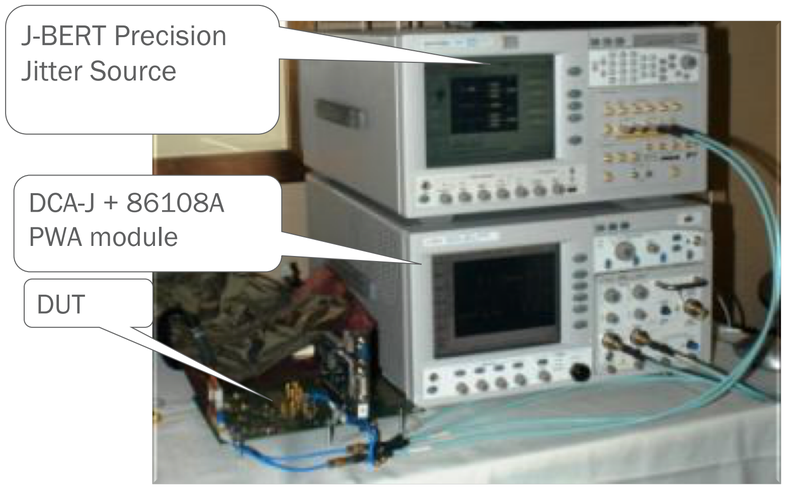

Recommended Equipment:

- Keysight 81150A, 81160A ARB or N4903B J-BERT
- Keysight 86100C DCA-J or N4877A CDR+DeMux
- Keysight 8108A Precision Waveform Analyzer (or 83496B Clock Recovery Module
- Access to TX and Reference Clock of DUT

Steps:

- Calibrate Jitter Source to CDR module.
- Set target device into Gen2 Compliance mode
- Run Keysight 86100CU-400 PLL and jitter spectrum measurement software
- Record PLL Response

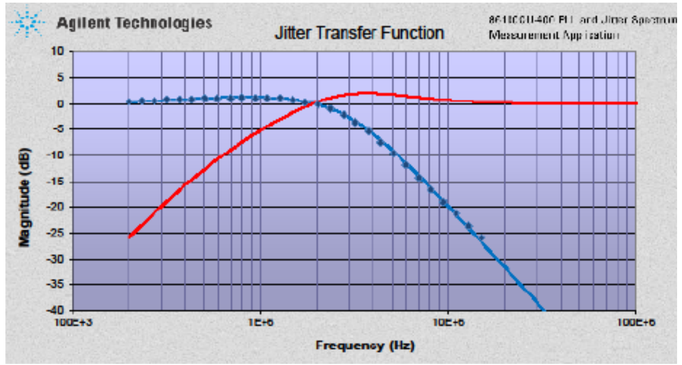

PCIe 3.0 DUT TX PLL Response

Standard	PLL BW	Peaking
PCIe 1.1	1.5MHz-22MHz	3dB
PCIe 2.0 (5GT/s)	5MHz-16MHz	1dB-3dB
PCIe 3.0 (8GT/s)	2MHz-5MHz	2dB-1dB

KEYSIGHT TECHNOLOGIES

PCIe 4.0 TX Measurement Test Setup - ASIC

- BASE Spec

Keysight Z-Series Real Time Oscilloscope

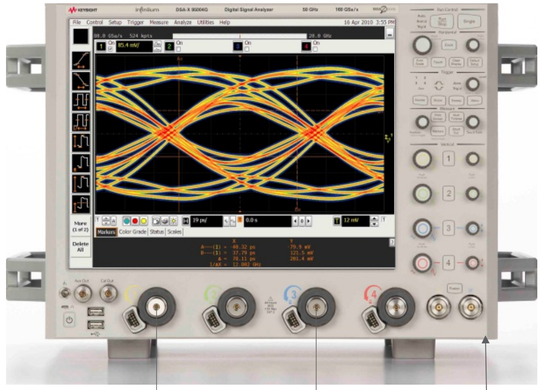

PCIe 4.0 ASIC/IC Custom Breakout Board

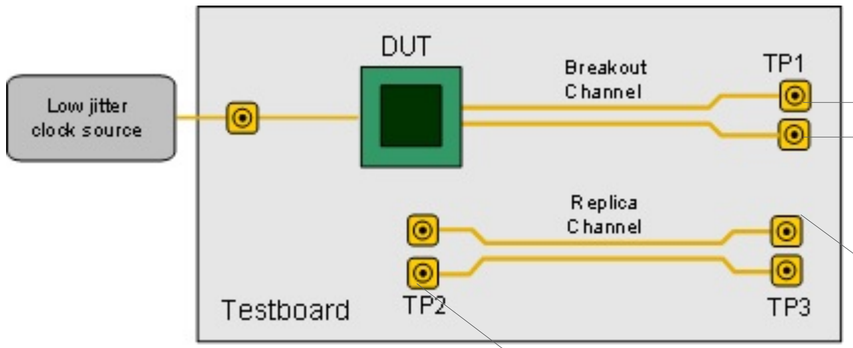

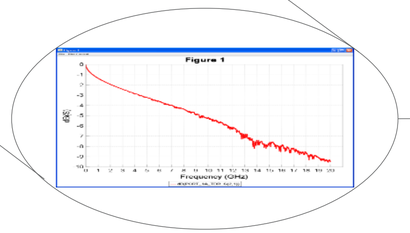

S-Parameters of Replica Ch.
Used to de-embed to pin or
Ref CTLE can be used (12dB).

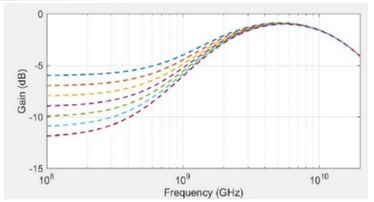

KEYSIGHT TECHNOLOGIES

Keysight PCIe 4.0 (Gen4) TX N5393F Test Applications

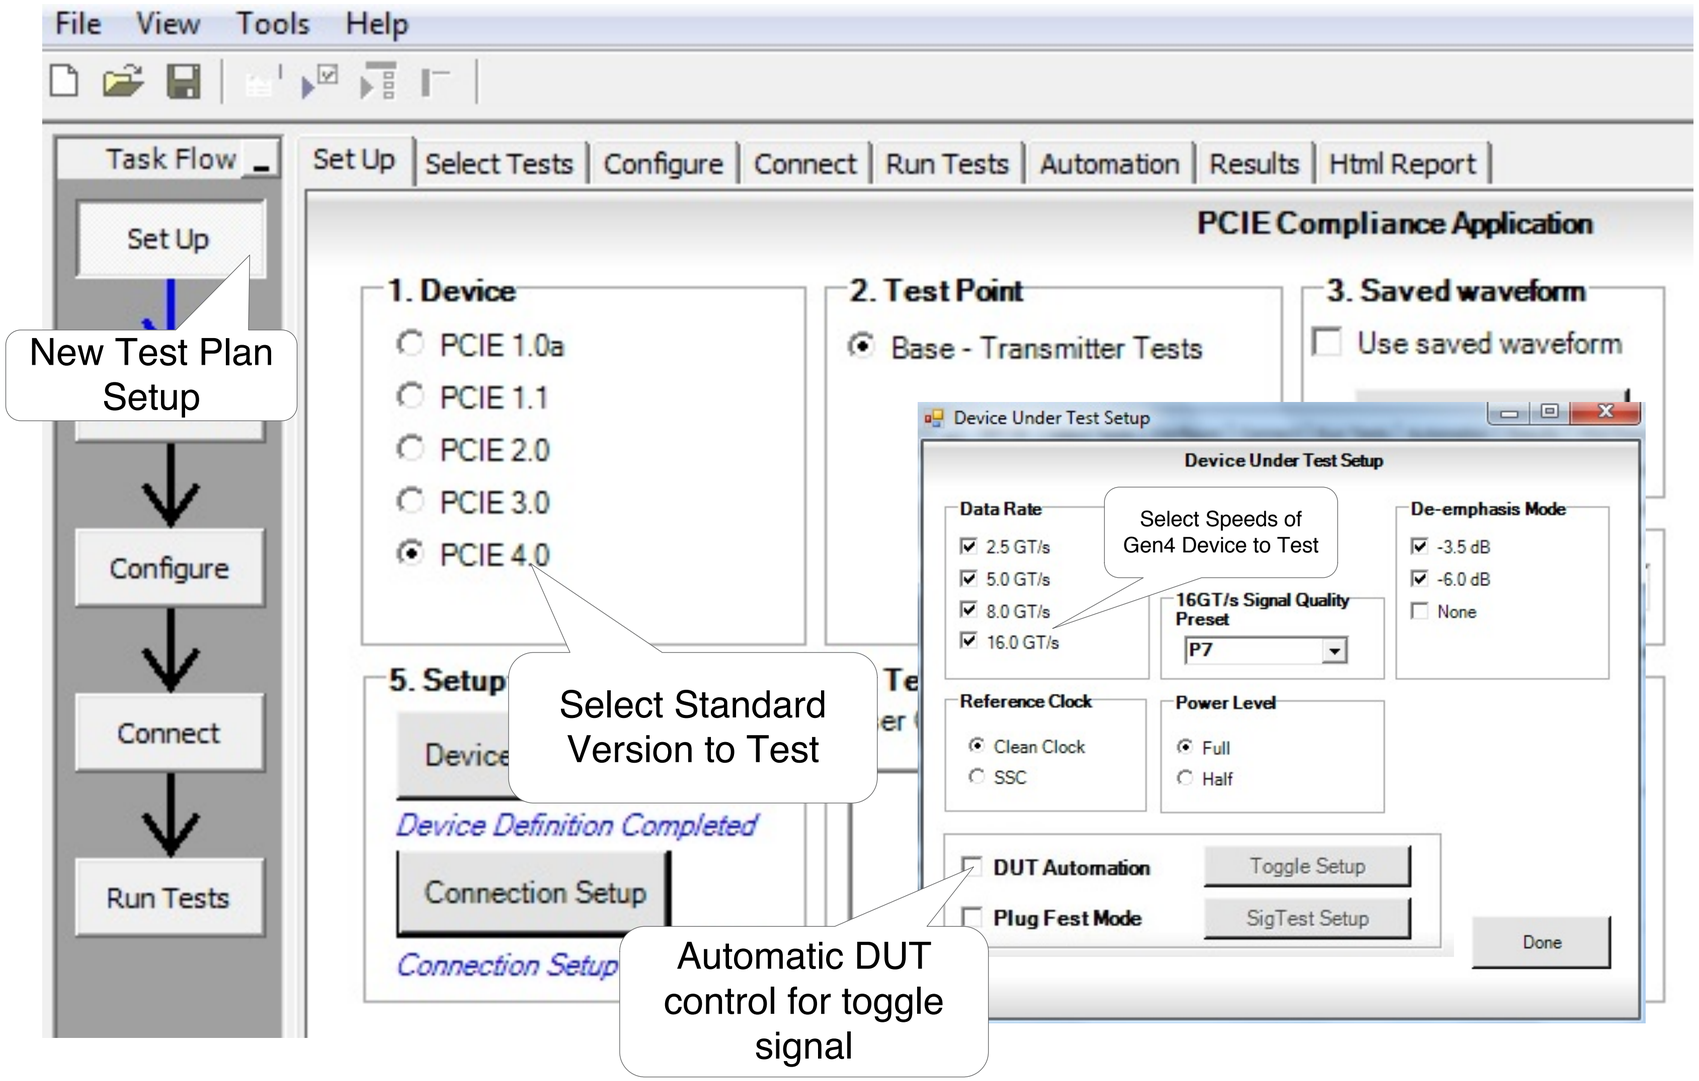

PCIe 4.0 RX Stress Signal Calibration

- 16 GT/s Receiver Stress Signal Calibration Setup – 2

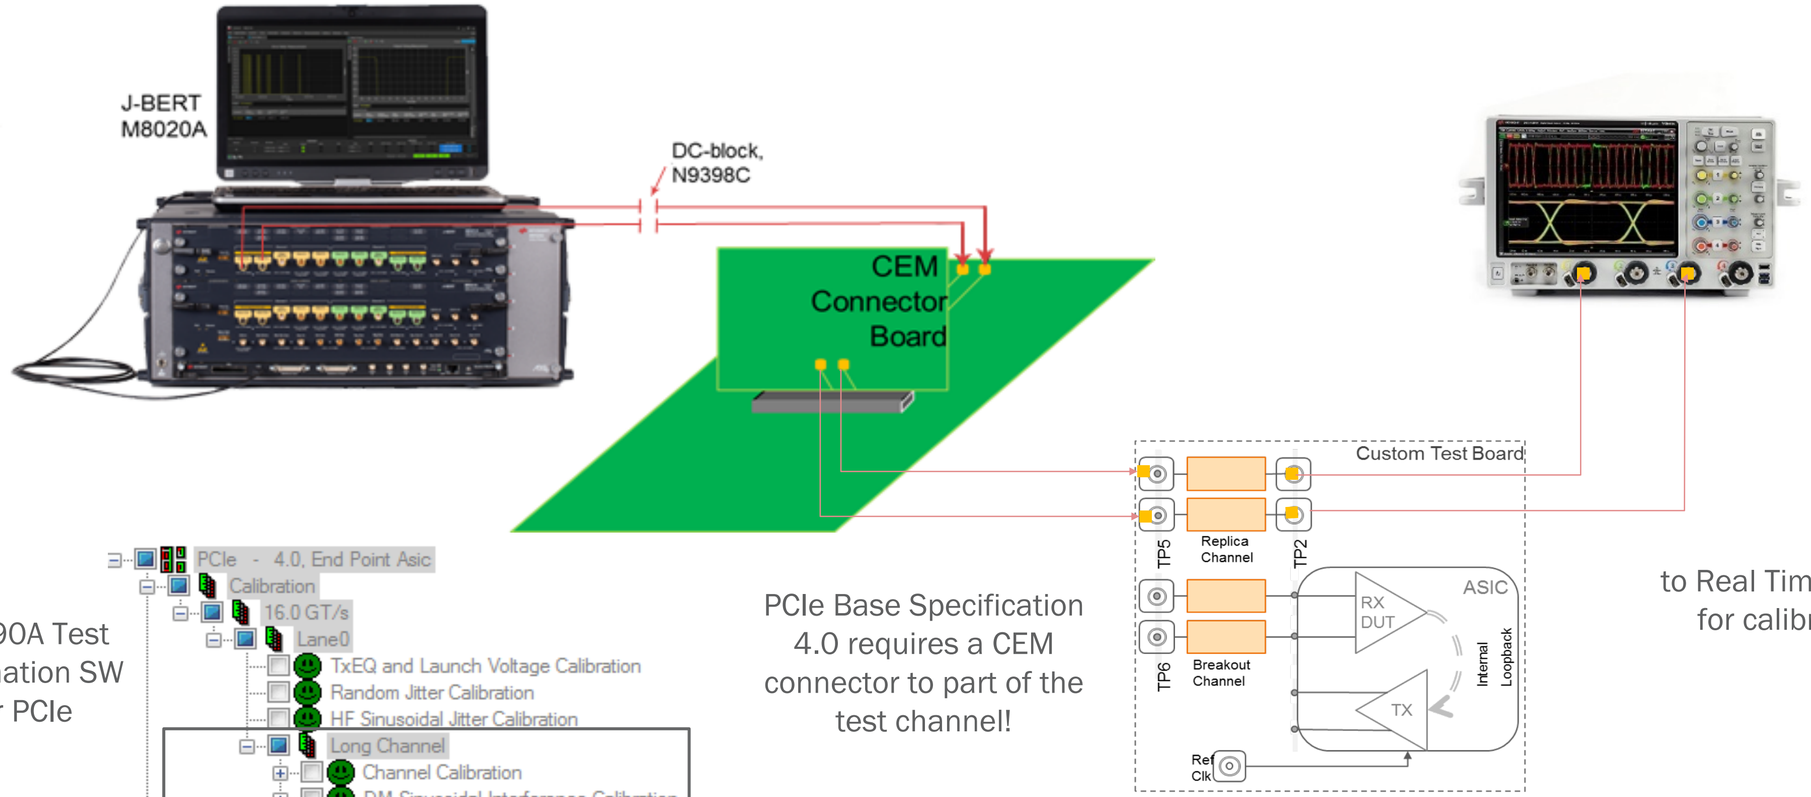

N5990A Test Automation SW for PCIe

PCIe Base Specification 4.0 requires a CEM connector to part of the test channel!

to Real Time Scope for calibration

PCIe 4.0 – 16 GT/s CEM Test Setup

- Calibration Setup for 16 GT/s RX

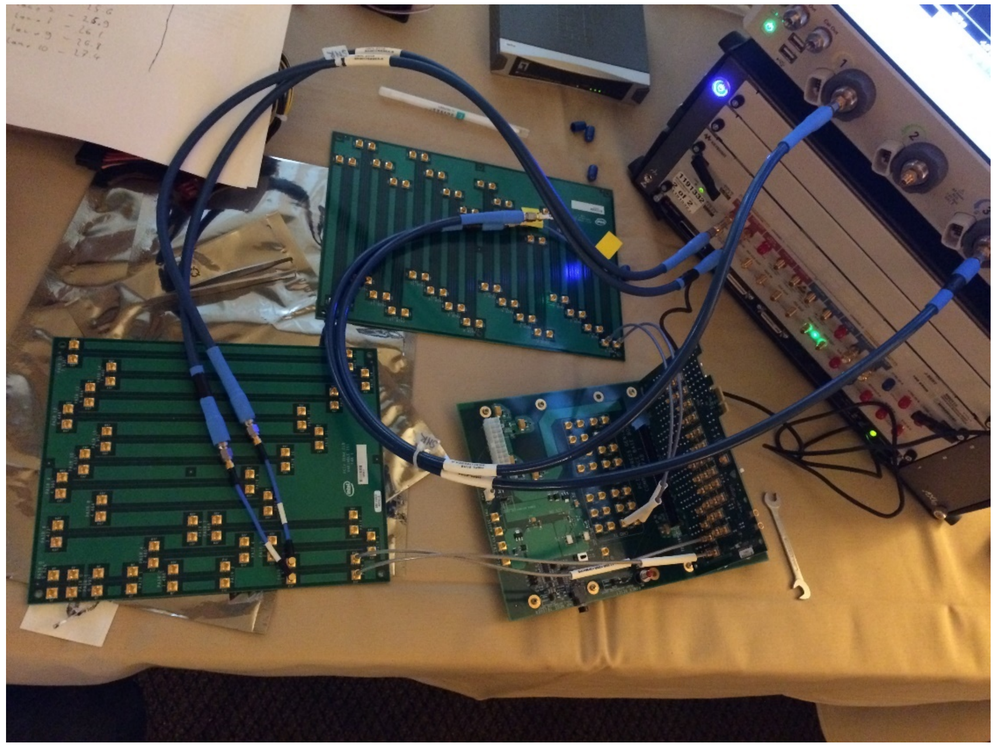

- CBB 4.0 as well as CLB 4.0 need to be combined with ISI trace boards
- CEM calibration procedure is very similar to base spec calibration but SIGTEST instead of SEASIM is mandatory
- J-BERT M8020A successfully tested most of the 16 GT/s AICs and systems at PCIe WS 101
- Many AICs and systems could be trained to loopback using the new LTSSM

KEYSIGHT TECHNOLOGIES

RX-TX Test Summary

- PCIe Gen4 RX and TX test tools available today (BASE Spec)
- PCIe Gen4 CEM test fixtures still awaiting full release from PCISIG. Early previews available from Keysight.

PCIe – Keysight Total Solution PCIe 3.0-4.0

Physical layer – interconnect design

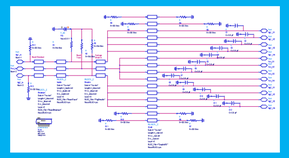

ADS design software

86100D DCA-J/TDR

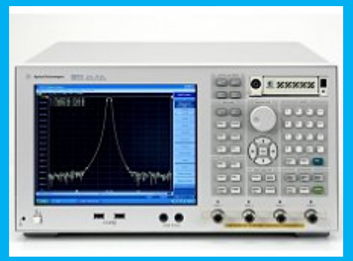

E5071C ENA option TDR

Physical layer-transmitter test

90000 V-Series oscilloscope

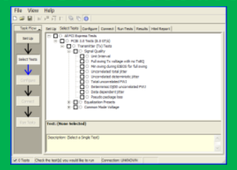

N5393F PCIe electrical compliance software

86100CU-400 PLL and Jitter Spectrum Measurement SW

Physical layer-receiver test

M8020A J-BERT High Perfformance BERT

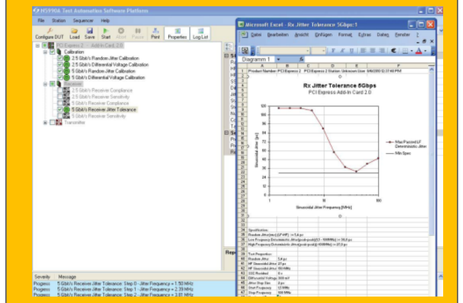

N5990A automated compliance and device characterization test software

Data link/transaction layer

Digital Test Console
• U4301A Protocol Analyzer

• U4305A Exerciser
•Protocol Test Card

• Multiple probes with ESP technology

Industry's lowest scope noise floor/sensitivity and trigger jitter

DSA-X Series & V Series Real-Time Oscilloscopes

Automated compliance software – accurate, efficient and consistent

X1 through x16 Analysis and Exerciser support, with industry's only ESP probing technology

KEYSIGHT TECHNOLOGIES

For further information

You will find more information on PCIe 3.0 and Keysight solutions for PCI Express at:

www.pci-sig.com	PCI-SIG Website, Specification, S/W Tools, Keysight Test Procedure
www.Keysight.com/find/pciexpress	Keysight tools to help you succeed with your PCI Express design such as the N5393C Compliance application.
www.Keysight.com/find/si	Keysight tools to help you master signal integrity challenges.
www.Keysight.com/find/PCIe_receiver_test	PCIe 3.0/4.0 Rx Test Information

KEYSIGHT TECHNOLOGIES

Future Work

Future Work

- OCP 3.0 NIC Design
- Validation
  - Thermal
  - Electrical
- PCIe Conformance Testing for OCP 3.0

OCP 3.0 NIC Design

- Continue driving specification to 1.0 by EOY 2018 closing open items
  - Thermal correlation with the test fixture (SFF an LFF)
  - Electrical characterization of 3.0 PCIe test fixtures
- Formalize OCP 3.0 channel guidelines (Target v0.9)
  - Platforms
  - Add-in Cards

Validation: Electrical: OCP 3.0

- Release Conformance guidelines
  - Communicate plan by end of March 2018
  - Incorporate details as part of v0.9 (Target August 2018)
    - Including acceptance criteria
    - Fixture design and release
- Engage 3rd Party Lab
  - UNH-IOL
- Enable Test automation
  - Similar to conventional PCI-CEM
    - AFR, embedding models, test equipment setup etc.

Validation: Electrical: OCP 2.0

- Release guidelines
  - Target: 2Q18
    - Establish acceptance criteria
    - Finalize conformance guidelines
- Fixture design
  - Schematics, layout
  - Channel simulations
  - BOM
  - Component data sheets
  - Fixture validation reports
  - Supporting documentation (test reports, specifications, datasheet etc.)
- Schedule
  - OCP 2.0: 2Q18

PCIe Conformance for OCP 3.0

- Leverage CEM 4.0
  - Testing methodology
  - Test Fixture design, only modified for Form Factor
- Define testability features to simplify/automate testing
- Enable 3rd party conformance testing worldwide
- Follow PCIe Gen 5 roadmap for future inclusion if technically feasible
  - 0.9 Gen 5 PCIe spec estimated at Q3 2019
  - Signaling rate: 32 GT/s NRZ
  - Proposed channel loss target is: -36dB @ 16GHz

Keysight tools for PCIe 5.0

TX Testing and RX Testing

DSAZ534A Z-Series Real Time Oscilloscope
- Bandwidth up to 63GHz
- Low noise

M8040A High Performance 64Gbaud BERT
- Data rates up to 64GBaud
- NRZ and PAM4 capable
- 4 tap de-emphasis
- Integrated Jitter injection

KEYSIGHT TECHNOLOGIES

Don’t miss the Demo

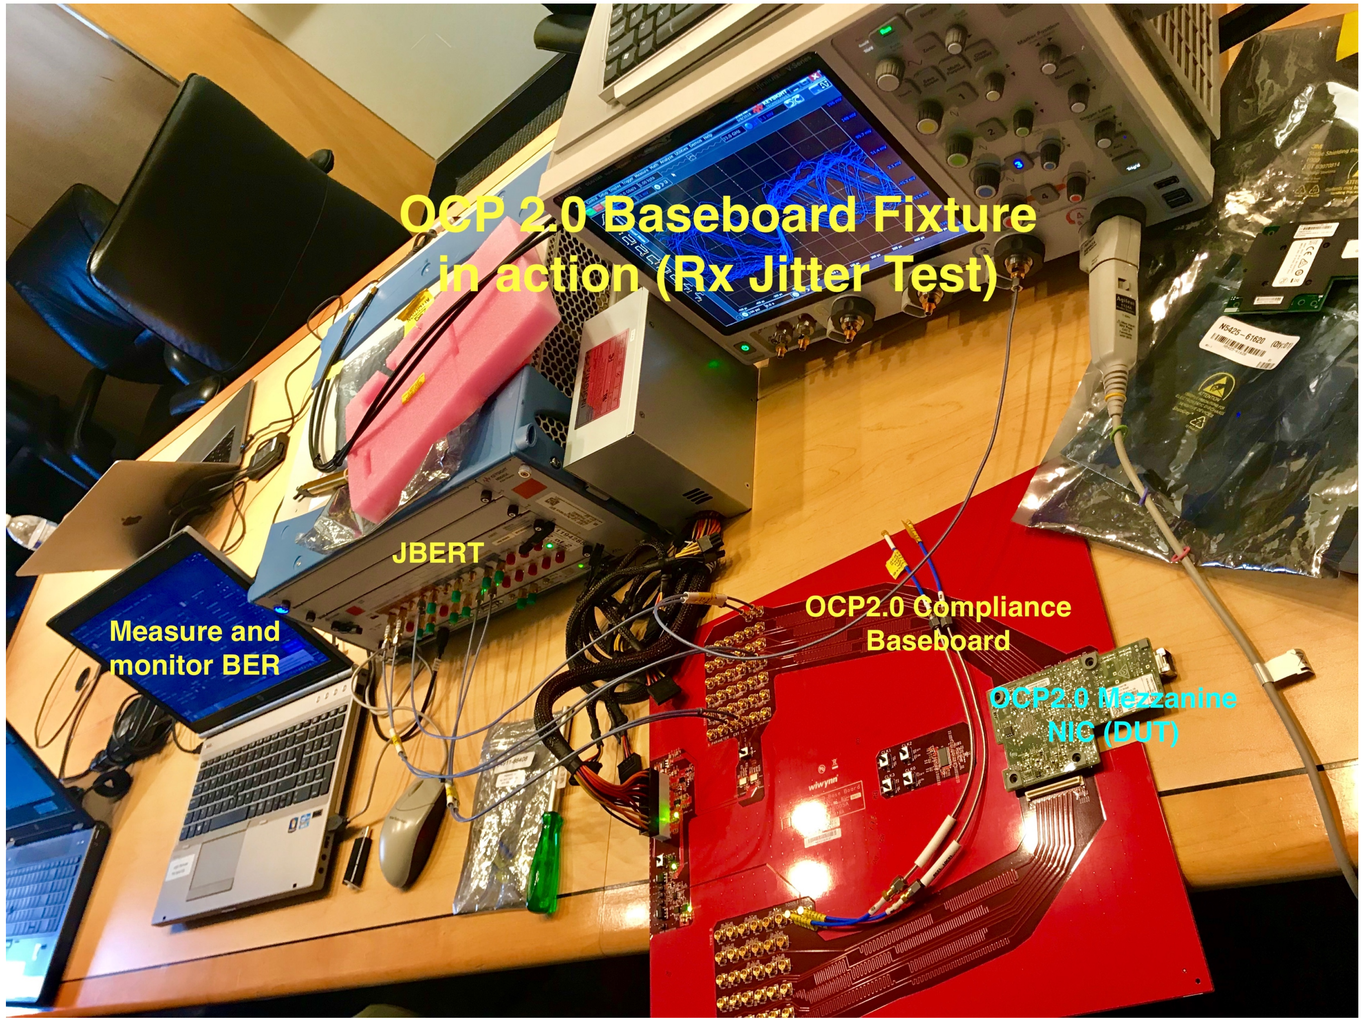

KEYSIGHT TECHNOLOGIES

OCP SUMMIT