SNIA SSSI PCIe SSD Committee
09DEC2013

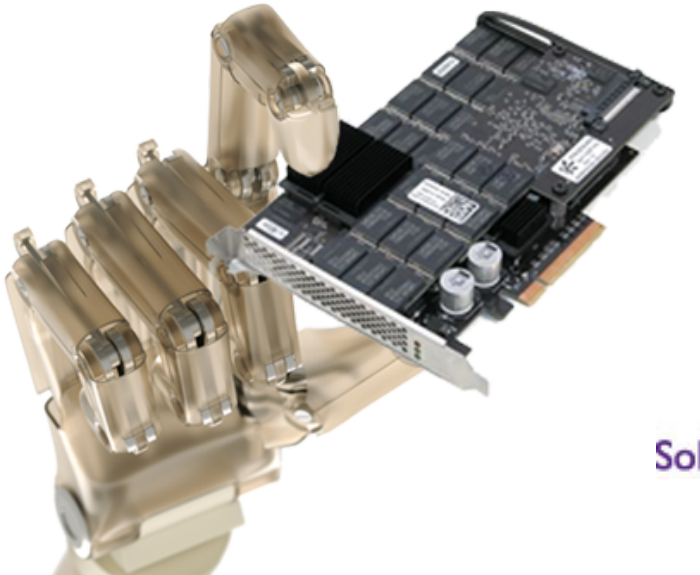

Toll-Free Access Number: 1-866-439-4480
Toll Access Number: 1-212-786-7191
Participant PIN Code: 57236696#

Webex: snia.webex.com
Mtg no: 797 289 257
Password: sssipcie

SNIA
Solid State Storage Initiative

Agenda

SNIA
Advancing storage & information technology

1.	Administrative	Roll Call; Minutes; Announcements	4:00 – 4:10
2.	NVMe	PCIe Protocol & Analysis tools	4:10 – 5:00
3.	Q & A	Open	5:00 – 5:15
4.	Close	Opens / AI	5:15 – 5:30

Time: 4:00 – 4:10

Administrative

SNIA
Advancing storage & information technology

Roll Call:

Eden Kim, Chair (Calypso)

John Wiedemeier (Teledynelecroy), Purush Gupta (Nutanix), Paul Wassenberg (Marvell), Don Deel (EMC), Dave Hiatt (SSSI), Mike Fitzpatrick (TAIS), Tom West (HyperIO),

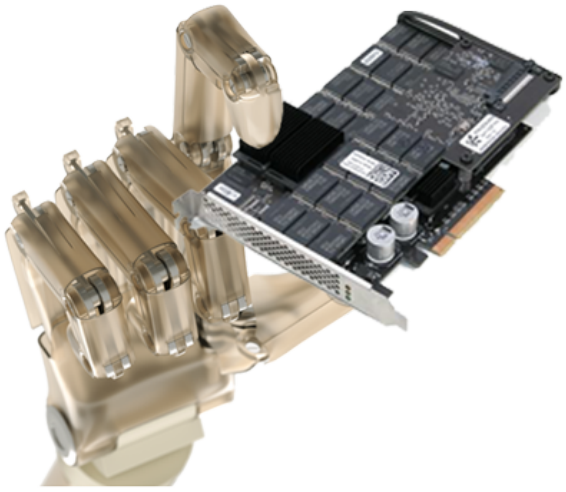

Time: 4:00 – 4:10

Administrative

SNIA
Advancing storage & information technology

A. Meeting Schedule:

Once per month through 2013

Next Mtg: Feb 10 , 2014 – TBD

B. SSSI Work Items – Update from Paul Wassenberg

C. Upcoming Events / Activities

1. Storage Visions - Las Vegas, Jan. 5-6, 2013 www.storagevisions.com
2. SNIA SNW Winter Symposium – San Jose, CA Jan. 28-31

Time: 4:10 – 4:15

PCIe Protocol and Analysis Tools

John Wiedemier
Teledynelecroy

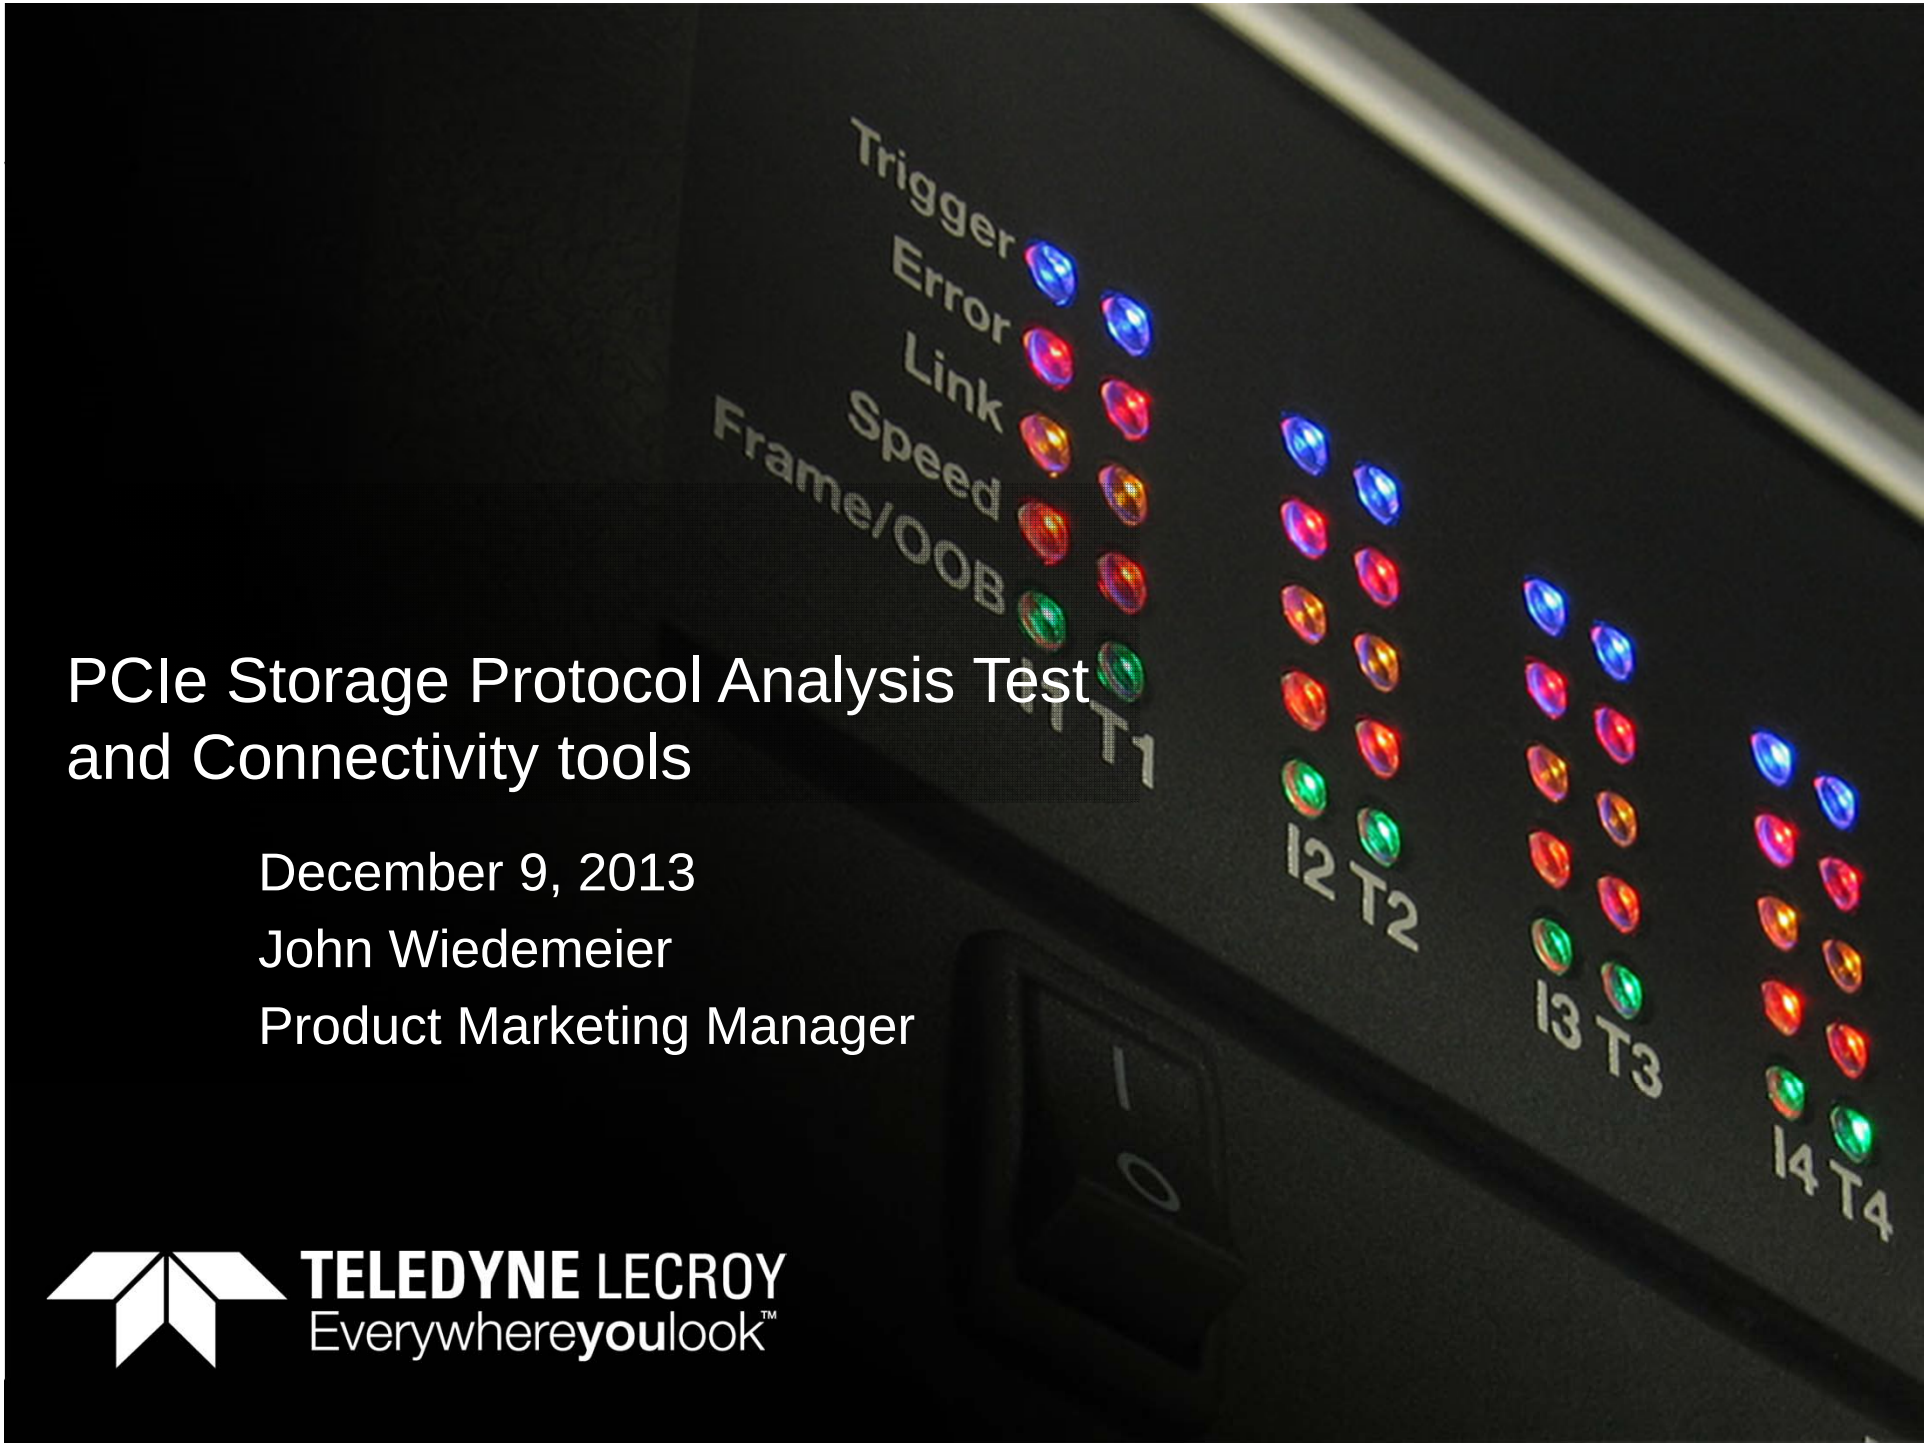

PCIe Storage Protocol Analysis Test and Connectivity tools

December 9, 2013

John Wiedemeier

Product Marketing Manager

TELEDYNE LECROY
Everywhereyoulook™

Teledyne LeCroy

- Teledyne LeCroy Corporation is the leader in Serial Data Test solutions
  - Founded in 1964
  - Headquarters in Chestnut Ridge, NY
- Teledyne LeCroy's Protocol Solutions Group (PSG) was formed in October 2004 with the acquisition of CATC Corporation
  - PSG Headquarters in Santa Clara, CA
  - Part of the Serial Data Division of LeCroy
  - Catalyst was acquired & added to PSG in 2006
- PSG specializes in providing complete protocol solutions for a wide range of serial data standards
  - Products range from production tools to full protocol analysis systems with intuitive user interfaces and complete traffic generation

PCIe SSDs Require New Testing Methodologies

- PCIe SSDs combine various storage protocol layers on top of the PCI Express electrical, link and transaction layers.
- The PCIe protocol is very different than previous storage protocols.
- New tools and testing methodologies are required to meet the challenges of high performance SSDs.

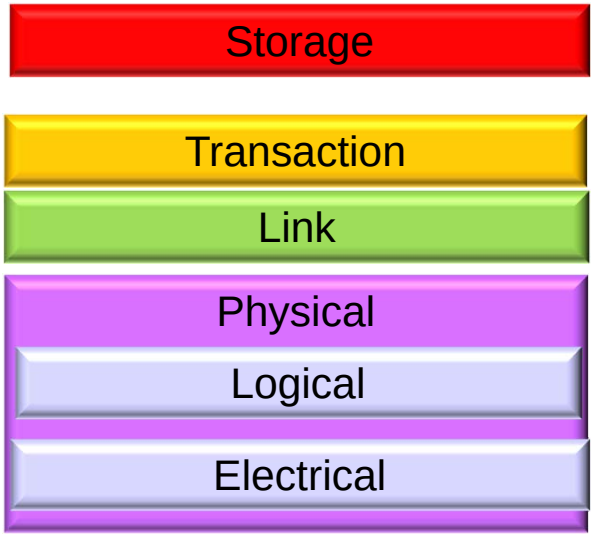

Testing Storage Interfaces

Test Points of Interest

1. System board and HBA
2. HBA and Storage Media
3. System board and PCIe Storage Media

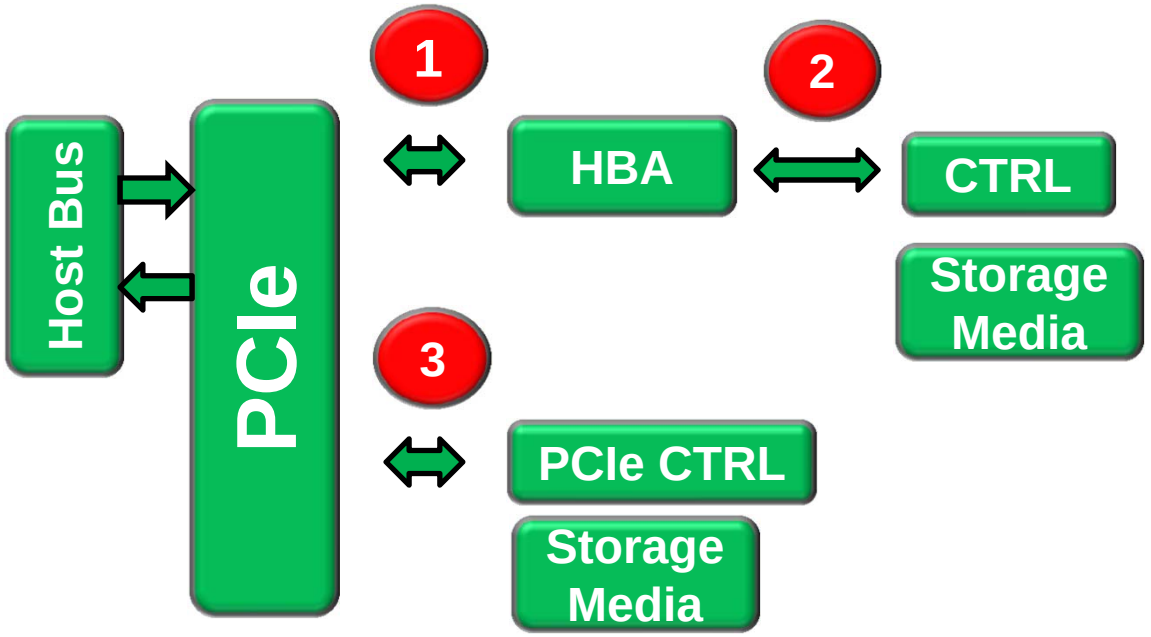

PCI Express Performance Tools

- **Measure and monitor PCIe SSD performance details**
  - PCIe Packet Metrics
  - Timing Calculator- Bandwidth, Link Utilization
  - Post Capture Bus Utilization Graph Tool
    - Latency, Throughput Views
  - Real Time Statistics Graph Tool

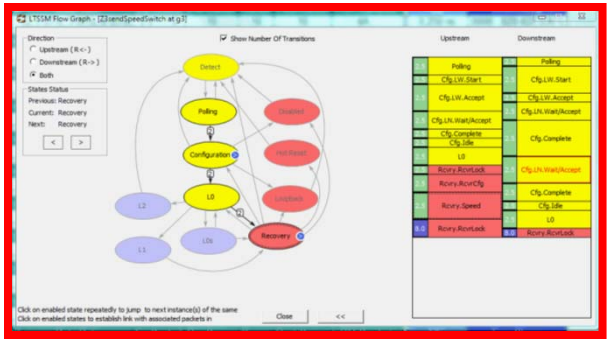

- **Understand link behaviors and improve SSD performance**
  - Flow Control View
  - Link Tracker
  - LTSSM State View

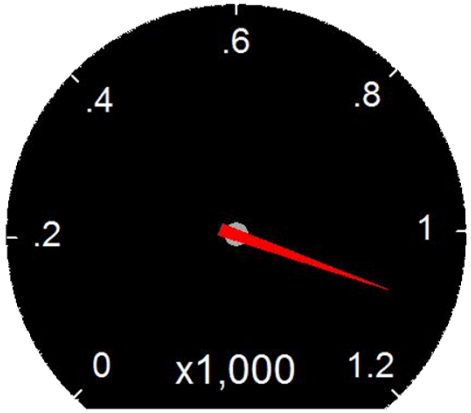

Bus Utilization	Upstream	Downstream
Link Utilization	45.127 %	44.546 %
Time Coverage	45.072 %	44.493 %
Bandwidth	9025.43 Mb/s	8909.10 Mb/s
Data Throughput	592.29 MB/s	598.72 MB/s
Packets/second	19696168.45	24833612.81

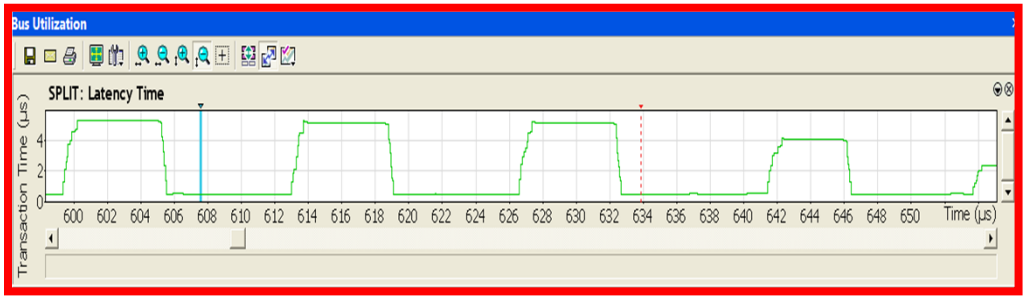

TELEDYNE LECROY
Everywhereyoulook™

PCI Express Protocol Analyzers

- **Summit T3-8**
  - PCIe 3.0
  - Speed: 8GT/s,
  - Lane width x1, x2, x4, x8

- **Summit T28**
  - PCIe 2.0
  - Speed: 5GT/s,
  - Lane width x1, x2, x4, x8

- **Summit T24**
  - PCIe 2.0
  - Speed: 5GT/s,
  - Lane width x1, x2, x4

SATA Express Decoded Trace

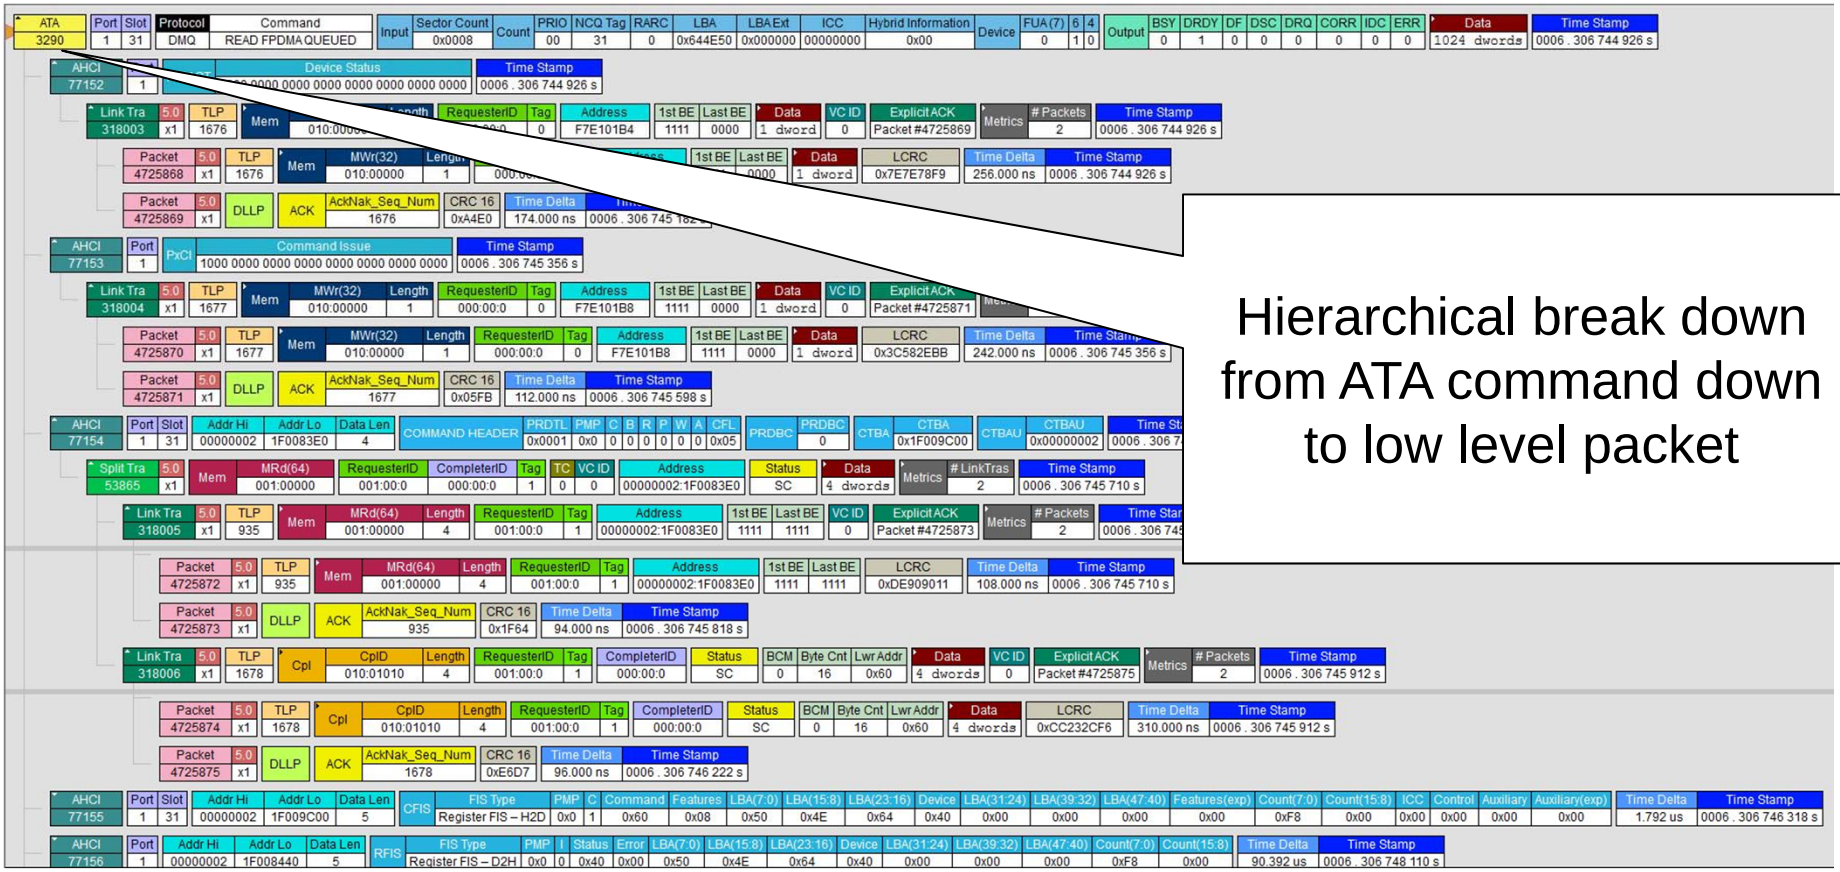

Hierarchical break down from ATA command down to low level packet

NVMe Decoded Trace

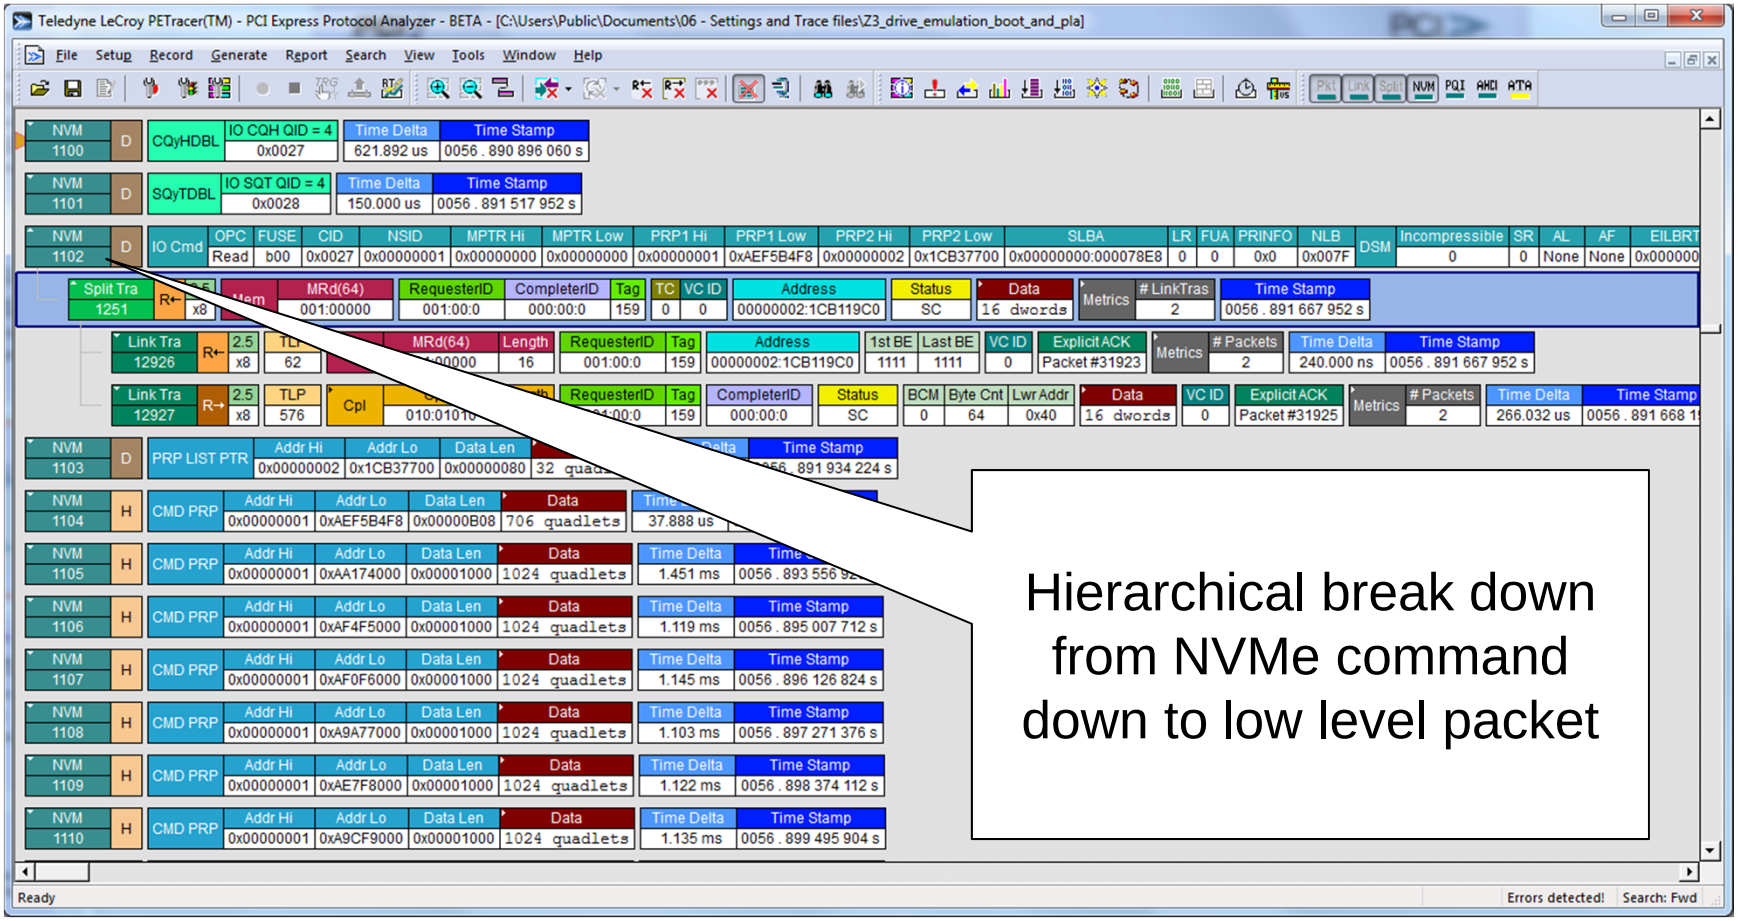

TELEDYNE LECROY
Everywhereyoulook™

SCSI Express Decoded Trace

PQI 3	D	ADMIN Queue Parameter	ADMIN IQ ELEMENTS	ADMIN OQ ELEMENTS	INTERRUPT MESSAGE	RsvdZ	Time Delta	Time Stamp		
			0x40	0x00	0x0000	0x00	9.352 us	0140 . 033 798 556 s		
PQI 4	D	CAPABILITY	MAX ADMIN IQ	MAX ADMIN OQ	AIQE LEN	AOQE LEN	MAX TO	Time Delta	Time Stamp	
			0x40	0x40	0x40	0x40	0x0000	9.368 us	0140 . 033 807 908 s	
PQI 5	D	CAPABILITY	MAX ADMIN IQ	MAX ADMIN OQ	AIQE LEN	AOQE LEN	MAX TO	Time Delta	Time Stamp	
			0x40	0x40	0x40	0x40	0x0000	9.464 us	0140 . 033 817 276 s	
PQI 6	H	ADMIN Queue Parameter	ADMIN IQ ELEMENTS	ADMIN OQ ELEMENTS	INTERRUPT MESSAGE	RsvdZ	Time Delta	Time Stamp		
			0x40	0x40	0x0000	0x00	176.000 ns	0140 . 033 826 740 s		
PQI 7	H	ADMIN IQ	RsvdZ	Offset Hi	Offset Lo	Time Delta	Time Stamp			
			0x00	0x00000000	0xBF709000	352.000 ns	0140 . 033 826 916 s			
PQI 8	H	ADMIN IQ CI	RsvdZ	Offset Hi	Offset Lo	Time Delta	Time Stamp			
			0x10	0x00000000	0xBF708190	304.000 ns	0140 . 033 827 268 s			
PQI 9	H	ADMIN OQ	RsvdZ	Offset Hi	Offset Lo	Time Delta	Time Stamp			
			0x00	0x00000000	0xBF70A000	328.000 ns	0140 . 033 827 572 s			
PQI 10	H	ADMIN OQ PI	RsvdZ	Offset Hi	Offset Lo	Time Delta	Time Stamp			
			0x30	0x00000000	0xBF7081B0	256.000 ns	0140 . 033 827 900 s			
PQI 11	D	Time Delta	Time Stamp							
		65.752 us	0140 . 033 828 156 s							
PQI 12	H	Time Delta	Time Stamp							
		1.001 ms	0140 . 033 893 908 s							
PQI 13	D	Time Delta	Time Stamp							
		11.048 us	0140 . 034 894 652 s							
PQI 14	D	ADMIN IQ PI	Offset Hi	Offset Lo	Time Delta	Time Stamp				
			0x00000000	0x00001000	18.856 us	0140 . 034 905 700 s				
PQI 15	D	ADMIN OQ CI	Offset Hi	Offset Lo	Time Delta	Time Stamp				
			0x00000000	0x00001008	42.200 us	0140 . 034 924 556 s				
PQI 16	H	IQ PI	Index	Time Delta	Time Stamp					
			0x0001	18.104 us	0140 . 034 966 756 s					

TELEDYNE LECROY
Everywhereyoulook™

Form Factor/Connectivity and Probing

- **CEM add-in card**
  - NVMe/SATA/SCSI
  - PCIe 2.0/3.0
  - x1, x4, x8
- **M.2 storage card** (i.e., NGFF)
  - NVMe/SATA
  - PCIe 2.0/3.0
  - X2, x4
- **SFF-8639 based drive**
  - NVMe/SATA/SCSI
  - PCIe 2.0/3.0
  - x2, x4

Interposer for CEM add-in card

Interposer for M.2 storage card

Interposer for SFF-8639 based drive

PCIe Single Port Usage Model

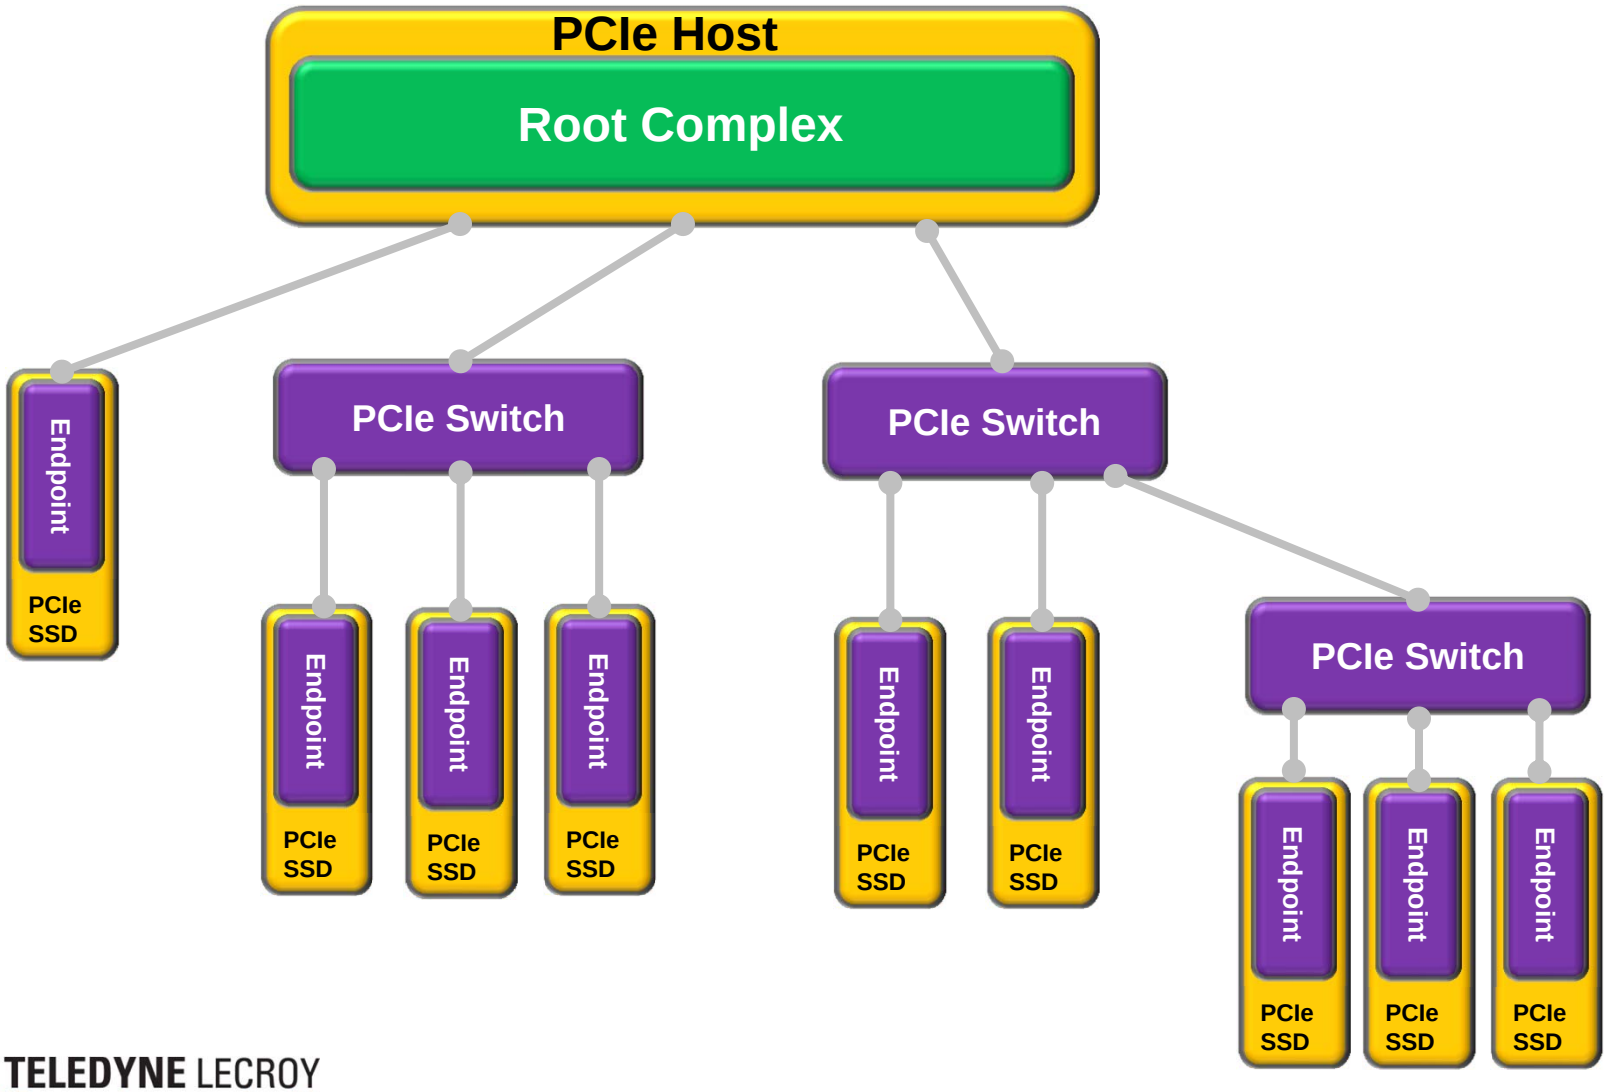

TELEDYNE LECROY
Everywhereyoulook™

Protocol Analysis with the Single Port SFF-8639 Interposer

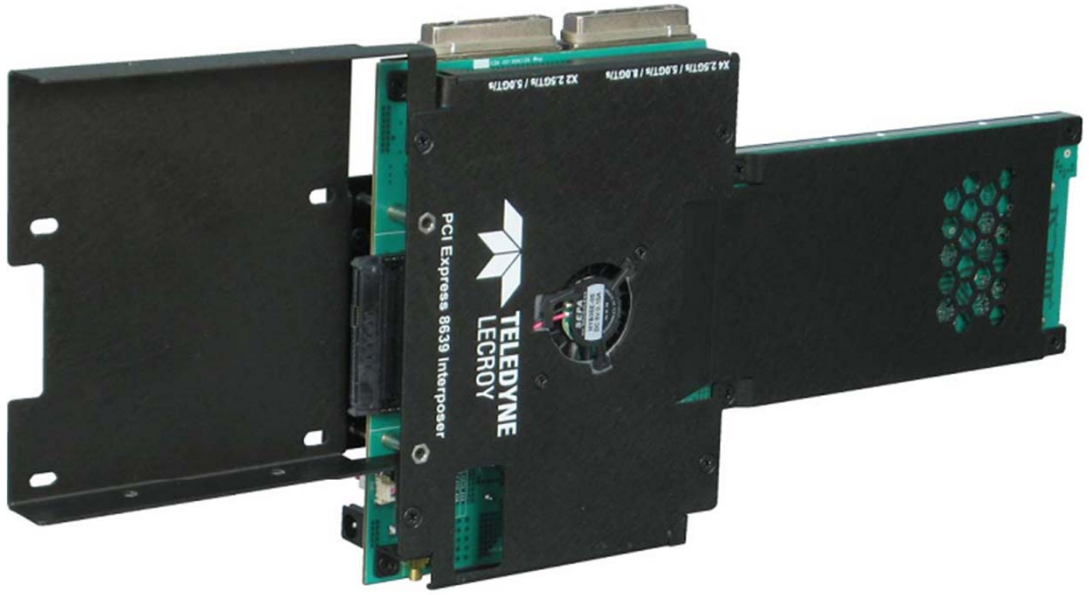

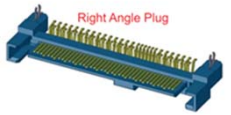

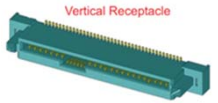

SFF-8639 Connector

Supports

- NVM Express G3x4
- SATA Express G3x2
- SCSI Express G3x4

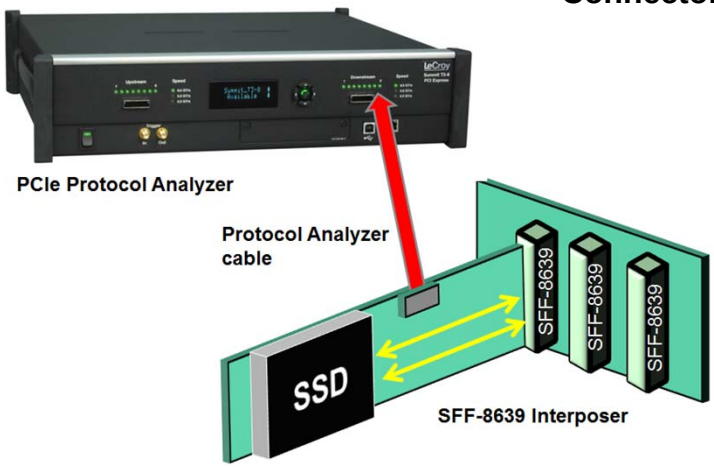

TELEDYNE LECROY
Everywhereyoulook™

SFF-8639 to PCIe Connector Adaptor

Supports Single Port Drives

- PCIe 1/2/3
- x1, x2, x4
- 2 ½ inch drives

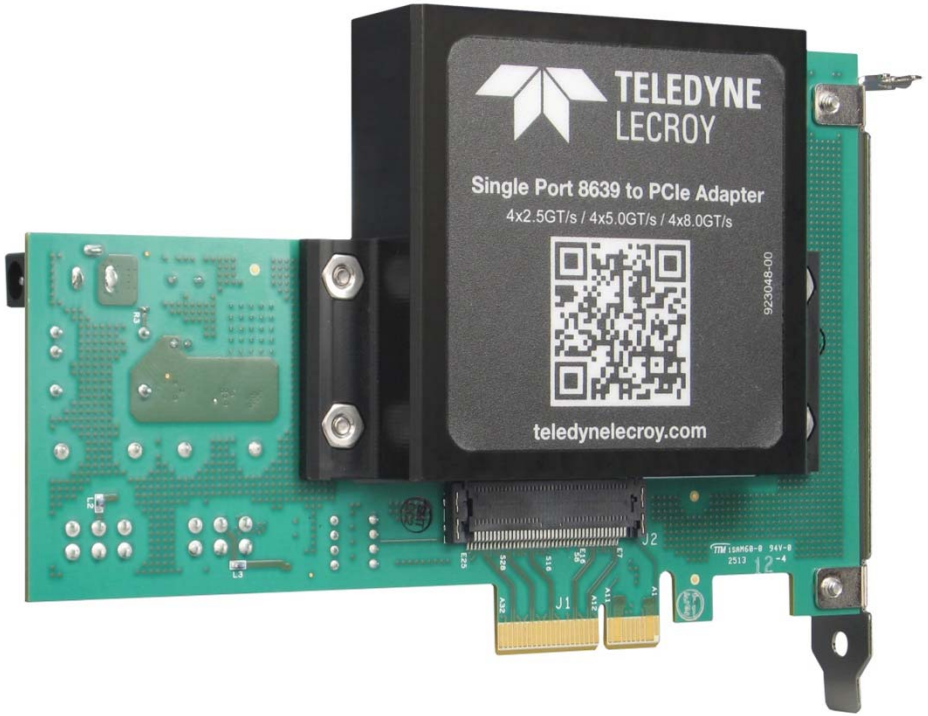

The answer to companies looking for 8639 backplanes for test

TELEDYNE LECROY
Everywhereyoulook™

Dual Port PCIe Usage Model

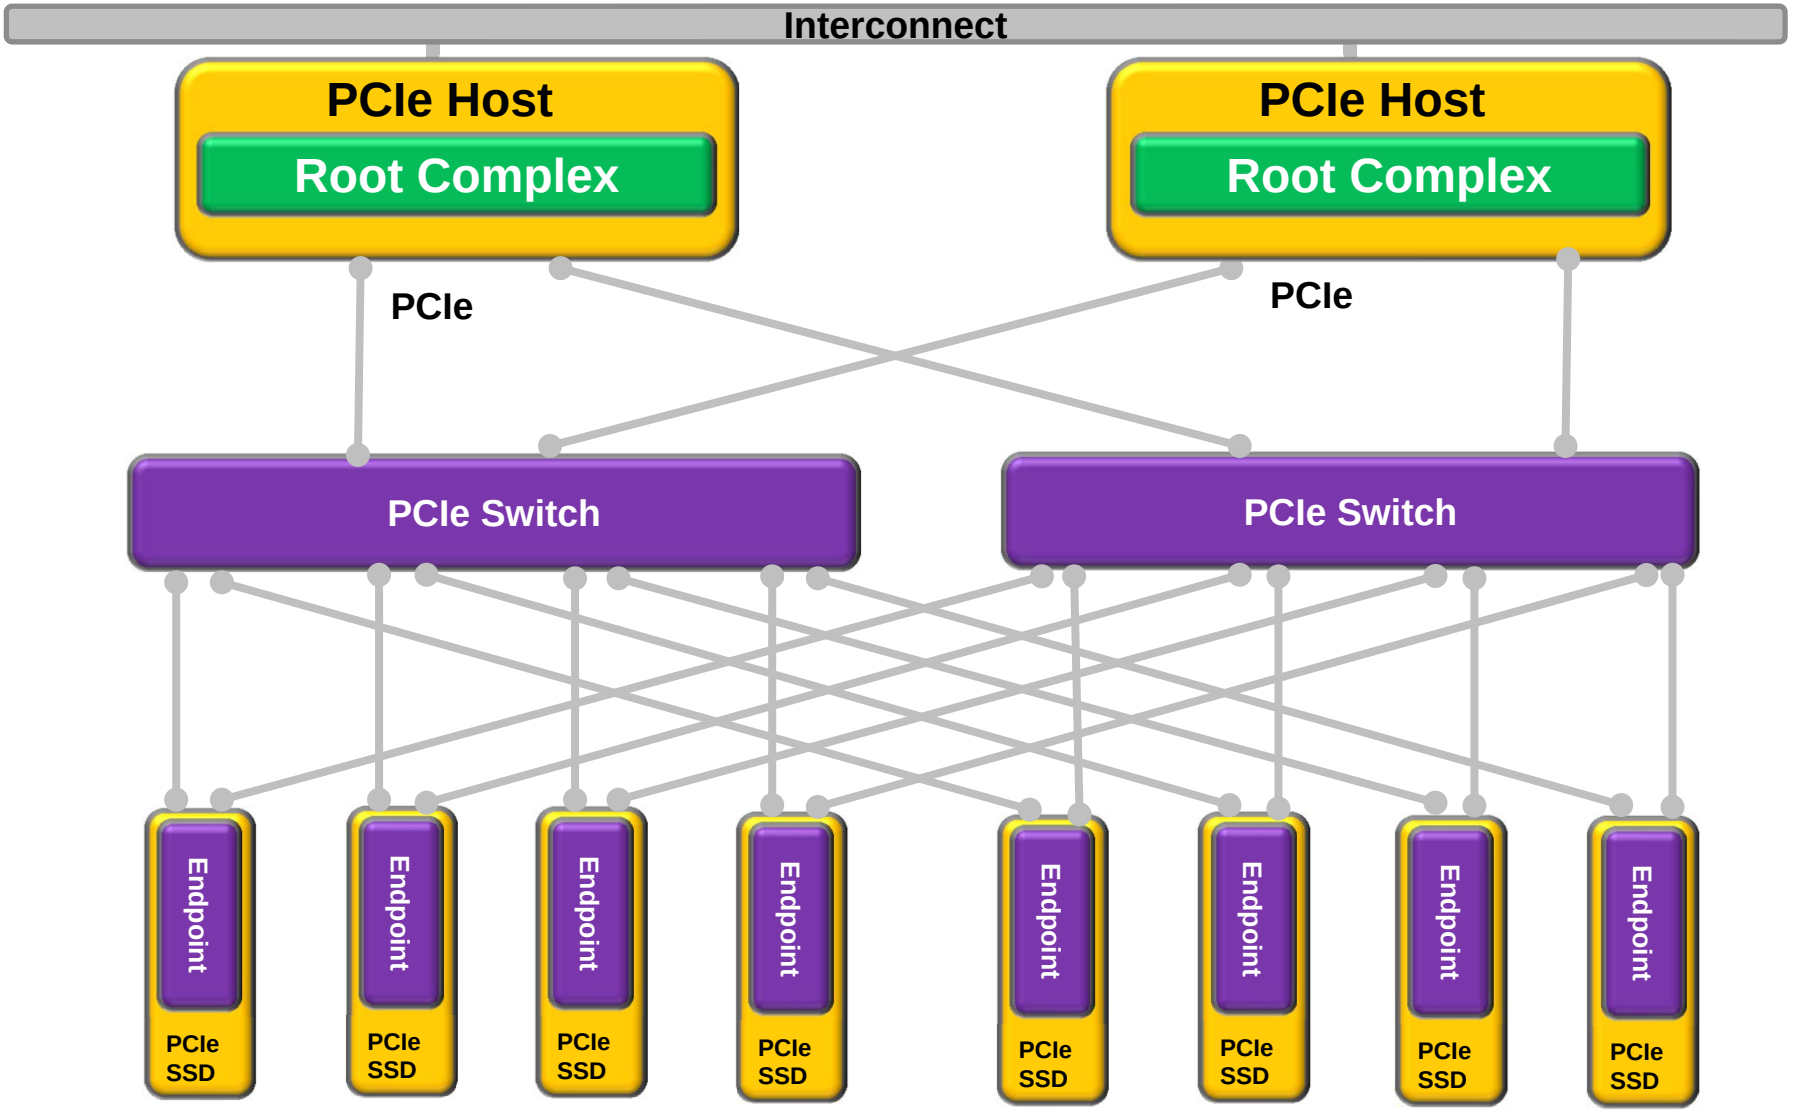

TELEDYNE LECROY
Everywhereyoulook™

Dual Port SFF-8639 Interposer

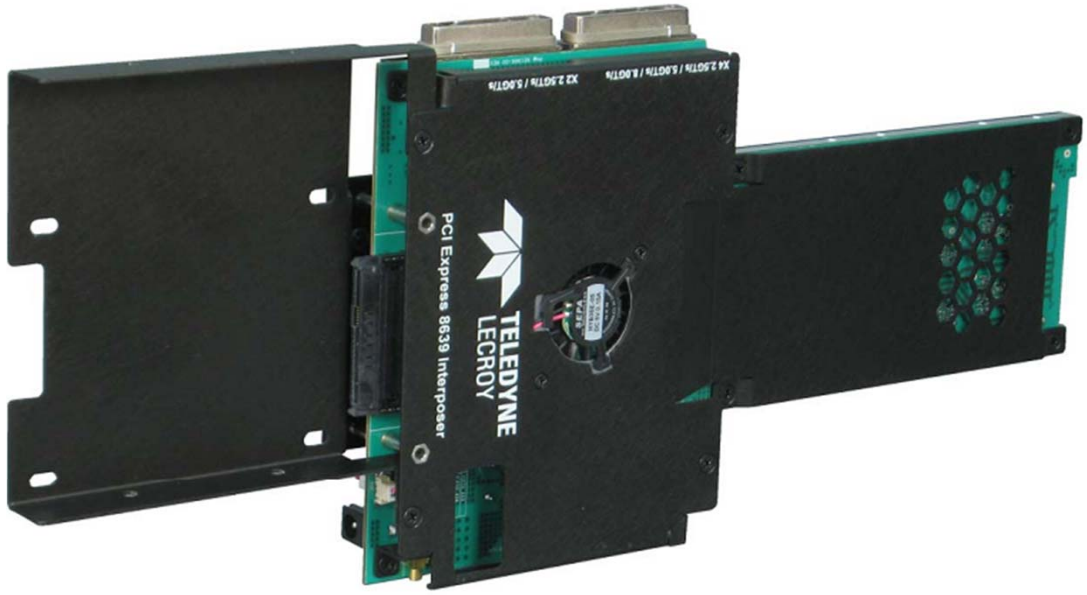

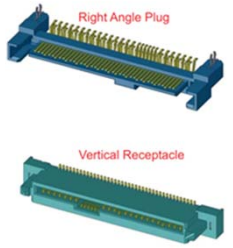

SFF-8639 Connector

Supports

- Gen3 Supported
- Host Interface Support
  - NVM Express Dual x2
  - SCSI Express Dual x2
- 2 PCIe link A/B port support
- 2 ½ and 3 inch drive support

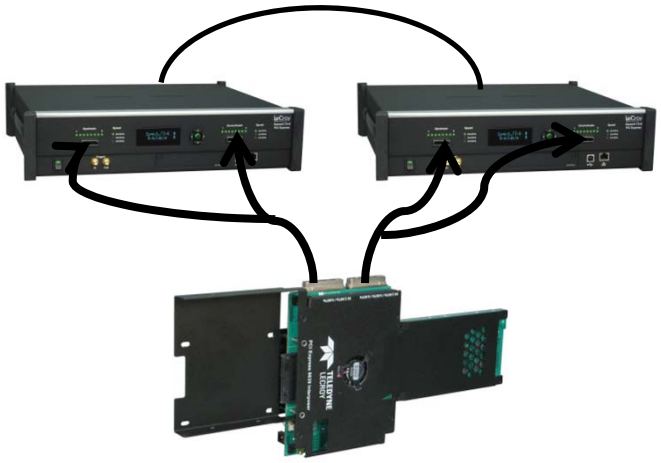

TELEDYNE LECROY
Everywhereyoulook™

SFF-8639 to PCIe Connector Adaptor

Supports Port B

- PCIe 1/2/3
- x1, x2, x4
- 2 ½ inch drives

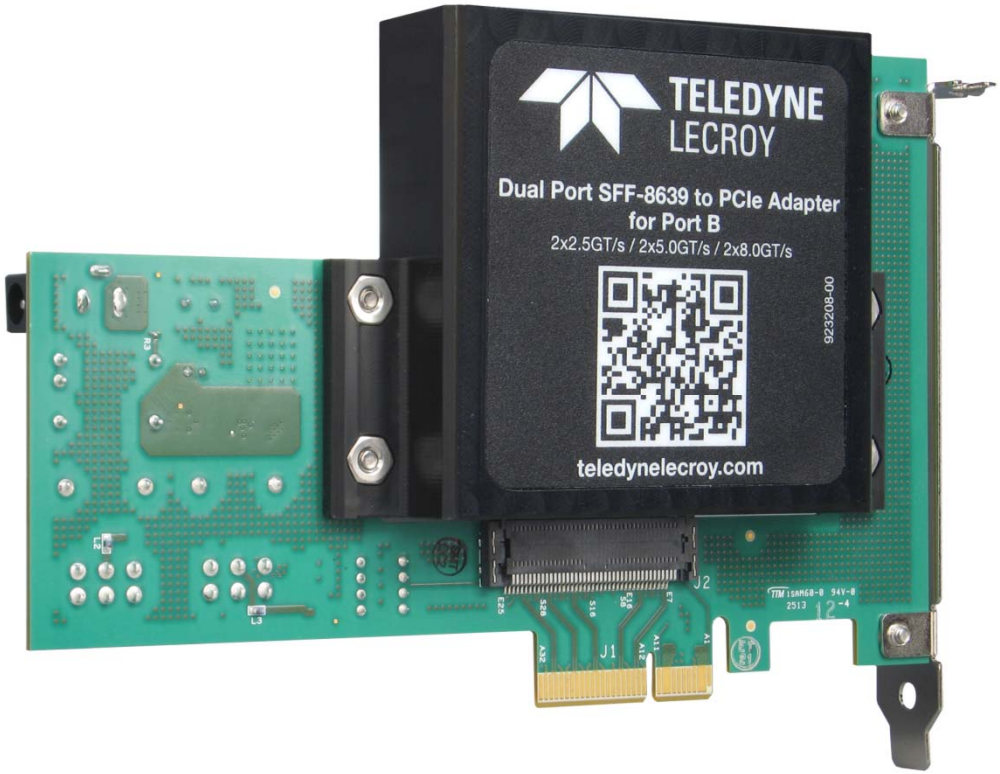

The answer to companies looking for 8639 backplanes for test

TELEDYNE LECROY
Everywhereyoulook™

SFF-8639 to PCIe Connector Adaptor Usage

- System Prototyping
- SSD Testing

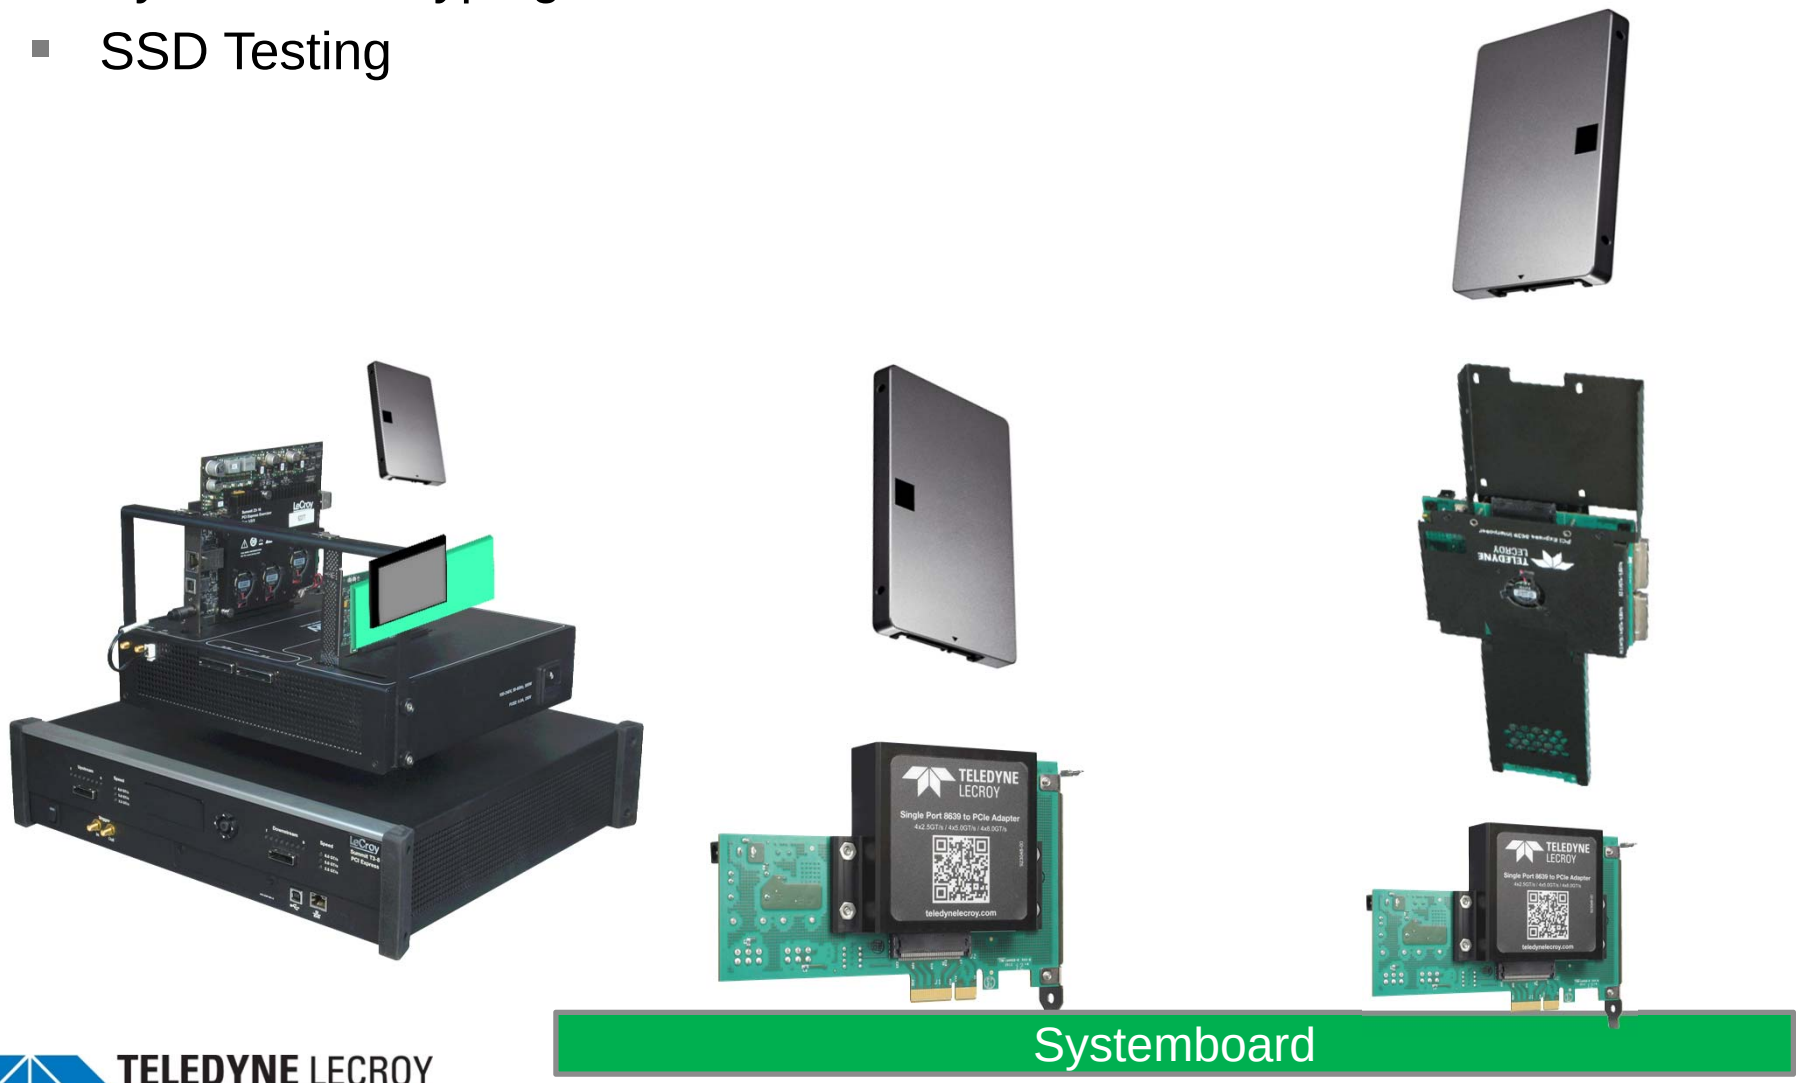

TELEDYNE LECROY
Everywhereyoulook™

M.2 Interposer for PCIe 3.0

- Speeds: PCIe 1/2/3
- Lane width: x2/x4
- Socket types: 2/3
- M.2 module lengths: 42mm x 22mm, 60mm x 22mm, 80mm x 22mm , and 110mm x 22mm
- SRIS supported
- Dimensions
  - Connector cable 18 inches
  - Carrier Board 5.13 x 9.05 inches
  - Interposer Board .824 x 1.65 inches

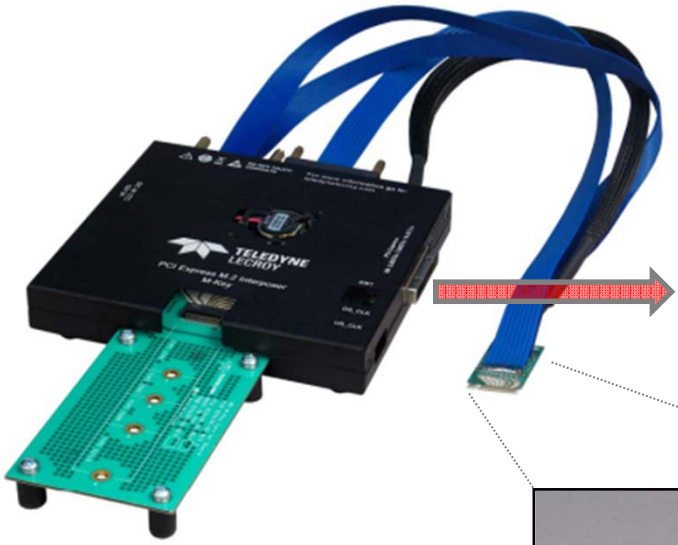

Protocol Analyzer

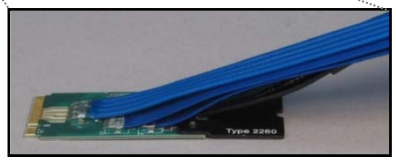

M.2 Interposer

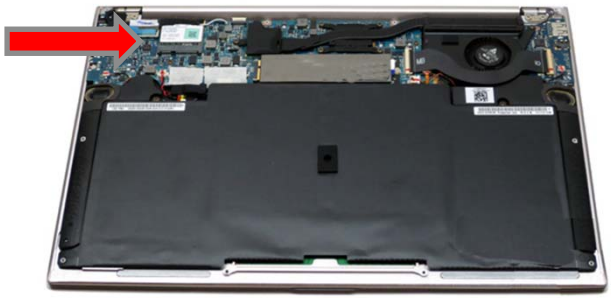

M.2 SSD

Laptop

TELEDYNE LECROY
Everywhereyoulook™

M.2 to PCIe Connector Adaptors

Two types of Interposer

- Socket 2 devices
  - B-type
  - PCIe 1/2/3
  - x1, x2
- Socket 3 devices
  - M-type
  - PCIe 1/2/3
  - X1, x2, x4

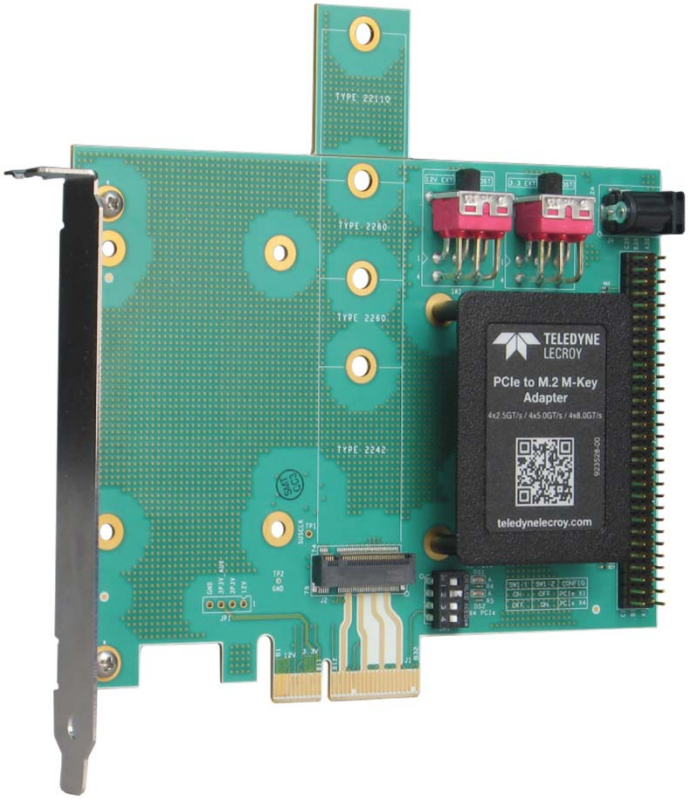

TELEDYNE LECROY
Everywhereyoulook™

Testing PCIe SSD and Systems

- Exercisers are the best way to test error handling between systems and storage devices in PCI Express.
- Emulation Scripts are run on exercisers and create low level bus traffic that can be accurately, repeatedly, and precisely created to assess performance in real-world conditions for a drive or host. No software driver required.

Summit Z3-16 Protocol Exerciser

SSD Test Setup

SSD Device Controller Test

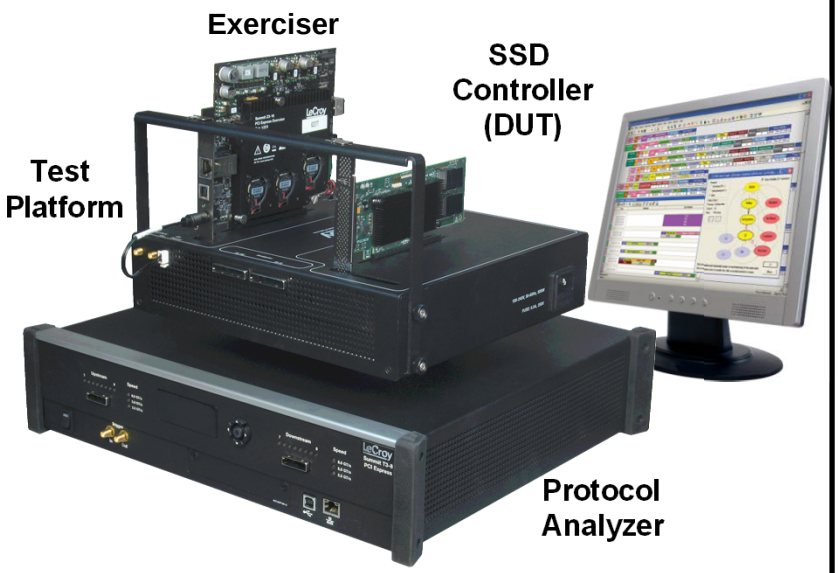

SSD Driver Host Test

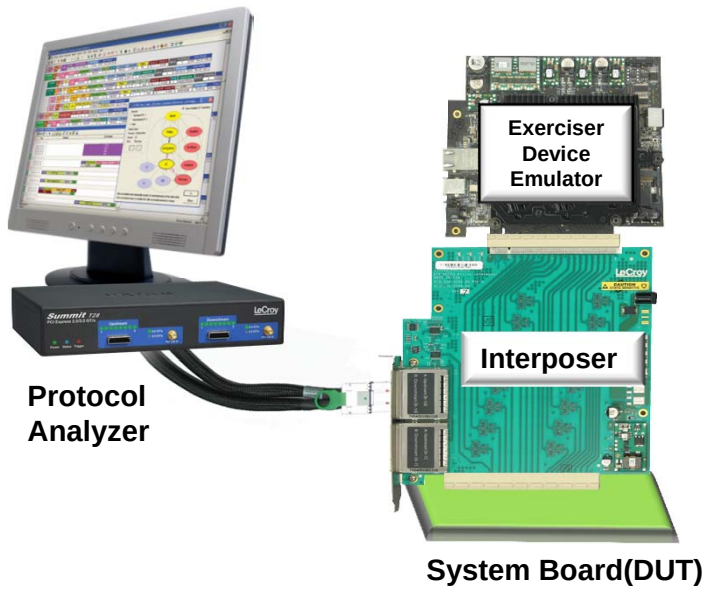

NVMe Conformance Testing

- Tests defined in NVMe Interop Test Suite document (publically available at www.iol.unh.edu)
- Test tools used to check proper construction and response to different NVMe Stimuli
- Admin Commands, NVM Commands, Controller Registers

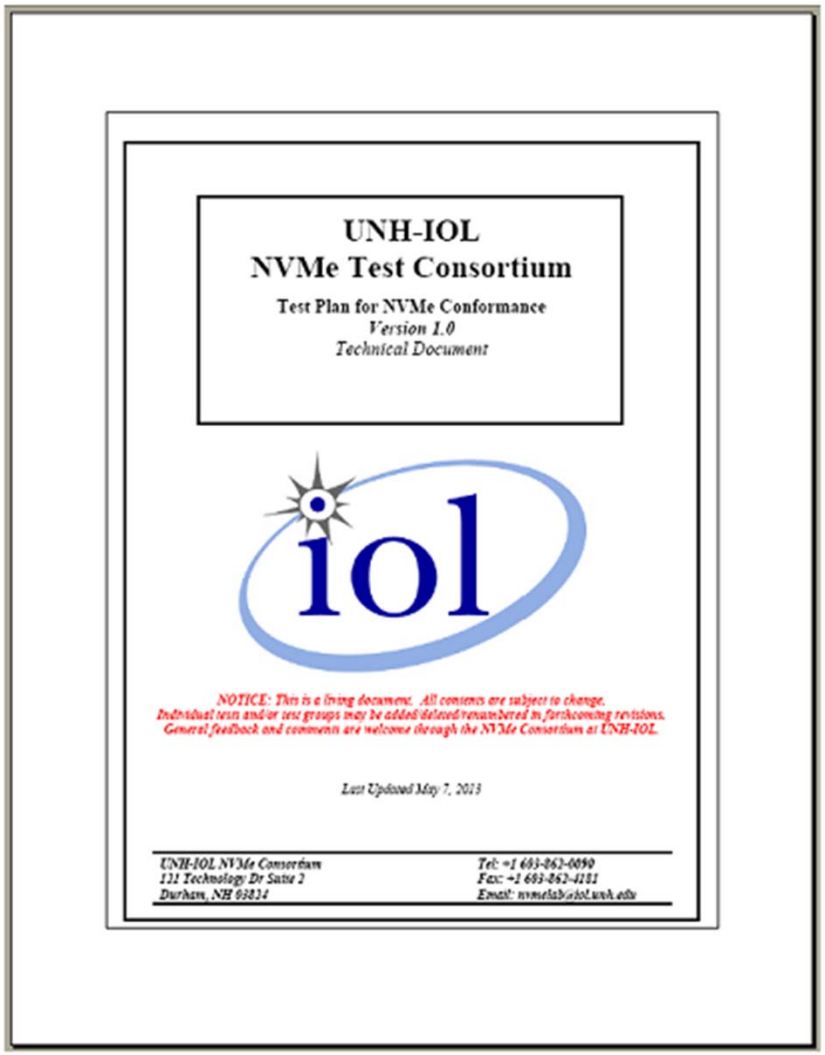

UNH-IOL
NVMe Test Consortium

Test Plan for NVMe Conformance
Version 1.0
Technical Document

iol

NOTICE: This is a living document. All contents are subject to change. Individual tests and/or test groups may be added/deleted/renumbered in forthcoming revisions. General feedback and comments are welcome through the NVMe Consortium at UNH-IOL.

Last Updated May 7, 2013

UNH-IOL NVMe Consortium
121 Technology Dr Suite 2
Durham, NH 03824

Tel: +1 603-862-0090
Fax: +1 603-862-4181
Email: nvmelab@iol.unh.edu

TELEDYNE LECROY
Everywhereyoulook™

NVMe Integrators List

- Hosted by UNH-IOL at ww.iol.unh.edu
- Opt-in list of qualifying NVMe products
- No PCIe component to qualification today
- UNH-IOL does offer PCIe testing to complement NVMe testing, but it is not a component of the NVMe IL qualification
- NVMe Host Qualification
  - Perform Interop Test against 4 different SSDs
  - Pass with 3 SSDs
  - Pass = data transfers without errors
- NVMe Device Qualification
  - Perform Interop Test against 4 different hosts, pass with 3 Hosts
  - One of the hosts must be either the Windows or Linux Reference Driver
  - Pass all conformance tests

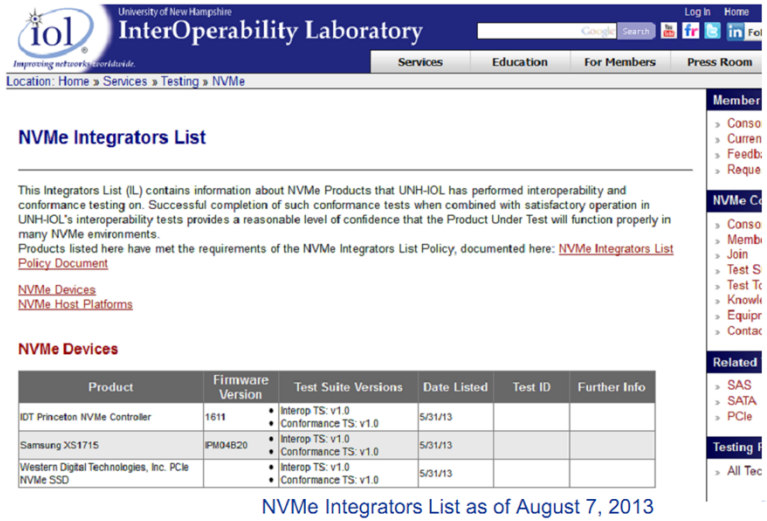
University of New Hampshire InterOperability Laboratory

Services | Education | For Members | Press Room

Location: Home » Services » Testing » NVMe

NVMe Integrators List

This Integrators List (IL) contains information about NVMe Products that UNH-IOL has performed interoperability and conformance testing on. Successful completion of such conformance tests when combined with satisfactory operation in UNH-IOL's interoperability tests provides a reasonable level of confidence that the Product Under Test will function properly in many NVMe environments.
Products listed here have met the requirements of the NVMe Integrators List Policy, documented here: NVMe Integrators List Policy Document

NVMe Devices
NVMe Host Platforms

NVMe Devices

Product	Firmware Version	Test Suite Versions	Date Listed	Test ID	Further Info
IDT Princeton NVMe Controller	1611	• Interop TS: v1.0 • Conformance TS: v1.0	5/31/13		
Samsung XS1715	IPM04B20	• Interop TS: v1.0 • Conformance TS: v1.0	5/31/13		
Western Digital Technologies, Inc. PCIe NVMe SSD		• Interop TS: v1.0 • Conformance TS: v1.0	5/31/13		

NVMe Integrators List as of August 7, 2013

NVMe Compliance and Interoperability Testing

- UNH-IOL(University Of New Hampshire IOL) and the NVM Express Promoters Group are collaborating to create an interop and conformance test program centered at UNH-IOL.
- They have created a multi-vendor test bed to help products prove interoperability, and conformance test services to prove that products follow the NVMe specification correctly.
- Feb 24, 2014.

University of New Hampshire InterOperability Laboratory

NVMe SSD Compliance Tester

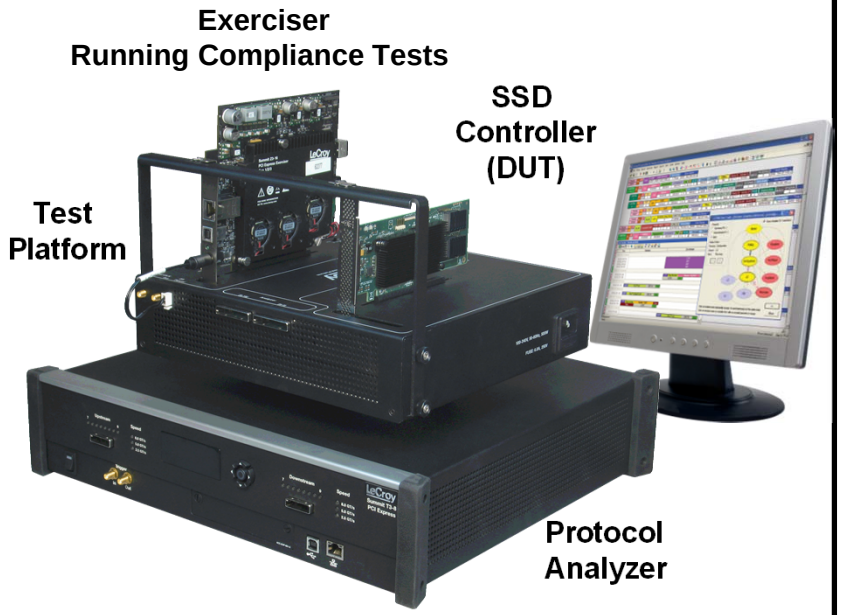

TELEDYNE LECROY
Everywhereyoulook™

Conformance Testing

- 2 types of test stations used
  - IOL Interact
  - Teledyne LeCroy Summit

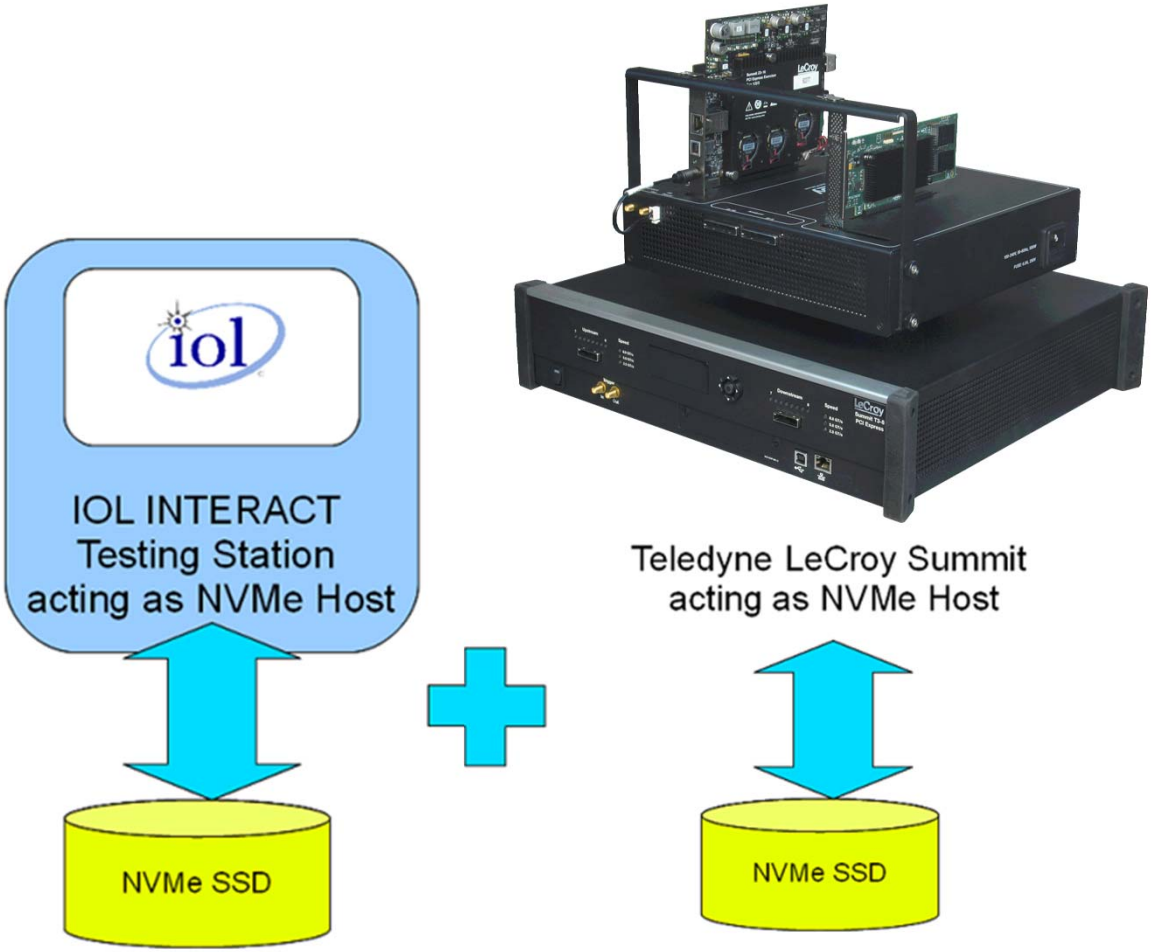

Conclusion

- New testing methodologies are needed for NVM Express, SCSI Express and SATA Express
- Protocol analysis and test tools can show developers useful details about PCIe SSD transactions between the storage host and controller.
- Host and device emulation can discover performance, error handling and protocol issues that affect the quality of products.

TELEDYNE LECROY
Everywhereyoulook™

Contact Teledyne LeCroy PSG

Fax(Sales&Service): 845-578-5985

Email
Sales: psgsales@teledynelecroy.com

Phone Support: 1-800-553-2769

Find us on Facebook

www.teledynelecroy.com

Close

SNIA
Advancing storage & information technology

A. Opens?
   - item

B. Discussion:
   - Item

C. Close

Time: 5:15 – 5:30