

SFF-8639 PCIe* SSD Ecosystem Readiness and Electrical Testing Update

Don Verner: Sr. Application Engineer, Intel NVM Solutions Group

Contributors:

Jonmichael Hands: Technical Program Manager, Intel SSDs

Dan Froelich: Sr. Staff I/O Architect, Intel Data Center Group

Intel Corporation



Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products

No computer system can provide absolute security. Requires an enabled Intel® processor, enabled chipset, firmware and/or software optimized to use the technologies. Consult your system manufacturer and/or software vendor for more information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to:
<http://www.intel.com/design/literature.htm>

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.


*Other names and brands may be claimed as the property of others.

Copyright © 2014 Intel Corporation. All rights reserved.

SSD Drive Connector


Current SATA Connector

- Uses legacy SATA pin pitch
- Keyed to preclude the insertion of a non-SATA drive




Current SAS Connector

- Added additional signaling pins for a secondary port option at with a tighter, modern, pin pitch
- Supports both SATA and SAS drives



SFF 8639 Connector

- Fills out all remaining pin capacity of the legacy form factor
- Designed to support many protocols
- Enterprise mapping supports legacy SATA, SAS, and modern PCIe* drives simultaneously
 - Both single port X4 and dual port X2 drives
 - PCI-SIG Specification now at rev. 0.7




Yellow: PCIe data, reference clock, and side band

SSF 8639 connector expected to meet same CEM electrical requirements as standard PCIe* connector


Basic PCIe* SSD Topologies

1 Connector




1. SFF-8639 connector directly on motherboard
2. Current designs using modified miniSAS HD to SFF-8639 single cable with external power

2 Connector



3. Current designs using miniSAS HD to miniSAS HD to a backplane. Required for hotplug.


3 Connector




Most using miniSAS HD now,
OCuLink cables in future

More Complex PCIe* Topologies

4 Connector: Using link extension device (retimer or switch) CEM x16 AIC lanes bifurcate to four x4 lanes




5 Connector: Using PCIe riser card plus a retimer or switch




miniSAS HD vs. OCuLink*

miniSAS HD internal connectors



OCuLink* Internal



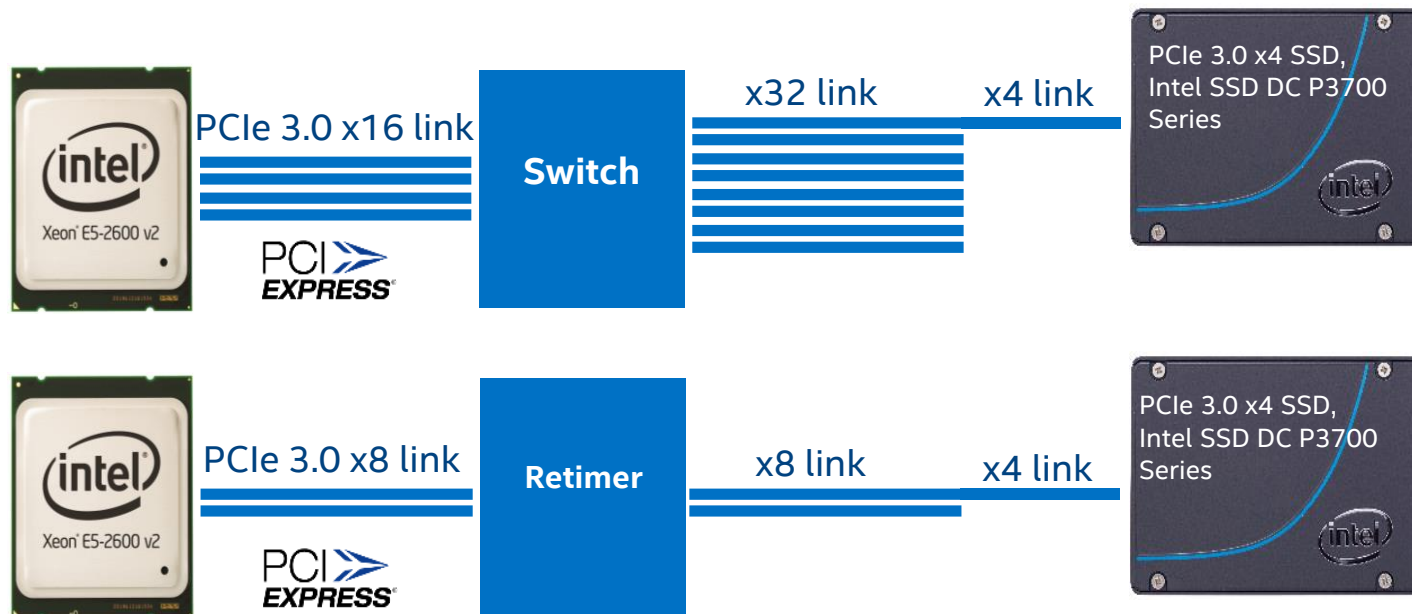
Category	miniSAS HD	OCuLink - Under Development
Standard	Not for PCIe*	Yes
Layout	Adequate footprint	Smaller footprint, easier routing
Signal Integrity	No advantage on loss dominated channels	Better crosstalk
PCIe* 4.0 ready	Made for 12Gbps SAS	16GT/s target
Clock, power	Requires custom cable, external power	Supports clock and 3.3/5V power

NOTE: See PCI Sig spec for additional Information on OCuLink

Link Extension Devices

When required to route PCIe* signals longer than current PCIe Signal Integrity allows, two options exist:

- **PCIe switches** allow for link extension with additional features, more ports
- **Re-timers** are software transparent but engage in PCIe protocol



PCIe* Switch & Retimer Comparison

SWITCHES

- Ease of implementation and hotplug support
- Less BIOS development needed
- Slot configurability
- Multiple PCIe lanes possible

RETIMERS

- Intel® supports the ECN in PCI Sig for retimers
- **Repeater**: A retimer or a Re-driver.
- **Re-driver**: Analog but not protocol aware
 - Not truly transparent to host system
- **Retimer**: Physical Layer protocol aware
 - Software transparent, Extension Device.
 - Forms two separate electrical sub-links

PCIe* 3.0 CEM Add-in Card Electrical Tools

■ Add-in Card Test Procedure

- ✓ CBB (Compliance Base Board) 3.0 Standard Test Fixture with add-in card to test connected environment
- ✓ Lane under test connected through fixture to oscilloscope
- ✓ Add-in card under test enters compliance mode
 - Fixture provides features to select different compliance speeds and de-emphasis levels
- ✓ Data lane sampled
 - 25 ps or smaller sample interval. At least 1 million UI.
- ✓ Standard Post Processing Analysis Software (Sigtest)
 - Supports all common RT Scope data formats
- ✓ Standard Test Procedures for specific test equipment



- Capture waveform on oscilloscope
- Run signal analysis software

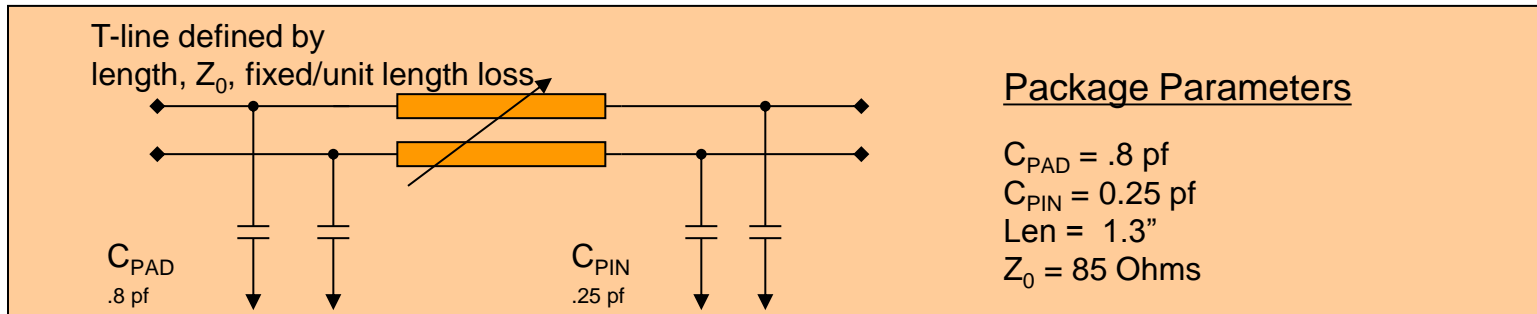
*Other names and brands may be claimed as the property of others.

Intel developing Sigtest, standard fixtures, and reference Tx/Rx test procedures for PCIe* Over SFF-8639 interface

PCIe* SFF-8639 CBB Design Changes

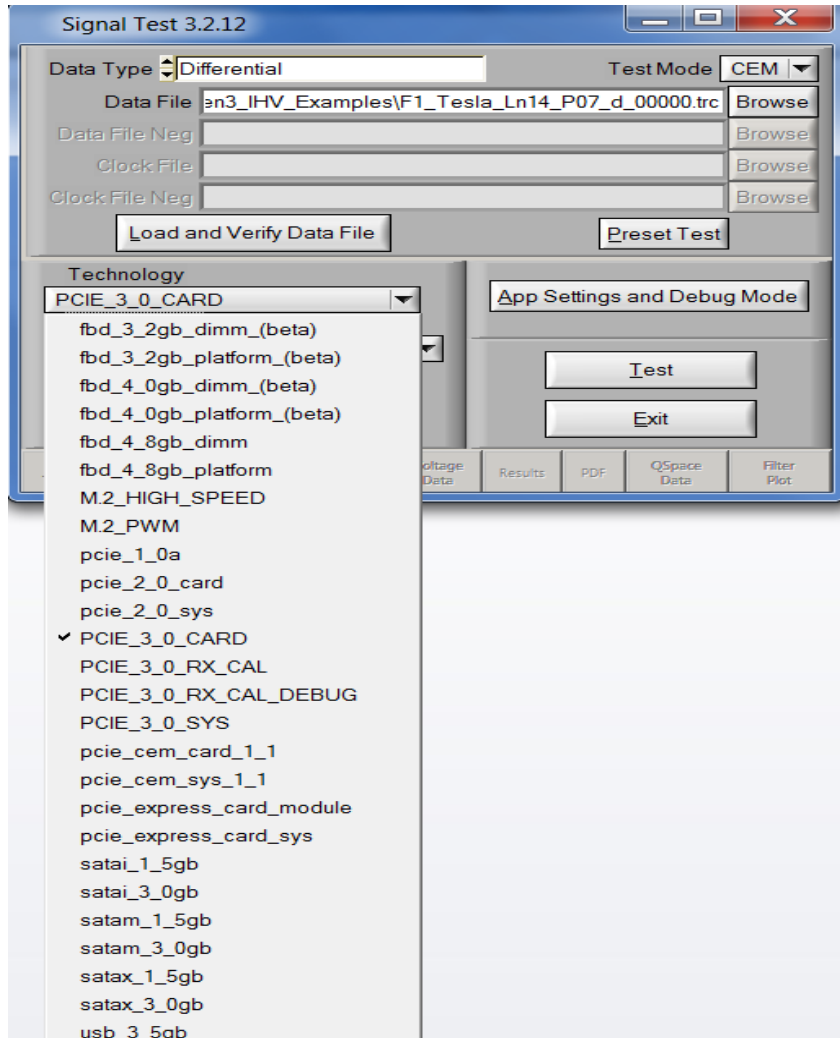
20" 85 Ohm reference channel

Base Spec RX Package Structure




- SMPs for all Tx and Rx Lanes
- Additional clock source for testing Dual Link modules
- 20" RX ISI channels with PCIe* 3.0 reference package structure
- DualLinkEn# control
- Tx lanes remain breakouts – reference channel embedded

Sigtest – Adding PCIe* SFF-8639 Support



- Tx and Rx Cal standard templates for Modules and SFF-8639 motherboards
- Built in embedding of SFF-8639 reference channel for Tx testing
- Rj/Dj jitter separation and Random/Deterministic voltage noise (eye height) separation
- PCIe* 3.0 reference equalizer (CTLE and DFE)
- Simultaneous clock/data analysis for motherboard TX test



THANK YOU!!



Backup Material




Flash Memory Summit Sigtest Overview

Windows* Application Consisting of:

- GUI
 - Sigtest.exe – Main GUI – written in LabWindows* CVI
 - GUI Options for each specification supported to perform standard pass/fail testing to TX parameters for that specification
- Analysis Libraries – written in C and compiled in Microsoft Visual Studio and using Intel Performance Primitives (IPP) math libraries for all basic math functions (DFT, IFT, etc.)
 - RjDj.dll
 - JitterEyedll.dll

Operates on any evenly sampled voltage waveform data in a variety of text and binary formats supporting most oscilloscopes

PCI* Express 2.0 CBB Baseline

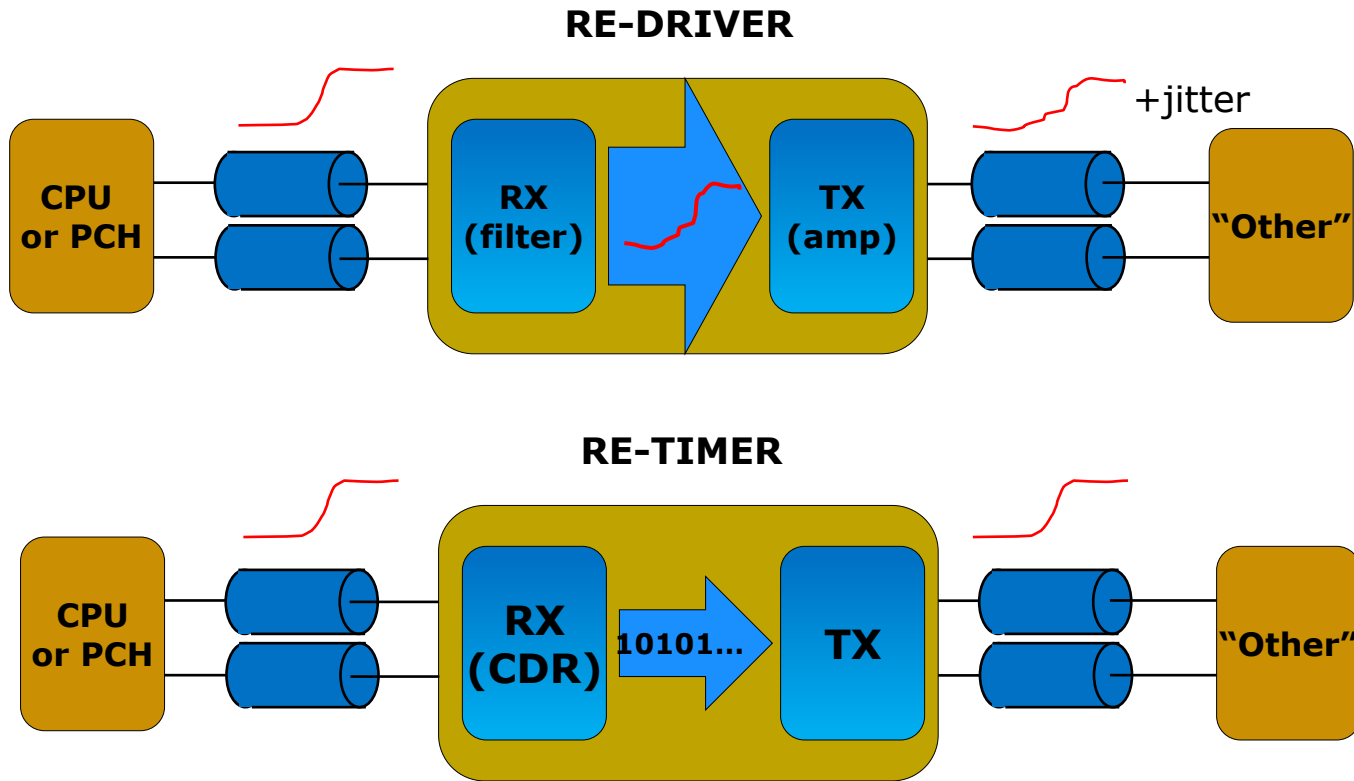




PCIe* SFF-8639 Form Factor Specification

- At revision .7 level in PCI-SIG workgroup
- Defines PCIe* electrical limits for modules and motherboards relative to the SFF-8639 connector
- Tx Limits (after PCIe* reference equalizer) and Reference Channels
 - Module
 - 20" of 85 ohm FR-4 and PCIe 3.0 Rx reference package structure
 - 34 mV eye height and .33 UI eye width
 - Motherboard
 - 4" of 85 ohm FR-4 and PCIe 3.0 Rx reference package structure
 - 34 mV eye height and .33 UI eye width
- Allow pass/fail Tx and RX PCIe electrical testing of whole module/motherboard.
- Root cause of failures requires additional debug.

PCI* Express Repeater Refresher



Primary difficulties are variations in devices and protocol awareness

Terminology

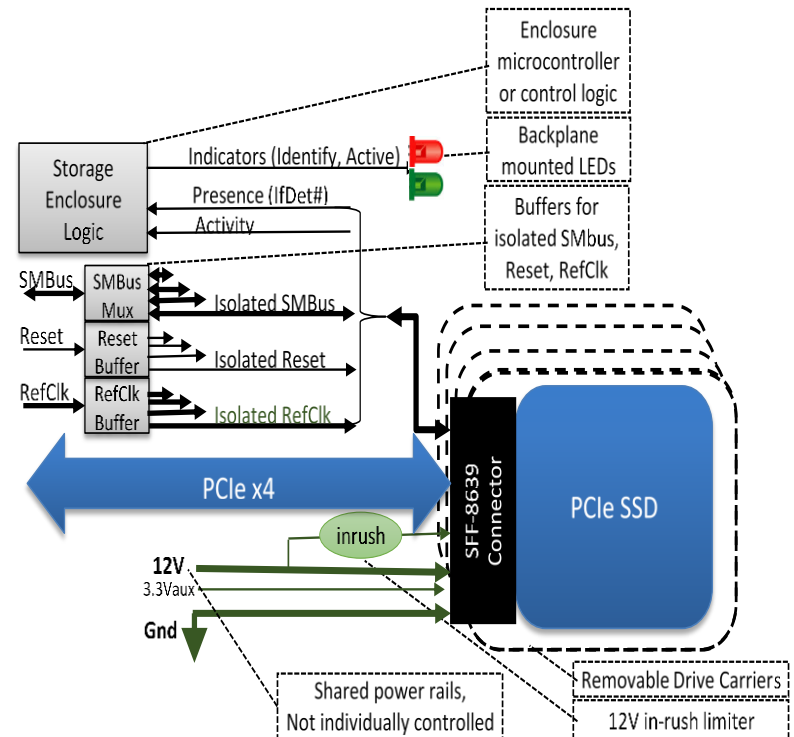
- **Hot Plug:** general term to describe adding and removing devices while system is running
- **Hot Add** – Also known as Hot Insertion
- **Hot Removal** – Software Managed Hot Removal (orderly)
- **Surprise Hot Removal** – possible outstanding IO transactions
- **Hot Swap (Hot Add + Removal)**

Requirements for Surprise Removal

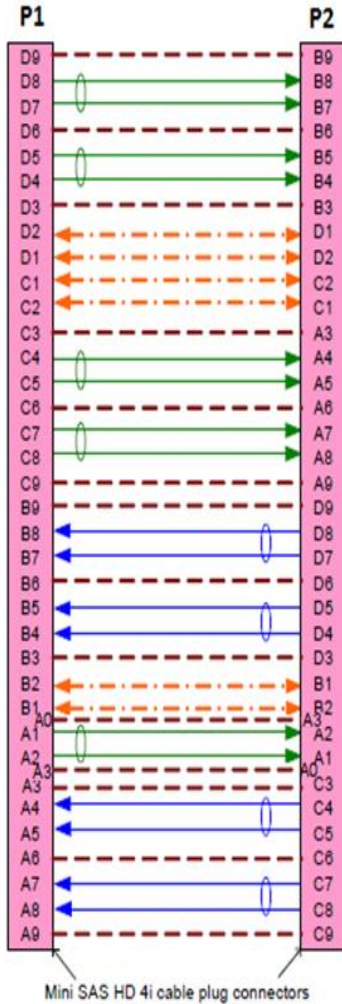
- **Hardware:** registers and drive status, master abort, and disable link
- **Software:** PCI Bus Driver and NVMe Driver
- **Drive:** Support unplanned power loss
- **LER, DPC, eDPC** – not required but make it easier to validate

PCI* Hot Plug Requirements – System



- PCIe* Slot Capability register: Hot Plug Capable and Hot Plug Surprise
- PCIe Slot Status: Presence Change Interrupt to notify PCIe bus driver
- Backplane, pre-charge circuit to limit in-rush current, isolated Reset, Refclk, and Smbus, presence detect via IfDet# (pin 4) and PRSNT# (pin10)
- Drive Identify and Fail Indicators
- PCIe Link Down Interrupt – for link down, uses PCIe AER
- BIOS: UEFI 2.3.1 or later, pre-allocate memory resources
- Pre-allocate slot resources (Bus IDs, interrupts, memory regions) using ACPI tables



miniSAS HD Cable and Connector



A0 is an internal only GND Pad (Yellow Circles) **This is NEW!**
A3 is an existing GND Pad (Green Circles)



Legend:

- High-speed serial differential pairs
- Signal return
- Sideband signal

(Modified for PCIe SSDs)