OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.

Datacenter Initiatives and Programs for Storage

Mark Carlson
Principal Engineer, Industry Standards
Toshiba Memory America

Initiatives

- Storage Interfaces
  - NVMe
  - T10/T13
- Storage Management
  - Redfish
  - Swordfish
- Form Factors
  - SFF

Storage Interfaces

- Hyperscalers have unique requirements for storage devices at datacenter scale
  - Eliminate/Reduce tail latency
  - Eliminate “noisy neighbor” interference with multiple apps using the same (large) drive
- These requirements impact the storage interfaces used for SSDs and HDDs
- For NVMe, I/O Determinism has resulted in changes to the NVMe standard
  - NVM Sets – divide an SSD into isolated smaller groups of media
  - Predictable Latency – time based Windowing for predictable reads without background task interference
- For T10/T13, changes are being discussed for fast fail and media types
  - OCP Storage – propose changes after working out details in subteam “streams”
  - Ad Hoc – open (to all) discussions with INCITS members

Storage Management

- Management of storage devices has issues in a scale out environment
  - Host based agents are largely used to gather data from attached devices
  - A conversion/adaptation from proprietary to common models is needed per vendor
- To scale out better, devices should accept and report core information in a common format. Vendor specific information should be available using the same protocol.
- DMTF has a standard called Redfish that has done this for systems management and it is gaining traction in datacenters
- Redfish has basic storage drive models that can be used for inventory and telemetry
- SNIA has an extension to Redfish for storage management called Swordfish
  - This can serve as a common model for SDS and other higher level storage software

Form Factors

- The U.2 form factor has wide adoption in datacenters and will likely continue to dominate hard drive devices for the foreseeable future
- However for SSDs, this may prove to be too limiting going forward
- M.2 form factor is becoming popular but limits the capacity of SSDs
  - Carrier cards (some in the new form factors) can extend M.2 into bigger FRUs
- New form factors are intended to serve for both carrier cards and single controller drives
- EDSFF has created a number of new SSD for factors and these have now been standardized by SNIA SFF
- JEDEC is working on NGSFF (leveraging M.2) and M.3

1U Long

- Standard available as SFF-TA-1007 from the SNIA.org website
- Developed by the EDSFF group and submitted to SNIA
  - Hyperscaler members include Facebook and Microsoft Azure
- Products are being announced based on this form factor
- Can be used by carrier cards to hold M.2 and 1U Short SSDs

Example system with 1U Long cards

1U Short

- Standard available as SFF-TA-1006 from the SNIA.org website
- Products are being announced based on this form factor
- Can be used in carrier cards or as a standalone add-in card for systems

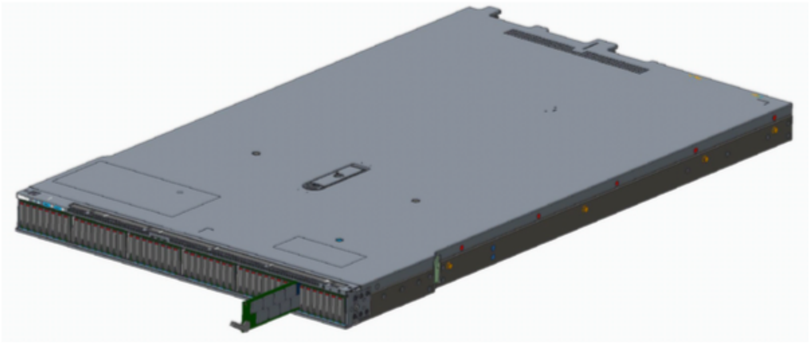

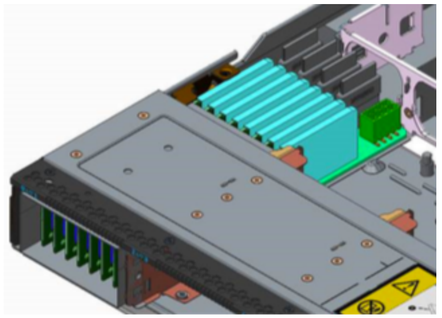

Example systems with 1U Short cards

3” Media

- Standard available as SFF-TA-1008 from the SNIA.org website
- Designed to either fit sideways in a 1U chassis or vertically in a 2U chassis
- Similar in capacity as tradition U.2 devices
- There are both short and long versions standardized
- Single (7.5mm) and double (16.8mm) width options

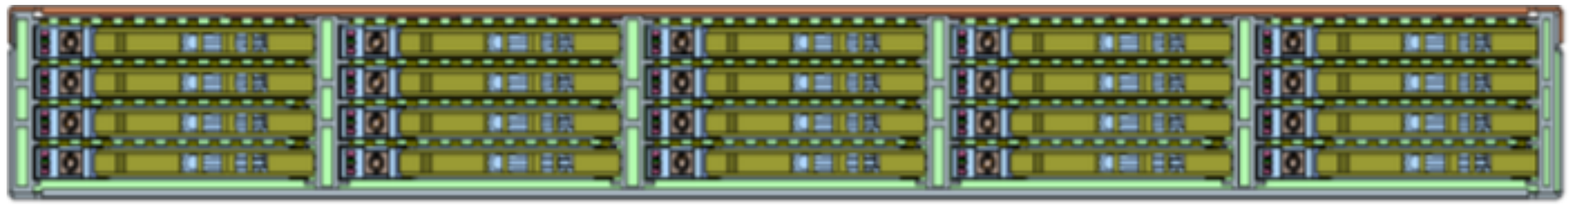

Example systems with 3” Media cards

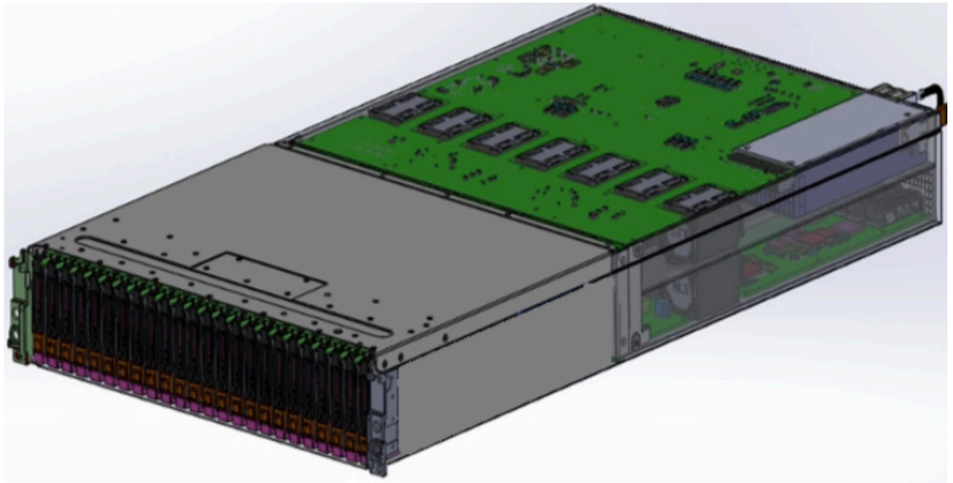

Connector for new Form Factors

- All of these new form factors share a common, next generation connector
- Standard available as SFF-TA-1002 from the SNIA.org website
- Options for x4, x8 and x16 PCIe lanes and future proofed for Gen4 and Gen5 speeds
- Should also work for future interconnect standards such as GenZ

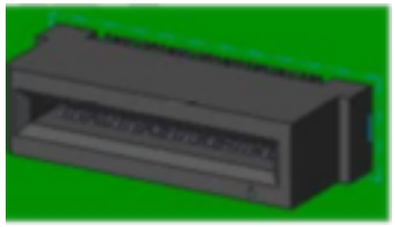

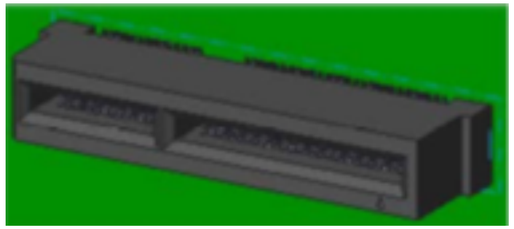

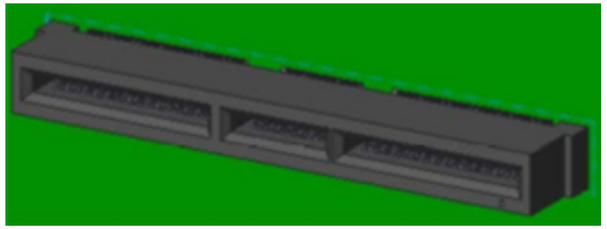

NGSFF

- Proposal for a form factor based on an expanded M.2 type connector
- Connector may accommodate M.2 cards
- Currently being worked in JEDEC
- Not yet publically available

Discussion

- These initiatives are slowly changing the storage industry to better address hyperscaler requirements for the datacenter
- Should also help tier two datacenter customers following the hyperscaler practices
- Recommend: Get involved!
  - OCP Storage, SNIA, DMTF, others
  - What other initiatives might we create to help solve datacenter problems?

OCP SUMMIT