



# NVMe™ Management Interface (NVMe-MI™) Workgroup Update

Sponsored by NVM Express® organization, the owner of NVMe™, NVMe-oF™ and NVMe-MI™ standards

# Speakers

Austin Bolen

**DELL**EMC

Myron Loewen



Flash Memory Summit

**nvm**  
EXPRESS®

# Agenda

- NVMe-MI™ Workgroup Update
- NVMe-MI 1.0a Overview
- What's new in NVMe-MI 1.1
  - In-band NVMe-MI
  - Enclosure Management
  - Managing Multi NVM Subsystem Devices
- Summary



# NVM Express®<sup>®</sup>, Inc. 120+ Companies defining NVMe™ together

## Board of Directors

13 elected companies, stewards of the technology & driving processes

*Chair: Amber Huffman*



## Marketing Workgroup

NVMexpress.org, webcasts, tradeshow, social media, and press

*Co-Chairs: Janene Ellefson and Jonmichael Hands*

## Technical Workgroup

NVMe Base and NVMe Over Fabrics

*Chair: Amber Huffman*

## Management Intf. Workgroup

Out-of-band management over SMBus and PCIe® VDM

*Chair: Peter Onufryk  
Vice Chair: Austin Bolen*

## Interop (ICC) Workgroup

Interop & Conformance Testing in collaboration with UNH-IOL

*Chair: Ryan Holmqvist*

facebook

Microsoft



CISCO

DELL EMC

SEAGATE

TOSHIBA



Micron

ORACLE

SAMSUNG

Microsemi

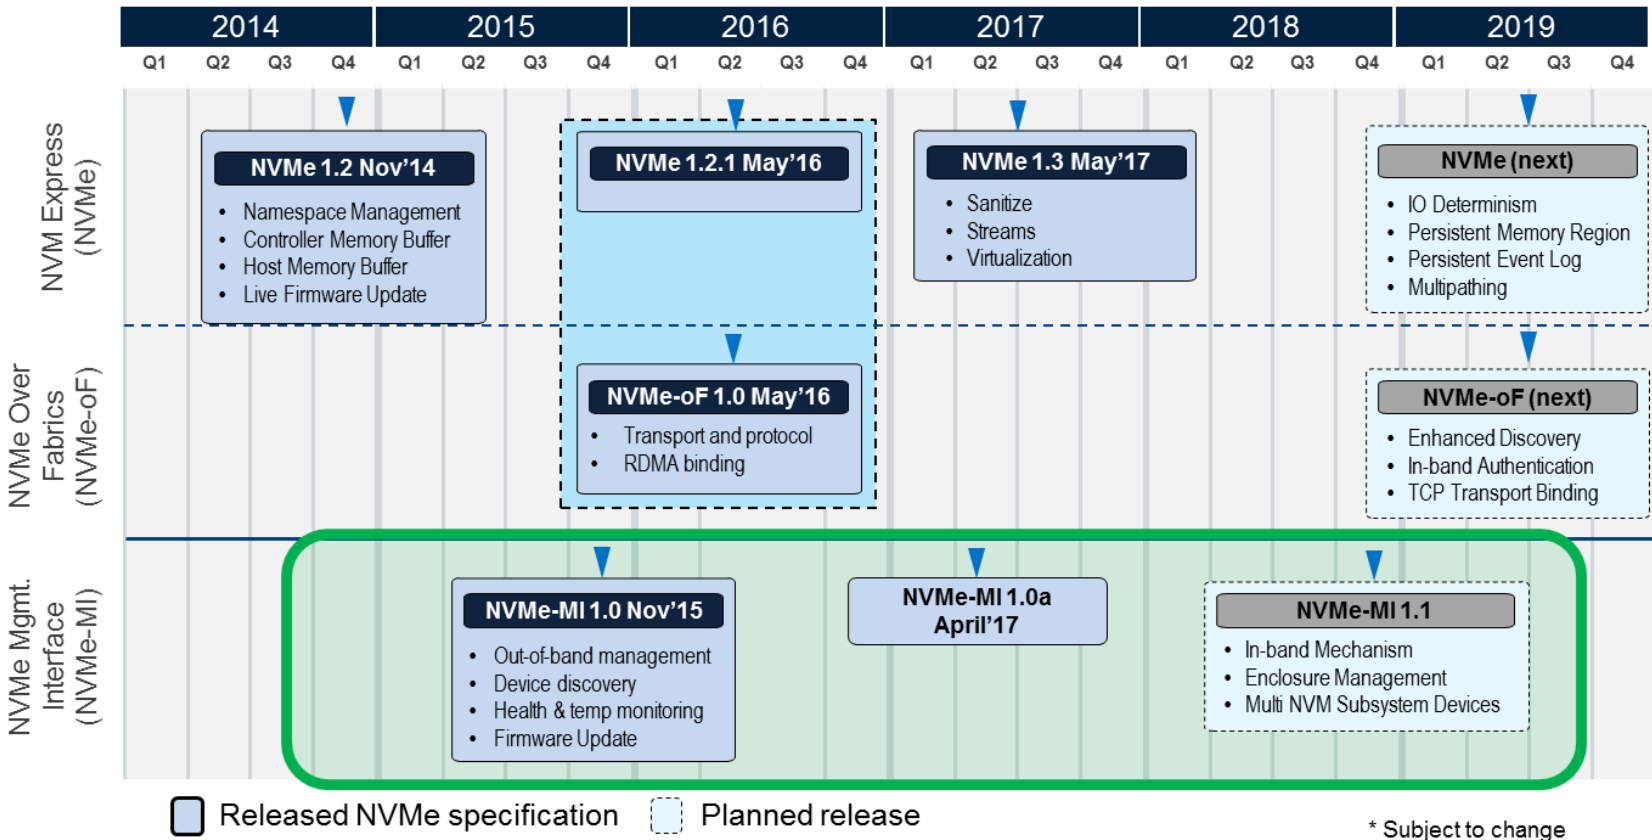
NetApp

Western Digital

Flash Memory Summit

nvm EXPRESS

# NVM Express™ Roadmap



# NVMe-MI™ Ecosystem

- Commercial test equipment for NVMe-MI
- NVMe-MI 1.0a compliance testing program has been developed
  - Compliance testing started in the May 2017 NVMe™ Plugfest conducted by the University of New Hampshire Interoperability Laboratory (UNH-IOL)
  - 7 devices from multiple vendors have passed compliance testing and are on the NVMe-MI Integrators List
- Servers are shipping that support NVMe-MI



Flash Memory Summit

**nvm**  
EXPRESS®

# What is the NVMe™ Management Interface 1.0a?

A programming interface that allows out-of-band management of an NVMe Storage Device Field Replaceable Unit (FRU)



Flash Memory Summit


**nvm**  
EXPRESS®

# Out-of-Band Management and NVMe-MI™

- **Out-of-Band Management** – Management that operates with hardware resources and components that are *independent of the host operating system control*

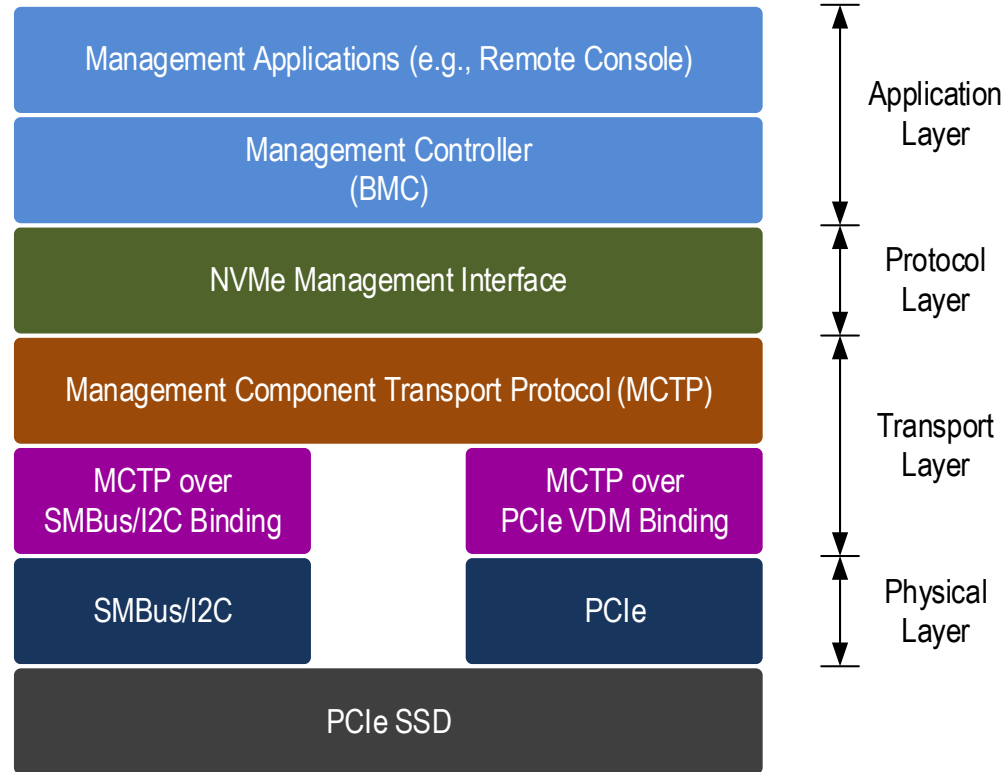
- **NVMe™ Out-of-Band Management Interfaces**

- SMBus/I2C
- PCIe Vendor Defined Messages (VDM)








# NVMe-MI™ Out-of-Band Protocol Layering

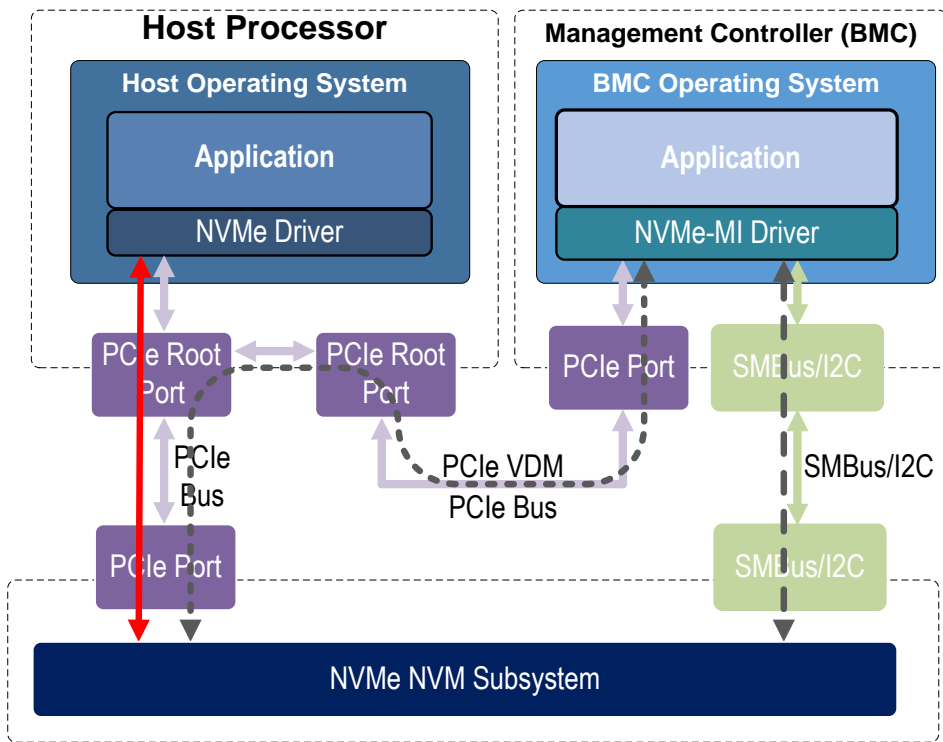


# NVMe™ Storage Device in 1.0a



- **NVMe Storage Device** – One NVM Subsystem with one or more ports, vital product data (VPD), and an optional SMBus/I2C interface

# In-Band Management and NVMe-MI™



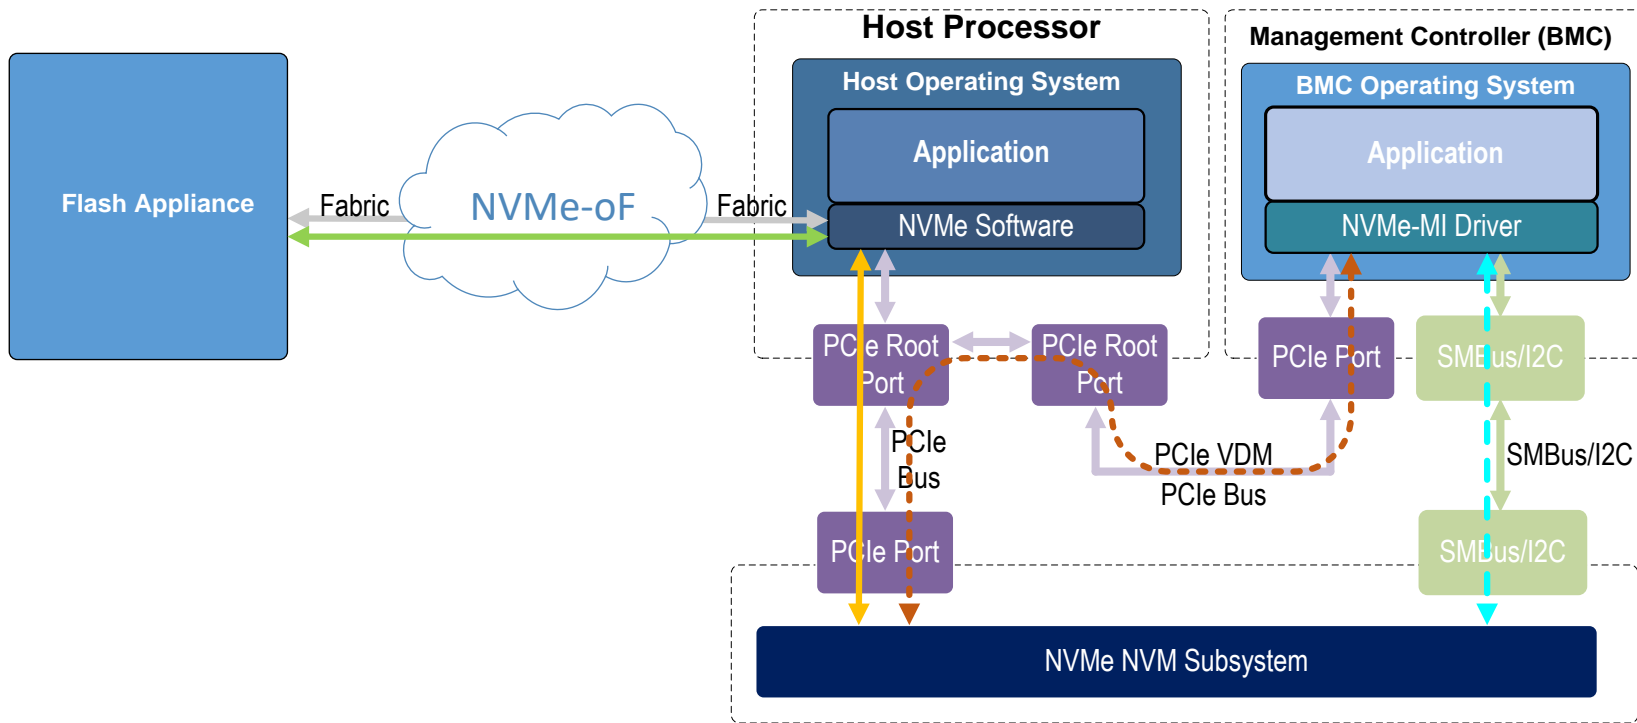
- In-band mechanism allows application to tunnel NVMe-MI commands through NVMe™ driver
  - Two new NVMe Admin commands
    - NVMe-MI Send
    - NVMe-MI Receive
- Benefits
  - Provides management capabilities not available in-band via NVMe commands
    - Efficient NVM Subsystem health status reporting
    - Ability to manage NVMe at a FRU level
    - Vital Product Data (VPD) access
    - Enclosure management



Flash Memory Summit

**nvm**  
EXPRESS®

# NVMe-MI™ over NVMe-oF™



Plumbing in place for NVMe-MI over NVMe-oF




Flash Memory Summit



# Enclosure Management

- SES Based Enclosure Management
  - Technical proposal developed in NVMe-MI™ workgroup
  - While the NVMe™ and SCSI architectures differ, the elements of an enclosure and the capabilities required to manage these elements are the same
    - Example enclosure elements: power supplies, fans, display or indicators, locks, temperature sensors, current sensors, voltage sensors, and ports
  - Comprehensive enclosure management that leverages SCSI Enclosure Services (SES), a standard developed by T10 for management of enclosures using the SCSI architecture



# Multi NVM Subsystem Management




Flash Memory Summit


**nvm**  
EXPRESS®

# NVMe-MI™ 1.0a NVMe™ Storage Device

- **NVM Storage Device** – One NVM Subsystem with one or more ports and an optional SMBus/I2C interface




Single Ported PCIe SSD




Dual Ported PCIe SSD with SMBus/I2C


# NVMe™ Storage Device with Multiple NVM Subsystems



M.2 Carrier Board from Amfeltec

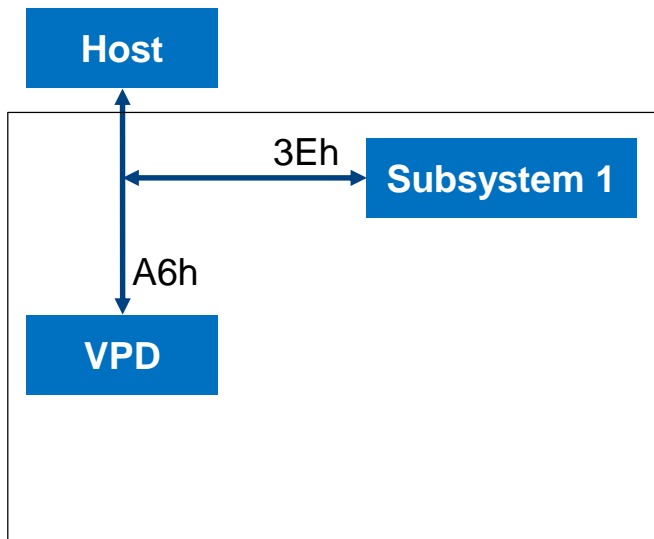


ANA Carrier Board from Facebook





# SMBus Topology for NVMe-MI™ 1.0

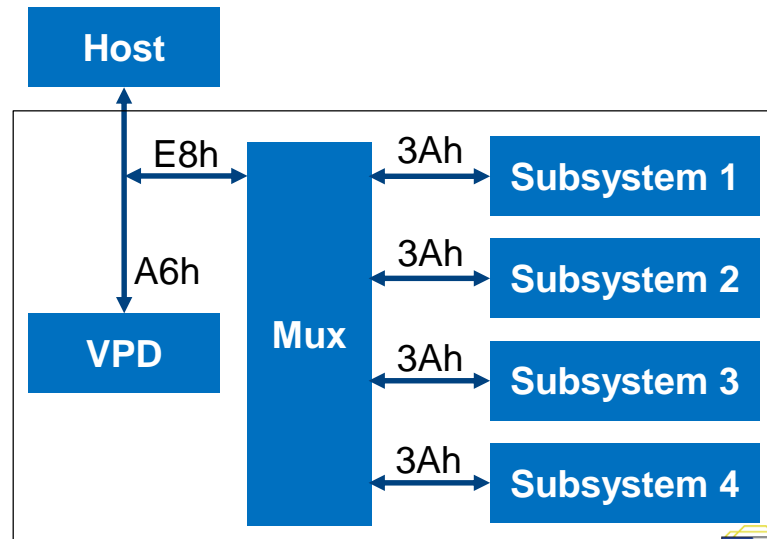
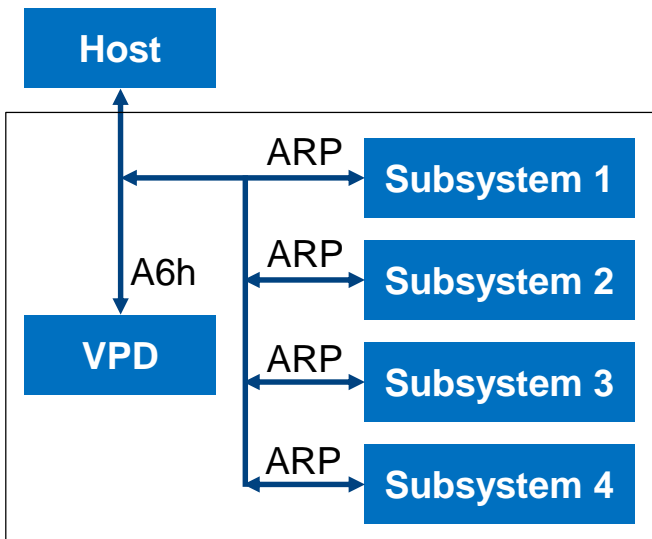


Flash Memory Summit

**nvm**  
EXPRESS®


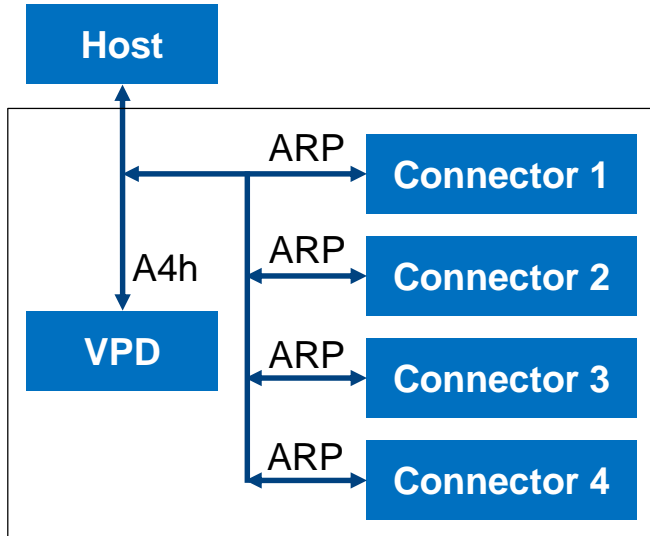
# Multiple NVM Subsystems on a single SMBus Port

- Describe topology in new VPD MultiRecord
- Add UDID types for additional devices like Mux




# Support Expansion Connectors

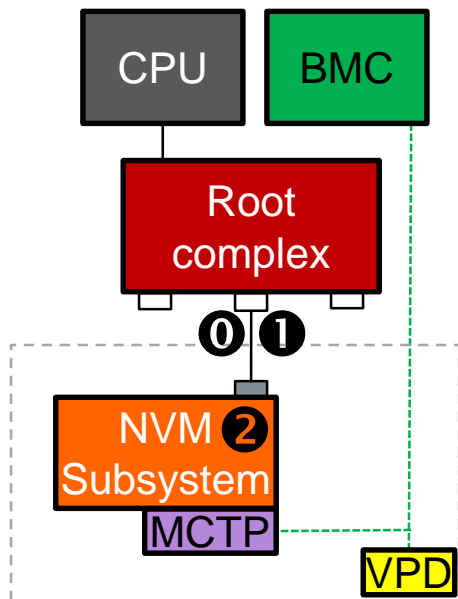
- New VPD address to avoid conflicts with plugged in devices
- Optional Labels for each connector to assist technicians



# A Connection Graph Between Element Types



# Single Port Example (35 bytes of 256B EEPROM)



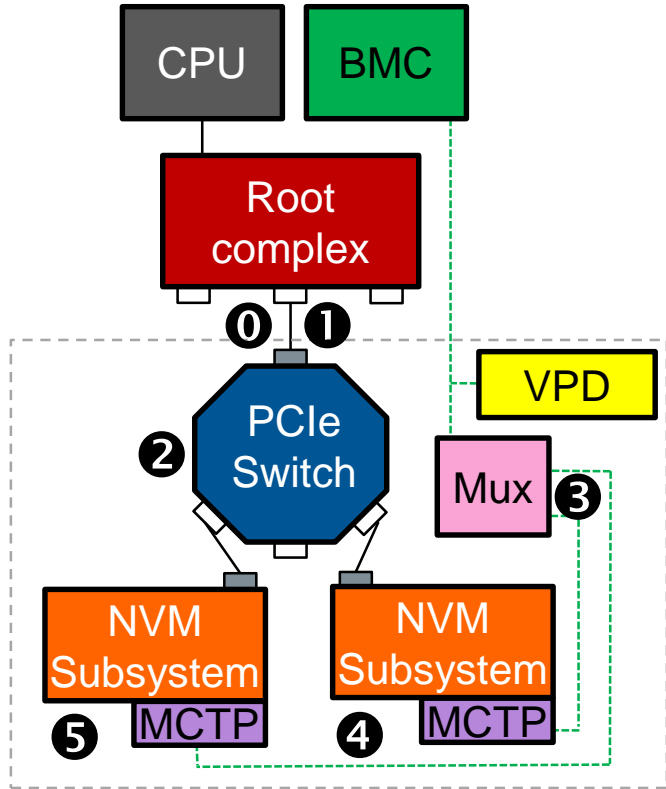
<b>Header</b>	Record: Element 0Dh	Record Format 82h	Record Length 23h	Record Chcksm 34h	Header Chcksm 75h	Version Number 00h	Rsvd 00h	Element Count 03h	
<b>Element 0</b>	Type: Host 01h	Element Length 08h	Form Factor 12h	SMBus Dest 02h	Link Options 00h	Link 0 Width 84h	Link 0 Start 00h	Link 0 Dest. 02h	
<b>Element 1</b>	Type: Power 02h	Element Length 08h	Thermal Load 0Fh	Vaux Load 32h	Rail Options 00h	Rail Voltage 78h	12V initial 08h	12V max 0Fh	
<b>Element 2</b>	Type: NVMe 09h	Element Length 13h	MCTP Address 3Ah	SMBus speed 01h	PCIe Ports 12h	Port 0 Speed 0Fh	Port 0 Flags 01h	Total NVM Capacity (MSB first) 000000000000000000000000h	



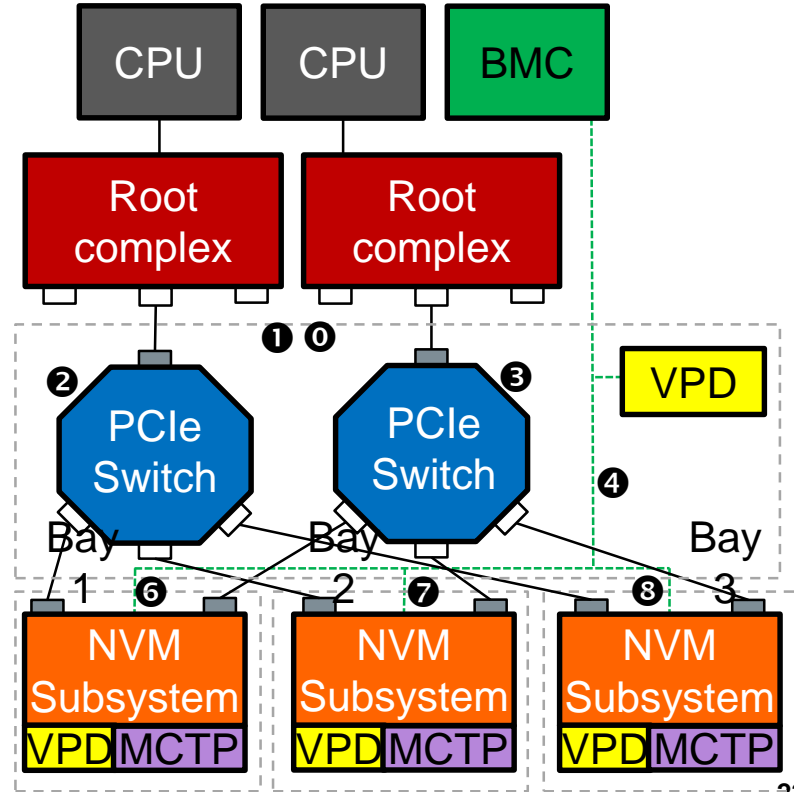
Flash Memory Summit

**nvm**  
EXPRESS®

# 2 NVM Subsystems with Mux (82)



# Dual Port with Expansion Connectors (78)



# Summary

- NVMe-MI™ 1.0a is gaining market acceptance and is available in shipping products
- NVMe-MI 1.1 is nearing completion
  - Significant new features
    - In-band mechanism
    - Enclosure management
    - Support for multi NVM subsystem management
- It is time to start thinking about anchor features for NVMe-MI 1.2



# Additional Material on NVMe-MI™

- BrightTALK Webinar
  - <https://www.brighttalk.com/webcast/12367/282765/the-nvme-management-interface-nvme-mi-learn-whats-new>
- Flash Memory Summit 2017
  - Slides: [https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2017/20170808\\_FA12\\_PartA.pdf](https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2017/20170808_FA12_PartA.pdf)
  - Video:
    - <https://www.youtube.com/watch?v=daKL7tlvNII>
    - <https://www.youtube.com/watch?v=Daqj-XqICo8>
- Flash Memory Summit 2015
  - Slides: [https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2015/20150811\\_FA11\\_Carroll.pdf](https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2015/20150811_FA11_Carroll.pdf)
- Flash Memory Summit 2014
  - Slides: [https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2014/20140804\\_SeminarF\\_Onufryk\\_Bolen.pdf](https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2014/20140804_SeminarF_Onufryk_Bolen.pdf)
- NVMe-MI Specification
  - <https://nvmexpress.org/resources/specifications/>



Flash Memory Summit

**nvm**  
EXPRESS®



# References

MCTP Overview: <http://dmtf.org/sites/default/files/standards/documents/DSP2016.pdf>

MCTP Base Spec: [https://www.dmtf.org/sites/default/files/standards/documents/DSP0236\\_1.3.0.pdf](https://www.dmtf.org/sites/default/files/standards/documents/DSP0236_1.3.0.pdf)

MCTP SMBus/I2C Binding:

[https://www.dmtf.org/sites/default/files/standards/documents/DSP0237\\_1.1.0.pdf](https://www.dmtf.org/sites/default/files/standards/documents/DSP0237_1.1.0.pdf)

MCTP PCIe VDM Binding:

[https://www.dmtf.org/sites/default/files/standards/documents/DSP0238\\_1.0.2.pdf](https://www.dmtf.org/sites/default/files/standards/documents/DSP0238_1.0.2.pdf)

IPMI Platform Management FRU Information Storage Definition:

<https://www.intel.la/content/www/xl/es/servers/ipmi/ipmi-platform-mgt-fru-infostorage-def-v1-0-rev-1-3-spec-update.html>



Flash Memory Summit

**nvm**  
EXPRESS®