nvm EXPRESS®

NVNe™ State of the Union

Peter Onufryk

intel

NVM Express, Inc.
120+ Companies Defining NVMe Together

Board of Directors

13 elected companies, stewards of the technology & driving processes

Chair: Amber Huffman

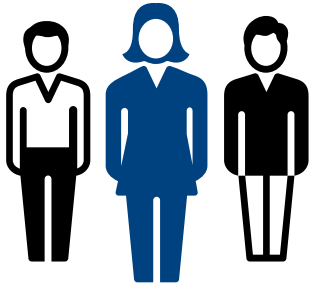

Marketing Workgroup

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

Co-Chairs: Jonmichael Hands and Cameron Brett

Technical Workgroup

NVMe Base and NVMe Over Fabrics

Chair: Peter Onufryk

Management Intf. Workgroup

NVMe Management

Co-Chairs: Austin Bolen and John Geldman

Interop (ICC) Workgroup

Interop & Conformance Testing in collaboration with UNH-IOL

Chair: Ryan Holmqvist

facebook Microsoft intel CISCO DELL EMC SEAGATE Toshiba Memory

Micron MARVELL SAMSUNG Microsemi NetApp Western Digital

Ten Years of NVMe

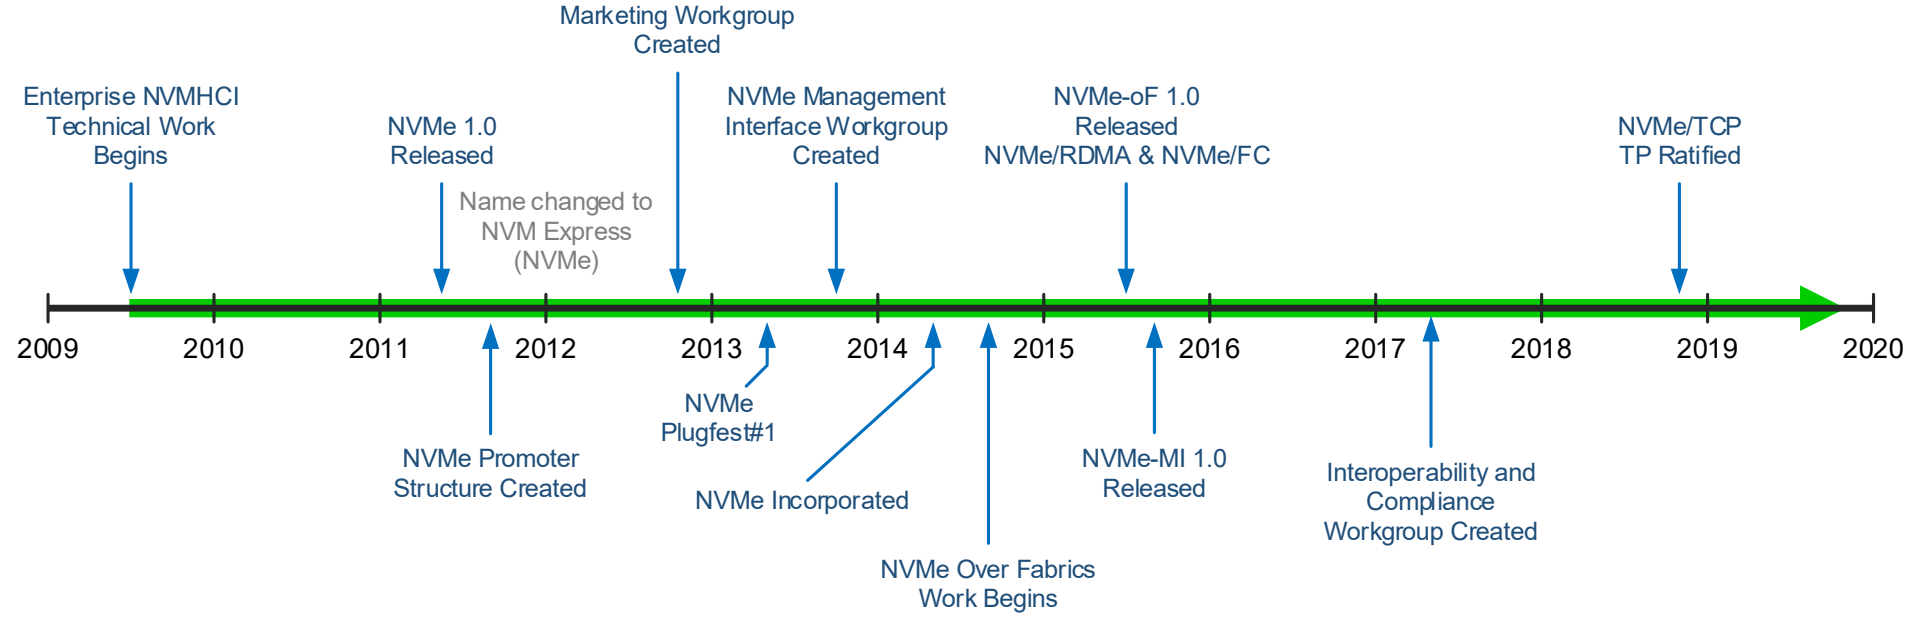

Strong Growth Across Segments

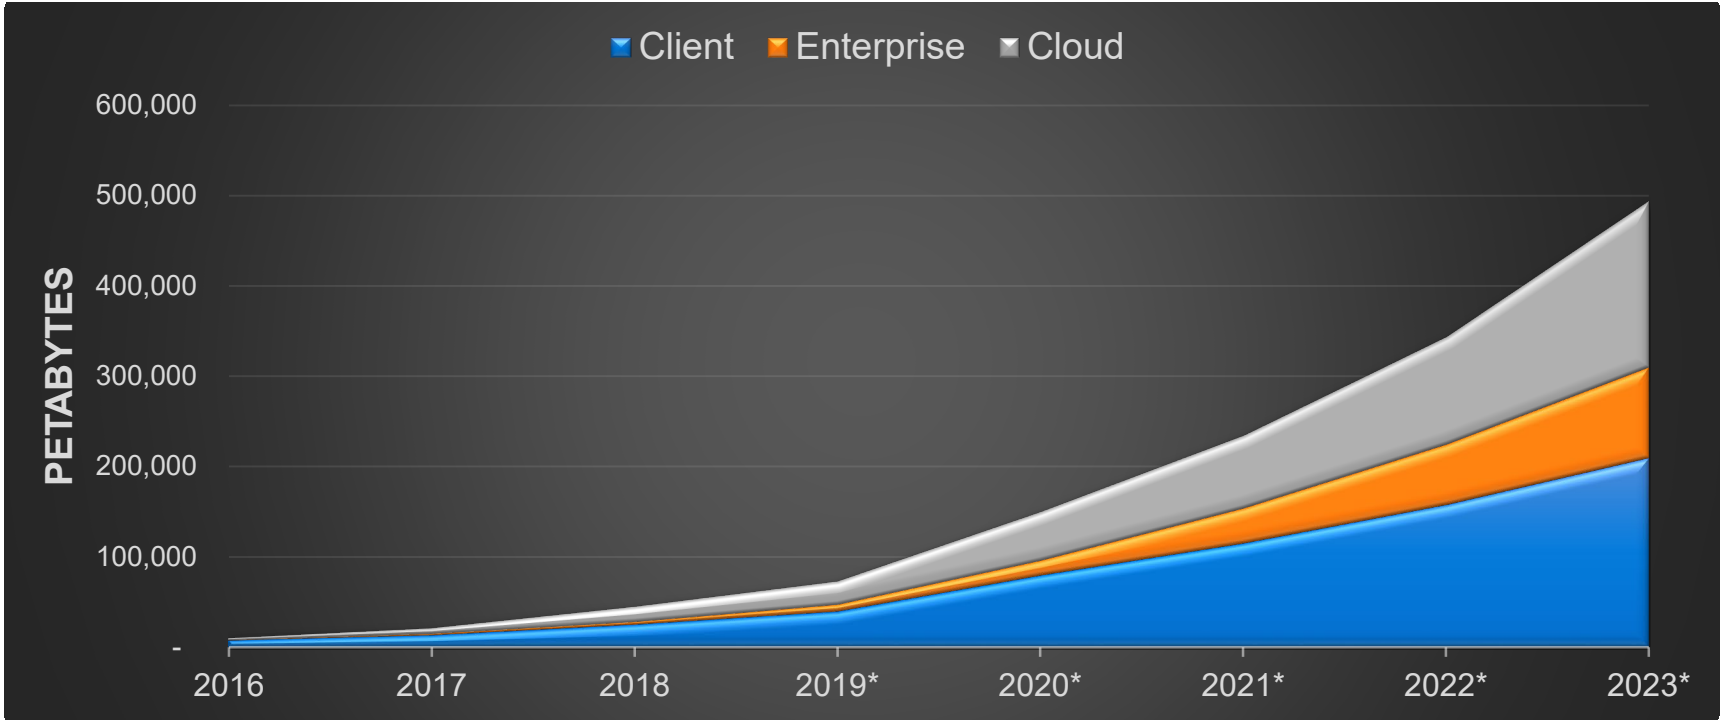

* Projections provided by Forward Insights Q2'19

With Millions of Units Shipping

K Units	2016	2017	2018	2019*	2020*	2021*
Enterprise	364	749	1,048	2,774	5,740	11,192
Cloud	2,051	3,861	10,231	17,338	25,891	31,050
Client	33,128	50,385	82,613	111,888	187,689	243,889

* Projections provided by Forward Insights Q2'19

NVMe is The New Language of Storage

NVMe SSDs	23 Companies Shipping 96 Models
NVMe Severs	13 Companies Shipping 93 Models
NVMe AFAs	11 Companies Shipping 21 Models
NVMe Appliances	8 Companies Shipping 21 Models
NVMe-oF HBAs/NICs/RNICs	5 Companies Shipping 53 Models
NVMe-oF Accelerated Adapters	6 Companies Shipping

* Data provided by G2M Research

2019 NVMe Deliverables

NVMe Base Specification	
NVMe 1.4	6/10/2019
NVMe 1.3	5/1/2017
NVMe 1.2	11/3/2014
NVMe 1.1	10/11/2012
NVMe 1.0	5/14/2008

NVMe Over Fabrics Specification	
NVMe-oF 1.1	45-day Review
NVMe-oF 1.0	6/5/2016

NVMe Management Interface Specification	
NVMe-MI 1.1	4/29/2019
NVMe-MI 1.0	11/17/2015

NVMe Plugfest	
Plugfest #11	6/24/2019
Plugfest #10	11/12/2018
⋮	
Plugfest #1	5/13/2013

NVMe Specification Roadmap

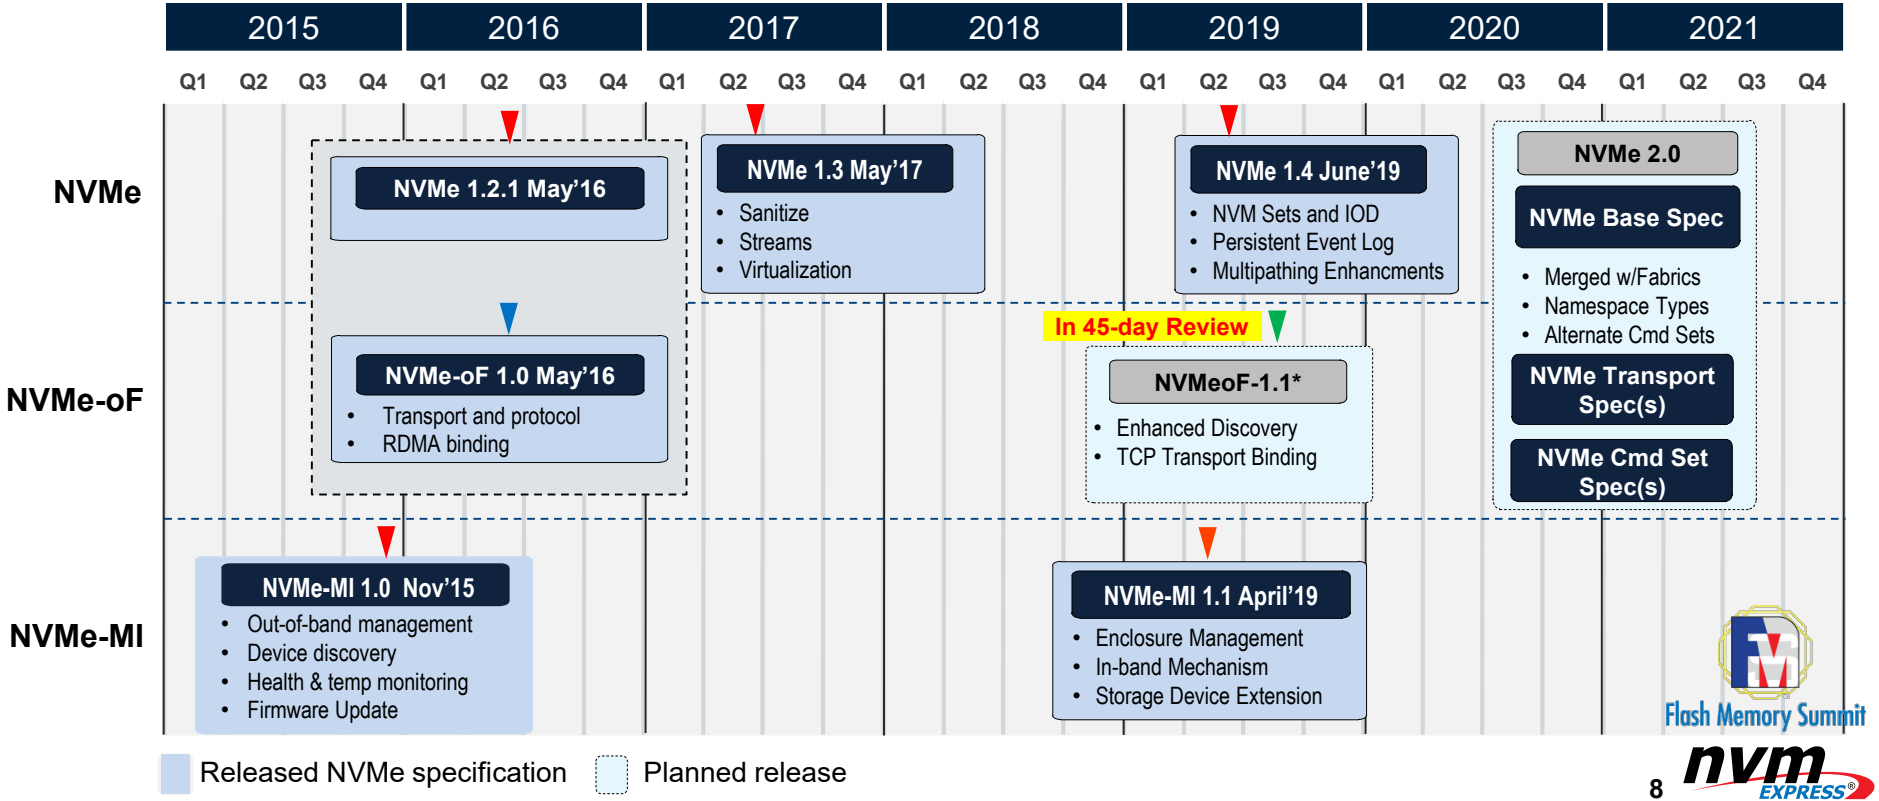

Three New Specifications for 2019

NVMe 1.4

NVM Sets and **IO Determinism** enable better performance, isolation, and QoS for hyperscale data centers. **Persistent event log** provides robust drive history for issue triage and debug. **Multipathing** provides optimal path for a namespace in multi-controller topologies

NVMe-oF 1.1

Enhanced Discovery for hosts to discover new NVMe devices. **TCP Transport Binding** NVMe/TCP enables efficient end-to-end NVMe operations with standard IP network with excellent performance and latency characteristics

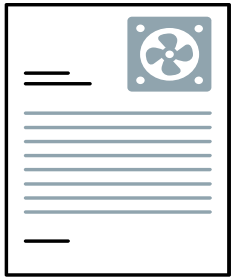

NVMe-MI 1.1

Enclosure Management enhances NVMe-MI for storage arrays for slot control, LED, and fans. **In-band Mechanism** opens up the NVMe-MI command set to standard NVMe driver (VPD, FRU). **Storage Device Extension** extends NVMe-MI to carrier cards and multiple controller devices

I/O Determinism – NVM Sets

- NVM Sets are QoS Isolated
  - Write to namespace A1 does not impact QoS associated with namespace B2
- NVM Subsystem may support one or more NVM Sets
- One or more Namespaces may be allocated to an NVM Set

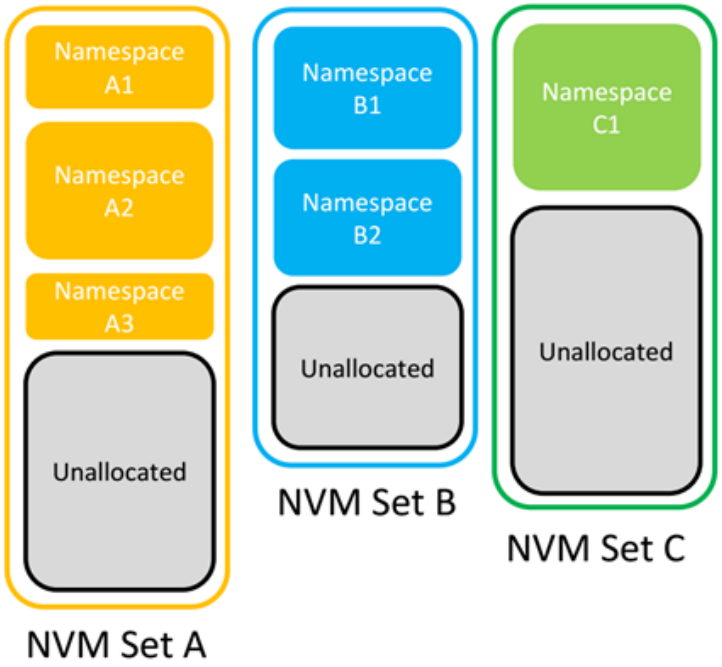

I/O Determinism – Predictable Latency Mode

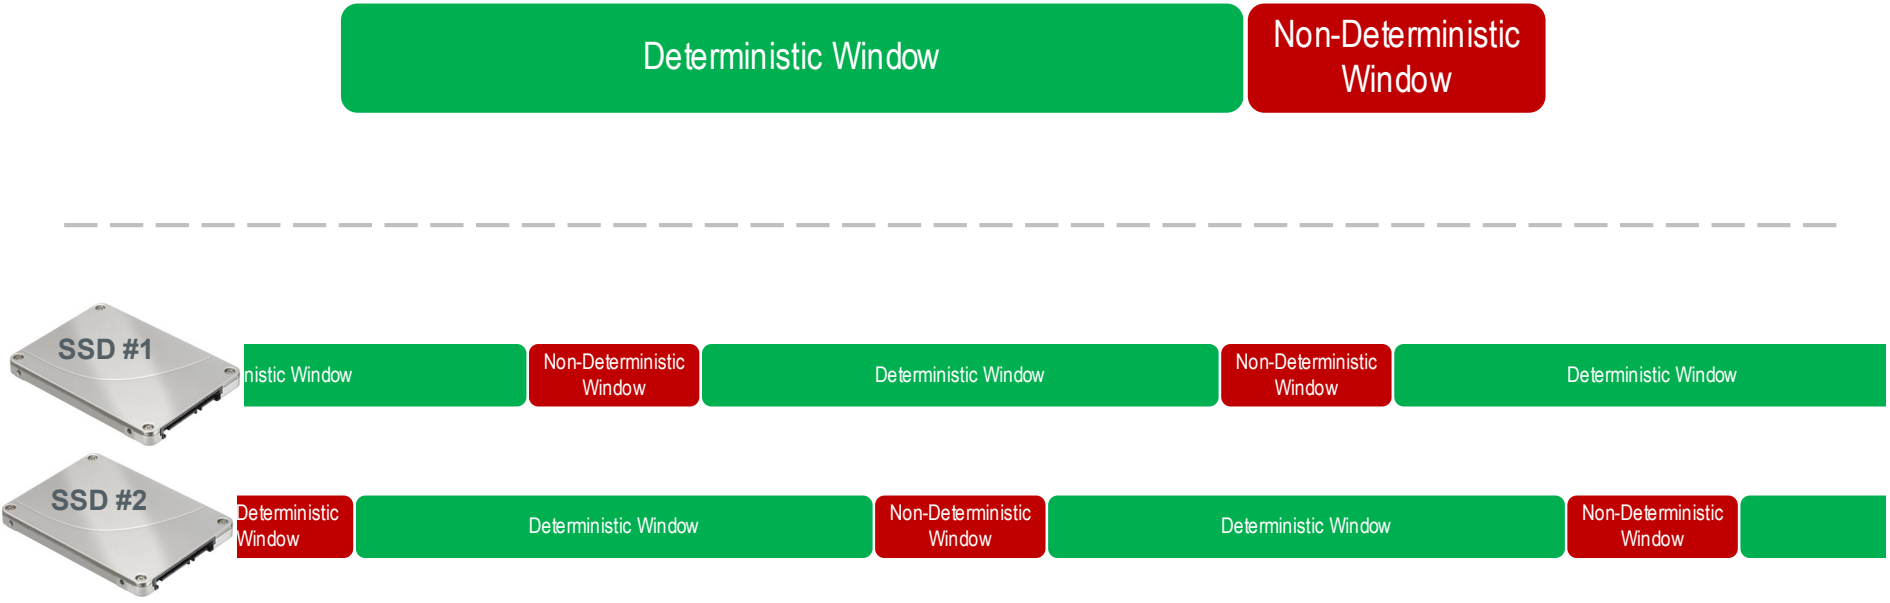

NVMe in High End Storage Systems

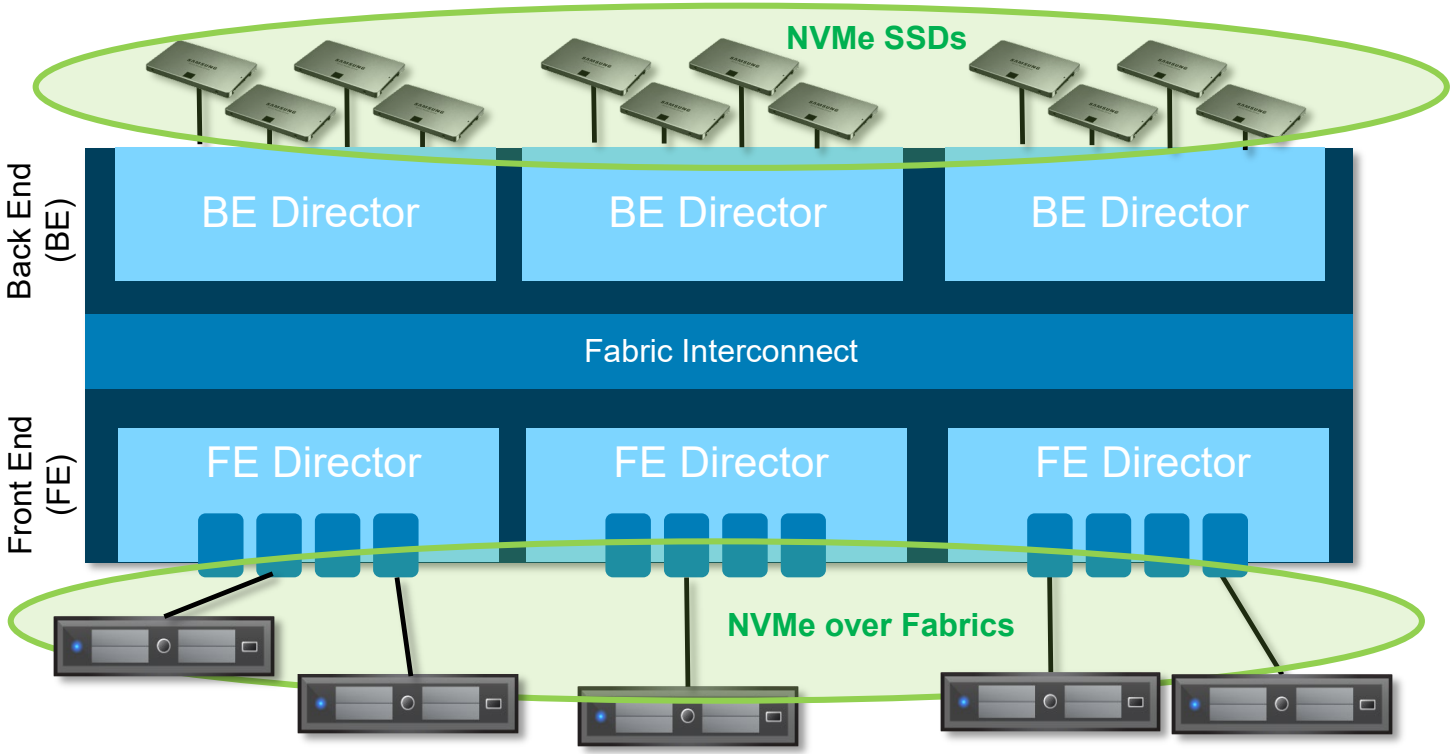

Slide credit: Dell EMC

Asymmetric Namespace Access (ANA)

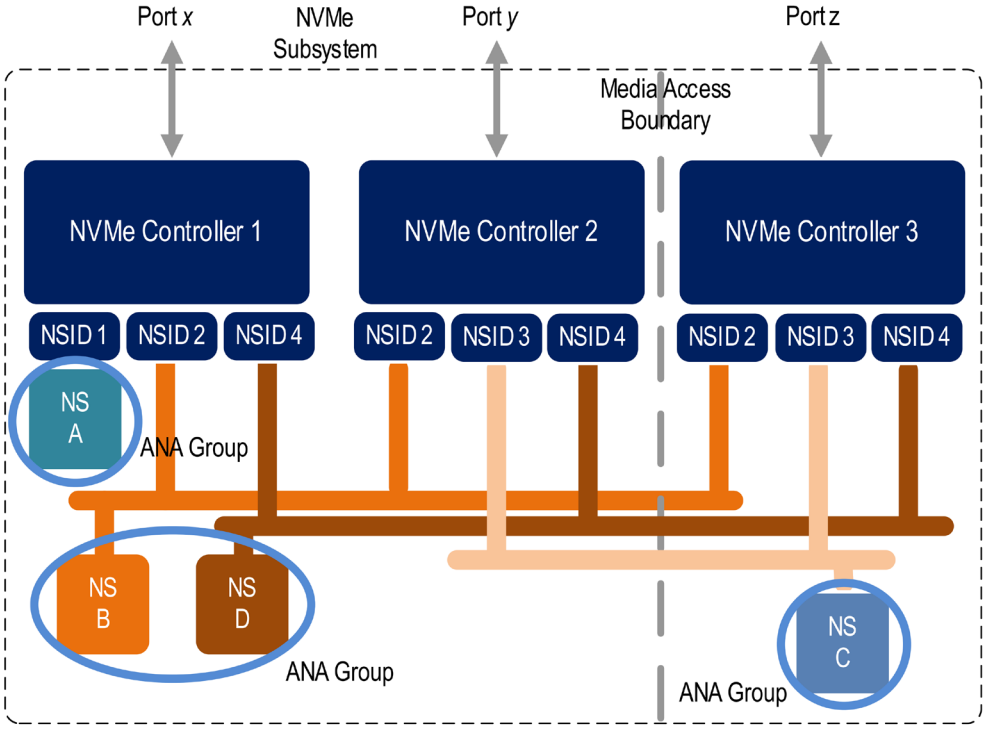

NVMe-oF TCP Transport Binding

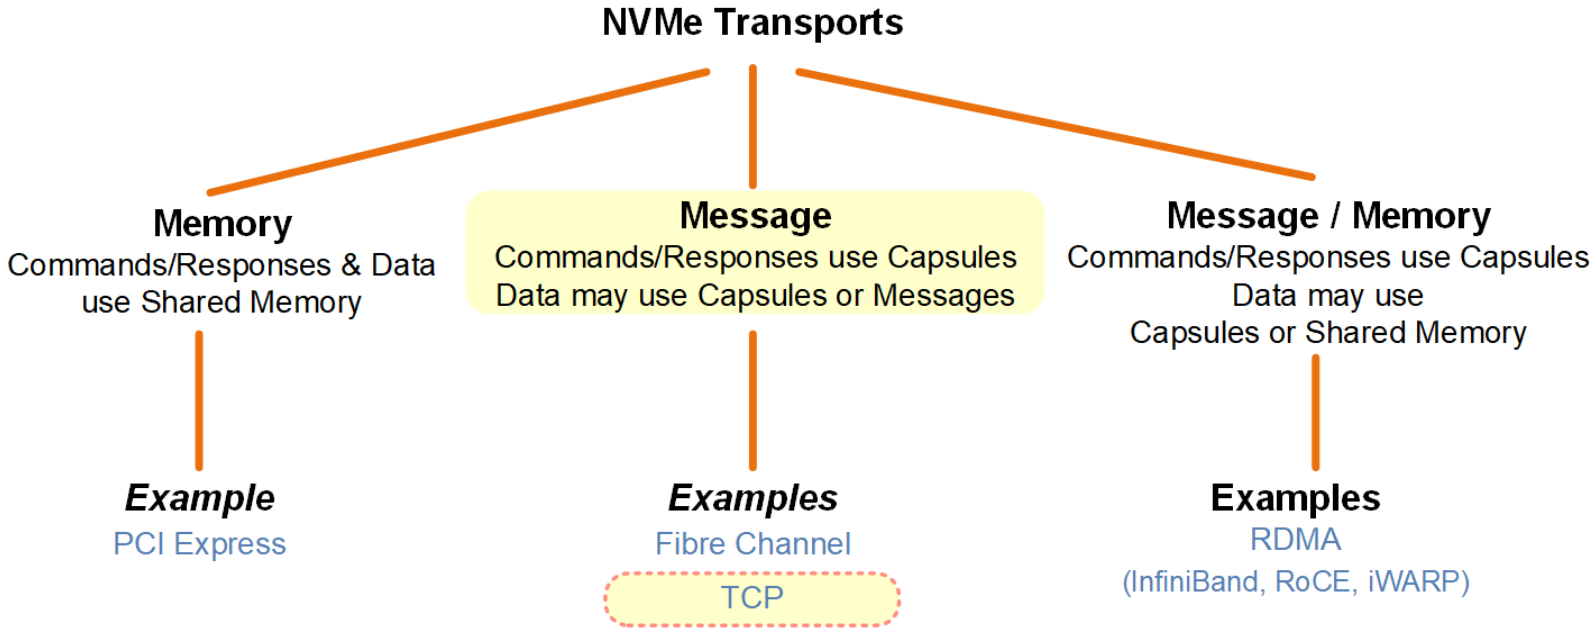

NVMe/TCP - Open Source Performance

Upstream Linux kernel NVMe™/TCP vs. NVMe/RDMA (added latency over direct attached PCIe® SSD)

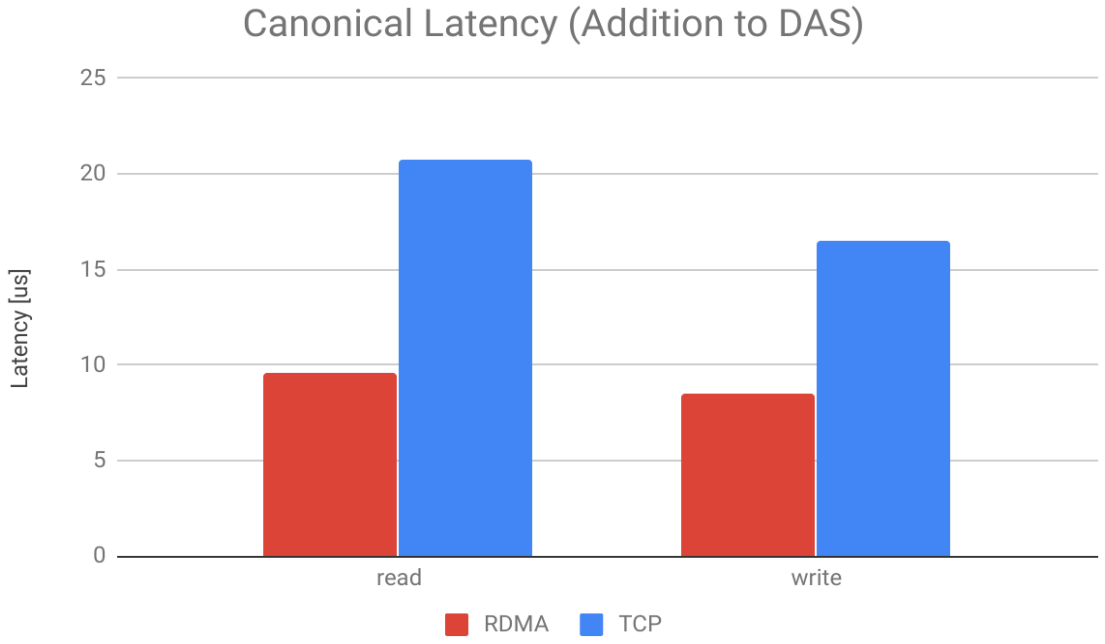

NVMe JBOFs

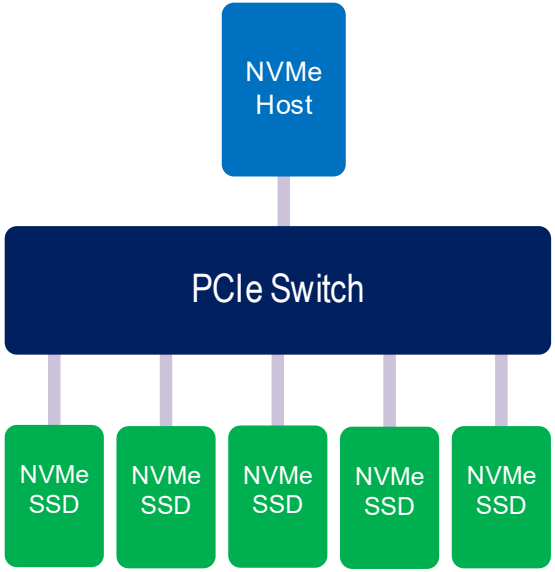

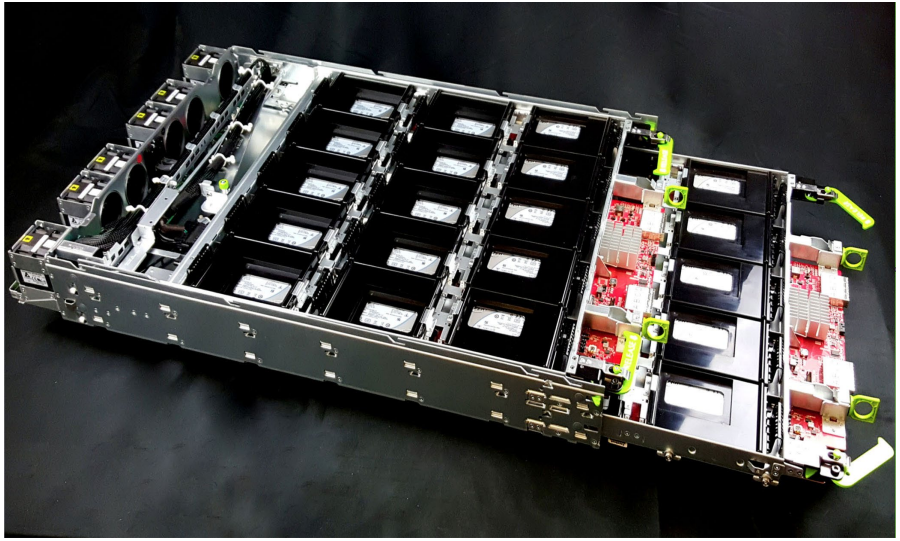

Facebook Lightning PCIe NVMe JBOF

NVMe Enclosure Management

- Native PCIe Enclosure Management (NPEM)
  - Submitted to the PCI-SIG Protocol Workgroup (PWG) on behalf of the NVMe™ Management Interface (NVMe-MI™) Workgroup
  - Approved by PCI-SIG on August 10th, 2017
  - Transport specific basic enclosure management
- SCSI Enclosure Services (SES) Based Enclosure Management
  - Technical proposal developed in the NVMe-MI workgroup
  - While the NVMe and SCSI architectures differ, the elements of an enclosure and capabilities to manage them are the same
    - Example enclosure elements: power supplies, fans, display or indicators, locks, temperature sensors, current sensors, voltage sensors, and ports
  - Comprehensive enclosure management for NVMe that leverages (SES), a standard developed by T10 for management of enclosures using the SCSI architecture

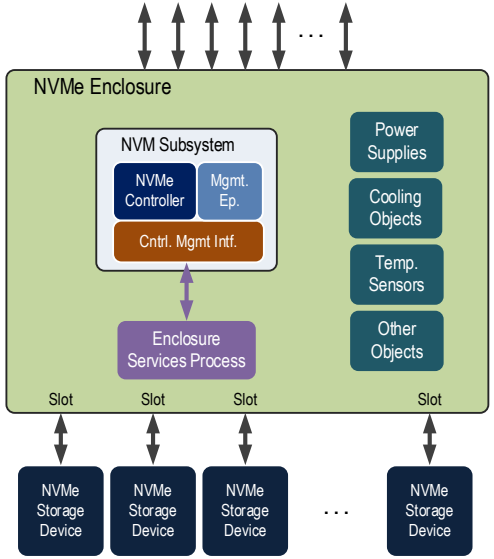

NVMe Storage Devices in NVMe-MI 1.0a

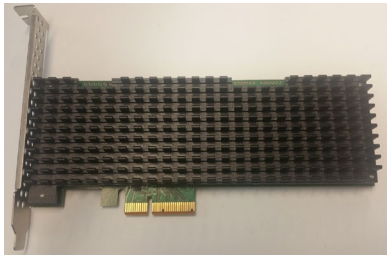

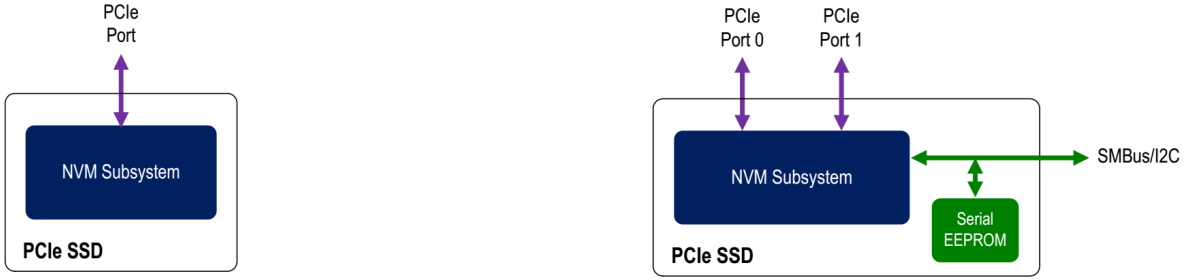

NVMe Storage Device – One NVM Subsystem with one or more ports, optional FRU Information Device, and an optional SMBus/I2C interface

NVMe Storage Devices in NVMe-MI 1.1

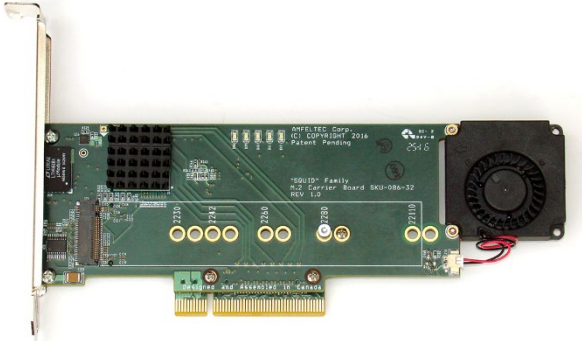

M.2 Carrier Board from Amfeltec

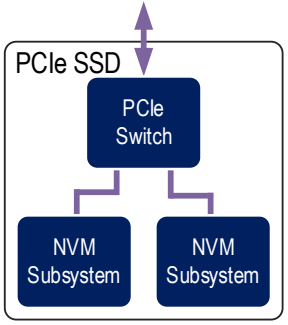

ANA Carrier Board from Facebook

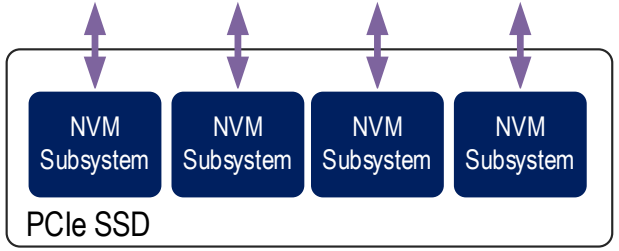

The Evolution of NVMe

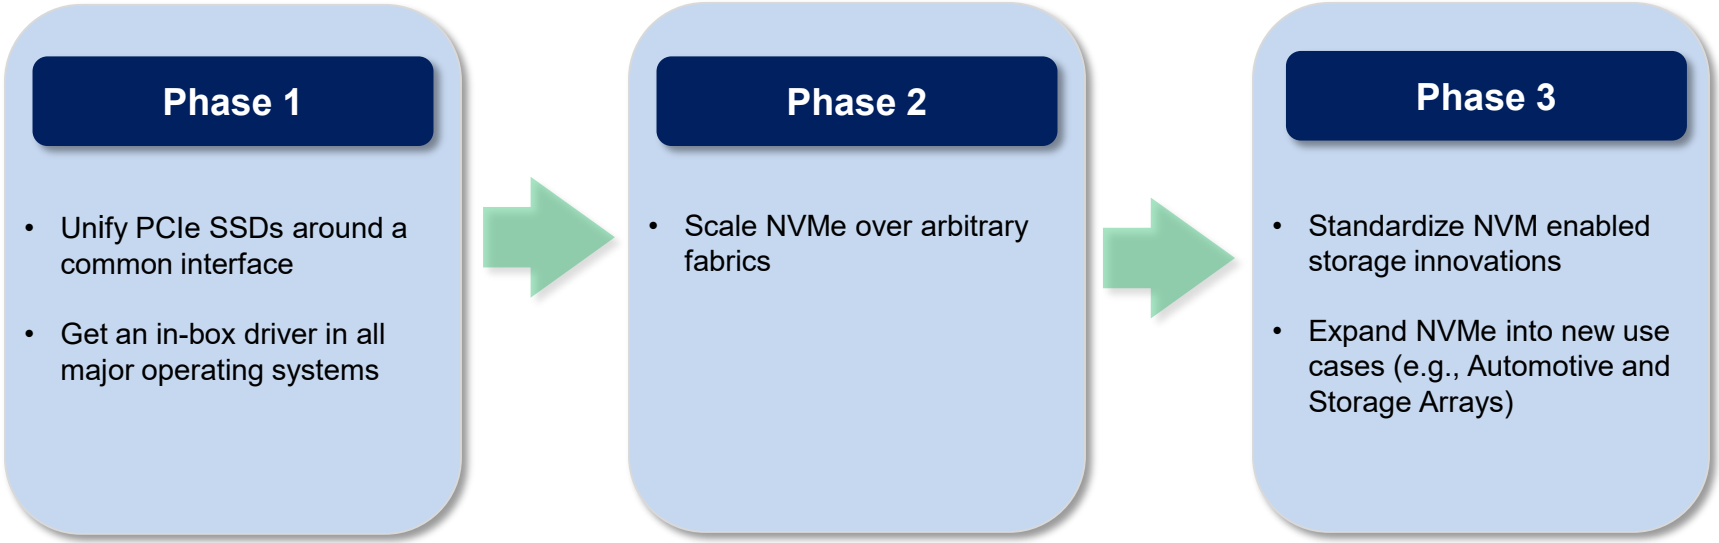

NVMe Continues to Drive Simplicity in A World of Complexity

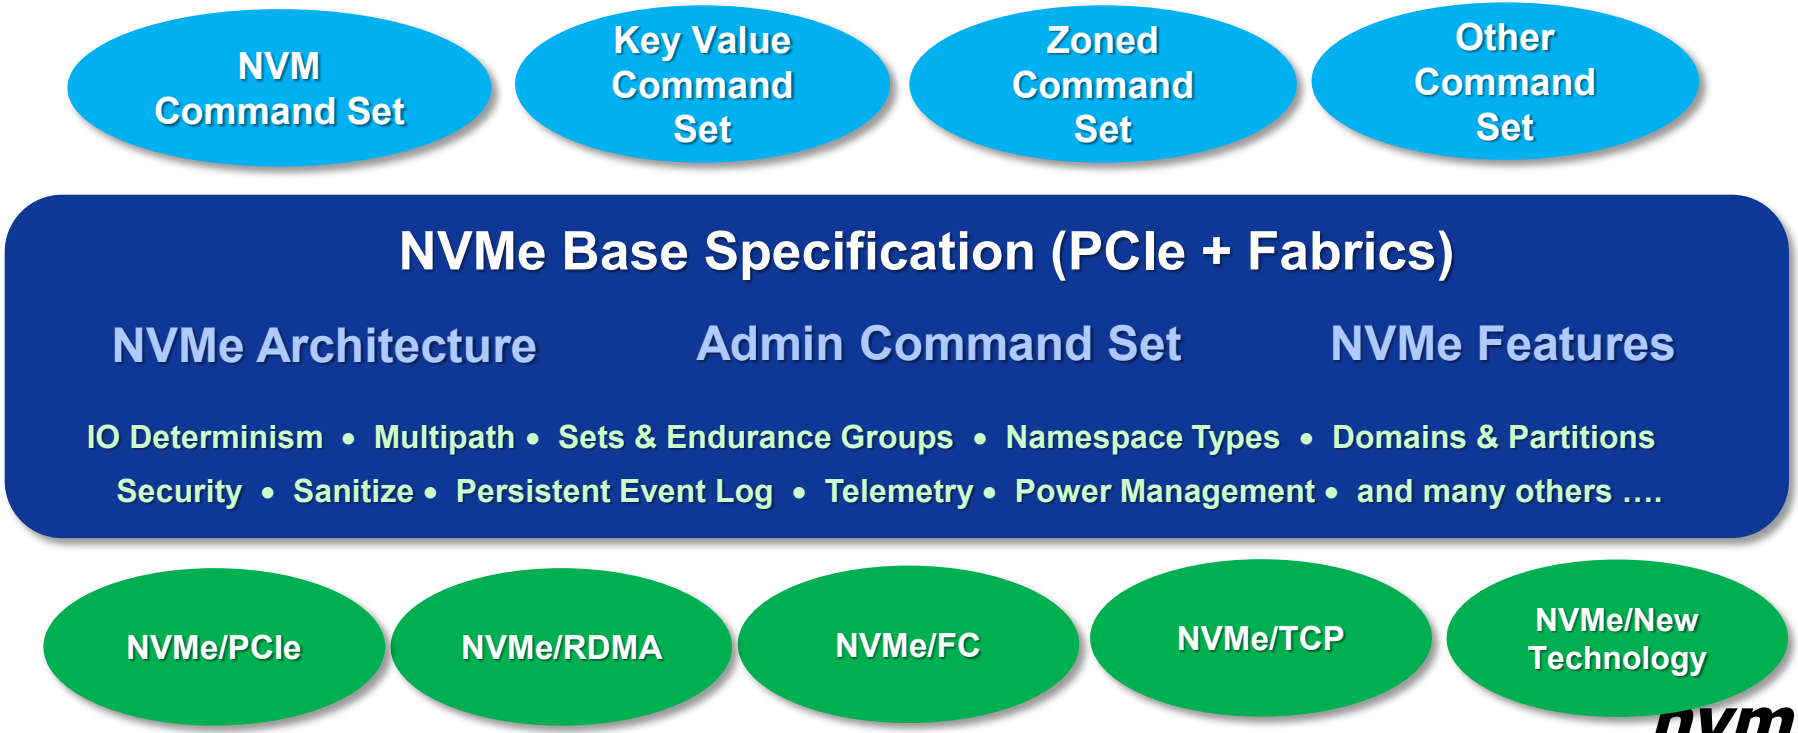

Evolution of the NVM Express Organization

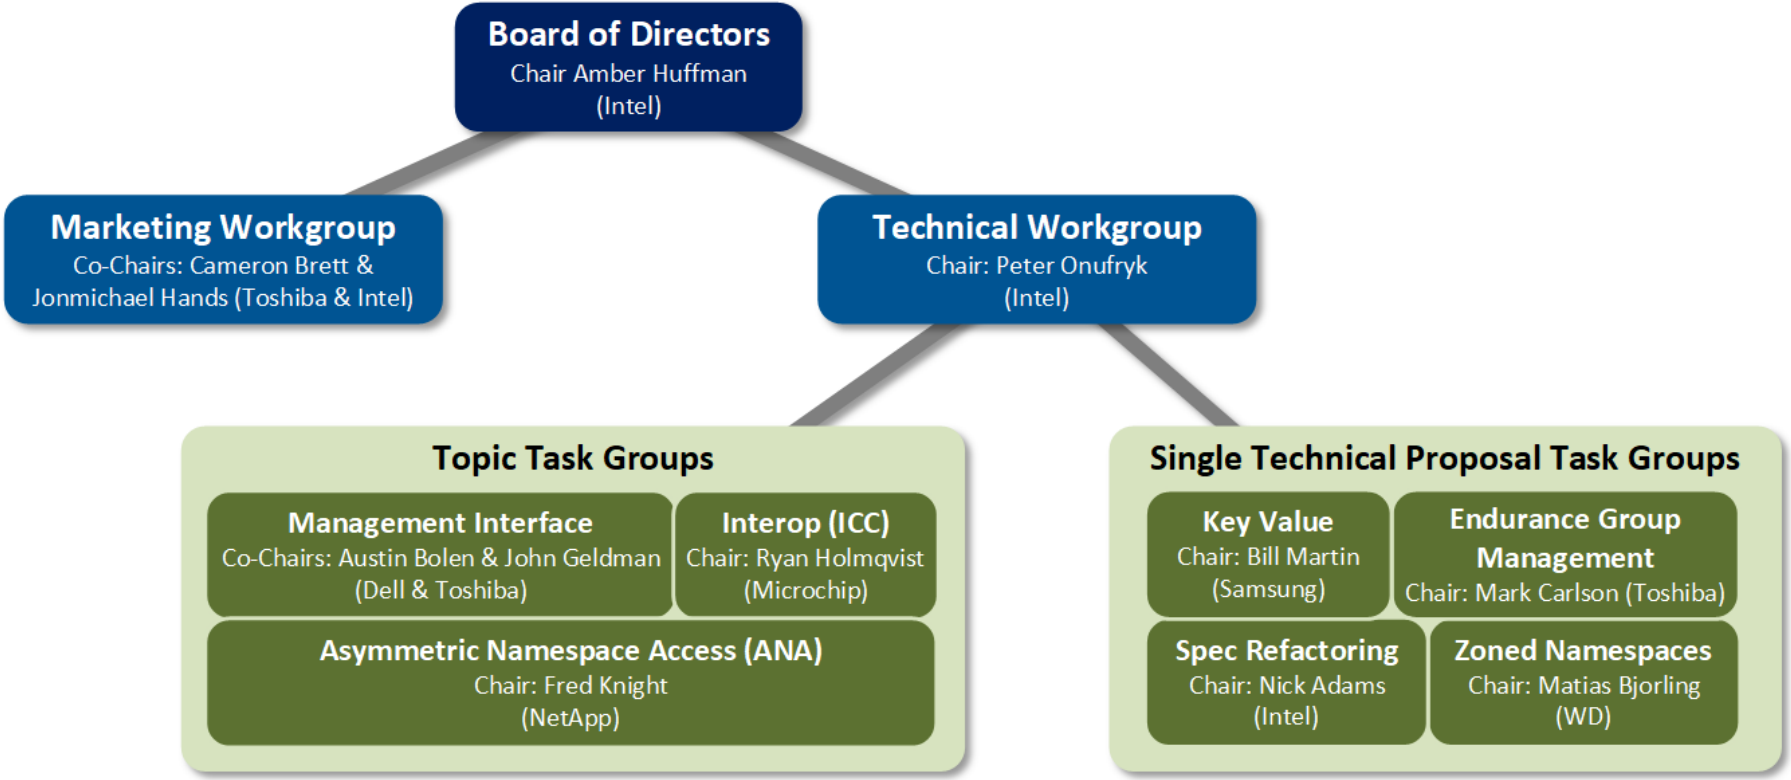

Increasing the Rate of Innovation Together with Greater Quality

- Formalized task groups with publicly published calendars and minutes
- Technical proposal phases with clear entries and exits
- Document repository with revision history
- Integrated draft specification always up to date
- Weekly electronic ballots

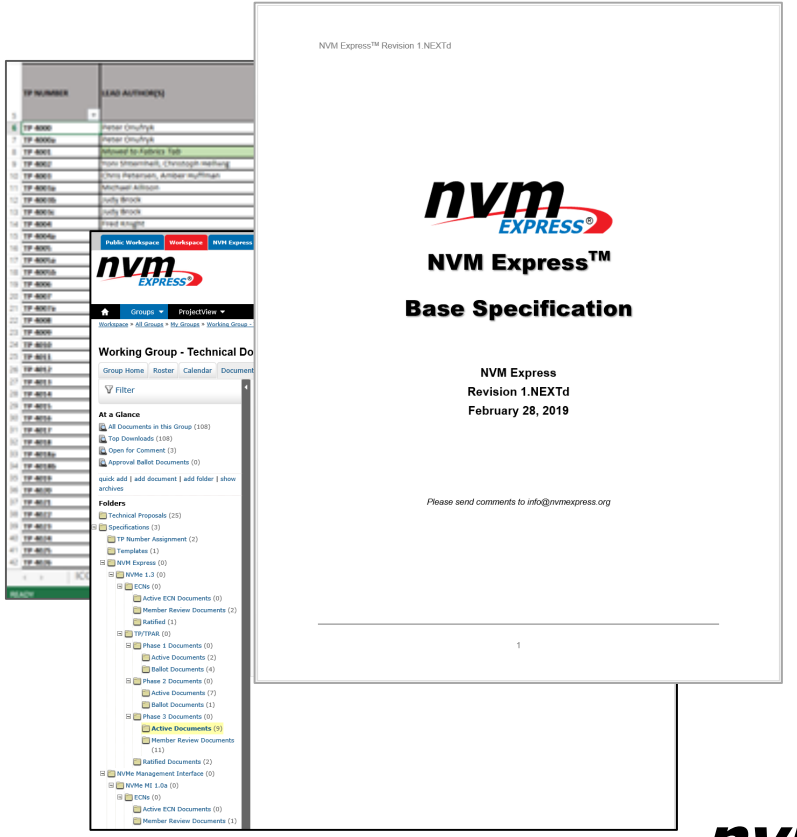

Summary

NVMe has unified client, cloud, and enterprise storage around a common command set and interface

The growth in NVMe adoption continues to accelerate

The NVMe organization has put in place processes and initiatives to support the increased rate of innovation enabled by NVM and new use cases

NVMe remains true to its core principles of simplicity and efficiency as it enters is second decade


nvm
EXPRESS®
Architected for Performance