NVMe® Technology

Powering the Connected Universe

Amber Huffman

Fellow & Chief Technologist of IP Engineering Group, Intel Corporation

President, NVM Express, Inc.

Agenda

- Fixing the Memory & Storage Hierarchy
- Refactoring for the Next Decade of Growth
- NVMe® Architecture Advancements
- What's Next: Computational Storage

Agenda

- **Fixing the Memory & Storage Hierarchy**
- Refactoring for the Next Decade of Growth
- NVMe® Architecture Advancements
- What's Next: Computational Storage

Memory PILLAR

WHAT WE SAID IN 2018

Exponential advances in all levels of memory hierarchy are needed to match the ever increasing compute demand

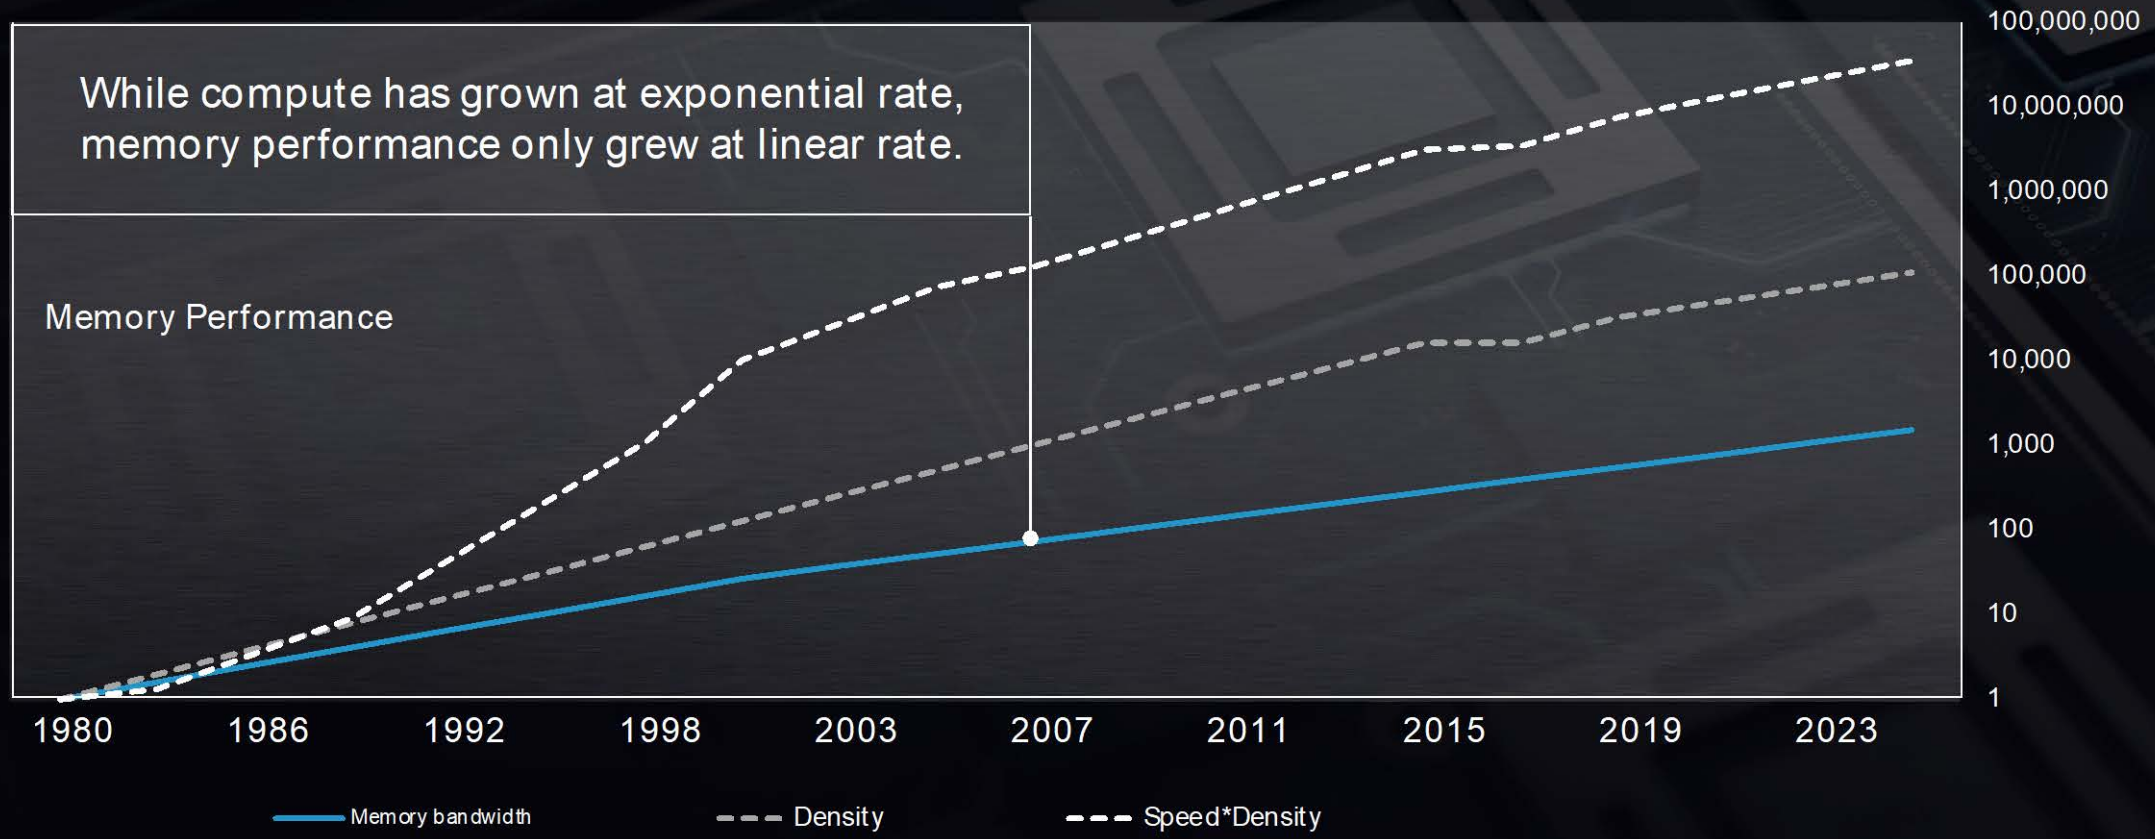

Memory and Storage Hierarchy Gaps

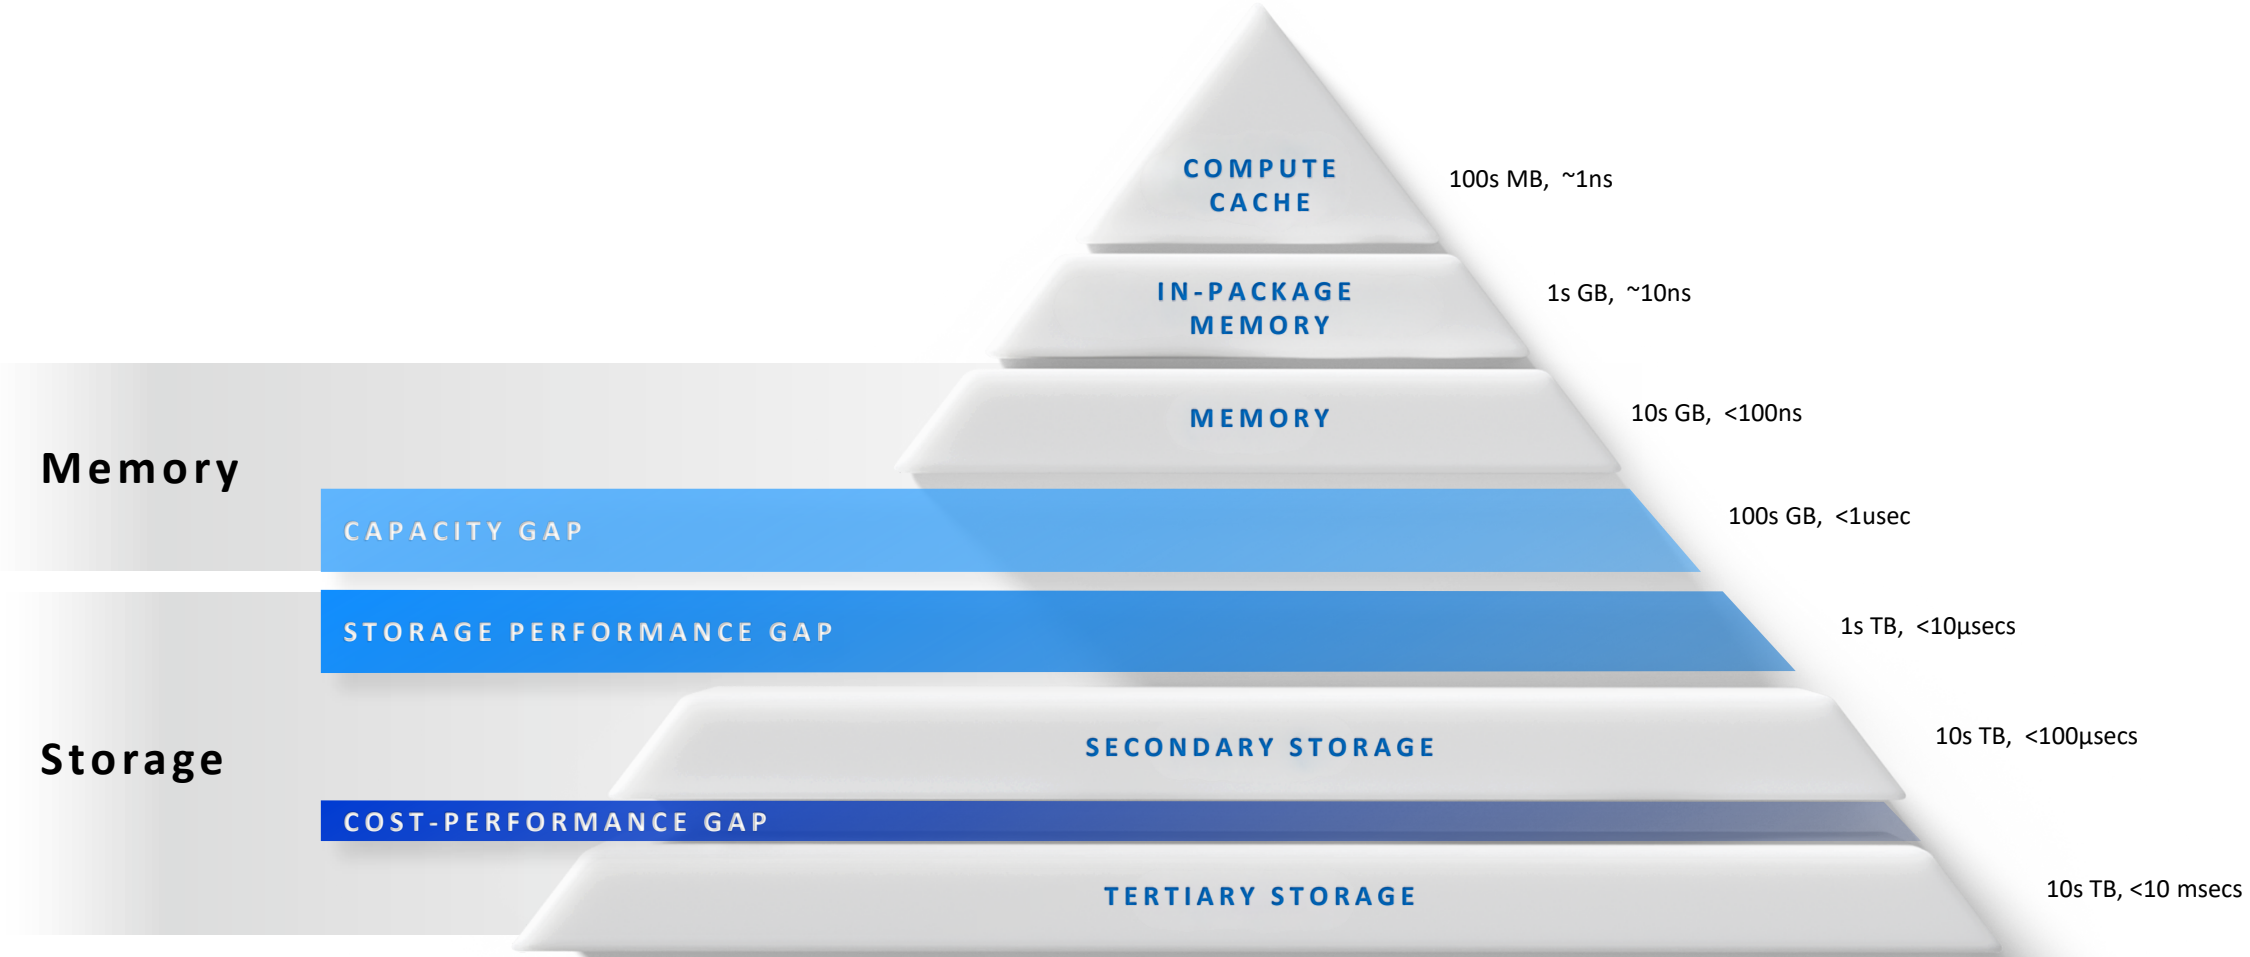

Memory and Storage Hierarchy Gaps

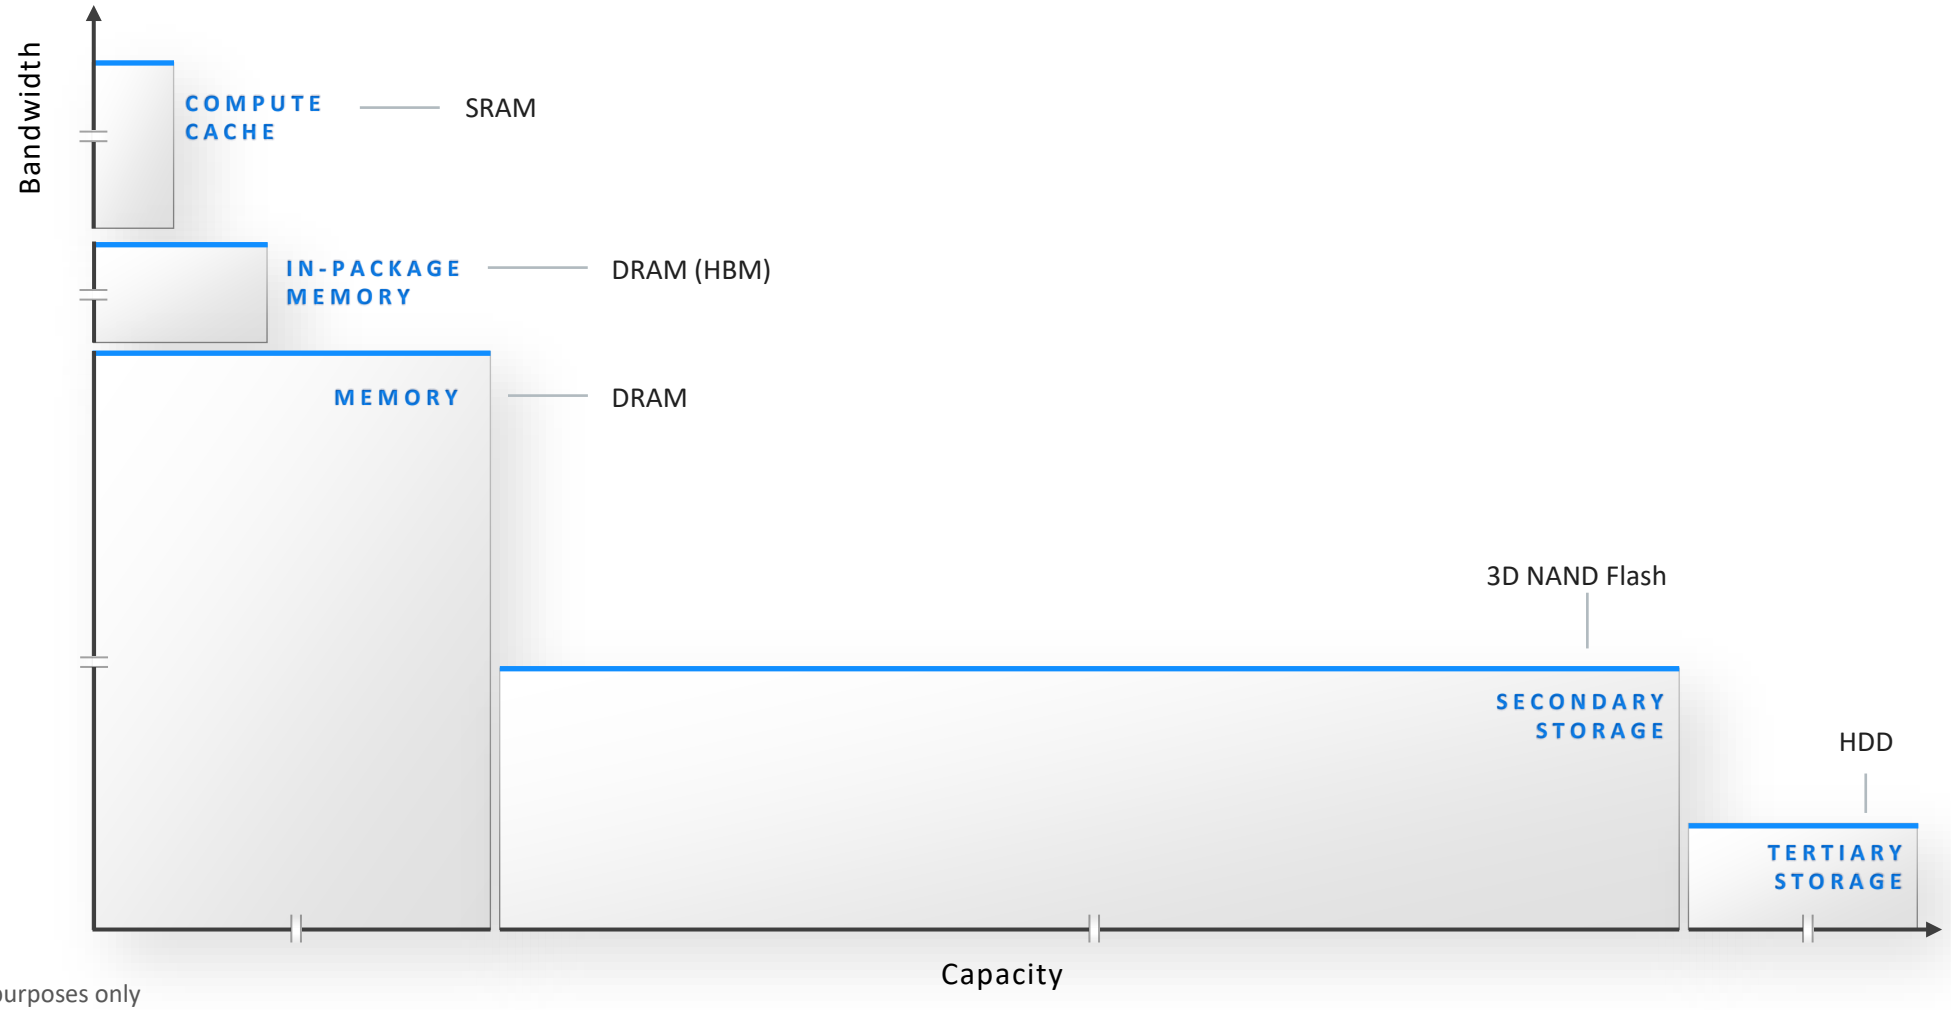

For illustrative purposes only

Types of Memory

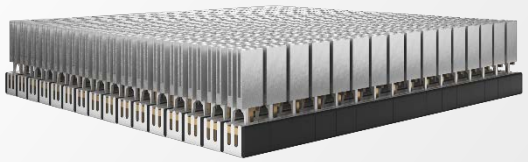

DRAM

One DRAM memory cell = 1 bit

3D XPOINT

One 3D XPoint memory cell = 1 bit

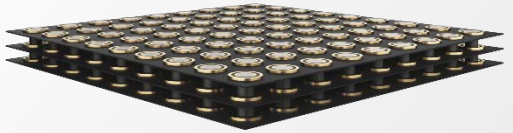

INTEL 3D NAND

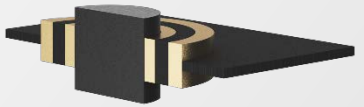

One 3D NAND memory cell = 1-4 bits

Types of Memory Compared

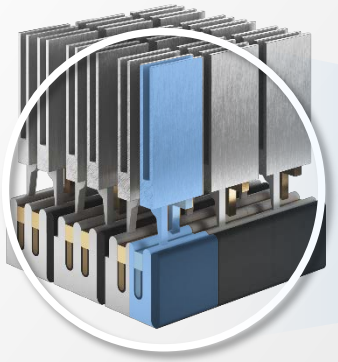

DRAM

One DRAM memory cell = 1 bit

3D XPOINT

One 3D XPoint memory cell = 1 bit

INTEL 3D NAND

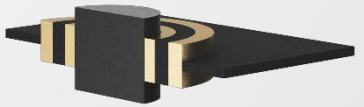

One 3D NAND memory cell = 1-4 bits

3D NAND Roadmap

3D NAND Roadmap

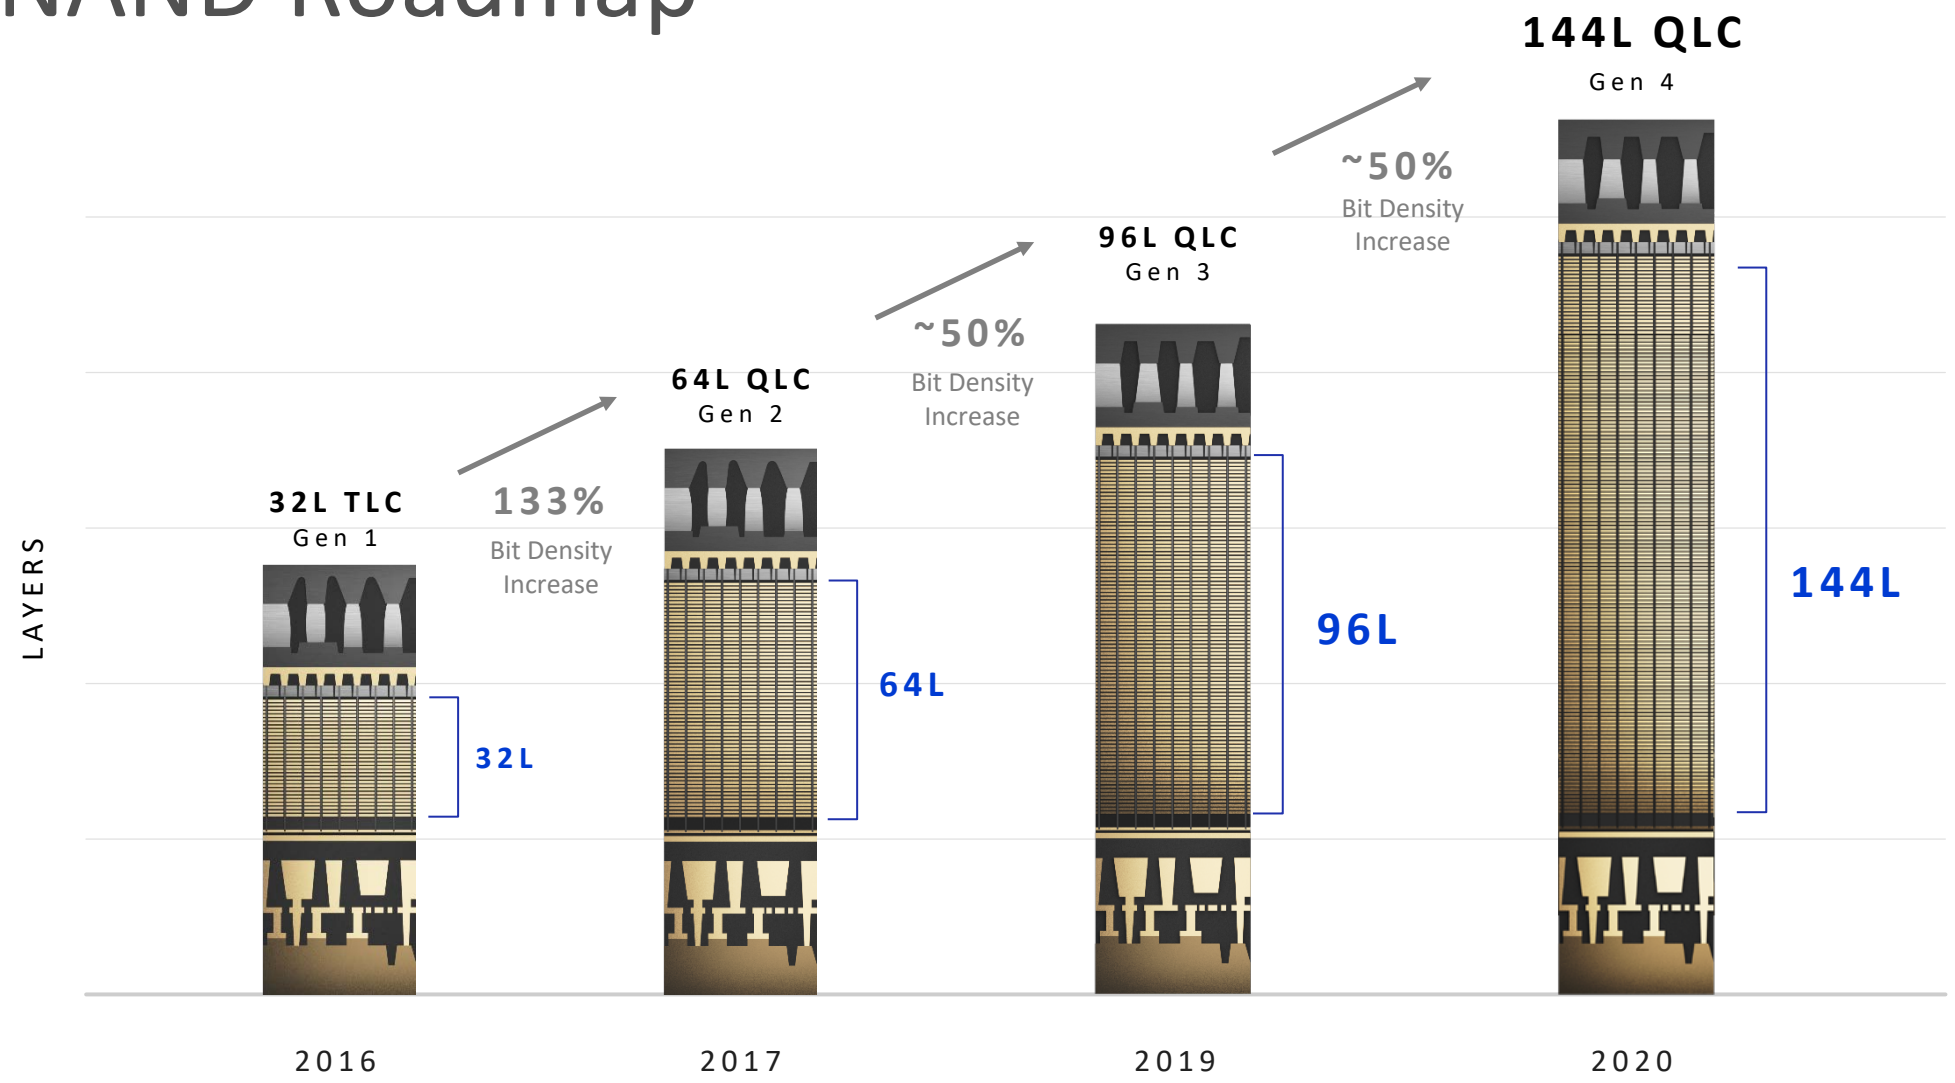

3D XPoint Memory Roadmap

1st Gen
2-Deck

2nd Gen
4-Deck

Multiple Millions of IOPS
on 2nd Generation Intel® Optane™ SSD

2017

2020*

*Target Production

Media Innovation Realized in Product

Innovation Powered by

nvm EXPRESS®

Fast

Simple

Scalable

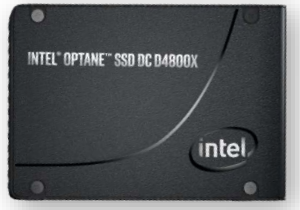

Intel Optane™ SSD

Intel Optane ™ Memory H10

Optane + 3D QLC NAND

Intel 3D NAND SSD

Memory and Storage Hierarchy Gaps

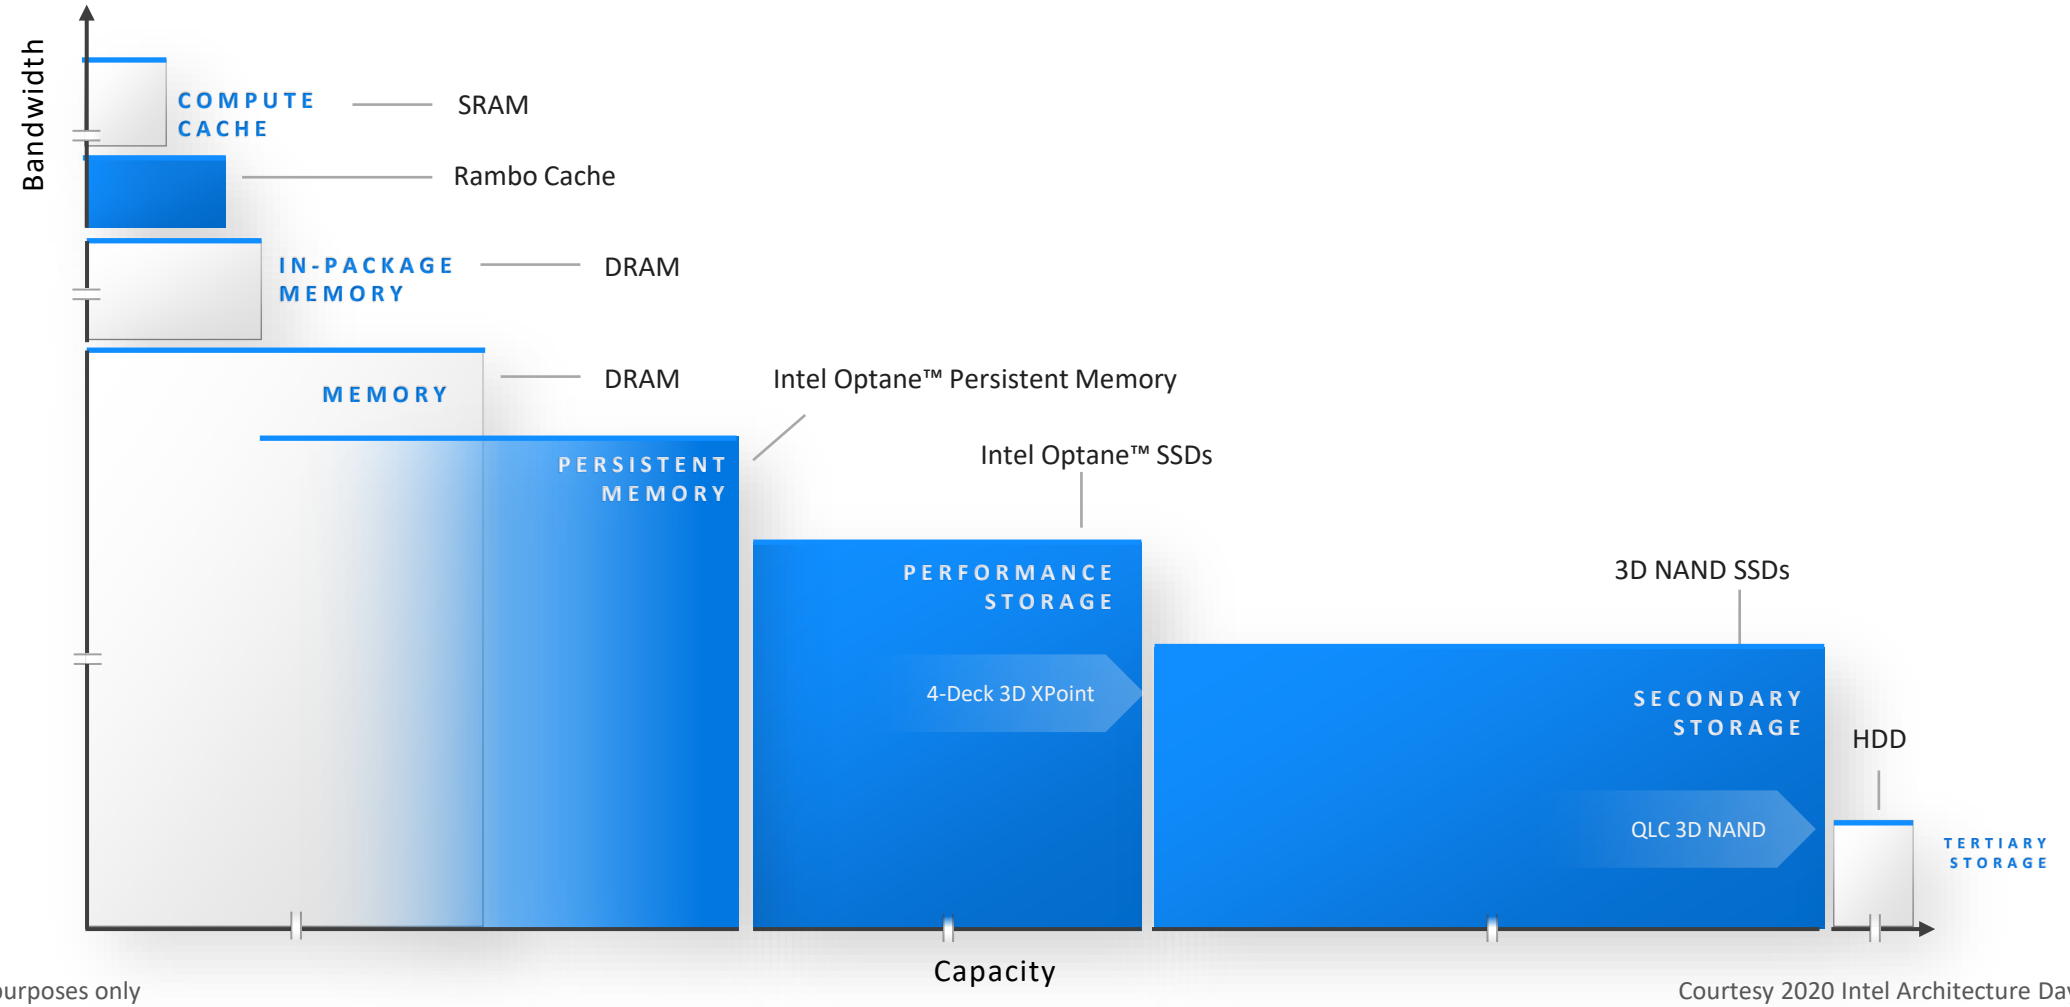

For illustrative purposes only

Courtesy 2020 Intel Architecture Day

Memory and Storage Hierarchy ~~Gaps~~ Solutions

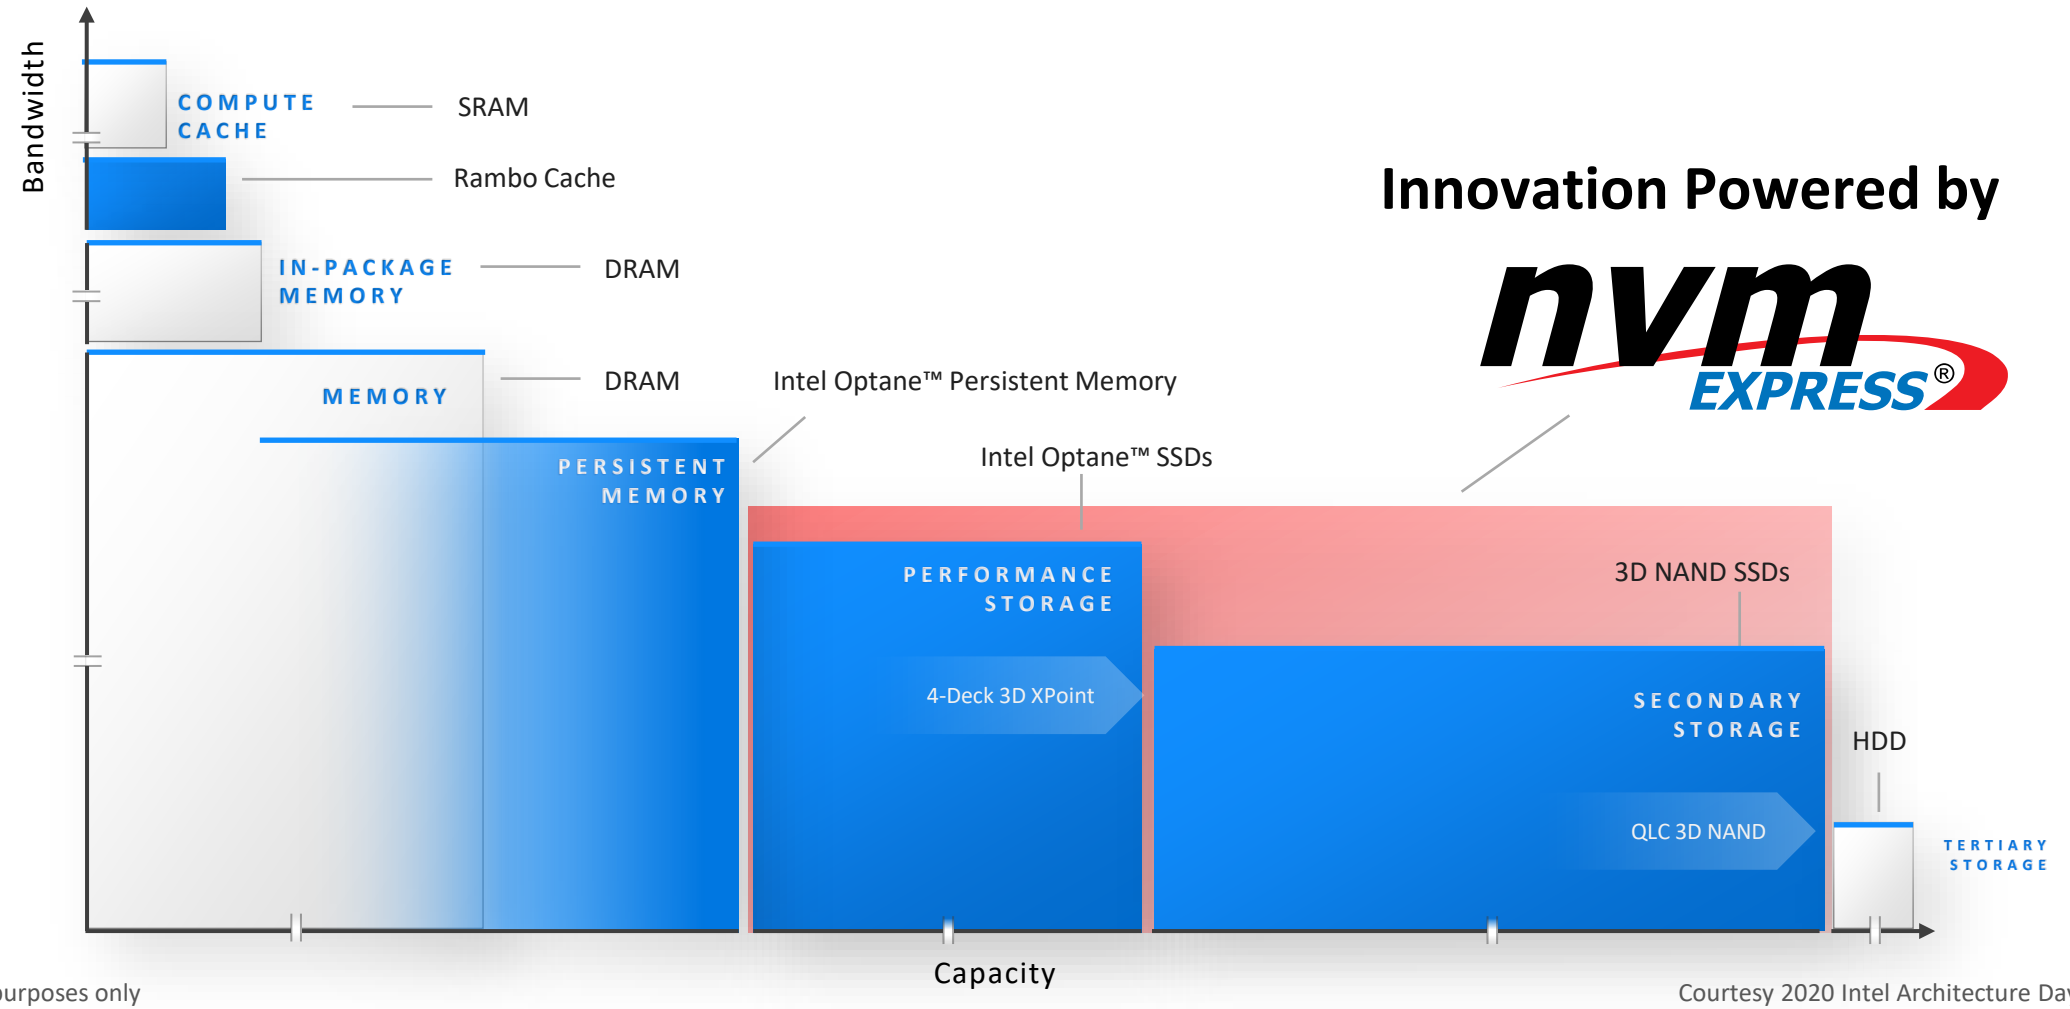

Innovation Powered by

nvm EXPRESS®

For illustrative purposes only

Courtesy 2020 Intel Architecture Day

NVMe® Technology Powers the Connected Universe

Units (Ku)	2016	2017	2018	2019	2020*	2021*
Enterprise	364	749	1,069	2,045	4,067	5,554
Cloud	2,051	3,861	10,369	12,276	18,982	21,999
Client	33,128	48,951	82,587	143,236	202,348	258,791

* Projections provided by Forward Insights Q2'20

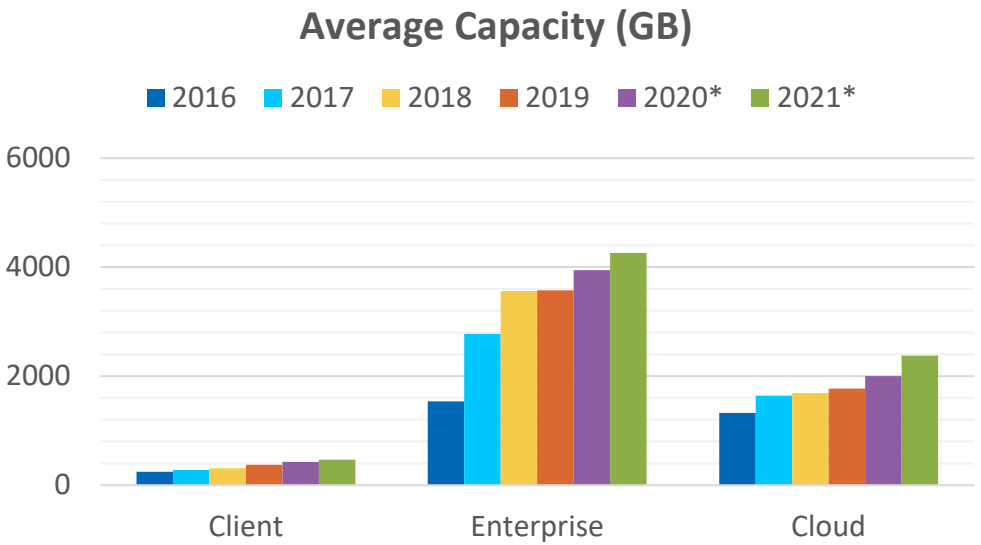

- NVMe technology grew from 3 Petabytes to 29 PB shipped per year from 2016 to 2019
- For 2020, the projection is 54 PB
- NVMe technology demand projected to remain strong in a post COVID world

Agenda

- Fixing the Memory & Storage Hierarchy
- **Refactoring for the Next Decade of Growth**
- NVMe® Architecture Advancements
- What's Next: Computational Storage

The Evolution of NVMe® Technology

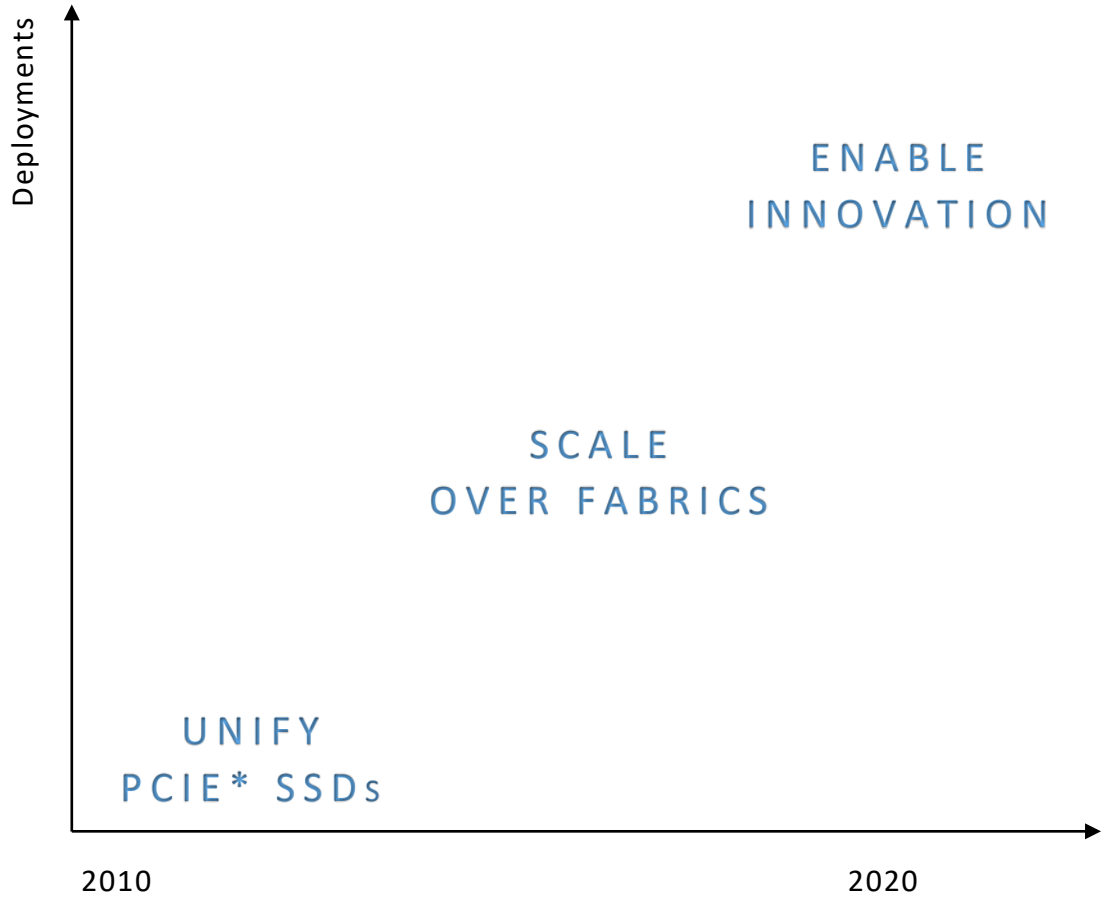

Driving Simplicity in a World of Complexity

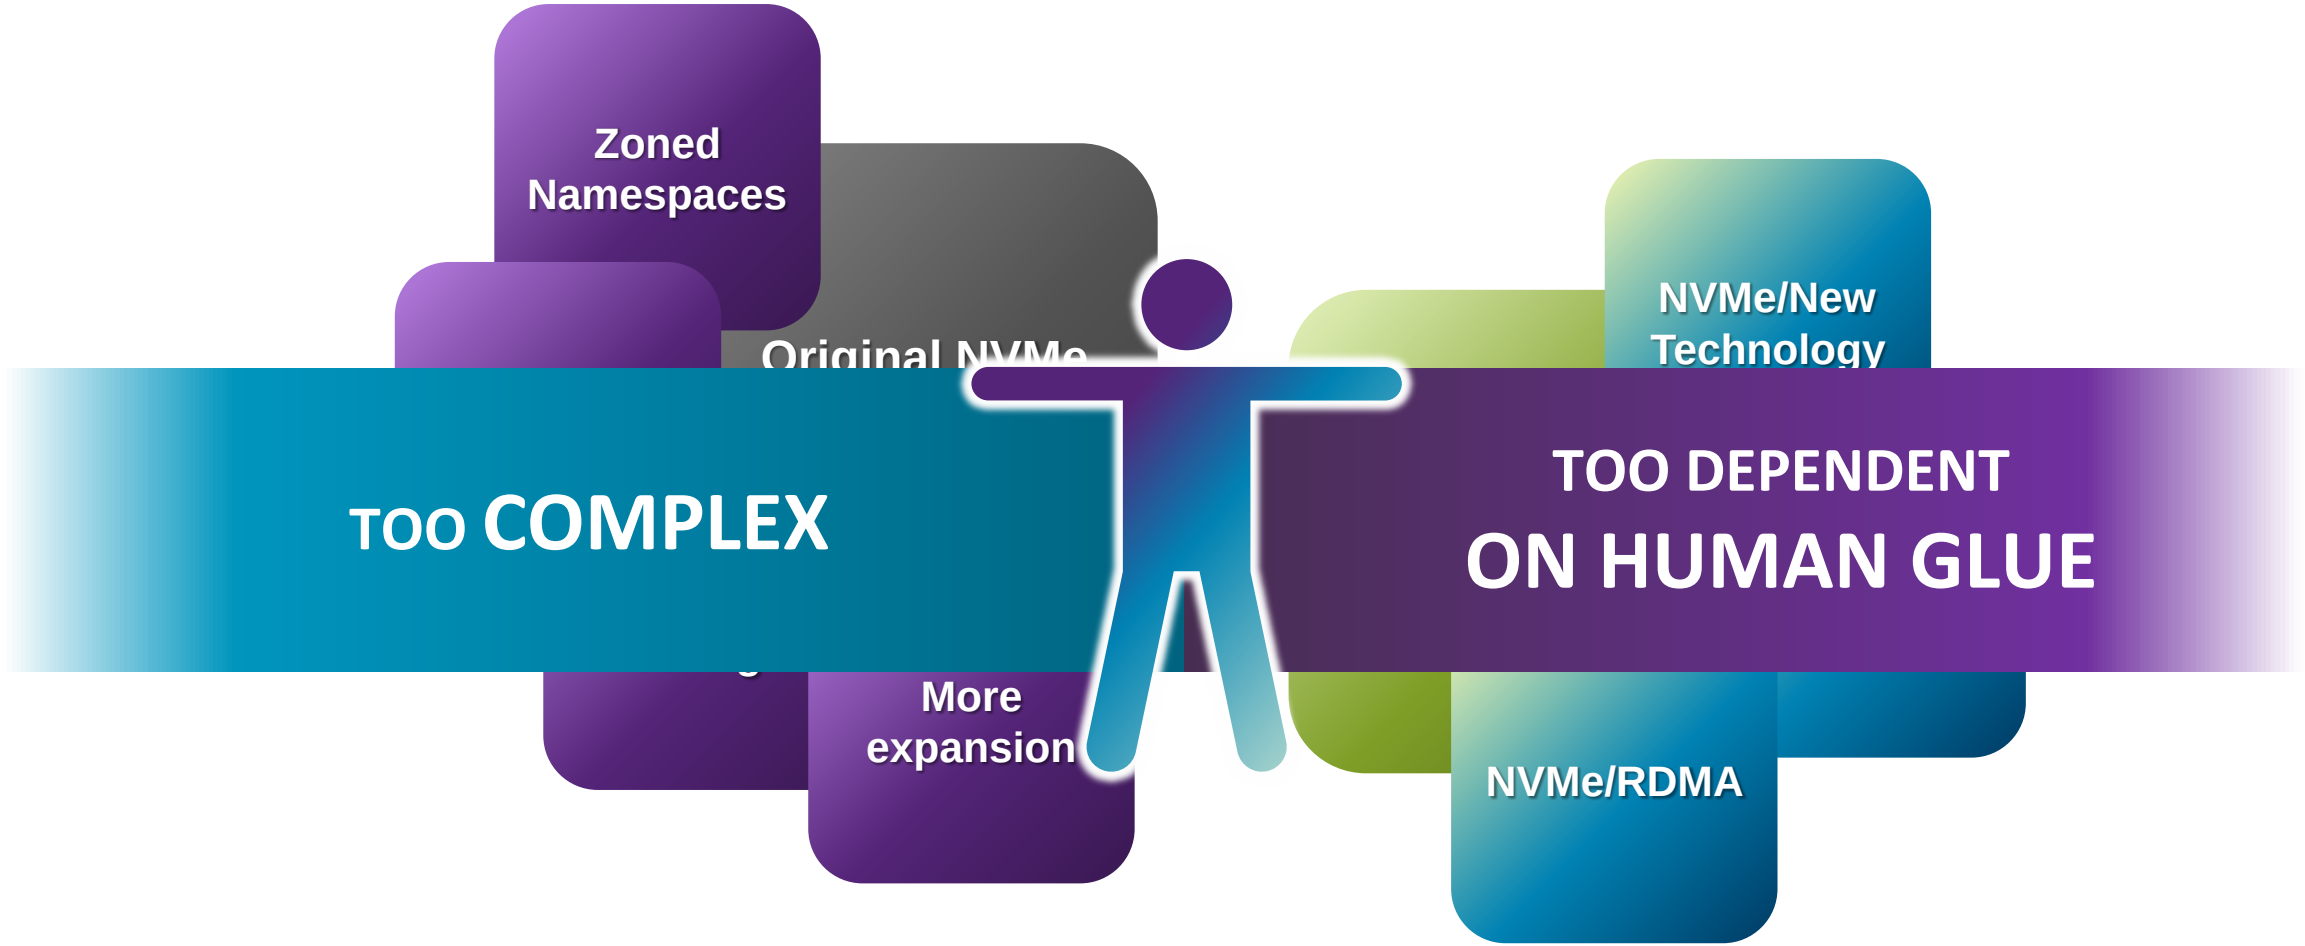

Driving Simplicity in a World of Complexity

- Back to core values… Fast, Simple, Scalable
- Foster areas of innovation AND avoid impact to broadly deployed solutions
- Create an extensible infrastructure that will take us through the next decade of growth

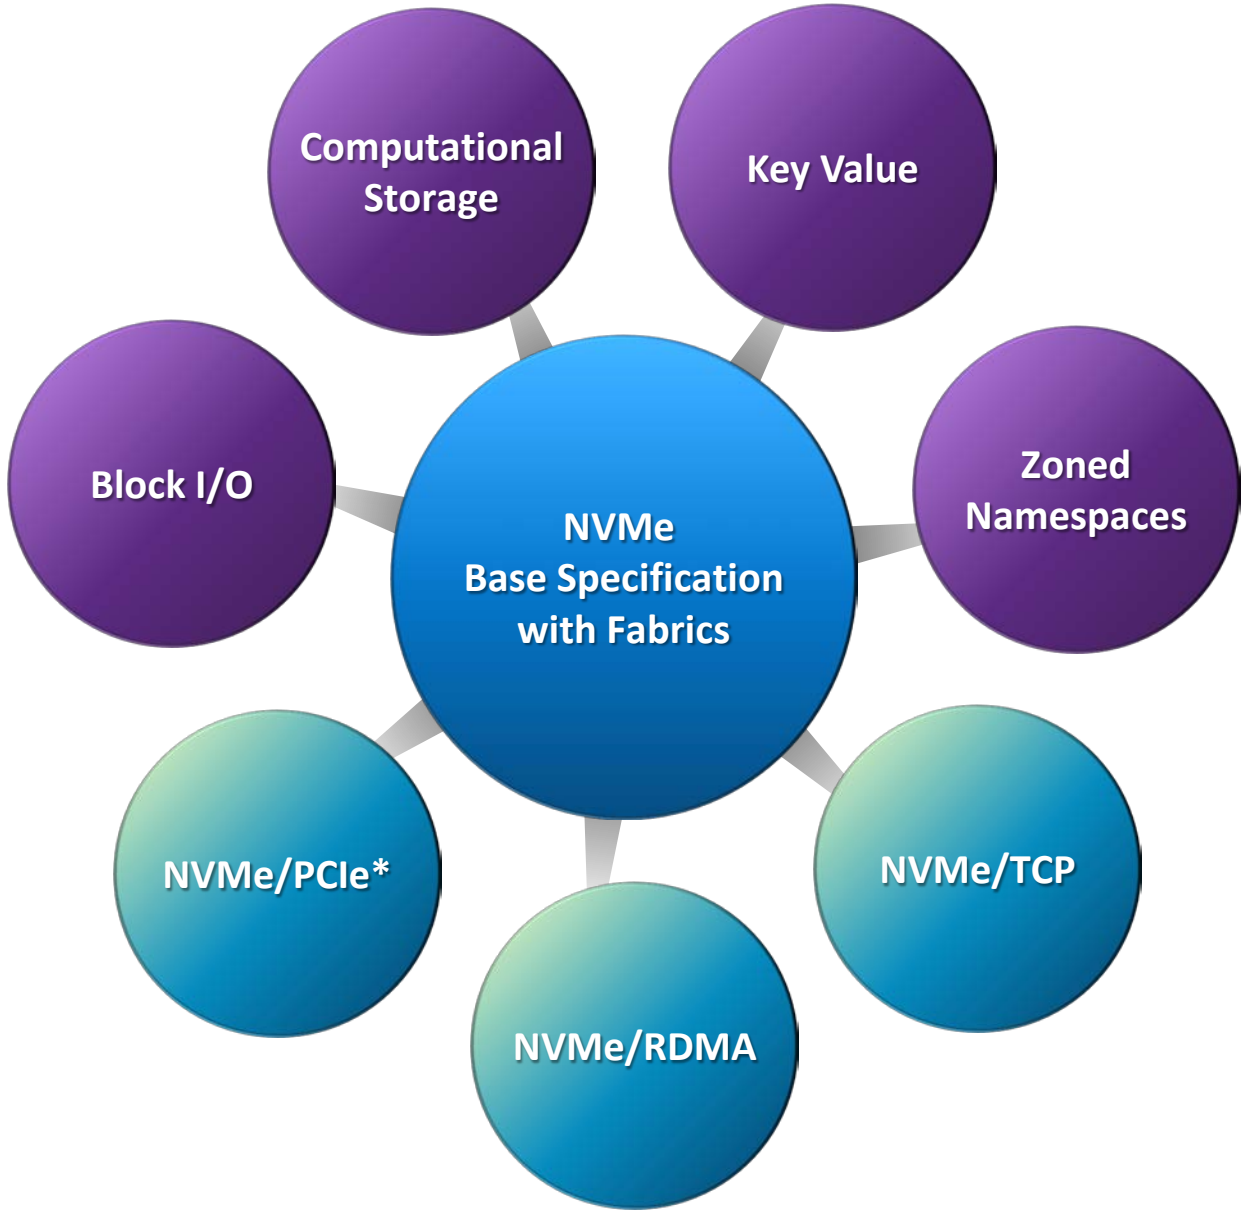

Specification Families

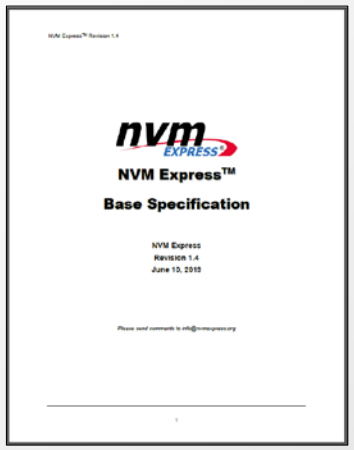

THE CORE

COMMAND SETS

TRANSPORT LAYERS

- The core of NVMe and NVMe over Fabrics integrated into a base specification
- Modular command set specifications (Block, Zoned Namespaces, Key Value, etc)
- Modular transport layer specifications (PCI Express*, RDMA, TCP)
- Maintain Management Interface as separate modular specification

Agenda

- Fixing the Memory & Storage Hierarchy
- Refactoring for the Next Decade of Growth
- **NVMe Architecture® Advancements**
- What's Next: Computational Storage

NVM Express Technology Specification Roadmap

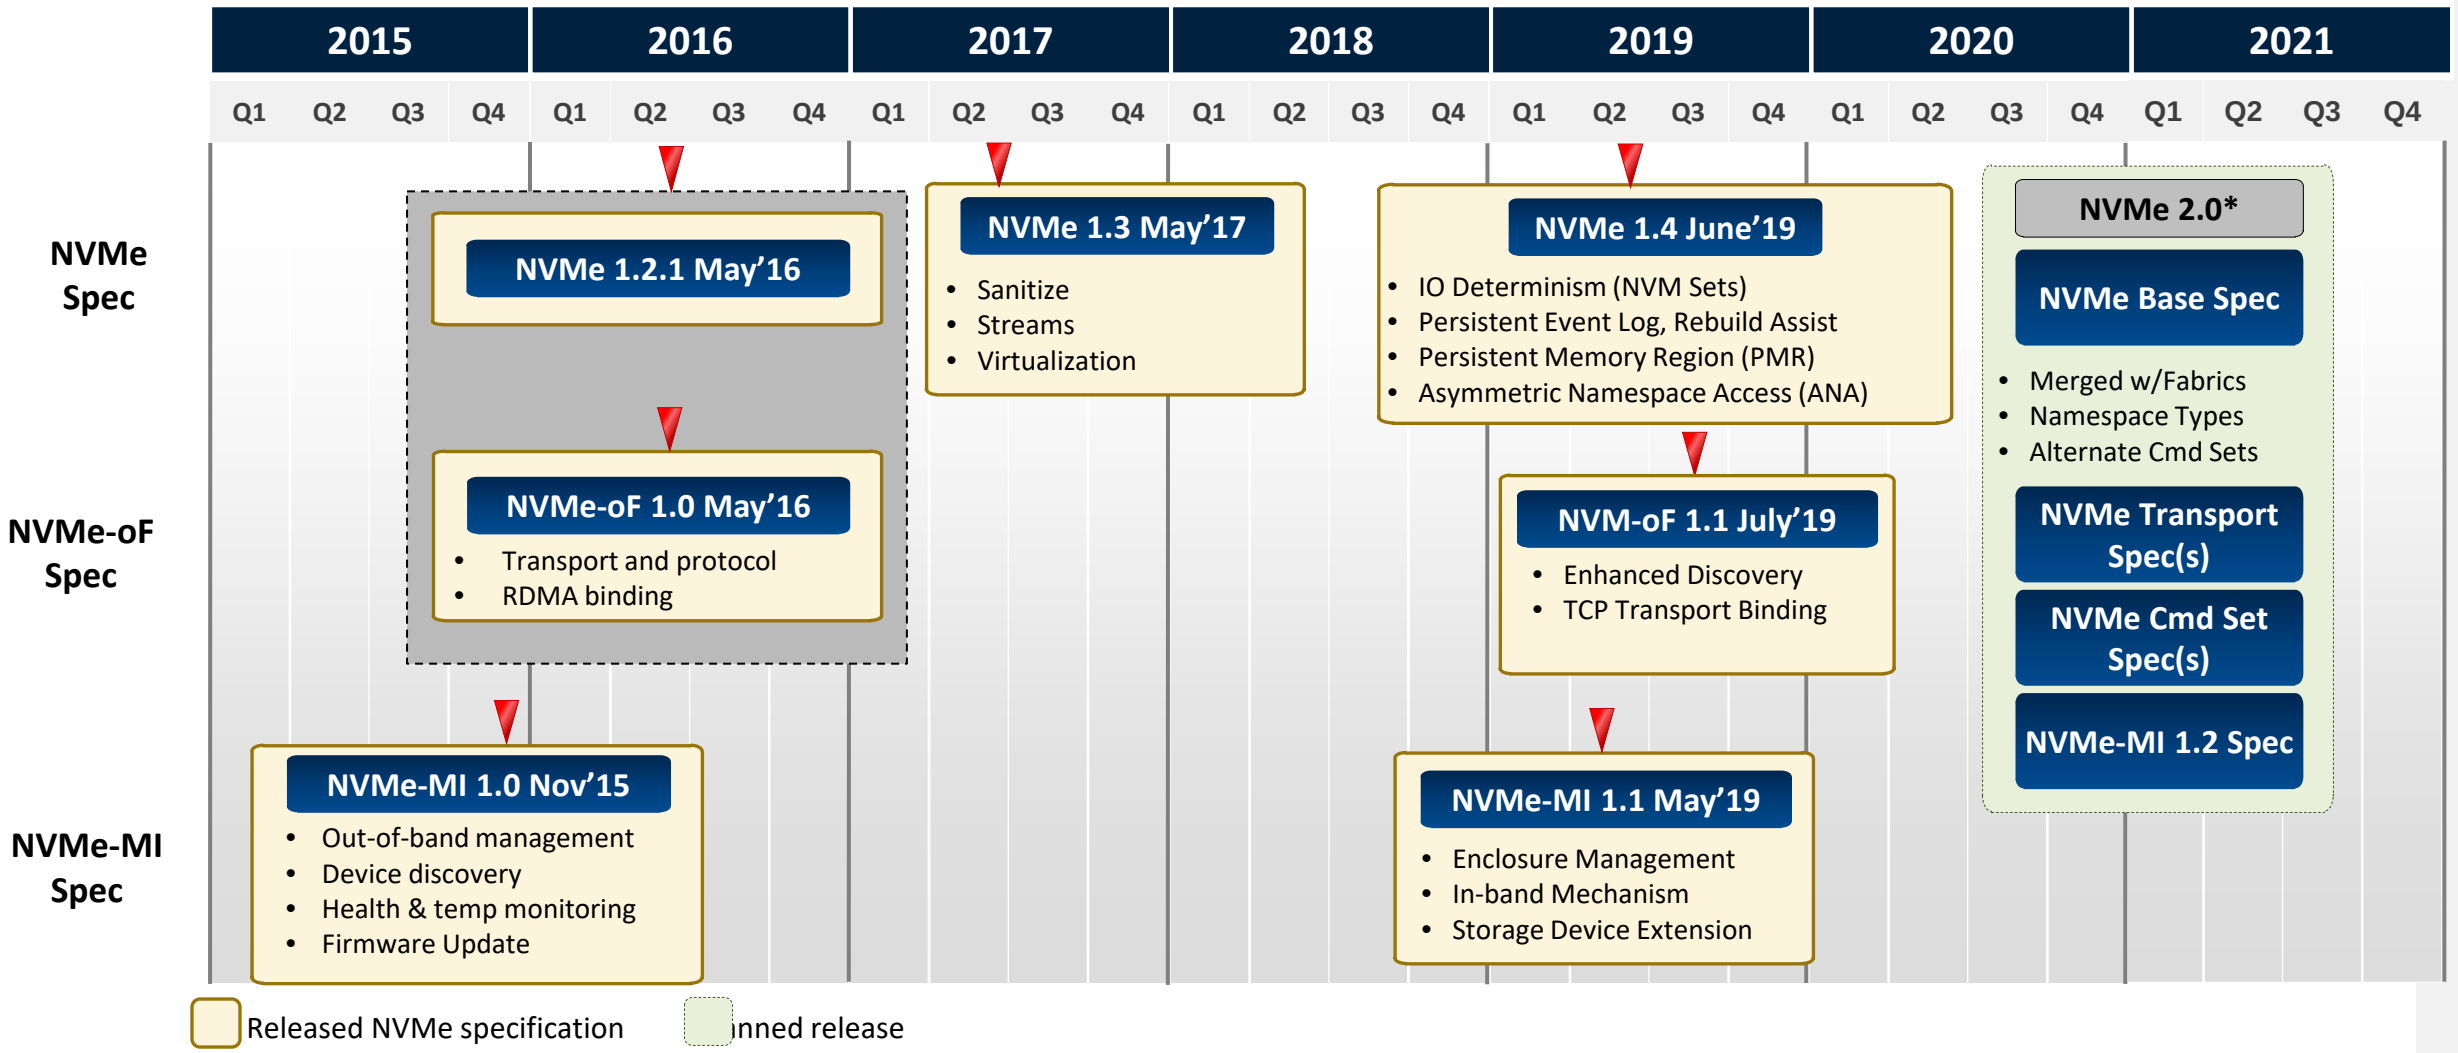

Namespace Types Enable Alternate Command Sets

- Discover Command Set supported for a Namespace
- Command Sets: Block I/O, Key Value, Zoned, < future >
- Extensible approach for future innovation

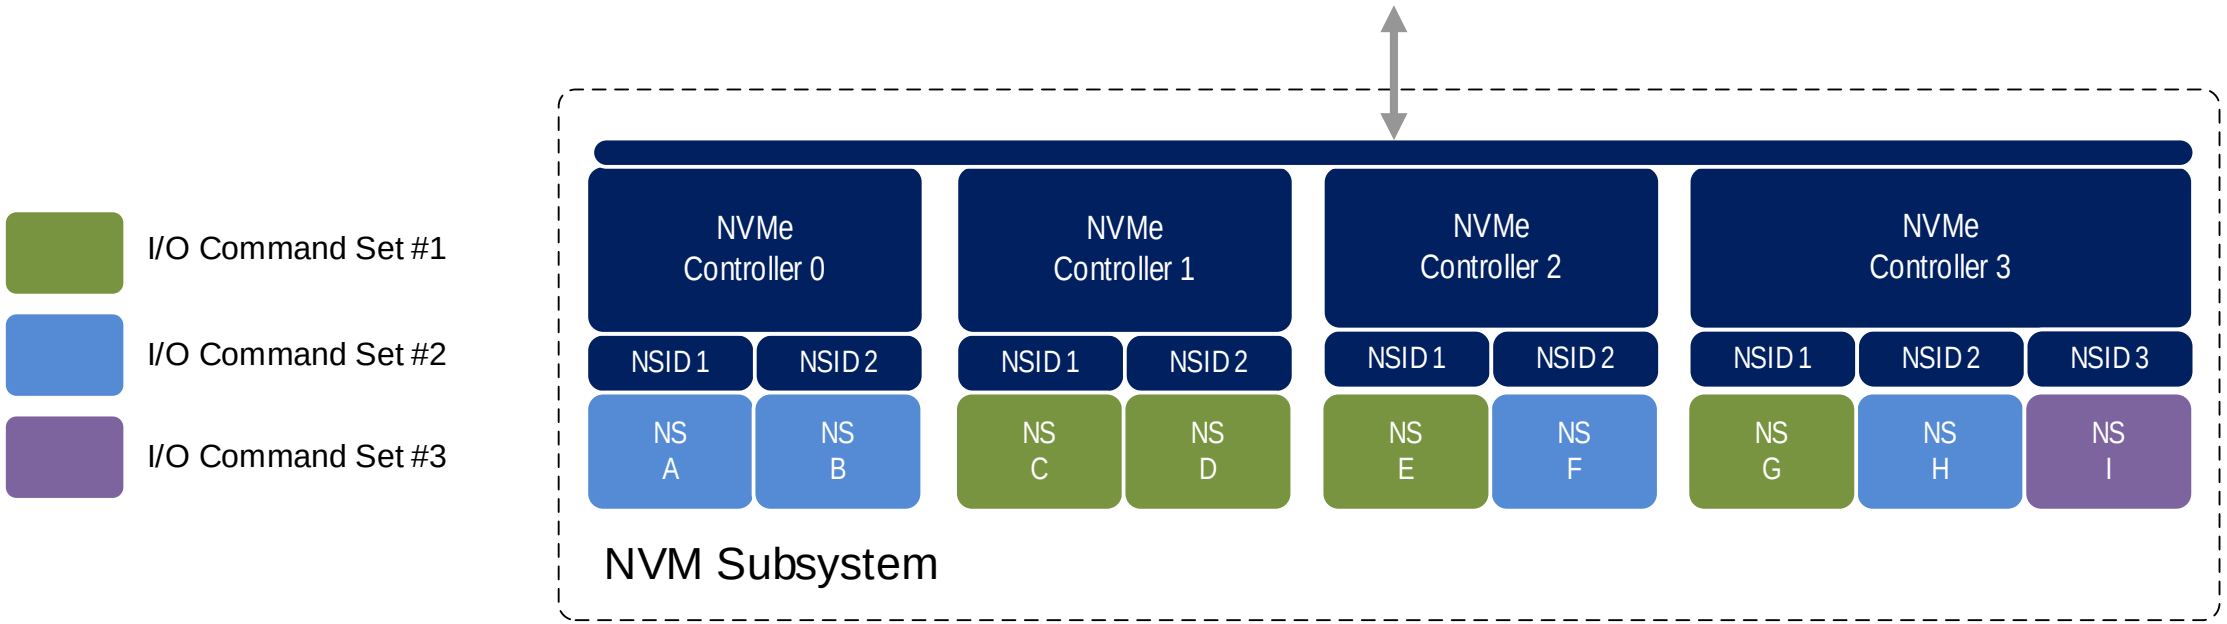

Zoned Namespaces

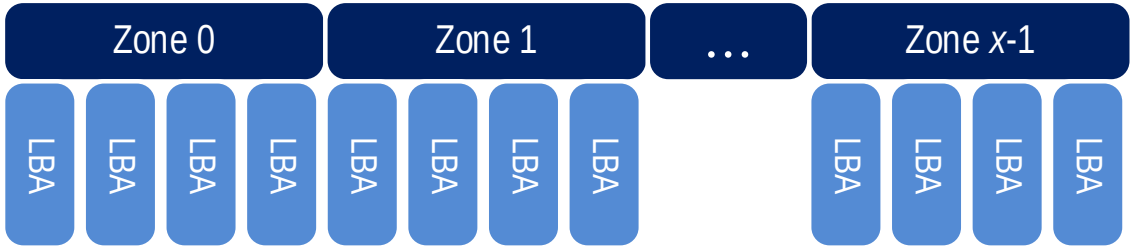

- NVMe® technology evolving to address underlying media changes with larger Erase Blocks and more
- Zoned Namespaces require Logical Blocks to be written sequentially in a Zone
- Reduces write amplification and overprovisioning

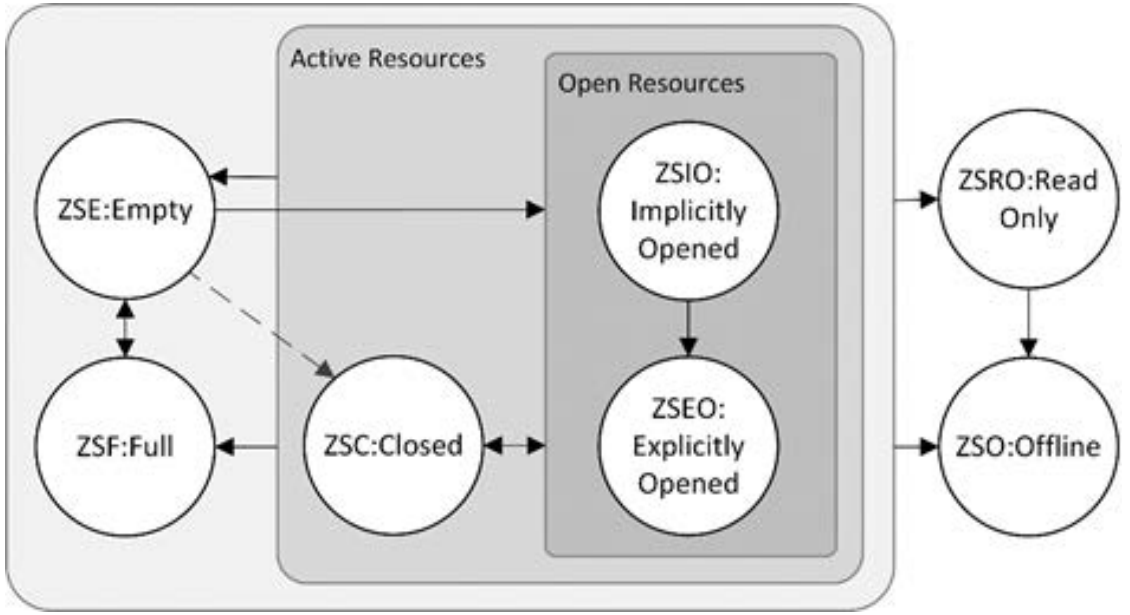

Endurance Groups

- Flexible capacity management model
- Create appropriate groupings based on access pattern, media type(s), and more

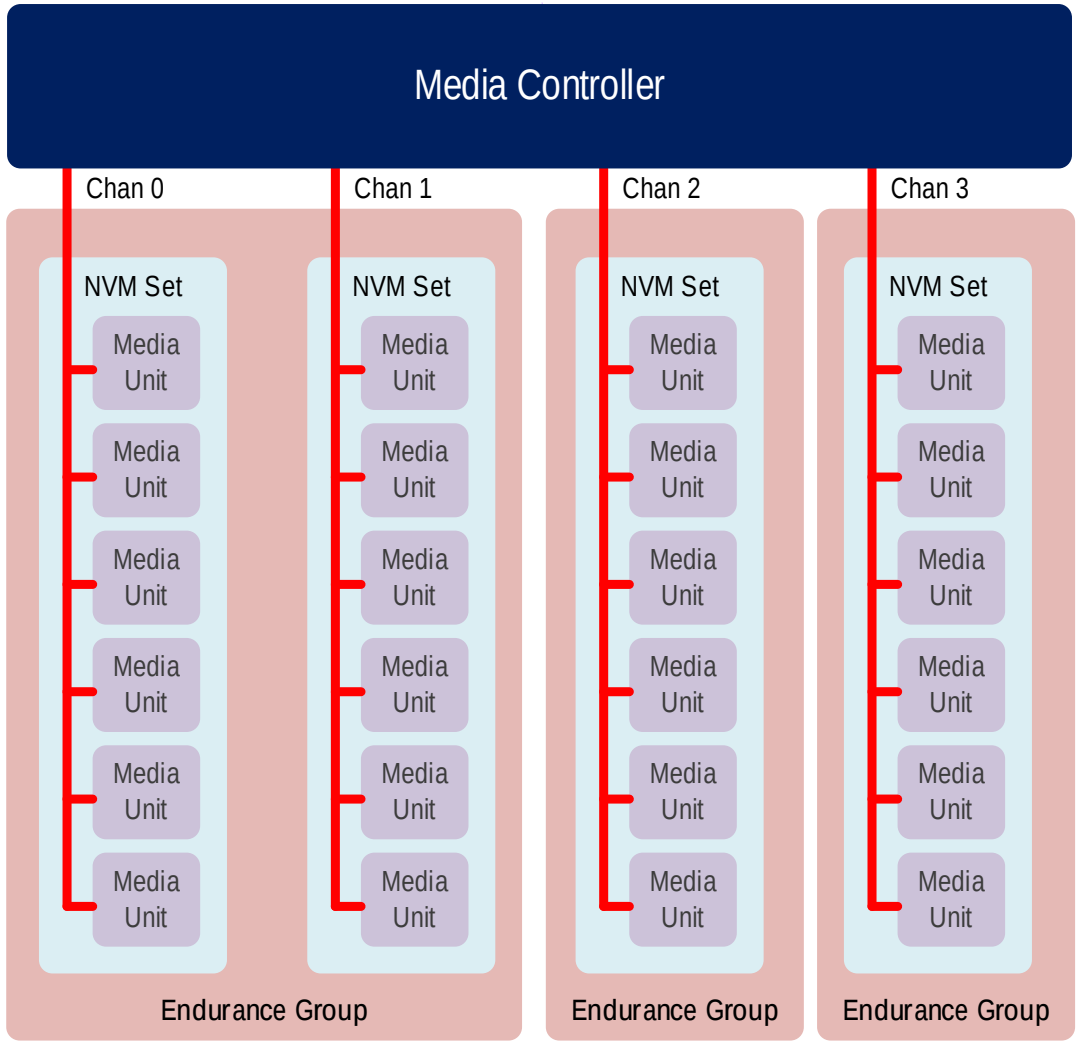

Domains and Partitions

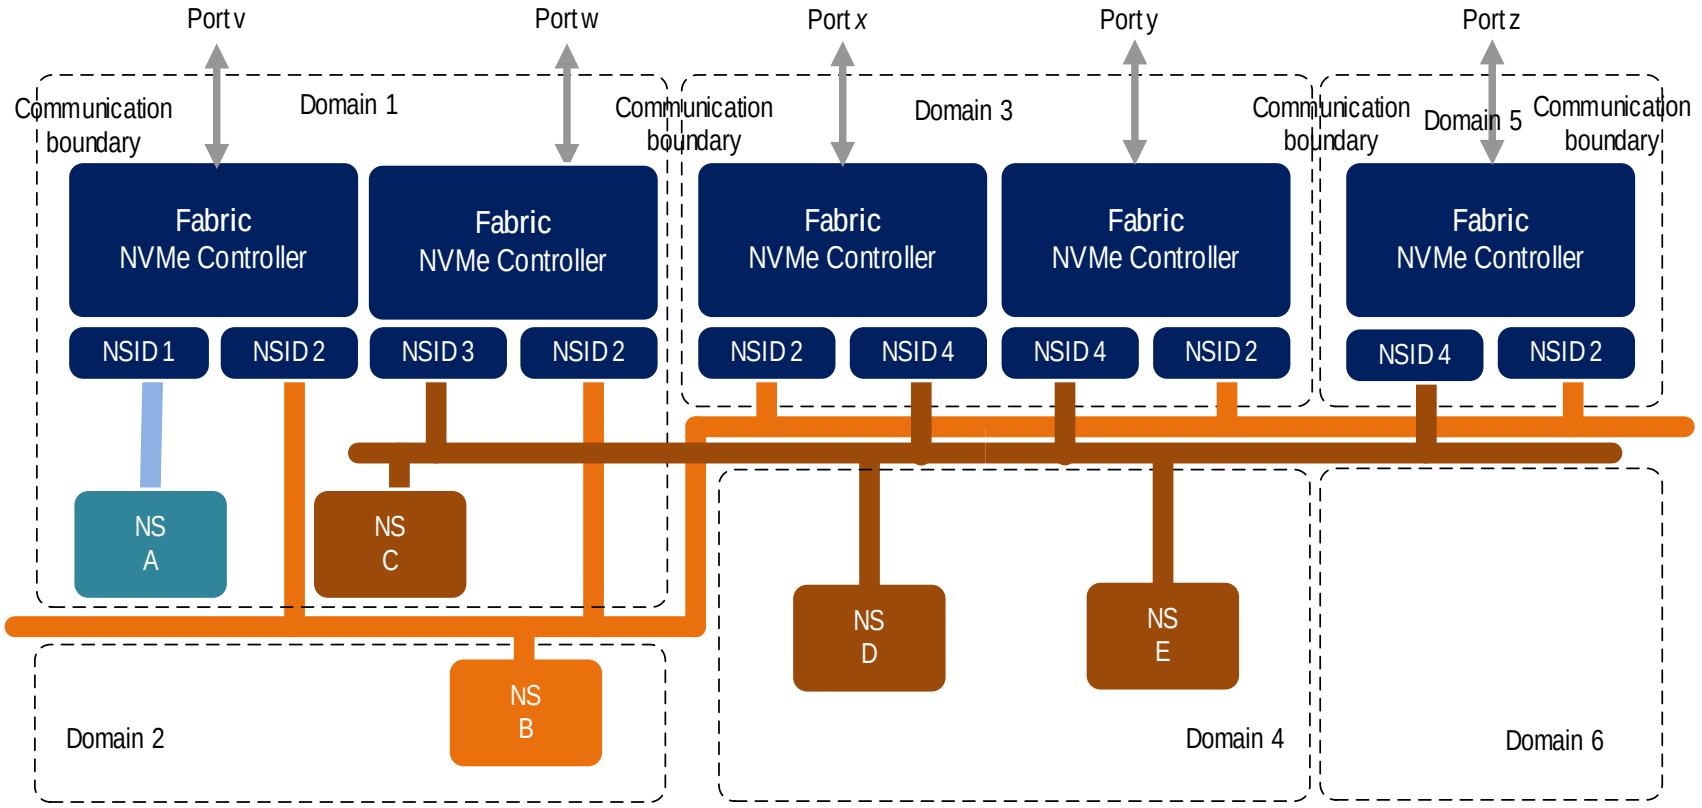

- Supports subdividing LARGE scale solutions – enabling partial operations & maintenance flows

Building Highways for Client, Cloud, & Enterprise

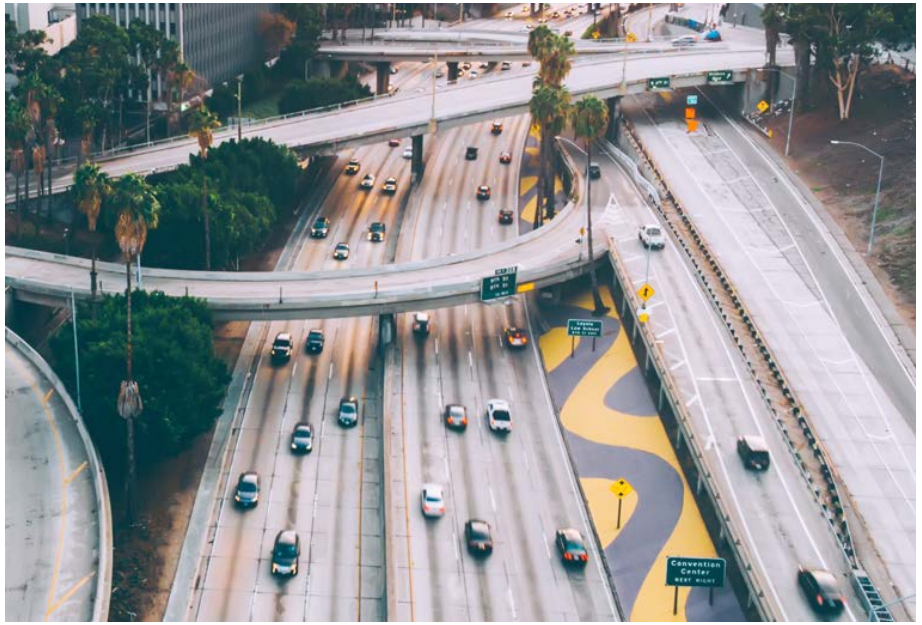

Infrastructure Enhancements
Simple Copy Command
Command Group Control Feature
Controller Memory Buffer Write Elasticity Status
Namespace Attachment Limits
Multiple Controller Firmware Update
Telemetry Enhancements
Command and Effect Log Enhancements
Non-"Maximum Data Transmit Size" Command Size Limits

- Use cases for NVMe® architecture continue to expand across Client, Cloud, and Enterprise
- Enhancements in Management, Telemetry and large storage systems meet the growing needs

Agenda

- Fixing the Memory & Storage Hierarchy
- Refactoring for the Next Decade of Growth
- NVMe® Architecture Advancements
- **What's Next: Computational Storage**

The Complex Database Universe

- Data Warehouses (Presto, SparkSQL, AWS Redshift AQUA, …) store **LOTS** of data
- Data stored in **LOTS** of (arbitrary) formats
- Data stored compressed and encrypted
- Formats and data constantly evolving

Compressed, Encrypted, Arbitrary Format

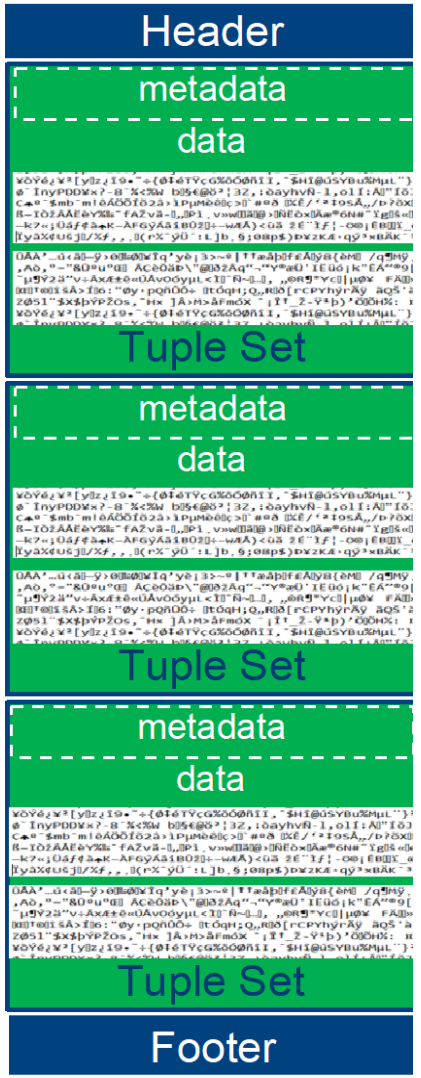

Stored Table

NAME	STATE
Bruce Banner	NH
Bruce Wayne	NY
Diana Prince	NJ
Jean Grey	NH
Jessica Jones	NY
Peter Parker	NY
Selina Kyle	CA
Tony Stark	CA

Finding the Needle in the Haystack

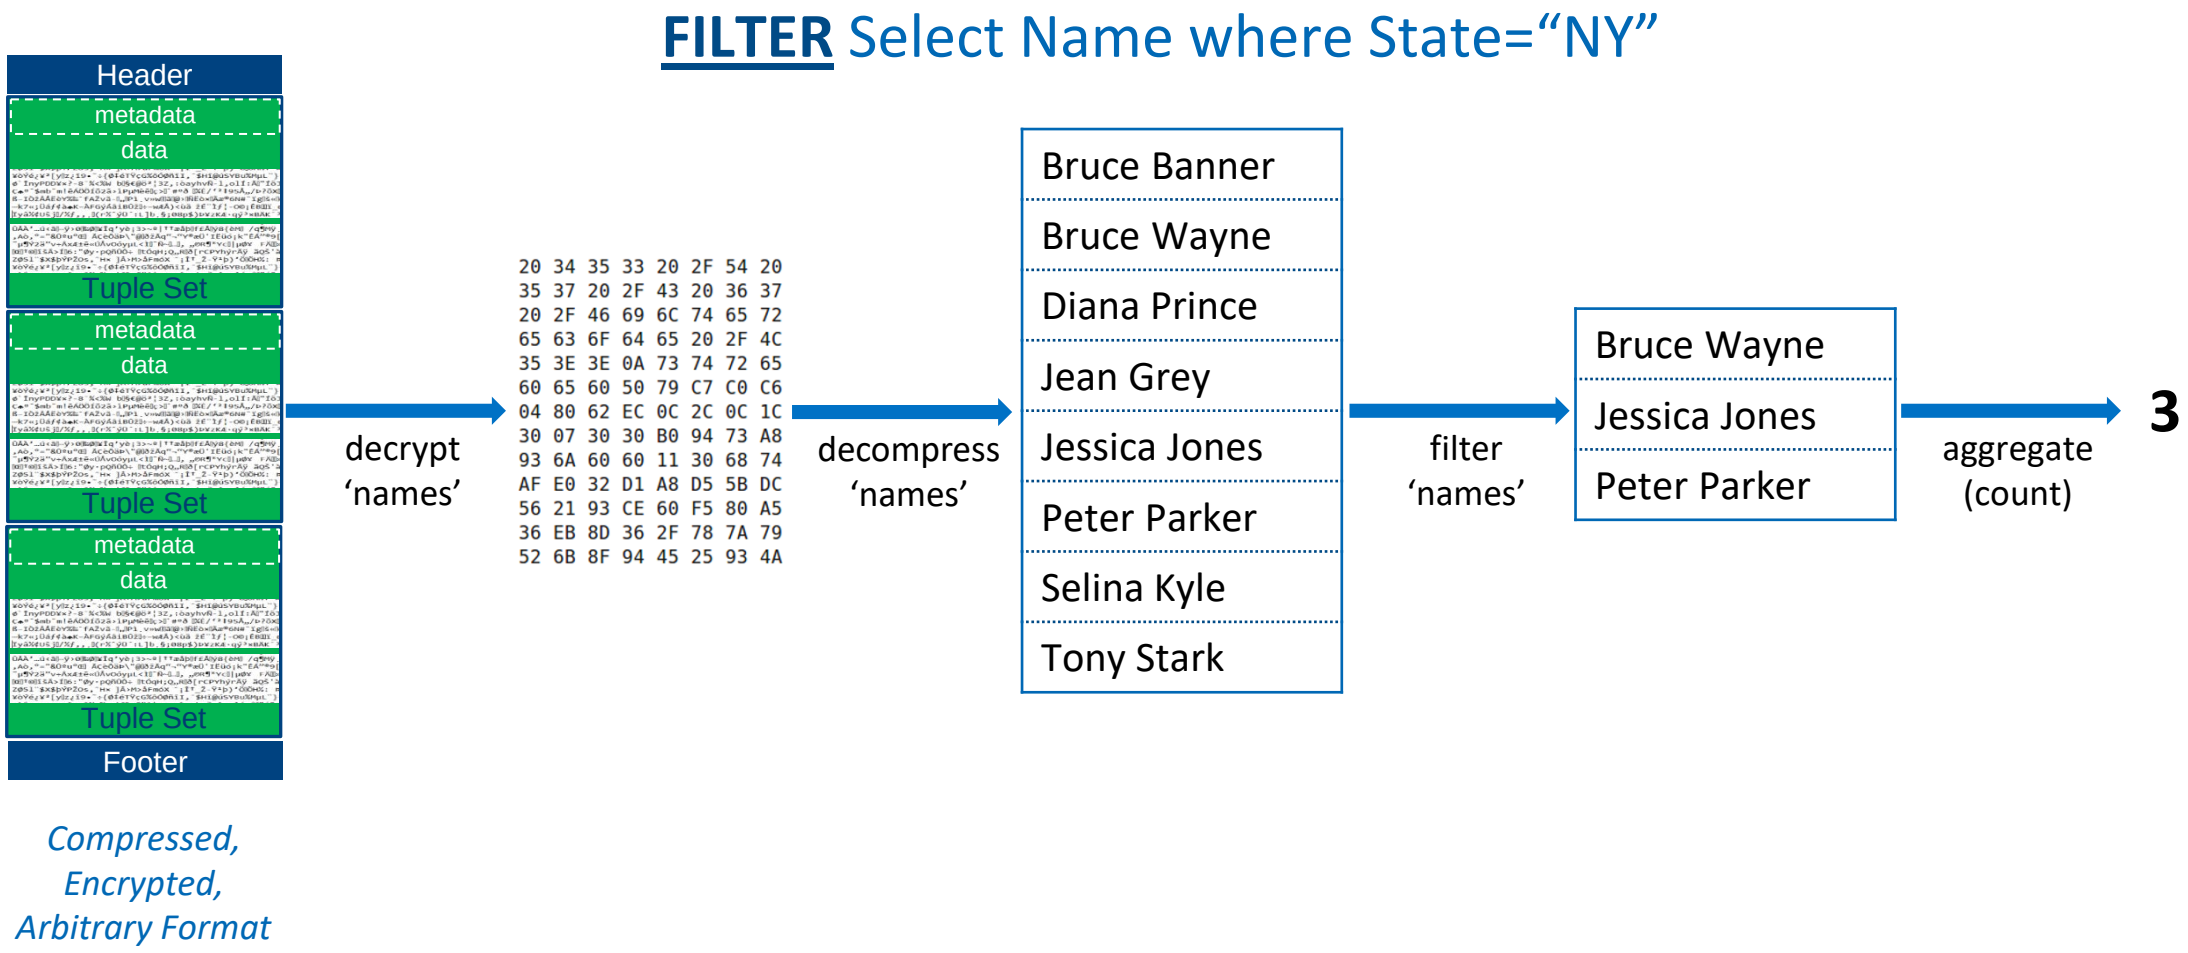

Programs as Computational Storage Offloads

- Programs invoked and used in standard way
- Programs in hardware agnostic bytecode and downloaded from host for later execution
- Device may offer fixed function programs
- Programs operate on data in on-device memory

Saving Power

Increasing Performance

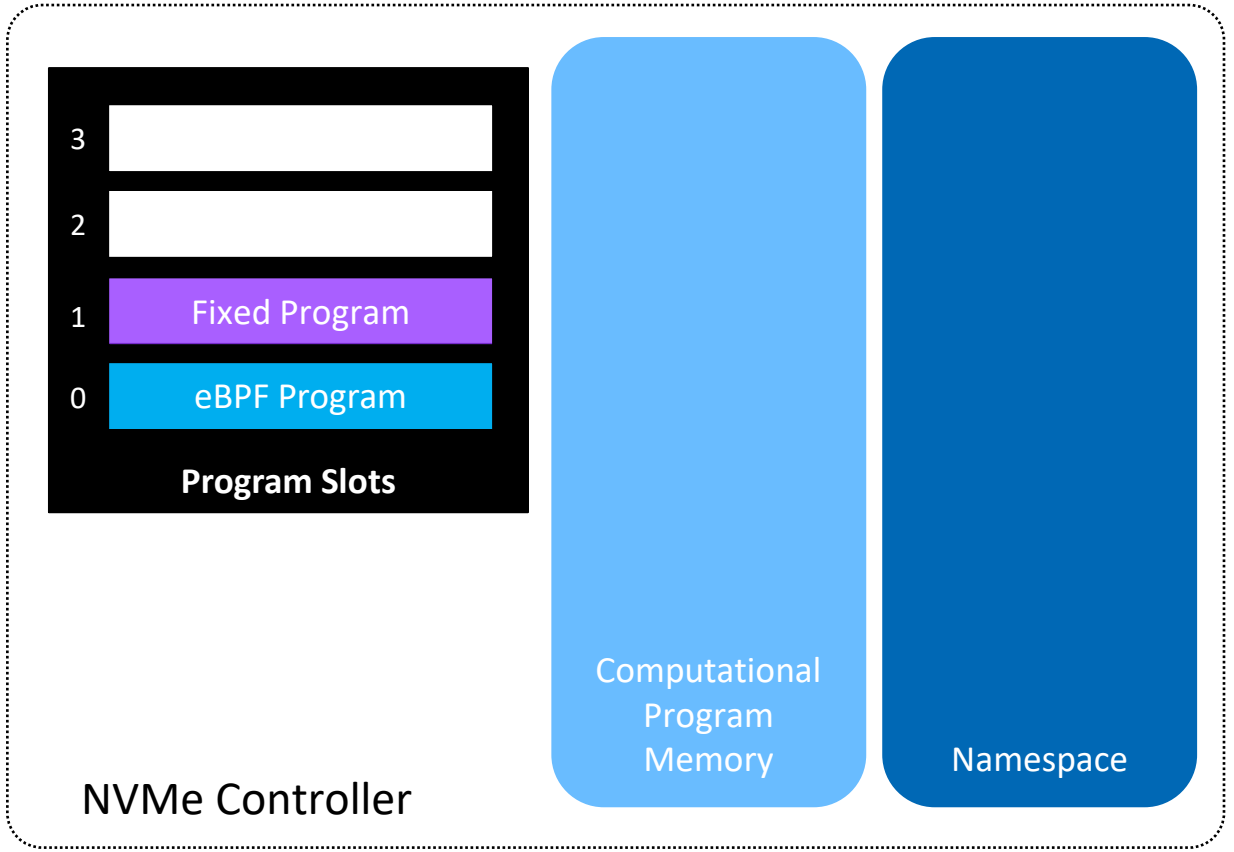

The Evolution of NVMe® Technology

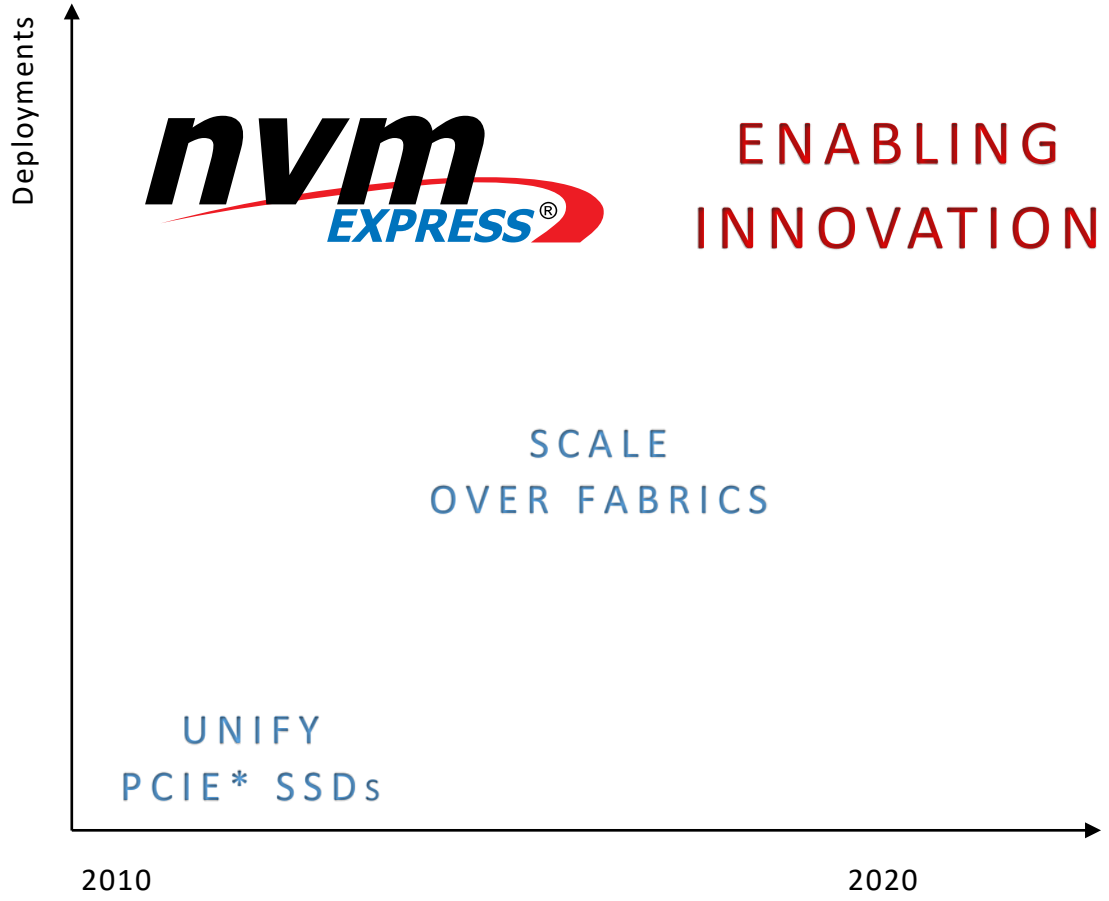


Q
A