Mixed Mode SSD

Mike Klemm,
Fellow, Technology and Storage Pathfinding
KIOXIA

FMS
the Future of Memory and Storage

KIOXIA

Mixed Mode SSD

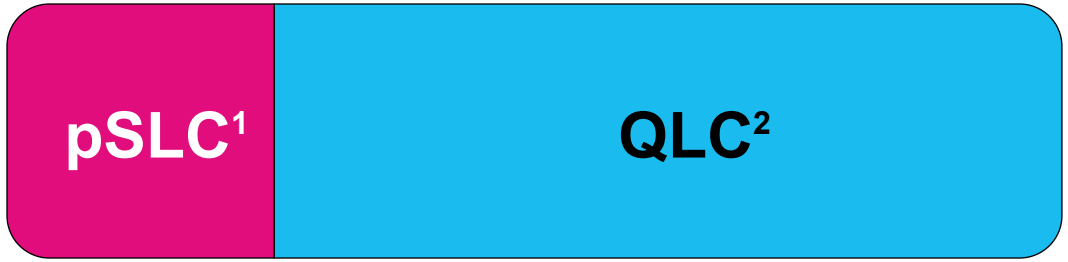

SSD delivering pseudo-SLC (pSLC) and QLC NAND delivering performance and endurance

Performance

- All SSDs provide input/output operations per second (IOPS) and bandwidth
  - vs Small-number discreet SLC SSD
- Balance PCIe® lane usage

Configuration Flexibility

- Custom pSLC:QLC ratio for each deployment
- pSLC random IOPS performance and durability with QLC capacity

FRUs

- Reduce Field Replaceable Units (FRUs)

Manage Wear

- Workload differences

1. Pseudo-single level cell (pSLC).
2. Quad-level Cell (QLC).

Usage need for Mixed Mode

- **Different IO Flows – Voice of Customer**
  - Small random vs large sequential writes
  - Metadata vs user data
  - Application differences
  - And many more…
- **Different Storage Needs**
  - Data tiering
  - Short-lived data vs long-lived data
- **Many Usages Need Flexible Space Configuration**
  - Ratio of pSLC:QLC capacity

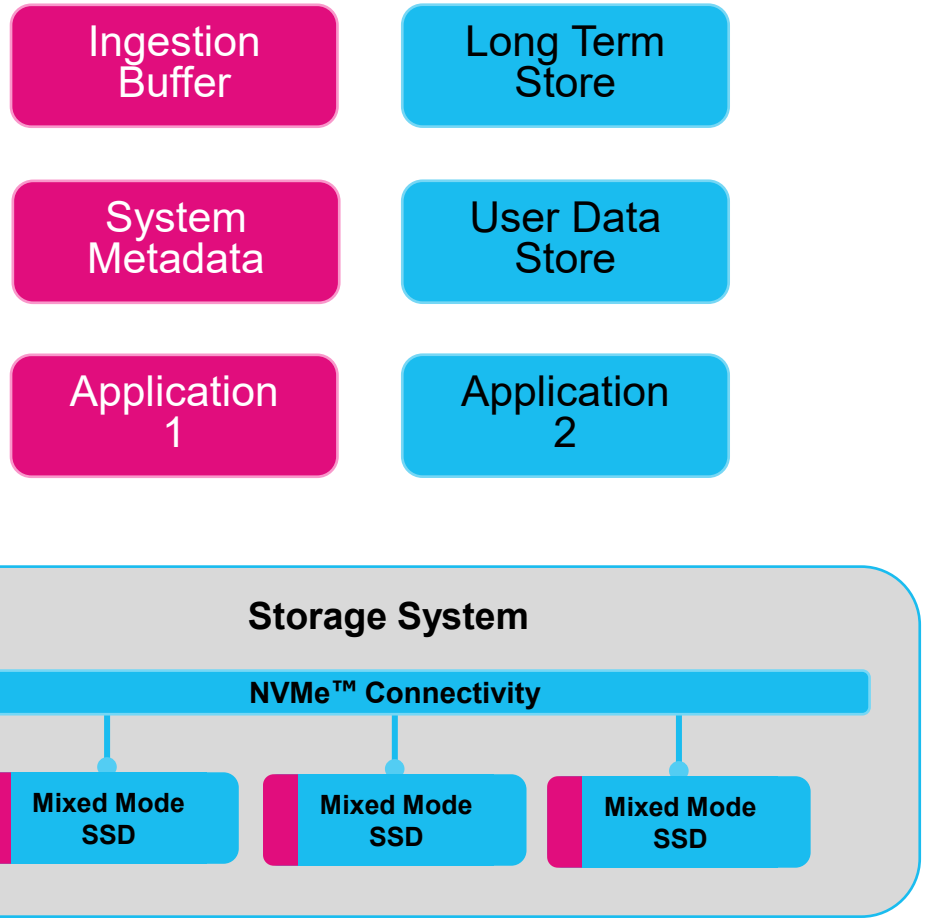

Relative Performance Estimation

- **Random IO Performance Driving Factor for Mixed Mode**
  - 4KB[1] or 16KB Indirection Unit (IU) size only matters for random write

Sequential Read
Consistent Performance
pSLC no advantage

Sequential Write
pSLC over QLC

Limited pSLC Capacity to sustain writes

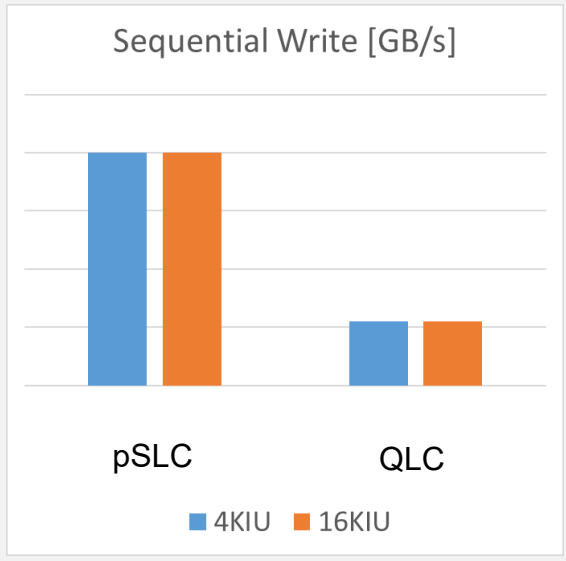

Random Read
pSLC over QLC
IU Size not Important

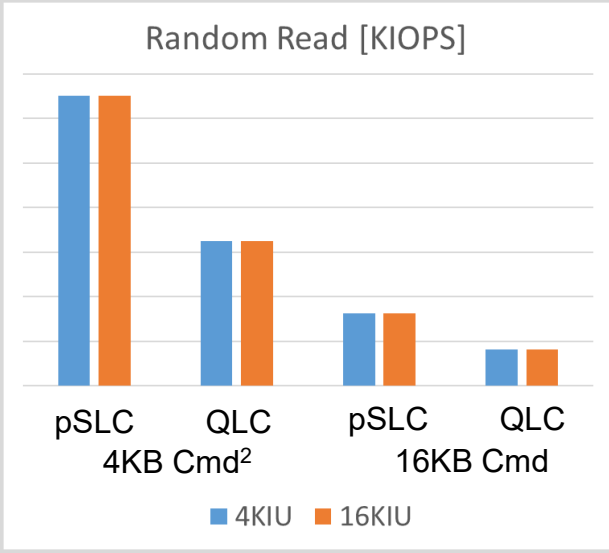

Random Write
pSLC over QLC
IU Size Matters

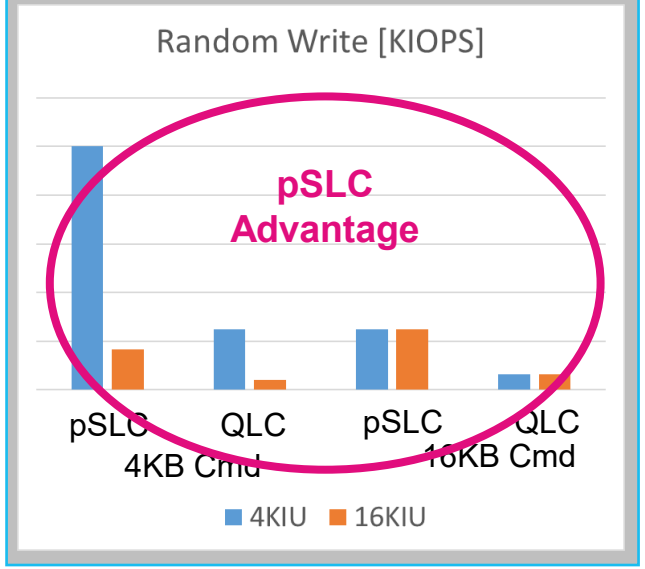

30 TB[3] SSD relative performance estimation. Units are not the same when comparing graphs.

Source: Tables and charts created by KIOXIA engineers. Images and/or graphics within this slide are the property of KIOXIA America, Inc. (KIOXIA) and are reproduced with the permission of KIOXIA. 1. Thousands of bytes (KB). 2. Command (Cmd) 3. Definition of capacity: KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes, a terabyte (TB) as 1,000,000,000,000 bytes and a petabyte (PB) as 1,000,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 2^30 bits = 1,073,741,824 bits, 1GB = 2^30 bytes = 1,073,741,824 bytes, 1TB = 2^40 bytes = 1,099,511,627,776 bytes and 1PB = 2^40 bytes = 1,125,899,906,842,624 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

Mixed Mode Capacities

- Deliver improved performance/endurance at the price of reduced capacity
- Voice of Customer:
  - Typically, pSLC ½% to 2% of maximum QLC capacity

Ratio	Example	Capacity Limit	Notes
1:2	TLC:QLC	50%	Limited Advantages
1:4	**pSLC:QLC** MLC:PLC	**25%**	Good beginning ratio for mixed mode
1:8	pSLC:PLC	12.5%	Future

Source: Tables and charts created by KIOXIA engineers.

Ratios

- **Naming**: Endurance NAND:Base NAND
- **Space**: 1:4 – Endurance Blocks:Base Blocks

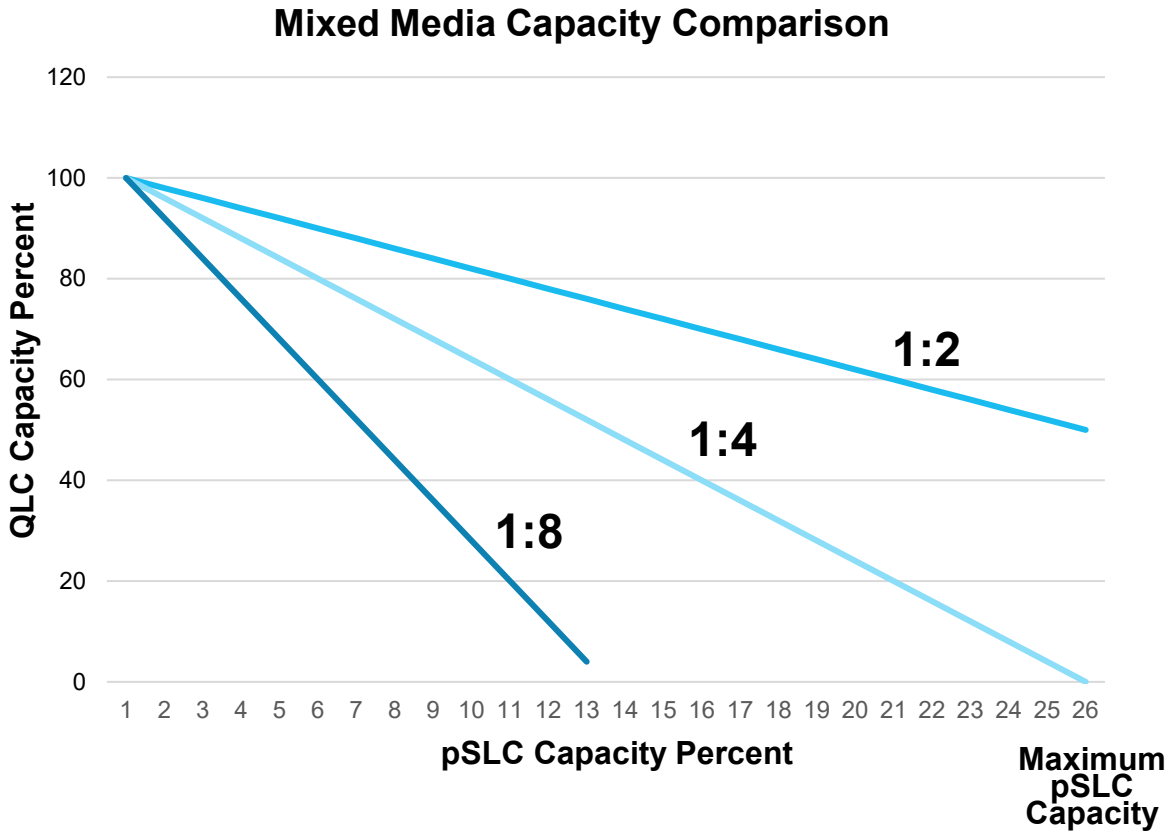

Creating SLC Capacity Discreet SLC vs pSLC

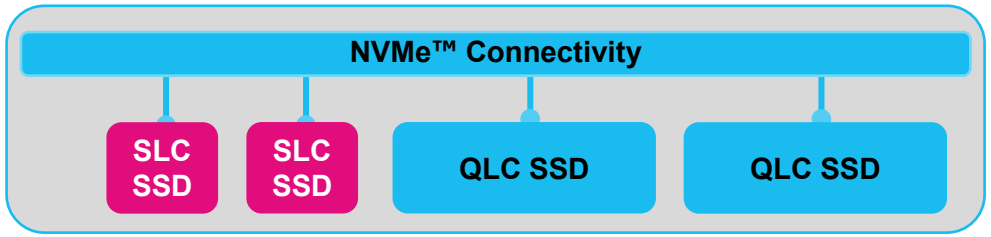

SLC Count	QLC	SLC Raw/Mirror
1	583 TB[1] / 95%	3.2 TB / 0 TB
2	552 TB / 90%	***6.4 TB*** / 3.2 TB
3	521 TB / 85%	9.6 TB / 4.8 TB
4	491 TB / **80%**	12.8 TB / **6.4 TB**

Limited number of SLC disks mirroring for protection

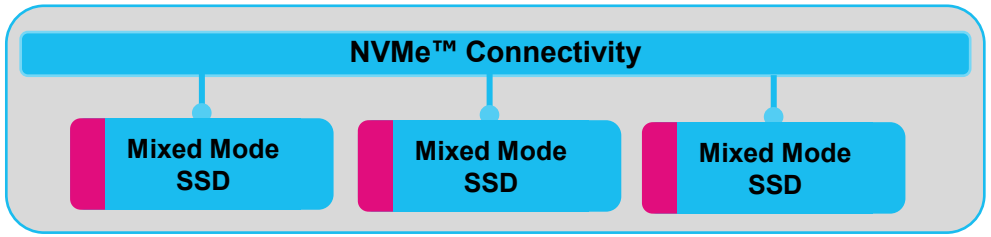

SLC Percent	QLC	SLC Parity
1%	589 TB / 96%	1.5 TB
2%	564 TB / 92%	3.1 TB
3%	540 TB / 88%	4.6 TB
4%	515 TB / **84%**	**6.1 TB**

pSLC space on all SSDs parity protection on pSLC

- **Mixed Mode efficiently delivers high endurance pSLC**
  - Slot Tax: SLC SSD consumes full drive slots reducing overall capacity, What about SLC spares?
  - Mixed Mode allows for parity protection in more configurations
- **Limit Field Replaceable Units (FRUs) – Single SSD model vs two for replacement**

Example SSD Capacities QLC: 30.7 TB, SLC: 3.2 TB, 20 Data SSDs (No Spares or Parity)

KIOXIA

Mixed Mode Configuration Flexibility

- **Define pSLC:QLC ratio at time of usage**
- **Configurable amount of space**
  - Initial vs expansion SSD capacity
  - Fine granularity
    - Dependent on system configuration
- **For performance, capacity or endurance**
  - Need the capacity for the system, but after certain number of SSDs it doesn't make sense

Example

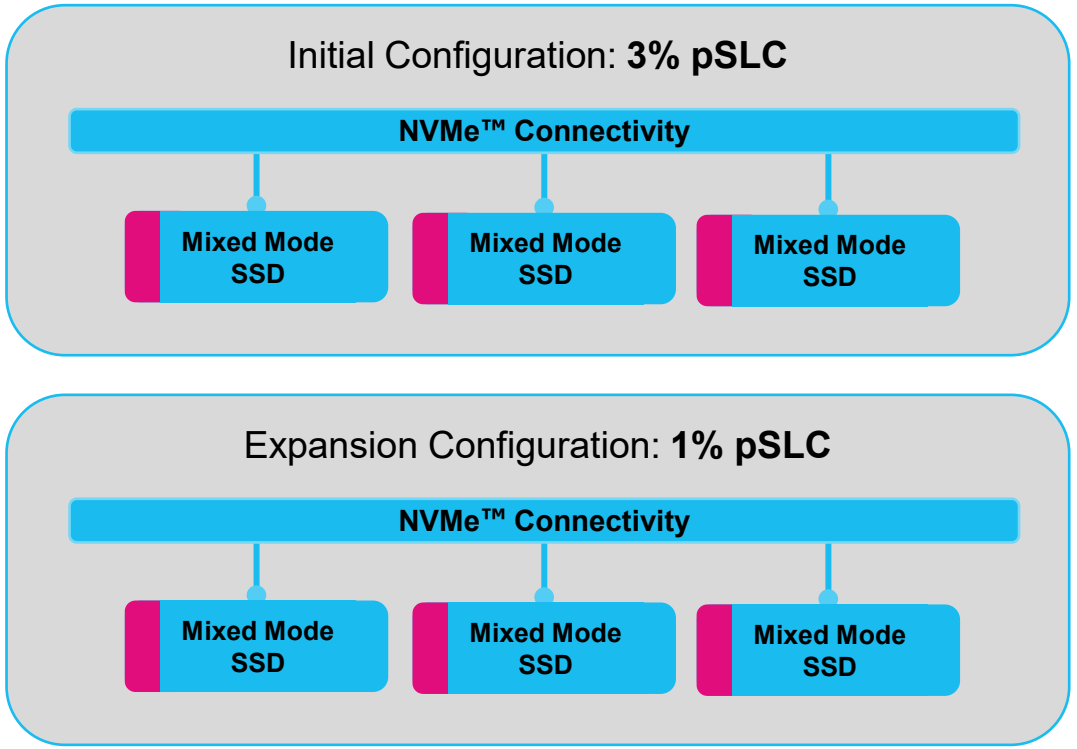

Presenting Mixed Mode Space

Initial pSLC:QLC ratio configuration defines permanent Namespace[1] and Endurance Group[2]

pSLC Namespace and Endurance Group are created when SSD is first used and cannot be deleted once created.

QLC Namespaces may be created and deleted as needed.

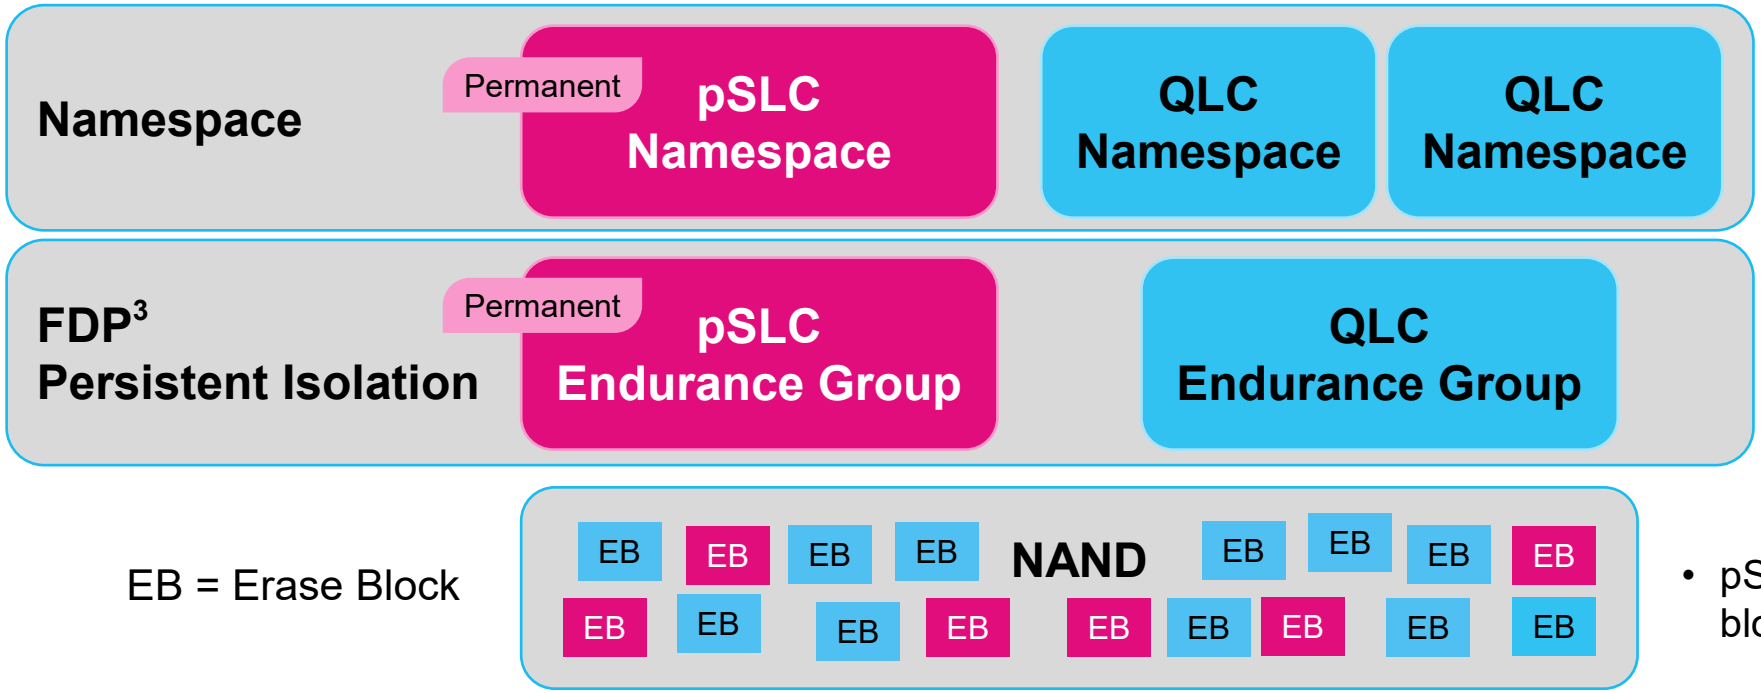

- Separate LBA Namespace
- Manage different address ranges

- Use FDP to avoid mixing writes

- pSLC uses all NAND die for performance blocks dedicated to pSLC or QLC

Once block is used by pSLC it cannot be used by QLC

 1. Namespace is a logical storage volume that acts as a separate storage space within the SSD. 2. Endurance Group refers to a feature where multiple NVM Sets can be grouped together, sharing endurance characteristics and undergoing wear leveling collectively. NMV Sets is a base specification feature under NVM Express®. NVM Express is a registered trademark of NVM Express, Inc. 3. Flexible Data Placement (FDP).

Balancing Performance with Mixed Mode

Discreet SSD Balancing

- Unbalanced network connections (IOPS & BW[1])
- Unbalanced PCIe® lanes
- Disrupted by expansion and replacement

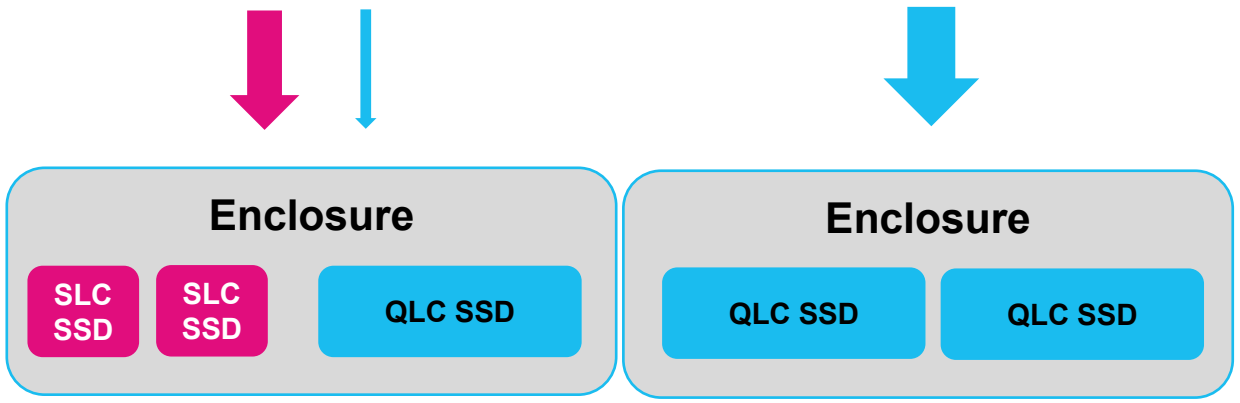

Mixed Media Balance IO Load

- Balance enclosure network connections
- Balance PCIe lanes
- Balances through expansion and replacement

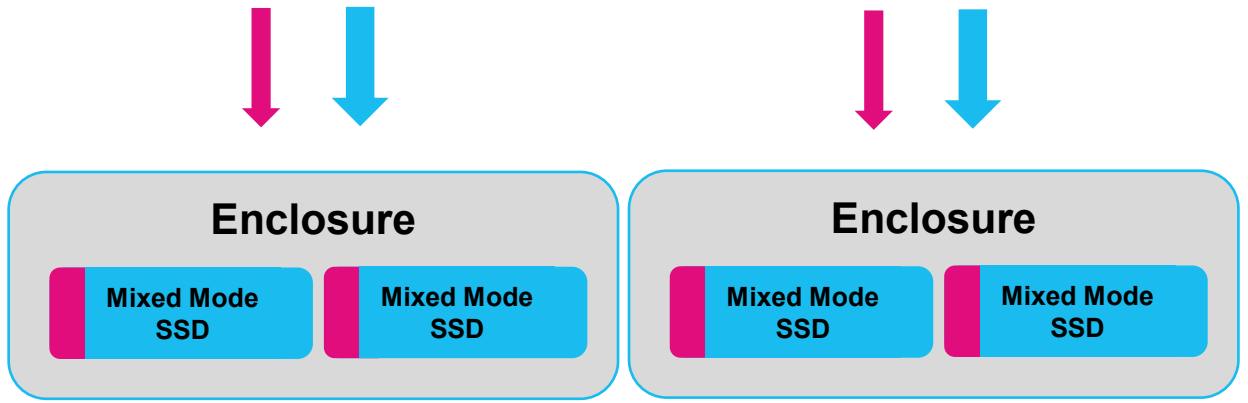

KIOXIA

 1. Bandwidth (BW). PCIe is a registered trademark of PCI-SIG

Mixed Mode Performance

- **NAND Write Cycles**
  - pSLC and QLC share same NAND dies
- **Gathering Requirements for Performance**
  - Looking for your feedback
  - What is the ratio of pSLC to QLC performance?
  - Likely application specific – customer configuration?
  - What are the minimums?

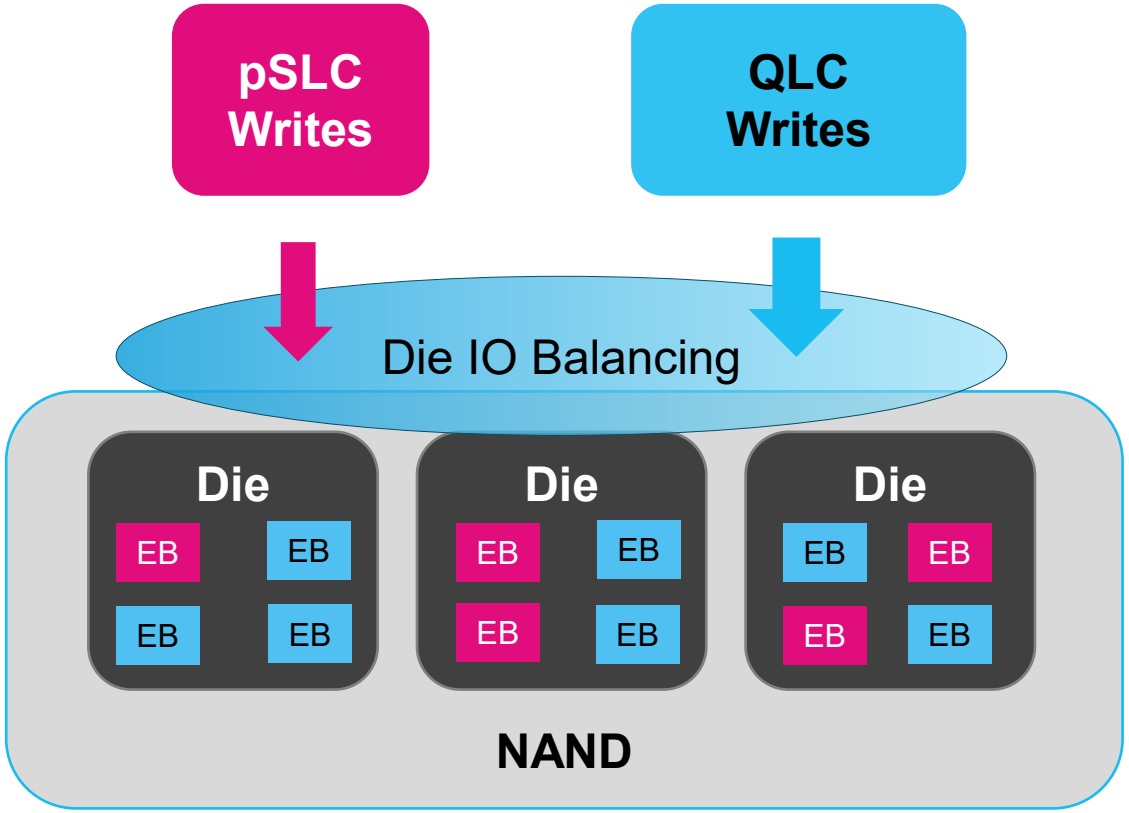

Mixed Mode Wear

- **Like Discreet Devices**
  - pSLC compares to SLC
  - QLC is the same
- **pSLC vs QLC**
  - 25+x write/erase (W/E) cycles
  - Cannot swap Erase Blocks usage
    - pSLC to QLC or QLC to pSLC

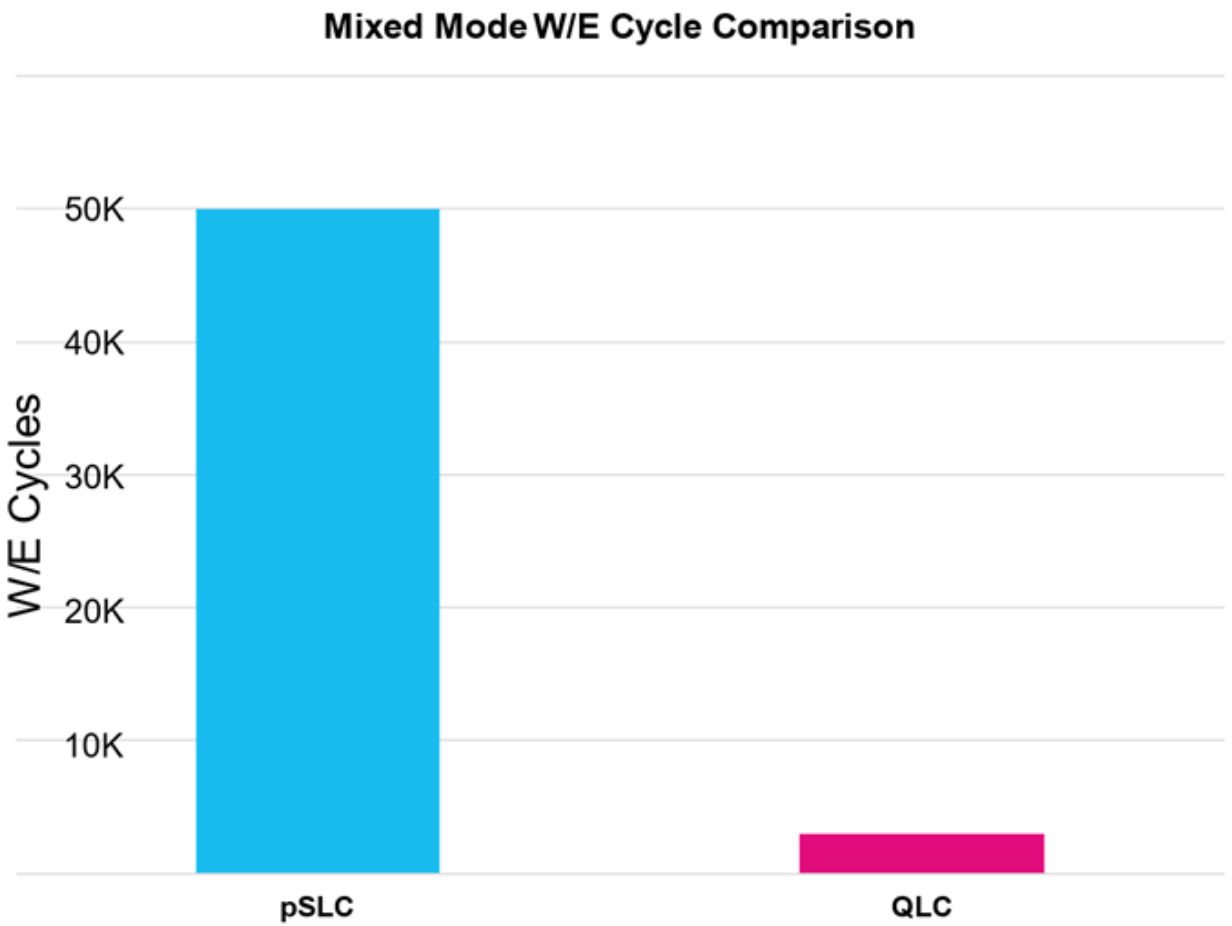

Standards

- **pSLC:QLC Ratio Configuration**
- **Permanent Namespace & Endurance Group mixed with user created**
  - pSLC Endurance Group cannot be deleted
  - pSLC Namespace cannot be deleted
  - Namespace attribute for pSLC vs QLC
- **Quality of Service (QoS)**
  - TP 4176
    - FMS Session: INDA-201-1: NVMe™ State of the Union, Configurable Device Security and Quality of Service (QoS)
    - Is Rate Limit Mode enough?

Summary

- **Mixed Mode delivers performance and capacity**
  - Connection balancing
  - Improved capacity
- **One-time configuration of pSLC:QLC ratio**
- **Standards work is needed**
  - Permanent Endurance Group and Namespace
  - QoS

KIOXIA