Using AI to Improve the Performance and Endurance of High-Capacity SSDs

Flash Memory Summit

Andy Mills, CEO/Co-founder
Enmotus, Inc.
andy.mills@enmotus.com

Santa Clara, CA
November 2020

Summary

- Advances in machine learning software has been shown to dramatically improve performance and endurance of QLC (and future PLC)
- QLC NAND segments may be programmed in a higher endurance pseudo SLC mode and accessed by intelligent drivers
- Allowing the host to control these segments instead of the SSD controller provides better application-storage media balancing

SSD Market Dynamics

- Solid State Disks (SSDs) market rapidly growing
  - $47BN in 2020 growing to $87BN by 2025*
  - 350 million SSDs shipped annually
  - Rapid shift to NVMe class SSDs

- SSDs are getting larger and more cost effective
  - PC OEMs replacing HDDs with SSDs in mainstream PCs/notebooks/servers
  - SSD vendors spending $billions to improve capacity/cost using 3D multi-layer NAND
  - Next generation Quad-Level-Cell (QLC) and Penta-Level-Cell (PLC) offers significant capacity and cost gains

But there is a problem….

* Source: TrendFocus, various Internet media sources

SSD NAND Trends and Tradeoffs

Newer Generations of SSD Storage Media driven by Lower Cost Per Bit

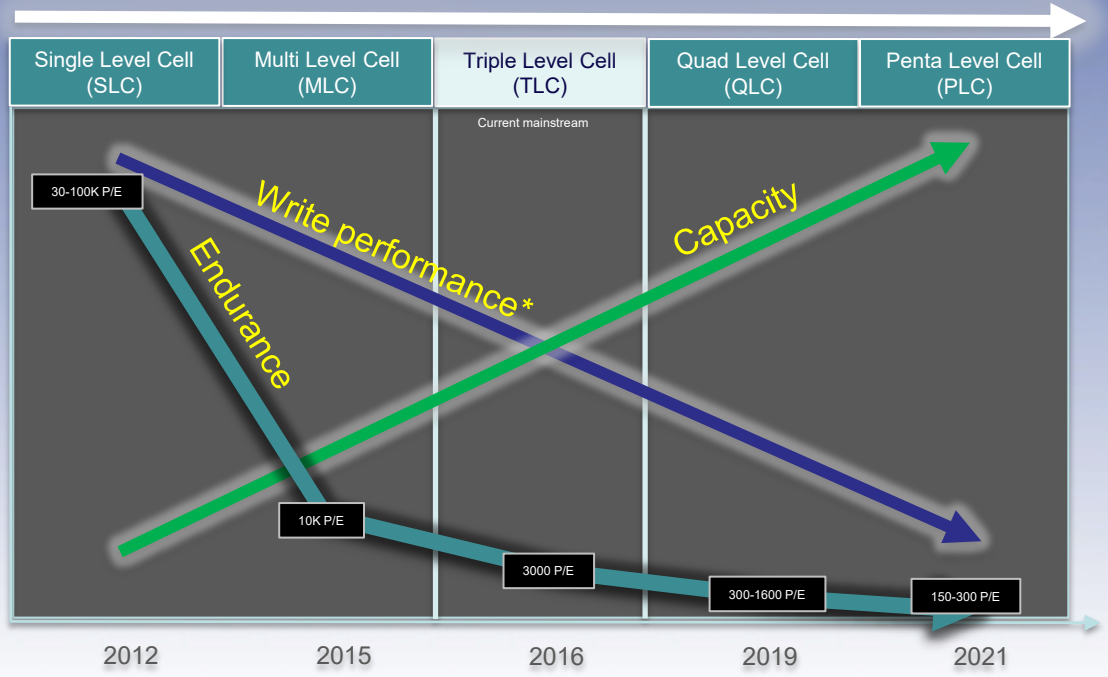

* Raw media performance shown without SLC caching

SSD storage media (NAND) reliability and performance is worsening with each generation.

High performance-endurance are small and expensive whereas next generation high capacity SSDs have low performance and endurance

Machine Learning SMART SSD Opportunity

- Annual spending on SSDs to grow by $40BN over next 5 years
  - TLC NAND will quickly hit cost per bit, capacity and product margin limits
  - QLC offers 33% capacity increase BUT has performance and endurance issues, holding back adoption

- Machine Smart SSDs – a new category of higher endurance SSD
  - Performance and endurance of SLC, capacity and cost of TLC/QLC/PLC
  - Adapts to OS/user workloads via innovative machine learning algorithms

- Smart SSD benefits
  - Demonstrated significant ***write attenuation*** to QLC in client PC workloads
  - Workload based zoning and allocation of different flash media type
  - Allows QLC and future low endurance NAND to be adopted in higher volume applications

Pseudo SLC

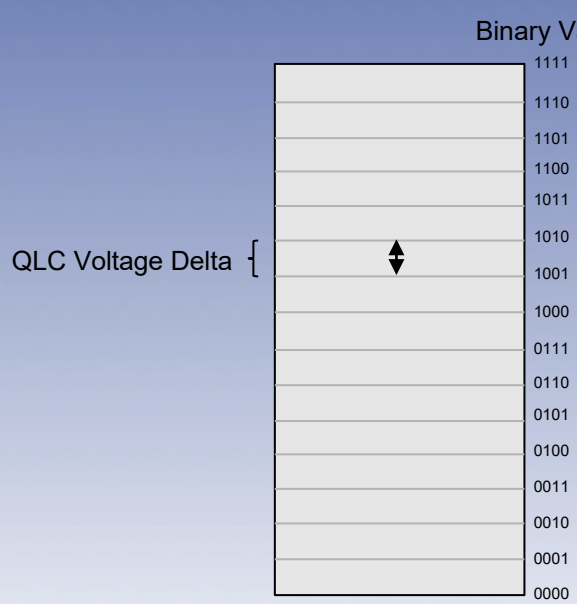

QLC = 16 voltage levels

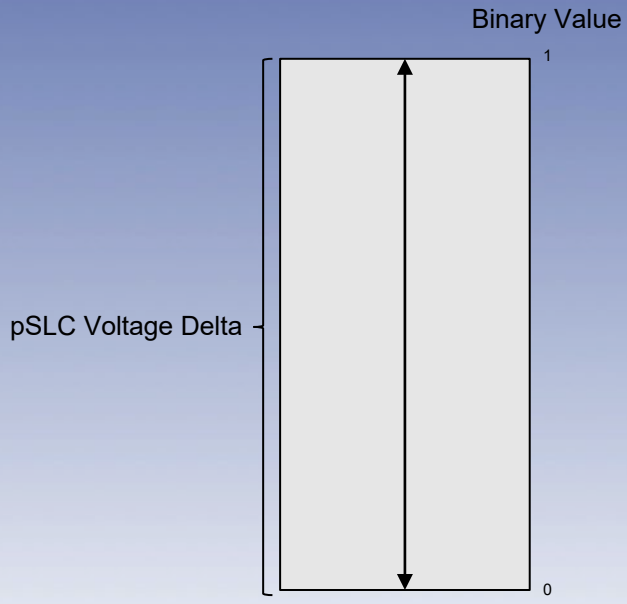

pSLC = single voltage level (max to min)

- Pseudo SLC (pSCL) uses the full cell voltage swing (or effectively 0000 and 1111 only) of a standard QLC cell
- May be programmed statically (30K P/E cycle) or dynamically (< 1600 P/E cycles)

QLC Allocation Options

AI Enabled QLC

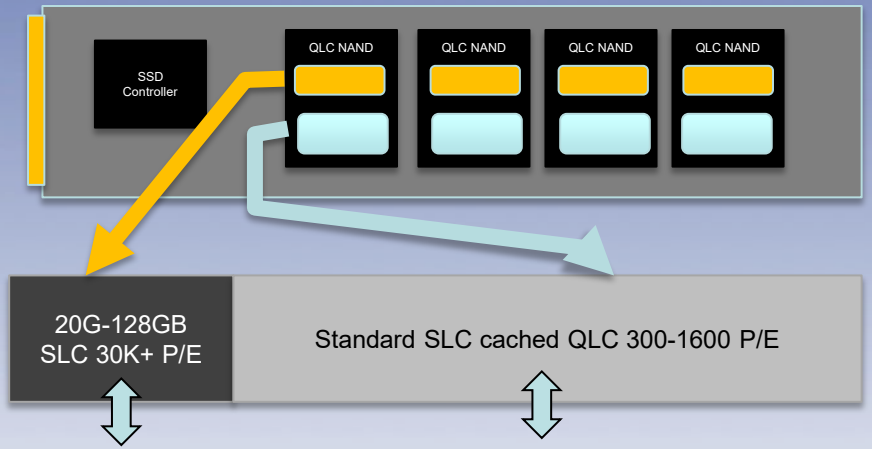

- Large use of visible ***static SLC*** as a host accessible region of the device
- Remainder of device operates as a standard SLC cached QLC device
- Independently accessible regions
- Static SLC region has very high endurance

Standard QLC

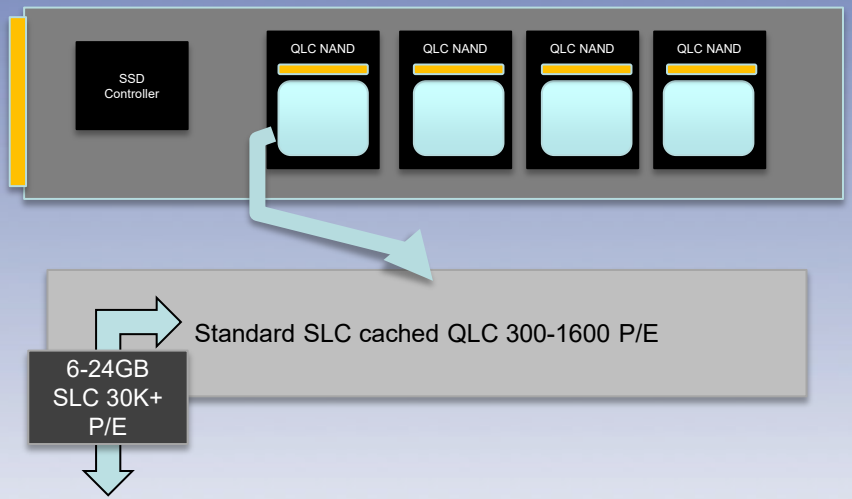

- Small use of hidden ***static SLC*** for buffering/caching of data
- Dynamic SLC mode while device < 20% full
- Data moved to QLC as part of housekeeping once > 20% full
- Dynamic SLC endurance is no better than QLC

How Smart SSDs Work

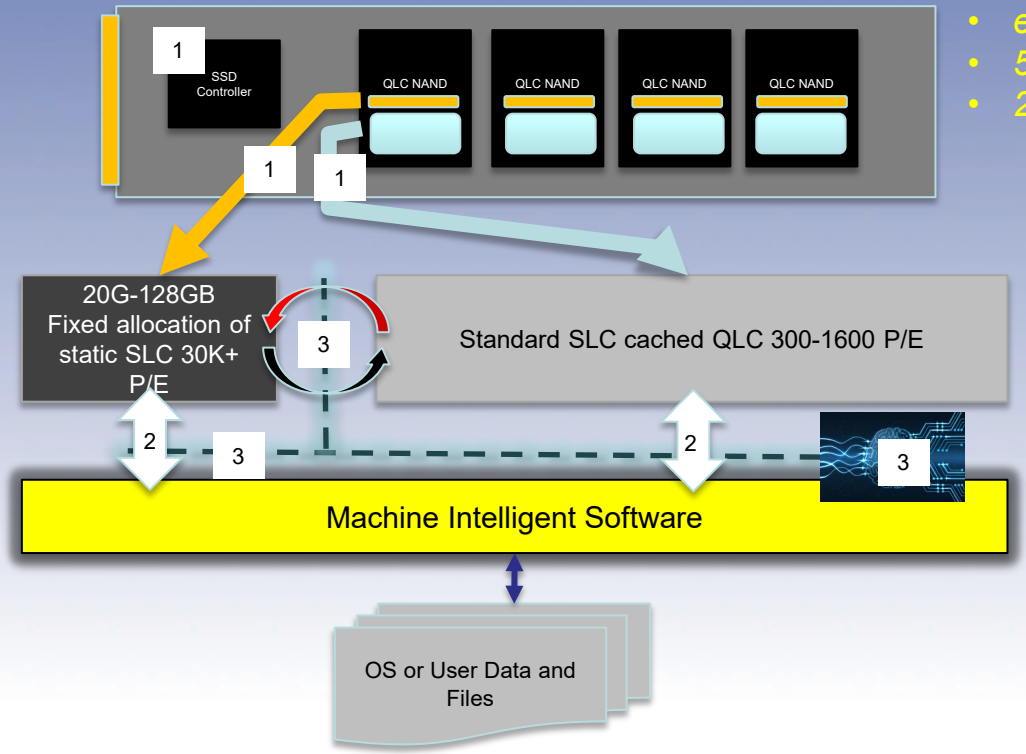

Lab tests show:

- *equivalent performance to high end MLC SSDs*
- *5x improvement in endurance over TLC*
- *25 improvement in endurance over standard QLC*

1

SSD NAND is divided into two pools: SLC and QLC by the SSD controller firmware

2

Smart SSD software has DIRECT access to the two types of NAND flash:

SLC = high performance and endurance

QLC = high capacity, low endurance

3

Data is intelligently and continuously balanced across SLC or QLC by MI software

Heavy traffic => SLC

Light traffic => QLC

SLC is smart provisioned on the fly

AI Storage Adapts to User Workloads

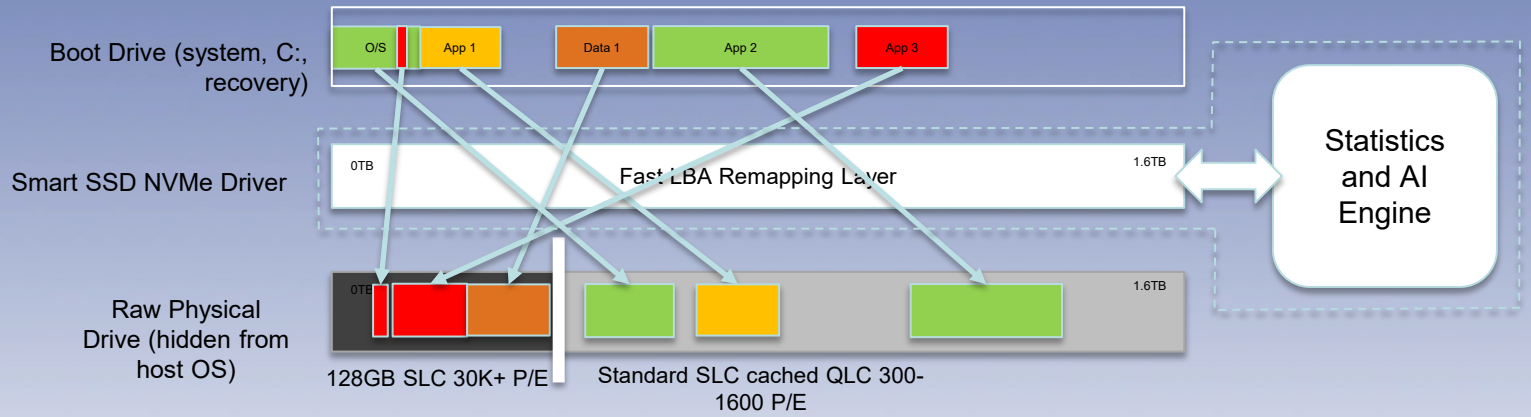

- AI engine constantly and passively analyzes user traffic over 100% of the volume
- Whole disk usage analysis allows better decisions to be made over time
- Active data is migrated in the background to the best available storage – in this case, SLC
- Results in significant write attenuation to the QLC (or write avoidance)

Performance: PCMark10 Quick Storage Test

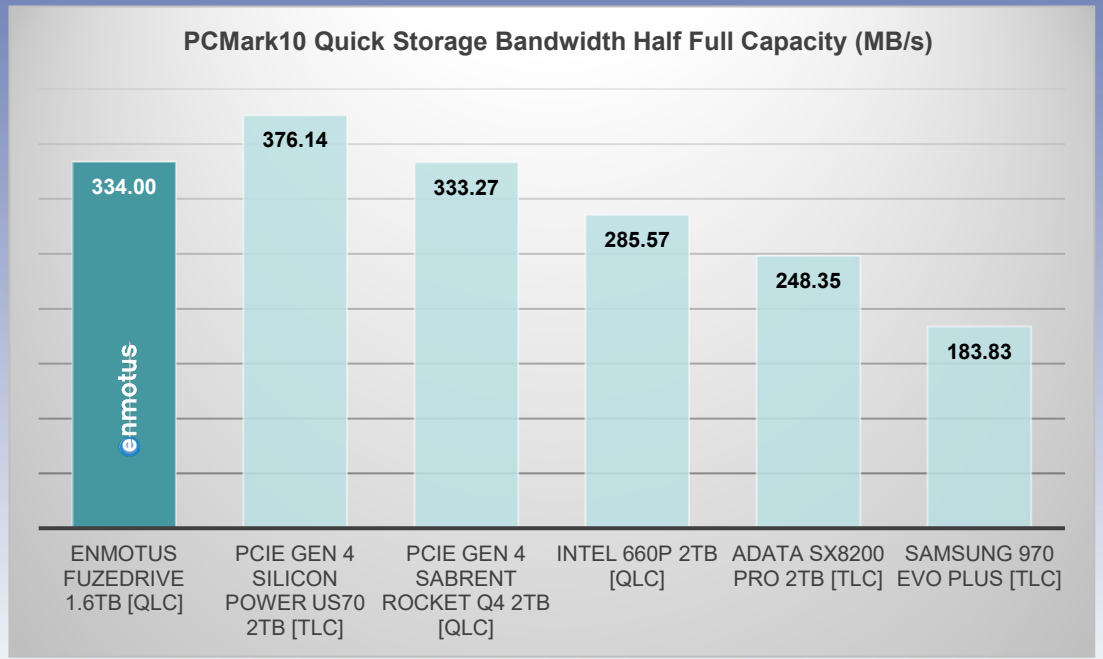

Source: "Sabrent Rocket Q4 NVMe 2TB SSD Review", Tweaktown, Jon Coulter

- AI adapts in real time to real world storage traffic
  - OS + data + streaming
- PCMark10 provides a useful way to verify real world performance
  - Up to 82% faster than TLC using blended workload types
- More AI friendly benchmarks are needed
  - All today's synthetic are based on single media assumptions

Maintaining Consistent Performance

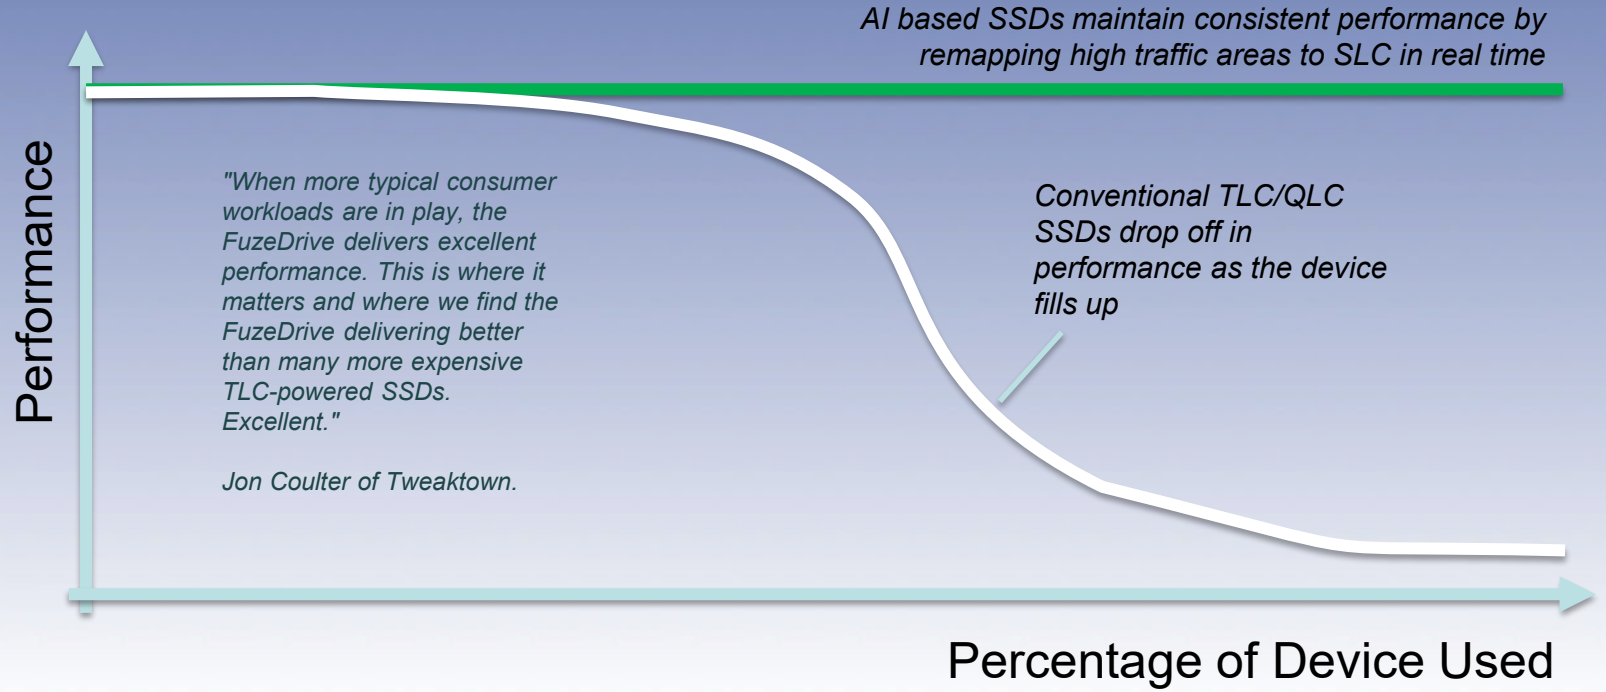

A New Endurance Classification

FUZEDRIVE			The World's Smartest SSD
SSD Endurance Levels			
ENDURANCE (terabytes written/GB)	> 1 TBW per GB	0.5 TBW to 1TBW per GB	Less than 0.5 TBW per GB

https://www.enmotus.com/hubfs/PDFs/White_papers/Fuzedrive-SSD%20endurance_v1.0%208.12.20.pdf

- TBW/JEDEC 219 specs normalized to a standard capacity to make it easier to compare between vendors
  - Calculated using TBW/GB e.g. 400 TBW for 2TB = 400/2000 = 0.2
  - Solves the complexity issue and increasing capacity causing unrealistic TBW numbers as SSDs get larger
- Enmotus offering free use of icons to any vendor/system vendor to help promote
  - Email us at info@enmotus.com

Endurance Comparison Chart (TBW)

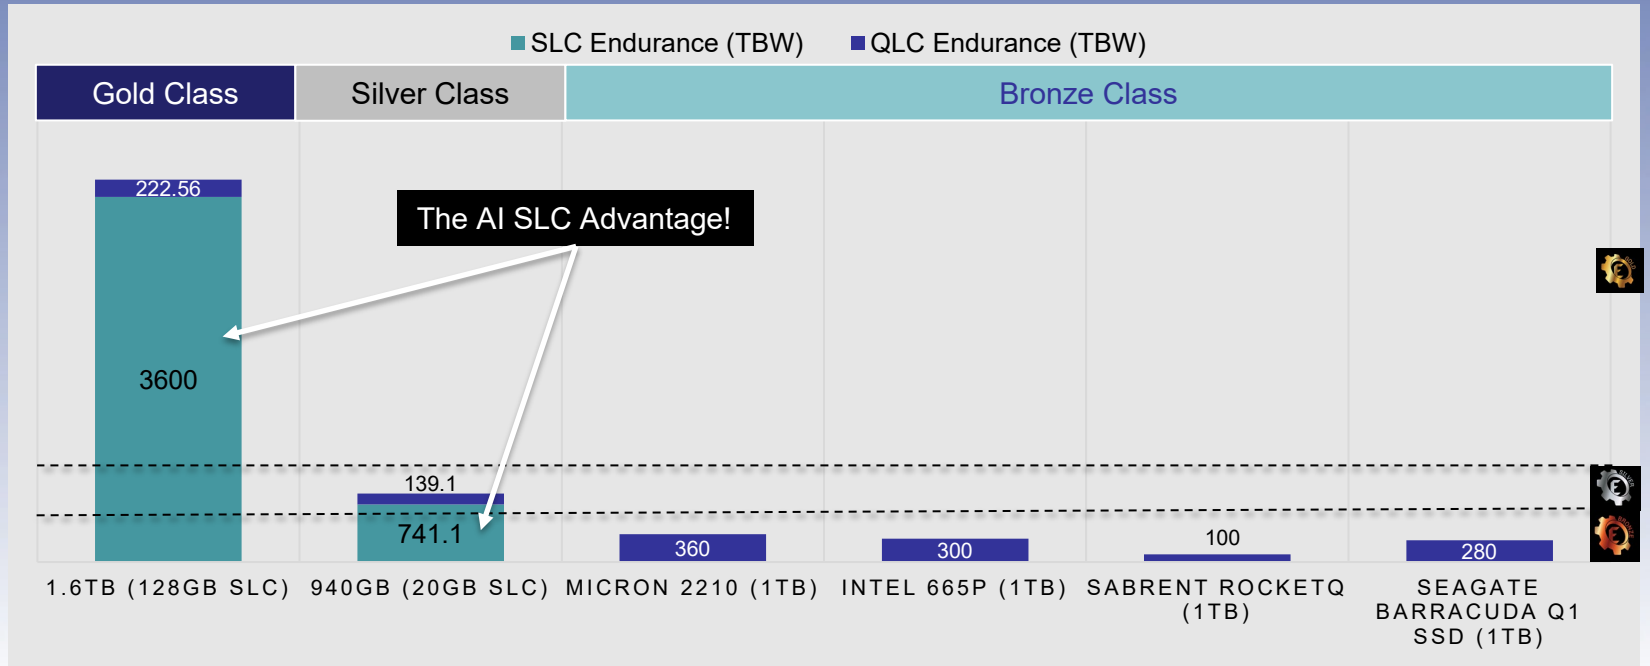

More info at: https://www.enmotus.com/hubfs/PDFs/White_papers/Fuzedrive-SSD%20endurance_v1.0%208.12.20.pdf

FuzeDrive Smart SSD

Developed in partnership with Phison – the industry's first SMART AI powered SSD

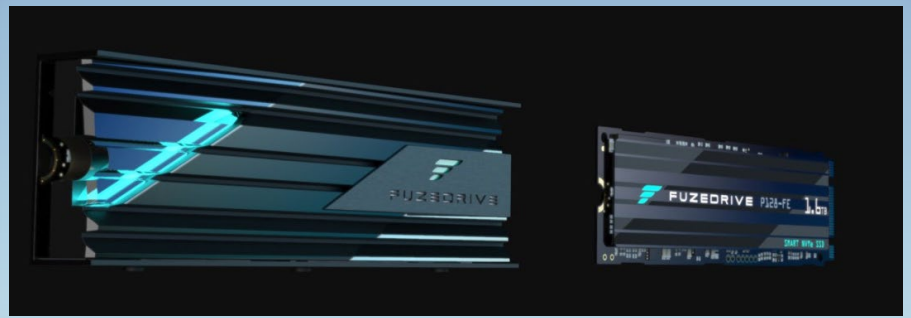

FuzeDrive SSD: 900GB, 1.6TB

Performance

Single cell NAND coupled to smart host software for maximum performance

Capacity

60% more capacity than performance TLC/MLC

Gold Endurance

Up to 25x higher endurance than regular QLC

Up to 5x higher endurance than TLC

Summary

- Using AI (machine learning), it is possible to greatly enhance the endurance of newer QLC and future PLC flash devices
- Takes advantage of pSLC modes in today's QLC NAND
- No new hardware – all off the shelf QLC controllers and NAND components
- Allows for a new class of high endurance "smart" consumer devices