It Takes Guts to be Great

Sean Stead, STEC

Tutorial C-11: Enterprise SSDs

Tues Aug 21, 2012 8:30 to 11:20AM

Who's Inside Your SSD?

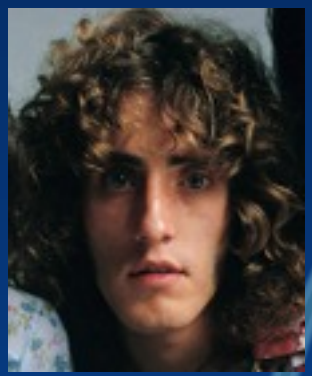

Wear Leveling Technology

Advanced Firmware Technology

Full Data Path Protection

Power Backup Technology

Advanced Flash Management

Host Interface

It's What's On The Inside That Counts…

- Controller "The Brains"
  - e.g.: Marvell, SandForce, Indilinx, STEC
  - types: ASIC, FPGA
- NAND "Flash" Chips for Storage

  e.g.: Toshiba, Samsung, Micron
- DRAM to Speed Up Data Transfer
- Interface to connect to system

  e.g.: SATA, SAS, PCIe, USB, Embedded

Flash Memory SUMMIT

Get Schooled

OLD

HDD

Vs

NEW

SSD

Inside a Hard Drive
What a mess…

Inside a Solid State SSD

Look ma, no moving parts…

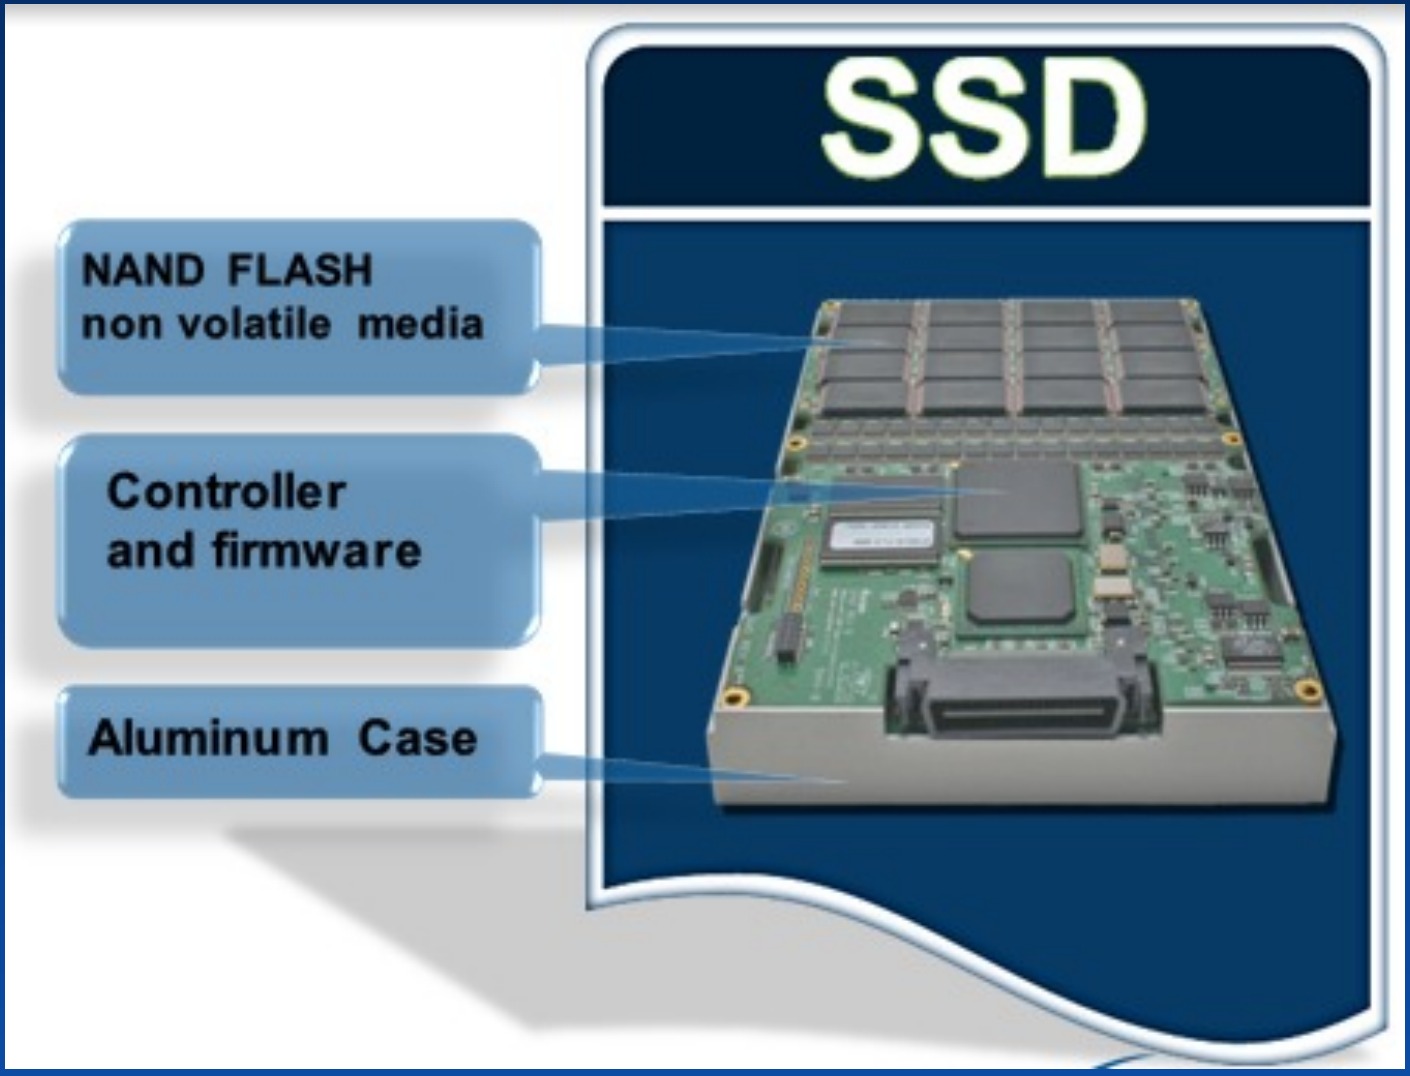

Importance of the Controller

Flash gets all the press…

but…it's static, not intelligent and does not interact with the host.

SSD is not Flash alone.

It is an intelligent system.

Enterprise SSD Must Haves…

Enterprise SSD Must Haves…

- Advanced Flash Management
- Power Backup Protection
- Full Data Path Protection
- Consistent High Performance
- Increased Write Endurance

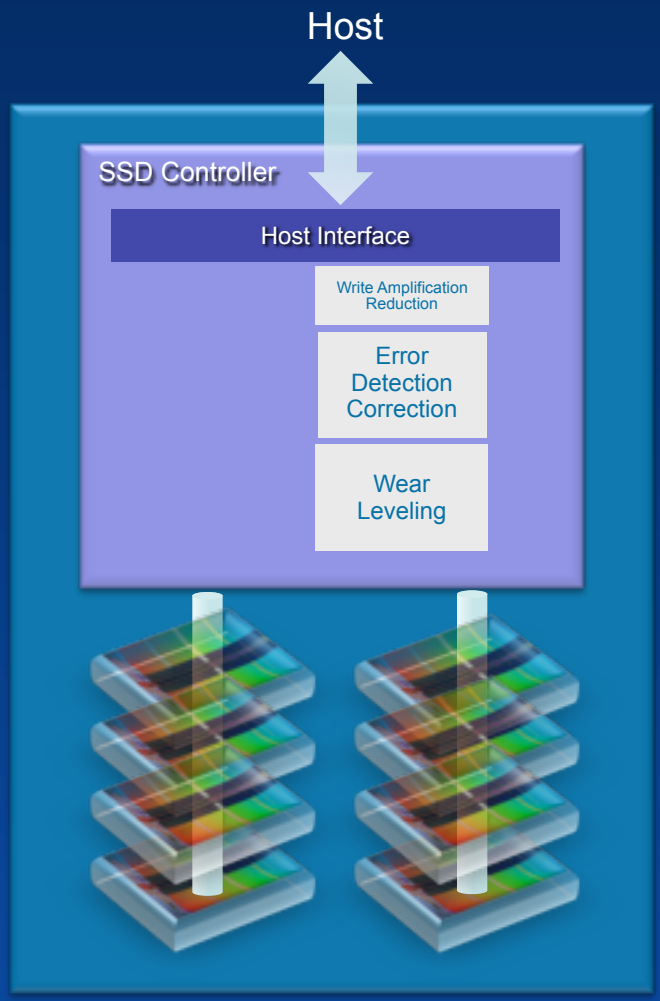

Enterprise SSD Must Haves…

- Advanced Flash Management
- Power Backup Protection
- Full Data Path Protection
- Consistent High Performance
- Increased Write Endurance

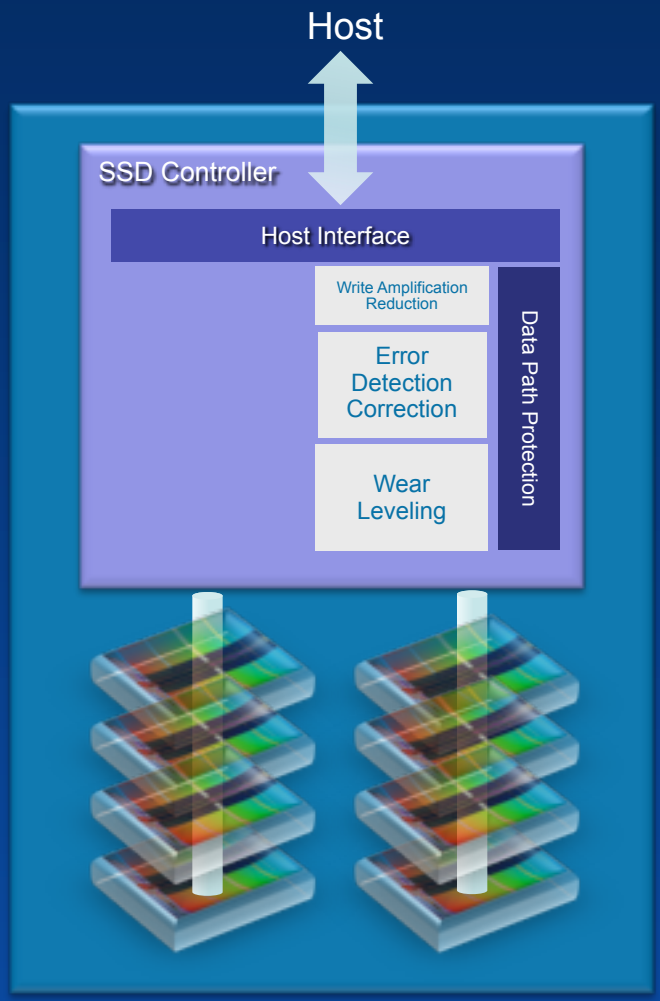

Enterprise SSD Must Haves…

- Advanced Flash Management
- Power Backup Protection
- Full Data Path Protection
- Consistent High Performance
- Increased Write Endurance

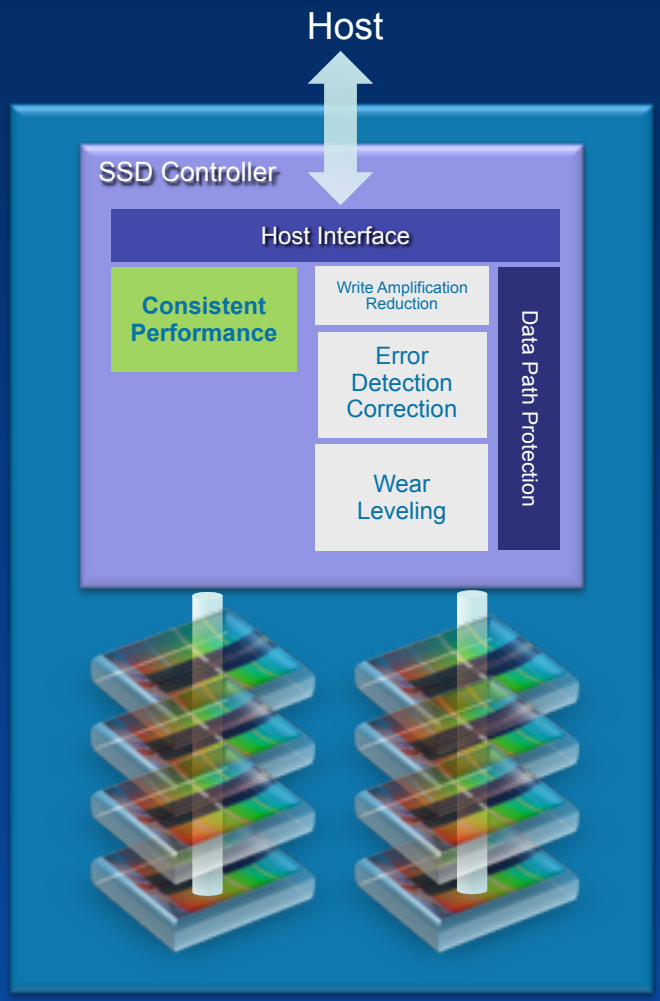

Enterprise SSD Must Haves…

- Advanced Flash Management
- Power Backup Protection
- Full Data Path Protection
- Consistent High Performance
- Increased Write Endurance

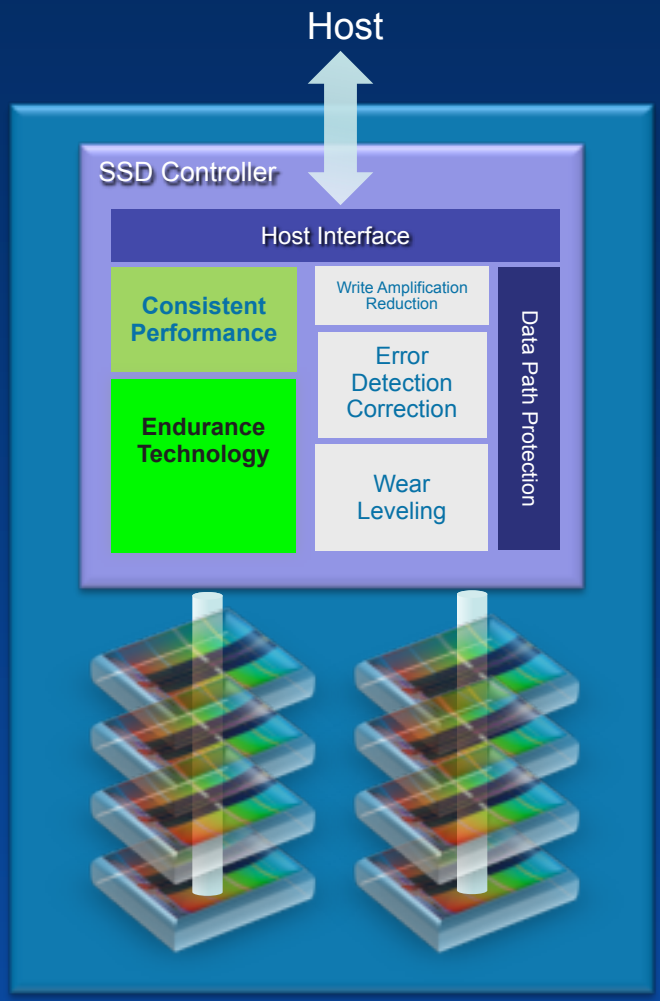

FLASH Media Trends

Flash Memory SUMMIT

Endurance and retention decrease significantly as process geometry shrinks

Improving ECC capability by itself cannot counteract this decrease in performance

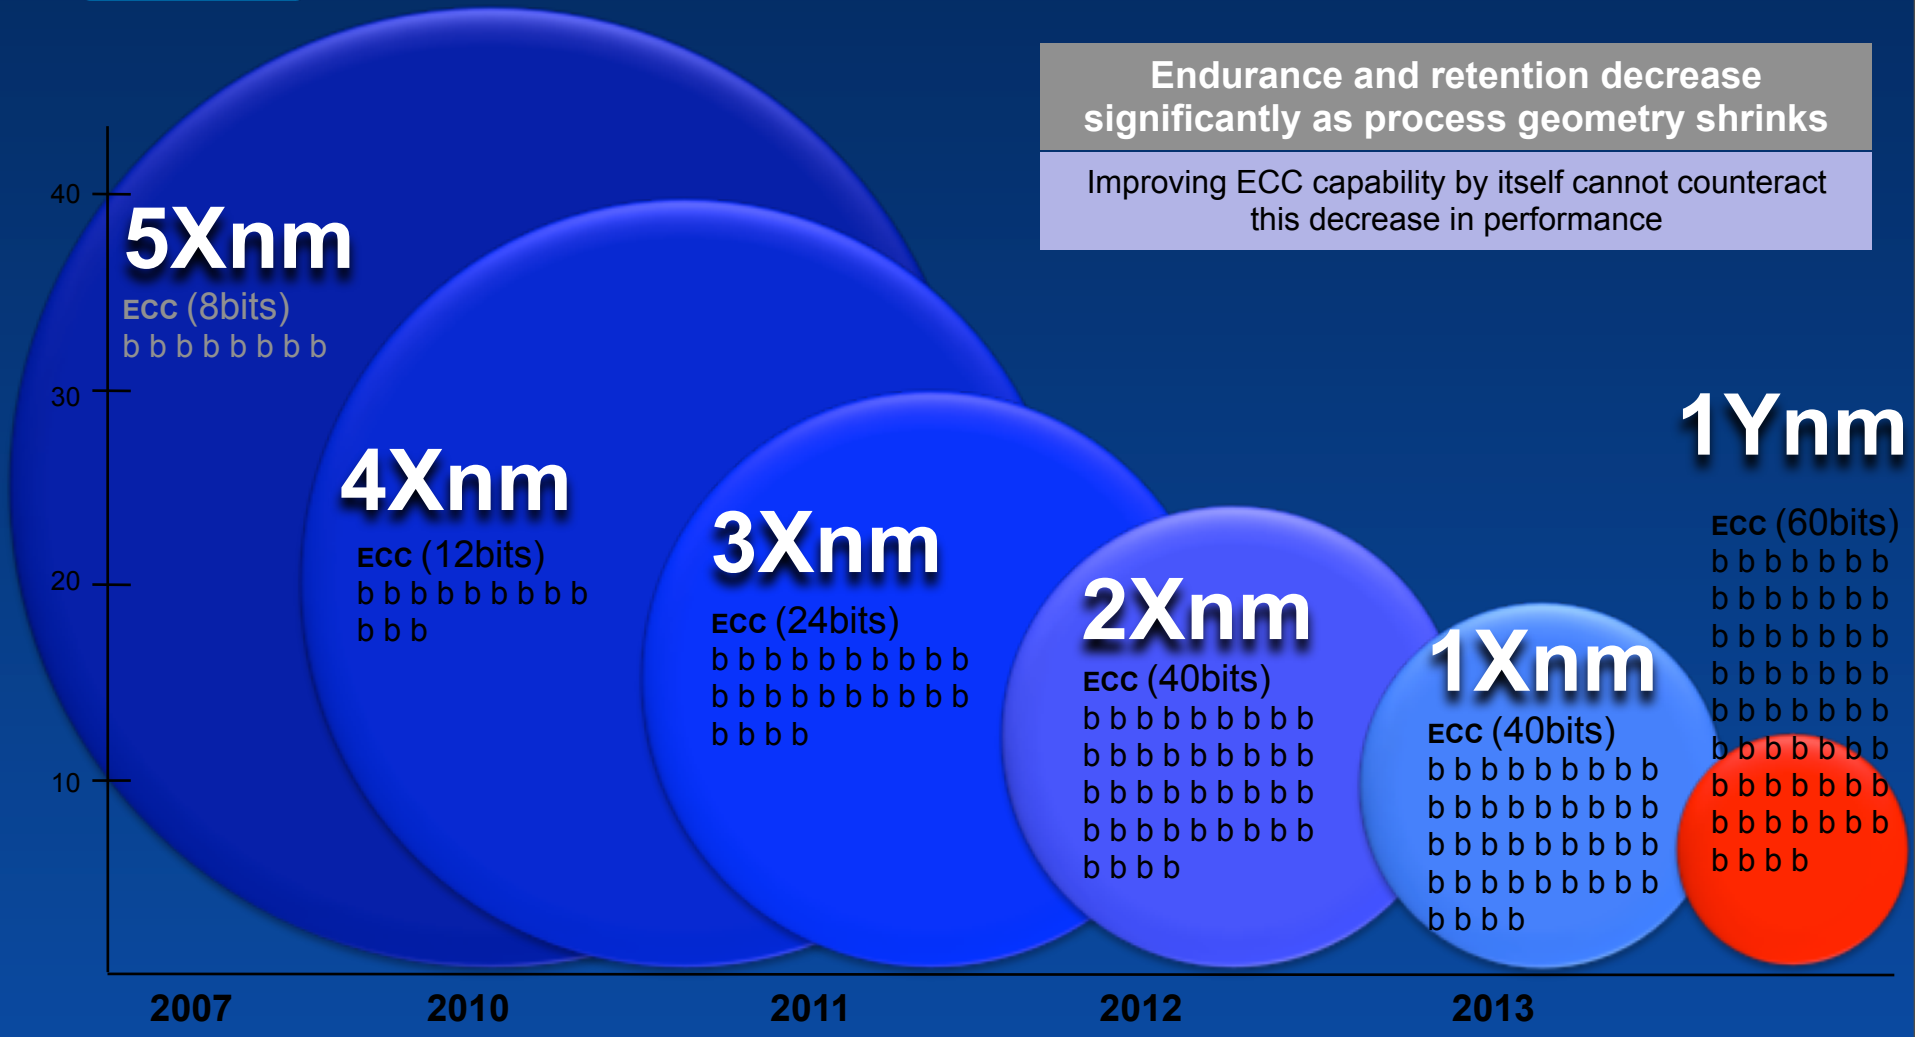

*Source: Average Supplier datasheets

Endurance Technology

Fills the Gap Between MLC and SLC Flash

Endurance Technology

Fills the Gap Between MLC and SLC Flash

Endurance Technology

Fills the Gap Between MLC and SLC Flash

Endurance Technology

Fills the Gap Between MLC and SLC Flash

Endurance Technology

Fills the Gap Between MLC and SLC Flash

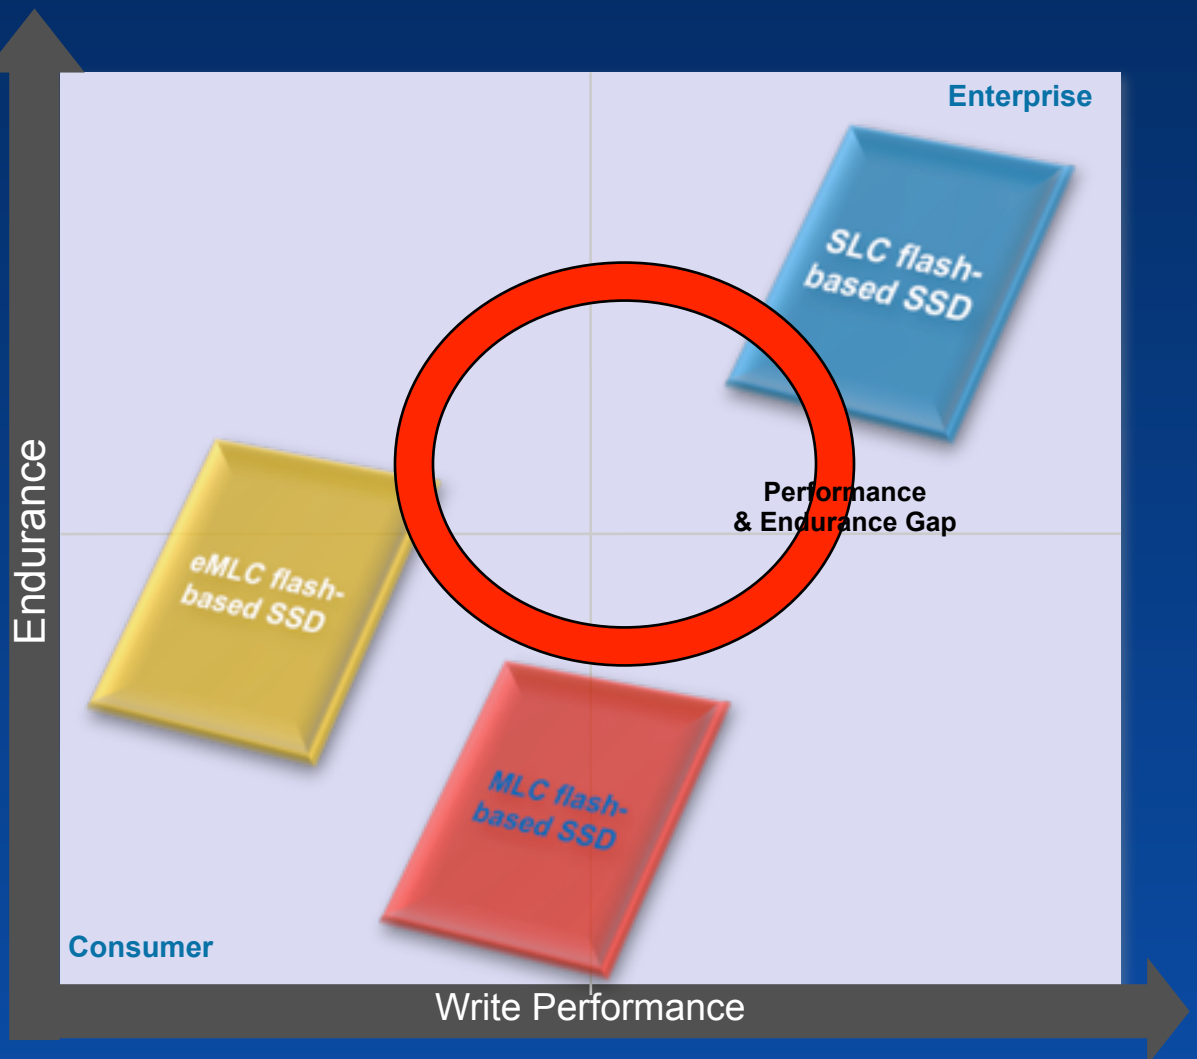

Endurance Technology

Fills the Gap Between MLC and SLC Flash

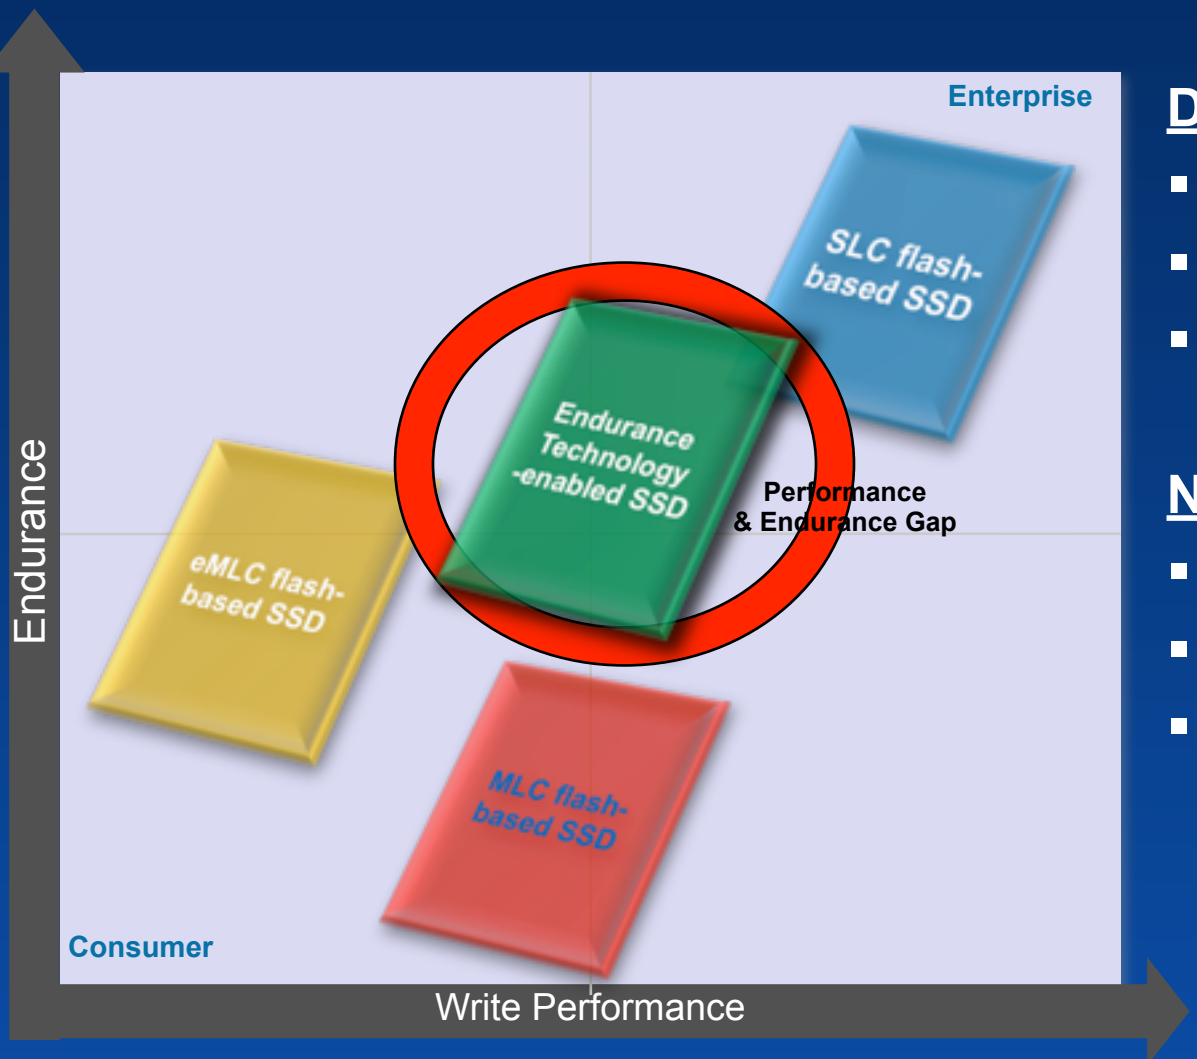

Data Centers Require:

- Performance
- Endurance
- Low **TCO**

Need to Enable MLC with:

- Endurance >1 year
- Consistent performance
- Increased reliability

Enterprise Hardware Only

Know Your Parts or
You'll Have No Parts

- Wear is measurable –
  - By tracking write activity you can measure how fast the SSD wears
- Wear is predictable –
  - Given use being measured, you can predict how long the SSD will live
- Example of "Service Life" for a 100GB SSD
  - Full 100% write data rate 24x7x365

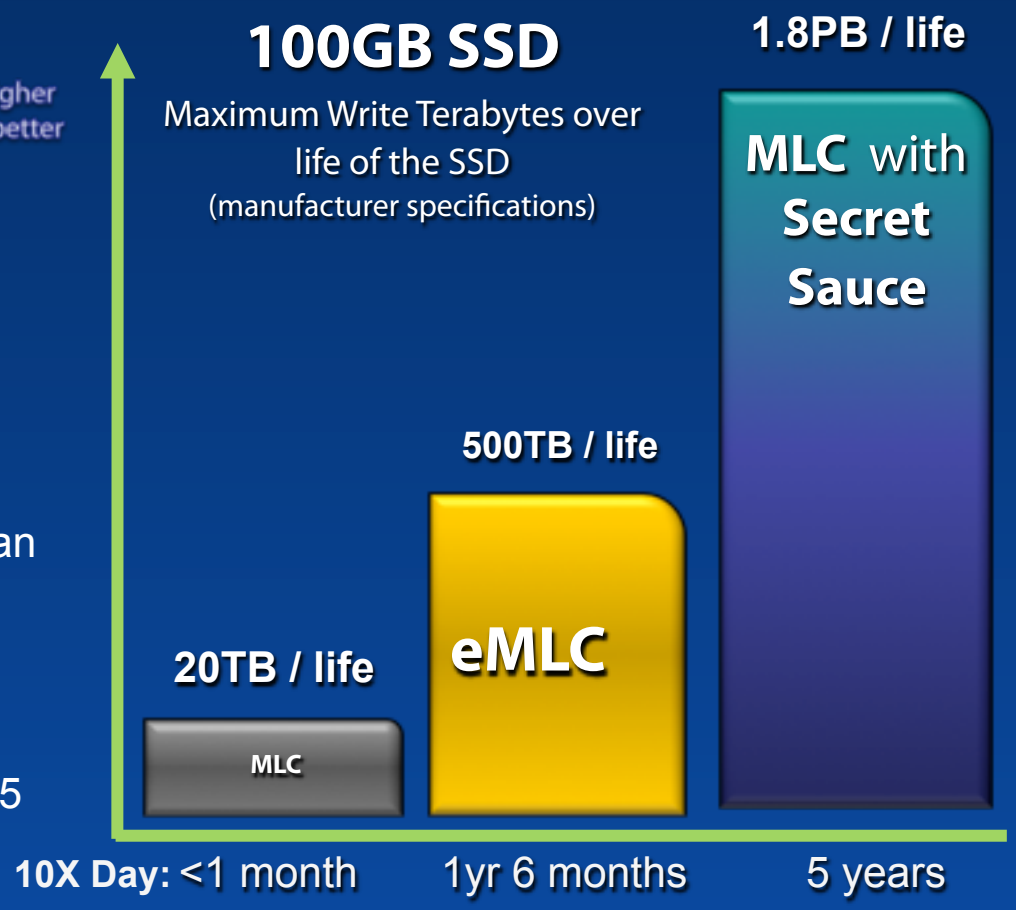

SSD Differences

Aren't all SSD's the same?

Not all SSD's are the same
There are fundamental design difference that affect performance, endurance and capability

Consumer SSD's

- Mobile and laptop devices
- Low duty cycles, infrequent write activity

Enterprise SSD's

- Performance
- 24x7x365 duty cycles
- Endurance and longevity expectations
- Data reliability expectations

Accelerate Access to Data™

www.stecssd.com
www.enterprise-ssd.techweb.com

Definitions

1. What is SSD over-provisioning?

The inclusion of extra storage capacity not accessible directly by the end-user in flash solid-state drive. That extra capacity is not visible to the host as available storage. SSD over-provisioning has one major benefit: It can increase the **endurance** of a solid-state drive by distributing the total number of writes and erases across a larger population of NAND blocks and pages over time.

2. What is Tier 0?

Tier 0 is the fastest and most expensive level of storage in the storage hierarchy. It's ideal for an enterprise that requires select applications to be quickly accessible. The storage hierarchy changed with the addition of tier 0 storage. It marked a change from simply moving less active data to slower, less expensive storage to focusing on efforts to move more active data to faster, more expensive storage and SSDs are driving that trend.

3. What are IOPS?

Input/output per second (IOPS) is the standard unit of measurement for the maximum number of reads and writes to non-contiguous storage locations.

4. What is solid-state storage program-erase cycle?

It's a sequence of events where data is written to solid-state NAND flash memory, then erased, and then rewritten. Program-erase cycles serve as a means for quantifying the endurance of a flash storage device. Flash memory devices have a limited number of PE cycles over time. The number of PE cycles a particular device can sustain before problems occur is varied with the type of technology.

5. What is TRIM?

TRIM is a SATA interface command that tells the NAND flash solid-state storage device which data to erase. The TRIM command allows the OS to notify the SSD which data in a set of pages can be overwritten, allowing the SSD controller to manage the erase process between the time when the host initiates a delete and the next write. By removing erase from the write process, writes occur faster.

Definitions

6. What is multi-level cell (MLC)?
MLC is flash memory that stores more than one bit per cell. It's less expensive than single-level cell (SLC) flash, which makes it a desirable option for solid-state storage. It does have a higher bit error rate than SLC flash because there are more opportunities for misinterpreting the cell's state. MLC should only be used in data center and enterprise environments with enterprise class SSDs.

7. What is a flash controller?
Flash controllers are the part of flash memory that communicate with the host device and manage the flash file directory. The controller is responsible for wear leveling, error correction, garbage collection and other critical functions.

8. Is there really any difference between consumer, client and enterprise SSDs?
Enterprise SSDs are different than consumer or "client" SSDs. Enterprise SSDs need to be able to handle many reads and writes for things such as database files, indexes, logs and other high-transaction types of operations. Enterprise SSDs are held to higher standards for performance, endurance, error recovery and continuous operation just as enterprise HDDs are held to similar higher standards.

9. Are there more differences regarding consumer vs enterprise storage?
Consumer-grade SSDs, like their desktop HDD counterparts, are not held to the same standards. For example, enterprise-class SSDs, like enterprise-HDDs, are expected to operate 24 hours per day and 7 days per week, while consumer-grade SSDs and HDDs are typically built to operate 8 to 10 hours per day and five days per week.

10. What are the standards typically used for SSDs endurance and performance?
Two of the industry standards bodies for SSDs, the Joint Electron Device Engineering Council (JEDEC) and the Storage Networking Industry Association (SNIA), have each published specifications for endurance (JEDEC) and performance (SNIA) that distinguish between enterprise and client SSDs.