Virtium®
Memory and Storage Solutions

Industrial Embedded Technology for an Interconnected World

Can You See Inside Your SSD?

Scott Phillips
Director, Product Marketing at Virtium

The Challenge

Limited visibility into how SSD is functioning

Standard SMART data is limited and often cryptic

SMART Details

C:\ 111.64 GB

ID	Description	Raw	Normalized	Threshold	Action
05	Re-allocated Sector Count	0	100	0	Ready for use.
09	Power-On Hours Count	894834	0	0	Ready for use.
0C	Power Cycle Count	12	100	0	Ready for use.
B5	Program Fail Count	0	0	0	Ready for use.
B6	Erase Fail Count	0	0	0	Ready for use.
C0	Unsafe Shutdown Count	6	100	0	Ready for use.
E1	Host Writes	129.31 GB	100	0	Ready for use.
E8	Available Reserved Space	0	100	10	Ready for use.
E9	Media Wearout Indicator	0	100	0	Ready for use.
F1	Total LBAs Written	129.31 GB	100	0	Ready for use.
F2	Total LBAs Read	70.72 GB	100	0	Ready for use.
F9	Total NAND Writes	102	100	0	Ready for use.

What is Really Happening?

Managing NAND reliability and endurance is tricky

- Application workload affects SSD usage acceleration (+/-), which impacts SSD available life
- Temperature affects reliability, endurance, and data retention

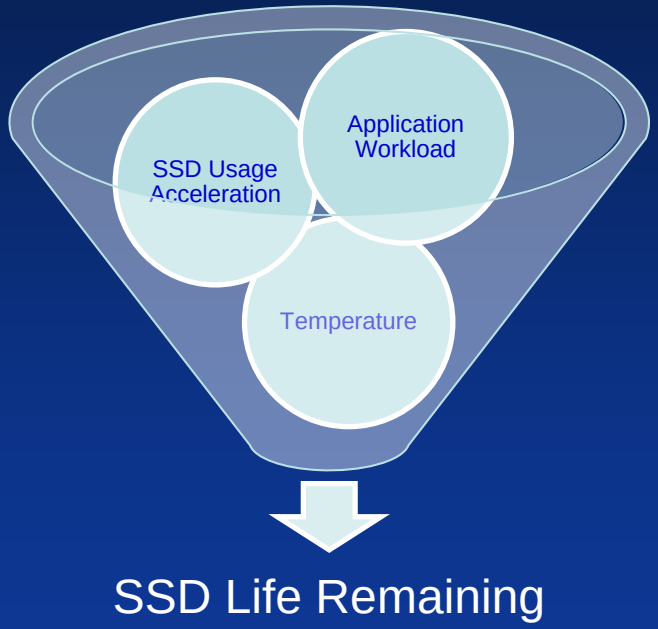

Why is it Happening?

Fundamental NAND architecture requires careful management

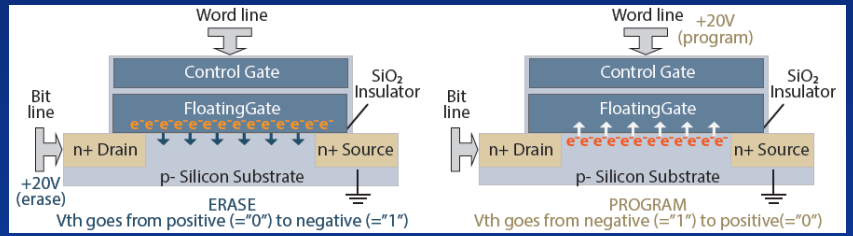

Planar NAND Flash Cell Structure

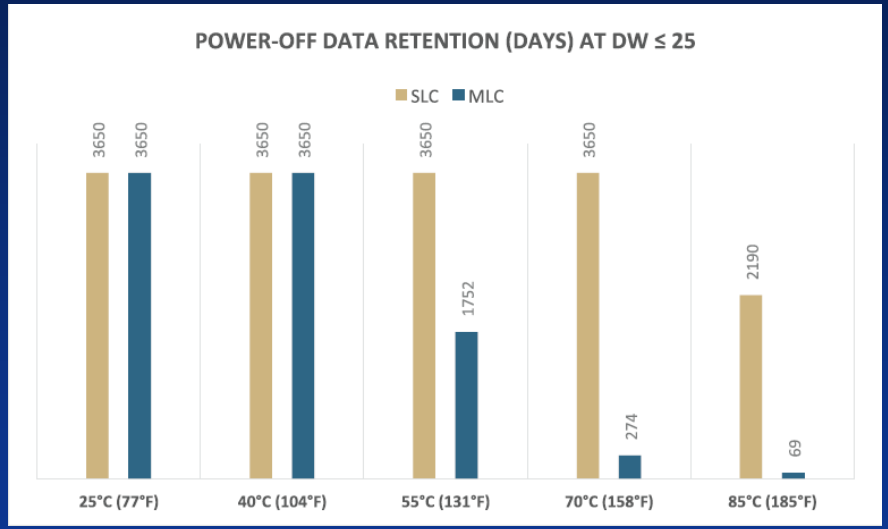

Temperature Effects on NAND Flash Data Retention

Why is it Happening?

NAND data structure requires even more management

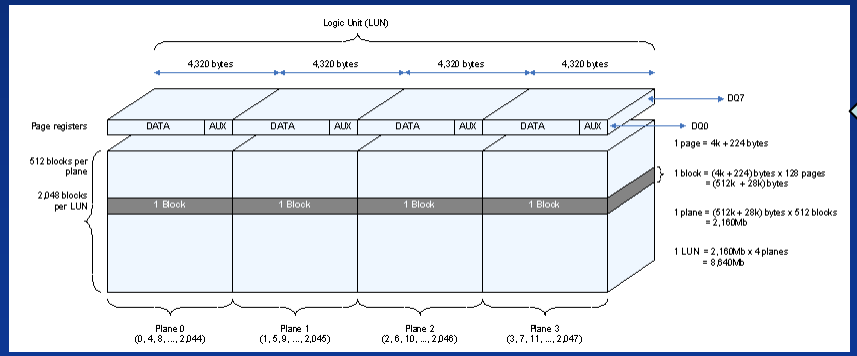

FIRMWARE SOFTWARE

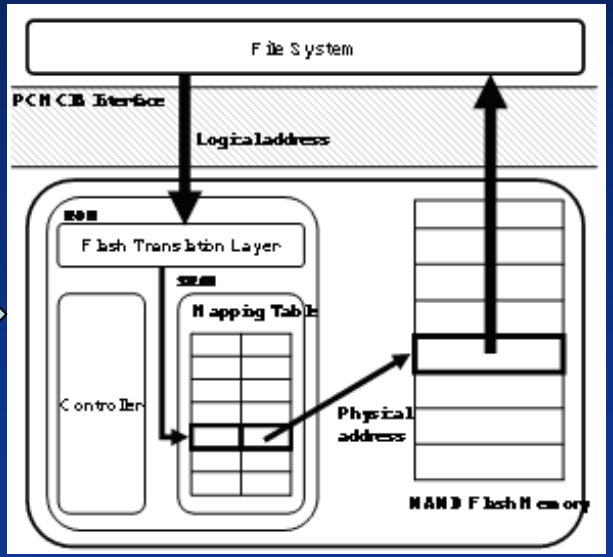

What to do About it

Provide ability to see what is happening inside the SSD

- ✓ Monitor workload and effects on SSD (e.g. life remaining)
- ✓ View usage acceleration impacts
- ✓ Develop “Smarter” SMART attributes
- ✓ View temperature vs. errors

Virtium’s vtView SSD Dashboard

Solving the Problem

Provide tools to optimize and control SSD operation

Optimize for Performance

- Firmware for specific application workloads

Optimize for Endurance

- Overprovisioning to lengthen available SSD life

Optimize for Reliability

- Built-in protection for high temp and power-loss

Optimize for Power

- Balance power and performance requirements

Proactive Notification

- Suggested maintenance, workload optimization

SSD Performance Tips

Data Alignment

- Understand the application's data transfer patterns
- Try to coalesce host data into larger files that are aligned on 4/8/16K page boundaries

Smaller, unaligned data causes increase in SSD usage acceleration, which in turn increases latency and wears the SSD more rapidly

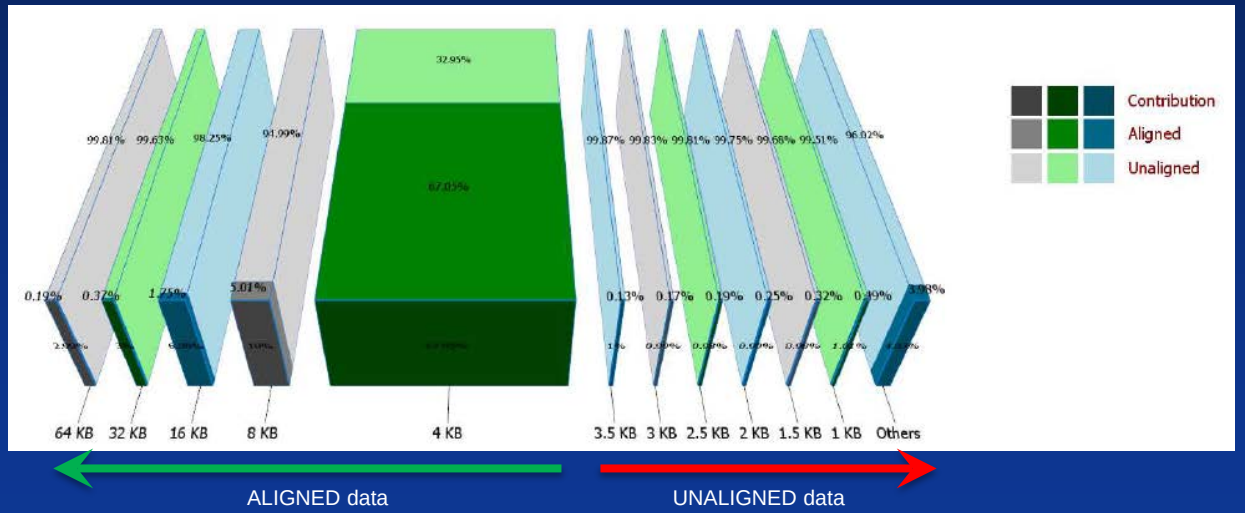

SSD Endurance Tips

Overprovisioning

- Selecting the right SSD configuration requires knowledge of application workload
  - Smaller, random data benefits greatly
  - Larger, sequential data benefits marginally

Effects of over-provisioning on SSD endurance and performance

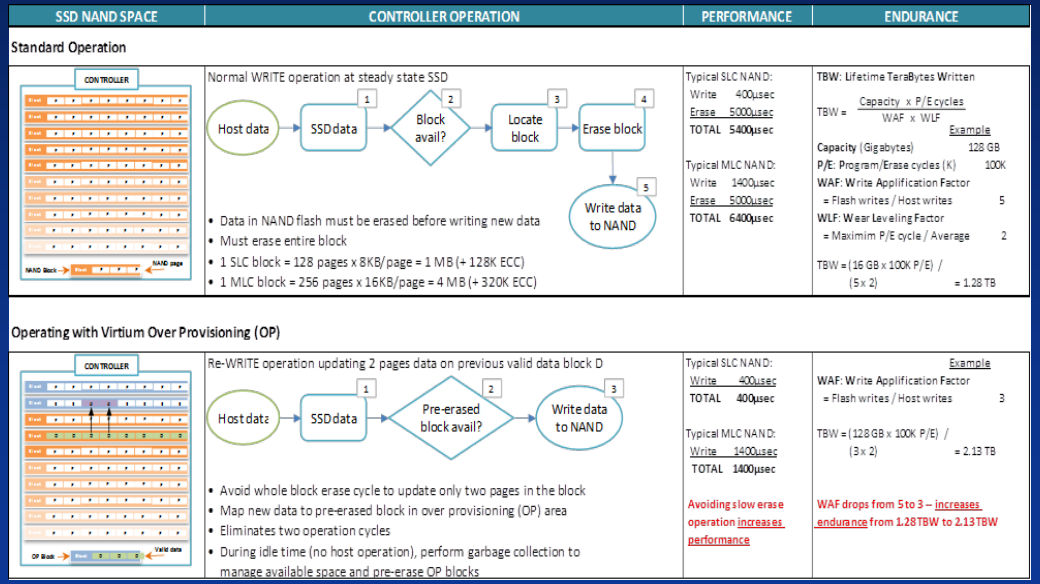

SSD Reliability Tips

Data protection during power loss

- In-flight data as well as firmware metadata at risk during power loss

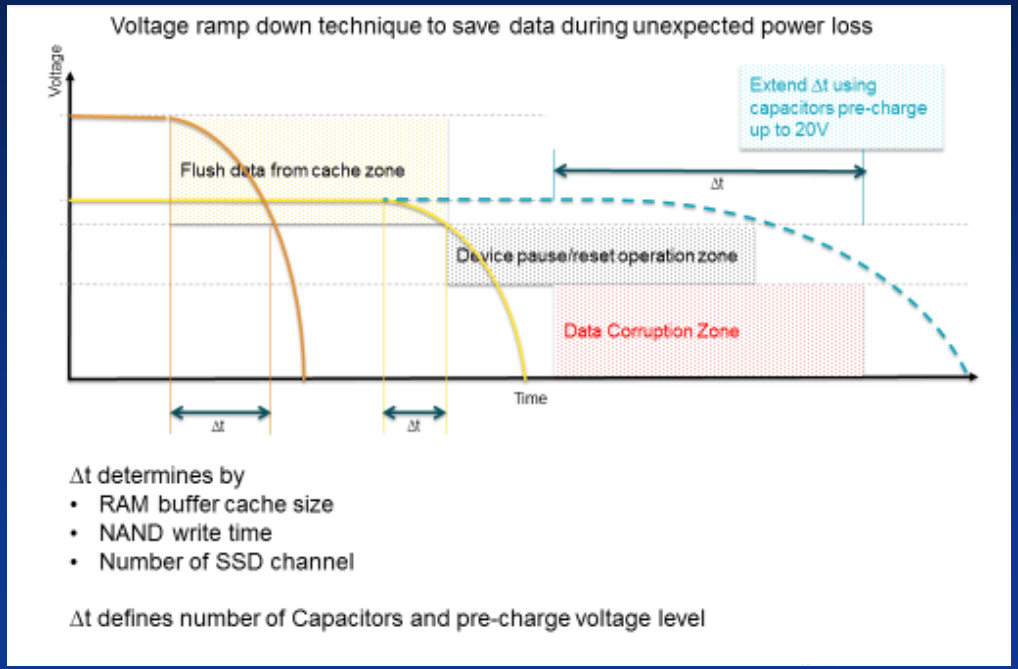

SSD Benchmarking Tips

Preconditioning

- Shows SSD's true "steady-state" performance
- **Goal:** Write to drives until a steady-state write pattern is achieved
- **Methodology:** Multiple full drive capacity writes with random, 4K, unaligned data

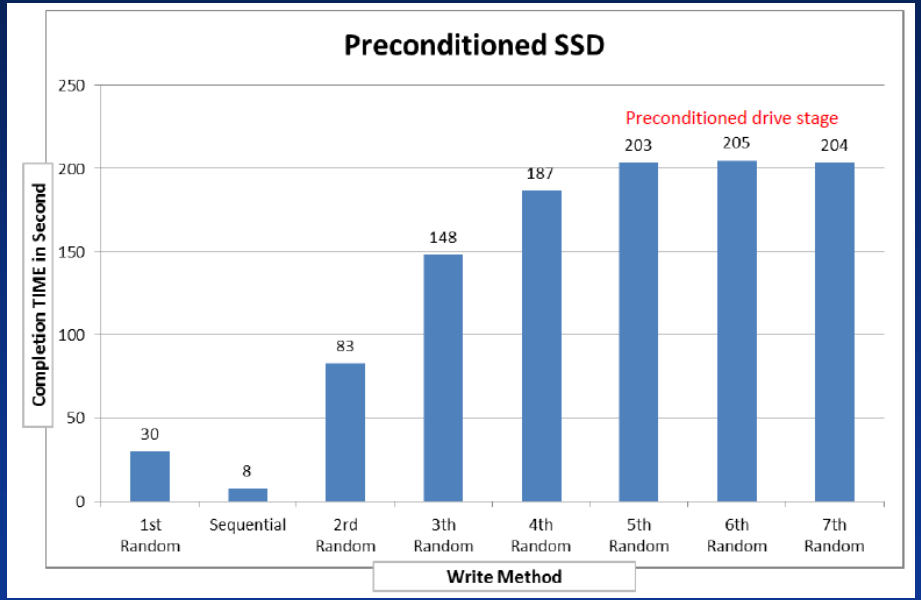

** Refer to JEDEC for published SSD benchmarking guidelines*

The Future of SSDs

“Software-defined storage” is coming

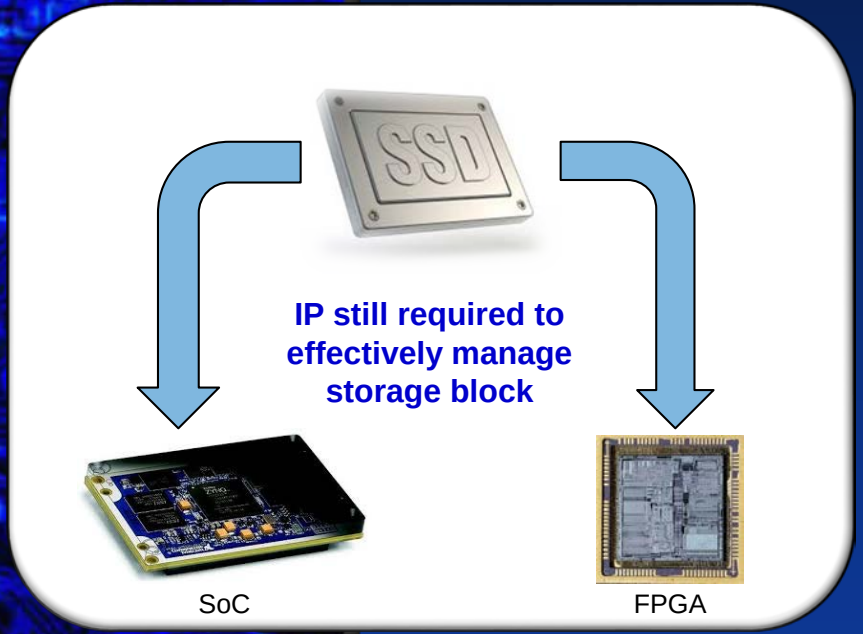

Will You be Ready?

Align storage needs with value-added capabilities

- Not enough to rely on “integrators” providing lowest cost
- Align with suppliers providing expertise and value-add technology:
  - Firmware development and flash management expertise
  - Software tools to monitor/manage SSD
  - Flexibility to adapt SSD to application workloads and environmental conditions

It’s a Wrap!

For more information, go to:

http://www.virtium.com/resources/product-collateral/

Virtium
Memory and Storage Solutions

Industrial Embedded Technology for an Interconnected World