



Benchmarking SSDs: The Devil is in the Preconditioning Details

Kent Smith
Sr. Dir. Product Marketing

Santa Clara, CA USA
August 2009





Agenda

- Preconditioning Explained
- System Setup Considerations
- Recycling (or Garbage Collection)
- “Time Zero” or Out-of-Box Drive State
- Sequential and Random Preconditioning
- Past Writes Affect Future Performance
- Conditioning Crossover
- Summary



What is Not Addressed Here

- Specific benchmark applications and settings
- Actual benchmark results
- Specific system setup steps

Preconditioning Explained

- HDD and SSD performance is affected by how data is transferred, e.g., sequential vs. random
- Unlike HDDs, the performance of an SSD is also dependent upon what has already been written to it
- Pre-conditioning an SSD assures repeatability and applicability of the test results



System Setup Considerations

Items to consider when benchmarking (or actually using) an SSD

- Advanced Host Controller Interface (AHCI) and Associated Driver
- Native Command Queuing (NCQ) and Queue Depth
- Offset and Alignment
- Operating System Background Operations
- Boot Drive vs. Secondary Drive

Recycling (or Garbage Collection)

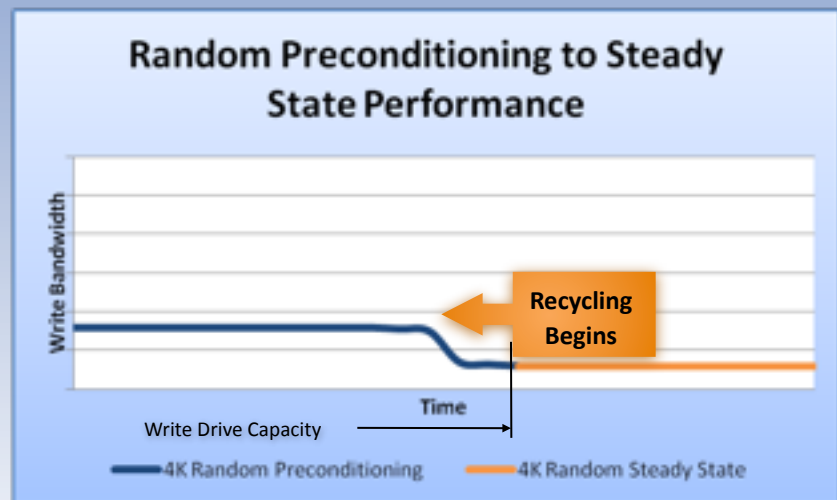
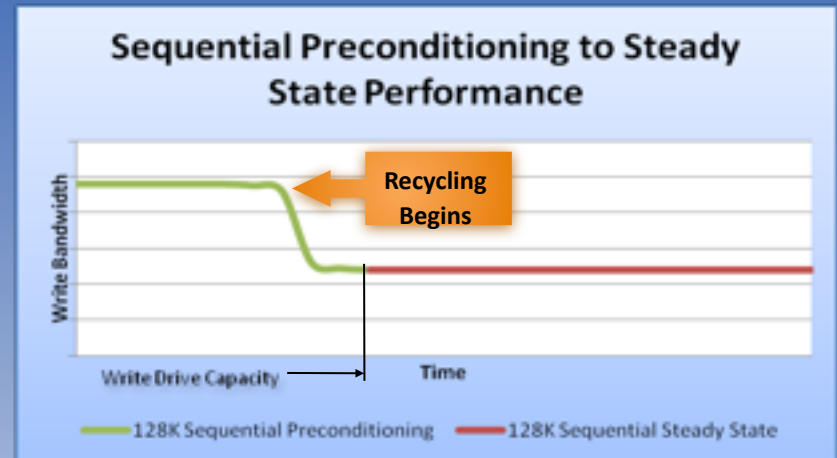
- NAND Flash must be erased in large blocks before it can be re-programmed
- When data is re-written, the original data is marked as “stale” and the new data is written to a different location
- When all blocks are written once, the drive must read all the good data around the stale data and move it to another location
- The old blocks are then erased and then new data is written to the newly freed-up blocks
- This activity causes large volumes of traffic on the flash bus and limits performance of new data from the host

“Time Zero” or Out-of-Box State

- When an SSD is first installed, all blocks are already erased and ready to be written
- No Recycling (or Garbage Collection) is required
- Initial performance measurements in this state are usually very high
- This is a poor indication of the real performance to be expected after each block is written one time (1st few hours or day)
- Some drives provide a “secure erase” utility to return a drive to this initial state, but it erases all data and is not indicative of true performance


Sequential & Random Preconditioning

- Write performance to an SSD will drop after recycling
- Green/Blue show performance before recycling
- Red/Orange is steady state operation after recycling turns on




Past Writes Affect Future Performance

- Unlike HDDs, when writing to an SSD the sequential or random nature of the writes will affect future performance
- Sequential writes will generally leave few large blocks of free space that make recycling faster
- Random writes will generally leave many small blocks of free space that make recycling slower




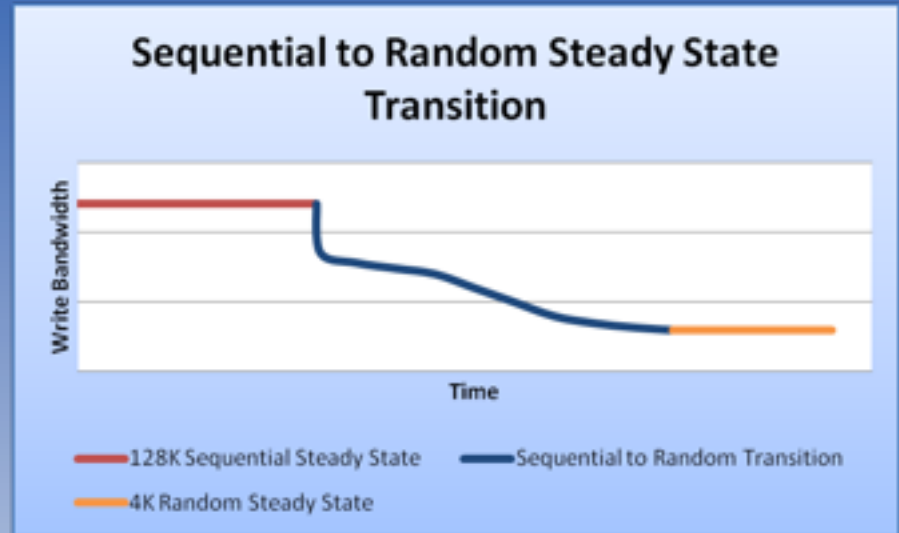
Free Space from Sequential Writes



Free Space from Random Writes

Conditioning Crossover

- When changing between random and sequential writes, be sure to wait until the writes are steady state
- Random writes look high initially after sequential writes
- Sequential writes look low initially after random writes





Summary

- Testing for short periods of time will not necessarily disclose steady state performance
- Predictable and steady state performance measurements require preconditioning between sequential and random writes



SandForce SSD Processors

**World-class reliability, performance, & power efficiency
for enterprise and client SSD applications**

- Come visit us at **booth #507** to see our DuraClass™ technology in action with the latest MLC flash technology
 - 30K+ sustained, random 4K write IOPS to a single SATA drive
 - Ultra-fast Windows 7 booting and user experience
 - 130K sustained, random 4K write IOPS through a PCIe card solution connected to four drives
 - 15K sustained, random 4K write IOPS through Emulex's next gen high-performance SAS-SATA embedded storage bridge technology