



The Myth of SSD Testing

– When It Is Done Right, The Result Can Be Astonishing and Instrumental for SSD Design

Frank Shu

Allion Test Lab, Inc.

(frankshu@allion.com)



Overview

Solid State Drive testing has been a hot topic for years, this is because testing SSD's quality (performance and reliability) in the way that truly represents its state and value is critical for the future of a SSD product.




SSD is expected to have a superlative performance, and yet the users of SSD also want low cost with a warranted reliability – this becomes the key motivation for SSD companies to be innovative and creative. At the end, to measure the real success of a SSD product must rely on a comprehensive testing methodology.

This presentation discusses some important aspects of SSD testing and astonishing data found in the SSDs on the market today. When the test is done right, the result will not only represent the true SSD quality, but also provide solid evident and focus for SSD quality and design improvement.



SSD Performance Testing


- Workload based performance testing
 - Typical (4K, 8K) R/W performance is interesting, but workload based performance is what counted by user experience
 - The performance can be different more than 100% between the SSDs on the market





SSD Performance Testing – Cont.

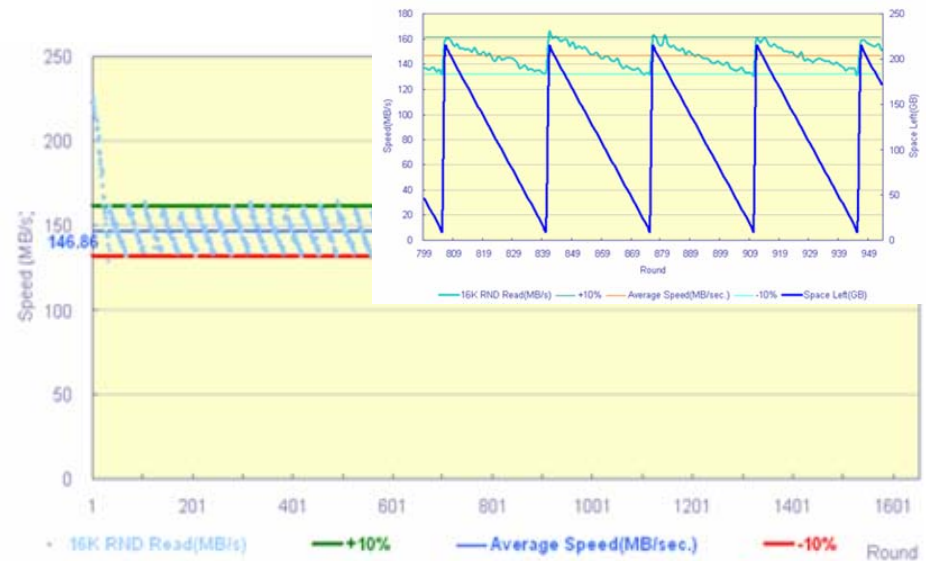
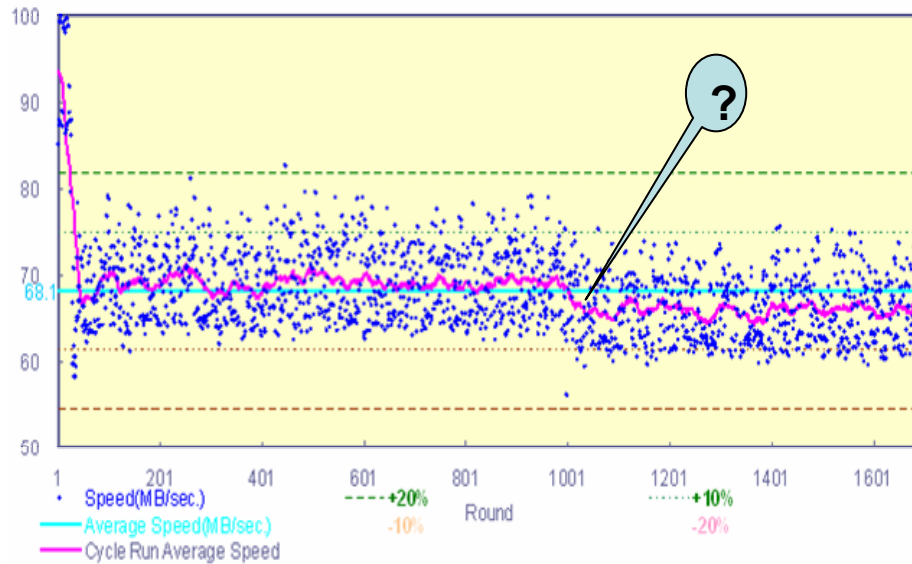
- Write cache can boost writes, but the device must handle Flush well to gain the full value.





SSD Performance Testing – Cont.

- Testing performance degradation
 - FOB is not the main focus here because
 - Native to NAND SSD
 - Well understood and easy to detect
 - To focus on long term performance degradation
 - Fragmentation
 - Swap space wear out
 - ...





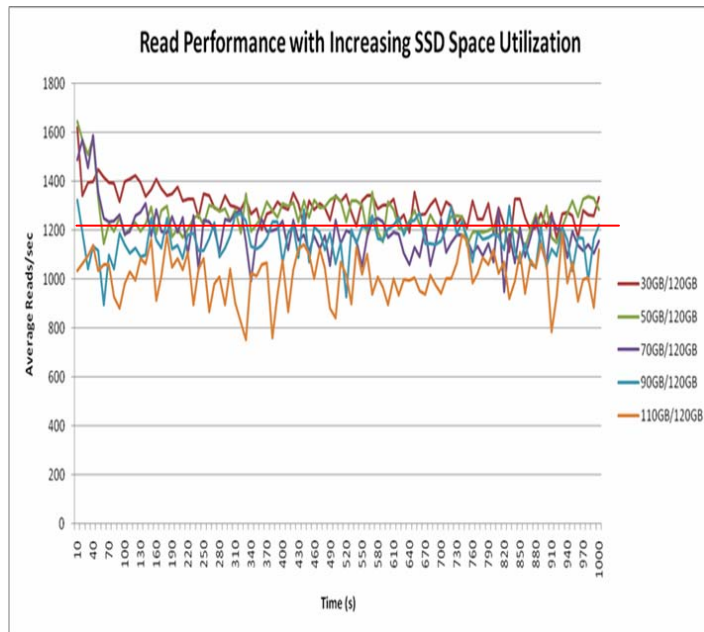
SSD Performance Testing – Cont.

- SSD IO queue (NCQ) depth impact
 - Device internal parallel operations
 - “Locality” in reference to SSD NAND flash page
- Data compression effectiveness
 - Compressible data
 - Compressing/De-compressing speed
- Write aggregation effectiveness
 - Transfer length
 - Random vs. Sequential




SSD Performance Testing – Cont.

- Making sure the implementation of Trim is right
Trim can recover the performance loss related to free space

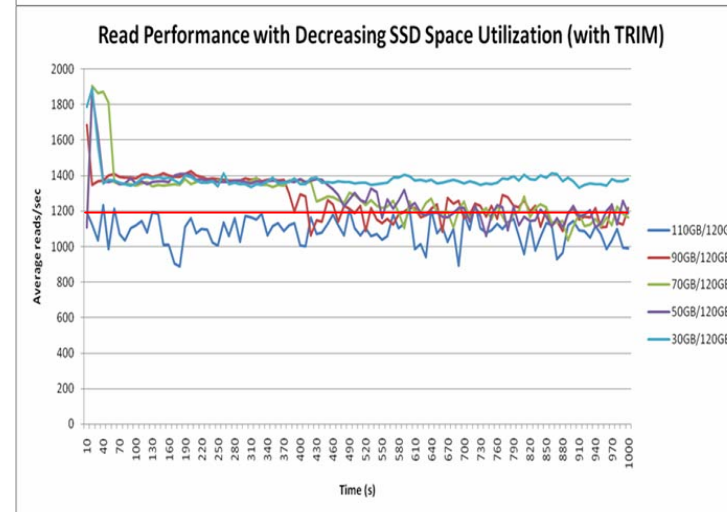


4k Random Read with 64K Random Write

No Trim



With Trim







SSD Performance Testing – Cont.

- Implementing Trim without pitfall

Trim operation can choke IOs, it can be detected





Testing to Ensure Data Integrity

- What happens and should happen with Flush
 - Factsheet – not all SSDs with volatile write cache honor Flush
 - Ignoring Flush can have a noticeable performance gain, but it could cause data corruption too
 - Data corruption is different from data loss; it is worse...
 - Even with the protection built in file system today, data corruption can happen
 - Fully test your implementation before product ship
 - Knowing the possible root cause(s)
 - Having a full test coverage with well designed test cases



Testing to Ensure Data Integrity – Cont.

- Do Trim right
 - Thinking the scenario like Snapshot – the protected area gets update before it is re-written
 - Be thoughtful when doing Self-Trim

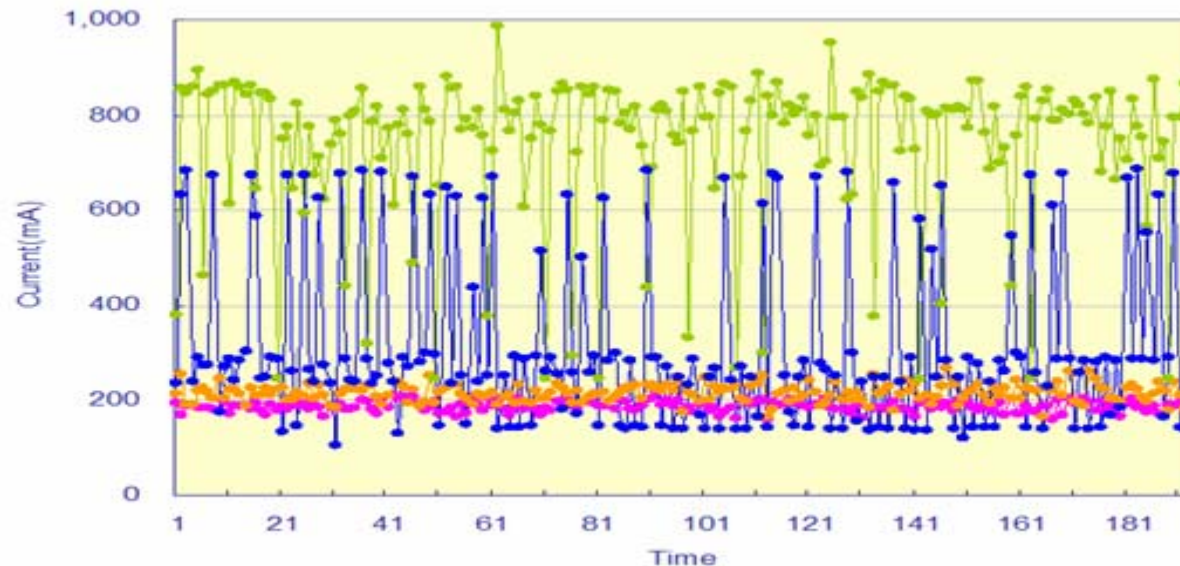
Trim notification is processed systematically by other part of OS and/or filter driver(s) through its layers before device gets it. Sniffing file system and initiated trim operation by third party apps or device itself can catastrophically damage user's data

 - Your last defense against possible data corruption is the validation test that covers all known scenarios

Power Consumption Analysis

- On-Off vs. On-All-Time, the test can tell the difference

Two implementations of power schema(e.g. 4K random write):



- Impact to battery time and battery life
 - Battery time – 10% saving; a extra hour battery time for a tablet PC with 10 hours battery time
 - Battery life – depending on the power regulation circuit



Reliability Testing

- Data retention test
 - Retention is an important aspect of SSD, it is a NAND level testing
 - The testing is workload independent per se
 - Easy to implement, but time consuming, even with high temperature acceleration
 - JEDEC specs are commonly followed by NAND vendors
 - Retention can be verified with third party testing



Reliability Testing – Cont.

- Endurance validation test

- SSD endurance can be validated
- Endurance testing is to test SSD's endurance, not just NAND
- ECC is counted as part of SSD endurance beside NAND quality, so do SSD controller and firmware
- Temperature does matter; most data miss-compare errors occurred at either low temperature (25°C) or high temperature (70°C)

A sample error report

```

*****Round 583***** at 25°C
2010/08/15 17:29:19 Create G:\_1
2010/08/15 17:29:19 .....
2010/08/15 17:30:19 .....
2010/08/15 17:30:19 .....
2010/08/15 17:30:24 Starting Compare C:\3GB test file and G:\_1
2010/08/15 17:31:41 .....
2010/08/15 17:31:41 Compare Log:
    - C:\3GB test file : G:\_1 -- includes left-only,right-only,differing files
      .\mpeg2-59f.mpeg different(both have identical times)          2316146e
        c4db02b4
      .\valkyrie.avi    different(both have identical times)        7a123902    654d89a9
      -                - 2 files listed
Flash Memory Summit 2011 2010/08/15 17:31:41 Compare FAILED
Santa Clara, CA *****

```



Thank You

(frankshu@allion.com)



Appendix – Table of Content

- SSD Performance Testing
 - Workload based performance testing
 - Does SSD Write-cache really help your SSD performance
 - Performance degradation
 - Queue-depth and Compression
- Testing to Ensure Data Integrity
 - What happens with Flush
 - Do Trim right
- Power Consumption Analysis
 - On-Off vs. On-All-Time
 - Battery time and battery life
- Reliability Testing
 - Data retention test
 - Endurance validation test