



Merits and Methods of I/O Traced-Based Performance Benchmarking of Client SSDs

Harry Pon

Intel Corporation: NVM Solutions Group
Folsom, California

August 21, 2012



Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Material in this presentation is intended as product positioning and *not* approved end user messaging.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

For more information go to <http://www.intel.com/performance>

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

"Ultrabook is a trademark of Intel Corporation in the U.S. and/or other countries"

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2012 Intel Corporation. All rights reserved.



Presentation Outline


- Position Statement
- Performance Measurement Score Card
- Propose Method of I/O Trace-Based Client SSD Testing
- Example Application of I/O Trace-Based Testing
- Summary

Position Statement

- SSDs Are the Marvel of PC Storage
 - Ultrabook™, notebook, tablets



- Hard Disk Drives are NOT the Same as Solid-State Drives
 - Should NOT be Tested the Same



- Correctly Test SSDs with I/O Traces and Methodologies
 - HDD benchmarks do not show SSD Goodness.

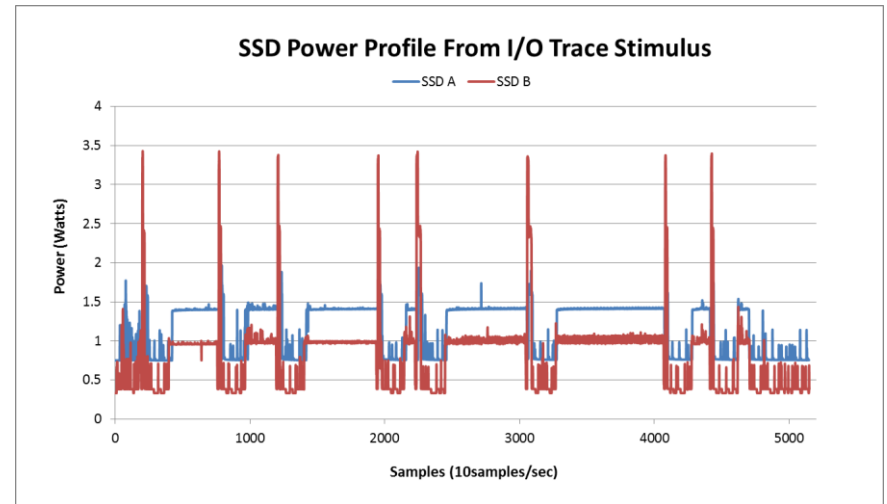
Performance Measurement Score Card

Measurement Method	Pro's	Con's
<p>Pure Synthetic: Algorithmic I/O Sequence</p>	<ol style="list-style-type: none"> 1. Point or Instance Analysis 	<ol style="list-style-type: none"> 1. Artificial; Not real life 2. Does not test the SSD as it is used 3. Overall Results are inconclusive
<p>Application Based: Custom Script Calling of Real Applications to some time attribute</p>	<ol style="list-style-type: none"> 1. Emulation is closer to real usage involving real user applications 	<ol style="list-style-type: none"> 1. Script or software coding is needed. 2. Purchase of Applications is Expensive 3. More Non-Storage Performance Variables Introduced Via the Hardware 4. Emulation of Idle times is inaccurate 5. Ambiguous Figure of Merit
<p>I/O Trace: Record User I/O Sequence and then Replay</p>	<ol style="list-style-type: none"> 1. Real Life Workload 2. Tests True SSD Performance 3. Accurate 4. No need to purchase new applications 5. Versatility 	<ol style="list-style-type: none"> 1. Limited Tracing Software Offerings 2. Existing Software Offerings Have. Limited Playback Capabilities

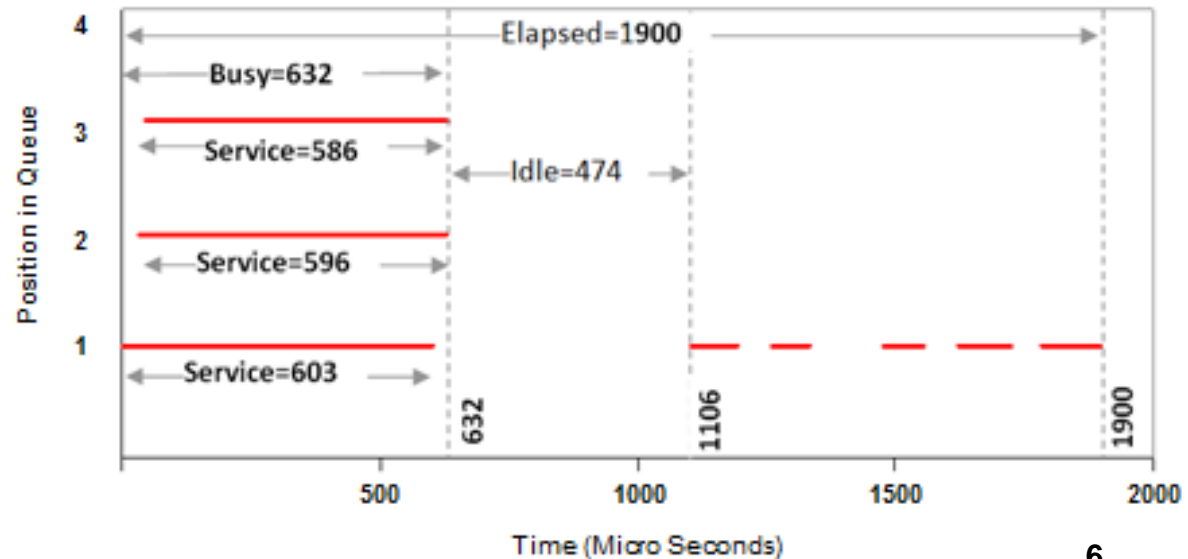
I/O Trace-Based SSD Testing is The Way to Go!

Example Application of I/O Trace Based Testing

- Focused & Reproducible Real Life Workload Methods to Optimize SSDs
 - Power Profiles

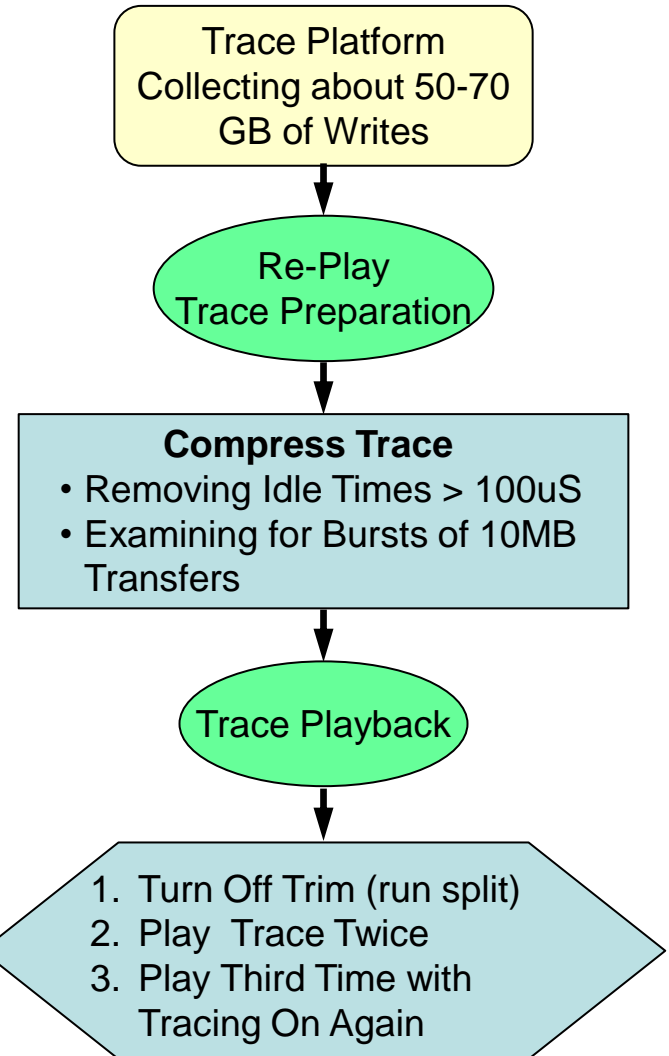


- Understand True SSD MB/s Performance Over an Elapsed, Busy, and Serviced Time Interval



Proposed Method of Trace-Based Testing of SSDs

- I/O Trace Tests SSD As Designed
- Intel is Developing I/O Trace-Based Test Methodology
 - 1) Drive Prep
 - 2) Trace Collection
 - 3) Trace Playback Preparation
 - 4) Trace Playback
 - 5) Analyzing Results
- Performance Comparison
 - Equivalent Trace Results in MB/s
 - Develop a New Figure of Merit Metric



***I/O Trace-Based
Performance Testing
Does Justice to SSDs***

- Artificial Synthetic Testing of Client SSDs Provides Inconclusive Benchmark Results
- I/O Trace-Based Testing Does Justice to SSDs
 - Real Life Workload
 - Accurate
 - Easy and Economical
 - Versatile

***I/O Trace-Based Testing of Client SSDs
is The Only Way to Go!***



Back Up

HDD vs. SSD Test Paradigm Shift



Moving Mechanical Parts
Magnetic Storage Medium

- Read Modified Writes

Slow Random/Sequential R/W
Not Rugged
High Power

No Moving Parts
NVM NAND Flash Storage Medium

- Program Pages & Erase Blocks

Fast Random/Sequential R/W
Rugged
Low Power

Why Test them the Same?