Capturing Client IO Workloads for SSD Performance Evaluation

You can join the effort to define a Client Composite Workload!

Eden Kim, CEO
Calypso Systems, Inc.
Chair, SNIA SSSI Technical Development Committee

The Case for Developing a Client Composite Workload

The Need

SSD Performance is affected by the user application and software/hardware environment

Designers/users need to understand Client workloads in order to optimize/select SSDs

The Problem

Most IO Performance testing is done using either synthetic benchmark or some IO operation trace based tests

These tests do not test SSDs to the user's actual workload

Test workloads often do not reflect user workloads at the SSD Device level

Solution & Application

Gather empirical IO operation performance metrics from actual Client users to define client workloads

Synthesize this data into a Composite Client Workload to test and compare client SSDs

Opportunity

Download the SSSI Workload IO Capture Program (WIOCP) software on to your Windows PC

Run the program for a week and send the IO metric logs to the SSSI

Help the SSSI define a Client Composite Workload for client SSD comparison & test

Footnote:

Workload = Access Pattern

What the heck is a "workload" anyway?

SSD Testing Working Definition

Workloads are logical stimuli – or access pattern(s) - that are applied to a mass storage device and are described by random or sequential accesses, data transfer sizes and read/write mixes

Access patterns

Access patterns are generally expressed as a combination of three attributes:
RND/SEQ access to the mass storage device; Block Size; Read/Write ratio

Access patterns are generally expressed as:

RND/SEQ %	Block Size (KiB)	Read Write Mix (Ratio %)
RND	4KiB	100% Write
75:25 RND/SEQ	8KiB	95:05 RW Mix
SEQ	128KiB	100% Read

Access Streams

Workloads can have a single (monotonic) access pattern (i.e. one stream) or be comprised of many (composite) access pattern (multiple streams)

The Need: Why is it so hard to define a Client SSD IO workload?

SSD NAND Flash Performance

- Performance changes over time
- Performance is affected by the SSD "write history"
- Performance depends on the type of workload
- There are a multitude of client user workloads

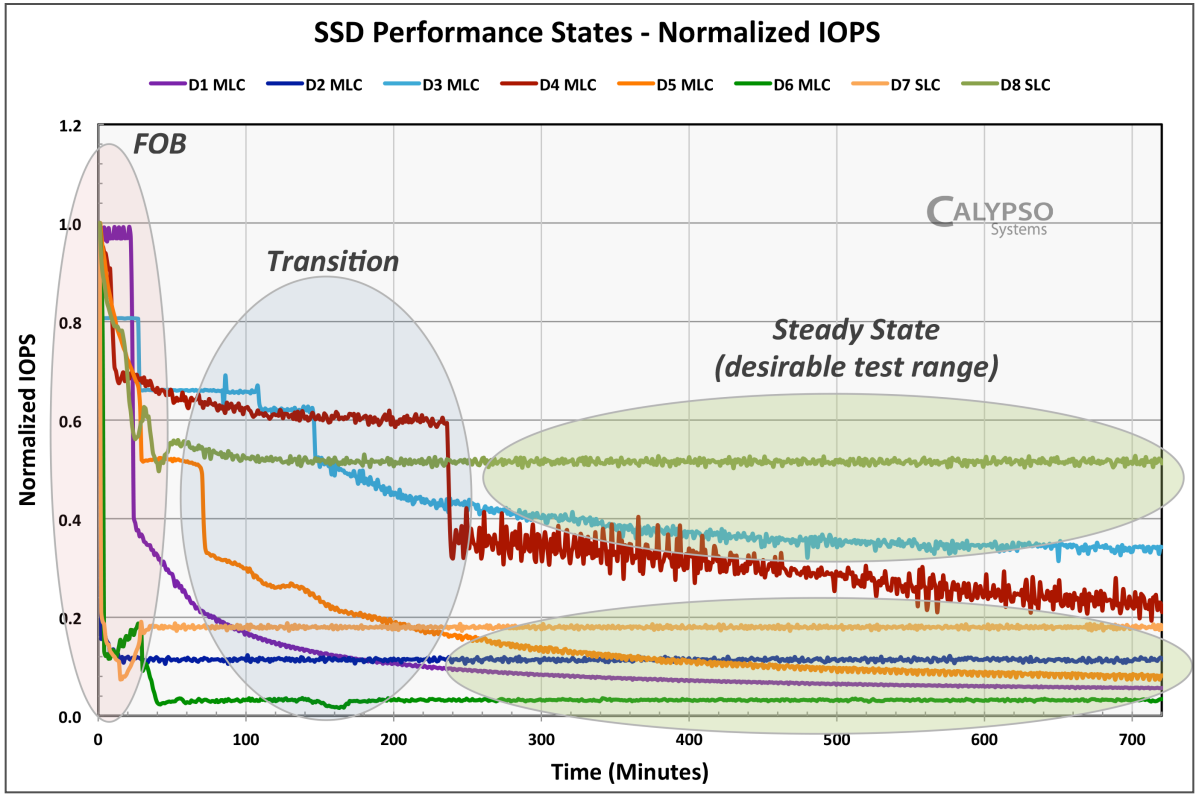

The Need: What is so hard about capturing a Client IO workload profile?

SSD Client User Workloads – Influence of the IO Stack

Performance is affected by the Hardware/Software IO stack

Measurements depend on where in the IO stack you take them

IOs can be cached, split or coalesced

Each Client User Workload consists of different software applications and different hardware platforms

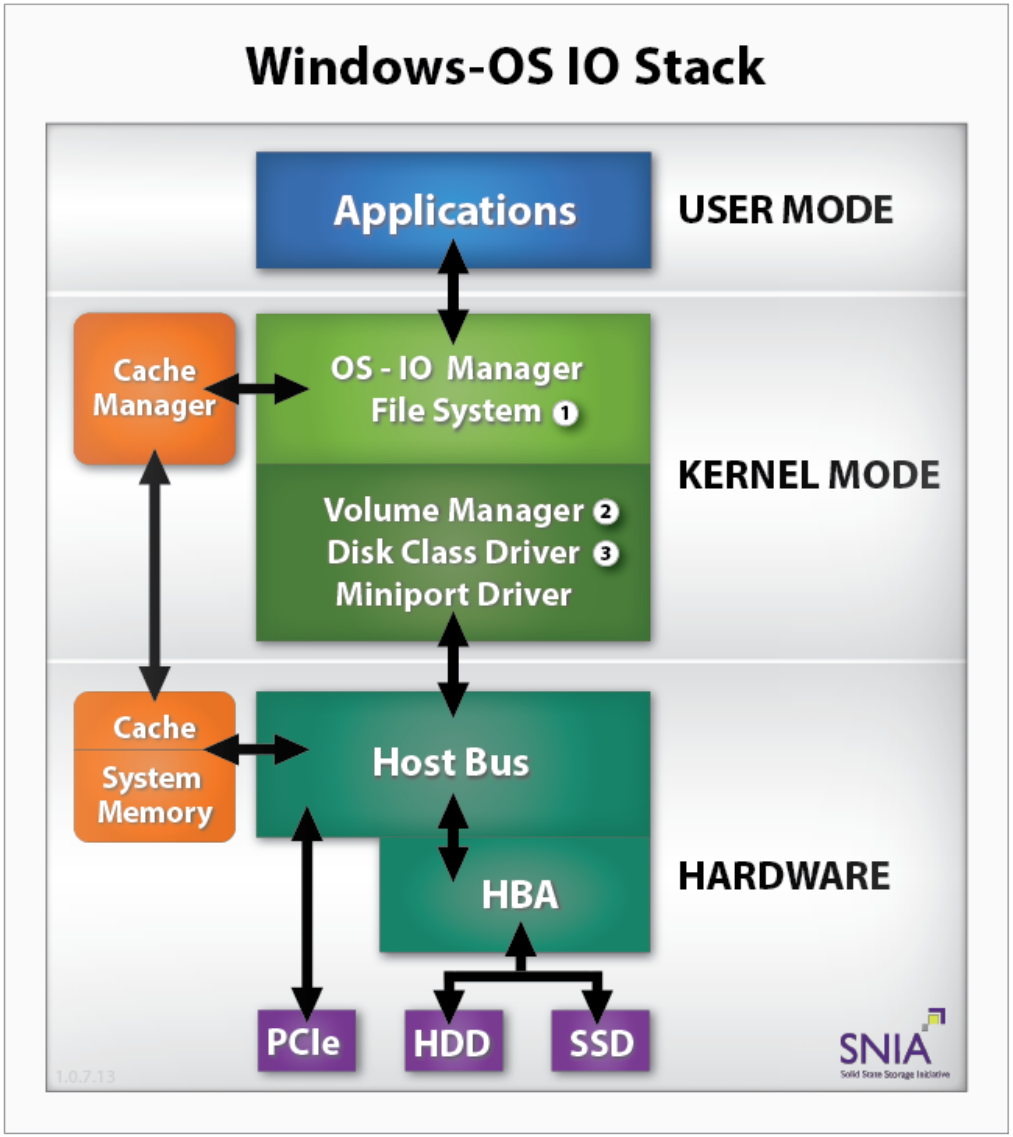

The Problem: Existing Client SSD workload tests miss the mark!

Synthetic Benchmarks

Each is based on an unique formula or recipe of access patterns

There is no empirical basis supporting these as being universal or describing “typical” user workloads

IO Operation Trace based Benchmarks

Traces are captured from HDD or SSD systems – HDD traces do not equal SSD workloads

Traces can be unique to the system they were captured on – which affects the trace itself

Issues related to timing in workload playback (e.g. how to treat long idle times)

Collected traces are extremely large size files – can make collection and replay impractical

Functional Streams

Concatenated stream of various client activities are often used as a playback test stimulus

They are relevant only to that specific stream and the order in which the workload segments are run

Corner Case Stress Tests

Use specific access patterns considered to be representative of user workloads (Read/Write intensive, Mixed or VOD)

By definition, corner case tests examine behavior outside the normal range of operation

A single SSD will respond differently to each corner case stimuli – thus how to merge the streams?

The Solution: Gather Empirical client user data!

WIOCP – Workload IO Capture Program

Captures empirical IO workload measurements during everyday usage

Metrics are concurrently collected at three different levels in the IO Stack

The metrics are processed to define a series of access patterns seen during the user capture period

SSSI posts workload captures to the SNIA IOTTA repository and are available for download at http://iotta.snia.org

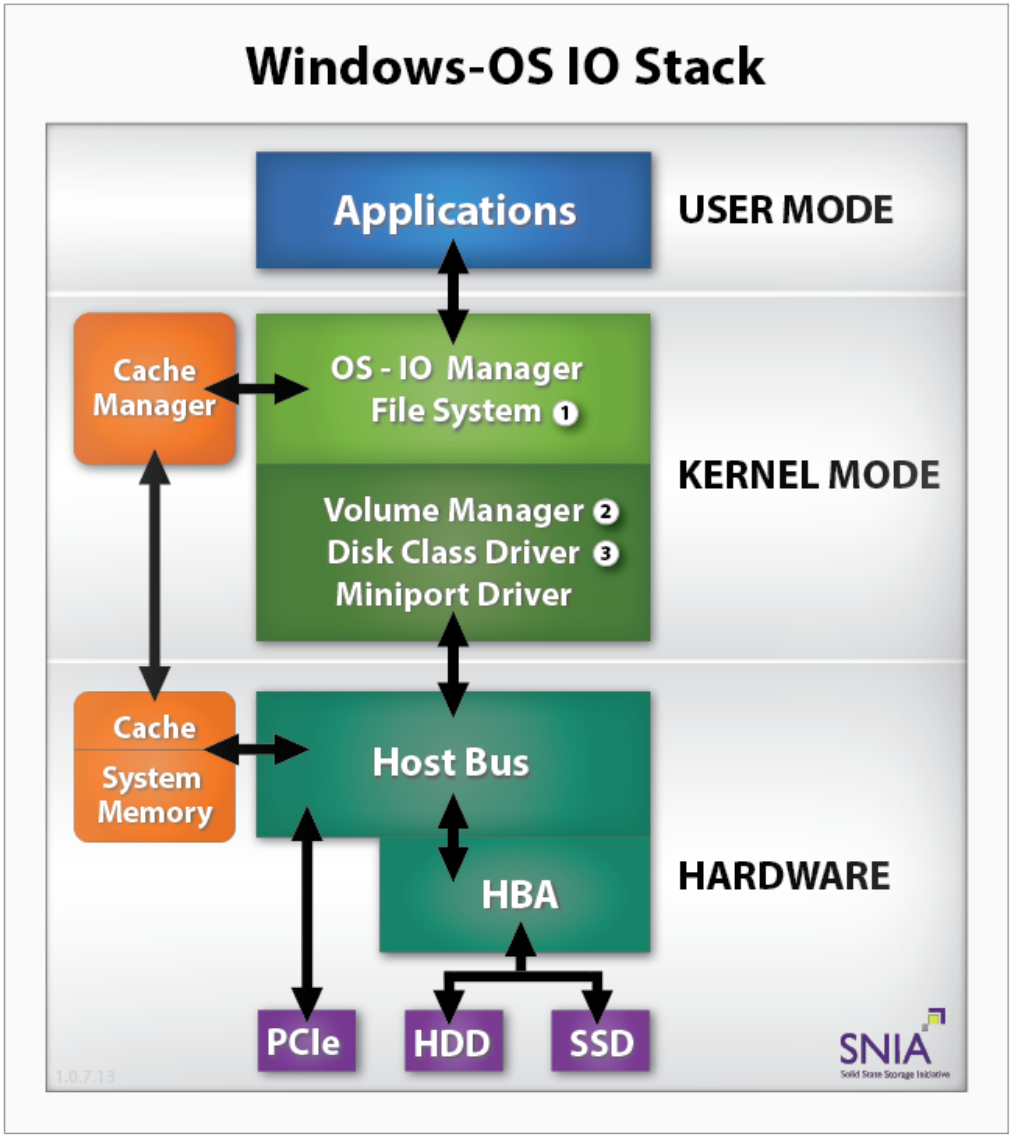

Note: The WIOCP tool and associated software is a product of hyperI/O, LLC and is being offered to the SSSI for its sole use. Contact hyperI/O at www.hyperIO.com for other uses.

Application: Creating a Client Composite Workload from a WIOCP IO capture

WIOCP Capture processed file:

Collection of user IO Statistics. No personally identifiable information or data content is captured.

Collection of Access Pattern statistics from the system boot HD / SSD.

Device level (Block IO) statistics captured to allow SSD comparison.

66 Streams. 66 different access pattern streams were observed during the observation period.

MonitoredName	Transfer Size	Access Percentage Overall	Overall Percentage Read	Overall Percentage Random	PTS CCW Read Random Access %	PTS CCW Read Seq Access %	PTS CCW Write Random Access %	PTS CCW Write Seq Access %
\Device\Harddisk0\DR(	512	0.47346%	0.12115%	0.38678%	0.04447%	0.07668%	0.34231%	0.00999%
\Device\Harddisk0\DR(	1024	0.27726%	0.22813%	0.19602%	0.14762%	0.08051%	0.04840%	0.00073%
\Device\Harddisk0\DR(	1536	0.26278%	0.19317%	0.18805%	0.15291%	0.04026%	0.03514%	0.03448%
\Device\Harddisk0\DR(	2048	0.15726%	0.13177%	0.12373%	0.09883%	0.03294%	0.02490%	0.00059%
\Device\Harddisk0\DR(	2560	0.15925%	0.13452%	0.12874%	0.10531%	0.02922%	0.02344%	0.00129%
\Device\Harddisk0\DR(	3072	0.15023%	0.09946%	0.12711%	0.07769%	0.02176%	0.04941%	0.00136%
\Device\Harddisk0\DR(	3584	0.14163%	0.10938%	0.09496%	0.06320%	0.04618%	0.03176%	0.00049%
\Device\Harddisk0\DR(	4096	30.40168%	15.61771%	24.42588%	10.98243%	4.63528%	13.44344%	1.34053%
\Device\Harddisk0\DR(	4608	0.02406%	0.01560%	0.01842%	0.01003%	0.00557%	0.00839%	0.00007%
\Device\Harddisk0\DR(	5120	0.09441%	0.02413%	0.06209%	0.01017%	0.01396%	0.05192%	0.01835%
\Device\Harddisk0\DR(	7168	0.00651%	0.00094%	0.00589%	0.00042%	0.00052%	0.00547%	0.00010%
\Device\Harddisk0\DR(	8192	5.18946%	3.06523%	4.65418%	2.76954%	0.29569%	1.88464%	0.23959%
\Device\Harddisk0\DR(	12288	1.00964%	0.45636%	0.71705%	0.25811%	0.19825%	0.45894%	0.09434%
\Device\Harddisk0\DR(	13312	0.00672%	0.00338%	0.00501%	0.00167%	0.00171%	0.00334%	0.00000%
\Device\Harddisk0\DR(	13824	0.00198%	0.00021%	0.00188%	0.00010%	0.00010%	0.00178%	0.00000%
\Device\Harddisk0\DR(	16384	50.75401%	19.33076%	29.26776%	2.74548%	16.58528%	26.52228%	4.90098%
\Device\Harddisk0\DR(	17408	0.00066%	0.00045%	0.00056%	0.00035%	0.00010%	0.00021%	0.00000%
\Device\Harddisk0\DR(	18944	0.00094%	0.00010%	0.00080%	0.00000%	0.00010%	0.00080%	0.00003%
\Device\Harddisk0\DR(	19456	0.00296%	0.00136%	0.00153%	0.00101%	0.00035%	0.00052%	0.00108%
\Device\Harddisk0\DR(	20480	0.02894%	0.00230%	0.02768%	0.00157%	0.00073%	0.02612%	0.00052%
\Device\Harddisk0\DR(	21504	0.00122%	0.00031%	0.00091%	0.00000%	0.00031%	0.00091%	0.00000%
\Device\Harddisk0\DR(	23040	0.00104%	0.00010%	0.00098%	0.00003%	0.00007%	0.00094%	0.00000%
\Device\Harddisk0\DR(	24576	0.16684%	0.07602%	0.14180%	0.05666%	0.01936%	0.08514%	0.00568%
\Device\Harddisk0\DR(	28672	0.05011%	0.04548%	0.03932%	0.03479%	0.01069%	0.00453%	0.00010%
\Device\Harddisk0\DR(	30720	0.00091%	0.00038%	0.00052%	0.00000%	0.00038%	0.00052%	0.00000%
\Device\Harddisk0\DR(	31232	0.00098%	0.00007%	0.00073%	0.00000%	0.00007%	0.00073%	0.00017%
\Device\Harddisk0\DR(	32768	6.22942%	3.68386%	5.30110%	2.89863%	0.78523%	2.40247%	0.14309%
\Device\Harddisk0\DR(	38400	0.00059%	0.00003%	0.00056%	0.00000%	0.00003%	0.00056%	0.00000%
\Device\Harddisk0\DR(	40960	0.01073%	0.00940%	0.00850%	0.00717%	0.00223%	0.00132%	0.00000%
\Device\Harddisk0\DR(	48128	0.00059%	0.00056%	0.00021%	0.00017%	0.00038%	0.00003%	0.00000%
\Device\Harddisk0\DR(	49152	0.14978%	0.08131%	0.11377%	0.06007%	0.02124%	0.05370%	0.01477%
\Device\Harddisk0\DR(	52736	0.00024%	0.00000%	0.00024%	0.00000%	0.00000%	0.00024%	0.00000%
\Device\Harddisk0\DR(	54784	0.00035%	0.00035%	0.00024%	0.00024%	0.00010%	0.00000%	0.00000%
\Device\Harddisk0\DR(	56320	0.00028%	0.00007%	0.00024%	0.00003%	0.00003%	0.00021%	0.00000%
\Device\Harddisk0\DR(	56832	0.00035%	0.00007%	0.00028%	0.00000%	0.00007%	0.00028%	0.00000%
\Device\Harddisk0\DR(	57344	0.00575%	0.00526%	0.00456%	0.00418%	0.00108%	0.00038%	0.00010%
\Device\Harddisk0\DR(	61440	0.10524%	0.10510%	0.02277%	0.02267%	0.08243%	0.00010%	0.00003%
\Device\Harddisk0\DR(	63488	0.00010%	0.00003%	0.00007%	0.00000%	0.00003%	0.00007%	0.00000%
\Device\Harddisk0\DR(	65536	2.33885%	1.17832%	1.27586%	0.30310%	0.87522%	0.97276%	0.18777%
\Device\Harddisk0\DR(	67072	0.00003%	0.00003%	0.00000%	0.00000%	0.00003%	0.00000%	0.00000%
\Device\Harddisk0\DR(	69632	0.01480%	0.01264%	0.00850%	0.00644%	0.00620%	0.00205%	0.00010%
\Device\Harddisk0\DR(	73728	0.01821%	0.01142%	0.01191%	0.00731%	0.00411%	0.00460%	0.00219%
\Device\Harddisk0\DR(	76800	0.00045%	0.00007%	0.00038%	0.00000%	0.00007%	0.00038%	0.00000%
\Device\Harddisk0\DR(	77824	0.00345%	0.00003%	0.00320%	0.00000%	0.00003%	0.00320%	0.00021%
\Device\Harddisk0\DR(	81920	0.01424%	0.00540%	0.01247%	0.00481%	0.00059%	0.00766%	0.00118%
\Device\Harddisk0\DR(	98304	0.01236%	0.00167%	0.01212%	0.00150%	0.00017%	0.01062%	0.00007%
\Device\Harddisk0\DR(	102400	0.01121%	0.00958%	0.00686%	0.00554%	0.00404%	0.00132%	0.00031%
\Device\Harddisk0\DR(	106496	0.00094%	0.00087%	0.00056%	0.00049%	0.00038%	0.00007%	0.00000%
\Device\Harddisk0\DR(	111104	0.00014%	0.00010%	0.00003%	0.00000%	0.00010%	0.00003%	0.00000%
\Device\Harddisk0\DR(	122880	0.02215%	0.01257%	0.01532%	0.00644%	0.00613%	0.00888%	0.00070%
\Device\Harddisk0\DR(	126976	0.19073%	0.18791%	0.06554%	0.06408%	0.12383%	0.00146%	0.00136%
\Device\Harddisk0\DR(	131072	1.02834%	0.95343%	0.20619%	0.14814%	0.80529%	0.05805%	0.01685%
\Device\Harddisk0\DR(	143360	0.00007%	0.00000%	0.00007%	0.00000%	0.00000%	0.00007%	0.00000%
\Device\Harddisk0\DR(	147456	0.01295%	0.00940%	0.00871%	0.00578%	0.00362%	0.00293%	0.00063%
\Device\Harddisk0\DR(	183808	0.00063%	0.00000%	0.00063%	0.00000%	0.00000%	0.00063%	0.00000%
\Device\Harddisk0\DR(	233472	0.00279%	0.00000%	0.00275%	0.00000%	0.00000%	0.00275%	0.00003%
\Device\Harddisk0\DR(	262144	0.32309%	0.20772%	0.15423%	0.09573%	0.11199%	0.05850%	0.05687%
\Device\Harddisk0\DR(	323584	0.00115%	0.00084%	0.00091%	0.00077%	0.00007%	0.00014%	0.00017%
\Device\Harddisk0\DR(	348160	0.00010%	0.00000%	0.00010%	0.00000%	0.00000%	0.00010%	0.00000%
\Device\Harddisk0\DR(	503808	0.00010%	0.00000%	0.00010%	0.00000%	0.00000%	0.00010%	0.00000%
\Device\Harddisk0\DR(	524288	0.01160%	0.00268%	0.00986%	0.00247%	0.00021%	0.00738%	0.00153%
\Device\Harddisk0\DR(	757760	0.00003%	0.00000%	0.00003%	0.00000%	0.00000%	0.00003%	0.00000%
\Device\Harddisk0\DR(	761344	0.00031%	0.00000%	0.00031%	0.00000%	0.00000%	0.00031%	0.00000%
\Device\Harddisk0\DR(	763904	0.00003%	0.00000%	0.00003%	0.00000%	0.00000%	0.00003%	0.00000%
\Device\Harddisk0\DR(	1048576	0.11962%	0.05338%	0.09378%	0.04238%	0.01100%	0.05140%	0.01483%
\Device\Harddisk0\DR(	2097152	0.00247%	0.00247%	0.00059%	0.00059%	0.00188%	0.00000%	0.00000%
		99.99920%	46.19507%	67.96197%	21.25113%	24.94394%	46.71084%	7.09329%

Application: Synthesizing a WIOCP file into a Client Composite Workload

WIOCP CCW test pattern:

13 streams. 66 streams reduced to 13 stream CCW 2.0.1 workload (stream probabilities of less than 1% are excluded)

Restricted LBA Zones are used & defined as:

- 50% of the IOs to the first 5% of the LBAs.
- 30% of the IOs to the next 15% of the LBAs, and
- 20% of the IOs to the remaining 80% of the LBAs

Write Saturation - CCW 2.0.1 workload is then applied to 4 SSDs for a 6 hour Write Saturation (WSAT) test after a Device PURGE

5 Predominant streams – 93.18% are 4/8/16K:

RND 4K 45:55 R/W	26.53%	SEQ 4K 78:22 R/W	6.49%
RND 8K 59:41 R/W	5.06%		
RND 16K 09:91 R/W	31.77%	SEQ 16K 77:23 R/W	23.33%
Total RND 4/8/16K	**63.36%**	**Total SEQ 4/16K**	**29.82%**

Access pattern used follows Client Composite Workload based on Windows 7 functioning with RND/SEQ mix, RW mix and block sizes of following probabilities:

--- 11.93% - Random, Reads, 4096 bytes (4K)
--- 3.01% - Random, Reads, 8192 bytes (8K)
--- 2.98% - Random, Reads, 16384 bytes (16K)
--- 3.15% - Random, Reads, 32768 bytes (32K)
--- 5.03% - Sequential, Reads, 4096 bytes (4K)
--- 18.01% - Sequential, Reads, 16384 bytes (16K)
--- 14.6% - Random, Writes, 4096 bytes (4K)
--- 2.05% - Random, Writes, 8192 bytes (8K)
--- 28.79% - Random, Writes, 16384 bytes (16K)
--- 2.61% - Random, Writes, 32768 bytes (32K)
--- 1.06% - Random, Writes, 65536 bytes (64K)
--- 1.46% - Sequential, Writes, 4096 bytes (4K)
--- 5.32% - Sequential, Writes, 16384 bytes (16K)

and uses restricted LBA zones:

1) 50% to first 5% (LBA group a)
2) 30% to next 15% (LBA group b)
3) 20% to remainder (LBA group c)

Application: Viewing 4 Corner Case workloads on a Single SSD

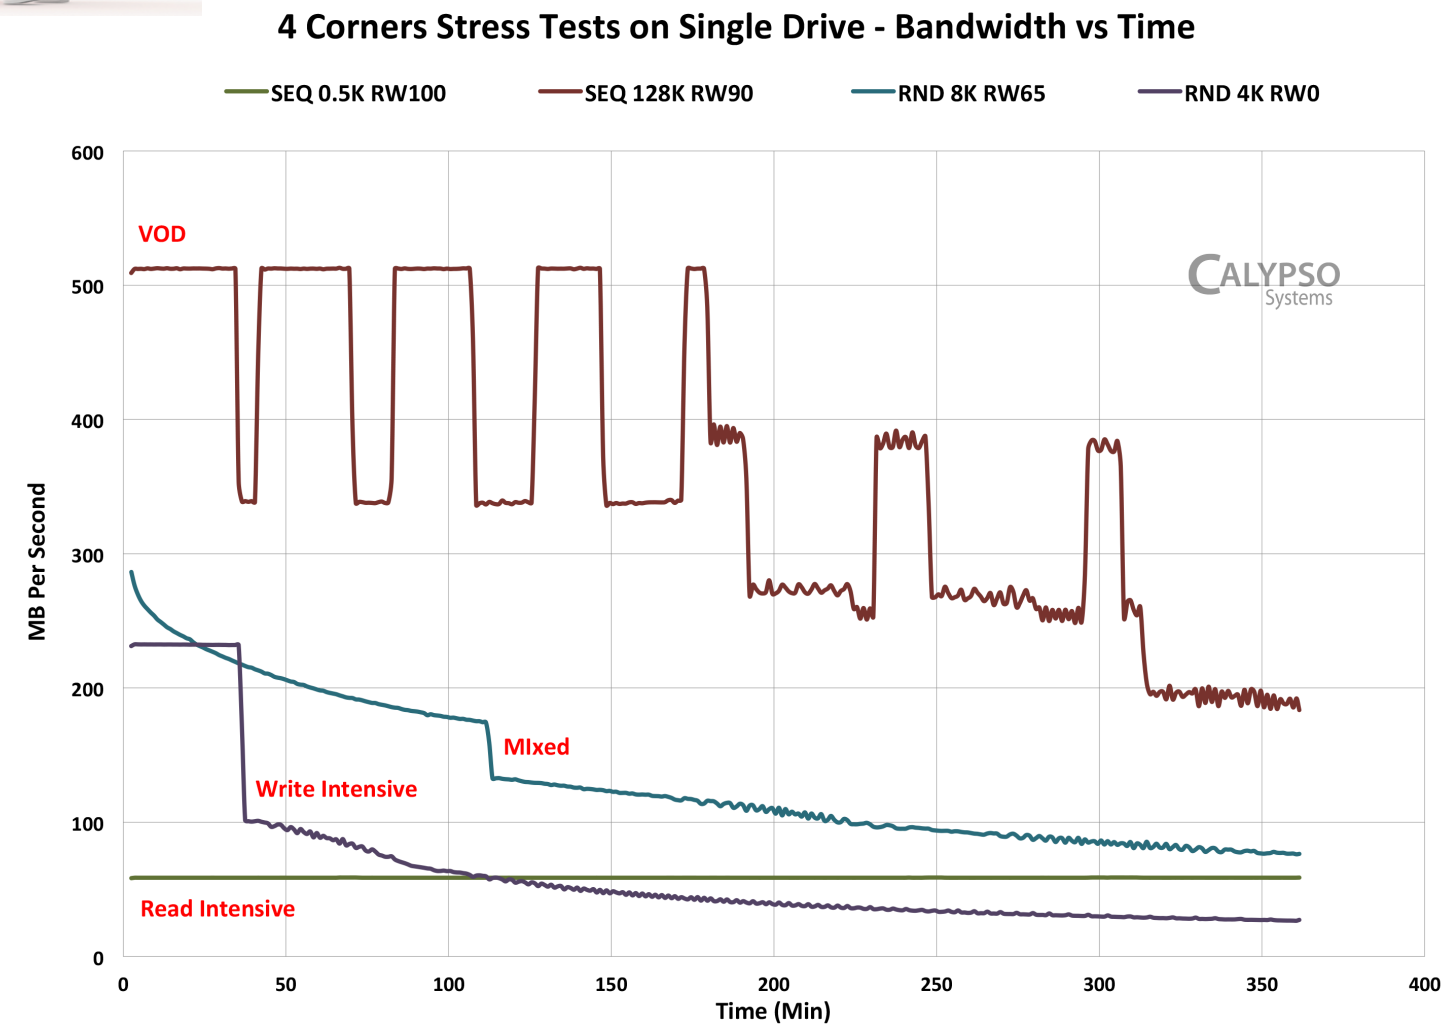

Application: Comparing 4 SSDs using a Client Composite Workload

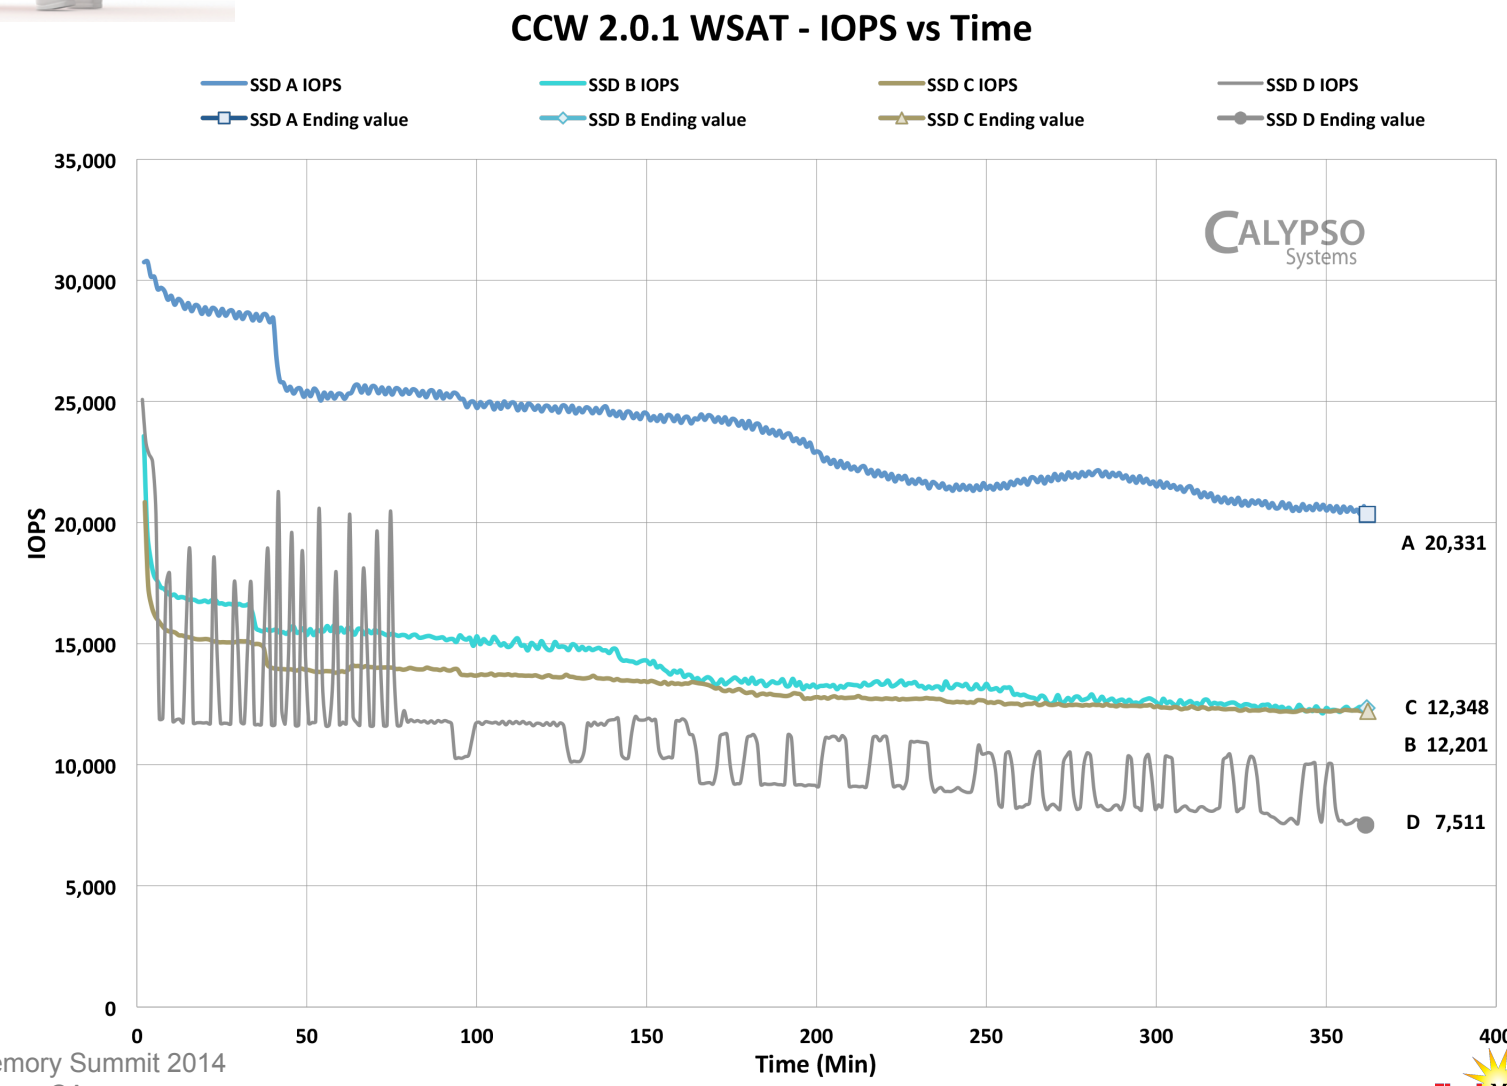

Application: Ranking Comparison – 4 Corner WSAT vs Client Composite 2.0.1

Ranking Comparison WSAT vs CCW 2.0.1 Workload									
RND 4K RW0		RND 8K RW65		SEQ 128K RW90		SEQ 0.5K RW100		CCW 2.0.1	
SSD Rank	IOPS	SSD Rank	IOPS	SSD Rank	IOPS	SSD Rank	IOPS	SSD Rank	IOPS
A	25,870	A	24,639	A	3,342	D	121,345	A	20,675
B	8,415	B	10,270	B	2,478	C	119,485	C	12,214
C	6,073	D	10,215	D	2,058	A	118,301	B	12,201
D	4,777	C	8,821	C	1,309	B	78,629	D	7,636

CCW 2.0.1 normalizes the difference between the drives and reduces the spread from Min to Max IOPS.

Note: All testing conducted on the Calypso RTP 3.0 / CTS 6.5 Reference Test Platform.
Client SSDs are commercially available and range from M.2 to SATA SSDs.

Next Steps:
More Captures! More Science!

Increase Sample Size

Gather more captures.

Seek volunteers to download the WIOCP capture tool and participate.

Conduct Specific Application Captures

Collect Captures from “general usage” sessions – e.g. user activity over a 1 week period

Collect Captures for “specific application” sessions – e.g. gaming, surfing, VOD, office work

Collect Captures for “different HW/SW” sessions – e.g. different PCs, different drives

Synthesize Captures into Test Workloads

Translate WIOCP IO metric capture files into Synthetic Client Workloads

Aggregate many captures from similar activities/hardware into “composite” workloads

Apply Tests to Client SSDs

Test a broad range of Client SSDs with the various Client Workloads

Compare relative rankings to rankings based on Corner Case test workloads

Investigate opportunity to define a “Universal Client Composite Workload”

Note: Latest WIOCP tool also can capture IO activity on enterprise Windows Server 2012 R2

Can we get there?
We need your help….

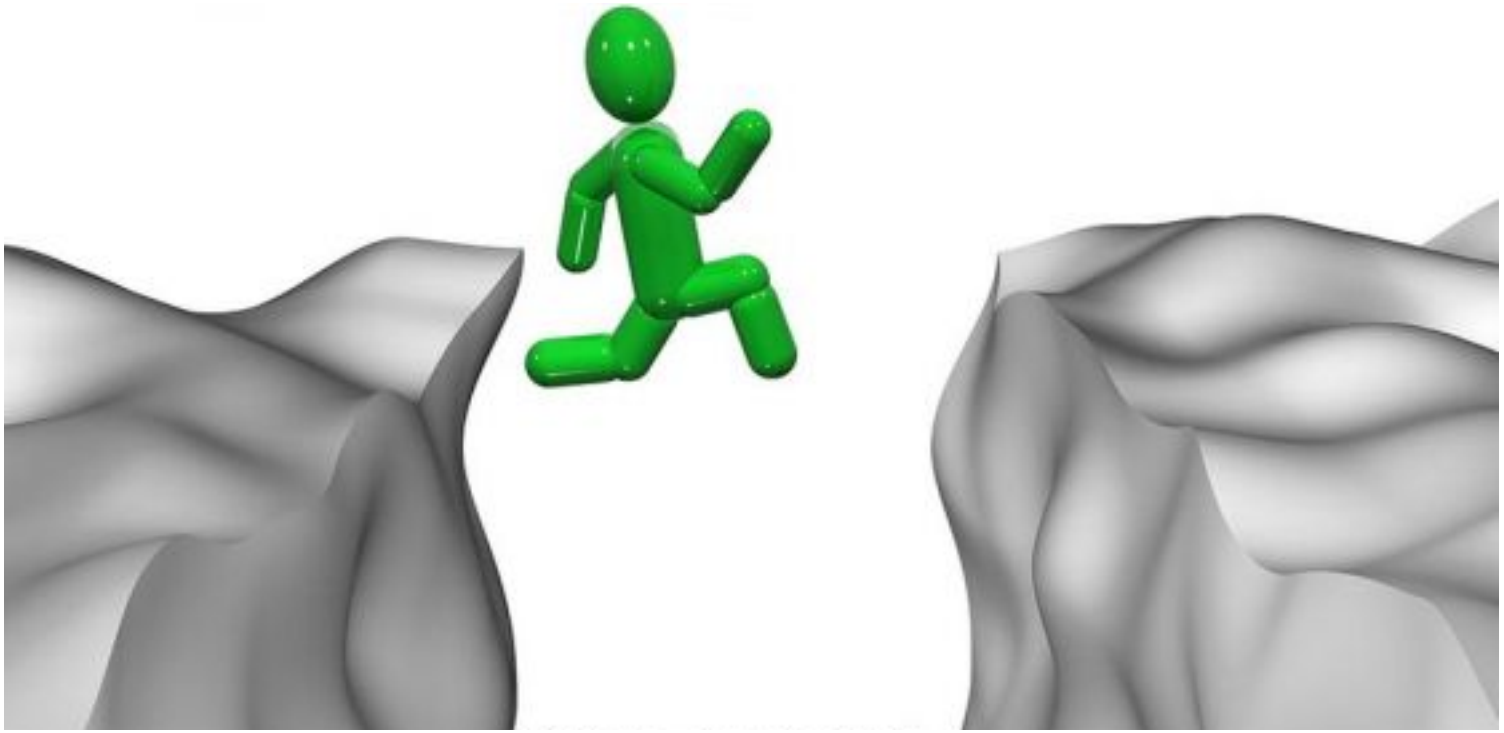

Opportunity: Join the WIOCP Program
Get a gift card and a chance to win a free SSD!!*

SNIA
Solid State Storage Initiative

Win an Amazon gift card & an SSD while helping develop next generation Solid State Drives!

Get with the **Workload I/O Capture** Program! Be a part of history in the industry and win some cool stuff while you're at it.

We are the Solid State Storage Initiative (SSSI) of the Storage Networking Industry Association (SNIA). SNIA is a well-respected organization supported by a number of industry leading companies such as HP, IBM, Intel, NetApp, and Microsoft.

The SSSI promotes SSD's in the marketplace, and in turn also develops performance measurement metrics for solid state storage devices. Our various programs contribute to improvements in PC and server performance, network efficiency and other benefits.

We are asking you to participate in a volunteer incentive effort to understand exactly what happens to disk drives when you use your computer for your everyday actions. It doesn't cost you anything. In fact, we will reward you for collecting data that will allow us to identify patterns of storage access by various applications. So, get on board and be a part of something exciting!

What exactly does the WIOCP capture?
The WIOCP (http://snia.org/forums/sssi/wiocp) captures I/O statistics unobtrusively and without compromising your PC's performance. No personal data or content is captured - only statistics on the *types of data transfers* that occur. This helps the SNIA SSSI & industry understand what actually takes place with your drive when you use your PC. Got more questions? Find the answers at http://tinyurl.com/wiocp-details So… what are you waiting for? Let's do this!

What do I do?
Download a simple program that runs in the background on your computer. Use your Windows PC (sorry, no Macs or Linux) as you normally would & forget that you ever installed anything. This program automatically collects statistics on your disk input/output activity. You are automatically notified when the 7-day collection period has ended & given steps on how to email us the statistics files & remove the program. Visit http://tinyurl.com/wiocp-program and follow the steps.

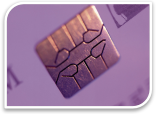

What do I get for participating?
This is an incentive program. So the more sets of statistics you return, the greater the potential incentive. If you participate and return one set of statistics, you qualify to win a $10 Amazon gift card and to be in a drawing for a free Intel SSD. The more results you submit, the better your chances of winning. Who doesn't want more stuff from Amazon or a screaming fast SSD?

Why does the SNIA SSSI do this?
Collecting I/O statistics helps computer scientists determine the type of workloads your drive is experiencing. By capturing statistics from a large number of computer users, designers can optimize both the drive and the host computer system to improve your overall computing experience. You can be a part of history!

Can I see the results of this WIOCP program?
Absolutely. The SNIA SSSI posts ***all*** results at the SNIA IOTTA by visiting http://iotta.snia.org/ You can see the statistics that are collected by the capture program before you submit them so that you can verify there is no confidential information. You can also see how the SNIA SSSI will use this data to compare client SSDs. By defining a standard set of I/O performance tests, SSDs can be ranked in order of performance. Ready, set, GO!

SNIA
Solid State Storage Initiative

intel

The SNIA SSSI thanks Intel for their generous donation of 2.5-inch SSD 520 Series 120GB SSDs!

2

Note: Raffle prize SSDs contributed by Intel Corp. to the SNIA SSSI WIOCP Program.
Limit one per person. See rules and eligibility. SNIA SSSI Governing Board Members not eligible for prizes.

Join the Effort!

Be a part of the future, sign up today!

http://www.snia.org/sites/default/files/wiocpflyer.FINAL_.5.1.14.pdf

For more information, contact Calypso Systems, Inc.

info@calypsotesters.com www.calypsotesters.com