Data Center Real World Storage Workloads

Comparison of Data Center Class SSD & SAS HDD Storage Performance

rev. 0.2

Eden Kim
Chair, SNIA Solid State Storage Technical Work Group
Calypso Systems, Inc.

Table of Contents

List of Figures

SNIA

Preface

Real World Application Workloads are becoming increasingly important to Data Center, Storage and IT Professionals because server and storage performance depends, in large part, on the IO Stream content of application workloads.

Unlike corner case benchmark tests, which are comprised of a synthetic stimulus with a single, or very few, IO Streams, Real World Workloads are comprised of a constantly changing combination of many IO Streams and changing Demand Intensity (or number of Users or Queue Depth) that occur when you actually run your applications on real world servers. This sequence of IO Stream combinations and the order of changing queue depths greatly affects how Data Center Storage and applications perform.

SNIA's Solid State Storage Technical Working Group (SSS TWG) has recently published a new Real World Storage Workload (RWSW) Performance Test Specification (PTS) v 1.0.7 for Data Center Storage. This RWSW PTS sets forth an industry standard methodology for the capture, analysis and test of Real World Application Workloads. The RWSW PTS sets forth four standardized tests to measure storage performance using workloads derived from real world captures. The RWSW PTS can be downloaded at https://www.snia.org/rwsw#.

RWSW tests can tell you how your storage will perform under real world use cases. When used in conjunction with industry standard benchmarks, such as those set forth by the SSS PTS v2.0.1, users can gain a comprehensive view of storage performance and understand how much performance, endurance or capacity is actually needed for their specific intended use case. The SSS PTS v1.0.7 can be downloaded at https://www.snia.org/sites/default/files/technical_work/PTS/SSS_PTS_2.0.1.pdf.

What do real world workloads look like and how can you capture them? Free tools for cross OS platform IO Capture and workload visualization are provided at www.TestMyWorkload.com. Here, users can download free IO Capture tools to capture their RWSW on any Mac, Windows or Linux system, view the captured real world workload and analyze various metrics, processes and events that occur during the real world capture.

The TestMyWorkload (TMW) site also lists reference IO Captures for use by the SSS TWG and public. Three example SNIA SSSI workloads are listed: a 24 hour Retail Web Portal, a 24 hour GPS Tracking Portal, and the SNIA Green Storage TWG reference workload. The content of each workload, as processed by the user, can be exported for free for use with any third party software.

This paper uses the SNIA SSSI 24 hour Retail Web Portal SQL server workload to show workload analysis, creation of a RWSW test stimulus, and the application of the RWSW to five sample drives: (1) SAS HDD and (4) Data Center class SSDs.

Results are provided for each drive and for comparative ordinal rankings. Readers can also view the sample 24 hour Retail Web Portal SQL server workload on the TMW site at http://testmyworkload.com/info/demo/#exampleKB24hr. Questions about the RWSW PTS v1.0.7 or PTS 2.0.1 can be sent to asksssi@snia.org.

Abstract

New IO Capture and analysis tools - such as those available for free at www.TestMyWorkload.com (TMW) - now enable users to capture, analyze and test ***Datacenter Real World Storage Workloads*** (RWSWs). RWSWs are very different from synthetic lab workloads. RWSWs are comprised of constantly changing combinations of IO Streams and Demand Intensity whereas lab workloads apply a fixed and constant workload. Datacenter storage performance depends on how well storage responds to these dynamic RWSWs and often does not align with published manufacturer specifications or single access pattern corner case tests.

IO Capture of RWSWs allows users to understand what IO Streams occur at the Block IO level and to analyze in-situ server storage performance that occurs during the IO Capture process. RWSWs also allow the user to validate software stack optimizations and to qualify server storage to the actual IO Streams that occur.

This paper uses the TMW ***IOPofiler*** toolset to capture and analyze a Block IO level capture on a 2,000-outlet retail chain web portal. The 24-hour IO Capture file is processed to extract SQL Server application IOs which are then used to present in-situ SQL Server performance. The extracted SQL Server workload is also used as the test stimulus in RWSW tests that compare storage performance of a sample pool of Datacenter class storage.

Ordinal rankings and performance to RWSW tests are listed for the sample pool drives. These RWSW test results vary significantly from manufacturer specification and single access pattern synthetic corner case tests due to the different IO Stream content and changing Demand Intensity of RWSWs.

In this study, five Datacenter class drives are tested and ranked to the sqlservr.exe workload: (4) 960GB SATA SSDs by Samsung, Seagate, Micron and Sandisk and (1) 143 GB Seagate SAS HDD.

Sample Pool

The original IO Capture was taken on a retail web portal running Windows Server 2008 R2 with two virtual physical drives. The server is a HDD JBOD. Test samples are five Datacenter class drives: four 960 GB SATA SSDs and a 143 GB SAS HDD. Details are below.

Manufacturer	Model	Capacity	Protocol	Type
Samsung	MZ7LM960HCHP-0E003	960 GB	SATA	SSD
Seagate	XF1230-1A0960	960 GB	SATA	SSD
Micron	Micron_5100_MTFDDAK960TBY	960 GB	SATA	SSD
Sandisk	SDLF1DAR-960G-1HA1	960 GB	SATA	SSD
Seagate	ST9146853SS	143 GB	SAS	HDD

Figure 1 – Test Drive Sample Pool

I. Introduction

Datacenter Designers, IT Professionals, Storage Server ODMs and SSD OEM/ODMs are increasingly interested in Real World Storage Workloads (RWSWs) because these workloads determine, in large part, the level of performance that will be provided by Datacenter storage. RWSWs, and their IO Streams, change as they traverse the Software (SW) Stack from User (application) space to storage and back.

This paper presents the Block IO level capture, analysis and test of a Real World Datacenter workload for a 2,000-outlet retail chain web portal SQL Server application running on a SAS HDD JBOD server. While IO Captures are taken at both the File System and Block IO levels, the data and tests herein focus on the extracted sqlservr.exe IO Streams observed at the Block IO level on Drive0. The evaluation of in-situ performance is achieved by using the IOProfiler Self-Test' feature. The Self-Test filters, or extracts, the desired application process IO Streams from the overall, or cumulative, workload IO Streams.

Here, sqlservr.exe IO Streams represent 79.9% of the total IOs that occur. The sqlservr.exe workload is next filtered to present those IO Streams that occur at least 3% or more of the time over the duration of the capture. This results in 6 dominant IO Streams which account for 78% of the total sqlservr.exe IO Streams. The metrics from these 6 IO Streams are then used to present in-situ SQL Server performance.

In-Situ Self-Test performance is presented in a variety of plots including:

- **IOPS and Throughput (TP) v Time**
- **IOPS and TP v Segment**
- **Response Times v Time**
- **Demand Intensity (Users or Queue Depth) v Time**
- **IO Streams by Quantity of IOs**
- **Compressibility Ratio & Duplication Ratio**

This same group of 6 major IO Streams (and associated Demand Intensity) is also used as the test stimuli in the RWSW tests. The following four RWSW tests are applied to the sample pool drives:

- **Multi-WSAT** - a fixed Composite 6 IO Stream workload applied for each test step
- **Individual Streams WSAT** - 6 independent and separate IO Stream tests
- **Replay-Native** - re-creating the sequence and combination of IO Streams and Demand Intensity that occurred during the capture for each test step
- **DIRTH** – or Demand Intensity Response Time Histogram - which applies the fixed Composite 6 IO Stream workload across a span of 1 to 1,024 Users to observe the range of performance and response time saturation.

The test sample pool consists of four 2.5" Datacenter class 960 GB SATA SSDs from Samsung, Seagate, Micron and Sandisk and a 2.5" 15K RPM SAS HDD from Seagate.

II. Background

The performance level of solid state storage depends, in large part, on the workload IO Streams that are applied to storage at the Block IO level. Storage and Datacenter professionals, therefore, need to know the IO Stream composition and Demand Intensity at specific levels in the Software (SW) stack to be able to validate SW abstractions (such as de-duplication, data reduction, back-up, data protection, snapshot, replication and storage tiering), to optimize storage firmware and to qualify storage systems and devices for use in Datacenters.

A. What are Real World Storage Workloads (RWSW)

RWSWs are constantly changing combinations of IO Streams generated by applications in User space. IO Streams are modified (appended, fragmented or coalesced) by Operating System (OS) and software abstractions as the IO Streams travel from User space to storage and back. See Figure 2 below.

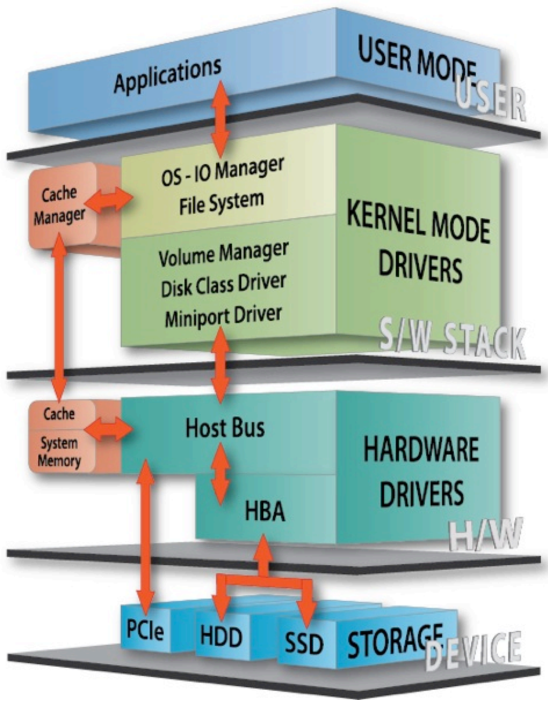

Figure 2 – Windows Software Stack

A specific IO Stream[1] is a unique Random or Sequential (RND/SEQ) access of a specific data transfer size (aka Block Size) of a Read or Write (Read/Write) IO with a given Demand Intensity (aka Queue Depth or Users) observed at a specific level in the SW Stack.

[1] An IO Stream, used in the context of RWSWs, should not be confused with 'data streams' used in relation to SSD Endurance where similar write operations are associated with a given 'data stream.'

B. What is an IO Capture?

An IO Capture is the statistical tabulation of IO Streams that occur at a given SW Stack level over time. No personal or other data is captured, only a table of binary IO metrics.

A given IO Stream, such as a SEQ 1.5K Write, can occur many times during a capture but is still tabulated as a single IO Stream. In Figure 3 below, the IO Statistics table for a single IO Capture step shows the IO Streams and metrics recorded for the capture step. Note that the SEQ 1.5K W IO Stream occurs 69 times during the step but is recorded as a single IO Stream.

Access Pattern	RND or SEQ	Block Size	Read/Write	% Occurrence	Quantity (IOs)
SEQ 0.5K W	SEQ	512	W	9.24	477
SEQ 1K W	SEQ	1024	W	4.32	223
SEQ 1.5K W	SEQ	1536	W	1.34	69
SEQ 4K W	SEQ	4096	W	22.31	1152
SEQ 16K W	SEQ	16384	W	14.25	736
RND 0.5K W	RND	512	W	1.99	103
RND 1K W	RND	1024	W	3.21	166
RND 1.5K W	RND	1536	W	1.74	90
RND 2K W	RND	2048	W	0.93	48
RND 2.5K W	RND	2560	W	0.74	38
RND 3K W	RND	3072	W	0.58	30
RND 3.5K W	RND	3584	W	0.62	32
RND 4K W	RND	4096	W	9.8	506
RND 4K R	RND	4096	R	0.91	47
RND 8K R	RND	8192	R	0.15	8
RND 8K W	RND	8192	W	2.73	141
RND 12K W	RND	12288	W	1.24	64
RND 16K W	RND	16384	W	15.63	807
RND 20K W	RND	20480	W	0.58	30
RND 28K W	RND	28672	W	2.03	105
RND 36K W	RND	36864	W	0.19	10

Figure 3 – IO Stream Statistics Table

IO Stream activity can be captured by various public and private IO Capture tools. Public tools include Perfmon for Windows (for captures at the File System level) and blktrace for Linux (for captures at the Block IO level). IOProfiler by Calypso provides cross platform IO captures for Linux, Windows, MacOS and FreeBSD at both the File System and Block IO levels.[2]

[2] IO Capture tools differ in the OS supported, IO metrics captured and data analytic features that are provided. IO Captures and metrics herein are taken using IOProfiler (IPF) software and IPF IO Capture Applets.

C. Why do I care about RWSWs?

RWSWs are different from synthetic lab corner case tests and present storage with a unique type of workload stimulus. Synthetic corner case and manufacturer published specifications are comprised of a single, or very few, discrete IO Streams (or access patterns) which vary significantly from the User's dynamically changing combination of RWSW IO Streams.

In addition, RWSWs change as they traverse the SW stack. Because storage performance depends on the type of IO Streams presented to storage, it is important to understand the content of your RWSW at specific SW Stack levels (such as at the Storage Block IO level).

IO metrics such as IOPS, Throughput, Latency, Queue Depth, De-duplication ratio, Compression Ratio, Disk Utilization, IO Bursts, IO Sequentiality and LBA Range Hit location can help the datacenter and storage professional assess SW stack and storage performance and help trouble shoot, design and qualify storage solutions.

D. What does a RWSW look like?

RWSWs can be visualized as an IO Stream Map that shows the changing combinations of IO Streams and associated metrics for each capture step[3].

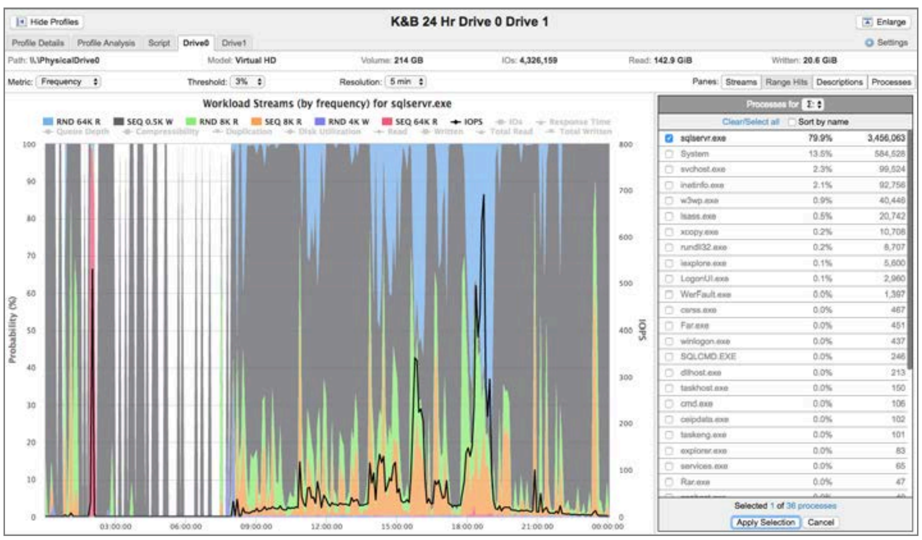

Figure 4 – IO Stream Map

[3] LBA Range Hit Maps and spatial locality of reference is outside the scope of this paper.

Figure 4 above shows the IO Stream Map used in this study. The IO Stream Map is filtered to show sqlservr.exe IOs that occur 3% or more of the time over the course of the 24-hour capture. Each time point shows the IO Streams and metrics that occur during the 5 minute step. The 6 dominant IO Streams are shown as colored data series, IOPS are the dominant black line with the sqlservr.exe IO filter selected in Processes box on the right.

E. How do I describe a RWSW?

RWSWs can be described as a summary of the application processes and IO Streams selected during a given capture time, event or process. Because RWSW are a collection of different combinations of IO Streams, RWSWs should be identified by the specific attributes of interest for the IO Capture or RWSW test workload.

For example, in this study the 3%, or higher, occurring sqlservr.exe IO Streams are extracted and presented in a list of 6 dominant IO Streams that are 79.9% of the total IOs.

This workload is generically referred to as the **'Retail Web Portal 24-hour SQL Server 6 IO Stream'** workload. Specific tests refer to different applications of the 6 IO Streams as in the:

- **Multi-WSAT:** SQL Server Composite 6 IO Stream workload
- **Individual Stream WSAT:** SQL Server 6 IO Stream workload
- **Replay-Native:** SQL Server 6 IO Stream workload
- **DIRTH:** SQL Server Composite 6 IO Stream workload

III. Results

A. SQL Server workloads

The 24-hour SQL Server RWSW represents a variety of applications, events and periods of varying storage utilization, each of which is comprised of different combinations of IO Streams. See Figure 4 IO Stream Map above.

For this study, two workloads are defined: a Composite 6 IO Stream workload with percentage occurrence for each IO Stream over the capture duration; and a Replay-Native workload where each combination of IO Streams is re-created for each step of the capture.

Figure 5 below lists the Composite 6 IO Streams with each respective percent of occurrence. This workload is used for the WSAT and DIRTH Tests while the Replay workload is used in the Replay-Native test and is the basis of the Self-Test in-situ plots.

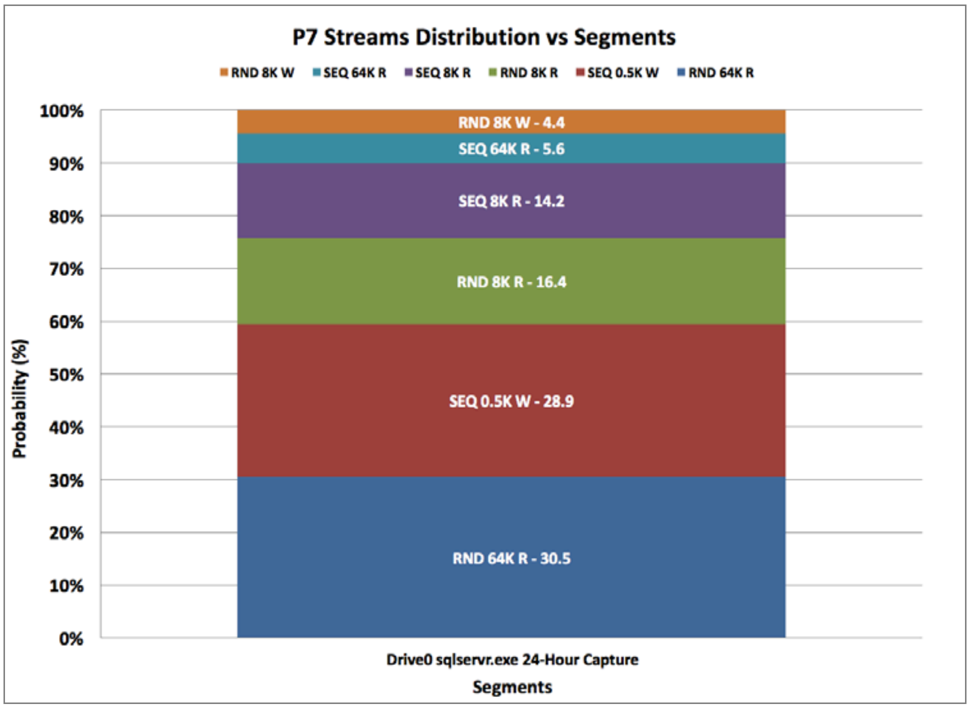

Figure 5 – Composite 6 IO Stream Workload

B. In-situ sqlservr.exe performance

In-situ performance can be analyzed by observing the IO combinations, QDs and IO metrics associated with the IO Streams. The ability to extract application specific IOs (such as the sqlservr.exe IOs in this study), for entire processes or sub-events, provides valuable data to IT professionals that allows them to assess and validate software stack optimizations.

The IO Stream Map in Figure 4 above presents the in-situ sqlservr.exe IOs as extracted from the overall 24-hour IO Capture workload. Figure 6 below shows how the IO Stream combinations, IO Stream percentages and QDs vary for different times: 2 am data back-up (SEQ 64K R/W), early morning hours (periods of low disk utilization), working hours (mixed RND/SEQ block sizes, high IOPS and high QDs) and for the total 24-hour sqlservr.exe IO workload.

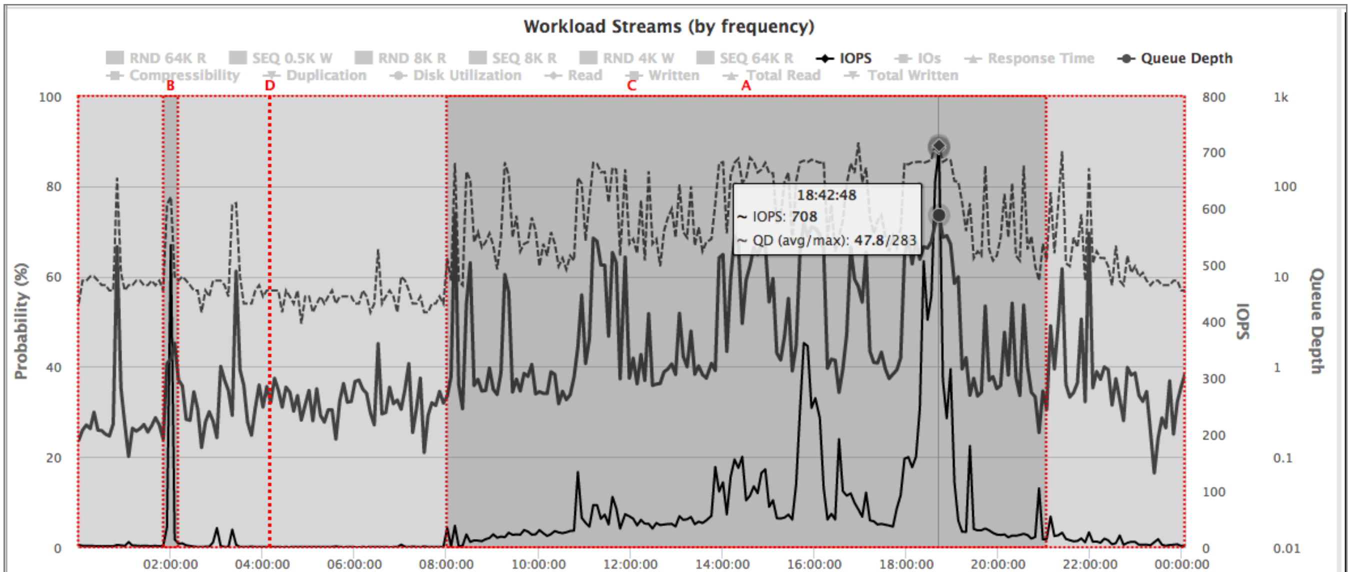

Σ Cumulative Workload		
RND 64K R	19.5%	841,495
SEQ 0.5K W	17.9%	775,072
RND 8K R	10.5%	454,786
SEQ 8K R	8.8%	381,168
RND 4K W	4.1%	179,391
SEQ 64K R	3.5%	152,723
RND 8K W	2.91%	125,959
RND 4K R	2.74%	118,560
Total IOs of 5,038 streams: 4,326,159		
Selected 6 streams: 2,784,635 (64%) E		

A Working Hours		
RND 64K R	21.0%	832,350
SEQ 0.5K W	18.5%	736,026
RND 8K R	11.1%	439,243
SEQ 8K R	9.4%	372,000
RND 4K W	3.8%	152,511
RND 8K W	3.0%	120,008
RND 4K R	2.57%	101,887
RND 16K R	2.01%	79,752
Total IOs of 4,979 streams: 3,970,471		
Selected 6 streams: 2,652,138 (67%)		

B 2 am Data Back-up		
SEQ 64K R	81.1%	144,396
RND 4K R	1.77%	3,150
RND 4K W	1.70%	3,034
RND 32K R	1.10%	1,956
RND 24K R	1.09%	1,940
SEQ 60K R	1.06%	1,883
SEQ 40K R	1.00%	1,772
SEQ 32K R	0.81%	1,436
Total IOs of 181 streams: 178,057		
Selected 1 stream: 144,396 (81%)		

C sqlservr.exe 24-Hr		
RND 64K R	23.7%	818,951
SEQ 0.5K W	22.4%	774,341
RND 8K R	12.7%	438,718
SEQ 8K R	11.0%	380,565
SEQ 64K R	4.4%	151,226
RND 8K W	3.4%	118,791
SEQ 512K R	1.73%	59,917
RND 24K R	1.60%	55,312
Total IOs of 395 streams: 3,456,063		
Selected 6 streams: 2,682,592 (78%) E		

D Early Morning		
RND 4K W	58.8%	57
SEQ 8K W	10.3%	10
SEQ 4K W	7.2%	7
RND 0.5K W	6.2%	6
SEQ 0.5K W	4.1%	4
RND 1K W	3.1%	3
RND 64K W	3.1%	3
RND 2K W	1.03%	1
Total IOs of 14 streams: 97		
Selected 7 streams: 90 (93%)		

Figure 6 – IO Stream Map: Varying IO Streams for different events

In-situ (Target Server Self-Test) performance is, as expected, lower than the sample pool RWSW. This is because the IO rate and QDs that occur on the target JBOD server are ‘throttled’ by low server application IOs, lower QDs, periods of low storage utilization and the fact that the target server is a HDD based JBOD.

By comparison, RWSW tests apply maximum IO accesses and/or apply a higher QD setting to achieve test drive saturation and thus show higher drive performance than in-situ captures. See Section F - Replay-Native that follows for a more detailed explanation.

C. Target Server Self-Test

The Target Server Self-Test results show a dynamic range of IOPS, TP and QD with different IO Stream composition at different times or events.

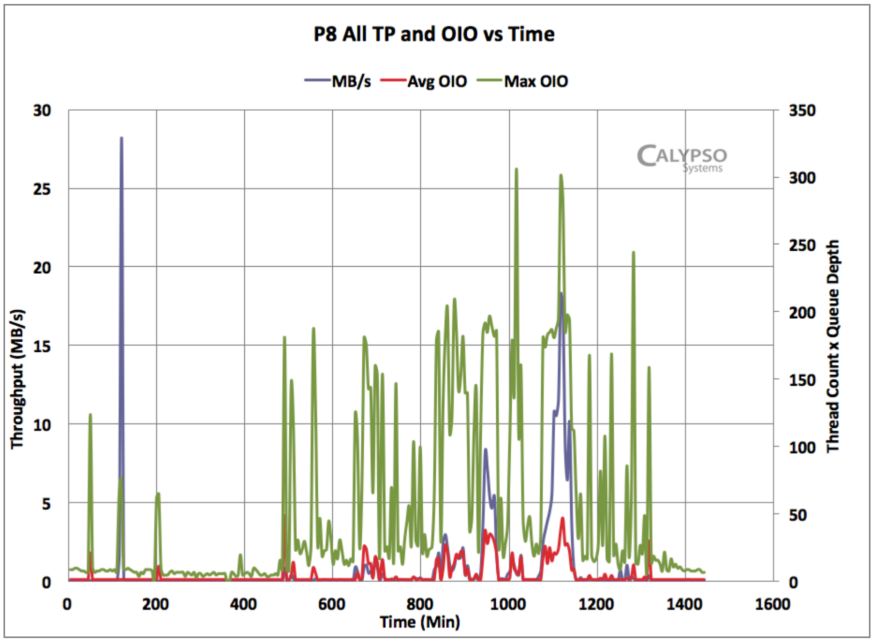

Figure 7 – Self-Test: Throughput & OIO: Average & Maximum Queue Depths

Figure 8 below shows the changing sqlservr.exe IO Stream combinations, IO count and IOPS over the 24-hour capture period.

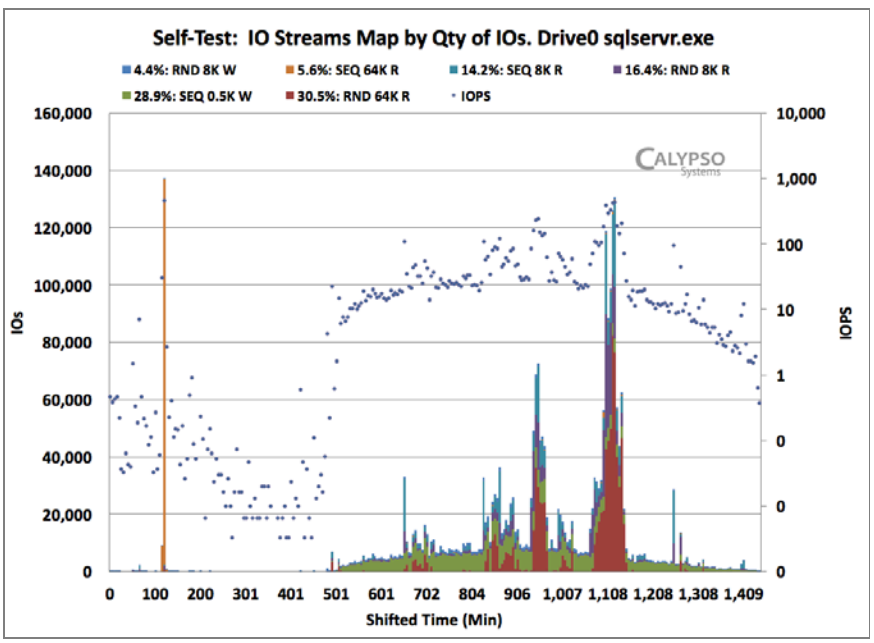

Figure 8 – Self-Test: IO Streams Map by Quantity of IOs – 24-Hour Drive0 sqlservr.exe

Figure 9 shows Average and Maximum Response Times. Response times are generally lower during periods of low IOs and higher during periods of high IOs.

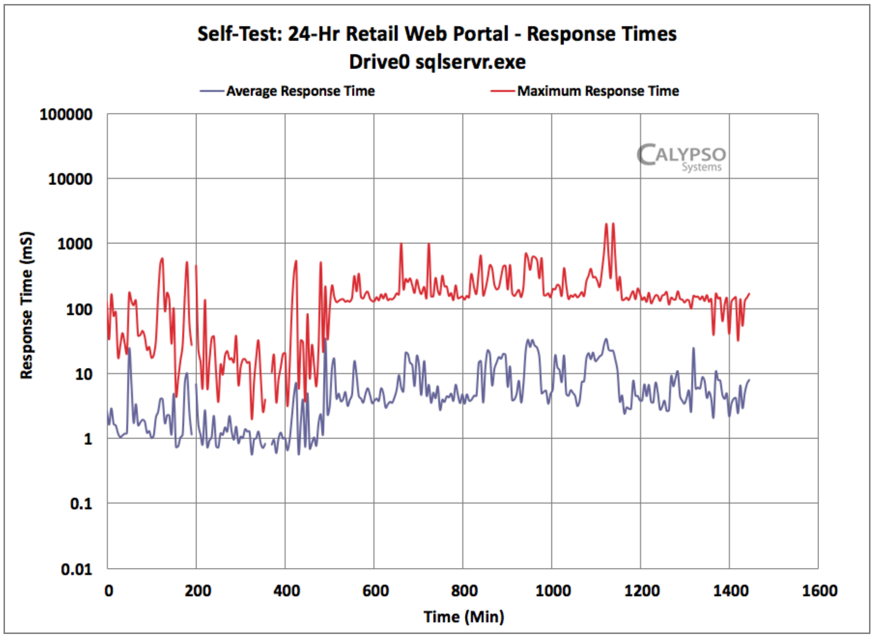

Figure 9 – Self-Test: Response Times

Compressibility Ratio analyzes data and calculates how much MORE the data can be compressed (4 times or 400%) while Duplication Ratio calculates what percent of blocks can be de-duplicated (28%).

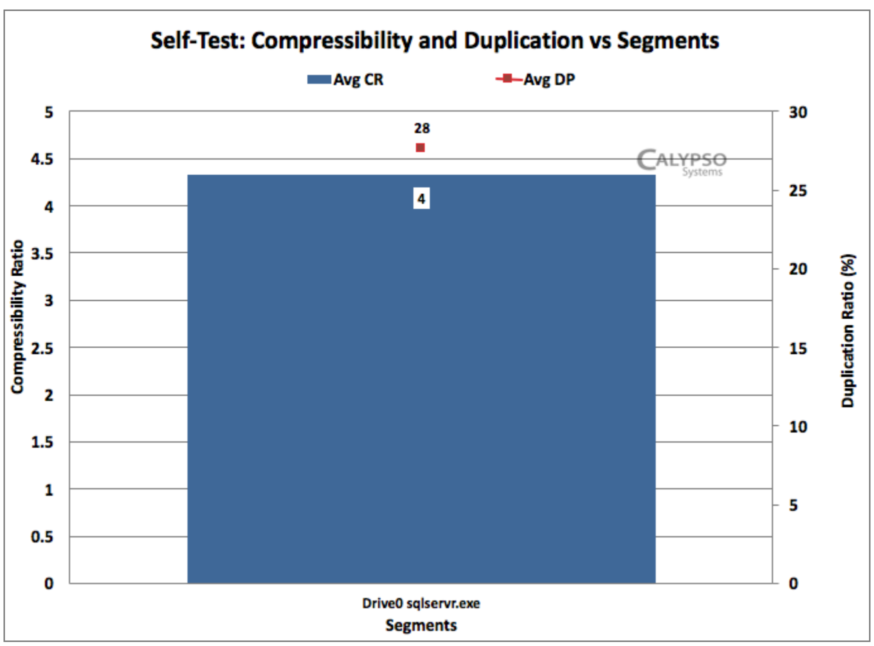

Figure 10 – Self-Test: Compressibility & Duplication Ratios

D. Multi-WSAT

For the Multi-WSAT Composite workload, Samsung SSD shows the highest performance followed by Seagate SSD, Micron SSD, Sandisk SSD and the Seagate SAS HDD. Higher IOPS are better.

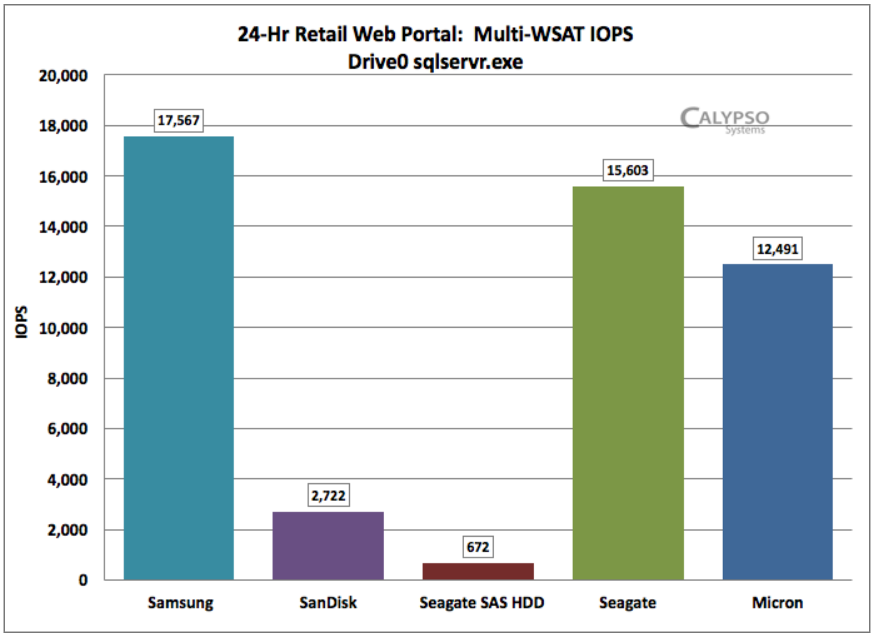

Figure 11 – Multi-WSAT: IOPS

Figure 12 below shows Multi-WSAT Steady State Throughput (TP). Higher TP is better.

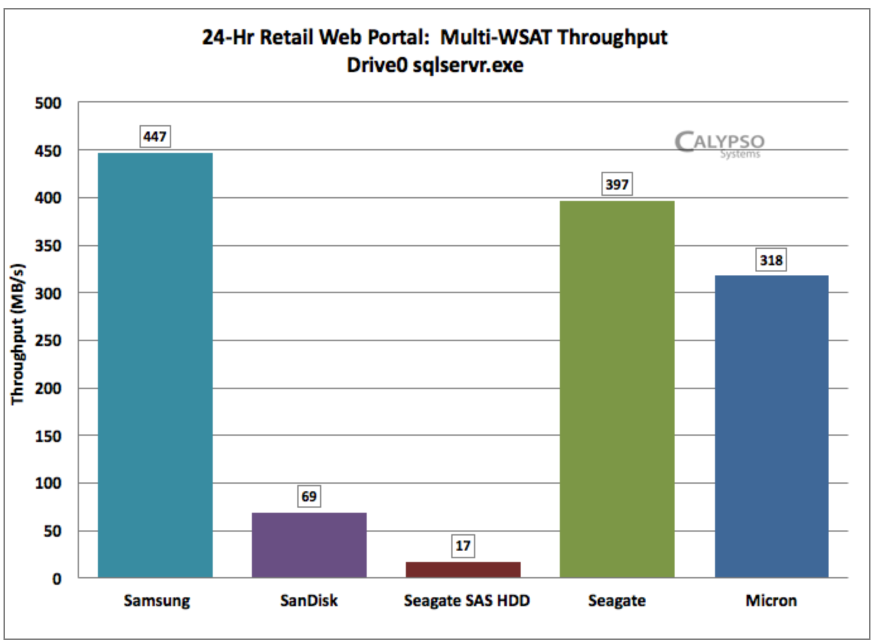

Figure 12 – Multi-WSAT: Throughput

Figure 13 below shows Multi-WSAT Average Response Times (ART). Lower ART is better.

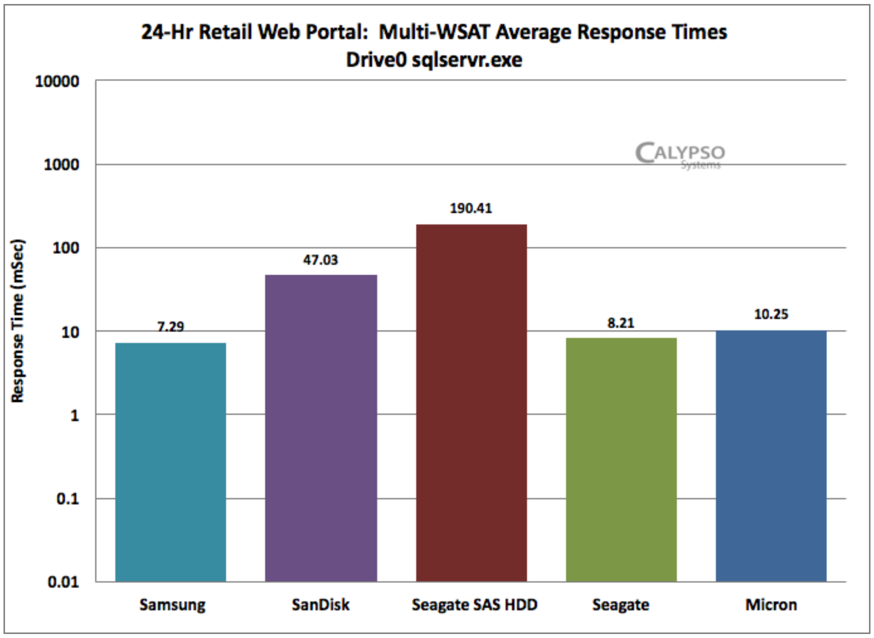

Figure 13 – Multi-WSAT: Average Response Times

Figure 14 below shows the 5 9s Response Time (RT) Quality of Service (QoS). 5 9s Response Time (or 99.999%) is the preferred Response Time figure of merit. Lower RT QoS is better.

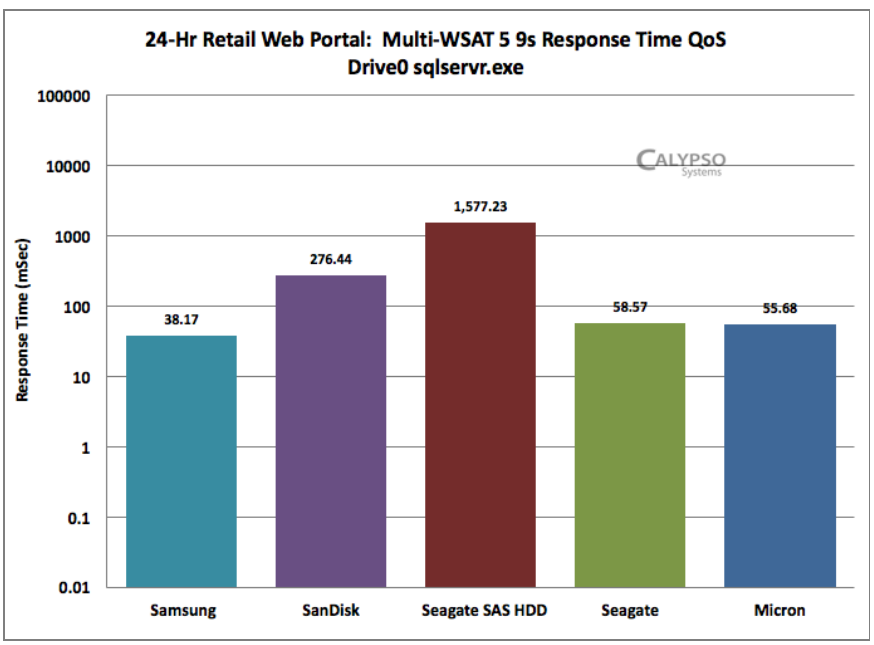

Figure 14 – Multi-WSAT: 5 9s Response Time Quality of Service

E. Individual Streams WSAT

Individual Streams WSAT tests each of the 6 dominant IO Streams as a separate test run to Steady State. Among the 6 IO Streams, the SEQ 0.5K W IO Stream component has the largest impact on performance (and impact on Multi-WSAT results above). Samsung shows 32,721 SEQ 0.5K W IOPS which is 2X Seagate, 10X Sandisk, 15X Micron and 30X SAS HDD.

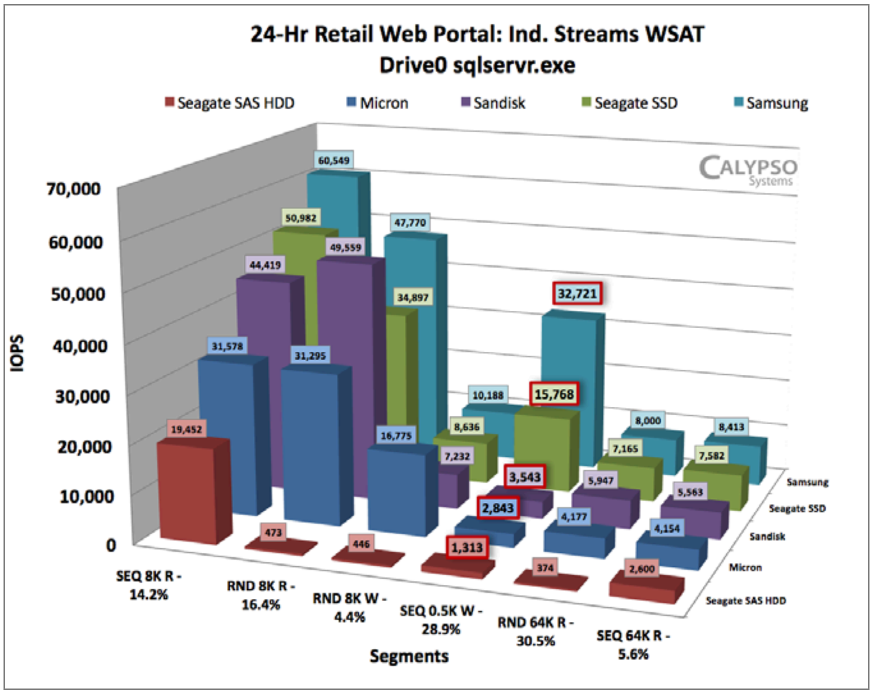

Figure 15 – Individual Streams WSAT: 6 IO Streams

F. Replay-Native v Self-Test

The Replay-Native and the in-situ Target Server Self-Test share the same series of IO Stream Combination and Queue Depth steps. However, Self-Test performance reflects in-situ throttling from periods of low application IOs and disk utilization. The Replay-Native test does not limit performance to the number of IOs observed during each capture step.

For example, a capture step may show 30,000 IOs of SEQ 0.5K W applied at a QD of 128. Self-Test IOPS are reported as 30,000 IOs over the course of one 5-minute step, or 100 IOs per Second (IOPS). In a Replay-Native test, sample pool drives show higher IOPS over a 5-minute step because the SEQ 0.5K W IOs are not throttled (limited) to 30,000.

In general, Datacenter storage workloads do not saturate Datacenter storage as application IOs are not applied in a limitless fashion, even at peak IO periods. Also, the previously mentioned periods of low application IOs and disk utilization results in lower in-situ performance.

Datacenter storage performance (IOPS, Sequentiality and Response Times) during peak IO Bursts is of great interest. However, IO Burst and IO Sequentiality analysis requires fine grain temporal and spatial resolution that is outside the scope of this study.[4]

G. Replay-Native Comparison Data

Replay-Native reproduces IO Capture IO Stream combinations for each test step. IOPS v Time Difference from Group Mean show Samsung (blue dotted line) and Seagate (green line) with the highest performance, Micron (blue dotted line) with low performance while Seagate HDD (red line) and Self-Test (red dotted line) show the lowest performance.

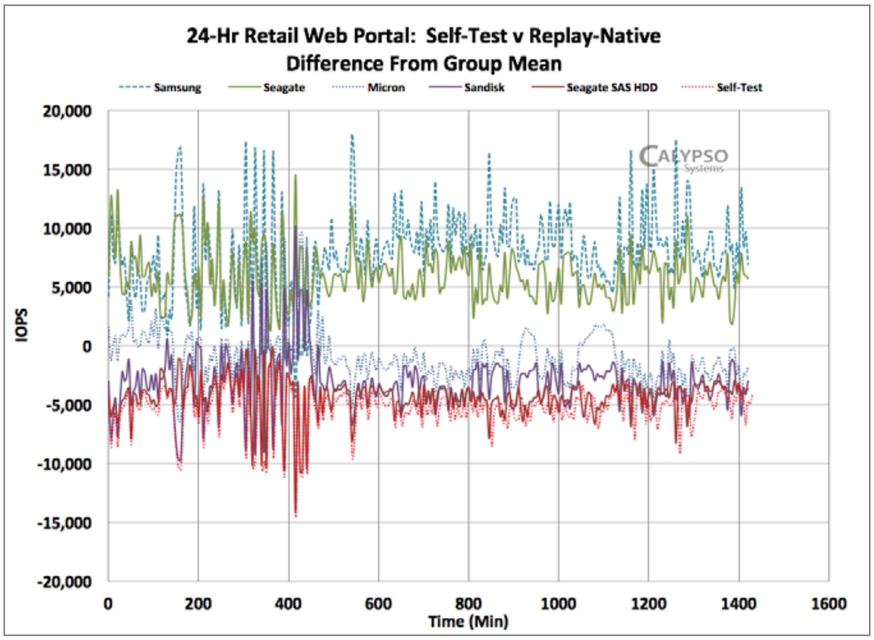

Figure 16 – Replay Test: IOPS v Time - Difference from Group Mean

Unlike Replay IOPS v Time, Figure 17 below shows IOPS averaged over the 24-Hour capture duration. Samsung and Seagate SSDs show the highest IOPS. Target Server Self-Test shows low IOPS due to limited application IOs occurring during the capture.

[4] While the IOProfiler IO Capture tools provide temporal resolution to 1 uSec and spatial resolution to 0.01%, this analysis involves different types of captures which are outside the scope of this study.

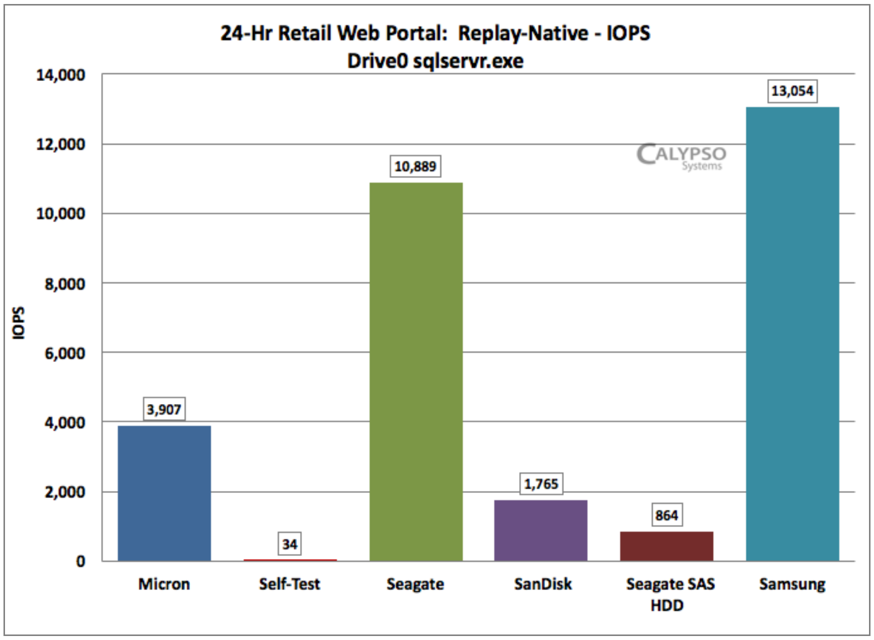

Figure 17 – Replay-Native: IOPS – 24 Hour Average

Figure 18 below shows Throughput (TP) averaged over the 24-Hour capture duration. Samsung and Seagate SSDs show the highest TP. Higher IOPS and TP is better.

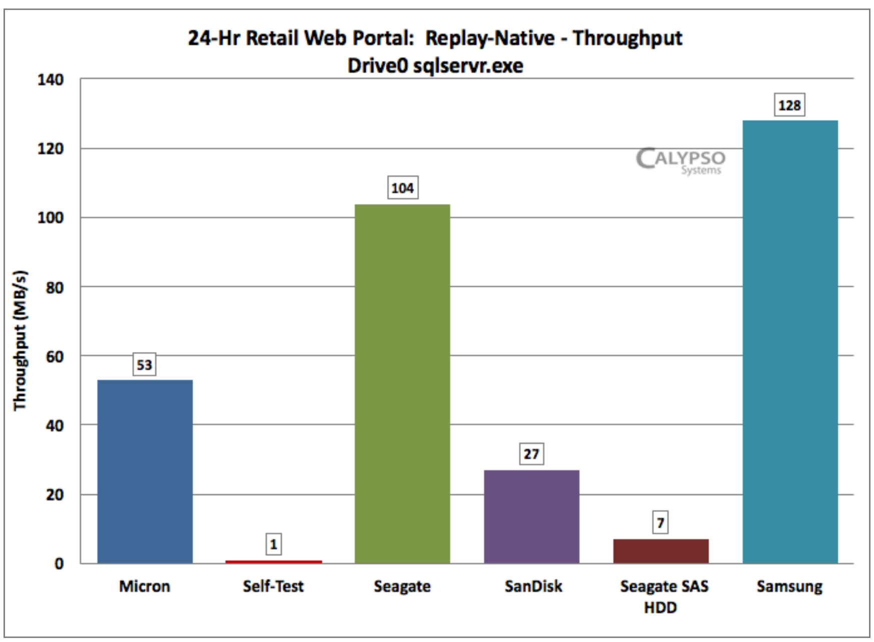

Figure 18 - Replay-Native: Throughput – 24 Hour Average

Figure 19 below shows Average Response Times (ART) over the 24-hour capture. Sample drives and the Self-Test show reasonable ARTs. Lower response times (RTs) are better.

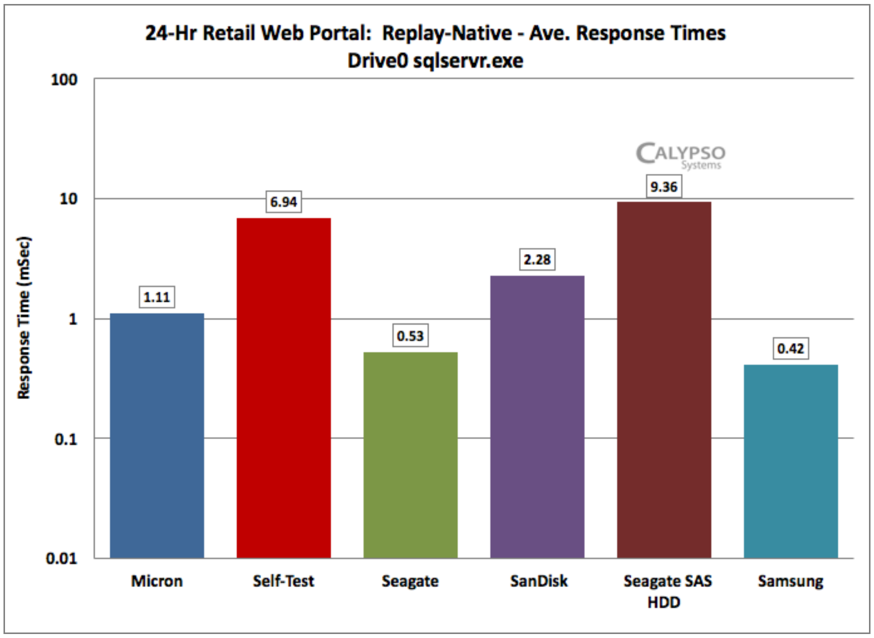

Figure 19 – Replay-Native: Average Response Times – 24-Hour Average

Figure 20 below shows Maximum Response Times (MRT) over the 24-Hour capture duration. MRT less than 100 mSec are reasonable. The Seagate SAS HDD shows extremely high MRT.

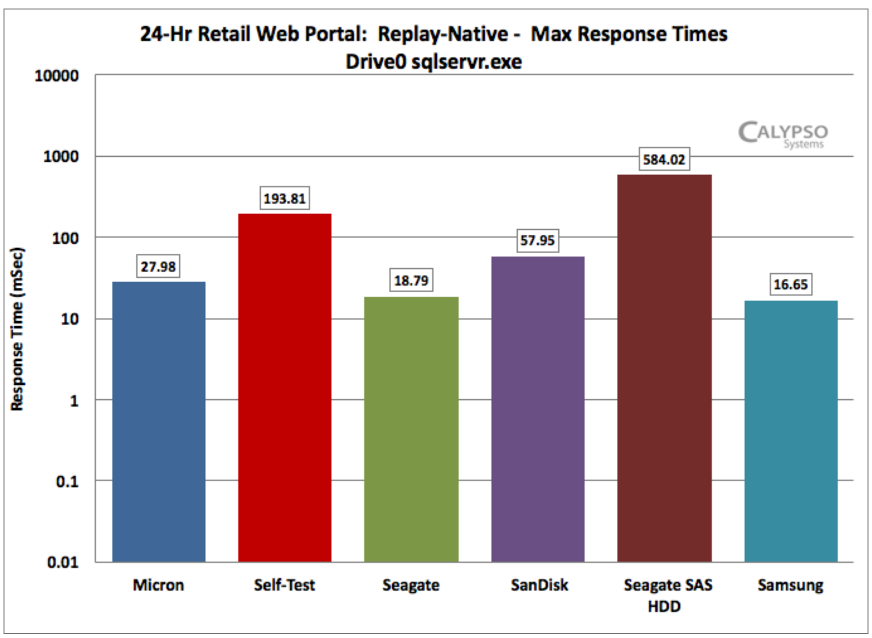

Figure 20 - Replay-Native: Maximum Response Times

Figure 21 shows 5 9s Response Time Quality of Service (QoS). No data appears for the Self-Test because this metric was not recorded during the capture. RT QoS of 20 mSec or less is desirable.

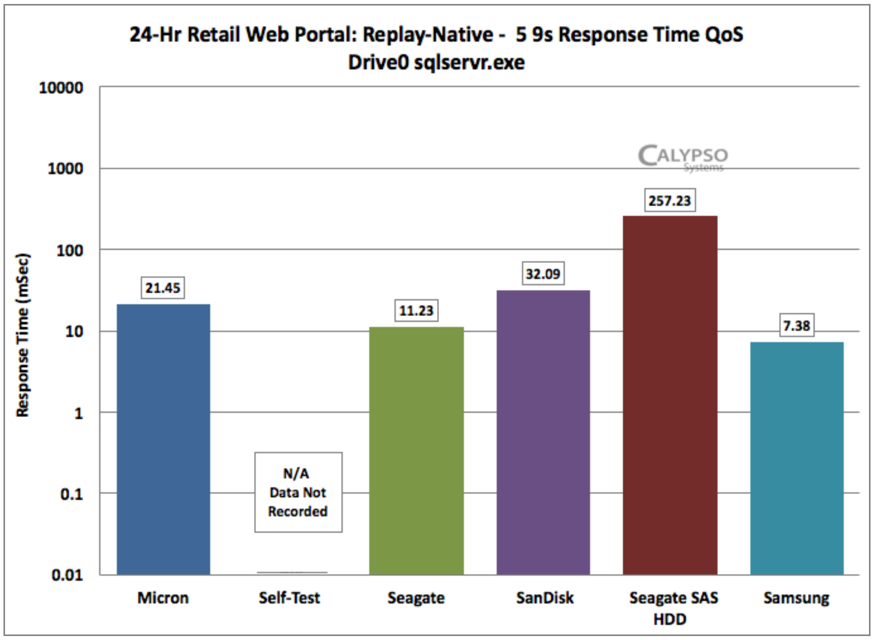

Figure 21 - Replay-Native: 5 9s Response Time QoS

Figure 22 shows Power Consumption. Lower power consumption is better. Seagate SAS HDD shows very high power consumption. No power data was collected during the Self-Test IO Capture.

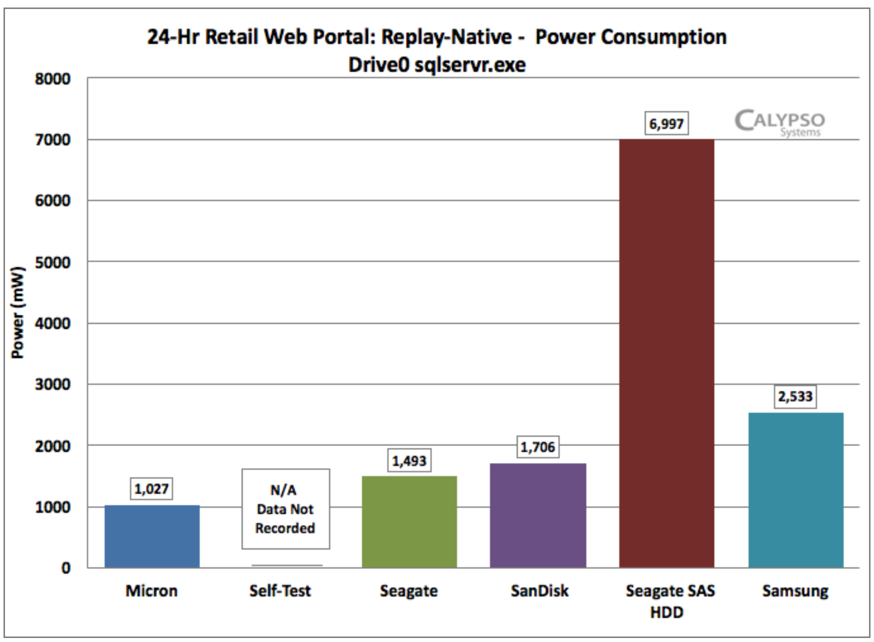

Figure 22 - Replay-Native: Power Consumption – Ave mW

G. DIRTH

The DIRTH test applies the Composite 6 IO Stream workload over a range of Users. The sqlservr.exe capture shows a QD range (brown highlight box) of QD 1 to 306 with an Average QD of 5.

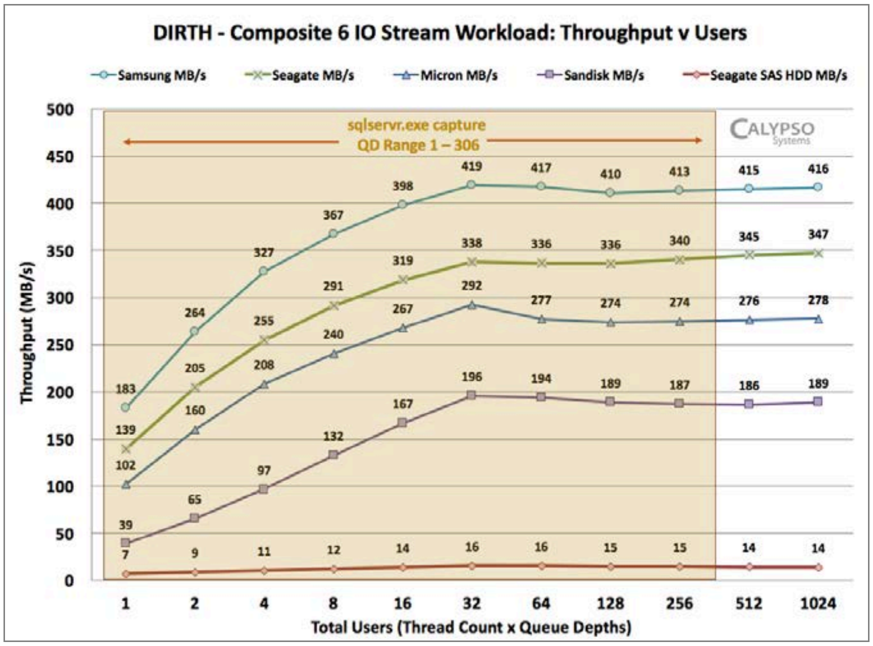

Figure 23 – DIRTH Composite Workload: Throughput v Users

Figure 24 shows 5 9s RT QoS for each drive at different Total Users (or Outstanding IOs).

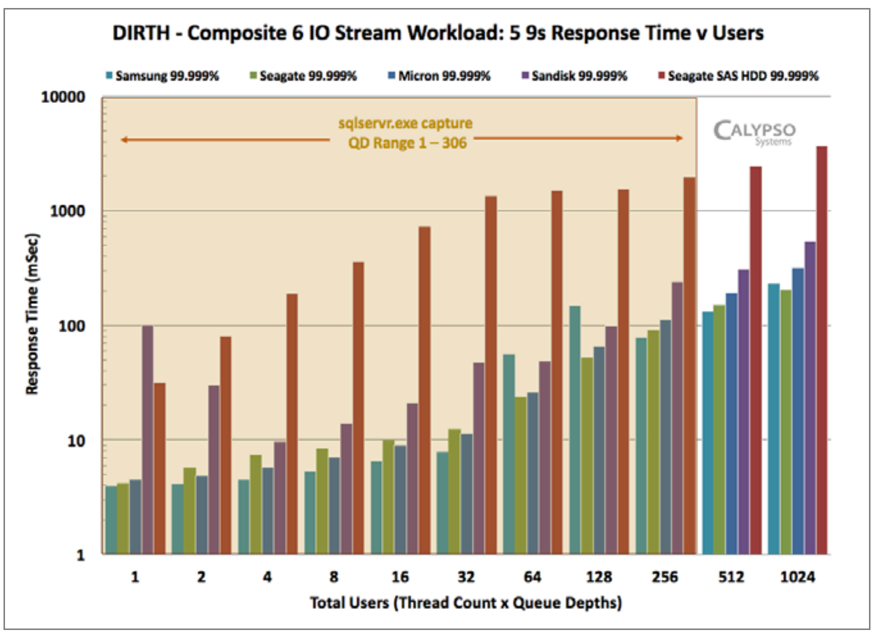

Figure 24 – DIRTH Composite Workload: 5 9s Response Time v Users

IV. Conclusions

In-situ sqlservr.exe Analysis

In-situ analysis provides data on the performance of the target storage during the IO Capture. The ability to extract and parse IO Capture workloads helps the IT professional to analyze data flow and validate software stack and storage optimizations.

Self-Test extracts sqlservr.exe IOs from the overall Block IO capture. Each capture step shows a changing combination of the 6 dominant IO Streams. Different IO Stream combinations are associated with different events and times (such as 2 am back-up, early morning hours and working hours). Queue Depths vary from an average OIO of 5 to maximum OIO of 306 with higher OIO associated with higher IO activity.

Two test workloads are created from the sqlservr.exe workload: a Composite 6 IO Stream workload with a fixed % of IO Stream occurrence; and a Replay workload that re-creates the overall sequence of changing IO Stream combinations and QDs. These workloads are used in the 4 RWSW tests applied to the sample drives.

While Compression and Duplication Ratios are also reported, analysis of these, and other advanced metrics, such as IO Bursts and IO Sequentiality (spatial and temporal locality of reference), are beyond the scope of this study.

Target Server Self-Test Results

Target Server Self-Test IO Streams and metrics can be used to create test workloads to qualify and validate storage. Self-Test metrics can also be used to compare the performance of test drive samples that are subjected to the Replay-Native workload.

Target Server Self-Test performance is typically lower than RWSW test results. This is because IO Capture performance is throttled by the number of application IOs that occur during the capture whereas RWSW apply unlimited IOs during test steps.

Multi-WSAT

The Multi-WSAT Composite 6 IO Stream workload applies the same cumulative combination of IO Streams that occur in the IO Capture for each test step. This allows the user to assess storage performance to the actual IO Streams that occur at the Block IO level. Using RWSW IO Streams is also more relevant than evaluating storage based on single access pattern corner case tests.

In this study, the Samsung SSD has the highest performance for IOPS, TP and Response Times followed closely by Seagate and Micron SSDs. Sandisk SSD and Seagate SAS HDD have lower performance.

Individual Streams WSAT

Individual Streams WSAT allows the reader to understand the individual component performance of each IO Stream of the Composite 6 IO Stream workload and to compare these values to manufacturer or single access pattern corner case results.

In the sqlservr.exe workload, the SEQ 0.5K W IO Stream occurs 28.9% and is a significant contributor to Multi-WSAT performance as test drives differ widely in their ability to process SEQ 0.5 K W IOs.

Samsung shows much higher SEQ 0.5K W performance compared to other test sample drives. Samsung also has higher performance for most other IO Stream components. Micron RND 8K W and Sandisk RND 8K R were faster than the Samsung SSD.

Replay-Native

Replay-Native applies the same sequence of IO Stream and QD steps that occur in the original IO Capture but, unlike Self-Test reports, the Replay-Native test steps are not throttled by the application IO count.

Re-creating the IO Capture IO Stream sequence also allows the user to test storage to the specific IO Stream workloads that occur at different times and for different events over the course of the capture.

For example, while Multi-WSAT applies the same cumulative Composite 6 IO Stream workload for each test step, Replay-Native applies the sequence of IO Streams and QDs that occur during the capture. This allows the user to see performance for each capture step workload and to specific capture events.

Replay-Native IOPS v Time plots show how storage responds to the specific events such as the 2 am back-up SEQ 64K R/W IO Streams, the low disk utilization and application IOs during early morning hours, and the higher IOs, QDs and mixed access patterns that occur during busy working hours.

Replay-Native Segment plots report performance averages over the 24-hour capture. These values can be compared to each other to evaluate comparative drive performance. They can also be compared to the fixed Composite IO Stream workload steps in the Multi-WSAT test by showing fine grain averages based on aggregating the individual test step values of differing IO Stream combinations.

DIRTH

DIRTH test evaluates drive performance (Throughput (TP) and 5 9s Response Times (RTs)) to the Composite 6 IO Stream workload across a range of 1 to 1,024 Users.

DIRTH shows TP performance for specific QDs as well as the range of TP over the Users. For example, Samsung shows TP of 183 MB/s at 1 user and 413 MB/s at 256 users and the TP values for each User setting between 1 and 256. The reader can also see at what point (User count) RTs begin to saturate and show large increases in RTs.

The sqlservr.exe workload Users range from 1 to 306 with an overall average QD (or Users) of 5. All drives maintain ordinal ranking across Users with Samsung showing the highest performance (highest TP and lowest 5 9s RTs) followed by Seagate, Micron and Sandisk SSDs. The Seagate SAS HDD showed the lowest TP and very high 5 9s RTs.

SUMMARY

This study shows how an IO Capture can be used to analyze in-situ performance and how to create test workloads from an extracted application workload. This multi IO Stream workload is applied both as a fixed composite as well as a collection of individual IO Streams.

RWSWs permit the user to assess how well storage responds to the multiple IO Stream combinations and changing QDs that are seen in the IO Captures and provides the industry with a more relevant and accurate way to assess and qualify Datacenter storage.

VI. Summary Results & Ordinal Rankings

A. Summary Results

Drive	Multi-WSAT			Replay-Native			DIRTH OIO=32	
	IOPS	TP MB/s	5 9s RT mS	IOPS	TP MB/s	5 9s RT mS	TP MB/s	4 9s RT mS
Samsung	17,567	447	38.17	13,054	128	7.38	419	7.9
Seagate	15,603	397	58.57	10,889	104	11.23	338	12.5
Micron	12,491	318	55.68	3,907	53	21.45	292	11.4
Sandisk	2,722	69	276.44	1,765	27	32.09	196	47.4
Seagate HDD	672	17	1,573.23	864	7	257.23	16	1,341.2

Drive	Individual Streams-WSAT					
Access Pattern % of Workload	RND 64K R 30.5%	SEQ 0.5K W 28.9%	RND 8K R 16.4%	SEQ 8K R 14.2%	SEQ 64K R 5.6%	RND 8K W 4.4%
Samsung	8,000	32,721	47,770	60,549	8,413	10,188
Seagate	7,145	15,768	34,897	50,982	7,582	8,636
Micron	5,947	3,543	49,559	44,419	5,563	7,232
Sandisk	4,177	2,843	31,295	31,578	4,154	16,775
Seagate HDD	374	1,313	473	19,452	2,600	446

Figure 25 – Summary Results

B. Ordinal Rankings

Drive	Multi-WSAT			Replay-Native			DIRTH OIO=32	
	IOPS	TP MB/s	5 9s RT mS	IOPS	TP MB/s	5 9s RT mS	TP MB/s	4 9s RT mS
Samsung	1	1	1	1	1	1	1	1
Seagate	2	2	3	2	2	2	2	3
Micron	3	3	2	3	3	3	3	2
Sandisk	4	4	4	4	4	4	4	4
Seagate HDD	5	5	5	5	5	5	5	5

Drive	Individual Streams-WSAT					
Access Pattern	RND 64K R	SEQ 0.5K W	RND 8K R	SEQ 8K R	SEQ 64K R	RND 8K W
% of Workload	30.5%	28.9%	16.4%	14.2%	5.6%	4.4%
Samsung	1	1	2	1	1	2
Seagate	2	2	3	2	2	3
Micron	3	3	1	3	3	4
Sandisk	4	4	4	4	4	1
Seagate HDD	5	5	5	5	5	5

Figure 26 – Ordinal Rankings

VII. Final Thoughts

New cross-platform IO Capture tools and workload analytics now make it easy and, yes, fun, for anyone to evaluate one's RWSW on any laptop, desktop or server. Even just using the free analytics at TestMyWorkload.com allows you to quickly see what is in your RWSW.

Do ***YOU*** know how many IOPS ***your computer*** uses or how many MB/s ***your storage*** is delivering? Do you know what kind of IOs access your storage when you run ***your applications***? Users can see how many IOPS, how much Bandwidth and what kind of accesses occur during their actual computer usage.

This knowledge can change preconceived notions about workloads and what is important for performance. IO Captures let you see the difference between File System IO activity (such as writes to cache, metadata small block appending and small block coalescing) and what IO Stream accesses make it to storage at the Block IO level (such as large block SEQ access fragmentation into smaller block concurrent RND accesses).

Using the more advanced metrics and analytics offered by the IOProfiler software, while outside the scope of this paper, allows you to observe software stack optimizations, validate data compression and de-duplication strategies, analyze spatial and temporal locality of reference (IO Bursts and Sequentiality), and extract specific workload IO Streams.

Test results become more relevant when one understands the IO Streams that comprise the RWSW. Some vendors optimize SSD performance to elevate performance of common benchmarks, such as RND 4K Writes, while neglecting other Block Size / RW mix accesses. Yes, they may actually game the specifications by optimizing a single 'glory' number. A cursory comparison of vendor specifications may show higher RND 4K W performance while ignoring the important IO Streams that are actually in your RWSW (such as SEQ 0.5K W for sqlservr.exe IO Streams as shown in this paper).

While every IO Capture is unique and depends on the target server HW/SW platform, OS, User applications and storage architecture, all RWSWs share the common characteristics of a dynamically changing combination of different IO Streams and Demand Intensity.

These new IO Capture tools and analytics can help Datacenter, IT and storage professionals understand what IO Streams actually occur on their storage systems. Analysis of in-situ performance for specific applications, events or time periods can help optimize, design and validate software stack optimizations and storage qualification.

The RWSW tests reported here introduce a repeatable, accurate and relevant test methodology and toolset that will allow storage and IT professionals to better understand how storage and Datacenter software stack designs can be improved, how costs can be reduced, and help troubleshoot and solve field failure events.

The sqlservr.exe capture analyzed in this paper can be viewed as Example No. 3 on the TestMyWorkload website at http://testmyworkload.com/info/demo/#exampleKB24hr. For more information, contact asksssi@snia.org.

About the Author

Eden Kim is the CEO of Calypso Systems, Inc., a pioneer in solid state storage test and measurement, supplier of industry standard SSD test equipment and software and the publisher of TestMyWorkload.com, the first industry Real World Storage Workload capture and analysis site.

Eden is also the Chair of the SNIA Solid State Storage Initiative (SSSI) Technical Development committee, Chair of the SNIA Solid State Storage Technical Working Group (TWG) and a member of the SSSI Governing Board. Mr. Kim has been leading the research into Real World Storage Workloads through the SSS TWG technical programs, work with leading Datacenter customers and through efforts in developing the IOProfiler toolset and TestMyWorkload.com website.

Mr. Kim regularly speaks at Industry Trade and Association events and recently presented papers and published articles on Real World Storage Workloads at the Flash Memory Summit Santa Clara 2018, the China Flash Summit in Beijing June 2017, Datacenter Storage Europe November 2016, Storage Developers Conference Santa Clara September 2017, Wuhan GSS 2018 show in China, the Flash Memory Summit in Santa Clara August 2018, and the Beijing Data Storage Summit. Mr. Kim matriculated from the University of California and holds a Juris Doctor degree from UC Davis and a Bachelor of Science degree from UCSB. Mr. Kim can be contacted at edenkim@calypsotesters.com.

About the SNIA

The Storage Networking Industry Association (SNIA) is a not–for–profit global organization, made up of member companies spanning the global storage market. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For more information, visit http://www.snia.org.

About the SNIA Solid State Storage Initiative

The SNIA Solid State Storage Initiative (SSSI) fosters the growth and success of the market for solid state storage in both enterprise and client environments. Members of the SSSI work together to promote the development of technical standards and tools, educate IT communities about solid state storage, perform market outreach that highlights the virtues of solid state storage, and collaborate with other industry associations on solid state storage technical work. SSSI member companies represent a variety of segments in the IT industry (see http://www.snia.org/forums/sssi/about/members).

For more information on SNIA's Solid State Storage activities, visit www.snia.org/forums/sssi and get involved in the conversation at http://twitter.com/SNIASolidState.