

How To Be A Part of the Real World Workload Revolution

A SNIA Solid State Storage Initiative Webcast

July 9, 2019

Today's Presenters



Eden Kim
Calypso Systems



Jim Fister
The Decision Place

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - Any slide or slides used must be reproduced in their entirety without modification
 - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

About SNIA



185

industry leading
organizations



2,000

active contributing
members



50,000

IT end users & storage
pros worldwide

Learn more: snia.org/technical



Today's Agenda

- Real World Workloads, Comparisons to Synthetic Workloads
- Measuring and Visualizing Real World Workloads
- Real World Workload Demonstration
- Testing and Sharing
- Benefits of Joining the Program

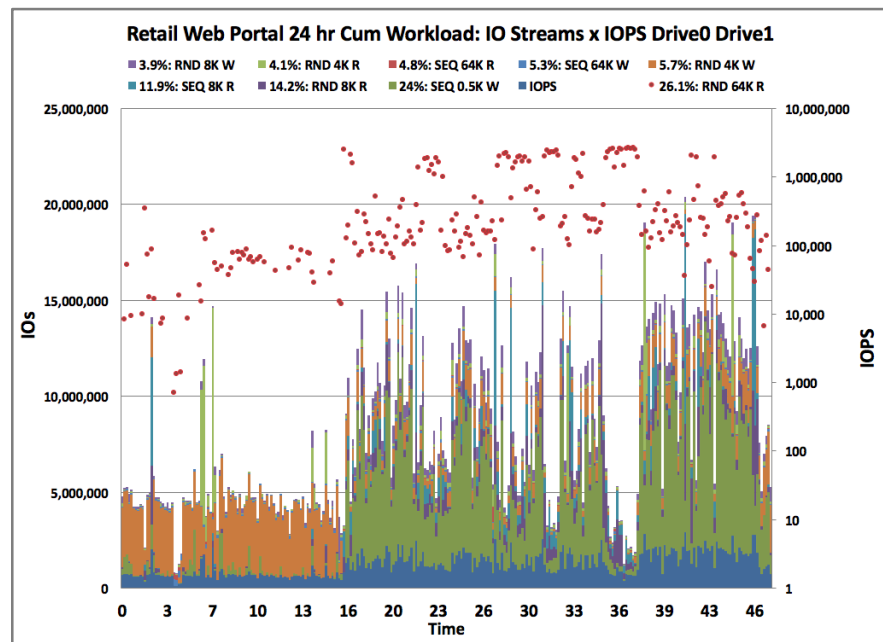
What are Real-World Workloads?

Real World Workloads are:

- A constantly changing sequence of
- IO Stream Combinations and QDs
- Captured from point A to B
- In the SW/HW Stack during actual use

Synthetic Benchmark Workloads are:

- A fixed and constant workload
- Applied to storage
- To test performance
- Outside the Range of Normal Usage



Retail Web Portal 24 Hr Workload: IO Streams & IOPS

Synthetic vs. Real-World Workloads

Synthetic SQL Workload:

- Typically defined as random 8K R/W
- 65:35 R/W ratio
- Limited I/O threads to storage

Actual Real-World Workload:

- 72:28 R/W ratio
- 10% 8K Read, 2.8% 8K Write
- Over 5000 unique I/O streams
- Factors in additional workloads on the system (i.e. Backup)

Σ	Cumulative Workload	
<input checked="" type="checkbox"/>	RND 64K R	18.5% 842,361
<input checked="" type="checkbox"/>	SEQ 0.5K W	17.0% 775,127
<input checked="" type="checkbox"/>	RND 8K R	10.0% 456,175
<input checked="" type="checkbox"/>	SEQ 8K R	8.4% 382,972
<input checked="" type="checkbox"/>	RND 4K W	4.0% 182,251
<input checked="" type="checkbox"/>	SEQ 64K W	3.7% 169,571
<input checked="" type="checkbox"/>	SEQ 64K R	3.4% 155,798
<input checked="" type="checkbox"/>	RND 4K R	2.92% 132,777
<input checked="" type="checkbox"/>	RND 8K W	2.78% 126,550
<input type="checkbox"/>	SEQ 4K R	2.20% 100,309
Total IOs of 5,086 streams: 4,551,062		
Selected 9 streams: 3,223,582 (71%) E		

A	2 am Back-up	
<input checked="" type="checkbox"/>	SEQ 64K W	45.1% 156,395
<input checked="" type="checkbox"/>	SEQ 64K R	41.6% 144,397
<input type="checkbox"/>	RND 4K W	1.83% 6,362
<input type="checkbox"/>	SEQ 2M R	1.40% 4,851
<input type="checkbox"/>	RND 4K R	0.95% 3,294
<input type="checkbox"/>	RND 32K R	0.75% 2,595
<input type="checkbox"/>	RND 24K R	0.56% 1,940
<input type="checkbox"/>	SEQ 60K R	0.54% 1,883
<input type="checkbox"/>	SEQ 40K R	0.51% 1,772
<input type="checkbox"/>	RND 8K R	0.44% 1,519
Total IOs of 298 streams: 346,825		
Selected 2 streams: 300,792 (87%) E		

Why are Real World Storage Workloads (RWSWs) Important?



RWSWs are key determinants in storage performance and have a significant impact on SW and storage optimizations

- Solid State Storage Performance depends, in large part, on RWSWs
- Unlike lab test workloads, RWSWs are comprised of dynamically changing combinations of IO Streams & Demand Intensity
- IO Stream content affects Optimization, Design, Validation & Failure Analysis
- IO Streams change at each layer of software abstraction in the SW Stack

What are Real World Storage Workloads?



A collection of IO Streams seen by Storage

Generated in Application space

During Real World computer usage

Now, Let's look at some Definitions:

What is an IO Stream?

Access Pattern	RND or SEQ	Block Size	Read/Write	Queue Depth Ave/Max	% Occurrence	Quantity (IOs)
SEQ 1.5K W	SEQ	1536	W	1/111	1.34	69
SEQ 1K W	SEQ	1024	W	1/111	4.32	223
SEQ 0.5K W	SEQ	512	W	1/111	9.24	477
SEQ 4K W	SEQ	4096	W	1/111	22.31	1152
SEQ 16K W	SEQ	16384	W	1/111	14.25	736
RND 4K W	RND	4096	W	1/111	9.8	506
RND 3.5K W	RND	3584	W	1/111	0.62	32
RND 3K W	RND	3072	W	1/111	0.58	30
RND 2.5K W	RND	2560	W	1/111	0.74	38
RND 8K R	RND	8192	R	1/111	0.15	8
RND 2K W	RND	2048	W	1/111	0.93	48
RND 1.5K W	RND	1536	W	1/111	1.74	90
RND 1K W	RND	1024	W	1/111	3.21	166
RND 0.5K W	RND	512	W	1/111	1.99	103
RND 8K W	RND	8192	W	1/111	2.73	141
RND 4K R	RND	4096	R	1/111	0.91	47
RND 12K W	RND	12288	W	1/111	1.24	64
RND 16K W	RND	16384	W	1/111	15.63	807
RND 20K W	RND	20480	W	1/111	0.58	30
RND 28K W	RND	28672	W	1/111	2.03	105
RND 36K W	RND	36864	W	1/111	0.19	10

IO Stream Table: 2 Minute Capture Step showing IO Stream Statistics

An IO Stream¹ is an Input/Output Operation (IO) that has a unique:

- Random or Sequential Access
- Block Size or Data Transfer Size
- Read or Write IO
- Queue Depth (QD)

A single IO Stream can occur many times during an IO Capture Step
Other Secondary Metrics can be associated with IO Streams

¹ Workload IO Stream definitions used here differ from SSD Endurance Streams where similar write operations are associated with a group of associated data

What is an IO Capture?


An IO Capture is the tabulation of statistics on IO Streams that are observed during a capture period



- IO Capture Tools gather statistics and metrics on IO Streams
- An IO Capture is NOT an IO Trace
- No data or private information is collected
- Only binary numeric tables are gathered

SEQ 128b R SEQ 4K W RND 4K W RND 16K W SEQ 0.5K W SEQ 16K W RND 128b R SEQ 1K W
RND 64K R SEQ 0.5K W RND 8K R SEQ 8K R RND 4K W SEQ 64K R RND 8K W RND 4K R RND 16K R
SEQ 4K R RND 32K R RND 24K R SEQ 512K R RND 40K R SEQ 16K R SEQ 1K W RND 12K R RND 56K R


What is the Software Stack?



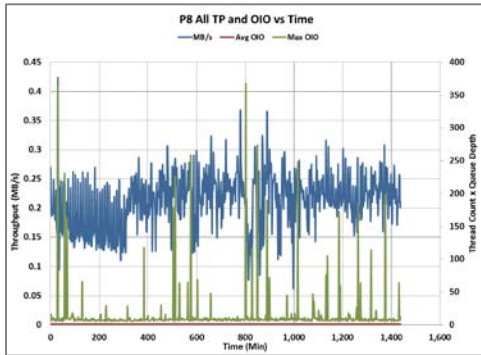
The Software (SW) Stack refers to the layers of software (OS, APIs, programs, drivers and abstractions) that exist between User space and storage

- IO Streams are generated in User space by software applications, run everywhere
- IO Streams traverse the SW State back and forth to storage
- IO Stream composition is different at different levels of the SW Stack

Why do RWSWs Affect Performance?



RWSW – IOPS & Changing Combinations of IO Streams



RWSW - Throughput & Changing Queue Depths

Solid State Storage Performance depends on how well storage responds to constantly changing combinations of IO Streams and Demand Intensity

- RWSWs are constantly changing combinations of IO Streams and QDs; Synthetic Lab tests are a fixed and constant
- IO Streams vary by Block Size, Accesses and R/W IOs
- Solid State Storage responds differently to the type of access (RND/SEQ), Block Size, and whether the IO is a Read/Write
- The type and combination of RWSW IO Streams and the Demand Intensity determines, in large part, the storage performance that is provided

What do RWSWs Look Like?

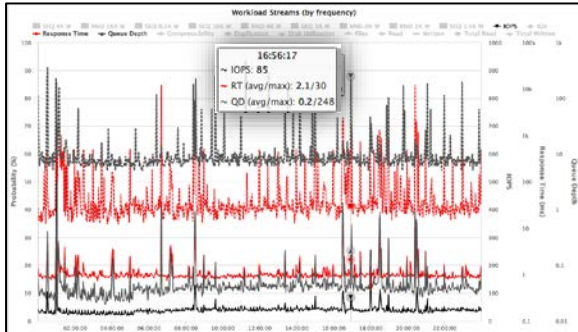
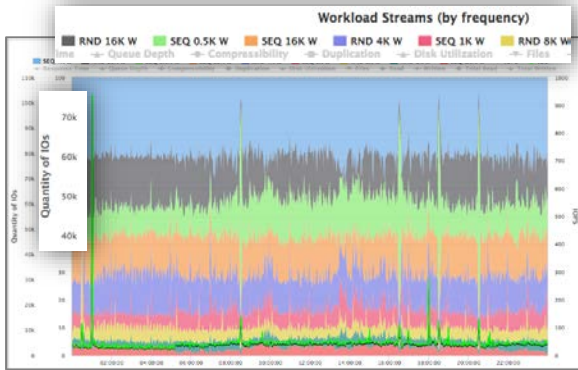


IO Stream Map: 24-hour Block IO Level (Drive0) Capture showing IOPS & IO Streams

RWSW can be visualized by creating an IO Stream Map that shows the changing combinations of IOs and metrics over Time

- IO Stream % probability of occurrence - different color data series
- IOPS - dominant black line
- Time - X-axis: 24-hour capture at 2 minute steps
- Secondary metrics captured by the Capture tool can be displayed

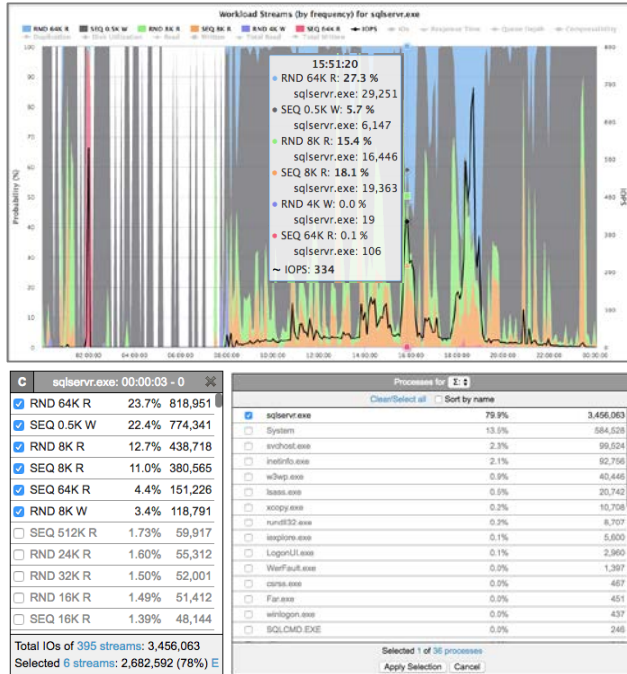
Viewing Secondary IO Metrics



Secondary IO Metrics captured by the IO Capture tool can be viewed on an IO Stream Map

- IO Capture tools can capture various Secondary Metrics
- IO Streams can be listed by RND/SEQ access, Block Size, R/W IO
- Average and Maximum Response Times and QDs are shown to the left
- Several additional IO Metrics can be shown


Viewing IO Streams by a Specific Application




IO Streams can be filtered, extracted or presented by specific Application IOs

- Here we extract only the sqlservr.exe IOs
- sqlservr.exe IOs that occur > 3% of the time over the 24-hours
 - The Cumulative workload - 395 Streams & 36 separate Processes
 - sqlservr.exe – 6 IO Streams are 78% of the Total IO Streams
 - sqlservr.exe application IOs - 79.9% of the Total IO Streams


What can RWSWs Tell Us?




IO Stream Composition of Selected Processes



Changing Combinations of IO Streams



Application Process Response Times




Throughput and Queue Depths v Time

RWSWs let us extract specific application IOs to analyze the IO Composition, Metrics and Performance that occur on the target server during the IO Capture. We can see:

- IO Stream Distribution of the selected Process(es)
- Changing combinations of IO Streams over time
- IO Process Average & Maximum Response Times
- IO Process Throughput and Avg/Max Queue Depth
- Secondary IO Metrics

Demonstration


RWSW Test Examples




Multi-WSAT



Individual Streams-WSAT



DIRTH



Replay-Native

1. **Multi-WSAT** – applies fixed 6 IO Stream composite workload for each test step to Steady State.
2. **Individual Streams-WSAT** – tests each individual IO Stream as a separate Steady State item
3. **DIRTH** (Demand Intensity Response time Histogram) - applies fixed 6 IO Stream composite workload across a range of 1 to 1,024 Users to measure IOPS & Response Time saturation.
4. **Replay-Native** - reproduces each capture step combination of IO Streams, Queue Depths and Idle Times for storage comparison to the original IO Capture server storage

SNIA SSSI Reference IO Capture Workloads

TestMyWorkload
My Captures Download [Demo](#) How To Use What It Means IPF Portal FAQ Support
Eden

Sample Capture: Corpor

Profile Details Drive0 Drive1

Name: Sample Capture: Corporate Portal During 24 hours

Platform: Windows Server 2012 R2

File Name: WIN-C255D607N82_20161602_0000.ipf

Started: 02/15/2016 11:00:04

Ended: 02/16/2016 11:08:46

Duration: 1 day

Steps:

Resolution: 5 min

Profiler Ver.: 1.0.43

Comp. Name: WIN-C255D607N82

OS: Windows version: 6.1.7601 Service pack: 1.0 64-bit OS

CPU: Intel(R) Xeon(R) CPU E5-2670 v2 @ 2.50GHz

RAM: 4.0 GB

SNIA SSSI Reference IO Captures:

Example 3. 24-Hr Retail Web Portal
Windows Server 2012R2
Drive 0 & Drive 1
Block IO Level Captures

Example 4. SNIA Green Storage TWG Workload
Linux OS - Block IO Level
IO capture of Synthetic Test Workload

Example 6. 24-Hr GPS Navigation Portal
Windows Server 2012R2
File System & Block IO Level

ad Streams (by frequency)

Path							
\\.\PhysicalDrive0	Virtual HD	214 GB	4,326,159	36	142.9 GiB	20.6 GiB	163.5 GiB
\\.\PhysicalDrive1	Virtual HD	193 GB	224,903	38	10.1 GiB	12.0 GiB	22.2 GiB
Total for 2 drives with data		407 GB	4,551,062	38	153.1 GiB	32.6 GiB	185.7 GiB

Comments: Capture of whole day block-level I/O activity of Windows server running corporate portal (Internet Information Services + MS SQL Server). Different load phases can be viewed very easily.

Benefits Across the Industry

Hardware Manufacturers

- Build/utilize an expanse of tests
- Characterize true needs of Solid State Storage/Memory
- Tune driver development

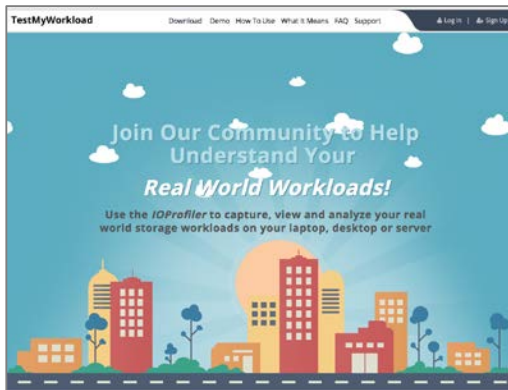
Software Developers

- Provide workloads to get hardware tuned for your performance
- Provide tools for purchasing decision makers
- Find race conditions and test new hardware more quickly

IT Organizations

- Provide your workload to get better results for your work
- Make the right choices in hardware/system specifications

Conclusions



- Understanding **Your** RWSW is **critical** for Datacenter, Storage Server and SSD Design and Optimization
- Know what you are optimizing (IO Streams & SW Stack levels)
- Find out what IO Streams are presented to Storage
- Try Demos, Free Captures & Analysis at TestMyWorkload.com
- See SNIA Reference Captures at TestMyWorkload.com
- View the SNIA Whitepaper and Calypso Whitepaper
- www.TestMyWorkload.com

Thanks for Attending

More Information:

Webcast slides: <https://www.snia.org/forums/sssi/knowledge/articles-presentations>

How To Be a Part of the Real World Workload Revolution white paper (English and Chinese versions):
<https://www.snia.org/forums/sssi/knowledge/whitepapers>

Webcast Q&A: <http://sniasssiblog.org>

Visit www.snia.org/sssi
for information on Solid State Storage