SNIA Global Education

Tutorial

Understanding Real World Storage Workloads

Eden Kim, Calypso Systems, Inc.

About the Instructor

- Eden Kim is CEO of Calypso Systems, Inc., a leader in SSD test and measurement equipment and services. Eden is also Chair of the SNIA Solid State Storage Technical Working Group

SNIA Legal Notice

- The material contained in this tutorial 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 Education Committee.
- 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.

Abstract

- Storage Optimization for Data Center IT Professionals, Storage designers and Storage OEM/ODMs

This session will appeal to Data Center IT Managers, IT Professionals, Storage Designers, Storage OEM/ODMs and those that are seeking a fundamental understanding of Real World Storage Workloads. The session will explain what Real World Workloads are, how they are captured and why they are important. Not limited to developers, the session will also bring a clear understanding of the value of Real World Storage Workloads to the IT community and Systems Integrators. The audience will receive the fundamental understanding of how and why Real World Storage Workloads are important in Datacenter storage, server and storage device design, optimization and qualification.

Real World Storage Workloads

SNIA Global Education

1. What are Real World Storage Workloads (RWSWs)?
2. Why are RWSWs Important?
3. What do RWSWs Look Like?
4. What can RWSWs Tell me?
5. How do I use RWSWs for Optimization, Design & Validation?
6. What is Next?
7. Questions & Answers

1. What are Real World Storage Workloads?

Real World Storage Workloads (RWSW) are:

- The collection of discrete IO Streams
- Observed at a specific level in the SW Stack
- That access physical or logical storage
- Over a given period of time

1.1 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 Stream[1] 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

See Table where a single SEQ 1.5K W IO Stream occurs 69 times

[1] **IO Stream definitions used here refer to Real World Storage Workloads which differ from Data Streams used in relation to SSD Endurance where similar write operations are associated with a given Data Stream.**

1.2 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
- IO Capture Tools are Operating System (OS) specific
- Public and private IO Capture Tools are available[2]
- IO Capture Tools vary by: OS, Level in the SW Stack and Metrics taken

[2] Examples of public tools include perfmon for Windows and blocktrace for Linux. Private tools include hiomon for Windows by hyperIO and IOProfiler for Windows, Linux, FreeBSD and MacOS by Calypso. IOProfiler captures and results were used for the data presented herein and can be seen as Demo No. 3 at TestMyWorkload.com

1.3 What is the Software Stack?

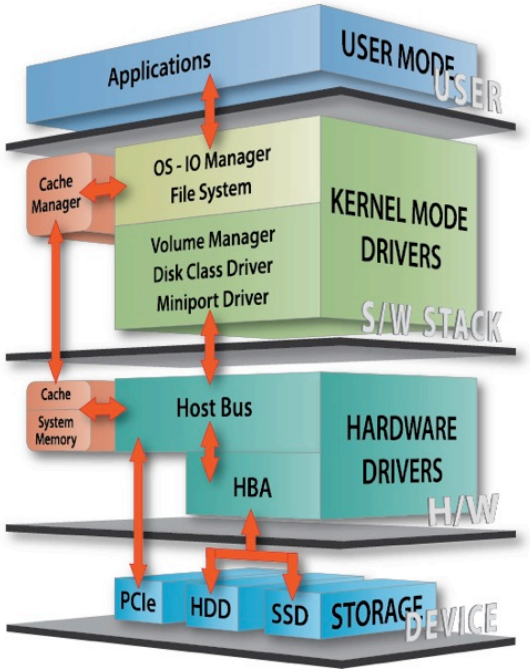

Windows Software Stack

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
- IO Streams traverse the SW Stack to storage and back
- IO Stream composition is different at different levels in the SW Stack

2. Why are 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 Streams change at each layer of software abstraction
- IO Stream content affects Optimization, Design, Validation & Failure Analysis

2.1 Why do RWSW Affect Storage Performance?

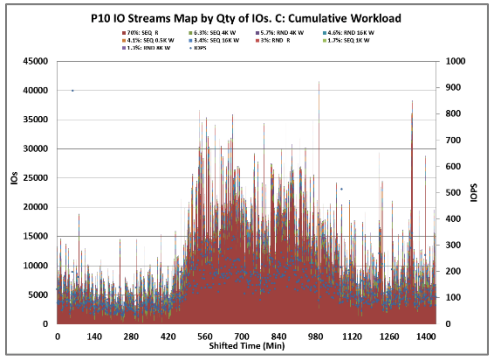

RWSW – IOPS & Changing Combinations of IO Streams

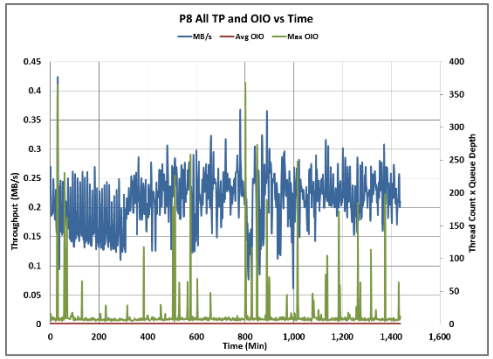

RWSW - Throughput & Changing Queue Depths

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
- IO Streams will have different Block Sizes, Accesses and Read or Write IOs
- Synthetic Lab tests are a fixed and constant workload
- Solid State Storage responds differently to the type of access (RND or SEQ), the Block Size, and whether the IO is a Read or a Write
- The type and combination of RWSW IO Streams and the Demand Intensity determines, in large part, the storage performance that is provided

2.2 RWSW - Key for Optimization, Validation & Qualification

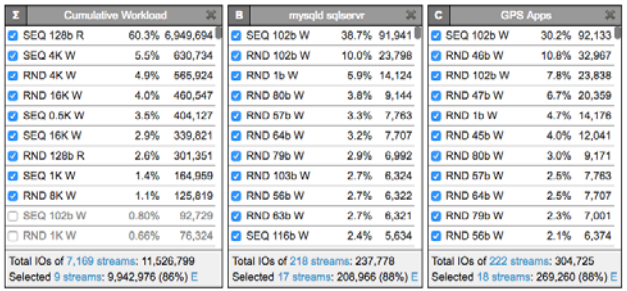

Σ Cumulative Workload		
SEQ 128b R	60.3%	6,949,694
SEQ 4K W	5.5%	630,734
RND 4K W	4.9%	565,924
RND 16K W	4.0%	460,547
SEQ 0.5K W	3.5%	404,127
SEQ 16K W	2.9%	339,821
RND 128b R	2.6%	301,351
SEQ 1K W	1.4%	164,959
RND 8K W	1.1%	125,819
SEQ 102b W	0.80%	92,729
RND 1K W	0.66%	76,324
Total IOs of 7,169 streams: 11,526,799		
Selected 9 streams: 9,942,976 (86%) E		

B mysqld sqlservr		
SEQ 102b W	38.7%	91,941
RND 102b W	10.0%	23,798
RND 1b W	5.9%	14,124
RND 80b W	3.8%	9,144
RND 57b W	3.3%	7,763
RND 64b W	3.2%	7,707
RND 79b W	2.9%	6,992
RND 103b W	2.7%	6,324
RND 56b W	2.7%	6,322
RND 63b W	2.7%	6,321
SEQ 116b W	2.4%	5,634
Total IOs of 218 streams: 237,778		
Selected 17 streams: 208,966 (88%) E		

C GPS Apps		
SEQ 102b W	30.2%	92,133
RND 46b W	10.8%	32,967
RND 102b W	7.8%	23,838
RND 47b W	6.7%	20,359
RND 1b W	4.7%	14,176
RND 45b W	4.0%	12,041
RND 80b W	3.0%	9,171
RND 57b W	2.5%	7,783
RND 64b W	2.5%	7,707
RND 79b W	2.3%	7,001
RND 56b W	2.1%	6,374
Total IOs of 222 streams: 304,725		
Selected 18 streams: 269,260 (88%) E		

File System Level – Dominant IO Streams

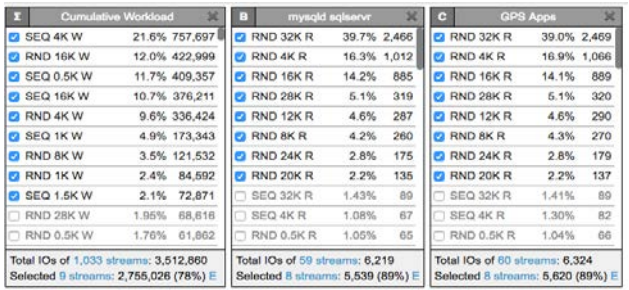

Σ Cumulative Workload		
SEQ 4K W	21.6%	757,697
RND 16K W	12.0%	422,999
SEQ 0.5K W	11.7%	409,357
SEQ 16K W	10.7%	376,211
RND 4K W	9.6%	336,424
SEQ 1K W	4.9%	173,343
RND 8K W	3.5%	121,532
RND 1K W	2.4%	84,592
SEQ 1.5K W	2.1%	72,871
RND 28K W	1.95%	68,616
RND 0.5K W	1.76%	61,862
Total IOs of 1,033 streams: 3,512,860		
Selected 9 streams: 2,755,026 (78%) E		

B mysqld sqlservr		
RND 32K R	39.7%	2,466
RND 4K R	16.3%	1,012
RND 16K R	14.2%	885
RND 28K R	5.1%	319
RND 12K R	4.6%	287
RND 8K R	4.2%	260
RND 24K R	2.8%	175
RND 20K R	2.2%	135
SEQ 32K R	1.43%	89
SEQ 4K R	1.08%	67
RND 0.5K R	1.05%	65
Total IOs of 59 streams: 6,219		
Selected 8 streams: 5,539 (89%) E		

C GPS Apps		
RND 32K R	39.0%	2,469
RND 4K R	16.9%	1,066
RND 16K R	14.1%	889
RND 28K R	5.1%	320
RND 12K R	4.6%	290
RND 8K R	4.3%	270
RND 24K R	2.8%	179
RND 20K R	2.2%	137
SEQ 32K R	1.41%	89
SEQ 4K R	1.30%	82
RND 0.5K R	1.04%	66
Total IOs of 60 streams: 6,324		
Selected 8 streams: 5,620 (89%) E		

Block IO Level – Dominant IO Streams

RWSW characterization is essential for software optimization, firmware validation, storage qualification & Failure Analysis.

- Each layer of abstraction in the SW Stack can change IO Streams
- Abstractions include metadata, data compression, encryption, data de-duplication, storage virtualization, storage tiers, back-up, snapshots and more
- IO Streams can be appended, fragmented, coalesced or written to cache
- File System level IO Streams tend to include smaller byte accesses to cache
- Block IO level IO Streams tend to include larger KB accesses to storage
- Understanding the IO Stream composition at different SW Stack layers is key to Software and Storage optimization

3. What do RWSWs look like?

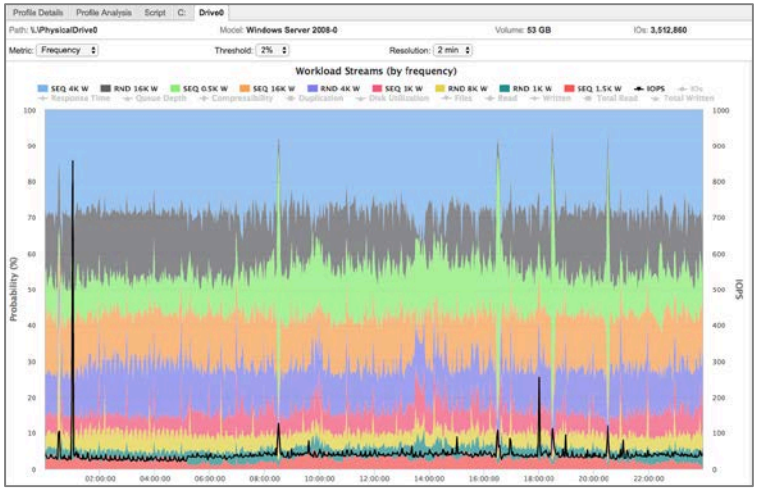

IO Stream Map: 24-hour Block IO Level (Drive0) Capture showing IOPS & IO Streams

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 is shown as different color data series and plotted against the Y-axis
- IOPS are shown by the dominant black line and are plotted against the secondary Y-axis
- Time is shown along the X-axis: 24-hour capture at 2 minute steps
- Secondary metrics captured by the Capture tool can be displayed

3.1 Viewing Secondary IO Metrics

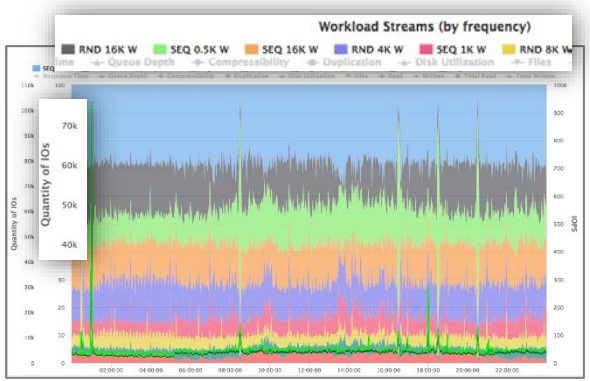

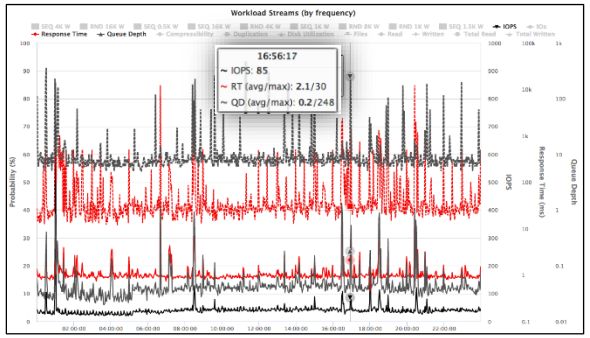

Secondary IO Metrics that are captured by the IO Capture tool can be viewed on an IO Stream Map.

- IO Capture tools can capture various Secondary IO Metrics
- IO Streams can be listed by RND/SEQ access, Block Size & R/W IO
- Average & Maximum Response Times and QDs can be shown
- Additional IO Metrics can be shown for IO Count, Duplication Ratio, Compression Ratio, Disk Utilization, Reads, Writes and more

3.2 Viewing IO Streams by Application / Process

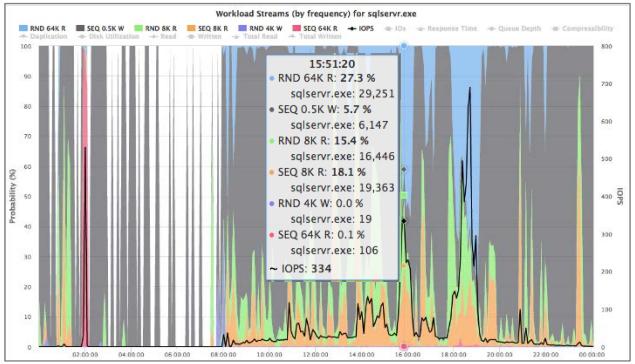

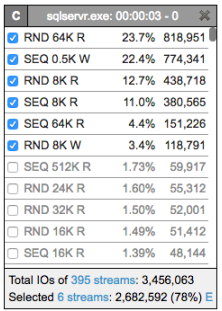

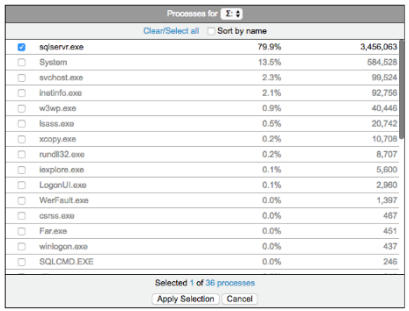

IO Streams can be filtered, extracted or presented by Application or Process IDs (PIDs)

- IO Capture Tools can associate PIDs to specific IO Streams
- The Cumulative IO Capture shows 395 Total Streams and 36 Total PIDs
- IOs can be filtered by Application, Process IDs and by IO Stream % of occurrence over the capture duration
- The IO Stream Map here shows sqlservr.exe PIDs that occur > 3% of the time over the course of the 24-hour IO Capture
- (6) sqlservr.exe Streams are 78% of the Total IO Streams
- (1) sqlservr.exe PIDs is 79.9% of the Total PID IO Streams

4. What can RWSWs tell me?

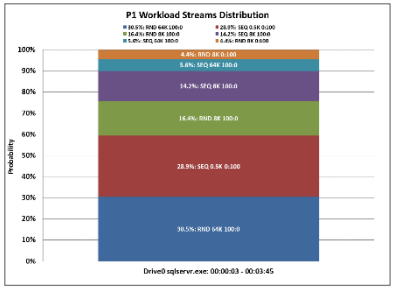

IO Stream Composition of Selected Processes

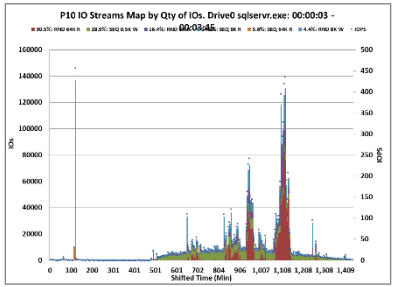

Changing Combinations of IO Streams

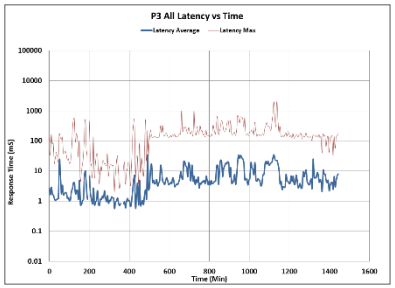

Application Process Response Times

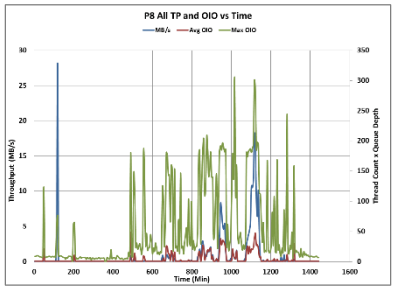

Throughput and Queue Depths v Time

RWSWs let you extract specific application IOs to analyze the IO Composition, Metrics and Performance that occur on the target server during the IO Capture. You 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 Average & Maximum Queue Depth
- Secondary IO Metrics gathered by the IO Capture tool

5. How do I use RWSWs for Optimization, Design & Validation?

RWSWs let you isolate/extract IO Stream workloads at different SW Stack levels to confirm the efficacy of software designs, the impact of abstractions and the content of workloads at the storage level.

- File System or Block IO level Captures
- Extract Application specific processes
- Examine in-situ server performance
- Evaluate Total Storage, Logical Units or Devices
- Replay for Failure Analysis or Storage Qualification

6. What's Next?

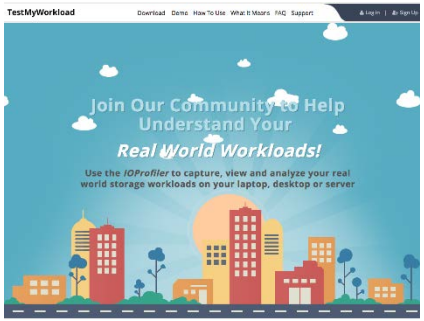

RWSWs make a difference! SNIA urges the Storage Community to actively participate in the capture and analysis of RWSWs. You can help to:

- Improve IO Capture Tools: Support OSes & IO Storage Metrics
- Increase Capture Database: Utilize Free Capture & Analysis at TestMyWorkload.com[3]
- Create Data Analytics: Analysis of Captures & SW Stack IO Streams
- Develop Methodologies: For Storage Qualification & Evaluation
- Establish Industry Standards: Participate in SNIA Technical Works

[3] IO Capture tools, images, analytics and data presented herein were taken using the Free tools available at www.TestMyWorkload.com - a SNIA SSSI collaborative site with Calypso Systems, Inc. The sample capture can be viewed as Demonstration Capture No. 3 'Corporate Web Portal – 24 hr capture' on the TestMyWorkload site at http://testmyworkload.com/info/demo/#exampleKB24hr

7. Questions & Answers

Contact SNIA for more information.

- SNIA Education: snia.org/education
- Solid State Storage Initiative: asksssi@snia.org
- SSSI TechDev Committee: sssi_techdev@snia.org
- Solid State Storage Technical Working Group: ssstwg@snia.org

Check out SNIA Tutorial:

Enterprise Applications:

How to Create a Synthetic Workload Test

Attribution & Feedback

The SNIA Education Committee thanks the following Individuals for their contributions to this Tutorial.

Authorship History

Eden Kim / August 2017

Updates:

Additional Contributors

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***