




How will your Data Center SSD Handle Real World Workloads?

Eden Kim, Calypso Systems, Inc.

What are Real World Datacenter Workloads?




- ◆ Real World Workloads are a **collection of IO Streams** that occur on a **deployed server** at a specific **software stack level** over a given **period of time**
- ◆ Real World Workloads are a **constantly changing combination of many, many IO Streams and Queue Depths**
- ◆ IO Streams are generated by **Applications, Operating System activities, software abstractions, User activities and more**
- ◆ **IO Streams change as they traverse the Software and Hardware Stack**




Flash Memory Summit

Why are Real World Workloads Important?

- ◆ **Workload Composition** – IO Stream Content
- ◆ **Load Balancing** – IO Traffic & Performance
- ◆ **Software Optimization** – Validate SW Abstractions
- ◆ **Interoperability** – Standard Replay Scripts
- ◆ **Failure Analysis** – Workload Capture & Replay
- ◆ **Server & Storage Qualification** – Replay Test Qualification




Viewing a Datacenter Real World Workload:





New SNIA Spec for Datacenter Real World Workloads



SNIA[®]
Advancing storage &
information technology

**Real World Storage Workload (RWSW)
Performance Test Specification for
Datacenter Storage**

Version 1.0.7

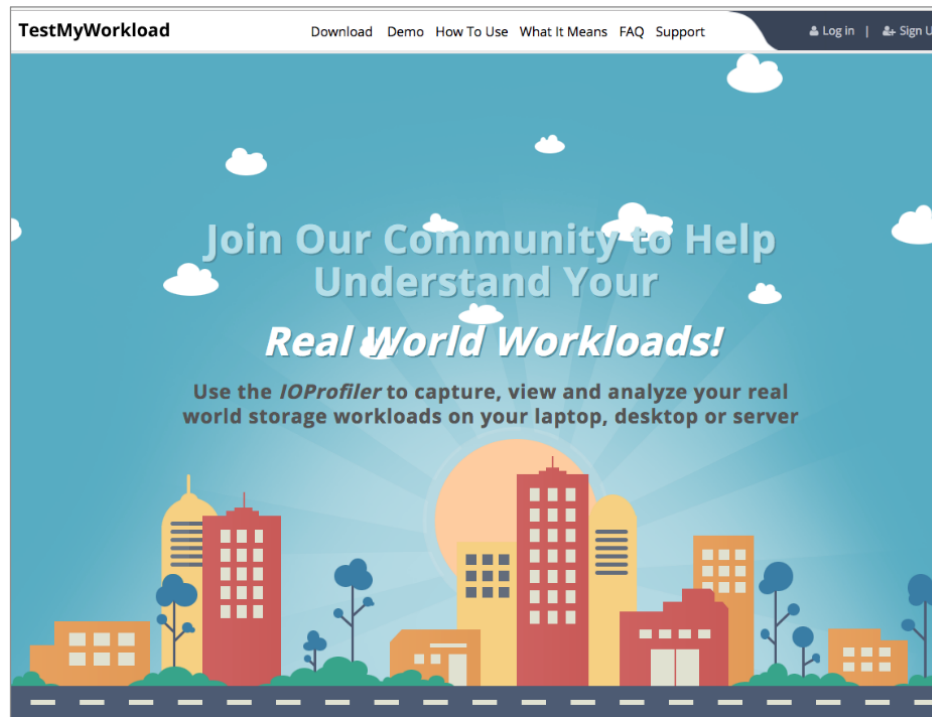
ABSTRACT: This document describes a Real-World Storage Workload (RWSW) IO capture, characterization, methodology, test suite and reporting format. It is intended to provide standardized analysis of in-situ target server application storage performance and standardized comparison and qualification of Datacenter storage when using Reference IO Capture Workloads as the test stimuli in RWSW tests.

This document has been released and approved by the SNIA. The SNIA believes that the ideas, methodologies and technologies described in this document accurately represent the SNIA goals and are appropriate for widespread distribution. Suggestions for revisions should be directed to <http://www.snia.org/feedback/>.

SNIA Technical Position
May 25, 2018

- **SNIA Technical Position**
- **RWSW PTS Spec**
- **IO Capture Process**
- **IO Stream Analysis**
- **Datacenter Storage**
- **Standard RWSW Tests**
- <https://www.snia.org/rsw>

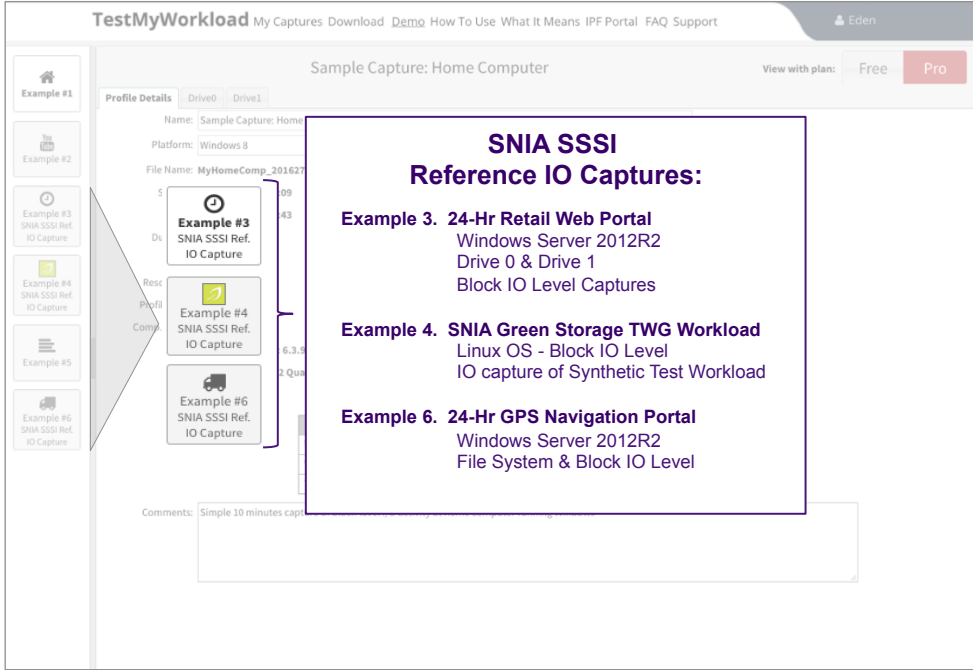
Free Capture Tools & Demos at TestMyWorkload.com



www.TestMyWorkload.com

- ◆ www.TestMyWorkload.com – SSSI reference site for Real World Workload captures and free cross platform IO Capture tools
- ◆ Free Capture Tools include:
 - ◆ Blk-Trace, Perfmon, D-trace, IOProfiler
- ◆ Free IOProfiler Capture tools at www.TestMyWorkload.com:
 - ◆ Windows, Linux, MacOS
 - ◆ Free visualization of captured workloads
 - ◆ Free export of IO data for use in 3d party software

SSSI Reference Captures No. 3, 4 & 6



The screenshot displays the TestMyWorkload web interface. The main content area is titled "Sample Capture: Home Computer" and shows a list of SSSI Reference IO Captures. A callout box highlights three specific examples:

- Example #3:** SNIA SSSI Ref. IO Capture
- Example #4:** SNIA SSSI Ref. IO Capture
- Example #6:** SNIA SSSI Ref. IO Capture

The callout box also lists the following details for these examples:

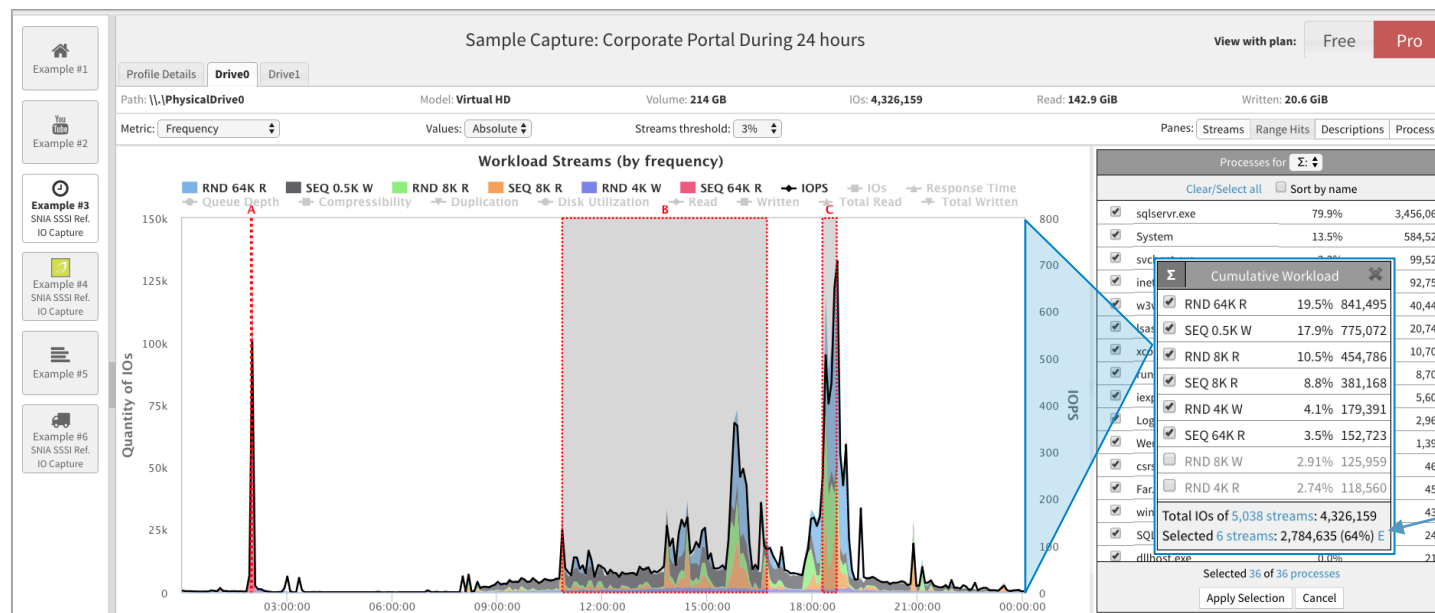
- Example 3. 24-Hr Retail Web Portal:** Windows Server 2012R2, Drive 0 & Drive 1, Block IO Level Captures
- Example 4. SNIA Green Storage TWG Workload:** Linux OS - Block IO Level, IO capture of Synthetic Test Workload
- Example 6. 24-Hr GPS Navigation Portal:** Windows Server 2012R2, File System & Block IO Level

www.testmyworkload.com/info/demo

- **Example 3:** 24 Hr Retail Web Portal
- **Example 4:** SNIA Green Storage TWG Workload
- **Example 6:** 24 Hr GPS Navigation Portal



IO Stream Map, Processes, Cumulative Workload



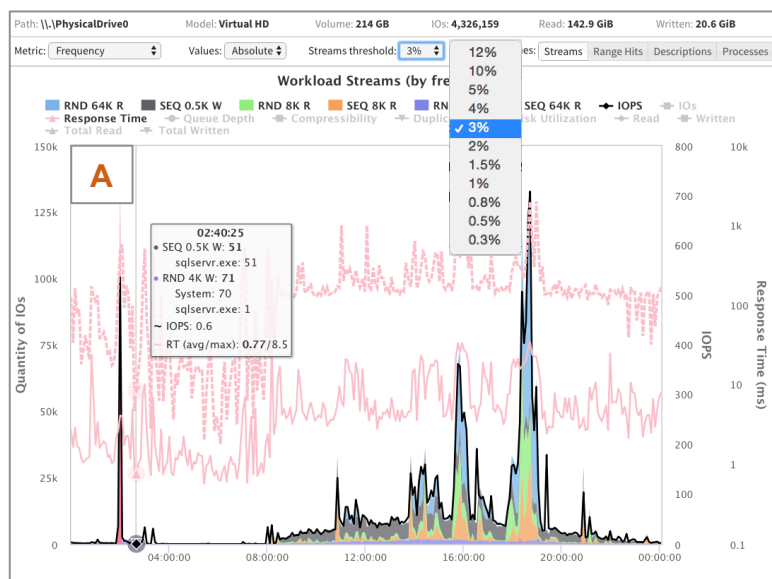
Example No. 3: IO Stream Map and Process IDs

Cumulative Workload:
6 IO Streams of 5,038 IO Streams
64% of IOs
2,784,635 IOs of 4,326,159 IOs

Export IO Stream Steps to csv file



Analytics: IOPS or MB/s, Latency, Queue Depth

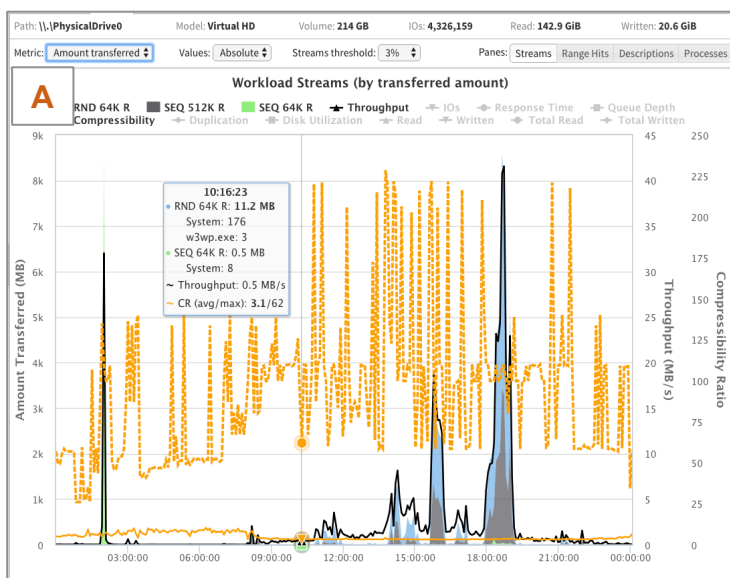


A IO Stream Map by Frequency – IO Rate in IOPS
IOPS, Ave/Max Response Times, 3% Threshold




B IO Stream Map by Amount Transferred – Throughput in MB/s
MB/s, Ave/Max Queue Depth

Analytics: Compression & Duplication Ratios




A Compression Ratio: How much more compressible is data
CR of 3.1 means data can be compressed 3.1 Times MORE

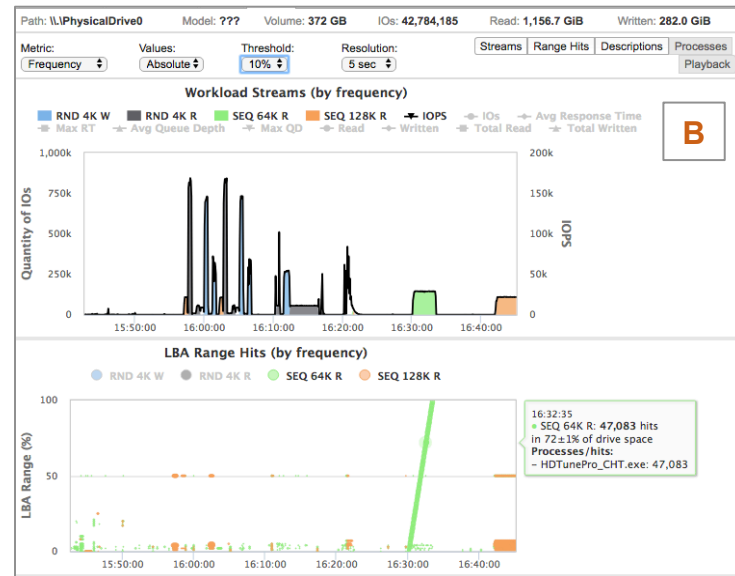


B Duplication Ratio – How many duplicative blocks are written
DR of 28% means that 28% of written blocks are duplicates

Analytics: Disk Utilization & IO Sequentiality

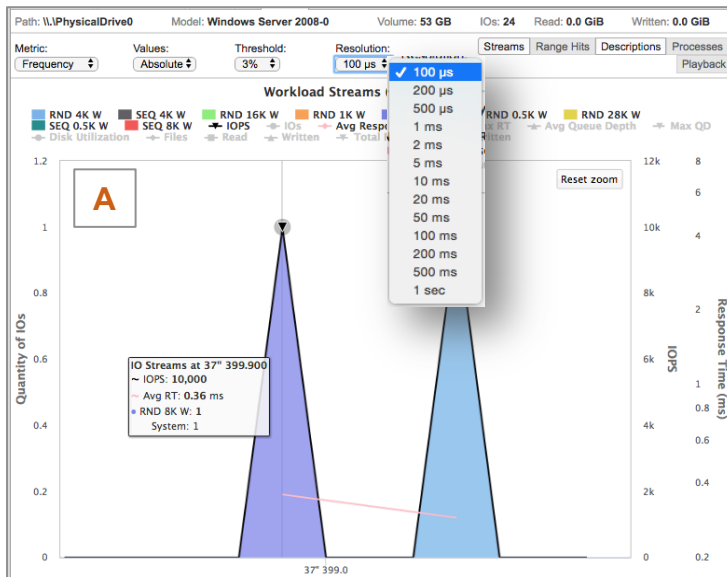


A Disk Utilization: IOs and Disk IO Idle Times
 Disk Utilization of 84% = 16% Disk IO Idle

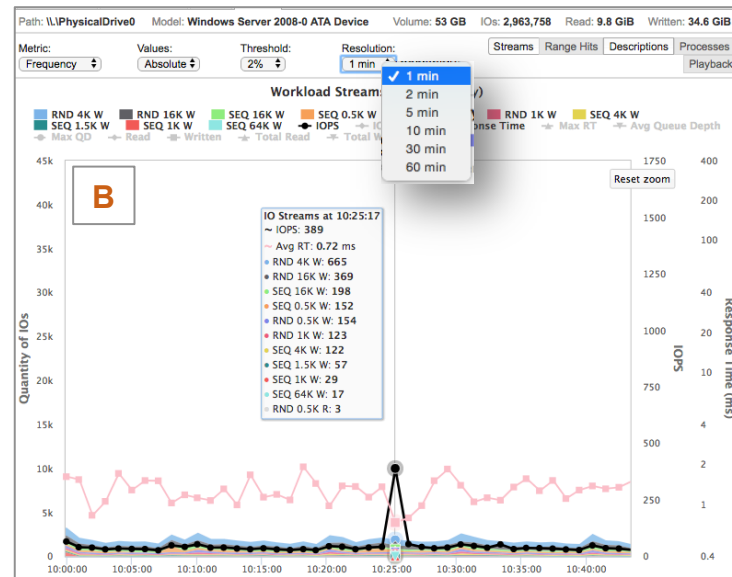


B IO Sequentiality – Adjacent LBA Range Hits
 Diagonal LBA Range Hit lines indicate Sequential IOs

Temporal Granularity: IO Capture Step Resolution

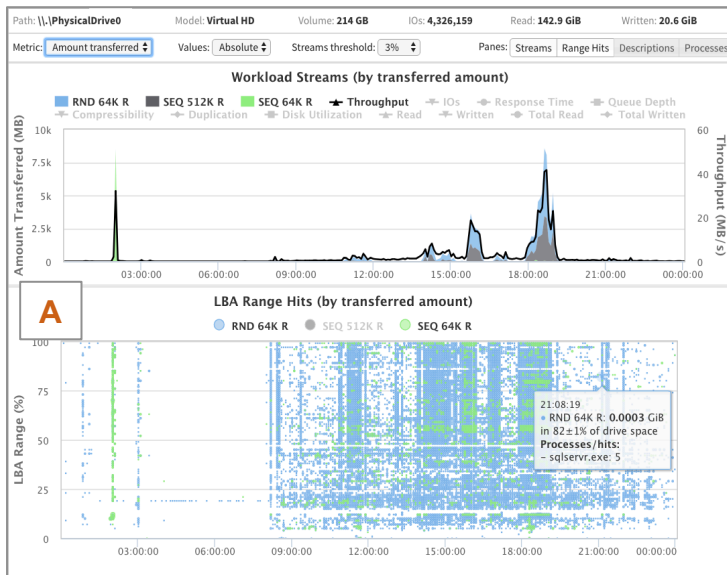


A Fine Grain: 100 μ s; 1 min Capture; Single IOs
IO Bursts, Disk Utilization, IO Specificity



B Coarse Grain: 1 Min; 24 hr Capture; 100's of IOs
Long Term workload characterization

LBA Range Hit Map: Spatial Locality of Reference



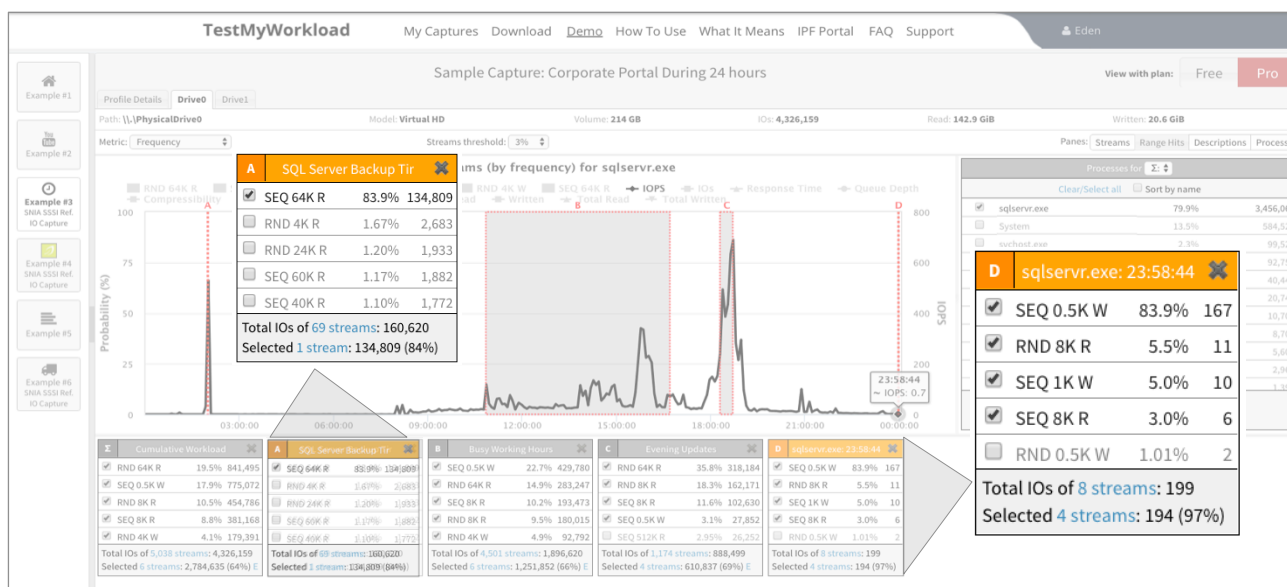
A LBA Hit Range – by Transferred Amount (MB/s)
 IO Hits selected by Size (Amount Transferred)



B LBA Hit Range – by Frequency (IOPS)
 IO Hits selected by IO Rate (Frequency)




Defining Workload Segments



- **A: SQL Server Back up** – 83.9% SEQ 64K R; 1 of 69 IO Streams
- **D: 24 hr sqlservr.exe** – 79.9% of Total IOs, 4 of 194 IO Streams


Creating a Workload & Script




- Filter & Select the IO Stream Workload
- List IOs and Metrics for each Capture Step
- Selected points illustrate changing IO Streams & QDs for each Step of the Capture




Tests from RWSW PTS



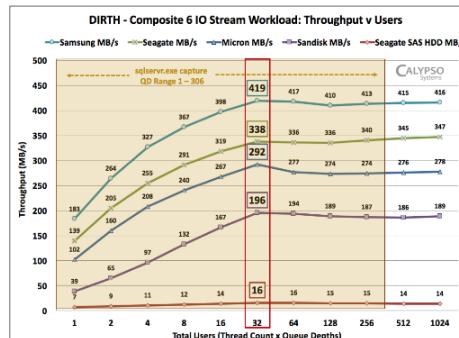
Replay-Native



Individual Streams-WSAT



Multi-WSAT



DIRTH

- 1. 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
- 2. Individual Streams-WSAT** – tests each individual IO Stream as a separate Steady State measurement
- 3. Multi-WSAT** – applies fixed 6 IO Stream composite workload for each test step to Steady State
- 4. 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



Flash Memory Summit

Questions?

Thank You!

Eden Kim

CEO, Calypso Systems, Inc.
Chair SNIA SSS Technical Working Group

edenkim@calypsotesters.com

www.TestMyWorkload.com