SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE
Innovation in Storage Products, Services, and Solutions
June 13-15, 2016 | Marriott San Mateo | San Mateo, CA

An Examination of User Workloads for SSDs

Eden Kim

Calypso Systems, Inc.

An Examination of User Workloads for SSDs

- Why SSD Performance Depends on the Workload
- Why We Test with Real World Workloads
- User Workload Examples:
  - Example A: Retail Store Web Portal 24 hr Capture
  - Example B: Mac OSX You Tube 20 min Capture
  - Example C: Windows 8 Home PC 20 min Capture
  - Example D: Court Room Video Surveillance 10 hr Capture
  - Example E: Anti Virus Boot Drive Scan

SSD Performance…

It's all about the workload

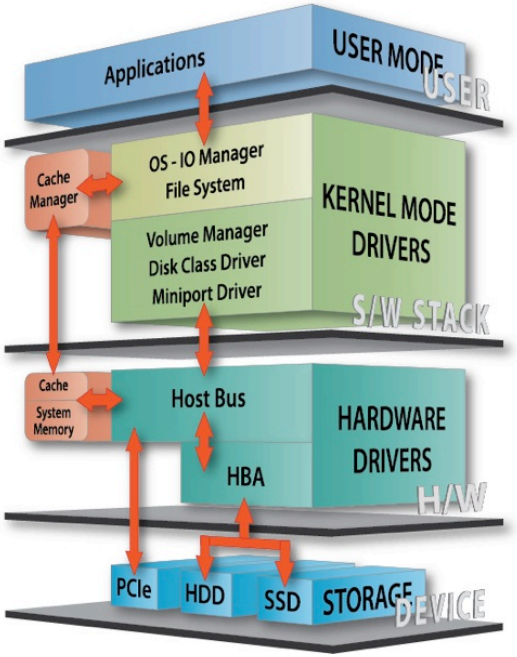

Where you measure it

Data ingress, application space, NIC packet traffic, LUN, Array or Device

IO Streams are affected at every level of abstraction

Metadata, data reduction, virtualization, coalescing, fragmentation, etc.

How you define it

Workloads are comprised of many, many IO streams

Differing Access Patterns, Data Transfer Sizes and RW mixes

Unique Data Content and Spatial and Temporal Locality of Reference

Different Demand Intensity (users and jobs) and the number of outstanding IOs

How you test it

Corner Case Benchmark, Synthetic Application or Real World workloads

Here, we examine several applications to determine what the storage actually 'sees'

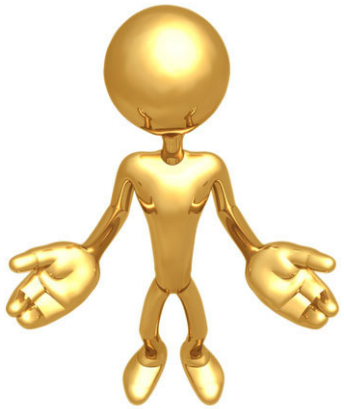

Different ways to test SSDs…

Each one has its place

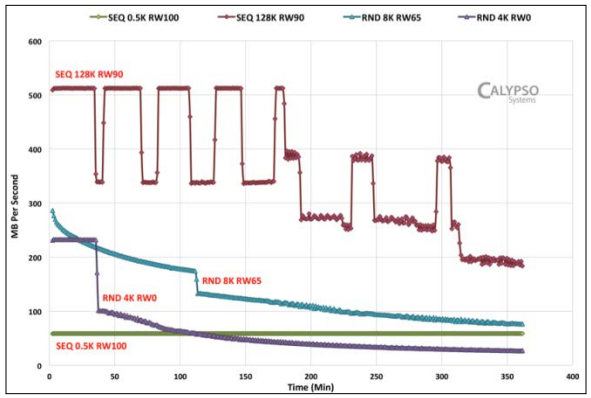

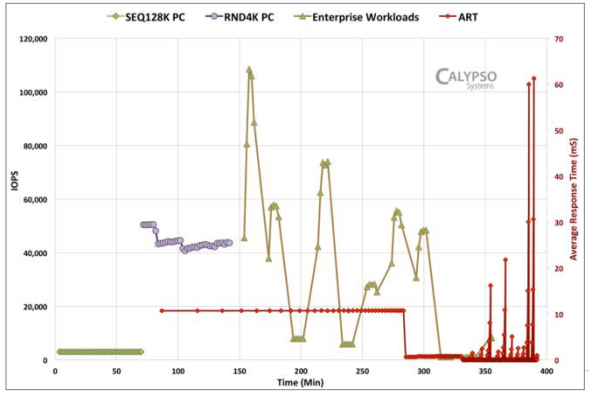

Corner Case Benchmark Tests

Convenient way to provide a quick comparison

SNIA Performance Test Specification (PTS-E/C) is a good way

However, real world workloads are never a single corner case workload

Synthetic Application Workloads

Synthetic approximations of commonly used applications

Allow more finely tuned test parameters and conditions

However, application workloads differ on different systems and at different times

Real World IO Capture Workloads

Creates specific test workloads based on IO captures of real world applications

IO Streams at the Data Center Storage LUN presents what the storage actually ‘sees’

However, they change over the course of a day and no two systems are identical

Why Test with Real World Workloads?

Because we can…

IO trace and capture tools are available to capture & characterize real world workloads

Test tools allow the creation of real world test workloads based on these trace captures

Real world workloads provide another dimension to evaluate storage

We need to know what our actual workloads look like

Everyone's deployed application workloads are unique to their hardware / software solution

Workload definition depends on where in the IO stack the workload is captured

Workloads change over the course of the day depending on the use case of the storage system

We want to test storage to the actual deployed workload

It is the best way to measure YOUR workload

Captures can confirm what IO streams actually are presented to the storage

Create or confirm Data Center Storage Tiering strategies - see what IO traffic goes where

Real World Workload Capture

EXAMPLE A

24 Hour 2,000 outlet retail webserver portal

Audience Participation!!

How many different IO Streams were observed in 24 hours?

A. 0 – 100

B. 101 – 500

C. 501 – 1,000

D. > 1,000

Quiz Answer is….

5,038!

How many different IO Streams were observed in 24 hours?

A. 0 – 100

B. 101 – 500

C. 501 – 1,000

D. > 1,000

Example A:
24 Hour 2,000 Store Webserver Portal Capture

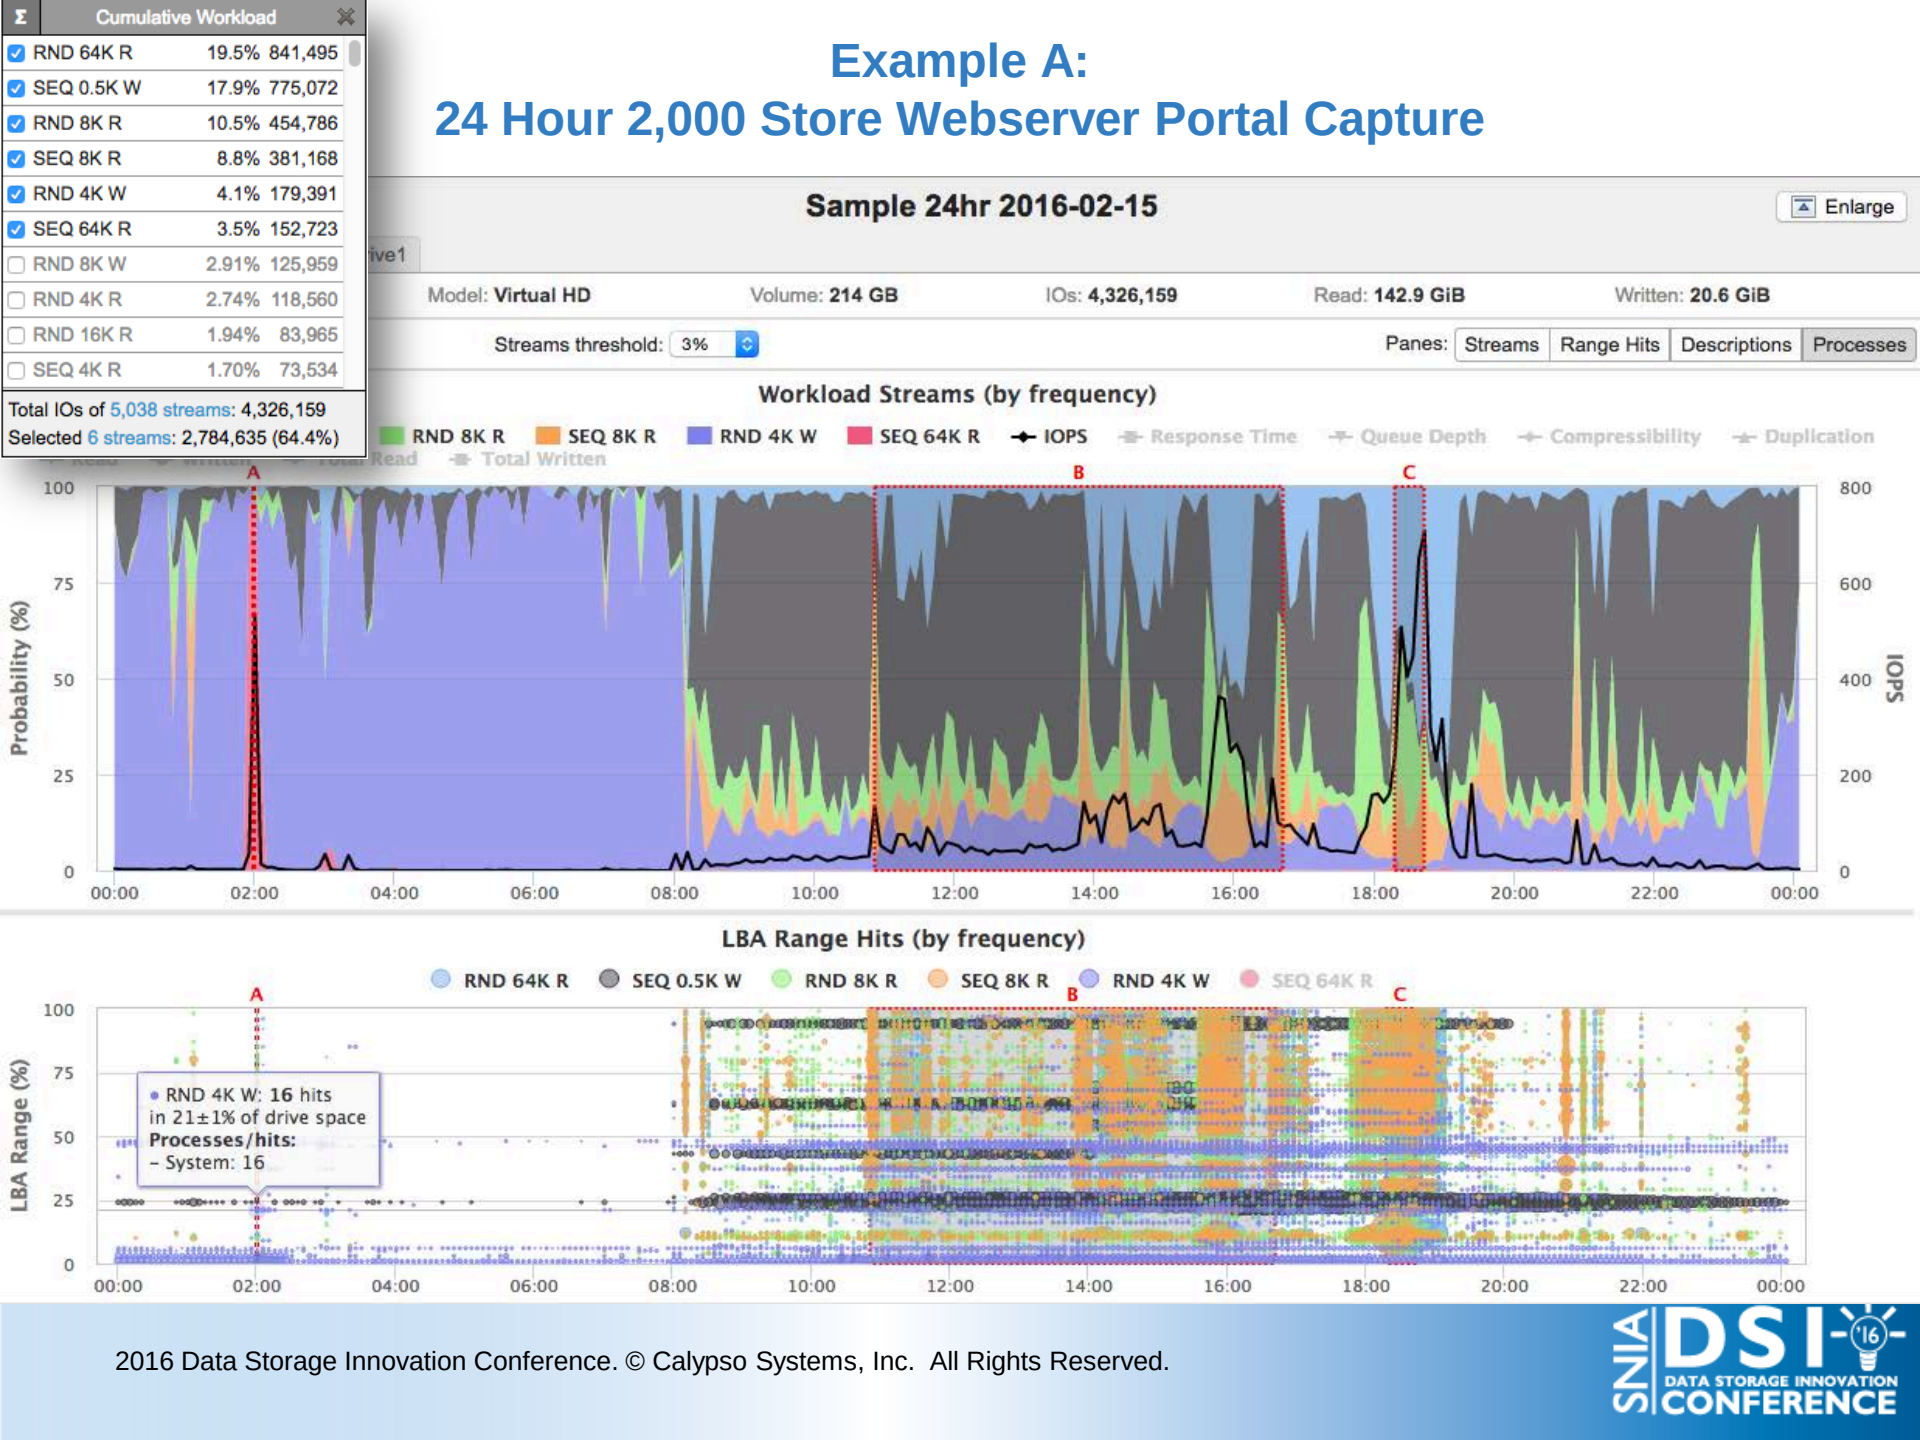

Example A: Retail Web Portal
Workload Segment Definition / Analysis

Hide Profiles | **Sample 24hr 2016-02-15** | Enlarge

Profile Details | Script | **Drive0** | Drive1

Path: **\\.\PhysicalDrive0** Model: **Virtual HD** Volume: **214 GB** IOs: **4,326,159** Read: **142.9 GiB** Written: **20.6 GiB**

Metric: Frequency Streams threshold: 3% Panes: Streams | Range Hits | Descriptions | Processes

Workload Streams (by frequency)

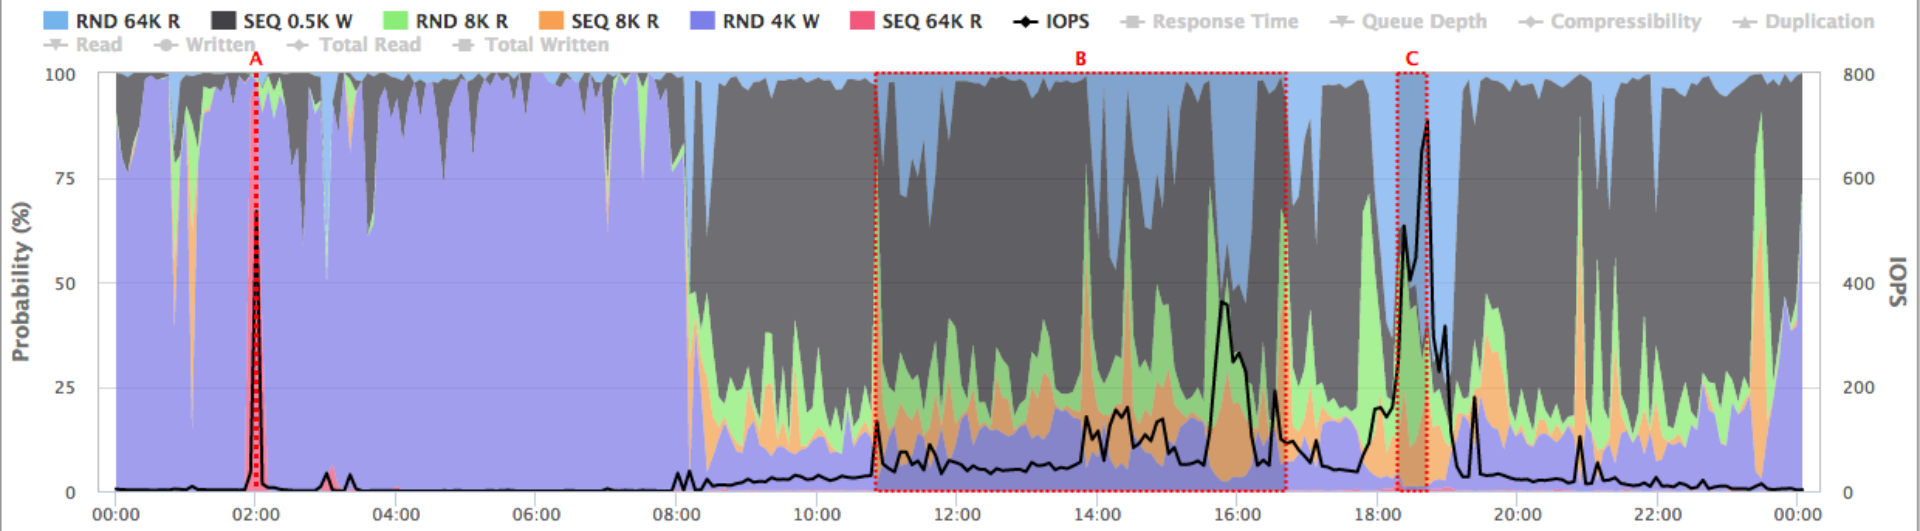

Σ Cumulative Workload

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
RND 16K R	1.94%	83,965
SEQ 4K R	1.70%	73,534

Total IOs of 5,038 streams: 4,326,159
Selected 6 streams: 2,784,635 (64.4%)

A SQL Server: 02:00:20

SEQ 64K R	83.9%	134,809
RND 4K R	1.67%	2,683
RND 24K R	1.20%	1,933
SEQ 60K R	1.17%	1,882
SEQ 40K R	1.10%	1,772
RND 32K R	0.88%	1,421
SEQ 32K R	0.84%	1,342
RND 8K R	0.82%	1,313
RND 12K R	0.79%	1,276
RND 36K R	0.69%	1,113

Total IOs of 69 streams: 160,620
Selected 1 stream: 134,809 (83.9%)

B Daily Operation

SEQ 0.5K W	22.7%	429,780
RND 64K R	14.9%	283,247
SEQ 8K R	10.2%	193,473
RND 8K R	9.5%	180,015
RND 4K W	4.9%	92,792
RND 8K W	3.8%	72,545
RND 4K R	2.81%	53,231
RND 16K R	1.99%	37,718
SEQ 4K R	1.63%	30,842
RND 32K R	1.51%	28,685

Total IOs of 4,501 streams: 1,896,620
Selected 6 streams: 1,251,852 (66%)

C Closing

RND 64K R	35.8%	318,184
RND 8K R	18.3%	162,171
SEQ 8K R	11.6%	102,630
SEQ 0.5K W	3.1%	27,852
SEQ 512K R	2.95%	26,252
RND 16K R	2.19%	19,482
RND 24K R	2.19%	19,422
RND 32K R	2.16%	19,234
SEQ 16K R	2.10%	18,690
RND 40K R	1.77%	15,733

Total IOs of 1,174 streams: 888,499
Selected 4 streams: 610,837 (68.7%)

Example A: Retail Web Portal
SQL Server over 24 Hours

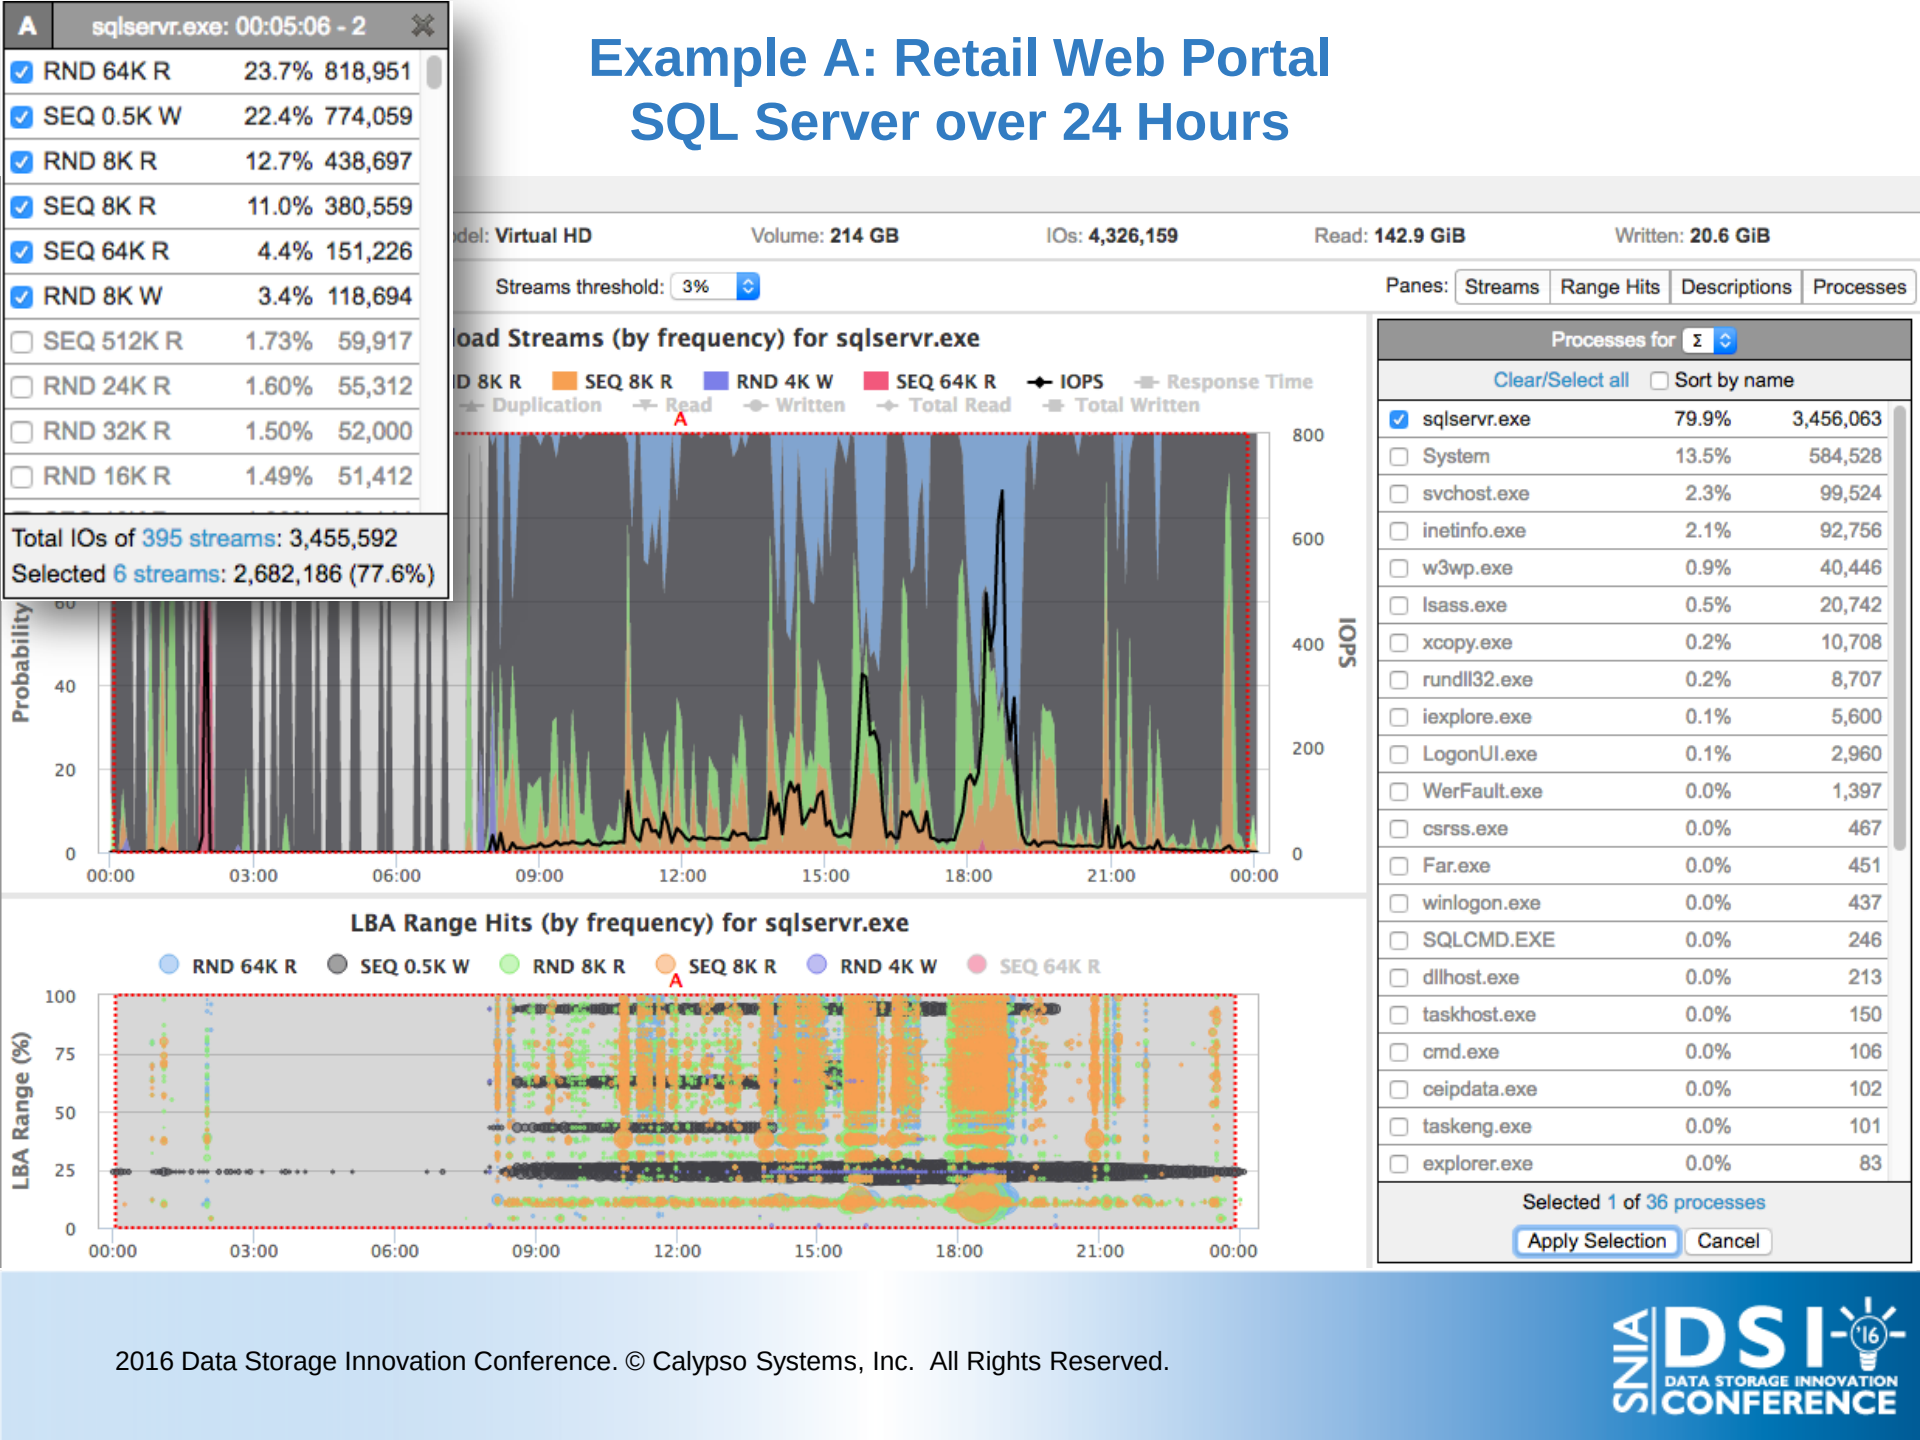

SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE

Real World Workload Capture

EXAMPLE B

Mac OSX You Tube 20 min Video

Example B1:
Mac OSX You Tube – Block IO

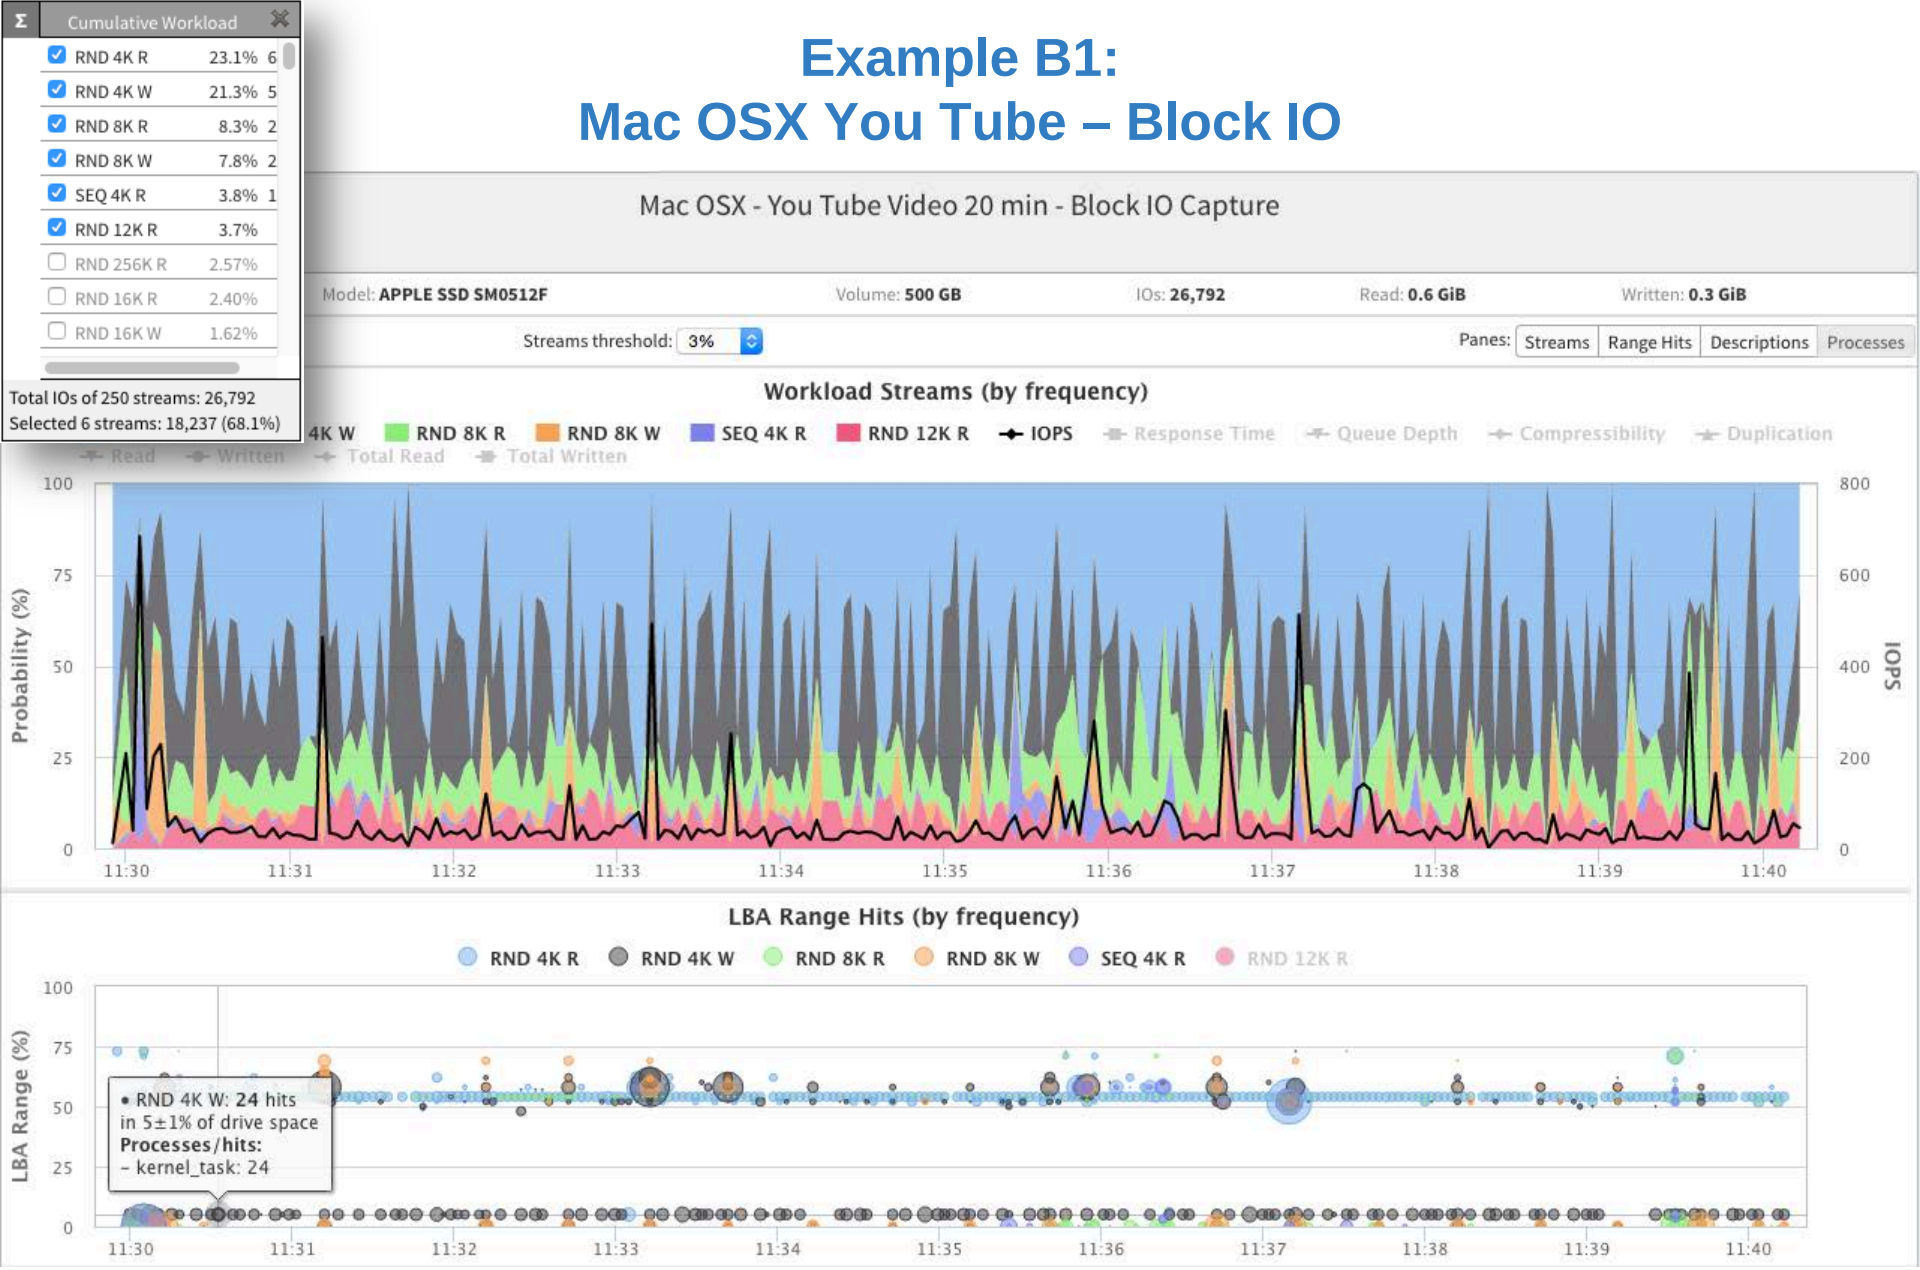

Example B2:
Mac OSX You Tube – File System

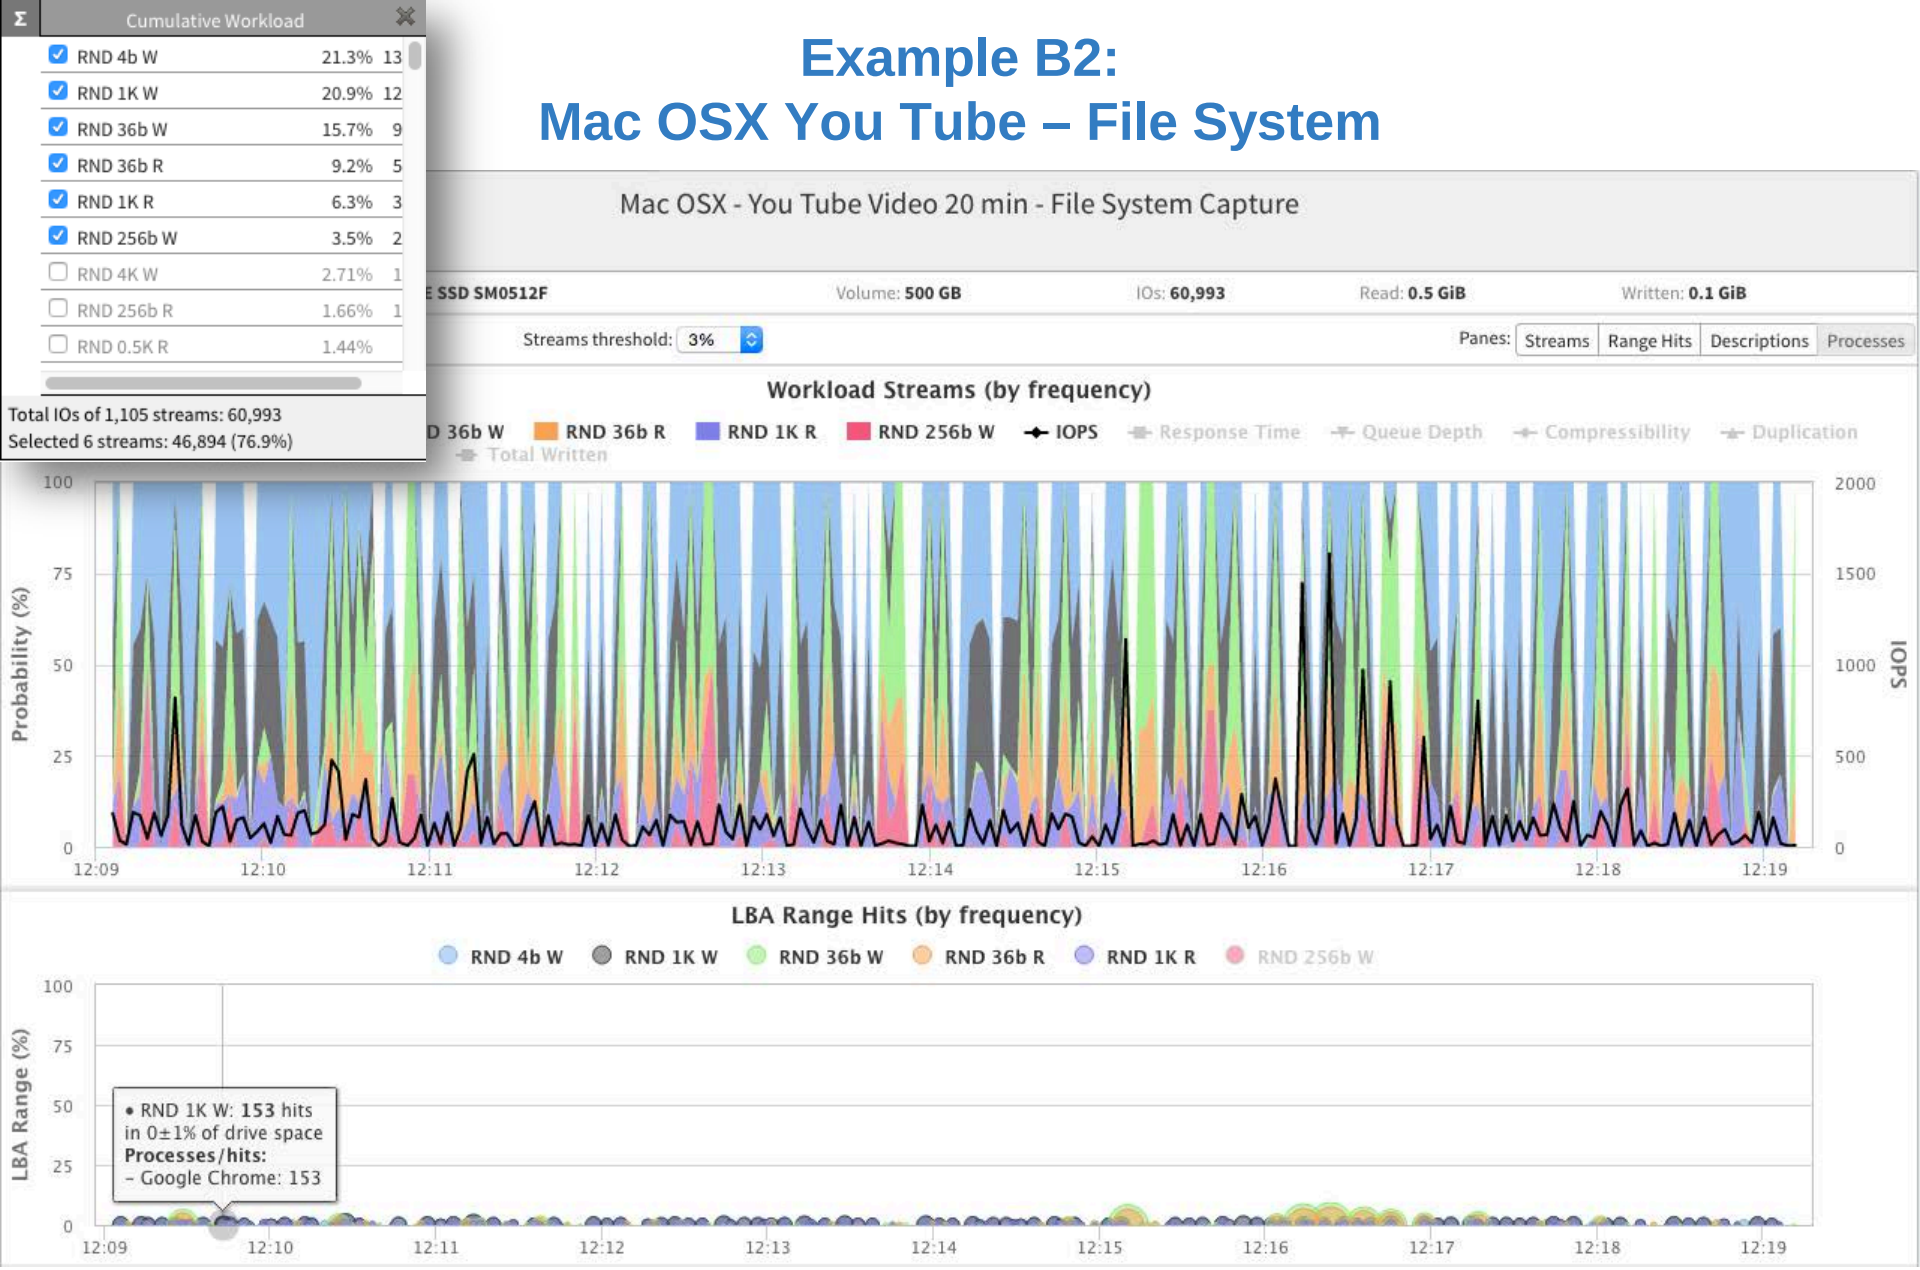

Real World Workload Capture

EXAMPLE C

Windows 8 Home Computer 20 min activity

Example C:
Home PC – Block IO 10 min Capture

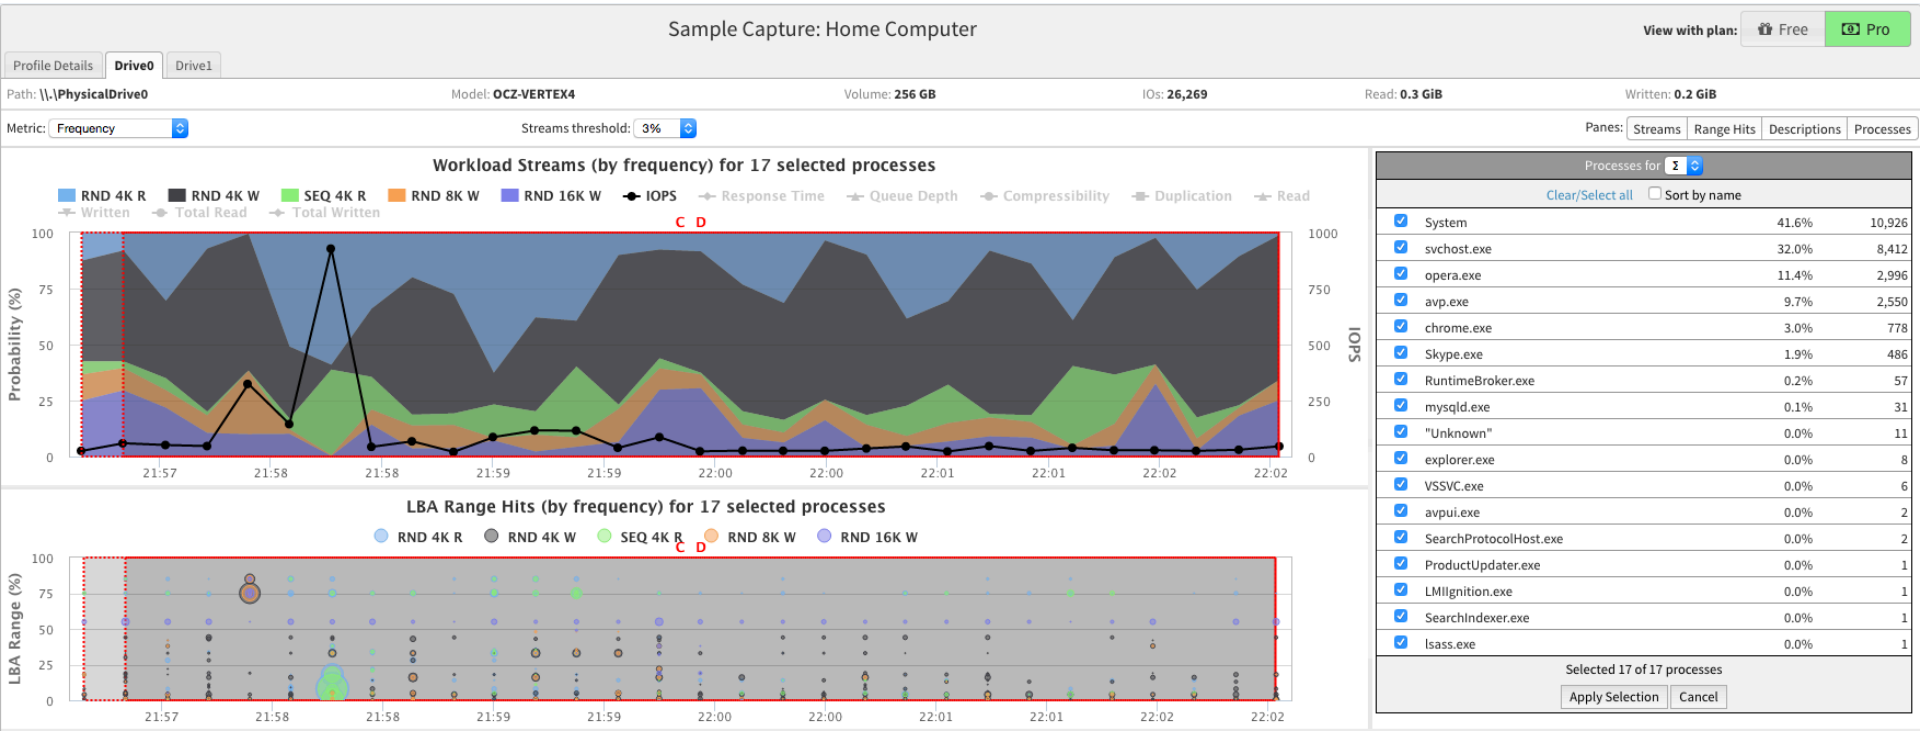

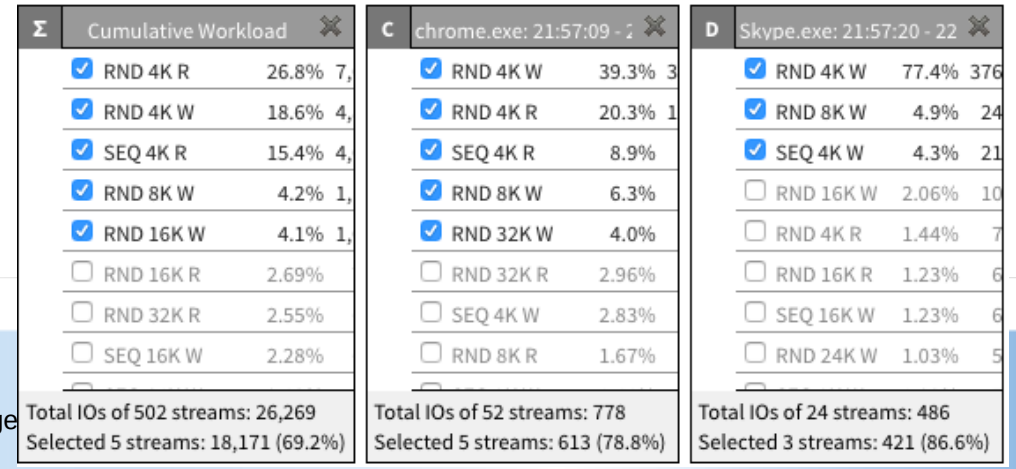

Real World Workload Capture

EXAMPLE D

Court Room Video Surveillance Cameras

Example D:
Court Room Video Surveillance – 10 hr Capture

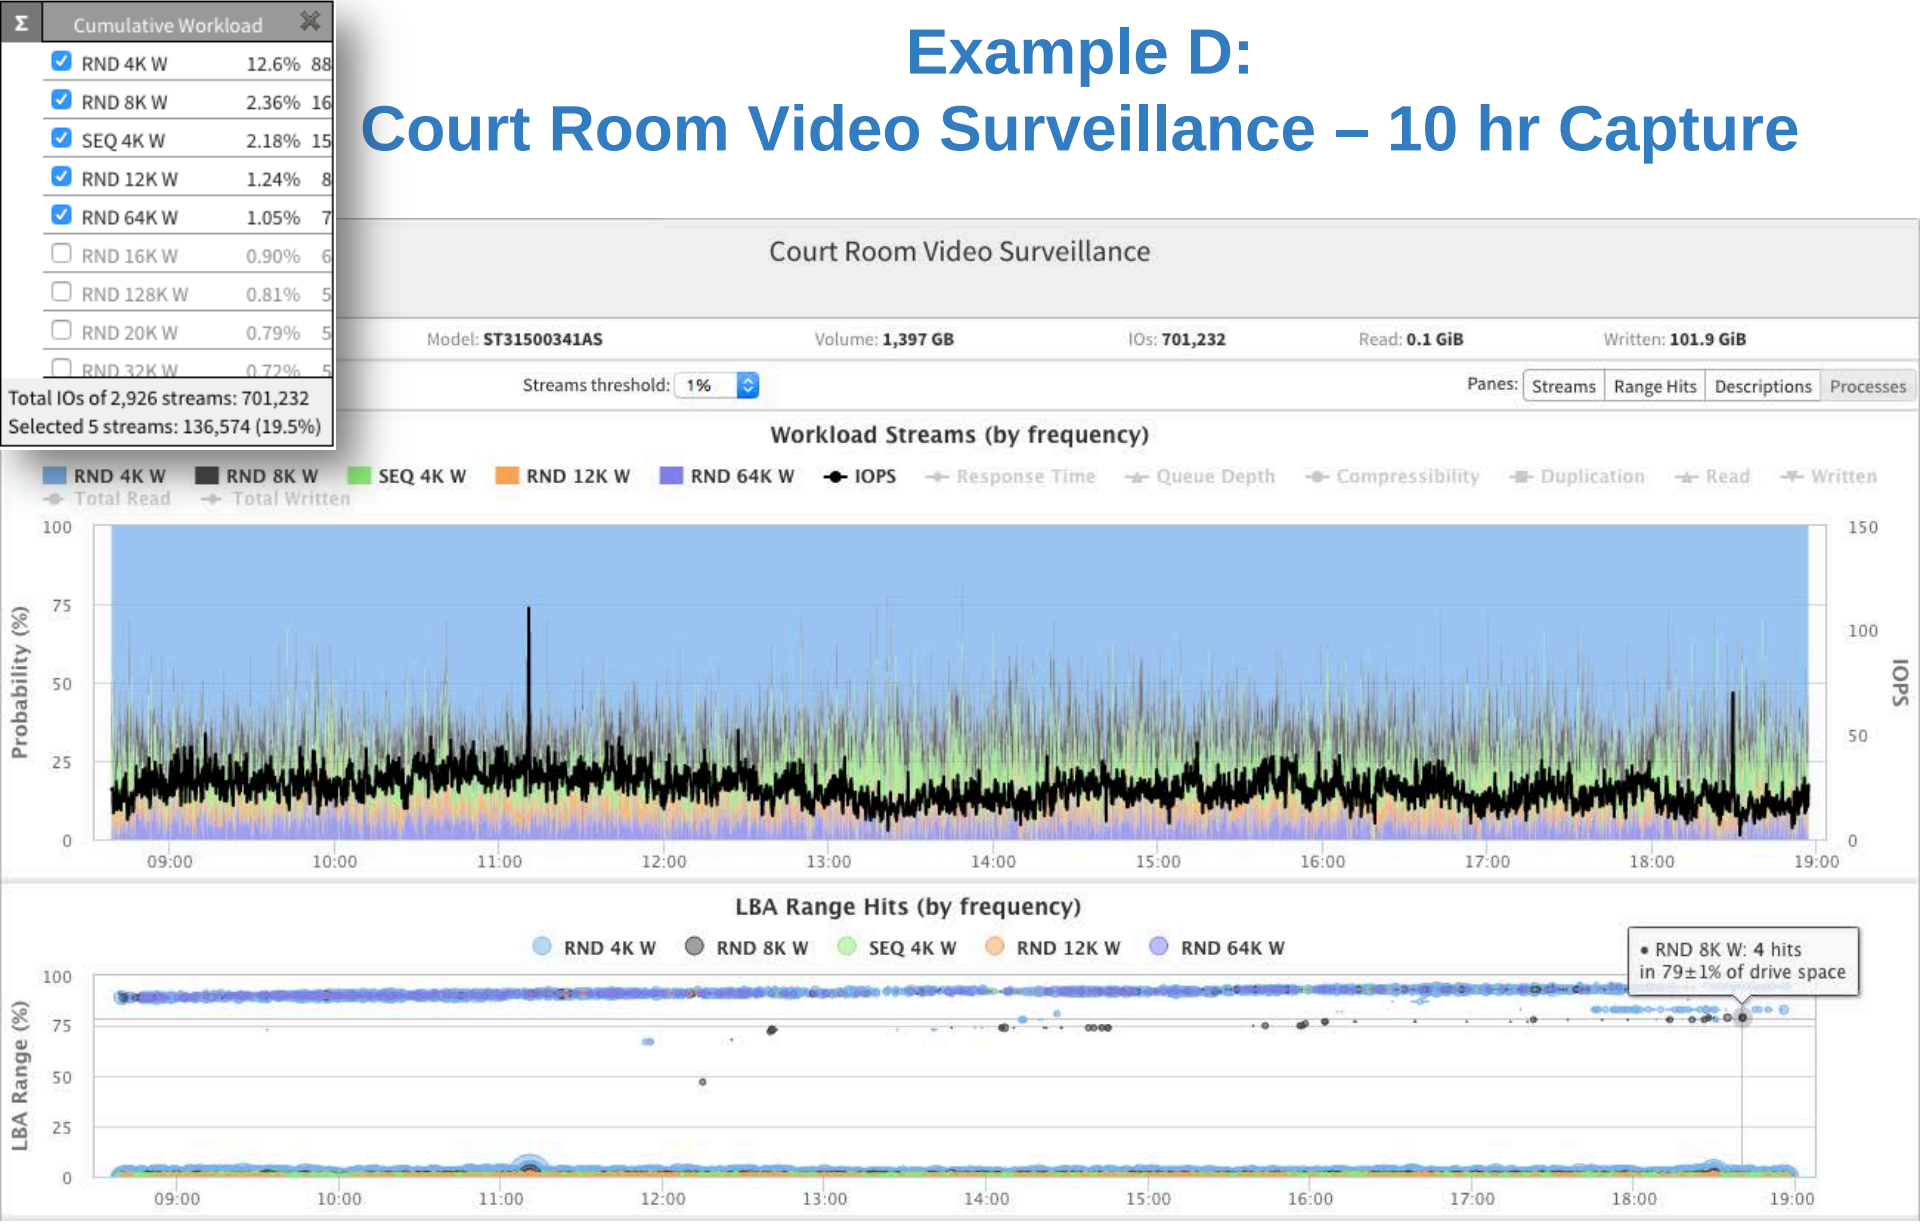

Real World Workload Capture

EXAMPLE E

PC Anti Virus Scan of HDD

Example E:
Anti Virus Scan – PC HDD

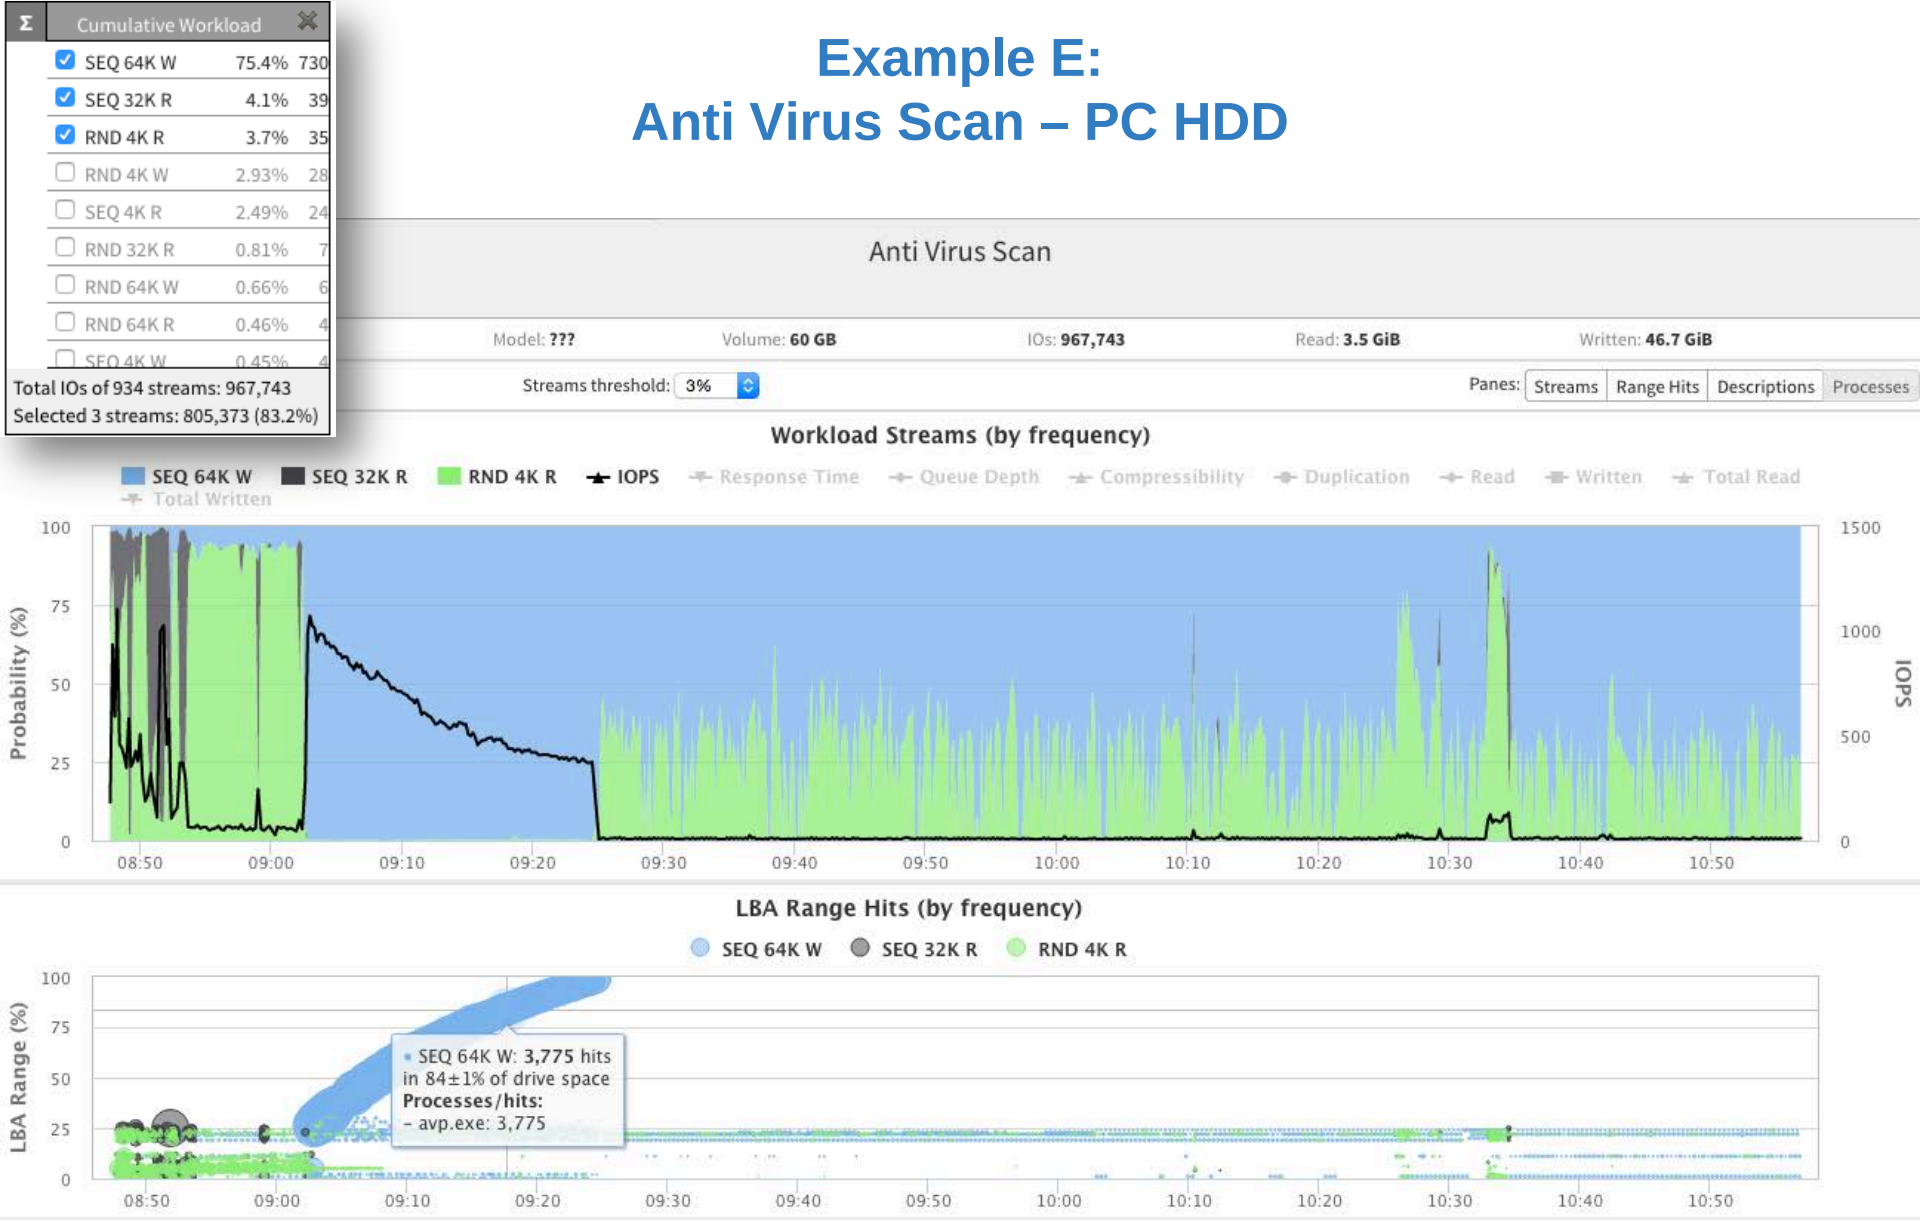

Take-Aways

- SSD Performance Depends on the SSD Workload
- IO Streams Change as they Traverse the SW Stack
- IOProfiler Captures IO Streams at the Block IO level
- See What IOs Actually get to the SSD Storage
- Be Sure you Buy the 'Right Amount of Performance'

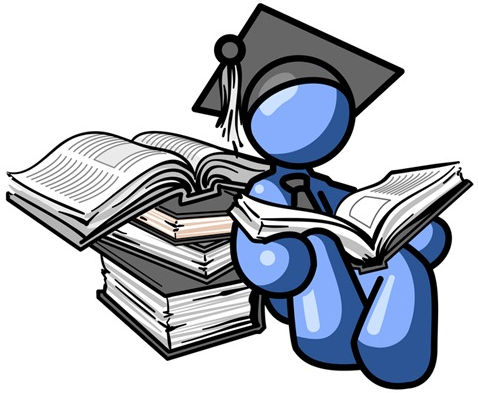

Understand Your SSD Workloads

To Analyze the Workloads Presented, Go to

TestMyWorkload.com

Capture & Analyze Your SSD

Real World Workloads Today!

For more information, contact Calypso Systems, Inc.

info@calypsotesters.com www.calypsotesters.com

SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE