SDC
STORAGE DEVELOPER CONFERENCE
SNIA SANTA CLARA, 2016

An Examination of User Workloads

Real World Storage Workload Capture

Eden Kim
Calypso Systems, Inc.

An Examination of User Workloads
Real World Storage Workload Capture

A. What Are Real World Storage Workloads
& Why Do We Care?

B. What Do Real World Storage Workloads Look Like?

C. Case Study: Retail Store Web Portal 24 hour Capture

D. Take Aways

E. Q & A

SDC 16

What Are Real World Storage Workloads and Why Do We Care?

Real World Storage Workloads are Unique and are the:

IOs that occur when *YOUR* application is running

IO Streams that traverse *YOUR* Hardware/Software Stack

IO Streams that present to *YOUR* HDD/SSD/Array Storage

SDC 16

IO Streams are YOUR Data Traversing the HW/SW Stack

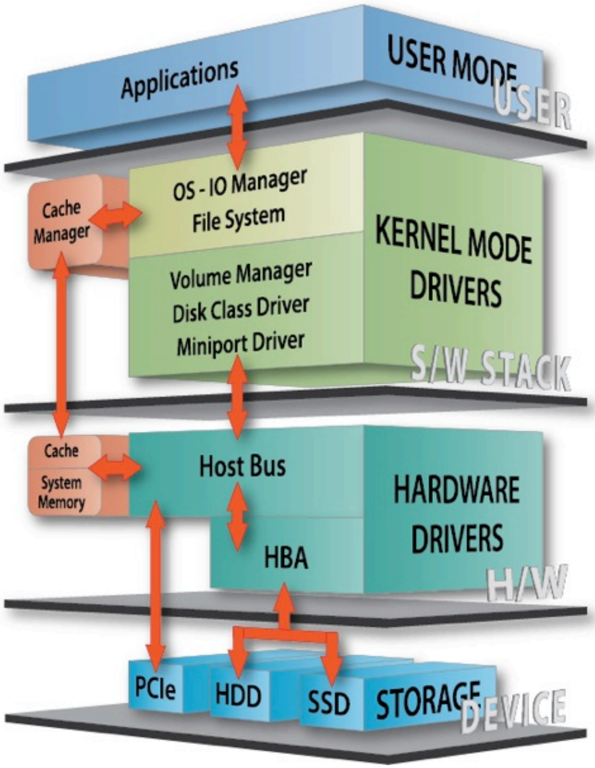

Many Different IO Streams each with a Unique:

Random or Sequential Access, Data Transfer Size
Read/Write Mix, Entropy and
Spatial & Temporal Locality of Reference

IO Streams Change at each level of SW Abstraction

Fragmented, Coalesced, Appended due to
Virtualization, Packetization, Encryption, Data Reduction
Compression, Deduplication, Storage Tiering and more

Because IO Streams Change, it is important to Capture IOs as close to the storage as possible – ideally at the Block IO level

Why Do We Care?

SSD Performance & Endurance Depends on it:

SSD Performance Depends on the Type of Workload

Workloads Affect Endurance (Write Amplification)

Storage Tiering Strategies are Based on Workload Assumptions

Workloads Determine What Type of SSD is Best for Your Application

Be Sure to Buy the Right Amount of Performance and Endurance!

What Do Real World Storage Workloads Look Like?

IO Streams Can Be Visualized As:

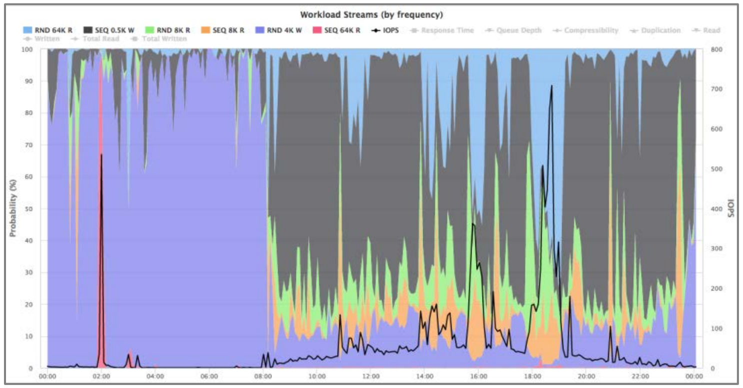

IO Stream Maps by Frequency over Time (IO Stream Map)

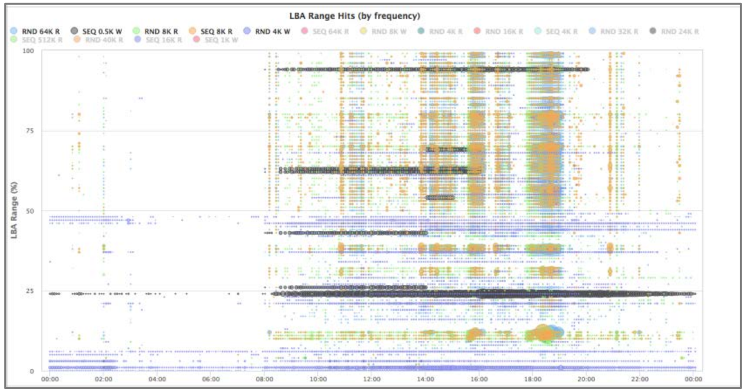

IO Stream Maps by Location over Time (LBA Hit Map)

A Tabular Distribution of IO Streams

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
☐ RND 32K R	1.63%	70,407

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

SDC 16

Key IO Stream Metrics Include:

IO Stream Map Frequency & Amount

LBA Hit Maps – Spatial & Temporal Locality of Reference

Specific Process ID for IO Streams

IO Stream Sequentiality, Queue Depth, Response Times

IOPS rate, Bandwidth, Access Patterns, Data Transfer Sizes

Reads, Writes, Amount Written

Compressibility Ratio

Deduplication Ratio

Workload Visualization – Example #3 at TestMyWorkload.com

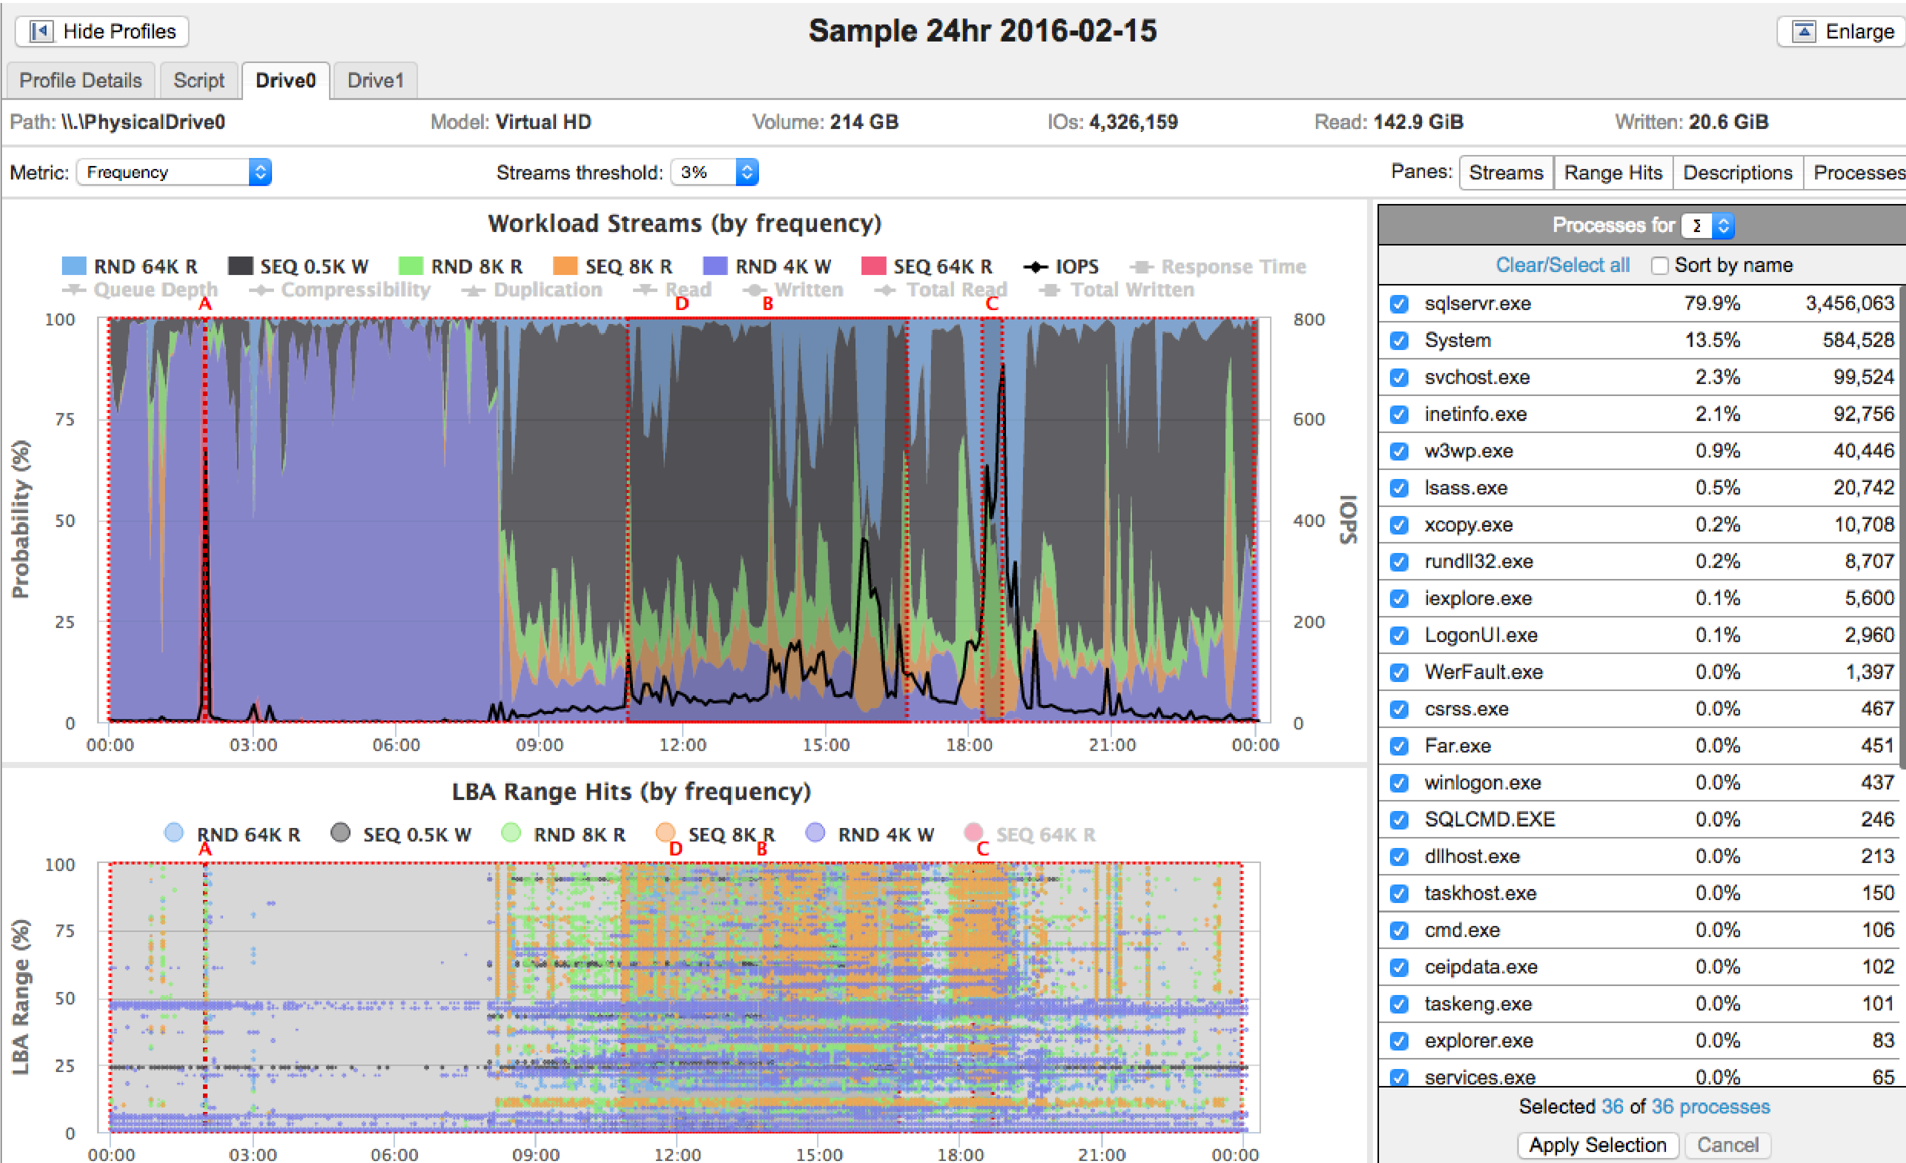

SDC 16

Case Study

Test Plan

- ✓ Create Workload Segments from IO Capture Data
- ✓ Test using the same OS and Software as Captured Data
- ✓ Apply Workloads to 3 Data Center SATA SSDs
- ✓ Compare Performance to Workload Segments
- ✓ Compare Performance to Workload Replay

Test Set Up

Hardware Platform

Calypso IOProfiler – Real World Workload Tester

OS: Hyper-V Server 2012 R2 (same as source capture)

Test Software: RTP BE ver 1.9.184

Motherboard: Intel SC2600COE, 32GB DDR3 ECC RAM

CPUs: Dual Intel XEON Eight Core W2687v2 3.1Ghz

HBA: 6Gb/s LSI 9212

Software Platform

Test Software: IPF 1.05 FE ver 1.18.11

OS: Windows 7 Pro

Capture Tool: IPF Win Capture applets

SDC 16

SSD Sample Pool

SSD A - Data Center 2.5” SATA SSD – 960 GB

SSD B - Data Center 2.5” SATA SSD – 800 GB

SSD C - Data Center 2.5” SATA SSD – 960 GB


RESULTS

TestMyWorkload IO Capture Demo #3 - Workload Segment Definition

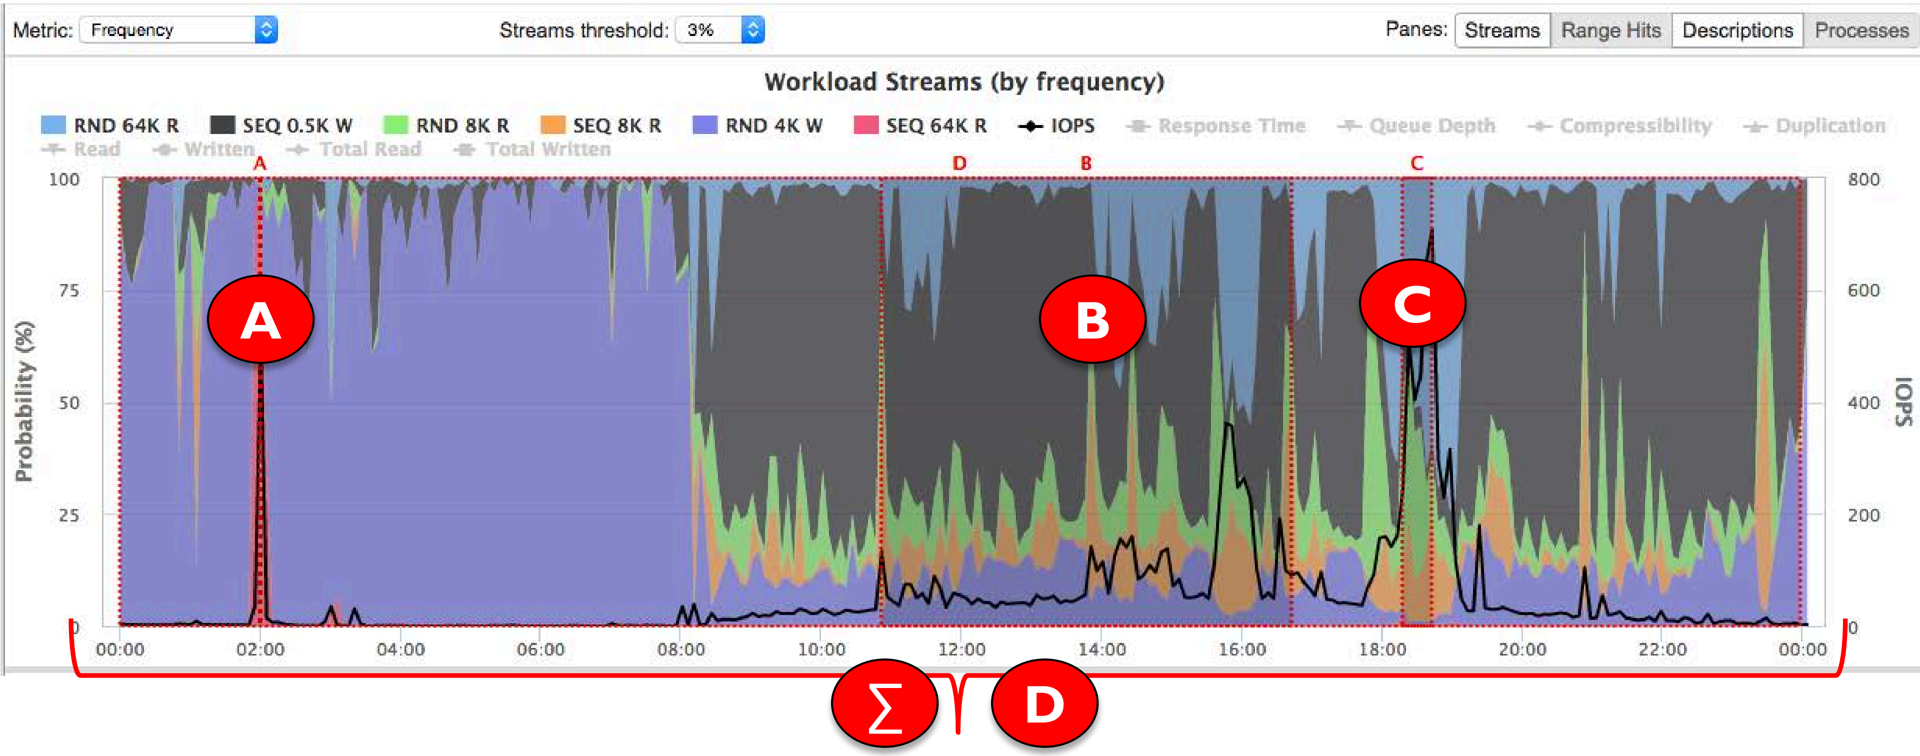

24 Hour Cumulative	2 am Back up	10 am - 4 pm Operations	6 - 7 pm Evening	24 Hour SQL
∑	A	B	C	D

∑	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 (64.4%)

A	sqlservr.exe: 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

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

B	10:51:29 - 16:42:16	
☑ 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

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

C	18:17:45 - 18:42:48	
☑ 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

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

D	sqlservr.exe: 00:00:04 - 2	
☑ RND 64K R	23.7%	818,951
☑ SEQ 0.5K W	22.4%	774,281
☑ 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%	118,749
☐ SEQ 512K R	1.73%	59,917
☐ RND 24K R	1.60%	55,312

Total IOs of 395 streams: 3,455,940
Selected 6 streams: 2,682,490 (77.6%)

Workload Segments - Net IO Streams Distribution % by Segments

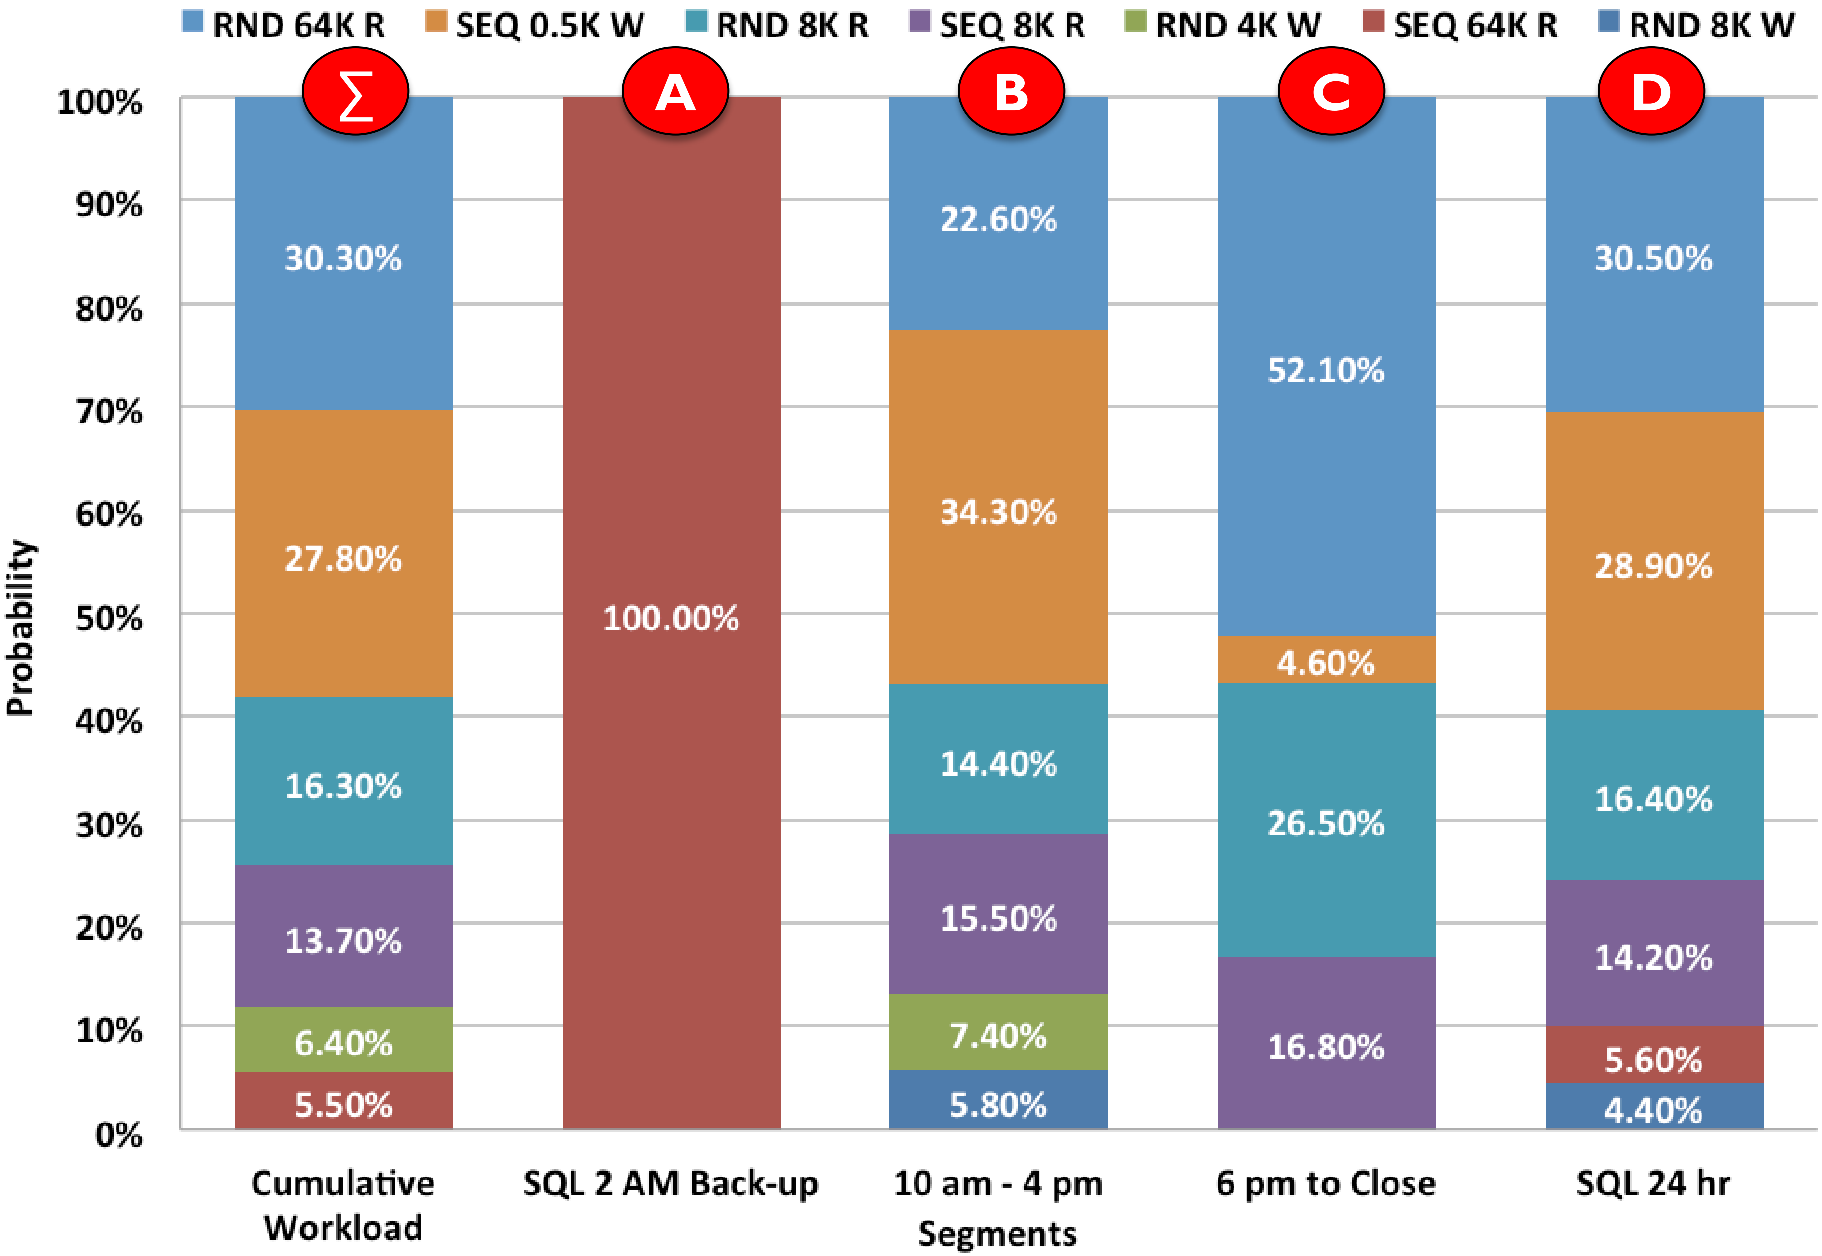

Replay - Throughput by Segments: Average MB/s Over Segment

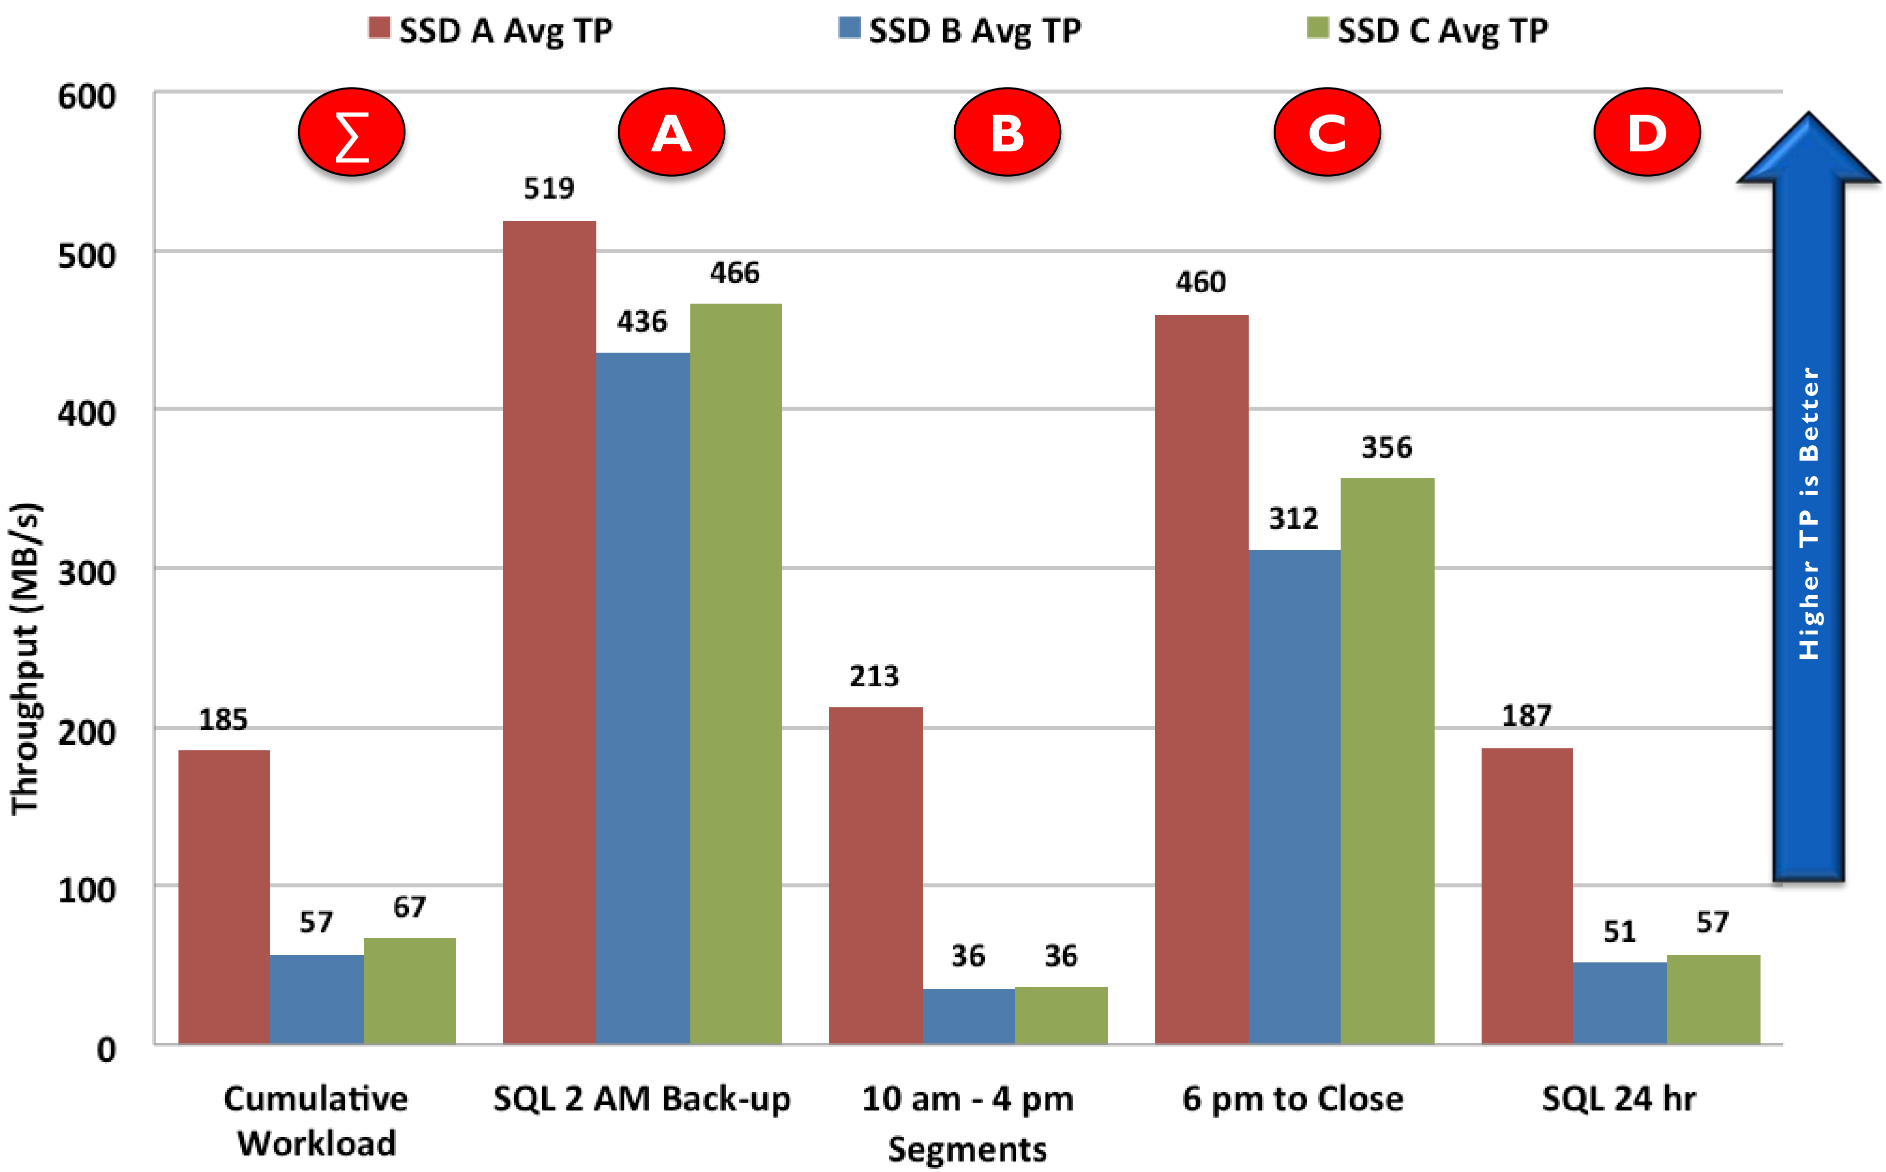

SDC 16

Replay - ART by Segments: Average Over Segment

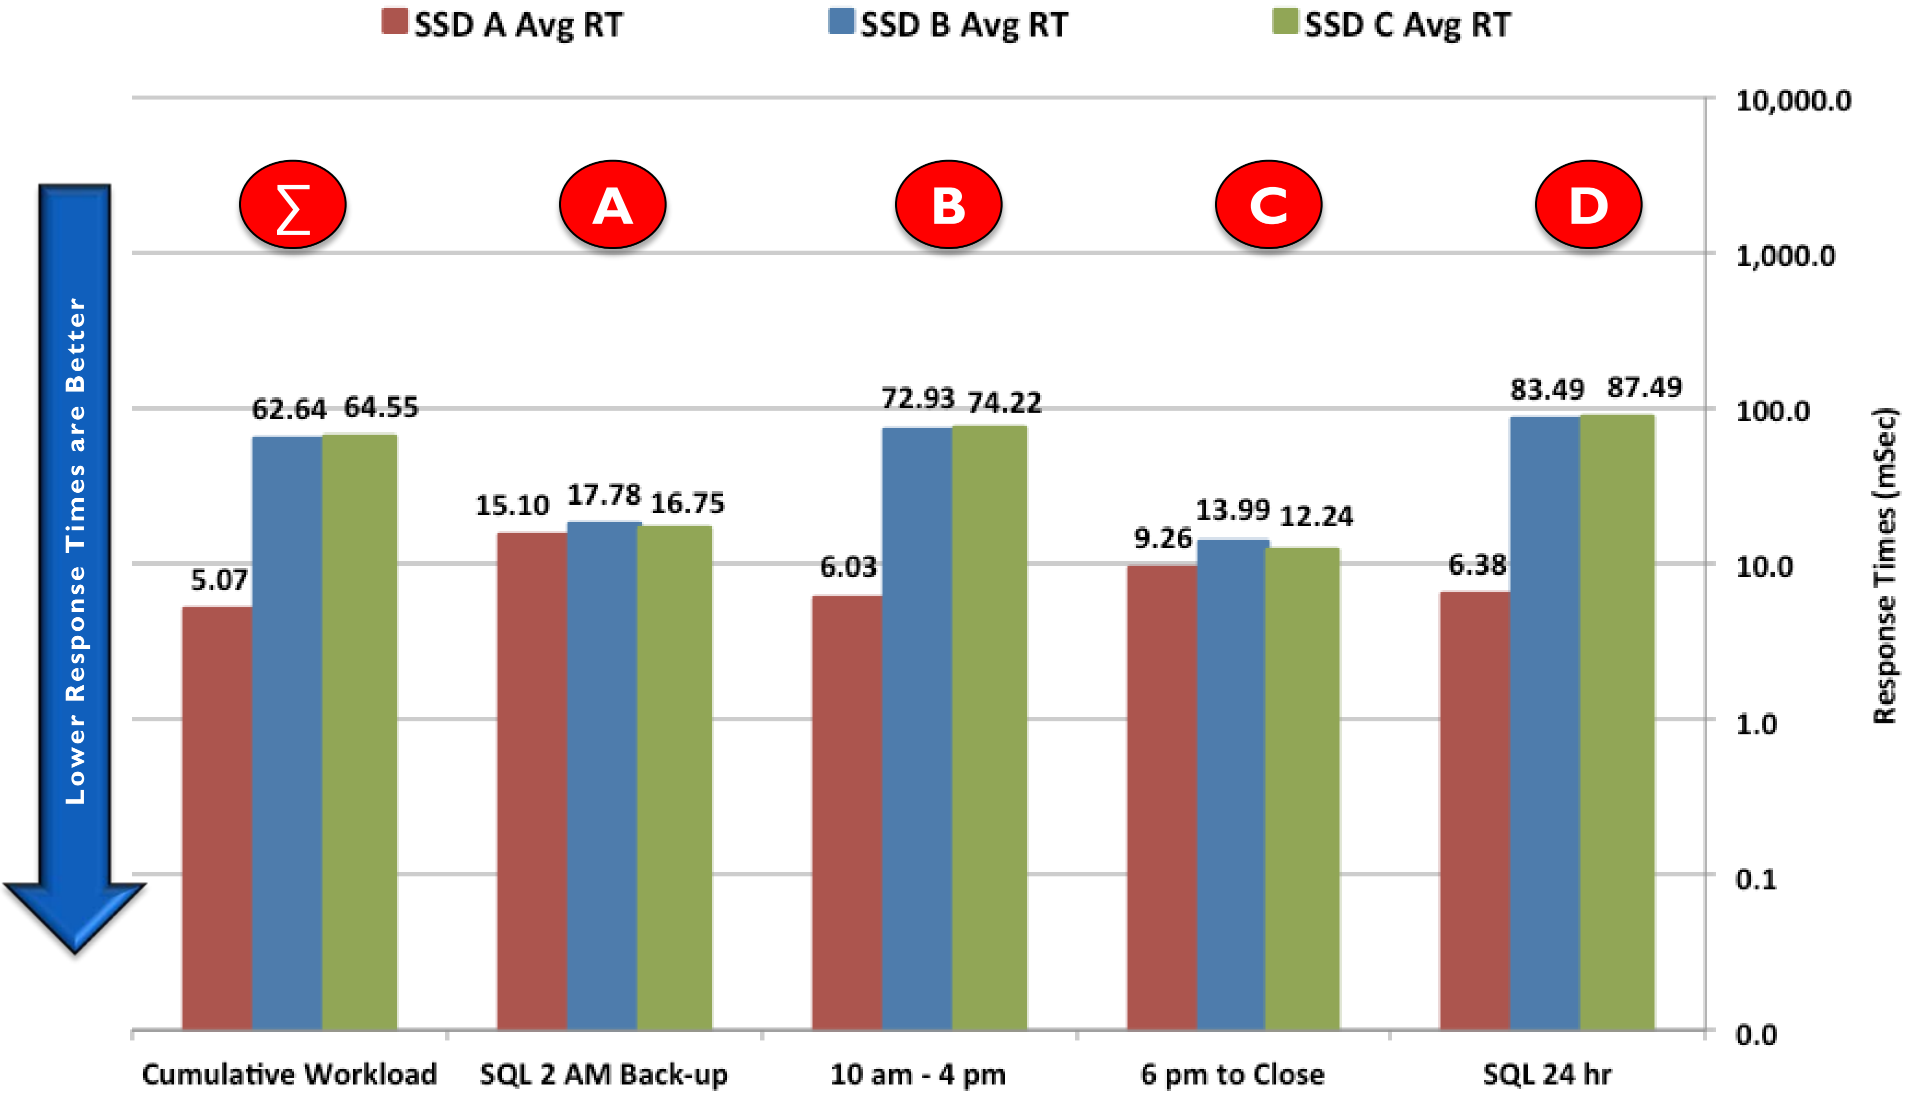

Replay - 5 9s Response Time by Segments: Average Over Segment

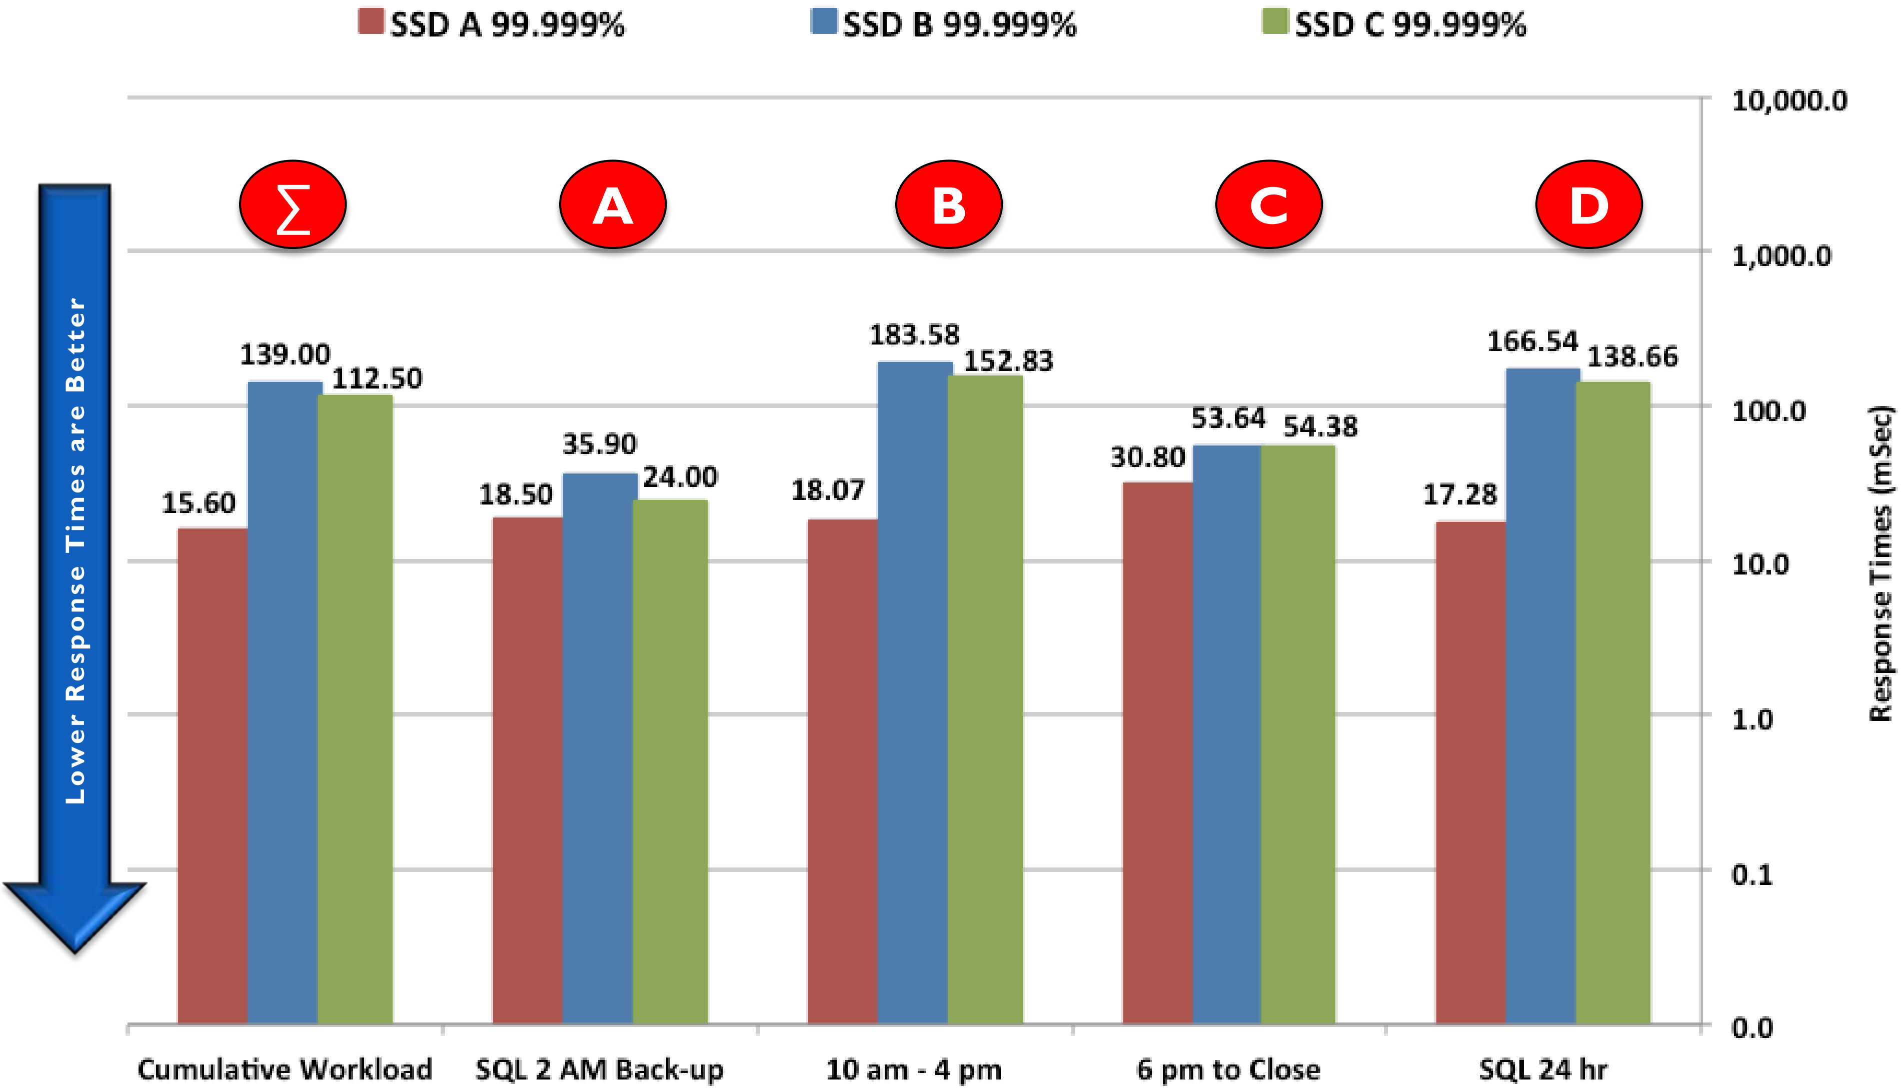

SDC 16

Power Consumption by Segments - Average Over Segment

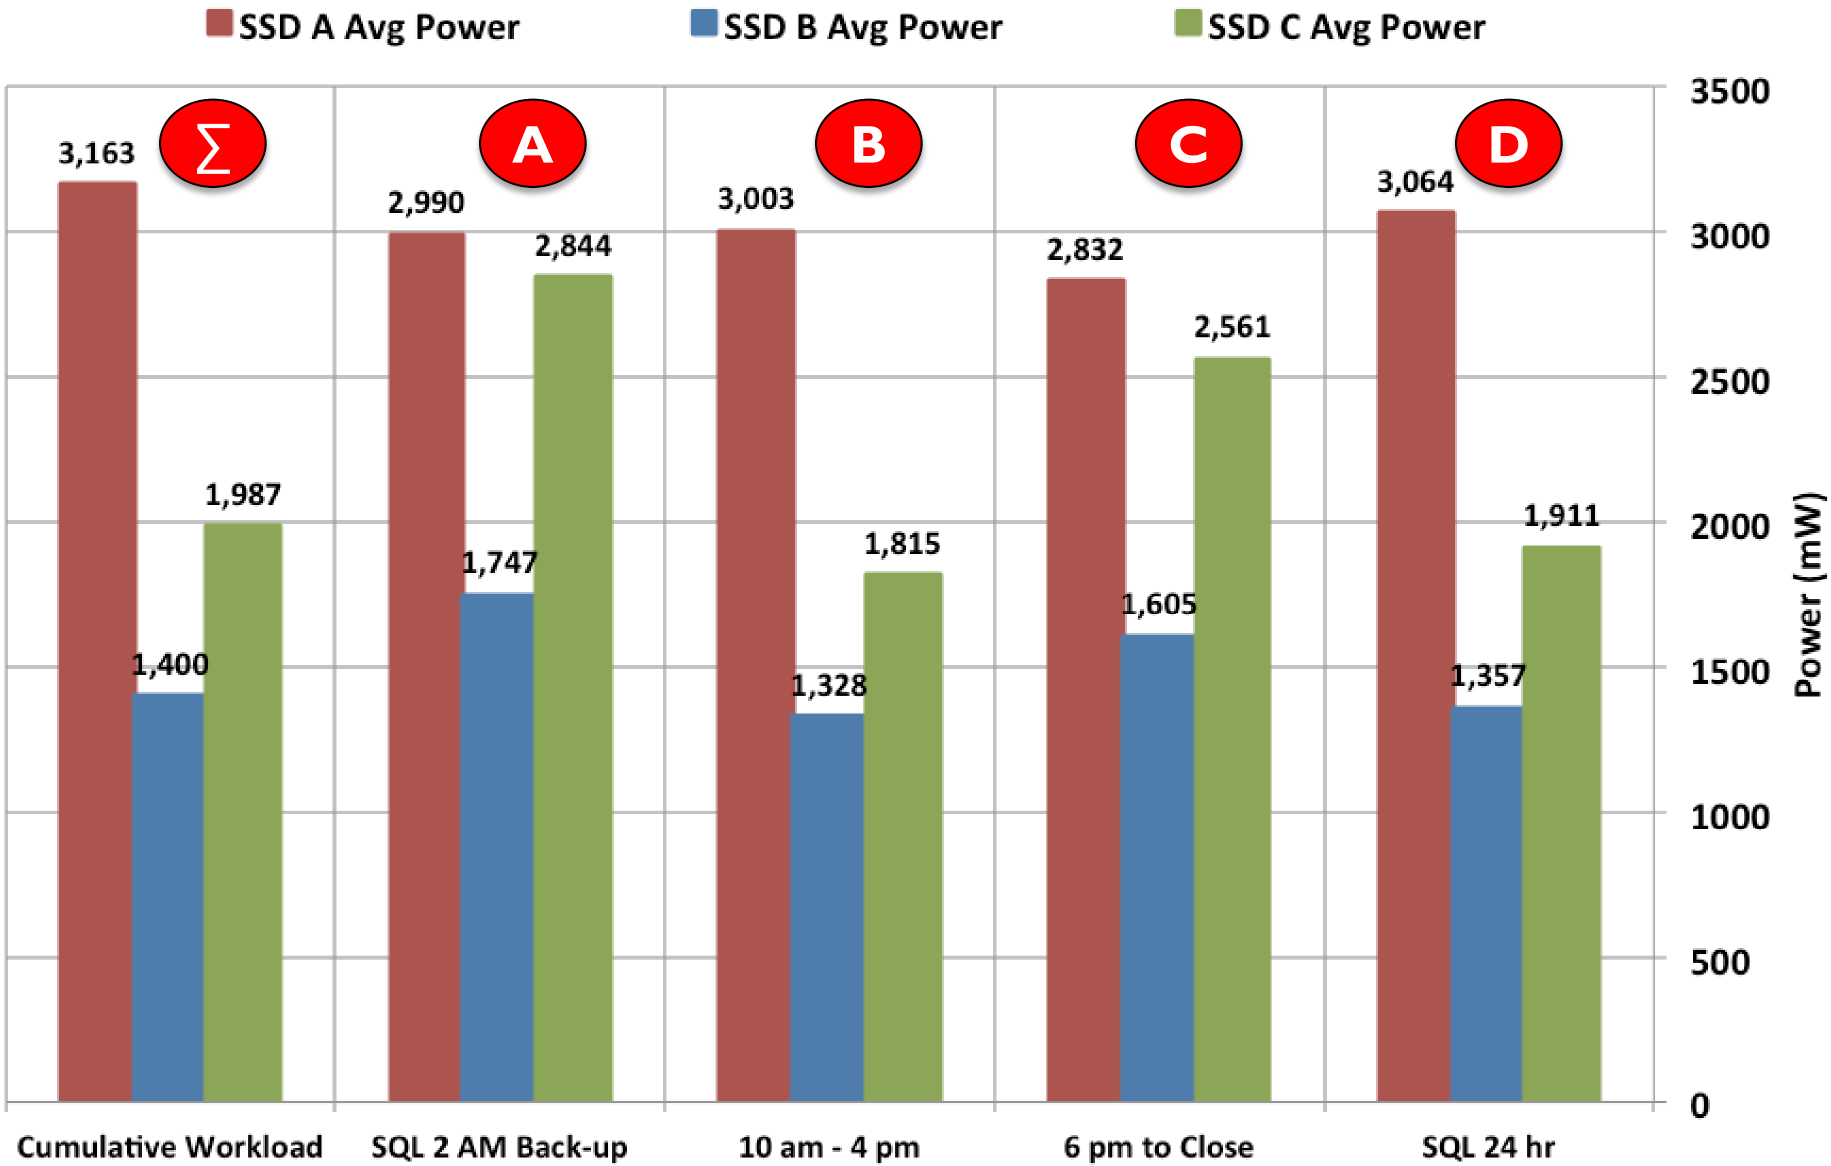

24 Hr Replay - Streaming IOPS

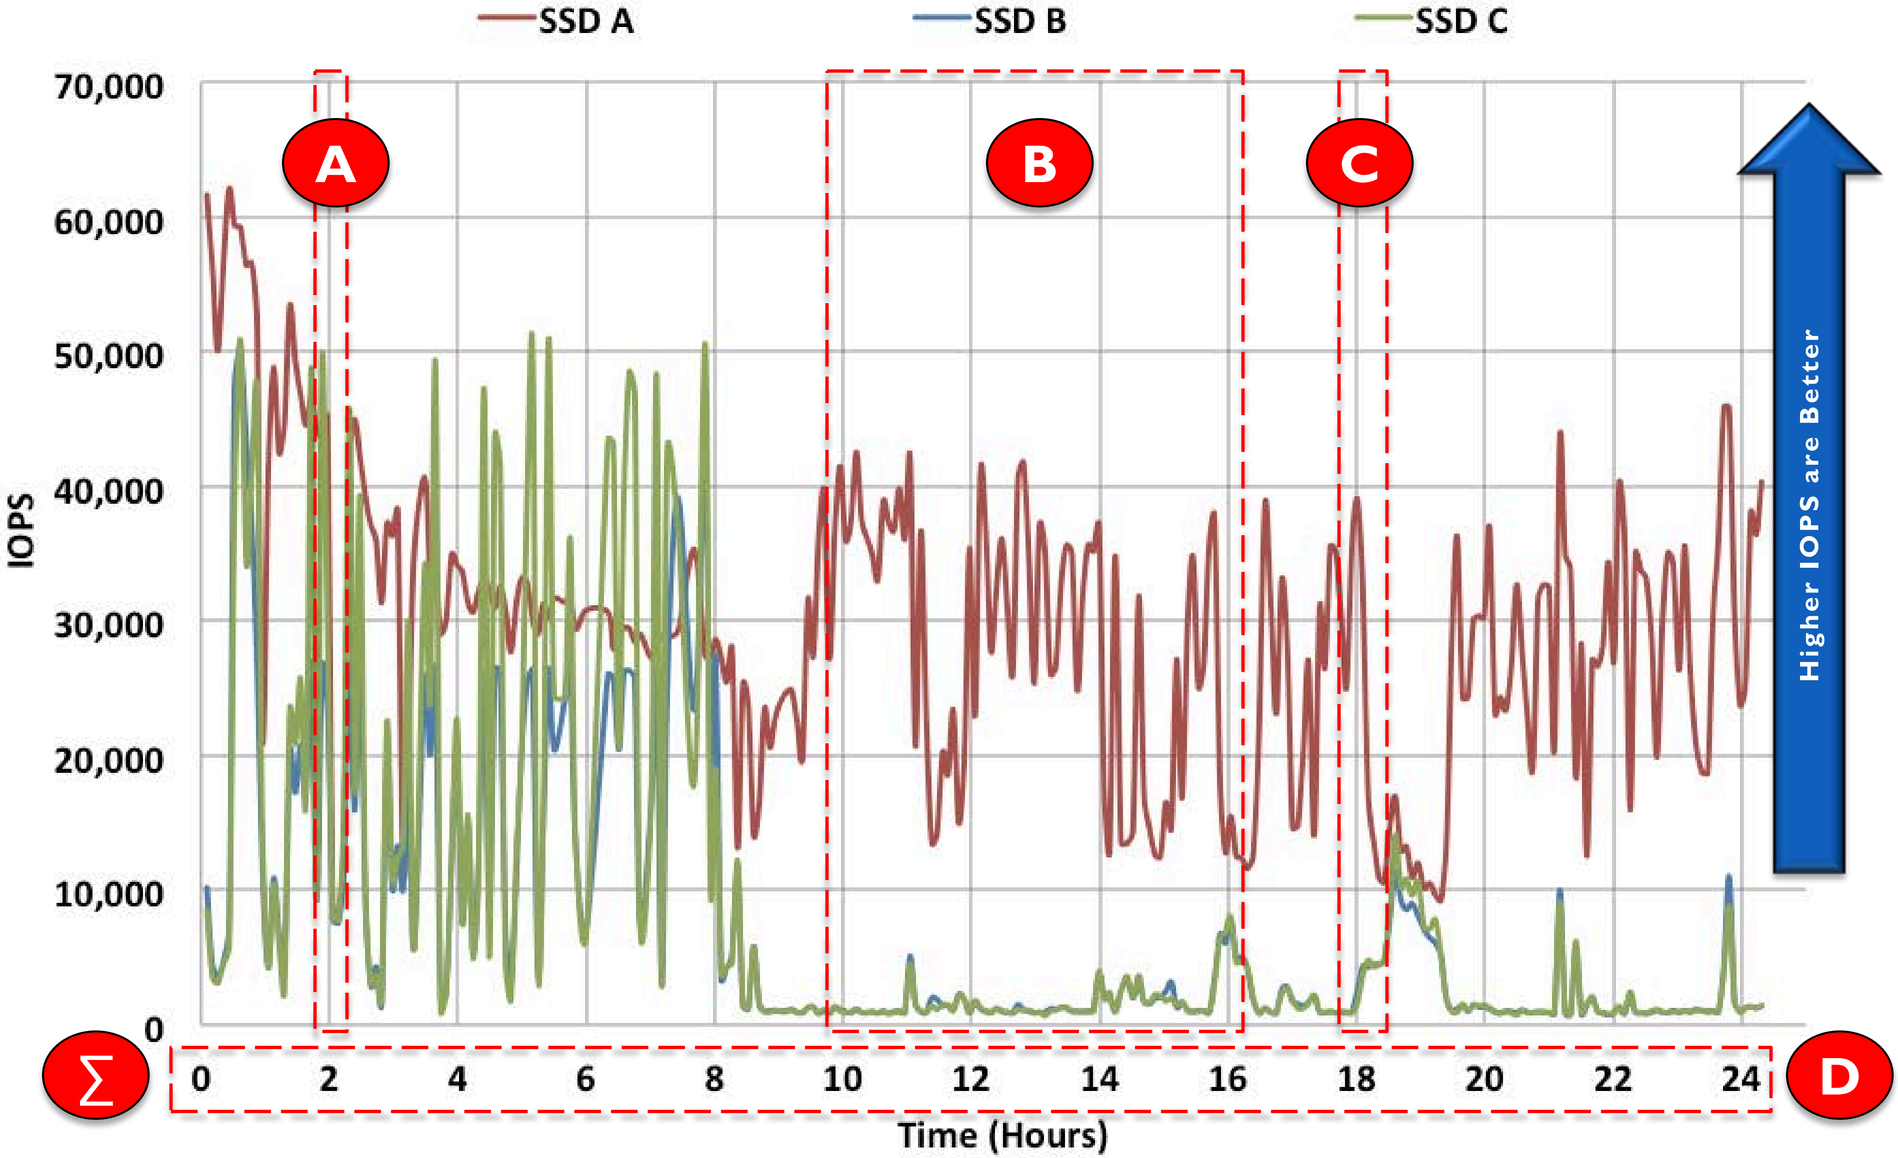

SDC 16

Replay - Throughput by Segments: Average MB/s Over Segment

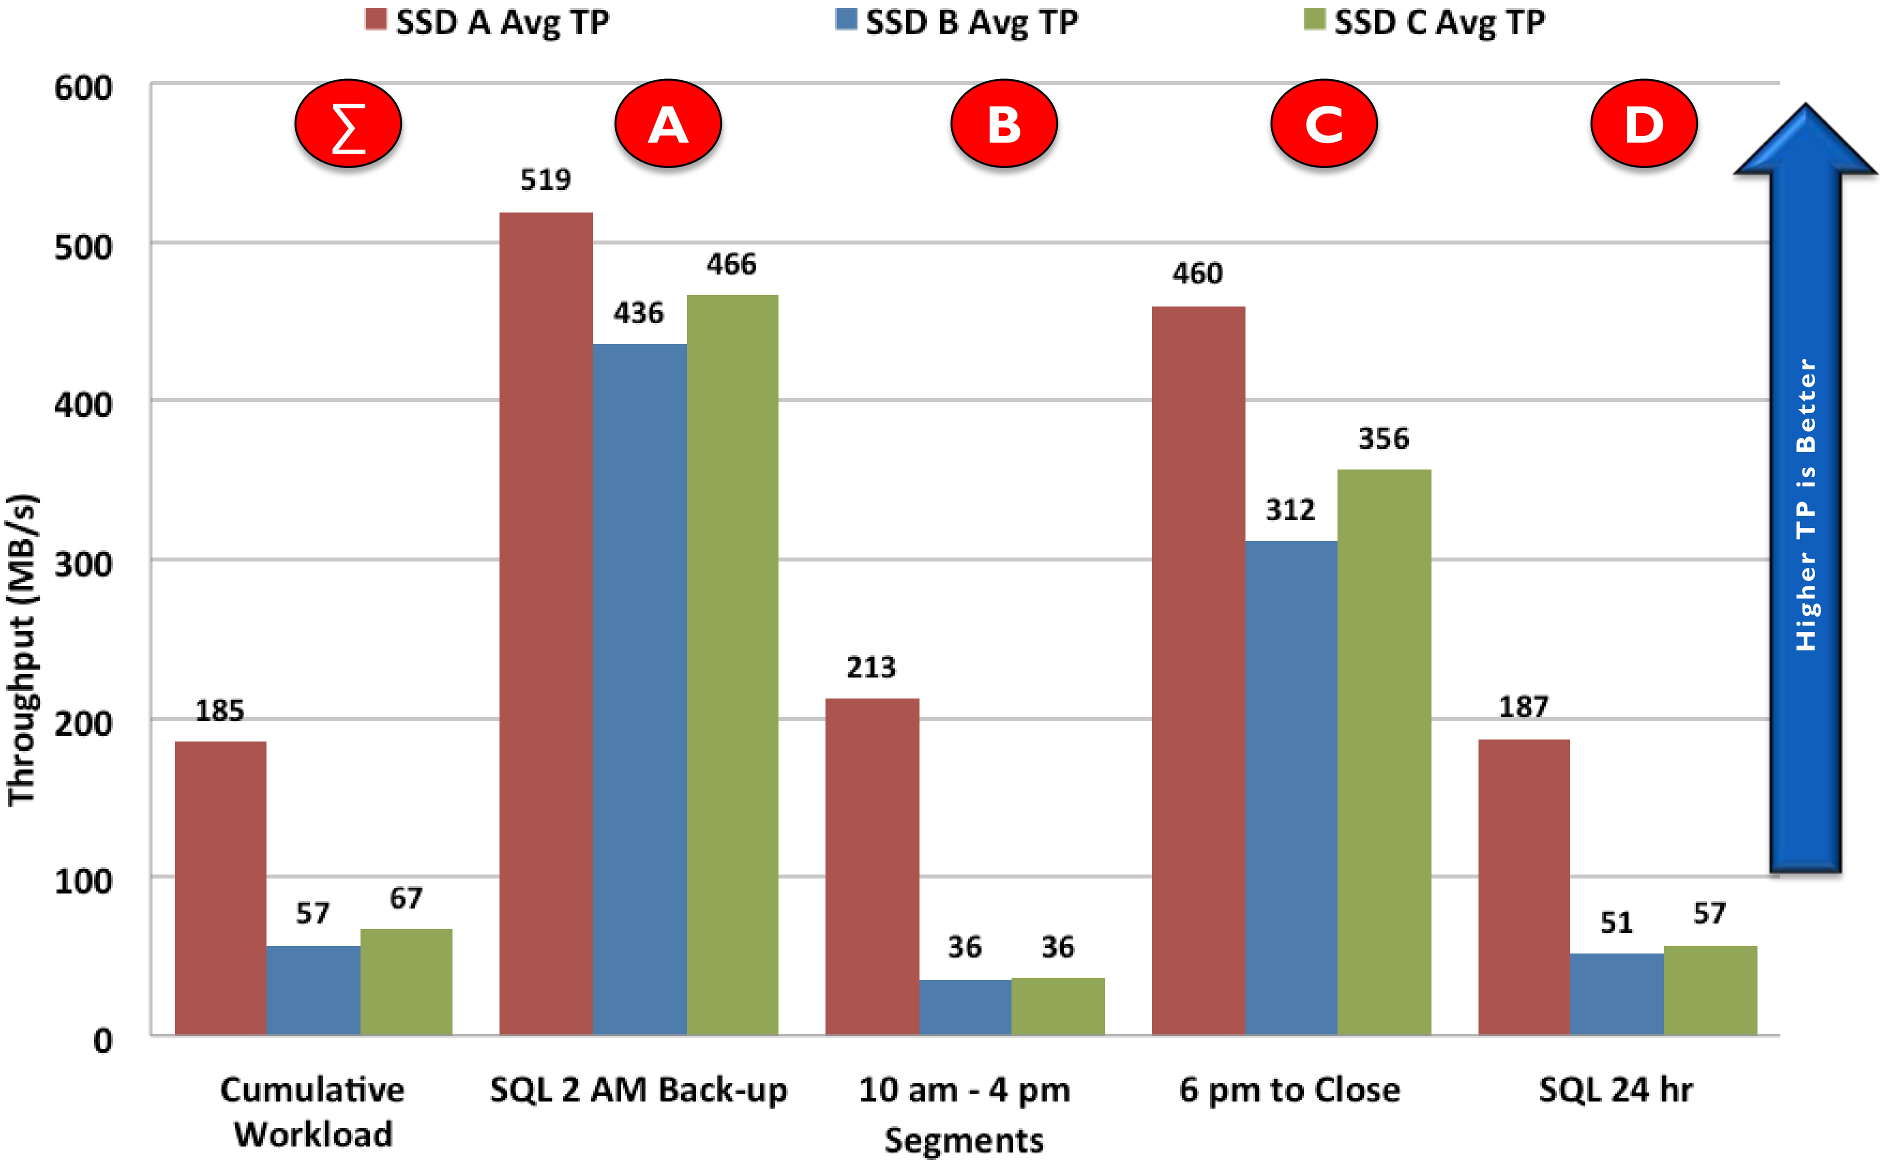

SDC 16

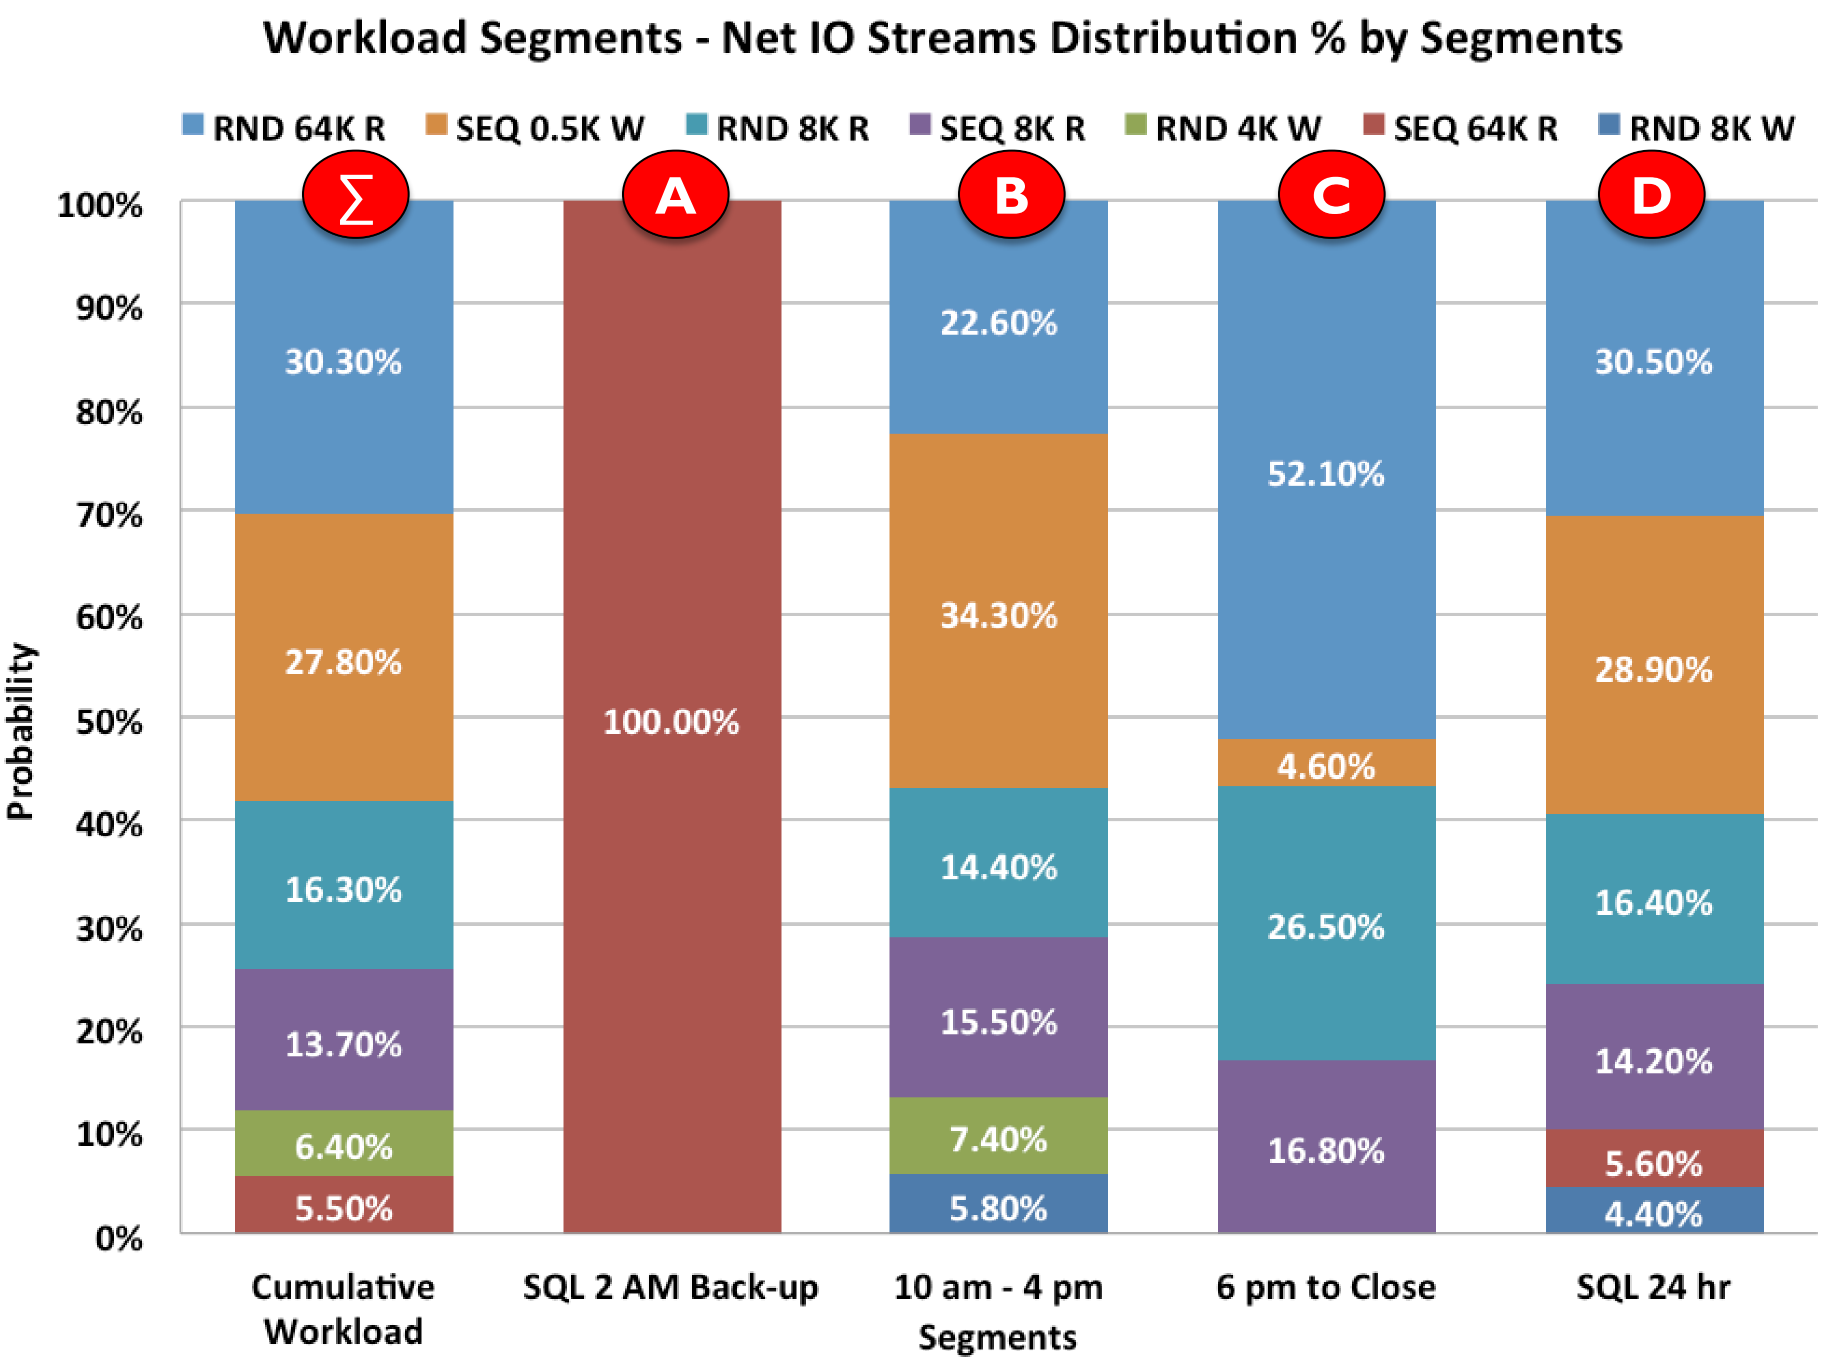

Workload Segments - Net IO Streams Distribution % by Segments
RND 64K R
SEQ 0.5K W
RND 8K R
SEQ 8K R
RND 4K W
SEQ 64K R
RND 8K W
∑
A
B
C
D
Probability
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
30.30%
27.80%
16.30%
13.70%
6.40%
5.50%
100.00%
22.60%
34.30%
14.40%
15.50%
7.40%
5.80%
52.10%
4.60%
26.50%
16.80%
30.50%
28.90%
16.40%
14.20%
5.60%
4.40%
Cumulative Workload
SQL 2 AM Back-up
10 am - 4 pm Segments
6 pm to Close
SQL 24 hr

6 Hr Write Saturation - Metadata SEQ 0.5K RW50 - IOPS vs Time

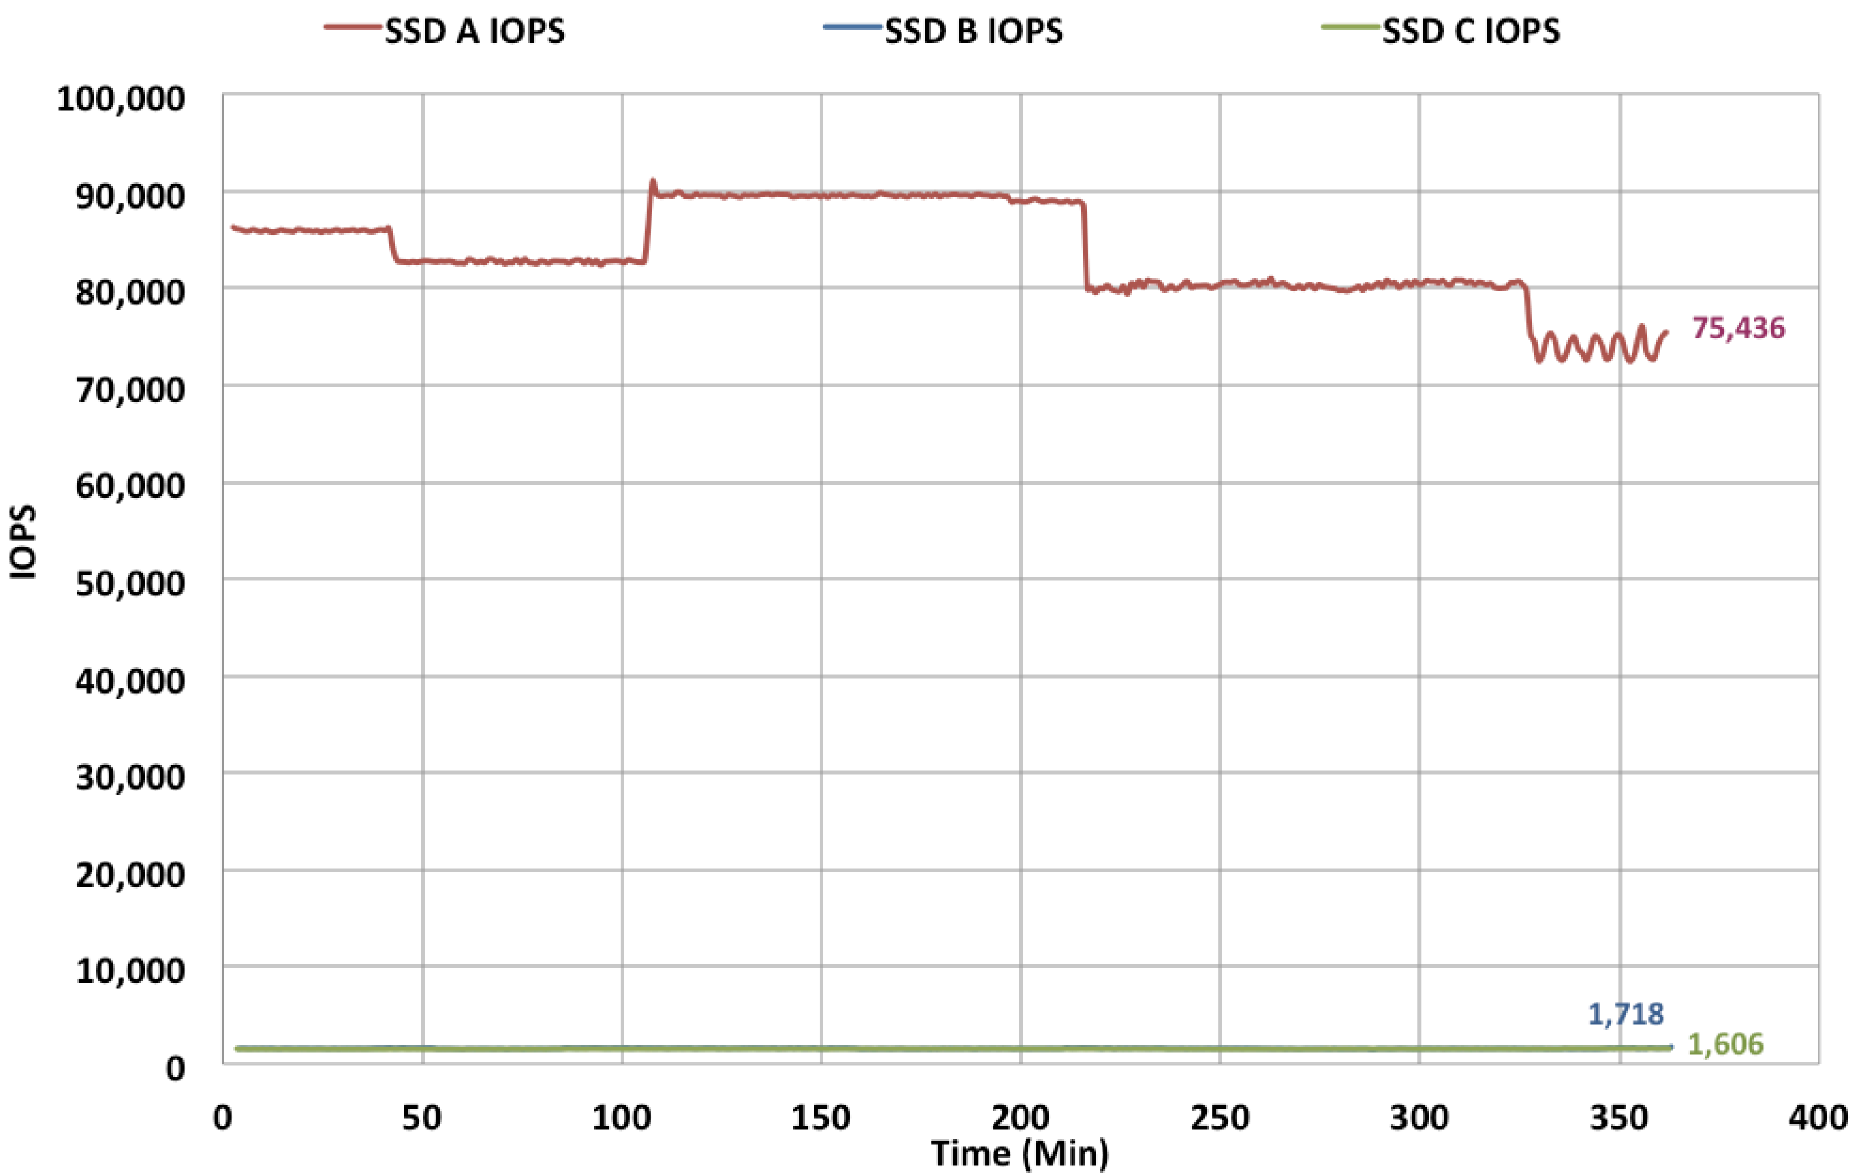

SDC 16

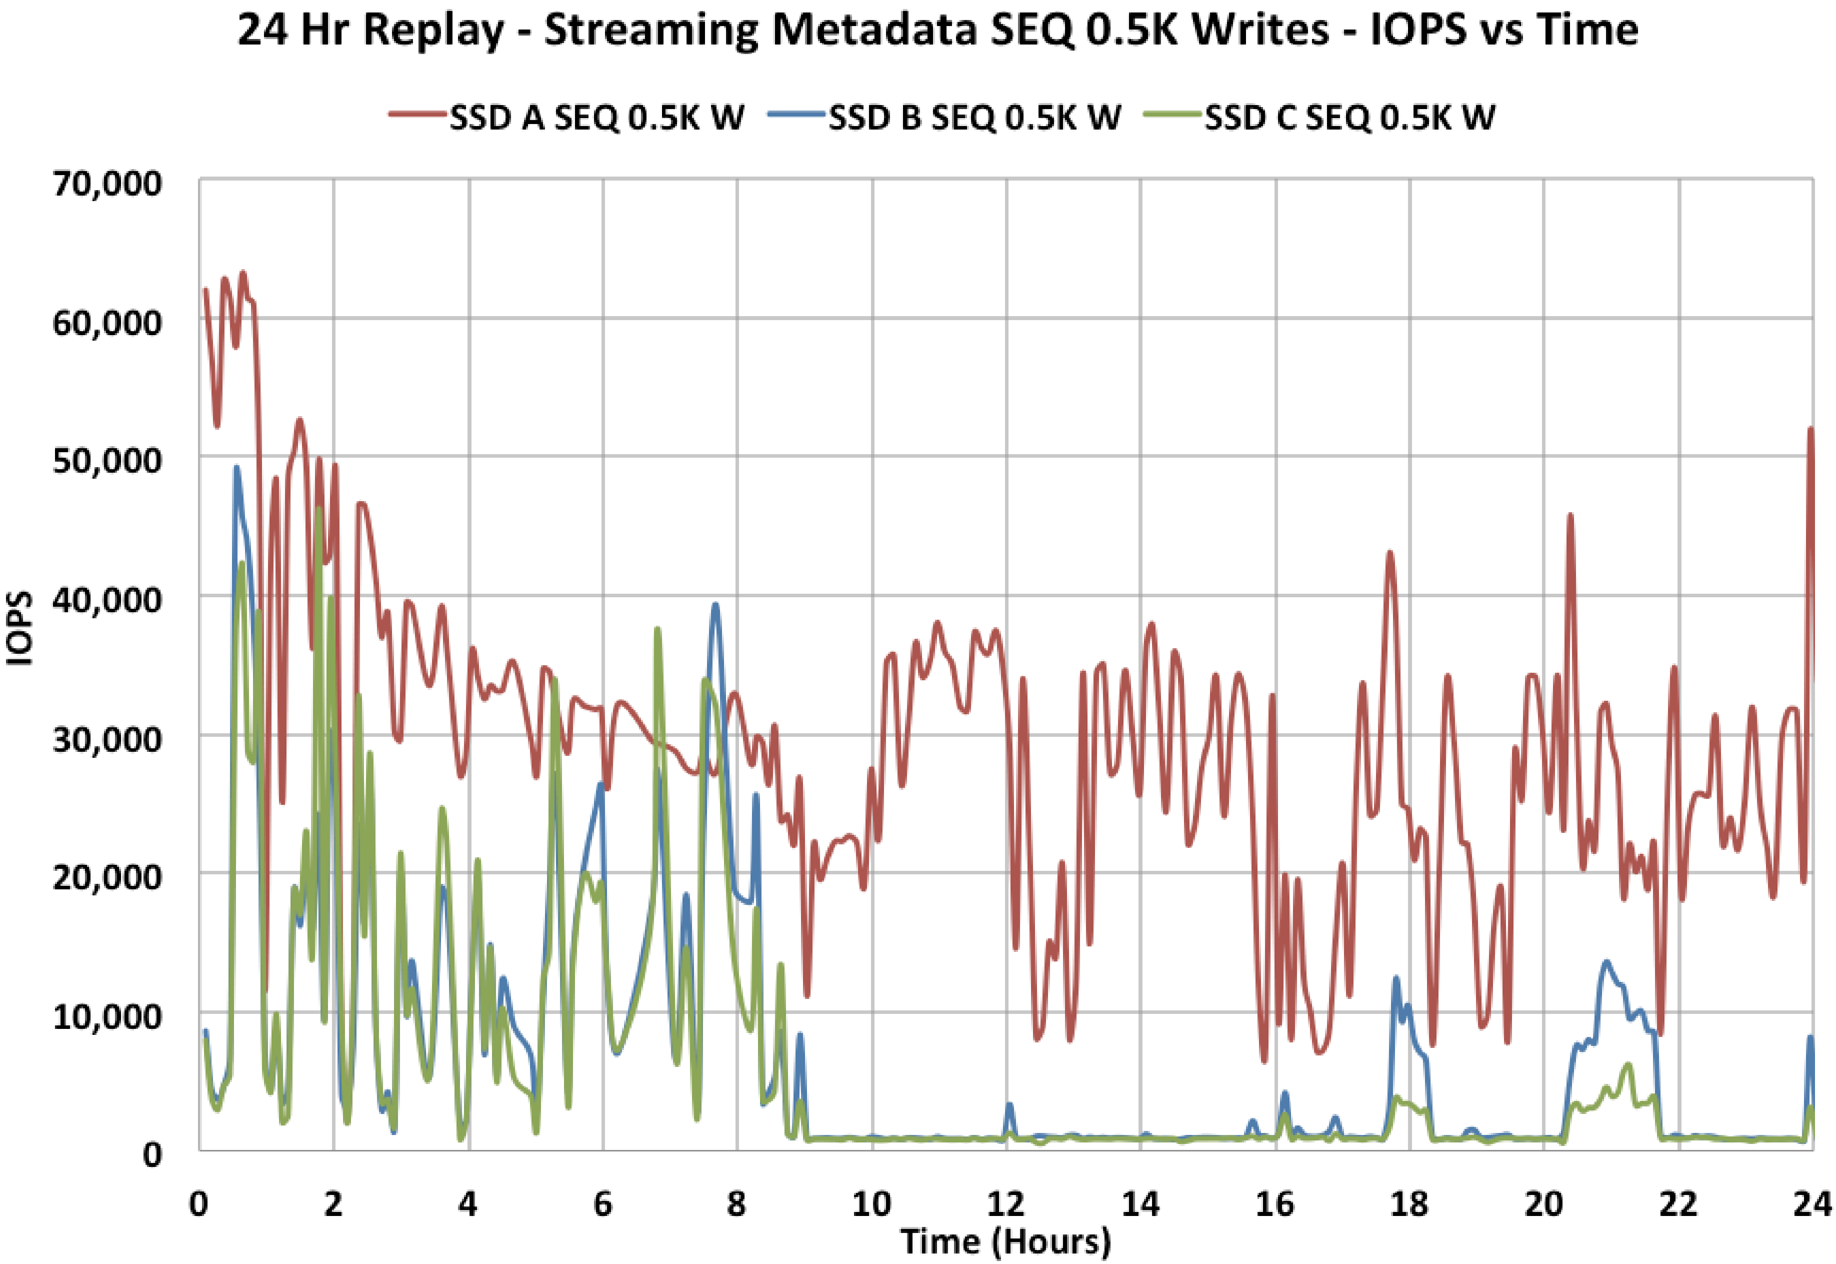
24 Hr Replay - Streaming Metadata SEQ 0.5K Writes - IOPS vs Time
SSD A SEQ 0.5K W
SSD B SEQ 0.5K W
SSD C SEQ 0.5K W
70,000
60,000
50,000
40,000
30,000
20,000
10,000
0
IOPS
0
2
4
6
8
10
12
14
16
18
20
22
24
Time (Hours)

SDC 16

Real World Workloads Are Changing Groups of IOs

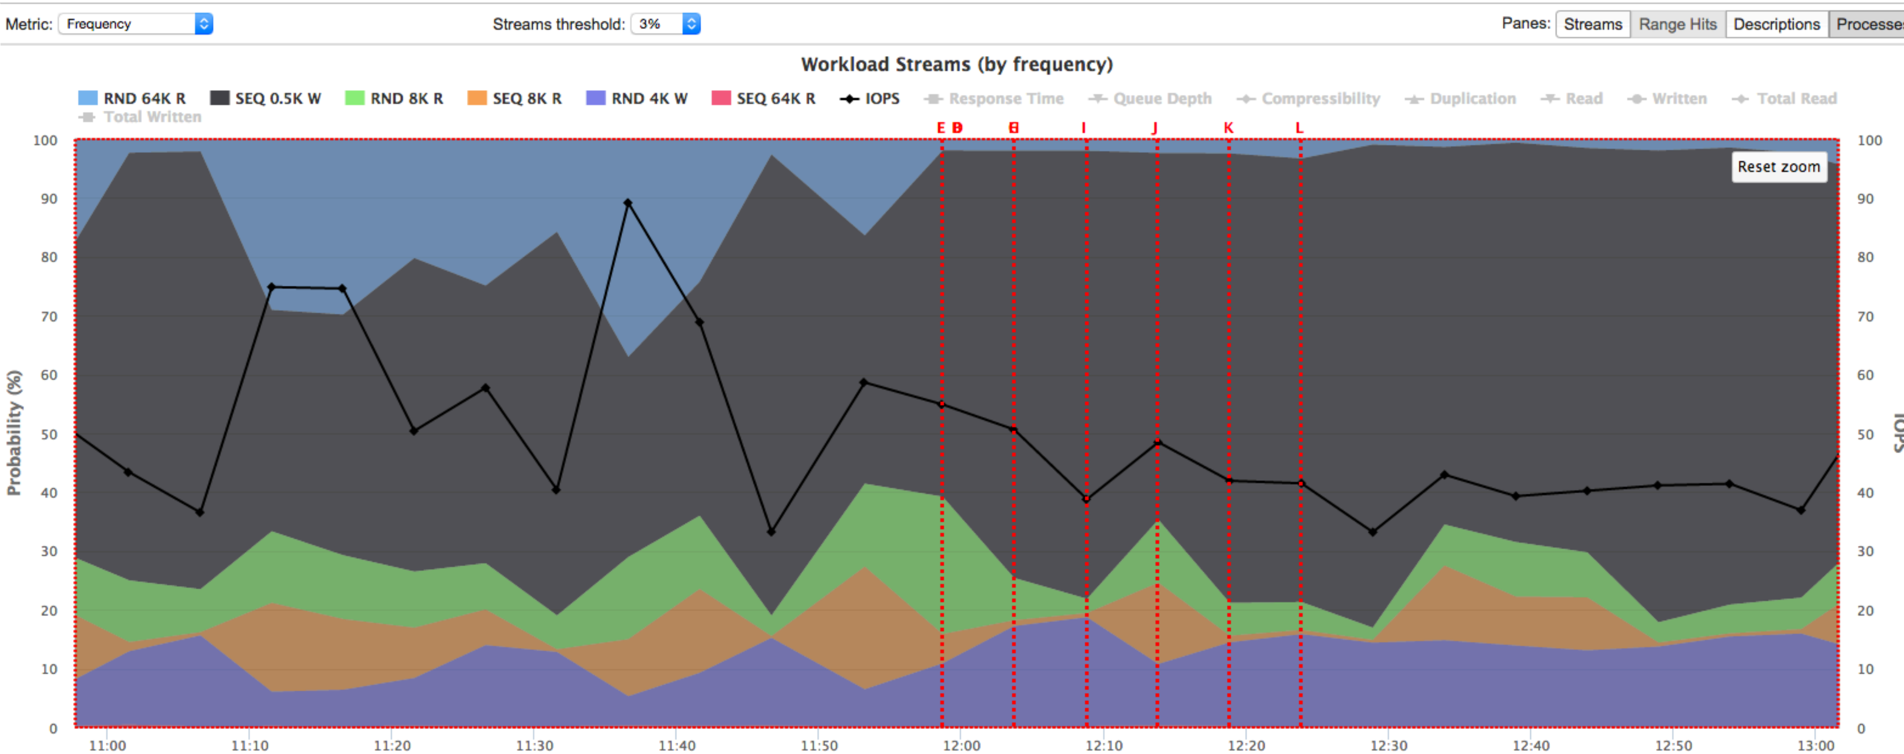

F	11:58:39	
☑ SEQ 0.5K W	40.5%	6,674
☑ RND 8K R	16.1%	2,650
☑ RND 4K W	7.3%	1,209
☑ RND 8K W	6.1%	1,008
☑ SEQ 8K R	3.4%	565
☐ RND 4K R	1.77%	291
☐ SEQ 1K W	1.71%	281
☐ RND 64K R	1.33%	219
☐ SEQ 4K W	1.19%	196

Total IOs of 591 streams: 16,469
Selected 5 streams: 12,106 (73.5%)

G	12:03:41	
☑ SEQ 0.5K W	32.3%	4,908
☑ SEQ 1K W	9.6%	1,464
☑ RND 4K W	7.6%	1,154
☑ RND 8K W	7.6%	1,147
☑ RND 16K R	3.2%	489
☑ RND 8K R	3.2%	487
☐ RND 4K R	2.67%	406
☐ RND 32K R	2.46%	373
☐ SEQ 64K W	2.17%	329

Total IOs of 547 streams: 15,190
Selected 6 streams: 9,649 (63.5%)

H	12:03:41	
☑ SEQ 0.5K W	32.3%	4,908
☑ SEQ 1K W	9.6%	1,464
☑ RND 4K W	7.6%	1,154
☑ RND 8K W	7.6%	1,147
☑ RND 16K R	3.2%	489
☑ RND 8K R	3.2%	487
☐ RND 4K R	2.67%	406
☐ RND 32K R	2.46%	373
☐ SEQ 64K W	2.17%	329

Total IOs of 547 streams: 15,190
Selected 6 streams: 9,649 (63.5%)

I	12:08:48	
☑ SEQ 0.5K W	44.8%	5,209
☑ RND 4K W	10.9%	1,265
☑ RND 8K W	7.2%	838
☑ SEQ 64K W	3.1%	365
☐ RND 4K R	2.78%	323
☐ RND 64.5K R	2.53%	294
☐ SEQ 1K W	1.77%	206
☐ RND 128.5K R	1.74%	202
☐ SEQ 4K W	1.57%	183

Total IOs of 519 streams: 11,636
Selected 4 streams: 7,677 (66%)

J	12:13:49	
☑ SEQ 0.5K W	38.1%	5,541
☑ SEQ 8K R	8.5%	1,229
☑ RND 8K R	6.5%	952
☑ RND 4K W	6.4%	930
☑ RND 8K W	5.9%	860
☑ RND 4K R	5.2%	760
☐ RND 12K R	2.85%	414
☐ SEQ 1K W	2.24%	326
☐ SEQ 4K R	2.21%	322

Total IOs of 456 streams: 14,544
Selected 6 streams: 10,272 (70.6%)

K	12:18:50	
☑ SEQ 0.5K W	47.5%	5,979
☑ RND 4K W	8.9%	1,113
☑ RND 8K W	7.1%	896
☑ RND 8K R	3.5%	436
☐ RND 64.5K R	2.50%	314
☐ RND 16K R	2.04%	256
☐ SEQ 1K W	2.02%	254
☐ RND 64K R	1.51%	190
☐ RND 4K R	1.36%	171

Total IOs of 477 streams: 12,575
Selected 4 streams: 8,424 (67%)

Single SSD WAF: Drive Fills, FTL and WAF for Different Workloads

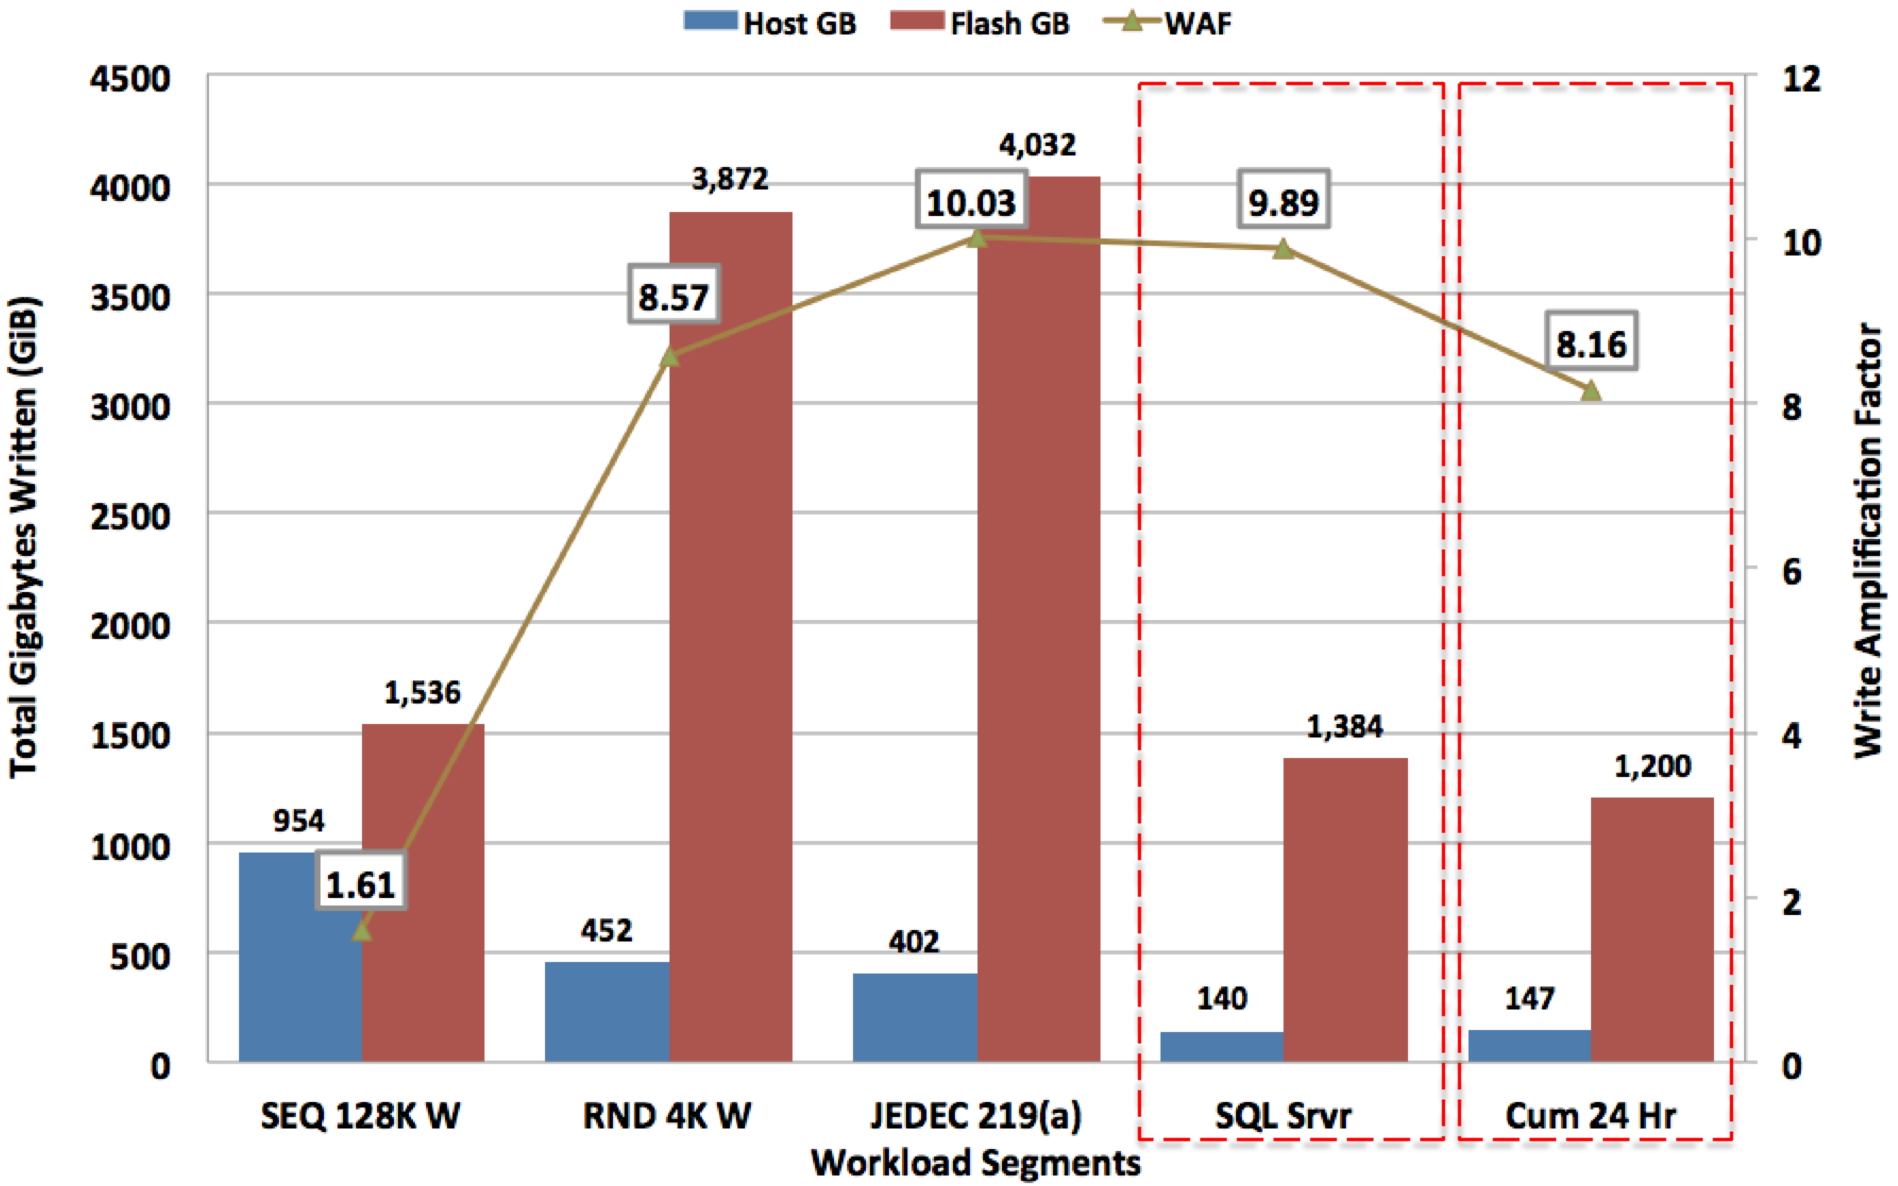

Take-Aways

- SSD Performance Depends on the SSD Workload
- IO Streams Change as they Traverse the SW Stack
- Replay Storage Workloads are Changing Groups of IO Streams
- Individual IO Streams Can Be Compared to Synthetic Benchmarks
- IO Workloads Result In Specific Performance & Endurance
- Each Real World IO Capture is Unique to its Time and System


AGENDA
E


Q&A

Understand Your SSD Workloads

The Example #3 Workload in this Presentation
Can Be Accessed as a Live Demo at

TestMyWorkload.com

Try FREE Capture & Analysis of Your SSD
Real World Storage Workloads Today!

For more information, contact Calypso Systems, Inc.

info@calypsotesters.com www.calypsotesters.com