Bring Your SSD Testing Up to Date

Tutorial – SNIA Performance Test Specifications

PTS 2.0.1 for Solid State Storage Devices
RWSW PTS 1.0.7 for Datacenter Storage

Eden Kim
Calypso Systems, Inc.
Chair, SNIA Solid State Storage Technical Working Group (SSS TWG)

Agenda

Introduction		1:00 – 1:15
Part 1:	PTS v 2.0.1 - Updates	1:15 – 2:00
	Overview	
	Updated Tests & Metrics	
Break		2:00 – 2:15
Part 2:	RWSW PTS v 1.0.7 - Capture & Analysis	2:15 – 3:30
	Overview - Real World Workloads	
	Capture – IO Captures & Reference Workloads	
	Analysis – IO Stream Map & LBA Range Hit Map	
Break		3:30 – 3:45
Part 3:	RWSW PTS 1.0.7 - Test	3:45 – 4:45
	In-situ Target Server Self-Test	
	Replay Native	
	Multi-WSAT	
	Individual Streams WSAT	
	RWSW DIRTH	
Questions		4:45 – 5:00

Introduction

- Two New SNIA SSS PTS Technical Positions
  - "PTS" - Performance Test Specifications for SSDs & Datacenter Storage
  - PTS v 2.0.1 Feb. 2018 - for SSDs
  - RWSW PTS v 1.0.7 May 2018 - for Datacenter Storage
- PTS v 2.0.1 – for Solid State Storage Devices
  - Focus is on the pre-conditioning of NAND Flash to Steady State
  - Workloads are corner case stress tests and synthetic emulations of applications
  - Tests are for Comparative Performance benchmarking
- RWSW PTS 1.0.7 – for Datacenter Storage
  - Focus is on the Capture, Analysis and Test of Real World Workloads
  - Real World Workloads are Derived from IO Trace Captures of User Applications
  - Tests are for the Analysis, Qualification & Validation of Datacenter Storage

Flash Memory Summit

Part 1: PTS v 2.0.1 for SSDs - Overview

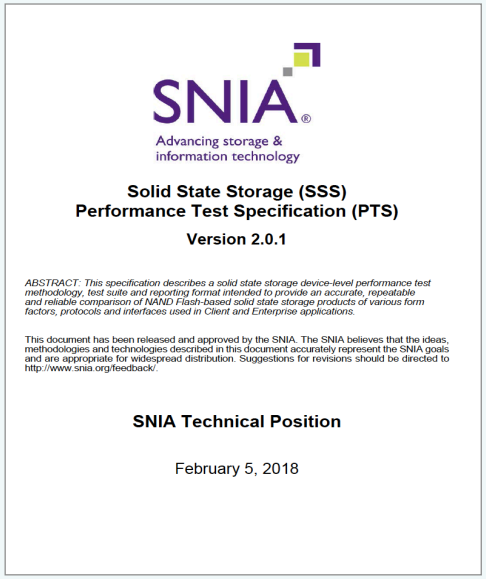
SNIA
Advancing storage & information technology

Solid State Storage (SSS)
Performance Test Specification (PTS)
Version 2.0.1

ABSTRACT: This specification describes a solid state storage device-level performance test methodology, test suite and reporting format intended to provide an accurate, repeatable and reliable comparison of NAND Flash-based solid state storage products of various form factors, protocols and interfaces used in Client and Enterprise applications.

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

February 5, 2018

www.snia.org/pts

- Consolidates PTS Client v1.2 and Enterprise v1.1 into 1 doc
- Updates IOPS Test to 3 Steady State Variables
- Adds WSAT, DIRTH and HIR to Client Tests
- Defines 6 synthetic workloads for WSAT & DIRTH:
  a. Meta Data & Journaling: SEQ 0.5K RW
  b. Write Intensive: RND 4K W
  c. Read Intensive: RND 4K R
  d. Database OLTP: RND 8K RW65
  e. VOD: SEQ 128K RW90
  f. Composite Block Size (CBS): JEDEC 219(a) Composite workload
- Adds Optional Secondary Metrics:
  a. Power consumption; IOPS/Watt
  b. Average/Maximum RTs and 5 9s Response Time Quality of Service
  c. CPU System Usage %; CPU IO Wait

PTS v2.0.1 – Tests & Test Flows

Test Family	Tests	Purpose	Workload Type	Pre-condition	Steady State
SNIA Basic	IOPS TP 128K, TP 1024K LAT	IO Rate Bandwidth (BW) Bandwidth (BW) Single IO Response Time (RT)	Mixed BS Loop	PURGE, WIPC, WDPC	20/10 5 Round Post Test Inspection
Saturation	WSAT	FOB IO, BW & RT Saturation	Single Access Patterns Composite BS Mix	PURGE, WDPC	Time, TGBW, Drive Fills Rounds with BTW
Optimized	DIRTH	IO, BW & RT Saturation Quality of Service Performance over Range of Users	Single Access Patterns Composite BS Mix	PURGE, WIPC, WDPC	Rounds with BTW
Sustained	XSR HIR	Garbage Collection Recovery Garbage Collection Recovery	Alternating Segments Alternating Segments	PURGE, WDPC PURGE, WDPC	Time Rounds with BTW

- ❖ Tests have Different Workload Types
  - Mixed Block Size Loops
  - Single Access Pattern
  - Composite BS Mix
  - Alternating Segments
- ❖ Tests use different PC
  - WIPC & WDPC
  - WDPC only
- ❖ Tests have Different SS Criteria
  - 20/10 % Round Formula
  - 5 Round with Between Round Writes
  - 25 Rounds w/ Post Test Inspection
  - Time, TGBW, Drive Fills

PTS v2.0.1 – Tests Workload Access Patterns & Demand Intensity

Test	Purpose	Workload Type	Workload Access Patterns	Demand Intensity
IOPS	IO Rate	Mixed BS Loop 56 RW elements	RW: 100R, 95:05, 65:35; 50:50; 35:65; 05:95; 100W BS: 1024K; 128K; 64K; 32K; 16K; 8K; 4K; 0.5K	Ent: T4Q32 Client: T2Q16
TP 128K; TP 1024K	IO, BW & RT Saturation	Single Access Patterns 2 RW elements	RW: 100R; 100W BS: [SEQ 128K; SEQ 1024K]	T1Q32
LAT	IO, BW & RT Quality of Service	Single Access Patterns 9 RW elements	RW: 100R; 65:35 RW; 100W BS: 8K; 4K; 0.5K	T1Q1
WSAT	FOB IO, BW & RT Saturation	Single Access Patterns Composite Block Size	RND 4K R; RND 4K W; RND 8K RW65 SEQ 128K RW90; SEQ 0.5K RW50 CBS 14 BS Composite (JEDEC 219(a))	Ent RND: T4Q32 Client RND: T2Q16 SEQ: T1Q32
DIRTH	IO, BW & RT Saturation RT QoS	Single Access Patterns Composite Block Size	RND 4K R; RND 4K W; RND 8K RW65 SEQ 128K RW90; SEQ 0.5K RW50 CBS 14 BS Composite (JEDEC 219(a))	Ent RND: T4Q32 Client RND: T2Q16
XSR	RT & BW Recovery from Super Saturation	Alternating Segments	SEQ 1024KW; RND 8K W; SEQ 1024K W	T1Q32; T4Q32; T1Q32
HIR	RT & BW Recovery Effects of Host Idles	RND 4K W Varying Host Idles	RND 4K W; 0,5,10,15,15 & 50 sec Host Idles	Ent: T4Q32 Client: T2Q16

PTS v2.0.1 Basic Test Process – PURGE

Basic Test Process:	
1.	PURGE
2.	Workload Independent Pre-conditioning
3.	Workload Dependent Pre-conditioning
4.	Types of Workloads
5.	Test Settings
6.	Steady State Determination
7.	Report Data from Steady State Window

1. PURGE[1]

- Return the Device to a state as if no writes have occurred by:
  - a) ATA: SECURITY ERASE
  - b) SCSI: FORMAT UNIT
  - c) NVMe: FORMAT namespace or
  - d) Vendor Specific Command / Methods
- Device PURGE creates a known and repeatable test starting point and facilitates a clear demonstration of the Steady State convergence behavior
- PURGE is a Required PTS 2.0.1 test step

[1] PTS 2.0.1 – p 21

PTS v2.0.1 – Basic Test Process: WIPC

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

2. Workload Independent Pre-conditioning (WIPC)[2]

- Apply a workload to facilitate convergence to Steady State
  - a) Write twice the stated User Capacity
  - b) Apply SEQ 128K W at T1 Q32 or
  - c) Apply SEQ 1024K W at T1 Q32
  - d) WIPC is 'independent' of the workload of interest ('dependent')
- WIPC is applied to facilitate convergence to Steady State by organizing LBA tables prior to application of WDPC
- NOTE: WIPC and/or WDPC by themselves do not put an SSD into a Steady State but prepares and applies the workload of interest for calculation of the Steady State determination criteria

[2] PTS 2.0.1 – p 17

PTS v2.0.1 – Basic Test Process: WDPC

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

3. Workload Dependent Pre-conditioning (WIPC)[3]

- Apply the Workload of Interest until convergence to Steady State
- WDPC applies the Workload of Interest for calculation of the Steady State determination criteria
- Steady State determination criteria depend on the type of workload

[3] PTS 2.0.1 – p 17

PTS v2.0.1 – Types of Workloads

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- There are three types of PTS 2.0.1 WDPC Workloads:[4]
  a) Mixed Block Size Loop (IOPS, TP, LAT)
  b) Single Access Pattern (WSAT, DIRTH, XSR, HIR)
  c) Fixed Composite (CBW WSAT, CBW DIRTH)

[4] PTS 2.0.1 – p 22

PTS v2.0.1 – Types of Workloads: Mixed BS Loops

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Mixed Block Size Loops – Nested RW & BS loops
- IOPS Test: 56 one minute RW/BS mix elements
  - a) (8) Read Write Mixes: 100% R, 95:05, 65:35, 50:50, 35:65, 05:95, 100% W
  - b) (7) Block Sizes: 1024K, 128K, 64K, 32K, 16K, 8K, 4K, 0.5K
- TP (Throughput) Test: 2 one minute RW/BS mix elements
  - a) (2) Read Write Mixes: 100% R, 100% W
  - b) (1) Block Size: [1024K TP test, 128K TP Test]
- LAT (Latency) Test: 9 one minute RW/BS mix elements
  - a) (3) Read Write Mixes: 100% R, 65:35 RW,100% W
  - b) (3) Block Sizes: 8K, 4K, 0.5K

PTS v2.0.1 – Types of Workloads: Mixed BS Loops

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- IOPS Test: One Round = 56 one minute RW/BS[5]
  - a) (8) Read Write Mixes: 100% R, 95:05, 65:35, 50:50, 35:65, 05:95, 100% W
  - b) (7) Block Sizes: 1024K, 128K, 64K, 32K, 16K, 8K, 4K, 0.5K

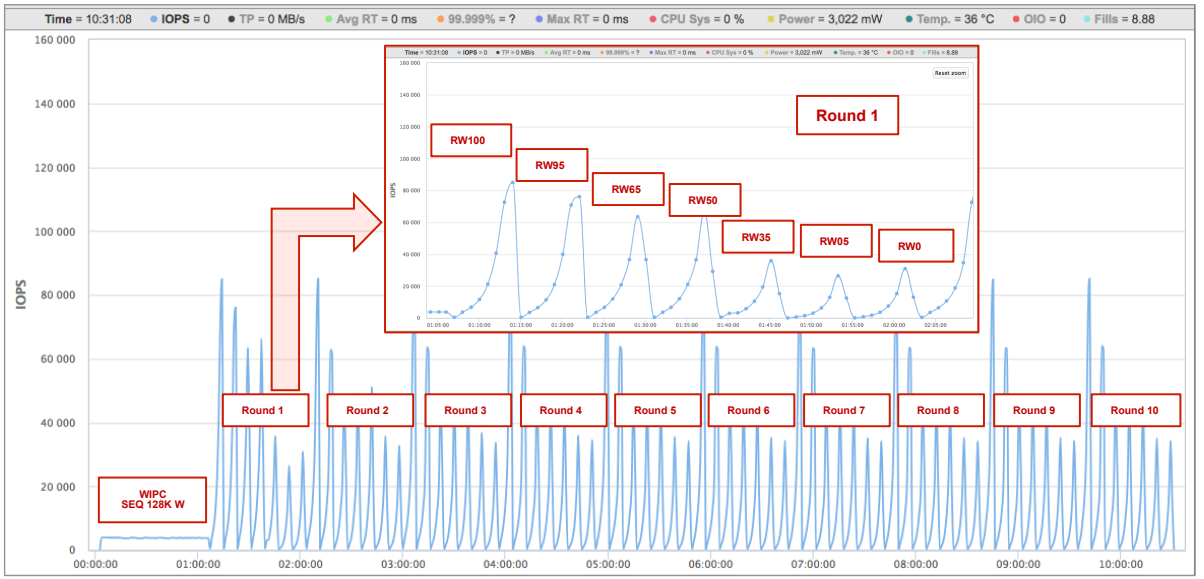

[5] PTS 2.0.1 – p 28

PTS v2.0.1 – Types of Workloads: Mixed BS Loops

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- TP (Throughput) Test: One Round = 2 one minute RW/BS[6]
  - a) (2) Read Write Mixes: 100% R, 100% W
  - b) (1) Block Size: [1024K TP test, 128K TP Test]

[6] PTS 2.0.1 – p 37

PTS v2.0.1 – Types of Workloads: Mixed BS Loops

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- LAT (Latency) Test: One Round = 9 one minute RW/BS[7]
  a) (3) Read Write Mixes: 100% R, 65:35 RW,100% W
  b) (3) Block Sizes: 8K, 4K, 0.5K

[7] PTS 2.0.1 – p 44

PTS v2.0.1 – Types of Workloads: Single Access Pattern

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Saturation[8] & DIRTH:[9] Single Access Pattern
  - a) Journal/Metadata: SEQ 0.5K R/W
  - b) Write Intensive: RND 4K W
  - c) Read Intensive: RND 4K R
  - d) Mixed/OLTP: RND 8K RW65
  - e) VOD: SEQ 128K RW90

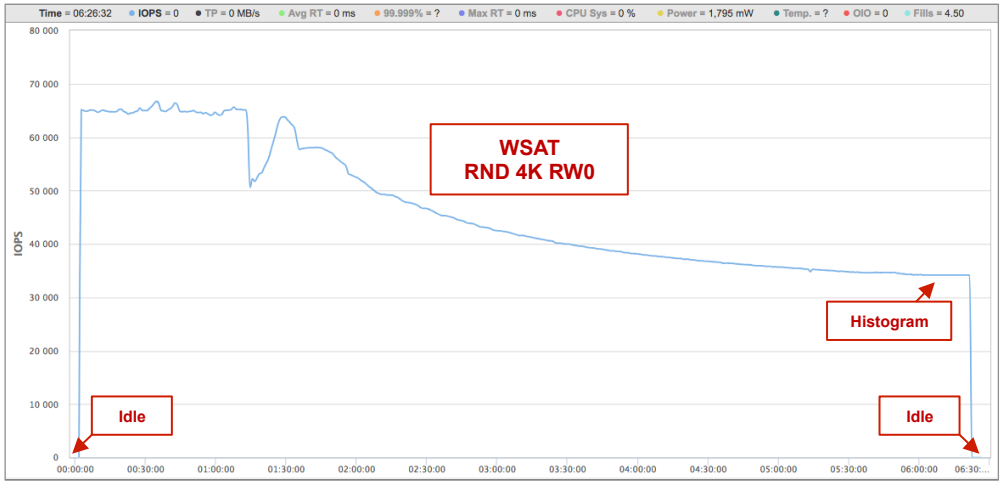

[8,9] PTS 2.0.1 – p 50, 80

PTS v2.0.1 – Types of Workloads: Single Access Pattern

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

❖ Host Idle Recovery:[10] Single Access Pattern RND 4K W

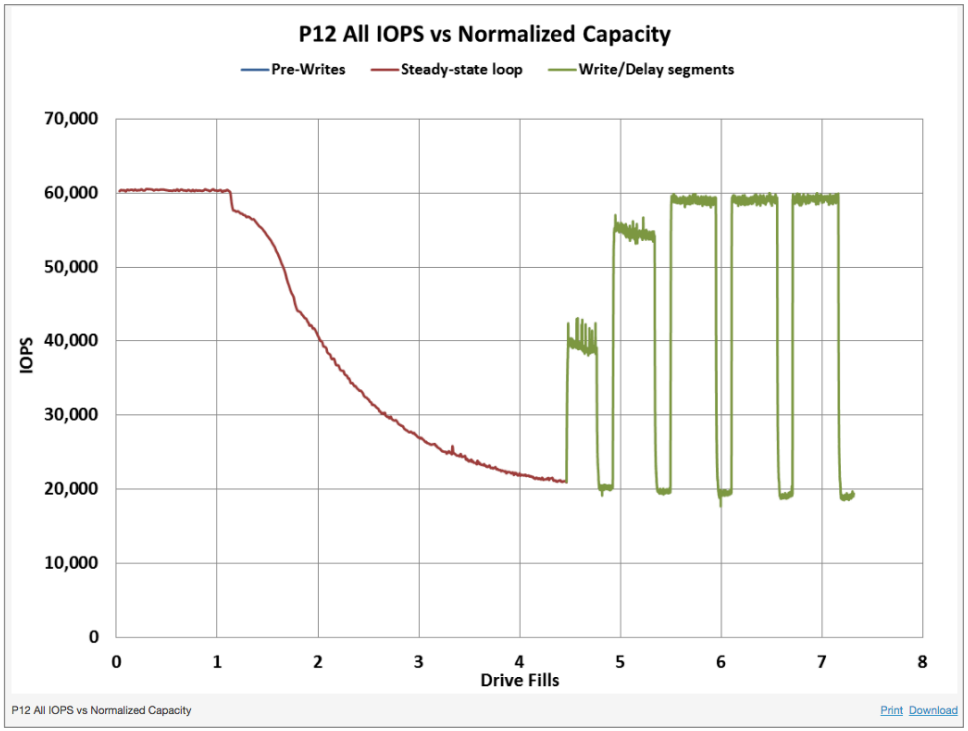

[10] PTS 2.0.1 – p 56

PTS v2.0.1 – Types of Workloads: Alternating Segment

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- XSR (Cross Stimulus Recovery):[11] Alternating Segments

SEQ 1024K W / RND 8K W / SEQ 1024K W

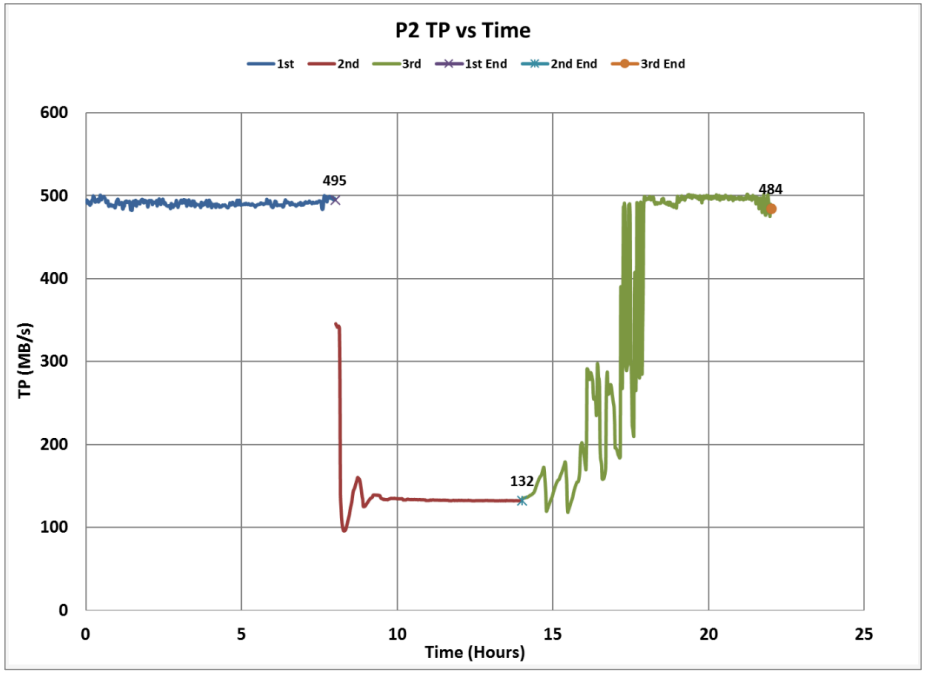

[11] PTS 2.0.1 – p 63

PTS v2.0.1 – Types of Workloads: Fixed Composite Workloads

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Saturation[12] & DIRTH[13]: Composite Block Size – JEDEC 219(a)
- Composite of 12 different RND Write Block Sizes

Block Size	Percent	Block Size	Percent
0.5K	4%	3.5K	1%
1.0K	1%	4.0K	67%
1.5K	1%	8.0K	10%
2.0K	1%	16.0K	7%
2.5K	1%	32.0K	3%
3.0K	1%	64.0K	3%

- Use of Restricted LBA Ranges

Restricted LBA Range Zones		
LBA Group A	50% of the IOs	To the first 5% LBAs
LBA Group B	30% of the IOs	To the next 15% LBAs
LBA Group C	20% of the IOs	To the remaining 20% LBAs

[12,13] PTS 2.0.1 – p 50,67

PTS v2.0.1 – Basic Test Process: Test Settings

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Test Settings are set forth for Enterprise (Ent) and Client tests
- Test Settings = Required or User Choice. Must be Disclosed
- Common Test Settings include:
  a) LBA Active Range: Ent = AR100; Client = AR75
  b) Volatile Write Cache: Ent = WCD; Client = WCE
  c) Binary Data Pattern: RND; Repeating; Binary File
- TC/QD Setting (Demand Intensity) for Specific Tests include:
  a) RND access workloads (e.g. IOPS, WSAT, HIR): Ent = T4Q32, Client = T2Q16
  b) SEQ access workloads (e.g. TP, XSR): T1Q32
  c) Single IO Response Time Measurement (e.g. LAT): T1Q1
  d) IO & RT Saturation (e.g. DIRTH): Outstanding IO range from T1Q1 to T32Q32

PTS v2.0.1 – Steady State Determination Methods

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. **Steady State Determination**
7. Report Data from Steady State Window

1. 20% Data Excursion/10% Slope 5 Round Formula[14]
2. 25 Round / Post Process Inspection[15]
3. 5 Round / 30 min Between Round Writes (BTW)[16]
4. Pseudo Steady State[17]
   a) Time
   b) Total GB Written
   c) Drive Fills

❖ All Reported Data Must from the 5 Round SS Range

[14,15,16,17] PTS 2.0.1 – p 24,24,50,50

Steady State: IOPS Test – Mixed Loop BS; 5 Round 20/10 SS

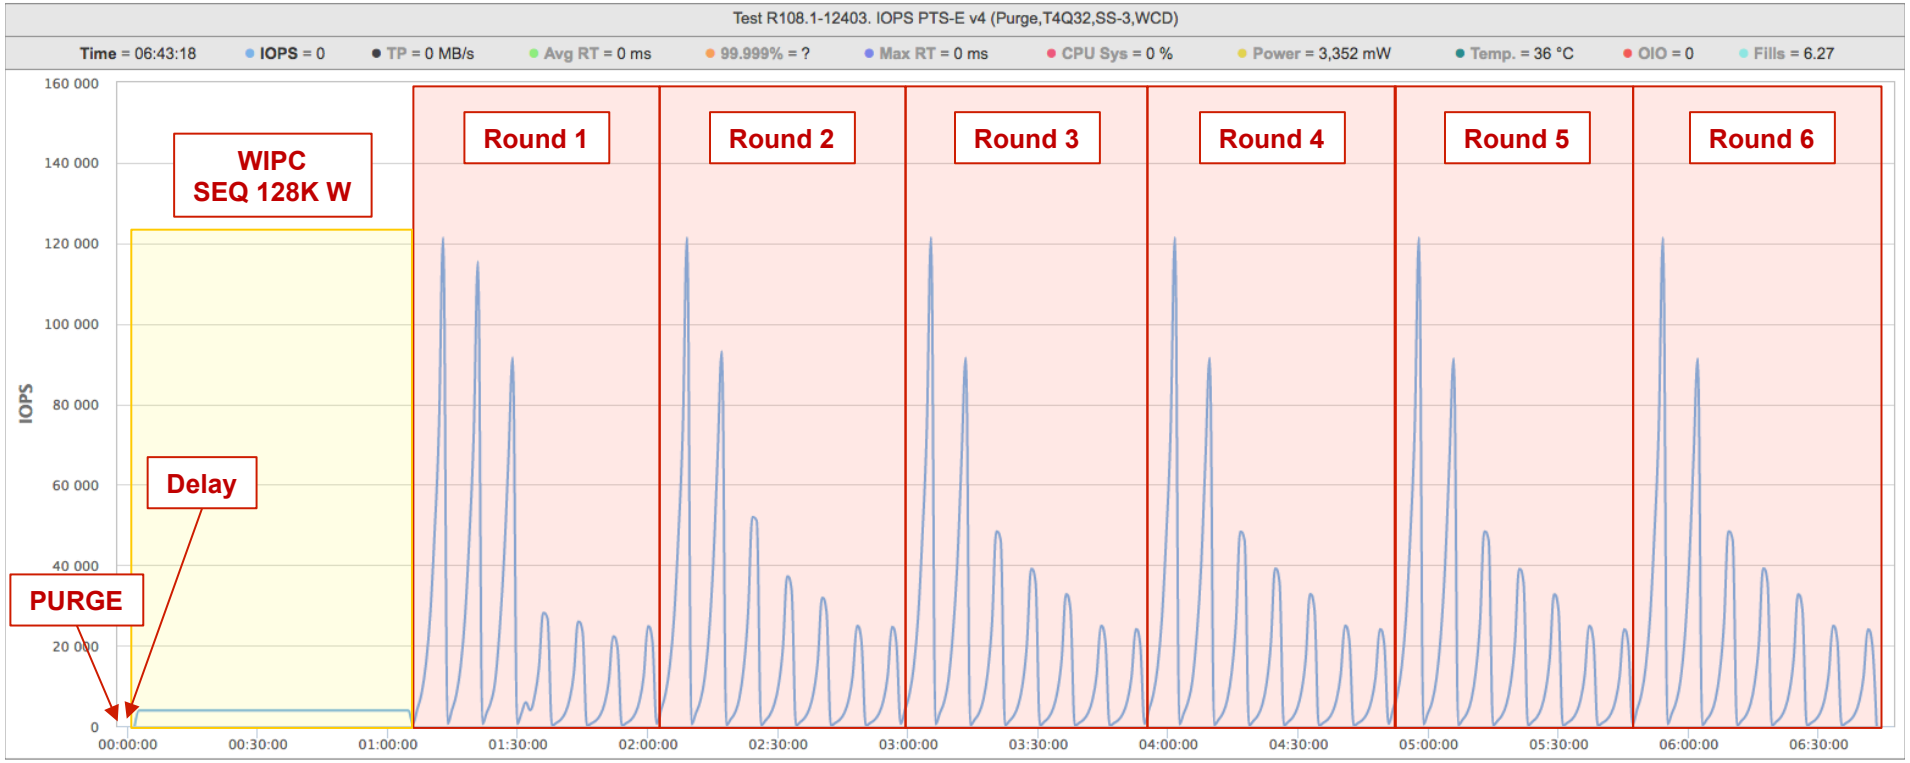

PTS v2.0.1 – "20/10" 5 Round Steady State Determination

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- "20/10" 5 Round Mixed Loop Steady State Determination
- (3) IOPS Steady State (SS) Determination Variables:[18]
  a) RND 4K W
  b) RND 64K RW65
  c) RND 1024K RW100
- Steady State Formula
  a) Five Consecutive Rounds
  b) 20% Data Excursion, 10% Slope Best Linear Fit Line
- All SS Variables Must be within 5 Round Range
- All Reported Data Must come from the 5 Round SS Range

[18] PTS 2.0.1 – p 28

PTS v2.0.1 – IOPS Mixed BS Loop SS Determination – SS Rounds 2-6

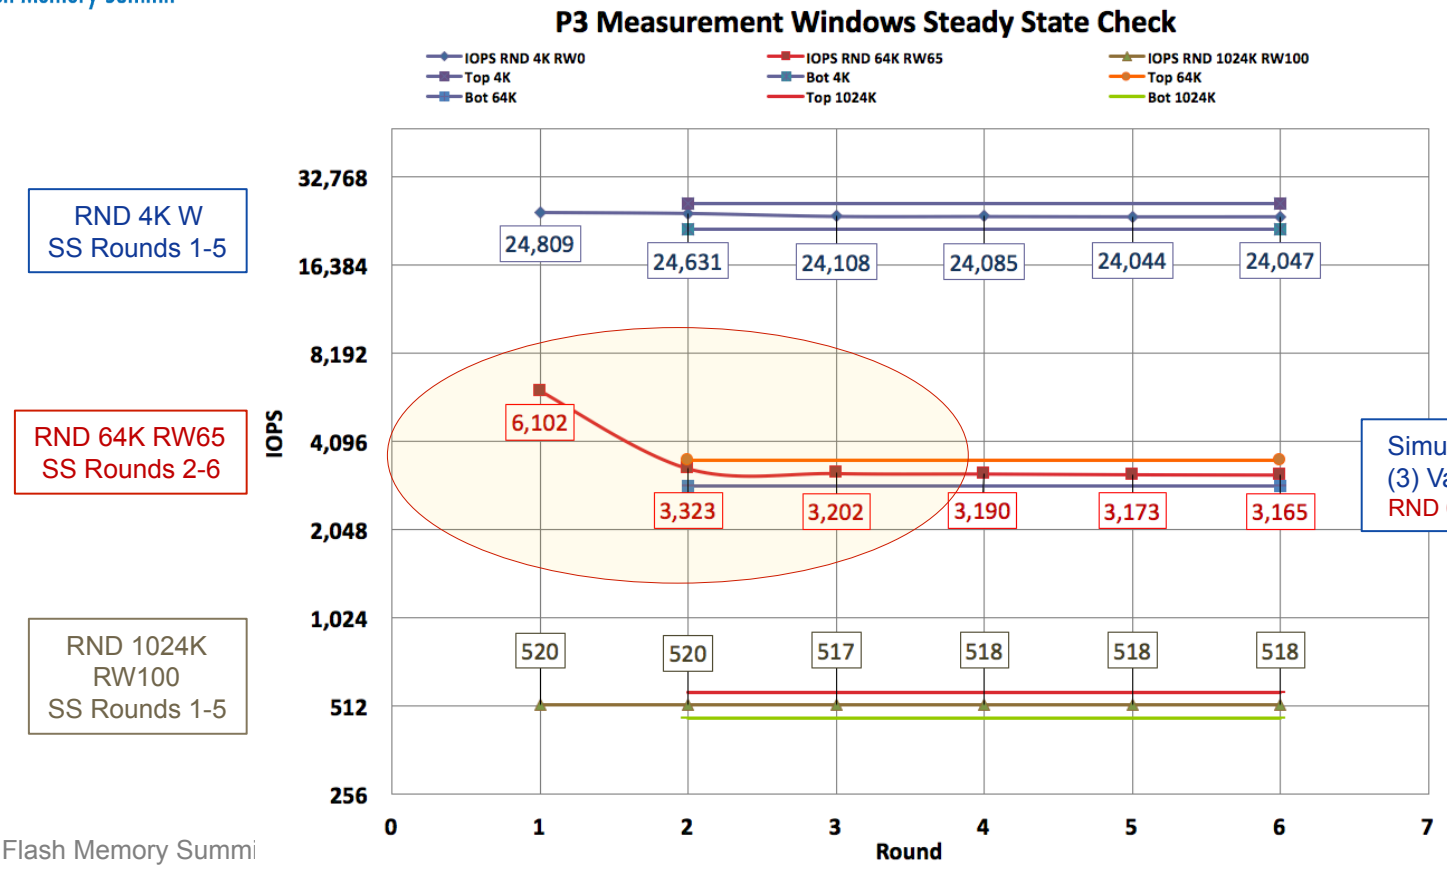

RND 4K W
SS Rounds 1-5

RND 64K RW65
SS Rounds 2-6

RND 1024K
RW100
SS Rounds 1-5

Simultaneous SS in Rounds 2-6
(3) Variable 20/10 (5) Round SS
RND 64K RW65 Round 1 Anomaly

PTS v2.0.1 – Rounds with Between Round Write SS Determination

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Single Access Pattern "5 Round/30 Min BTW" SS Determination
- 1 Min Rounds with 30 Min Between Round Writes (BTW):
  a) 1 Min Measurement Rounds
  b) 30 Min Between Round Writes
  c) Run Rounds until 5 Round "20/10" SS Determination is Met
- Single Access Pattern SS can be used for:
  a) DIRTH (Demand Intensity Response Time Histogram)
  b) HIR (Host Idle Recovery)
  c) WSAT (Write Saturation)
- All Reported Data Must come from the 5 Round SS Range

PTS v2.0.1 – Single Access Pattern Steady State Determination

16 (1) Min Rounds separated by 30 Min Between Round Writes (BTW)

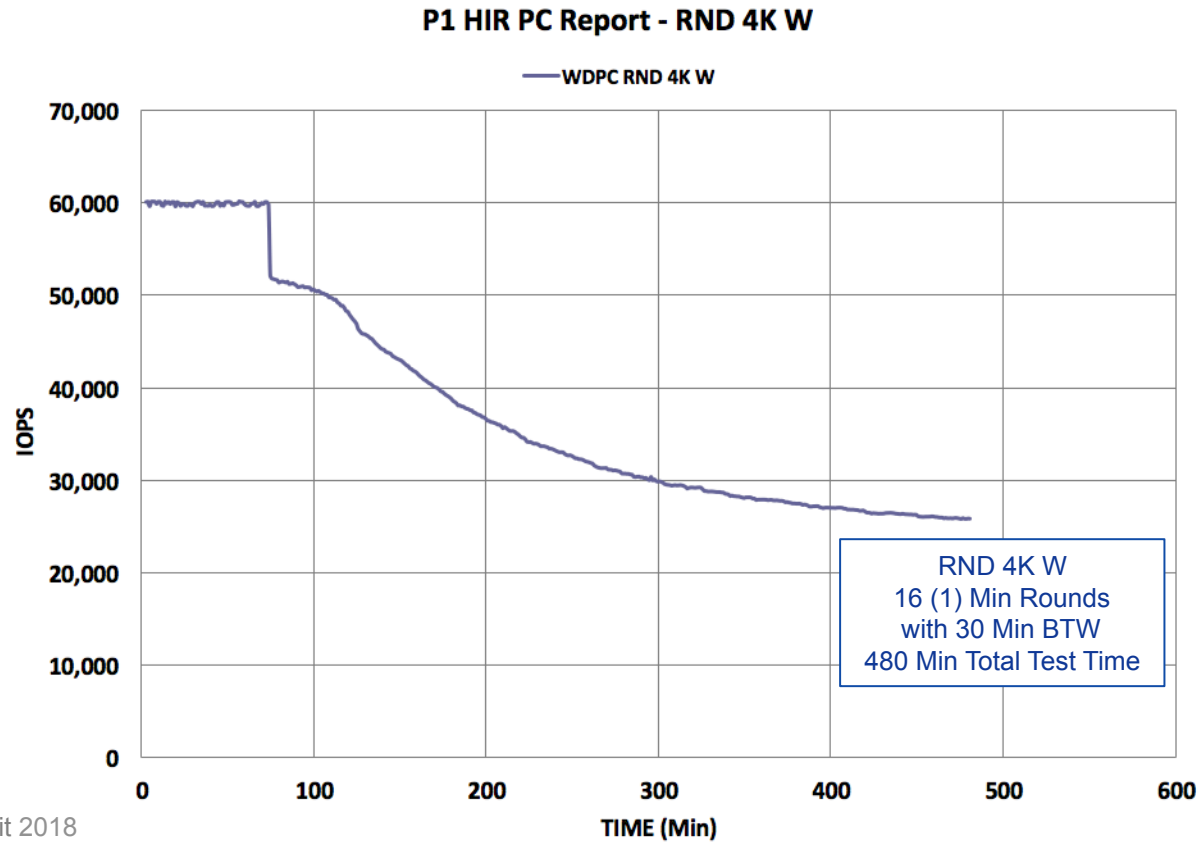

PTS v2.0.1 – Single Access Pattern Steady State – Rounds 12-16

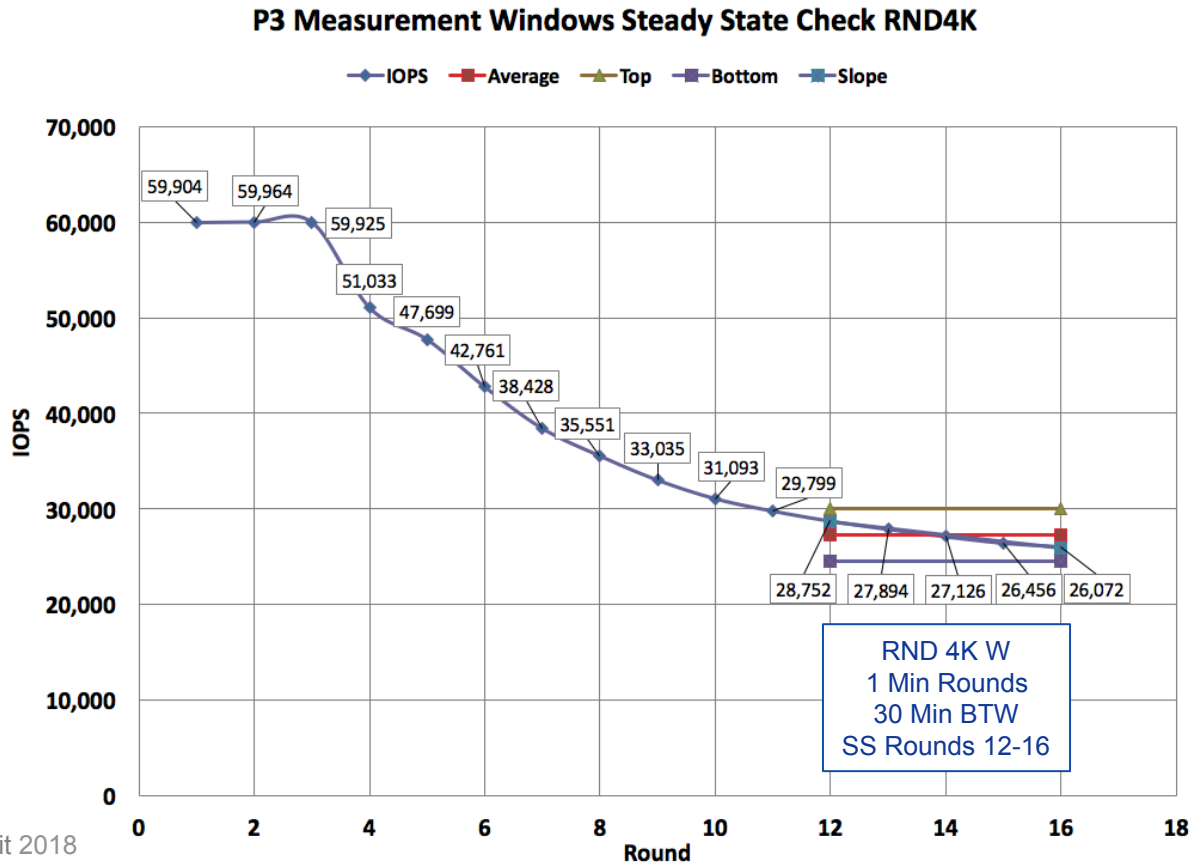

PTS v2.0.1 – Pseudo Steady State Determination

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Pseudo Steady State is appropriate when there are no post SS steps
  a. Appropriate for WSAT, IOPS, TP & LAT
  b. Post SS Rounds saturate and affect Response Times (DIRTH)
- Pseudo Steady State can be by:
  a) Time (WSAT)
  b) Total GB Written (WSAT)
  c) Drive Fills (WSAT)
- 25 Round Post Process determination:
  a) Run 25 Rounds
  b) Post process inspection for Steady State (IOPS)

PTS v2.0.1 – Steady State Reporting Data

Basic Test Process:

1. PURGE
2. Workload Independent Pre-conditioning
3. Workload Dependent Pre-conditioning
4. Types of Workloads
5. Test Settings
6. Steady State Determination
7. Report Data from Steady State Window

- Updated v2.0.1 Reporting: Primary & Secondary Metrics
- Required - Primary Metrics include:
  a) IOPS, TP & Response Times per individual test requirements
- Optional - Secondary Metrics include:
  a) Average/Maximum Response Times, 5 9s Response Time Quality of Service
  b) Power consumption; IOPS/Watt; Internal Temperature
  c) CPU System Usage; CPU IO Wait
- Secondary Metrics - may become Required in Future PTS Revisions
- See PTS 2.0.1 SNIA Report Format

PTS v2.0.1 – SNIA PTS 2.0.1 Report Header

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

Test Run Date:	01/19/16 08:11 PM	Report Run Date:	11/07/16 11:17 AM

IOPS Test (REQUIRED) - Report Page

SNIA SSS TWG	Solid State Storage Performance Test Spec (PTS)	IOPS - Block Size x RW Mix Matrix	Rev.	PTS-E 2.0
			Page	1 of 8

Vendor:	Mfg A	SSD Model:	SSD A - NVMe SSD	TEST SPONSOR	SNIA

Test Platform		Device Under Test		Set Up Parameters		Test Parameters	
Ref Test Platform	Calypso RTP 3.0	Mfgr	MFGR A	Data Pattern	RND_ONCE	Data Pattern	RND_ONCE
Motherboard		Model No.	ABCD	AR	100%	AR & Amount	100%
CPU	Intel(R) Xeon(R) CPU E5-2687W v2 @ 3.40GHz	S/N	0123456	AR Segments		Test Stimulus 1	IOPS Loop
Memory	31.3 GB	Firmware ver	0a1b2c3d	Pre Condtion 1	SEQ 128KiB W	RW Mix	Outer Loop
Operating System	CentOS 6.5	Capacity	2000 GB	TOIO - TC/QD	TC 1 / QD 32	Block Sizes	Inner Loop
Test SW	CTS 6.5 1.23.17	Interface	NVMe	Duration	2 x User Capacity	TOIO - TC/QD	TC 8 / QD 32
Test SW Info	1.23.17/1.9.267-el6	NAND Type	MLC	Pre Condtion 2	IOPS Loop	Steady State	3 - 7
Test ID No.	R84-9014	PCIe NVM	U.2	TOIO - TC/QD	TC 8 / QD 32	Test Stimulus 2	
HBA	SAS9212-4i4e	Purge Method	PURGED	SS Rounds	3 - 7	TOIO - TC/QD	
PCIe	Gen 3	Write Cache	Disabled	Note		Steady State	

Flash Memory Summit

PTS v2.0.1 – RND 4K W Steady State

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

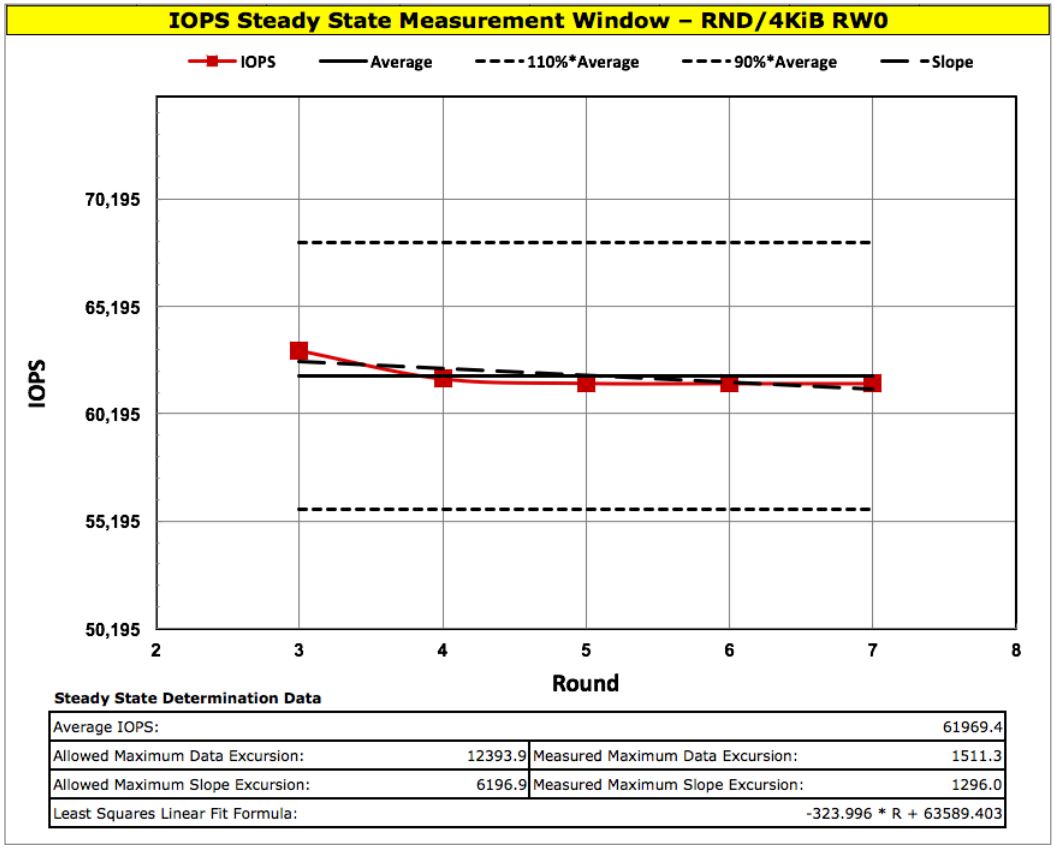

Steady State Determination Data

Average IOPS:			61969.4
Allowed Maximum Data Excursion:	12393.9	Measured Maximum Data Excursion:	1511.3
Allowed Maximum Slope Excursion:	6196.9	Measured Maximum Slope Excursion:	1296.0
Least Squares Linear Fit Formula:			-323.996 * R + 63589.403

PTS v2.0.1 – RND 64K RW65 Steady State

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

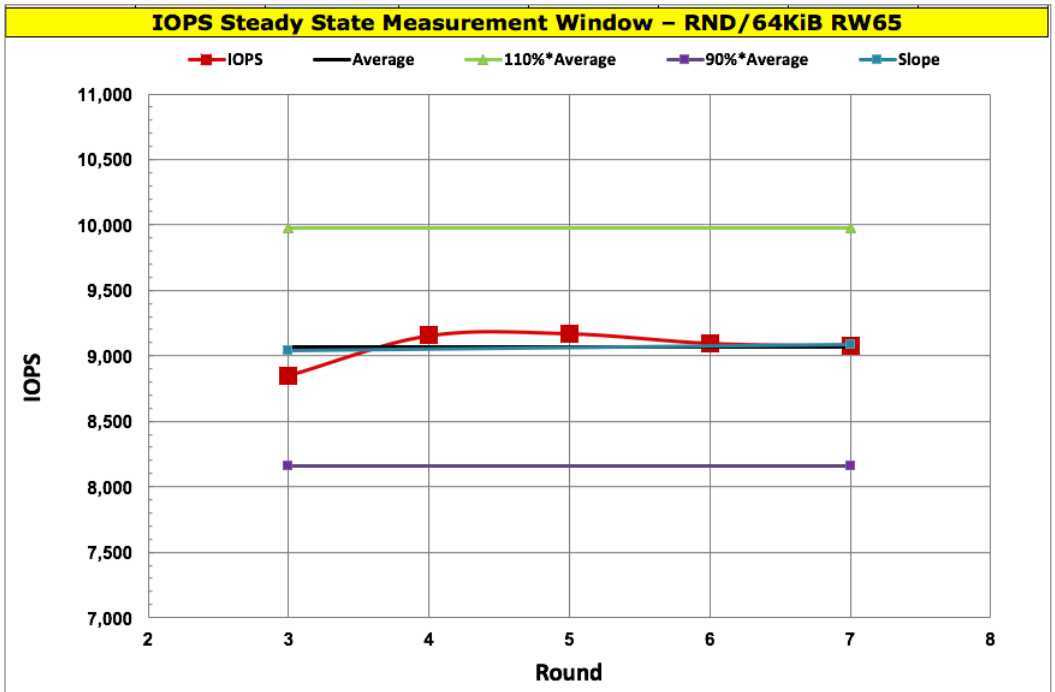

Steady State Determination Data

Average IOPS:			9070.6
Allowed Maximum Data Excursion:	9978.0	Measured Maximum Data Excursion:	1511.3
Allowed Maximum Slope Excursion:	6196.9	Measured Maximum Slope Excursion:	1296.0
Least Squares Linear Fit Formula:			-323.996 * R + 63589.403

PTS v2.0.1 – RND 1024K R Steady State

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

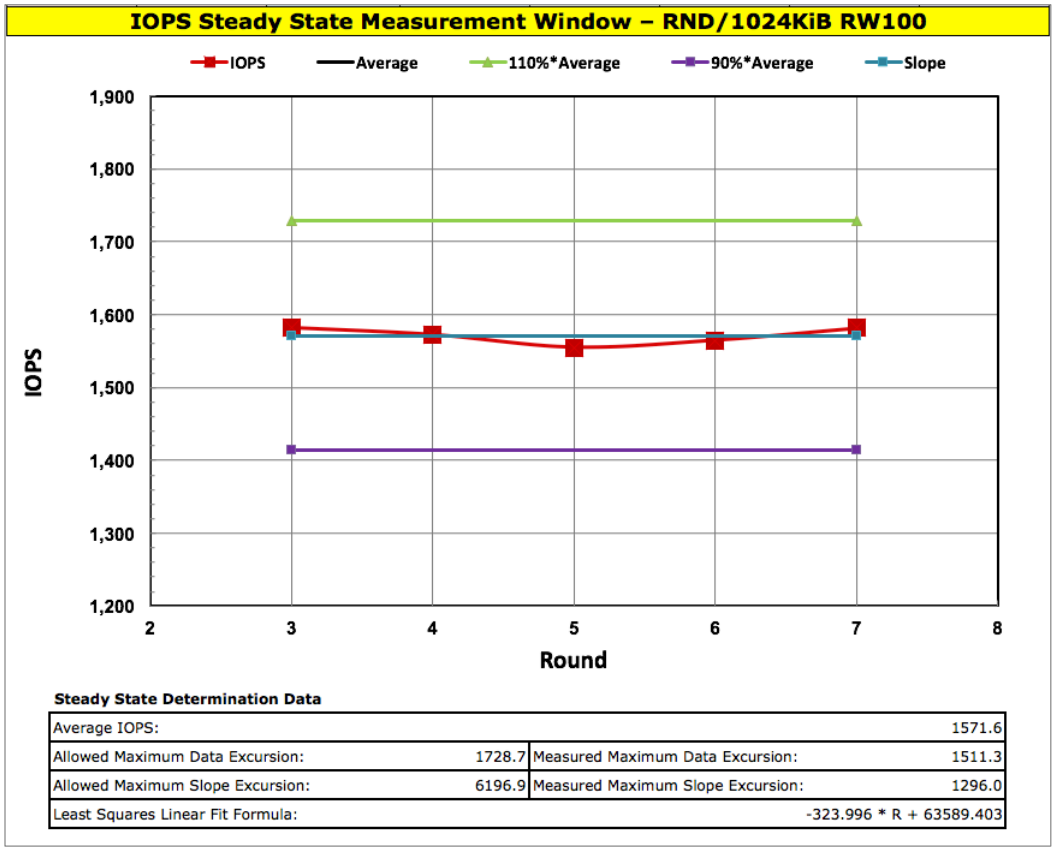

Steady State Determination Data

Average IOPS:			1571.6
Allowed Maximum Data Excursion:	1728.7	Measured Maximum Data Excursion:	1511.3
Allowed Maximum Slope Excursion:	6196.9	Measured Maximum Slope Excursion:	1296.0
Least Squares Linear Fit Formula:			-323.996 * R + 63589.403

PTS v2.0.1 – IOPS Table

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

Test Run Date:	01/19/16 08:11 PM	Report Run Date:	11/07/16 11:17 AM

IOPS Test (REQUIRED) - Report Page

SNIA SSS TWG	Solid State Storage Performance Test Spec (PTS)	IOPS - Block Size x RW Mix Matrix	Rev.	PTS-E 2.0
			Page	6 of 8

Vendor:	Mfg A	SSD Model:	SSD A - NVMe SSD	TEST SPONSOR	SNIA

Test Platform		Device Under Test		Set Up Parameters		Test Parameters	
Ref Test Platform	Calypso RTP 3.0	Mfgr	MFGR A	Data Pattern	RND_ONCE	Data Pattern	RND_ONCE
Motherboard		Model No.	ABCD	AR	100%	AR & Amount	100%
CPU	Intel(R) Xeon(R) CPU E5-2687W v2 @ 3.40GHz	S/N	0123456	AR Segments		Test Stimulus 1	IOPS Loop
Memory	31.3 GB	Firmware ver	0a1b2c3d	Pre Condtion 1	SEQ 128KiB W	RW Mix	Outer Loop
Operating System	CentOS 6.5	Capacity	2000 GB	TOIO - TC/QD	TC 1 / QD 32	Block Sizes	Inner Loop
Test SW	CTS 6.5 1.23.17	Interface	NVMe	Duration	2 x User Capacity	TOIO - TC/QD	
Test SW Info	1.23.17/1.9.267-el6	NAND Type	MLC	Pre Condtion 2	IOPS Loop	Steady State	3 - 7
Test ID No.	R84-9014	PCIe NVM	U.2	TOIO - TC/QD	TC 8 / QD 32	Test Stimulus 2	
HBA	SAS9212-4i4e	Purge Method	PURGED	SS Rounds	3 - 7	TOIO - TC/QD	
PCIe	Gen 3	Write Cache	Disabled	Note		Steady State	

IOPS - ALL RW Mix & BS – Tabular Data

Block Size (KiB)	Read / Write Mix %						
	0/100	5/95	35/65	50/50	65/35	95/5	100/0
0.5	47,139.1	48,885.0	70,740.4	89,676.2	118,331.1	338,953.4	499,097.2
4	61,969.4	65,324.4	90,078.8	108,316.3	104,022.5	310,634.6	407,395.3
8	31,086.5	32,730.9	44,174.2	43,065.4	67,758.7	171,091.2	207,067.5
16	15,612.4	16,417.3	22,832.9	24,658.0	37,991.9	89,209.8	104,000.5
32	7,813.4	8,237.5	11,537.4	14,568.4	18,790.9	46,002.2	52,149.4
64	3,904.7	4,126.0	5,739.5	7,519.0	9,071.1	22,856.6	26,054.5
128	1,964.8	2,065.9	2,843.5	3,725.5	4,465.9	10,992.1	12,948.0
1024	248.6	262.2	360.8	469.6	636.5	1,659.1	1,571.6

Secondary Metrics: IOPS, Ave Power mW, Temperature

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

IOPS | ART | 5 9s | MRT | Power | Temp.

IOPS								
R108.1-12403		PTS-E	WCD	T4/Q32	DP=RND	AR100	AIn=4K	v4
Data	Block Size	RW 0	RW 65	RW100	Block Size	RW 0	RW 65	RW100
IOPS	0.5 KiB	20,392	47,096	121,431	32 KiB	3,049	6,560	14,666
	4 KiB	**24,183**	**49,072**	**73,714**	64 KiB	1,525	3,211	7,820
	8 KiB	12,193	26,032	47,470	128 KiB	760	1,591	4,057
	16 KiB	6,112	13,410	26,901	1024 KiB	94	203	518
Avg Power mW	0.5 KiB	3,732	3,558	2,349	32 KiB	3,805	3,702	2,794
	4 KiB	**3,767**	**3,656**	**2,350**	64 KiB	3,807	3,632	2,879
	8 KiB	3,804	3,734	2,580	128 KiB	3,806	3,576	2,952
	16 KiB	3,809	3,745	2,691	1024 KiB	3,805	3,573	3,023
Temp. Celsius	0.5 KiB	39	38	32	32 KiB	39	38	35
	4 KiB	**39**	**39**	**32**	64 KiB	39	38	36
	8 KiB	39	39	33	128 KiB	38	38	36
	16 KiB	39	39	34	1024 KiB	38	36	37

- Secondary Metrics are Optional
  - Summary Table: 3 RW Mixes x 8 BS
  - Highlighted: RND 4K W, RND 4K RW65 & RND 4K R
  - IOPS; Average Power mW; Internal SSD reported Temp Celsius

Secondary Metrics: Average Response Time, 5 9s RT, Max RT

SNIA Report Format:

1. PTS 2.0.1 Header
2. RND 4K W SS
3. RND 64K RW65 SS
4. RND 1024K R SS
5. IOPS Table
6. Secondary Metrics

IOPS | ART | 5 9s | MRT | Power | Temp.

IOPS								
R108.1-12403		PTS-E	WCD	T4/Q32	DP=RND	AR100	AIn=4K	v4
Data	Block Size	RW 0	RW 65	RW100	Block Size	RW 0	RW 65	RW100
ART mSec	0.5 KiB	6.28	2.72	1.05	32 KiB	41.94	19.52	8.73
	4 KiB	**5.29**	**2.61**	**1.74**	64 KiB	83.80	39.84	16.36
	8 KiB	10.50	4.92	2.70	128 KiB	167.76	80.32	31.52
	16 KiB	20.93	9.56	4.76	1024 KiB	1,335.89	623.58	245.51
99.999% mSec	0.5 KiB	62.37	18.56	9.96	32 KiB	396.22	121.92	84.64
	4 KiB	**51.04**	**17.38**	**16.68**	64 KiB	812.70	253.47	156.37
	8 KiB	97.54	31.94	26.04	128 KiB	1,454.64	499.71	286.04
	16 KiB	207.22	59.95	45.87	1024 KiB	5,260.68	2,236.18	739.21
MRT mSec	0.5 KiB	79.48	21.23	13.32	32 KiB	445.54	128.28	112.90
	4 KiB	**71.75**	**20.61**	**21.91**	64 KiB	924.70	275.87	182.35
	8 KiB	118.33	36.44	33.06	128 KiB	1,677.05	531.65	350.00
	16 KiB	250.46	69.43	59.28	1024 KiB	6,079.57	2,483.48	743.89

- Secondary Metrics are Optional
  - Summary Table: 3 RW Mixes x 8 BS
  - Highlighted: RND 4K W, RND 4K RW65 & RND 4K R
  - Average Response Time, 5 9s RT, Maximum RT in mS

BREAK

2:00 – 2:15

Part 2: RWSW PTS v 1.0.7 - Overview

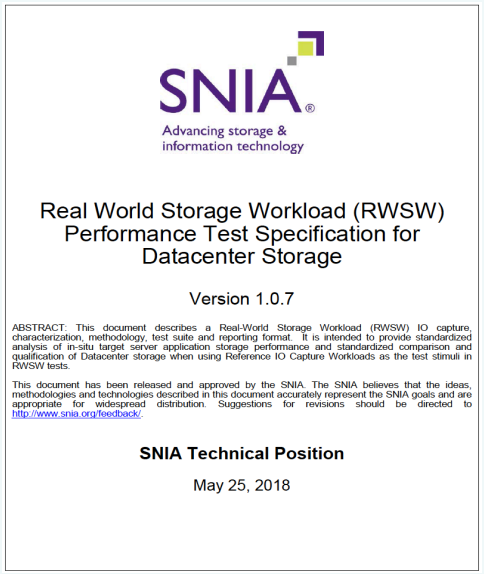
SNIA
Advancing storage & information technology

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

www.snia.org/rwsw

- RWSW PTS v1.0.7 describes methodologies for Real World IO Capture, Analysis & Test for Datacenter Storage
- IO Captures present IO Streams & Metrics that actually occur during real world application usage – not synthetic emulations
- RWSWs show that:
  a. Server Storage Performance depends, in large part, on the workload
  b. RWSWs are constantly changing combinations of many IO Streams and QDs
  c. IO Streams change as they traverse the HW/SW stack
  d. RWSW Replays can improve Failure Analysis & enable Storage Qualification
- RWSW PTS defines (4) RWSW Storage Tests designed to test and qualify storage to SSSI Reference and/or User captured RWSWs

Overview - Real World Storage Workloads

Overview of RWSWs:

1. What are Real World Storage Workloads?
2. Data Center SQL Server RWSW
3. Why are RWSWs Important?

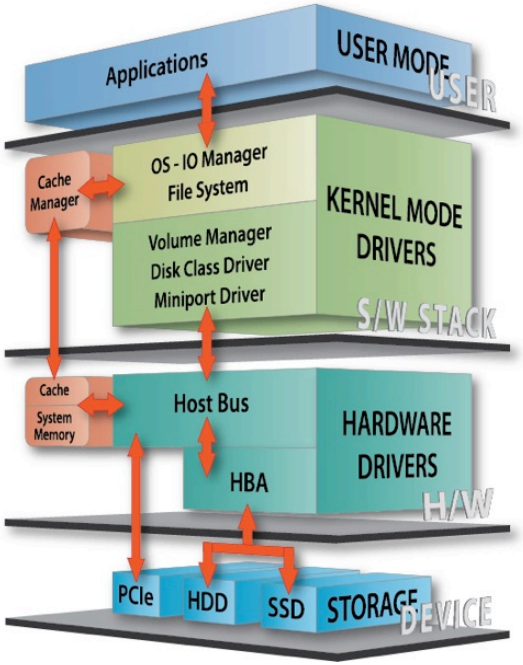

Windows Software Stack

Real World Storage Workloads (RWSWs) are:[19]

- IO Streams generated by Applications that traverse the SW stack from User space to Data Center Storage and back
- IO Streams that present to Storage at the File System, Block IO or other specified software layer
- IO Streams are modified at each layer of software abstraction by coalescing, fragmentation, appending & merging
- Data Center storage includes: SAN, NAS, DAS, JBOF, JBOD, SDS, Open, Virtualized, Object, LUN and SSD

[19] RWSW PTS v1.0.7 – page 12

Data Center 24 Hour RWSW - Retail Web Portal[20]

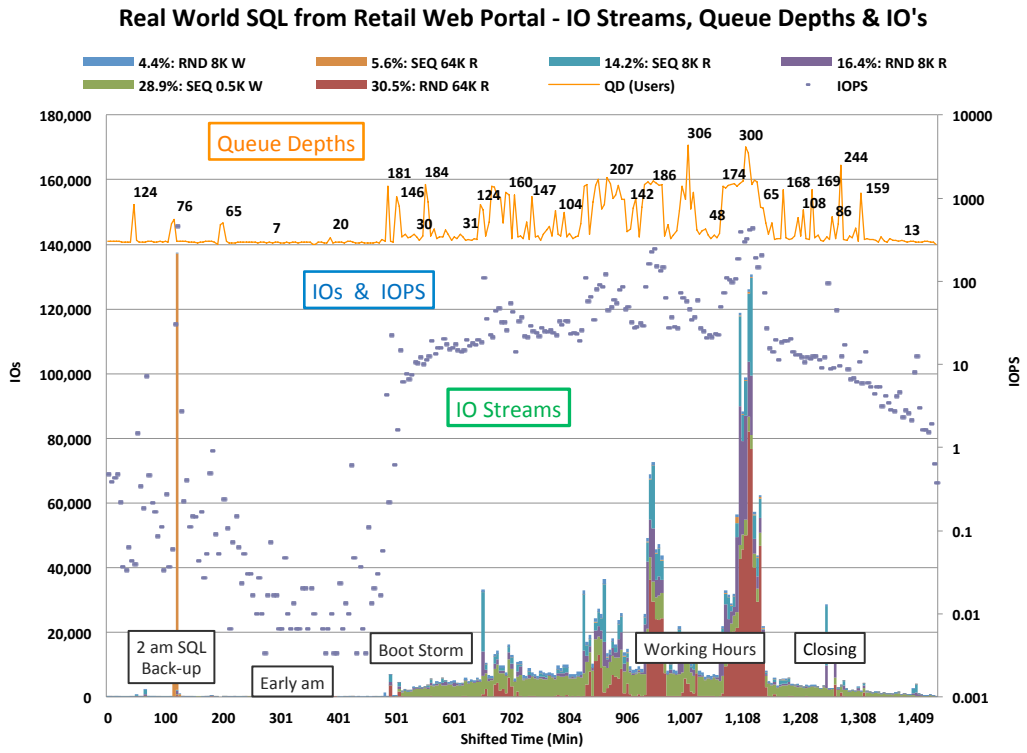

Real World Workloads:

- Constantly changing combinations of IO Streams & range of QDs
- IO Streams and QDs change with Time, Events and Processes
- IO Rates are throttled by real world Applications and Users

[20] SSSI Reference Capture No. 3

Why are RWSWs Important?

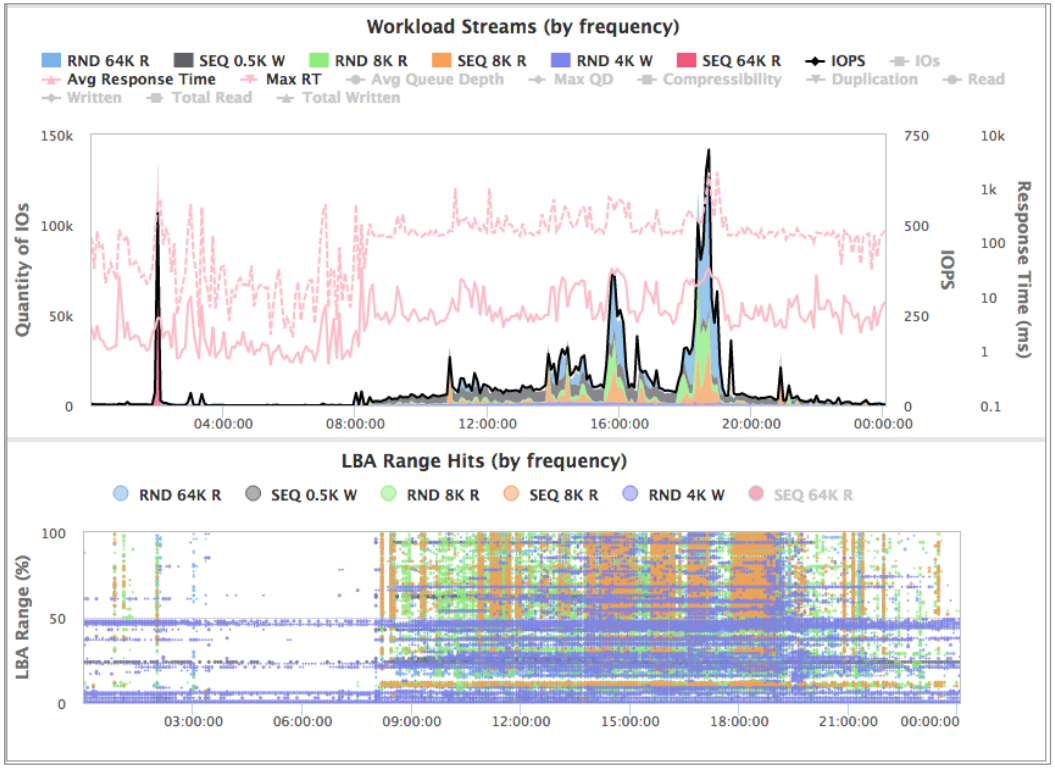

Understanding RWSWs Makes a Difference:

- Know what, where & when IO Streams actually occur
- Observe in-situ performance during the IO Capture process
- Evaluate RWSWs for:
  - software optimization
  - dev ops
  - load balancing
  - Interoperability
  - failure analysis replay
  - storage server & SSD qualification

IO Captures

IO Captures:

1. What is an IO Capture?
2. What is an IO Stream?
3. What is an IO Capture Step?
4. What is an IO Stream Map?
5. IO Capture Tools
6. SSSI Reference Workloads – TestMyWorkload.com

- An IO Capture is:[21]

  a) The tabulation of IO Stream statistics, observed at a given level in the SW Stack, during the IO Trace Capture period

  b) Comprised of tables of IO statistics and metrics in binary data form – no private or personal data is captured

  c) Derived from continuous IO Trace data that is parsed into ***Steps*** for visualization, analytics & playback

  d) Comprised of IO Capture Steps that allow for flexibility in visualization granularity, capture length & file size

[21] RWSW PTS v1.0.7 – page 15

What is an IO Stream?

IO Captures:

1. What is an IO Capture?
2. What is an IO Stream?
3. What is an IO Capture Step?
4. What is an IO Stream Map?
5. IO Capture Tools
6. SSSI Reference Workloads – TestMyWorkload.com

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 is a distinct:[22]
  - a) Read or Write IO (Input / Output) Operation
  - b) RND or SEQ Access
  - c) Data Transfer Size (Block Size)
- A single IO Stream can occur many times during a single IO Capture step
- Other metrics and data are associated with IO Streams (such as Response Times, Process IDs, Queue Depths)

[22] RWSW PTS v1.0.7 – page 30

What are IO Capture Steps?

IO Captures:

1. What is an IO Capture?
2. What is an IO Stream?
3. What are IO Capture Steps?
4. What is an IO Stream Map?
5. IO Capture Tools
6. SSSI Reference Workloads – TestMyWorkload.com

```
### Replay representation for workload 'Cumulative Workload'

## Step 1
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455580803.613    302.219  1233      9192960  0.15 5      2.71 125.89    4.97 44.9683
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
SEQ 0.5K W    SEQ 512  W    4.87 60    0.33 30720
SEQ 8K R      SEQ 8192       R    0.08 1     0.09 8192
RND 4K W      RND      4096      W    47.61      587  26.15      2404352
RND 8K R      RND      8192      R    1.22 15    1.34 122880

## Step 2
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455581105.832    300.547   798  4341248  0.20 9      1.66 33.66     4.99 41.3918
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
RND 4K W      RND      4096      W    53.26      425  40.10      1740800
SEQ 0.5K W    SEQ 512  W    13.41      107  1.26 54784
RND 8K R      RND      8192      R    0.13 1     0.19 8192
RND 64K R     RND      65536     R    0.13 1     1.51 65536

## Step 3
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455581406.379    300.656   761  9361920  0.23 9      2.93 163.45    4.69 42.0124
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
SEQ 0.5K W    SEQ 512  W    15.37      117  0.64 59904
RND 4K W      RND      4096      W    50.46      384  16.80      1572864
RND 64K R     RND      65536     R    0.66 5     3.50 327680

## Step 4
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455581707.035    300.578   726  4136448  0.21 10     1.71 76.96     5.28 35.9942
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
RND 4K W      RND      4096      W    50.41      366  36.24      1499136
SEQ 0.5K W    SEQ 512  W    10.47      76   0.94 38912
RND 8K R      RND      8192      R    0.69 5     0.99 40960
SEQ 8K R      SEQ 8192       R    0.96 7     1.39 57344

## Step 5
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455582007.613    300.578   671  3803136  0.31 10     1.60 88.29     5.61 42.4012
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
RND 4K W      RND      4096      W    58.27      391  42.11      1601536
SEQ 0.5K W    SEQ 512  W    8.49 57    0.77 29184
RND 8K R      RND      8192      R    0.15 1     0.22 8192

## Step 6
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455582308.207    300.594   635  3642368  0.20 9      1.21 18.16     5.32 43.0401
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
RND 4K W      RND      4096      W    63.62      404  45.43      1654784
SEQ 0.5K W    SEQ 512  W    0.94 6     0.08 3072

## Step 7
# Timestamp  Delta (sec)    IOs  Amount (bytes)     Avg QD  Max QD  Avg RT (ms)   Max RT (ms)   Compress.Ratio    Dupl.Ratio (%)
1455582608.801    300.656   521  3213312  0.20 8      1.06 27.00     5.98 40.5197
# Stream  RND/SEQ      Size (bytes)    R/W Quantity (%)   Quantity (IOs)  Amount (%)    Amount (bytes)
RND 4K W      RND      4096      W    64.30      335  42.70      1372160
SEQ 0.5K W    SEQ 512  W    0.38 2     0.03 1024
```

IO Capture Steps: Capture Steps Aggregate IO Stream Statistics

- IO Capture Steps are:[23]
  a) The aggregation of IOs and metrics into discrete time intervals
  b) Metrics that are averaged over the interval and reported as steps
- Step Resolution:
  a) Can be widened to optimize file size for long duration captures
  b) Can be narrowed to observe IO Bursts, Disk Utilization, IO Sequentiality and Quality of Service
- Unlike continuous IO Trace data, IO Captures are a series of discrete time interval steps
- IO Capture Steps may appear continuous or discrete depending on temporal resolution

[23] RWSW PTS v1.0.7 – page 15

Flash Memory Summit

What is an IO Stream Map?

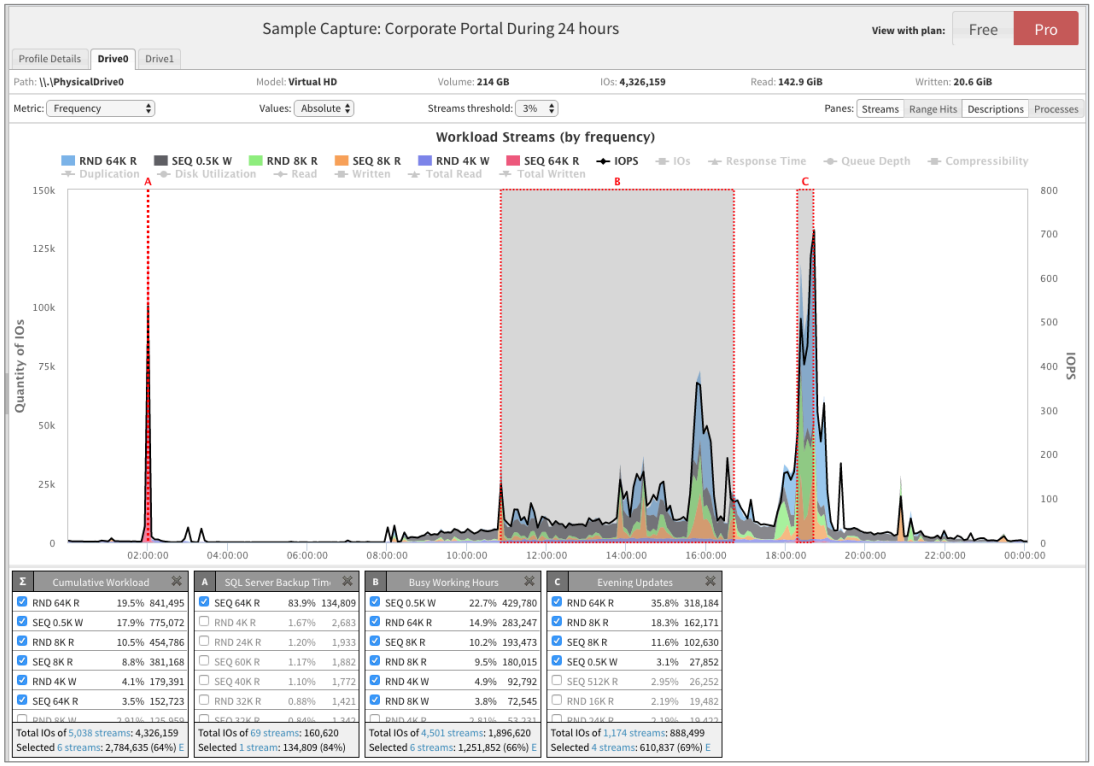

IO Stream Maps are visual representations of IO Capture Steps where:[24]

- Each IO Step is plotted as a time point on the x-axis showing IO Streams by Frequency (IOPS) or Amount Transferred (MB/s)
- IO Metrics are presented as data series along the y-axis
- Metrics include:
  - Process IDs (PID) for each IO
  - IO Stream Composition
  - IOPS, MB/s, RTs, QDs
- Workloads can be parsed by Time, Event, PID, IO Stream Threshold and other criteria

[24] RWSW PTS v1.0.7 – page 30

What are some IO Capture Tools?

IO Captures:

1. What is an IO Capture?
2. What is an IO Stream?
3. What are IO Capture Steps?
4. What is an IO Stream Map?
5. IO Capture Tools
6. SSSI Reference Workloads – TestMyWorkload.com

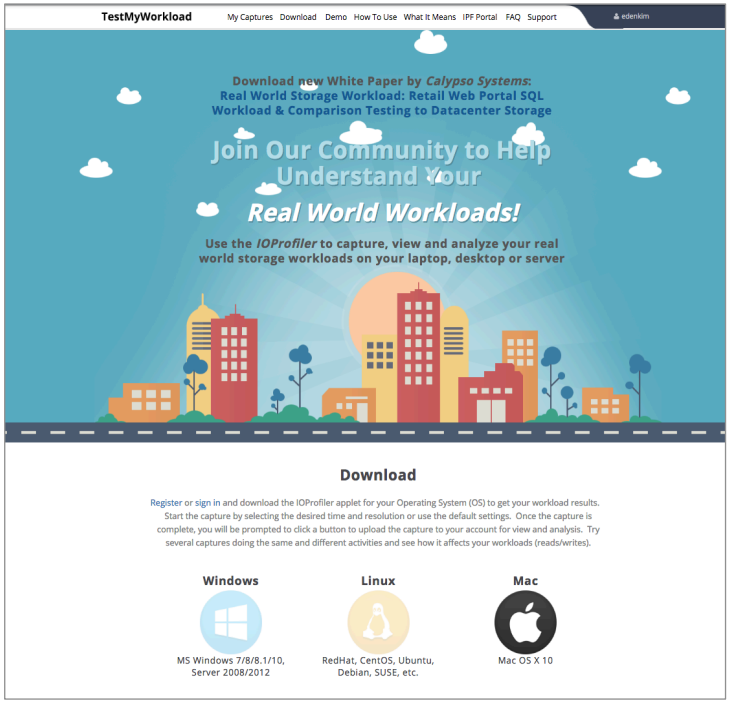

www.TestMyWorkload.com

- IO Capture tools are:[25]
  a) Public or private, fee or free
  b) Specific for OS and software layer
  c) Designed to capture IO traffic
  d) Capture various IO metrics
- Free tools include blk-trace for Linux, Perfmon for Windows and IOProfiler for cross OS platform
- IOProfiler IO Capture Applets are:
  a) Free at TestMyWorkload.com
  b) Windows, Linux or MacOS
  c) File System or Block IO layer
  d) Free upload and visualization at TestMyWorkload.com

[25] RWSW PTS v1.0.7 – page 27

SSSI Reference IO Capture Workloads

IO Captures:

1. What is an IO Capture?
2. What is an IO Stream?
3. What are IO Capture Steps?
4. What is an IO Stream Map?
5. IO Capture Tools
6. SSSI Reference Workloads – TestMyWorkload.com

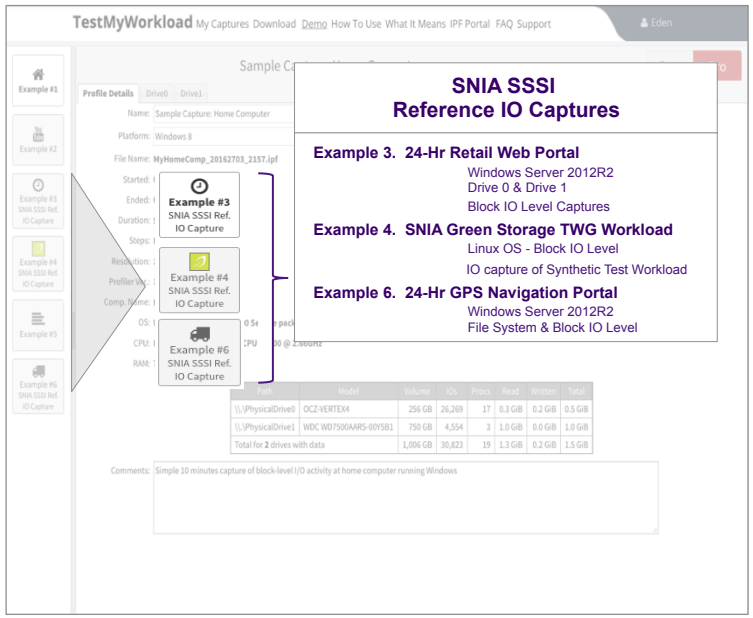

www.testmyworkload.com/info/demo

- SSSI Reference IO Captures at TestMyWorkload.com:
  1. 24-Hr Retail Web Portal
  2. SNIA Green Storage TWG Workload
  3. 24-Hr GPS Navigation Portal
- Free Data Analytics
- Free Export of IO Capture Steps for use with 3d Party software tools
- Additional Workloads posted with ongoing SSS TWG research

Analysis - Overview

Analysis:

1. Overview
2. IO Stream Composition
3. IO Stream PIDs
4. Compare File System Drive C & Drive 0
5. Compare Frequency (IOPS) v Amount Transferred (MB/s)
6. Compare Block IO on Drive 0 & Drive 1
7. LBA Range Hit Map
8. Examples: Data Analytics

- Workload Visualization with IO Stream Maps:
  1. IO Stream Maps present the IO Streams and Metrics that occur during the IO Capture
  2. Resolution can be adjusted to show long duration captures with reasonable file sizes
  3. Fine grain resolution can enable micro second analysis of IO Capture events
- Identification of IO Process IDs:
  1. IO PIDs show IO association with individual processes (applications)
  2. PID enable load balancing, interoperability and failure analysis
- Observation of Specific Software layer IO Traffic:
  1. Enables validation of software optimizations
  2. Confirms how IO Streams change as they traverse the software stack
  3. Allows tracking of individual IOs to storage architecture layers and logical storage

Flash Memory Summit

IO Stream Composition

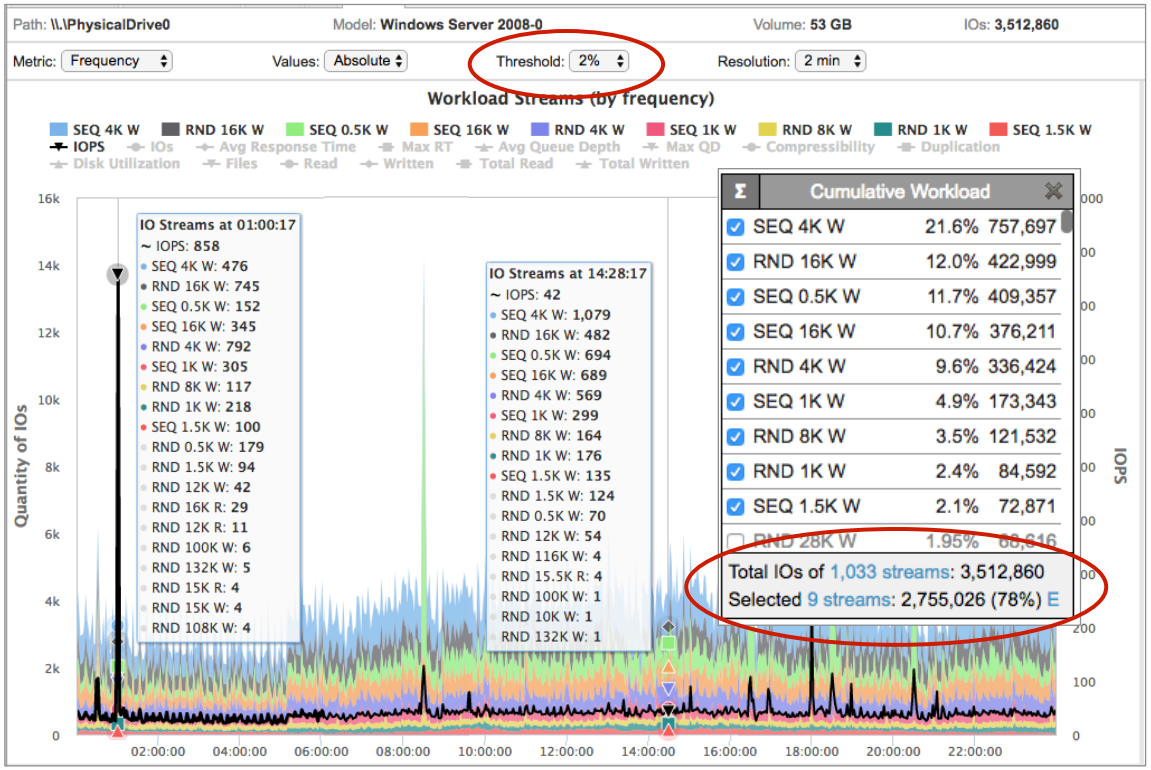

IO Stream Map displays IO Stream Composition for each Step

- Each x-axis point is an IO Step
- Each IO Step is comprised of a unique combination of IO Streams
- Cumulative Workload:
  - Presents IO Streams by percent occurrence over the entire IO Capture duration
  - The IO Streams presented can be filtered by the ‘Threshold’ value
  - 2% Threshold shows all IO Streams that occur 2% or more of the time over the entire capture
  - At 2% Threshold, 9 IO Streams represents 78% of the total IO Streams (1,033 IO Streams)

IO Streams - Process IDs

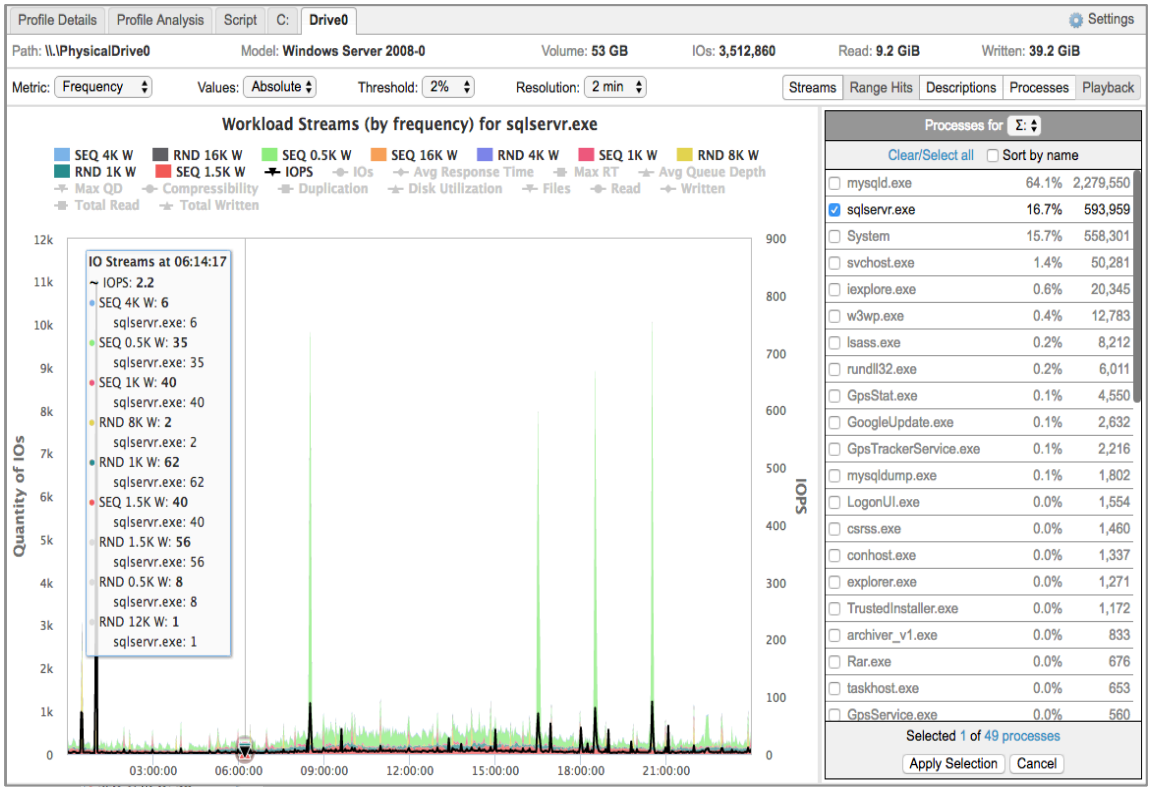

IO Stream Process IDs (PID):

- PIDs can be shown for each IO
- Cumulative Workload shows 49 PIDs
- IO Stream Map can be filtered by:
  - IO Streams by PID – e.g. SQL IOs
  - Time point or Range
  - Event
- IO Streams at Time: 06:16:17
  - IOs can show multiple PIDs (sqlserver.exe, mysqld.exe, lsass.exe)
  - IOs can show single PIDs (sqlserver.exe)

Cumulative Workload - by Frequency v Amount Transferred

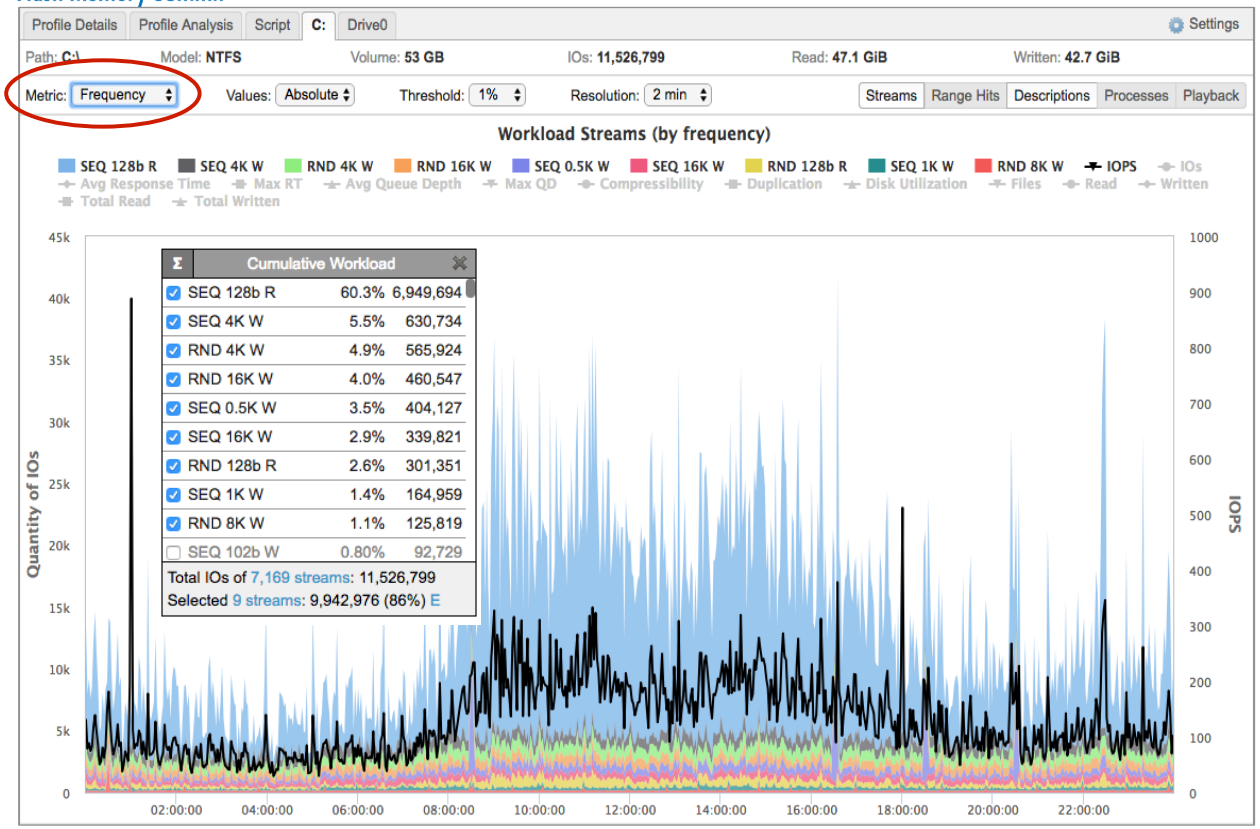

24-Hour GPS Navigation Portal

Example No. 6 TestMyWorkload.com

- IO Streams by Frequency (IOPS)
  1. Drive C File System level
  2. 24-Hour Capture
- Workload Settings
  1. IO Streams Threshold – 1%
  2. Temporal Resolution – 2 Min
- Cumulative Workload Box:
  1. IO Streams at 1% Threshold
  2. 9 IO Streams/86% of Total IOs
  3. 9.942M IOs of 11.526M IOs
- IO Stream Combinations vary for each IO Step

Cumulative Workload – Frequency v Amount Transferred

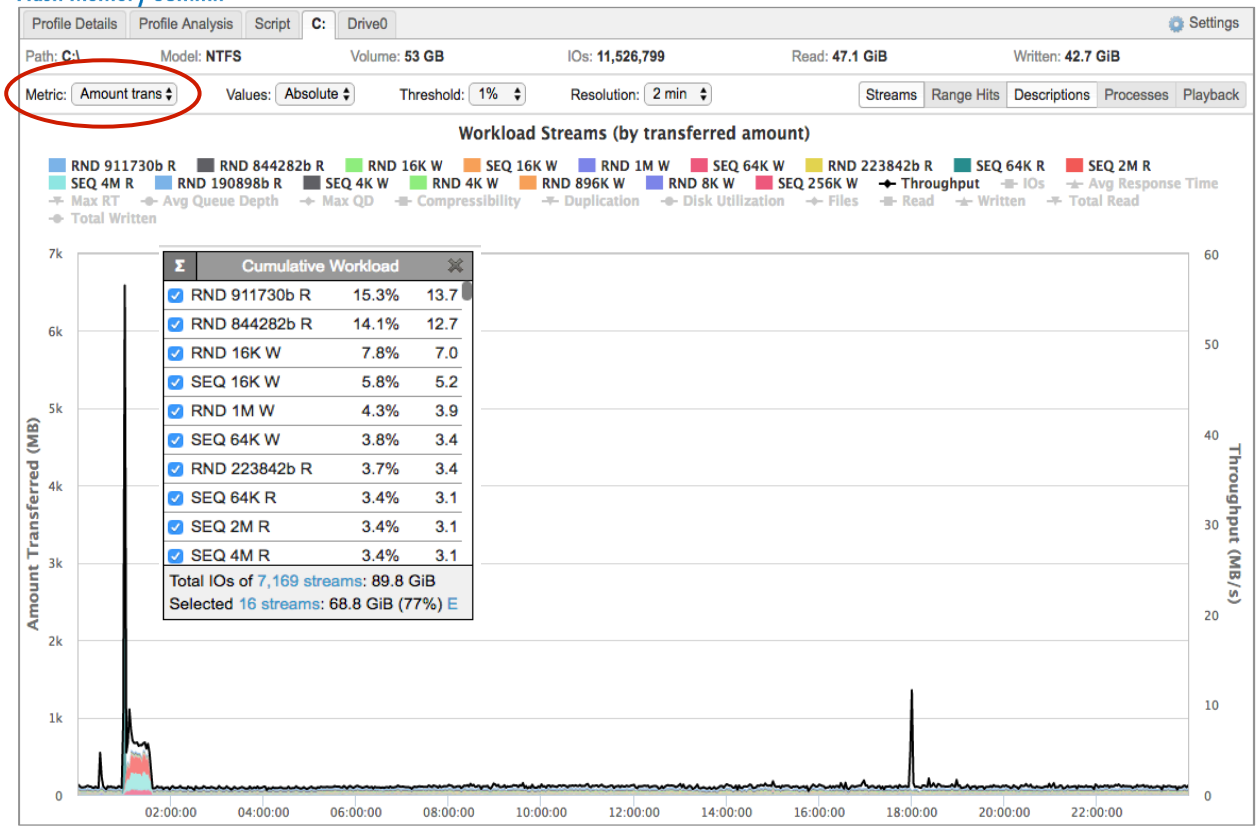

24-Hour GPS Navigation Portal

Example No. 6 TestMyWorkload.com

- IO Streams by Amount Transferred (MB/s)
  1. Drive C - File System level
  2. 24-Hour Capture
- Workload Settings
  1. IO Streams Threshold – 1%
  2. Temporal Resolution – 2 Min
- Cumulative Workload Box:
  1. IO Streams at 1% Threshold
  2. 16 IO Streams/69.8% of Total IOs
  3. 68.8 GiB of 89.8 GiB
  4. Note: 911,730b = 1,780.7KB

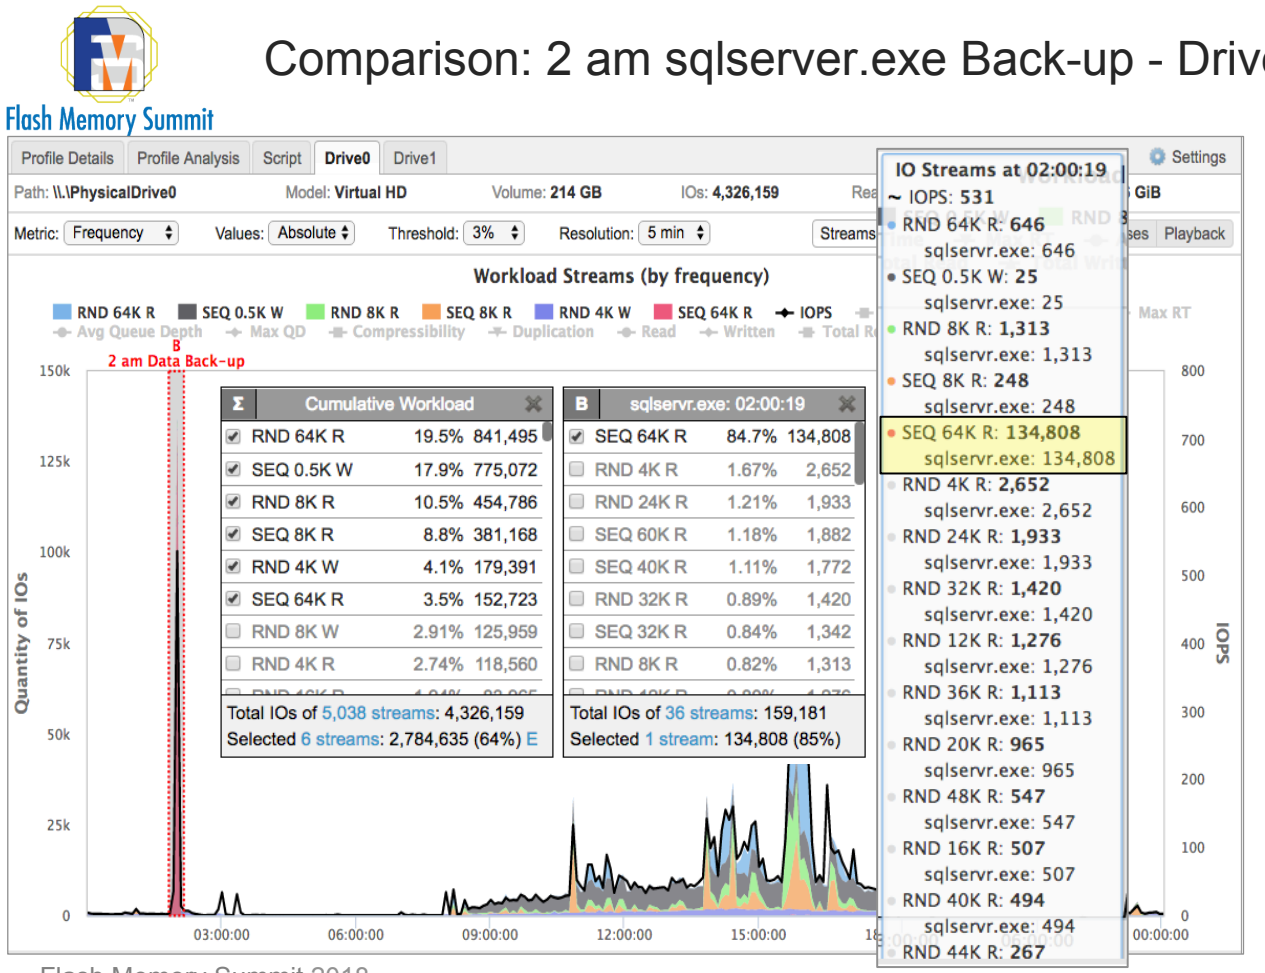

Comparison: 2 am sqlserver.exe Back-up - Drive 0 & Drive 1

24-Hour Retail Web Portal

Example No. 3 TestMyWorkload.com

- Cumulative Workload by IO Stream Threshold
  1. Drive 0 – Block IO
  2. Resolution – 5 min
  3. IO Streams Threshold – 3%
  4. IO Streams – 6 Streams; 64%
- 2 am Data Back-up by IO Stream Threshold
  1. IO Stream Threshold – 3%
  2. IO Streams – 1 Stream; 81.1%
  3. SEQ 64K R, 134,808 IOs
- 2 am by Process ID (sqlserver.exe)
  1. PID - sqlserver.exe
  2. SEQ 64K R, 134,808 IOs

LBA Range Hit Map: SEQ IOs

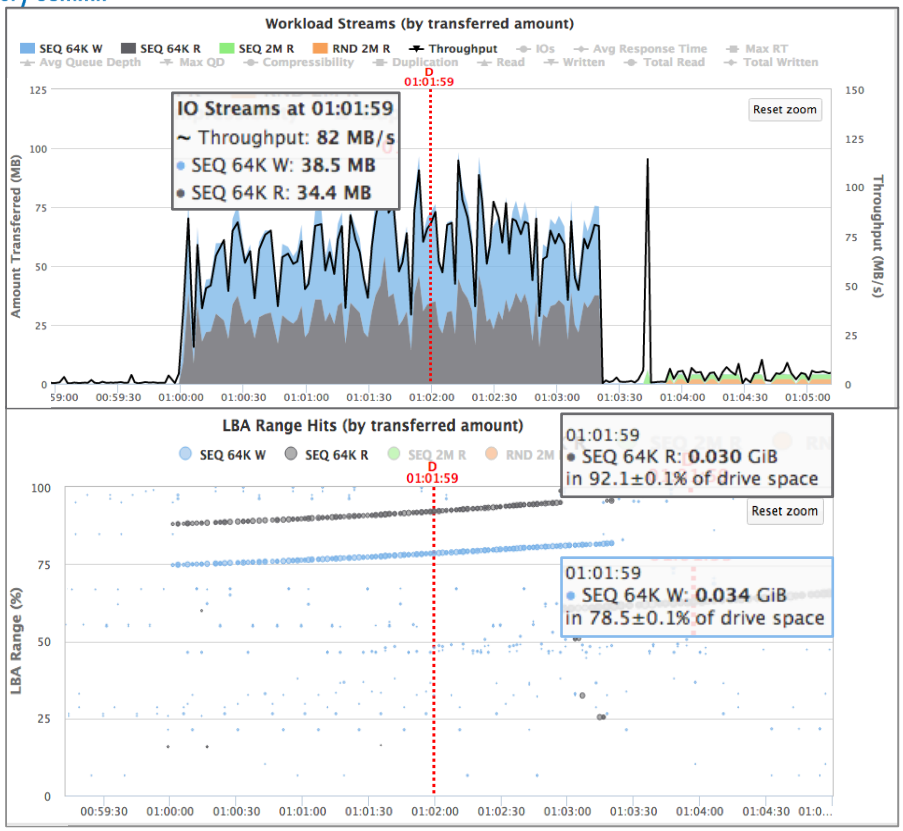

Drive C: Back-up Activity

- IO Stream Map at Time = 1:01:59
  1. Throughput: 82 MB/s
  2. SEQ 64K W: 38.5 MB
  3. SEQ 64K R: 34.4 MB
  4. Other IOs: 9.1 MB
- LBA Range Hit Map at Time 1:01:59
  1. SEQ 64K R: 0.030 GiB at LBA 92.1% +/- 0.1%
  2. SEQ 64K W: 0.034 GiB at LBA 78.5% +/- 0 .1%
- LBA Map Range Hits
  1. Gray Line = 64K R
  2. Blue Line = 64K W
  3. Diagonal Lines indicate SEQ Accesses

Data Analytics: IOPS, Amount Transferred, RTs & Queue Depths

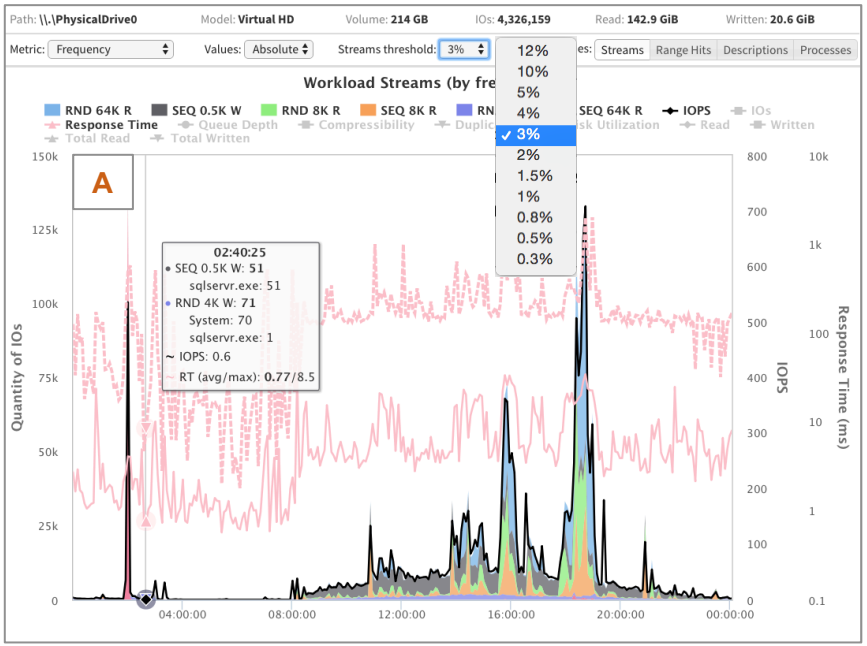

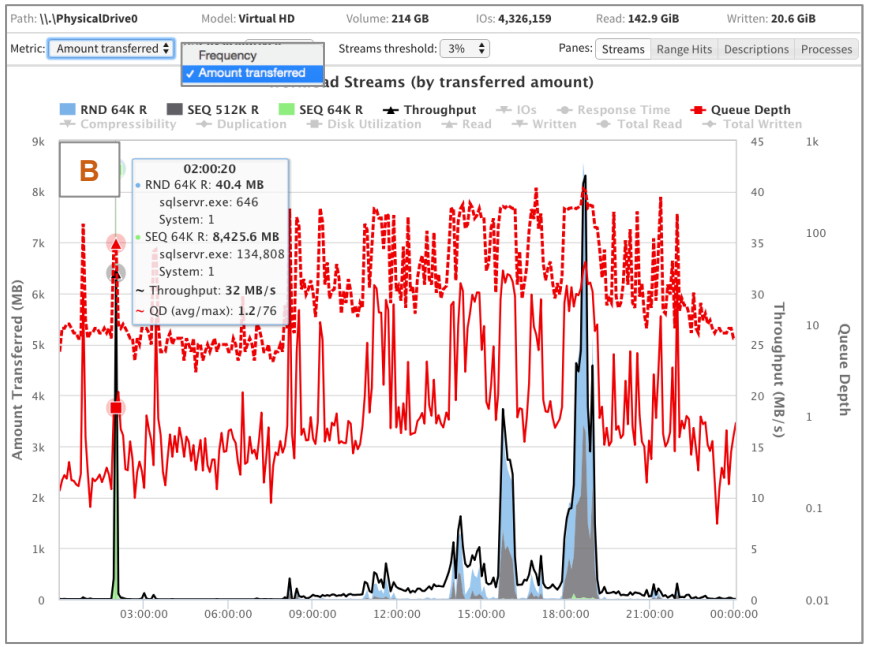

A **IO Stream Map by Frequency – IO Rate in IOPS**

IOPS, Ave/Max Response Times, 3% Threshold

B **IO Stream Map by Amount Transferred – Throughput in MB/s**

MB/s, Ave/Max Queue Depth

Data Analytics: Compression & Duplication Ratios

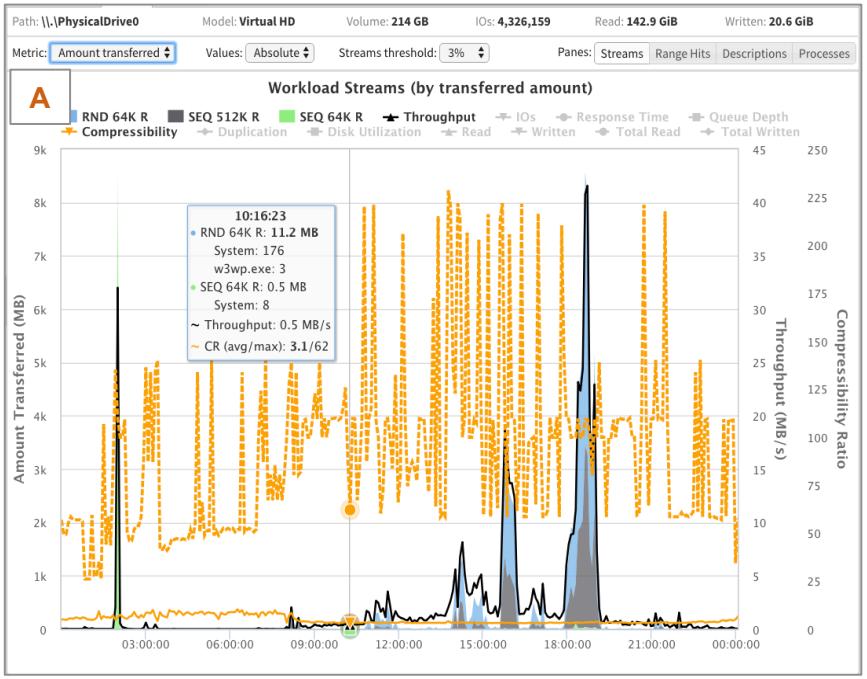

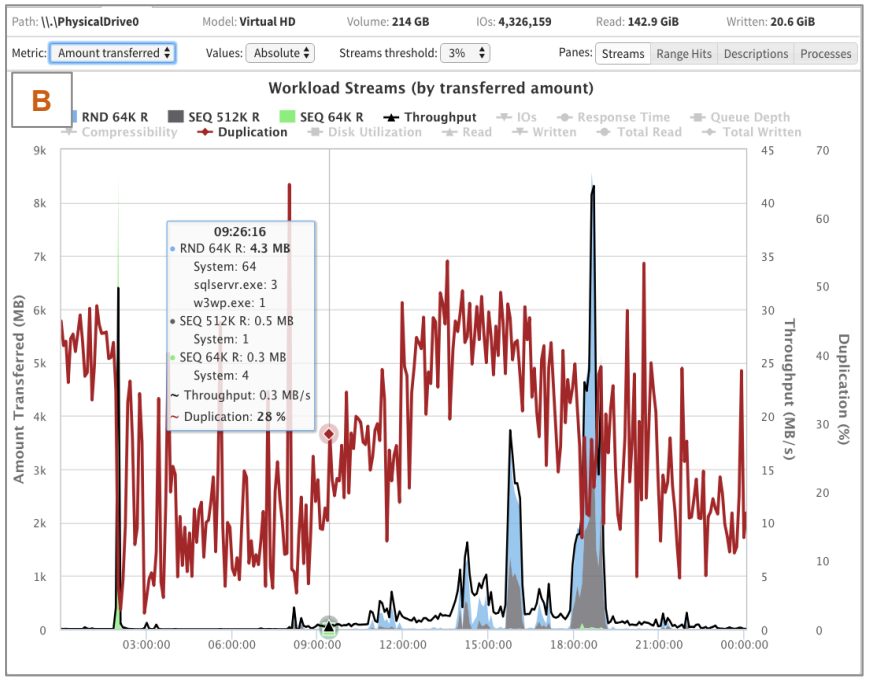

A Compression Ratio: How much more compressible is data

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

Data Analytics: Disk Utilization & IO Sequentiality

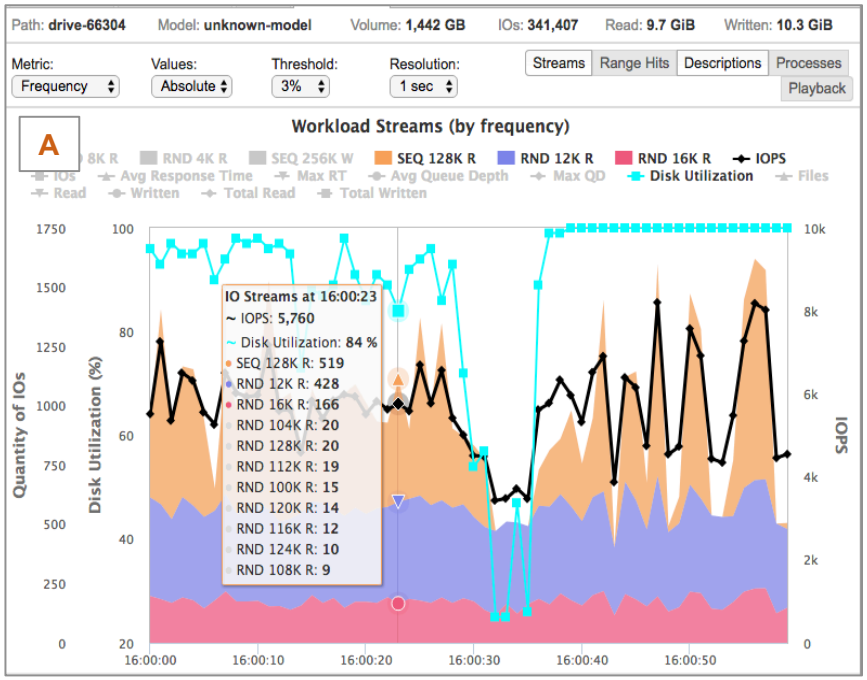

A Disk Utilization: IOs and Disk IO Idle Times

Disk Utilization of 84% = 16% Disk IO Idle

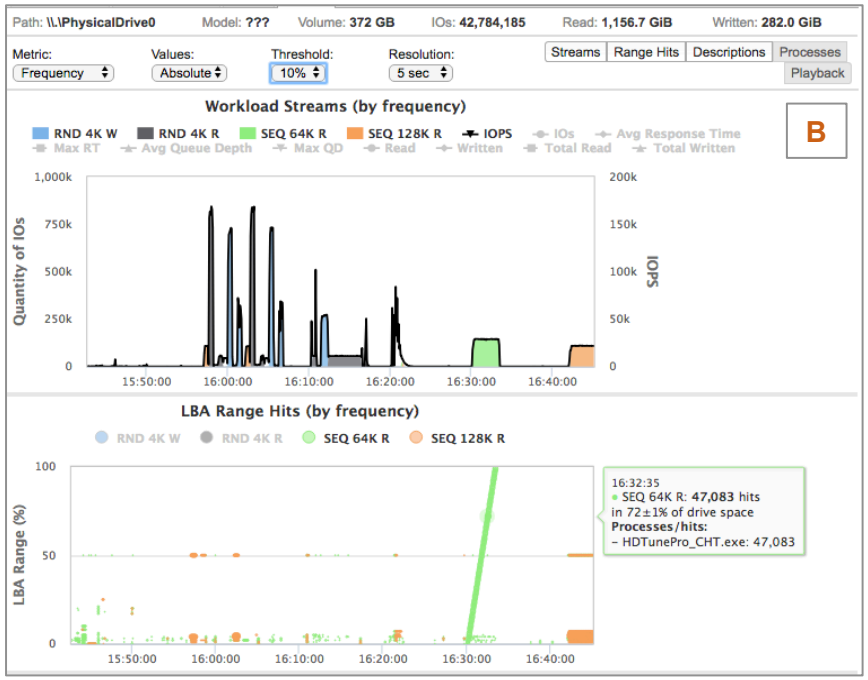

B IO Sequentiality – Adjacent LBA Range Hits

Diagonal LBA Range Hit lines indicate Sequential IOs

Data Analytics: Temporal Granularity & Step Resolution

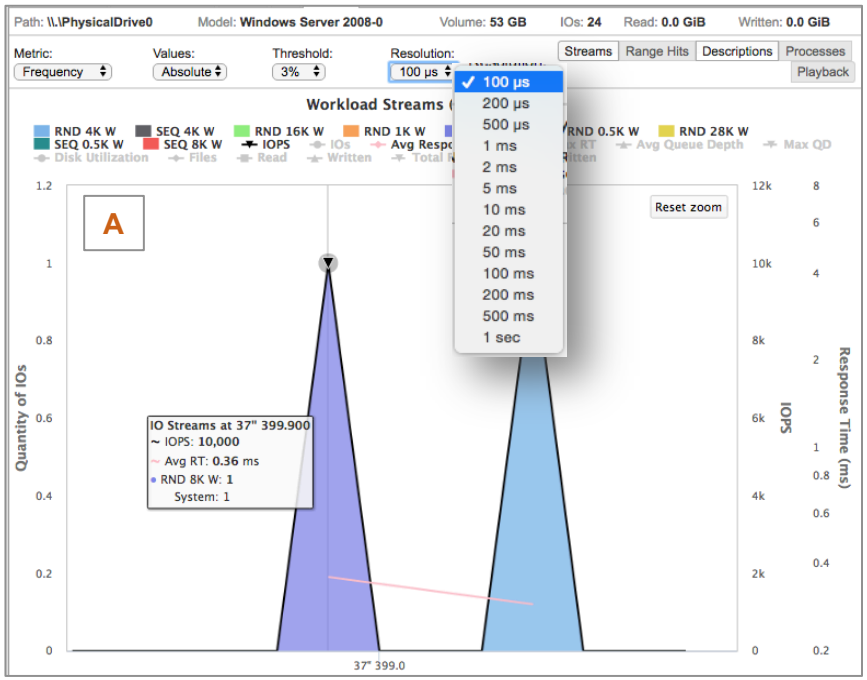

A Fine Grain: 100 uS; 1 min Capture; Single IOs

IO Bursts, Disk Utilization, IO Specificity

B Coarse Grain: 1 Min; 24 hr Capture; 100's of IOs

Long Term workload characterization

BREAK

3:30 – 3:45

Flash Memory Summit

Part 3: RWSW PTS v 1.0.7 - Tests

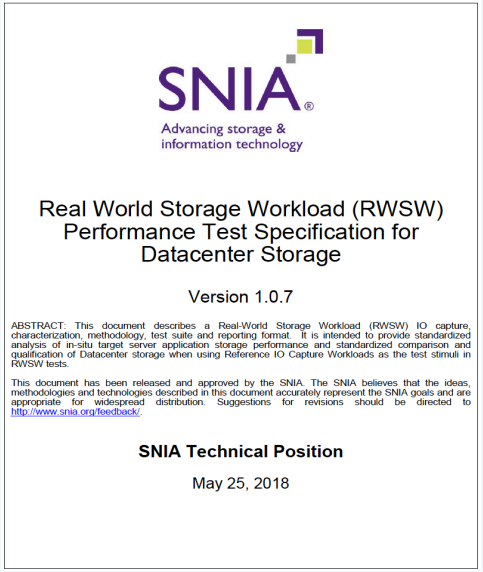
SNIA
Advancing storage & information technology

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

www.snia.org/rwsw

- Overview
- RWSW Tests:
  1. Replay
  2. Individual Streams WSAT
  3. Multi-WSAT
  4. RWSW DIRTH

Overview - Key Test Process Concepts[26]

Overview - RWSW Tests:

1. Key Test Process Concepts
2. Define the Test Workload
3. Applied Test Workloads
4. Basic Test Flow
5. Test Settings

- RWSW Tests are intended to:
  1. ***Analyze*** in-situ performance of the target server during the IO Capture
  2. ***Optimize*** Data Center Storage performance
  3. ***Validate & Qualify*** Data Center Storage
- Test Operator shall select and disclose the following:
  1. OS, IO capture tool used & the IO capture layer (File System, Block IO or other)
  2. The RWSW (which can be an SSSI Reference or User selected workload)
  3. The Applied Test Workload (which is filtered/parsed from the RWSW)
  4. The Data Center Storage to be tested (logical storage recognized by OS)
- Test Storage Preparation:
  1. It is recommended to apply the RWSW to Steady State (SS) when possible
  2. It may be impractical to PURGE, Pre-condition (PC) or bring RWSW to SS
  3. Test operators shall select an appropriate PC regime and disclose it in the test results

[26] RWSW PTS v1.0.7 – page 24

Define the Test Workload

Overview - RWSW Tests:

1. Key Test Process Concepts
2. Define the Test Workload
3. Applied Test Workloads
4. Basic Test Flow
5. Test Settings

- Run IO Capture to obtain RWSW at desired software level:[27]
  1. File System
  2. Block IO
  3. User selected (virtualized, cluster, LUN, hypervisor, custom, other)
- Create / Select Real World Storage Workload:
  1. SSSI Reference Workload
  2. User selected/defined RWSW
- Create Applied Test Workload:
  1. Filter/parse the selected IO Capture as desired (by Time, Event, PID, etc.)
  2. Set the IO Stream Threshold
  3. Define the Applied Test Workload & IO Stream Distribution

[27] RWSW PTS v1.0.7 – page 12

Applied Test Workloads

Overview – RWSW Tests:

1. Key Test Process Concepts
2. Define the Test Workload
3. Applied Test Workloads
4. Basic Test Flow
5. Test Settings

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

Total IOs of 1,033 streams: 3,512,860
Selected 9 streams: 2,755,026 (78%) E

Cumulative Workload Selected from SSSI Ref. Capture No. 6
9 IO Streams = 78% of Total IOs

Σ0	Drive0 Cumulative W	
SEQ 4K W	21.6%	27.3
RND 16K W	12.0%	15.4
SEQ 0.5K W	11.7%	14.9
SEQ 16K W	10.7%	13.7
RND 4K W	9.6%	12.2
SEQ 1K W	4.9%	6.3
RND 8K W	3.5%	4.4
RND 1K W	2.41%	3.1
SEQ 1.5K W	2.07%	2.65

Applied Test Workload
9 IO Streams Normalized to equal 100%

Create the Applied Test Workload:[28]

- First, the IO Capture workload is filtered, or parsed, as the user desires – often the Cumulative Workload is selected
- The filtered IO Streams and metrics are used to create an Applied Test Workload
- IO Stream percentages of the desired workload are normalized to equal 100%
- The Applied Test Workload is used in the RWSW tests as specified, but can be:
  1. A replay of the IO Capture sequence
  2. A fixed composite of the IO Streams
  3. Individual IO Streams tested to SS
  4. A fixed composite DIRTH test

[28] RWSW PTS v1.0.7 – page 32

Overview – RWSW Tests:

1. Key Test Process Concepts
2. Define the Test Workload
3. Applied Test Workloads
4. Basic Test Flow
5. Test Settings

Basic Test Flow

- RWSW PTS Basic Test Flow:[29]
  a) PURGE
  b) Apply Pre-Conditioning
  c) Run to Steady State
  d) Apply RWSW
  e) Report Measurements

[29] RWSW PTS v1.0.7 – page 24

Overview – RWSW Tests:

1. Key Test Process Concepts
2. Define the Test Workload
3. Applied Test Workloads
4. Basic Test Flow
5. Test Settings

Test Settings

- PURGE & Pre-Conditioning:[30]
  a) It may be impractical to PURGE and run Pre-conditioning
  b) Target Data Center Storage may be LUNs or other logical storage
- Steady State
  a) Apply RWSW until performance meets the Steady State Criteria
  b) Measurement values should be relatively time invariant
- Active Range: AR=100 for Enterprise; AR=75 for Client
- Data Pattern: DP= Random
- Test Operator may select other test settings so long as settings are Disclosed

[30] RWSW PTS v1.0.7 – page 24

RWSW Tests

RWSW Tests:

1. Target Server Self Test
2. Replay Native
3. Individual Streams WSAT
4. Multi-Stream WSAT
5. RWSW DIRTH

Note:

Sample test data in this section is based on SSSI Reference IO Capture Workload No. 3 – 24-Hour SQL Server Retail Web Portal. Demo and example located at www.testmyworkload.com/info/demo

- Purpose of Tests:[31]
  1. To analyze in-situ target server storage performance using IO Captures
  2. To characterize and test Data Center storage performance using RWSWs
- RWSW PTS sets forth five tests:
  1. Target Server Self Test – Pseudo test that reports metrics during IO Capture process
  2. Replay Native – reproduces the sequence and combinations of IO Streams for testing
  3. Individual Streams WSAT – tests each individual IO Stream to Steady State
  4. Multi-Stream WSAT – applies the fixed composite of IO Streams for every test step
  5. RWSW DIRTH – applies the fixed composite IO Stream steps across a range of OIO
- Results Reporting:
  1. All Test Settings, Workload Composition and Test Variables shall be disclosed
  2. Additional settings, metrics & reports are Optional
  3. Test Operator may select SSSI Reference Workloads or apply User selected RWSWs

[31] RWSW PTS v1.0.7 – page 13

Target Server Self-Test: In Situ Performance

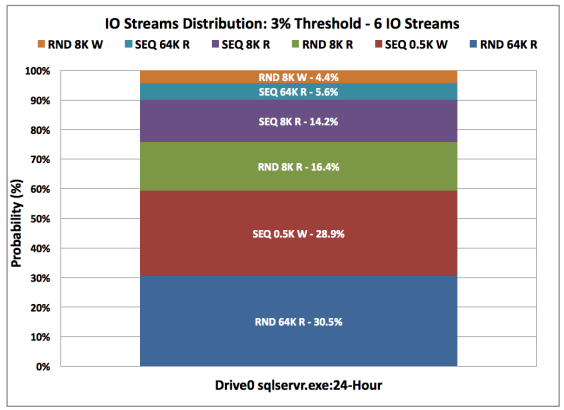

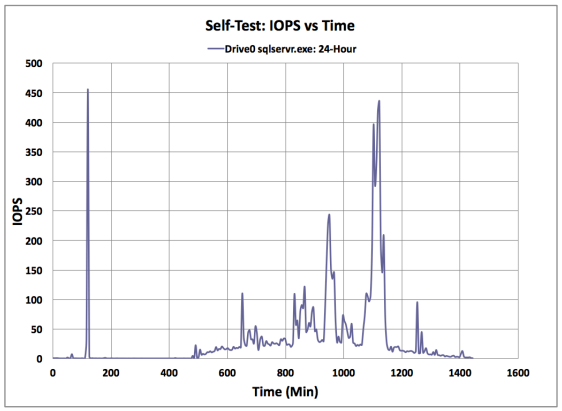

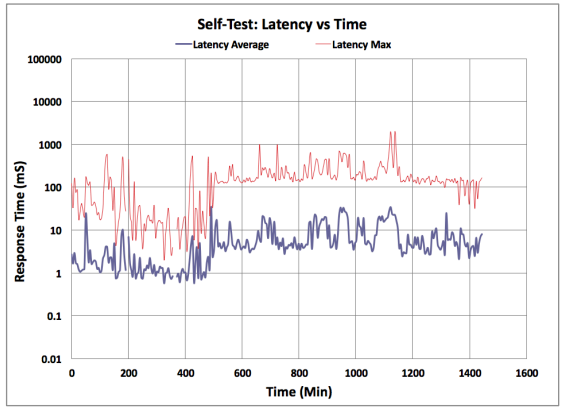

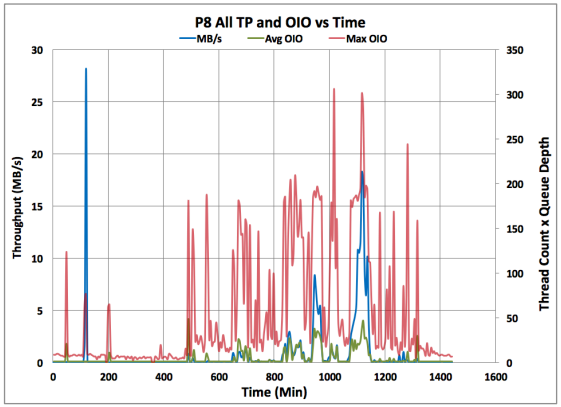

Target Server Self-Test presents performance of the Target Server during the IO Capture Process[32]

- Target Server Self-Test:
  1. Is a pseudo-test, not an actual test
  2. Is a compilation of measurements based on metrics taken during the IO Capture
  3. Server Performance is throttled by the Users, HW/SW & applications (and thus can be lower than lab tests that use the RWSW as the test workload)
- Target Server Self-Test Reports include:
  1. IO Streams Distribution v Segments
  2. IOPS & TP v Time
  3. Response Time Latency v Time
  4. Outstanding IO (OIO) v Time Thread Count x Queue Depth

[32] RWSW PTS v1.0.7 – page 33

Replay Native Test

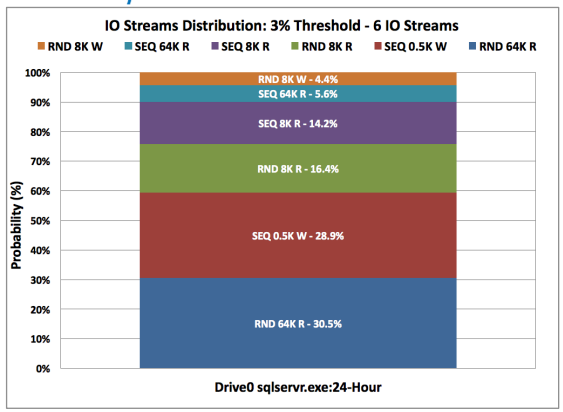

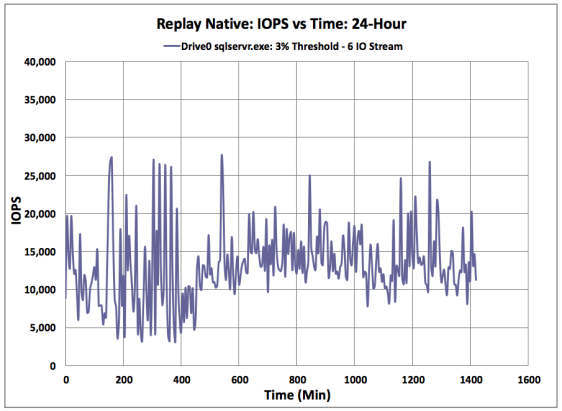

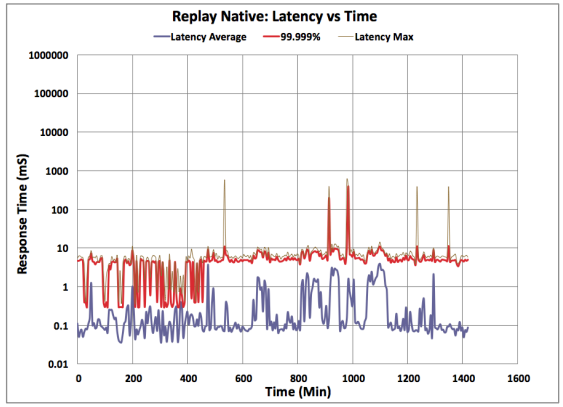

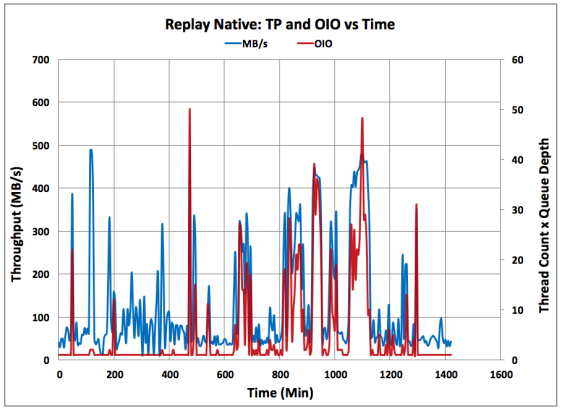

Replay Native reproduces the sequence and combinations of IO Streams & QDs of the IO Capture[33]

- Replay Native Workload:
  1. Sequence of IO Streams & QDs observed during the capture are applied to each test step
  2. RWSW DIRTH test or other RWSW Test recommended to be used as PC as PC & SS are difficult w/ Replay (due to changing step combinations)
- Replay Native Reports include:
  1. IO Streams Distribution v Segments (Cumulative Workload listed as listing every step composition is impractical)
  2. IOPS & TP v Time
  3. Response Time Latency v Time
  4. Outstanding IO (OIO) v Time Thread Count x Queue Depth

[33] RWSW PTS v1.0.7 – page 39

Individual Streams WSAT Test

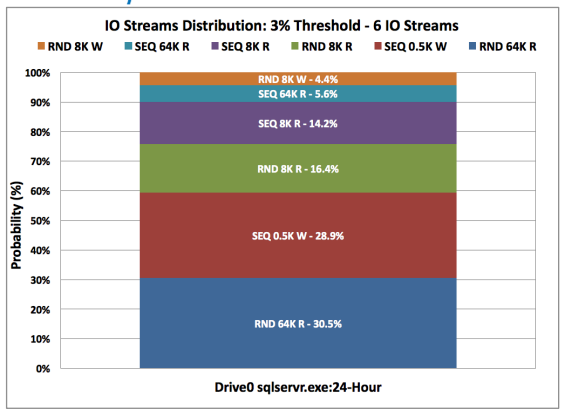

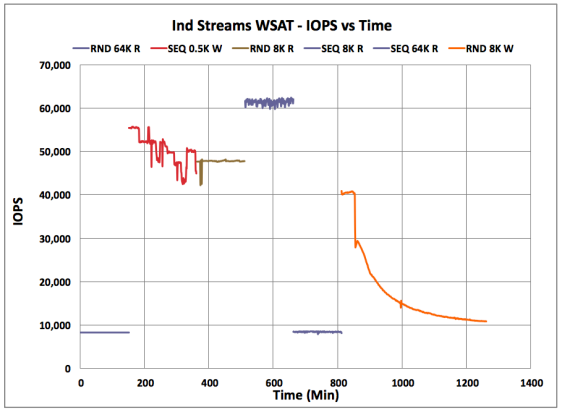

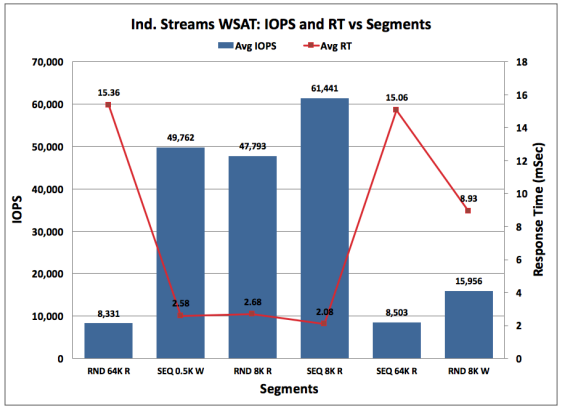

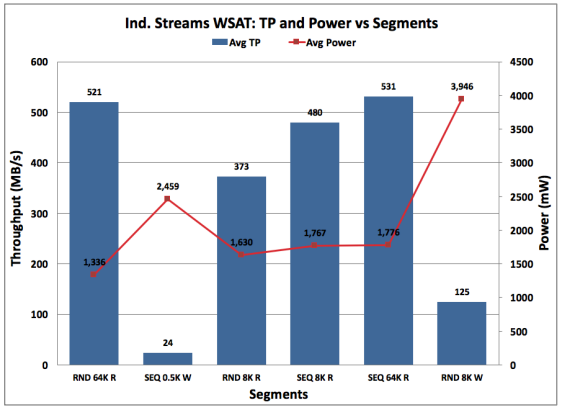

Individual Streams WSAT applies each individual IO Stream to Steady State[34]

- Individual Streams WSAT Workload:
  1. Each IO Stream is applied to SS
  2. Ind. Stream SS can be compared to mfgr benchmark tests
- Individual Streams WSAT Reports include:
  1. IO Streams Distribution v Segments
  2. Segment IOPS v Time (run to SS)
  3. IOPS & Response Time by Segment
  4. Throughput & Power by Segment

[34] RWSW PTS v1.0.7 – page 54

Multi-Streams WSAT Test

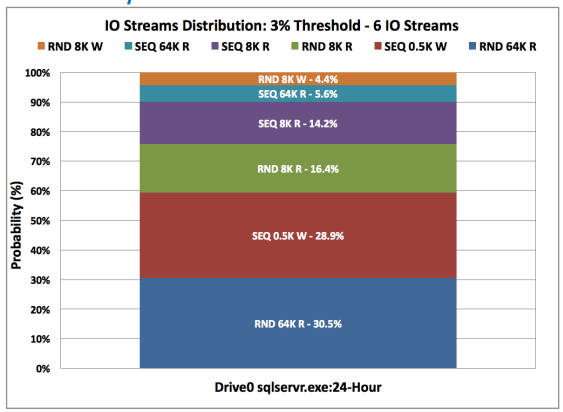

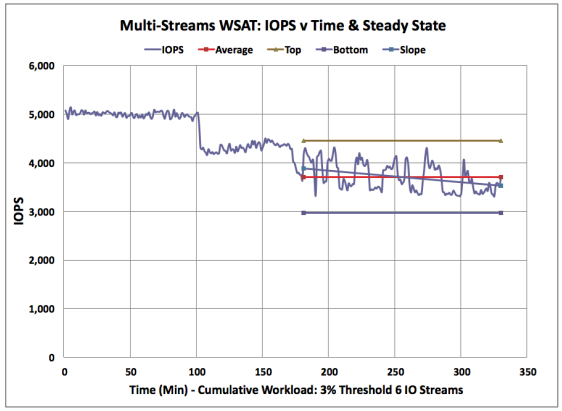

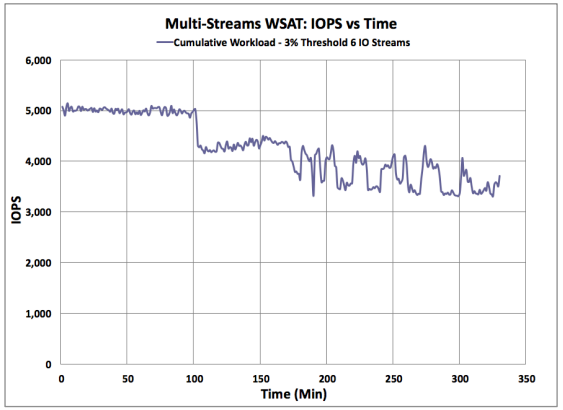

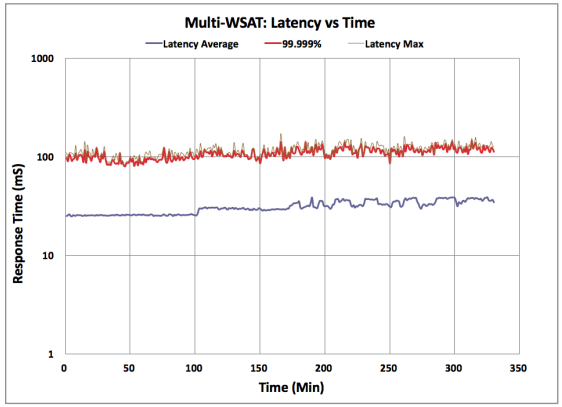

Multi-Streams WSAT applies the fixed composite of IO Streams to Steady State[35]

- Multi-Streams WSAT Workload:
  1. The fixed composite of the IO Streams is applied to SS
  2. PURGE & Pre-conditioning are optional, SS is required
- Multi-Streams WSAT Reports include:
  1. IO Streams Distribution v Segments
  2. Segment IOPS v Time (run to SS)
  3. IOPS & Response Time by Segment
  4. Throughput & Power by Segment

[35] RWSW PTS v1.0.7 – page 48

RWSW DIRTH Test

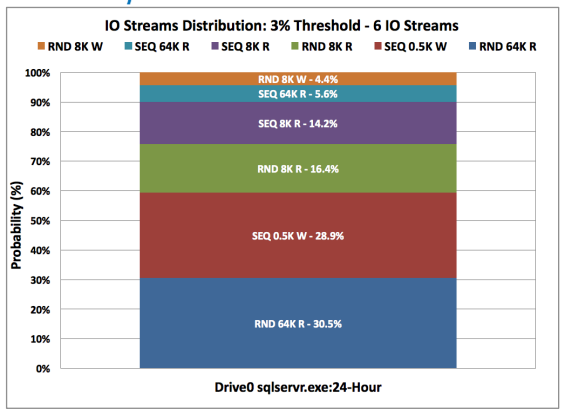

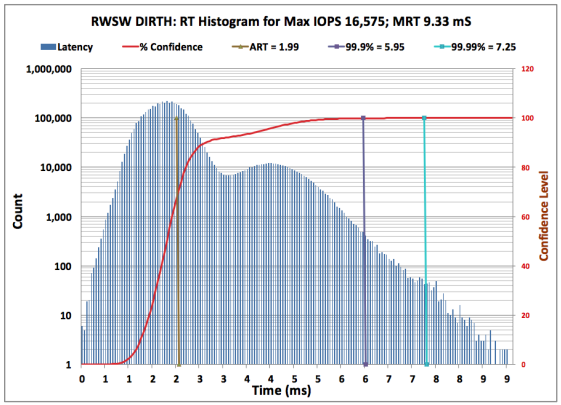

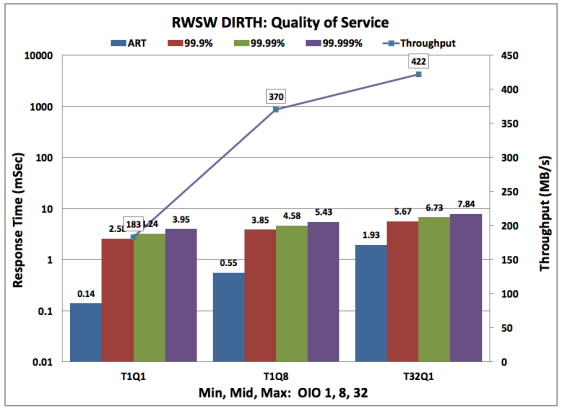

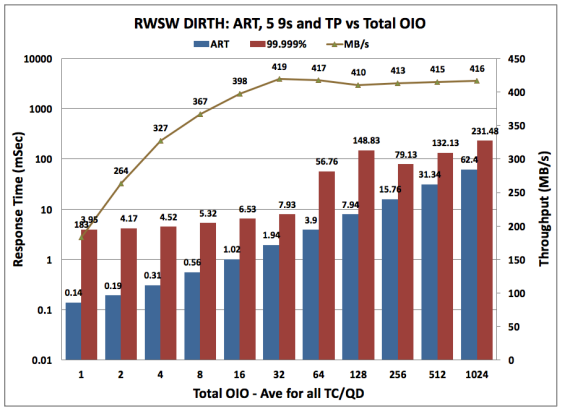

RWSW DIRTH applies the fixed composite of IO Streams to Steady State then varies OIO[36]

- RWSW DIRTH Workload:
  1. Applies the fixed composite of IO Streams to SS followed by OIO loop
  2. Shows IO & Bandwidth saturation across a range of OIO from 1 to 1,024
- RWSW DIRTH Reports include:
  1. IO Streams Distribution v Segments
  2. RT Histogram for Max IOPS 16,575
  3. Quality of Service: OIO 1, 8, 32
  4. ART, 5 9s & TP v Total OIO Range

[36] RWSW PTS v1.0.7 – 63

Flash Memory Summit

Thank You

www.TestMyWorkload.com

edenkim@calypsotesters.com