Education

SNIA Solid State Storage Performance Test Specification

Easen Ho
CTO, Calypso Systems, Inc.

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- SNIA Solid State Storage Performance Test Specification
  - This session will appeal to End Users, OEMs, Drive Manufacturers, System Integrators, as well as those that are seeking a fundamental understanding of the SNIA's Solid State Storage Performance Test Specification.

SSS PTS Enterprise Draft 1.0

- Available now for public review
- Download the spec:
  www.snia.org/tech_activities/publicreview
- Upload Feedback:
  www.snia.org/tech_activities/feedback
- Updates to Spec:
  www.snia.org/forums/sssi

Table of Contents

Motivation

- No Industry Standard
  - No standard methodology, common terminology, nor test environment for measuring SSS performance
- Myriad of Applications on Various Platforms
  - SSS makers/reviewers use different applications, OS and hardware; produces and uses selected metrics
- Market Confusion
  - Difficult to compare test results from different sources; difficult to ensure accuracy in representing SSS products to end users; box-top numbers are pretty useless

SSD Performance Varies Over Time

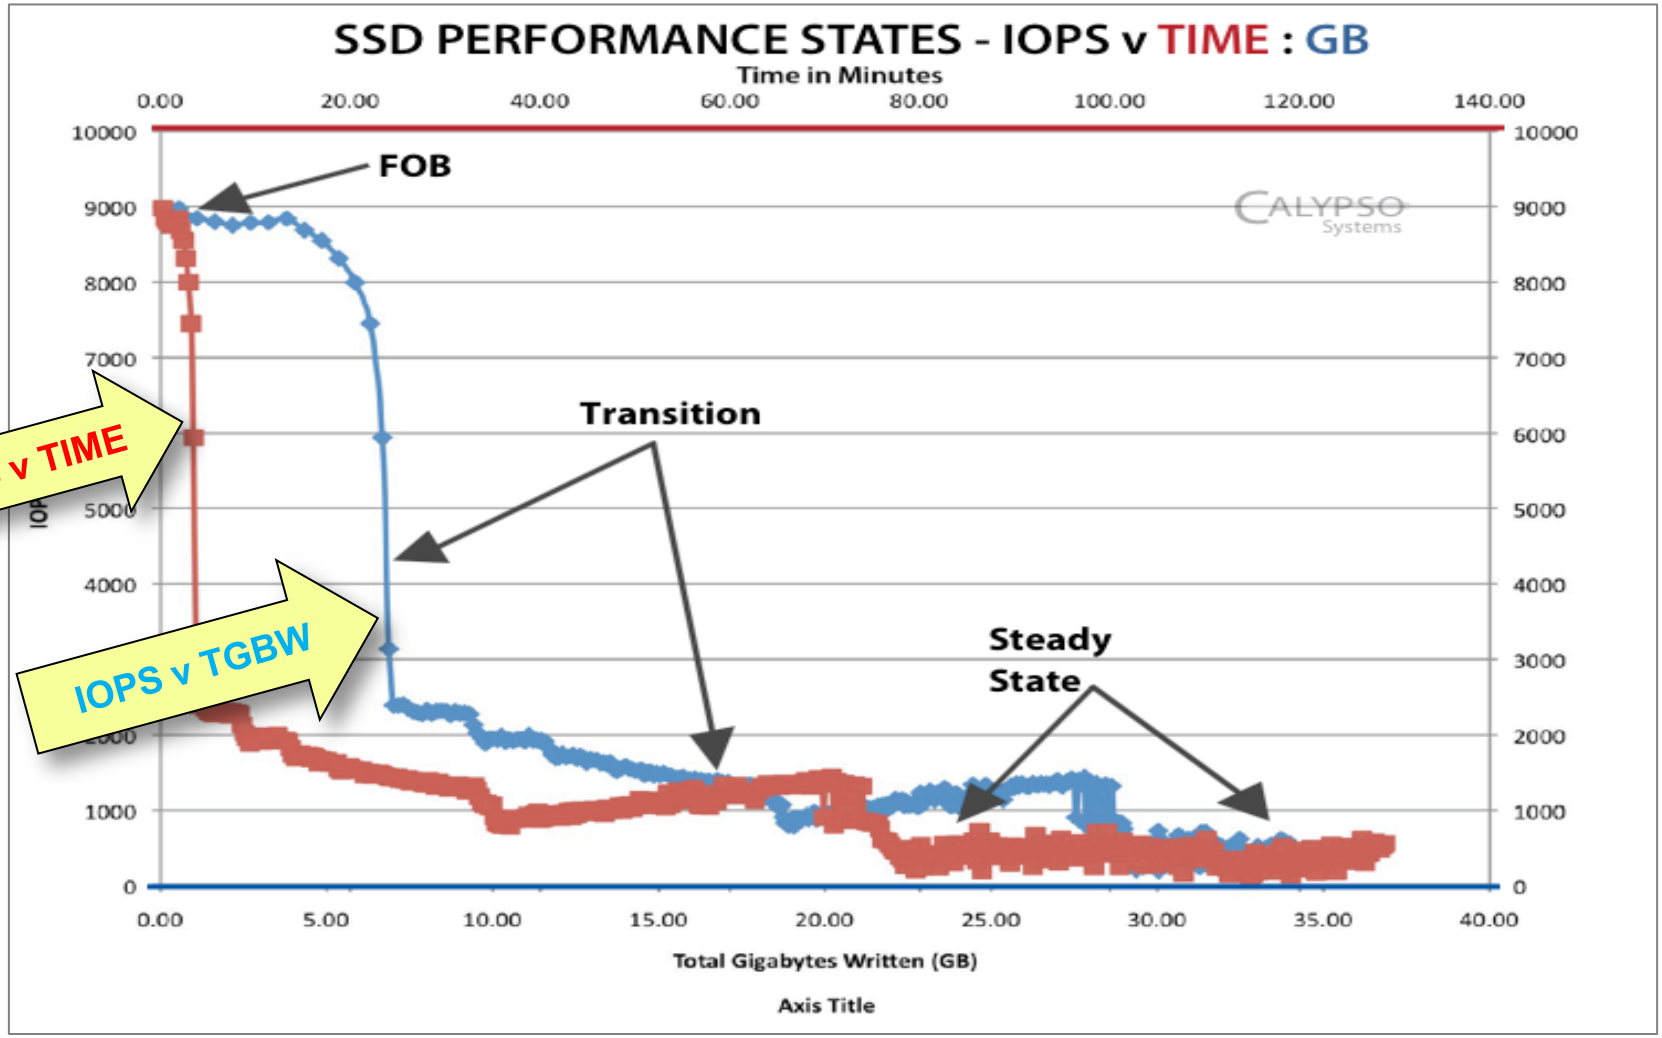

Many Different Kinds of Behaviors

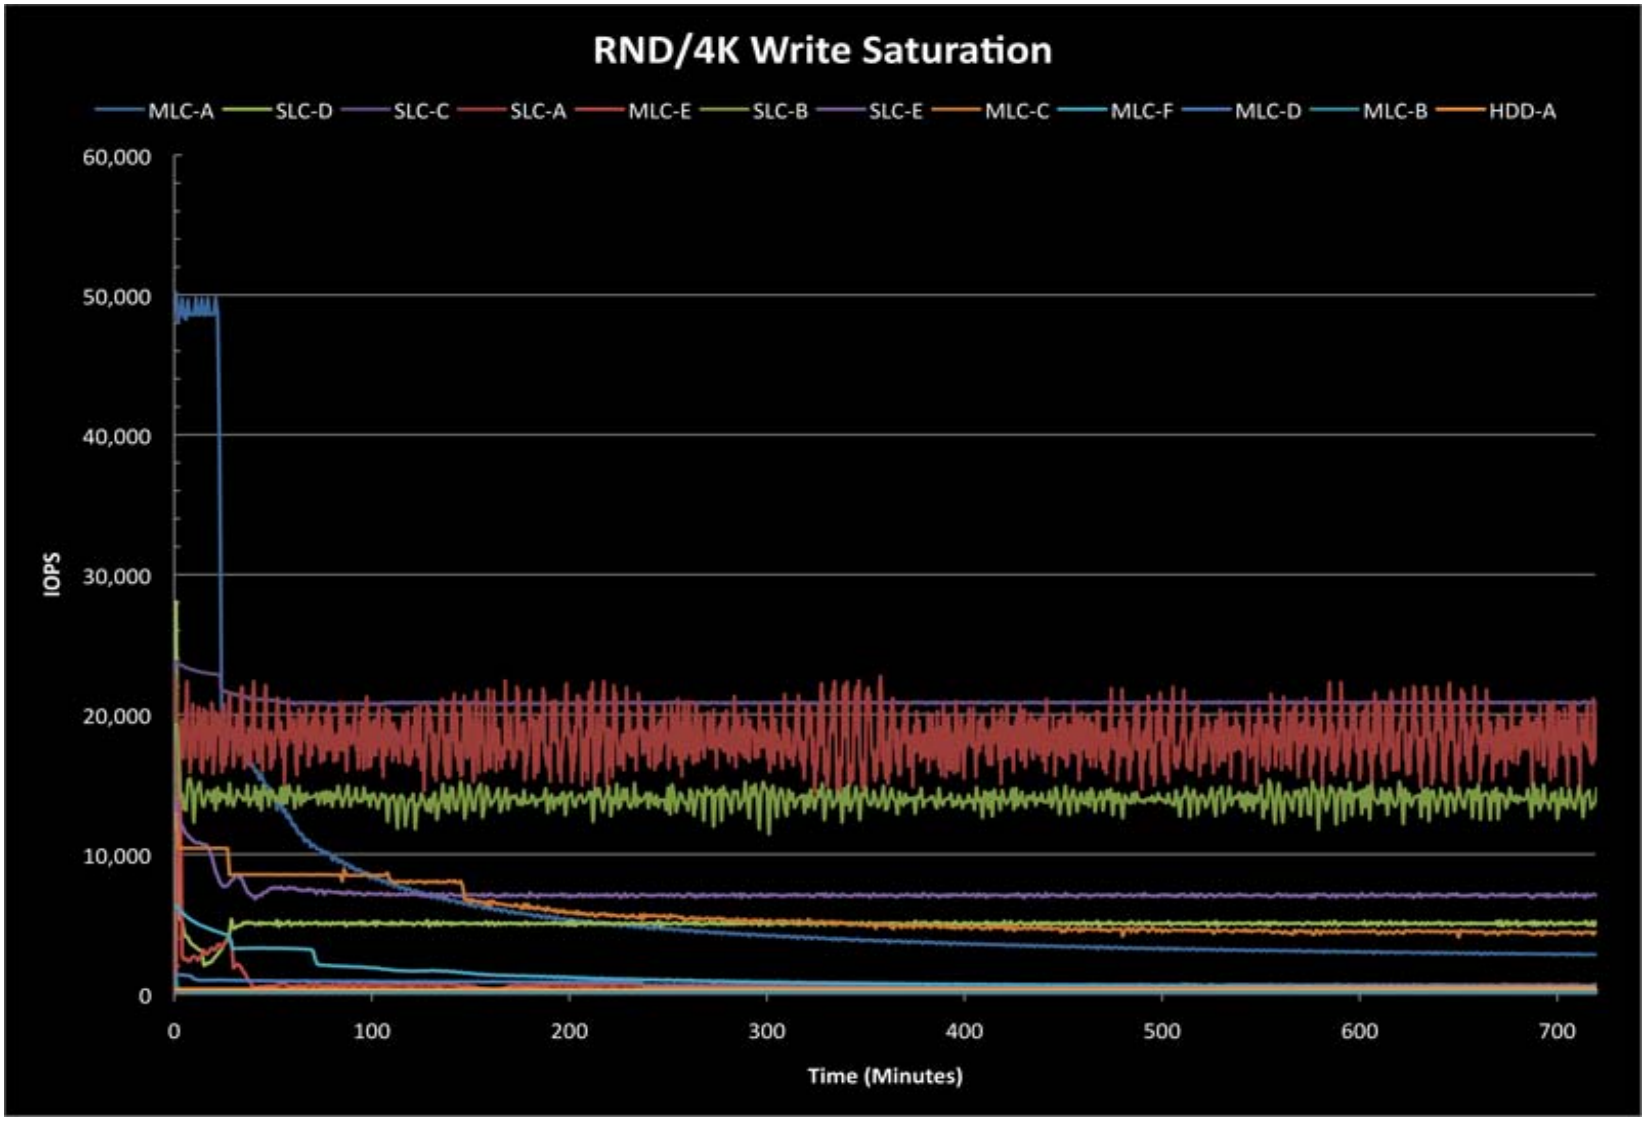

SSD's are Complicated Devices

- NAND-based SSD's are quite different from HDDs
  - No moving parts
  - Generally no direct overwrite
  - Not true "random access devices" → page read/writes, block erase
  - Limited cell cycle life → wear leveling required
  - Various NAND flavors: SLC, MLC, XLC, SLC+MLC….
  - Internal data flow management key to performance differentiator

Items Impacting SSD Performance

Education SNIA

Write History	PC Active Range	Test Active Range	Data Pattern
• What was previously written	• Where data was previously written • Trim effects	• Where and how much data is written	• What is the content

Access Pattern	Demand Intensity	Throttling	?
• Manner in which data is being accessed	• How hard apps are driving the device	• How fast data is allowed to be written	

A Solution is in Emerging

- SSS PTS
  - SNIA SSSI & TWG Solid State Storage Performance Test Specification (SSS PTS) Enterprise Draft v 1.0 –Public Review
- Standardized Tests & Methodologies
  - Effectively measure device performance of SSS products
- Fair Comparisons
  - Using a standardize test methodology and reporting requirements, performance can be more easily compared, particularly done using a reference environment

Key Concepts

- The SSS PTS is grounded in a few key concepts:
  - ***Common Starting Point*** **-** start test by first placing the drive into a known, repeatable state
  - ***Pre-Conditioning*** – from the common starting point, pre-conditioned the drive to a "used" state
  - ***Steady State*** – measurements are taken only when key performance metrics are relatively time invariant
  - ***Required Reporting*** – establishes required testing conditions and results reporting

The SSS PTS

- Agnostic to Test Platforms
  - no specific test environment
  - however, a Reference Test Platform (RTP) is outlined to facility direct comparisons
- Generic Test Tool Requirements
  - sets minimum requirements a test application and the test environment must be able to do
- Standardized Test Report Format –
  - Draft test report format is proposed with required conditions, variables & data formats

SSD Reference Test Platform

- The SNIA SSS TWG approved a "Reference Test Platform", which specified a set of hardware and options for software to allow direct comparisons
- Calypso has developed a RTP product based on this recommendation
- A significant portion of the data used as input to inform the formation of the Spec is taken on the Calypso RTP
- The Calypso RTP has been used extensively to validate the current Draft V1.0 Specification
- Calypso's RTP is fully PTS-compliant, and has been used by Calypso for 2010 Blind Survey of SAS/SATA SSDs & other 3d Party Comparison Reports

SSSI Reference Test Platform

Education SNIA

- Intel S5520HC
- Single Intel W5580, 3.2GHz, Quad-core CPU
- 12GB, 1333MHz, ECC DDR3 RAM
- LSI 9212-4e4i 6Gb/s SAS HBA
- Intel ICH10R 3Gb/s SATA
- 8X Gen-II PCI-e

- CentOS 5.5
- Calypso RTP Backend V1.5
- Calypso Test Suite (CTS) V6.5

R/W=65/35, Various Block Sizes

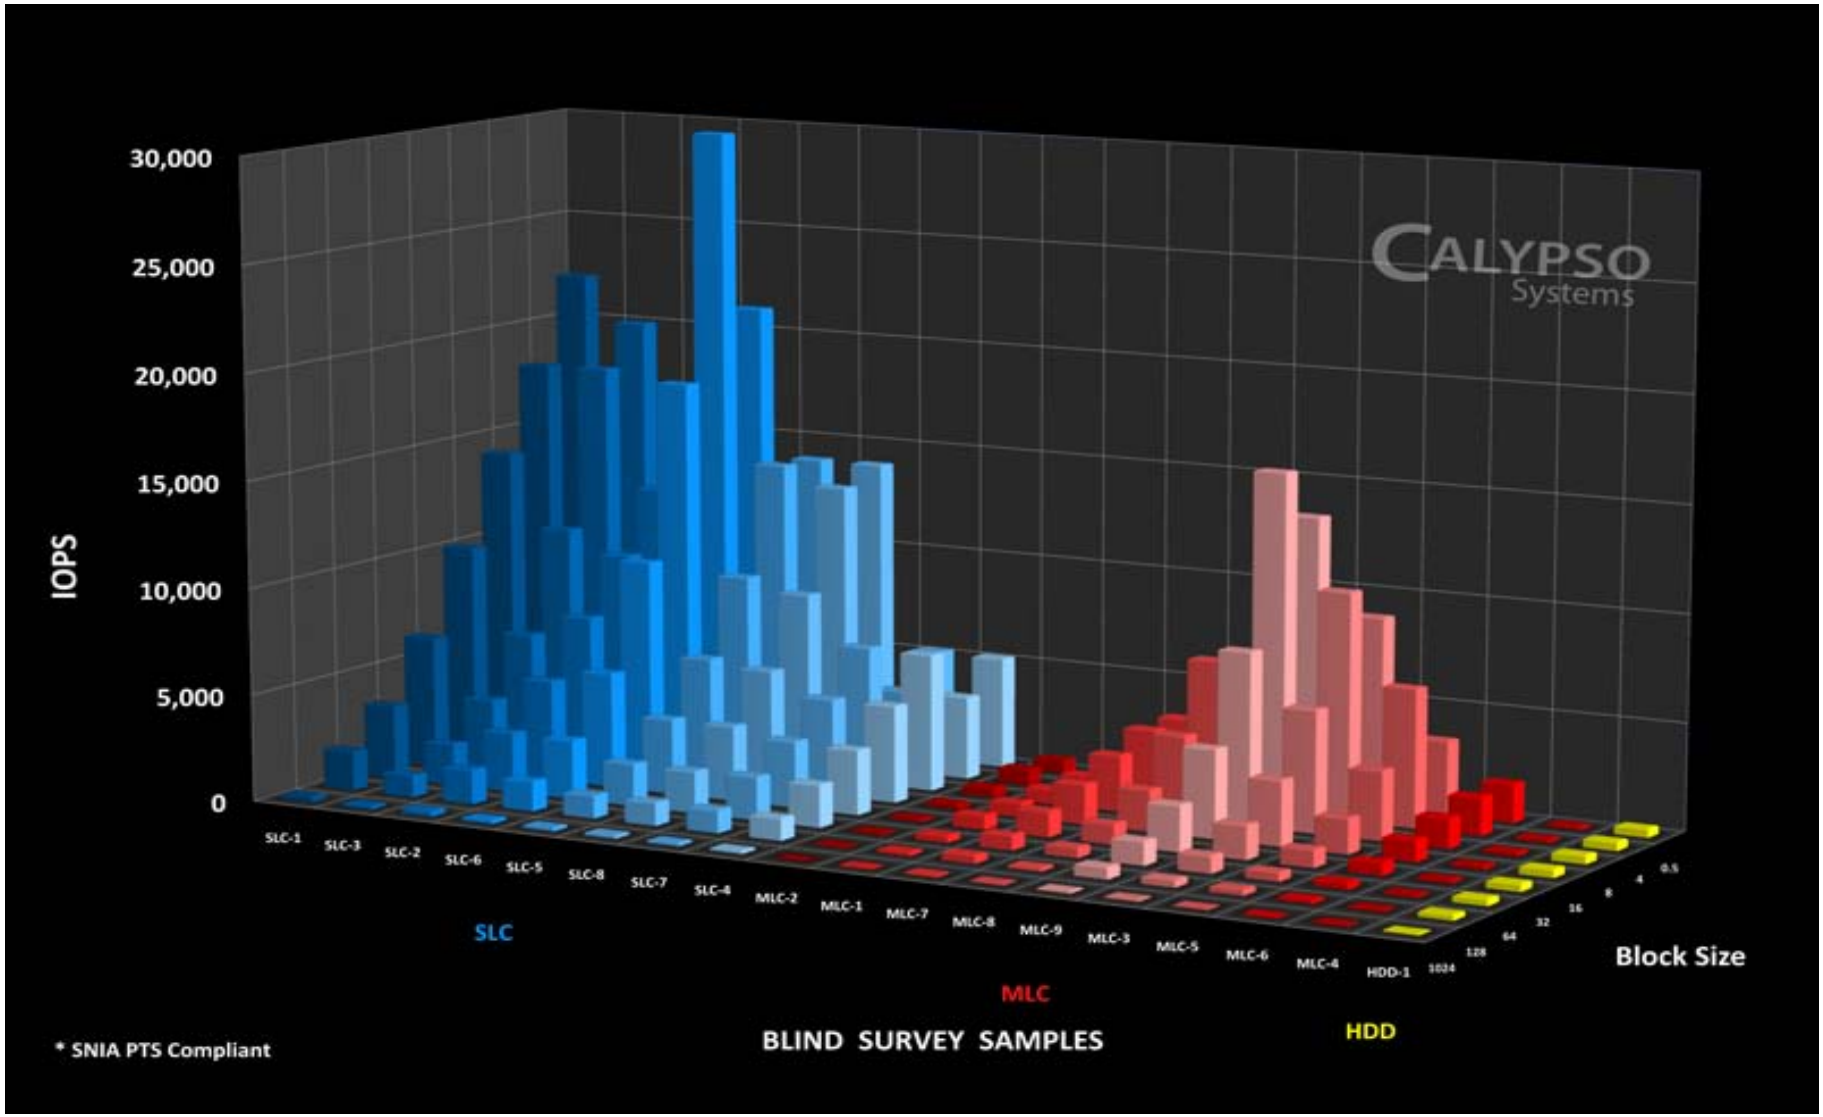

Table of Contents

- Overview
  - Motivation
  - Key Issues Considered
  - Test Environment
- PTS Specification
  - Enterprise PTS v1.0 – Purpose, Scope, Exclusions
  - Test Setup, Purge, Steady State
  - Tests Contained in the Enterprise PTS v1.0
  - Examples: Enterprise IOPS and Enterprise Latency Tests
- PTS Roadmap
  - Follow-on Work In Progress or Consideration
- SNIA Organization and the SSSI
  - Feedback & Involvement

Purpose

"…This Specification defines a set of
device level tests and ***methodologies***
to enable *comparative testing*
of Solid State Storage (SSS) devices
in Enterprise systems."

- Performance Test Specification v1.0 – Section 1.1

Tests Contained In Draft V1.0 Spec.

- The V1.0 Specification encompasses:
  - A suite of basic SSS performance tests
    - Write Saturation
    - IOPS
    - Throughput
    - Latency
  - Preconditioning and Steady State requirements
  - Standard test procedures
  - Standard test reporting requirements

What Is NOT Covered In the Spec

- Application workload tests
- Matching to user workloads
- Energy efficiency
- Required test platform (HW/OS/Tools)
- Certification
- Device endurance, availability, data integrity

- Performance Test Specification v1.0 – Section 1.4

The SNIA PTS Enterprise Draft V1.0

Education SNIA

Write Saturation	Enterprise IOPS	Enterprise TP	Enterprise Latency

• Random Access
• R/W: 100% Writes
• BS: 4K
	• **Random Access**
	• **R/W:**
	  • 100/0, 95/5, 65/35, 50/50, 35/65, 5/95, 0/100
	• **BS:**
	  • 1024K, 128K, 64K, 32K, 16K, 8K, 4K, 0.5K
		• **Sequential Access**
		• **R/W:**
		  • 100/0, 0/100
		• **BS:**
		  • 1024K, 64K, 8K, 4K, 0.5K
			• **Random Access**
			• **R/W:**
			  • 100/0, 65/35, 0/100
			• **BS:**
			  • 8K, 4K, 0.5K

Basic Test Flow

1. Purge	• Security Erase, Sanitize, Format Unit, other proprietary methods where indicated
2. Set Conditions	• Set user selectable test parameters, such as Active Range, Data Pattern, Demand Intensity
3. Pre-Condition	• Workload independent (WIPC) • Workload dependent (WDPC)
4. Run Until SS	• Reiterate loops until Steady State is reached, or run to a prescribed maximum number of loops
5. Collect Data	• Collect data from Steady State Measurement Window
6. Generate Reports	• Use standard report formats and include required and optional elements

Basic Test Flow

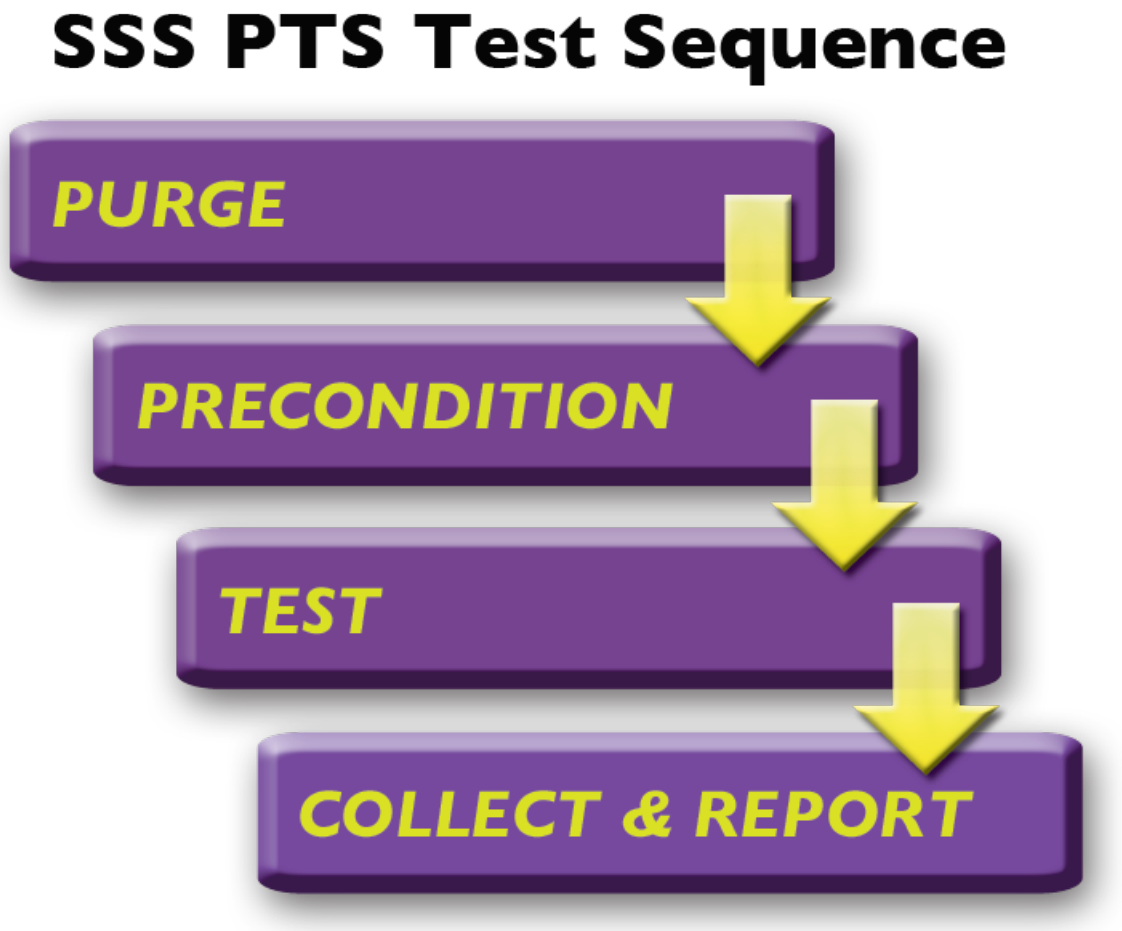

Key Concepts Used in the Spec.

- A. Purge
- B. Pre-Condition
  - Workload independent
  - Workload dependent
- C. Active Range
  - Pre-conditioning
  - Test
- D. Steady State
  - Measurement window
  - Data excursion condition
  - Slope excursion condition

A: Purge

- As per the PTS V1.0 Specification, purge is defined as:

  " The process of returning an SSS device to a state in which subsequent writes execute, as closely as possible, as if the device had never been used and does not contain any valid data"

- Example implementation includes: ATA Security Erase, Sanitize, SCSI Format Unit

B: Pre-Conditioning

- Pre-Conditioning is a key requirement in getting repeatable, representative results
- Goal is to put drive into "Steady State", using:
  - Workload independent – *PTS v1.0 Section 3.3*
    - Use a prescribed workload unrelated to the test loop
    - Write 2X user capacity using SEQ/128KiB blocks
  - Workload dependent – *PTS v1.0 Section 3.3*
    - Run test workload itself as pre-conditioning (self pre-conditioning)

C: Active Range

- As per the PTS V1.0 Specification, Active Range is defined as:

"… ActiveRange is the range of LBA's that may be accessed by the preconditioning and/or test code..."

- They are normally defined as % of the maximum LBA available to the user
- Note Pre-conditioning and Test can have different Active Ranges

D: Steady State Definition

- Premise is that reported data should be take only AFTER the test loop results shows the drive has reached and maintained "Steady State"
- The Measurement Window is the interval, measured in Rounds, when the test results have entered and maintained Steady State for 5 Rounds

D: Steady State Definition

- Steady State is reached only if BOTH of the following conditions are satisfied (assuming "*y*" is the variable being tracked):
  1. Variation of *y* within the Measurement Windows is within 20% of the Average
     - *" Max(y)-Min(y) within the Measurement Window is no more than 20% of the Ave(y) within the Measurement Window; and "*
  2. Trending of *y* within the Measurement Windows is within 10% of the Average
     - *" [Max(y) as defined by the linear curve fit of the data within the Measurement Window] – [Min(y) as defined by the best linear curve fit of the data within the Measurement Window] is within 10% of Ave(y) within the Measurement Window. "*

D: SS Measurement Window

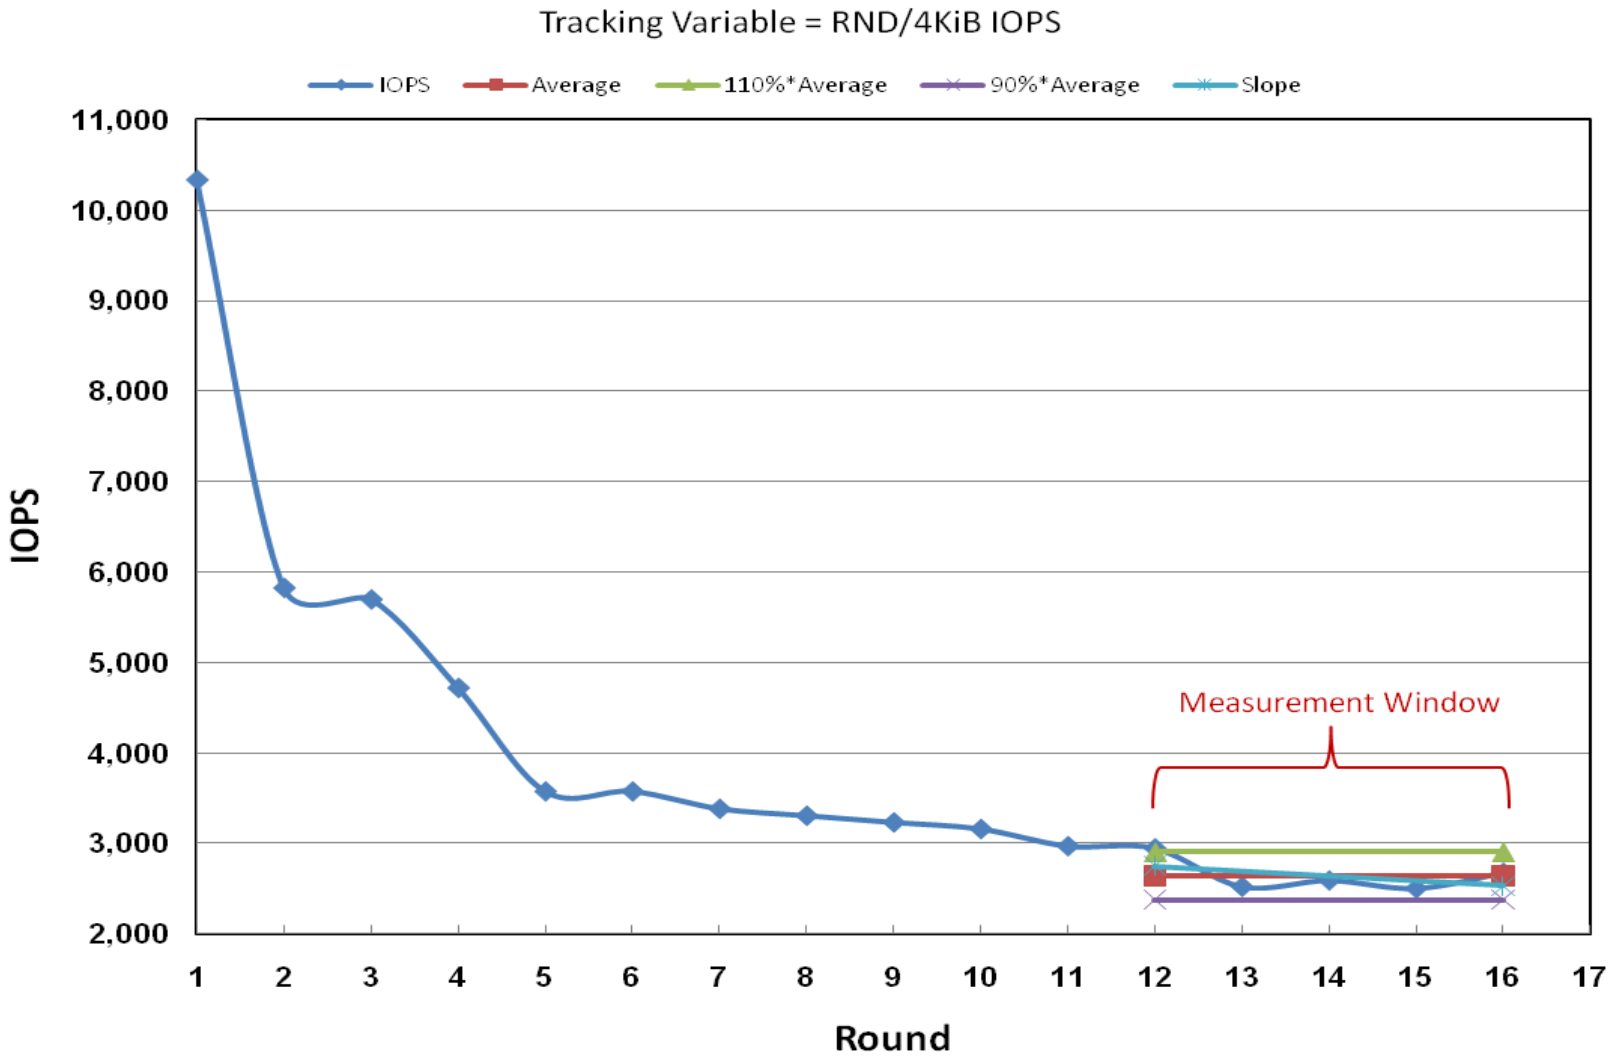

D: SS Measurement Window

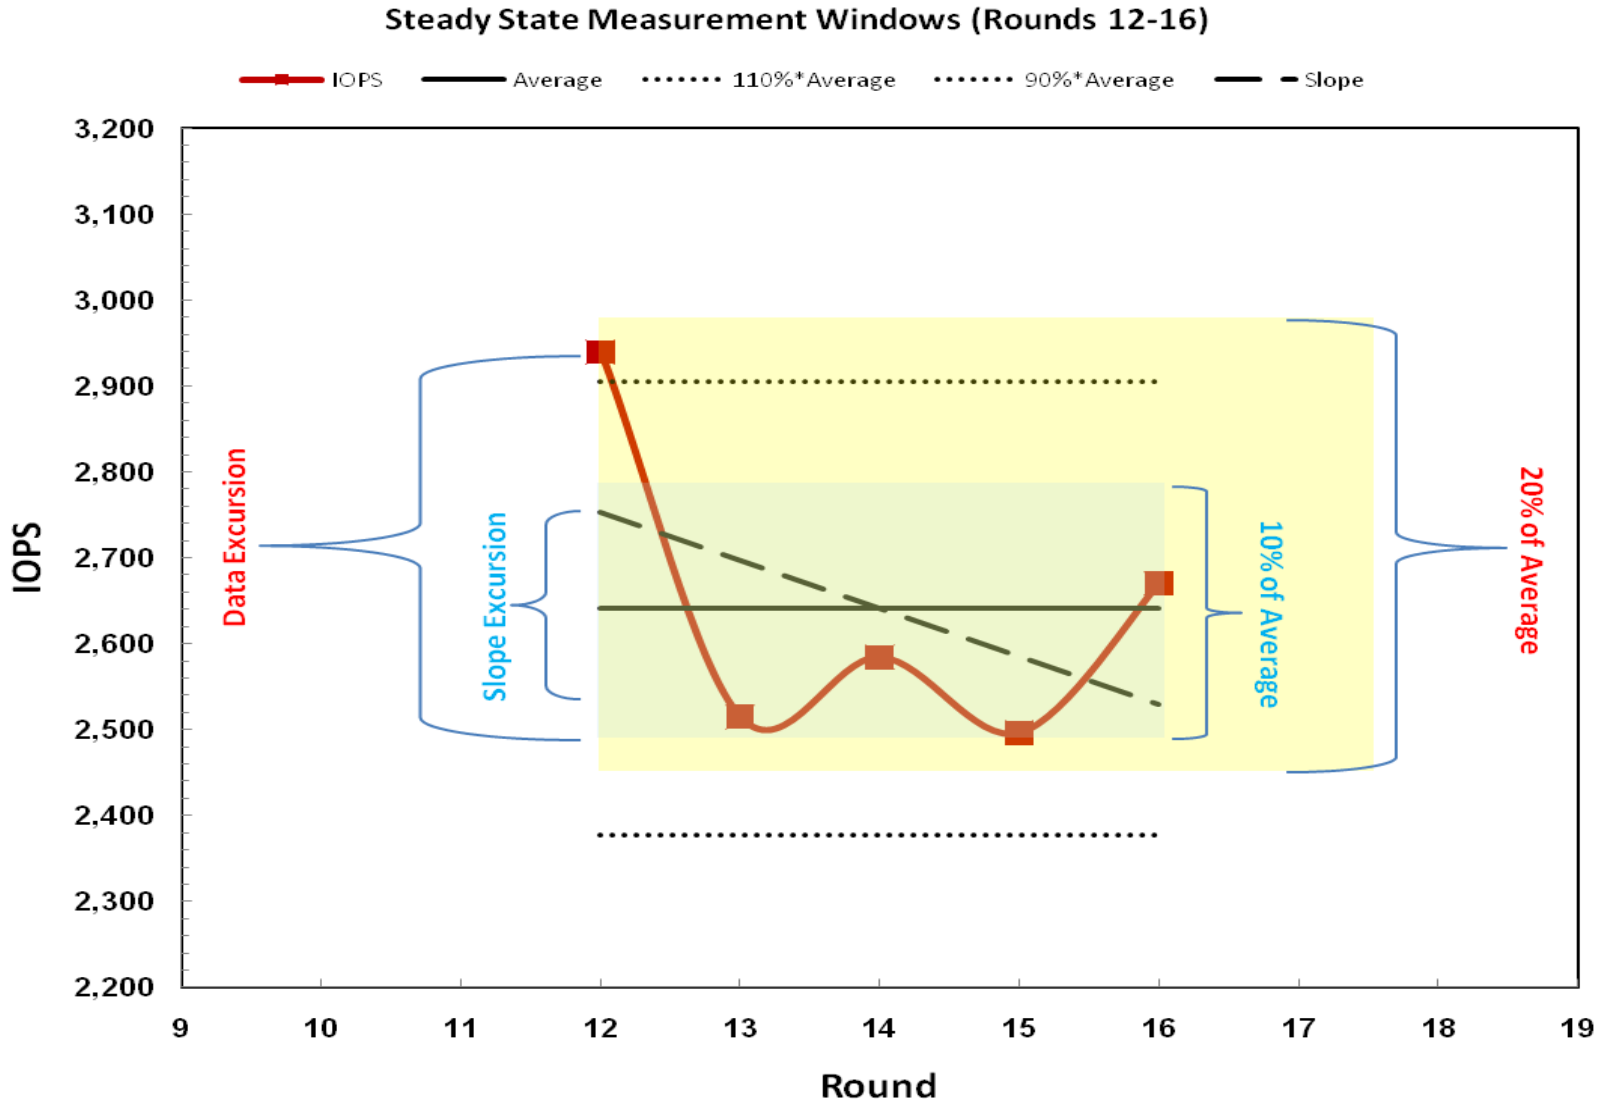

D: Steady State

- Compare
  - [Data Excursion] with [20% of Average]
  - [Slope Excursion] with [10% of Average]
- Note
  - This method is slightly more tolerant than +10% and – 10% data excursion method and +5% and – 5% slope excursion method

D: How Good is the Steady State

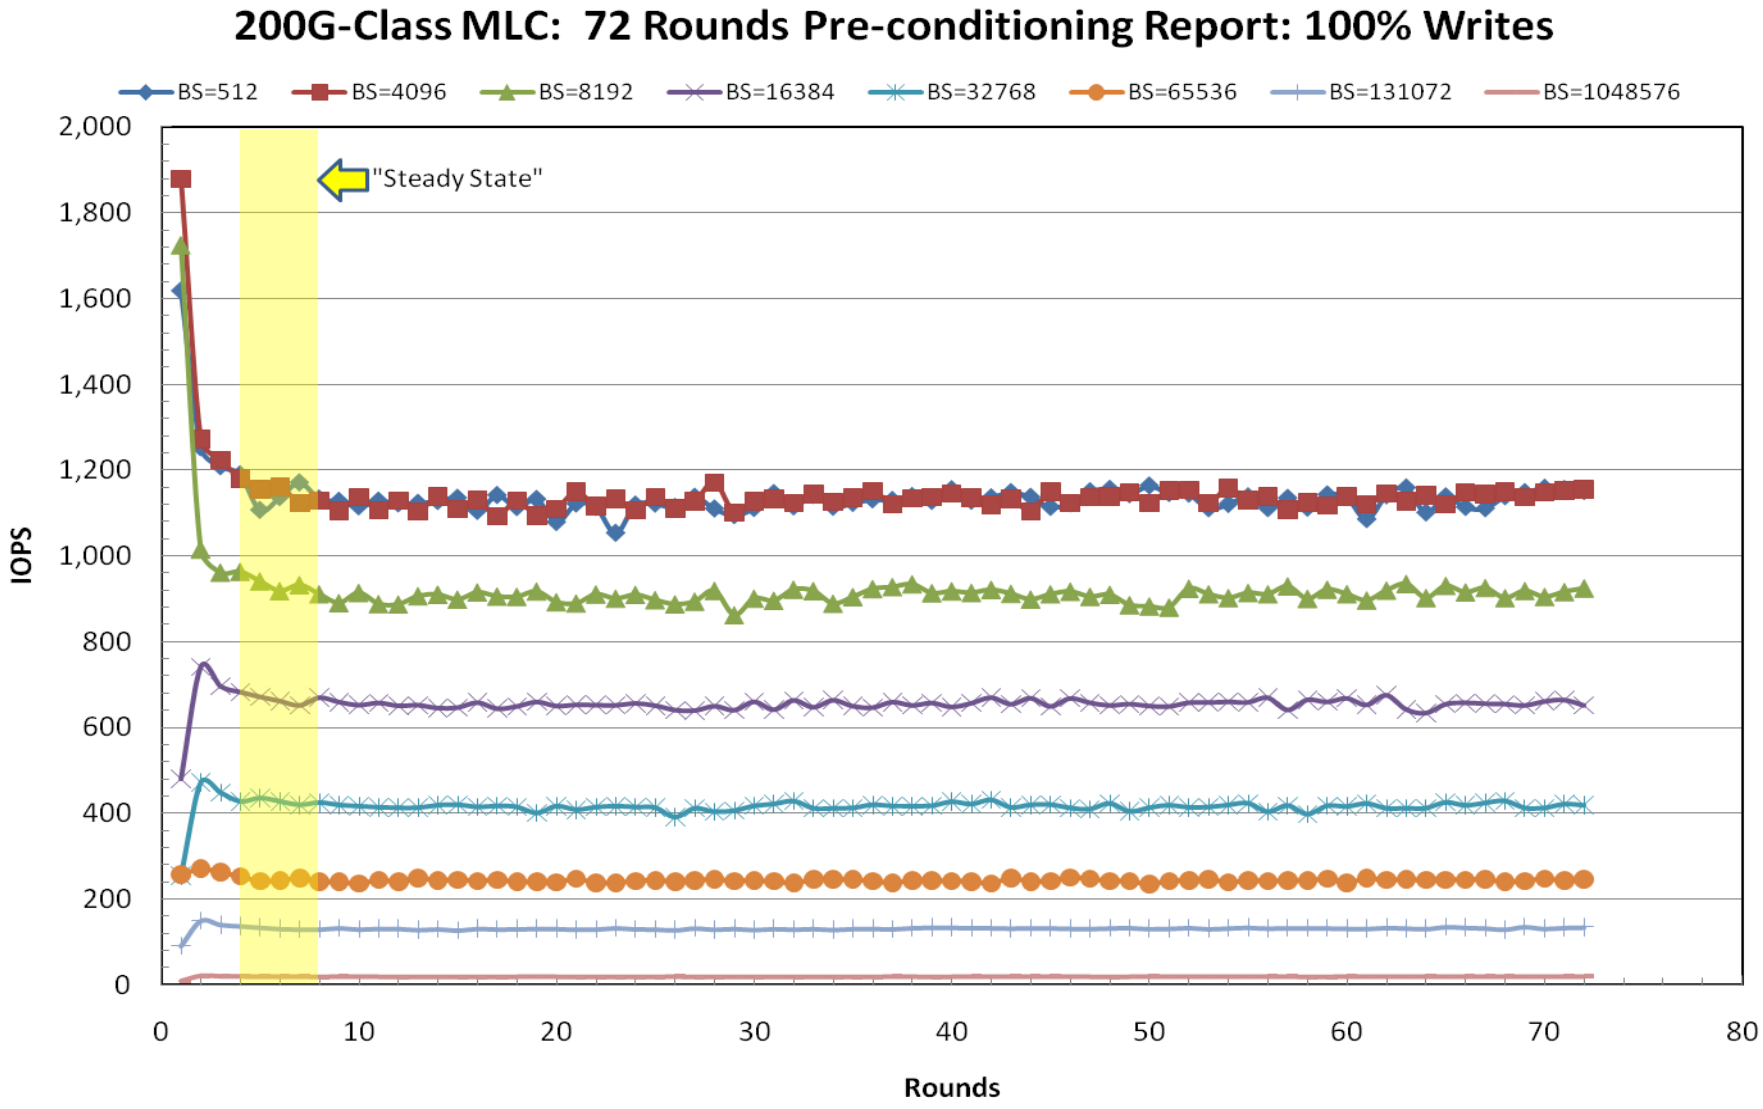

Workload Schematics

Write Saturation	Enterprise IOPS	Enterprise TP	Enterprise Latency

• Random Access
• R/W: 100% Writes
• BS: 4K
	• **Random Access**
	• **R/W:**
	  • 100/0, 95/5, 65/35, 50/50, 35/65, 5/95, 0/100
	• **BS:**
	  • 1024K, 128K, 64K, 32K, 16K, 8K, 4K, 0.5K
		• **Sequential Access**
		• **R/W:**
		  • 100/0, 0/100
		• **BS:**
		  • 1024K, 64K, 8K, 4K, 0.5K
			• **Random Access**
			• **R/W:**
			  • 100/0, 65/35, 0/100
			• **BS:**
			  • 8K, 4K, 0.5K

IOPS RW/BS Sequence

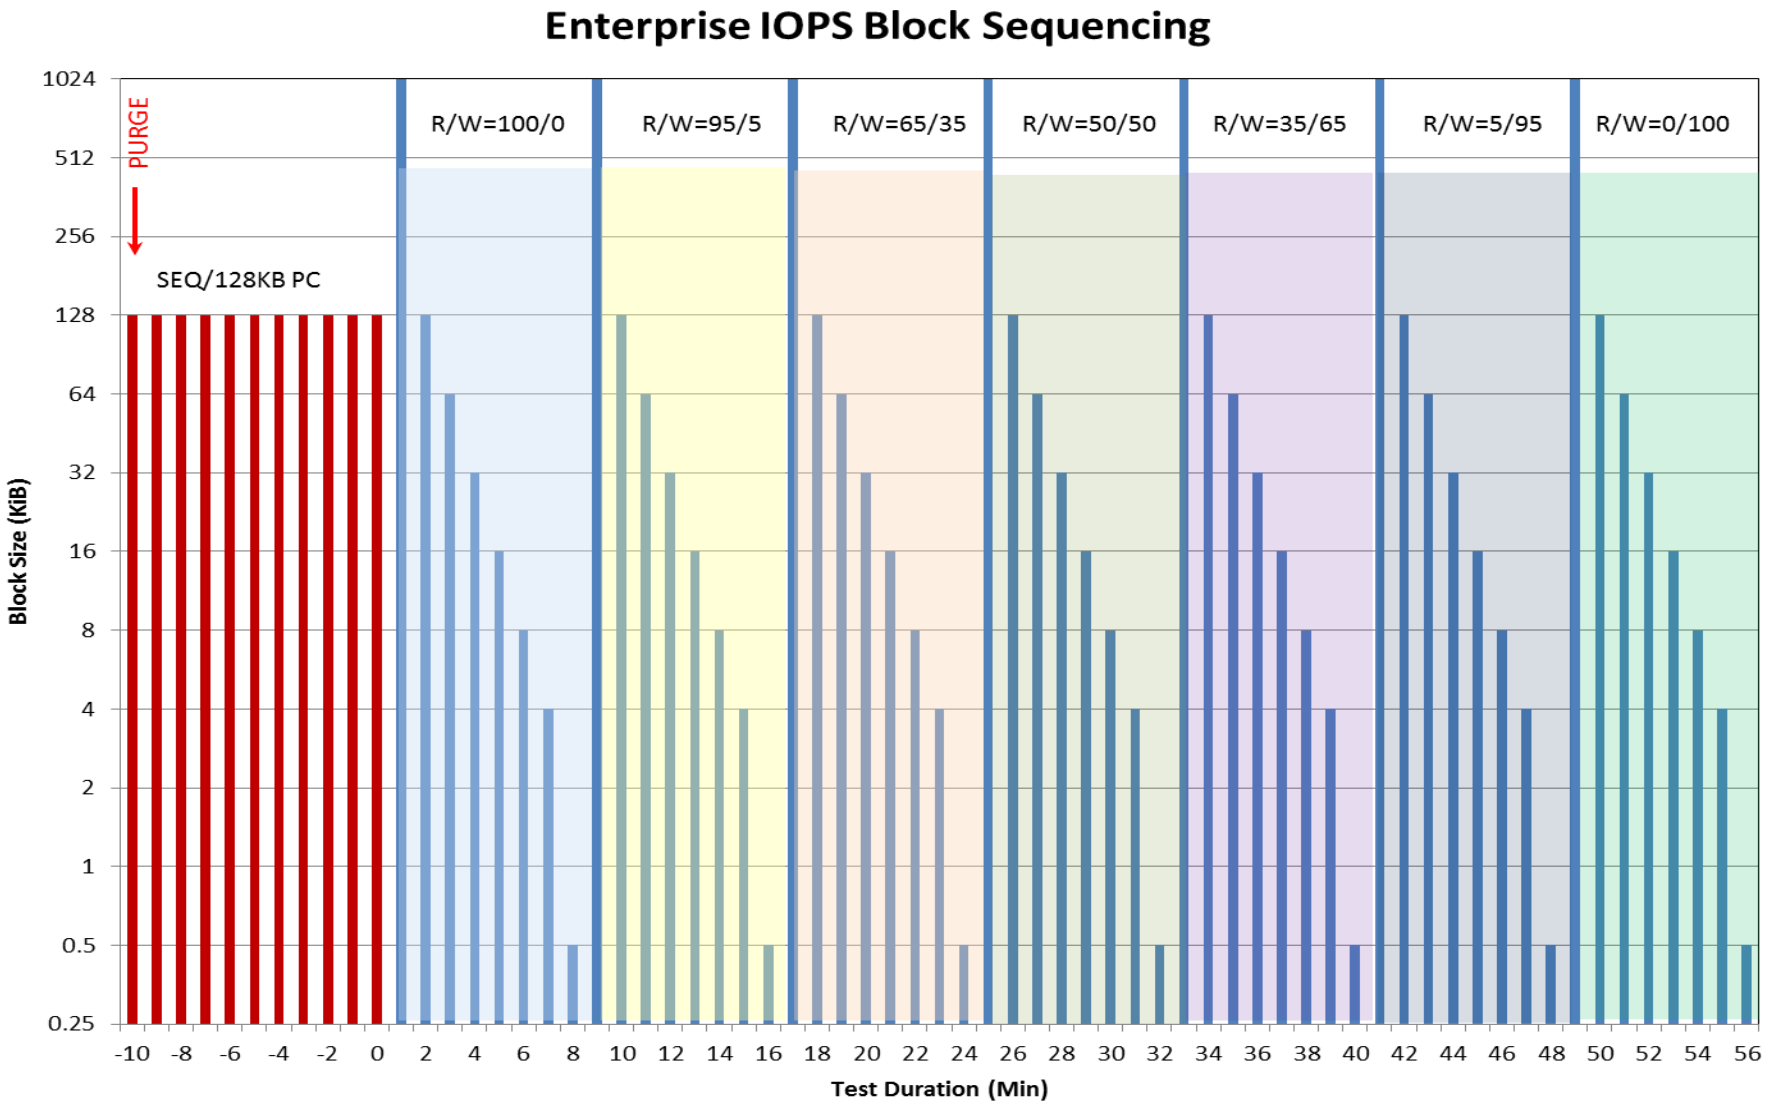

Enterprise IOPS RW/BS Sequence

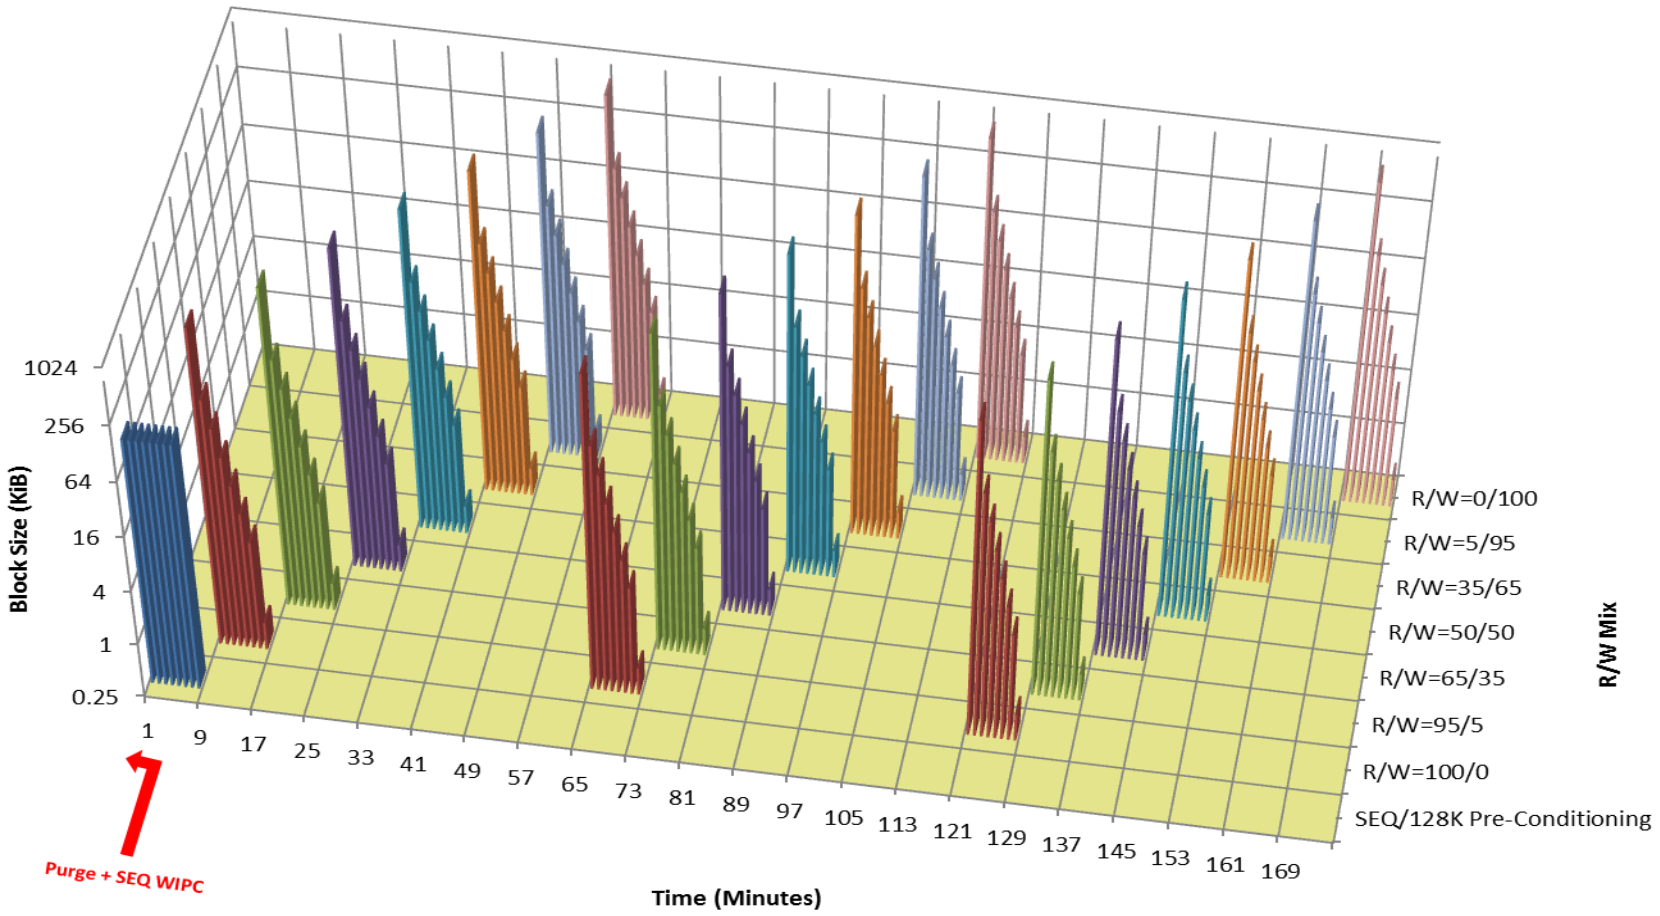

TP RW/BS Sequence

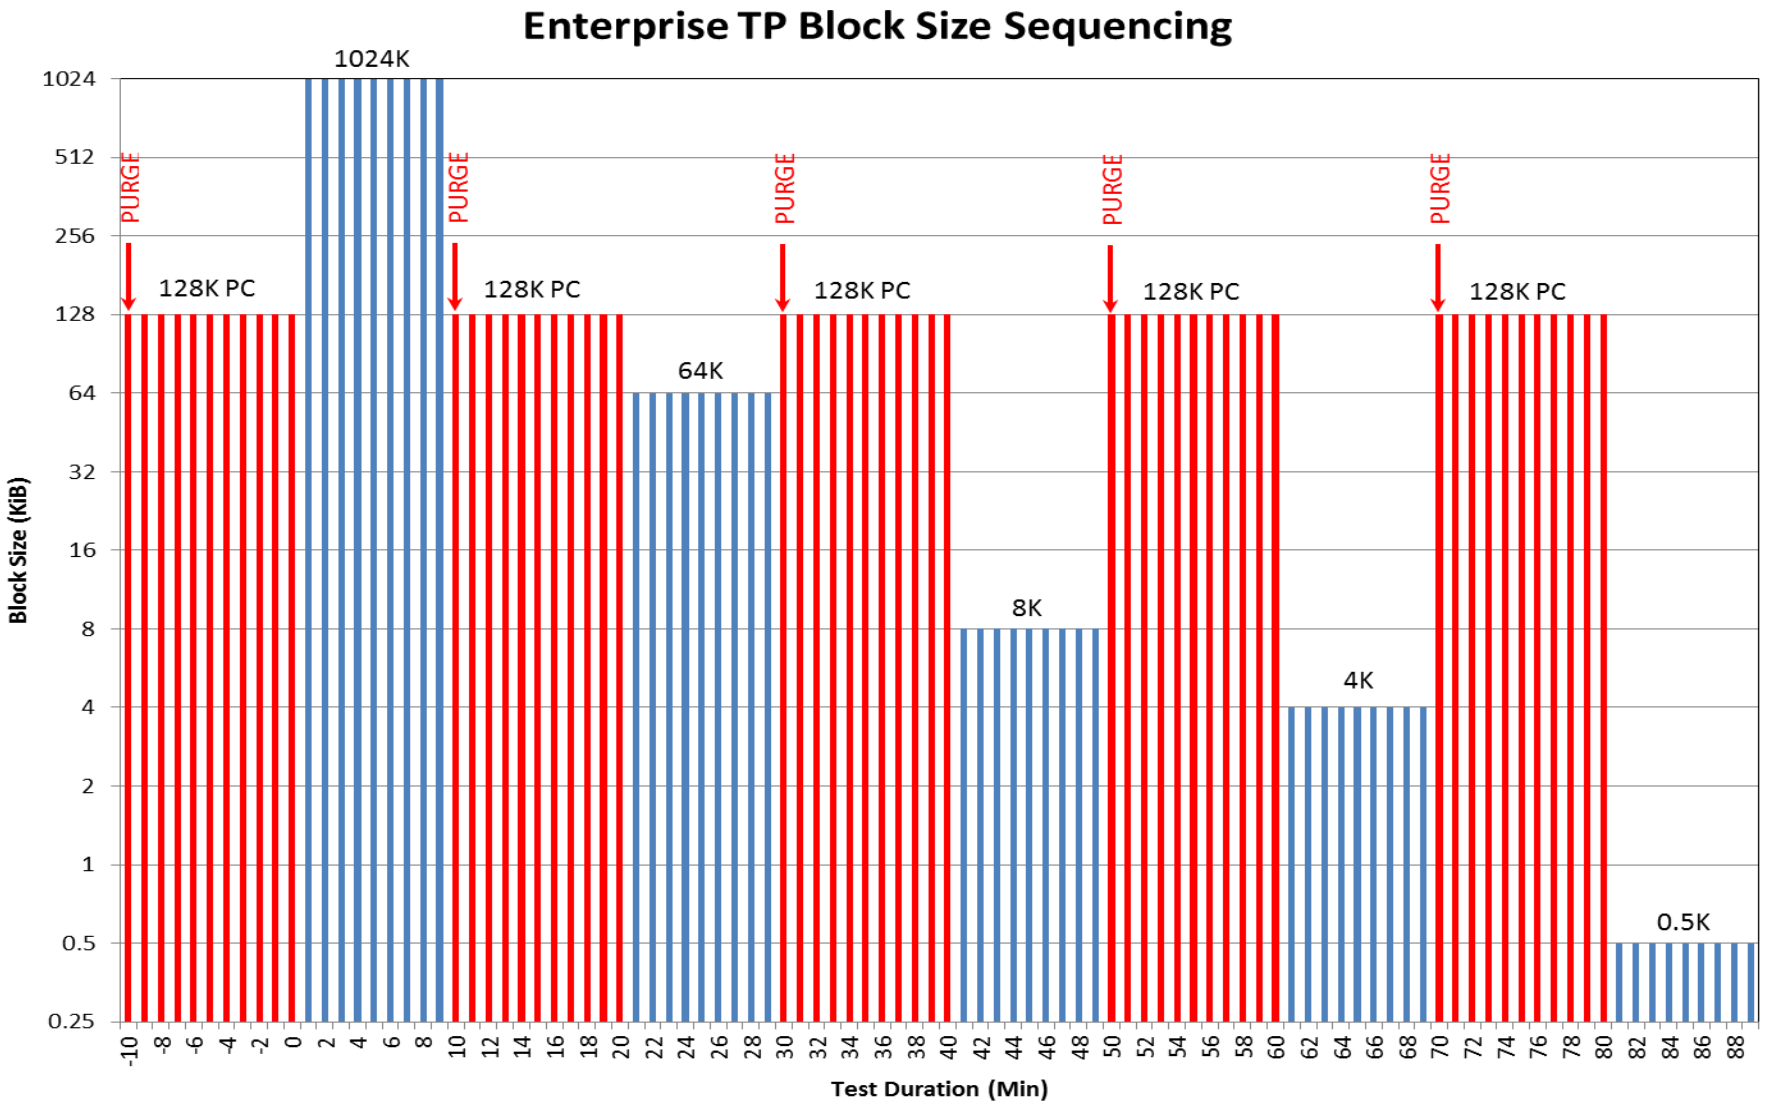

Enterprise Latency RW/BS Sequence

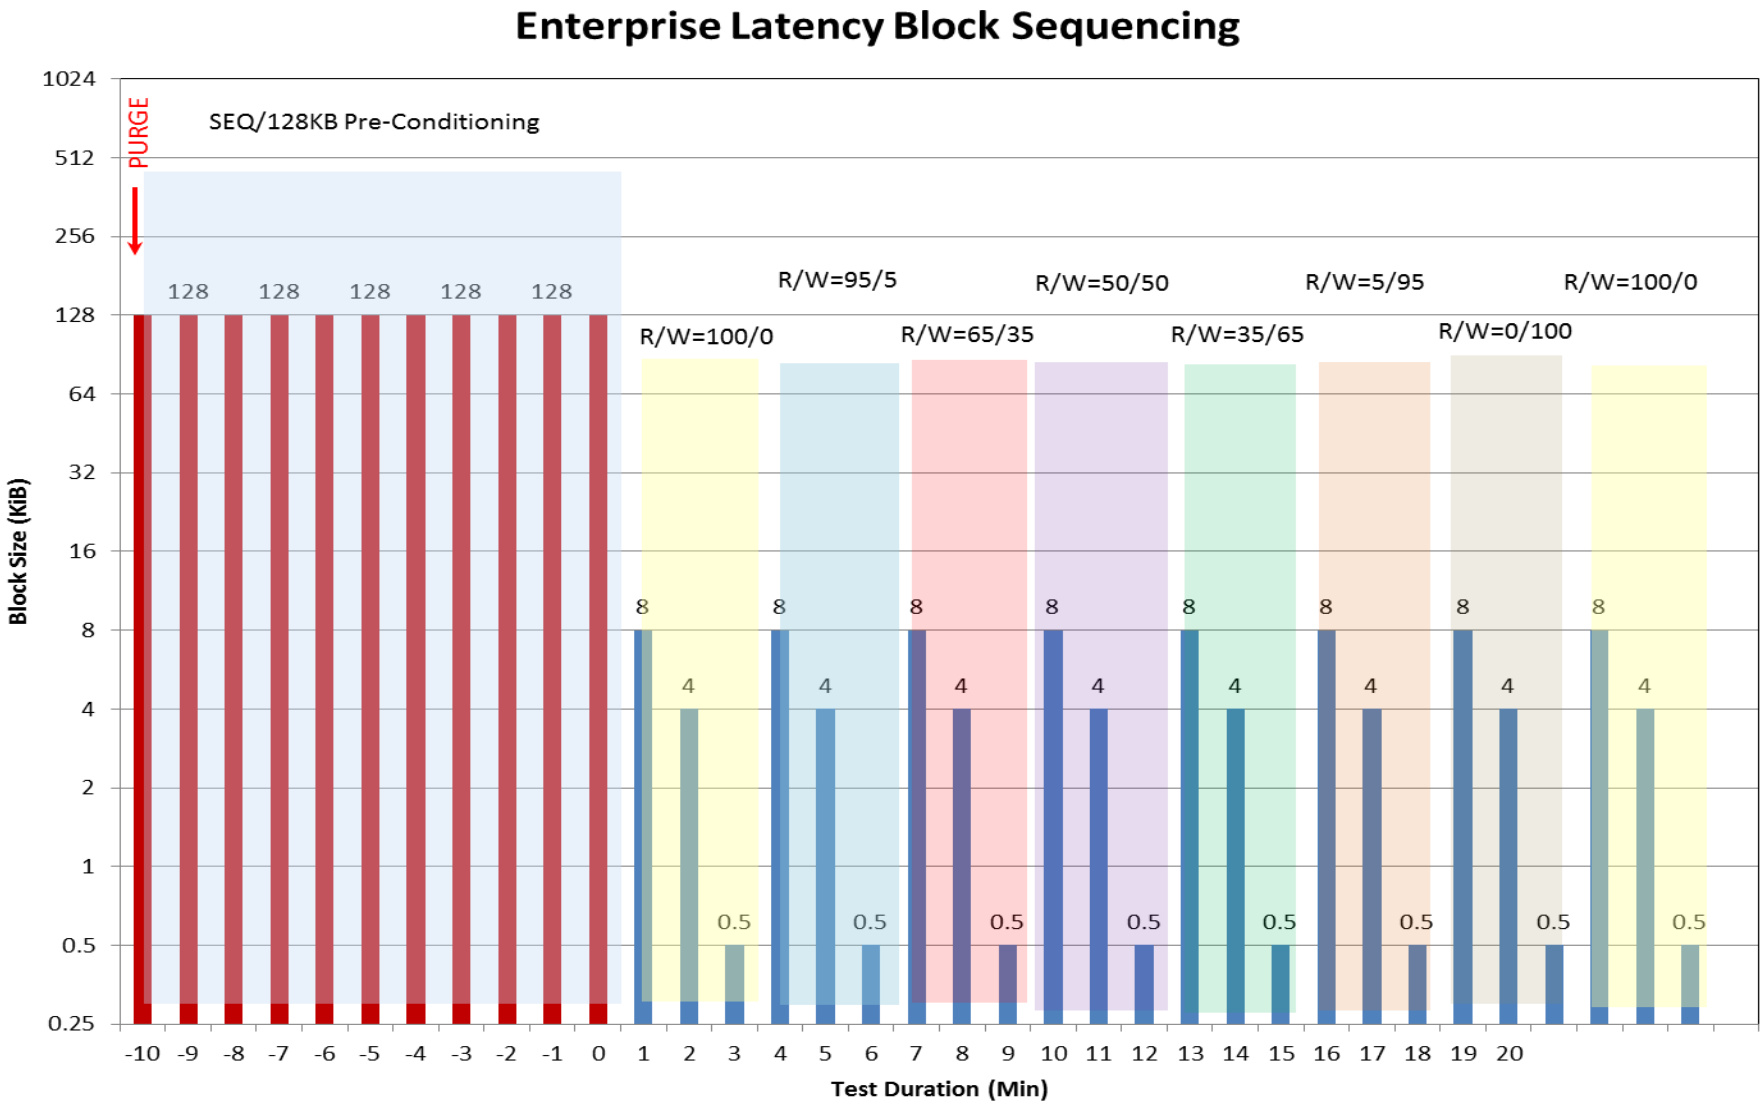

Two Examples

- Enterprise IOPS
- Enterprise Latency

Enterprise IOPS

- DUT:
  - 100GB-Class Enterprise SLC drive
- Test Parameters:
  - Active Range = [0,100%]
  - Thread Count=2
  - Queue Depth (Outstanding IO/Thread)=16
  - DP=RND

Enterprise IOPS Draft Formatted Report, 1/6

Enterprise IOPS Draft Formatted Report, 2/6

Enterprise IOPS Draft Formatted Report, 3/6

Enterprise IOPS Draft Formatted Report, 4/6

Enterprise IOPS Draft Formatted Report, 5/6

Enterprise IOPS Draft Formatted Report, 6/6

Enterprise Latency

- DUT:
  - 100GB-Class SLC drive
- Test Parameters:
  - Active Range = [0,100%]
  - Thread Count=1
  - Queue Depth (Outstanding IO/Thread)=1
  - DP=RND

Enterprise Latency Draft Formatted Report 1/6

Enterprise Latency Draft Formatted Report 2/6

Enterprise Latency Draft Formatted Report 3/6

Enterprise Latency Draft Formatted Report 4/6

Enterprise Latency Draft Formatted Report 5/6

Enterprise Latency Draft Formatted Report 6/6

Table of Contents

SNIA SSSI PTS Follow-On Work

Enterprise SPEC

- power measurement
- cross stimulus recovery
- response time histograms
- …

Client SPEC

- Constraints to better match client usage cases → active footprint, active span, segmentation, etc.
- Trace-based workloads
- …

SNIA SSSI PTS Follow-On Work

Education SNIA

IOPS/W	• Measure total W used over a period of time, and deriving the number of IOPS that can be achieved at a given block size and access pattern with unit power • Measure the power efficiency of the device
Client Active Range Restriction	• Perform PC or Test in a restricted LBA range or ranges • To better simulate the environment of client usage
Cross Stimulus Recovery	• Measure performance metrics when changing between RND/SEQ and small block/large block stimulus • To see how drive handles switching between sustained access patterns
Demand Intensity	• Measure performance metrics with various outstanding IOS from the test application • understand the trade-off between achieving maximum IOPS vs acceptable maximum response time criteria

SNIA SSSI PTS Follow-On Work

Response Time Histogram	• Get detailed response time statistics during specific stimulus • to provide better insight into a drive's response time performance beyond a single average response time number
Task-Based Synth. Workloads	•Synthetic approximation to IO Trace playback based on understanding of access characteristics of specific tasks, such as video streaming, office productivity, etc. •Build a library of well-studied synthetic stimulus that can be used to form more complex user cases
SSD Figure of Merit	• Derive simplified metric(s) from data resulting from various PTS tests • To allow simply comparison between drives to aid marketing
Industry Requests?	• ?
?	• ?

Table of Contents

- Overview
  - Motivation
  - Key Issues Considered
  - Test Environment
- PTS Specification
  - PTS v1.0 – Purpose, Scope, Exclusions
  - Test Setup, Purge, Steady State
  - Tests Contained in the PTS v1.0
  - Test Drive: Sample Run using Various SSD
- PTS Roadmap
  - Follow-on Work In Progress or Consideration
- SNIA Organization and the SSSI
  - Feedback & Involvement

SNIA – Organization Chart

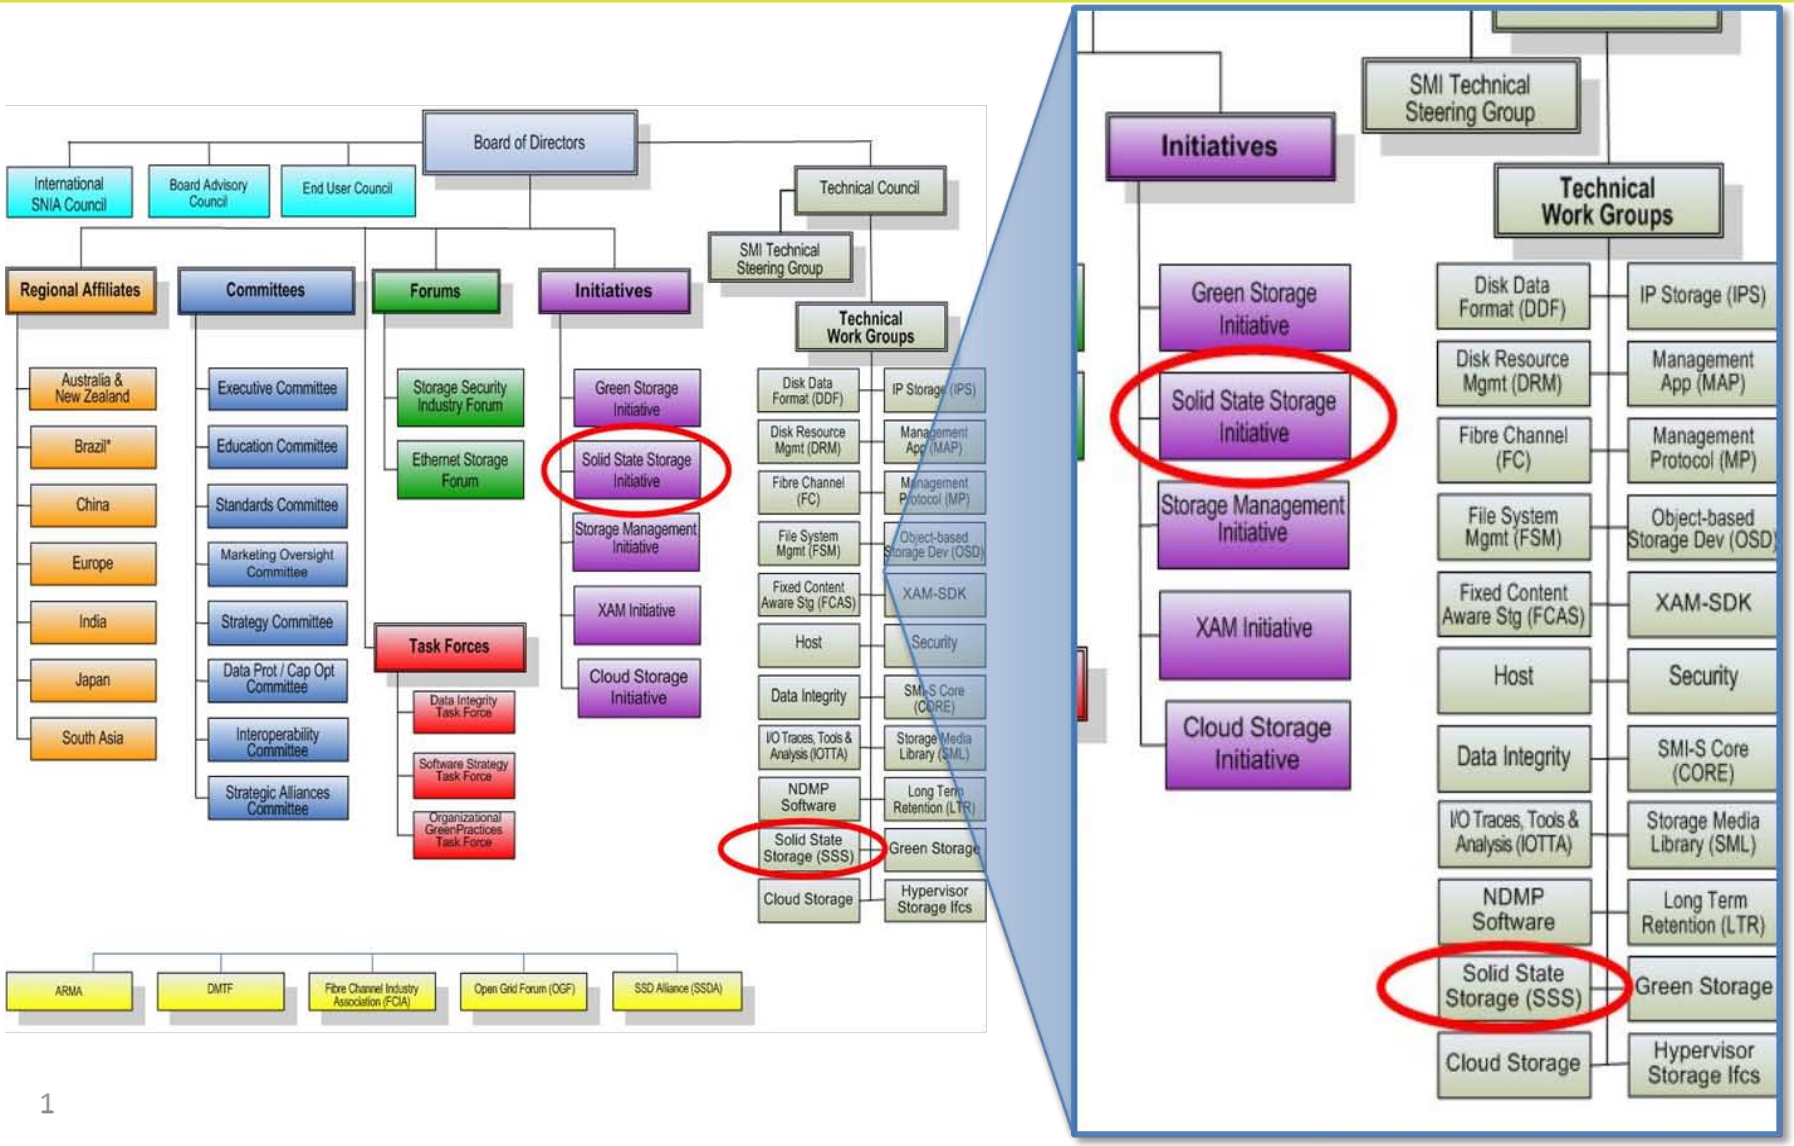

1

Charter & Mission Statements

SNIA – *Storage Networking Industry Association*

"Lead Industry Standards for information storage management"

SSSI – *Solid State Storage Initiative*

"Foster the success of Enterprise & Client SSS markets"

SSS TWG – *Solid State Storage Technical Working Group*

"Develop SNIA Technical Specifications & Standards"

Feedback & Involvement

Download The Spec.	• www.snia.org/publicreview
Upload Feedbacks	• www.snia.org/feedback
Updates to The Spec.	• www.snia.org/forums/sssi
SSS White Papers	• www.snia.org/forums/sssi/knowledge/education
Get Involved: SSSI	• www.snia.org/forums/sssi
Get Involved: SSSI TWG	• www.snia.org/apps/org/workgroup/ssstwg/

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: tracksolidstate@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Eden Kim