Enterprise Applications
How to create a synthetic workload test

Eden Kim, CEO 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

- Test Engineers & Marketing Managers responsible for SSD test, deployment and qualification.
  - This session will appeal to Data Center Managers, Development Managers, and those that are seeking a fundamental understanding of the SNIA's Solid State Storage Performance Test Specification (PTS).
  - The session will investigate Demand Intensity / Outstanding IOs, Confidence Level Plots / Response Time Histograms, and Synthetic Application Workloads.
  - The audience learn how to evaluate IOPS and Response Time Saturation, how to conduct a Response Time Confidence Level Sensitivity Analysis, and how to build a test template for synthetic workloads.
  - A case study will be presented using a synthetic database OLTP workload.

Synthetic Enterprise Application Workloads

- What are they?
- Why are they used?
- How do I use them?
- What are some examples?
- Case Study
  - *Two 2014 released SATA Enterprise class SSDs*
  - *OLTP workload tested to saturation*
  - *Examination of IOPS and Response Times*

Learning Objectives

- What are:
  - Demand Intensity Levels & Outstanding IOs
  - Confidence Level Plots & Response Time Histograms
  - Synthetic Application workloads
- How to:
  - Evaluate IOPS and Response Time Saturation
  - Do a Sensitivity Analysis of Confidence Levels
  - Build a test template for synthetic workloads

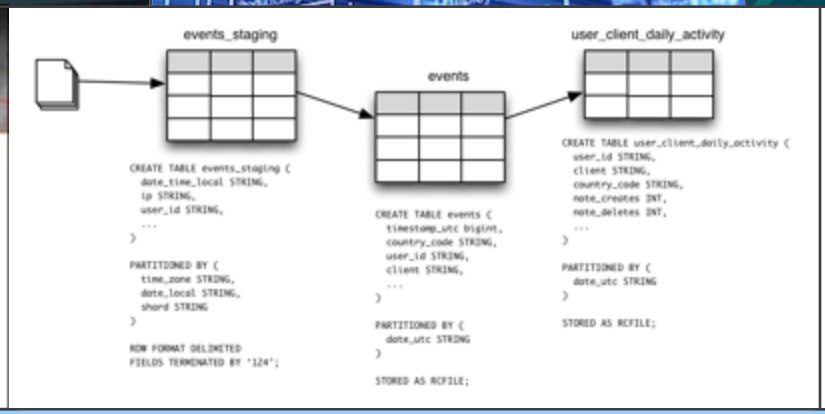

Synthetic Enterprise Application Workloads

WHAT ARE THEY?

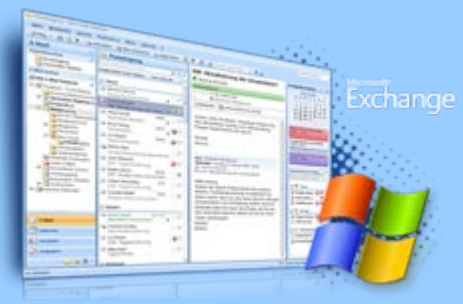

Enterprise Applications

Applications common in the Enterprise include:

- OLTP (On Line Transaction Processing)
- OLAP (On Line Analytical Processing)
- VOD (Video on demand)
- OS Paging
- Webserver / Exchange mail
- Logging (web server, SQL server logs)
- DSS (Decision Support Systems)
- Medical Imaging

TAKE AWAY
Different workloads apply different access patterns to the SSD.

Enterprise Application Workloads are Different

Test Profile – IOPS & RT levels of different synthetic workloads

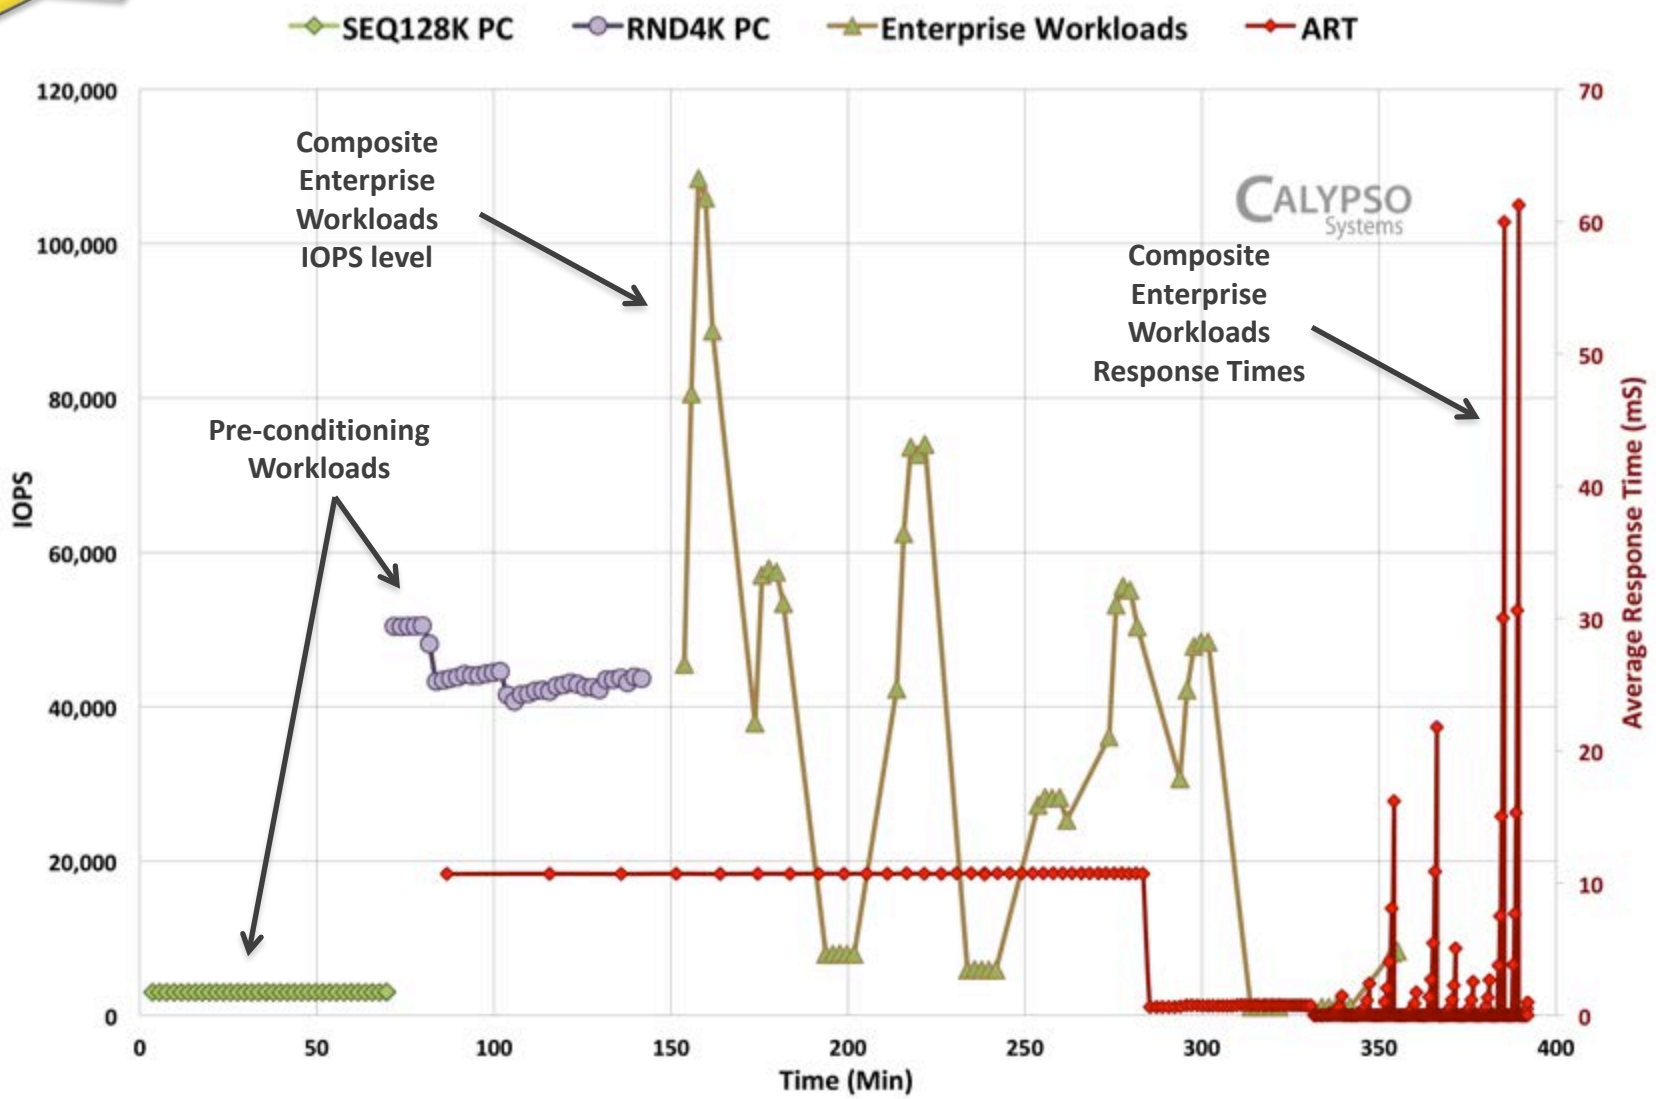

TAKE AWAY
HW/SW stack affects IOs – use standardized test platform to isolate SSD Device Performance.

Workloads are affected by the IO Stack

Application Workloads are generated in user space and traverse the IO stack to the SSD

Testing wants to measure workloads as close to the SSD as possible

(Block IO level)

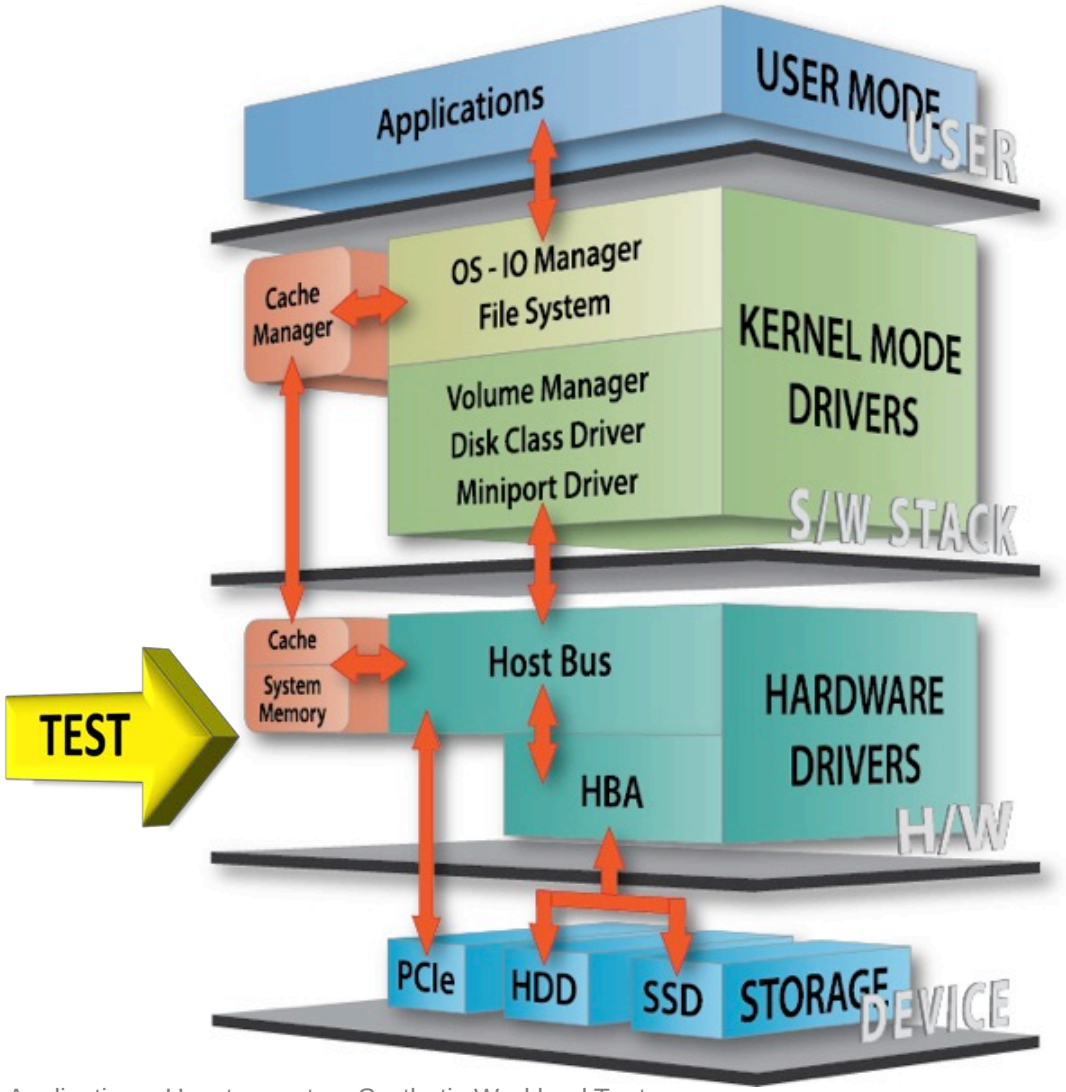

Workloads are described by Access Patterns

Degree of Randomness	Data Transfer Size	Read Write Ratio
Random or Sequential	Block Size	Read/Write Mix
Workloads are a series of Access Patterns over an observation period		
Examples of Access Patterns for different workloads:		
RND	4KiB	RW0
RND	8KiB	RW65
SEQ	128KiB	RW90
SR75	64KiB	RW95
SR25	128KiB	RW05

Workloads can be:

Monotonic or Composite stream of Accesses
Synthetic or Real World workloads

TAKE AWAY
IO Traces captures are specific to the system HW and drive tested.

Real-world Workloads

IO Trace Captures are specific to the system

Contain many many streams of different access patterns

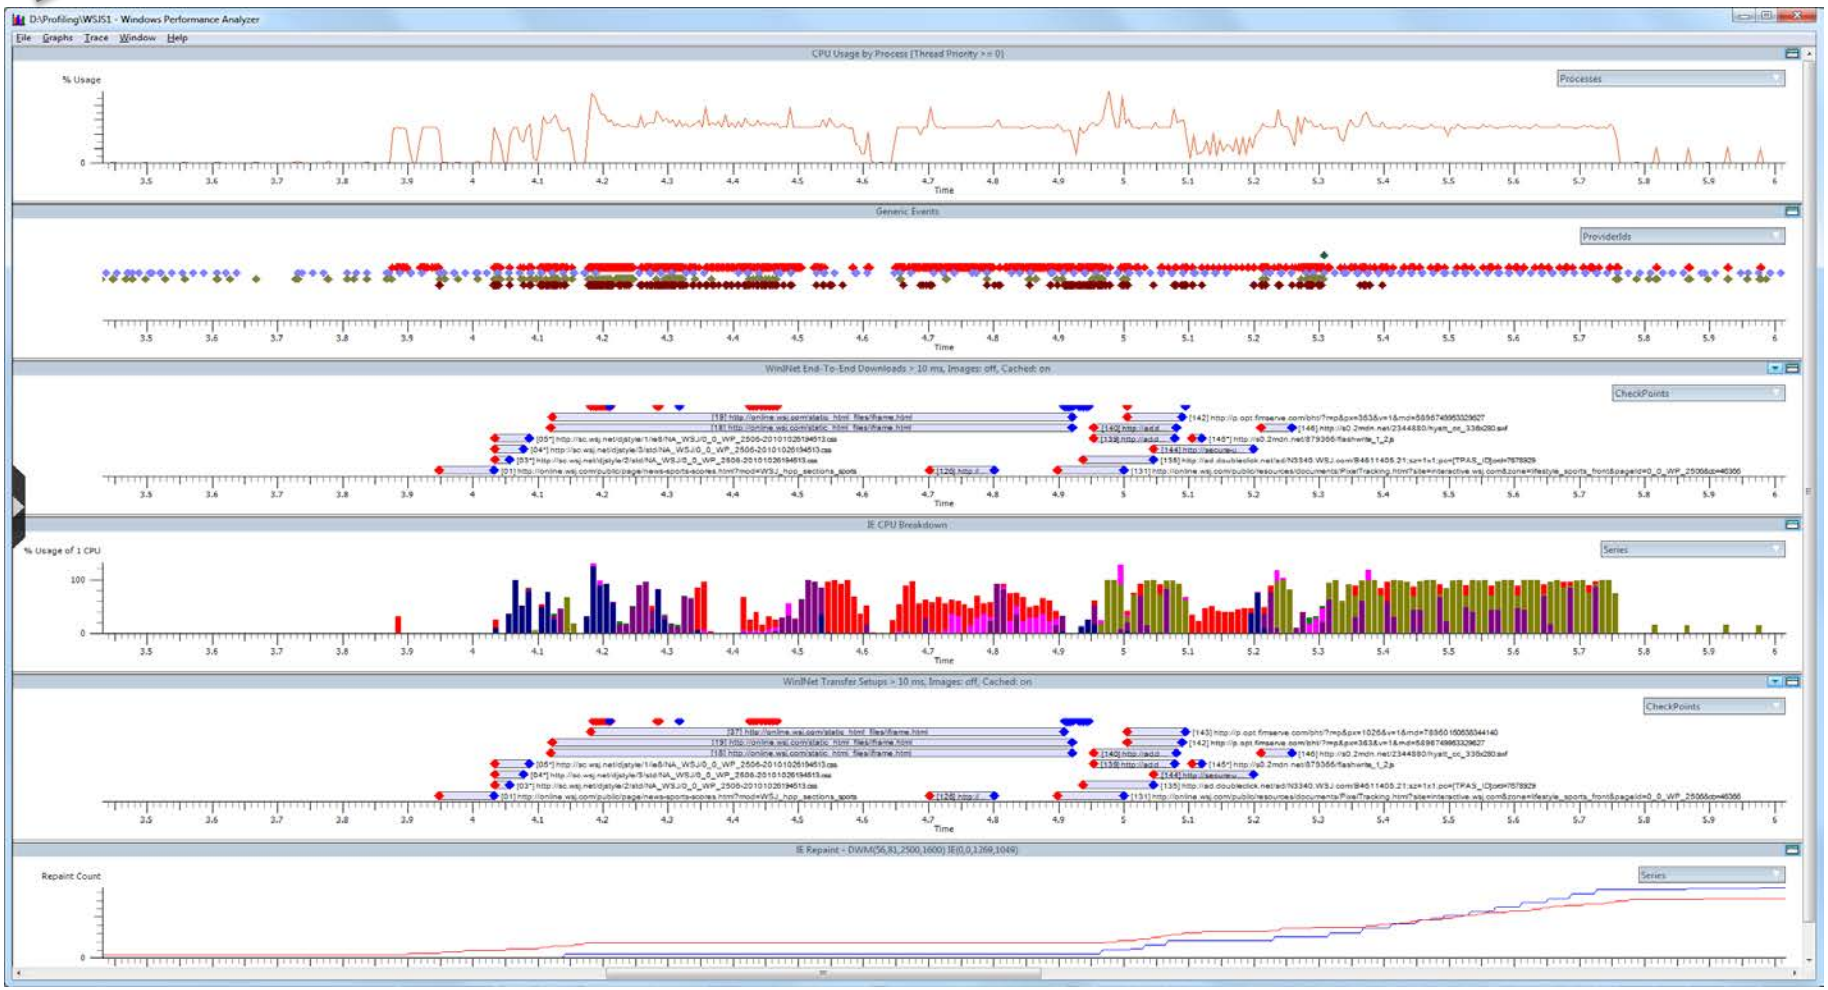

TAKE AWAY
Standard synthetic workloads can be used with all testers and applied to all SSDs.

Synthetic Workloads – Known & Repeatable Stimulus for Standardized Test

Four "corner case" workloads applied to a single SSD - Bandwidth

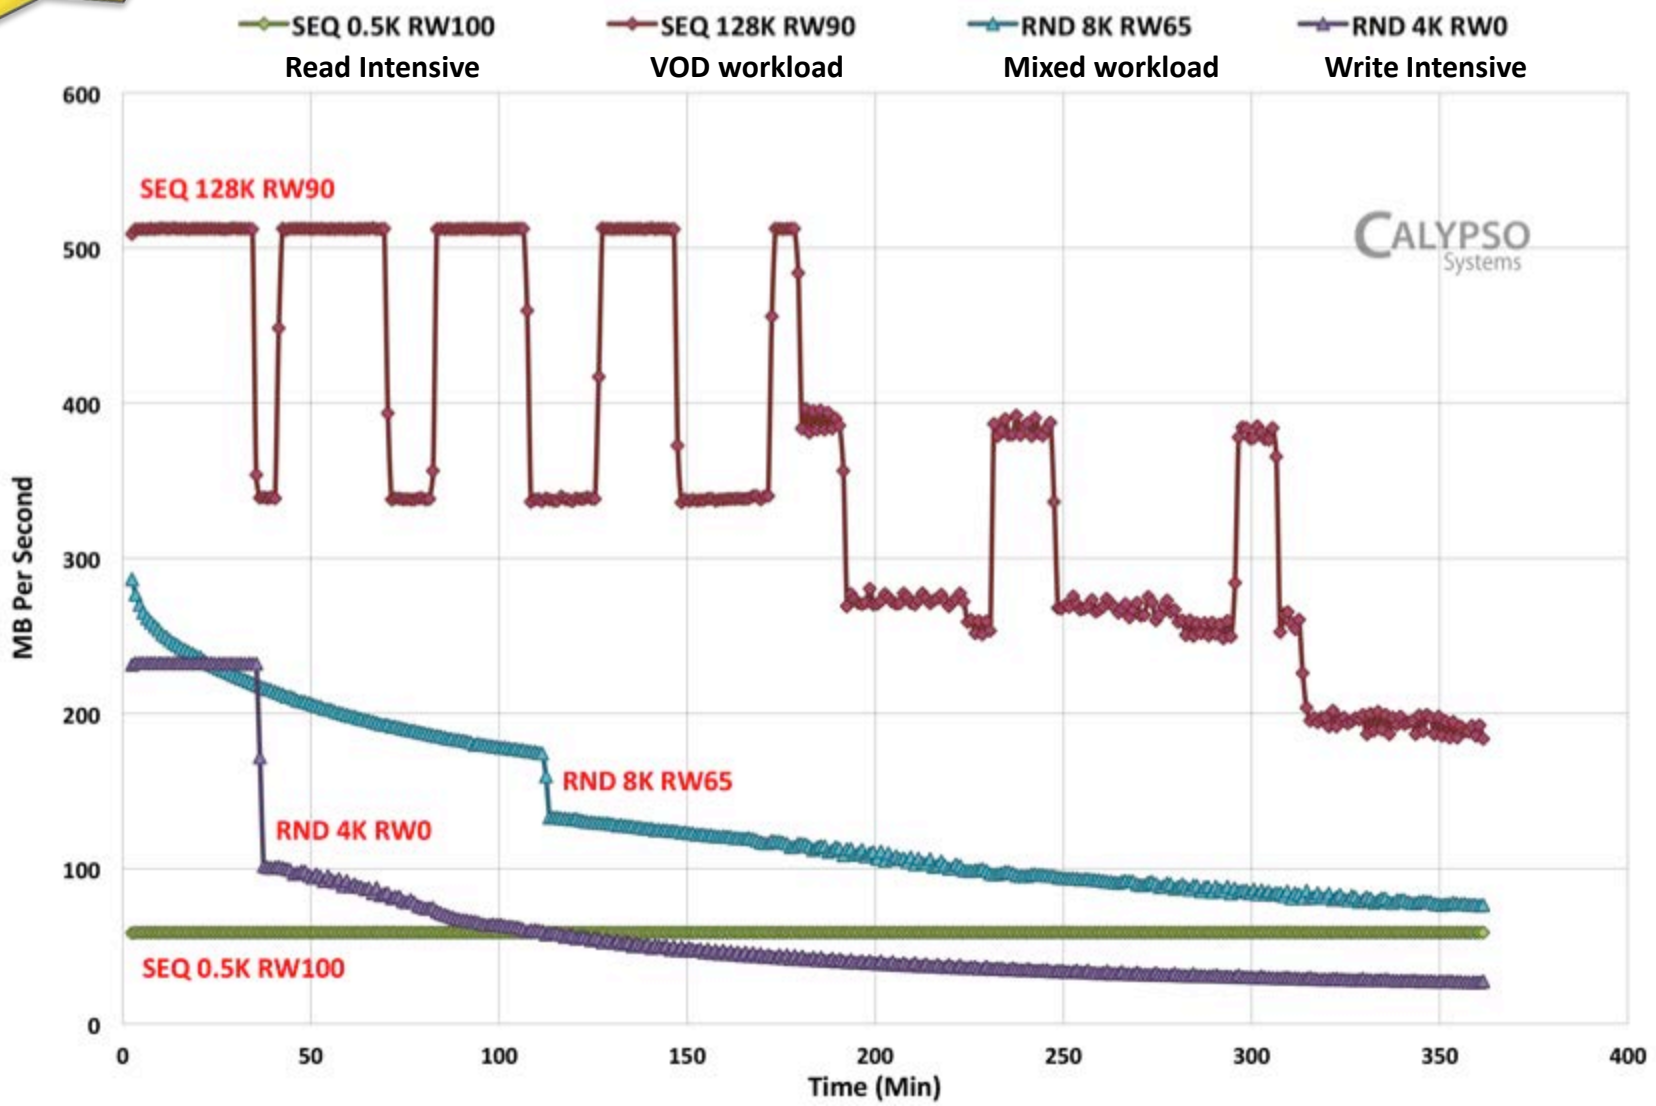

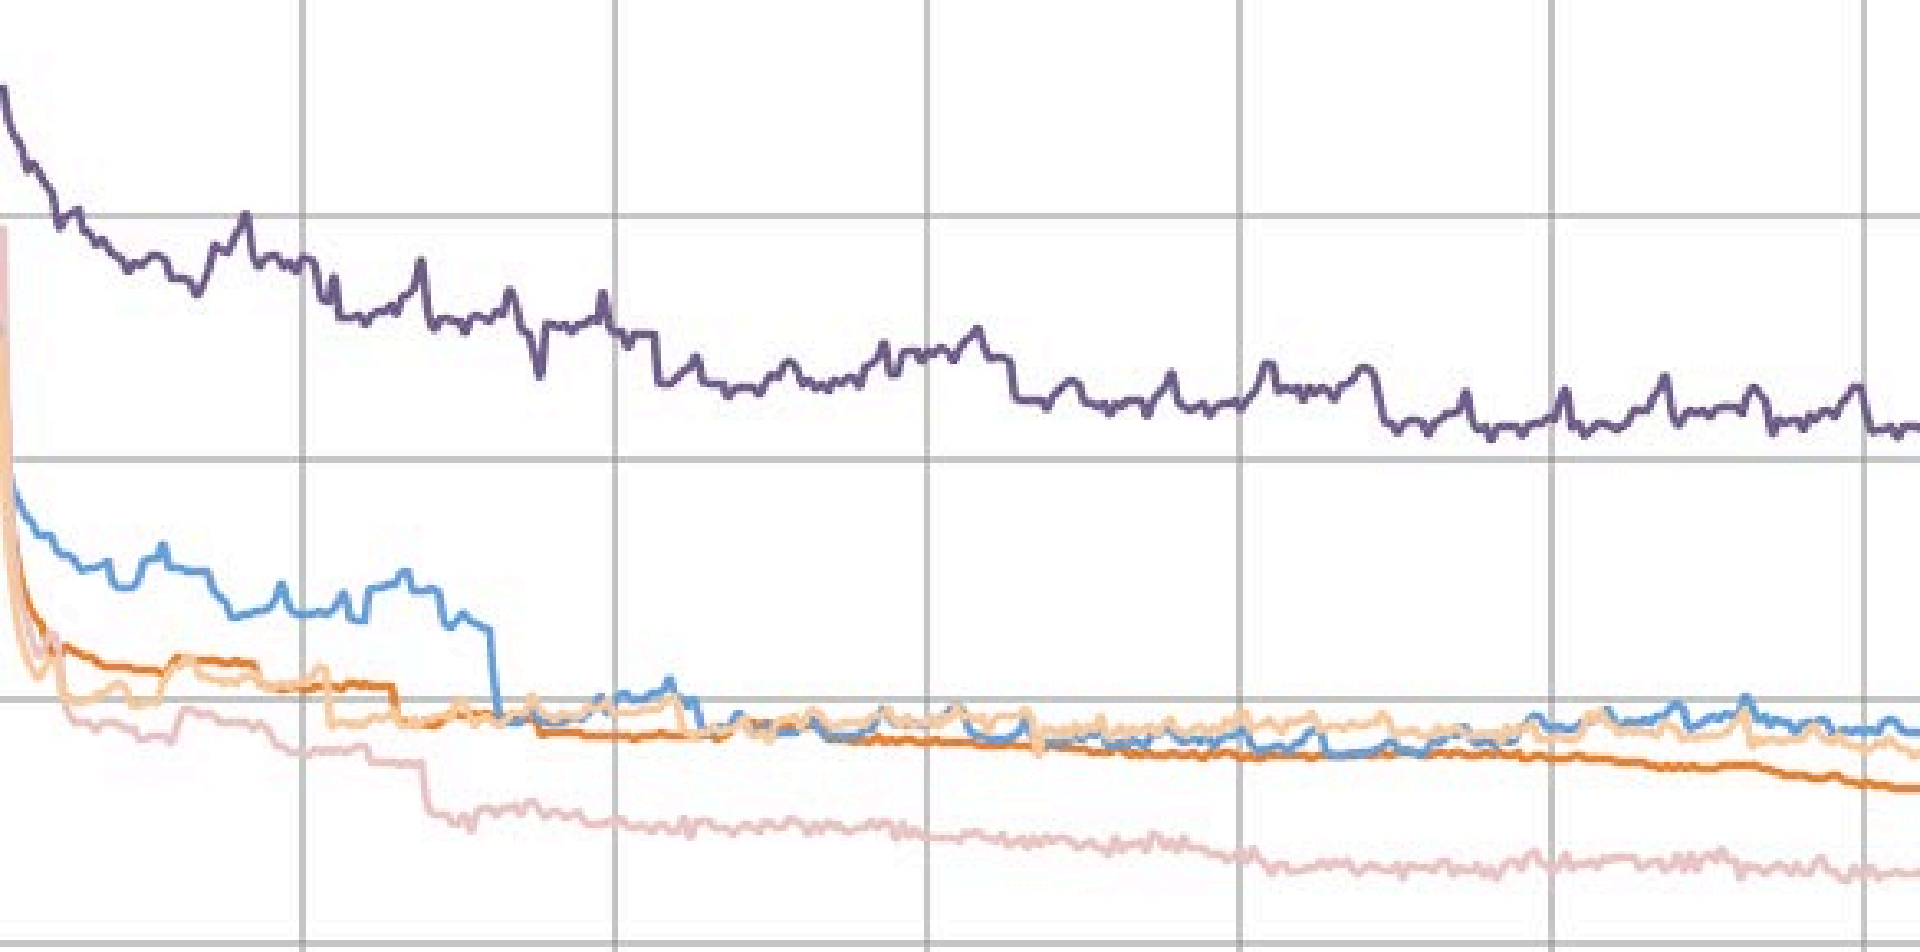

Synthetic Enterprise Application Workloads

WHY ARE THEY USED?

TAKE AWAY
Synthetic Workloads select representative accesses for SSD Performance Comparison.

Synthetic workloads provide a reasonable basis for SSD Comparison

Real World (Trace based / IO based) Workloads:

- Are specific only to the system on which it was captured (***Apples to Oranges***)
- Difficult to account for "idle times"
- Have a large number of streams (can be 50 or more discrete access patterns)

Synthetic Workloads:

- Are repeatable and of known content (***Apples to Apples***)
- Can be standardized for testing on a few known patterns
- ***Provide a basis for SSD Performance comparison***

Webserver 8K	8 KiB	75:25	95:5
Webserver 64K	64 KiB	75:25	95:5
Exchange Mail 4K	4 KiB	100:0	67:33
Exchange Mail 64K	64 KiB	100:0	67:33
Web Server Logs	8 KiB	0:100	0:100
File Servers	8 KiB	75:25	90:10
DB OLTP	8 KiB	100:0	70:30
SQL Logs	64 KiB	0:100	0:100
Media Streaming	64 KiB	0:100	98:2
Archive	2048 KiB	95:5	55:45

Synthetic Enterprise Application Workloads

WHAT ARE SOME EXAMPLES?

Examples of Common Synthetic Enterprise Application Workloads

Application Name	General Description	Example Access Pattern
OLTP VDI	Small Block Mixed RND Workload	RND 4KiB RW65 RND 8KiB RW67 RND 8KiB RW70
VOD	Video Edge Server large block SEQ	SEQ 128KiB RW90 RND 128KiB RW90
Webserver/ MS Exchange Mail	Application page sizes of 64K	SEQ25 64KiB RW95 RND 64KiB RW0
Webserver Logs	Logging workloads for RDBMS	SEQ 8KiB RW0
Decision Support Service (DSS)	Heavily Indexed large table structures for real time access	RND 64KiB RW100
OS Paging Media Streaming Medical Imaging	OS data requests from storage in medium block SEQ transfers	SEQ 64KiB RW90 SEQ 64KiB RW98 SEQ 1,024KiB RW05

Workloads can be a single access pattern or a composite of many different access patterns

Synthetic Enterprise Application Workloads

CASE STUDY

Building a synthetic workload test

Objectives

- Define a procedure to test and measure ***Synthetic application workloads***
- Use ***SNIA PTS test methodologies*** for Pre-conditioning & Steady State
- ***Isolate variables*** for IOPS & RT Confidence measurement
- ***Create a template*** to modify / substitute synthetic workloads of interest

Test Plan

1. Define a synthetic OLTP database workload
2. Run the OLTP workload to steady state
3. Run Demand Intensity (DI) Sensitivity workload loops
4. Measure IOPS, ART & MRT at each QD for each loop
5. Plot Confidence Level Histograms for each OIO point
6. Determine IOPS & Response Time Saturation point(s)
7. Compare OIO “sweet spots” for different DI workload loops

Test Set Up

Hardware Platform – PTS RTP 3.0

- PTS 1.1 Reference Test Platform
- OS: CentOS 6.5
- Test Software: CTS BE ver 1.9.184
- Motherboard: Intel Gen3 E5
- CPUs: Dual S2697W 3.1Ghz 8 core
- HBA: 12Gb/s LSI 9300

Software Platform – CTS 6.5

- Test Software: CTS FE ver 1.18.11
- OS: Windows 7 Pro

Reference Test Platform for the SSS PTS specification
RTP/CTS used to develop and validate the PTS
Certified SNIA SSSI PTS test labs RTP

Test Flow

1. **PURGE** (Security Erase)
2. **Settings**: Data Pattern=RND; LBA Range=(0,100)
3. **Pre-condition** – WIPC 2x user capacity SEQ 128KiB RW0; T1Q32
4. **Run Workload Dependent PC** – RND 8KiB RW65; T4Q8
5. **Run to Steady State**
   - 5 consecutive one minute Rounds
   - Each Round separated by 29 min of WDPC pre-writes
   - Least squares linear fit no greater than 20% data excursion nor exceeding a 10% slope
6. **Set Restricted LBA zones**
   - 50% of the IOs to the first 5% of the LBAs
   - 30% of the IOs to the next 15% of the LBAs
   - 20% of the IOs to the last 80% of the LBAs
7. **Run Workload Segment loops** w/ varying OIO:
   - Thee TC loops per drive: T2, T4, T8 or T4, T8, T16
   - Step QD 2,4,8,16,32
8. **Plot IOPS, ART, MRT** at each OIO
9. **Plot Confidence Level Histogram & RT Ceiling** for each OIO
10. **Compare Optimal OIO** sweet spots between different TC workload segment loops

What are Confidence Levels?

Response Time (RT) Histograms

Confidence Level Plots (CLP)

CLP Comparison Plots (CLPC)

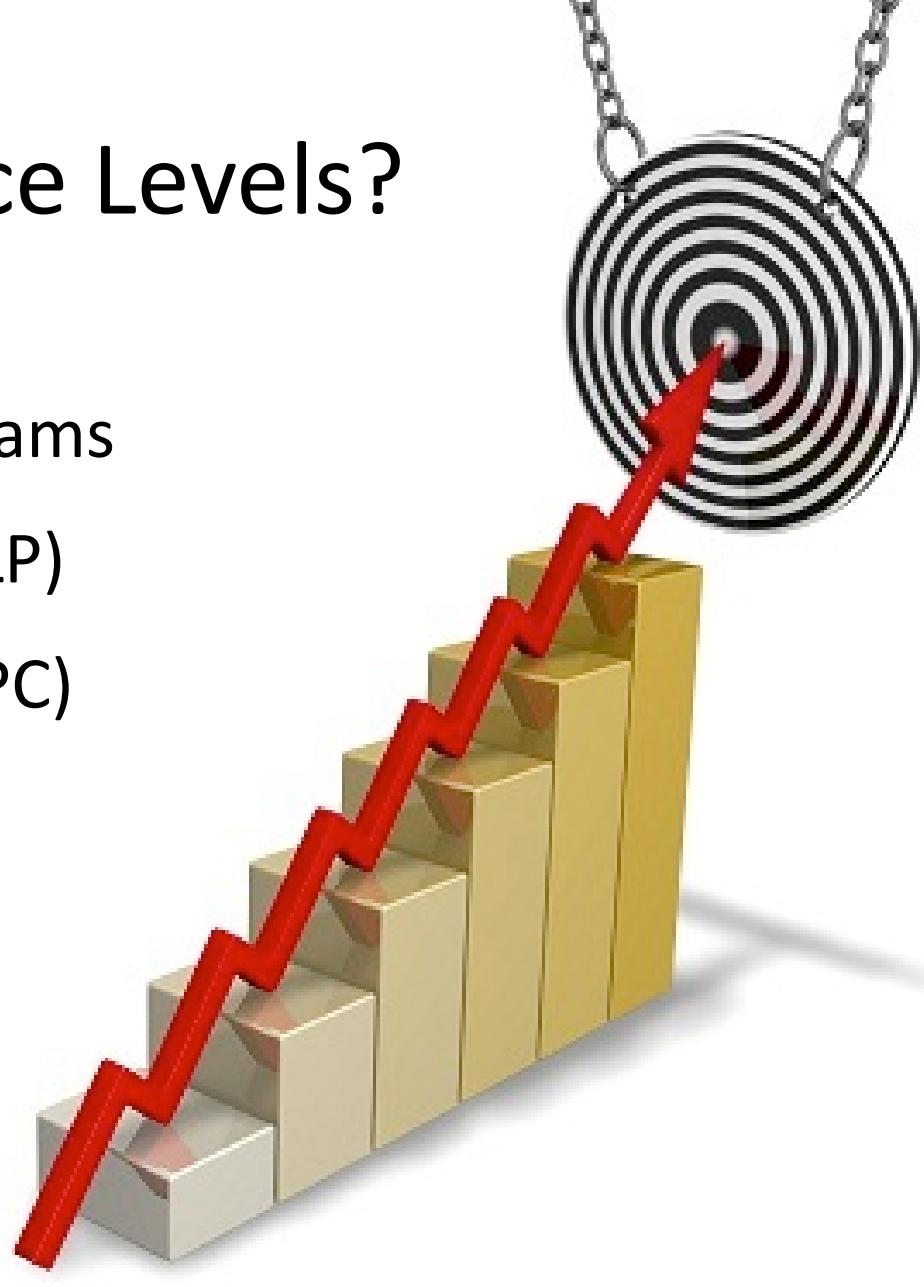

Confidence Levels – Quality of Service ("Q o S")

IO Rate Per Sec	Measurement Period	Total IOs	% Confidence Level	Q o S No. "9's"	No. of Dropped IOs
One Minute Measurement Period					
10,000	One Minute	600,000	99.999%	5 9's	6
10,000	One Minute	600,000	99.99%	4 9's	60
10,000	One Minute	600,000	99.9%	3 9's	600
Ten Minute Measurement Period					
10,000	Ten Minutes	6,000,000	99.999%	5 9's	60
10,000	Ten Minutes	6,000,000	99.99%	4 9's	600
10,000	Ten Minutes	6,000,000	99.9%	3 9's	6,000

"Five 9's" = 99.999% or 99,999 of 100,000 events

Response Time Histogram with Confidence Level Percentages

TAKE AWAY
Figures of Merit are the Response Time Bars that show cumulative IOs at a given mS response time.

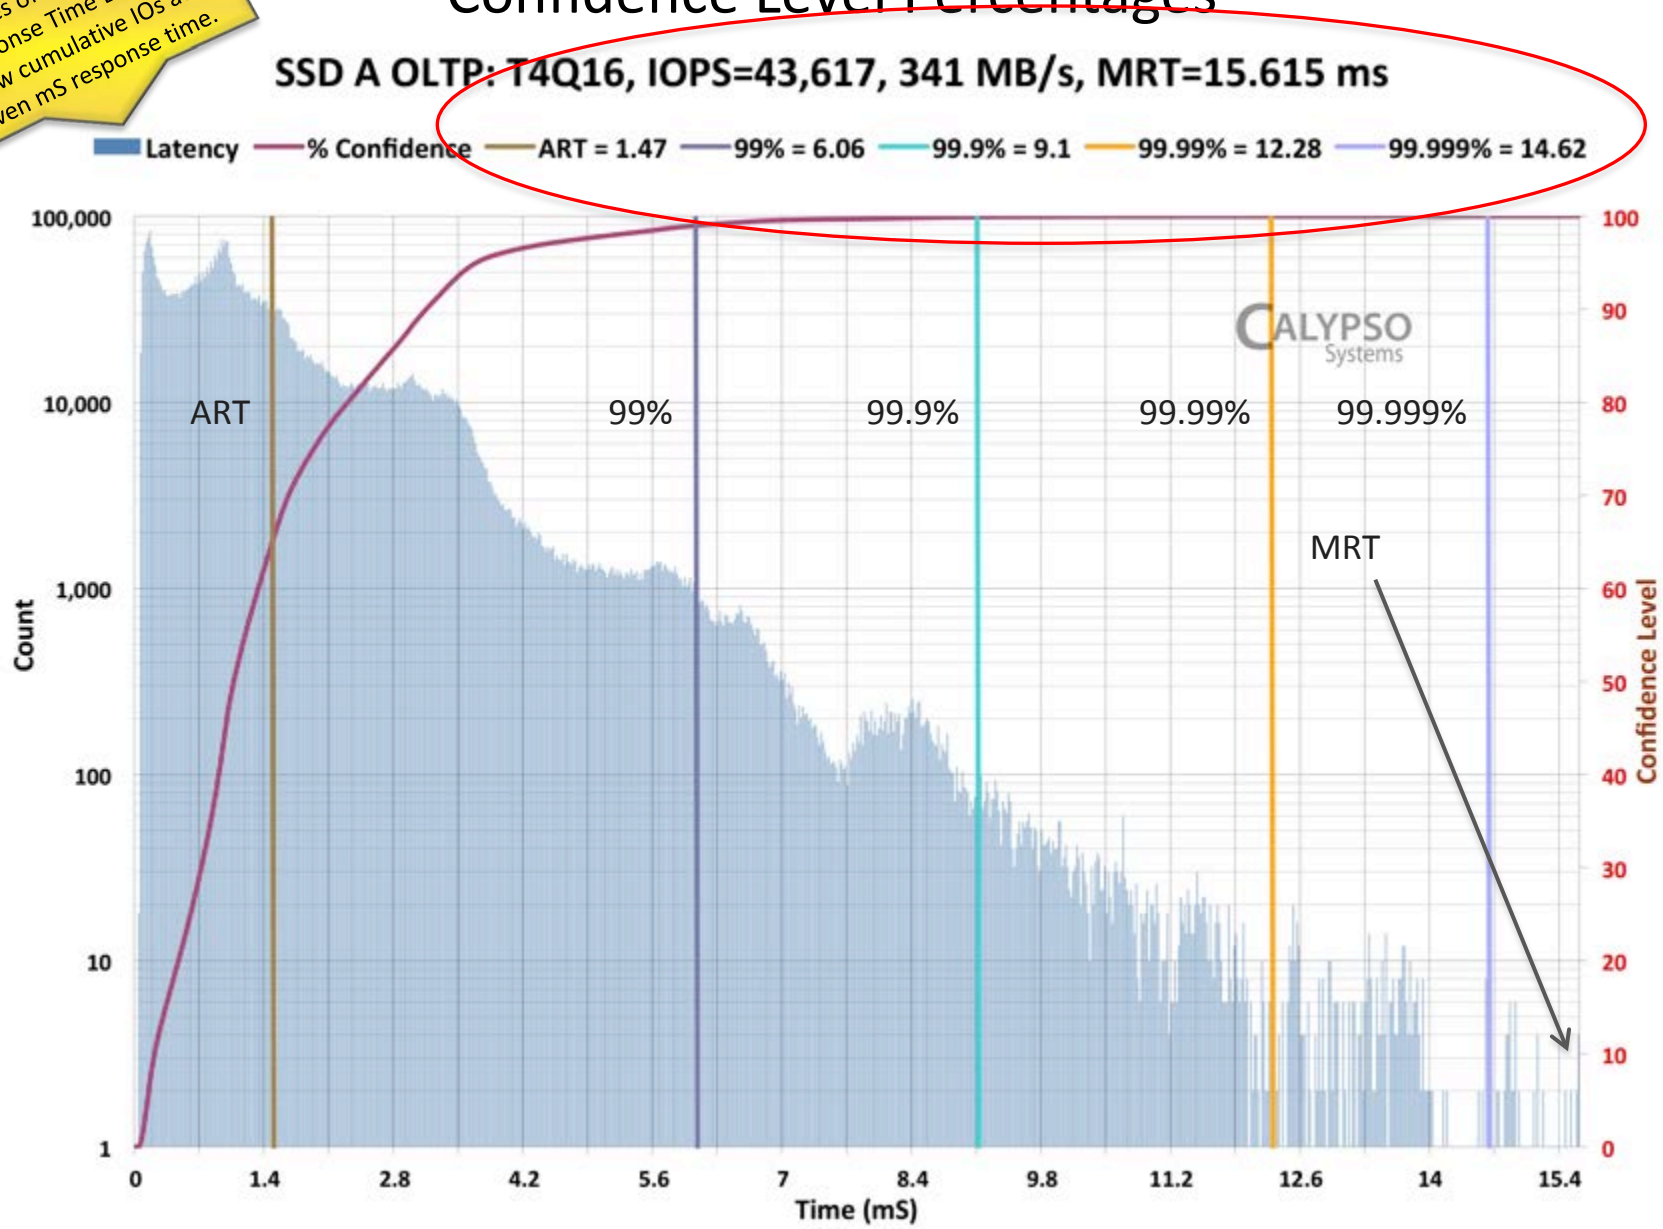

TAKE AWAY
Histogram data can be plotted and compared on a Bar chart. RT Ceiling is set per application.

Histogram Compare Plot with RT Ceiling

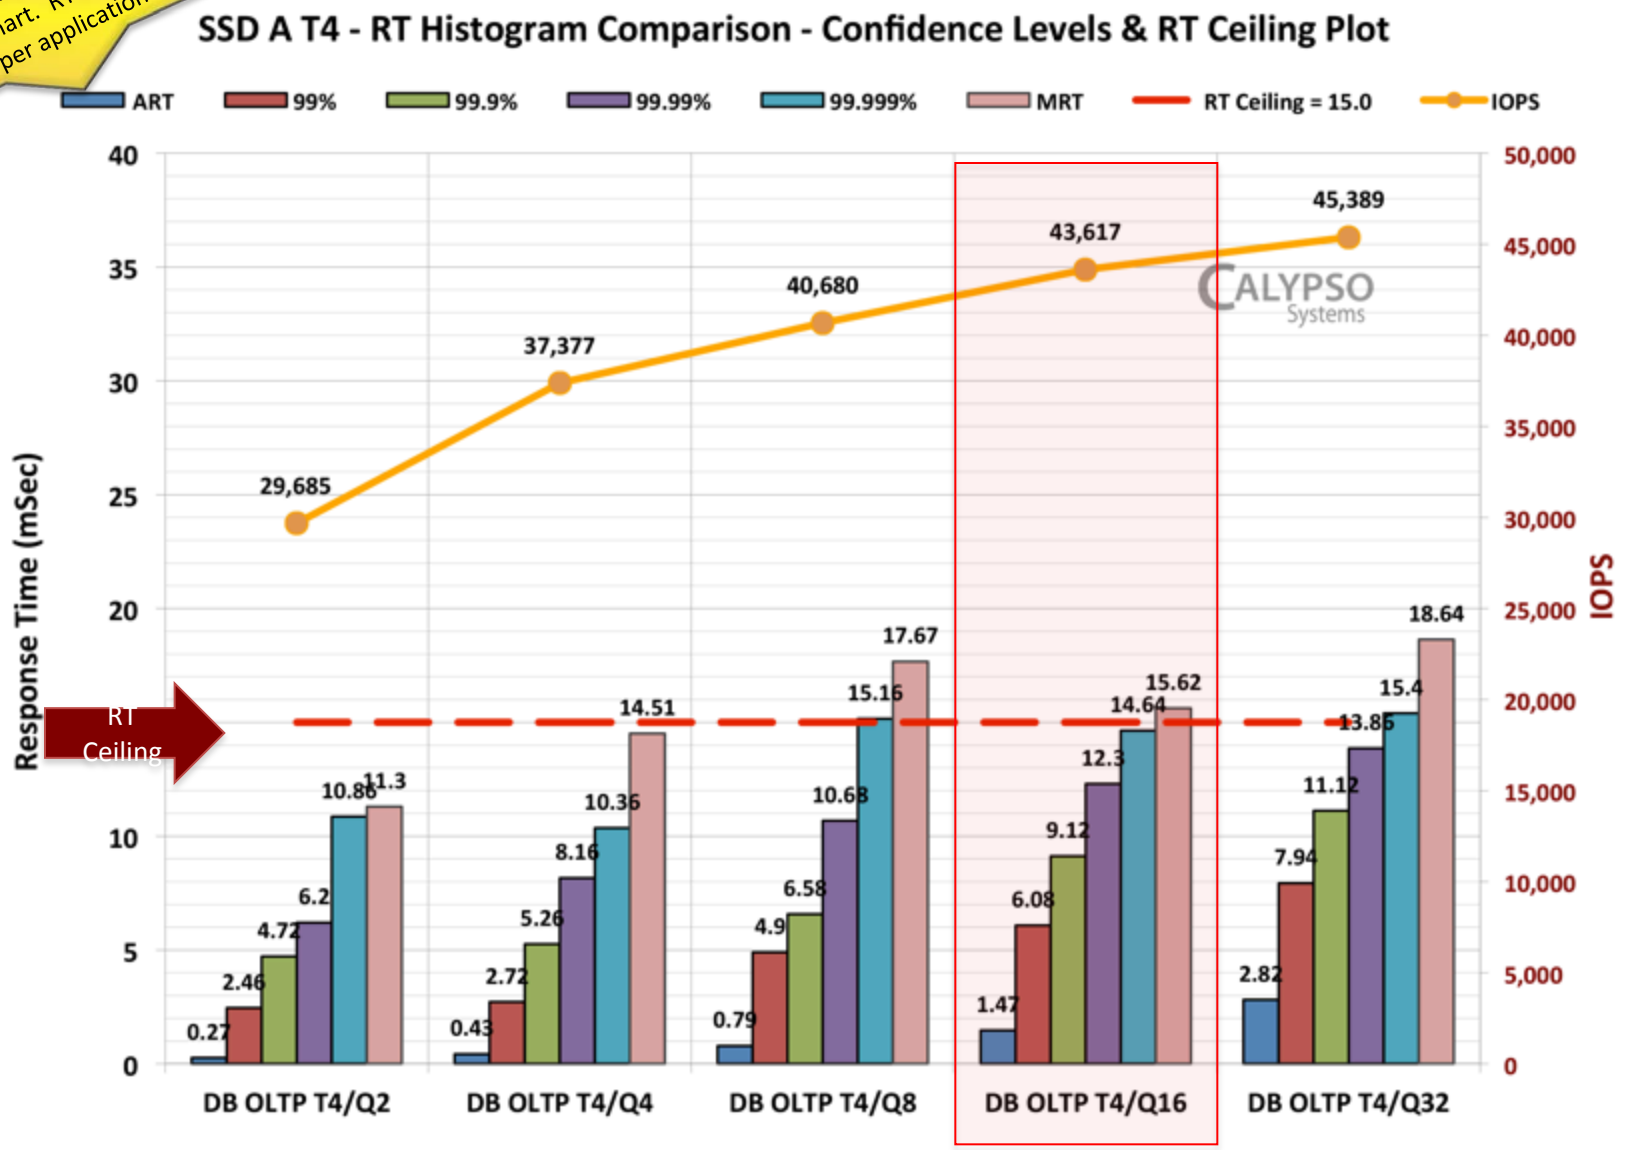

OLTP Saturation Test with Demand Intensity (DI) Sensitivity Analysis

Workload loops with Varying Demand Intensity

SSD A: T4, T8, T16

SSD B: T2, T4, T8

Find your operating Sweet Spot!

800 GB Enterprise class 2.5” SATA

SSD A

TAKE AWAY
DI Sensitivity test measures workload segments at increasing TC

SSD A – Test Profile: IOPS v Time

Workload Segments T4, T8, T16

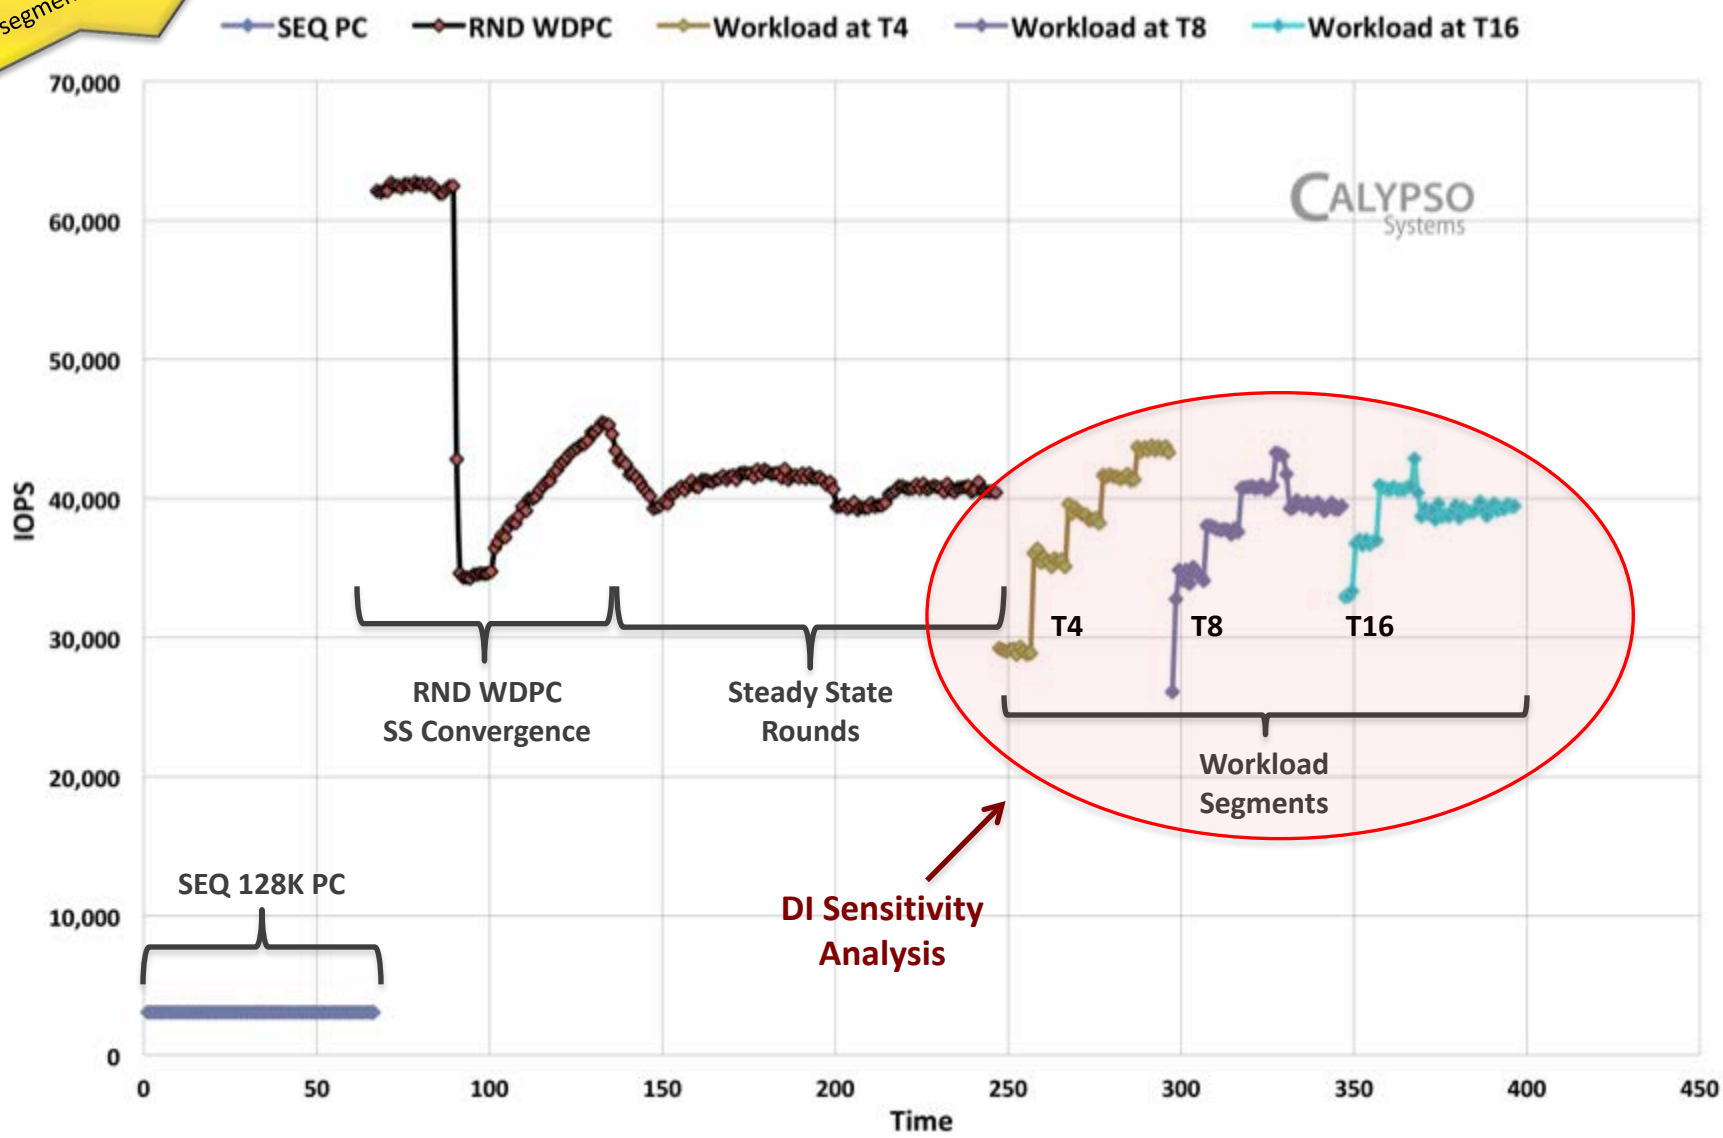

SSD A: SEQ128K PC

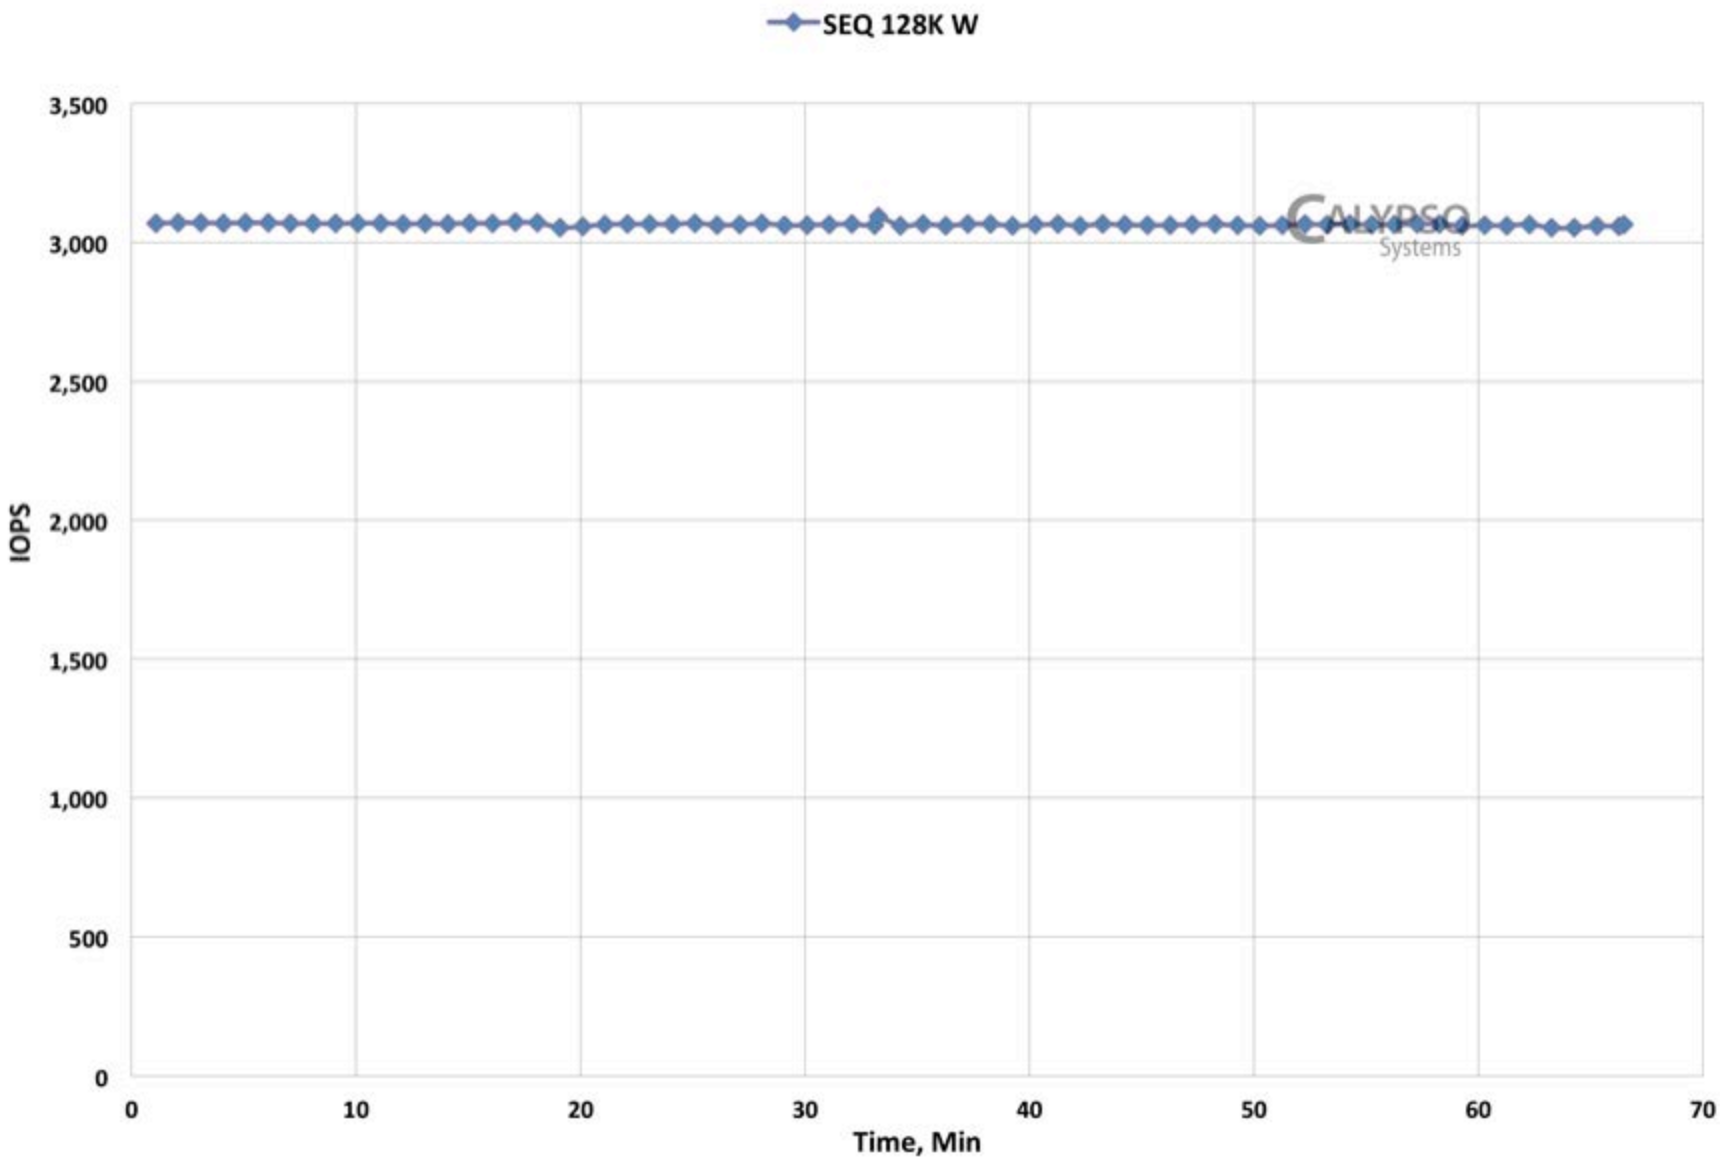

SSD A: WDPC & Steady State

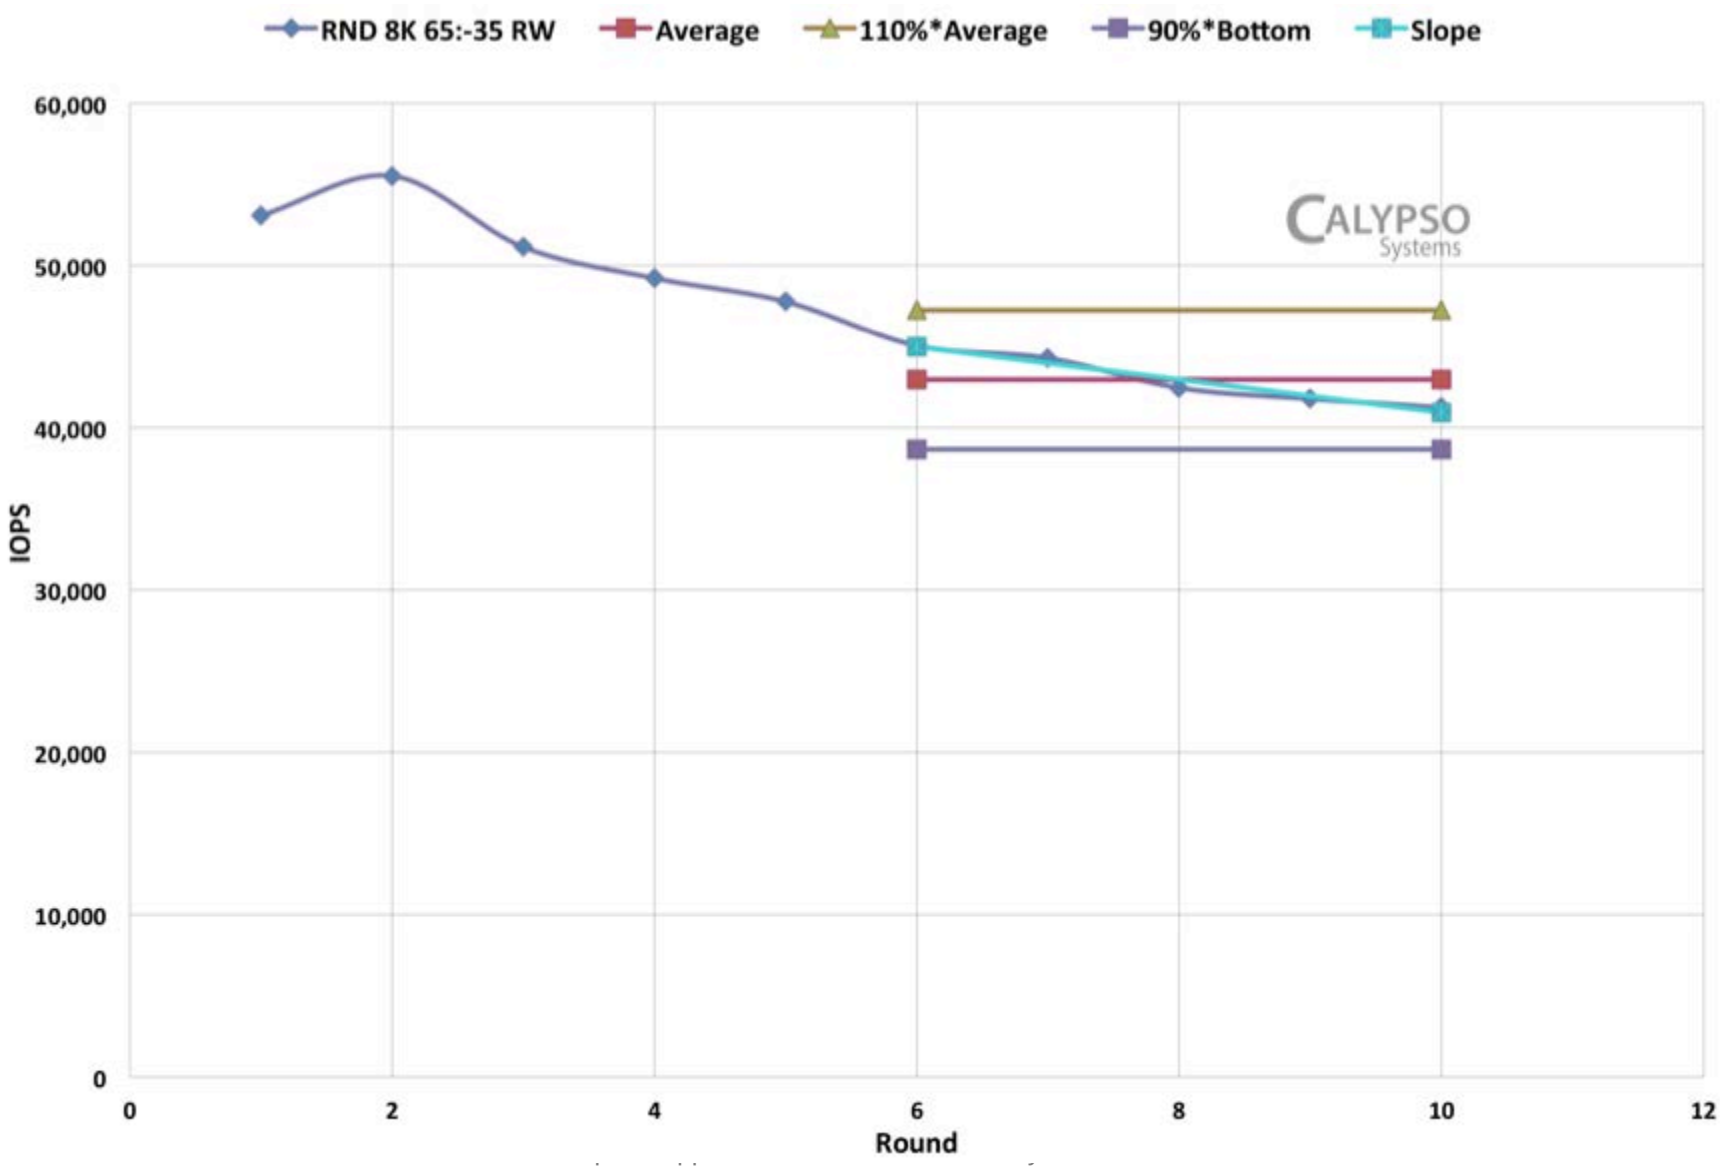

TAKE AWAY
T4 is not enough OIO to Saturate.

SSD A T4: QD Histogram Compare

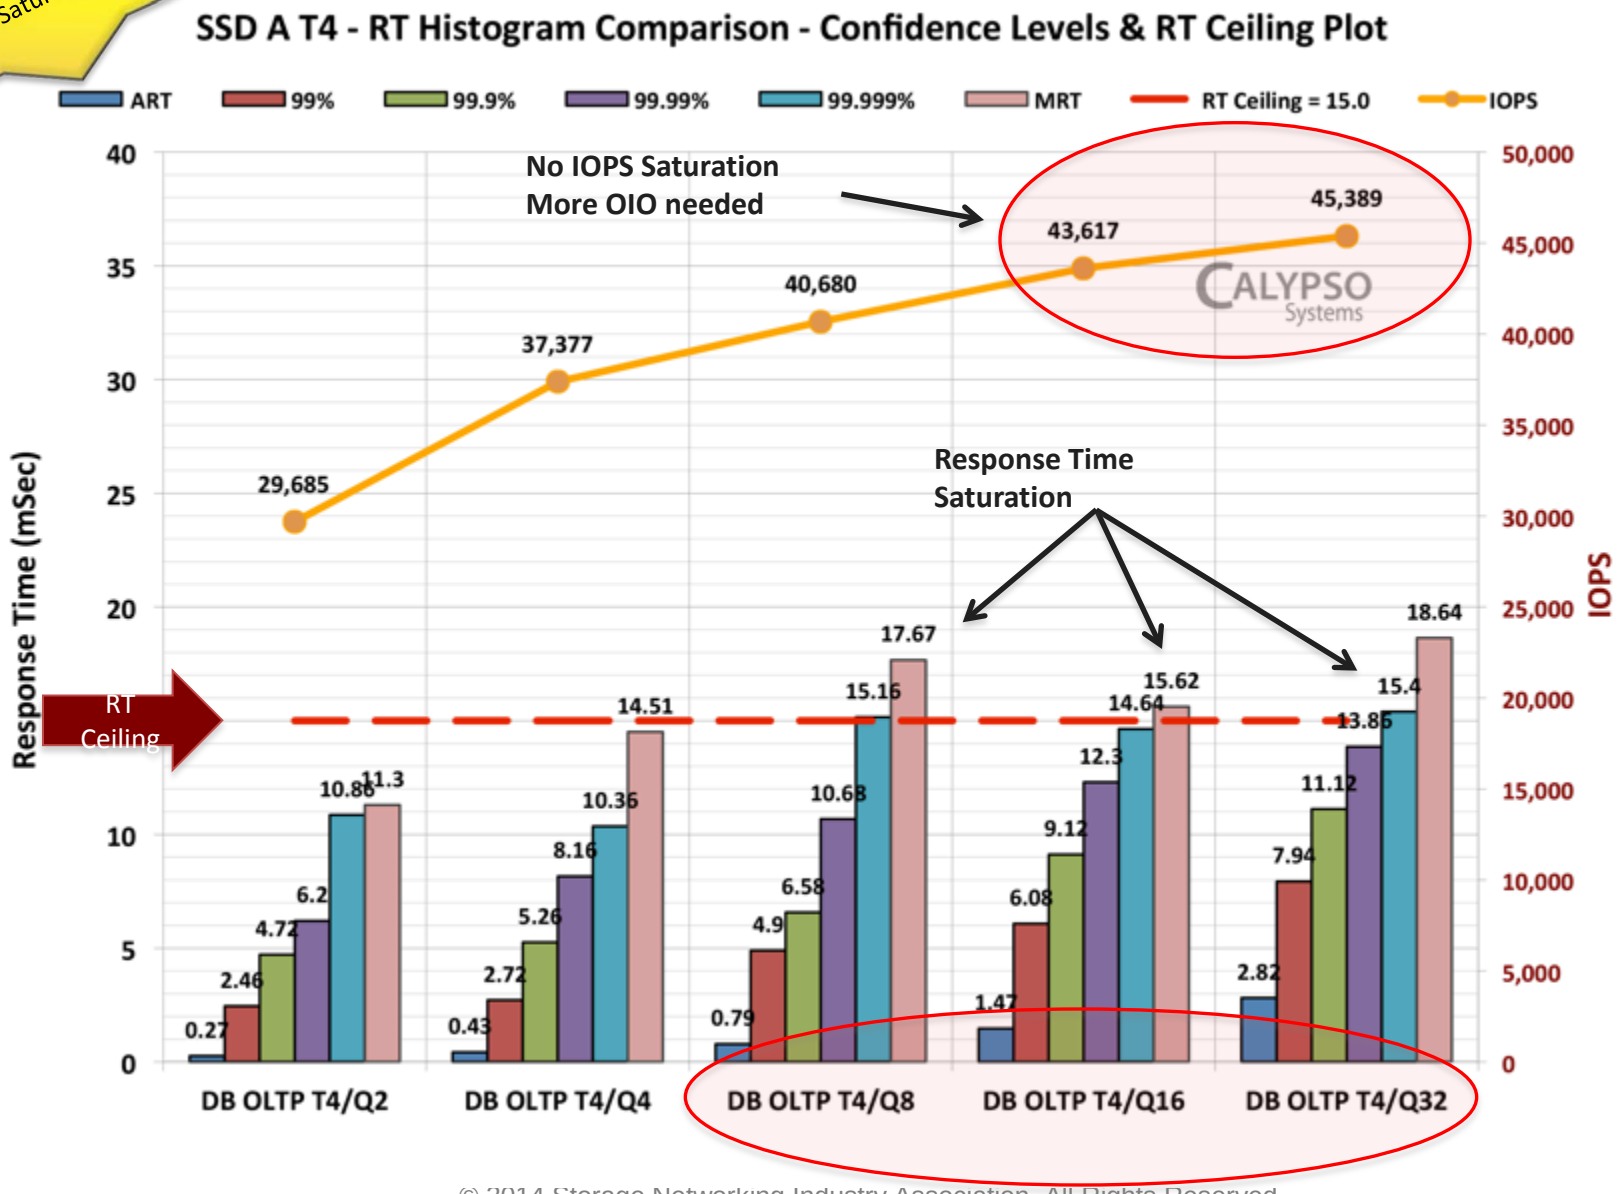

TAKE AWAY
IOPS Saturation begins at T8Q16 RT Saturation at 99.999% T8Q32.

SSD A T8: QD Histogram Compare

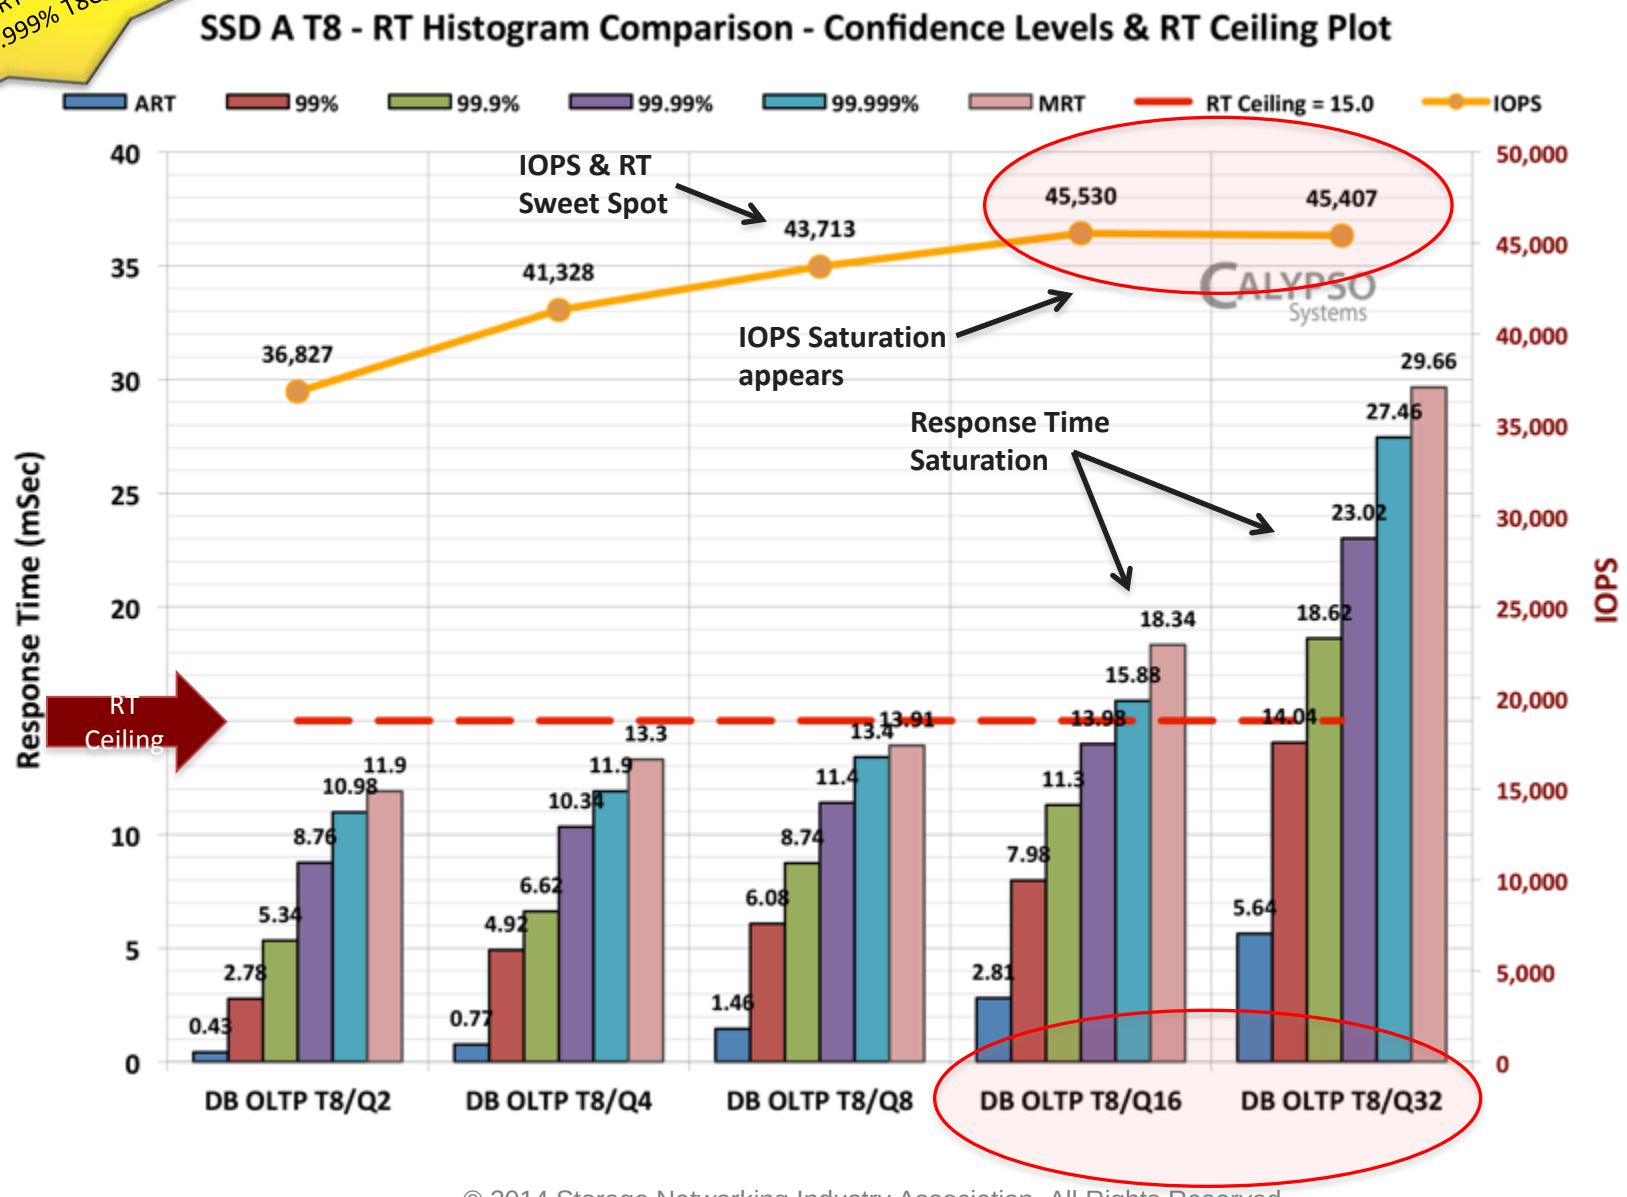

TAKE AWAY
SSD at T16 fully Saturates.
Note increase in RT scale and decrease in IOPS.

SSD A T16: QD Histogram Compare

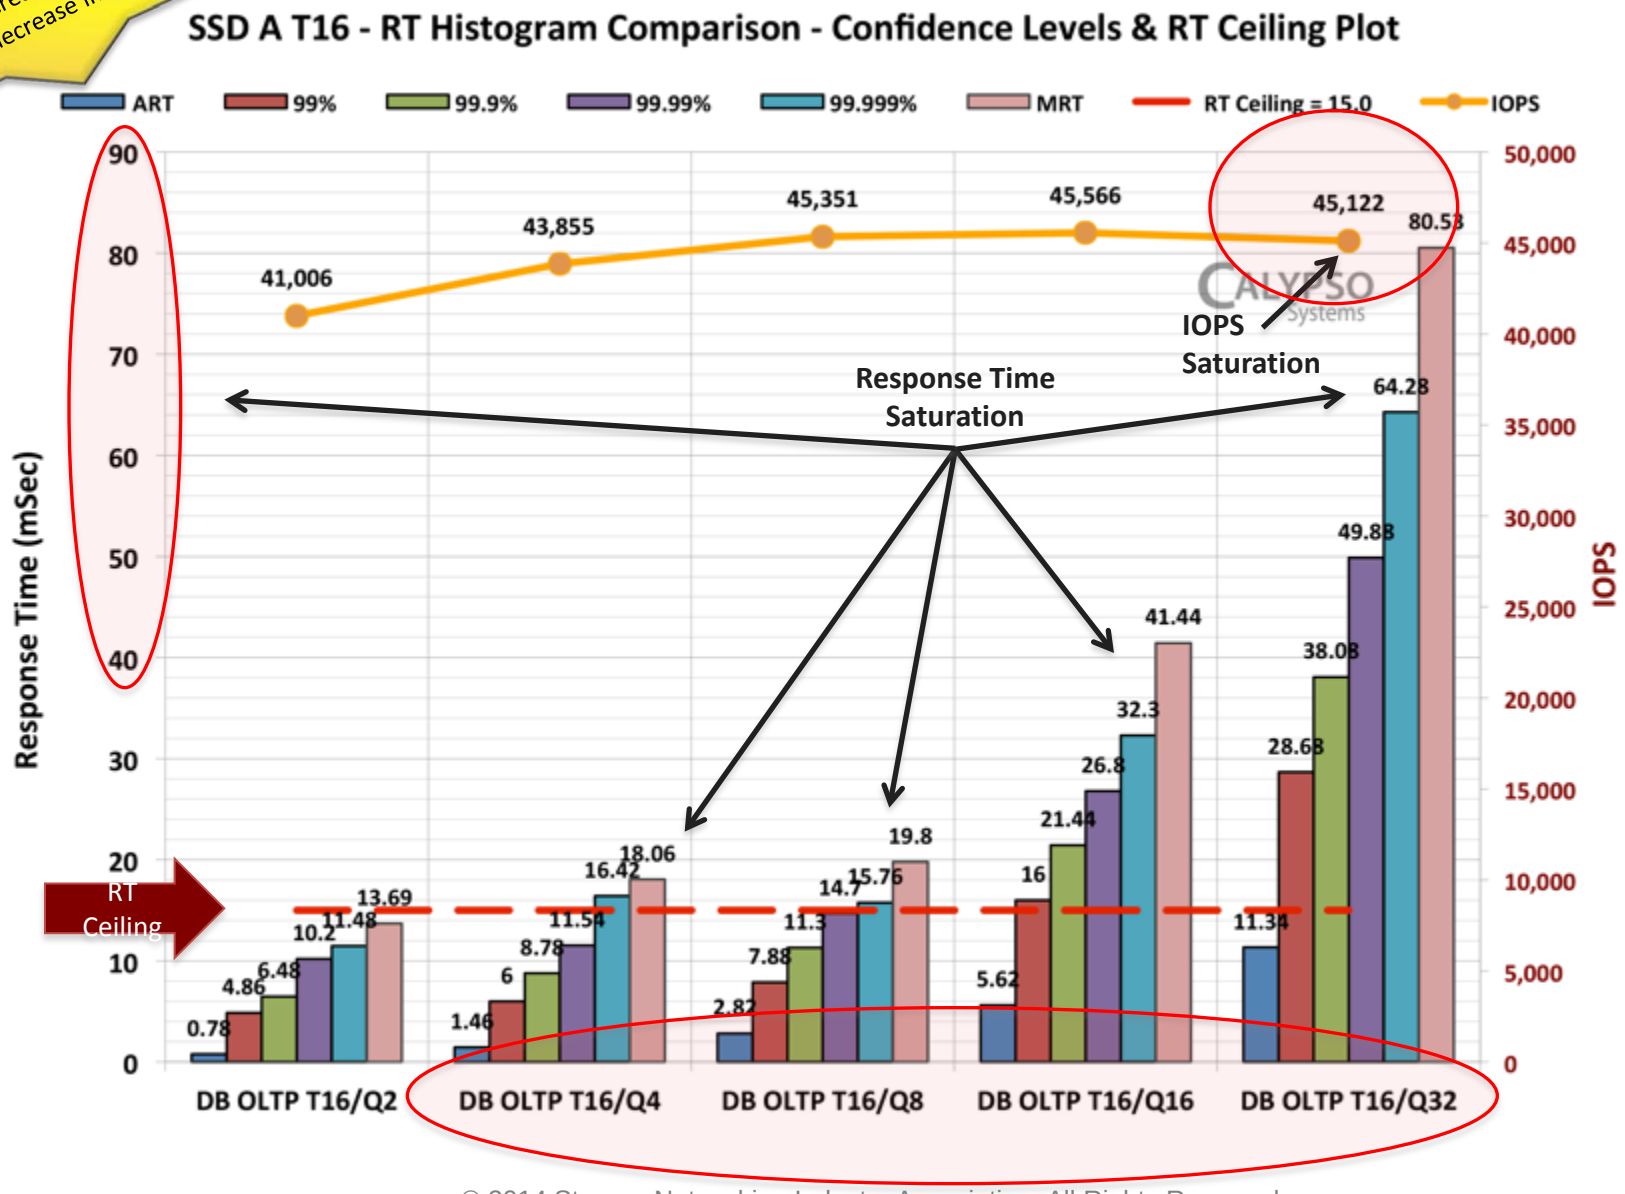

200 GB Enterprise class 2.5” SATA

SSD B

TAKE AWAY
SSD B DI Sensitivity is lower and tests will set OIO at T2, T4 and T8.

SSD B – Test Profile: IOPS v Time

Workload Segments T2, T4, T8

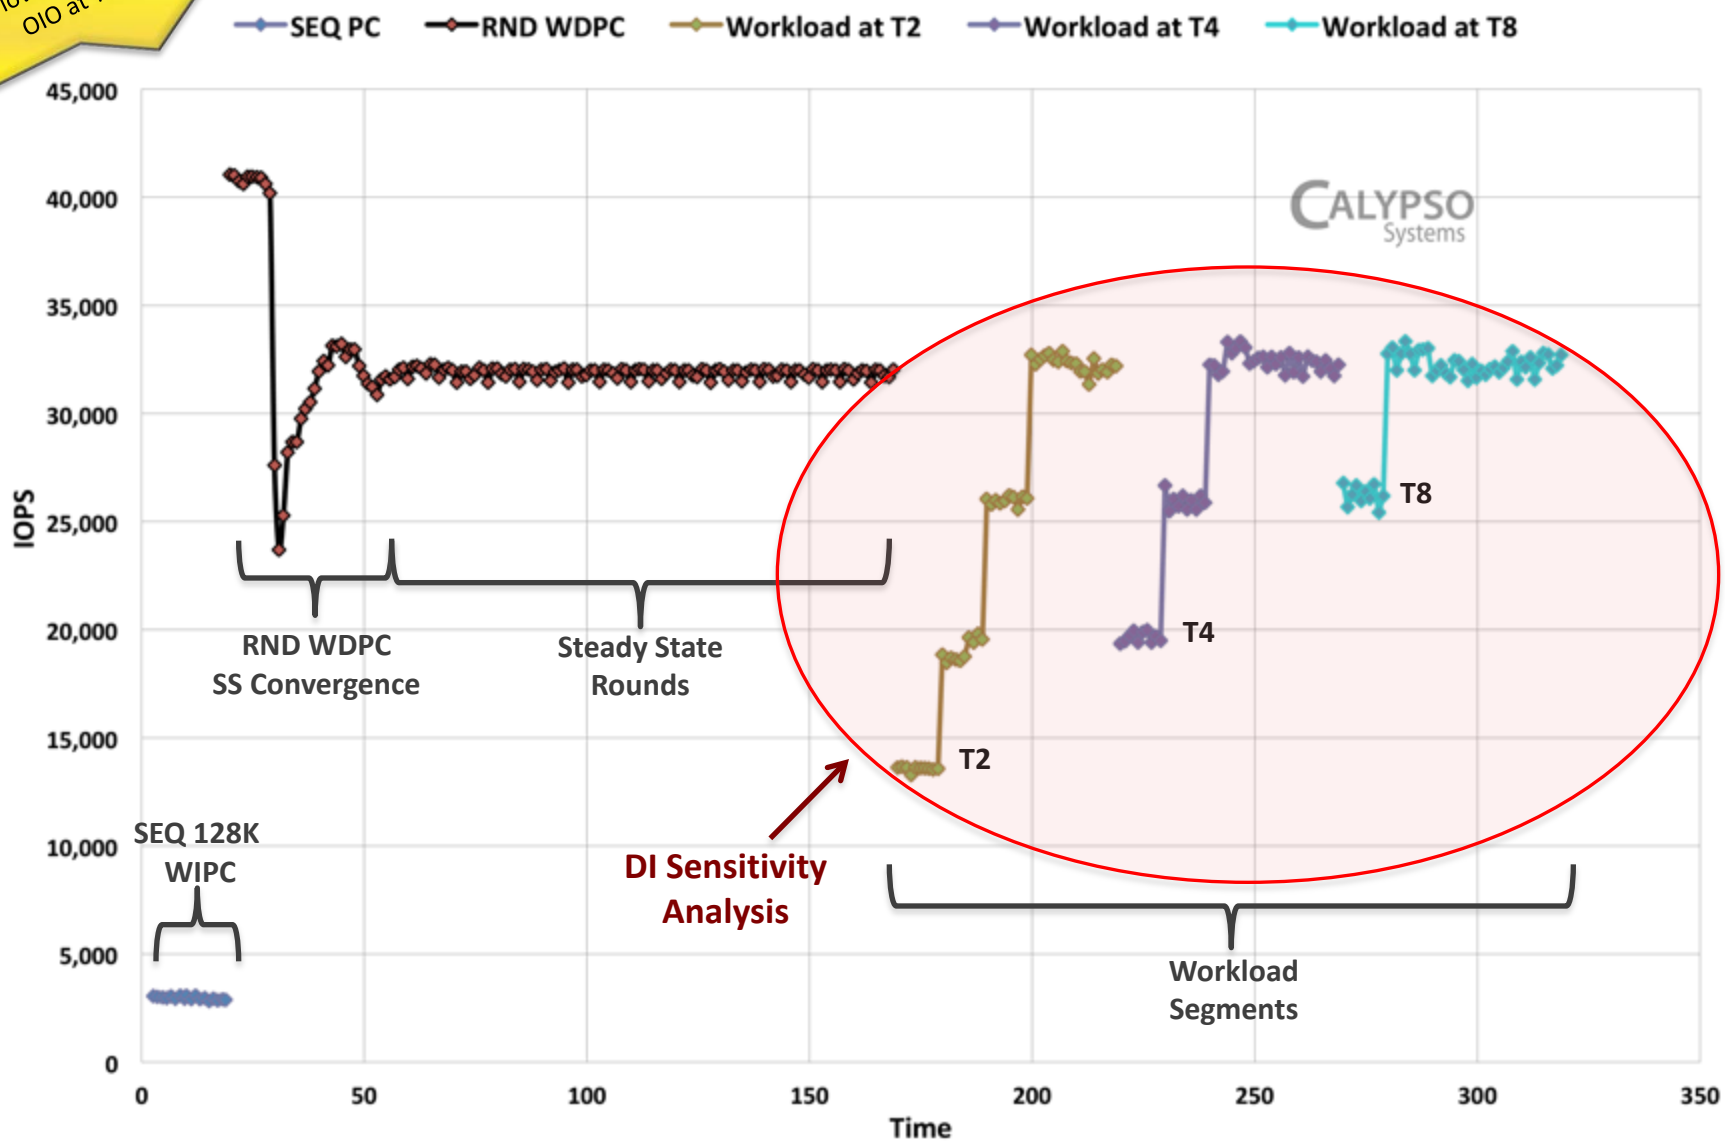

TAKE AWAY
SSD B OIO T2 shows Saturation. Note low OIO at T2Q2.

SSD B T2: QD Histogram Compare

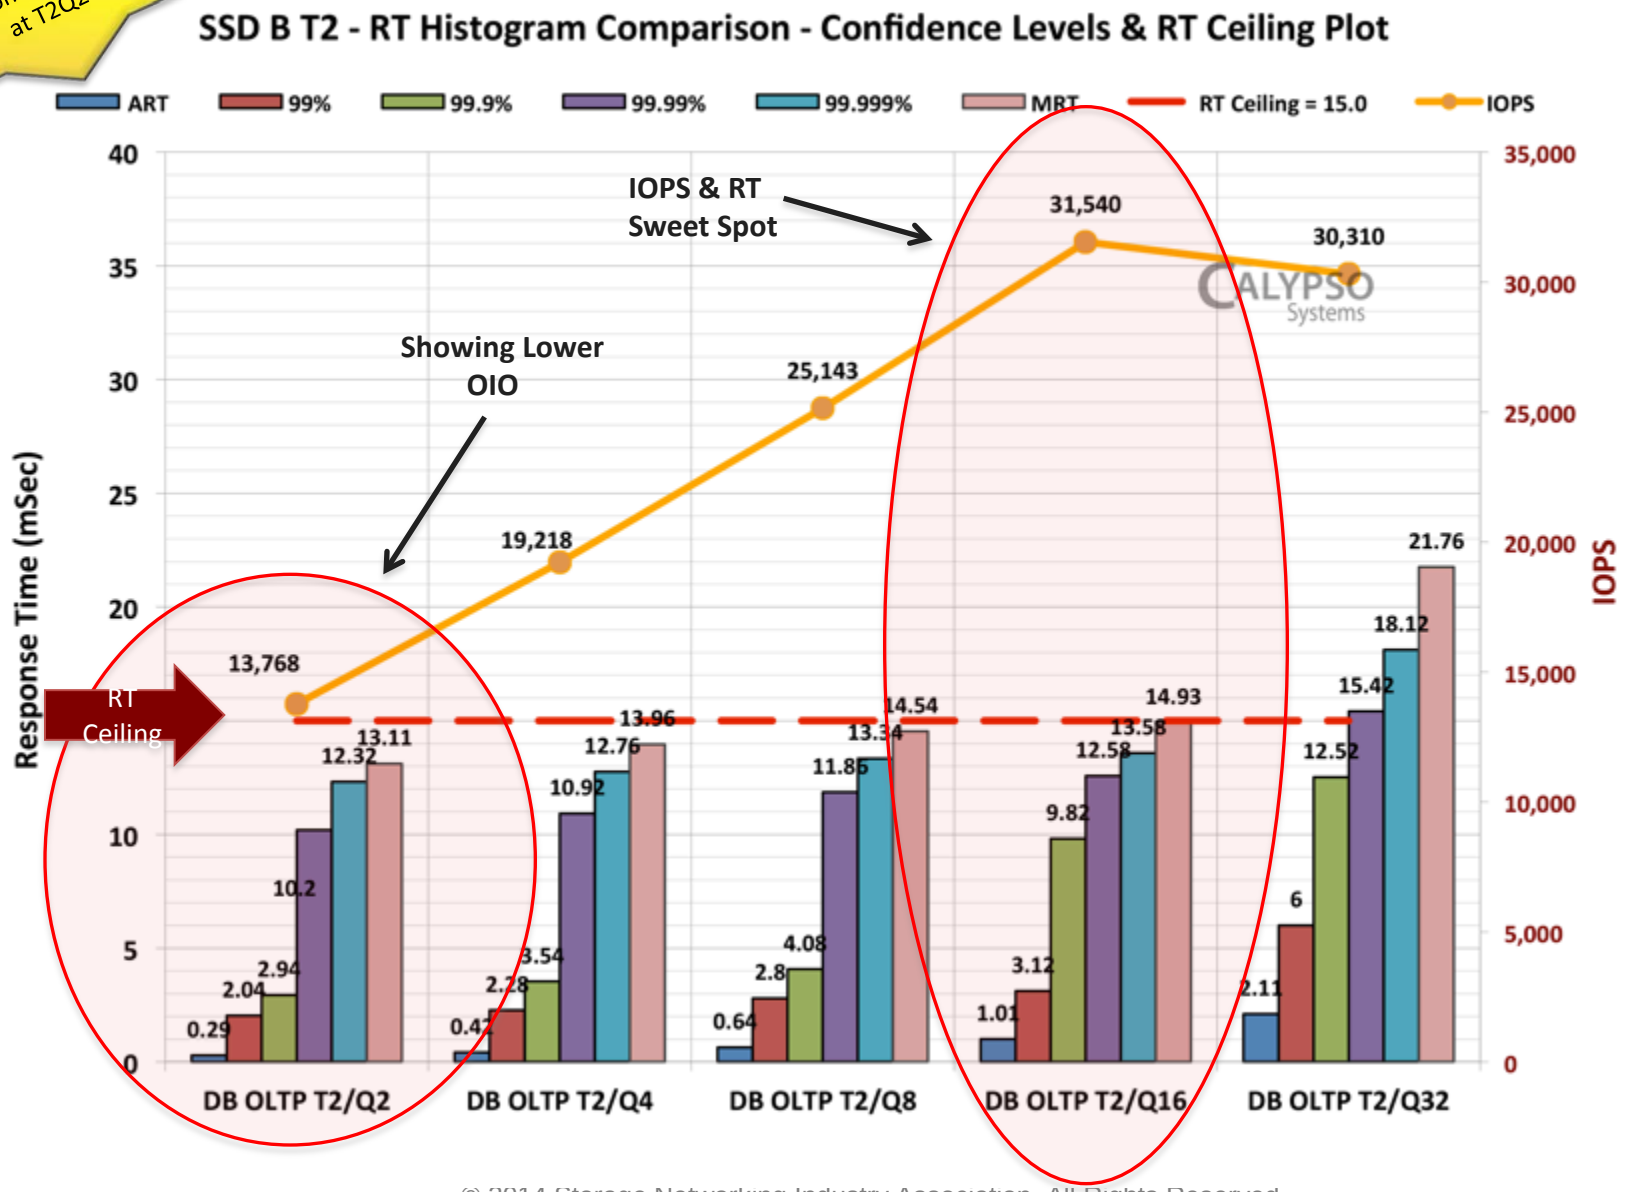

TAKE AWAY
SSD B at T4 shows Saturation. Note OIO Sweet Spot T4Q8.

SSD B T4: QD Histogram Compare

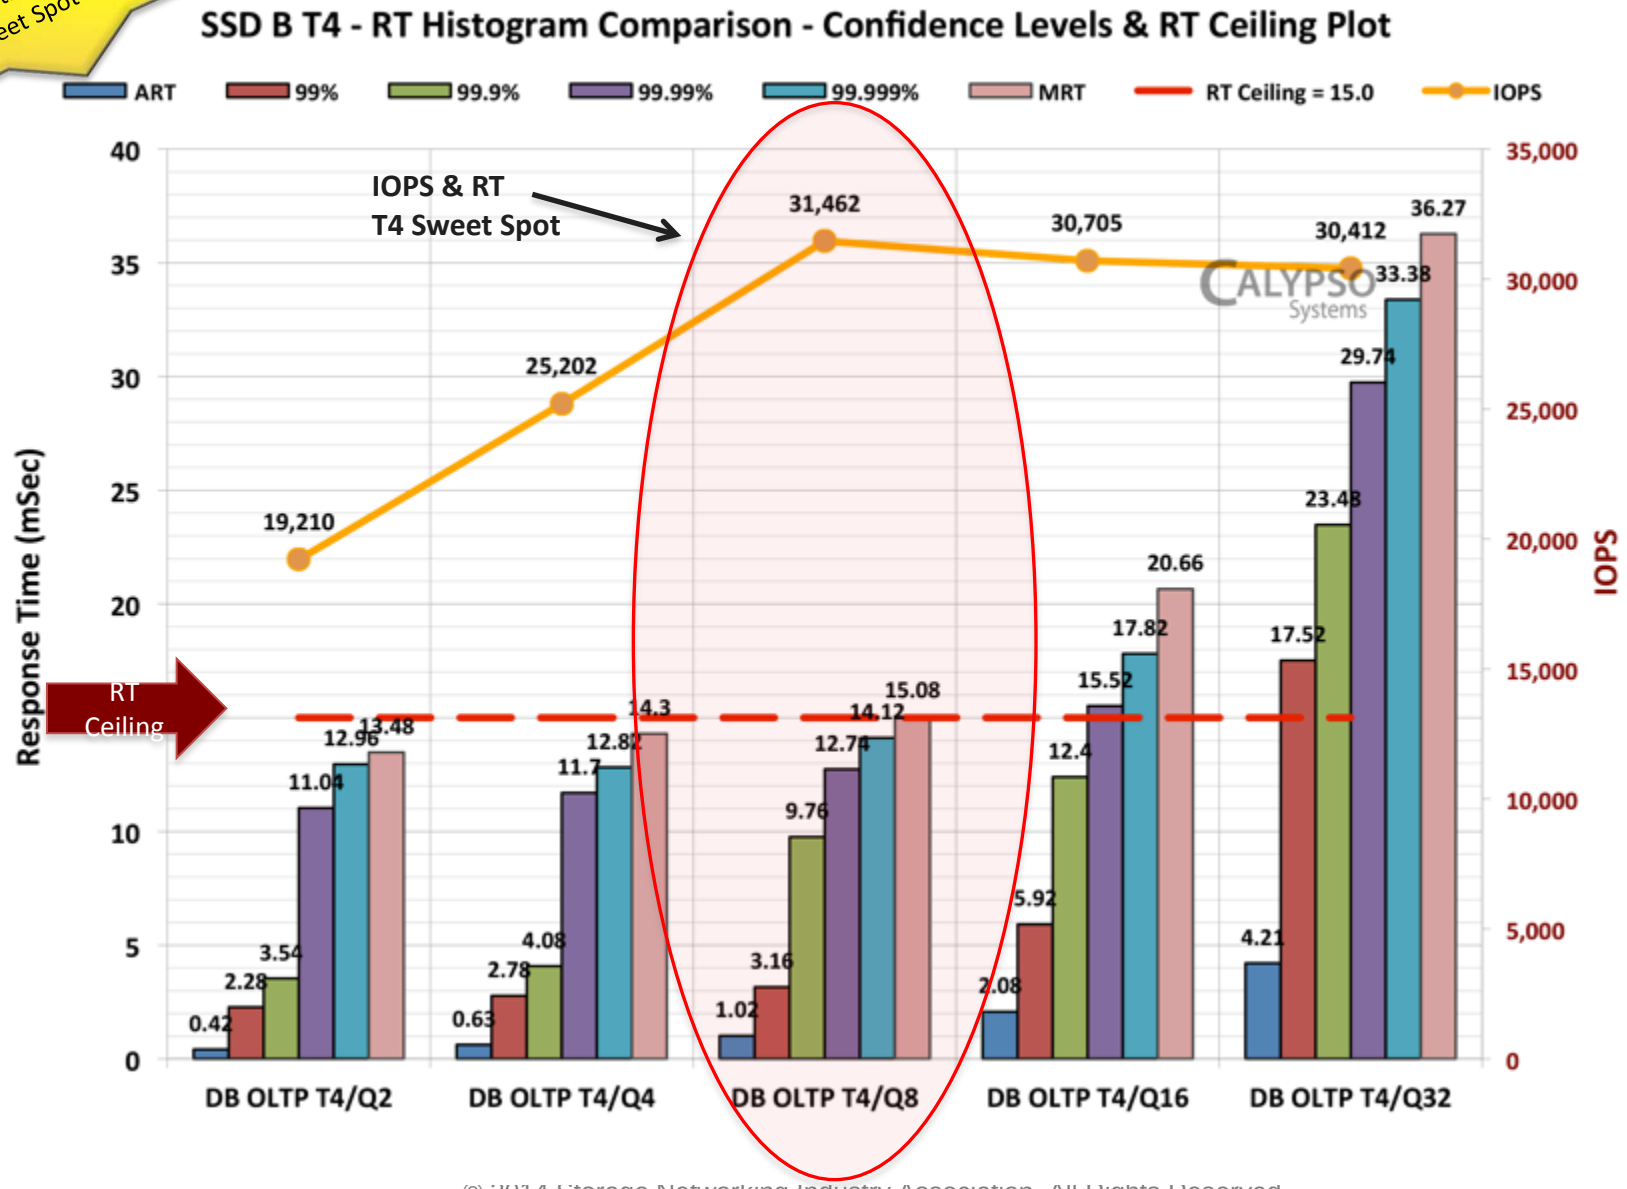

TAKE AWAY
SSD B at T8 verifies Saturation. Note T8 Sweet Spot at T8Q4.

SSD B T8: QD Histogram Compare

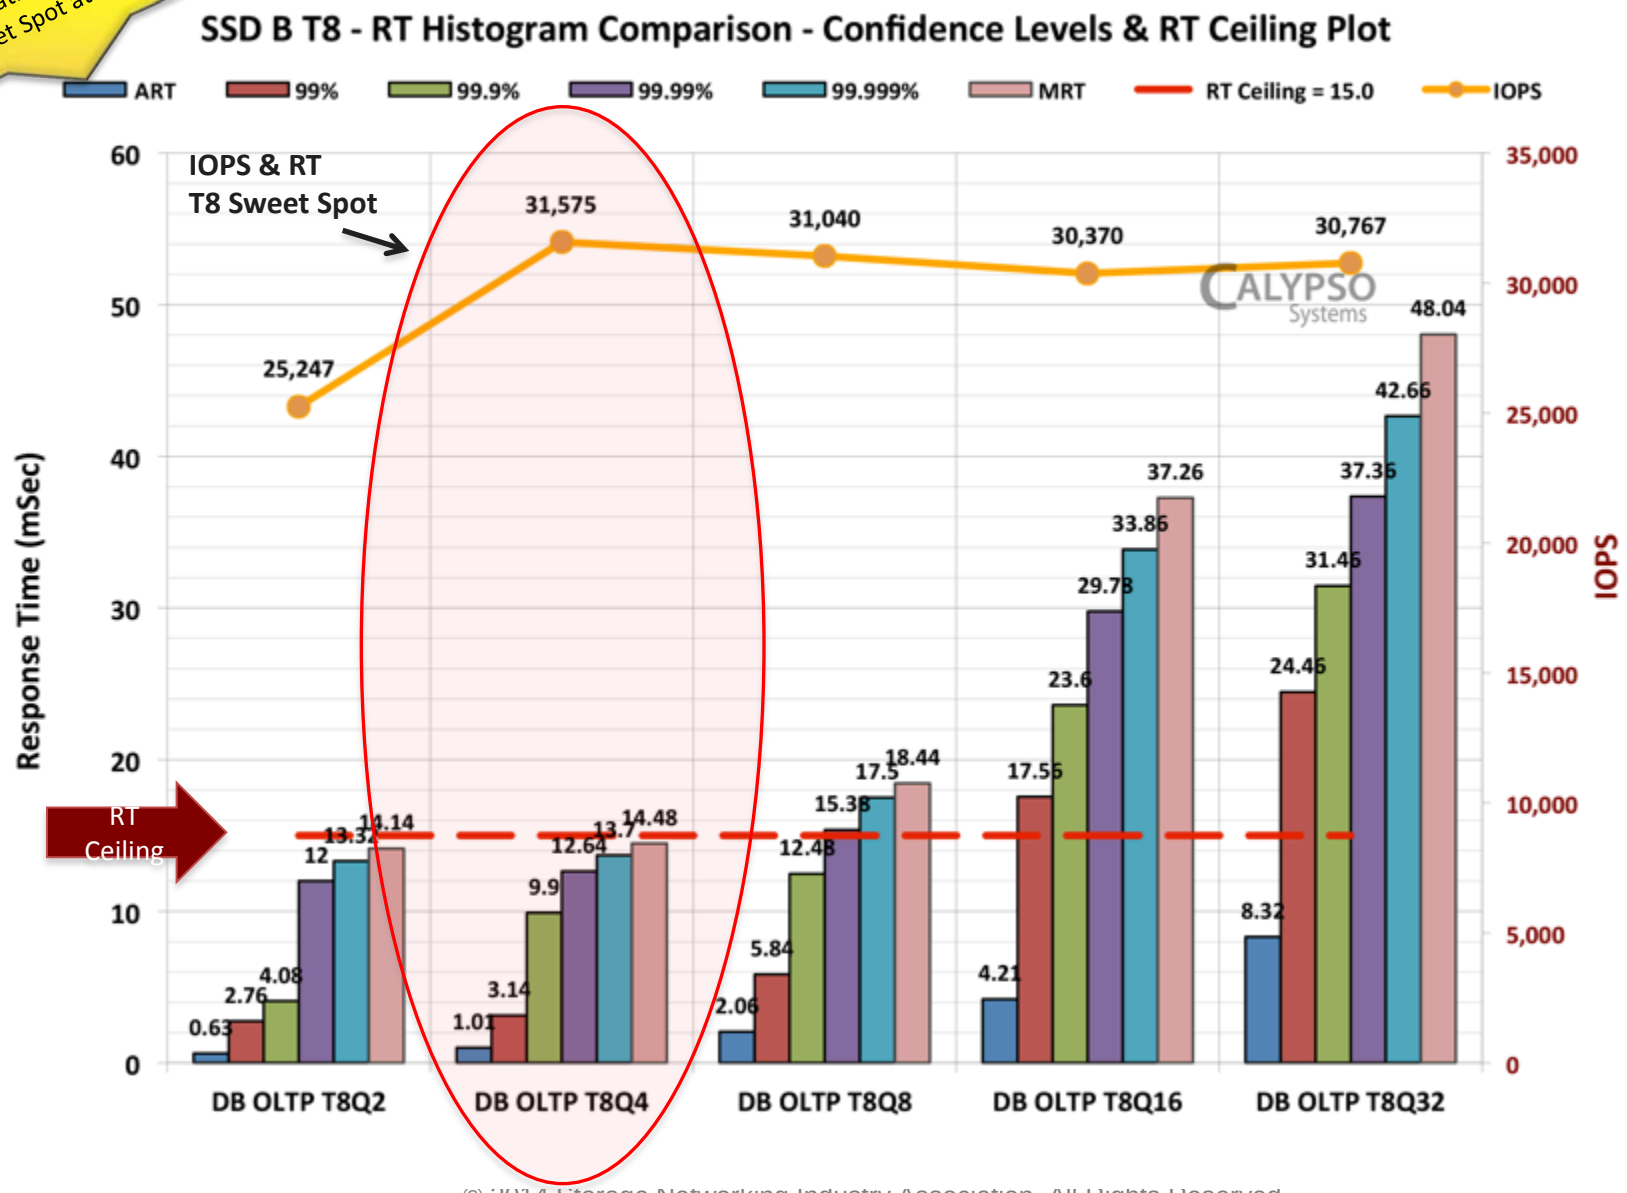

Findings

What are the optimal operating point(s)?

TAKE AWAY
SSD A 800 GB shows good performance at OIO 16, 31, & 64.

SSD A: Histogram Compare T4 T8 T16

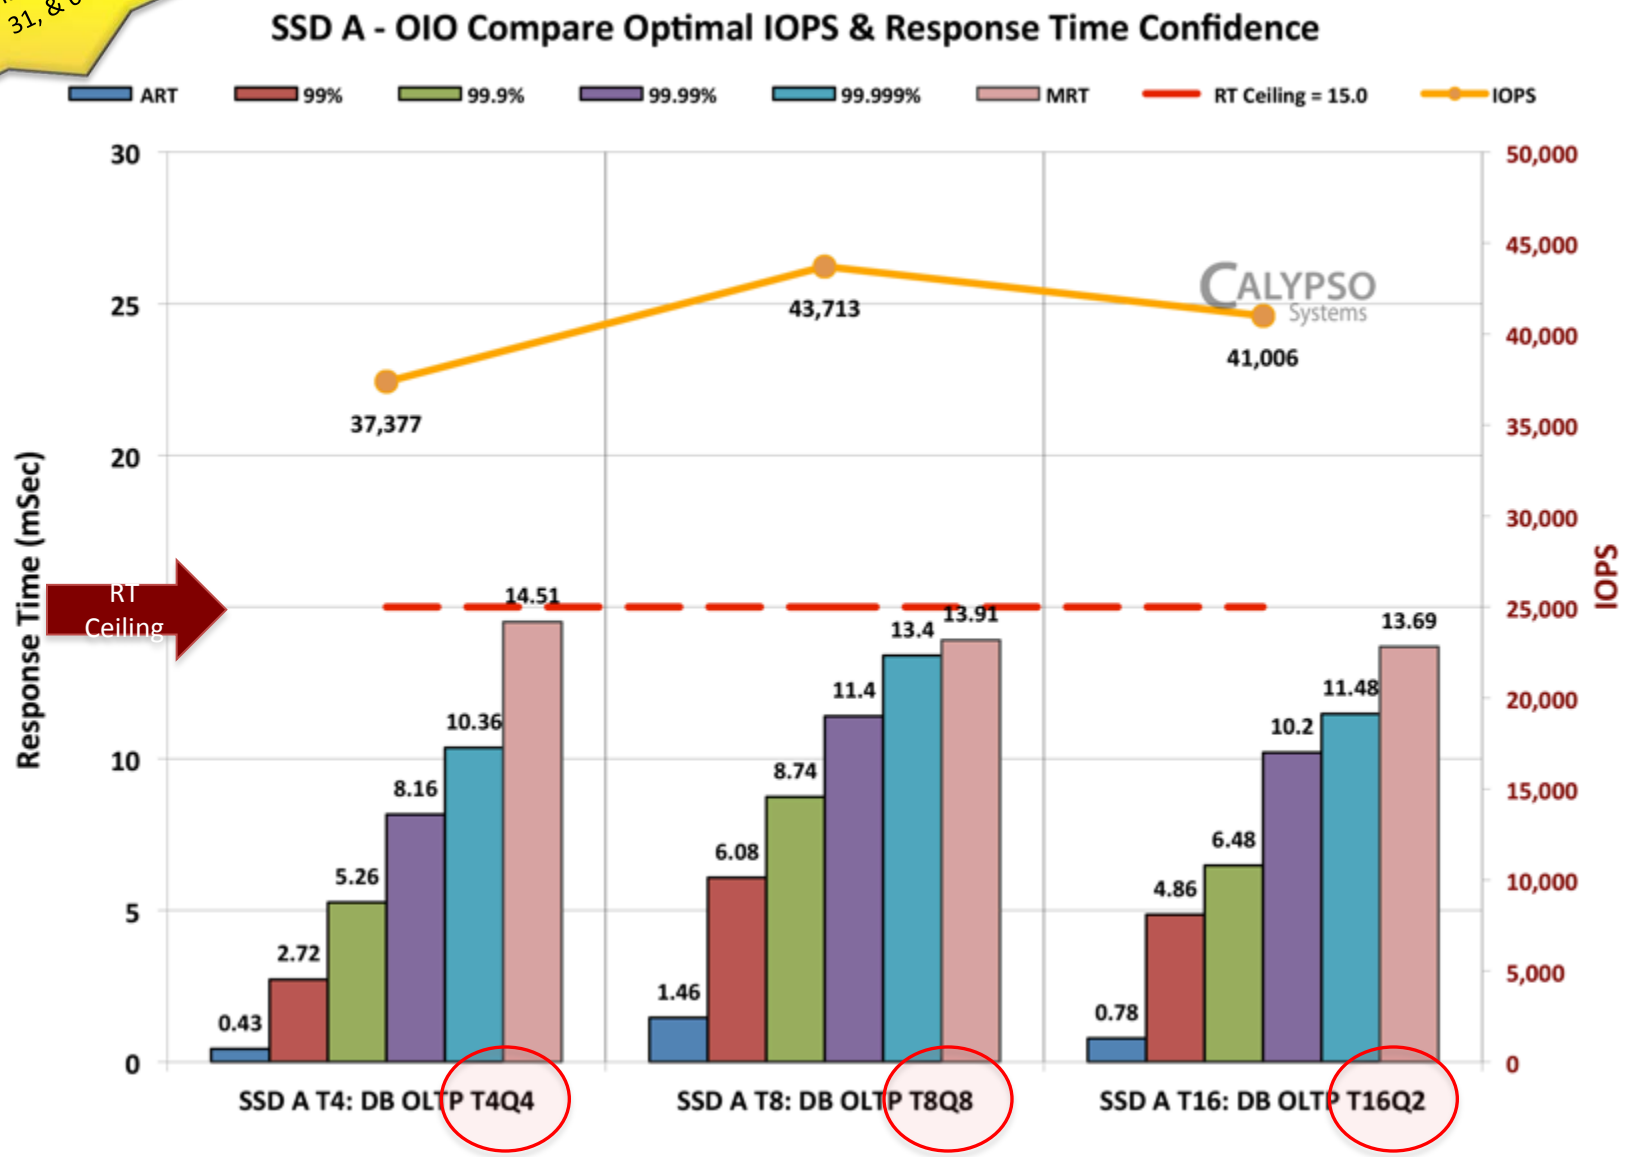

TAKE AWAY
SSD B 200 GB optimal OIO at any combination of OIO=32: T2Q16, T4Q8 or T8Q4.

SSD B: Histogram Compare T2 T4 T8

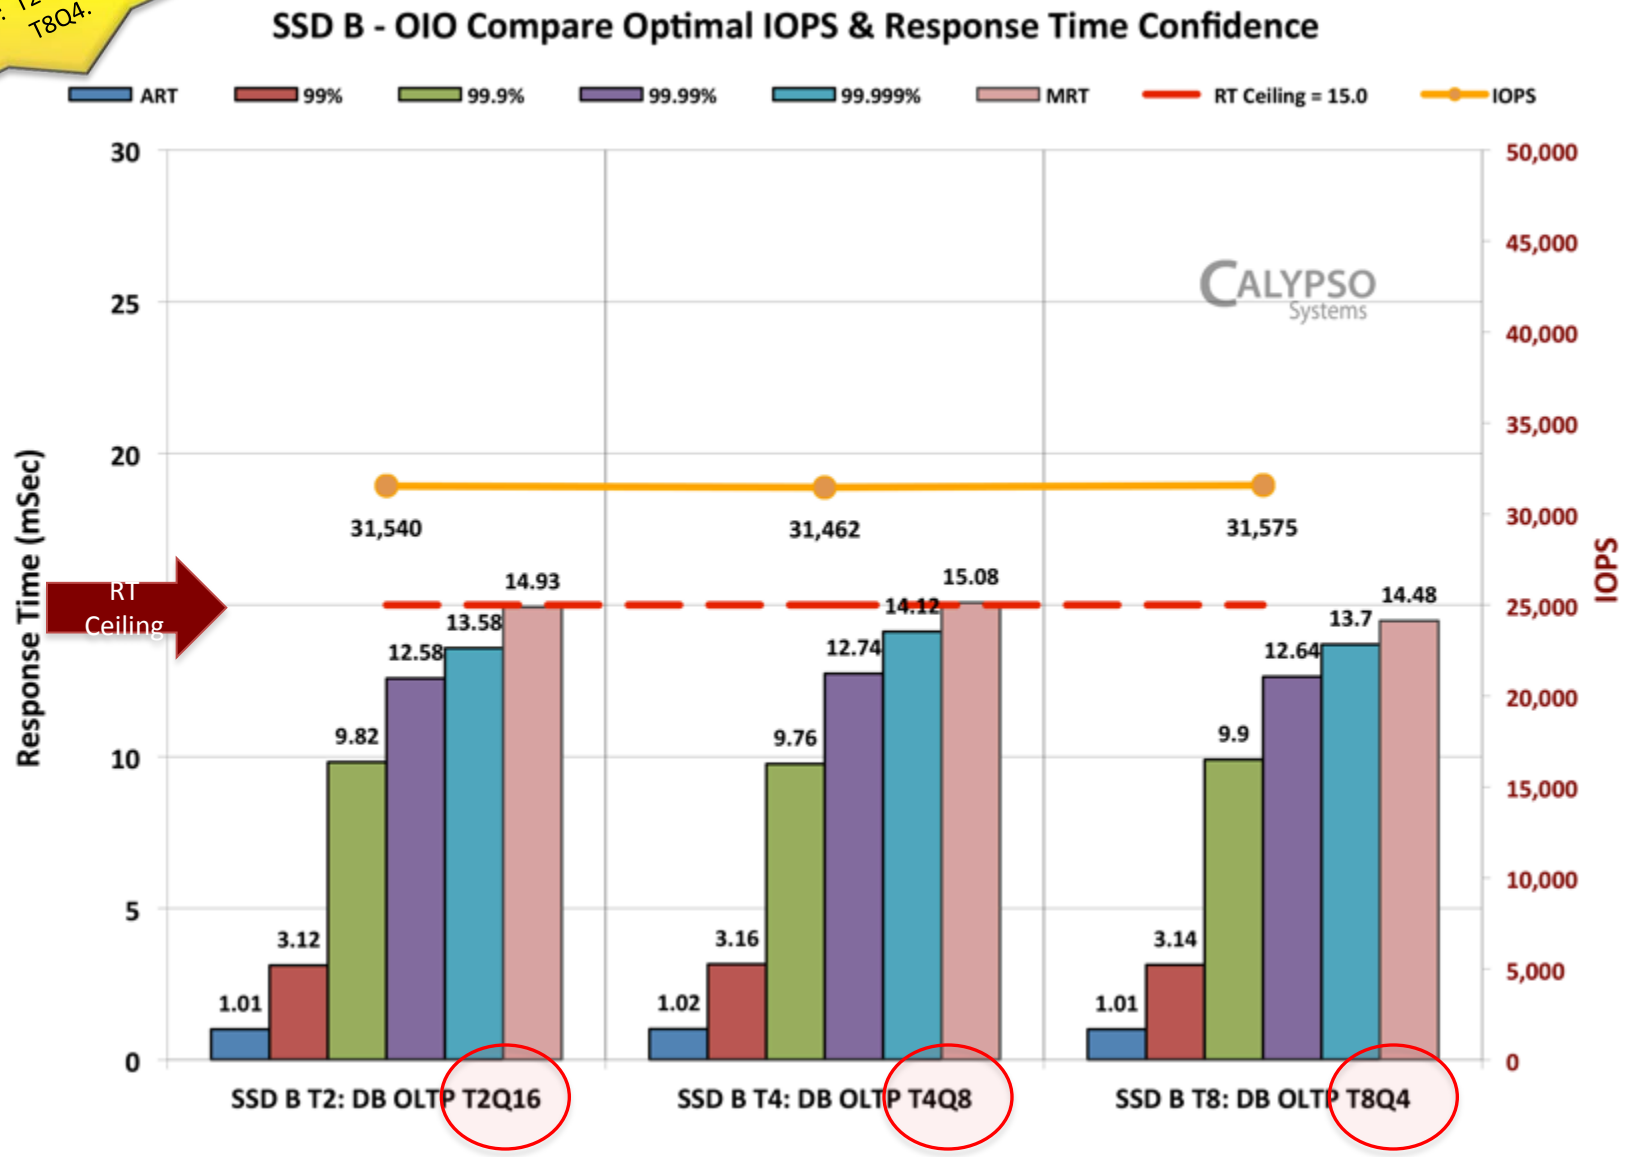

Conclusion

- Application(s) may require low or high OIO
- Drive Design may emphasize low or high OIO
- Drives may be designed for specific (different) workloads

Understanding your SSD's deterministic behavior helps SSD & system design optimization.

Know the "cost" of more IOPS in terms of Response Times.

For more information, contact Calypso Systems, Inc.

info@calypsotesters.com www.calypsotesters.com

Attribution & Feedback

The SNIA Education Committee thanks the following individuals for their contributions to this Tutorial.

Authorship History

Eden Kim, June 30, 2014

Updates:

- **Name/Date**
- **Name/Date**
- **Name/Date**

Additional Contributors

Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here
Name of contributor here

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***