



# PCIe SSD Edge Card Blind Survey

Session 101-A

## Blind Comparison of PCIe SSD Performance

Eden Kim, CEO

Calypso Systems, Inc.



# Emergence of PCIe SSD

- PCIe SSDs – have higher performance & faster response times
- Gen 3 PCIe emerging for Client & Enterprise environments
- Enterprise Server SSDs - edge card & 2.5” form factor (SFF 8639)
- Client Module SSDs – M.2, mSATA, MO-297





# Cost of Higher PCIe Performance?

- Increased Latencies
- Higher CPU usage
- Higher Power Consumption
- More NAND Flash Dies
- Higher Price

Blind Survey 2013 focuses on Response Times & CPU Usage

Testing based on SNIA PTS-E 1.1, RTP 3.0 & CTS 7.0

# Blind Comparison of Performance



Calypso PCIe SSD Card Blind Survey  
5 leading PCIe SSD Edge Cards

Testing to SNIA PTS-E 1.1 and CTS Tests  
Corner Case Synthetic Tests  
Application Workload Tests

Test Platform: Calypso RTP 3.0  
Test Software: Calypso CTS 7.0

Test Run Dates: June 1 – August 31, 2013

# Sample Pool

Enterprise Class PCIe SSD Edge Cards		
PCIe Card	Capacity	Note
Card 1	602 GB	Performance Mode
Card 2	350 GB	Performance Mode
Card 3	1,847 GB	Performance Mode
Card 4	1,600 GB	Performance Mode
Card 5	200 GB	Single Mode
Card 6	800 GB	RAID x4

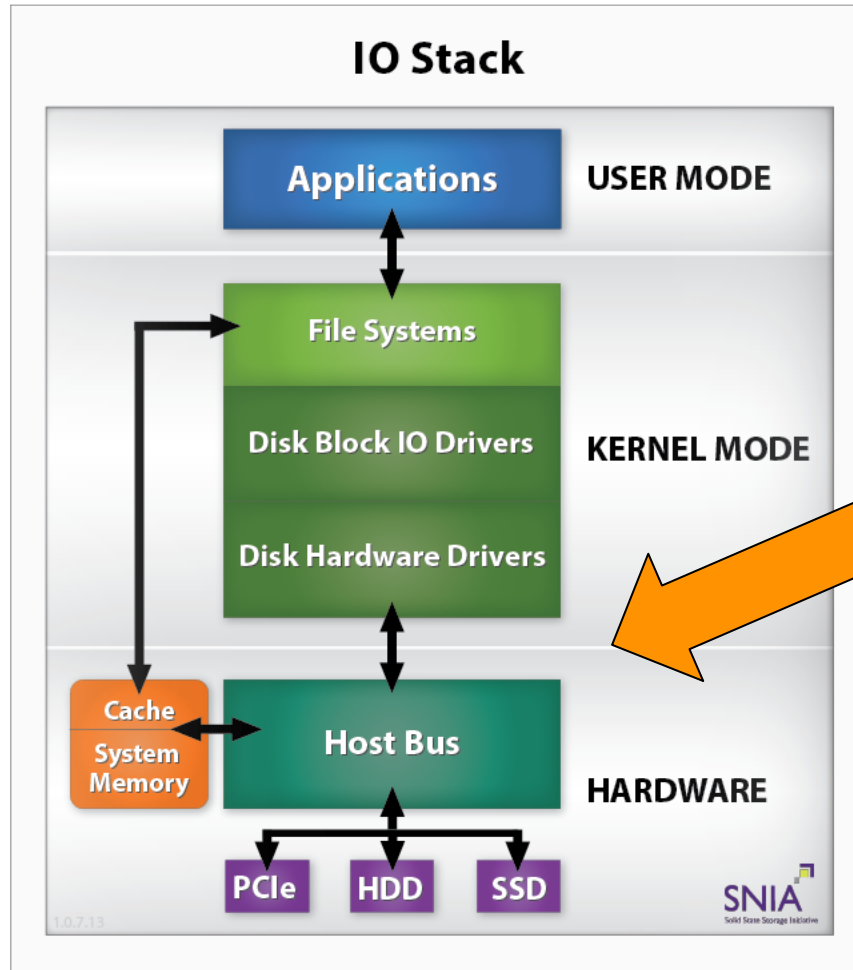
# Test Plan & Report

Blind Survey Test Plan			
Test Phase	Description	Test	Note
<b>Phase I</b>	Synthetic Corner Case	WSAT	Monotonic PC
<b>Phase II</b>	SNIA PTS – Basic	Steady State	PTS-E 1.0
<b>Phase III</b>	SNIA PTS - Advanced	Steady State	PTS-E 1.1
<b>Phase IV</b>	Workload Based Synthetic	Application Specific	Complex PC
Blind Survey Report			
Release Date	Report Type	Data Sets	Note
August 2013	Summary Data Sets	Select	Summary from each Phase – includes SNIA IOPS Report
September 2013	Full Report	Full Comparison	Over 250 pages of detailed individual and comparative reports

# Survey Tests

Blind Survey – Test List			
Test Phase	Description	Test	Note
<b>Phase I</b>	WSAT - RND	4KiB W; 8KiB 65:35	Evolution over Time
	WSAT - SEQ	4KiB W; 128KiB W	Evolution over Time
<b>Phase II</b>	SNIA PTS – Basic	Steady State	PTS-E 1.0
	Steady State	IOPS, TP, LAT	SNIA WDPC/WIPC
<b>Phase III</b>	SNIA PTS - Advanced	Steady State	PTS-E 1.1
	Demand Intensity Response Time Histograms	RND 128KiB 90:10 RND 8K 65:35	SNIA WDPC SNIA Hot Spots
<b>Phase IV</b>	Select Workload Based	Application Specific	Complex PC
	Modified WSAT w/ complex Pre-Conditioning	TBD from: Webservers, Exchange Mail, Media Streaming, File Servers, database OLTP, OS paging, VOD, webserver logs, SQL logs, archiving, medial imaging and more	

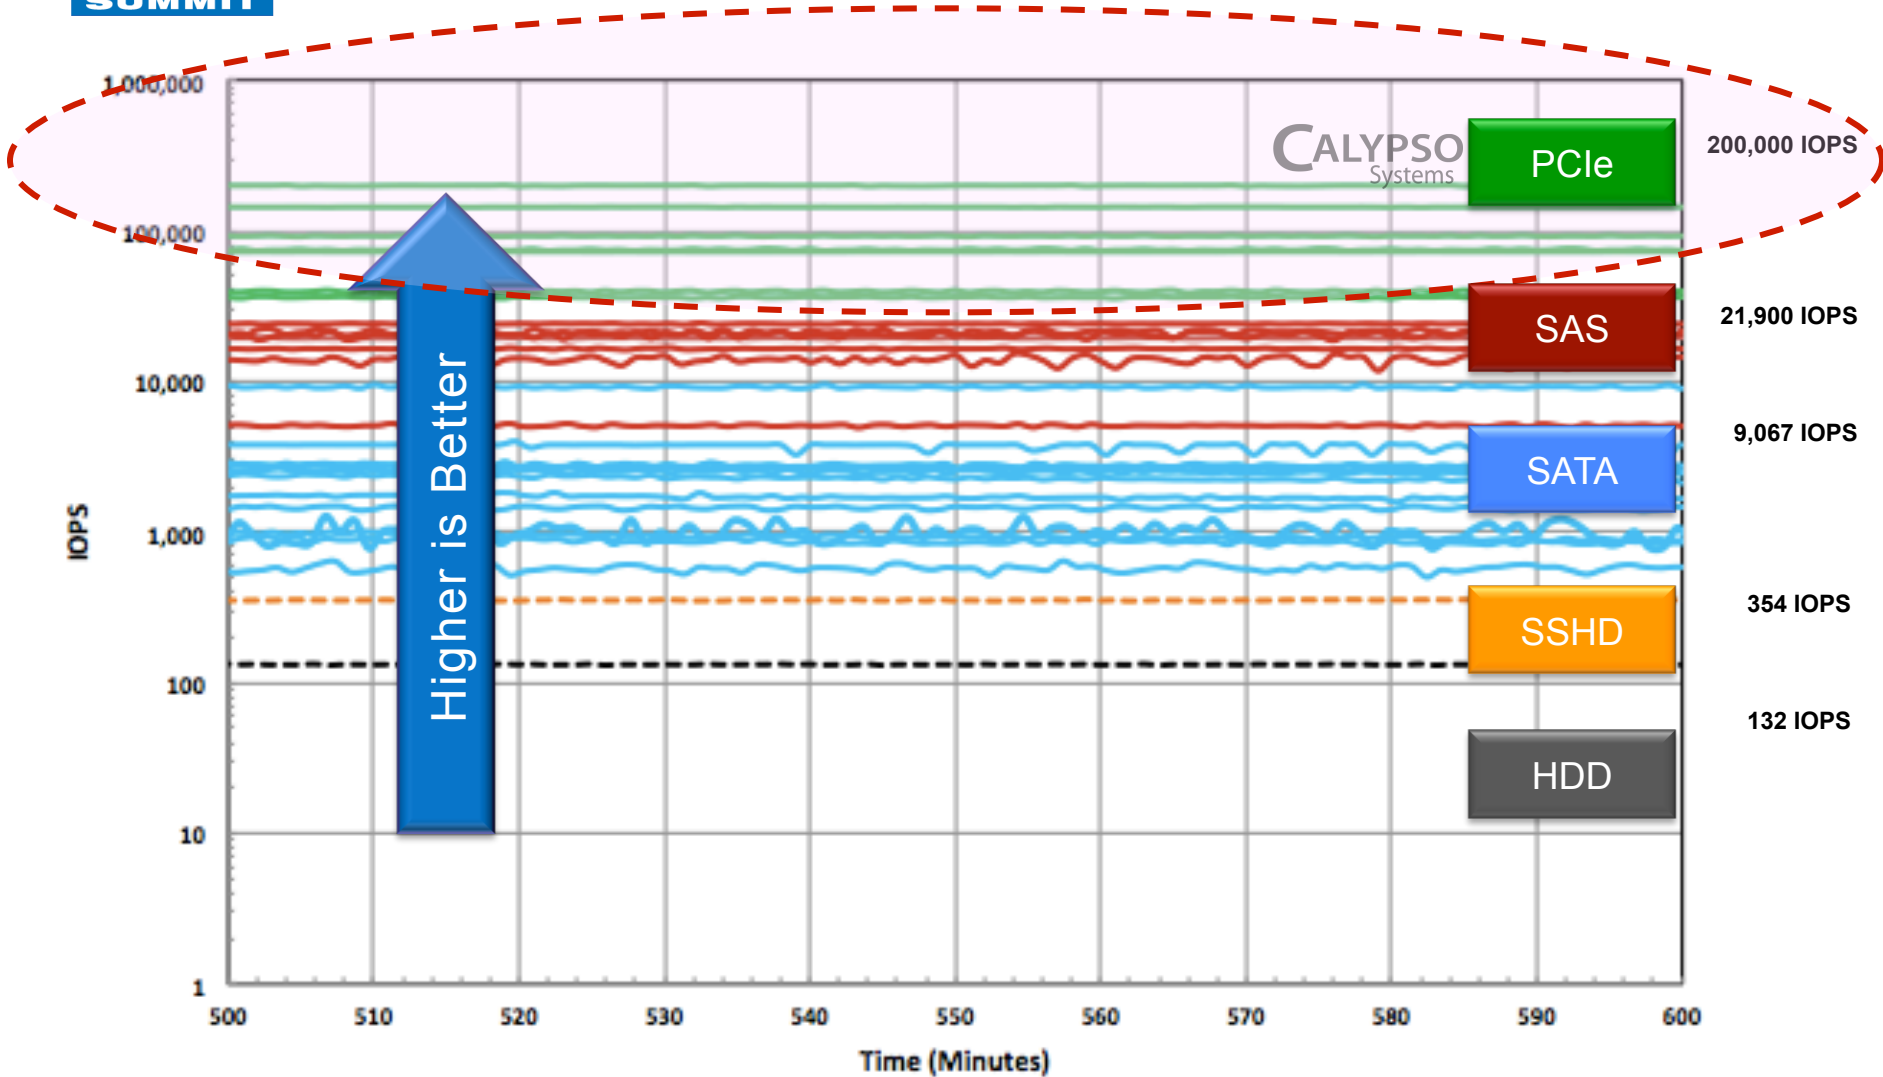
# Device Level Testing – Block IO



**CTS Testing  
Device level  
Steady State**



# WSAT RND 4KiB W – By Storage Class



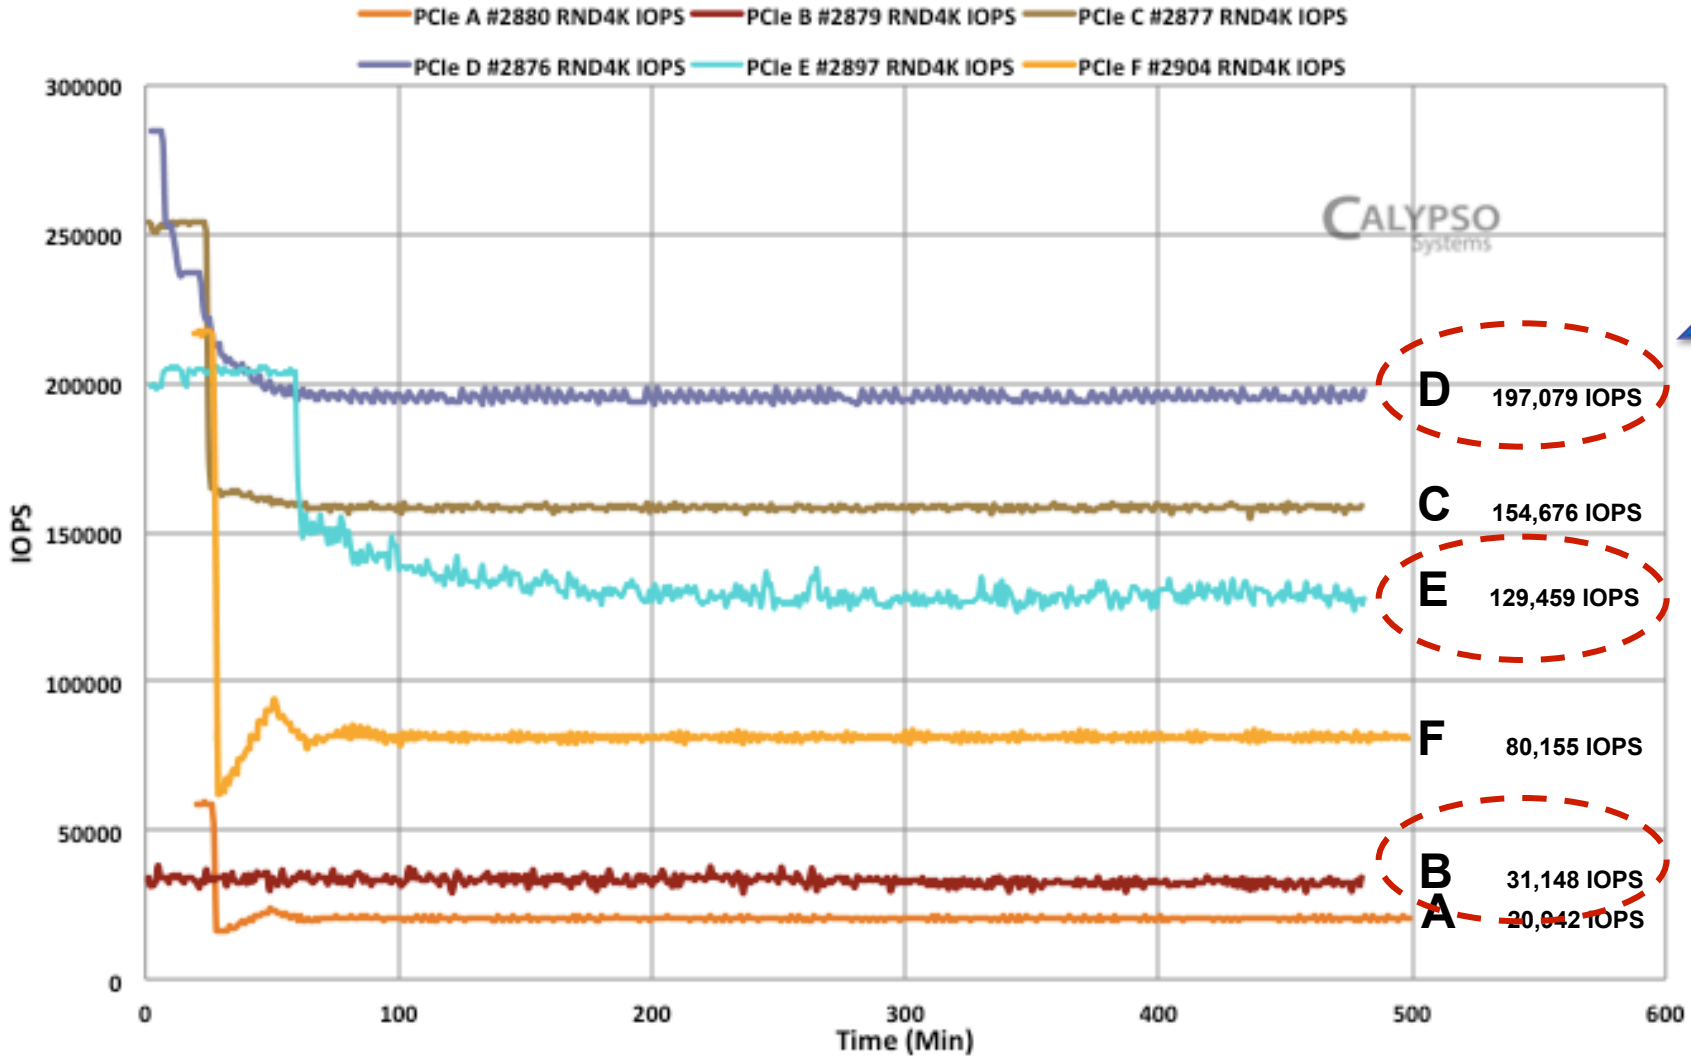


# Blind Survey 2013 – Preliminary Data


## Phase I: WSAT IOPS

WSAT RND 4KiB W  
WSAT SEQ 128KiB W

# Phase I: WSAT RND 4KiB W – IOPS v Time



# Phase I: WSAT SEQ 128KiB W – Bandwidth



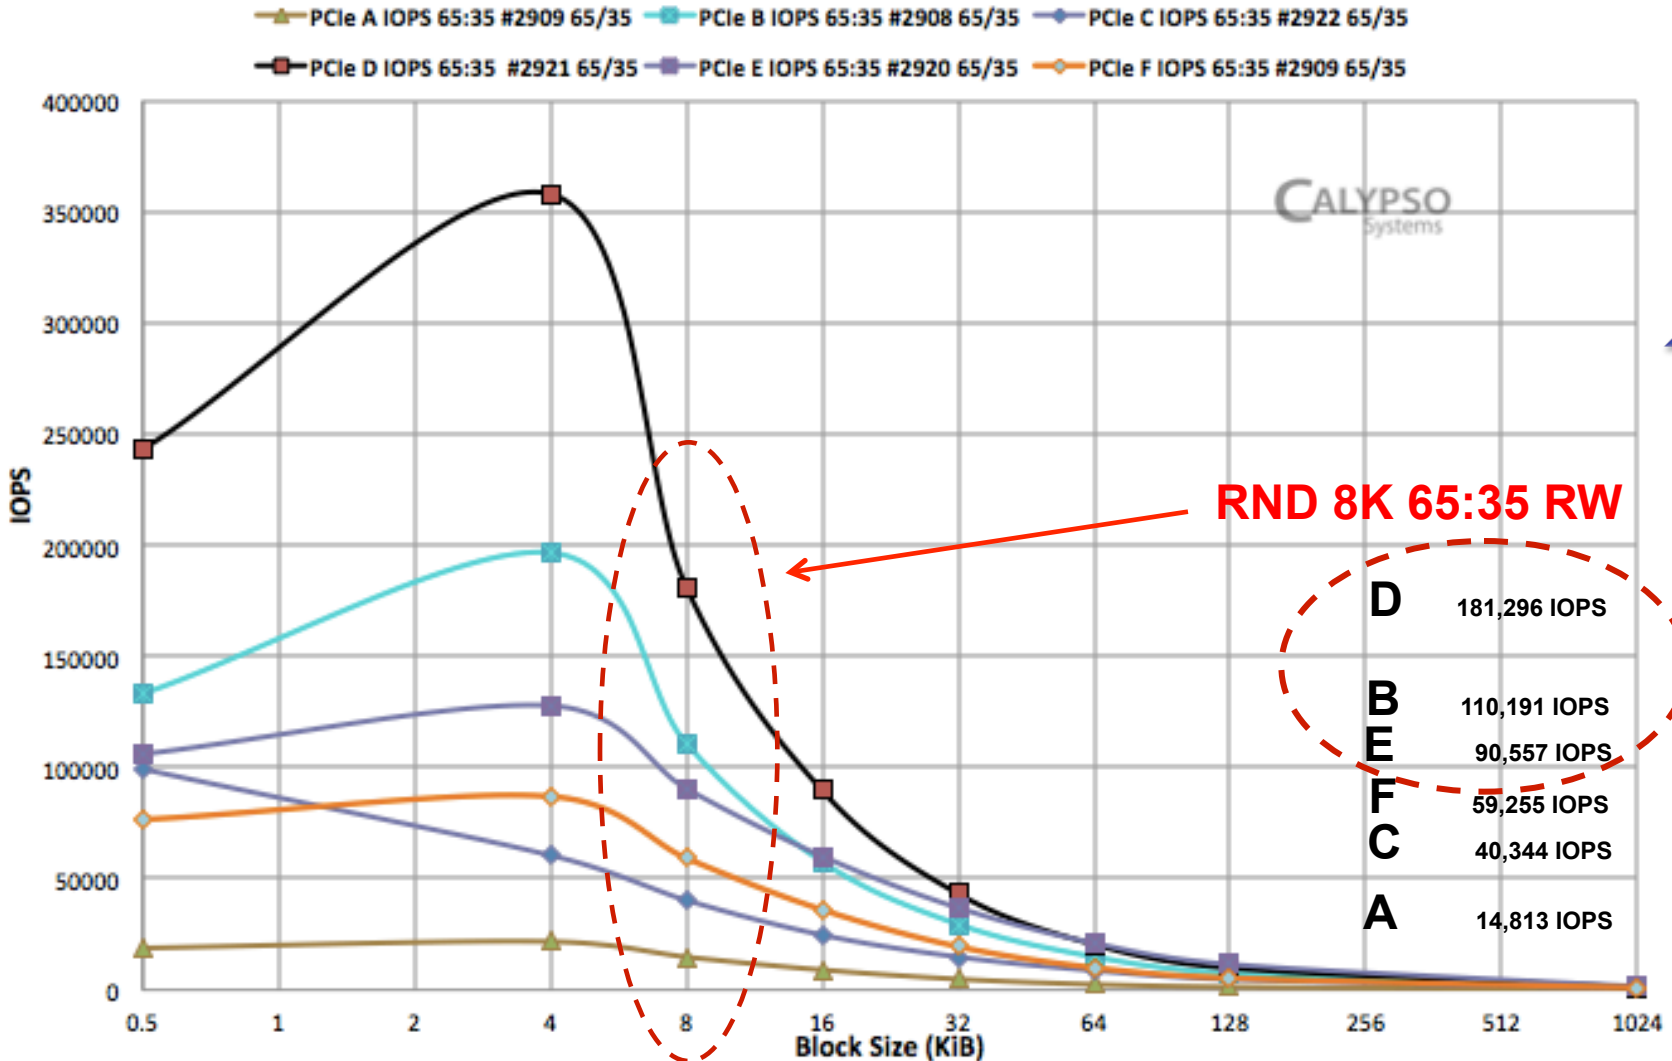


# Blind Survey 2013 – Preliminary Data

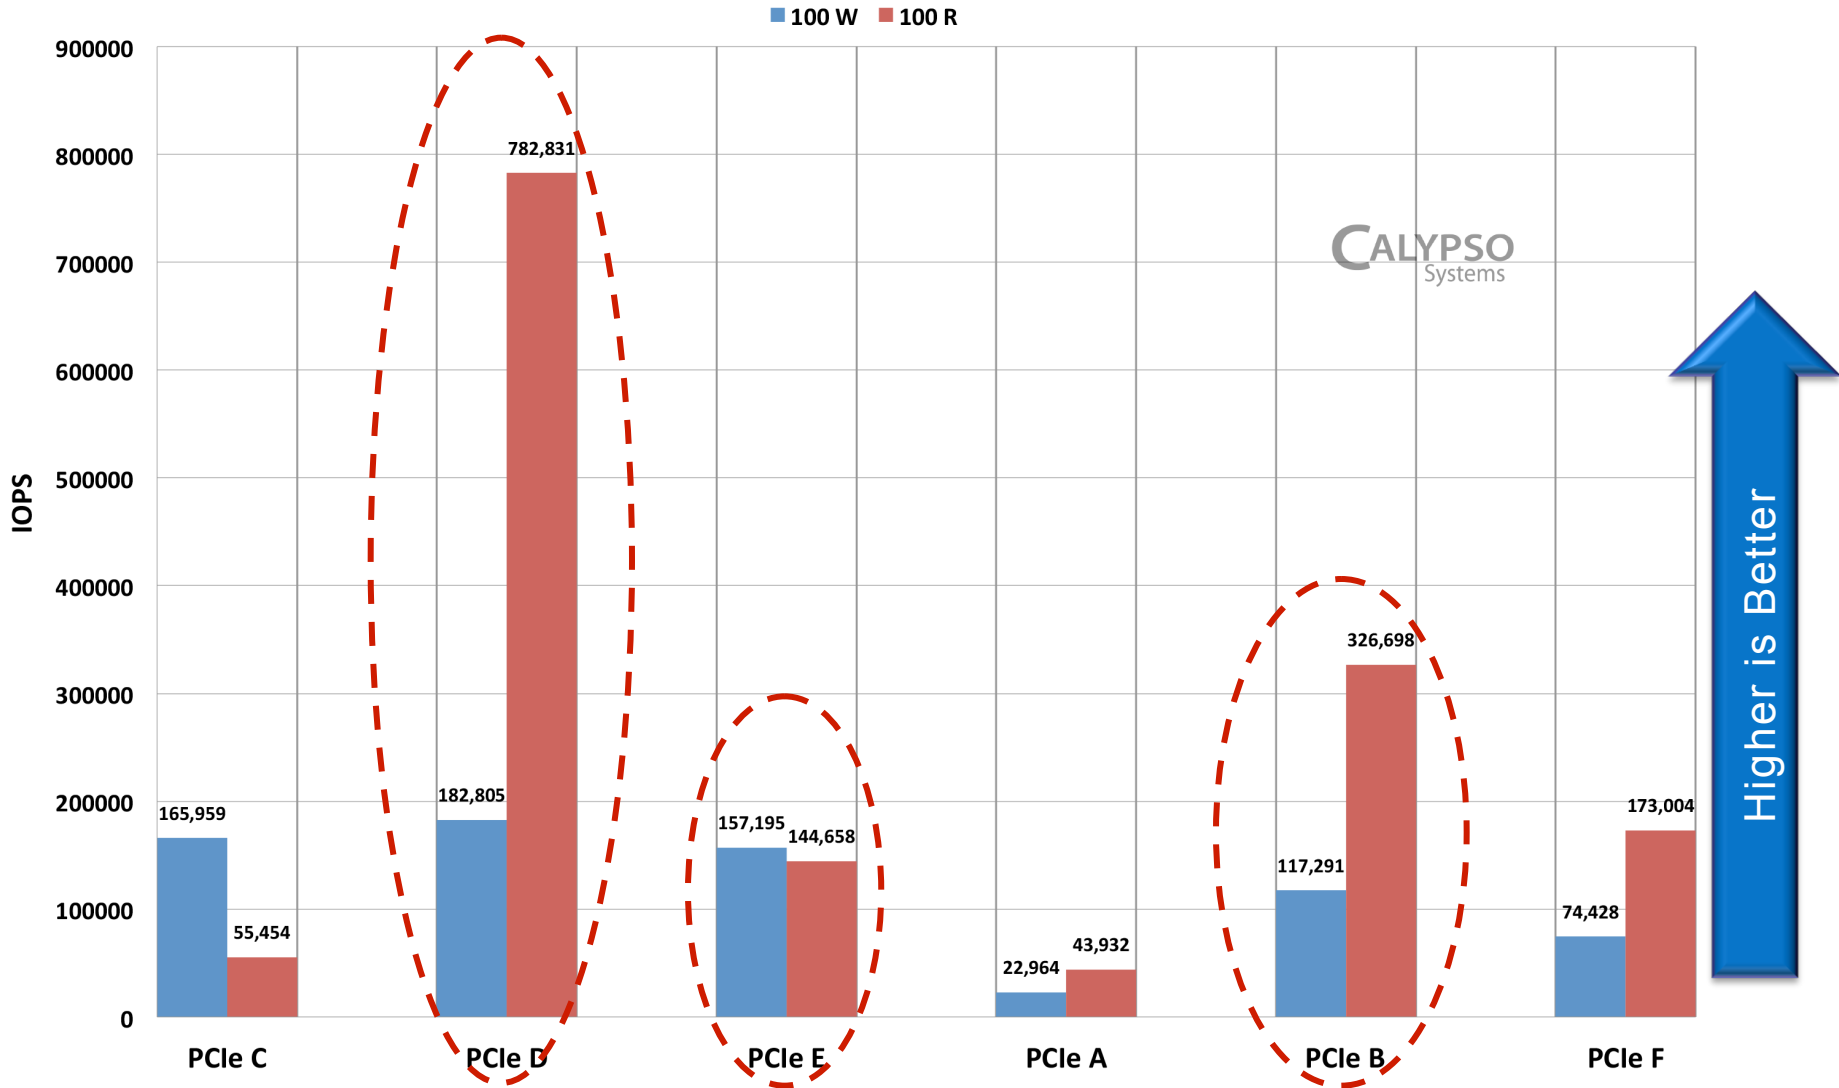
## Phase II: Steady State IOPS

SS IOPS Block Size x 65:35 RW  
SS IOPS RND 4KiB R & W

# Phase II: Steady State RND IOPS – 65:35 RW mix



# Phase II: Steady State IOPS – RND 4K R/W



# Blind Survey 2013 – Preliminary Data

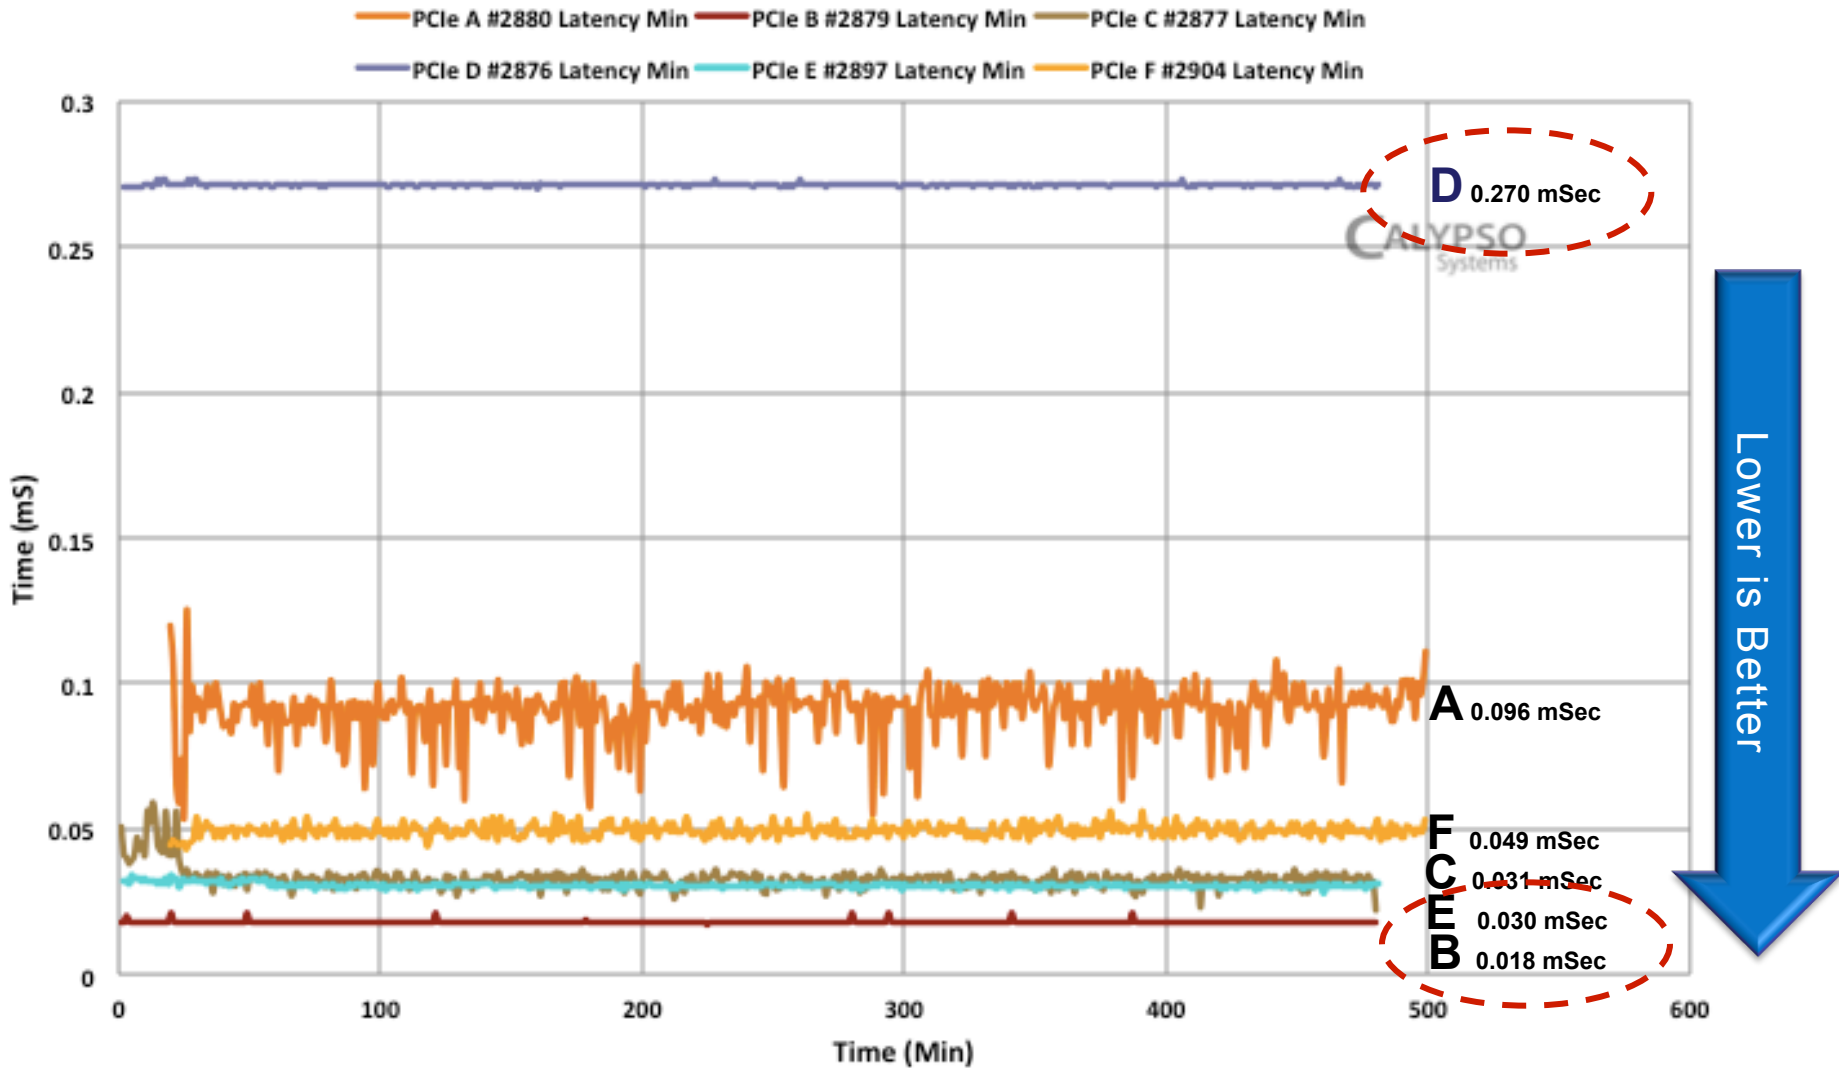
## Phase I: Response Times

WSAT RND 4KiB W  
WSAT SEQ 128KiB W

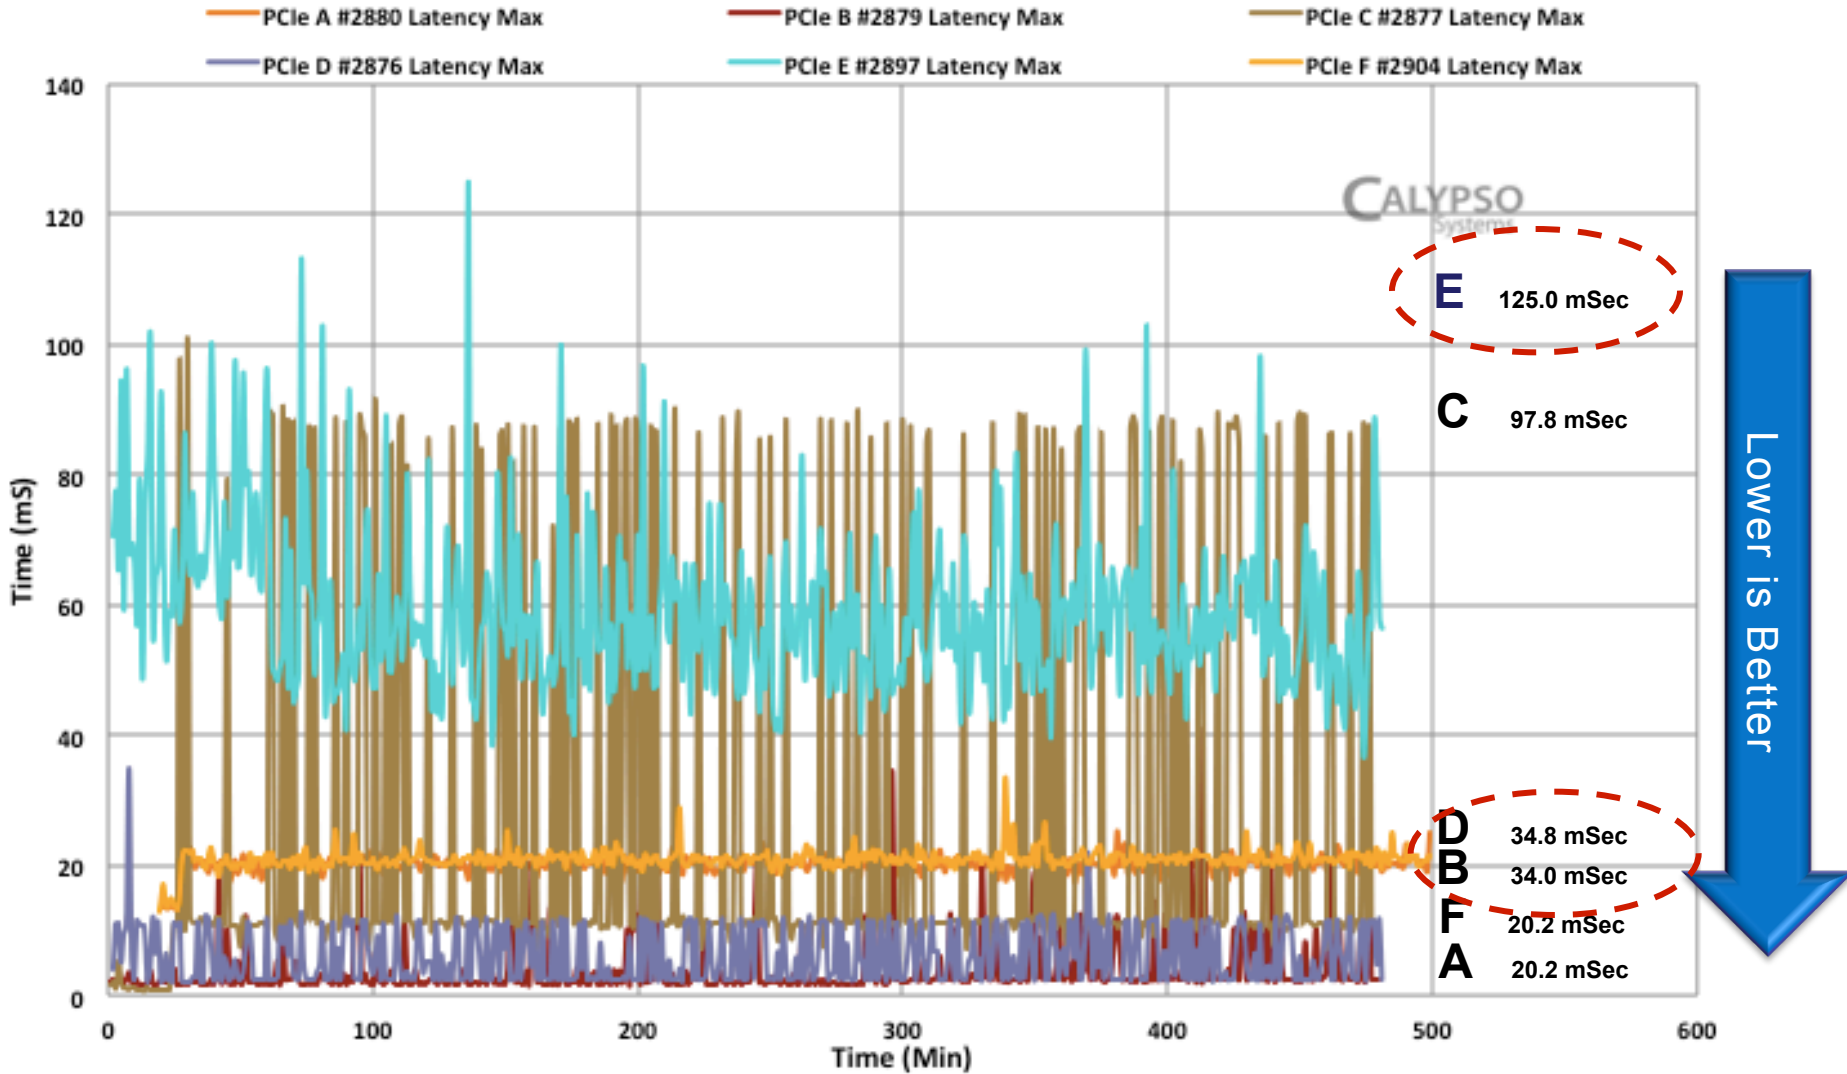
Response Time Histogram  
Histogram Confidence Plot



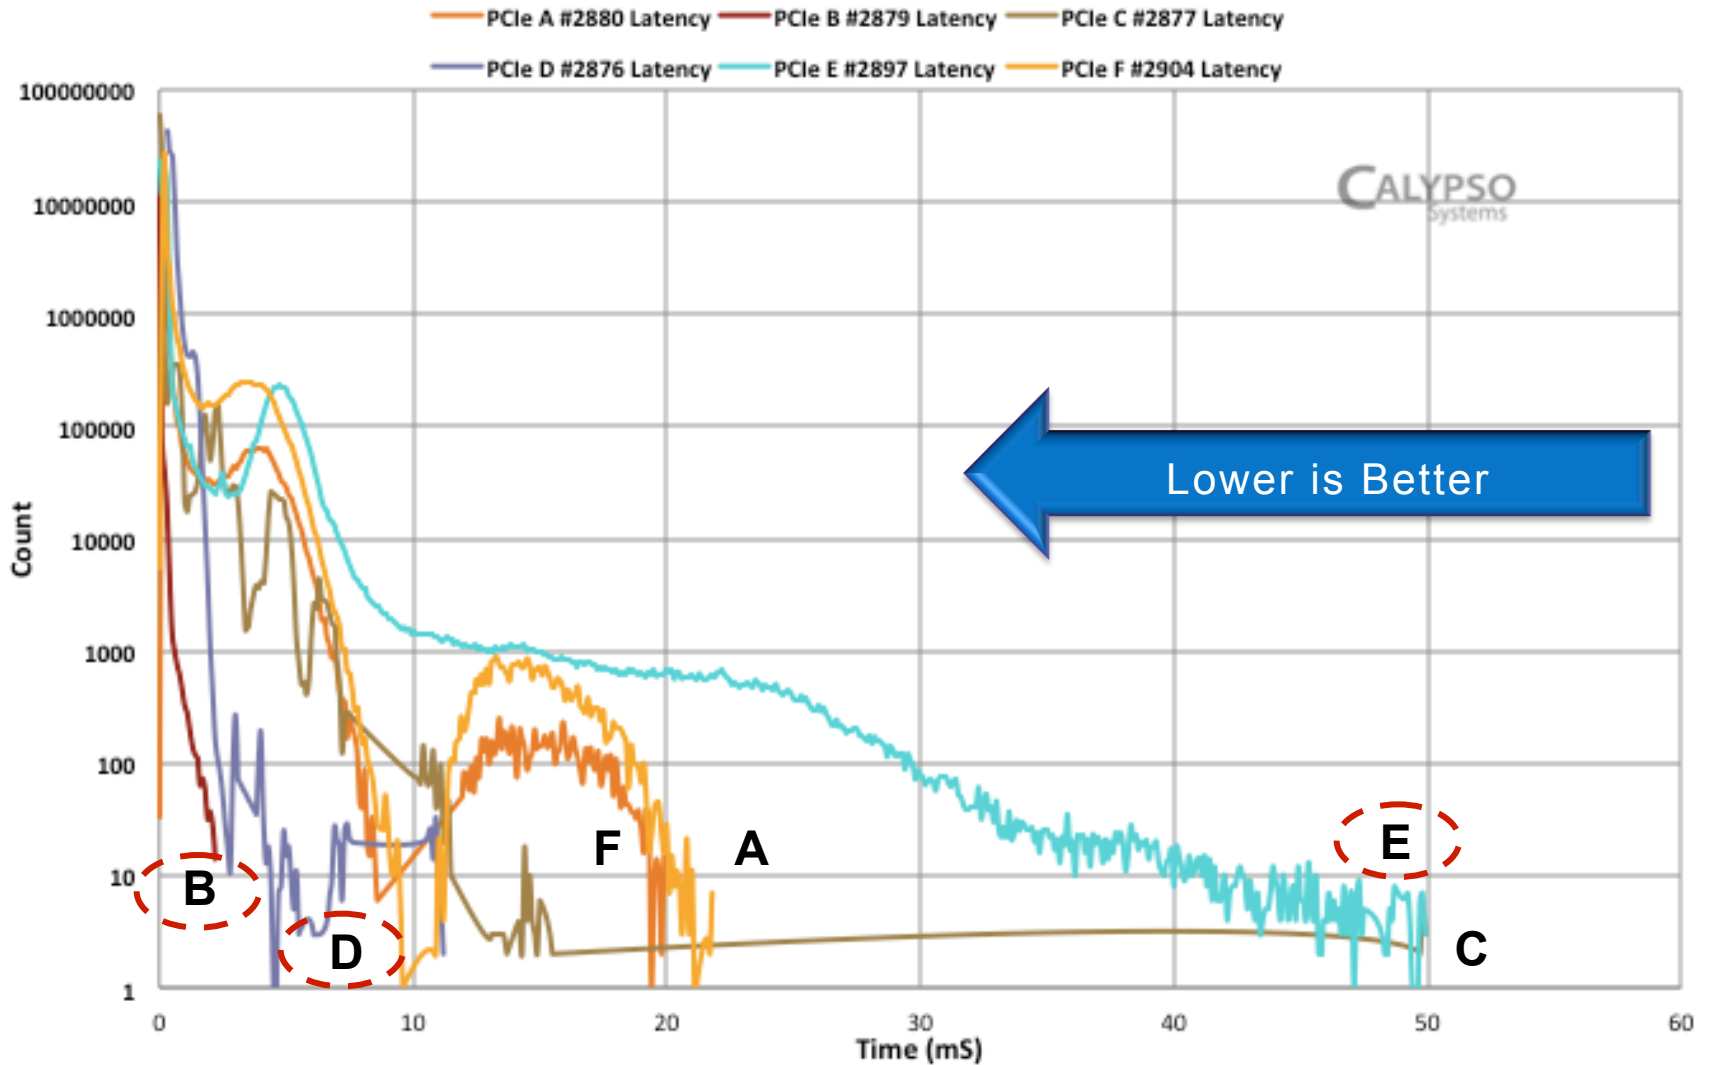
# WSAT RND 4KiB W – Min Response Times



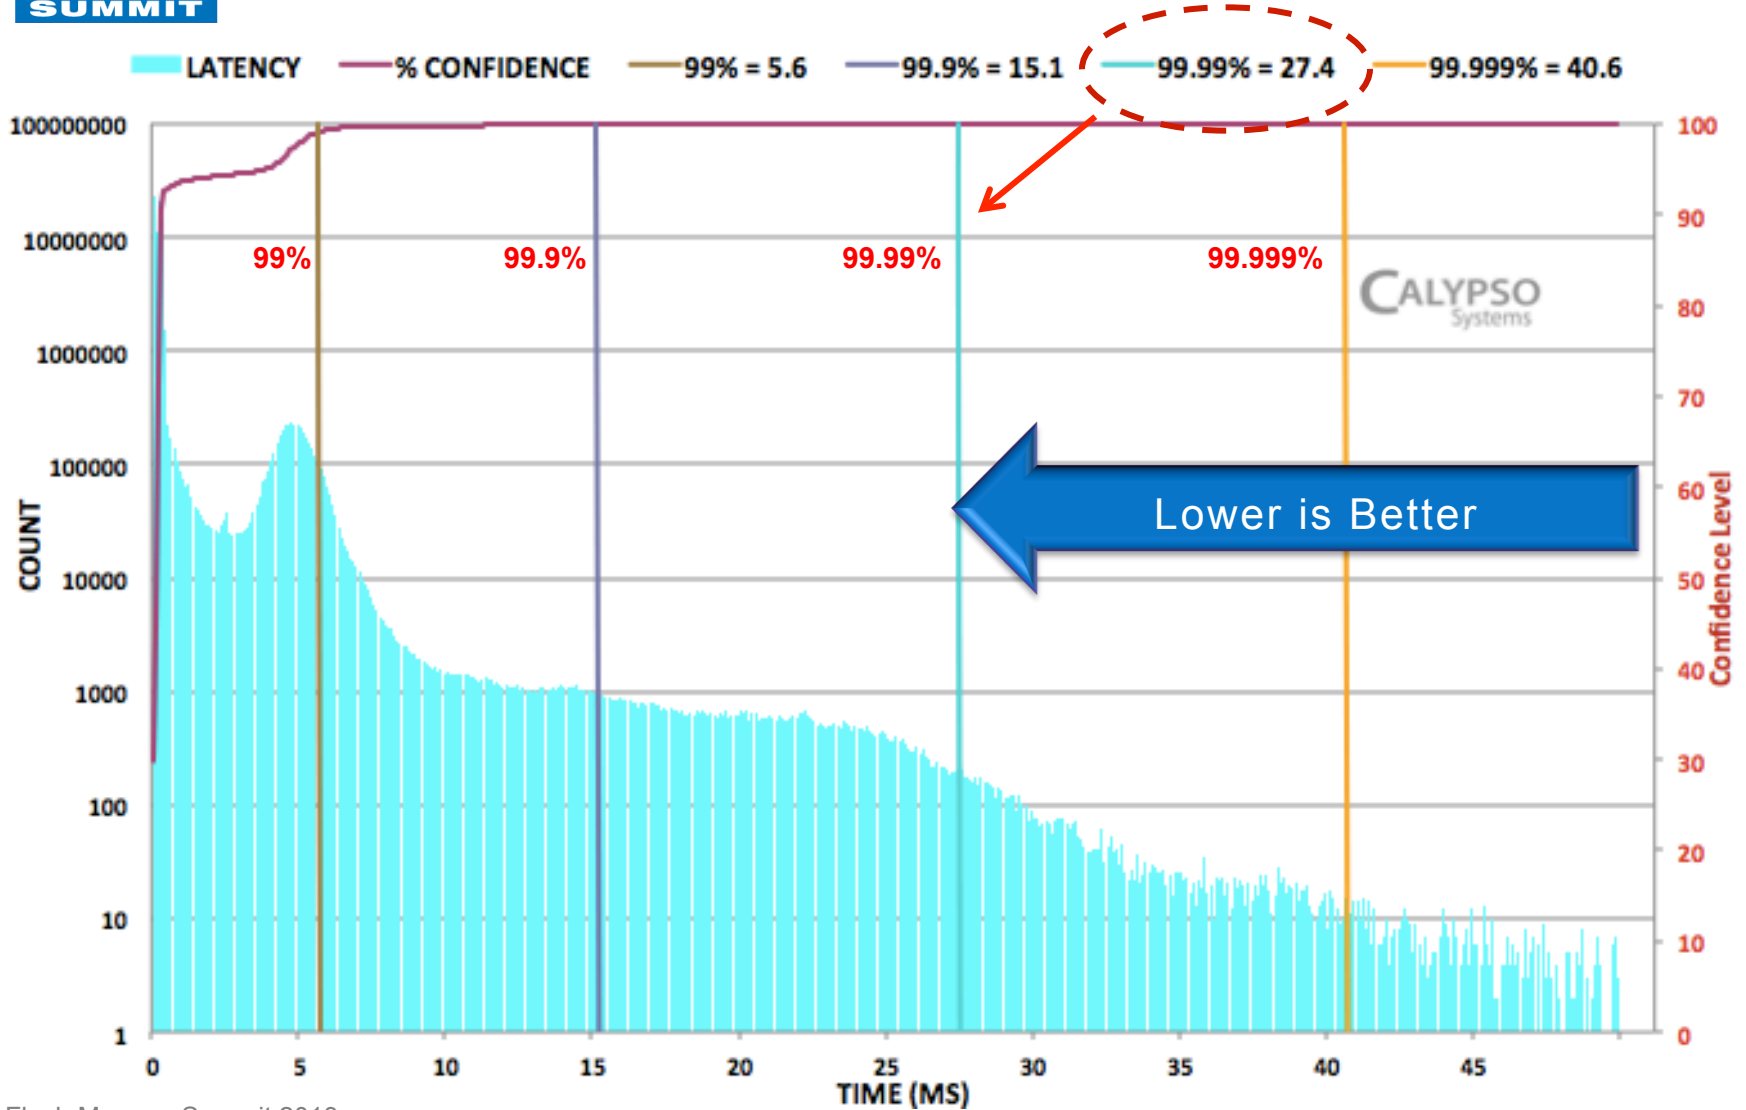
# WSAT RND 4KiB W – Max Response Times



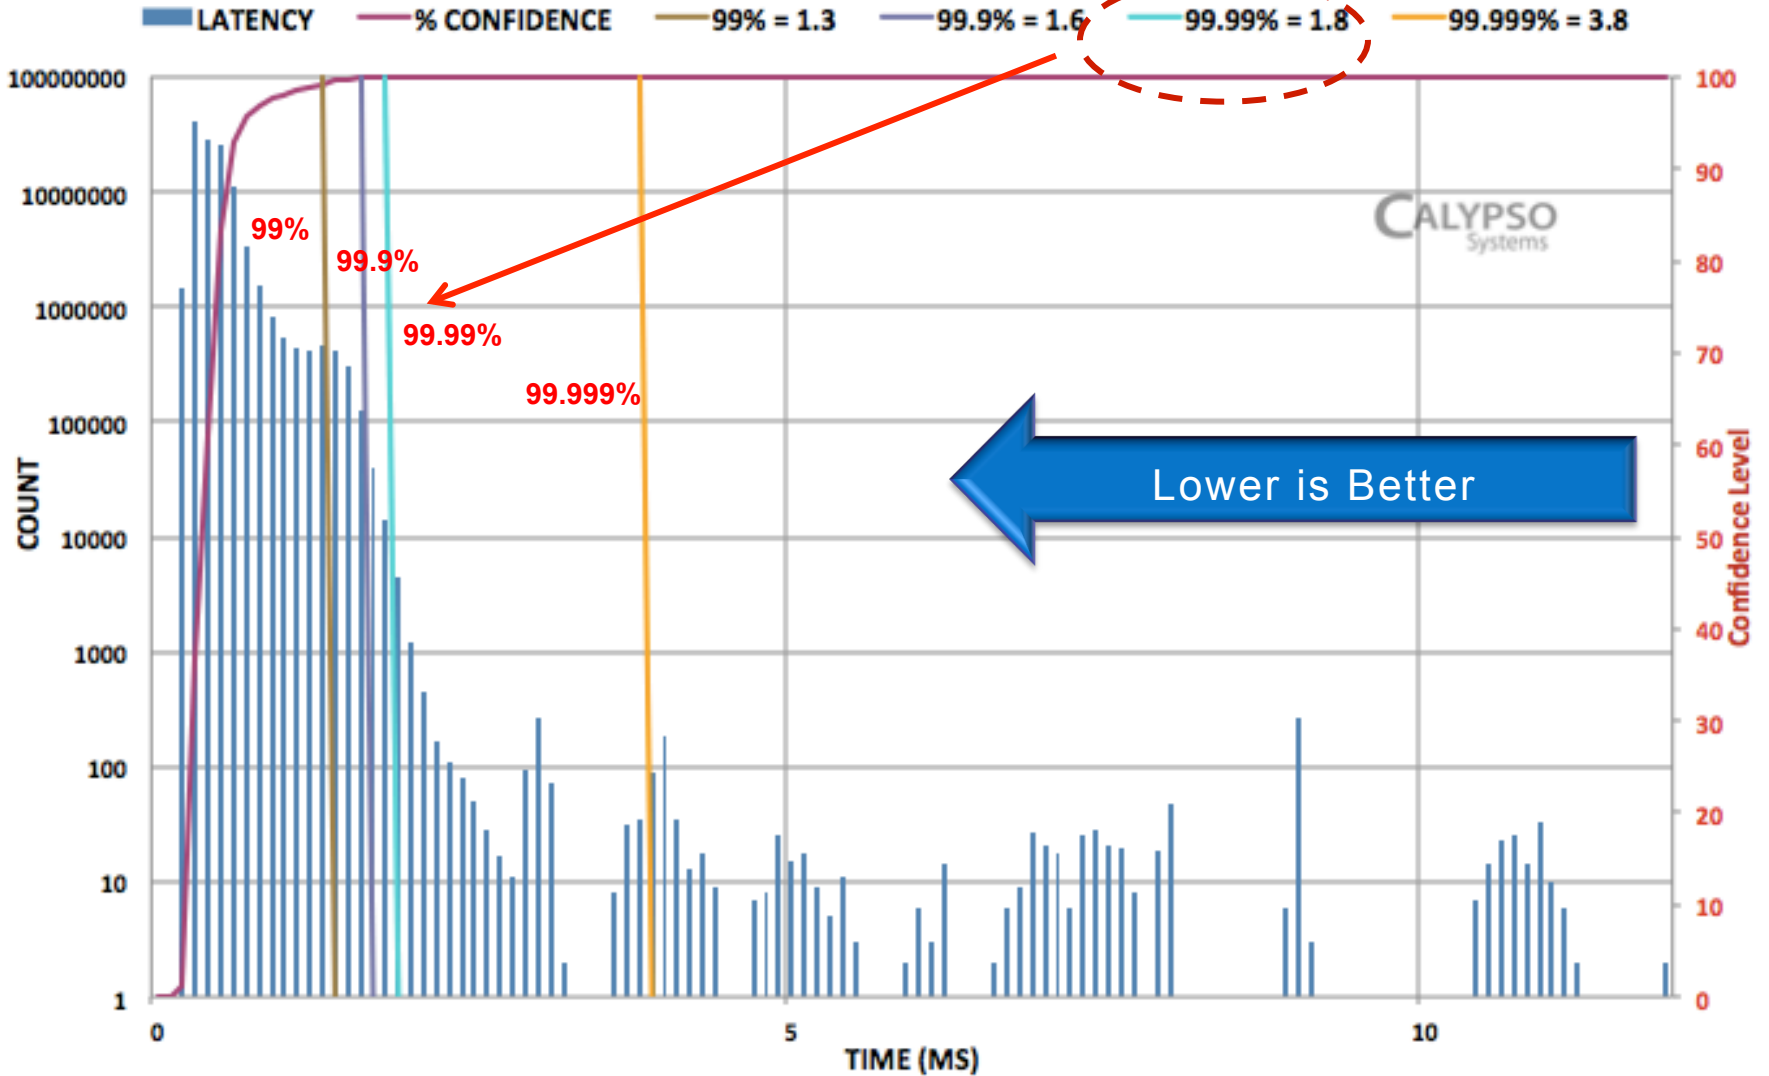
# WSAT RND 4KiB W – Response Time Histogram



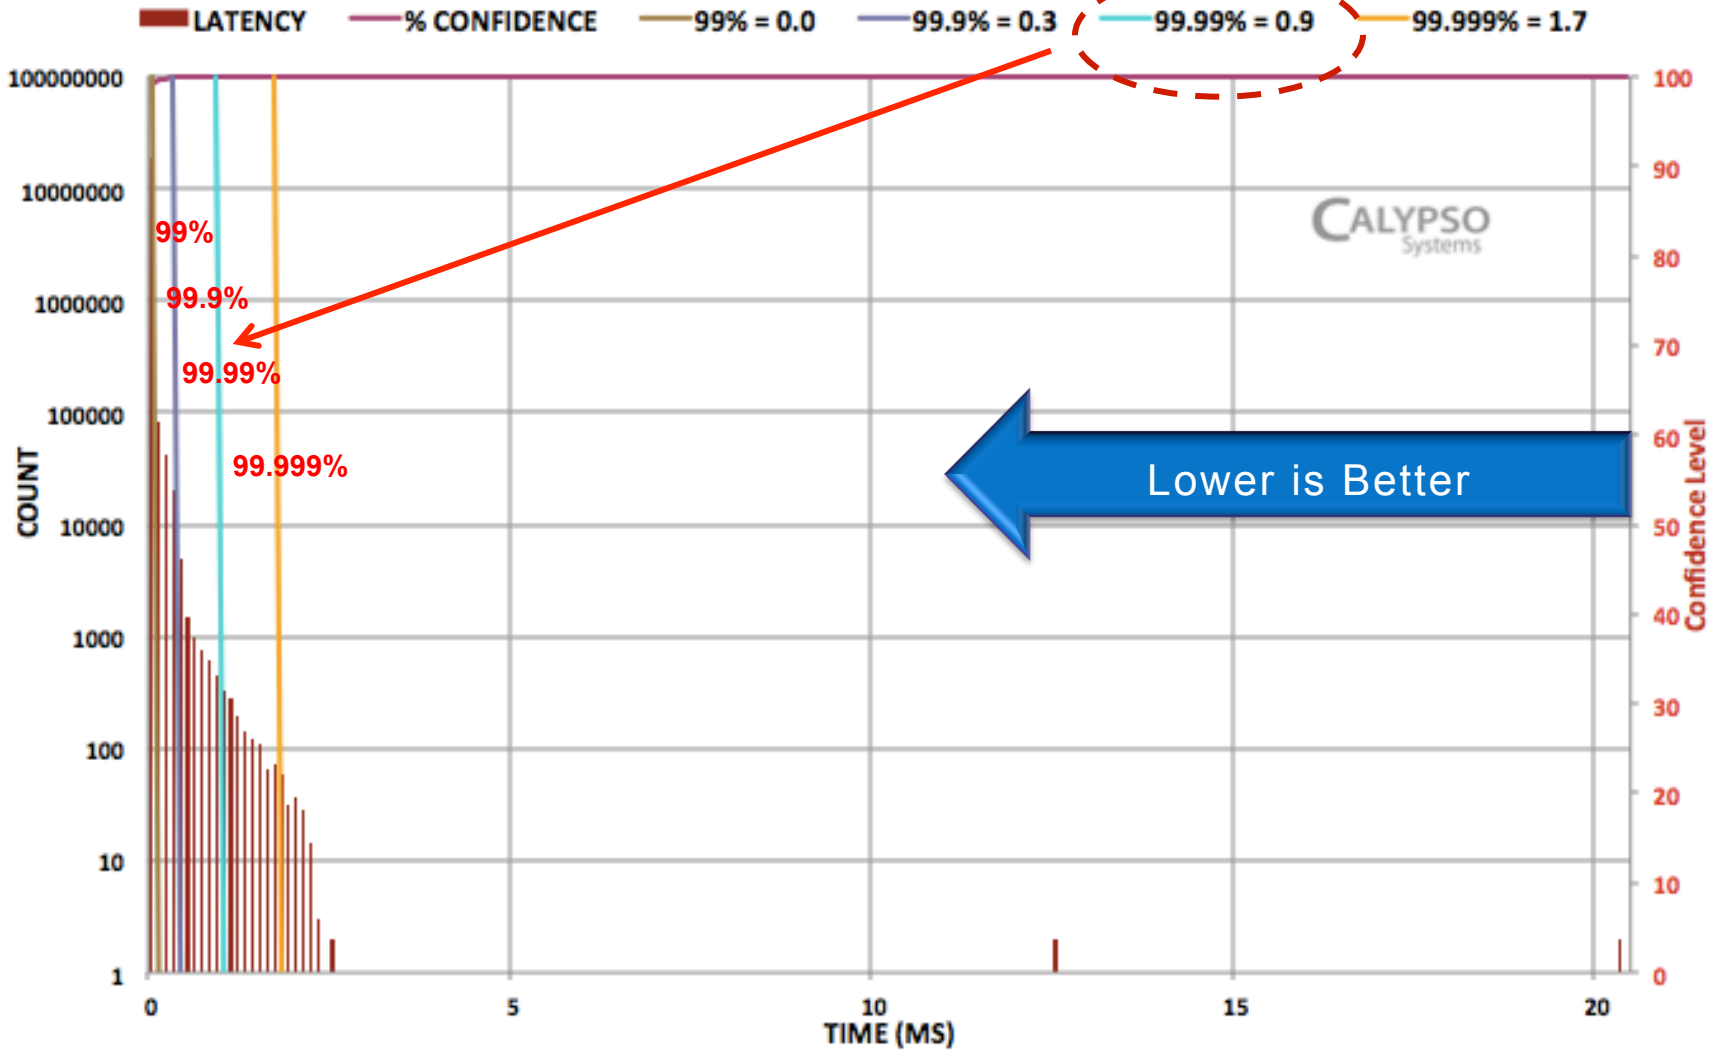
# PCIe E – Histogram Confidence Plot



# PCIe D – Histogram Confidence Plot



# PCIe B – Histogram Confidence Plot



# Phase II Tests – SNIA PTS Basic

## IOPS -

- RND IOPS
- 8 Block Size
- 7 RW Mix

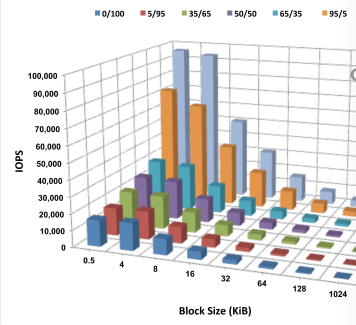
## Throughput –

- SEQ 1024KiB
- 100R & 100W

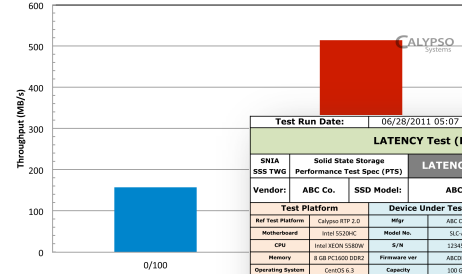
## Latency –

- .5, 4 & 8KiB
- 100R, 65:35, 100W

Test Run Date:		11/02/2011 02:56 PM		Report Run Date:		3/4/2013 8:43:00 AM		
<b>IOPS Test (REQUIRED) - Report Page</b>								
SNIA SSS TWG	Solid State Storage Performance Test Spec (PTS)		IOPS - Block Size x RW Mix Matrix				Rev. 5 of 6	PTS-E 1.1
Vendor:	ABC Co.	SSD Model:	ABC Co. SLC-A 100	TEST SPONSOR	CALYPSO Systems			
Test Platform	Device Under Test	Set Up Parameters		Test Parameters				
Ref Test Platform	Calypso RPT 2.0	Mfg	ABC Co.	Data Pattern	RND	Data Pattern	RND	
Motherboard	Intel S520DC	Model No.	SLC-A	AR	100%	AR & Amount	100%	
CPU	Intel Xeon S3500	S/N	123456	AR Segments	N/A	Test Stimulus 1	IOPS Loop	
Memory	8 GB PC10600 DDR2	Firmware ver	ABCDEF	Pre Condition 1	SEQ 128K W	RW Mix	Client IOPS	
Operating System	CentOS 6.3	Capacity	100 GB	TOID - TC/QD	TC 2/ QD 16	Block Size	Inner Loop	
Test SW	CTS 6.5.13.8	Interface	SATA GEN5	Duration	Twice User Capacity			
Test SW Info	1.10.72.8.16	RAND Type	SLC	Pre Condition 2	IOPS Loop			
Test ID No.	RS-856	PCIe MUX	N/A	TOID - TC/QD	TC 2/ QD 16			
HBA	LSI 9211-4641	Purge Method	Format LUN	SR Roundtrips	1 - 5			
PCIe	Gen 2 x 8	Write Cache	WCD	Note				
<b>IOPS - ALL RW Mix &amp; BS - 3D Columns</b>								



Test Run Date:		12/04/2011 08:21 AM		Report Run Date:		3/04/2013 10:03 AM		
<b>Throughput Test (REQUIRED) - Report Page</b>								
SNIA SSS TWG	Solid State Storage Performance Test Spec (PTS)		TP - SEQ 1024KiB				Rev. 5 of 10	PTS-E 1.1
Vendor:	ABC Co.	SSD Model:	ABC Co. SLC-A 100	TEST SPONSOR	CALYPSO Systems			
Test Platform	Device Under Test	Set Up Parameters		Test Parameters				
Ref Test Platform	Calypso RPT 2.0	Mfg	ABC Co.	Data Pattern	SEQ	Data Pattern	SEQ	
Motherboard	Intel S520DC	Model No.	SLC-A	AR	100%	AR & Amount	100%	
CPU	Intel Xeon S3500	S/N	123456	AR Segments	N/A	Test Stimulus 1	SEQ 1024KiB	
Memory	8 GB PC10600 DDR2	Firmware ver	ABCDEF	Pre Condition 1	SEQ 128K W	RW Mix	100R / 100W	
Operating System	CentOS 6.3	Capacity	100 GB	TOID - TC/QD	TC 1/ QD 1	Block Size	SEQ 1024KiB	
Test SW	CTS 6.5.13.8	Interface	SATA GEN5	Duration	Twice User Capacity	TOID - TC/QD	TC 2/QD 16	
Test SW Info	1.10.72.8.16	RAND Type	SLC	Pre Condition 2	SEQ 128K W	Steady State	1 - 5	
Test ID No.	RS-897	PCIe MUX	N/A	TOID - TC/QD	TC 2/ QD 16	Test Stimulus 2	SEQ 128K W	
HBA	LSI 9211-4641	Purge Method	Format LUN	SR Roundtrips	1 - 5	TOID - TC/QD	TC 2/QD 16	
PCIe	Gen 2 x 8	Write Cache	WCD	Note		Steady State	1 - 5	
<b>Throughput - ALL RW Mix &amp; BS - 2D Plot 1024KiB</b>								



Test Run Date:		06/28/2011 05:07 AM		Report Run Date:		3/04/2013 08:47 AM		
<b>LATENCY Test (REQUIRED) - Report Page</b>								
SNIA SSS TWG	Solid State Storage Performance Test Spec (PTS)		LATENCY - Response Time OIO=1				Rev. 4 of 6	PTS-E 1.1
Vendor:	ABC Co.	SSD Model:	ABC Co. SLC-A 100	TEST SPONSOR	CALYPSO Systems			
Test Platform	Device Under Test	Set Up Parameters		Test Parameters				
Ref Test Platform	Calypso RPT 2.0	Mfg	ABC Co.	Data Pattern	RND	Data Pattern	RND	
Motherboard	Intel S520DC	Model No.	SLC-A	AR	100%	AR & Amount	100%	
CPU	Intel Xeon S3500	S/N	123456	AR Segments	N/A	Test Stimulus 1	LAT Loop	
Memory	8 GB PC10600 DDR2	Firmware ver	ABCDEF	Pre Condition 1	SEQ LUN W	RW Mix	Client Loop	
Operating System	CentOS 6.3	Capacity	100 GB	TOID - TC/QD	TC 1/ QD 1	Block Size	Inner Loop	
Test SW	CTS 6.5.13.8	Interface	SATA GEN5	Duration	Twice User Capacity	TOID - TC/QD	TC 1/QD 1	
Test SW Info	1.10.91.9.33	RAND Type	SLC	Pre Condition 2	LAT Loop	Steady State	1 - 7	
Test ID No.	RS-523	PCIe MUX	N/A	TOID - TC/QD	TC 1/ QD 1	Histogram	N/A	
HBA	LSI 9211-4641	Purge Method	Format LUN	SR Roundtrips	3 - 7	TOID - TC/QD	N/A	
PCIe	Gen 2 x 8	Write Cache	WCD	Note				
<b>Average and Maximum Response Time - ALL RW Mix &amp; BS - Tabular Data</b>								

Average Response Time (ms)				
Read / Write Mix %				
Block Size (KiB)	0/100	65/35	100/0	
0.5	0.0811486	0.1307696	0.1470862	
4	0.0865518	0.1386792	0.155532	
8	0.1300188	0.1697318	0.171237	

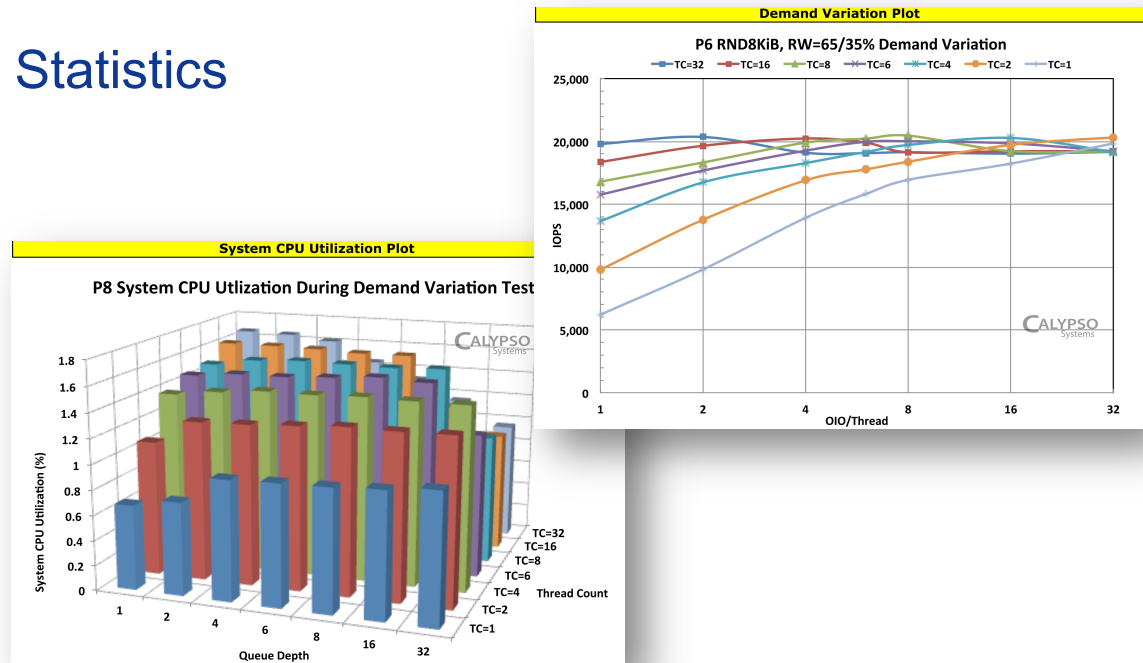
  

Maximum Response Time (ms)				
Read / Write Mix %				
Block Size (KiB)	0/100	65/35	100/0	
0.5	56.8668	39.4566	5.7996	
4	32.6272	46.7808	6.2596	
8	30.4264	39.3802	12.8666	

# Phase III Tests – SNIA PTS Advanced

## Demand Intensity Response Time Histogram

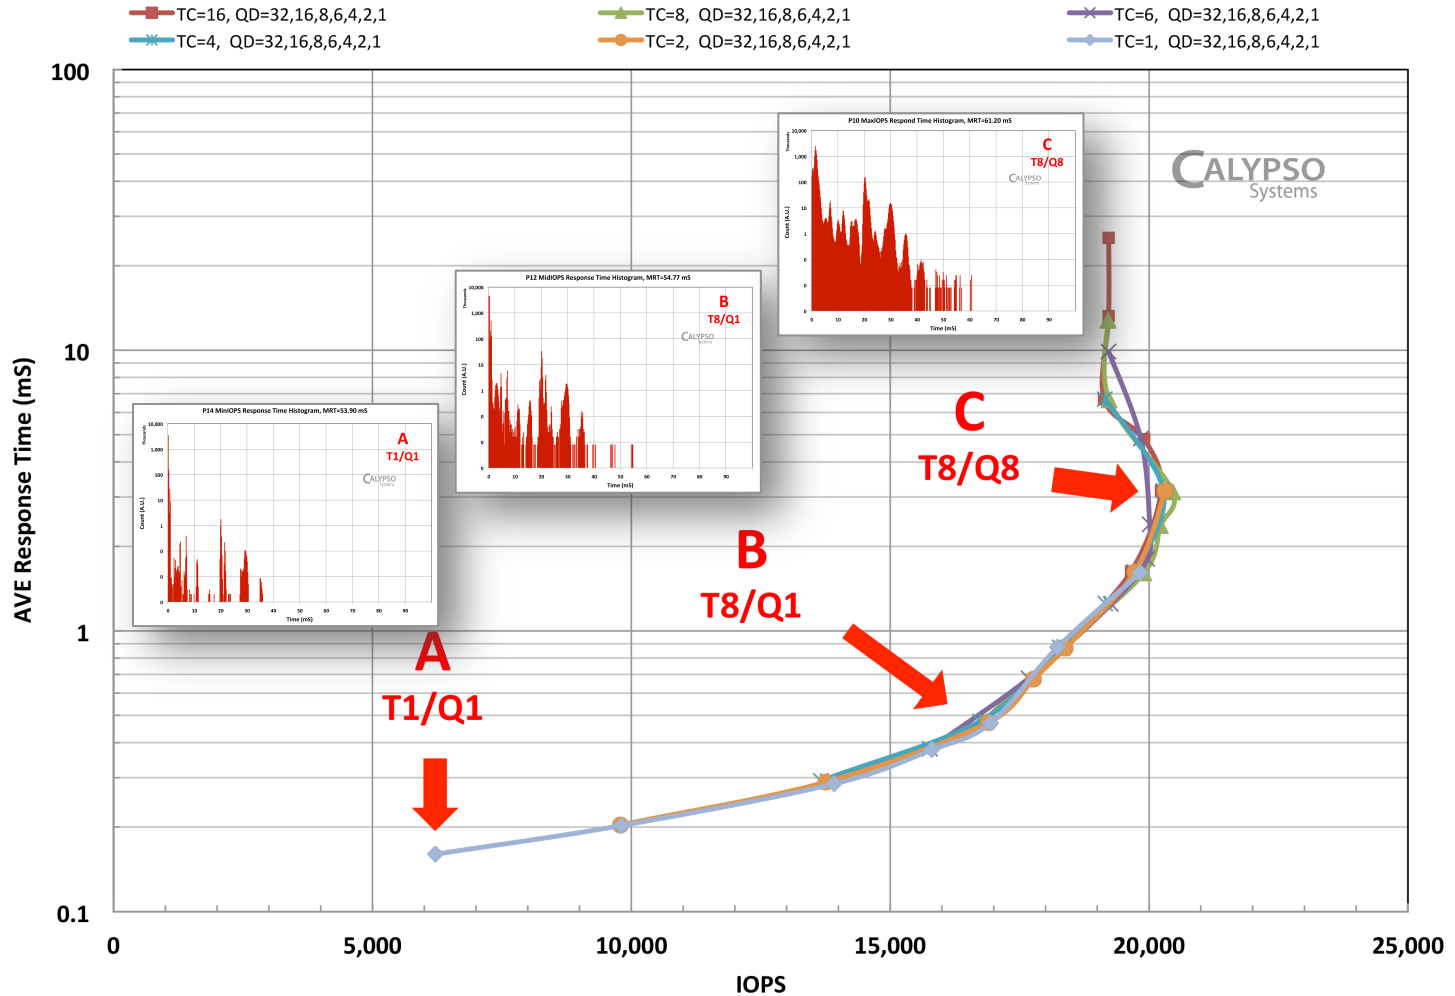
- Select / Set Access Pattern Workload (e.g. OLTP / VOD)
- Run IOPS while Varying Outstanding IO (Thread & Queue)
- Plot Demand Variation
- Plot CPU Usage Statistics





# Phase III Tests – DIRTH

## P7 RND/8KiB, R/W=65/35% Demand Intensity



CALYPSO Systems

# Phase IV Tests – Workload Based

## Specific Application Workload Tests

- Selected Single Workload Access Patterns
- Complex Pre-conditioning (mixed workloads)
- Complex LBA zones (hot spots)
- Examples of potential Workloads:

Webservers, Exchange Mail, Media Streaming,  
File Servers, database OLTP, OS paging, VOD,  
webserver logs, SQL logs, archiving, medial imaging and more

## PCIe SSDs are High Performance

- “IOPS & Bandwidth are easy, Latencies are hard”
- Examination of Response Times & CPU Usage is key
- Be sure to Define the Test Workload
- Use a Standardized & Normalized Test Environment



Thank you

Contact Calypso for Blind Survey reports at

[info@calypsotesters.com](mailto:info@calypsotesters.com)