Flash Memory SUMMIT

Tutorial T1C
Testing/Performance/Endurance

Changing Dynamics of Flash Performance Benchmarks

Bob Weisickle – OakGate Technology
Mike Engbretson – Granite River Labs

August 10, 2010

OakGate Technology
Granite River Labs

Agenda

- Market Trend - update
- Changes in Performance Landscape
- Benchmarking Methods & Standards
- Challenges beyond just Performance
- Ways OakGate Technology and Granite River Labs can help

SSD Market Trend

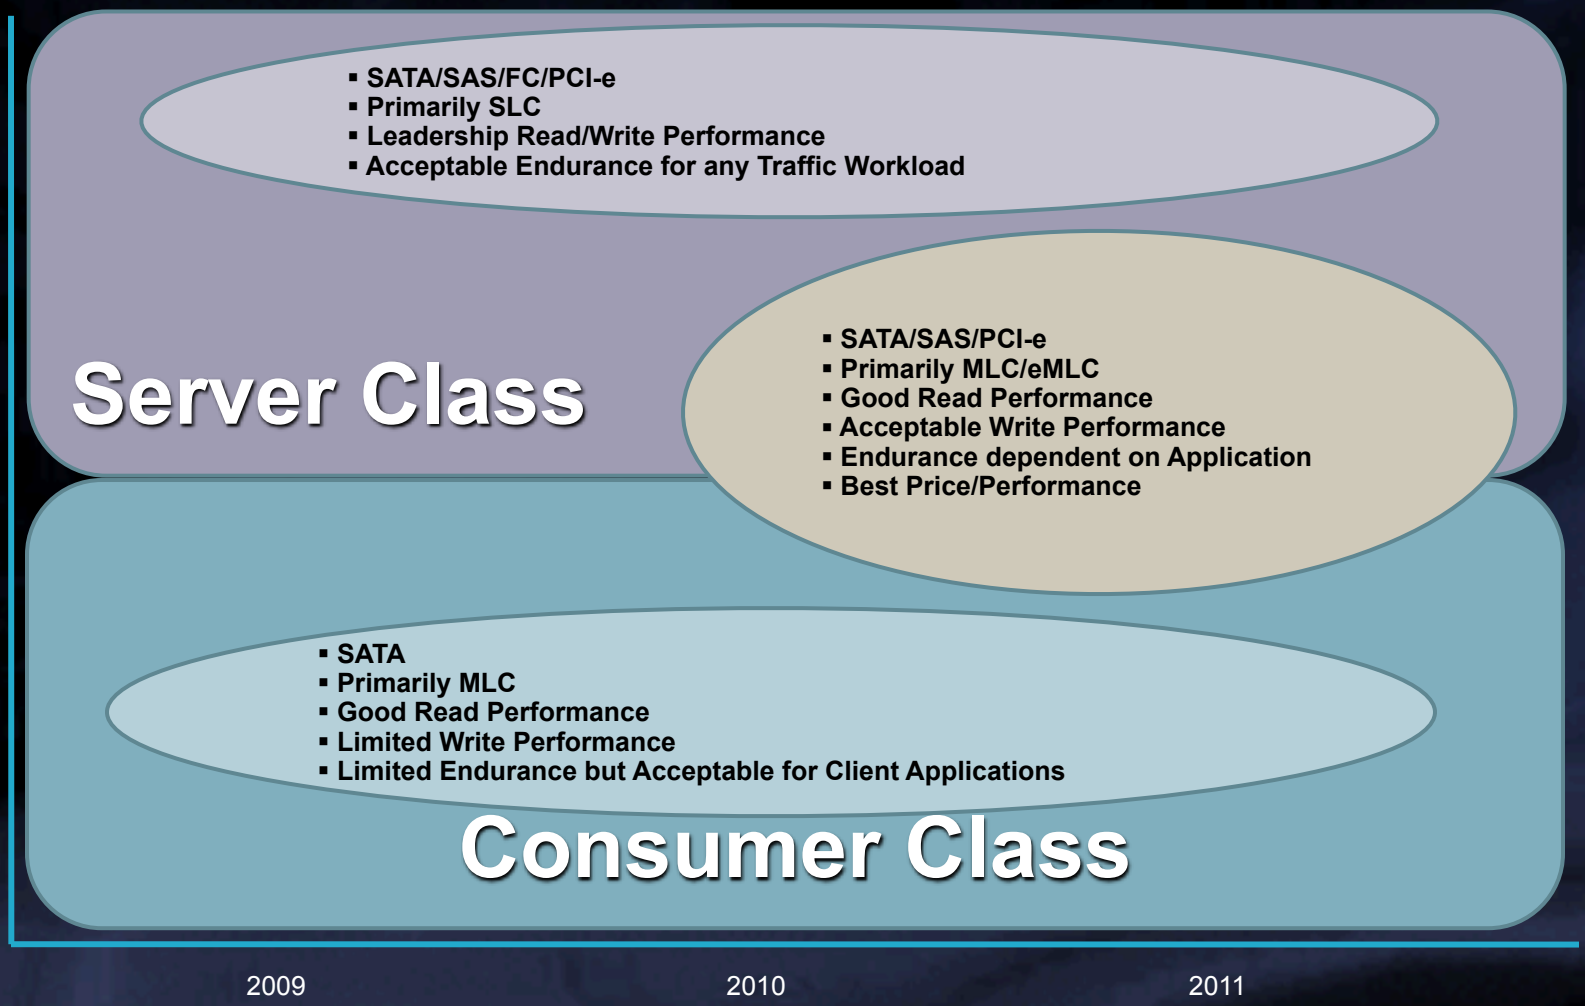

Performance Landscape

- 3rd / 4th Generation of Controllers
- Improved Performance (especially writes)
- Improved OTB versus Steady State Performance
- Refinements in Wear-leveling and Garbage Collection algorithms
- New FLASH (eMLC) that improves endurance
- Use of Data Compression to improve Write Performance and Write Amplification
- 4K IO Optimization – general trend to 4K Sectors

OTB versus Steady State

MLC FLASH – 2 hour pre-conditioning
2010 2011

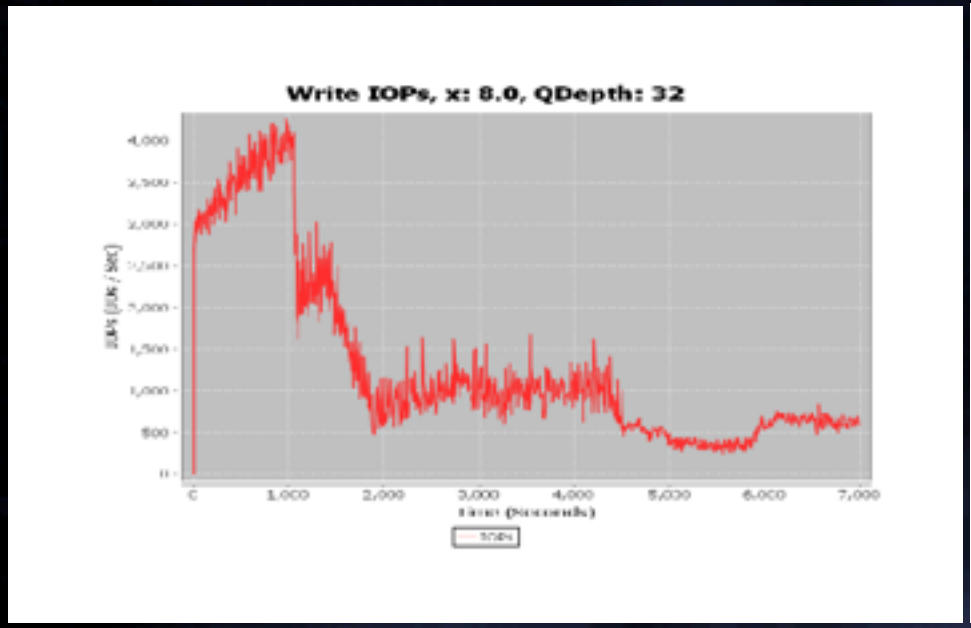

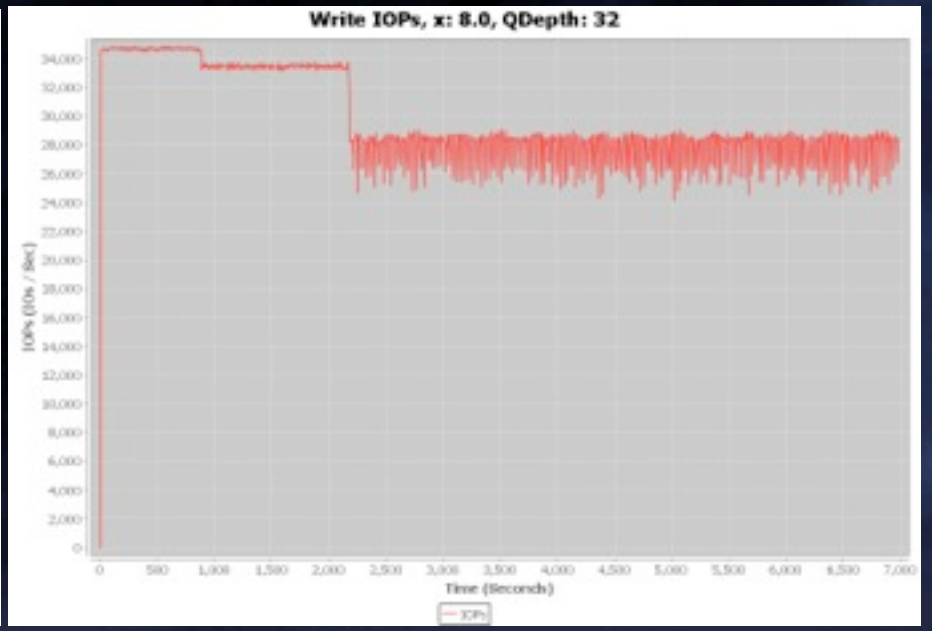

- Higher Performance
- Quicker transition to Steady State

OakGate Technology
Granite River Labs

Impact of Data Compression

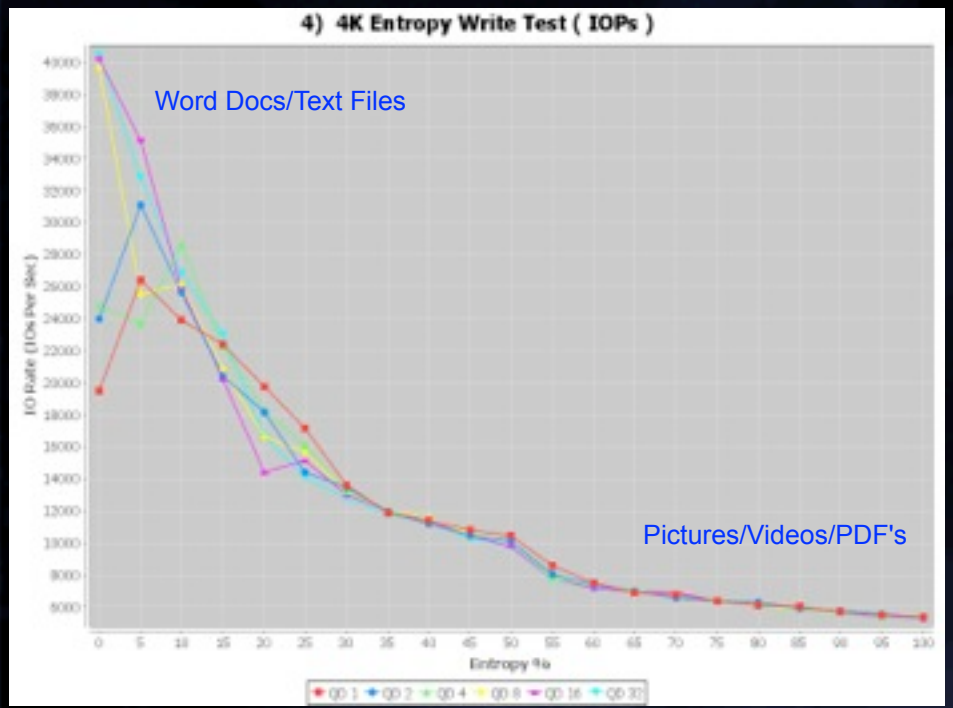

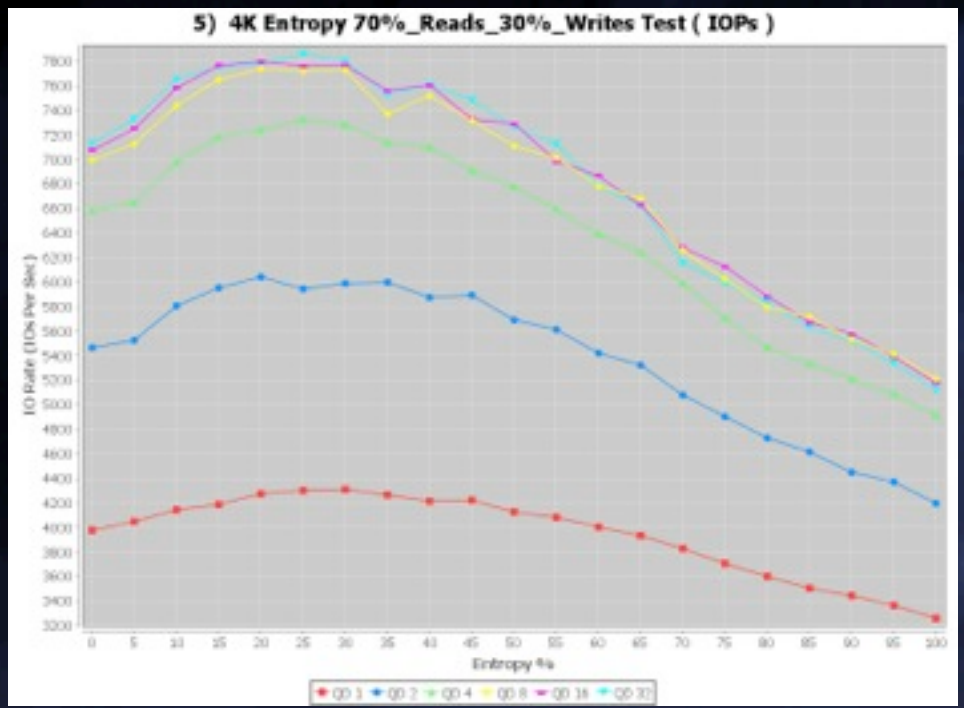

- Need to understand the Entropy of the Real Data

OakGate Technology
Granite River Labs

4K IO Optimized

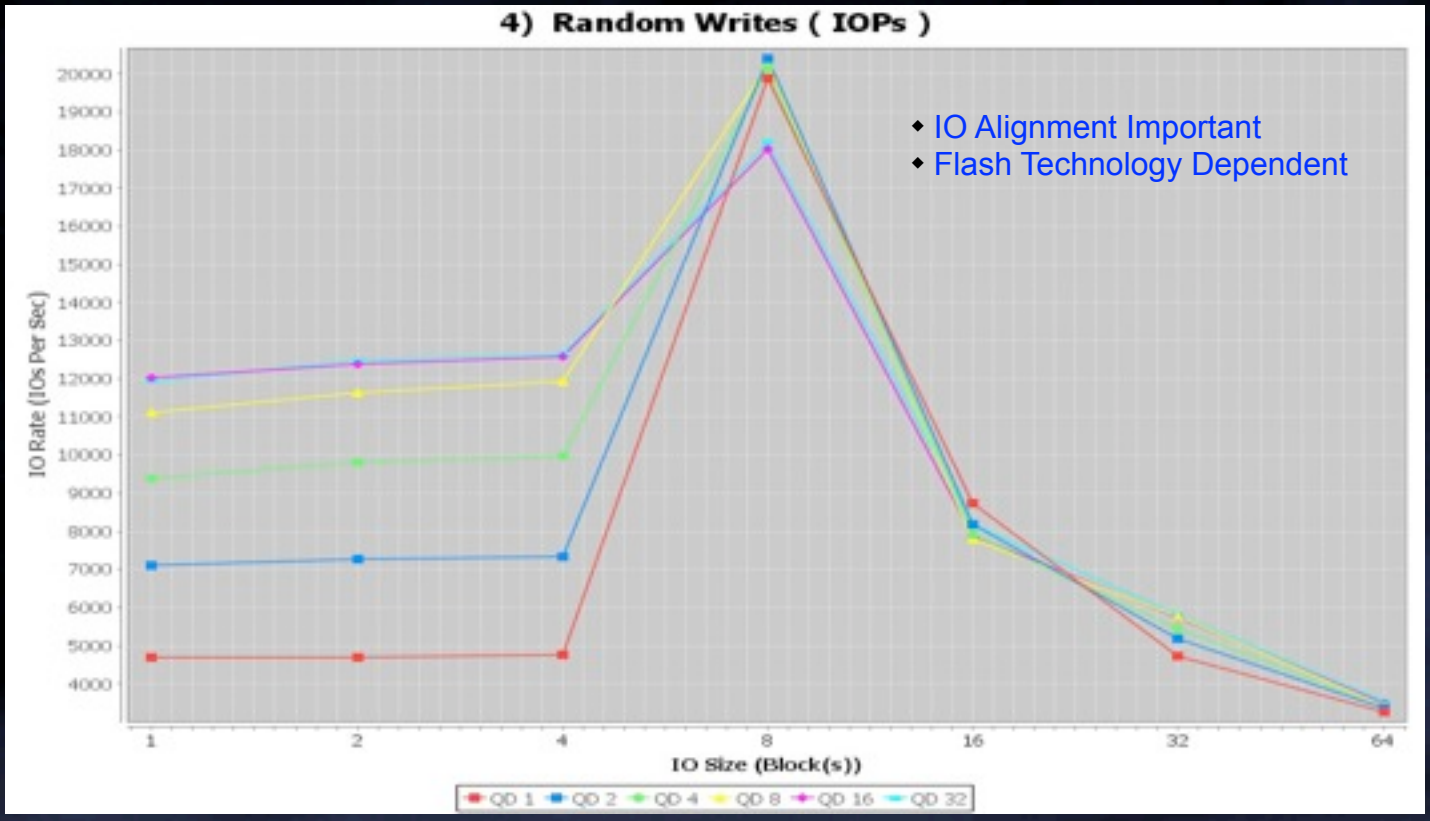

OakGate Technology
Granite River Labs

Benchmarking Methods

- Standards Based Performance/Endurance
  - SNIA – Solid State Storage Performance Test Specification Enterprise 1.0
  - JEDEC – JESD218A and JESD219
- User Defined Performance Measurement
  - Application Specific
  - Synthetic workloads
  - Captured Traffic workloads
- Functionality Validation
  - Conformance and Error Injection

Benchmarking Methods

- SNIA – SSS-PTS-Enterprise Ver. 1.0

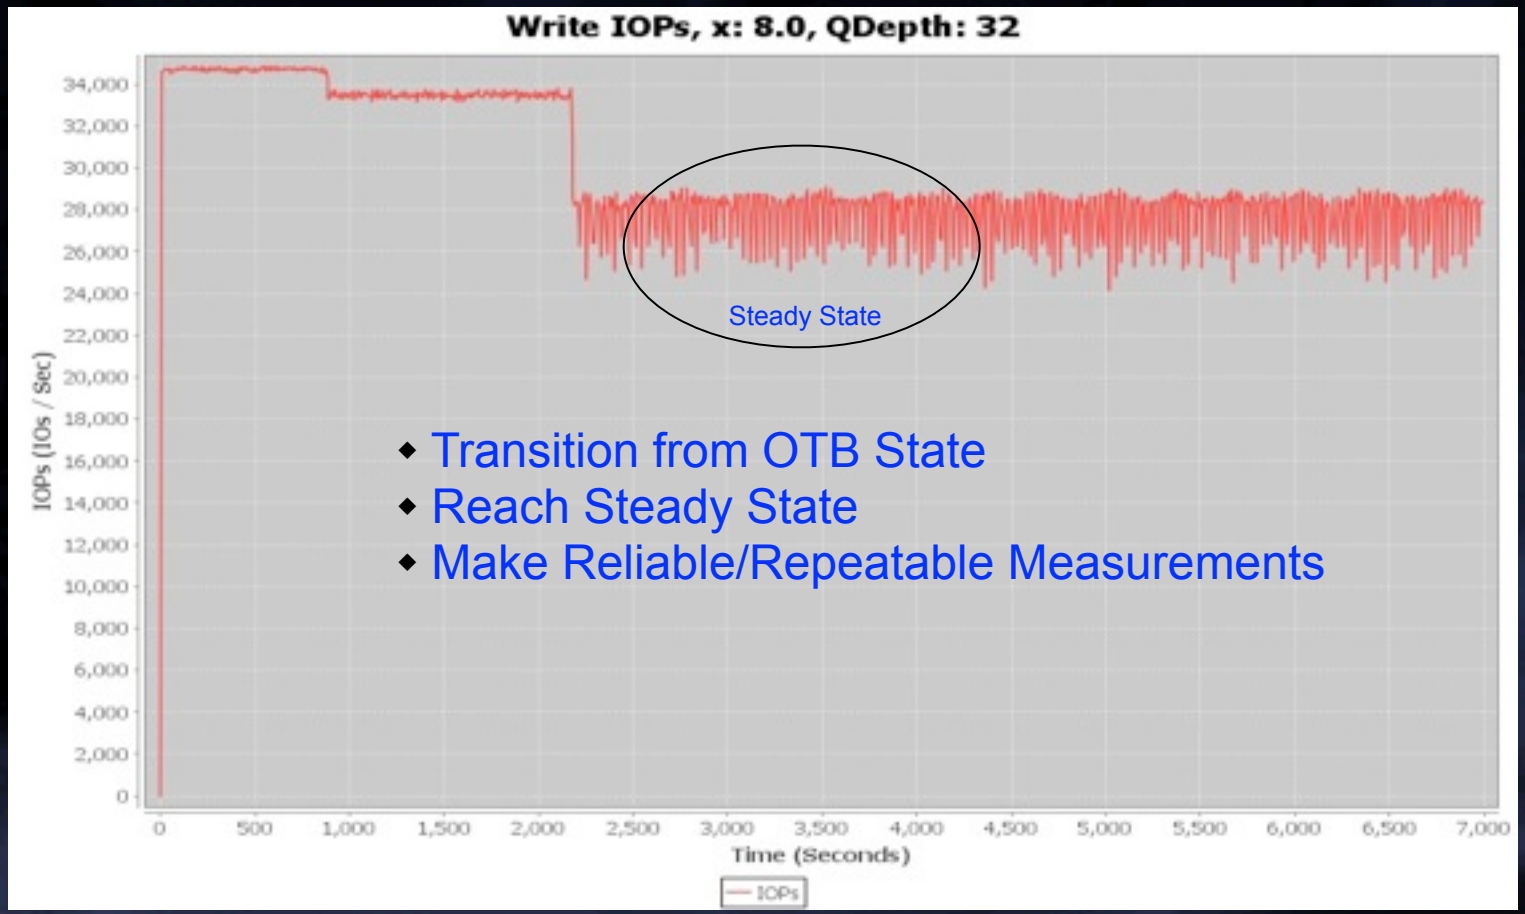

Benchmarking Methods

- SNIA – SSS-PTS-Enterprise Ver. 1.0
  - Measurement Rounds
  - Measurement Convergence

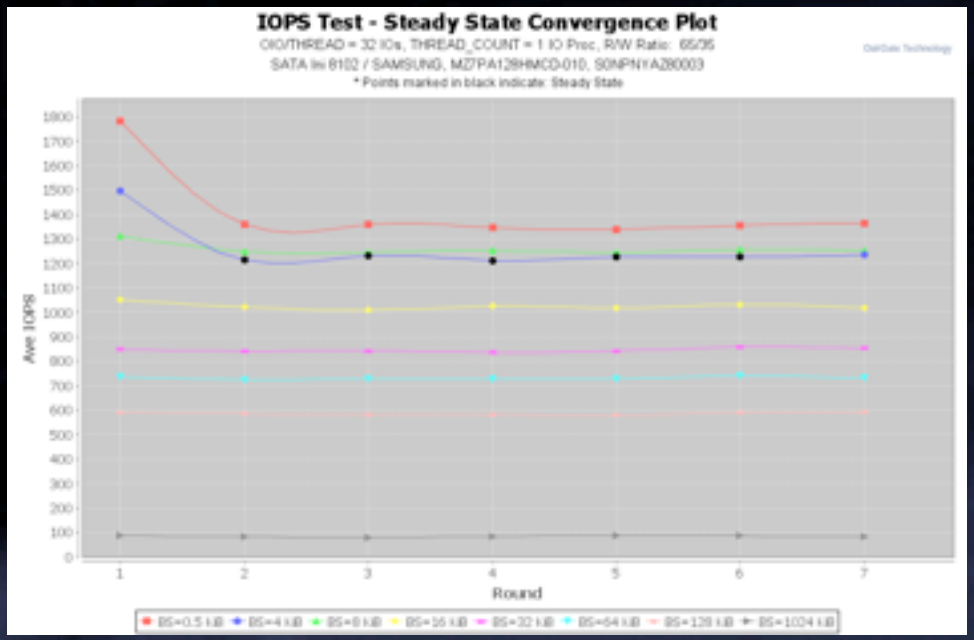

Blk Size	0/100	95/5	65/35	50/50	35/65	5/95	100/0
0.5 KiB	1392	3782	1366	1605	1990	1411	30825
4 KiB	987	2963	1235	1466	1441	1029	22455
8 KiB	815	2418	1252	1347	1166	839	15619
16 KiB	607	1819	1020	1167	880	653	9902
32 KiB	395	1300	852	1001	558	408	5780
64 KiB	238	894	731	418	345	249	3187
128 KiB	134	603	592	242	189	138	1675
1024 KiB	174	158	83	73	71	119	218

OakGate Technology
Granite River Labs

Challenges beyond Performance

- Data Validation
  - Data Integrity Checking (including stale data and missed writes)
  - Data commit during power failure/recovery
- Endurance Prediction
  - Using SMART Attributes
  - JEDEC
- Latency
  - IO Latency Distribution
  - Understand impact on application
- Protocol/Command Robustness
  - Device Software Reliability

Latency Distribution
Example 1

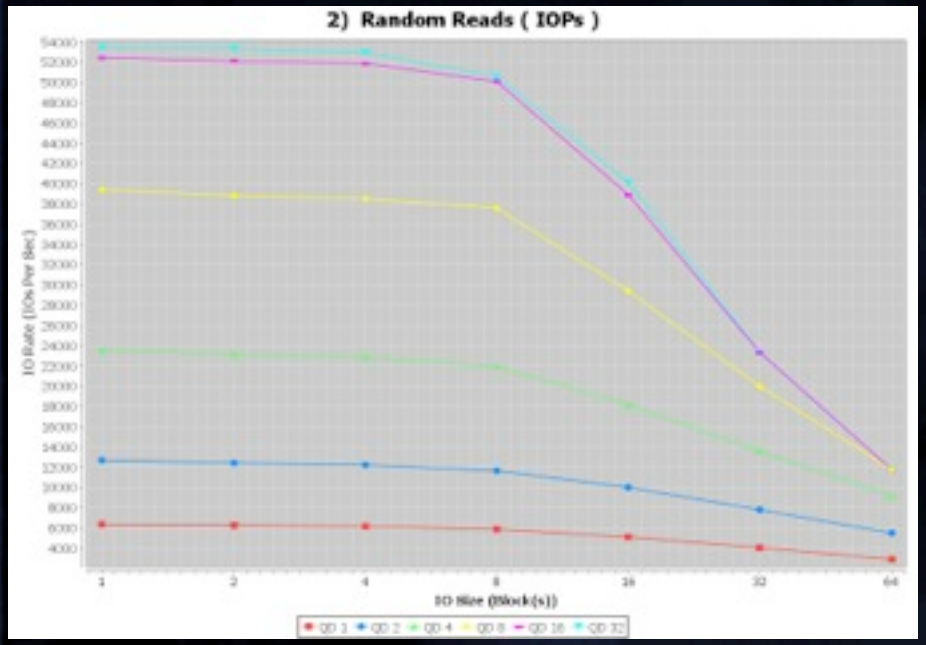

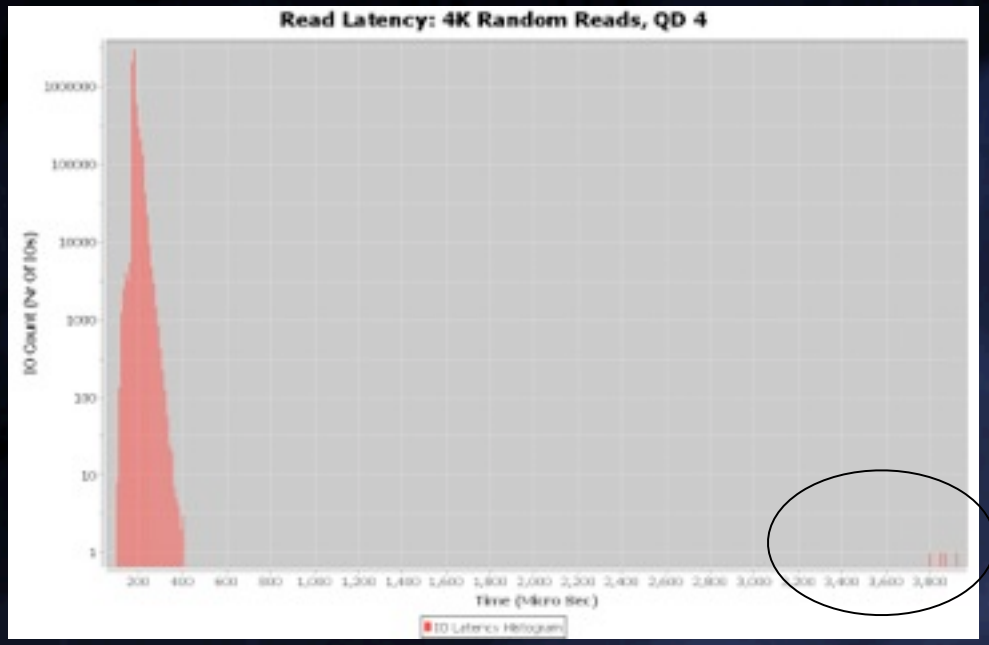

- A few longer latency IO's but in general very good distribution

OakGate Technology
Granite River Labs

Latency Distribution
Example 2

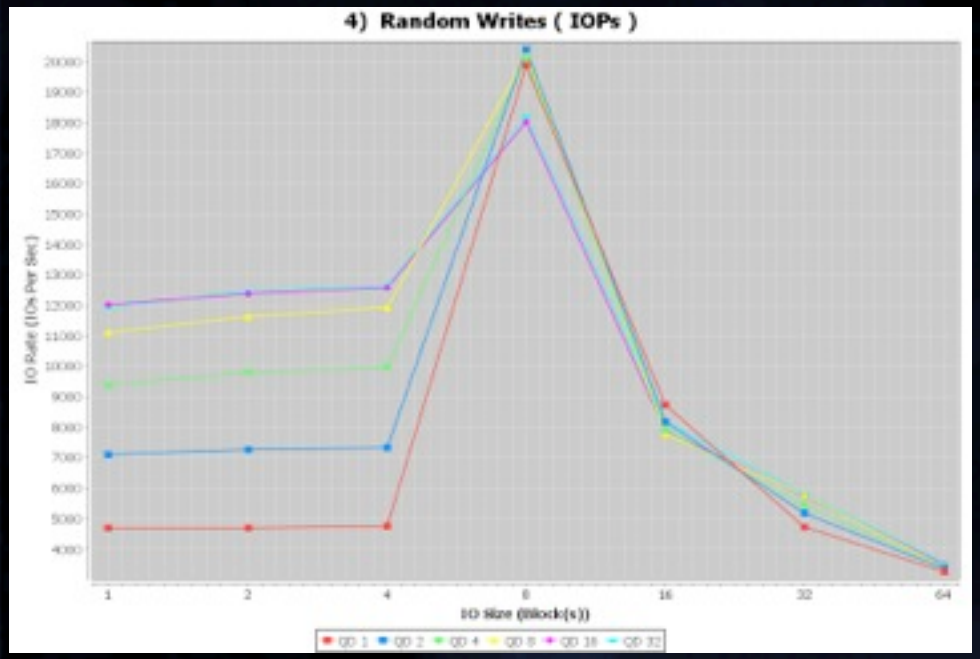

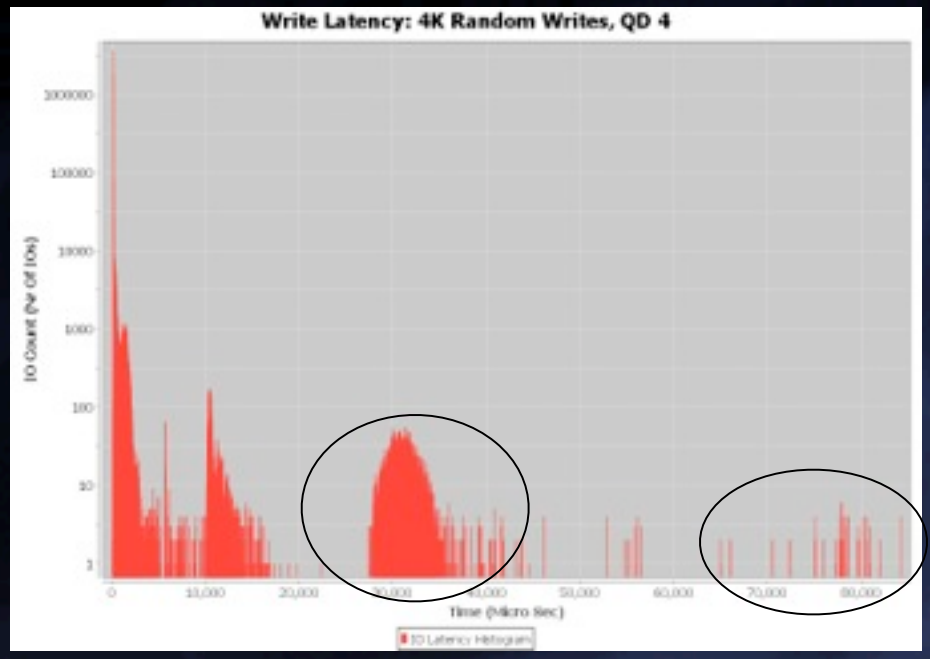

- Large groupings of long latency IO's
- Impact on Application Performance
- Impact on Raid Controller Performance

OakGate Technology
Granite River Labs

Smart Attributes (Endurance Related)

ID	Attribute Name	A	B	C	D	E	F	G	H
5	Retired Block Count – Life Indicator	X		X	X	X	X	X	X
171	Program Fail Block Count – Life Indicator			X	X		X		X
172	Erase Fail Block Count – Life Indicator			X	X		X		X
177	1) Wear Range Delta 2) Wear Leveling Count		2		2		2	2	1
231	1) SSD Life Left 2) Temperature (degC)				2		2	2	1
232	Vendor Unique	X					X		
233	Media Wear-out Indicator	X	X		X		X	X	X
241	Lifetime Writes from HOST				X		X	X	X

Data Integrity/Power Fail Testing

- Why this is Important
  - Young and Maturing Technology (FLASH, Controllers, Super Caps)
  - New Software Algorithms
  - Data is not Stationary (wear-leveling)
- What should be tested
  - Super Cap or Equiv Hold up time – did all data get committed in time
  - Extended Data Checking (correct data and from correct location/time)
  - Extended Run periods to stress wear-leveling
  - Power off – data retention intervals

OakGate Technology/Granite River Labs

- Validation & Performance System
  - SNIA and Custom Performance Benchmarks
  - JEDEC Compliant Endurance test suite
  - User definable/customized Benchmarks and Validation test suites
  - Data Integrity/Power Fail Test Suite
  - Full API for fully vendor unique tests development
- Full set of Services
  - SATA/SAS physical layer compliance
  - SATA Interop Testing
  - SAS/SATA Device Benchmarking
  - Data Integrity and Power Cycle Testing

SATA/SAS Compliance - Overview

What SATA Compliance & SAS Conformance DO address:

Standard	Official Logo Program?	PHY	Digital	System Interop	Mechanical	Notes
SATA	Y – Administered by SATA-IO	Y	Y	Y	Y	"Building Blocks" cert program available for IC components DOS-based scripts used for System Interop tests
SAS	N – "Conformance" based on test methodologies developed by UNH	Y	Y	Not defined	Not defined	Receiver PHY jitter tolerance requirements defined but no conformance test spec RX/TX (return loss/ impedance) test accepted practice includes "gating" out the connector

SATA/SAS Compliance - Overview

What SATA Compliance & SAS Conformance DO NOT address:

PHY Stress Testing	"How much input jitter can my receiver tolerate and still pass compliance?"
PVT Characterization	"How much margin does my product have in meeting spec under a range of PVT conditions and what are my points of failure?"
Extensive Interop & System Validation	"Will my device/host interoperate with a wide range of products and system environments?"
Functional Stress Tests	"How well does my product handle a variety of real-world and corner case test conditions?"
Performance Benchmarking	"How does my product stack up against industry benchmarks and competitors' products?"

OakGate Technology/Granite River Labs

- Contact OakGate Technology, Inc
  - bob.weisickle@oakgatetech.com
  - www.oakgatetech.com

- Contact Granite River Labs
  - Mike Engbretson, Chief Technology Engineer mikeen@graniteriverlabs.com
  - Quintin Anderson, COO qanderson@graniteriverlabs.com
  - www.GraniteRiverLabs.com