Flash vs. Disk Storage: Testing Workloads is Key

Len Rosenthal

VP of Marketing

SwiftTest

Flash Memory SUMMIT

SwiftTest Overview

- The leader in Storage Performance Validation.
- Our Mission: *To provide storage experts with critical insight to cost-effectively develop and deliver high performance, high availability storage infrastructure.*
- Founded in 2008; venture-backed

Azure Capital Partners

BGV

FLOODGATE

- Headquarters in Santa Clara, CA

EMC² NetApp hp skyera NEXENTA Enterprise class storage for everyone panasas CISCO

DELL ORACLE IBM HUAWEI AVERE CacheIQ HITACHI DATA SYSTEMS

ADP GE GoDaddy.com ELLIE MAE WELLS FARGO VISA

Key Customer Questions when Considering SSDs vs. HDDs

- Which applications will benefit from flash?
- How much faster would flash be over my current HDD storage arrays?
- How much head room do I have with a specific flash configuration?
- How much flash does my app really need?
- Hybrid or all-flash?
- Which vendor offers the best value to support my workload demands and budget constraints?
- How can I evaluate flash without impacting my production storage?
- How do I validate the performance impact of caching, compression, dedupe & tiering?

SwiftTest

Simulating a Production Workload

PRODUCTION STATS
(Workload Analyzers, Netstat, NFSstat, etc)

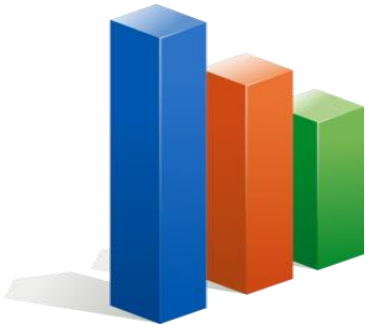

PACKET CAPTURES
(PCAP, etc)

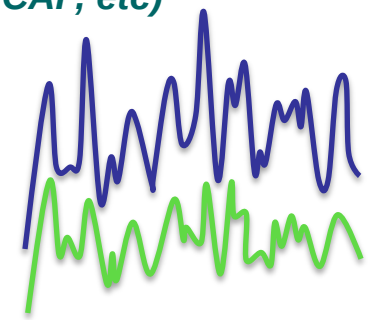

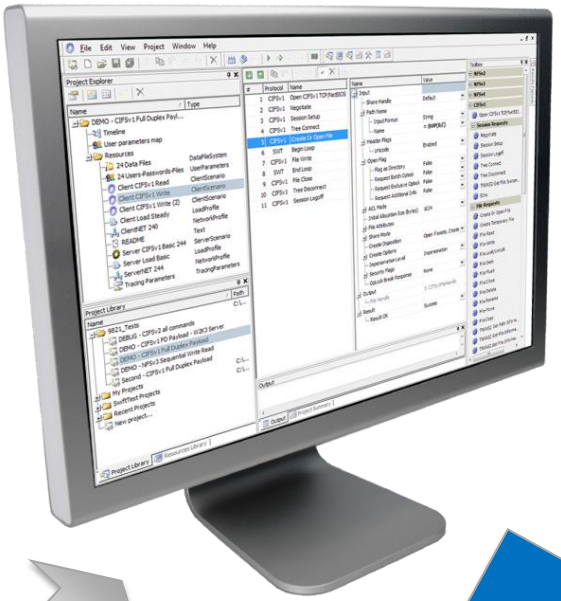

PRE-BUILT TEST SUITES

ACCURATE, REALISTIC WORKLOAD MODEL

SwiftTest

Realistic Workload Simulations that Don't Require Banks of Servers/VMs to Drive Are Key

	GoDaddy.com In Production	SwiftTest Simulation
Total NFS ops	~65K	~66K
Avg. Latency • Read • Write • Metadata Ops	1.5 ms 10 ms 0.5 ms 0.5 ms	1.4 ms 11.5 ms 0.6 ms 0.4 ms

Op-Mix
• Reads
• Writes
• Metadata Ops
  • Getattr
  • Lookup
  • Access
	
	5%
	1%
	94%
	62%
	11%
	17%
		
		5%
		1%
		94%
		63%
		13%
		14%

Avg. CPU Utilization	81%	80%
Max. Disk Utilization	55%	54%

Workload models proved that a Hybrid SSD/HDD approach based on an open server platform would support workloads at less than 50% of the cost

Workload Models help Evaluate New Flash Storage Products without impacting Production Storage Systems

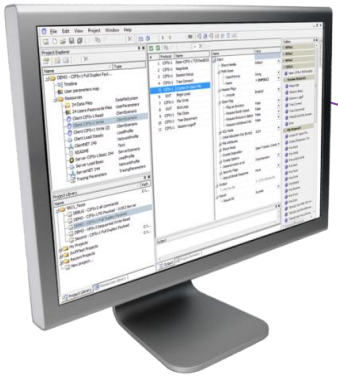

Test Development Environment
(File, Block, Object)

Massively scalable Performance Validation Appliance (NFS, FC, CIFS/SMB, iSCSI, HTTP)

Switch

Storage Pilot Candidate(s)

Solution A

Solution B

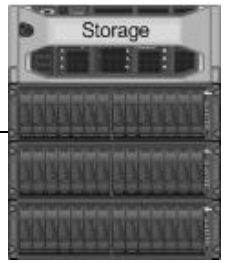

Production Test Example

SwiftTest

How Can You Dramatically Compress Time to Insight

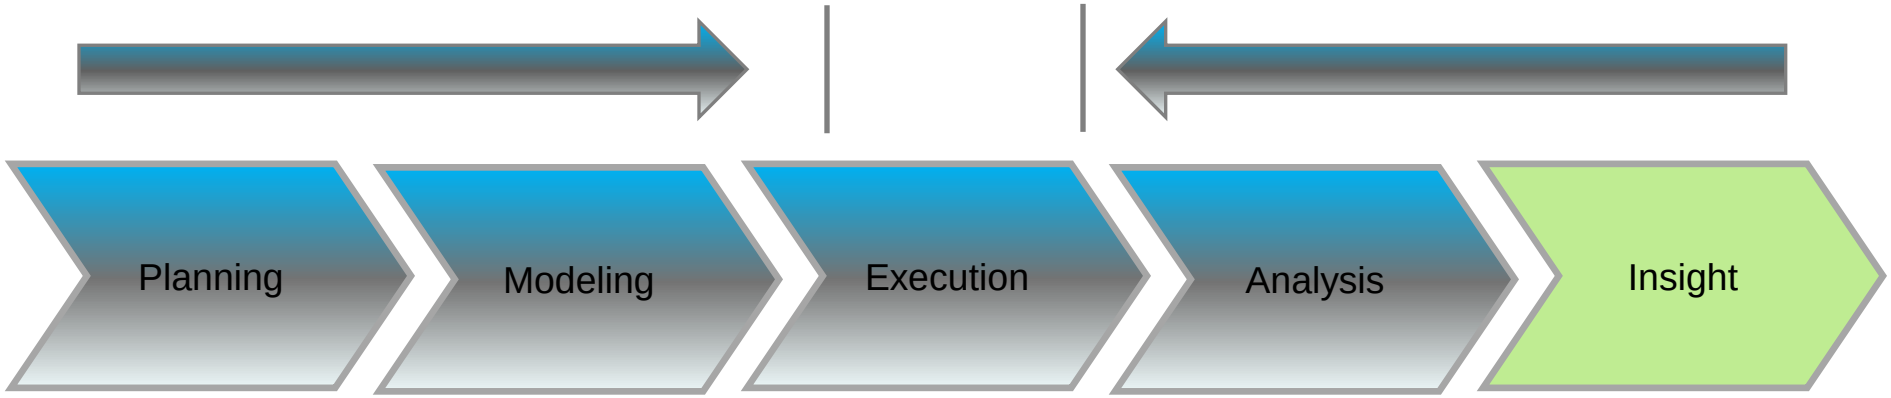

- What to test?
- What workloads should be used?
- What forms of testing?

- Building test cases
- Configuring DUT/SUT
- Workload/pattern modeling

- Running tests
- Building Regression suites
- PASS/FAIL criteria

- Measurement Interpretation
- Building PASS/FAIL criteria

- **Where are my performance limits?**
- **What are the optimal configurations**
- **Where are issues and defects?**
- **How many bugs?**

Dramatically compress time to insight by introducing:

1. **Modeling tools – statistical workload models**
2. **Simulation/Automation tools – ability to automate test workflow and run large numbers of test simulations**
3. **Analytics tools – ability to generate results against specific methodology objectives – not merely report metrics**

Flash Storage Validation Functions

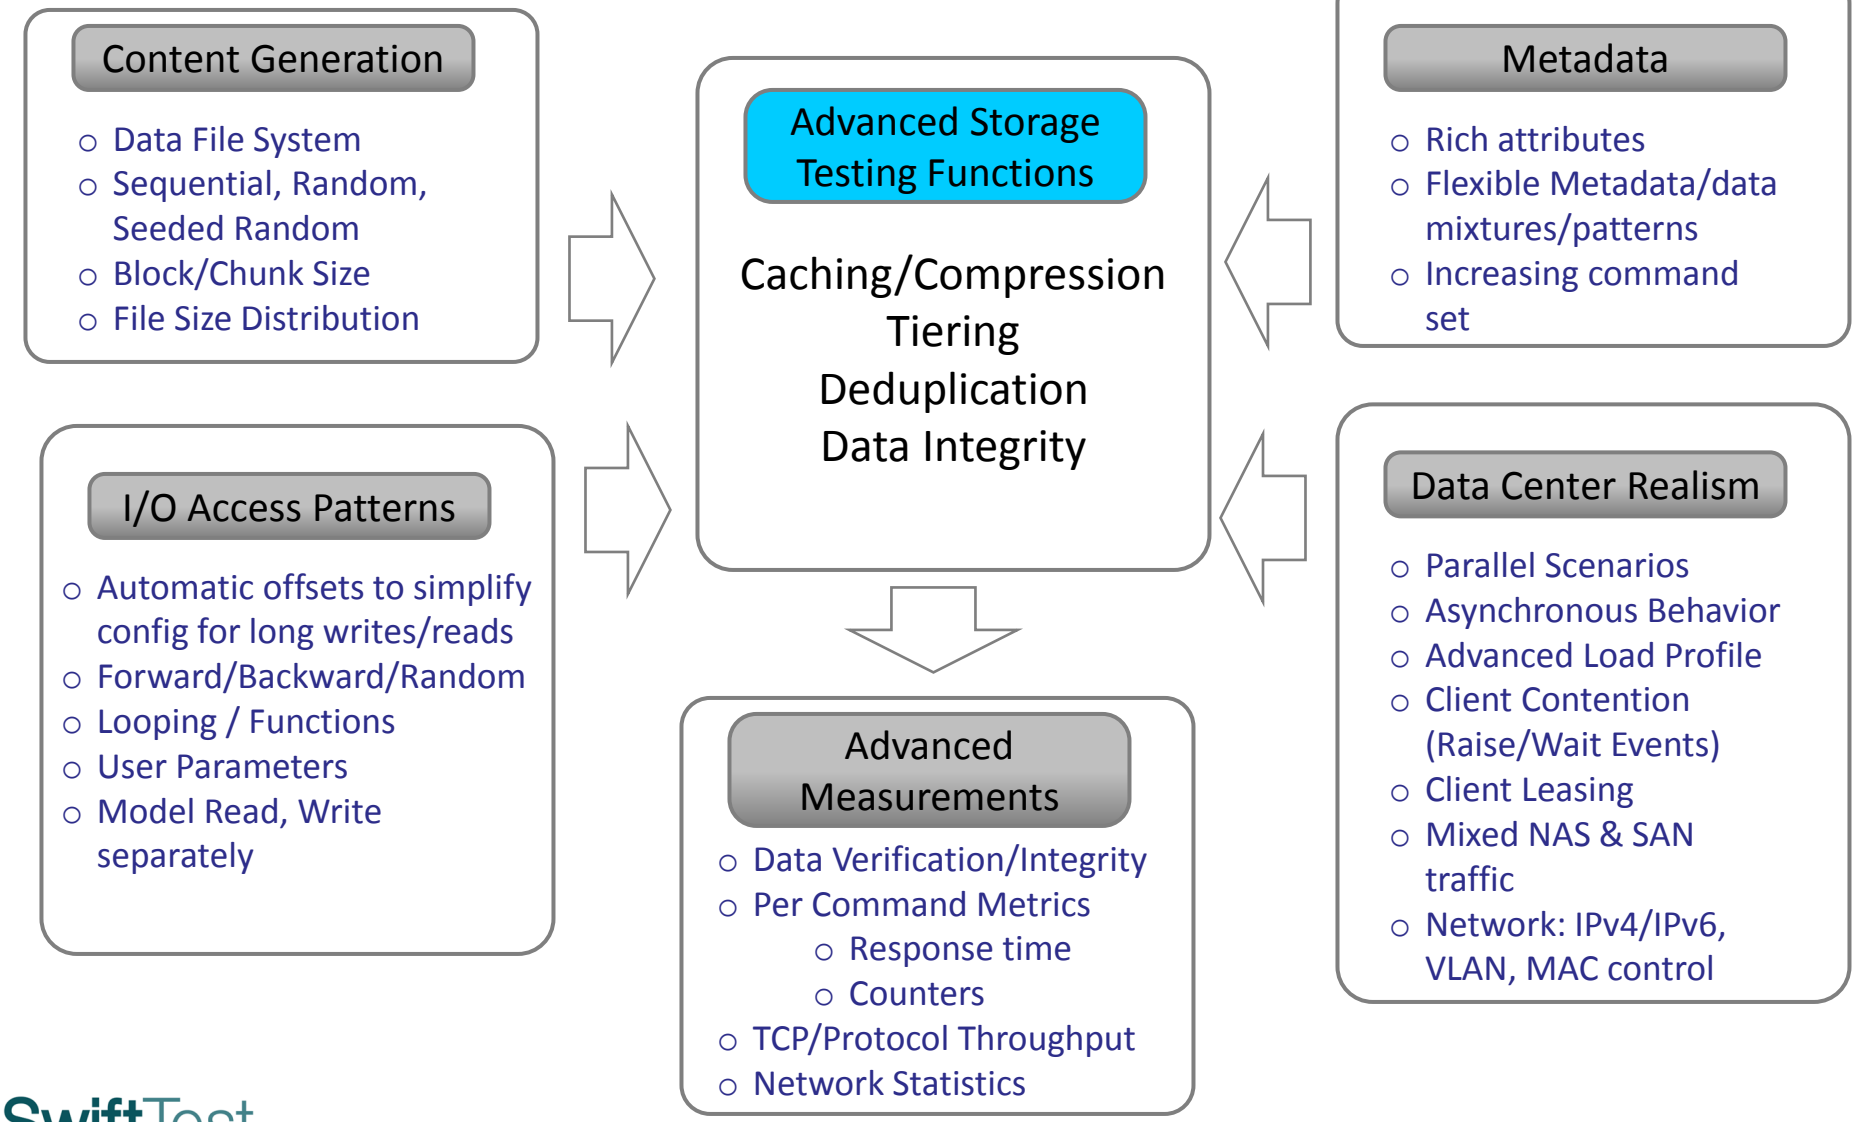

Simplifying & Automating the Workload Modeling Process

Workload Model

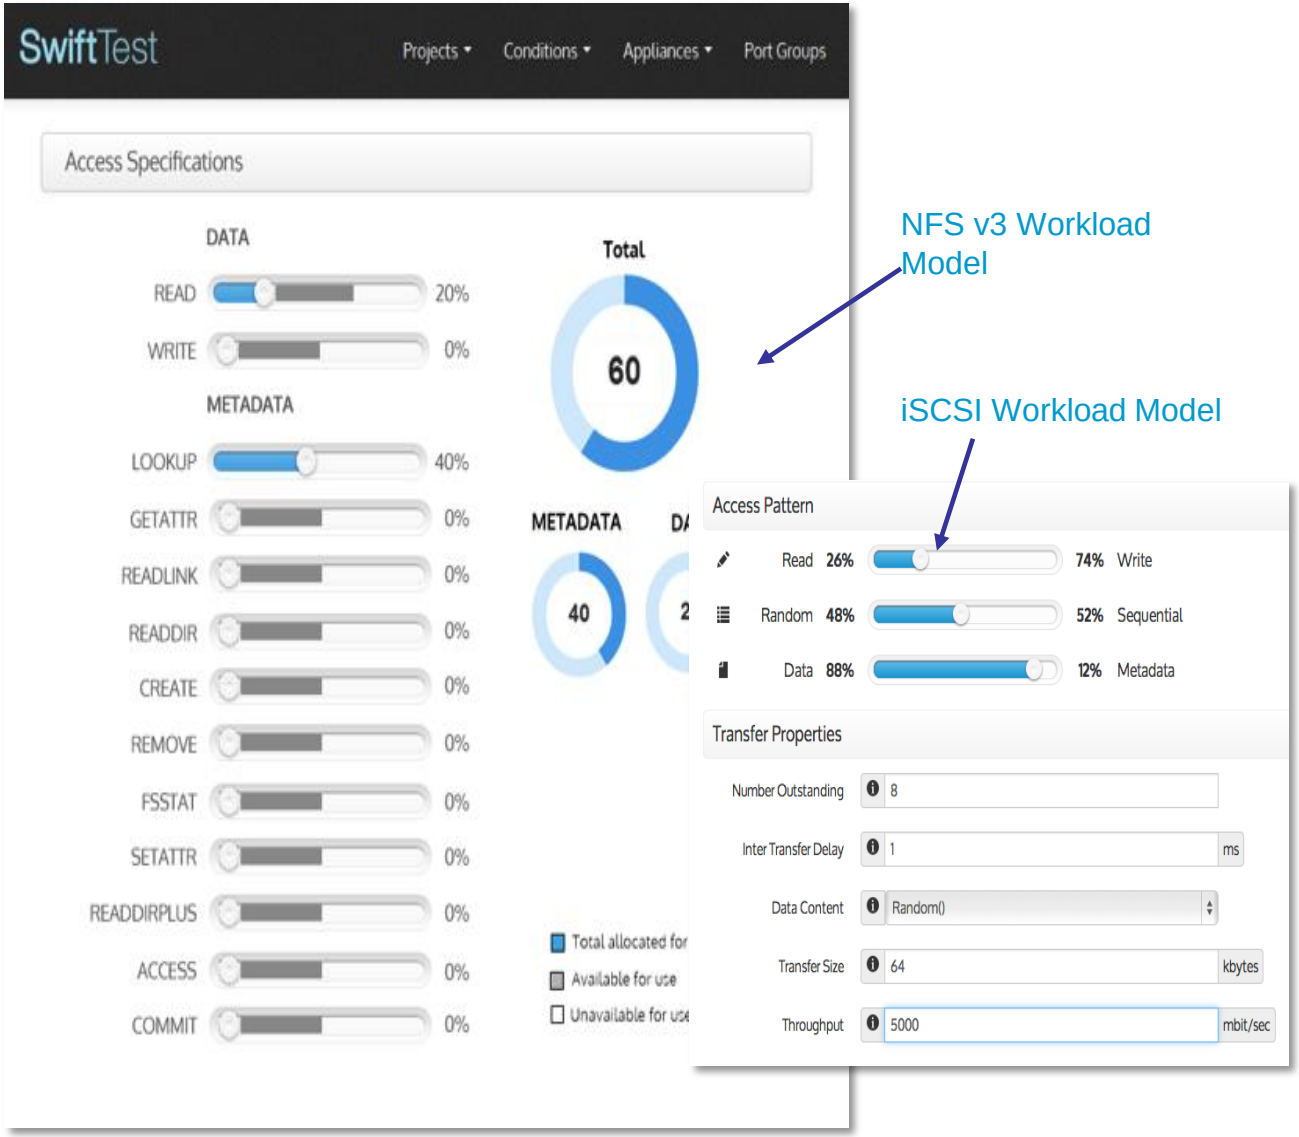

Measuring iSCSI IOPS by block size

iSCSI IOPS

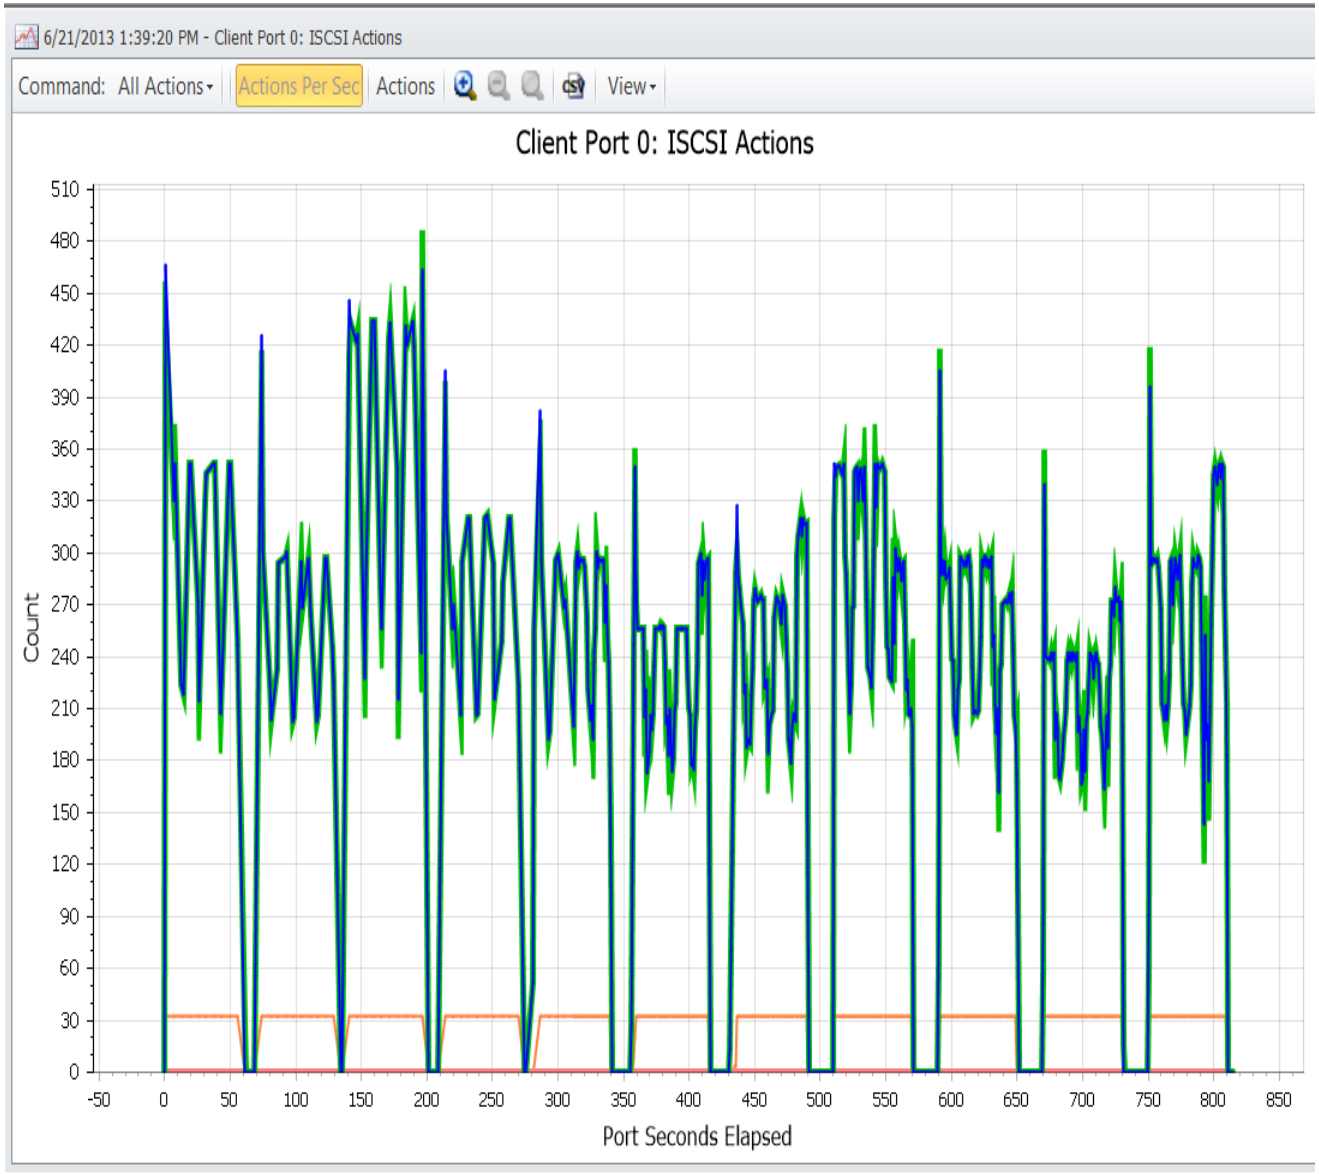

Measuring NFS Throughput

NFSv3 Throughput

Flash Memory SUMMIT

Benefits of Storage Performance Validation

Cost Optimization
Select most cost effective solution for your workload

Increased Reliability
Ensure uptime under load

Performance Assurance
Confidence in the scalability of new products

Engineering Productivity
More insight, less resources devoted to testing

SwiftTest

SwiftTest

Thank you!

- **Len Rosenthal**
- (408) 755-2043
- lrosenthal@swifttest.com