Infrastructure Performance Analysis

Peter Murray
Virtual Instruments

Hard Problems You're Trying to Understand

- How to most accurately test all solid state storage arrays
- Approaches for assessing storage performance
- How to select the best methodology for real application(s)
- Find proven strategies to size arrays and help avoid over-design and over-provisioning

Application Emulation

- The best way to test all solid state arrays is to emulate real applications
- Temporal locality
  - When data is written/read
- Spatial locality
  - Where data is written/read
- Data content patterns
  - Random or compressible
  - Pattern repetition
- Bursts
  - Bursts are present in applications

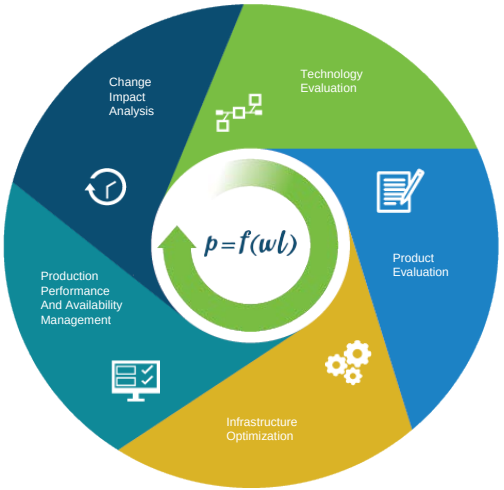

The Journey: How Did we Get Here?

- Early storage testing was oversimplified
  - Testing programs designed for disk drives
    - Did not represent actual applications
  - Could not emulate temporal or spatial locality
  - Did not emulate Data Content
  - Did not burst
- Difficult or impossible to:
  - Emulate varying load on many LUNs
  - Configure metadata and data structures required to emulate file-based applications

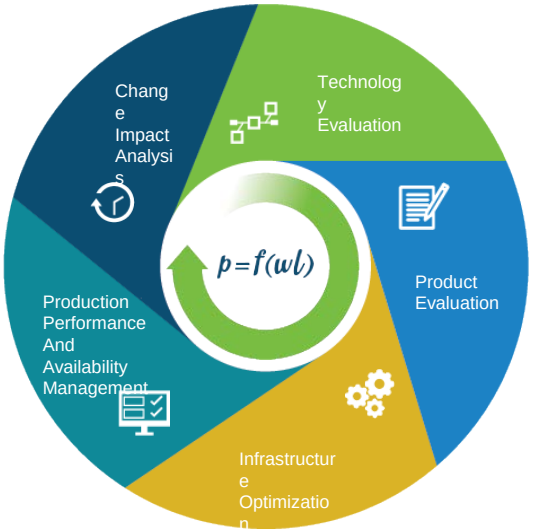

SS Arrays Require New Storage Testing Methods

- Applications exhibit spatial and temporal locality
  - Modern solid state arrays are designed with this in mind
- Much data content can be reduced
  - Data is random or compressible
  - Data can also be de-duplicated
  - All content types are present in most applications
- Some arrays must be tested with locality and content
  - Data reduction is a key feature - can't be turned off
- Data is sent in bursts
  - It's a feature, not a bug

The World has Changed

- Disk drives have not increased in performance in years
  - For years, drives were the primary network storage bottleneck
  - Processors, networks, architectures got better
  - Short-stroking and other techniques had limited effect
  - Data reduction caused unacceptable transaction times
- Solid state memory technologies change this model
  - "Reading is free"
    - Read access time is uniform
    - Unaffected by data location

Writing to Solid State Arrays

- Typically slower than reading
  - Write cycles are limited, so writing often performed at a page level
  - Inserting or appending existing data causes write amplification
  - So, modern solid state storage arrays avoid writing
  - Some always append, then reclaim stale memory pages
- Flash write access time is implementation dependent
  - Sequential writing speed may be impacted
  - Random writing can impact garbage collection
- Data reduction may require post-processing
- Sustained writes may affect write performance over time

Writing is Hard

“Bursty” Writes are Inherent to Application Traffic

- All application traffic is sent in bursts
  - Clients send at full rate - transactions are not metered
  - They occur during both quiet and busy periods
    - Applications complete transactions as quickly as possible
- Test tools that send metered traffic are worse than unhelpful
  - Testing that does not include bursts:
    - Cannot demonstrate how a storage array will behave in production
    - Should not be used to evaluate storage arrays
- More information: http://tinyurl.com/zg5x3dk

How is Flash Different?

- Addressable storage space usually less than raw space
  - May help avoid performance issues during garbage collection
  - Other methods are available to avoid performance issues
  - Can help increase flash life
- Dedupe & compression decrease app storage requirements
  - More storage per nominal byte
  - But may impact performance

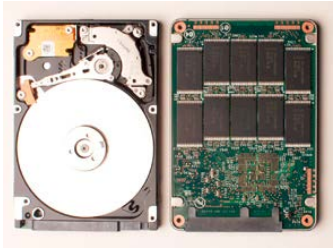

How Else is Flash Different?

- Metadata processing makes it harder to fill an array
  - Testing at near full capacity required to understand array performance
  - May require using multiple workloads or scaling workload very high
- Garbage collection or metadata processing may affect performance
- Testing with hotspots crucial to modeling application behavior
- Software services & protocols often run differently on SSD than on HDD

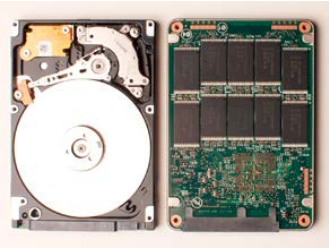

Storage Performance

- Vendor marketing have good stories
  - But don't confuse marketing with reality
- Vendors endorse performance testing with your workloads, derived from production environments, via synthetic workloads
- Vendors and standards organizations produce benchmarks, but they are guidelines at best
- Benchmarks don't offer configuration guidance – and don't represent your workloads

Typical Performance Testing Questions

- Which is the best storage technology for a specific requirement?
  - All Solid State (AFA, NVMe)?
  - Hyper-converged array?
  - Hybrid array?
  - All-disk array?
  - Private or public cloud?
- Which is the best vendor / product for a specific requirement?
- What is the optimal configuration for a specific requirement?
- Does performance degrade with enterprise features enabled?
  - Deduplication
  - Compression
  - Snapshots, Clones, Replication

Other Performance Testing Questions

- What are the performance limits of a particular configuration?
  - How will an application grow over time?
- How does an array behave when it reaches its performance limit?
- Does array performance degrade over time?

Traditional Storage Testing Approaches

- Limits finding
- Functional testing
- Error Injection
- Soak testing

Storage Performance Validation

2 core methodologies

Workload Modeling
Simulate the I/O profiles of your production environment

Performance Profiling
Fully characterize performance of arrays under wide variety of load parameters

Performance Profiling

- Performance Profiling
- Characterization under a wide range of workload conditions
- Understand sweet spots and weaknesses of an array
- Sometimes referred to as “4 corners” or “limits” testing, but you can do much more than that
- Vendors need these tests to validate portions of a storage array
- Applications don’t act like performance profiles
  - Some exceptions; e.g. block sizes, queue depth, outstanding commands
- Extraordinarily difficult with free tools

Performance Profiling

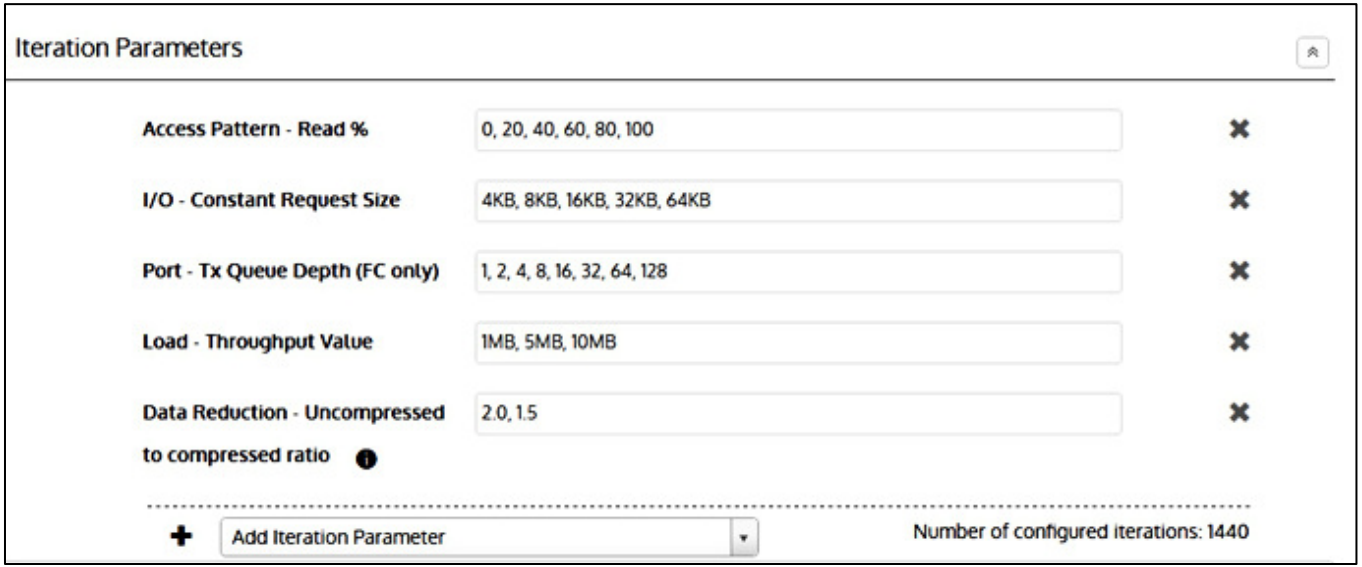

Performance Profiling
(continued)

Fibre channel performance — Finished

Started by admin

1:00:05:15 — 00:00

Last Log Record: 2015-01-26 11:38:35 AM | Success | Test Suite finished — Logs

Iteration Results Export to CSV

#	Status	Duration	Access Pattern - Read %	I/O - Constant Request Size	Port - Tx Queue Depth (FC only)	Load - Throughput Value	Data Reduction - Uncompressed to compressed ratio	SCSI Throughput (average)	SCSI IOs Succeeded/sec (average)	SCSI Average Response/Latency Time (average)
48	Finished	01:01	0	4KB	128	10MB	1.5	8.3 MB/sec	2115.387	6 ms
47	Finished	01:00	0	4KB	128	10MB	2	8.0 MB/sec	2044.602	.7 ms
42	Finished	01:00	0	4KB	64	10MB	1.5	7.5 MB/sec	1921.051	.5 ms
41	Finished	01:00	0	4KB	64	10MB	2	7.2 MB/sec	1837.487	.9 ms
36	Finished	01:00	0	4KB	32	10MB	1.5	6.5 MB/sec	1663.073	.3 ms
288	Finished	01:00	20	4KB	128	10MB	1.5	6.5 MB/sec	1657.239	.5 ms
35	Finished	01:00	0	4KB	32	10MB	2	6.3 MB/sec	1612.252	.5 ms
282	Finished	01:00	20	4KB	64	10MB	1.5	6.2 MB/sec	1586.806	.8 ms
281	Finished	01:00	20	4KB	64	10MB	2	6.1 MB/sec	1554.01	.1 ms
287	Finished	01:01	20	4KB	128	10MB	2	6.1 MB/sec	1545.593	.7 ms

Workload Modeling

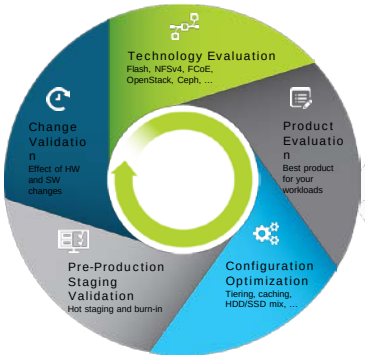

Workload Modeling
Simulate the I/O profiles of your production environment

Performance Profiling
Fully characterize performance of arrays under wide variety of load parameters

Workload Modeling

- Stresses an array using a realistic simulation of specific production workload/s
  - For IT customers, from your current environment
  - For vendors, using customer examples or "dog food"

- Realisic I/O profiles are the most accurate way to test
- Packet traces offer limited utility in testing
  - Huge volume of data
  - Short duration
  - Security concerns

Workload Modeling

(continued)

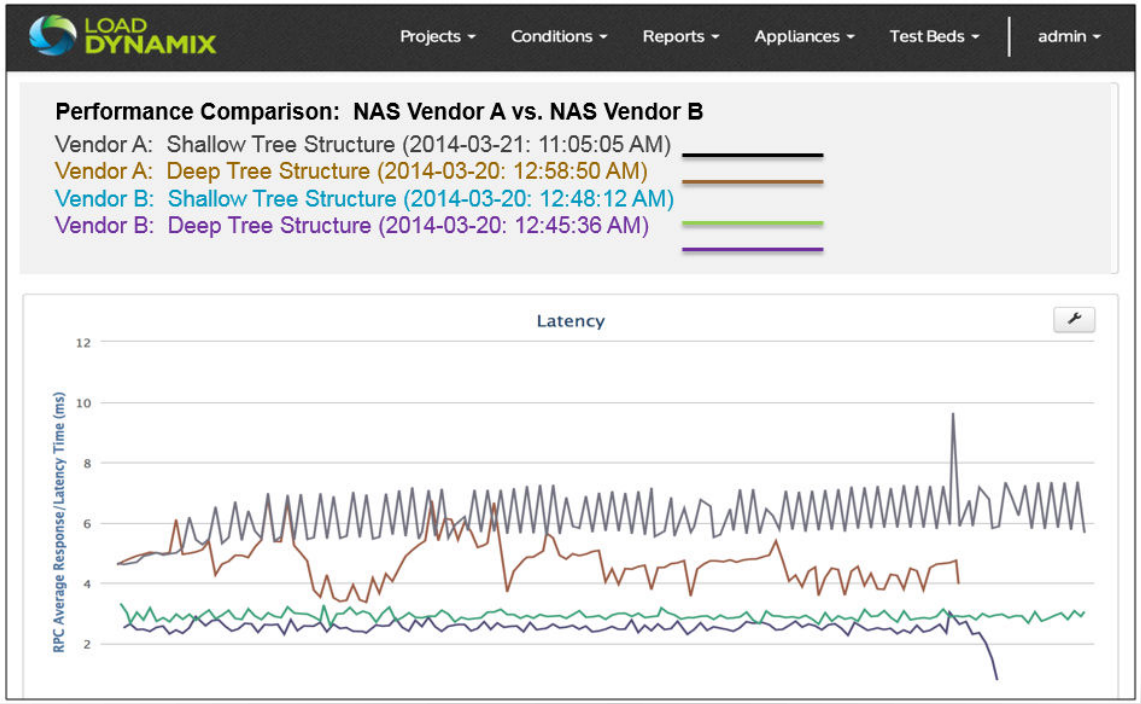

Where Does Workload Modeling Come From?

- Customers ask for workload models
  - IT customers want models of their workloads
  - Vendors want "the" workload
    - Oracle, Exchange, etc.
- IT customers ask to help make better decisions about:
  - Upgrading storage hardware or software
  - Changing storage network configuration
- Vendors ask for help to:
  - Test customer examples/issues
  - Find realistic scaling limits to test app growth over time

Result: A New Modeling Method

- Cloud-based workload modeling
- Community-based workload sharing
- Workload model that be ingested into Virtual Networks load generation
- More realistic and scalable than benchmarks

Workload Central Beta

- WorkloadCentral is a free cloud-based analytics platform and community that allows you to understand analyze, create and share workloads.
- Available at: www.workloadcentral.com
- Key Features:
  - Free workload analysis & creation
  - Advanced workload analytics
  - Workloads for validation, testing & benchmarking
  - Workload Library, community & discussion

Uploading Your Workload Data

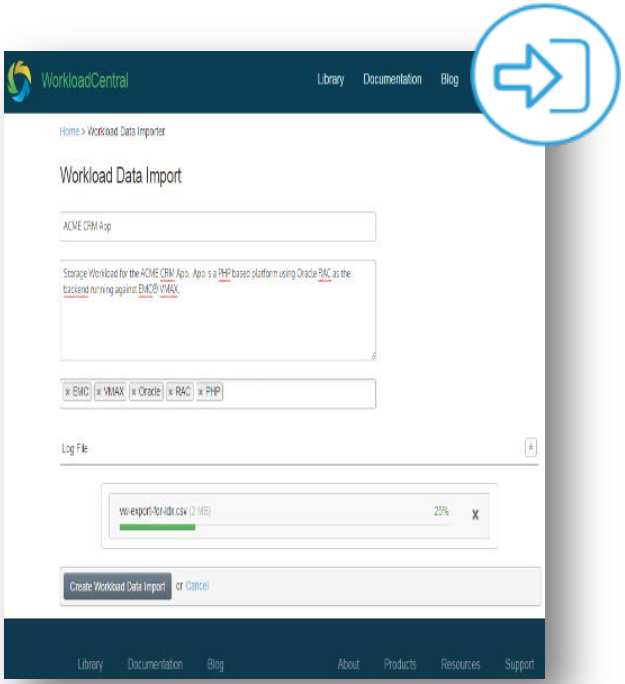

- The Workload Importer offers:
- Ability to upload data from any vendor or environment
- Out of the box import policies
- Analysis policies provide flexibility to define different workloads

Visualizing Your Data with the Workload Analyzer

- A free downloadable, printable report and dashboard that provides:
  - Workload access pattern
  - Workload behavior characteristics
  - Workload performance
  - Workload creation

Running a Block-Based Workload Model

Workload Modeling

Simulate the I/O profiles of your production environment

1. Characterize Workload I/O

- Per-LUN I/O:
  - Read-Write Mix
  - Random or sequential access
  - Hot spots and hot spot drift
- Data Content
  - Randomness
  - Compressibility
  - Unique vs. duplicated blocks

2. Determine Data Content Patterns

- Consist of repeating and non-repeating patterns
  - Random or compressible
- Consist of varying pattern lengths
- Data content patterns
  - Create during preconditioning
- Data content streams
  - Create during preconditioning
  - Replay during testing

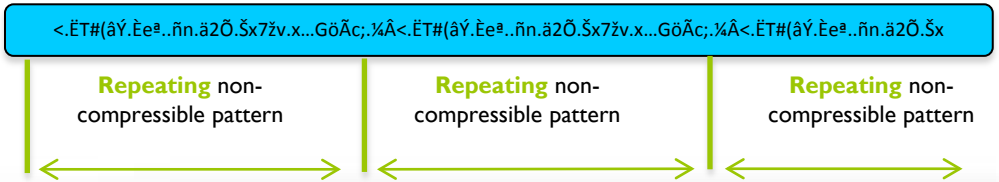

3. Build I/O Models

- Decide when to model
- Boot storm
- Everyday office load
- Backups
- End of period processing
  - Month, Quarter, year end
- Test primary models individually
- Test periodic models on top of everyday load
- Magnify load to test expected maximums

4. Run Workload Models

- Run most common model(s) first
  - Bootstorm
  - ”Normal” daily workload
  - Daily backup processing
  - Shutdown
  - Use to determine baseline performance
- Add periodic models to common model
  - E.g. end of period processing
- Combine apps if appropriate and test together

5. Test Array Features

- Test effects of MPIO
- Test effect of maintenance/other management
  - Snapshots, clones, replication, etc.
- Test at or near full capacity
- Test effect of QoS

Test in an Iterative Manner

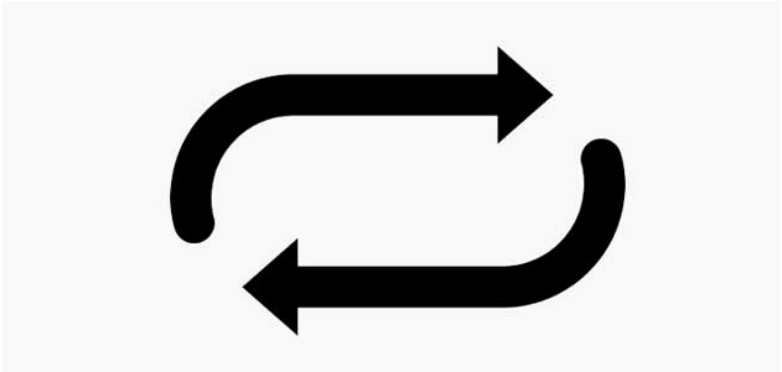

- Run
- Analyze
- Repeat as necessary
  - Change testing to reflect expected business conditions

Summary

Benefits

- Performance assurance
- Reduced storage costs
- Increased uptime
- Acceleration of new application deployments

Summary

- Application Testing is now mandatory
  - Black art has become repeatable
- Testing with bursts is mandatory
- No synthetic workload is perfect
  - But is the best approach available
  - This will only improve over time
- Customers can see:
  - How closely the model emulates apps
  - A realistic view of how an array operates
- This new model is changing storage testing