AFA Storage Performance Testing and Validation Methodology

Peter Murray

Load DynamiX

Agenda

- Introduction
- Load DynamiX Testing Methodologies
- Performance Profiling
- Workload Modeling
- Case Study
- Summary

Introduction:
Flash Memory and All-Flash Arrays

- **Flash memory tiers in storage arrays is mainstream**
- **All-flash Arrays (AFA) are moving to mainstream**
- **Advanced AFAs are appropriate for Tier 1 applications today**
  - Appropriate for some combined applications
- **Advanced AFAs implement new and unique features**
- **Testing arrays with advanced features enabled is required**
  - Testing these arrays, in particular, requires a new approach

How is Flash Unique?

- **Addressable storage space may be less than raw space**
  - Aids SSD wear leveling
  - Increases flash life
- **Deduplication and compression may increase available space**
  - More storage per nominal byte
- **Advanced metadata processing makes it hard to saturate array**
  - Array must be tested at near full capacity to understand performance
- **Testing with hotspots helps model application behavior**
  - Vendor garbage collection may affect performance
- **Software services – software runs differently than on HDD**

Why Performance Testing is Important
Typical AFA Questions

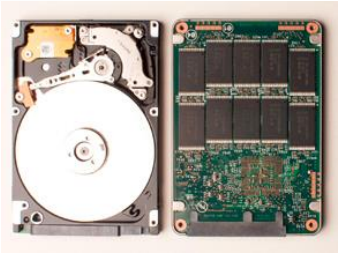

- **Which workloads run best on AFAs?**
- **Which is the best vendor / product for my needs?**
- **What is the optimal configuration for my array?**
- **How much does performance degrade with enterprise features:**
  - Deduplication
  - Compression
  - Snapshots, Clones, Replication
- **Where are the performance limits of a potential configuration?**
- **How does an AFA behave when it reaches its performance limits?**
- **Does performance degrade over time?**
- **Which workloads are best for an AFA? A hybrid storage array?**

Agenda

- Introduction
- **Load DynamiX Testing Methodologies**
- Performance Profiling
- Workload Modeling
- Case Study
- Summary

Load DynamiX Testing Methodologies

Load DynamiX Enterprise software

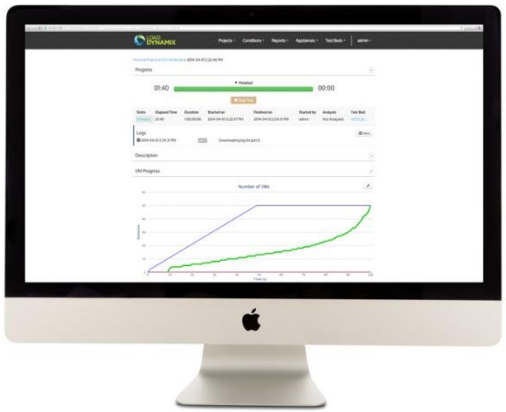

Workload Modeling & Performance Profiling

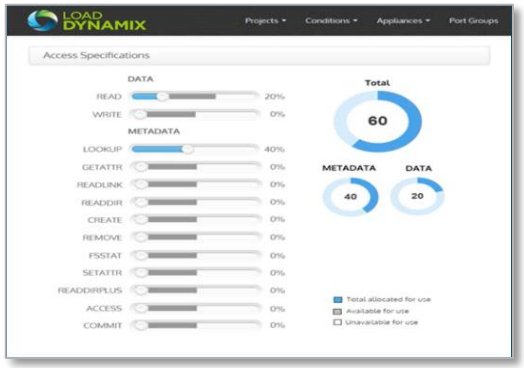

Performance Analytics

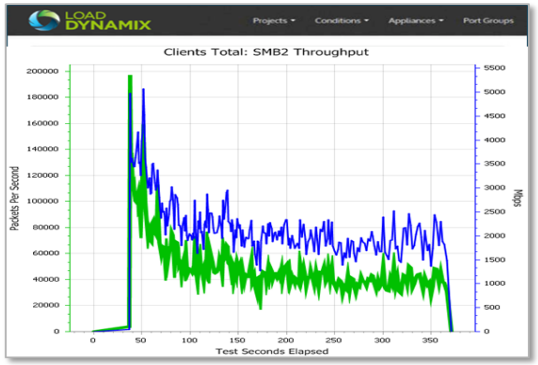

Load DynamiX Performance Validation Appliance

Product or Configuration A

File, Block, or Object storage

Product or Configuration B

Storage Performance Validation:
2 core methodologies

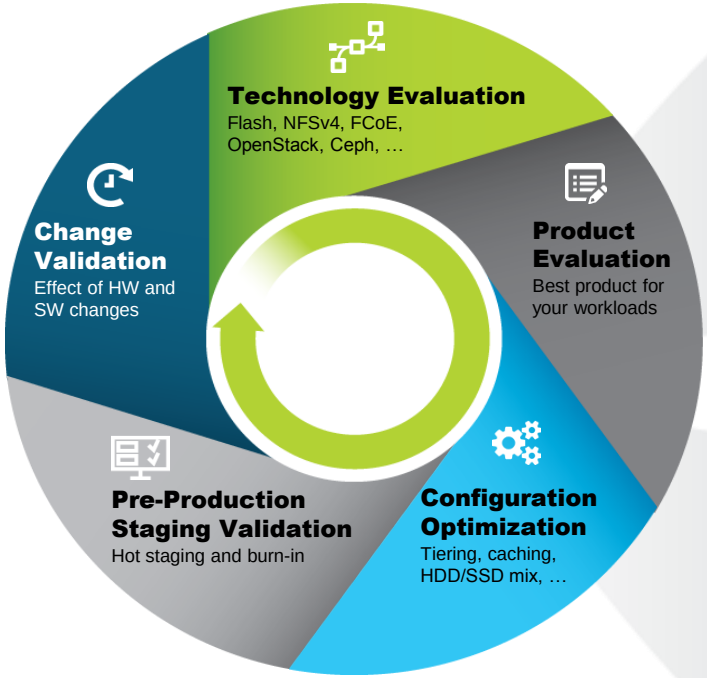

Workload Modeling
Simulate the I/O profiles of a production environment

Performance Profiling
Fully characterize performance of an array under wide variety of load parameters

Typical Storage Testing Approaches

- **Limits finding** – determining the workload conditions that drive performance below minimal thresholds, and the documenting of storage behavior at failure point
- **Enterprise feature functional testing** – the investigation under simulated load of various functions of the storage system (e.g. snapshots, clones, replication, backup, etc.)
- **Error Injection** – the investigation under simulated load of specific failure scenarios (e.g., fail-over when an array controller or individual drive fails)
- **Soak testing** – the observation of the storage system under load sustained over significant time (e.g. 2 days, 1 week)

Agenda

- Introduction
- Load DynamiX Testing Methodologies
- **Performance Profiling**
- Workload Modeling
- Case Study
- Summary

Performance Profiling

- Performance characterization using a wide range of conditions
- Demonstrates array-level sweet spots, bottlenecks
- Enabled by LDX-E Iteration engine

Example Test Configuration

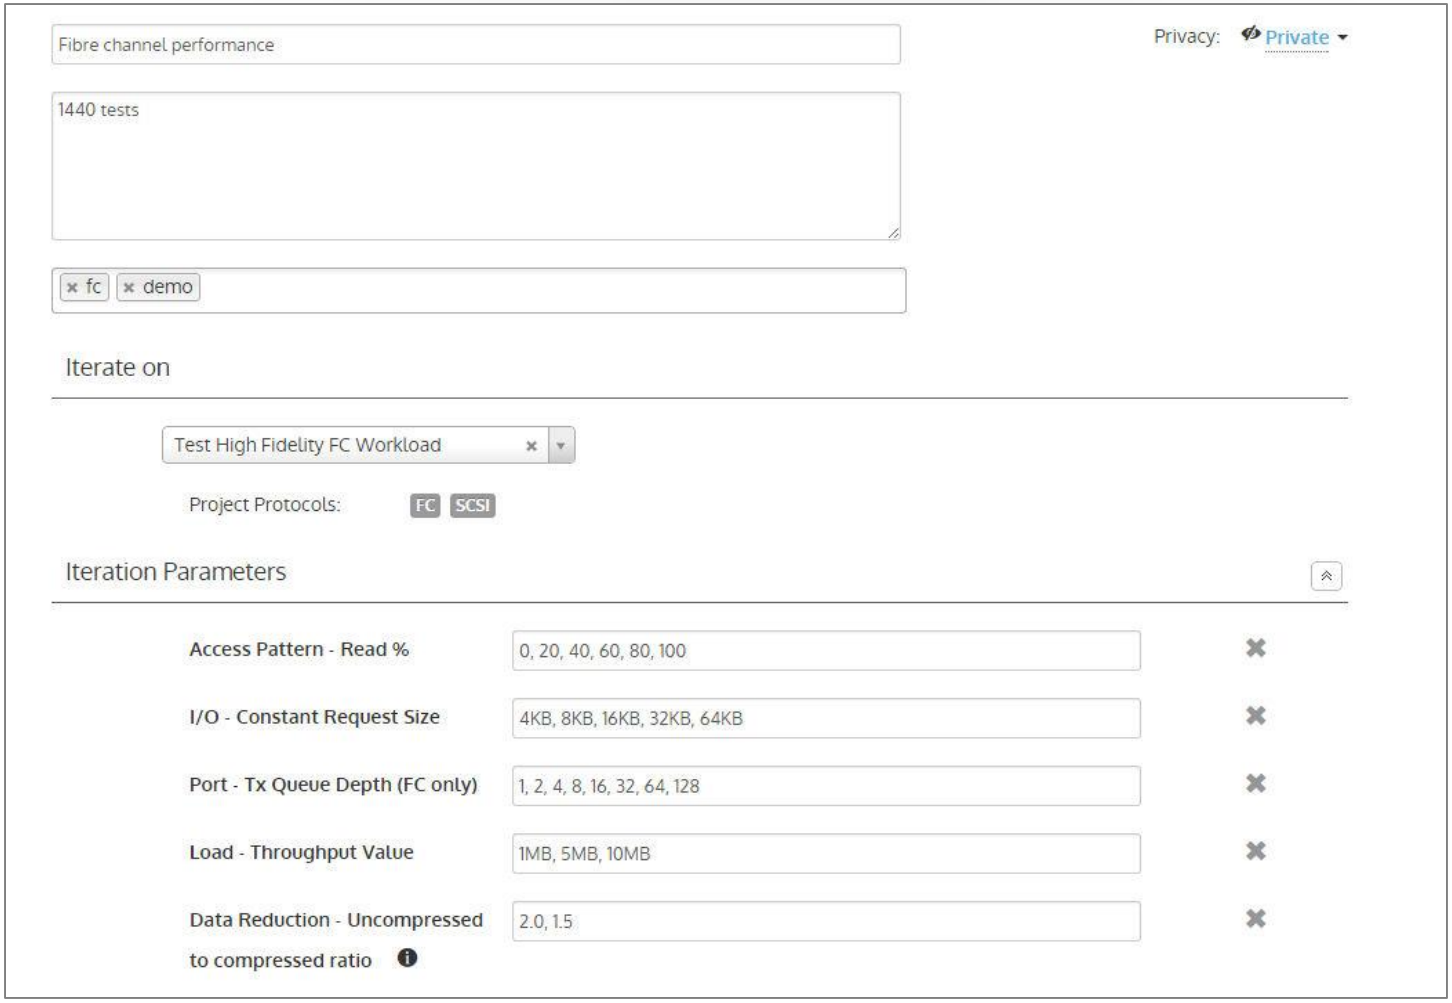

Example Results: Multiple Test Runs

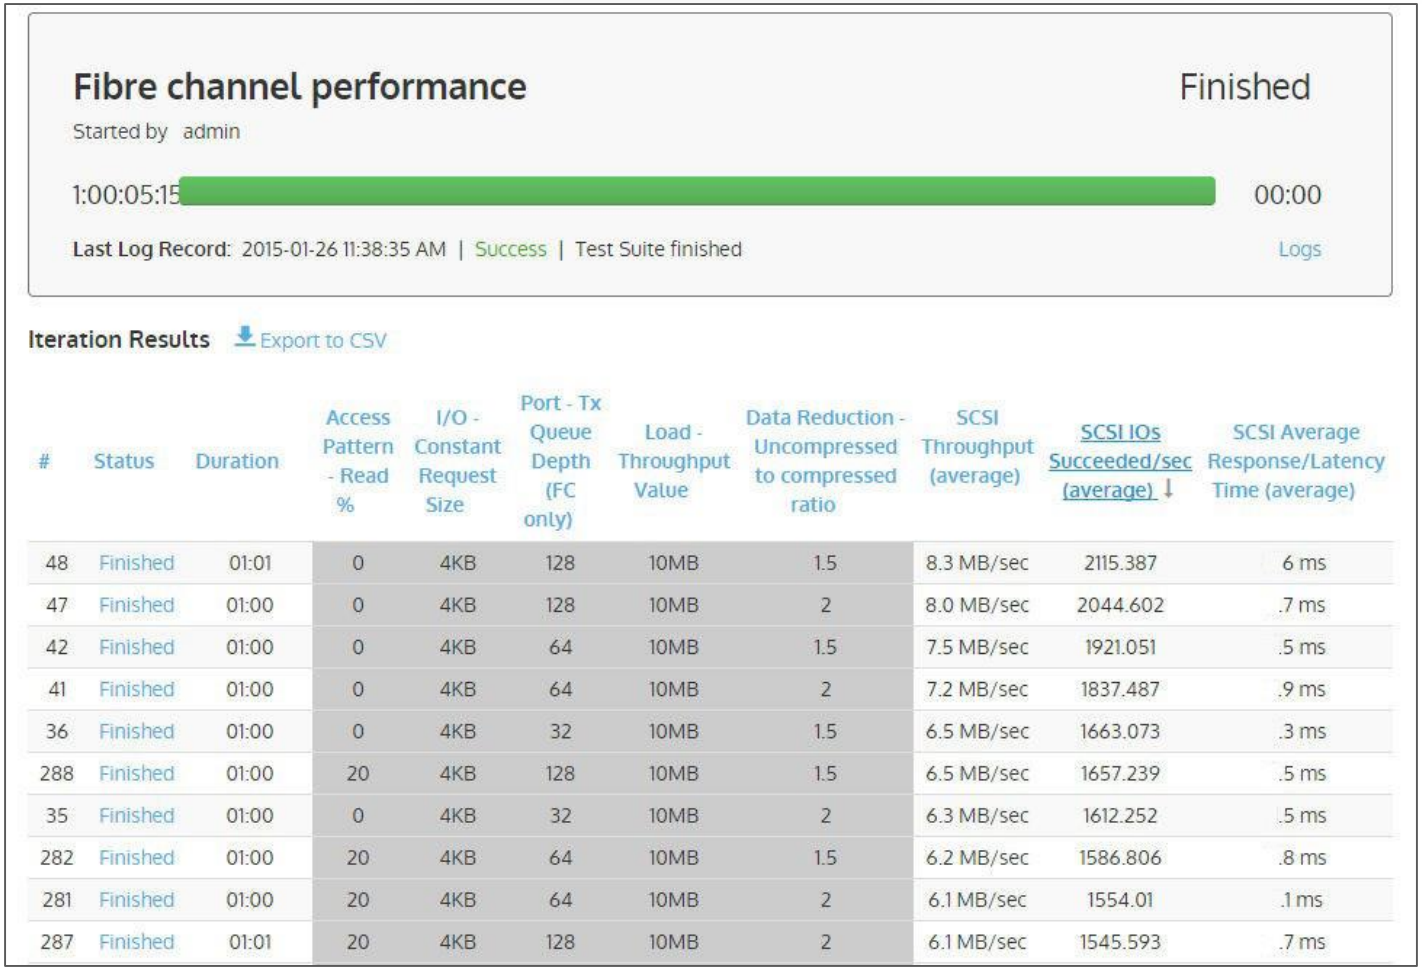

Fibre channel performance — Finished

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

Agenda

- Introduction
- Load DynamiX Testing Methodologies
- Performance Profiling
- **Workload Modeling**
- Case Study
- Summary

Workload Modeling

- Stressing an array using a realistic simulation of one or more specific production workloads
- Realism is paramount
- Enabled by the LDX-E Composite Workload function

Workload Modeling Example

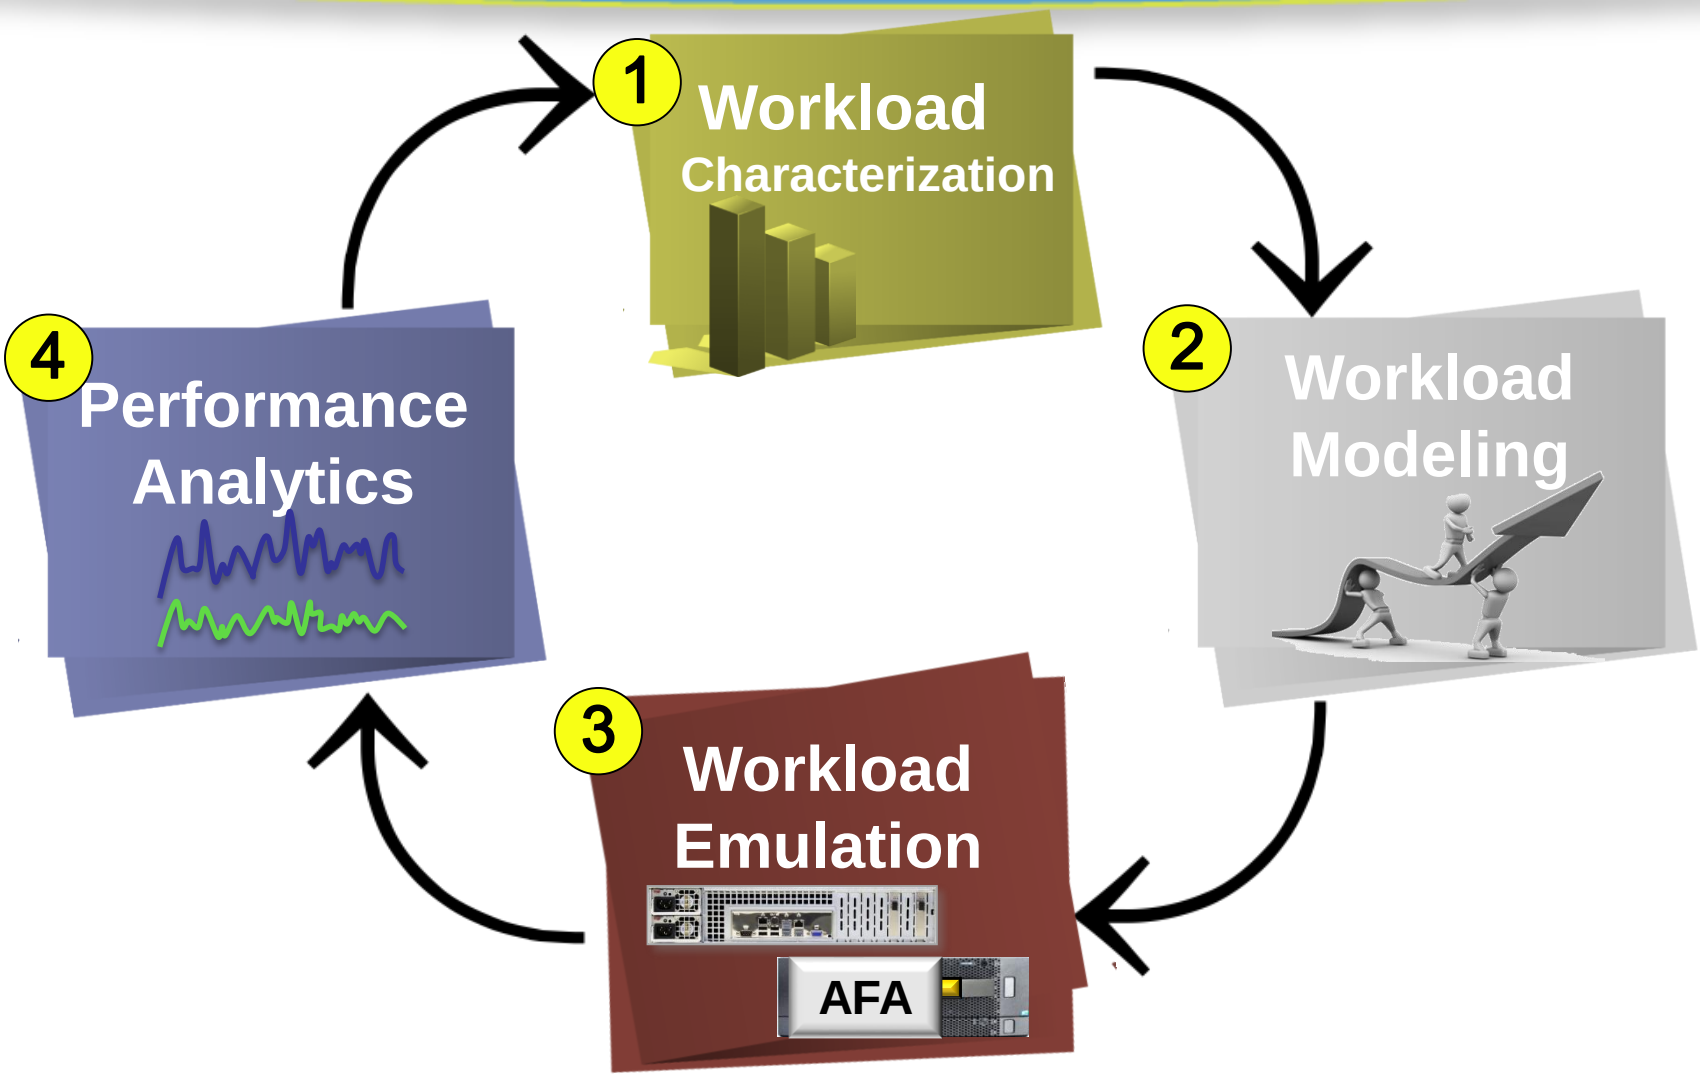

1: Characterize a Workload Model

Example Vendor Workload Statistics

Name	Host IOs/sec	% Writes	%Reads	Avg I/O Size	Capacity (GB)	%RR	%SR	%RW	%SW
dbf1	522.1	0.4	99.6	19	256	95	4	0	0
dbf2	448.5	0.1	99.9	16	256	94	6	0	0
dbf3	316.6	5.2	94.8	19	256	84	11	5	0
dbf4	297	0.9	99.1	29	100	99	0	1	0
dbf5	235.6	4.8	95.2	17	256	84	11	5	0
dbf6	220.2	5.6	94.4	20	256	84	11	5	0
dbf7	165.7	5.1	94.9	19	200	91	3	5	0
dbf8	91.9	6.2	93.8	17	100	82	11	6	0
dbf9	90.3	27.7	72.3	48	200	73	1	26	0
dbf10	7.6	17.8	82.2	105	256	81	1	18	1
dbf11	201.4	1.5	98.5	237	256	94	4	1	0
redo1	70.2	96.9	3.1	28	32	3	0	88	9
redo2	68.1	99.6	0.4	14	32	0	0	90	9
quest	6.3	88.5	11.5	13	10	9	2	86	4
arc	2.8	98.4	1.6	347	256	0	0	93	7
oraex	1.5	17.7	82.3	2	33	82	0	17	2
dbf	**2395.5**	7.38	92.62	**30.90**	213.60	**87**	**6**	**7**	**0**
dbf11	**201.4**	1.5	98.5	**237**	256	**94**	**4**	**1**	**0**
redo	**138.3**	98.25	1.75	**21**	32	**2**	**0**	**89**	**9**
other	**10.6**	68.2	31.8	**120.7**	99.7	**30**	**1**	**65**	**4**

Determining Data Content Patterns

- **Data content patterns**
  - Created before testing
- **Data content streams**
  - Written during testing
- **Repeating and non-repeating patterns**
  - Random
  - Compressible
- **Varying pattern lengths**

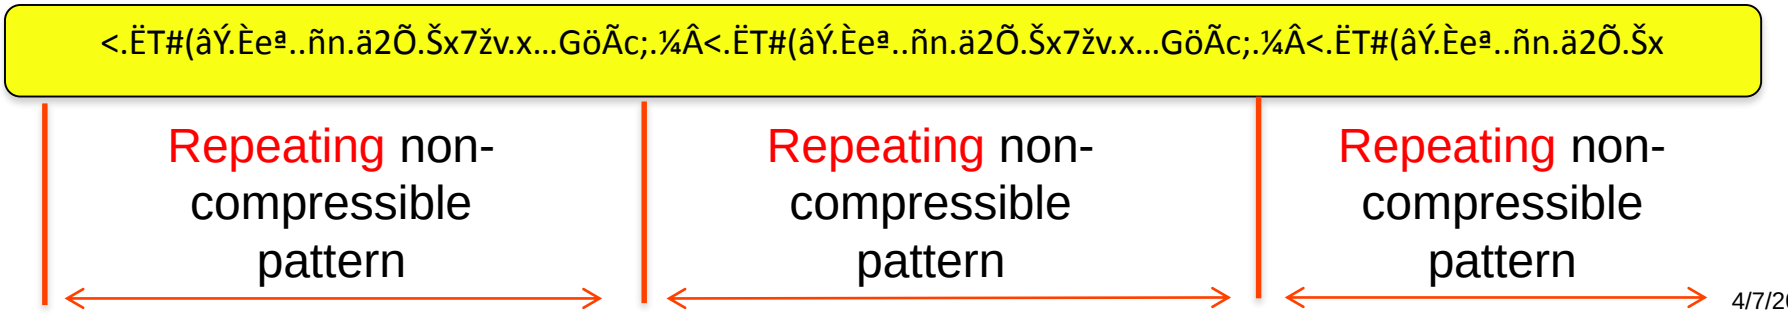

4/7/2015

18

2: Create a Workload Model

PRE-BUILT TEST SUITES
(VDI, etc)

PRODUCTION STATS
(Perfstats, .nar, .btp, NFSstat, etc)

ACCURATE, REALISTIC WORKLOAD MODEL

Example Test Bed

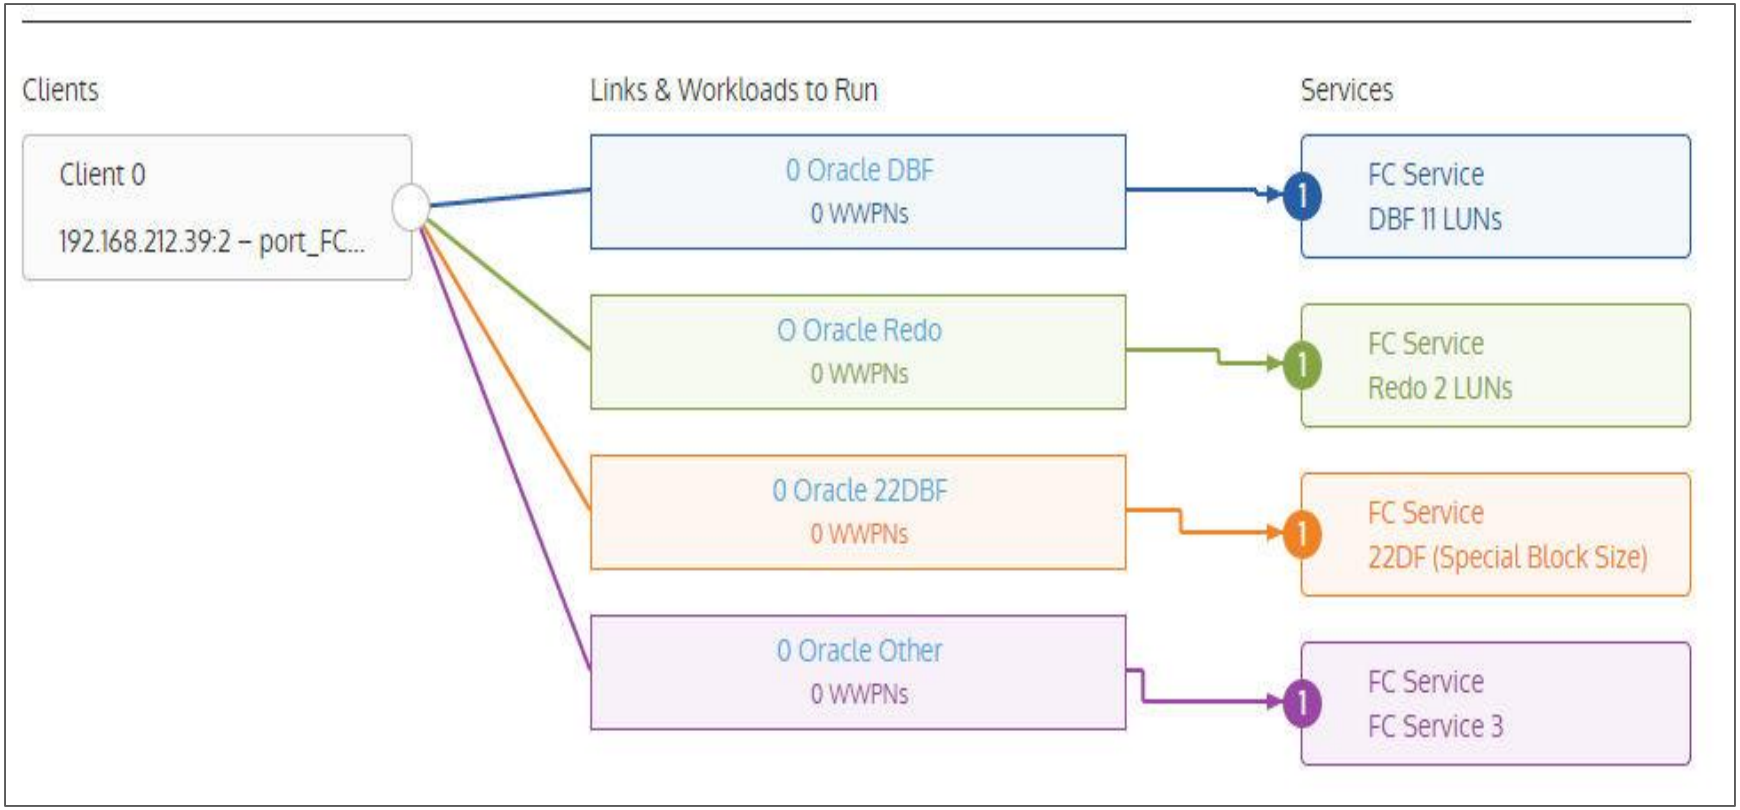

Example Workload Test Components

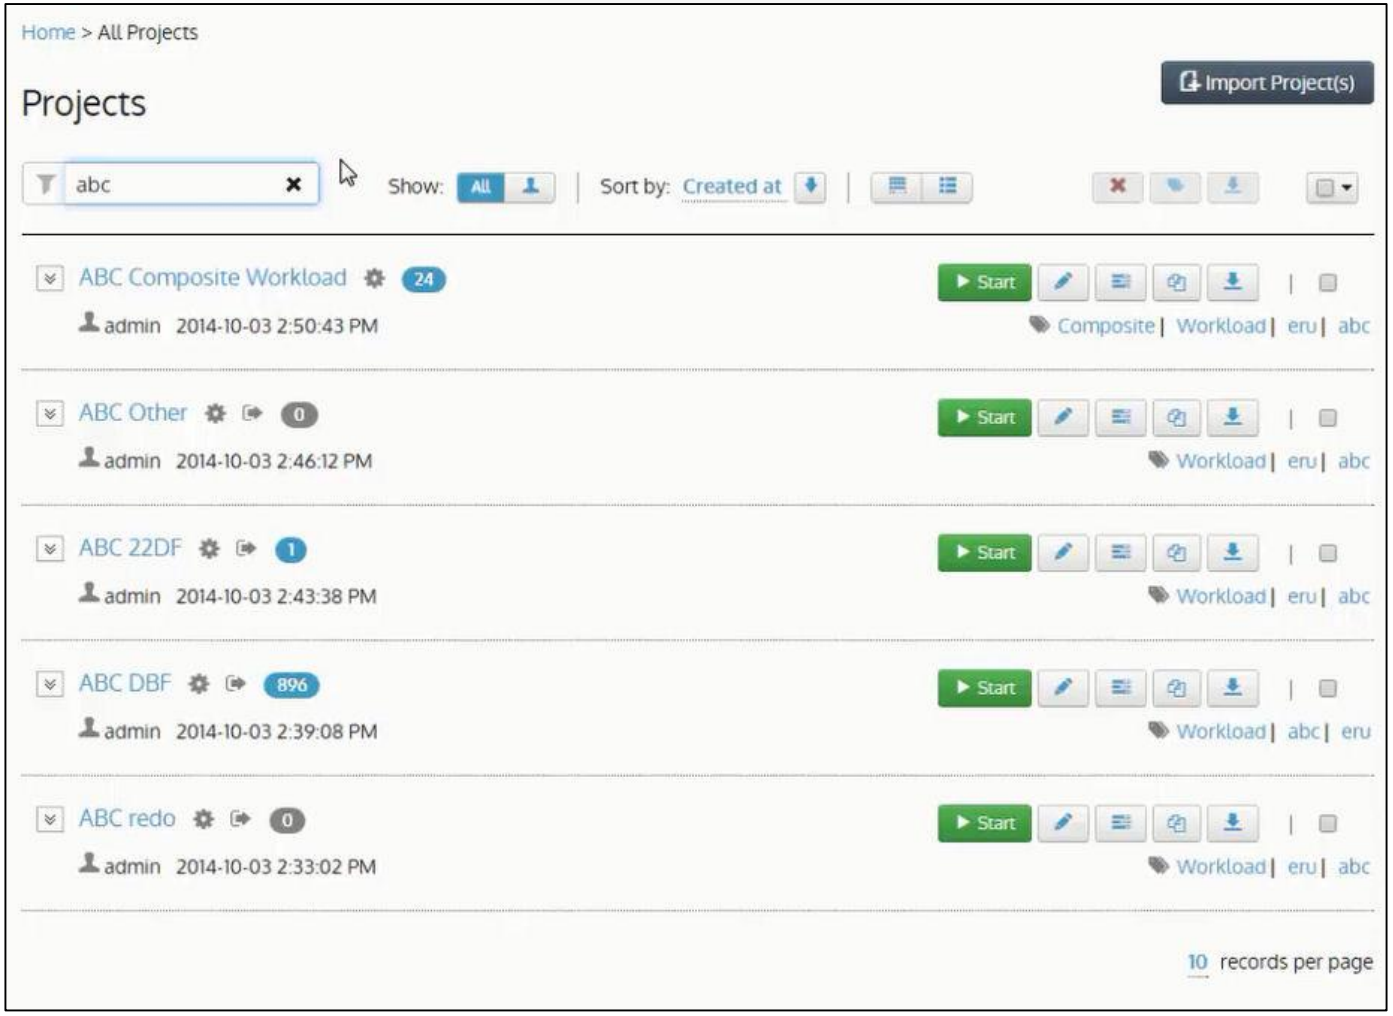

Example Test Configuration #1

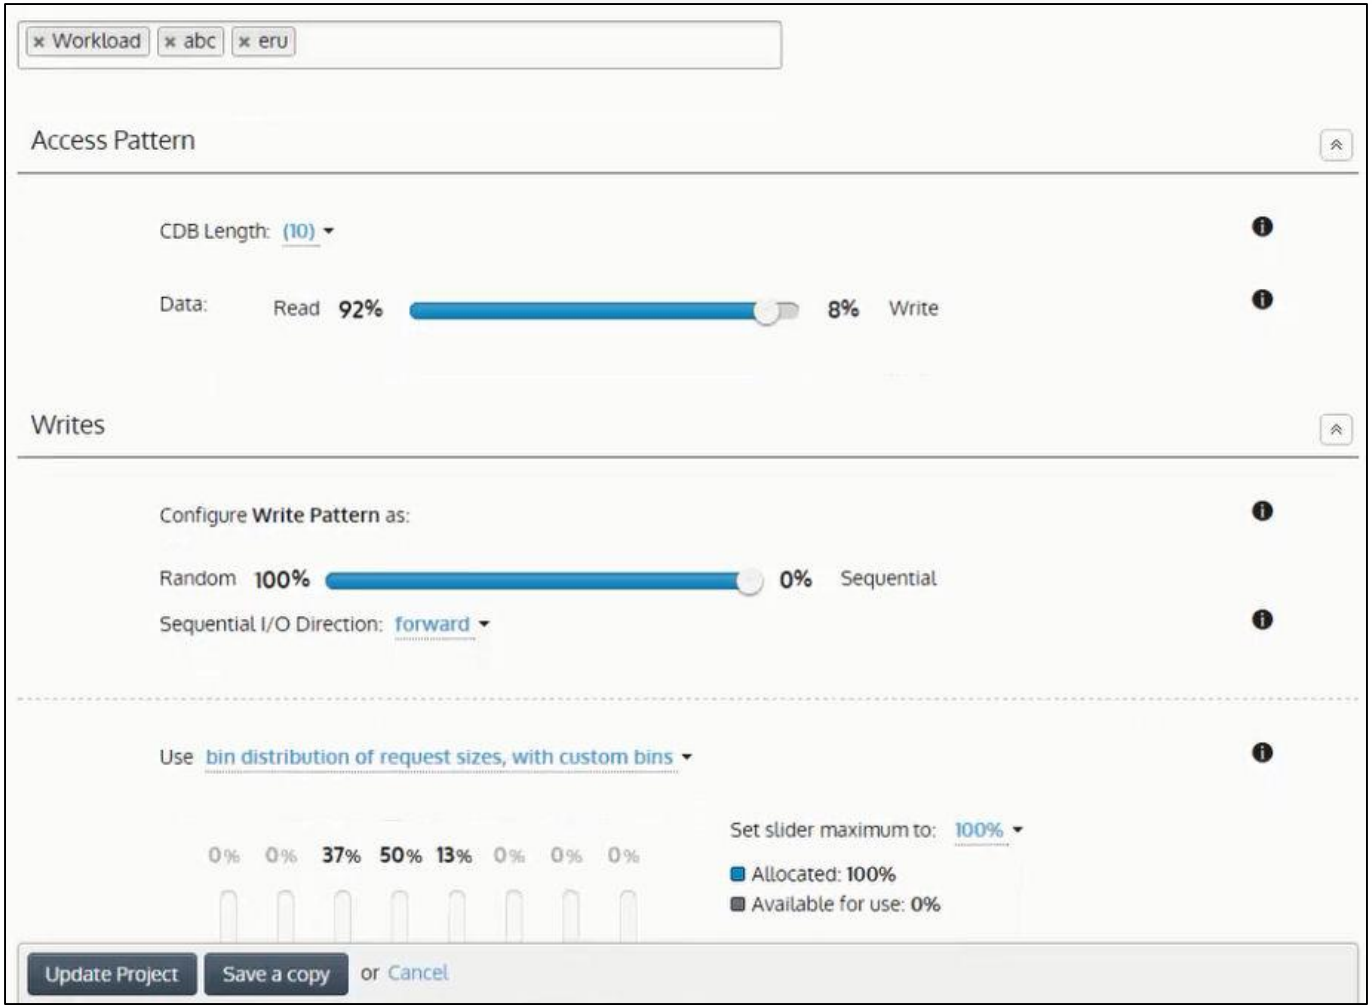

Example Data Reduction Configuration

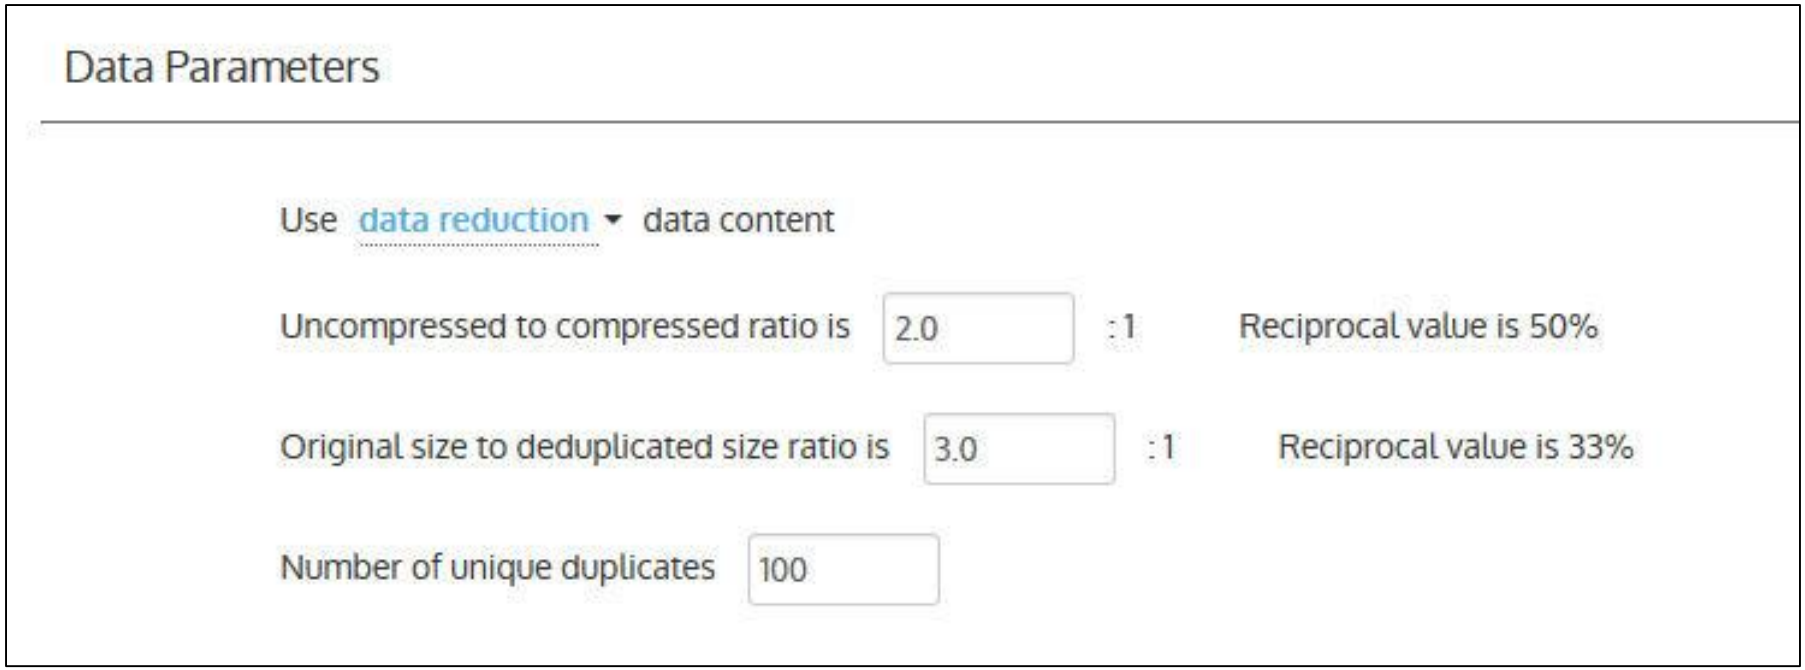

Example Test Configuration #2

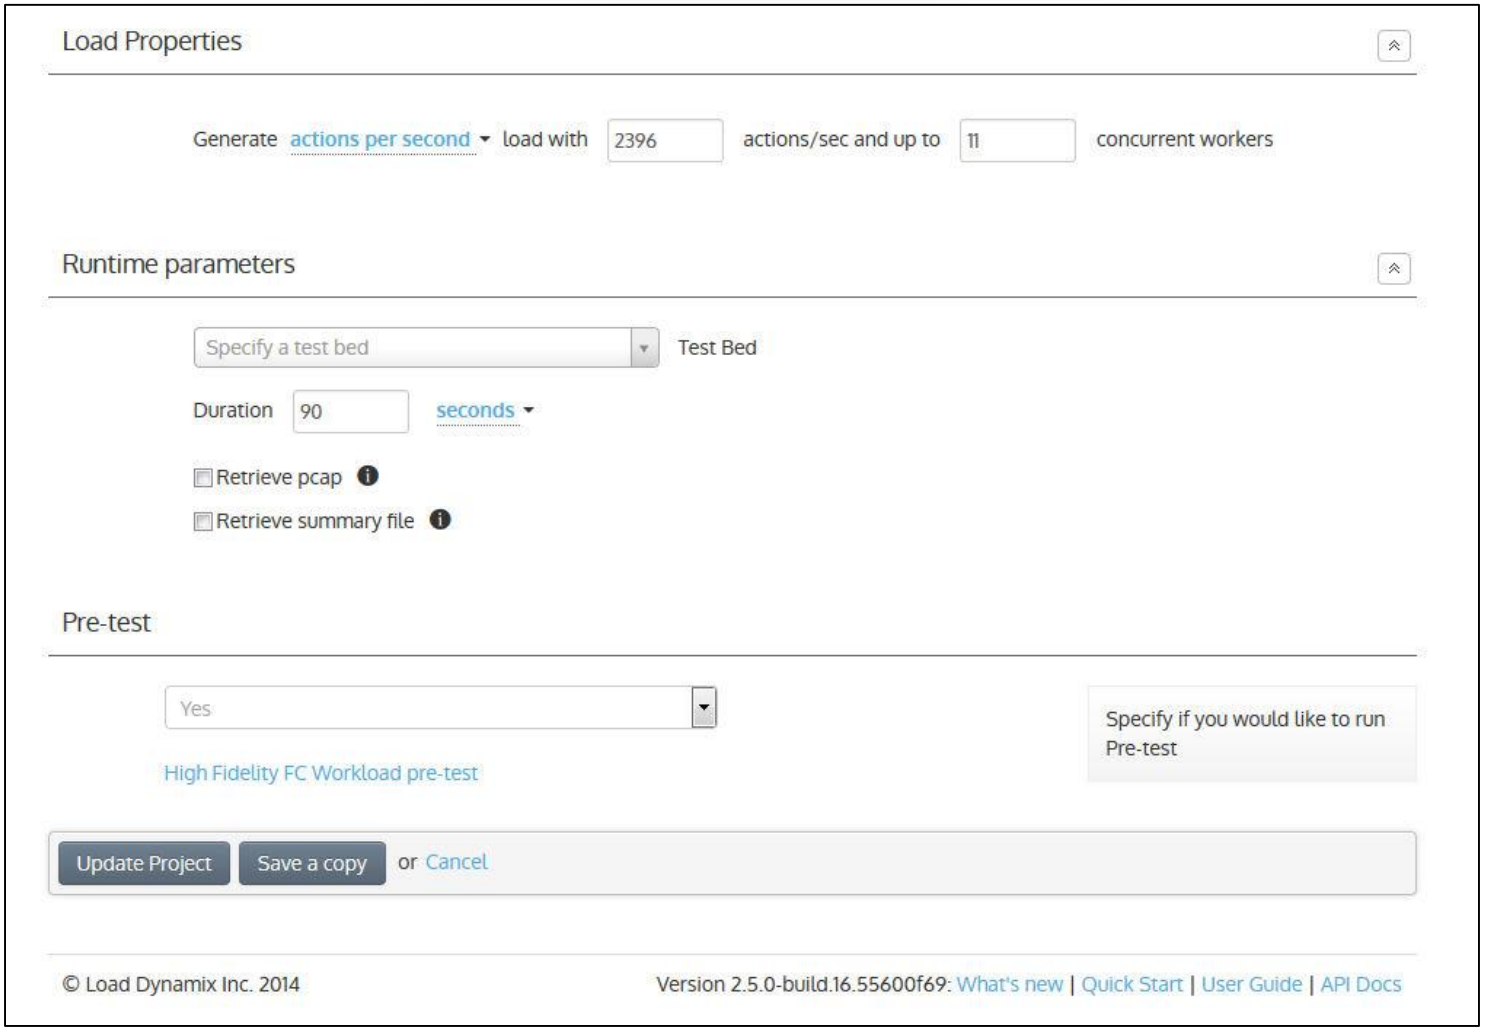

3: Deploy Test Configuration, Run Emulations

Load DynamiX Enterprise

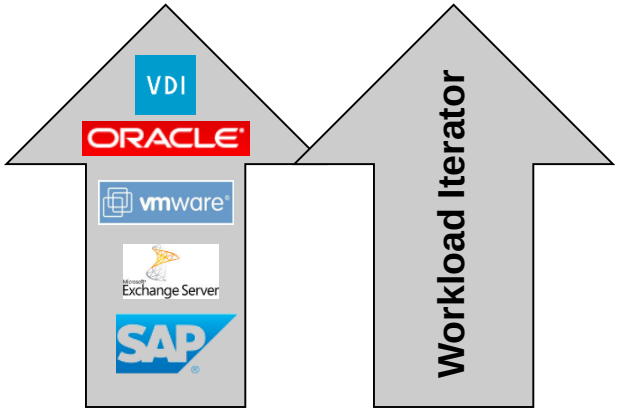

FC
iSCSI
SMB
NFS
S3
HTTP-S
Swift
CDMI

File, Block, Object

Load DynamiX Performance Validation Appliance

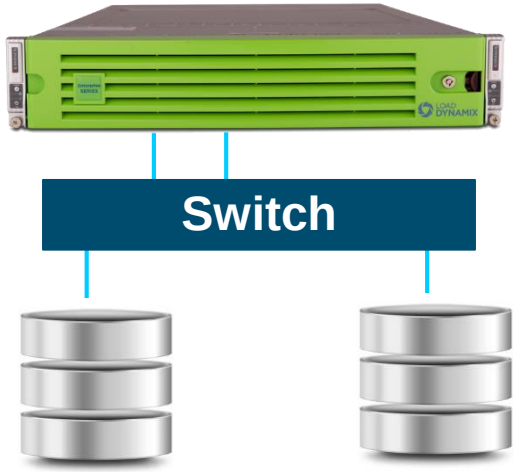

Product or Configuration A

Product or Configuration B

Example Test Run

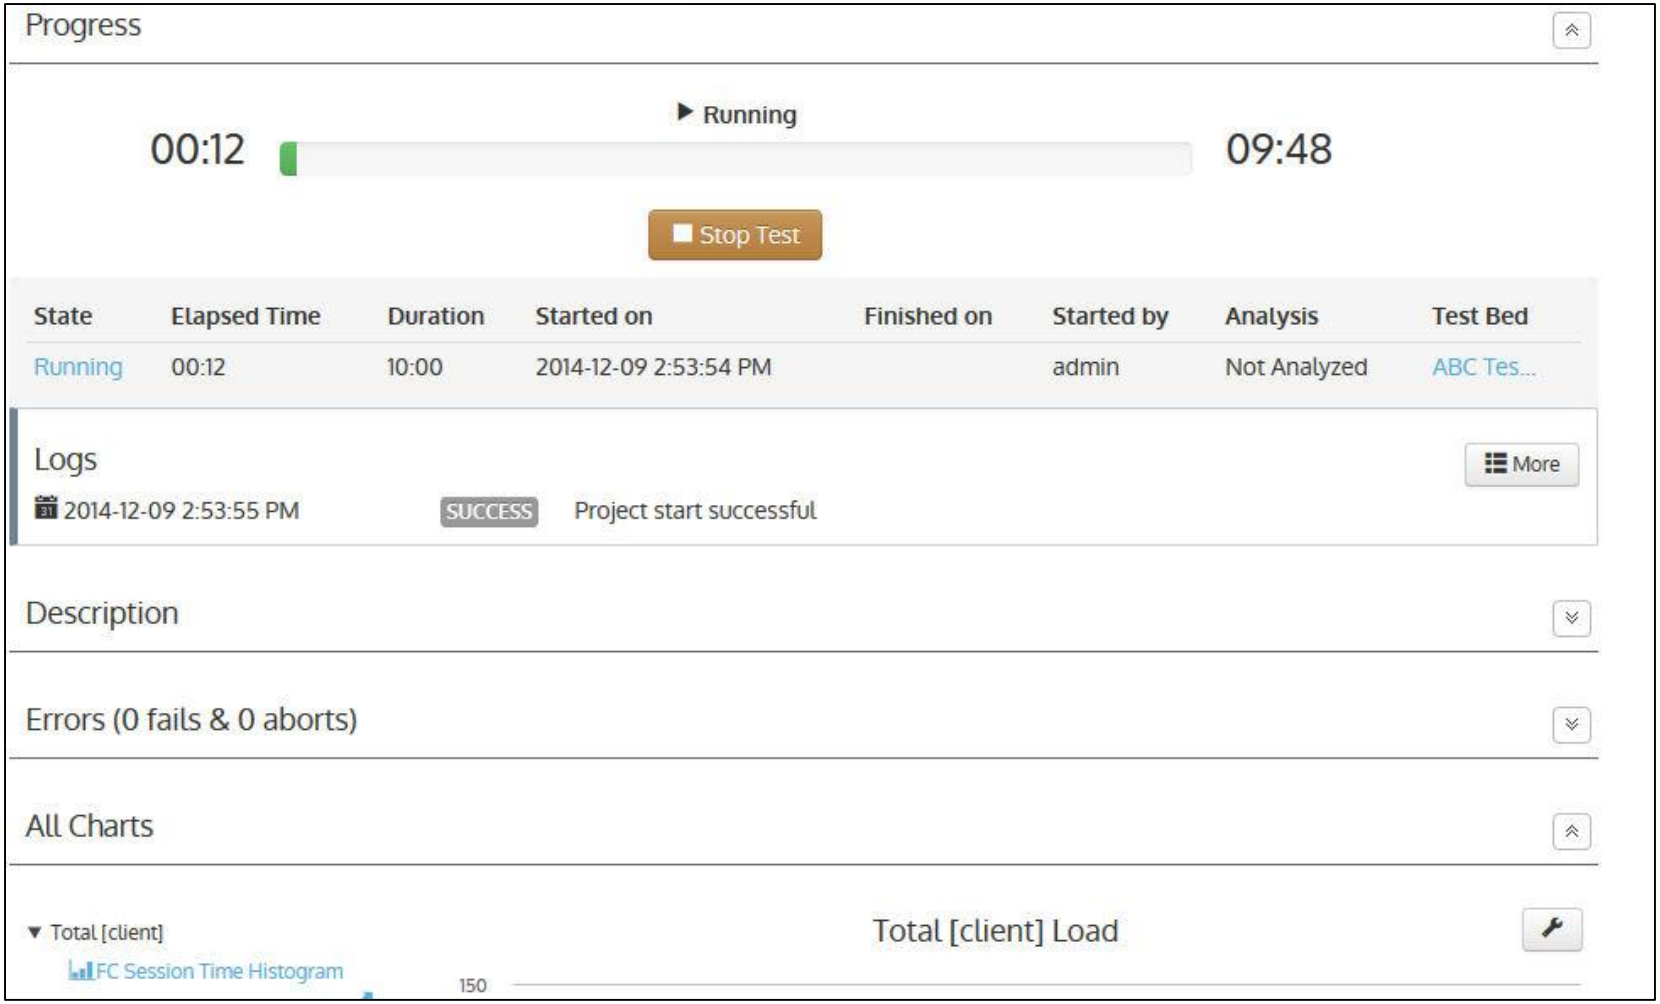

4: Analyze Results

Analytics

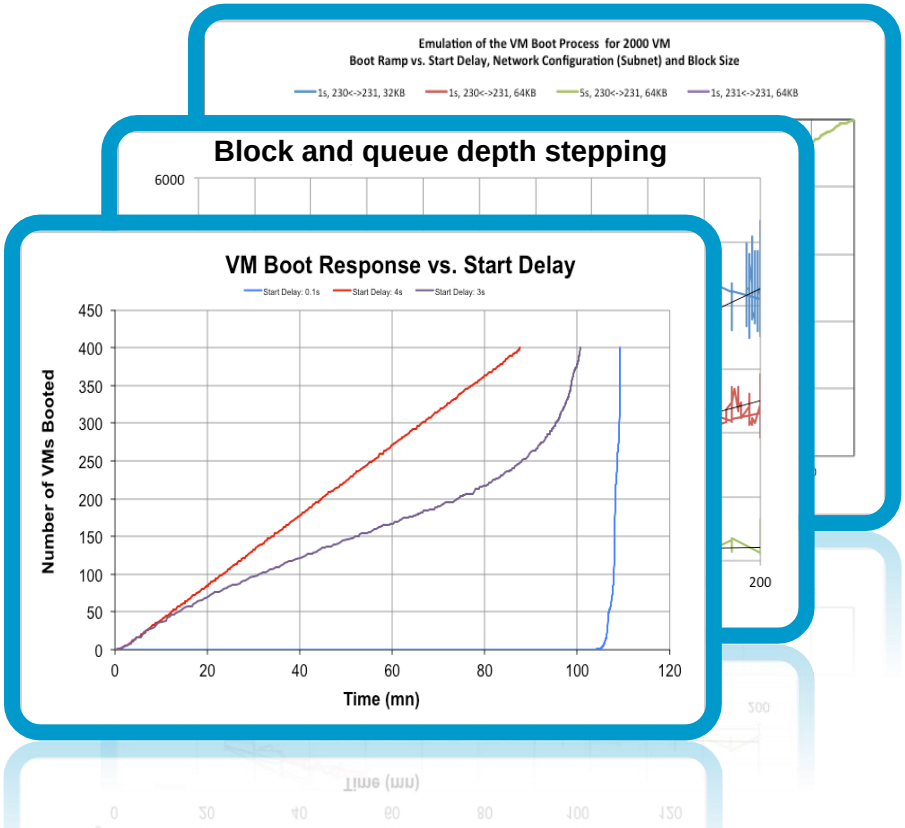

Insight

- Technology Evaluation
- Product Evaluation
- Configuration Optimizations
- Pre-Production Staging Validation
- Change Validation

Example Statistics

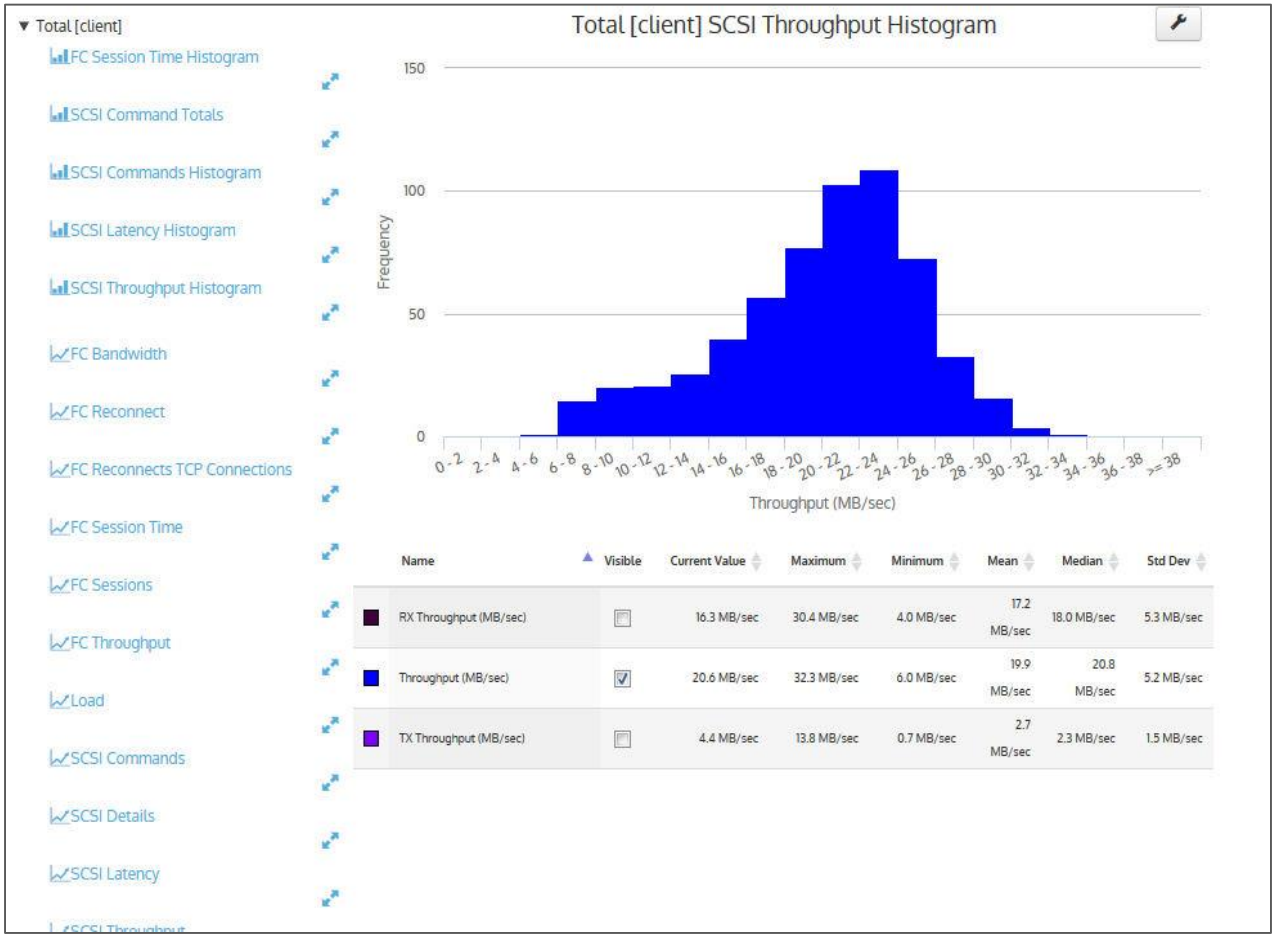

Agenda

- Introduction
- Load DynamiX Testing Methodologies
- Performance Profiling
- Workload Modeling
- **Case Study**
- Summary

Flash Storage Product Selection
NYP case study

Customer Overview
- Large healthcare organization
- >20,000 staff supporting >2,000 beds

IT Challenge
- Determine best price / performance flash storage systems to support current and expected workloads
- Finding the right tools / methodology to use to test with

Solution
- Load DynamiX

Benefits
- Data to enable decision to select the best vendor / array
- Validated the configurations that would support the workloads
- Full confidence that the storage systems from the new vendor array can address performance issues

"We wanted to test certain 'smart' storage arrays, that support dedupe & compression; therefore Iometer and similar tools aren't appropriate."

Storage Engineer
NY Presbyterian Hospital

Use Case
Storage product selection – flash storage

- Comparative performance of all-flash and hybrid systems
- Run workloads that reflect specific applications
- Determine optimal price/performance using performance profiling and workload modeling

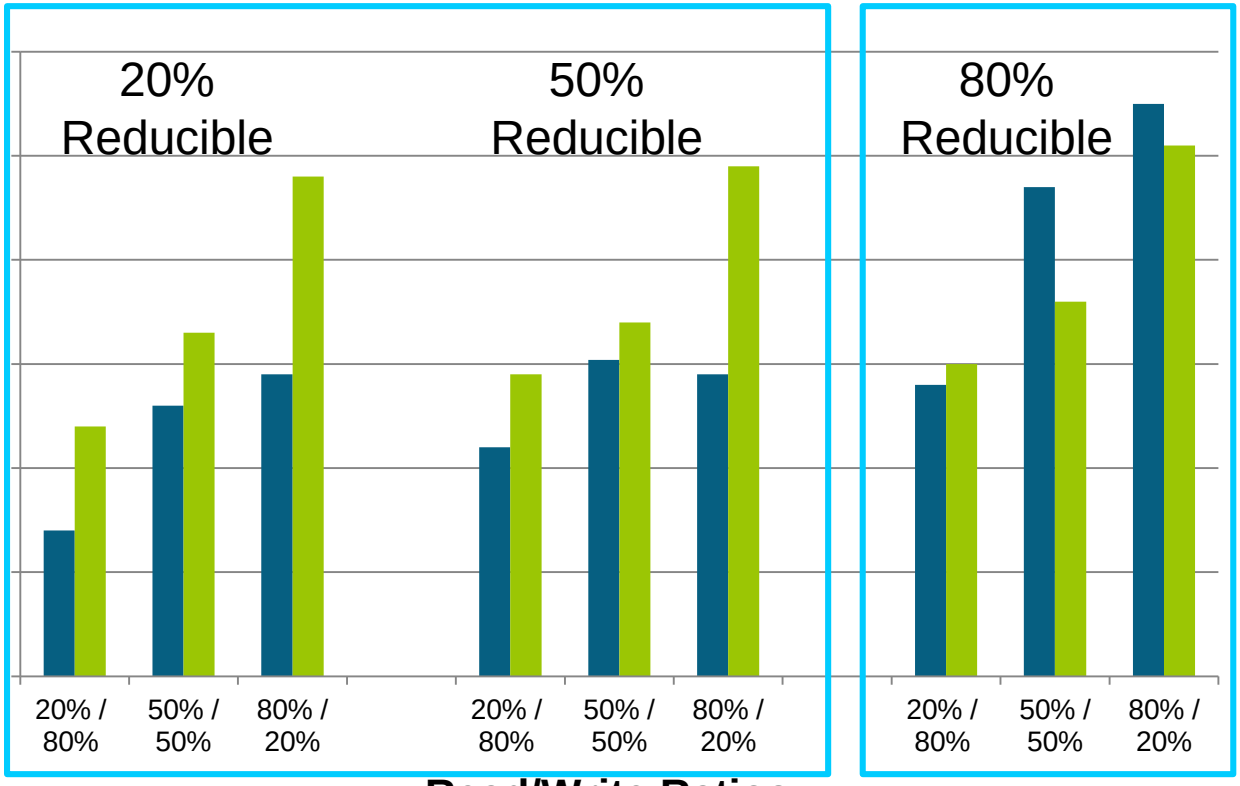

Dedupe Effects: Vendor B does best when data is 50% or less compressible

Dedupe Effects: Vendor A does better when data is highly compressible

Case Study
Storage product selection – flash storage

Primary business-critical application:

- Array A performance (left hand graph) exceeded workload profile requirement, achieving 240 MBytes/sec throughput, consistent with the workload profile requirement. Spikes matched application load spikes
- Array B array reached approximately 120Mbytes/sec, failing to meet required application throughput requirement
- Array A was selected for this business critical application

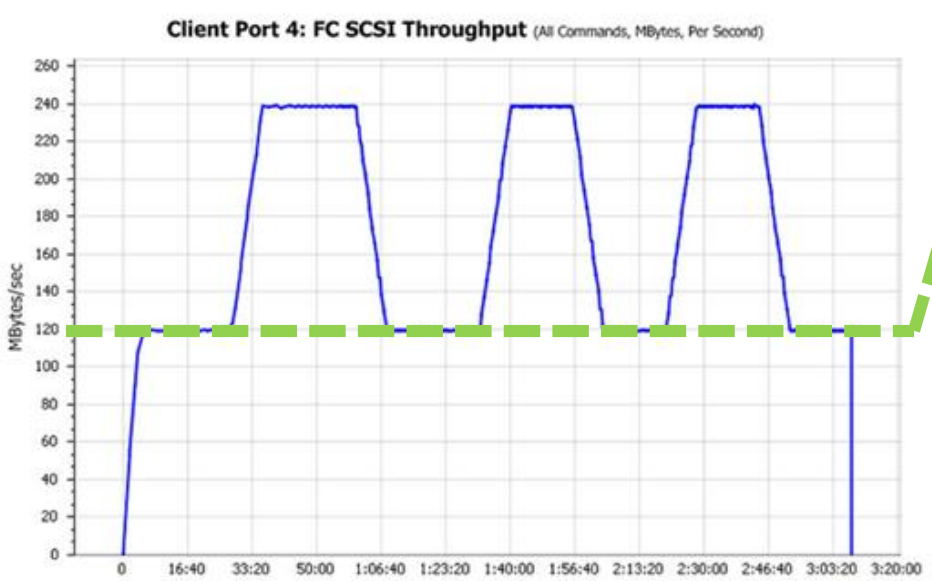

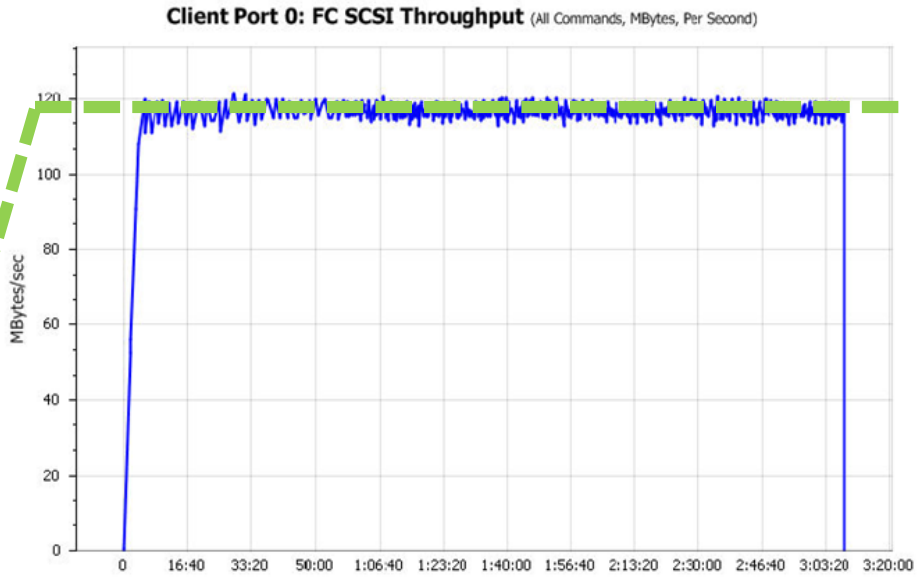

Agenda

- Introduction
- Load DynamiX Testing Methodologies
- Performance Profiling
- Workload Modeling
- Case Study
- **Summary**

Summary

- Performance assurance: You can ensure your AFA solutions will meet performance SLAs under your specific workloads. Chose the optimal solution for your workloads.
- Reduced storage costs: you can reduce over-provisioning and choose the lowest cost AFA systems for specific workloads. Quantify the benefit and effects of AFA systems.
- Increased uptime: You can identify problems in the development lab prior to production deployment; validate all infrastructure changes against workload requirements and troubleshoot more effectively by re-creating failure-inducing workload conditions in the lab.
- Acceleration of new application deployments: You can accelerate time to market by validating new applications on your AFA systems; making deployment decisions faster and more confidently.

QUESTIONS?