SDC
STORAGE DEVELOPER CONFERENCE
SNIA SANTA CLARA, 2017

Getting it Right II: Testing Storage Arrays The Way They'll be Used
An Example Implementation

Peter Murray
Virtual Instruments

Typical Performance Testing Questions

- Which is the best technology for my needs?
- Which is the best vendor / product for my needs?
- What is the optimal configuration for my array?
- Does performance degrade with enterprise features:
  - Deduplication
  - Compression
  - Snapshots, Clones, Replication
- What are the performance limits of a potential configuration?
- How does an array behave when it reaches its performance limit?
- Does performance degrade over time?
- How does a new firmware version affect array performance?
- Does it cause regression in failover, replication, snapshots?

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
- IT customers do not generally benefit from this testing
  - Applications don’t act like performance profiles
  - Some exceptions; e.g. queue depth or outstanding commands

Performance Profiling

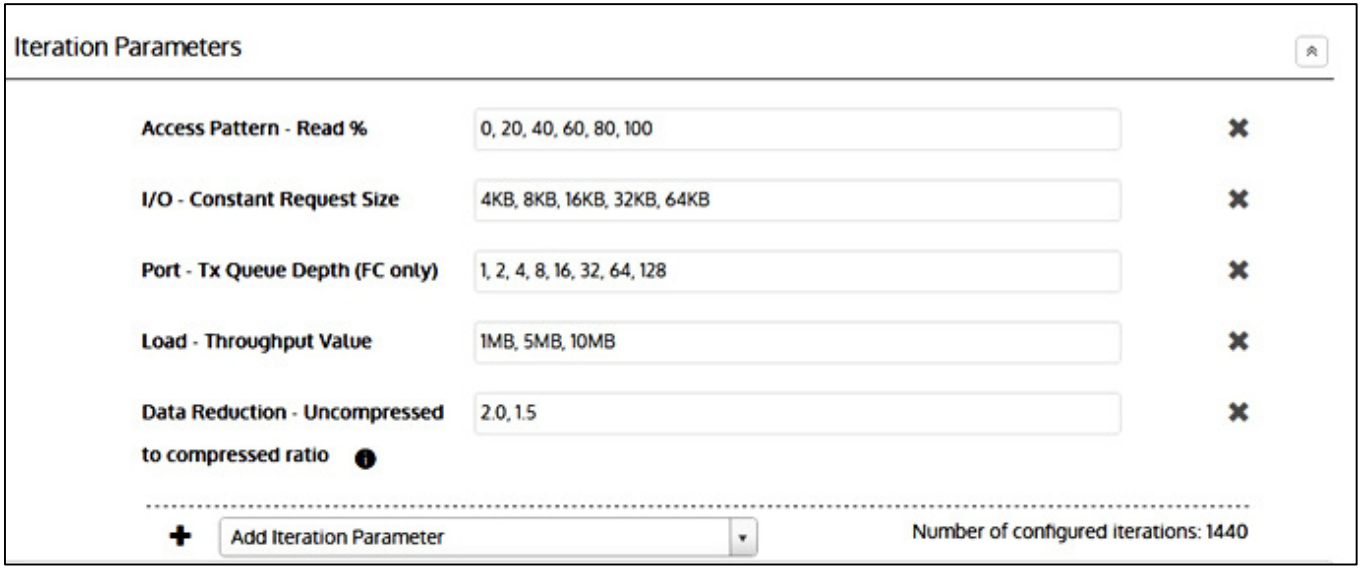

Performance Profiling
(continued)

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

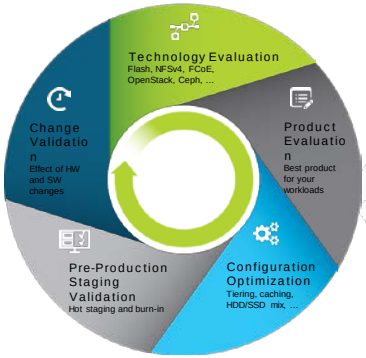

Workload Modeling
Simulate the I/O profiles of your production environment

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

Workload Modeling

Virtual Instruments combines a storage workload modeling application – Load DynamiX Enterprise, with purpose-built load generation appliances and data capture probes, to help storage architects and engineers to accurately characterize storage performance.

Where Does Workload Modeling Come From?

- Customers ask for workload models
  - Vendors want “the” application workload
    - Oracle, Exchange, etc.
  - IT customers want “their” workload models
- Vendors ask for help to:
  - Find realistic array and application scaling limits
  - Test customer examples/issues
- IT customers ask to help make better decisions about:
  - Upgrading storage hardware or software
  - Changing storage network configuration
  - Consolidating workloads

Workload Modeling

- Stresses an array using a realistic simulation of specific production workload(s)
  - For vendors, using customer examples or “dog food”
  - For IT customers, from the current environment
- Realism is paramount – running better I/O profiles
- Method produces better realism without packet trace limitations
  - Workload profile is smaller
  - Longer duration
  - Fewer security concerns

Workload Modeling Process

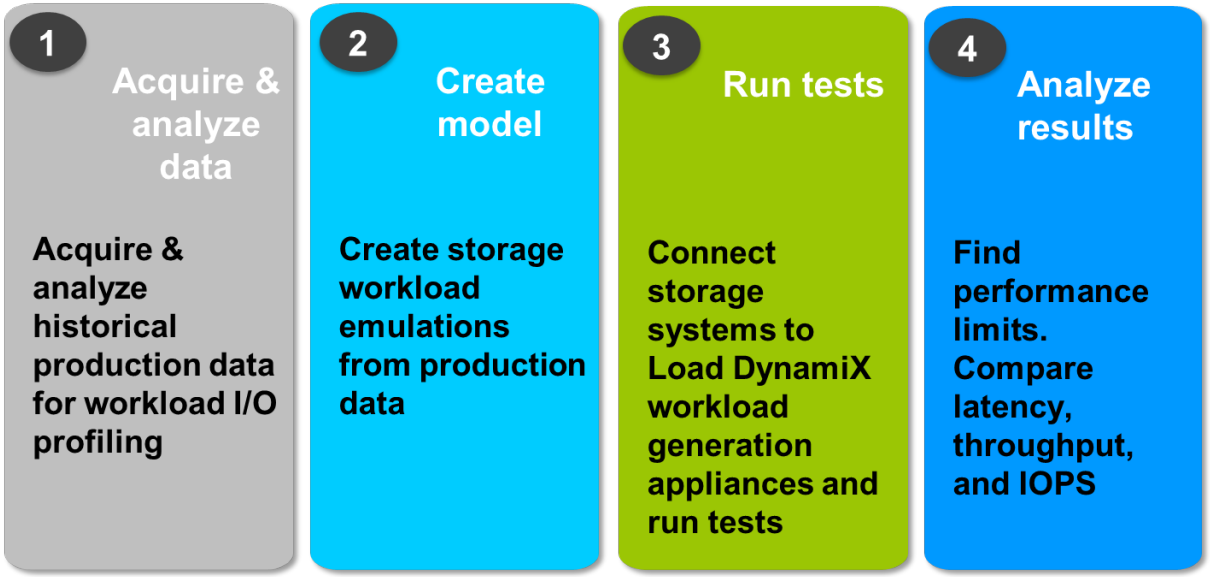

Acquiring and Analyzing Data: Sources

- Wire data from taps
- Storage array performance tools
- Traces
- Guesstimates

Acquiring and Analyzing Data: Taps

- Advantages
  - Data directly related to an application
  - Sometimes, very detailed
  - Few security concerns
- Disadvantages
  - May not be fine-grained enough

Acquiring and Analyzing Data: Perf Tools

- Advantages
  - Often, data directly from one application
  - May be highly detailed
  - Few security concerns
- Disadvantages
  - Resolution highly varies according to vendor
  - Not always public (may be difficult to obtain)
  - Perf tool may summarize data across applications

Acquiring and Analyzing Data: Traces

- Advantages
  - Highly specific and accurate
- Disadvantages
  - Short duration
  - Difficult to pinpoint exact conditions
  - High security concerns

Acquiring and Analyzing Data: Guesstimates

- Advantages
  - Simple to set up
- Disadvantages
  - Not reflective of any specific application

Acquiring and Analyzing Data

Name	Host IOs/sec	Read Response Time (ms)	Write Response Time (ms)	% Hit	% Writes	% Reads	% Read Miss	WP Count	Avg I/O Size	Capacity (GB)	% Used	%I/O Avg		Member	%RR	%SR	%RW	%SW	
arc	2.8	1.7	7.7	100	98.4	1.6	2.9	0	347	256	98	0.10	2745.9	4			93		7
dbf1	522.1	2.7	0.8	51.5	0.4	99.6	48.7	10.9	19	256	95	19.01	2745.9	4	95	4			
dbf2	448.5	2.9	0.8	51.3	0.1	99.9	48.7	2.4	16	256	100	16.33	2745.9	4	94	6			
dbf3	316.6	1.8	1.2	82.5	5.2	94.8	18.4	96.8	19	256	100	11.53	2745.9	4	84	11	5		
dbf4	297	1.8	0.9	42.9	0.9	99.1	57.6	8	29	100	100	10.82	2745.9	2	99		1		
dbf5	235.6	1.4	1.2	87.2	4.8	95.2	13.4	65.8	17	256	100	8.58	2745.9	4	84	11	5		
dbf6	220.2	1.7	1	83.9	5.6	94.4	17	58.9	20	256	100	8.02	2745.9	4	84	11	5		
dbf7	201.4	3.3	1.4	82.7	1.5	98.5	17.6	17.2	237	256	95	7.33	2745.9	4	94	4	1		
dbf8	165.7	3.1	1.1	66.2	5.1	94.9	35.6	35.1	19	200	83	6.03	2745.9	4	91	3	5		
dbf9	91.9	1.3	2.2	88.3	6.2	93.8	12.4	24.7	17	100	100	3.35	2745.9	2	82	11	6		
dbf10	90.3	3.3	2.3	71.6	27.7	72.3	39.1	145.7	48	200	99	3.29	2745.9	4	73	1	26		
dbf11	7.6	5.4	1.3	57.9	17.8	82.2	51.3	6.3	105	256	100	0.28	2745.9	4	81	1	18		1
oraex	1.5	3.6	0.7	62.6	17.7	82.3	42.2	1.4	2	33	86	0.05	2745.9	1	82		17		2
quest	6.3	0.8	1.4	98.5	88.5	11.5	7.2	13.5	13	10	40	0.23	2745.9	1	9	2	86		4
redo1	70.2	6	0.7	87.9	96.9	3.1	20.3	63.6	28	32	93	2.56	2745.9	4	3		88		9
redo2	68.1	0.5	0.8	88	99.6	0.4	0.9	68.4	14	32	93	2.48	2745.9	4			90		9

Clustering

- Typical application load varies widely across LUNs
  - A few LUNs handle much of the traffic
  - Many operate at much lower rates
- Using uniform load across multiple LUNs is unrealistic
  - Overtaxes array, understates array performance
  - Load should be divided into groups (clusters) to reflect the relative traffic levels and content

Why Cluster? An example

- Early on, we tested by treating all LUNS the same
  - One workload repeated for hundreds of LUNs
  - Used customer's busiest time to record the workload
- With this configuration, the test experienced 130ms latency
- When test re-ran using clustering, latency dropped dramatically to < 4ms, with high throughput and IOPs

Implementing Clustering

- Choose a limited number of example workloads as input
  - Busiest and least busy – often use 8 workloads as models
- Distribute example workloads proportioinally among total number of LUNs to be tested
- Produce a model from this cluster

Analyzing Acquired Data

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**

Importing a Data Workload

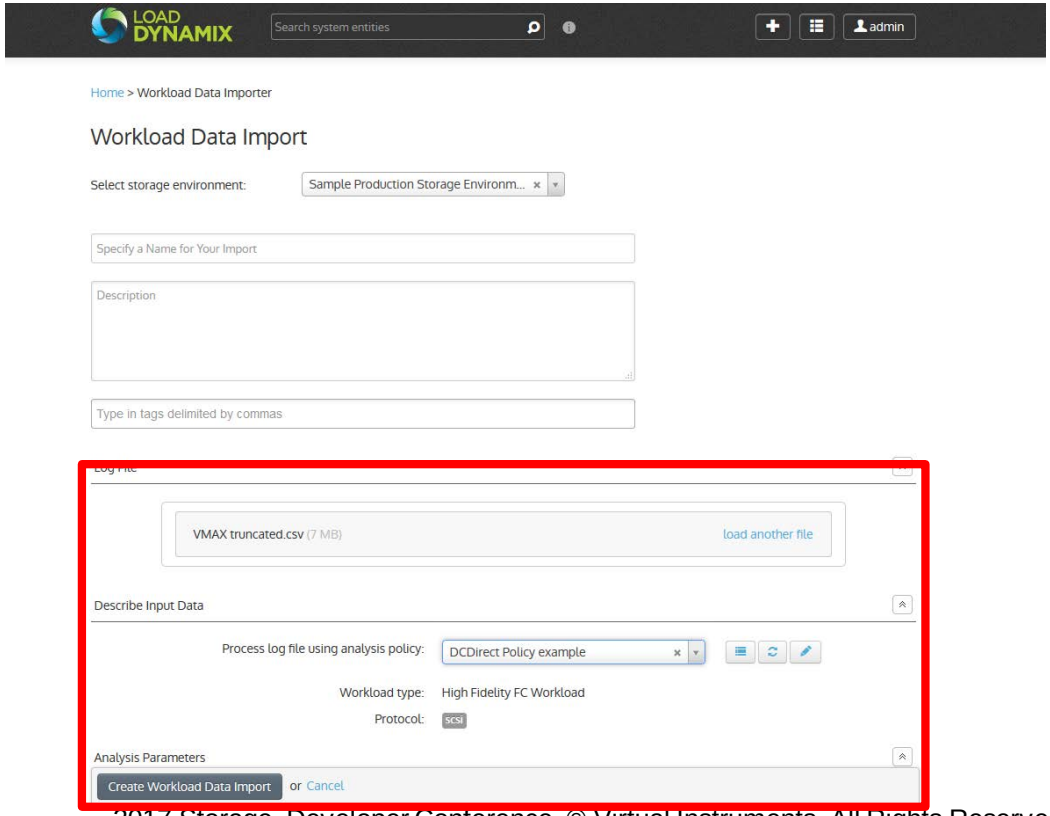

Parsing the File

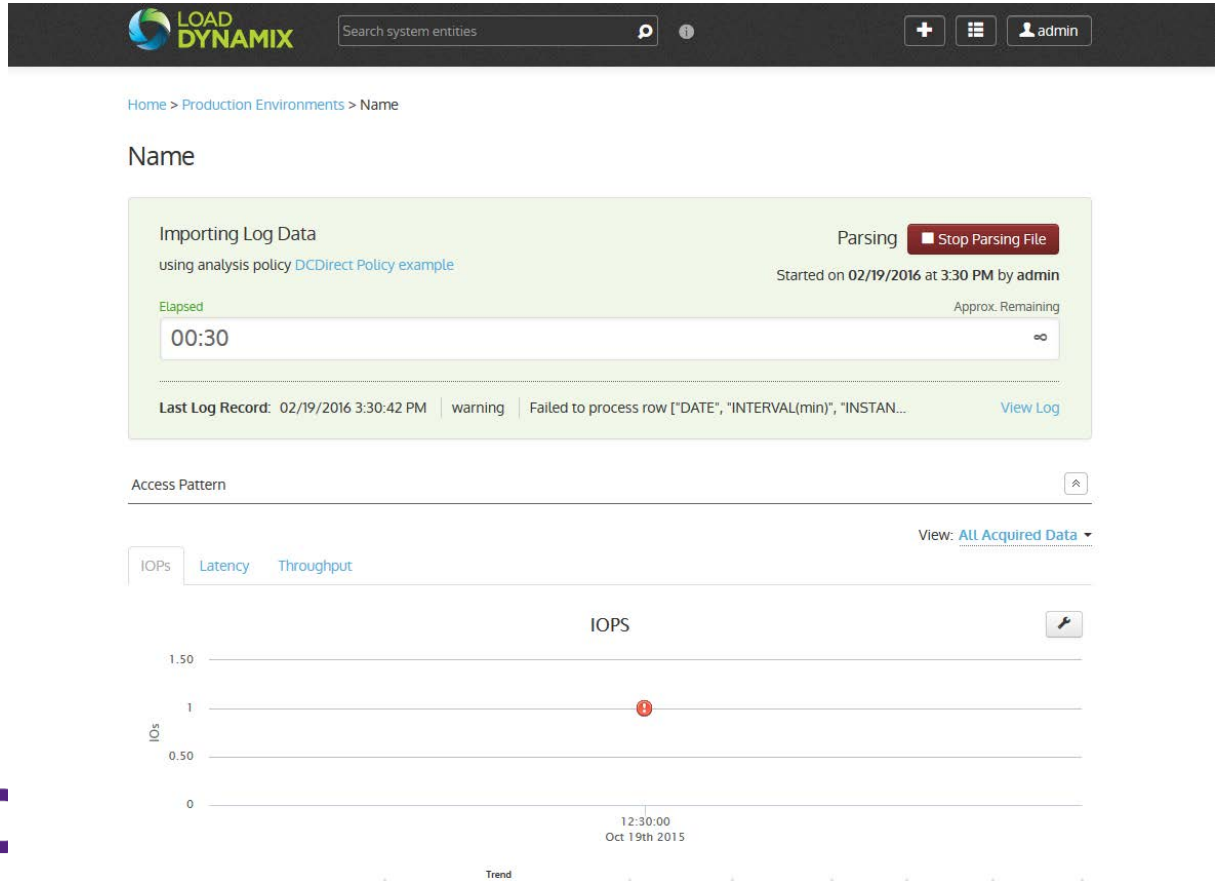

Viewing the Results of the Import

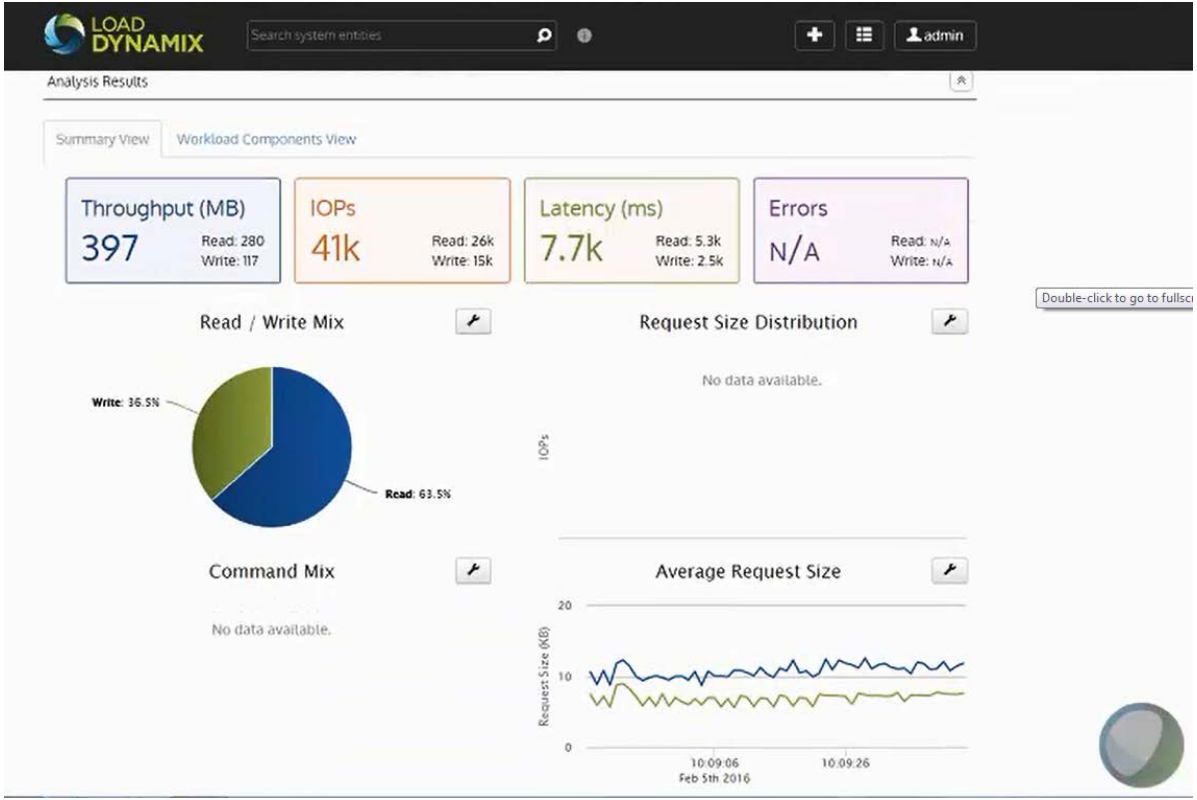

Workload Central Beta

- 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

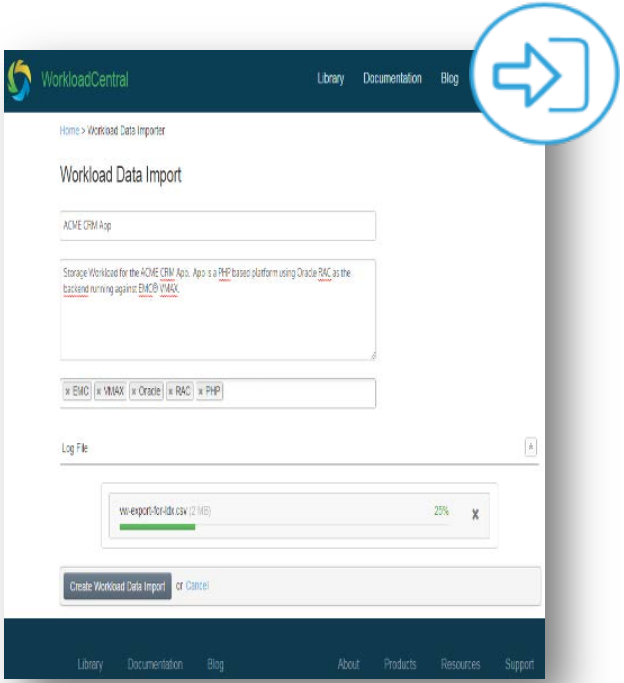

- The Workload Importer offers:
- Ability to upload data from any vendor or environment
- Out of the box import policies
- Analysis policies provide flexibility to define different workloads

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

Considerations for Other Testing Tools

- This method is designed for use with Load DynamiX Enterprise and the Load DynamiX appliances
- With other tools:
  - When testing using VMs, use the exact same OS version, Hypervisor Version, drivers and network configuration when comparing systems
  - Enable bursts!
  - Ensure that the Test OS/hypervisor and HBA Drivers all support manipulating queue depth and are set the same
  - Make sure to test with multiple profiles to accurately represent the application

Summary

- This process enables vendors and customers to:
  - Upload many workload types
  - From multiple vendors
  - Test using this data to:
    - Compare software versions
    - Compare products
    - Ensure apples-to-apples comparisons

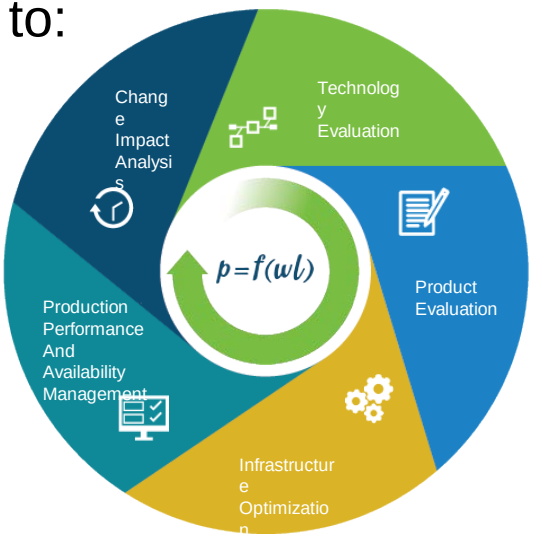

Thank You

For more information:

peter.murray@virtualinstruments.com

www. virtualinstruments.com