

Object Storage Performance Validation

Tim Van Ash

VP, Products, Load Dynamix

Twitter: @tvanash

Email: tvanash@loaddynamix.com



**LOAD
DYNAMIX**


Sept 2014

Agenda


- State of Object Testing
- Things to consider when testing Object storage
- What is Load DynamiX doing today for Object storage
 - Easy to use Object Storage Validation
 - Comprehensive Unified Storage Validation
 - Granular control at the client & protocol level
- Conclusion

State of the Nation for Object Storage

Customers are actively evaluating Object storage



In the last 12 months we have seen a significant increase in customers doing Object Storage performance validation



The need for Object Storage Validation

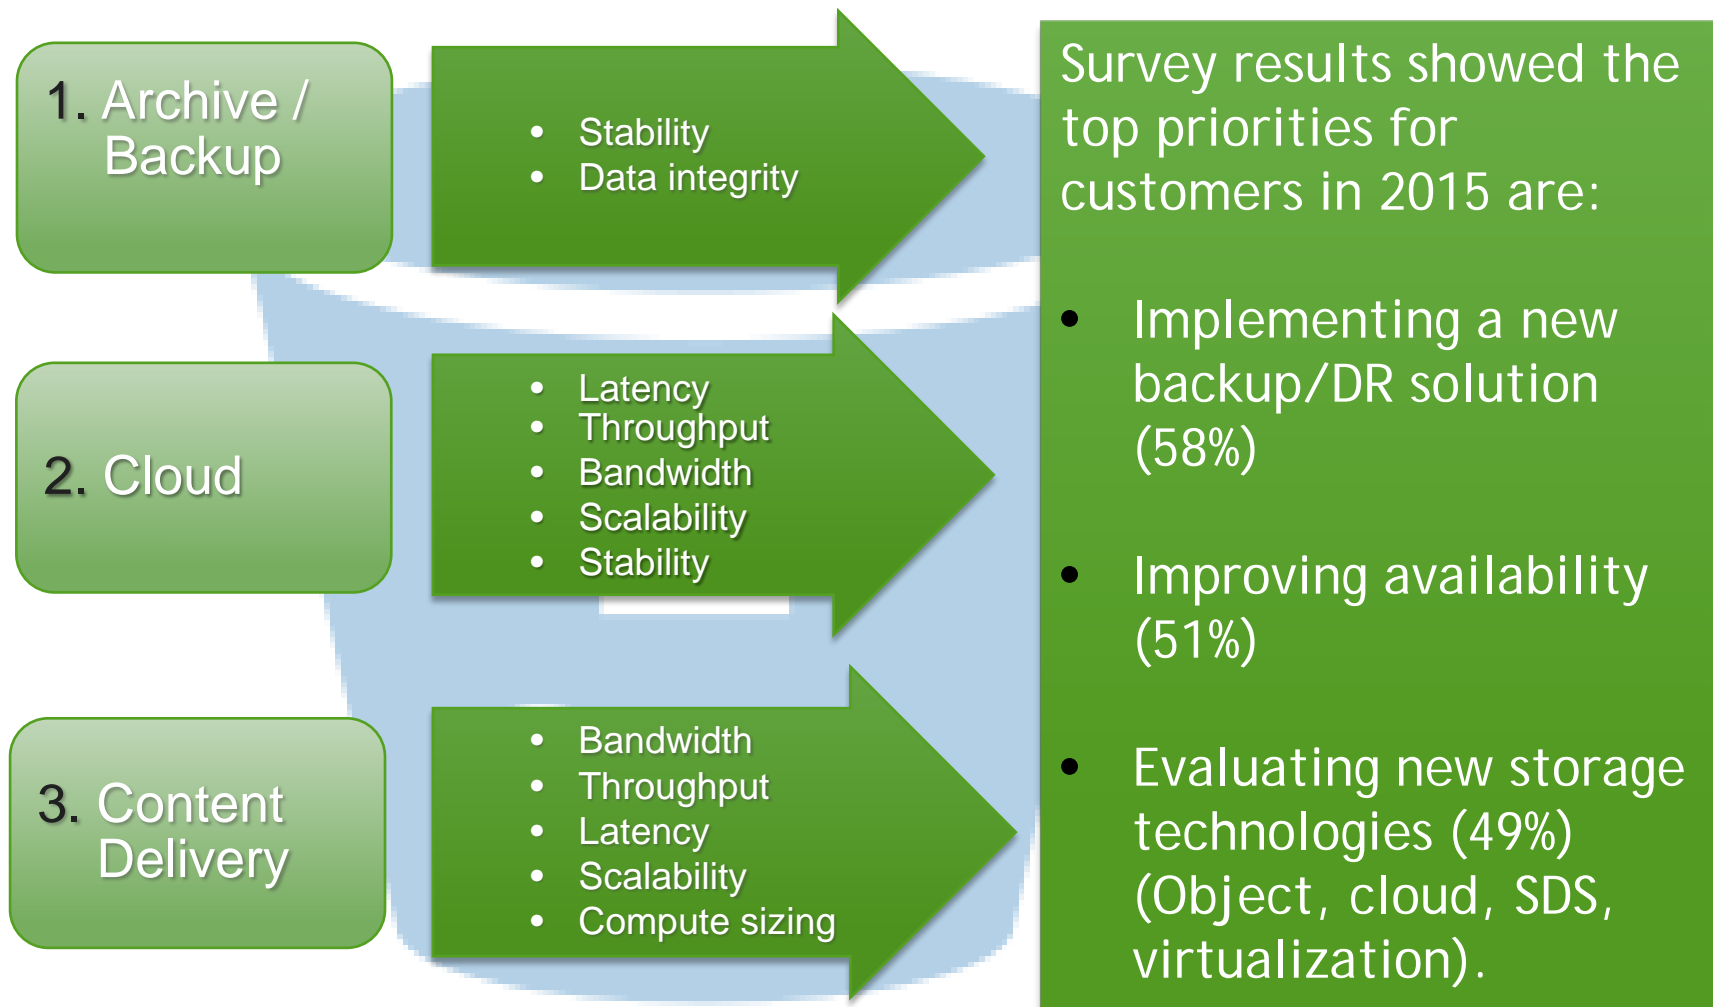
The disconnect we see...



- Customers are looking to reduce storage costs
- Open source solutions are being evaluated by every customer
- Different levels of maturity across the common platforms:
 - ▶ S3 well defined and strong adherence to documented API behavior.
 - ▶ Swift evolving fast; often gaps in documentation
 - ▶ Commercial implementations differ in supported operations & integrations
- There is often a disconnect between vendors offerings & customers use cases

Top use cases we are seeing...

Performance & Stability are inherent challenges



Object Performance Validation

Things to consider when testing performance

■ Overall Performance

- ▶ Validate latency for different object sizes. Vendors tend to optimize for large objects. Customers workloads are mixed.
- ▶ Need to be able to validate per service, per cluster **and per for** read/write throughput and aggregate across. Line rate is rarely achieved by clusters.
- ▶ Unified storage models need to validate file, object and block interfaces concurrently under load.

■ Meta-Data Performance

- ▶ Meta data operations performance can vary significantly under load. Different protocols handle meta-data differently. Swift places it in the http header, S3 in the http body.
- ▶ Eventual consistency in Amazon S3. There can be considerable latency when making changes to buckets that affects meta data performance.

Object Performance Validation

Things to consider when testing availability

■ Authentication

- ▶ Multiple authentication protocols and versions need to be validated (ie S3 v2/4 or Keystone 2/3)
- ▶ Time skew will cause authentication failures in S3 & Swift
- ▶ Token expiration in Swift can be problematic

■ Availability/Reliability

- ▶ Stability of distributed nodes, and the overall cluster under load is often a challenge.
- ▶ Complex service availability. We often see proxy architectures respond even though core services are down. Creates weird errors
- ▶ Multi-tenancy adds another layer that needs to be validated across tenants running different load patterns

Object Performance Validation

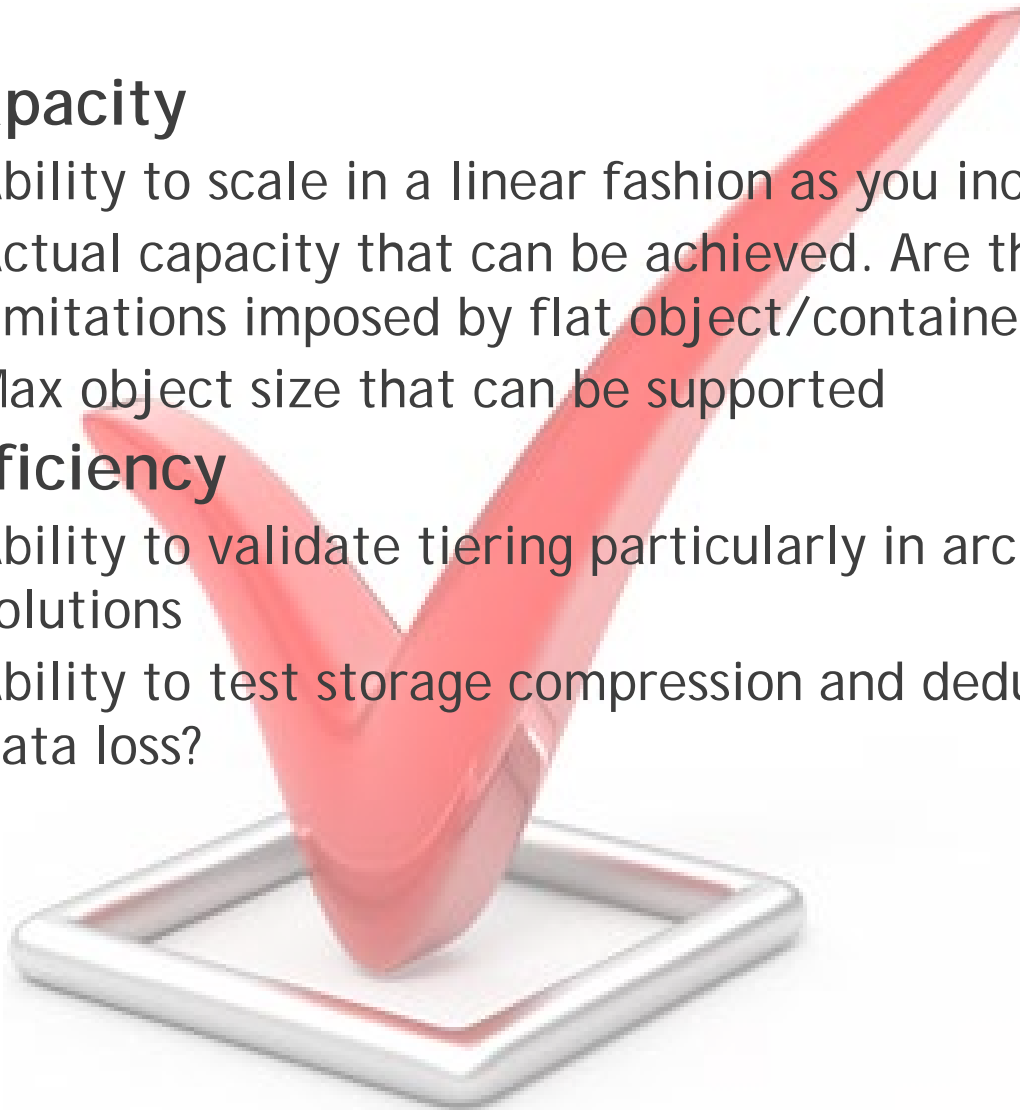
Things to consider when testing capacity

■ Capacity

- ▶ Ability to scale in a linear fashion as you increase nodes
- ▶ Actual capacity that can be achieved. Are there any, limitations imposed by flat object/container structures?
- ▶ Max object size that can be supported

■ Efficiency

- ▶ Ability to validate tiering particularly in archiving or backup solutions
- ▶ Ability to test storage compression and deduplication. Is there data loss?



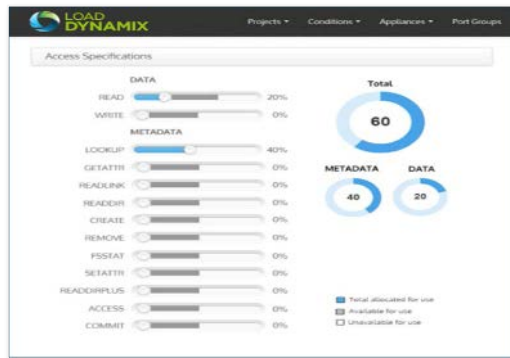


The Load DynamiX Solution


Industry Leading Object Storage Performance Validation

Load DynamiX
Enterprise software

Workload I/O Profiles



Performance Analytics



Load DynamiX Performance
Validation Appliance



Switch



Product or
Configuration A

File, Block,
or Object
storage



Product or
Configuration B

IOPS, Throughput, Latency

Easy to use, Enterprise Ready

Unified Object Workloads out of the box for Swift & S3



Search system entities



admin

Swift Object Workload

Privacy: Public

Object storage operations for Amazon S3 and OpenStack Swift protocols

Created by: admin
2015-07-23 4:48:46 PM

× Workload × Object × swift

Access Pattern

Configure access pattern as [OpenStack-Swift commands distribution](#)

Data

WRITE OBJECT 20%
RETRIEVE OBJECT 30%
COPY OBJECT 5%

Metadata

DELETE OBJECT 7%
WRITE ACCOUNT 2%
RETRIEVE ACCOUNT 10%
DELETE ACCOUNT 1%
WRITE CONTAINER 5%
RETRIEVE CONTAINER 5%

Setting data and meta data ratios & commands are easy

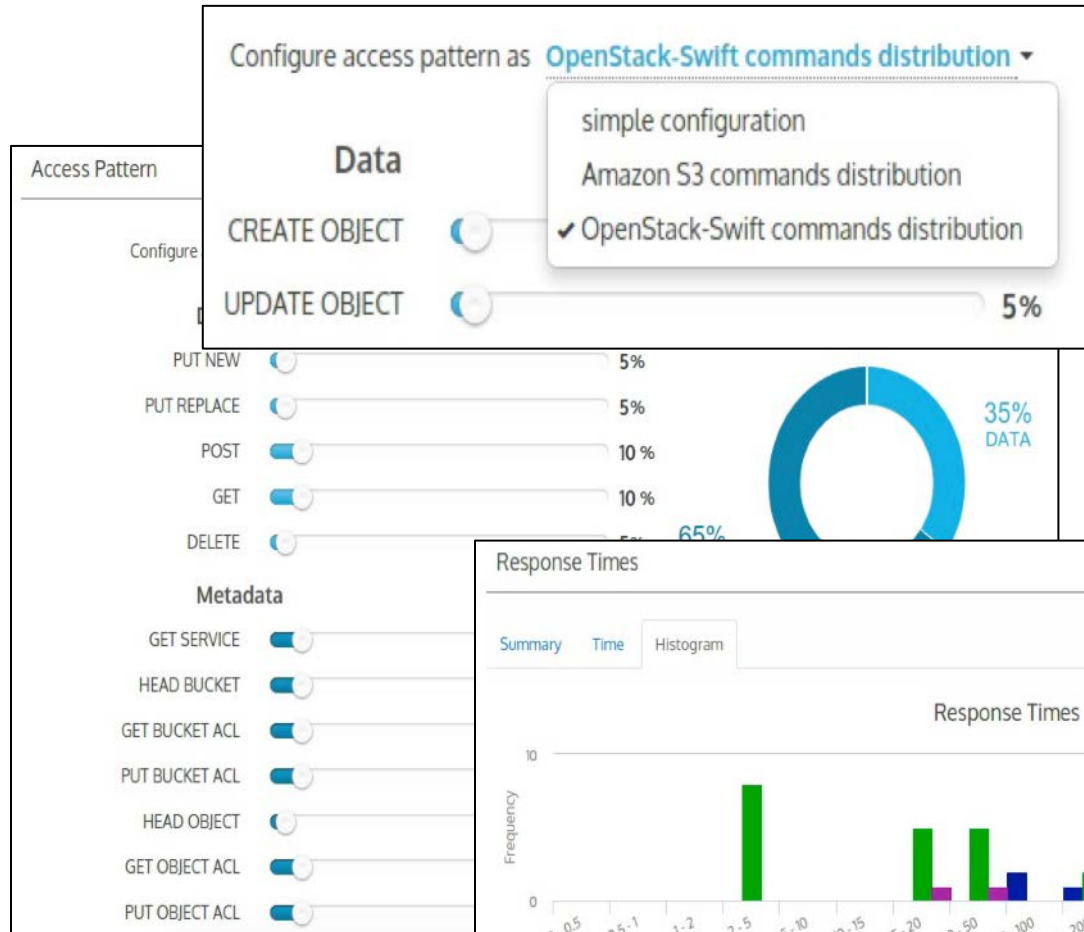


Legend

- Total allocated for use
- Available for use
- Unavailable for use

Load Dynamix Enterprise provides an enterprise platform for modeling, validating and managing your file, object & block storage validation.

Out of the box Object Workload Models For Swift & S3

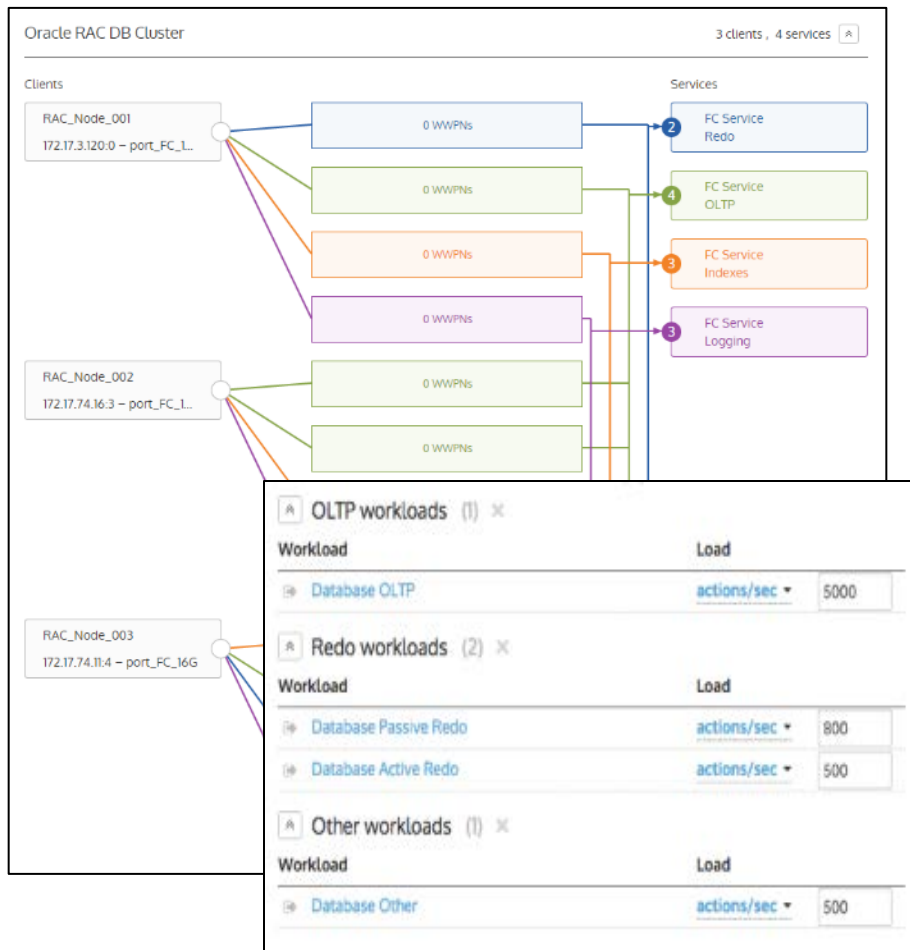


- Provides intuitive configurations to model object workloads for S3 & Swift
- Analytics provide the data to optimize object storage infrastructure



Web scale Performance Validation

Mixed Workloads modeling made Easy



Oracle Real-Application Cluster with multiple access patterns

What does a Composite Workload simulate?

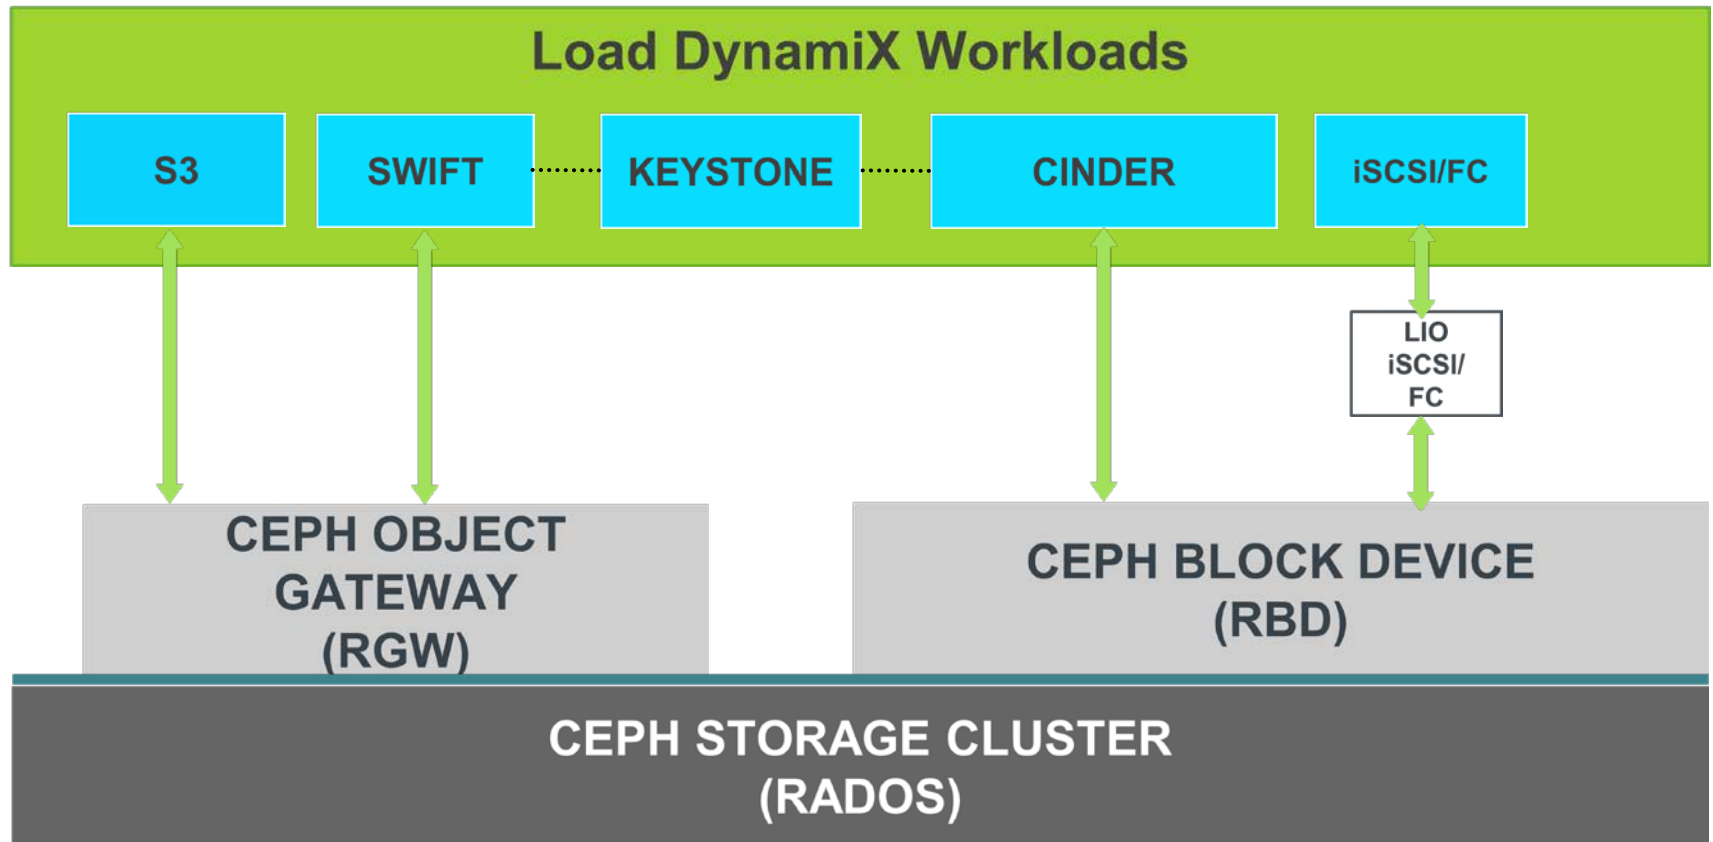
- Applications with multiple access patterns
- Across multiple tenants or Virtualized environments
- Multiprotocol / unified storage deployments

Customer Benefits

- Easy modeling of complex workloads
- The most realistic workload models available today

Unified Storage Performance Validation

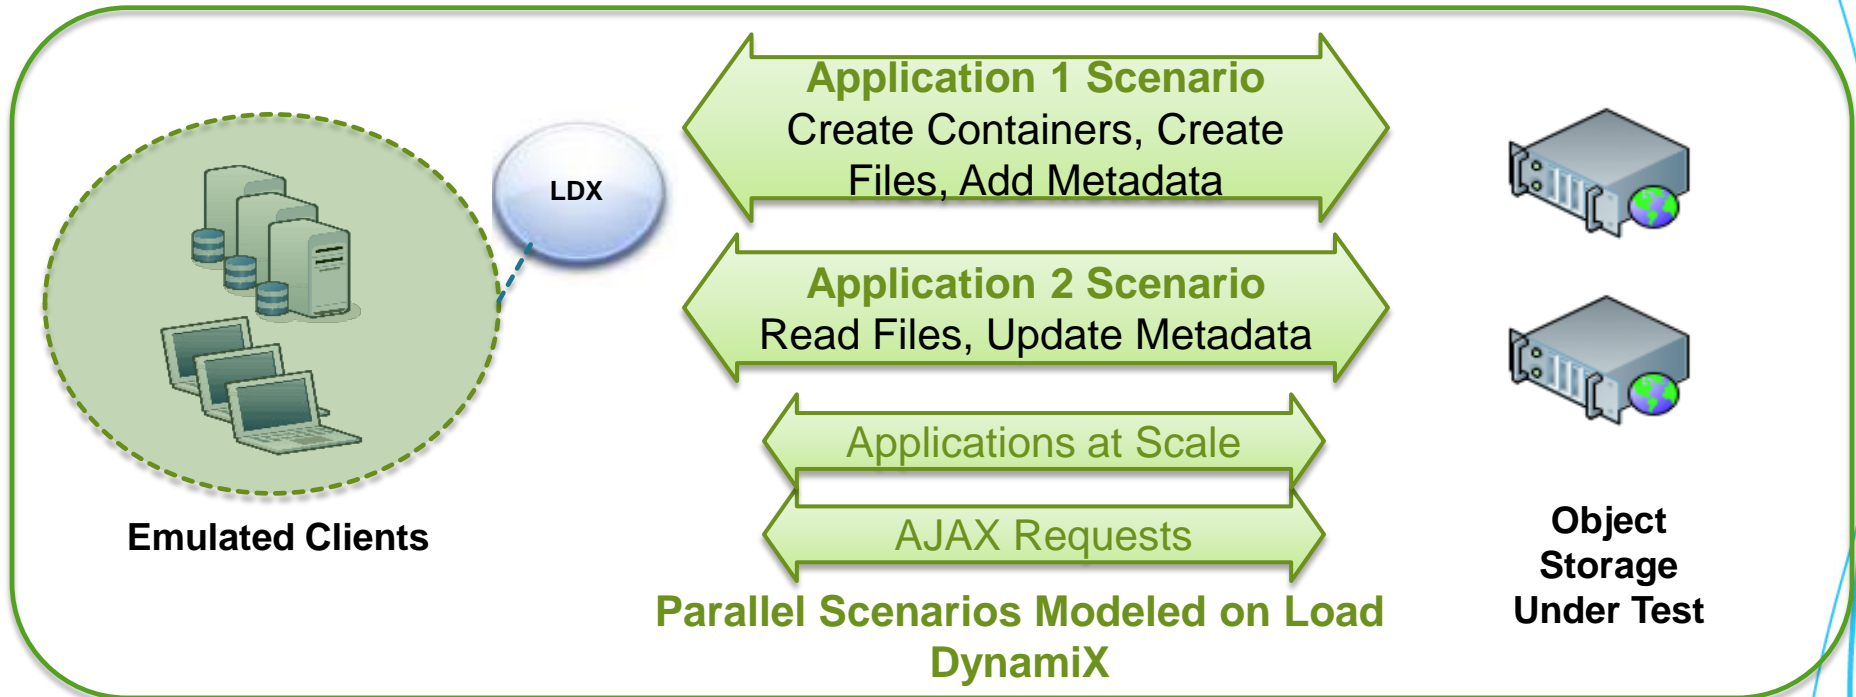
Ceph Performance Validation



Load DynamiX provides scalable unified performance validation for Object, File and Block

Object Storage Client Modeling

Granular control over client behavior



- Key Configurable Parameters:
 - ▶ Credentials, URI/URLs, HTTP Header/Body Content, HTTPS/Certificates, Load Profiles...
- Key Performance Indicators
 - ▶ Concurrent Applications/Connections, Scenarios/Users/Connections Per Second, Throughput, Time to First Byte, Min/Max/Avg Response Times...

Precise control when needed

Hyper scale performance validation for Swift, S3 & CDMI



1	HTTP	Open HTTP Connection
2	#	Container Operations
3	CDMI	Create Container
4	CDMI	Read Capabilities Object
5	CDMI	Create Container
6	CDMI	Read Container
7	CDMI	Update Container
8	CDMI	Read Container
9	#	Data Object Operations
10	CDMI	Create Data Object
11	CDMI	Read Data Object
12	CDMI	Update Data Object
13	CDMI	Read Data Object
14	#	Deleting all objects
15	CDMI	Delete Data Object
16	CDMI	Delete Container
17	CDMI	Delete Container
18	HTTP	Close HTTP Connection

Precise control
over Swift, S3 &
CDMI operations

The screenshot shows the 'Request Headers' configuration window. It has two tabs: 'Input' and 'Output'. The 'Input' tab is active, showing a list of headers with their values. The 'Output' tab shows 'Extract Headers' and 'Extract Body'. Below the tabs is the 'HTTP Content Editor' with a 'Request Headers' section. This section contains a table with headers and their values.

Name	Value
Host	Server
X-Account-Meta-Test	= @STRING(someparam) + @UP(0,CONTAINER)
X-Account-Meta-Test2	= @STRING(anotherParam) + @SCENARIOCOUNTER(3,999)

Deep definition of request headers for all operations and authentication



Rich statistics for in-depth
analysis

Object Storage Performance Validation

Scalable Workload Generation from Load DynamiX

User/Client Emulation

- Session Persistence
- Parallel Scenarios
- Object & Meta-data operations (and APIs)
- Authentication (S3 – v2/4), Keystone 2/3)
- Bucket/Container and Object ACL validation

Content Generation

- Compressible & dedupeable content
- HTTP/S Parsing: Content Insert/Extraction
- Dynamic HTTP Body Templates

Object Storage Web Services

- OpenStack Swift, Amazon S3, SNIA CDMI, other...
- Authentication (S3 v2/4, Keystone 2/3)
- RESTful Interfaces
- Form Submissions
- Multi-part Uploads
- AJAX emulation
- Page Load Times
- Large File Transfers
- Load Balancing

Scalability

- 19Gbps HTTP per port (stateful)
- 1M+ Open Connections
- 200K+ Connections/Sec
- High throughput 1/10GE
- Appliance Stacking

Insight/Debug

- REST Method Metrics
- L4-7: TCP/Protocol
- L2-3: Network Statistics
- Inline Tracing, Logging
- Data Verification/Integrity

Conclusion

- Customers are looking for better ways to meet storage and cost needs
- Object, open source and unified storage is seen as key to meeting those objects
- Vendor and customer use cases do not always align
- Performance validation requires testing across multiple dimensions
- Load Dynamix is the leader in Object and Unified storage validation for your customers workloads

Thank you!

