



Storage Performance Validation for Panzura

Ensuring seamless cloud storage
performance for Panzura's Quicksilver
Product Suite

Table of Contents

Background on Panzura	3
Storage Performance Challenges	4
Load DynamiX: Solving Storage Performance Challenges.....	4
Storage performance validation of Panzura's product infrastucture	5
Solving storage performance challenges at the "edge"	6
Accurately measuring compression and de-duplication	6
Addressing translation challenges	6
Conclusion	6

Background on Panzura

Founded in 2008 and headquartered in Campbell, California, Panzura is a venture-backed storage software company whose vision is to enable cloud storage to be faster and more secure than traditional, local tier 1 storage.

Panzura's Quicksilver Cloud Storage Controller provides a global, cloud-integrated file system. Going beyond existing appliances that offer simple common file access, or backup and archive support, Panzura brings the cloud gateway to the enterprise by providing file locking, access control and data encryption, among other features, with their unique Cloud Operating system and Global File System.

Users access data from edge appliances via NFS or SMB clients, just like a traditional NAS. But unlike a traditional NAS, the Quicksilver Cloud Storage Controller syncs with a remote public or private cloud service, writes data to and reads data from the service over HTTP and maintains the associated SMB or NFS metadata as if the remote service ran locally as a NAS server. It also caches new and frequently accessed data locally on high-speed flash storage to reduce latency typically experienced on distributed file system metadata processing. Global inline data deduplication and compression is implemented to more efficiently use WAN bandwidth.

The Quicksilver Cloud Storage Controller connects remotely to a central cloud storage pool (the core), supporting most public and private cloud providers and standards, removing silos of remote, under or overprovisioned storage at one of the individual storage components (the "edge").

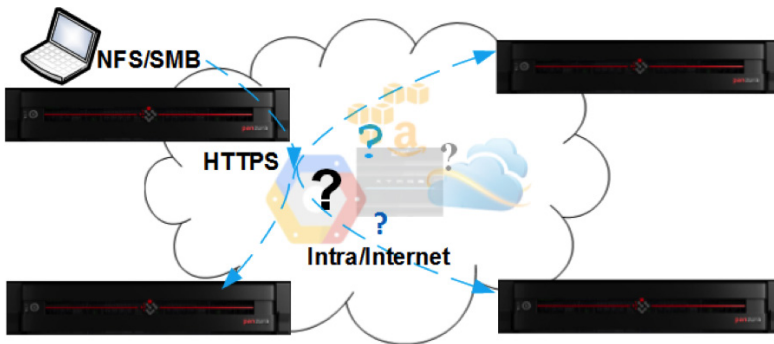


Figure 1: Architecture overview of the Panzura cloud storage product suite.

Storage Performance Challenges

However, an infrastructure of such large scope and a multitude of components require careful planning from the solution architect wishing to implement it. Most immediately, the performance requirements for the individual edge appliances needs to be understood. What is a realistic expectation of IOPs and throughput provided by the Quicksilver Controller? Is it sufficient for the number of clients wishing to connect to it at a particular site? How well will it scale in the future as that branch office grows? Also, for a given data set, how effective is the global de-duplication and compression? Will the solution require additional bandwidth between the central core and edge sites? An even greater consideration for the solution architect is the end-to-end latency and throughput. With the large number of options possible, how does one decide on the storage

How Load DynamiX Works



Create Model

Create accurate emulations of storage workloads by gathering usage data and applying storage statistics from your production environment. Or use Load DynamiX sample workloads.



Run Tests

Connect new storage systems or configuration candidates to the Load DynamiX appliance and run the tests.



Analyze Results

Isolate bottlenecks, investigate future workloads, determine rainy day vs. sunny day scenarios, and optimize storage limits.



Knowing that the latest product generation works as expected before it's deployed into the real world by our customers is worth everything. Thank you, Load DynamiX.

John Taylor
STORAGE ENGINEER,
GO DADDY

vendor and type of cloud storage for the back-end? In addition – because a change is made at the edges from a file based protocol to an object based protocol – how can the solution architect certify that data corruption is not inadvertently being introduced? Finally, how can one validate how modifying the solution infrastructure will impact overall solution performance? From (supposed) minor firmware updates to entire reconfigurations, must changes be made on blind faith or speculation in a production environment? Or is there a better, less disruptive way?

Load DynamiX: Solving Storage Performance Challenges

Load DynamiX’s product suite provides performance and functional testing solutions to storage technology vendors like Panzura. Its technology combines the industry’s deepest and most accurate emulation of storage protocols and workloads with the highest load generating performance available.




Figure 2: Load DynamiX’s products validate the storage performance of the Panzura cloud storage solution including: Appliance IOPs, throughput, dedupe ratio, and client capacity.

Panzura leveraged Load DynamiX’s load generation appliance to validate performance on the individual components within its product including the client and application access to the appliance, the appliance access to the backend storage, the various networks connecting them all, and finally, a comprehensive validation of the entire solution.

Unlike traditional benchmarking and load generation software, a single Load DynamiX appliance is capable of producing millions of file operations and emulating tens of servers or initiators, as well as thousands of clients, all in a 2U appliance space. This is much simpler than attempting to setup a dozen or more servers to generate a similar sized load. In addition, unlike benchmarks, scripts or freeware emulation, Load DynamiX Enterprise workload modeling software allows the creation of realistic workloads mirroring the type generated by user applications

and usage patterns. The combination of extreme load generation and true to life workload modeling allows an accurate prediction of the Quicksilver controllers’ throughput and client capacity, regardless of file protocols used (NFSv2, v3, v4, v4.1 or CIFS/SMB, v2, v3). This performance validation scenario can then be scaled up to see how it performs under forecasted future load and worst case conditions.

Storage performance validation of Panzura’s product infrastructure

The Load DynamiX load generation appliance can be used to validate the backend object storage as well as the Quicksilver controller itself. Because modeled workload tests can be repeated, the solution architect can make exact comparisons between various cloud storage vendors and validate their throughput as well as the scalability of edge appliances attaching to it. These repeatable tests also mean consistent testing of infrastructure changes and reconfigurations after deployment. The Load DynamiX appliance can also be used to test for data corruption by writing and then reading large datasets through an edge appliance to the cloud storage. Finally, end-to-end latency and throughput may be measured by attaching the load generation appliance to a local and remote Quicksilver controller and emulating a realistic workload through the entire solution. Particular writes to the local device can be tied in to the time it takes to read the same file on the remote device. Benefits include:

- Performance of a variety of public or private cloud arrays may be tested with the same appliance in order to select an optimal vendor and product for the backend storage
- Time spent troubleshooting and optimizing configuration changes is considerably shortened with a portable, repeatable and easily modified workload.
- Performance between any points in the solution with Ethernet connectivity may be easily tested with a single appliance, giving an accurate understanding of how and when data is replicated through the Panzura cloud storage solution.




Figure 3: Load DynamiX validating the combined impact of Panzura’s cloud storage and appliance components.

In addition to testing Panzura's entire storage infrastructure, Load DynamiX can also validate the components and other important aspects of their product suite.

Solving storage performance challenges at the "edge"

Utilizing the native storage array, operating system, or application performance statistics, realistic workloads can be defined and then applied with Load DynamiX load generation appliances to the Quicksilver Cloud Storage controllers measuring latency and throughput. These workloads may then be scaled up to determine maximum client capacity for proper provisioning at a given site.

Accurately measuring compression and de-duplication

Using real data sets or synthesizing data with known patterns with seeded random I/O generation, the deduplication or compression algorithms of the Quicksilver controllers can be accurately tested or bypassed altogether using purely random data streams in order to test the backend storage alone.

Through the emulation of data pattern repetition, the effectiveness of the Quicksilver's caching algorithms can be matched up against a particular data pattern. The control of deduplicability and compressibility of emulated data content on the Load DynamiX appliance can take the Quicksilver's dedupe and compression functions through its paces to accurately determine its bandwidth efficiency.

Addressing translation challenges

Through the emulation of data pattern repetition, the effectiveness of the Quicksilver's caching algorithms can be matched up against a particular data pattern. The control of deduplicability and compressibility of emulated data content on the Load DynamiX appliance can take the Quicksilver's dedupe and compression functions through its paces to accurately determine its bandwidth efficiency.

Conclusions

When using the combined Load DynamiX load generation appliance and its enterprise workload modeling software, architects can design the exact type of global distribution mechanism or tiered cloud storage infrastructure needed by his or her organization. From basic cloud archiving to providing a wide area, low latency mechanism for distributing millions of cross referenced files for collaboration such as used in CAD/AEC or other EDA environments, the solution architect can proceed with confidence through the entire project life cycle. From initial planning, design, procurement and deployment, to troubleshooting, optimization, and non-disruptive change control, they will have the tools and an ongoing change validation process that will provide consistent and realistic workload testing, and accurate, comprehensible analysis with the Load DynamiX storage infrastructure validation solution.



Sales
sales@virtualinstruments.com
1.888.522.2557

Training
training@virtualinstruments.com

Website
virtualinstruments.com