

Making Hybrid Cloud Storage a Reality for Life Science Workflows

Challenges

- **Rapid data growth** – Life sciences data is rapidly increasing every year and traditional NAS storage deployments struggle to keep up with this growth.
- **Storage solution complexity** – Hybrid cloud storage solutions can be complex and often require the installation of proprietary software on all clients.
- **Not impacting user experience** – Tiering done wrong can result in poor end-user experience and force scientists to become experts in data management.
- **Budget busting costs** – Cloud object storage is perfect for inactive data but there has been no way to transparently move data without workflow disruption.

Highlights

- **Immediate ROI** – Real-time analytics automatically identify and migrate inactive data from expensive primary storage to cost effective cloud object storage
- **Zero-impact** – Files always appear local to users regardless of their actual location
- **High-performance** – Accelerates metadata operations for demanding life sciences workloads
- **Massively scalable** – Solution scales to billions of files and petabytes of capacity

Solution

Western Digital and InfineteIO deliver a true hybrid cloud storage solution, accelerating performance and intelligently placing inactive data onto a more affordable tier of storage. Support for billions of files and petabytes of storage with zero disruption.

Reclaim valuable capacity on your NAS and reduce costs as primary capacity purchases are delayed or reduced as inactive data is moved to cloud object storage.

Only Western Digital and InfineteIO can deliver performance with cost savings to meet the demands of storage growth in Life Sciences.

Realizing Hybrid Cloud Storage in Life Sciences with Western Digital and InfineteIO

Data is at the heart of research. The data created by Life Sciences organizations is growing rapidly, out pacing other industries, and is becoming increasingly more difficult to store and manage affordably using traditional Network Attached Storage (NAS). Web-scale cloud infrastructure provides a more cost-efficient and scalable methodology to address this growth, but many of the workflows continue to use files, not objects, making deployment of cloud object storage challenging as an archive tier.

With cloud object storage from Western Digital and InfineteIO – this all changes. InfineteIO installs just like a network switch between research scientists and their existing NAS. InfineteIO accelerates file system performance, and at the same time identifies inactive files on the NAS before automatically migrating them into ActiveScale™, a cost-efficient, reliable and massively scalable disk archive based on cloud object storage from Western Digital.

All files, migrated or not, continue to appear local to clients. Researchers see zero changes to their workflow. For the storage administrator, inactive data is non-disruptively and automatically migrated to a more affordable and reliable tier of storage, freeing up valuable space on the primary NAS for more regularly accessed research data.

Real-Time Analytics as the Basis for Automated Intelligence

An InfineteIO node or cluster installs just like network switches in front of primary NAS storage, and then performs a one-time scan of all files making a copy of metadata. InfineteIO learns hot data on the fly and responds to all metadata requests on behalf of the storage. Metadata is kept consistent with the NAS and there is no need to continuously scan the file system as metadata is always kept consistent and available.

This patent-pending technology is the foundation for real-time assessment of inactive files and instant migration. InfineteIO identifies files that are inactive and can immediately migrate them to ActiveScale freeing up capacity. The return on investment will begin to be realized as soon as InfineteIO and ActiveScale are installed together.

Accelerate Performance Across All Tiers of Storage

A recent review of a Life Sciences customer relying on both InfineteIO and ActiveScale and revealed that 86% of file operations were metadata operations. These operations were driven by workflows specific to Life Sciences, with many small files per directory very common, meaning many millions of metadata operations for even a simple directory listing.

This would typically lead to huge performance bottlenecks in other hybrid cloud solutions that are not metadata-aware. InfineteIO is optimized to serve metadata requests from the network, offloading and accelerating responses. As metadata responses are accelerated, primary NAS performance increases as CPU, memory, and disk contention is off-loaded. InfineteIO accelerates all metadata responses regardless of whether the file is on NAS or

has already been migrated to ActiveScale. InfiniteIO and ActiveScale is one of the only hybrid cloud solution that accelerates metadata performance for Life Sciences workloads and migrates data to a more affordable tier of storage.

Let the Researchers Worry About the Science and Not IT

From installation through to migration and recall, the combined InfiniteIO and ActiveScale solution is transparent to end users with no disruption to any workflow. There is no proprietary software to install on clients or systems, and no stubs, symbolic links or constant scanning of file systems that slows workflows down.

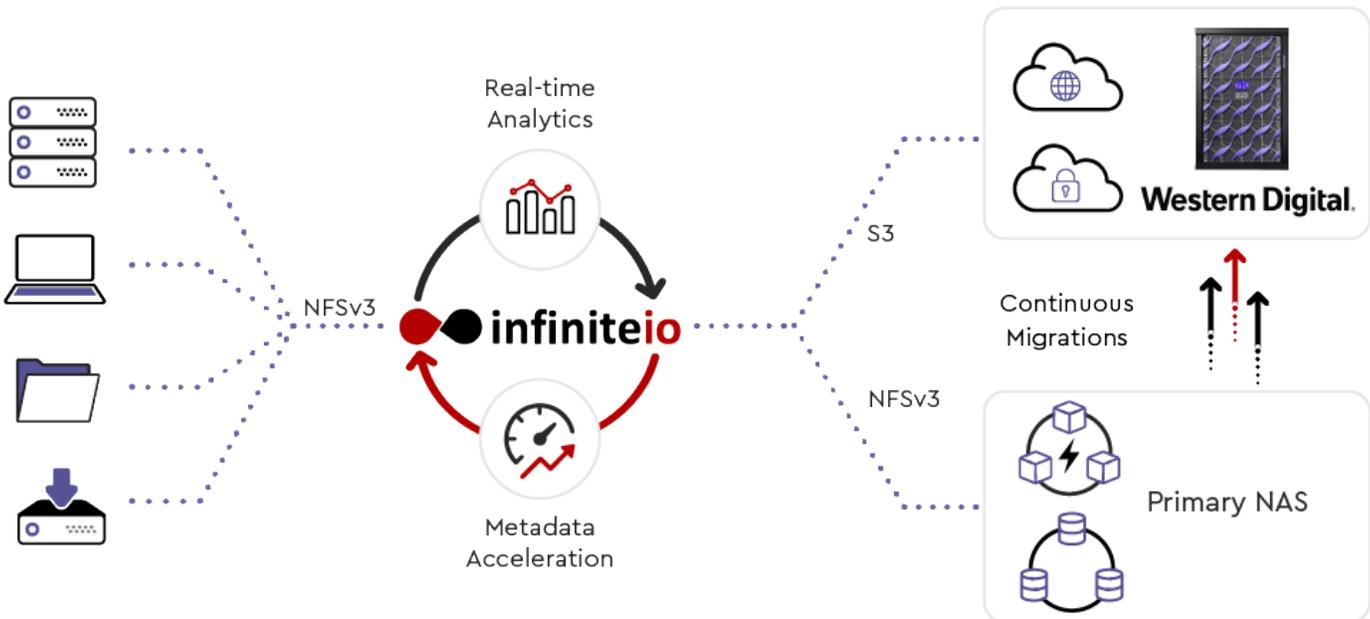
Inactive data is automatically identified and continuously migrated to ActiveScale. Migration is non-disruptive and transparent to users, who will continue to see migrated data as if it were local to them. All metadata operations on migrated data continue to be accelerated by InfiniteIO leaving the data in place on ActiveScale. When a file does need to be accessed, InfiniteIO recalls the file from ActiveScale and places it back to its original location allowing the NAS to respond.

Deploy Your Hybrid Cloud Object Storage Solution Without Compromise

As data continues to grow well into petabyte levels, traditional methods of data protection and IT purchasing models break down. Cloud object storage is ideally suited for inactive data placement with its high durability, low cost per useable gigabyte and huge scale. When deployed together, InfiniteIO and ActiveScale provide a powerful, scalable and cost effective solution for managing data growth.

Conclusion

Western Digital and InfiniteIO provide the only life sciences solution that addresses the unchecked growth of unstructured data, and increases performance with no disruption to existing workflows.



Western Digital.



5601 Great Oaks Parkway
San Jose, CA 95119, USA
US (Toll-Free): 800.801.4618
International: 408.717.6000

© 2019 Western Digital Corporation or its affiliates. 5601 Great Oaks Parkway, San Jose, CA 95119. All rights reserved. Western Digital, the Western Digital logo, and ActiveScale are registered trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. InfiniteIO and the InfiniteIO logo are registered trademarks of InfiniteIO or its affiliates in the US and/or other countries. All other marks are the property of their respective owners. References in this publication to Western Digital products, programs, or services do not imply that they will be made available in all countries.

www.westerndigital.com