



SOLUTION BRIEF

Meeting the Storage Challenge of Exponential Data Growth in Genomics

Highlights

- **Budget Constraints.** Need affordable high-density system to implement an enterprise-class storage cloud
- **Collaboration.** Enable research partners across the country or around the world to share data regardless of device platform
- **Scalability/Serviceability.** Deliver non-disruptive capacity and performance scaling, data repairs and upgrades
- **Availability Requirements.** Ensure easy access to data regardless of where it is stored, how old it is or even if a data center goes offline
- **Complexity Sprawl.** Simplify overall deployment and management of storage infrastructure and operations

Solution

- Western Digital ActiveScale™ is a complete, easy to install, easy to use object storage solution
- Pay-as-you-grow capacity and performance expansion that does not strain the budget
- Continuous access to data regardless of its physical location or the accessing device type
- Reduced time-to-data and improved collaboration with rich metadata tagging and a single namespace
- Consolidated backup, archive and restore processes for simplified management
- Disaster prevention through the geographic distribution of data

The life sciences and healthcare industries are in the midst of a dramatic transformation that will make personalized medicine common-place. Completing the sequencing of the first human genome in 2003 was a key breakthrough that took more than 10 years and cost over US\$3 billion. Since then, the speed at which genomes can be sequenced has more than doubled, easily outpacing Moore's Law. Today's high-throughput sequencing machines can process the human genome in a matter of hours at a cost approaching US\$1 thousand. These advancements have allowed researchers to run more sequencing tests in less time, greatly increasing the pace of scientific discovery. Clinicians, too, have begun leveraging genome sequencing to more effectively treat patients with medications tailored to the patient's unique genetic makeup. The result has been an explosion in genomic data, and organizations are now looking for ways to affordably store and manage this avalanche of data.

Traditional Archival Systems Have Limitations

Archival systems are often comprised of complex disk and magnetic tape storage subsystems organized into performance tiers. Disk tiers rely on RAID for data protection and do not scale to levels suitable for life science applications. System performance can be severely impacted when a single multi-terabyte drive fails. The risk of data loss significantly increases if a second drive fails while the data from the first drive is being reconstructed.

ActiveScale for Genomics Architecture Example

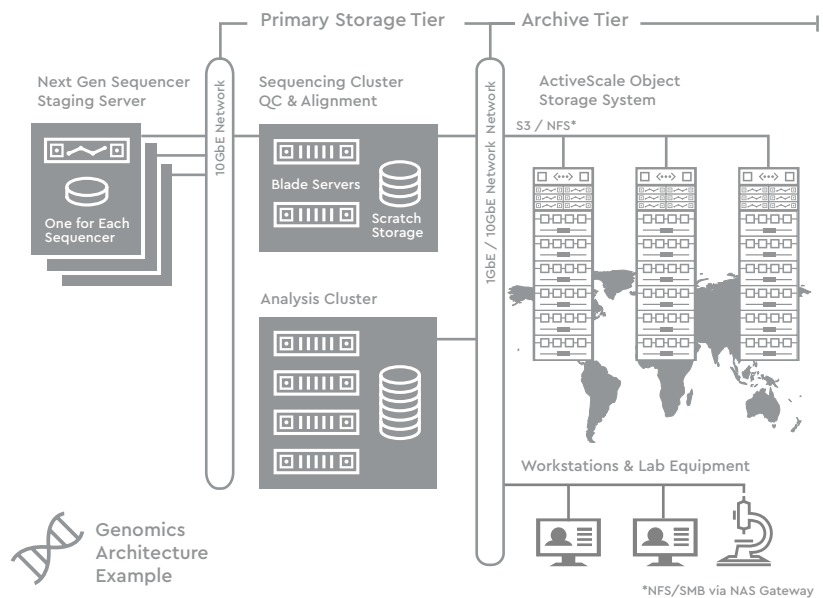


Figure 1. Simplified Enterprise Cloud Storage Infrastructure

As data is used less frequently, it is moved to magnetic tape. Tape archival systems consume little energy and can be cost effective, but the tradeoff is longer retrieval times and added management complexity. Tapes must be stored both onsite and offsite for true disaster recovery, which means a duplicate investment in tape media and ongoing offsite storage costs. While tape systems are designed for long term archival, they offer no means of preventing the natural data degradation that occurs over time, and autoloaders become a single point of failure in the event of a mechanical failure. This means that there is no guarantee that the data will be accessible or readable when it is needed.

Object Storage is an Ideal Solution

Life science involves some of the most complex analyses found in scientific research. It's not surprising that researchers have very unique needs in terms of compute and storage performance, scalability and availability. Clustered file-based storage is frequently used to support the high performance compute clusters used during the analysis phase of the workflow. Utilizing such expensive primary storage for archiving is not practical or cost effective. These architectures must rely on expensive software to replicate data to a remote location in order to provide disaster resilience. Implementing a cloud-scale archive that does not rely on replication software is a much more cost effective solution.

Up and Running Quickly – Easy to Use

ActiveScale object storage removes the challenges of architecting, purchasing and operating cloud-scale storage. As an integrated rack-level system, it's up and running in minutes. Put it in place, connect the power, configure the network connections and it's online, presenting an S3-compliant object interface and global namespace.

Highest Levels of Data Protection and Availability

ActiveScale was architected with industry-leading scalability and data durability as primary design goals. ActiveScale object storage software that ensures valuable research data is well protected and always available. Patented technology delivers 19-nines durability, keeping data consistent and accessible even during a full data center outage in a multi-site implementation. Through background data

integrity checking, data remains bit-perfect and highly available. The system automatically detects and corrects data degradation without user disruption or IT intervention.

Compelling TCO

Research organizations no longer have to weigh the risks of a public cloud against the cost to build a private cloud or design around the limitations of existing storage architectures. Equally important, there is no need to incur the expense or burden of transferring large data sets to the cloud only to pay a second time to retrieve them. Through innovation and integration, the ActiveScale delivers cloud-like scale and efficiency at low \$/TB and watts/TB.

Replication for disaster recovery is common practice with traditional storage systems. ActiveScale reduces the total storage capacity required by distributing a single copy of data across multiple data centers. Eliminating the need for expensive replica copies and replication software licenses can save up to 60% in extra storage costs. Further savings can be realized through the system's use of patented Western Digital helium-filled hard drives, which consume 60% savings in power and cooling than air-filled hard drives.

The fully integrated design delivers the highest capacity per square foot in the industry, allowing you to optimize expensive data center floor space and delay data center expansion. Leveraging high capacity hard drives with HelioSeal® technology, a single rack holds 4.5 petabytes of usable capacity and delivers up to 8GB/sec total throughput. Performance scales linearly with added capacity to meet the needs of large and growing genomic data sets.

Expensive primary storage should be dedicated to the processing and analysis of genomic data, while archiving and backup and disaster recovery can be consolidated into a single archive tier. Figure 1 shows what a simplified two-tier storage infrastructure might look like. By reducing the cost and complexity of managing several tiers of storage and leveraging the pay-as-you-go capacity expansion model, ActiveScale ensures that more money goes to research than to infrastructure.

To learn more visit www.wdc.com/dc-systems

Western Digital.

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

© 2015–2018 Western Digital Corporation or its affiliates. All rights reserved. Produced 6/15, Rev. 8/18. Western Digital, the Western Digital logo, HelioSeal, and ActiveScale are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. All other marks are the property of their respective owners.

One megabyte (MB) is equal to one million bytes, one gigabyte (GB) is equal to 1,000MB (one billion bytes), one terabyte (TB) is equal to 1,000GB (one trillion bytes) and one petabyte (PB) is equal to 1,000TB when referring to storage capacity. Usable capacity will vary from the raw capacity due to object storage methodologies and other factors.