



Fred Hutchinson Cancer Research Center

SwiftStack Enables Cost Effective Cloud Storage in Medical Research HPC Facility

Fred Hutchinson Cancer Research Center (Fred Hutch) is a major medical research center conducting cancer and HIV research. They needed a cost effective cloud storage platform to replace hundreds of isolated pockets and silos of legacy storage. The facility has 1.5 million square feet of offices and laboratories in 13 buildings on a sprawling 15 acre campus.

Like many IT groups, skilled resources are at a premium. Although the entire research center has about 80 IT staff members only 3 were dedicated storage admins who needed to meet the capacity, availability, and round the clock support requirements of 2,700 employees and 220 faculty members in 13 different research programs. And each resource program and department had widely differing requirements and use cases, but only limited budgets and internal IT resources to supplement what was provided by the central Scientific Computing group. Keeping up with the round-the-clock data capture from a large number of research projects wasn't easy.

Challenge

Since data is the lifeblood of research, the researchers needed a reliable but also highly cost-effective storage infrastructure platform to help them acquire, create, exchange, receive, and archive large datasets.

Researchers across campus were struggling with high storage chargebacks in the face of budgetary pressures. To make matters worse, their aging storage arrays had a scalability ceiling that couldn't support the expected 40% growing in scientific data each year, leaving the IT group facing expensive forklift upgrades every few years.

Another major concern for researchers was the need to archive data for far longer than typical in enterprise settings. Research data would need to be accessible for decades, a requirement hardware-based storage platforms could not meet.

The existing legacy storage silos also didn't address one of the primary concerns of any research environment - how to efficiently share data with colleagues. Storage practices at Fred Hutch had often left this valuable data "hidden" away on USB drives, unsecured consumer cloud storage services, or trapped in isolated and obsolete storage silos. Even if the data could be collected, existing solutions like their centralized scale-out NAS are far too costly and complex to handle the current and future data.

All of these factors created a demand for a single pool of storage that could provide more scalability, accessibility, and most importantly, cost efficiency.

USES

- Single pool of storage for research data across the campus
- Seamless connectivity to AWS for cloud bursting
- Feeds data to HPC cluster for analysis
- Backup target to protect thousands of servers and laptops

RESULTS

- 21% of legacy NAS costs, and 48% less than public cloud
- Reduced storage chargebacks by an astonishing 90%
- 600 MB/s throughput compared to 200 MB/s with NAS

HARDWARE

- Heterogeneous server nodes containing consumer-grade 4-6 TB drives
- Nodes spread across 3 buildings on campus
- 10 Gb Ethernet connectivity between storage nodes

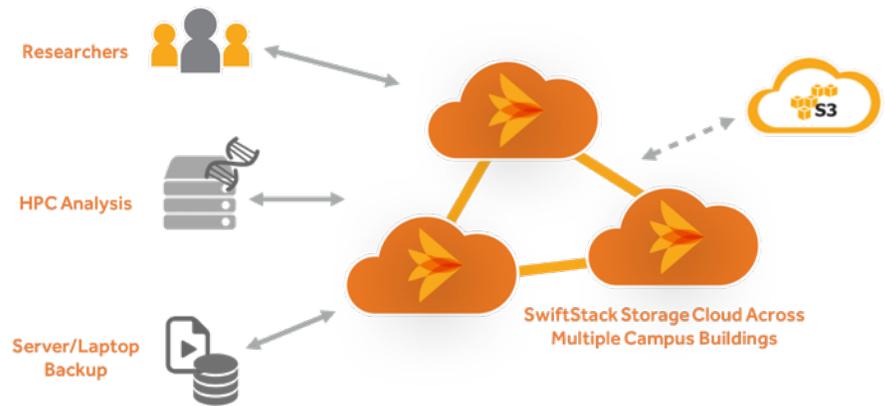
CLOUD FACTS

- Started at 266 TB and has scaled to almost 2 PB
- Researchers use Swift Commander and other tools that were built in-house (<https://libraries.io/github/FredHutch>)

CASE STUDY



"Our researchers needed a new pooled private storage cloud that we could scale up indefinitely and easily. SwiftStack let us deliver it, with low costs and great flexibility, including supporting a bring-your-own-storage model across all of our departments." - Dirk Petersen, Director of Scientific Computing



SwiftStack Solution

Taking all of these requirements and challenges into account quickly lead them first to open source solutions, then to software-defined object storage, and finally to SwiftStack. The ability to select optimal commodity server components allowed storage chargebacks to be reduced by an astonishing 90%. Leveraging SwiftStack's automated deployment and monitoring capabilities, a campus-wide storage cloud was quickly up-and-running and was initially used for infrequently used datasets. Now on this multi-year journey with SwiftStack, the storage cloud at Fred Hutch has grown from 266 TB to almost 2 PB.

Since a SwiftStack cluster can be comprised of mixed models of servers with different drive types and capacities, this enables departments at Fred Hutch to contribute to the private cloud's expansion via a bring-your-own-server model rather than just with budget dollars. SwiftStack's policy-driven functionality allows for department data to be dedicated to their hardware but enables easy sharing with others since the entire cloud has a single namespace.

To make their new storage environment easier to use and more powerful for their researchers, they authored the open source Swift Commander. This is a wrapper for various command line tools that makes interaction with a storage cloud much like a Unix file system, which they were used to. At Fred Hutch, they embrace the use of metadata and provide researchers an easy way to tag their data, making all of their assets fully text searchable.

Since they also utilize Amazon Web Services (AWS), an Avere appliance sits in front of SwiftStack to make their private cloud and public cloud (Amazon S3) environments more seamless. Now, S3-based workloads are available on-premises and SwiftStack-based workloads are in AWS. Also, their server and laptop data protection systems leverage the SwiftStack cloud as a local backup target. This shows the true value of cloud storage, where they were able to start small with a single use and scale the environment to support many more. A drastic reduction in storage cost has allowed them to focus more energy on the critical research-at-hand instead of the supporting infrastructure.

To try SwiftStack for free, go to swiftstack.com/try-it-now.

For additional assistance or to learn more, always feel free to contact us. We're here to help.

Phone - (415) 625-0293

Email - contact@swiftstack.com

Chat - Just go to swiftstack.com and look

for the chat pop-up in the bottom right