

NCSA SELECTS VERSITY FOR ARCHIVING

SOLUTION FOCUS

- On-premises glacier
- Academic sector
- Collaboration

CHALLENGE

- Cost efficient solution for researchers to store data
- Seamless, transparent access to data via file, S3, and Globus for internal and external researchers to archive and retrieve data
- Easy and accurate space usage information
- High throughput, clusterable system that fully utilizes hardware investments

KEY BENEFITS

- Scalable, high throughput S3 object archiving to tape
- Transparently archiving and retrieving both object and file data
- All data is visible from file, object, and Globus interfaces
- Significant cost savings compared to public cloud and legacy internal systems
- High throughput
- Accounting data that enables chargeback

NCSA chooses the cutting-edge Versity solution to manage 60 PB “Granite” data collection.

Overview

The National Center for Supercomputing Applications (NCSA) is a leader in computing, data, networking, and visualization research and services. Established in 1986 at the University of Illinois Urbana-Champaign, NCSA enables a diverse collective of scientists and engineers from around the world to run simulations that process and generate massive amounts of data. This is achieved through funding by the State of Illinois, the University of Illinois, and the NSF along with other federal agencies.

The Challenge

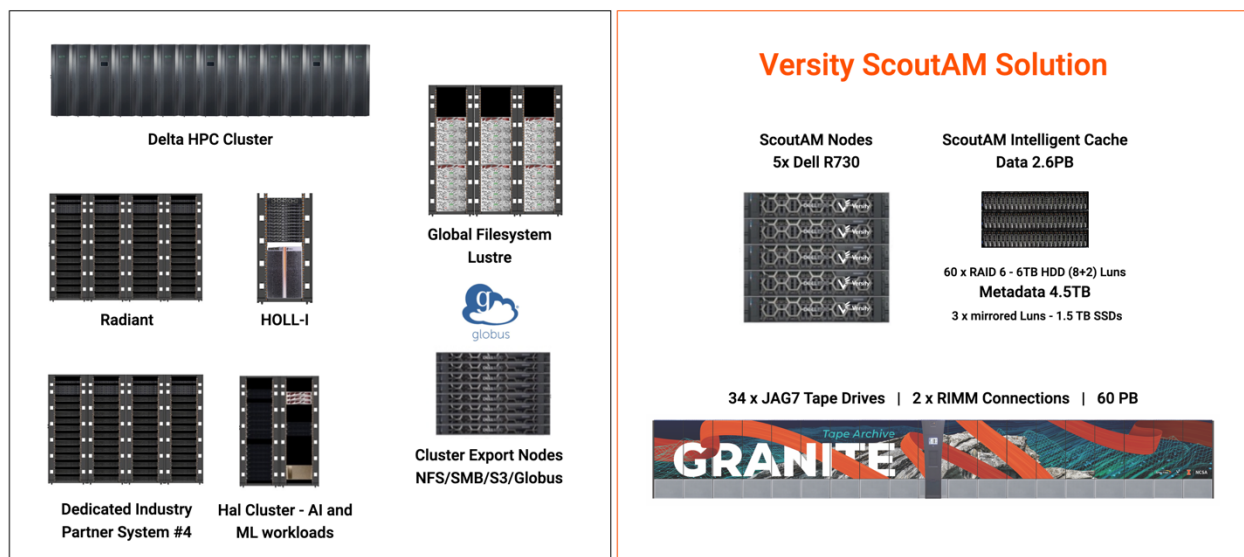
Compute, memory, and data-intensive simulations and calculations are performed every second at NCSA. Scientists and researchers use a multitude of applications and programs to explore everything from artificial intelligence, engineering, and astrophysics to humanities, health, digital agriculture, and environmental sciences. All this work produces a deluge of data that needs to be effectively managed and stored for a diverse set of users and stakeholders across the globe. The first storage for researchers at NCSA is their central Lustre filesystem, Taiga. In an effort to provide an economical solution for inactive and long-term data, NCSA was searching for a scalable, modular, and cost-efficient system to provide researchers capacity for storing vast amounts of data. With such a diverse set of applications and workloads from varying global locations, NCSA required a solution that could support multiple workflows and access methods such as file, S3, and Globus. Based on past experiences with legacy systems and tape, NCSA wanted something easy to maintain, with minimal operational overhead, that could deliver usage and consumption statistics for financial accounting.

The Solution

After a thorough and comprehensive analysis, NCSA chose to partner with Versity to launch their new Granite tape archive capable of storing up to 60PB of data. Versity was able to provide a solution that is cost-efficient and easy to use for their diverse user base. The archive is built using Versity's ScoutAM product, giving scientists a single archive namespace from which to access their data. Access to the Granite system is available directly via common tools and protocols the scientists already know how to use, such as SCP, Globus, and S3. NCSA utilizes Versity's high availability and intelligent archiving features to protect its end users from any hardware failures. Data written to Granite is replicated to two tapes for mirrored protection in case of tape failure. One additional benefit is that the solution increased collaboration among the NCSA team because the system provided a common space for collaboration. The team wasn't fractured among several systems anymore while trying to communicate about projects.

Implementation

A 2.5PB Intelligent Cache runs Versity's ScoutFS filesystem which allows Granite to provide scientists a single archive namespace from which to performantly access their data using common tools and protocols such as sftp, Globus, S3 and NFS. Versity's ScoutAM manages the archive operations with a 19-Frame Spectra Logic TFinity tape library using TS1140 tape technology. ScoutAM's powerful CLI and GUI management interface allows for simple customization of archiving workflows, replication, scheduler management and job prioritization. Granite uses ScoutAM's high availability and intelligent archiving to provide a reliable and performant low-cost storage resource to a variety of different customers with extremely variable workloads and data profiles.



Key Benefits from Versity

With the Versity solution, NCSA receives accounting data that enables chargeback to storage consumers. They can track both the usage and consumption of the Granite storage system so that they can charge a proportionate amount to each customer group. Utilizing these accounting features enables NCSA to offer the archive as a service which funds the archive independently and indefinitely. This way data is not lost if there are any lapses in funding or if funding for a specific system ends. One of the key benefits of the Versity system is scalability, high-throughput performance, and support for file and S3 object archiving to tape. Now, NCSA can access data at higher speeds than ever before. The Versity solution is modular and can be easily scaled up without having to change the fundamental design of the system. ScoutAM is easy for scientists and researchers to use. The system does not require any special training or administrator intervention. ScoutAM transparently archives and retrieves both object and file data. All data is visible from both the file and object interfaces to those who have proper access. Versity's Open-source

metadata combined with open-source data formats gives NCSA complete control over data collections and ensures NCSA will always have access to their data. When compared to public cloud and legacy internal systems, NCSA easily identified significant cost savings. Utilizing a third party to store high volumes of data can cost tens of thousands of dollars per year more than using a private cloud. Another advantage of the Versity solution is maintenance and administrative efficiency. Compared to legacy systems that use dated programming languages and obsolete technology, the Versity solution is easy to install, and easy to operate.

Recap

Versity enables NCSA to archive large volumes of data at high speed with a scalable, high throughput file and S3 object archiving capability. The organization can archive and retrieve both files and objects transparently. The cost savings of implementing the ScoutAM are enormous when compared to saturated public offerings and outdated legacy systems. Obtaining accounting data from legacy systems to enable chargeback to customers was difficult and time consuming. The Versity solution has made this task quick and easy which has been a major time saver for NCSA. NCSA now has a high-performance long term mass storage solution that fully utilizes its hardware investments.

About Versity

Versity is the leading global independent provider of scalable, modular, and efficient mass storage and large archive systems. Organizations leverage Versity's Scale Out Archive Manager (ScoutAM) to implement data preservation strategies for long-term storage, and retrieval of massive data stores, both on-premises and in the cloud. Versity currently manages over 1 Exabyte of archival storage across a global customer base including leading public and private entities in defense, financial services, education, research, aerospace, energy, entertainment, telecommunications, consumer technology, and publishing. Versity is a privately held company founded in 2011 and headquartered in Hood River, Oregon. For more information visit www.versity.com.



Versity gives you clarity and control over your archival storage,
so you can allocate more resources to your core mission.

Contact
sales@versity.com
1-844-726-8826