Research Data Storage for Columbia's Zuckerman Mind Brain Behavior Institute

*An experience/position paper for the Workshop on Research Data Management Implementations**,
September 14-15, 2017, Arlington

Rajendra Bose, Ph.D., Director, Research Computing, Zuckerman Institute
Bruno Scap, Manager, Research Computing Infrastructure, Zuckerman Institute
Columbia University, New York NY
August 21, 2017

Summary

The aim of this experience/position paper is to provide an example of how research data storage planning and strategy has evolved in the first years of a nascent research institute at Columbia University. Ultimately, we plan to complement the reliable but higher cost enterprise-grade storage system we initially procured with future tape-based "deep archives" for long-term stores of imaging and other neuroscience-related research data. These future archives will be enabled by maturing object storage technologies and our existing tape infrastructure.

Introduction

The Mortimer B. Zuckerman Mind Brain Behavior Institute assembles a team of neuroscientists, engineers, statisticians, psychologists and other scholars from across Columbia University. Zuckerman researchers explore how the brain develops, performs, endures, and recovers to gain critical insights into human health and behavior. The Institute is being brought to life in the new Jerome L. Greene Science Center—a state-of-the-art facility on Columbia's new Manhattanville campus which will house up to 1000 research faculty, staff and students in 2018.

Planning for research data storage and computing resources for the Institute commenced over five years ago. Early activities included rounds of interviews with a subset of the roughly 50 total faculty principal investigators to gather information about the research workflows and data flows, computing environments and technical support needs of their laboratories. In addition, a daylong visiting advisory committee arranged in 2013 featured scientific computing representatives from eight regional institutions.

Guidance provided by the advisory committee addressed long-term planning strategies for a dedicated, on-site research computing data center, and underscored the importance of designing a robust file system to provide the foundation for what could be multiple iterations of local computational clusters in the first years of the Institute.

An initial system size of two petabytes (PB) of storage, consisting of several different performance levels, was based on requirements estimated from further faculty and lab interviews during 2014-2015. This initial size, plus growth trajectories of 2-3 PB per year experienced by several peer institutions over the past few years, informed the scope of a request for proposal (RFP) process from vendors during 2015-2016 for the acquisition of a storage system suitable to meet the immediate needs of the Institute as well as provide rapid, multi-petabyte growth if needed during its first years of operation.

Our storage system was purchased and installed during 2016 as labs began to relocate from elsewhere into the newly completed Jerome L. Greene Science Center building. The storage system includes tape infrastructure to allow for a tape-based offsite disaster recovery solution. A dedicated Research Computing team includes two infrastructure personnel that are tasked with managing the data center, planning and provisioning the storage service, and designing and implementing additional, future services provided to Institute researchers.

Advanced Imaging Technology Storage Needs

In addition to the Research Computing team, a number of other scientific resource groups within the Zuckerman Institute support the research activity of faculty and their labs. Two of these groups in particular, Cellular Imaging and Magnetic Resonance (MR) Imaging, will be close partners that drive our storage strategies and our plans for future sustainability.

The Zuckerman Institute *Cellular Imaging* team provides access to state-of-the-art imaging instruments, such as microscopes that can scan large groups of cells in an instant with a sheet of light. The Cellular Imaging facility provides imaging and image analysis expertise, and develops new techniques for illuminating the workings of cells. Specialized instruments run by this group will collect, process and analyze an estimated one terabyte (TB) of imaging data a day in the near future. These anticipated large file sizes—orders of magnitude larger than current practice—reflect vastly improving image resolution.

Proposed projects will deploy new technologies and approaches to observe, map, and analyze neural activity. The Research Computing and Cellular Imaging groups are investigating software for managing, visualizing and analyzing microscopy images and associated metadata, and will draw upon the practices, formats and resources established by colleagues in the maturing neuroscience data community. The ultimate goal of Zuckerman researchers is real-time sharing of cellular imaging data and algorithms with Institute and external colleagues across multiple sets of experiments.

Similar goals hold for the data generated by the *MR Imaging* group. This team includes engineers, physicists and systems operators who work with researchers to design and implement new experiments using MR technology and assist with the analysis and interpretation of resulting data. Research Computing group is assisting the MR Imaging group to install and configure both vendor-supplied and open source software for MR imaging data capture, processing and analysis used for research studies.

As with cellular imaging, requirements for MR image data storage will increase in coming years. Today, typical processing pipelines for raw data for an individual subject study results in a final data product on the order of 50 gigabytes (GB) in size. Multiple scans per study, multiple studies per day and multiple MRI machines in operation translate into MR imaging requirements of 0.5-1PB over the first year or so of operation.

The Research Computing group, in concert with the imaging teams, are fostering close working relationships with our labs to introduce more nuanced data management strategies in traditional workflows within labs, which often do not distinguish between active, archive, disaster recovery and other levels of storage.

Evolving Data Storage Technology and Design Strategies

The relationship between data storage price, performance and technology is often relatively straightforward. On one end of the spectrum, one can purchase an enterprise-class data storage system that is reputable and stable, but expensive. On the other end, the data storage solution could be built in-house with more cost effective, off-the-shelf hardware components, and open source software. These options are not mutually exclusive – a number of different storage systems can co-exist, from the low cost / low performance to high end, more expensive, and well supported solutions. When the requirements for data storage capacity are not very demanding, these options continue to work well.

Because the Zuckerman Institute research computing infrastructure needs the capability to potentially store and manage tens of PB of data in the coming years, however, solutions beyond the traditional storage options need to be considered. Design points guiding our strategy include:

- Supporting data sets consisting of very small and very large files
- Managing the primary storage back up for large amounts of data
- Expanding data storage capacity without large interruptions to service and high costs
- Designing for flexibility to enable future integration with new technologies and the addition of data storage capabilities
- Minimizing the vendor lock-in

We know our researchers value leading edge technology tools, and need to be able to deploy the latest software tools quickly and efficiently. Some also develop their own software. Centralized infrastructure, on the other hand, changes at a more gradual pace. Budget constraints and the need to carefully plan the integration and migration of technologies increase the time required to design and implement technologies that will work for everyone in the Institute.

Our centralized infrastructure will need to support many different workflows. The storage technologies implemented today should ideally consist of standard, open and flexible components to enable future capabilities and services that may not exist now, but will be required in the future.

Storage technologies are constantly evolving and at the same time the prices are dropping. Changes are causing some storage performance tiers to collapse, and new options are becoming available as current ones are no longer able to keep up with the large amounts of data. Object storage technology, only suitable for early adapters over the past few years, has matured into a suitable option for storing large amounts of data.

Vendor choices and competition are also increasing the cost effectiveness of storage. The rapid development has delivered new features and benefits, as well as increased stability. Start-ups are not the only players any longer, as larger, more established companies like IBM, EMC, NetApp, HPE, Hitachi, Red Hat and others have begun to offer object storage products. There are a variety of private on-premise, public cloud and hybrid solutions available.

Concluding Remarks: Plans for Sustainability

With the changing vendor landscape, it is important to decide whether proposed solutions are temporary trends, or core technologies that will persist over time. For centralized storage environments, the potential for integration is a critical consideration. Vendors often limit use of products outside of their own offerings. But most data storage infrastructure will include a number of different vendors and technologies. Integration between different storage products and the compute infrastructure is critical when providing research infrastructure capabilities as a service.

After consulting with faculty, we are planning for a future research storage infrastructure that will maintain a relatively small footprint of hardware with performance suitable for active, day-to-day research data storage, while we expand our object storage capabilities. Active storage is the primary storage provided by NAS and configured with three tiers: performance, general use, and longer-term infrequent access.

The longer-term infrequent access tier will be replaced by the object storage tier. Object storage systems enable storing large amounts of unstructured data and scale well. In addition, we will investigate a data storage deep archive tier based on object storage technology and going directly to tape.

The imaging and other research data generation and processing activities within the Institute will guide our data storage strategy, but will present a challenge for our storage planning as we serve the aggregate research and collaboration needs of our researchers. Some research groups expect to keep their acquired data in perpetuity, because they plan to analyze previously collected data in new ways in the future. Certainly, the ability to interpret historical data collected over many decades could bring unique insights as new discoveries are made in the future. We believe a large volume of deep archive storage, coupled with a smaller, stable footprint of higher performance storage for active lab use, will meet the future needs of Zuckerman Institute researchers.