

# Research Computing Infrastructure Requirements: immediate, and longer term

---

Harry Mangalam, Research Computing Support, OIT  
version 1.0; November 16, 2015

## Table of Contents

- [1. Storage.](#)
- [2. The Problem](#)
- [3. The Recommendation](#)
- [4. The Rationale](#)

**This is a strawman version intended to provoke discussion. It is not the final version.**

## 1. Storage.

---

This is a summary of a [more extensive description of the problem](#), but that document provides no explicit recommendations. This one does.

## 2. The Problem

---

As machines of all kinds become more digital they produce more digital data, and the techniques to digest and analyze this data require additional storage as well, often 10X the raw data. Especially in research, this leads to problems in handling that research for people who are experts in their fields, but not at Information Technology. Therefore it is up to the organizations that do understand this to provide solutions that provide the optimum balance of speed, flexibility, cost, scalability, reliability, and especially, easy of use.

## 3. The Recommendation

---

UCI should immediately provide a local storage system similar to [Perdue's DataDepot](#) of 120TB (2 half-populated storage servers + metadata server + Input/Output node (IO node)), where local researchers can rent local storage for:

- active files
- short term backup
- web distribution of files
- sync & share with collaborators

OIT should buy the required chassis hardware with the actual disks being paid for by storage rental (at a slight premium to cover the next chassis and maint costs). The cost for raw disk is currently ~ \$50/TB for disks of decent reliability.

This storage can be accessed in a variety of ways via the [caching](#) IO nodes that provide the actual services, so that the users rarely access the filesystems directly.

The filesystem should probably be a commercial appliance, but the IO nodes can be run by OIT to provide the services required by researchers. The IO node(s) should provide CIFS & NFS access, web services.

Optimally, the same storage would support spinning disks for bulk storage, SSDs for complex access, [encrypted volumes](#) for data needing extra security, parallel access for high-speed RW, and support for [SMB/CIFS](#), [Network File System](#), [WebDAV](#), and other protocols as needed.

Users who require the very fast parallel access to the filesystem can purchase the specialized clients to do so. Such clients could be the compute clusters on campus or specialized machines at the microscope facility.

## 4. The Rationale

---

From both internal and external surveys, by far the most critical resource that faculty desire is storage in various forms. Storage is needed for actively used files, short-term backup, and long-term archives. This storage needs to be *medium to high performance* on many metrics, from streaming reads & writes (RW) as in bioinformatics & video editing, to small high-jitter RW operations with rapidly changing offsets (relational databases, access to zillions of tiny files (ZOTfiles)).

Researchers particularly need storage for terms appropriate for their publishing cycles which is typically 1-2 years in the life sciences, longer in some social sciences. They also need backup for these files, since their loss can often terminate the project with the fiscal loss to the campus often in the range of \$10K to multiple \$M. The administrative contribution to this system is provided by the overhead of successful grants, themselves considerably assisted by the availability of this storage.

Research now often involves access to these files both on campus and off, by lab members and external collaborators. Some of this data is restricted by intellectual property agreements or privacy concerns and therefore requires special protections of permission, local firewalls, and even on-disk encryption.

Commercial cloud storage can provide some of these resources, and therefore offload the cost of technical oversight. Services like Amazon Glacier or Univ Oklahoma's Petastore are hard to beat for archiving data, which would otherwise require a ~\$100K investment in hardware. For unrestricted data that has been published and no longer has to be accessed quickly, cloud archiving is the recommended solutions.

However, locality is attractive for data that has critical latency, bandwidth, security, and legal/ownership requirements. Even peering with another UC campus (such as UCLA's CASS) cannot provide the kind of resources that many local users require. For example, it cannot provide untunneled SMB/CIFS file services to local users due to the insecurity of that protocol and the encryption adds even more latency to the connection. Additionally, any operation that requires frequent communication between client and server will take longer with the longer packet round trip times due the increased number of network hops. While opening a single small file for editing is hardly more noticeable from CASS, unpacking a 100MB archive from a client at UCI to CASS at UCLA takes about 700X longer than the local operation.

The latest version of this document [can be found here](#).