

(Old) Archival Storage

The Center for High Performance Computing includes Research Object Storage System (nicknamed ROSS) as a low-cost option for archiving data. Although there is a fee for using the object store, the cost of the storage is heavily subsidized by the university, according to current plans.

Each HPC user can request private space on the object store, tied to their HPC account, for personal use, for example to back up their home directory on Grace's cluster storage, or to hold research data sets when they are not being actively used. In addition, labs, departments, or other entities can request shared space that can be utilized by all of their users. ROSS is currently based on the Dell EMC ECS (Elastic Cloud Storage) system. The ECS system is now available for use, and is jointly managed by the HPC administration team and UAMS IT (i.e. UAMS IT runs the data center infrastructure in which ROSS sits and the networking infrastructure needed to access it, while the HPC team manages access to the ECS itself).

- Per the plan set out by Dean Morrison, archival storage on ROSS may be purchased for \$70 per TB, with an expected lifespan of 5 years, with no free storage available. Note that this cost is considerably less expensive than commercial-cloud-based archival storage (e.g. Amazon Glacial storage).
 - The \$70/TB charge is for the amount of storage reserved (i.e. the quota limit), and is not based on actual use.
 - Unlike some cloud providers, UAMS does not impose additional data transfer (access, egress or networking) fees for accessing the archival storage.
 - Please coordinate Robin at DBMI to arrange an IDT to the storage core to pay for storage quota requests.
- Once you've picked the size of your storage pool and have arranged the financial details with Robin, one of the administrators will set up a namespace and quotas for you.
- As part of this setup, you may designate one or more namespace administrators who would manage users of and set permissions for your namespace, as well as create and maintain buckets.
- Although the EMC ECS storage pools have 12+4 error coding (i.e. redundantly stores data) to protect data against failures, there currently is no offsite backup. (We are actively working to rectify this.)
 - Users who need offsite backup could, for example, send backup copies to Amazon Glacier, Box, or similar systems, with the hope of never having to ever retrieve them except in dire circumstances. However, the users are responsible for the off site backup costs.
 - We are looking into an option that would allow researchers to send copies of their archival data to the NFS-sponsored OURRStore project, a write once, read seldom research archive at the University of Oklahoma in Norman. If this pans out, users would only be charged for the media (currently LTO-7/8 tapes, at least 2, preferably 3). Tapes with 9 TB capacity currently run \$50 to \$75 each (i.e. \$100 to \$225 for a set of 2-3, or about \$12-25 per TB unformatted). The expected lifetime of this media is 15 years, with a minimum expected lifetime of 8 years for the equipment needed to read the media. Taking into account the media lifetime, this storage is about one third the cost of storage on ROSS. However, OURRStore will not be in production mode until mid 2021, according to the current schedule.
- Users who need automatic offsite backups or better file performance can still request space on the Research NAS that UAMS IT manages (i.e. the EMC ECS system is not the only game in town).

Once you have your namespace in play, here are some hints for using it, in addition to the [technical info about the object store, explained in a separate article in this wiki](#).

- Objects in the archive are stored in buckets that belong to namespaces. Your namespace administrator may create as many buckets as desired within your namespace. Note that it is recommended that a namespace have no more than 1000 buckets.
 - On an ECS system like ROSS, a bucket may be accessed with either the S3 or Swift protocol interchangeably. Of course, certain features available in one object storage API might not be available on the other.
 - On bucket creation, your namespace administrator can also configure a bucket for file access (in addition to object access) using either the NFS or HDFS protocol. Changing the file access option after bucket creation currently requires re-creating the bucket and copying data from the old to the new bucket. Note that you do not lose object access by enabling file access, but enabling file access on a bucket may have some minor impacts on object access. Note that there are other, potentially faster methods for POSIX-style file access to object storage that do not depend on enabling the NFS/HDFS access built into ROSS.
 - The ECS system also offers EMC-proprietary bucket formats (Atmos or CAS) which we are not actively supporting, and which do not offer cross-head support (i.e. you can't access them with other protocols like S3 or Swift). They are there for compatibility purposes for older systems/software, if needed.
 - It is also possible to enable CIFS, or SMB access to a bucket set up for file access. (CIFS / SMB is often used for Windows directory shares. But since CIFS/SMB access goes through a secondary server, performance likely suffers, so we are not recommending it for heavy use. There are tools available that can allow Windows users to access buckets as if they were mounted file systems.
- Once your administrator sets up the bucket, object users designated by your namespace administrator may use the object APIs, and if you enabled file access, mount the buckets to any system inside the UAMS firewall.
 - It currently is not possible to access the EMC ECS system from outside the UAMS firewall.
- Grace, the HPC, has data movers that can assist in staging data to/from Grace's cluster storage for running HPC jobs. If you plan on using this feature, please discuss with HPC staff to get it set up for you.