

Quasar archive for data collections

The Quasar archive is a cold, tape-based archive for storing curated data collections that have an indefinite lifetime. It is not designed to serve data or to store data that will be frequently accessed, overwritten, or deleted. (Active data should be on GLADE or Campaign Storage rather than on Quasar.)

Before requesting access, please review the following information regarding how to archive files, the minimum and maximum file sizes, and related use policies.

Page contents

- [Storing data](#)
 - [File size requirements](#)
 - [File reads and data/metadata change frequency](#)
- [Disaster recovery](#)
- [Policies](#)

Storing data

Users store data on Quasar by transferring files via the Globus mapped collection named **NCAR Quasar**. For documentation about how to use Globus, see [Globus file transfers](#).

A note about verifying Globus transfers

Using the Globus **checksum** sync option when transferring files can result in "operation timed out" error messages when it causes file recalls from tape, which can be slow.

To avoid such errors when doing an incremental backup, use a different sync level – **exists**, **mtime** or **size**, for example – when making the transfer.

To verify that a just-completed transfer did not encounter any corruption, do the checksum immediately to complete it before files are transferred to tape and purged from the disk cache.

File size requirements

The following requirements apply to files stored in a project's high-level allocation (RDA or EOL, for example).

- The *maximum* file size is 5 TB.
- The *target* file size for Quasar is 1 GB or larger, so use a tool like tar to combine multiple smaller files into a larger file or files before storing them.
- At least 90% of a project's files must be at least 100 MB.
- Up to 10% of a project's files may be smaller than 100 MB.
- The length of a file name and its full path name cannot exceed 1022 characters.

File reads and data/metadata change frequency

The system is designed to support large file writes effectively. As a tape-based archive, however, it is not designed to support frequent read activity. File reads should be infrequent, and data and metadata changes should also be rare.

Under normal operational use, no more than 10% of your files should be read, rewritten, renamed, or deleted during any 12-month period. If a special case arises – a recovery operation, for example – and you anticipate more activity, please contact the [NCAR Research Computing help desk](#).

Disaster recovery

Disaster recovery storage is available to approved projects. When a disaster recovery account is approved, a secondary directory tree is made available for users' data

The data are written to a separate pool of tapes from the primary data copies. The disaster recovery tapes are moved from the TS4500 library as they fill up and are stored in a fireproof vault in Cheyenne, Wyoming. See [Quasar system specifications](#) for details.

Policies

- The system is not backed up.
- Vendor support for the system is 9 a.m. to 5 p.m. next business day, so problems that occur outside of those hours may need to wait to be resolved.
- CISL does not enforce file size at writing time, but when files smaller than the minimum size are found on the system, you may be asked to relocate the holdings to more appropriate storage such as the NCAR Campaign Storage file system or Stratus object storage system.
- If excessive read, rewrite, or metadata change activity is detected, you may be asked to relocate the holdings to more appropriate storage such as the NCAR Campaign Storage file system or Stratus object storage system.