

Minerva Archival Storage Policy 1.0
Scientific Computing and Data
Icahn School of Medicine at Mount Sinai
August 28th, 2024

The Scientific Computing and Data team offers archival storage capabilities for infrequently used large files that need to be stored over a long period. This archival storage service is distinct from a backup service. A backup service is optimized for making a snapshot of files that are currently in use and change frequently.

On the Minerva HPC system, Tivoli Storage Manager (TSM) is the archival storage solution. When you store a file in TSM, one copy is written to tape and remains on campus. A second copy of the file is written to tape and transferred to an offsite location. All tapes are encrypted. The quality of the tape may degrade over time, and files may not be recoverable.

TSM is a shared system. There are two activities that cause extreme delays for all users: (1) archiving files <1 GB in size, and (2) archiving or retrieving >40 terabytes per week.

To ensure that all users receive consistent performance to archive and/or retrieve files, the following policy is now in effect:

- 1 GB is the minimum file size to be transferred to TSM. Files <1 GB must be aggregated with tar/zip, resulting in a file size of 1 GB or larger.
- The maximum amount of data archived or retrieved from TSM is limited to 40 terabytes per week.
- TSM is for long-term file storage only, with each file stored once.
- Data archived in TSM has a retention time of six years. We will notify you by email three months in advance before your file will be deleted from TSM.
- Failure to adhere to this policy may result in your archival or retrieval processes being killed, and may ultimately result in removal of your access to TSM.