Granite Usage Guide

Archive Architecture

The archive is essentially made up of four main parts, the Granite servers, the disk cache, the Library, and its export nodes. The Granite servers connect to the library via multiple 2Gb FC connections, the cache via 12G SAS, and the NFS Export nodes via 2x100 GbE.

Block Size on Cache

File System Block Size: 1 **MB**

For a balance of throughput performance and file space efficiency, a block size of 1MB has been chosen for the *ScoutFS* file system that acts as the Granite front end disk cache.

Section Size on Tape

The ScoutAM software allows for section sizing media into chunks to mitigate writing small files to tape. We currently have a section size of 10GB.

Accessing Granite:

There are currently three different methods that will be implemented to store and retrieve data, Globus, SCP, and a DR Solution the SET team provides.

Globus:

Globus is a third party service NCSA uses to leverages across a number of it's storage systems.

Granite is accessible via Globus at the endpoint name "NCSA Granite" and the endpoint is open to the public internet for transfers from anywhere with another Globus endpoint. Authentication to the endpoint is handled by NCSA's CILogon service and requires two-factor via Duo. For shared collections, or other questions submit a ticket to help+globus@ncsa.illinois.edu. Information about Globus can be found at their site https://www.globus.org

SCP:

Users will be able to target granite-scp.ncsa.illinois.edu from any public or private-routed resource within NPCF to send data to or retrieve data from tape. Any retrieval or storage from non-NCSA vetted resources will require 2FA authentication.

DR-Backup:

NCSA's Storage Team has developed and maintains DR software to facilitate the backup of large chunks of POSIX file systems, especially those that see lower levels of churn (eg. not /scratch). Projects that leverage and pay for these DR services are able to use Granite as a backend target for these backups and the software has been configured to optimize storage on *Granite*.

Usage and Quota controls

Capacity to inode Ratio

Ratio: **1TB Quota to 10,000 iNodes**

An inode is a record that describes a file, directory, or link. To ensure that good streaming performance of the tape archive subsystem and increased performance for file recovery this quota is enforced for all projects.

Controlling Quota

Quota is currently controlled *physically*. ScoutAM and the filesystem it's based off (ScoutFS) currently only allows to assign a pool of tapes to users /groups. Each tape is 4TB so quota's will be set in 4TB chunks up to and slightly over (1-3GB) of your allocated quota. We intend for this to be controlled more precisely at the filesystem level in the near future.