


# Granite Architecture

**Note: Granite's Disk Cache & Movers are expected to be upgraded during its September 2022 PM window**



## Components

### Tape Library

Granite is made up of a single Spectra TFinity library. This 19-Frame library is currently capable of holding ~32PB of replicated storage within the unit, leveraging 30 TS1140 tape drives to transfer data to over 9000 IBM 3592-JC tapes.

### Granite Mover Nodes

Four nodes currently form Granite's server infrastructure; Each node connects over direct Fiber Channel connections to 8 tape interfaces on the tape library (30 Drives - 2 Robotic controllers).

Each Granite node is also connected at 2x100GbE to the SET aggregation which also allows a 2x100GbE the NPCF core network. This is also leveraged as the NFS mount point that will server Globus data movers and other mount points to access the archive.

### Disk Cache

The archive disk cache is where all data will need to land to be ingested or extracted from the archive. This is currently made up of three NetApp E6000 controller units and one expansion tray, each using 60 6TB drives. One controller unit operates as the storage device for archive metadata, and the other two controllers and expansion operate as the intelligent disk cache. These are redundantly connected to the Granite nodes via 12G SAS links.

### Cluster Export Nodes

The tape archive is mounted via NFS on to the Globus export nodes directly so they have direct access to the archive. Since Granite shares its export nodes with *Taiga*, this allows for quicker and more direct Globus transfers between *Taiga* and *Granite*.