

Xrootd Administration

Unknown macro: 'html'

- [XRootD Access](#)
- [Fermi XRootD Cluster Setup](#)
- [Server Setup](#)
- [Authentication and Authorization](#)
- [Tasks](#)

Unknown macro: 'html'

XRootD Access

Two xrootd setup are used for Fermi. The *system-test* xrootd consists of a single data server that provides access to some nfs and test data directories. The main Fermi xrootd cluster contains multiple PB of disk space and holds all of Fermi's data. It is used fro reading and writing files. There are three entry points to this cluster. The *production redirector* is used by users and production. The *test redirector* is used to test new xrootd versions and configuration. It runs on the same data servers as the production xrootd. A proxy server is also available that allows writing files from remote sides (e.g.: IN2P3) to SLAC.

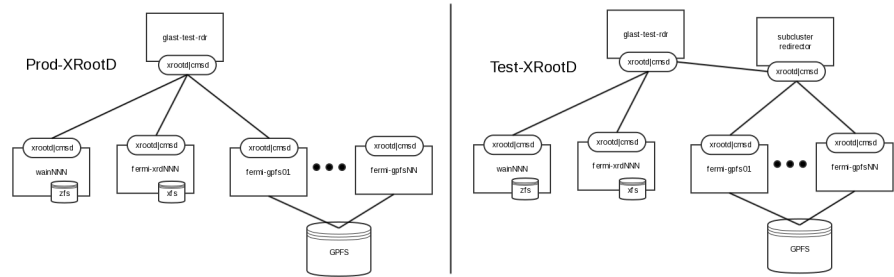
type	host	alias	subcluster	admin	comment
system tests	fermilnx-v07	glst-xrootd01		taylor	
redirector (production)	fermilnx-v02, fermilnx-v12	glst-rdr	fermilnx-v01, fermilnx-v03	taylor	
test redirector	fermilnx-v03, fermilnx-v06	glst-test-rdr	fermilnx-v03, fermilnx-v06	ansible	xrootd name: gltst, gltg (for subcluster)
proxy	fermilnx-v06	glst-xrd-xfer		taylor	

For the Fermi servers see: [NFS/GPFS and Xroot Disk Assets](#)

Fermi XRootD Cluster Setup

The production xrootd cluster consists of a set of Solaris servers (wainNNN), Linux servers with a local files system (fermi-xrdNNN) and xrootd servers that access a GPFS file system. The current setup for the production xrootd use only one of the gpfs servers connecting directly to the redirector. As shown in the figure each fermi-gpfs server has access to the whole GPFS space so that all see the same files. In this case if all fermi-gpfs servers would connected to the redirector files on the GPFS file system would show up as multiple copies.

In order to handle the shared file system a subcluster was introduced in is currently used by the test xrootd. The subcluster redirector is aware that the xrootd data server export a shared file system. The setup is shown in the figure below. Clients still connect first the the main redirector (glst-test-rdr). If a file is on GPFS the client will be first redirected to the subcluster redirector which subsequently redirects it to one of the gpfs data server.



Server Setup

- [Remove server permanently from cluster](#)
- [Prepare xrootd for an outage](#)
- Outage for XrootD redirector

Authentication and Authorization

Access to the Fermi Xrootd cluster requires the users authentication. The authentication and authorization is based on the users name and uses the xrootd unix-authentication module. The authorization information, which directory path a user can read from and write to, is kept in a file that a xrootd server reads and periodically checks for updates.

- All Fermi users have read access to all files
- Special productions accounts (glast, glastraw,...) have write access (to all path, including the science-group directories)
- Science groups directories in */glast/ScienceGroups/<groupName>* are writeable by all users that are associated with this group
- There are a few extra non-Fermi users that have read access granted (for example myself)

The authorization file is created in two steps:

1. The first step is a cron job that creates the auth-file from the fermi users list and the list of users in the science groups.
The program *ScaAuthfileUpdate* creates the auth-file and runs as a trsrontab (as wilko).
2. Each data server runs a cron job that updates the local auth-file that xrootd uses if the master file (from step 1) has been updated.

Tasks

- [Purge Old Recon Files](#)