

xrootd space

There are two high performance xrootd spaces - The Tier 2 space and the Tier 3 space. These space are primarily for storing ROOT files and ATLAS datasets. Creating a new file or directory on a Xrootd space takes at least 5 seconds. So they are not good for large number of small files such as logs, codes, and small text files.

On ATLAS interactive nodes atlint0[1-4] (and only on them), one can see these spaces as mounted filesystems:

```
$ ssh atlint01.slac.stanford.edu
$ df -h /atlas /xrootd/atlas
Filesystem      Size  Used Avail Use% Mounted on
xrootdfs        741T  682T   60T  93% /atlas
xrootdfs        2.8P  1.9P  893T  69% /xrootd/atlas
$ ls -l /atlas
total 15
drwxrwxrwx 1 daemon daemon 140 Oct  7 11:43 dq2/
drwxrwxrwx 1 daemon daemon   6 Oct 25 2012 group/
drwxrwxrwx 1 daemon daemon 4096 Oct  7 12:27 local/
drwxrwxrwx 1 daemon daemon 4096 Dec 23 2014 output/
drwxrwxrwx 1 daemon daemon 4096 Aug  4 11:28 proof/
drwxrwxrwx 1 daemon daemon   37 Jan 23 2012 test/
$ cat /atlas/dq2/site-size
{"sizes": {"total": 813980054454272, "available": 65063350558720}, "time" : "2015-11-18T19:30:01"}
$ cd /atlas/local
$ rucio download ...
```

/xrootd/atlas is the Tier 2 space. /atlas is the Tier 3 space. Avoid using atlint03 for large number of filesystem access.

Tier 2 space

The Tier 2 space is for Tier 2 only so users should not write to it. However, SLAC ATLAS users are encouraged to use [R2D2](#) to transfer official ATLAS datasets to the WT2 (the SLAC Tier 2). The corresponding DDM endpoint names are SLACXRD_*DISK spaces (space token: USERDISK, GROUPODISK, LOCALGROUPODISK, SCRATCHDISK). Once they are at the Tier 2 storage, they are accessible (readonly of course) from all SLAC interactive nodes and batch nodes.

- On atlint0[1-4] this space is mounted at /xrootd/atlas/
- In your batch jobs, you can access them via the xrootd protocol (explained below). Files in xrootd have a URL-type format. The accessing URL is root://atl-xrdr//xrootd/atlas... (/atlas/xroot/XXX also works)

Tier 3 space

SLAC ATLAS group owns a private cluster. It is actually a batch cluster and a xrootd storage cluster (the Tier 3 space). It was also a proof cluster but the proof function is obsolete.

- On atlint0[1-4] the xrootd storage cluster is mounted at /atlas .
- You can use [R2D2](#) to transfer official ATLAS datasets to the space under directory /atlas/dq2. The Tier 3's name in R2D2 is SLAC-ATLAS-T3_GRIDFTP.
- ~~Files under /atlas/proof and /atlas/output are used by the proof jobs.~~
- Space under /atlas/local are for users to read and write.
- In your batch jobs, you can access them via the xrootd protocol (explained below). The accessing URL is root://atlprf01:11094//atlas/...

Access the xrootd spaces

To read input data in batch jobs, we recommend either directly read the ROOT files using the xrootd protocol, or by copying files (ROOT file and non-ROOT file) via "xrdcp" to batch nodes' /scratch space.

For output files, write the output to batch node's /scratch space first. At the end of your batch job, use "cp" to copy them back to the NFS space or "xrdcp" to copy them to the Tier 3 xrootd space.

To read an xrootd files in Athena using the xrootd protocol:

```
filelist = []
filelist += ["root://at1-xrdr//atlas/xrootd/atlasuserdisk/rucio/data15_13TeV/07/53/DAOD_MUON0.06783725._000004.
pool.root.1"]
filelist += ["root://at1-xrdr//atlas/xrootd/atlasuserdisk/rucio/data15_13TeV/4f/ca/DAOD_MUON0.06783725._000001.
pool.root.1"]
...
athenaCommonFlags.PoolESDInput=filelist
```

The Tier 2 space supports FAX global logical file name (gLFN), so the above two files can be simplified as:

```
filelist = []
filelist += ["root://at1-xrdr//atlas/rucio/data15_13TeV:DAOD_MUON0.06783725._000004.pool.root.1"]
filelist += ["root://at1-xrdr//atlas/rucio/data15_13TeV:DAOD_MUON0.06783725._000001.pool.root.1"]
...
athenaCommonFlags.PoolESDInput=filelist
```

Experimental: If the files are not at SLAC but are available in an ATLAS site that provides xrootd service, you can also access them by replacing "at1-xrdr" by "atlfax" (see below).

```
filelist = []
filelist += ["root://atlfax//atlas/rucio/data15_13TeV:DAOD_MUON0.06783725._000004.pool.root.1"]
filelist += ["root://atlfax//atlas/rucio/data15_13TeV:DAOD_MUON0.06783725._000001.pool.root.1"]
...
athenaCommonFlags.PoolESDInput=filelist
```

To read an xrootd file in ROOT:

```
root[0] TXNetFile f("root://at1prf01:11094//atlas/local/myFile.root");
```

To copy files in and out of xrootd:

```
xrdcp myFile.root root://at1prf01:11094//atlas/local/myFile.root
xrdcp root://at1prf01:11094//atlas/local/myFile.root myFile.root
```