

SLAC Analysis Computing Facility

Introduction

SLAC hosts a shared analysis computing facility for the US ATLAS members. The center provides CPU (interactive, including Jupyter, and batch), disk space and software tools to support both Grid and non-Grid based physics analysis activities.

Getting started to obtain a SLAC computer account

This information is for users who want direct access to SLAC computers, not accessing SLAC computing resources via the GRID. The steps listed here may take days to complete here so plan early.

Register as a SLAC user and obtain a SLAC computer account

Please refer to instructions at: <https://atlas.slac.stanford.edu/atlas-support-center>, under "USER INFORMATION".

Subscribe to e-group

Please go to [CERN e-group blocked URL](#) and (search and) subscribe to `atlas-us-slac-acf`. We will use this e-group for announcement and for user discussion specific to the SLAC-ACF. If you do not have a CERN account to subscribe to this e-group by yourself, please e-mail yangw@slac.stanford.edu to have your email address added to the group.

Login to SLAC

SLAC provides a pool of login nodes with CVMFS and Grid tools. You can access them by ssh to `centos7.slac.stanford.edu`. Assuming your unix shell is `/bin/bash`, you may use the following as a template of your `$HOME/.bashrc` file

```
# setup ATLAS environment
export ALRB_localConfigDir=/gpfs/slac/atlas/fs1/sw/localconfig
export RUCIO_ACCOUNT="change_me"
export ATLAS_LOCAL_ROOT_BASE=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase
alias setupATLAS='source $ATLAS_LOCAL_ROOT_BASE/user/atlasLocalSetup.sh'
...
```

upon login, you can run the alias "setupATLAS" to set up CVMFS env, which will print out the available commands such as `asetup`, `rucio`, and their usage.

Please refer to [Analysis Software Tutorial](#) for more info.

Jupyter Notebooks

Please refer to this page for SLAC's new Jupyter resources: <https://usatlas.github.io/tier3docs/>. The following info are obsolete.

OLD: Jupyter notebook environment is available at <http://jupyter.slac.stanford.edu>. Login with your SLAC Unix account, and you can choose any JupyterLab environments (container images) available from that page (some of them will allow access to GPU resources). Only the ATLAS images will mount your GPFS home directory and data directory. The ATLAS images provide PyROOT and JupyROOT (ROOT C++) notebooks with capability to access remote data via the `xroot` (`root`) protocol or `webdav` (`http`) protocol. It includes `uproot` and several Machine Learning software packages. It also allows storing notebooks to Google Drive for portability.

- This is an area that SLAC and the Analysis Computing Facility is constantly looking to improve. Please give us your feedback by sending e-mail to SLAC ACF mailing list atlas-us-slac-acf@cern.ch
- If your home directory is on AFS (which is not visible from the JupyterLab), you can still request a GPFS data directory - send your request to unix-admin@slac.stanford.edu, indicating that "I want a GPFS data directory under `/gpfs/slac/atlas/fs1/d/<my_user_name>`".
- CVMFS is available in the JupyterLab
- With JupyROOT (ROOT C++) and remote data access (via `xroot` or `webdav/http` protocol), it is possible to run [ProofLite](#) even inside a JupyROOT notebook

Remote X window access

Please refer to [SLAC's FastX page blocked URL](#) for detail instructions.

Disk space

SLAC provides to **new** ATLAS users two personal storage spaces

- Home directory `/gpfs/slac/atlas/fs1/u/<user_name>`. The quota is 100GB. This space is backed up to tape. The most recently backups are also available online in the `/gpfs/slac/atlas/fs1/u/.snapshots` directory.

- Data directory /gpfs/slac/atlas/fs1/d/<user_name>. The initial quota is 2TB. To request quota increase to up to 10TB, please send e-mail to unix-admin@slac.stanford.edu. Data in this area is not backed up.

Some users already have computer accounts at SLAC. Their home directories may be on AFS and they don't have the data directories on GPFS. Those users can request moving their home directories to GPFS and/or creating data directories for them by sending e-mail to unix-admin@slac.stanford.edu.

SLAC also provides Xrootd based storage that are managed by RUCIO. The data are read-only to users. However, it is possible to use [R2D2blocked URL](#) to request ATLAS datasets to be transferred to those space.

- From the interactive login machines, one can browse the Xrootd storage by `cd /xrootd/atlas` on RHEL6 interactive login nodes
- From batch jobs, one can't list directory in the Xrootd storage. But can access files via `root://atlrdr1/xrootd/atlas/...`
- Using the rucio tools, one can list ATLAS datasets that are already in the Xrootd storage system via command `rucio list-datasets-rse <RSE>`, where RSE can be SLACXRD_DATADISK, SLACXRD_LOCALGROUPDISK, SLACXRD_SCRATCHDISK.

Submit batch jobs

SLAC uses LSF batch system. LSF replica your current environment setup when submitting jobs. This includes your current working directory and any Unix environment variable setups. The following are examples of using LSF:

Submit a job

```
$ cat myjob.sh
#!/bin/sh
#BSUB -W180
#BSUB -R 'centos7'
pwd
echo "hello world"
```

```
$ bsub < myjob.sh
Job <96917> is submitted to default queue <medium>.
```

This will submit a job to LSF. The "pwd" command will print out the job's working directory, which should be the same the directory where this job is submitted.

Manage jobs

Use `bjobs`, or `bjobs -I <JOBID>` to get detailed info about a specific job. Use `bkill <JOBID>` to kill a job

- For details about those LSF commands and their options, please refers to the man pages of "bsub", "bjobs", "bkill".

More info on LSF

- [ATLAS specific info on LSF](#)
- For information regarding high performance computing clusters, including LSF related documents and best practices, please refer to [the SLAC High Performance Computing pageblocked URL](#).

Resource monitoring

Coming soon