

JLab Auger

Auger is the interface to JLab's data analysis cluster ("the farm")

- Controls batch job submissions
 - Manages input/output from jobs
 - Provides details on job status
 - Gathers job statistics
- Code written in-house at JLab
 - Connects to open source PBS/Torque/Maui resource manager and scheduler
 - Underlies JLab's SWIF workflow tool (Chris' talk)

JLab Auger (2)

Auger commands

- `jsub` – submits job described by text or xml file
- `jobstat` – checks job statistics
- `jkill` – ends running jobs
- `jobinfo` – displays running job information
- `farmhosts` – displays farm node information

Auger commands reside in `/site/bin`, available from all JLab CUE (Common User Environment) systems.

Interactive user systems are **ifarm.jlab.org**

JLab Auger (3)

Auger requires specifying a valid project, track, and command

- Project - select appropriate from <https://scicomp.jlab.org/scicomp/#/projectName>
- Track – describes job type, submits to appropriate queue

Track	Batch Queue	Description
debug	priority	For debugging use
reconstruction	prod64	Reconstruction of raw data
analysis	prod64	Analysis jobs
one_pass	prod64	Combined reconstruction/ analysis
simulation	prod64	Simulation jobs
test	idle	Test run of code
theory	longJob	Running long theory jobs



JLab Auger (4)

Auger options include specifying other than defaults for

OS
COMMAND_COPY
JOBNAME
MAIL
TIME
OPTIONS
INPUT_FILES
SINGLE_JOB
MULTI_JOBS
OTHER_FILES
INPUT_DATA
OUTPUT_DATA
OUTPUT_TEMPLATE
CPU
DISK_SPACE
MEMORY

Details for these options are at https://scicomp.jlab.org/docs/text_command_file

JLab Auger (5)

Examples

Text configuration file

```
ifarm/ > jsub <config file> f
```

```
PROJECT: MyProject  
TRACK: MyTrack  
JOBNAME: MyJob  
COMMAND: ls
```

XML configuration file

```
ifarm/ > jsub -xml <XML config file> f
```

```
<Request>  
<Email email="user@jlab.org" request="false" job="true"/>  
<Project name="MyProject"/>  
<Track name="MyTrack"/>  
<Name name="MyJob"/>  
<Command><![CDATA[  
  ls  
]]></Command>  
  
<Job>  
</Job>  
  
</Request>
```

JLab Auger (6)

Standard out and standard errors

- stored in your `~/.farm_out` directory
- named `<job-id>.out` , `<job-id>.err`

Exit codes

- User defined exit codes are passed by Auger as the final job exit code; if the job is killed by PBS for out of resource or other, the code is $128 + \text{the signal number}$
- Exit code of 0 means the job completed successfully

JLab Auger (7)

Fairshare and Priorities

Auger Account	FS percentage
halla	10%
hallb	20%
hallc	10%
halld	60%
other	10%

Physics Computing Coordinators can request adjustments based on current priorities

JLab Auger (8)

Useful tips...

- When you need resources other than defaults, request enough but not too much, especially
 - MEMORY – wastes resources if overspecified
 - WALLTIME – if the job is short, tell the scheduler; likely it will be able to fit small jobs in sooner
- Only one user runs jobs on a node at a time – use the whole machine because it's all yours
- Run multi-threaded jobs – more efficient than the same number of single-threaded jobs doing the same work
- I/O usage – Lustre likes large files (Chip's talk)
- Use SWIF, JLab's workflow tool (Chris' talk)

Jlab Auger (9)

Viewing information about jobs

<http://scicomp.jlab.org>

From the command line,

```
ifarm> jobstat -u ton
```

JOB_ID	USER	STAT	QUEUE	EXEC_HOST	JOB_NAME	SUBMIT_TIME	CPU_TIME
38065376	ton	Q	prod64	--	...eplay_1_1898	May 17 12:18 --	halla
38065379	ton	Q	prod64	--	...eplay_2_1899	May 17 12:18 --	halla
38065381	ton	Q	prod64	--	...eplay_3_1885	May 17 12:18 --	halla
38065383	ton	Q	prod64	--	...eplay_4_1889	May 17 12:18 --	halla

```
ifarm> ls ~/.farmout
```

```
-rw-r--r-- 1 philpott scicomp 117 Feb 2 12:00 FarmJob.33168932.out
-rw-r--r-- 1 philpott scicomp 438 Feb 2 12:00 FarmJob.33168932.err
-rw-r--r-- 1 philpott scicomp 117 Feb 2 12:05 FarmJob.33169006.out
-rw-r--r-- 1 philpott scicomp 22 Feb 2 12:05 FarmJob.33169006.err
```

JASMine

JASMine

- Jefferson Lab's Asynchronous Storage Manager
- Interface between disk and tape storage
- In house written software
- Used with Auger and SWIF for efficient tape I/O
- Manages tape library and data
 - currently ~12 PB
 - IBM TS3500 tape robot with 14 LTO drives
 - stores all “raw” CEBAF beam data, with a duplicate copy in the vault, plus all other “production” data from processing on the farm

JASMine (2)

JASMine commands

- jput – put files to tape
- jget – get files from tape
- jqueue – info on read and write tape jobs
- jcancel – cancel jasmine jobs

File system /mss – stub files of the tape data

```
ifarm> ls -alL /mss
```

```
drwxrwxr-x 9 larrieu epics 4096 Jul 27 2016 accel
dr-xr-xr-x 6 jasmine scicomp 4096 Jan 10 1997 ccc
dr-xr-xr-x 62 jasmine nobody 4096 Feb 18 02:53 clas
drwxrwxr-x 9 clas12 clas12 4096 Jan 26 10:37 clas12
drwxrwxr-x 4 pavel muscn 4096 Mar 22 2016 ehsq
dr-xr-xr-x 3 jasmine nobody 4096 Nov 11 2014 eic
dr-xr-xr-x 60 jasmine nobody 4096 Feb 15 11:03 halla
dr-xr-xr-x 10 jasmine nobody 4096 Feb 7 14:11 hallb
dr-xr-xr-x 40 jasmine nobody 4096 Mar 6 08:24 hallc
dr-xr-xr-x 20 jasmine nobody 4096 Jan 3 15:39 halld
dr-xr-xr-x 369 jasmine nobody 12288 Mar 30 16:36 home
drwxrwxr-x 36 jasmine lattice-1 4096 Apr 20 03:37 lattice
```