

# LSF at SLAC

Using the SLAC/LCLS Offline Batch Cluster

Neal Adams

SLAC National Accelerator Laboratory

[neal@slac.stanford.edu](mailto:neal@slac.stanford.edu)

# What is LSF?

- Load Sharing Facility (LSF) product by Platform Computing Corporation.
- Allows queuing and scheduling of batch jobs.
- Provides scheduling of jobs based on load conditions and resource requirements specified by the user.

## What is a batch job?

- "A unit of work run in the LSF system."
- A batch job can be a script, command or program.

Example: `bsub hostname`

# LCLS Offline LSF Servers

- LSF commands for querying and job submission can only be performed from licensed LSF hosts.
- LCLS Offline interactive servers licensed for LSF.  
**psdev**  
**pslogin**
- LCLS Offline LSF batch servers.  
**psana1201-1210**  
**psana1301-1320**  
**(psana1101-1120 are still used for interactive jobs and will be added later to LSF)**

# The LCLS batch queue.

```
QUEUE: lclsq
  -- LCLS Offline Processing Farm

PARAMETERS/STATISTICS
PRIO NICE STATUS          MAX JL/U JL/P JL/H NJOBS  PEND  RUN SSUSP USUSP  RSV
115   0 Open:Active      -   -   -   16   30   0   30   0   0   0

```

SCHEDULING PARAMETERS

	r15s	r1m	r15m	ut	pg	io	ls	it	tmp	swp	mem
loadSched	-	-	-	-	-	-	-	-	-	-	-
loadStop	-	-	-	-	-	-	-	-	-	-	-

  

	lammpi_load	scratch
loadSched	-	-
loadStop	-	-

```
USERS: lclsgrp/
HOSTS: psanafarm/
```

## Useful LSF Commands

<b>bsub</b>	submit a batch job to LSF
<b>bjobs</b>	display batch job information
<b>bkill</b>	kill batch job
<b>bmod</b>	modify job submission options
<b>bqueues</b>	display batch queue information
<b>users</b>	displays information about batch users
<b>lshosts</b>	display LSF host information
<b>lslload</b>	display LSF host load information

For more details use: [\*man <command\\_name>\*](#).

# Useful LSF Commands

- **bqueues**

```
[neal@pslogin ~]$ bqueues
QUEUE_NAME      PRIO STATUS          MAX JL/U JL/P JL/H NJOBS  PEND  RUN  SUSP
...
lcls q          115 Open:Active      -  -  -  -  384   0   384   0
...
short          185 Open:Active      -  -  -  -   0    0    0    0
medium         180 Open:Active      -  -  -  -  153  102   51    0
long           175 Open:Active      -  -  -  -  897  757  140    0
xlong          170 Open:Active      -  -  1  2 1636 1359  277    0
genmpiq        168 Open:Active      -  -  -  -   0    0    0    0
xxl            165 Open:Active     160 64  -  1   56   1    55    0
...
```

- **busers**

```
[neal@pslogin ~]$ busers
USER/GROUP      JL/P  MAX  NJOBS  PEND  RUN  SSUSP  USUSP  RSV
neal             -    -    0      0     0    0      0      0

[neal@pslogin ~]$ busers perazzo
USER/GROUP      JL/P  MAX  NJOBS  PEND  RUN  SSUSP  USUSP  RSV
perazzo         -    -   384    0    384  0      0      0
```

# Useful LSF Commands

- lshosts

```
[neal@pslogin ~]$ lshosts -R psana
```

```
HOST_NAME      type    model  cpuf  ncpus  maxmem  maxswp  server  RESOURCES
psana1201      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1202      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1203      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1204      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1205      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1206      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1207      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1208      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1209      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
psana1210      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
[...]
```

```
[neal@pslogin ~]$ lshosts psana1320
```

```
HOST_NAME      type    model  cpuf  ncpus  maxmem  maxswp  server  RESOURCES
psana1320      LINUX  INTEL_29  14.6   16  24098M  4095M   Yes (bs linux linux64 rhel50 psana)
```



## Using bsub

- To submit batch jobs to the SLAC/LCLS LSF cluster use the *bsub* command.

*bsub* [*bsub options*] *command* [*arguments*]

For example:

*bsub -o outputfilename date -u*

# Using bsub

Example of a simple bsub:

```
[neal@pslogin ~]$ bsub -q lclsq hostname  
Job <945166> is submitted to queue <lclsq>.
```

```
[neal@pslogin ~]$ bjobs  
JOBID  USER  STAT  QUEUE      FROM_HOST  EXEC_HOST  JOB_NAME  SUBMIT_TIME  
945166  neal   PEND  lclsq      psdev      hostname   hostname  Mar  4 19:17
```

```
[neal@pslogin ~]$ bjobs  
JOBID  USER  STAT  QUEUE      FROM_HOST  EXEC_HOST  JOB_NAME  SUBMIT_TIME  
945166  neal   RUN   lclsq      psdev      psana1202  hostname  Jun 28 15:13
```

```
[neal@pslogin ~]$ bjobs 945166  
JOBID  USER  STAT  QUEUE      FROM_HOST  EXEC_HOST  JOB_NAME  SUBMIT_TIME  
945166  neal   DONE  lclsq      psdev      psana1202  hostname  Jun 28 15:13
```

# Using bsub

## Output from my simple batch job:

```
Job <hostname> was submitted from host <pslogin> by user <neal>.
Job was executed on host(s) <psana1202>, in queue <lclsq>, as user <neal>.
</reg/neh/home/neal> was used as the home directory.
</reg/neh/home/neal> was used as the working directory.
Started at Mon Jun 28 15:13:27 2010
Results reported at Mon Jun 28 19:13:32 2010
Your job looked like:
```

```
-----
# LSBATCH: User input
hostname
-----
```

```
Successfully completed.
Resource usage summary:
CPU time   :      0.06 sec.
Max Memory :        2 MB
Max Swap   :       16 MB
Max Processes :      1
Max Threads :        1
```

```
The output (if any) follows:
psana1202
```

## A few useful bsub options.

- Submit with a CPU limit (normalized): `bsub -c`  
example: `bsub -q lclsq -c 24:00 date`
- Submit with a RUN limit (wallclock): `bsub -W`  
example: `bsub -q lclsq -W 24:00 date`
- Submit with a jobname: `bsub -J "job_name"`  
example: `bsub -q lclsq -J "Date_job" date`
- Submit a job array: `bsub -J "job_name[array-elements]"`  
example: `bsub -q lclsq -J "amedeo[1-100]" my_script`

# The LCLS Offline Cluster

- LCLS servers psana1201-1220, 1301-1320 ([psanafarm](#))
  - 40 Supermicro Blade servers each with Intel(R) Xeon(R) CPU @ 2.27GHz
  - 640 cores (job slots)
- Dedicated LSF queue ([lclsq](#))
  - Access controlled via LSF user group ([lclsgrp](#))

# Good Practice

- Specify output files for batch job output. (bsub with -o or -oo options).  
Make sure the file path exists and that you have the appropriate permissions.
- Before submitting 100s of jobs to LSF, please try submitting a smaller number to ensure that you get the expected results.
- Everything required by the batch job (incl. binary) needs to be visible from the batch nodes.
- Use local disk space on the LSF servers for job files and output files for better performance and copy files to your user or group space at job completion.
- LSF can handle tens of thousands of jobs. However we would prefer that not all of them are yours.

# Batch Job Exit Codes

- Job exit codes 1-128 are from whatever the user is running while those exceeding 128 are the signal values modulo 128.

Example:

A job exit code of 137 would indicate that the job was sent SIGKILL ( $137-128=9$ ) or kill signal 9.

A job exit code of 152 would indicate that the job was sent SIGXCPU ( $152-128=24$ ) or kill signal 24.

- To determine the signal name and number use *man*.  
Linux: `man 7 signal`

# Is LSF having problems?

- SLAC's LSF cluster can be very busy at times causing the LSF master to respond slowly to your command requests (bsub, bjobs, etc). You may see the following messages in response to your LSF batch commands when this occurs. These can also occur briefly when we have initiated an LSF reconfiguration for administrative purposes.

```
batch system daemon not responding ... still trying
batch system daemon not responding ... still trying
batch system daemon not responding ... still trying
```

This does not effect jobs already running or pending in the LSF cluster.  
It only affects LSF's ability to talk to you. The commands will eventually complete.

- If you see these messages [Monday through Thursday between 19:35 and 19:55 \(7:35-7:55PM\)](#) we automatically run an LSF reconfiguration during those times.
- Scheduled outages of the LSF cluster are normally announced via the SLAC Computing Outages web page <https://www-internal.slac.stanford.edu/comp-out>.



# LSF Documentation

- SLAC specific LSF documentation.

<http://www.slac.stanford.edu/comp/unix>

Click on "High Performance"

- Platform LSF documentation.

<http://www.slac.stanford.edu/comp/unix/package/lsf/currdoc/html/index.html>

<http://www.slac.stanford.edu/comp/unix/package/lsf/currdoc/pdf/manuals/>

# Problem Reporting

Send email to:

[pcds-help@slac.stanford.edu](mailto:pcds-help@slac.stanford.edu)