

Martinos Center Compute Clusters

Why n How
Feb 6 2014

Jon Kaiser

Martinos Center Compute Clusters

Intro

- What are the compute clusters
- How to gain access
- Housekeeping

Usage

- Log In
- Submitting Jobs
- Request CPUs/vmem
- Queues
 - GPU
- Email Status
- I/O
- Interactive
- Dependencies
 - Daisy Chain
 - Wrapper Script
 - In Progress

Job Status

Running Jobs

- Show Job Status
- See Standard Output

Completed Jobs

Failed Jobs

Delete Jobs

Idle Jobs

Misc Commands

- nodecount
- nodeusage
- usercount
- qselect

Summary

Martinos Center Compute Clusters

Intro - What are the compute clusters?

launchpad

At Needham Data Center

127 nodes

~115 “normal” nodes

Two 64 bit Intel Xeon quad cores

56 GB RAM

~12 GPU nodes

Available exclusively for GPU jobs

tensor

In CNY “cold room”

107 nodes

89 “normal” nodes

Dual cores with 4GB RAM

18 “big” nodes

Dual cores with 8GB RAM each

I will only talk about launchpad today, but the commands for each are the exact same. Use tensor if launchpad is full, if launchpad resources are overkill, or if the data for your jobs live locally and are slowed by the I/O problem.

Martinos Center Compute Clusters

Intro - How to gain access

Email me: [kaiser \[at\] nmr.mgh.harvard.edu](mailto:kaiser@nmr.mgh.harvard.edu)

Let us know who/what/why/how you need access.

Martinos Center Compute Clusters

Intro - Housekeeping

Questions?:

- Any user-specific questions can be sent to me
 - extend Walltime
 - permission to use MATLAB toolboxes or a high priority queue
- General questions should be sent to the group for advice and help:
batch-users [at] nmr.mgh.harvard.edu

Limits:

- We recommend each user use up to 150 job slots during normal usage.
- Evenings/weekends you may use up to 200 slots of CPU/vmem
- While there is a queue, we request you only use ~75 CPU/vmem slots

Do not run anything directly on launchpad. Submit your jobs. Any programs found running on the master node will be killed, no exceptions.

Martinos Center Compute Clusters

Intro - Housekeeping

MATLAB:

There is a limited amount of matlab licenses for the entire center (~120). For this reason, we recommend any matlab code submitted for execution should be “compiled” ahead of time. Please see the URL to the article on how to do it. When the program is compiled, it doesn't use a matlab license and is no longer under a matlab restriction.

Courtesy of coutu:

<http://nmr.mgh.harvard.edu/martinos/itgroup/deploytool.html>

If you receive a matlab license error, the limited amount of licenses are all occupied.

To see the distribution of matlab licenses, run:

```
lmstat -f MATLAB
```

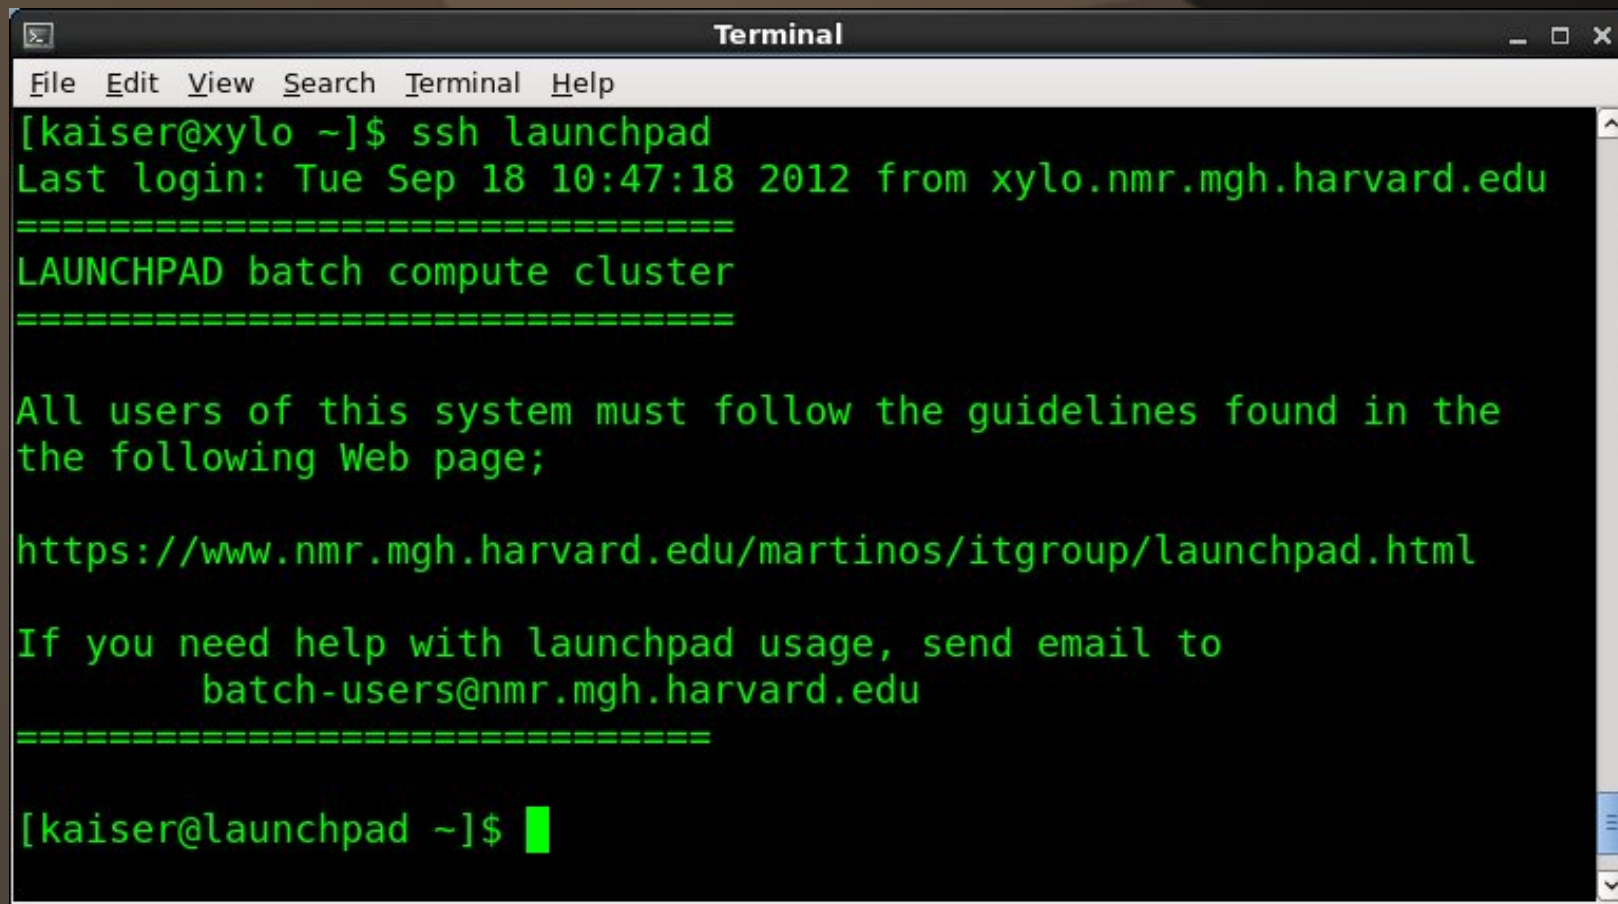
To see the distribution of toolboxes as well, run:

```
lmstat -a
```

Please note, licenses for individual users at their own workstations are given priority, ahead of launchpad users. If users complain we will have to kill a job in order to recover a license.

Martinos Center Compute Clusters

Usage - Log In



```
Terminal
File Edit View Search Terminal Help
[kaiser@xylo ~]$ ssh launchpad
Last login: Tue Sep 18 10:47:18 2012 from xylo.nmr.mgh.harvard.edu
=====
LAUNCHPAD batch compute cluster
=====

All users of this system must follow the guidelines found in the
the following Web page;

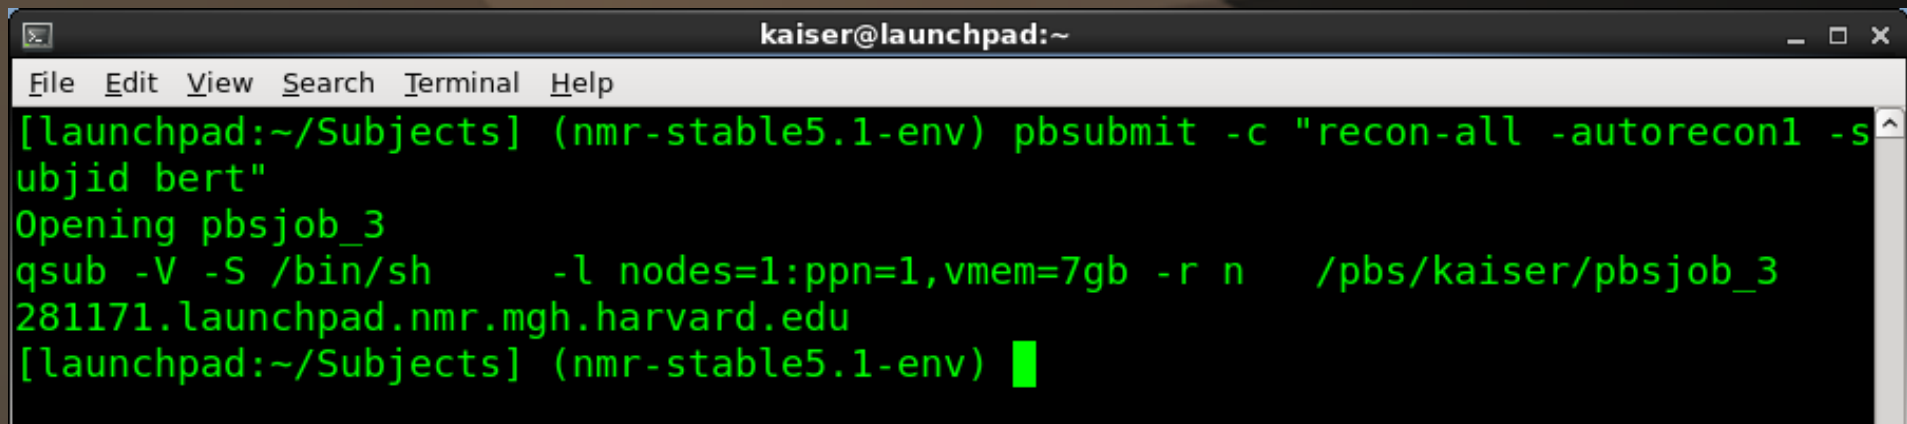
https://www.nmr.mgh.harvard.edu/martinos/itgroup/launchpad.html

If you need help with launchpad usage, send email to
    batch-users@nmr.mgh.harvard.edu
=====

[kaiser@launchpad ~]$ █
```

Martinos Center Compute Clusters

Usage – Submitting Jobs



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -c "recon-all -autorecon1 -s  
ubjid bert"  
Opening pbsjob_3  
qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_3  
281171.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

pbsubmit is a wrapper script that:

- formats the command that is executed (/pbs/kaiser/pbsjob_3)
- automatically selects the default settings (unless overridden)
 - number of nodes (nodes=1)
 - number of CPUs (ppn=1)
 - amount of virtual memory (vmem=7gb)
- submits the job using the qsub command

pbsjob_3

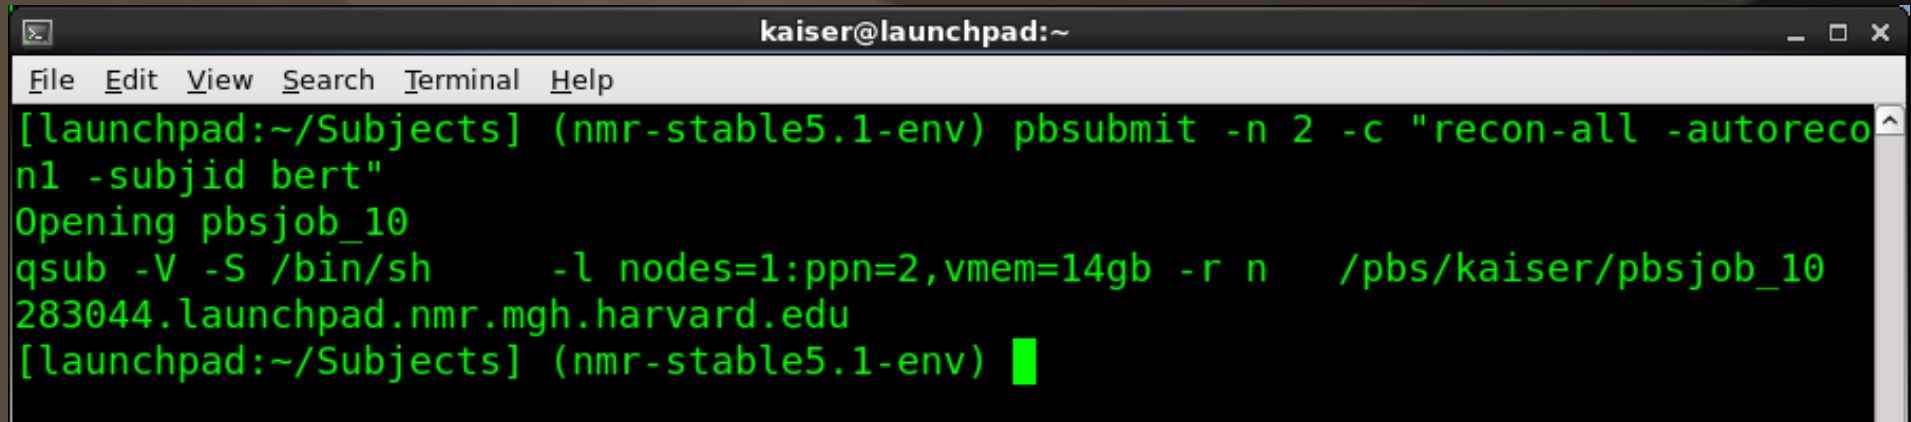
281171.launchpad.nmr.mgh.harvard.edu

is the Job Number

is the Job ID

Martinos Center Compute Clusters

Usage - Request CPUs/vmem



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -n 2 -c "recon-all -autoreco  
n1 -subjid bert"  
Opening pbsjob_10  
qsub -V -S /bin/sh -l nodes=1:ppn=2,vmem=14gb -r n /pbs/kaiser/pbsjob_10  
283044.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Only request more CPUs or Virtual Memory if you need them.

CPUs

- You should only request extra CPUs if the program you are running is multi-threaded.
- If you aren't sure if the program is multi-threaded, it probably isn't.

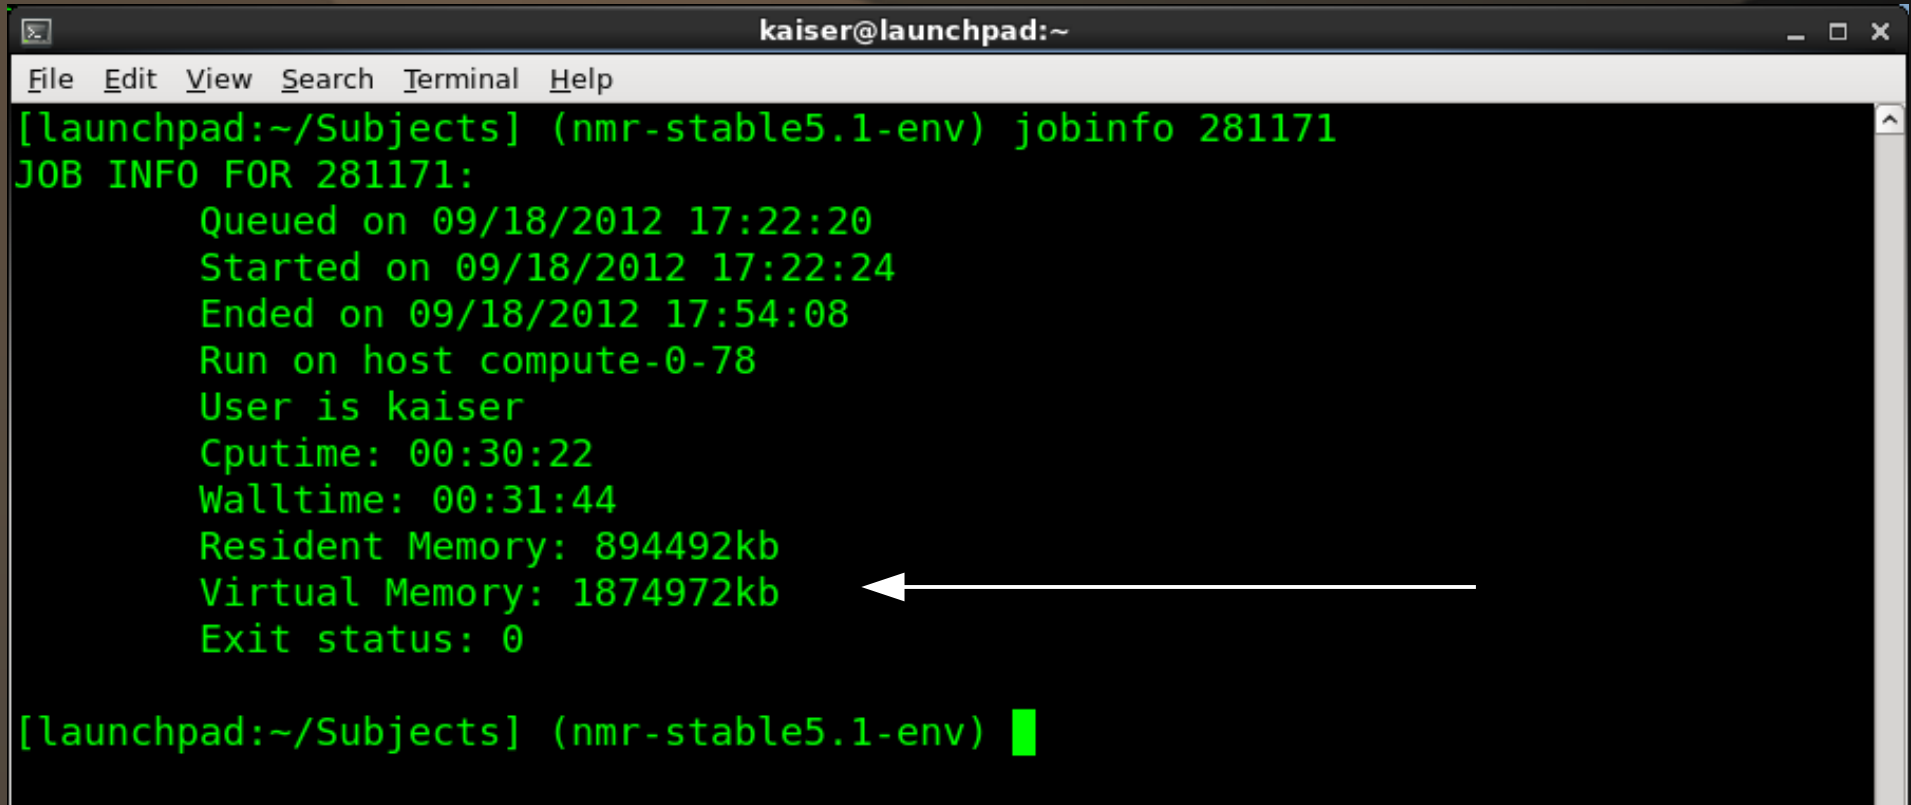
Virtual Memory

- Only request as much as you need.
- If you aren't sure how much you'll need, run a single test case. Start with the default of 7GB of vmem. If it fails due to a lack of memory, request 14GB. Then 21GB etc...

So, how much virtual memory did the job use?

Martinos Center Compute Clusters

Usage – Request CPUs/vmem



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) jobinfo 281171  
JOB INFO FOR 281171:  
    Queued on 09/18/2012 17:22:20  
    Started on 09/18/2012 17:22:24  
    Ended on 09/18/2012 17:54:08  
    Run on host compute-0-78  
    User is kaiser  
    Cputime: 00:30:22  
    Walltime: 00:31:44  
    Resident Memory: 894492kb  
    Virtual Memory: 1874972kb  
    Exit status: 0  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Only used 1.9GB of virtual memory. Safely under the default request of 7GB. No need to ask for more.

Limits – Reminder that we prefer each user to only use ~150 job slots during the day.
A job that requests 1 CPU and 14GB of vmem counts as two slots worth of resources.
Submit the jobs to the max75 queue ('-q max75') to self-regulate.

Martinos Center Compute Clusters

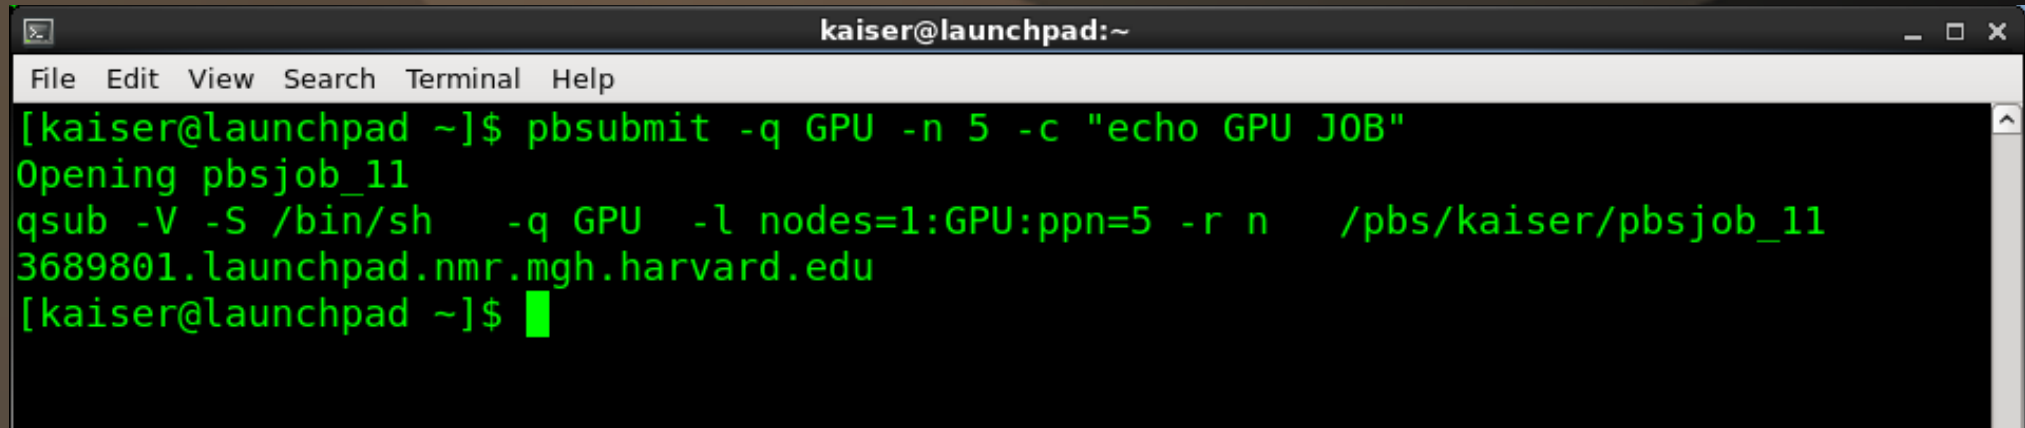
Usage - Queues

Queue	Priority	Max CPU/User	Description
default	101000	150	Walltime of 96 hours
p20	102000	Unlimited	
p30	103000	Unlimited	
GPU	90	Unlimited	GPU nodes
extended	80000	50	Walltime of 196 hours
matlab	101000	20	Limit of ~120 matlab licenses for the Center
max10	101000	10	
max20	101000	20	
max50	101000	50	
max75	101000	75	
max100	101000	100	
max200	80000	200	

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -q max100 -c "recon-all -aut  
orecon1 -subjid bert"  
Opening pbsjob_4  
qsub -V -S /bin/sh -q max100 -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbs  
job_4  
281184.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Martinos Center Compute Clusters

Usage – Queues - GPU



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ pbsubmit -q GPU -n 5 -c "echo GPU JOB"  
Opening pbsjob_11  
qsub -V -S /bin/sh -q GPU -l nodes=1:GPU:ppn=5 -r n /pbs/kaiser/pbsjob_11  
3689801.launchpad.nmr.mgh.harvard.edu  
[kaiser@launchpad ~]$
```

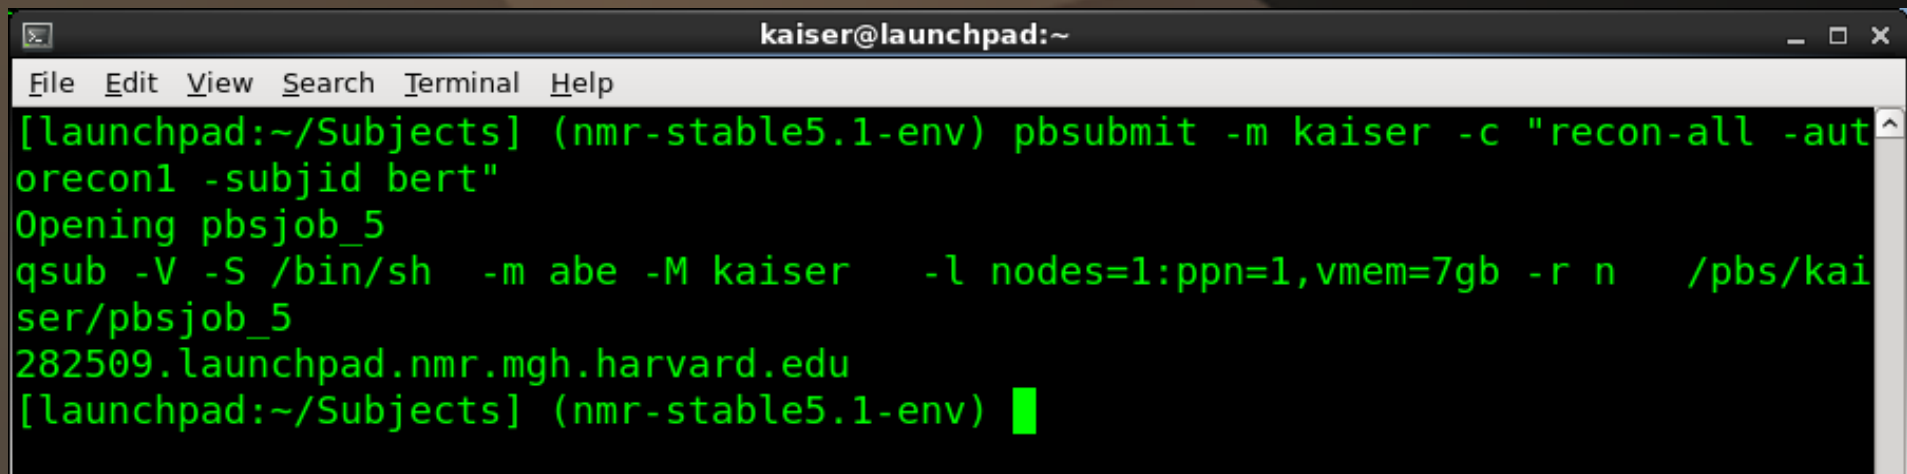
For a GPU job, just submit the job to the GPU queue and it will automatically be sent to the GPU nodes.

Since each node only has one GPU (but 8 CPUs) we recommend that users request 5 CPUs for their job. If you choose the default of 1 CPU, then multiple GPU jobs could be sent to the node which will then fail because the GPU is already in use.

The reason for requesting 5 CPUs, and not 8, is if the cluster is overloaded and there are several jobs waiting idle, we want the option to run some regular jobs on the GPU nodes.

Martinos Center Compute Clusters

Usage - Email Status



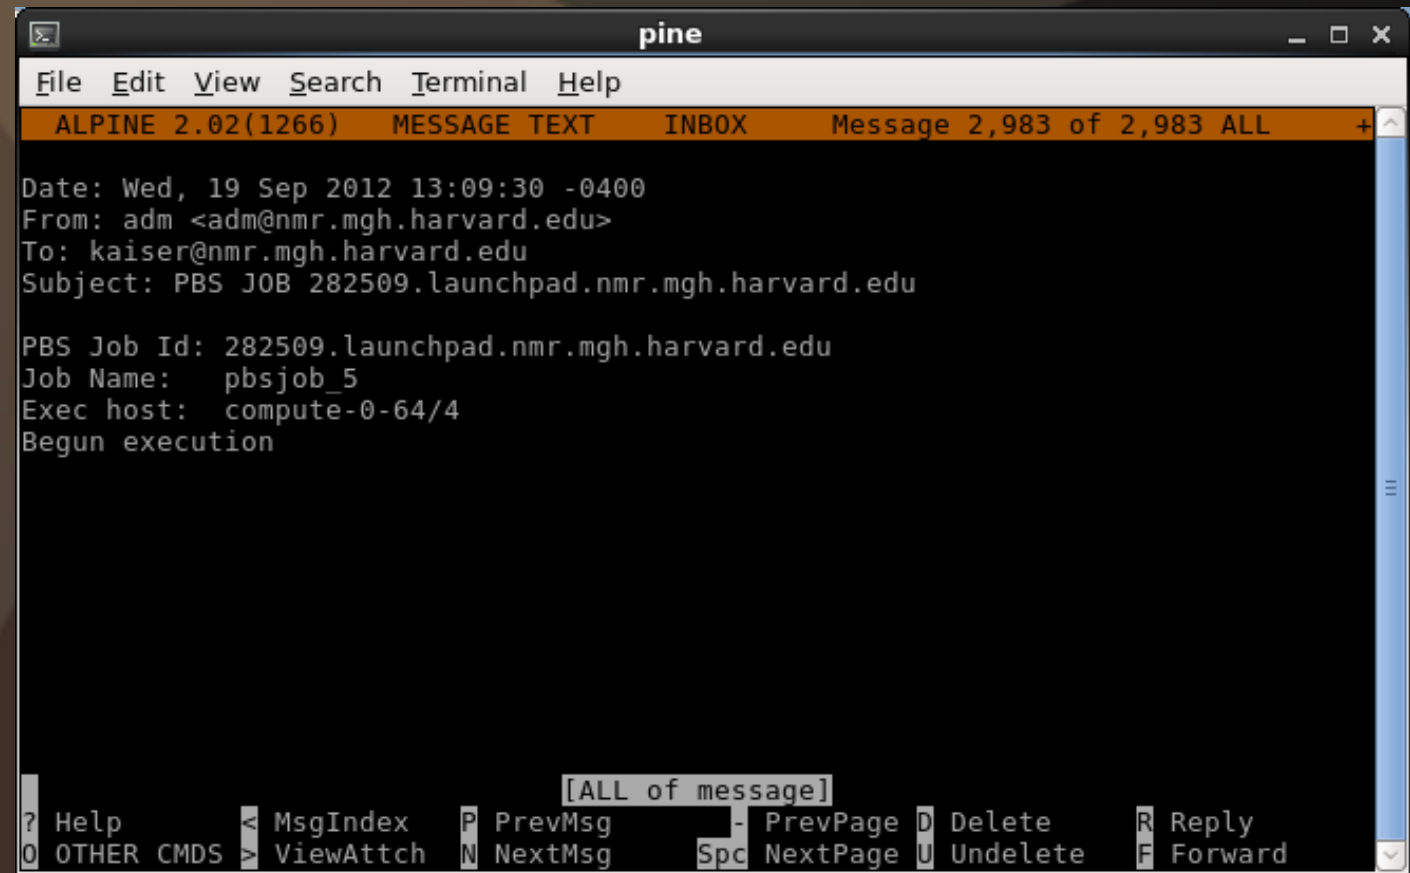
```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -m kaiser -c "recon-all -aut  
orecon1 -subjid bert"  
Opening pbsjob_5  
qsub -V -S /bin/sh -m abe -M kaiser -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kai  
ser/pbsjob_5  
282509.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

- Sends email to user (**replace 'kaiser' with your username**) on job start and finish
- To receive email only if job completes with an error, append '-e' to command line
 - To receive email only upon job completion (error or no error), append '-f' to command line

Martinos Center Compute Clusters

Usage - Email Status

Start Execution:



```
pine
File Edit View Search Terminal Help
ALPINE 2.02(1266) MESSAGE TEXT INBOX Message 2,983 of 2,983 ALL
Date: Wed, 19 Sep 2012 13:09:30 -0400
From: adm <adm@nmr.mgh.harvard.edu>
To: kaiser@nmr.mgh.harvard.edu
Subject: PBS JOB 282509.launchpad.nmr.mgh.harvard.edu

PBS Job Id: 282509.launchpad.nmr.mgh.harvard.edu
Job Name: pbsjob_5
Exec host: compute-0-64/4
Begun execution

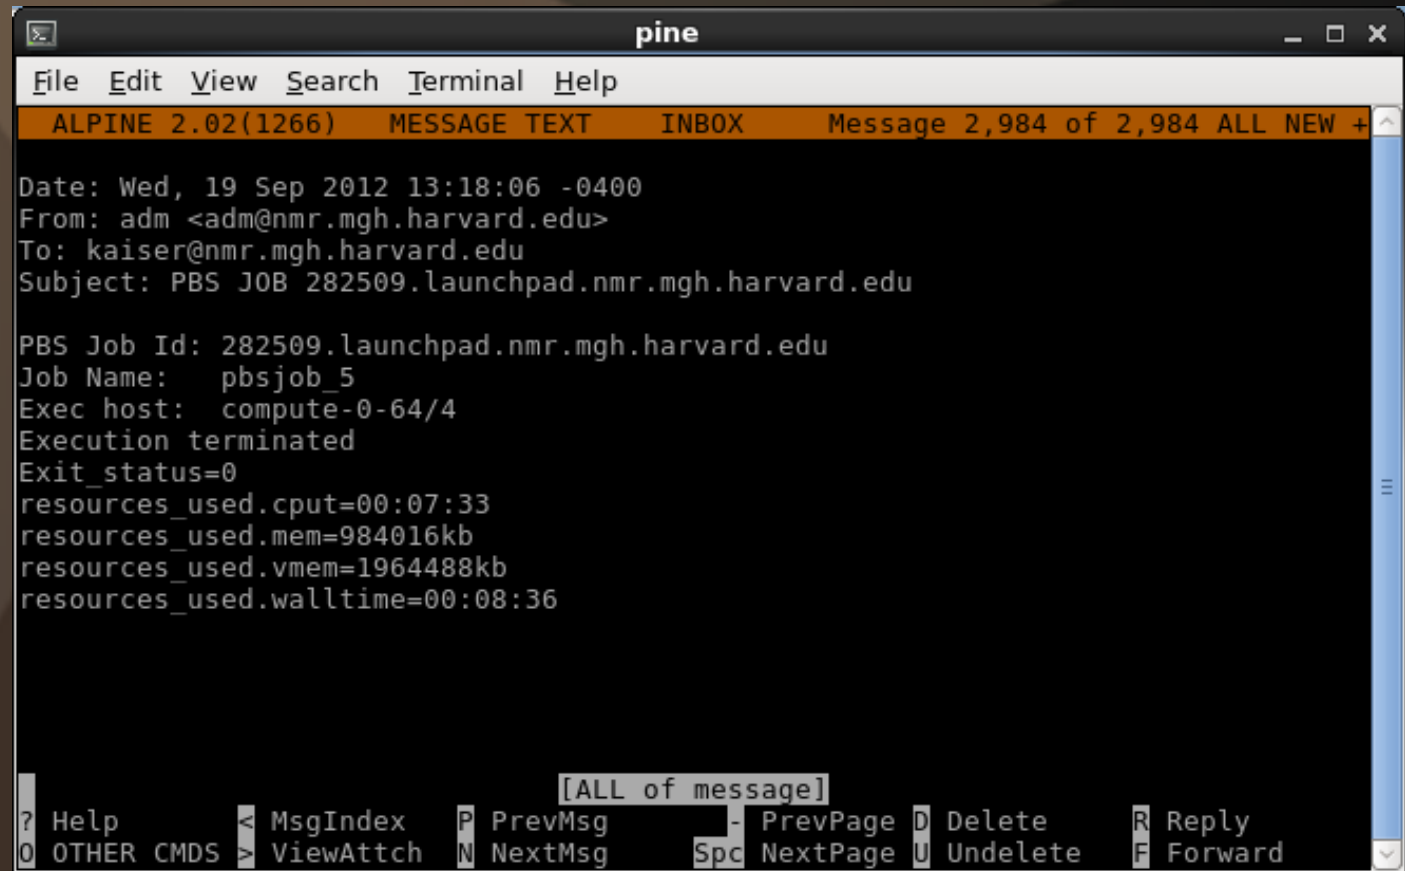
[ALL of message]
? Help      ^ M msgIndex  P PrevMsg    - PrevPage  D Delete     R Reply
0 OTHER CMDS v ViewAtch   N NextMsg    Spc NextPage  U Undelete  F Forward
```

Identifies the JobID, Job Number, and the node it is running on

Martinos Center Compute Clusters

Usage - Email Status

Finish Execution:



```
pine
File Edit View Search Terminal Help
ALPINE 2.02(1266) MESSAGE TEXT INBOX Message 2,984 of 2,984 ALL NEW +
Date: Wed, 19 Sep 2012 13:18:06 -0400
From: adm <adm@nmr.mgh.harvard.edu>
To: kaiser@nmr.mgh.harvard.edu
Subject: PBS JOB 282509.launchpad.nmr.mgh.harvard.edu

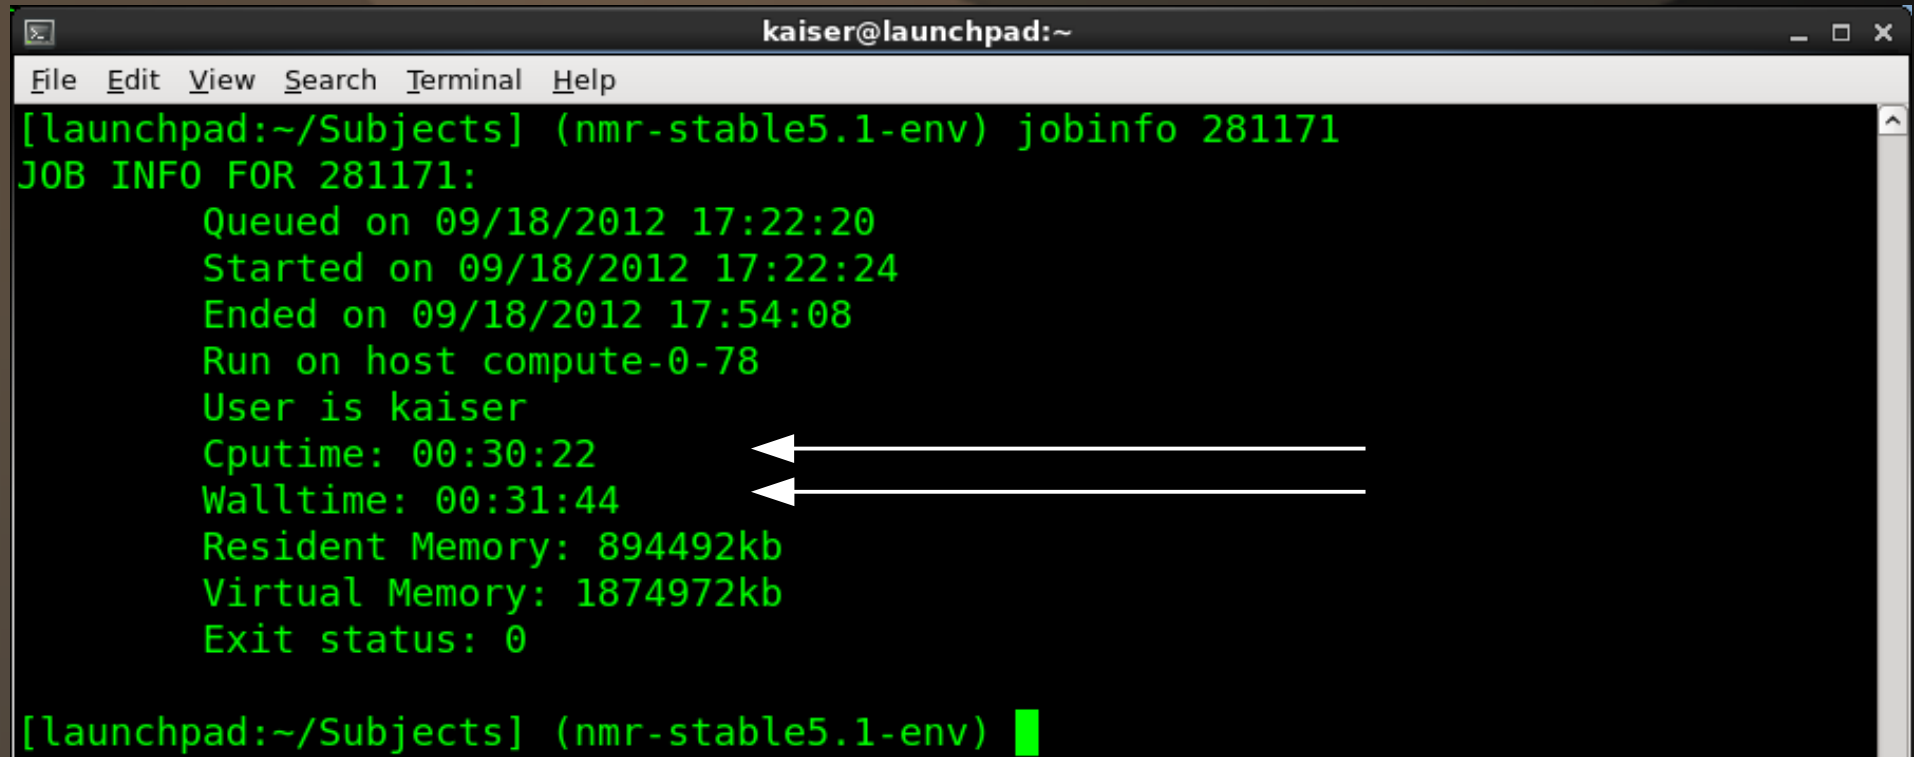
PBS Job Id: 282509.launchpad.nmr.mgh.harvard.edu
Job Name: pbsjob_5
Exec host: compute-0-64/4
Execution terminated
Exit_status=0
resources_used.cput=00:07:33
resources_used.mem=984016kb
resources_used.vmem=1964488kb
resources_used.walltime=00:08:36

[ALL of message]
? Help      M msgIndex  P PrevMsg    - PrevPage  D Delete     R Reply
0 OTHER CMDS V ViewAtch  N NextMsg    Spc NextPage  U Undelete  F Forward
```

Shows you the exit status, CPU time, walltime and the virtual memory used

Martinos Center Compute Clusters

Usage - I/O



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) jobinfo 281171  
JOB INFO FOR 281171:  
    Queued on 09/18/2012 17:22:20  
    Started on 09/18/2012 17:22:24  
    Ended on 09/18/2012 17:54:08  
    Run on host compute-0-78  
    User is kaiser  
    Cputime: 00:30:22  
    Walltime: 00:31:44  
    Resident Memory: 894492kb  
    Virtual Memory: 1874972kb  
    Exit status: 0  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

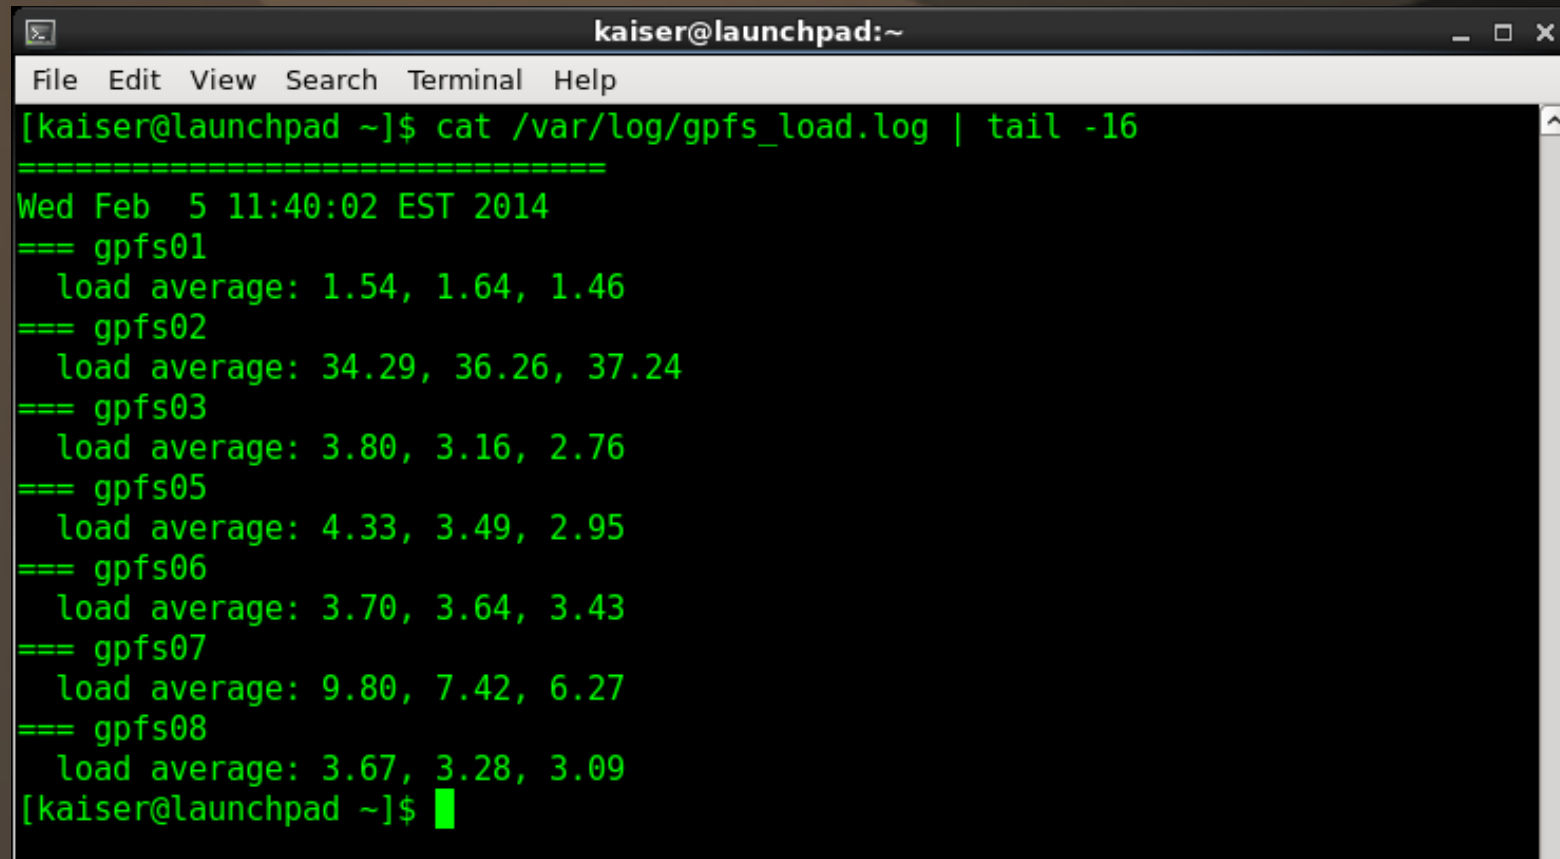
Compare CPUtime and Walltime. If Walltime is larger than CPUtime, time was wasted in I/O.

This job was run using data from my local machine. Over one minute was wasted transferring data back and forth between launchpad (in Needham) to my computer at Martinos.

Martinos Center Compute Clusters

Usage - I/O

How to check the I/O load on the cluster

A terminal window titled 'kaiser@launchpad:~' showing the command 'cat /var/log/gpfs_load.log | tail -16'. The output displays system load averages for various gpfs nodes. The node 'gpfs02' has a notably high load average of 34.29, 36.26, and 37.24, which is significantly higher than the other nodes shown.

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ cat /var/log/gpfs_load.log | tail -16  
=====  
Wed Feb  5 11:40:02 EST 2014  
=== gpfs01  
  load average: 1.54, 1.64, 1.46  
=== gpfs02  
  load average: 34.29, 36.26, 37.24  
=== gpfs03  
  load average: 3.80, 3.16, 2.76  
=== gpfs05  
  load average: 4.33, 3.49, 2.95  
=== gpfs06  
  load average: 3.70, 3.64, 3.43  
=== gpfs07  
  load average: 9.80, 7.42, 6.27  
=== gpfs08  
  load average: 3.67, 3.28, 3.09  
[kaiser@launchpad ~]$
```

If the load average is above 30 (gpfs02), the load is too high. We are automatically notified and have to investigate the reason. If the problem is caused by I/O on launchpad we find the user with the largest I/O need, delete all their jobs and monitor the situation.

Martinos Center Compute Clusters

Usage - I/O

Is it your jobs causing the problem?

ssh to a node
your jobs are
running on and
run the command
'top'

```
kaiser@compute-0-105:~  
File Edit View Search Terminal Help  
top - 11:48:04 up 8 days, 18:50, 1 user, load average: 8.09, 7.49, 4.53  
Tasks: 272 total, 2 running, 270 sleeping, 0 stopped, 0 zombie  
Cpu(s): 13.8%us, 2.1%sy, 0.0%ni, 83.8%id, 0.2%wa, 0.0%hi, 0.1%si, 0.0%st  
Mem: 32877400k total, 5127336k used, 27750064k free, 235876k buffers  
Swap: 67108856k total, 0k used, 67108856k free, 2140420k cached  
  
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND  
23293 mueller 20 0 1218m 1.0g 848 R 100.0 3.2 0:29.88 gauss_nifti  
9056 root 0 -20 5282m 586m 20m S 22.3 1.8 30:37.81 mmfsd  
22928 mueller 20 0 341m 74m 20m D 2.0 0.2 0:10.90 MATLAB  
22825 mueller 20 0 341m 76m 20m D 1.7 0.2 0:10.89 MATLAB  
22886 mueller 20 0 341m 76m 20m D 1.7 0.2 0:10.92 MATLAB  
22573 mueller 20 0 341m 78m 20m D 1.3 0.2 0:12.16 MATLAB  
22659 mueller 20 0 341m 74m 20m D 1.3 0.2 0:11.15 MATLAB  
22769 mueller 20 0 341m 74m 20m D 1.3 0.2 0:10.92 MATLAB  
8868 root 20 0 0 0 0 S 0.3 0.0 0:01.22 gpfsSwapdKproc  
9097 root 20 0 49036 27m 2660 S 0.3 0.1 4:57.15 pbs_mom  
22714 mueller 20 0 340m 76m 20m D 0.3 0.2 0:10.56 MATLAB  
23351 kaiser 20 0 21476 1552 1084 R 0.3 0.0 0:00.04 top  
1 root 20 0 25684 1660 1320 S 0.0 0.0 0:02.18 init
```

Look at the %CPU and S (state) columns. If your jobs are using less than 100% of it's CPU and is in the 'D' state, the job is stuck waiting for I/O. In this case, only run 20 jobs at a time. Use the max20 queue.

Martinos Center Compute Clusters

Usage - I/O

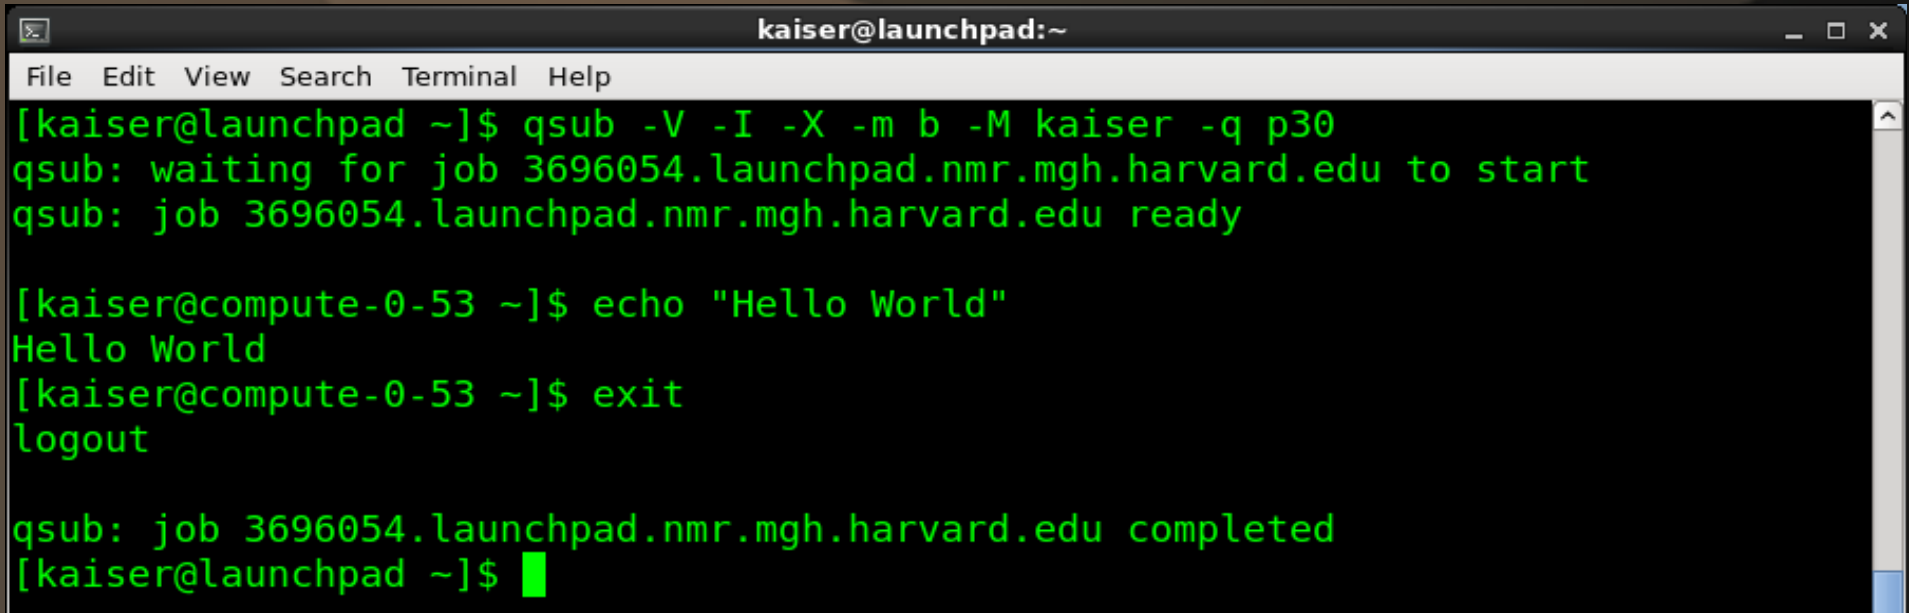
Tips on dealing with this bottleneck:

- Copy/move local data to /cluster/ directories before running jobs.
- Have scripts/programs write temp data to /cluster/scratch/
- Instead of launchpad, use tensor which lives in CNY.
- Space out submission of jobs so they don't all have large I/O needs at the same time.

For instance, the nu_correct step in FS uses CNY local data. If you have 150 jobs all reach this point at the same time, there are 150 separate requests for CNY data. This load is unmanageable. Since this step is in the beginning of the stream and only takes a minute, space out each job submission by a couple minutes. Or submit ten jobs at a time, wait ten minutes, then submit the next batch.

Martinos Center Compute Clusters

Usage – Interactive Jobs

A terminal window titled 'kaiser@launchpad:~' with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal shows the following sequence of commands and output:

```
[kaiser@launchpad ~]$ qsub -V -I -X -m b -M kaiser -q p30
qsub: waiting for job 3696054.launchpad.nmr.mgh.harvard.edu to start
qsub: job 3696054.launchpad.nmr.mgh.harvard.edu ready

[kaiser@compute-0-53 ~]$ echo "Hello World"
Hello World
[kaiser@compute-0-53 ~]$ exit
logout

qsub: job 3696054.launchpad.nmr.mgh.harvard.edu completed
[kaiser@launchpad ~]$
```

Use the `qsub` command to start an interactive job using the high priority `p30` queue. You will receive an email when the job begins execution. **Replace 'kaiser' with your username!** Actively wait until the job is slated for execution. Don't immediately leave for lunch.

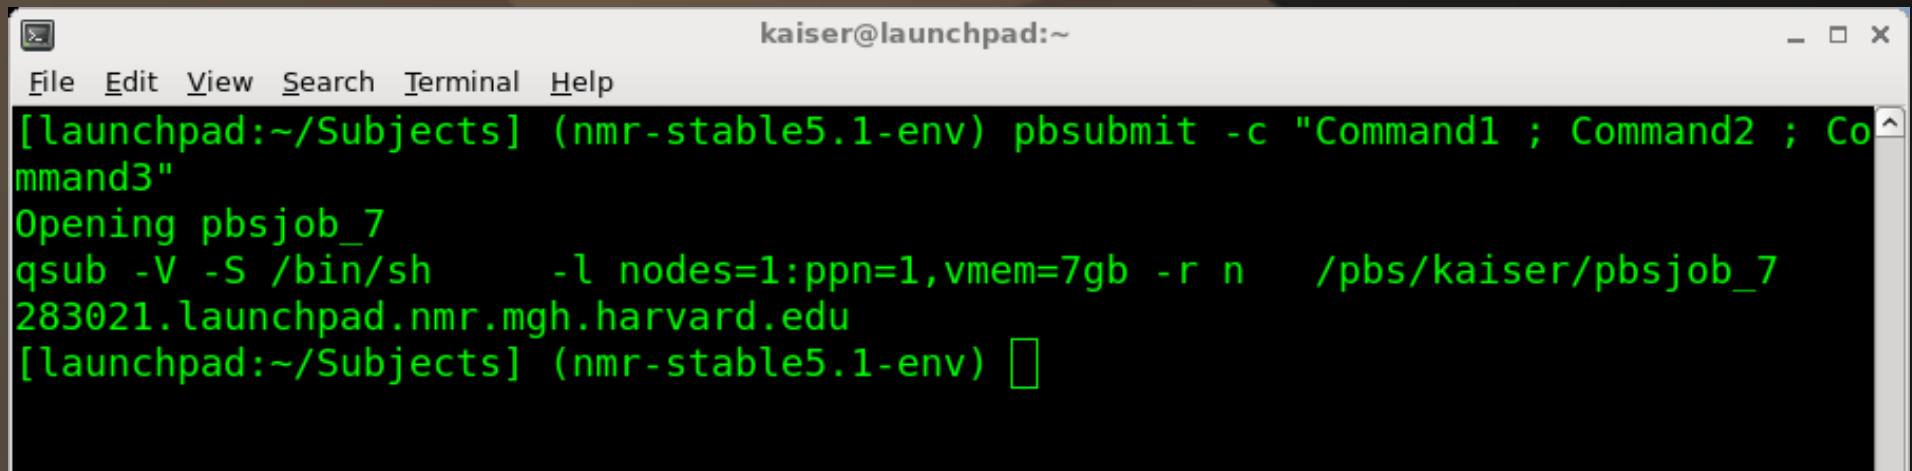
1. As soon as a slot becomes available, the job is assigned a Job ID and you are ssh'd to the node where your job will execute.
2. Run your commands...
3. When completed, exit out of the node. Your job will not be completed until you exit.

Please attend to an interactive session. As soon as the job begins and you are ssh'd into the node, you take up a job slot. Exit out of the node as soon as your commands are done. You will continue to take up a job slot until you exit out of the node.

Martinos Center Compute Clusters

Usage – Dependencies – Daisy Chain

If you have a series of commands that you want to execute in a row (one after another). The easiest way to do it is to daisy chain the commands together on the command line:

A terminal window titled 'kaiser@launchpad:~' with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal shows a command being entered: [launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -c "Command1 ; Command2 ; Command3". The output shows 'Opening pbsjob_7' and 'qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_7 283021.launchpad.nmr.mgh.harvard.edu'. The prompt returns to [launchpad:~/Subjects] (nmr-stable5.1-env) with a cursor.

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -c "Command1 ; Command2 ; Command3"  
Opening pbsjob_7  
qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_7  
283021.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

The commands are separated on the command line by a semicolon (;).

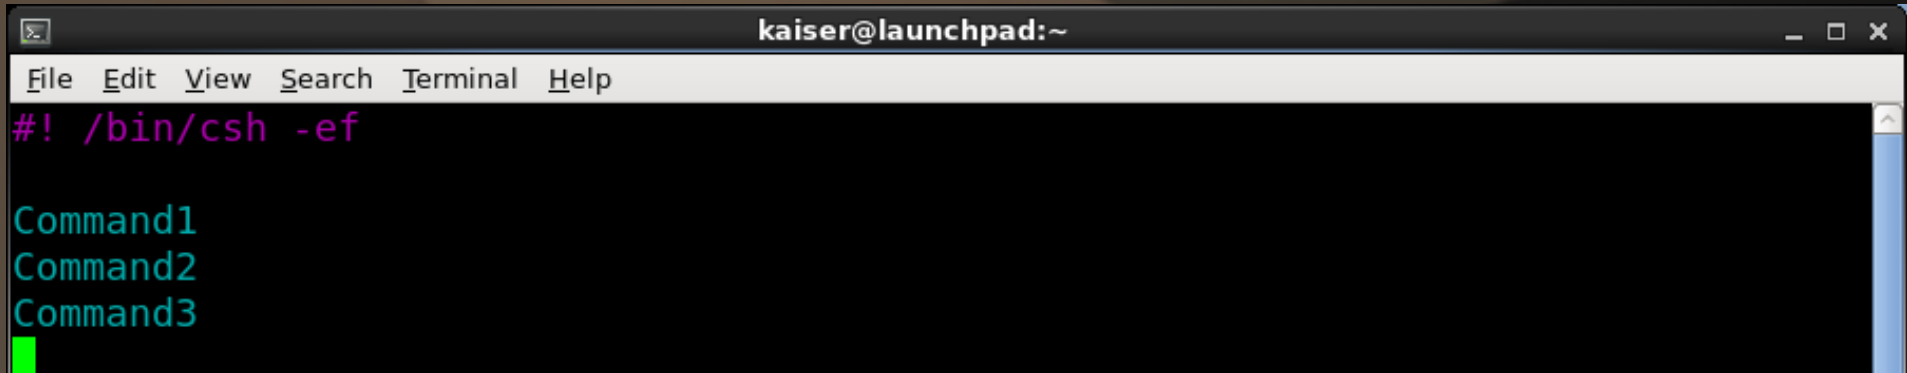
Each command will run even if the one before it failed.

Replace Command1, Command2, Command3 with the specific commands you want to run

Martinos Center Compute Clusters

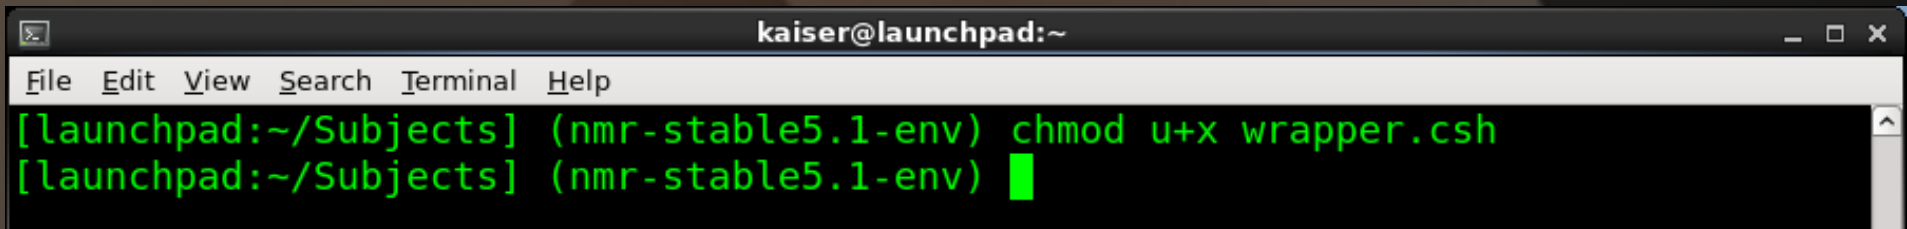
Usage – Dependencies – Wrapper Script

A more elegant way to do it is to write a wrapper script. Use a text editor to create a file called wrapper.csh with these contents:



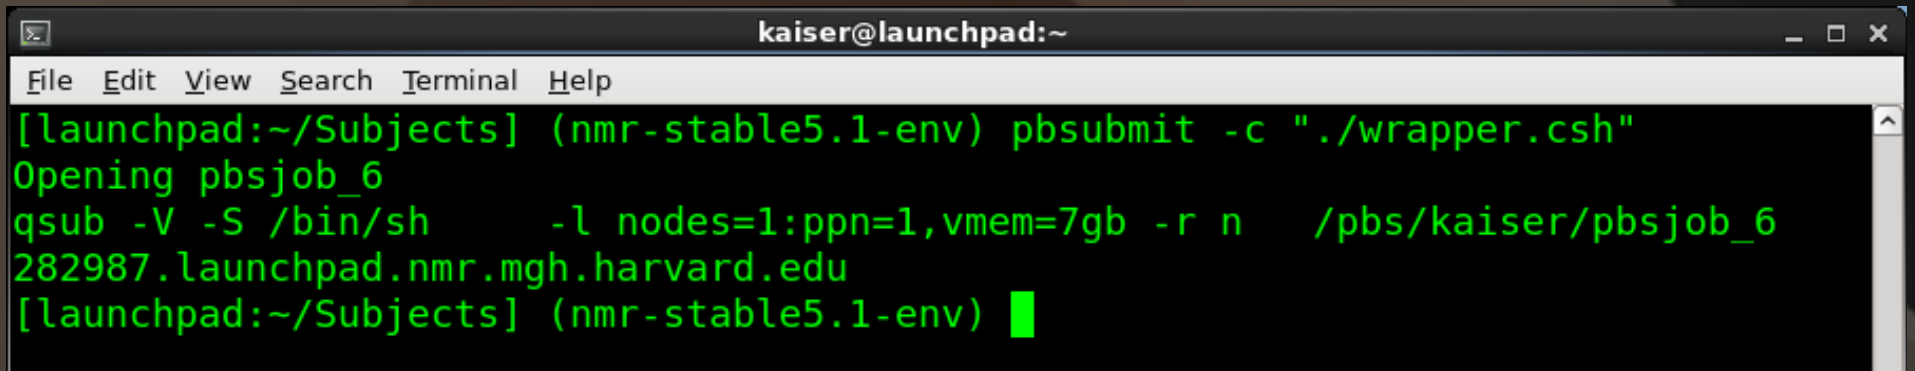
```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
#!/bin/csh -ef  
  
Command1  
Command2  
Command3  
█
```

The -e flag above instructs the script to exit if any of the individual commands exit with an error. Make the script executable:



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) chmod u+x wrapper.csh  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Submit the script for execution:

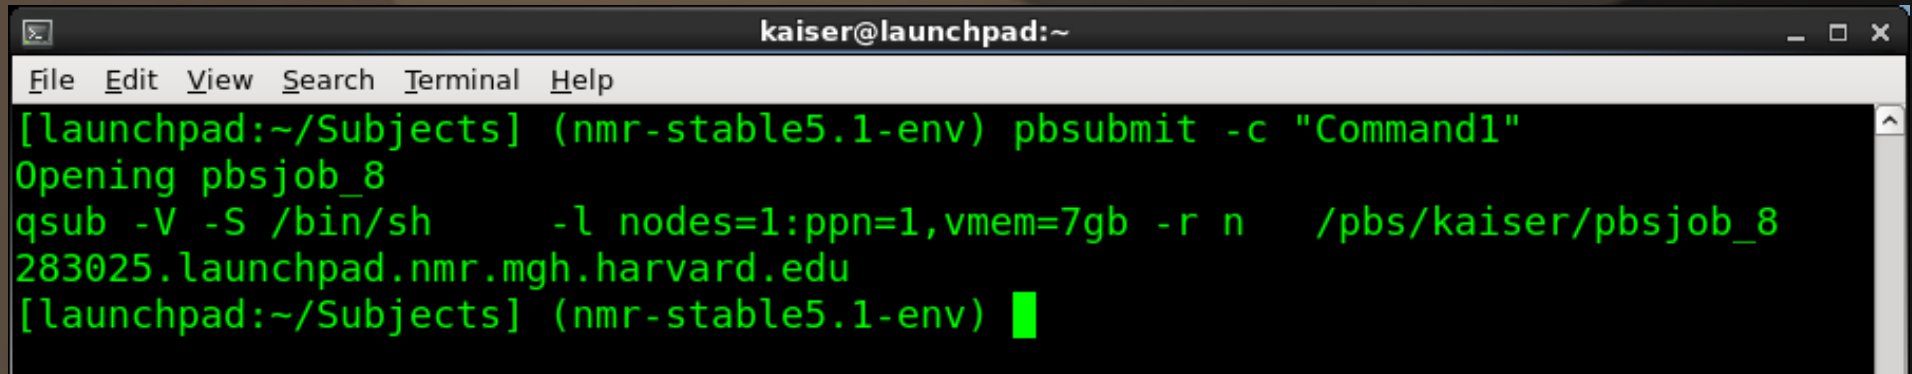


```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -c "./wrapper.csh"  
Opening pbsjob_6  
qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_6  
282987.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Martinos Center Compute Clusters

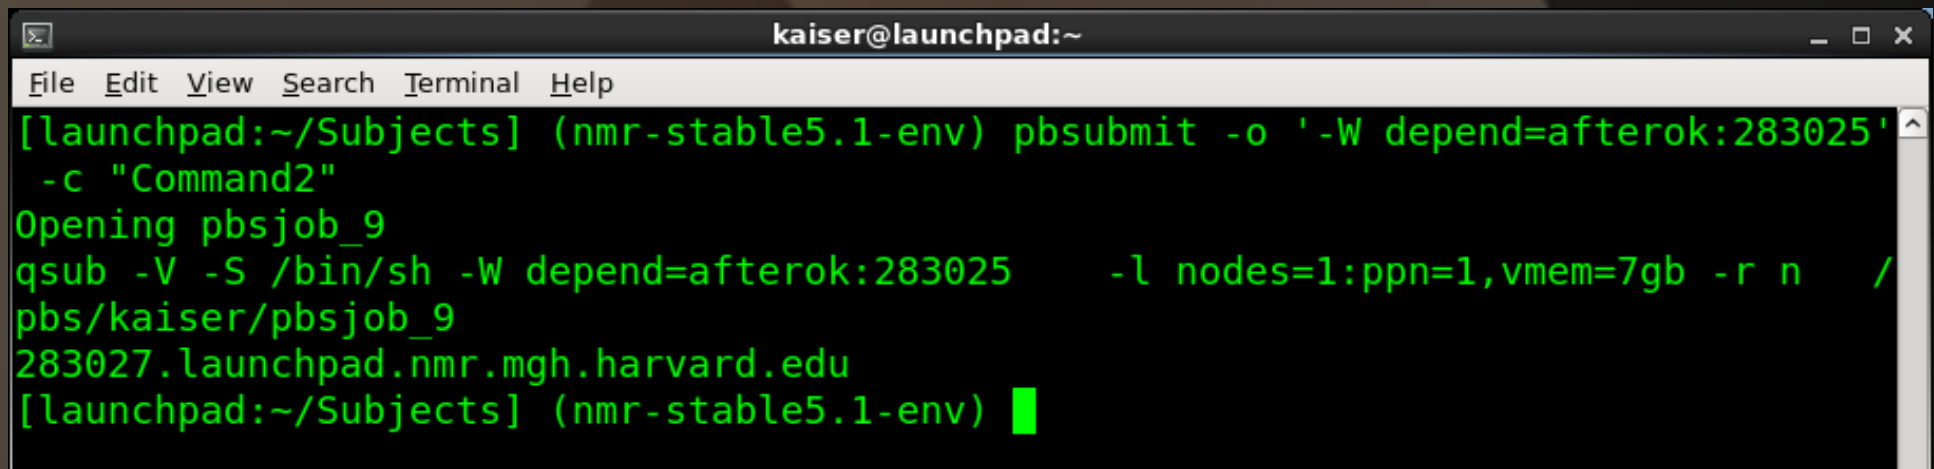
Usage – Dependencies – In Progress

If you already have a job running....



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -c "Command1"  
Opening pbsjob_8  
qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_8  
283025.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

And you want to start another job that will run immediately after the first job completes without error:



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -o '-W depend=afterok:283025'  
-c "Command2"  
Opening pbsjob_9  
qsub -V -S /bin/sh -W depend=afterok:283025 -l nodes=1:ppn=1,vmem=7gb -r n /  
pbs/kaiser/pbsjob_9  
283027.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

This second job will be held until the first one completes without error. If the first job exits with an error, the second job will not run.

Martinos Center Compute Clusters

Job Status - Running Jobs – Show Job Status

Job ID - Job Name - User - CPUtime - State - Queue

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) qstat  
283568.launchpad pbsjob_201 krienen 02:57:33 R default  
283569.launchpad pbsjob_252 khoa 02:46:16 R default  
283604.launchpad pbsjob_34 spaeth 01:47:00 R default  
283605.launchpad pbsjob_35 spaeth 01:45:58 R default  
283617.launchpad ...0_real011.txt slbowen 00:34:47 R default  
283672.launchpad pbsjob_2 kbickart 00:27:59 R max10  
283673.launchpad pbsjob_3 kbickart 00:20:10 R max10  
283675.launchpad pbsjob_4 kbickart 00:20:22 R max10  
283676.launchpad pbsjob_5 kbickart 00:20:26 R max10  
283677.launchpad pbsjob_6 kbickart 00:18:20 R max10  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Additional options:

To see just your jobs:

```
qstat -u <username>
```

```
qstat | grep -w <username>
```

To get all your running and queued jobs:

```
qstat | grep -w <username> | grep -w R
```

```
qstat | grep -w <username> | grep -w Q
```

States:

[R]unning

[Q]ueued

[H]eld

Martinos Center Compute Clusters

Job Status - Running Jobs – Show Job Status

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ showq  
ACTIVE JOBS-----  
JOBNAME          USERNAME          STATE  PROC  REMAINING          STARTTIME  
3679935          shou             Running  1    1:06:30:37  Mon Feb  3 01:58:40  
  
IDLE JOBS-----  
JOBNAME          USERNAME          STATE  PROC  WCLIMIT           QUEUETIME  
3689589          ville            Idle    8    4:00:00:00  Wed Feb  5 16:24:10  
  
BLOCKED JOBS-----  
JOBNAME          USERNAME          STATE  PROC  WCLIMIT           QUEUETIME  
3686742          wachinge        Idle    1    4:00:00:00  Wed Feb  5 04:14:17  
[kaiser@launchpad ~]$ █
```

Another Job Status command is 'showq':
It shows the Active (running), Idle
(queued) and Blocked (held) jobs

To see just your jobs:
showq -u <username>

To see all your running and idle jobs:
showq -r -u <username>
showq -i -u <username>

Martinos Center Compute Clusters

Job Status - Running Jobs – See Standard Output

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ pbsubmit -c "sleep 120 ; echo 'Hello World' ; sleep 120"  
Opening pbsjob_13  
qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_13  
3689806.launchpad.nmr.mgh.harvard.edu  
[kaiser@launchpad ~]$
```

Job is running:

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ qstat | grep -w kaiser  
3689806.launchpad pbsjob_13 kaiser 00:00:00 R default  
[kaiser@launchpad ~]$
```

Check on the standard output of the job:

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ jobinfo -o 3689806.launchpad  
Hello World  
[kaiser@launchpad ~]$
```


To see the standard error of an actively running job; 'jobinfo -e <Job ID>'

Martinos Center Compute Clusters

Job Status - Completed Jobs

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ jobinfo 3689806  
JOB INFO FOR 3689806:  
    Queued on 02/05/2014 18:17:52  
    Started on 02/05/2014 18:20:48  
    Ended on 02/05/2014 18:24:50  
    Run on host compute-0-106  
    User is kaiser  
    Cputime: 00:00:00  
    Walltime: 00:04:02  
    Resident Memory: 3496kb  
    Virtual Memory: 321564kb  
    Exit status: 0  
[kaiser@launchpad ~]$
```

Check the Exit Status of the job. Zero means it completed without errors.



The job script, standard output, standard error and the exit status are all saved as separate text files in your pbs directory:

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ ls /pbs/kaiser/pbsjob_13*  
/pbs/kaiser/pbsjob_13          /pbs/kaiser/pbsjob_13.o3689806  
/pbs/kaiser/pbsjob_13.e3689806 /pbs/kaiser/pbsjob_13.status  
/pbs/kaiser/pbsjob_13.env  
[kaiser@launchpad ~]$
```

Martinos Center Compute Clusters

Job Status - Failed Jobs

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) pbsubmit -c "BadCommand"  
Opening pbsjob_13  
qsub -V -S /bin/sh -l nodes=1:ppn=1,vmem=7gb -r n /pbs/kaiser/pbsjob_13  
283853.launchpad.nmr.mgh.harvard.edu  
[launchpad:~/Subjects] (nmr-stable5.1-env) jobinfo 283853  
JOB INFO FOR 283853:  
    Queued on 09/20/2012 11:34:11  
    Started on 09/20/2012 11:34:14  
    Ended on 09/20/2012 11:34:14  
    Run on host compute-0-48  
    User is kaiser  
    Cputime: 00:00:00  
    Walltime: 00:00:00  
    Resident Memory: 0kb  
    Virtual Memory: 0kb  
    Exit status: 127  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

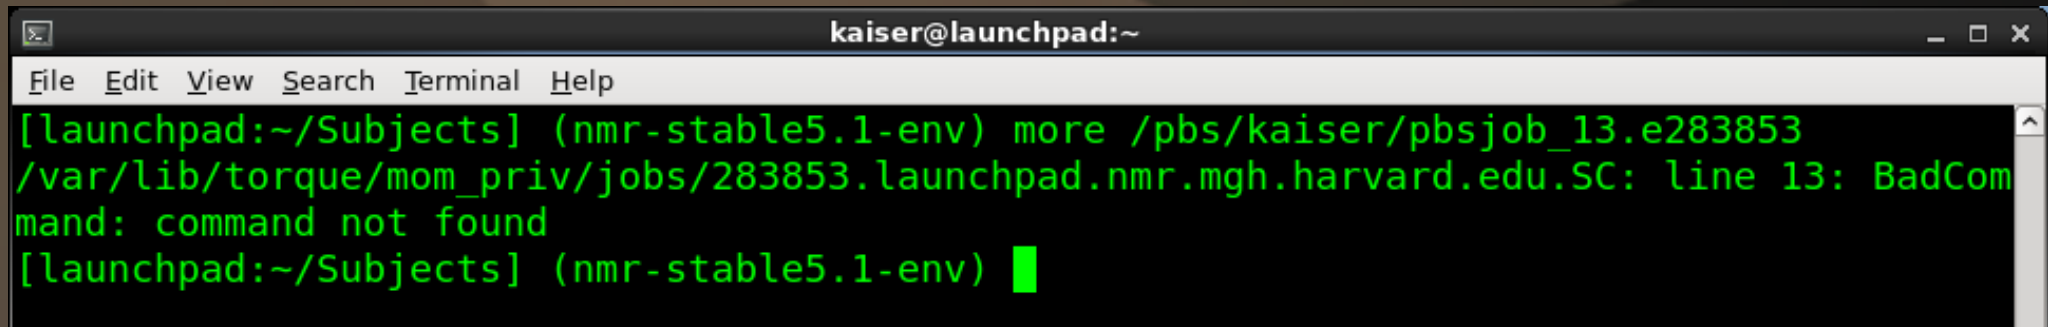
Ack! My job finished with an Exit Status of 127.

How do I troubleshoot???

Martinos Center Compute Clusters

Job Status - Failed Jobs

Check the standard error and standard output files for any hints:



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[launchpad:~/Subjects] (nmr-stable5.1-env) more /pbs/kaiser/pbsjob_13.e283853  
/var/lib/torque/mom_priv/jobs/283853.launchpad.nmr.mgh.harvard.edu.SC: line 13: BadCom  
mand: command not found  
[launchpad:~/Subjects] (nmr-stable5.1-env) █
```

Other Possible Hints:

Resource Related

Check vmem is under the requested amount (default: 7GB)

Check walltime is under the requested amount (default: 96 hours)

Command Related

Check standard error and standard output files!!

Check standard error and standard output files (again)!!

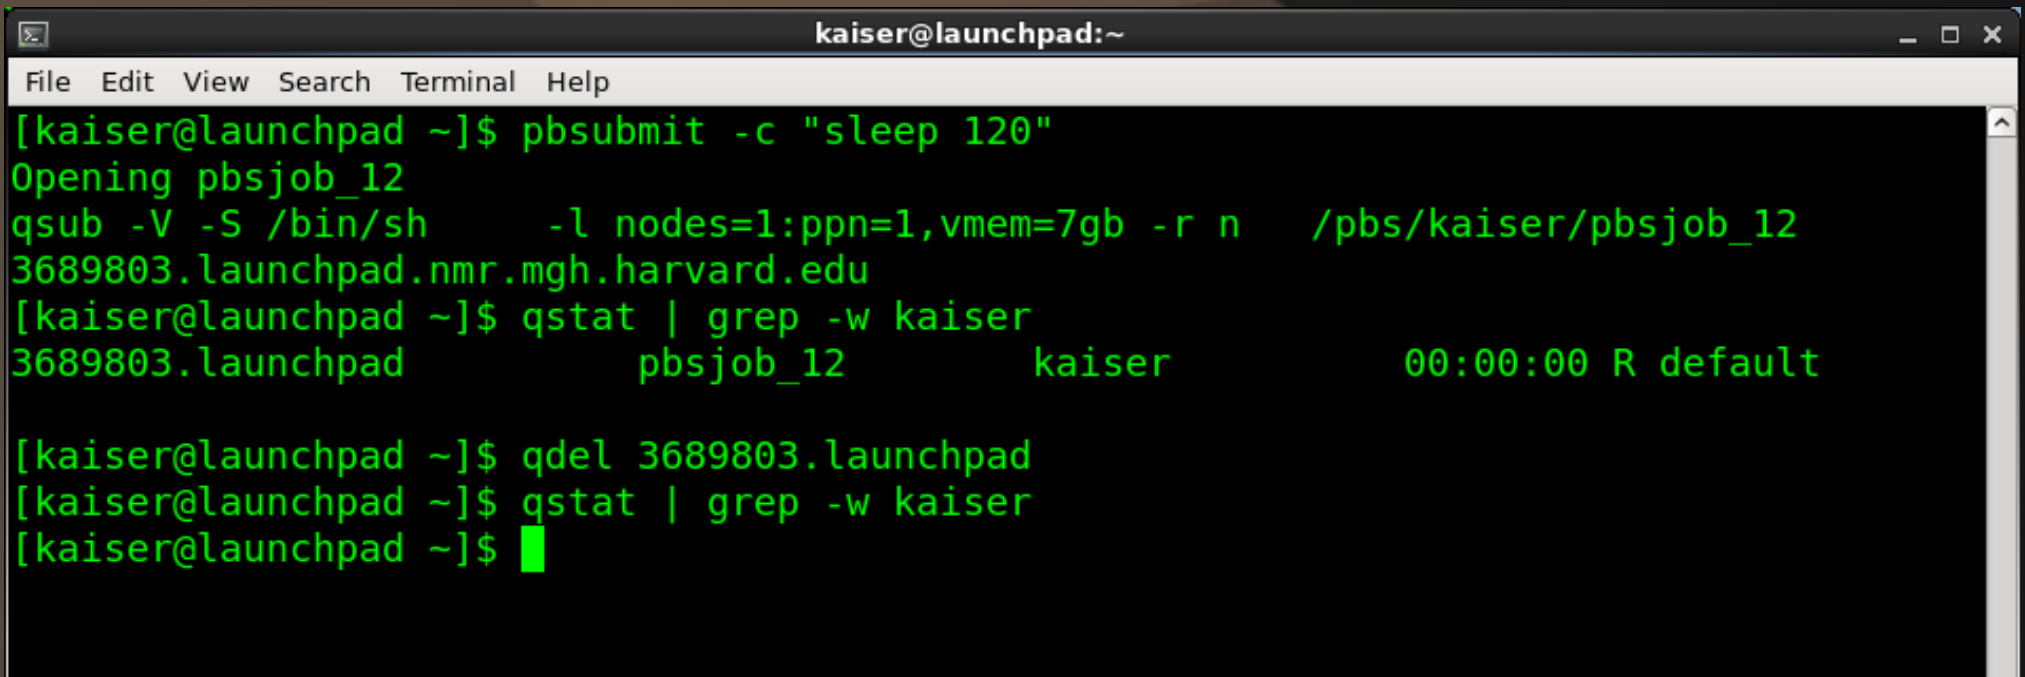
If the program is home-made, was it compiled for the launchpad architecture?

Test-run the command locally. If it breaks, the problem is probably not with the cluster.

Martinos Center Compute Clusters

Job Status - Delete Jobs

You submit a job, realize there is a mistake and want to delete it:

A terminal window titled 'kaiser@launchpad:~' with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal shows the following sequence of commands and output:

```
[kaiser@launchpad ~]$ pbsubmit -c "sleep 120"
Opening pbsjob_12
qsub -V -S /bin/sh      -l nodes=1:ppn=1,vmem=7gb -r n    /pbs/kaiser/pbsjob_12
3689803.launchpad.nmr.mgh.harvard.edu
[kaiser@launchpad ~]$ qstat | grep -w kaiser
3689803.launchpad      pbsjob_12      kaiser      00:00:00 R default

[kaiser@launchpad ~]$ qdel 3689803.launchpad
[kaiser@launchpad ~]$ qstat | grep -w kaiser
[kaiser@launchpad ~]$
```

Martinos Center Compute Clusters

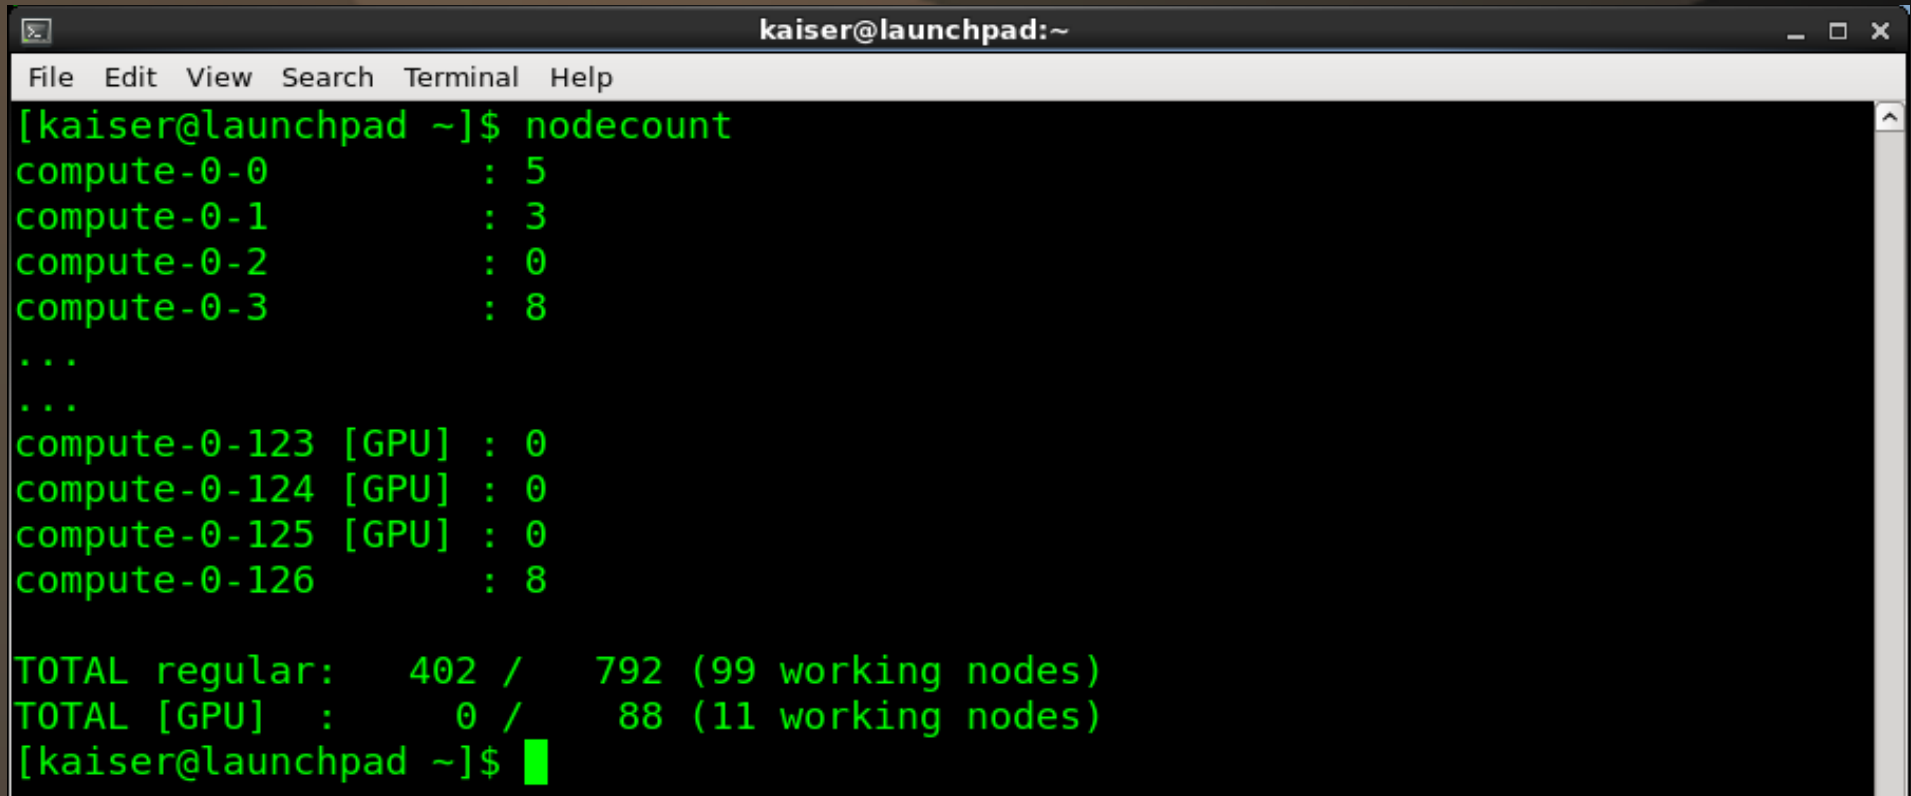
Job Status - Idle Jobs

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ showq -i  
JobName      Priority  XFactor  Q      User      Group  Procs   WCLimit   Class      SystemQueueTime  
3689589*     101145   1.0     -      ville     ville   8       4:00:00:00 default    Wed Feb  5 16:24:10  
3689591     101145   1.0     -      ville     ville   8       4:00:00:00 default    Wed Feb  5 16:24:24  
3689596     101139   1.0     -      ville     ville   8       4:00:00:00 default    Wed Feb  5 16:29:59  
3689607     101134   1.0     -      ville     ville   8       4:00:00:00 default    Wed Feb  5 16:34:46  
3686736     101018   1.0     -      wachinge  wachinge 1       4:00:00:00 max50     Wed Feb  5 18:30:36  
3686737     101018   1.0     -      wachinge  wachinge 1       4:00:00:00 max50     Wed Feb  5 18:30:36  
3686738     101014   1.0     -      wachinge  wachinge 1       4:00:00:00 max50     Wed Feb  5 18:35:14  
3687580     101004   1.0     -      clarsen   clarsen   1       4:00:00:00 max50     Wed Feb  5 18:45:08  
3687581     101004   1.0     -      clarsen   clarsen   1       4:00:00:00 max50     Wed Feb  5 18:45:08  
3687582     101004   1.0     -      clarsen   clarsen   1       4:00:00:00 max50     Wed Feb  5 18:45:08  
3687583     101004   1.0     -      clarsen   clarsen   1       4:00:00:00 max50     Wed Feb  5 18:45:08  
3686739     101002   1.0     -      wachinge  wachinge 1       4:00:00:00 max50     Wed Feb  5 18:46:34  
3689811     80006    1.0     -      slbowen   slbowen   1       4:00:00:00 max200    Wed Feb  5 18:42:50  
3689812     80006    1.0     -      slbowen   slbowen   1       4:00:00:00 max200    Wed Feb  5 18:42:50  
3689813     80006    1.0     -      slbowen   slbowen   1       4:00:00:00 max200    Wed Feb  5 18:42:50  
3689814     80006    1.0     -      slbowen   slbowen   1       4:00:00:00 max200    Wed Feb  5 18:42:50  
Jobs: 16  Total Backlog: 4224.00 ProcHours (4.63 Hours)  
[kaiser@launchpad ~]$ █
```

Shows the Idle Queue. Jobs in this state are waiting for slots to open up. They are sorted by 'Priority'. Priority is determined by the queue (default, max50 etc.) and how long it has been waiting.

Martinos Center Compute Clusters

Misc Commands - nodelist



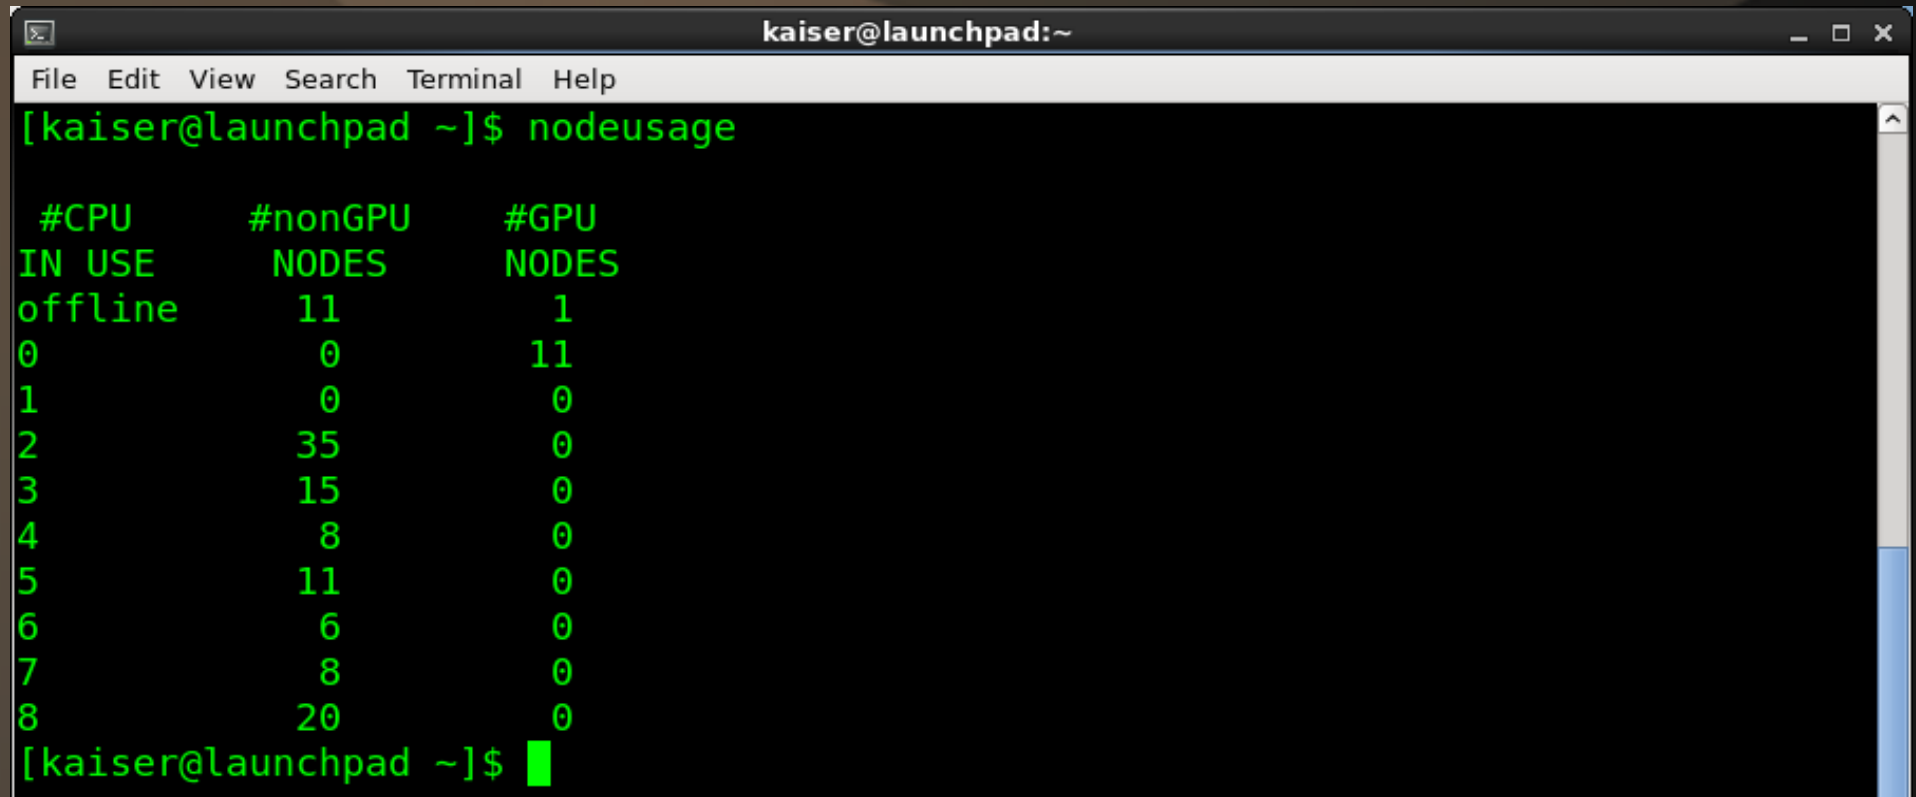
```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ nodelist  
compute-0-0          : 5  
compute-0-1          : 3  
compute-0-2          : 0  
compute-0-3          : 8  
...  
...  
compute-0-123 [GPU] : 0  
compute-0-124 [GPU] : 0  
compute-0-125 [GPU] : 0  
compute-0-126      : 8  
  
TOTAL regular:    402 / 792 (99 working nodes)  
TOTAL [GPU]      :    0 / 88 (11 working nodes)  
[kaiser@launchpad ~]$
```

Shows the number of CPUs that are occupied for each compute node.

Helpful in seeing how busy the cluster is and how many job slots are available.

Martinos Center Compute Clusters

Misc Commands - nodeusage



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ nodeusage  
  
#CPU      #nonGPU    #GPU  
IN USE    NODES      NODES  
offline   11         1  
0         0         11  
1         0         0  
2        35         0  
3        15         0  
4         8         0  
5        11         0  
6         6         0  
7         8         0  
8        20         0  
[kaiser@launchpad ~]$
```

Condenses the output of `nodecount` from displaying the number of CPUs in use for each node, into a summary of the number of nodes that have a number of CPUs occupied.

For example, there are zero nodes with zero jobs running (completely free) and there are 35 nodes with only two CPUs in use (can support six job slots).

A similar command is called 'freenodes' (thanks Doug Greve). Instead of counting the nodes with CPUs that are busy, it counts the number of nodes that have free CPUs.

Martinos Center Compute Clusters

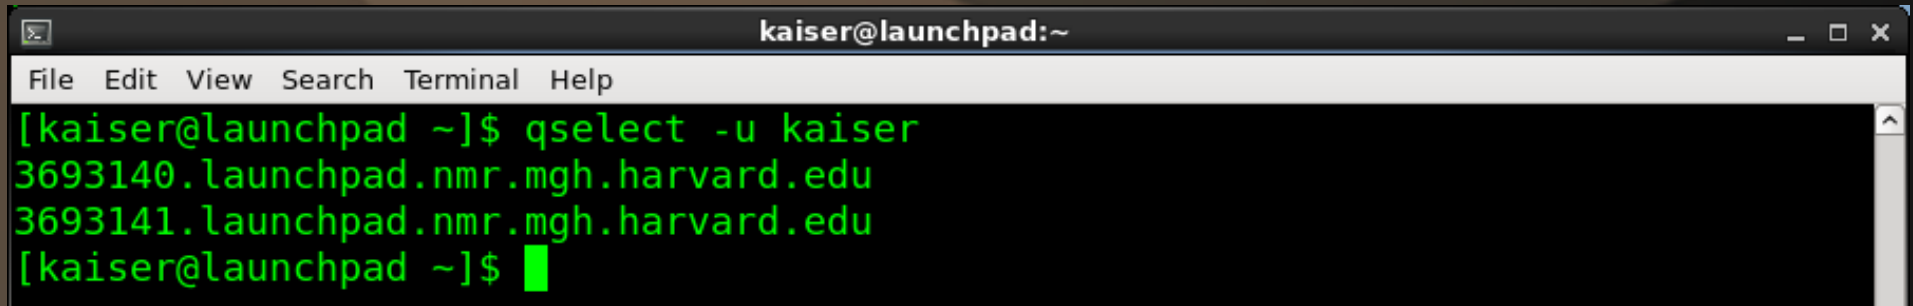
Misc Commands - usercount

```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ usercount  
                Running          Queued  
USER           Jobs / CPUS      Jobs / CPUS  
adagley         2/5                0/  
clarsen         51/51             548/548  
ganz            10/10             0/  
jonp            1/8                0/  
lzollei        8/8                0/  
martab         1/8                0/  
pwilkens       6/6                0/  
ryu            8/8                0/  
seesaw         2/2                0/  
shou           83/83             0/  
slbowen        13/16             64/64  
ville          2/16             10/80  
vinke          2/2                0/  
wachinge       50/50            149/149  
xnat           192/192           0/  
[kaiser@launchpad ~]$
```

Displays the number of jobs and CPUs being used by each user. Also counts the idle (queued) ones as well.

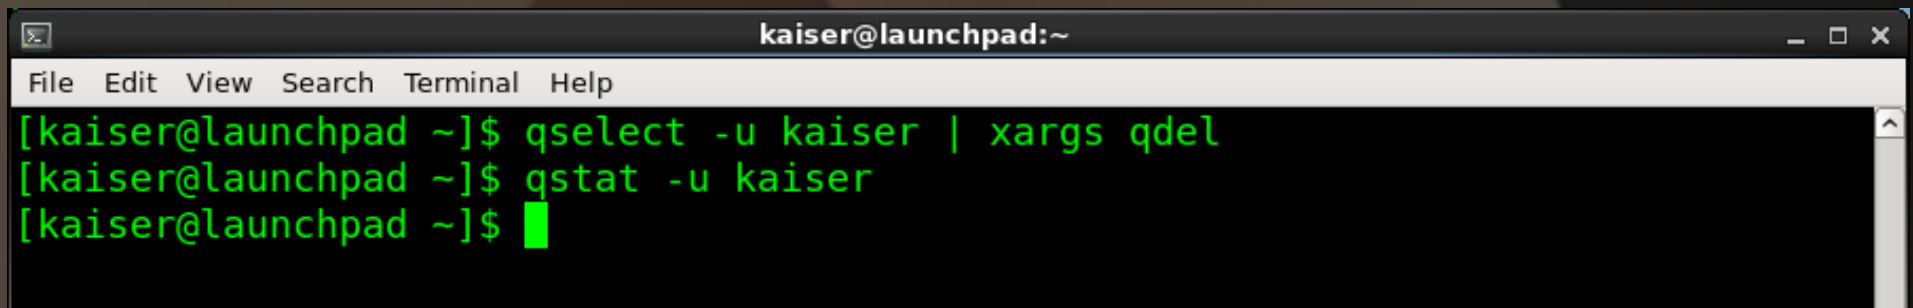
Martinos Center Compute Clusters

Misc Commands - qselect



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ qselect -u kaiser  
3693140.launchpad.nmr.mgh.harvard.edu  
3693141.launchpad.nmr.mgh.harvard.edu  
[kaiser@launchpad ~]$ █
```

Displays all the JobIDs for a user. This command can accept many parameters to filter out certain job states, resources, queues etc.



```
kaiser@launchpad:~  
File Edit View Search Terminal Help  
[kaiser@launchpad ~]$ qselect -u kaiser | xargs qdel  
[kaiser@launchpad ~]$ qstat -u kaiser  
[kaiser@launchpad ~]$ █
```

It can be very useful if you need access to a set of JobIDs.

Martinos Center Compute Clusters

Summary

Troubleshooting steps:

If a job exits in error, check the standard output and standard error files in your pbs directory.

Search the launchpad page for help and guidelines:
<http://www.nmr.mgh.harvard.edu/martinos/userInfo/computer/launchpad.php>

Try running the command locally on your own machine. If you receive the same errors, the problem is not with the cluster.

Send email to batch-users mailing list

Any Questions?

Martinos Center Compute Clusters

Adios