CGRB Computational Infrastructure

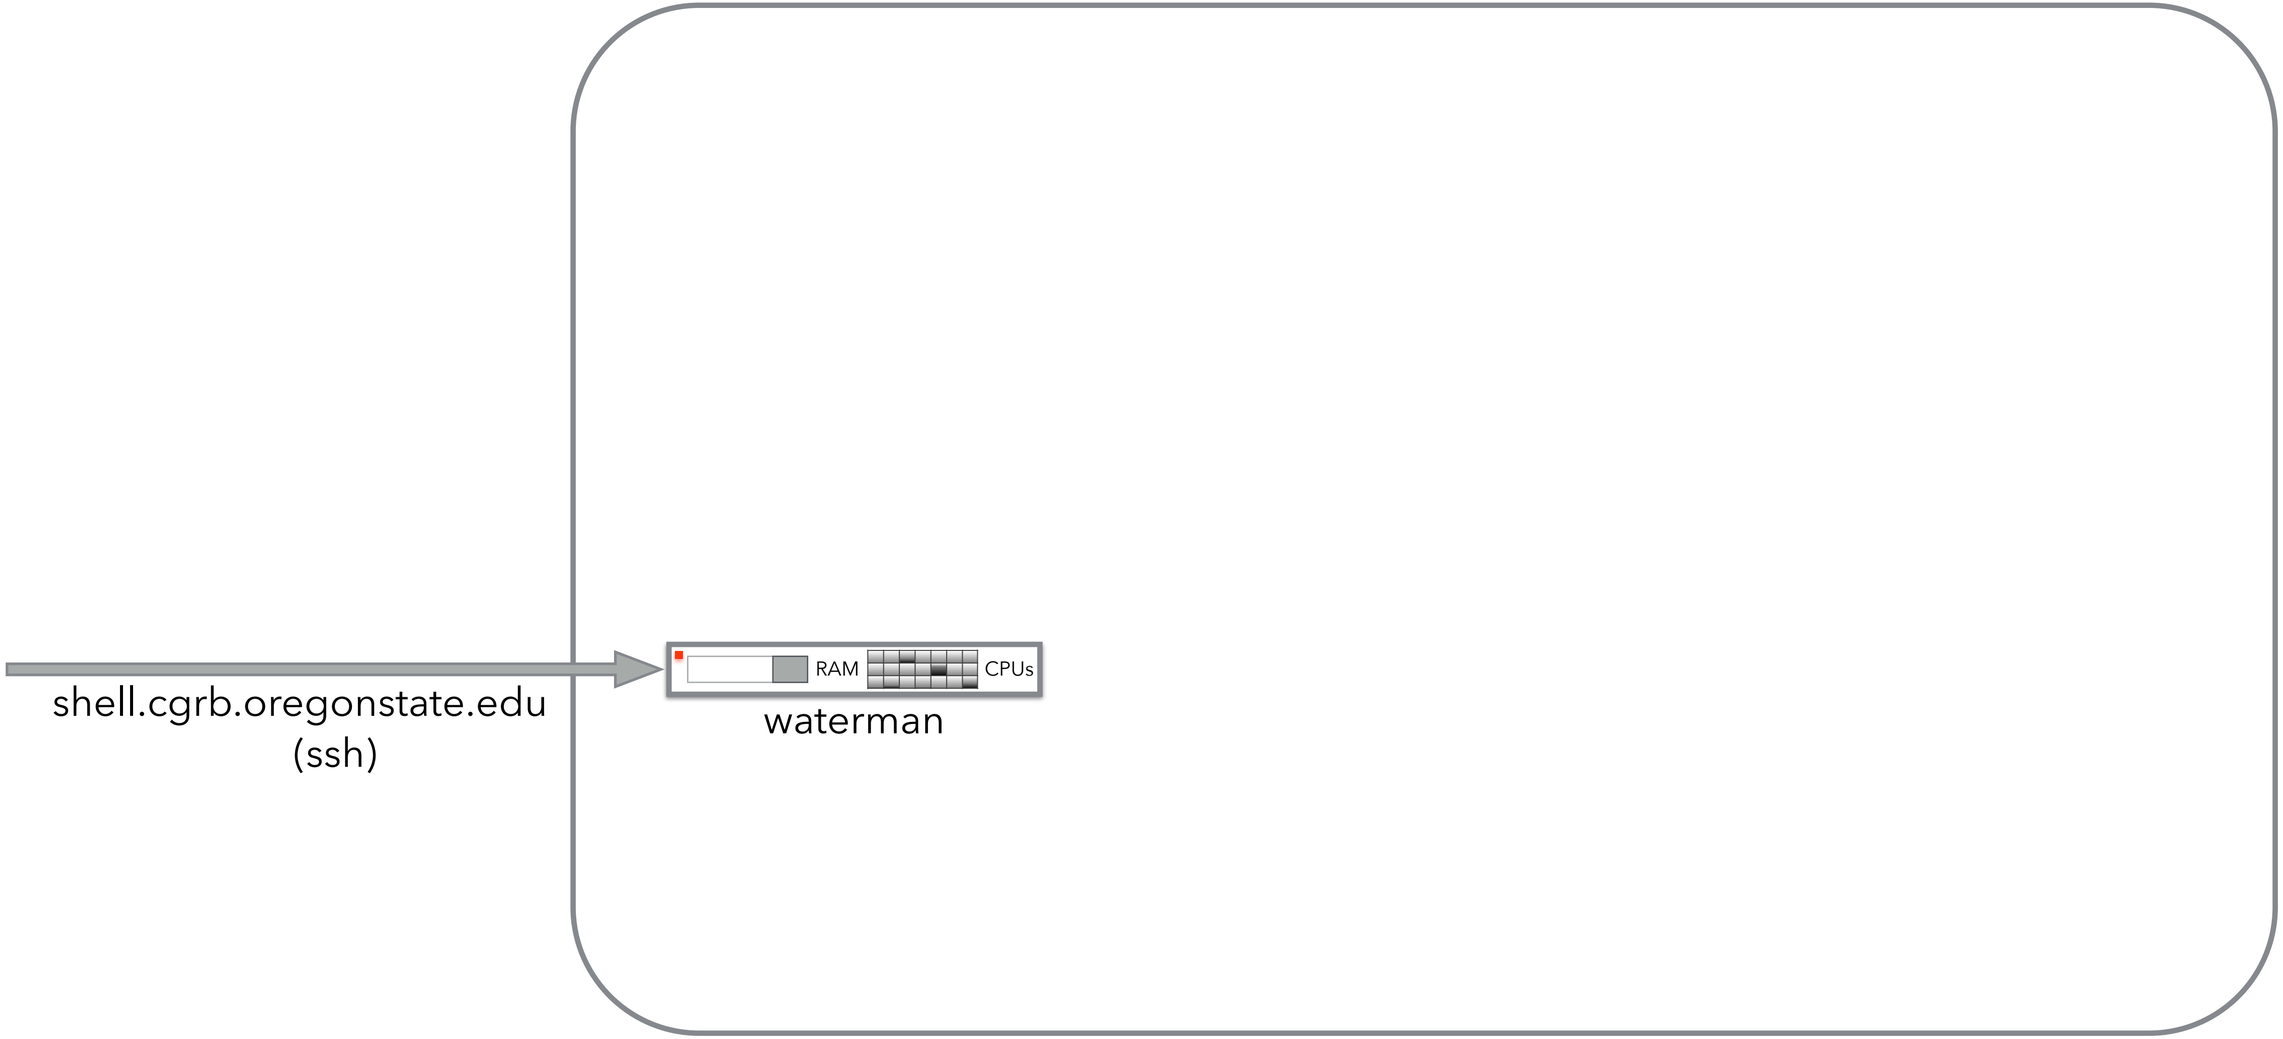

CGRB Computational Infrastructure

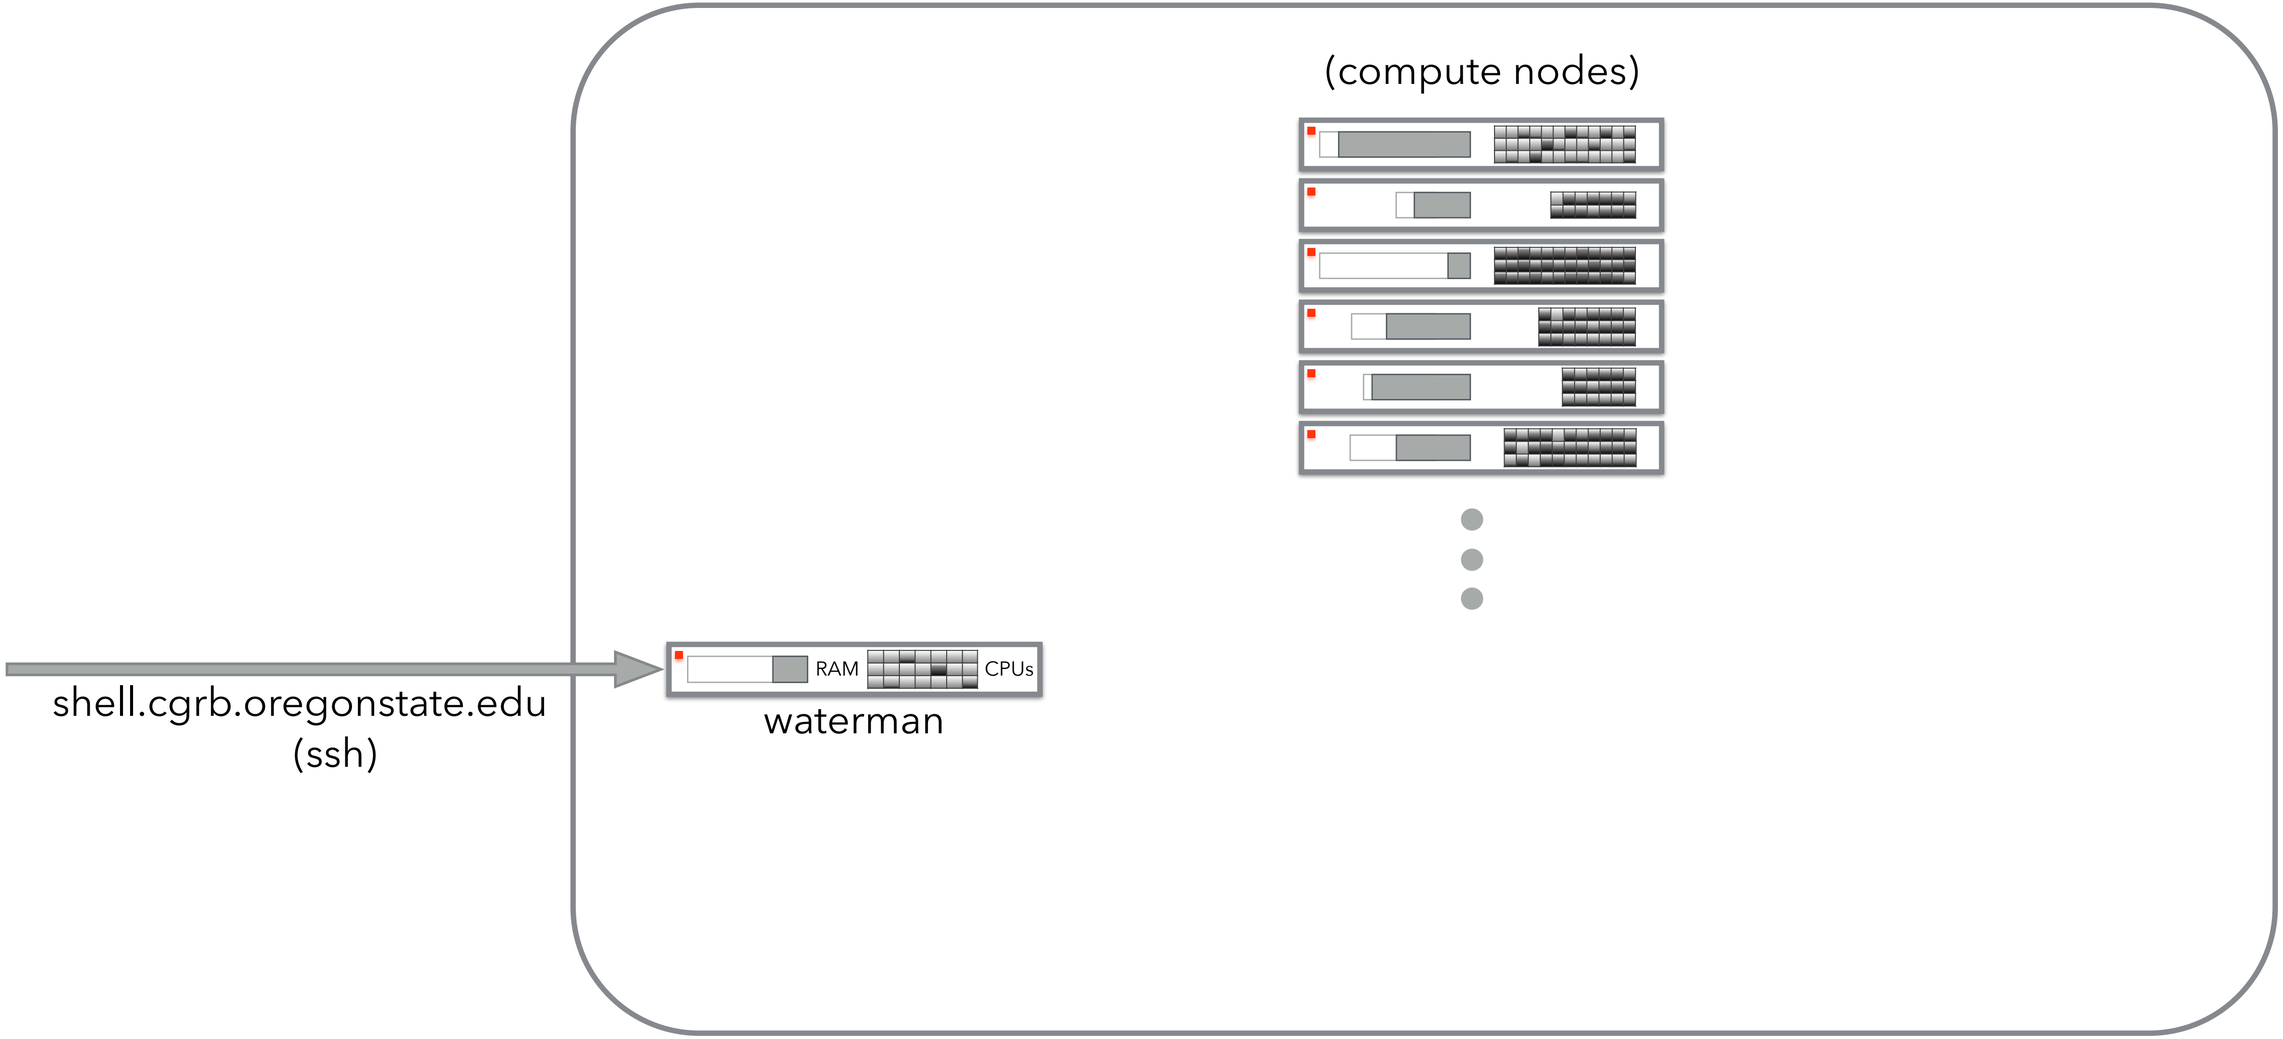

CGRB Computational Infrastructure

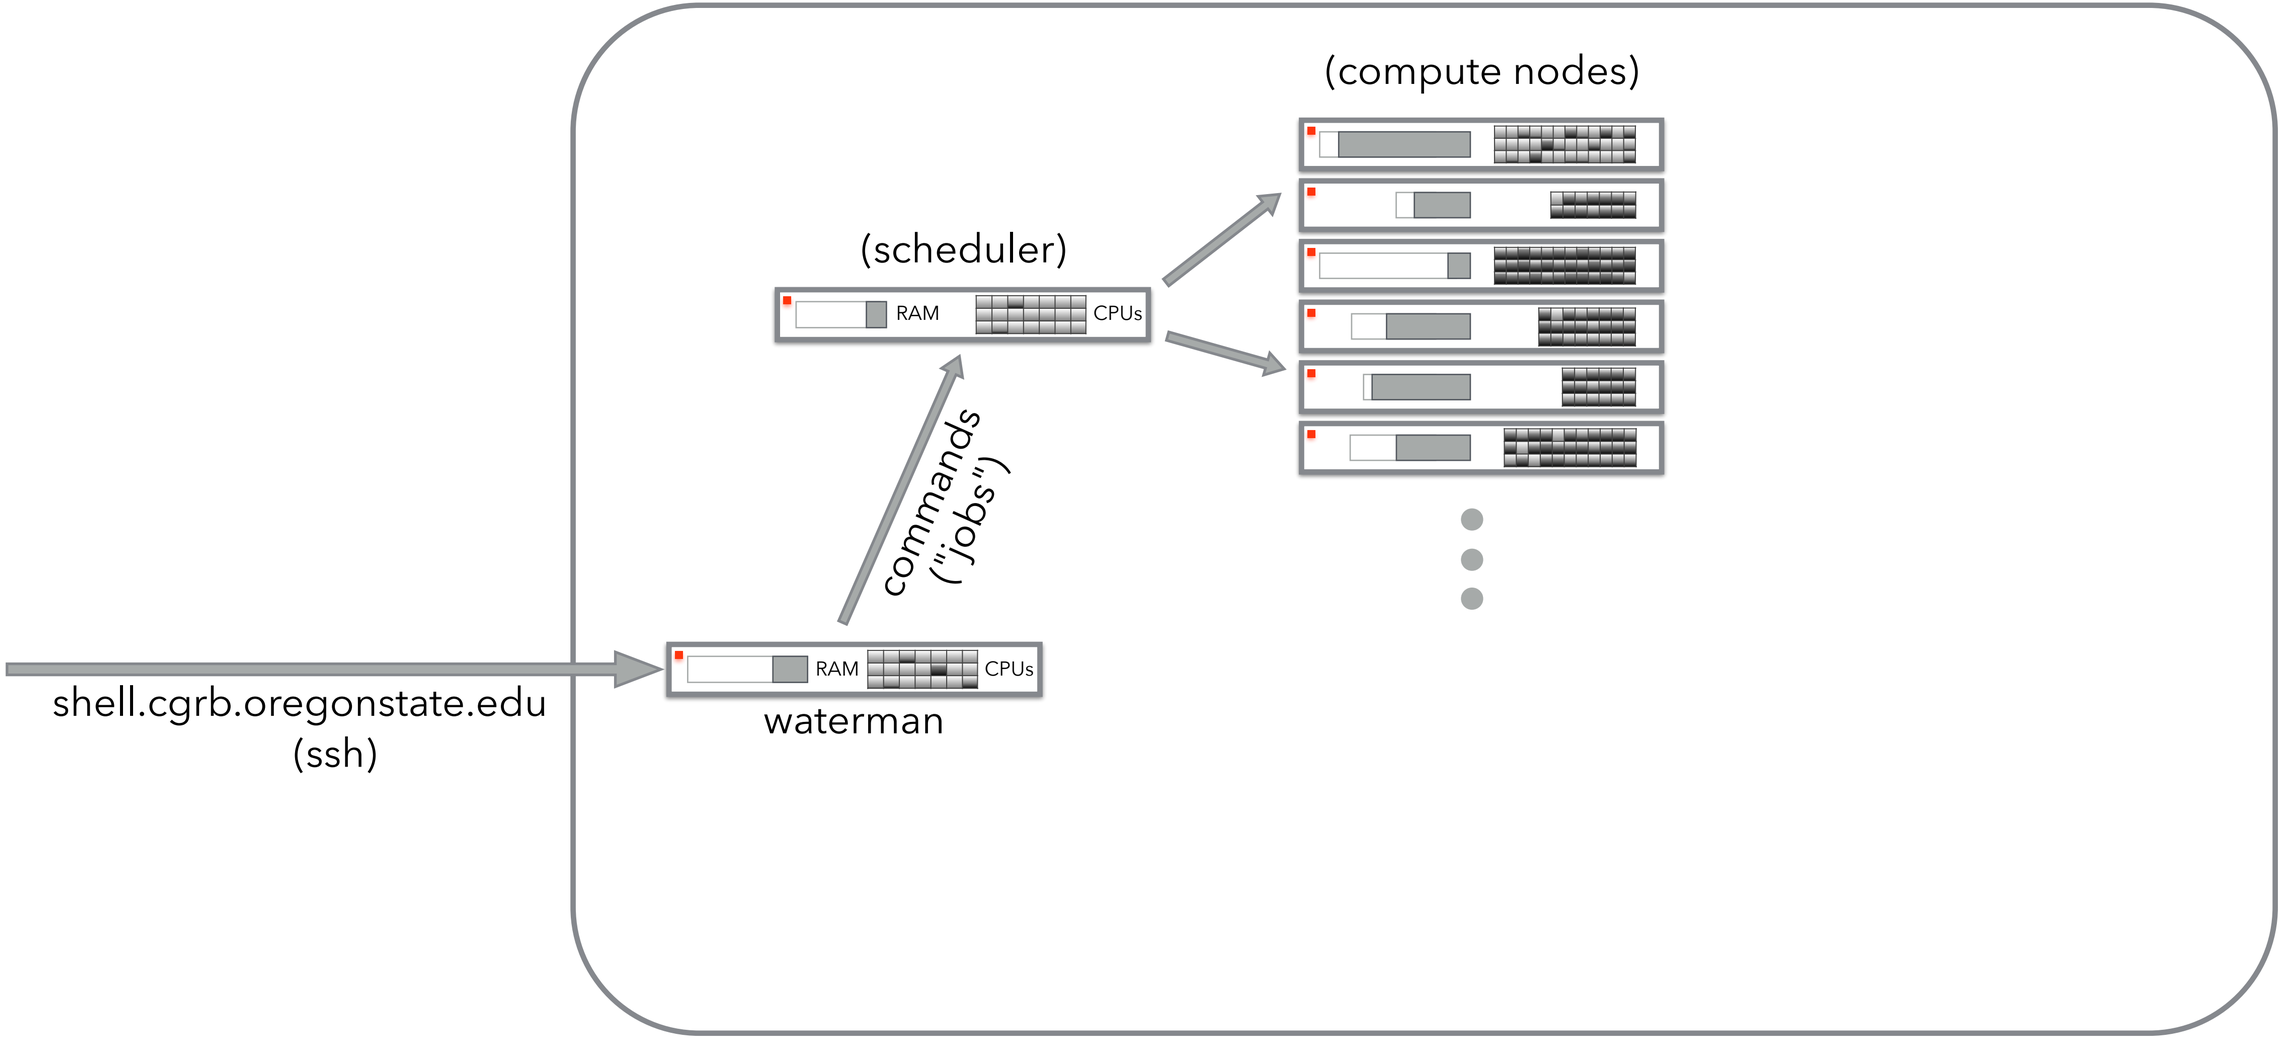

CGRB Computational Infrastructure

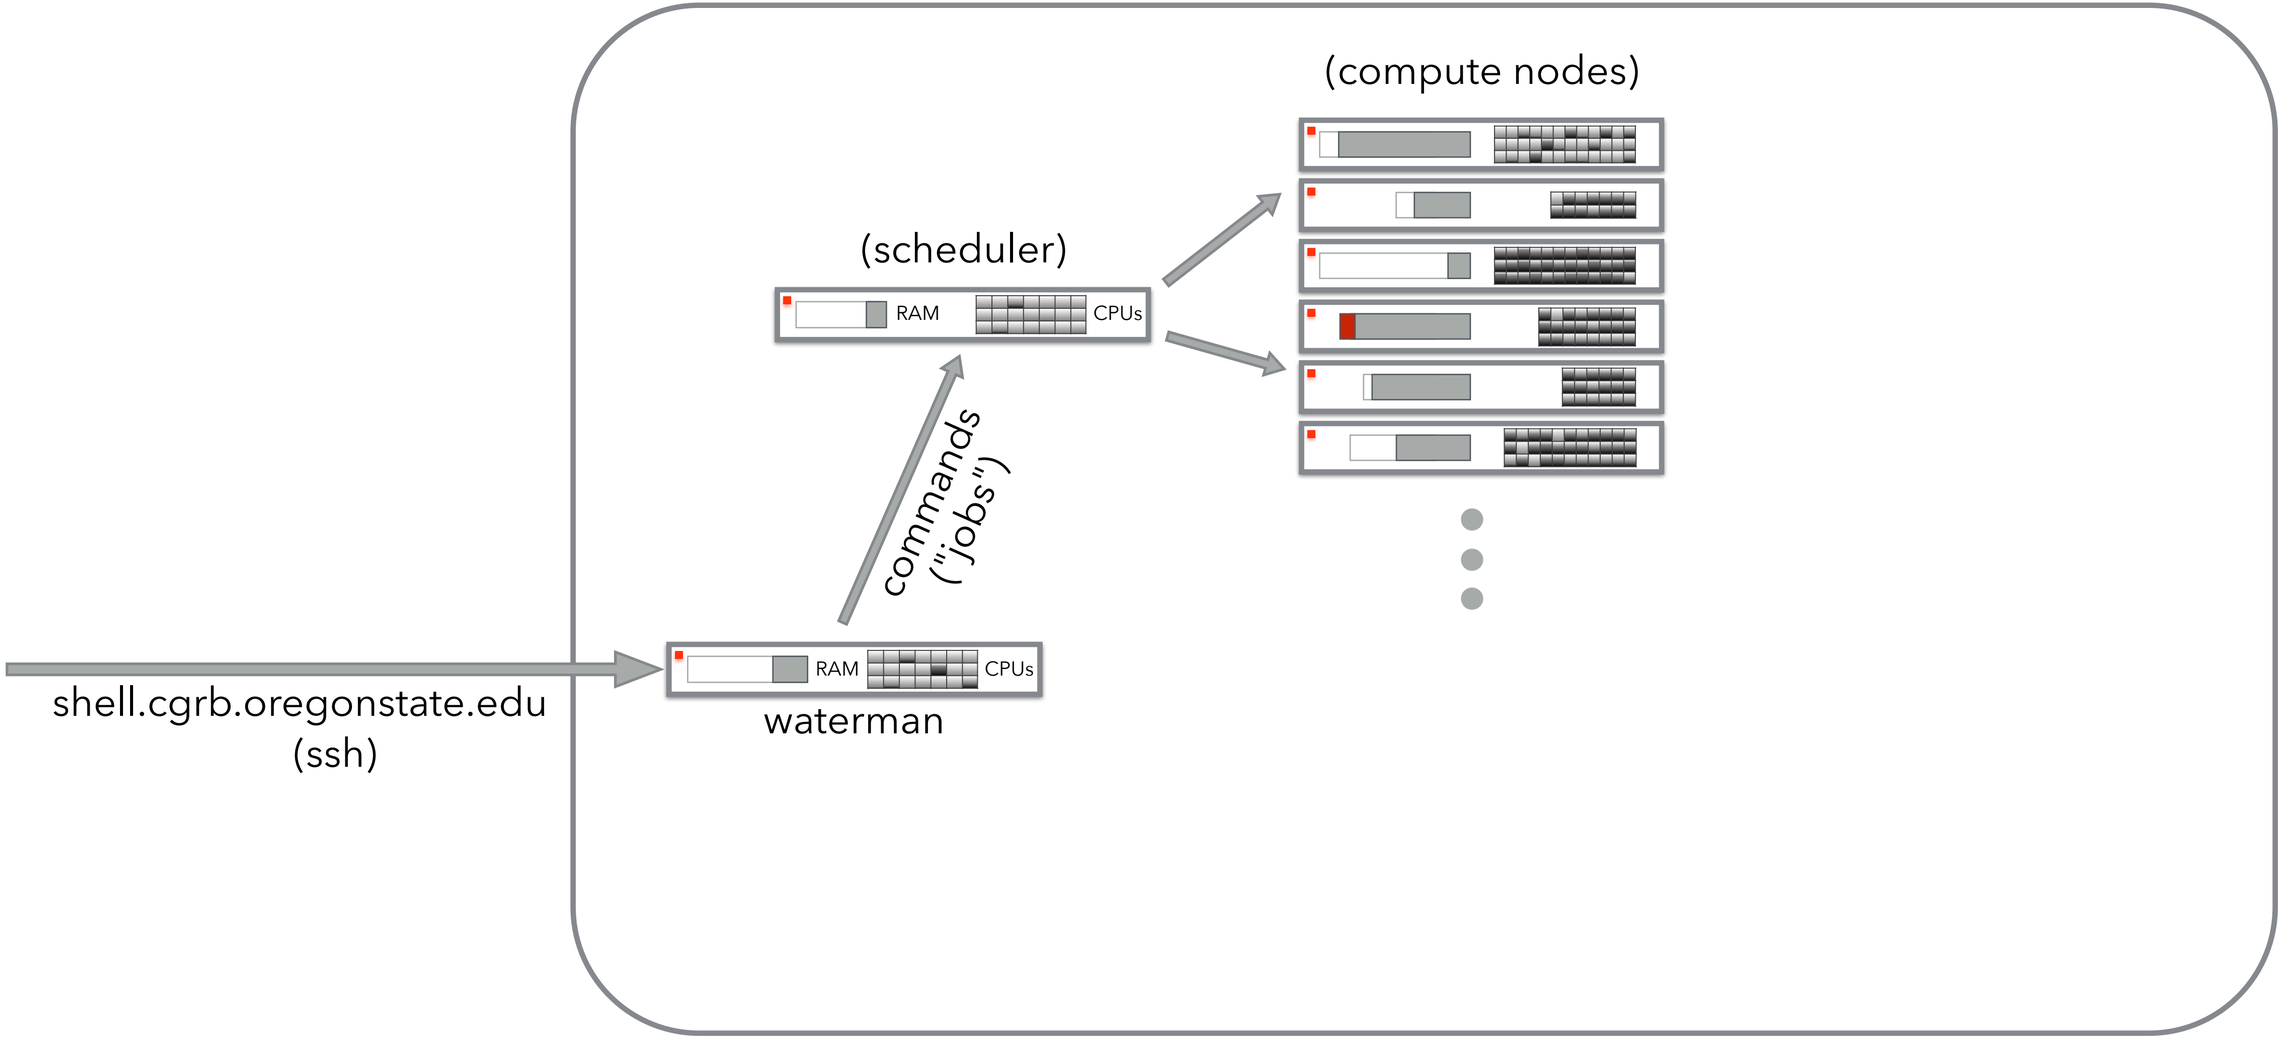

CGRB Computational Infrastructure

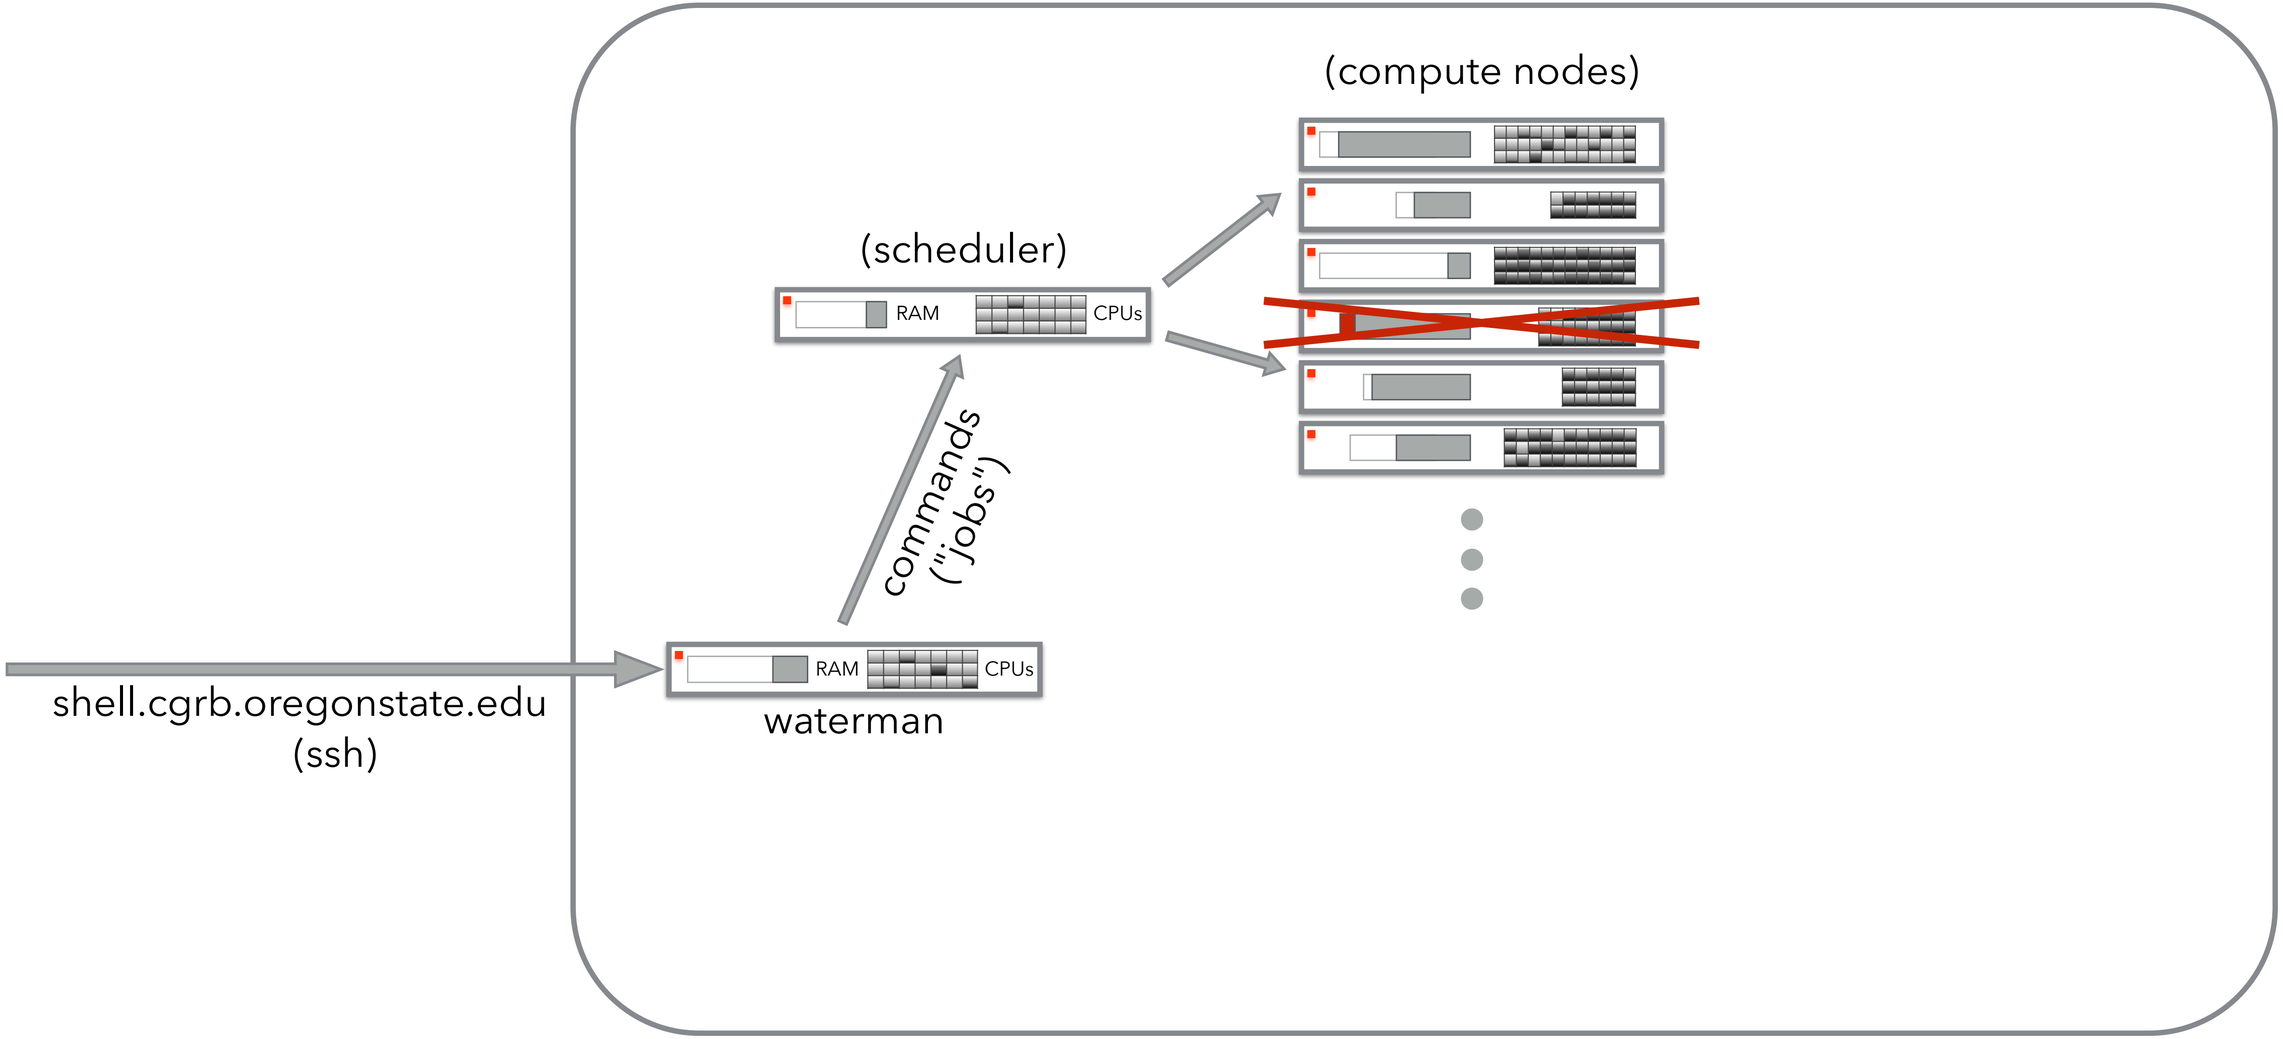

CGRB Computational Infrastructure

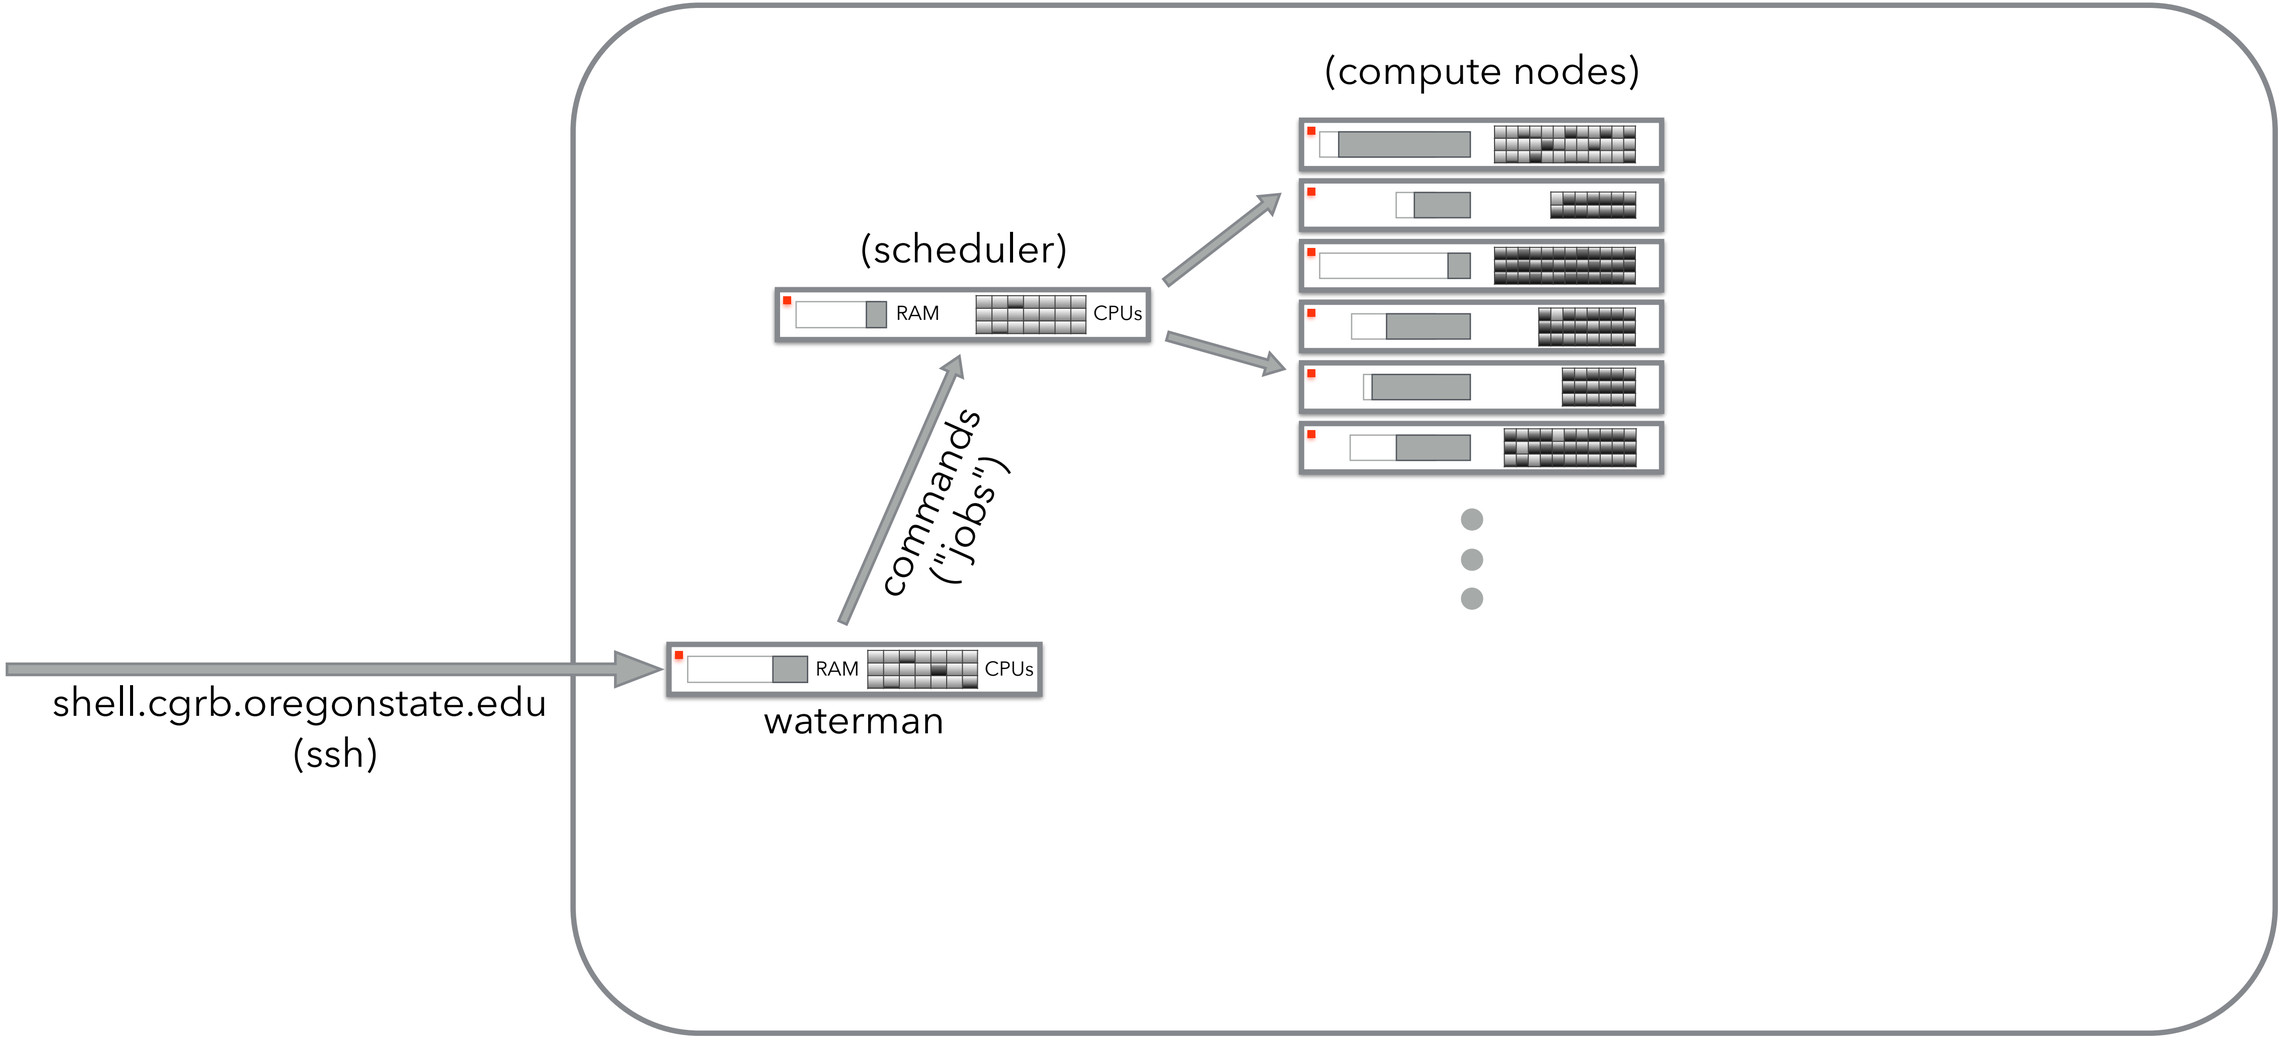

CGRB Computational Infrastructure

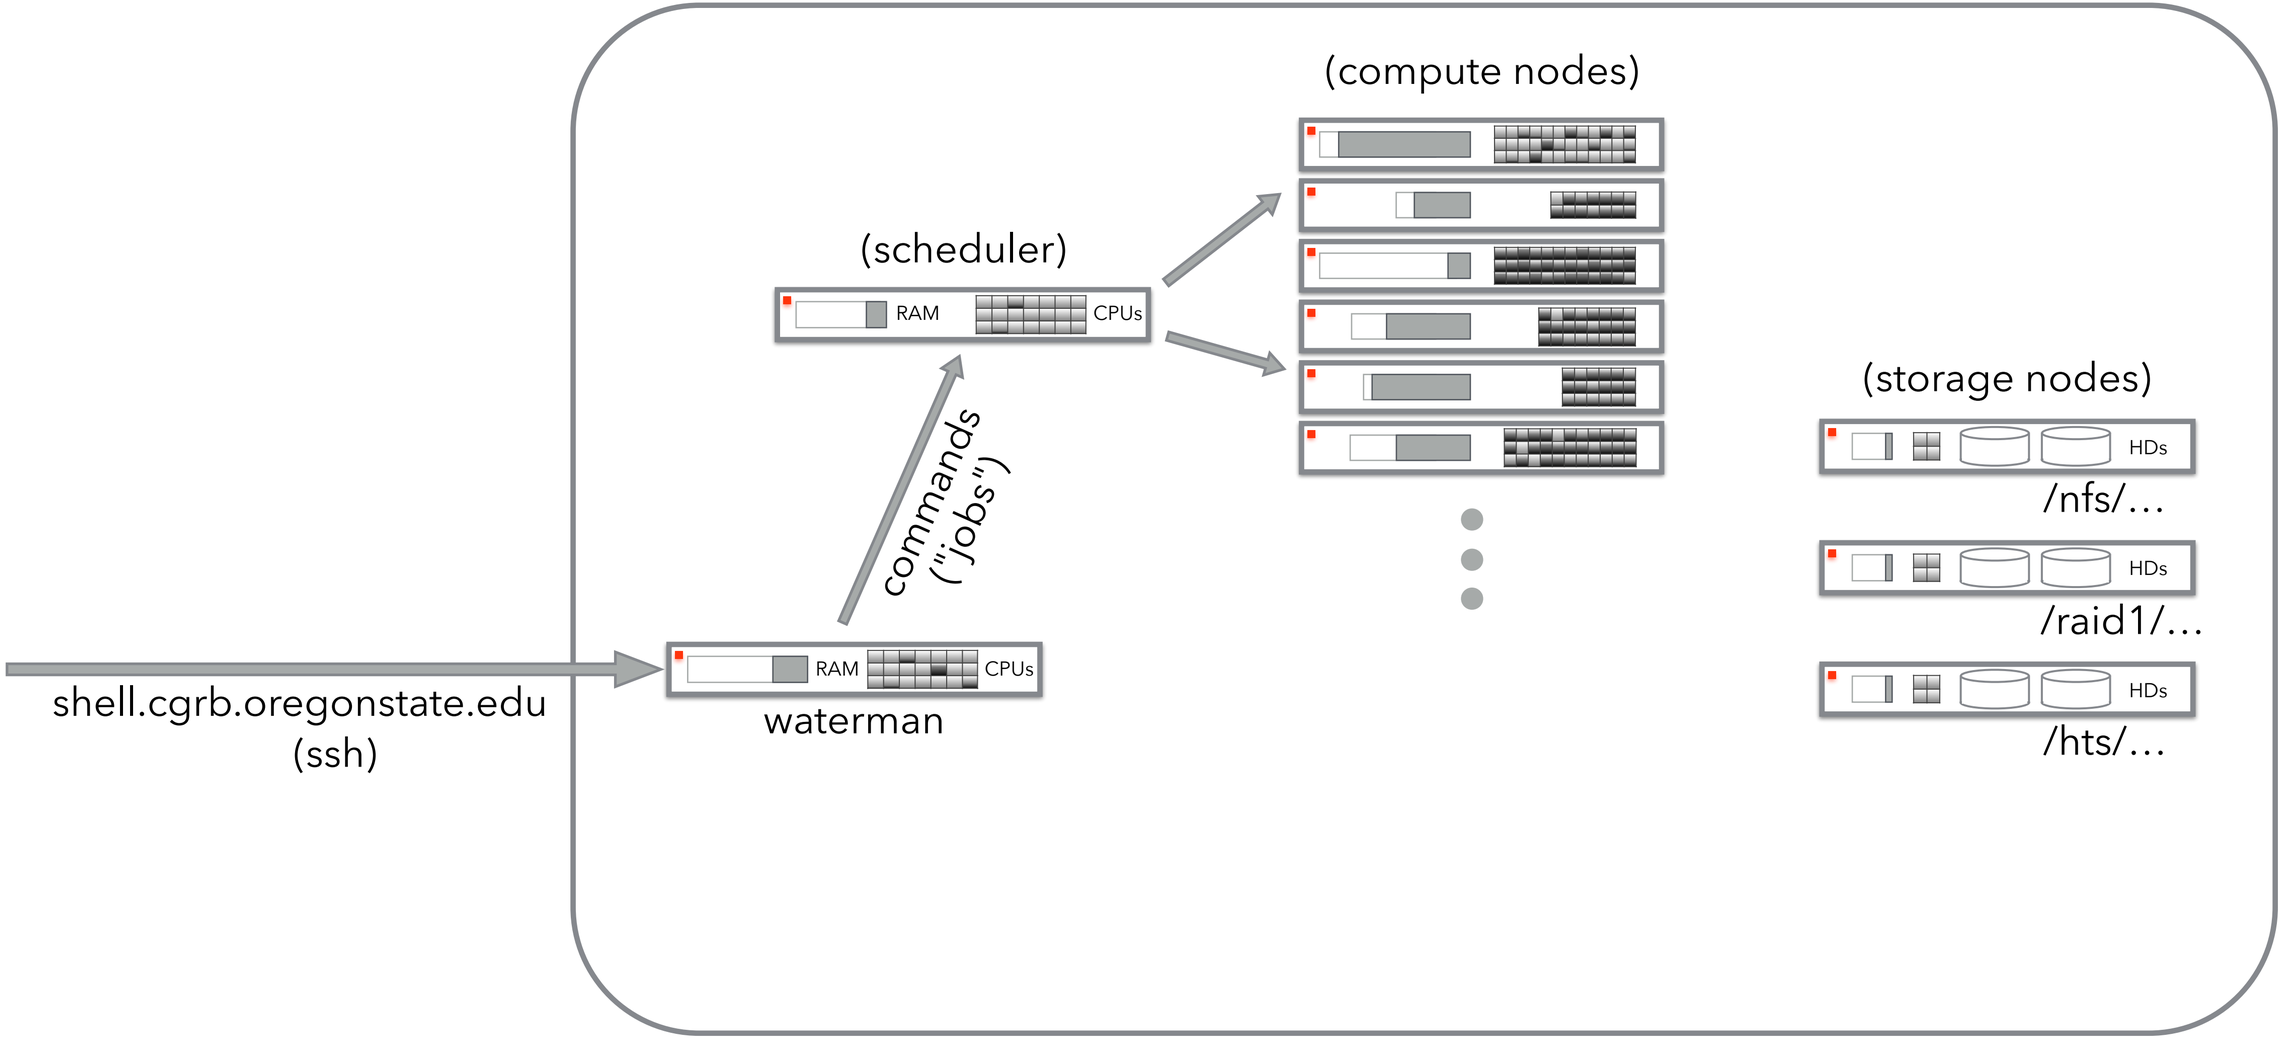

CGRB Computational Infrastructure

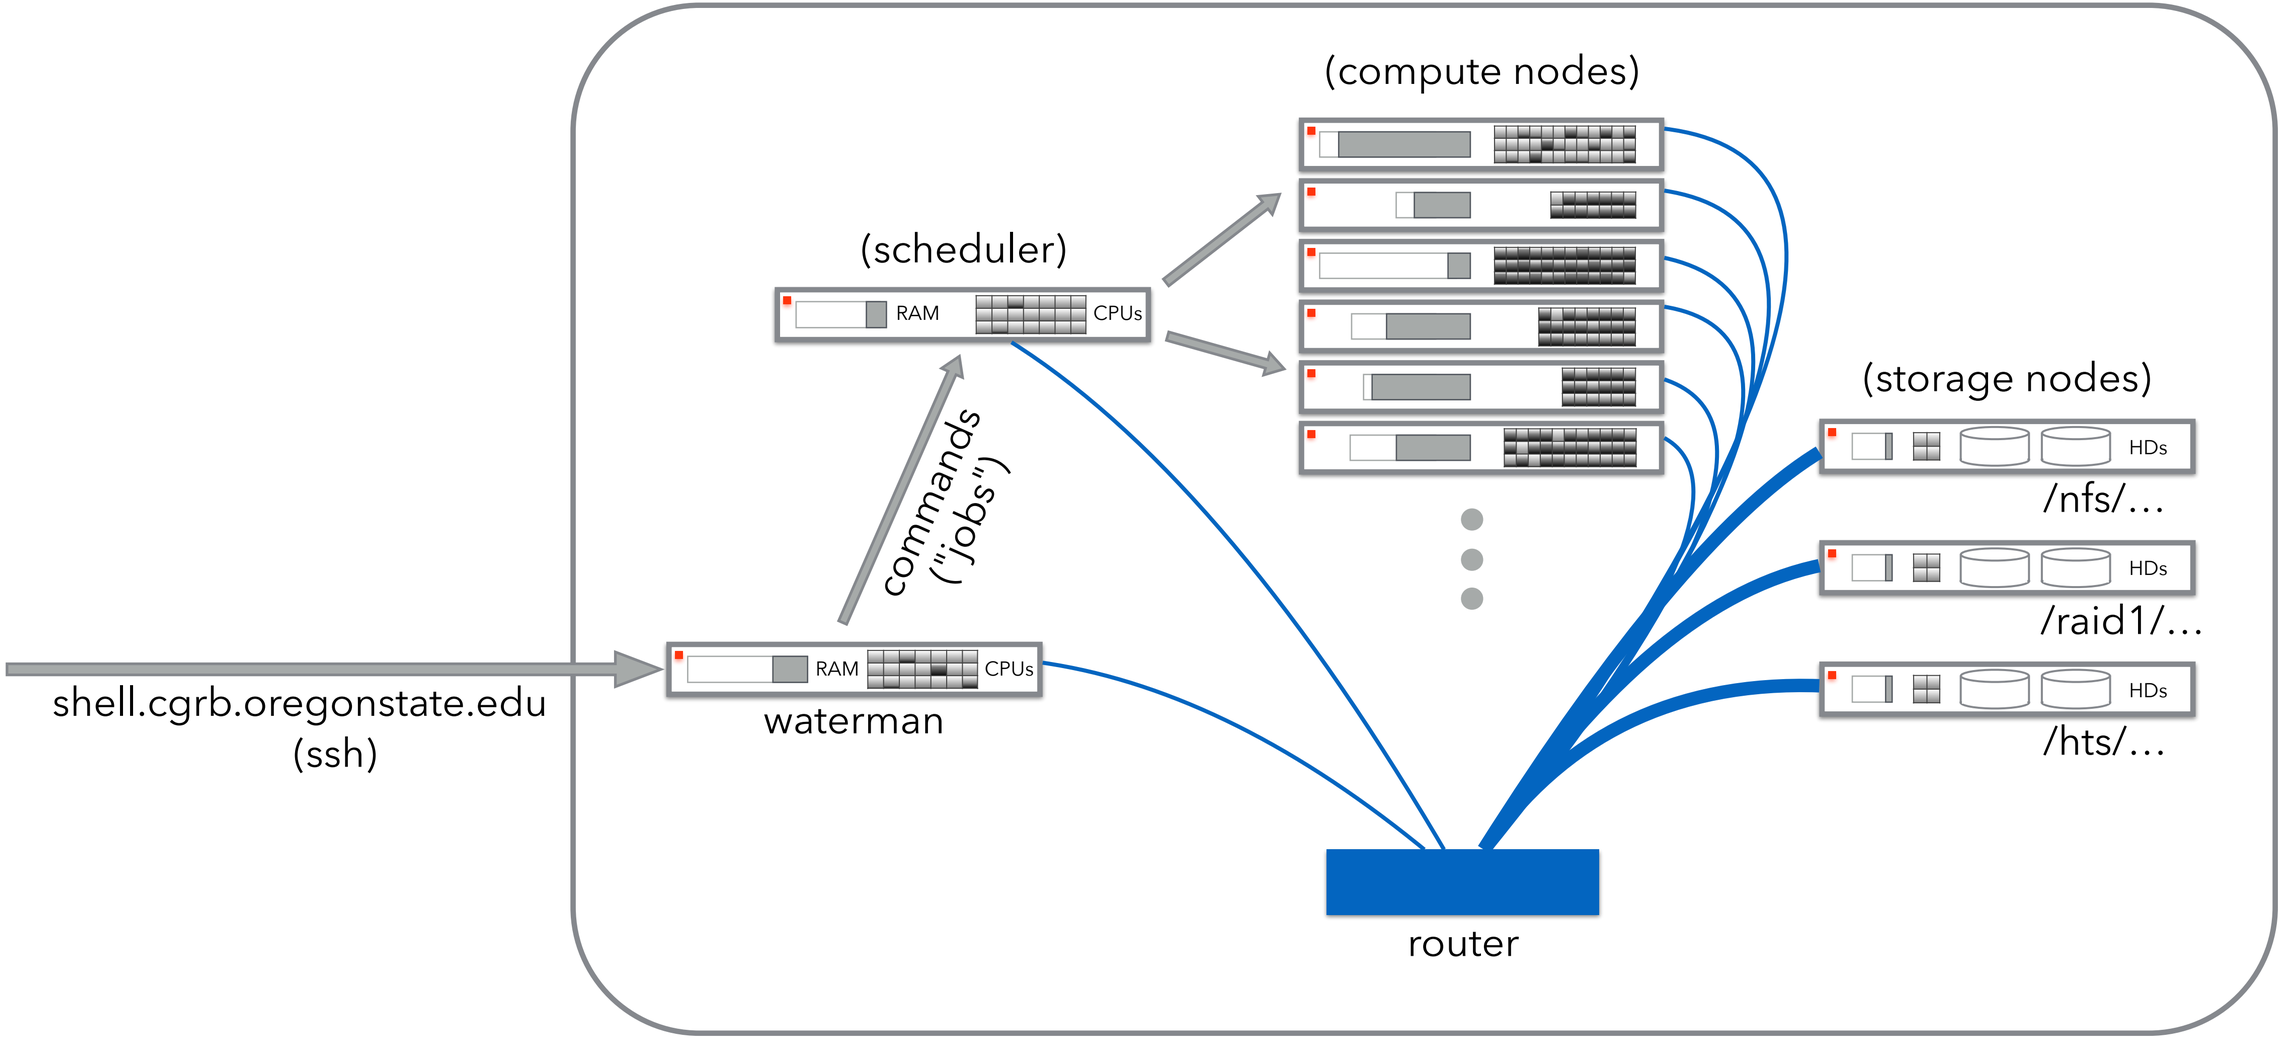

CGRB Computational Infrastructure

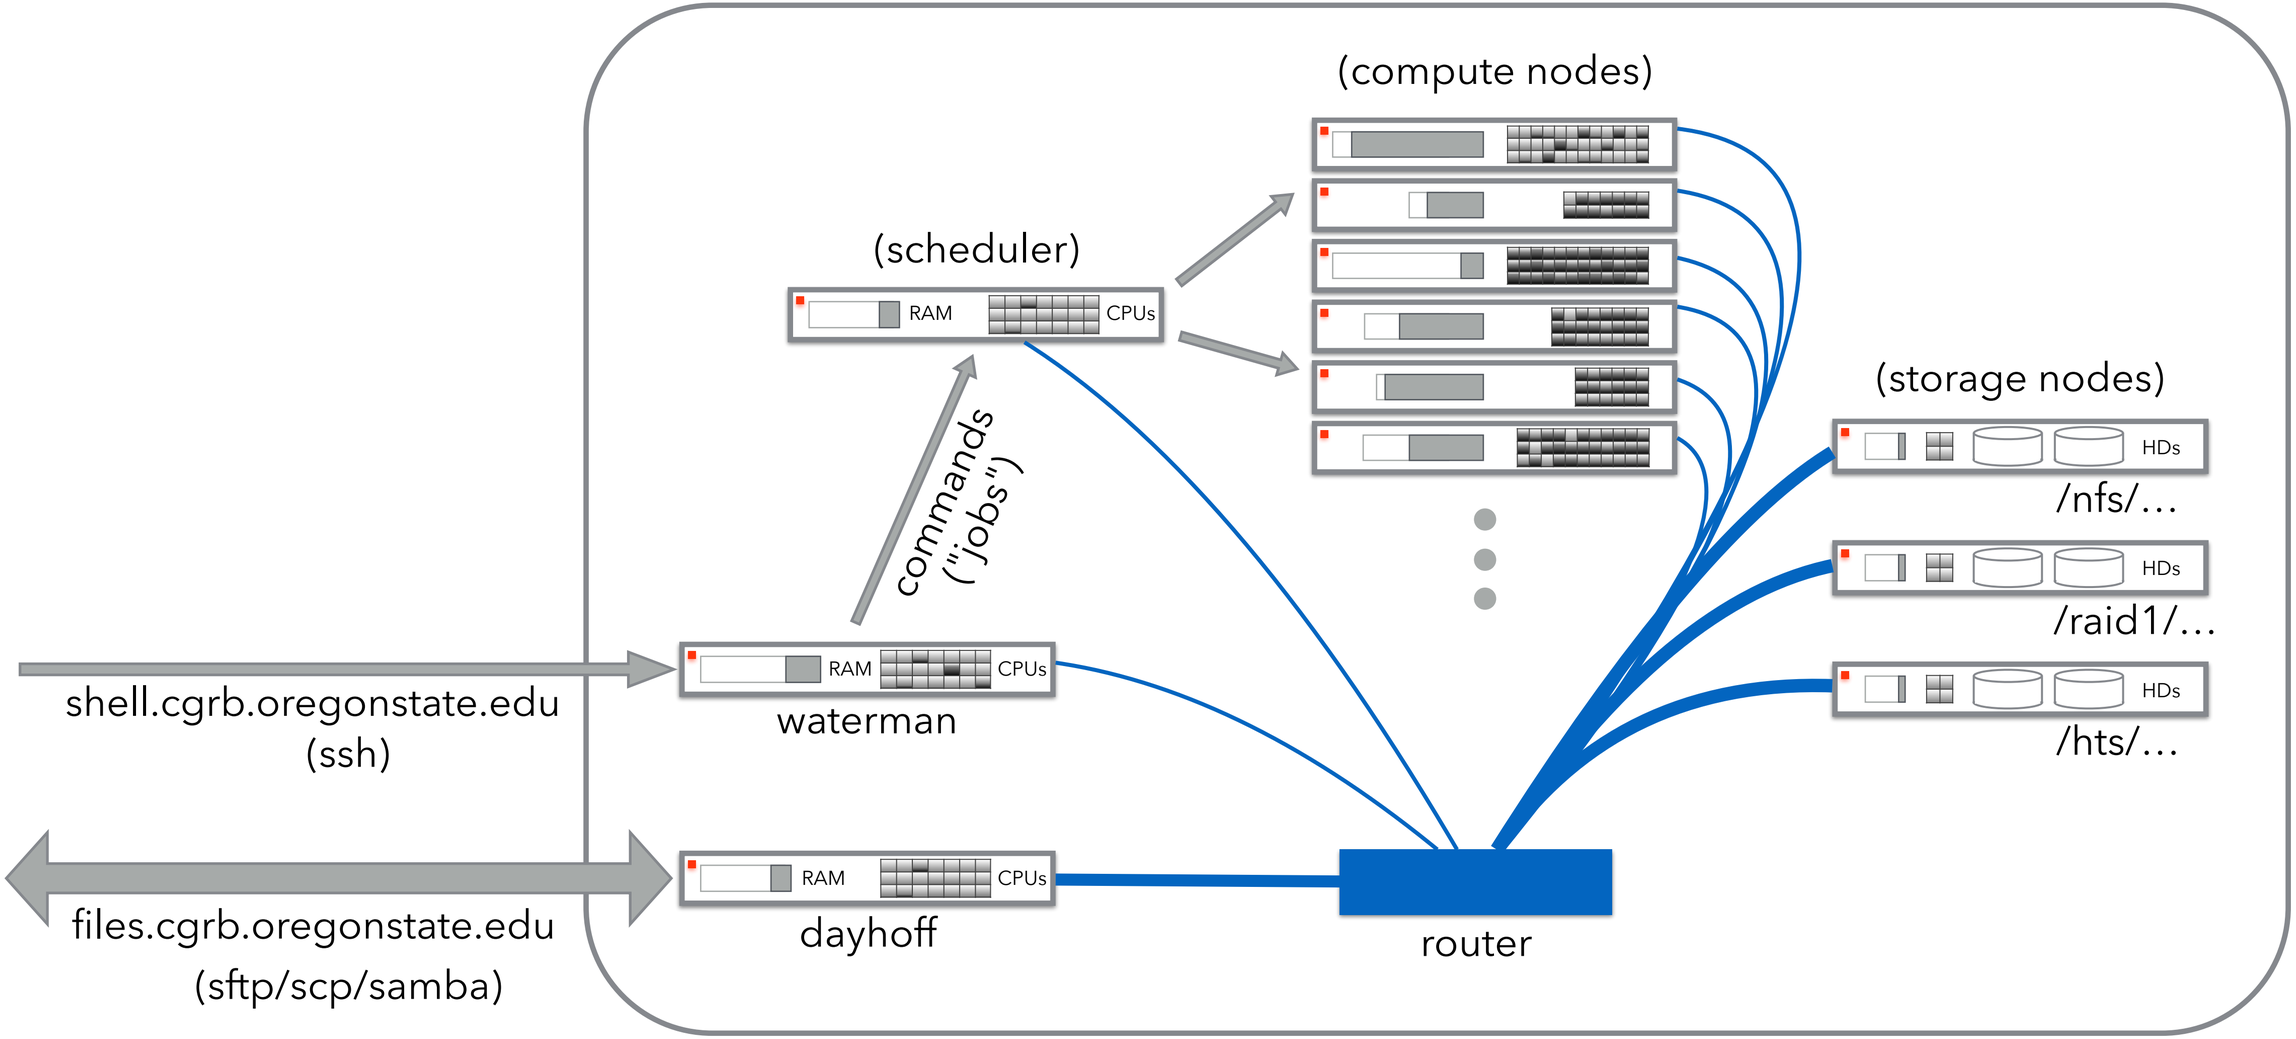

Kinda secret

```
[soneil@mbp ~]$ ssh -p XXX oneils@shell.cgrb.oregonstate.edu
[Linux@waterman ~]$ cd bug
/raid1/home/cgrb/oneils/bug
[Linux@waterman ~/bug]$ ls
[Linux@waterman ~/bug]$ date ; sleep 10 ; date > result.txt
Wed Jan 17 10:07:24 PST 2018
[Linux@waterman ~/bug]$ ls
result.txt
```

Some bioinformatics command that creates an output file

Create this directory for logs and printed output

My command (in quotes)

```
[Linux@waterman ~/bug]$ SGE_Batch -c 'date ; sleep 10 ; date > result.txt' -r log_dir -P 4 -m 10G

    * Beginning the Data run
        RunID = log_dir
        Dir = log_dir

    * Your job 2863594 ("log_dir") has been submitted

[Linux@waterman ~/bug]$ qstat
job-ID  prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
2863594 0.00000 log_dir    oneils       qw    01/17/2018 10:25:45                                  4
[Linux@waterman ~/bug]$ qstat
job-ID  prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
2863594 0.50976 log_dir    oneils       r     01/17/2018 10:25:56 roots@roots1.cgrb.oregonstate.     4
[Linux@waterman ~/bug]$ qstat
[Linux@waterman ~/bug]$ ls
log_dir/  result.txt
```

I'll need 4 CPUs and 10G RAM

(only start job if this much is free when it starts; kill it if it tries to use more than this)

How much time + RAM did that job take?

```
[Linux@waterman ~/bug]$ SGE_Plotdir log_dir
#RAM(Gigs)  Time(Seconds)
0.000166893 0.00
[Linux@waterman ~/bug]$ cd log_dir/
/raid1/home/cgrb/oneils/bug/log_dir
[Linux@waterman log_dir]$ ls
log_dir_1_sge.sh  log_dir.e2863594  log_dir.o2863594  log_dir.pe2863594  log_dir.po2863594
[Linux@waterman log_dir]$ cat log_dir.e2863594
    Full Command:                   date
    Memory (kb):                    700
    # SWAP  (freq):                 0
    # Waits (freq):                 3
    CPU (percent):                  50%
    Time (seconds):                 0.00
    Time (hh:mm:ss.ms):             0:00.00
    System CPU Time (seconds):      0.00
    User   CPU Time (seconds):      0.00
[Linux@waterman log_dir]$ cat log_dir.o2863594
  Started on:             roots1
  Started at:             Wed Jan 17 10:25:56 PST 2018
Wed Jan 17 10:25:56 PST 2018
  Finished at:             Wed Jan 17 10:26:06 PST 2018
```

Inspect the “std error” log of the command (none to show here)

Inspect the “std output” log of the command (result of “date” printed)

What do I have access to, right now?

```
[Linux@waterman ~/bug]$ SGE_Avail
#HOST   TOTRAM FREERAM    TOTSLOTS          Q  QSLOTS  QFREESLOTS   QSTATUS     QTYPE
megraw3   503.7   503.7          64     megraw      64          64        au       BIP
megraw2   503.7   503.7          64        bpp      64          64        au       BIP
roots1    503.7   484.9          64      roots      64          56    normal       BIP
plant0     86.4    86.4           8      plant       8           8        au       BIP
otter     503.7   493.8          64      otter      64           7    normal       BIP
```

Machine roots1 has 503.7G of RAM, currently 484.9G available.

It has 64 total CPUs.

64 of those are allocated to the “roots” queue

56 of *those* are available

Machine + queue are in “normal” state (au or E: not good)

B: SGE_Batch-able
I: qrsh (login)-able
P: multi-CPU (-P option)-able

Resources available to burnsg at this moment

```
[Linux@waterman ~/bug]$ SGE_Avail -u burnsg
                #HOST  TOTRAM FREERAM    TOTSLOTS           Q  QSLOTS  QFREESLOTS   QSTATUS     QTYPE
              beagle   913.2   900.9          64      beagle      64          56    normal       BIP
               reef   125.7    62.8          48        reef      12          12    normal       BIP
             chrom22     3.7     3.7           4      all.q       4           4         u        BP
             chrom20     3.7     3.7           4      all.q       4           4         u        BP
             chrom18     3.7    3.49           4      all.q       4           3    normal        BP
               epi1    31.3    23.2           8      all.q       6           2    normal        BP
```

Practice…

epi1 has 6 CPUs in the all.q, 2 are available;
notice no ‘I’ - *not* qrsh (login)-able

all.q: for everyone!
other queues: for specific lab, department members

Log me in interactively, somewhere

```
[Linux@waterman log_dir]$ qrsh
[Linux@roots1 ~]$ exit
logout
[Linux@waterman log_dir]$ qrsh -pe thread 8
[Linux@roots1 ~]$ SGE_Avail
               #HOST  TOTRAM FREERAM    TOTSLOTS             Q  QSLOTS  QFREESLOTS   QSTATUS     QTYPE
             megraw3   503.7   503.7          64        megraw      64          64        au       BIP
             megraw2   503.7   503.7          64           bpp      64          64        au       BIP
              roots1   503.7   486.9          64         roots      64          48    normal       BIP
              plant0    86.4    86.4           8         plant       8           8        au       BIP
               otter   503.7   492.7          64         otter      64           7    normal       BIP
[Linux@roots1 ~]$ qstat
job-ID  prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
2885254 0.51611 QRLOGIN    oneils       r     01/17/2018 10:54:43 roots@roots1.cgrb.oregonstate.     8
[Linux@roots1 ~]$ exit
logout
[Linux@waterman log_dir]$ qrsh -pe thread 8 -q roots@roots1
[Linux@roots1 ~]$
```

Same, but allocate me 8 CPUs

Same, but I specifically want the roots queue, on the roots1 machine

(-q roots would get me the roots queue on any machine that has it)

If using qrsh, do not forget to log back out!

You're telling the scheduler you are using those resources, so others can't.

I want to run a *lot* of jobs….

Per-job requirements

Similar to SGE_Batch

Only run 10 at a time

```
[Linux@waterman log_dir]$ head -n 5 lotsa_commands.txt
sleep 2; echo job1 > result1.txt
sleep 2; echo job2 > result2.txt
sleep 5; echo job3 > result3.txt
sleep 1; echo job4 > result4.txt
sleep 3; echo job5 > result5.txt
[Linux@waterman ~/bug]$ SGE_Array -c lotsa_commands.txt -r log_dir -P 4 -m 10G -b 10
[Linux@waterman ~/bug]$ qstat
job-ID  prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 1
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 2
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 3
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 4
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 5
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 6
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 7
2893513 0.50976 log_dir    oneils       r     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 8
2893513 0.50976 log_dir    oneils       t     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 9
2893513 0.50976 log_dir    oneils       t     01/17/2018 11:08:12 roots@roots1.cgrb.oregonstate.     4 10
2893513 0.00000 log_dir    oneils       qw    01/17/2018 11:08:07                                    4 11-50:1
```

10 Are running (or starting); the other 40 are waiting!

```
[Linux@waterman ~/bug]$ SGE_Plotdir log_dir
#ArrayJob#  GigsRAM      Seconds
1           0.000350952  2.00
2           0.000350952  2.00
3           0.000350952  5.00
4           0.000349998  1.00
5           0.000350952  3.00
...
```

How much time/RAM did those jobs use?

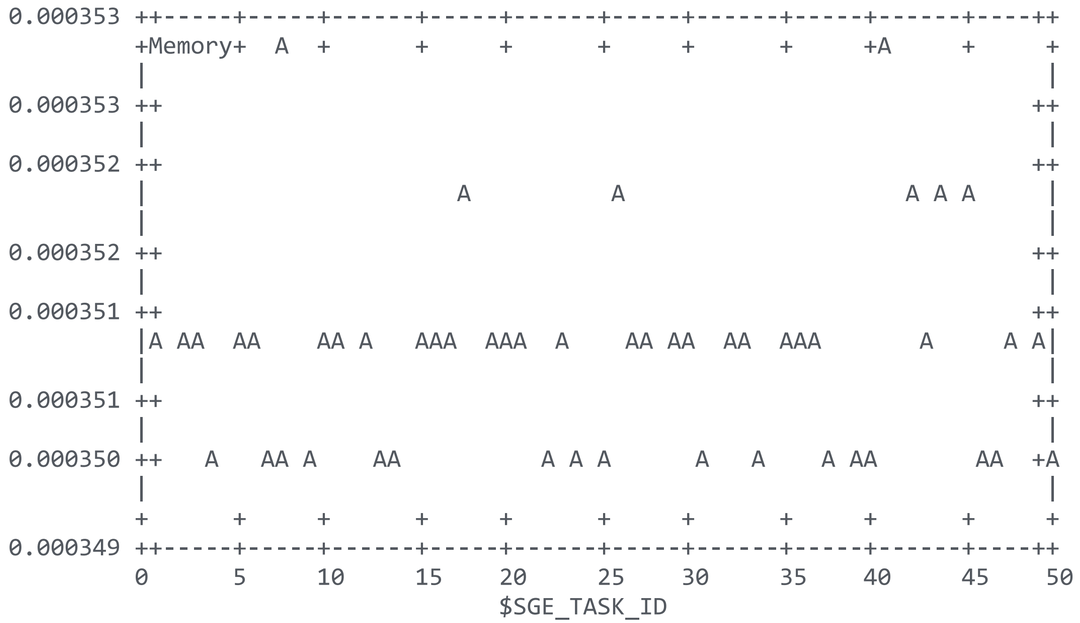

RAM plot per job (gigs)

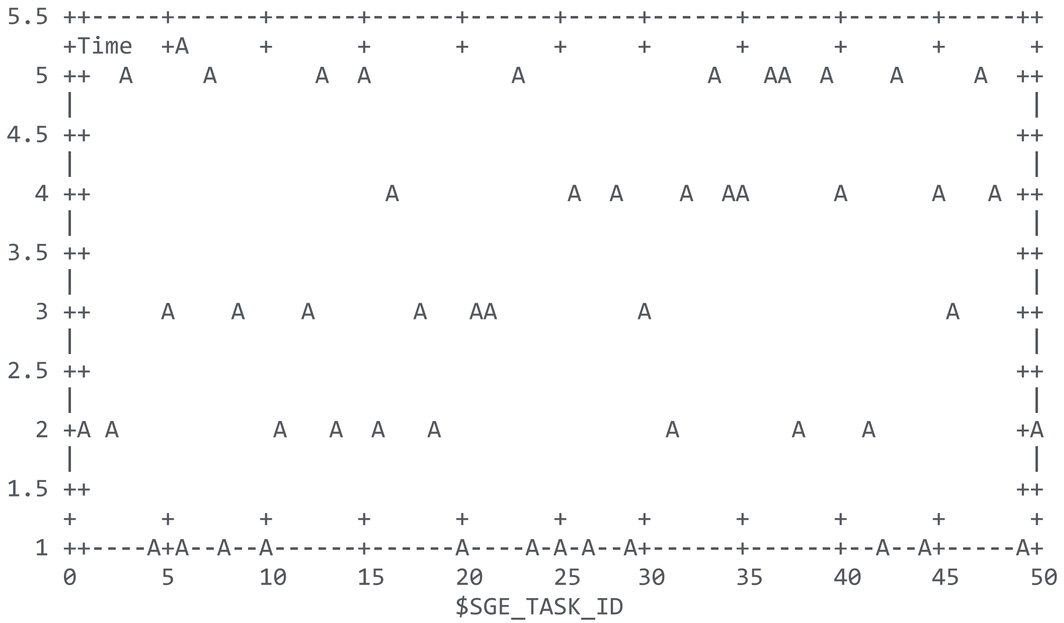

time plot per job (seconds)

Delete job # 2899950 (if it belongs to me)

```
[Linux@waterman ~/bug]$ qdel 2898850
```

Delete all my jobs (oh %*!$)

```
[Linux@waterman ~/bug]$ qdel -u oneils
```

See all jobs by all users

```
[Linux@waterman ~/bug]$ qstat -u '*'
```