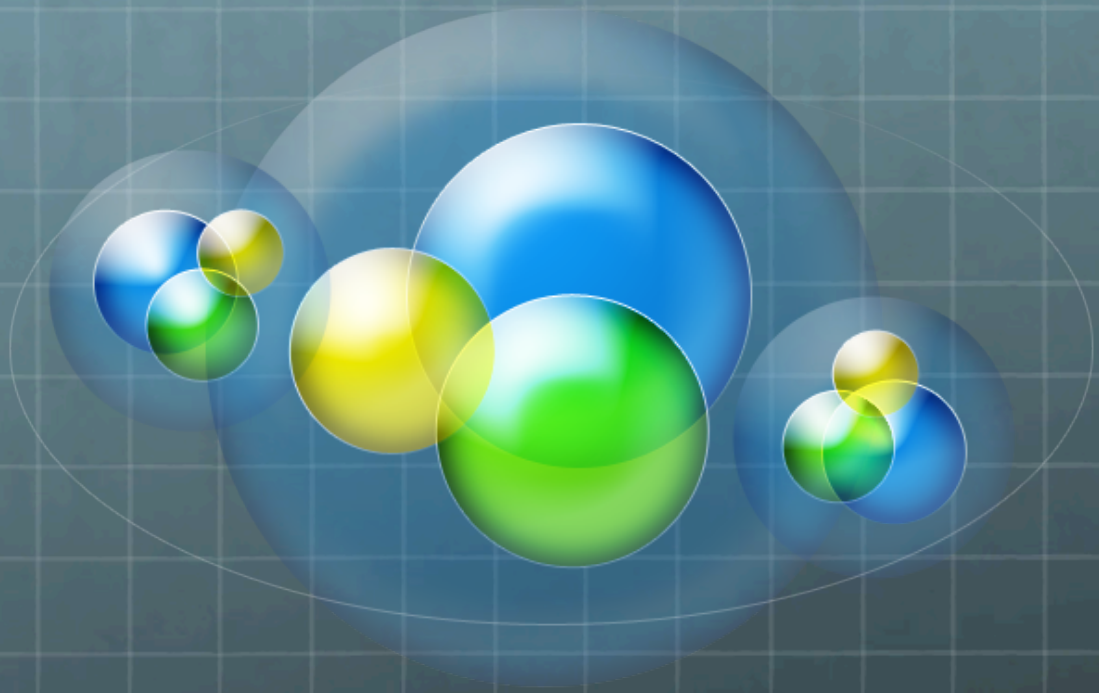

Concepts in Cluster Computing – Introduction to Biocomputing

Ken Weiss
MSIS Computational Research Consulting Division
August, 2016

Roadmap

- Flux overview
- Logging in to Flux
- PBS script creation
- Job submission
- Job monitoring
- Advanced topics

Flux

Flux

- Flux is a university-wide shared computational discovery / high-performance computing service.
- Interdisciplinary
  - Provided by Advanced Research Computing at U-M (ARC)
  - Operated by ARC Technical Services (ARC-TS)
  - Hardware procurement, software licensing, billing support by U-M ITS
  - Used across campus
- Collaborative since 2010
  - Advanced Research Computing at U-M (ARC)
  - College of Engineering's IT Group (CAEN)
  - Information and Technology Services (ITS)
  - Medical School
  - College of Literature, Science, and the Arts
  - School of Information

http://arc-ts.umich.edu/systems-services/

A Flux node

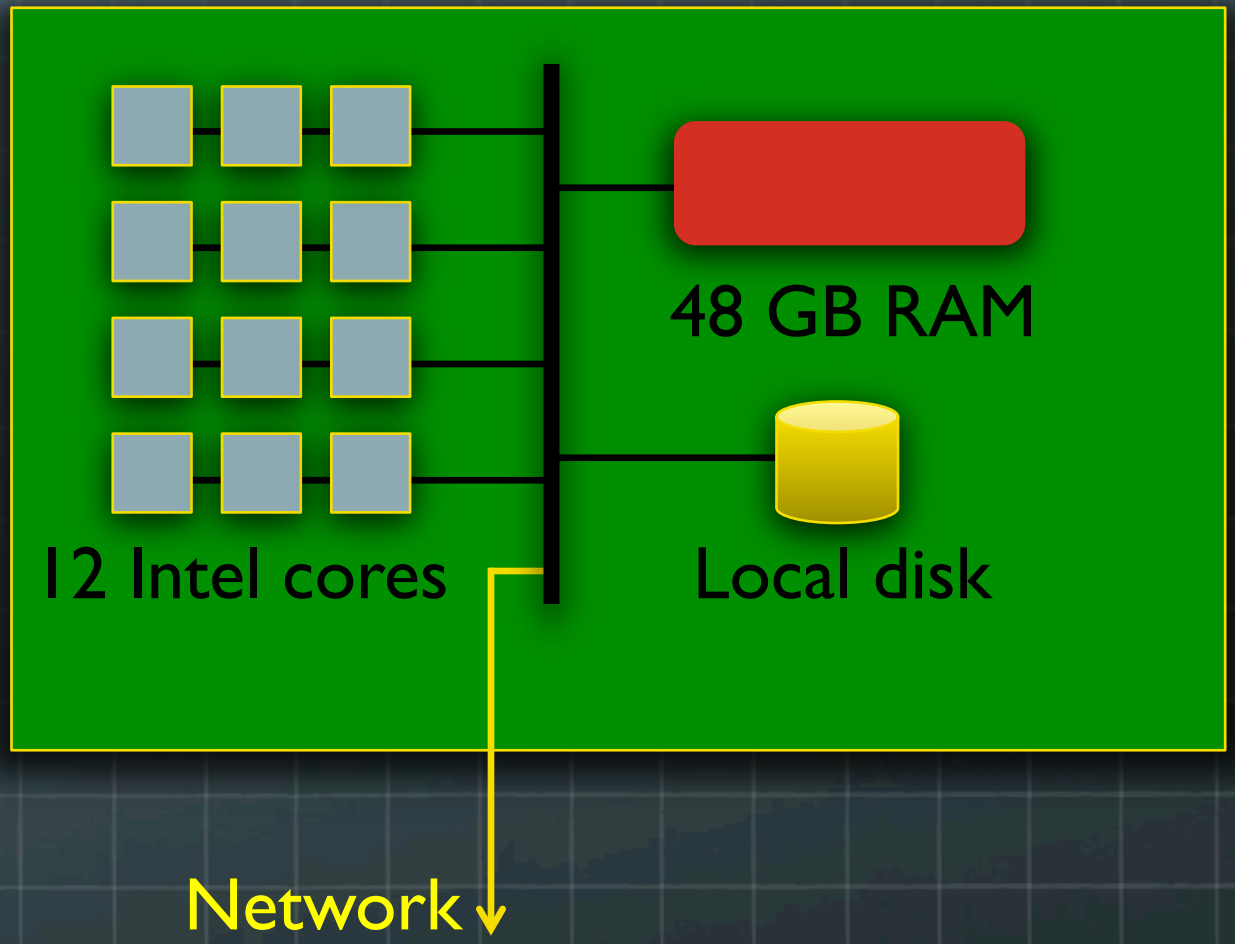

The Flux cluster

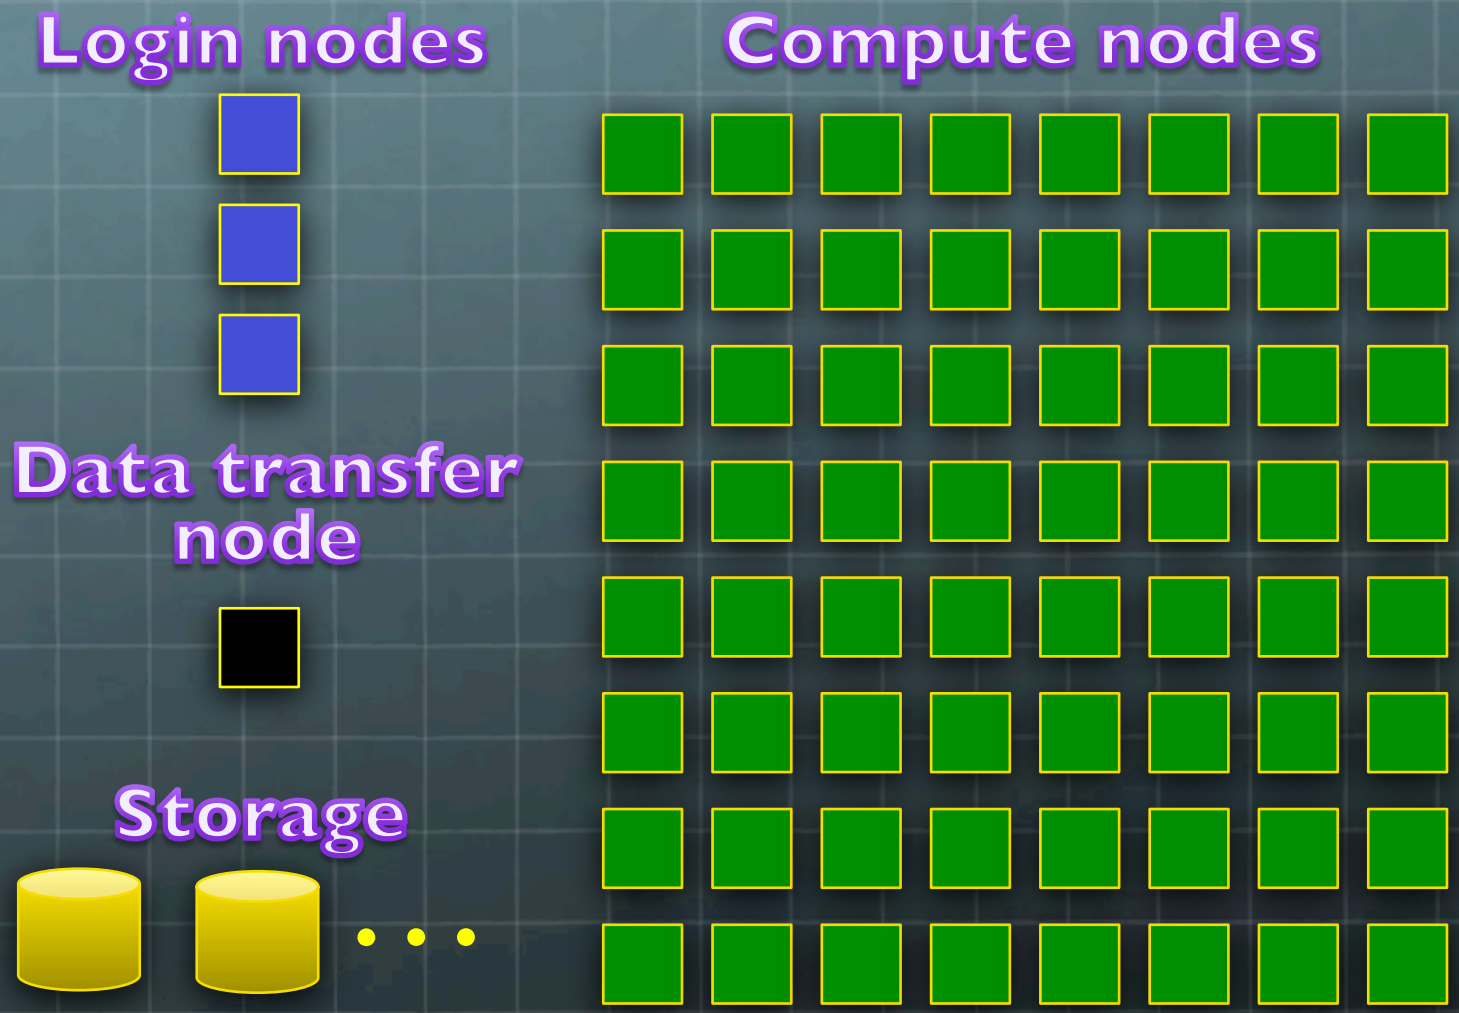

Logging In To Flux

Connecting via ssh

- Terminal emulators
  - Linux and Mac OS X
    - Start Terminal
    - Use `ssh` command
  - Windows
    - MobaXterm
      http://mobaxterm.mobatek.net/
    - U-M PuTTY/WinSCP (U-M Blue Disc)
      https://www.itcs.umich.edu/bluedisc/
    - PuTTY
      http://www.chiark.greenend.org.uk/~sgtatham/putty/

Logging in to Flux

Flux user account, Duo authentication, and Level 1 password are required

- For Mac or other Linux workstation, from a terminal window type: ssh *uniqname*@flux-login.arc-ts.umich.edu
- On a PC, start MobaXterm. If you have not attached to Flux and saved your session, click on the "Session" button. In the Session settings box, select SSH.

  Remote Host:
  flux-login.arc-ts.umich.edu

  In the Advanced SSH settings
  Check:
  X11-Forwarding, Compression

  Uncheck:
  Display SFTP browser

  Remote environment:
  Interactive shell

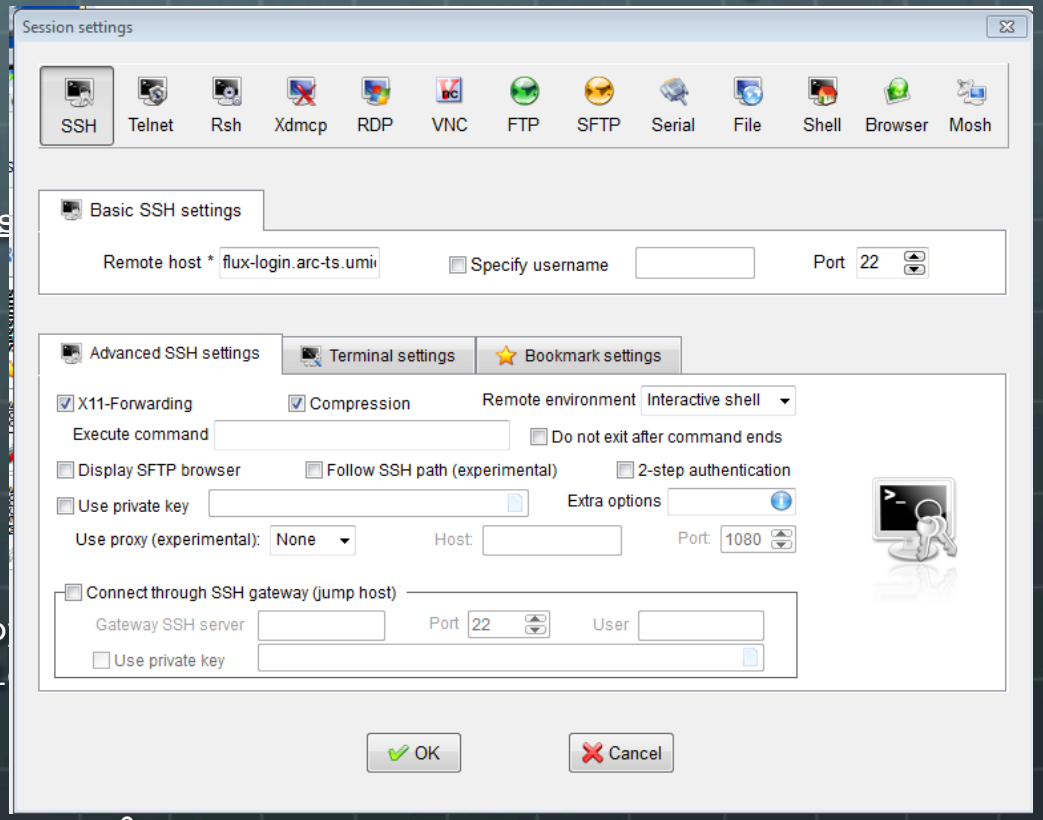

Click OK. A terminal window will o
the prompt 'Login:' enter your uni
and press enter.

Logging in to Flux

- You will be prompted, “Password:”

  Enter your Level 1 password and press enter.

- You will be prompted: “Enter a passcode or select one of the following options:”

  Select your Duo method of two factor authentication and press enter.

- You will be randomly connected a Flux login node
  - Currently `flux-login1`, `flux-login2` or `flux-login3`

A shell on a login node

- Now you have a shell on the login node
- Your shell prompt has changed to something like `[uniqname@flux-login1 ~]$`
- Try these simple commands here!
  - ~$ date ~$ id ~$ ps ~$ df -kh
  - ~$ who ~$ top # type Control-c or q to exit
- Flux is for HPC work
  - The login nodes are reserved for copying data, editing and compiling programs and scripts, and short test runs
  - For HPC computation, you should run on the compute nodes
    - A batch job will always run on the compute nodes
    - An *interactive* batch job will get you a shell on a compute node

Why command line?

1. Linux was designed for the command line
2. You can create new Linux commands using the command line, without programming
3. Many systems provide only the command line, or poorly support a GUI interface
   - Such as most HPC systems
4. Many things can be accomplished only through the command line
   - Much systems administration & troubleshooting
5. You want to be cool

Access to Software

- Software modules have been installed on Flux to make it easy to use the software which you need. Use the module command to setup the software for your session.

  module avail

  module list

  module load *software_package*

  module unload *software_package*

  module swap *software_package*

- You can create your own modules in your /home directory if you wish to install software yourself.

PBS Script Creation

PBS Script Creation

- PBS = Portable Batch System
- With so many users on the system and a finite number of resources, the scheduler takes job requests and allocates resources, assigns which compute node(s) the job will run on and maintains 'fair sharing' of the cluster.
- The scheduler also maintains job logs which are useful for:
  - selecting the proper amount of resources to request
  - troubleshooting issues with a user's workflow
  - identifying hardware issues with the compute nodes

A Sample PBS Script

```
#PBS -N Rcoinflip
#PBS -A biobootcamp_fluxod
#PBS -q fluxod
#PBS -l nodes=1:ppn=1,pmem=3GB,walltime=1:00:00,qos=flux
#PBS -M kgweiss@umich.edu
#PBS -m abe
#PBS -j oe
#PBS -V
```

Scheduler Directives

```
#  Include the next three lines always
if [ "x${PBS_NODEFILE}" != "x" ] ; then
   cat $PBS_NODEFILE   # contains a list of the CPUs you were using if run with PBS
fi

# Always add this line to change to the working directory of the program. Otherwise
# the script will be run from the users /home directory
cd $PBS_O_WORKDIR
echo $PBS_O_WORKDIR
```

```
#  Put your job commands after this line
R CMD BATCH --no-save --no-restore coinflip.R
```

Your workflow goes here

/scratch

- What is /scratch?

  /scratch is a high performance file system connected to fast networking which allows for high throughput reads/writes
- Why should I use /scratch?

  In a word… performance
- What should I do with my files on /scratch

  Once your program is done running, copy your results files to your /home directory or other storage you have access to
- Your /scratch space for this class is:

  /scratch/biobootcamp_fluxod/*uniqname*
- WARNING: /scratch is NOT backed up. Make sure that files contained on /scratch are not your only copy.

Job Submission

Job Submission

- For a job to run on a cluster, you must submit the job to the scheduler.
- The command to send jobs to the scheduler is:

  qsub *PBS_script_name*

- Useful commands for working with queued jobs:

  qdel <job ID #> - removes the job from the queue (will work also if the job is running)

  qhold <job ID #> - puts the job on hold (pause)

  qrelease <job ID #> - releases a held job (resume)

  qalter <job ID #> - alters the PBS directives for a submitted job

Job Monitoring

Job Monitoring

- Often when you have very long jobs and/or very many jobs, you wish to know the status of your job(s). There are a variety of ways you can monitor your jobs.

  qstat -u <uniqname> - list job info for a user

  showq –w user=<uniqname> - details regarding the status of blocked, eligible and active jobs

  qstat –f <job ID #> - list all the attributes for a given job

  checkjob -vv <job ID #> - list job details for a given job

Job Submission Example

- First, let's get the files into your /scratch directory.

  type: cd /scratch/biobootcamp_fluxod/*uniqname*

  cp ../kgweiss/coinflip.* .

- Next, modify coinflip.pbs to send email to your uniqname.

  type: nano coinflip.pbs

  modify: #PBS -M *uniqname*@umich.edu

  save and exit

- Now, let's load the R module so that our program will run

  type: module load R

- Finally, submit the job

  type: qsub coinflip.pbs

Advanced Topics

Advanced Topics

- Job dependencies – Start a job after some condition
  qsub -W depend=afterok:<job ID #> program2.pbs
- Job arrays
  #PBS -t 1-10 (add to your PBS directives)
  ./myprogram -input=file-${PBS_ARRAYID} (command to execute)
- Interactive jobs
  qsub -I -X -V -l procs=2 -l walltime=8:00:00 -l mem=3GB
  -A *allocation_flux* -l qos=flux -q flux
- Globus - another, very useful way to copy files
  See https://www.globus.org/ for details

Useful Links

- http://arc.umich.edu/
- http://arc-ts.umich.edu/resources/
- http://arc-ts.umich.edu/software/torque/
- http://arc-ts.umich.edu/software/
- http://fluxhpc.blogspot.com/
- https://twitter.com/ARCTS_UM/
- http://www.mais.umich.edu/mtoken/
- https://www.globus.org/

Contact Info

- hpc-support@umich.edu (Flux questions)
- msishelp@umich.edu (non-Flux questions)
- Personal Contact Information
  - kgweiss@umich.edu
  - dmontag@umich.edu