GACRC

Introduction to HPC Using zcluster at GACRC

On-class BINF8940

Georgia Advanced Computing Resource Center
University of Georgia
Suchitra Pakala
pakala@uga.edu
Slides By: Zhuofei Hou

Outline

- What is GACRC?
- What is zcluster?
- How does zcluster operate?
- How to work with zcluster?

What is GACRC?

Who Are We?

- **G**eorgia **A**dvanced **C**omputing **R**esource **C**enter
- Collaboration between the Office of Vice President for Research (**OVPR**) and the Office of the Vice President for Information Technology (**OVPIT**)
- Guided by a faculty advisory committee (GACRC-AC)

Why Are We Here?

- To provide computing hardware and network infrastructure in support of high-performance computing (**HPC**) at UGA

Where Are We?

- http://gacrc.uga.edu (Web)
- http://wiki.gacrc.uga.edu (Wiki)
- http://gacrc.uga.edu/help/ (Web Help)
- https://wiki.gacrc.uga.edu/wiki/Getting_Help (Wiki Help)

GACRC

What is zcluster?

- Cluster Structural Diagram
- General Information
- Computing Resources
- Storage Environment

Some Keywords…

- NODE – a single computer
- JOB – unit of work defined by a script that is run on cluster
- QUEUE – order in which the submitted jobs are run
- CLUSTER - set of computers connected together so that, in many respects, they can be viewed as a single system.

GACRC Linux HPC Zcluster Structural Diagram

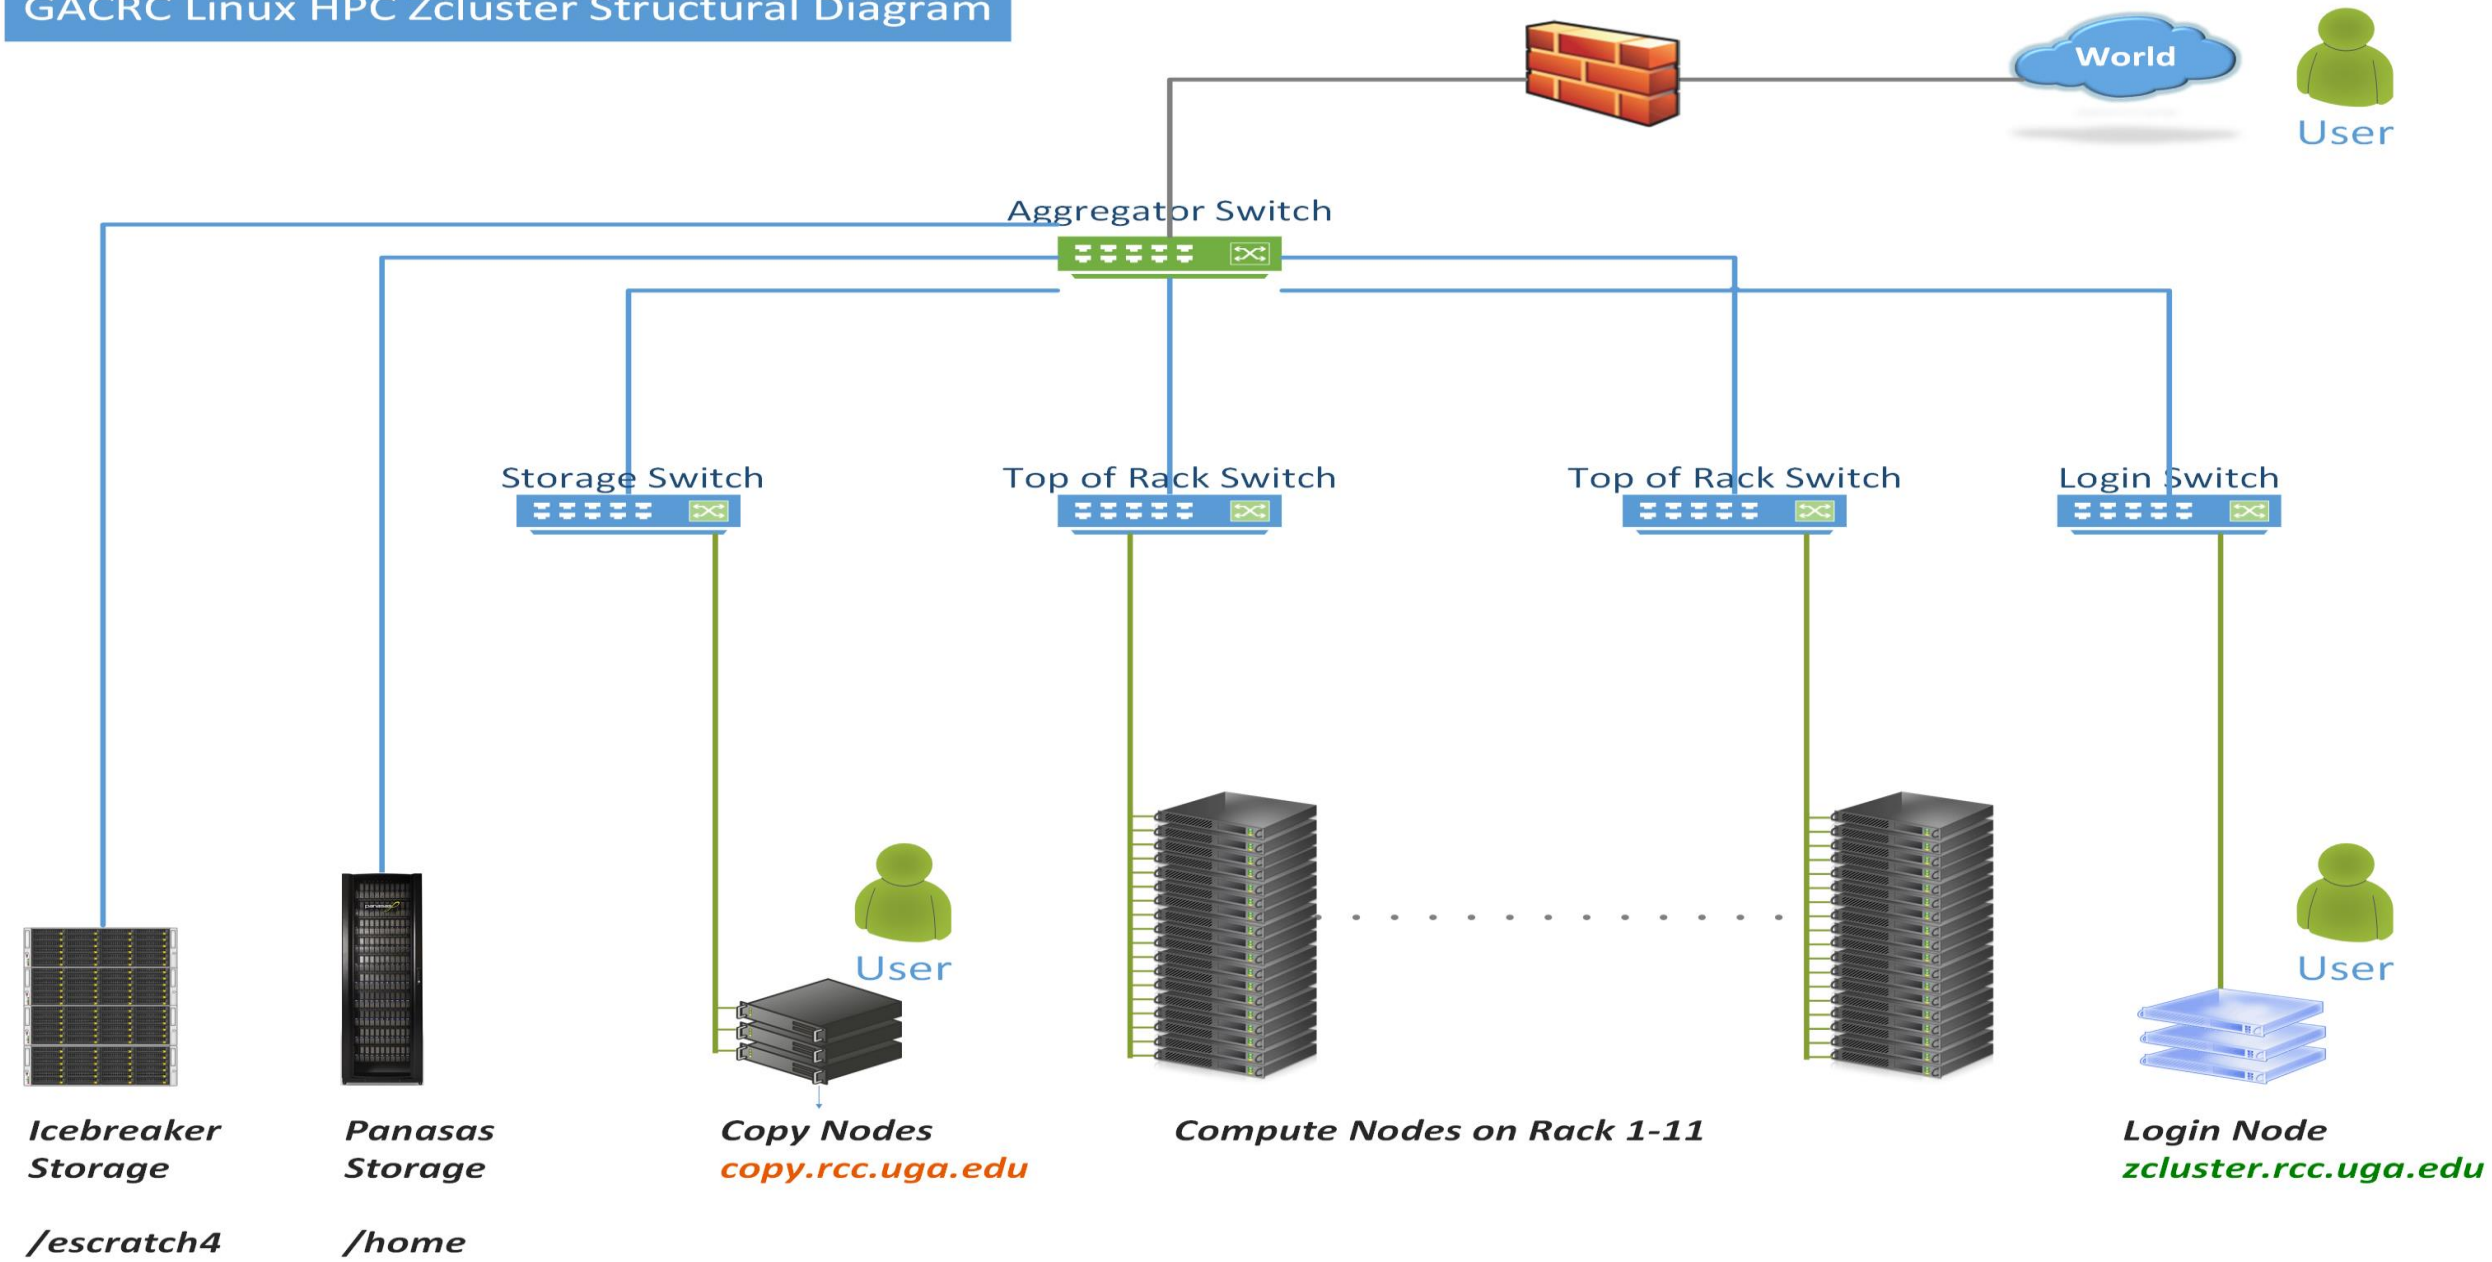

What is zcluster – General Information

GACRC zcluster is a Linux high performance computing (HPC) cluster:

- Operating System: 64-bit Red Hat Enterprise Linux 5 (RHEL 5)
- Login Node: zcluster.rcc.uga.edu
  zcluster.rcc.uga.edu `qlogin` → Interactive Node: compute-14-7/9
- Copy Node: copy.rcc.uga.edu

What is zcluster – Computing Resources

	Queue Type	Queue Name	Nodes #	CPUs on Node	RAM(GB) on Node
1	Regular	rcc-30d	45	12	48
			150	8	16
2	Class	fsr3	1	12	256
Permission 3	High Memory	rcc-m128-30d	1	8	128
			4	8	192
			10	12	256
		rcc-m512-30d	2	32	512

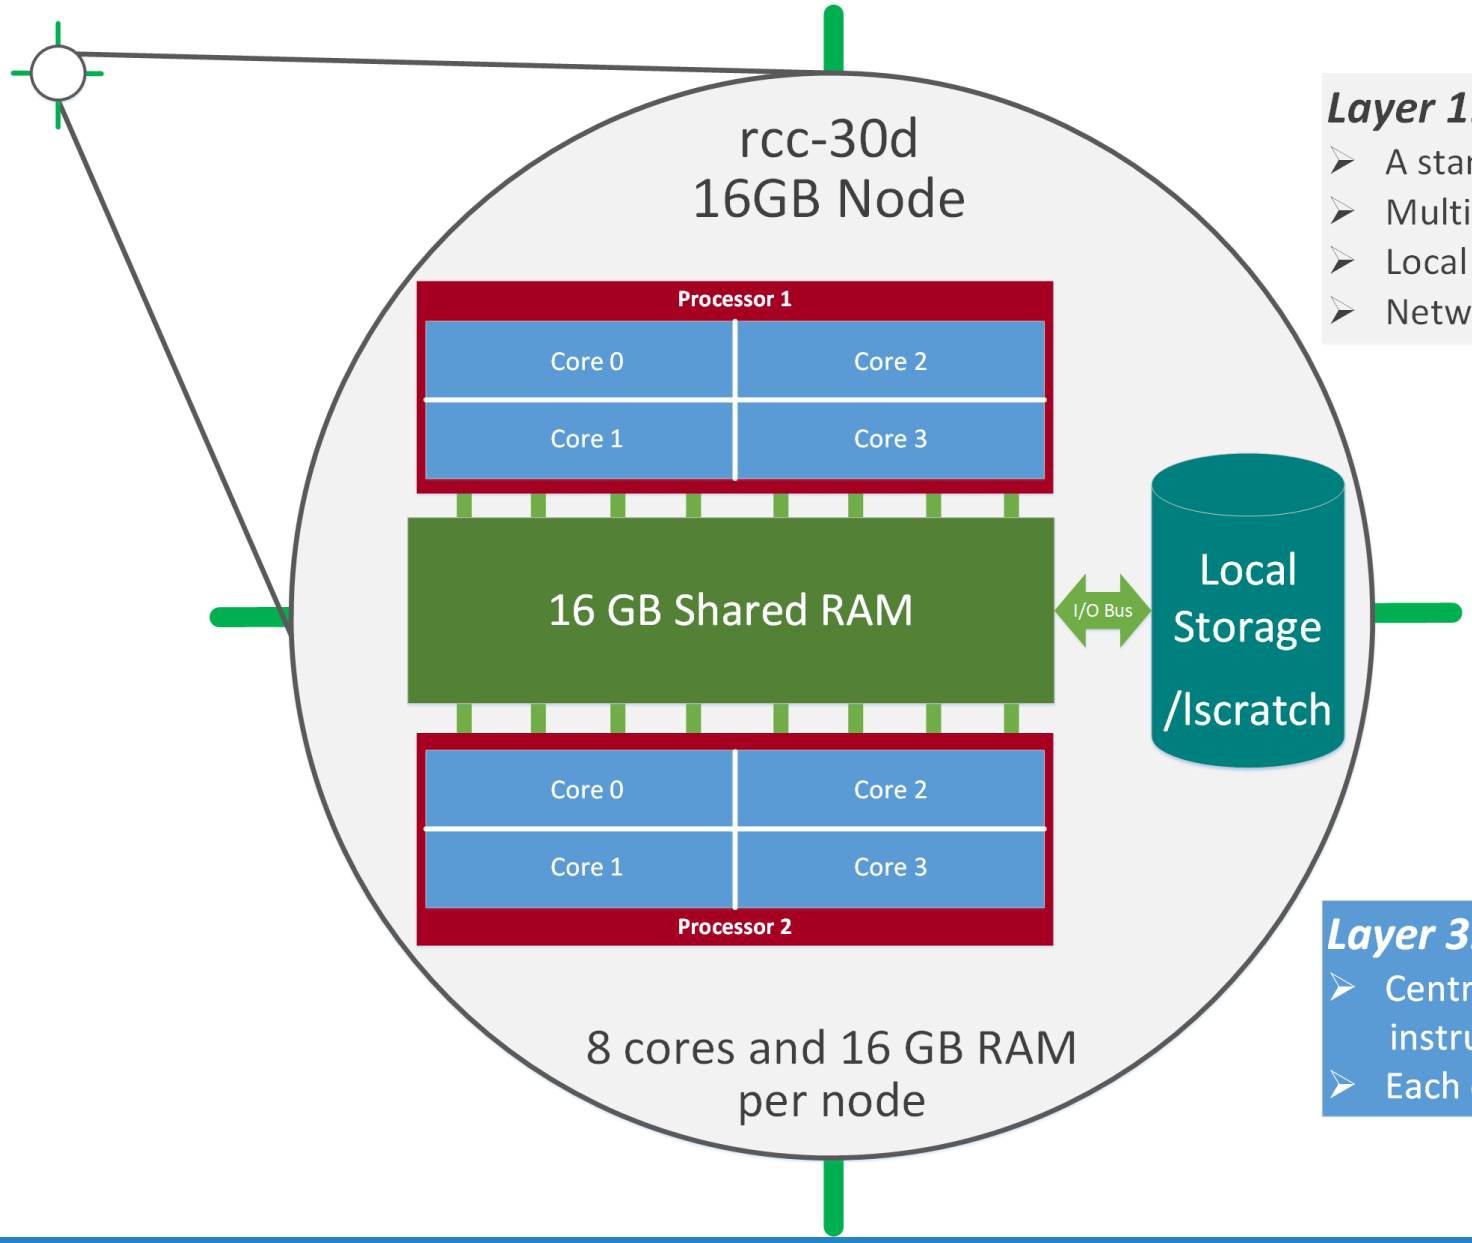

Layer 1: Node

- A standalone "computer in a box"
- Multiple processors, e.g. 2, sharing memory
- Local disk storage, network interface, etc.
- Networked into a cluster

Layer 2: Processor

- A single computing component
- Multicore processor, e.g. 4 cores

Layer 3: Core

- Central processing unit (CPU) reading and executing instructions independently
- Each core is assigned to a software thread

GACRC

What is zcluster – Storage Environment

- **Home directory** ➔ /home/student/binf8940/s_110
  - Mounted and visible on all nodes, with a quota of ~100GB

- **Ephemeral Scratch** ➔ /escratch4/s_110/s_110_Jan_28
  - Create with `make_escratch` command
  - Visible to all nodes with a quota of 4TB
  - To be deleted after 37 days

GACRC

What is zcluster – Storage Environment

2 Filesystems	Role	Quota	Accessible from	Intended Use
/home/student/binf8940/ s_110	Home	100GB	zcluster.rcc.uga.edu (Login) copy.rcc.uga.edu (Copy) Interactive nodes (Interactive)	Highly static data being used frequently
/escratch4/s_110	Scratch	4TB		Temporarily storing large data being used by jobs

3 Main Functions	On/From-Node	Related Filesystem
Login Landing	Login or Copy	/home/student/binf8940/s_110 (Home) (Always!)
Batch Job Submitting	Login or Interactive	/escratch4/s_110 (Scratch) (Suggested!) /home/student/binf8940/s_110 (Home)
Data Archiving , Compressing and Transferring	Copy	/escratch4/s_110 (Scratch) /home/student/binf8940/s_110 (Home)

How does zcluster operate?

Next Page

GACRC zcluster Operational Diagram

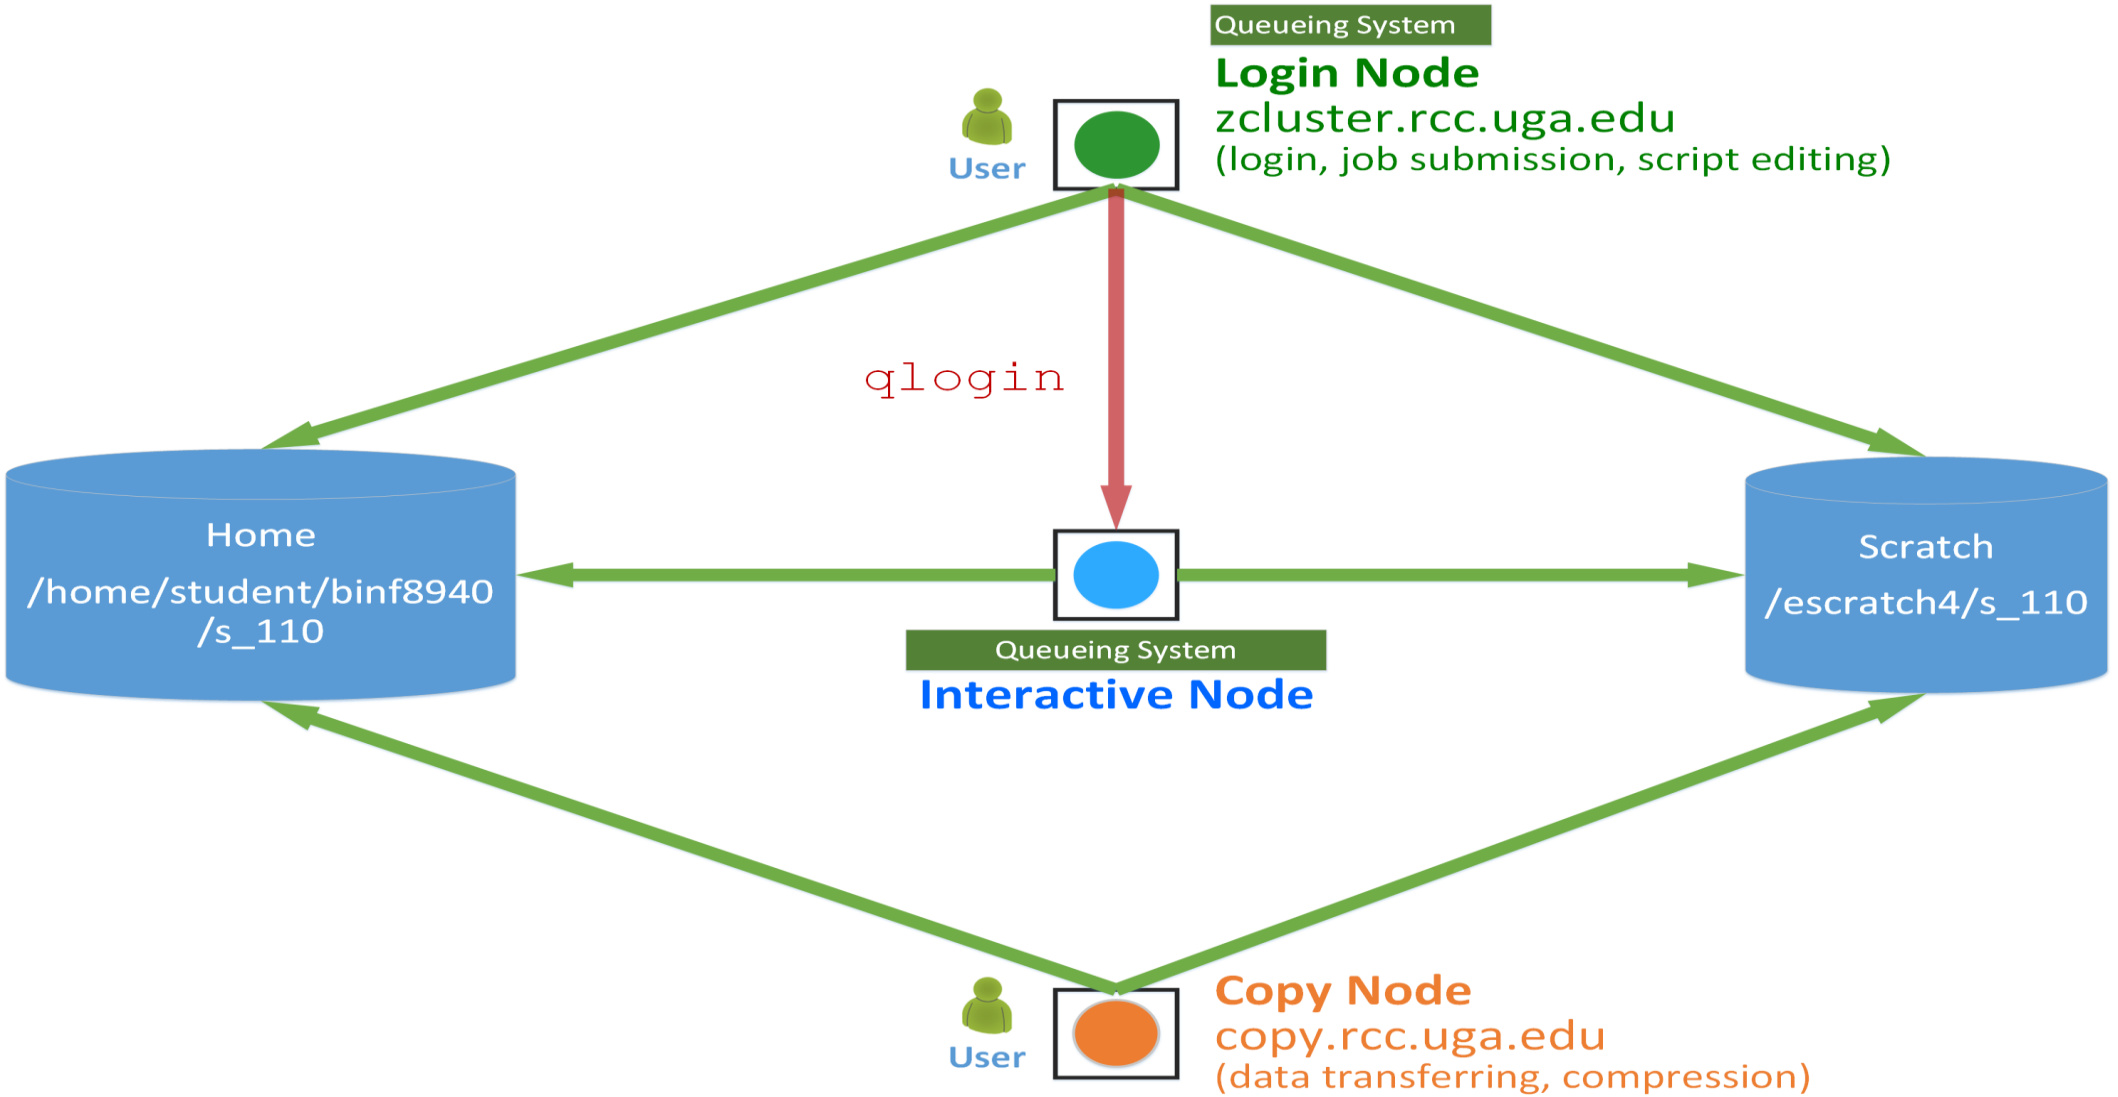

How to work with it?

- Cluster's CCOC:

On cluster, you are not alone……. Each user is sharing finite resources, e.g., CPU cycles, RAM, disk storage, network bandwidth, with other researchers. *What you do may affect other researchers on the cluster*.

6 rules of thumb to remember:

- NO jobs running on login node
- NO multi-threaded job running with only 1 core requested
- NO large memory job running on regular nodes
- NO long job running on interactive node
- NO small memory job running on large memory nodes
- Use the copy node for file transfer and compression

GACRC

How to work with zcluster?

- Start with zcluster
- Connect & Login
- Transfer Files
- Software Installed
- Run Interactive Jobs
- Submit Batch Jobs
  - How to submit ***serial***, ***threaded***, and ***MPI*** batch jobs?
  - How to check job status and cancel a job?
  - How to check memory usage by jobs?

How to work with zcluster – Start with zcluster

- You need a class user account : s_110@zcluster.rcc.uga.edu
- Use `passwd` to change temporary password to a permanent one

How to work with zcluster – Connect & Login

- Open a connection: Open a terminal and `ssh` to your account

```
ssh s_110@zcluster.rcc.uga.edu
```

or

```
ssh -X s_110@zcluster.rcc.uga.edu
```

([1] `-X` is for X windows application running on the cluster to be forwarded to your local machine

[2] If using Windows, use SSH client to open connection, get from UGA download software page)

- Logging in: You will be prompted for your zcluster password

```
s_110@zcluster.rcc.uga.edu's password:
```

([3] On Linux/Mac, when you type in the password, the prompt blinks and does not move)

- Logging out: `exit` to leave the system

```
s_110@zcluster:~$ exit
```

How to work with zcluster – Transfer Files

User's local  

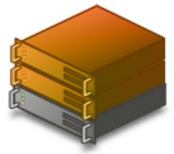 Copy node (copy.rcc.uga.edu)

- On Linux, Mac: `scp [Source] [Target]`

  E.g. 1: working on local machine, Local ➔ zcluster

```
scp file1 s_110@copy.rcc.uga.edu:/escratch4/s_110/s_110_Jan_28/
```

```
scp *.dat s_110@copy.rcc.uga.edu:/escratch4/s_110/s_110_Jan_28/
```

  E.g. 2: working on local machine, zcluster ➔ Local

```
scp s_110@copy.rcc.uga.edu:/escratch4/s_110/s_110_Jan_28/file1 ./
```

```
scp s_110@copy.rcc.uga.edu:/escratch4/s_110/s_110_Jan_28/*.dat ./
```

- On Window: SSH Secure Client File Transfer, WinSCP, FileZilla etc.

How to work with zcluster – Software Installed

- Perl, Python, Java, awk, sed, C/C++ and Fortran compilers
- Matlab, Maple, R, AMBER, GKIN
- Many Bioinformatics applications: NCBI Blast+, Velvet, Trinity, TopHat, MrBayes, SoapDeNovo, SAMtoolsetc.
- RCCBatchBlast (RCCBatchBlastPlus) to , Mafft, RAxML, PASTA, MrBayes and MP-EST distribute NCBI Blast (NCBI Blast+) searches to multiple nodes.
- Many Bioinformatics Databases: NCBI Blast, Pfam, uniprot, etc.
https://wiki.gacrc.uga.edu/wiki/Bioinformatics_Databases
- For a complete list of applications: https://wiki.gacrc.uga.edu/wiki/Software

How to work with zcluster – Run Interactive Jobs

- From login node, qlogin command to open an interactive session:

Login Node: zcluster.rcc.uga.edu —qlogin→ Interactive Node: compute-14-7/9

```
s_110@zcluster:~$ qlogin
Your job 1391816 ("QLOGIN") has been submitted
waiting for interactive job to be scheduled ...
Your interactive job 1391816 has been successfully scheduled.
...
compute-14-7.local$
```

← Now I am on compute-14-7, which is an interactive node

- Current maximum runtime is 12 hours
- When you are done, `exit` to log out!

How to work with zcluster – Submit Batch Jobs

- Components you need to submit a batch job:
  - Software already installed on zcluster
  - Job submission script to run the software, and
    - ✓ Specify working directory
    - ✓ Export environment variables, e.g.,
      - PATH (searching path for executables)
      - LD_LIBRARY_PATH (searching paths for shared libraries)
- Common commands you need:
  - qsub with specifying queue name, threads or MPI processes
  - qstat, qdel
  - qacct, qsj, etc.

How to work with zcluster – Batch *Serial* Job

- **Step 1**: Create a job submission script *ens.sh* running ens:

```
#!/bin/bash                          ➔ Linux default shell (bash)
cd /escratch4/s_20/s_20_Jan_28       ➔ Specify and enter (cd) working directory (/escratch4/s_110/s_110_Jan_28)
time ./ens                           ➔ Run ens with 'time' command to measure amount of time it takes to run
```

- **Step 2**: Submit *ens.sh* to the queue:

```
$ qsub -q rcc-30d  ens.sh
```

Submit a job

job submission script

to the queue rcc-30d
with 16GB RAM/Node

OR

```
$ qsub -q rcc-30d -l mem_total=20g ens.sh
```

to the queue rcc-30d
with 48GB RAM/Node

OR

```
$ qsub -q fsr3 ens.sh
```

to the queue fsr3 with 256GB RAM/Node

How to work with zcluster – Batch *Threaded* Job

- **Step 1**: Create a job submission script *bt.sh* running bowtie2:

```
#!/bin/bash

cd /escratch4/s_110/s_110_Jan_28
export PATH=/usr/local/bowtie2/latest/bin:$PATH
time bowtie2 -p 4 [options]
```

➔ export the searching path for bowtie2

➔ run bowtie2 with 4 threads (-p 4)

Number of Threads = Number of Cores Requested

- **Step 2**: Submit *bt.sh* to the queue:

```
$ qsub -q rcc-30d -l mem_total=20g -pe thread 4 bt.sh
```

to the queue rcc-30d with 48GB RAM/Node

4 cores requested

Note:
Please use the rcc-mc-30d queue, If using threads more than 8!

How to work with zcluster – Batch *MPI* Job

- Step 1: Create a job submission script *raxml.sh* running RAxML:

```
#!/bin/bash

cd /escratch4/s_110/s_110_Jan_28

export MPIRUN=/usr/local/mpich2/1.4.1p1/gcc 4.5.3/bin/mpirun          ➔ Define and export environment variable (MPIRUN)

$MPIRUN -np $NSLOTS /usr/local/raxml/latest/raxmlHPC-MPI-SSE3 [options]  ➔ Run RAxML with 20 MPI processes (-np $NSLOTS )
```

- Step 2: Submit *raxml.sh* to the queue:

```
$ qsub -q rcc-30d -pe mpi 20 raxml.sh
```

20 cores requested,
$NSLOTS will be assigned to 20 automatically, before the job submission script is interpreted

How to work with zcluster – Check and Cancel Jobs

- To check the status of your jobs: qstat

```
$ qstat
job-ID     prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
9707321    0.50766 sub1.sh    s_110        r     01/28/2016 13:39:23 rcc-30d@compute-7-12.local     1
9707322    0.50383 sub2.sh    s_110        Eqw   01/28/2016 13:39:23 rcc-30d@compute-7-12.local     1
9707323    0.00000 sub3.sh    s_110        qw    01/28/2016 13:39:28                                1
```

- To cancel your job with a JobID: qdel

```
$ qdel 9707322
job-ID     prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
9707321    0.50766 sub1.sh    s_110        r     01/28/2016 13:39:23 rcc-30d@compute-7-12.local     1
9707323    0.00000 sub3.sh    s_110        qw    01/28/2016 13:39:28                                1
```

How to work with zcluster – Check Memory Usage

- For running jobs: qsj

```
$ qsj 9707368

==========================================
job_number:              9707368
owner:                   s_110
cwd:                     /escratch4/s_110/s_110_Jan_28
hard_queue_list:         rcc-30d
script_file:             sub.sh
......
usage    1:              cpu=00:01:27, mem=0.96498 GBs,
                         io=0.00014, vmem=73.734M,
                         maxvmem=73.734M
```

- For finished jobs: qacct

```
$ qacct -j  9707323

==============================
qname        rcc-30d
hostname     compute-7-12.local
jobname      sub.sh
jobnumber    9707323
......
cpu          183.320
mem          2.021
io           0.000
maxvmem      5.230G
```

Total Memory

Thank You!

To get support from us: https://wiki.gacrc.uga.edu/wiki/Getting_Help

Please tell us details of the question or problem you have, including but not limited to:

- Your user name, e.g., s_110
- Your job ID
- Your working directory
- The queue name and command you used to submit the job