

# Introduction to HPC Using zcluster at GACRC

On-class BCMB8330

---

Georgia Advanced Computing Resource Center

University of Georgia

Zhuofei Hou, HPC Trainer

[zhuofei@uga.edu](mailto:zhuofei@uga.edu)

# Outline

---

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

# What is GACRC?

---

## Who Are We?

- Georgia **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 (**OVPIIT**)
- 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?

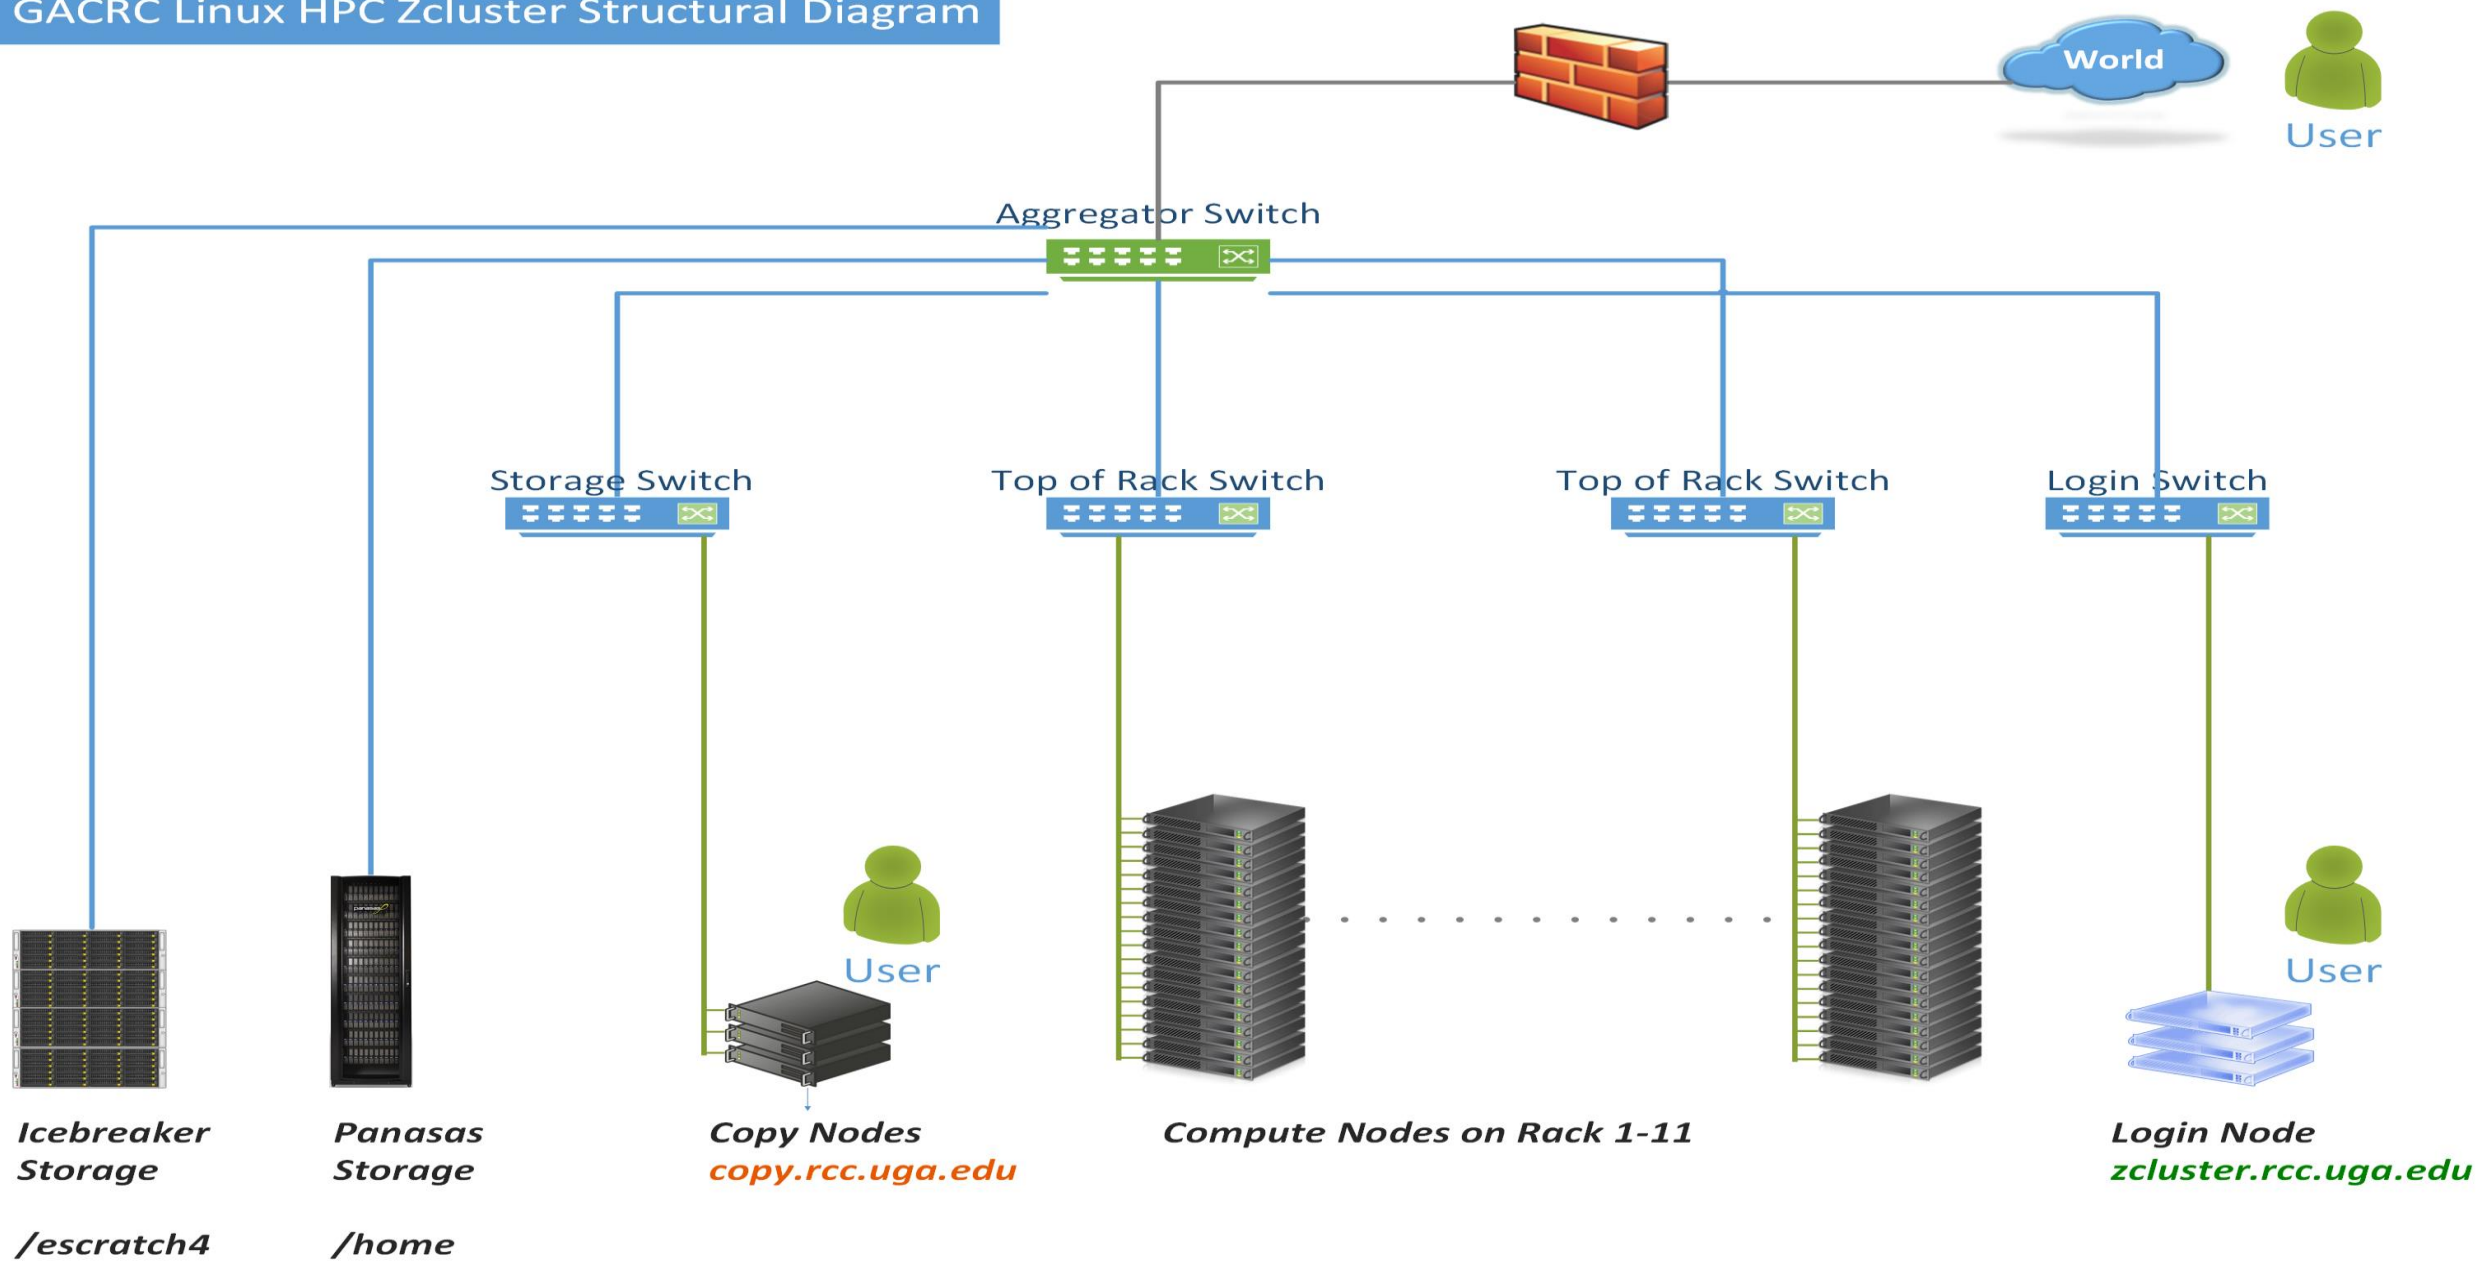
- <http://gacrc.uga.edu> (Web) <http://wiki.gacrc.uga.edu> (Wiki)
- <http://gacrc.uga.edu/help/> (Web Help)
- [https://wiki.gacrc.uga.edu/wiki/Getting\\_Help](https://wiki.gacrc.uga.edu/wiki/Getting_Help) (Wiki Help)

# What is zcluster?

---

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


# GACRC Linux HPC Zcluster Structural Diagram



# What is zcluster – General Information

---

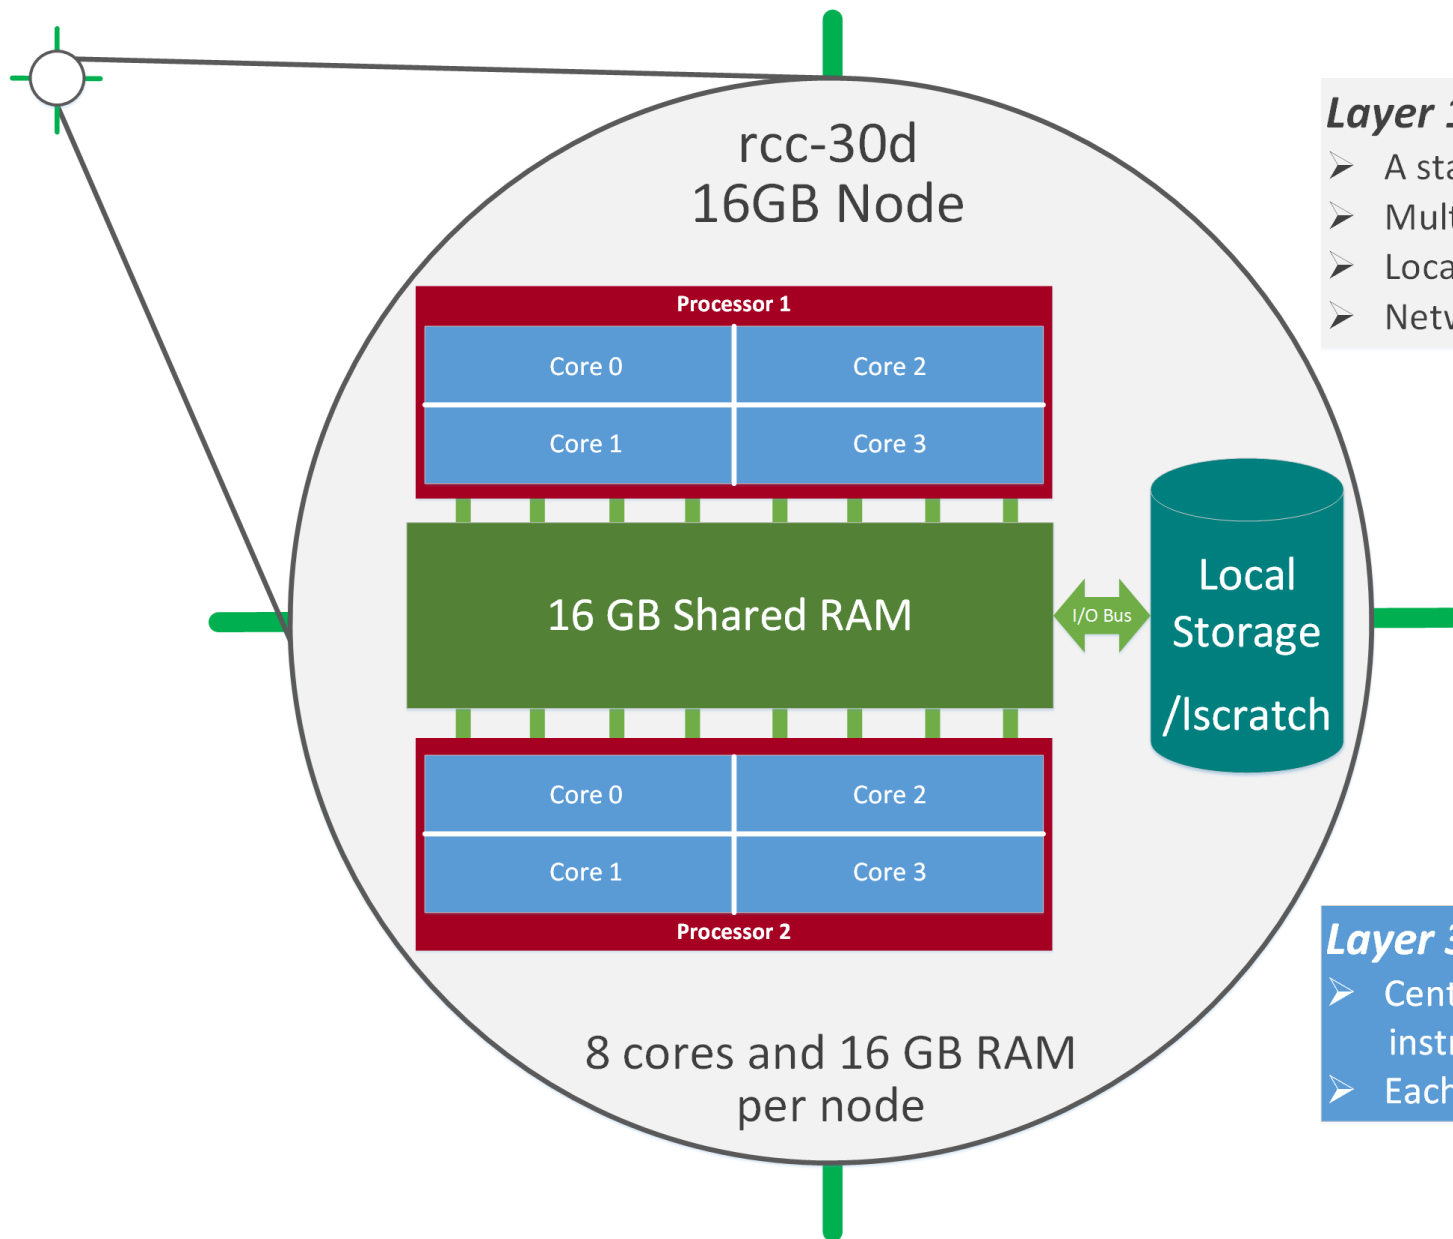
GACRC zcluster is a Linux high performance computing (HPC) cluster:

- Operating System: **64-bit Red Hat Enterprise Linux 5 (RHEL 5)**
  - Login Node: **zcluster.rcc.uga.edu**  Interactive Node: **compute-14-7/9**  
The arrow is a thick, dark gray arrow pointing to the right. The word "qlogin" is written in red text above the arrow.
- Copy Node: **copy.rcc.uga.edu**

# What is zcluster – Computing Resources

Queue Type	Queue Name	Nodes	Processor	CPU Cores/Node	RAM(GB)/Node	NVIDIA GPU
Regular	rcc-30d	45	Intel Xeon	12	48	N/A
		150		8	16	
High Memory	rcc-m128-30d	1	Intel Xeon	8	128	N/A
		4		8	192	
		10		12	256	
	rcc-m512-30d	2		32	512	
Multi Core	rcc-mc-30d	6	AMD Opteron	32	64	N/A
Interactive	interq	2	AMD Opteron	48	132	N/A
GPU	rcc-sgpu-30d	2	Intel Xeon	8	48	4 Tesla S1070 cards
	rcc-mgpu-30d	2		12	48	9 Tesla (Fermi) M2070 cards
	rcc-kgpu-30d	4		12	96	32 Tesla (Kepler) K20Xm cards
	fsr12	1		12	96	8 Tesla (Kepler) K20Xm cards





### 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

8 cores and 16 GB RAM per node



# What is zcluster – Storage Environment

---

- **Home directory** → `/home/student/bcmb8330/s_01`
  - Mounted and visible on **all nodes**, with a quota of **~100GB**
  
- **Ephemeral Scratch** → `/escratch4/s_01/s_01_Feb_9`
  - Create with `make_escratch` command
  - Visible to **all nodes** with a quota of **4TB**
  - To be deleted after **37 days**

# What is zcluster – Storage Environment

2 Filesystems	Role	Quota	Accessible from	Intended Use
/home/student/bcmb8330/s_01	Home	100GB	zcluster.rcc.uga.edu (Login) copy.rcc.uga.edu (Copy)	Highly static data being used frequently
/escratch4/s_01	Scratch	4TB	Interactive nodes (Interactive)	Temporarily storing large data being used by jobs


3 Main Functions	On/From-Node	Related Filesystem
Login Landing	Login or Copy	/home/student/bcmb8330/s_01 (Home) (Always!)
Batch Job Submitting	Login or Interactive	/escratch4/s_01 (Scratch) (Suggested!) /home/student/bcmb8330/s_01 (Home)
Data Archiving , Compressing and Transferring	Copy	/escratch4/s_01 (Scratch) /home/student/bcmb8330/s_01 (Home)

# How does zcluster operate?

---

Next Page





# How to work with zcluster?

---

*Before we start:*

- To get zcluster to be your best HPC buddy

**GACRC Wiki:** <http://wiki.gacrc.uga.edu>

**GACRC Support:** [https://wiki.gacrc.uga.edu/wiki/Getting\\_Help](https://wiki.gacrc.uga.edu/wiki/Getting_Help)

# How to work with zcluster?

---

## *To submit your ticket to GACRC?*

### ➤ **Job Troubleshooting:**

Please tell us details of your question or problem, including but not limited to:

- ✓ Your user name
- ✓ Your job ID
- ✓ Your working directory
- ✓ The queue name and command you used to submit the job

### ➤ **Software Installation:**

- ✓ Specific name and version of the software
- ✓ Download website
- ✓ Supporting package information if have

#### Note:

It's **USER's** responsibility to make sure the **correctness of datasets** being used by jobs!



# How to work with it?

---

- 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



# How to work with zcluster?

---

- Start with zcluster
- Connect & Login
- Transfer Files
- Software Installed
- Run Interactive Jobs
- Submit Batch GPU Job with AMBER
  - How to submit a job
  - How to check job status, cancel a job, etc.



# How to work with zcluster – Start with zcluster

---

- You need a **Class User Account** : `s_01@zcluster.rcc.uga.edu`
- Use `passwd` command to change initial 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_01@zcluster.rcc.uga.edu
```

or

```
ssh -X s_01@zcluster.rcc.uga.edu
```

<sup>(1)</sup> `-X` is for X windows application running on the cluster to be forwarded to your local machine

<sup>(2)</sup> 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_01@zcluster.rcc.uga.edu's password: █
```

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

- Logging out: `exit` to leave the system

```
s_01@zcluster:~$ exit
```

# How to work with zcluster – Transfer Files



- On Linux, Mac or cygwin on Windows : `scp [Source] [Target]`

*E.g. 1:* On local machine, do Local → zcluster

```
scp file1 s_01@copy.rcc.uga.edu:/escratch4/s_01/s_01_Feb_09/
```

```
scp *.dat s_01@copy.rcc.uga.edu:/escratch4/s_01/s_01_Feb_09/
```

*E.g. 2:* On local machine, do zcluster → Local

```
scp s_01@copy.rcc.uga.edu:/escratch4/s_01/s_01_Feb_09/file1 ./
```

```
scp s_01@copy.rcc.uga.edu:/escratch4/s_01/s_01_Feb_09/*.dat ./
```

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

# How to work with zcluster – Run Interactive Jobs

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

Login Node: **zcluster.rcc.uga.edu** <sup>qlogin</sup>  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 – Softwares Installed

---

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

# 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, GPU cards number, etc.
  - **qstat, qdel**
  - **qacct, qsj**, etc.

## How to work with zcluster – Batch GPU Job with AMBER

- Step 1: Create a job submission script *subgpu.sh* running Amber14:

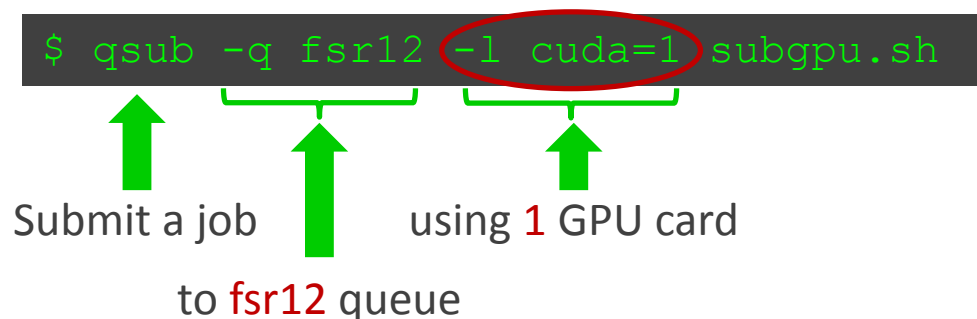
```
#!/bin/bash           → Linux default shell (bash)
cd working_directory → Specify and enter (cd) working directory (e.g., /escratch4/s_01/s_01_Feb_09)
source /usr/local/amber/14/amber.sh

export LD_LIBRARY_PATH=/usr/local/cuda/5.0.35-gcc447/lib64:${LD_LIBRARY_PATH}

/usr/local/amber/14/bin/pmemd.cuda -O -i mdin -o mdout -p prmtop -c inpcrd
```

- Step 2: Submit *subgpu.sh* to **fsr12** queue:

```
$ qsub -q fsr12 -l cuda=1 subgpu.sh
```



Submit a job

to **fsr12** queue

using **1** GPU card

# How to work with zcluster – Check and Cancel Jobs

- To check the status of all queued and running jobs: **qstat**

```

qstat           → shows your job in the pool
qstat -u "*"    → shows all the jobs in the pool
qstat -j 12345  → shows detailed information, e.g., maxvmem, about the job with JOBID 12345
qstat -g t      → list all nodes used by your jobs
  
```

- To cancel a queued or running job: **qdel**

```

qdel -u zhuofei → deleted all your jobs
qdel 12345      → deletes your job with JOBID 12345
  
```

- To list detailed information about a job: **qsj, qacct**

```

qsj 12345       → shows information, e.g., maxvmem, about the RUNNING job with JOBID 12345
qacct -j 12345  → shows information, e.g., maxvmem, about the ENDED job with JOBID 12345
  
```



Thank You!

A solid blue horizontal bar at the bottom of the slide.