Georgia Advanced Computing Resource Center
UNIVERSITY OF GEORGIA

GACRC User Training:

Migrating from Zcluster to Sapelo

The GACRC Staff
Version 1.0

Discussion Points

I. Request Sapelo User Account

II. Systems

III. Transfer Files

IV. Configure Software Environment

V. Job Submission Workflow

VI. Job Submission Scripts

VII. Check Job information and Cancel Job

VIII. Appendices

Request Sapelo User Account

Sapelo cluster user account: MyID@sapelo1.gacrc.uga.edu

Note: A valid official UGA MyID is a MUST to create a user account!

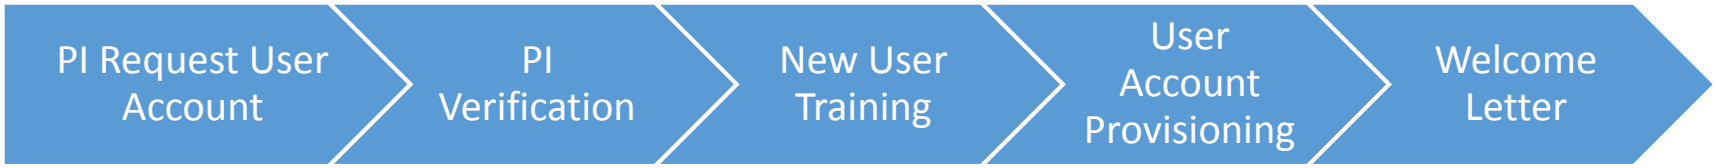

1. The UGA PI uses the GACRC online form http://help.gacrc.uga.edu/userAcct.php to request a user account for a group member.
2. Once we received the request, we will verify it with the PI.
3. After verification by the PI, the new user will be required to attend a training session.
4. After the user attended training, we will provision a Sapelo account for the user.
5. A welcome letter is sent to the user once user account is ready.

Systems – Diagram

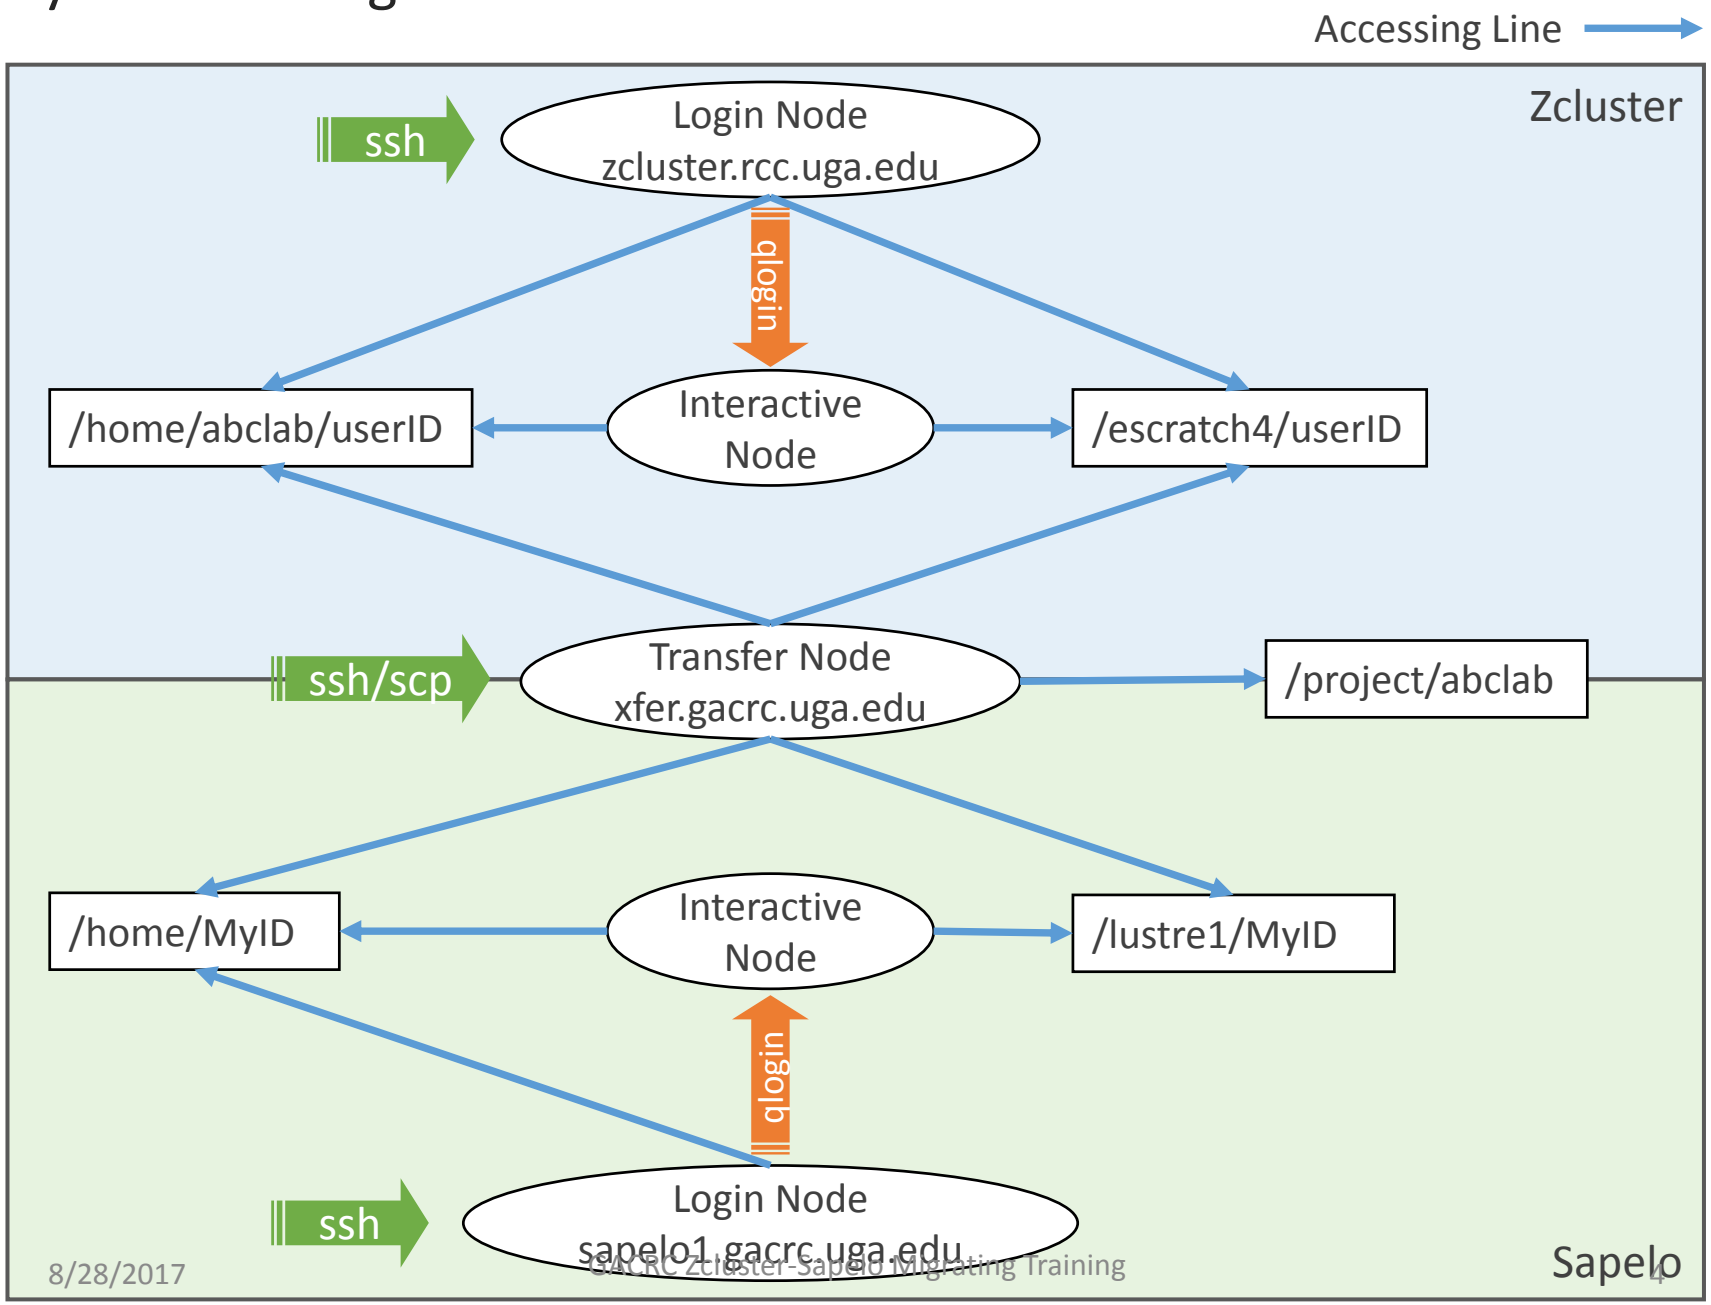

Systems – Overview

Log on to Zcluster Login Node:	***'ssh userID@zcluster.rcc.uga.edu'*** (using zcluster password)
Open Zcluster Interactive Node:	type ***'qlogin'*** command from login node
Zcluster Home Directory:	/home/abclab/userID
Zcluster Global Scratch:	/escratch4/userID (4TB quota)
Create Working Directory on Zcluster Global Scratch:	type '***make_escratch***' command from login node
Log on to Sapelo Login Node:	'***ssh MyID@sapelo1.gacrc.uga.edu***' (using UGA MyID password)
Open Sapelo Interactive Node:	type ***'qlogin'*** command from login node
Sapelo Home Directory:	/home/MyID (100GB quota)
Sapelo Global Scratch:	/lustre1/MyID
Create Working Directory on Sapelo Global Scratch:	use ***'mkdir'*** command in /lustre1/MyID from interactive node
Log on to GACRC Transfer Node:	***'ssh MyID@xfer.gacrc.uga.edu'*** (using UGA MyID password)
Group Project Space:	/project/abclab ("abc" are initials of the PI of the group)

Sapelo Cluster Diagram

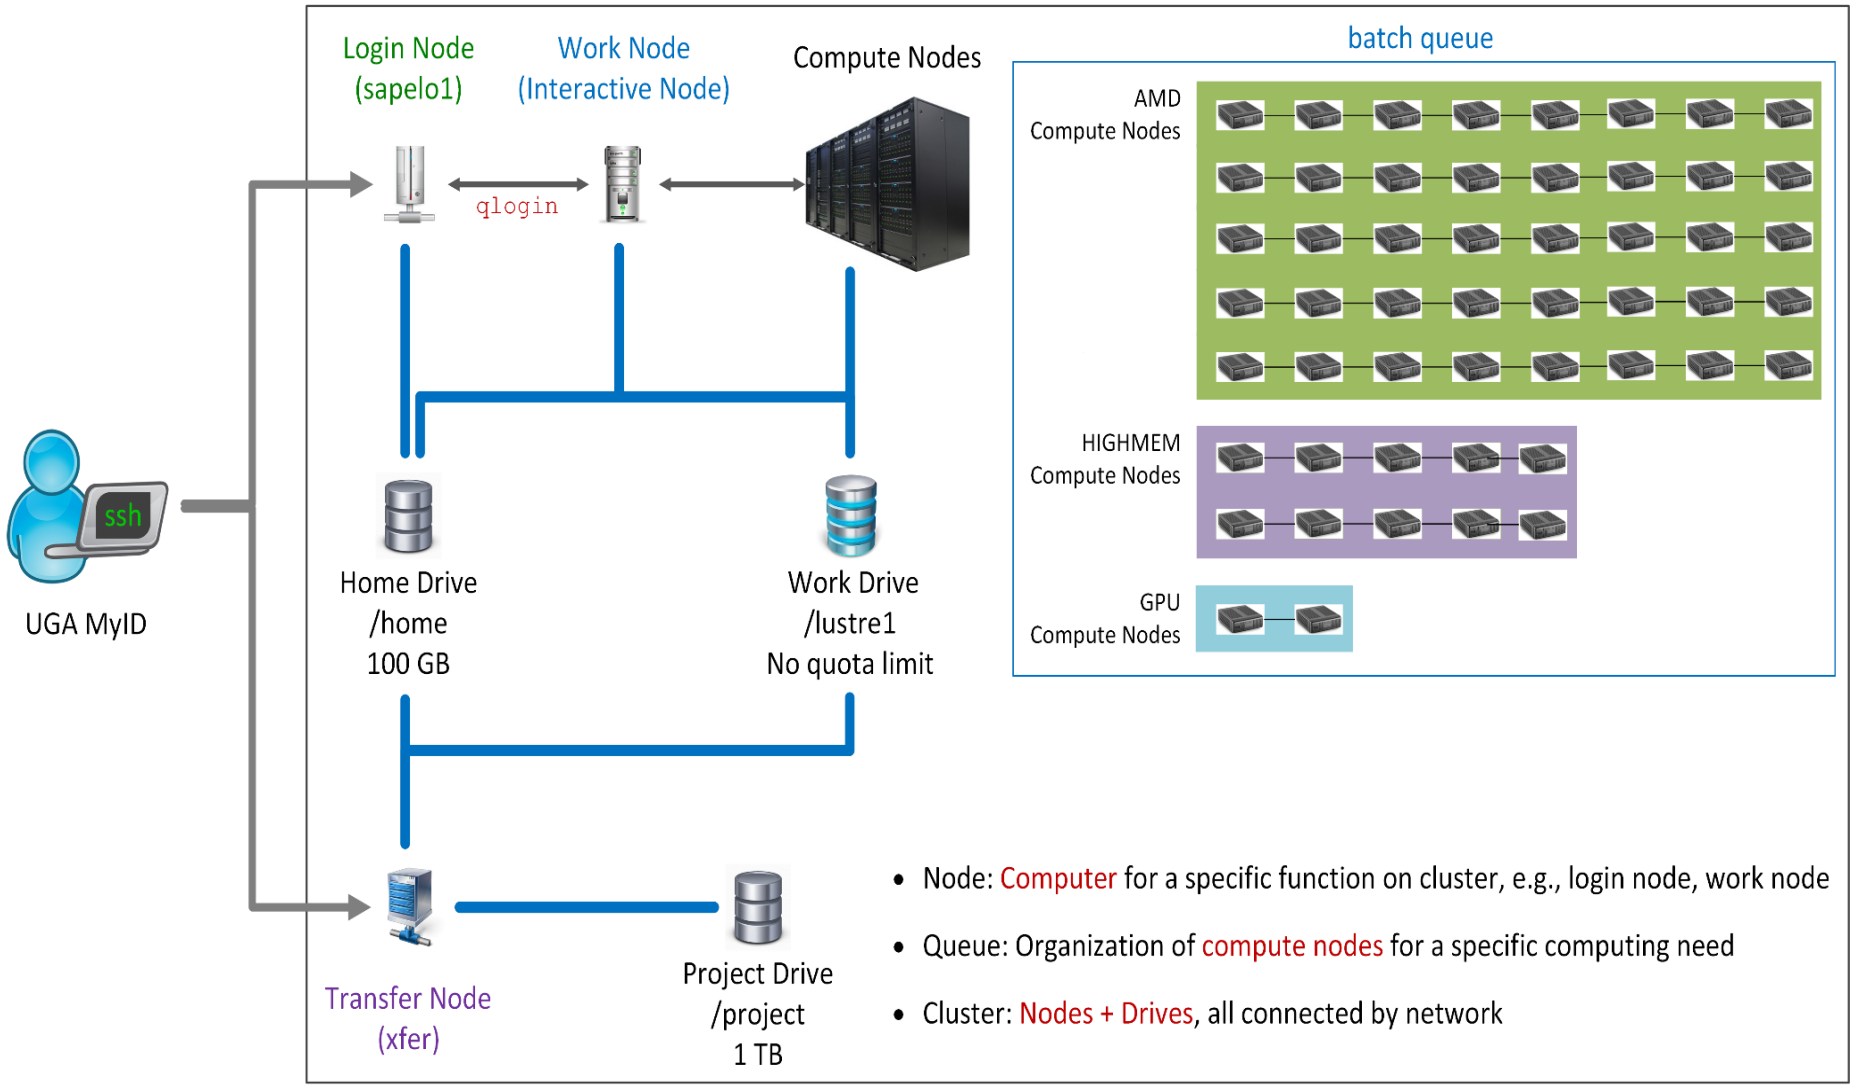

- Node: Computer for a specific function on cluster, e.g., login node, work node
- Queue: Organization of compute nodes for a specific computing need
- Cluster: Nodes + Drives, all connected by network

Systems – Storage Environment

https://wiki.gacrc.uga.edu/wiki/Disk_Storage

Name	Role	Purpose	Zcluster	Sapelo
Home	ssh landing spot	static data, e.g., scripts, software	/home/abclab/userID	/home/MyID (100GB quota)
Global Scratch	Job working space	Current Job data being used by running jobs	/escratch4/userID (4TB quota)	/lustre1/MyID (no quota – limited by capacity)
Project	Active data storage	Non-current active data for future analysis	/project/abclab (1TB initial quota)	

Systems – Computing Resources

https://wiki.gacrc.uga.edu/wiki/Systems

Zcluster	150x 8-core 16GB RAM nodes with Intel Xeon processors (total 1200 cores) 45x 12-core 48GB RAM nodes with Intel Xeon processors (total 540 cores)
	4x 8-core 192GB RAM nodes with Intel Xeon processors (total 32 cores) 10x 12-core 256GB RAM nodes with Intel Xeon processors (total 120 cores) 2x 32-core 512GB RAM nodes with Intel Xeon processors (total 64 cores)
	6x 32-core 64GB RAM nodes with AMD Opteron processors (total 192 cores)
	4x 12-core 96GB RAM nodes with Intel Xeon processors, 8 NVIDIA K20Xm GPU cards each
Sapelo	112x 48-core 128GB RAM nodes with AMD Opteron processors (total 5376 cores) 4x 48-core 256GB RAM nodes with AMD Opteron processors (total 192 cores)
	6x 48-core 512GB RAM nodes with AMD Opteron processors (total 288 cores) 1x 48-core 1TB RAM node with AMD Opteron processors (total 48 cores) 3x 28-core 1TB RAM node with Intel Xeon processors (total 84 cores)
	2x 16-core 128GB RAM nodes with Intel Xeon processors, 8 NVIDIA K40m GPU cards each
	85+ buy-in nodes with variable configurations

Systems – Sapelo Computing Resources

https://wiki.gacrc.uga.edu/wiki/Systems

Total Nodes	Cores /Node	RAM(GB) /Node	Processor Type	Queue Name	Node Feature	NVIDIA GPU	InfiniBand
112	48	128	AMD Opteron	batch	AMD	N/A	Yes
4		256					
6		512			HIGHMEM		
4		1024 (1)					
	28	1024 (3)	Intel Xeon				
2	16	128			GPU	16 K40m	
85	variable	variable	variable		abcnode (buy-in)	variable	

Transfer of Files

https://wiki.gacrc.uga.edu/wiki/Transferring_Files

All Sapelo and zcluster users, as well as all GACRC PIs, can access the GACRC transfer nodes (xfer.gacrc.uga.edu) and using UGA **MyID** and **MyID password** to authenticate.

Note that a user's home directory on transfer node is the same as the user's Sapelo home directory, which is not the same as the user's Zcluster home directory.

Transfer nodes can access the following file systems using their **full paths**:

Sapelo Home Directory:	***cd /home/MyID***
Zcluster Home Directory:	***cd /panfs/pstor.storage/home/abclab/userID***
Zcluster Global Scratch:	***cd /escratch4/userID***
Sapelo Global Scratch:	***cd /lustre1/MyID***
Group Project Folder	***cd /project/abclab***

Detailed instructions can be found at the link above

File Transferring Tree

https://wiki.gacrc.uga.edu/wiki/Transferring_Files

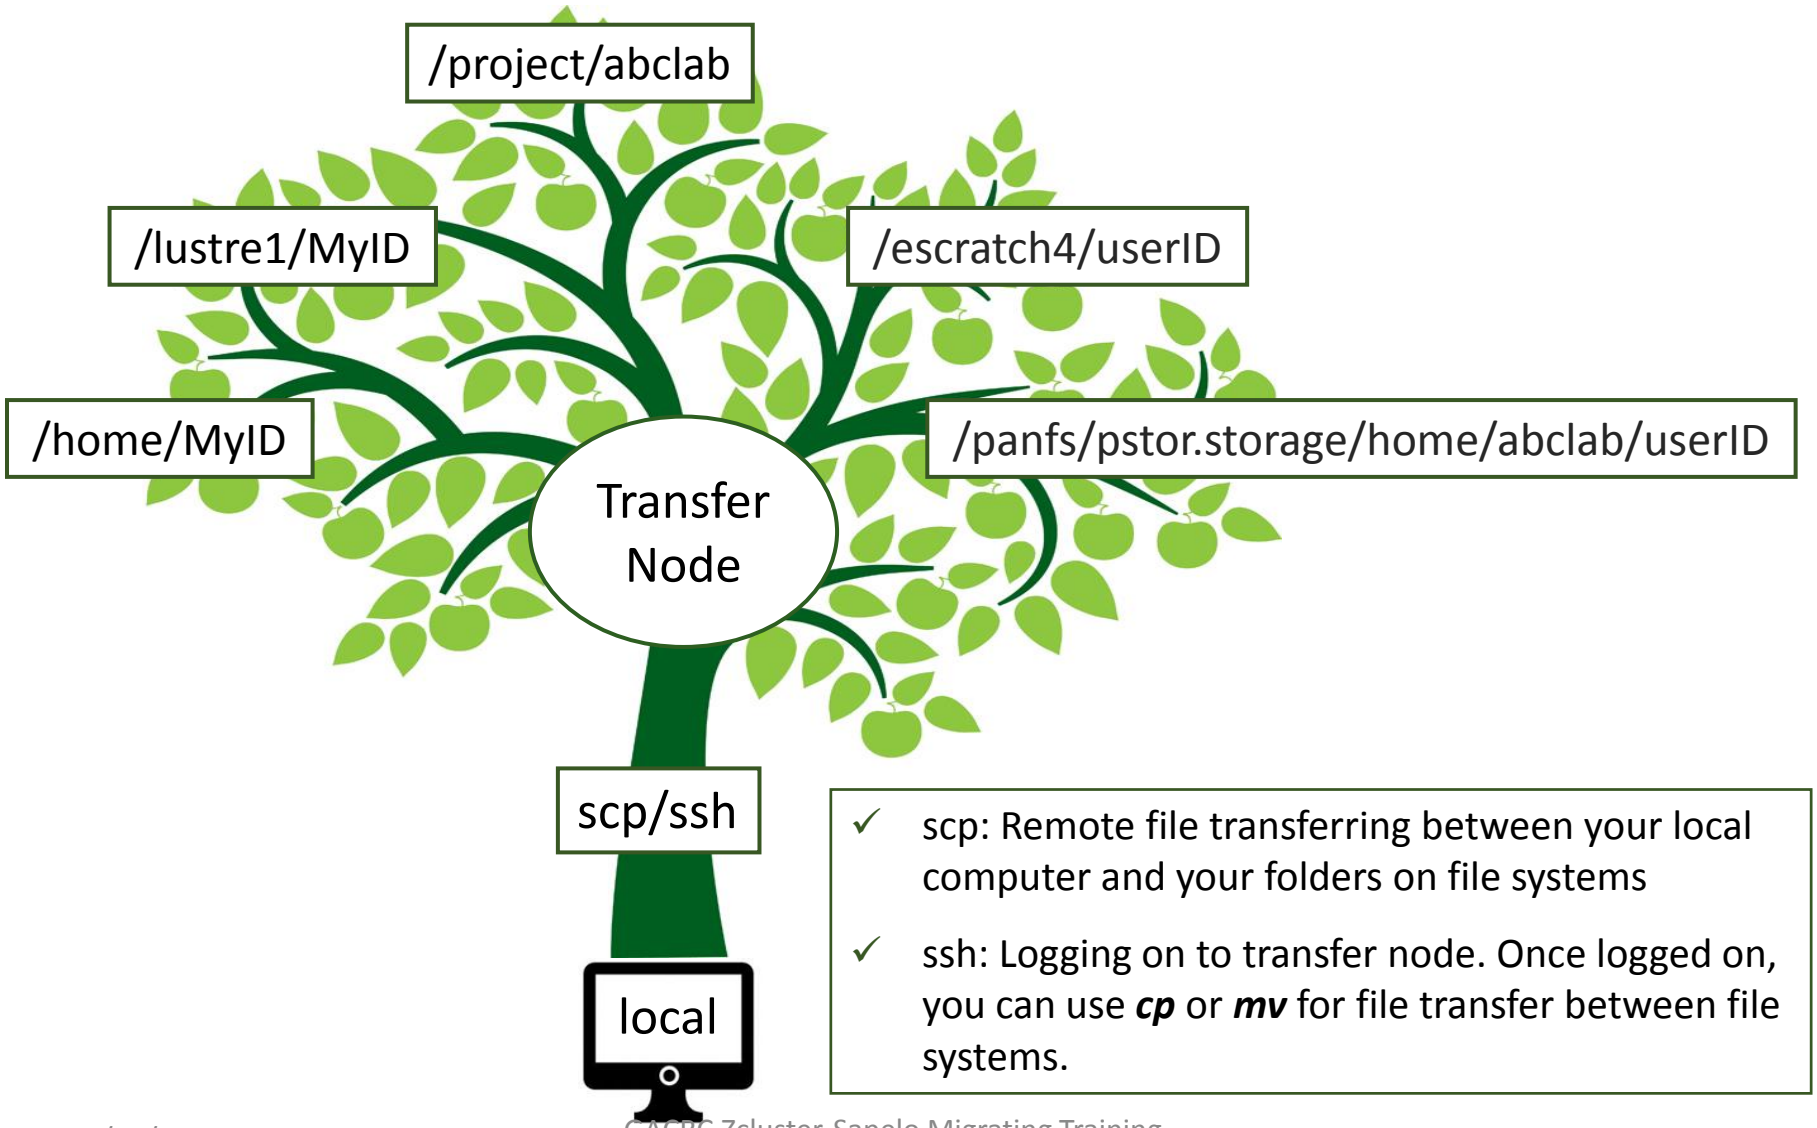

- ✓ scp: Remote file transferring between your local computer and your folders on file systems
- ✓ ssh: Logging on to transfer node. Once logged on, you can use ***cp*** or ***mv*** for file transfer between file systems.

Configure Software Environment

https://wiki.gacrc.uga.edu/wiki/Software

Zcluster

User "manually" exports various paths for software packages (in a job submission script or on the command line of an interactive session), for example:

export PATH=……:$PATH

Sapelo

User issues module commands to configure various paths for software packages (in a job submission script or on the command line of an interactive session):

1. ***module avail*** : show all modules available on cluster
2. ***module load moduleName/version*** : load the needed module "moduleName" with a version
3. ***module list*** : list all modules you already loaded

For example:

$ ***module load python/2.7.8***
$ ***module load R/3.2.3***
$ ***module list***
Currently Loaded Modules:
1) moab/8.1.1 2) StdEnv 3) python/2.7.8 4) R/3.2.3

Job Submission Workflow (general)

Zcluster	Sapelo
1. Log on to zcluster login node: ***ssh userID@zcluster.rcc.uga.edu*** 2. Create a working dir on global scratch: ***make_escratch*** 3. Change to working dir on global scratch: ***cd /escratch4/userID/userID_mmm_dd*** (mmm, dd: date abbreviations) 4. Create a zcluster job submission script: ***nano ./sub.sh*** 5. Submit job: ***qsub -q rcc-30d ./sub.sh***	1. Log on to Sapelo login node: ***ssh MyID@sapelo1.gacrc.uga.edu*** 2. Go to Sapelo Interactive node: ***qlogin*** 3. Change to Sapelo global scratch: ***cd /lustre1/MyID*** 4. Create a working dir on global scratch: ***mkdir ./workDir*** 5. Change to working dir on global scratch: ***cd ./workDir*** 6. Create a Sapelo job submission script: ***nano ./sub.sh*** 7. Submit job: ***qsub ./sub.sh***

Job Submission Scripts (NCBI Blast as example)

sub.sh on Zcluster	sub.sh on Sapelo
#!/bin/bash *cd `pwd`* *PATH=/usr/local/ncbiblast/latest/bin:$PATH* *export PATH* *blastall -p blastn -d /db/ncbiblast/latest/nt * *-a 4 -i QUERY * *-o out.QUERY*	*#PBS -S /bin/bash* *#PBS -q batch* *#PBS -N testBlast* *#PBS -l nodes=1:ppn=4:AMD* *#PBS -l mem=20gb* *#PBS -l walltime=480:00:00* *cd $PBS_O_WORKDIR* *module load ncbiblast+/2.2.26* *blastall -p blastn -d /db/ncbiblast/nrte/latest/nt * *-a 4 -i QUERY * *-o out.QUERY > outputfile 2>&1*
qsub -q rcc-30d -pe thread 4 ./sub.sh	***qsub ./sub.sh***

Sapelo Serial Job Submission Script (NCBI Blast as example)

#PBS -S /bin/bash	➔ Linux shell (bash)
#PBS -q batch	➔ Queue name (batch)
#PBS -N testBlast	➔ Name of the job (testBlast)
#PBS -l nodes=1:ppn=1:AMD	➔ Number of nodes (1), number of cores/node (1), node type (AMD)
#PBS -l mem=20gb	➔ Maximum amount of physical memory (20 GB) used by the job
#PBS -l walltime=48:00:00	➔ Maximum wall clock time (48 hours) for the job, default 6 minutes
cd $PBS_O_WORKDIR	➔ Use the dir from which the job is submitted as the working dir
module load ncbiblast+/2.2.29	➔ Load the module of ncbiblast+, version 2.2.29
time blastn [options] >outputfile 2>&1	➔ Run blastn with 'time' command to measure the amount of time it takes to run the application

Sapelo Threaded Job Submission Script (NCBI Blast as example)

```
#PBS -S /bin/bash
#PBS -q batch
#PBS -N testBlast
#PBS -l nodes=1:ppn=24:AMD
#PBS -l mem=20gb
#PBS -l walltime=480:00:00

#PBS -M jSmith@uga.edu
#PBS -m ae
#PBS -j oe

cd $PBS_O_WORKDIR

module load ncbiblast+/2.2.29

time blastn -num_threads 24 [options] >outputfile 2>&1
```

➔ Number of nodes (1), number of cores/node (24), node type (AMD)
Number of threads (24) = Number of cores requested (24)

➔ Email address to receive a notification for computing resources
➔ Send email notification when job aborts (a) or terminates (e)
➔ Std error file will be merged into stdout file

➔ *Run 24 threads (-num_threads 24)*

Sapelo MPI Job Submission Script (RAxML as example)

```
#PBS -S /bin/bash
#PBS -q batch
#PBS -N testRAxML
#PBS -l nodes=2:ppn=48:AMD
#PBS -l walltime=480:00:00
#PBS -l mem=20gb

cd $PBS_O_WORKDIR
module load raxml/8.1.20
module load intel/15.0.2
module load openmpi/1.8.3/intel/15.0.2

mpirun raxmlHPC-MPI-AVX [options] >outputfile 2>&1
```

➔ **Number of nodes (2), number of cores/node (48), node type (AMD)**

Total cores requested = 2 × 48 = 96

Number of MPI Processes ≤ Number of cores requested (96)

➔ **Run with default 96 MPI processes**

Sapelo MPI Job Submission Script (RAxML as example)

```
#PBS -S /bin/bash
#PBS -q batch
#PBS -N testRAxML
#PBS -l nodes=2:ppn=27:AMD
#PBS -l walltime=480:00:00
#PBS -l mem=20g

cd $PBS_O_WORKDIR

# Context Sharing
CONTEXTS=$(/usr/local/bin/set_contexts.sh $PBS_NUM_PPN)
if [[ "$?" -eq "0" ]] ; then
  export PSM_SHAREDCONTEXTS_MAX=$CONTEXTS
fi

module load raxml/8.1.20
module load intel/15.0.2
module load openmpi/1.8.3/intel/15.0.2

mpirun –np 50 raxmlHPC-MPI-AVX [options] >outputfile 2>&1
```

➔ ppn number (27) fewer than 48 MUST be a multiplier of 3!

copied from GACRC Wiki.
4 lines are REQUIRED if you run a MPI job with a ppn number smaller than 48!

➔ Run with 50 MPI processes (–np 50), default 54

Check Job Information (qstat) and Cancel a Job (qdel)

```
$ qstat
Job ID                    Name             User            Time Use S Queue
------------------------------------------------------------------------
481929.pbs                testJob1         jSmith          900:58:0 C batch
481931.pbs                testJob2         jSmith          04:00:03 R batch
481934.pbs                testJob3         jSmith                 0 Q batch
```

Job status:
R : job is running
C : job completed (or canceled or crashed) and is not longer running. Jobs stay in this state for 1h
Q : job is pending, waiting for resources to become available

```
$ qdel 481934
$ qstat
Job ID                    Name             User            Time Use S Queue
------------------------------------------------------------------------
481929.pbs                testJob1         jSmith          900:58:0 C batch
481931.pbs                testJob2         jSmith          04:00:03 R batch
481934.pbs                testJob3         jSmith                 0 C batch
```

Check Computing Resources Used by a Job

Option 1: **`qstat -f JobID`** for ***running jobs*** or ***finished jobs in 1 hour***

Option 2: **`showjobs JobID`** for ***finished jobs over 1 hour, but ≤ 7 days***

Option 3: Email notification from ***finished jobs (completed, canceled, or crashed)***, if using:

```
#PBS -M MyID@uga.edu
#PBS -m ae
```

Georgia Advanced Computing Resource Center
UNIVERSITY OF GEORGIA

Outputs of `qstat -f JobID`

```
$ qstat -f 1426778
Job Id: 1426778.pbs.scm
    Job_Name = BS2_GFF_range_methylator5-10.py
    Job_Owner = gent@uga-2f0f976.scm
    resources_used.cput = 76:54:46
    resources_used.energy_used = 0
    resources_used.mem = 44136kb
    resources_used.vmem = 90200kb
    resources_used.walltime = 76:50:04
    job_state = C
    queue = batch
            .
    Error_Path = uga-2f0f976.scm:/home/gent/scripts/BS2_GFF_range_methylator5-
10.py.e1426778
    exec_host = n2/7
    Output_Path = uga-2f0f976.scm:/home/gent/scripts/BS2_GFF_range_methylator5-
10.py.o1426778
            .
    Resource_List.mem = 2gb
    Resource_List.nodect = 1
    Resource_List.nodes = 1:ppn=1:AMD
    Resource_List.walltime = 480:00:00
    .
    Variable_List = PBS_O_QUEUE=batch,PBS_O_HOME=/home/zhuofei, …… ,
                    PBS_O_WORKDIR=/home/zhuofei/MPIs,
```

Outputs of `showjobs JobID`

```
$ showjobs 1426778
Job Id             : 1426778.pbs.scm
Job Name           : BS2_GFF_range_methylator5-10.py
Output File        : uga-2f0f976.scm:/home/gent/scripts/BS2_GFF_range_methylator5-
10.py.o1426778
Error File         : uga-2f0f976.scm:/home/gent/scripts/BS2_GFF_range_methylator5-
10.py.e1426778
Working Directory  : /home/gent/scripts
Home Directory     : /home/gent
Submit Arguments   : BS2_GFF_range_methylator5-10.py.sh
User Name          : gent
Group Name         : krdlab
Queue Name         : batch
Wallclock Limit    : 20:00:00:00
Wallclock Duration: 3:04:50:04
CPUTime            : 3:04:54:46
Memory Used        : 43.1Mb
Memory Limit       : 2gb
vmem Used          : 88.1Mb
Submit Time        : Mon Mar 27 20:51:01 2017
Start Time         : Tue Mar 28 03:06:41 2017
End Time           : Fri Mar 31 07:58:32 2017
Exit Code          : 0
Master Host        : n2
```

Georgia Advanced Computing Resource Center
UNIVERSITY OF GEORGIA

Email notification from finished jobs

```
PBS Job Id: 700009.pbs.scm
Job Name:   testJob
Exec host:  n1/4-27
Execution terminated
Exit_status=0
resources_used.cput=00:05:12
resources_used.energy_used=0
resources_used.mem=410984kb
resources_used.vmem=6548516kb
resources_used.walltime=04:00:59
Error_Path: uga-
2f0f976.scm:/home/zhuofei/MPIs/testJob.o700009
Output_Path: uga-
2f0f976.scm:/home/zhuofei/MPIs/testJob.o700009
```

```
PBS Job Id: 700097.pbs.scm
Job Name:   testJob
Exec host:  n5/4-27
Execution terminated
Exit_status=271
resources_used.cput=00:11:22
resources_used.energy_used=0
resources_used.mem=412304kb
resources_used.vmem=6548524kb
resources_used.walltime=05:00:41
Error_Path: uga-
2f0f976.scm:/home/zhuofei/MPIs/testJob.o700097
Output_Path: uga-
2f0f976.scm:/home/zhuofei/MPIs/testJob.o700097
```

Appendix 1: Common Command Summary

https://wiki.gacrc.uga.edu/wiki/Command_List

Tasks	Zcluster	Sapelo
Common file/dir operations	ls, cd, pwd, file, cat, less, more, cp mv, mkdir, rm, rmdir, dos2unix, mac2unix, etc.	
Transfer remote files	scp, sftp or SSH Secure Client File Transfer, FileZilla, WinSCP	
Open interactive session	qlogin	
Create working directory on global scratch	make_escratch	mkdir subdirectoryName
Edit script files	nano sub.sh or vim sub.sh	
Submit a batch job	qsub –q queueName –l … –pe … sub.sh	qsub sub.sh
check job status	qstat	
Check detailed job info	qsj jobID, qstat –j jobID, qacct –j jobID	qstat –f jobID, showjobs jobID
Cancel a job from queue	qdel jobID, qdel –u userID	qdel jobID

Appendix 2: Current process for a UGA PI to request (re)activation or creation of a MyID for an external collaborator

To keep a user's MyID active ***after*** he/she leaves UGA, or to create a MyID for an external research collaborator:

1. The PI accesses the EITS helpdesk portal at https://uga.teamdynamix.com/TDClient/Home/
2. Click on "MyID Account Request" on the right side of the page in "Popular Services" section.
3. Use MyID credentials to log on the service page.
4. In the MyID Request From, enter his/her information and select "GACRC Remote Collaborator" for the Account Type field.
5. After verification by the GACRC Director, EITS will generate collaborator's MyID
6. If collaborator is visiting UGA campus, MyID can be obtained through OVPR's Visiting Researcher & Scholar program – https://research.uga.edu/visiting/

Appendix 3: Useful Links

Georgia Advanced Computing Resource Center
UNIVERSITY OF GEORGIA

- **GACRC Website:** http://gacrc.uga.edu
- **GACRC Wiki:** https://wiki.gacrc.uga.edu/wiki/Main_Page
- **Systems:** https://wiki.gacrc.uga.edu/wiki/Systems
- **How to Transfer Files:** https://wiki.gacrc.uga.edu/wiki/Transferring_Files
- **Current Software List:** https://wiki.gacrc.uga.edu/wiki/Software
- **How to Submit a Job:** https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_Sapelo
- **How to Get Help:** http://help.gacrc.uga.edu/
- **GACRC Training Opportunities:** https://wiki.gacrc.uga.edu/wiki/Training
- **How to register a new GACRC New Lab:** http://help.gacrc.uga.edu/labAcct.php
- **How to register a new GACRC User:** http://help.gacrc.uga.edu/userAcct.php