Practical Example: High Performance Computing

Ivan Gesteira Costa, James Nagai, Martin Manolov, Kai Peng, Mehdi Joodaki Institute for Computational Genomics



RWTH Compute Cluster

High-Performance-Computing infrastructure

Linux-based

Highly parallelizable



- All RWTH affiliates are granted access, but with limited resource allocation
- More information on the following link.



RWTH RegApp

 Use RWTH RegApp (<u>https://regapp.itc.rwth-aachen.de/</u>) to create and manage an HPC-account to use the cluster



You have already registered with the following services:

RWTH High Performance Computing

The IT Center hosts one of the fastest supercomputers in Germany. The High Performance Computing group (HPC) supports users from all German universities including institutions from RWTH Aachen University in the efficient use of the central high-performance computing systems.

Registry info	>
Set service password	>
Set SSH Key	→

Computational Genomics 01011011010 10100100101



RWTH RegApp 2FA Tokens or SSH-Key

- Additionally to establishing a password, a user should establish an MFA Token and/or SSH Key. We will proceed in this tutorial to use MFA Tokens (<u>Link</u>).
- After RegApp Login, navigate to Index -> My Tokens
- * Use your configured tokens to authenticate and create a new one for the HPC
- If you don't have any tokens configured, follow the instructions to create a TAN List to use to authenticate once (and don't use it again unless its an emergency)
- Alternatively you can set up an SSH-Key and log in passwordless. (Link)



Server-name	OS	Purpose
lect.hpc.itc.rwth-aachen.de lect2.hpc.itc.rwth-aachen.de	Rocky Linux 8	Dedicated to attend lectures
login23-N.hpc.itc.rwth- aachen.de	Rocky Linux 8	Front-End Dialogue System CLAIX2023 N={1,2,3,4}
login23-x-N.hpc.itc.rwth- aachen.de	Rocky Linux 8	CLAIX2023 Remote Desktop sessions N = {1,2}
login23-g-1.hpc.itc.rwth- aachen.de	Rocky Linux 8	CLAIX2023 GPU-System

* Rocky Linux 9 currently rolling out on most machines.



SSH and Cluster Login

- Use established Secure Shell Protocol (ssh) connection to log in to the front-end cluster nodes with your username and password
- \$ ssh <username>@<server-address>

- In the case of the RWTH cluster:
- \$ ssh <TIM>@login23-N.hpc.itc.rwth-aachen.de



SSH and Cluster Login

• Special nodes for intensive I/O operations

copy23-1.hpc.itc.rwth-aachen.de	2 x 100 Gbit/s
copy23-2.hpc.itc.rwth-aachen.de	2 x 100 Gbit/s

• Use these nodes if you need to transfer large or many files



Cluster Available File Systems

Name	F	Path	Backup	Quota (file)	Quota (#files)
\$HOME	/home/ <tim·< td=""><td>-Kennung></td><td>yes</td><td>250 GB</td><td>1.000.000</td></tim·<>	-Kennung>	yes	250 GB	1.000.000
\$WORK	/work/ <tim-i< td=""><td>Kennung></td><td>no</td><td>250 GB</td><td>1.000.000</td></tim-i<>	Kennung>	no	250 GB	1.000.000
\$HPCWORK	/hpcwork/ <t< td=""><td>IM-Kennung></td><td>no</td><td>1000 GB</td><td>1.000.000</td></t<>	IM-Kennung>	no	1000 GB	1.000.000
\$BEEOND	BeeGFS On	Demand	no	-	-
Source code Configuration f	iles	Output files Working data		IO intensive large files. SSD storag	e job, Provides je pool
				Institute for Computational Genomics 01011011010101010101010101010101010101	WITH AACH UNIVERS

Cluster - mount drive

- Use *sshfs* to mount remote directory to local machine:
 - \$ sshfs <TIM>@copy23-1.hpc.itc.rwth-aachen.de:<Path> <local Path>

- Unmount with:
 - \$ sudo umount -I <local Path>

- Access to cluster only from eduroam or through RWTH VPN *
 - <u>https://help.itc.rwth-aachen.de/en/service/vbf6fx0gom76/article/6a2cfd0933604cd28eaaa69194ff8d16/</u>
 - * VPN also requires 2FA



Cluster - Batch Jobs

- Job handling system SLURM
- SBATCH [option] command [arguments]
- Slots

Parameter	Function
-c,cpus-per-task <numcpus></numcpus>	Number of threads/processes for an OpenMP/ Hybrid script
-n,ntasks <numtasks> ntasks-per-node <numtasks></numtasks></numtasks>	Number of threads/processes for an MPI job
-N,nodes <numnodes></numnodes>	Number of nodes/hosts for the job

- List of slurm commands:
 - https://help.itc.rwth-aachen.de/en/service/rhr4fjjutttf/article/85b21b312bfb48b290043083d2a34b8f/



• Further job parameters

Parameter	Function
-Jjob-name= <jobname></jobname>	A name for the current job
mem-per-cpu= <size></size>	Required RAM per allocated CPU
-o,output= <filename></filename>	Name for a report file, containing the standard output of the job
-t,time=d-hh:mm:ss	Time for job execution, after this time the job is killed
-A,account= <project></project>	Submit a job for a specific project
gres=gpu: <type>:2</type>	Requesting two GPUs per node

Batch documentation - <u>https://slurm.schedmd.com/sbatch.html</u>



Cluster - Batch Jobs

- **sbatch** job.sh sends job, declared in job.sh, to the scheduler
- After defining the slurm parameters, don't forget to write the script you want to execute
- You job might need to load some modules, update environmental variables or activate condo environments.
 - module load Python/3.9.6
 - module avail # Lists all available modules
 - export PATH=~/.local/bin:\$PATH
 - conda activate myenv



Cluster - Example Scripts

#!/bin/bash

Job name #SBATCH --job-name=MYJOB

File for the output
#SBATCH --output=MYJOB_OUTPUT

Time your job needs to execute, e. g. 15 min 30 sec #SBATCH --time=00:15:30

Memory your job needs per node, e. g. 1 GB #SBATCH --mem=1G

The last part consists of regular shell commands: ### Change to working directory cd /home/usr/workingdirectory

Execute your application
myapp.exe



Cluster - Monitor Jobs and Resources

- Use **squeue** to monitor current jobs in progress:
 - squeue -u <TIM> # to only view your jobs
- Use scancel to stop current jobs from executing
 - scancel -u <TIM> # cancels all your jobs
 - scancel -n <JOBID> # cancels job with ID=JOBID
- Use quota to see how much space you're occupying
- Use **r_wIm_usage -q** to view your accounting information
- More info:
 - https://help.itc.rwth-aachen.de/service/rhr4fjjutttf/article/13ace46cfbb84e92a64c1361e0e4c104/



Cluster - GPU Nodes

- Login to a GPU node
- See GPU usage with nvidia-smi

NVIDIA-SMI 565.57.01	Driver	Version: 565.57.01	CUDA Version: 12.7
GPU Name Fan Temp Perf	Persistence-M Pwr:Usage/Cap	Bus-Id Disp.A Memory-Usage 	Volatile Uncorr. ECC GPU-Util Compute M. MIG M.
0 NVIDIA H100 N/A 36C P0	On 69W / 700W	 00000000:1B:00.0 Off 5MiB / 95830MiB 	0 0% Default Disabled
1 NVIDIA H100 N/A 36C P0	On 69W / 700W	0000000:2C:00.0 Off 5MiB / 95830MiB	0 0% Default Disabled
2 NVIDIA H100 N/A 35C P0	On 73W / 700W	00000000:9D:00.0 Off 107MiB / 95830MiB	On N/A E. Process Enabled
3 NVIDIA H100 N/A 37C P0	On 69W / 700W		0 0% E. Process Disabled
MIG devices:		·	+

MIG	devi	ces:									
GPU	GI ID	CI ID	MIG Dev	Memory-Usage BAR1-Usage	 SM 	Vol Unc ECC	CE	ENC	Share DEC	d OFA	JPG
2	7	0	0	16MiB / 11008MiB 0MiB / 16383MiB	+ 16 	0	1	0	1	0	1
2	8	0	1	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
2	9	0	2	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
2	10	0	3	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
2	11	0	4	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
2	12	0	5	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
2	13	0	6	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1

+								
F	roc	esses:						
i	GPU	GI	CI	PID	Туре	Process	name	GPU Memory
İ		ID	ID					Usage
==		=======						
	No	running	processes	found				
L								

Institute for Computational Genomics 01011011010 10100100101



Cluster - GPU Nodes

NVID	IA-SMI 5	65.57.01	1	0	Driver	Version:	565.5	7.01	CUDA Versio	on: 12.7
GPU Fan	Name Temp	Perf	Persi: Pwr:U:	ster sage	nce-M e/Cap	Bus-Id	Mem	Disp.A ory-Usage	Volatile GPU-Util 	Uncorr. ECC Compute M. MIG M.
0 N/A	NVIDIA 36C	H100 P0	69W	/	0n 700W	+======== 000000 5M	===== 00:1B iB /	:00.0 Off 95830MiB	-+====================================	0 Default Disabled
1 N/A	NVIDIA 36C	H100 P0	69W	/	0n 700W	000000 5M	00:2C iB /	:00.0 Off 95830MiB	0%	0 Default Disabled
2 N/A	NVIDIA 35C	H100 P0	73W	/	0n 700W	000000 107M	00:9D iB /	:00.0 Off 95830MiB	 N/A 	On E. Process Enabled
3 N/A	NVIDIA 37C	H100 P0	69W	/	0n 700W	000000 5M	00:AD iB /	:00.0 Off 95830MiB	 0%	0 E. Process Disabled

Main GPU Table, featuring 4 GPUs:

- 0 and 1 in Default mode (multiple processes allowed) with disabled MIG
- 2 in exclusive process mode with MIG Enabled
- 3 in exclusive process mode with disabled MIG

GPU	GI	CI	MIG	Memory-Usage		Vol			Share	d	
	ID	ID	Dev	BAR1-Usage	SM	Unc ECC	CE	ENC	DEC	OFA	JP
2	7	0	0	16MiB / 11008MiB 0MiB / 16383MiB	16	+===== 	1	0	1	0	:
2	8	0	1	16MiB / 11008MiB 0MiB / 16383MiB	16	0	1	0	1	0	
2	9	0	2	16MiB / 11008MiB 0MiB / 16383MiB	16	0	1	0	1	0	
2	10	0	3	16MiB / 11008MiB 0MiB / 16383MiB	16	0	1	0	1	0	
2	11	0	4	16MiB / 11008MiB 0MiB / 16383MiB	16	0	1	0	1	0	
2	12	0	5	16MiB / 11008MiB 0MiB / 16383MiB	16	0	1	0	1	0	
2	13	0	6	16MiB / 11008MiB 0MiB / 16383MiB	16	0 	1	0	1	0	

+ 	Proc GPU	cesses: J GI	CI	PID	Туре	Process	name	GPU Memory	
I		ID	ID					Usage	
Ŀ	=====								=
I	No	running	processes	found					1

Institute for Computational Genomics 01011011010 10100100101



Cluster - GPU Nodes

+										
NVIDI	IA-SMI 5	65.57.01			Driver	Version:	565.5	7.01	CUDA Versi	on: 12.7
GPU Fan 	Name Temp	Perf	Persi: Pwr:U:	ste sag	nce-M e/Cap	Bus-Id	Mem	Disp.A ory-Usage	Volatile GPU-Util 	Uncorr. ECC Compute M. MIG M.
0 N/A 	NVIDIA 36C	H100 P0	69W	/	0n 700W	000000 5N	000:1B 4iB /	:00.0 Off 95830MiB	 0%	0 Default Disabled
1 N/A 	NVIDIA 36C	н100 Р0	69W	/	0n 700W	000000 5N	000:2C 4iB /	:00.0 Off 95830MiB	 0%	0 Default Disabled
2 N/A 	NVIDIA 35C	н100 Р0	73W	/	0n 700W	000000 107N	000:9D 4iB /	:00.0 Off 95830MiB	 N/A 	On E. Process Enabled
3 N/A 	NVIDIA 37C	H100 P0	69W	/	0n 700W	000000 5N	000:AD 4iB /	:00.0 Off 95830MiB	0%	0 E. Process Disabled

+	MIG	devi	ces:									$\overline{}$
+ 	GPU	GI ID	CI ID	MIG Dev	Memory-Usage BAR1-Usage	+ SM 	Vol Unc ECC	CE	ENC	Share DEC	d OFA	JPG
	2	7	0	0	 16MiB / 11008MiB 0MiB / 16383MiB	+===== 16 	0	1	0	1	0	1
+ 	2	8	0	1	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
	2	9	0	2	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
	2	10	0	3	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
I	2	11	0	4	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
	2	12	0	5	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	1
	2	13	0	6	16MiB / 11008MiB 0MiB / 16383MiB	16 	0	1	0	1	0	X

+										-+
Pr	oc	esses:								I
G	ΡU	GI	CI	PID	Туре	Process	name	GPU	Memory	1
Í		ID	ID					Usag	e	Ì
====	==					========			======	:
N	0	running	processes	found						İ

Institute for Computational Genomics 01011011010 10100100101



MIG Table, featuring sections of GPU 2:

Each section has 11 GB VRAM

Cluster - Exercise - GPU with python

- Sing in to the cluster
 - \$ **ssh** <TIM>@login23-g-2.hpc.itc.rwth-aachen.de
- Check if you're successfully in the group
 - \$ groups #should see lect0139 among others
- · Load modules and install some libraries locally
 - \$ module load Python/3.9.6
 - \$ module load cuDNN/8.6.0
 - \$ module load CUDA/11.8
 - \$ export PATH=~/.local/bin:\$PATH
 - \$ pip3 install --user tensorflow pandas scikit-learn
 - \$ pip3 install --user matplotlib tensorrt numba

Cluster - Exercise - GPU with python

- Create a new directory and copy the files needed:
 - \$ cd ~/
 - \$ mkdir BioInfo
 - \$ rsync -rp /home/lect0139/sample BioInfo
 - \$ cd BioInfo/sample
- Edit (vim) your email address and submit job (sbatch run.sh)
- Check status regularly (\$ squeue -u <TIM>)



Practical Example: Linux Basics



The Operating System UNIX

- UNIX Development
 - 60s: MULTICS: MULTiplexed Information and Computing Service
 - Developed by MIT, Bell Labs and General Electrics
 - Written in Programming Language 1 (PL1) and Assembly
 - Thompson and Ritchie later create UNIX, written in C



• command parameter1 parameter2 ...

 man – manual. Displays manuals for linux packages (try man man). Displays useful information about how to use the package.

pwd – print working directory. Displays the current working directory



- *Is* list. Displays the content of a directory
 - Is -I ~/Pictures

-rw-rr	1	martin	martin	101685	Apr 29	13:23	Selection_001.png
-rw-rr	1	martin	martin	39835	Apr 29	14:37	Selection_002.png
-rw-rr	1	martin	martin	76985	Apr 29	14:48	Selection_003.png
-rw-rr	1	martin	martin	27096	May 2	14:04	Selection_004.png
-rw-rr	1	martin	martin	62783	May 3	11:58	Selection_005.png
-rw-rr	1	martin	martin	65201	May 6	16:11	Selection_006.png
drwxrwxr-x	2	martin	martin	4096	Apr 24	11:41	Temp

Access LinksOwners/SizeLastFilenamerightsGroupsModified



• ls - list. Displays the content of a directory

-rw-r--r-- 1 martin martin 101685 Apr 29 13:23 Selection_001.png

- Access rights:
 - Type: (d)irectory, (l)ink, (-) a file
 - Rights: (r)ead, (w)rite, e(x)ecute
 - Sequence: user group anybody
- Owner/Group:
 - Owner first column: *martin*
 - Group second column: martin
- Either can be changed with chmod and chown/chgrp respectively
- Run executable with full path or from directory with "./file"



 cd – change directory. Switches to a new directory, supplied as a parameter

mkdir – make directory. Creates a new directory

rm – remove. Removes a specified file or directory ("-r")

- cp copy. Copy a file or directory ("-r")
- *scp* **s**ecure **copy**. Copy a file to or from a remote source
- *rsync* remote sync. A fast copying tool. Also works for remote copy



• *In* – link. Link a file into a new directory

• *echo* – Prints a string to standard output

 cat – concatenate. The content of a file is printed to standard output

wc – word count. Counts the number of words, rows
 ("-l") pr characters ("-c") in a file.



pipe ("I") - connect commands

 Standard output (">") and input ("<") for a specific program. Output can also concatinate to existing content, without deleting (">>")



Linux – Environment Variables

Variable	Description
PATH	Colon separated list of directories, which will be searched through when entering a name of executable
HOME	The pathname of the home directory.
SHELL	The currently used shell program
USER	The current username

\$ echo \$PATH /usr/local/sbin

\$ export PATH=~/.local/bin:\$PATH
\$ echo \$PATH
/home/martin/.local/bin:/usr/local/sbin



- A further development of Vi
- Open-Source editor for use inside the Terminal
- Open file with "\$ vim filename"
 - Type "i" for insert (edit)
 - ESC to go back to entry menu
 - :q exit without saving
 - :wq save and exit
 - :q! force exit without saving



Further Readings

- Linux tutorials:
 - <u>https://www.tutorialspoint.com/unix/</u>
 - <u>https://ryanstutorials.net/linuxtutorial/</u>

- SLURM batch system
 - <u>https://slurm.schedmd.com/documentation.html</u>
 - <u>https://doc.itc.rwth-aachen.de/display/CC/</u> <u>Using+the+SLURM+Batch+System</u>
 - <u>https://hpc-wiki.info/hpc/SLURM#Jobscript_Examples</u>

