

ilifu Online Training

Session 1: Introduction to Slurm

25 August 2026


Dane Kennedy · ilifu Bioinformatics Support
University of Cape Town

[View on github.io](#)

[printable version](#)

Job Scheduling & Cluster Management Tool

An open-source **job scheduling & cluster management tool**. You submit work from the login node; Slurm allocates compute resources and runs your jobs when they're available.



- Login Node
 - Accessed via ssh (`$ ssh <username>@slurm.ilifu.ac.za`)
 - Submit jobs and basic file management
- Scheduler and Accounting Database
 - Manage jobs, Partitions and Nodes
 - Accounting
- Compute Nodes
 - Where your analysis / simulation runs (inside of a slurm job)
 - Software available via singularity containers or modules

Node / Partition	Use Case	Resources
<code>slurm.ilifu.ac.za</code>	SLURM & bash commands (<code>cd</code> , <code>mkdir</code> , <code>sbatch</code>)	—
<code>compute-001</code> / <code>Devel</code>	Development, debugging, testing, interactive jobs	1 × (32 cores, ~232 GiB)
<code>jupyter-0[01-14]</code> / <code>Jupyter</code>	JupyterLab — interactive development/analysis	14 × (32 cores, ~232 GiB)
<code>compute-0[02-86]</code> / <code>Main</code>	Stable, heavy computation	85 × (32 cores, ~232 GiB)
<code>highmem-00[1-8]</code> / <code>HighMem</code>	High memory jobs	6 × (32 cores, 500 GiB) + 1 × (32 cores, 1 TiB) + 1 × (96 cores, 1.5 TiB)
<code>gpu-00[1-7]</code> / <code>GPU</code>	GPU-accelerated workloads	7 × (32 cores, ~232 GiB + NVIDIA GPUs)

Available Resources

Web Access

<https://jupyter.ilifu.ac.za>

 Home Admin

daneLogout

Launch JupyterLab

Hi dane. Remember to try and choose the smallest profile that fits your task. This helps us to make sure that everyone has access to the resources they need. Please visit the [user documentation](#) to learn more about Jupyter on ilifu. If you have any more questions, please send an email to [ilifu support](#).

The following table shows the job profiles available on the ilifu cluster (as at 2026-02-13 20:53):

Job Profile	Available Jobs
GPU Session (16 cores, 1 GPU)	3
Minimum Session (1 core, dedicated)	49
Small Session (2 cores, dedicated)	24
Medium Session (4 cores, dedicated)	11
Large Session (8 cores, dedicated)	5
Half-Max Session (16 cores, dedicated)	2
Max Session (32 cores, dedicated)	0

Select a job profile:

Development Session - 2 core, 3 GB RAM, shared, 18 hrs idle timeout, max 14 days lifespan

Start

0:00 / 0:47





Remember to end your session: File → Hub Control Panel → Stop My Server.

<https://jupyter.ilifu.ac.za>



CLI Access

```
ssh <username>@slurm.ilifu.ac.za
```

```
dane@dane-latitude-7640:~$ ssh dane@slurm.ilifu.ac.za
```



00:00

Logging into ilifu

SSH Keys

```
1 ssh <username>@slurm.ilifu.ac.za # connect to login node
2 sinfo                             # show partition information
3 squeue                             # show all job information
4 squeue -u $USER                   # show only your jobs
5 sbatch your_job_script.sh         # submit a job
6 scancel                           # cancel a job
7 man sbatch                         # show the manual page for a command
8 sbatch --help                     # show the help page for a command
```

```
1 #!/bin/bash
2
3 module add python/3.11.2
4 python hello_world.py
```

Bash Scripting

```
1 #SBATCH --time 0-03:00:00 # 0 days + 3 hours
2 #SBATCH --mem 3G          # 3 GiB
3 #SBATCH --ntasks 1        # one task
4 #SBATCH --nodes 1         # one node
5 #SBATCH --partition Main
6 #SBATCH --account <your default>
```

[sbatch parameters](#)

```

1  #!/bin/bash
2  #SBATCH --job-name=tutorial2_R_container
3  #SBATCH --time=00-00:01:00
4  #SBATCH --mem=4G
5  #SBATCH --partition=Main
6  #SBATCH --output=R_container-%j.stdout
7  #SBATCH --error=R_container-%j.stderr
8  #SBATCH --mail-user=YOUR_EMAIL_ADDRESS
9  #SBATCH --mail-type=BEGIN,END,FAIL,TIME_LIMIT_80
10 #SBATCH --account=ACCOUNTING_GROUP
11
12 singularity exec /software/common/containers/RStudio2023.06.1-524-R4.3.1.sif Rscript hello_world.R

```

R_container.sbatch

Submit the job:

```

sbatch R_container.sbatch

```

[Slurm sbatch documentation](#)

Partition	Node names	Default CPUs	Max CPUs	Default Memory (GiB)	Max Memory (GiB)	Default wall-time	Max wall-time
Main	compute-[002-021]	1	32	3	232	3 hours	14 days
Main	compute-[101-105]	1	48	3	232	3 hours	14 days
Main	compute-[201-260]	1	32	3	251	3 hours	14 days
HighMem	highmem-[001-002]	1	32	15	503	3 hours	14 days
HighMem	highmem-003	1	96	15	1508	3 hours	14 days
HighMem	highmem-[004-007]	1	32	15	503	3 hours	14 days
HighMem	highmem-008	1	32	15	1007	3 hours	14 days
GPU	gpu-[001-004]	1	32	7	232	3 hours	14 days
GPU	gpu-005	1	24	7	232	3 hours	14 days
GPU	gpu-006	1	48	7	354	3 hours	14 days
GPU	gpu-007	1	48	7	354	3 hours	14 days
GPU	highmem-008	1	32	7	1007	3 hours	14 days
Devel	compute-001	1	32	-	-	3 hours	5 days

Maximum Allocation

[minimal.sbatch](#)

[maximal_python.sbatch](#)

[hello_world.py](#)

[maximal_R.sbatch](#)

[hello_world.R](#)

Remember that the ilifu cluster is a shared resource with limited resources. It is important to be mindful of this when requesting resources.

✓ Do


- Use `sbatch` for production work
- Identify resource requirements first:
 - Number of nodes
 - Cores per node
 - Ram per core/node
 - Expected runtime
- Clean up temp files (e.g., `/scratch3`)
- Use Singularity containers for software
- Use `transfer.ilifu.ac.za` for data transfers

✗ Don't

- Run heavy processes on the login node
- Place large files in `/users` (home dir)
- Use `scp` / `rsync` on login node
- Request way more resources than needed

Thank you for your time!

We hope this presentation was helpful.



Contact Support

support@ilifu.ac.za



Documentation

docs.ilifu.ac.za