

Reference

Quick reference guides, cheat sheets, and troubleshooting for TAU HPC

- [Common Slurm Commands](#)
- [PBS to Slurm Migration](#)
- [Troubleshooting](#)
- [Security Installations](#)

Common Slurm Commands

Quick reference for the most common Slurm commands on the TAU HPC cluster.

Submitting Jobs

Command	Description
<code>sbatch job.sh</code>	Submit a batch job script
<code>srun --pty bash</code>	Start an interactive session
<code>sbatch --depend=afterok:JOBID job.sh</code>	Submit job after another completes

Monitoring Jobs

Command	Description
<code>squeue -u username</code>	Your running and pending jobs
<code>squeue</code>	All jobs on the cluster
<code>scontrol show job JOBID</code>	Full details of a job
<code>sacct -j JOBID --format=JobID,JobName,State,MaxRSS,Elapsed</code>	Job accounting and memory usage
<code>sattach JOBID</code>	Attach to a running job's output

Managing Jobs

Command	Description
<code>scancel JOBID</code>	Cancel a specific job
<code>scancel -u username</code>	Cancel all your jobs

Cluster Information

Command	Description
<code>sinfo</code>	Partition and node status
<code>scontrol show partition PARTITION</code>	Partition details and limits
<code>check_my_partitions</code>	Your available partitions and accounts
<code>features</code>	Available node constraints/features

Environment Modules

Command	Description
<code>module avail</code>	List all available modules
<code>module avail NAME</code>	Search for a specific module
<code>module spider NAME</code>	Detailed module info including dependencies
<code>module load NAME</code>	Load a module
<code>module list</code>	List loaded modules
<code>module unload NAME</code>	Unload a module
<code>module purge</code>	Unload all modules

Common SBATCH Directives

Directive	Description
<code>#SBATCH --job-name=NAME</code>	Job name
<code>#SBATCH --account=ACCOUNT</code>	Account name
<code>#SBATCH --partition=PARTITION</code>	Partition/queue
<code>#SBATCH --qos=QOS</code>	Quality of Service
<code>#SBATCH --time=HH:MM:SS</code>	Max run time
<code>#SBATCH --ntasks=N</code>	Number of tasks
<code>#SBATCH --nodes=N</code>	Number of nodes
<code>#SBATCH --cpus-per-task=N</code>	CPU cores per task
<code>#SBATCH --mem-per-cpu=NG</code>	Memory per CPU
<code>#SBATCH --mem=NG</code>	Total memory
<code>#SBATCH --gres=gpu:N</code>	Number of GPUs

Directive	Description
#SBATCH --constraint=FEATURE	Node constraint/feature
#SBATCH --array=1-N	Job array
#SBATCH --output=FILE_%j.out	Output file (%j = job ID)
#SBATCH --error=FILE_%j.err	Error file
#SBATCH --mail-user=EMAIL	Notification email
#SBATCH --mail-type=END,FAIL	When to notify

PBS to Slurm Migration

Quick reference for users migrating from the old PBS/Torque system to Slurm.

Command Equivalents

PBS/Torque	Slurm	Description
<code>qsub job.sh</code>	<code>sbatch job.sh</code>	Submit a job
<code>qsub -I</code>	<code>srunk --pty bash</code>	Interactive session
<code>qstat</code>	<code>squeue</code>	View jobs
<code>qstat -u username</code>	<code>squeue -u username</code>	Your jobs
<code>qdel JOBID</code>	<code>scancel JOBID</code>	Cancel a job
<code>pbsnodes</code>	<code>sinfo</code>	Node status
<code>qstat -f JOBID</code>	<code>scontrol show job JOBID</code>	Job details

Directive Equivalents

PBS/Torque	Slurm	Description
<code>#PBS -N name</code>	<code>#SBATCH --job-name=name</code>	Job name
<code>#PBS -q queue</code>	<code>#SBATCH --partition=partition</code>	Queue/partition
<code>#PBS -l nodes=1:ppn=8</code>	<code>#SBATCH --nodes=1 --cpus-per-task=8</code>	Nodes and cores
<code>#PBS -l mem=4gb</code>	<code>#SBATCH --mem=4G</code>	Memory
<code>#PBS -l walltime=02:00:00</code>	<code>#SBATCH --time=02:00:00</code>	Wall time
<code>#PBS -o output.log</code>	<code>#SBATCH --output=output.log</code>	Output file
<code>#PBS -e error.log</code>	<code>#SBATCH --error=error.log</code>	Error file
<code>#PBS -M email</code>	<code>#SBATCH --mail-user=email</code>	Email
<code>#PBS -m abe</code>	<code>#SBATCH --mail-type=ALL</code>	Mail events
<code>#PBS -V</code>	<code>#SBATCH --export=ALL</code>	Export environment
<code>#PBS -t 1-10</code>	<code>#SBATCH --array=1-10</code>	Job array

Environment Variables

PBS/Torque	Slurm	Description
<code>\$PBS_JOBID</code>	<code>\$SLURM_JOB_ID</code>	Job ID
<code>\$PBS_JOBNAME</code>	<code>\$SLURM_JOB_NAME</code>	Job name
<code>\$PBS_NODEFILE</code>	<code>\$SLURM_JOB_NODELIST</code>	Node list
<code>\$PBS_ARRAYID</code>	<code>\$SLURM_ARRAY_TASK_ID</code>	Array task ID
<code>\$PBS_NP</code>	<code>\$SLURM_NTASKS</code>	Number of tasks
<code>\$PBS_O_WORKDIR</code>	<code>\$SLURM_SUBMIT_DIR</code>	Submission directory

Key Differences

- Slurm requires `--account` and `--qos` — run `check_my_partitions` to find yours
- Slurm uses `--cpus-per-task` instead of `ppn`
- Memory in Slurm is per-CPU (`--mem-per-cpu`) or total (`--mem`)
- Slurm does not automatically change to the submission directory — add `cd` `$SLURM_SUBMIT_DIR` to your script if needed

Troubleshooting

Common errors and solutions for job submission and cluster usage.

Job Submission Errors

No partition specified

```
srun: error: Unable to allocate resources: No partition specified or system default partition
```

Always specify a partition. Run `check_my_partitions` to find yours.

Invalid account or partition

```
sbatch: error: Batch job submission failed: Invalid account or account/partition combination specified
```

Your account and partition combination is incorrect. Run `check_my_partitions` and make sure both match.

QOS not permitted

```
sbatch: error: Batch job submission failed: Job violates accounting/QOS policy
```

The QOS you specified doesn't match your account/partition. Run `check_my_partitions` to see valid combinations.

Job Failures

Out of Memory (OOM)

```
sacct -j JOBID -o JobID,JobName,State%20
```

```
JobID      JobName          State
-----
71         my_job          OUT_OF_MEMORY
```

Your job used more RAM than allocated. Resubmit with a higher `--mem` or `--mem-per-cpu`. To estimate needed memory:

```
sacct -j JOBID --format=JobID,JobName,MaxRSS,Elapsed
```

Timeout

Job state shows `TIMEOUT` — your job exceeded the time limit. Resubmit with a longer `--time` value.

Job stuck in Pending (PD)

Check the reason:

```
squeue -u username -o "%.18i %.9P %.8j %.8u %.2t %.10M %.6D %R"
```

Common reasons:

- **Resources** — cluster is busy, wait for nodes to free up
- **QOSMaxCpuPerUserLimit** — you've hit your CPU quota, wait for running jobs to finish
- **InvalidQOS** — wrong QOS, check `check_my_partitions`
- **ReqNodeNotAvail** — requested node is down, remove `--nodelist` constraint

NFS / Storage Issues

Job hangs or freezes on file operations

May indicate an NFS mount issue. Check if your home directory is accessible:

```
ls ~
```

If it hangs, contact HPC support — do not kill the job manually as it may cause further issues.

Disk quota exceeded

```
bash: cannot create temp file: Disk quota exceeded
```

Your home directory is full. Move large files to scratch space or contact HPC support for a quota increase.

Module Issues

Module not found

```
module avail MODULE_NAME
```

Check the exact module name. Use `module spider MODULE_NAME` for a broader search including partial matches.

Getting Help

If you can't resolve an issue, contact HPC support at hpc@tauex.tau.ac.il. Include:

- Your username
- Job ID
- The command you ran
- The full error message

Security Installations

Required security software for TAU workstations and servers.

NAC — Forescout

Network Access Control client required for connecting to the TAU network.

[ForeScoutSecureConnector_64_visible_daemon.tar.gz](#)

Installation

```
tar -zxvf ForeScoutSecureConnector_64_visible_daemon.tar.gz
cd secure_connector
./install.sh
systemctl start SecureConnector.service
```

EDR — CrowdStrike Falcon

Endpoint Detection and Response client.

- [falcon-sensor_7.18.0-17106_amd64.deb](#)— Ubuntu 22.04 / 24.04
- [falcon-sensor-7.17.0-17005.el9.x86_64.rpm](#)— Rocky / RHEL

Installation — Ubuntu

```
dpkg -i falcon-sensor_7.18.0-17106_amd64.deb
systemctl restart falcon-sensor.service
```

Installation — Rocky / RHEL

```
rpm -ivh falcon-sensor_7.17.0-17005.el9.x86_64.rpm  
systemctl restart falcon-sensor.service
```

Registration (new installations only)

```
/opt/CrowdStrike/falconctl -s --cid=<cid-code>
```

To obtain the CID code, contact infosec@tauex.tau.ac.il.