

Standard Operating Procedure (SOP)

INDUS HPC – Access and Job Submission Guide

Field	Details
Document Version	1.0
Cluster Environment	Rocky Linux 8.7
Scheduler	Slurm Workload Manager
Master Node	172.30.3.10
Fallback Master Node	172.30.3.11

1. Introduction

This document provides a complete Standard Operating Procedure (SOP) for users to access and utilize the IIT HPC cluster. It covers login procedures, the environment module system, directory structure, storage quotas, and job submission and monitoring through Slurm.

The cluster provides:

- Rocky Linux 8.7 operating environment
- Slurm-based job scheduling
- A module system for pre-installed compilers, libraries, and applications
- Shared scratch and home storage

All workloads must be submitted through Slurm. Users must connect to the campus network before accessing the cluster.

2. Login Instructions

2.1 Network Requirement

Users must first connect to the campus Wi-Fi / network before attempting to log in to the cluster.

2.2 Linux / macOS

Open a terminal and connect using SSH:

```
ssh username@172.30.3.10
```

Enter your password when prompted to complete the login.

2.3 Windows

Users may log in using one of the following methods:

- MobaXterm (Portable)
- Windows Terminal

- PowerShell

Using MobaXterm Portable:

- Download [MobaXterm_Portable_v24.3.zip](#) and extract it.
- Double-click the application icon and allow the requested permissions.
- Click the Session icon (top-left corner), choose the SSH option, and enter the server IP address.
- Click OK, then enter your password when prompted to log in.

Using Windows Terminal / PowerShell:

```
ssh username@172.30.3.10
```

Enter your password when prompted to complete the login.

2.4 macOS (Terminal)

Use the built-in Terminal application with your assigned username and password:

```
ssh username@172.30.3.10
```

3. Module System

The cluster uses environment modules to manage compilers, libraries, and applications. Use the following commands to locate, list, and load software modules.

Command	Purpose
module use /opt/modulefiles	Adds the directory the module system searches for modulefiles.
module avail	Lists all modules currently available to load.
module load <modulename>	Loads the specified module/application into the environment.
<modulename> --version	Checks the version of a loaded application.

Example output of module avail:

```
[root@master01 ~]# module use /opt/modulefiles/
[root@master01 ~]# module avail
```

/opt/modulefiles

advisor/2025.4	dal/latest	glibc/2.34	mk1/2025.2	openmpi2/4.1.8
advisor/latest	debugger/2025.2.0	gromacs/2025.3	mk1/2025.3	phonopy/2.12
anaconda/2025	debugger/2025.3.0	gromacs/2025.3	openmpi/2025.3	mk1/latest
ansys/2025R1	debugger/latest	hdf5/1.12.2	molpro/2015.1	python/3.12.4
ansys/emconfig	dev-utilities/2025.2.0	intel/2021.3	molscat/1.0	python/python3
ansys/lsdyna	dev-utilities/2025.3.0	intel_ipp_intel64/2022.2	molscat/mpi	qe/7.0
ansys/mapdl/2025R1	dev-utilities/latest	intel_ipp_intel64/2022.3	mpi/2021.7	qe/7.4.1

ansys/rsm/2025R1	dnnl/3.8.1	intel_ipp_intel64/latest	mpi/2021.16	qe/7.5
ansys/workbench/2025R1	dnnl/3.9.1	intel_ippcp_intel64/2025.2	mpi/2021.17	quantumatk/2026.03
ccl/2021.16.1	dnnl/latest	intel_ippcp_intel64/2025.3	mpich/3.0.4	R/4.6.0
ccl/2021.17.0	dpct/2025.2.0	intel_ippcp_intel64/latest	namd/2.14	scalapack/2.1.0
ccl/latest	dpct/2025.3.0	ishmem/1.3.0	namd/3.0.1	siesta/5.4.1
cmake/3.28.2	dpct/latest	ishmem/1.5.0	namd/3.0.2	siesta/transiesta
compiler-intel-llvm/2025.2.1	dpl/2022.9	ishmem/latest	namd/3.0.2-g++	tbb/2022.2
compiler-intel-llvm/2025.3.0	dpl/2022.10	jasper/1.900.1	netcdf-fortran/4.4.2	tbb/2022.3
compiler-intel-llvm/latest	dpl/latest	jpeg-9e/9e	netcdf-fortran/4.6.1	tbb/latest
compiler-rt/2025.2.1	expresso/7.1	laamps/22jul2025	netcdf/4.5.0	umf/0.11.0
compiler-rt/2025.3.0	expresso2/7.0	lammps/22july2025	netcdf/4.9.2	umf/1.0.2
compiler-rt/latest	expresso3/7.5	lammps/2025	openblas/0.3.27	umf/latest
compiler/2025.2.1	fftw/3.3.10	lapack/3.12.0	openfoam/2312	
compiler/2025.3.0	fluent/2025R1	libgridxc/libgridxc-openmpi	openfoam/bak	
compiler/latest	g16/16	libpng/1.2.50	openmpi-intel/5.0.8-intel	vaspkit/1.5.0
comsol/6.3	g16/16-linda	libtool/2.5.4	openmpi/2.1.1	vtune/2025.7
comsol/6.4	gaussian/1.0	libxc/6.2.2	openmpi/3.1.6	vtune/latest
dal/2025.8	gcc/7.3.0	mathematica/wolfarm/14.2	openmpi/4.1.8	xcrysden/1.6.2
dal/2025.9	gcc/12.2.0	matlab/2024a	openmpi/5.0.8	zlib/1.2.11

/usr/share/Modules/modulefiles

ccl/2021.16.1	dal/2025.8	dot	intel_ippcp_intel64/2025.2	module-info	tbb/2022.2
compiler-intel-llvm/2025.2.1	debugger/2025.2.0	dpct/2025.2.0	ishmem/1.3.0	modules	umf/0.11.0
compiler-rt/2025.2.1	dev-utilities/2025.2.0	dpl/2022.9	mkl/2025.2	mpi/2021.16	use.own
compiler/2025.2.1	dnnl/3.8.1	intel_ipp_intel64/2022.2	module-git	null	

/usr/share/modulefiles

open/new-mpi	pmi/pmix-x86_64
--------------	-----------------

Important: Keep the terminal session open for as long as you are accessing the HPC cluster.

4. Directory Structure and Storage Quota

4.1 Directory Paths

Directory	Path
User Home Directory	/scratch/home/<username>
User Data Directory	/scratch/user_data/<department_name>/<username>
Operational Scripts	/scratch/software/scripts

Directory	Path
Job Instructions Reference	/scratch/software/scripts/job-instructions.txt

Users should read the job instructions file carefully before submitting jobs, as it contains cluster-specific guidance maintained by the administration team.

4.2 Storage Quota

Location	Default Quota
Home Directory	50 GB
Scratch Directory	500 GB

5. Submitting Jobs

All workloads must be submitted through Slurm using a job script. The template below (based on a GROMACS molecular dynamics job) illustrates the standard structure used on the cluster.

```
#####
# Template slurm script to submit a job on HPC
#####
#!/bin/bash
#SBATCH --job-name=gmx_md          # Job name, appears in queue
#SBATCH --output=gmx_md.out        # Standard output file
#SBATCH --error=gmx_md.err         # Standard error file
#SBATCH --partition=small          # Partition: small (default) or highmemory
#SBATCH --qos=small                # QOS: small, medium, highmemory
#SBATCH --account=small_acct       # small: max 1 node, 10 days
#                                   # medium: max 2 nodes, 4 days
#SBATCH --nodes=1                 # Number of nodes to allocate
#SBATCH --ntasks-per-node=64       # MPI tasks per node
#SBATCH --time=48:00:00            # Job runtime (HH:MM:SS)

# Load modules
module use /opt/modulefiles/
module purge                       # Clear all loaded modules
module load gcc/12.2.0
module load openmpi/5.0.8
module load gromacs/2025.3-openmpi

# Add FFTW path (if not picked up automatically)
export LD_LIBRARY_PATH=/opt/fftw-3.3.10/lib:$LD_LIBRARY_PATH
export PATH=/opt/fftw-3.3.10/bin:$PATH
export OMPI_MCA_coll_hcoll_enable=0

# GROMACS env
export OMP_NUM_THREADS=1
```

```
export GMX_MAXCONSTRWARN=-1

# Check GROMACS executable
which gmx_mpi
gmx_mpi --version

# Run multi-node simulation
srun --mpi=pmix gmx_mpi mdrun -deffnm md_0_10
```

Submit the job script using:

```
sbatch <scriptname>
```

6. QOS / Partition System

The cluster is organized into three partitions, each with its own resource limits and maximum runtime.

Partition	Max Nodes	Max CPU Cores	Max Jobs/User	Max Runtime
small	1	64	2	10 days
medium	2	128	1	4 days
highmemory	2	128	1	1 day

6.1 Partition Usage Directives

Small partition:

```
#SBATCH --partition=small
#SBATCH --qos=small
#SBATCH --account=small_acct
```

Medium partition:

```
#SBATCH --partition=medium
#SBATCH --qos=medium
#SBATCH --account=medium_acct
```

Highmemory partition:

```
#SBATCH --partition=highmemory
#SBATCH --qos=highmemory
#SBATCH --account=highmemory_acct
```

7. Monitoring Cluster and Job Status

7.1 View Cluster Status

```
sinfo -l
```

Example output:

```
(base) [root@master01 ~]# sinfo
PARTITION  AVAIL  TIMELIMIT  NODES  STATE  NODELIST
small*     up     10-00:00:0  4      mix    cn[04-05,07,19]
small*     up     10-00:00:0  13     alloc  cn[01-03,06,08-16]
small*     up     10-00:00:0  2      idle   cn[17-18]
medium     up     4-00:00:00  4      mix    cn[04-05,07,19]
medium     up     4-00:00:00  13     alloc  cn[01-03,06,08-16]
medium     up     4-00:00:00  2      idle   cn[17-18]
highmemory up     1-00:00:00  2      idle   hm[01-02]
```

7.2 View Job Queue

```
squeue
```

7.3 Submit a Job

```
sbatch <scriptname>
```

8. Master Node Access – Failure Scenario

In case the primary master node is not accessible, or users are unable to log in to the master node, users must connect to the fallback master node using SSH:

```
ssh <username>@172.30.3.11
```

9. Useful Commands Reference

Task	Command
Add module search path	module use /opt/modulefiles
List available modules	module avail
Load a module	module load <modulename>
Check application version	<modulename> --version
Submit job	sbatch <scriptname>
Check job queue	squeue
View detailed partition/node list	sinfo -l

10. Support

For issues related to login, storage quota, module availability, or Slurm job submission, please contact the HPC Administration Team and reference the job instructions file at `/scratch/software/scripts/job-instructions.txt` for cluster-specific guidance.