User Guide for the ITP Linux Cluster

1. Introduction

The ITP Linux Cluster provides computational resources for numerical simulations at the Institute for Theoretical Physics (ITP). It consists of 20+ compute nodes and one central server for user authentication, shared storage, and essential network services.

Users are expected to share resources responsibly, monitor node availability via Ganglia, and follow best practices for data storage and fair use.

Content of this guide

1. Introduction
2. Getting an account
3. Accessing the cluster
4. Running programs
5. Organization and storage
6. Services
7. Nodes
8. Support

---

2. Getting an account

Accounts on the ITP Linux Cluster can be created for ITP students and researchers. To request an account:

1. Find the official user <u>authorization</u> ("Benutzungsbewilligung") that was sent to you by ZID, which provides the license to use the ZID services along with registration data (including username, initial password, and email address).
2. Share a copy of this document with your system admin (see Contacts).

User authentication is managed via the Network Information Service (NIS) hosted on Mungo. The same credentials (c-number and password) allow access to all nodes of the cluster. The initial password is printed on your "*Benutzungsbewilligung*".

To change your password: (1) log into the central server (node name: Mungo; IP: 138.232.67.8) and run *yppasswd*, (2) insert your current password, then (3) insert your new password (twice, with no visual feedback). Password updates automatically propagate to all cluster nodes (see Sec. 7.2) and do not affect other ZID services.

---

3. Accessing the cluster

For security reasons, the ITP Linux Cluster is accessible only from UIBK's IP addresses via SSH. To access it from other networks you need a VPN such as Cisco Secure Client (see ZID's guide).

3.1 Connecting using a VPN

Cisco Secure Client is easy to set up and use, as described in the above guide by ZID. You can download a free version of the client from UIBK's software download area.
Note that the VPN client will route traffic through the UIBK network, so you will appear to be connecting from UIBK and may experience performance degradation with external services.

3.2 Connecting via SSH

Both Linux/maxOS and Windows users can access the ITP Linux Cluster via SSH.

- **3.2.1 Connecting from Linux/macOS**
  In Linux, log into the node of your choice (see Sec. 7.2) by running the following command in your local shell: ssh username@<node>.uibk.ac.at

- **3.2.2 Connecting from Windows**
  You need a terminal emulator that supports the SSH protocol. We recommend installing PuTTY from UIBK's free software download area.
  To connect to the cluster, (1) open PuTTY's graphical interface, (2) enter the IP or fully qualified domain name (<node>.uibk.ac.at) of the desired login node into the "Host Name" field, then (3) click "Open" to get access to a new Terminal window where (4) you are asked to enter your username and password.

---

4. Running programs

4.1 Software modules

Software on the Linux Cluster is managed using Environment Modules in the two-level structure <software>/<version>. Software is periodically updated: To streamline the system, up to three versions per software are made available at the same time. New software can be installed upon request (see Sec. Contacts).

Below is a list of commands to navigate the environment modules:

Goal	Command
List all available modules	*module avail*
See what a module does to your env	*module display <software>/<version>*
Load a module	*module load <software>/<version>*
Unload a module	*module unload <software>/<version>*
List all currently loaded module	*module list*
Unload all modules at once	*module purge*

4.2 Choosing a node

To choose the node(s) where you will launch your scripts:

1. Visit **Ganglia's web interface** at http://mungo.uibk.ac.at/ganglia/ from within UIBK's network or VPN (see Sec. 3.1).
2. Identify the node(s) with low CPU and RAM usage (blue or green panels).
   Note: Nodes have different hardware (see Sec. 7.2) but the same configuration.
3. Connect to the chosen node(s) via SSH (see Sec. 3.2).

4.3 Resource usage

- To ensure a smooth and resilient operation of the whole cluster, **running computational tasks on Mungo is not allowed**. Any violation is recorded and results in the immediate interruption of the process.
- **Use free nodes** (see Secs. 4.2 and 7.2) whenever possible.
- **Monitor your tasks**' usage and status with *top*, *htop,* or *free -m*.

5. Organization and storage

Files and directories in the ITP Linux Cluster can be either local (on compute nodes) or shared (from the central server Mungo). Understanding where you and your files are is essential to ensure the best operation of the cluster.

5.1 Filesystem overview

/home – All home directories */home/<username>* reside on Mungo and are shared with each compute node via NFS. This means that when you log into a compute node and access your home directory, you are actually working with files stored on Mungo. Since Mungo also manages other critical services, it must not be used for running any computational tasks (see Sec. 4). Files in */home/<username>/backup* on Mungo are automatically backupped daily by a centrally-scheduled TSM service. Home directories have user quotas to ensure stability and fair use (see

Sec. 5.2). While your home directory is suitable for scripts and personal files, disk space is limited: For large or temporary files, please use the scratch folders described below.

/scratch – Shared scratch folder for temporary storage. The /scratch directory resides on Mungo and is available on all nodes at the same path /scratch. It provides a large, centralized location for storing and accessing data across the cluster. Since this /scratch folder is NFS-mounted over the network, any file you save in your /scratch/<username> directory is actually written to Mungo (without visual indication). For this reason, writing or reading large files from /scratch can be slower, especially when many users are active at once. To reduce load and avoid possible issues, consider using the additional local scratch space described below.

/scratch-local – Local scratch folder on each compute node. This directory is stored on the local disk of each node and offers faster access for I/O-intensive tasks. We recommend using /scratch-local when your scripts generate or access large files frequently, or when performance is critical. Note that this space is node-specific: files saved here are not visible from other nodes. See Sec. 7.2 for details on the compute nodes and their /scratch-local folder.

A summary of the above filesystems is shown in this table:

Mount point	Purpose	Backup policy
/home	Login directory	Daily backup for */home/<username>/backup*
/scratch	Shared storage	No automatic backup
/scratch-local	Node-specific storage	No automatic backup

5.2 Quotas

Quotas are active on Mungo (20 GB/user). To check your disk quota, run *quota* at the command line. The *quota(1)* command displays the current disk usage along with your limits for disk space (blocks) and number of inodes (files). The soft limit (quota) can be temporarily exceeded for a grace period of 7 days, while the hard limit (limit) is an absolute upper bound.

If you run out of disk space, try to (i) remove unused dvi, aux, log files; (ii) clear your browser's cache; (iii) tar and gzip directories; (iv) avoid keeping huge data files if possible;

When your quota is exceeded, login to any node is denied after the grace time. If this happens: (1) login to a node with a different account, (2) switch account (*su <username>*), then (3) follow the above tips to shrink your disk usage.

5.3 Best storage practices

- Store critical files in your personal /home/<username>/backup.
- Use /scratch/<username> to store large data, periodically cleaning up old files.
- Use /scratch-local/<username> for fast, temporary storage to reduce network load.

6. Services

6.1 Extra services

Below is a list of relevant services running on the cluster, in addition to the software available for numerical simulations.

Service	Purpose
Ganglia	Monitor cluster resource usage (link)
Tivoli Storage Manager	Centralized backup system
Nagios	System health monitoring
Superdoctor	Hardware monitoring tool
Quotas	Stability and fair use
GitLab CE	ITP instance of GitLab for collaborative projects
Overleaf CE	ITP instance of Overleaf for the online LaTeX editor

Please contact the system admin to request access to GitLab CE and Overleaf CE.

6.2 Using graphical applications

Note: You can very likely skip the whole Sec. 6.2!

While the ITP Linux Cluster does not have a graphical desktop environment, some graphical applications can still be used via X11 forwarding.

6.2.1 Using GUI from Linux/macOS
To enable X11 forwarding, necessary to run programs with graphical interfaces, add the -X option (uppercase X) to the command line above: ssh username@<node>.uibk.ac.at

6.2.2 Using GUI from Windows
You need an X-Windows server installed on your local Windows machine to open graphical interfaces for applications running on the cluster. We recommend using Xming.

o **Install Xming as local X server:** The XMing X Server and Xming fonts can be obtained from UIBK's software download area. We recommend to avoid an outdated, defective Public Domain release of Xming at sourceforge.net. Here you can find details to create a shortcut.

o **Enabling X11 forwarding with PuTTY:** X Windows applications only work if X11 forwarding via SSH was enabled and the X server is running. To configure it in PuTTY proceed as follows:

1. Select the 'Connection → SSH → X11' panel from the left-hand selection menu.
2. Enter 127.0.0.1:0 into the 'X display location' dialog box (display number 0 is typical).
3. Check the 'enable X11 forwarding' check box.
4. Open a SSH session and run "xclock" to test your X11 enabled connection.

7. Nodes

The ITP Linux cluster consists of a central server (Mungo) and 20+ compute nodes. The OS of all nodes is Ubuntu-Server 24.04.3 LTS, the latest version as of September 2025.

7.1 Mungo (central server)

Mungo (mungo.uibk.ac.at) is the central server of the cluster. It provides network services for user authentication, shared storage, and software modules. RAID (Redundant Array of Independent Disks) protects data on Mungo by distributing it across multiple drives.

Login to Mungo is only necessary for password updates.

Below is a list of specification for the server.

IP:	138.232.67.8
CPU	2x Intel Quad-Core Xeon 5345, 2.33 GHz (Clovertown)
RAM	32 GB
Disk size	29 T (RAID 6 protection)

7.2 Compute nodes

Compute nodes have the same OS but different specifications (see table below).

Name	IP	Processor	Cache	RAM	Scratch-local
bozon	138.232.67.5	4 x Eight-Core Opteron 6128, 2.0 GHz (Magny-Cours)	8 x 512 kB L2, 12288 kB L3	64 GB	580 GB
dopey	138.232.67.15	2 x Intel Xeon 8-Core E5-2640V3, 2.6 GHz (Haswell)	8 x 1024 kB L2, 20 MB L3	64 GB	806 GB
electron	138.232.67.24	4 x Twelve-Core Opteron 6174, 2.2 GHz (Magny-Cours)	12 x 512 kB L2, 12288 kB L3	64 GB	126 GB
fermion	138.232.67.4	2 x Intel Xeon 8-Core E5-2640V3, 2.6 GHz (Haswell)	8 x 1024 kB L2, 20 MB L3	64 GB	806 GB
gnu	138.232.67.7	4 x Twelve-Core Opteron 6174, 2.2 GHz (Magny-Cours)	12 x 512 kB L2, 12288 kB L3	64 GB	359 GB
graviton	138.232.67.6	4 x Quad-Core Opteron 8380, 2.5 GHz (Shanghai)	4 x 512 kB L2, 6144 kB L3	128 GB	806 GB
grunt	138.232.67.33	2 x Intel Xeon 8-Core E5-2640V3, 2.6 GHz (Haswell)	8 x 1024 kB L2, 20 MB L3	64 GB	807 GB
hbar	138.232.67.25	4 x Six-Core Opteron 8431, 2.4 GHz (Istanbul)	6 x 512 kB L2, 6144 kB L3	64 GB	807 GB
ito	138.232.67.19	4 x Twelve-Core Opteron 6174, 2.2 GHz (Magny-Cours)	12 x 512 kB L2, 12288 kB L3	64 GB	807 GB
mcavity	138.232.67.14	4 x Quad-Core Opteron 8380, 2.5 GHz (Shanghai)	4 x 512 kB L2, 6144 kB L3	64 GB	807 GB
muon	138.232.67.37	4 x Twelve-Core Opteron 6174, 2.2 GHz (Magny-Cours)	12 x 512 kB L2, 12288 kB L3	64 GB	71 GB
neutrino	138.232.67.35	4 x Sixteen-Core Opteron 6274, 2.2 GHz (Interlagos)	6 x 512 kB L2, 6144 kB L3	512 GB	807 GB
noise	138.232.67.26	2 x Intel Xeon 8-Core E5-2640V3, 2.6 GHz (Haswell)	8 x 1024 kB L2, 20 MB L3	64 GB	806 GB
obiwan	138.232.67.44	Intel Quad-Core Xeon X3360, 2.83 GHz (Yorkfield)	2 x 6144 kB L2	4 GB	353 GB
panic	138.232.67.12	4 x Sixteen-Core Opteron 6274, 2.2 GHz (Interlagos)	16 x 1024 kB L2, 16384 kB L3	256 GB	71 GB

Name	IP	Processor	Cache	RAM	Scratch-local
photon	138.232.67.1	4 x Quad-Core Opteron 8380, 2.5 GHz (Shanghai)	4 x 512 kB L2, 6144 kB L3	64 GB	807 GB
proton	138.232.67.28	4 x Quad-Core Opteron 8380, 2.5 GHz (Shanghai)	4 x 512 kB L2, 6144 kB L3	64 GB	807 GB
qbit	138.232.67.47	2 x Intel Xeon 8-Core E5-2640V3, 2.6 GHz (Haswell)	8 x 1024 kB L2, 20 MB L3	64 GB	806 GB
quantum	138.232.67.3	2 x Intel Xeon 8-Core E5-2640V3, 2.6 GHz (Haswell)	8 x 1024 kB L2, 20 MB L3	64 GB	806 GB
sk-c705	138.232.67.48	Intel Core i7-3930K, 3.2 GHz (Sandy Bridge-E)	6 x 256 kB L2, 12 MB L3	32 GB	353 GB
vacuum	138.232.67.21	4 x Quad-Core Opteron 8380, 2.5 GHz (Shanghai)	4 x 512 kB L2, 6144 kB L3	64 GB	807 GB
vortex	138.232.67.38	4 x Sixteen-Core Opteron 6274, 2.2 GHz (Interlagos)	16 x 1024 kB L2, 16384 kB L3	128 GB	353 GB
yedi	138.232.67.51	4 x Sixteen-Core Opteron 6274, 2.2 GHz (Interlagos)	16 x 1024 kB L2, 16384 kB L3	256 GB	807 GB

The following nodes are currently offline:

Name	IP	Processor	Cache	RAM	Scratch-local
axion	138.232.67.5	4 x Eight-Core Opteron 6128, 2.0 GHz (Magny-Cours)	8 x 512 kB L2, 12288 kB L3	64 GB	
neutron	138.232.67.29	4 x Six-Core Opteron 8431, 2.4 GHz (Istanbul)	6 x 512 kB L2 6144 kB L3	64 GB	
quaxo	138.232.67.10	4 x Quad-Core Opteron 8380, 2.5 GHz (Shanghai)	4 x 512 kB L2, 6144 kB L3	64 GB	
positron	138.232.67.58	4 x Twelve-Core Opteron 6174, 2.2 GHz (Magny-Cours)	12 x 512 kB L2 12288 kB L3	64 GB	

This is how the filesystems are organized on each compute node (command: *df*). Note how the home directories, scratch, and software are actually shared from Mungo.

	Filesystem	Size	Accessible at
local	/dev/mapper/ubuntu--vg-ubuntu--lv		/
	/dev/mapper/ubuntu--vg-scratch--lv	Node-specific	/scratch-local
network	mungo:/home	10 TB	/home
	mungo:/scratch	18 TB	/scratch
	mungo:/opt/local/cluster	1 TB	/net/apps64

8. Support

Getting started with Linux – Internet is full of great resources for Linux, but it's easy to get lost in the process! Tools are also rapidly changing and books are being complemented by powerful platforms and chatbots. If you look for a good place to start, don't hesitate to ask around or discuss it with the system admin (ITP, 3rd floor, office 4S14).

To report issues, request software updates, or for general enquires, please also contact the system admin or just stop by at the office above.