



ScaDS.AI

DRESDEN LEIPZIG

CENTER FOR SCALABLE DATA ANALYTICS AND
ARTIFICIAL INTELLIGENCE

Machine Learning on HPC – Introduction

NHR Summer School

Lena Jurkschat, Christoph Lehmann, Elias Werner

Dresden, 11 June 2024



GEFÖRDERT VOM



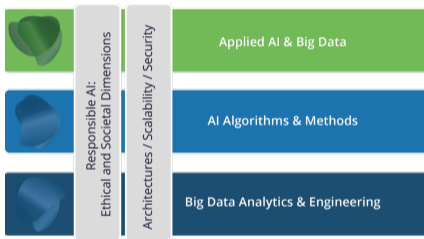
Bundesministerium
für Bildung
und Forschung

SACHSEN



Diese Maßnahme wird gefördert durch die Bundesregierung aufgrund eines Beschlusses des Deutschen Bundestages. Diese Maßnahme wird mitfinanziert durch Steuermittel auf der Grundlage des von den Abgeordneten des Sächsischen Landtags beschlossenen Haushaltes.

Instructors



<https://scads.ai/>



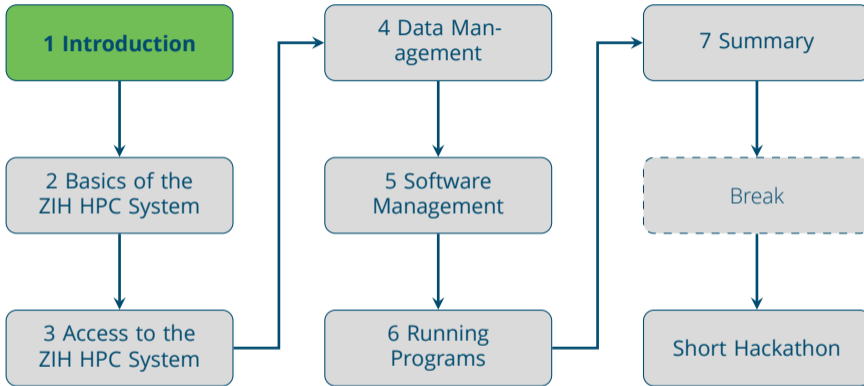
Lena Jurkschat,
Research Associate, ScaDS.AI
LLMs, Distributed Machine Learning, Performance Analysis



Christoph Lehmann,
Senior Researcher, ScaDS.AI
Statistics, Deep Learning, HPC



Elias Werner,
Research Associate, ScaDS.AI
Parallelization and performance analysis of
data-intensive application



Intention of this Training

- Making things understandable esp. for new users without a background in computer science
- Complete workflow examples are shown which are based on Python (can be used as a blueprint for other software/tools)
- From the user's perspective: What are the most important things to work with ML on an HPC cluster?

Intention of this Training

- Making things understandable esp. for new users without a background in computer science
- Complete workflow examples are shown which are based on Python (can be used as a blueprint for other software/tools)
- From the user's perspective: What are the most important things to work with ML on an HPC cluster?

Note

- We do **not** want to show everything possible. We want to show what is needed to get started.
- We do **not** want to show the perfect HPC-ML workflow. We want you to comprehend ML principles with an HPC machine starting with concrete examples.

Intention of this Training

- Making things understandable esp. for new users without a background in computer science
- Complete workflow examples are shown which are based on Python (can be used as a blueprint for other software/tools)
- From the user's perspective: What are the most important things to work with ML on an HPC cluster?

Note

- We do **not** want to show everything possible. We want to show what is needed to get started.
- We do **not** want to show the perfect HPC-ML workflow. We want you to comprehend ML principles with an HPC machine starting with concrete examples.

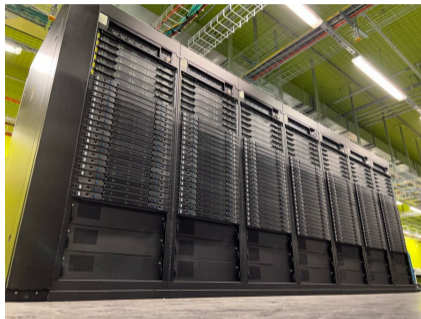
Hint

Please interrupt and ask immediately if something is not clear.

What is an HPC machine?

Terminology

- **Compute or login Node:** An individual computer, part of an HPC cluster
- **CPU, Core:** Central Processing Unit. A modern CPU is composed of numerous cores
- **Cluster:** A group of machines interconnected in a way that work together as a single system



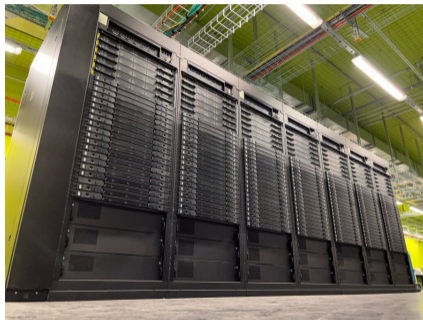
What is an HPC machine?

Terminology

- **Compute or login Node:** An individual computer, part of an HPC cluster
- **CPU, Core:** Central Processing Unit. A modern CPU is composed of numerous cores
- **Cluster:** A group of machines interconnected in a way that work together as a single system

Note

HPC cluster is a relatively tightly coupled collection of compute nodes, the interconnect typically allows for high bandwidth, low latency communication. Access to the cluster is provided through a login node(s). A resource manager and scheduler provide the logic to schedule jobs efficiently on the cluster.



HPC system

- **Facility:** Center for Information Services and High Performance Computing (ZIH)
- **Setup:** 5 clusters tailored to different types of workloads



HPC system

- **Facility:** Center for Information Services and High Performance Computing (ZIH)
- **Setup:** 5 clusters tailored to different types of workloads
- **Details:**
 - ▶ More than 100 000 cores,
 - ▶ 500 GPUs (NVIDIA - A100 and V100),
 - ▶ Flexible storage hierarchy with about 16 PB total capacity,
 - ▶ Linux (RHEL 8.7), batchsystem Slurm,
 - ▶ Perfect platform for highly scalable, data-intensive and compute-intensive applications.
 - ▶ New system with Nvidia H100 currently in planning.



The ZIH System



Local Machine

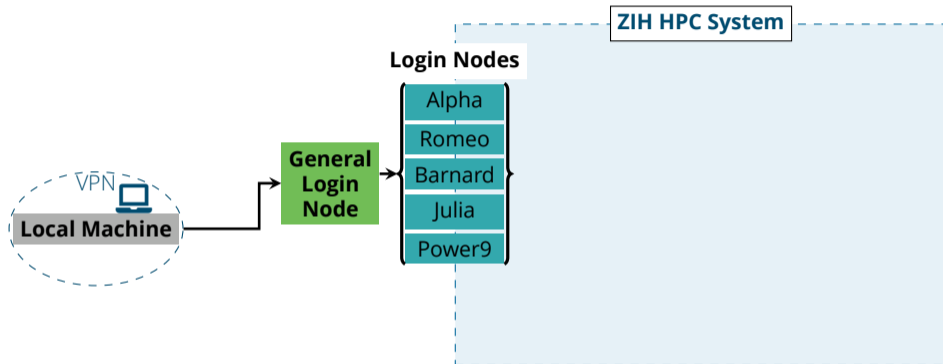
ZIH HPC System

The ZIH System

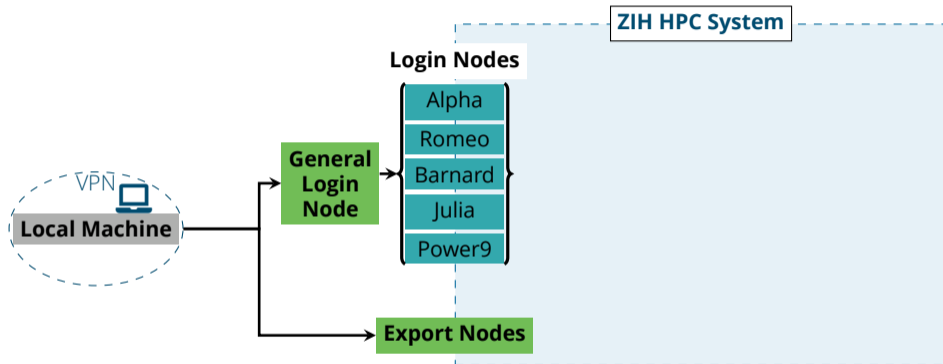


ZIH HPC System

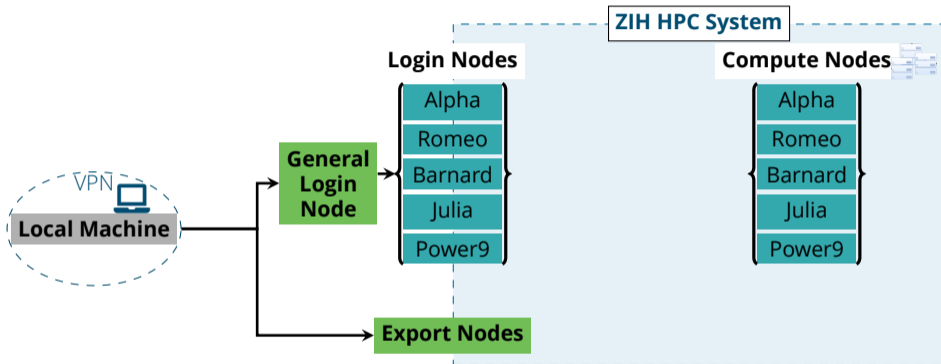
The ZIH System



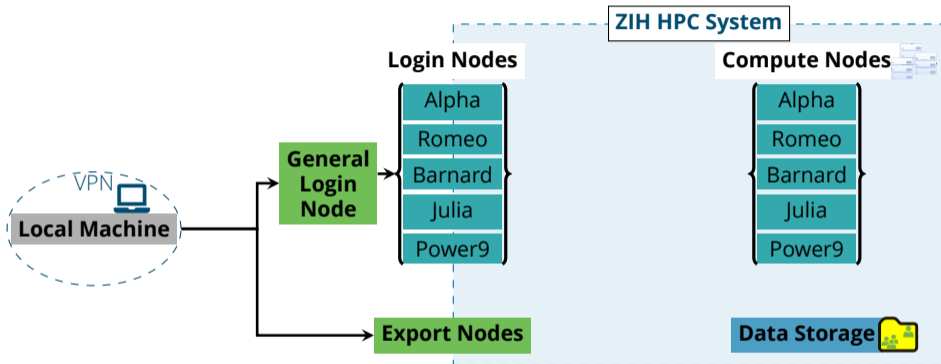
The ZIH System



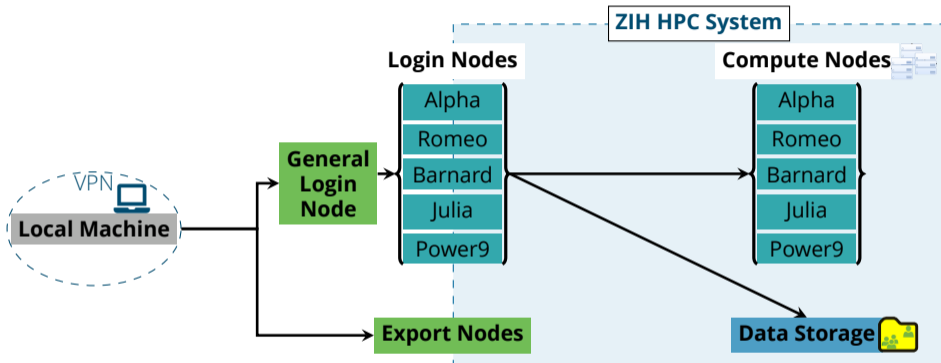
The ZIH System



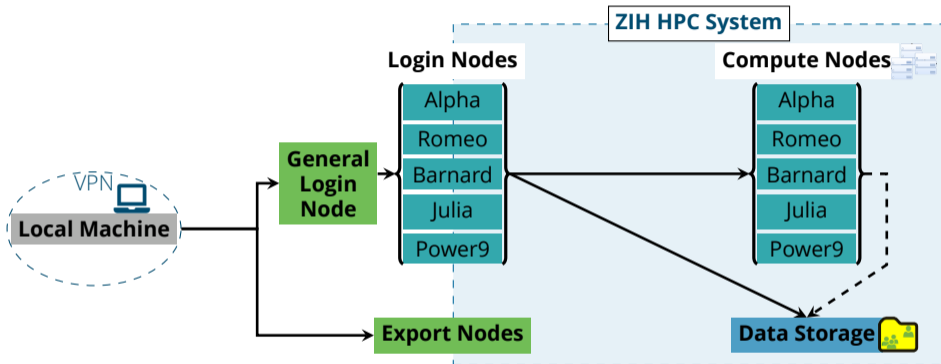
The ZIH System



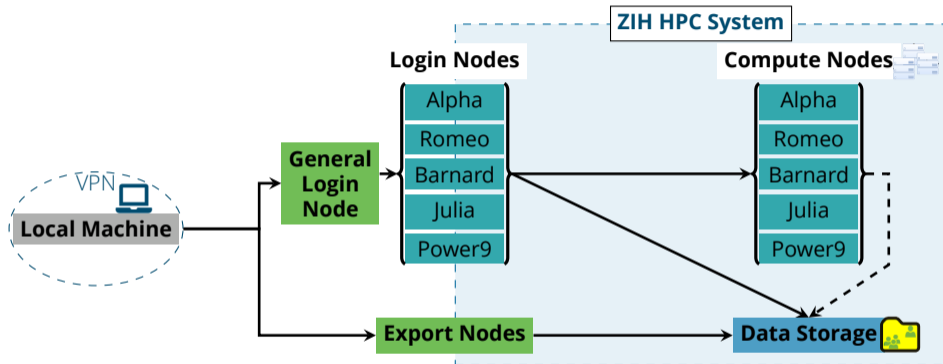
The ZIH System



The ZIH System



The ZIH System

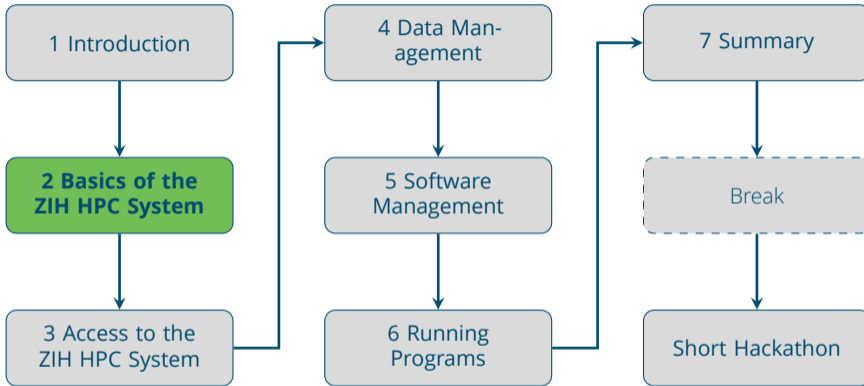




Survey

Please take part in our 2min survey:
<https://tud.link/uk7r2f>





Typical Workflow

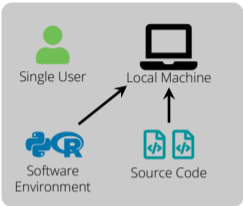
Recommended workflow to achieve a productive command line (CLI) based pipeline (some steps are optional)

1. Local machine: development of a pipeline as prototype (maybe with virtual environments like virtualenv or conda), usage of IDE or within jupyter notebook
2. Local machine: fully working CLI-based pipeline for small model/data
3. HPC machine: switch to a compute cluster with need for larger resources
4. HPC machine: use graphical front-end as jupyter notebook for testing purposes (software, hardware, algorithms)
5. HPC machine: fully working CLI-based pipeline with full data

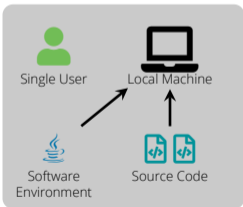
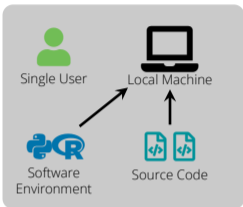
Hint

See the official docs for the ZIH system: <https://compendium.hpc.tu-dresden.de/>.

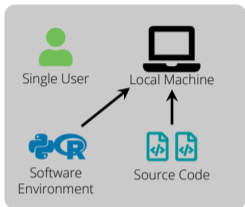
Switching from Local Machines to HPC Systems



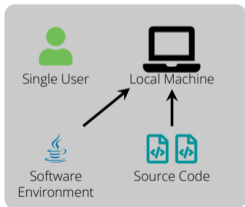
Switching from Local Machines to HPC Systems



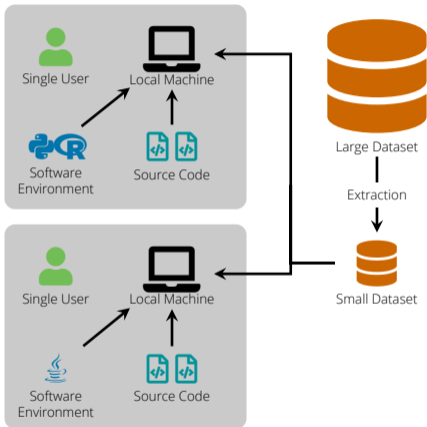
Switching from Local Machines to HPC Systems



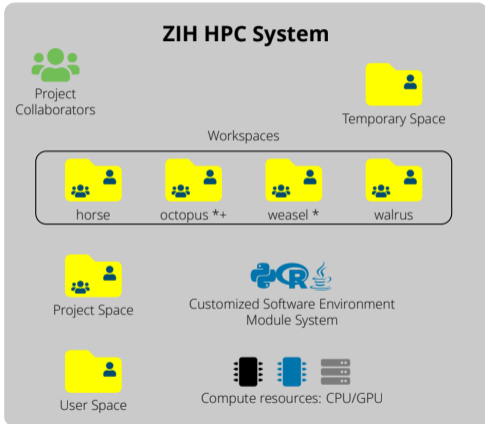
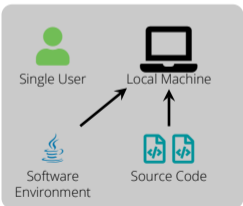
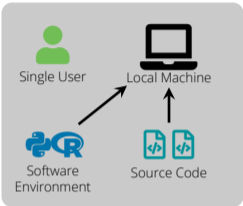
Large Dataset



Switching from Local Machines to HPC Systems

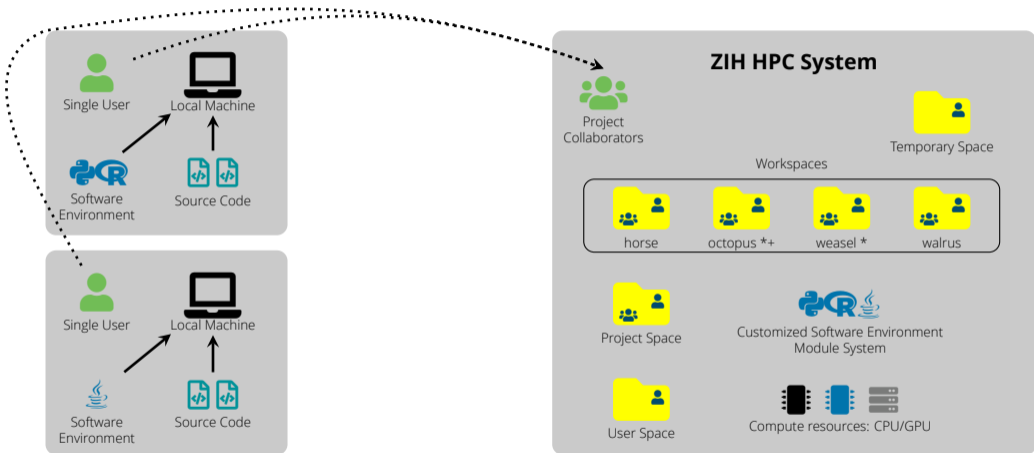


Switching from Local Machines to HPC Systems



* Not available currently.
 + Will only be available on Alpha Centauri cluster.

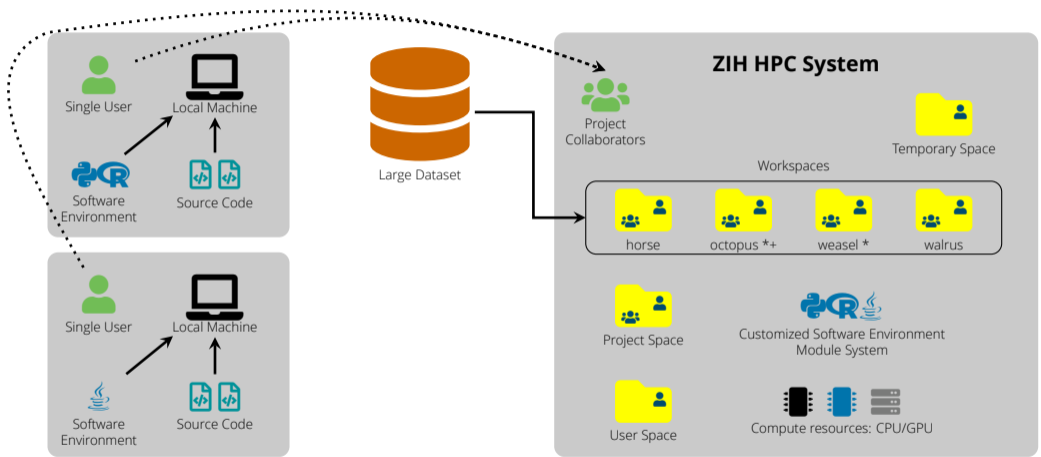
Switching from Local Machines to HPC Systems



* Not available currently.

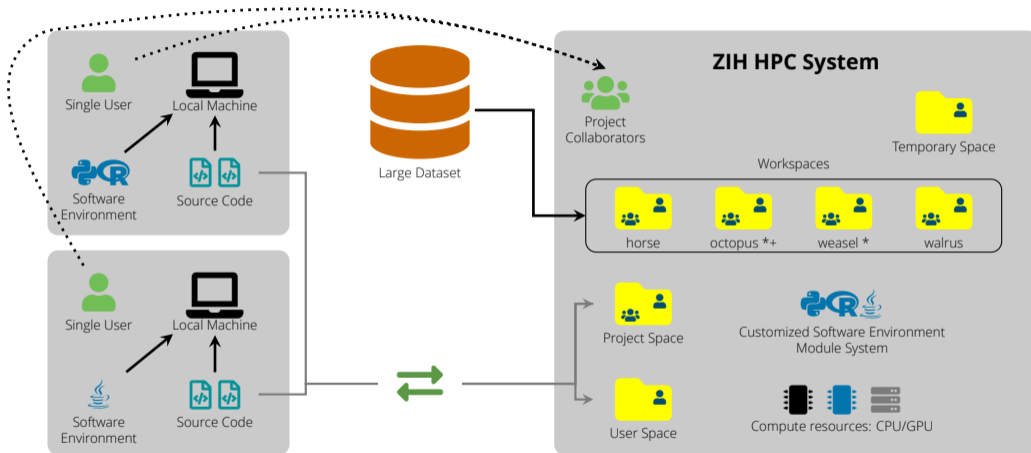
+ Will only be available on Alpha Centauri cluster.

Switching from Local Machines to HPC Systems



* Not available currently.
 + Will only be available on Alpha Centauri cluster.

Switching from Local Machines to HPC Systems

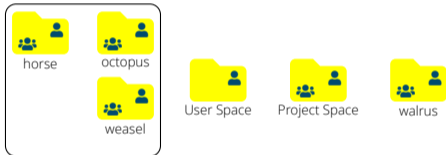


* Not available currently.

+ Will only be available on Alpha Centauri cluster.

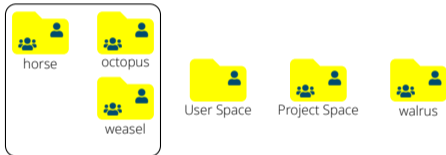
Collaborative Working

- We can achieve a collaborative working environment at two different levels based on the access/permission restrictions on the cluster:



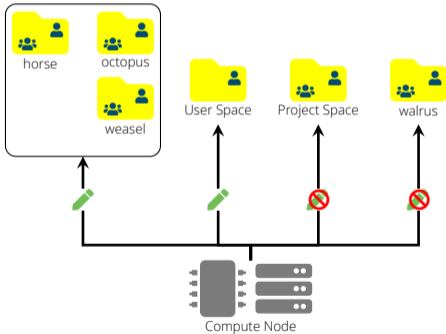
Collaborative Working

- We can achieve a collaborative working environment at two different levels based on the access/permission restrictions on the cluster:
 - ▶ Consideration 1: setting permissions to **different users** in your project (esp. important for data and customized software environments)



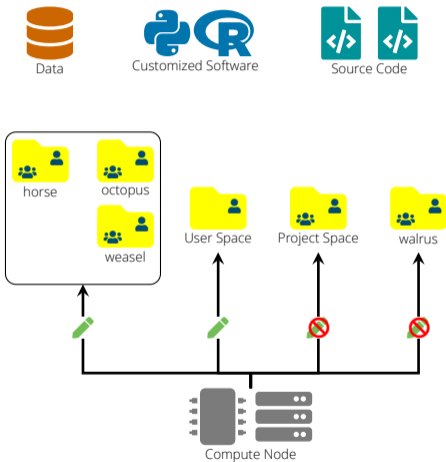
Collaborative Working

- We can achieve a collaborative working environment at two different levels based on the access/permission restrictions on the cluster:
 - ▶ Consideration 1: setting permissions to **different users** in your project (esp. important for data and customized software environments)
 - ▶ Consideration 2: Defined access possibilities based on **which node** you are on (login or compute) and on **which storage filesystem** you are trying to access (e.g. horse vs archive)



Collaborative Working

- We can achieve a collaborative working environment at two different levels based on the access/permission restrictions on the cluster:
 - ▶ Consideration 1: setting permissions to **different users** in your project (esp. important for data and customized software environments)
 - ▶ Consideration 2: Defined access possibilities based on **which node** you are on (login or compute) and on **which storage filesystem** you are trying to access (e.g. horse vs archive)
- It is important to distinguish usage of data, software environment, source code.

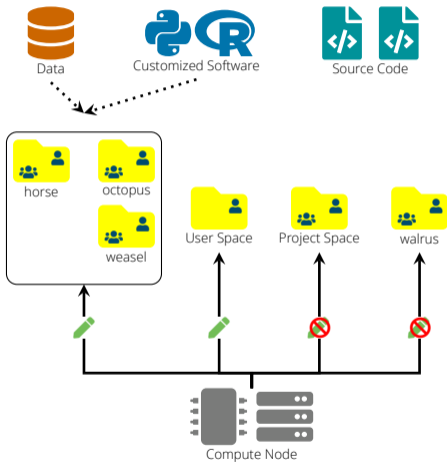


Collaborative Working

- We can achieve a collaborative working environment at two different levels based on the access/permission restrictions on the cluster:
 - ▶ Consideration 1: setting permissions to **different users** in your project (esp. important for data and customized software environments)
 - ▶ Consideration 2: Defined access possibilities based on **which node** you are on (login or compute) and on **which storage filesystem** you are trying to access (e.g. horse vs archive)
- It is important to distinguish usage of data, software environment, source code.

Hint

- **Data:** in workspaces with group access
- **Software environment:** module system or workspaces with group access

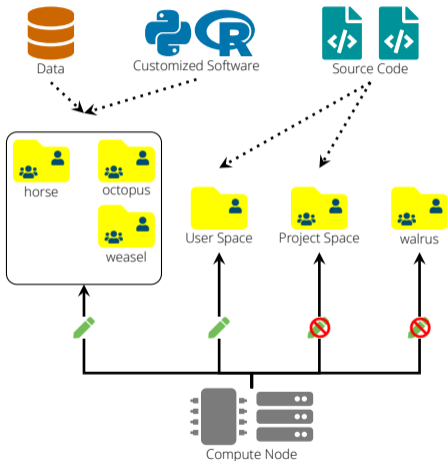


Collaborative Working

- We can achieve a collaborative working environment at two different levels based on the access/permission restrictions on the cluster:
 - ▶ Consideration 1: setting permissions to **different users** in your project (esp. important for data and customized software environments)
 - ▶ Consideration 2: Defined access possibilities based on **which node** you are on (login or compute) and on **which storage filesystem** you are trying to access (e.g. horse vs archive)
- It is important to distinguish usage of data, software environment, source code.

Hint

- **Data:** in workspaces with group access
- **Software environment:** module system or workspaces with group access
- **Source code:** user space or user specific directory in the project space



HPC project application

The ZIH system is structured by HPC projects. A HPC project on the ZIH system includes:

- project directory
- project group (linux group)
- project members (at least project leader and project administrator)
- resource quotas for compute time (CPU/GPU hours) and storage

HPC project application

The ZIH system is structured by HPC projects. A HPC project on the ZIH system includes:

- project directory
- project group (linux group)
- project members (at least project leader and project administrator)
- resource quotas for compute time (CPU/GPU hours) and storage

There are different possibilities to work with the ZIH HPC:

- create a new project
 - ▶ fill in the application form: <https://hpcprojekte.zih.tu-dresden.de/application/scads>
 - ▶ find additional information on the wiki: https://compendium.../application/project_request_form/ 
- join an existing project: e.g. new researchers in an existing project, teaching purposes

1 Introduction

2 Basics of the
ZIH HPC System

**3 Access to the
ZIH HPC System**

4 Data Man-
agement

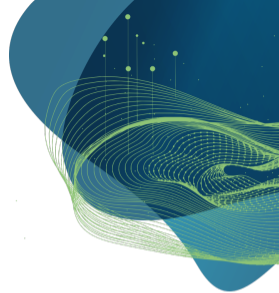
5 Software
Management

6 Running
Programs

7 Summary

Break

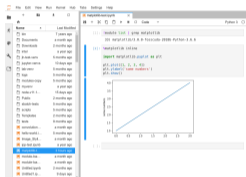
Short Hackathon



Access the ZIH HPC System

JupyterHub

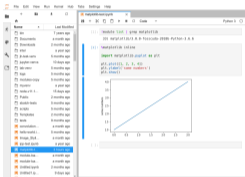
- browser based approach
- easiest way for beginners



Access the ZIH HPC System

JupyterHub

- browser based approach
- easiest way for beginners



SSH connection (CLI)

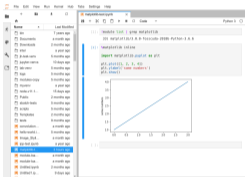
- "classical" approach
- command line interface (CLI) knowledge is necessary

The screenshot shows a terminal window with a dark background. It displays the output of an SSH command connecting to a remote host. The prompt `root@zih01:~#` is visible, followed by the execution of a command that results in a large block of text output, likely the contents of a file or the output of a system command.

Access the ZIH HPC System

JupyterHub

- browser based approach
- easiest way for beginners



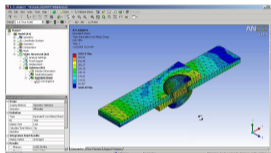
SSH connection (CLI)

- "classical" approach
- command line interface (CLI) knowledge is necessary

The screenshot shows a terminal window with a dark background. It displays the command to connect via SSH: `ssh user@hostname`. The output shows the remote host's fingerprint and the user's confirmation to proceed. The terminal also shows the user's prompt and some system information.

Desktop visualization

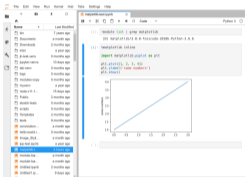
- esp. in the case of using GUI-based software
- e.g. Ansys, Vampir,...



Access the ZIH HPC System

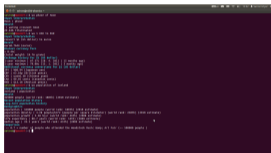
JupyterHub

- browser based approach
- easiest way for beginners



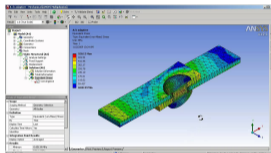
SSH connection (CLI)

- "classical" approach
- command line interface (CLI) knowledge is necessary



Desktop visualization

- esp. in the case of using GUI-based software
- e.g. Ansys, Vampir,...



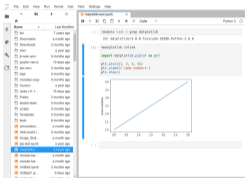
Note

- When working from **outside of the university network**, the ZIH HPC system can be accessed **only via Virtual Private Network (VPN)**
- Please use **provided** login information for this training

Access the ZIH HPC System

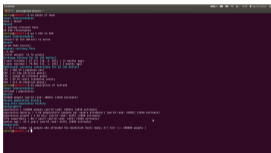
JupyterHub

- browser based approach
- easiest way for beginners



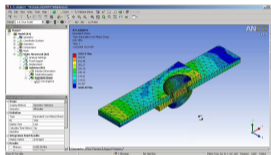
SSH connection (CLI)

- "classical" approach
- command line interface (CLI) knowledge is necessary



Desktop visualization

- esp. in the case of using GUI-based software
- e.g. Ansys, Vampir,...



Note

- When working from **outside of the university network**, the ZIH HPC system can be accessed **only via Virtual Private Network (VPN)**
- Please use **provided** login information for this training



Hint

More info on VPN: [🔗](#)

Job Scheduler Basics: Running Jobs

JupyterHub

Login to JupyterHub (slide 18)

Set parameters and load the required modules (slide 18)

Server Options

Select a job profile: Alpha - 1 core, 1.5 GB, 1 CPU, 1 hour

Advanced

Present: Alpha GPU NVIDIA Ampere A100

Cluster: alpha

Nodes (N) - nodes: 1

Number of tasks (n) - tasks: 1

CPUs per task (c) - cpus-per-task: 2

Memory per CPU (m) - mem-per-cpu: 2048

Generic resources (g) - gpus: 1

Runtime (L) - time: 01:00:00

Reservation (r) - reservation: no reservation

Project (A) - account: default

Workspace scope (s) - NotebookApp notebook_dir: default (your home directory)

Run

Running the application either using
Console or **Jupyter Notebook**



Interactive

Login to ZIH HPC System login-
node using console (slide 51)

Using **srun** to allocate resources
and attach shell using **--pty bash -l**

```
marie@login$ srun <params> --pty bash -l
```

Load the required modules (slide 61)

```
marie@compute$ module load <module>
```

Run the application

```
marie@compute$ python my_script.py
```

Batch

Login to ZIH HPC System login-
node using console (slide 51)

Create a sbatch script

```
#!/bin/bash
#SBATCH --partition=alpha # Use alpha partition
#SBATCH --nodes=1 # Use one node
#SBATCH --cpus-per-task=2 # Use 2 threads per task
#SBATCH --mem=8000 # Use 8gb of RAM
#SBATCH --gres=gpu:1 # Use 1 GPU per node
#SBATCH --time=00:05:00 # five minutes should be enough

# Loading required modules
module load modenv/hiera CUDA/11.1.1 OpenMPI/4.0.5 PyTorch/1.9.0

# Run the application
python my_script.py >> output
```

Run the application by submitting the
sbatch script using **sbatch** command

```
marie@login$ sbatch my_sbatch_script.sbatch
```

Access via JupyterHub

1. Connect to TU Dresden network via VPN, if you are accessing outside university network

Access via JupyterHub

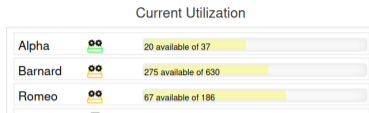
1. Connect to TU Dresden network via VPN, if you are accessing outside university network
2. Go to following JupyterHub link using any browser:
<https://jupyterhub.hpc.tu-dresden.de>

Access via JupyterHub

1. Connect to TU Dresden network via VPN, if you are accessing outside university network
2. Go to following JupyterHub link using any browser:
<https://jupyterhub.hpc.tu-dresden.de>
3. Login using ZIH-ID and password

Access via JupyterHub

1. Connect to TU Dresden network via VPN, if you are accessing outside university network
2. Go to following JupyterHub link using any browser:
<https://jupyterhub.hpc.tu-dresden.de>
3. Login using ZIH-ID and password
4. Allocate resources via the spawner interface, and click on "Start" to start a job



Server Options

Select a job profile:

Alpha - 1 core, 1.5 GB, 1 GPU, 1 hour

Advanced

Preset: Alpha GPU (NVIDIA Ampere A100) Save preset Delete preset ...

Cluster: alpha info

Nodes (-N, --nodes): 1

Number of tasks (-n, --tasks): 1

CPUs per task (-c, --cpus-per-task): 2

Memory per CPU (--mem-per-cpu): 2048

Generic resources (--gres): gpu:1

Runtime (-t, --time): 01:00:00 (hh:mm:ss)

Reservation (--reservation): no reservation

Project (-A, --account): default

Workspace scope (--NotebookApp.notebook_dir-): default (your home directory)

Start



Access via JupyterHub

1. Connect to TU Dresden network via VPN, if you are accessing outside university network
2. Go to following JupyterHub link using any browser:
<https://jupyterhub.hpc.tu-dresden.de>
3. Login using ZIH-ID and password
4. Allocate resources via the spawner interface, and click on "Start" to start a job

Hint

More information in the wiki:
<https://compendium.../access/jupyterhub> 

Current Utilization

Alpha	 	20 available of 37
Barnard	 	275 available of 630
Romeo	 	67 available of 186

Server Options

Select a job profile:

Alpha - 1 core, 1.5 GB, 1 GPU, 1 hour

Advanced

Preset Alpha GPU (NVIDIA Ampere A100)

Save preset Delete preset ...

Cluster: alpha

info

Nodes (-N, --nodes): 1

Number of tasks (-n, --tasks): 1

CPUs per task (-c, --cpus-per-task): 2

Memory per CPU (--mem-per-cpu): 2048

Generic resources (--gres): gpu:1

Runtime (-t, --time): 01:00:00 (hh:mm:ss)

Reservation (--reservation): no reservation

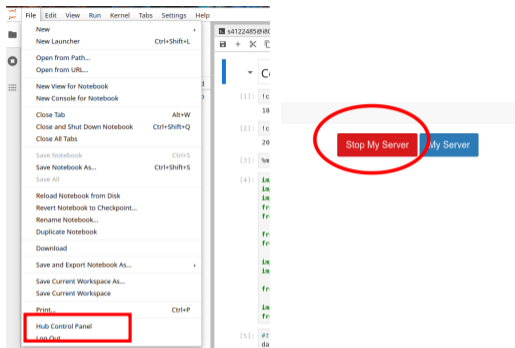
Project (-A, --account): default

Workspace scope (--NotebookApp.notebook_dir): default (your home directory)

Start

Access via JupyterHub

- When using JupyterHub do not forget to **stop your server!** Otherwise, the resources will not be available to others and will be included in your CPU quota!



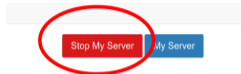
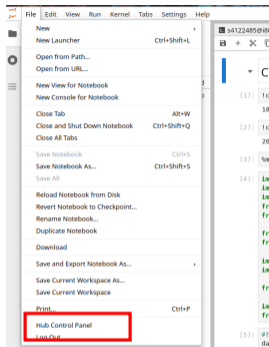
Access via JupyterHub

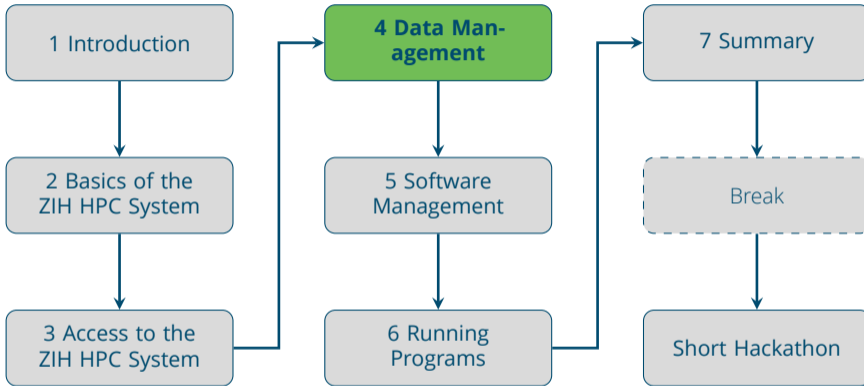
- When using JupyterHub do not forget to **stop your server!** Otherwise, the resources will not be available to others and will be included in your CPU quota!

Hint

In case of any issues while accessing/launching Jupyter Notebook, it is recommended to:

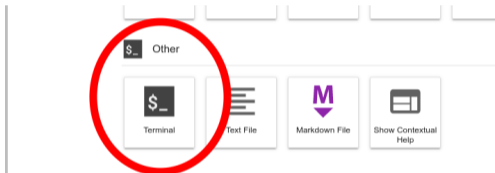
- clear browser cookies and cache, and/or
- use incognito/private browser mode





Data and Softwaremanagement on HPC: CLI commands

- As of today, CLI commands are required to manage data and software on HPC
- Jupyter offers the Terminal feature for interacting with such commands



```
LSTM_PCA_FixedForceLevel_1 X Terminal 2 X
Module Python/3.8.6, GCCcore/10.2.0 and 10 dependencies unloa
(python-env-python3.8.6-20210715-1421) ws_find
Usage: ws_find [options] [workspace]

(python-env-python3.8.6-20210715-1421) python Location_DNN.ip:
python: error while loading shared libraries: libpython3.8.so
(python-env-python3.8.6-20210715-1421) ls
2_RandForcePCA.ipynb          GRU_1FixedForce-Labels.py
ANN_1_FixedForceLevel_UR5.ipynb  GRUFixedForce-MA.py
CNN_EX1.ipynb                GRU_LOWEMI.ipynb
datax.hdf5                    jupyter-session-29232985.log
datay.hdf5                    jupyter-session-29401685.log
DIR                            jupyter-session-29401907.log
Features_1FixedF.ipynb        kfold.ipynb
GRU_1FixedForce-Labels.ipynb    K Fold Task A-Multi class.ip:
GRU_1FixedForce-Labels-MA.ipynb Location_DNN.ipynb
(python-env-python3.8.6-20210715-1421) █
```

Filesystems

- There are different areas for storing your data on the ZIH HPC system, called **filesystems**
- The filesystems have different properties (available space, time limit, size, permission rights). Different filesystems are available for different partitions. Therefore, choose the one that fits your project best.
- You need to create a workspace for your data on one of these filesystems.

Filesystems

- There are different areas for storing your data on the ZIH HPC system, called **filesystems**
- The filesystems have different properties (available space, time limit, size, permission rights). Different filesystems are available for different partitions. Therefore, choose the one that fits your project best.
- You need to create a workspace for your data on one of these filesystems.

Note

For this presentation: a filesystem refers to a "space/place" to store data and a workspace refers to the "access" you created to that filesystem

Filesystems

- There are different areas for storing your data on the ZIH HPC system, called **filesystems**
- The filesystems have different properties (available space, time limit, size, permission rights). Different filesystems are available for different partitions. Therefore, choose the one that fits your project best.
- You need to create a workspace for your data on one of these filesystems.

Note

For this presentation: a filesystem refers to a "space/place" to store data and a workspace refers to the "access" you created to that filesystem

scope	mount point	speed	size	duration
local	/tmp	+++	---	---
Global temporary	/weasel*	++	-	--
Global temporary	/horse	+	+++	+
Global permanent	/projects /home	--	++	++
archiving	/walrus	---	+++	+++

Filesystems

- There are different areas for storing your data on the ZIH HPC system, called **filesystems**
- The filesystems have different properties (available space, time limit, size, permission rights). Different filesystems are available for different partitions. Therefore, choose the one that fits your project best.
- You need to create a workspace for your data on one of these filesystems.

Note

For this presentation: a filesystem refers to a "space/place" to store data and a workspace refers to the "access" you created to that filesystem

Note

Detailed recommendations for the file system usage can be found [here](#)

scope	mount point	speed	size	duration
local	/tmp	+++	---	---
Global temporary	/weasel*	++	-	--
Global temporary	/horse	+	+++	+
Global permanent	/projects /home	--	++	++
archiving	/walrus	---	+++	+++

* Not available currently

Workspaces

- On the ZIH system, data has a **limited lifetime** depending on the filesystem
- User creates a workspace on filesystems with defined **expiration date**
- Data is **deleted** automatically after expiration.

Workspaces

- On the ZIH system, data has a **limited lifetime** depending on the filesystem
- User creates a workspace on filesystems with defined **expiration date**
- Data is **deleted** automatically after expiration.

Note

`/projects` and `/home` are permanent filesystems **without expiration** date

Workspaces

- On the ZIH system, data has a **limited lifetime** depending on the filesystem
- User creates a workspace on filesystems with defined **expiration date**
- Data is **deleted** automatically after expiration.



Note

`/projects` and `/home` are permanent filesystems **without expiration** date

Some commands to manage workspaces on the ZIH system:

CLI Command	Description
<code>ws_find --list</code>	Find available workspace filesystems
<code>ws_allocate -F <filesystem> <name_of_your_ws> <duration_of_your_ws></code>	Allocate workspace
<code>ws_list</code>	List your workspaces and get information
<code>ws_extend -F <filesystem> <name_of_your_ws> <duration></code>	Extend workspace

Workspaces

- On the ZIH system, data has a **limited lifetime** depending on the filesystem
- User creates a workspace on filesystems with defined **expiration date**
- Data is **deleted** automatically after expiration.



Note

`/projects` and `/home` are permanent filesystems **without expiration** date

Some commands to manage workspaces on the ZIH system:

CLI Command	Description
<code>ws_find --list</code>	Find available workspace filesystems
<code>ws_allocate -F <filesystem> <name_of_your_ws> <duration_of_your_ws></code>	Allocate workspace
<code>ws_list</code>	List your workspaces and get information
<code>ws_extend -F <filesystem> <name_of_your_ws> <duration></code>	Extend workspace

Hint

More info at: https://compendium.../data_lifecycle/workspaces/

Workspaces

Example

```
1 marie@login$ ws_find --list
2
3 marie@login$ ws_allocate -F horse myworkspace 30
4
5 Info: creating workspace.
6 /horse/ws/1/marie-myworkspace
7 remaining extensions : 2
8 remaining time in days: 30
9
10 marie@login$ ws_list
11
12 id: myworkspace
13 workspace directory : /horse/ws/1/marie-myworkspace
14 remaining time : 29 days 23 hours
15 creation time : Mon Oct 19 09:00:00 2023
16 expiration date : Wed Nov 18 08:00:00 2023
17 filesystem name : horse
18 available extensions : 2
```

Workspace: JupyterHub

- On JupyterHub, the working directory can be set in the spawner options
- Set **Workspace scope** parameter to the full path of your workspace or location
- Specified workspace will be set as the default directory for the notebook execution

Server Options

Select a job profile:

Alpha - 1 core, 1,5 GB, 1 GPU, 1 hour

Advanced

Presets: Alpha GPU (NVIDIA Ampere A100) Save preset Delete preset ...

Cluster: alpha info

Nodes (-N, --nodes): 1

Number of tasks (-n, --tasks): 1

CPUs per task (-c, --cpus-per-task): 2

Memory per CPU (--mem-per-cpu): 2048

Generic resources (--gres): gpu:1

Runtime (-t, --time): 01:00:00 (hh:mm:ss)

Reservation (--reservation): no reservation

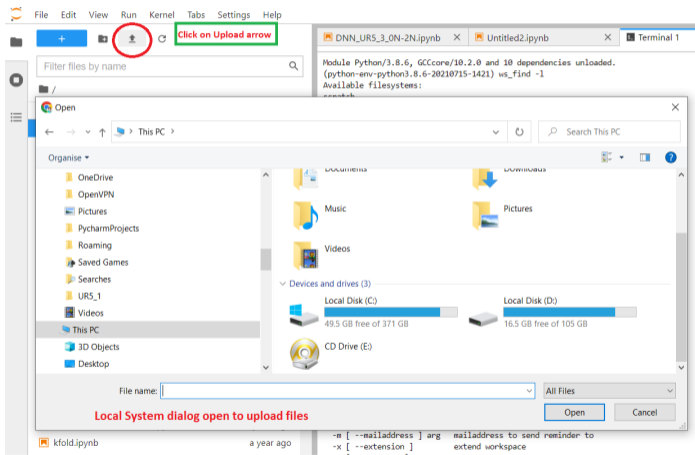
Project (-A, --account): default

Workspace scope (--NotebookApp.notebook_dir-): default (your home directory)

Start

Data Transfer: From Local Machine via JupyterHub

With JupyterHub GUI, data can be uploaded, downloaded to workspace or home directory.



Data Transfer: From External Sources

- Any file from external sources can also be downloaded on ZIH HPC system directly using commands like `wget` .

Example

```
1 #Use wget command to get dataset from the web
2 marie@compute$ wget --directory-prefix=/horse/ws/1/marie-myworkspace
   https://cloudstore.zih.tu-dresden.de/index.php/s/bm9HMTRGbN9ibrn/download/myfile.txt
3
4 #Check the content of the file with the nano editor
5 marie@compute$ nano /horse/ws/1/marie-myworkspace/myfile.txt
```

Data Transfer: Within the ZIH System

- Linux commands like `cp` and `mv` can be used transfer small data

Example

```
1 # Copy the training data into horse workspace directory
2 marie@login$ cp <my-file> /horse/ws/1/marie-myworkspace/
```

Data Transfer: Within the ZIH System

- Linux commands like `cp` and `mv` can be used transfer small data

Example

```
1 # Copy the training data into horse workspace directory
2 marie@login$ cp <my-file> /horse/ws/1/marie-myworkspace/
```

- Special data transfer nodes for moving large data between different filesystems. Commands for data handling and transfer are prefixed with `dt`.
 - ▶ `dtcp`, `dtls`, `dtmv`, `dtrm`, `dtrsync`, `dttar`
 - ▶ `dtqueue --me` to check the status of the data transfer
 - ▶ These commands create a **Slurm job** with dedicated resources to conduct the data handling/transfer.

Example

```
1 # Copy the training data into horse workspace directory
2 marie@login$ dtcp <my-file> /horse/ws/1/marie-myworkspace/
3
4 # Check copy job status
5 marie@login$ dtqueue --me
```

Data Transfer: Within the ZIH System

- Linux commands like `cp` and `mv` can be used transfer small data

Example

```
1 # Copy the training data into horse workspace directory
2 marie@login$ cp <my-file> /horse/ws/1/marie-myworkspace/
```

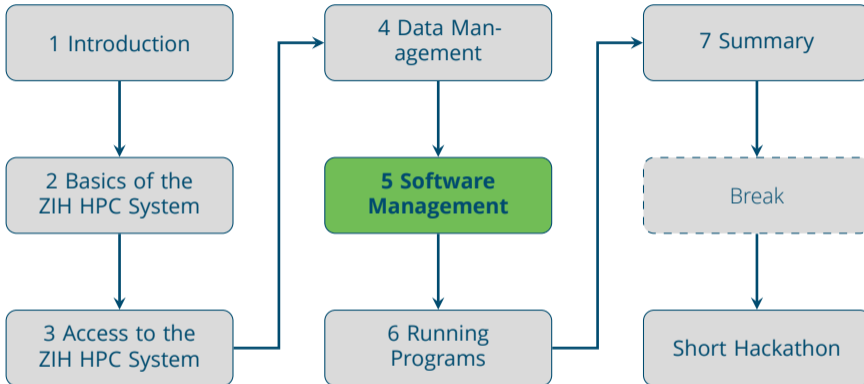
- Special data transfer nodes for moving large data between different filesystems. Commands for data handling and transfer are prefixed with `dt`.
 - ▶ `dtcp`, `dtls`, `dtmv`, `dtrm`, `dtrsync`, `dttar`
 - ▶ `dtqueue --me` to check the status of the data transfer
 - ▶ These commands create a **Slurm job** with dedicated resources to conduct the data handling/transfer.

Example

```
1 # Copy the training data into horse workspace directory
2 marie@login$ dtcp <my-file> /horse/ws/1/marie-myworkspace/
3
4 # Check copy job status
5 marie@login$ dtqueue --me
```

Hint

More info at: https://compendium.../data_transfer/datamover 




Software Management in JupyterHub

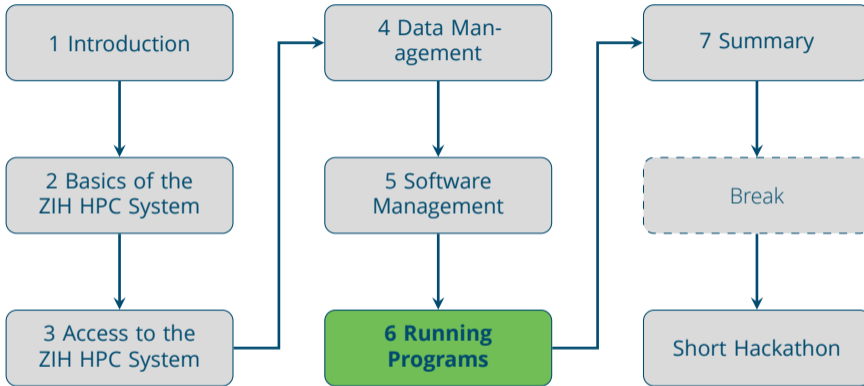
- On the ZIH HPC systems, modules provide interfaces to interact with pre-installed software

💡 Hint

User-based module management not necessary for JupyterHub in most cases.

For JupyterHub:

- Kernels interact with modules and provide additional Python packages in virtual environments
- E.g. we provide kernels for TensorFlow or PyTorch
- `pip install` also works in this environments
- Customized kernels can be created https://compendium.../access/jupyterhub_custom_environments 



Workload Management System

- User-defined calculations, programs etc. are **not run directly** and **interactively** on an HPC system (as commonly on a personal workstation or laptop)

Workload Management System

- User-defined calculations, programs etc. are **not run directly** and **interactively** on an HPC system (as commonly on a personal workstation or laptop)
- Running some program on an HPC machine means running it within a **temporary resource allocation**

Workload Management System

- User-defined calculations, programs etc. are **not run directly** and **interactively** on an HPC system (as commonly on a personal workstation or laptop)
- Running some program on an HPC machine means running it within a **temporary resource allocation**
- The definition of a program and a resource allocation is a **job**

Workload Management System

- User-defined calculations, programs etc. are **not run directly** and **interactively** on an HPC system (as commonly on a personal workstation or laptop)
- Running some program on an HPC machine means running it within a **temporary resource allocation**
- The definition of a program and a resource allocation is a **job**

Note

Most crucial issue esp. for HPC beginners: You need to specify in advance compute, memory, and time resources according to your program's needs. We will refer to this later.

Workload Management System

- User-defined calculations, programs etc. are **not run directly** and **interactively** on an HPC system (as commonly on a personal workstation or laptop)
- Running some program on an HPC machine means running it within a **temporary resource allocation**
- The definition of a program and a resource allocation is a **job**

Note

Most crucial issue esp. for HPC beginners: You need to specify in advance compute, memory, and time resources according to your program's needs. We will refer to this later.

- On the ZIH system: job scheduling and workload management with **slurm**
 - ▶ **Manages** jobs and provides an **interface** for the users to submit their jobs
 - ▶ **Evaluates** resource requirements and **priorities, distributes** jobs to suitable compute nodes -> "job queue"



Workload Management System

- User-defined calculations, programs etc. are **not run directly** and **interactively** on an HPC system (as commonly on a personal workstation or laptop)
- Running some program on an HPC machine means running it within a **temporary resource allocation**
- The definition of a program and a resource allocation is a **job**

Note

Most crucial issue esp. for HPC beginners: You need to specify in advance compute, memory, and time resources according to your program's needs. We will refer to this later.

- On the ZIH system: job scheduling and workload management with **slurm**
 - ▶ **Manages** jobs and provides an **interface** for the users to submit their jobs
 - ▶ **Evaluates** resource requirements and **priorities, distributes** jobs to suitable compute nodes -> "job queue"



More details about job schedulers can be found at: https://hpc-wiki.info/hpc/Scheduling_Basics

Slurm Parameters: Example

- alpha partition: `--partition=alpha`
- one node (no distributed application): `--nodes=1`
- two CPUs for preprocessing: `--cpus-per-task=2`
- 8GB memory for the data: `--mem=8000`
- one GPU : `--gres=gpu:1`
- run the job for 1 hr: `--time=01:00:00`
- account: `--account=p_nhr_summerschool` (only for this login)
- reservation: `--reservation=p_nhr_summerschool_TBD` (only for today)

Hint

- Typically, the most crucial parameters are `--gres` , `--cpus-per-task` , `--mem` , `--time` .
- Choosing the right parameters needs application-specific knowledge and experience.
- For more insights on resources use the performance monitoring tool PIKA (see slide 46).

Slurm Parameters: JupyterHub

- Many slurm parameters are usable directly in the JupyterHub GUI (“Advanced” mode)
- Access overview is provided on slide 18.
- Note the choice of workspace scope option.
- Once the JupyterHub is spawned, open the required notebook and execute the code in it.

Server Options

Select a job profile:

Alpha - 1 core, 1,5 GB, 1 GPU, 1 hour

Advanced

Preset	Alpha GPU (NVIDIA Ampere A100)	Save preset	Delete preset	...
Cluster:	alpha	info		
Nodes (-N, --nodes):	<input type="range"/>	1	⌵	
Number of tasks (-n, --tasks):	<input type="range"/>	1	⌵	
CPUs per task (-c, --cpus-per-task):	<input type="range"/>	2	⌵	
Memory per CPU (--mem-per-cpu):	<input type="range"/>	2048	⌵	
Generic resources (--gres):	gpu:1			
Runtime (-t, --time):	01:00:00		(hh:mm:ss)	
Reservation (--reservation):	no reservation			
Project (-A, --account):	default			
Workspace scope (--NotebookApp.notebook_dir-):	default (your home directory)			

Start

Job Scheduler Basics: Running Jobs

JupyterHub

Interactive

Batch

Job Scheduler Basics: Running Jobs

JupyterHub

Login to JupyterHub (slide 18)

Set parameters and load the required modules (slide 18)

Server Options

Select a job profile:
Alpha - 1 core, 1.5 GB, 1 GPU, 1 hour

Advanced

Present Alpha GPU (NVIDIA Ampere A100) Start preset Cancel preset

Cluster: alpha OK

Nodes (N, -nodes): 1 OK

Number of tasks (i.e., -tasks): 1 OK

CPUs per task (i.e., -cpus-per-task): 2 OK

Memory per CPU (i.e., -mem-per-cpu): 2048 OK

Generic resources (-gres): gpus:1

Priority (L, -time): 01:00:00 OK

Reservation (-reservation): no reservation

Project (A, -account): default

Workspace scope (i.e., NotebookApp.notebook_dir): default (your home directory)

Run

Running the application either using
Console or **Jupyter Notebook**



Interactive

Batch

Job Scheduler Basics: Running Jobs

JupyterHub

Login to JupyterHub (slide 18)

Set parameters and load the required modules (slide 18)

Server Options

Select a job profile:

Alpha - 1 core, 1.5 GB, 1 GPU, 1 hour

Advanced

Present: Alpha GPU (NVIDIA Ampere A100)

Cluster: alpha

Nodes (N, --nodes): 1

Number of tasks (i.e., --tasks): 1

CPUs per task (i.e., --cpus-per-task): 2

Memory per CPU (i.e., --mem-per-cpu): 2048

Generic resources (i.e., --gres): gpu:1

Partition (i.e., --time): 01:00:00

Reservation (i.e., --reservation): no reservation

Project (i.e., --account): default

Workspace scope (i.e., --workspace): default (your home directory)

Run

Running the application either using
Console or **Jupyter Notebook**



Interactive

Login to ZIH HPC System login-
node using console (slide 51)

Using **srun** to allocate resources
and attach shell using **--pty bash -l**

```
marie@login$ srun <params>
```

Load the required modules (slide 61)

```
marie@compute$ module load <module>
```

Run the application

```
marie@compute$ python my_script.py
```

Batch

Job Scheduler Basics: Running Jobs

JupyterHub

Login to JupyterHub (slide 18)

Set parameters and load the required modules (slide 18)

Server Options

Select a job profile:

Alpha - 1 core, 13 GB, 1 GPU, 1 hour

Advanced

Parent: Alpha GPU (NVIDIA Ampere A100)

Cluster: alpha

Nodes (N, -nodes): 1

Number of tasks (i.e., tasks): 1

CPUs per task (i.e., -cpus-per-task): 2

Memory per CPU (i.e., -mem-per-cpu): 2048

Generic resources (-gres): gpu:1

Partition (-time): 01:00:00

Reservation (-reservation): no reservation

Project (A, -account): default

Workspace scope (i.e., NotebookApp.notebook_dir): default (your home directory)

Run

Running the application either using
Console or **Jupyter Notebook**



Interactive

Login to ZIH HPC System login-node using console (slide 51)

Using **srun** to allocate resources and attach shell using **--pty bash -l**

```
marie@login$ srun <params> --pty bash -l
```

Load the required modules (slide 61)

```
marie@compute$ module load <module>
```

Run the application

```
marie@compute$ python my_script.py
```

Batch

Login to ZIH HPC System login-node using console (slide 51)

Create a sbatch script

```
#!/bin/bash
#SBATCH --partition=alpha # Use alpha partition
#SBATCH --nodes=1 # Use one node
#SBATCH --cpus-per-task=2 # Use 2 threads per task
#SBATCH --mem=8000 # Use 8gb of RAM
#SBATCH --gres=gpu:1 # Use 1 GPU per node
#SBATCH --time=00:05:00 # five minutes should be enough

# Loading required modules
module load modenv/hiera CUDA/11.1.1 OpenMPI/4.0.5 PyTorch/1.9.0

# Run the application
python my_script.py >> output
```

Run the application by submitting the sbatch script using **sbatch** command

```
marie@login$ sbatch my_sbatch_script.sbatch
```

Job Scheduler Basics: Running Jobs

JupyterHub

Login to JupyterHub (slide 18)

Set parameters and load the required modules (slide 18)

Server Options

Select a job profile:

Alpha - 1 core, 1.5 GB, 1 GPU, 1 hour

Advanced

Project: Alpha GPU (PVE/M-Ansys A100)

Cluster: alpha

Nodes (N, nodes): 1

Number of tasks (i.e., tasks): 1

CPU's per task (i.e., cpus-per-task): 2

Memory per CPU (i.e., mem-per-cpu): 2048

Generic resources (gres): gpus.1

Partition (i.e., limit): q1-00:00

Resource(s) (resources): cu-rtxsmem

Project (i.e., account): default

Workspace scope (i.e., NotebookApp notebook_dir): default (your home directory)

Running the application either using
Console or **Jupyter Notebook**



Interactive

Login to ZIH HPC System login-
node using console (slide 51)

Using **srun** to allocate resources
and attach shell using **--pty bash -l**

```
marie@login$ srun <params> --pty bash -l
```

Load the required modules (slide 61)

```
marie@compute$ module load <module>
```

Run the application

```
marie@compute$ python my_script.py
```

Note

JupyterHub job is inherently a batch job. Since it is a little different approach than batch job submission using CLI, it is shown here as a separate method.

Batch

Login to ZIH HPC System login-
node using console (slide 51)

Create a sbatch script

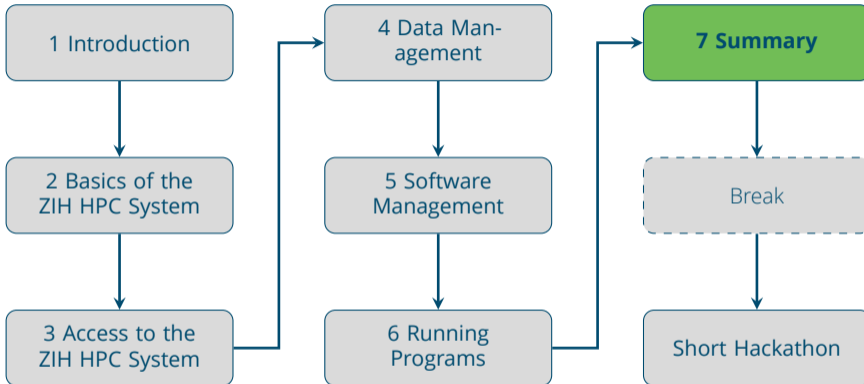
```
#!/bin/bash
#SBATCH --partition=alpha # Use alpha partition
#SBATCH --nodes=1 # Use one node
#SBATCH --cpus-per-task=2 # Use 2 threads per task
#SBATCH --mem=8000 # Use 8gb of RAM
#SBATCH --gres=gpu:1 # Use 1 GPU per node
#SBATCH --time=00:05:00 # five minutes should be enough

# Loading required modules
module load modenv/hiera CUDA/11.1.1 OpenMPI/4.0.5 PyTorch/1.9.0

# Run the application
python my_script.py >> output
```

Run the application by submitting the
sbatch script using **sbatch** command

```
marie@login$ sbatch my_sbatch_script.sbatch
```



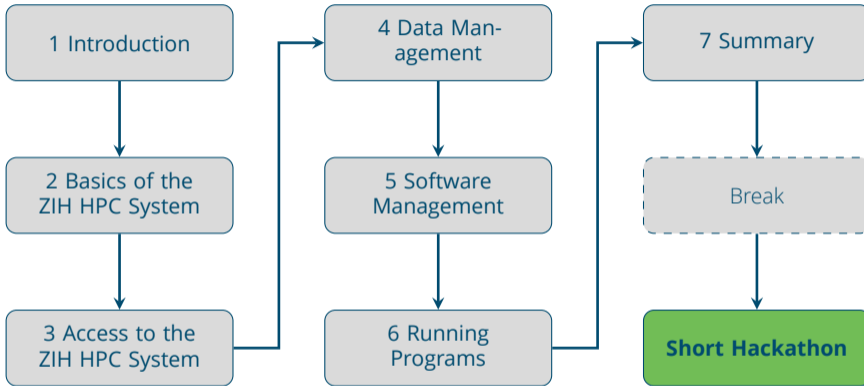
7 Summary

- There are three ways to work on HPC: JupyterHub, interactive, non-interactive.
- Do not use too much resources, **make resources free after you are done.**
- For more details read <https://doc.zih.tu-dresden.de/>
- Contact us with your own ideas, experiences and wishes!
- For getting support, esp. on ML, contact our ScaDS.AI service center <https://scads.ai/>.
- Logins (scads0xx) will work next 9 days, afterwards please apply for your own HPC project



Enjoy HPC!

Contact us on:
lena.jurkschat@tu-dresden.de
christoph.lehmann@tu-dresden.de
elias.werner@tu-dresden.de





Workflow: Using git

Git is a version-control system, commonly used for managing source code and coordinating the work of multiple contributors on a software project.

Please consider the following basics while using git on the ZIH system:

- git is available on all cluster nodes
- can be used on HPC in the same way as in the shell of your local machine, e. g. `git clone`, `git add`, `git commit`, `git push` etc.
- every project collaborator is working on his own copy and changes are pushed to the remote repository
- no binary files are contained in the repo (this keeps the storage requirement low)
- cloned repo copies are located in user space

Note

Using the same repo copy for different users will cause permission and access issues - DON'T DO THAT.



Workflow: Using your favorite IDE

Instead of using the personal favorite IDE¹ (e.g. Eclipse, Code::Blocks, PyCharm, VisualStudio etc.) on an HPC system editors as vi or vim are available.

- Without establishing any SSH session, directories from the ZIH system can be mounted into the filesystem of the local machine with SSHFS:

Example

```
1 marie@local$ sshfs  
    <loginname>@login1.alpha.hpc.tu-dresden.de:<path-of-your-workspace-on-zihsystem>  
    <path-on-your-local-machine> -o nonempty
```

- after calling SSHFS files are accessible from local file browsers and can be opened with your locally installed IDE

Hint

For running source code on the ZIH system just open a CLI in an interactive session and execute the code there. Do not forget to save in your IDE and therewith writing changes to the ZIH filesystem.

¹Integrated Development Environment

CI/CD + HPC

- MLOps is an essential practice for modern ML research. MLOps is not only for cloud solutions!
- MLOps for ML research on HPC:
 - ▶ Source control: GitHub, GitLab; DVC - data control
 - ▶ Experiment logging: MLFlow, Weights & Biases
 - ▶ Model registry: MLflow Model Registry, DVC
 - ▶ Orchestration/Containerisation: Docker, Singularity, Kubernetes
 - ▶ CI-CD: **GitLab CI/CD**

```

1  stages:
2    - linting
3    - test
4    - run
5
6  test-black:
7    stage: linting
8    tags:
9      - docker
10   image: registry.gitlab.com/pipeline-components/black/latest
11   script:
12     - black --check --verbose --diff --color --exclude=other/
13   rules:
14     - if: $CI_PIPELINE_SOURCE == "push"
15     - if: $CI_PIPELINE_SOURCE == "merge_request_event"
16
17  unittests:
18   stage: test
19   tags:
20     - docker
21   image: python:3.10.6-slim
22   script:
23     - pytest --vv tests/
24   rules:
25     - if: $CI_PIPELINE_SOURCE == "merge_request_event"
26
27  test_merge_request:
28   stage: run
29   tags:
30     - hpc
31   script:
32     - source /home/ai/0/vars-beef6_work/venv/ekslearn_python/bin/activate
33     - python /home/ai/0/vars-beef6_work/classif_slearn.py
34   variables:
35     SCHEDULER_PARAMETERS: "p alpha--gpus:1 --cpu-per-task=6 --mem=2000 --time=01:00:00"
36   rules:
37     - if: $CI_PIPELINE_SOURCE == "merge_request_event"

```

ScaDS.AI / Proseminar 2023 Genomics Language Models / Pipelines / #130470

Update .gitlab-ci.yml

✖ Failed Andrei Politov created pipeline for commit 98233b6a finished just now

For [test_ci](#)

latest 👁 5 Jobs 🕒 22 seconds, queued for 2 seconds

Pipeline Needs Jobs 5 Tests 0

linting

✔ test-black 🔄

✔ test-flake8 🔄

! test-pylint 🔄

test

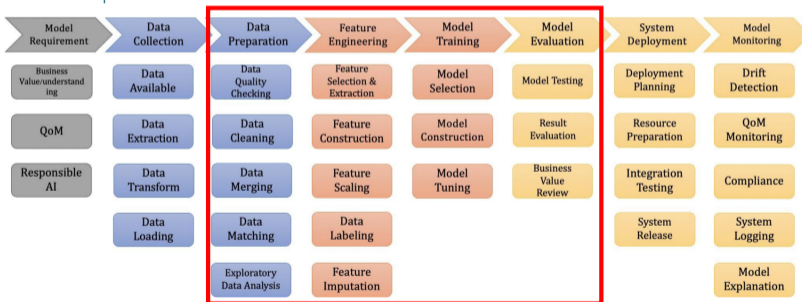
✔ unittests 🔄

run

✖ test_merge_request 🔄

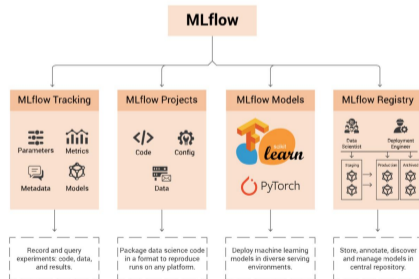
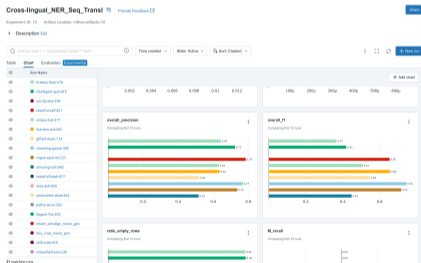
MLOps: General Terms

- The role of MLOps in the ML pipeline lies at the intersection of model development and model operations and directly includes both: solving the ML task itself and deploying, monitoring and updating the model.
- Task: Transfer the model and metrics from Jupyter Notebook to the ready solution.
- Goal: Make the whole ML pipeline more **reliable, efficient, reproducible and useful** for other developers and users.



Maximizing HPC: GitLab CI/CD, MLflow and even more

- Increased Efficiency: Utilize **HPC** capabilities for faster processing, improved resource utilization, and streamlined workflows
- Advanced Analytics: Leverage **MLflow** on ZIH HPC for efficient model development, tracking, and collaboration.
- Streamlined Development: Accelerate software delivery with GitLab CI/CD, enabling automated testing and training/inference pipelines.



Big Data Processing on HPC

Big Data Processing Engines,

- are software designed to make **scaling of big data processing** easy
- can **distribute work** across multiple actors efficiently



Flink



kafka



HPC clusters,

- make it **easy to get** powerful compute hardware
- are specialized for very **fast inter-node communication and I/O**



- For more information, visit: https://compendium.../software/big_data_frameworks/
- **Big Data Processing on HPC** training: [link to training info](#)

Performance Analysis – Overview

- Performance analysis: basis for optimizing parallel programs on HPC w.r.t. run-time and efficiency
- challenging task for complex applications
- analysis of an application: collecting relevant information at runtime
- different approaches to collect info:
 - ▶ easiest way: simple monitoring of CPU/GPU processes and e.g. RAM (simple tools on OS level).
 - ▶ advanced possibilities: using advanced tools level for tracing and profiling (e.g. analysis with Score-P)
- user-friendly starting point: simple monitoring using PIKA for answering questions as:
 - ▶ Are supposedly parallelized parts of the program actually executed in parallel?
 - ▶ Is the GPU actually used?

Performance Analysis – Job Monitoring with PIKA

- Simple approach for overview and to check used resources for a job: PIKA (hardware performance monitoring stack)
- monitoring during and after runtime: access PIKA web interface at https://selfservice.zih.tu-dresden.de/l/index.php/hpcportal/jobmonitoring/z../jobs_and_resources

The screenshot shows the PIKA web interface. The top navigation bar includes the logo of Technische Universität Dresden and the path: Startseite » ... » ZIH » Self Service Portal » HPC-Portal. Below this, there are two main sections: 'SELF SERVICE PORTAL:' and 'HPC-Portal:'. The 'HPC-Portal:' section has three tabs: 'Live', 'Job', and 'Footprint'. The 'Live' tab is highlighted with a red circle. Below the tabs, there is a search bar and a 'Total Jobs: 1' indicator. The main content area shows a table with columns for 'Name' and 'Project', containing one entry: 'bash' with project 'p_scaads'. At the bottom, it shows the date 'Stand: 29.10.2021 12:53'.

The screenshot shows a calendar interface for October 2021. The date 2021-10-29 is highlighted in blue. The calendar is part of a larger interface, likely for job scheduling or monitoring, with a 'Live' button visible below the calendar.

Choose between “Live” (for running jobs) and “Jobs” (finished jobs with filtering by date/time if needed)

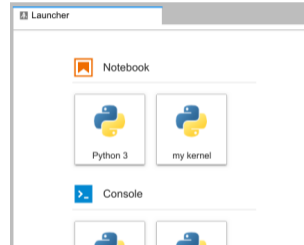
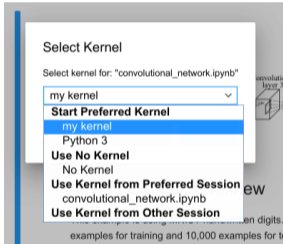
Hint

More info about PIKA at: <https://compendium.../software/pika/> 

Custom Jupyter Kernel

Allows you to install your own preferred Python packages and use them in your notebooks.

- You can switch kernels of existing notebooks in the menu
- Your kernels are listed on the launcher page



Hint

To create and use your own kernel, see here:

https://compendium.../access/jupyterhub_custom_environments/ 

https://compendium.../access/jupyterhub_teaching_example/ 

Python Virtual Environment + Kernel

- Virtual environments are isolated run-time environments and allow users to install additional Python packages
- For managing virtual environments on the ZIH HPC system, `virtualenv` is preferred to be used, which is a part of the Python modules

Example

```
1 # Load default Python
2 marie@compute$ module load Python
3
4 # Install virtual environment
5 marie@compute$ virtualenv --system-site-packages <path-to-desired-location>/my-env
6
7 # Activate the virtual environment
8 marie@compute$ source <path-to-desired-location>/my-env/bin/activate
9
10 # User can now install additional packages using pip command
11 (env) marie@compute$ pip install -r requirements.txt
12
13 # register and name your kernel
14 (env) marie@compute$ pip install ipykernel
15 (env) marie@compute$ python -m ipykernel install --user --name my-env --display-name="my kernel"
```

Preparing and Sharing Jupyter Kernels

- You can share such virtual environments with your collaborators
- Creates a unified software environment for collaborative working
- Helps to keep the system and your project clean without software duplicates

Hint

The `chown` and `chgrp` command change the owner and group of the directory/file

Example

```
1 # Give read and execute access to the group of the file
2 marie@compute$ chmod g+rx -R <path-to-your-workspace>/marie_dir
3 # Check permissions of directory
4 marie@compute$ ls -l <path-to-your-workspace>
5 # Output
6 drwxr-x--- 1 marie marie_group 4096 19. Oct 11:12 marie_dir
```

Shared Python Environment and Jupyter Kernels

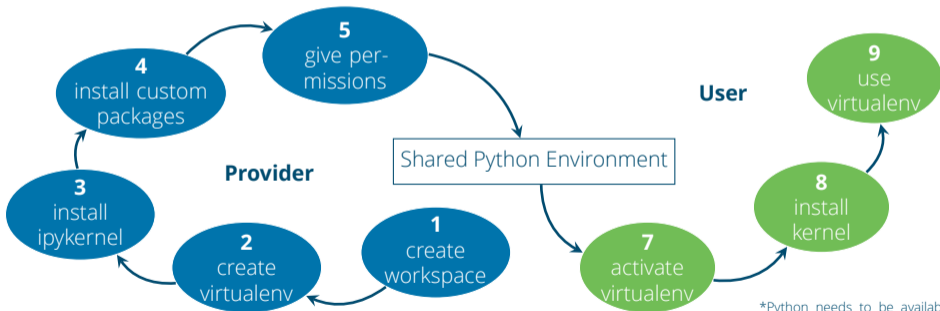
We assume two roles for collaboration:

- **Provider:**

- ▶ Provide a consistent Python environment
- ▶ Share environment with collaborators
- ▶ e.g. teacher, researcher

- **User:**

- ▶ uses a Python environment to reproduce results or studying
- ▶ e.g. pupil, research collaborator



*Python needs to be available for following the workflow

Access via SSH connection (CLI)

ZIH systems can be accessed via CLI using SSH connection.

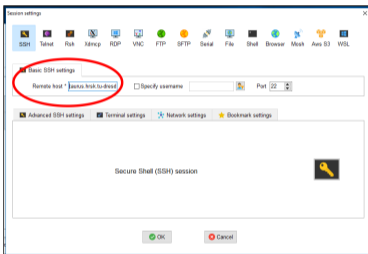
1. Connect to TU Dresden network via VPN, if you are accessing outside university network
2. Launch terminal
3. Connect to ZIH HPC login-nodes via `ssh` using ZIH-ID, login-node address (`login1.alpha.hpc.tu-dresden.de`) and ZIH account password
4. Allocate resources via `srun` or `sbatch`

Hint

More info at: https://compendium.../access/ssh_login/ 

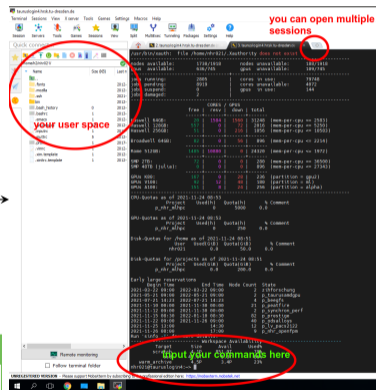
Access via SSH connection (CLI) (contd.)

- Accessing the ZIH system from a **Windows system**  is done via SSH client
- We recommend using **MobaXterm** (download free version at <https://mobaxterm.mobatek.net/download.html>)



Remote host: `login1.alpha.hpc.tu-dresden.de`
username: <ZIH-ID>

successful login →



Hint

More info at: https://compedium.../access/ssh_mobaxterm/ 

Slurm Parameters: Interactive Job

Example

```
1 # Allocate resources via slurm and run a bash interactively (--pty)
2 marie@login$ srun --nodes=1 --partition=alpha --gres=gpu:1 --cpus-per-task=2 --mem=8000 --time=
    01:00:00 --account=p_nhr_summerschool --reservation=p_nhr_summerschool_TBD --pty bash -l
3
4 # Now, within our temporary resource allocation, we run our interactive job
5 # [..do what you want and use the resources...]
6
7 marie@compute$ module load modenv/hiera GCC/10.2.0 CUDA/11.1.1 OpenMPI/4.0.5 PyTorch/1.10.0
    tqdm/4.56.2
8
9 marie@compute$ cd /beegfs/ws/1/marie-myworkspace/src
10
11 marie@compute$ python myscript.py
```

Slurm Parameters: Batch Job

Create the a batch script (`myscript.sbatch`) with the any editor:

Example

```

1 #!/bin/bash
2 #SBATCH --job-name=example
3 #SBATCH --nodes=1          # Number of nodes
4 #SBATCH --ntasks=1        # Run on a single CPU
5 #SBATCH --cpus-per-task=2  # use 2 threads per task
6 #SBATCH --gres=gpu:1       # 1 GPU per node
7 #SBATCH --time=0-00:05:00  # d-hh:mm:ss
8 #SBATCH --partition=alpha  # use alpha partition
9 #SBATCH --mem=2GB          # Memory per node
10 #SBATCH --output=%j.out    # Standard output and error log
11 #SBATCH --reservation=p_nhr_summerschool_TBD
12 #SBATCH --account=p_nhr_summerschool
13
14 # From here code is executed line by line
15
16 module load modenv/hiera GCC/10.2.0 CUDA/11.1.1 OpenMPI/4.0.5 PyTorch/1.10.0 tqdm/4.56.2
17 cd /beegfs/ws/1/marie-myworkspace/src
18 python myscript.py

```


Slurm Parameters: Batch Job

Run batch job files using command: `sbatch <jobfile>`




Example

```
1 # run sbatch job file via sbatch command and retrieve the job id
2 marie@login$ sbatch myscript.sbatch
3 Submitted batch job 20815837
4 #show status of submitted job
5 marie@login$ squeue --me
```

Job Scheduler Basics: Slurm

- Some basic commands used to manage Slurm jobs

Type	Command	Description
Job submission	<code>srun <params> <command-to-run></code>	Allocate resources and execute an application (interactive)
	<code>sbatch <params> <command-to-run></code>	Run a command, script etc. with dedicated hardware resources (non-interactive)
Job management	<code>scancel <jobid></code>	Cancel job
	<code>squeue --me</code>	List own jobs in queue and retrieve jobid
	<code>sinfo <params></code>	View information about nodes and partitions

- Hardware limits can be found at: https://compendium.../jobs_and_resources/slurm_limits/?h=memory\#slurm-resource-limits-table 
- More slurm parameters can be found at: https://compendium.../jobs_and_resources/slurm/\#options 
- Additional commands in the docs can be found at: https://compendium.../jobs_and_resources/slurm/ 

Access Permissions

Permission management is done by the 'chmod' command via CLI:

```
1 chmod [class][operator][mode] myobject
2 chmod [ugoa][-+=][rwx] myobject
```

where:

- class = [ugoa] = **u**ser/owner, **g**roup, **o**ther, **a**ll three classes
- operator = [-+=] = remove(-), add(+), set(=) the specified modes for the specified classes
- mode = [rwx] = **r**ead, **w**rite, **e**xecute

Example

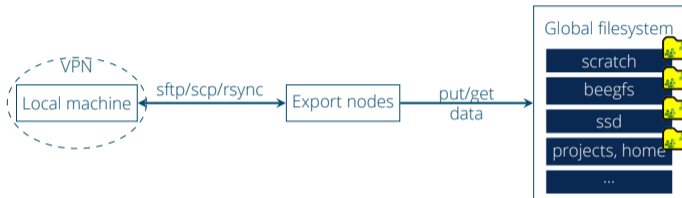
```
1 # check permissions of myobject i.e. the workspace
2 marie@login$ ls -l /beegfs/ws/1/marie-myworkspace/
3
4 # give "read, write and execute" access to the group that owns myobject
5 # -R means recursively i.e. on all files and directories at myobject
6 marie@login$ chmod g+rwx -R /beegfs/ws/1/marie-myworkspace/
```

Hint

More info at:  and https://compendium.../data_lifecycle/data_sharing/  or `man chmod`

Datatransfer: Local Machine and ZIH System

- Connection with local machine in VPN to export nodes via `scp`, `sftp`, `rsync`
- Transfer data to your target directory via export node



Example

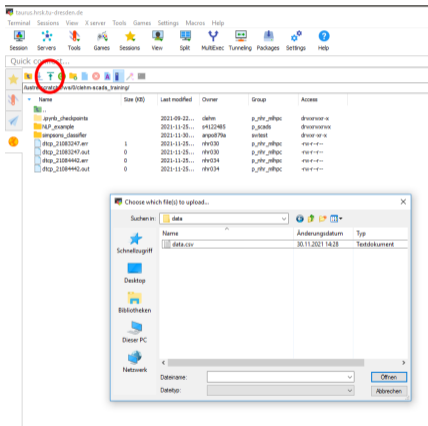
```
1 marie@local$ scp myscript.py marie@taurusexport.hrsk.tu-dresden.de:/home/marie/
2 marie@local$ rsync -avz myscript.sbatchr marie@taurusexport.hrsk.tu-dresden.de:/home/marie/
3 marie@local$ rsync -avz requirements.txt
  marie@taurusexport.hrsk.tu-dresden.de:/beegfs/ws/1/marie-myworkspace/
```

Hint

More info at: https://compendium.../data_transfer/export_nodes/ 

Datatransfer: Windows

With an SSH client (such as MobaXterm), data can be copied with a GUI



Datatransfer: Large vs Small Data

- Runtime of applications on login nodes is **limited** to 10min
- Data transfer that takes longer than that, will be **canceled**
- Therefore, it is recommended to use **export nodes** for large data transfer (might take longer)

Note

For very small data (up to 100MB), you can also use the login nodes (i.e. a single script file)

Module Management: CLI

After allocating resources, setting up your software can be done using following module commands:

Command	Description
<code>module avail/spider <your software></code>	Find required software
<code>module show <module name></code>	Show additional module information
<code>module load <module name></code>	Load module
<code>module list</code>	List all loaded modules
<code>module rm <module name></code>	Unload module
<code>module purge</code>	Unload all module
<code>module save</code>	Save modules for next login

Example

```
1 marie@compute$ module load modenv/hiera GCC/10.2.0 CUDA/11.1.1 OpenMPI/4.0.5 PyTorch/1.9.0 tqdm/4.56.2
   matplotlib/3.3.3
2
3 marie@compute$ module list
```

Module System

- On the ZIH system, software is organized in modules.
- A **module** is a user interface, that:
 - ▶ allows you to easily switch between different versions of software
 - ▶ dynamically sets up user's environment (PATH, etc.) and loads dependencies

CLI Command	Description
<code>module spider</code>	List of all available modules
<code>module spider <your software></code>	Find required software
<code>module show <module name></code>	Show additional module information
<code>module load <module name></code>	Load module
<code>module list</code>	List all loaded modules
<code>module rm <module name></code>	Unload module
<code>module purge</code>	Unload all module
<code>module save</code>	Save modules for next login

Module Management: Example

Example

```
1 #get an overview of available TensorFlow modules
2 marie@compute$ module spider TensorFlow
3 [...]
4
5 Versionen:
6 TensorFlow/2.8.4
7 [...]
8
9 #Check the requirements for a particular TensorFlow module
10 marie@compute$ module spider TensorFlow/2.8.4
11
12 [...]
13 Sie müssen alle Module in einer der nachfolgenden Zeilen laden bevor Sie das Modul "TensorFlow/2.8.4"
    laden können.
14
15 release/23.04 GCC/11.2.0 OpenMPI/4.1.1
16 [...]
17
18 #Load requirements and TensorFlow
19 marie@compute$ module load release/23.04 GCC/11.2.0 OpenMPI/4.1.1 TensorFlow/2.8.4
```