PP&B – Computer Support Group

FRITZ-HABER-INSTITUT
MAX-PLANCK-GESELLSCHAFT

HIGH PERFORMANCE COMPUTING AT FHI

THE VIPER SUPERCOMPUTER & THE FHI COMPUTE CLOUD – AN OVERVIEW

Simeon D. Beinlich[1,*] and Heinz Junkes[1]

[1] PP&B – Computer Support Group
Fritz-Haber-Institut der Max-Planck-Gesellschaft, Berlin, Germany

*beinlich@fhi.mpg.de

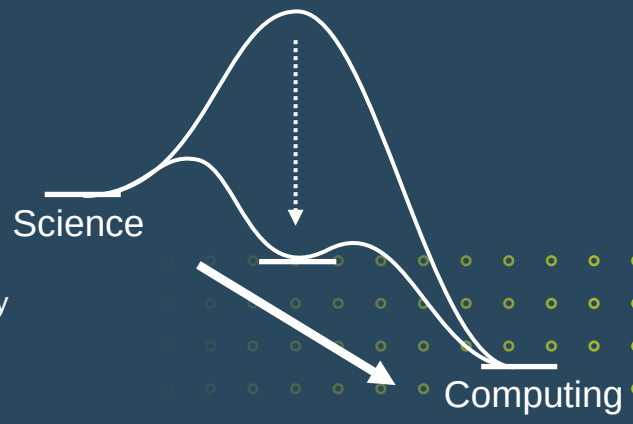

COMPUTING REQUIREMENTS IN SCIENCE

More data

More complex data

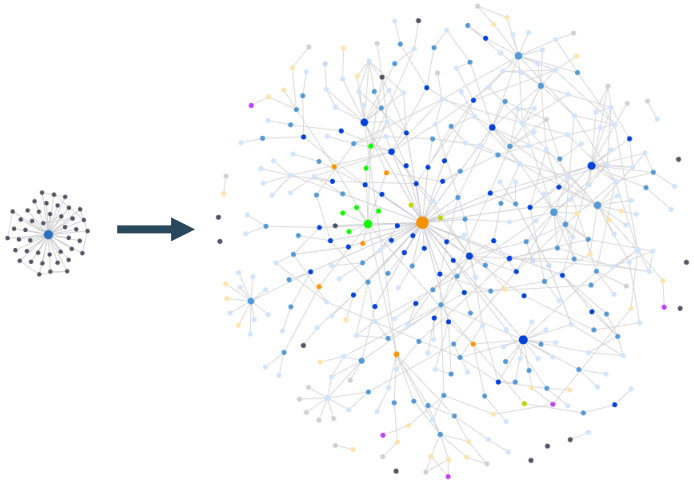

More complex analyses

More complex setups

More requirements

COMPUTING REQUIREMENTS IN SCIENCE

More data

More complex data

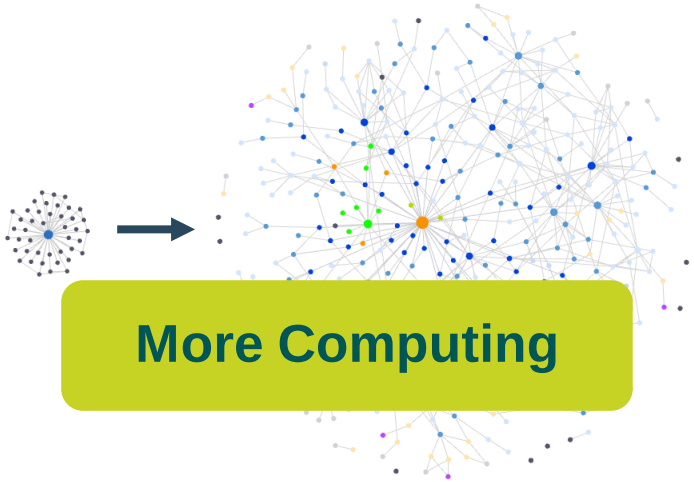

More complex analyses

More Computing

More complex setups

More requirements

COMPUTING REQUIREMENTS IN SCIENCE

Laptop / Desktop

Computing Infrastructure

Laptop

Run 24/7 for 1 month

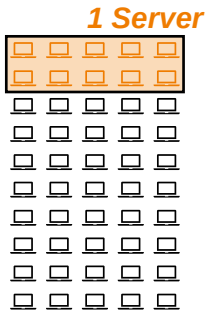

FHI COMPUTE
x 50

Run 3 days (1 server)

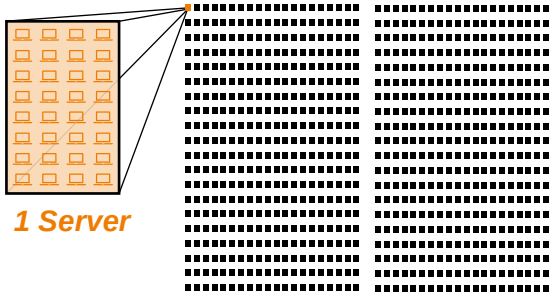

VIPER
x
25000

Run 22 hours @ 1 server
or 2 hours @ 10 servers

HPC INFRASTRUCTURE AT FHI (& MPG)

HPC INFRASTRUCTURE AT FHI (& MPG)

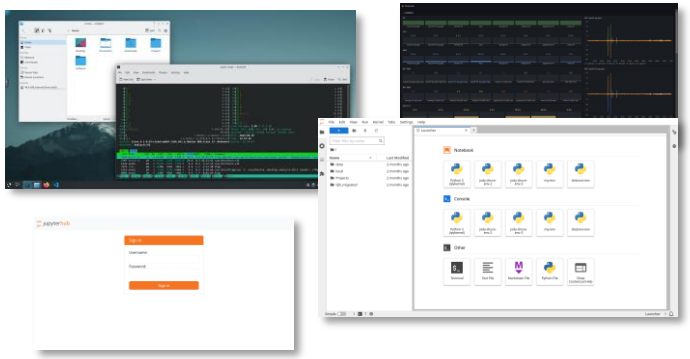

FHI (PPB)

FHI COMPUTE Servers

- Large Compute Servers (6x):
  Large CPU + Memory + GPU
- **Virtual Desktops** (Graphical Applications)
- **Jupyterhub** (Python & Matlab Server)
- Terminal Access
- **Direct Execution** (no Scheduling)
- User friendly

Anything that requires more than a desktop PC

MPG (MPCDF)

VIPER Supercomputer

- Huge Compute Servers (768x):
  Huge CPU + Memory + GPU
- Terminal Access
- **Job Submission** (Scheduling)
- Advanced

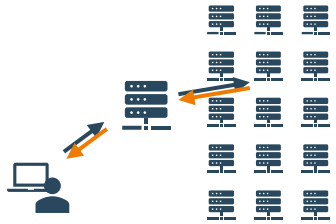

Large Scale Analyses / Processing & Ab-Initio Simulations

HPC INFRASTRUCTURE AT FHI (& MPG)

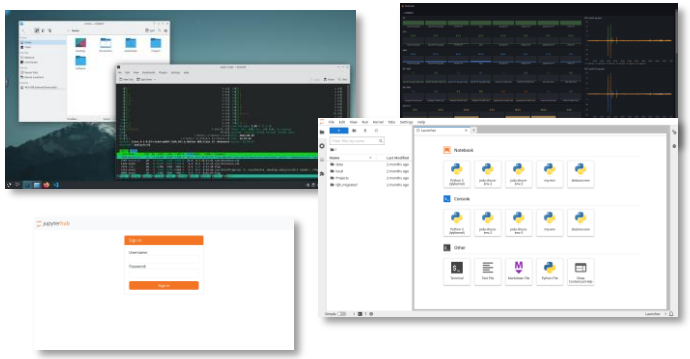

FHI (PPB)

FHI COMPUTE Servers

- Large Compute Servers (6x):
  Large CPU + Memory + GPU
- **Virtual Desktops** (Graphical Applications)
- **Jupyterhub** (Python & Matlab Server)
- Terminal Access
- **Direct Execution** (no Scheduling)
- User friendly

Anything that requires more than a desktop PC

MPG (MPCDF)

VIPER Supercomputer

- Huge Compute Servers (768x):
  Huge CPU + Memory + GPU
- Terminal Access
- **Job Submission** (Scheduling)
- Advanced

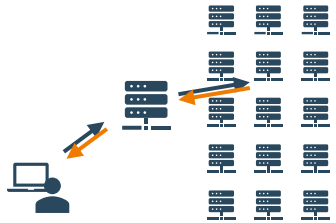

+ RAVEN, SuperMUC, JUWELS, …
(department-level / grant-based)

Large Scale Analyses / Processing & Ab-Initio Simulations

HPC INFRASTRUCTURE AT FHI (& MPG)

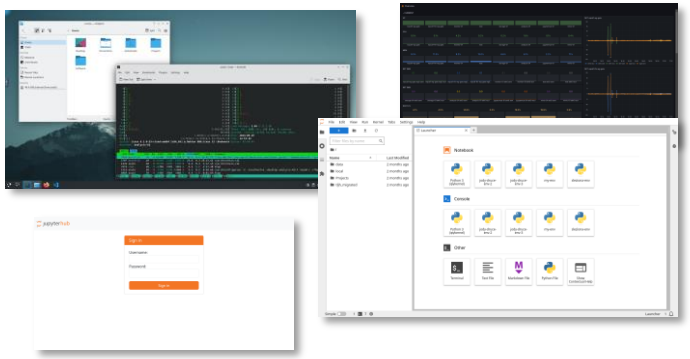

FHI (PPB)

FHI COMPUTE Servers

- Large Compute Servers (6x):
  Large CPU + Memory + GPU
- **Virtual Desktops** (Graphical Applications)
- **Jupyterhub** (Python & Matlab Server)
- Terminal Access
- **Direct Execution** (no Scheduling)
- User friendly

Anything that requires more than a desktop PC

MPG (MPCDF)

VIPER Supercomputer

- Huge Compute Servers (768x):
  Huge CPU + Memory + GPU
- Terminal Access
- **Job Submission** (Scheduling)
- Advanced

+ RAVEN, SuperMUC, JUWELS, …
(department-level / grant-based)

Large Scale Analyses / Processing & Ab-Initio Simulations

FHI COMPUTE SERVERS

FHI COMPUTE SERVERS – OVERVIEW

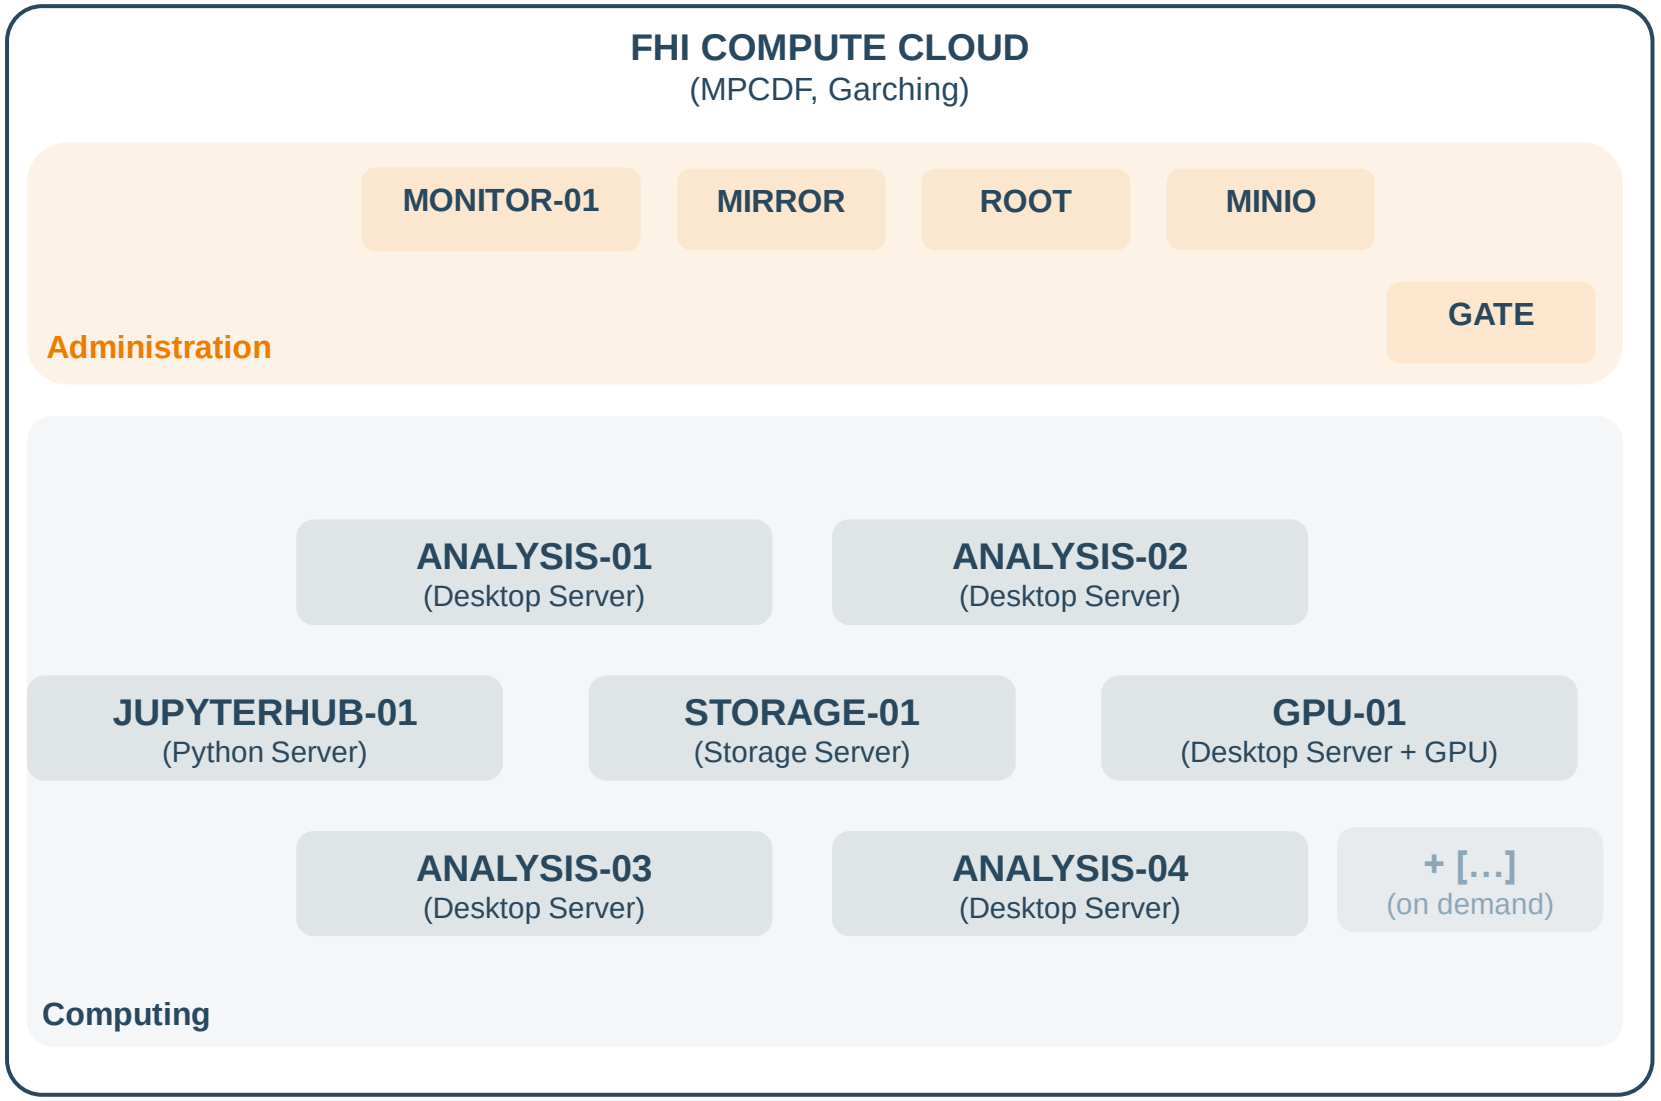

FHI COMPUTE SERVERS – OVERVIEW

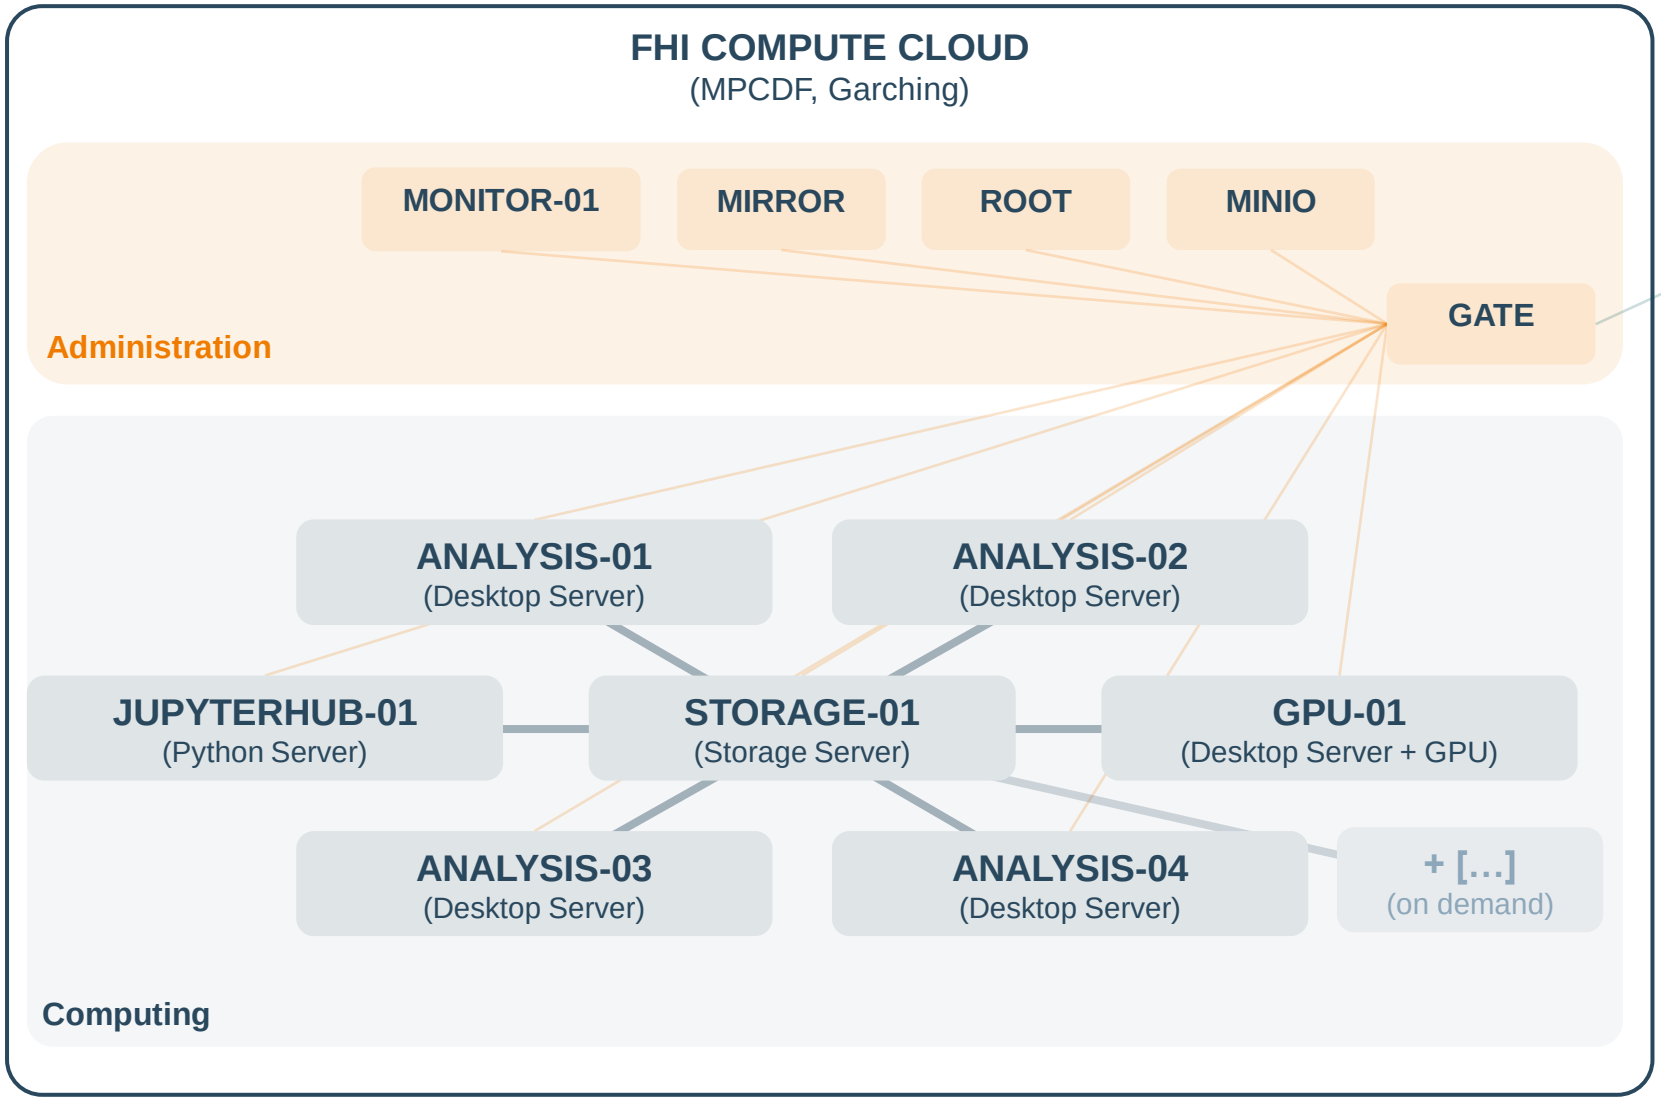

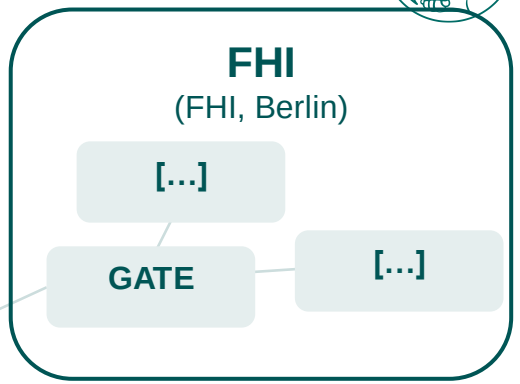

- 6 Compute Servers
  up to 32c128g + NV Tesla A100
- Full access from FHI network
  - Storage (Windows shares)
  - Desktop / terminal access
- Central storage
- Modular, scalable setup
- PPB administrated (full control)
- Custom software on request
- Shared project accounts possible
- No MPCDF account required
- Direct execution (no scheduler)

User friendly

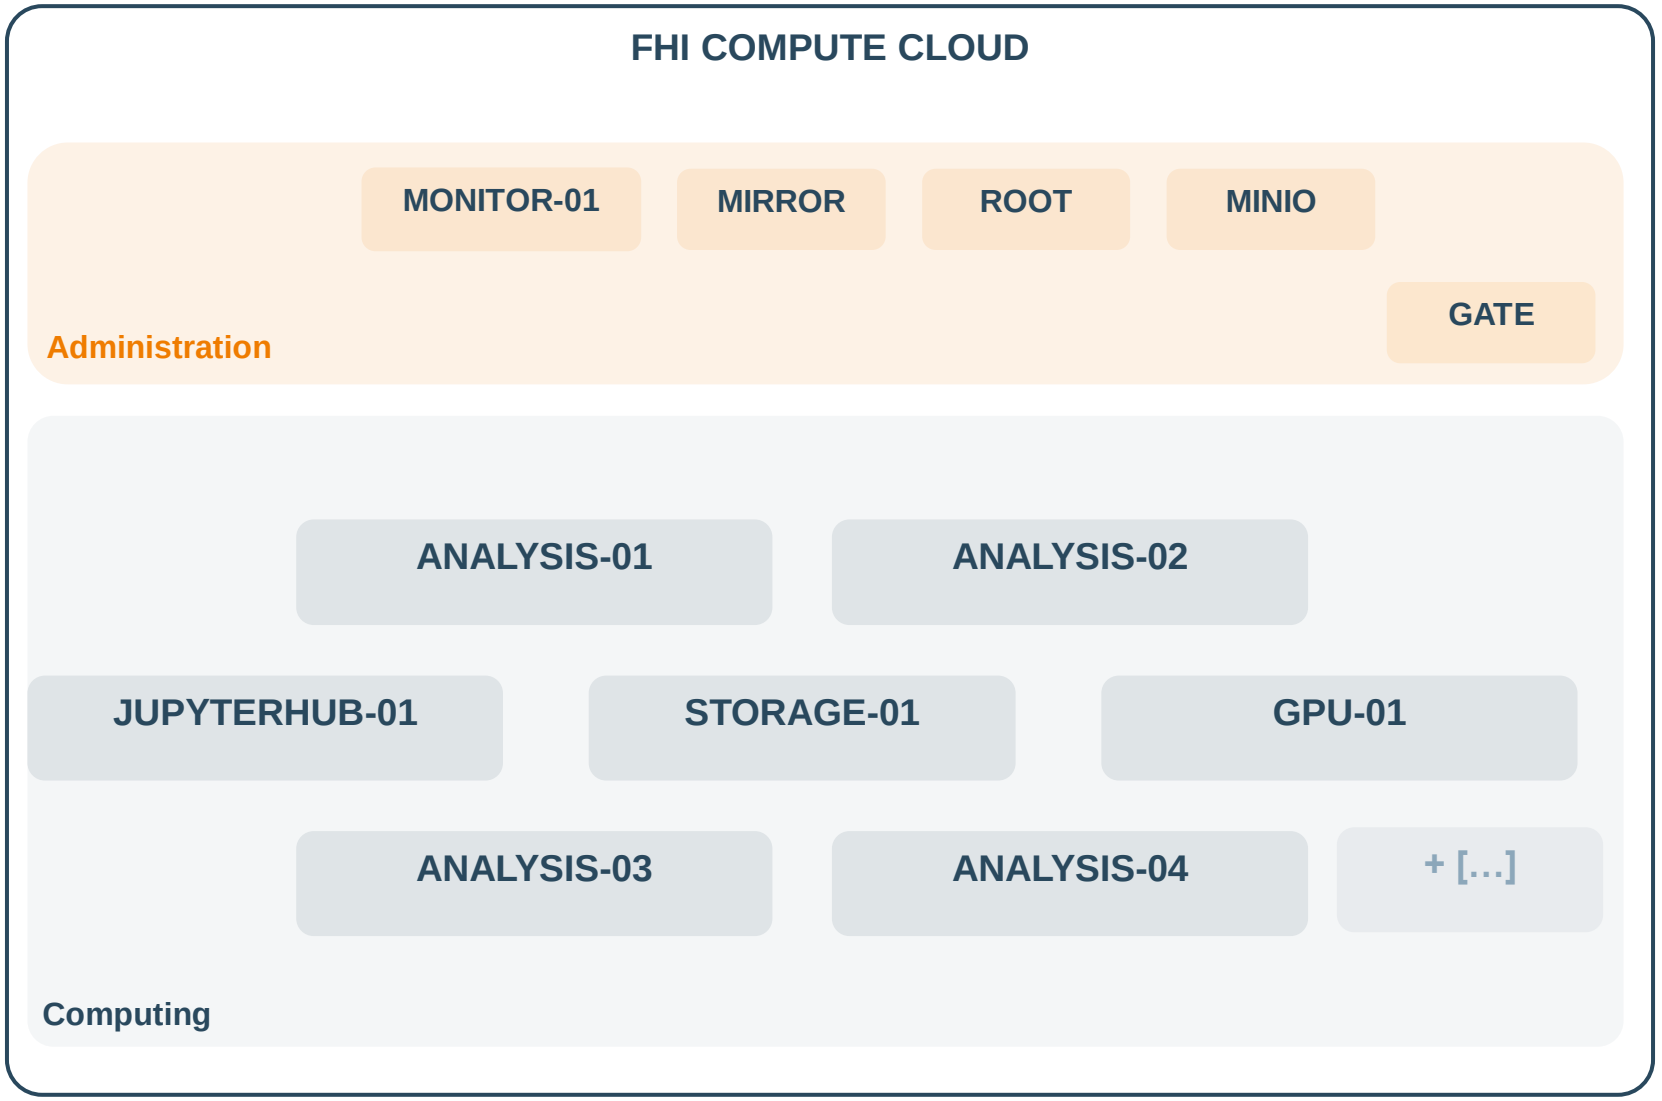
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ [...]
Computing

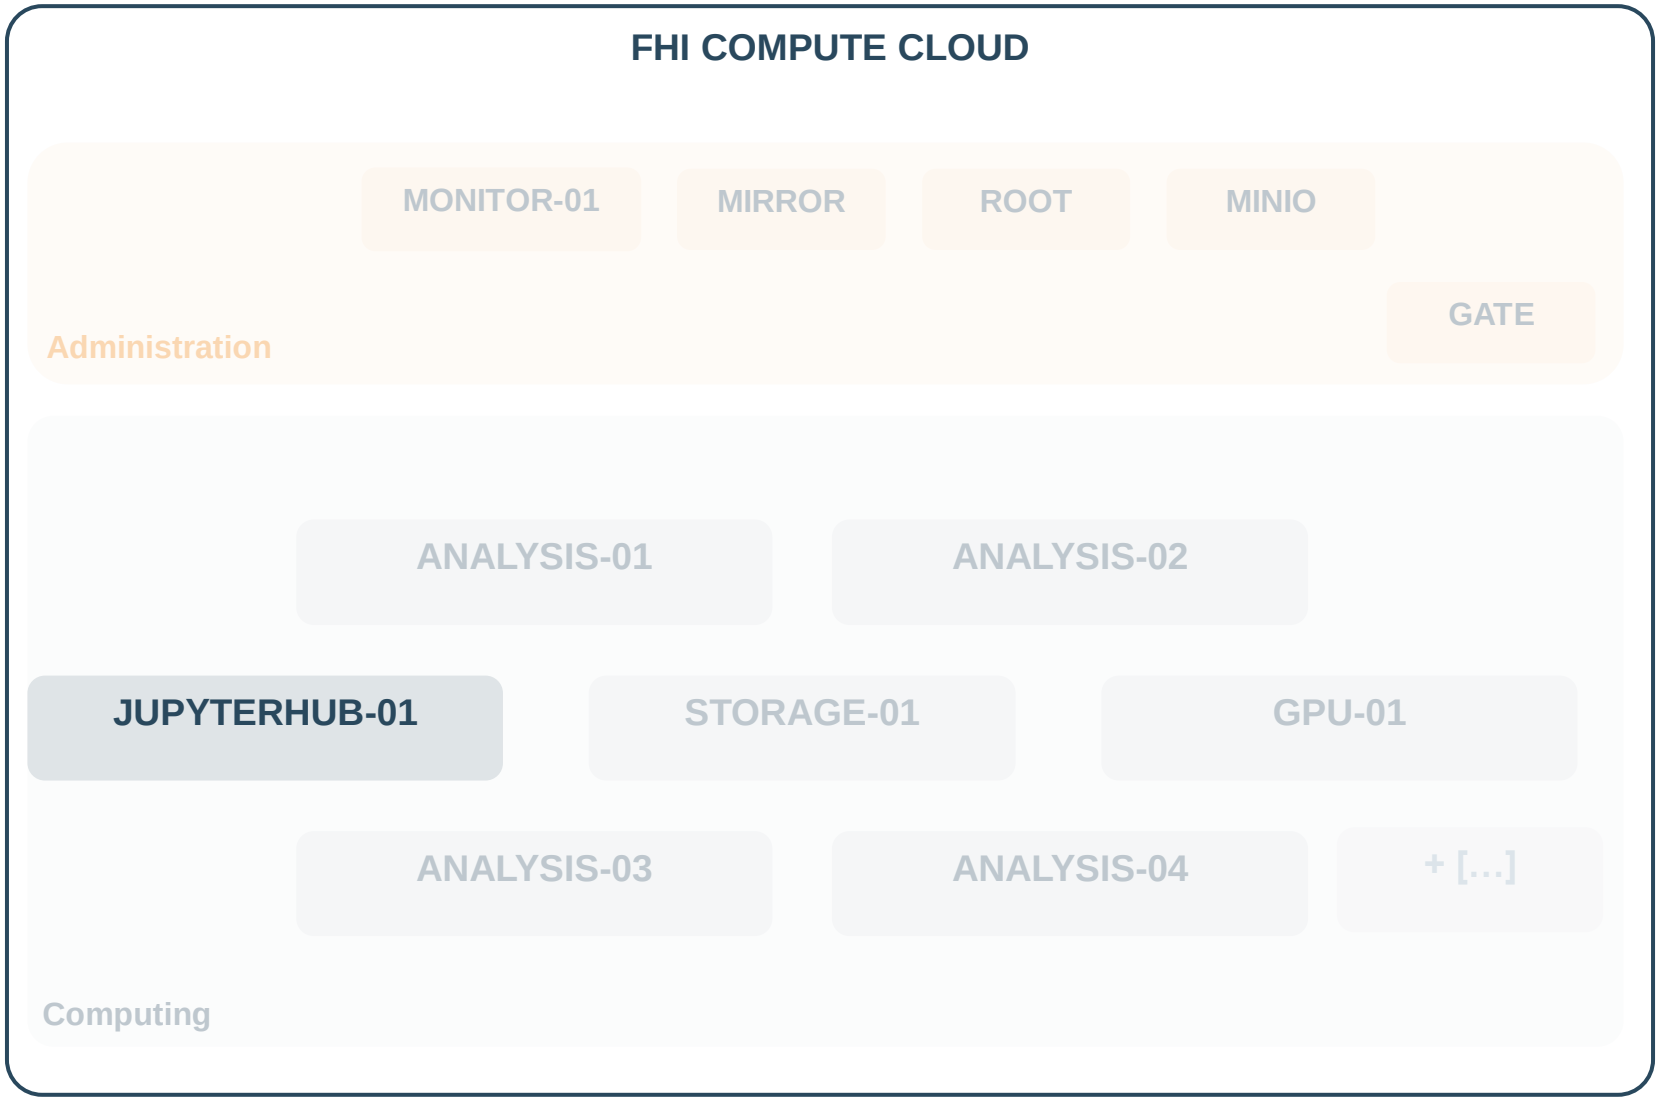
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

JUPYTERHUB-01 – PYTHON & MATLAB SERVER (AND MORE)

Use case

- Python / Matlab Shells
- Python / Matlab Scripts
- Python / Matlab Notebooks
  “interactive script with notes”
- Terminal

Resources

- 32 cores, 128GB RAM
- Central, FHI accessible storage
  - 2TB Home
  - 10TB Data
- Fast local Storage (1TB)

Advantages

- Large resources
- No local installation required (browser/terminal access)
- Custom setups for single users
- Multiple, independent Python setups
- Uniform setups for multiple users (e.g. for groups)
- Directly use ready-made notebooks!

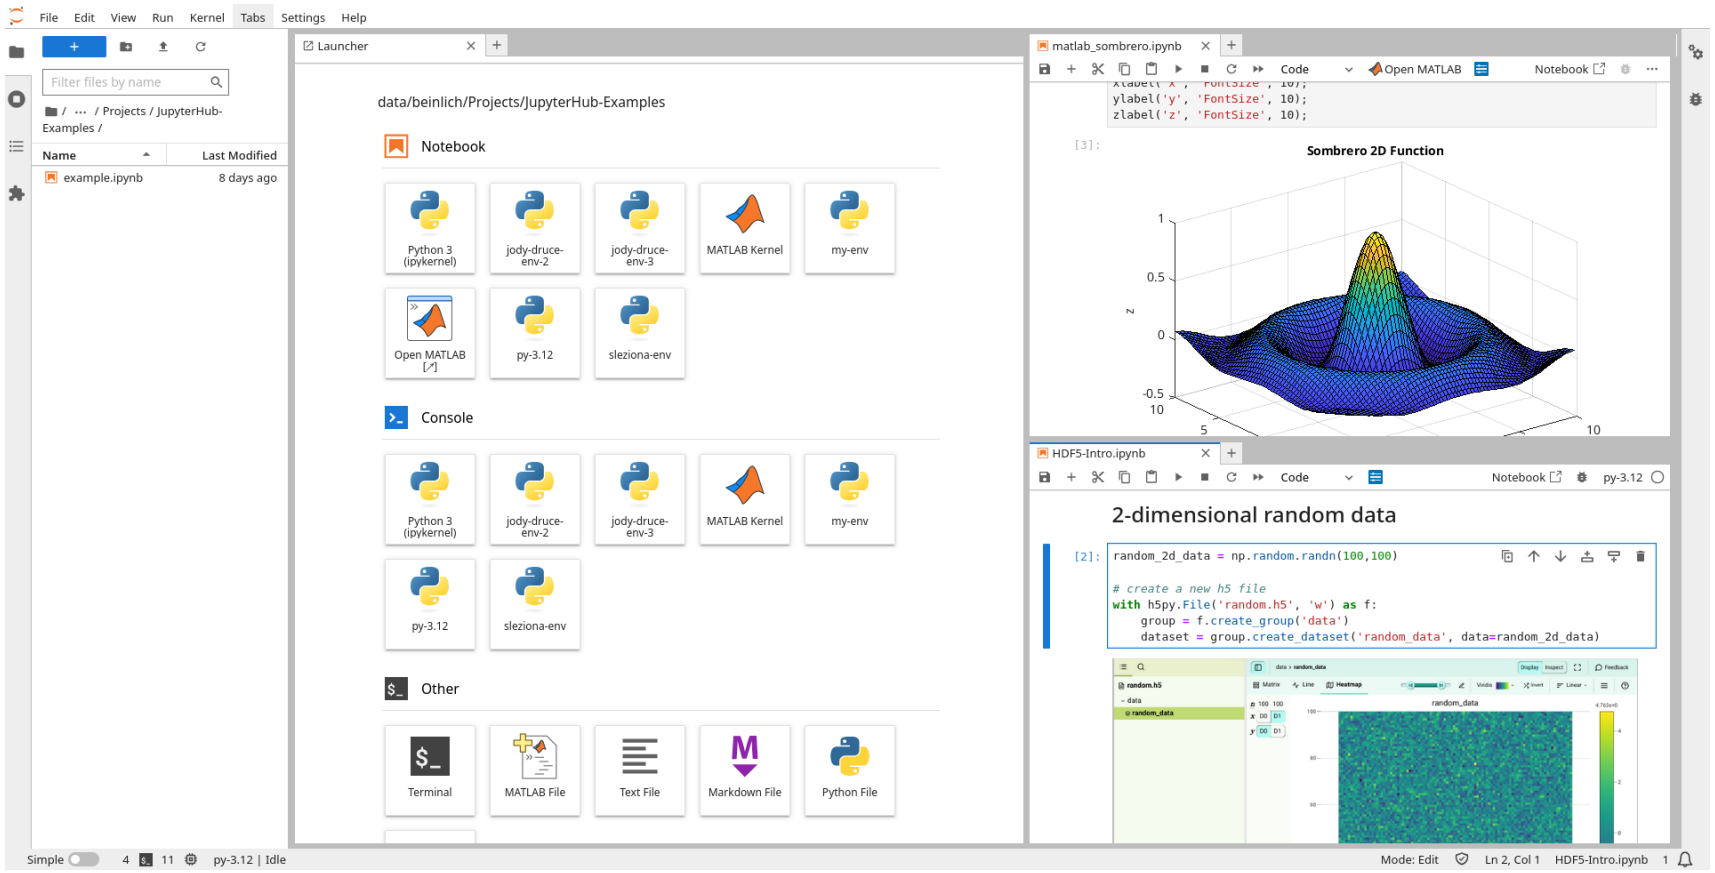

Requires:

- Compute Account
- Browser *(Chrome recommended)*

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/Compute-Jupyterhub
- GIT Wiki (large): https://gitlab.fhi.mpg.de/fhi-hpc/jupyterhub-compute/-/wikis/HOME

JUPYTERHUB-01 – PYTHON & MATLAB SERVER (AND MORE)

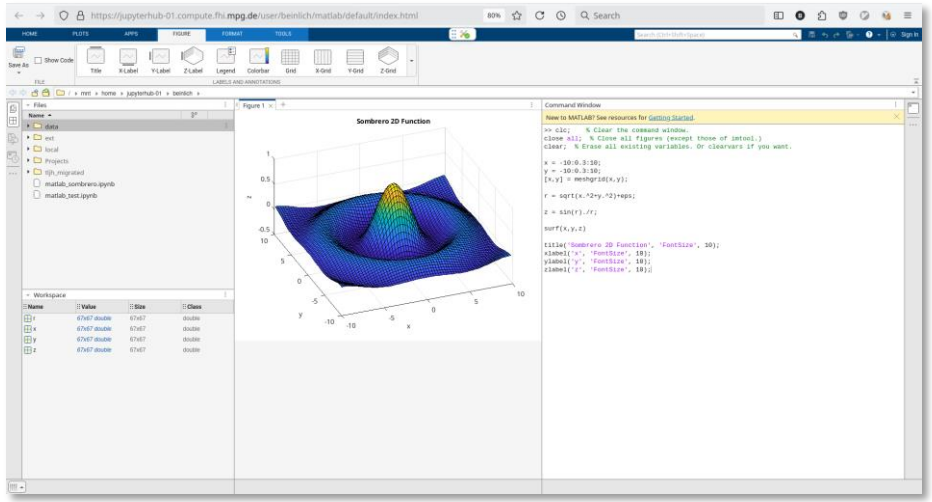

Matlab in a Browser

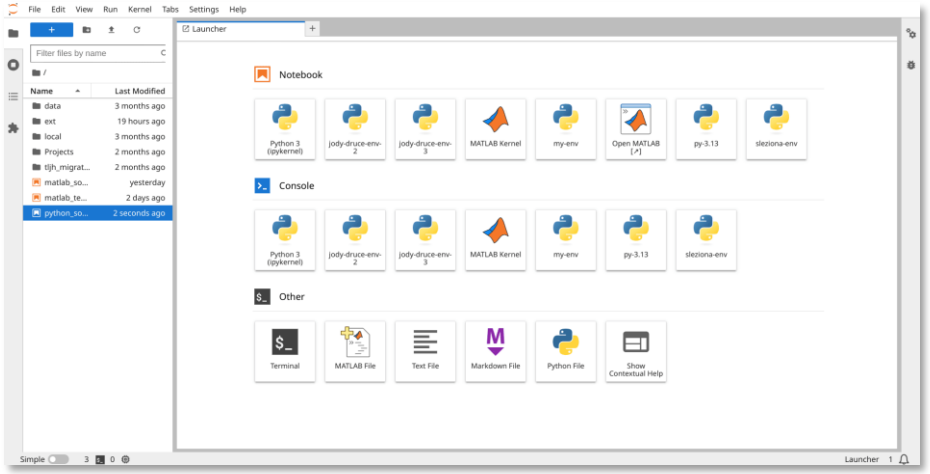

Python in a Browser

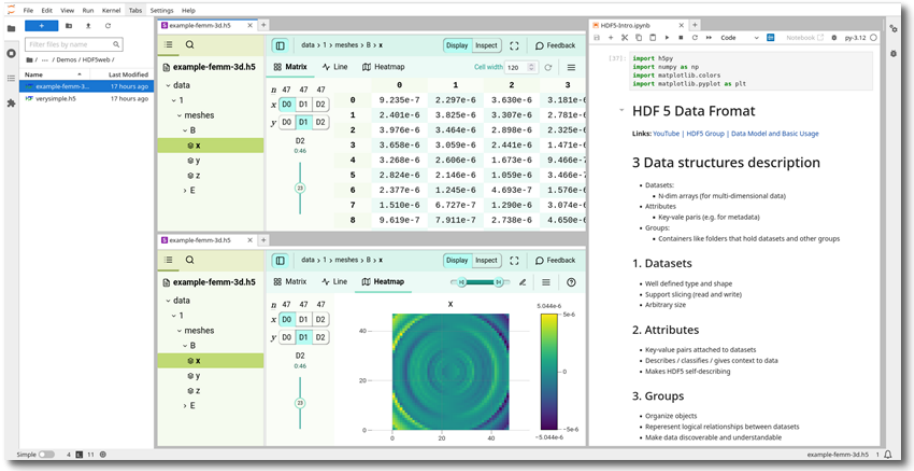

HDF5web in a Browser

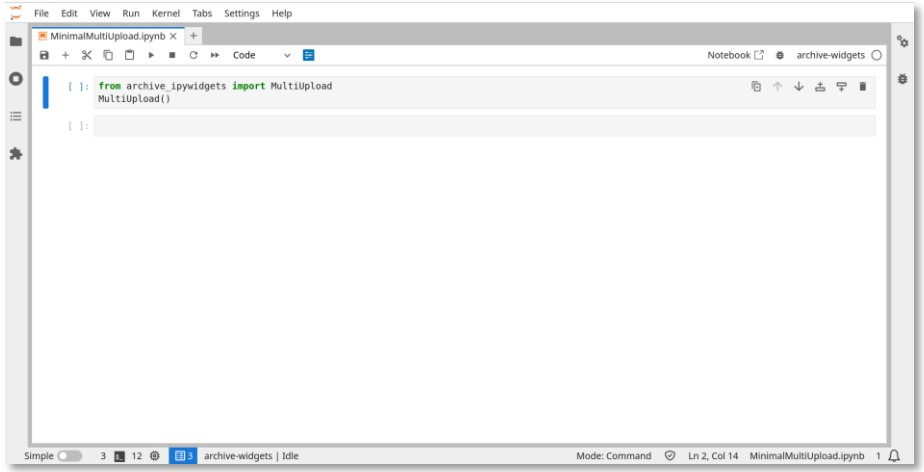

Custom Widgets (e.g. Archive)

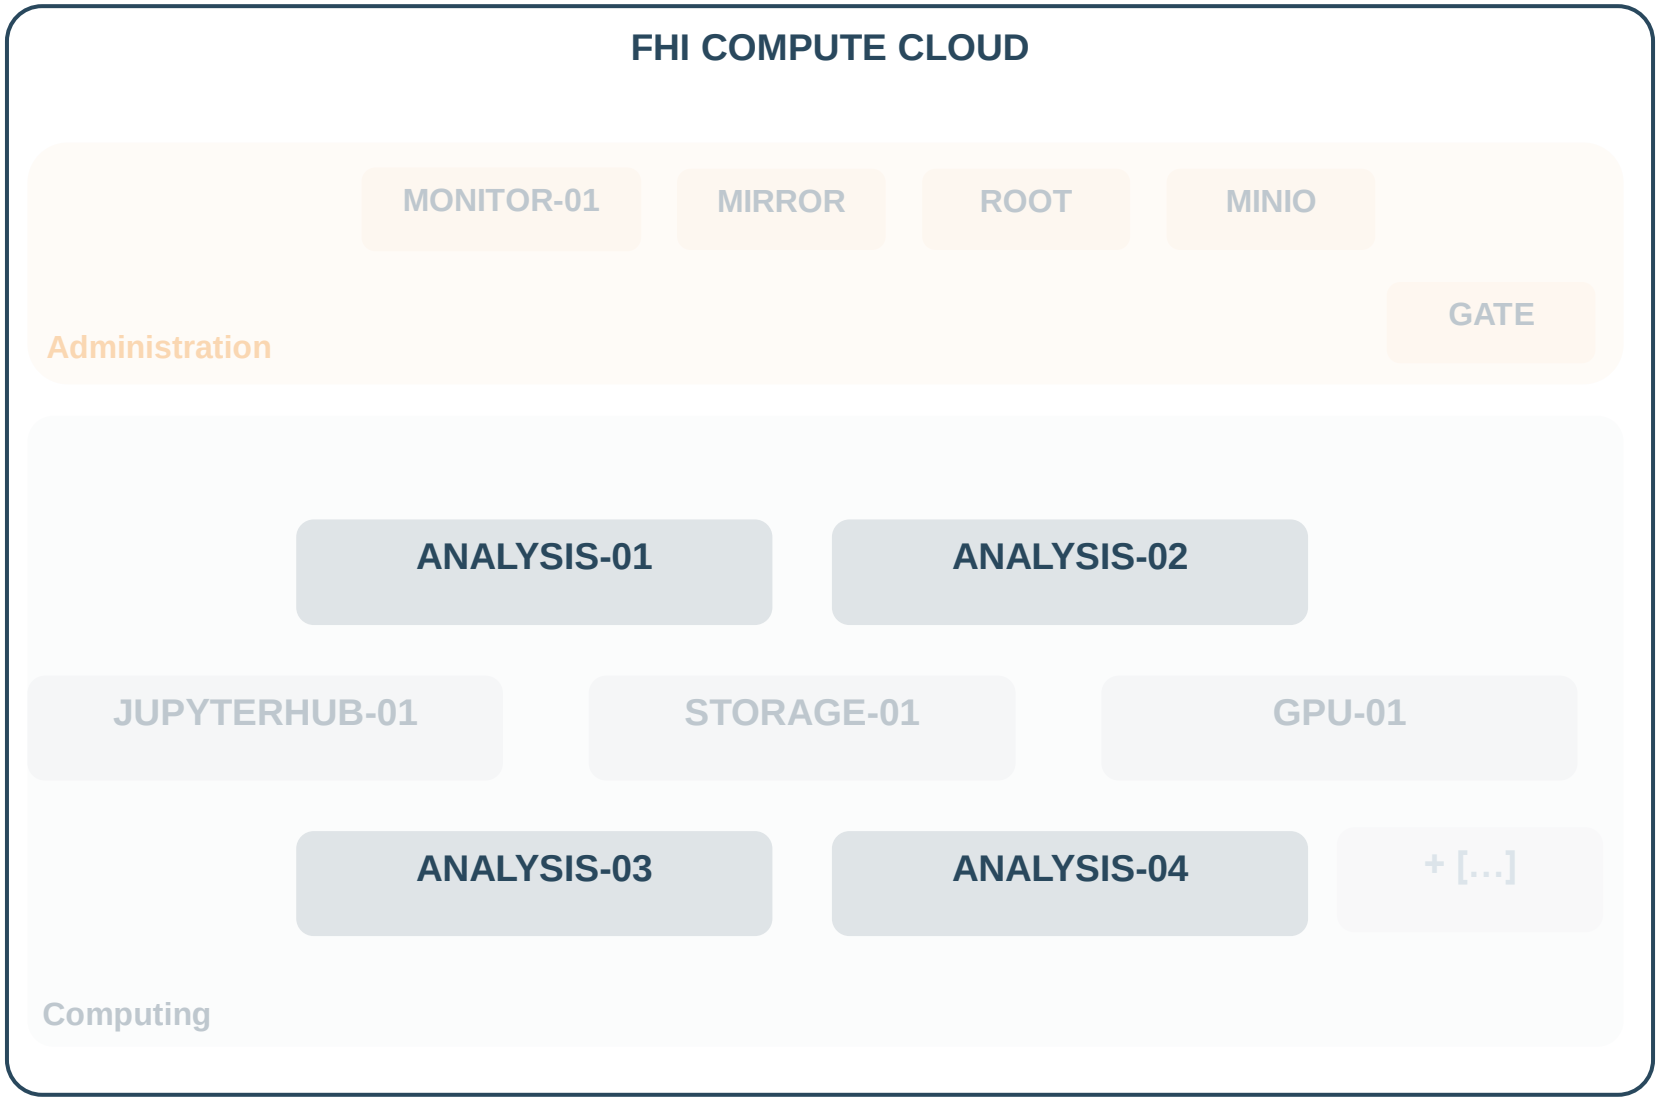
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ [...]
Computing

ANALYSIS-0X – VIRTUAL DESKTOP SERVERS

Use case

- Graphical Applications (MATLAB, Gaussview, VESTA, *maybe no video editing…)*
- Terminal Applications (DFT, Molecular Dynamics)

Resources

- 32 cores, 128GB RAM (max.)
- Central, fhi accessible storage
  - 2TB Home
  - 10TB Data
- Fast local Storage (1TB)

Advantages

- Large resources
- Graphical Applications
- No local installation required (browser sufficient)
- Any Linux-compatible software (debian)
- No terminal / programming experience required.

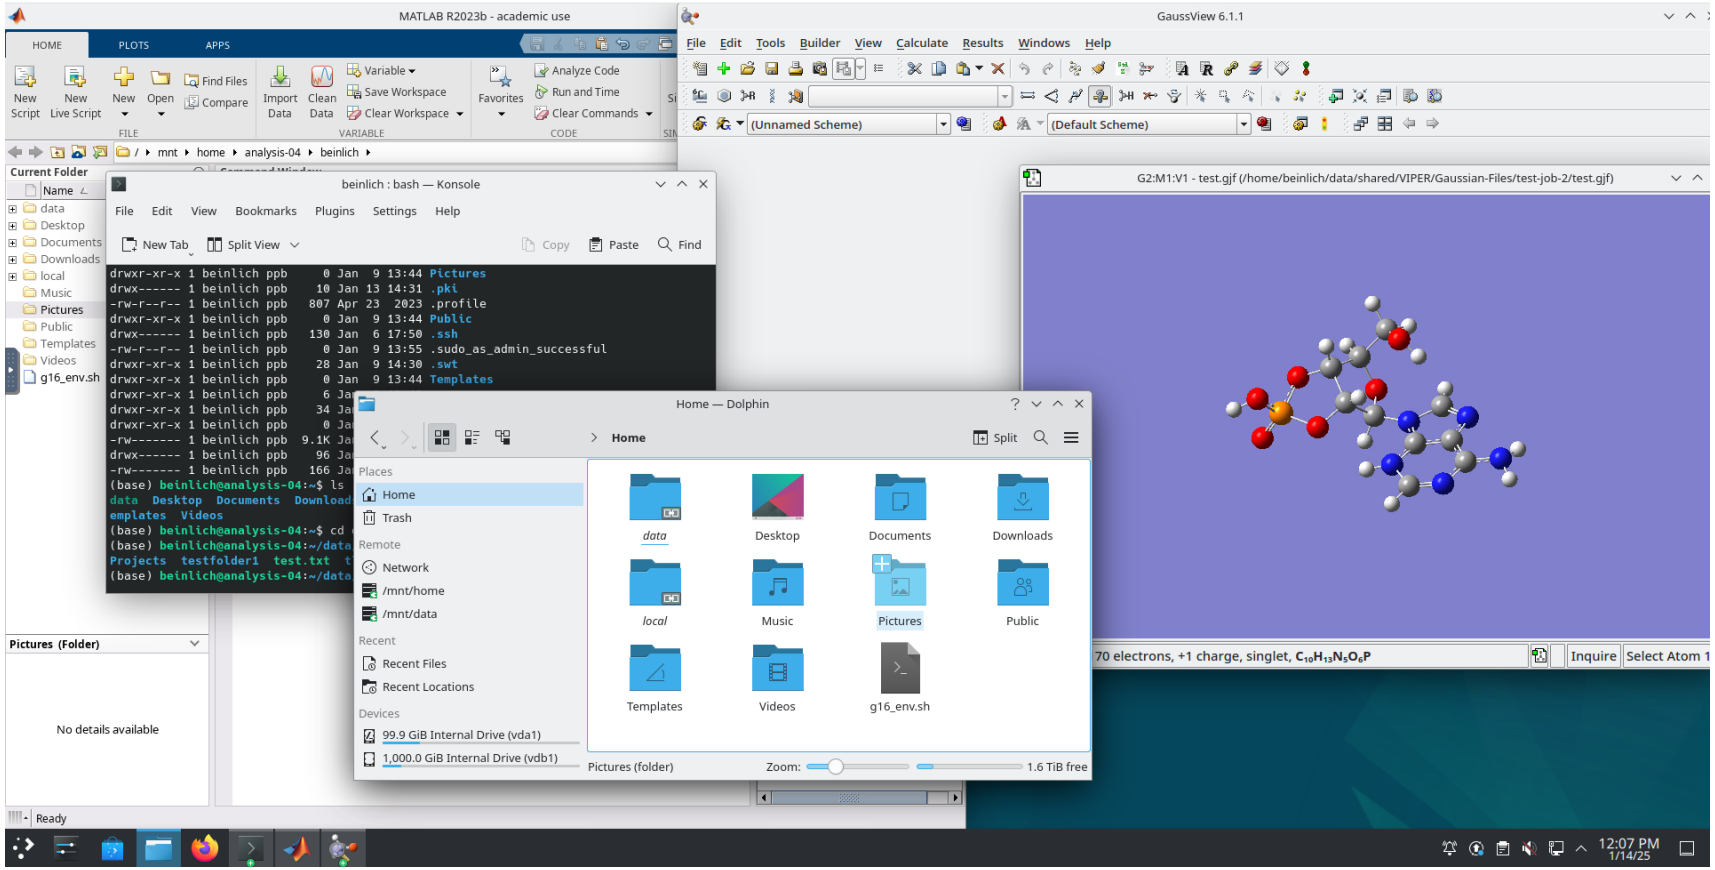

Requires:

- Compute Account + Desktop Session
- Browser *(Chrome recommended)* **or** SSH *(Powershell / Terminal recommended)* *+ VNC Viewer (tigervnc, tightvnc, realvnc)*

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/Compute-Analysis
- GIT Wiki (large): https://gitlab.fhi.mpg.de/fhi-hpc/analysis-compute/-/wikis/HOME

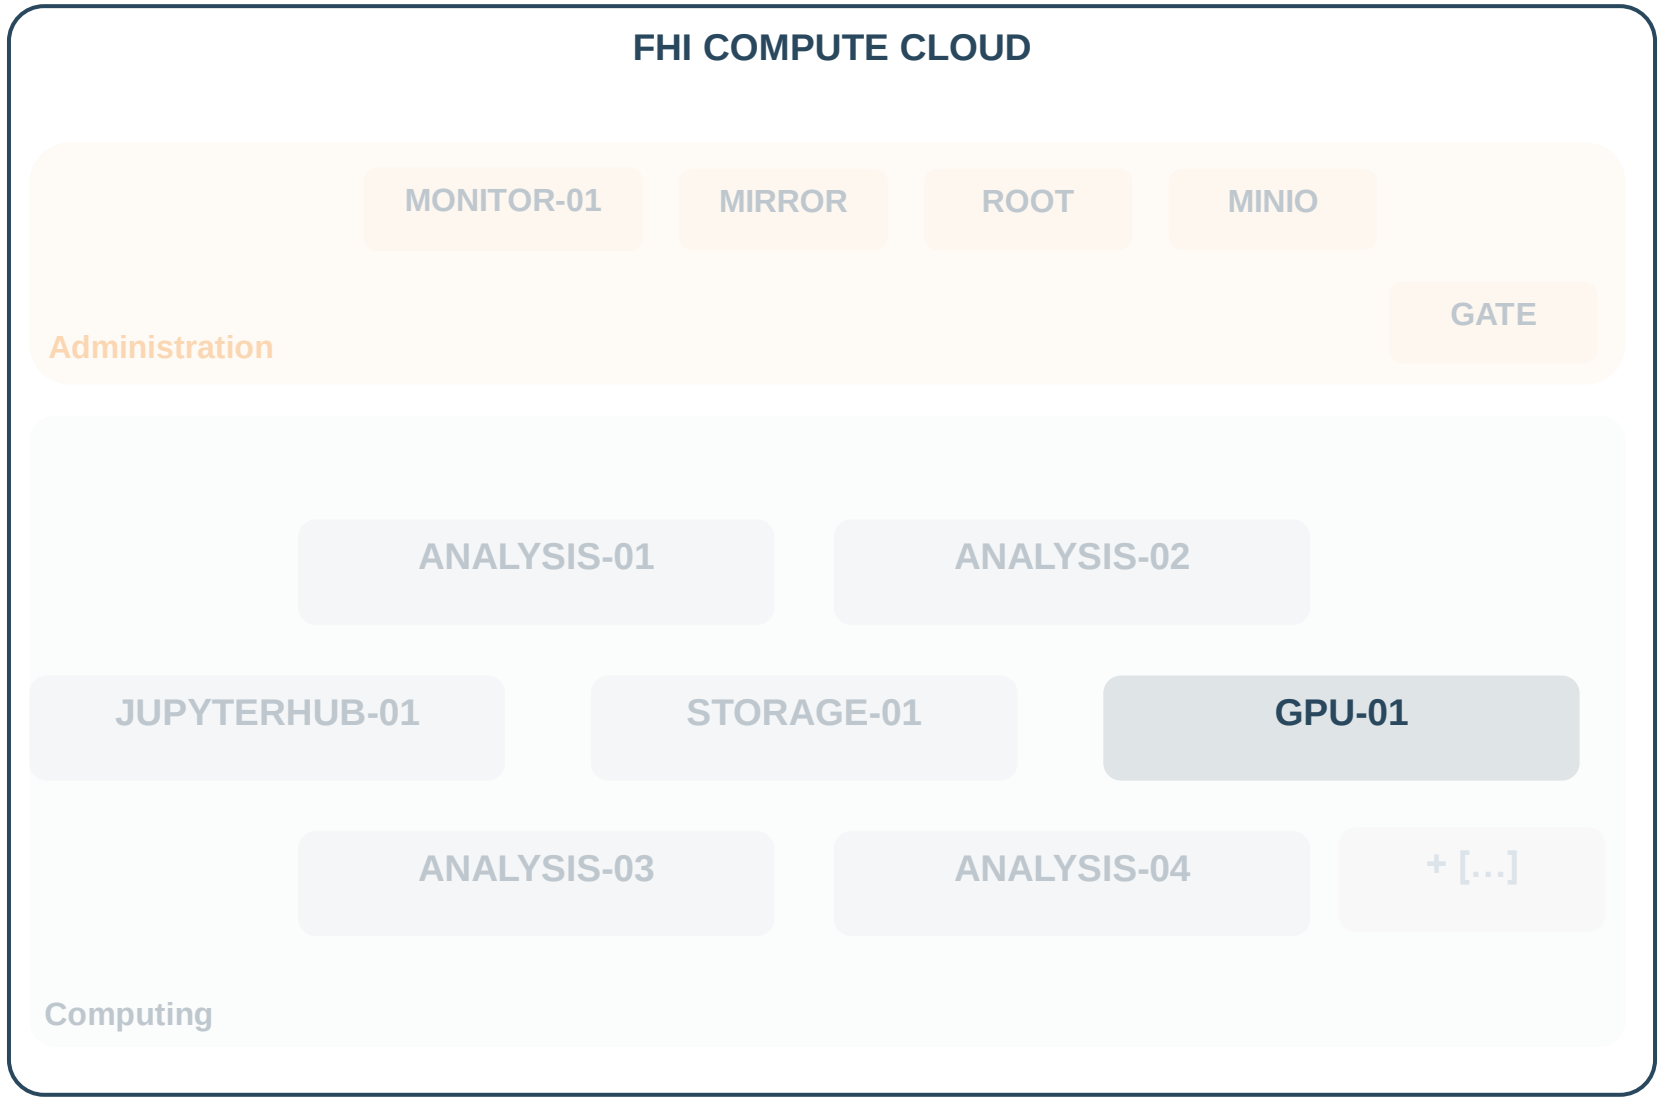
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

GPU-01 – VIRTUAL DESKTOP + NVIDIA TESLA A100 GPU

Use case

- see Analysis-0X
- + Machine Learning / AI
- + Accelerated Computing (GPU-accel.)

Resources

- 32 cores, 128GB RAM,
  NVIDIA Tesla A100 GPU (80GB)
- Central, fhi accessible storage
  - 2TB Home
  - 10TB Data
- Fast local Storage (1TB)

Advantages

- see Analysis-0X
- + AI / ML / GPU acceleration

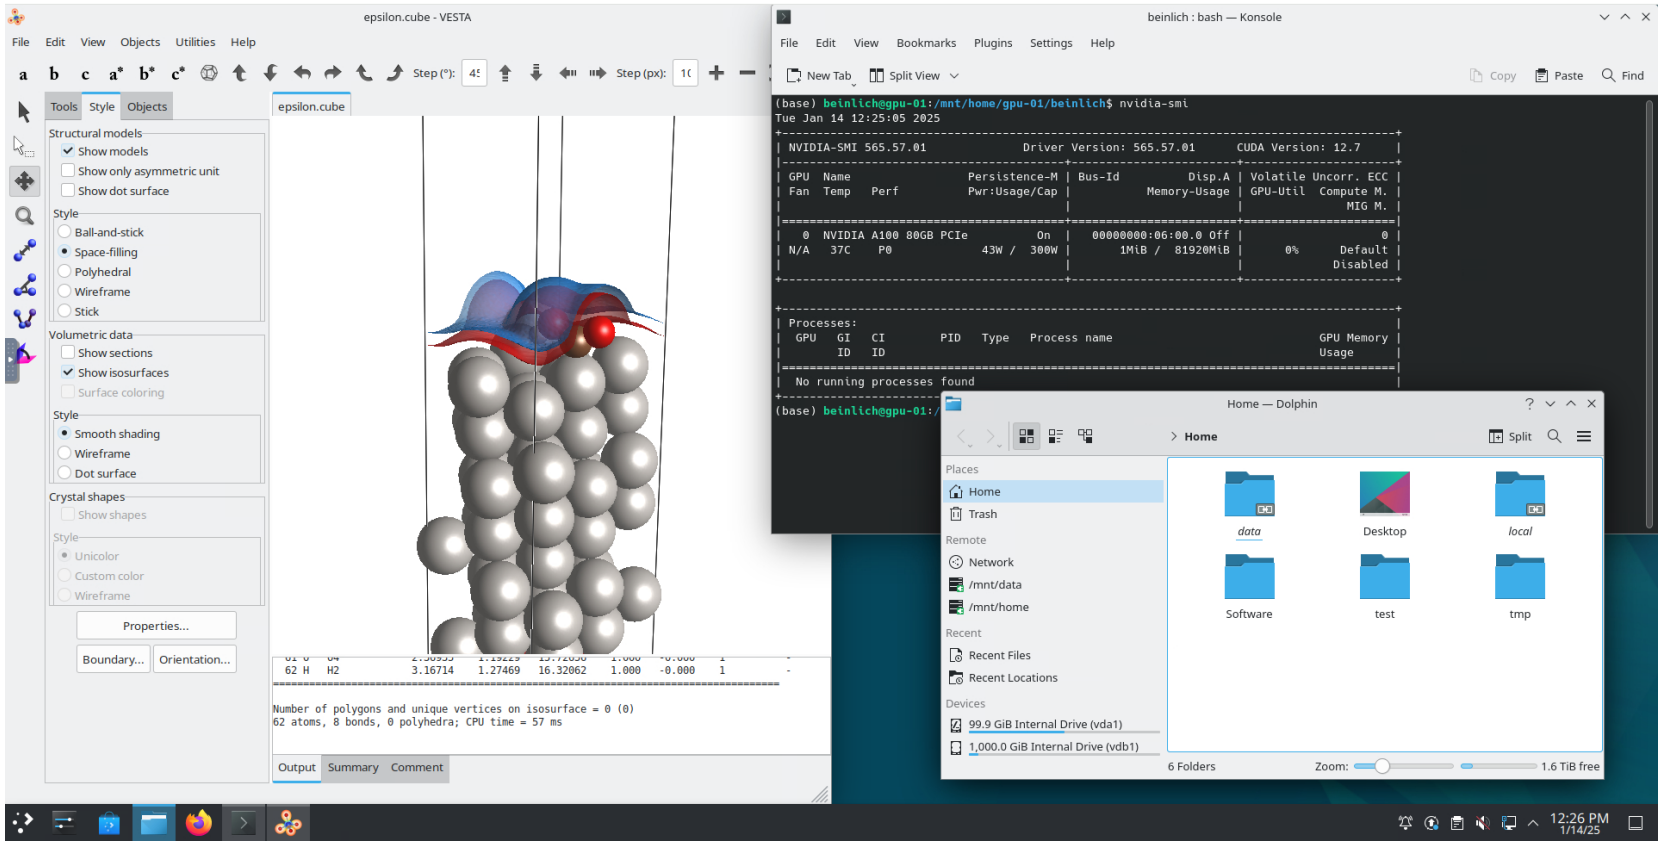

Requires:

- Compute Account + Desktop Session
- Browser *(Chrome recommended)*
  or SSH *(Powershell / Terminal recommended)*
  + VNC Viewer (tigervnc, tightvnc, realvnc)

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/Compute-Analysis
- GIT Wiki (large): https://gitlab.fhi.mpg.de/fhi-hpc/analysis-compute/-/wikis/HOME

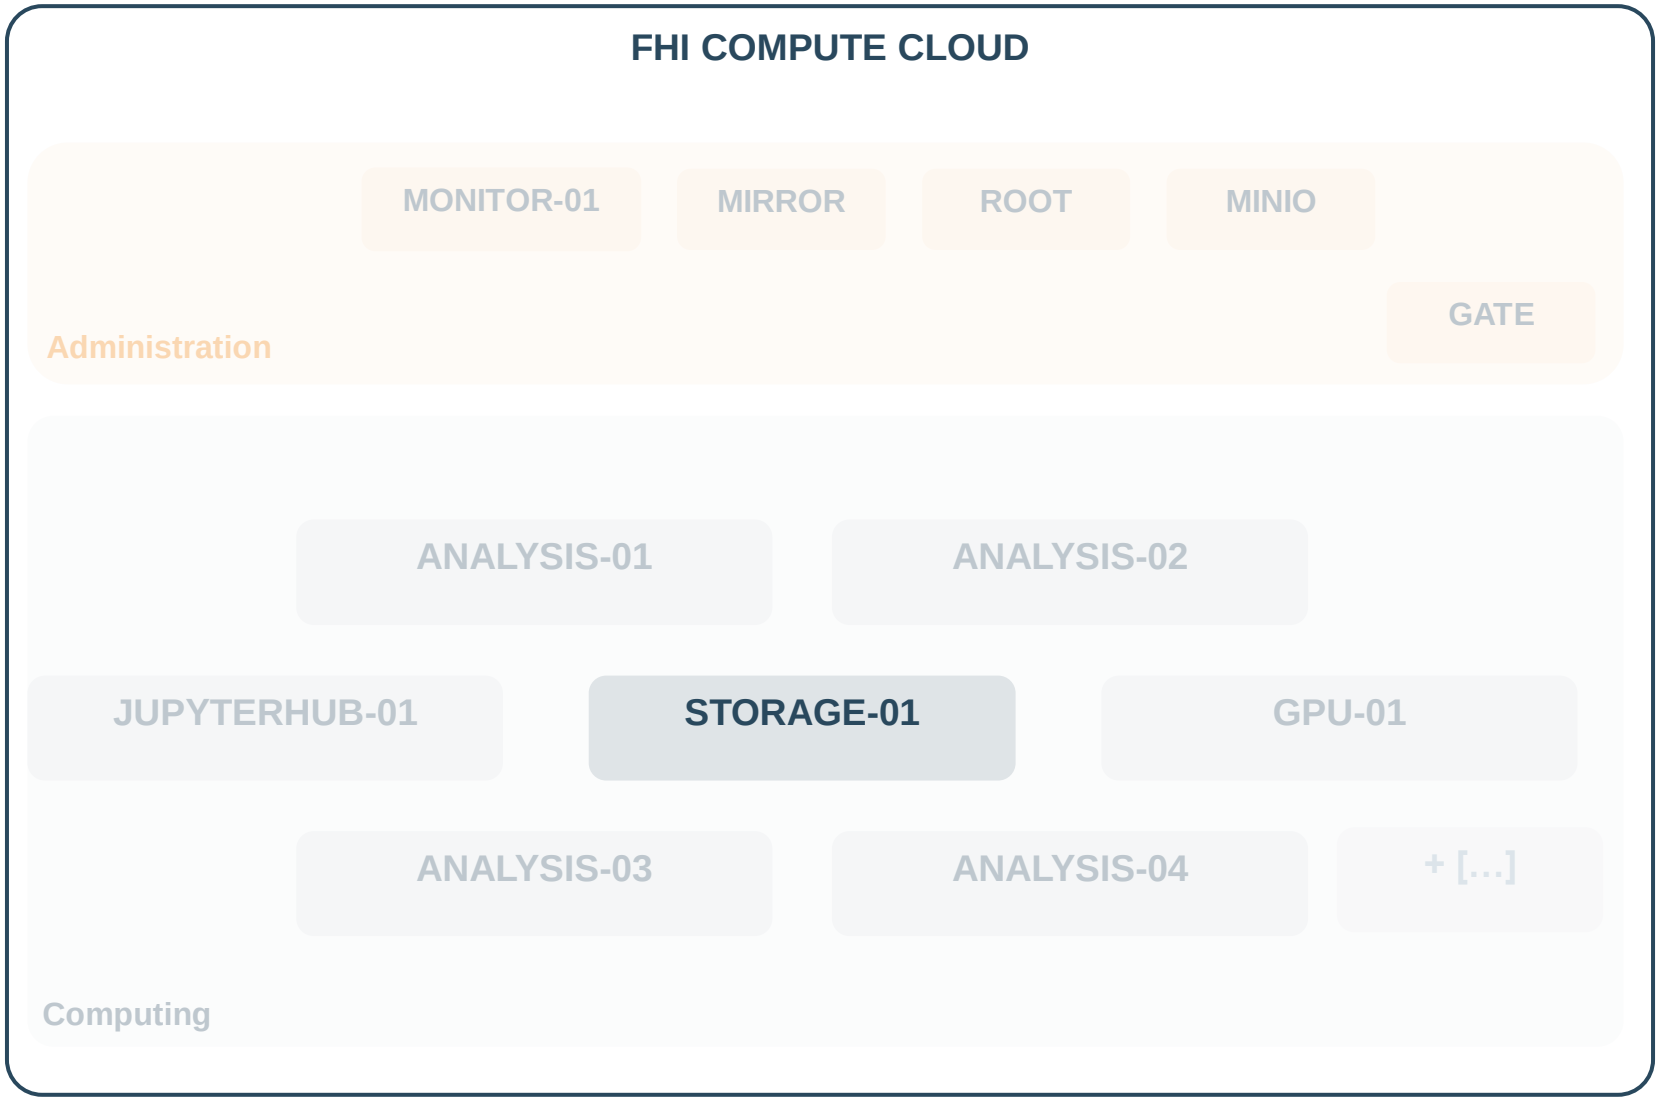
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

STORAGE-01 – CENTRAL STORAGE FOR COMPUTE SERVERS

Use case

- Transfer files to/from the Compute Servers
- Manage all your compute server files in one place

Resources

Shared filesystems (snapshotted, backed-up)

- 2TB Home
- 10TB Data
- X TB Storage Extensions (>10TB, on demand)

SMB/SFTP/SSH access from FHI

Advantages

- Simple File/Data Management
- Backed up data
- Integration of large data buckets possible on demand (>100TB)

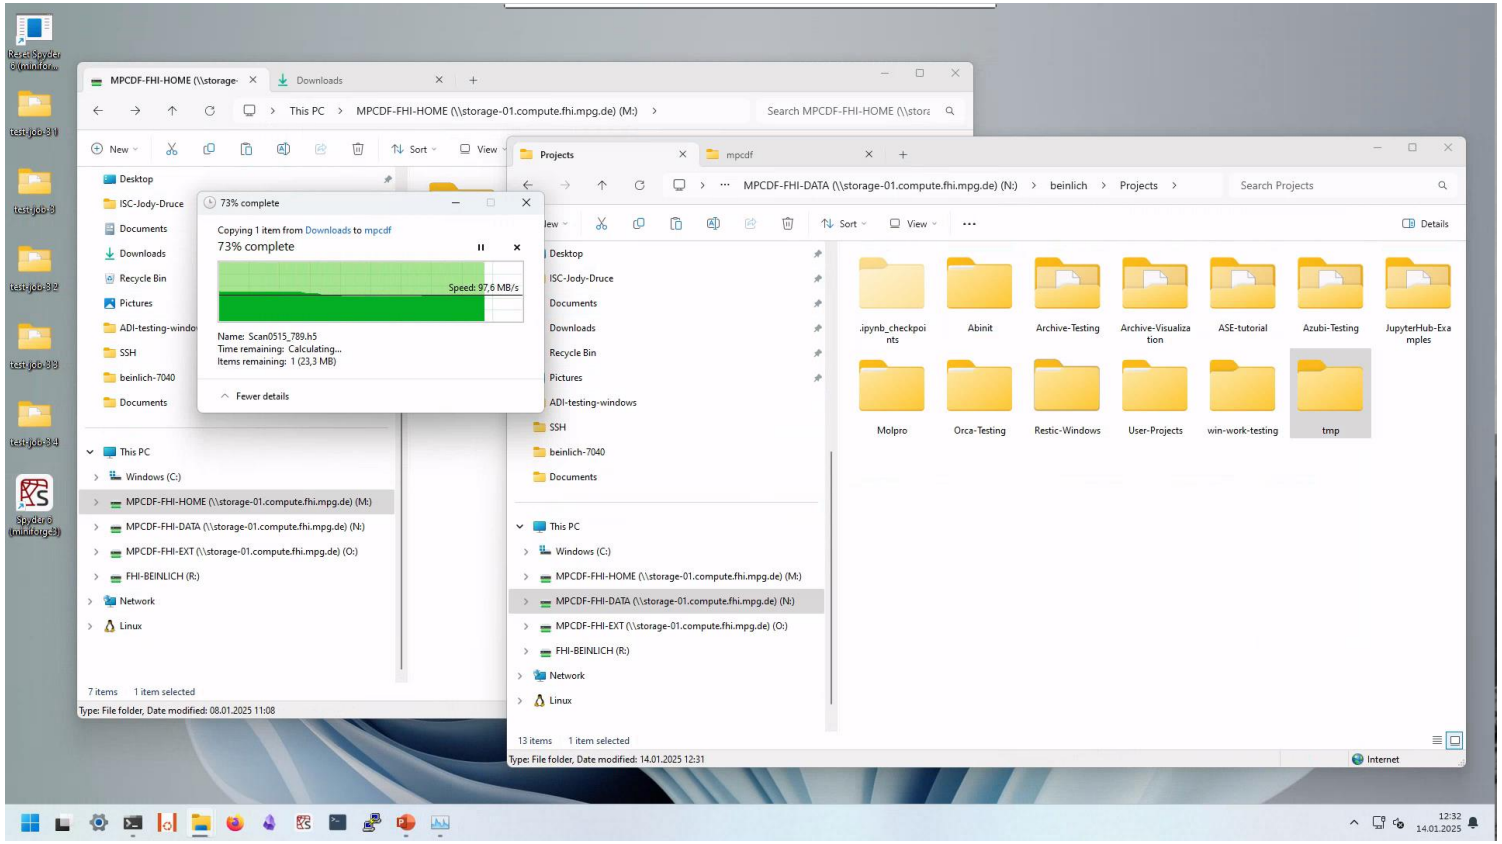

Requires:

- Compute Account
- File Explorer *(Windows Explorer or SMB/SFTP-enabled file explorer)*

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/Compute-Storage
- GIT Wiki (large): https://gitlab.fhi.mpg.de/fhi-hpc/storage-compute/-/wikis/HOME

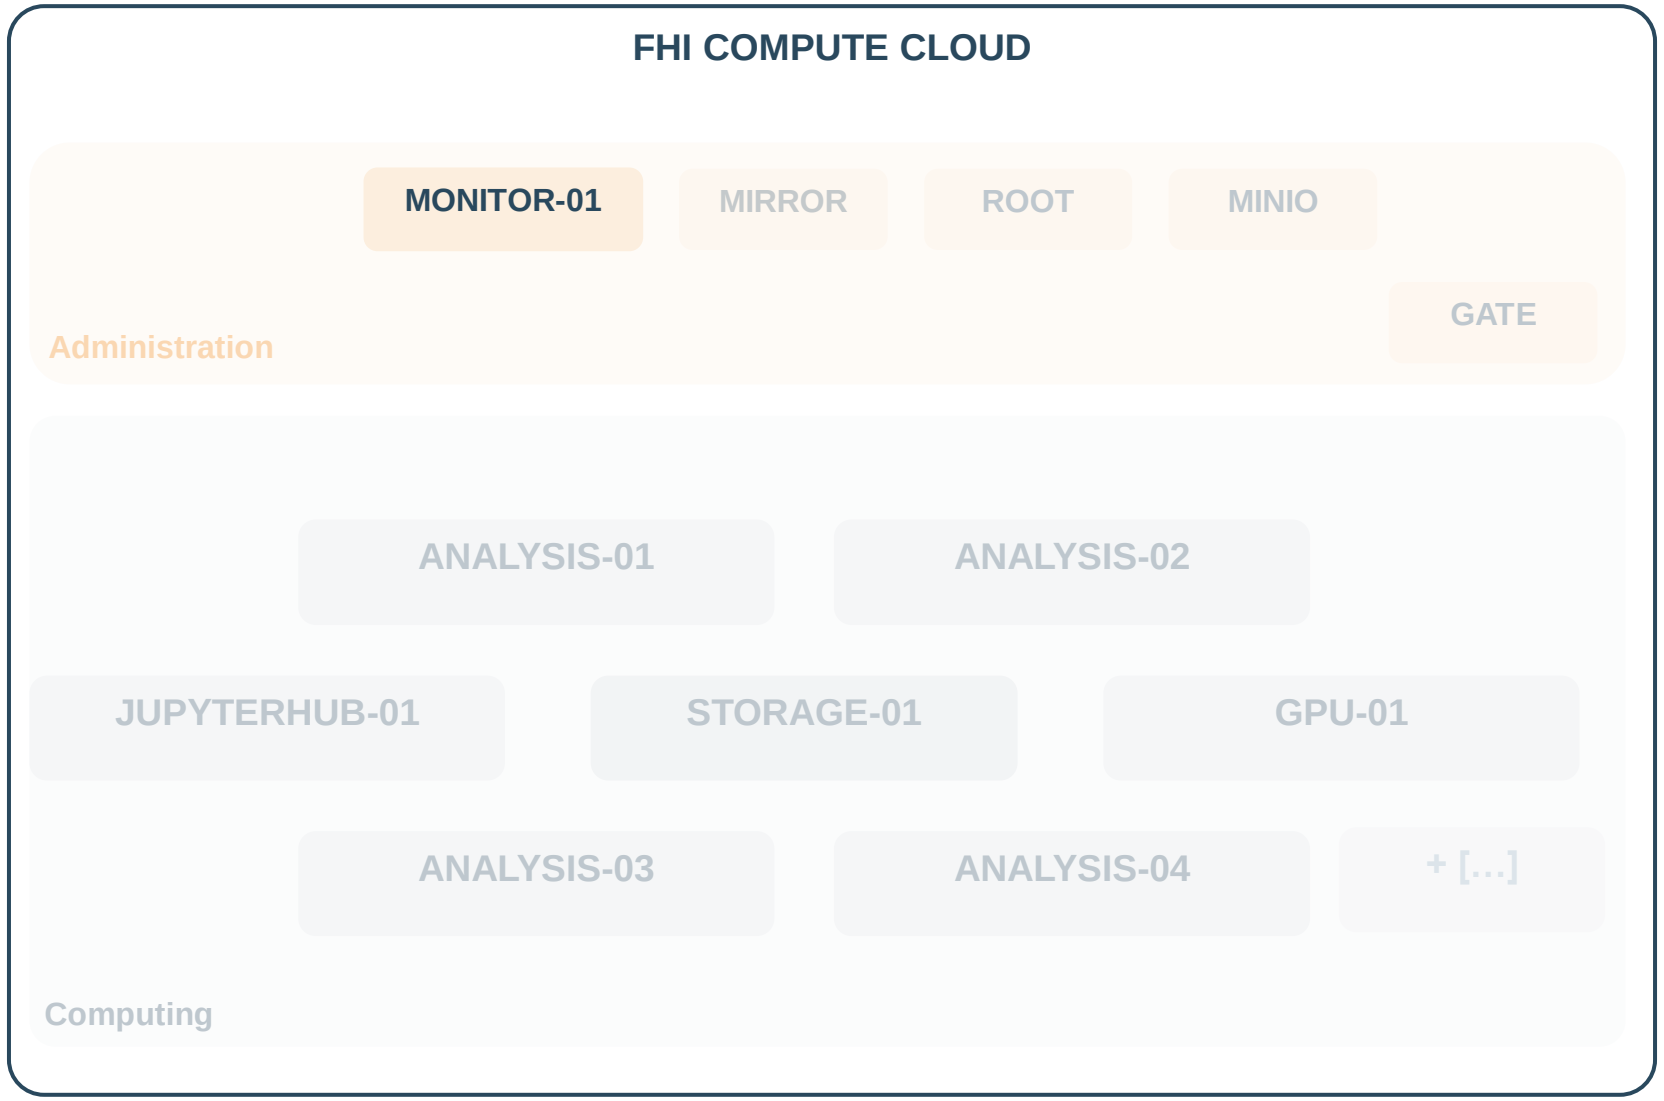
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

MONITOR-01 – CENTRAL RESOURCE MONITORING

Use case

- Monitor current state and usage of compute servers and storage
- Monitor current usage of your sessions
- Monitor Viper Compute Budgets

Resources

- Compute Server Monitoring:
- UP state, CPU & Memory usage, Network traffic, Storage Space
- Viper Usage Monitoring

Advantages

- Easy monitoring of cpu and memory consumption to prevent issues or optimize your processes
- Check that you stay within the Viper budget

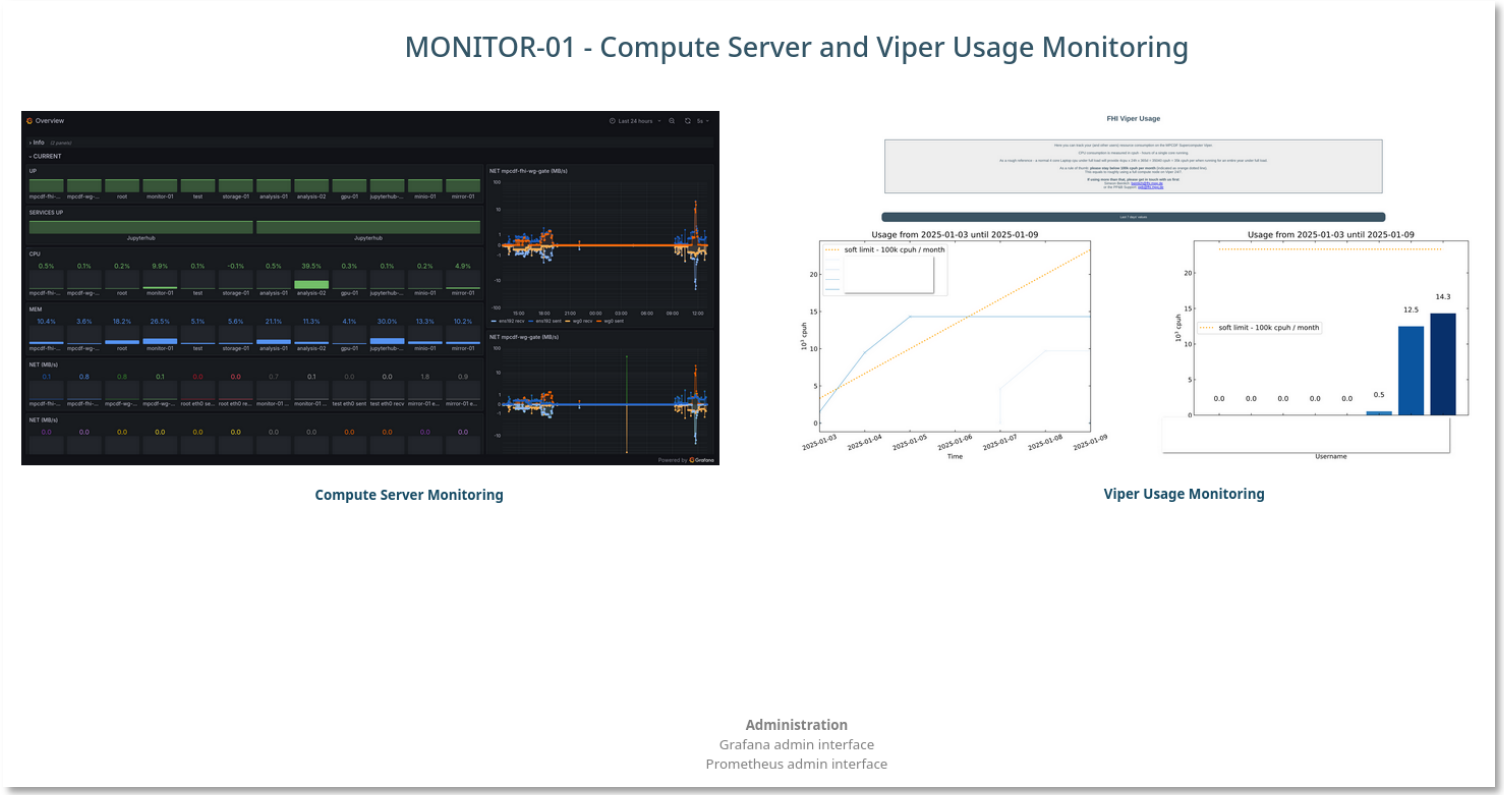

Requires:

- Browser *(Chrome recommended)*

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/HPC

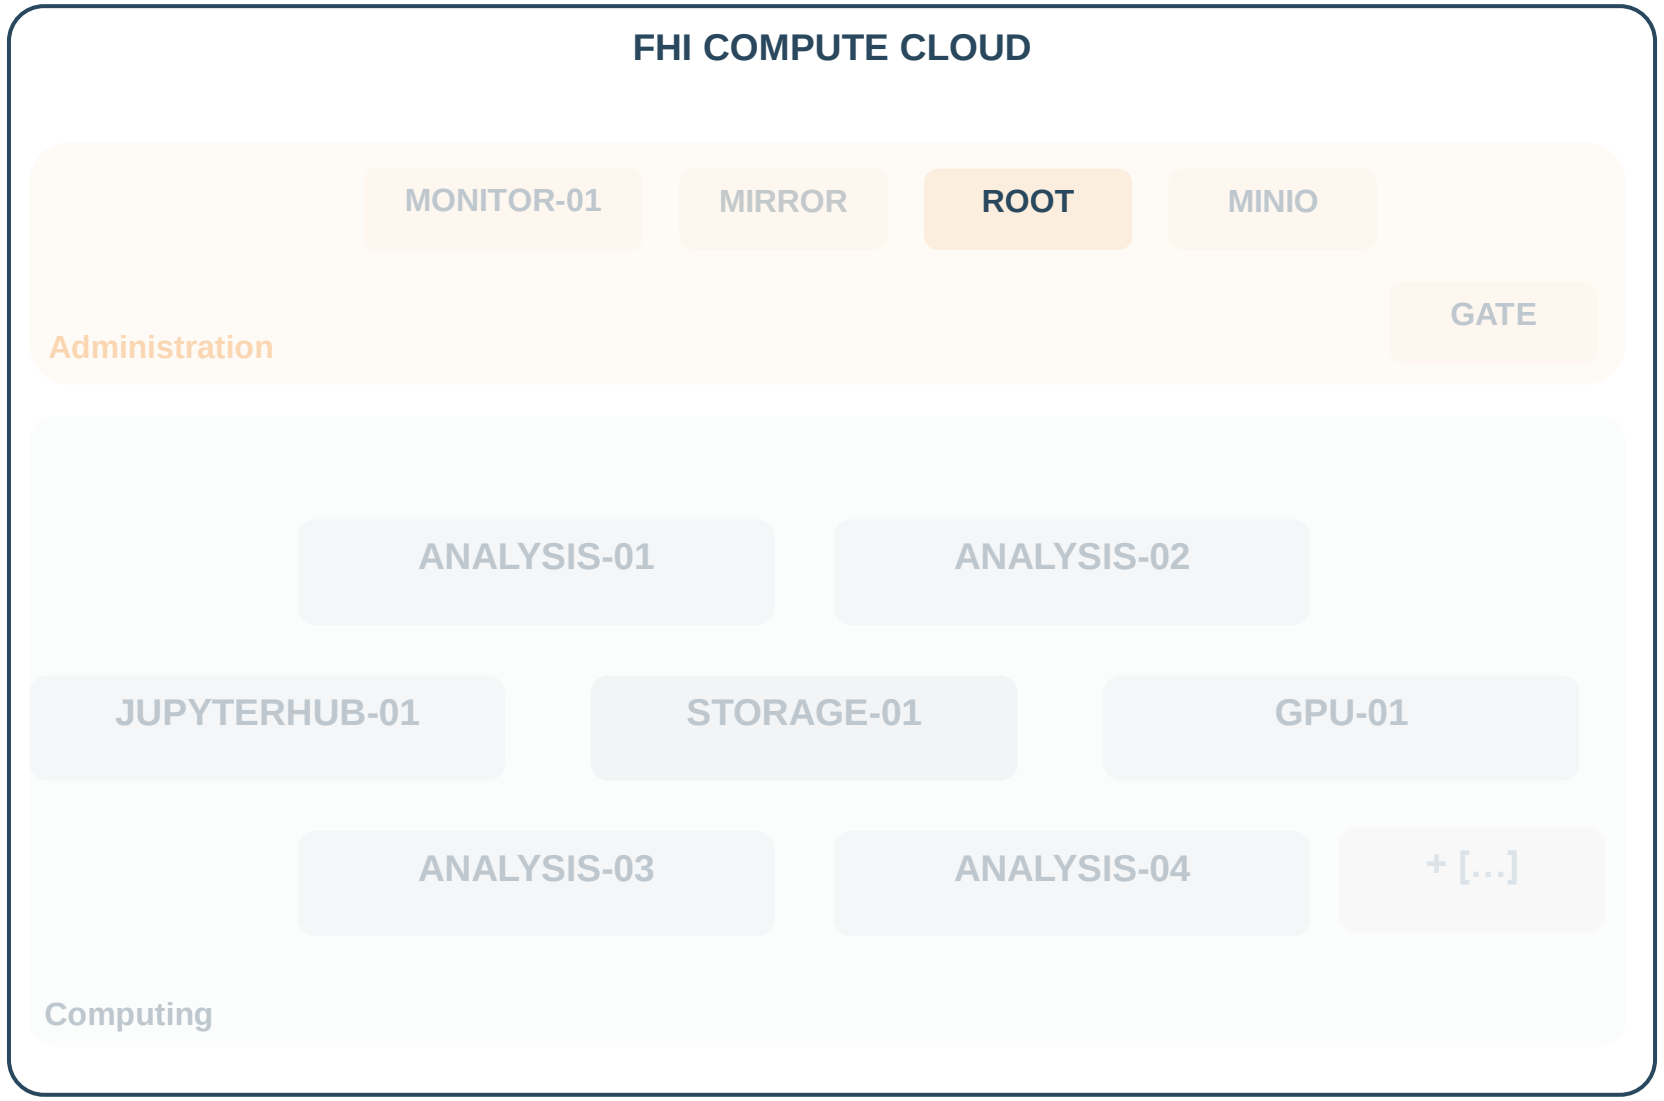
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

ROOT – CENTRAL ACCOUNT MANAGER

Use case

- Account Creation
- Password Change / Reset
- Desktop Session (Re-)Start

Advantages

- Simple restarting of VNC servers
- Simple Password Change for all compute services

Requires:

- Browser *(Chrome recommended)*

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/HPC

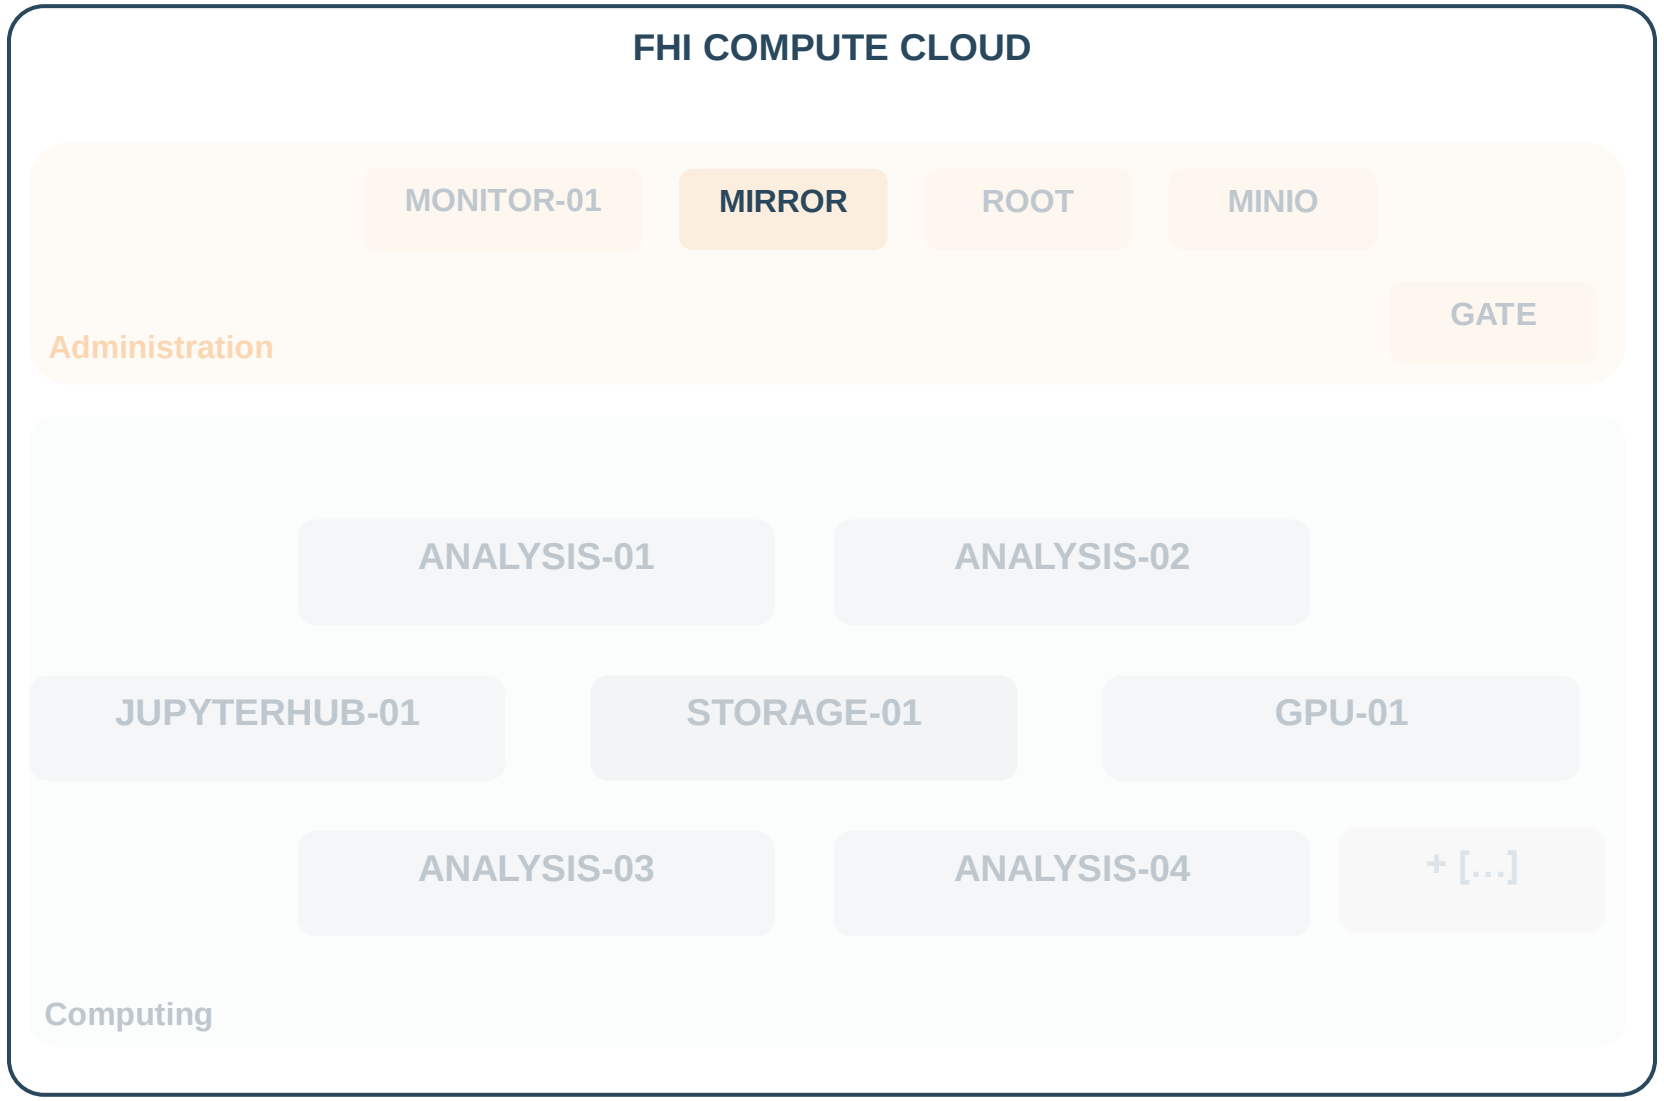
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

MIRROR – LARGE-SCALE FILE TRANSFER (ASYNCHRONOUS S3 SYNC)

Use case

- Near real-time, fast synchronization of FHI S3 buckets to FHI Compute Storage (one-way, ~2.5-8GBit/s)
- E.g. for experiments that produce large data sets that should be processed at Compute Servers
- Large Data Stores (> 100TB), Large Files (>10MB)
- Full Data Access at Compute Servers
- Single- or multi-user access on entire data set or subfolder
- FHI → MPCDF or MPCDF → FHI

Resources

- Based on S3 Storage (200TB+)
- Local SSD Cache possible (r/w, soon)
- Filesystem-like access from Compute Servers
- SMB (Windows network storage) convenience access from FHI (<1GBit/s)

Advantages

- Fast, large data store

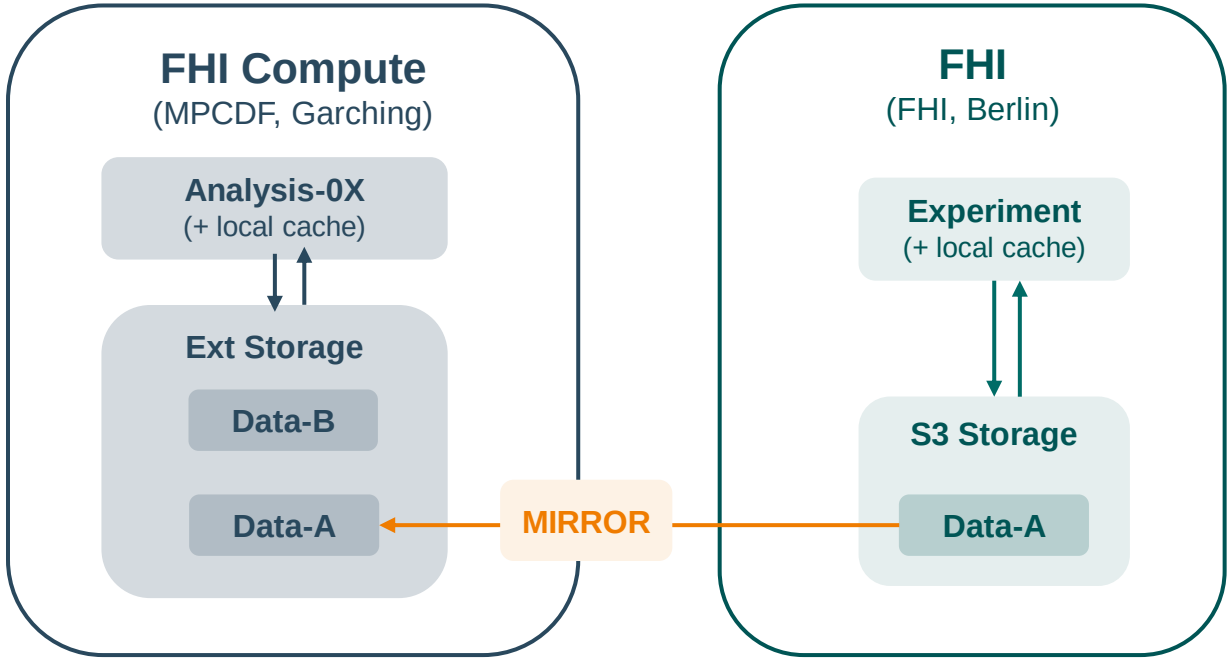

Requires:

- FHI S3 Bucket *(on request)*
- MPCDF Ext Storage Access *(on request)*

Documentation:

- TBA

TESTING PHASE !

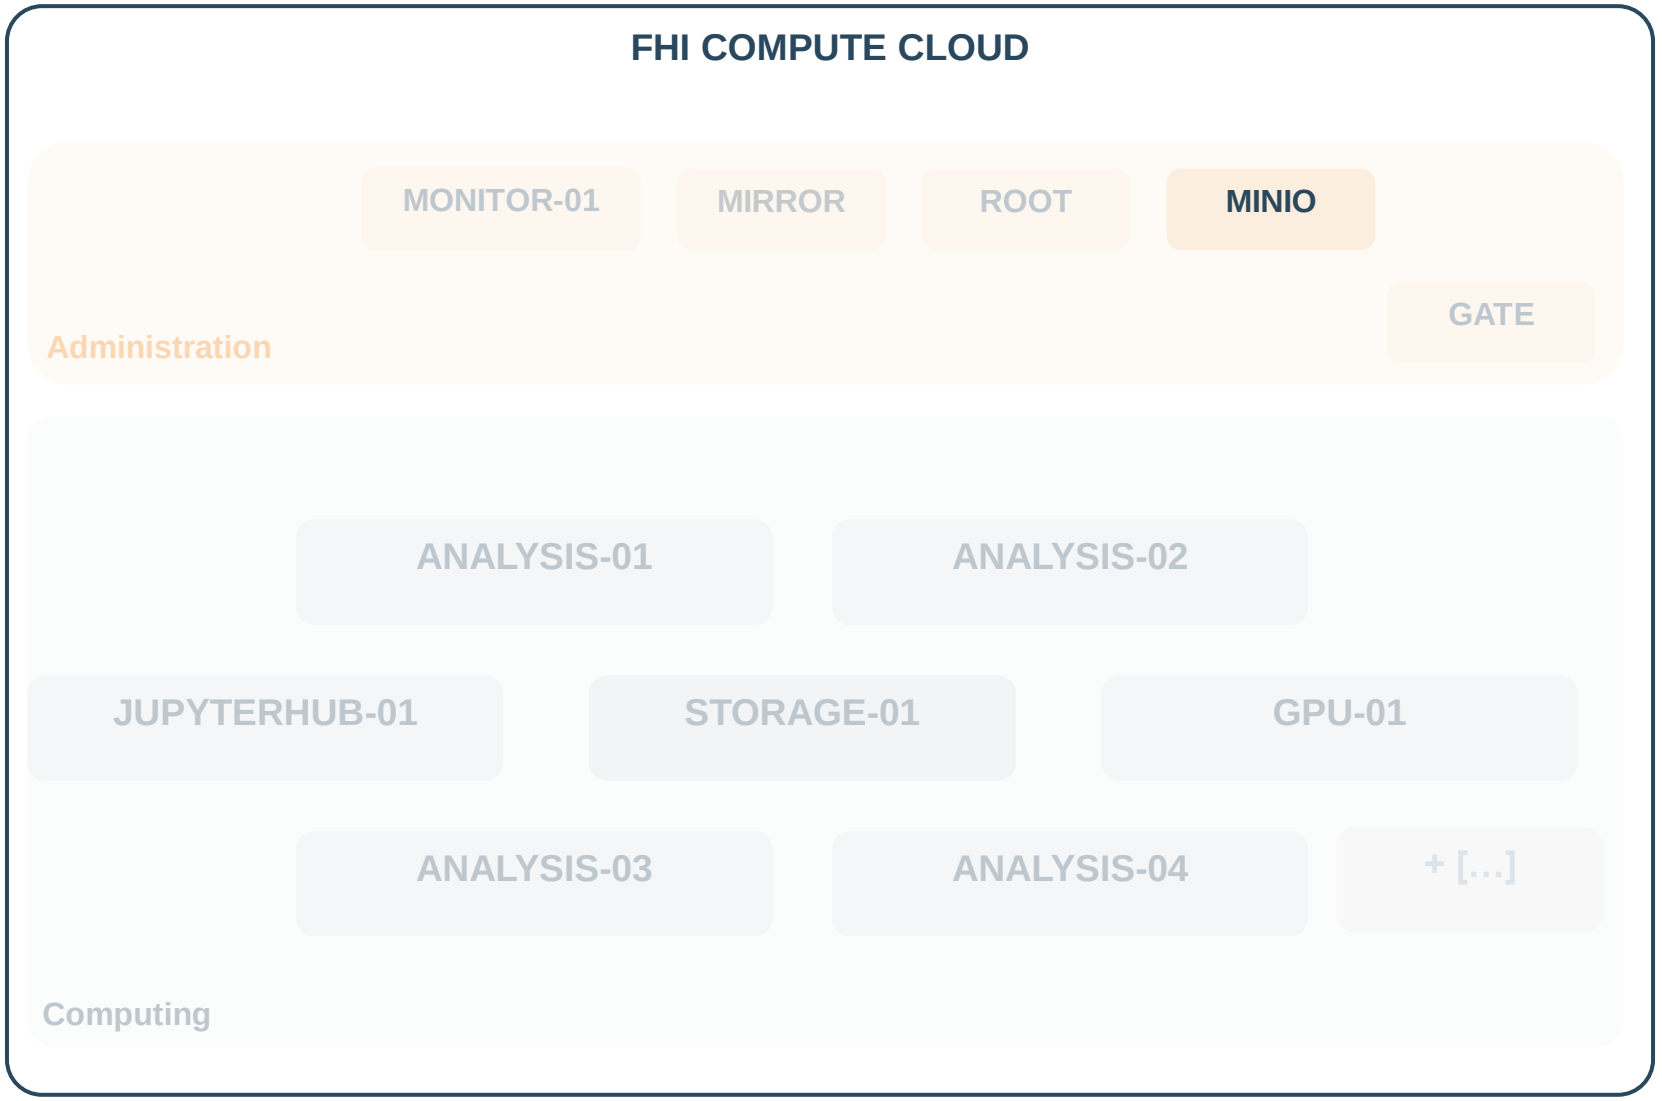
FHI COMPUTE CLOUD
MONITOR-01
MIRROR
ROOT
MINIO
GATE
Administration
ANALYSIS-01
ANALYSIS-02
JUPYTERHUB-01
STORAGE-01
GPU-01
ANALYSIS-03
ANALYSIS-04
+ […]
Computing

MINIO-01 – S3 USER BUCKETS

Use case

- S3 User/Project buckets similar to FHI S3 buckets

Resources

- 80TB (more on demand)

Advantages

- Known S3 interface as with the FHI S3 Storage
- No File Size restrictions
- Direct S3 API and Webinterface access with per-user credentials

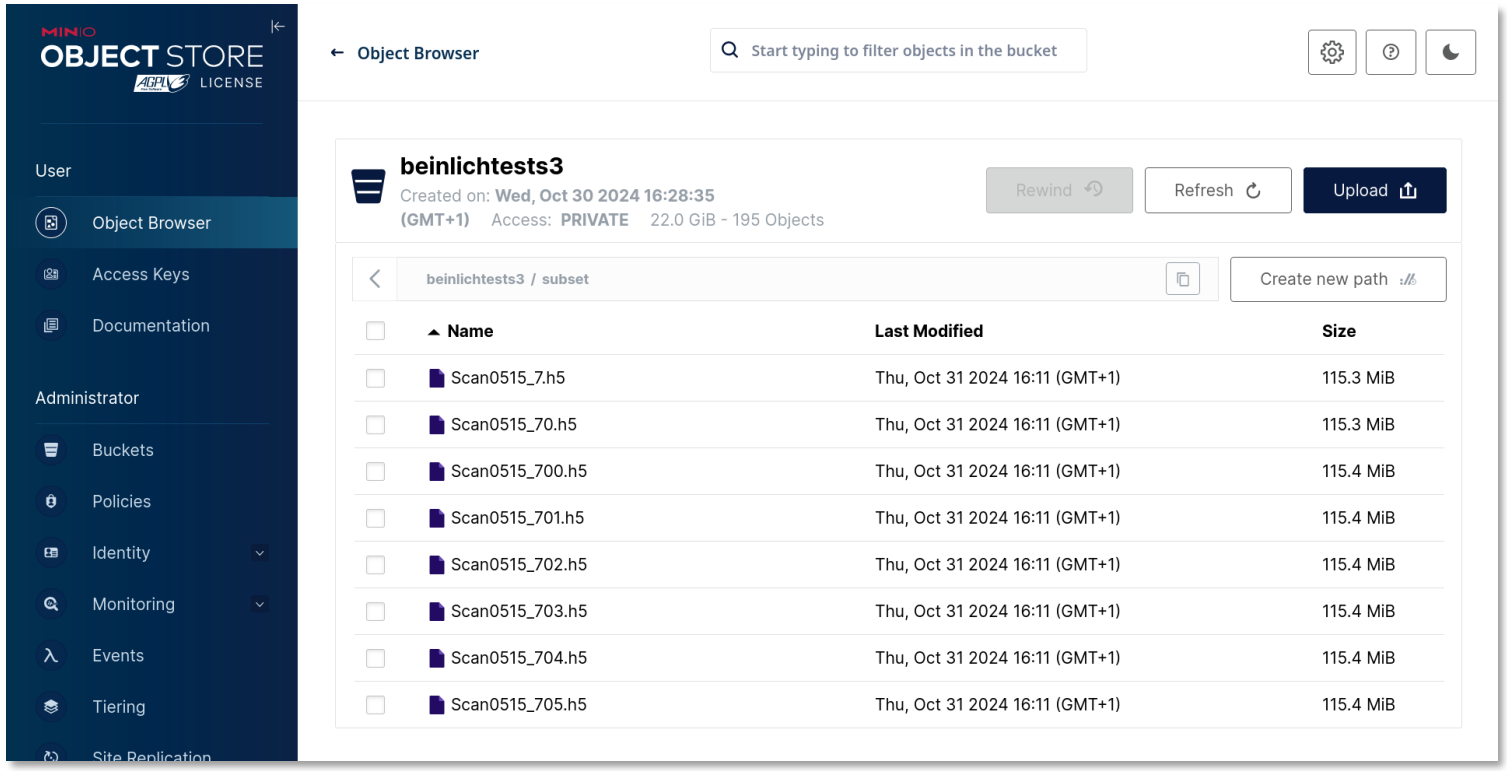

Requires:

- Browser / Minio Client / S3 client / rclone / restic / …
- FHI MINIO-01 Bucket *(on request)*

Documentation:

- TBA

VIPER SUPERCOMPUTER

HPC INFRASTRUCTURE AT FHI (& MPG)

FHI (PPB)

FHI COMPUTE Servers

- Large Compute Servers (6x):
  Large CPU + Memory + GPU
- **Virtual Desktops** (Graphical Applications)
- **Jupyterhub** (Python Server)
- Terminal Access
- **Direct Execution** (no Scheduling)
- User friendly

Anything that requires more than a desktop PC

MPG (MPCDF)

VIPER Supercomputer

- Huge Compute Servers (768x):
  Huge CPU + Memory + GPU
- Terminal Access
- **Job Submission** (Scheduling)
- Advanced

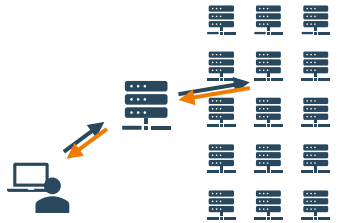

+ RAVEN, SuperMUC, JUWELS, …
(department-level / grant-based)

Large Scale Analyses / Processing
& Ab-Initio Simulations

VIPER SUPERCOMPUTER – MPCDF

Use case

- Large-Scale Data Analysis & Processing
- Machine Learning / AI
- Ab-initio Simulations (DFT etc.)
- Anything massively parallelizable

Resources

- SLURM scheduler
- Massive module system
  (Compiler, Software, …)
- 768 CPU compute nodes:
  - 609 x 128c/512g
  - 90 x 128c/768g
  - 66 x 128c/1024g
  - 3 x 128c/2304g
- 340 APU compute nodes:
  (GPU accelerated CPUs)
  - 2 x AMD MI300A + 2 x 128g HBM3

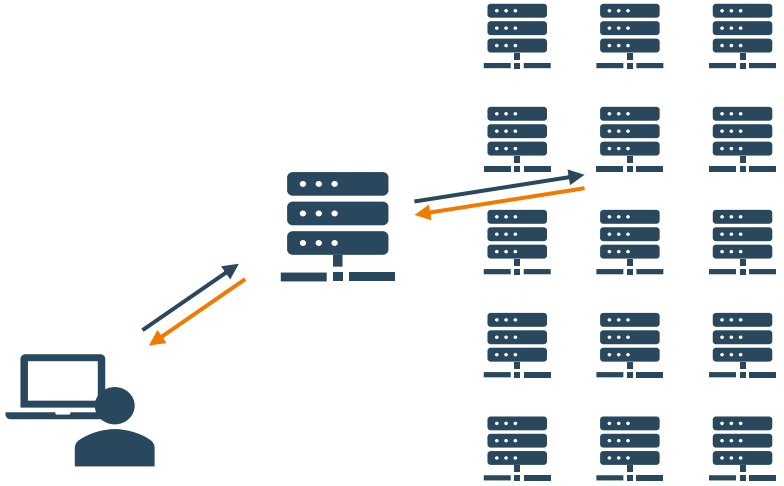

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/HPC
- GIT Wiki (large): https://gitlab.fhi.mpg.de/fhi-hpc/viper/-/wikis/home

VIPER SUPERCOMPUTER – MPCDF

Advantages

- Huge computing power
- State-of-the art software & hardware stack

Common Workflow

- Transfer data to Viper
  e.g. *WinSCP, Rsync, …*
- Log in via SSH
  e.g. *Powershell, Putty, …*
- Create a job script
  e.g. *job.sh*
- Submit the job
  e.g. *sbatch job.sh*
- Check the Queue until the job is done
  e.g. *squeue -u myusername*
- Transfer the results back to your PC
  e.g. *WinSCP, Rsync, …*

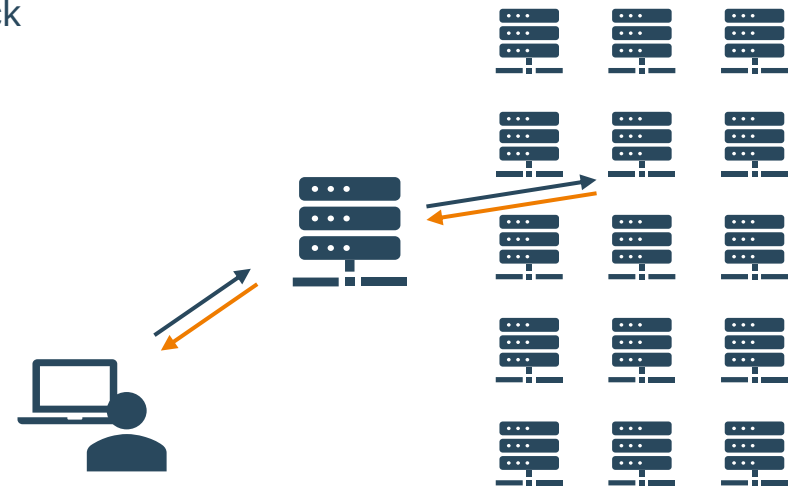

```
#!/bin/bash

#SBATCH --job-name=my-job
#SBATCH --no-requeue
#SBATCH --nodes=4
#SBATCH --ntasks-per-node=128
#SBATCH --time=02:00:00
#SBATCH --output=std.out
#SBATCH --error=std.err

QE_root="/u/sbeinlich/Software/QE/qe-7.3.1/intel_standard/"
inputfile="struc.pwi"
outputfile="struc.pwo"

#QE_flags=""
QE_flags="-nk 4"

srun_flags=""
#srun_flags="-n 512"


module purge

unset LD_LIBRARY_PATH

module load intel/2024.0
module load impi/2021.11
module load mkl/2024.0

source "$MKLROOT/env/vars.sh"

# This should ideally not be needed, but QE had trouble linking these libraries correct
ly...

export LIBS="-lmkl_intel_lp64 -lmkl_sequential -lmkl_core"

# magic environment variable settings similar to raven (might not be needed but unlikel
y to do harm)

srun $srun_flags $QE_root/bin/pw.x $QE_flags -in $inputfile > $outputfile
~
~
"job.sh" 38L, 828B written                                    38,1          All
```

QuantumESPRESSO job.sh

Documentation:

- PPBwiki (small): https://ppbwiki.fhi.mpg.de/Main/HPC
- GIT Wiki (large): https://gitlab.fhi.mpg.de/fhi-hpc/viper/-/wikis/home

VIPER SUPERCOMPUTER – MPCDF

Laptop / Desktop

Computing Infrastructure

1 node on Viper
~ 4 FHI Compute Servers
~ 32 Laptops

5 jobs x 10 nodes on Viper
~ 200 FHI Compute Servers
~ 1600 Laptops

Laptop

Run 24/7 for 1 month

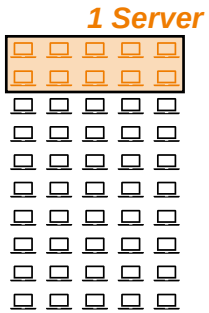

FHI COMPUTE
x 50

Run 3 days (1 server)

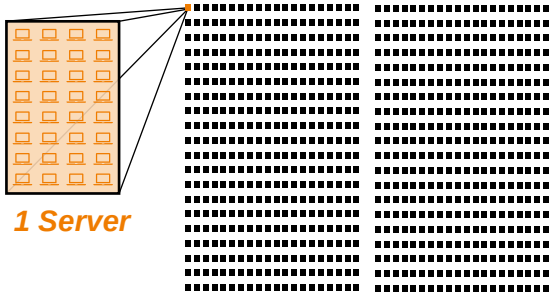

VIPER
x
25000
Run 22 hours @ 1 server
or 2 hours @ 10 servers

OTHER MPCDF SERVICES

OTHER MPCDF SERVICES

user-based, requires MPCDF account

- **Remote Visualization Services**
  (small, shared compute cloud)
  - Remote Desktop (VNC): 12c/256g, shared
  - Remote Visualization (VNC+GPU): 12c/256g/A30, shared
  - Jupyter Lab / Juyter Notebook: 12c/64g, shared
  - Jupyter Lab + GPU: 12c/64g/A30, shared
  - R Studio
- **DataShare: OwnCloud 200GB**
- **Nexus-S3: S3 Bucket 1TB**
- **Gitlab**
- **Archiving (Tape Backup)**
- **Globus Data Transfer**
- **…**

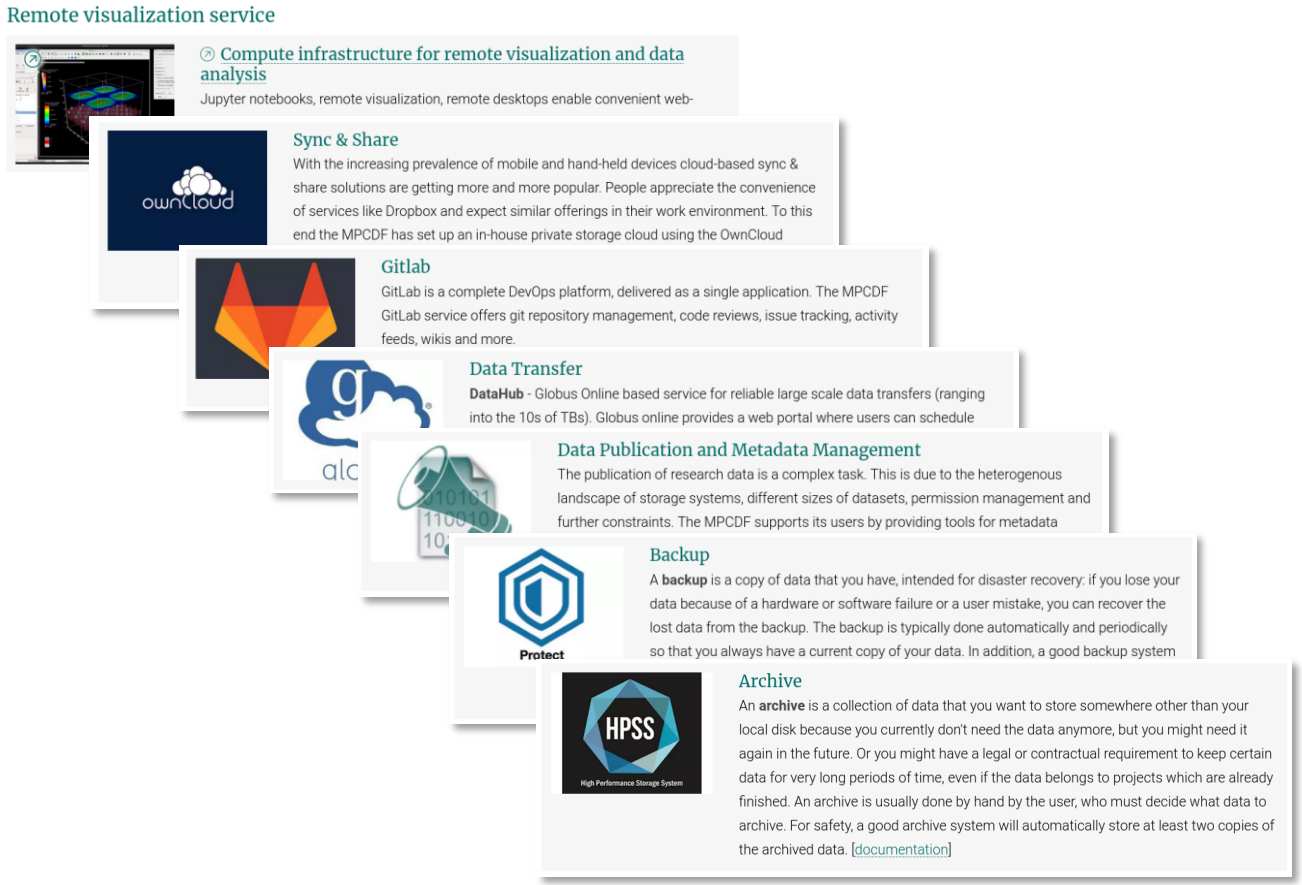

SUMMARY

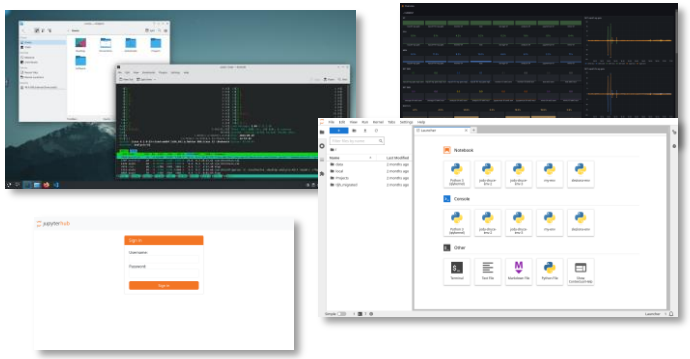

FHI (PPB)

FHI COMPUTE Servers

- Large Compute Servers (6x):
  Large CPU + Memory + GPU
- **Virtual Desktops** (Graphical Applications)
- **Jupyterhub** (Python & Matlab Server)
- Terminal Access
- **Direct Execution** (no Scheduling)
- User friendly

Anything that requires more than a desktop PC

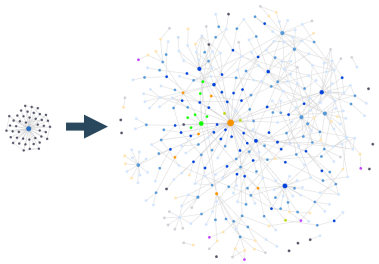

THANK YOU!

MPG (MPCDF)

VIPER Supercomputer

- Huge Compute Servers (768x):
  Huge CPU + Memory + GPU
- Terminal Access
- **Job Submission** (Scheduling)
- Advanced

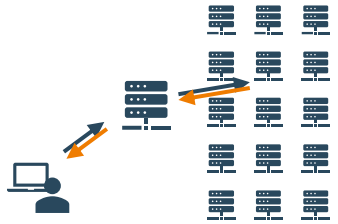

+ RAVEN, SuperMUC, JUWELS, …
(department-level / grant-based)

Large Scale Analyses / Processing & Ab-Initio Simulations