More Sustainable HTC Computing

How to better utilize compute resources.

NAF Admin Team
Christoph Beyer, Kruno Sever, Martin Flemming, Jan Hartmann, Thomas Hartmann, Yves Kemp
DESY IT

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

RF2.0 RESEARCH FACILITY

DESY.

Two sessions

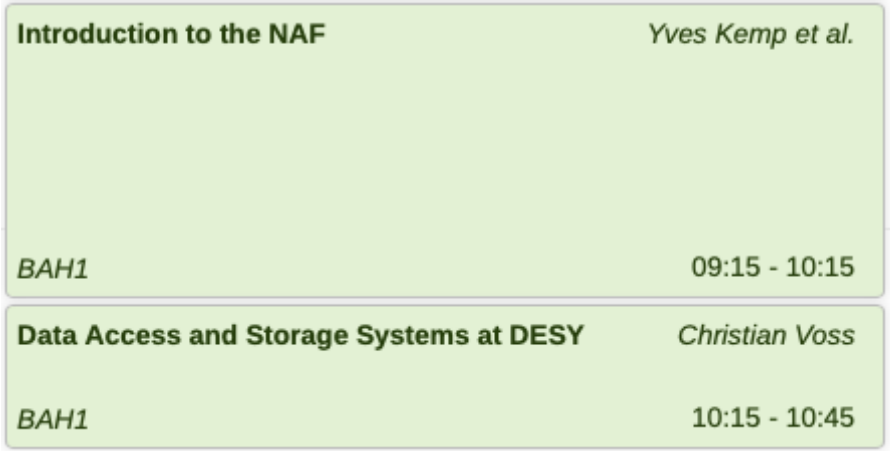

- General sustainability considerations
- What is the NAF? A (short) introduction
- The different building blocks
- Software on the NAF
- Storage on the NAF
- Remote storage
- Getting support

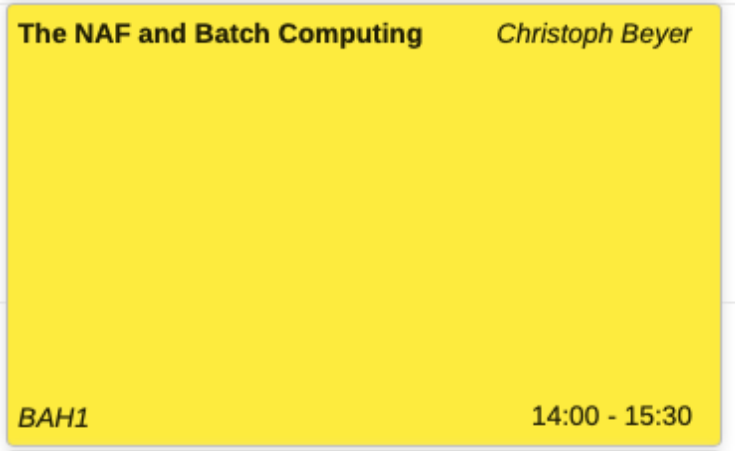

- Basic concepts of batch clusters
- basic usage: Hello-world jobs
- more advanced jobs, multiple jobs, DAG
- interactice jobs
- information on running/past jobs, debugging
- Sustainability in batch

Some more qualitative aspects

- Computers consume energy
  - We estimate that in average, computing one hour on one CPU core (with access to data store)
  - consumes 0.28 kWh of energy
  - produces 0.014 kg of CO2 (according to the German power mix)
  - which is equivalent to that produces on 0.095 km on a VW Golf

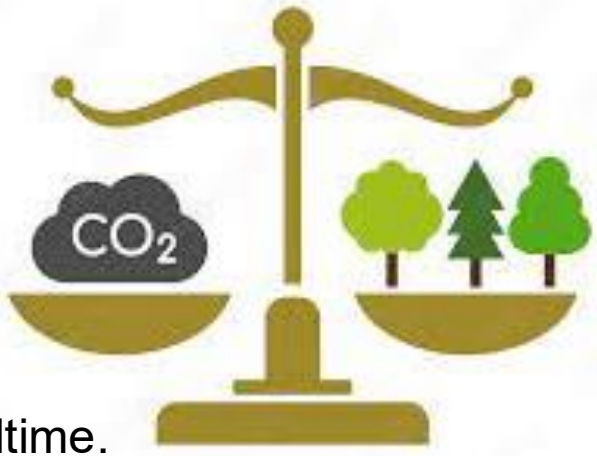

- … the NAF batch system has ~9000 CPU cores and operates 24x7x365
- e.g. jobs are getting killed by the batch system because they exceeded their walltime.
  - e.g. in the 7 days before 29.8.2023:
  - equivalent to ~20.000 core-h of walltime
  - 5600 kWh of energy … or 33 kW of average power
  - 280 kg of CO2 … would need 1200 large trees to capture that (large tree captures around 12 kg CO2/year)
  - 1900 km on a VW Golf

Sustainability research at IT

- Our batch farms are not always fully occupied
  - Shut down the machines when not in use
  - Or throttle the frequency
  - Or use them parasitically for something else
- In some future, energy prices will be flexible also for us. Currently, our model is still flat price … but we had a situation in July where power peaks over noon should be diminished
  - Grid, NAF, Maxwell participated – testing different approaches
  - Up to 110 kW less power in peaks
  - ~3.300 kWh energy saved over 12 days
- More efforts: Digital twin, frequency optimization, server lifecycle prolongation, intelligent scheduling, …

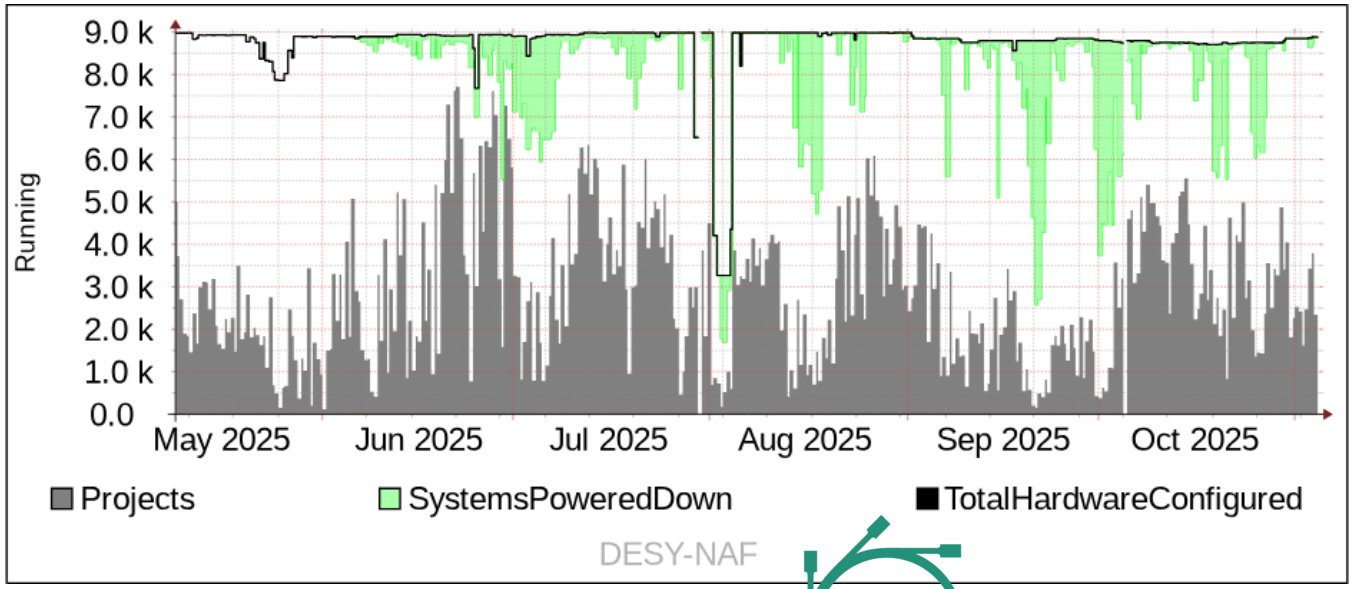

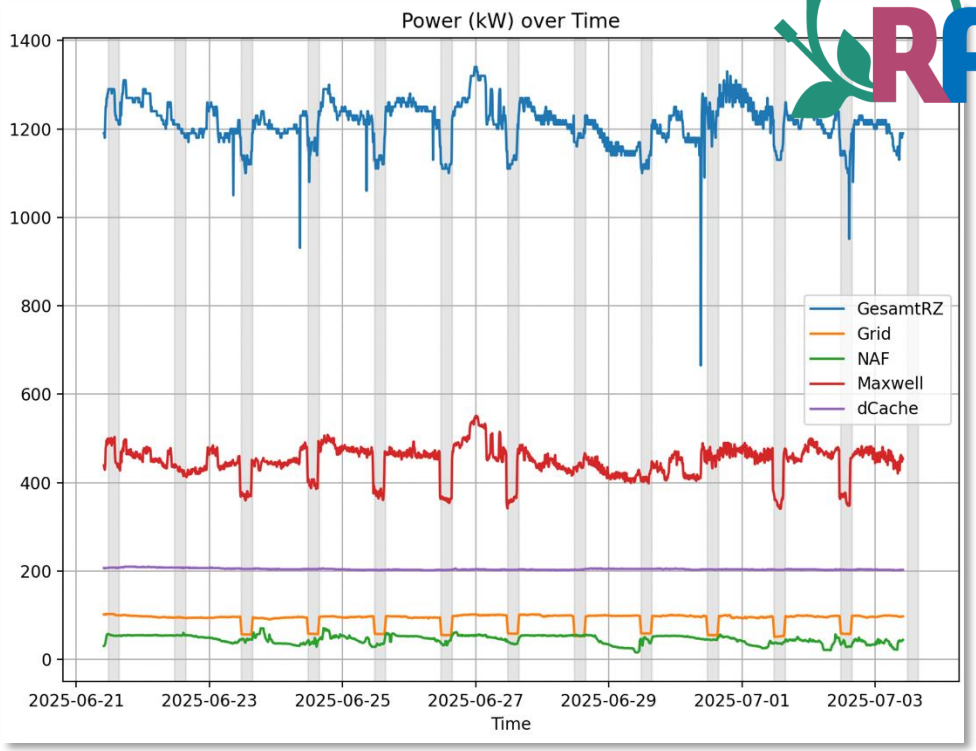

RF2.0 RESEARCH FACILITY

Who are you?

What do you expect from this workshop?

… but first some toughts about sustainability

- IT is a major energy consumer
- IT equipment has a footprint beyond energy

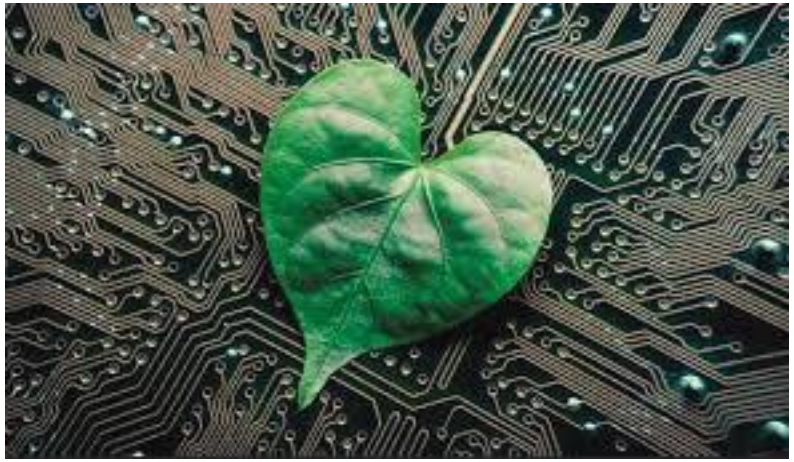

- Users:
  - Most efficient usage of resources
    - → Optimize workflows, optimize algorithms, optimize code
- Resource providers:
  - Chose efficient hardware, optimize setup
  - Well think about lifecycle of systems
  - Reduce non-IT energy consumption, reuse waste heat
- Research infrastructures, energy providers, society
  - Provide more renewable energy, use more renewable energy, use renewable energy when it is available

First part:
Introduction to the NAF

What is the NAF?

- NAF stands for „National Analysis Facility“
  - *National* means: For people working in institutes in the Terascale Alliance
  - *Analysis* means: Analysing data taken in the Terascale Alliance
  - *Facility* means: Something where you can do real work
- Basically: The NAF is a facility where **YOU** can do your analysis (and stuff around your analysis)

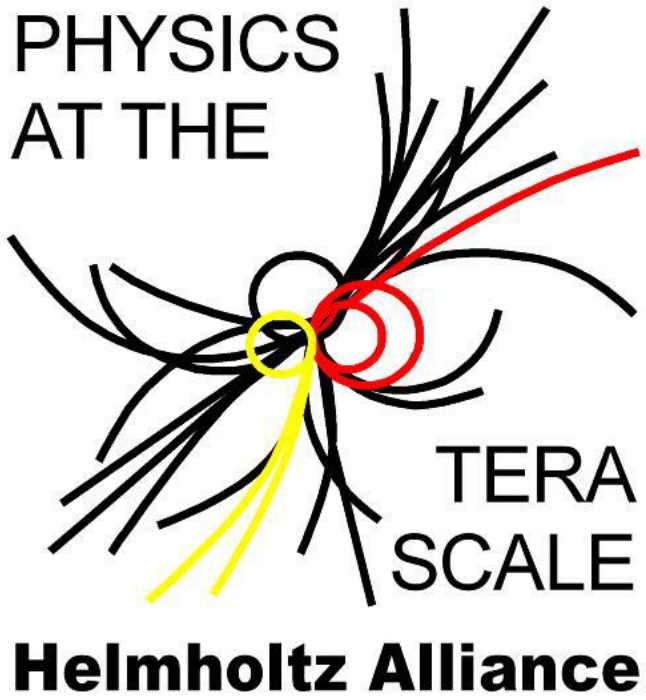

Who can use the NAF?

- The political answer:
  - You are from a Terascale Alliance institute
  - You are in LHC, doing analysis (ATLAS, CMS, LHCb)
  - You are in Belle, ILC, legacy HERA
- The technical answer:
  - You need a DESY account from
- Each user group has (a set of) login machines, e.g.: naf-atlas.desy.de
  - Access to this login machines (→ your group admin)
- You need the “batch” resource to be able to submit jobs (→ your group admin / UCO)

You – and the LHC compute & storage at DESY

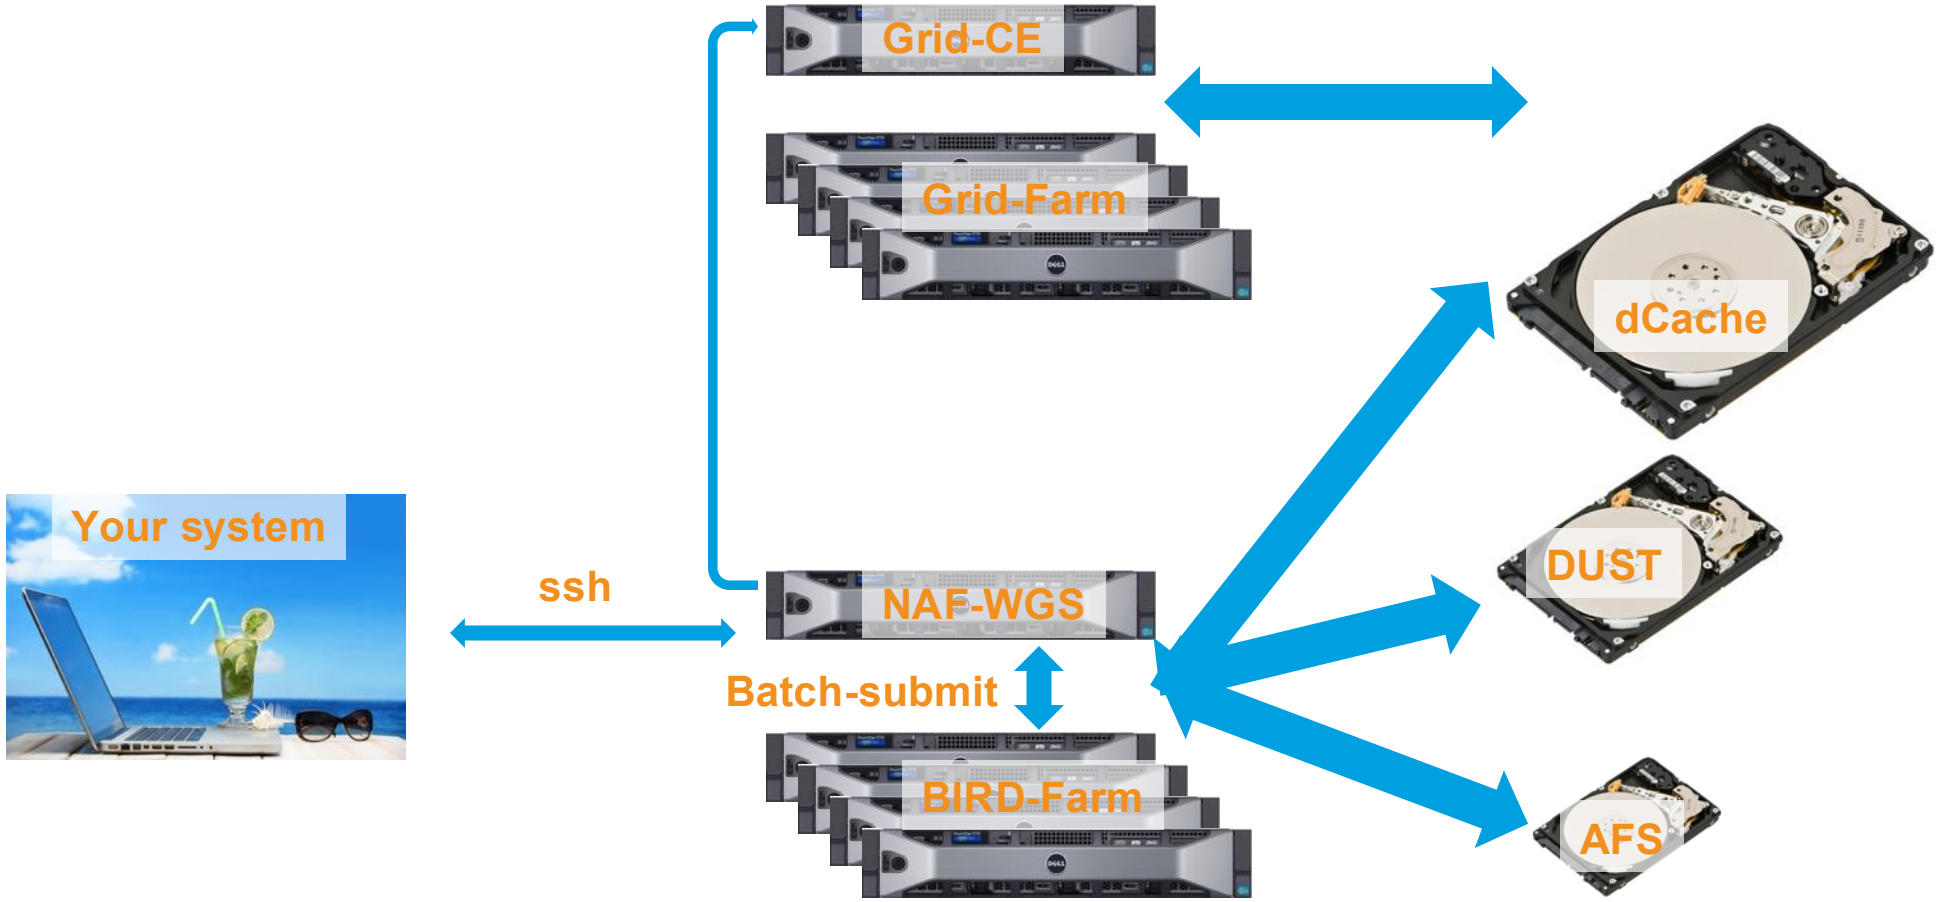

Workgroupserver to login to

- try logging into a WGS:
- open a terminal on your laptop
- ssh USERNAME@naf-NNN.desy.de

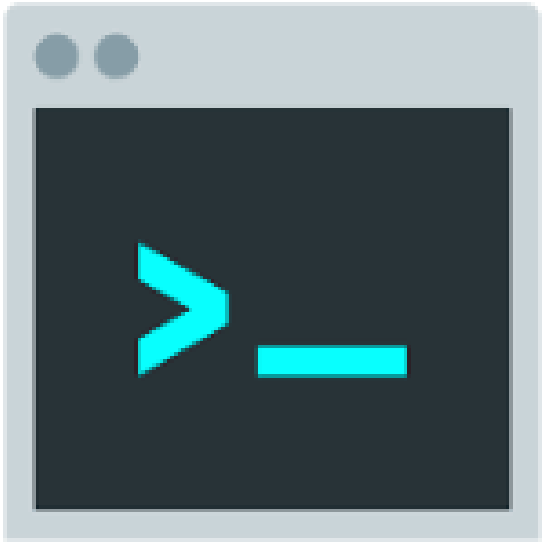

The WGS: Login Node, WorkGroupServer, Interactive Node, …

- Each group has a set of own WGS
- Why own?
- How many?
- How many resources?

- **Usage of WGS (from the NAF docu)**
- **Be reasonable** when you use the WGS as you are **sharing the resources** with your colleagues and any (mis)use will affect them. So, please be careful and nice to your colleagues. As a rule of thumb, use the WGS for development**, short (<10min) test jobs, graphical things**, … Do not use the WGS for **longer lasting processes**, or processes that do not require interactive return. Such processes are best placed onto the batch system - which has much more resources and resource control mechanisms.
- We check WGS on a daily basis, and if we find long running jobs or jobs consuming huge amounts of memory, we write an automated email.
- We have to reboot systems to activate new kernels. Using load balancing and alias mechanisms, we try hard to keep logins alive for at least one week. Sometimes we must act faster, though.

Other access means

- Graphical access: FastX to the NAF remote desktop

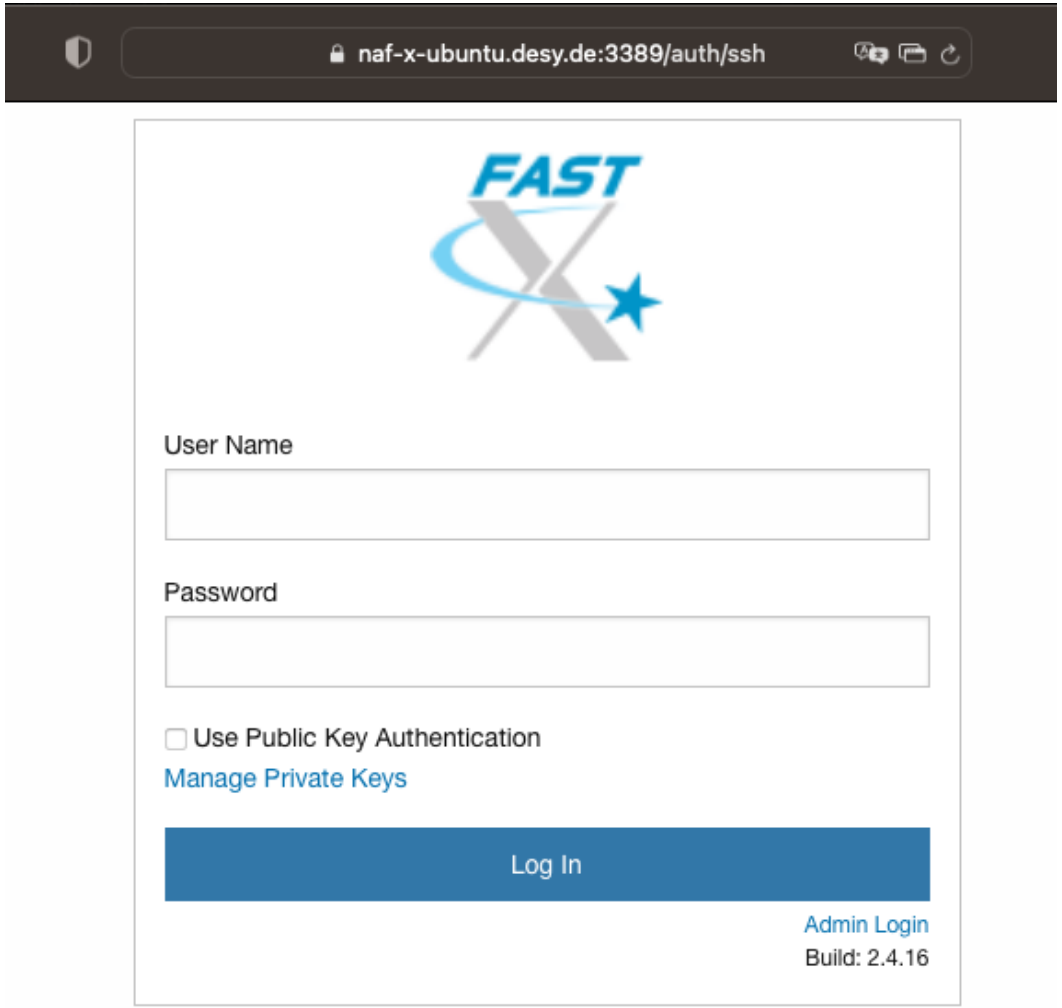

- JupyterHub

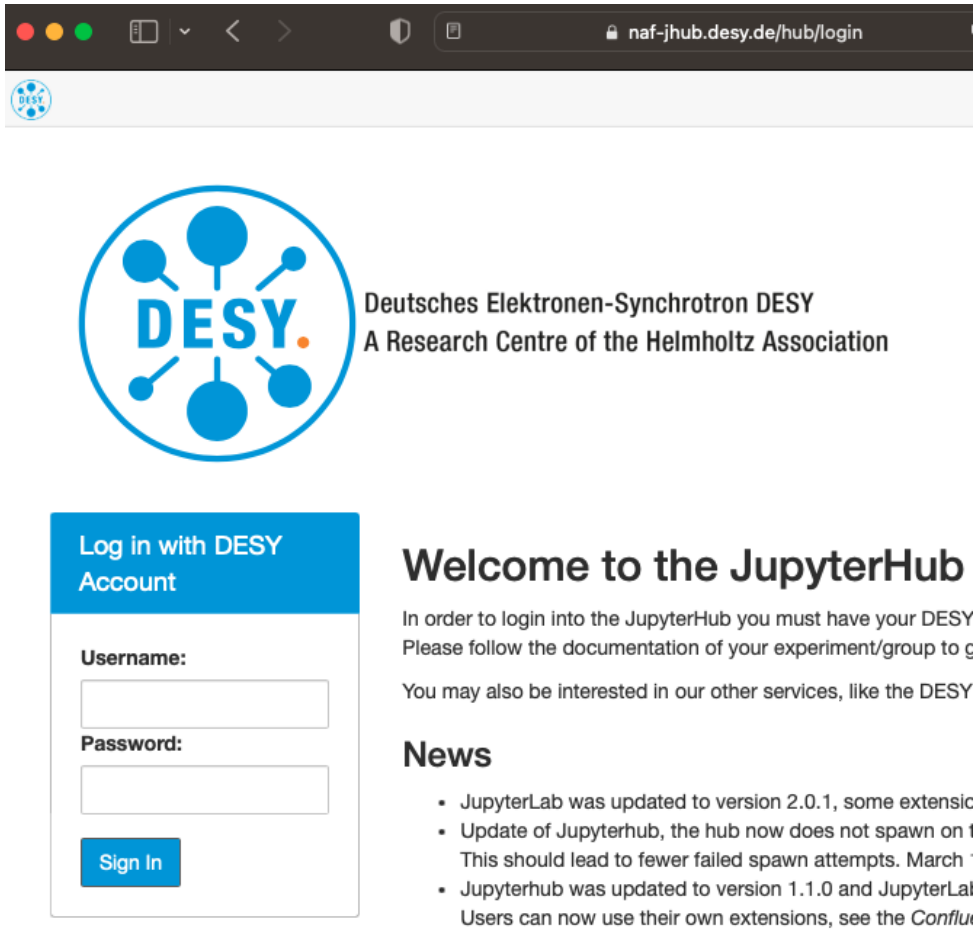

What is this?

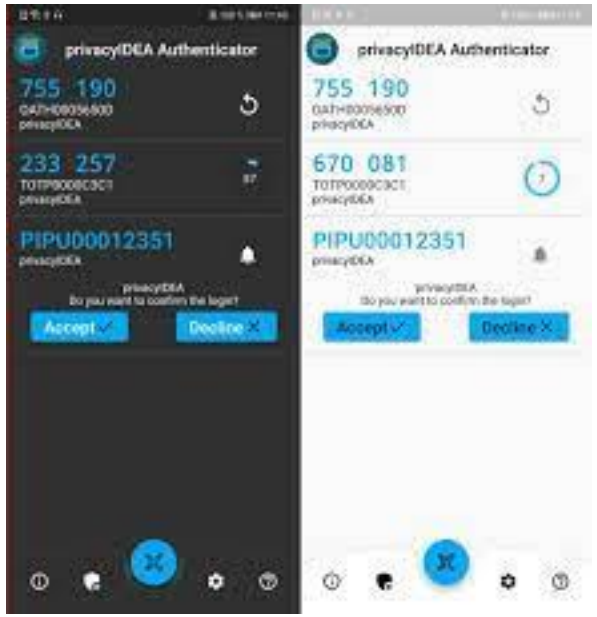
privacyIDEA Authenticator
755 190
233 257
PIPU00012351
privacyIDEA
Do you want to confirm the login?
Accept
Decline
privacyIDEA Authenticator
755 190
670 081
PIPU00012351
privacyIDEA
Do you want to confirm the login?
Accept
Decline

Why is the WGS slow? What is the footprint of my job? …

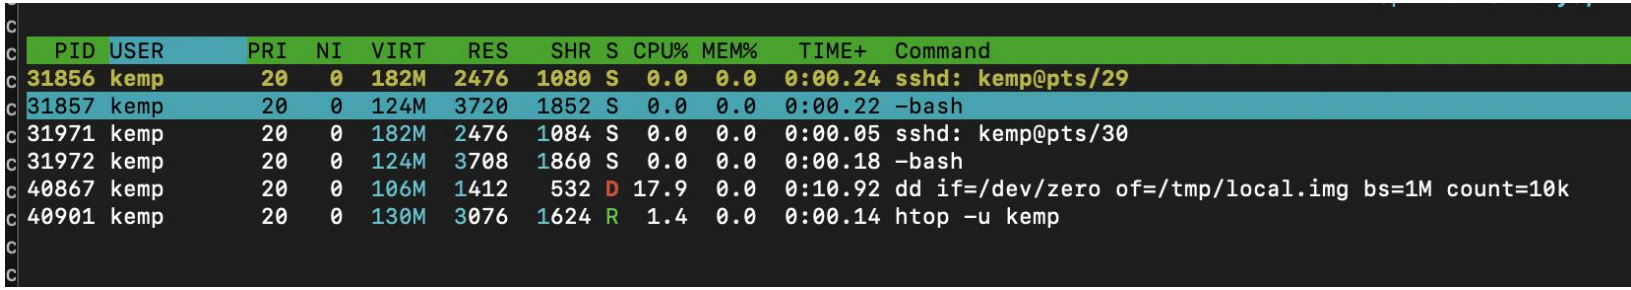

```
  PID USER      PRI  NI  VIRT   RES   SHR S CPU% MEM%   TIME+  Command
31856 kemp       20   0  182M  2476  1080 S  0.0  0.0  0:00.24 sshd: kemp@pts/29
31857 kemp       20   0  124M  3720  1852 S  0.0  0.0  0:00.22 -bash
31971 kemp       20   0  182M  2476  1084 S  0.0  0.0  0:00.05 sshd: kemp@pts/30
31972 kemp       20   0  124M  3708  1860 S  0.0  0.0  0:00.18 -bash
40867 kemp       20   0  106M  1412   532 D 17.9  0.0  0:10.92 dd if=/dev/zero of=/tmp/local.img bs=1M count=10k
40901 kemp       20   0  130M  3076  1624 R  1.4  0.0  0:00.14 htop -u kemp
```

- What goes on on a server?
  - top or (nicer) htop (htop –u kemp)
- NFS traffic:
  - nfsiostat 1
  - nfsiostat 1 /data/dust/group/atlas
- general, including some IO:
  - sar 1
  - vmstat 1
- Discussion 1: Interpretation:
  - Is 10% IO wait a lot? acceptable? negligible?
- Discussion 2:
  - What is *my* job doing? Am I responsible for a slow system?

- State “D” : Waiting for Disk (or network disk)
  - Might indicate a problem if observerd for a longer time
  - e.g. NFS server problem
  - e.g. close to 100% quota
- State “Z” : Zombie
  - Is definitely a problem … if consuming lots of CPU or RAM → write to helpdesk, only reboot can fix
- IT informs per email if cpu> or RAM > …
  - sometimes a simple “kill $PID” does not help
    - difference SIGTERM and SIGKILL

Can’t login to the WGS? … and some other WGS problems

- ssh naf-atlas13.desy.de
  - Connection refused … or the like
  - Have you tried naf-atlas.desy.de (without the 13 ?)
  - → We have many WGS, and might be doing maintenance on this special one. Preferably use the generic one.

```
Could not chdir to home directory /afs/desy.de/user/x/xxx: Connection timed out
/usr/bin/xauth:  error in locking authority file /afs/desy.de/user/x/xxx/.Xauthority
-bash: /afs/desy.de/user/x/xxx/.bash_profile: Connection timed out
[xxx@naf-cms14 /]$ pwd
/
```

- Problem with your AFS – home
  - Drop naf-helpdesk a message, with your account, and the WGS in question (naf-cms14 in this case)
- WGS freezing randomly:
  - Sometimes, AFS/NFS is slow to react when many files are in a directory
  - If no clue: inform naf-heldesk with [user,time,WGS,path] … we might or might not find a reaon, but good for us to know anyhow

RAM story

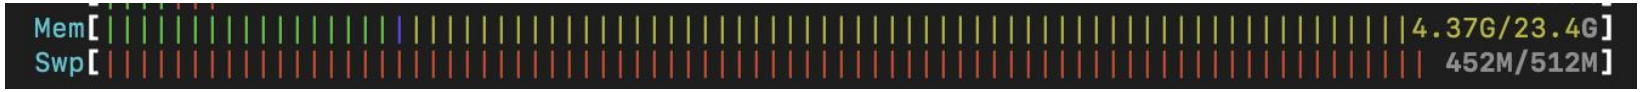

- RAM is used as file system cache
  - high RAM occupancy is not necessarily a problem
- High Swap is OK: unused MEM pages are stored on disk
  - … unless the swapd has lots of CPU time (or vmstat swap-in/swap-out is increased)

- The Kernel kills processes when RAM is getting too tight
  - over the years, it has gotten better at spotting the right culprit

Software on the NAF

- Local system installed
- cvmfs
- module (and AFS)
- group software (if any)

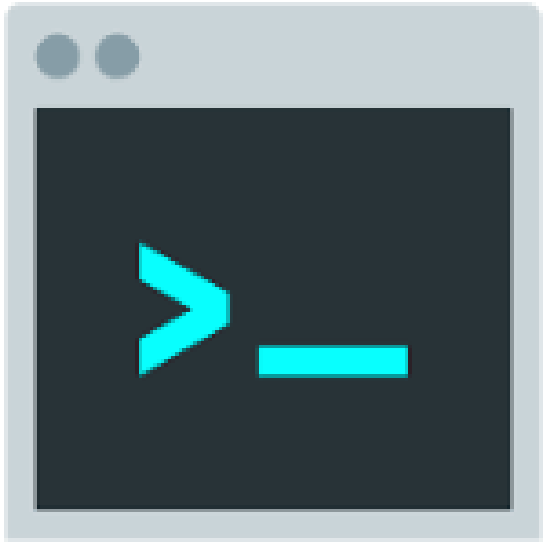

Storage in the NAF

Local and Shared Mount Namespaces

More details: Next presentation

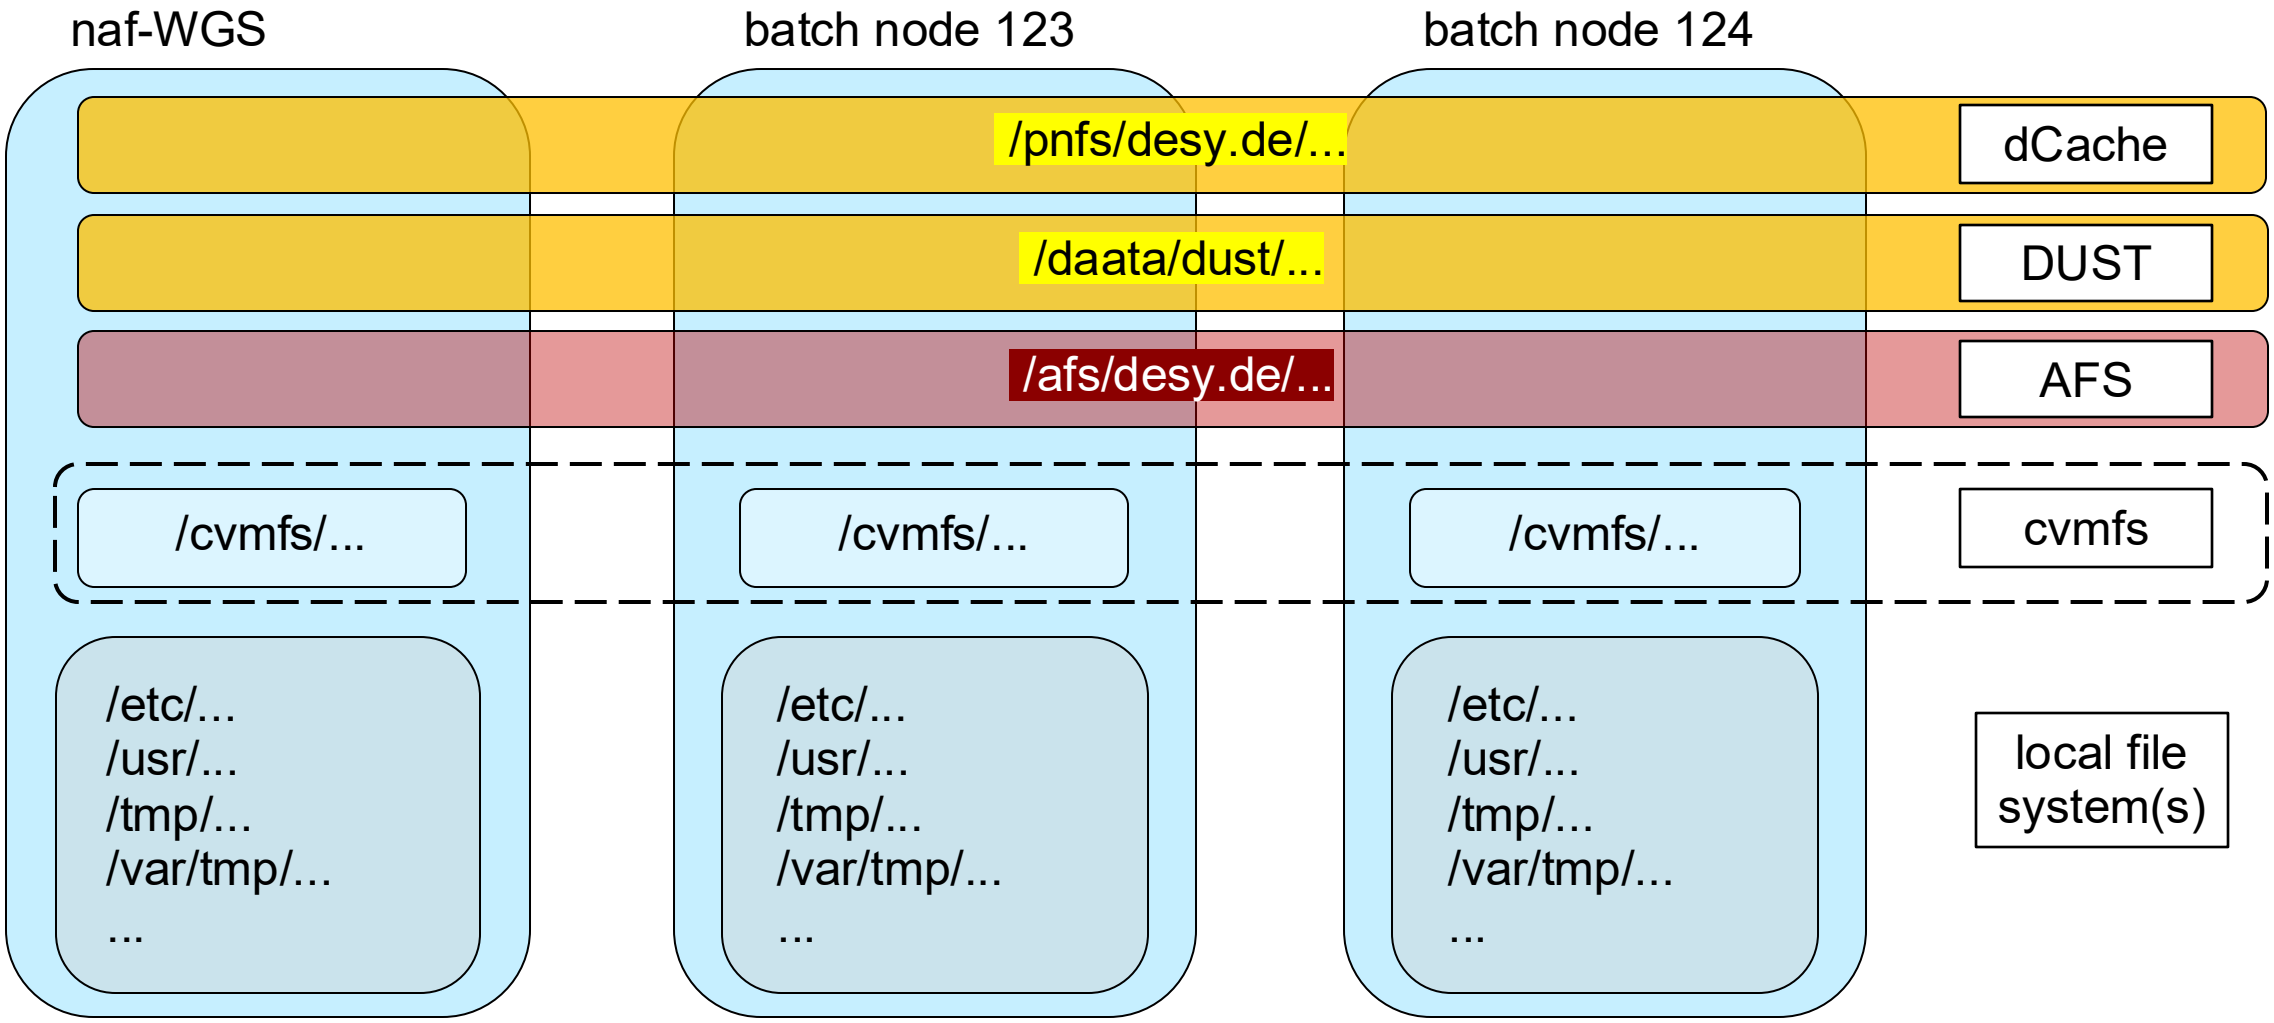

Why so many of them?

More details: Next presentation

- Different in scope and availability
- Different in available size of installation
- Different in available user directory size
- Different in connection to backup
- Different in performance
- Different in protocols and access methods

Storage dos and don’ts

More details: Next presentation

… or: https://docs.desy.de/naf/documentation/filesystems/

- AFS
  - Backup … IO limited, size limited, some restrictions w.r.t number of files and file names
- DUST:
  - No backup
  - Few, large files are nicer than many, small files: Files <128k contribute more to quota
  - Recommendation: Limit yourself to base ASCII (and best no space) … no Emojis, UTF-8
- dCache (see next slide)
- CVMFS
  - … very large files or container images might cause troubles

https://docs.desy.de/naf/documentation/filesystems/

Network filesystems are mounted EVERYWHERE on NAF: No need to scp/rsync/… between components!

dCache network storage

More details: Next presentation

Do's und Don'ts:

- Keep the numbers of jobs accessing the same files to a minimum. Don't have 1k jobs read the same file over and over again
- Storage does not scale linearly into infinity; 2 Jobs reading the same file might half the time to wait for results, 1k jobs do not reduce the runtime by a factor of 1k
- Be aware, that the storage is also a shared resource; excessively reading the same file again and again will have a negative impact on your colleagues reading data from the same storage node or blocking the worker nodes for your colleagues and other NAF communities
- Keep file sizes reasonable, merge output into few larger files; many small files can over the long run cost performance in file listings
- **Tape:** Some NAF groups (ILC, Belle and CMS/ATLAS in the future) have data only stored on tape. These file need to be staged in advance. If you get a permission denied error when reading files; check the locality of the file. In order to prevent issues on the NAF, reading a file that is only on tape is forbidden

Storage tips

More details: Next presentation

- How much space do I have?
- DUST:
  - statistics collected once per hour
- AFS:

```
naf-atlas13::kemp>my-dust-quota
Fileset Name         Usage (TB)  Limit (TB)  Use (%)  File Usage
user.af-atlas.kemp   0.534       1.0         53.43    124899
naf-atlas13::kemp>my-dust-quota -g
Fileset Name             Usage (TB)  Limit (TB)  Use (%)  File Usage
user.af-atlas.kemp       0.534       1.0         53.43    124899
group.af-atlas.atlas-d   0.369       1.0         36.95    218866
```

```
naf-atlas13::~>fs lq -hu
Volume Name                    Quota       Used %Used   Partition
user.kemp                      16.0G      10.0G   63%         47%
naf-atlas13::~>fs lq -hu xxl/
Volume Name                    Quota       Used %Used   Partition
xxl.kemp                        7.6G       7.3G   96%<<       76%    <<WARNING
```

- No such tool for dCache
  - Experiments might control usage

```
naf-atlas13::~>df -h /tmp/
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda6       59G  3.8G   52G   7% /tmp
```

- Local file systems:
  - (Network filesystem gives unmeaningful information)
- Information on one directory (I have read permissions on)

```
naf-atlas13::kemp>du -sch gcc-11.4.0*
365M    gcc-11.4.0
937M    gcc-11.4.0-non
133M    gcc-11.4.0.tar.gz
1.5G    total
naf-atlas13::kemp>du -sch --inodes gcc-11.4.0*
11K     gcc-11.4.0
28K     gcc-11.4.0-non
1       gcc-11.4.0.tar.gz
39K     total
```

Remote storage (outside of DESY)

- Through XrootD, Grid-FTP or WebDAV data can be read from remote sites, e.g. CERN, FNAL or BNL
- Keep in mind due to limited bandwidth accessing/transferring the data can waste CPU time
- Ideally, try to transfer the data first, then submit jobs; or keep in mind the lower performance when requesting job slots on NAF

support & documentation

- **Infrastructure support**
  - Multi-level support within IT
  - Dedicated Queues in our Ticket system
  - naf-helpdesk@desy.de
- **Experiment support**
  - Organized by experiments
  - Mailing lists
- Good communication between experts from IT and experiments needed
  - e.g. via the NUC (users committee)

- **Documentation:**
- Lots of stuff under:
  - https://naf.desy.de/
  - and probably some experiment-own documentation