NHR FAU

FAU Friedrich-Alexander-Universität Erlangen-Nürnberg

HPC systems at NHR@FAU

https://hpc.fau.de/systems-services/documentation-instructions/

Which cluster should I use?

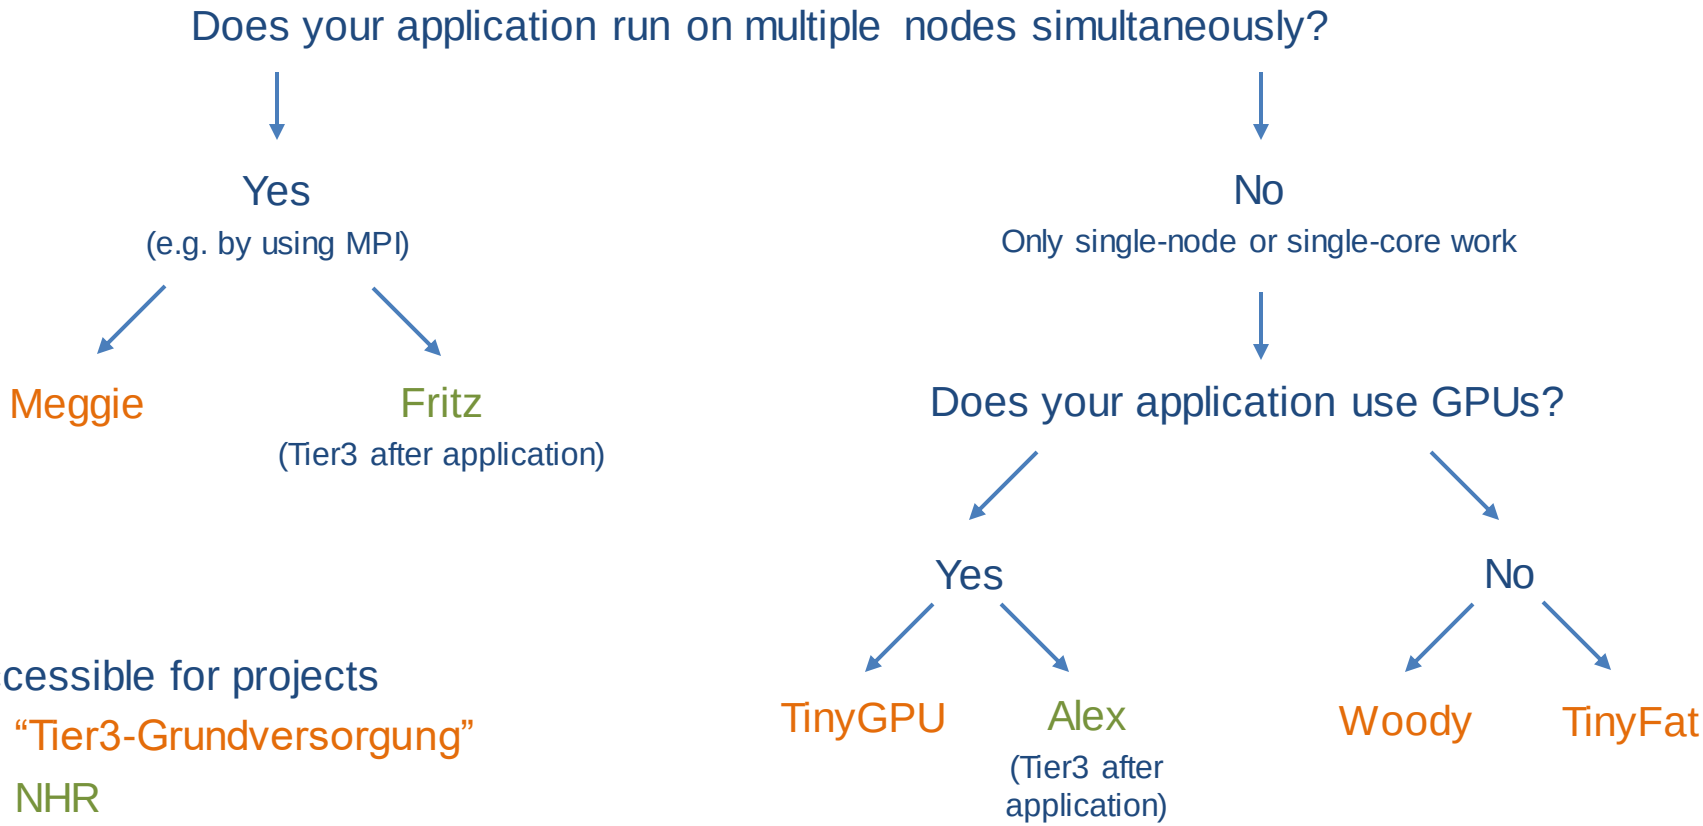

Accessible for projects

- "Tier3-Grundversorgung"
- NHR

“Fritz” cluster

NHR parallel cluster, open for Tier3 users after application

- 992 compute nodes (71.424 cores)
  - 2 Intel Xeon Platinum 8360Y “Ice Lake” 2.4 GHz (36 cores)
  - 256 GB main memory per node
- 64 huge-memory nodes
  - 2 Intel Xeon Platinum 8470 “Sapphire Rapids” 2.0 GHz (52 cores)
  - 1 or 2 TB of main memory per node
- Blocking (1:4) HDR100 Infiniband network, up to 100 GBit/s
- Parallel file system with 3,5 PB capacity

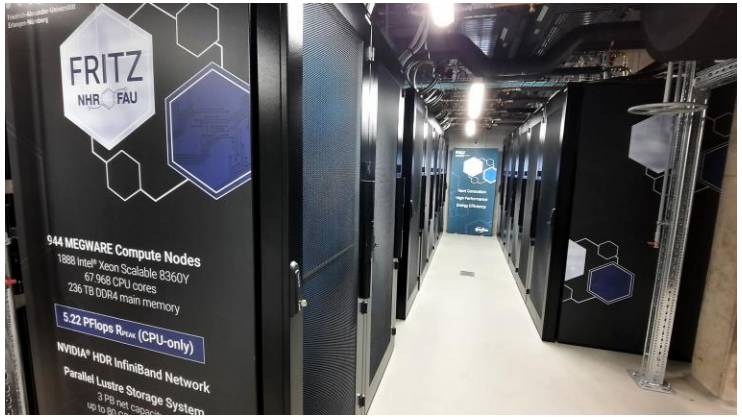

“Meggie” cluster

current main cluster for parallel jobs, intended for highly parallel jobs (Tier3)

- 728 Compute nodes (14.560 cores)
  - 2 Intel Xeon E5-2630 v4 (Broadwell) 2.2 GHz (10 cores)
  - 20 cores/node
  - 64 GB main memory
- No local disks
- Intel OmniPath network: Up to 100 Gbit/s

“Alex” cluster

NHR GPGPU cluster, open for Tier3 users after application

- 44 nodes with
  - 8x NVIDIA A100 (each 40 GB / 80GB HBM2)
  - 2x AMD EPYC 7713 “Milan” 2.0 GHz, 1024 GB / 2048 GB of main memory
  - 14TB local NVMe SSD
  - HDR200 Infiniband network
- 38 nodes with
  - 8x NVIDIA A40 (each with 48 GB DDR6)
  - 2x AMD EPYC 7713 “Milan” 2.0 GHz, 512 GB of main memory
  - 7 TB local NVMe SSD

“TinyGPU” cluster

for GPU workloads – not all nodes always generally available (Tier3)

- 12 nodes with 2x “Skylake” @ 3.2 GHz, 96 GB RAM, 1.8 TB SSD, 4x RTX 2080Ti
- 4 nodes with 2x “Skylake” @3.2 GHz, 96 GB RAM, 2.9 TB SSD, 4x Tesla V100
- 7 nodes with 2x “Cascade Lake” @2.9 GHz, 384 GB RAM, 3.8 TB SSD, 8x RTX3080
- 8 nodes with 2x AMD Rome 7662 @2.0 GHz, 512 GB RAM, 5.8 TB SSD, 4x Volta A100

“Woody” cluster

Main workhorse for throughput and single-node jobs (Tier3)

- 176 nodes with 4 cores and high clock frequency (3.5/3.7 GHz) Intel Xeon E3-1240 v? processors
  - 64x Intel Skylake, 32 GB RAM
  - 112x Intel Kaby Lake, 32 GB RAM
- 70 nodes with 2x Intel Xeon Gold 6326 (32 cores total @2.9 GHz, 256 GB RAM)
- at least 960 GB local HDD/SSD
- Gbit network only

NHR FAU

FAU Friedrich-Alexander-Universität Erlangen-Nürnberg

File Systems

File systems overview

Available file systems differ in size, redundancy and how they should be used

Mount point	Access	Purpose	Technology	Backup	Snap-shots	Data lifetime	Quota
/home/hpc	$HOME	Source, input, important results	NFS	YES	YES	Account lifetime	50 GB
/home/vault	$HPCVAULT	Mid-/long-term storage	NFS	YES	YES	Account lifetime	500 GB
/home/{woody, saturn, titan, janus, atuin}	$WORK	General-purpose, log files	NFS	NO	NO	Account lifetime	500 GB (Tier3) NHR project quota
/lxfs /lustre	$FASTTMP (fritz+alex)	High performance parallel I/O	Lustre via InfiniBand	NO	NO	High watermark	Only inodes
/???	$TMPDIR	Node-local job-specific dir	SSD/ ramdisk	NO	NO	Job runtime	NO

Working with large datasets containing small files

Best case: use a container file format (HDF5, Parquet, …)

- if necessary, copy to node-local SSD

Alternative:

- pack small files into archive
- Do not unpack archive to **`$HOME/$HPCVAULT/$WORK`**
- Unpack files to node-local SSDs only and use them from there

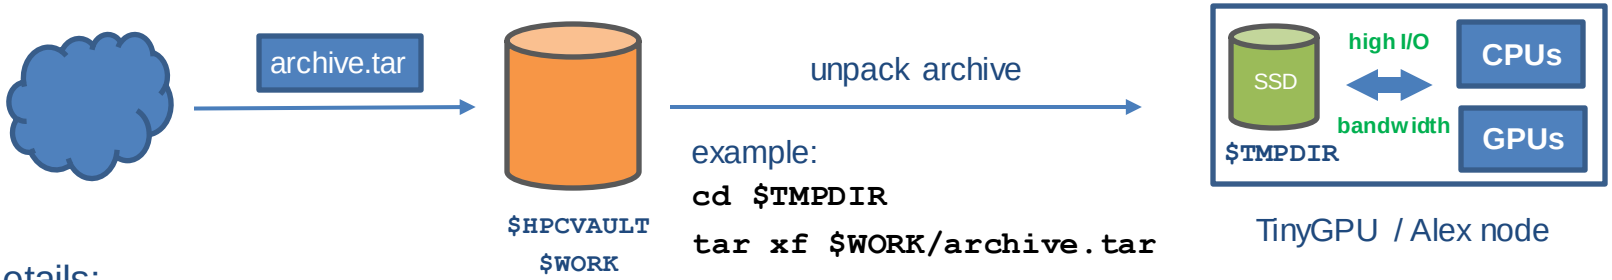

More details:

https://hpc.fau.de/files/2022/01/2022-01-11-hpc-cafe-file-systems.pdf

https://www.fau.tv/clip/id/40199

File system quotas

- File system may impose quotas on data volume and/or number of files
- Quotas may be set per user or per group (or both)
- Soft quota
  - Can be exceeded temporarily (7 days)
  - Turns into hard quota at end of grace period
- Hard quota:
  - absolute upper limit, cannot be exceeded

```
$ quota -s            # generic command

$ shownicerquota.pl   # only on RRZE systems
   Path               Used     SoftQ    HardQ    Gracetime  Filec    FileQ    FiHaQ    FileGrace

   /home/hpc             5.7G    52.5G    104.9G       N/A      72K     500K   1,000K        N/A
   /home/woody           112G   333.0G    499.5G       N/A     188K                          N/A
```

Data transfer

- **$HOME**, **$HPCVAULT**, **$WORK** are mounted on all HPC systems

- **scp** / **rsync** is used to transfer files from and to the outside

```
scp [options] source destination
# source/destination:  local-path | username@host:remote-path
```

```
# local file to remote
# -r  recurse into directories
scp -r code user@cshpc.rrze.fau.de:/home/woody/group/user
# remote file to local
scp user@cshpc.rrze.fau.de:results/output.dat .
```

- Additionally on Windows: WinSCP, MobaXTerm