UNIVERSITÄT PADERBORN

Short Introduction to Otus

Paderborn University, Germany
Paderborn Center for Parallel Computing

NATIONALES HOCHLEISTUNGS RECHNEN

Paderborn Center for Parallel Computing

Topics for Today

1. What is Otus
2. Timeline
3. How to Access
4. Software
5. Details on Nodes

What is Otus?

Otus: The new HPC System at PC2 (Otus-NHR)

- 743 compute nodes, 108 NVIDIA H100-94GB GPUs
- 6 PB GPFS parallel file system
- Infiniband 200/400 Gbit/s interconnect
- Funded by the NHR Alliance (nhr-verein.de)

© Paderborn University / Stefan Rohde

Nodes	Partition	Count	
Compute nodes	normal	636	2xAMD 9655, 192 cores, 768 GiB RAM
Large-memory nodes	largemem	48	2xAMD 9655, 192 cores, **1536 GiB RAM**
GPU-Nodes	gpu	27	2xAMD 9655, 192 cores, 768 GiB RAM, **4xNVIDIA H100-94GB NVLINK**
FPGA-Nodes	fpga	32	2xAMD 9655, 192 cores, 768 GiB RAM, Planned: **FPGAs and NVIDIA A40**
(Huge-memory nodes in planning)			(Planned >=3 TB RAM)

NATIONALES HOCHLEISTUNGS RECHNEN

More details http://pc2.uni-paderborn.de/go/otus

Timeline

Otus-NHR

- 2024: procurement process
- Spring 2025: delivery and setup of system
- August 2025: acceptance of the system and test operation
- September 2025: general user availability
- From October 2025: production operation
- (End of 2025: Retirement of Noctua 1)

NATIONALES HOCHLEISTUNGS RECHNEN

One HPC system!

Otus-Tier3 (HPCFachCluster.nrw)

More details https://hpc.dh.nrw/de/hpcfachcluster

- Mid 2025: DFG Proposal submitted
- End of 2025: procurement process
- Mid of 2026: production operation

Access

Otus-NHR

- September 2025: general user availability
  - Access: any compute time project
  - Compute contingent:
    - small getting-started contingent
    - NHR projects can request to move quota from Noctua 1/2 to Otus

- From October 2025: production operation
  - Access: only NHR compute projects (NHR-starter, NHR-normal, NHR-large) due to funding (simply apply for an NHR project! Details later today)
  - Compute quota:
    - NHR projects can request to move quota from Noctua 1/2 to Otus

- After Extension (Otus-Tier3) 2026:
  - Access: NHR (starter, normal, large) and Tier3 (small)

Otus-Tier3

NATIONALES HOCHLEISTUNGS RECHNEN

Access

Summary of useable systems

	Tier-3 projects: (small)	NHR projects (starter, normal, large):
Noctua 1	Till end of 2025	Till end of 2025
Noctua 2	As usual	As usual
Otus	In Sept. 2025	From Sept. 2025
	From mid of 2026 (extension Otus-Tier3, HPCFachCluster.nrw)	

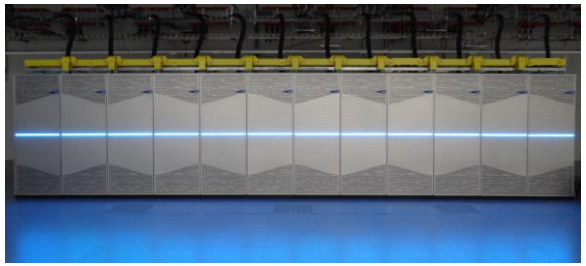

Why different Project Categories?

Different needs for HPC are served/finded by different tiers:

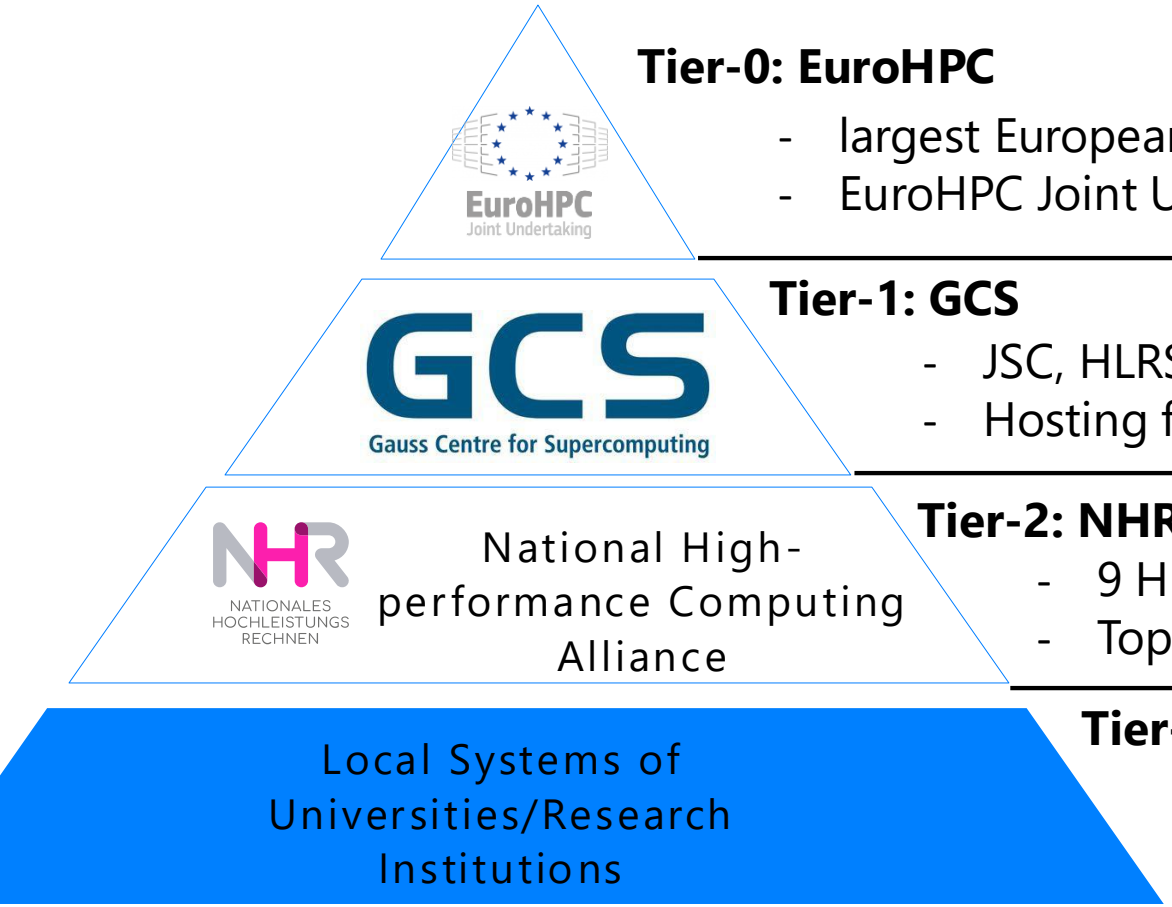

Tier-0: EuroHPC

- largest European systems
- EuroHPC Joint Undertaking (JU)

Tier-1: GCS

- JSC, HLRS, LRZ
- Hosting for EuroHPC-Systems

Tier-2: NHR

- 9 HPC centers in Germany (including PC2)
- Topically focused

Tier-3: local

- Single institute or local university
- DFG Forschungsgroßgeräte

Paderborn Center for Parallel Computing

Overview: Possible Project Categories

	Tier-3 (small)	NHR-Starter	NHR-Normal	NHR-Large
Resources per year	up to **4M** CPU-core-h, up to **10k** GPU-h or FPGA-h	up to **8M** CPU-core-h, up to **10k** GPU-h or FPGA-h	up to **30M** CPU-core-h, up to **100k** GPU-h or FPGA-h	**> 30M** CPU-core-h or **> 100k** GPU-h or FPGA-h
Location of PI	Paderborn or NRW (Tier-3 supplementary service)	Germany		
Project Duration	One year	One year	One year or multiyear	One year or multiyear
Proposal submission	**Any time**	**Any time**	**Any time**	**4 fixed dates** (first of Jan., Apr., Jul., Oct.)
Project start	a few days	a few days	preliminary start about one week full project starts 4-8 weeks	preliminary start about one week full project starts 3 months
Simplifications	Only short online form!	Only short online form!	1) Multiyear project 2) Scientifically reviewed project 3) well-known programs	1) Multiyear project 2) well-known programs

See also https://pc2.uni-paderborn.de/system-access
If unsure, please contact us.

Possible Simplifications for NHR-Projects

1. Using well-known programs
   - List at https://pc2.uni-paderborn.de/go/well-known-programs
   - **-> No benchmark or scalability required**

2. Already scientifically reviewed project (NHR-Normal)
   - by German federal ministry (BMBF, BMWK, BMUV, BMEL,...), DFG, GCS, European Union (EU projects), Volkswagen Stiftung, another NHR center
   - e.g. TRR, CRC, ExIn, GRK, Emmy Noether,....
   - Needed: Information about granted project (project id, GEPRIS-link,...) but NOT the review
   - **-> much shorter detailed description, shorter processing time**

3. Multi-year proposals
   - **-> one proposal for multiple years of project runtime**

4. Extensions of NHR-Normal/Large projects
   - **-> reduced detailed description**

See also https://pc2.uni-paderborn.de/go/simplifications

Access: Login and Partitions

In practice just like Noctua 1 and Noctua 2:

- SSH: `fe.otus.pc2.uni-paderborn.de` then `ssh otus`
- File systems:
  - **HOME**: same on all systems
  - **PC2DATA** (permanent, /pc2/groups/hpc-prf-.../): same on all systems
  - **PC2PFS** (temporary, /scratch): parallel file system of Otus
- Job submission: Same SLURM configuration as in N1/N2

Nodes	Partition	Count		Note
Compute nodes	normal	636	2xAMD 9655, 192 cores, 768 GiB RAM	
Large-memory nodes	largemem	48	2xAMD 9655, 192 cores, **1536 GiB RAM**	
GPU-Nodes	gpu	27	2xAMD 9655, 192 cores, 768 GiB RAM, **4xNVIDIA H100-94GB NVLINK**	`--gres=gpu:h100:[count]`
FPGA-Nodes	fpga	32	2xAMD 9655, 192 cores, 768 GiB RAM, Planned: **FPGAs and NVIDIA A40**	

Software

In practice just like Noctua 1 and Noctua 2:

- Full list available at https://pc2.uni-paderborn.de/go/software
- "`find_module`" to search
- In job script

```
module reset
module load ...
```

- Please let us know if you are missing
  - Software/versions/libraries
  - Usage examples, e.g. example job scripts

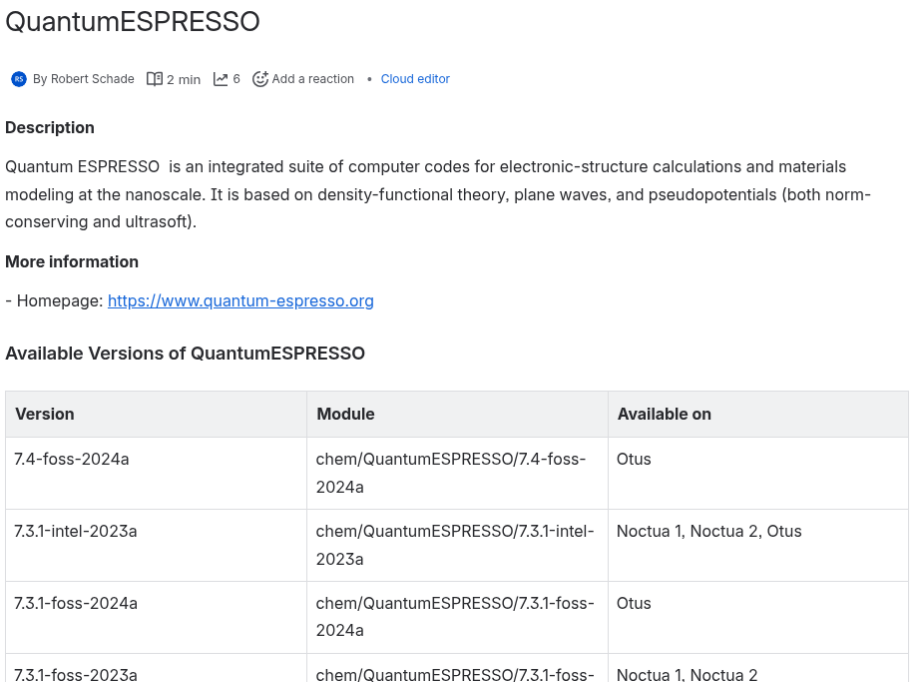

QuantumESPRESSO

By Robert Schade 2 min 6 Add a reaction • Cloud editor

Description

Quantum ESPRESSO is an integrated suite of computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory, plane waves, and pseudopotentials (both norm-conserving and ultrasoft).

More information

- Homepage: https://www.quantum-espresso.org

Available Versions of QuantumESPRESSO

Version	Module	Available on
7.4-foss-2024a	chem/QuantumESPRESSO/7.4-foss-2024a	Otus
7.3.1-intel-2023a	chem/QuantumESPRESSO/7.3.1-intel-2023a	Noctua 1, Noctua 2, Otus
7.3.1-foss-2024a	chem/QuantumESPRESSO/7.3.1-foss-2024a	Otus
7.3.1-foss-2023a	chem/QuantumESPRESSO/7.3.1-foss-	Noctua 1, Noctua 2

In case you need/want to compile yourself:

Please use

- For GCC/Gfortran/AOCC:
  - `-march=znver5` (or at least `-march=znver4`), -march=native if compiling on Otus
- For Intel compiler:
  - `-march=core-avx512` (but NOT `-xHost`)
- For nvcc/GPUs:
  - Compute capability 9.0: `-arch=compute_90 -code=sm_90`

BLAS/Lapack:

- Best performance: AOCL-BLAS
- In practise: use Flexiblas module
  (wrapper around different BLAS/Lapack implementations)

Support

Support for any questions/problems, e.g.:

- Accessing the HPC systems at PC2
- Submitting compute jobs to the workload manger
- Available software packages and development tools
- System architecture and specifications
- Data management and additional services
- ….

Scientific Support:

- Compute time proposals
- Porting and Optimizing your compute jobs
- Organization of development workflows for your scientific code
- Code porting, debugging, optimization,…
- Algorithm design, implementation, and benchmarking
- Performance modelling
- Workflow organization and research data management
- Organization of scientific events like conferences and workshops
- ….

Main contact:
pc2-support@uni-paderborn.de or via https://portal.pc2.uni-paderborn.de/support

Weekly online support/consultation hour, Thursdays 13:00-14:00,
https://pc2.uni-paderborn.de/go/consultationhour

HPC courses/events/schools:
https://pc2.uni-paderborn.de/go/training

Details on Nodes: CPU Nodes

Improvements to Noctua 2:

- More powerful CPU cores: Zen5 (AVX512) vs. Zen3 (AVX2)
  - --> recompile your software!
- More cores per Node: 192 vs. 128
  - Adapt your job scripts!
- More memory available:

	Otus	Noctua 2
Normal	768 GiB	256 GiB
Largemem	1536 GiB 3 TB local NVME SSDs	1024 GiB
Hugemem	In planning	2048 TiB

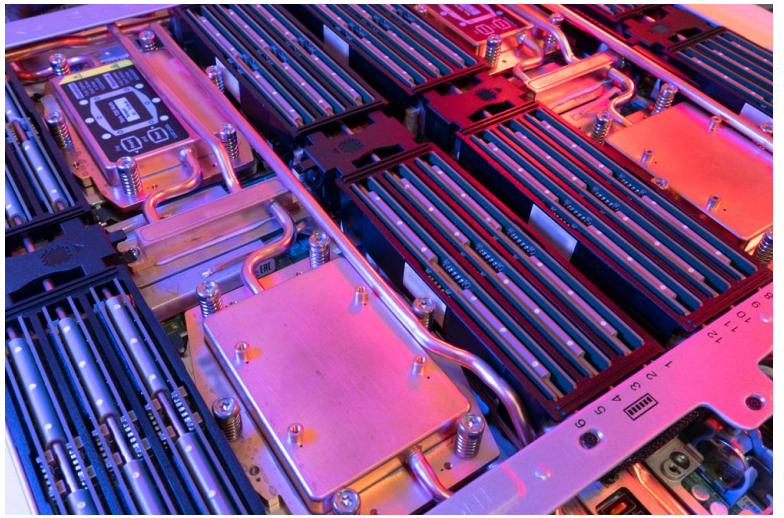

© Paderborn University / Stefan Rohde

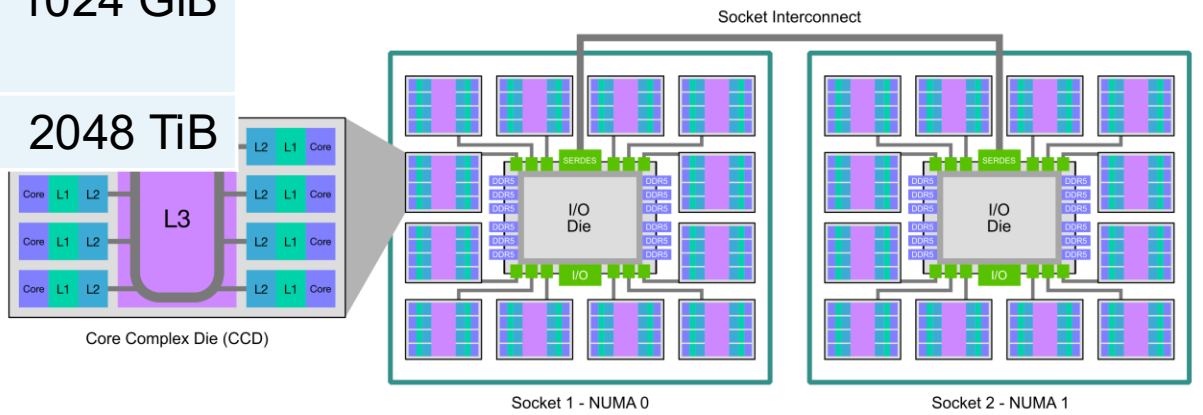

Otus node architecture

Details on Nodes: CPU Nodes

Improvements to Noctua 2:

	Otus	Noctua 2	Increase
Memory Bandwidth	970 GB/s	360 GB/s	2.7x
Floating-point Performance per Node (FP64)	13.2 TFlop/s	3.95 TFlop/s	3.3x
Interconnect	200 Gbit/s	100 Gbit/s	2x
VASP CuC_VdW	95 s	237 s	2.5x
CP2K H2O_512	224 s	674 s	3.0x
Quantum Espresso GRIR443	410 s	1139 s	2.8x

Increase by more than **2x in all per-node metrics and benchmarks**

-> your calculations should run **2-3 times as fast** with the same number of compute nodes
(if not, there is likely an issue with compiler settings/libraries, please contact us)

Details on Nodes: GPU Nodes

NVIDIA GPU	H100 (Otus)	A100 (Noctua 2)
GPUs per Node	4	4
GPU Arch	Hopper	Ampere
NVLINK	900 GB/s	600 GB/s
PCIe	Gen5 128 GB/s	Gen4 64 GB/s
Memory	94 GiB HBM2e	40 GiB HBM2
Bandwidth	2.40 TB/s	1.56 TB/s
Flops FP64 (TC)	67 TFlop/s	19.5 TFlop/s
Power	Up to 700 W	Up to 400 W

```
#SBATCH --gres=gpu:h100:[count]
```

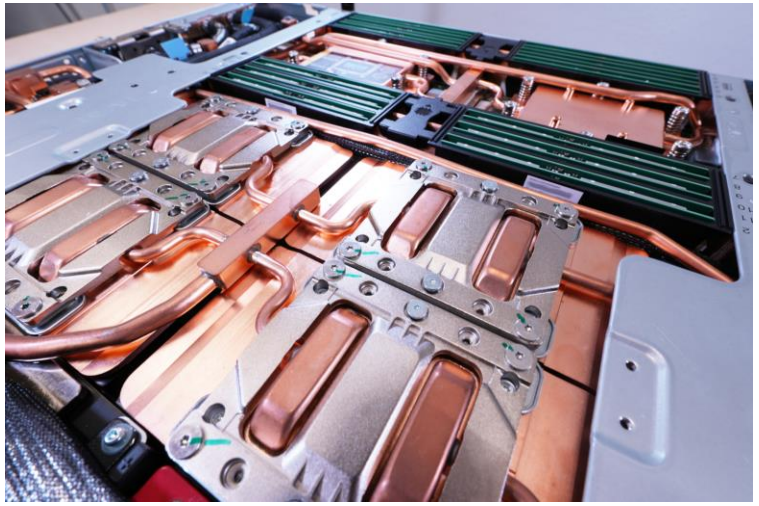

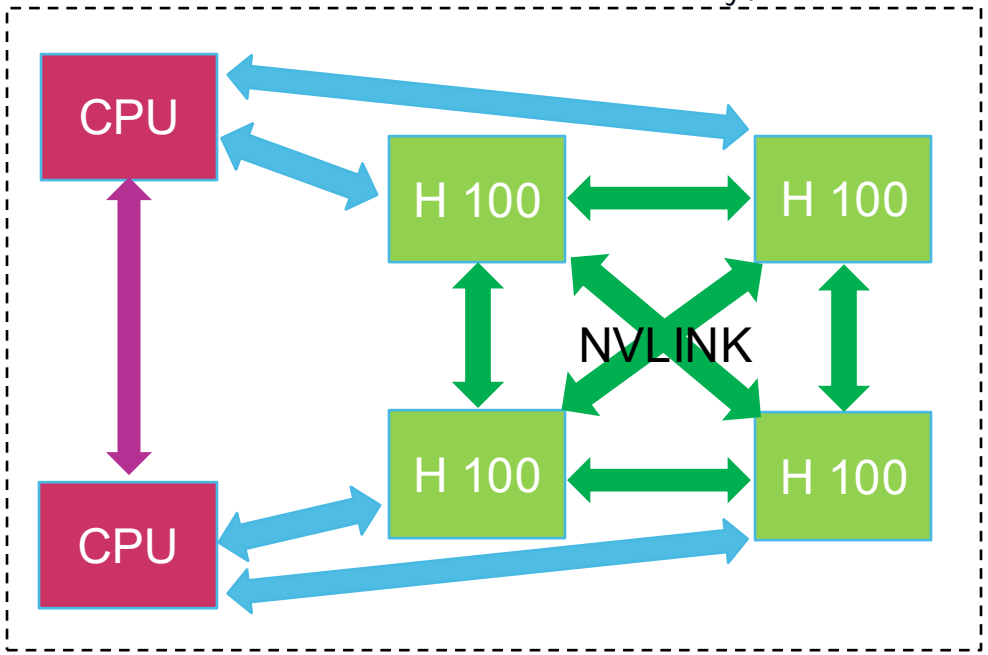

GPU Node

Details on Nodes: FPGA Nodes

	FPGA Partition	
Number of Nodes	32	
Accelerator Cards	AMD V80	IBEX IPAC-1000
	PCIe Gen5 x8	PCIe Gen5 x16
	Up to 32	Up to 32

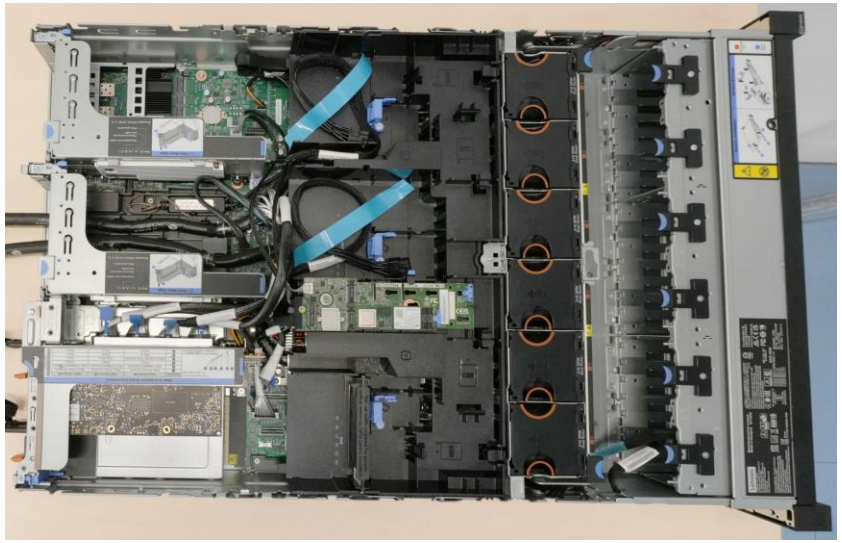

Currently in Pilot phase
See also https://doku.pc2.uni-paderborn.de/pages/607125508/Otus+FPGA+Pilot+Phase

Parallel File System

Noctua 1: ClusterStor with Lustre

- All-in-all smooth operation

Noctua 2: DDN Exascaler with Lustre

- Often problems with slow metadata operations (open/close/stat) due to design of Exascaler
- SSD and HDD storage tier management was challenging

Otus: IBM Spectrum Scale with GPFS

- More advanced storage tier management
- Parallel file system of Noctua 2 accessible on login nodes via `/scratch-n2`

Noteworthy Differences to Noctua 2

- **192 cores** per node instead of 128 cores
- **AVX512** instead of AVX2
- configured with **one NUMA-Domain per CPU** (instead of 4 for Noctua 2)
- More **memory** and **memory bandwidth**
- More powerful **GPUs**
- (hopefully) **more stable parallel file system**

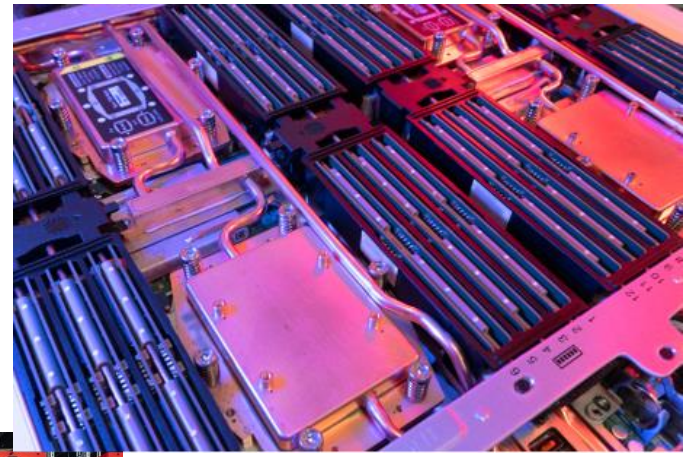

Job Monitoring

Performance metrics of your compute jobs:

- Online: https://jobmon.pc2.uni-paderborn.de/
- Check for example:
  - Usage of CPU cores in job
  - Memory usage and memory bandwidth
  - File IO
  - GPU usage

We are happy to optimize your compute jobs or support you!

Job Monitoring

Example: LAMMPS Job on GPUs

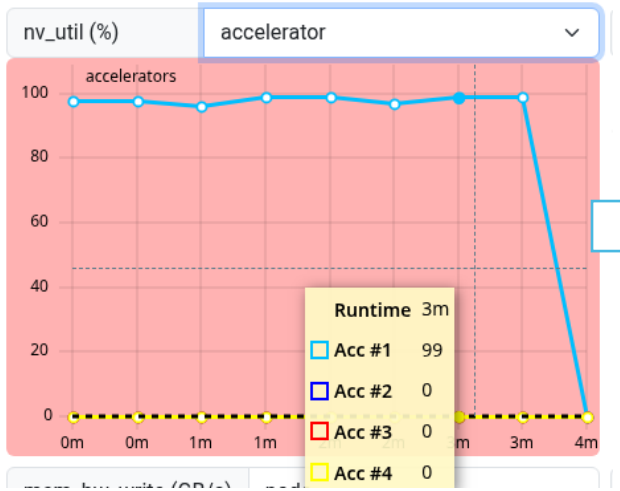

Only **one** of the four allocated GPUs is used!

```
#!/bin/bash
#SBATCH -N 1
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=48
#SBATCH -t 1:00:00
#SBATCH -p gpu
#SBATCH --gres=gpu:h100:4

export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK
export OMP_PLACES=cores
export OMP_PROC_BIND=True

module reset
module load chem/LAMMPS/27Jun2024-foss-2023b-kokkos-CUDA-12.5.0
srun lmp -k on g 1 -sf kk -pk kokkos neigh half comm device
   neigh/qeq full newton on -in in.reaxc.hns -nocite -log out.log
```

Example: Quantum Espresso Job on CPUs

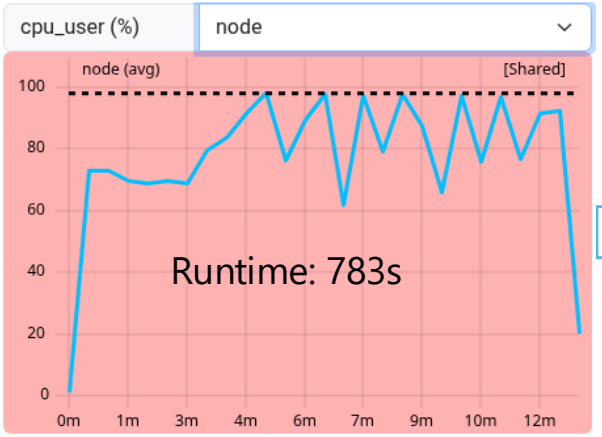

No full utilization of allocated cpu cores (here due to too many threads: 48 threads per MPI-process)

With 2 threads per MPI-process

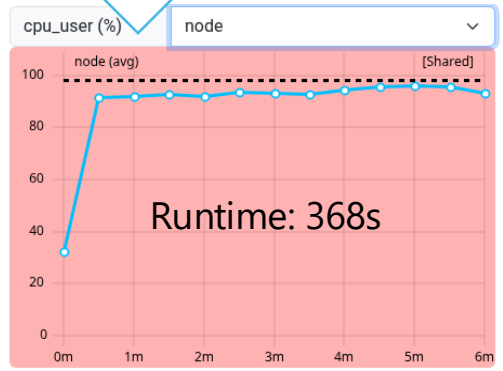


```
#!/bin/bash
#SBATCH -N 1
#SBATCH --ntasks-per-node=4 #MPI-processes per node
#SBATCH --cpus-per-task=48 #OpenMP-threads per MPI-process
#SBATCH -t 1:00:00

export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK
export OMP_PLACES=cores
export OMP_PROC_BIND=True

module reset
module load chem/QuantumESPRESSO/7.4-foss-2024a
srun pw.x -input grir443.in 2> err > out
```

What is Otus?

Otus: The new HPC System at PC2 (Otus-NHR)

- 743 compute nodes, 108 NVIDIA H100-94GB GPUs
- 6 PB GPFS parallel file system
- Infiniband 200/400 Gbit/s interconnect
- Funded by the NHR Alliance (nhr-verein.de)

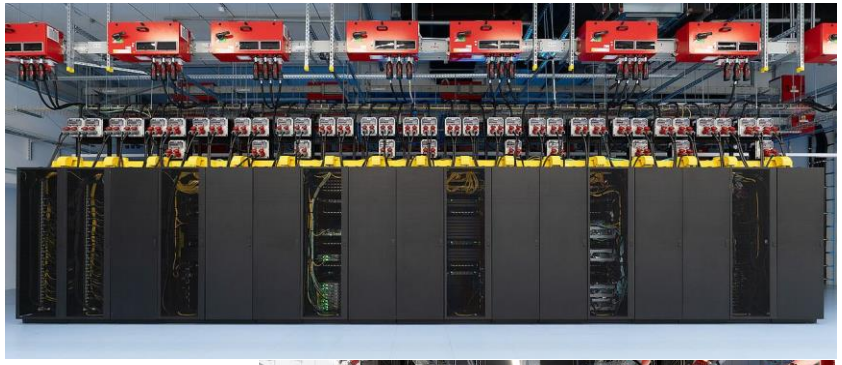

© Paderborn University / Stefan Rohde

Nodes	Partition	Count	
Compute nodes	normal	636	2xAMD 9655, 192 cores, 768 GiB RAM
Large-memory nodes	largemem	48	2xAMD 9655, 192 cores, **1536 GiB RAM**
GPU-Nodes	gpu	27	2xAMD 9655, 192 cores, 768 GiB RAM, **4xNVIDIA H100-94GB NVLINK**
FPGA-Nodes	fpga	32	2xAMD 9655, 192 cores, 768 GiB RAM, Planned: **FPGAs and NVIDIA A40**
(Huge-memory nodes in planning)			(Planned >=3 TB RAM)

NATIONALES HOCHLEISTUNGS RECHNEN

More details http://pc2.uni-paderborn.de/go/otus

Additional Material: Compute Time Proposals

Why Three different Project Categories?

Different needs for HPC are served/finded by different tiers:

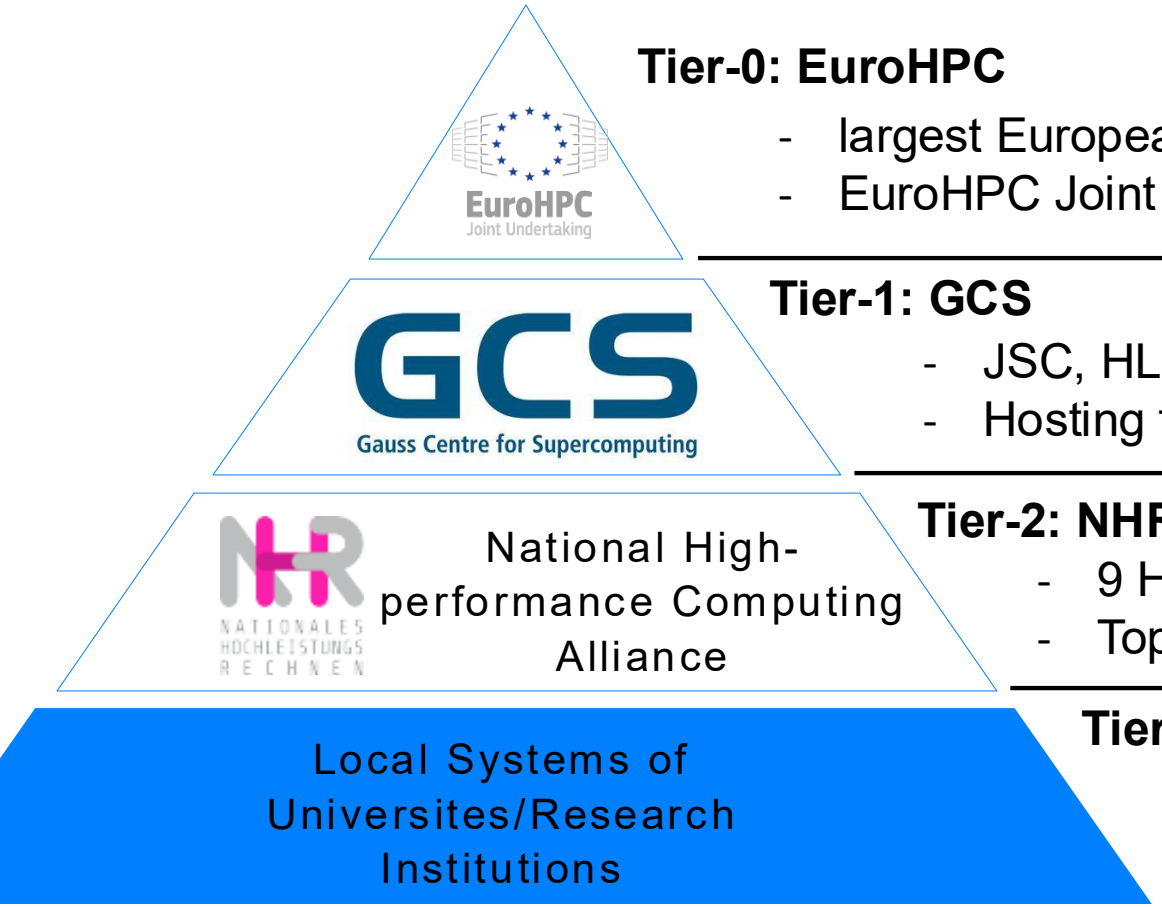

Tier-0: EuroHPC

- largest European systems
- EuroHPC Joint Undertaking (JU)

Tier-1: GCS

- JSC, HLRS, LRZ
- Hosting for EuroHPC-Systems

Tier-2: NHR

- 9 HPC centers in Germany (including PC2)
- Topically focused

Tier-3: local

- Single institute or local university
- DFG Forschungsgroßgeräte

Paderborn Center for Parallel Computing

What is the NHR Alliance?

National High-Performance Computing Alliance

- NHR, **N**ationales **H**ochleistungs**R**echnen
- funded by GWK and federal states, first funding phase 2021-2030
- NHR compute time regulations: https://www.nhr-verein.de/en/computing-time

Mission: strong research infrastructure for all researchers of universities im Germany with need for computing resources

- operating suitable and efficient HPC systems
- support and training for using of HPC resources
- scientific support and consulting in focus areas
- special support for young scientists
- advancing HPC and its use in research with collaborative research with users

- Focus areas of PC2:
  - Condensed matter physics, esp. solid-state physics and computational chemistry
  - Efficient HPC with accelerators like GPUs and FPGAs

NHR
NATIONALES HOCHLEISTUNGS RECHNEN

How to Choose a Suitable Follow-up Project?

Important aspects to consider:

- Resource Needs
- Formal Eligibility
- Effort for Proposal

How to Choose a Suitable Follow-up Project?

Resource Needs

- **CPU-Core-hours:**
  - 1 CPU-core-hour = allocating one CPU-core for one hour
  - Usually in millions of CPU-core-hours
  - CPU-cores are allocated exclusively to a job, i.e., different jobs can't share a core
  - memory usage is included, i.e., allocating all memory of a node in jobs but only one core will "cost" you the core-hours of the whole node
- **GPU-hours or FPGA-hours:**
  - 1 GPU-hour = allocating one GPU for one hour
  - Usually in thousands of GPU-hours
  - GPUs are allocated exclusively to a job, i.e., different jobs can't share a GPU

-> Needed is a resource need per year

Additional:

- File system usage (space and number of files)
- Required main memory, maximal job runtime,...

How to Choose a Suitable Follow-up Project?

Possible Routes to Estimate the Required Resources:

1. **Personell-based estimation**
   - "I have N project members that each need M compute nodes or GPUs most of the time"
   - Fine as a rough rule of thumb
   - Example: 3 people using 4 GPUs most of the time = ~100k GPU-hours per year

How to Choose a Suitable Follow-up Project?

Possible Routes to Estimate the Required Resources:

1. **Personell-based estimation**
2. **Looking up what you have used in the NHR-Starter so far and extrapolating:**
   - On the cluster: pc2status
     (this shows only the usage on the current cluster)
   - HPC-Portal: https://portal.pc2.uni-paderborn.de/
     log in and go to Dashboard

```
Start: 01.04.2025
End: 31.03.2026
State: enabled

Phases and granted resources:
    01.04.2025 - 31.03.2026:
        CPU-Core hours                  : 4,000,000
        NVIDIA A100 GPU hours           : 10,000
Current granted resources:
        CPU-Core hours                  : 4,000,000
        NVIDIA A100 GPU hours           : 10,000
Used Resources in active phase:
        CPU-Core hours                  : 3,680.66
        NVIDIA A100 GPU hours           : 124.9
```

File System Quota

hpc-prf-autumn

File System	Cluster	Limit	Current Usage (%)
PC2PFS - Data	Noctua 1	5 TiB	4.5 GiB (0%)
PC2PFS - Files	Noctua 1	1,048,576	254 (0%)
PC2PFS - Data	Noctua 2	10 TiB	407.9 GiB (3%)
PC2PFS - Files	Noctua 2	10,000,000	674,743 (6%)
PC2DATA - Data	all Clusters	200 GiB	0 KiB (0%)
PC2DATA - Files	all Clusters	unlimited	0 (-)

Ressource	Cluster	U30*	C30
CPU-Core hours	Noctua 2	3,681	252,055
NVIDIA A100 GPU hours	Noctua 2	125	630

U30=usage in last 30 days

How to Choose a Suitable Follow-up Project?

Possible Routes to Estimate the Required Resources:

1. **Personell-based estimation**
2. **Looking up what you have used in the NHR-Starter so far and extrapolating**
3. **From the required calculations:**
   - (Number of Calculations) x (Number of cores used for the calculation) x (typical runtime of one calculation)
     1000 calculations x 256 cores x 48 hours = 12.3 mio. CPU-core-hours
   - Most accurate but also not easy in practice
   - this kind of estimation is required for the NHR proposals

Recommended estimation:

1. First try personell-based estimation and look as past usage
2. If now large changes are planned and if the result is clear within a project category, go for it
3. If not, try to estimated planned calculations

If you could use help with the estimation, contact us via pc2-support@uni-paderborn.de, consultation hours,...

How to Choose a Suitable Follow-up Project?

How precise does the estimate have to be?

- Requested/granted resources are **no hard limit** for usage but only determine the priority (details at https://doku.pc2.uni-paderborn.de/pages/1902070/Quality-of-Service+QoS+and+Job+Priorities)
- **Changes afterwards are possible:**
  - **Tier-3 (small) project:** resource changes up to the upper limit is possible by informal request to pc2-support@uni-paderborn.de
  - **NHR-normal/large project:** resources can be increased by 25% once by informal request to pc2-support@uni-paderborn.de

How to Choose a Suitable Follow-up Project?

Formal Eligibility

- **Principal investigator** must have a **Ph.D/Dr.** and is **member of a university** (incl. HAWs/FHs)
- **NHR projects:** any university in Germany

- **Tier-3 project (small):**
  - Local: Paderborn University
  - University in NRW:
    - Till end of 2025 Tier-3 supplementary services: if the university has no or no sufficient resources/services
    - But from 2026: new structure of state-wide Tier-3 HPC systems
      - hpcBasisCluster.nrw: general state-wide Tier-3 HPC system in Cologne
      - hpcFachCluster.nrw: science-specific state-wide Tier-3 HPC systems in Aachen (Engineering, AI), Paderborn (Physics, Chemistry, AI) and not yet determined third site
      - Universities (incl. HAWs/FHs) must contribute to be allowed access, more infos at https://hpc.dh.nrw/de/hpcfachcluster

How to Choose a Suitable Follow-up Project?

Effort for Proposal

- Rule of thumb: the more resources requested, the higher the effort for the proposal
- Two components:
  - Online form
  - Detailed description of the project
    - Using a template (latex and docx available)
    - Only needed for NHR-projects, for scientific review of the request
    - Simplifications possible:
      - For extensions of existing projects
      - For already reviewed projects (DFG/federal ministries/EU/Volkswagen Stiftung)
      - For projects using well-known programs
      - For multiyear-projects (one proposal for multiple years of resource usage)

How to Choose a Suitable Follow-up Project?

Example

Resource Needs

- Personell-based estimation: 2 people using on average 8 nodes of Noctua 2 (~18 mio CPU-Core-h)
- Looking up what you have used in the NHR-Starter so far and extrapolating: ~16 mio. CPU-core-h used is last year

-> probably NHR-Normal

- From the required calculations: 800 calculations x 512 cores x 48 hours = 19.6 mio. CPU-core-hours

-> clearly NHR-Normal (up to 30 mio. CPU-core-hours)

Formal Eligibility

- PI is member of university (incl. HAWs/FHs) in Germany

Effort for Proposal

- Online form: 30 minutes
- Detailed description with simplifications: afternoon to a day

Why all these Procedures and Regulation?

Short answer: The institutions (DFG, MWK, GWK,…) that fund the systems and services demand it.

Longer answer: Compute resources cost money
(investment, maintenance, power cost, storage system, …)

Rough estimates:
(details depend on cooling, hardware types,…)

- 1 mio. CPU-core-hours ≈ 2000-5000 Euro
- 1 thousand GPU-hours ≈ 500-1000 Euro

-> typical NHR-normal project:

- 20 mio. CPU-Core-hours+20k GPU-hours per year ≈ 50k-120k Euro worth

Overview: Possible Project Categories

	Tier-3 (small)	NHR-Normal	NHR-Large
Resources per year	up to **4M** CPU-core-h and up to **10k** GPU-h or FPGA-h	up to **30M** CPU-core-h and up to **100k** GPU-h or FPGA-h	**> 30M** CPU-core-h or **> 100k** GPU-h or FPGA-h
Location of PI	Paderborn or NRW (Tier-3 supplementary service)	Germany	Germany
Project Duration	One year	One year or multiyear	One year or multiyear
Proposal submission	**anytime**	**anytime**	**4 fixed dates** (first of Jan., Apr., Jul., Oct.)
Review	internal technical review Internal scientific review	internal technical review external scientific review	internal technical review external scientific review
Project start	a few days	preliminary start about one week full project starts 4-6 weeks	preliminary start about one week full project starts 3 months
Simplifications		1) Multiyear project 2) Scientifically reviewed project 3) well-known programs	1) Multiyear project 2) well-known programs

See also https://pc2.uni-paderborn.de/system-access

Possible Simplifications for NHR-Projects

1. Using well-known programs

- List at https://pc2.uni-paderborn.de/go/well-known-programs
- **-> No benchmarks or scalability required**

Chemistry/solid-state-physics/atomistic simulations				
Abinit	Fleur	LAMMPS	Octopus	VASP
Amber	Gamess	Lumen	ORCA	xtb/tblite
CP-PAW	GAUSSIAN	ms2	Psi4	Yambo
CP2K	Gromacs	Molpro	Quantum Espresso	
Dalton	GPAW	NAMD	Salmon	
FHI-aims	ls1 mardyn	NWChem	Turbomole	

Engineering			
Abaqus	ANSYS CFX/Fluent	CST Microwave	OpenFOAM
Adams	Comsol	LS-Dyna	StarCCM

High-energy physics
SIMULATeQCD

FPGA Development Tools		
Intel OneAPI FPGA Toolkit	OpenCL SDK for FPGAs	Vitis

Machine Learning			
DeepSpeed	Jax	PyTorch	Tensorflow/Keras

See also https://pc2.uni-paderborn.de/go/simplifications

Possible Simplifications for NHR-Projects

1. Using well-known programs

- List at https://pc2.uni-paderborn.de/go/well-known-programs
- **-> No benchmarks or scalability required**

2. Already scientifically reviewed project (NHR-Normal)

- by German federal ministry (BMBF, BMWK, BMUV, BMEL,...), DFG, GCS, European Union (EU projects), Volkswagen Stiftung, another NHR center
- e.g. TRR, CRC, ExIn, GRK, Emmy Noether,....
- Needed: Information about granted project (project id, GEPRIS-link,…) but NOT the review
- No full external scientific review required, only check for suitability of resource request
- **-> much shorter detailed description, shorter processing time**

3. Multi-year proposals

- **-> one proposal for multiple years of project runtime**

4. Extensions of NHR-Normal/Large projects

- **-> reduced detailed description**

See also https://pc2.uni-paderborn.de/go/simplifications

How to Apply for a Follow-up Project?

1. **Select suitable project category and simplifications** ✓
2. **Detailed description (NHR-Normal/Large)**
3. **Online form**
4. **Sign proposal form and send**
5. **Project proposal review**
6. **(preliminary) Project start/extension**

(step-by-step guides for the individual project types are availble at https://pc2.uni-paderborn.de/system-access)

How to Apply for a Follow-up Project?

Steps

1. **Select suitable project category and simplifications** ✓
2. **Detailed description (NHR-Normal/Large)**

- use templates (latex, docx, pdf) available at https://pc2.uni-paderborn.de/go/proposaltemplates
- Feel free to send us a draft if you are unsure
- Simplification shorten the description

1) Introduction (0.5-1 page)
2) Preliminary Work (0.5-1 page) -> only 0.5 page
3) Description of the Project (2-4 pages)
4) Numerical Methods and Algorithms (0.5-1 page)
5) Related Computational Work (0.5-1 page)
6) Personal and Previous HPC Experience (0.5 page)
7) Computational Methods and Programs
    - 7.1) program summary (table)
    - 7.2) parallel efficiency and scaling
    - 7.3) workflow
8) Justification of Requested Resources
    - 8.1) Estimation of Resources (table)
    - 8.2) Schedule for Resources Usage (table)
9) Special Resource and Support Requirements (optional)
10) References

How to Apply for a Follow-up Project?

Steps

1. **Select suitable project category and simplifications**
2. **Detailed description (NHR-Normal/Large)**

- use templates (latex, docx, pdf) available at https://pc2.uni-paderborn.de/go/proposaltemplates
- Feel free to send us a draft if you are unsure
- Simplification shorten the description
- **Using codes from the well-known programs list** (https://pc2.uni-paderborn.de/go/well-known-programs)
- **Already scientifically reviewed project** (by German federal ministry (BMBF, BMWK, BMUV, BMEL,...), DFG, GCS, European Union (EU projects), Volkswagen Stiftung, another NHR center

1) Introduction (0.5-1 page) -> only 0.5 page
~~2) Preliminary Work (0.5-1 page) -> only 0.5 page~~
~~3) Description of the Project (2-4 pages)~~
~~4) Numerical Methods and Algorithms (0.5-1 page)~~
~~5) Related Computational Work (0.5-1 page)~~
6) Personal and Previous HPC Experience (0.5 page)
7) Computational Methods and Programs
    7.1) program summary (table)
    ~~7.2) parallel efficiency and scaling~~
    7.3) workflow
8) Justification of Requested Resources
    8.1) Estimation of Resources (table)
    8.2) Schedule for Resources Usage (table)
9) Special Resource and Support Requirements (optional)
10) References

How to Apply for a Follow-up Project?

1. **Select suitable project category** ✓
2. **Detailed description (NHR-Normal/Large)** ✓
3. **Online Form**

   - **Crucial fields:** runtime/multiyear, resource volume (CPU-Core-h, GPU-h/FPGA-h), upload of detailed description, project information for already reviewed projects
   - **Important, but can be changed later easily:** cluster selection, distribution of compute resources to clusters, storage requirements
   - **Purely informational fields:** job characteristics (job duration, main memory demand, GPUs per job,...), details about applications used, ...

How to Apply for a Follow-up Project?

1. **Select suitable project category and simplifications** ✓
2. **Detailed description (NHR-Normal/Large)** ✓
3. **Online form** ✓
4. **Sign proposal form and send**
   - after submission of the online form, you will get a pdf to sign (by hand or electronically)
   - please send it to pc2-support@uni-paderborn.de

How to Apply for a Follow-up Project?

1. **Select suitable project category and simplifications** ✓
2. **Detailed description (NHR-Normal/Large)** ✓
3. **Online form** ✓
4. **Sign proposal form and send** ✓
5. **Project proposal review**
   - Formal and technical check of the proposal
   - Scientific review of the proposal (NHR normal: typically 2 reviewers, NHR Large: at least 2 reviewers, see http://pc2.uni-paderborn.de/go/reviewerguide-nhr)

How to Apply for a Follow-up Project?

Timelines

Tier-3 (small)

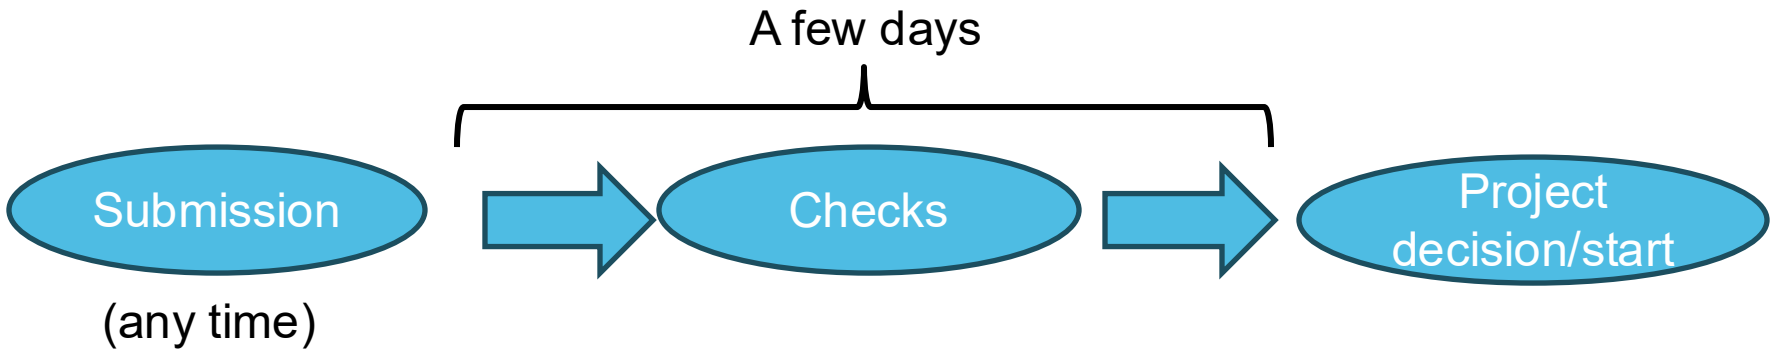

How to Apply for a Follow-up Project?

Timelines NHR Normal:

NHR-Normal:

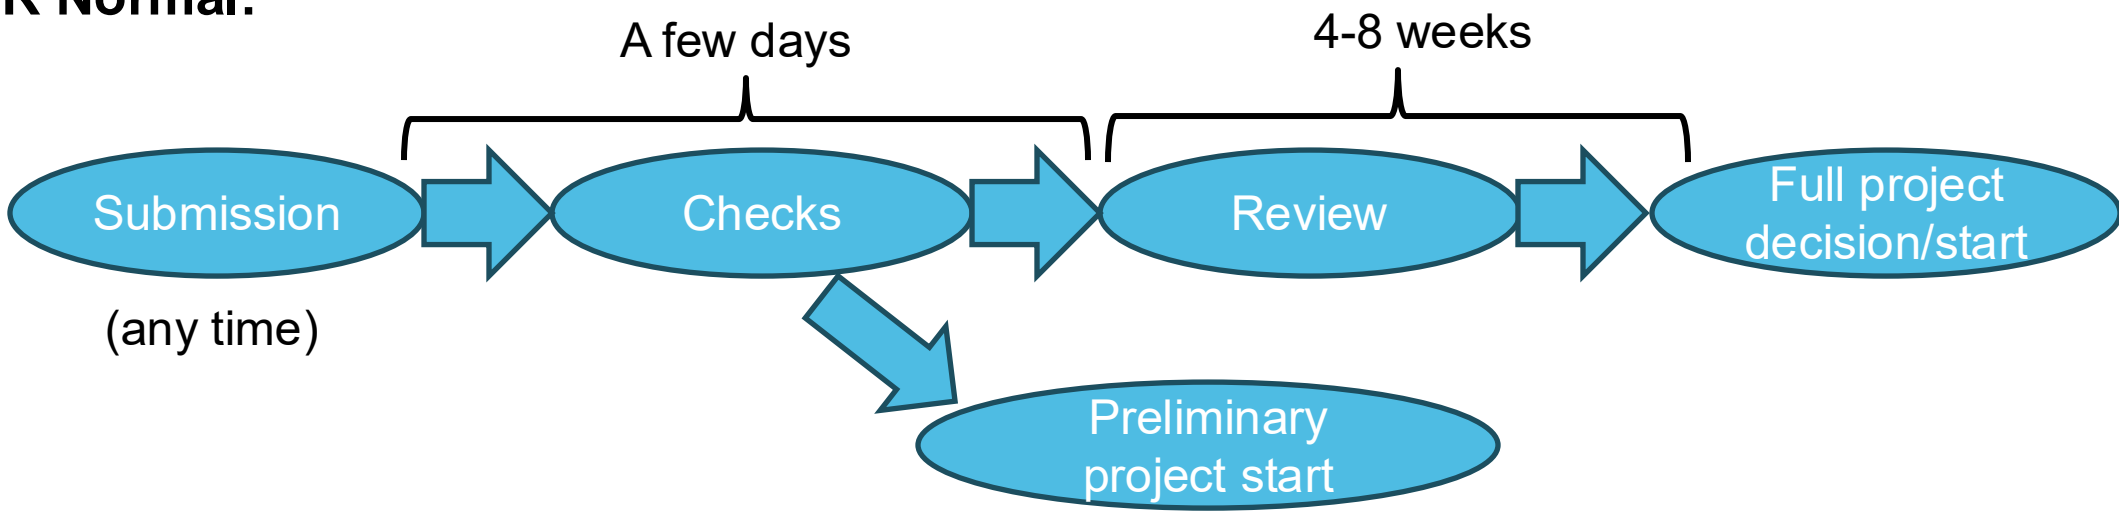

NHR-Normal already reviewed:

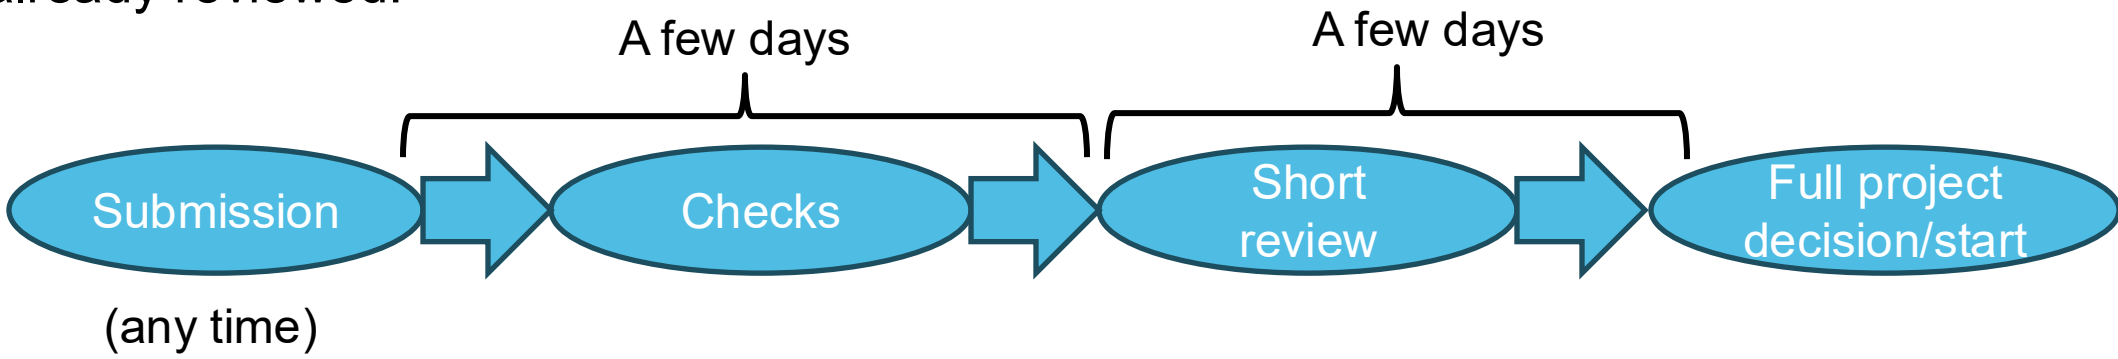

How to Apply for a Follow-up Project?

Timeline NHR Large:

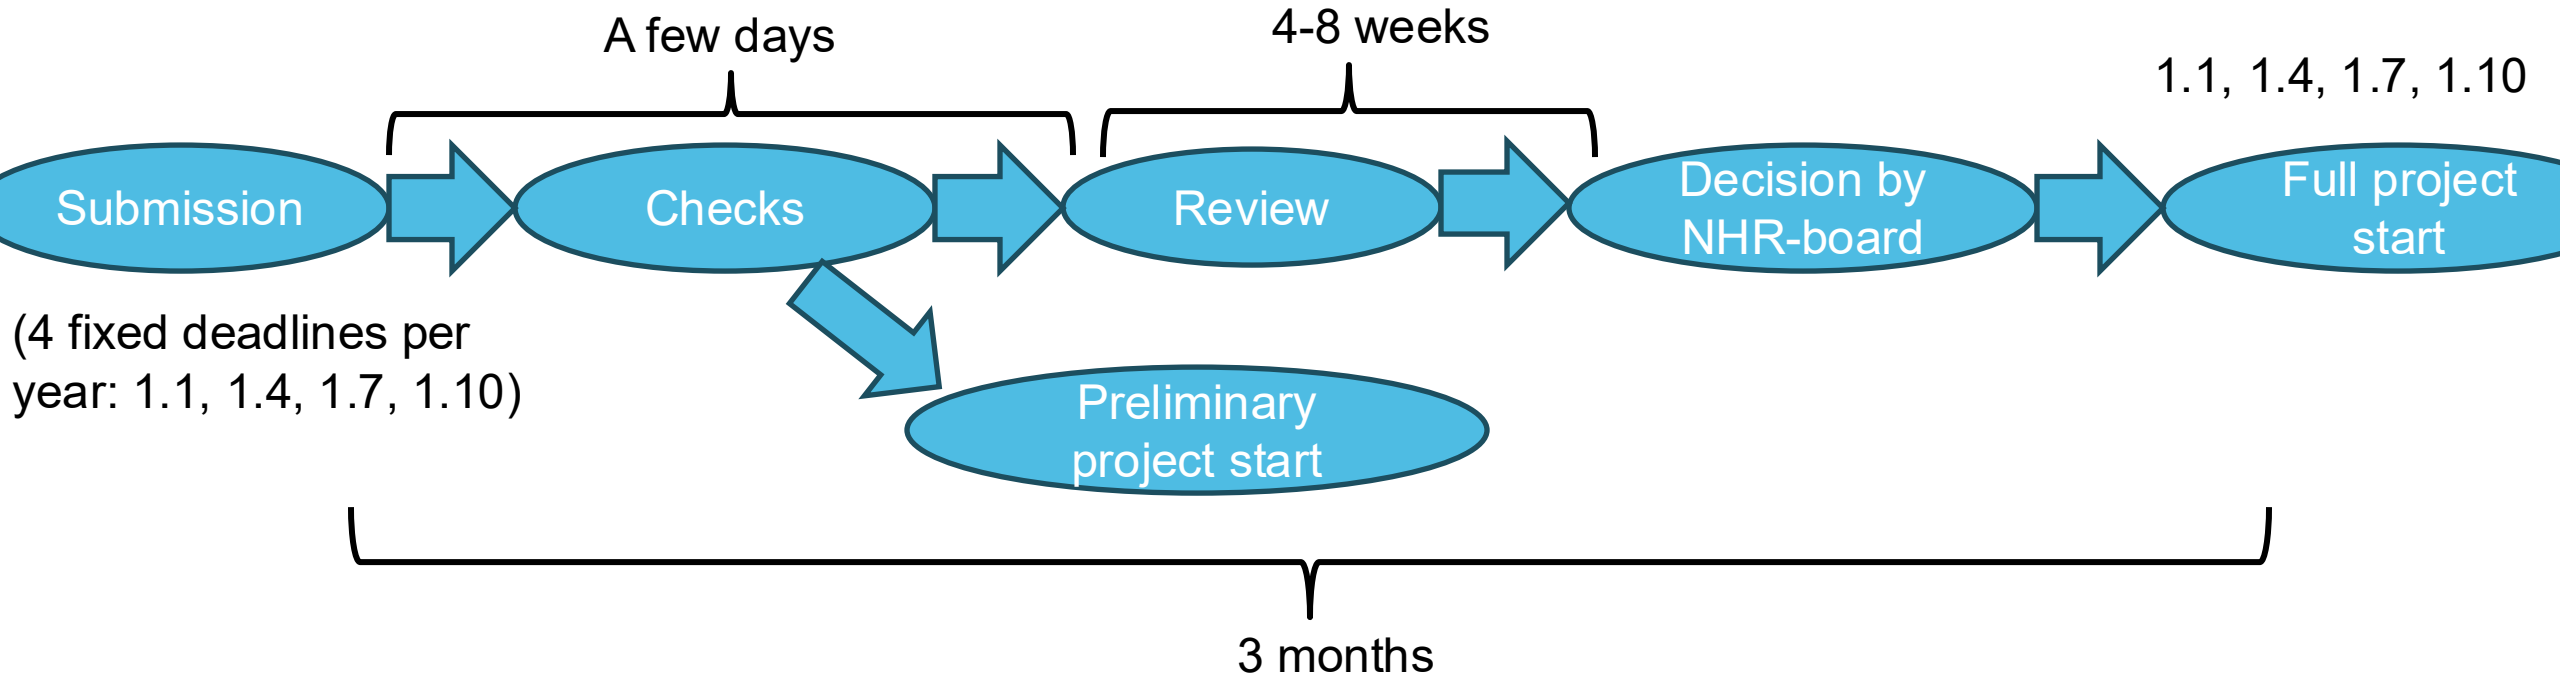

Note: we also accept proposals submitted a few days after the official deadlines

How can you get Support?

- Contact us via pc2-support@uni-paderborn.de

- Ask us in our regular weekly online consultation hour see https://events.uni-paderborn.de/category/13/

- We are also happy to schedule individual meetings to support you (in-person, online or phone).

For compute time proposals we can, for example, help you with:
- Estimating needed resources
- Choosing suitable project track
- Benchmarking
- Questions for online form and detailed description
- …