

# NHR centres in WLCG Workflows

Sebastian Wozniewski

PUNCH4NFDI AAI Meeting, Virtual – 27.06.25

# National High Performance Computing (NHR)

- association of **large, multi-disciplinary HPC centres**
- founded in 2020
- funded by the federal government and most of the federal states (~60 Mio. Euro / year)
- **provides compute time to university research groups** passing the review of a scientific committee (projects on a yearly basis)
- centres have individual scientific profiles and offer dedicated consulting, see <https://www.nhr-verein.de/informationen-zur-zentrenauswahl> (profiles not exclusive and not very sharply defined though)



# Transformation towards HL-LHC

**Gradual transition** from university-based Tier-2 centres to **NHR (CPU)** and **Helmholtz-Centres (mass storage)** towards beginning of HL-LHC, stretched over 5 years.

Chose 3 NHR centres with local WLCG site and work group (more important than NHR site profile):


- HoreKa@KIT (ATLAS/supervised by FR; CMS/supervised by KIT)
- EMMY@Göttingen (ATLAS/supervised by GÖ)
- CLAIX@Aachen (CMS/supervised by AA)

**Local ATLAS/CMS groups** keep supervising the NHR resources and apply for compute time on behalf of all involved German institutes.

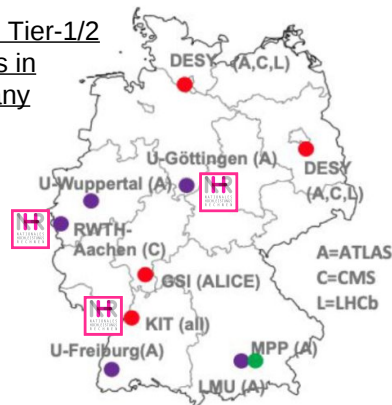
*NOT A STANDARD NHR USE CASE - REQUIRED PREPARATION IN ADVANCE AT BOTH MANAGEMENT AND SITE LEVEL*

Strategy paper by KET from 2022:

[https://www.ketweb.de/sites/site\\_ketweb/content/e199639/e312771/KET-Computing-Strategie-HL-LHC-final.pdf](https://www.ketweb.de/sites/site_ketweb/content/e199639/e312771/KET-Computing-Strategie-HL-LHC-final.pdf)



WLCG Tier-1/2 centres in Germany



Helmholtz Centres  
Max-Planck-Institute  
Universities

# Applications for Compute Time @ NHR

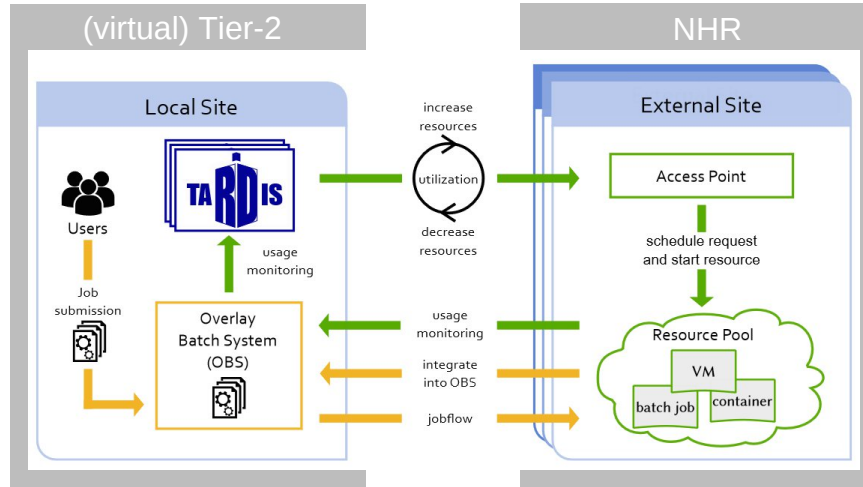
Standard application procedure:

- Different project categories:
  - Test account with O(100k CPU hours/Quarter) - can be requested any time
  - Medium/Large project (definition and quota depends on centre) with quarterly calls and review process for projects starting in the subsequent quarter (rolling call at some sites) - project duration is usually and maximally one year

WLCG case (after negotiations with NHR):

- So far sticking to yearly applications at each site but suboptimal for us (objectives don't change that fast in WLCG)
- Handing in applications 9 months earlier -> receive approval early enough to pledge resources to WLCG
- Selected sites according to our preferences (local contact and preparation, WLCG storage nearby for now), not strictly according to NHR profiles (we do not depend much on their research specific consultation anyways)

# Technical setup - Virtual Worker Nodes (Drones)



*Basic concept being implemented for all three NHR centres (twice for NHR@KIT)*

**As adopted within PUNCH:** Nodes are turned into virtual worker nodes (**drones**) of an overlay batch system (**OBS**). One main reason is the whole-node-scheduling policy of the NHR centre.

Required preparational discussions and negotiations (more or less at each site) with respect to:

- enabling unprivileged user namespaces (required for drones)
- outbound network connections for data transfers and OBS (often just via proxy or restricted to certain IP addresses)
- max. job lifetime, i.e. drone lifetime

# Further Comments

- Formally, we apply for compute time on behalf of the German ATLAS/CMS institutes in order to contribute to the computing of the ATLAS/CMS collaborations.
- Technically, we receive and process workloads via WLCG (simulation, data processing, user analysis...) => Workloads/project are not further specified in advance than the general physics goals of ATLAS/CMS. Keeping goals general is certainly helpful for PUNCH as well. The question if all PUNCH workloads could run under a single project (I doubt it) would be part of negotiations with NHR.
- Granularity of projects (e.g. by use case or research field) would need to be reflected in AAI groups and OBS conditions.