The PUNCH4NFDI Consortium

Particles, Universe, NuClei and Hadrons for the NFDI

Baida Achkar II. Institute - GAU for the PUNCH4NFDI Consortium

TA6-WP5 Meeting

19.10.2023

CONTENT

- **TA6-WP5-4: Deliverable details.**
- **TA6-WP5-4: Short term solution.**
- **TA6-WP5-4: Work in Progress: Acceptance Use Policy.**

PUNCH4NFDI@GAU/TA6-WP5-4/Deliverable details

- Involved people- Georg-August-University Goettingen:
  - Main contact: baida.achkar@phys.uni-goettingen.de
  - sebastian.wozniewski@uni-goettingen.de
- Due by 31 Dec 2025

1. Interfaces to the “HLRN – High Performance Computing in Northern Germany” in Goettingen will be developed to the PUNCH4NFDI communities.
2. Interfaces to the GPU cluster in Goettingen will be developed and provided for education and development purposes to the PUNCH4NFDI communities.
3. A fraction of the GoeGrid grid computing cluster in Goettingen will be provided to PUNCH4NFDI and beyond for analysis of the CERN open data to *users without explicit CERN or experiment affiliation.*

PUNCH4NFDI@GAU/TA6-WP5-4/External users offer

- The offer is meant to be available to anyone. Authentication/Authorization via university/institution affiliation is not a requirement.
- Closely related to TA4 for the long term solution with *central PUNCH accounts* and to the *Acceptance Use Policy AUP for user groups management*.
- Related to TA2-WP2 for technical implementation: restrict users to run only a dedicated container(wlcg-wn for instance), once we know how people get access to resources.
- TA7 may use the offer to provide educational material.
- Intermediate Solution Setup of HTCondor Overlay Batch System OBS with exclusively own (GoeGrid - Goettingen) resources:
  - Technical implementation within GoeGrid infrastructure: HTCondor OBS implemented, locally tested through running Compute4PUNCH tutorials that use CERN open data (ATLAS, CMS). Local user accounts used for the tests. The service is running and stable. *Details in the backup.*
  - **Work in progress: AUP draft.**

PUNCH4NFDI@GAU/TA6-WP5-4/Dynamic User Documentation

AUP WORK IN PROGRESS

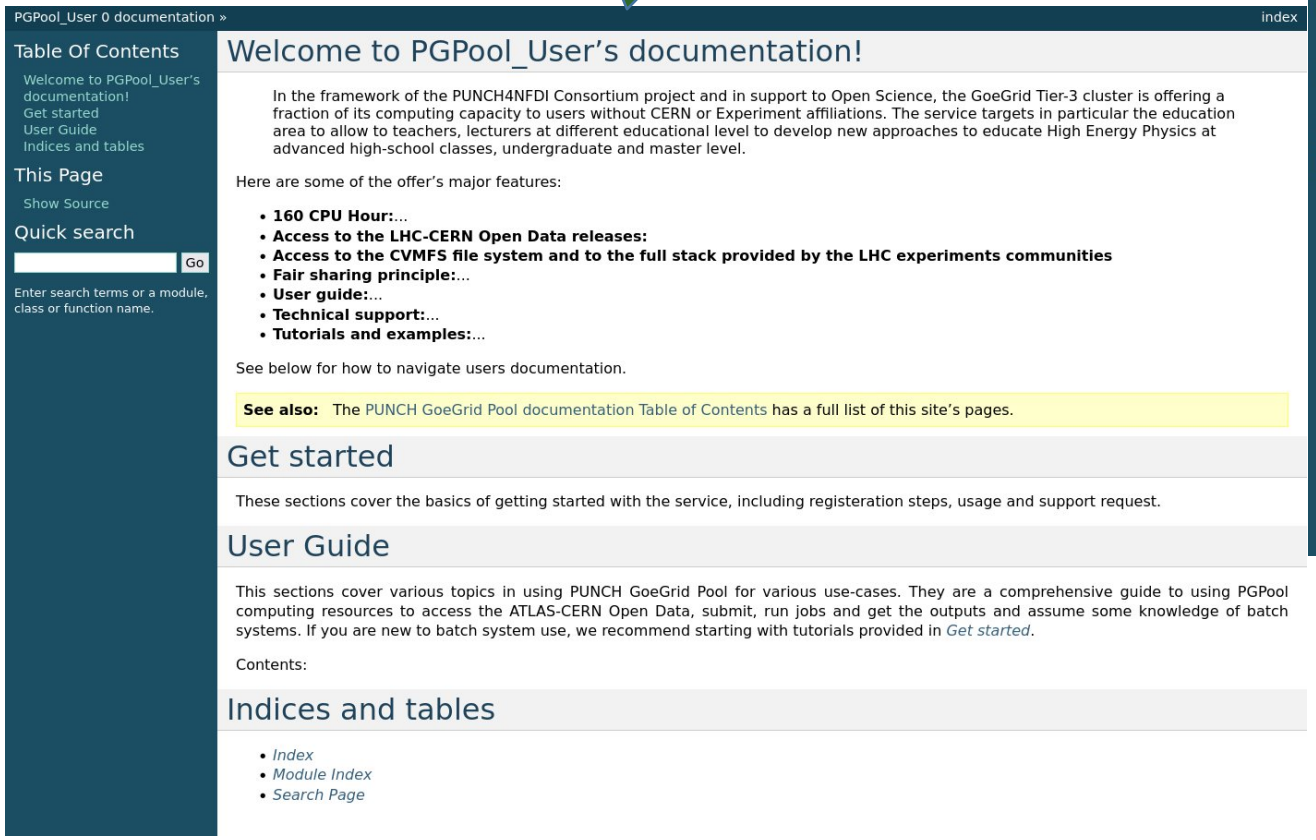

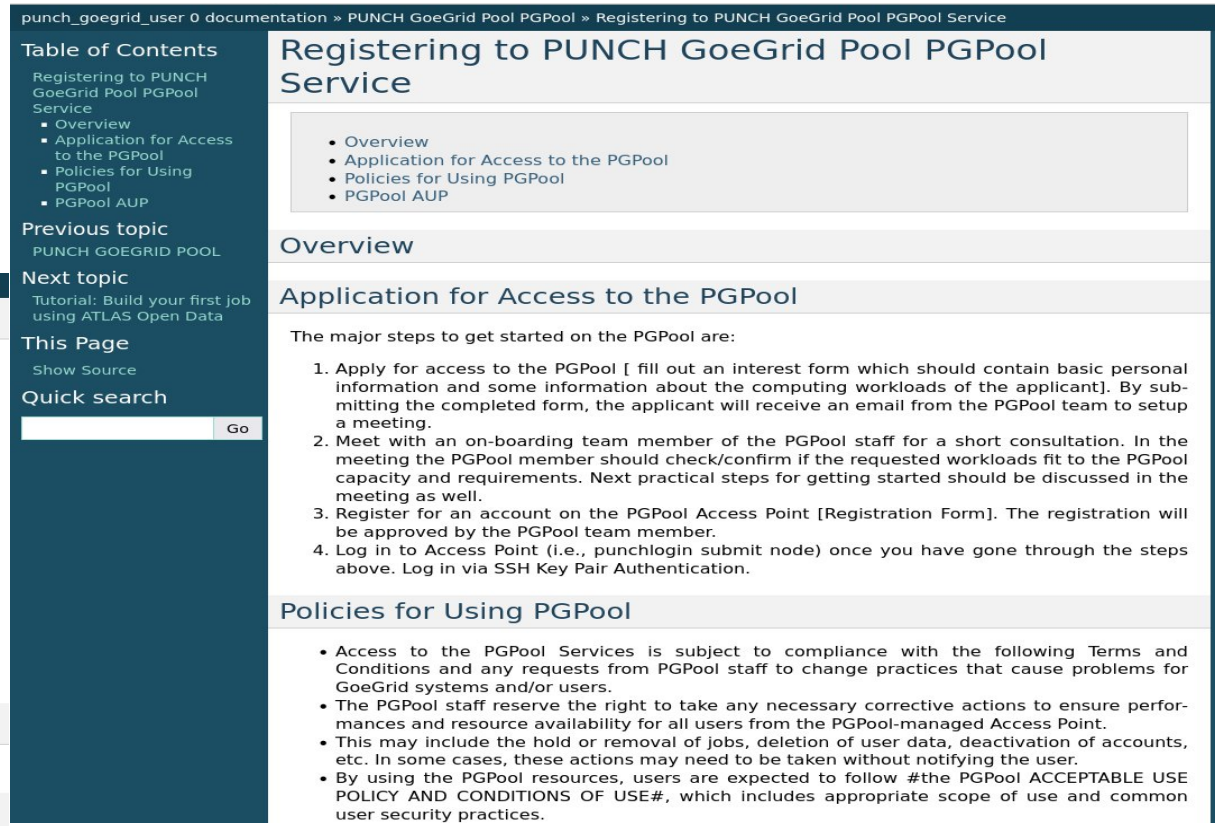

Backup

PUNCH GoeGrid Pool

PUNCH4NFDI@GAU / PUNCH GoeGrid Pool PGPool Setup

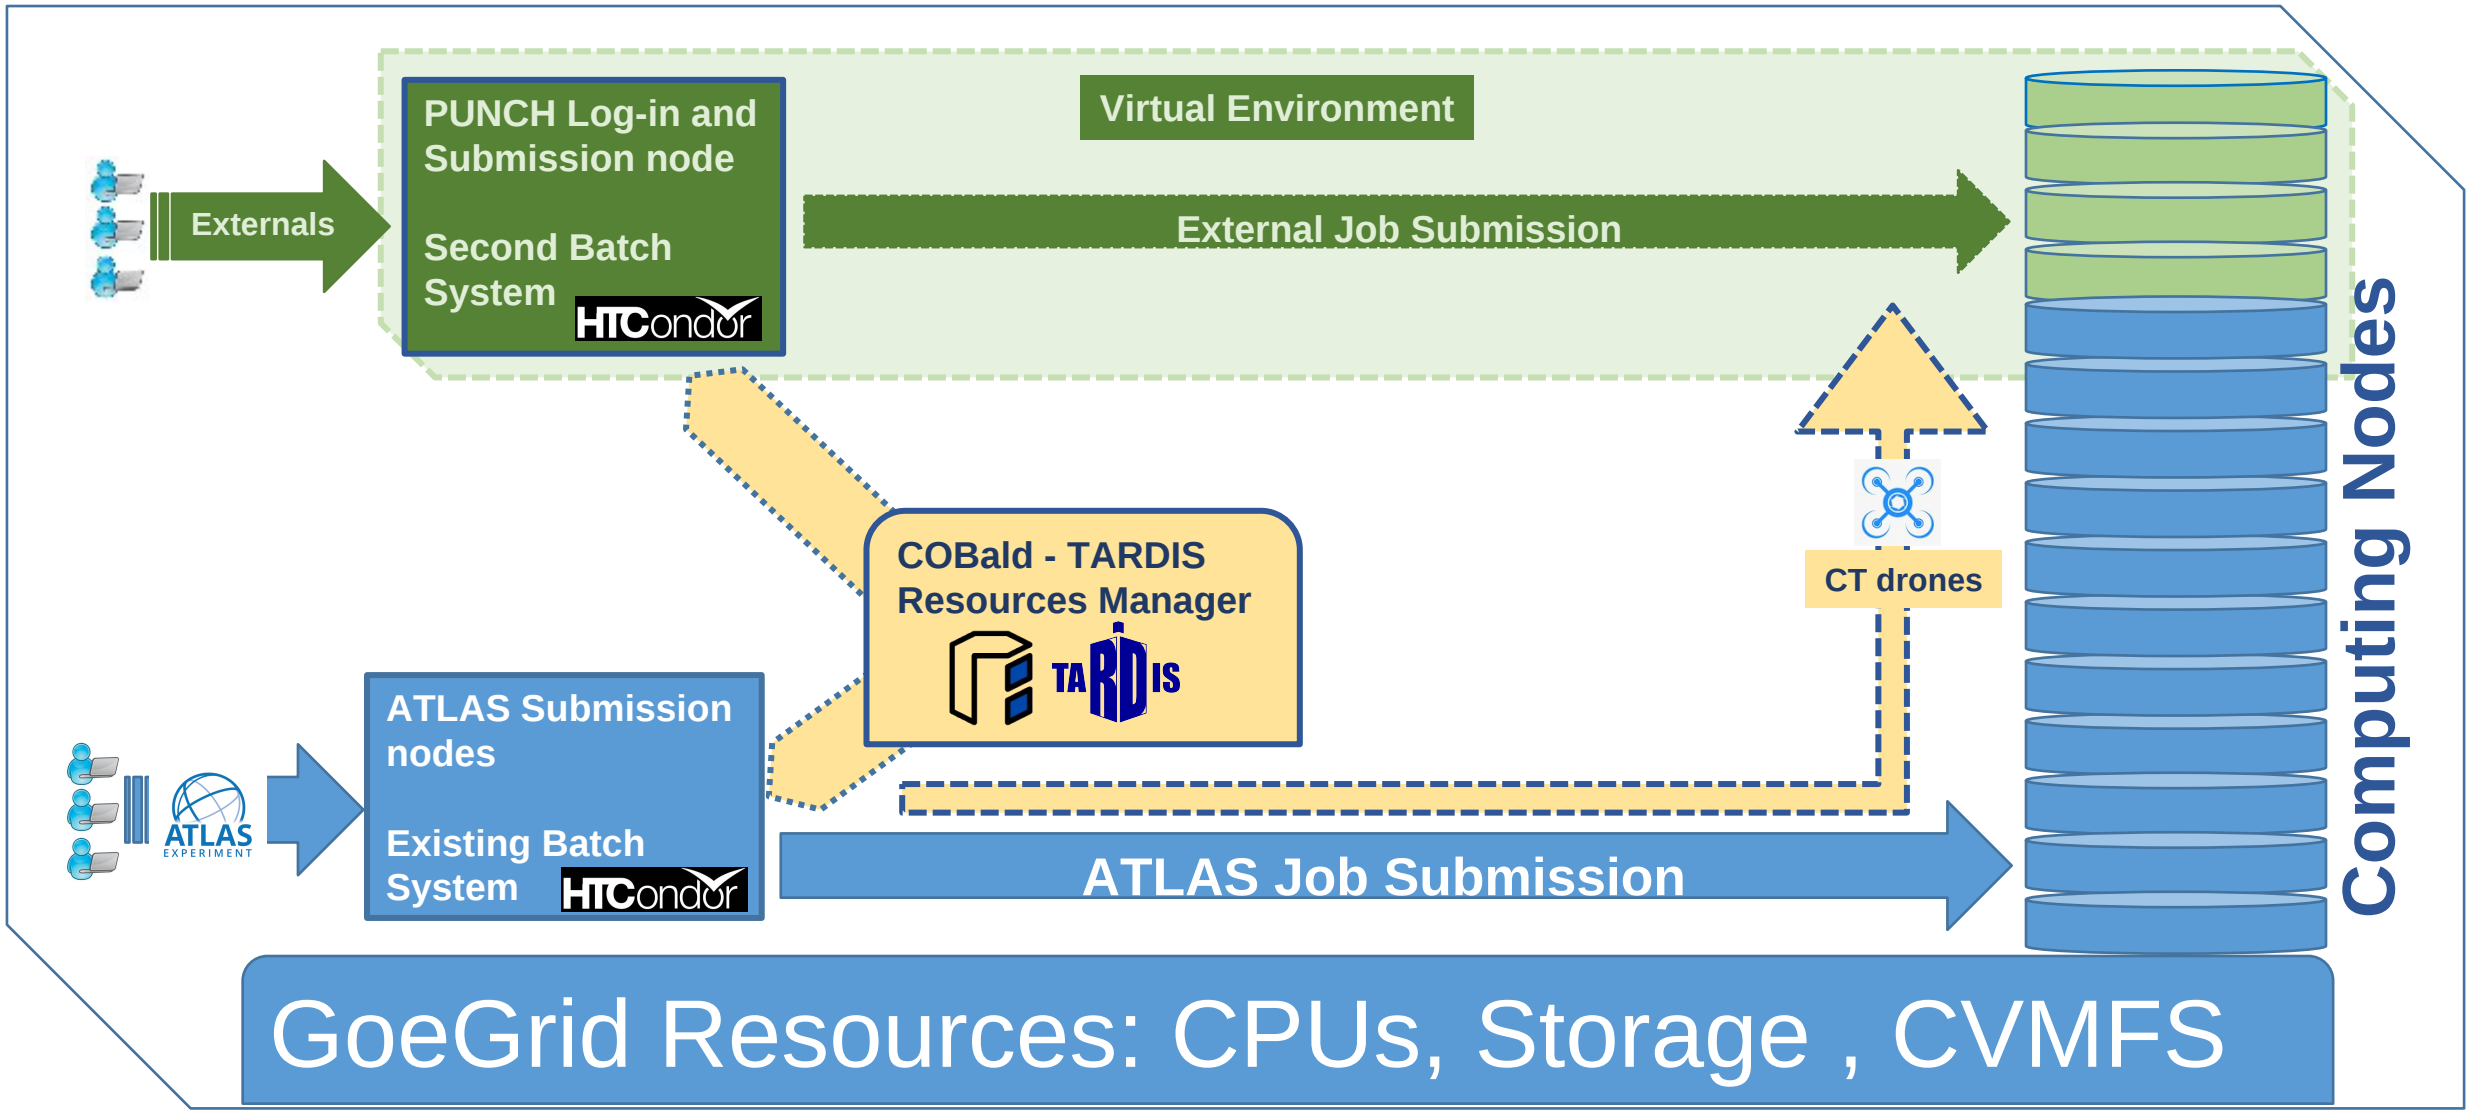

PUNCH4NFDI@GAU / PGPool Infrastucture

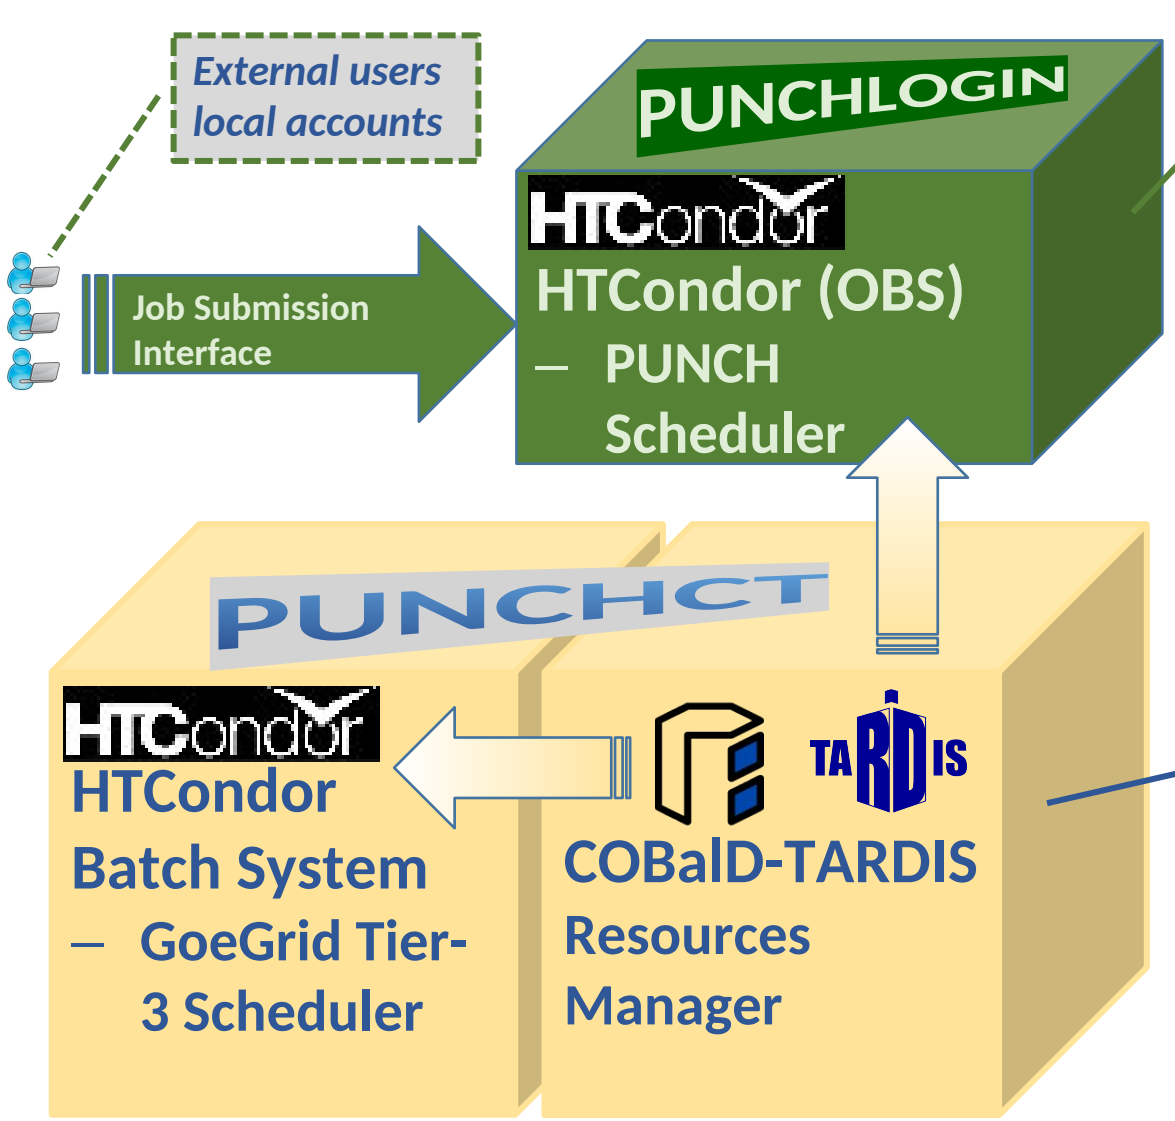

HTCondor Overlay Batch System OBS:
MASTER, COLLECTOR, NEGOTIATOR,
SCHEDD & GANGLIAD

- ✓ A COBalD-TARDIS instance runs as a system service:
  systemctl start cobald@cobald_obs_site
- ✓ Configured with:
  - ✓ HTCondor Batch System Adapter,
    - ✓ Options: pool: punchlogin.local
    - ✓ HTCondor Site Adapter:
      quota: 160
      drone_heartbeat_interval: 60
      drone_minimum_lifetime: 3600
- ✓ TARDIS starts drones (via Vanilla pilot.jdl jobs), which then are integrated themselves into the OBS using the HTCondor Batch System Adapter.
- ✓ Pilot jobs run under the accounting group Group_ATLAS.punchct on GoeGrid WNs

PUNCH@GAU / HTCondor-WN-CVMFS container

HTCondor-wn-cvmfs.SIF = 627M
Master + Startd

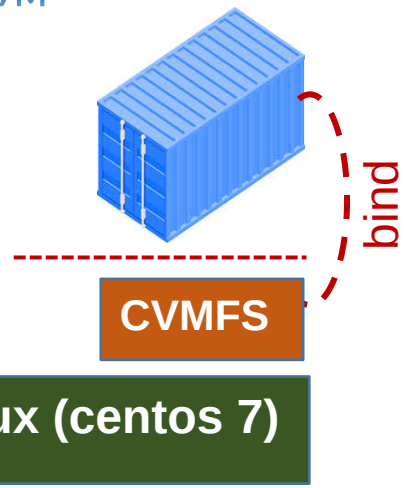

- ✓ HTCondor-wn drone built (using Singularity) based on **centos7 + htcondor 9.0.17**
- ✓ Configured using **ansible** + condor-git-config **hook:** Master + Startd
- ✓ To access the files stored within /CVMFS, the CVMFS directory is bind mounted from the the container host to a directory within the container. Environment variables adjusted
- ✓ Vanilla pilot job starts the drone to run on the GoeGrid working nodes
- ✓ **Pool status accessible via condor_status –pool punchlogin.local**
- ✓ **Ganglia service running for Pool web-based monitoring**
- ✓ **Drone Features:**

MachineTypes:
- **eightcore**

MachineTypeConfiguration:

eightcore:

jdl: /home/condor/htcondor-wn/pilot.jdl

MachineMetaData:

eightcore:

Cores: 8

Memory: 20

Disk: 160

```
[root@punchct htcondor-wn]# condor_status -pool punchlogin.local

Name                                      OpSys      Arch   State     Activity LoadAv Mem    ActvtyTime
slot1@goegrid-edbe37b7b7@compute-11-15.local LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:47
slot1@goegrid-b0280c1f2e@compute-12-2.local  LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:09:44
slot1@goegrid-d65020cb77@compute-15-4.local  LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:46
slot1@goegrid-e691adab5d@compute-15-7.local  LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:47
slot1@goegrid-720cdc2c2d@compute-15-13.local LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:49
slot1@goegrid-227ed94469@compute-15-14.local LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:47
slot1@goegrid-5a5bc70cba@compute-15-15.local LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:47
slot1@goegrid-cdd0eb9b0a@compute-16-11.local LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:47
slot1@goegrid-89624e331f@compute-16-14.local LINUX      X86_64 Unclaimed Idle      0.000 20480  0+00:14:47
                     Total Owner Claimed Unclaimed Matched Preempting Backfill  Drain
  X86_64/LINUX         9     0       0         9       0          0        0      0
         Total         9     0       0         9       0          0        0      0
```

PGPool Ganglia Monitoring

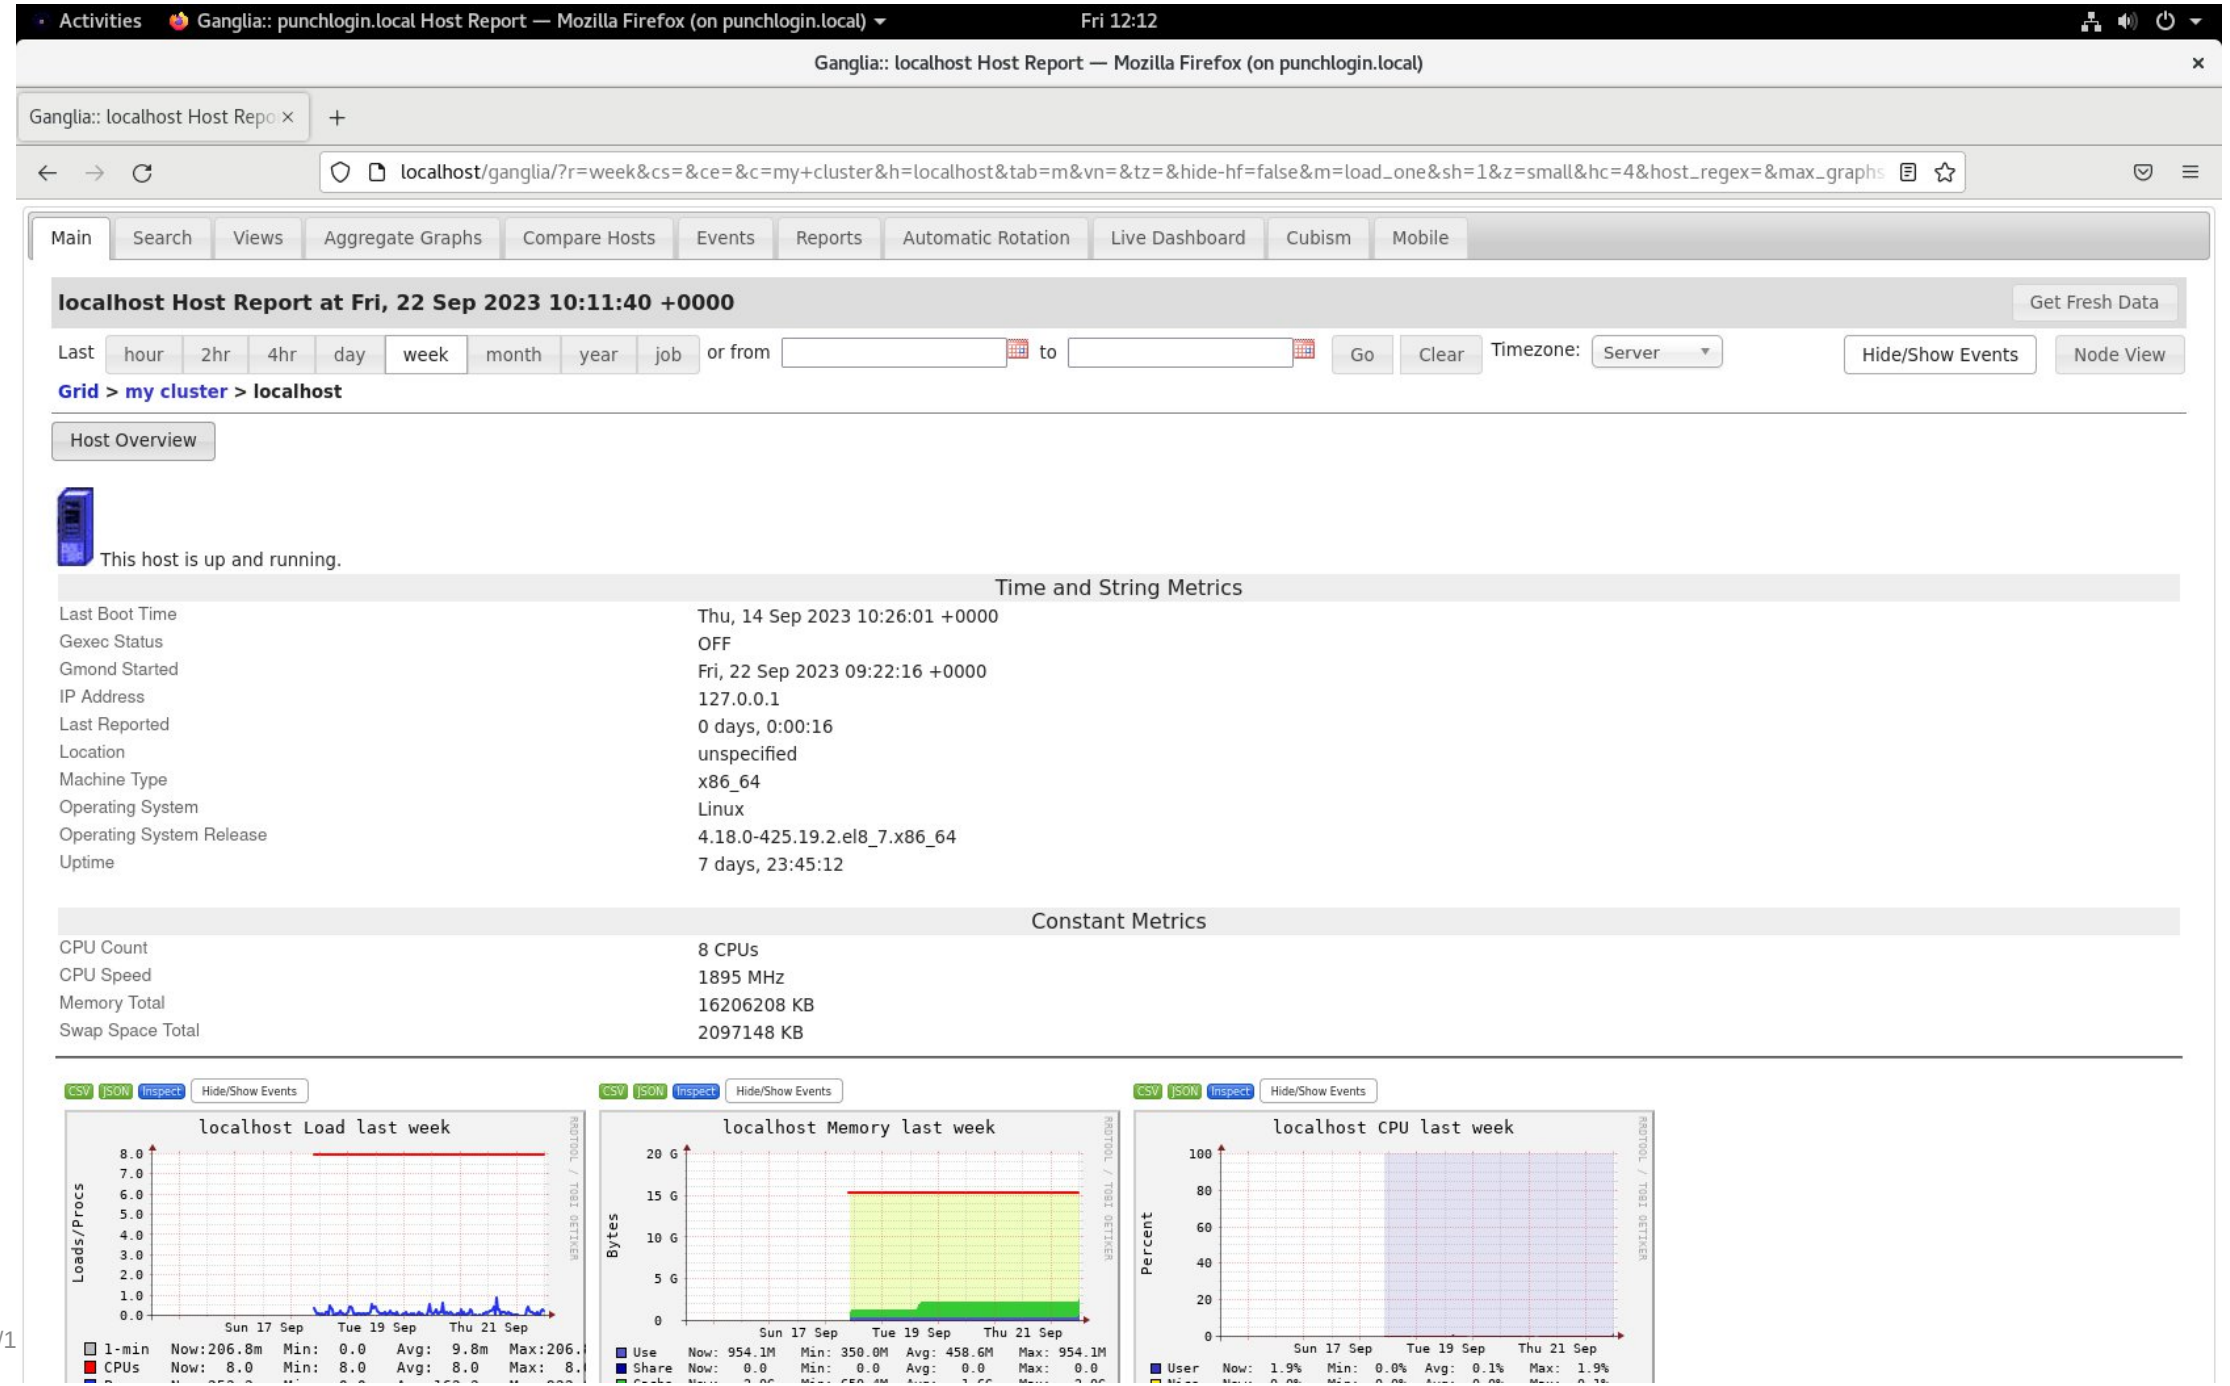

PGPool Ganglia Monitoring

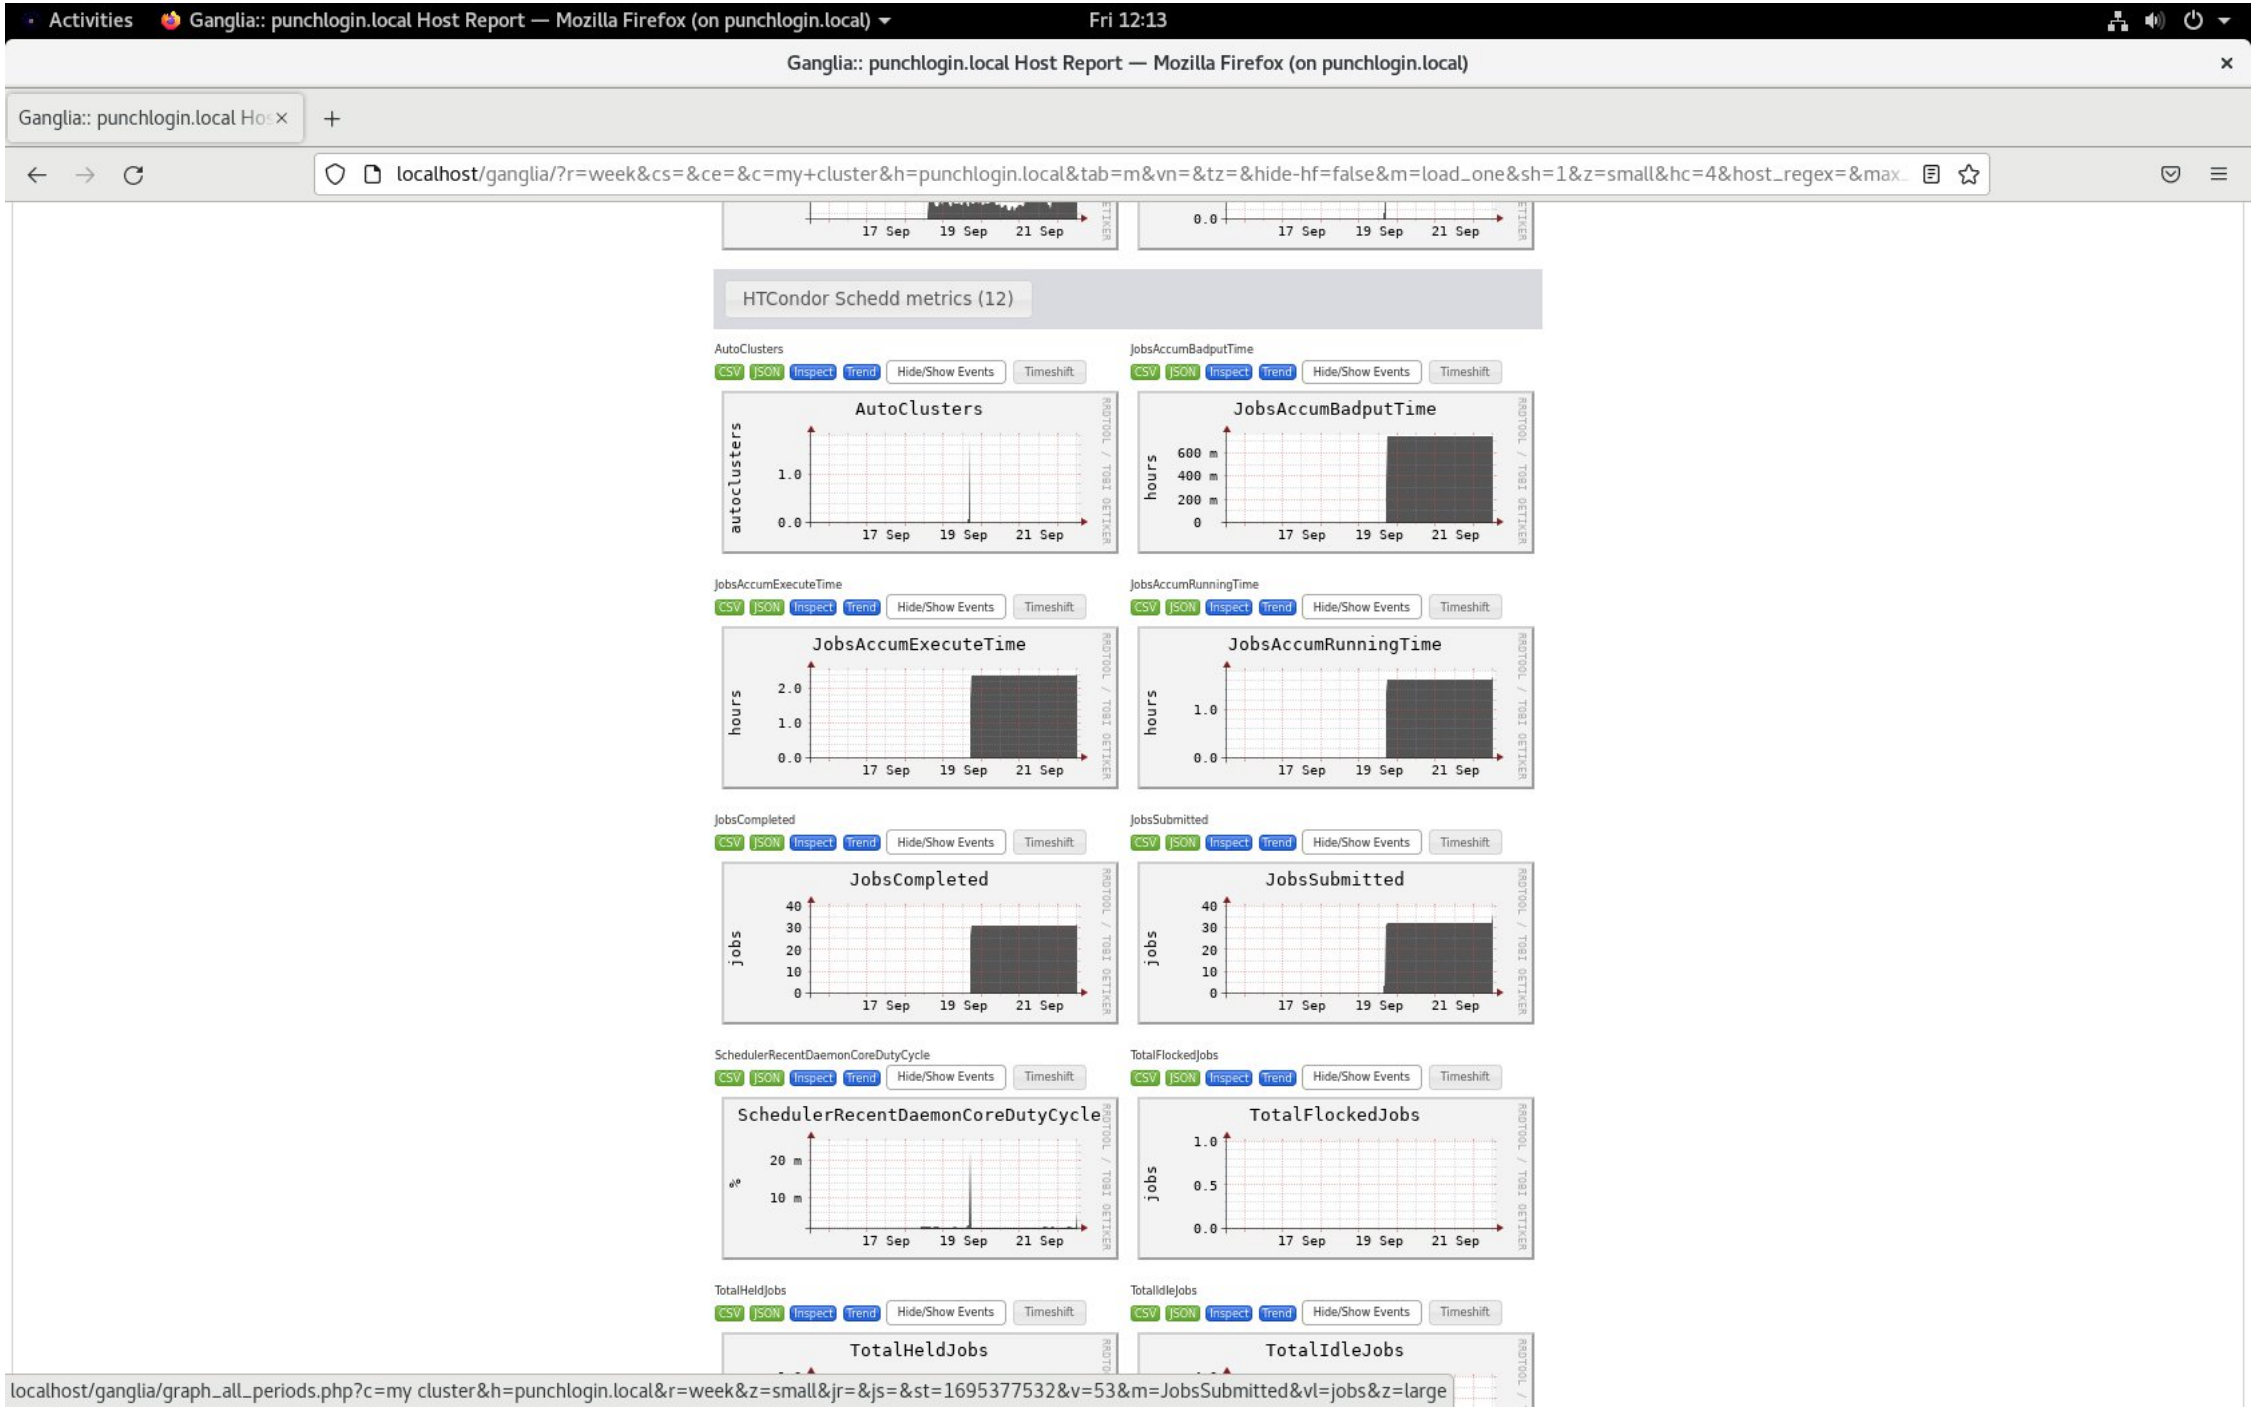

PUNCH@GAU / PGPool Submission Tests

Executable	compute4punch/tutorials/CMS_DiMuonSpectrum2012		**Ttbar_atlas tutorial** compute4punch/tatlas_open_data_ttbar
Data Set	**63673328** events. **5371** files. **20.6 TB** in total. CMS/Dimuon-2012		**58 Gb** opendata.atlas.cern/release/2020
External image (from Docker,..)?	Apptainer used to pull a docker root:latest image + Image used to run the executable	No	No
CVMFS access	No	Yes: setup root environment	
File transfer	Yes, the executable		Yes, executable, headers
Output	Output, error, log, . Spectrum as pdf		.root files (560K), png (240K)
Remarks	M Usage < 2000M D Usage< 900M	M Usage <1400M D Usage < 500K	Queue cmd in the jdl file used to run 29 processes / 1 cluster (1 proc per data set) M < 600M, D < 600K per process

PUNCH@GAU / PGPool Wrap up

- ✓ PGPool Quota set to 160 [20 drones x 8 CPU hour]. ~ 1% of GoeGrid Tier-2 capacity.
- ✓ COBalD-TARDIS resources manager starts pilot jobs to launch PGPool drones acting as virtual punch working nodes. /CVMFS available within the container. Pilot jobs run under a specific accounting group "Group_ATLAS.punchct" in GoeGrid.
- ✓ condor_status –pool punchlogin.local command line enabled on punchlogin and punchct machine for monitoring purposes.
- ✓ Ganglia web-based monitoring running on punchlogin node, basic + condor-specific metrics are visualized through firefox browser
- ✓ Access to the ATLAS/CMS-CERN Open Data , C4P tutorials, and analysis jobs submitted and completed successfully on PGPool.
- ➢ **First draft of PGPool dynamic user's documentation (including registration rules + AUP) is under development using Sphinx 5.02 tool.**

AUP draft presentation @TA6-WP5 meeting on 19th October!

- ➢ **Next: Integrate the setup into Compute4PUNCH infrastructure: Users can be restricted to use a dedicated container, once C4P knows how people outside get access to the resources**

AUP Draft

Registering to PUNCH GoeGrid Pool PGPool Service

AUP-WIP

Overview

Application for Access to the PGPool

The major steps to get started on the PGPool are:

1. Apply for access to the PGPool [fill out an interest form which should contain basic personal information and some information about the computing workloads of the applicant]. By submitting the completed form, the applicant will receive an email from the PGPool team to setup a meeting.
2. Meet with an on-boarding team member of the PGPool staff for a short consultation. In the meeting the PGPool member should check/confirm if the requested workloads fit to the PGPool capacity and requirements. Next practical steps for getting started should be discussed in the meeting as well.
3. Register for an account on the PGPool Access Point [Registration Form]. The registration will be approved by the PGPool team member.
4. Log in to Access Point (i.e., punchlogin submit node) once you have gone through the steps above. Log in via SSH Key Pair Authentication.

Policies for Using PGPool

- Access to the PGPool Services is subject to compliance with the following Terms and Conditions and any requests from PGPool staff to change practices that cause problems for GoeGrid systems and/or users.
- The PGPool staff reserve the right to take any necessary corrective actions to ensure performances and resource availability for all users from the PGPool-managed Access Point.
- This may include the hold or removal of jobs, deletion of user data, deactivation of accounts, etc. In some cases, these actions may need to be taken without notifying the user.
- By using the PGPool resources, users are expected to follow #the PGPool ACCEPTABLE USE POLICY AND CONDITIONS OF USE#, which includes appropriate scope of use and common user security practices.

AUP-WIP

This Page

Show Source

Quick search

Go

Enter search terms or a module, class or function name.

PUNCH GoeGrid Pool PGPool AUP

By registering as a user you declare that you have read, understood and will abide by the following conditions of use:

1. You can use the services and resources of the service only for activities that relate to the work that you described in the form when you applied for access.
2. Your registration can be suspended when you exceed the resources that were allocated for you during membership approval, or when your membership expires. (Membership is initially 6 months and can be extended for another 6 months).

No need for this point, OR: You agree on the fair-share CPU jobs scheduling adapted by the PGPool operator for ensuring that CPU resources are allocated fairly and efficiently among different users, and you will not complain about the scheduling time length.

3. You shall only use the resources/services to perform work, or transmit or store data consistent with the stated goals, policies and conditions of use as defined by the resources provider and by the PUNCH4NFDI Science Data Platform operator.
4. You shall provide appropriate acknowledgment of support or citation for your use of the resources/services provided. The following acknowledgment text can be used for this purpose (e.g. in scientific publications):
5. You shall not use the resources/services for any purpose that is unlawful and not (attempt to) breach or circumvent any administrative or security controls.
6. You shall respect intellectual property and confidentiality agreements.
7. You shall protect your access credentials (e.g. private keys or passwords).
8. You shall keep all your registered information correct and up to date.

AUP-WIP

9. You shall immediately report any known or suspected security breach or misuse of the resources/services or access credentials to baida.achkar@uni-goettingen.de and to the relevant credential-issuing authorities.
10. You use the resources/services at your own risk. There is no guarantee that the resources/services will be available at any time or that their integrity or confidentiality will be preserved or that they will suit any purpose.
11. You agree that logged information, including personal data provided by you for registration purposes, may be used for administrative, operational, accounting, monitoring and security purposes. You agree that this logged information may be disclosed to other authorized participants via secured mechanisms, only for the same purposes and only as far as necessary to provide the services.
12. You agree that the body or bodies granting you access and resource/service providers are entitled to regulate, suspend or terminate your access without prior notice and without compensation, within their domain of authority, and you shall immediately comply with their instructions.

You agree that the PGPool operator/ administrator is entitled to regulate, suspend or terminate your access without prior notice and without compensation, within their domain of authority, and you shall immediately comply with their instructions.

13. You are liable for the consequences of your violation of any of these conditions of use, which may include but are not limited to the reporting of your violation to your home institute and, if the activities are thought to be illegal, to appropriate law enforcement agencies. You are liable for the consequences of your violation of any of these conditions of use, which may lead, if the activities are thought to be illegal, to appropriate law enforcement agencies.

Note: These policies are derived from AUP-EGI-Document