TA2 Status, Highlights and Plans
PUNCH general meeting, 19 November 2025

Particles, Universe, NuClei and Hadrons for the NFDI

Alexander Drabent on behalf of the TA 2 workforce

Storage4PUNCH

- Based on two storage technologies
  - dCache (Instances at DESY and KIT
  - XrootD (Instances at U Bonn & GSI)
- Token based access using PUNCH AAI
- AAI is still limited, however a paid feature request is on the way to enable some new capabilities

- Rucio
  - Commonly used DM tool in HEP and elsewhere
  - Provides file catalog and rule based data replication between storage instances
  - Testbed service for PUNCH setup by HIFIS with still limited capabilities

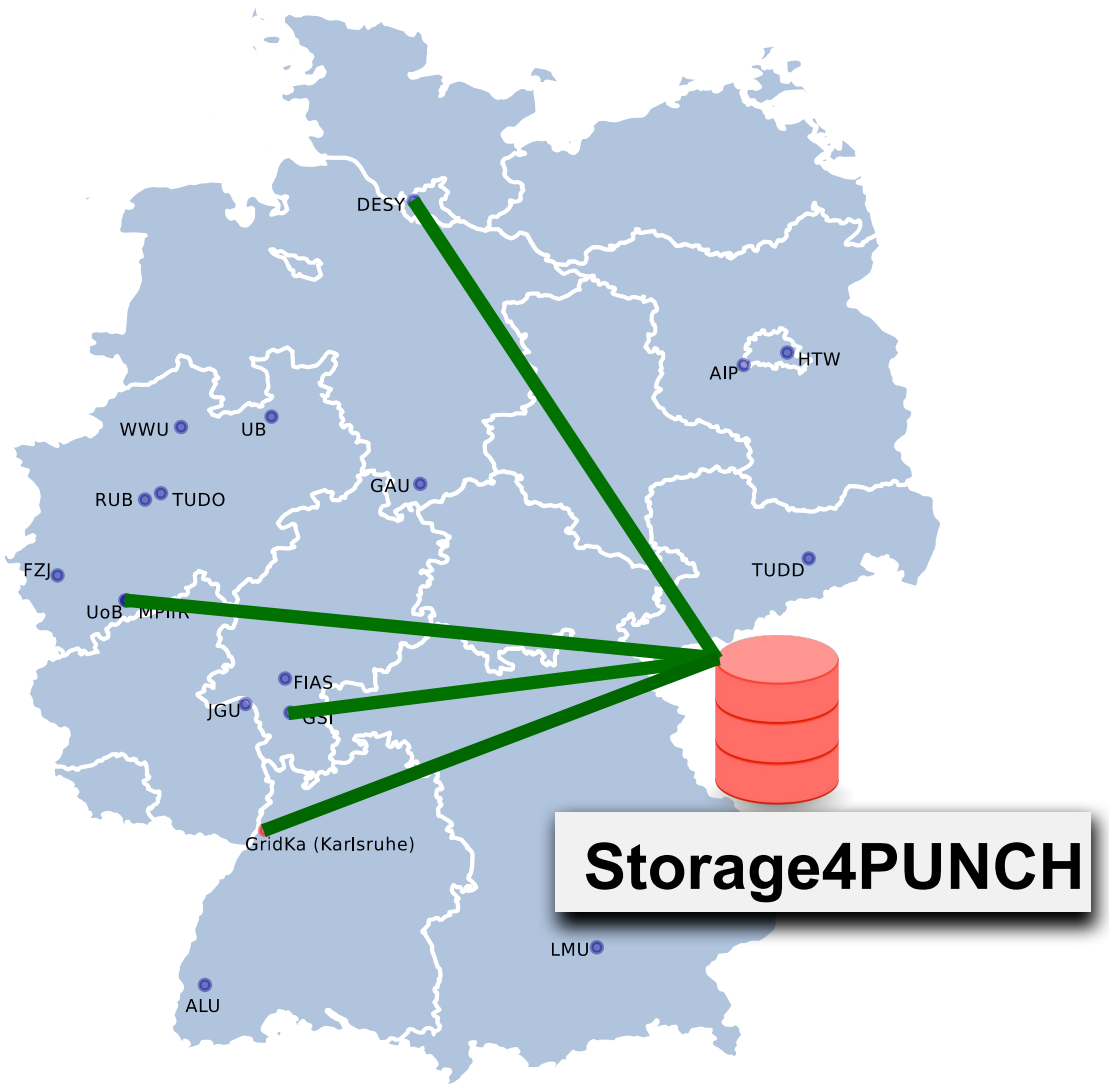

DFG Call for NFDI Storage “Datenspeicher”

- Should address needs of all consortia
- Only universities eligible for funding
- Significant amount of storage **requested**
  - Funding decision is pending
- Planned storage capabilities beyond present Storage4PUNCH functionality
  - Object storage, e.g. S3
  - POSIX access
- Storage4PUNCH prototyping exposed a lot of challenges. Many of them remain to be addressed in PUNCH2.0

Provider	**Capacity (TB)**	**Share in 3-5y (%)**
GAU Goettingen	4.000	8
Goethe-Uni Frankfurt	10.000	21
JGU Mainz	5.000	11
KIT	7.000,75	15
LMU/ TUM	8.000	17
RWTH Aachen	4.000	8
TU Dresden	3.000	6
Total		**86**

AAI

PUNCH-AAI based on Helmholtz AAI

- Implemented using the UNITY AAI toolkit
- Provided by HIFIS

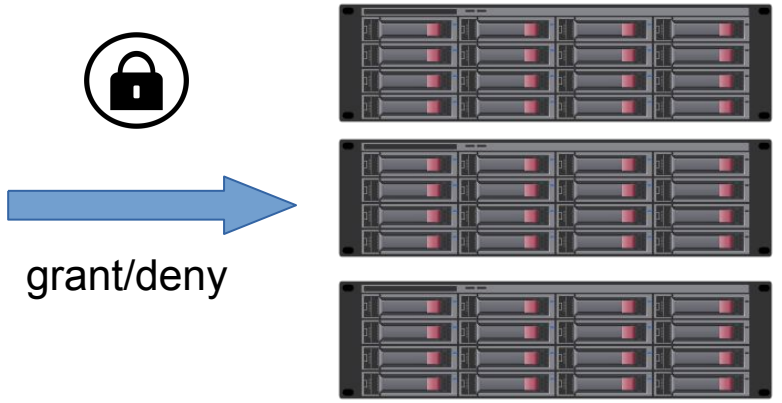

Present system lacks some features beneficial for PUNCH

- (Paid) Development request towards UNITY
  - Claim filter: Select which capability get embedded into token. Use tokens with just the required permissions
  - Extended claims: AAI manages access rights for directories, permissions for certain resources
  - Interface to a Policy Engine, e.g. Open Policy Engine to maintain access rights
  - Expected to be finished in Dec ‘25
- Adoption of client and services to cope with more complex tokens

Quite some work can realistically only be in PUNCH2.0

Evolution of Compute4PUNCH

- C4P fully integrated into REANA production instance hosted at AIP
- Dynamic integration of various heterogeneous compute resource of providers established
  - First successful job flocking to Mainz (HPC w/o outgoing connectivity). Working on production deployment
- Monitoring data of available resources is now transmitted to InfluxDB at AIP (REANA Dashboard)
- Container Registry 2.0 for enhanced container provisioning
- Some features in Helmholtz AAI are worked on (compute/storage scopes, MFA, data processing agreement → JupyterHub)

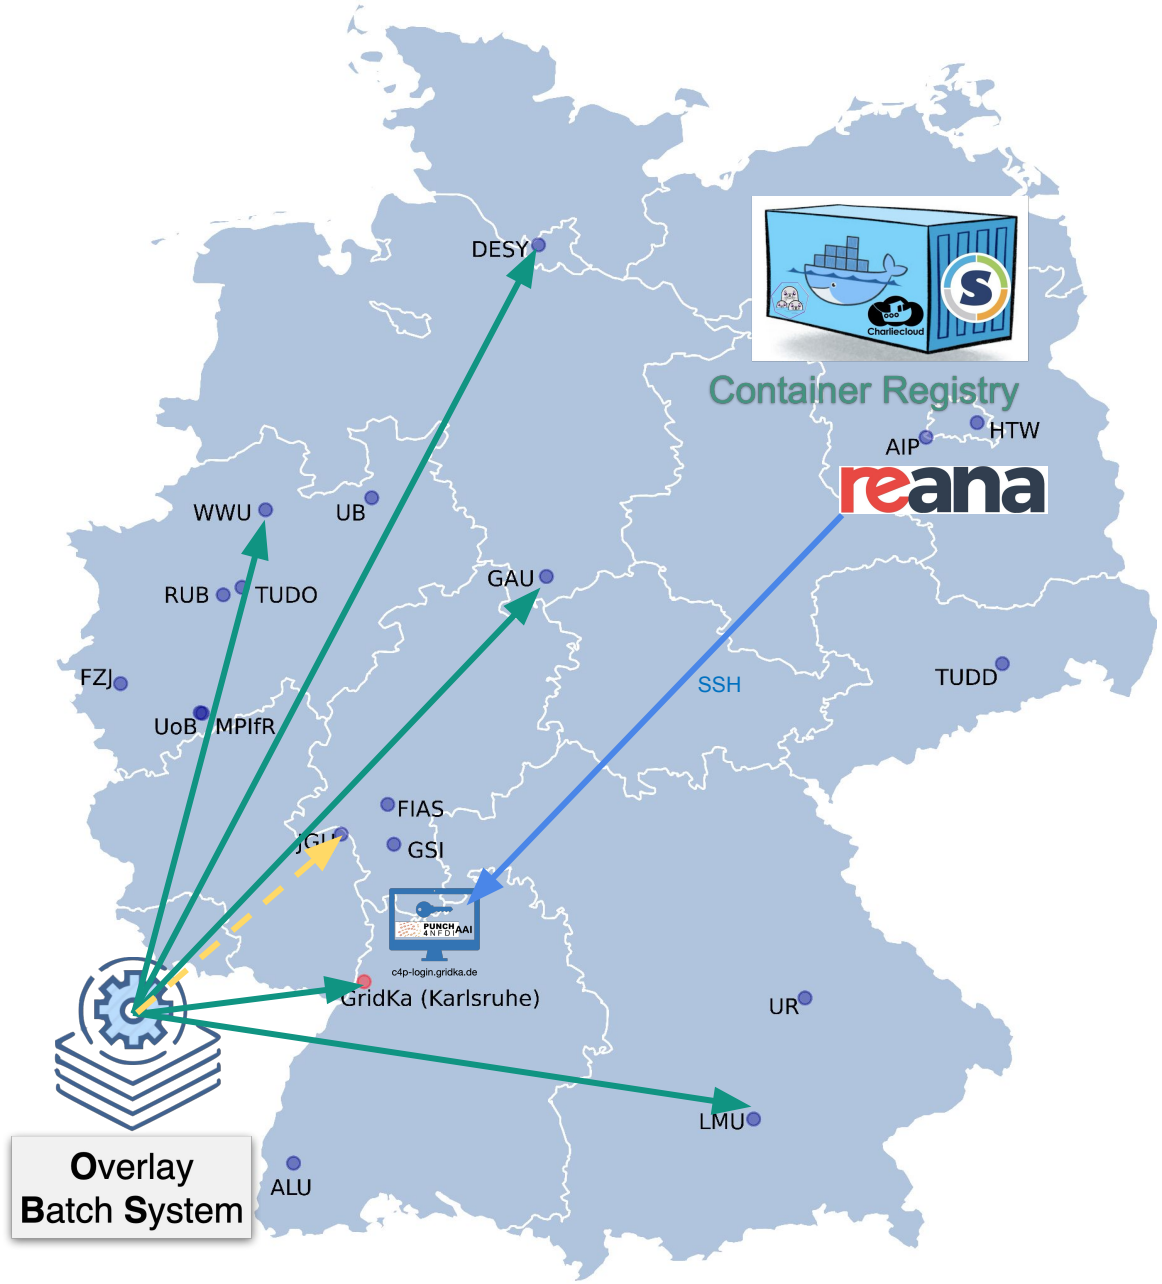

Compute4PUNCH Monitoring Dashboard

Grafana Dashboard

- Hosted at KIT
- Authentication via Helmholtz AAI to prevent anonymous data scraping
- Displays detailed information on Compute4PUNCH compute resources and jobs running
- Enables external data processing and analysis
- Technical Stack
  - InfluxDB for time-series data storage
  - Telegraf for data collection and metrics forwarding
- Data can be written to multiple InfluxDB instances simultaneously
- Worker node data also storages in AIP InfluxDB for upcoming REANA dashboards

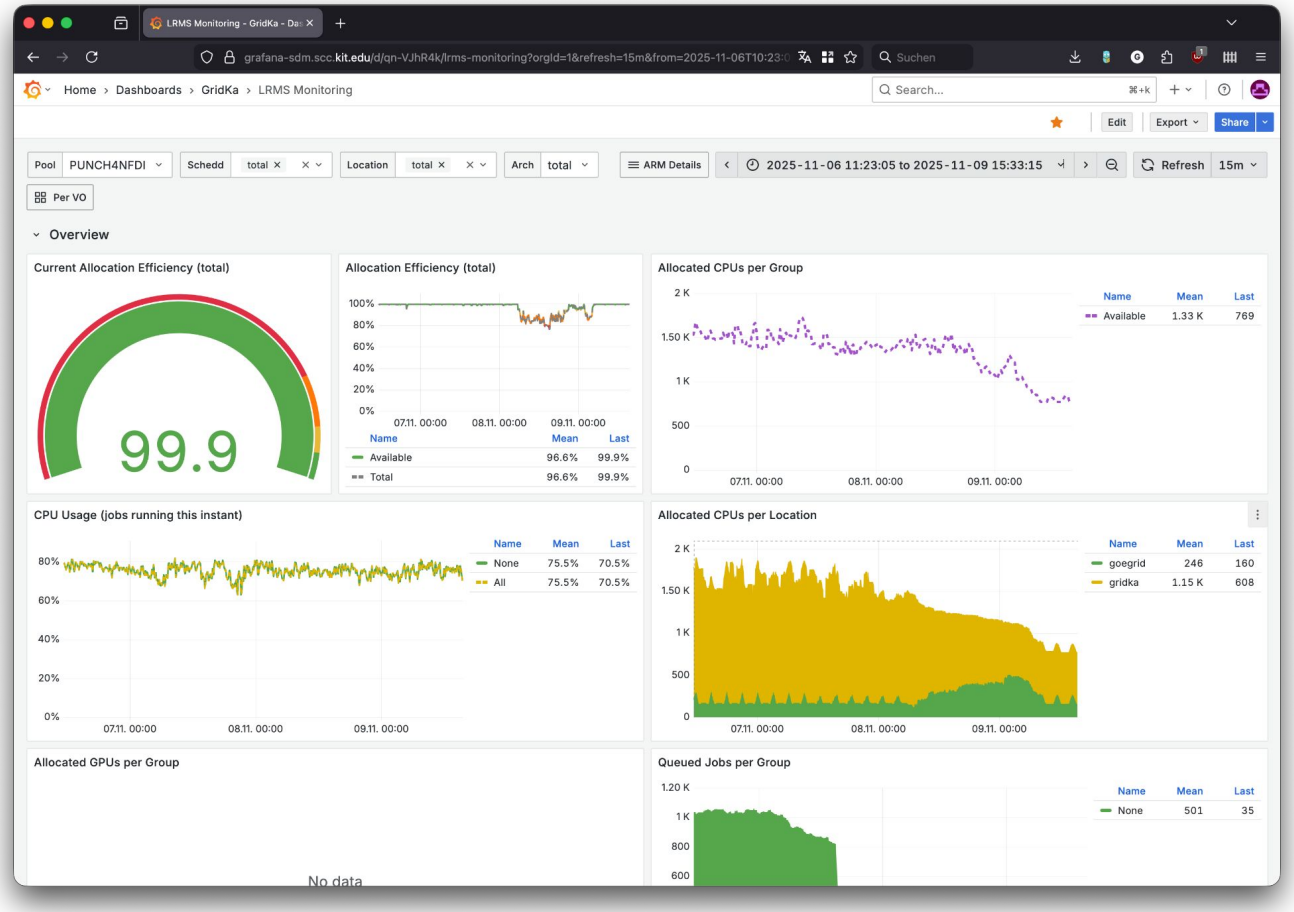

https://grafana-sdm-open.scc.kit.edu/login

Publicly Accessible via Helmholtz AAI

Survey on PUNCH Services

- Survey on PUNCH service
  - Initiated by TA6 and User Board
  - Still running: LINK
- Feedback on C4P
  - Known to most participants
  - Used by about half of participants knowing about it
  - Reasons for not using C4P
    - Access to other CPU resources
    - Need for HPC like resources

https://survey.hifis.dkfz.de/721319?lang=en

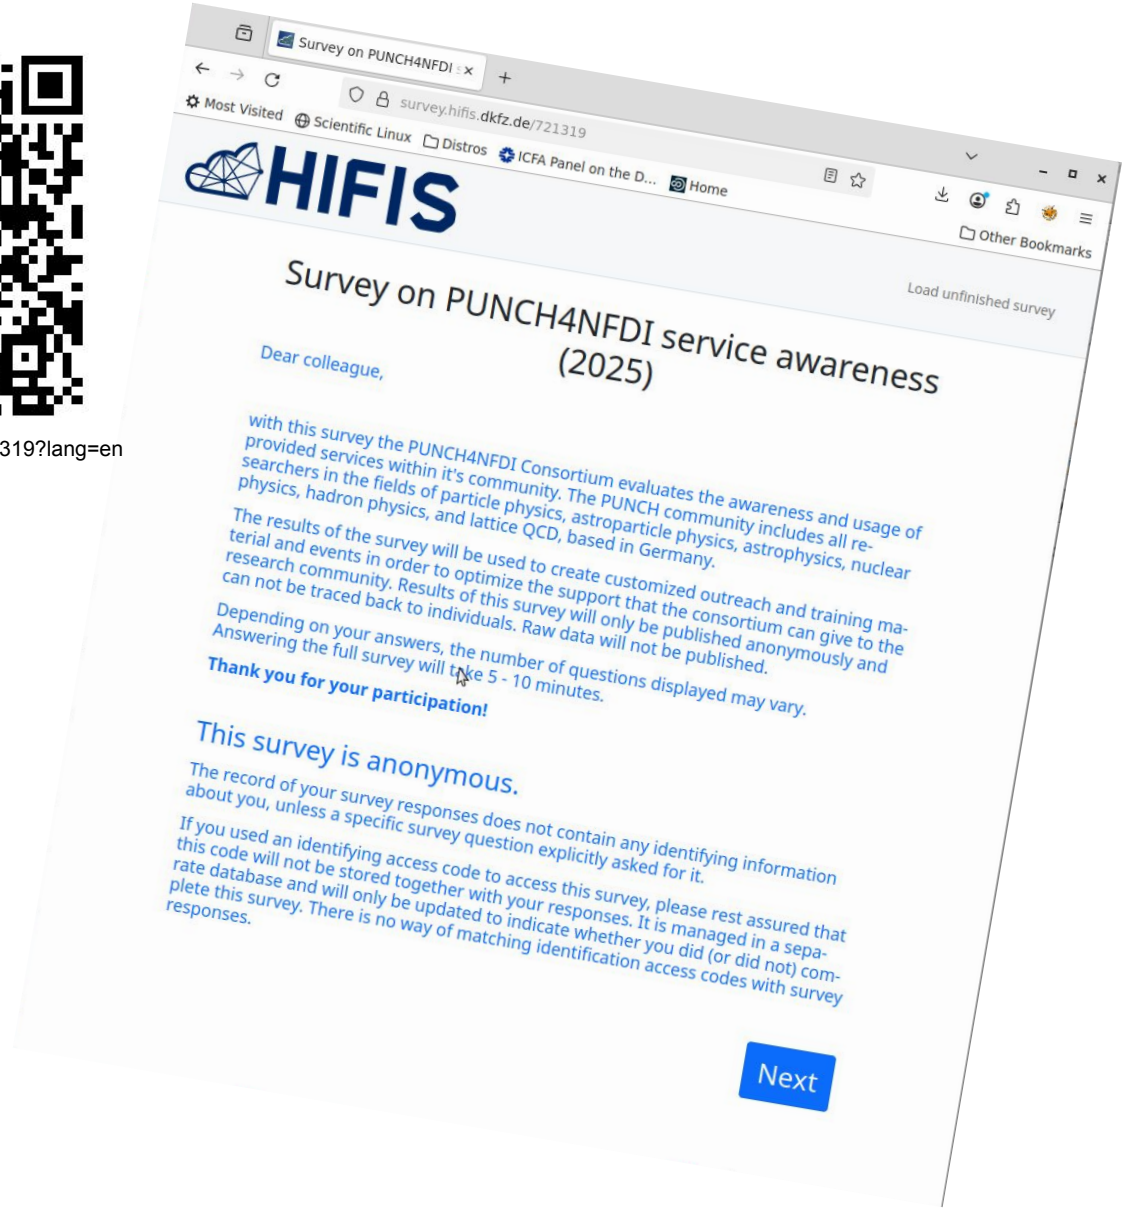

Metadata Catalog for ILDG & Astro

Status Quo

- Support for arbitrary XML schemas, tested against ILDG and IVOA schemas
- MDC+FC use token-based authentication and access control (Indigo IAM)
- Postgres-DB for good performance and stability (formerly eXist DB was used)
- Quick search functionality via JSON queries (XSLT-translation from XML to JSON)

Recent and ongoing developments

- Content-based sorting / pagination for the Astro Use-Case
- API finalization for current ILDG and Astro stakeholders
- In talk with further interested projects (e.g. Alps)
- Metadata preparation for GLOW pulsar data release

Outlook

- Support for PUNCH-AAI (when MDC requirements are fully supported)
- Generic GUI (for markup and search)
- Broader support for other communities through PUNCH 2.0

EXPLORE Service @ GAU Göttingen
Hosting, System Architecture, Status

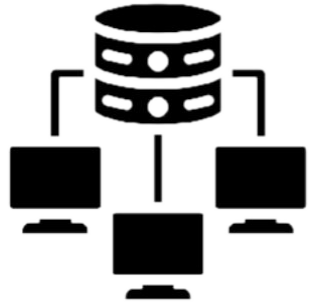

- **Aim:** Provide computing resources for **CERN Open Data analysis** to users without CERN experiment, or university affiliation. Registration is required.
- **Hosted & Operated by GoeGrid Cluster, Göttingen - II. Physikalisches Institut – GAU Göttingen**
- **Deployment:** Runs within the **WLCG GoeGrid Göttingen cluster** (Tier-3 part).
- Technical implementation via **COBalD/TARDIS** like Compute4PUNCH. Resources can technically become a C4P resource
- **Software & Environment:**
  - **CVMFS & Apptainer:** Provide users with containerized environments, ensuring flexible, reproducible, and scalable analyses
  - Users receive pre-configured operating systems and software tailored for CERN Open Data analysis
- EXPOLRE was topic of a PUNCH lunch seminar in September 2025
- **Service will be presented and used** in a dedicated **hands-on session** at the upcoming ATLAS open Data Workshop (24-27 Nov) at CERN, participants will run two workflows form **ATLAS Open Data Release 2025** on EXPLORE: https://indico.cern.ch/event/1564767/
- PUNCH4NFDI members are welcome to register and try it out

Our Well-Known Use Cases

MeerKLASS

Multi-petabyte MeerKAT survey using new OTF imaging software to process large sky areas on LRZ HPCs integrated into C4P

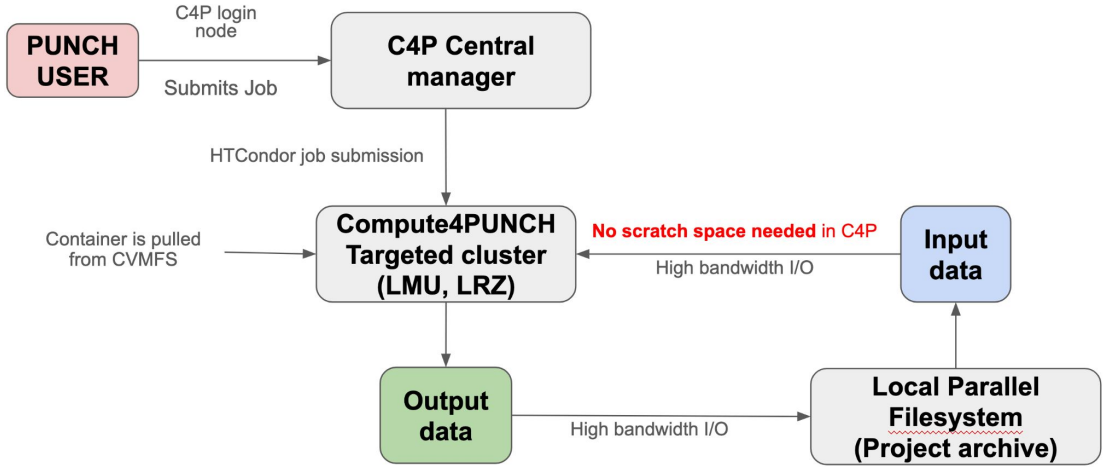

Cosmological Simulation Post Processing

Post-processing cosmological simulations leveraging PUNCH Federated Infrastructure for compute and storage

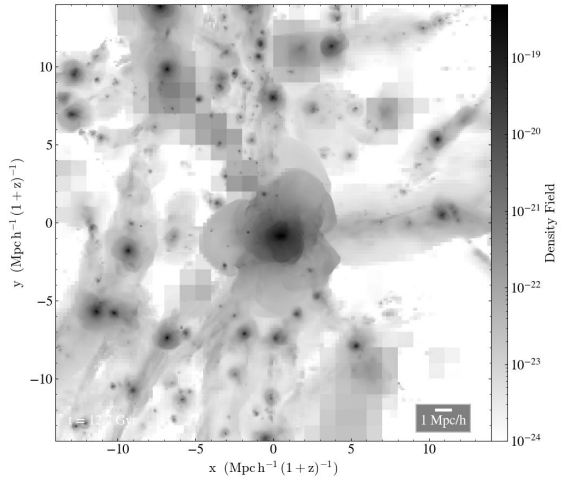

CERN Open Data

First-ever event-level combined ATLAS + CMS Open Data analysis utilizing PUNCH Federated Infrastructure for compute and storage based on REANA workflows

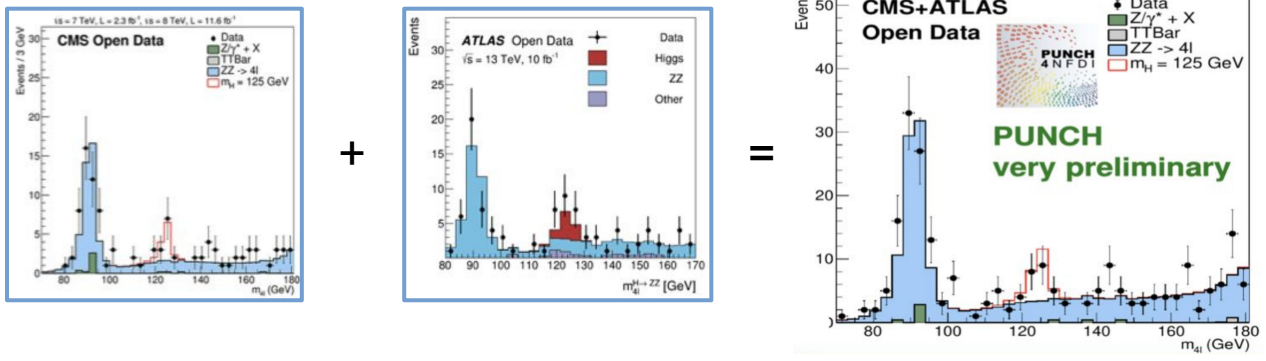

Use case:
Processing the LOFAR Two-Metre Sky Survey (LoTSS)

- Sky survey covers northern hemisphere @ 144 MHz
- **LOFAR** radio interferometric data is large (~ 1 TB/hr)
- Complex interplay of data calibration and sky brightness reconstruction („imaging“)
  → Mixture of high-throughput and high-performance computing

Special requirements:
- Data are stored on tapes (long-term archive)
- Very large data volumes → many copies unfeasible
- Current software design requires powerful (single) nodes and sufficient scratch space (~5 … 20 TB)

Realisation with S4P and C4P @ KIT:
- User only needs to provide URLs of staged archival data
- prototype development of semi-automated service to manage job inputs/outputs and submission
- 3rd party transfer between archive and storage
- Limit job submission to local C4P resources @ KIT
- Output is written back to local S4P resources @ KIT

LOFAR LTA

dCache instance

JÜLICH Forschungszentrum

davix-cp

GRID-certificate + BEARER_TOKEN

KIT dCache

KIT Karlsruher Institut für Technologie

KIT compute resources

submit

access token renewal

c4p-login.gridka.de

Storage4PUNCH

Thüringer Landessternwarte Tautenburg

Radio Sky Image from LOFAR

LOFAR ERIC

Extracting our homeland galaxy properties

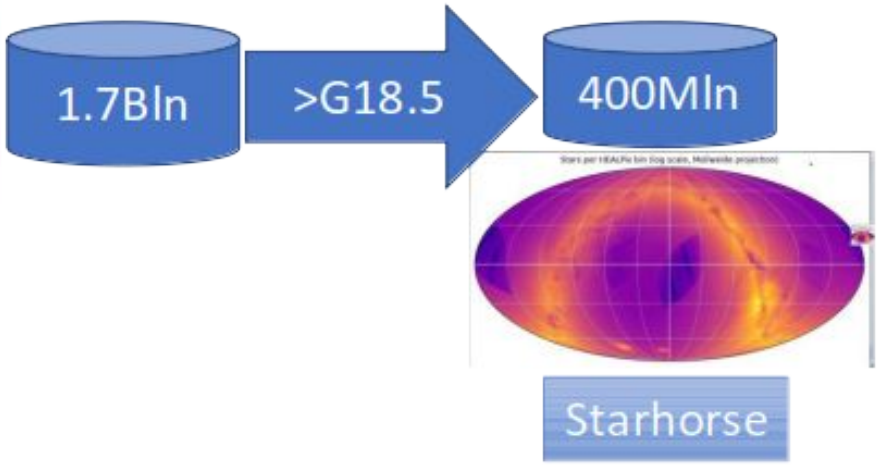

The Gaia mission has provided an invaluable wealth of astrometric data for almost two billions of stars in our Galaxy.

The synergy with the Bayesian isochrone-fitting code **StarHorse**, allowed to derive additional stellar properties, like distance and dust extinction (Anders+2021, Queiroz+2023).

Then a regression technique, **xgboost**, trained with spectroscopically derived data from a variety of stellar surveys, was used to determine additional stellar parameters: **extinction, effective temperature, metallicity, surface gravity, and mass** (Khalatyan+2024).

Many more details in a [presentation from Arman]

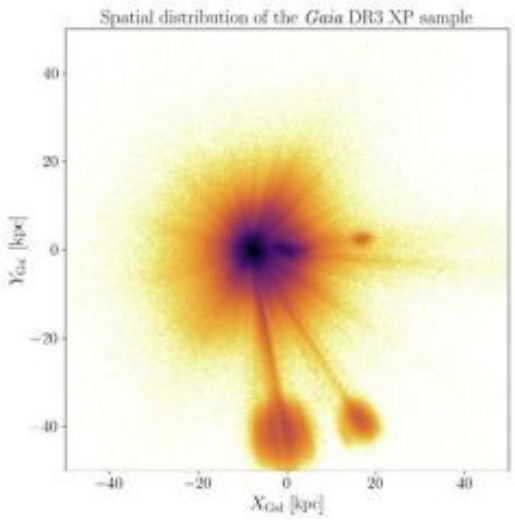

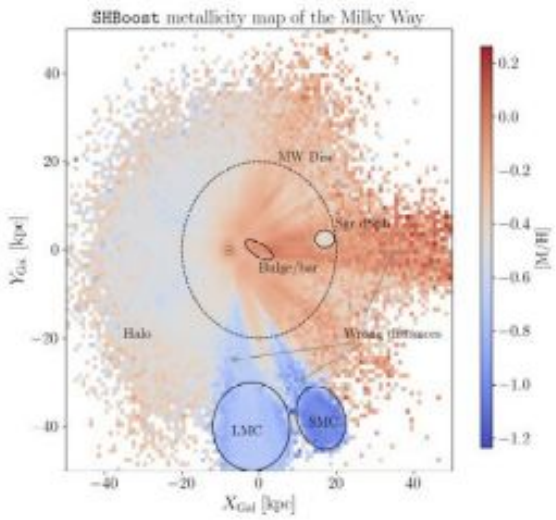

Extracting our homeland galaxy properties on C4P

Demanding task

- 12288x2 files to analyze
  - Sorted in Parquet or FITS format
  - Host on S3 storage at AIP
  - Total volume ~50GB
  - File size: 1MB-100MB
- Problem can (mostly) parallelized
  - One job per file
  - Duration 2min to 50h

Scale testing for C4P and REANA

- A number of issues identified and (mostly) fixed
- Up to 2,500 cores used simultaneously at C4P sites GridKa and GoeGrid
- About 1 Million Core Hours consumed in Oct. 2025

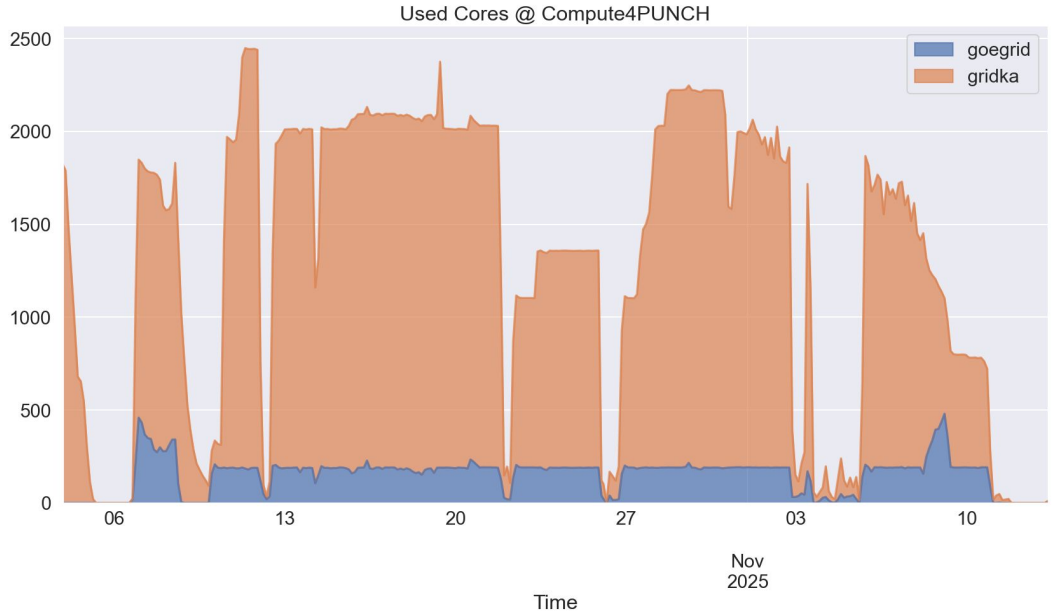

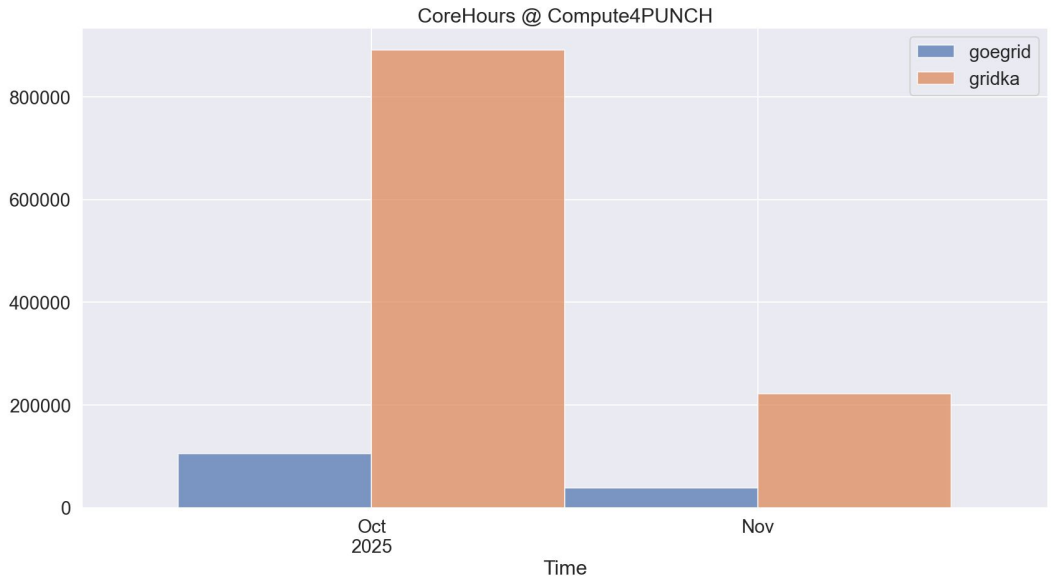

Summary

- With major contribution from TA2 PUNCH4NFDI set up componentes for federated computing
  - Compute4PUNCH
  - Storage4PUNCH
  - Container delivery via CVMFS
  - Rucio demonstrator
  - Flexible file catalog, in production use for ILDG, also serving as file catalog
- A number of scientific applications was ported to the federated PUNCH infrastructure in a joined approach between several TAs
  - Scale testing and harding of sub-components
  - Many aspects for improvements found
  - Valuable input to the AAI development could be given
- Especially user related (Frontend, UI) topics get addressed in PUNCH2.0

Backup

CI/CD for Container Registry

Containers are maintained on docker container registry at AIP (gitlab-p4n.aip.de) using CI/CD pipeline

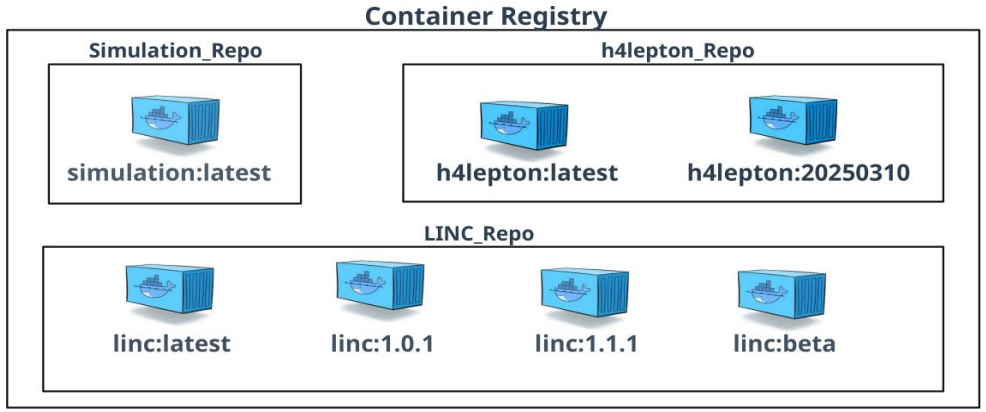

CI/CD Framework v1.0

- Different docker files managed in directories on single GitLab project (compute4punch/container-stacks)
- CI/CD pipeline triggers everyday or on any change in `container-stacks` repository
- Changes to one project triggers rebuild of all containers, can overload CI/CD pipeline
- Everyday, a new container image generated for each docker file, with `:DATE` and `:latest` as tag
- No control on versioning the docker containers

CI/CD Framework v2.0

- Separate repository for each docker file within the GitLab project group (compute4punch/container-stacks-v2)
- A master CI/CD pipeline, can linked to individual projects repository
- CI/CD pipeline triggers only in the repository where any change occur; gives more control on versioning of docker containers
- Tag container image with `:DATE` and `:latest`, multiple versions
- Update/modification to the README file will not trigger the CI/CD
- Common CI/CD pipeline template available in Master CI/CD documentation

RUCIO4PUNCH - The Testbed

- Tool to manage scientific data across a distributed infrastructure
- Very popular in High Energy Physics (HEP)
    - Originally developed in ATLAS
    - Now used by other major HEP experiments like CMS, Belle II, but also smaller ones and even non-HEP experiments
- HIFIS is providing a RUCIO (test-) instance
    - Running at DESY
    - PUNCH has been enabled via PUNCH-AAI (selected accounts so far)
    - 3 Storage4PUNCH instances are configured (Bonn, DESY, KIT) as RUCIO Storage Elements (RSE)
    - Underlying FTS (File Transfer Service) provided by CERN
- RUCIO testbed is actually a deliverable in TA6

MeerKLASS Use-case

- 4 PB of raw data recorded at the MeerKAT array (~2000 hours observation time)
- 10,000 deg^2 in UHF band, 300 deg^2 in L band

Science goals:

  - Neutral hydrogen intensity mapping for baryon acoustic oscillations
  - Continuum imaging of active galactic nuclei, radio galaxies

- New software enables On The Fly (OTF) interferometric imaging
- Evolving requirements to image 1 deg^2 sky patch
  - 128 cores, 500 GB RAM, Runtime ~7-8 hours
  - Up to 4 TB intermediate data produced!
- Requires integration of dedicated LRZ HPC cluster hosting data on local parallel FS into C4P
  → Prototype integration of CM4 functional

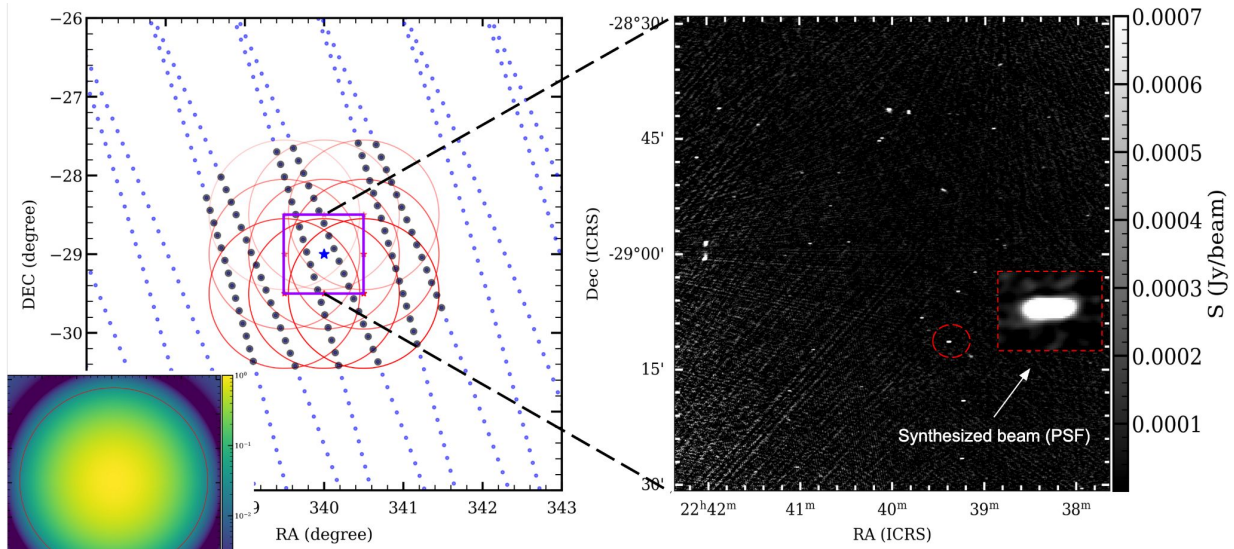

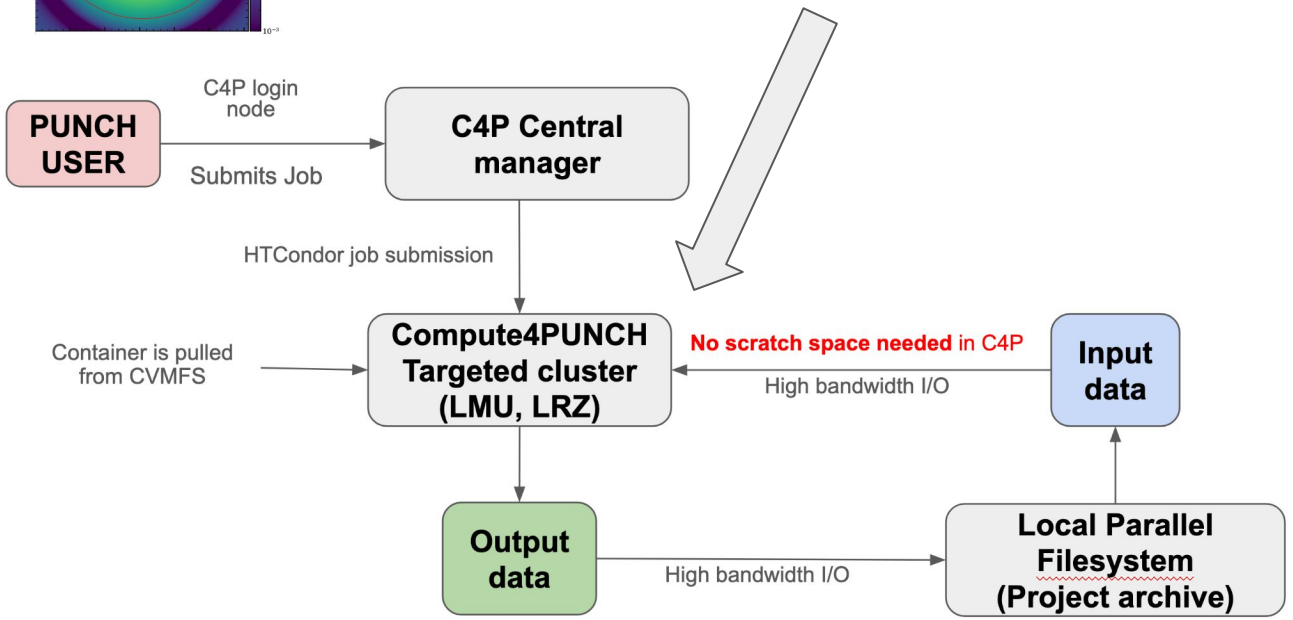

Post-processing of Cosmological Simulations Use-case

Demands large storage: Typical cosmological simulation occupies ~500 GB to 10 TB of disk size

Post-processing of such a large volumetric data requires intensive computing resources

It involves: Projection plot, Volume rendering, Time-series plot, Finding dark matter halos/galaxy clusters

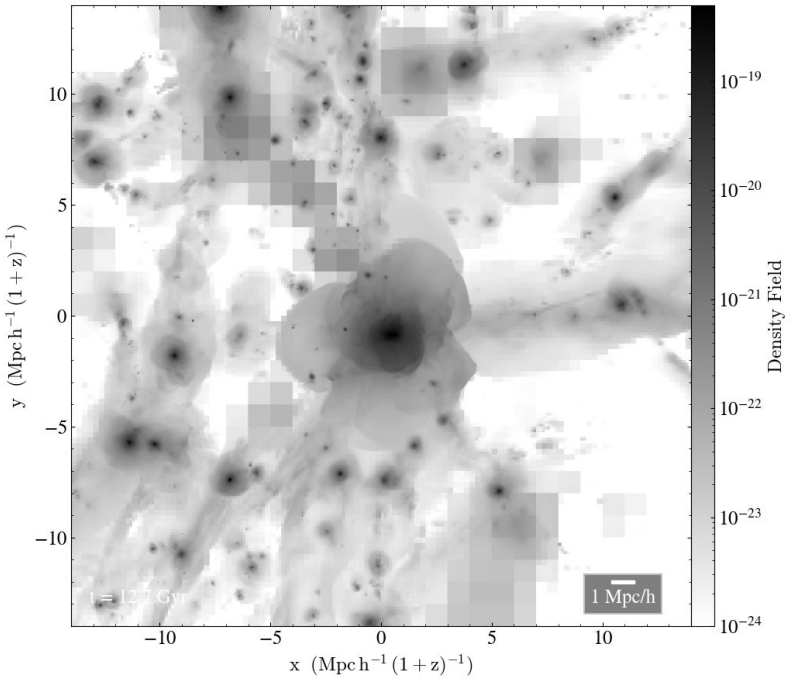

The workflow for post-processing of the private cosmological simulations handled with

- ReAna4PUNCH workflow manager
- Simulated data stored on Storage4PUNCH
- Post-processing on Compute4PUNCH backend
- Results uploaded on Storage4PUNCH

Ongoing:

- Post-processing of publicly available cosmological simulation snapshots, i.e., Illustris
- Able to query data with API-requests; authentication stored as ReAna-secrets

Working example available at:
https://gitlab-p4n.aip.de/punch/tutorials/Simulation_PostProcessing_usage_example

CERN Open Data Use-case

Original CMS legacy research data
(2 PB via CERN Open Data portal)
(2010 data 100%, 2011/12 data 70%)

CMS legacy software from public github via CERN Open Data portal)

Produce histograms
(many CPU months)

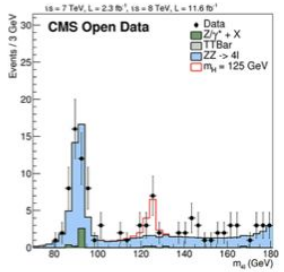

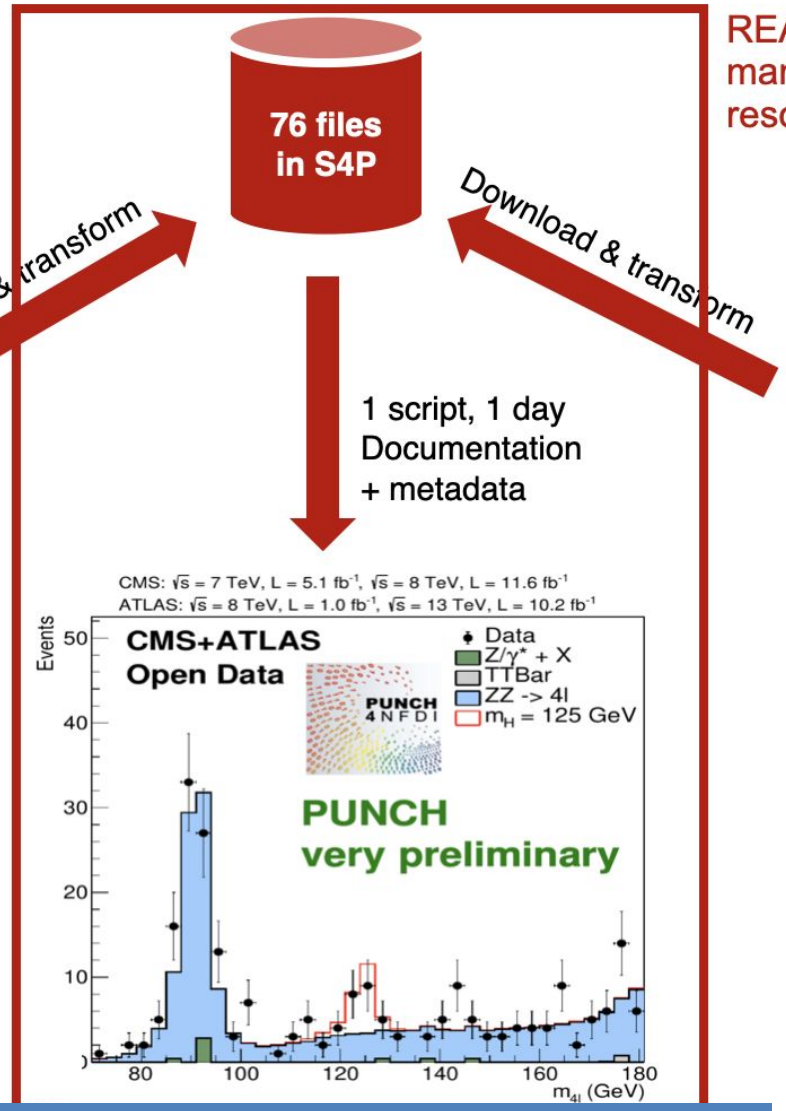

REANA for workflow management and resource organisation

Non-public ATLAS legacy data

Simplified educational data sets 2012 / 2016 via CERN Open Data portal or ATLAS Open Data portal [8]

VM with dedicated software package or Jupyter notebook

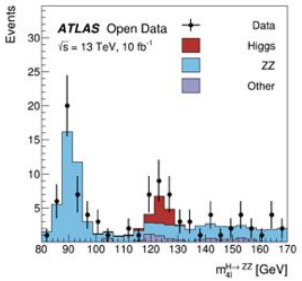

First-ever event-level combined ATLAS + CMS Open Data analysis