D-MeerKAT-II: Summary of on-going activities and access to GLOW HPC cluster

Radio 2021 — Garching

SUPPORTED BY

Bundesministerium
für Bildung
und Forschung

Mayumi Sato, Jörn Künsemöller & Dominik Schwarz (Bielefeld University)

WP1: Federated Science Data Centre and Big Data Challenges [Bochum, FZ Jülich, Bielefeld, TLS]

1a: Data Transfer, Storage and Computing
1b: Pipeline Development and Scalability for Future Facilities

GLOW HPC cluster

Jörn Künsemöller, Aritra Basu, Peter Kamphuis

- System installation done, ready for use (→ later slides)

Peter Kamphuis, Dominik Bomans, Ralf-Jürgen Dettmar

CARACal (Containerized Automated Radio Astronomy Calibration) pipeline on GLOW HPC

- Installation complete, as a virtual environment (Singularity-in-Singularity is being tested)
- Preliminary tests ongoing using MeerKAT open-time data
- Preparing scripts accessible to all users and autoload proper paths
- Python package version issues, solving them for the HPC
- Visualisation and interoperability → CARTA to be tested next

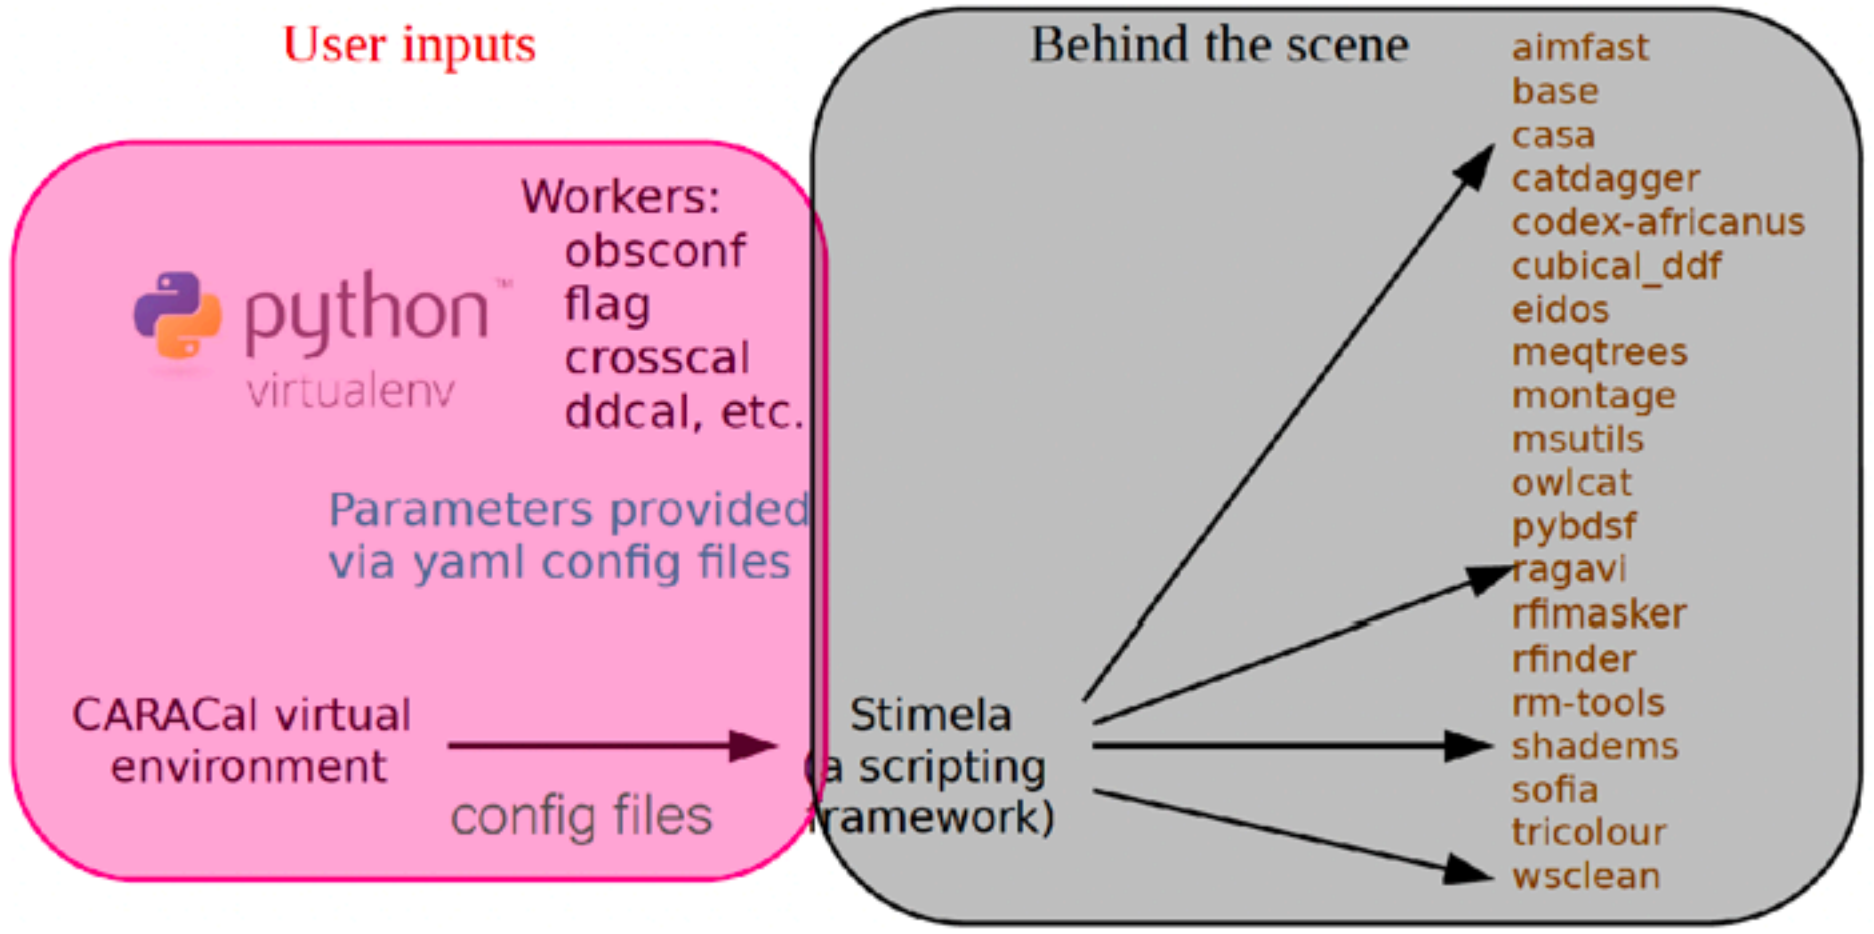

WP2: MeerKAT Imaging and Signal Processing [TU Münich, Hamburg, MPA]

2a: Ionospheric Effect

Feng Gao, Virginia Cutti, Francesco de Gasperin

2d: RM Synthesis Imaging

Major progress: implemented **full-Stokes** deconvolution in DDFacet in Högbom CLEAN mode

- Caveat: no specific PSF for each Stokes, using PSF_I (mainly computationally limited)
  - primary beam model and self-calibration expected to catch any mis-match of PSF between Stokes

In progress: testing full-polarisation primary beam/DDE correction

- on 3C286 MeerKAT raster scan data
- Average beam model taken into account in DDFacet
- DDE correction provided by package killMS (currently Stokes I only)

Next step: merge RM synthesis inside imaging

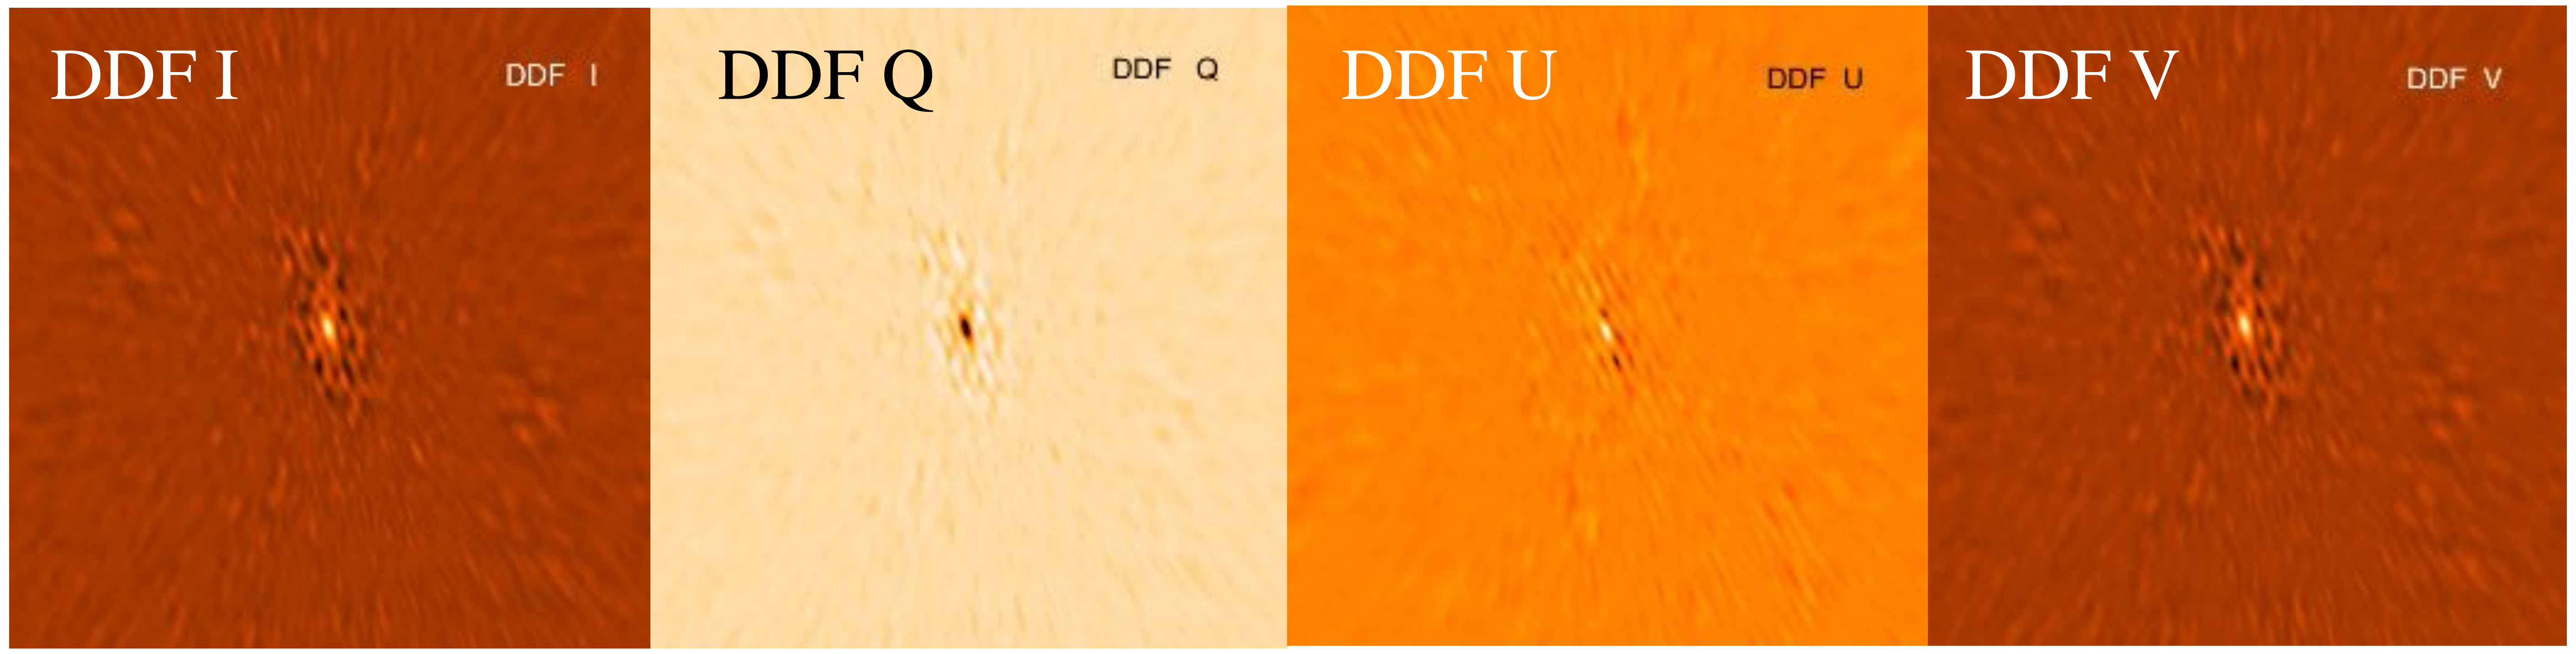

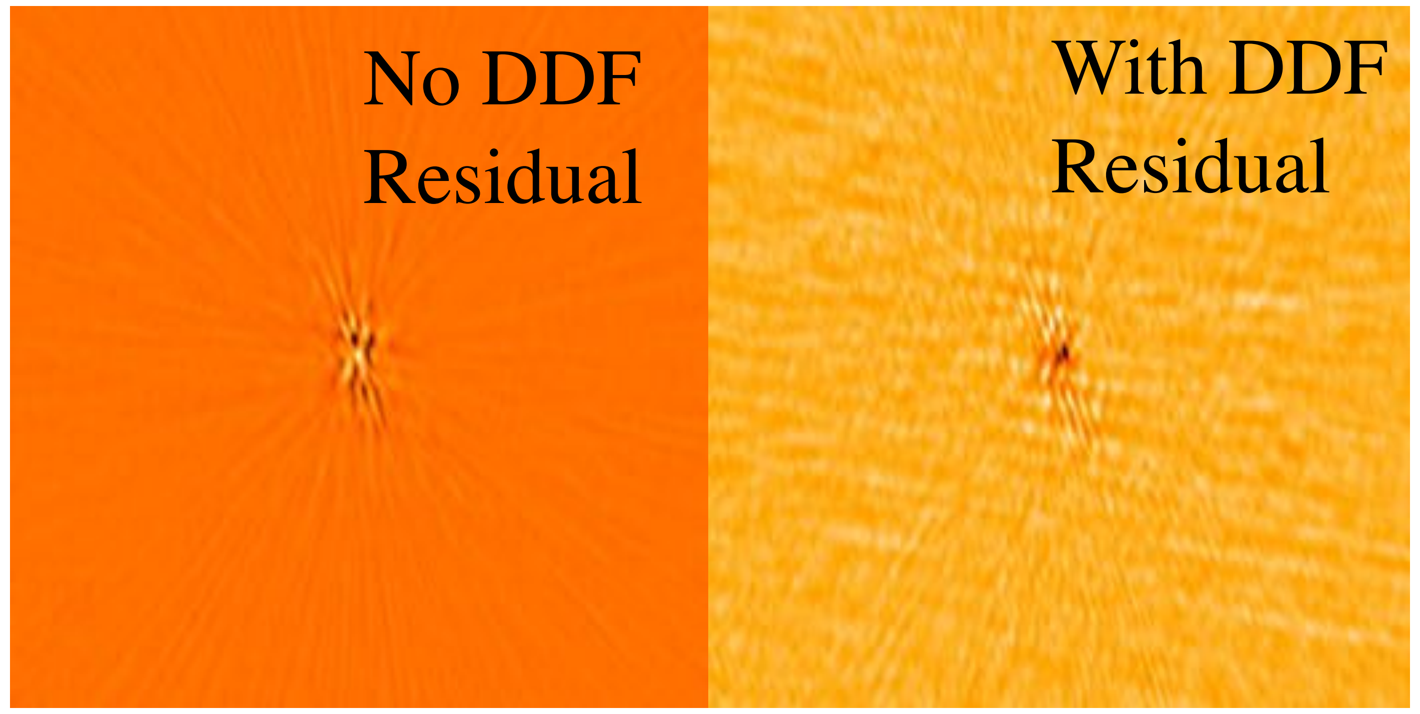

WP2: MeerKAT Imaging and Signal Processing [TU Münich, Hamburg, MPA]

2a: Ionospheric Effect

2d: RM Synthesis Imaging

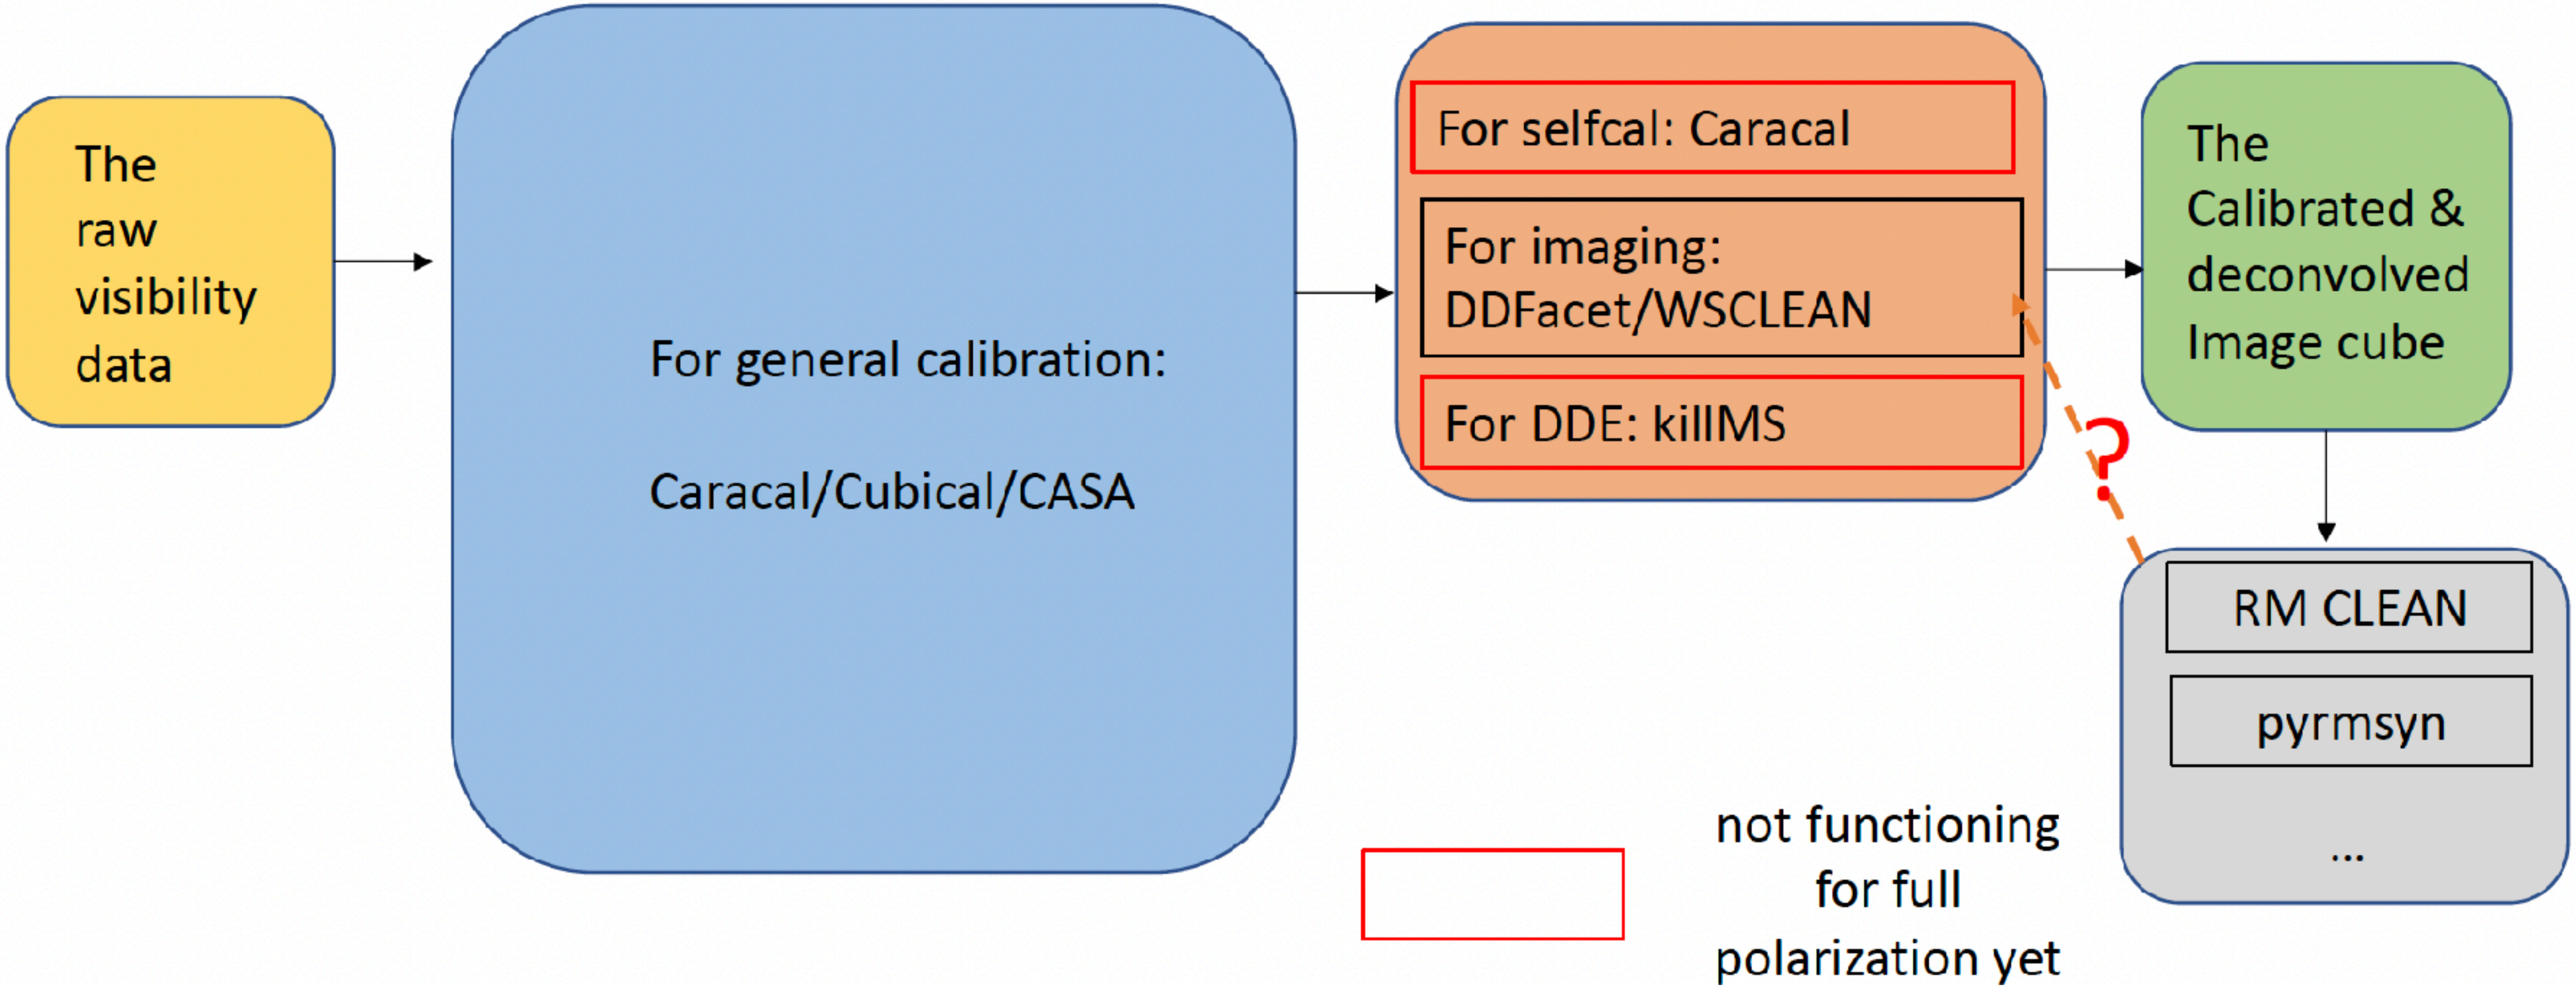

WP2: MeerKAT Imaging and Signal Processing [TU Münich, Hamburg, MPA]

2b: Direction Dependent Polarisation Calibration

2c: Multi-frequency Polarimetric Imaging

Current projects

Philipp Arras, Torsten Enßlin, Rüdiger Westermann

- Calibrate MeerKAT from scratch
- Direction-dependent calibration
- Bayesian mosaicing
- Bayesian multi-frequency models
- EHT imaging challenge
- Beam modelling
- General NIFTy and resolve development

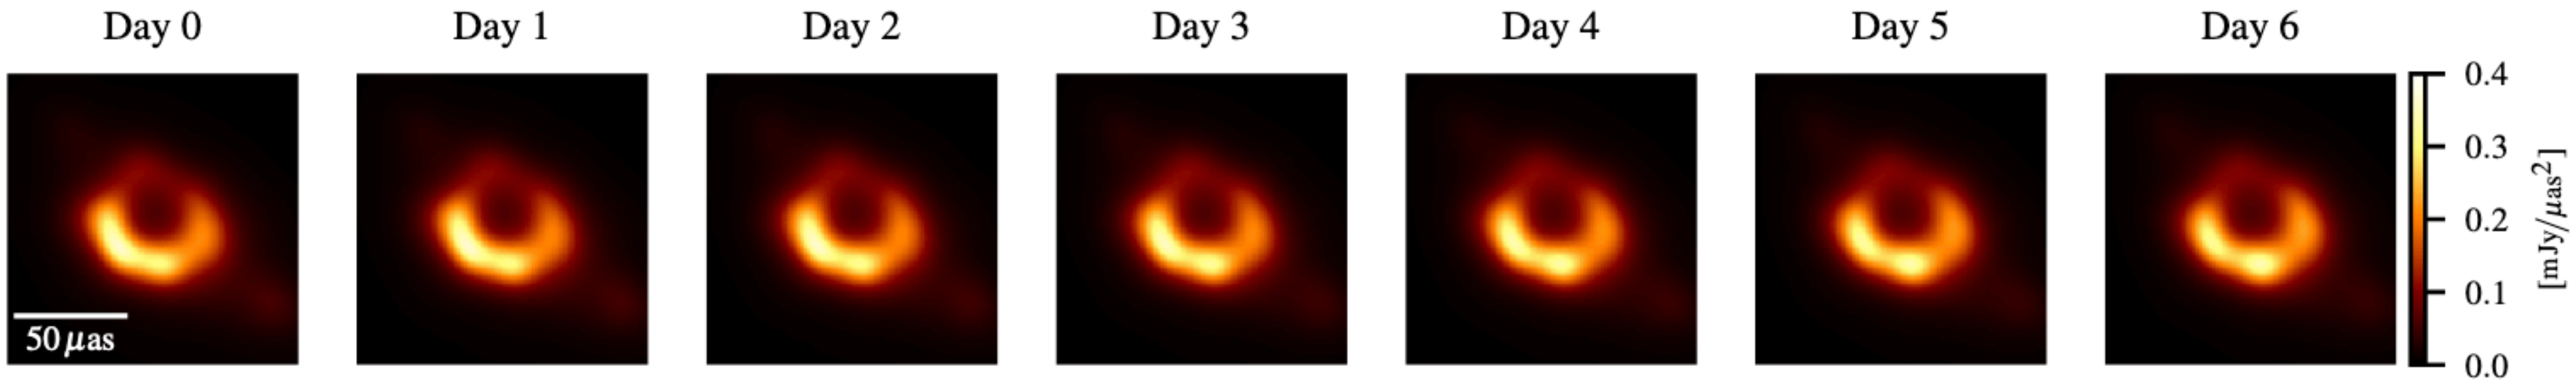

Event Horizon Telescope: M87* (Arras et al, arXiv:2002.05218)

WP2: MeerKAT Imaging and Signal Processing [TU Münich, Hamburg, MPA]

2b: Direction Dependent Polarisation Calibration

2c: Multi-frequency Polarimetric Imaging

Calibrate MeerKAT from scratch

- Take actual raw data and process it with Bayesian calibration algorithm

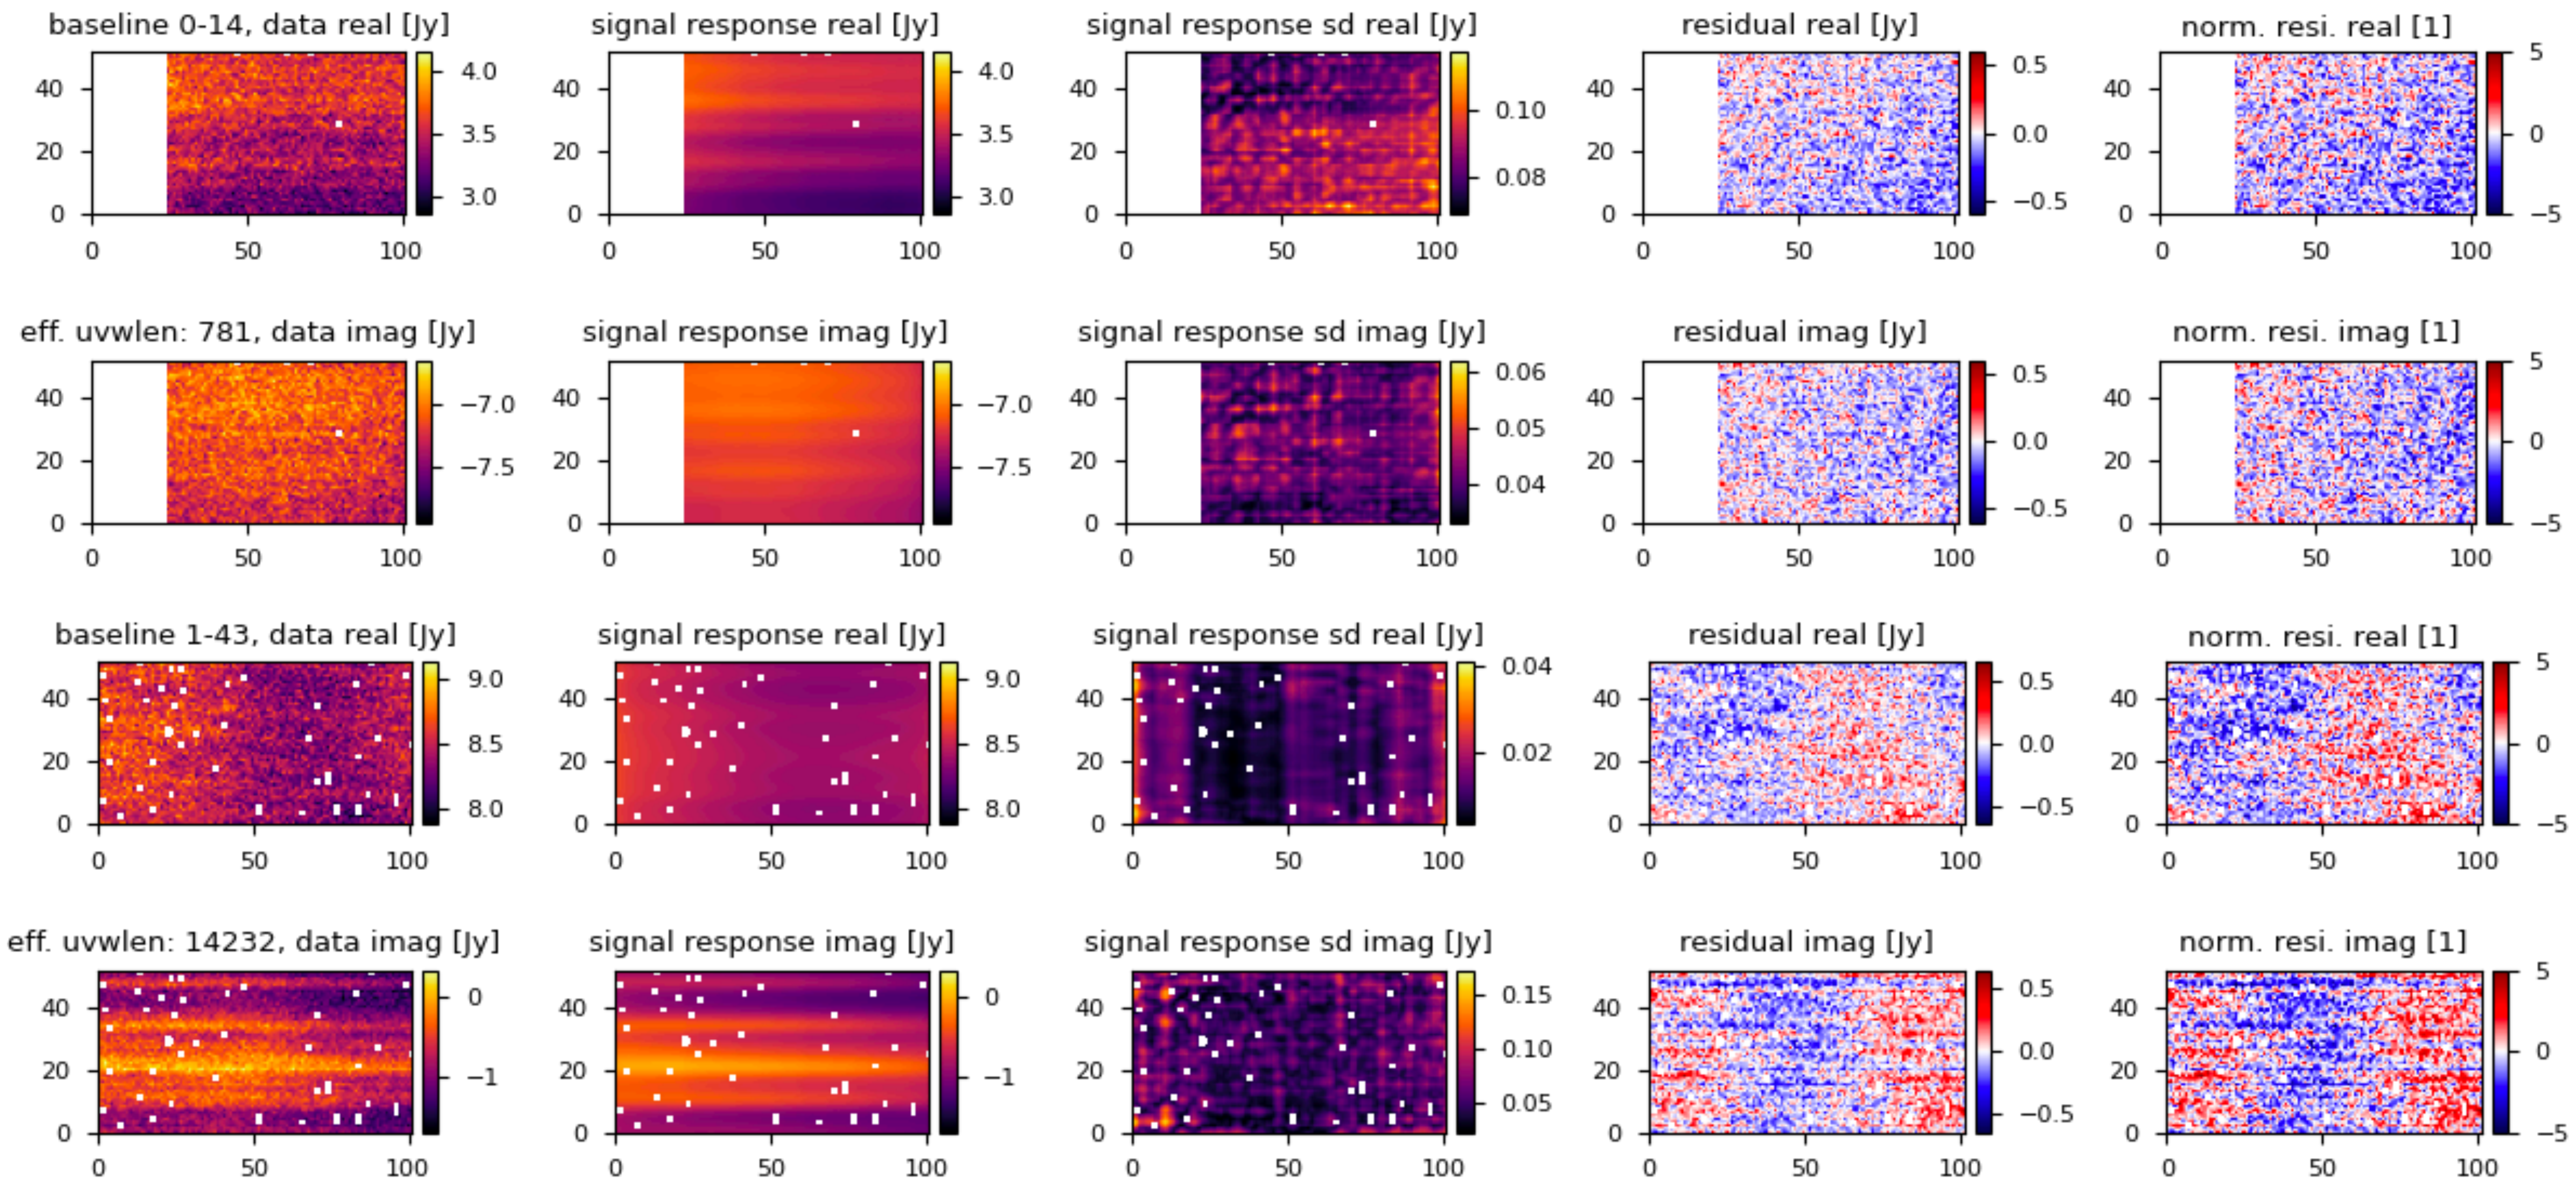

Reduced χ^2 vs. baseline length

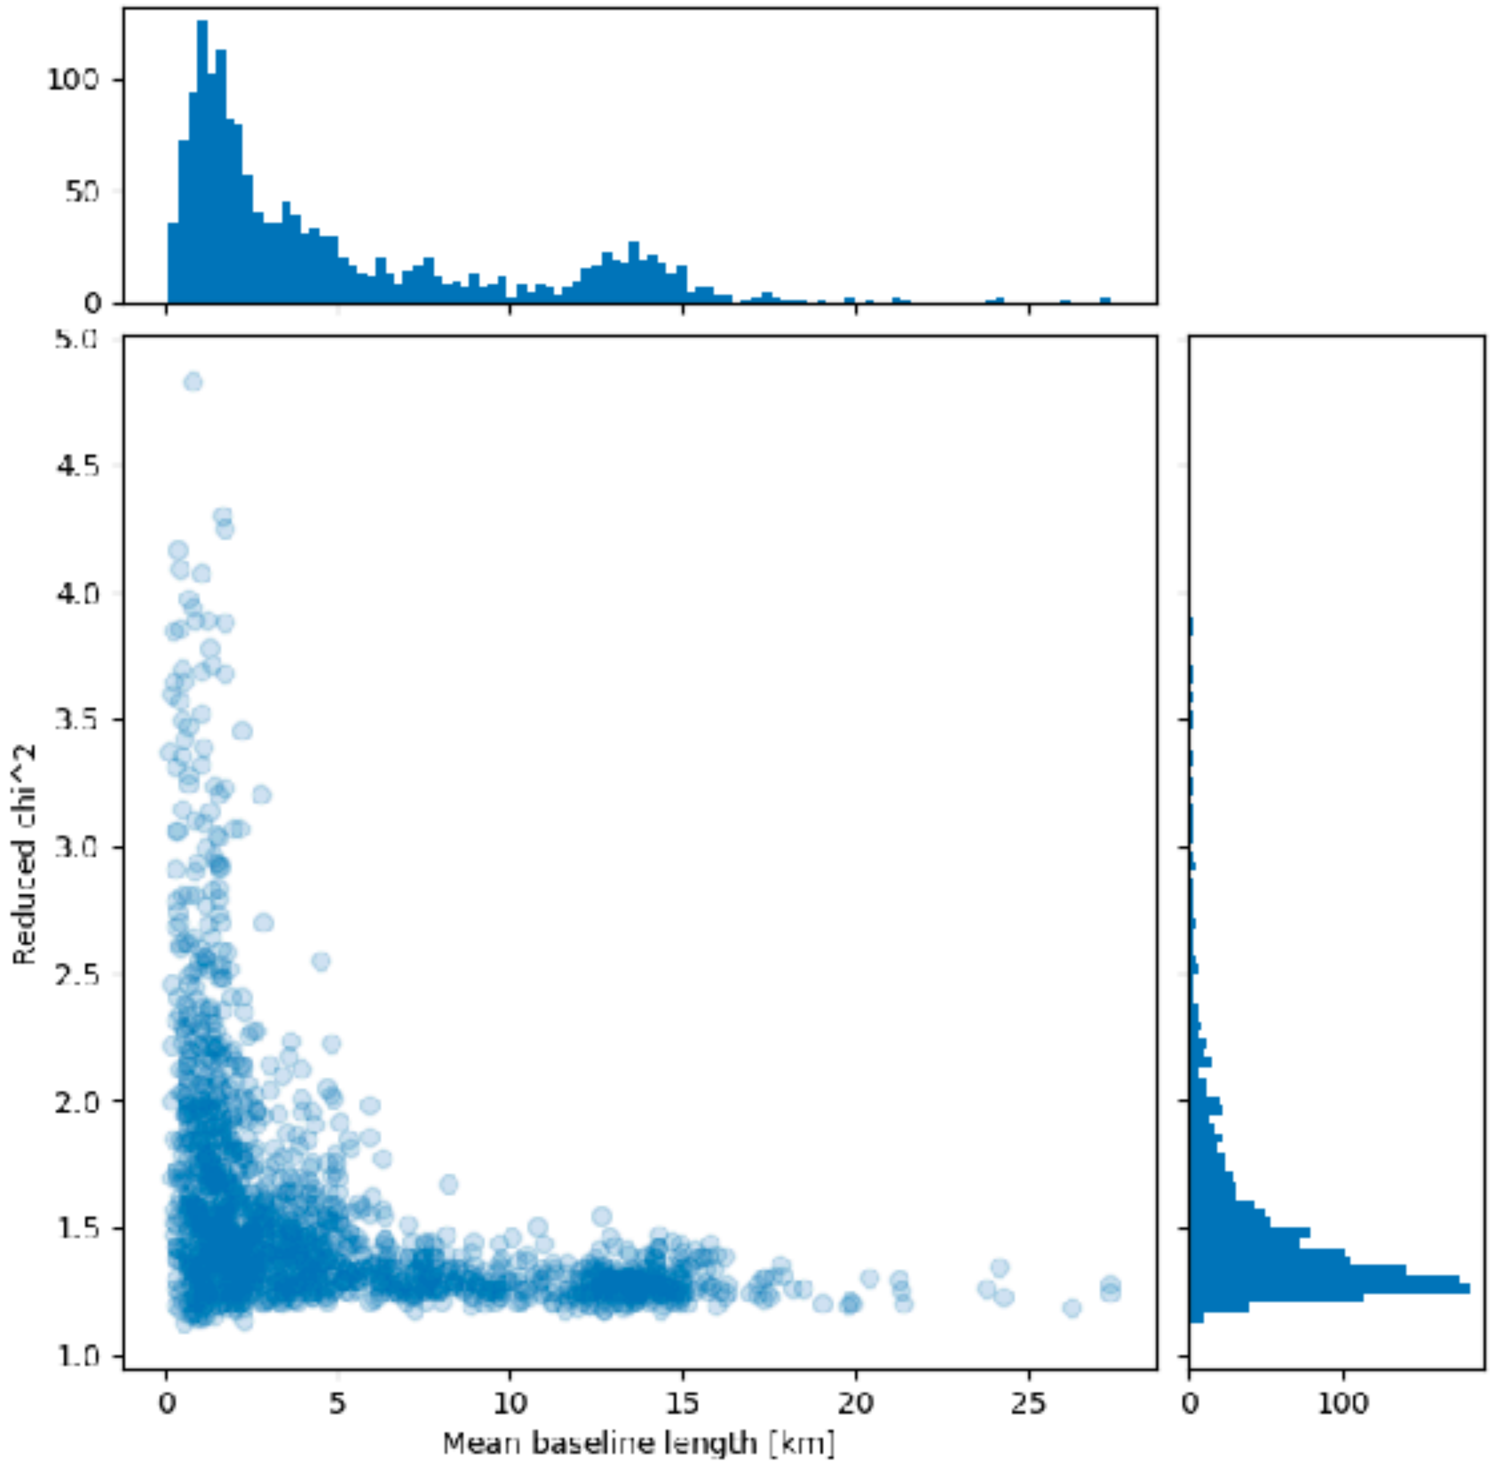

WP3: Real-time Screening and Transient Identification [Heidelberg, Bielefeld, MPIfR]

Triggering pipeline:

VOEvent (Virtual Observatory Event) within MeerTRAP

- Internal – to trigger transient buffer to get an image of the transient field for localisation
- External – dissemination of information about transients detected by MeerTRAP to the outside world to enable multi- frequency follow-up
- Triggering external VOEvents (of they are in the field of view)

RFI mitigation (→ see later slide)

Marina Berezina, Felix Jankowsky, Stefan Wagner, Lars Künkel, Joris Verbiest

WP3: Real-time Screening and Transient Identification [Heidelberg, Bielefeld, MPIfR]

MeerTRAP triggering pipeline

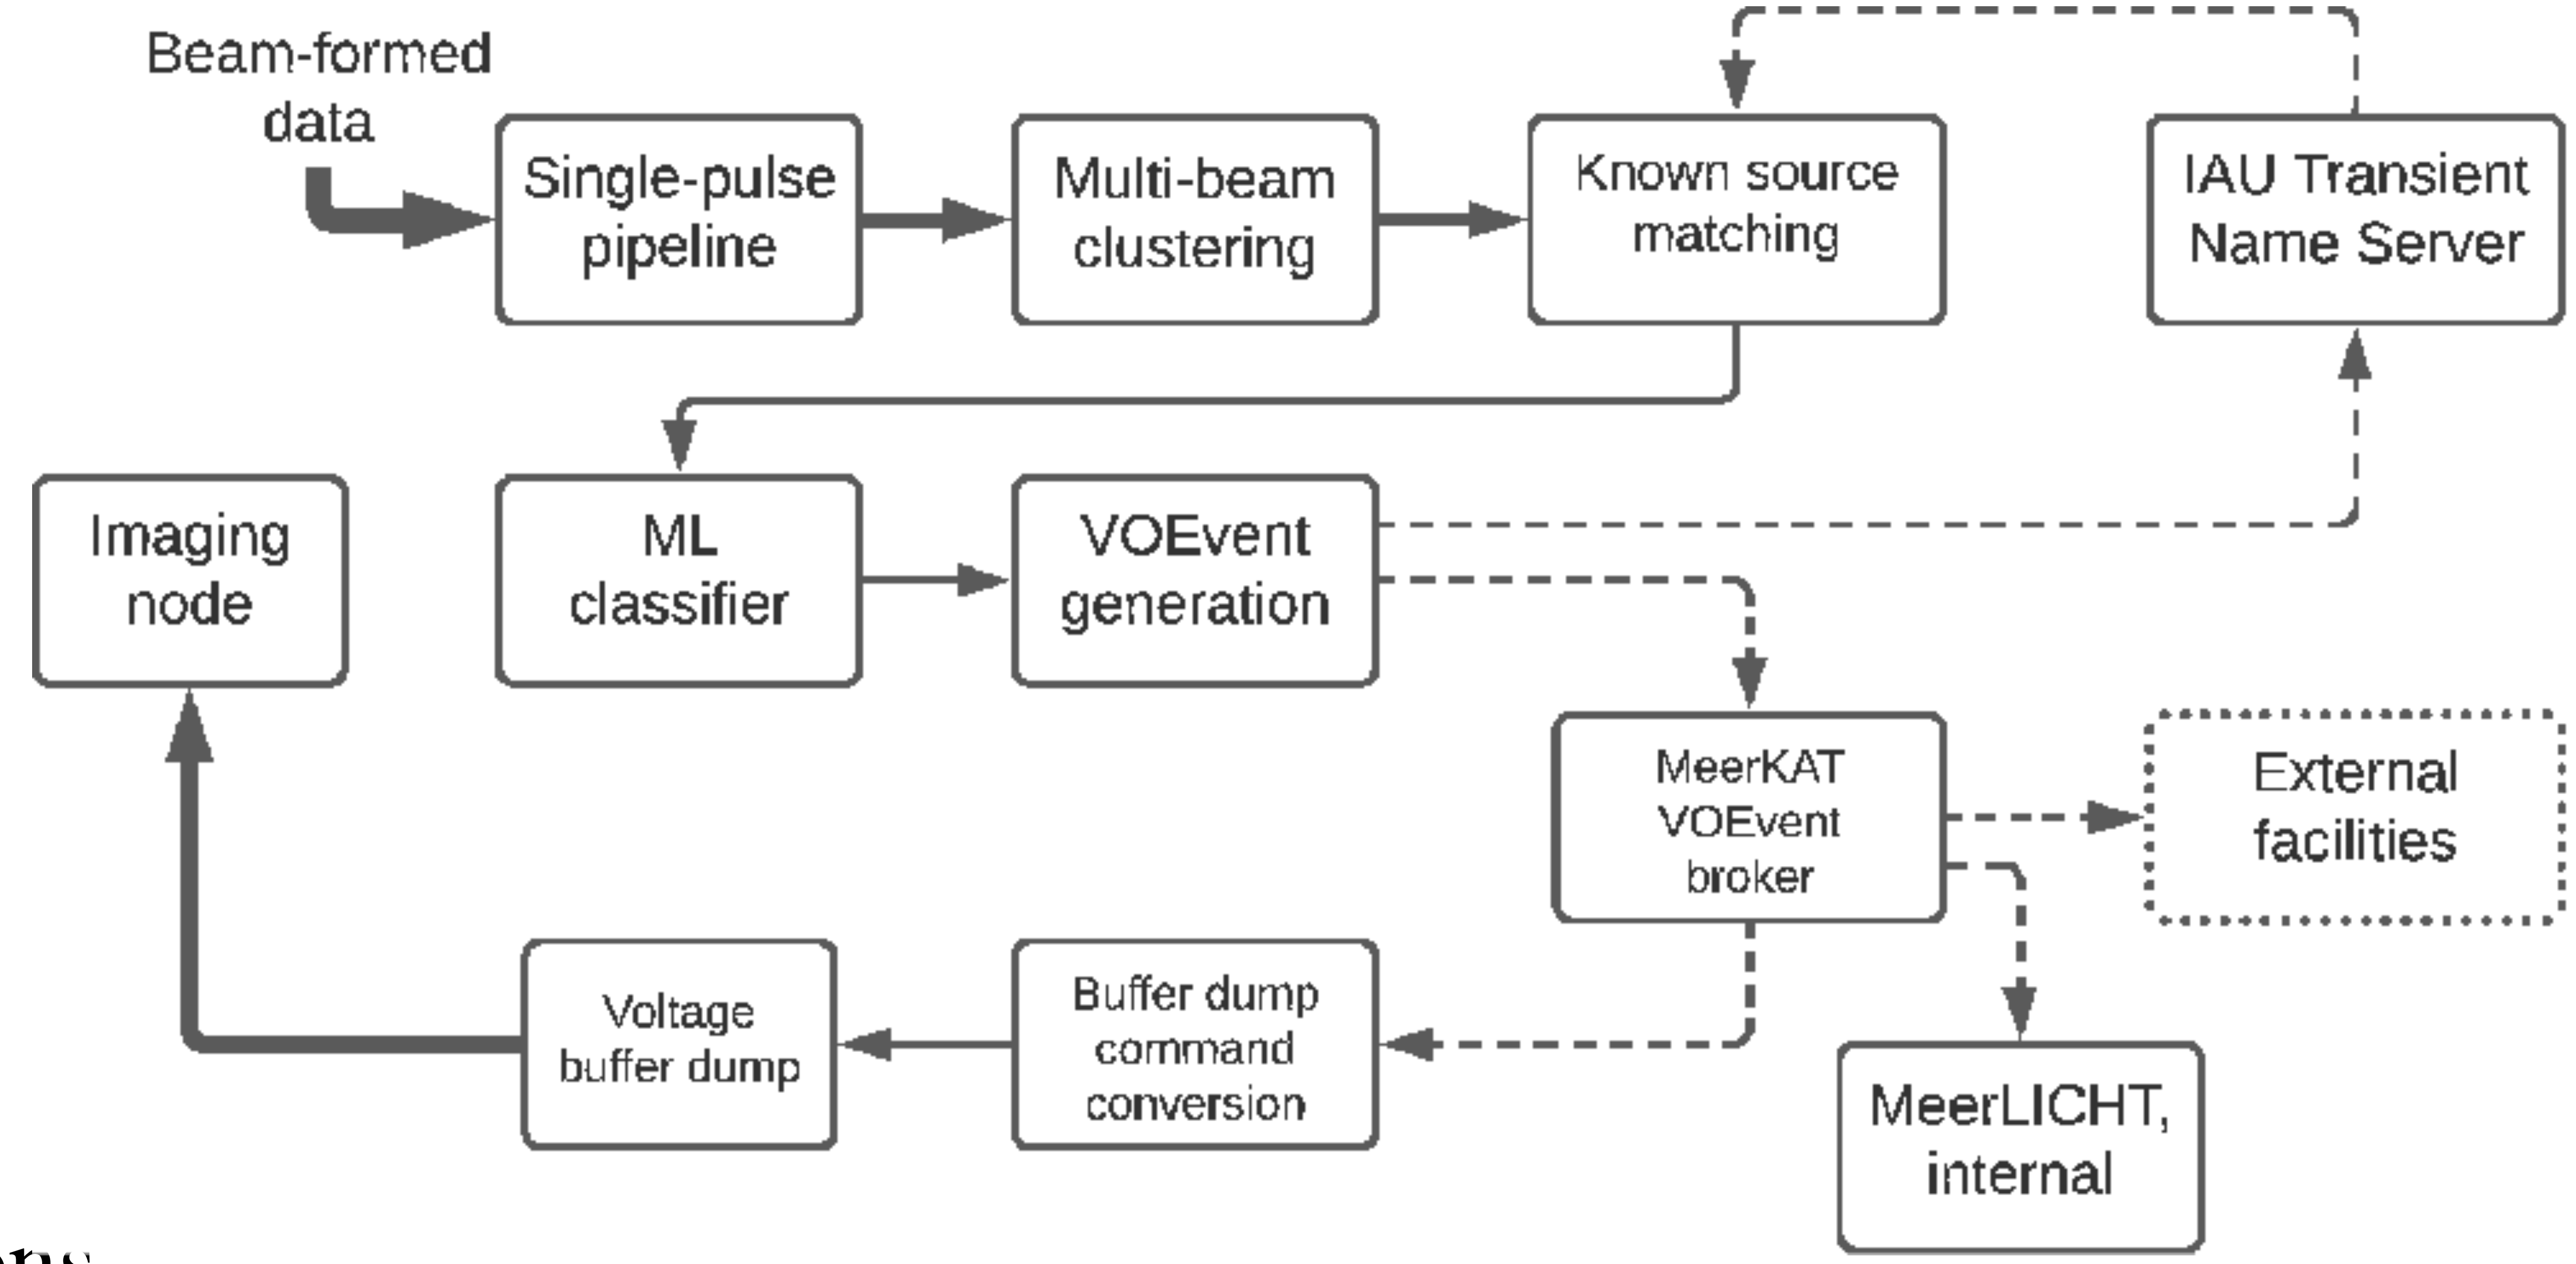

(Jankowski, Berezina et al. 2021)

- Commensal (~20000h - planned)

  N_{beams}=768

  L-band (856 MHz-bw)

  τ=76.56μs
- Real-time single-pulse
- FRBs (FRB-like signals), RRATs, Galactic transients
- Currently ~7.7 h/day of commensal observations
- A real-time question
  - Trigger internally (transient buffer with ~40 seconds of channelized complex voltage data) in real time
  - Up to 10 times per day with some events (not exactly reliable – with different importance)
  - False alarm rate of the ML classifier >1% which is still high for producing "publicly trusted" VOEvents
  - A bit hard to disseminate events and trigger external facilities

RFI mitigation: satellite prediction

- Predict a satellite pass during the observation
  - Two-line element files (TLE): satellite ephemerides (from spacetrack.org, celestrak.com, etc.)
  - Skyfield python library
- Mask frequency channels corresponding to satellite emission/mask time samples of the passage
- Tried on the L-band Galactic plane pulsar survey (MLGPS) data
- Create “birdie” lists with satellite periodicities

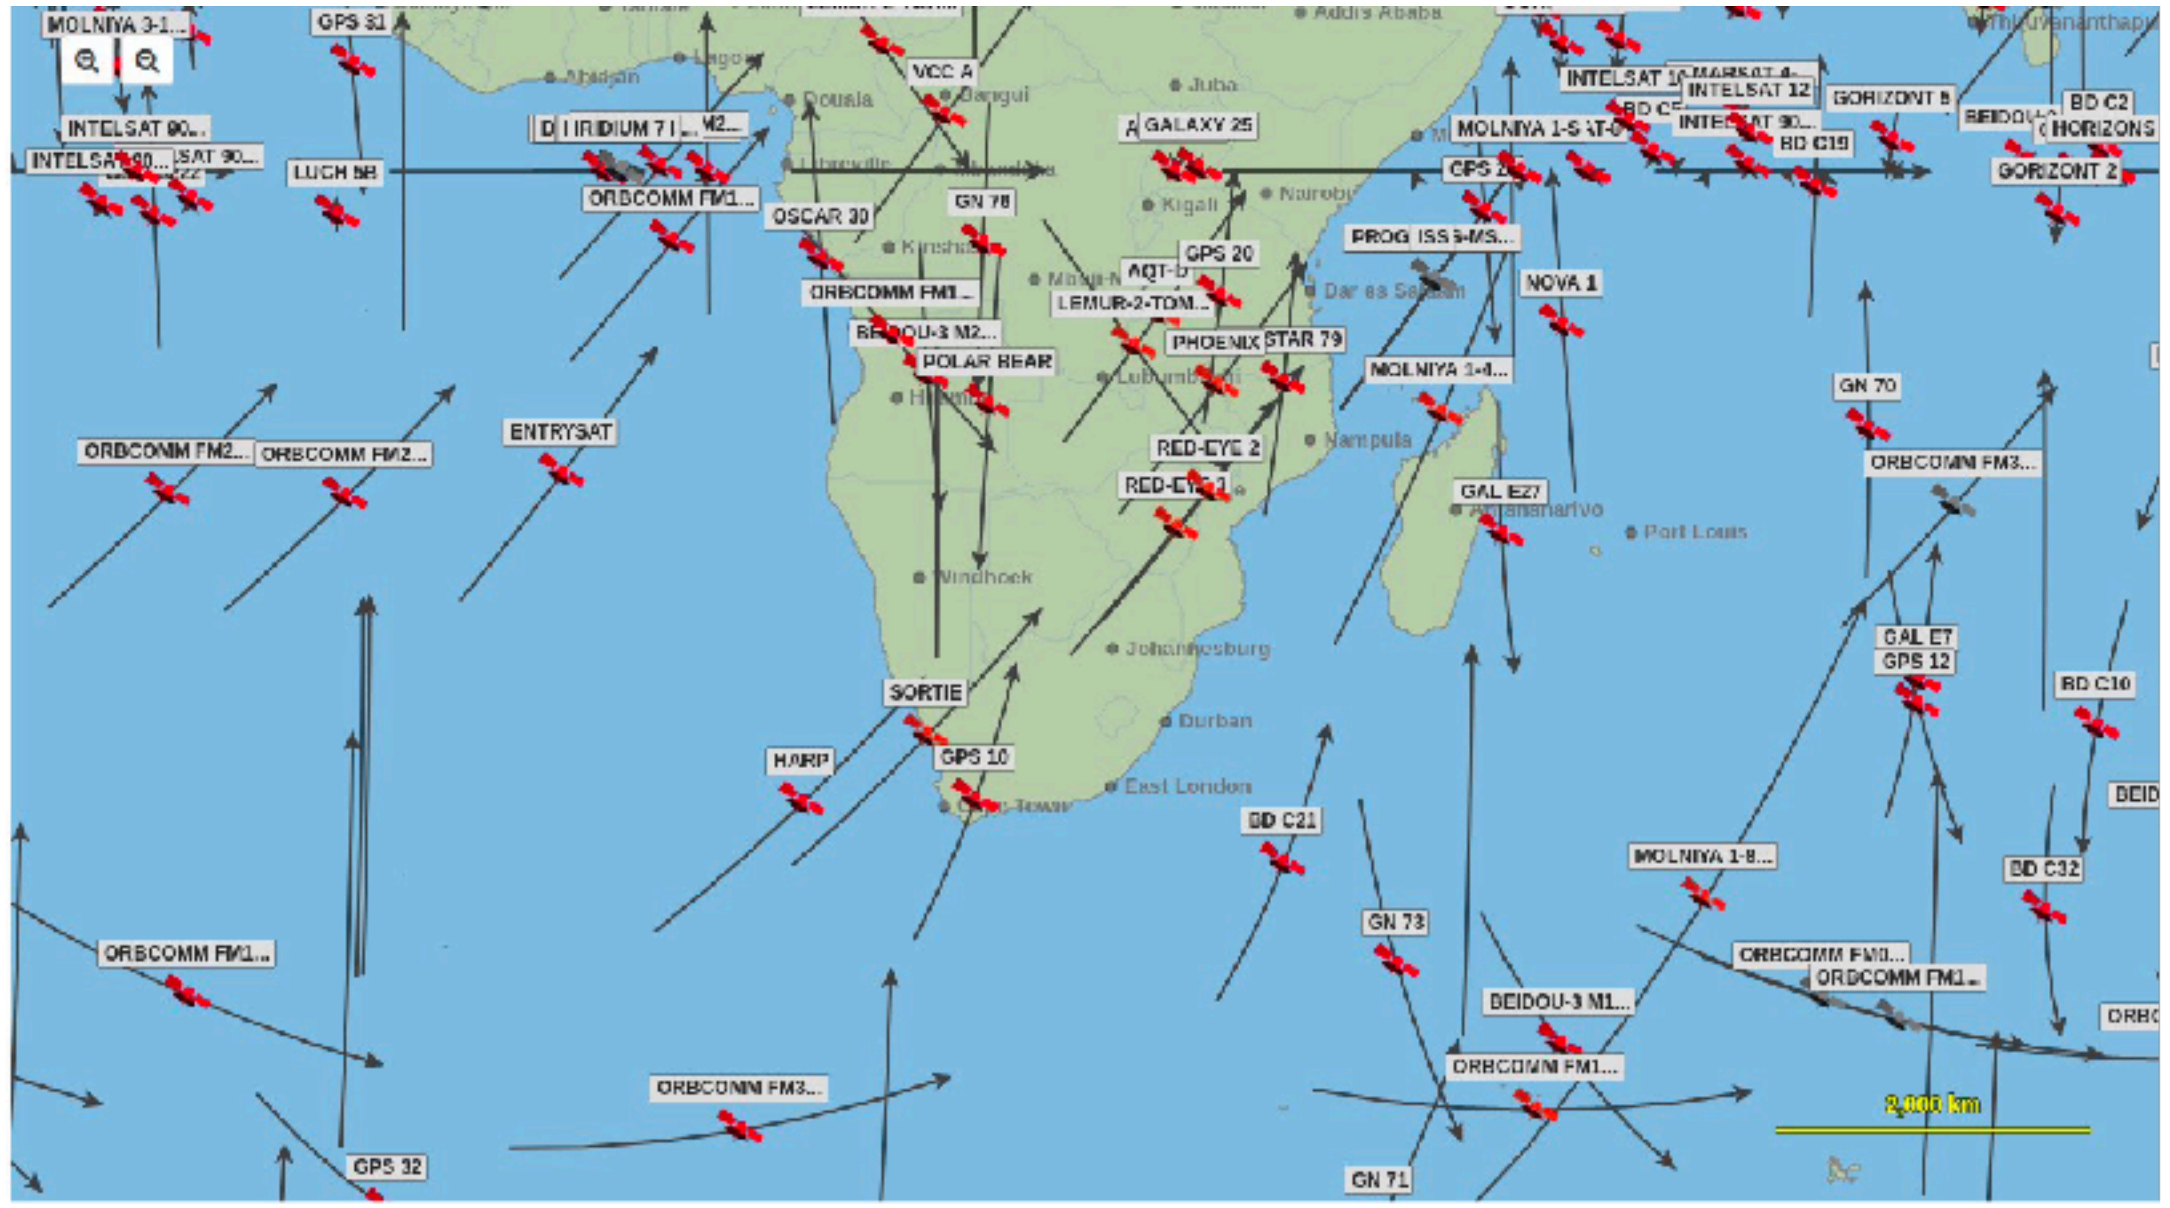

(Credit: in-the-sky.org)

WP4: On-the-fly Interferometry during Scanning Mode Observations with MeerKAT [LMU, TLS]

4a: Calibrated HI and Continuum Imaging for Scanning Mode Observations

- MIGHTEE survey (→ talk by Natasha Maddox)
- MeerKLASS (Santos et al. 2017 @arxiv) planned HI intensity mapping survey using autocorrelation (single dish) mode
  - L-band wide-field (~4000 deg^2) survey using a scanning observing strategy
  - Primarily an HI intensity mapping (IM) experiment for BAO detection and redshift space disortions
- Whitepaper on on-the-fly PB correction for MeerKLASS in preparation

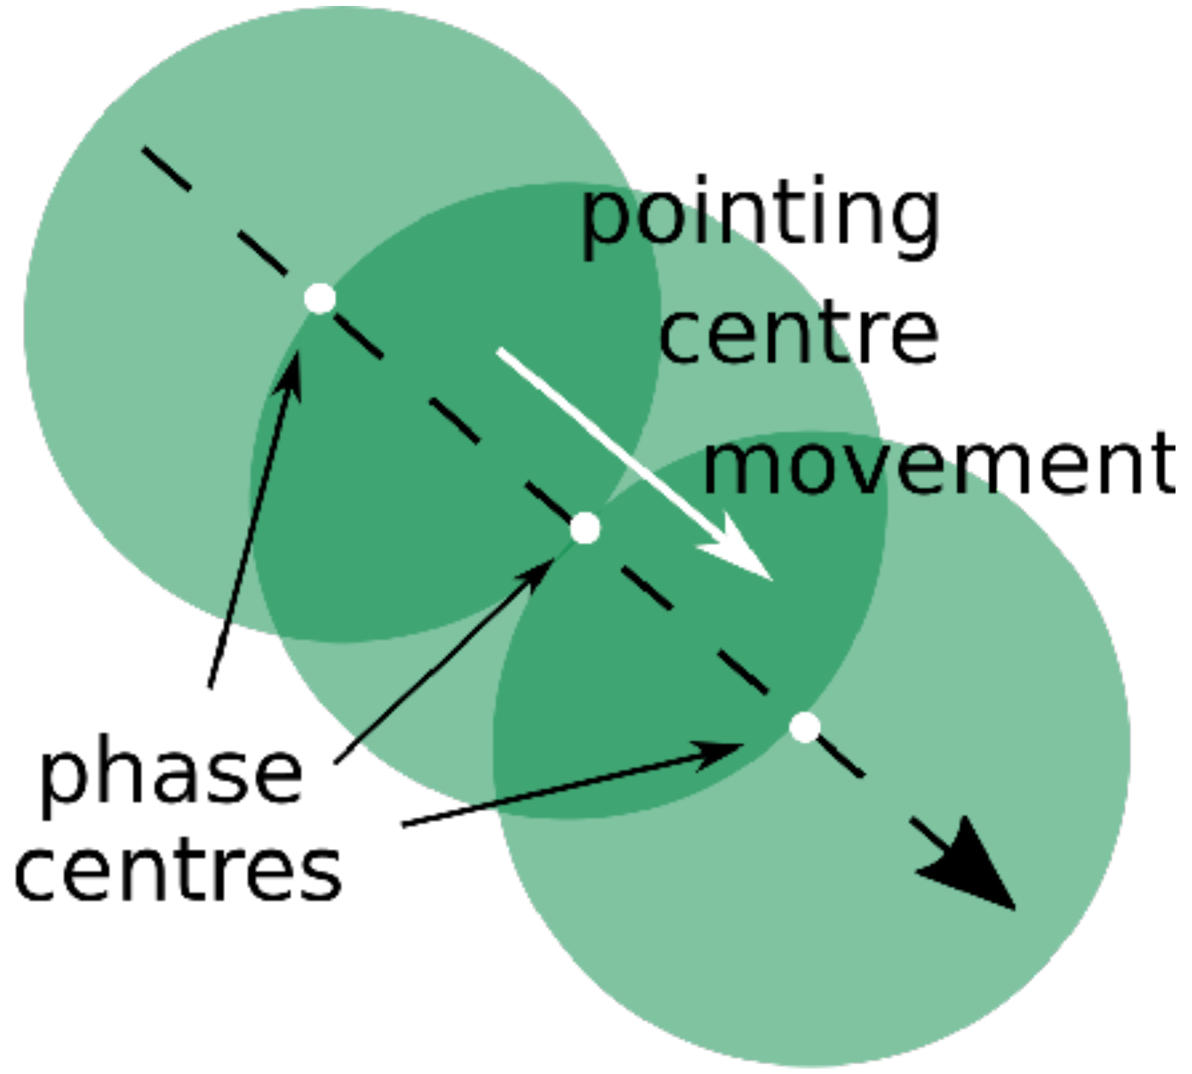

Natasha Maddox, Kristof Rozgonyi, Joe Mohr, Matthias Hoeft

5a: Enabling Science with the SKA-MPG Telescope

5b, 5c: Towards Robotic Operations:

Automated Scheduler and Data Quality Control from Metadata

(Credit: SKAO)

SKA-MPG telescope (→ talks by Hans-Rainer Klöckner, Ferdinand Jünemann)

- Hardware installation complete, currently under final evaluation
- Software prototype is ready for observations and processing
- Publication on system design and commissioning in preparation
- Characterisation of polarised beam

Automated scheduler in progress

Data Quality Control from Metadata:

- Apply data stream algorithms and general data mining methods to telescope's data stream and environment sensor data
- Testing on example data
- Planning to integrate with the automated scheduler

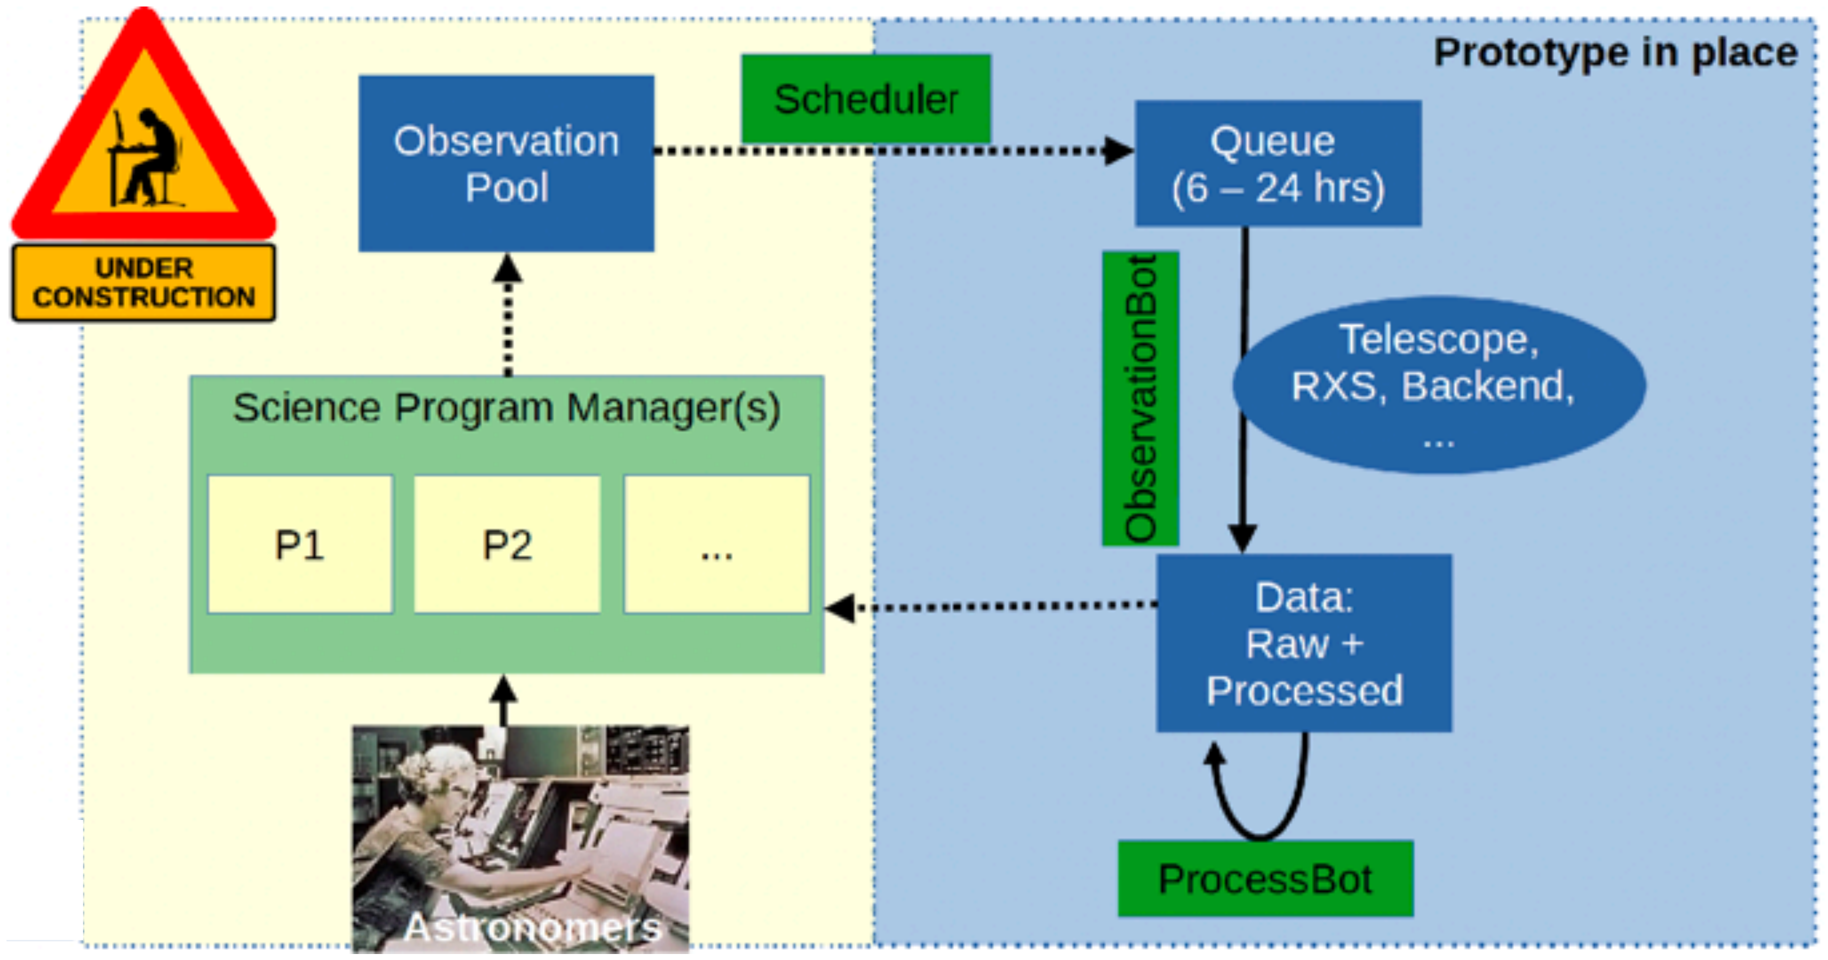

Tobias Winchen, Ferdinand Jünemann, Hans-Rainer Klöckner,
Mayumi Sato, Aritra Basu, Lena Linhoff, Dominik Elsässer, Felix Jankowsky, Stefan Wagner, Dominik Schwarz

Access to GLOW HPC cluster

GLOW HPC Cluster @ FZ Jülich

2016: Established cluster for LOFAR GLOW mode

- Housing-agreement with FZ Jülich
- Original purpose:
  data taking and preprocessing for German LOFAR stations

2019: 8 CPU nodes + 2 storage nodes

2020: HPC Upgrade → Real HPC cluster with fast interconnect

- 2GPU 4CPU nodes with large memory
- Central SSD cache and HDD storage
- Purpose:
  - Central processing resource for GLOW
  - Prepare pipelines for scalable HPC (→ D-MeerKAT WP1)

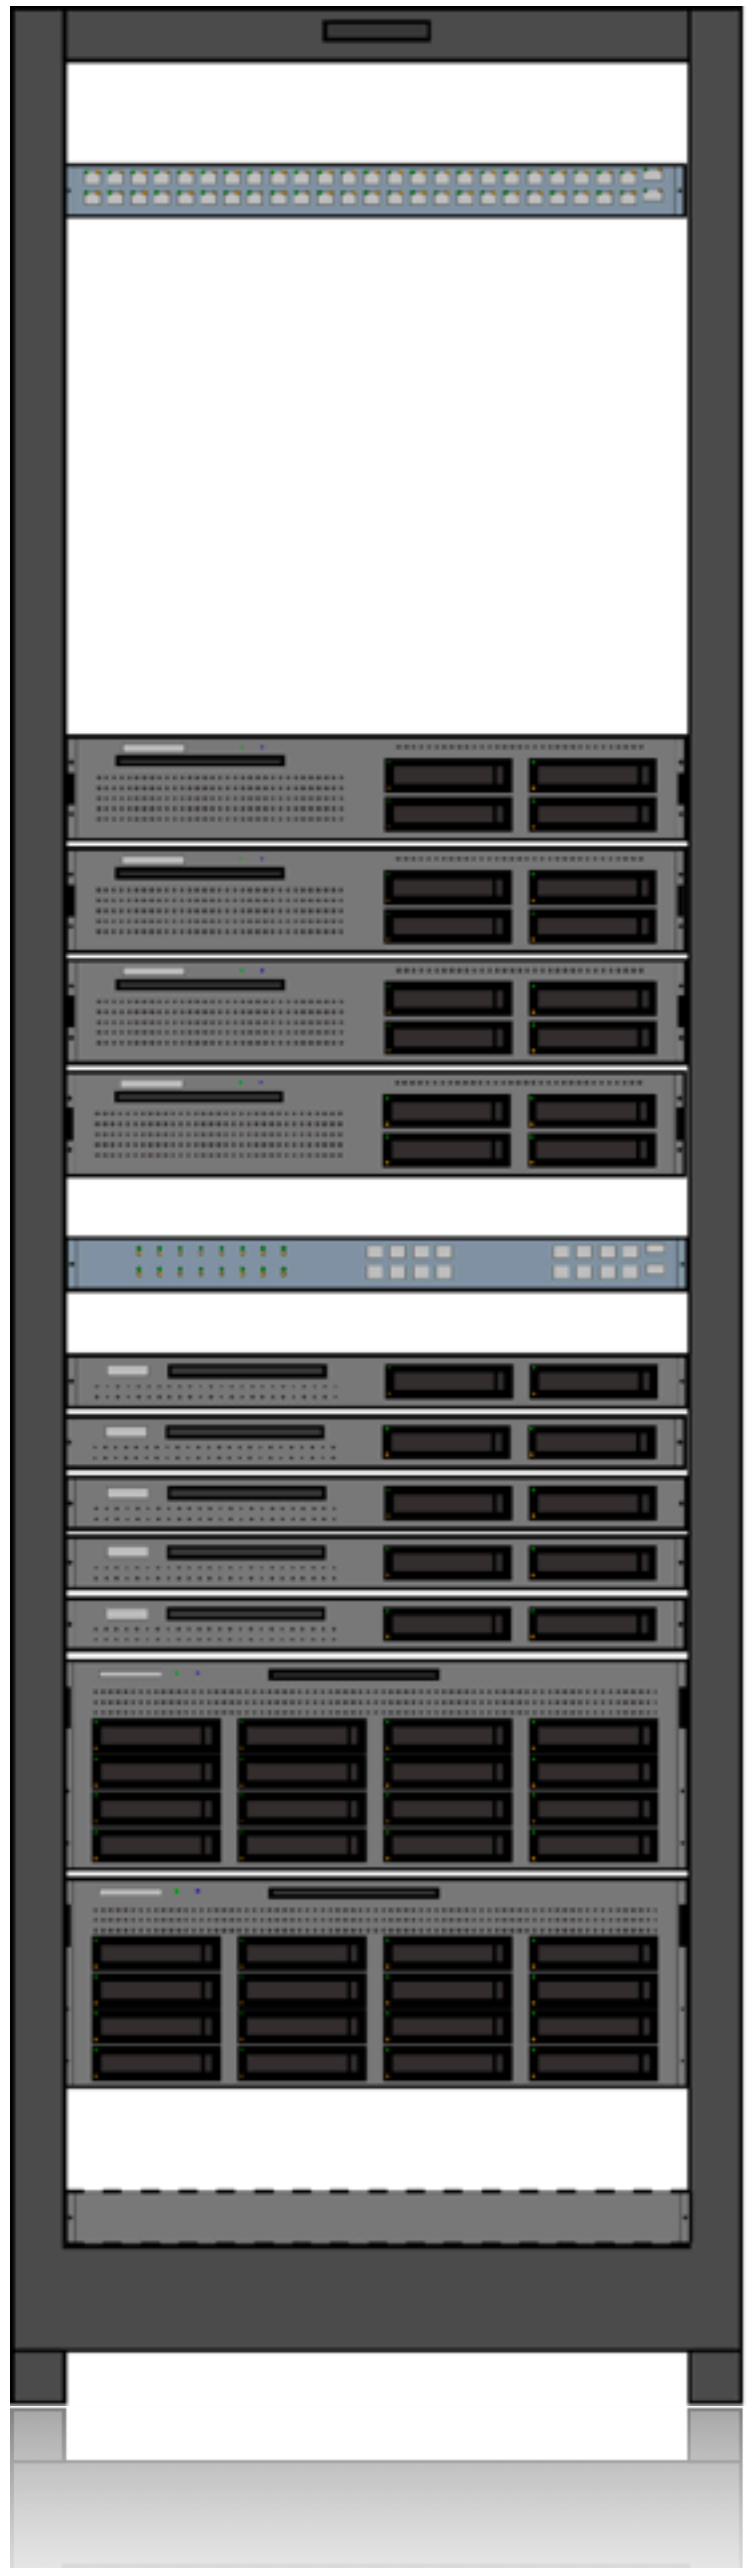

GLOW HPC Cluster @ FZ Jülich

Ready for use!

✓ Physical setup and base OS installation done (Jul 2020)

✓ Storage cluster installed (April 2021)

✓ User management + batch scheduling system installed (Jul 2021)

- Ongoing work (known issues)
  - X-forward of GUIs started via Slurm
  - Ceph performance (data placement on SSD cache)
  - GLOW cluster gateway solution is not ideal
  - Accounting/quota in Slurm to be configured

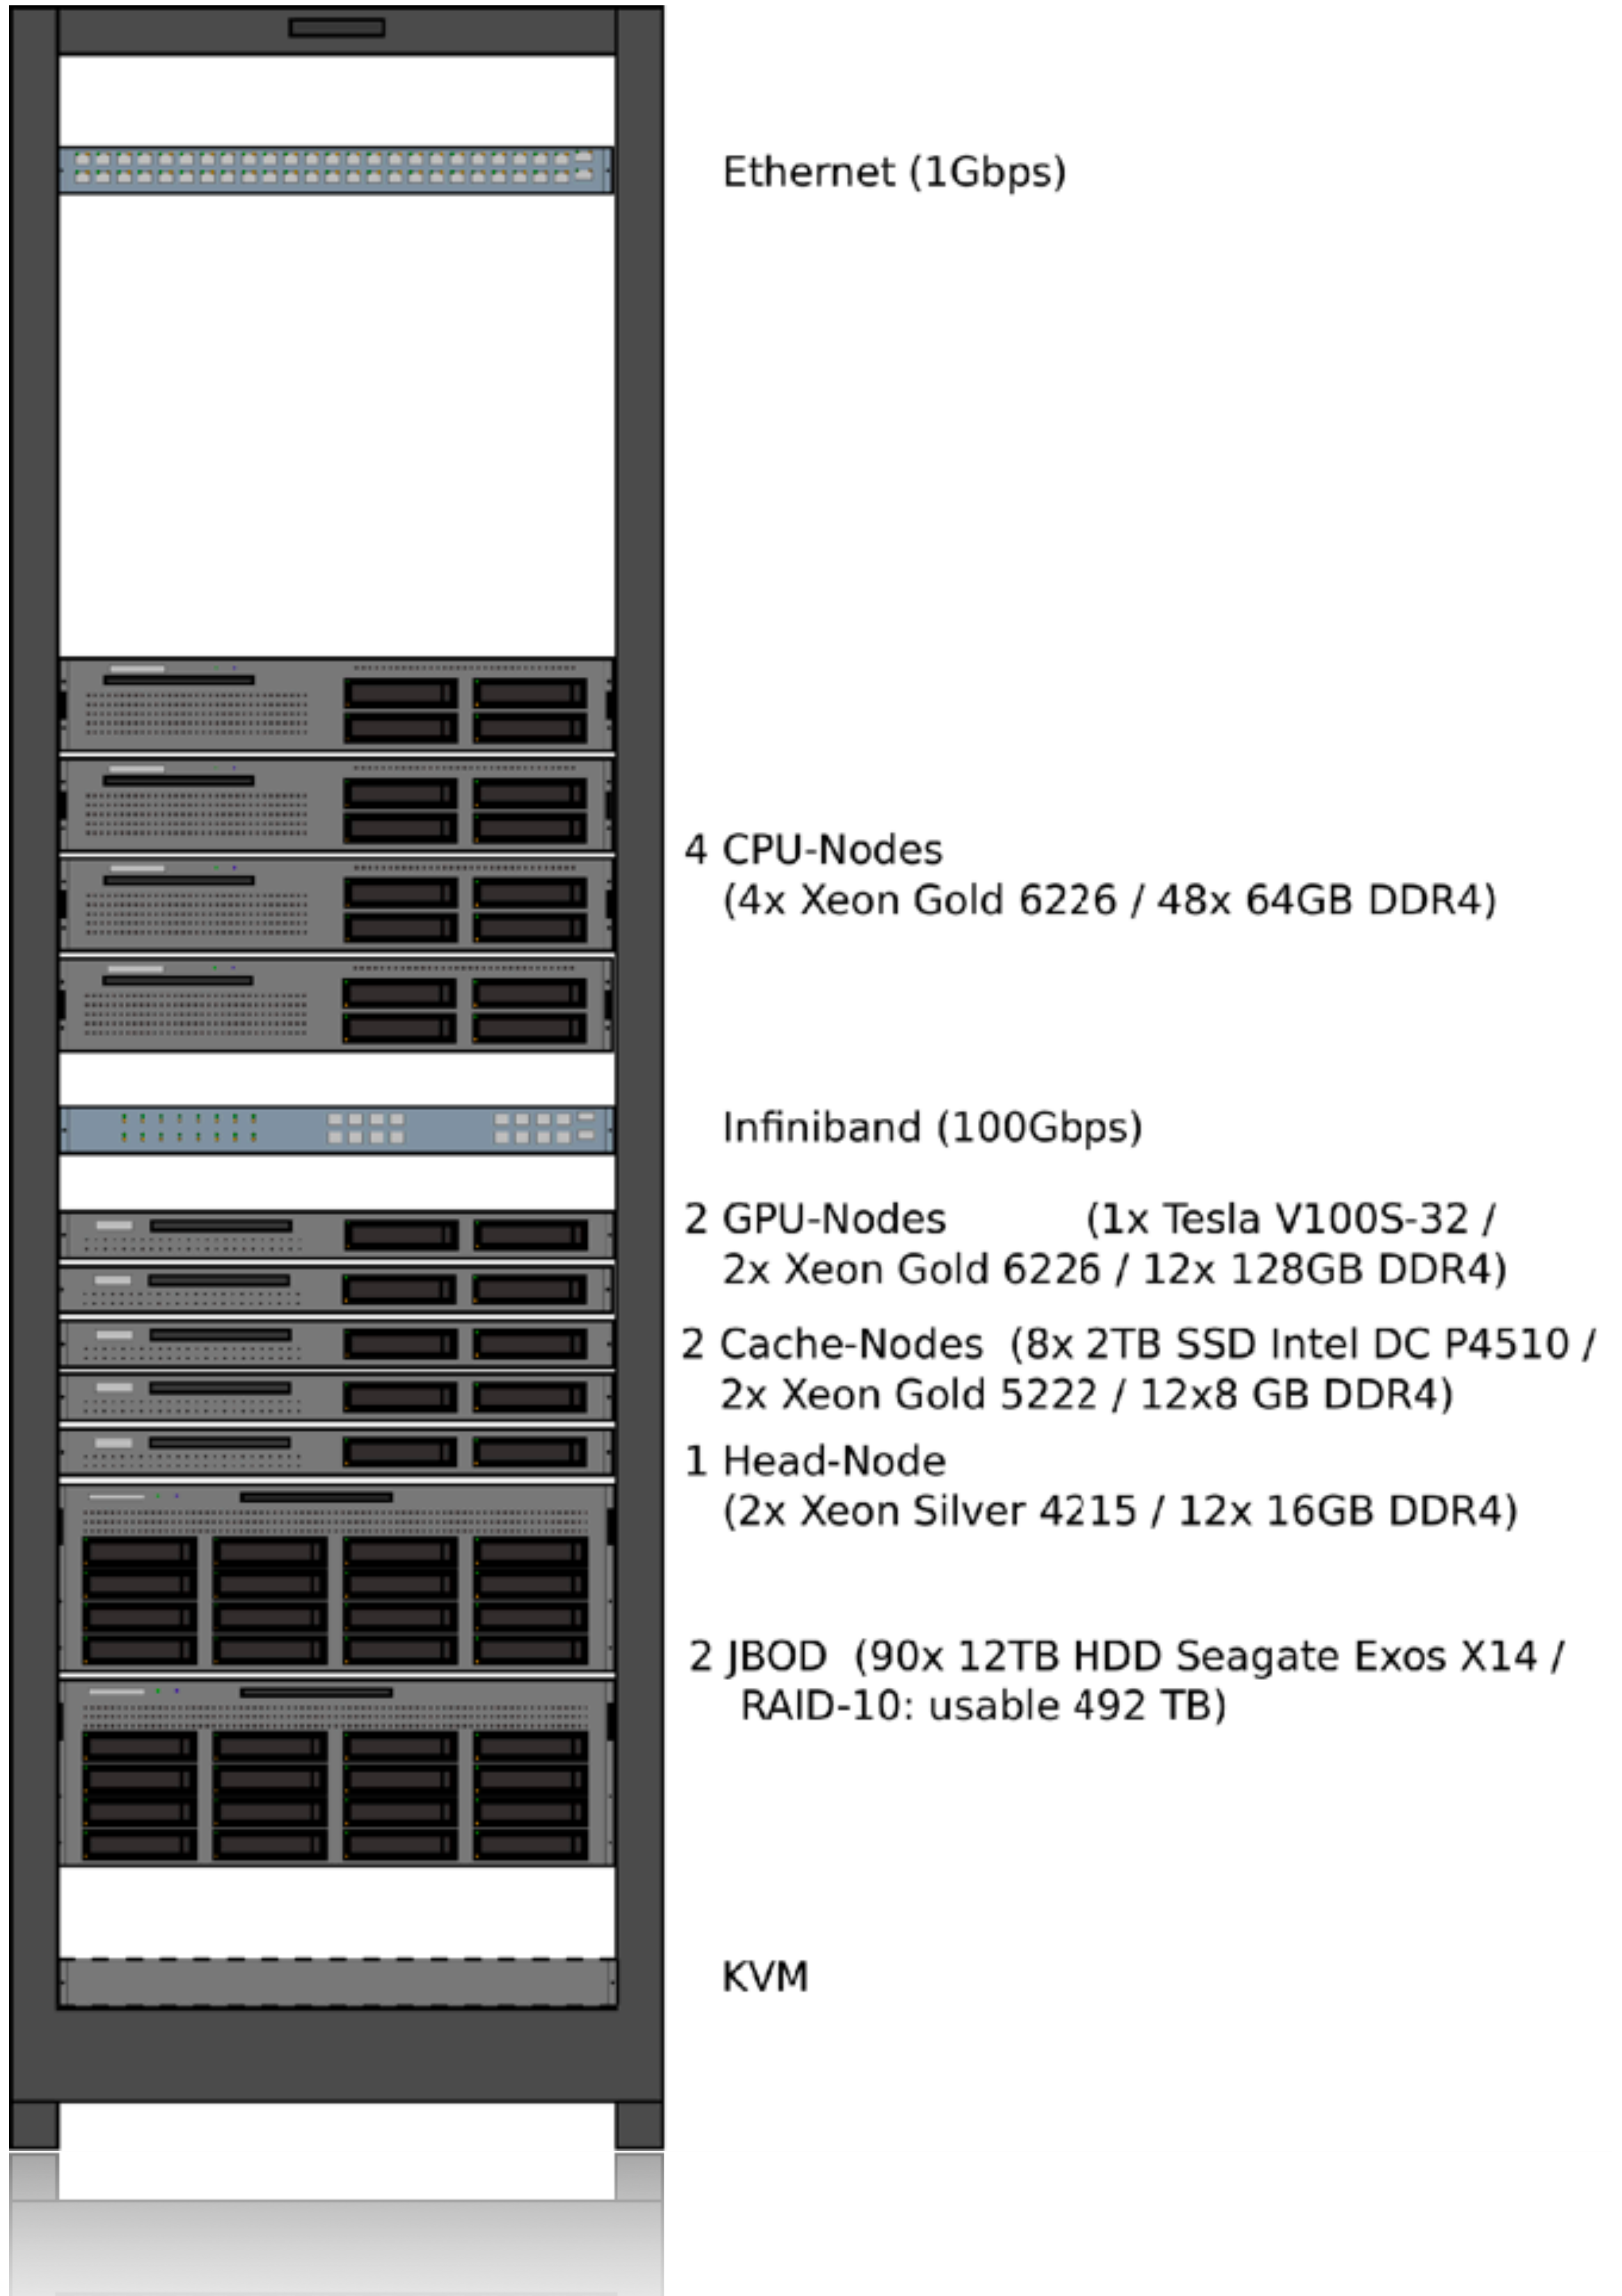

How to Access to GLOW HPC Cluster

SSH Access

- Send a **public** SSH key (and your home institution's gateway IP) to Jörn Künsemöller:

  jkuensem@physik.uni-bielefeld.de

- Linux/Mac: If you have one of these files, you can just send that:

  ~/.ssh/id_rsa.pub

  ~/.ssh/id_ed25519.pub

  Important! Under no circumstances share the corresponding *private key* (without .pub in the end) with anyone!

- If you do not have one of those files, you can generate a new key pair like this from a Linux/Mac terminal:

  ```
  ssh-keygen -a 100 -t ed25519 -f ~/.ssh/id_ed25519
  ```

  When asked for a *passphrase*, please specify one.
  It will then later be required to unlock the key when that is used.

  Important! Make sure that the specified filename *does not exist*!
  Otherwise you will overwrite an existing key, and you cannot use that any more for machines that authorised that existing key!

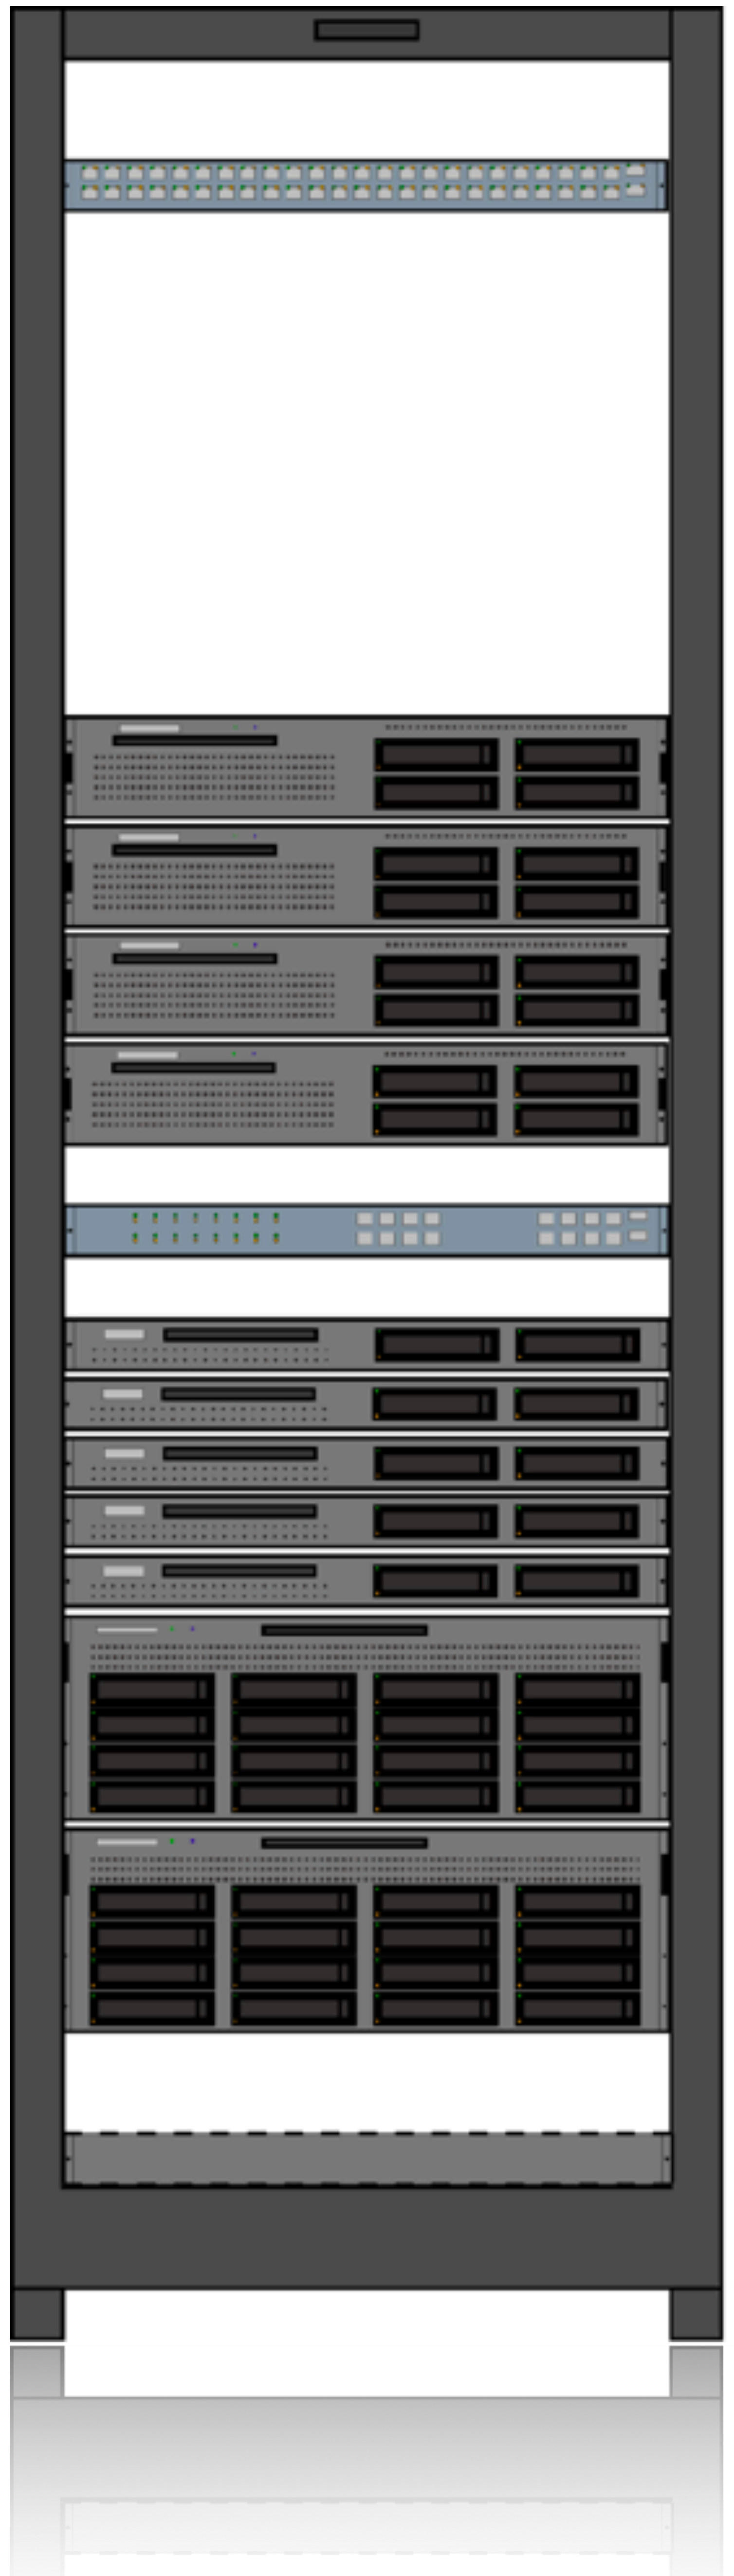

How to Access to GLOW HPC Cluster

SSH Access

- The HPC head node (*glowhead01*) is behind a cluster gateway (*glow.fz-juelich.de*; you cannot login here directly) and a firewall with an IP filter (only accessible from known IP addresses)
- Jump through GLOW cluster gateway, e.g.:

  ssh -J glow.fz-juelich.de glowhead01
  rsync -e "ssh -J glow.fz-juelich.de" example.tar.gz glowhead01:

- From home, jump through machine with whitelisted IP first, e.g.:

  ssh -J entry3.physik.uni-bielefeld.de,glow.fz-juelich.de glowhead01

- (optional) configure in ~/.ssh/config:

```
Host glowhpc
  Hostname glowhead01
  ProxyJump entry3.physik.uni-bielefeld.de,glow.fz-juelich.de
```

  With that, you can simply do:

  ssh glowhpc
  rsync example.tar.gz glowhpc:

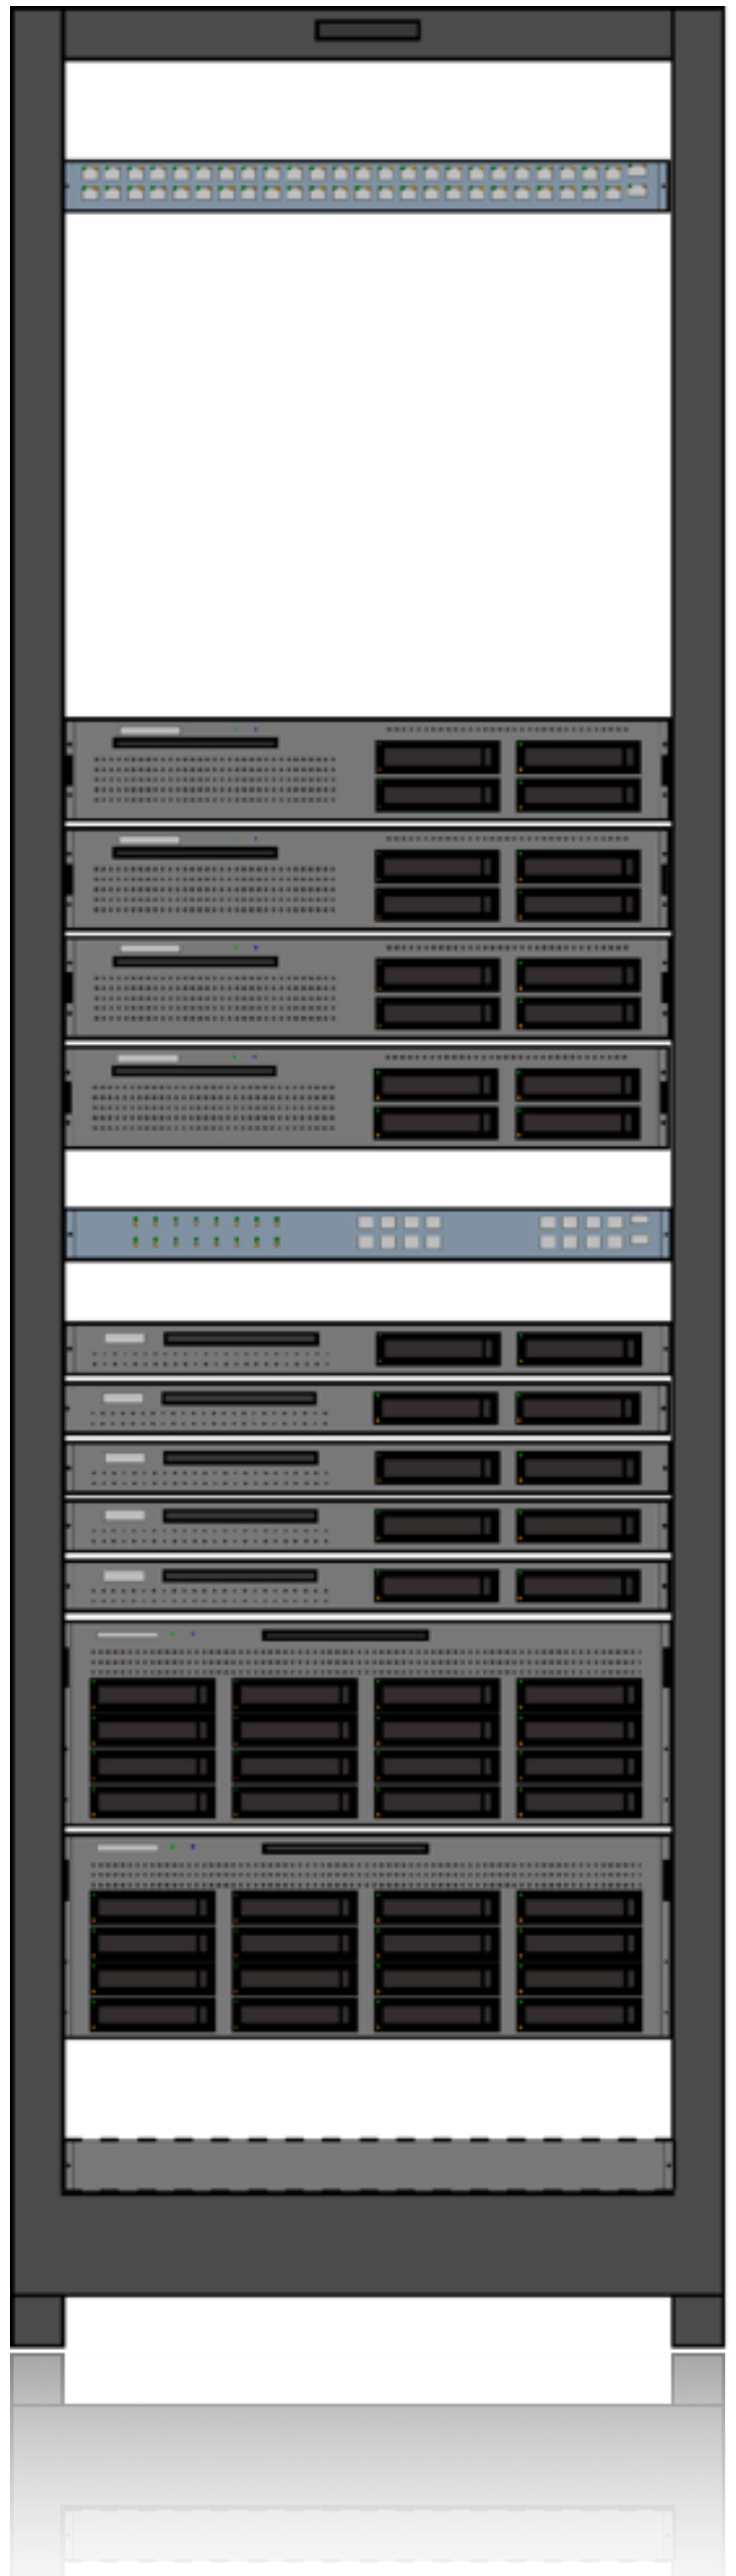

How to Access to GLOW HPC Cluster

GLOW HPC Basics

- Do **not** do any heavy lifting on the head node (glowhead01)
- Switch to worker nodes for interactive work or schedule jobs (both via Slurm → see next slide)
- Home areas live on cluster file system → same on all nodes
- Shared storage areas: /data/, /containers/, /virtualenvs/
- Common end user software to be containerized (or put in venv)

Singularity Basics

- Images are files that contain applications and all their dependencies
- That means: a base OS, all libraries, python packages, etc.
- Portable setup that is independent from host OS / installation
- Containers are running instances of these images, e.g. run:

    singularity shell /containers/astropy.sif

- Own filesystem! Mount data in, e.g.: --bind /data:/data

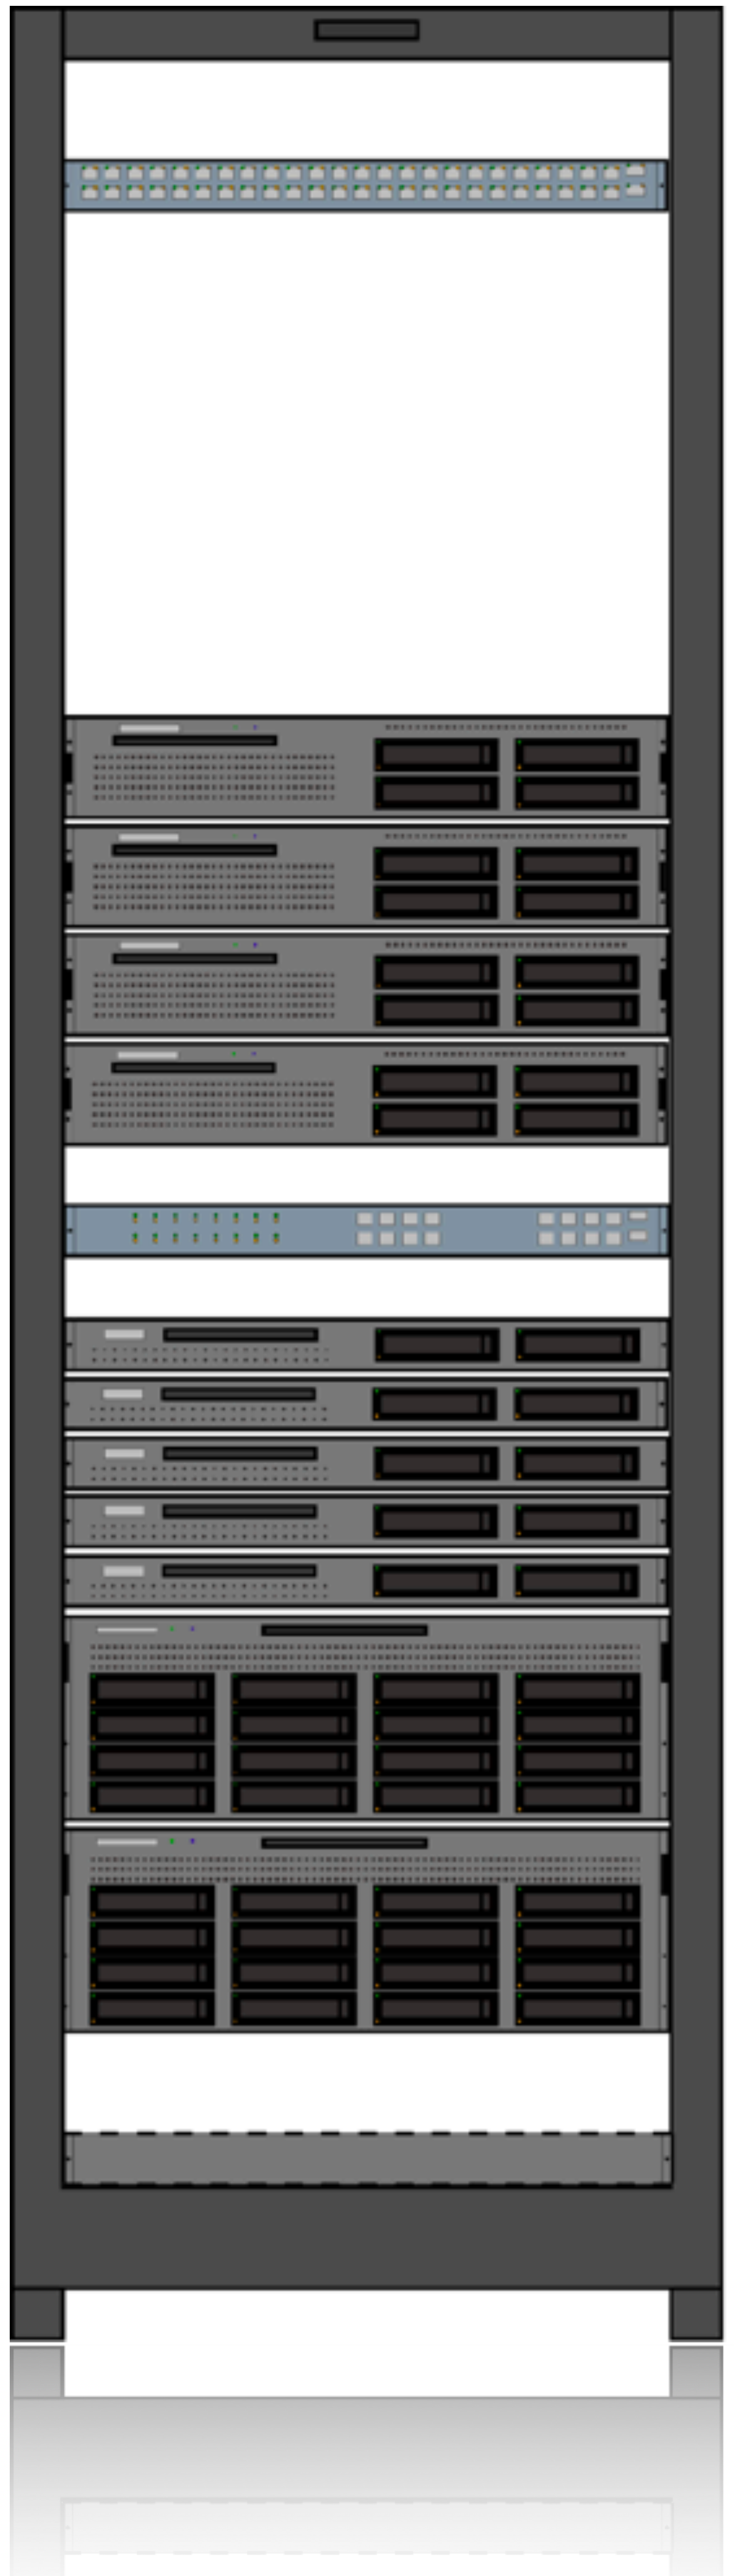

How to Access to GLOW HPC Cluster

Slurm Workload Manager

- Places your jobs / you on a worker node when/if resources available
- For interactive work (or direct execution), use srun, e.g.:

```
srun --nodes=1 --cpus-per-task=12 --partition=cpu --pty /bin/bash
```

- For scheduling a job for later execution, use sbatch, e.g.:

  First define myjob.txt:

```
#!/bin/bash
#SBATCH --job-name=myjob
#SBATCH --output=myjob.out
#SBATCH —cpus-per-task=12
#SBATCH --partition=gpu
singularity exec --nv /containers/tensorflow.sif nvidia-smi
```

  Then submit the job:

```
sbatch myjob.txt
```

- Other useful commands: sinfo, squeue, scancel (HPC) / sshfs (client)

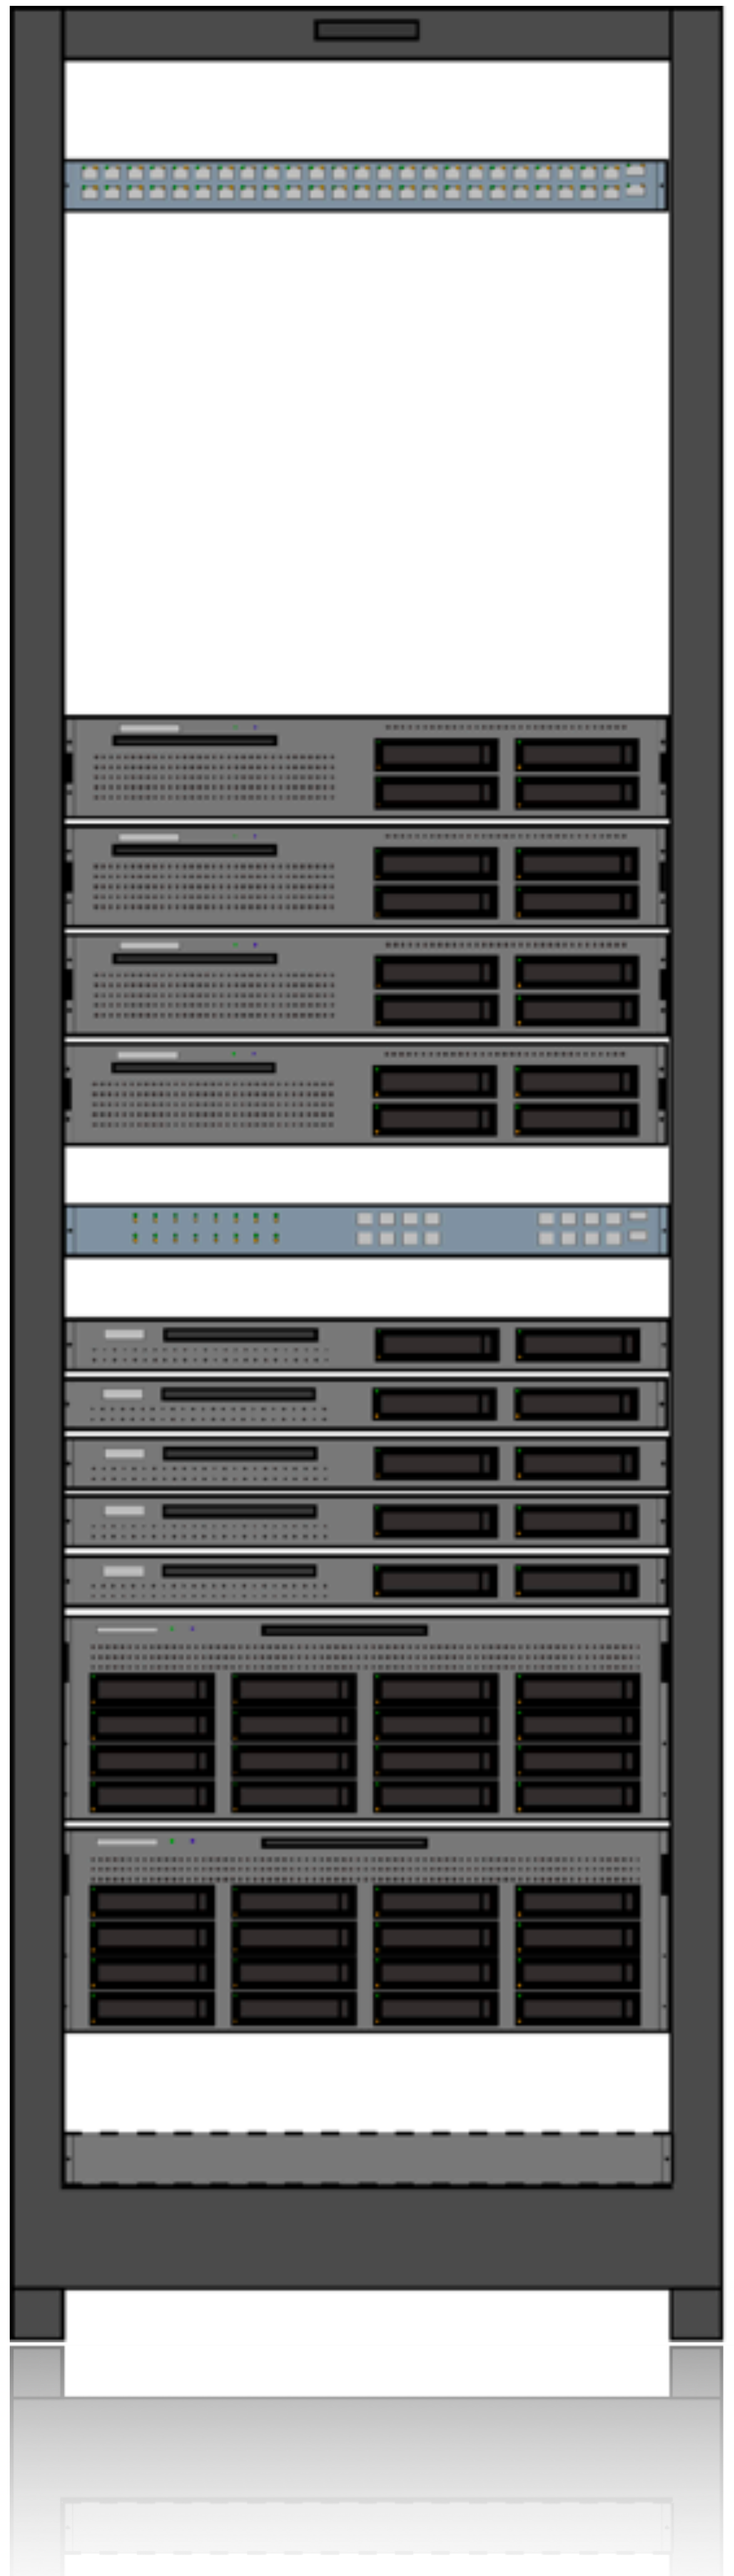

Discussion and questions on issues related to a radio astronomy science data centre in Germany