Moving to digital: technical developments in radio astronomy using the 100m Effelsberg and MeerKAT telescope as an example

E. Barr, G. Wieching

17.03.2026

MeerKAT S-Band: A challenge

S-Band Specifications

- frequency range 1.75 – 3.50 GHz:
- sensitivity T_{sys} < 18 K (T_{rec} < 6 K)
- cryogenic cooled: LNA and OMT <20K
- Highly reliable operation in harsh environments
- Operating temperature -5 - 40DegC
- weight max 80 kg
- full-fills MeerKAT and SKAO RFI requirements

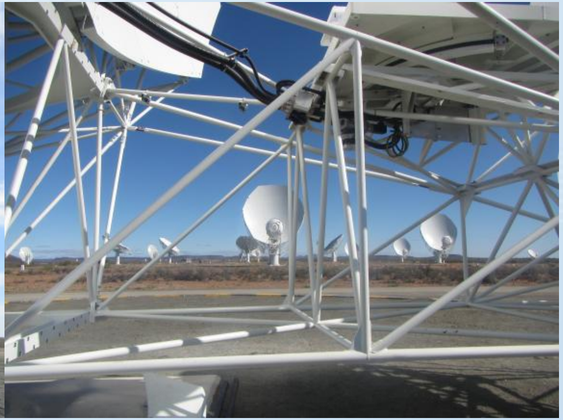

Results in a full integrated receiver:

- 2 polarisation RF Processor
- LNA (production)
- Digitizer (producing 84 GB/sec)
- Packetizer generating 70 GB/sec) ethernet packets
- Time and Reference System (Rx and Tx)
- Active Temperature Control System to 20mK
- Dewar Temperature/Vacuum readout
- Signal/Power Distribution
- Rx Control System
- MMIC Bias Supply
- Dewar Heater System
- Cryocooler Control
- Vacuum Control
- Power supply

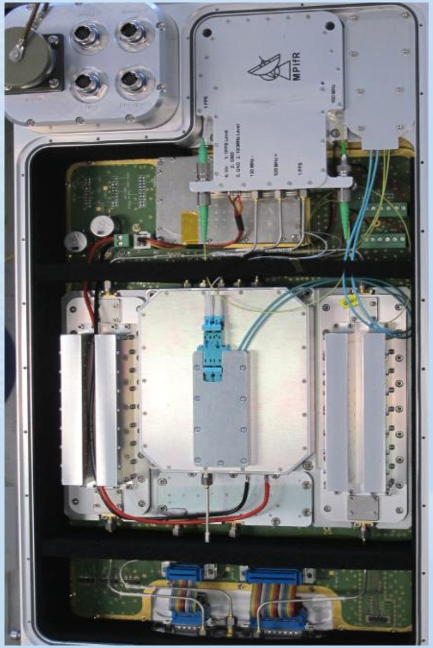

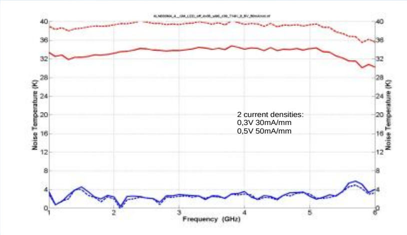

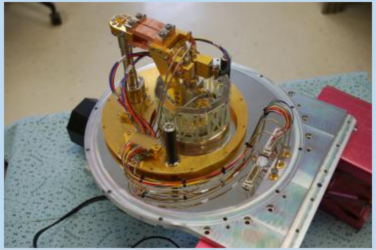

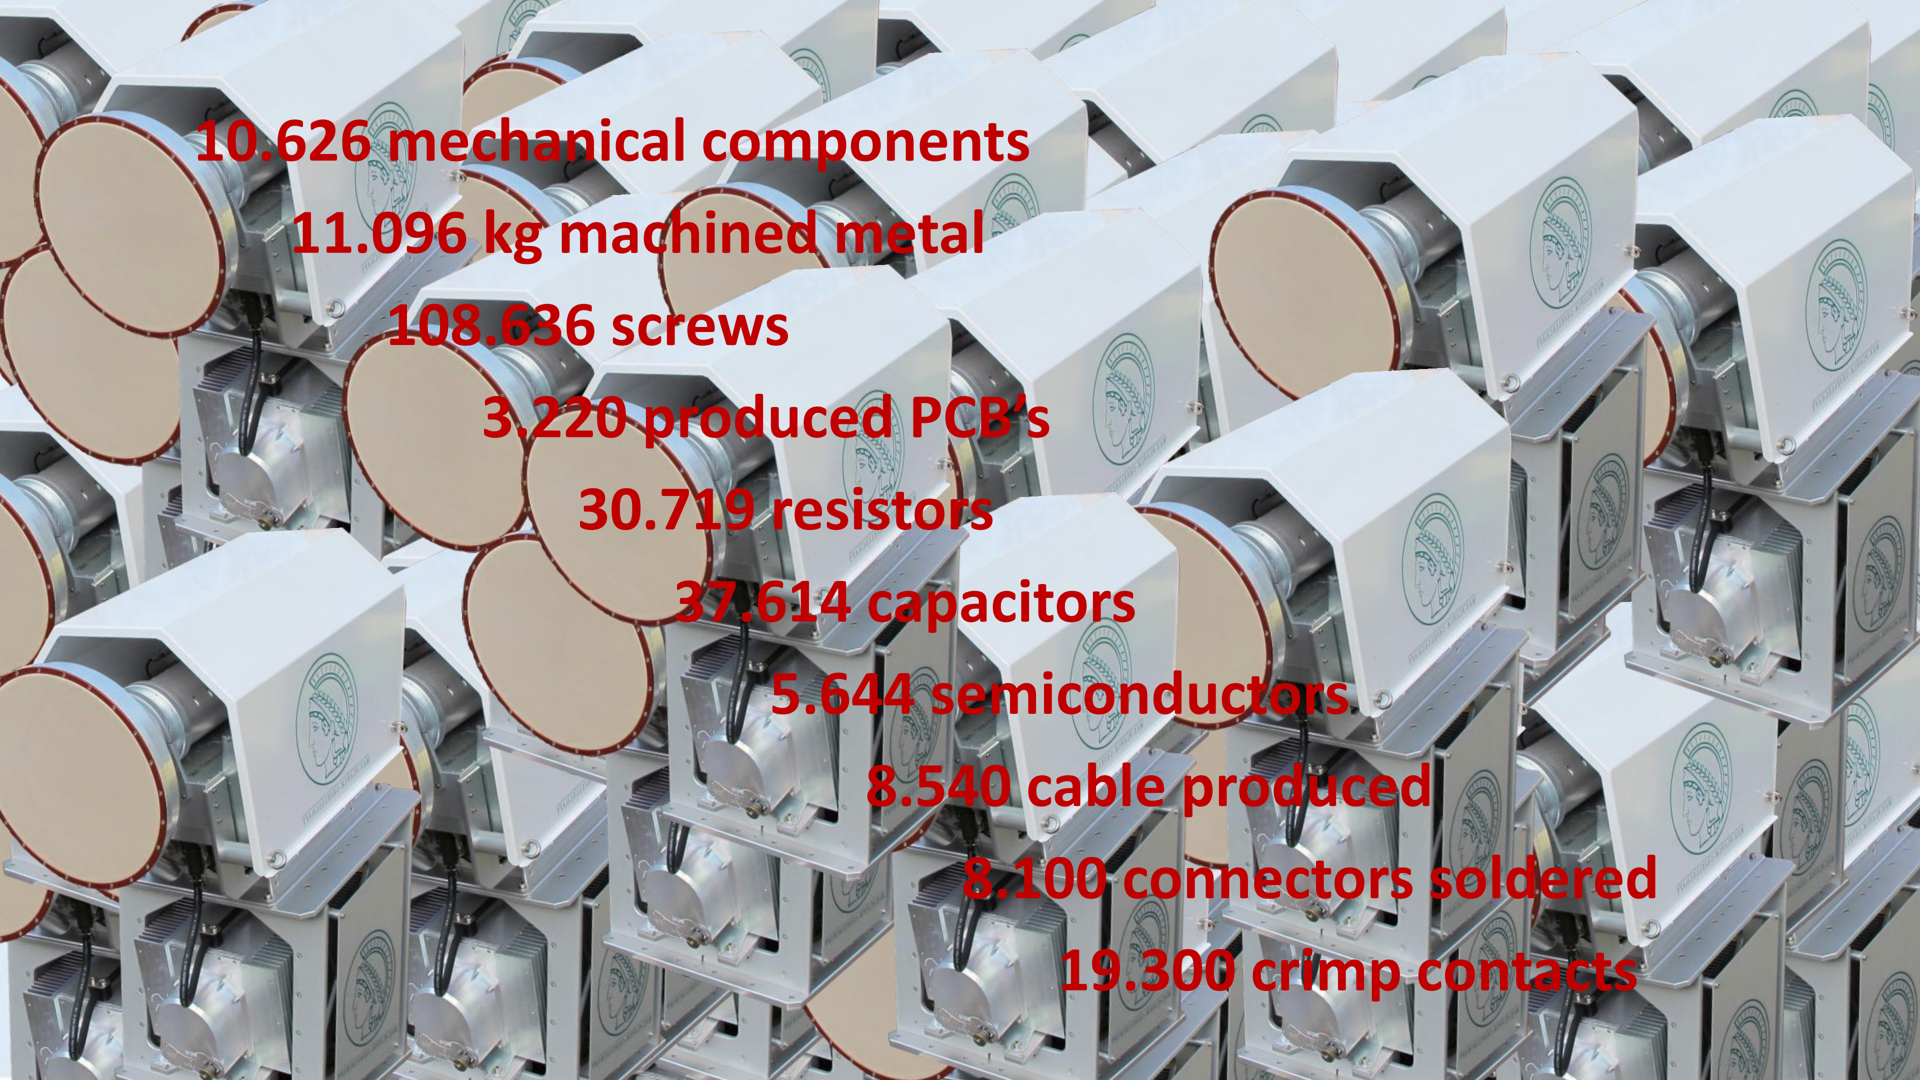
10.626 mechanical components
11.096 kg machined metal
108.636 screws
3.220 produced PCB's
30.719 resistors
37.614 capacitors
5.644 semiconductors
8.540 cable produced
8.100 connectors soldered
19.300 crimp contacts

MeerKAT S-Band: A challenge!

- A technological challenge
  - Provide performance in a difficult environment
  - especially RFI and engineering
- Production challenge:
  - Requires different skills:
    - system engineering
    - documentation system
    - MTBF analysis
    - operations and maintenance scheme
    - management software
- But we had time! Approx 5 years

INAF MeerKAT Band 5b - 68 receivers

Band 5b Specifications

- frequency range 8.3 – 15.4 GHz
- sensitivity Trec < 12K average (measured 8-10K)
- cryogenic cooled: LNA and OMT <20K
- highly reliable operation in harsh environments
- environmental Operating temperature -5 - 40DegC
- full-fill SKAO RFI requirements

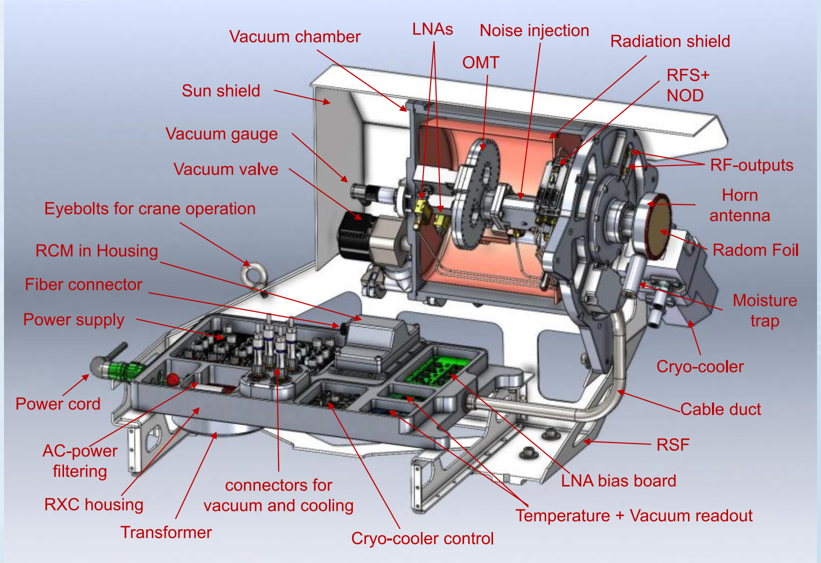

Challenging Timeline (18month):

- Call for tender (April 2024)
- Tender submission (May 2nd 2024)
- Tender awarded (June 2024)
- Kick-off meeting (July 2nd 2024)
- Prototype FAT (Mid Nov. 2024)
- Prototype SAT (Mid Dec. 2024)
- Start Production (January 2025)
- All orders placed (February 2025)
- All PCBs and >90% of all components are available (May 2025)
- 1st delivery of 41 units (September 2025)
- 2nd delivery of 10 units (October 2025)
- final delivery of 16 units (November 2025)

and the next system in preparation

SKAO Band 5ab (in cooperation with OHB DC)

- 86 receiver systems
- frequency range 4.6 – 8.5 GHz and 8.3 – 15.4 GHz
- cryogenic cooled: LNA and OMT <20K
- Highly reliable operation in harsh environments
- Operating temperature -5 – 40DegC
- full-fill SKAO RFI requirements and documentation

Will single-dish telescopes still have a place at the forefront of scientific research?

Yes, but there is a need to invest in industrial-grade multi-feed receivers and massive data processing to be competitive with array.

Phased Array Feeds

- Phased Array Feeds (PAFs) are multi-element receiver systems placed at the focal plane of a radio telescope.
- Electronically form multiple simultaneous beams to expand the effective field of view.
- Digital beamforming enhances sensitivity, survey speed, and RFI suppression.
- Applications include wide-field surveys, transient detection, and improved imaging performance.
- Requireds high-performance digital processing

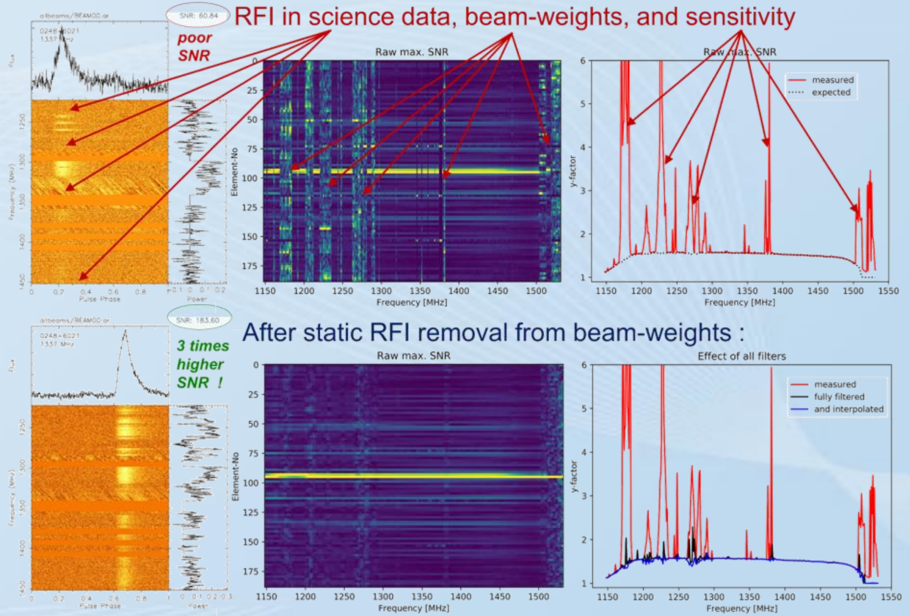

Spatial and intensity-selecting filtering on array covariance matrices followed by beam-weight interpolation removes most of the RFI in beam-forming observations

➔ much improved S/N on science observations

Cryo-PAF for Effelsberg

focal plane layout and performance

- number of elements : 253
  - elements in x-polarization : 124
  - elements in y-polarization : 129
  - spacing (within one polarization) : 42 mm
  - arrangement : squared
- focal-plane diameter : 600 mm
- number of formed beams : <128
  - FPGA based beam former
  - GPU based backend system
- target sensitivity : Tsys/η < 40 K
  - highly integrated cryogenic frontend
  - LNA developed in cooperation with IAF
- RF-bandwidth (monolithic) : 2.65 to 3.85 GHz
  - RFI filter in front of the LNA
  - Digitizer directly attached to analog signal chain
  - full digital signal transport from frontend to backend

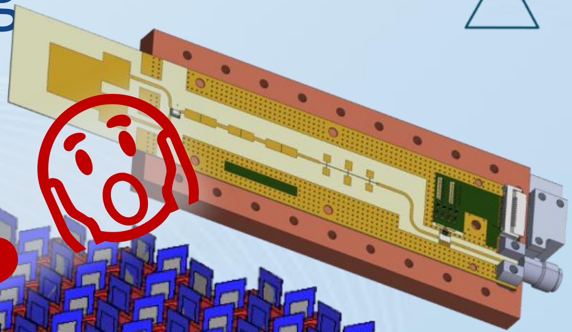

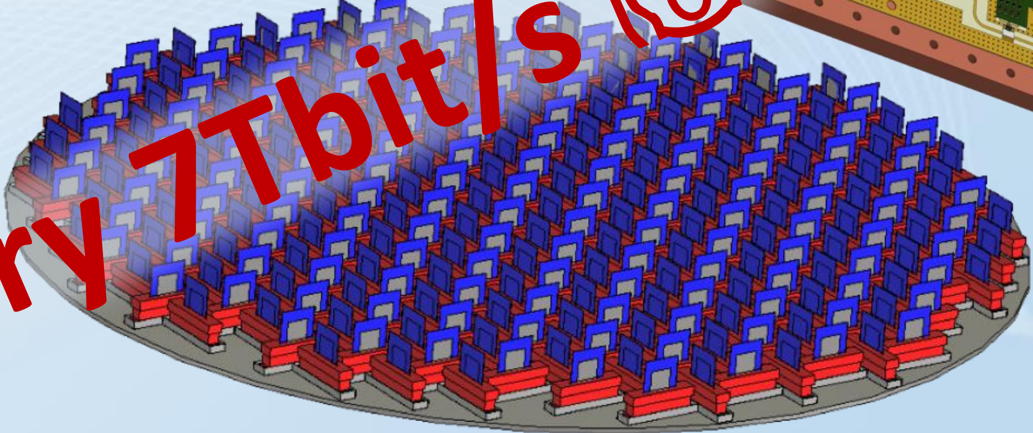

scary 7Tbit/s

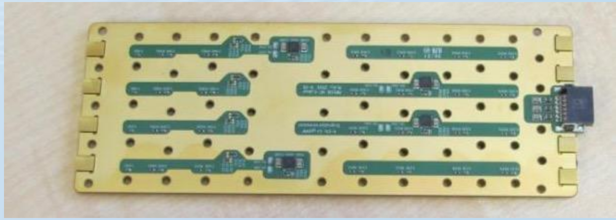

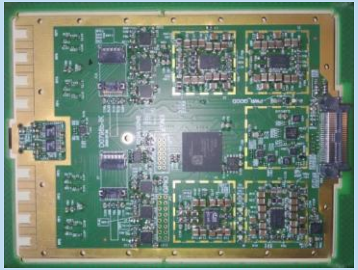

FX-Correlator/Beamformer

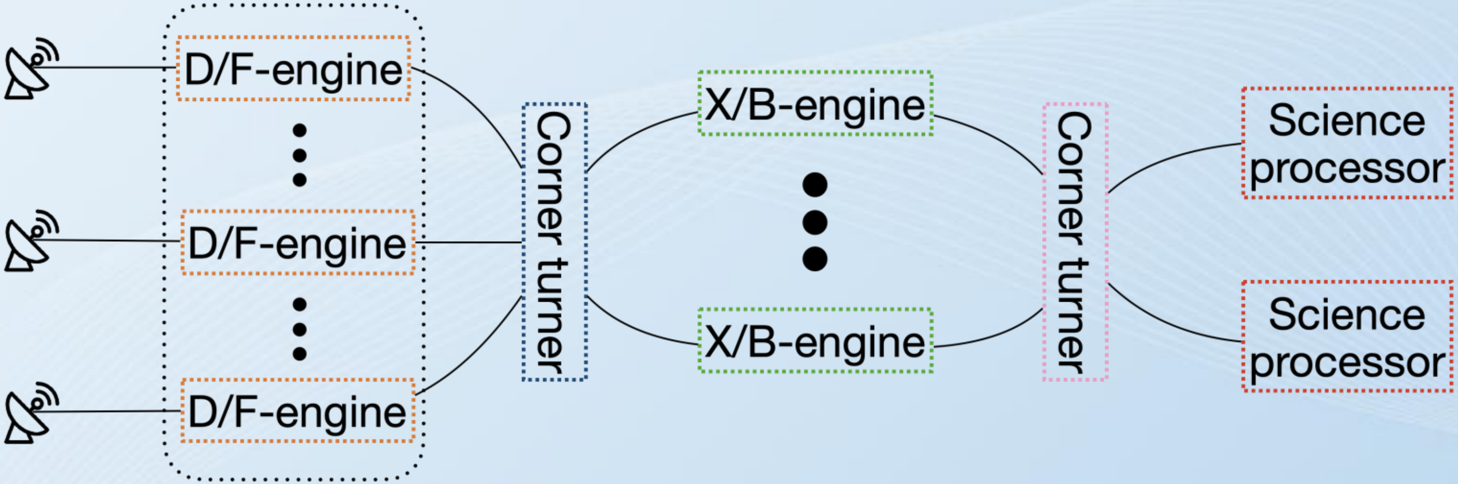

D/F-engines digitise and channelise the antenna signals

X/B-engines perform correlation and beamforming

Perform all downstream processing for observatory science cases

First corner turner transposes to put common frequencies together

Second corner turner reconstructs the band for all visibilities and beams

Radio Astronomy has Become a Digital Challenge

From disk to dedicated real-time

FFTS

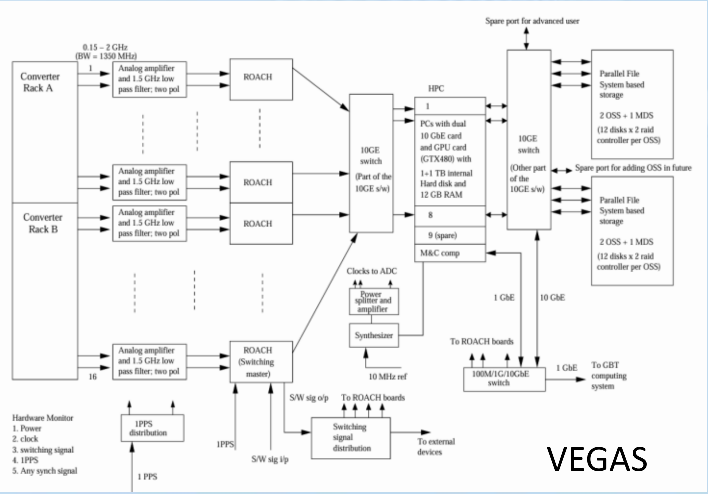

VEGAS

- Early 2000s: FPGA/DSP-based baseband processors like FFTS spectrometers, Berkeley Pulsar Processors, ATNF DFBs emerged.
- Instruments dedicated to specific tasks (e.g., spectropolarimetry, coherent dedispersion) with limited flexibility.
- These instruments usually required custom hardware and had limited versatility / flexibility.
- Late 2000s: Hybrid FPGA+GPU backends debuted, including CASPSR (Parkes), GUPPI, and VEGAS (10 GHz bandwidth) @ GBT.

From dedicated real-time to universal real-time

- Driving forces behind GPU backends: **versatility**, **scalability**, **extensibility and availability of hardware**.
- Universal backends are now being realised: **Medusa** @ Parkes (Hobbs et al. 2020) and the **Effelsberg Direct Digitisation (EDD)** backend @ Effelsberg, SKA-MPI and the TNRT.
- These provide commensal spectroscopy, continuum, pulsar search/timing, transient search, VLBI recording and baseband recording capability on commodity off-the-shelf hardware.

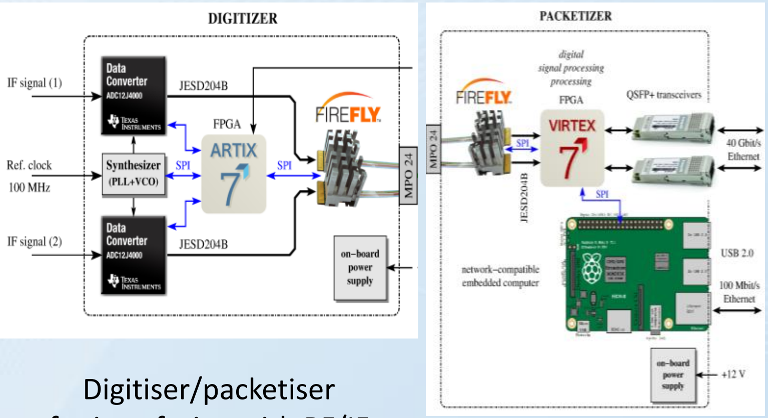

Digitiser/packetiser
for interfacing with RF/IF

High speed multicast Ethernet
networks for data distribution

General purpose GPU
compute clusters
with large storage capacity

FX-correlator Implementation

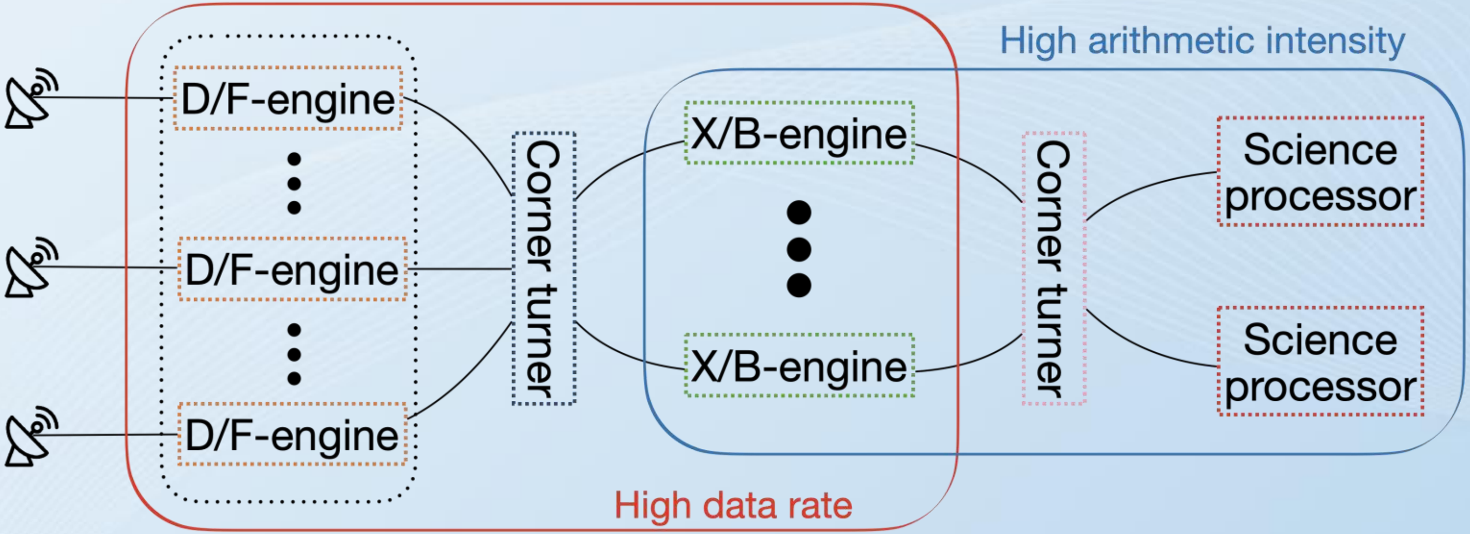

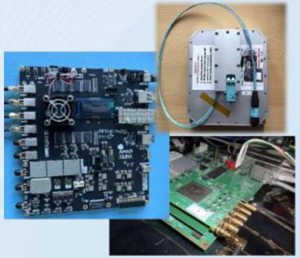

Low compute, high throughput domain well suited to **FPGAs, RFSoC, and ASICs**

High compute, high/low throughput domain well suited to flexible **GPU systems**

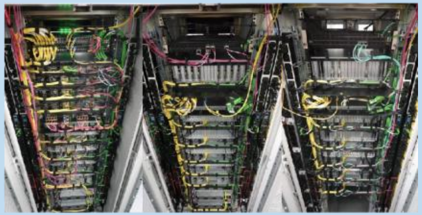

EDD Backend System

The EDD Backend

- Real-time fast transient search
- Modular design
- Any-to-any data routing
- High-speed networking
- Grafana influxdb Monitoring & visualisation
- Correlation Beamforming
- Pulsar timing & search
- Virtualisation
- Commodity hardware
- puppet Configuration management
- Code
- Polarimetry
- Continuum spectroscopy
- VLBI
- Networked data storage
- redis PostgreSQL Metadata management
- docker PROXMOX Container/VM orchestration
- CI/CD & Packaging


Extensibility via plugin architecture
Reusable core for all telescopes
Versatility through module selection
Modular design
Any-to-any data routing
Multicast UDP for data distribution
Dynamic routing to processing nodes
Multiple science cases, same data
Flexible deployment and rollback
Provides pipeline isolation
Ensures reproducibility
Virtualisation
Technologies
High-speed networking
Up to 800 GbE and NDR Infiniband
Supports extreme data rates from next-generation receivers
COTS GPUs, FPGAs, servers, switches
Support many hardware configuration
Simplified support and maintenance
Commodity hardware
Networked data storage
Massive IO at scale
Supports PB-scale volumes
Access from all nodes/pipelines

Tensor Core correlator and beamformer
Includes calibration routines for PAFs
Scales to 1000s of beams

DSPSR/PSRCHIVE/TEMPO2 stack
Folding, search and baseband recording
Direct sampling and PFB support

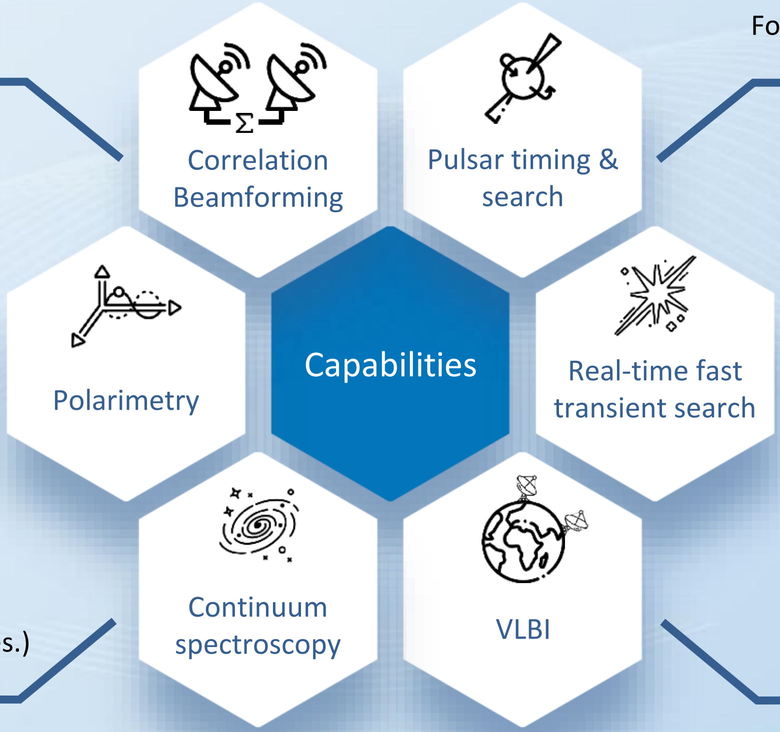

IQUV pulsar timing and search
XX,YY and IQUV spectroscopy

TransientX-based pipeline
Optimised CPU implementation
Interactive candidate viewing

Noise diode gated spectrometer
Up to 32 million channels (< 100 Hz res.)
“Zoom bands” for reduced data rates

GPU-based DDC implementation
VDIF-formatted outputs
e-VLBI streaming with jive5ab

Monitoring points recorded in InfluxDB
Dynamic Grafana dashboards
Loki for logging and fault finding

Cluster configuration via Puppet
EDD configuration via Ansible
Configuration as code (YAML + JSON)

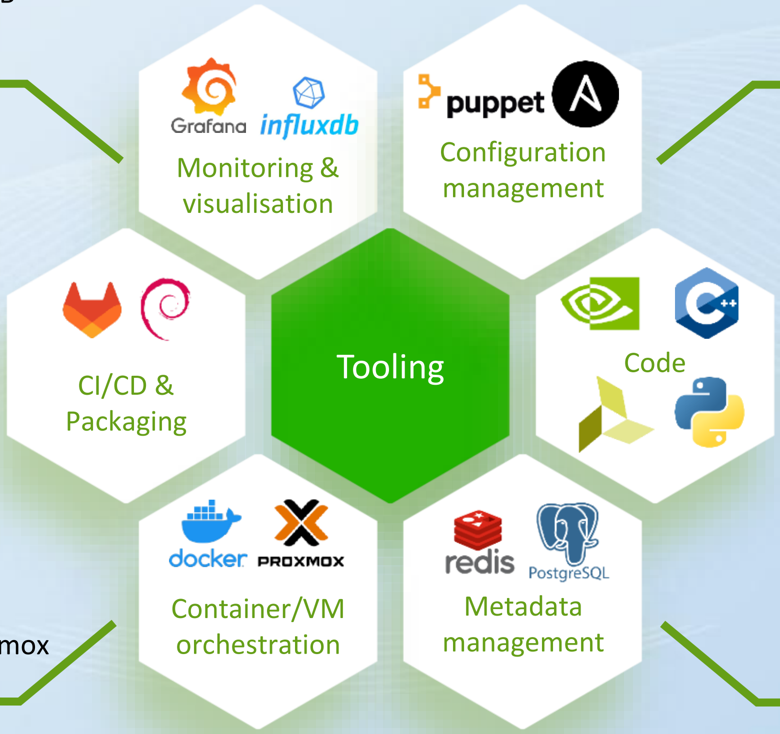

Continuous integration and deployment with GitLab
Aptly for Debian packaging

CUDA/C++ for GPU applications
Vivado HLS for FPGA applications
Python control layer

Infrastructure virtualisation with Proxmox
Application virtualisation with Docker

Metadata caching via Redis
Deployment and measurement tracking via PostgreSQL

EDD @ Effelsberg

MAX-PLANCK-INSTITUT FÜR RADIOASTRONOMIE

EDD @ Effelsberg

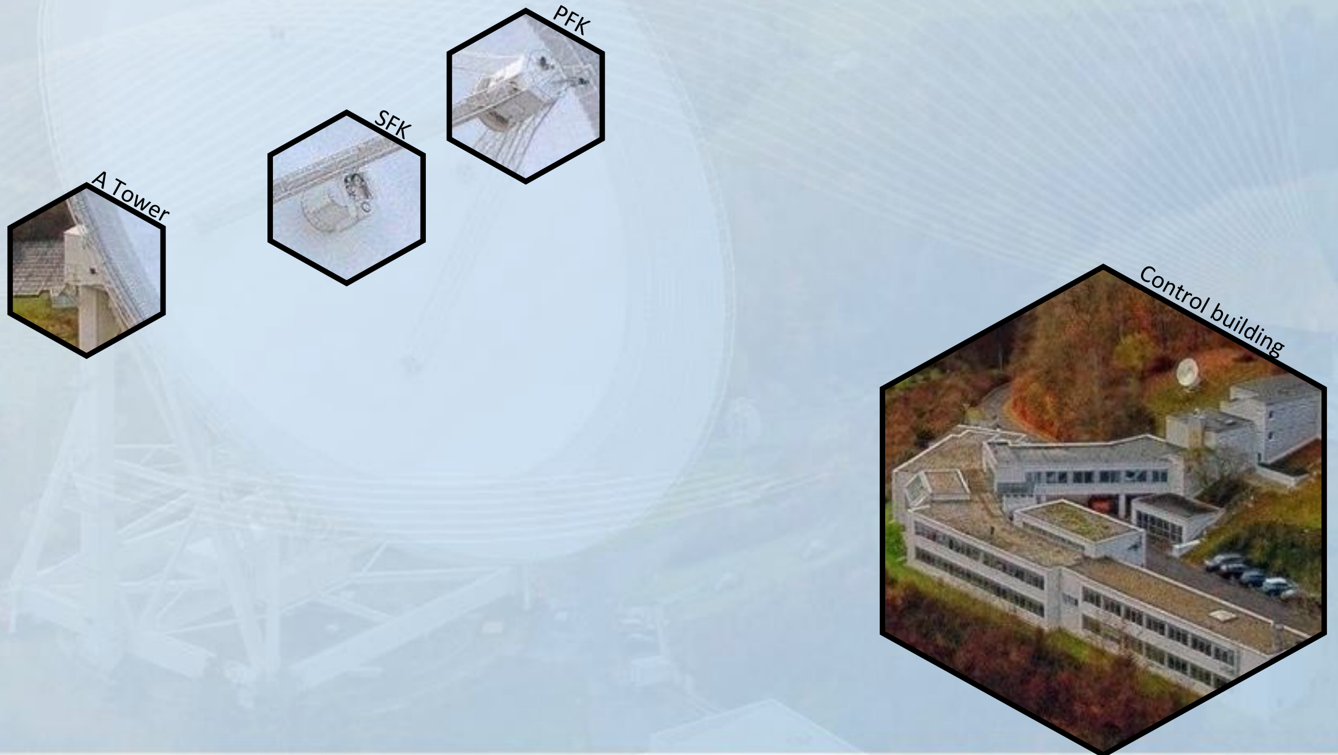

EDD @ Effelsberg

PFK SFK A Tower

16 x 6 Gsps dual-poln packetisers

11 x 4 Gsps dual-poln packetisers

256 x 2.6 Gsps CryoPAF channelisers

3 x Nvidia SN3700c 100-GbE network switches

472 GHz aggregate bandwidth

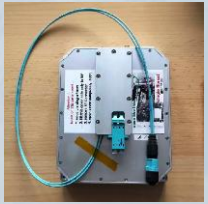

EDD sampler unit

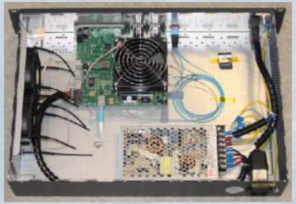

EDD packetiser

Control building

3 x 4 Gsps dual-poln digitiser/packetiser pairs (legacy Rx support)

Nvidia SN5600 **800-GbE** switch

Nvidia SN5400 **400-GbE** switch

Nvidia **400-Gb/s NDR** IB switch

8 x virtual machine hosts

32 x GPU servers (64 x N**vidia L40**)

FPGA server (8 x Xilinx **Alveo U55c**)

4 x Storage servers
(**4 PB** capacity, **68 GB/s** write speed)

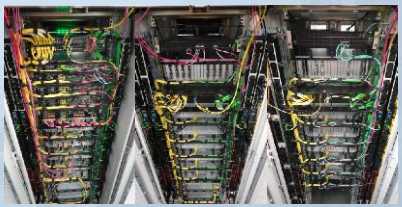

GPU servers and network switches

VM farm

Storage node

EDD Backend @ Effelsberg

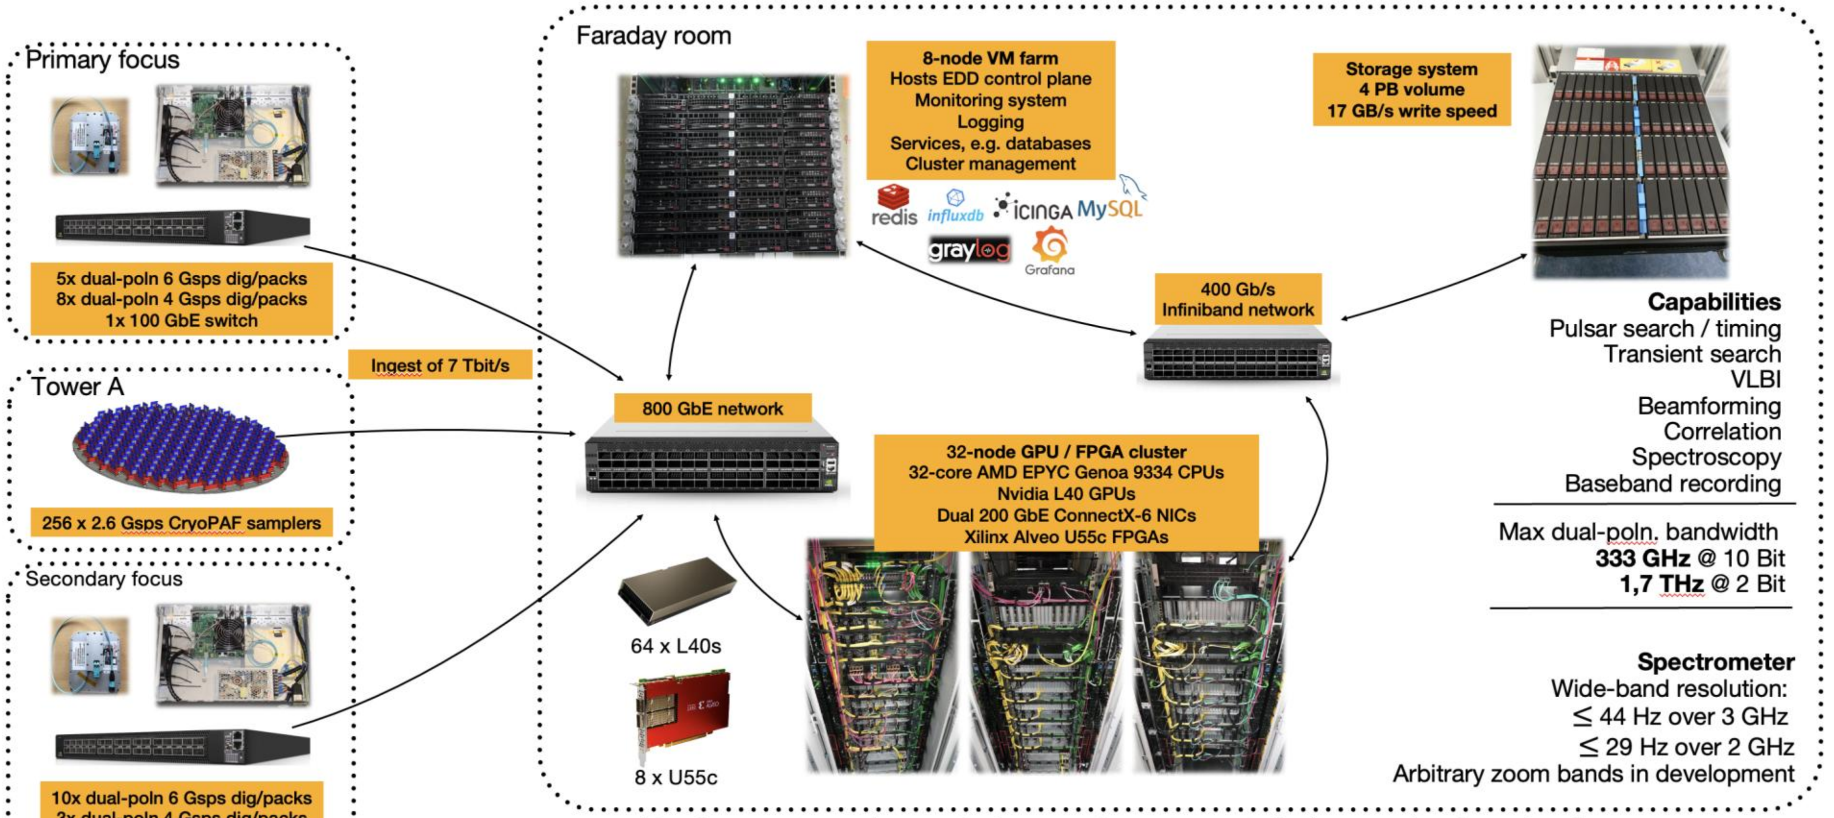

Deployment can be scaled up or down to match the needs of each telescope or receiver

EDGAR/EDD Ecosystem

Layer 3: Pipeline provisioning, observation monitoring and user processing

EDD Core + Plugins | Grafana | Dash by plotly | influxdb | Flask

Containerised science pipelines
Pipeline monitoring Web-UI
Network traffic mapping

HTCondor
User jobs

Layer 2: Backend build, deployment and personality management

Backend installer Container registries
Plugin and 3rd-party dependency packaging
Metadata and observation tracking databases

PostgreSQL | redis | ANSIBLE | HTCondor | APTLY | docker | EDD Core

Mostly site agnostic

Mostly site specific

Layer 1: Managed server installation and configuration

FOREMAN | iCINGA | PROXMOX | puppet | telegraf

Networking NTP Drivers Local file system
Virtualisation SSH Authentication Logging
Distributed file system Monitoring/alerting

Layer 0: Physical hardware

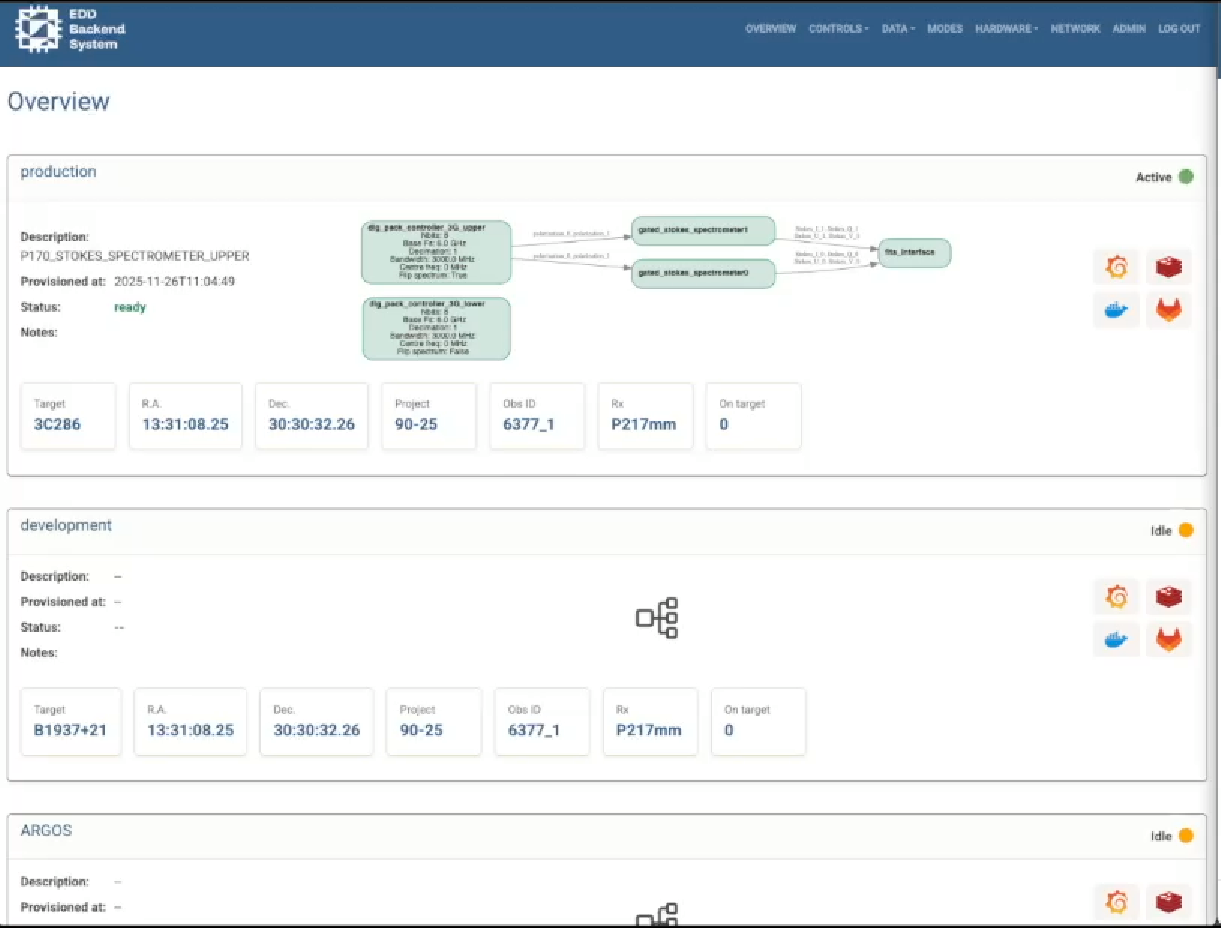
EDD Backend System
OVERVIEW
CONTROLS
DATA
MODES
HARDWARE
NETWORK
ADMIN
LOG OUT
Overview
production
Active
Description:
P170_STOKES_SPECTROMETER_UPPER
Provisioned at: 2025-11-26T11:04:49
Status: ready
Notes:
Target 3C286
R.A. 13:31:08.25
Dec. 30:30:32.26
Project 90-25
Obs ID 6377_1
Rx P217mm
On target 0
development
Idle
Description: –
Provisioned at: –
Status: --
Notes:
Target B1937+21
R.A. 13:31:08.25
Dec. 30:30:32.26
Project 90-25
Obs ID 6377_1
Rx P217mm
On target 0
ARGOS
Idle
Description: –
Provisioned at: –

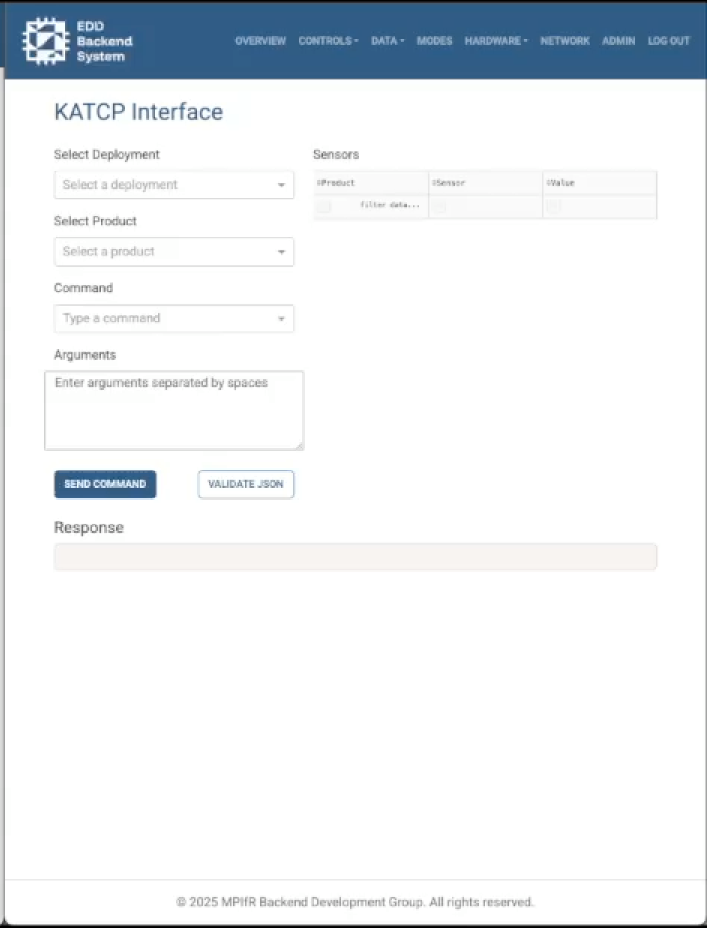
EDD Backend System
OVERVIEW
CONTROLS
DATA
MODES
HARDWARE
NETWORK
ADMIN
LOG OUT
KATCP Interface
Select Deployment
Select a deployment
Select Product
Select a product
Command
Type a command
Arguments
Enter arguments separated by spaces
SEND COMMAND
VALIDATE JSON
Response
Sensors
Product
Sensor
Value
filter data...
© 2025 MPIfR Backend Development Group. All rights reserved.

General use case of EDGAR

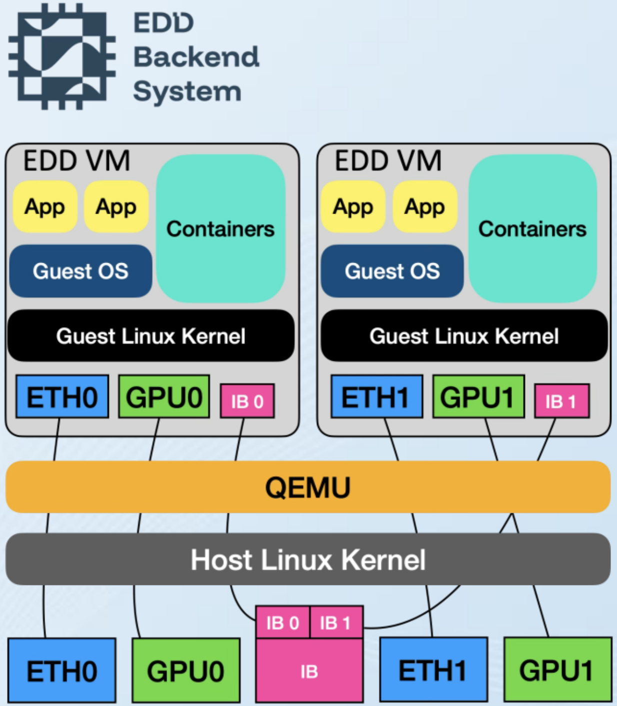

- One VM per NUMA node each with 100 GbE and a GPU
- PCIe passthrough for hardware components to VMs
- 24 cores, 256 GB RAM per VM
- RDMA enabled storage access over infiniband

EDD Mode

HTCondor Mode

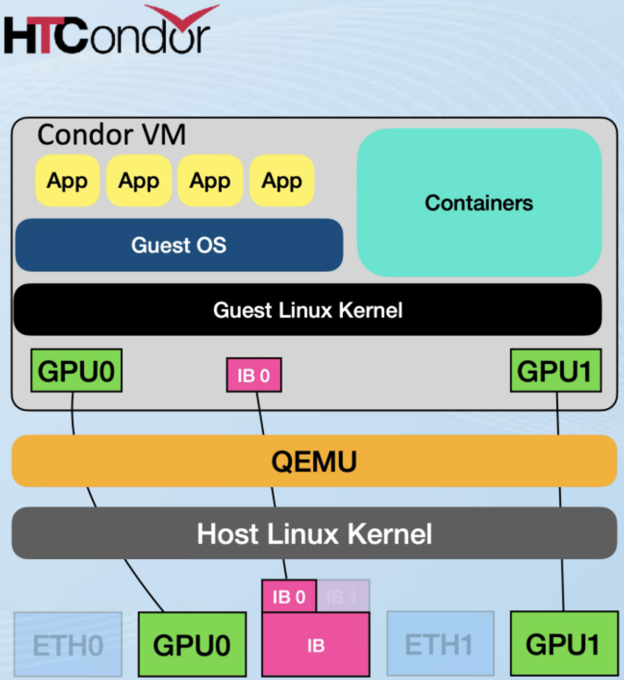

- One VM with two GPUs
- PCIe passthrough for hardware components to VM
- 60 cores, 600 GB RAM per VM
- RDMA enabled storage access over infiniband

EDGAR legacy RX usage

32 physical hosts

EDGAR next gen wideband RX usage

32 physical hosts

EDGAR CryoPAF usage

32 physical hosts

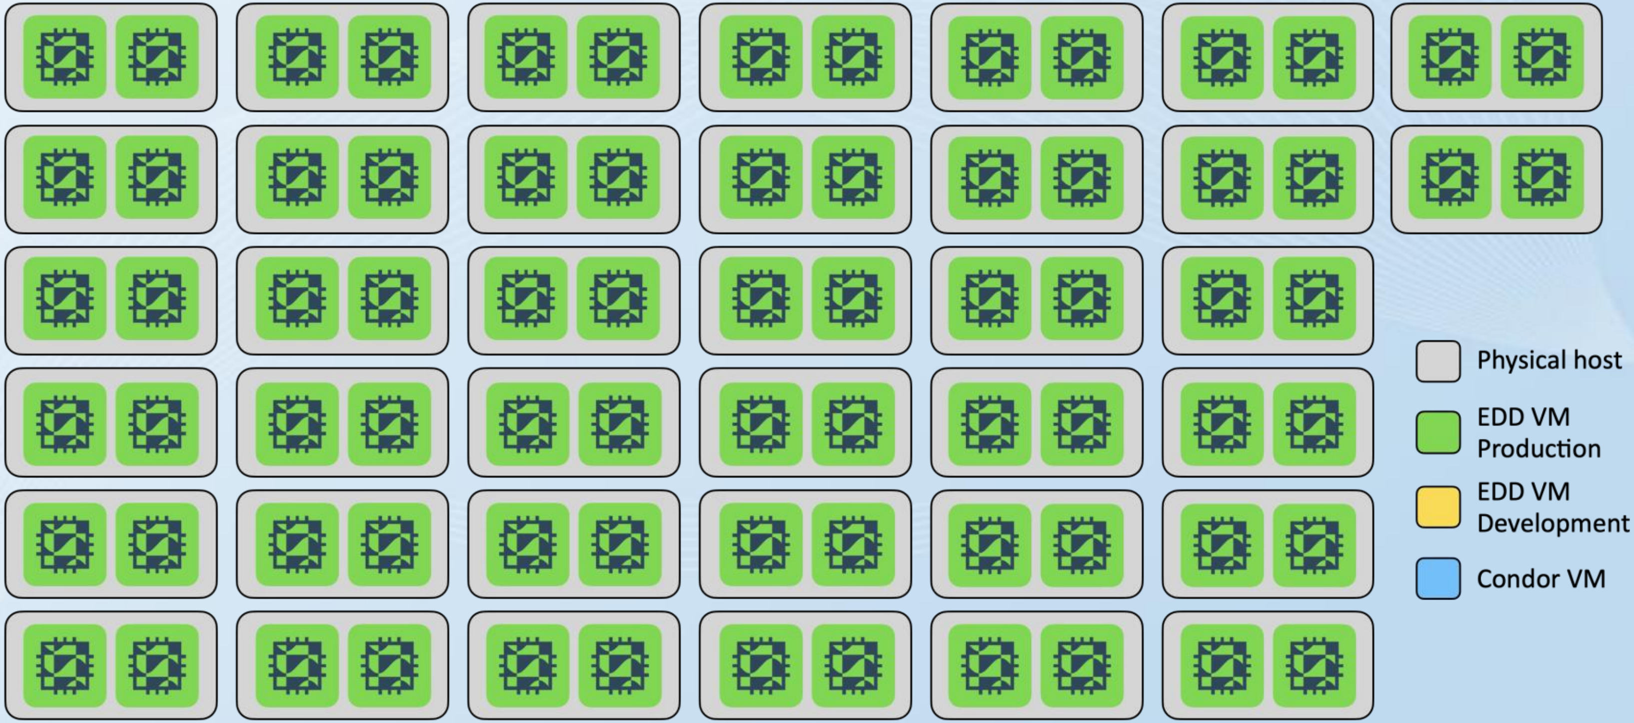

Conclusion

- Developments on:
  - Proven ability to handle mass(ive) construction
  - early digitisation (even in the frontend)
  - full control on the data is again available to the astronomer
  - massive scaling possibilities
  - Implementation of semi-industrial standards
  - Cost savings through the use of COTS hardware, e.g. for processing
  - Cm receiver technologies have evolved from a few tens of GHz into the hundred GHz (120 GHz is mature, while above is under evaluation).
- Real-time, truly coherent data processing
  - Provision of diverse data products and rapid implementation of new observing modes
  - Scalable and commensal observations
- Similar developments will also become available for mm and submm

Outlook

Riding the AI ~~bubble~~ wave?

What might the future ???

Array Scaling Challenges

	Driver	Challenge	Solution
Higher sensitivity	Probe higher distances and lower luminosity populations	Larger number of receiving elements, larger more complex corner turners	Faster networking standards allowing for smaller more manageable networks
Wider bandwidth	Higher sensitivity, more lines of interest	Higher required throughputs per element, larger more corner turners	
Larger fields-of-view + higher spatial resolution	Faster survey speeds, instantaneous sky coverage, lower confusion limits	Higher compute load from finer gridding and more beams	More and faster compute with high-speed interconnects
Higher time/frequency resolution	Short duration transients, pulsar timing, velocity resolution, interference excision	Increased network, I/O and compute requirements	

Network roadmap

- **Ethernet landscape as of 2025**
  - Current state-of-the-art: **800 GbE**
  - **1.6 TbE** standardization in progress (IEEE 802.3df)
  - OSFP transceivers: **17 W for 800 GbE**, projected **30 W for 1.6 TbE** → thermal limits for switches

- **Silicon photonics** maturing → lower power, CMOS-compatible
- Switch ASICs exceeding **100 Tb/s** (e.g., Broadcom Tomahawk 6, Nvidia Spectrum6)
- **Co-packaged optics (CPO)**: integrated transceivers near ASIC → lower latency/ power, not field-replaceable

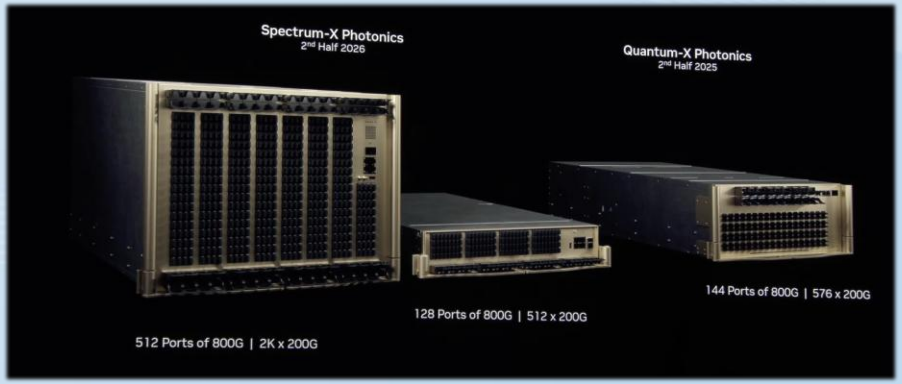

Compute roadmap

- GPU hardware tailored to inference problems good for correlation/beamforming problems
- Nvidia currently dominates the high-end GPU market (> 90% market share)
- Trend is towards high-density rack systems with integrated CPUs/GPUs which share memory

- **Nvidia roadmap:**
  - Blackwell (now) → Rubin (2026) → Rubin Ultra (2027) → Feynman (2028)
  - Expected x1.5 - x2.0 performance for each generation (less for Feynman)
- **AMD roadmap:**
  - MI350 (now) → MI400 (2026)→ MI500 (2027)
  - Similar performance uplift to Nvidia expected, but limit information available

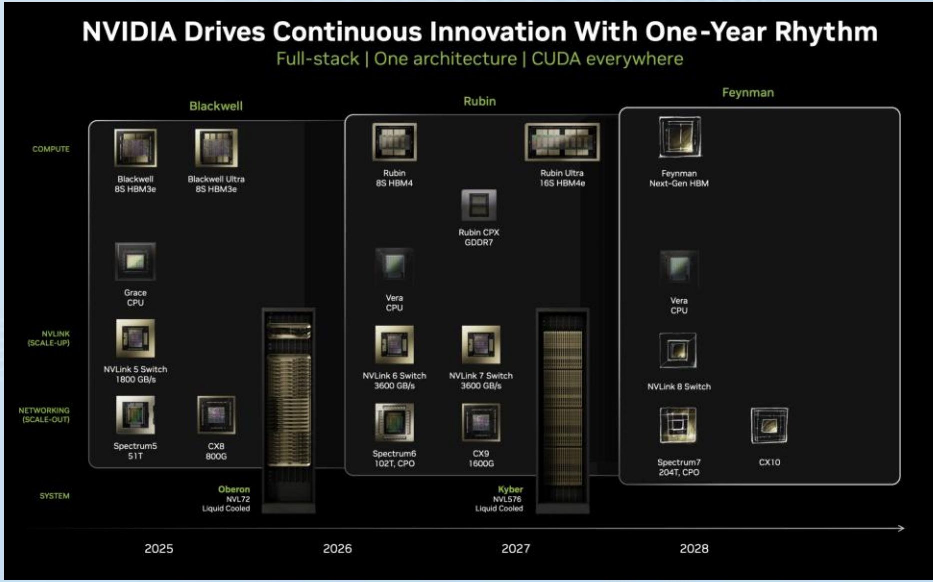

Compute roadmap

- **Nvidia GB200 NVL72 "Oberon"**:
  - 36 Grace CPUs and 72 Blackwell GPUs
  - 1.8 TB/s GPU-to-GPU NVLink backplane
  - Integrated liquid cooling
- **AMD "Helios"**
  - EPYC "Venice" CPUs and 72 MI450 GPUs
  - 800 GB/s GPU-to-GPU Infinity Fabric (IF) backplane
  - Integrated liquid cooling

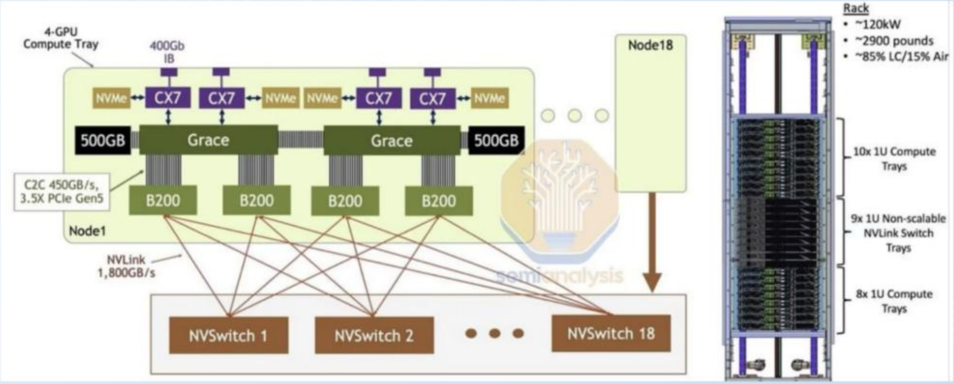

- Nvidia roadmap for 2026-27 includes the **576 GPU, 600 kW** rack system **"Kyber"** which will use the Rubin Ultra architecture GPUs and require a **800 V DC** or high-voltage AC power distribution system.

Context

- A single GB200 NVL72 rack has:
  - **15-20x** the compute capacity of the current **MeerKAT beamformer**
  - the ability to ingress **7x** the raw **MeerKAT data rate** (14+ Tb/s)
- Scaling to **SKA-Mid-like telescope** with dual-polarisation feeds and 1 GHz of instantaneous bandwidth (at 8-bit resolution), a single rack could ingress **200-400 antennas**
- System would require **1-2x 64-port 800 GbE** switches for corner turners (accounting for logical port allocations and assuming 100 GbE to each antenna)
- How big can we go?

Example 2026 Array

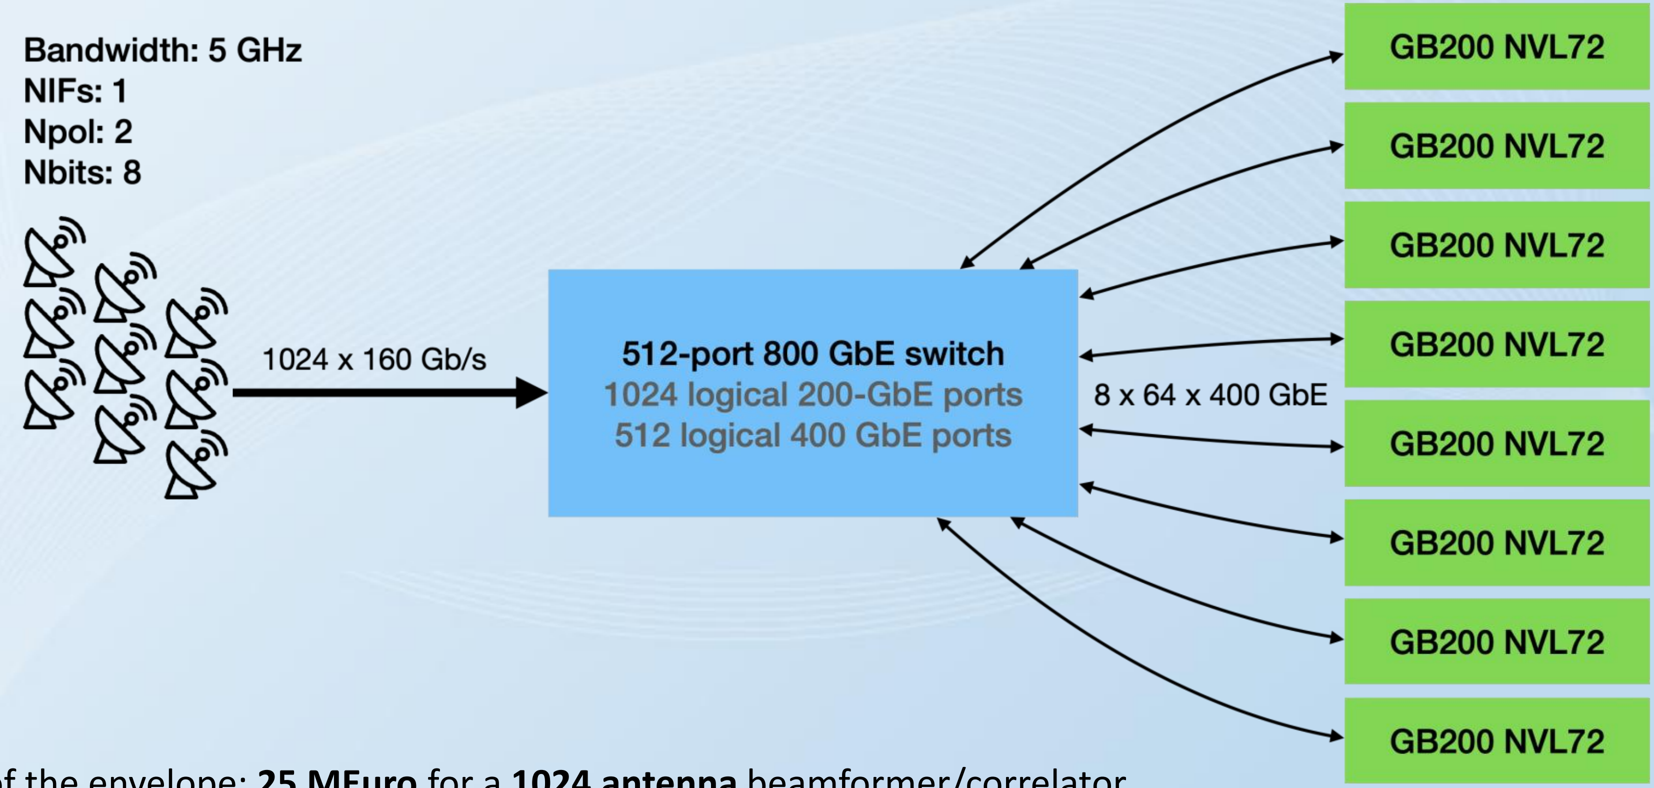

Back of the envelope: **25 MEuro** for a **1024 antenna** beamformer/correlator

Thank you for your attention.

PAF 2026 - International Phased Array Feed (PAF) & Advanced Receiver Workshop
June 1-5, 2026 @ Guiyang
https://conf.koushare.com/conf/PAF2026/