A long term storage solution for Tango attribute data at SKAO

Mauricio Zambrano, Thomas Juerges, Johan Venter

ICALEPCS 2025,

SKA Telescope - Overview

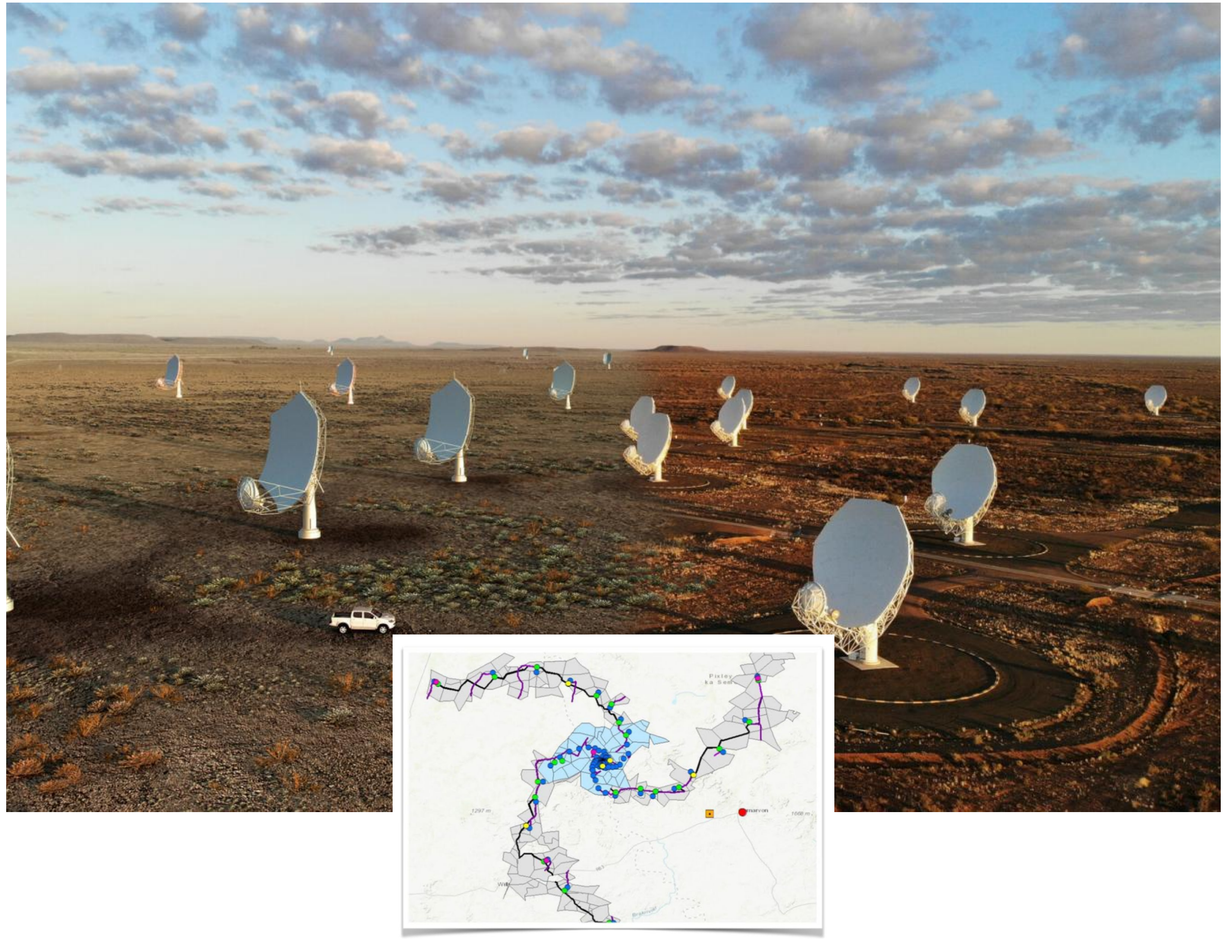

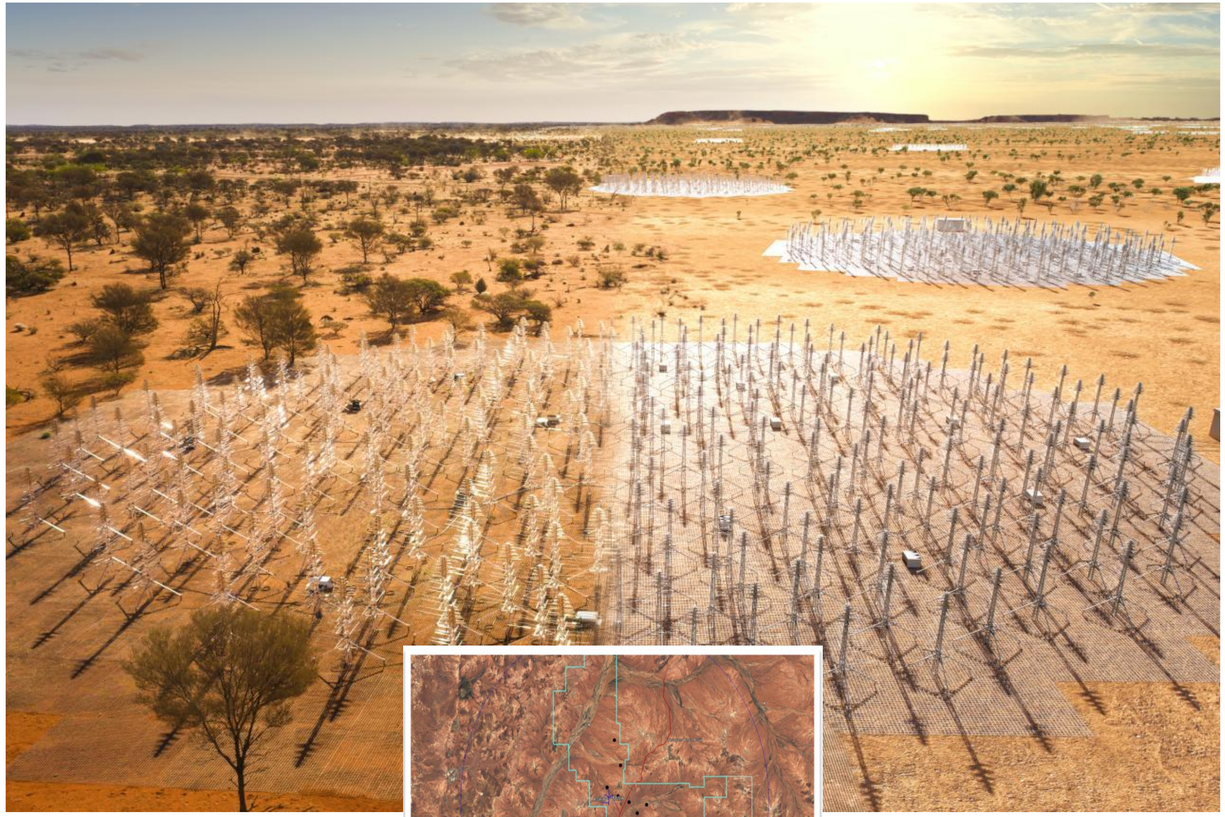

197 fully steerable dishes
Maximum baseline 150 km
Frequency range 350 MHz - 15.4 GHz

512 stations
Maximum baseline 70 km
Frequency range 50 MHz - 350 MHz

One Observatory Two Telescopes Three Continents

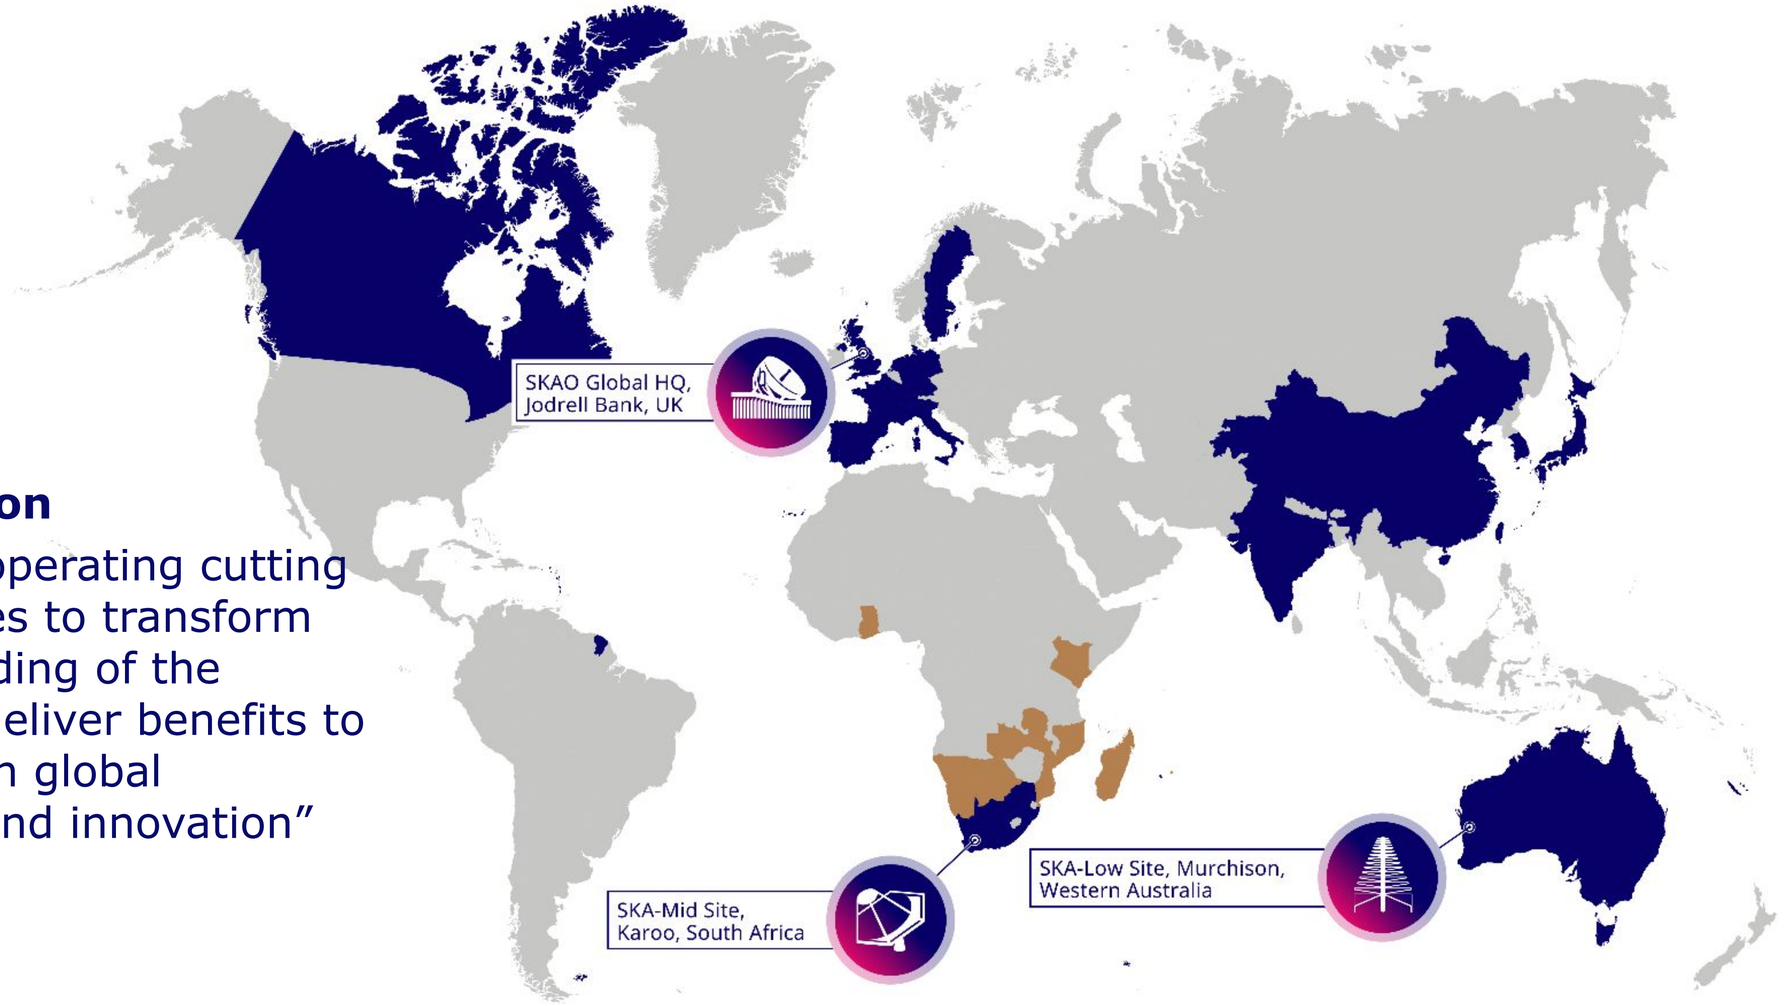

SKAO's Mission

"To build and operating cutting edge telescopes to transform our understanding of the universe and deliver benefits to society through global collaboration and innovation"

Australia
Canada
China
Germany
India
Italy
The Netherlands
Portugal
South Africa
Spain
Sweden
Switzerland
United Kingdom

SKAO Partnership - includes SKAO Member States* and SKAO Observers (as of August 2025)

African Partner Countries

Modern telescopes and accelerators

Complex distributed systems

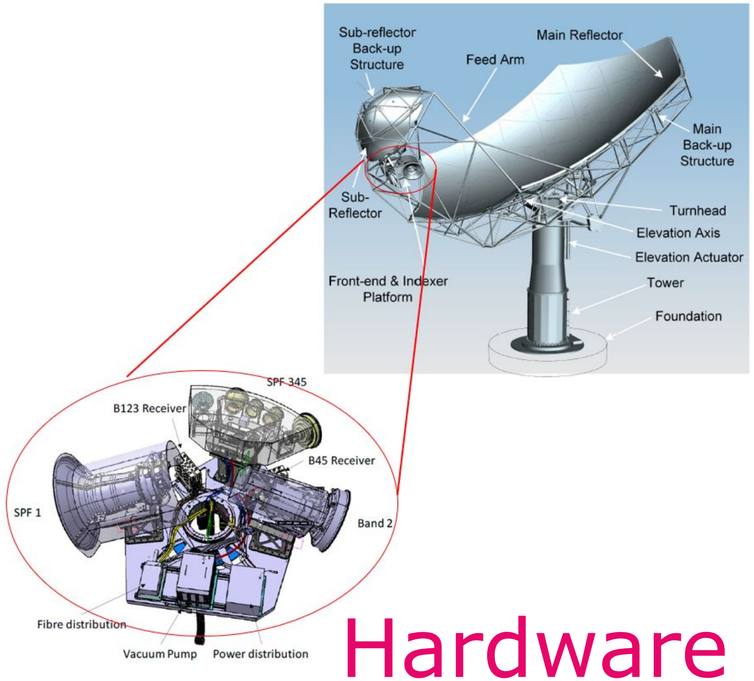

Hardware

Control System

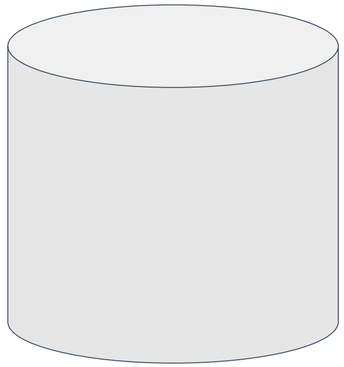

Engineering Data Archive

Image credit: Alice Pellegrini, SKAO

Functions of the Engineering Data Archive (EDA)

- **Operational Reliability & Troubleshooting**
- **Commissioning & Calibration**
- **Performance Optimisation**
- **Applications:**
  - **Root cause analysis**
  - **Data driven maintenance**

The MeerKAT Telescope

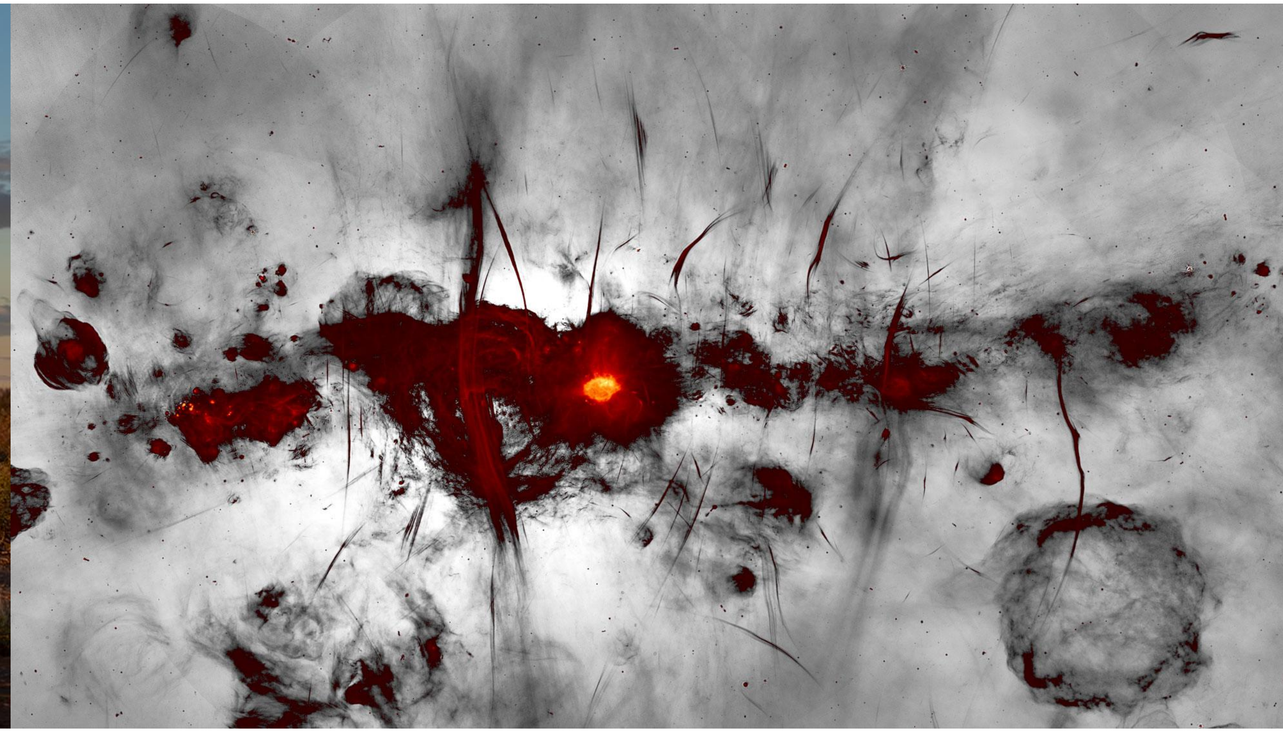

- 64 antennas
- ~6000 samples/s
- KATCP devices

Image credit: Alice Pellegrini, SKAO

The MeerKAT EDA

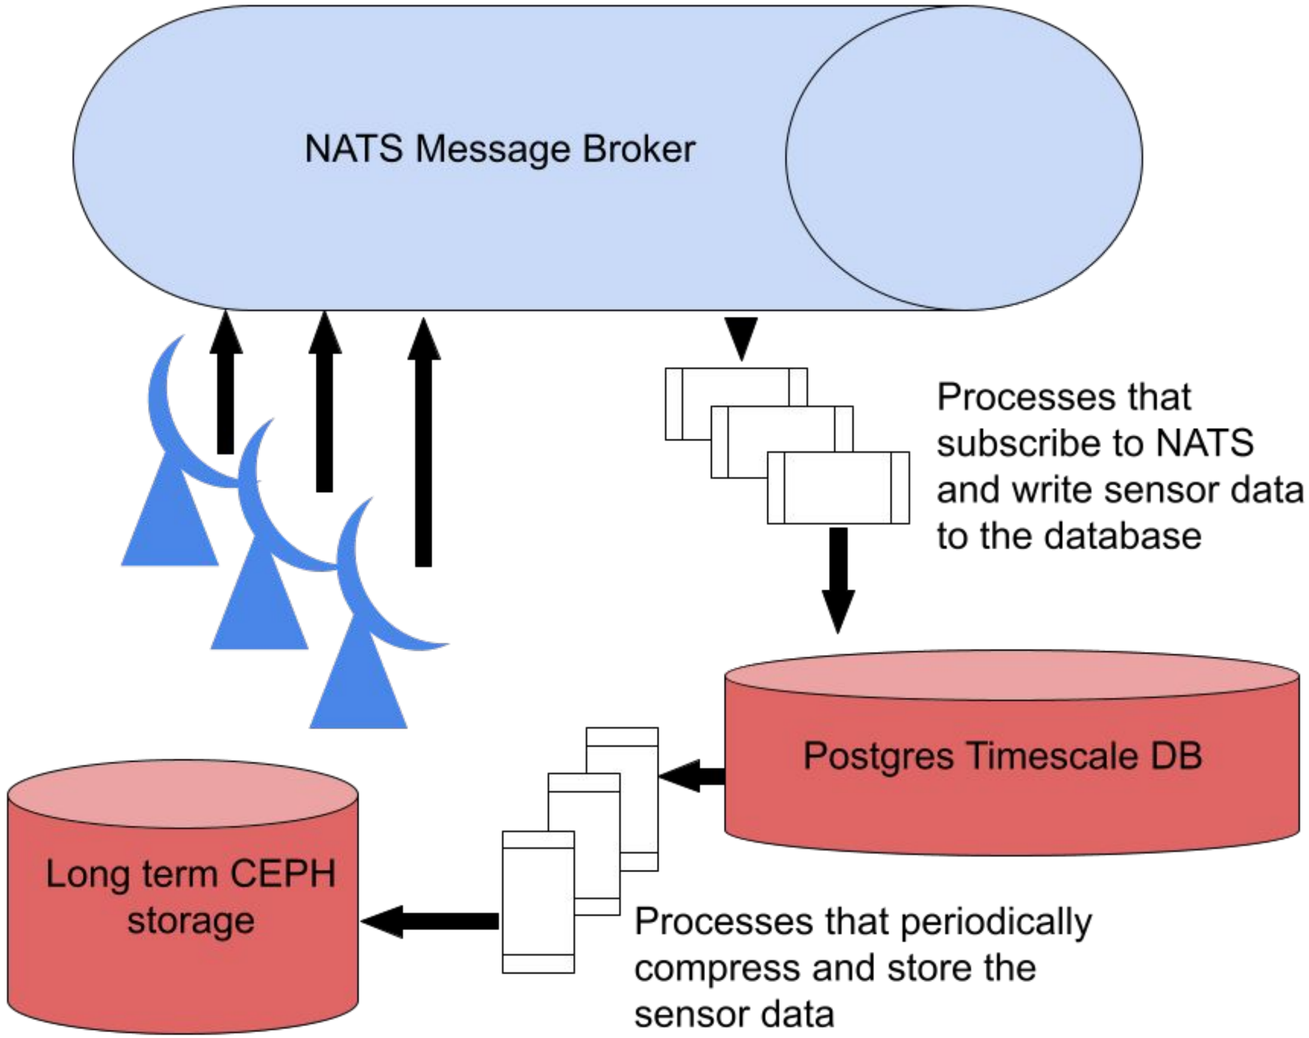

- NATS
- PG/TimescaleDB
- CEPH block storage ~ 40TB

The SKAO - current situation

- Construction activities at sites
- Teams are testing, validating, calibrating and commissioning

SKAO Computing environment

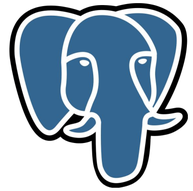

- Software runs on Kubernetes, StackGres for
- The control system is Tango:
  - Archiving system: HDB++

HDB++ components:

- Configuration manager
- Event subscriber
- A PG/TimescaleDB inserter implementation
- A database

Ceph is our main storage. Keep the data for 50 years

TimescaleDB, a time series extension for PostgreSQL

- Partitioning: chunk, time period
- Reordering: relevant data contiguous - chunk re-rewrite
- Compression: compress chunks after a period of time
- Aggregation: provide average of period
- Decimation: delete data older than x days

You need to tune your TimescaleDB parameters!

- Chunk size
- Worker jobs
- WAL size, background jobs, temp file space

One environment: MCCS@LOW

- Four stations, 256 antennas per station
- 16.031 monitoring points, configured by AIV
- ~600GB of data per month, started on 2024

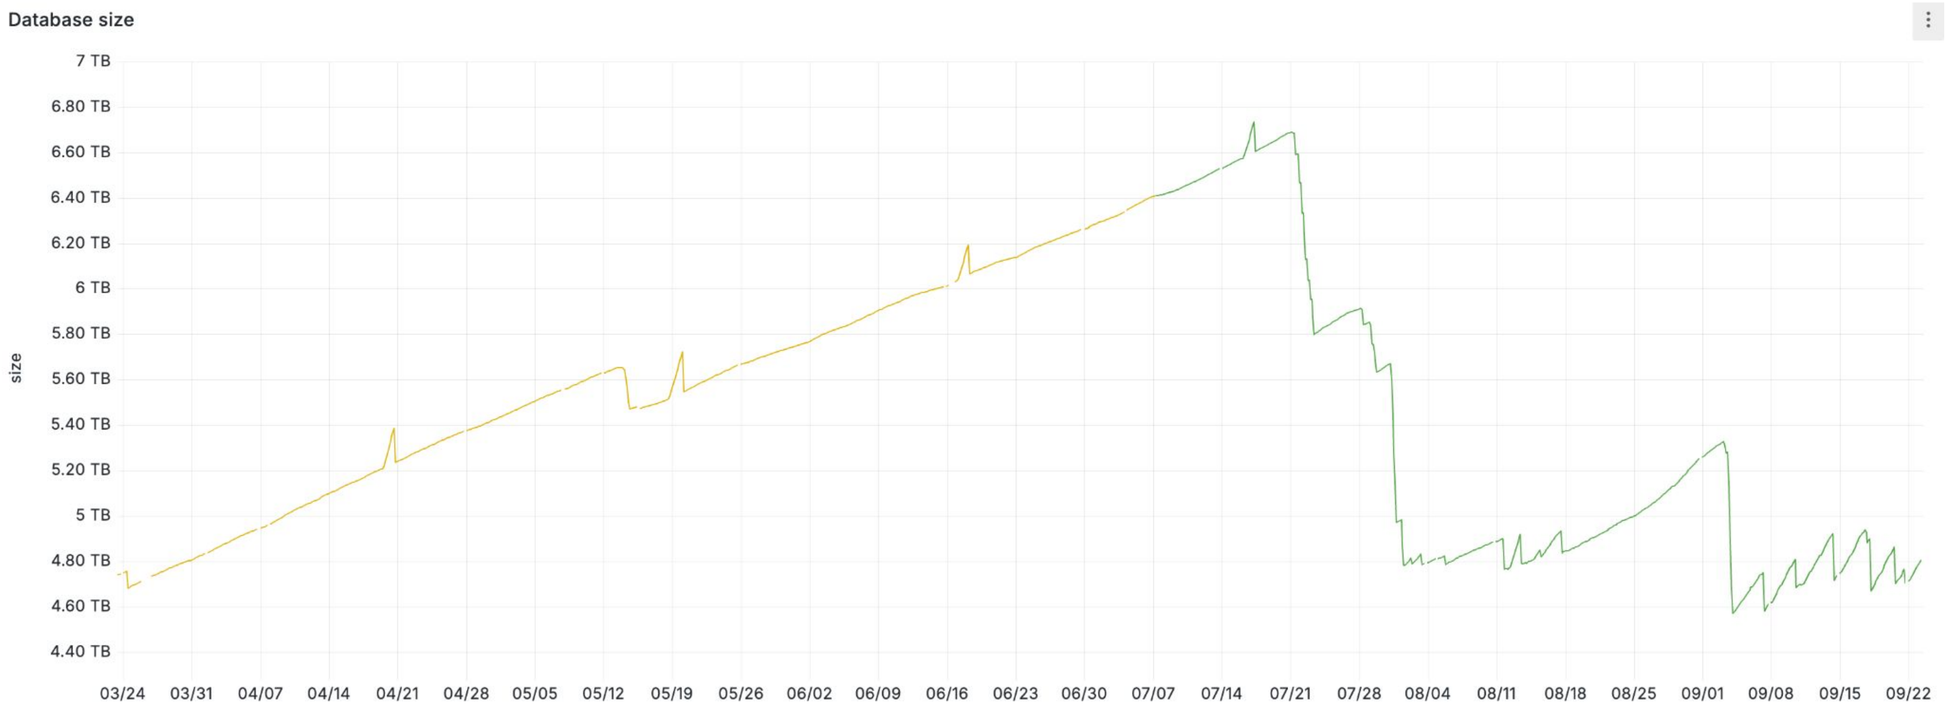

A Data Life Cycle Management solution for the SKAO EDA

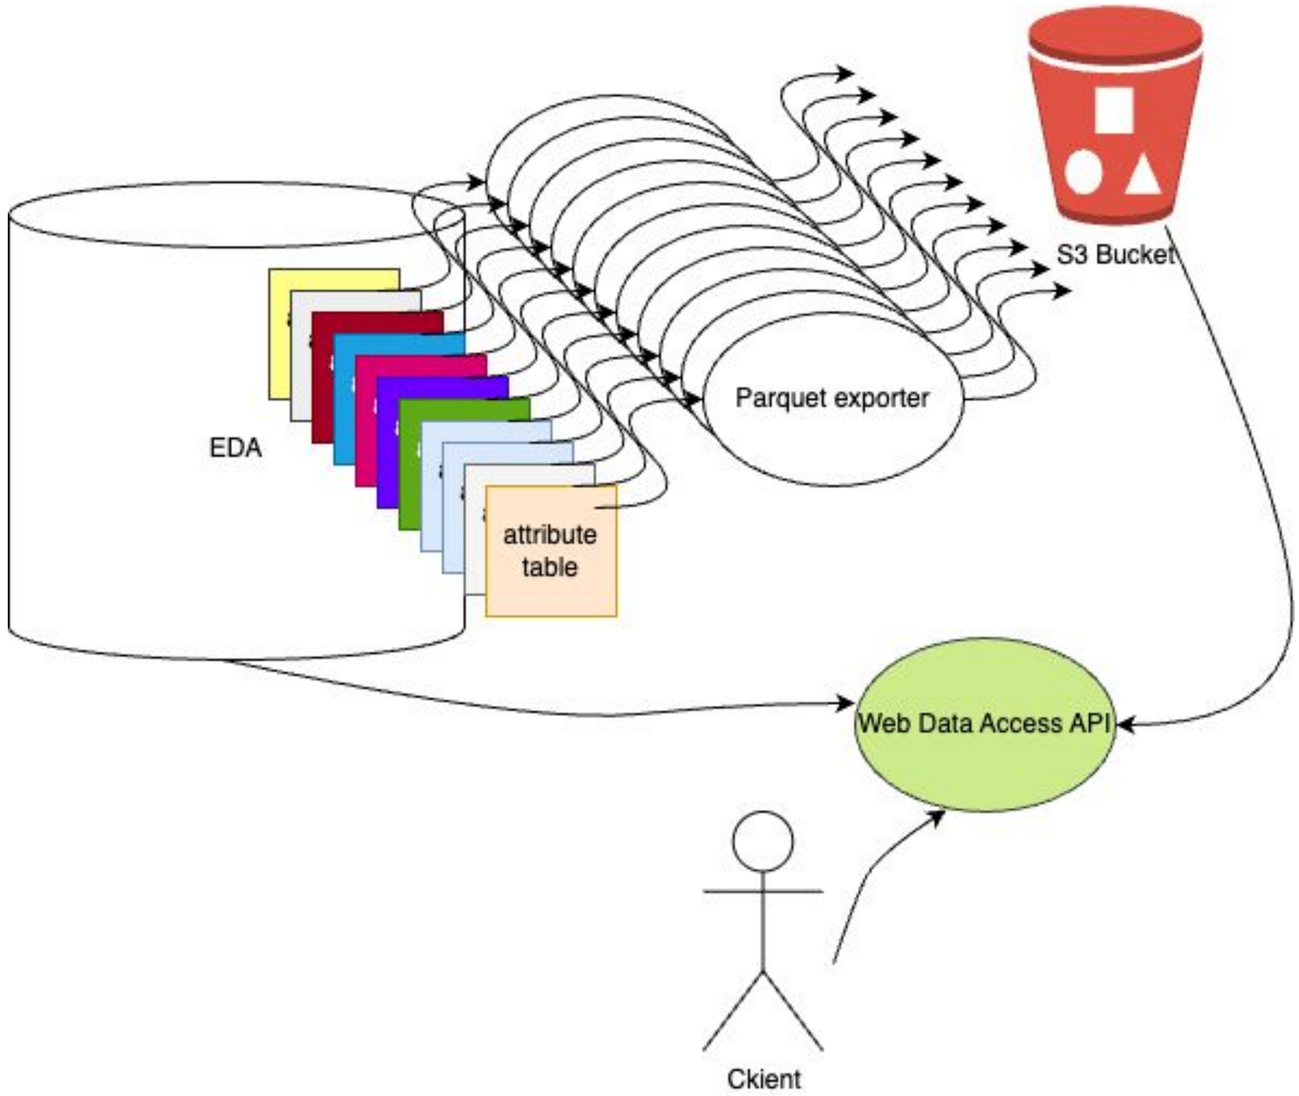

- Backups on S3 as parquet files
- PG2PQ2S3
- One job per table
- Execute it periodically
- Backups are queriable:
  - API
  - Duckdb, Clickhouse, PG

Future work

- Deploy the solution, test its performance, test with end users, create data product from parquet files.
- Provide a Tango based library for parquet files
- Monitor ingestion rates
- Enable aggregation
  - Can we use it?
- Discussion on decimation
  - If full data is on backups then it should be fine

Conclusions

There is no DLM solution for PostgreSQL - BYODLM

Parquet seems appropriate

TimescaleDB provides valuable tools to tame growth

Moving data from the database to a storage system seems the way to go

Questions?

We recognise and acknowledge the Indigenous peoples and cultures that have traditionally lived on the lands on which our facilities are located.

SKAO

NRF National Research Foundation | SARAO South African Radio Astronomy Observatory

www.skao.int