National Research Foundation | SARAO | South African Radio Astronomy Observatory

South African SRC Landscape

PRESENTER: Simon Ratcliffe

MeerKAT – 64 antennas

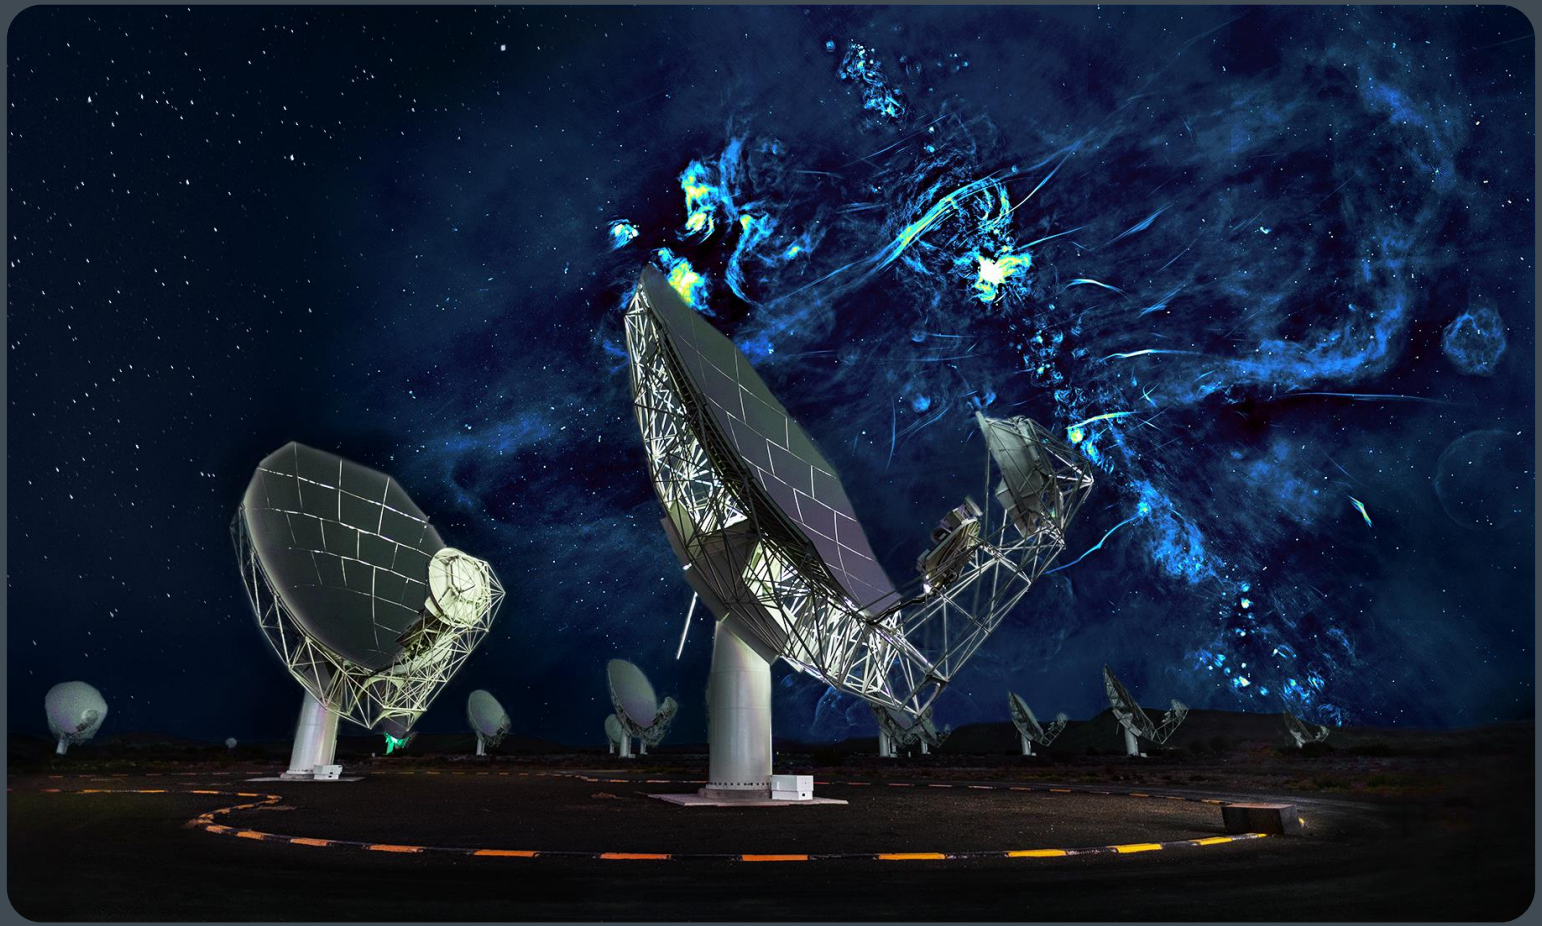

https://www.nature.com/articles/s41586-019-1532-5

Current MeerKAT Data Products

Calibrated, flagged, full time and spectral resolution visibility data (medium term storage) + 10x reduced product (indefinite)

Full res is 0.5 Hz / 32,768 channels + per vis flags and weights (typically 0.125Hz / 16,384 channels)

MVF4 (MeerKAT Visibility Format v4) native with MSv2 export. Combines SDM and object visibility storage. Will track MSv3.

Calibration Tables including bandpass, gain, delay and cross-pol terms.

Raw capture of F- and B-engine data

Continuum image pipeline: baseline subtraction, self-cal solutions, and best effort images.

Spectral line pipeline: image cubes (up to 100 hrs joint)

Deployed Infrastructure

Realtime Mesos Cluster : Ingest + Cal

100 TFLOPs, 6TB RAM

Batch Cluster : Spectral + Continuum Imaging

1.5 PFLOPs, 4TB RAM, 1 PB buffer (CEPH)

Object Storage : Vis Data + Science Products

cluster1 (CEPH): 5.4 PiB – decommisioning
cluster3 (CEPH): 12.2 PiB – production
tape: 20.3 PiB – cold visibility

Pipeline Status

Calibration pipeline :

deployed: B, G, K, K-cross, 2D Flagging, cal report

development: further quality metrics, pol cal

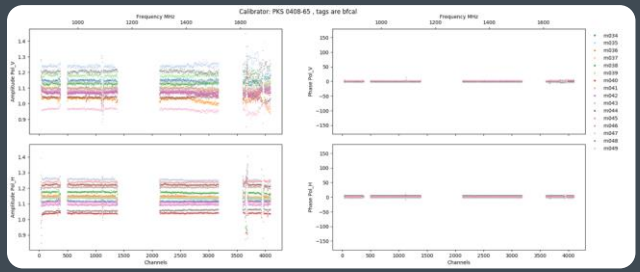

Imaging pipelines (primarily for Quality Assessment):

Total execution time ~ observation time

Continuum:

- Based on Obit
- SPEAD stream to AIPS disk (BDA in use)
- Primary products are clean components and self-cal solutions, but images also stored.

Spectral:

- Built from scratch – mostly to optimize efficiency
- Hybrid w-projection and w-stacking
- Compute efficiency above 50%

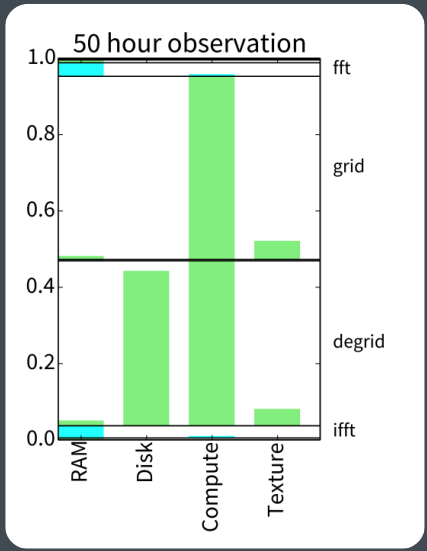

Pipeline Data Flow

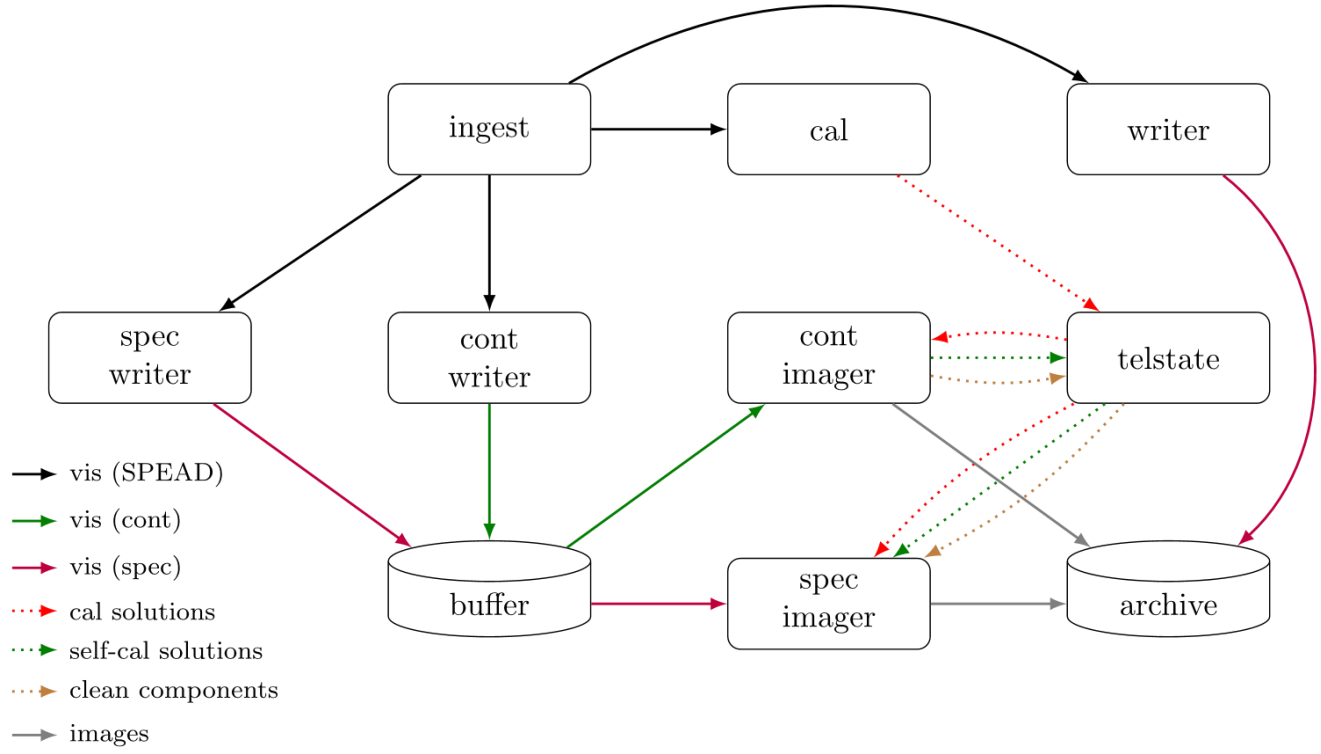

Hands-free Imaging (M83)

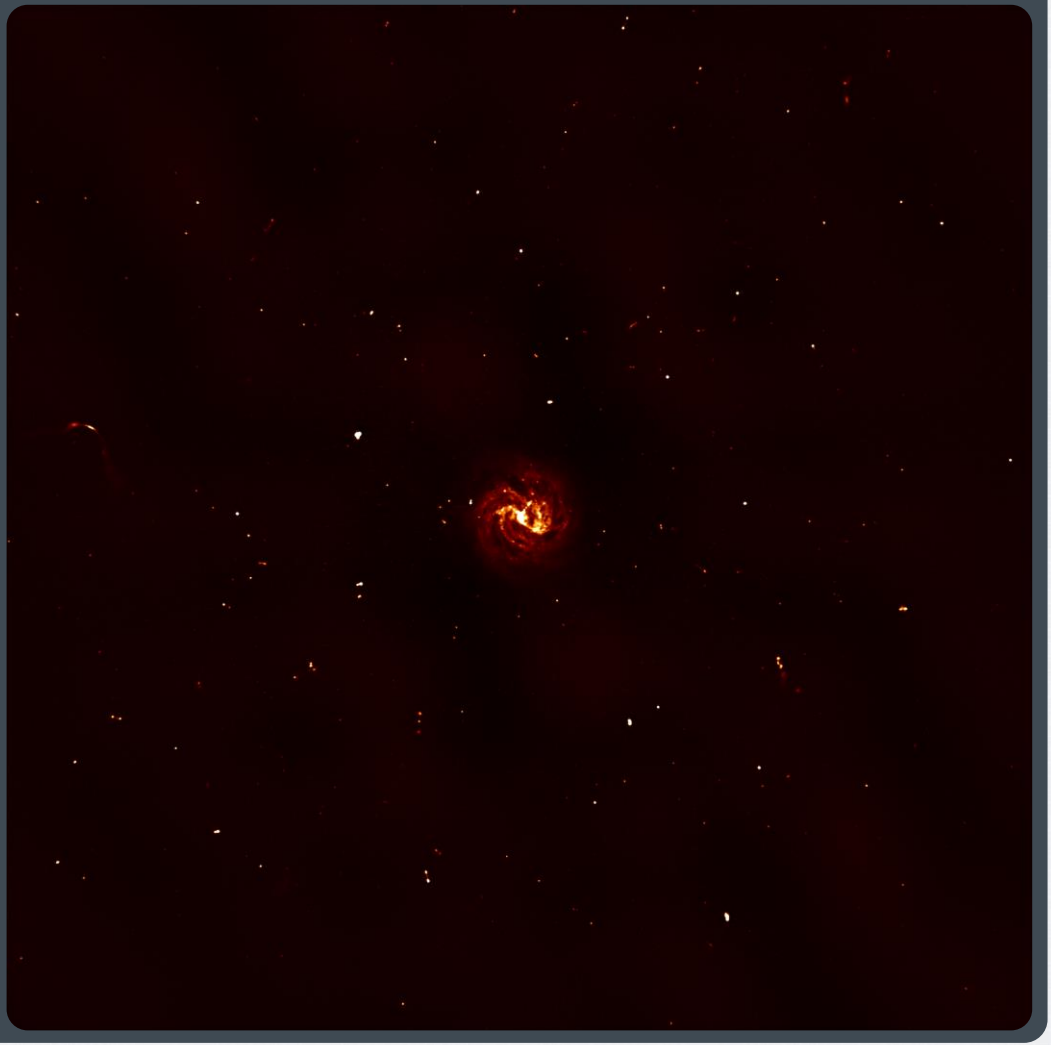

Hands-free Imaging (M83)

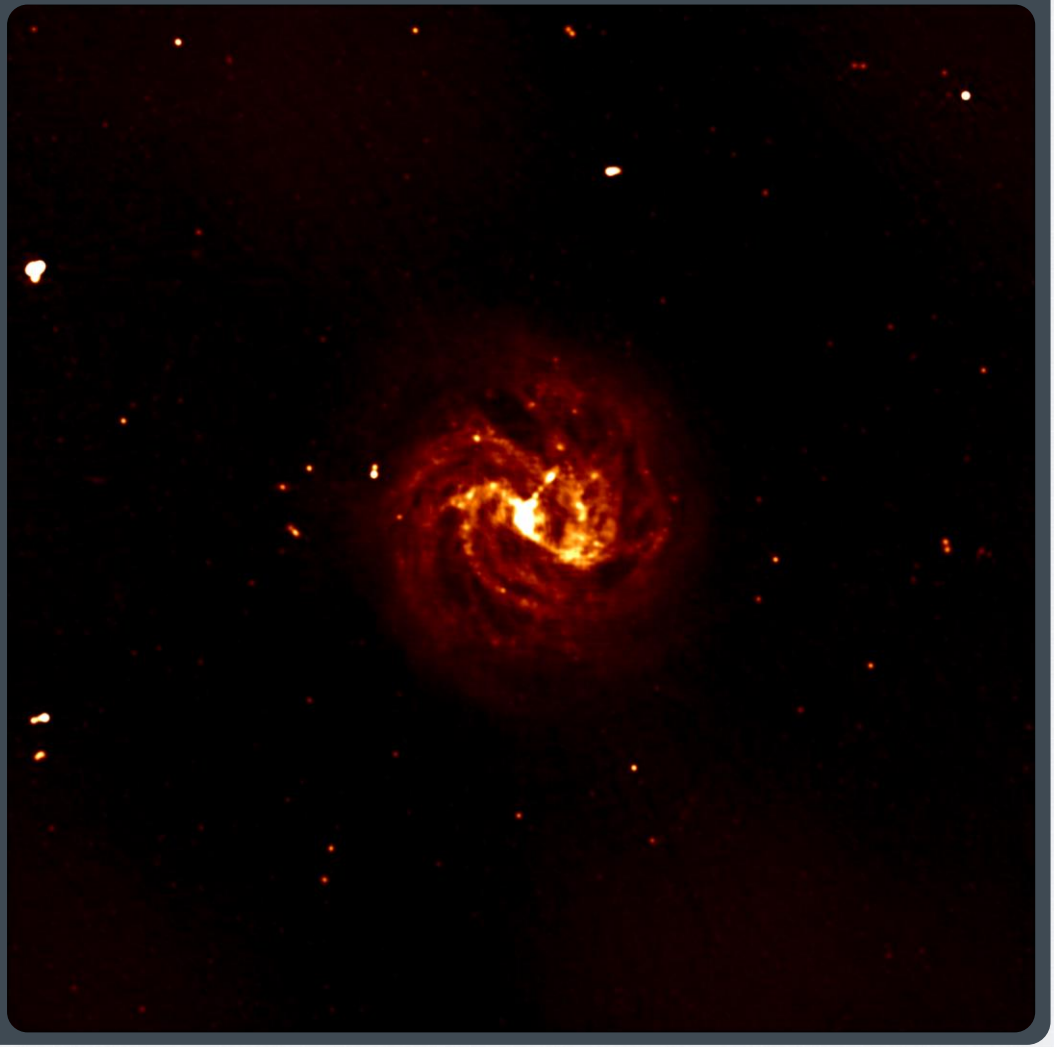

J0358-8103

Single 4k Channel

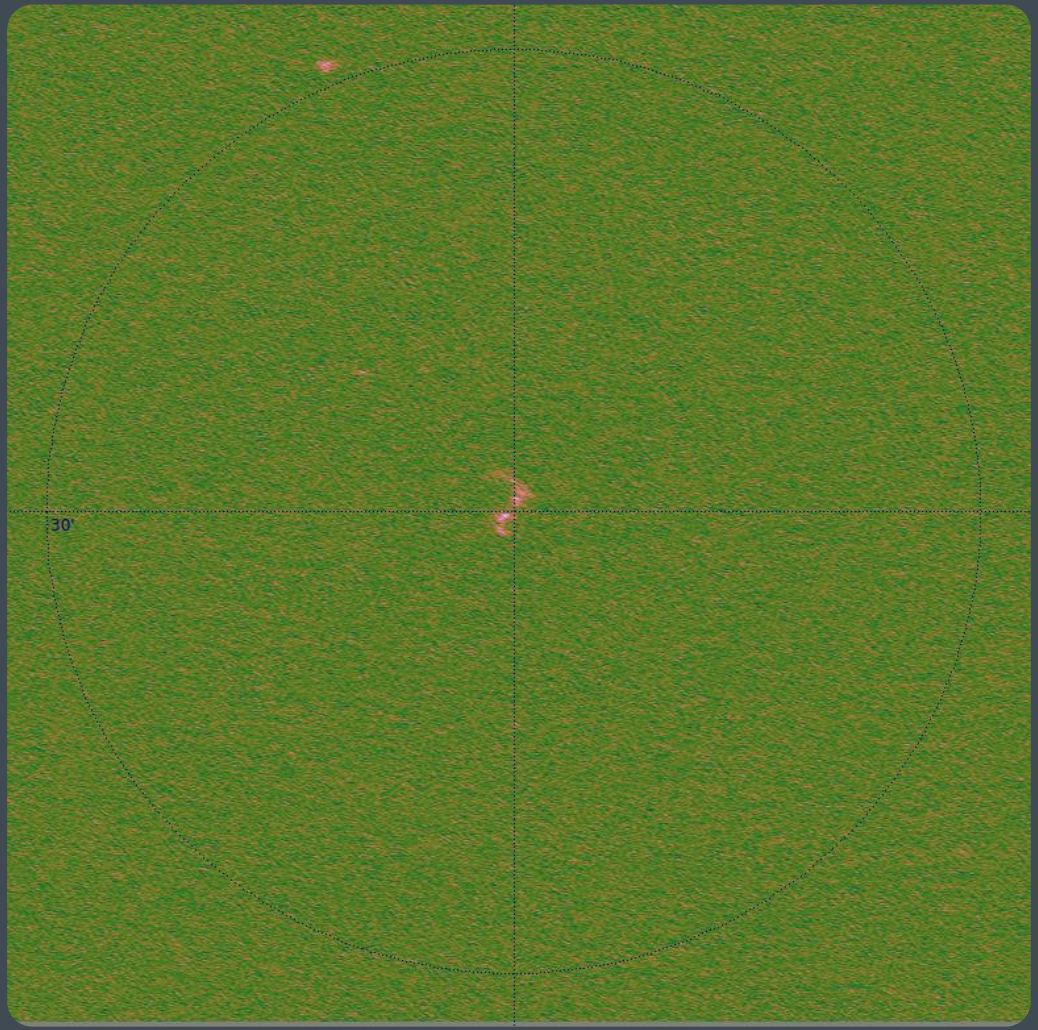

…but what do the users actually want ?

Decreasing Trust

- Raw visibility data
- Flags (1D & 2D)
- Base Cal Solutions (B,G,K, etc..)
- Calibrated Visibility Data
- Self-Cal Solutions
- Image Cubes (QA)
- Image Cubes (Science
- Catalogues

MeerKAT Q4 2019

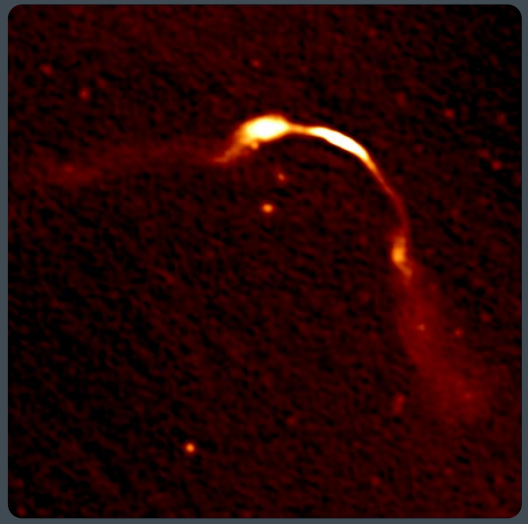

Current line of attack

Multiple pipelines deployed across the community:

SARAO: Internal QA pipelines with focus on wallclock and compute efficiency.

RARG: MeerKATHI and others – prime focus on correctness and new algorithm development.

IDIA: Hybrid of the two – needs to service the end users effectively but also start to push the science boundaries.

Trust takes time, but we see more and more convergence between these three, which is ultimately the path to acceptance.

Data Flow and Tiering

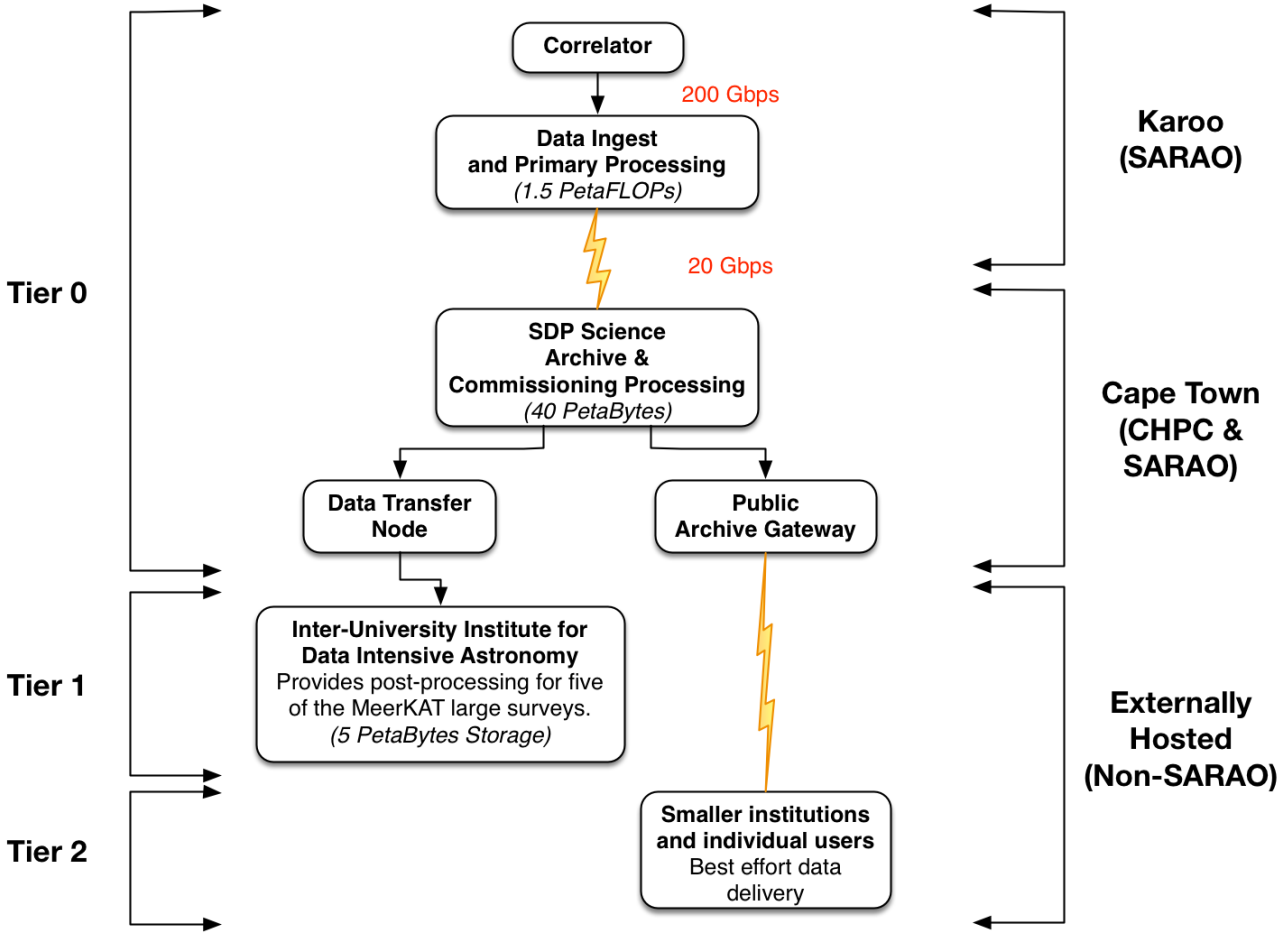

Ilifu Update

Ilifu = an Openstack cloud computing environment

- Most users interact with it via SLURM and/or Jupyter Notebooks
- Over 220 users so far
- Supporting a wide array of astronomy + bioinformatics projects
- Also used for various training workshops

Expecting a small number of new storage nodes within a month

- Provide some “breathing room”
- Enable gradual rollout of a more robust ceph (storage) configuration

Prepping for a much larger storage expansion beginning Feb (approx.)

Moving towards a federated Openstack environment

- EGI Check-in working testbed <-> testbed

CARTA Collaboration

Cube Analysis and Rendering Tool for Astronomy

IDIA (South Africa) – NRAO (US) – ASIAA (Taiwan) – U of A (Canada)

- Cloud-based Visual analytic of remote large image cubes
- Beta release v 1.2, 15 August 2019
- To replace CASA viewer and deploy at ALMA Regional Science Centres

CARTAvis

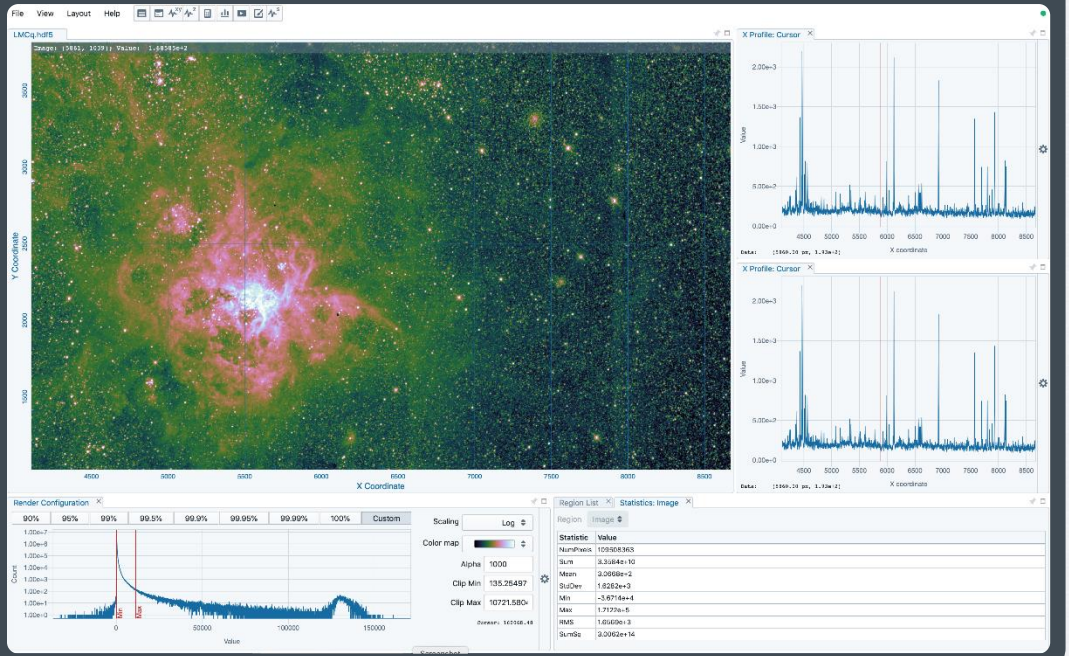

IDIA MeerKAT Pipeline

Brad Frank, Jordan Collier, Srikrishna Sekhar, Russ Taylor

- V1.0 (released March 2019) under performance testing by LSPs

Full Stokes calibration in CASA

- Continuum images + polarisation cubes + spectral line cubes

Parallelised package for HPC processing (SLURM + cluster)

- Uses multi-measurement sets (MMS) to parallelise across a cluster

Robust, generic, fast implementation of a'priori calibration

Easy to use, transparent, reproducible

Builds and submits pipeline jobs to SLURM

- Input measurement set, build / run your config file, request resources
- Optionally insert your own scripts, specify containers and MPI wrappers

Aim: T(cal) ~ T(obs)

MeerKAT Extension

Adds another 20 antennas to the existing 64

New dishes to be built on future SKA locations

Includes L- and S-band feeds (not UHF)

Joint venture between SARAO, MPG and China

Biggest impact is on SDP
(8km -> 18km max baseline)

MeerKAT Extension Timeline

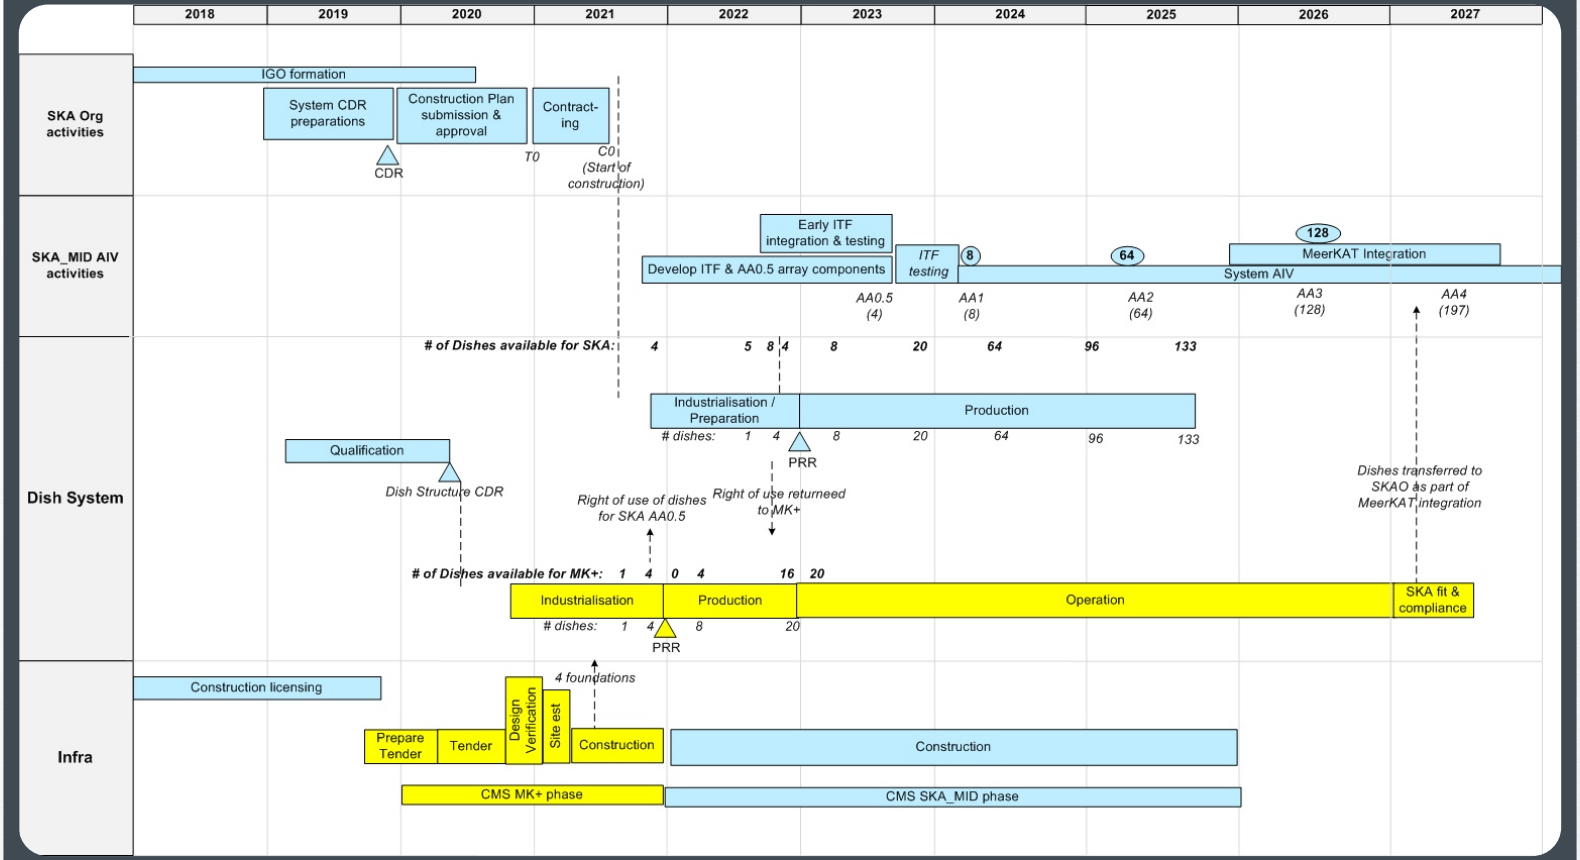

4 year “filling the gap”

Dish Construction Plan

1) Industrialisation Phase
- 4 dishes
- Construction Jun'21 - Feb'22
- Funded by China & SARAO
- Transferred to SKAO
- Used for AA0.5 (EPA)

2) MeerKAT Extension Production
- 16 Dishes
- Construction: Feb'22 - Dec'22
- Funded by MPG & SARAO
- Used for MeerKAT Extension Operation

3) SKA Production
- 4 Dishes
- Construction period: TBD
- Funded by SKA Observatory
- Transferred to MeerKAT Extension
- Used for MeerKAT Extension Operation

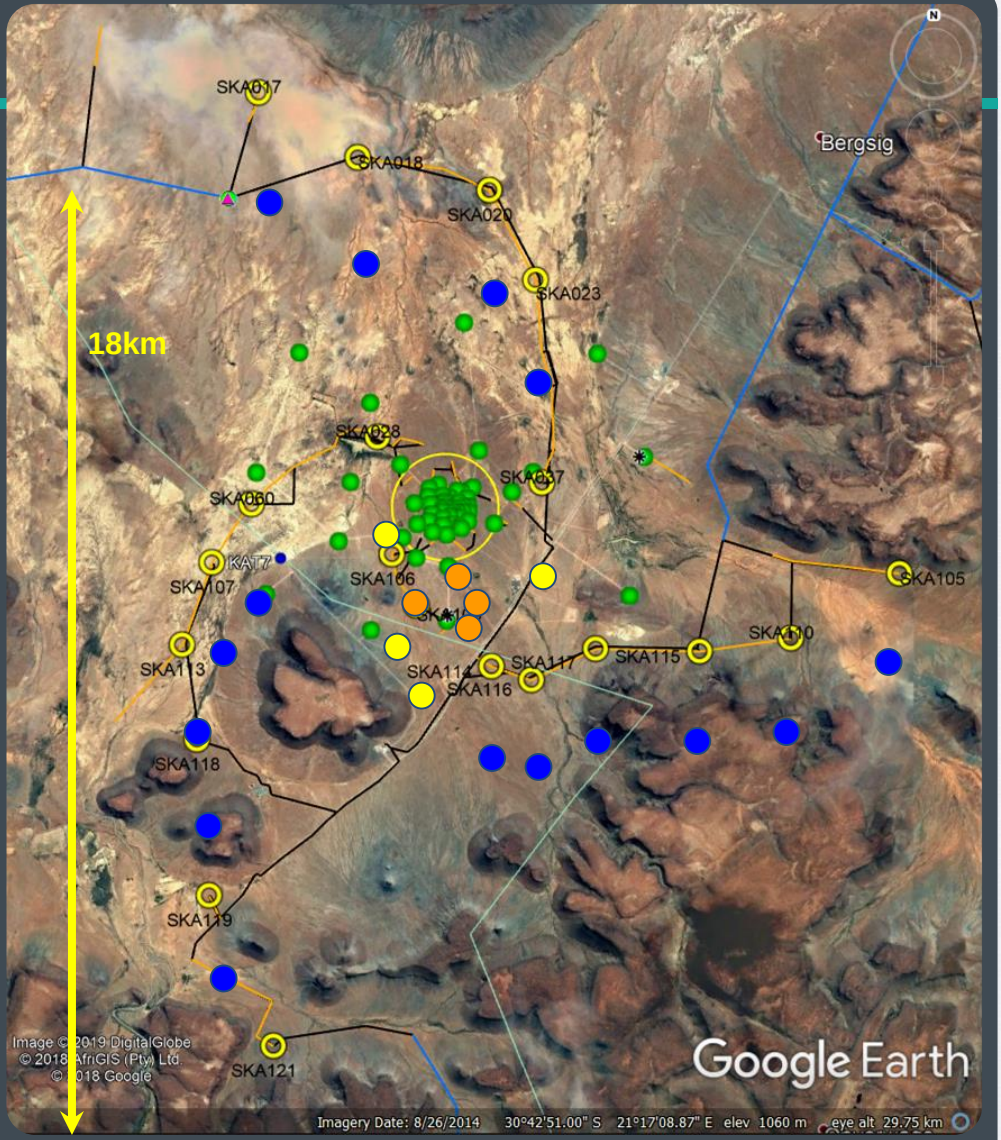

MeerKAT Extension SDP Scaling

Data Rates

84 antenna (CBF dependent) => 1.7 x baselines

Max baseline (8 -> 18) => integration time reduced to as little as 0.4s (less core dense, 2.25 baseline length increase) => 5 x data rate

Peak case is thus as much a 8x higher than MKAT

Data Volumes

Visibility data => Up to 6x higher (assuming channel count remains)

Imaging data => 5x higher for full resolution cube

50% duty cycle assumed => 2x reduction

Reduced lifespan compared to MKAT => 2x reduction

Overall assume a doubling of current archive

MeerKAT Extension SDP Scaling

Data Ingest and Calibration

At least 4x ingest servers to handle peak

CBF bandwidth (16 x 40GbE) is sufficient

Calibration servers will likely need 8x increase in memory

Data Processing

Imaging challenge nearly cubic in baseline terms => 10x increase over MKAT

Mixed array a substantial compute and scientific challenge (as yet this is an unsolved)

Overall we assume a 7x increase in compute to around 10 PFLOPs

Useful scale comparison to SKA

Science Processing Center

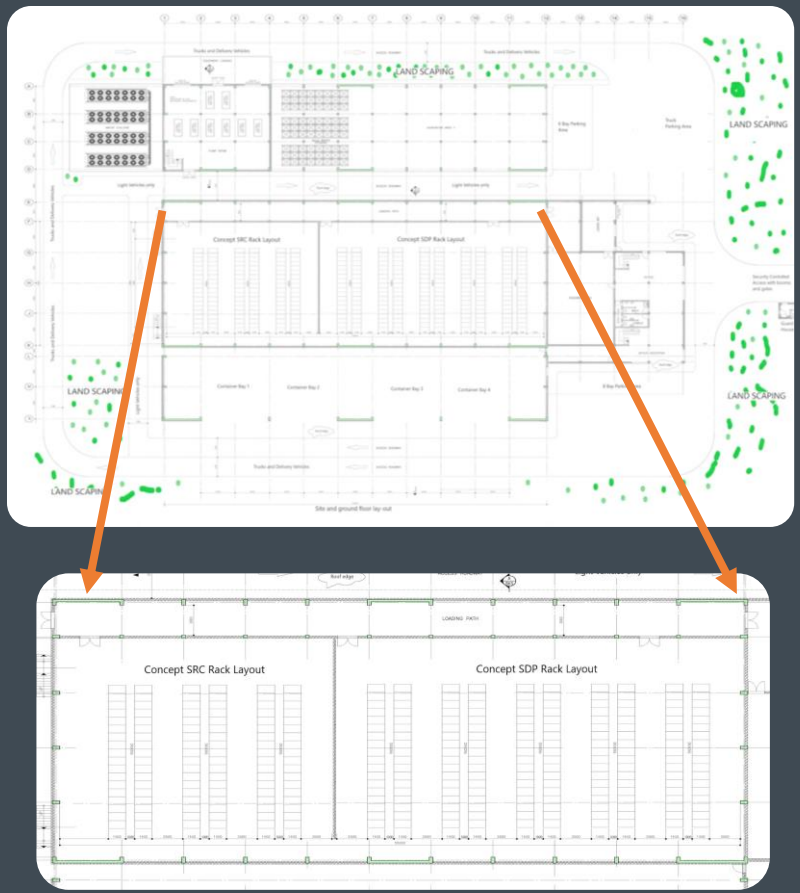

Science Processing Center (SPC) detailed planning underway.

South African SRC is back to back with MID SDP within this building.

Potential for higher data rates (Tbps) to the ZA-SRC, particularly for visibility storage.

Funding request for prototype phase
(-2023) in preparation (~ $4m)

Planning activities now include using the MeerKAT extension for prototype work.

Possible ZA-SRC Model

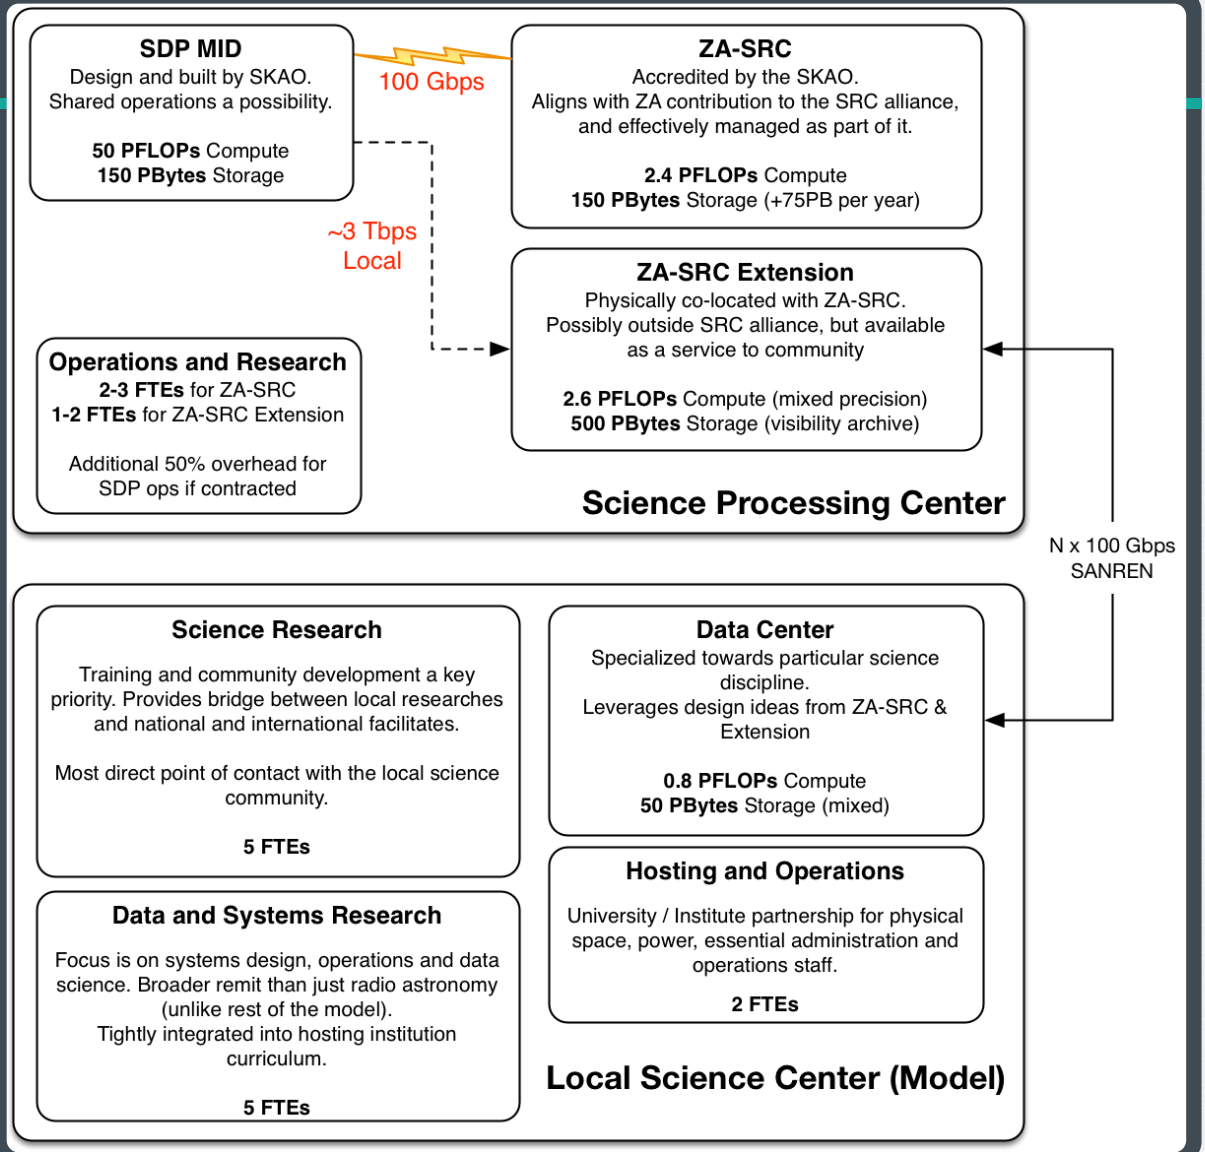

Onwards…