ASTRON

Netherlands Institute for Radio Astronomy

European SKA Regional Science Centre
Concept and Planning

Michael Wise
Head, ASTRON Astronomy Group

SKA Workshop on SDP&HPC
Shanghai, May 11, 2016

ASTRON is part of the Netherlands Organisation for Scientific Research (NWO)

LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

European SKA Regional Science Centre
Concept and Planning

Talk Outline

Context for SKA Regional Centres
LOFAR as a case study
Capabilities and functionality of SRCs
The AENEAS project and a European SRC

ASTRON is part of the Netherlands Organisation for Scientific Research (NWO)

LOFAR


eMERLIN

Data Intensive Astronomy

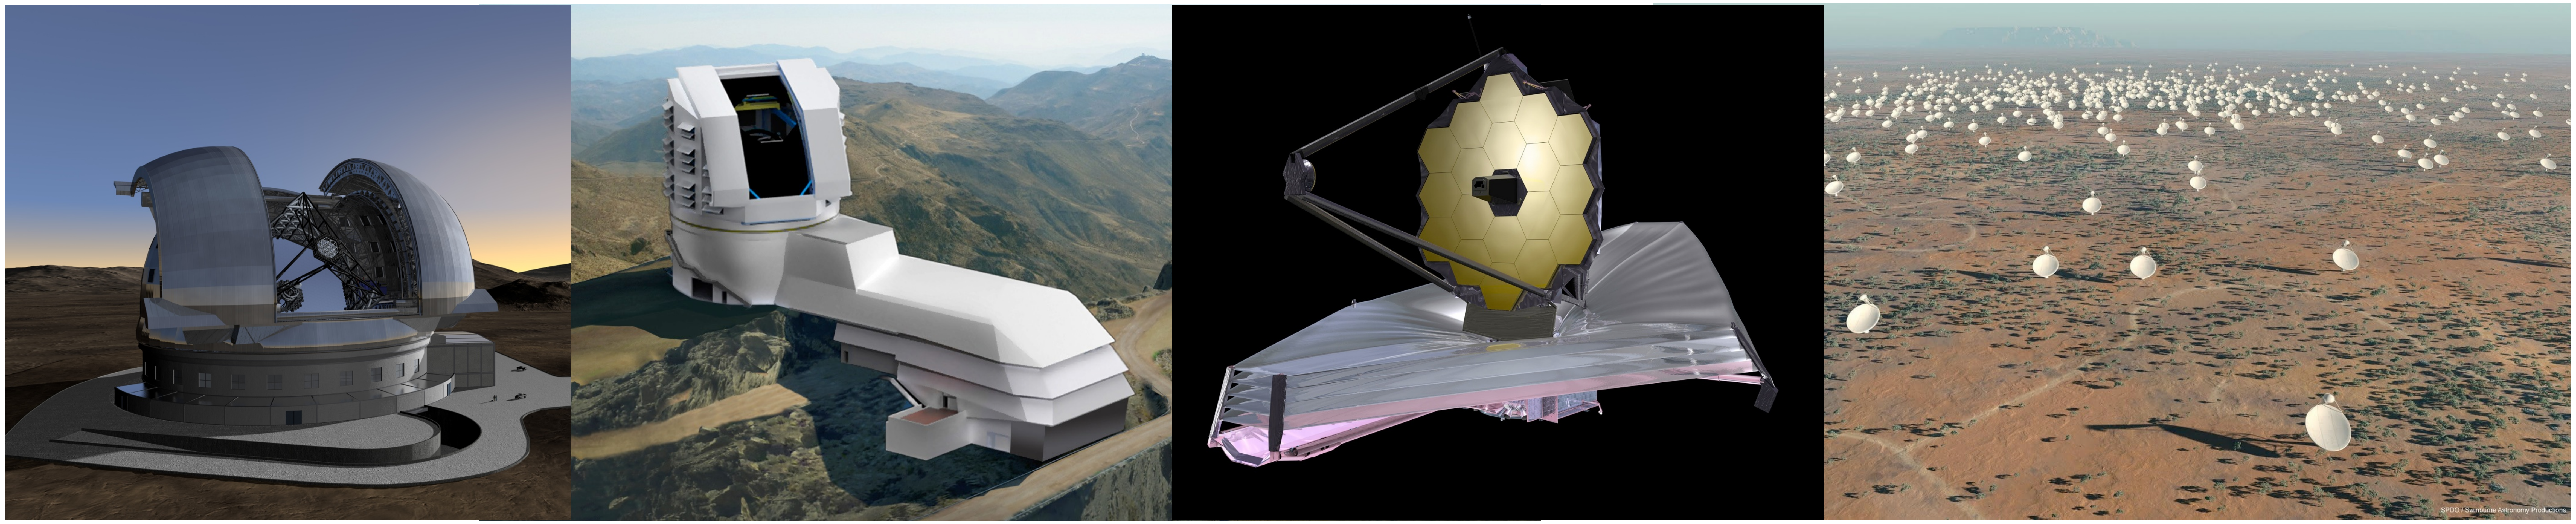

- Science is increasingly driven by large data sets
- Massive data collections and large scientific collaborations
- Most science extraction is based on the archived data
- Current instruments already producing petascale datasets

SKA will ultimately produce exabytes of science-ready data!

Impact of Science Archives

Science archives are a multiplier for total science output

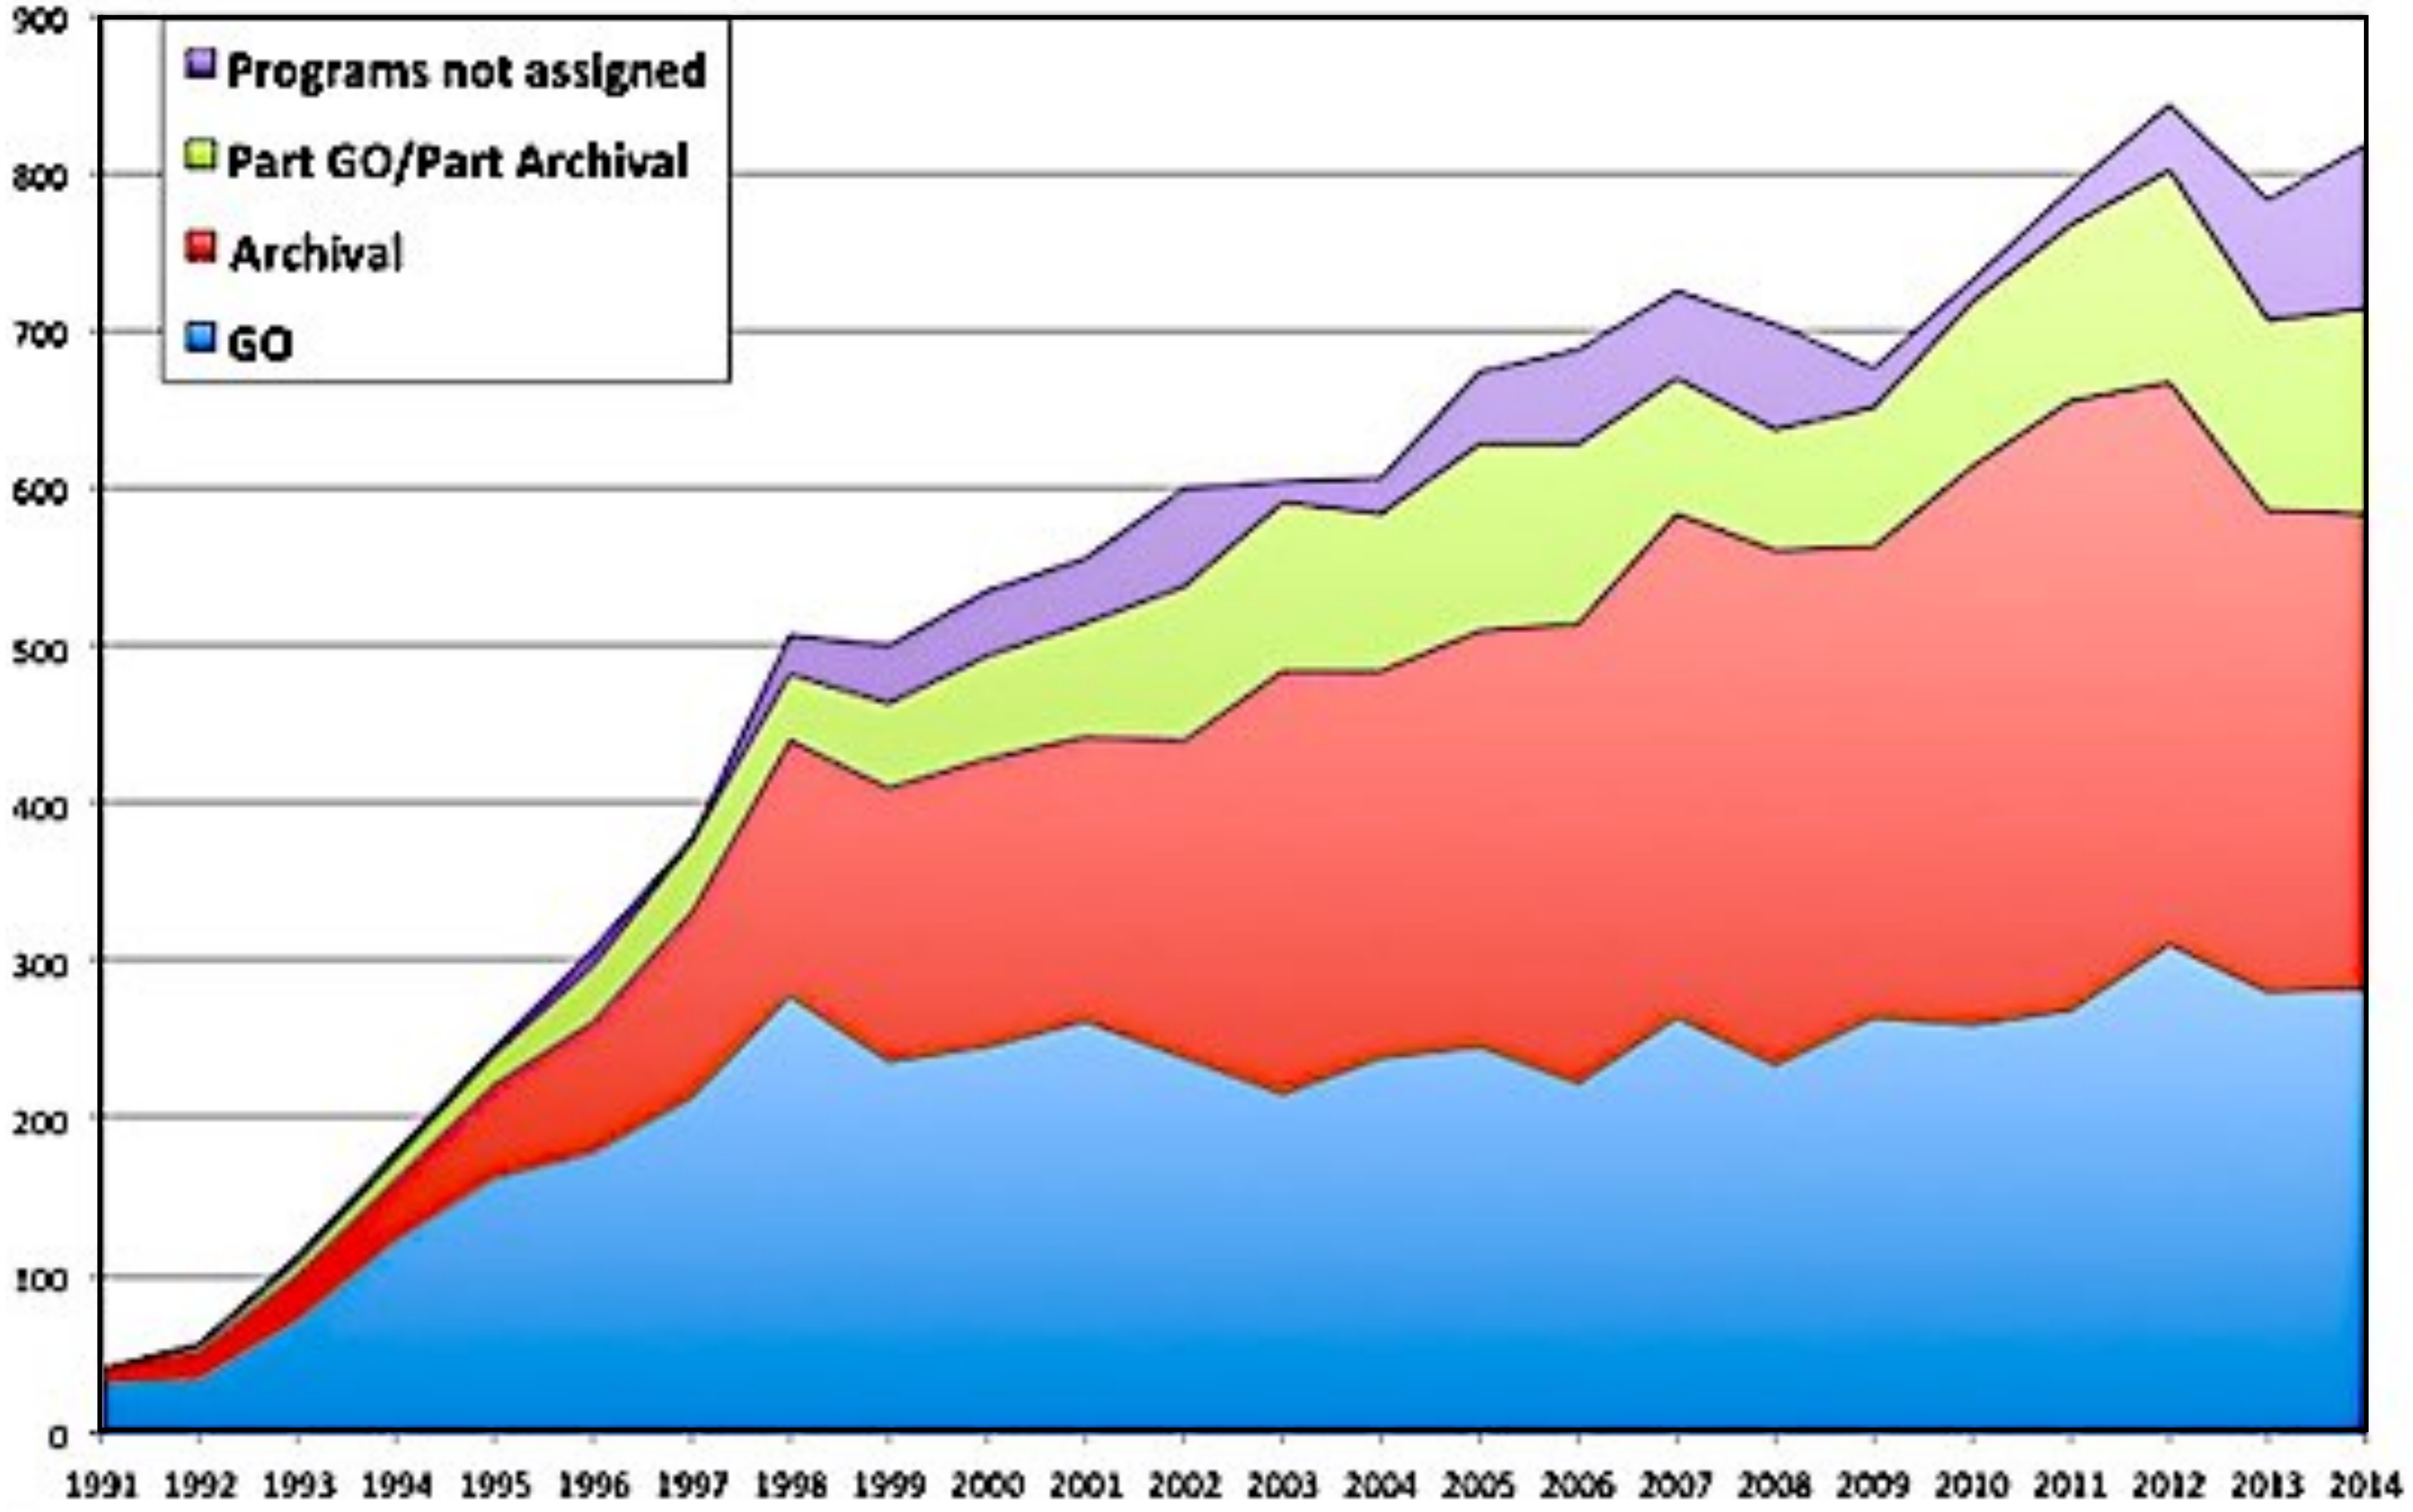

HST Publication Rate

- Assumes the archives are persistent and maintained
- Assumes archival data is open and accessible
- Assumes users retrieving data have resources to process to a science result
- Assumes data products stored are appropriate for general use

Evolution of Science Archives

Data centers are the new "large facilities"

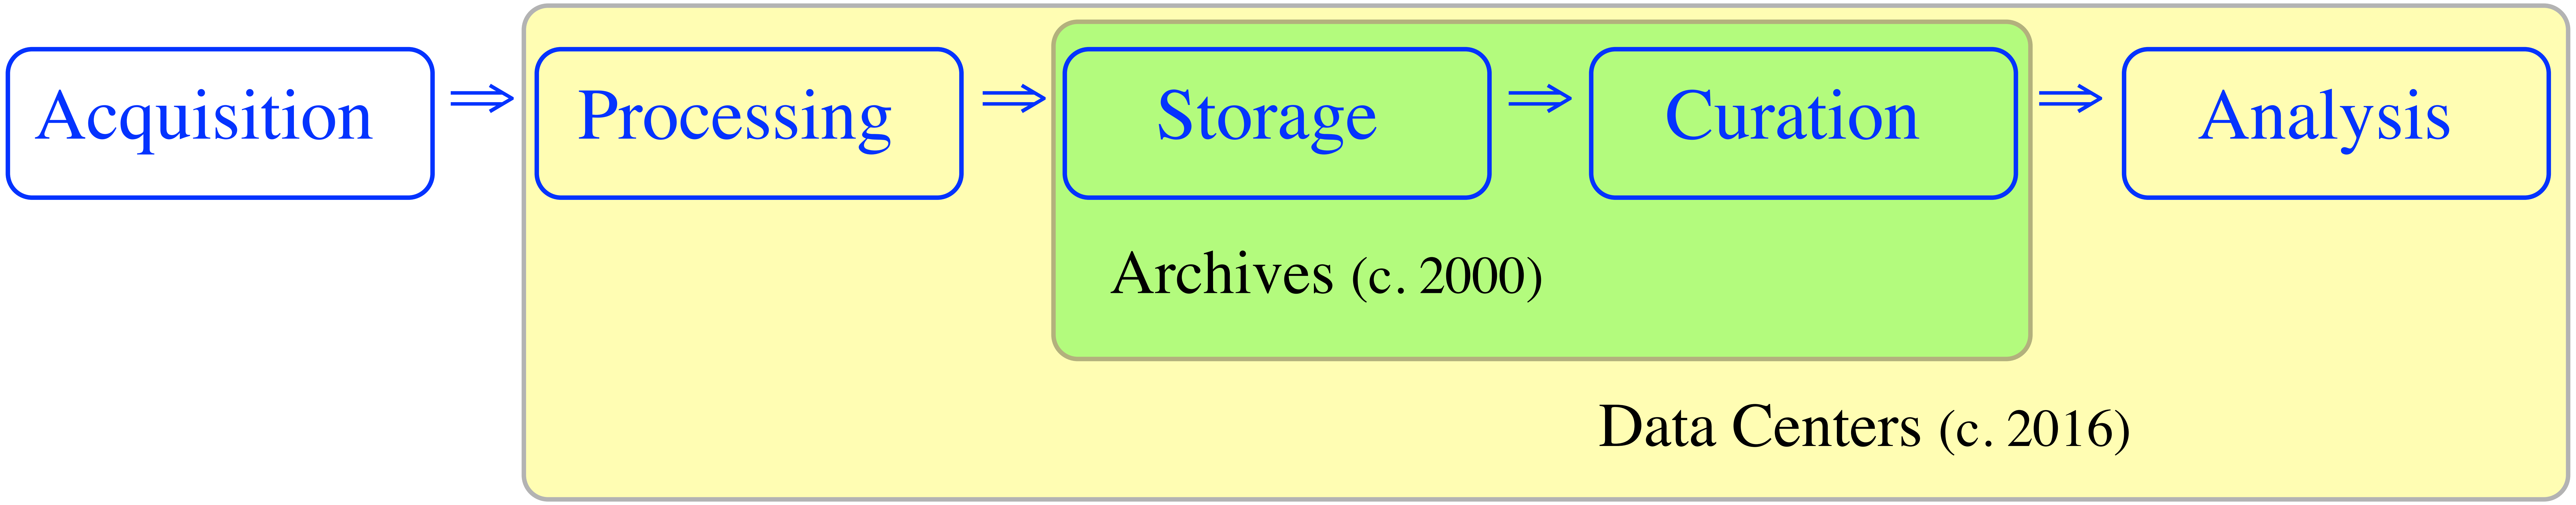

- Extreme data volumes require minimized data movement
- Scale of user analysis tasks exceeds individual or group compute resources
- Archives must evolve to include processing and analysis components
- Science Data Centers poised to become the "working surface" for users

International LOFAR Telescope

LOFAR Key Science Projects

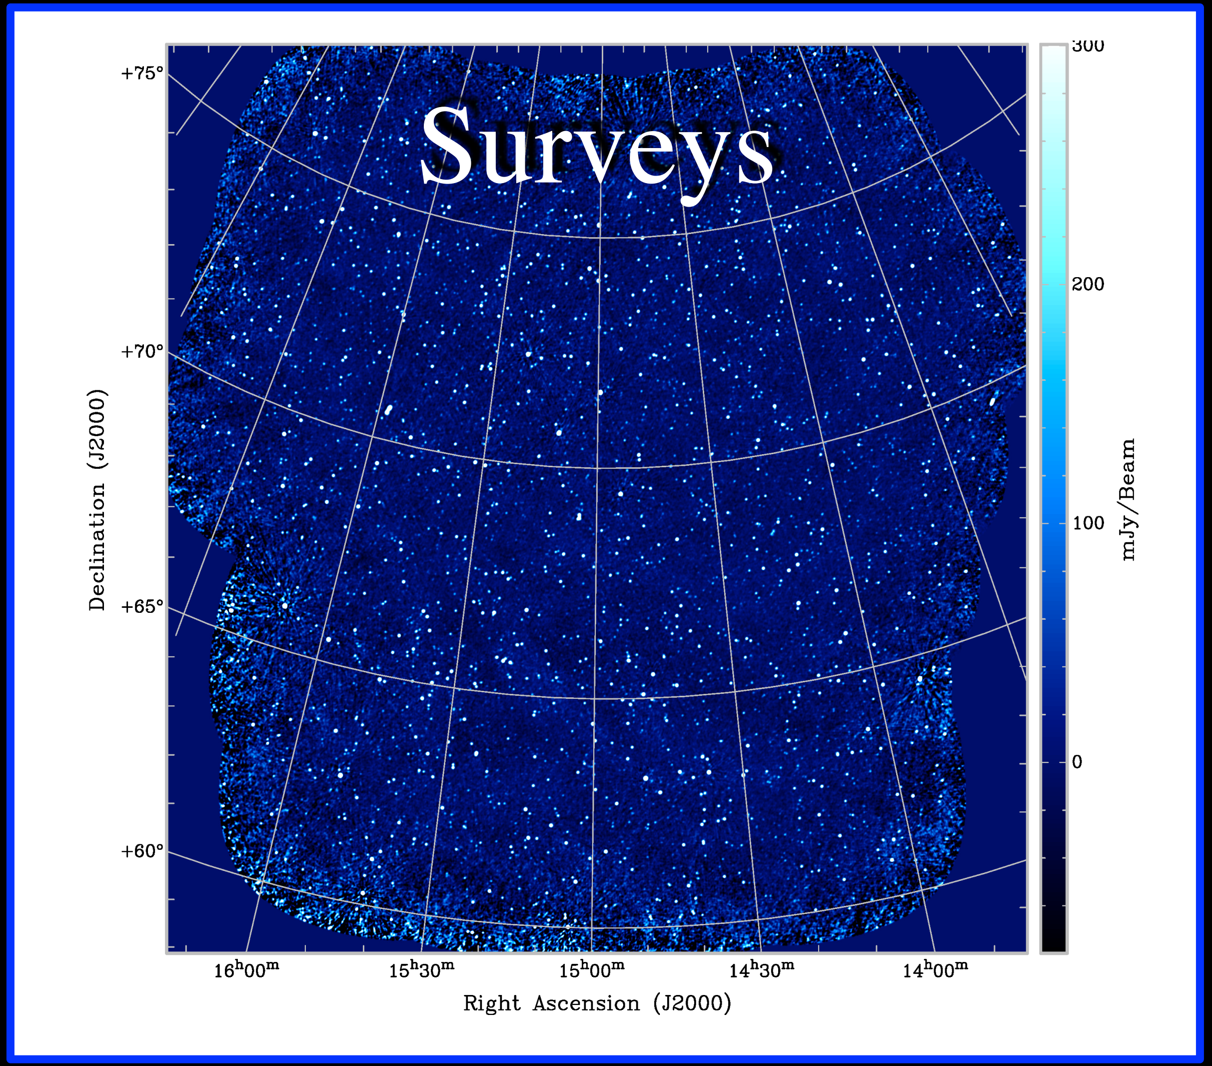

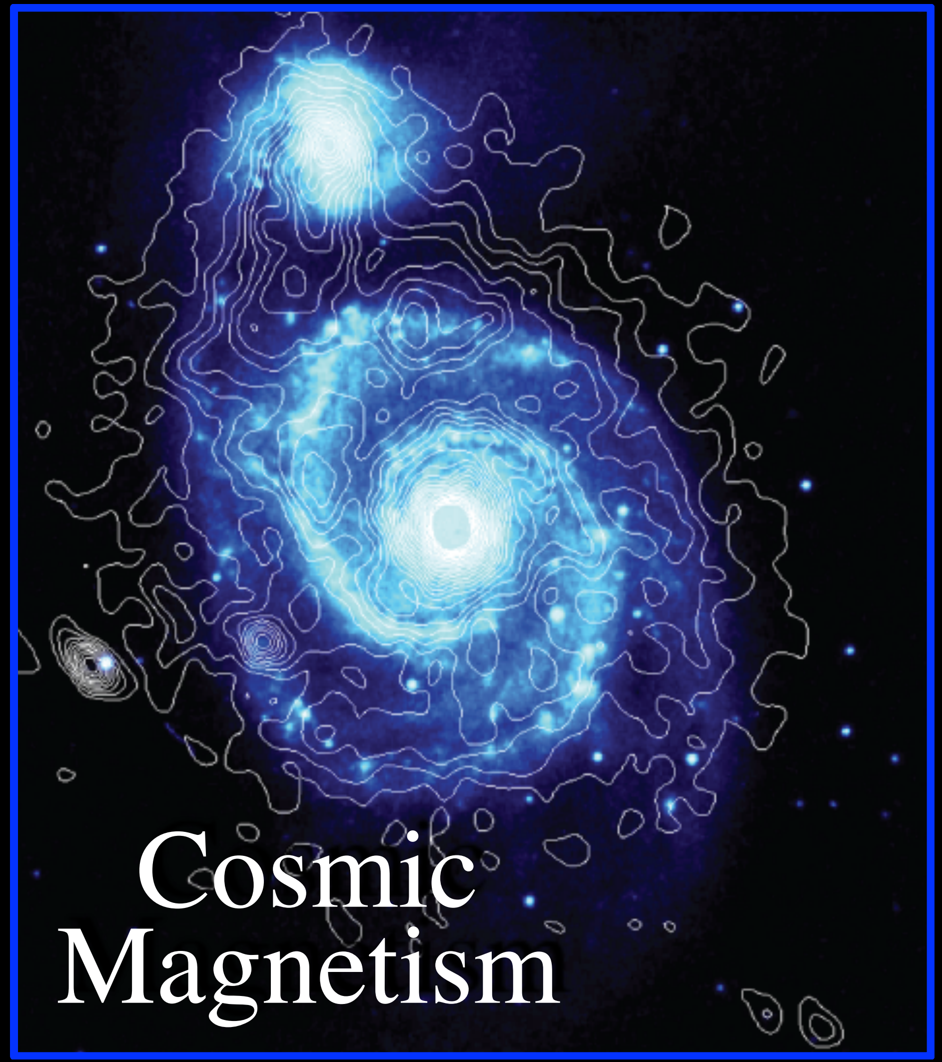

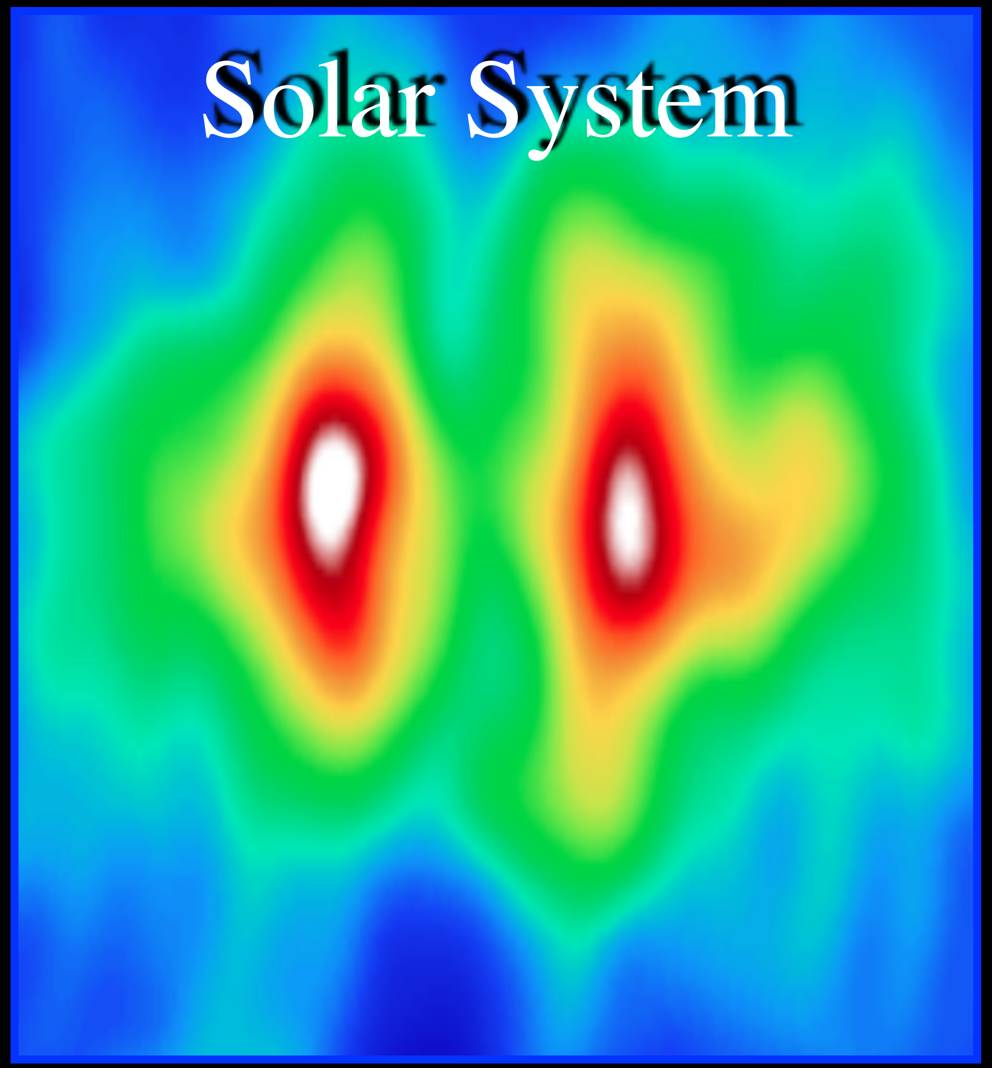

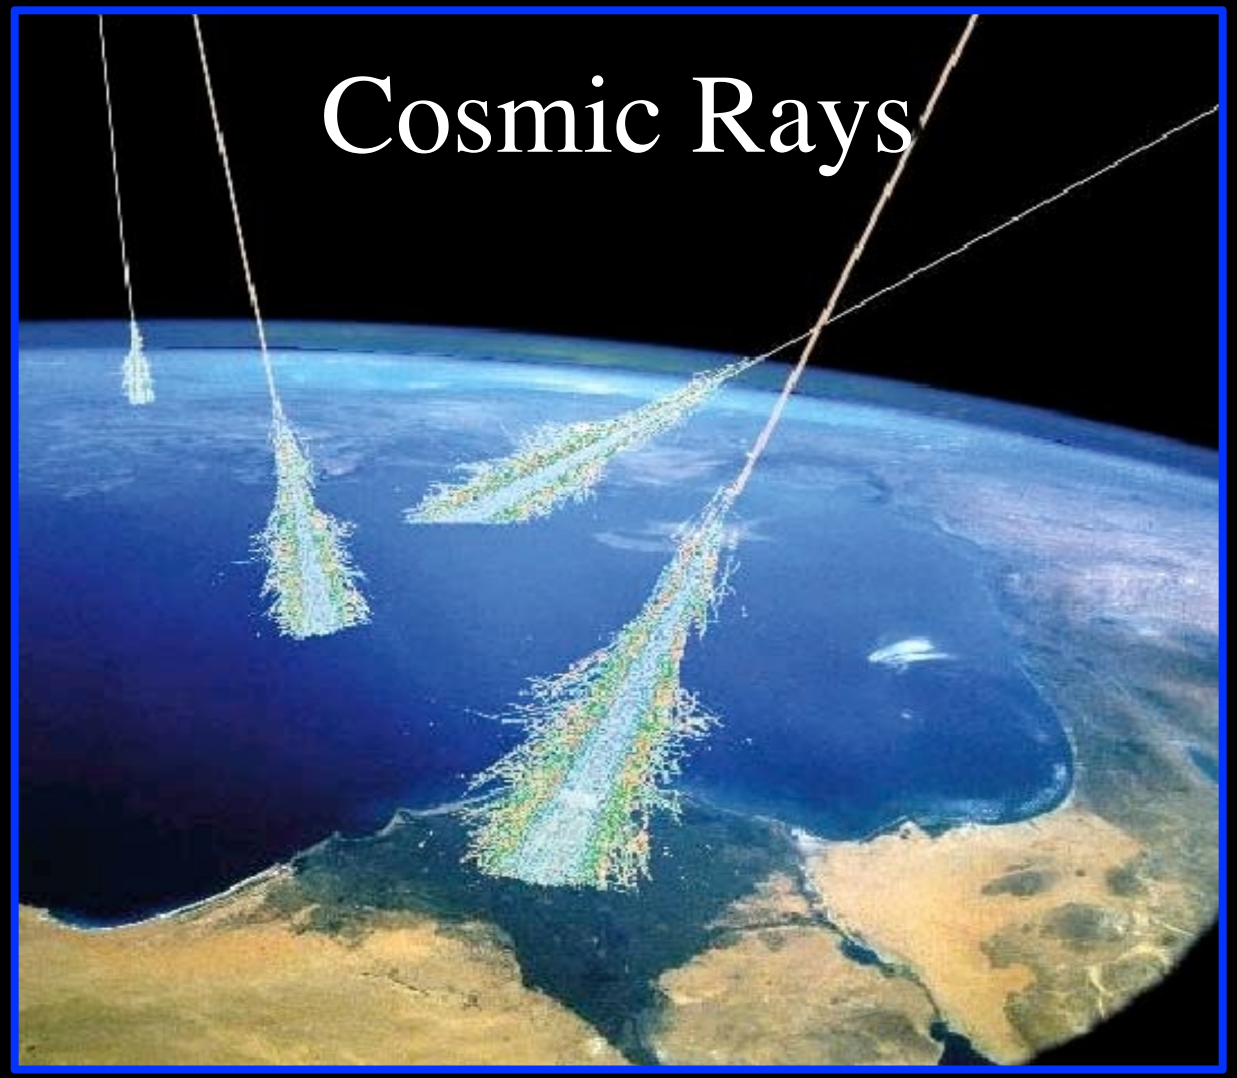

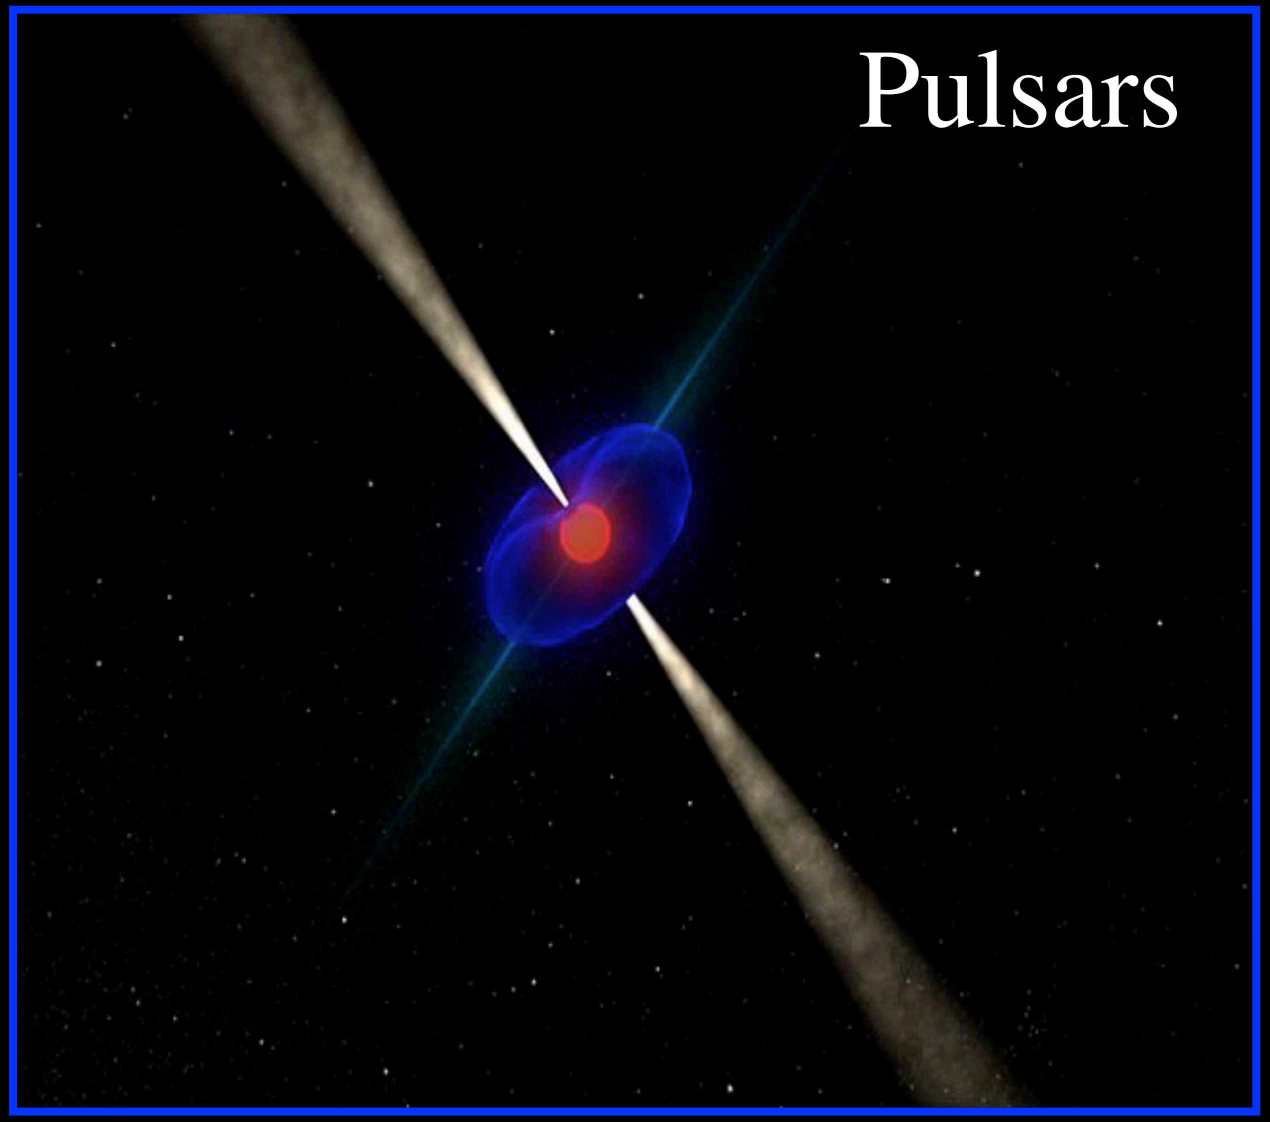

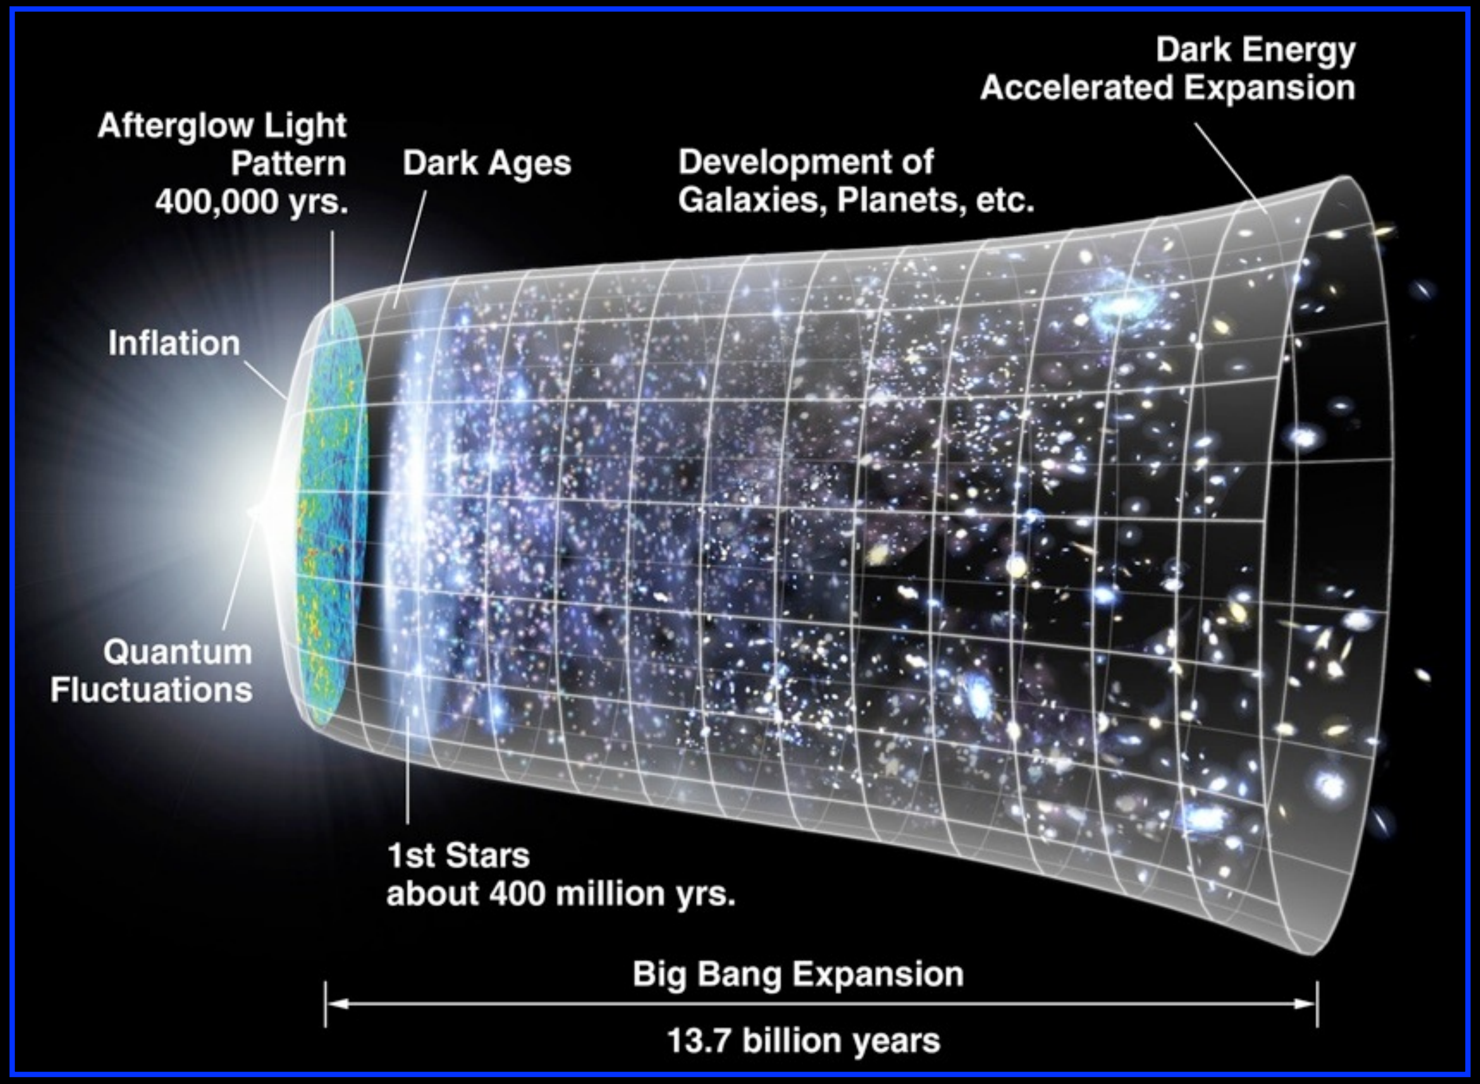

Epoch of Reionization

LOFAR Data Accumulation

- Data Storage
  - *23 Petabytes*
  - *3 PB/yr growth*
  - *4 sites, 3 countries*
  - *300 TB/month ingest*
  - *100 TB/month staged*

- Contents
  - *Over ~10^7 products*
  - *10^9 individual files*
  - *Visibilities, images, and BF data*
  - *Does not include raw visibilities*

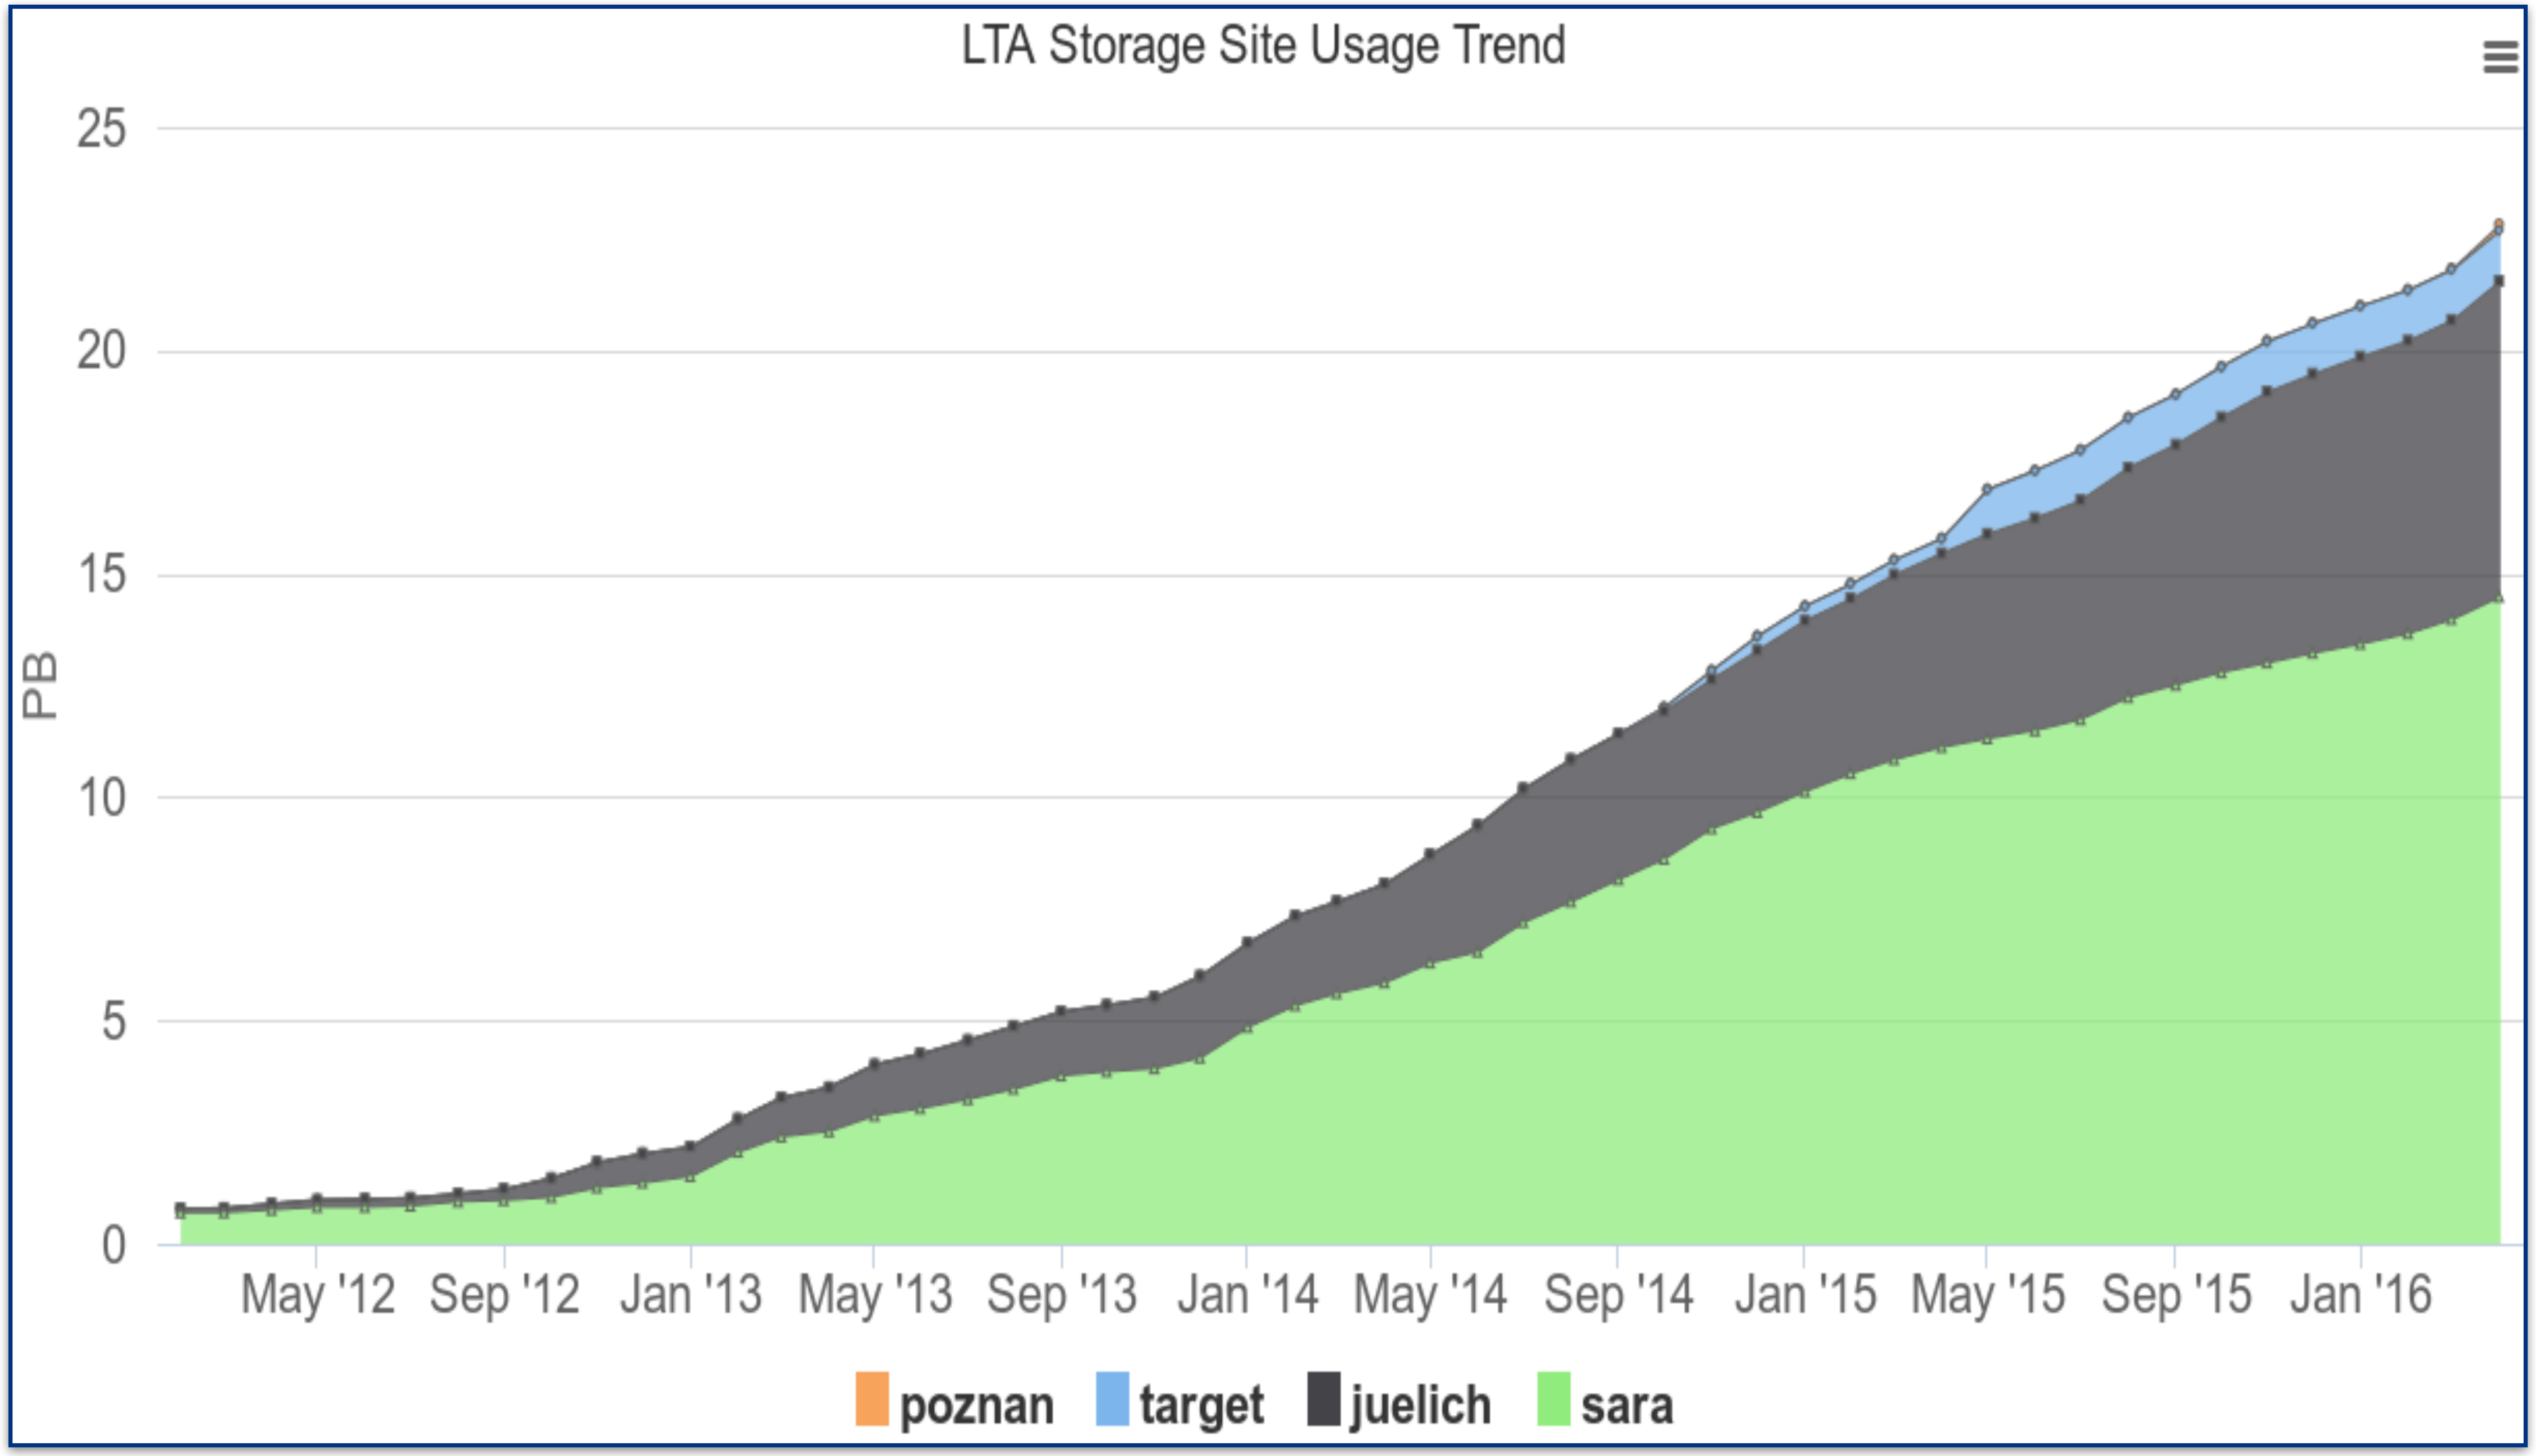

LOFAR Long-Term Archive (LTA) is federated over four locations.

Data Access Patterns

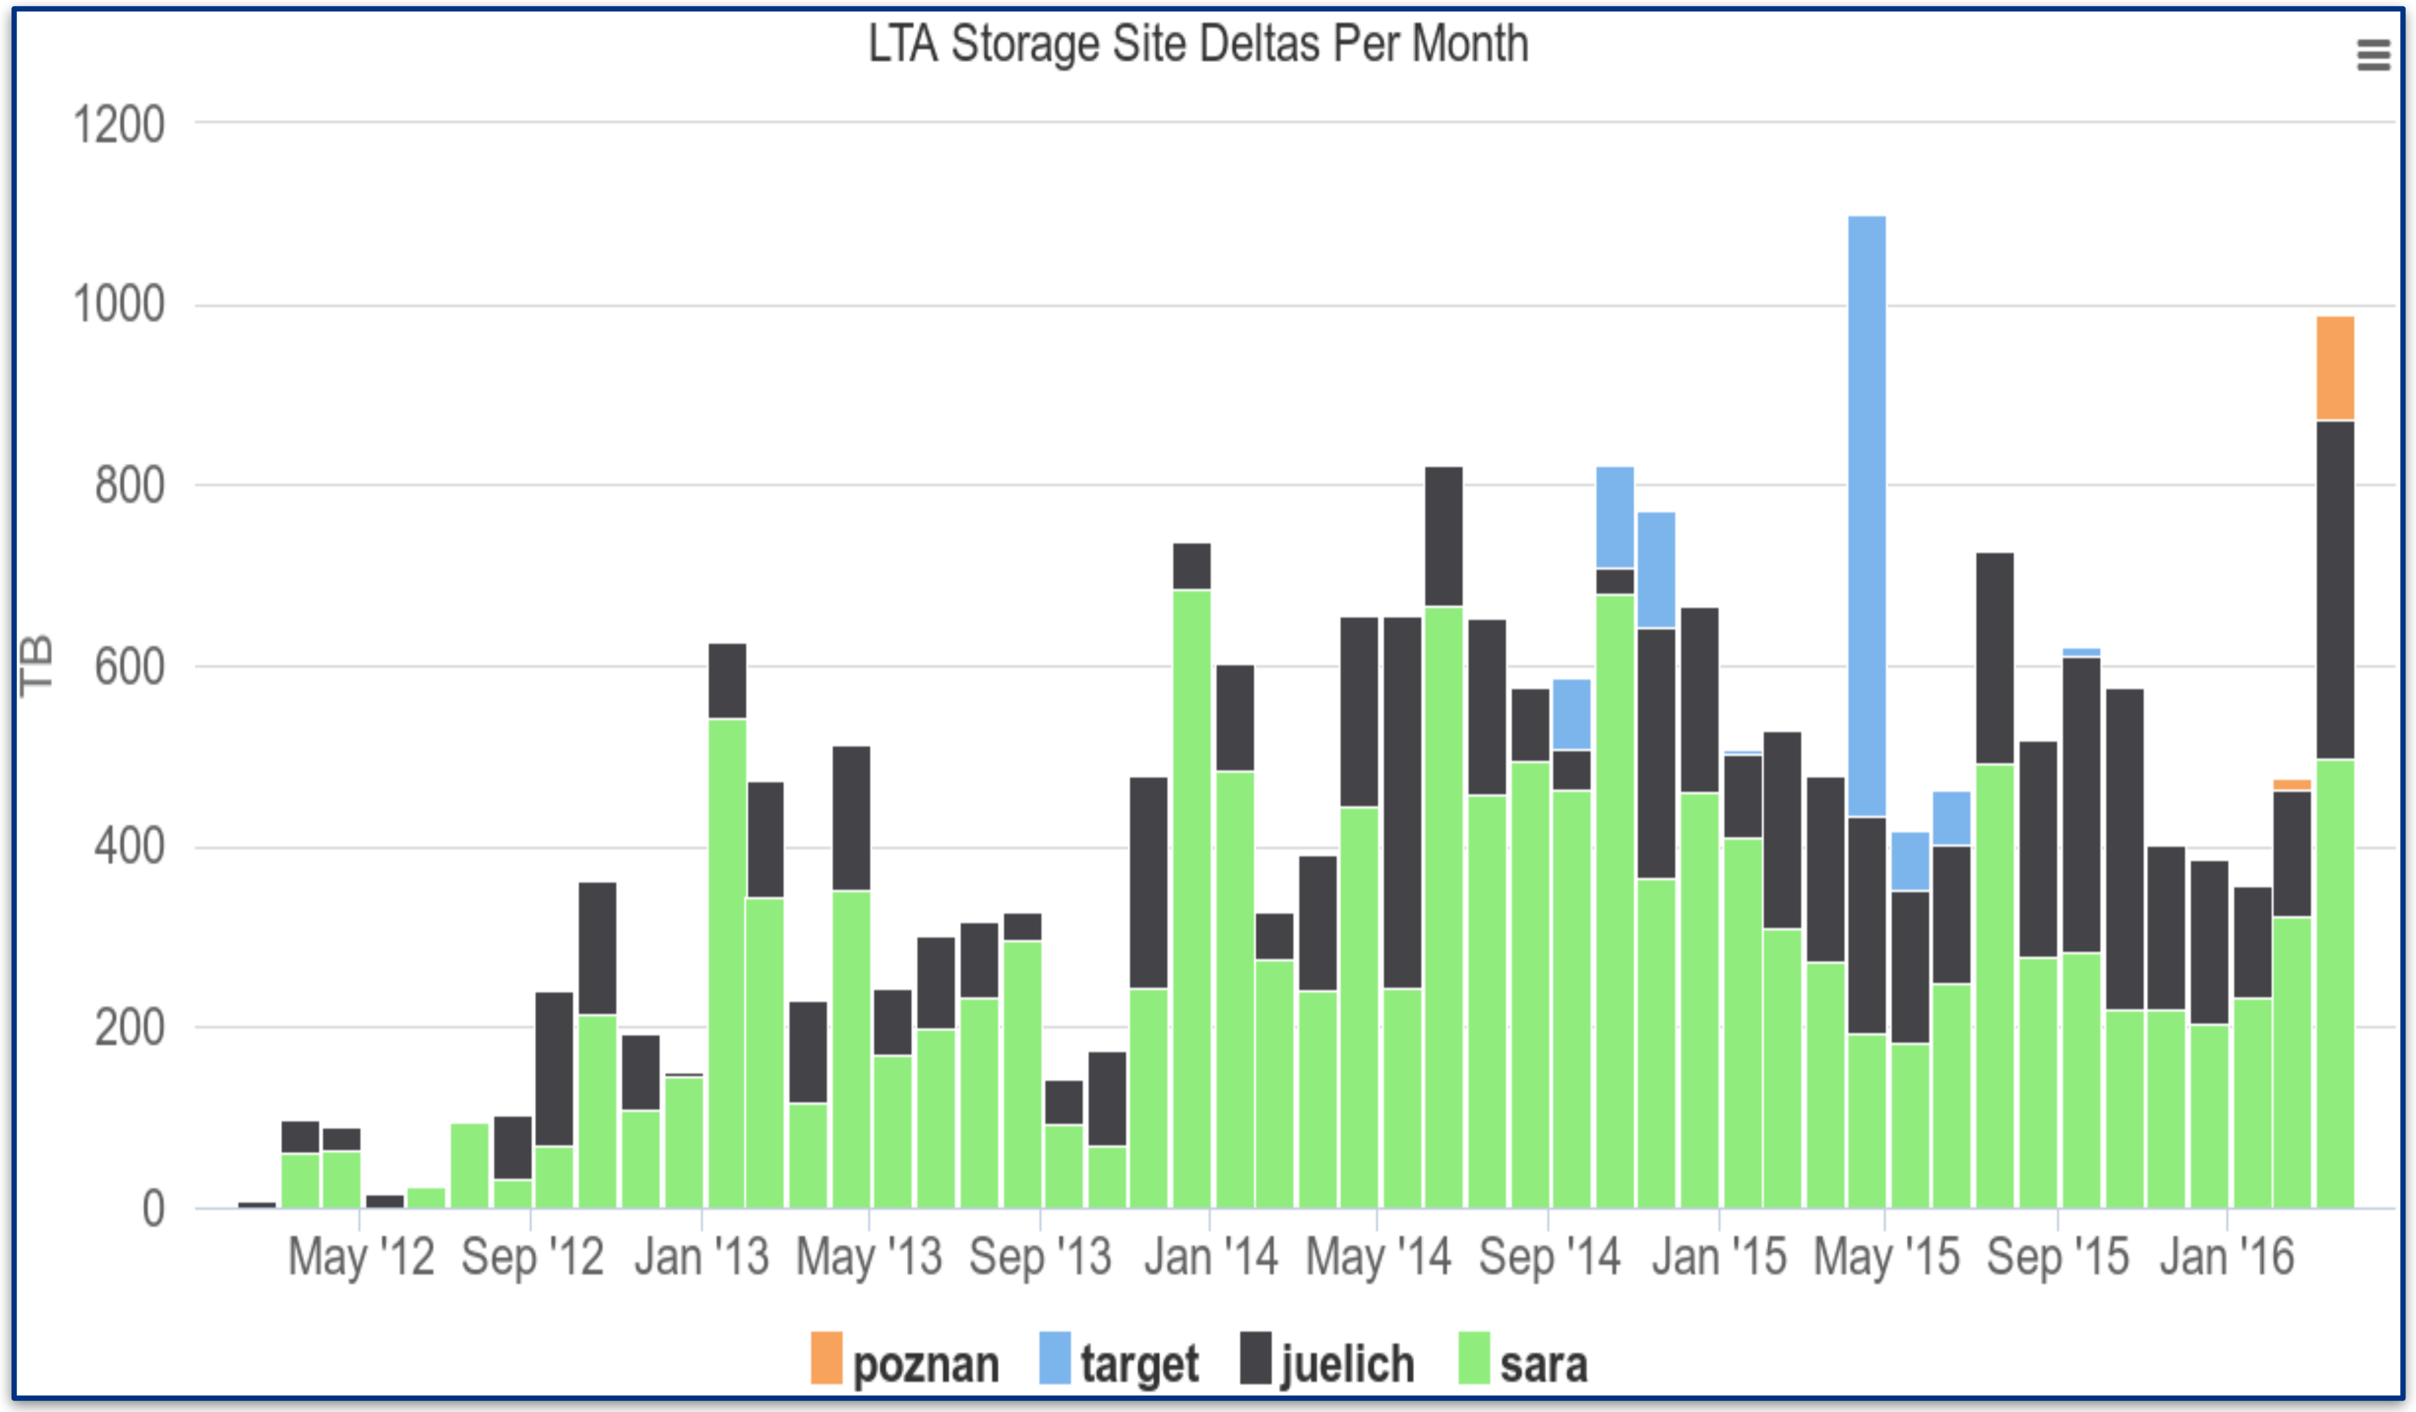

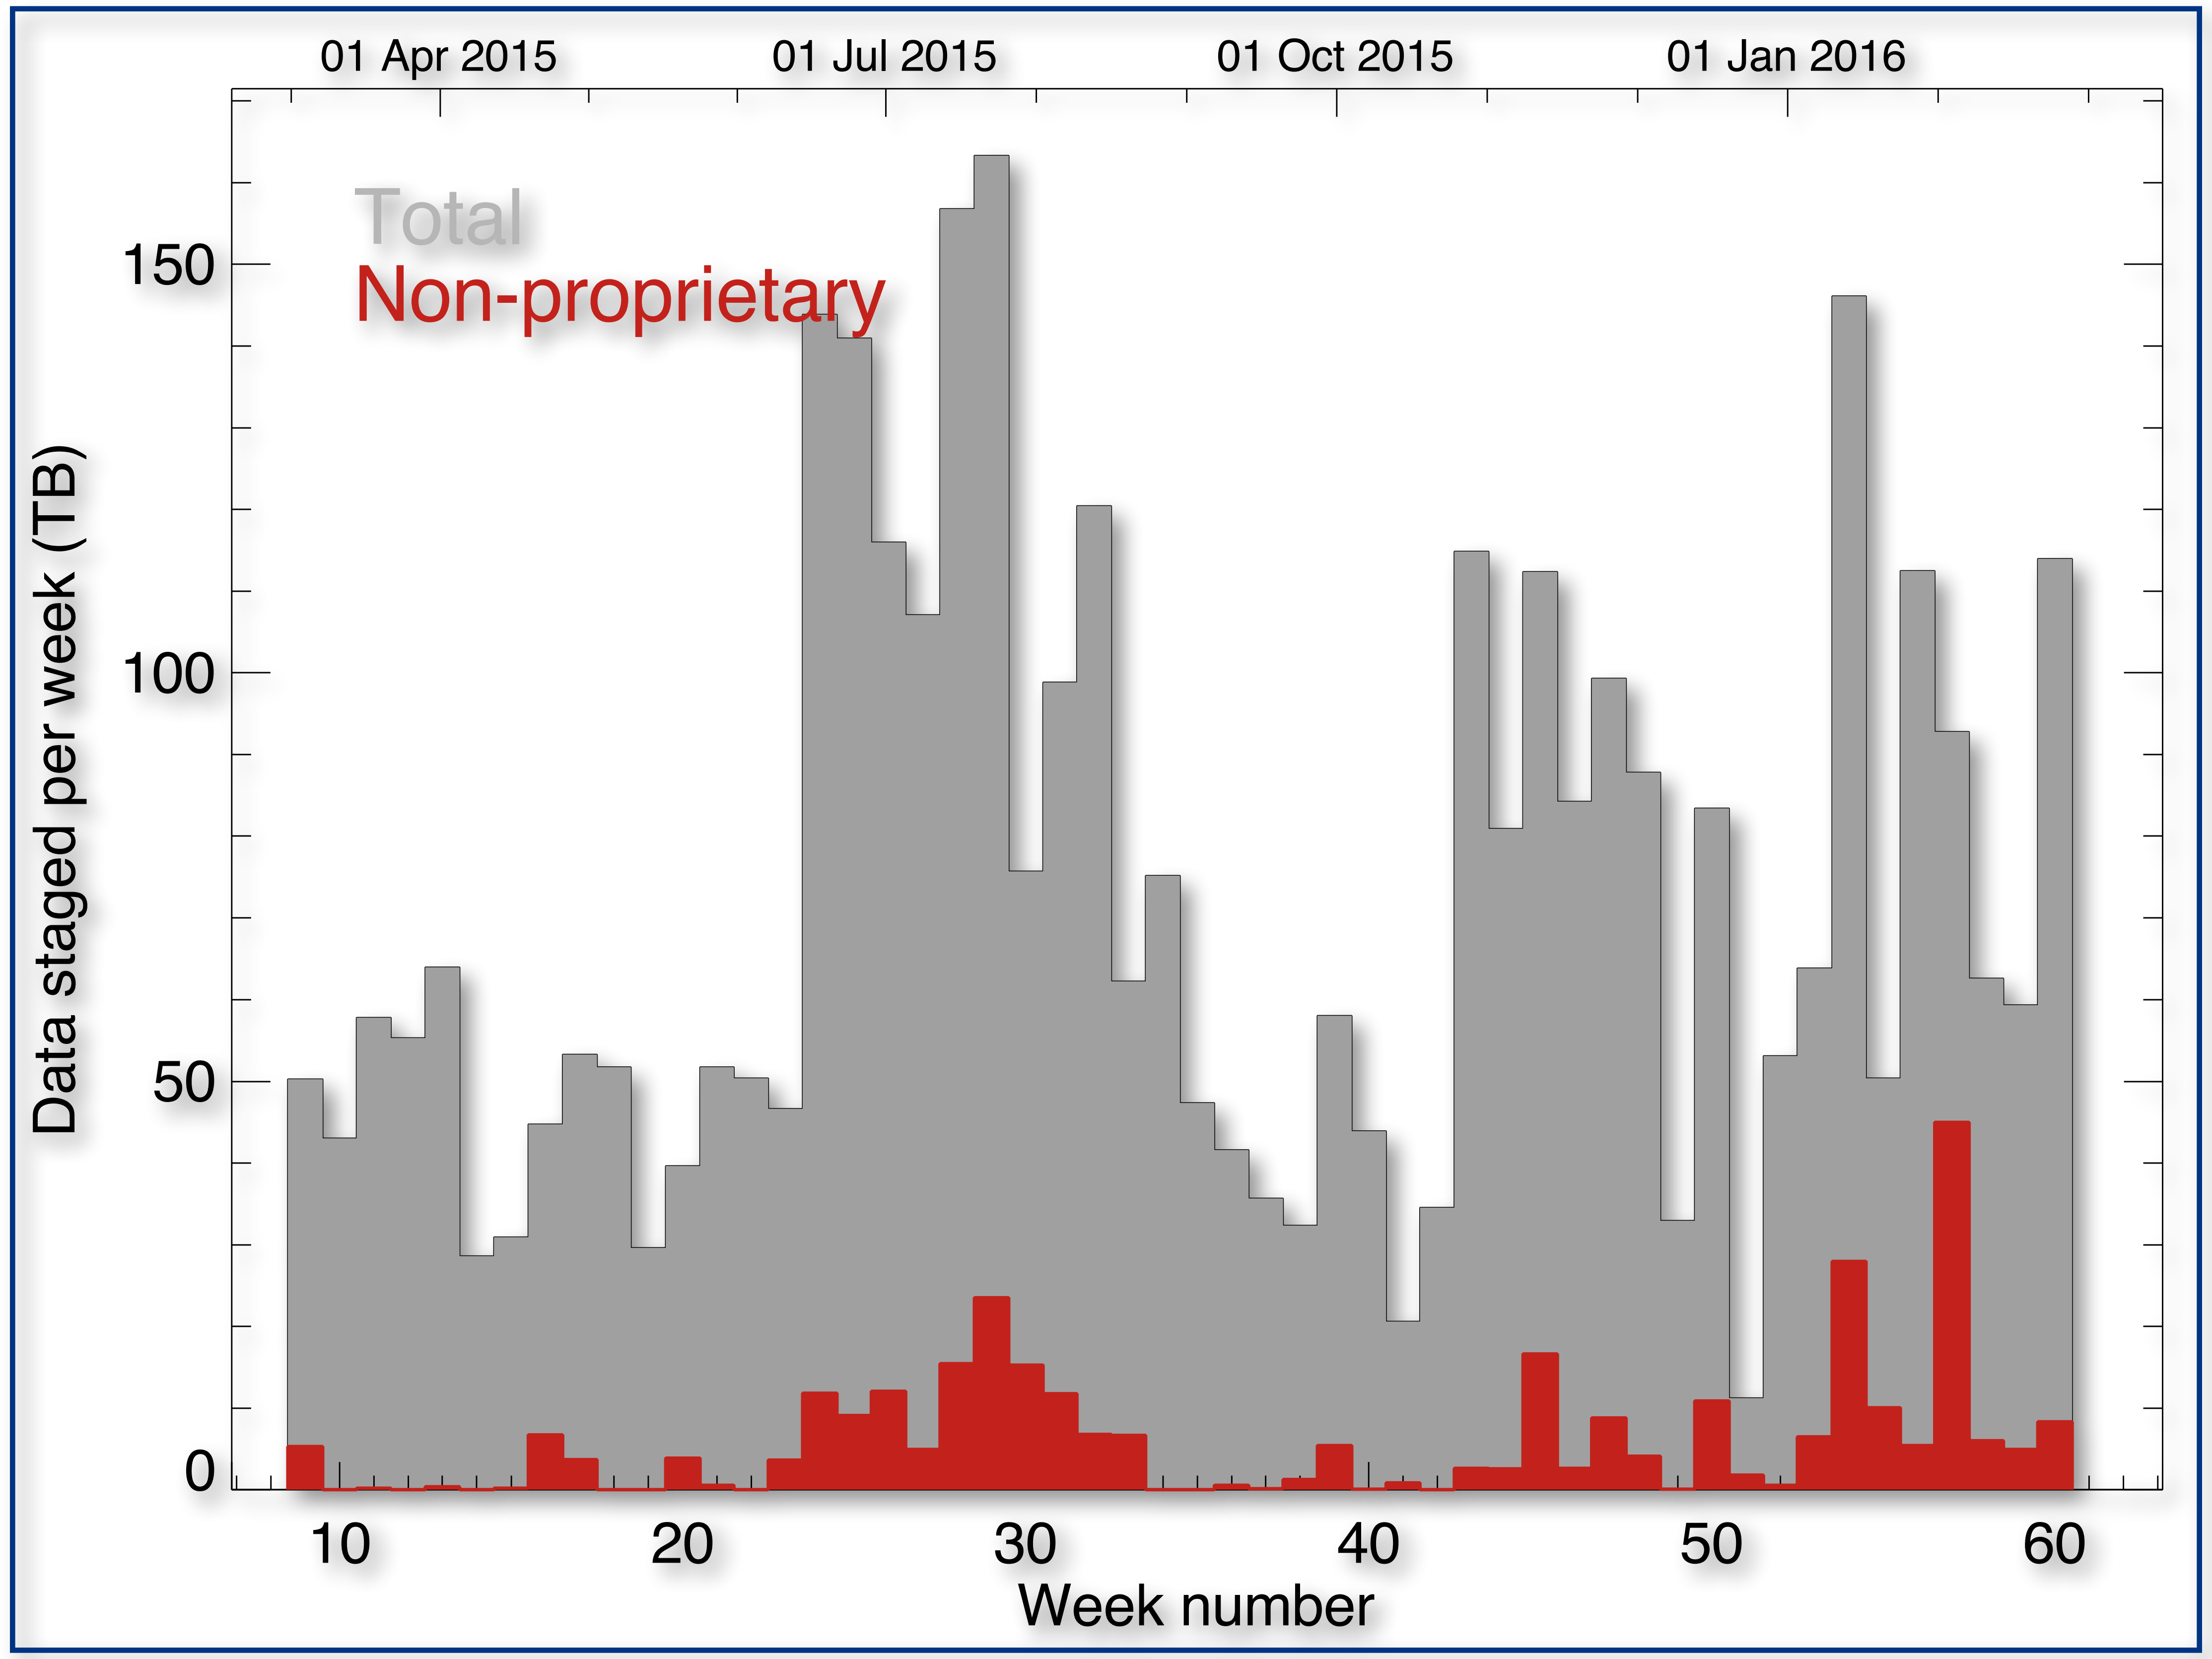

Typical 5 node clusters (320 Gb, 120 cores, 250 Tb) at individual research institutes are NO longer sufficient:

→ *Data transfer from archive to institutes too slow: ~ 10 Mb/s*

→ *Current P/O for a single observation too high: 10 - 100*

Typical data size is 10-100 Tb

Problematic for many researchers!

Processing in a Distributed Archive

- National NL GRID Resources
  - *14 data centres (3 large grid clusters, 11 smaller ones)*
  - *approx. 10,000 compute cores*
  - *12 PB disk, 170 PB tape*
- Global GRID Resources
  - *170 data centres in 36 countries*
  - *more than 330,000 compute cores,*
  - *500 PB disk, 500 PB tape*

For LOFAR ⇒ Standardized pipelines
Integration with catalog & UIs
Processing close to the data

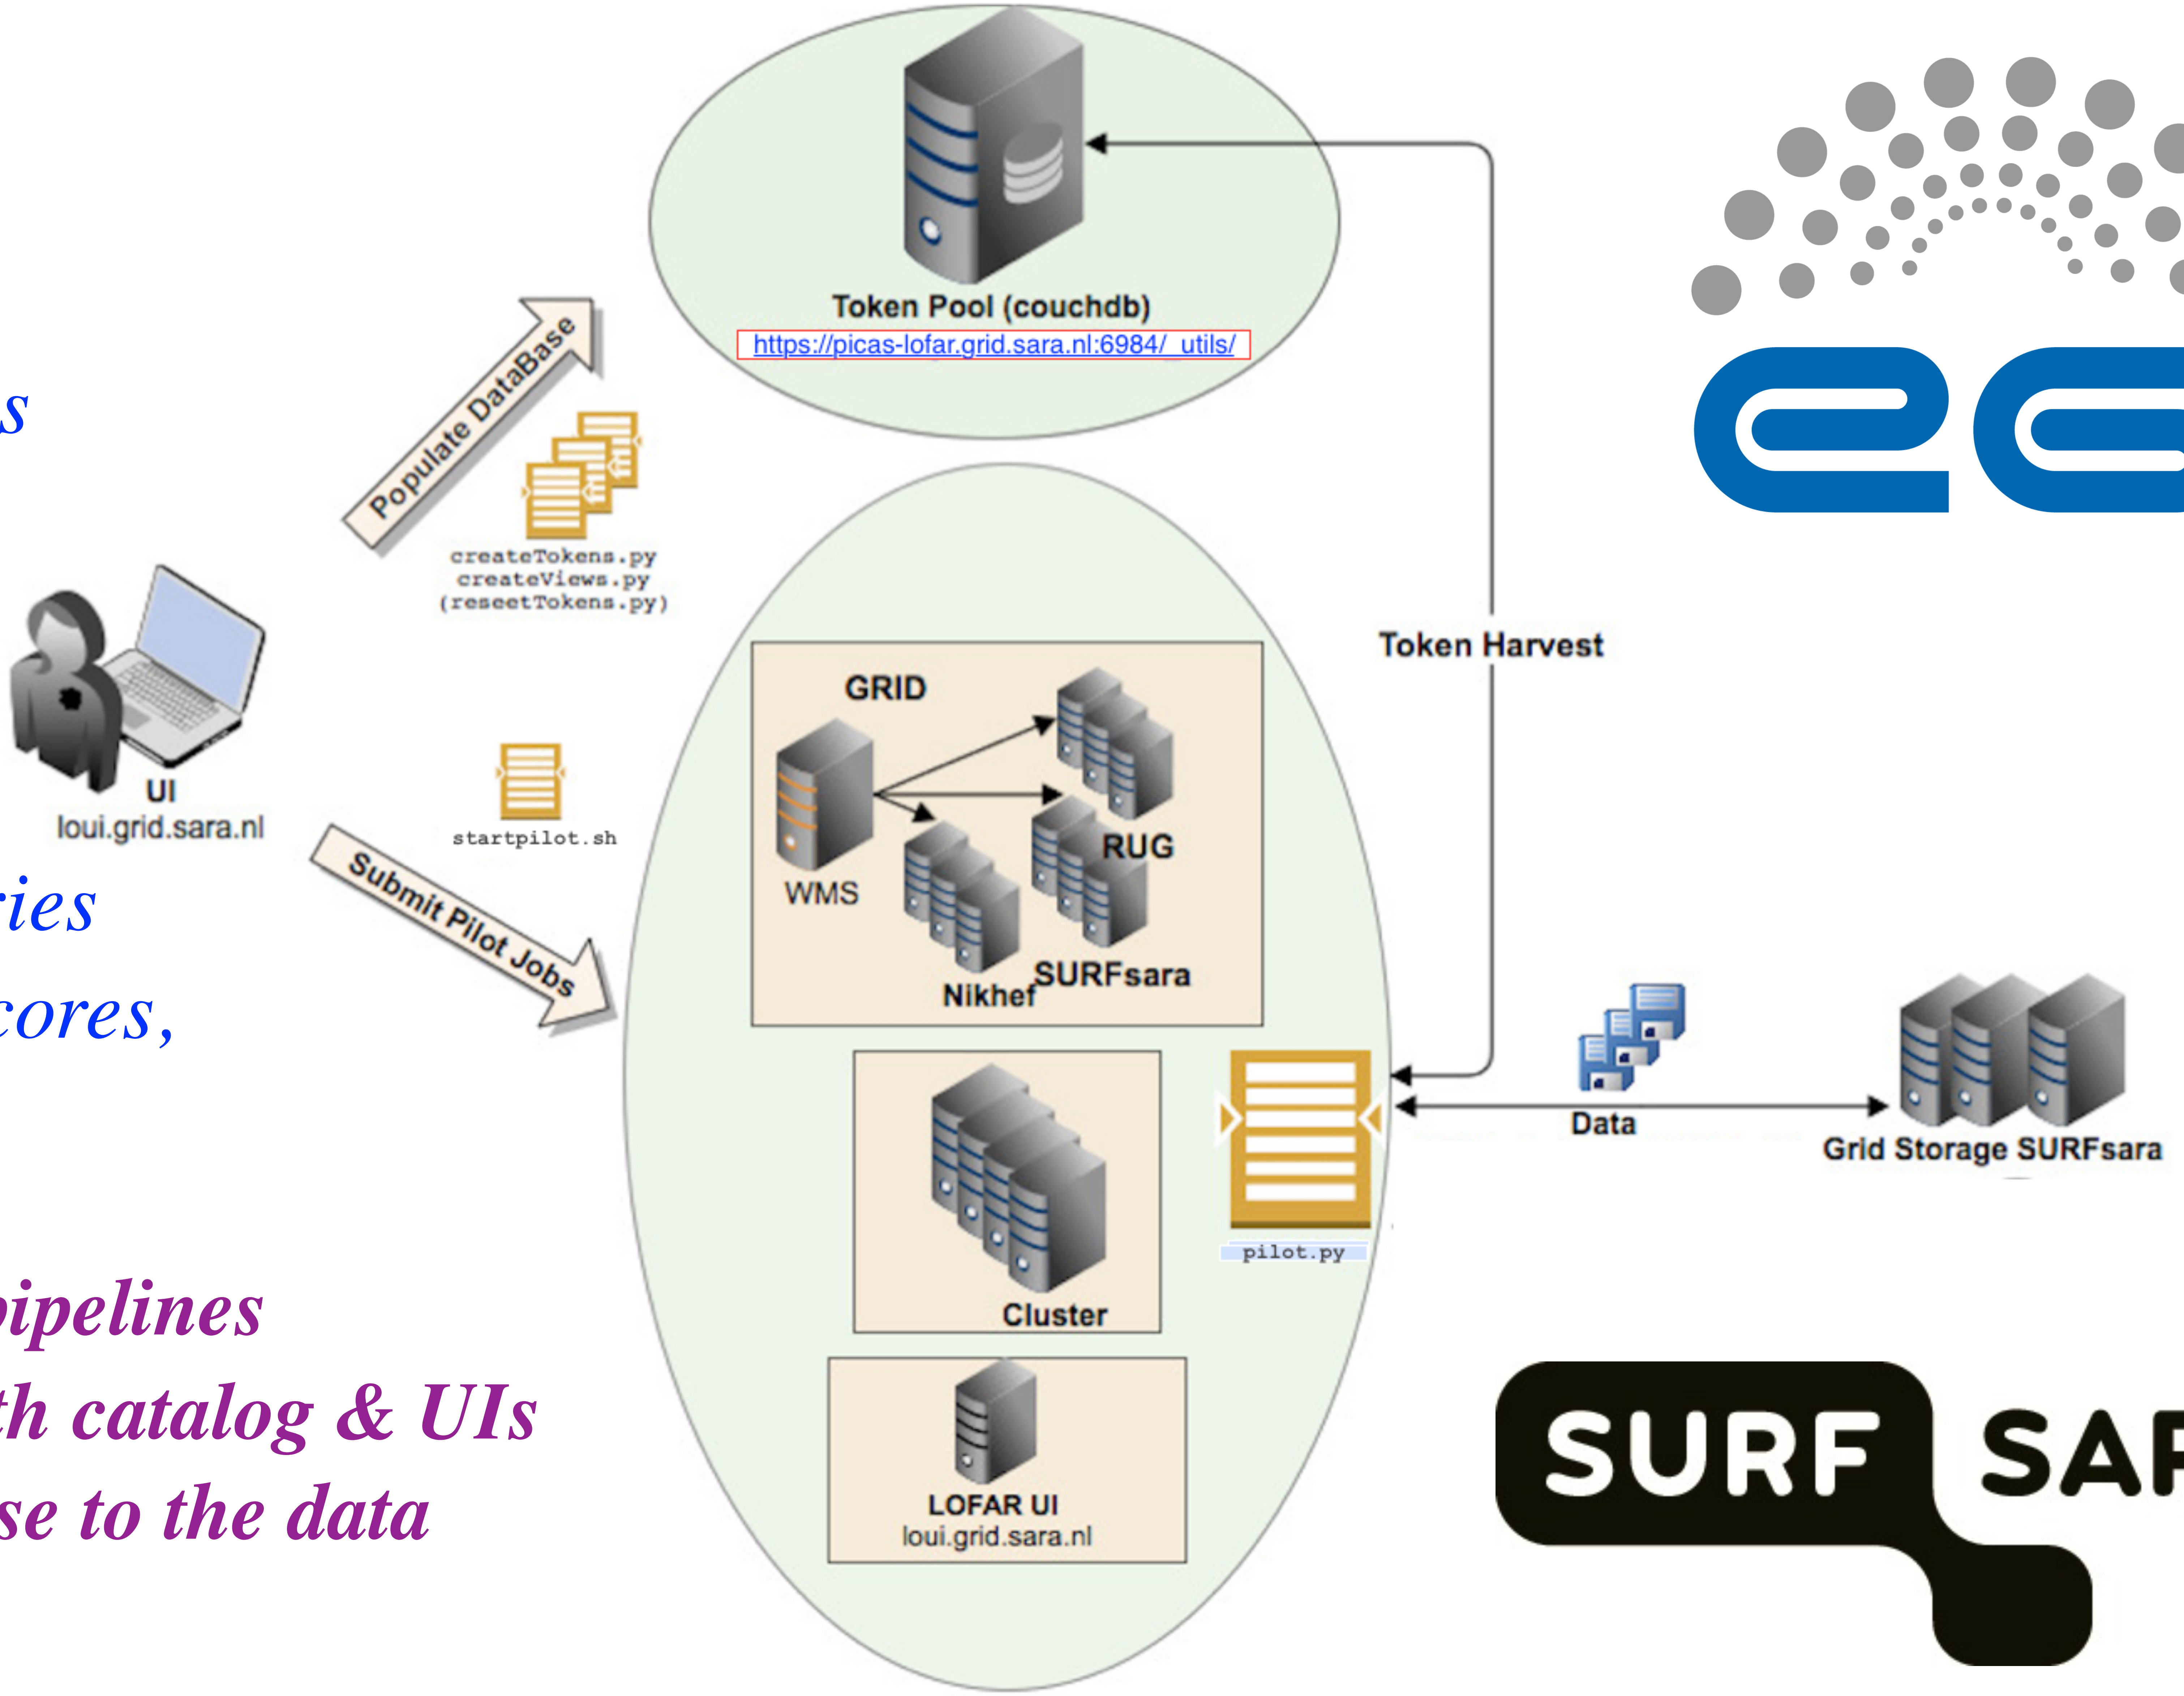

EGI Federated Cloud for calibrating and analysing Radio-Astronomy data

Susana Sanchez-Expósito [1,*], José Sabater [2], Daniele Lezzi [3], Julian Garrido [1], Lourdes Verdes-Montenegro [1], José Enrique Ruiz [1], Pablo Martín [4]. Raül Sirvent [3], Rosa M. Badia [3], Antonio Ruiz-Falcó [4]

(1) Instituto de Astrofísica de Andalucía – CSIC
(2) University of Edinburgh
(3)Barcelona Supercomputing Center
(4) Fundación de Supercomputación de Castilla y León

Use Case 1: LOFAR data calibration

Fig.1. LOFAR datacube representation

An interferometer as LOFAR correlates the signals from several antennas, generating the so-called measurement sets. They are a kind of **datacubes** (3D data)**:** two Fourier spatial coordinate axes plus a spectral axis.

A datacube can reach several **terabytes,** depending on factors as the amount of involved antennas, the observation time, as well as the amount of observed subbands – i.e. frequency intervals-. LOFAR telescope allows **up to 488 subbands**, which can reach several GBs. **Each subband can be processed independently what allows the parallelization of the whole datacube calibration**.

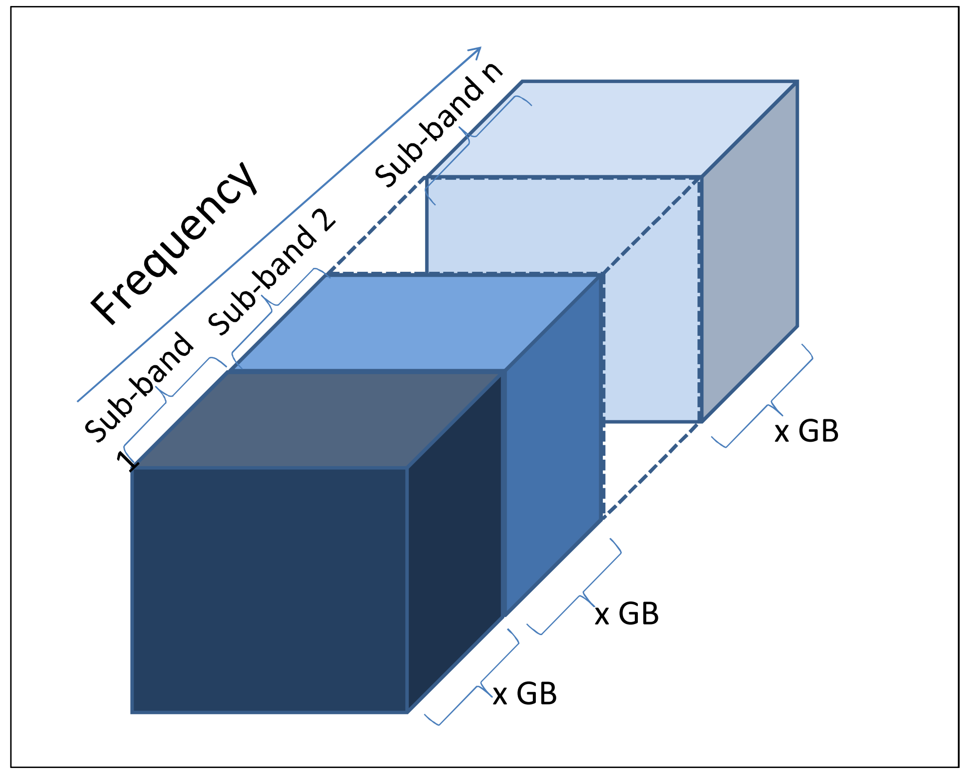


```
import subprocess
import sys
import os
from pycompss.api.task import task
from pycompss.api.parameter import *

@task( script_name = FILE)
def iter_calib(script_name):
   os.chmod(script_name,0744)
   subprocess.call(script_name)
   print "end executiong"
if __name__ == "__main__":
   args = sys.argv[1:]
   DATA_PATH=args[0]
   TEMPLATE_FILE=args[1]
   f=open(TEMPLATE_FILE,'r')
   content=f.read()
   f.close()
   list_f=os.listdir(DATA_PATH)
   for directory in list_f: # Iterate over the data inputs
      if os.path.isdir(DATA_PATH+"/"+directory):
         new_content=content.replace("INPUTDATAPATH",directory)
         script_name="job"+directory+".sh"
         f=open(script_name,"w")
         f.write(new_content)
         f.close()
         iter_calib(script_name)
```

Fig.2. COMPSs application. It iterates over the subbands, executing for each one a COMPSs task that calls the LOFAR software. Through a simple interface for describing the methods, COMPSs is able to analyse the dependencies among them, to match their requirements with the available resources and to orchestrate their execution on VMs.

Fig.3. Execution time. This figure shows the results of different tests in which the application has been configured to calibrate from 1 to 4 sub-bands, and its tasks have been distributed either on a high capacity VM (mammoth=32GB memory + 8 cores) or on several smaller VMs (large=8GB memory + 4 cores). Since each subband is processed in parallel, **the executions for calibrating one subband take approximately the same time than those for calibrating several subbands**. The results also reveal that **the execution time for the whole COMPSs application (in red) is slightly higher than for the tasks (in blue)**. Thus we can state that the time to start and contextualize the VMs is not significant. In addition, the time for the applications running on mammoth is higher than the applications whose tasks have been distributed on smaller VMs. This would mean that **distributing the tasks among several small VMs is more efficient than gathering them in a VM with high memory capacity and amount of cores.**

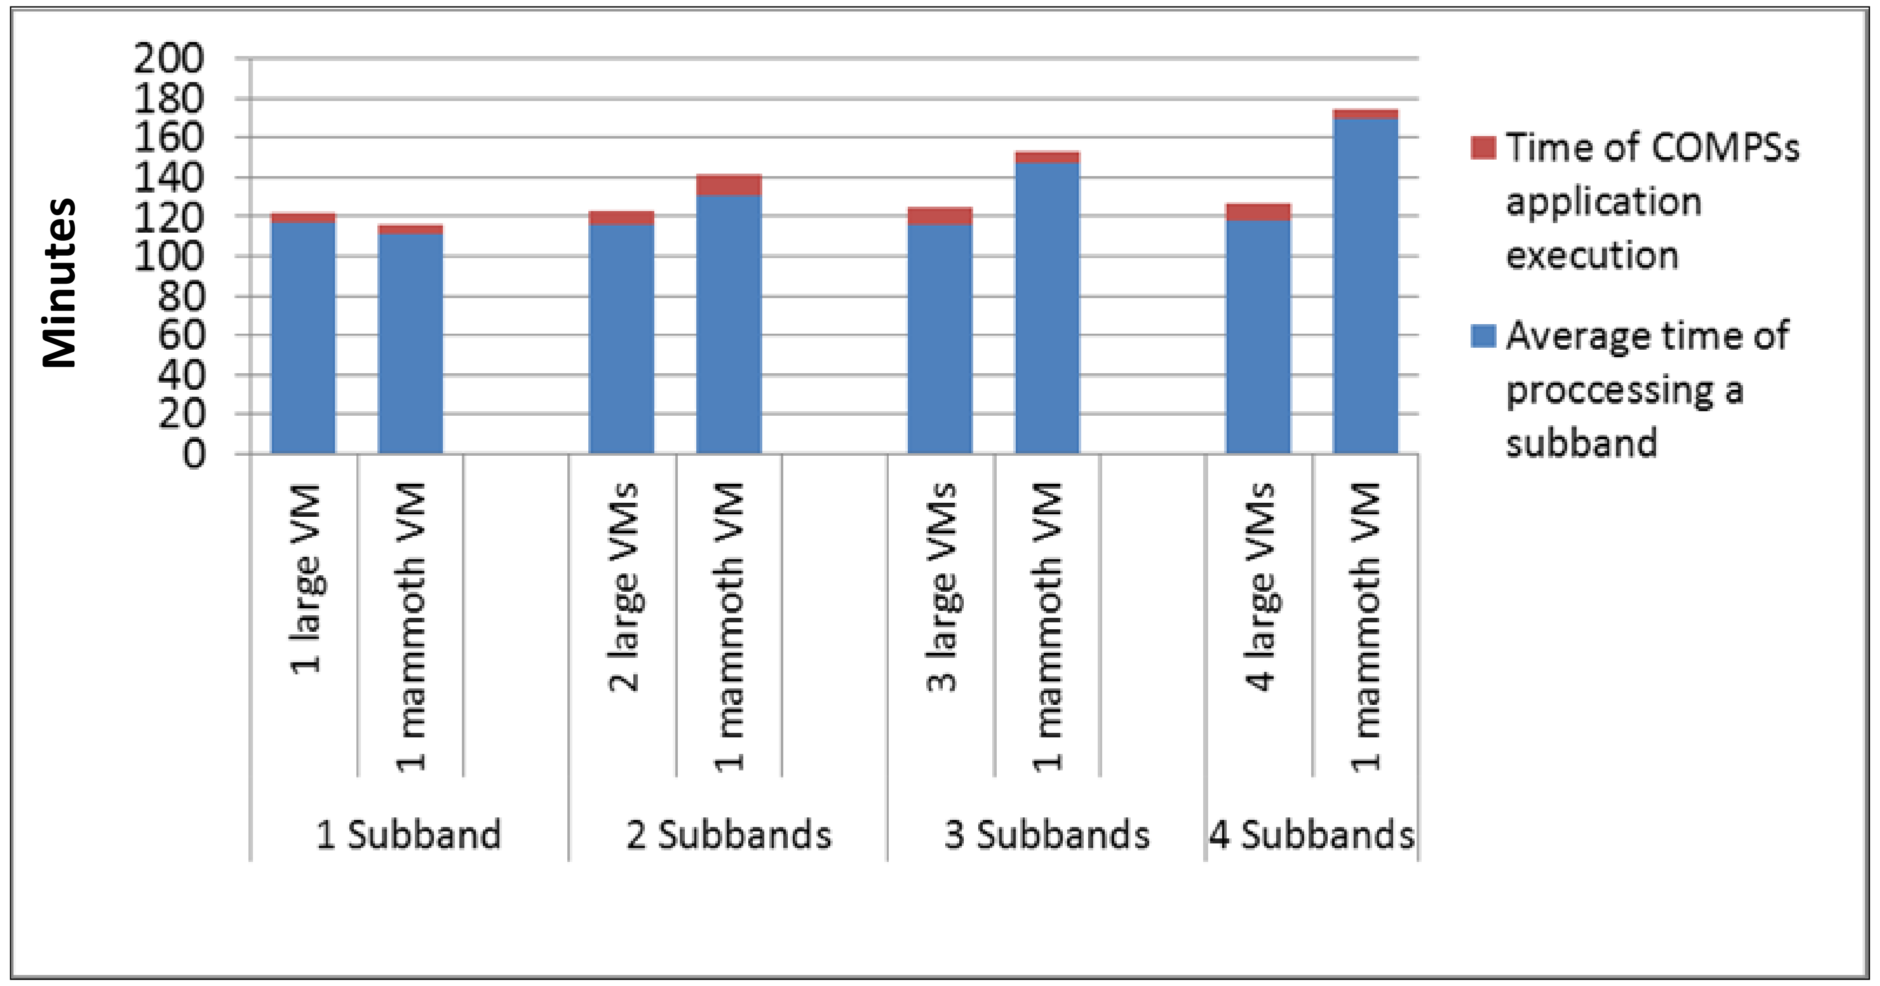

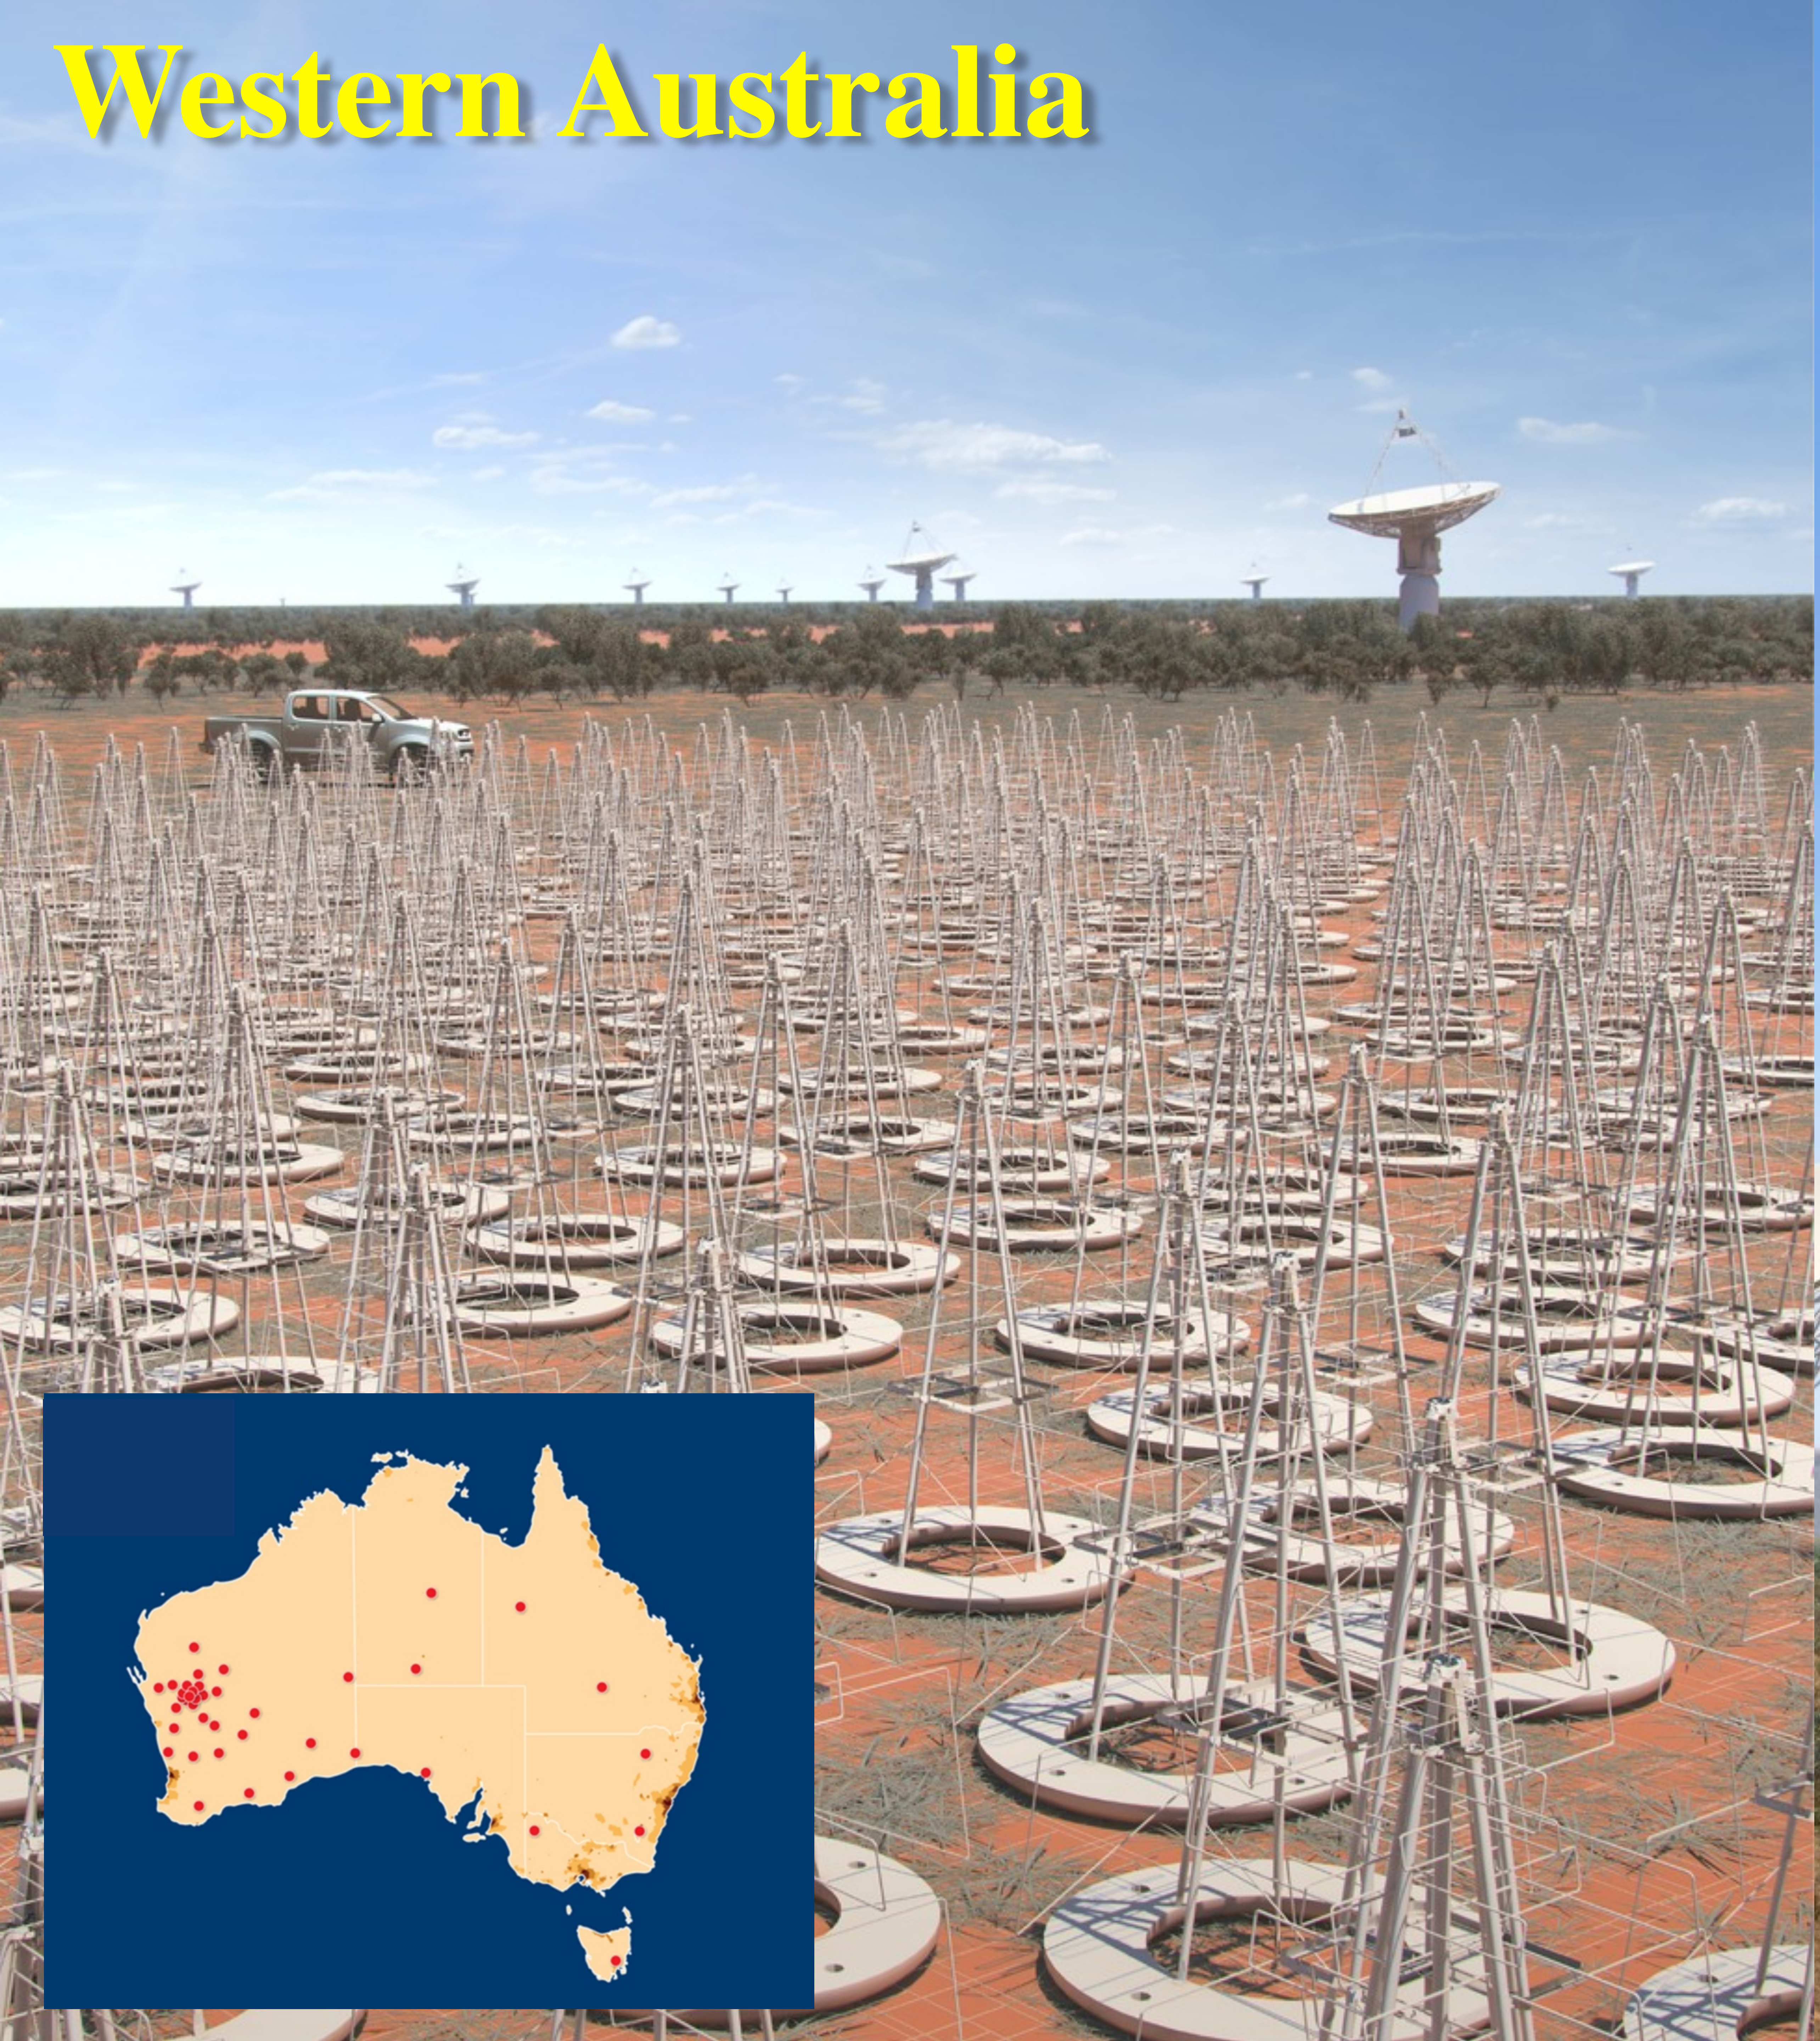
Western Australia

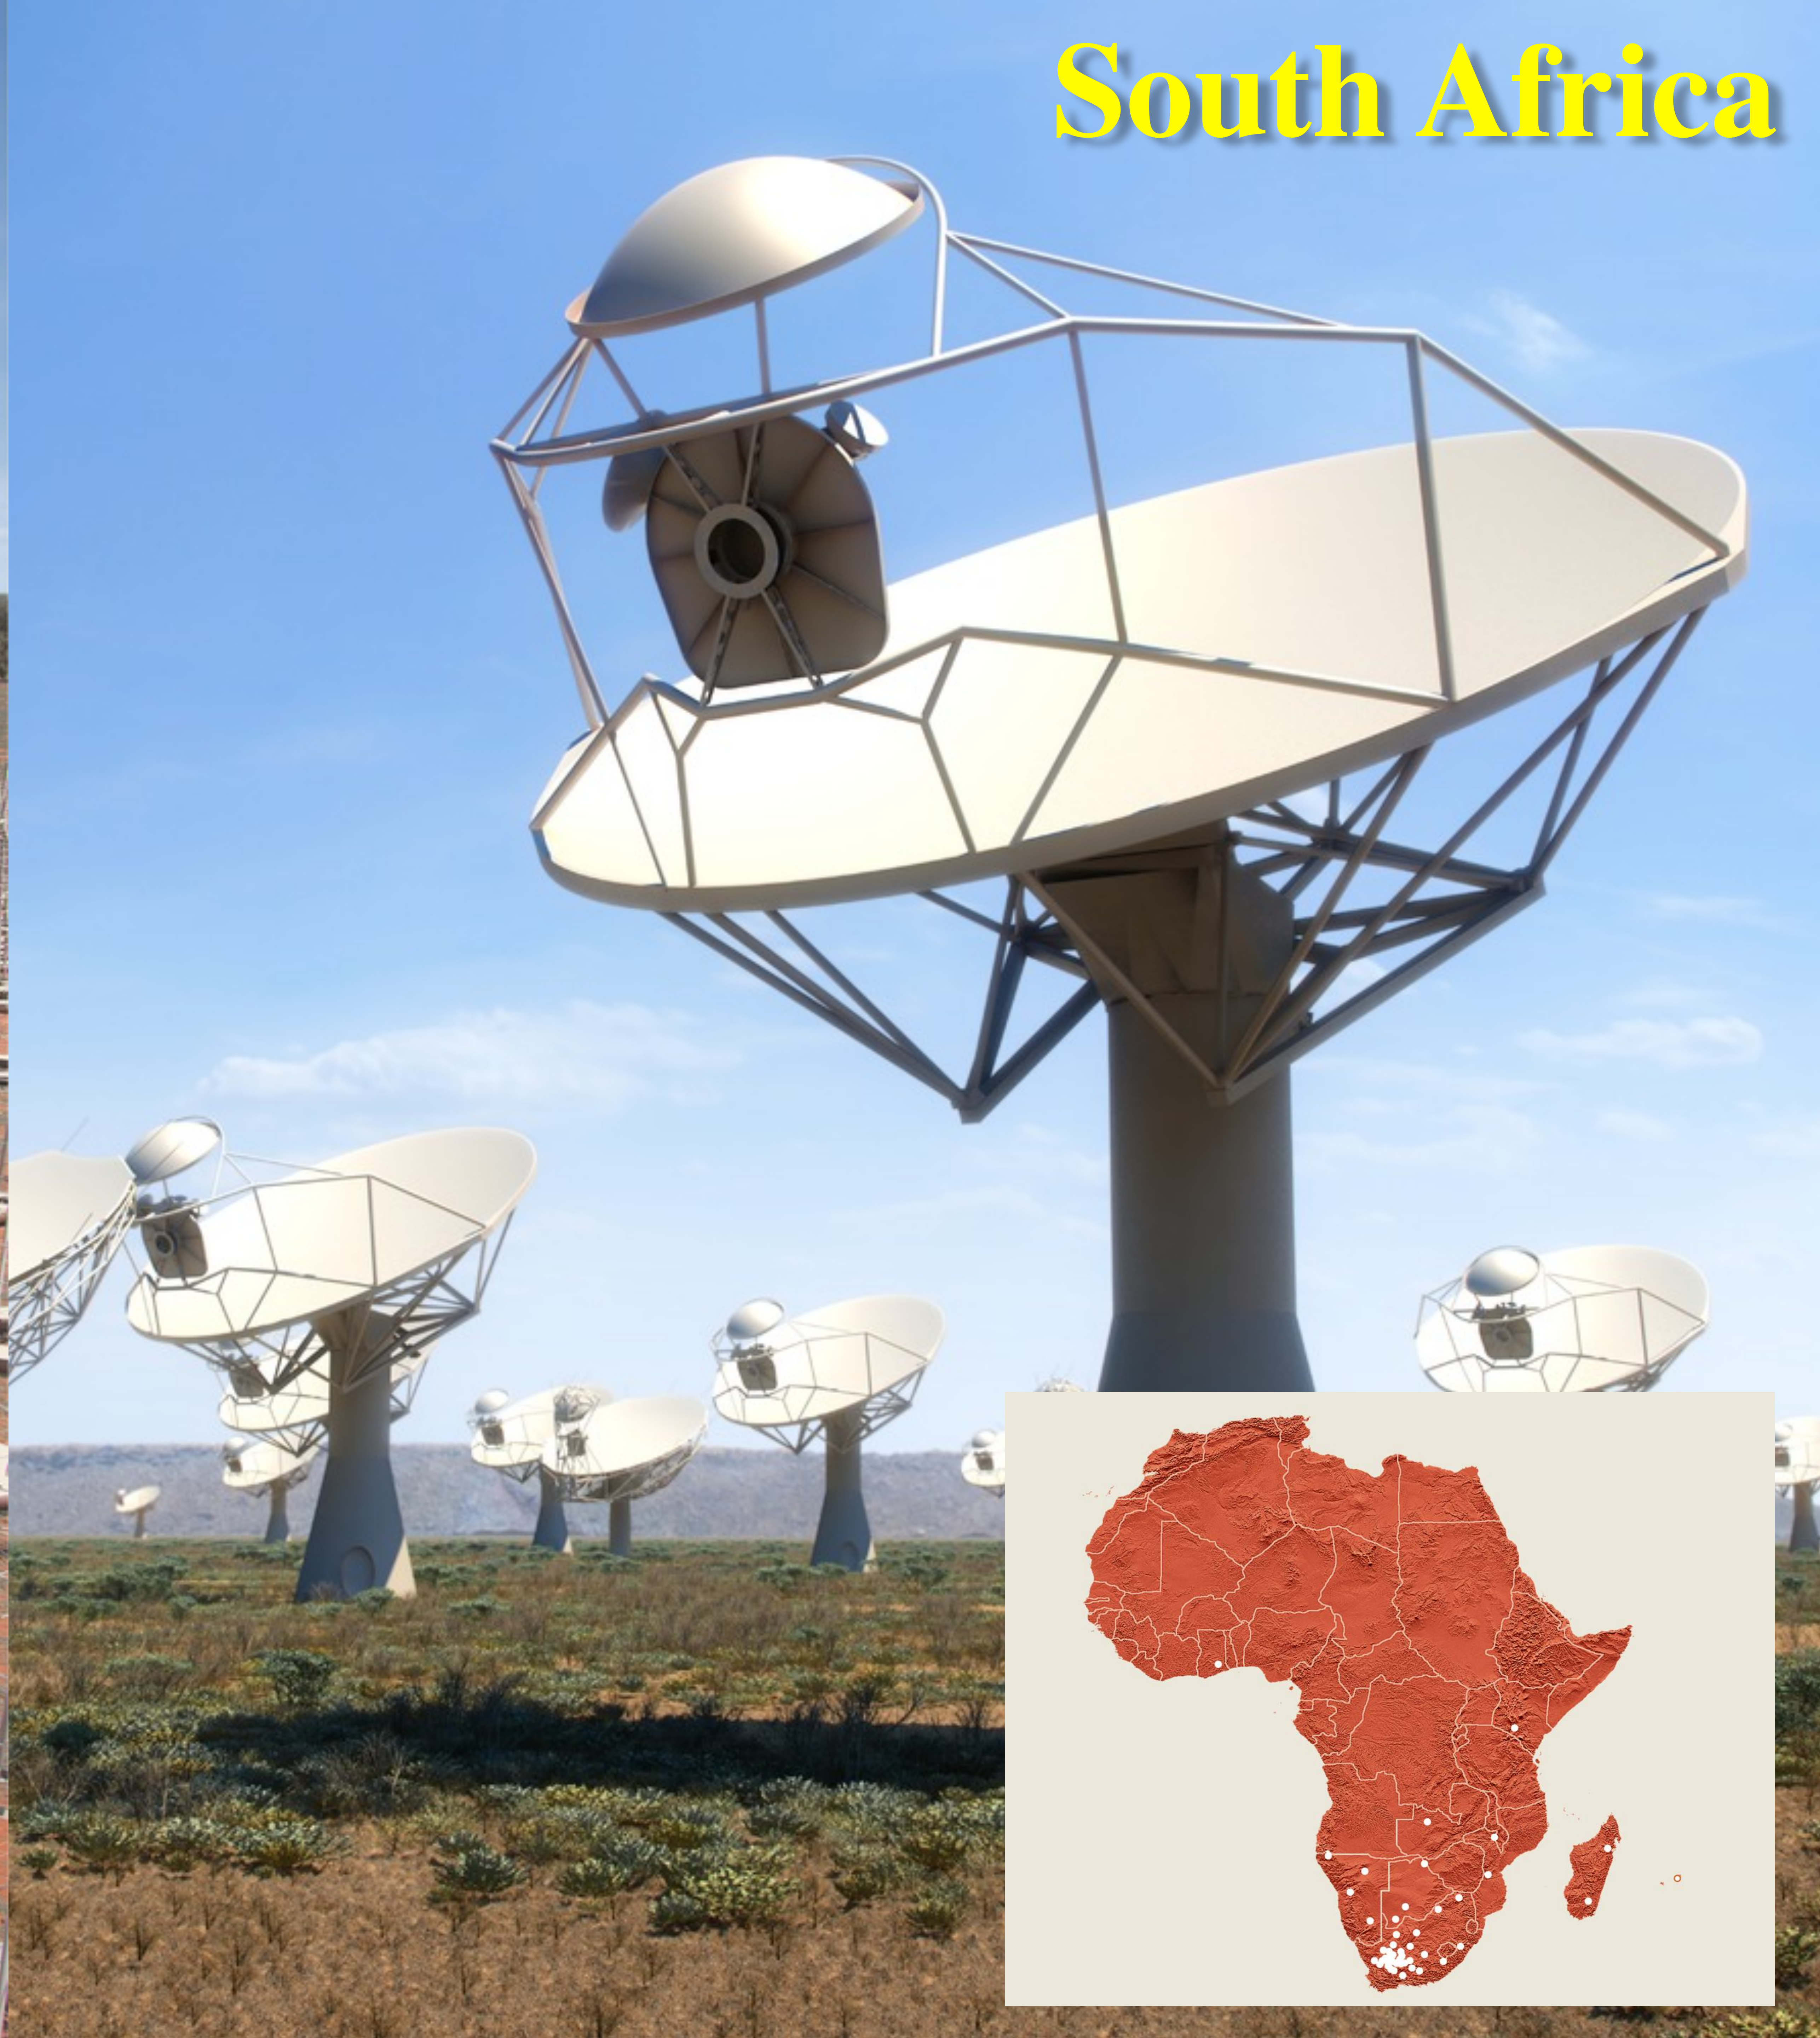
South Africa

SKA System Data Flow

Antennas

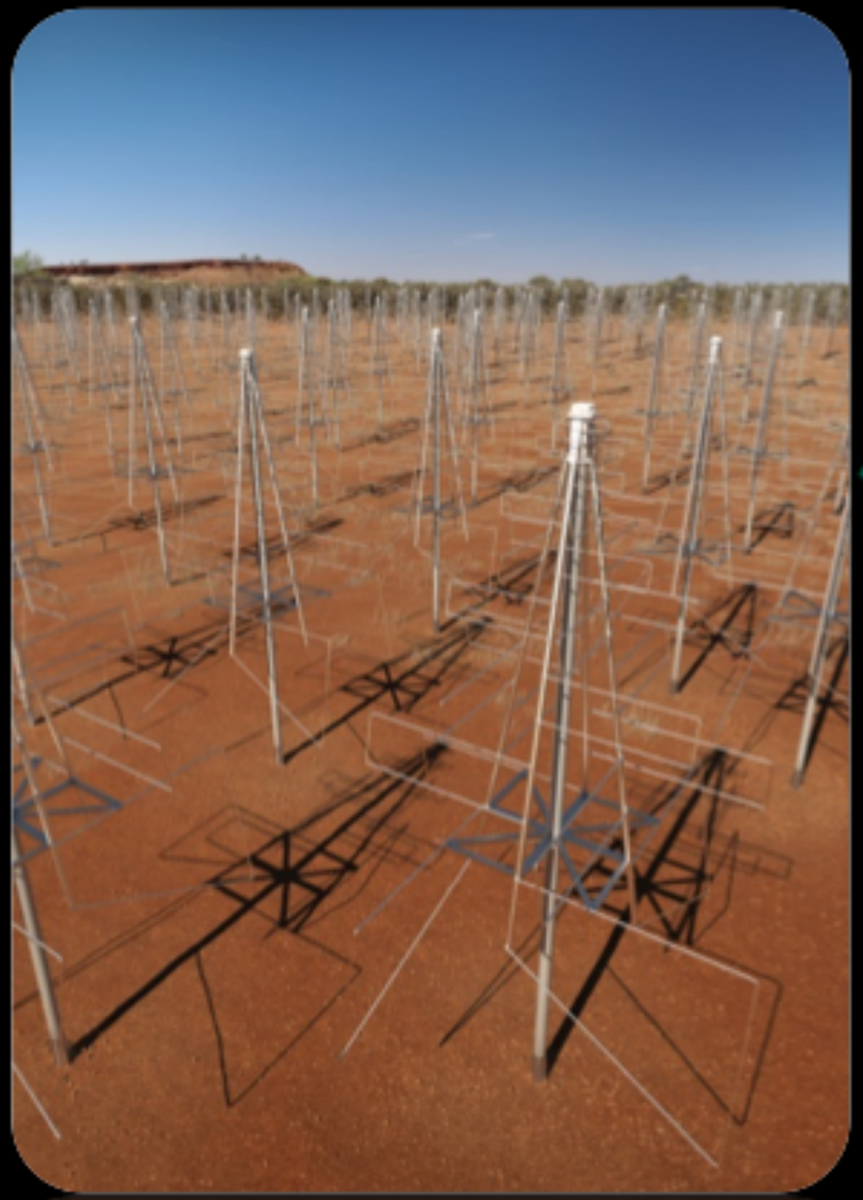

Digital Signal Processing (DSP)

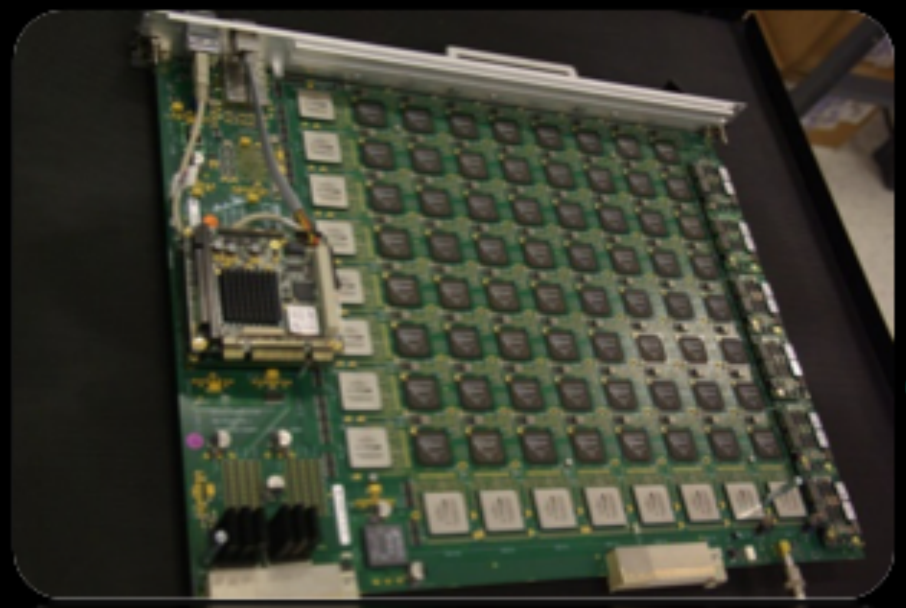

Transfer antennas to DSP
2020: 5,000 PBytes/day
2030: 100,000 PBytes/day

Over 10's to 1000's kms

To Process in HPC
2020: 50 PBytes/day
2030: 10,000 PBytes/day

Over 10's to 1000's kms

High Performance Computing Facility (HPC)

HPC Processing
2020: 300 PFlop
2028: 30 EFlop

30 – 300 Pbytes / year of fully processed science data products

SKA Science Archive

PER YEAR

• 1 Petabyte

Full Cost of Science Extraction

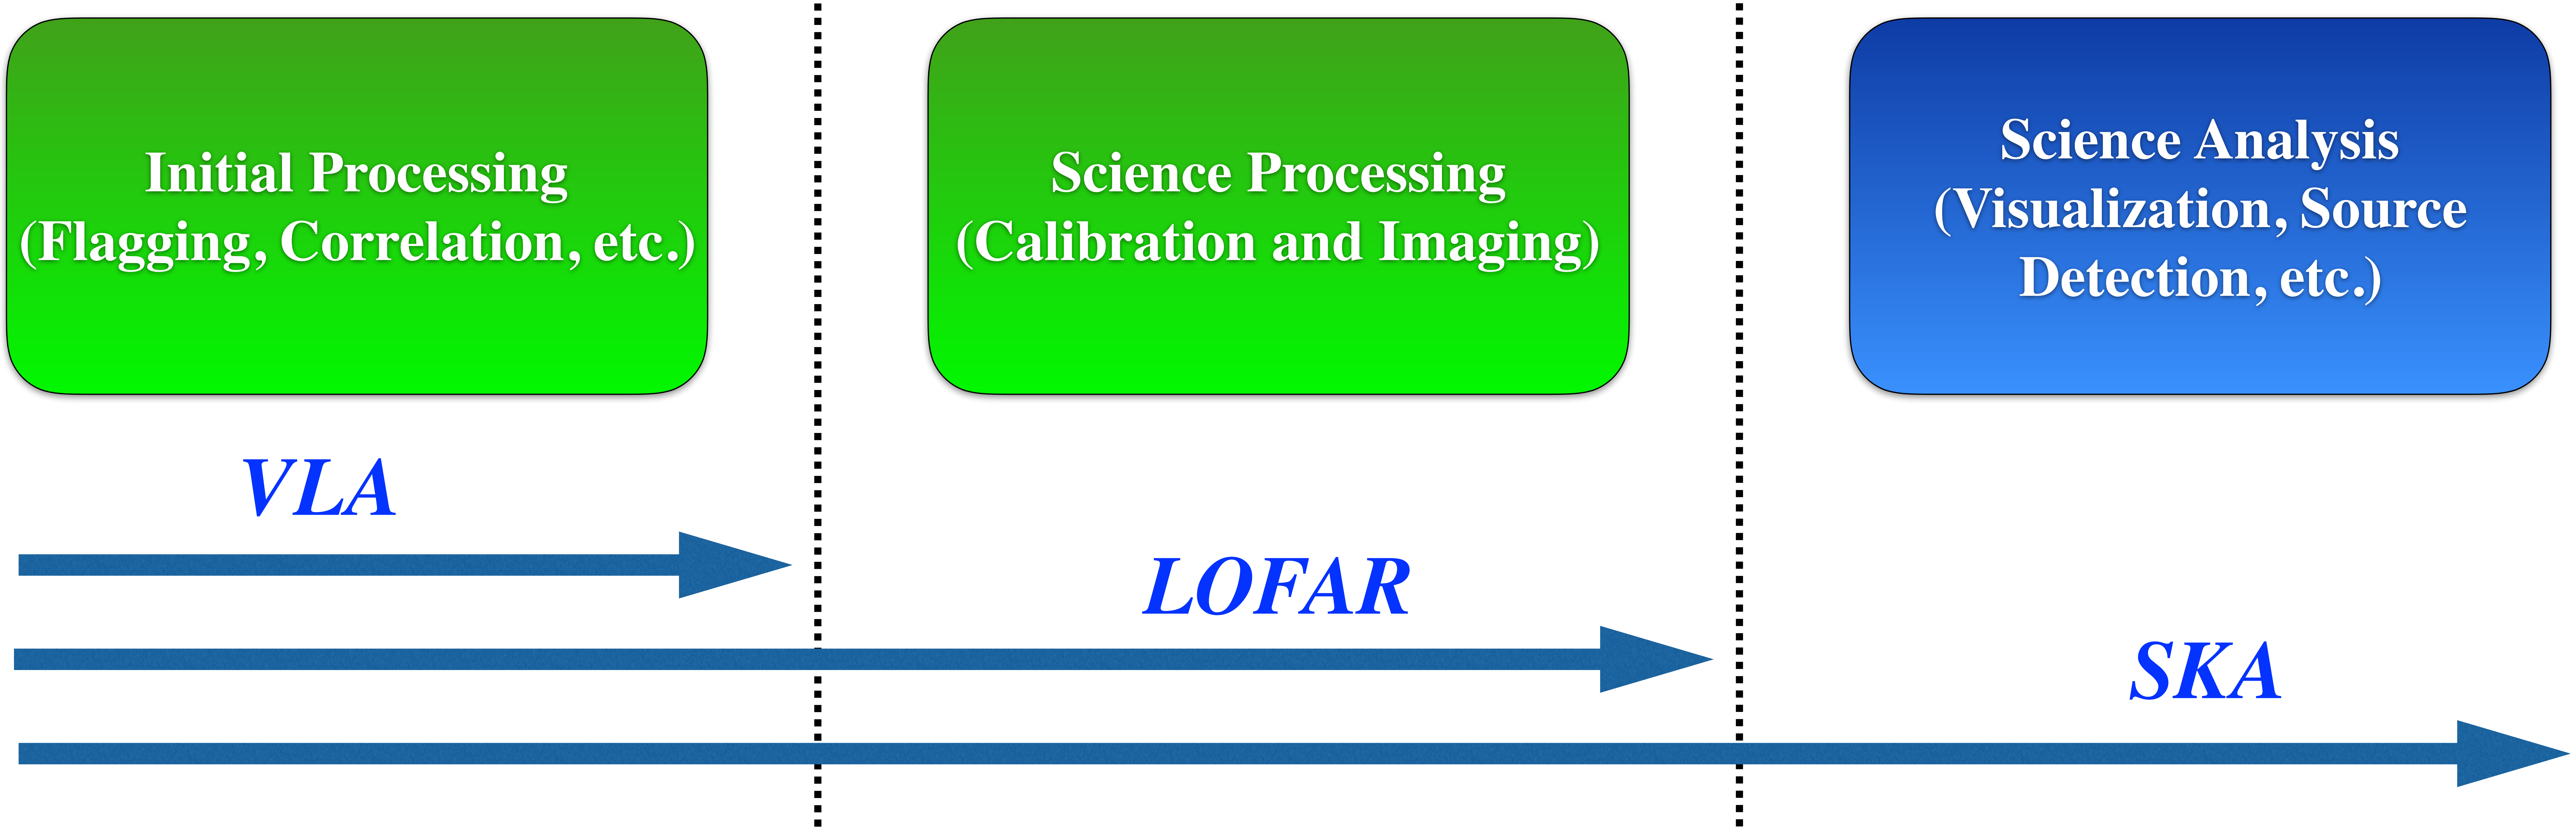

Desktop Systems (simple tools)	***Moderate Clusters (pipelines)***	***Science Data Centres (smart systems)***

Resources for Science Extraction

SKA Data Flow Advisory Panel (DFAP) - March 2016

⇒ *"Within the current construction costs, no provision is made for the distribution of data to users, nor for computational facilities to enable users to undertake further data analysis…"*

⇒ *"Full scientific exploitation of the SKA will require a research ecosystem be in place for efficiently translating the large data volume into science results."*

Not covered in current SKA scope:

- Computational capacity for re-processing and science analysis
- Long-term capacity for archiving standard and derived data products
- Local user support for post-processing and science analysis

SKA Regional Centres

- Science Data Centres (SDCs) will likely host the SKA science archive
- Provide access and distribute data products to users
- Provide access to compute and storage resources for users
- Provide analysis capabilities
- Provide user support
- Multiple national SRCs, locally resourced

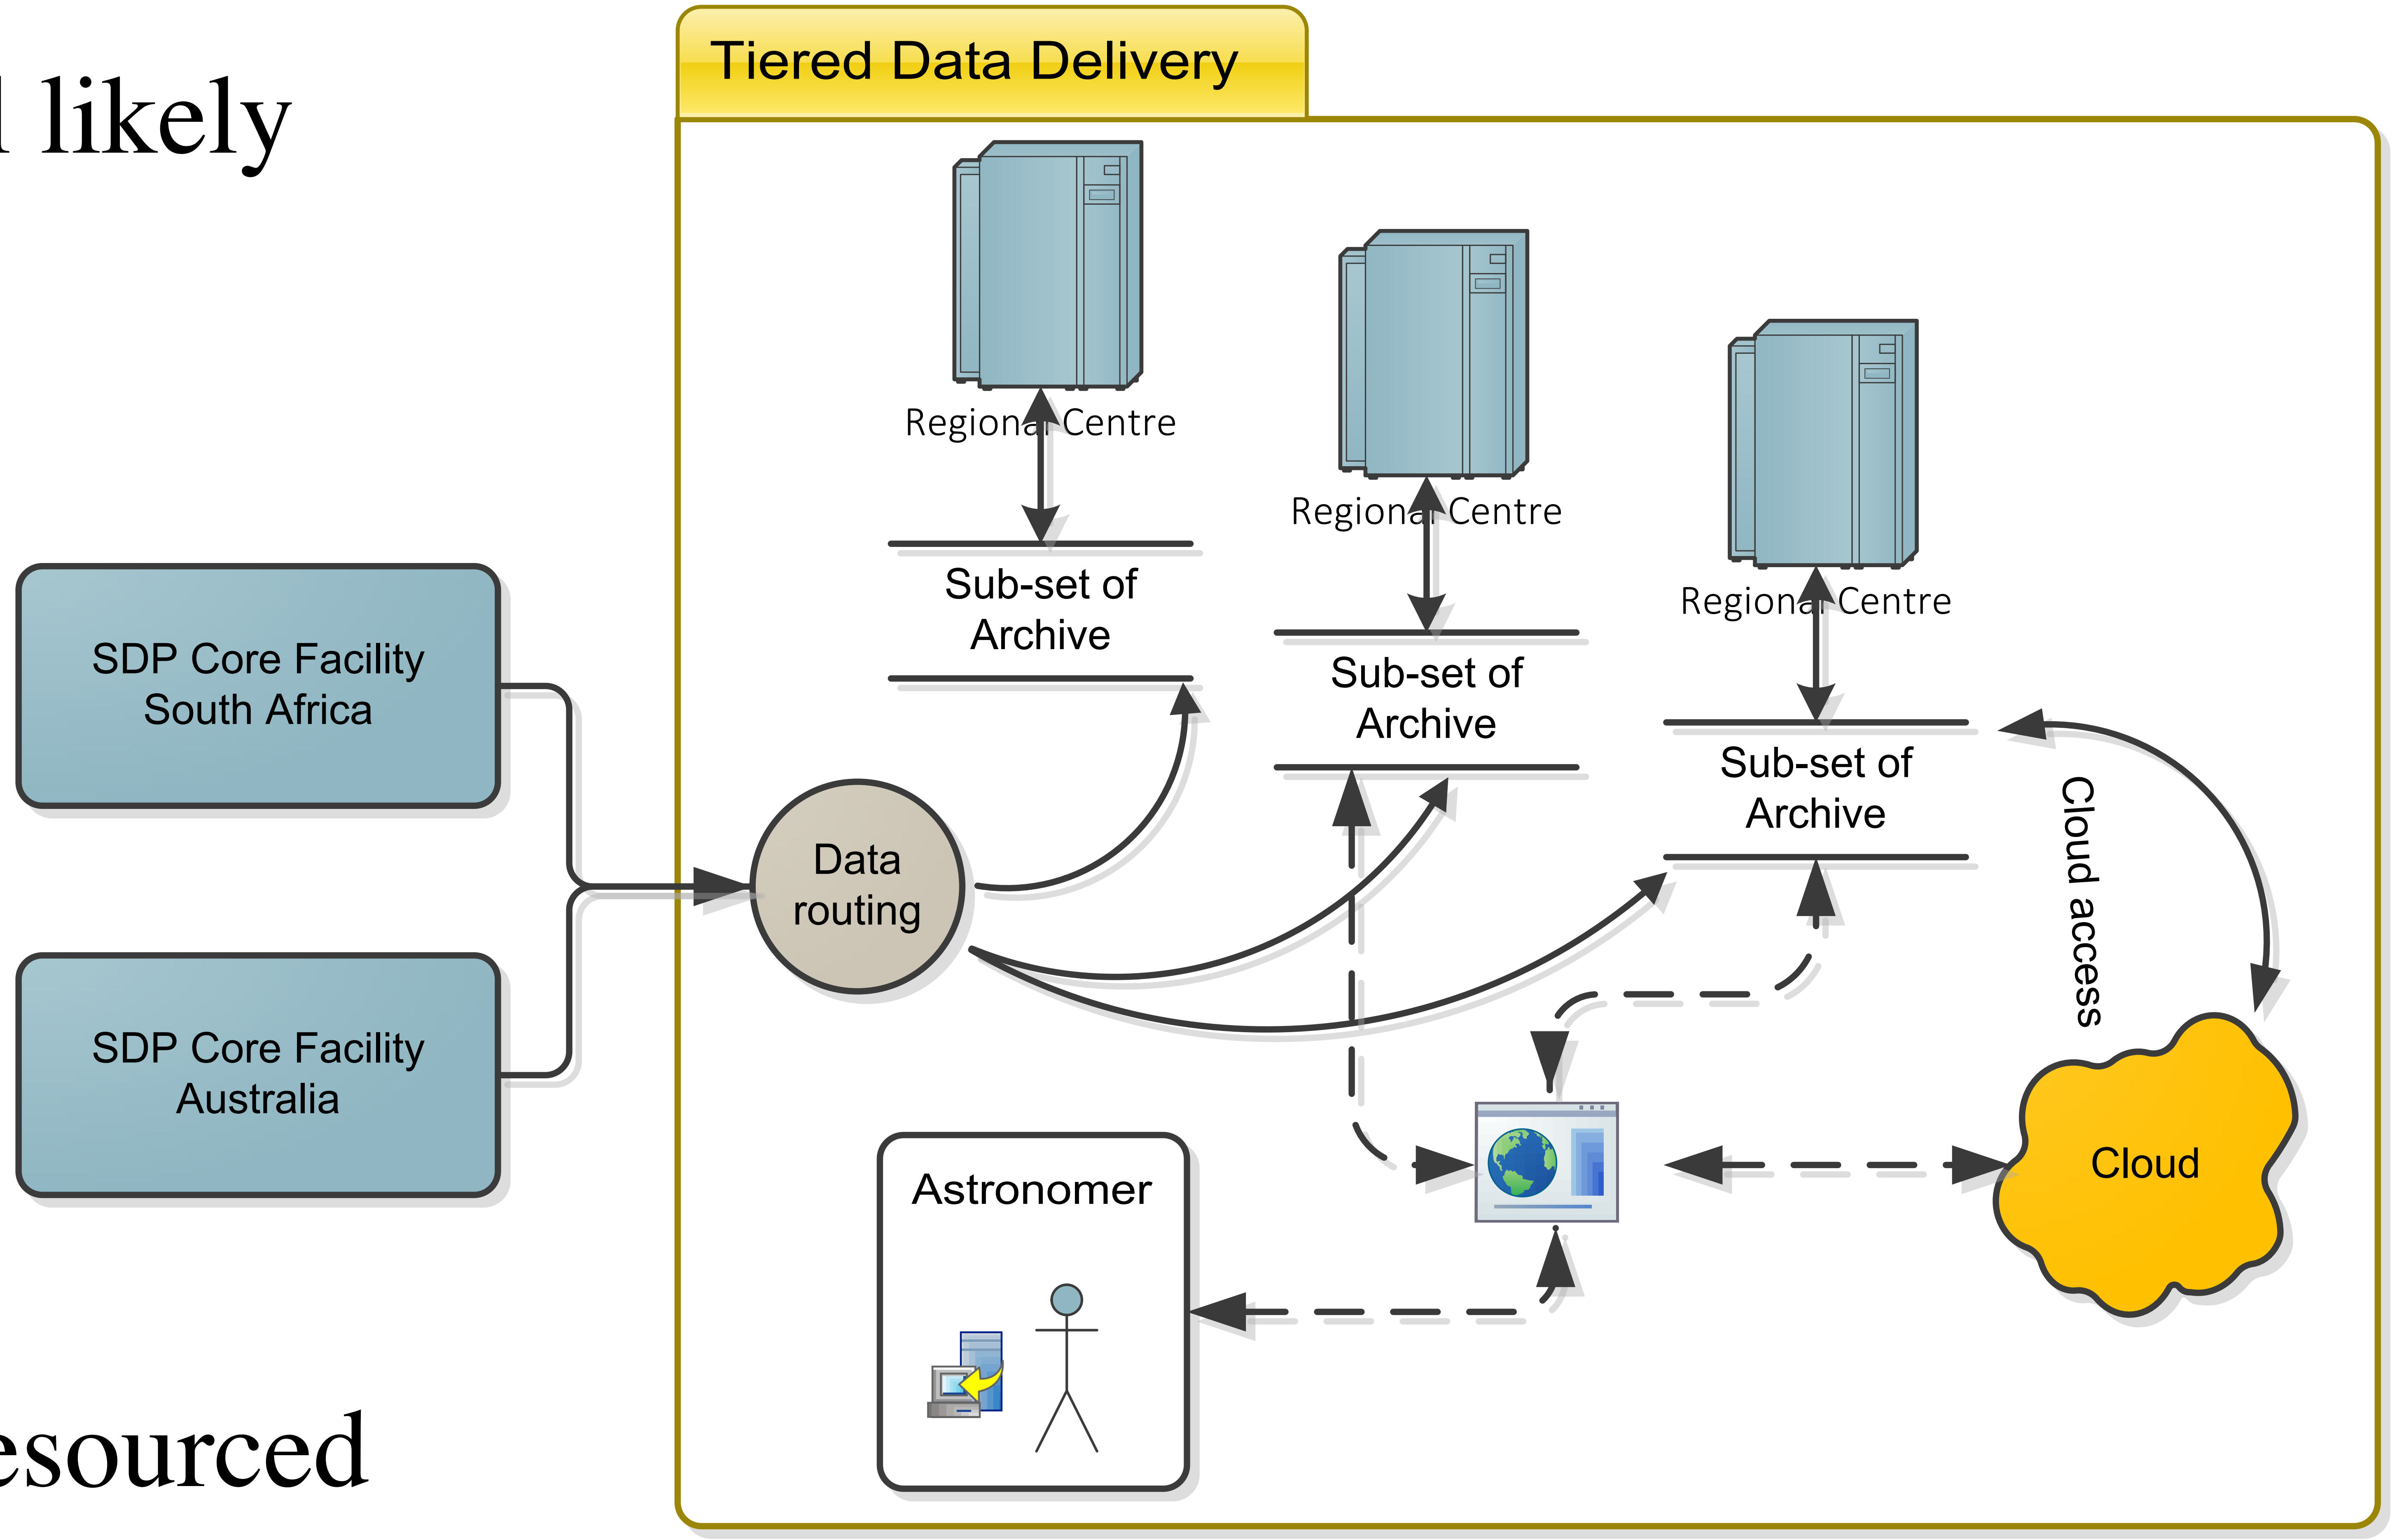

SRCs will be primary interface for astronomers to extract SKA science!

Regional Centres Functionality

Data Discovery

- Observation database
- Associated metadata
- Quick-look data products
- Flexible catalog queries
- Integration with VO tools
- Publish data to VO

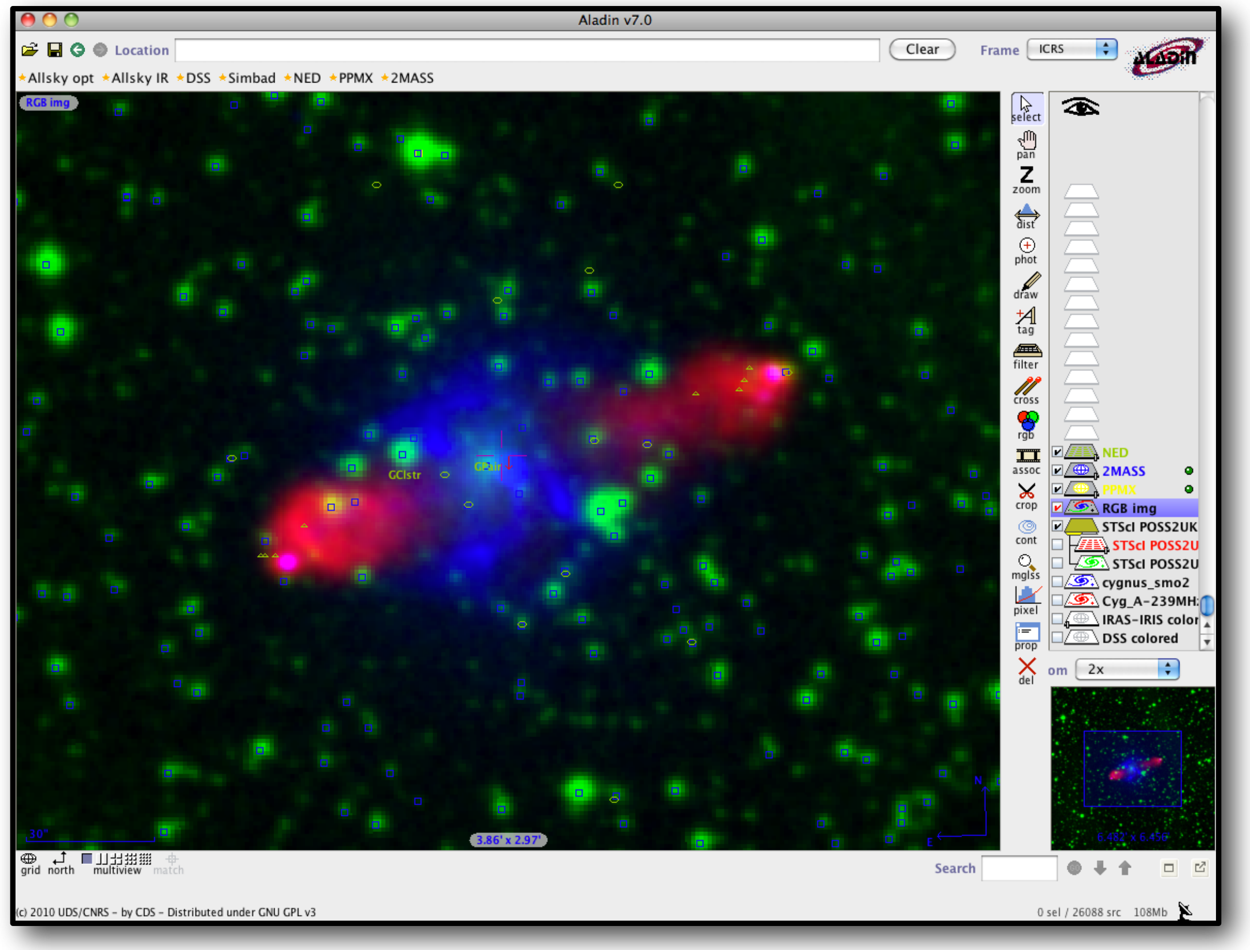

Data Processing

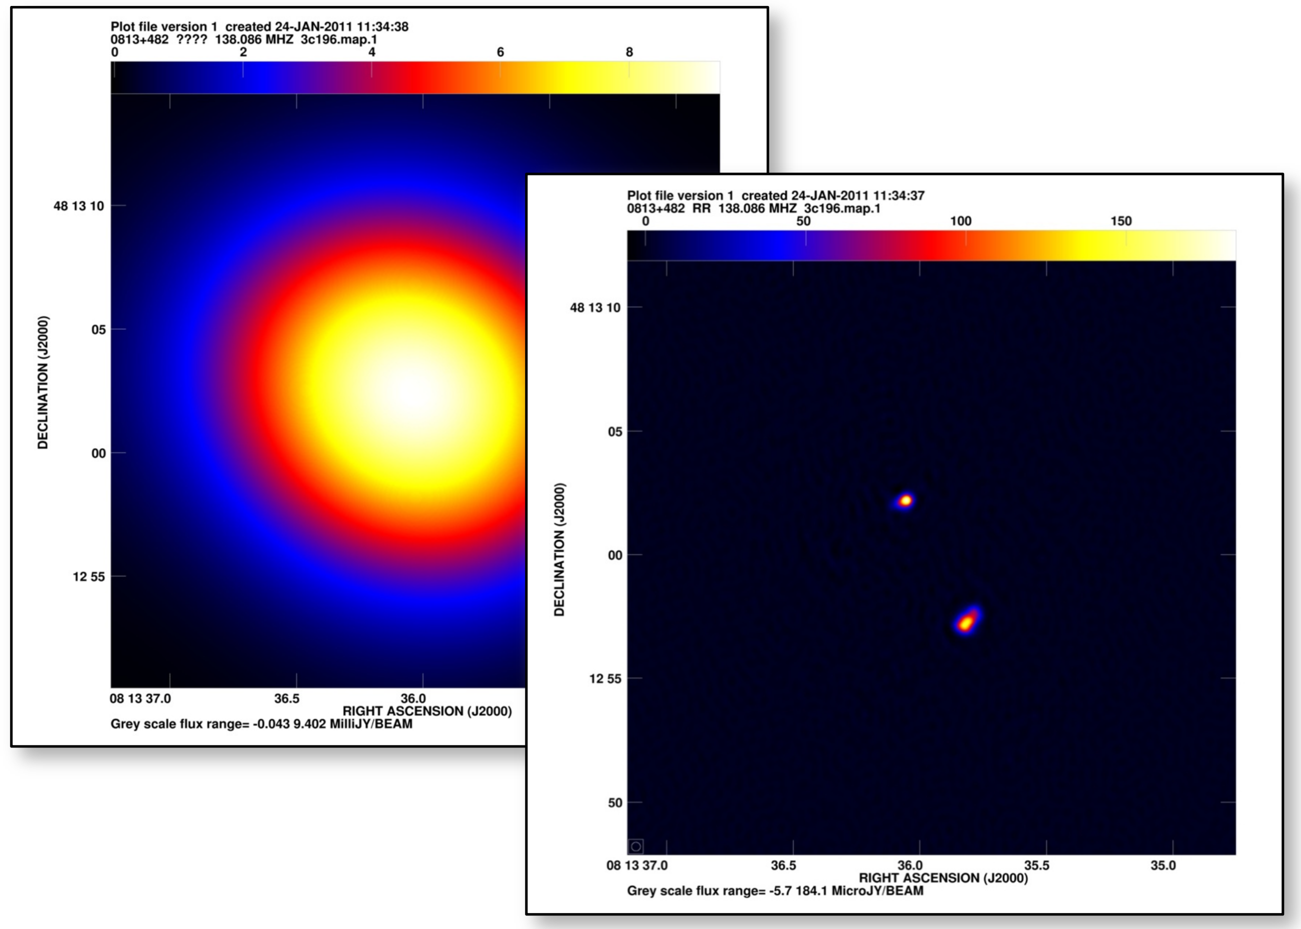

- Reprocessing and calibration
- High resolution imaging
- Mosaicing
- Source extraction
- Catalog re-creation
- DM searches

Data Mining

- Multi-wavelength studies
- Catalog cross-matching
- Light-curve analysis
- Transient classification
- Feature detection
- Visualization

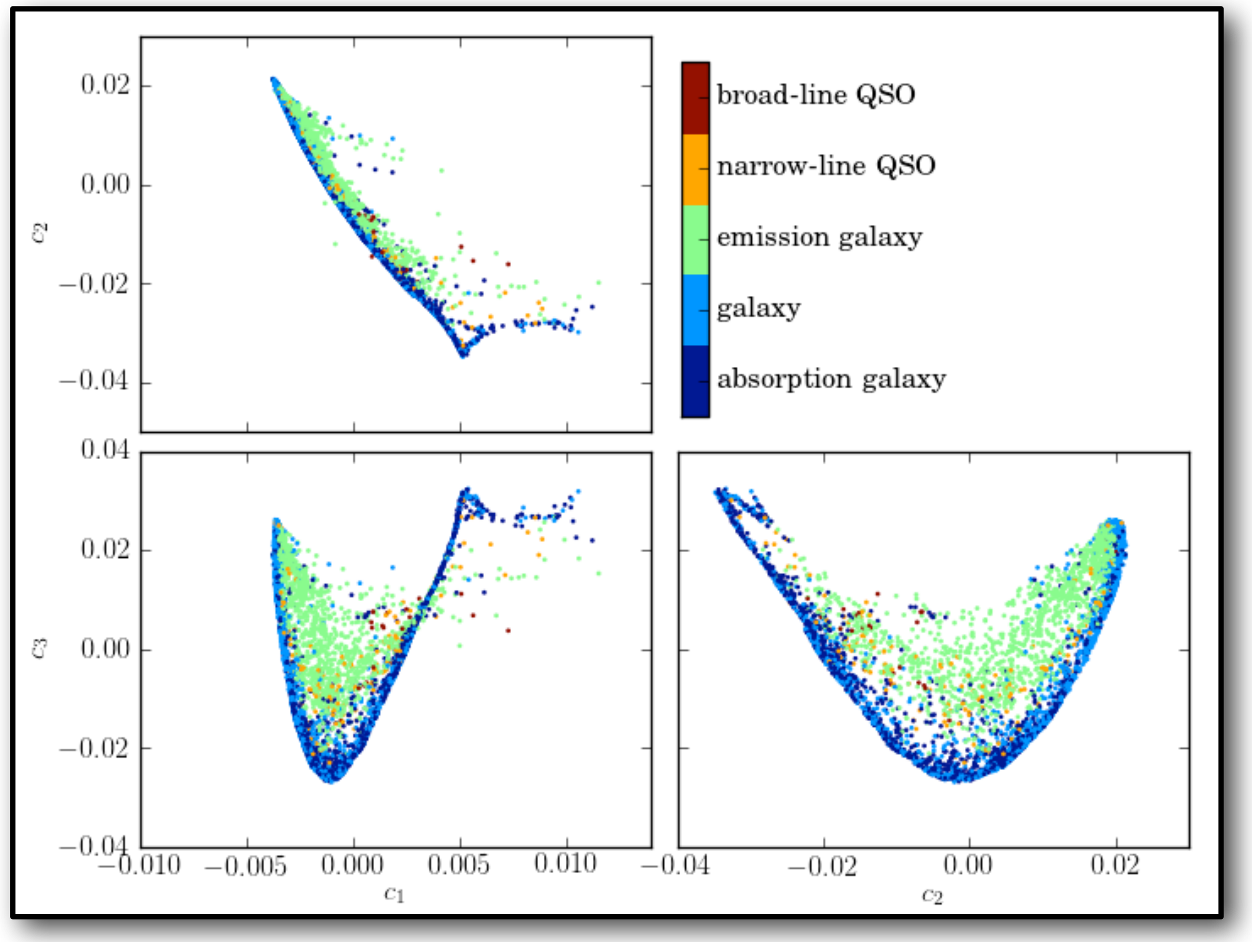

Regional Centres Services

Data oriented operations:

- Data archiving and curation
- Data management, discovery, and access
- Automated processing and reprocessing
- Generation and storage of science products
- Continued pipeline development

Science oriented operations:

- Portal-based data product access
- Interface to processing pipelines
- Interface to VO discovery and analysis tools
- Support for custom user analysis
- Development of new algorithms and tools
- End-to-end astronomer support
- Community education & outreach
- Face-to-face user support
- 24/7 help desk

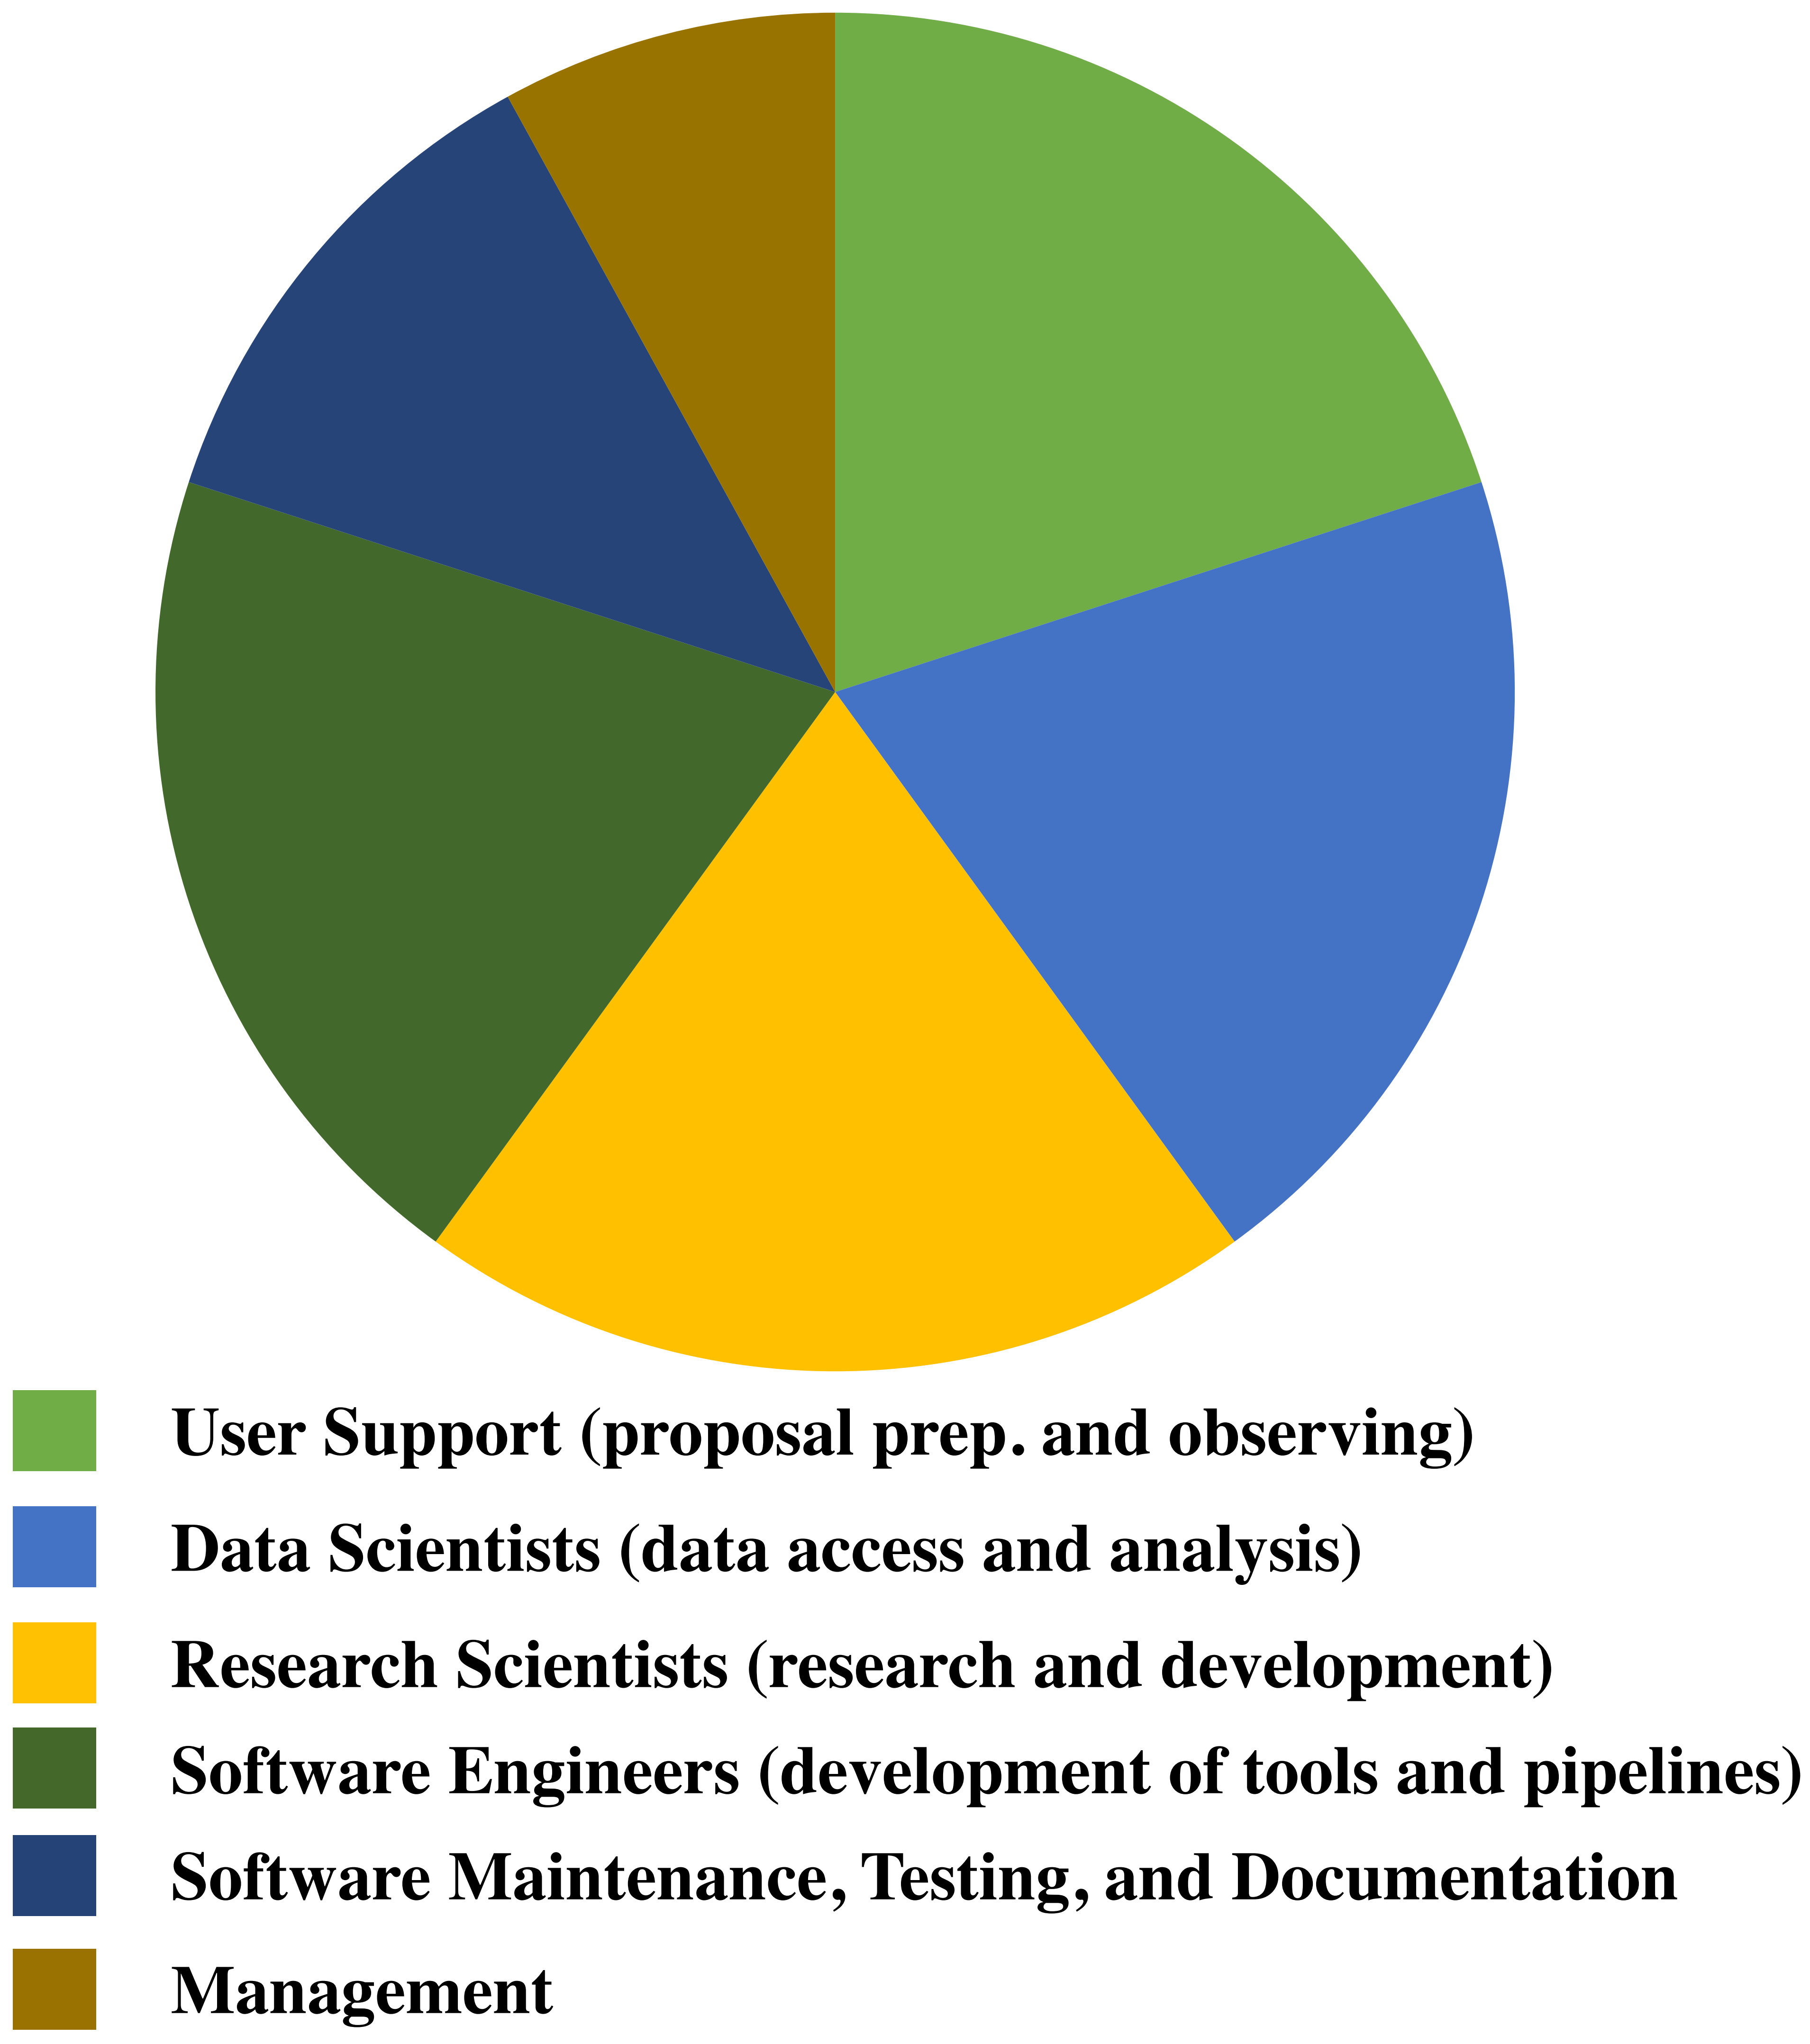

Wider range of user support required

European SKA Regional Centre

- Create a European-scale federated Science Data Center for the SKA
- Provide resources for SKA science extraction to users
- Coordinated engagement with national ICT communities, industry, and service providers
- Facilitate shared development, interoperability, accessibility and innovation
- European counterpart for engagement with other Science Data Centres internationally

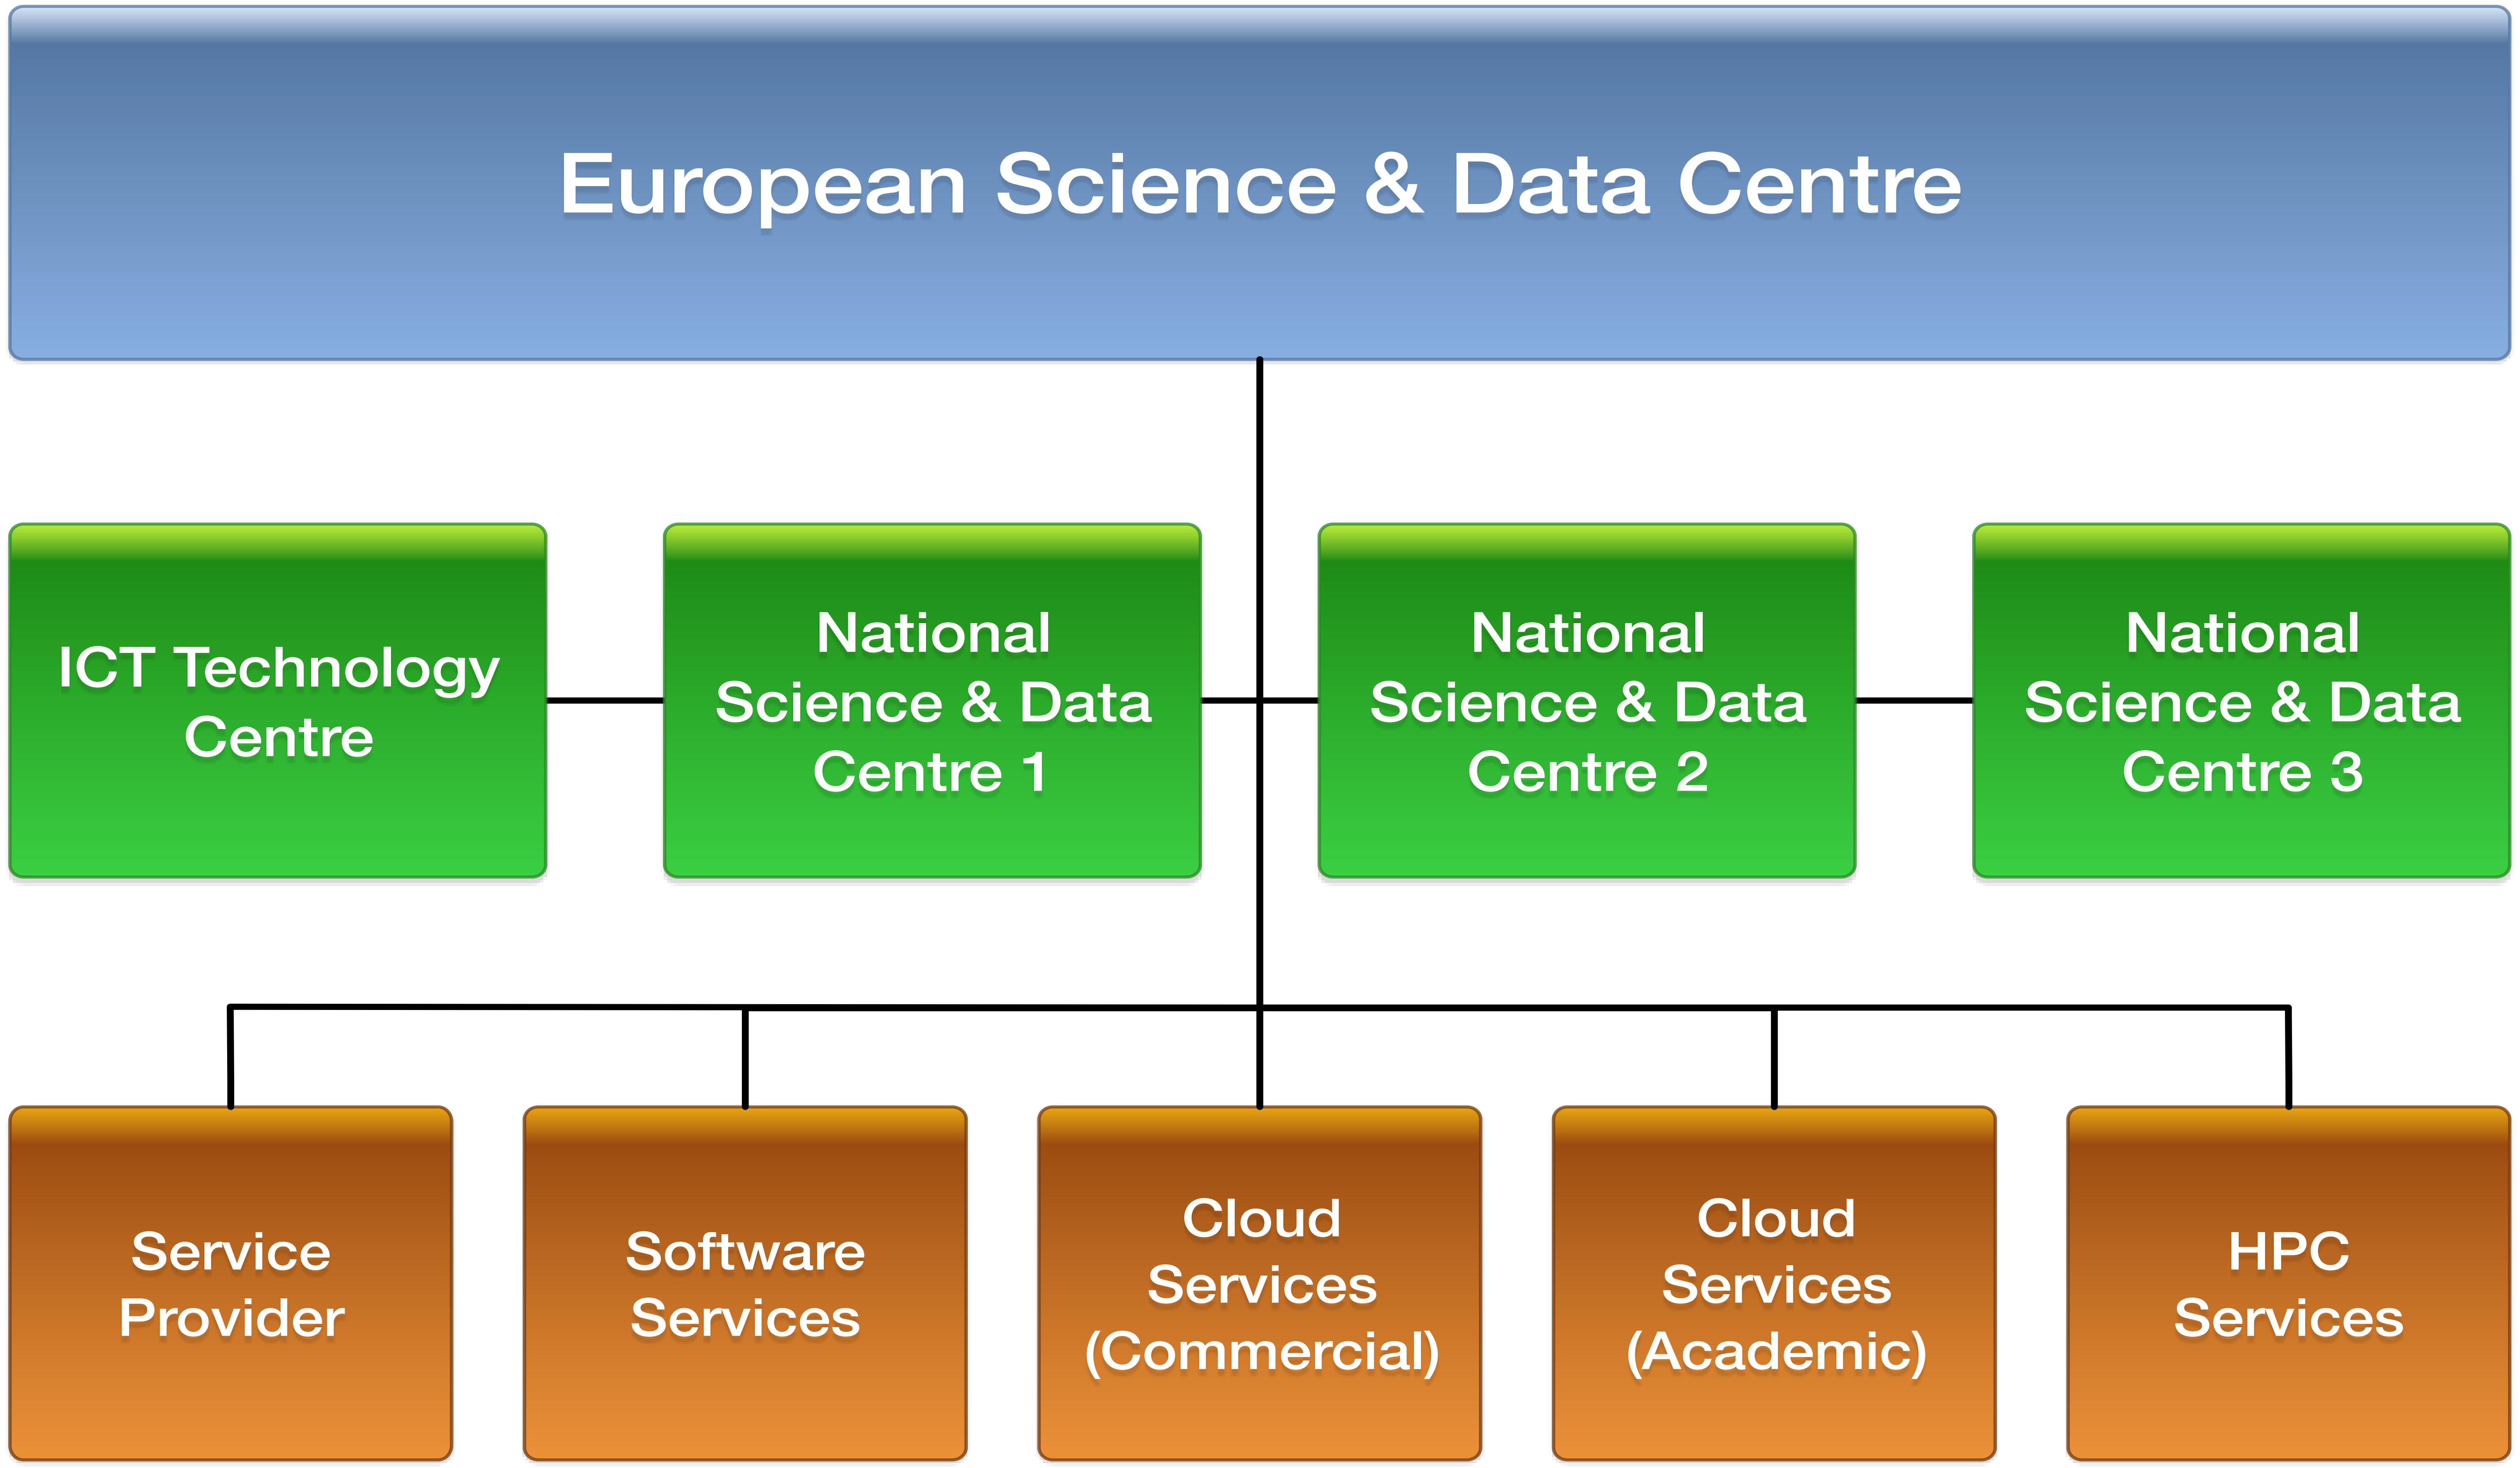

The AENEAS Project

Advanced European Network of E-infrastructures
for Astronomy with the SKA

Design and specification of a distributed, European Science Data Centre (ESDC) to support the pan-European astronomical community in achieving the scientific goals of the SKA

EC Horizon 2020 (€3 million)

13 countries, 28 partners, SKAO, host countries, e-infrastructures (EGI, GÉANT, RDA), NREN's

Three year project (2017-2019)

- WP2: ESDC Governance Structure and Business Models
- WP3: ESDC Computing and Processing Requirements
- WP4: SKA Data Transport and Optimal European Storage Topologies
- WP5: User Data Access and Knowledge Creation

Global Network of Centres

Summary

- *Computing resources to extract science are a bottleneck*
- *Need to consider the full cost of science extraction*
- *Current SKA scope does not cover data delivery or analysis*
- *SKA Regional Centres can provide these capabilities*
- *Scientific success of the SKA will depend on these centres*