SKAO

one observatory
two telescopes
three continents

SKA Observatory & Regional Centre Software development using IRIS

Rohini Joshi, Ben Mort, Piers Harding (SKAO)

IRIS TWG Meeting 26-04-22

Overview

- **The SKA Observatory**
- **SKA Observatory Software**
  - Organisation, architecture, & schedule
  - Challenges
  - Use of IRIS resources
- **SKA regional Centre Activities**

Software will drive all aspects of the SKA observatory …

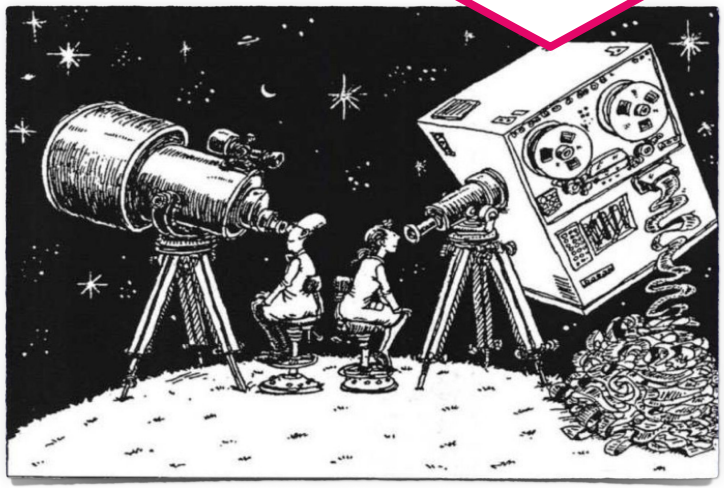

The challenge of operating a high-availability, multi-site system, capturing and processing unprecedented data rates makes use of resources such as IRIS critical to the success of delivering SKA science data products for our users

The SKA Observatory

Our mission is to build and operate cutting-edge radio telescopes to transform our understanding of the Universe and deliver benefits to society through global collaboration and innovation.

Observatory Headquarters
UK: Cheshire

MID Frequency Array
South Africa: Karoo & Cape Town

LOW Frequency Array
Australia: Murchison & Perth

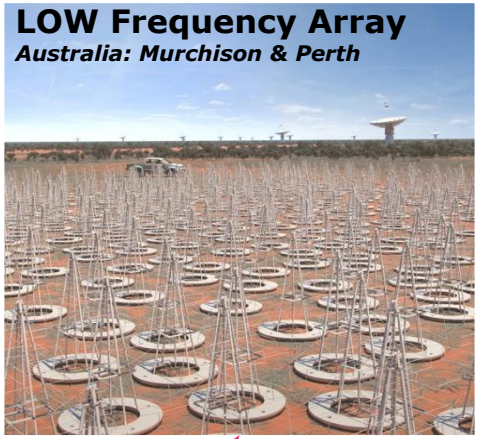

The SKA Observatory

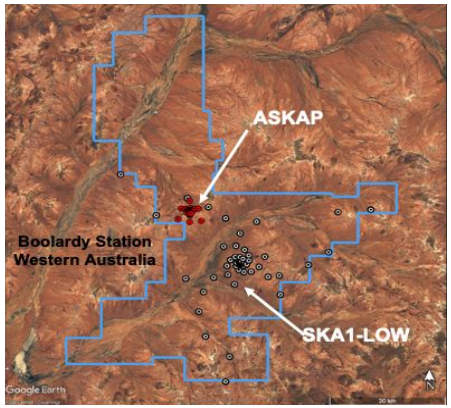

SKA LOW

- 512 aperture array stations (256 antennas)
- 50 - 350 MHz
- 65 km max baseline

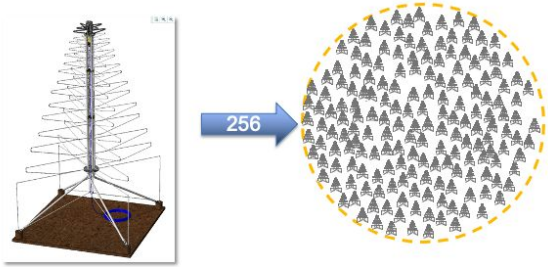

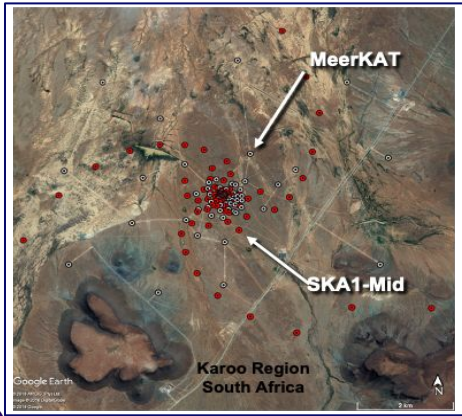

SKA MID

- 133 x 15m (SKA) + 64 x 13.5m (MeerKat) dishes
- 0.35 - 15.4 GHz
- 150 km max baseline

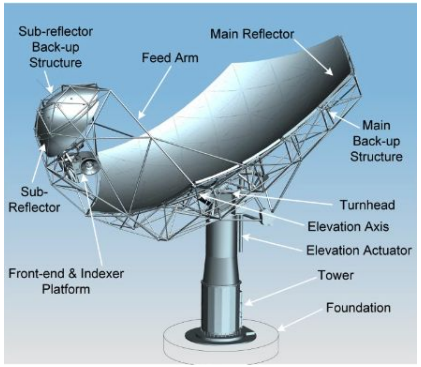

SKA Science!


Testing General Relativity
(Strong Regime, Gravitational Waves)
Cosmic Dawn
(First Stars and Galaxies)
Broadest range of science of any facility, worldwide!
Galaxy Evolution
(Normal Galaxies z~2-3)
Cradle of Life
(Planets, Molecules, SETI
Cosmology
(Dark Energy, Large Scale Structure)
Cosmic Magnetism
(Origin, Evolution)
Exploration of the Unknown

SKA Observatory Software

Organisation, architecture, schedule

SKA Software development

- Construction started in June 2021
- Coordinated effort of 21 globally distributed teams!
- Following the Scaled Agile Framework (SAFe™)
- Three Agile Release Trains (ARTs)
- Planning increment of 3-months with features managed in JIRA
- Work shared with stakeholders through regular series of system demos

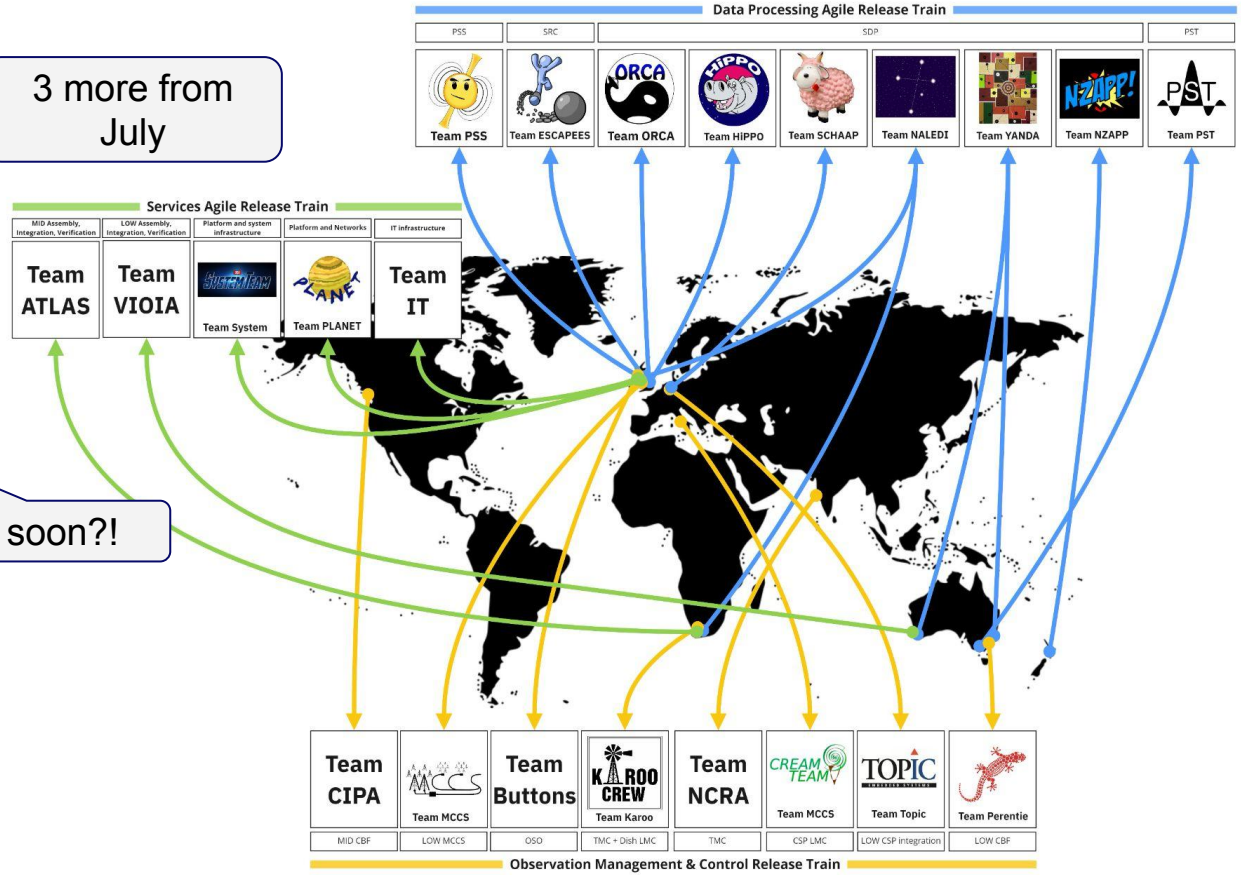

SKA Data Flow

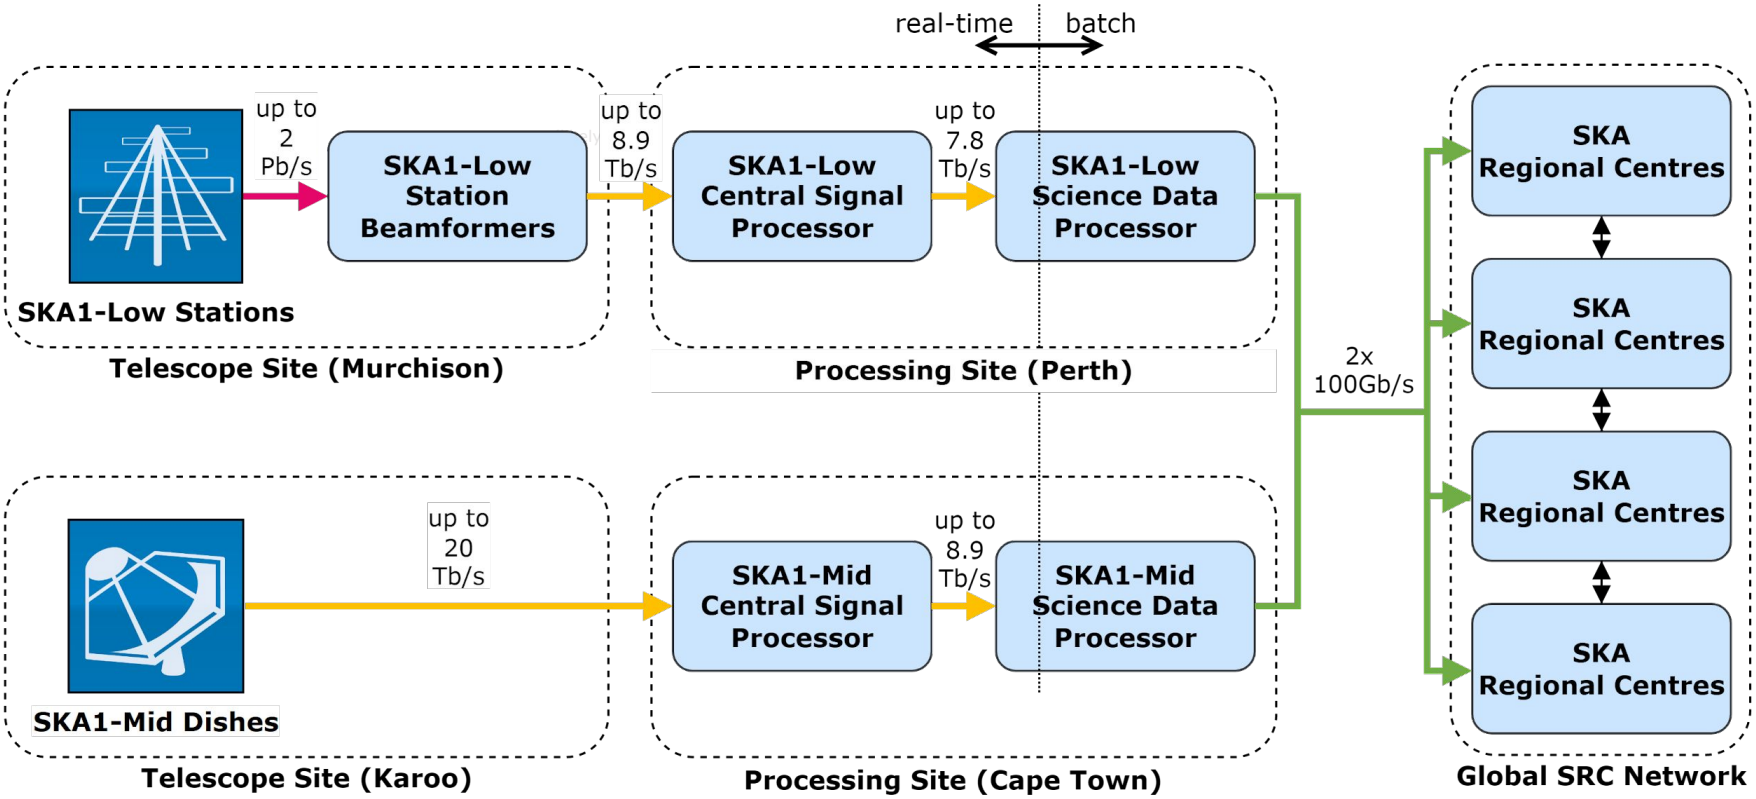

SKA Software - Key components

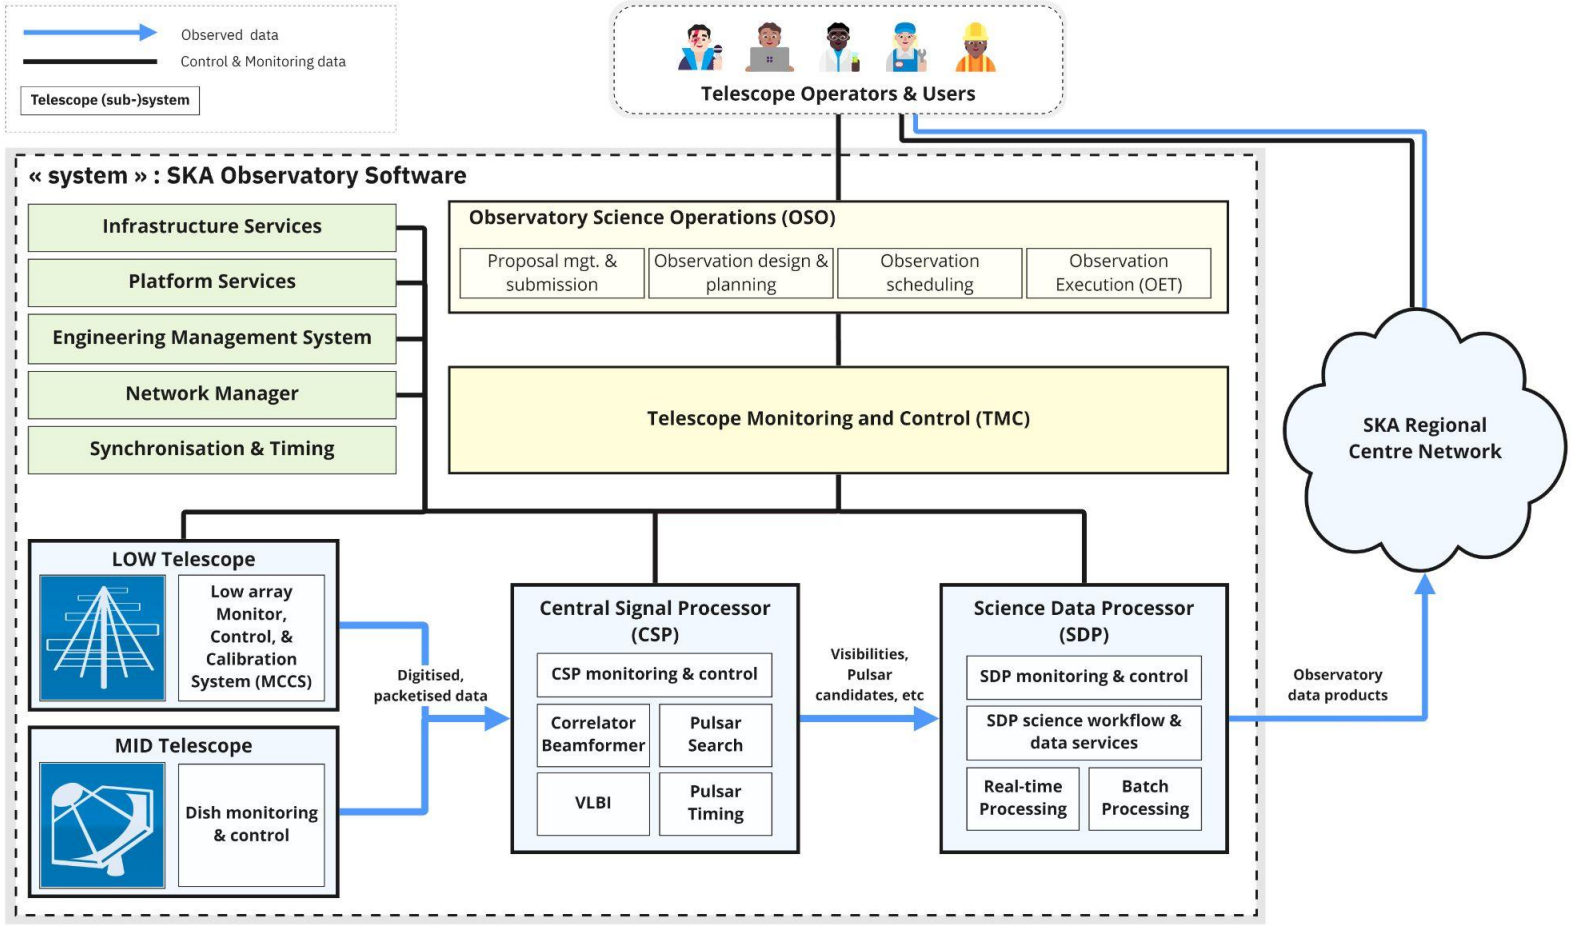

The Target Platform – Cloud Native

Universal Application Abstraction

kubernetes

Abstract Network and Storage

Calico

Rook.io

Universal Storage

ceph

Universal Logging and Monitoring

Prometheus

ElasticStack

Continuous Integration & Deployment

GitLab CI

Cloud Infrastructure Layer

openstack.

RED HAT ANSIBLE Automation

CI/CD Pipeline - GitLab and Kubernetes Centric

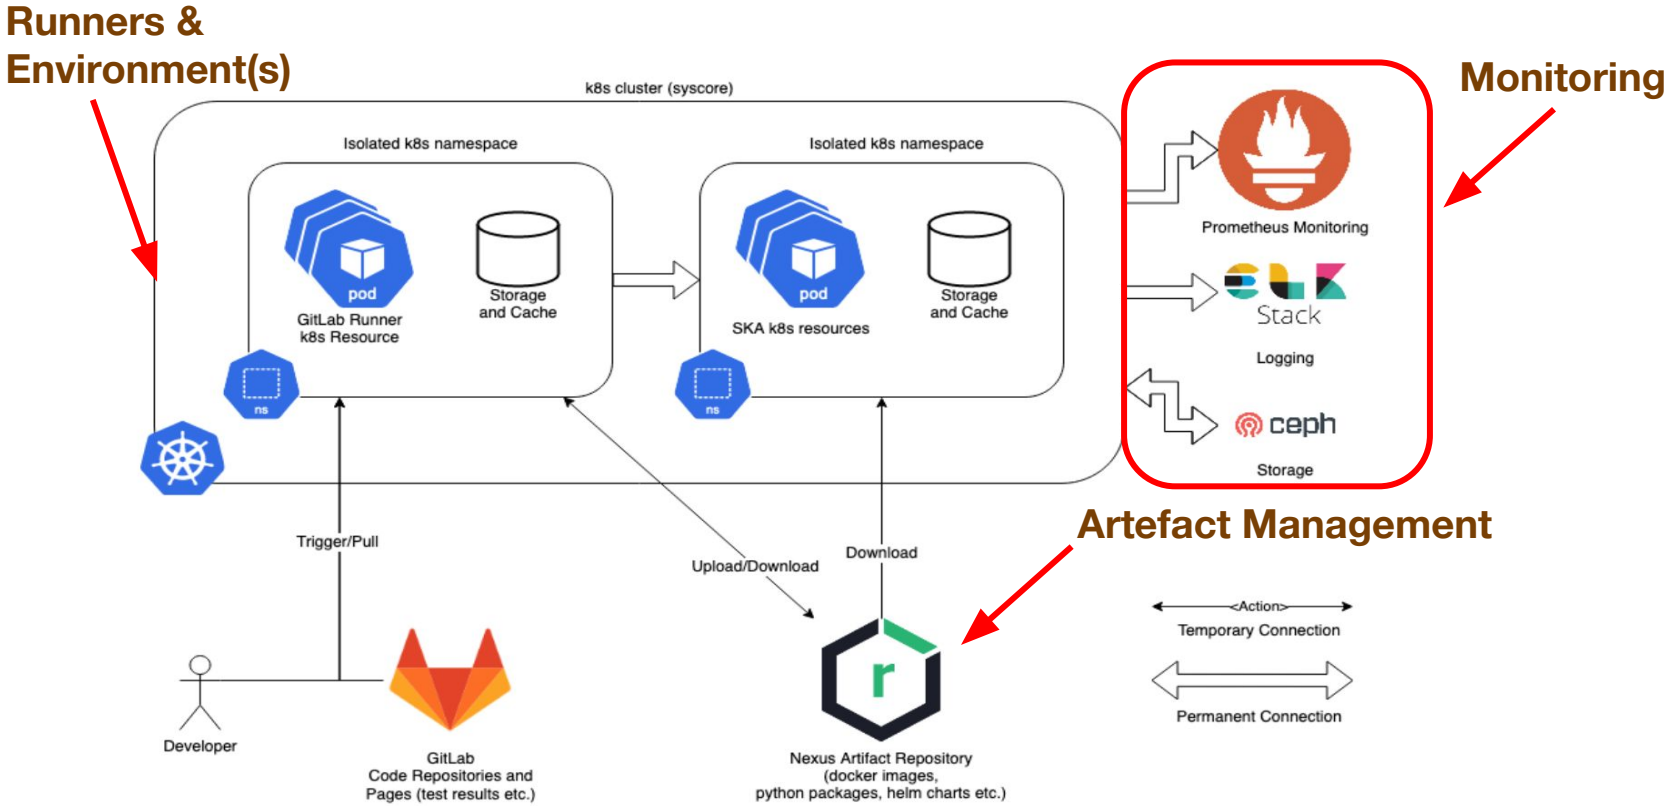

We cannot deliver everything at once!
→ → Array Assemblies (AAs)

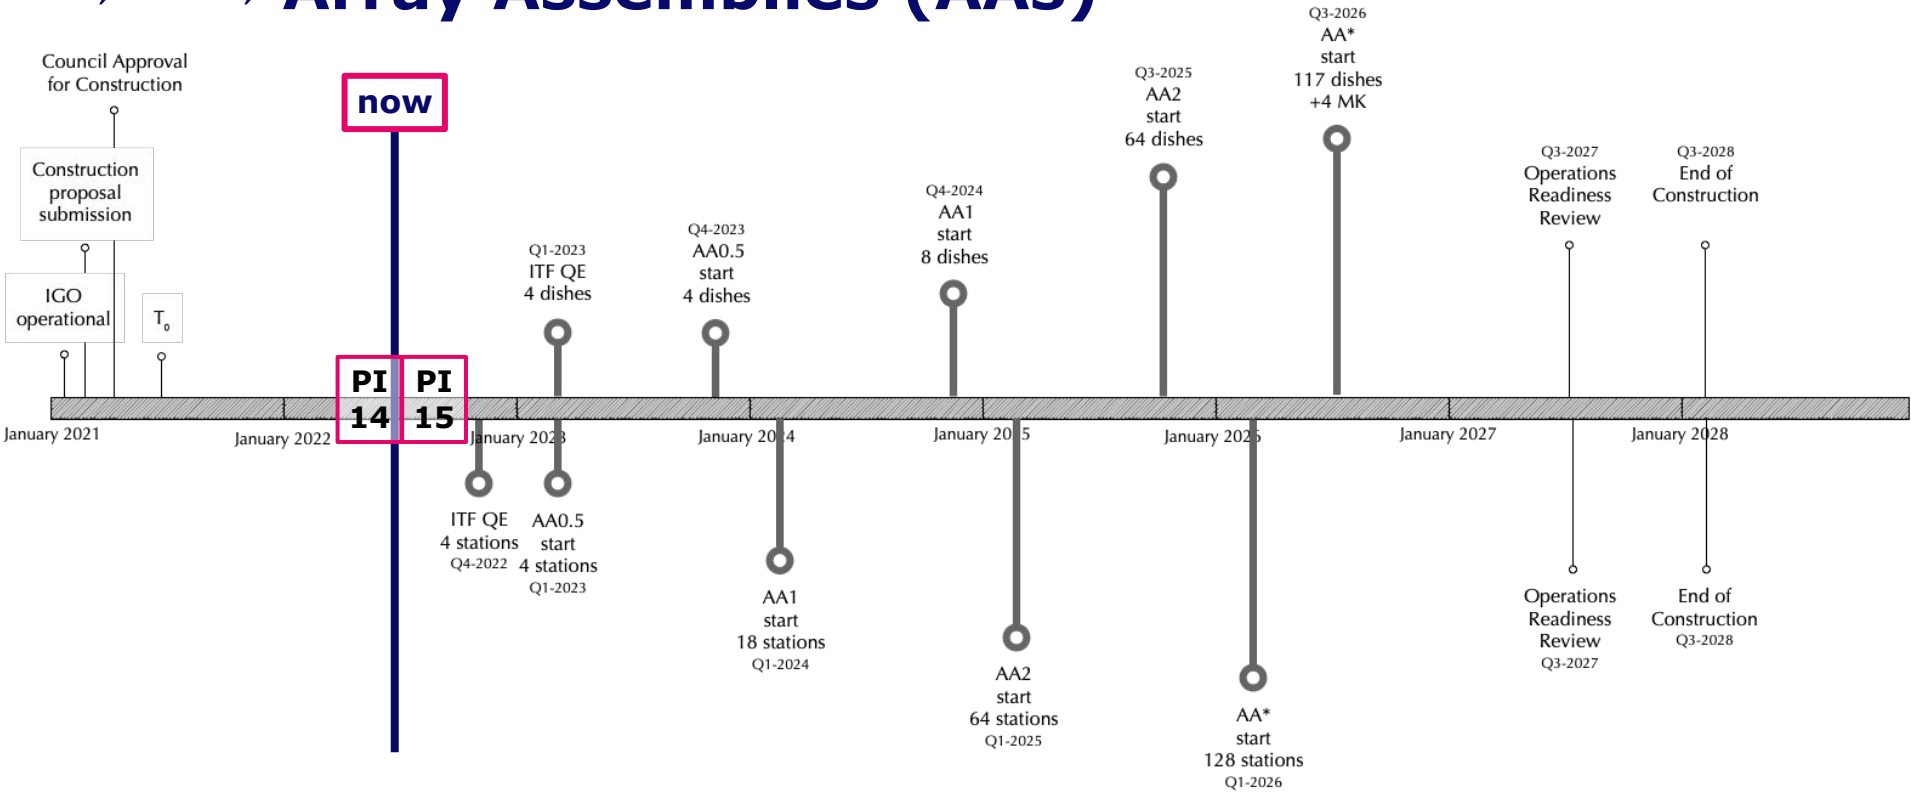

vertical scales are logarithmic wrt Array Assembly size

SKA Observatory Software

… and the challenges IRIS is helping us address!

SKA software challenges IRIS is helping us solve

1. How do we build a **consistent development & deployment platform** to a deliver robust, resilient, distributed software system?

2. How do we **scale SDP data processing for full-array operations** meet the science performance needs of the telescope users?

3. **How will** the network of **SKA regional centres** to **enable users to manage and analyse of telescope data** products?

Challenge 1: Consistent development & deployment platform

- **How do we provide a consistent platform API & UX for a range of users?**
  - Developers, AIV Engineers, & Operators
- How do we roll out a consistent, common platform with multiple infrastructure services and on environments?
  - Dev and test envs: IRIS, EngageSKA
  - Integration envs: PSI, ITF
  - Telescope envs: MID & LOW AAs

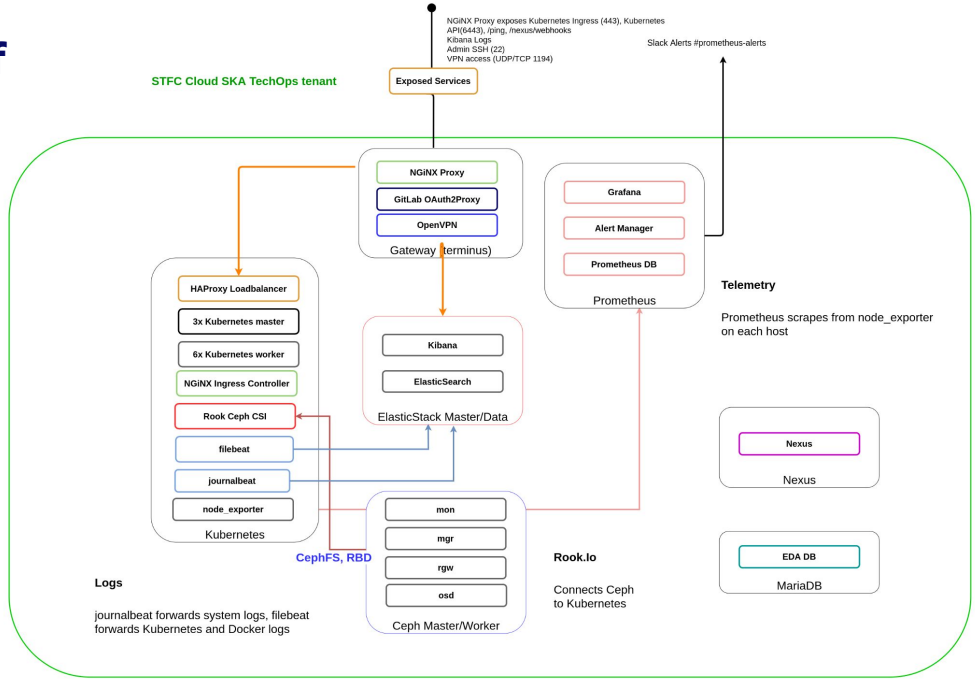

SKA STFC Cloud deployment footprint

Challenge 2: The SDP crunch is coming!

SDP complexity back-loaded:

- **Last stage** in signal path
- **Real-world telescope data** needed to inform development
- **Worst for scaling** (~50x within 17 months!)

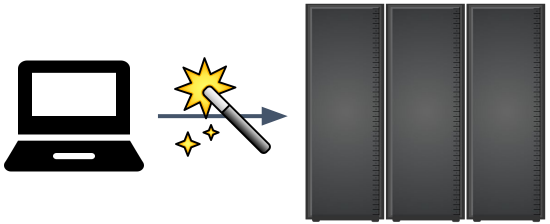

How to prepare?

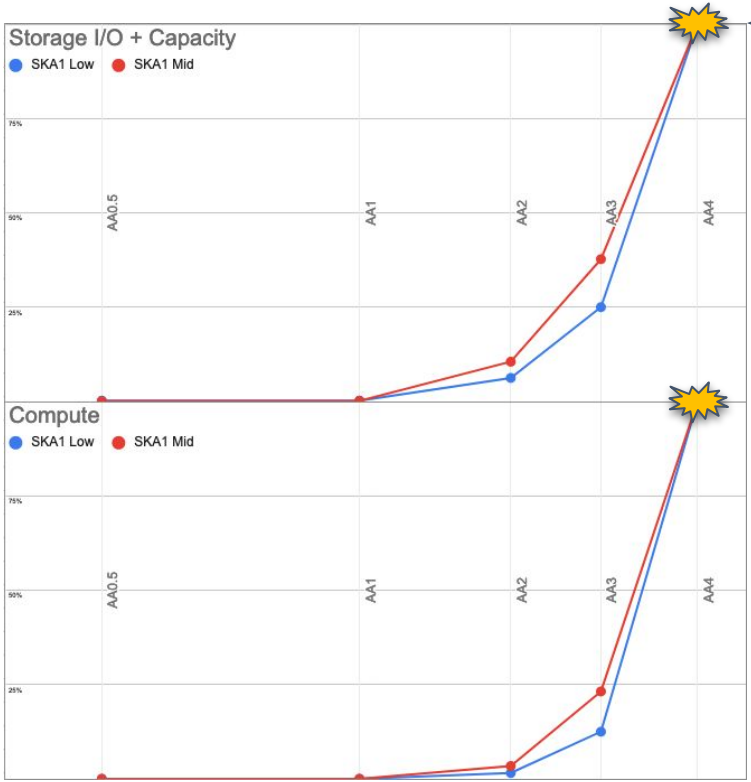

~ 40 PB storage
~ 4 TB/s I/O
(per site)

~ 10 Pflop/s
(effective,
per site)

Challenge 3: Prototyping the SRC network

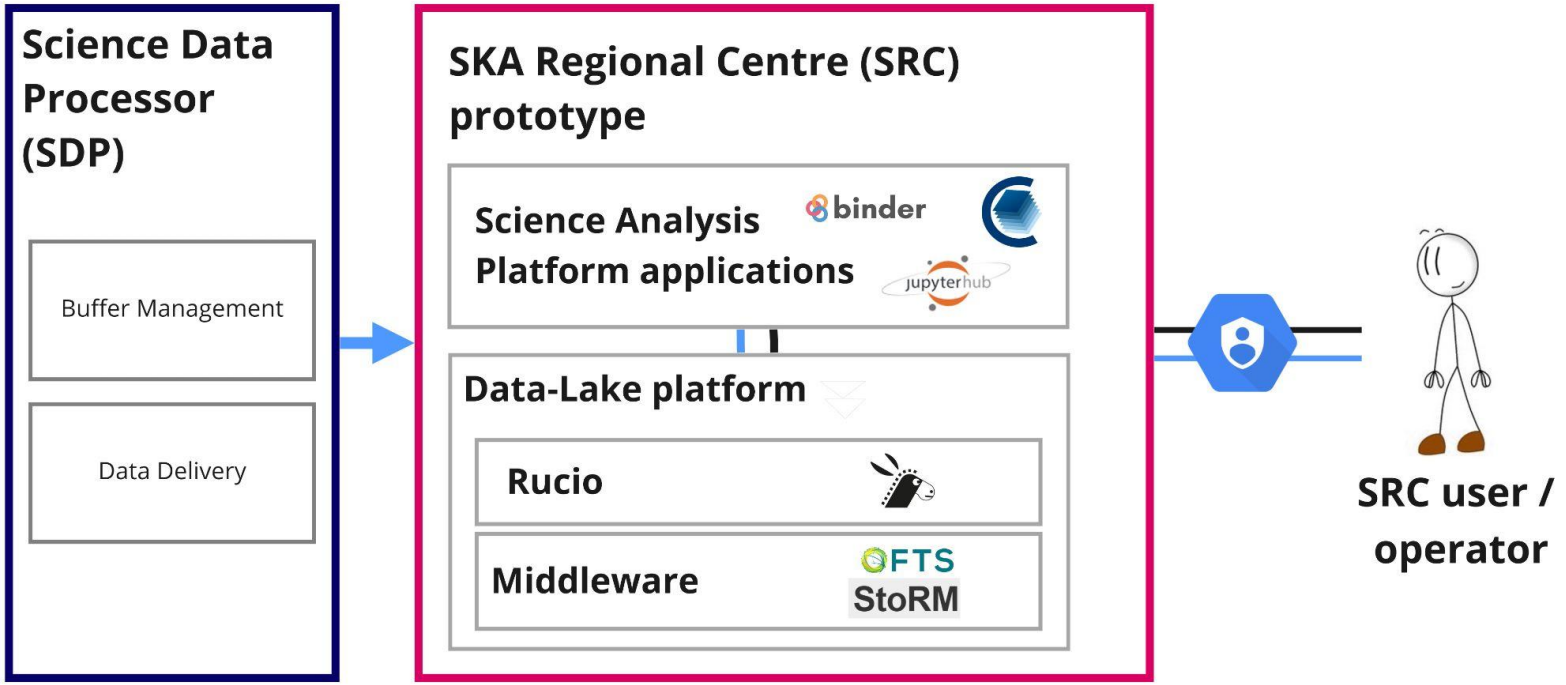

.... more details in the SKA Regional Centre section

SKA Observatory Software

Use of IRIS Resources

Deployment and use of IRIS resources for SKA Observatory Software

- SKA Observatory Software has **two IRIS projects**:
  1. SKA Operational System integration & testing
  2. SDH&P (Data Processing) workflow development
- These directly address challenges 1 & 2
- Platform for both projects deployed **using software defined infrastructure**
  - Use of Ansible for common, configurable deployment solution
  - ***Technical details for follow up presentation?***

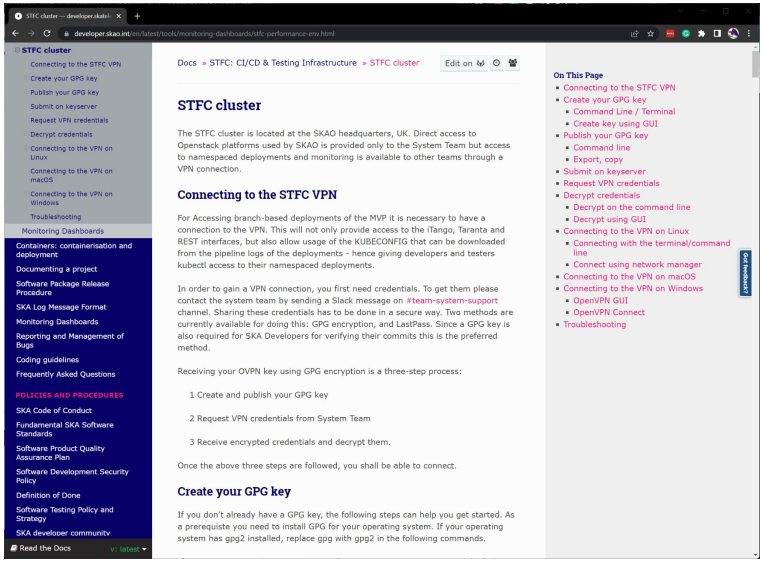

IRIS STFC cluster instructions in the SKA developer portal

1. SKA Operational System Integration & Testing

Current Use (1)

- **Gitlab CI/CD runners**
  - Support for GPUs enables testing for SKA simulators and SDP processing function library development
- **Automated linting, testing, packaging**, and deployment of SKA Software
  - ~300 repos, ~1000 jobs/day
  - 50/day integration env (20 concurrent)
- **Traceability** between JIRA XRay defined BDD tests and CI Job execution
  - Enables verification and regression analysis of Array Assemblies engineering requirements

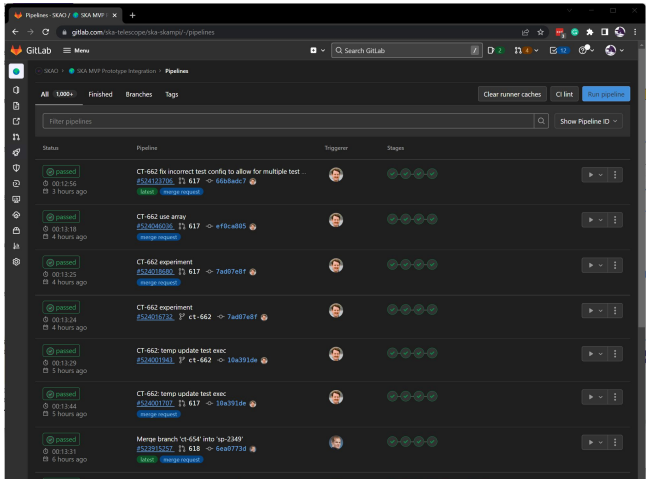

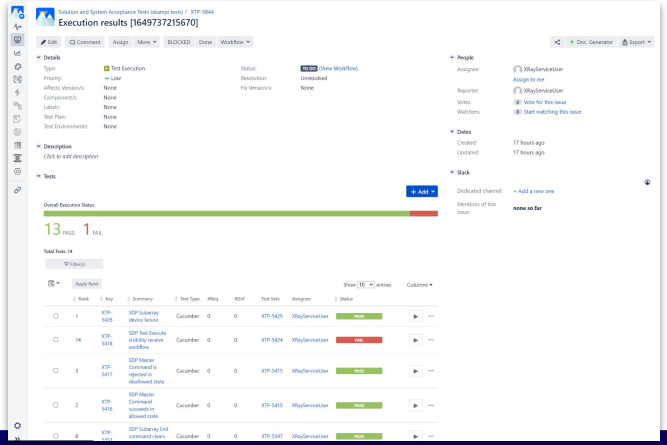

1. SKA Operational System integration & testing

Current Uses (2)

- **K8s cluster for staging and testing environments**
  - Deployments managed using Helm configured by CI/CD jobs
  - Multi-tenancy & job isolation achieved using k8s namespaces
- **Nexus central artefact repo** providing source of truth for SKA packages and containers
  - Provides automation to validate published artefacts
  - Artefacts traceable to original commit
  - Automated processing and checking of package/image dependencies

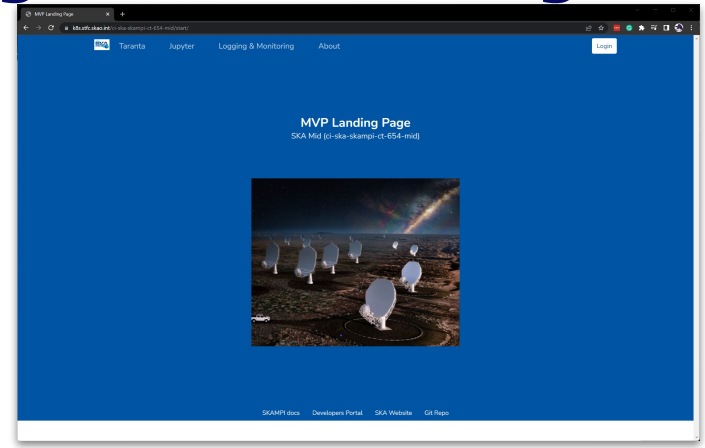

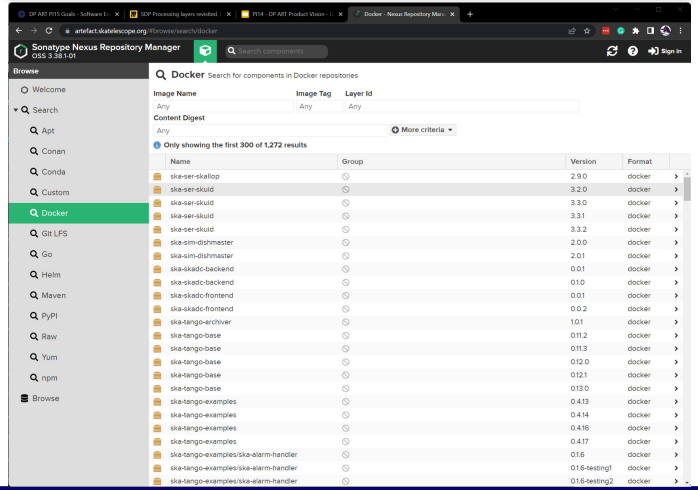

1. SKA Operational System integration & testing

Next Steps

- **Next 3-9 months: Ramp up in testing load for AA0.5 commissioning tests**
  - Expecting increased CI/CD load as we move from prototypes to production systems and add additional system capabilities
  - Creating clusters on demand for persistent secure testing environments for demos, exploratory testing, and operator (user) feedback (eg. k8s Cluster API)
  - On the part to support deployment and operation of platform for Array Assembly 0.5 (Q1/2 2023)
    → IRIS provides training ground for this

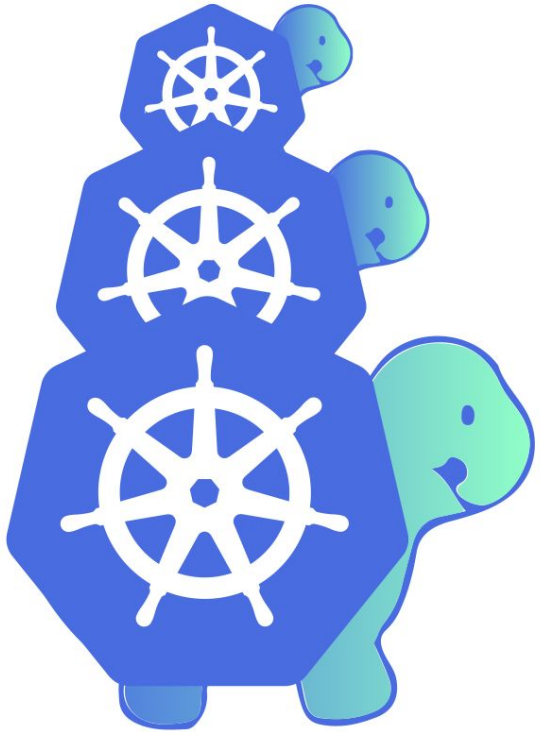

2. SDH&P workflow development

Current Uses (1)

- **Still experiential & ramping up** (first deployed in Dec 21)
  - Deployed as GPU enabled K8s cluster
- **Now completed initial set-up & shakedown** for use as a shared test platform for SDP workflow development
  - Collaboration between Services & Data Processing ART
  - Used existing prototype distributed Dask continuum imaging workflow, deployed using Helm
  - Successful, but still some usability and stability improvements noted for follow up work

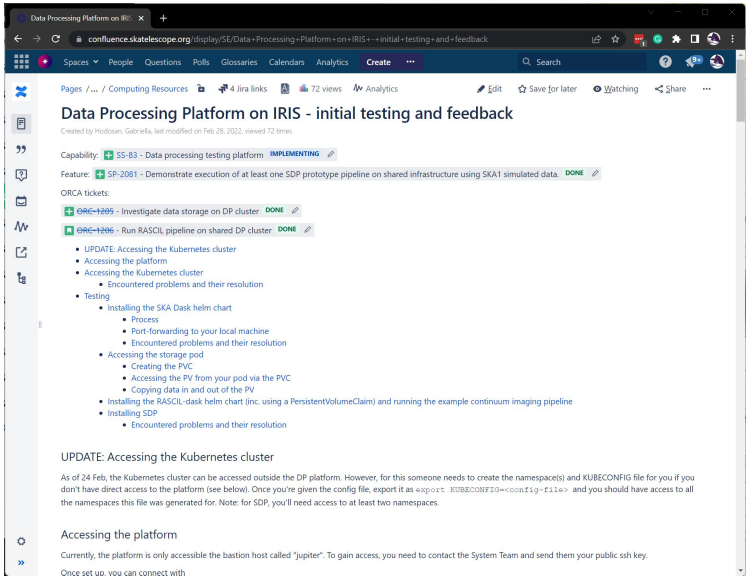

2. SDH&P workflow development

Current Uses (2)

- **Simulations for analysis of end-to-end telescope** imaging performance
  - IRIS has enabled SKAO system scientists to run SKA simulation software developed by DP ART software teams
- **Example simulation**: Exploration of beam performance with pseudo-random and Vogel station layouts
  - Use of and developed the OSKAR simulator - tested and packaged with SKA IRIS CI/CD infrastructure
  - Benefited greatly from 4xA100 GPU (g-a100.4) nodes

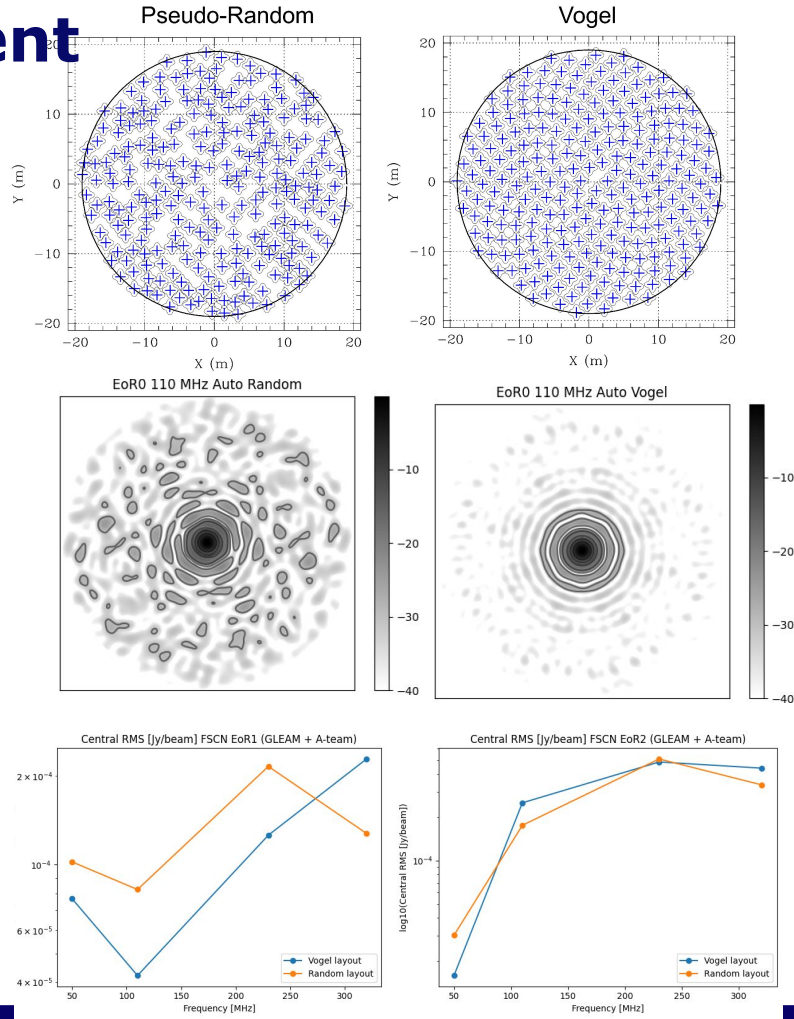

2. SDH&P workflow development

Current Uses (3)

- Develop of a **non-trivial I/O intensive SDP data reduction workflow**; Algorithm: “distributed FT”
  - Distributed image ↔ grid space transform
  - Aims to address processing huge images and data sets that don’t fit in system memory
  - By distributing both data and image subgrids without having to hold the full data set in any one node.
- Goal of initial phase of work: **Assess suitability an approach to using Dask and IRIS resources for SDP scaling** tests
  - Initial live demo to stakeholders last week (right) using 3 large IRIS nodes
  - Performance results and analysis available soon!

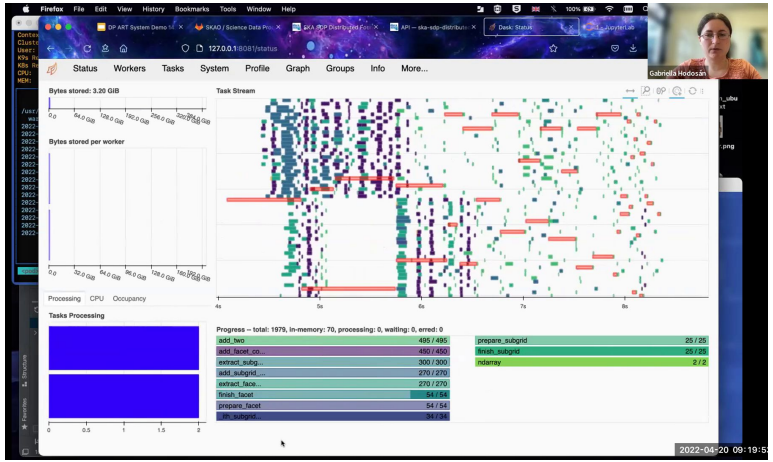

Distributed Dask scaling demo shown at a recent DP ART System demo (Credit: Team ORCA)

2. SDH&P workflow development

Next Steps

- **Next 3-9 months: Persistent SDP deployment**
  - Gain operational experience needed for running commissioning tests for early Array Assemblies (AAs)
  - Develop persistent API for integrating & testing new SDP workflows
- **Next 9-18 months: SDP workflows for AA2+**
  - SDP workflow data challenges for AA2 scale using simulated and SKA precursor telescope data
  - Scaling
    from: 2 CPU nodes, ~4k^2 pixels, ~10GB data
    to: 20 GPU nodes, and 65k^2 pixels, ~10TB data
  - Set-up and Tear-down of resources on demand (depends on underlying GPU resource availability and reliability, use of Cluster-API for autoscaling?)
  - Support for burst interactive operation (manual benchmarking) and regular testing using CI/CD

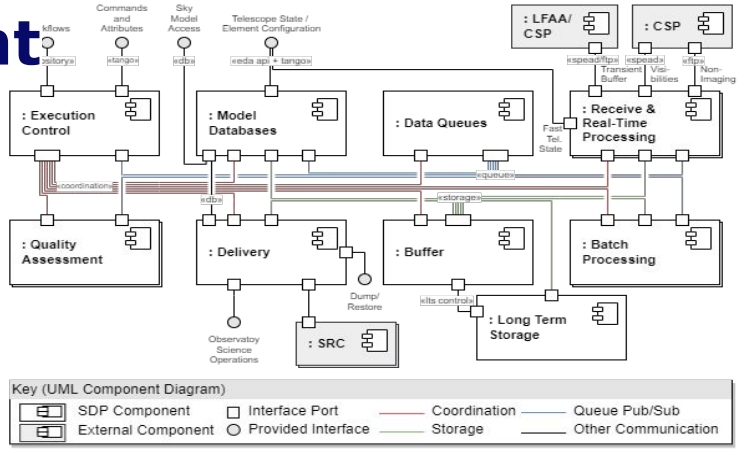

High-level SDP components

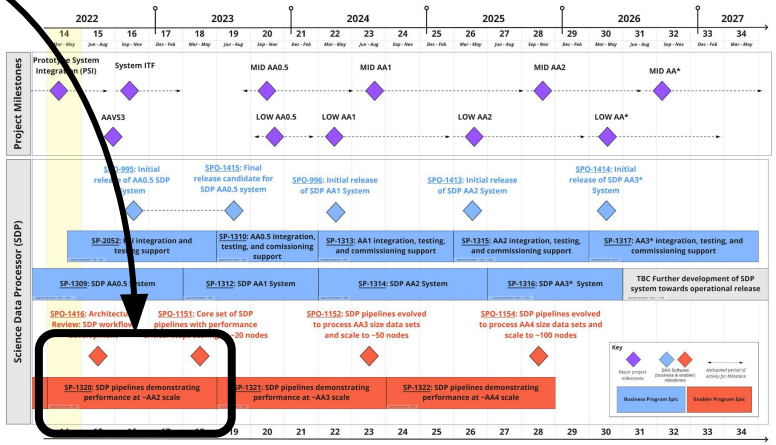

SDP Long term strategic roadmap

SKA Regional Centre activities

Overview

- SKA Regional Centres (SRCs)
  - SRC Capabilities
- Activity streams
- What’s next
- SRCNet Prototyping phase

SKA Regional Centres: SKAO data processing stages

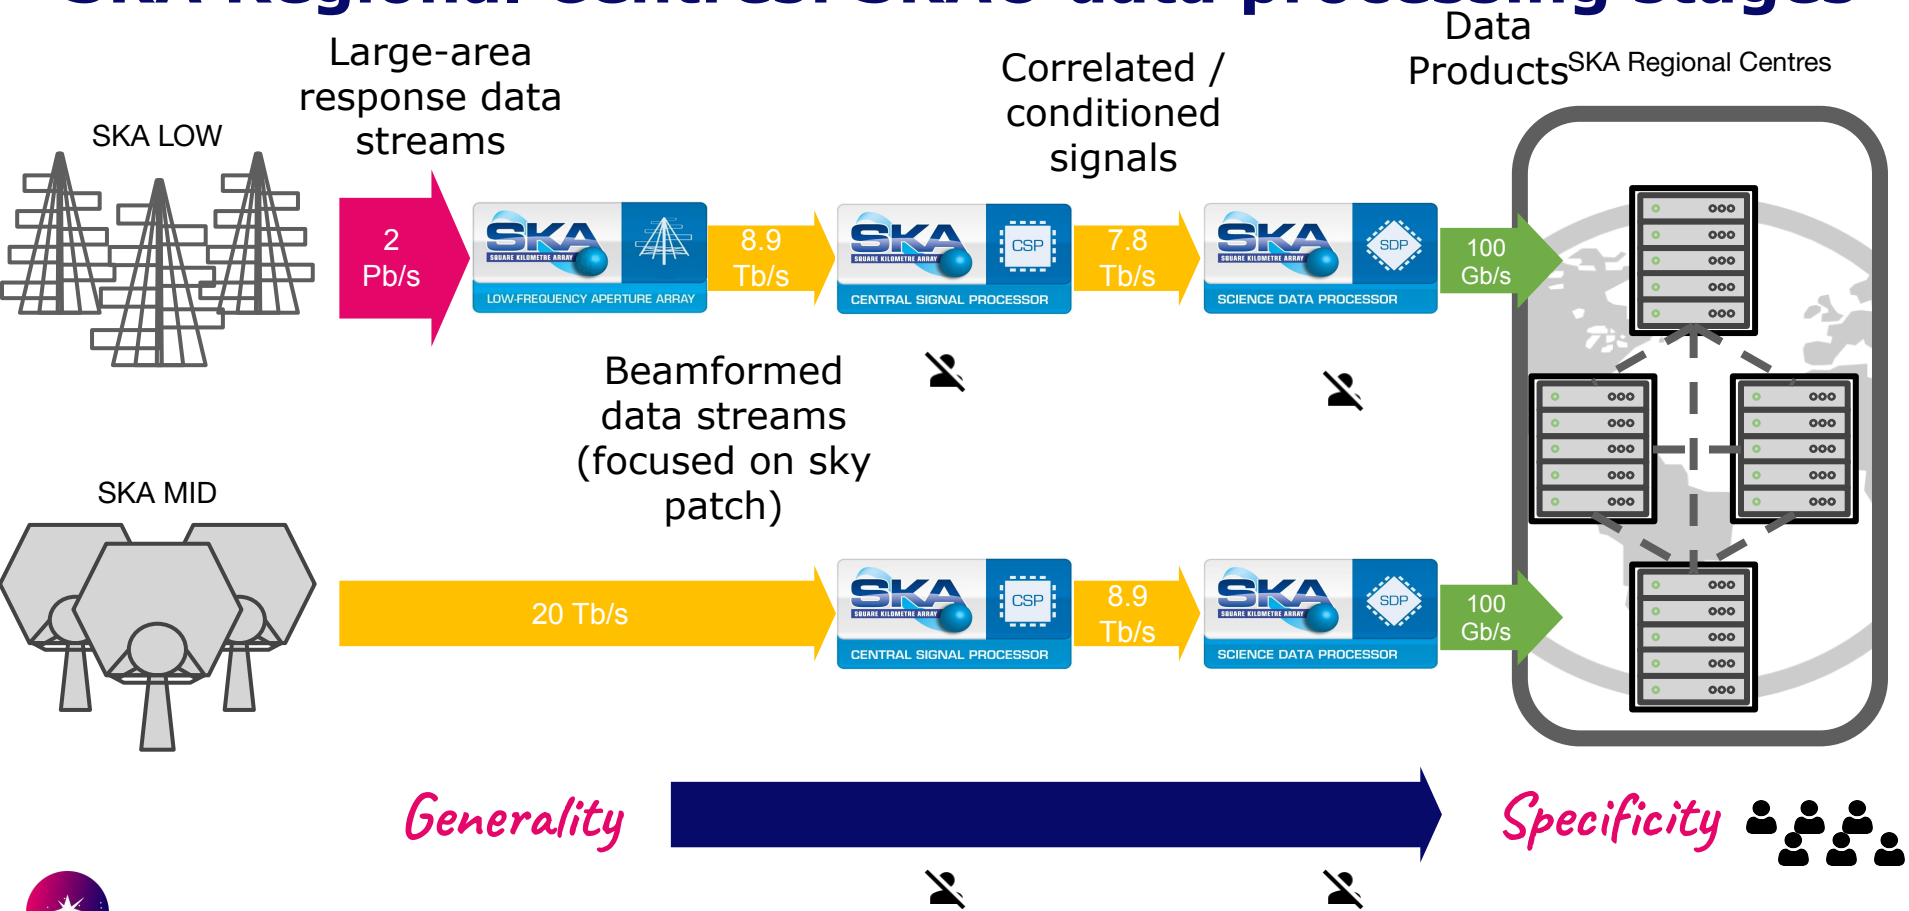

SKA Regional Centre Capabilities

Science Enabling Applications
Analysis Tools, Notebooks, Workflows execution Machine Learning, etc

Distributed Data Processing
Computing capabilities provided by the SRCNet to allow data processing

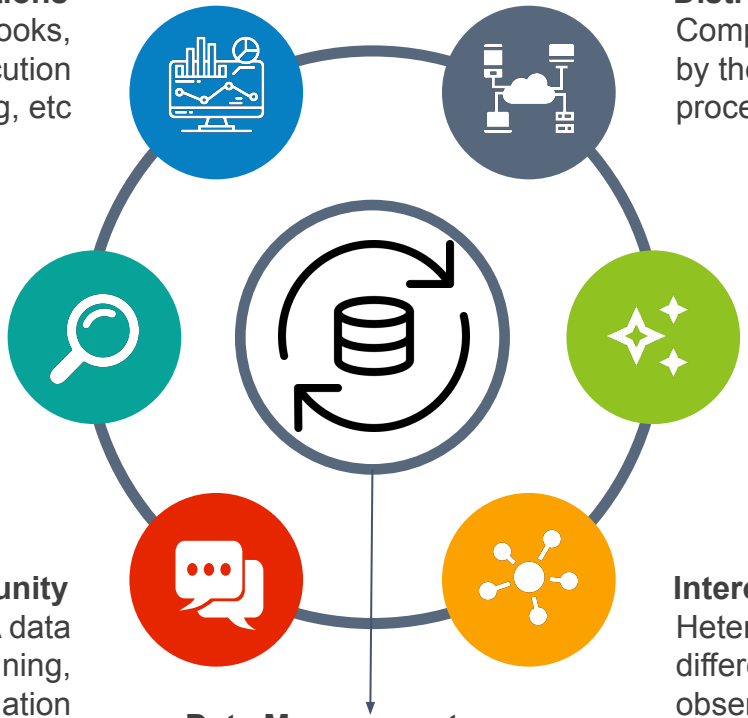

Data Discovery
Discovery of SKA data from the SRCNet, local or remote, transparently to the user

Visualization
Advanced visualizers for SKA data and data from other observatories

Support to Science Community
Support community on SKA data use, SRC services use, Training, Project Impact Dissemination

Interoperability
Heterogeneous SKA data from different SRCs and other observatories

Data Management
Dissemination of Data to SRCs and Distributed Data Storage

What have we been doing*?

*on IRIS resources under the SRC tenant

Activity Streams

- ESCAPE
  - Rucio
    - Rucio Server, storage
    - ESCAPE monitoring
  - ESFRI Science Analysis Platform
    - JupyterHub/BinderHub
- SDC
  - Scoring service
  - (Resources for teams)
- User engagement/enabling activities

Deployed services

Kubernetes clusters, Grafana dashboards

Ceph cluster (ceph on ceph), Rook, StoRM-webdav

JupyterHub, Binderhub,

Keycloak service, Database, Submission handler, dashboard

CARTA

What have we been doing*?

*on IRIS resources under the SRC tenant

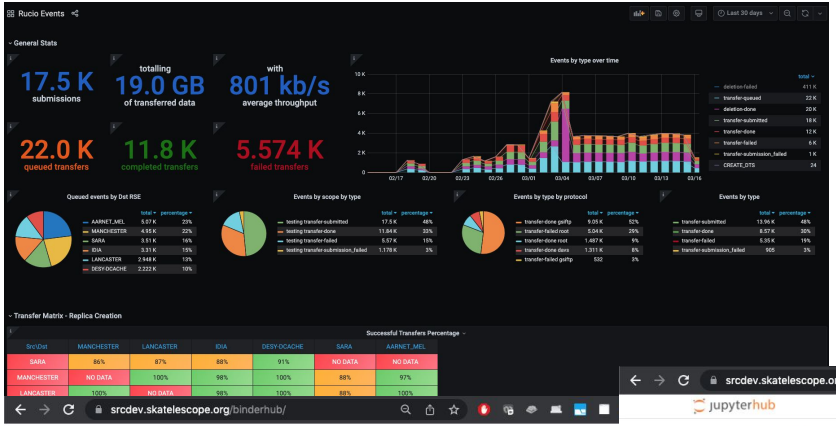

Kube
dashl

Ceph

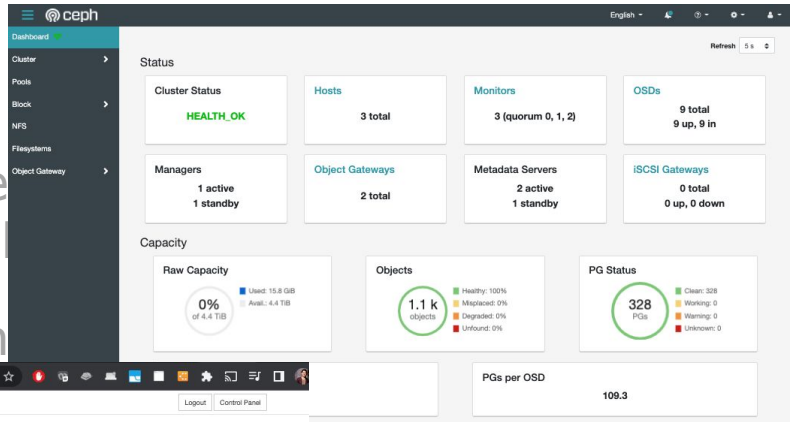

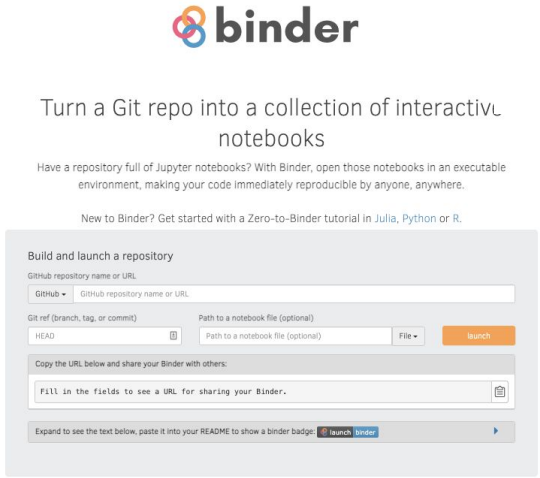

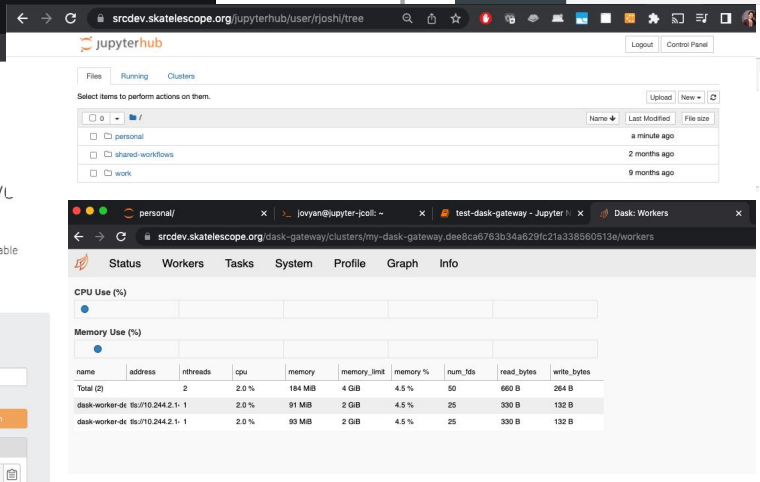

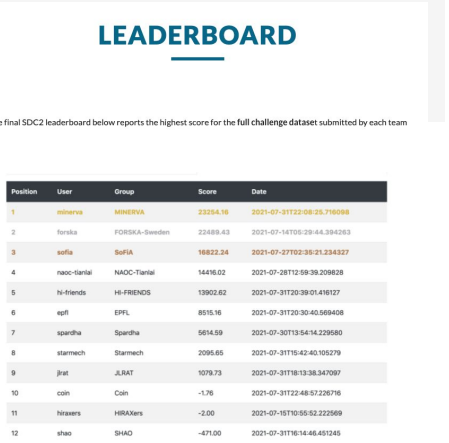

What’s next

- Storage
  - Data lake storage
    - More endpoints
    - Learning about token-based storage
    - Data management system design
  - Improvements to the way storage is provisioned for k8s
- Technical debt
  - Cluster consolidation
  - External footprint reduction
- Accounting/monitoring of resources

SRCNet Prototyping phase

- Shifting from a working group structure to an agile way of working in order to build the network of Regional Centres.
- This will aim to bring together regional SRC efforts and improve alignment among protoSRCs but also between the SKAO and the SRCs.
- Teams identified and being refined
- Prototyping phase Kickoff meeting held on April 19th
- Upcoming events include Backlog development & refinement, SoS, Planning events
- UK-SRC efforts being lead by Jeremy Yates