SKA MASUM 2022 meeting

Software

11 May 2022

Adam Avison, Marco Bartolini, Alan Bridger, Piers Harding, Peter Lewis, Vivek Mohile, Ben Mort, Aris Noutsos, Sonja Vrcic, Peter Wortmann

Overview

Software overview & Roadmaps

- SKA Software Solution overview
  - DP ART Roadmaps
  - OMC ART Rodmaps

Software & Computing Architecture

- Overview of the overall software architecture
- Software architecture at AA0.5

Software & computing impacts almost all aspects of the SKA observatory! ...

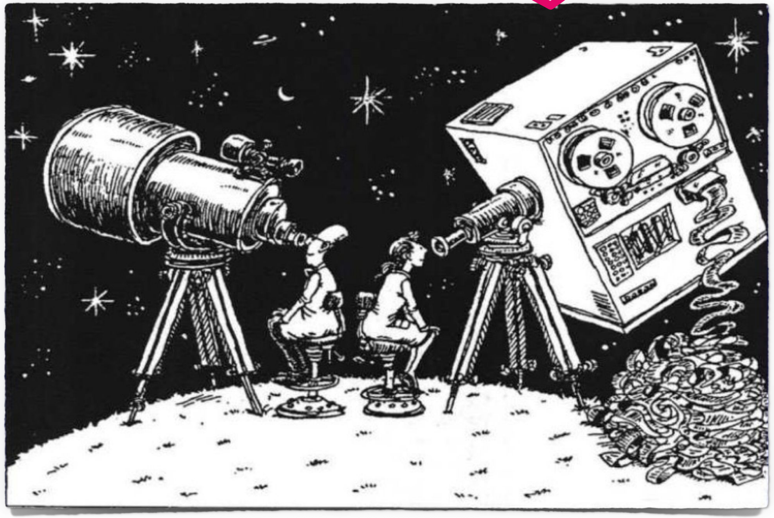

SKA Software Solution Overview

SKA Software - Key components

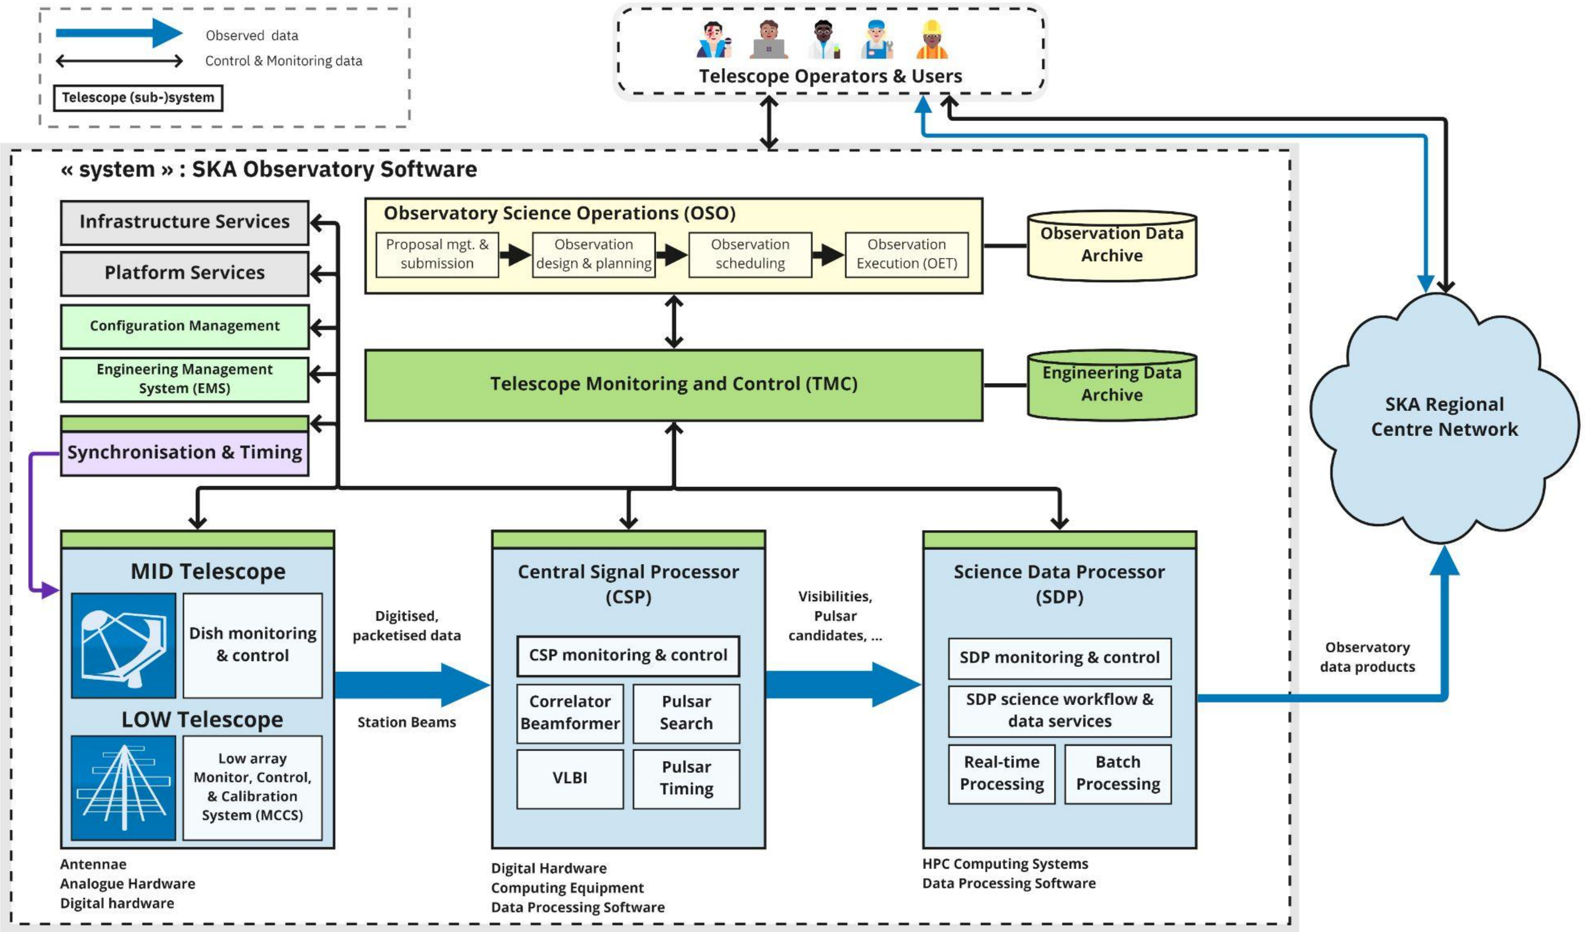

Organisation of SKA Software & Computing

- Coordinated effort of 21→ 23 globally distributed teams!
- Guided by the Scaled Agile Framework (SAFe™)
- Three Agile Release Trains (ARTs) + SRC ART now ramping up for start in June
- Alignment between ARTs managed by the Solution Team
- All ARTs follow a 3-month planning and delivery cadence → Planning Increment (PI)
- Work shared with stakeholders through regular series of system demos

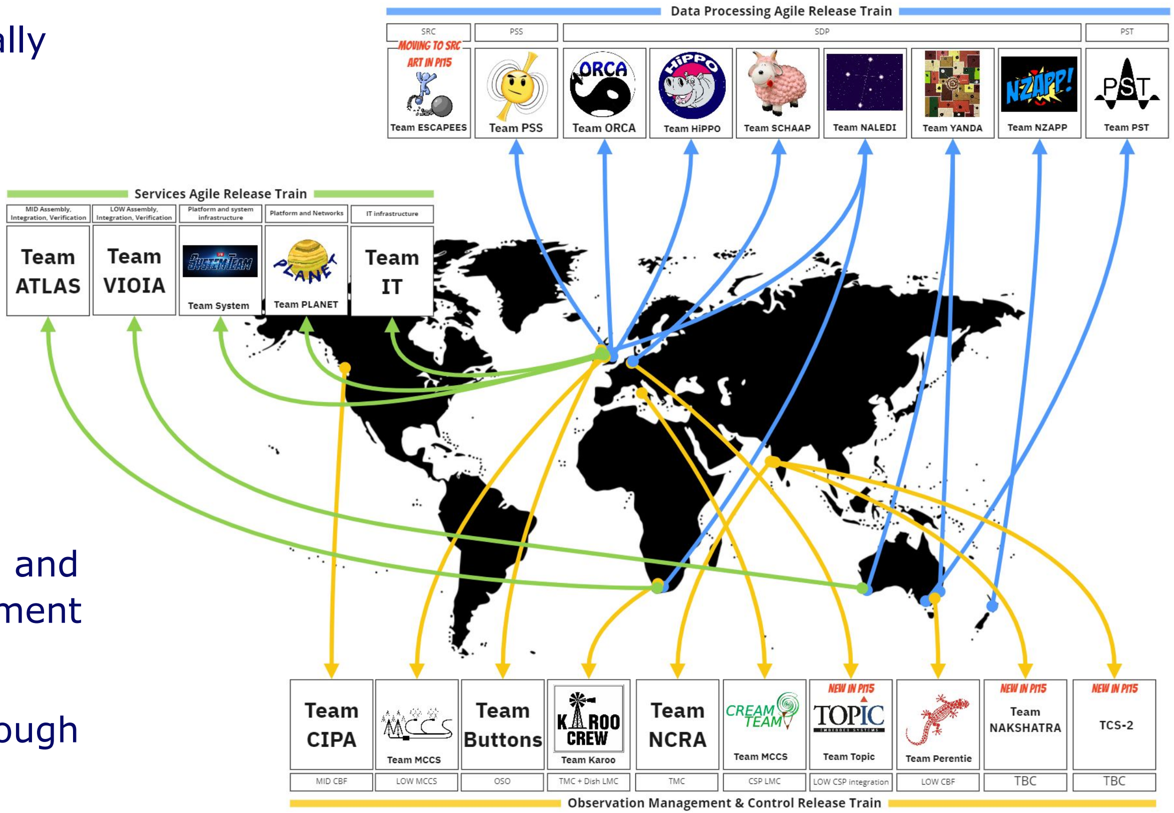

We cannot deliver everything at once! → Array Assemblies (AAs)

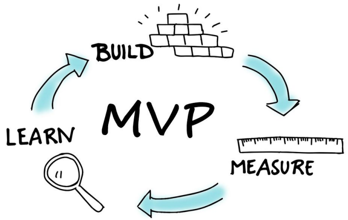

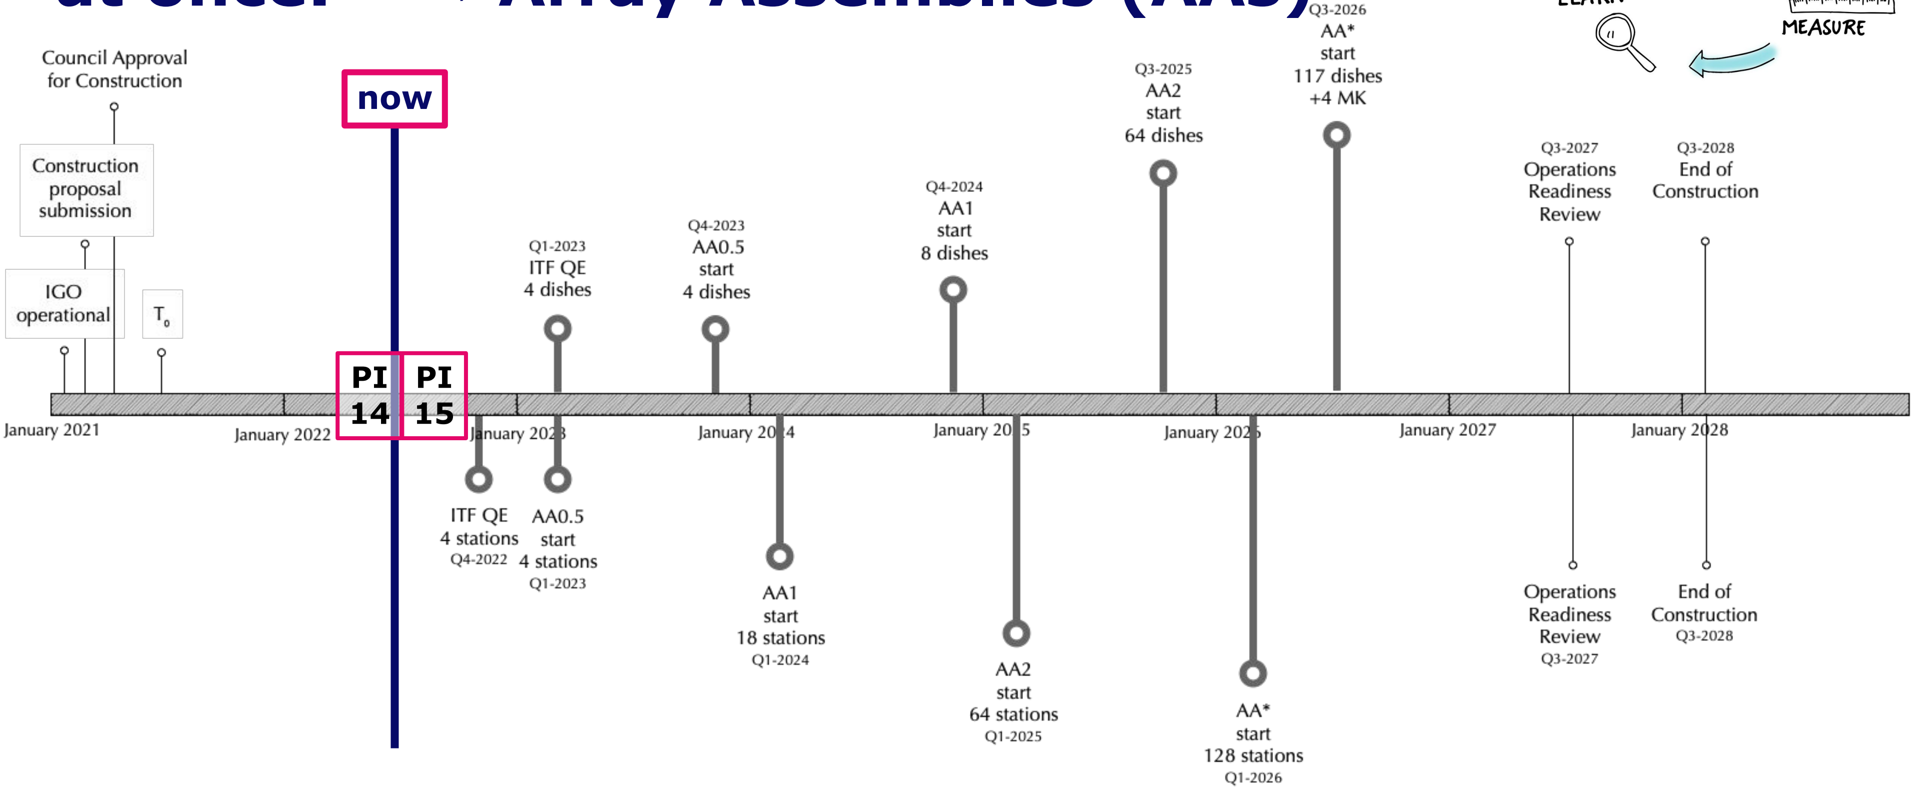

vertical scales are logarithmic wrt Array Assembly size

Strategic, Long-Term Solution Roadmap

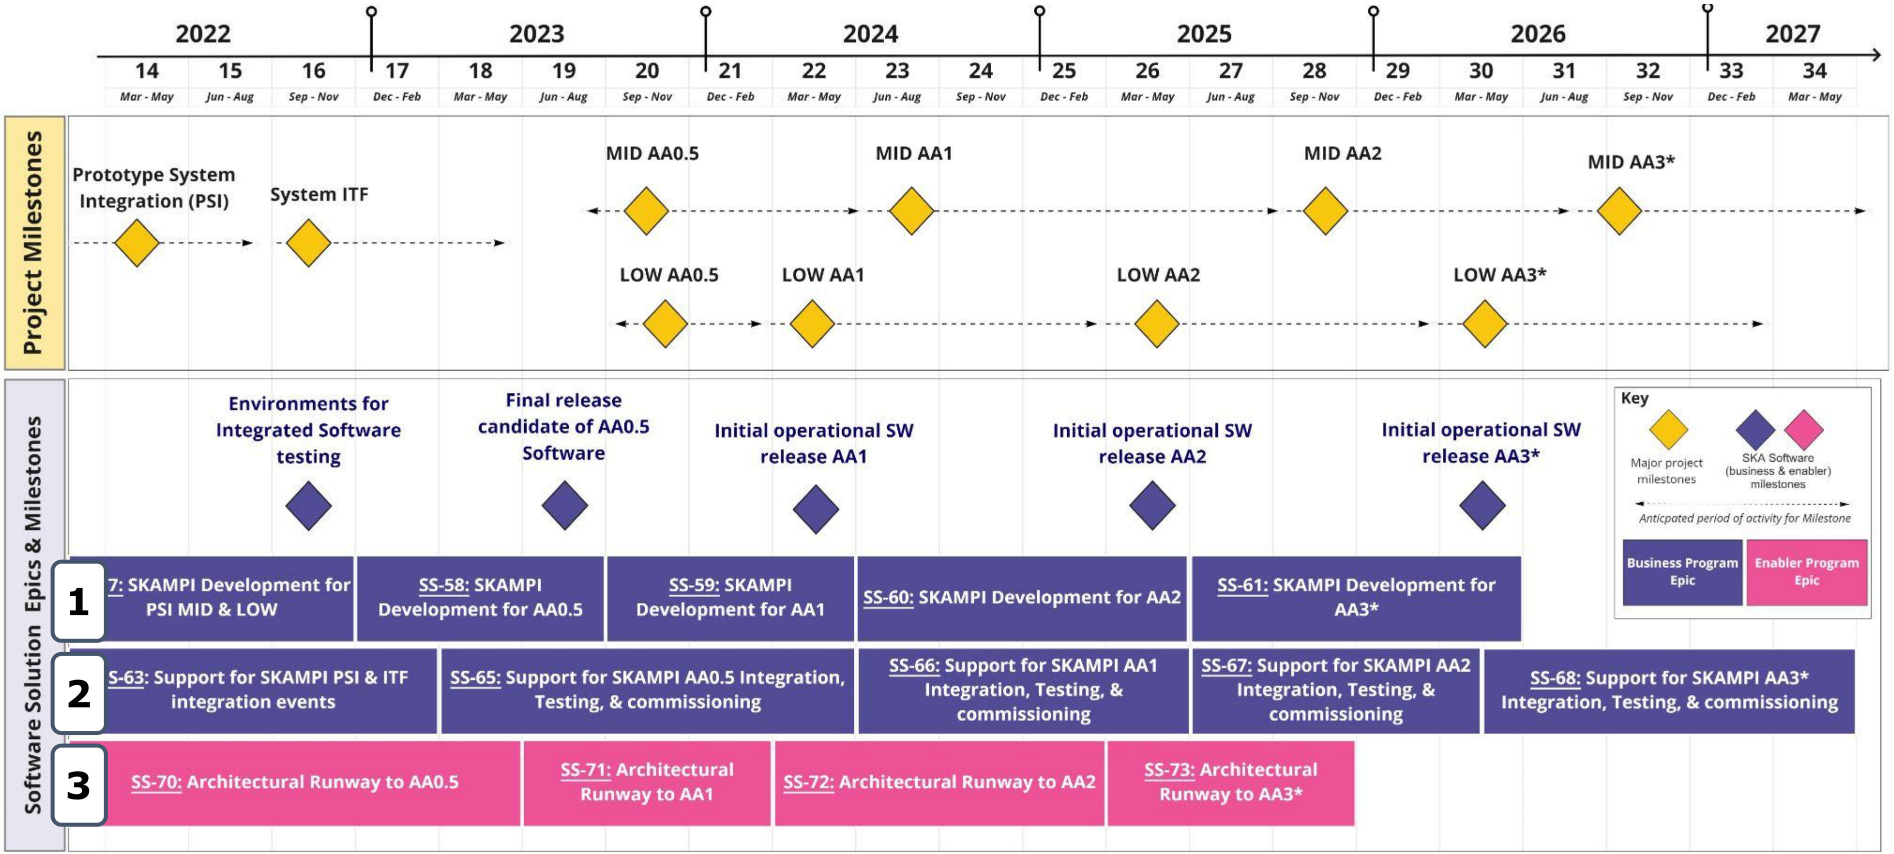

Use Cases and Case Case flows

https://confluence.skatelescope.org/display/SWSI/Solution+Use+Cases

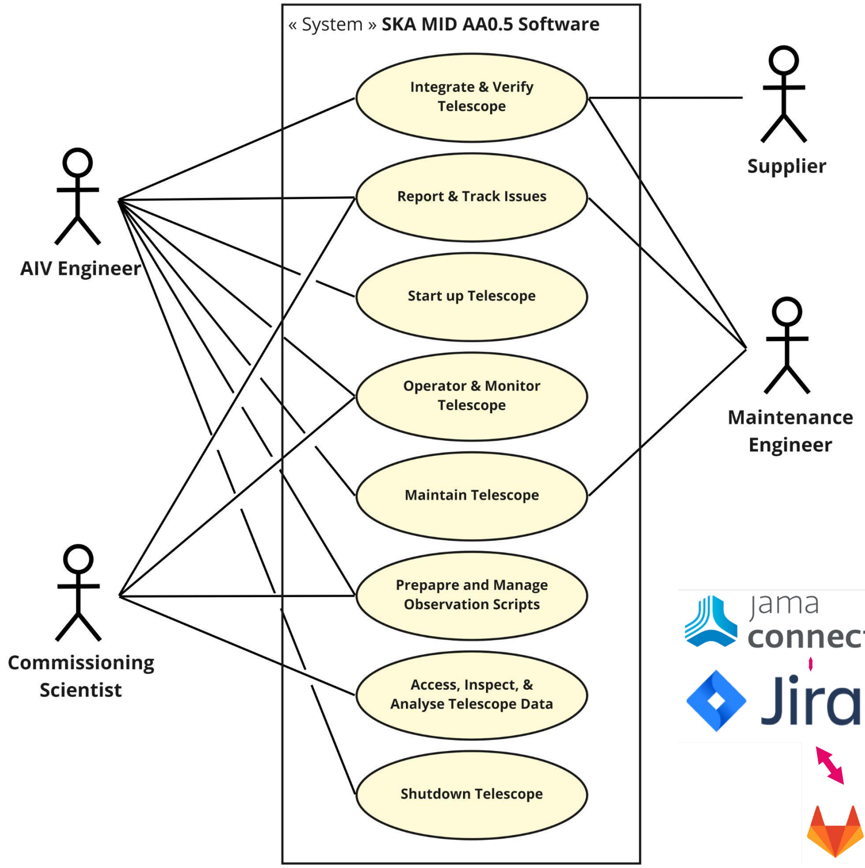

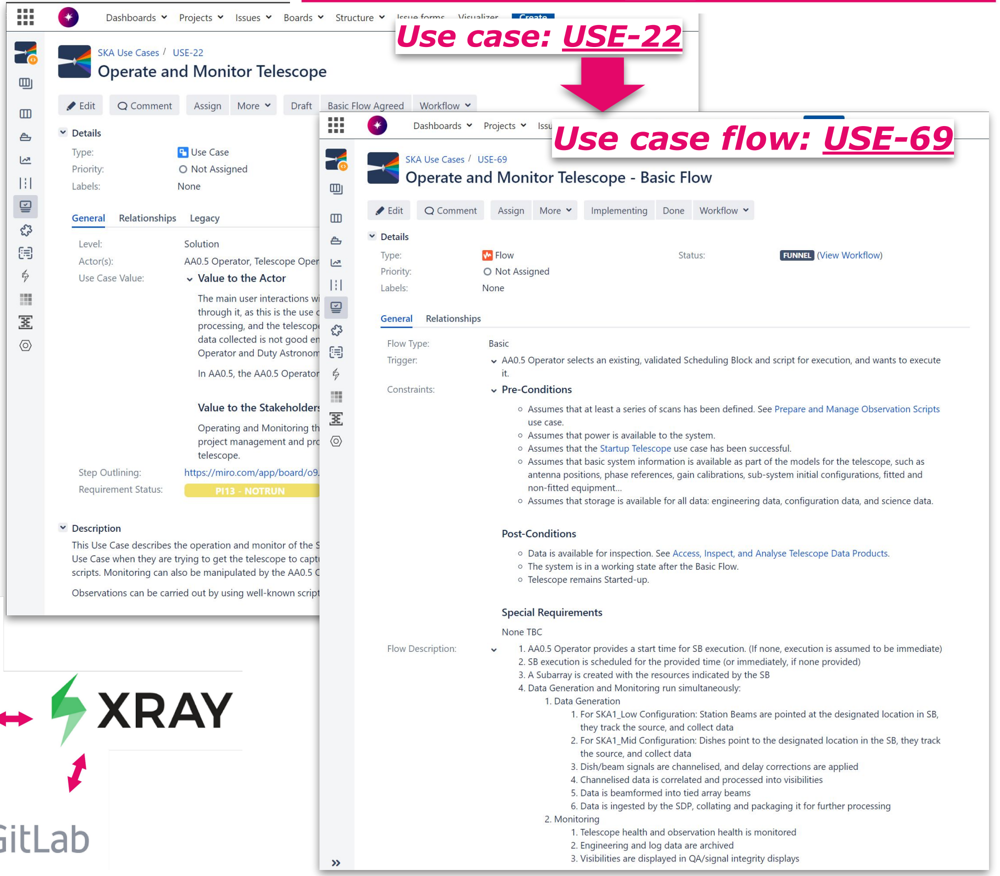

From Use Cases to "Story Mapping" …

- Complements use cases, mapping:

User Activities (goals) → **Narrative Flow** → **Features** organised by **Release Slice**

- Adds context for reasoning about feature and release prioritisation

- Being used from PI15

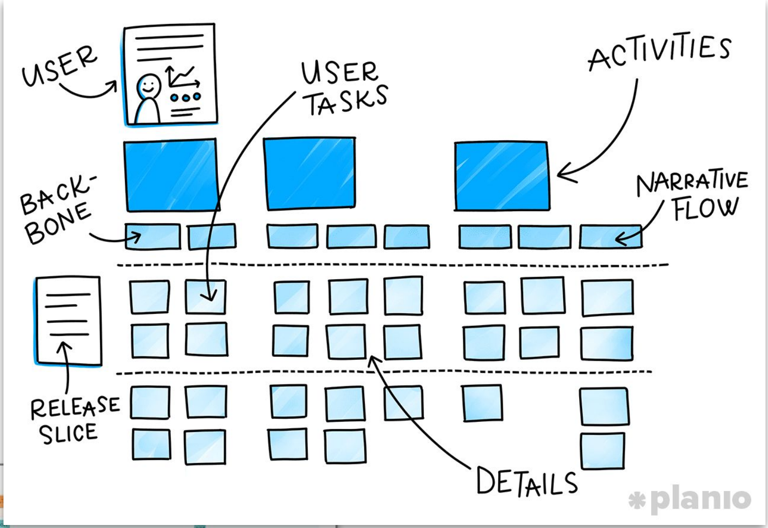

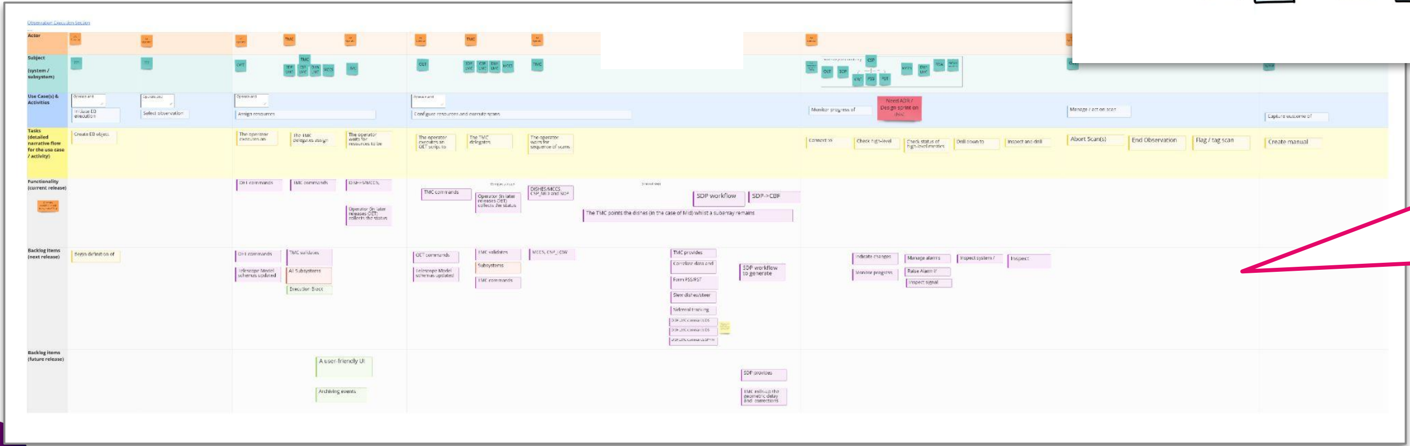

Story mapping board **for "Operator & Monitor"** use case.

Outcome of the PI14 [Roadmapping workshop]()

Current Solution Focus

Priorities for next 6-12 Months (2-4 PIs)

1. Robust, resilient, observable distributed control and monitoring system → user feedback from exploratory testing
2. Incremental delivery of integrated AA0.5 capabilities, needed for engineering and science commissioning, ahead of full AA0.5 deployment
   → Collaboration with AIV for system acceptance testing

Short term goals (PI15→16)

1. Exercise and test named release(s) of integrated software deploying AA0.5 MVP in established environments:
   - PSI-MID & PSI-LOW
   - Per-telescope, user accessible, sandbox/staging environment
2. Computing platform, including network and storage, for AA0.5 is well defined
3. System-level interfaces for tired-array beamformer calibration

SKAO

Roll-Out Plan for SKA1 MID

SKA-TEL-AIV-2410001	Revision 10
Classification:	Unrestricted
Document type:	MPL
Date:	2021-08-04
Status:	Released

Role	Name	Designation	Affiliation	Signature	Date
Author	R.T. Lord, et al.	AIV Senior Systems Engineer	SARAO		2021-08-06
Author	D. Gammon	AIV Systems Engineer	SARAO		2021-08-06
Author	P. Hekman	MID AIV Lead Engineer	SARAO		2021-08-09
Owner	P. Hekman	MID AIV Lead Engineer	SKAO		2021-08-09
Approver	G. Swart	MID Telescope Engineer	SKAO		2021-08-09
Released by	J. McMullin	Programme Director	SKAO		2021-08-08

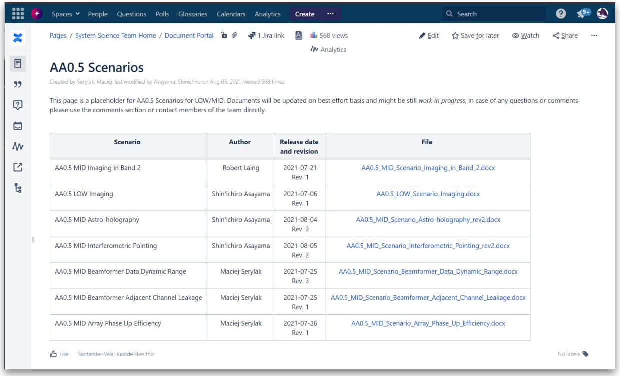

Pages / System Science Team Home / Document Portal

AA0.5 Scenarios

This page is a placeholder for AA0.5 Scenarios for LOW/MID. Documents will be updated on best effort basis and might be still *work in progress*, in case of any questions or comments please use the comments section or contact members of the team directly.

Scenario	Author	Release date and revision	File
AA0.5 MID Imaging in Band 2	Robert Laing	2021-07-21 Rev. 1	AA0.5_MID_Scenario_Imaging_in_Band_2.docx
AA0.5 LOW Imaging	Shin'ichiro Asayama	2021-07-06 Rev. 1	AA0.5_LOW_Scenario_Imaging.docx
AA0.5 MID Astro-holography	Shin'ichiro Asayama	2021-08-04 Rev. 2	AA0.5_MID_Scenario_Astro-holography_rev2.docx
AA0.5 MID Interferometric Pointing	Shin'ichiro Asayama	2021-08-05 Rev. 2	AA0.5_MID_Scenario_Interferometric_Pointing_rev2.docx
AA0.5 MID Beamformer Data Dynamic Range	Maciej Serylak	2021-07-25 Rev. 3	AA0.5_MID_Scenario_Beamformer_Data_Dynamic_Range.docx
AA0.5 MID Beamformer Adjacent Channel Leakage	Maciej Serylak	2021-07-25 Rev. 1	AA0.5_MID_Scenario_Beamformer_Adjacent_Channel_Leakage.docx
AA0.5 MID Array Phase Up Efficiency	Maciej Serylak	2021-07-26 Rev. 1	AA0.5_MID_Scenario_Array_Phase_Up_Efficiency.docx

Data Processing (DP) ART Roadmaps

Team PSS	Team PST	Team ORCA	Team HiPPO	Team SCHAAP	Team NALEDI	Team YANDA	Team NZAPP
PSS	PST	SDP					

SDP long-term roadmap

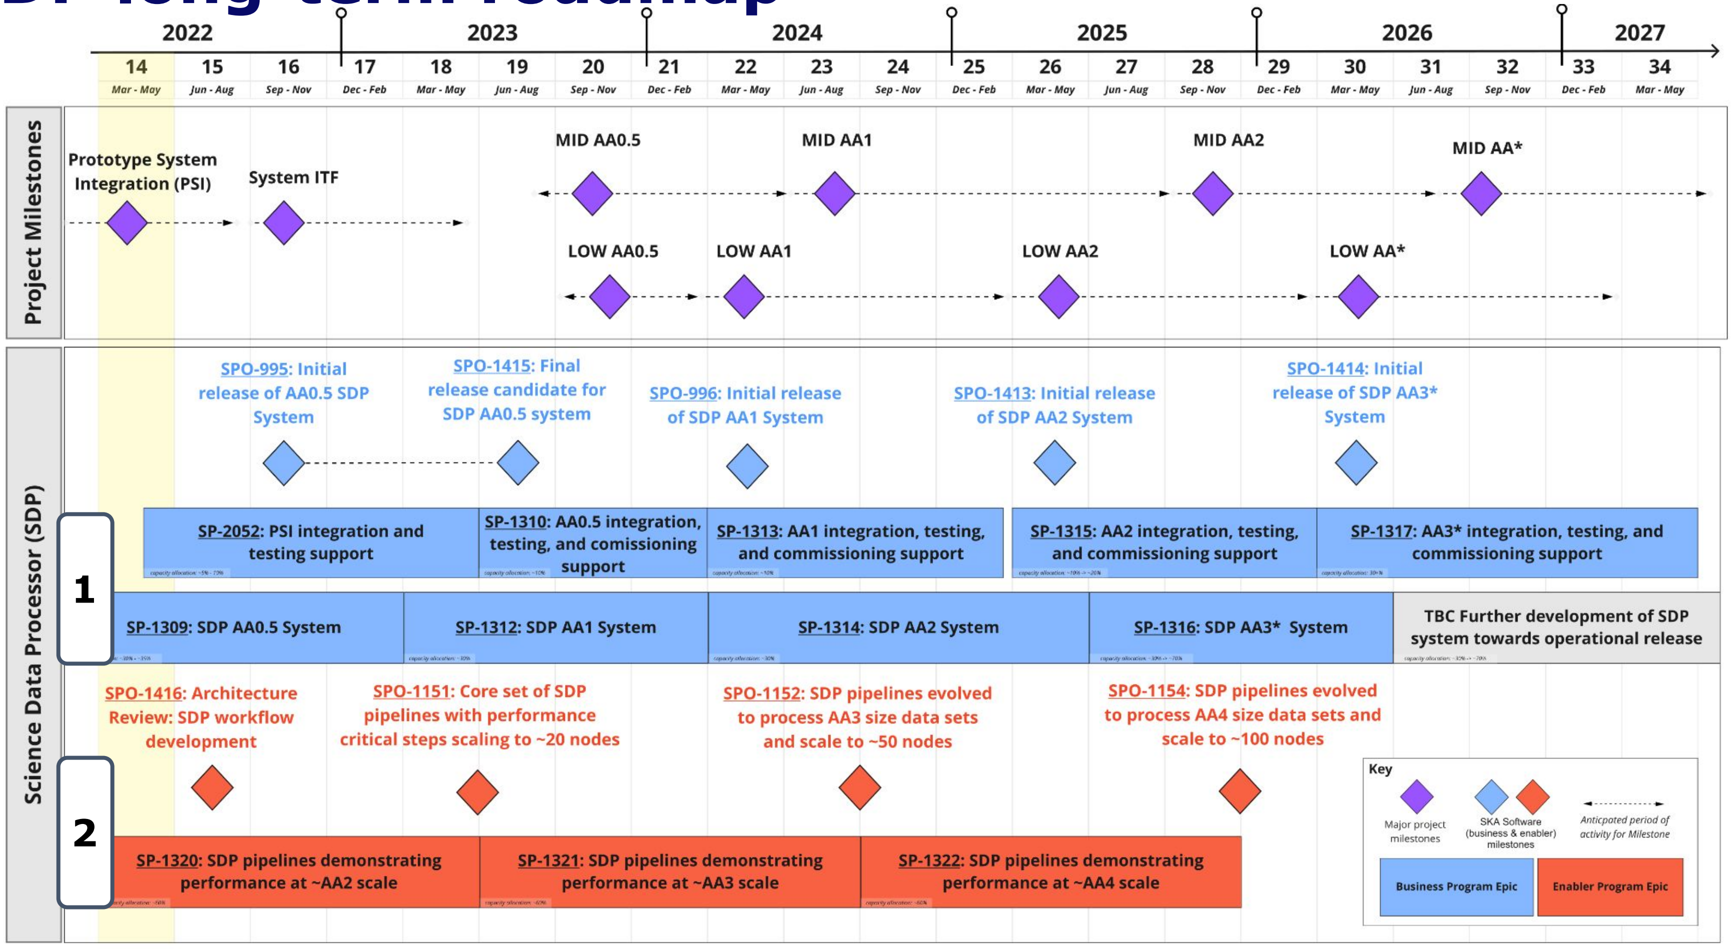

SDP Short Term Roadmap: Towards AA0.5

1. **SDP system for capturing visibility data at AA0.5**
   - Collaboration with OMC on operator interfaces
   - Aiming for early operator feedback & exploratory testing starting from PI15
2. **Data management** for test data and data products
3. Pseudo-real-time **tied-array beamformer calibration**
4. **MID pointing calibration**
5. **Support for PSI tests**

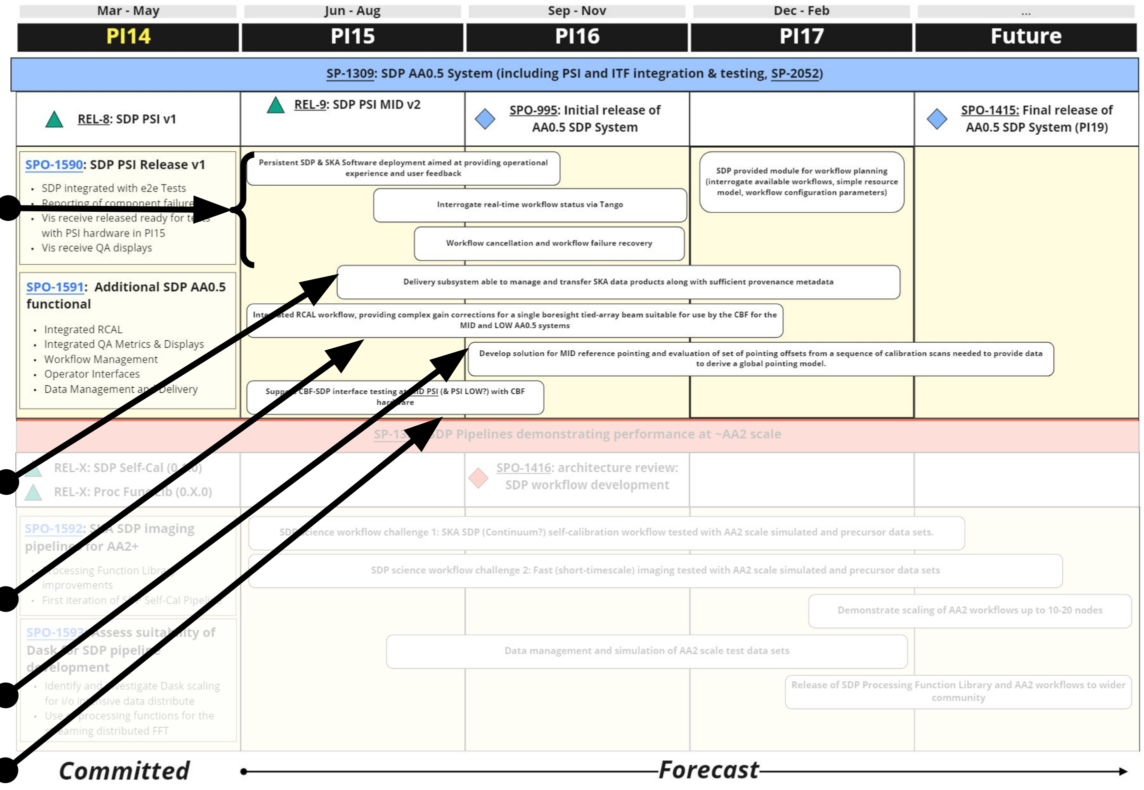

SDP Short Term Roadmap: Workflows @ AA2 scale

1. Establish approach to SDP workflow development initially targeting AA2 scale data sets
   - SKA1 simulations & precursors
   - Establish clean levels of abstraction in software

2. **Starting with self-cal & Fast Imaging** to exercise batch and real-time compute and science performance considerations

3. Collaboration with Services ART for infrastructure & platform (potential for co-design)

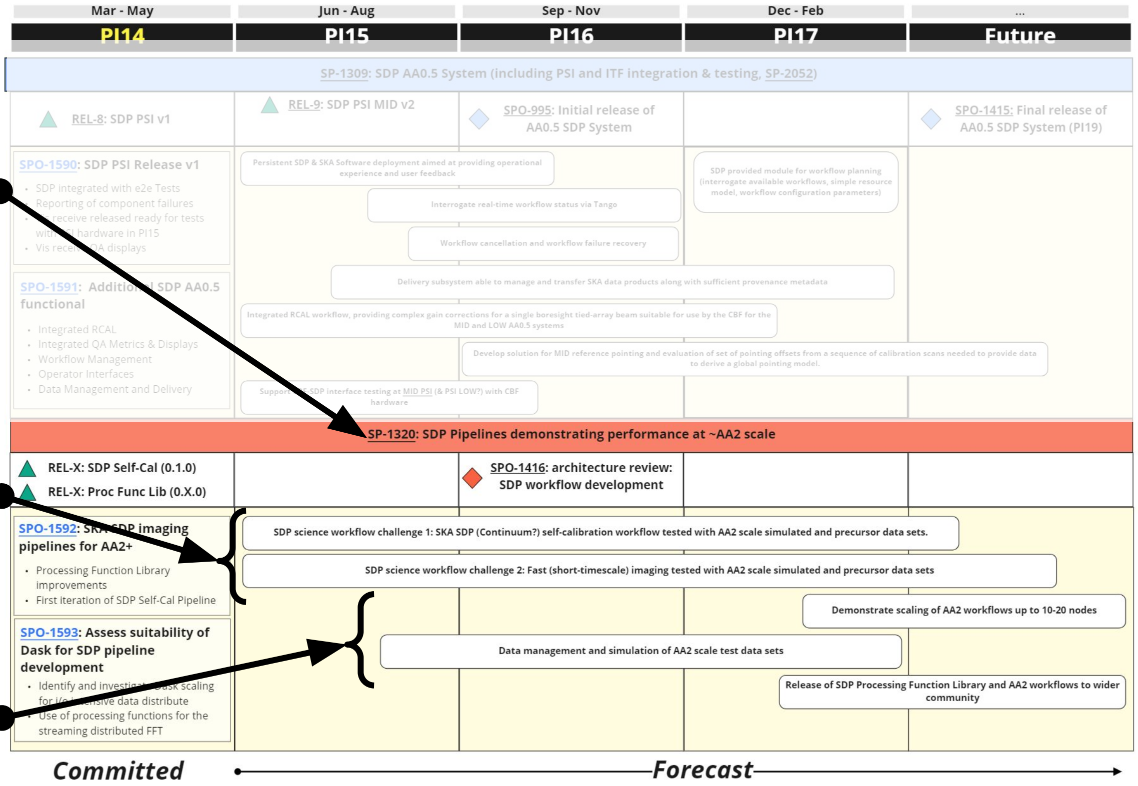

PSS Roadmap to AA0.5... and beyond

- **Pipeline testing with MeerKAT** observations of transients
- **Pipeline testing on MeerKAT servers** (with GPUs/FPGAs)

Testing

- **Continuous Integration (CI)** pipelines
- **Integration/verification** activities at the **Prototype System Integration (PSI)** facilities (AUS/NL/CAN)

Integration

Pipeline functionality:
- Raw data capture
- Single Pulse Search (SPS)
- Acceleration Search (FDAS)

Development

FPGAs/GPUs:
- Optimisation
- Testing
- Pipeline integration (with SPS & FDAS)

Tech research

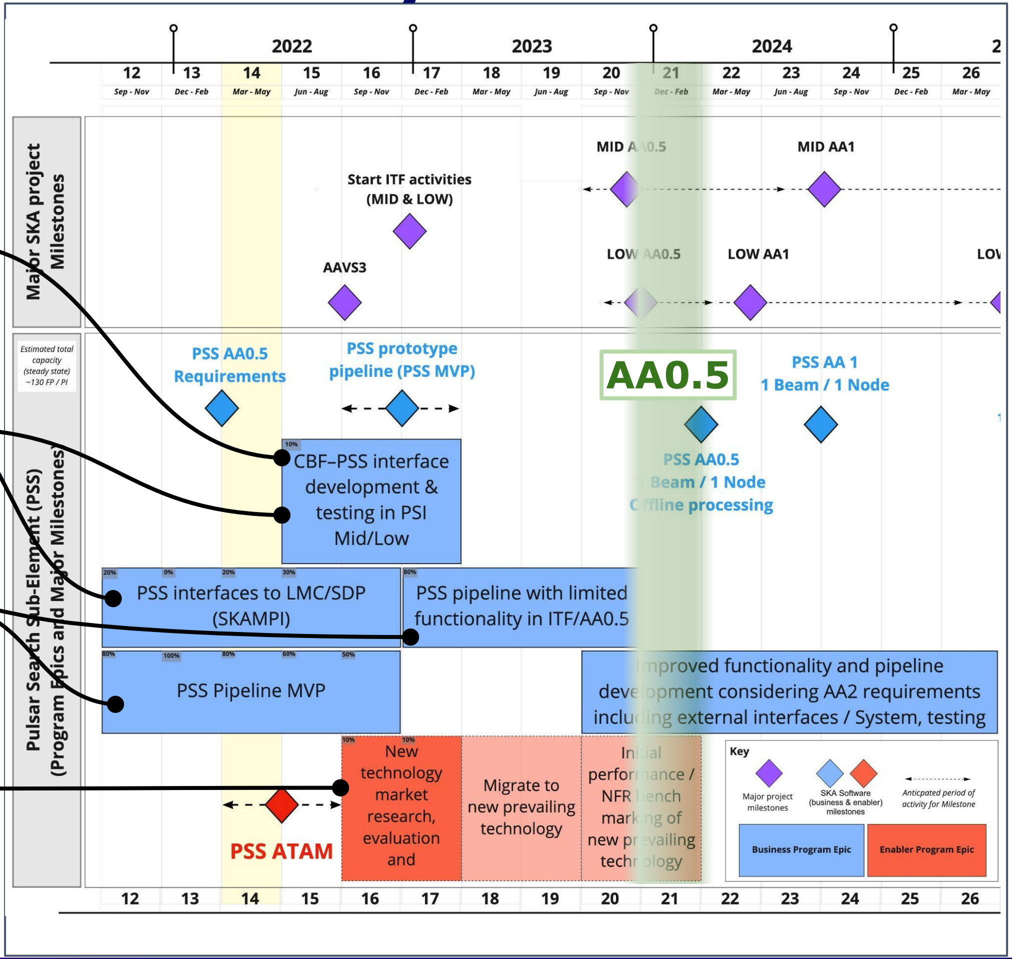

PST Roadmap to AA0.5… and beyond

Low PSI deployment & testing

Testing

Pulsar timing pipeline:
- Voltage recording
- Dynamic spectrum recording
- Pulsar timing

Development

Monitoring & Control SW

Integration

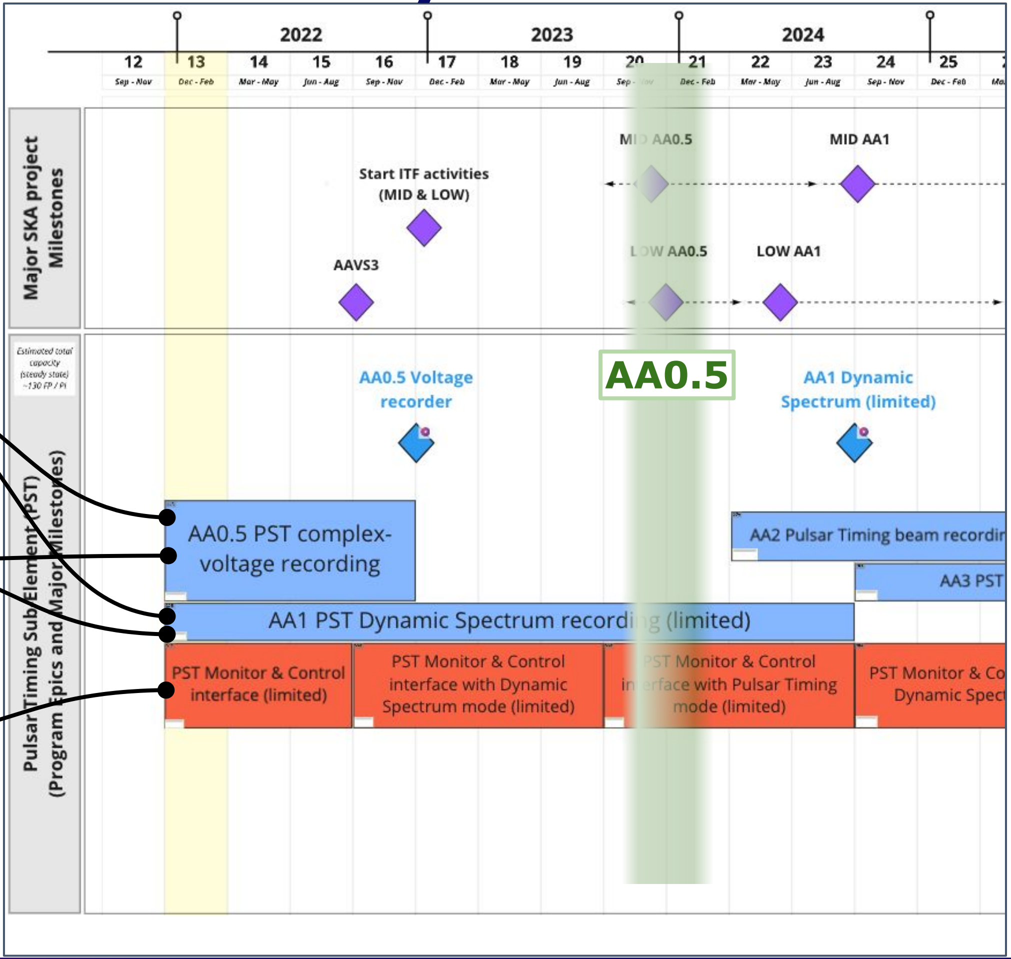

PSS & PST SW activities for SKA Mid & Low

Development

- **SKA Mid/Low receptor modules** for PSS–CBF interface
- **PSS Stokes/Complex voltage writers** for writing out raw data from the Mid/Low CBF to disk (AA0.5–AA1)
- **PST voltage recorder** pipeline for capturing Complex Voltages to disk from PST Mid/Low CBF

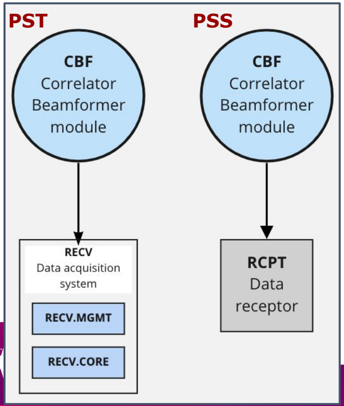

Testing

- **PSS single-pulse pipeline testing on MeerTrap** (MeerKAT) nodes with real data (RRAT observations)
- **PSS testing on ProtoNIP HW** (with FPGAs/GPUs) at MeerKAT
- **PST testing on Low PSI (CSIRO)** with Perentie prototype CBF HW
- **PSS pipeline testing with synthetic test vectors** (pulled from test-vector server in Manchester)

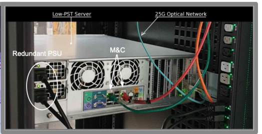

PST server (pst-beam1) at the Low PSI

PSS server (ProtoNIP)at MeerKAT

Integration

- **PSS Continuous Integration (CI) pipeline** running on Manchester server (Kelvin), for continuously executing and testing new versions of the code merged with the PSS repository
- **CSP (CBF, PSS, PST) integration and verification activities** at PSI Low (NL) by the TOPIC team, using CBF HW
- **PSS & PST Monitoring & Control SW** being developed by OMC and PST in order to integrate pulsar processing with TMC

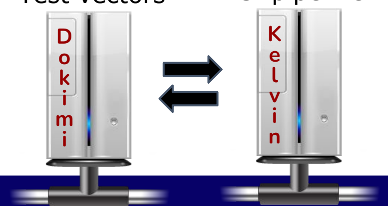

Tech research

- **PSS: FPGA and GPU HW accelerators** are being optimised and benchmarked in order to select the ones that can deliver the required performance at AA4
- **GPU-accelerated tests of the PSS** single-pulse pipeline at MeerKAT

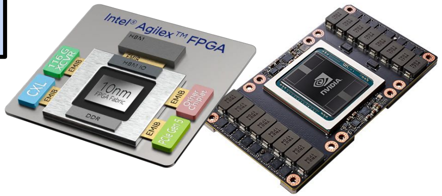

Observation Management and Control (OMC) ART Roadmaps

Team CIPA | Team MCCS | Team Buttons | Karoo Crew – Team Karoo | Team NCRA | Cream Team – Team MCCS | Team Perentie | TOPIC Embedded Systems – Team Topic | Team NAKSHATRA

PMs: Adam Avison, Vivek Mohile, Pamela Klaassen, Gerhard le Roux & Giorgio Brajnik
Architects: Alan Bridger & Sonja Vrcic

OMC Long-Term Roadmap

- OMC Roadmap updated in a more goal-oriented way.
- **Currently:** Goal focussed, split into "epics" based loosely around subsystems.
- **Coming Soon:** Use Case / User functionality focussed structure.

Link to Miro Board

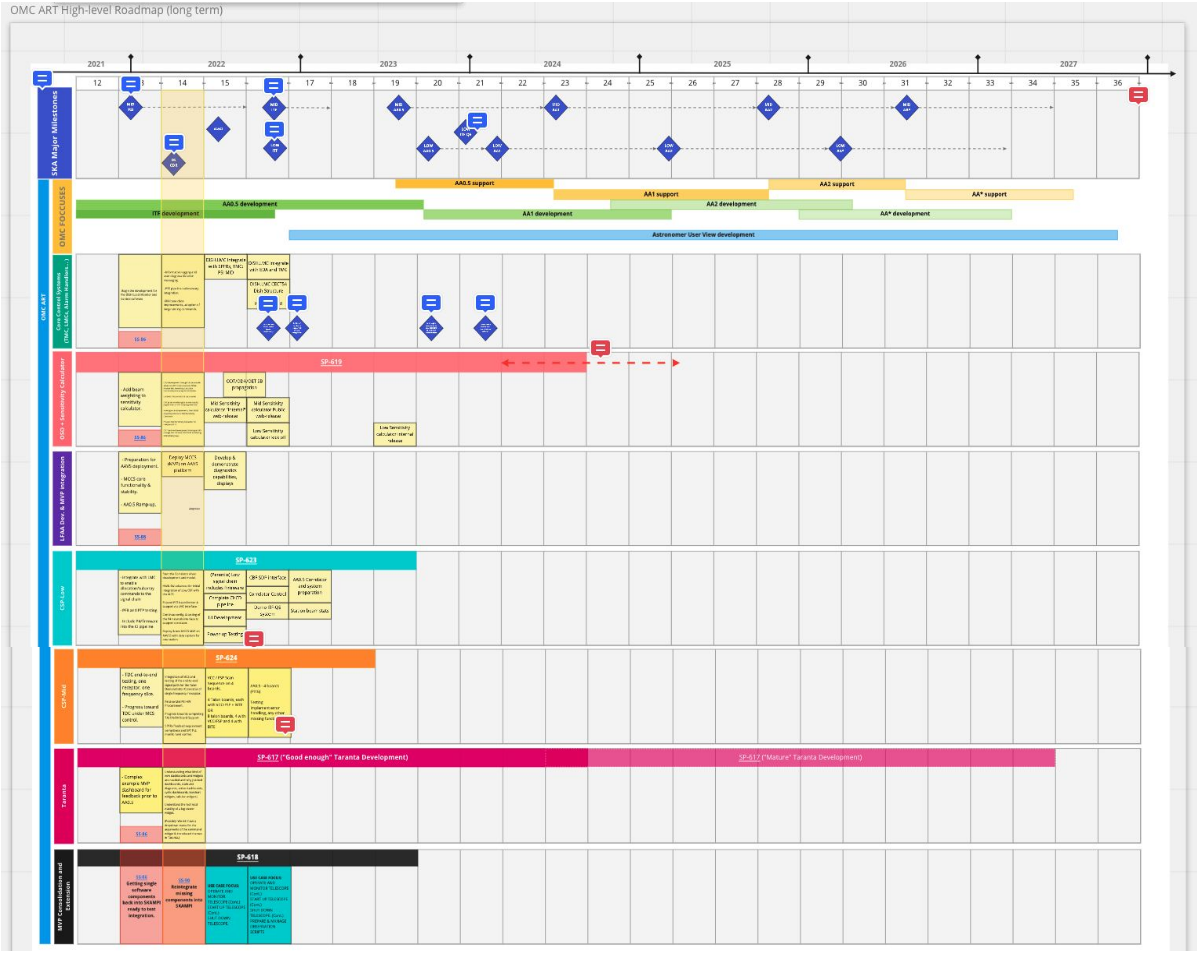

What are we targeting at AA0.5

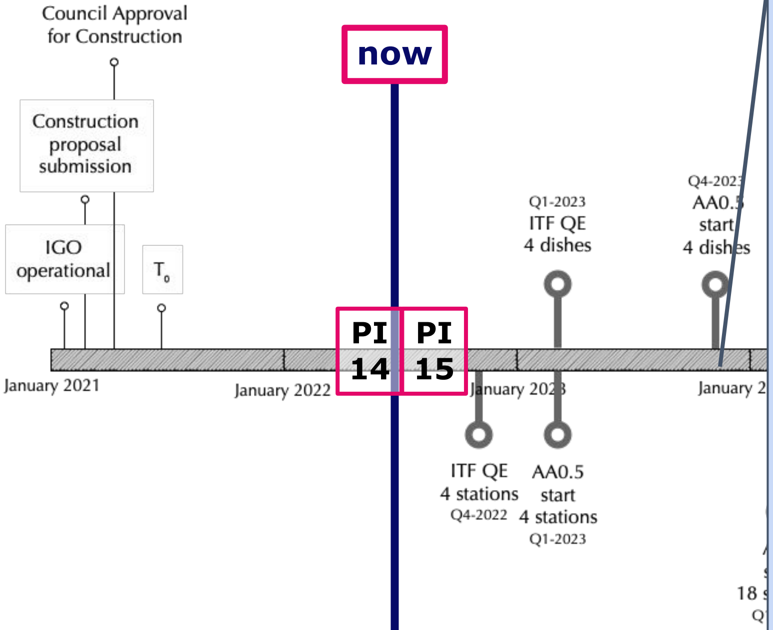

At AA0.5 OMC train / teams are focused on supporting* :

- Software for deployment of minimal arrays on-site
- Primary goal: end-to-end test of interferometry (and beam-forming)
- Initial control system primarily to support commissioning
- Validate key interfaces
- Verify fundamentals of system performance in a realistic operating environment (e.g. RFI, wind, temperature, ...)
- Reduce risk and identify potential failures including functional & non functional (scalability, reliability etc.)
- Develop AIV, Commissioning teams and procedures

*Adapted from a slide in Robert Laing's presentation at PI14 planning

verti

wrt Array Assembly size

Highlights of some items we are focused on near term

- **Central Signal Processor (CSP):** End to end testing of the signal chain with simulated data
- **Observatory Science Operations (OSO):** Integration of vertical slice from ODT to OET using the ODA, to propagate Scheduling Blocks
- **Telescope Monitor and Control (TMC):** Re-integration and testing (after some critical refactoring of the software) of a full scan functionality from On/Off through to Assign resources to sub-array, Configure Scan and Scan

ODT = Observation Design Tool,
OET = Observation Execution Tool,
ODA = Observation Data Archive

OMC: MID Specific Roadmap zoom-in

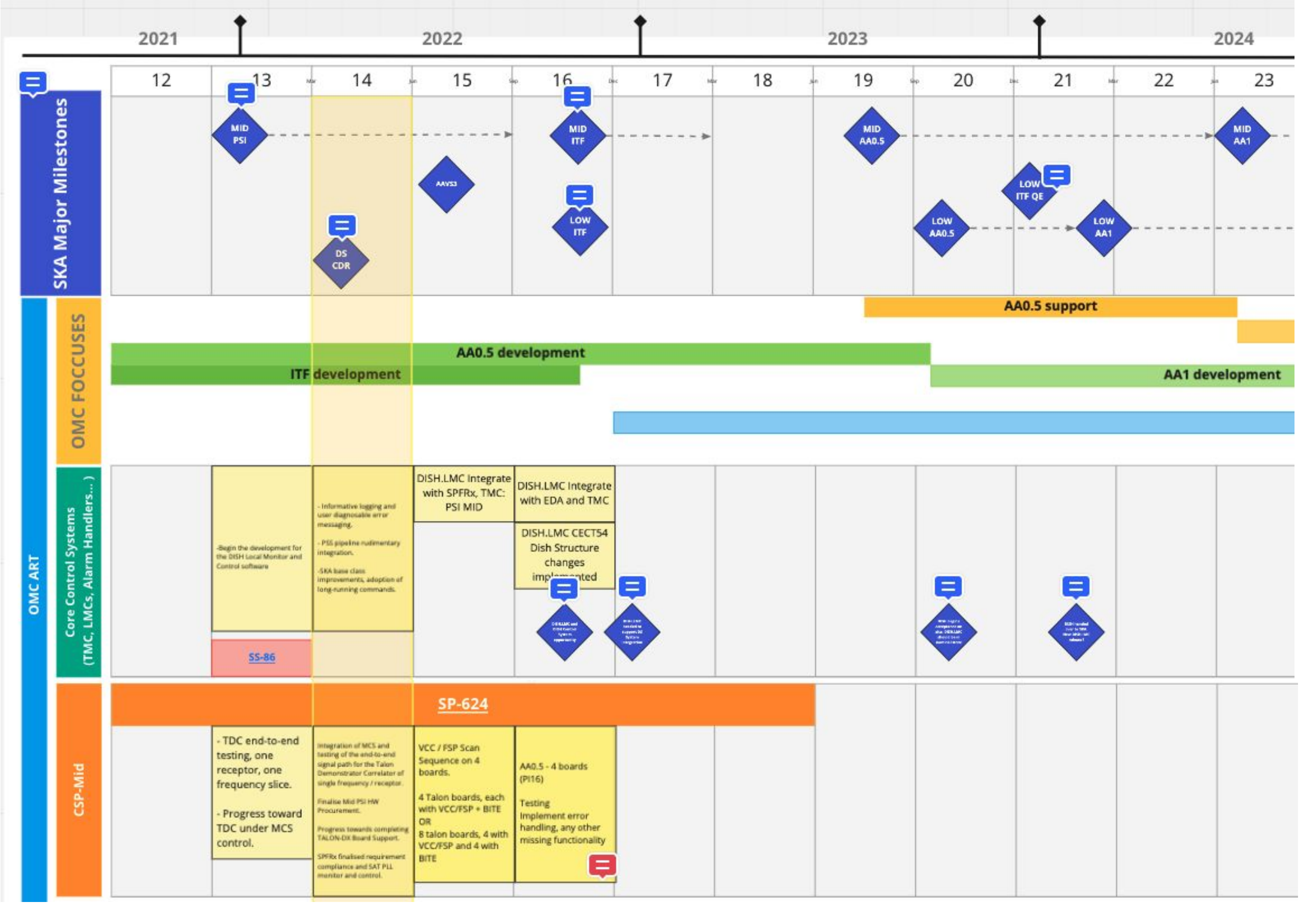

CSP.MID for AA0.5

- Testing of end-to-end signal chain. Targeting full AA0.5 signal processing capabilities by end of PI16.
- Integration of TDC with MCS and later CSP.LMC control.
- AA0.5 capabilities available on the required number of boards.

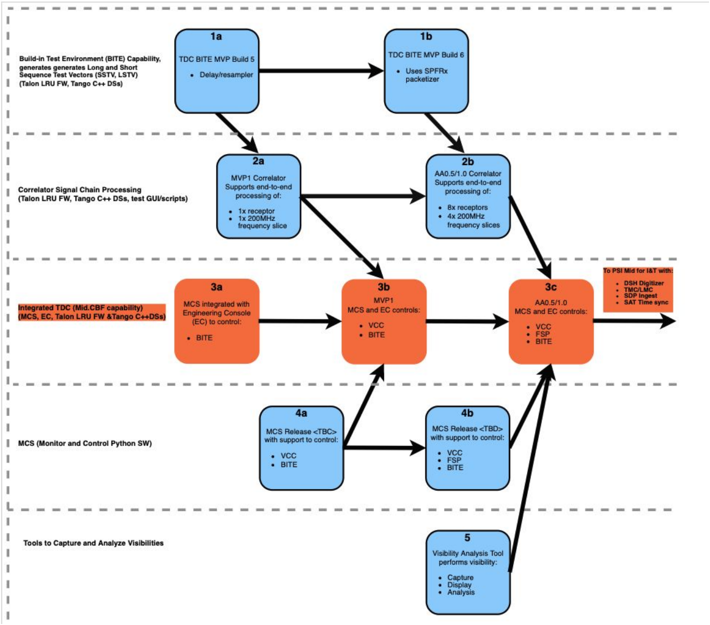

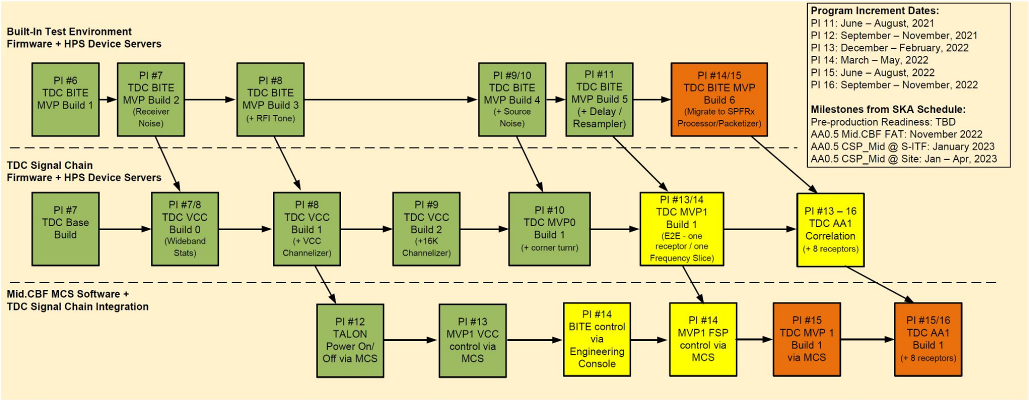

Thanks to CIPA team (Nicolas Loubser) https://confluence.skatelescope.org/display/SE/TDC+AA0.5+Rollout+Plan and https://confluence.skatelescope.org/display/SE/TDC+Talon+LRU+based+Releases+for+AA0.5

Dish LMC roadmap

For AA0.5 modify early 2019 C++ version of Dish LMC software.

Later gradual migration to Python implementation based on ska-tango-base.

Drivers for the approach (See Dish LMC Roadmap)

At AA 0.5:

- TANGO framework allows easy integration of TANGO devices that are written in different languages (in this case C++ and Python)
- Maturity of Dish LMC software (in terms of functions supported and issues resolved after having integrated with other Dish subsystems),
- Less time for developing mature Python based Dish LMC at AA0.5

Later Migration:

- Benefit brought by using the Python implementation of ska-tango-base, and making Dish LMC compliant to latest SKA Software Standards and Control System Guidelines
- Python competency of SKA SAFe teams.

Dish LMC roadmap

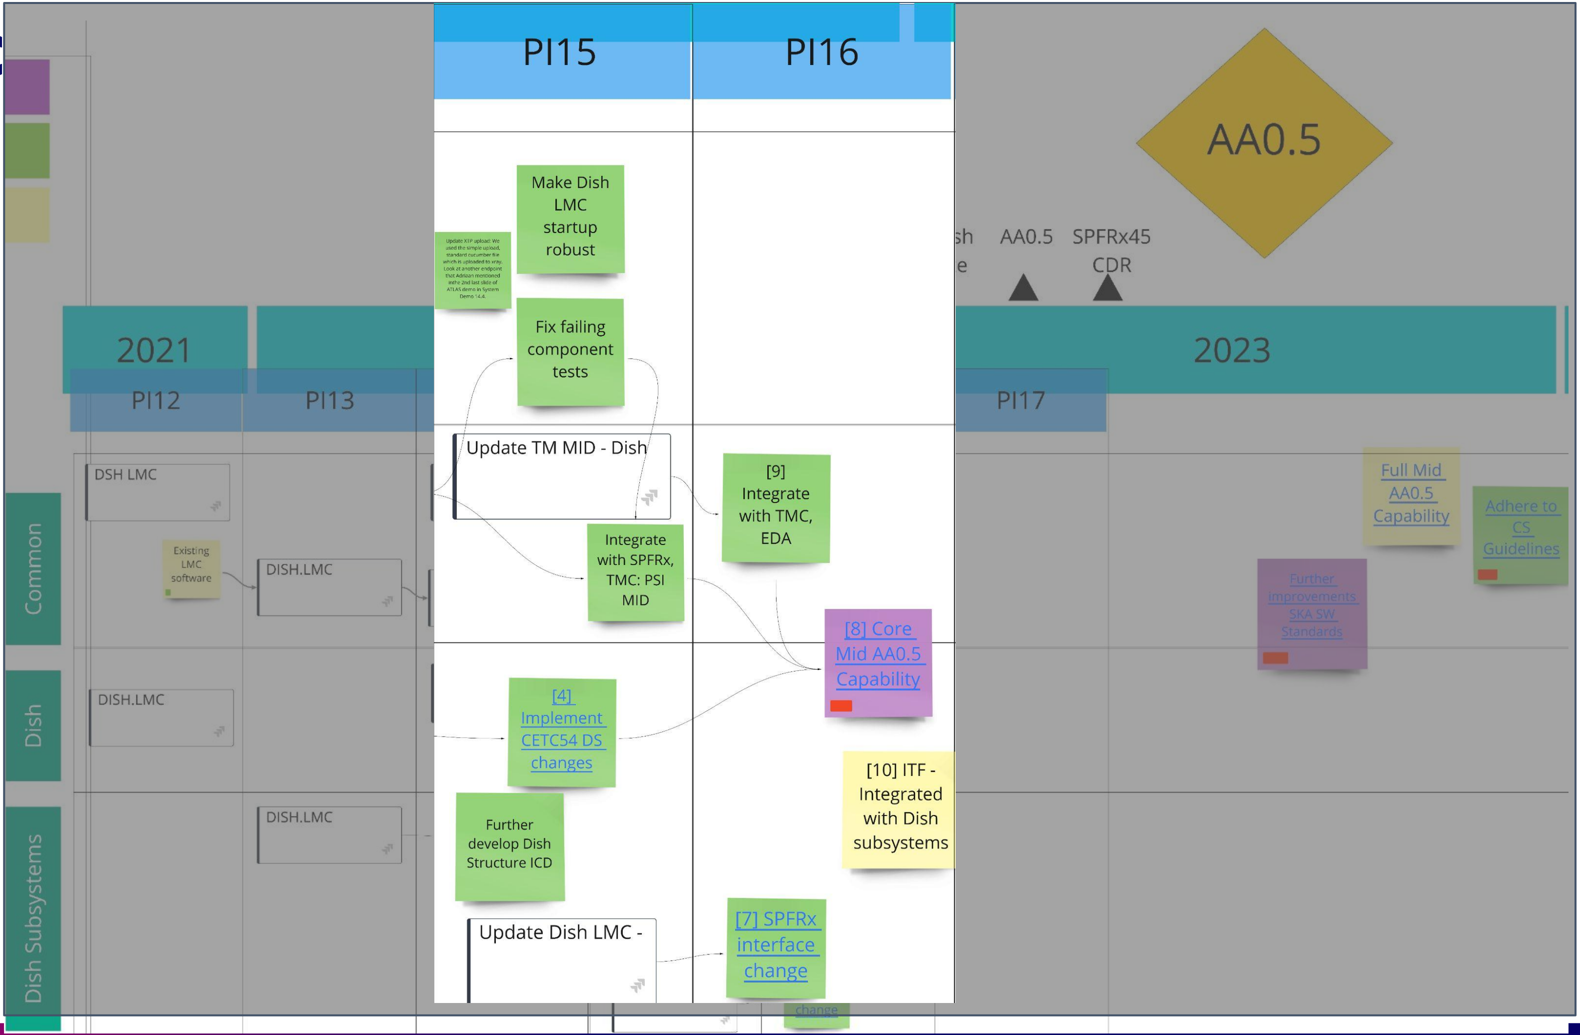

Common OMC Products at AA0.5

- **TMC:** Will be a reliable package enabling the basic operation and monitoring of the telescopes.

- **OSO:**
  - Allows the creation/editing/storage of Scheduling Block Definitions (SBDs) for AA0.5 capabilities.
  - Allows creation/editing/storage of Observing Scripts.
  - OET allows generation and execution of SB Instances (SBIs) derived from SBDs, but can drive observations without SBD/SBIs
  - Initial Shift Log tool available.

Dish LMC roadmap PI15 PI16 plans

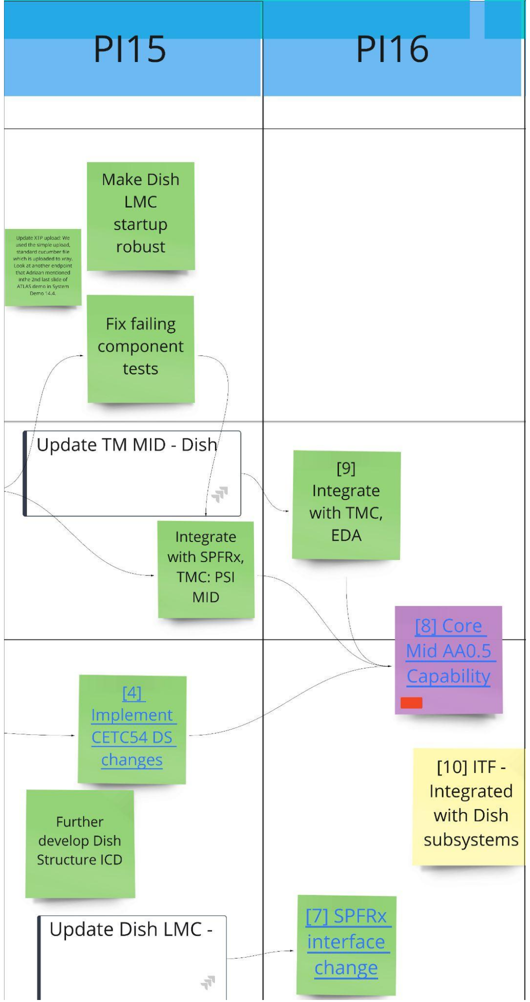

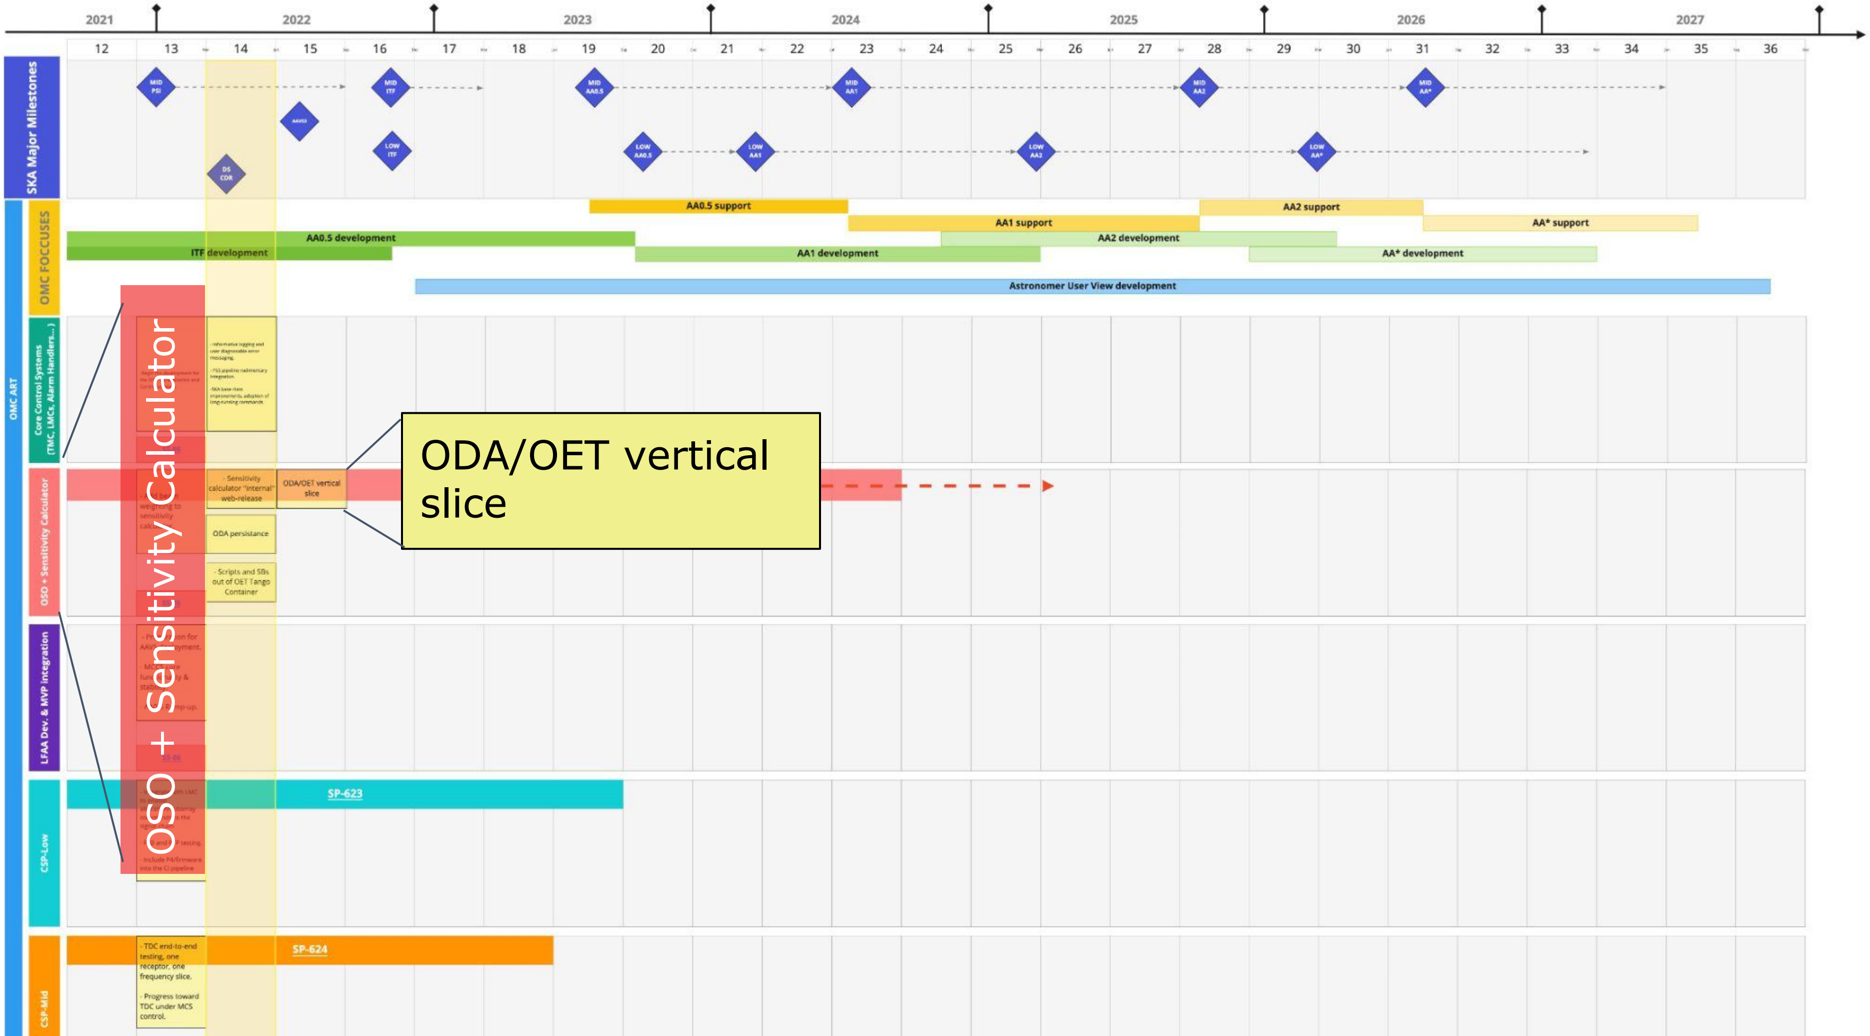

2021
2022
2023
2024
2025
2026
2027
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
SKA Major Milestones
OMC FOCCUSES
OMC ART
Core Control Systems (TMC, LMCs, Alarm Handlers…)
OSO + Sensitivity Calculator
LFAA Dev. & MVP integration
CSP-Low
CSP-Mid
AA0.5 support
AA2 support
AA1 support
AA* support
AA0.5 development
AA2 development
ITF development
AA1 development
AA* development
Astronomer User View development
- Sensitivity calculator "internal" web-release
ODA/OET vertical slice
ODA persistance
- Scripts and SBs out of OET Tango Container
SP-623
SP-624
- TDC end-to-end testing, one receptor, one frequency slice.
- Progress toward TDC under MCS control.
ODA/OET vertical slice
OSO + Sensitivity Calculator

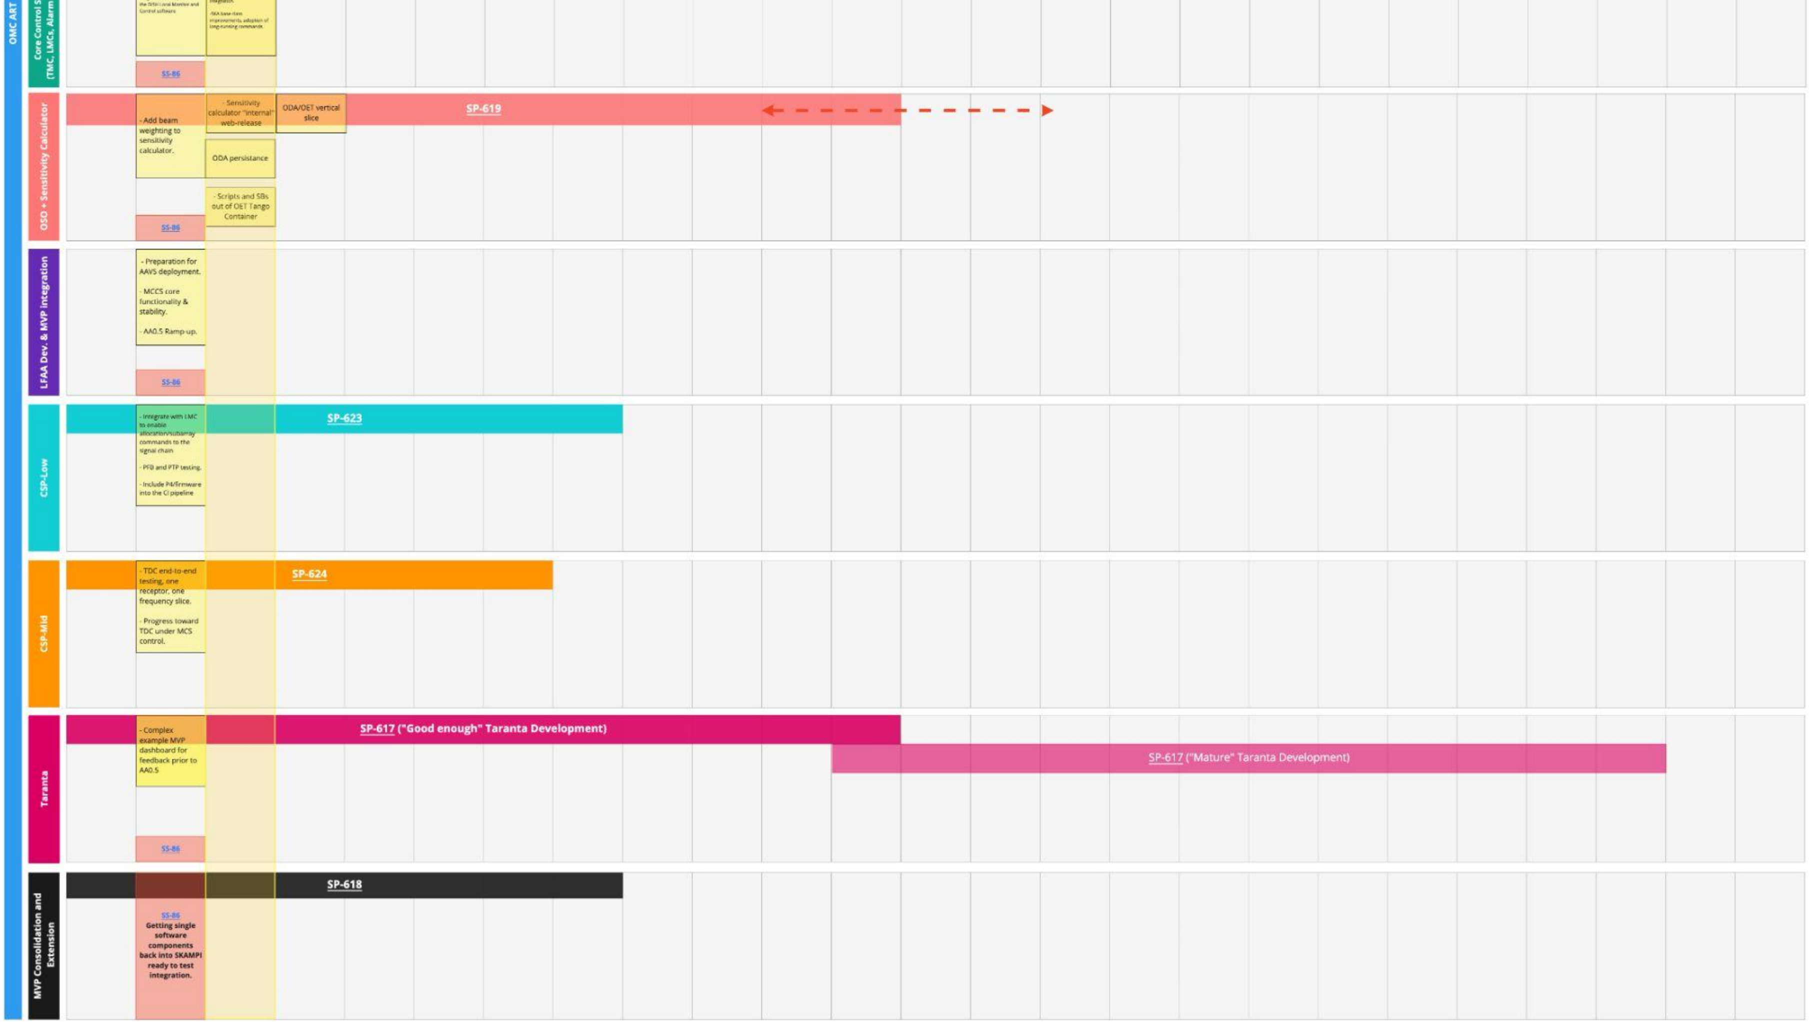
OMC ART
Core Control S (TMC, LMCs, Alarm
SS-86
OSO + Sensitivity Calculator
- Add beam weighting to sensitivity calculator.
- Sensitivity calculator "internal" web-release
ODA/OET vertical slice
SP-619
ODA persistance
- Scripts and SBs out of OET Tango Container
SS-86
LFAA Dev. & MVP integration
- Preparation for AAVS deployment.
- MCCS core functionality & stability.
- AA0.5 Ramp-up.
SS-86
CSP-Low
- Integrate with LMC to enable allocation/subarray commands to the signal chain
- PFB and PTP testing.
- Include P4/firmware into the CI pipeline
SP-623
CSP-Mid
- TDC end-to-end testing, one receptor, one frequency slice.
- Progress toward TDC under MCS control.
SP-624
Taranta
- Complex example MVP dashboard for feedback prior to AA0.5
SP-617 ("Good enough" Taranta Development)
SP-617 ("Mature" Taranta Development)
SS-86
MVP Consolidation and Extension
SP-618
SS-86
Getting single software components back into SKAMPI ready to test integration.

SKA MID Software Architecture

SKA Software - Key components

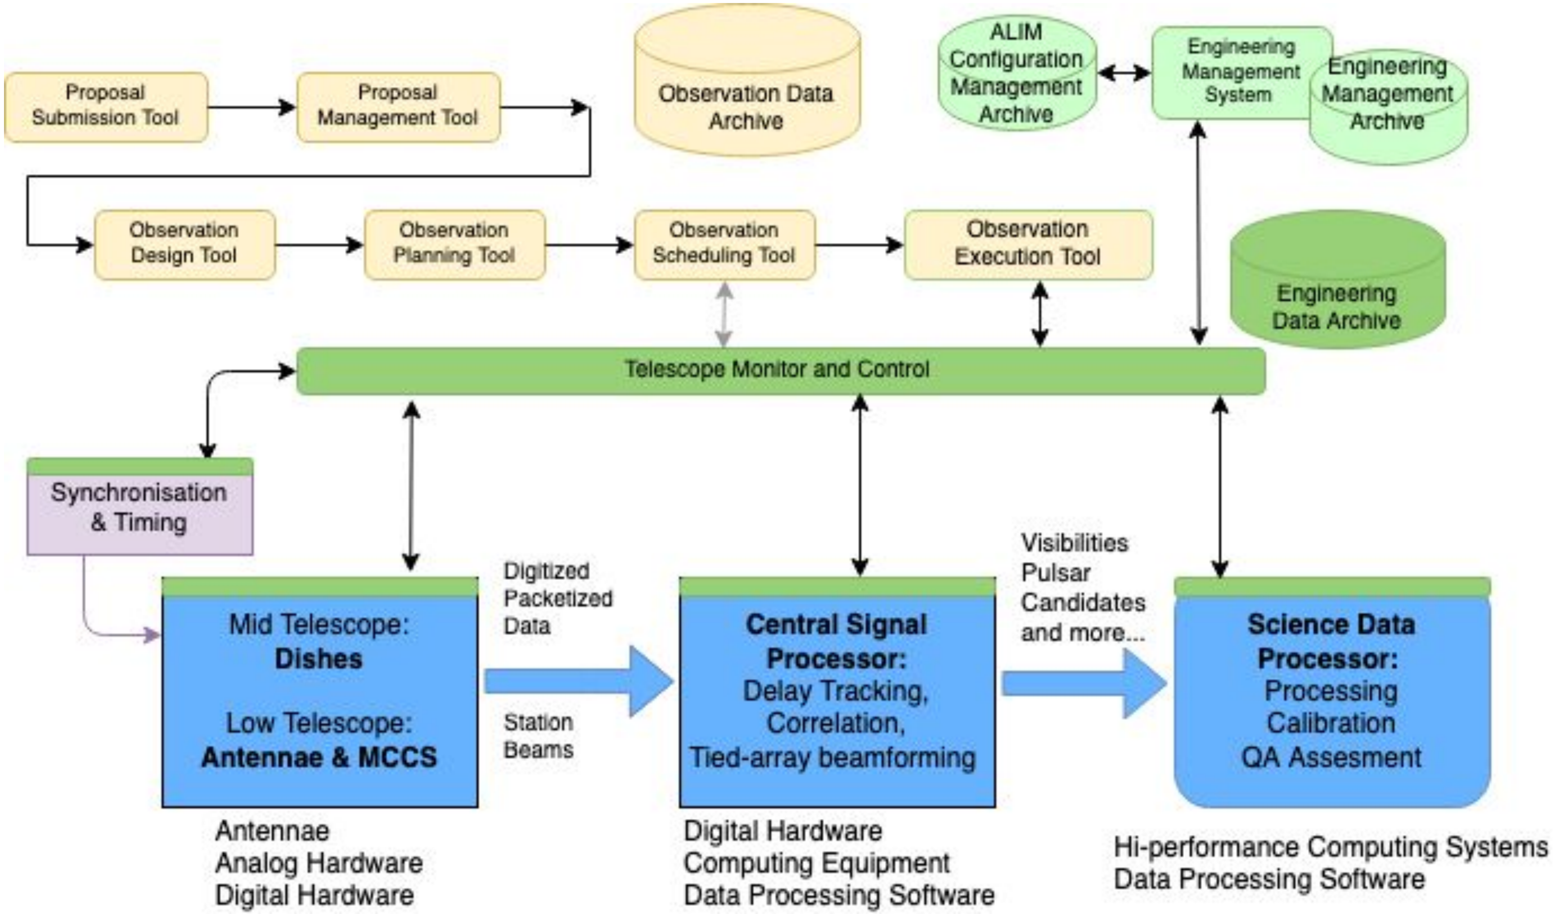

SKA Software - Key components

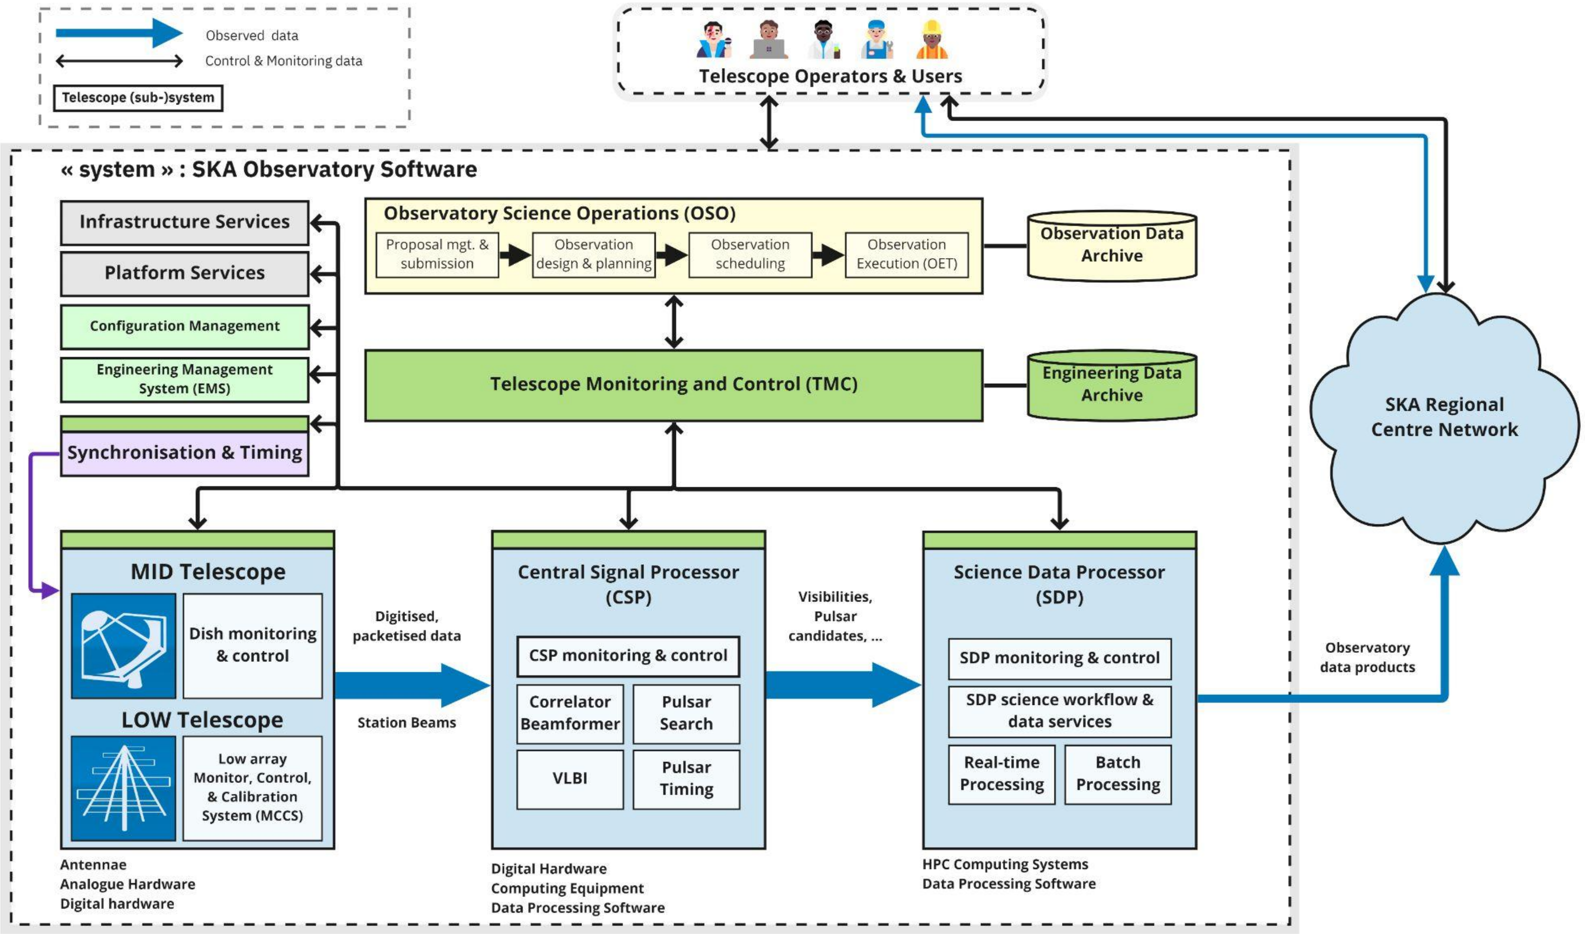

SKA Software modules view detail

https://confluence.skatelescope.org/display/SWSI/Views%3A+Module

Observation Management

Telescope Control System

Data Processing

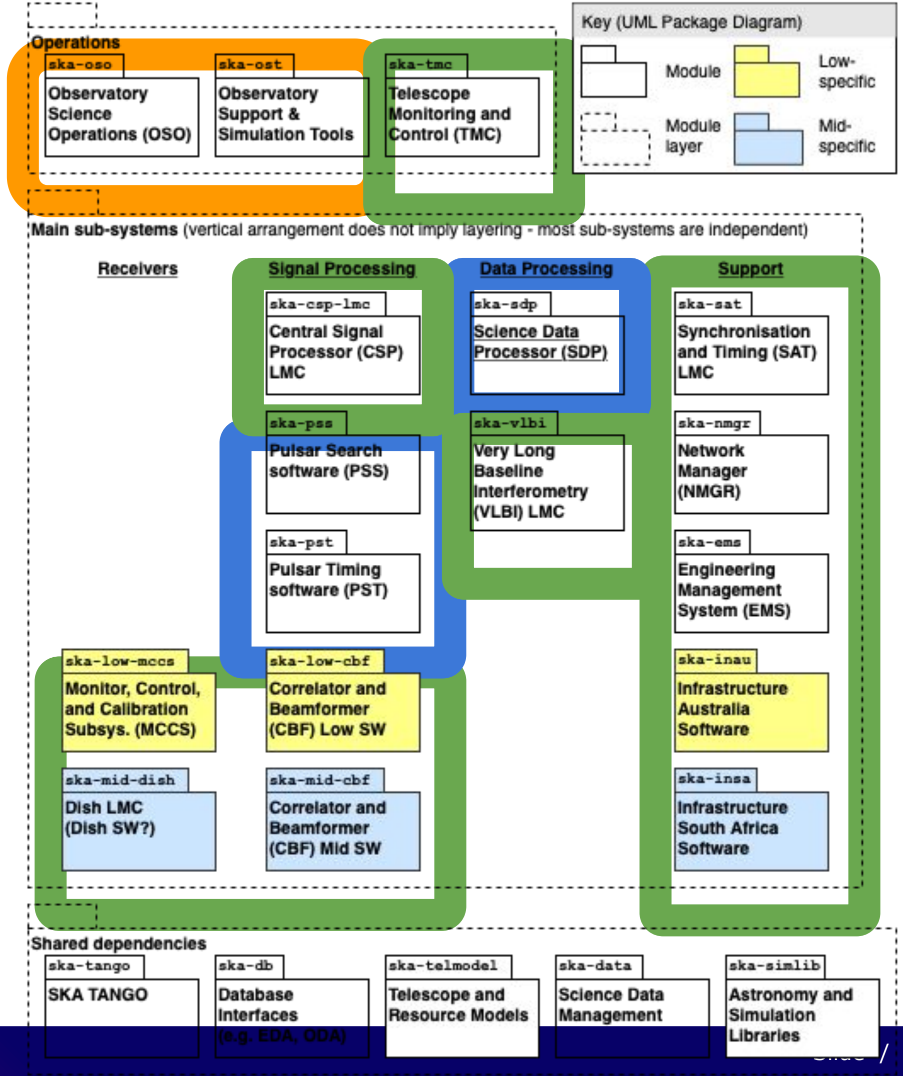

Observation Monitoring and Control

- OSO - Observatory Science Operations
- TMC - Telescope Monitoring and Control

- What is the purpose of the OSO and TMC?
  - Manage the design, scheduling and execution of all observing on both telescopes.
  - Correctly operate and monitor the observatory and telescopes.

- Why are they important for the telescope to function correctly?
  - Key interface between the science community users and the observatory.
  - Ensure efficient, complete and reliable data and meta-data collection.

Proposal Observation Flow

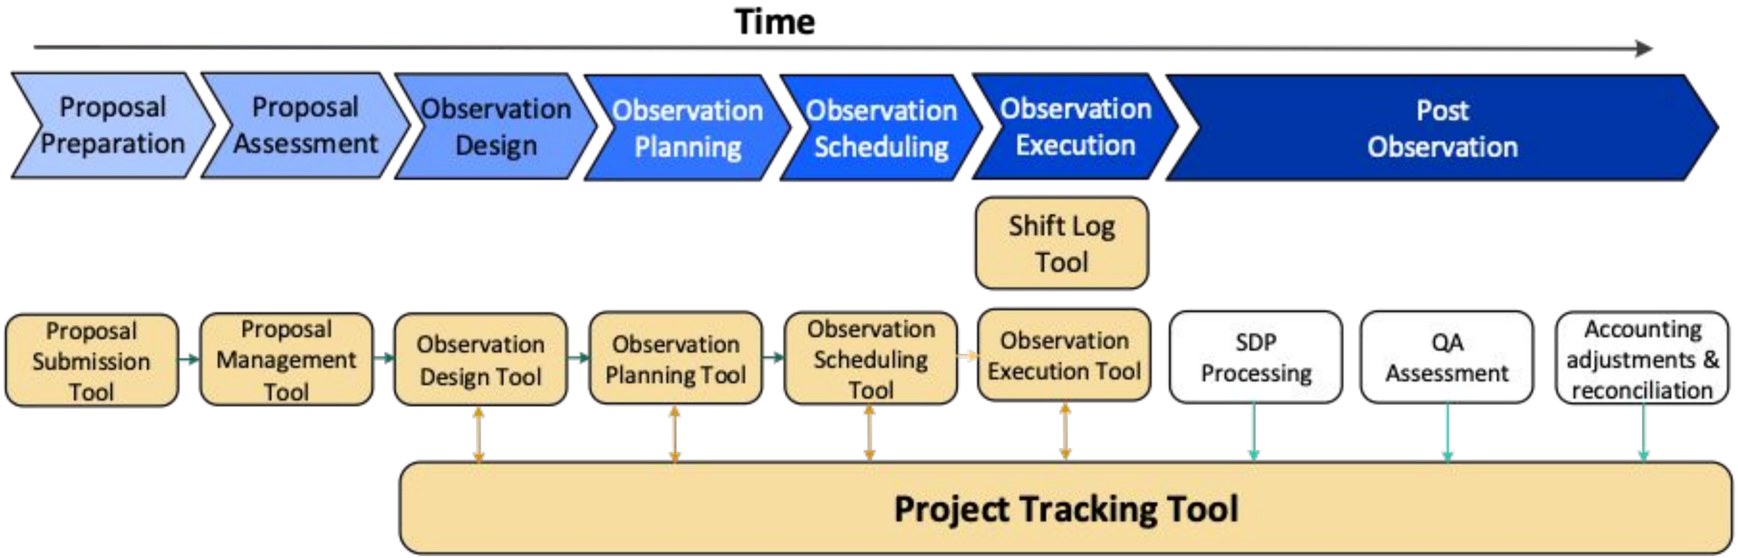

OSO is the suite of software tools that support the entire science process from proposal submission to data-product delivery.

Telescope Control System

The role of the Telescope Control System may be classified in the three key areas:

- ❑ Instrumental monitor and control
- ❑ Observation configuration, execution and monitoring
- ❑ Support for sub-arraying

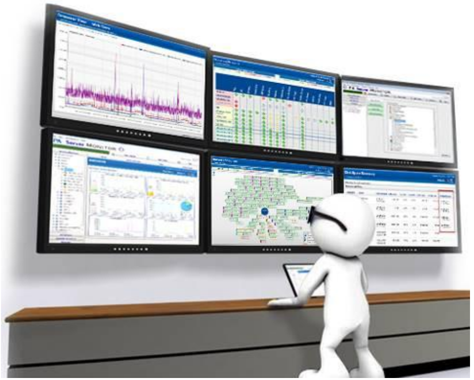

Implementation is based on the

TANGO controls framework TANGO

Science Data Processor

SDP performs "main" data reduction (~factor 40) - challenges:

Performance & Scalability

- Compute, I/O & Storage
  - >10 Pflop/s effective
  - ~0.77 TB/s ingest rate
  - ~4 TB/s into processing
  - >40 PB tiered buffer
- Need to scale
  - Trivial (e.g. Ingest) and expensive (e.g. ICAL) workflows co-exist
  - SKA ">1" will be even harder on SDP

Modifiability & Maintainability

- Long lifespan (>50 yrs)
- Software changes
  - Execution Engines
  - Science Workflows
  - Processing Components
  - Data Models
- Hardware changes
  - Processing
  - Storage
  - Network

Buildability, Affordability

- COTS components
  - Get going quickly
  - Externalise/share maintenance
- Support agile development
  - Small functional increments
  - Parallel work of different teams

Testability

- Continuous Integration
  - Test outside SDP
  - Support different development speeds
- Ensure scientific validity
  - Must trust pipelines with autonomous analysis

Portability

- SKA Low, SKA Mid, SRCs

SKA MID Software in AA0.5

- **TMC** needs to be reliable and enable operations and monitoring
- **OSO** needs to provide enough functionalities to enable commissioning activity
- **SDP** is needed to enable calibration, quality monitoring and data ingest and storage

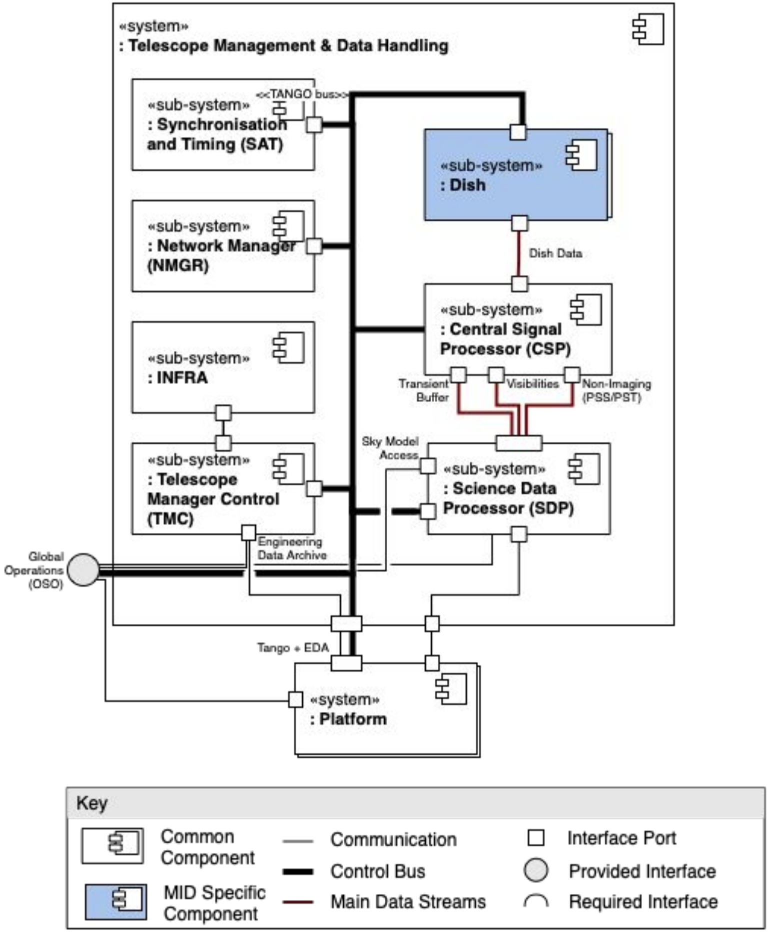

SKA Software Architecture AA0.5 - Telescope Management and Data Handling

AA0.5 Deployment

- Current baseline
- Showing also some essential external components

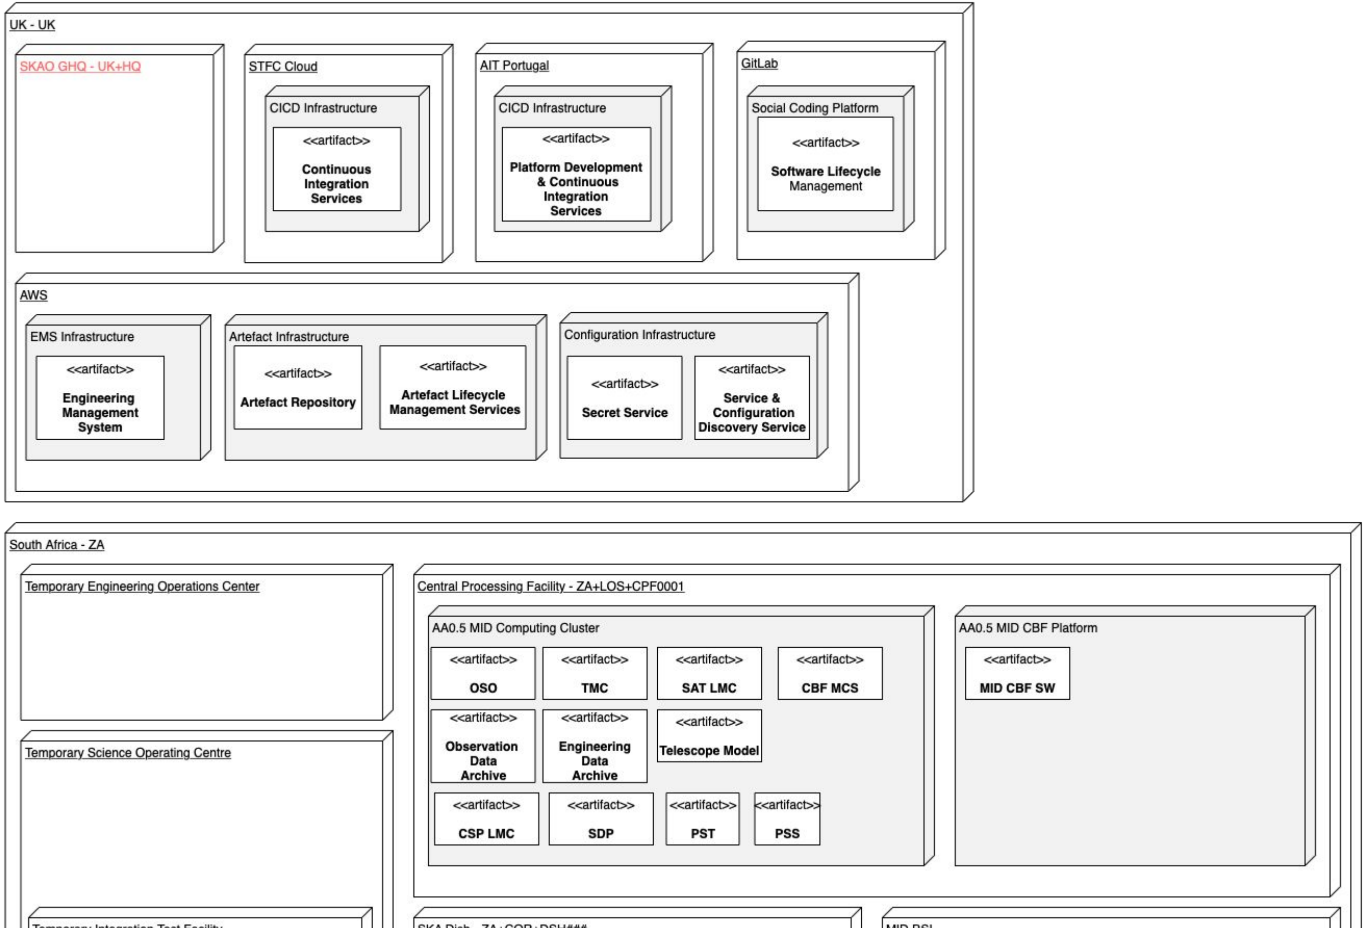

[SKA Software Architecture AA0.5 - Deployment View]()

OSO capabilities at AA0.5

- Observation Design Tool (ODT) ability to create/edit/store SB Definitions consistent with the AA0.5 system capabilities.
  - Expected to be first version editors (i.e. simple)
- Ability to create/edit/store observing scripts
- Observation Execution Tool (OET) ability to retrieve and execute SB Instances derived from these definition
- OET ability to execute observing scripts outside SB context
- Initial version of Shift Log Tool (SLT) available to log operator comments.

OSO Deployment

- Assume Observation Data Archive (ODA) storage at temporary CPF, but could be cloud-based.

- Assumed all UI clients (ODT, OET, SLT) available at temporary CPF. ODT and SLT clients additionally at SOC (and possibly SKAHQ?)

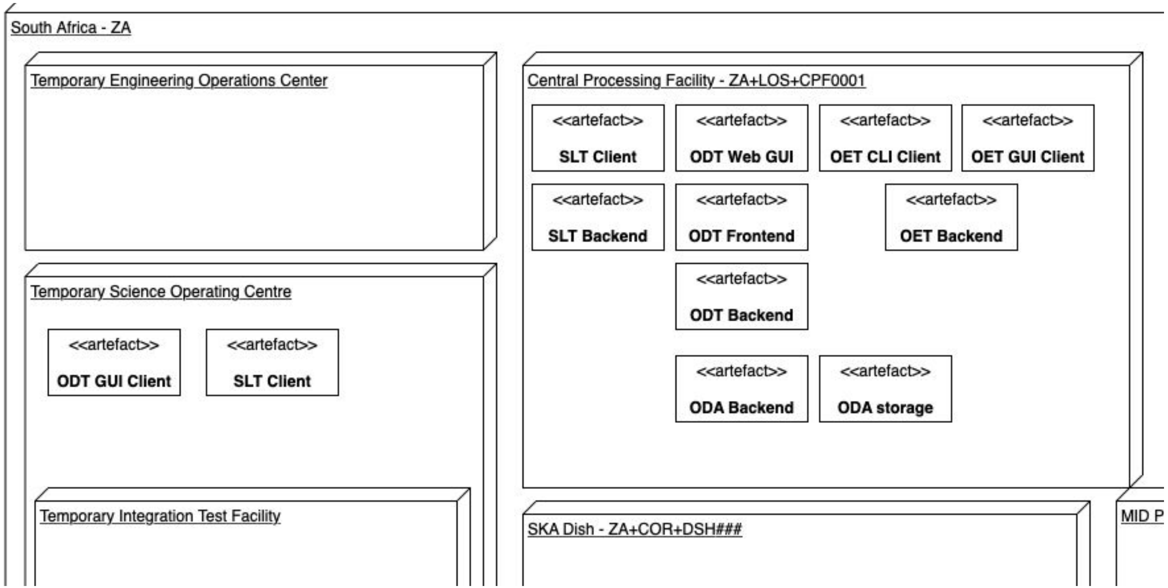

OSO Architecture AA0.5 Deployment View

Overall TMC architecture AA0.5

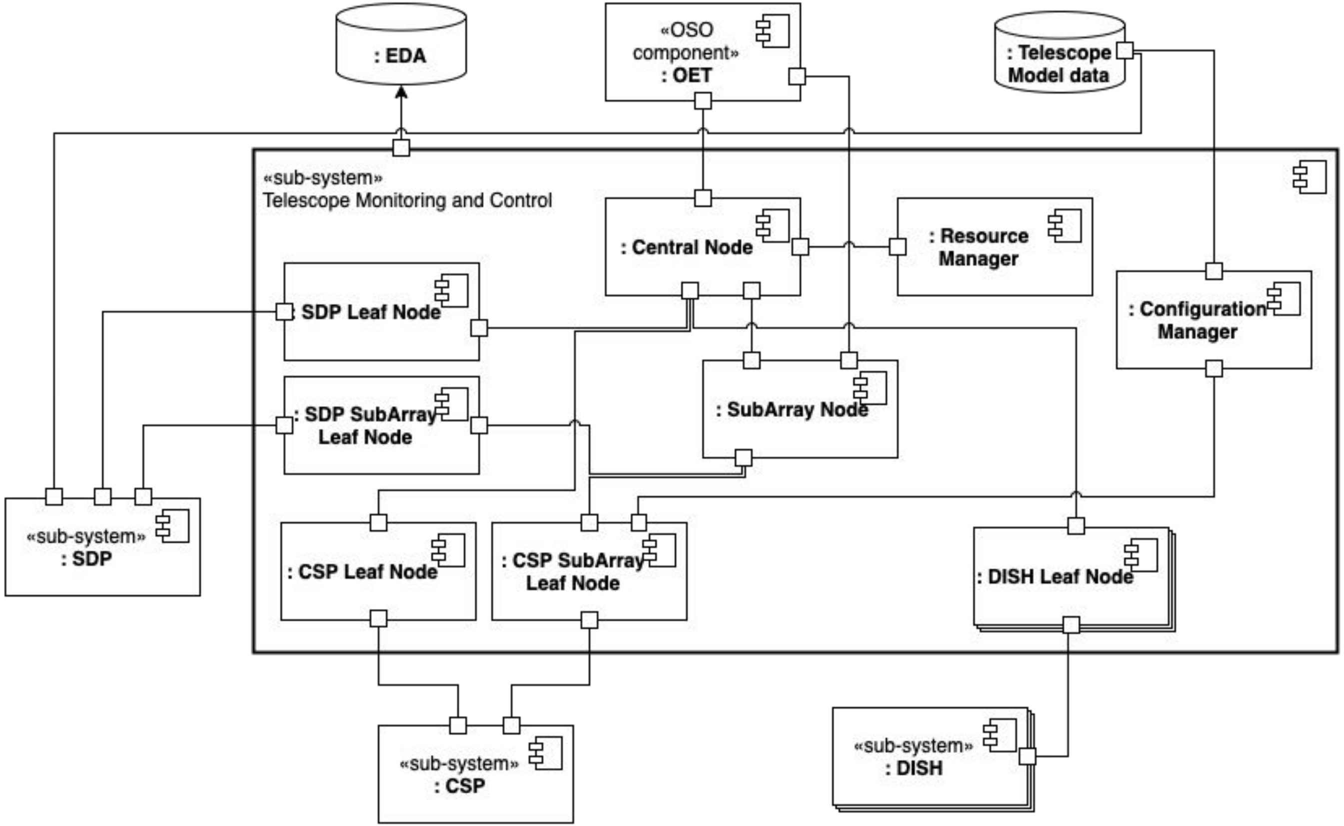

[TMC Architecture AA0.5 - C&C View]

Dish LMC

- Actively being implemented
- Captures new DS design and interfaces
- Great progress demonstrated in the Intermediate DS review, but highlights future areas of work:
  - Verification strategy
  - DS remote ICD details

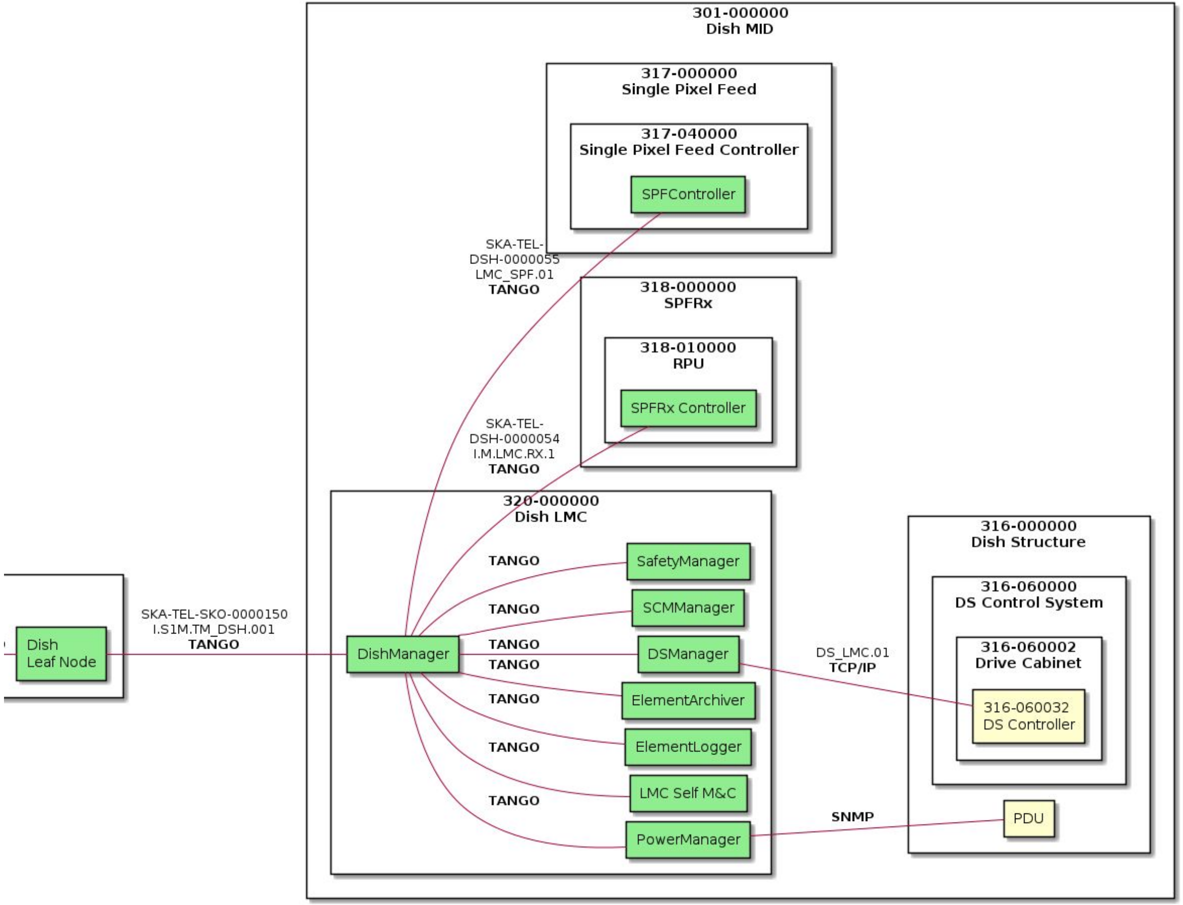

MID Dish Control - Solution Intent pages

Overall SDP Architecture AA0.5

Notable features

- Batch processing
  Simple CLI
- Storage
  Shared storage
- Delivery
  Synchronise with external location(s)?
- Quality information displays

Robust fallbacks (e.g. control via CLI)

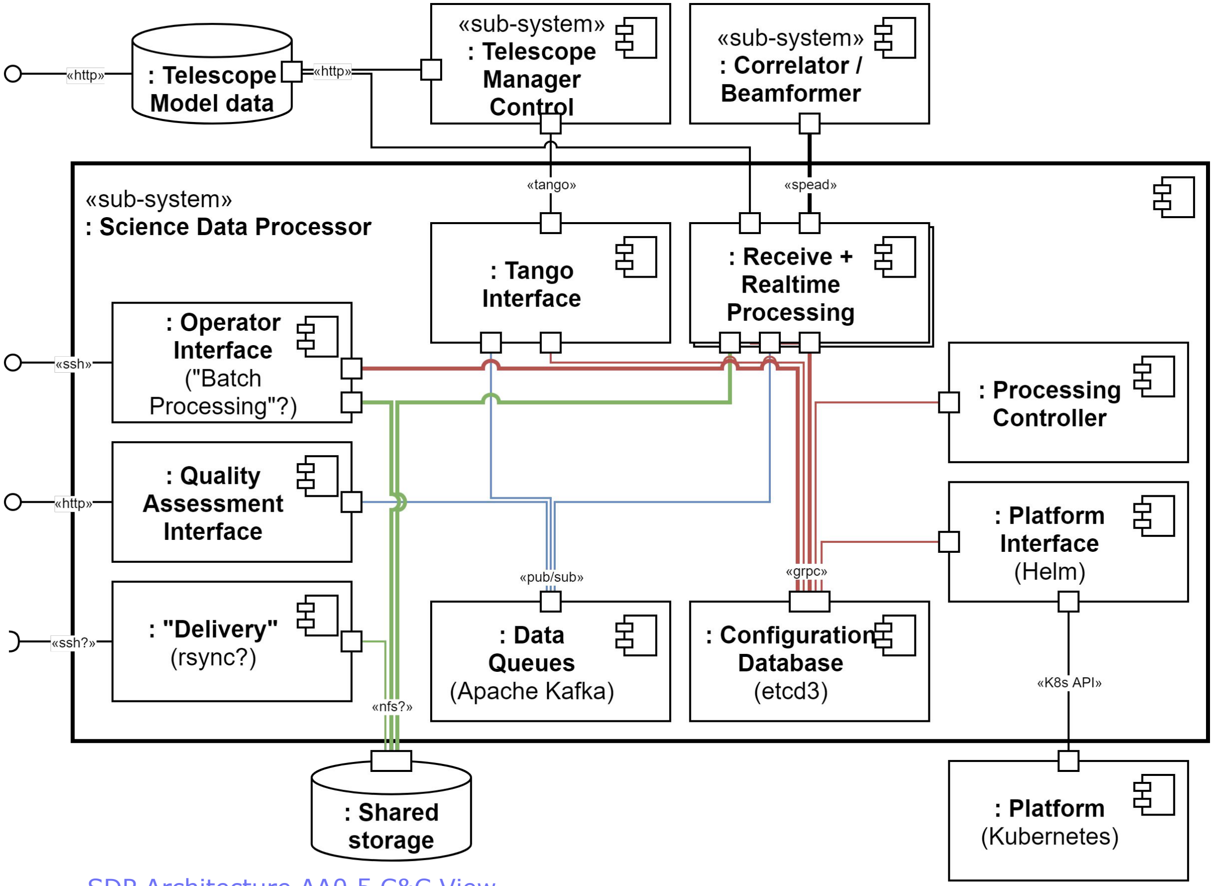

SDP Architecture AA0.5 C&C View

Pulsar Processing for AA0.5

- Main driver at AA0.5 is to verify beamforming functionalities
- Missing link to SDP
- Processing can happen offline, needs storage functionality

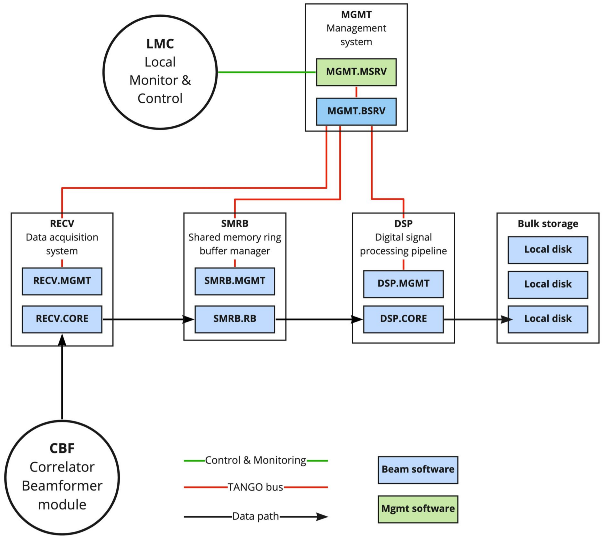

PST Architecture at AA0.5

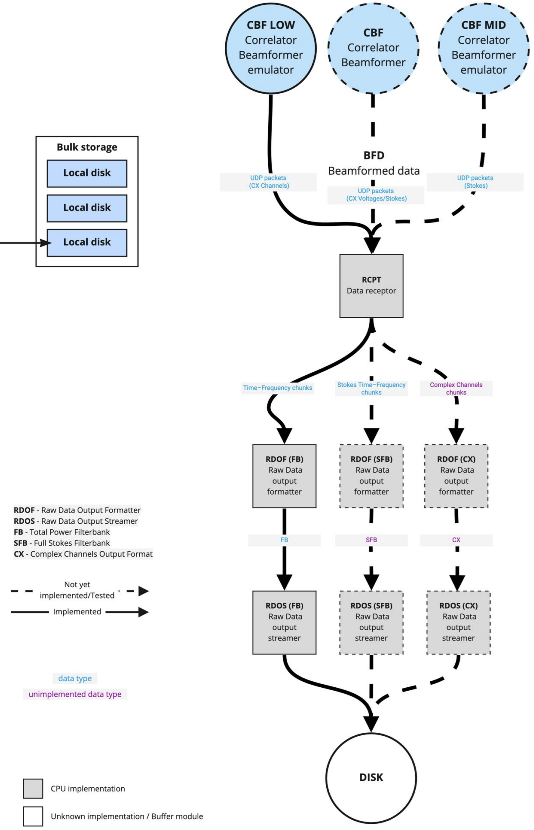

PSS Architecture at AA0.5

Developing the system

- Addressing detailed design decisions and gaps in the design of the software system
- Evolving our understanding of the system in successive iterations
  - ADR-57 Decide functional allocation of MID dish pointing corrections DECIDED
  - ADR-56 CBF to SDP Interface for Early Array Releases ANALYZING

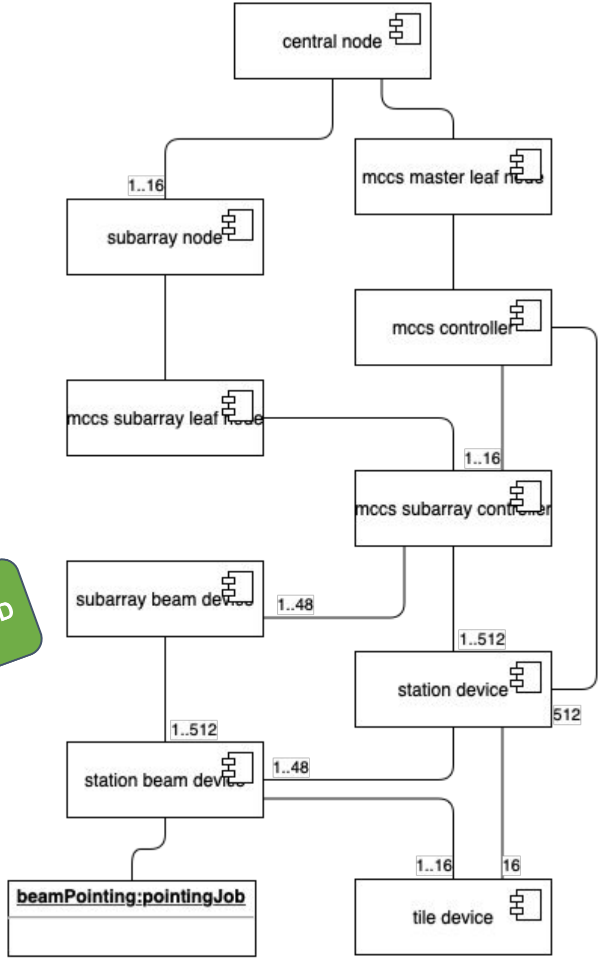

ADR-23 component view

Software and Firmware Development Standards

- SKAO Organisational Standard
  - Applicable to all software and firmware, including all embedded software and PLC systems, such as those described in the IEC 16113-3 standard.
- Providing developer guidance on the [SKA developer portal]()
- Providing software development infrastructure

SQUARE KILOMETRE ARRAY

FUNDAMENTAL SKA SOFTWARE & HARDWARE DESCRIPTION LANGUAGE STANDARDS

Document Number	SKA-TEL-SKO-0000661
Document Type	STD
Revision	03
Author	Marco Bartolini, Nick Rees
Date	2021-04-28
Document Classification	UNRESTRICTED
Status	Released

Name	Designation	Affiliation	Signature	
Authored by:				
Marco Bartolini	Lead Software Architect	SKAO	Marco Bartolini	
			Date:	Apr 28, 2021
Owned by:				
Marco Bartolini	Lead Software Architect	SKAO	Marco Bartolini	
			Date:	Apr 28, 2021

SKAO

latest

Search docs

GETTING STARTED

Onboarding: Welcome to the SKA developer community

Docs » SKA telescope developer portal

Edit on

SKA telescope developer portal

Welcome to the Square Kilometre Array software documentation portal. Whether you are a developer involved in SKA or you are simply one of our many users, all of our software processes and projects are documented or linked to in this portal.

The portal is frequently updated as the project evolves; if you feel that something is missing, please have a look at our guide to contributing to the developer portal

CI/CD Pipeline - GitLab and Kubernetes Centric

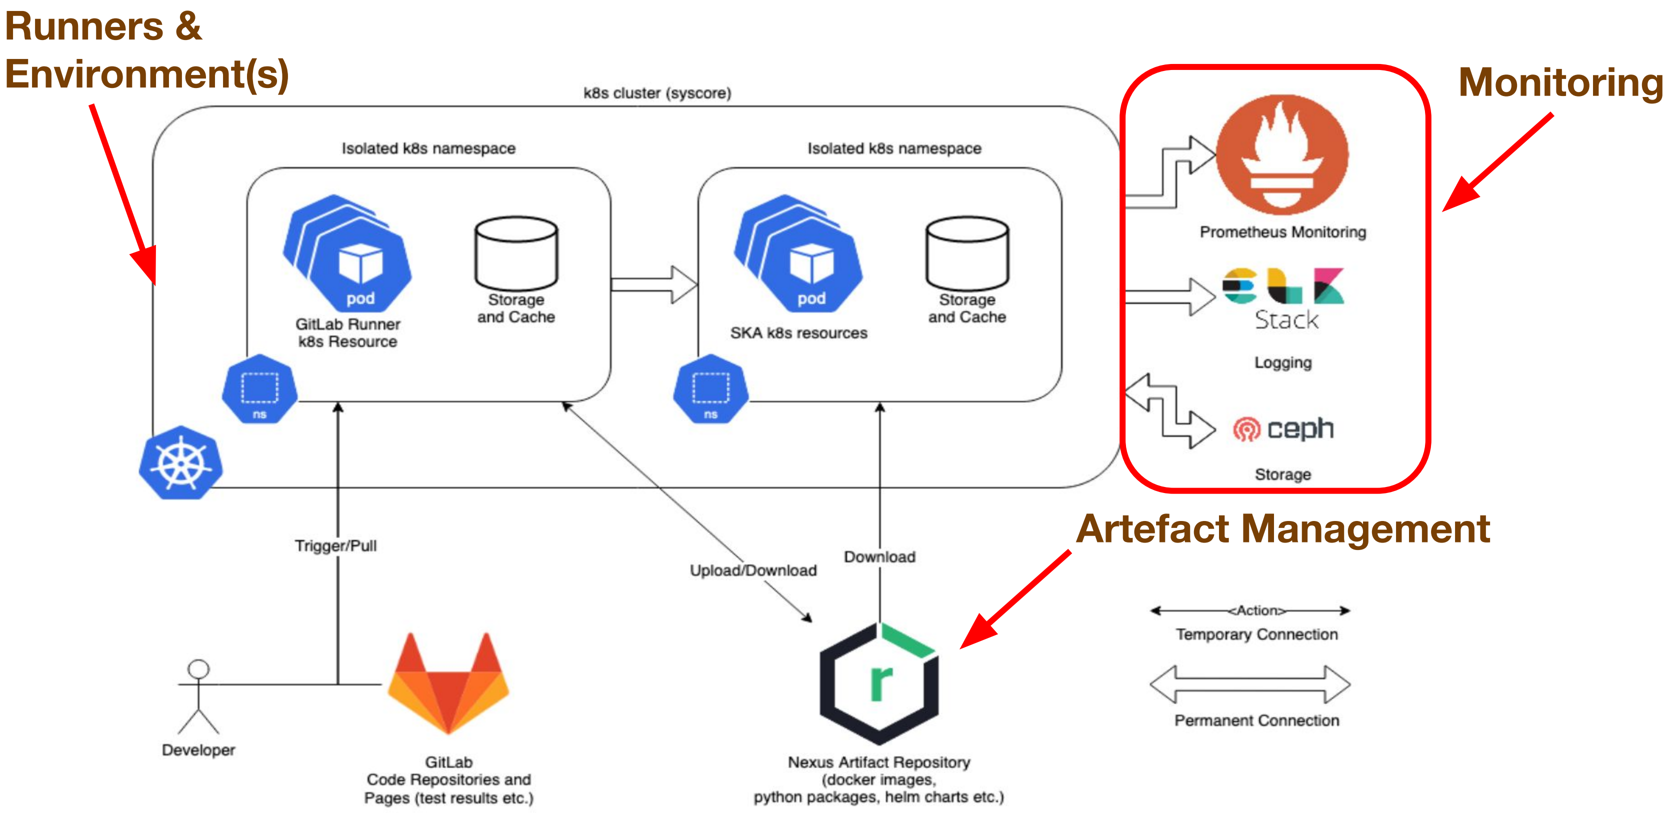

Platform - Infrastructure Management

- Managing Compute and Storage as a programmable service using Software Defined Infrastructure technologies
  - ADR-43 Determining the appropriate computing infrastructure management solution
  - Developing consistent management resources across all SPC facilities (and potentially ITF, PSI, CPF etc.)
  - Supported by a dedicated team - focused on AA0.5/1.0

ANALYZING

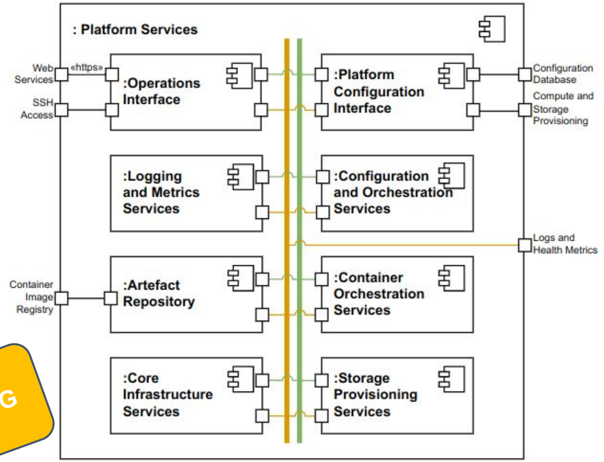

Platform C&C View

ANSIBLE Terraform openstack

The Target Platform - Cloud Native

Universal Application Abstraction

kubernetes

Universal Logging and Monitoring

Continuous Integration & Deployment

Abstract Network and Storage

Calico

Rook.io

Prometheus

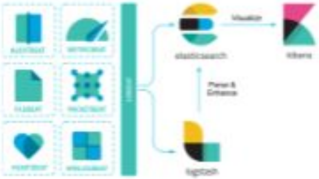

ElasticStack

GitLab CI

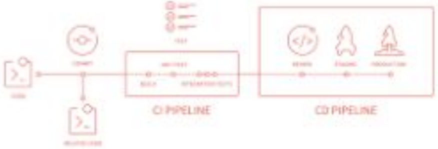

Universal Storage

ceph

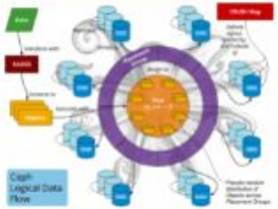

Cloud Infrastructure Layer

openstack.

RED HAT ANSIBLE Automation

System Dashboards and user interfaces

- Web based
- Integrated with the TANGO system

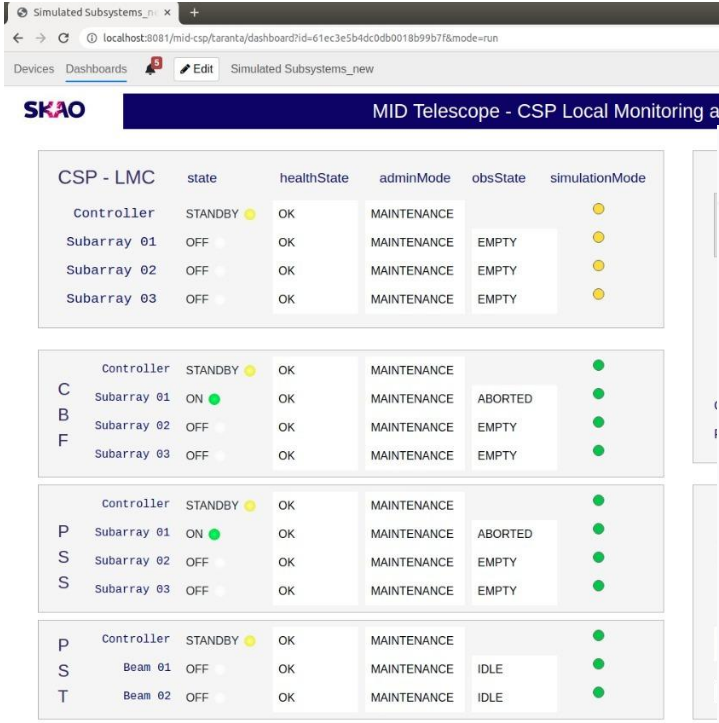

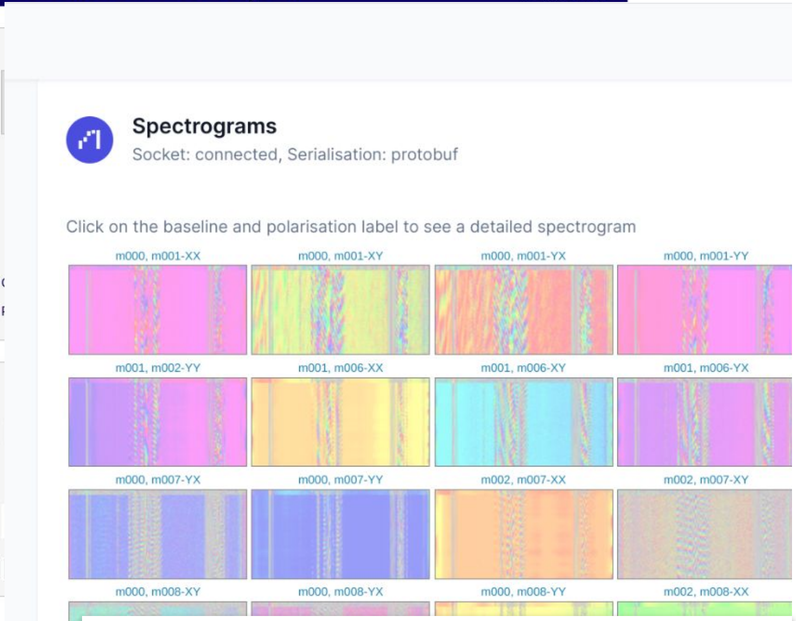

MeerKAT data being played through the SKA SDP QA displays

SKA MID AA0.5 Network

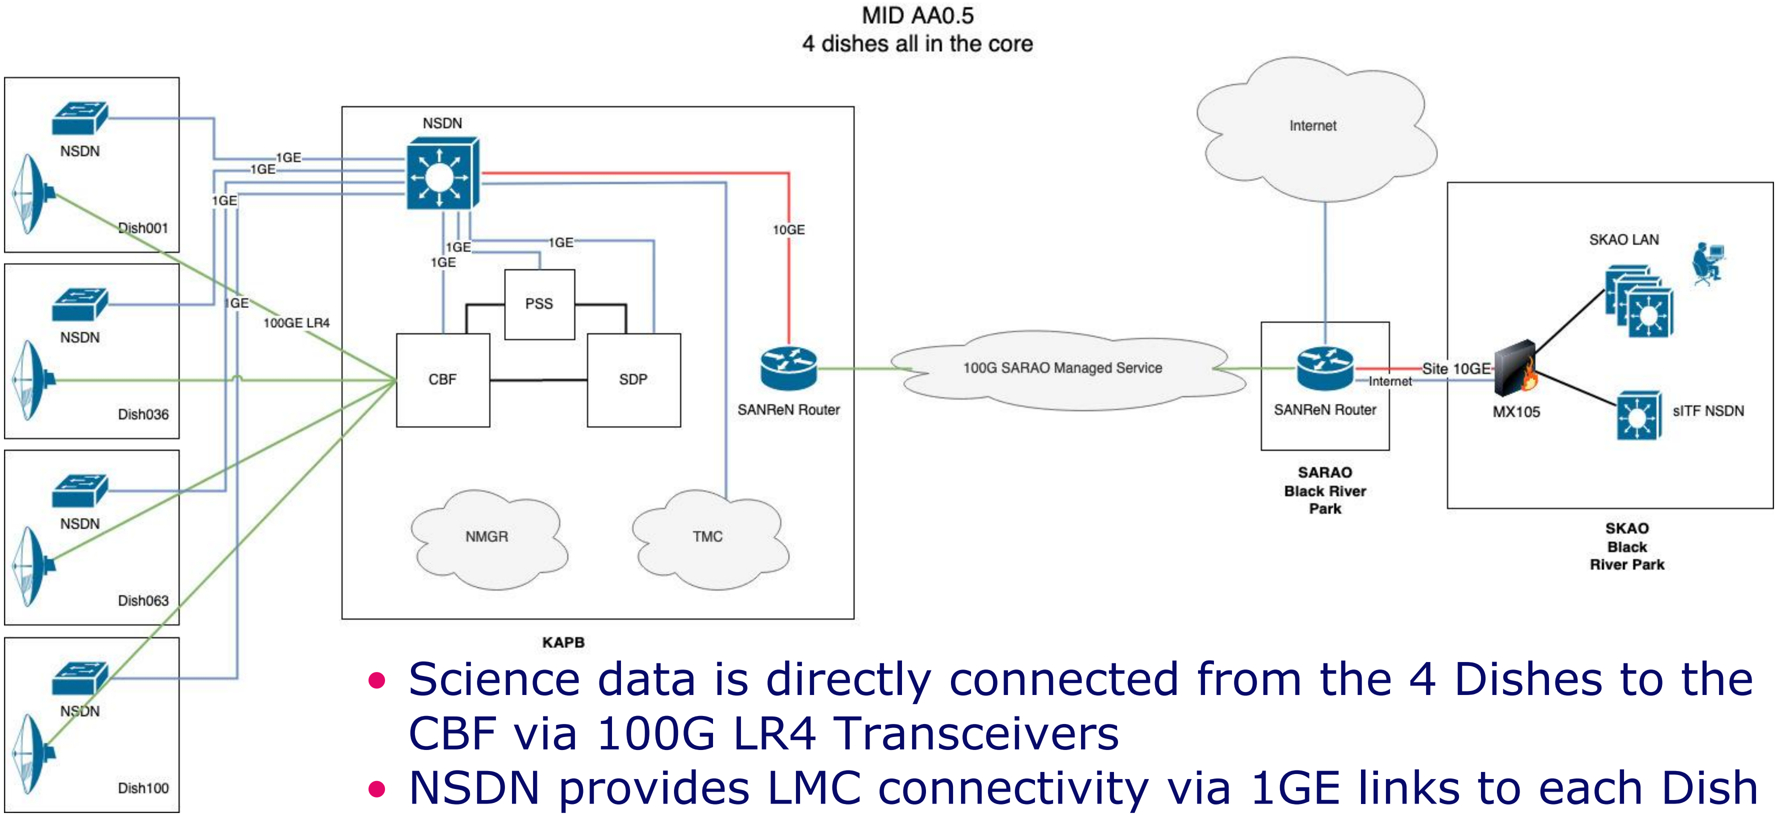

- Science data is directly connected from the 4 Dishes to the CBF via 100G LR4 Transceivers
- NSDN provides LMC connectivity via 1GE links to each Dish
- Access from site to Cape Town via SARAO managed 100G link

SKA MID AA0.5 in Netterrain

- AA0.5 will be represented in Netterrain
- Real equipment and fibre routing
- Dish and fibre routing from the SPFRx to the CBF shown here.

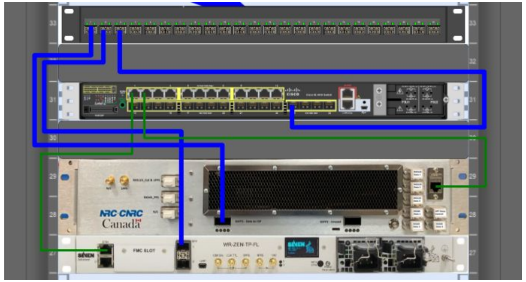

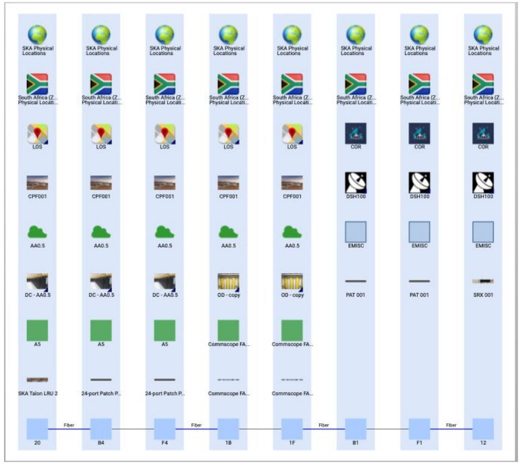

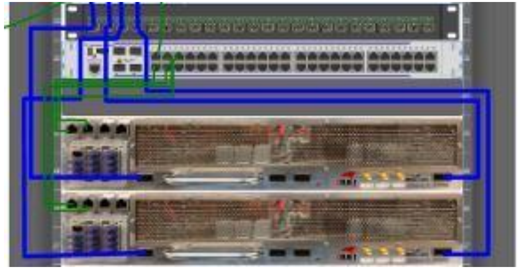

https://netterrain-development.skao.int/Diagram/24000000003630
https://netterrain-development.skao.int/Diagram/24000000033831

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

SKAO

www.skao.int

For reference: Robert's slide on AA0.5

- Deployment of minimal arrays on-site as early as possible
- Primary goal: end-to-end test of interferometry (and beam-forming)
- (Almost) **all** sub-systems (including initial control and data processing software)
    - Includes Dish/Station (not tested in ITF)
- Verify fundamentals of system performance
    - realistic operating environment (e.g. RFI, wind, temperature, ...)
- Interfaces
- Develop AIV, Commissioning teams and procedures
- Identify failures to meet requirements, lack of reliability
- Reduce risk by fixing problems as soon as possible, ideally before mass production
- Verify the supply chain