SKA SDP Critical Design Review

Architecture, CDR Outcome and Australia contribution during Bridging

J.C. Guzman and the SKA SDP Architecture Team

C4SKA 2019 – 14 – 15 Feb 2019, Auckland NZ

CSIRO ASTRONOMY AND SPACE SCIENCE
www.csiro.au

The Road to the Critical Design Review (CDR)

- Delta-PDR – May 2016
- Major project structure during 2017
- Software Engineering Institute (SEI) Architectural Approach
  - Views and Beyond
- Pre-CDR – June 2018
  - SEI's Architecture Trade-off Analysis Method (ATAM) to review the SDP Design (used by Telescope Manager as well) – 4 days
  - Partial Architecture presented and reviewed – Science pipeline workflows and Data Models missing
- CDR Submission – Oct 2018
- Revised SDP Architecture presented to SKAO in Nov 2018
- CDR – 15 – 18 Jan 2019

Architecture Context

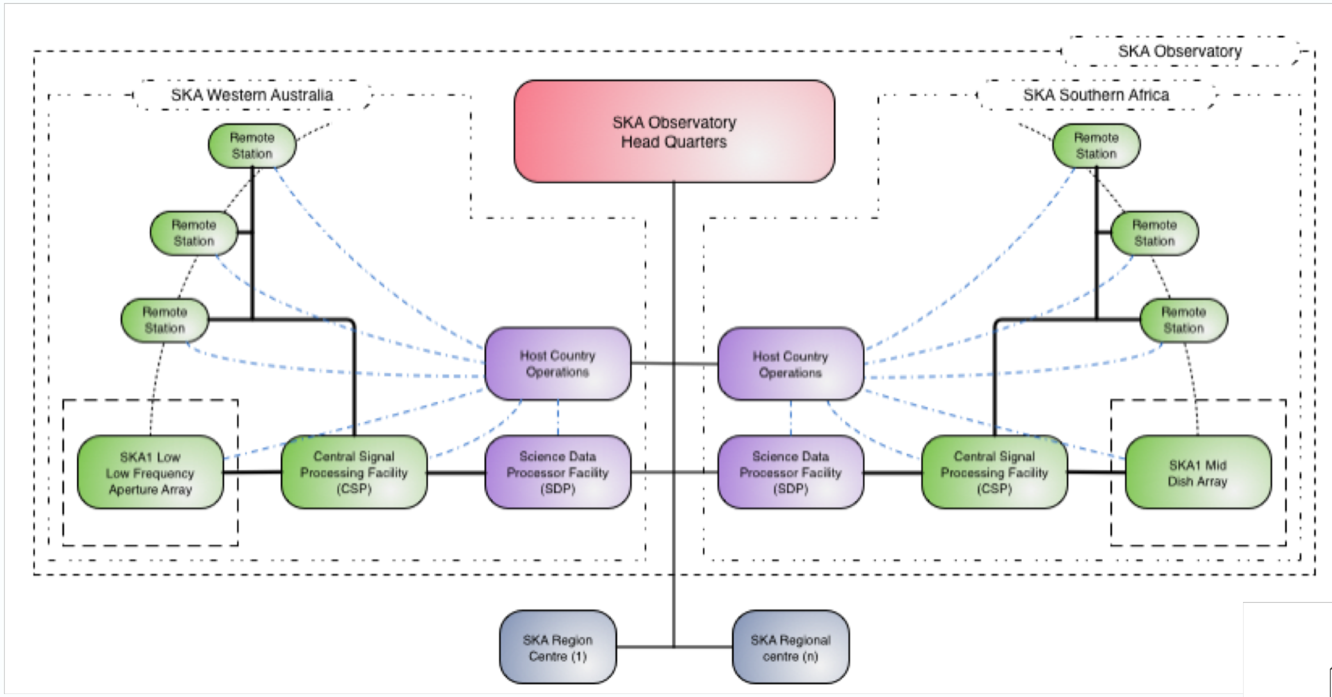

Single SDP Deployment

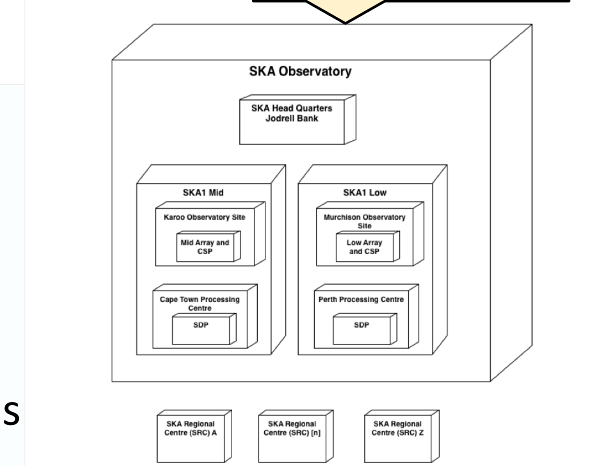

Principal functions of SDP sub-system:

1. Controlled by SKA observatory control
2. Ingest measurements from CSP/LFAA
3. Perform a variety of calibration functions
4. Perform batch processing to build Data Products
5. Deliver Data Products to Observatory & Regional Centres

Slide Credit: SDP Architecture Team

Architecture Context: Data Handled

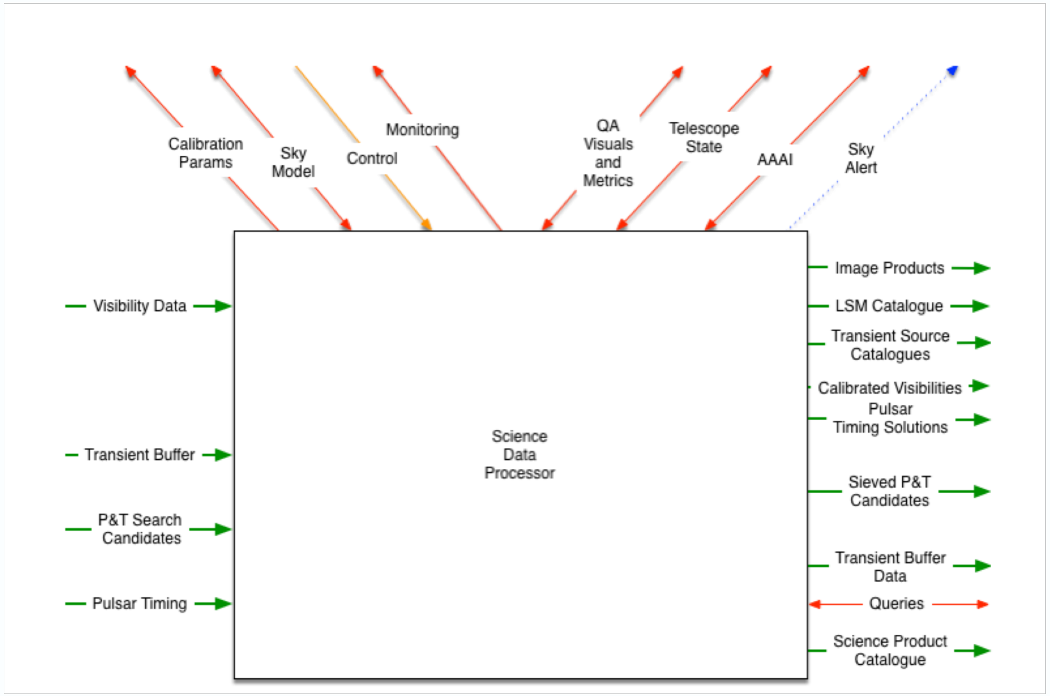

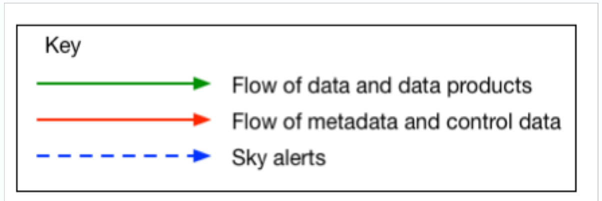

Types of data consumed and produced by SDP:

Control: Control & Monitoring, Sky & Telescope information, Quality Assessment, Science Event Alerts

Ingest: Visibilities, Pulsar Search + Timing and Transient B.s

Delivery: Data Products (many types), Catalogue & Query Interface

SDP Architecture – SEI's Views and Beyond

New views and view packets shown in bold

High Level Architecture Documents

- **SKA1 SDP High Level Overview**
- SKA1 SDP Architecture Reading guide

Data Model Views

- System Data Model View
  - **Execution Control**
  - **Processing**
    - **Visibilities**
    - **Gridded Data**
    - **Image**
    - **Calibration**
    - **Sky Model**
    - **Transient Source Catalogue**
    - **Non-Imaging**
  - **Science Data Model**
  - **Science Data Product Catalogue**

Workflow Views

- **Science Pipeline Workflow View**
  - **Workflow scripts C&C / deployment**
  - **Receive**
  - **Pre-processing**
  - **Real-time Calibration**
  - **Fast Imaging**
  - **ICAL**
  - **Instrumental Calibration**
  - **Model Partition Calibration**
  - **Imaging**
  - **Deconvolution**
  - **Pulsar Search & Single Pulse**
  - **Pulsar Timing**

Module Views

- System Module View
  - **Execution Control**
  - **Science Pipeline Workflows**
  - Processing Components
  - **Delivery**
  - **Platform**

Component & Connector Views

- Operational System C&C View
  - **Execution Control**
  - Processing
    - **Dask Execution Engine**
    - **DALiuGE Execution Engine**
    - **MPI Execution Engine**
  - **Buffer and Long-Term Storage**
    - **Buffer Data Lifecycle**
  - **Model Databases**
  - Delivery
- Platform C&C View
- **Software Management C&C View**

Use Case Views

- Science Pipeline Management Use Case View

Other Views

- Functional View
- Hardware Decomposition View
- Security View

Slide Credit: SDP Architecture Team

SDP Architecture: Driving Attributes

Main tension within SDP architecture:

Performance & Scalability ⟺ **Modifiability & Maintainability**

Performance & Scalability

- Compute, I/O & Storage
  - >10 Pflop/s effective
  - ~0.4 TB/s ingest rate
  - ~4 TB/s into processing
  - >40 PB tiered buffer
- Need to scale
  - Trivial and expensive workflows co-exist
  - SKA “>1” will be even harder on SDP

Modifiability & Maintainability

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

Slide Credit: SDP Architecture Team

SDP Architecture: Driving Attributes (2)

Further attributes driving the architecture:

Availability

- Support observations
- Provide Data Products

Buildability, Affordability

- COTS components
- Support agile development

Reliability

- Store measurements & especially Data Products
- Perform processing steps

Testability

- Isolated processing and services

Portability

- 2 observatory deployments
- SRC deployments

Usability

- Workflow modification
- Quality assessment
- Data Product catalogue

SDP Architecture: Scalability & Performance

Scalability & Performance:

- Implement raw performance requirements
- Keep pace with hardware and workflow evolution

Strategies:

1. **Data Driven Software Architecture**
    1. **Scalable stores**
       Buffer, Data Queues, Configuration
    2. **Workflows**
       Tailoring of data distribution to pipeline
2. **Hardware Platform**
   Provides infrastructure enabling scalability

Slide Credit: SDP Architecture Team

SDP Architecture: Scalability & Performance

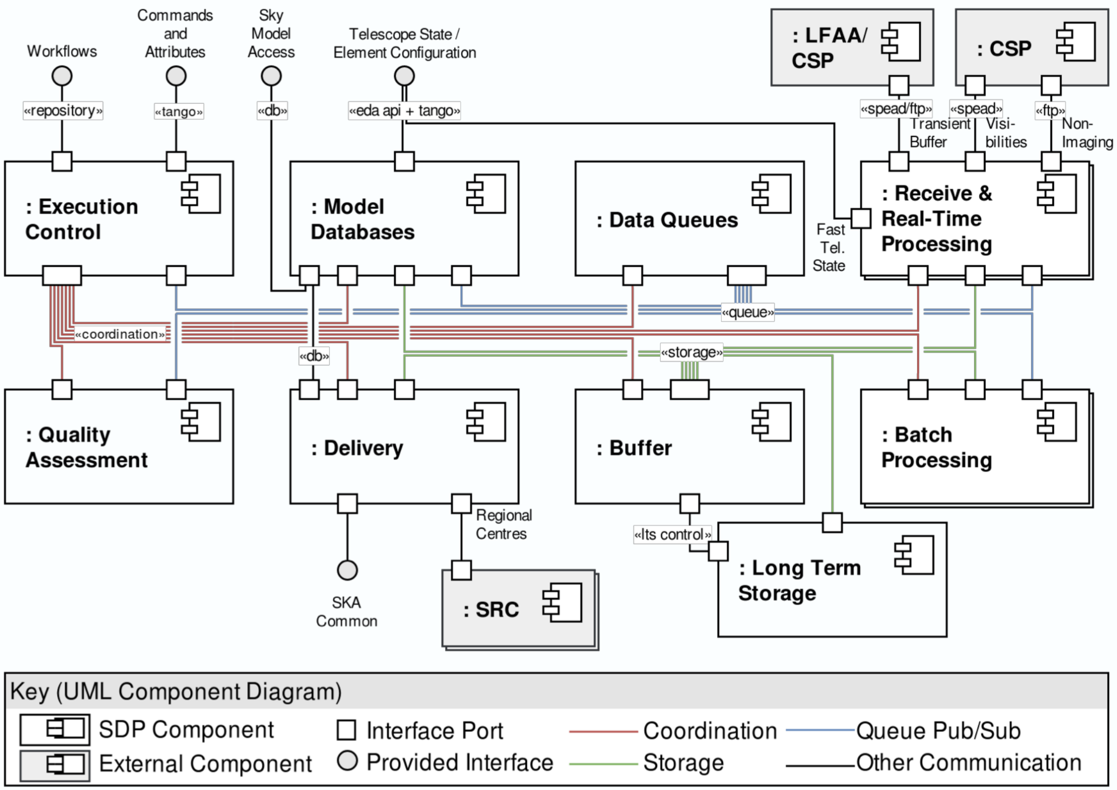

Shared Data & Publish/Subscribe:

- **Buffer (Storage)**
  high throughput
  high capacity
  requires locality!
- **Data Queues**
  low latency
  good throughput
  scales globally
- **Coordination**
  high reliability
  low latency (read)
  limited access

Slide Credit: SDP Architecture Team

SDP Architecture: Scalability & Performance

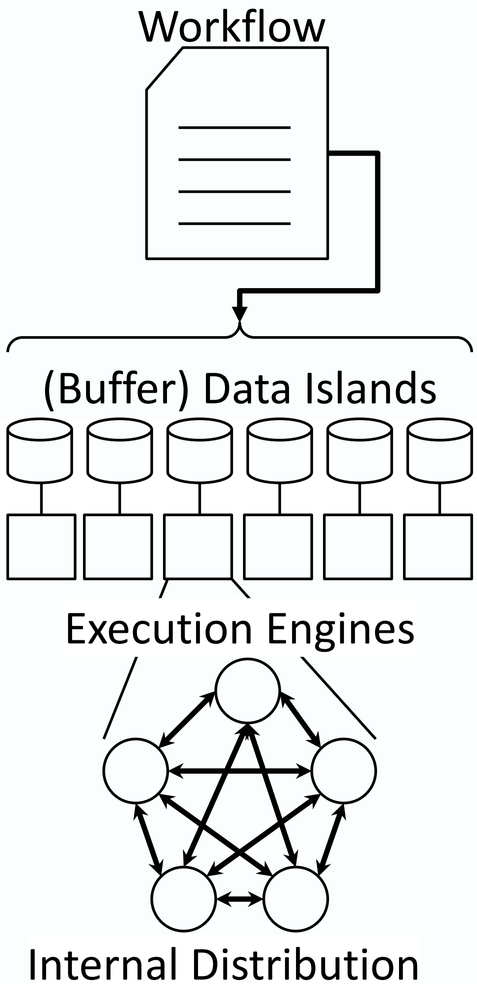

Workflow Processing Stages
use two types of data distribution:

1. **"Embarrassing" distribution**
   (e.g. frequency, time axis):

   Data Islands + Execution Engines offer strong performance isolation

2. **"Complex" distribution**
   (e.g. spatial axis):

   Internally Execution Engines tailor approach, e.g. all-to-all communication

Number + scale of data islands changes according to workflow!

Slide Credit: SDP Architecture Team

SDP Architecture: Scalability & Performance

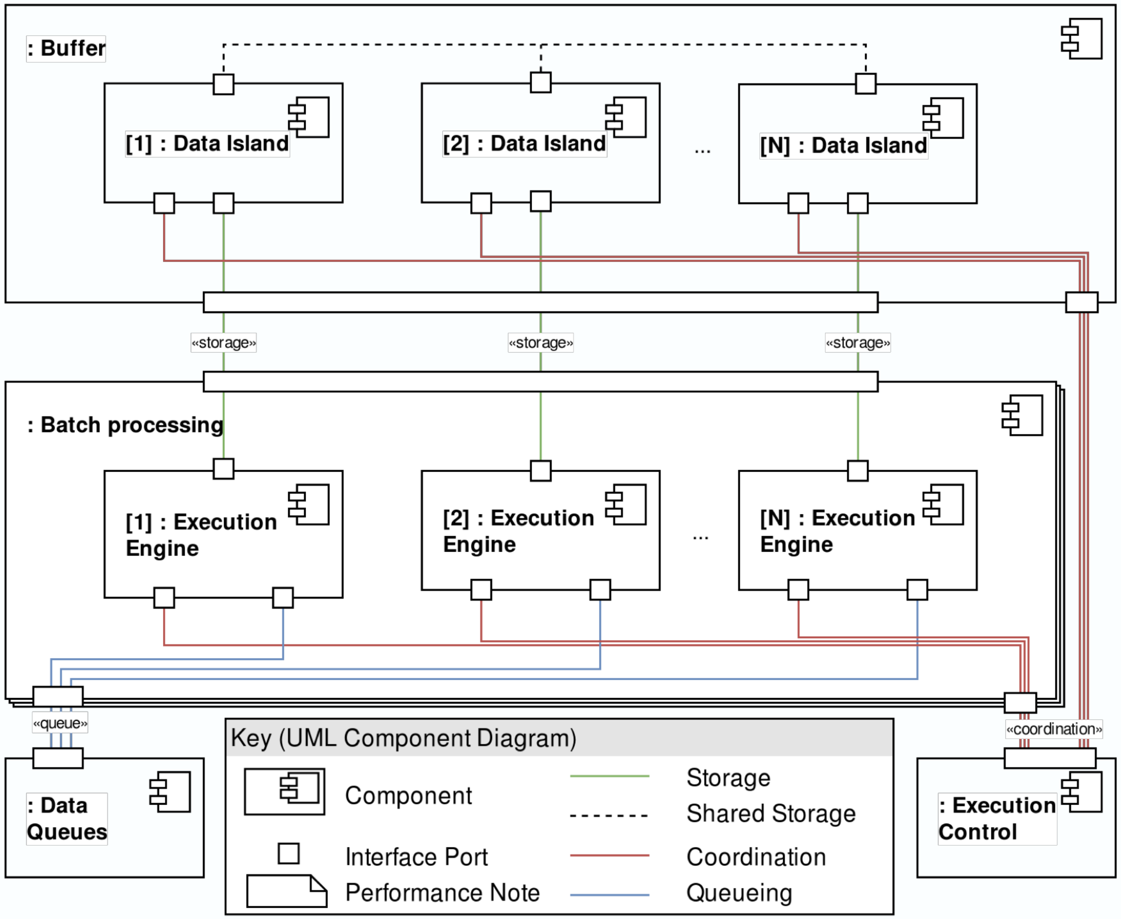

Data Islands ensure storage locality:

- Tailored towards Processing needs
- Dedicated file metadata services
- Can share across islands if needed

Execution Engines

- isolated from environment
- Data Queues low-rate pub/sub (e.g. calibration, QA)

Slide Credit: SDP Architecture Team

SDP Architecture: Scalability & Performance

Hardware & Platform scaling concepts:

Full nonblocking network infrastructure
⇒ Required for flexibility in Data Island structure

Heterogeneous server resource pool
⇒ Support "flavours"
⇒ Take advantage of future hardware

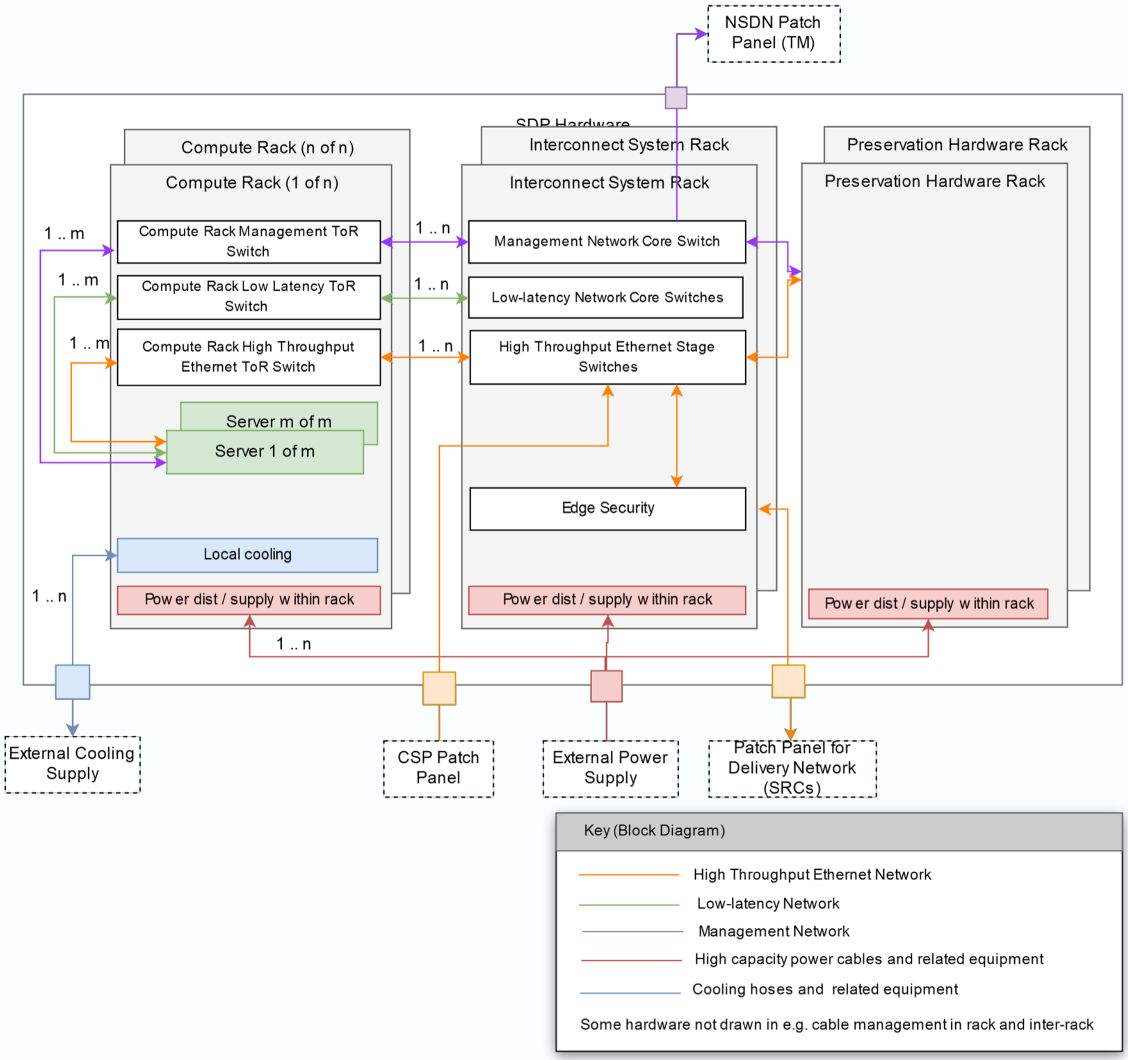

Slide Credit: SDP Architecture Team

SDP Architecture: Maintainability & Modifiability

Maintainability & Modifiability:

- Plan for long lifetime
- Allow software + hardware changes

Strategies:

1. **Loose Coupling**
    1. **Services** vs **Processing**
    2. **Workflows / Execution Frameworks** from each other
    3. **Processing Components** from each other
2. **Modification as Use Case**

SDP Architecture: Maintainability & Modifiability

Modules structured into "pillars":

- **Service Pillar** handles control, inputs and outputs
- **Processing Pillar** for core scientific functions

Workflow (Libraries) main connection point:

- Interface to Execution Control
- Coordinates Services

Communication using Data Models

⇒ Strong decoupling at top level

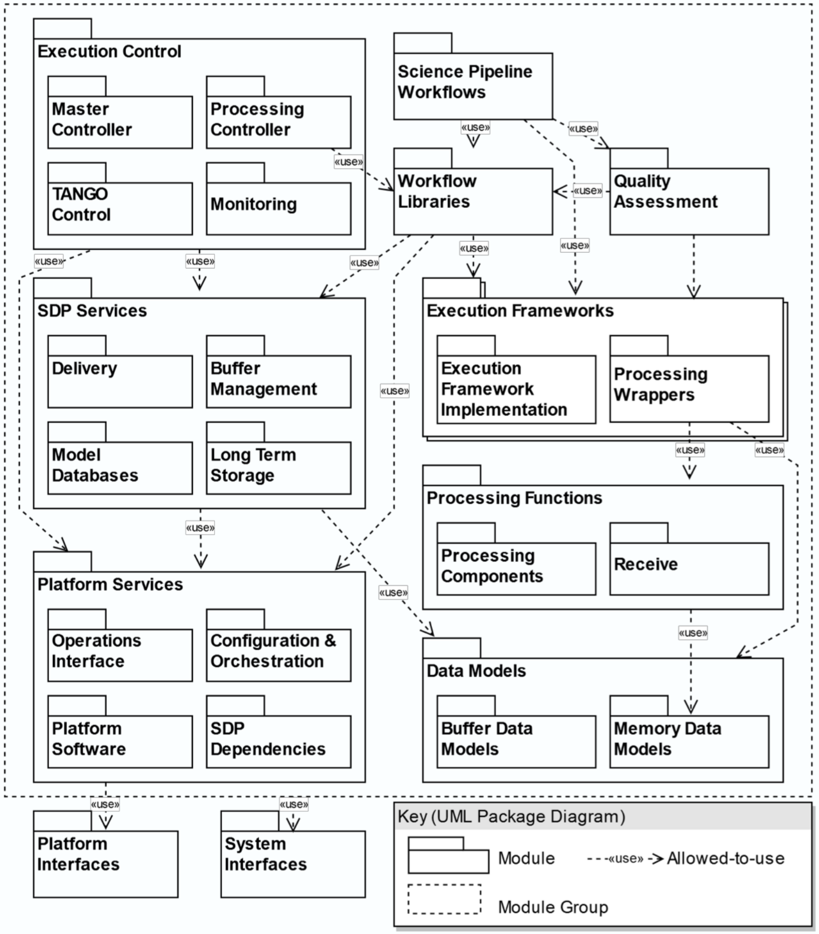

Slide Credit: SDP Architecture Team

SDP Architecture: Maintainability & Modifiability

Processing Layers

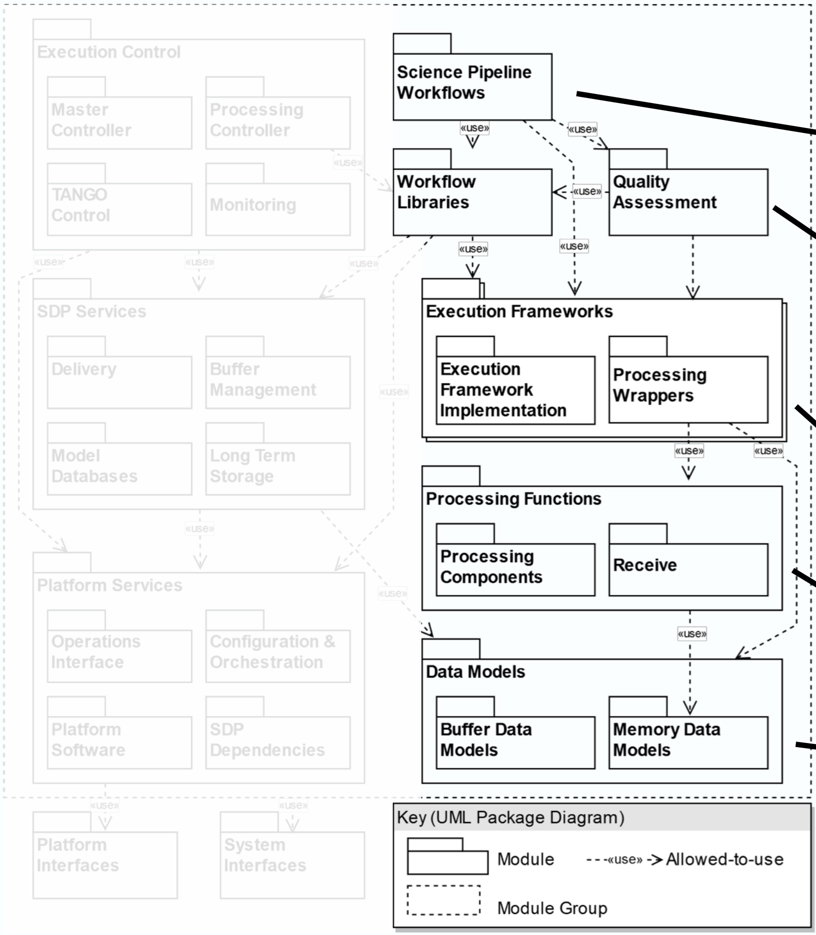

Flexible processing architecture to serve varied scientific needs:

- **Workflows** implement high-level **pipelines**
- **Workflow Libraries + Quality Assessment** capture workflow patterns
- **Execution Engines** handle I/O and distribution
- **Processing Components** domain-specific kernels
- **Data Models** describe shared data representations

Typical pipeline architecture, adapted for scalability

CDR: ATAM Process

- 15 - 18 Jan 2019
  - ATAM “Round 2” – 2 days followed by OAR discussion the next 2 days
  - Scenarios updated (or new) and prioritized

	A	E	F	G	H	I	K
1	**#ID**	**Stimulus** *A condition that requires a response when it arrives at the system*	**Environment** *The stimulus occurs under certain conditions. The system may be in an overload condition or in normal operation or some other relevant state*	**Response** *How the system should respond to the stimulus. The response consists of the responsibilities that the system or the developers should perform in response to the stimulus*	**Importance (H,M,L)**	**Difficulty (H,M,L)**	**Review Order (1-10)**
2	9	TM issues request to observe an imaging SB including reference pointing and SDP responds that the resources are available. Once the visibilities are processed by SDP into Observatory Data Products, they are archived into Long Term Storage and delivered to an SKA Regional Centre.	Normal Operations	The observing commences and the data are ingested, flagged, averaged, direction independent calibration ensues followed by a direction dependent-calibration imaging pipeline. The provenance is recorded and updated for all steps of the processing and the quality assessment is calculated and recorded throughout. Finally the calibrated Observatory Data Products are staged for delivery to SRCs, entered into the Data Product Catalogue, and archived into Long Term Storage.	H	L	1
3	4	A project requires the processing of visibilities using the SDP instance on the SRC platform, using new algorithms developed by the project team. This may include the processing of new and old visibilities into new products, and perhaps from different telescopes (e.g. single dish).	Commissioning and Operations	The SDP at the SRCs support generation of different instances allowing new algorithms and workflows to be combined with existing capabilities. The operations capabilities and staff must be able to support the entire workflow in collaboration with the SRC, while still permitting other large-scale observing programs.	H	H	2
	6	The parameters for an imaging	Commissioning and/or	In each case, changes to the PB should	H	H	3

CDR Outcome

- ATAM process covered mainly operational, interfaces (M&C) and non-domain part of the system
- OAR discussion covered the algorithmic/science domain of the system
- Overall the panel an outcome was very positive (PASS) for SDP
- Panel commended the depth level of the Construction planning and team work
- Some issues will be inevitably carried forward to Bridging
- Finding and Recommendations:

Architecture	SKA Common
SDP – SRC Interfaces	Staffing Ramp-Up
End-To-End Testing	Reliability, Availability and Maintainability
Algorithm Development	Security
Commissioning and Early Operations	Adoption of SAFe

CDR Outcome

- ATAM process covered mainly operational, interfaces (M&C) and non-domain part of the system
- OAR discussion covered the algorithmic/science domain of the system
- Overall the panel an outcome was very positive (PASS) for SDP
- Panel commended the depth level of the Construction planning and team work
- Some issues will be inevitably carried forward to Bridging
- Finding and Recommendations:

Architecture	SKA Common
SDP – SRC Interfaces	Staffing Ramp-Up
End-To-End Testing	Reliability, Availability and Maintainability
Algorithm Development	Security
Commissioning and Early Operations	Adoption of SAFe

CDR Outcome: Architecture

- SDP has a complex conceptual architecture, in comparison to existing radio astronomy processing and archiving systems
  - Appropriate to the scope and scale of SKA SDP.
  - Contains features to address the functional, performance, and growth
  - Has been exercised analytically, and by prototyping in key areas.
- The panel endorse the planned development of an initial simple system (MVP) and the realization of the full architecture incrementally through the SAFe process
- Further work in the abstract/conceptual level is unlikely to add value before construction
- Only by down-selecting technologies, implementing SDP, and testing it, will the architecture be truly validated.

Recommendation: Begin concrete construction activities as soon as possible.

When: Bridging and beyond

CDR Outcome: SDP – SRC Interfaces

- The lack of a formal SDP – SRC ICD prevents a comprehensive elaboration of the DELIV component of SDP.
- Development of the SRC concept is converging slowly

Recommendation: Develop a “reference” SDP – SRC interface document as soon as possible as a design and planning vehicle, and maintain it until it can be baselined.

When: Bridging

CDR Outcome: End-To-End Testing

- Milestones sufficient to manage the SDP development, integration, and commissioning contained in the Construction Plan.
- It will be informative to conduct an end-to-end processing of a representative volume of pre-cursor or simulated data.
- Objective is to enable assessment of the scientific quality of the algorithm output, and the computational performance of the algorithms.
- May expose need/desirability for changes to non-domain components as well.

Recommendation: The SDP team should incorporate at least one or two “end-to-end” Test Milestones into the plan to test algorithm scalability and scientific quality.

When: Before SDP Closeout (?)

CDR Outcome: Algorithm Development

- Refinement of domain algorithms for a specific telescope or observatory requires very specialized skills.
- Some gaps in planning and missing resources for:
  - advancing the technical readiness level of algorithms
  - tailoring of these algorithms for application for the SKA telescopes.
- Larger research community will continue to develop and evolve algorithms
  - May not address the specific challenges of the SKA.

Recommendation:

- Add staff to the project plan dedicated to the advancement and tailoring of algorithms for the SKA.
- Develop an implementation plan for the science workflows to satisfy Commissioning, AIV, and end-to-end testing.

When: System CDR

CDR Outcome: Commissioning and Early Operations

- Design provides an efficient and automated system. ○ Suitable for routine operations.
  - Automation is appropriate and necessary for steady state operations
- May be counter-productive during commissioning and early science activities.

Recommendation: Ensure that the system supports non-standard use cases through suppression of undesired automated behaviors. *Identify use cases (tools required) for SDP during commissioning and early operations – Pre-cursor Experience is vital!!*

When: Bridging

CDR Outcome: Others

SKA Common

- Concept is overloaded and needs to be clarified (Needs work from All Software Elements)

Staffing Ramp Up

- Incorporate a Ramp Up period in the Construction Schedule from 0 to ~42 people

RAM

- Document/analysis of failure modes in some key areas without full FMECA

Security

- The SKAO should define the System level Security program and requirements, based on existing formal standards and processes.

Adoption of SAFe

- Be careful to routinely evaluate the effectiveness of SAFe in addressing issues of varying types.

SDP Bridging Activities in Australia

ICRAR/CSIRO Joint Project, codename “Rialto 2” – 2 years, funded by the Federal Govt.

Project goals are:

- Contribution to the SKA CDR submission and close
- Evolution of the newly developed integration prototype into a production quality software product capable of processing SKA1 and precursor science data. This activity aligns with the proposed Bridging statement of work Science Data Handling & Processing
- Extensions/improvements to DAliuGE and ASKAPsoft

Adoption of SAFe (Essential SAFe)

- 1 Agile Release Train, 2 Teams, ~15 people (11 ICRAR, 4 CSIRO)

International Centre for Radio Astronomy Research

Summary

- CDR Successful – Lots of good work during pre-construction and many organizational challenges & changes
- Time to develop software and re-use as much as possible
- Developed SDP capability in Australia and New Zealand
  - Cover key aspects of the SDP system during bridging
- Opportunity to work closer together during Bridging
- Alignment with SKAO SAFe Trains

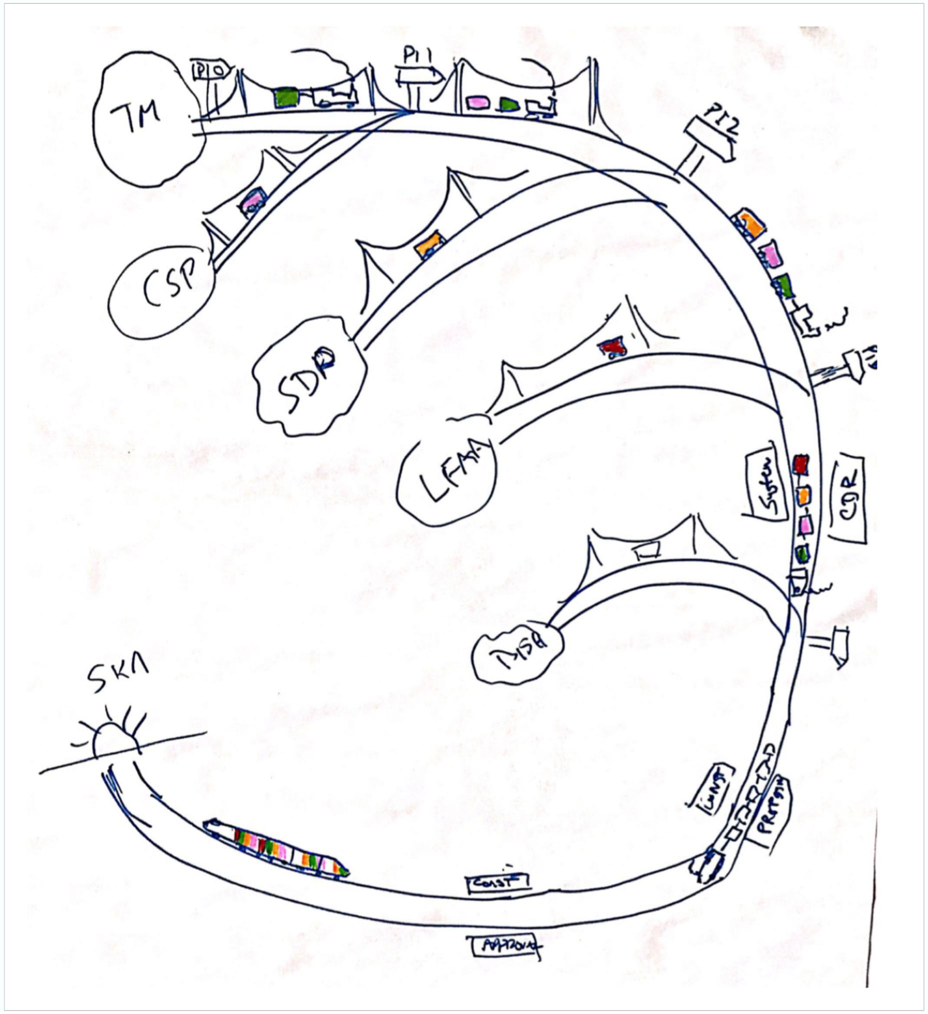

Credit Image: SKAO

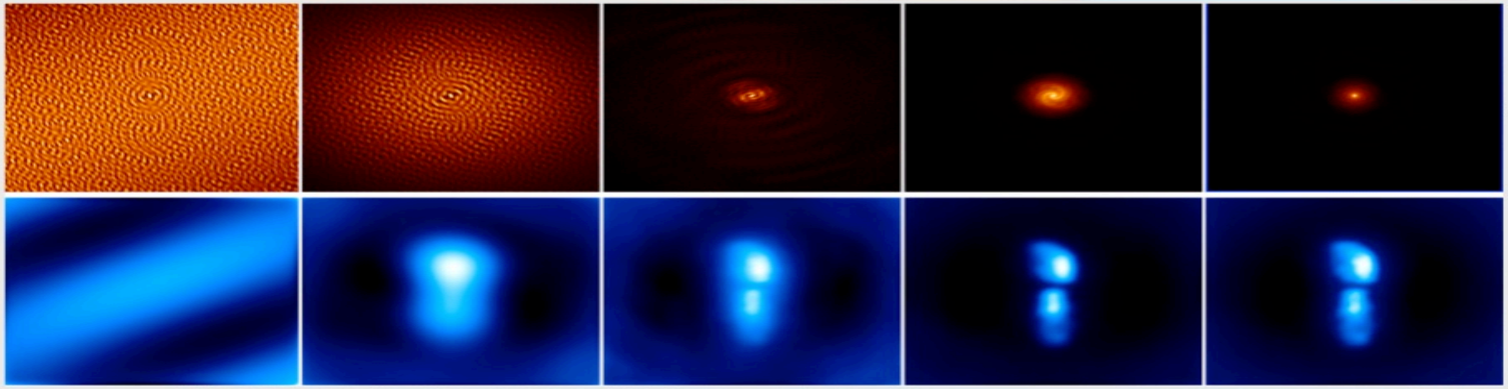

Thank You!

CSIRO Astronomy and Space Science (CASS)
Juan Carlos Guzman
Head of Software and Computing

t +61 8 6436 8569
e juan.guzman@csiro.au
w www.atnf.csiro.au

SKA SQUARE KILOMETRE ARRAY SDP SCIENCE DATA PROCESSOR

ICRAR International Centre for Radio Astronomy Research

CSIRO