SDP Architectural Overview

Document number.. SKA-TEL-SDP-0000013
Document Type.. REP
Revision... 05
Author.. P. Alexander, B. Nikolic, V. Allan
Release Date.. 2018-04-25
Document Classification.. Unrestricted
Status.. Draft

Lead Author	Designation	Affiliation
Paul Alexander	SDP Project Lead	University of Cambridge
Signature & Date:	Paul Alexander Paul Alexander (Apr 23, 2018)	

Released by	Designation	Affiliation
Paul Alexander	SDP Project Lead	University of Cambridge
Signature & Date:	Paul Alexander Paul Alexander (Apr 23, 2018)	

Table of Contents

List of Abbreviations

AAAI	Authorization, Access, Authentication and Identification
ALMA	Atacama Large Millimeter Array
C&C	Component & Connector
CASA	Common Astronomy Software Applications package
CDR	Critical Design Review
HPC	High Performance Computing
ILS	Integrated Logistics Support
JVLA	Jansky Very Large Array
L2	Level 2
NREN	National Research and Education Network
P&T	
PBS	Product Breakdown Structure
PSS	PulSar Search
PST	PulSar Timing
QA	Quality Assessment
RAM	Reliability, Availability, Maintainability
SDP	Science Data Processor
SEI	Software Engineering Institute
SKA	Square Kilometre Array
SKAO	Square Kilometre Array Organisation
SRC	SKA Regional Centre
TM	Telescope Manager
Vis	Visibility
WAN	Wide Area Network

1. The SKA-SDP System Overview

This architectural documentation concerns the Science Data Processor (SDP) of the SKA. This section explains the scope of the SDP and its relationship to the overall SKA system.

The SDP is an element of the overall SKA system (see Figure 1) [AD01]. The SDP is the element of the telescope responsible for the processing of various observed data into the required data products, the long term preservation of these data products, and the delivery of these products to the SKA Observatory, and delivery to SKA Regional Centres (SRCs).

The SKA-SDP will function in a way unlike conventional HPC or Cloud systems:

- The SDP is under the overall control of the SKA control and monitoring system and some SDP activities are scheduled in the same way as the rest of the telescope. For example, some data acquisition and processing will be triggered by transient astronomical events; and some maintenance activities will need to be aligned with the maintenance of the rest of the telescope.
- data must be ingested from the telescope at a high rate and some initial processing performed to feedback information to the telescope. Both are time-critical activities.
- The ingested data must be further processed in a batch-oriented mode so that the throughput of the SDP keeps up with the overall telescope operations.

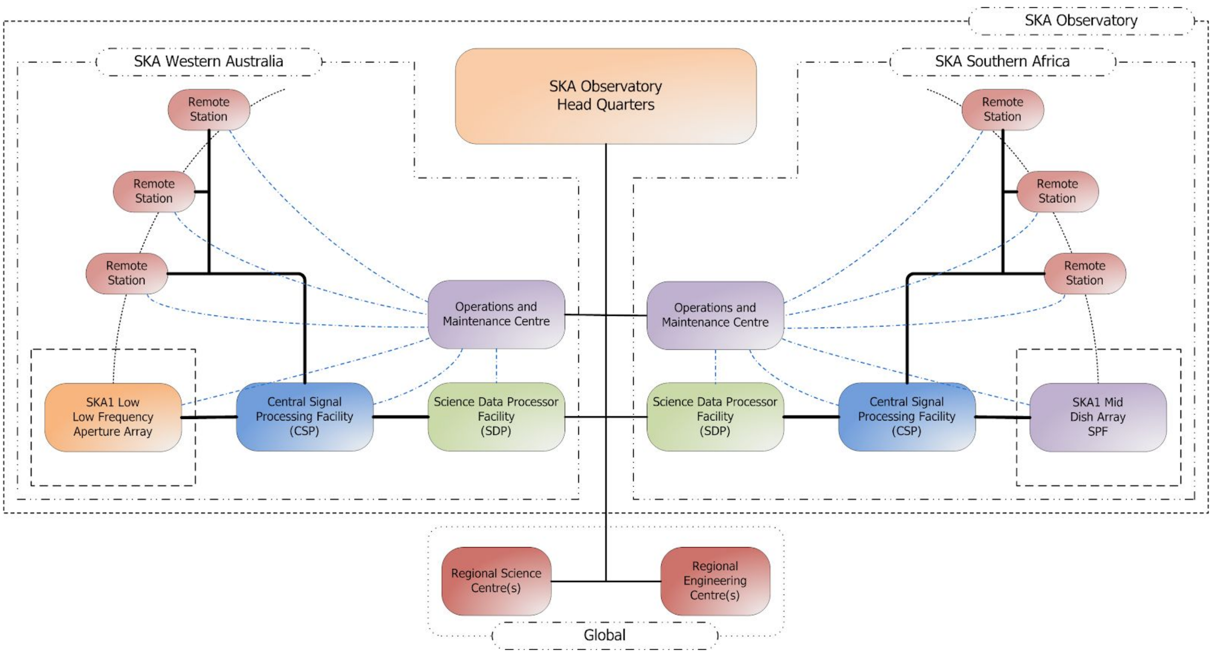

Figure 1: *Overview of the SKA system showing the relationship of the SDP to the rest of the SKA system.*

The SKA Observatory is distributed across three physical sites (see Figure 2). The headquarters are located in the UK and the two observatory sites in South Africa (SKA1-Mid) and Australia (SKA1-Low). The SDP will be deployed to each of the observatory sites: there will be separate hardware deployments at the SDP data centres in Cape Town and Perth; however, the SDP *software* is a common deployment to both of these platforms.

The science data products may be queried by duly authorised users at each SDP site with further processing and analysis being performed at SKA Regional Centres (SRCs).

The SDP is also responsible for [AD02, AD03]:

- Computing calibration information required by the SKA system including time-critical calibration solutions;
- Alert generation;
- Providing additional metadata to describe the provenance of data;
- Quality Assessment information for evaluating the efficacy and scientific quality of the processing.

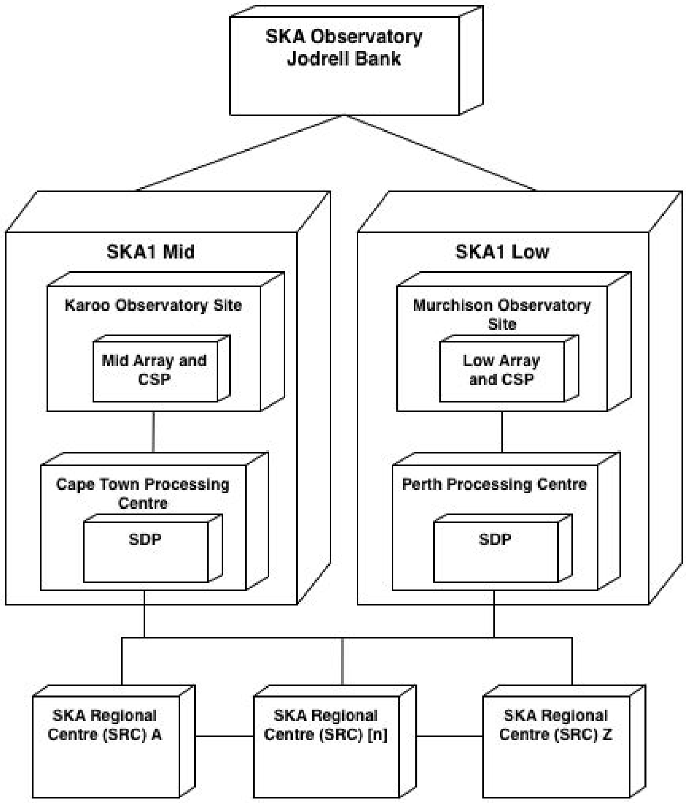

Figure 2: *Deployment of the SDP within the SKA Observatory. The SKA Observatory has its headquarters in the UK at Jodrell Bank and two sites in South Africa and Western Australia. There will be two physical platform deployments to processing centres in Cape Town (South Africa) and Perth (Australia). The interconnects between the physical deployments are shown by solid lines. The SDP software will be common across the observatory and will be deployed to each of the SKA processing centres associated with SKA1 Mid and SKA1 Low. The SKA Regional Centres (SRCs) form a group of data and science support centres: SKA data products may be moved between them.*

The SDP ingests observed (raw) data from the correlator, pulsar-search candidates or pulsar timing solutions. On command, the SDP is also required to ingest raw "voltage data" associated with each collecting element. The SDP interfaces to the SKA control system and time-critical processing is directly scheduled by the SKA observatory. Production of non-time-critical data products is performed in a batch-oriented processing mode: the overall science scheduling of the telescopes is linked to the available compute and data-storage resource of the SDP (determined by a system sizing model and cost constraints) so that the overall throughput of the processing does not result in the telescope being unable to observe.

The SKA has adopted a Tiered model for data delivery with SKA Regional Centres (SRCs) playing a formal role of accepting/requesting SKA-SDP data products and making these available to astronomers together with processing resources and support. The SKA Regional Centres will enforce the SKA data access policies.

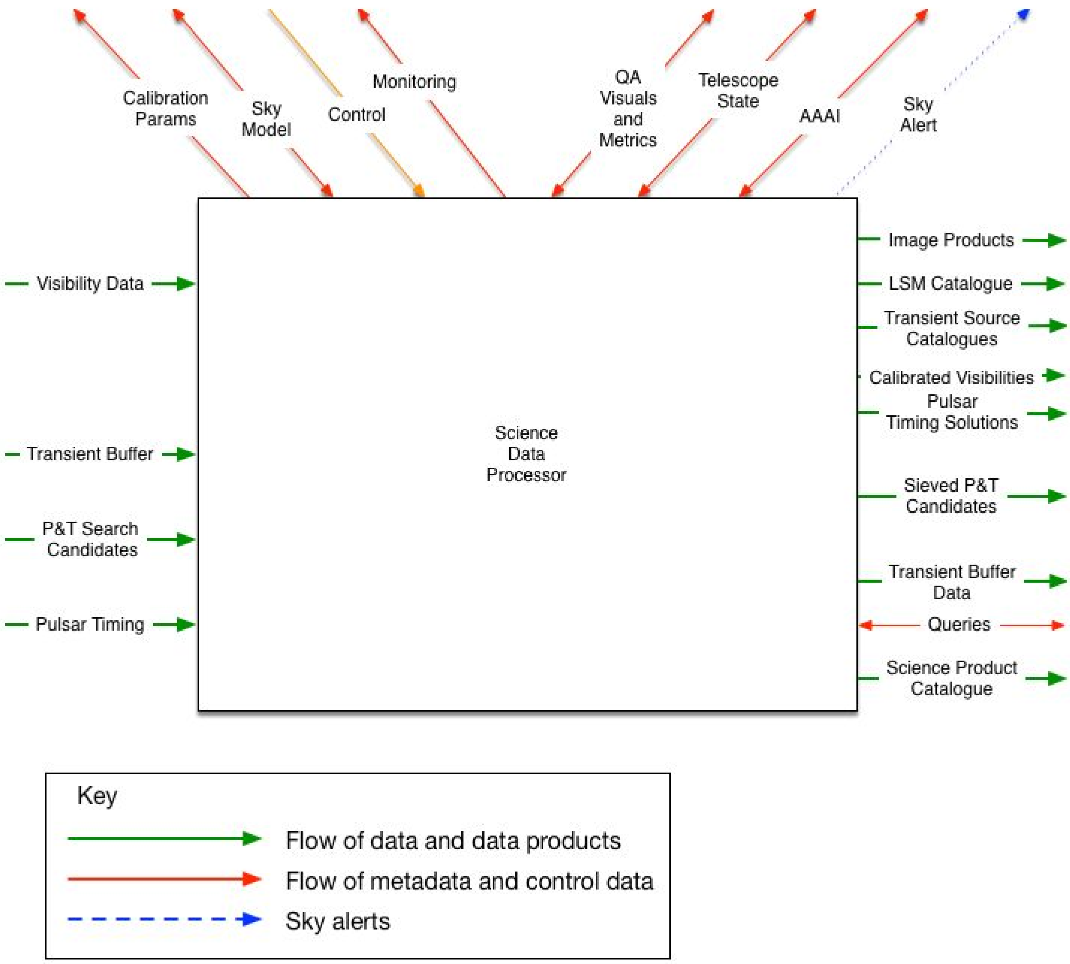

Figure 3: *Diagram depicting the scope of the SDP in terms of the information and data to be communicated. All data are communicated via interfaces to other elements of the SKA telescopes or SKA Regional Centres (see system overview below). Raw astronomical data are input and combined with data and metadata labelled Telescope Model and State. These data are processed and output as Science Products which can be queried. Red arrows indicate the flow of metadata and control information; green arrows indicate the flow of data and data products; blue arrows represent sky alerts.*

The scope of the SDP in terms of the nature of interactions is illustrated in Figure 3.

The SDP receives two types of observational data. These are visibility data, received as a continuous flow to be processed and imaged, and non-imaging data (Transient buffer, Pulsar and Transient Search Candidates and Pulsar Timing Data) received as discrete chunks. The SDP provides the ability to process both these data independently (as may be the case during commissioning) or commensally (as expected during operation).

Additionally, the SDP supports commensal operation and the processing of multiple observational programmes each of which utilise a subset of the telescope — the so-called sub-arraying of or multi-beaming for aperture arrays.

Along with observational data, Control, Monitoring, Feedback and Event information is provided from and to the rest of the SKA system as shown in Figure 3 [AD05]. These data and metadata (operating under different cadence depending on the processing being performed) provide information on:

- real-time **Calibration** solutions for updating the SKA system
- providing access to the **Sky Model** for use elsewhere in the SKA system and to enable maintenance of the sky model by the SKA observatory staff
- **Control** information for the SDP system (via TANGO) and for processing workflows (scheduling blocks)
- **Monitoring** information describing the health-state of the SDP and providing resource information for managing the execution of the scheduling blocks (via TANGO)
- **Quality Assessment** data in the form of visual information for Operator intervention and metrics for Quality Assessment .
- metadata describing the **Telescope State** and **Telescope Model** including configuration information, empirical parameters etc.
- **Alerts** which can be generated and promulgated from the SDP, via the control system, to the SKA to permit follow-on processing.

The SDP allows appropriately privileged external users and users within the Observatory to query the metadata associated with the data products. The required Authentication, Authorisation, Allocation and Identity (AAAI) management information about users is requested from TM. The result of such a query may be the bulk transfer of data to an SKA Regional Centre (SRC) over International WANs and National Research and Education Networks (NRENs).

The SKA Regional Centres will be required to support a query client and to support a function to receive data from the SDP. The SDP produces a number of standard data products which may be maintained at both SDP sites.

2. Documentation Roadmap

2.1 Documentation Scope

The architectural documentation referenced by this document is preliminary and is released for review at the pre-CDR project milestone for SDP. The views are therefore not complete and further development of the views included in this review is planned to CDR.

The documentation approach taken is a mix of the “Views and Beyond” approach of the Software Engineering Institute (SEI) and the SKA System Engineering documentation standards. The configured documents are discussed in Sections 2.2 and 2.3. Additionally supporting and referenced documents are referenced. For supporting documents, the SDP has maintained a memo series in a similar fashion to similar development projects in astronomy.

The documentation is submitted to the SKAO as a set of signed documents.

2.2 How the Documentation is Organised

The documentation is organised into a set of views. Where appropriate a given view is organised hierarchically via view packets.

Additional documentation exists beyond the architecture documentation and includes the following system engineering configured documents

Non-Architecture Documents	
Construction & Verification Plan	SKA-TEL-SDP-0000047 Rev 03C
Construction Schedule	SKA-TEL-SDP-0000047 Rev 03C
Requirements specification and Compliance Matrix [AD03]	SKA-TEL-SDP-0000033 Rev 02C
RAMs analysis	SKA-TEL-SDP-0000115 Rev C
Risk Register	SKA-TEL-SDP-0000052 Rev 07C
Cost Breakdown	SKA-TEL-SDP-0000046 Rev 03C SKA-TEL-SDP-0000114 Rev 01
Operations Plan	SKA-TEL-SDP-0000081 Rev 01C
ILS Plan	SKA-TEL-SDP-0000050 Rev 03C
Configuration Items list	SKA-TEL-SDP-0000048 Rev 02C
PBS	SKA-TEL-SDP-0000064 Rev 02C

Table 1: System Engineering artefacts and other configured items document list

Supporting documentation is provided via the SDP memo series. The SDP memo series may be found at the following web address:
https://confluence.ska-sdp.org/display/SDPDocumentLibrary/SDP+Memos

2.3 View Overview

The SDP Architecture is documented according to the SEI recommendations for Architectural documentation as a series of “Views” and “View Packets”. The views form a hierarchical structure onto the documentation.

The entry point for most stakeholders will be the Component and Connector view. This view gives a view onto the instantiated run-time components of the system, their interfaces and multiplicity.

The following views form the architecture for the pre-CDR review of SKA-SDP

- Component and Connector Views
  - SDP Operational System [RD01]
    Component and connector views are a standard SEI style indicating the instantiation of components within the system at run-time and interfaces between them. Multiplicity is also shown which is a key aspect of the distributed parallel SDP architecture.

    View packets:
    - Processing [RD02]
      Here the scalability of the processing architecture is addressed
    - Delivery [RD03]
      Delivery provides a critical interface to the data products produced by SDP and includes the cataloguing, indexing and management of these data products both within the observatory, but also the technical aspects of product management in relation to their delivery to SKA Regional Centres.

  - Platform [RD04]
    The platform is a key element of the SDP architecture and is conceived around the architecture of cloud-like environments to provide large aspects of the resilience, reliability and portability of the SDP architecture.

- Module Views
  - SDP System [RD05]
    The system module view most clearly demonstrates the maintainability requirement for the SDP and also how algorithms will be supported as well as the ability of the architecture to make use of multiple (and evolving) execution frameworks which may have different characteristics.

    View packets
    - Processing components [RD06]
      This view demonstrates the requirements for support of current best practice algorithms and also the maintainability and extensibility of the architecture from the point of view of new and emerging algorithm support.

- Data Model Views
  - SDP Operational System [RD07]
    This view illustrates the main data models that underpin the SDP architecture, their relationship and the relationship to SKA-system data models.

- Security View

- System [RD08]
  System-level security view for SDP

- Computation and scaling view
  - Performance model [RD09]
    The performance model provides a textual description of the analysis and model which is used to calculate the required computational performance for the workflows to be run on the SDP. The model has been developed for the most computationally intensive workflows and steps therein and does not cover all the possible workflows that will run on the SDP system.

- Functional architecture view
  - Allocation of functional decomposition to products [RD10]
    The L2 requirements are allocated to functions and these functions are then allocated to products. The mapping from requirements to functions is in general many to few. The L2 to functional mapping is available in the full L2 requirements specifications.

- Use Case Views
  - Science pipeline management [RD11]
    The use case view shows how the Science Operations team as "Actors" interact with the SDP to create, modify and submit science pipeline workflows.

- Hardware Views
  - Hardware decomposition view [RD12]
    The hardware view is a specific view to represent required abstractions of the structure of the hardware in order to achieve the required performance of the system.

2.4 How Stakeholders Should Use the Documentation

The key stakeholders in the SDP documentation including all documents are:

- Detailed design and architecture team [Des]
- Expert technical reviewers [Rev]
- SKAO Senior Leadership Team [SLT]
- SKAO System Architecture Team [SAT]
- SKAO Computing and software team [SST]
- User: Telescope operator [Ops]
- User: Platform operational support team [Plt]
- User: Science operations [Sci]
- Development teams [Dev]
- Hardware integrators [HW]

	Des	Rev	SLT	SAT	SCT	Obs	Plt	Sci	Dev	HW
Non-architecture documents										
Construction and Verification Plan			X	X	X				X	
Requirements specification and Compliance Matrix	X	X	X	X	X				X	
RAMs analysis		X					X			X
Risk Register	X	X	X	X	X					
Cost Breakdown	X		X	X						
Operations Planning		X			X	X		X		
ILS Plan		X			X	X	X			X
Configuration Items list		X			X		X		X	
Product Breakdown Structure		X			X		X		X	
Views										
Component and Connector Views										
SDP Operational System	X	X		X	X	X	X	X	X	
Processing	X	X			X				X	
Delivery	X	X			X				X	
Platform	X	X		X	X				X	
Module Views										
SDP System	X	X		X	X	X	X	X	X	
Work flows	X	X		X	X			X	X	
Data Model Views										
SDP Operational System	X	X		X	X				X	
Security View										
System	X	X		X	X	X	X		X	
Functional architecture View										
Allocation of functional decomposition to products	X	X		X	X				X	
Computation and Scaling View										
Performance Model	X	X		X	X			X		
Use Case Views										
Science pipeline management		X		X	X	X		X		
Hardware, Deployment and Allocation Views										X
Hardware Decomposition View	X			X	X		X			X

Table 2: Mapping documents to stakeholders

2.5 View overview planned for CDR

As discussed in Section 2.1 these documents represent a snapshot of the architecture development suitable for a pre-CDR review. The planned list of views for the complete CDR is listed below - new views or view packets are shown italicised. The list is subject to modification as the architecture work completes.

- Component and Connector Views
  - SDP Operational System
    View packets:
    - Processing
    - *Buffer and data management (initial coverage in SDP Operational System)*
    - Delivery
    - *Execution control (initial coverage in SDP Operational System)*
    - *Quality assessment (initial coverage in SDP Operational System)*
    - *Model databases (initial coverage in SDP Operational System)*
    - *Queues (initial coverage in SDP Operational System)*

  - Platform

- Module Views
  - SDP System
    View packets
    - Processing components
    - *Platform*

- Data Model Views
  - SDP Operational System
    V*iew packets*
    - *Visibility raw data*
    - *Non-visibility raw data*
    - *Science data model*
    - *Intermediate data products*
    - *Data products*
    - *Science data product catalogue*

- Data flow and management
  - *Data flow view*
  - *Data management view*

- Security View
  - System

- Computation and scaling view
  - Performance model

    - *Computational Scaling*

- Functional architecture view
    - Allocation of functional decomposition to products

- Use Case Views
    - Science pipeline management
    - *SDP Control*
    - *SDP maintainability*

- Hardware Views
    - Hardware decomposition view
    - *Deployment view*

3. How a View is Documented

We have adopted a standard template for each view based on recommendations from the Software Engineering Institute [RD13]. Most views adhere to the following template:

---

1. Primary Representation

2. Element Catalogue
- 2.1. Elements and Their Properties
- 2.2. Relations and Their Properties
- 2.3. Element Interfaces
- 2.4. Element Behavior

3. Context Diagram

4. Variability Guide

5. Rationale

6. Related Views

7. References
- 7.1. Applicable Documents
- 7.2. Reference Documents

8. Version History

The performance and hardware decomposition views are not a standard view within the SEI documentation and we have adopted an ad-hoc structure appropriate to these views.

4. System Overview

4.1 Context of SKA-SDP and Variability [RD01]

The context of the SDP within the SKA telescope is shown in Figures 4a and 4b. While there is one SDP software architecture, it is deployed to two SKA telescope sites. The context of these two instances differs slightly, with SDP interfaces interfacing to different telescope elements in the two instances.

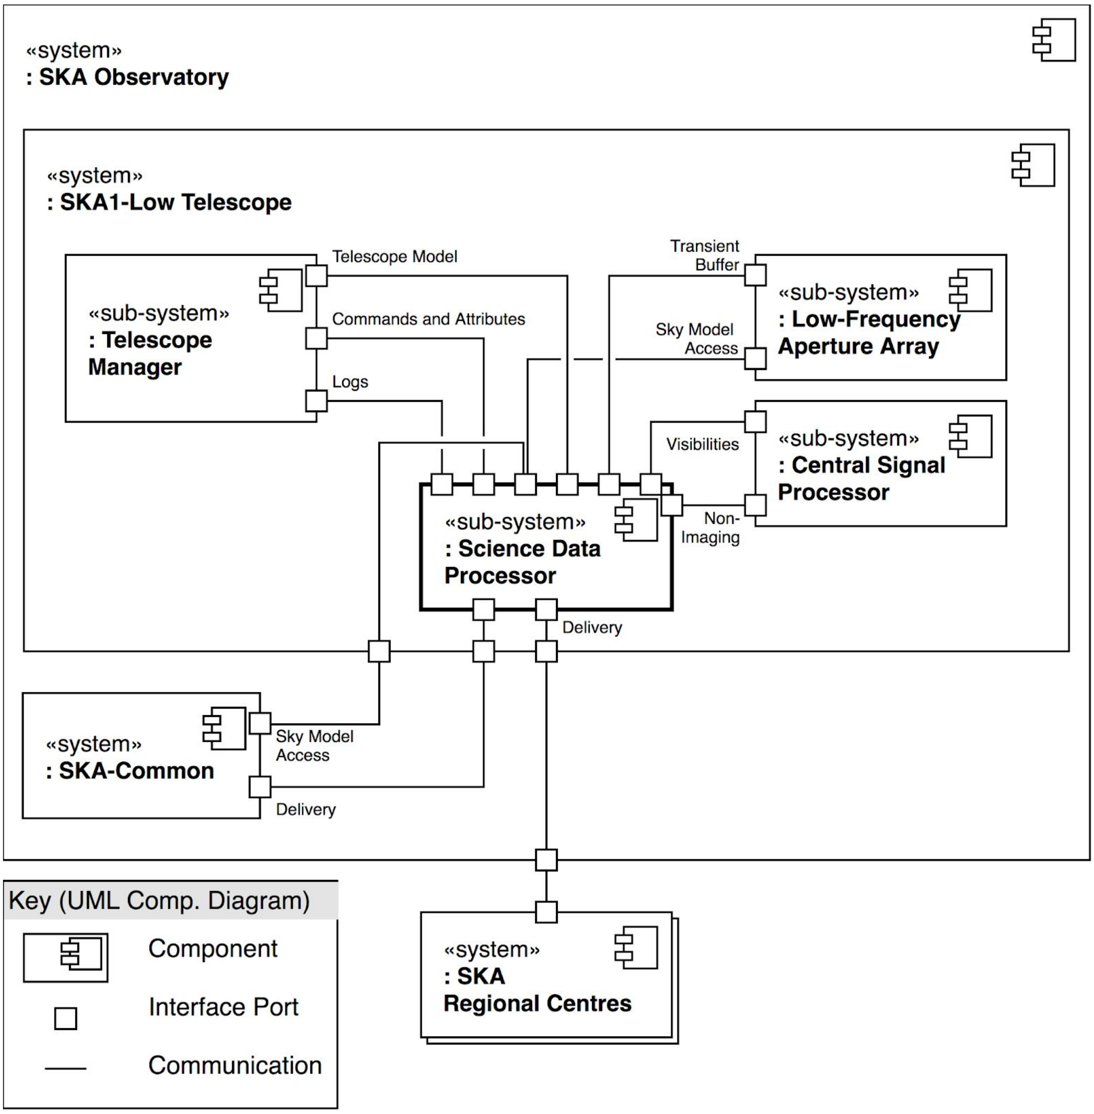

Figure 4a: *Context Diagram depicting the scope of the SDP within the SKA1-Low telescope. Raw data are ingested into SDP from the Central Signal Processor: Visibility data, **Vis**; Pulsar Search Candidates **PSS**; pulsar timing data, **PST**; **transient buffer.** SDP makes the sky model available to other elements of the telescope and enables updating by the SKA observatory staff. Control, monitoring and quality assessment information is communicated between Telescope Manager and SDP via the **TANGO** interface. SDP queries from **Telescope Manager** information on the **Telescope Model** and **Telescope State** and updates the telescope state with calibration solutions as required. Monitoring information in the form of logs is also passed to **Telescope Manager**. Interfaces to the Observatory are provided to for maintenance of the sky model and specifying delivery policies.*

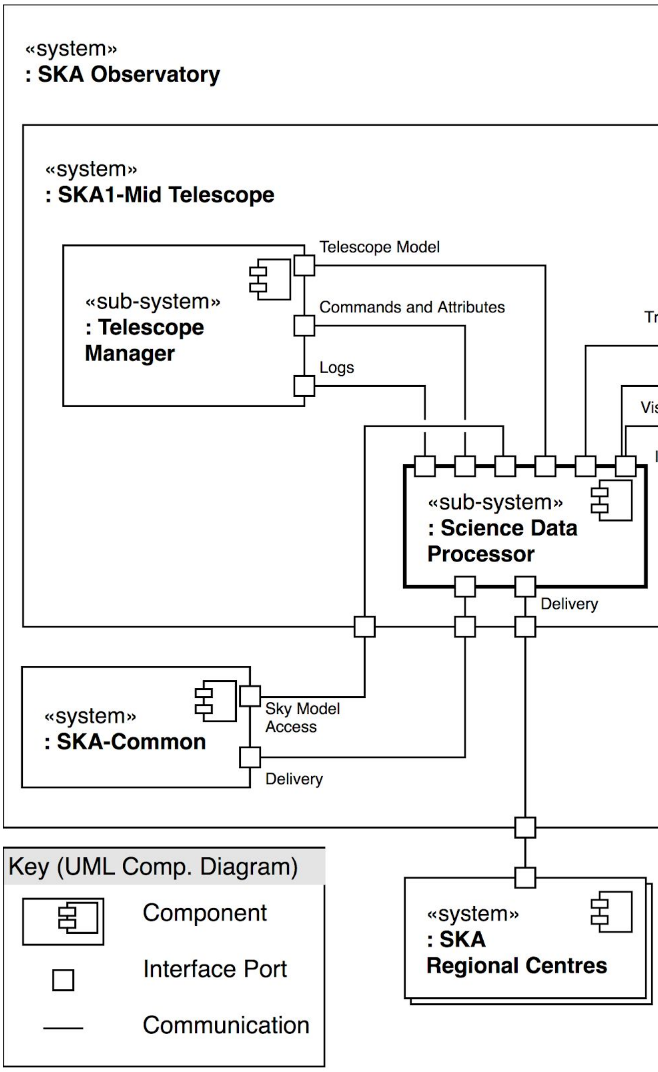

Figure 4b: *Context Diagram depicting the scope of the SDP within the SKA1-Mid telescope. Raw data are ingested into SDP from the Central Signal Processor: Visibility data, **Vis**; Pulsar Search Candidates **PSS**; pulsar timing data, **PST**; **transient buffer.** SDP makes the sky model available to other elements of the telescope and enables updating by the SKA observatory staff. Control, monitoring and quality assessment information is communicated between Telescope Manager and SDP via the **TANGO** interface. SDP queries from **Telescope Manager** information on the **Telescope Model** and **Telescope State** and updates the telescope state with calibration solutions as required. Monitoring information in the form of logs is also passed to **Telescope Manager**. Interfaces to the Observatory are provided to for maintenance of the sky model and specifying delivery policies.*

4.2 Control and Operational Concepts [AD04, AD05]

Unlike a normal HPC or Cloud-like infrastructure the SDP is under the control of the overall telescope Control and Monitoring system often referred to in current documentation as the Telescope Manager, TM.

The SDP control concept provides top-level control over the SDP resources, both hardware and software, to ensure that the required capabilities of the overall telescope can be realised. Thus, the SDP control concept encompasses both the internal control and monitoring of the individual components of the SDP and the interface(s) between the SDP and other telescope subsystems.

The telescope has an overall schedule of observations that is constructed taking into account the availability of resources of all of the elements of the telescope. The SDP responds to the schedule it receives and, commensurate with the actual observations, must ingest data in real time. At the same time, SDP performs other real-time processing which is required for the operation of the telescope (real-time calibration), and for time-critical science operations. In the case of SDP, this includes assessing system resource availability which is affected by the processing of data from completed observations. In planning the schedule of the telescope, the Telescope Manager system uses modules supplied by SDP to determine current and projected resource allocations.

SDP receives a series of scheduling blocks from Telescope Manager which contain both the observations to be conducted but also, as a sub-component of the scheduling block, one or more processing blocks which instruct SDP how to analyse the data. Typically a processing block will instruct SDP to perform one of a number of predetermined workflows using parameters appropriate the observational data, and also define the data products to be produced and added to the science data product catalogue.

There is additional overall control of SDP from the Telescope Manager giving a number of defined commands typically to instruct SDP to transition between well-defined states. The SDP element must also produce monitoring information both of the hardware and software, but also of the progress and quality of reduction the workflows achieve. The latter are distinct and are handled by the Quality Assessment functions. Communication of the control and monitoring information is in large part managed by interfaces which support TANGO which is the adopted control and monitoring layer of the telescope.

4.3 Architectural aspects common across all views

The requirements of the SDP mean that there is no one architectural pattern used to implement the SDP. Instead a number of different architectural patterns are used. However a number of architectural considerations are common across the architecture. These include:

- Scalability. The raw required peak performance of the SDP across both sites is of order 250PFlop. The ingest rate of raw data into the system is of order 1 TByte/s and the sustained read rate from persistent storage is of order ten times this value. Scalability of the software architecture is essential so that an appropriately sized hardware system can be utilised.
- Cost control and software reuse. Although the SDP presents a new software architecture there is a significant requirement not to “reinvent the wheel” wherever possible. This is highly desirable for cost control but also schedule control and

reliability of reusing, where possible, tested software modules. The system architecture has therefore been designed with this in mind - for example the architecture of the platform services closely follows open-source products such as OpenStack while not making a downselection to a particular technology so that an appropriate solution can be adopted at the time of construction.

- Modifiability. The detailed workflows which the SDP will run must reflect state of the art scientific algorithms at the time of telescope operation. It is therefore essential to designing for these workflows to be upgraded. Enabling modified and new workflows to be developed by observatory staff (and potentially others) is a significant consideration behind the architecture. This is reflected architecturally by separating the running of the workflows and the services they require from the workflows themselves (which can be thought of as scripts linking computationally intensive components).
- Availability. The system is required to have high levels of availability and resilience. This has been the fourth underlying architectural consideration, designing the system around, wherever possible, loosely coupled components with stateless control - control state is managed via a resilient configuration database.

5. Mapping Between Views

5.1. SDP system-level module decomposition mapping to SDP C&C view

Allocation matrix:
https://docs.google.com/spreadsheets/d/1xaCXRVO9GvBdx89w5MlLbjthZDKcs39uJW-pmXcJWX8

Table 3: System-level Module to Component Allocation matrix. The association depicted in this matrix is Module *implements* Component or Component *is implemented by* Module.

6. Rationale

6.1 Why a new software for radio astronomy

The SKA-SDP is a new architecture for a platform to deliver analysis of radio astronomy data. Other such software have been and are actively being developed, most notably CASA for the JVLA and ALMA which is a general package, but also many more specialised software packages. We therefore consider why a new architecture is required.

The Science Data Processor (SDP) challenge has aspects which, when considered together, make it unique among comparable systems in astronomy. These include:

1. The SDP is an intrinsic element of the SKA telescopes and not a separately scheduled, remote processing facility. Hence:
    - The SDP will need to be scheduled as an integral part of the observatory, i.e., the data ingest, the raw data storage and processing (into science data products) will need to be carefully coordinated. In contrast, in typical

observatories the data ingest and data processing are largely decoupled by an archive that permanently stores all of the raw data.
   - It is also very different from standard HPC facilities, which do not usually need to manage near real-time systems with long distance data delivery and very high data delivery rates.
2. The SDP processes the incoming data via a set of pipelines. The computational requirements to process this incoming data into scientifically useful data products are significantly greater (by approximately two orders of magnitude) than the largest systems currently used in astronomy and must be able to operate largely autonomously.
   - For this reason the capital costs and operational costs associated with the compute hardware become very important considerations. In order to minimise the storage requirements, as well as meet the real-time requirements where applicable, the processing of each observation needs to be parallelised. This is in contrast to the typical situation at radio-interferometric facilities where many observations are indeed processed in parallel but with a limited degree of parallelism in the processing of each observation.
3. The incoming data rate is so high that raw data are unlikely to be kept permanently. Also, the temporary storage of raw data will need to be minimised (to perhaps as short as 6 hours). This has the implication that data processing and Quality Assessment will need to be automated with little or no possibility for intervention by operators or scientists.
4. The SDP will need to perform some of the data processing within strict deadlines (e.g., around 15 s for real-time calibration).
5. The SKA telescopes are novel and very large facilities. Past experience with (at their time) similarly ground-breaking facilities has shown that once the SDP is online, considerable scientific benefits can be achieved through modifying, improving and adding to the algorithms exploited in the SDP. This means that the SDP must have sufficient flexibility to allow such long-term improvement.
6. The requirements for the SDP element are evolving and will continue to do so into the operational phase of the telescope. During the designed 50 years lifetime of the Observatory the key science objectives will almost certainly change significantly and thus the requirements for the SDP element will evolve as well.
7. The lifetime of the telescope, the lifetime of the compute hardware and the need to minimise power consumption are such that the hardware element of the SDP will need to be refreshed, or completely replaced, on a relatively frequent timescale. The software is very likely to need corresponding updates.

There are existing packages which individually meet one or a few of the above requirements but none meet all. One of the main architectural drivers for SDP is maintainability and extensibility, and for this reason we have selected an architecture which uses a single integrated code base of reusable processing components, and such an architecture is not possible with without a substantial new development (as opposed to selective upgrades to existing packages).

6.2 Main requirements driving the architecture and adopted architectural approach

The notable high-level architectural decision we have made are:

1. Buffer as a single, top-level, component with a file-system interface.
   Rationale: This architecture de-couples data acquisition, data processing and reprocessing, and delivery of science data using a generic component with many available implementations. It offers a standard interface allowing early start of implementation of modules using it.
2. Queues as a high-level communication mechanism
   Rationale: Queues with sharding are a proven highly scalable and efficient architecture for frequently updated data, e.g., in very large-scale web applications. We enforce this communication pattern to ensure scalability of the SDP processing.
3. Multiple concurrent real-time and batch-processing components on a common, flexible, SDP hardware platform
   Rationale: The driver is the need to minimise the size of the SDP computing system, and for this reason the SDP (and SKA) have adopted a flexible, load-balanced over a number of days, scheduling of the telescope and the processing of the acquired data. The chosen architecture best supports such a mode of operation
4. Multiple execution engines using a common set of processing components
   Rationale: The driver is minimisation of maintenance cost and enable extensibility. A single, integrated, code base of processing components minimises the long term maintenance costs.
5. Two-tiered model for execution engines: a high-level workflow tier which orchestrates data movement and process execution, and a lower tier with support for potentially highly-connected processing.
   Rationale: The driver for this decision is scalability. The pattern allows SDP to exploit coarse-grained data parallelism where it exists within a processing block (e.g. because it contains multiple sub-bands to be processed independently) using highly scalable workflow execution engines.

7. References

7.1. Applicable Documents

The following documents are applicable to the extent stated herein. In the event of conflict between the contents of the applicable documents and this document, **the applicable documents** shall take precedence.

[AD01]	SKA-TEL-SKO-0000002 SKA1 System Baseline Design V2, Rev 03
[AD02]	SKA-TEL-SKO-0000008 SKA1 Phase 1 System Requirement Specification, Rev 11
[AD03]	SKA-TEL-SDP-0000033 SDP Requirements Specification and Compliance Matrix, Rev 02C

[AD04] SKA-TEL-SKO-0000307 SKA1 Operational Concept Documents, Rev 02

[AD05] 000-000000-010 SKA1 Control System Guidelines, Rev 01

7.2. Reference Documents

The following documents are referenced in this document. In the event of conflict between the contents of the referenced documents and this document, **this document** shall take precedence.

[RD01] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Operational System Component and Connector View

[RD02] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Operational System C&C Processing View Packet

[RD03] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Operational System C&C Delivery View Packet

[RD04] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Platform Component and Connector View

[RD05] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP System Module View

[RD06] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP System Module Workflows View Packet

[RD07] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05:SDP Operational System Data Model View

[RD08] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP System Security View

[RD09] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Computation and Scaling View

[RD10] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Functional Architecture View

[RD11] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Science Pipeline Management Use Case View

[RD12] SKA-TEL-SDP-0000013 SDP Architecture, Rev 05: SDP Hardware Decomposition View

[RD13] Clements, P. *et al.*, Documenting Software Architectures: Views and Beyond, 2nd edition, 2010

8. Version History

Version	Date of Issue	Prepared by	Comments
05	2018-04-23	P. Alexander et al.	Prepared for M20 Pre-CDR submission