



SKA1 SDP OPERATIONS PLAN

Document Number..... SKA-TEL-SDP-0000081
 Document Type..... MPL
 Revision..... 02
 Author..... J. Taylor, M. Ashdown, J Coles, R. Simmonds
 Date..... 2018-10-31
 Document Classification..... Unrestricted
 Status..... Released

Name	Designation	Affiliation	Signature
Authored by:			
John Taylor	SDP Team	University of Cambridge	 John Taylor (Oct 29, 2018)
			Date: <input style="width: 80%;" type="text"/>
Owned by:			
			Date: <input style="width: 80%;" type="text"/>
Approved by:			
			Date: <input style="width: 80%;" type="text"/>
Released by:			
Paul Alexander	SDP Project Lead	University of Cambridge	 Paul Alexander (Oct 30, 2018)
			Date: <input style="width: 80%;" type="text"/>

DOCUMENT HISTORY

Revision	Date Of Issue	Engineering Change Number	Comments
C	2016-03-24	-	Prepared for SDP delta-PDR review
D	2016-05-24		Prepared for SDP delta-PDR close-out
01	2016-07-20		Prepared for SDP delta-PDR final sign-off
01C	2018-04-25	ECP-SDP-180001	Prepared for SDP pre-CDR review
02	2018-10-31	ECP-SDP-180002	<p>Prepared for SDP CDR review (M21)</p> <p>SDP PRECDRs addressed:</p> <p>Major:</p> <p>SDPPRECDR-78 SDP benchmark suite</p> <p>SDPPRECDR-220 Operations Plan Timeline issues</p> <p>SDPPRECDR-239 SDP SDP scheduling & interface with TM</p> <p>SDPPRECDR-240 start-up/shut-down process</p> <p>SDPPRECDR-411 Observatory Support Tools Requirements</p> <p>Minor:</p> <p>SDPPRECDR-80 SDP_REQ-822</p> <p>SDPPRECDR-146 Sentence doesn't parse</p> <p>SDPPRECDR-147 Don't understand argument</p> <p>SDPPRECDR-150 AD numbering</p> <p>SDPPRECDR-151 Affordability requirement</p> <p>SDPPRECDR-152 Portability requirement</p> <p>SDPPRECDR-217 Incomplete operational success metrics</p> <p>SDPPRECDR-218 SciOps role in SDP operation</p> <p>SDPPRECDR-219 Simulator/emulator milestone confusion</p> <p>SDPPRECDR-242 prevent operational impact</p> <p>SDPPRECDR-293 Space for in place Upgrades</p> <p>Typo:</p> <p>SDPPRECDR-79 typo</p> <p>SDPPRECDR-241 typo</p> <p>SDPPRECDR-339 http/https</p> <p>SDPPRECDR-340 workflow management systems as standards</p>

DOCUMENT SOFTWARE

	Package	Version	Filename
Word processor	Google Docs		SKA-TEL-SKO-0000000-01_GenDocTemplate
Block diagrams			
Google docs Add-ons	Cross Reference Table of contents List of figures		<p>Used for figure & table numbering and references.</p> <p>Used for heading numbering.</p> <p>Used to generate list of figures and tables</p>

ORGANISATION DETAILS

Name	SDP Consortium
Lead Organisation	The Chancellor, Masters and Scholars of the University of Cambridge The Old Schools Trinity Lane Cambridge CB1 1TN United Kingdom
Website	www.ska-sdp.org

© Copyright 2018 University of Cambridge



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

TABLE OF CONTENTS

1 INTRODUCTION	6
1.1 Purpose of the document	6
1.2 Scope of the document	6
1.2.1 Scientific success metrics	7
1.2.2 Operational success metrics	7
2 REFERENCES	8
2.1 Applicable documents	8
2.2 Reference documents	8
3 SDP OVERVIEW	9
3.1 Processor Platform	10
4 SCIENCE OPERATIONS	11
5 ENGINEERING OPERATIONS	12
5.1 Availability	12
5.2 Standards	13
5.3 Operational performance monitoring	13
5.3.1 Health and Status Monitoring	13
5.3.2 Fault Handling and Status Reporting	14
5.3.3 Error Detection	15
5.3.4 Error Recovery	16
5.4 Maintenance Process	16
5.4.1 DevOps	16
5.4.2 Maintenance Plan	17
5.4.3 Support Tools and Instrumentation	18
5.4.4 Hardware Lifecycle Management	18
5.4.5 Upgrades	19
5.4.6 Acceptance Process	20
6 OPERATIONS PLAN TIMELINE	22
6.1 SDP Operation	22
7 APPENDIX 1 SDP HARDWARE	27
8 APPENDIX 2 SDP SOFTWARE MODULES	29
9 APPENDIX 3 L2 OPERATIONAL REQUIREMENTS	29

LIST OF FIGURES

Figure 1	Platform Services Component and Connector View	10
Figure 2	Component and Connector View showing External Interfaces to SDP	14

LIST OF TABLES

Table 1	Tasks Underlying the Maintenance Strategy	17
Table 2	Acceptance Tests and Location	19
Table 3	Preliminary Overview of Acceptance Tests in line with Construction Plan	25
Table 4	Product Breakdown Structure and LRU assignation (Preliminary)	26
Table 5	High Level Software Modules and High Availability Strategy	28
Table 6	L2 Operational Requirements	31

LIST OF ABBREVIATIONS

AAAI	Authentication, Authorisation, Access and Identification
Ai	Inherent Availability
AIV	Assembly Integration and Verification
AR	Array Release
BOM	Bill Of Material
C&C	Component and Connector
CDR	Critical Design Review
CICD	Continuous Integration Continuous Delivery/Deployment
COTS	Commercial off the Shelf
CSP	Central Signal Processing
DCIM	Data Centre Infrastructure Management
ECC	Error Correcting Memory
FAT	Factory Acceptance Tests
FMEA	Failure Modes and Effects Analysis
FMECA	Failure Modes, Effects and Criticality Analysis
FRACAS	Failure Reporting and Corrective Action System
FV	Functional Verification
H/A	High Availability
ILS	Integrated Logistic Support
iOBL	Interim Operational Baseline
LOW	Low Frequency : Telescope
LRU	Line Replaceable Unit
LSA	Logistic Support Analysis
MID	Mid Frequency : Telescope
MTBcF	Mean Time Between Critical Failures
MTBF	Mean Time Between Failure

MTTR	Mean Time To Repair
NIC	Network Interface Card (typically Ethernet)
NSDN	Non-Science Data Network
OBL	Operational Baseline
OEM	Original Equipment Manufacturer
PBytes	Peta Bytes
PBS	Product Breakdown Structure
PHS&T	Packaging, Handling, Storage & Transport
PV	Performance Verification
UAT	User Acceptance Tests
RAM	Reliability, Availability, Maintainability
SAFe	Scalable Agile Framework environment
SAT	Systems Acceptance Tests
SKA	Square Kilometre Array
SKAO	SKA Observatory
SIT	Systems Integration Tests
SPC	Science Processing Centre
SRU	Shop Replaceable Unit
TM	Telescope Manager

1 INTRODUCTION

1.1 Purpose of the document

This document provides an Operations Plan for the Science Data Processor (SDP) for the SKA1 Observatory (SKAO) based on the operational context of the SKAO taking into account the criteria for success and the operating constraints of the Science Processing Centres in which the SDP will be housed. Primary drivers for the SDP Operations focus on the availability criteria and ensuring that the SDP maintains optimal efficiency and performance to meet the higher-level observing efficiency, which in turn will influence observing project completion. SDP Operations, like other elements, will be bound by constraints imposed by hosting agreements with the AUS/RSA authorities, relevant Health & Safety [SDP_REQ-755] and Environmental considerations, and Operational budget constraints.

1.2 Scope of the document

The Operations Plan will evolve over time as the SDP Commissioning and AIV support system, together with feedback from SKA system-level experiences, provide valuable input into the SDP Operations Plan plan, in particular with respect to the necessary tools for maintenance, configuration management, performance monitoring, continuous integration methods, benchmarks and documentation.

The document refers to SDP Level 2 Operational Requirements as listed in the Operational Requirements section at the end of the document (Appendix 3). Operational Requirements are labelled [SDP_REQ-XXX], where “XXX” is the number relating to the specific requirement.

For consistency with the SDP Hardware PBS held in the Configuration Item Data List [AD10] specific terminology is used in this document. The term “server” is used to refer to a compute server. Where an allocation of the server to a specific component or attribute is inferred, the term “node” is used. For example where a particular server will support a component of Platform Services, this is referred to as a Platform Services node.

The operational model described in this document is designed to maximise the operational success of the SKA Observatory. Success, in this context, is defined by the science and engineering metrics described below taken from the SKA1 Operational Plan Concept [AD1] and the methods proposed for Operational Performance.

1.2.1 Scientific success metrics

Several scientific success metrics will be monitored once the observatory becomes operational:

- the over-subscription of observing time is a measure of community demand for access to the facility. This metric will be determined during each time allocation cycle;
- the number of publications, subject to defined acceptance criteria including peer review, is a measure of the observatory’s productivity. This metric will be tabulated at least annually through a combination of web searches and manual reviews of the literature;
- the number of citations to included publications is a measure of scientific impact. This metric will be tabulated as required, primarily through web services such as the NASA Astrophysical Data System;
- the number of publications or citations per unit cost is a measure of value for money.

This list is not exhaustive, and will likely evolve over time, it nevertheless encapsulates the primary drivers for the SKA1 operational model [AD1].

1.2.2 Operational success metrics

Scientific success and operational success are intimately linked: a highly-efficient observatory, for example, will enable more science time on sky, which should in turn produce more science papers per unit time.

The following operational success metrics will be monitored once the Observatory becomes operational:

- system down time due to faults;
- system down time due to unavailability of computational resources;
- system down time due to planned maintenance;
- operational availability;

- operational availability of specific capabilities (specific bands and specific observing modes);
- observing efficiency (integration time per unit available time);
- production of scientifically correct results;
- observing project completion;
- and safety record.

This list is not exhaustive, and will likely evolve over time, it nevertheless encapsulates the primary drivers for the SDP Operational Model [AD1].

2 REFERENCES

2.1 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.

- [AD1] SKA-TEL-SKO-0000307 SKA1 Operational Concept Document
- [AD2] SKA-TEL-SDP-0000013 Rev 06 SDP Architecture Rev 06
- [AD3] SKA-TEL-SDP-0000047 SDP Construction and Verification Plan
- [AD4] SKA-TEL-SKO-0000915 Engineering Report Nov 2017 to Mar 2018
- [AD5] SKA-TEL-SDP-0000050 SDP Integrated Logistics Support Plan
- [AD6] SKA1 MID TM - SDP ICD 300-000000-029 Rev 05
- [AD7] SKA1 LOW TM - SDP ICD 100-000000-029 Rev 05
- [AD8] SKA-TEL- SKO-0000484 SDP TO INFRA-AUS AND SKA SA Rev 02
- [AD9] SKA-TEL-SDP-0000033 SDP L2 Requirements Rev 03
- [AD10] SKA-TEL-SDP-0000048 SDP Configuration Item Data List Rev 03
- [AD11] SKA-TEL-SDP-0000046 SDP Costing Basis of Estimate Rev 04

2.2 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.

- [RD1] <https://www.openstack.org/science/>
- [RD2] <https://docs.openstack.org/openstack-ansible/latest/>
- [RD3] <https://github.com/SKA-ScienceDataProcessor/p3-appliances>
- [RD4] <https://wiki.openstack.org/wiki/Ironic>
- [RD5] SKA-TEL-SDP-0000137 SKA1 System Integration Prototype Report
- [RD6] SKA-TEL-SDP-0000115 SDP RAM Report
- [RD7] SKA-TEL-SDP-0000151 P3-AlaSKA Prototyping Report
- [RD8] SKA-TEL-SDP-0000165 SDP Memo 068 P3-ALaSKA Monitoring & Logging
- [RD9] <https://rundeck.org>

3 SDP OVERVIEW

In some respects the SDP is consistent with traditional High Performance Computing (HPC) systems in terms of the types of applications it needs to support, the components it needs to achieve the performance it requires (of order 100s of Pflops coupled with many PBytes of multi-tiered storage). Unlike traditional HPC facilities however, the SDP has to operate as an integral part of the SKA observatory and be scheduled and coordinated in real time with the other elements of the SKA. This close coordination of data ingest and data processing, storage and delivery gives rise to important operational requirements on the SDP. In particular the SDP must be able to quickly configure itself to accept input data from the telescopes under instruction from the Telescope Manager, it must be able to calculate in advance how long it will take to process certain observations and derive Scheduling and Processing Blocks, and subsequently it must be bounded within that estimated time, unless instructed otherwise by TM. Additionally, in order to achieve availability requirements of the SKA telescope, it is necessary that maintenance schedules of the SDP are aligned with maintenance schedules for the rest of the SKA telescope with steady-state maintenance being conducted while the SDP is available. Finally, it must offer a Delivery service to the overall Observatory for access to Observatory Data Products via a Preservation product. These aspects are discussed in the SDP High Level Overview [RD11].

These requirements together have led to an architectural representation of the SDP couched in a framework best exemplified by the System Level Module and Dependency View [AD2], such that services required by the Execution Frameworks under the auspices of Execution Control are clearly separated in terms of services and interfaces to support:

- Science Pipeline Workflows;
- Platform Services providing configuration management together with general purpose cluster management ;
- SDP services such as Delivery and Model Databases;
- System Interfaces providing underlying access to O/S, Storage and Accelerator services.

This makes for an architecture that:

- promotes adoption of standard practice in cloud and DCIM techniques automated provisioning of compute and storage infrastructure for multiple tenants, fault isolation and recovery, monitoring and logging, resource reservation, elasticity and automation of self-test;
- promotes best practice in modern data centre operations such as the ability to collect, collate and report logging information from the compute, storage and networking devices to centralised dashboards [SDP_REQ-757]

By adoption of these, significant operational efficiencies can be achieved and result in a reduction of Operational Costs (SDP_REQ-745). Such properties are clearly identified by the quality attributes of these services in [AD2] and define distinct areas for development for both domain and non-domain functionality as described in [AD3]. Maintaining this strategy further ensures that hardware refreshes do not present obstacles to the porting of SDP on such hardware.

3.1 Processor Platform

A key aspect of the Operational aspects of the SDP rests with Platform Services and the Platform C&C View [AD2] (reproduced in Figure 1 for context) and corresponding Operational States of the SDP described therein.

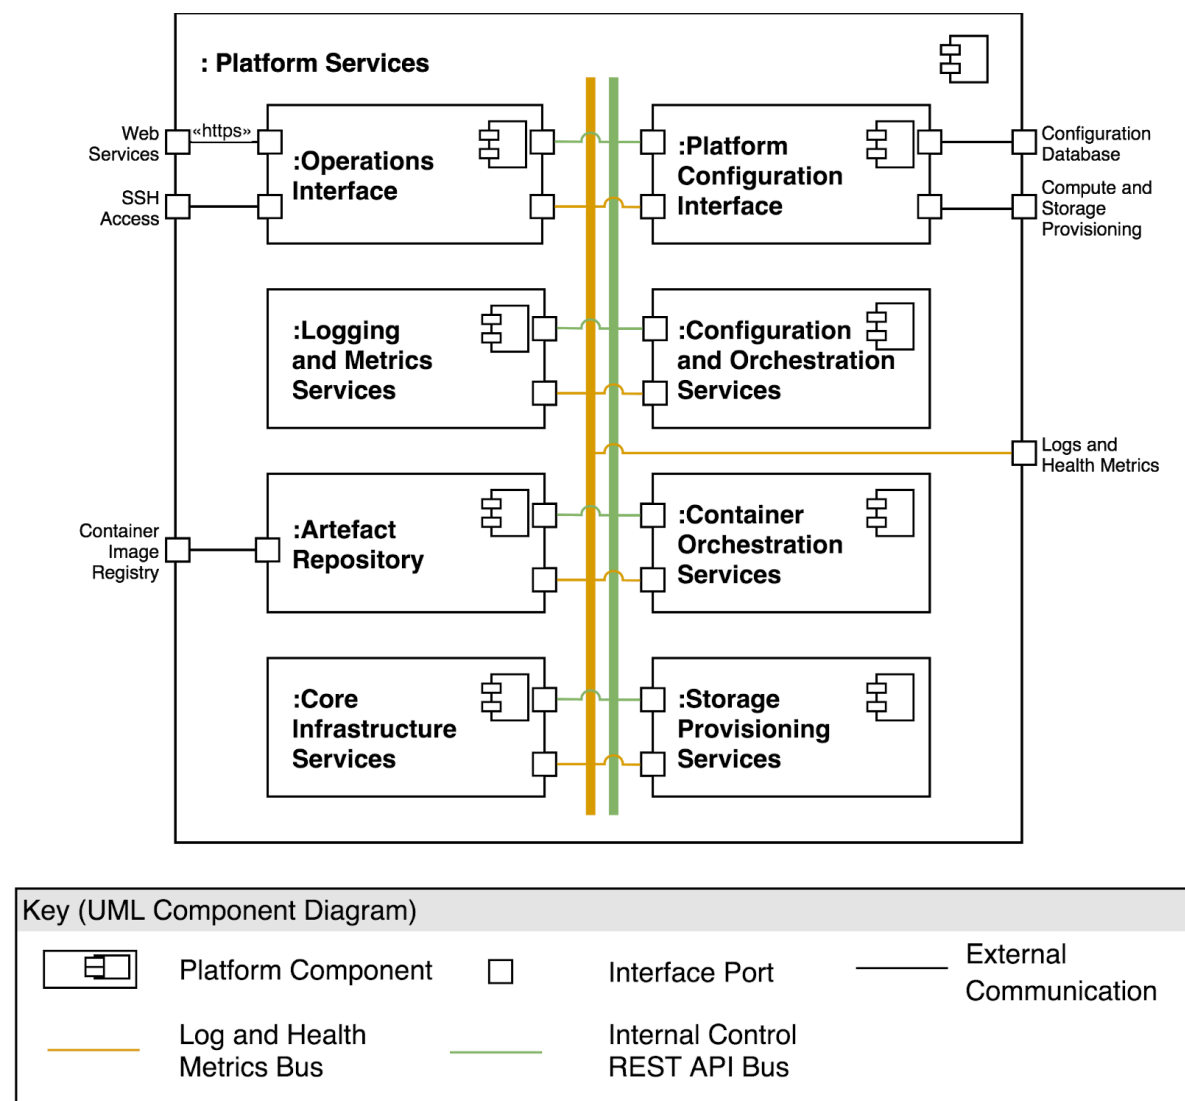


Figure 1: Platform Services C&C View from [AD2]

Platform Services provide critical services for the operation of the SDP in conjunction with Execution Control. As such these services and any persistent state require allocation to the appropriate hardware in order to maintain availability [SDP_REQ-821]. Hardware allocated to the support of these critical services will require appropriate back-up power to avoid any power interruption

[SDP_REQ-782]. On start-up of the SDP the initial state of the SDP will be in low-power mode [SDP_REQ-783] as defined by the Platform Services C&C View [AD2].

Platform Operators have access to the SDP via standard mechanisms [SDP_REQ-285], such as ssh (see SDP Security View [AD2]), with the addition that Health Monitoring and Logging are accessible from https interfaces which provide appropriate dashboards onto the SDP. [SDP_REQ-4, 723]

Core Infrastructure Services is the component that handles the provisioning, management and control of the SDP hardware (software-defined compute, storage and networking). The scale of the SDP systems is such that automation of these functions is essential and we are currently tracking developments in cloud software stacks such as OpenStack [RD1] and infrastructure-as-code environments, such as OpenStack-Ansible [RD2]. Examples of Ansible-Playbooks are already available on the SDP Github [RD3] which will evolve into a suite of CICD tools for use by Platform Operators. In addition to the developments of Cloud ecosystems, the SKAO is collaborating with CERN on OpenStack which could provide shared opportunities in the development of Operation aspects and tooling of the SDP and other sub-elements [AD4]. These aspects are described further in the [P3-AlaSKA Prototyping Report \[RD7\]](#) and related Memos.

These tools will be made available for Platform Operators to ensure that optimal performance of the system can be achieved. This varies from benchmarking tools for various parts of the system (compute, data transport, buffer, etc.) to network analysers and low-level debugging tools. Mostly these are the tools that any competent system administrator and software developer will be aware of and trained to use. In addition, these also include tools to locally access and analyse the performance counters and sensor information available in modern hardware.

4 SCIENCE OPERATIONS

Science Operations encompass those activities which relate to the delivery of science. There is substantial interaction between the overall operations of the SKA and the Science Operations of the SDP as much of the former activity has either a direct or indirect impact on SDP.

The overall operation of SDP, how it is scheduled, which data products are produced etc. is discussed in the SDP High Level Overview [RD11]. Details of how SDP deals with scheduled processing is discussed in the behaviour section of the SDP Operational System Component and Connector View [AD2].

The science operations activities specific to SDP operations are:

- Preparation, maintenance, modification and testing of new workflows / algorithms;
- Development, maintenance and testing of new processing components to be used in workflows;
- Scientific assessment of data quality, algorithmic and telescope performance via data mining and visualisation;
- Delivery of science data to SKA Regional Centres for astronomer access. An analysis of specific use-cases has been made and their impact is described in the Delivery C&C View [AD2].

Science algorithms are defined by parameterised workflows formed from scripts/programs that drive the execution frameworks which themselves make use of processing components to perform all processing. Developing, programming and testing these workflows is a core aspect of the science operations of the SDP. The SDP architecture has been developed with a major architectural driver being the maintainability of these workflows by the observatory science, programming and operations team. This is achieved architecturally by the use of execution frameworks which provide the framework for highly scalable data-driven workflows. Processing components can then be written, typically, to run on individual nodes without the complexity of developing software for a highly parallel, distributed, environment. Only validated workflows and components are run by the SDP, although in early science operations, operator access may require the deployment of software for debugging purposes. The operations activity of developing, modifying and maintaining workflows is addressed in the architecture documentation Science Pipeline Management Use Case View [AD2].

A specific example of the development of workflows is to allow for the scientific assessment of data quality together with the visualisation and analysis of derived information. This is achieved by the development and deployment of workflows which include Quality Assessment components. These quality assessment components provide the mechanism by which operators and science staff can analyse and visualise information produced during a workflow (e.g. statistics on flagging, or quality of the gain solution etc.). Modifying workflows to include Quality Assessment components is a standard aspect of science operations. Additionally science operations staff may access SDP data products via Delivery and perform analysis on these products outside of the SDP within the environment provided by Observatory Support Tools. Modification of workflows to define new data products (perhaps by identifying intermediate data items as data products for the purposes of science operations) is an additional and complementary route to analysis.

Access to data products whether by the observatory staff is via the Delivery interface to SDP. This interface also provides the mechanism to send data products to SKA Regional Centres. Science and operations staff also have an interface to define the policies on data products. These aspects are discussed in more detail in the Delivery C&C View [AD2].

5 ENGINEERING OPERATIONS

5.1 Availability

Engineering Operations comprises those activities necessary for the maintenance and development of the telescopes and the SKA infrastructure in the host countries. The scientific success of the SKA, as described above, requires that the telescopes be available for science observations for as much of the time as possible within the constraint of available resources as defined by [SDP_REQ-762].

Achieving this level of availability requires that all telescope systems (including both hardware and software) be designed for high reliability (high MTBF or more importantly high MTBcF) and simple maintainability (low MTTR). The scale of the SDP demands an industrial level of logistics engineering

covered in [AD5]. The trade-off between MTBF, MTBcF and MTTR is a delicate balance, given the desire to minimise manual intervention, subject to the appropriate level of maintenance, and the cost implicit to gain higher reliability. For example, one aspect of this may be the over-provisioning of spares, which may well be hot-spares, compute servers such that in the event of a fault, the server is not replaced or fixed, but taken out-of-service until a more convenient time. Such practises are widely adopted in HPC centres and given the tolerance of the SDP architecture to such failures, this may well provide a more cost-efficient practise. For other services, such as storage related, appropriate use of redundancy, for example RAID or erasure coding together with H/A configurations will be employed. These will be coupled with recovery of components through the configuration database Operational System C&C View[AD2]. Failures of hardware in this respect may require lower MTTR to ensure availability criteria are met. These are also identified in the RAM Analysis [RD6].

5.2 Standards

The wide adoption of suitable standards [SDP_REQ-745] is a means to mitigate diversity and the inevitable increases in costs for operation as a result. The SDP is active in the investigation of common approaches to solutions that could be adopted both within the SDP sub-systems as well across sub-elements. Of particular concern for operational aspects are:

- Workflow Management tools;
- automated configuration [RD2] and deployment [RD4] tools;
- containerised application environments such as Docker and Singularity;
- Platform Management tools and Cloud-based environments such as found in OpenStack and previously identified in the quality attributes associated with the SDP Module View;
- System Services based, for example, on widely available technologies such as Linux, Object and File Storage Systems, etc.;
- Customised SDP Services.

5.3 Operational performance monitoring

These are the primary metrics that will be used to measure the operational performance of the SDP. The operational requirements to track these metrics are covered elsewhere in this section. In summary:

- System downtime due to faults - tracked through the observing logs and fault database; consistent with the requirement for planned maintenance downtime of 2 hours per year [SDP_REQ-874]
- Meeting Operational availability through implementation of the ILS Plan [AD5] and monitoring to prevent critical failures [SDP_REQ-764];
- Maintain an SDP Performance Model [SDP_REQ-749] to ensure that scheduled observations can be processed within known constraints;
- Ratio of processing to observing time.

5.3.1 Health and Status Monitoring

Platform Services will provide a wide range of standard monitoring points that will be provided to Execution Control to allow TM [AD6 and AD7] to have sufficient system wide visibility to make informed decisions on the operation of the SDP [SDP_REQ-723]. The SDP will handle monitoring in a hierarchical and consequently more manageable fashion by querying aggregate information from the SDP sub-elements such as the Compute and Interconnect Racks, Delivery and Preservation. This information will provide an overall health state with respect to the SDP, together with the actual status of the SDP in terms of usage and operational state. The following diagram (Figure 3) provides a view on the Security View [AD2] representation showing component interfaces between Observatory Support and the NSDN.

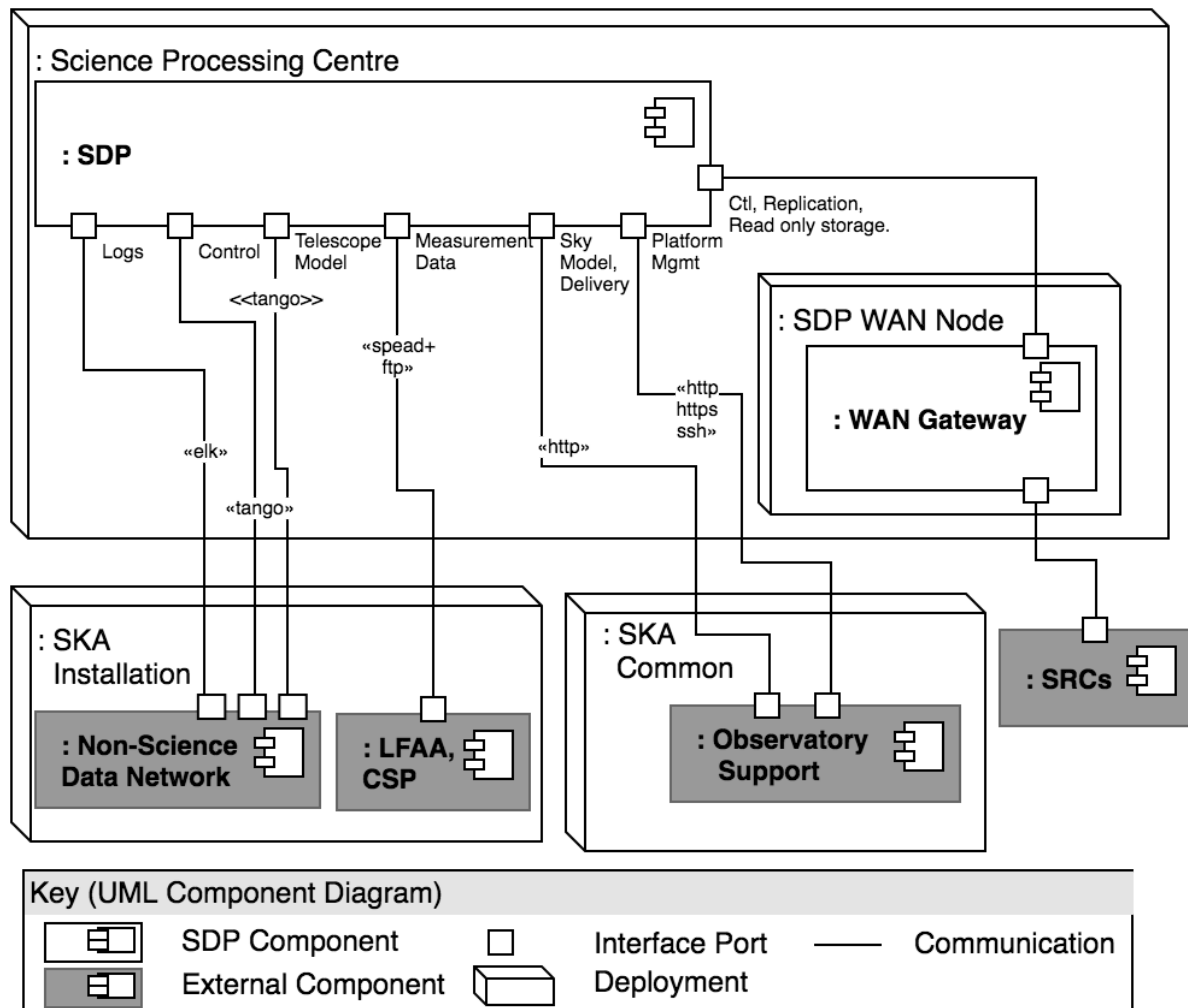


Figure 2: Component and Connector View showing External Interfaces to SDP.

5.3.2 Fault Handling and Status Reporting

The logging of faults will be detected by Platform Services and the nature of the fault automatically described in the log from the error condition and/or fault code allowing information to flow to the Help Desk for Operator intervention which may lead to a Maintenance action. The appropriate course of action will depend entirely on the severity and impact of the fault and defined in the Failure Reporting and Corrective Action System (FRACAS) which will be incorporated as part of the ILS Plan. In exceptional circumstances faults may trigger specific events where risk of damage is, or

may be about, to occur, e.g. where temperature of a hardware component has breached a soft or hard threshold, loss of power or communications. Such faults will preempt any actions by the Operator and render the SDP to a failsafe state [SDP_REQ-52,755].

The following provides a list of Fault handling and status reports which should be integrated with the SKA-wide fault database. At the present time a FMECA has not been performed and will be undertaken as part of the ILS Plan [AD5] and is discussed in the Follow-up Action and Recommendations section of the SDP RAM Analysis [RD6].

- Be informed that a fault has occurred, performed automatically by Platform Services [SDP_REQ-821].
- Identify the nature and cause of the fault.
 - A fault database of currently open and closed faults, and their resolution (or otherwise), should be accessed if the fault or its solution is not familiar to the operator. (e.g. a compute server failure may warrant no further action until an appropriate point of time - scheduled maintenance, taking into account impact on Operational Availability and RAM analysis [RD06]). A software, or configuration, fault may well be able to be corrected while in operation.
 - Attempt limited troubleshooting, depending on the nature and severity of the fault (e.g reboot server, run diagnostic self-test, etc) limited to 10 mins [SDP_REQ-818, 819].
- Isolate the faulty system from the rest of the SDP if this is not done automatically.
 - E.g. Soft-errors in a memory sub-system may warrant shut-down of a server and maintenance.
 - Quality Assessment reports should inform the operator whether there is any significant degradation to the data compared to expectations. [SDP_REQ-821]
- Continue processing if possible.
 - If it is not possible to continue with the present Processing Block, based on the appropriate policies to be decided, then the next feasible Processing Block in the plan should be commenced. In most cases, this will occur automatically.
 - The SDP status will be updated automatically once a critical fault is realised so that the telescope manager can flexibly and dynamically change its observing plan to choose the next feasible observation to execute utilising the available resources or escalate appropriately [SDP_REQ-764].
- Annotate the fault report with information that should contain:
 - A narrative describing the fault, its characteristics and the impact to observations not already described by the automated report;
 - Any corrective actions taken and the amount of observing time lost, if any.
- Notify appropriate personnel of the faults.

5.3.3 Error Detection

Platform Services will be responsible for the aspects of error detection not available through standard OS-level tools by providing methods to combine standard monitoring of the sub-systems alongside Quality Assessment mechanisms implicit to the pipelines. Firstly, errors and state information are broadly classified into alarms and events (SDP_REQ-35). Alarms are events that

without further scope broadening represent a critical failure [SDP_REQ-763] and escalate to a telescope fault and require operator notification as described above (although intervention may still be automatic within the SDP). Events are points of interest, that when taken in a broader system view, perhaps via automatic aggregation, could reveal a higher level issue.

It is also important to be able to mark alarms, and other errors, with a variety of states, such as 'known failure', to manage and prioritize faults. This is especially important for alarms that get triggered at a regular cadence (e.g. ECC memory errors or temperatures above a soft threshold which if left uncorrected may lead over time to a critical failure) – the system must be able to suppress further warnings when directed, but still track the alarm events themselves. Such alarms will be stored and updated within the engineering database and engineering support function requirements [AD4]. Where faults will potentially impact on the inherent Availability (viz. critical) these will be monitored in order to prevent operational impact by preventative maintenance. [SDP_REQ-821].

5.3.4 Error Recovery

The primary assumption of the SDP is that, as much as possible, sufficient loose-coupling exists at the level of the Processing together with SDP and Platform Services to provide roll-back without becoming limited by any particular serial process or global synchronisation [SDP_REQ-820] and discussed in the Rationale in the Operational System Level C&C View [AD2] and Processing Component Module View [AD2]. This requirement is also fundamental to resilience and error recovery. In effect, and where possible, errors will be detected, and that particular operation will simply be discarded until such time as the error can be corrected. Where appropriate as identified in the PBS (see Appendices I and II), redundancy, high availability strategies will be required to ensure that appropriate mitigations of such failures do not lead to critical failures and hence loss of service and to the unavailability of the SDP [SDP_REQ-723,728].

5.4 Maintenance Process

5.4.1 DevOps

DevOps tools are expected to fit into one or more of these categories, reflective of key aspects of the development, support and maintenance process, many have which have been prototyped and discussed in [RD7] together with specific SDP Memos as detailed below:

- Code - development and review, source code management tools, code merging;
- Build - continuous integration (CI) tools and build status such as Jenkins methods ;
- Deployment methods, such as RunDeck [RD9], to automate post-build process and deployment verification.
- Test - continuous delivery/deployment tools allied to the CI tools provide feedback on risks;
- Package - package management for application pre-deployment staging, this may be afforded by use of Containers, for example, Docker and Singularity as discussed in the SIP Prototyping Report [RD5];
- Release - change management, release approvals and release automation, such as described in the P3-AlaSKA OpenStack Prototyping Memo [RD10] which has demonstrated the

piecemeal upgrade of the OpenStack control plane from previous release to current release, by use of Kolla-Ansible;

- Configure - infrastructure configuration and management, Infrastructure as Code tools such as, for example, OpenStack Ansible discussed further in the P3-AlaSKA OpenStack Prototyping Memo [RD10] ;
- Monitor - applications performance and infrastructure monitoring such as provided by, for example, SDP Memo 068: Performance Prototype Platform (P3-ALaSKA) Monitoring & Logging [RD8].

The development of the DevOps environment is expected, by its very nature, to be an ongoing process, although should be in place at the beginning of SDP Operations by Q3 2020 [AD3]. Further discussion with the SKAO will be required as input into the SKAO Operations Plan.

The SDP Architecture makes widespread use of open source software technologies and over the lifetime of the Observatory these open source software technologies could become obsolete and incur technical debt. Throughout the operational lifetime of the SDP Software, technology choices (past and future) should be continuously evaluated in order to appropriately manage technical debt.

5.4.2 Maintenance Plan

A detailed maintenance plan will be an output of Integrated Logistics Support Plan [AD5]. By adoption of suitable tools and processes the maintenance strategy will be able to diagnose faults down to LRU level so that maintenance can be scheduled appropriately [SDP_REQ-756]. For the most part, maintenance will be able to be performed while the rest of the Telescope is in operation by judicious use of redundancy in the design together with appropriate LRU specified in the Product Breakdown Structure [AD10]. This process is identified by the following tasks and discussed in detail in the document:

Maintenance Planning	<p>Define the actions and all the support aspects necessary to ensure that the SDP attains the specified system readiness objectives with minimum Life Cycle Cost;</p> <p>State specific maintenance tasks to be performed on the machine consistent with SDP_REQ-759];</p> <p>Define levels of replacements, repair, task times, testability requirements, support equipment and automatic test equipment needs, training, manpower skills and facilities;</p> <p>Develop the preventive maintenance programme in accordance with observation periods and refine this from experience gained;</p> <p>Analyse the proposed work environment on the health and safety of operations team, and define relevant qualifications and/or training requirements as well as safety equipment;</p> <p>Minimize the use of hazardous materials and take into account local regulations for disposal methods for hazardous, recyclable and non-recyclable waste;</p> <p>Obsolescence Management. This will comprise of regular meetings with technology suppliers to ensure lifetime of current equipment and roadmaps for components are well understood. An equipment asset database will be maintained to track the duty-cycle of the Bill-of-Materials and suppliers should provide an obsolescence management plan (OMP) for all equipment supplied.</p> <p>See sections below on Hardware Lifecycle Management.</p>
----------------------	--

Support Technical Documentation	Technical manuals; Spare parts list; Preventative (scheduled and unscheduled) maintenance instructions; Corrective maintenance instructions including fault tracking; Installation, test, commissioning, acceptance and change management procedures and reports; Drawings and part lists; Specifications; System Software.
Support Tools and Instrumentation	As much as possible automated and Interactive performance measurement tools to assess individual components (e.g. functionality, power consumption, performance, reliability) and sub-systems. Where possible these should reflect the aspects of the actual processing pipelines but also consist of custom and standard industry tests; Equipment for in-situ and ex-situ diagnosis and repair of SDP components (LRU and SRU);
Support Personnel	An assumption is made that Support Personnel are Contracted out to External Organisation(s) - although in the SDP Costing Basis of Estimate [AD11] a number of dedicated support personnel to ensure the operational success of the SDP are identified. These staff will augment the 3rd Party support organisation and provide additional capability in fully understanding the complexity of the SDP from a systematic point-of-view rather than fulfilling maintenance activities.

Table 1: Tasks Underlying the Maintenance Strategy

5.4.3 Support Tools and Instrumentation

As has been noted, the SDP has some similarities with traditional HPC systems in particular around the use of High Performance Compute Server components and high-end networking in order to meet the challenges of multi-Pflop computational capability. This capability typically pushes the solution into the realm of state-of-the-art processing technologies and can lead to the early adoption of components which have very little field data and use. Such circumstances can lead to failures which are particularly difficult to root-cause and can lead to undesirable maintenance situations. Acceptance criterion and strict acceptance tests, both in the Factory (FAT) where typically only suitable modular hardware can be tested and at the SPC (SAT) where final integration happens, are thus paramount in order to avoid these situations. Such circumstances should be taken into consideration for SDP Operational reasons and the AIV programme of work too, as discussed in the Construction Plan. In addition to the acceptance tests, close collaboration with systems vendors and 3rd party suppliers (e.g. accelerator providers) should be maintained. This may involve the use of vendor specific simulators to assess performance well-ahead of general availability of these components.

5.4.4 Hardware Lifecycle Management

The following process will be taken into account:

- A database as part of the Core Infrastructure Services to store all hardware attributes as part of the inventory, e.g. Manufacturer, product revision, firmware version, etc.¹ as provided by OpenStack Ansible Manage Inventory plug-ins.;
- Flexible and complete Hardware introspection;
- Flexible API to add/query server attributes (e.g. OpenStack Nova);
- Burn in and acceptance process (CICD);
- Define when resources are available to application through State workflow as described in the Platform Services C&C View [AD2];
- Policy needs to allow segregated access to the different teams (operators, developers), which may involve multi-tenancy;
- Clear retirement procedure or obsolescence management.

5.4.5 Upgrades

Although the exact nature of the specific hardware design is, of course, not known at present, the SDP Hardware Decomposition View [AD2] and PBS [AD10] define a highly modular view of the proposed SDP hardware. This aspect provides not only aids the design of the Software Architecture in terms of the multi-faceted nature of the SDP, as shown by the SDP Operational C&C View [AD2] but also, this modularity provides clear advantages in terms of upgradeability of the SDP hardware.

Upgrades should be considered both in terms of a proposed roll-out of the initial SDP system, meeting the cost goals and capabilities of the telescope and upgrades of the system in terms of embracing new, more cost-effective hardware technologies in conjunction with obsolescence of hardware from a particular vendor and will be factored by aspects of the SDP Infra ICDs [AD8]

The Hardware Decomposition View [AD2] proposes modularity based on (the high level representation is included here for reference)

- A unit of compute and storage based on a Compute (or Server Rack) where the personality of the server in terms of processing, storage or service function is described in addition to the first stage of networking for global communication and control. The rack, therefore provides a unit of upgradeability, which can be tested in isolation before performing a level of system integration.
- A unit of global connectivity or networking based on an Interconnect Rack.
- A unit of Long Term Storage

The amount of additional space, etc. will be determined by the hardware roll-out, procurement and refresh strategy during operations since this will only occur during operations.

¹ This is in addition to the SKA CM management tool

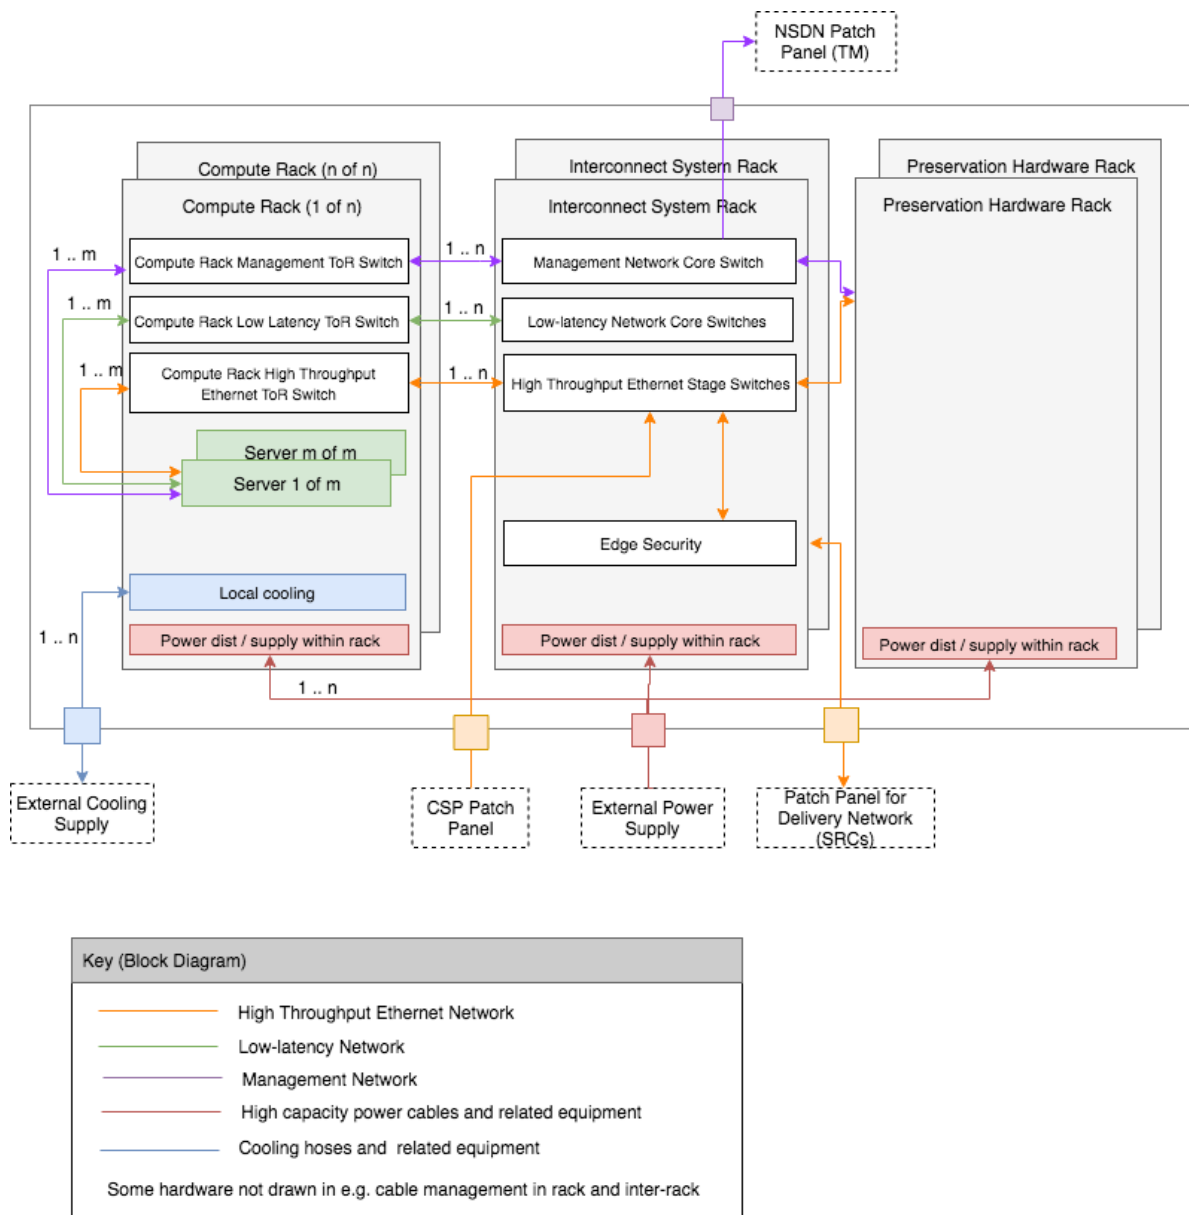


Figure 3: SDP hardware block diagram (from [AD2])

The cadence of upgradeability between these may be different. For example, the second stages of networking providing global connectivity across the whole system may require to be supported over a longer time period than compute server technology and upgrades scheduled appropriately. Similarly the Long Term Storage is typically supported over a 10-15 year life-cycle and specific attention made to the obsolescence and preserving of data as storage technology evolves.

5.4.6 Acceptance Process

The acceptance testing strategy will be defined to stress the individual sub-systems that comprise the SDP PBS and its implementation, alongside the Construction Plan. Tests assume that a certain level of vendor supplied tests (burn-in) have already been performed to remove failed components early on in the system-build (preferably in the factory but also repeated to a lesser extent at the respective SPC to flush-out Packaging, Handling, Storage and Transport (PHS&T) effects) and then

performance tests to ensure that the system is behaving in a manner that meets expectations and satisfies requirements. These performance tests will be governed by the performance verification suite commensurate with the Construction Plan [AD3]. These will consist of both domain specific tests and industry standard tests to exercise specific hardware (Storage, Compute Server and Networking). The activities to be conducted during Acceptance Testing of each Construction phase (as appropriate - see Table 2) will include, but not be limited to, the following:

- BOM Consistency and Configuration Management;
- Vendor-supplied system diagnostic tests to verify that all components of the system are functioning within specifications.;
- Interface testing;
- The system’s capacity and capability measured using an SDP Benchmark Suite [AD3]. (At the time of writing this does not exist but it could be included as part of the DevOps built-in test environment this is expected to evolve during the next phase of the project as we approach Construction);
- A workload composed of components of the SDP Benchmark Suite;
- Testing of the remote management and other operational and administrative attributes of the system, including placing simulated maintenance and support requests to demonstrate that service levels are in place;
- Specific timed and untimed activities to evaluate particular aspects of the system, including but not limited to: measuring system shutdown and cold-start time (including fail-safe state), removal and addition of hardware from/to an operational system, time required to perform a complete backup and subsequent full bare-metal restore of a server’s software and configuration to an operational state as described in the Platform Services C&C View [AD2];
- Simulated failover of and recovery of scientific workflows operating under Execution Frameworks.

The following Table provides a description of potential Acceptance Criteria and appropriate Tests. These tests will be developed as part of the CI DevOps Built-in Test framework.

Acceptance	Test	Description/Examples	Place
FAT	Specification	Document Defining the solution and specification of the system	Factory
	Inventory	MAC Address, Firmware Version, Model and Serial Number, Platform Services	Factory
	Performance	Limited set of tests to allow shipment	Factory
SAT	Verification	Connectivity (IP Address), Operating System, SDP Libraries, System Service, SDP Services, Size (e.g Bytes), Visual Inspection, Repeat Factory Acceptance to trigger payments	SPC
	Performance	Industry Standard, SDP Benchmark, CI Tests,	SPC

		Platform and SDP Services	
	Failover and Recovery	Customised Tests for H/A Services together with low-power mode and graceful degradation	SPC
SIT	Integration	Operator Interface, Dashboard Services	SPC
	Maintenance	Fault Identification Logging and Repair, Safety	SPC
UAT	Operator	Documentation, Fault Reporting, Training	SKAO

Table 2: Acceptance Tests and Location

Once the acceptance process has been finalised, steady-state operations will prevail.

6 OPERATIONS PLAN TIMELINE

The following table provides a preliminary assessment of the Operational Tests as applied to the Hardware and Software Product Breakdown [AD10]. Tests will be refined and updated during each of the Milestones specified as part of the SAFe engineering process.

6.1 SDP Operation

Mile stone	Key Engineering Goals	Operational Test Software (Appendix II)	Operational Test Hardware (Appendix I)
Tools	SAFe & tools training, testing of tools. DevOps and test platforms in place. SAFe management processes in place and roles filled. Power-up, shutdown and failsafe	Support Tools and Instrumentation. Configuration Management and DevOps see Section 5.1 covering Platform Services and System Services System shutdown and power verification including failsafe.	P.2.1.1.2 Service Nodes to support Devops on P.2.1.1 Compute Rack
Simulators/ emulators	Provide a SDP-TM interface simulator for early testing by TM SDP-CSP Interface Receive visibilities	Verification and Tests available for Receive Server and High Throughput Ethernet together with Platform , System and SDP Services infrastructure	P.2.1.1.2 Receive Server and P2.1.1.3 High Throughput Ethernet Network
PV1	Demonstration of Model Databases and Buffer component scalability. This is an early demonstration of	Verification and Tests available for Receive Server and Capacity	P.2.1.1.2 Receive Server and P2.1.1.2.7.1 (Capacity

	scalability of two critical long-lead time components which are key to the performance of the complete system.	Storage, High Throughput Ethernet together with SDP Services infrastructure	Storage) P2.1.1.3 High Throughput Ethernet Network
First MVP	This would be the first meaningful solution level integration from SDP. Goals: - Working high-level architecture; - External interfaces; - Evaluation & correction of existing code base (prototypes).	Initial Systems Integration and Verification Tests	P.2.1.1.2 Service Nodes to support Devops on P.2.1.1 Compute Rack
PV2	Demonstration of Receive & Real-time processing component performance. This needs to be an early sub-system performance demonstration (against a model, given the hardware requirements) since this sub-system needs to cross over into the commissioning strand quite early on.	Verification and Tests available for Receive Server and Capacity and Performance Storage, High Throughput Ethernet and Low Latency Network together with SDP Services infrastructure. Test Failover and Recovery	P.2.1.1.2 Receive Server and P2.1.1.2.7.1 (Capacity Storage) P2.1.1.3 High Throughput Ethernet Network
PV3	Demonstration of Data Queue and Quality Assessment component scalability. These are grouped together since the underlying property is solubility of stream-like data aggregation and distribution solutions.	Verification and Performance Tests available for SDP Services specified. Test Failover and Recovery	P.2.1.1.2 Service Nodes to support SDP Services
PV4	Performance and functional test of Execution Control and TM interaction.	Verification and Integration Tests available for Execution Control interfaces with SDP and Platform Services infrastructure. Test failover and recovery	P.2.1.1.2 Service Nodes to support SDP and Platform Services
ITF QE product handover	Certain external interfaces: - Basic TM M&C interface (TANGO) - CSP visibility interface (SPEAD) Functionality supported at AA1	Functional Tools for - Ingest (receive & pre-processing) - Generate Ingest QA metric data	P.2.1.1.2 Service Nodes to support Devops on P.2.1.1 Compute Rack (minimal configuration)

	<p>scale:</p> <ul style="list-style-type: none"> - Visibility data Ingest (receive & pre-processing) - Generate Ingest QA metric data - Store data on disk in suitable format (i.e. MS) for processing with existing (3rd party) tools - Process data with existing (3rd party) tools for ITF signal displays (subset of commissioning displays). 	<p>3rd Party Tools supporting:</p> <ul style="list-style-type: none"> - Store data on disk in suitable format (i.e. MS) for processing with existing tools - Process data with existing tools for ITF signal displays <p>Basic Platform Services (Provisioning & Management)</p> <p>System shutdown and power verification including failsafe.</p>	<p>Milli-SDP (0.1% of the size of the final system) hardware deployment at ITF.</p>
PV5	<p>Demonstration of the execution engine performance at full system scale, with a focused test that tests scheduling, task-startup, memory management and data transfer overheads.</p> <p>This will demonstrate execution engine performance at full scale, against a model of how fast the execution should taken given the available hardware.</p> <p>As a sub-system test, it won't use real computational components, or other sub-systems such as the buffer.</p>	<p>Performance Tests for Receive and Processing Node.</p> <p>Extend SDP and Platform Services infrastructure verification testing</p>	<p>P.2.1.1.2 Service Nodes to support SDP, Platform and System Services with P.2.1.1.2 Receive and Processing Nodes</p>
FV at modest scale	<p>End-to-end processing from data block to calibrated image data with simplified calibration & imaging requirements.</p>	<p>Extend System Integration Test</p>	<p>P.2.1.1 Compute Rack</p>
PV6	<p>Test (or investigate) the architectural decomposition of a hierarchical scalable system. Test to see if the scalability within data islands and between data islands is as required for scalability to the full system. Use a workflow that exercises different communication requirements, e.g. Direction-Dependant calibration features.</p>	<p>Verification and Tests of Low Latency and High Throughput Ethernet Networks tests incorporated into Systems Performance Testing Framework.</p> <p>Extend Platform and System Services</p>	<p>P.2.1.1 Compute Rack and P.2.1.2 Interconnect System</p>
AA1	<p>Provides complete commissioning</p>	<p>Functional and</p>	<p>P.2.1.1.2 Service Nodes</p>

product handover	<p>and ITF/AIV support functionality including functionality required to support AA2.</p> <p>Certain external interfaces:</p> <ul style="list-style-type: none"> - Basic TM M&C interface (TANGO) - CSP visibility interface (SPEAD) <p>Functionality supported at AA1 & AA2 scale:</p> <ul style="list-style-type: none"> - Visibility data Ingest (receive & pre-processing) - Generate Ingest QA metric data - Basic real-time calibration - Store data on disk in suitable format (i.e. MS) for processing with existing (3rd party) tools - Process data with existing (3rd party) tools for: <ul style="list-style-type: none"> - commissioning signal displays - basic imaging - basic time series analysis 	<p>Commission Tools for</p> <ul style="list-style-type: none"> - Ingest (receive & pre-processing) - Generate Ingest QA metric data - Basic real-time calibration <p>Commission Tools for 3rd Party Tools supporting:</p> <ul style="list-style-type: none"> - Imaging - Time Series Analysis <p>Store data on disk in suitable format (i.e. MS) for processing with existing tools</p> <ul style="list-style-type: none"> - Process data with existing tools for ITF signal displays <p>Enhanced Platform Services</p>	<p>to support Devops on P.2.1.1 Compute Rack (minimal configuration) Milli-SDP (0.1% of the size of the final system) hardware deployment at MID and LOW sites.</p>
PV7	<p>Testing scalability of interactions between SDP performance components with a workflow including at least DD CAL and MFMS features.</p> <p>The focus of this milestone will be the efficiency of interaction between the execution engine, model databases, data queues and the batch processing.</p>	<p>Hot Buffer verification and Performance Tests incorporate into Systems Performance Testing Framework. Extend Platform, SDP and System Services. Failover and Recovery</p>	<p>P.2.1.1 Compute Rack and P.2.1.2 Interconnect System</p>
PV8	<p>Demonstration of whole system performance, at the full required scale, when running the ingest, real-time calibration and fast imaging pipelines.</p> <p>This should be a separate (and earlier) milestone to PV10 since these features have quite different requirements to PV10 and are likely to be used at their full scale earlier in the scientific operations of the telescope. Although a system test, this milestone mostly tests the performance of the real-time processing components.</p>	<p>Systems Performance Testing Framework Completed</p>	<p>P.2.1.1 Compute Rack and P.2.1.2 Interconnect System with focus on P.2.1.1.2 Receive Server P.2.1.1.2 Storage Node P.2.1.1.2.7.1 Capacity Storage</p>
Software	<p>Software suite for SKA to evaluate</p>	<p>Systems Acceptance</p>	<p>N/A</p>

suite for hardware vendors	(benchmark) hardware. Required for procurement of AA4 hardware therefore s/w needed 12 months before AA4 SDP deployment.	Testing Framework Released	
PV9	Demonstration of whole system linear (or near linear) scalability up to full required scale with a complex feature set, e.g., MFMS synthesis with DD calibration. This is the scalability (but not performance) precursor to milestone PV10.	Hot Buffer Verification and Performance Tests	P.2.1.1 Compute Rack and P.2.1.2 Interconnect System Exercising P.2.1.1.2 Processing Node
SPC MID integration	Integration of SDP system with solution level once the long distance SaDT links (from site to SPC) are available (AA2). This would be the first time SDP systems are not running at ITF or on site, but in the SPC. This is ahead of SDP hardware deployment for AA3 at the SPC. Look at issues with respect to messaging, packet loss, etc.	High Throughput Ethernet verification tests complete. Extend Platform Services	P.2.1.1 Compute Rack and P.2.1.2 Interconnect System with focus on P.2.1.1.2 Receive Server P.2.1.1.2 Storage Node P.2.1.1.2.7.1 Capacity Storage
SPC LOW integration	Integration of SDP system with solution level once the long distance SaDT links (from site to SPC) are available (AA2). This would be the first time SDP systems are not running at ITF or on site, but in the SPC. This is ahead of SDP hardware deployment for AA3 at the SPC. Look at issues with respect to messaging, packet loss, etc.	High Throughput Ethernet performance tests complete. Failover and Recovery	P.2.1.1 Compute Rack and P.2.1.2 Interconnect System with focus on P.2.1.1.2 Receive Server P.2.1.1.2 Storage Node P.2.1.1.2.7.1 Capacity Storage
SDP AA3 s/w product handover	Sufficient functionality to perform real-time calibration and basic imaging workloads in the context of a system controlled by TM with appropriate control, monitoring and reporting. Minimal set of functionality for the AA3 release of the SDP Operational System to work at AA3 scale.	Systems Integration and User Acceptance	P.2.1.1.2 Service Nodes
Centi-SDP H/W ready (AA3)	Integrate SDP hardware deployment (at the MID & LOW SPC) at solution level & integrate SDP software with centi-SDP hardware before AA3.	Systems Integration Complete	P.2.1 SDP Compute Hardware

PV10	<p>Demonstration of performance of whole SDP system at the required scale and performance level when running a spectral line imaging pipeline.</p> <p>This is a natural precursor to PV11 below and while this will not test the intricacies of data movements and scheduling of the most complex pipelines it will be a good whole system throughput test. Although this tests the complete system, the feature being tested will most exercise the Buffer component and the efficiency of the computational tasks. The stress on the execution engine will be small.</p>	Buffer and Processing Server Performance and Verification	P.2.1 SDP Compute Hardware
SDP 1st operational deployment. (AA4 product handover)	<p>Integrate SDP hardware deployment (at the MID & LOW SPC) at solution level & integrate SDP software with SDP hardware for AA4.</p> <p>Full software and hardware functionality and performance available at both SPCs.</p>	System Integration Tests Complete	P.2.1 SDP Compute Hardware
PV11	<p>Demonstration of SDP whole-system performance at the required scale and performance level and with a full feature set. E.g.: full-scale MFMS synthesis with DD calibration stages; full-scale fast-imaging and calibration test.</p> <p>This would be the final performance milestone: essentially a complete demonstration that SDP is achieving the performance that is required (or expected).</p>	System Acceptance Test, Systems Integrations Test, UAT	P.2.1 SDP Compute Hardware

Table 3: Preliminary Overview of Tests in line with Construction Plan. PV = Performance Verification. FV=Functional Verification

7 APPENDIX 1 SDP HARDWARE

The intent is that SDP hardware as defined by [AD10] will be the same for both MID and LOW Telescopes, thus for convenience the following Table decomposes the hardware for MID only. For reference the Configuration Item for the major hardware component is provided and subdivided further in terms of previous Product codes used to identify inventory holding.

Configuration Item [AD10]	Previous PBS Reference	Inventory	Spares
304-000001	P.2.1 SDP Compute Hardware MID	N/A	N/A
304-000003	P.2.1.1 Compute Rack	SRU	
	P.2.1.1.1 Compute Rack Management Switch	LRU	Yes
	P.2.1.1.2 Server [Service, Receive, Processing, Storage] P.2.1.1.2.1 Latency optimized cores P.2.1.1.2.2 Main memory P.2.1.1.2.3 Throughput optimized cores P.2.1.1.2.4 High Throughput Ethernet NIC P.2.1.1.2.5 Low Latency network NIC P.2.1.1.2.6 Out-of-band management unit (BMC)	LRU SRU LRU SRU LRU LRU SRU	Yes
	P.2.1.1.2.7 Storage P.2.1.1.2.7.1 Capacity Storage P.2.1.1.2.7.2 Performance Storage	LRU	Yes
	P.2.1.1.3 Compute Rack High Throughput Ethernet Switch	LRU	Yes
	P.2.1.1.4 Compute Rack Low Latency Switch	LRU	Yes
	P.2.1.1.5 Cabling	LRU	Yes
	P.2.1.1.6 Racks Infrastructure P.2.1.1.6.1 PDU P.2.1.1.6.2 Bulk power supply P.2.1.1.6.3 Cable Management P.2.1.1.6.4 Local cooling P.2.1.1.6.5 Rack	SRU LRU LRU LRU LRU SRU	Yes
304-000004	P.2.1.2 Interconnect System P.2.1.2.1 Low Latency Network Core Switch P.2.1.2.2 Management Network Core Switch P.2.1.2.3 High Throughput Ethernet Core Switch P.2.1.2.3.1 Pluggable optics P.2.1.2.3.2 Fibre P.2.1.2.3.3 Copper cabling P.2.1.2.3.4 Patch panel	LRU	Yes Yes Yes Yes Yes Yes
	P.2.1.2.4 Edge Security P.2.1.2.5 Interconnect System Rack	LRU SRU	Yes

304-000005	P.2.1.3 Inter-rack Infrastructure P.2.1.3.1 Hot Aisle/Cold Aisle equipment P.2.1.3.2 High capacity power equipment P.2.1.3.3 Liquid cooling equipment (hoses, etc) P.2.1.3.4 Cross-rack cable management	SRU	
304-000002	P.2.2 SDP Preservation Hardware P.2.2.1 Hierarchical Storage Management P.2.2.2 Intermediate Storage P.2.2.3 Long Term Storage	SRU LRU	Yes

Table 4: Product Breakdown Structure and LRU assignation (Preliminary)

8 APPENDIX 2 SDP SOFTWARE MODULES

The Operational System Level C&C View [AD2] and Processing Component Module View [AD2] discuss software reliability in more detail. Table 6 highlights aspects of the reliability strategy for software modules.

Module	Reliability
Execution Control	H/A Configuration
Platform Services	H/A Configuration
SDP Services	H/A Configuration
System Services	Handled by Platform Services
Processing and Other Modules	Handled by Execution Control

Table 5: High Level Software Modules and High Availability Strategy

9 APPENDIX 3 L2 OPERATIONAL REQUIREMENTS

A subset of the Operational Requirements are provided here with a brief description and a pointer to the relevant section where they are discussed. Full details are tabulated in [AD9]. Where an item is marked as TBD further references are provided.

No.	Name	Description	Covered In Section X of document
SDP_REQ-4	SDP Resource Reporting	The SDP shall, on request from the TM, provide a summary of overall resource availability or resource availability for the requested scheduling block.	3.1

SDP_ REQ- 35	Alarms	The SDP shall support the generation of notification events, in the form of alarms, to the TM based on predefined conditions.	5.3
SDP_ REQ- 52	Failsafe	The SDP shall actively ensure that internal failures do not result in a hazardous situation to the systems and personnel with which it interfaces.	5.3
SDP_ REQ- 285	Accessibility	The SDP shall enable per user access to SDP resources (hardware and software) using the Authentication and Authorisation facilities provided by the SKA (as per EN 50600-2-5. Data centre facilities and infrastructures. Part 2-5. Security systems).	3.1
SDP_ REQ- 723	State logging	The SDP shall log its operational states at all times.	5.3
SDP_ REQ- 728	Master controller failover latency	The SDP master controller shall support failover to an operational state within 1s.	5.3
SDP_ REQ- 745	Affordability	The SDP shall be affordable, i.e., the chosen architecture should ideally minimise expenditure on capital and operational costs while meeting the requirements. The cost associated with technology readiness, risk and uncertainty needs to be taken into account when minimising capital and operational expenditure. This may be in conflict with the other design principles and in that case the selected architecture should not cost significantly more than the cost of other possible architectures.	3 & 5.2
SDP_ REQ- 755	Electronic equipment safety	All SDP equipment shall conform to South African and Australian occupational health, safety and environmental standards for electronic equipment as per local legislation in South Africa and Australia.	5.3
SDP_ REQ- 756	SDP part identification	All SDP parts, cables and connection panels shall be uniquely identified down to LRU level.	5.4
SDP_ REQ- 757	SDP equipment installation	All SDP equipment shall follow standard data centre or HPC centre installation practices.	3

SDP_ REQ- 759	SDP Software update downtime	The SDP shall not require the telescope to be offline while performing software updates. Major software updates shall be performed during engineering and maintenance down time periods of the telescope.	5.4
SDP_ REQ- 762	SDP Inherent Availability (Ai)	The SDP shall have an Inherent Availability (Ai) higher than or equal to 99.9%.	5.1
SDP_ REQ- 763	SDP Critical failure identificatio n	The SDP shall identify more than 99% of all critical failures and report them to the TM.	5.3
SDP_ REQ- 764	SDP Isolation of critical failures	The SDP shall isolate 95% of all critical failures and report it to TM.	5.3
SDP_ REQ- 782	Power interruption survivability	The SDP shall survive a main power interruption, at an arbitrary time.	3.1
SDP_ REQ- 783	Start-up power mode	On start-up, the SDP shall enter low power mode until commanded otherwise.	3.1
SDP_ REQ- 818	Software Reboot Time	Software failures of the SDP (TBC-084) software that requires rebooting in order to recover from the failure, shall have a MTTR (recover time) of less than or equal to 10 minutes.	5.3
SDP_ REQ- 819	Software Maximum Allowable Recovery Time	Software failures of the SDP (TBC-0085) software that require a software fix, shall fall back on to a previous working state or isolate the problem in such a way to achieve a MTTR (recovery time) of less than or equal to TBD (>10 mins <<8h) minutes.	5.3 [RD6]
SDP_ REQ- 820	Software failure requiring Rollback Recovery	The following design concerns were identified to enable quick rollback recovery from failures: - Design the system to minimise the need for schema modification to any shared databases. Addition to the schema is easier to roll-back. - Loose coupling between SDP components so that can roll back on a sub-system basis, reducing impact and time to rollback.	5.2

		- There shall be a process to allow rollback which includes a change to the DB schema.	
SDP_ REQ- 821	Failure detection to Achieve Ai	The SDP shall detect failures to allow recovery within the time windows specified in SDP_REQ-818 and SDP_REQ-819.	5.3

Table 6: L2 Operational Requirements.