

# Science Data Processor SKA Engineering Meeting 2015

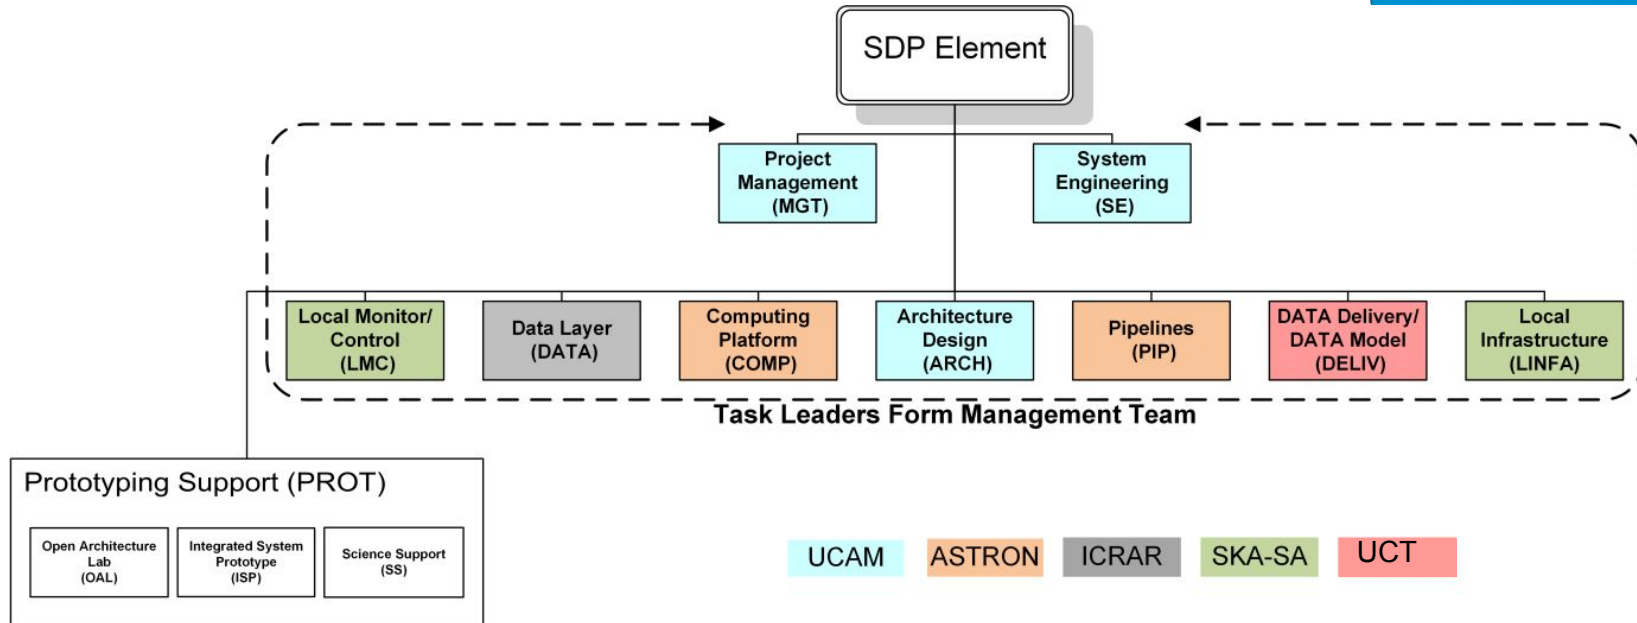
Presented by: **Jasper Horrell**

Wednesday 11th November 2015



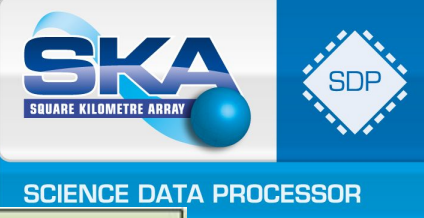
- The SDP Consortium
- High-level task and approach
- Challenges
- Current Status
- Stage-1 highlights and PDR observations
- Approach going forward
- Issues and considerations

# Top Level WBS and Management Team



- |                 |                |          |                 |
|-----------------|----------------|----------|-----------------|
| • Lead:         | Paul Alexander | • COMP:  | Chris Broekema  |
| • PM:           | Jeremy Coles   | • PIP:   | Ronald Nijboer  |
| • Deputy PM:    | Ian Cooper     | • DATA:  | Andreas Wicenec |
| • PE/Architect: | Bojan Nikolic  | • DELIV: | Rob Simmonds    |
| • SE:           | Ferdi Graser   | • LMC:   | Shagita Gounden |
| • PS:           | Rosie Bolton   | • LINFA: | Jasper Horrell  |

# Consortium members & wider engagement



Institute/Member	Partner status
Auckland University of Technology, NZ (AUT)	Full
Canadian Data Centre (CADC)	Full
(Canadian Universities Collaboration)	Full
Centre for High Performance Computing	Full
Chinese Universities Collaboration (UPRC)	Full
Commonwealth Scientific and Industrial Research Organisation (CSIRO)	Full
Forschungszentrum Jülich GmbH: JUELICH	Full
International Centre for Radio Astronomy Research (ICRAR)	Full
Max Plank Institute for Radio Astronomy	Full
Netherlands Institute for Radio Astronomy (ASTRON)	Full
Pawsey Supercomputing Centre (formerly known as IVEC)	Full
SKA South Africa (SKA-SA)	Full
UK Science and Technology Facilities Council - Hartree	Full
University of Cambridge	Full
University of Manchester	Full
University of Oxford	Full
Victoria University of Wellington, NZ (VUW)	Full
Barcelona Supercomputing Centre - Centro Nacional de Supercomputación	Associate
Instituto de Telecomunicações	Associate
University College London	Associate
Institute of Space Science & Astronomy (ISSA) (part of the University of	Associate
Instituto de Astrofísica de Andalucía & Consejo	Associate
Superior de Investigaciones Científicas (IAA-CSIC)	
Fundacion Centro de Supercomputación de Castilla y Leon (FCSCCL)	Associate
University of Cape Town	Joining
University of Groningen	In discussion
Science and Technology Facilities Council (STFC) - UK Astronomy Techno	Interested
Universidad de Chile	Interested
Instituto Nazionale De Astrofisica	Interested

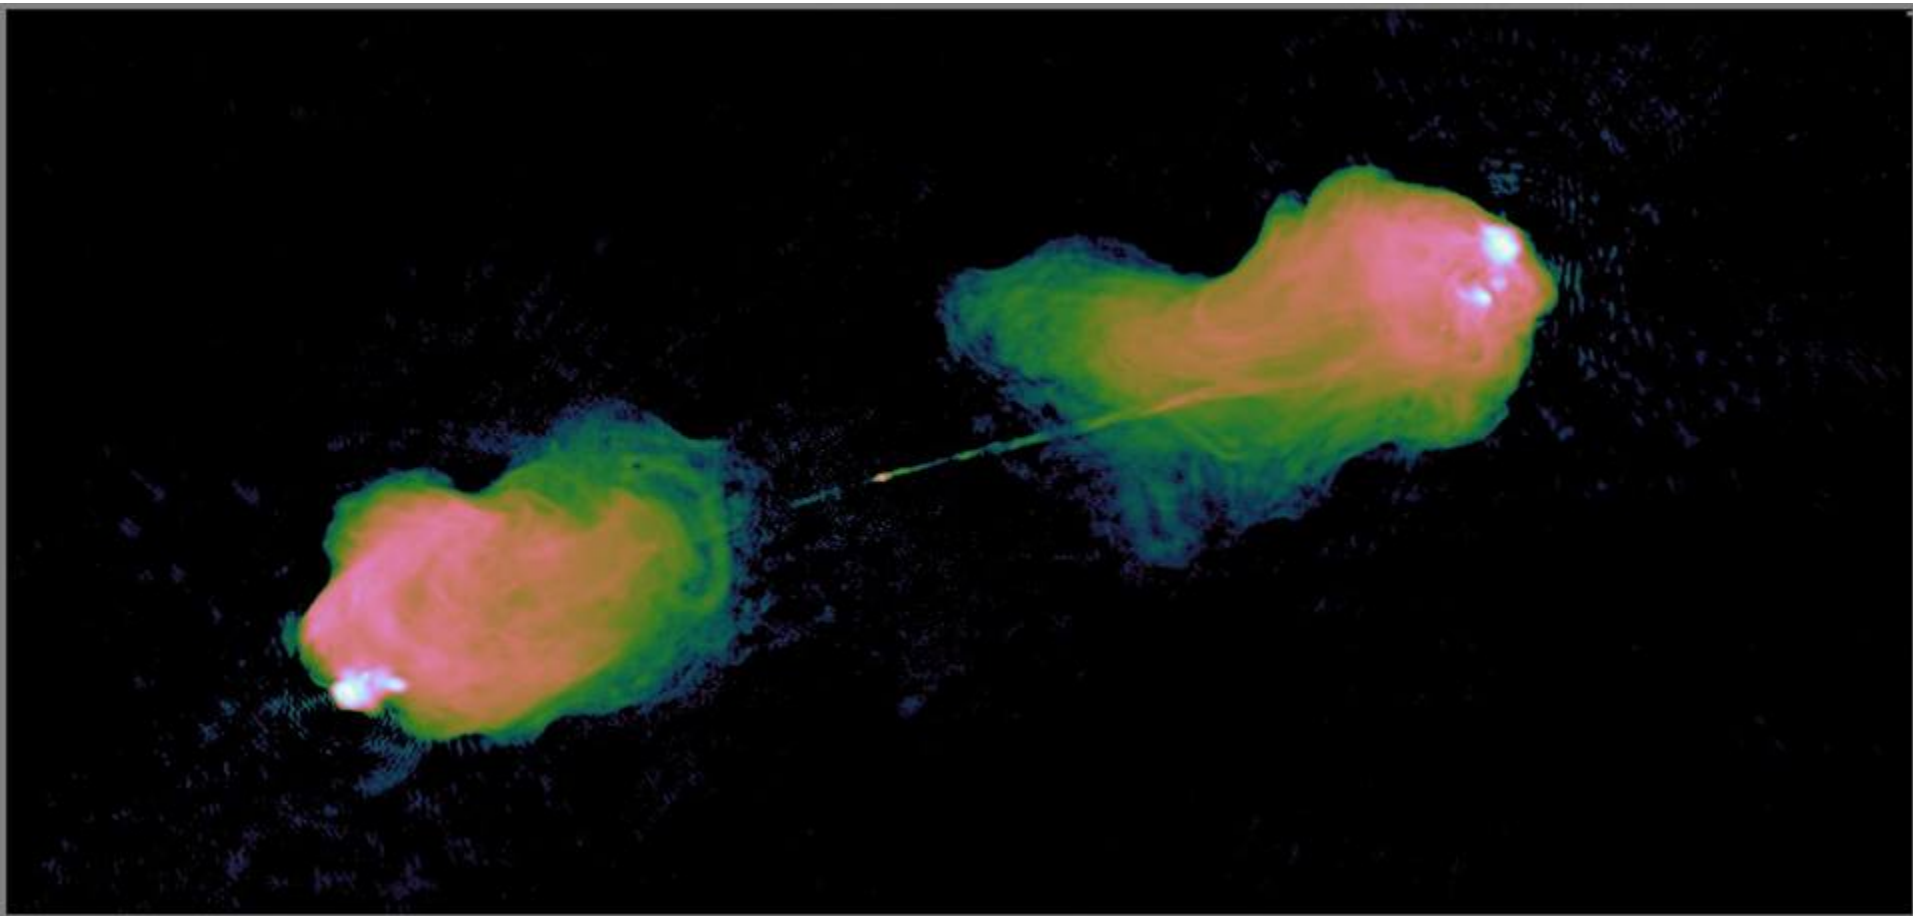
**Strong industry involvement:**  
Joint ventures/labs & self-funded contributors include

- IBM
- Seagate Systems (UK)
- Intel
- Nvidia

Industry contracts include:

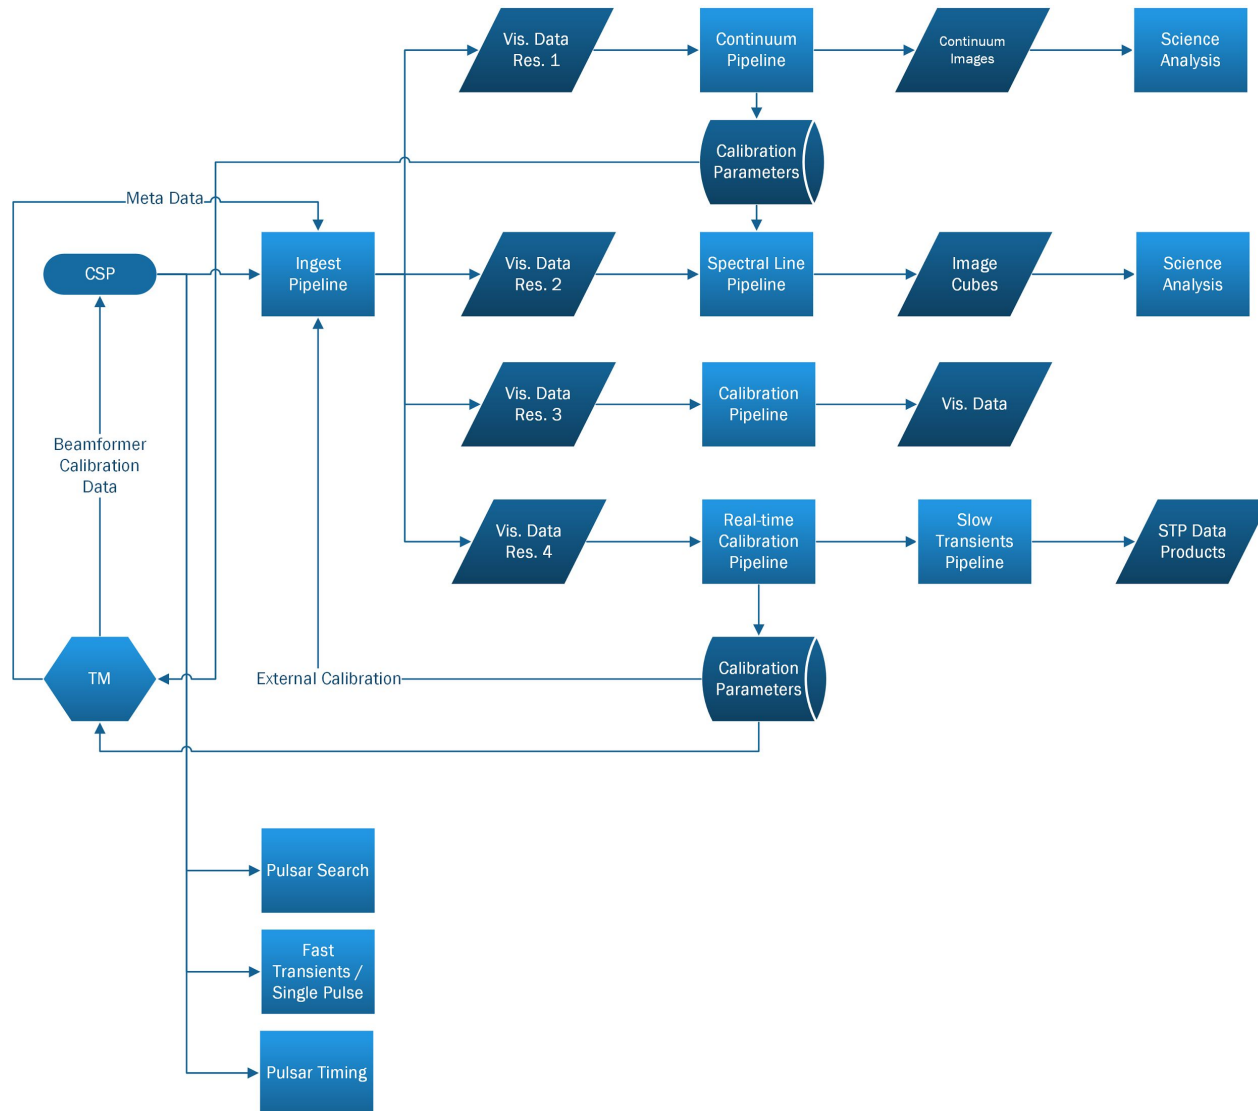
- Braam Consulting
- Nvidia
- Dell
- Mellanox
- Parallel Scientific
- Calsoft
- HPC Consultancy

# What we aim to produce



*Image credit: Oleg Smirnov (SKA SA/Rhodes Univ) / Rick Perley (NRAO)*

# SDP Pipelines



# Layout with example rates



Telescope Manager

Science Data Processor

SDP Local Monitor & Control

Data Processor

High Performance

~100 PetaFLOPS

Data Intensive

~100 PetaBytes/observation  
(job)

Partially real-time

~10s response time

Partially iterative

~10 iterations/job (~3 hour)

Archive

High Volume &  
High Growth Rate

~100 PetaByte/year

Infrequent Access

~few times/year max

Delivery  
System

Data Distribution

~100 PetaByte/year  
from Cape Town &  
Perth to rest of World

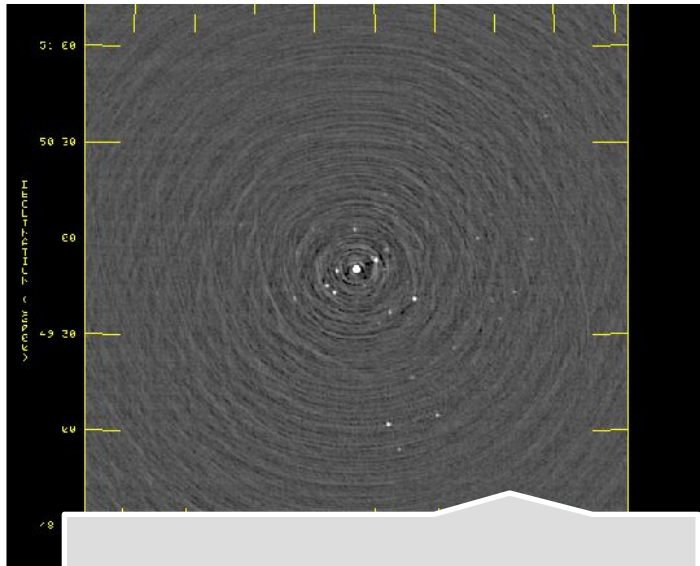
Data Discovery

Visualisation of 100k  
by 100k by 100k  
voxel cubes

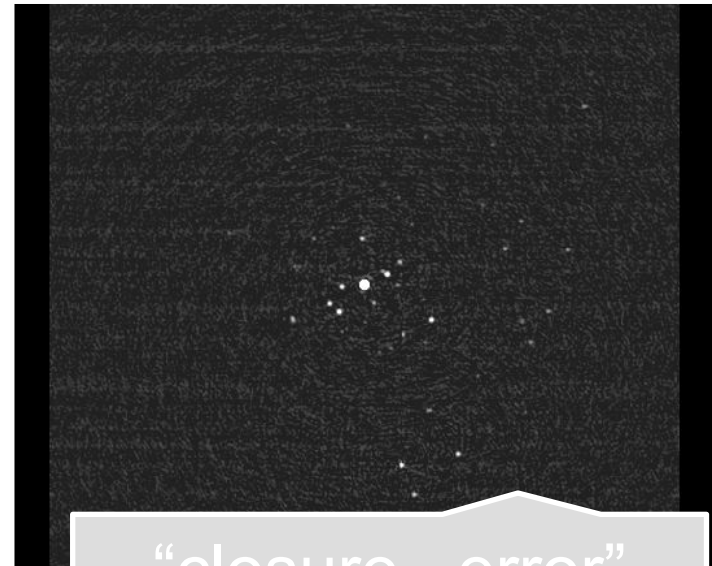
C  
S  
P

Regional Centres & Astronomers

# Iterative solution essential



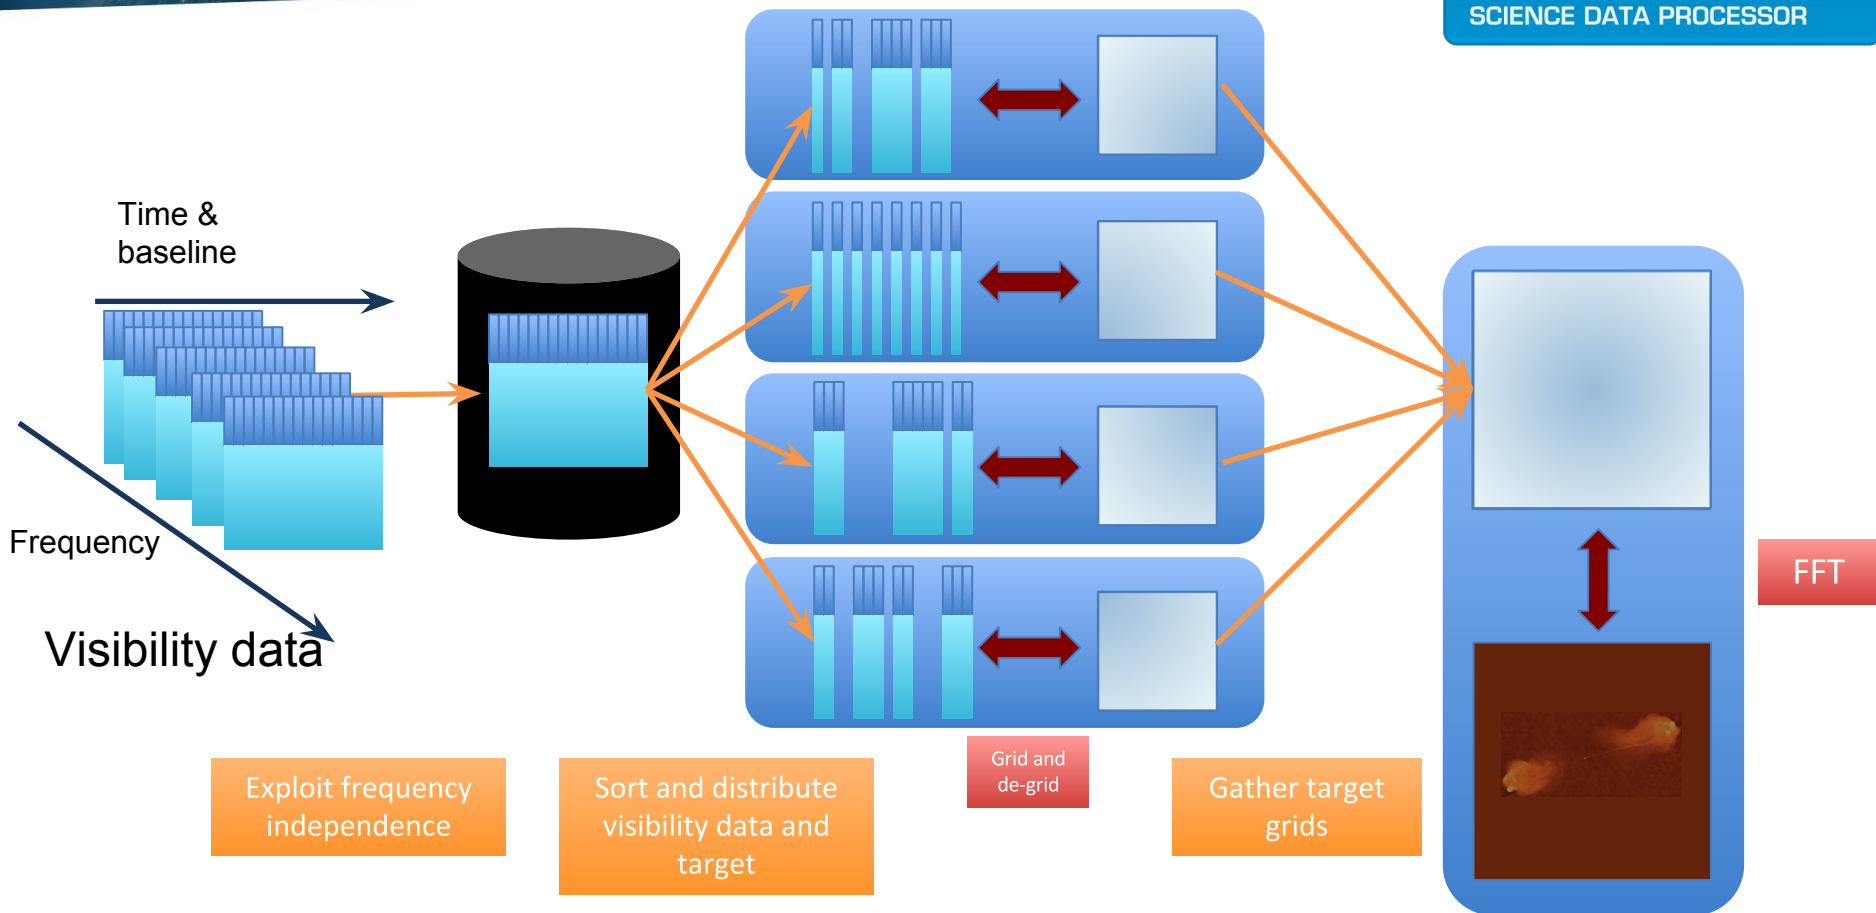
“Self-Calibration”



“closure –error”  
calibration

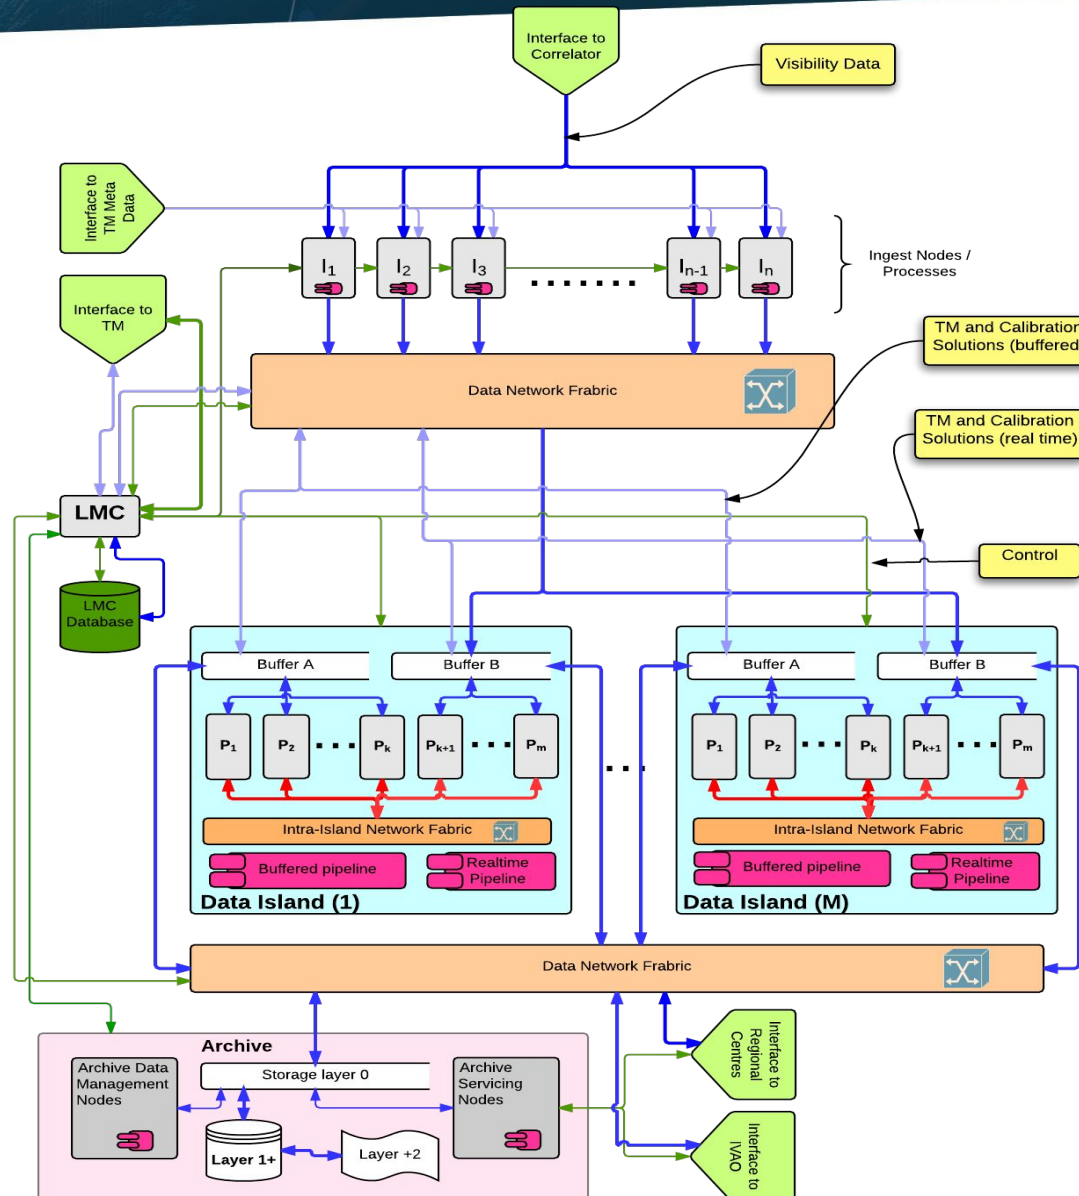
Rick Perley & Oleg Smirnov: “High Dynamic Range Imaging”, [www.astron.nl/gerfeest/presentations/perley.pdf](http://www.astron.nl/gerfeest/presentations/perley.pdf)

# Approach: Exploit data parallelism



- Data parallelism: Dominated by frequency. Provides dominant scaling  
Further data parallelism in locality in UVW-space
- Use to balance memory bandwidth per node
- Some overlap regions on target grids needed
- UV data buffered either on a locally shared object store or locally on each node

# Architectural view of data and information flow



## Hardware

- Power efficiency;
- Getting the right mix of storage processing and networking capabilities;
- Evolving hardware needs to support reasonably standard programming model and programming environment (cannot keep rewriting all the software!)

## Software

- Achieving high computational efficiency: op-ex
- Dealing with failures
- Adaptability to future system and node architectures
- Minimise development & delivery risks
- Maintenance and Enhancement over planned 50yrs observatory lifetime.

- Our initial requirements analysis based on current L1 is done, and our L2 requirements decomposed
  - Major assumptions and gaps identified and documented
- Level 1 and 2 functional decomposition complete
- Parametric model developed
- Preliminary data-driven architecture developed
- Cost model developed linked to performance requirements
- Initial pipeline analysis complete
- Hardware and system prototyping
  - Tests on existing HPC platforms
  - Analysis of File systems / Object store
  - Informal and formal review by industry partners
  - Initial kernel performance analysis/reports from industry partners
- RBS deliverables completed on time
- PDR Submission

## Quotes from the report:

- Uncertainties in the L1 requirements are a major risk for the SDP element system-level integration and commissioning process,
- SDP design and costing was performed (by direction from the project) without consideration of commissioning
- Given the current level of maturity of the SDP design, the panel is concerned that the presented schedule leading to a CDR in 2016 is exceedingly aggressive.
- Given, SDP's "downstream" position in the overall SKA architecture and data flow, it is natural that SDP's schedule be slightly offset and later than the other SKA element
- Inter-system interface definition is incomplete
- Establishing a proof of concept baseline, and sufficient prototyping to validate the architecture is required
- A baseline for technology choices needs to be established, the prototyping plan reviewed to include TRL assessment methods



- Established:
  - Overall requirements, system interfaces
  - System decomposition
- Provisional
  - Sub-system requirements, internal interfaces
- Next steps
  - More detailed architecture (**being advanced this week!**)
  - Identifying baseline technologies for all sub-systems
  - Identifying stable APIs and technologies for verification before commencement of construction.

# Stage-2 milestones



... taking on board PDR observations

Milestone	Description	Due Date
	Close out Stage 1	May 2015
M12	Design for prototyping	31 <sup>st</sup> July 2015
M13	Prototyping internal CoDR	30 <sup>th</sup> November 2015
M14	Delta-PDR and Requirements Review	25 <sup>th</sup> March 2016
M15	Prototyping and Risk Report and Review 1	24 <sup>th</sup> June 2016
M16	Prototyping and Risk Report and Review 2	23 <sup>rd</sup> September 2016
M17	Design maturity review [External/SKAO Review]	16 <sup>th</sup> December 2016
M18	Prototyping and Risk Report and Review 3	27 <sup>th</sup> January 2017
M19	Prototyping and Risk Report and Review 4	29 <sup>th</sup> September 2017
M20	CDR Submission [External Review]	15 <sup>th</sup> December 2017
M21	Close out Stage 2	23 <sup>rd</sup> March 2018

## M12 – Design for prototyping

- described the SDP product tree to SDP Level 3;
- contain an initial risk assessment for every element of the product tree;
- show a selection matrix capturing options considered/prototyped;
- demonstrate how Technical Readiness Levels will be used to retire risk;
- present an analysis supporting an overall architectural choice;
  - A look at Compute Islands
  - Similarities/differences with industry analytics platforms
  - The Data Flow approach
- show planning to ensure prototyping coverage of the architectural choice.

### Important changes

1. Changed product tree to make internal interfaces simpler.
2. Developed a systematic approach to capture implementation considerations and options for each product (or grouping).
3. Introduced Technology Readiness Levels.

# Current work: populate products + plan

## 2. Each product has corresponding Confluence page based on template

### Summary

JIRA ticket	<a href="#">PT-17 - C.1.1.1 Compute Island</a> IN PROGRESS
Owner	@Andrew Ensor
Developers	@Shaohua Wu
TRL	
Summary	

- Summary
- Links to Documents
  - Tickets
- 1.0 Introduction
  - 1.1 Background and Strategic Fit
  - 1.2 Context
  - 1.3 Behaviour
  - 1.4 Interfaces
    - 1.4.1 Logical
    - 1.4.2 Physical Interfaces
    - 1.4.3 Software Interfaces
- 2.0 Requirements
  - 2.1 Performance Requirements
  - 2.2 Functional Requirements
  - 2.3 Cost
  - 2.4 Schedule
- 3.0 Select Candidate Solution
  - 3.1 Architectural Drivers
  - 3.2 Candidate Solutions
  - 3.3 Concept Selection Table
  - 3.4 Risk Assessment Table
  - 3.5 Select Preferred Option(s)
- 4.0 Critical Technology Element Selection
- 5.0 Technology Readiness Level Assessment
- 6.0 List of TBDs

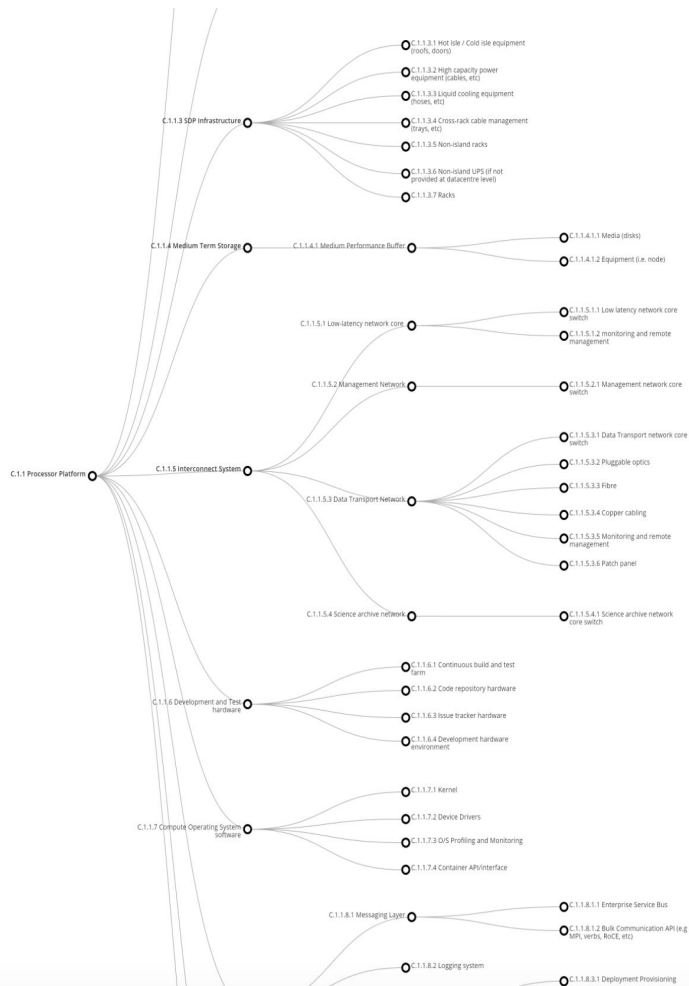
### Links to Documents

(Insert links to any relevant Google Docs, SysML models, items in Configuration management tools etc.)

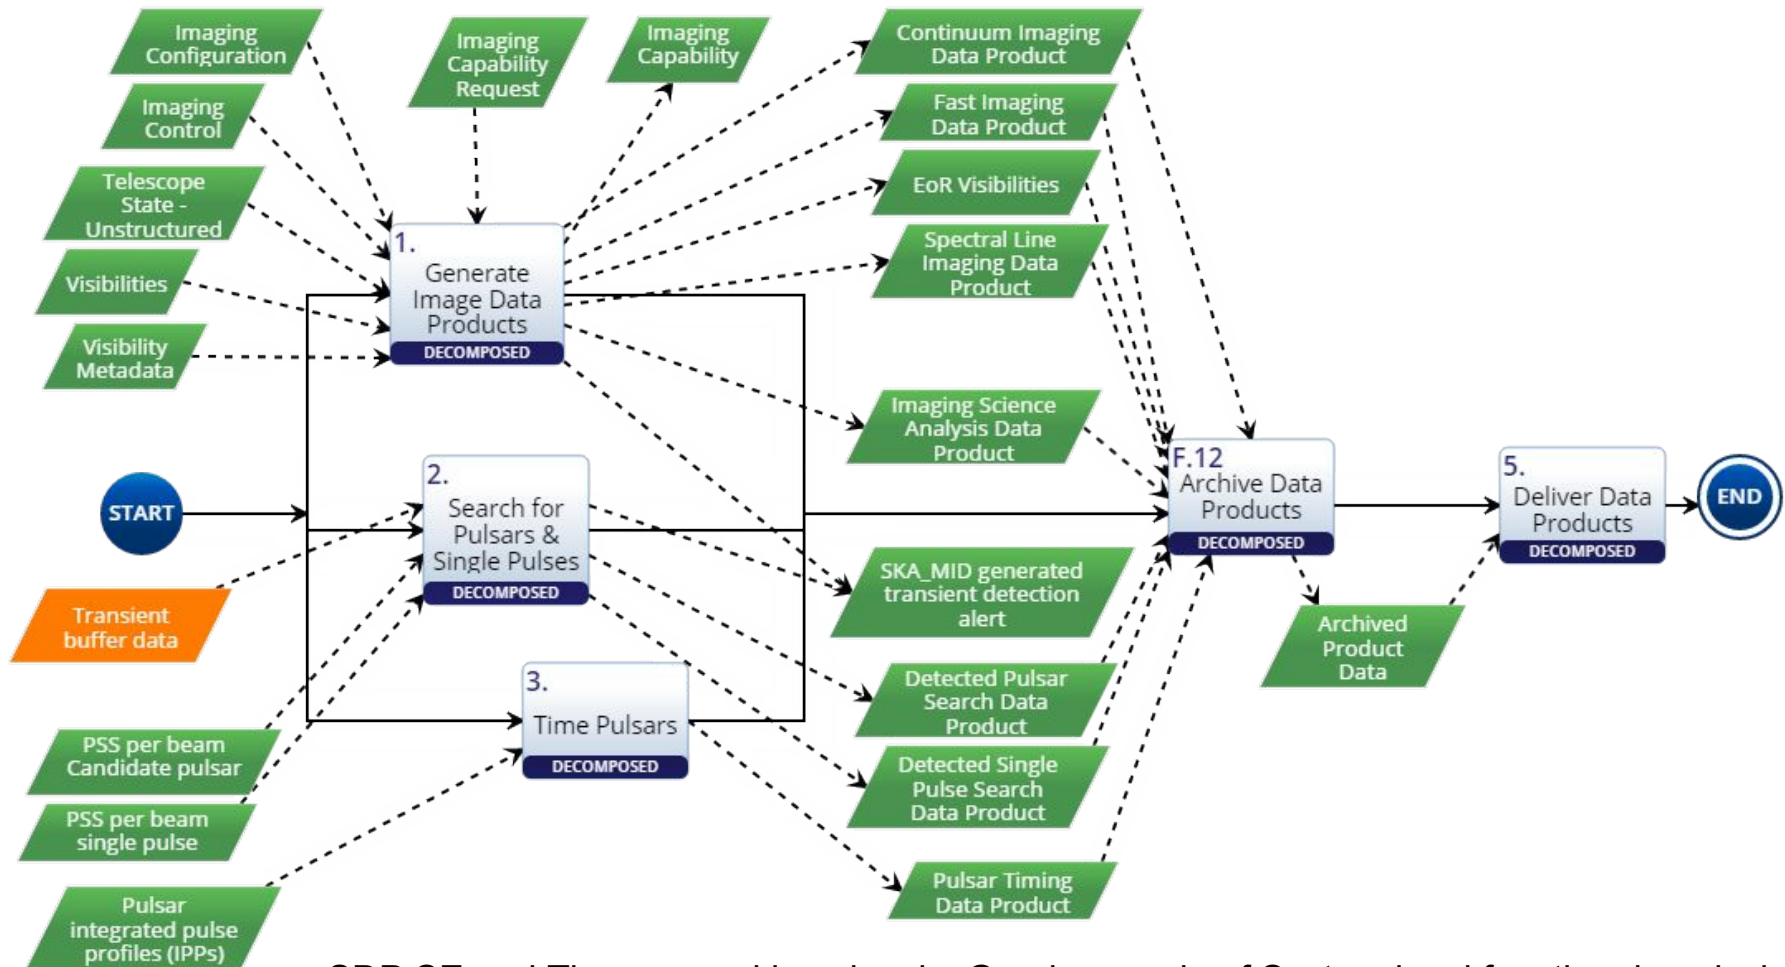
- ✓ C.1.1.1.1 Management Compute Island
  - C.1.1.1.1.1 Compute Island Management Network Switch
  - C.1.1.1.1.2 Compute Island Master Node
  - C.1.1.1.1.3 Compute Island Service Node
  - C.1.1.1.1.4 Compute Island Software Repository
- ✓ C.1.1.1.2 Compute Node
  - C.1.1.1.2.1 Latency Optimised Cores
  - C.1.1.1.2.2 Main Memory \*
  - C.1.1.1.2.3 Throughput Optimised Cores
  - C.1.1.1.2.4 Bulk Data Transport Network \*
  - C.1.1.1.2.5 Low Latency Network HCA \*
  - C.1.1.1.2.6 Archive Network NIC \*
  - C.1.1.1.2.7 Out-of-Band Management Unit (BMC) \*
- ✓ C.1.1.1.3 Compute Island Bulk Data Network
  - C.1.1.1.3.1 Compute Island Bulk Data Network Switch
- ✓ C.1.1.1.4 Compute Island Low Latency Network
  - C.1.1.1.4.1 Compute Island Low Latency Network Switch
- ✓ C.1.1.1.5 Compute Island Science Archive Network
  - C.1.1.1.5.1 Compute Island Science Archive Network Switch \*
- ✓ C.1.1.1.6 Cabling \*

## 3. Prototyping plan<sub>17</sub>

## 1. Product tree developed



# Current work: Allocate functions



SDP SE and Thomas working closely. Good example of System level functional analysis working well with SDP functional analysis .... Colour coded items added from System level SE work. Telescope team work revealed new function (in orange)

# Some issues & what SDP needs to make progress



- Internally SDP not currently benefitting from all resource. SDP is a broad and diverse consortium and working effectively across all areas remains a high priority.
- The power/cost budget constraints and outlook are a concern. The main cost to SDP is operational cost. In discussion with SKAO about total cost of ownership (power/upgrades/operations...) and what is right for the SKA.
- Understanding required reliability and maintainability. SDP is in need of realistic constraints to design against as this impacts costs.
- Telescope Model is a key element for SDP but is a work in progress. Being worked on now by SKAO. SDP needs to make sure it is fully resourced to respond.
- Some missing information (e.g. in requirements and calibration) creates a challenge. A more collaborative approach is being pursued. SDP will take a pragmatic internal view where information is missing.
- The detailed requirements for the pipelines requires improved science input. SDP needs a formal relationship with the science teams.



- In March SDP will have its delta-PDR.
- We need to firm up the SDP architecture and tie this more strongly into the system level engineering (including requirements).
- Internal and external interfaces will be better elucidated.
- We will address questions of reuse from other projects at all levels.
- Our Product Tree analysis work will be continued to identify and reduce the highest risk areas through clear and prioritised prototyping activities that are better compartmentalised (to improve overall engagement).
- SDP will enter an intensive prototyping period.
- Where appropriate (to address high-risk areas) we will be taking forward multiple prototyping options in the same area.



- The project boundaries, in terms of data layers and distribution have not been clear.
- The SKA Board has set up a Data Flow panel, which has a good representation from SDP, to help define the boundaries of the SKA-SDP element properly.
- What is coming into or moving out of SDP scope (to Regional Centres) may not be certain until after March. Therefore SDP may have to make assumptions for its delta-PDR.

# Other things we need to think about



- How should SDP support (data) Quality Assurance?
- Requirements on SDP during commissioning were previously excluded from SDP scope but need to be developed. SDP will then have to consider how to provide some early functionality to fit with commissioning needs.
- The SDP Consortium was set up to take the design to CDR and we need to give some input into developing the SDP parts of the procurement models. For example, can the SDP be procured / delivered in parts? What delivery phasing is required to support the AIV plans?

# What SDP will deliver for CDR



- We need to decide on what is an appropriate CDR documentation set for an ICT project of the size and nature of SDP (with a software focus and hardware elements).
- Following the SDP PDR it was agreed that the Consortium would produce an 'end-to-end prototype' by CDR. Prototyping the system/design rather than just prototyping in support of the design process is a large and fundamental change. We need to define the scope of this prototype appropriately and consider options for reuse of software where possible (but we note that no other telescope currently does automated Quality Assurance).
- One deliverable will be the construction plan

- Within SDP there is good internal and external engagement. Work is progressing at many levels - including to clarify terminology in the high-level architecture.
- A lot of material was produced for PDR which is now being consolidated and enclosed in a more systematic Product Tree analysis and Prototyping Plan
- The Consortium is taking a pragmatic approach to a number of open issues and working closely with SKAO to clarify scope.