



# AENEAS, UK focus

Rosie Bolton  
Anna Scaife  
Michael Wise



The University of Manchester

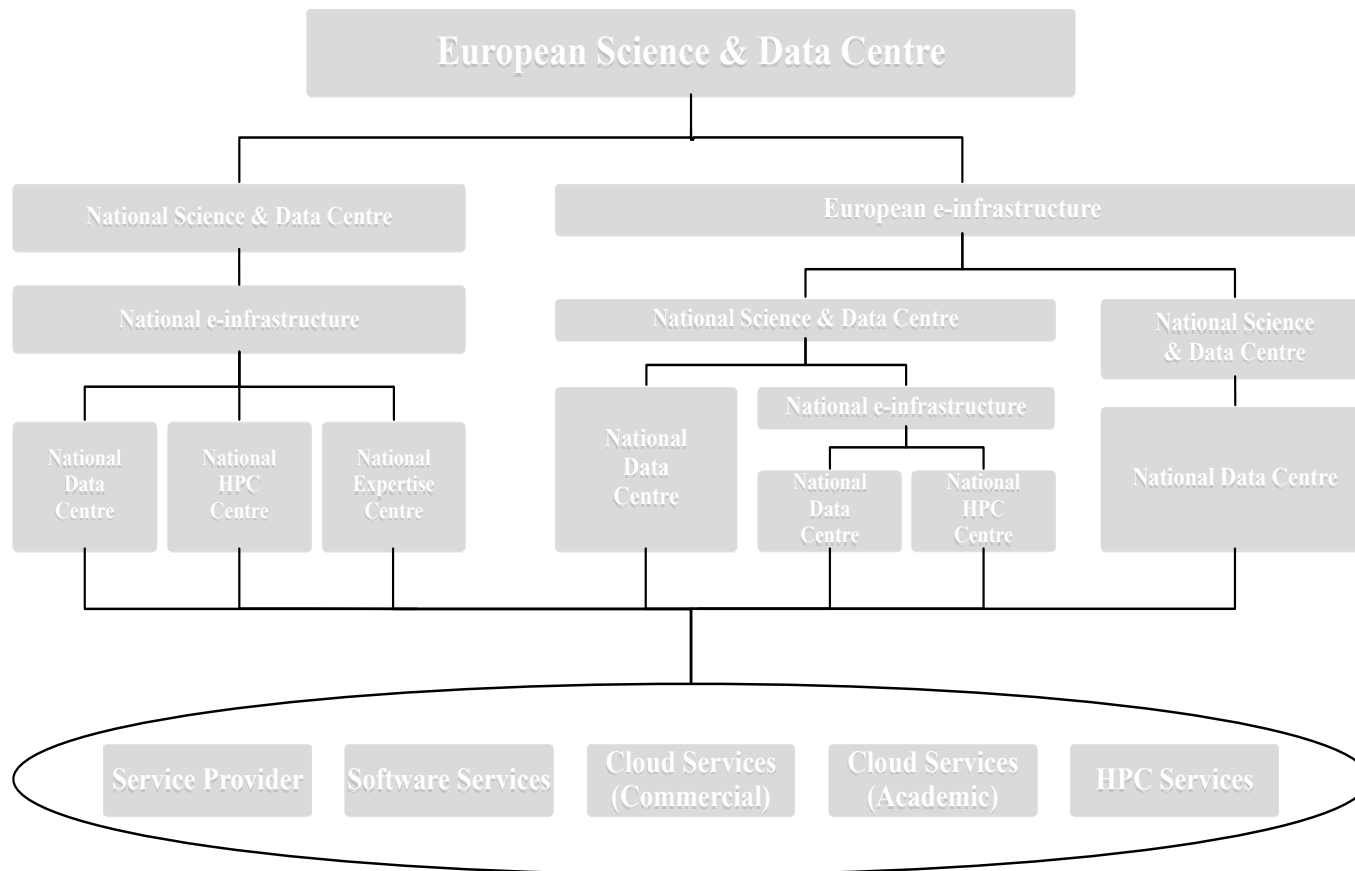


UNIVERSITY OF  
CAMBRIDGE

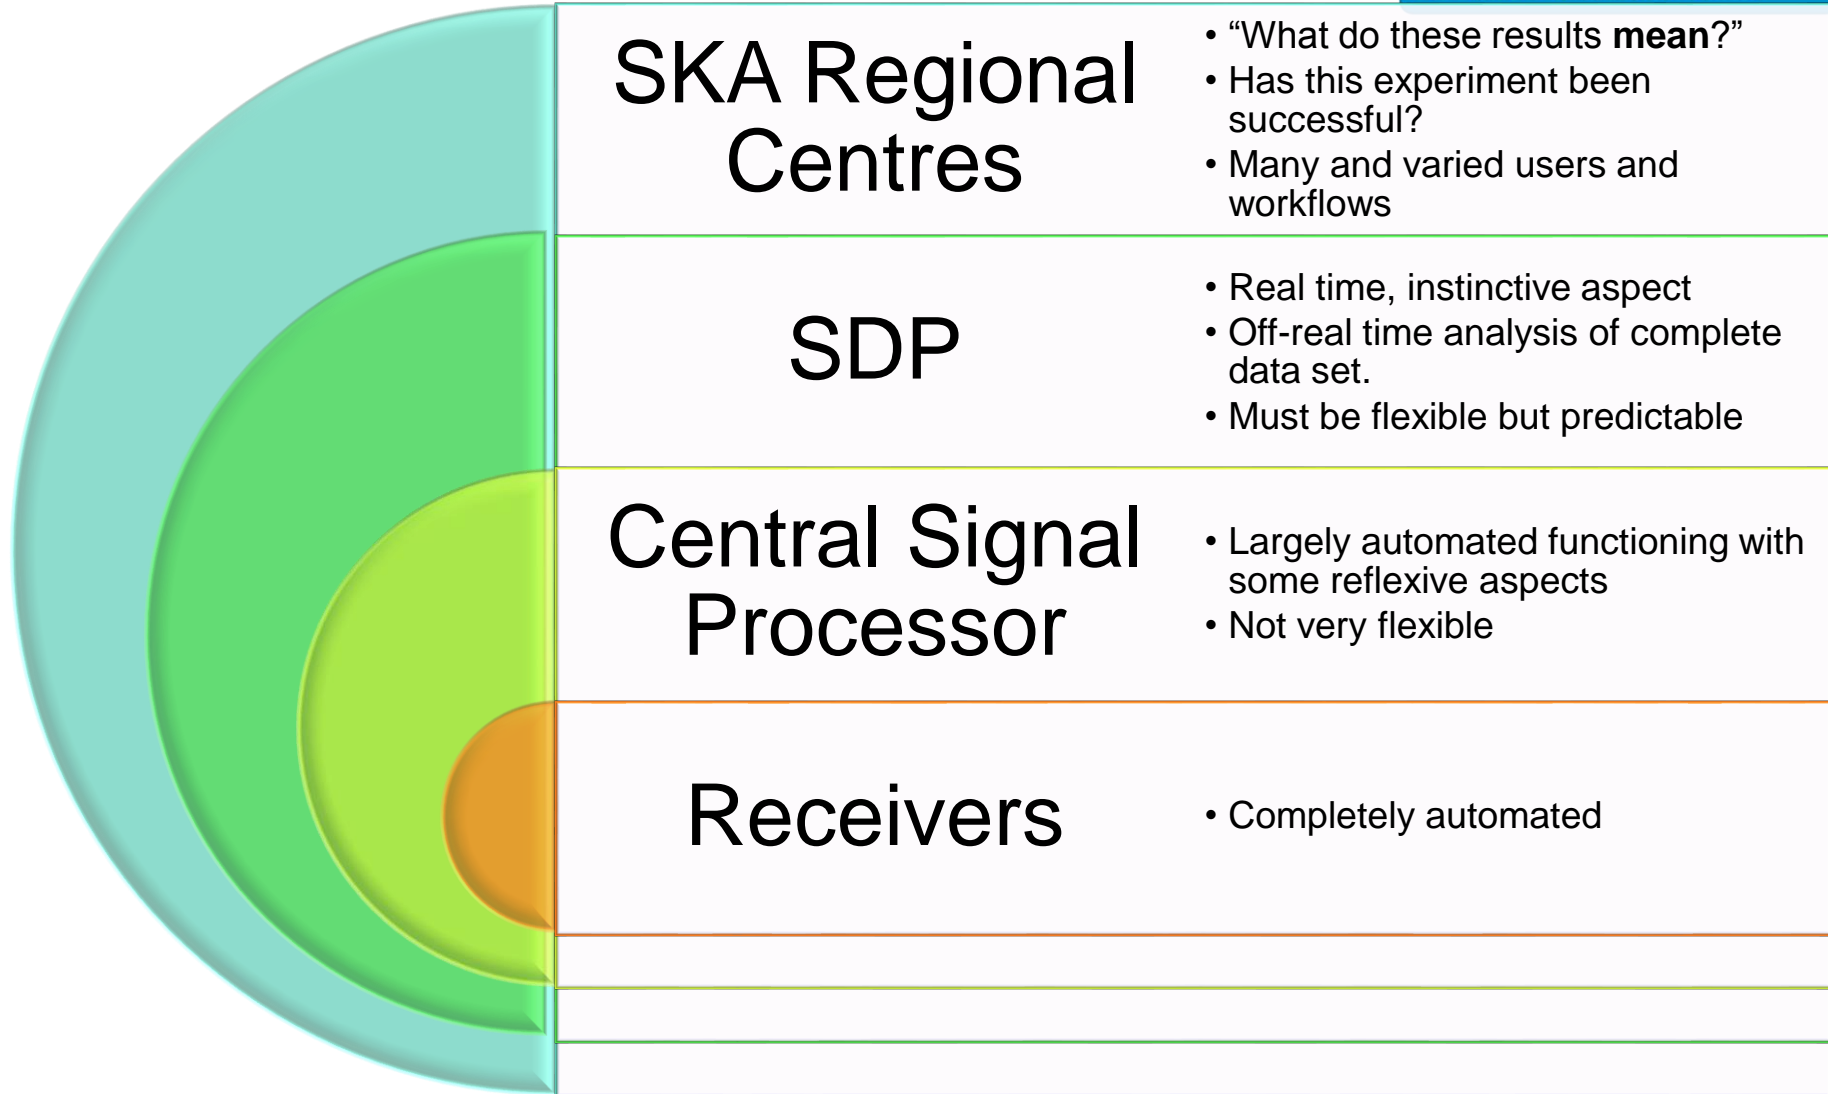


- Develop a concept and design for a distributed, federated European Science Data Centre (ESDC) to support the astronomical community in achieving the scientific goals of the SKA
- Include the functionality required by the scientific community to enable the extraction of SKA science
- Integrate the necessary underlying infrastructure not currently provided as part of the SKA Observatory to support that extraction

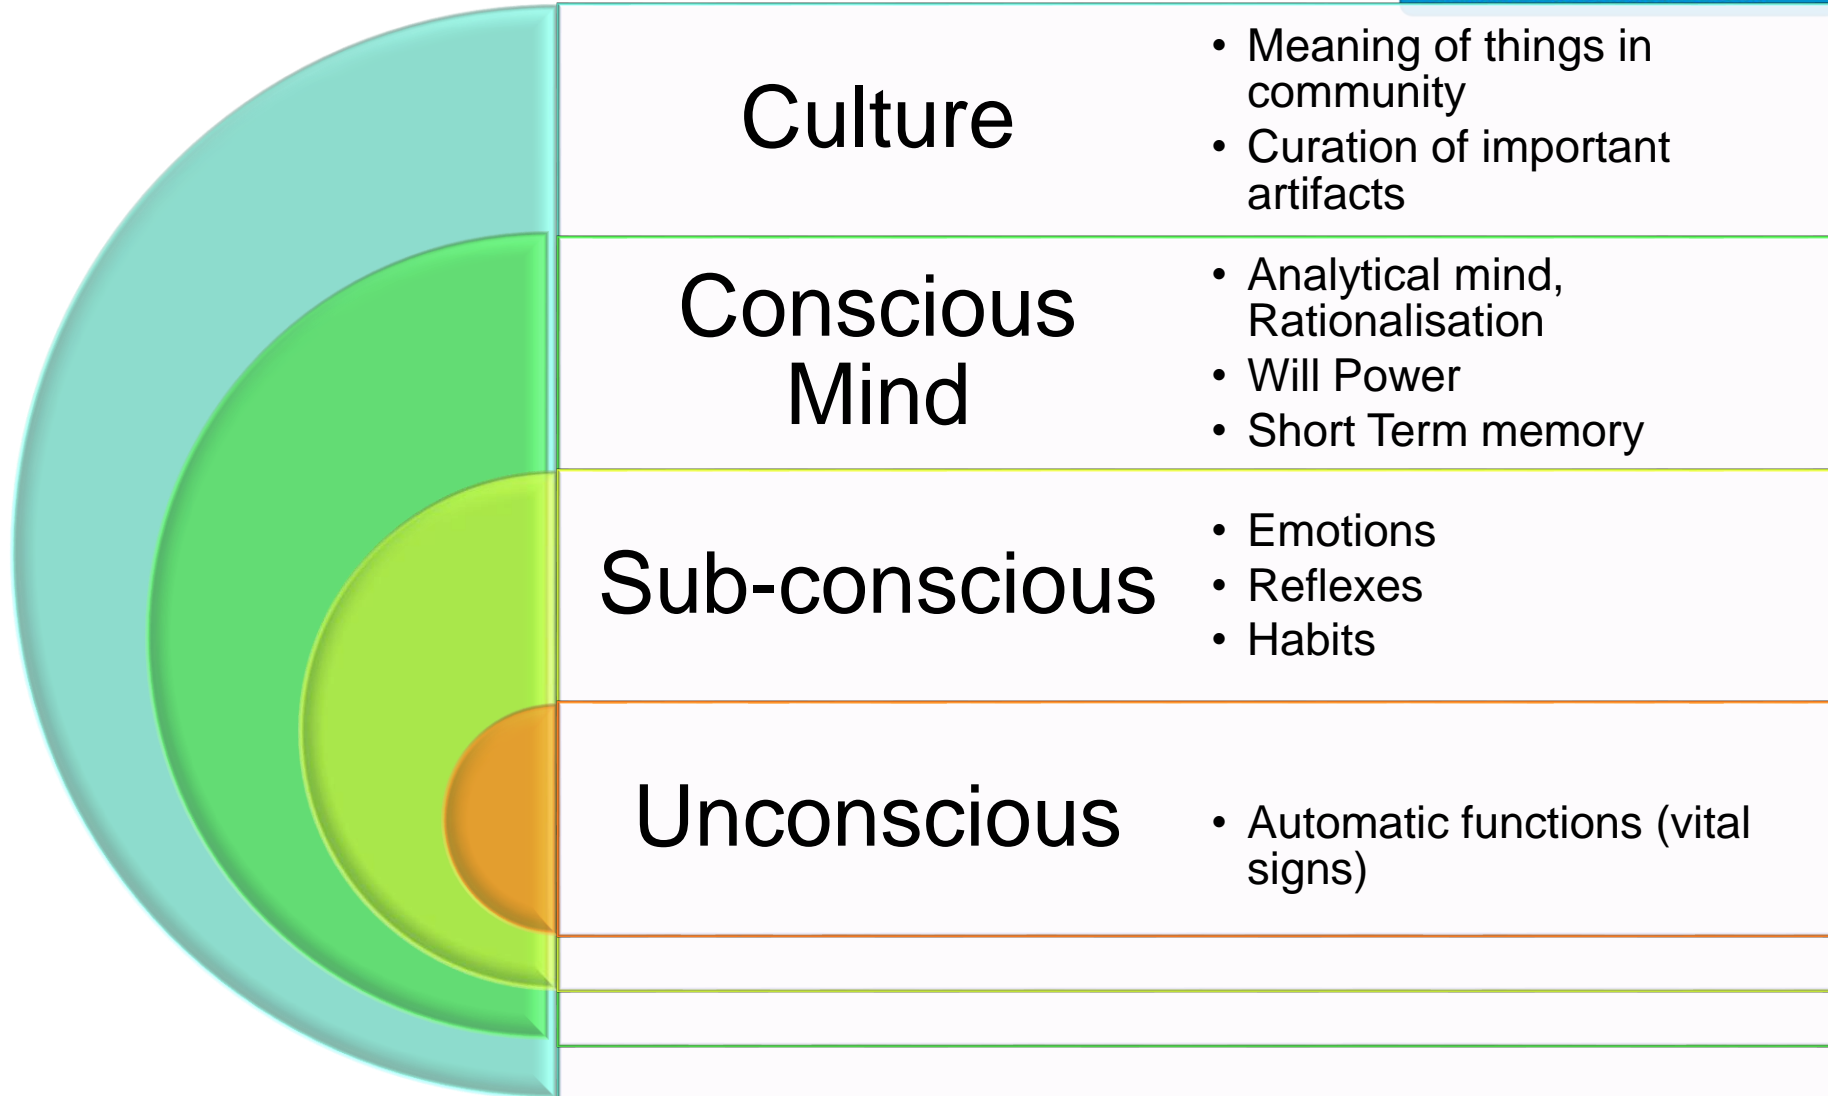
# AENEAS – ESDC view



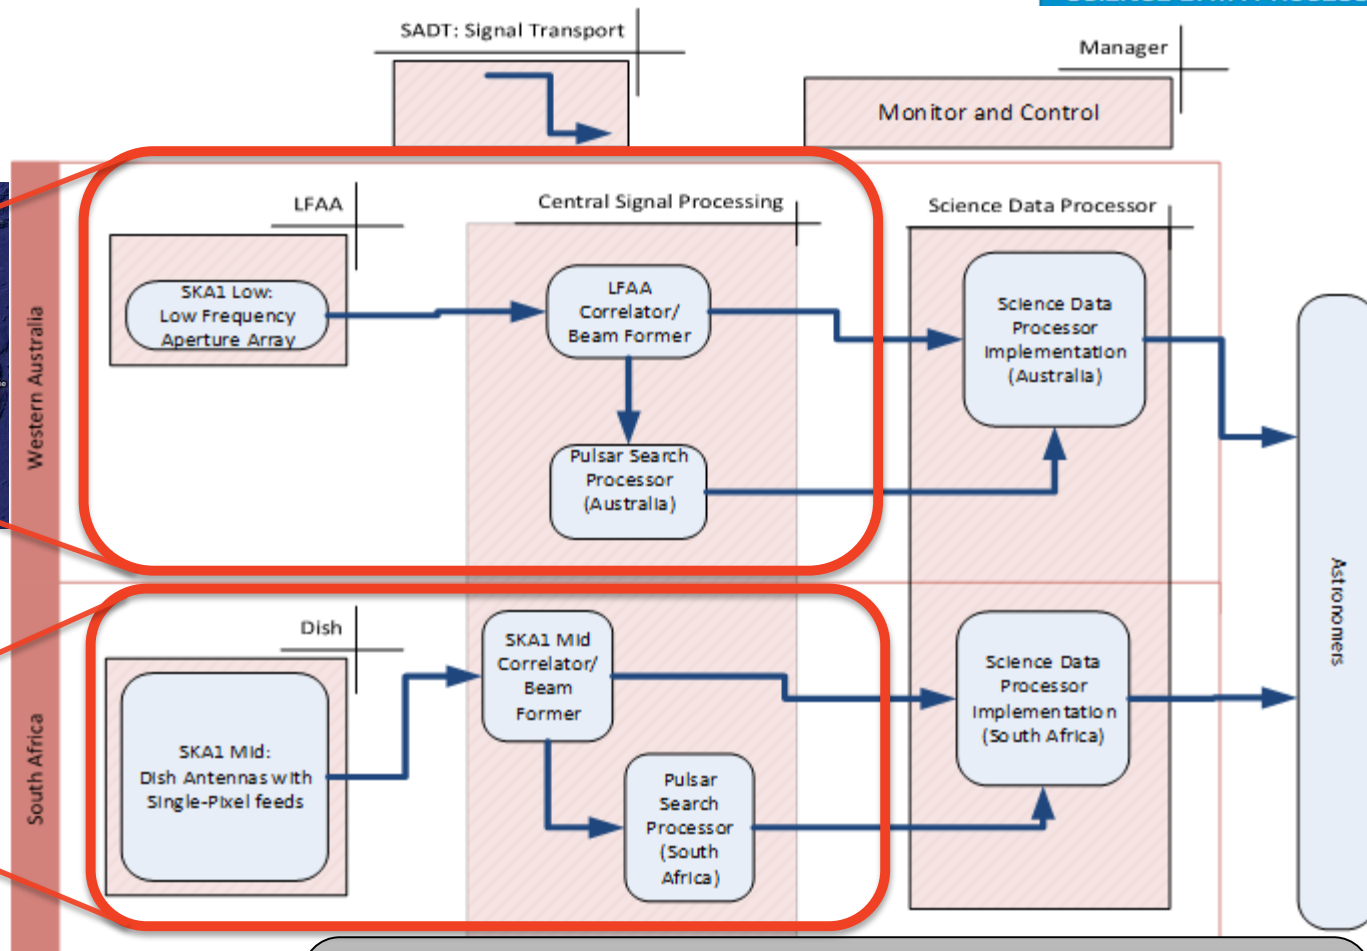
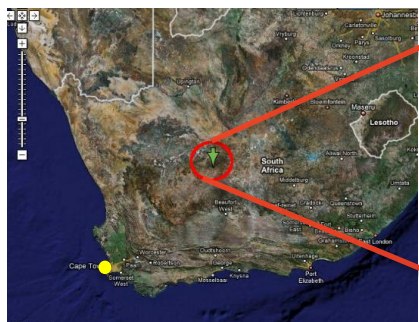

# SKA-SDP in context: What is the Science Data Processor?



# SKA-SDP in context: What is the Science Data Processor?



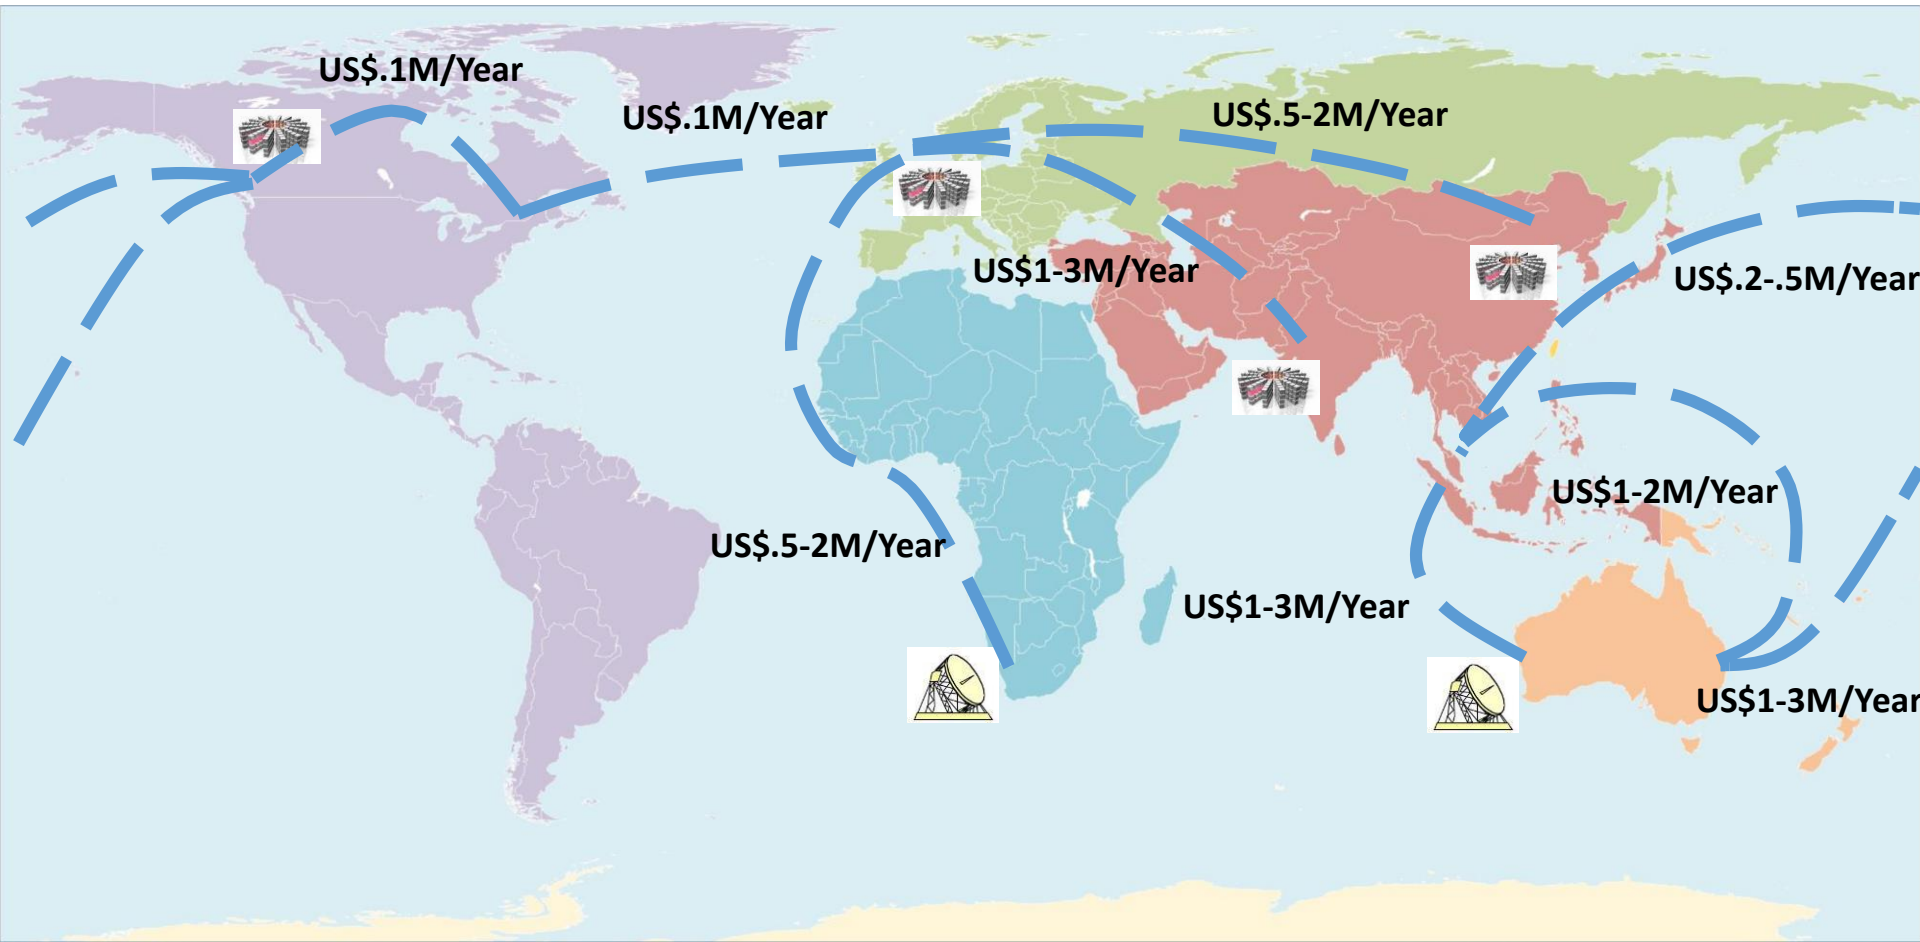
# SKA-SDP Context Diagram



SDP facilities are off-site! (In Perth & Cape Town). Cannot be in the desert. Dedicated HPC facilities.

# And Finally Regional Centre Network

- 10 year IRU per 100Gbit circuit 2020-2030
- Guesstimate of Regional Centres locations



# Telescope Manager

## Science Data Processor

### SDP Local Monitoring & Control

#### Data Processor

##### High Performance

- ~400 PetaFLOPS

##### Data Intensive

- ~30 PetaByte/observation (job)

##### Partially real-time

- ~10s response time

##### Partially iterative

- ~10 iterations/job (~6hour)

#### Data Preservation

##### High Volume & High Growth Rate

- ~100-300 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

Observatory & Regional Centres

CSA

2 x 400 Gbyte/s

2 x 10 Gbyte/s

2 x 10 Gbyte/s

NB Numbers are approximate!



# SDP Challenges:



Input

- 400 Gbyte/s INGEST

Distribute

- 400 million tasks in the graph
- Half an Exaflop/s total peak

Generate  
(and destroy)

- 1.3 ZettaBytes intermediate data products


Preserve and  
ship

- 1 PetaByte per day of Science Data Products

# SKA Science Working groups




- **Astrobiology (“The Cradle of Life”)**
- **Continuum**
- **Cosmic Magnetism**
- **Cosmology**
- **Epoch of Reionisation & the Cosmic Dawn**
- **HI Galaxies**
- **Pulsars (“Strong field tests of gravity”)**
- **Transients**



# AENEAS



- **A whole range of different scientific applications and techniques.**
- **Definition of the Key Science Projects is underway**
- **KSPs will make up the bulk (~70%) of time allocated to SKA**
- **Scientists involved in KSPs will interact with their data via Regional Centres**
  - What are the workflows?
  - How much processing is needed?
  - What balance between processing, memory, and long term preservation / archive functionality?
  - How is data transfer likely to work best?
  - What simulation / modelling work is required in the SRCs?
  - How can we manage federated access, authorisation and bookkeeping?
  - What should the policies be?
  - How do these use cases fit into national facility planning roadmaps in Europe?



# ANEAS UK: This meeting



- **UK AENEAS members**
  - UCAM, UMAN, STFC, GEANT
  - SKAO (unfunded)
- **UK contribution to**
  - WP2 “Development of ESDC Governance Structure and Business Models”
  - WP3 “Computing Requirements”
  - WP4 “Analysis of Global SKA Data Transport and Optimal European Storage Topologies”
  - WP5 “WP5: Access and Knowledge Creation”



UNIVERSITY OF  
CAMBRIDGE



MANCHESTER  
1824



Science & Technology  
Facilities Council

The University of Manchester

# Aims of this meeting (Rosie's views)



- **Better understanding of the existing expertise in academic HPC within the UK**

Oh, you guys faced and solved these problems already?

- **How can we collaborate effectively, without duplicating effort or missing opportunities to utilise existing work?**

Where can I get access to this detailed knowledge? Who can help me? We need an environment where it's OK to ask questions and with some shared resources.

- **Do we need a formal collaboration?**

Does this help with funding and reporting?

- **Better understanding of SKA / GridPP / HPC language**

Where's the dictionary?

- **Where is funding coming from (and how many FTEs, and who) to progress this work over the next decade? What is the UK view?**

# AENEAS WP3 details

(Thursday)

3.1 & 3.2

- ***Inventory of SKA science cases and post-SDP computing and data storage requirements***

3.3

- ***Evaluation of existing HPC, cloud and distributed computing technologies***

3.4

- ***Design and costing for distributed ESDC computing architecture***

3.5

- ***Requirements for interfaces to SKA Science Archives & Other Repositories***

3.6

- ***Validation, Verification & Proof of concept activities utilizing SKA pathfinder and pre-cursor facilities***



## ***Inventory of SKA science cases and post-SDP computing and data storage requirements***

The SKA has developed a list of 13 High-Priority Science Objectives (HPSOs) which are being used to generate survey strategies for the SKA in its first several years of observations. These are large projects with many thousands of observing hours each. Since these will be made up of tens to hundreds of separate data sets substantial processing and manipulation of the SKA data products will be required in the regional data centres to deliver the survey science at anticipated fidelity. This task will focus specifically on the delivery of these key experiments and their compute processing requirements and data storage and access requirements, and provide a basis on which to proceed with the sizing and costing efforts.

Once these large surveys are complete, enormous benefits will be available if we can combine data from other observatories (e.g. LSST, Euclid). Using results and insights from the Asterics programme we will make estimates for ESDC resources needed to support these efforts and maximise scientific return on the ESFRI astronomy projects.

Further work will investigate whether specific Science Use Cases (more representative of open time programmes) could have significantly different ESDC compute requirements. We will also consider the options for “Discovery Products” which would be generic products not covered by specific experiments, but piggy-backing on observing time.

The output of this task will be a series of system-sizing and functional requirements to appear in deliverable 3.1.



### ***Evaluation of existing HPC, cloud and distributed computing technologies***

This task will enumerate the key elements needed for the software infrastructure required for the ESDC, and evaluate options for fulfilling them, these include:

1) **Middleware** – i.e. infrastructure to support distributed compute models within, for example a cloud-like environment although also including HPC facilities. Different software products and middleware solutions for allowing access to distributed computing facilities and capabilities will be analysed and compared to the data analysis requirements collected in Tasks T3.1 and T3.2. This will include products and solutions for compute (cloud compute, HTC (High Throughput Computing), HPC (High Performance Computing) and container-based cloud compute).

**This will include an evaluation of OpenStack**, and other cloud middleware stacks with the aim of ensuring portability of data and applications in the distributed environment to be implemented in the ESDC. It will also include an analysis of available replica management and data transport organisation tools such as PHEDEX & PANDA.

2) **Elements required for a federated ESDC**, including services. The ESDC must provide resources to users in a way which combines many different computing resources but presents these in a harmonized way to each user, and which can validate users' requests for data access and keep accounts of computing and storage resources for each user or use group, while avoiding un-necessary data movement between sites. Some of this federation functionality may be present in a selected middleware layer, but other aspects may not and these must be considered: for example Authentication and Authorization Infrastructure (AAI), efficient movement of data based on policies; integration with HPC software stacks so that (as needed) the ESDC is able to utilise HPC resources for processing; accounting elements and ensuring proper and fair use of resources.

3) **The top-level software stack providing an environment for efficient distributed analysis of data** – the task should certainly look at the possibility of building on top of industry-standard Big Data / data-science stacks such as Spark (now beginning to take over from Hadoop).

The deliverable of this task will be D3.2: "Report on suggested solutions to address each of the key software areas associated with running a distributed ESDC", which will include a list of options, each assessed for suitability.

### ***Design and costing for distributed ESDC computing architecture***

Based on input from the evaluations in D3.1 and D3.2 this task will provide a top-level architecture and functional design for the ESDC.

To proceed we will make use of the inventory of national roadmaps from WP2 (D 2.1) and determine the potential for incorporation / co-use of existing or planned facilities to achieve economies of scale.

We will develop a costing of additional resources needed (over and above existing facilities) to bring about a functioning ESDC, considering the full SKA observatory lifecycle from commissioning as the SKA is built and well into full operations as the SKA observatory develops and undergoes upgrade cycles.

The outputs of this task will be a 1) A preliminary system sizing estimate (D3.3) and 2) a documented design for a ESDC model



### ***Requirements for interfaces to SKA Science Archives & Other Repositories***

The work in this task includes the assessment of existing policies for interactions between science facilities and data centres, incorporating an evaluation of policy items with respect to their technical applicability in the SKA case, as well as a gap analysis for SKA-specific needs. This technical assessment will feed into more general policy recommendations.

The ESDC will incorporate multiple interfaces, both functional and digital (data IO). This task will assess the requirements both for ensuring controlled and managed ingest of data across these interfaces, and of the subsequent storage strategy. This requires an assessment of existing data-moving tools and protocols (commercial and academic – for example WLCG (Worldwide LHC Computing Grid), their compatibility with an ESDC architecture, as well as verified assessments of data ingest from global sources including the SKA Science Data Processors, other nodes within the ESDC, and external archives (e.g. LSST, EUCLID, JWST etc). It will also ensure the compatibility of recommended ESDC standards with the widely used VO standards. In doing this it will also form recommendations on minimum meta-data requirements for ESDC-held data, inline with analyses from the ASTERICS project.

A major functional interface within an ESDC will be the mapping of user specifications onto data processing work flows, from ingest to delivery. This mapping should incorporate both a translation between user-defined parameters (data product specific) and processing parameters (function specific) as well as the impact of different parameter choices mapped onto different types of compute system (data access patterns, distribution of processing etc). This task will also inform policy decisions governing the persistence of user work-flows to enable reproducibility of results or regeneration of data.



### ***Validation, Verification & Proof of concept activities utilizing SKA pathfinder and pre-cursor facilities***

This task contains the technical work required to verifying the design recommendations developed for the ESDC in T3.1 - T3.5 using, where appropriate, data from precursor and pathfinder instruments.

The work in this task includes the provision of a standardised set of appropriate test data, incorporating output from existing facilities and pathfinder instruments. The task also includes the provision of prototype software blocks to verify and validate the functional requirements derived in T3.3, as well as the incorporation of these software blocks into pilot workflows to verify recommendations in T3.3 on applicability of different middleware environments. This work will specifically address the potential distribution of functionality, given a particular processing need, as well as the required data access patterns and the evaluation of appropriate replica managers.

This task will provide technical effort to address a number of technical interface requirements between WP3 and WPs 4&5. This task will verify that user interface requirements from WP5 can be mapped effectively to workflow models for ESDC processing needs (see Table 3.22), as well as evaluating the ingest requirements are met for data transfers utilising data moving tools assessed jointly between WP3 and WP4 (joint milestone).

Furthermore, this task will provide technical work to verify the scaling of critical elements for the system sizing in T3.4. This will involve prototyping system elements identified as critical by T3.4 and verifying the sizing of these elements.

# Discussion on UK work in AENEAS



- *Can we identify individuals to contribute?*
- *Who are we missing from the current assumed contributors?*
- *Where is there in-kind contribution available? (This is likely to depend very much on the individuals concerned and overlap with existing commitments)*
- *What work can we get started ahead of time?*



SCIENCE DATA PROCESSOR

END

