ERIDANUS

Exascale Research Infrastructure For Data In Asia-Pacific Astronomy Using The SKA

Theme 9 Report on Dissemination

Data Dissemination at SKA Scale and VO Software and Technologies

1. THE PROBLEM

The SKA baseline design includes a storage system for science ready data products, the Preservation System. It does not, however, include a user accessible data archive and has no means to support the generation of advanced (survey) data products and their distribution. This provides a business case for SKA Regional Centres (SRC). Similar considerations led to the establishment of ALMA regional centres, CERN computing and analysis centres as well as EUMETSAT satellite application facilities and regional implementation centres.

SKA's data product generation rate is a key design driver for the dissemination of science-ready data products from the observatory to external SRCs and their user community. For the sake of simplicity and to establish the order of magnitude, one can assume a dedicated 100 Gbps network link to the world, which gives an upper limit for the dissemination rate of 300 PB/year (Figure 1).

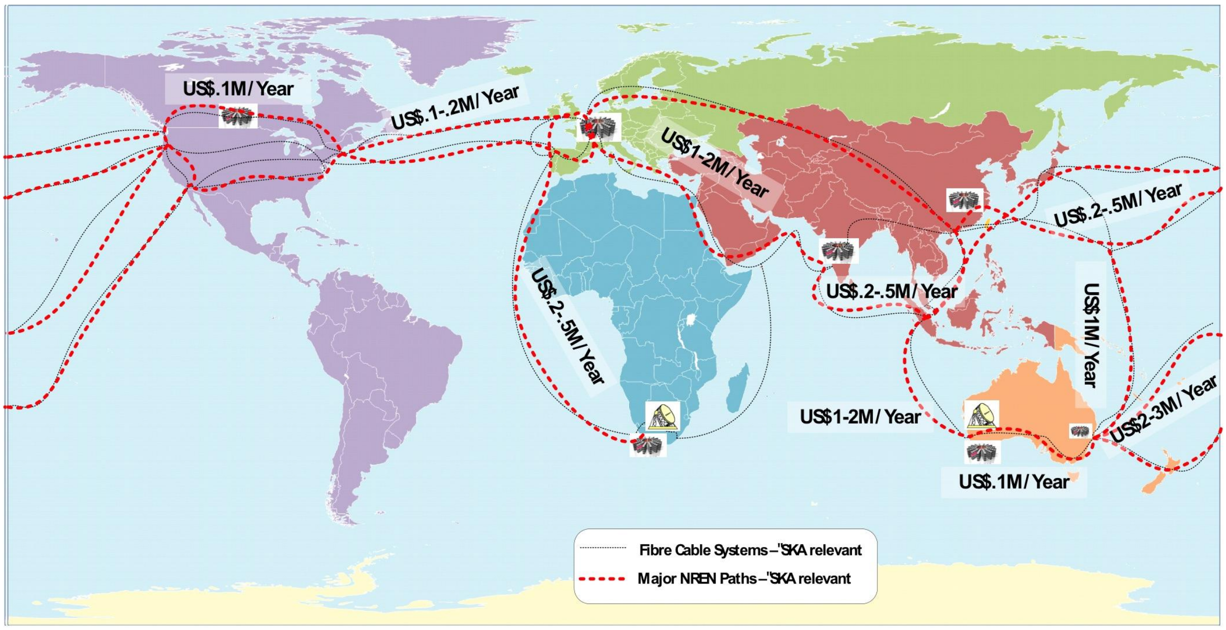

Figure 1: Estimated to world cost per 100 Gbit circuit 2020-2030 and fictional SRC locations. Courtesy SaDT/SKAO.

Controlling the data stream is just an initial, necessary step to the full exploitation of SKA's scientific potential. Collaborative partnerships, co-investment, geographic disbursement and operational efficiency add complexity. Lowering of both the technology barrier and the reduction in the time to "insight" is more important than ever before. The deployment of standardised workspaces, hybrid infrastructures, shared common code libraries on a global scale, and the fervent pursuit of innovation has led a transformation over the last decade with a move towards "transparent insight", or removing the technological restraint of how big the questions are that we ask, and how quickly we can answer them.

2. GROUP ACTIVITIES

The working group (fltr): C. Jin, R. Shen, M. Dolensky, K. Monks, S. Gillard, C. Shi; not on picture S. Gaudet

Experts from industry, academia and government organizations worked through the items of the bullet list below. There was plenty of interaction with other workshop themes and observers including members of the SKA Regional Centre Coordination Group (SRCCG):

- Identification of science data product types and SRC services,
- identification of suitable (VO) technologies,
- dissemination strategy across sites and centres,
- SRC services vs data management environments,
- lessons learned from precursors and pathfinders,
- definition of data challenges based on above,
- and prioritisation.

In addition, there was discussion of two novel approaches in the database field:

1. Negative Database, an indexing scheme for time series data developed for MUSER at KMUST
2. A Multi-Version Data Management system, a concept introduced by C. Jin at ECNU

2.1 Data Products and VO Technology

The SKA science data product types were taken from the high level architecture document of the SDP pre-construction consortium. An analysis of suitable IVOA standards and protocols conducted prior to the preliminary design review milestone was discussed and can be summarized as follows:

SKA Data Product Type	IVOA Candidate Standard Protocols and Models		
	Data Discovery	Data Access	Data Model
Image cube	TAP, SIAv2, DataLink	SODA	ObsCore, NDimCubeDM
UV-grid	TAP, SIAv2, DataLink	SODA	ObsCore, NDimCubeDM
Calibrated visibilities	TAP, SIAv2, DataLink	SODA	ObsCore, NDimCubeDM
GSM/LSM catalogue	TAP, DataLink	TAP	STC, CharDM, PhotDM
Pulsar timing solution	TAP, DataLink	TAP	ObsCore, TimeSeriesDM
Pulsar & transient candidate	TAP, DataLink	TAP	(PSRFITS)
Science product catalogue	TAP, DataLink	TAP	RegTAP, STC, CatalogDM, PhotDM

Table 1: Virtual Observatory support

There is generally good coverage at the level of data discovery, access and modelling. The biggest challenge are operations on large data cubes and to some the degree the fact that there is very few public VO services in the radio regime compared to numerous implementations in other wavelength domains. It was noted that catalogue support is underrepresented in Table 1. It turns out that this is due to the way the science product list was derived rather than a lack of requirements and suitable tools.

In conclusion, it is advisable to experiment with cube access protocols and services and to explore the respective capabilities of the SODA and DataLink protocols for both, subsetting large data products as well as packaging complex datasets.

2.2 FAIR Data Principle for SKA

The focus was on guiding principles for managing the scientific data products. For the physical network aspect of dissemination refer to theme one. A subgroup analysed the FAIR data principles published by FORCE 11 (https://www.force11.org/group/fairgroup/fairprinciples): **Findable, Accessible, Interoperable** and **Re-usable**. The intent is that these data principles may act as a guideline for those wishing to enhance the reusability of their data holdings. Distinct from peer initiatives that focus on the human scholar, the FAIR Principles put specific emphasis on enhancing the ability of machines to automatically find and use the data, in addition to supporting its reuse by individuals.

The FAIR data principle was officially released in 2016 and it has already been endorsed by EU. For example, the European Commission Directorate-General for Research&Innovation. The H2020 programme published the "FAIR data management in Horizon 2020" guidelines (http://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/hi/oa_pilot/h20 20-hi-oa-data-mgt_en.pdf). Other countries also consider to support the FAIR data principles. E.g. The US National Science Foundation has set up a working group to create specific guidelines. Australia ARC is working closely with Australian National Data Service. Thus, it is important to analyse how to make SKA data easily **Findable, Accessible, Interoperable** and **Re-usable**.

To better support the data dissemination at SKA scale, the intention of adapting FAIR data principles apply not only to 'data' in the conventional sense, but also to the algorithms, tools, and workflows that led to that data. The initial group discussion of adapting FAIR data principle at SKA scale can be summarized at the following items:

- Findable
  - Data are assigned a globally unique identifier.
  - Data are described with rich metadata.
  - Data/metadata are searchable.
- Accessible
  - Data are retrievable by their identifier
  - Protocol allows for an authentication and authorization procedure
  - Metadata are accessible.
- Interoperable
  - Data/metadata use controlled vocabularies.
  - (meta)data include qualified references to other (meta)data.
- Reusable
  - Data have a clear license
  - Data are associated with their provenance
  - Data meets astronomy standard

Possible next steps include:

- Using MWA and ASKAP data as a case study to continue refining SKA FAIR data principle
- Consult the draft SKA FAIR data principle with other partners
- Work with relevant parties to endorse SKA FAIR data principle
- Work with relevant parties to apply SKA FAIR data principle to better disseminate SKA data

2.3 SRC Services and Data Management Environments

A service level agreement (SLA) between potential SRCs and SKAO was sketched out following the industry standard ITIL. It was based on a set of requirements for a strawman data access policy (Table 2):

Short Name	Strawman Requirement

Release Status	Each data product shall have a release status and the system shall support scheduled changes to this status information.
Consistent Release State	The system shall provide methods for consistently setting and updating the release status of shared data products in the context of commensal observing.
User Roles	Support an extensible list of defined user roles for the purpose of authorisation.
Data Product Types	Support an extensible set of data product types.
Access Rules	The system shall provide an operator interface that allows the creation and subsequent maintenance of access rules based on user role and product type.
Meta Access	It shall be possible to define access to metadata and data products separately.
Access Log	There shall be a logging facility keeping track of access on a per user and on a per data product basis.

Table 2: Strawman requirements for science archive data access

The derived service catalogue in Table 3 includes also other potential services. It is a starting point and neither complete nor the only possible mapping to services.

Service	Comment
Science Archive	storage, data discovery interface, user portal
Processing Framework	re- and post-processing capability
Backup & Disaster Recovery	for SKAO preservation system
Authorisation, Access Control	based on a data access policy
Logging	usage statistics
…	

Table 3: Sample service catalogue

The general service levels for each service and the organisational responsibilities of SKAO and SRC need to be defined and captured, for instance, using the respective templates available from the Information Technology Infrastructure Library ITIL. Following the ITIL process, a service description includes interface standards, availability target, capacity, access management, responsibilities, type of support, procedures for handling exceptions, to name a few. A domain specific aspect is the curation of data products. It is then a matter of selection suitable key performance indicators (KPIs). It is important that these can be actually measured and turned into system requirements such that the quality of the service can be monitored. KPIs can be categorized into operational, managerial, etc. Examples of KPIs are average incident resolution time (operations KPI), measured availability compared to agreed availability and deviation of predicted capacity from actual (management KPIs).

2.4 Lessons learned from Precursor MWA

Firstly, data access patterns are changing as an archive ages and the telescope progresses from commissioning to operations and further matures and gets enhanced. Exploiting available bandwidth requires fine tuning and knowledge about where to optimize. It is important to ensure reversibility given that the impact of upgrades may not be easily testable upfront. File size does matter, lots of small files are detrimental to transfer efficiency. Checksums are important to detect data corruption. Paper analysis of hardware and software systems is not sufficient: Benchmark, simulate, evaluate.

2.5 SRC Services and Data Management Environments

Progressive systemisation and automation within the scientific insight pipeline, represents one of the potentially game changing" opportunities for the research community. Within the Enterprise space (last 5 years), the application of HPC, AI, Machine Learning on large scales to solve and automate repeatable processes and protocols are underway, and are paving the way of efficiency that is enabling organisations to spend more time innovating and generating insight.

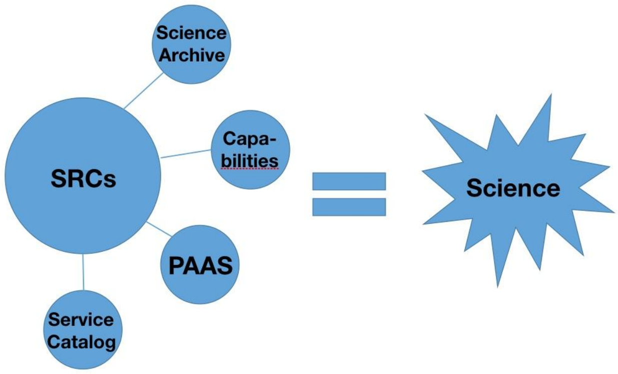

Sharing “workbench / tool benches, and playbooks between geographically dispersed teams is bringing about a homogeneity of approach, a decrease in bespoke “reinventing the wheel” scenarios, that are opening large windows for innovation and creativity. By speaking to common playbooks, or guides of proven approaches and pipeline design, we are creating a culture of common vernacular, and standardised process approaches that not only enable deep customisation, but also pipeline reuse amongst geographically disparate teams.

Taking a “Mechano” modular approach to pipeline design and deployment also decreases the management overhead of system engineering support tasks around processing and storage PAAS environments. By standardising “containerised” Docker / VM based resources, not only is pipeline fragmentation minimised, which typically accounts for 20-30% of project cost over runs.

One of the greatest technological changes in operational flexibility has been underway with the Enterprise move to the use of dedicated, public or hybrid clouds. Sustainable economic and operational efficiency are immediately available with the move to no longer “rolling your own”, but instead using a burstable model for compute resource. Empowering users with “Resource Calculators” that enable Science users to model the Tech resources required to deliver a result within a specific function of (Time or Cost) empowers the team in their prioritisation and optimising of Resource allocation.

For long term research projects, and the cyclical nature of experiments, burstable resource can significantly decrease the costs (move from Capex to Opex) expenditure for organisations. Research organisations are in the science business, the need to be Data Centres is becoming less relevant with each passing day.

Furthermore, by abstracting the infrastructure technology layer from innovation teams, Insight teams are finding the opportunity to practise their science, rather than facilitating it. This one fact has been one of the revolutionary elements of the rise of Data Science and Machine learning in the market, and poses one of great possibility to projects such as SKA.

3. Next Steps

The group concluded that setting up a SDP Preservation System in Australia and linking it up with an SRC in China is the most effective way to identify requirements, explore the technology and build bonds between the Sino-Australian and further Asia-Pacific partners of the emerging Eridanus initiative. The goals of the next steps can be summarized like this:

- foster partnerships within the Asia-Pacific region in the context of Eridanus
- explore organisational aspects of SRCs
- inform and consult with SRC working groups and potential stakeholders
- lower the technology barrier for the user community
- analyse international network connectivity to/from China

Each of the seven scenario elements outlined below does contribute to at least one of the goals and the deployment of individual elements is not tied to specific geographic locations.

1. A Data Product store that
2. gets populated with public MWA Precursor Data and
3. a has a Science Product Catalogue accessible through a VO Portal and

4. supporting FAIR Data Principles.
5. There is a Workbench for processing the stored Data Products.
6. Capabilities are captured in a Service Catalogue (SLA) and
7. a Token Framework provisions and tracks resource usage.

For the purpose of dissemination there are at least two instances of the Data Product store (#1). One instance emulates the SKA Preservation System and the others correspond to SKA Regional Centres (SRC). Public MWA Precursor Data (#2) are a sizeable test data set (~1 Petabyte) for transferring and processing. The setup itself is not MWA specific. There is interest in a similar exercise for ASKAP by creating a CASDA (CSIRO ASKAP Science Data Archive) clone at SHAO and further data sets may be added, provided there is sufficient capacity. The Science Product Catalogue and VO portal can be adapted from existing science archives and VO software tools. FAIR Data Principles (#4) are a source of requirements and set the scope for the functionality exposed by the portal (#3). The Workbench (#5) exposes a processing environment and provides pre-packaged data reduction codes, for instance, as Docker images. One element of #5 is the SDP logical graph editor, which is a prototype for defining complex workflows. It has a web GUI and default recipes. This limits the required domain knowledge needed when defining and running large complex tasks. The service catalogue (#6) is an inventory and a first step towards (service level) agreements between the observatory and external contractual partners. In a non commercial environment the notion of a token budget (#7) is meant to improve transparency at all organizational levels. The capacity of a service can be expressed in terms of available tokens. Using a service consumes tokens. The resource allocation when defining a workflow (#5) on the workbench results in a token spending plan. The seven elements are loosely coupled so that hitting a snag in one area does not block the other activities during this explorative phase.

Acknowledgement
Special thanks go to

- Stefan Gillard, EngineRoom.io, for sharing insights on many aspects of state of the art processing environments,
- Robert Shen, AAL, for raising awareness of FAIR data principles,
- Kim Monks, Systemic, for filling us in on service level agreements,
- Chequing Jin, ECNU, for presenting the concept of a multi-version data management system,
- Congming Shi, KMUST, for introducing the Negative Database system,
- Chen Wu, ICRAR, for providing lessons learned from the MWA precursor telescope,
- James Dempsey, CSIRO, and Minh Huynh, CSIRO/ICRAR, for a shared session on SKA data products and requirements, and last but not least to
- Nick Rees, SKAO, Shenghua Yu, NAOC, and Séverin Gaudet, CADC, all members of the SKA Regional Centre Coordination Group, for feedback and advice.