

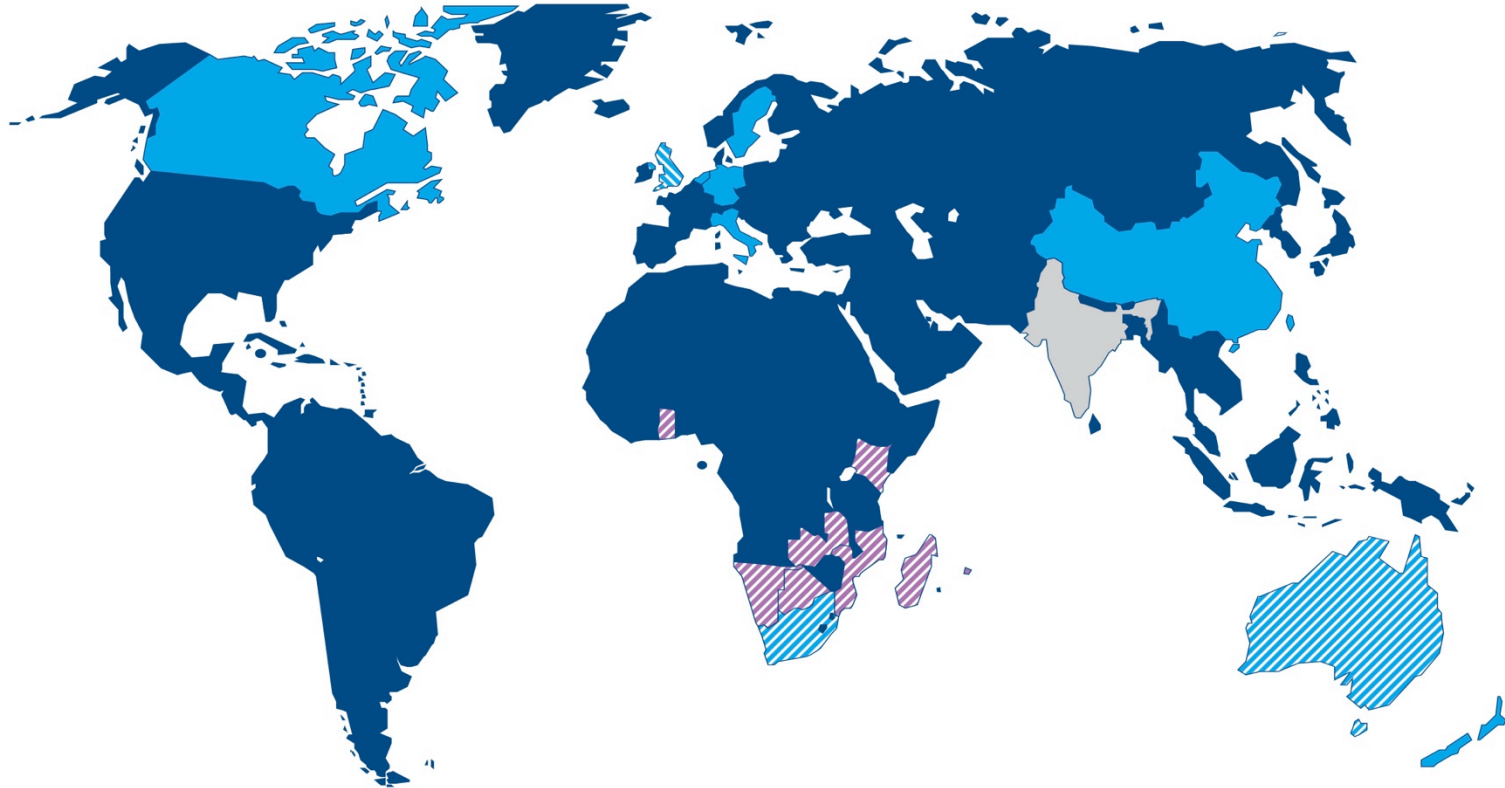
# Overview

---



- SKA project
- SKA telescope design
- SKA SDP package
- SDP data delivery element – DELIV
- DELIV testbed

# SKA Global Partnership



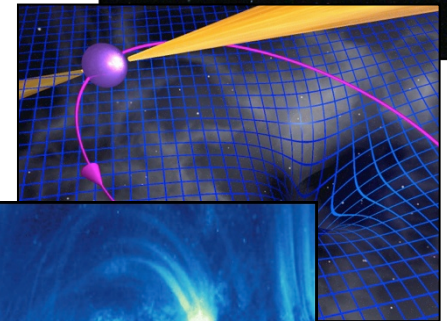



- 
- Full members
  - Associate members
  - ▨ Member SKA Phase 1 and Phase 2 host countries
  - ▨ Non-member SKA Phase 2 host countries
  - ▨ SKA Headquarters host country

# SKA Key Science Questions



- **Probing the Dark Ages**
  - When & how were the first objects in the universe formed?
- **Cosmology and Galaxy Evolution**
  - Nature of Dark Energy and Dark Matter
  - Formation and evolution of galaxies and structure
- **Gravity**
  - What happens to space-time under extreme conditions?
  - Gravity-wave astronomy (pulsar timing arrays)
- **Origin & Evolution of Cosmic Magnetism**
  - Where does magnetism come from?
  - What is its role in cosmic evolution?
- **Cradle of Life**
  - What and where are the conditions for life?
  - Does intelligent life exist elsewhere?



# SKA Timeline




South Africa

€ 70 M

**SKA1**  
Pre-construction

**SKA1**  
250 15-m Dishes

**SKA**  
2500 15-m Dishes



1%

10%

100%

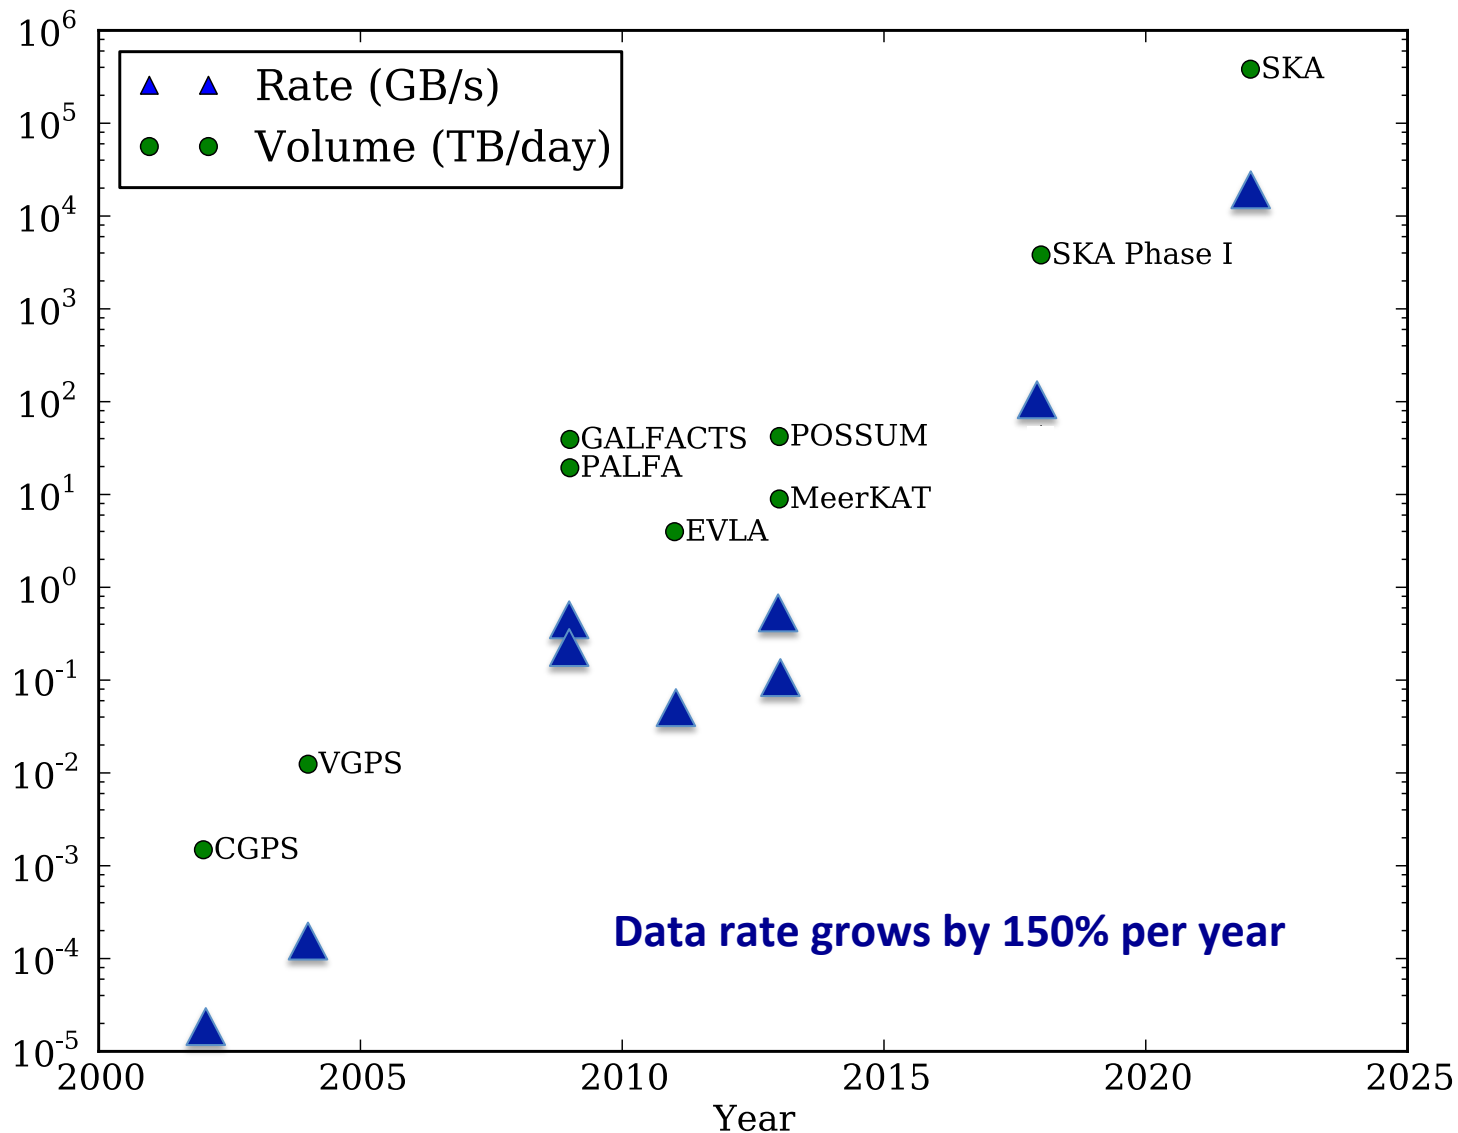
Australia

**SKA**  
Low-frequency antenna array





# Data Rates and Volumes to Scientists

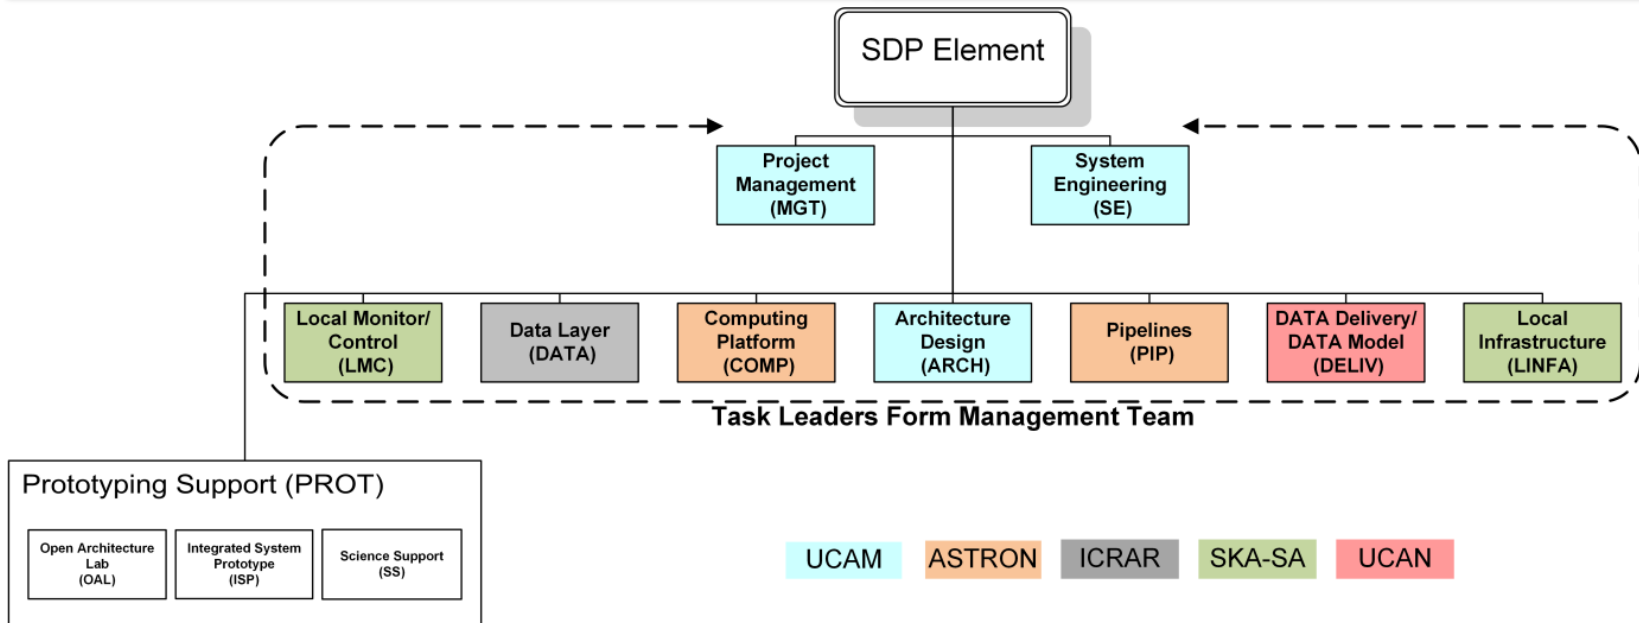


# SKA Design



- Three telescopes: Low, Mid and Survey
  - Each with computing resources
- Design work packages:
  - Assembly, Integration and Verification (AIV)
  - Central Signal Processor (CSP)
  - Dish (DSH)
  - Infrastructure Australia and Africa (INFRA AU/INFRA SA)
  - Low-Frequency Aperture Array (LFAA)
  - Mid-Frequency Aperture Array (MFAA)
  - Signal and Data Transport (SaDT)
  - Science Data Processor (SDP)
  - Telescope Manager (TM)
  - Wideband Single Pixel Feeds (WBSPF)

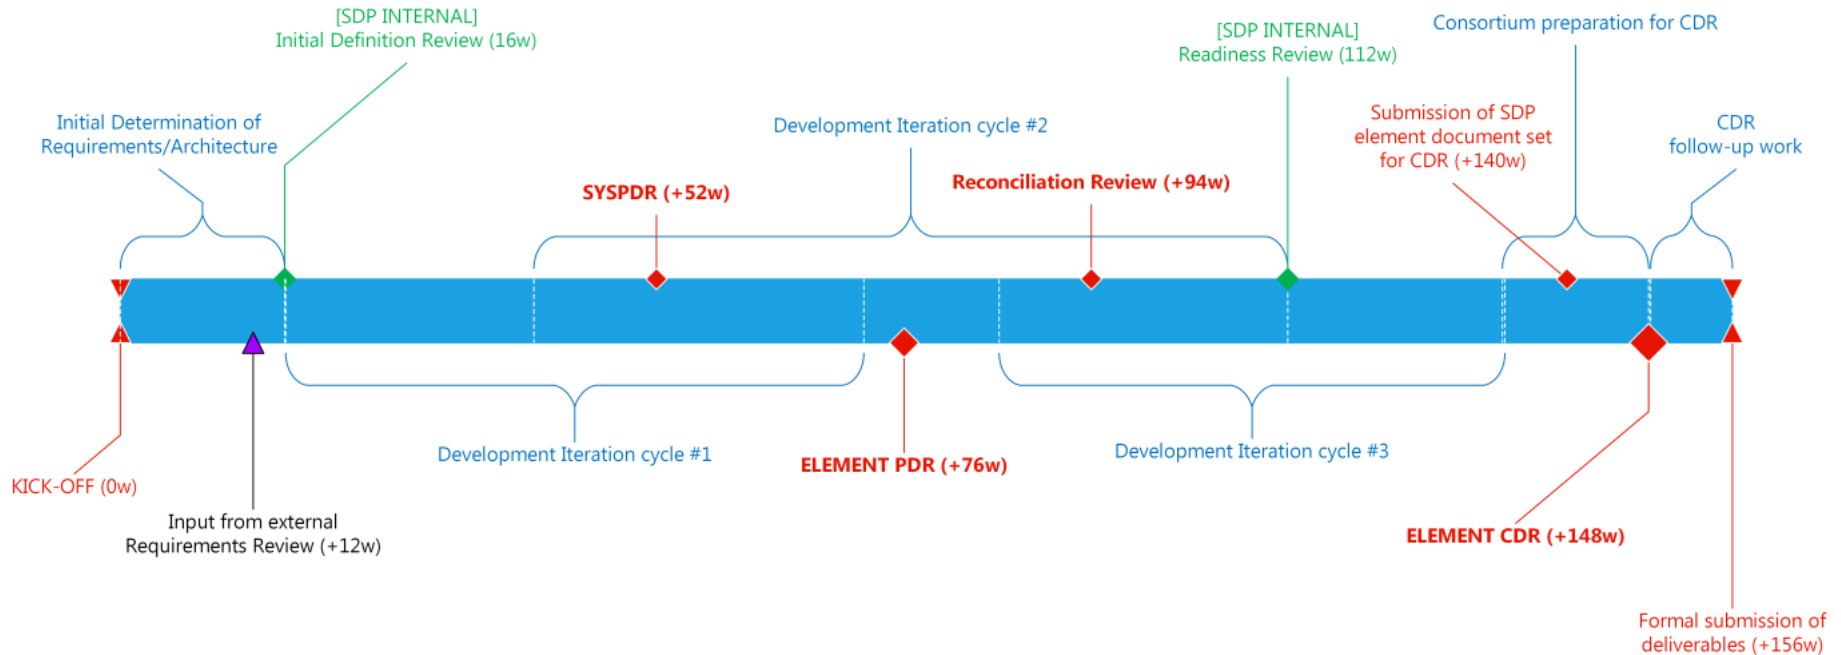
# SDP Top Level WBS and Management Team



- Lead: Paul Alexander
- PM: Ian Cooper
- PE/Architect: Bojan Nikolic
- SE: Ferdl Graser
- PS: Rosie Bolton

- COMP: Chris Broekema
- PIP: Ronald Nijboer
- DATA: Andreas Wicenec
- DELIV: Rob Simmonds
- LMC: Simon Ratcliffe
- LINF: Jasper Horrel

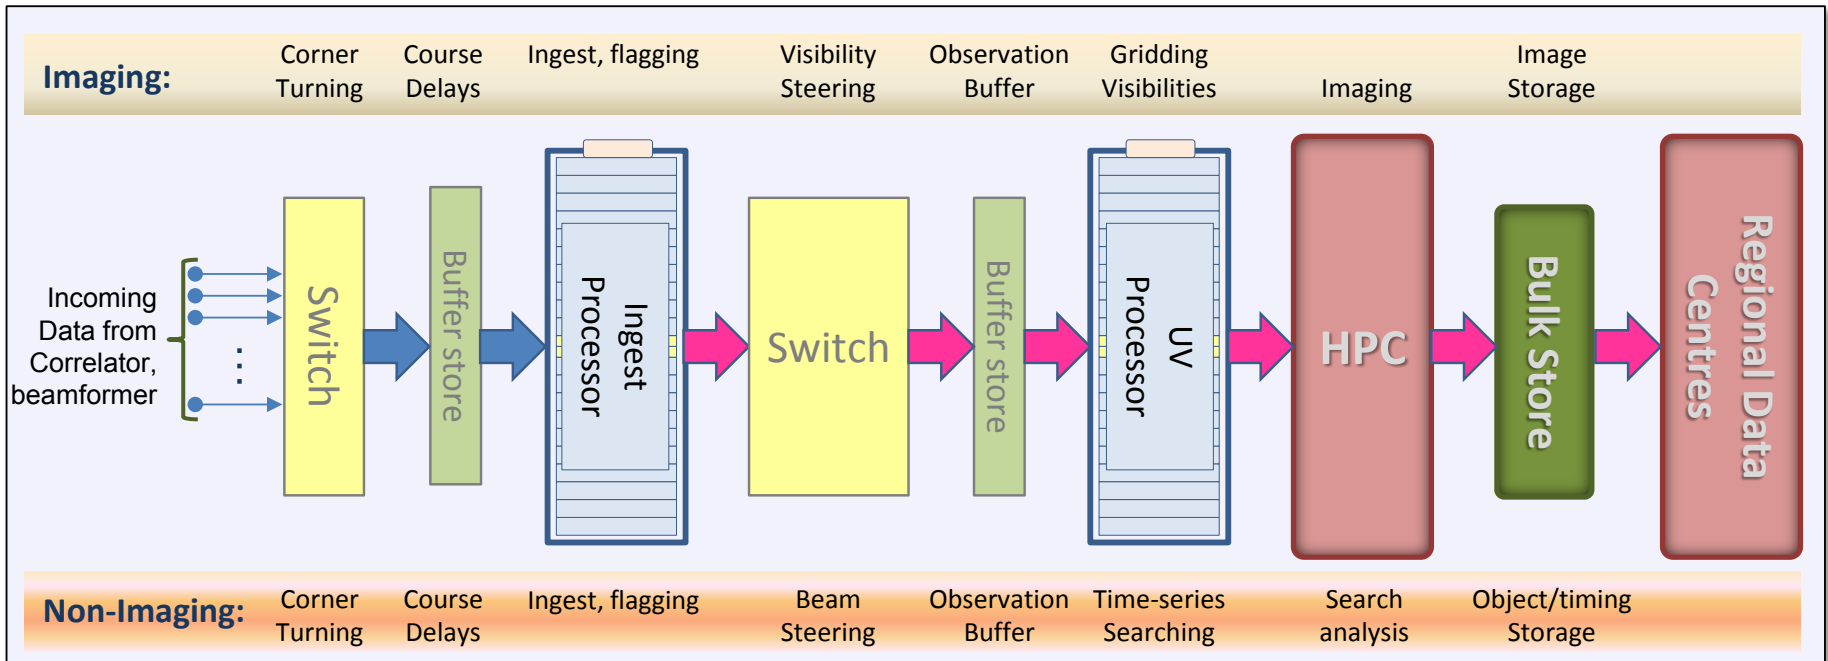
# SDP Design Schedule



- Expectation: major system changes following re-baselining
- Follow standard SE approach of element PDRs following system PDR
- Baseline element architecture at PDR



# SDP: Overall Architecture



- Heterogeneous hardware architecture
- Homogeneous software stack

# Example Data Rates and Data Products



- Baseline Design

Element	BL Design (Ingest) (GB/s)	Use Case (maximal) (GB/s)
LFAA	842	245
Survey	4670	995
Mid	1800	255

- Aperture Array Line experiment (e.g. EoR)
  - 5 sq degrees; 170000 channels over 250 MHz bandwidth
    - ~ 30 GB/s reducing quickly to ~ 1GB/s
    - Up to 500 TB UV (Fourier) data; Images (3D) ~ 1.5 TB
- Imaging experiment with long baselines
  - 50 km baseline with the low-frequency AA or SKA1\_Survey
    - 1.5 TB/s reducing to ~ 50 GB/s
    - Up to 1000 TB/day to archive if we archive raw UV data
    - Images (3D) ~ 27 TB

# DELIV



- 
- Provide science products to end users
  - Access data online and move data to where users can perform analysis
  - Enable searching for data products
  - Explore roles for Regional Centres to support SKA research
  - Manage the distribution of data to make best use of available network capacity

# DELIV partners

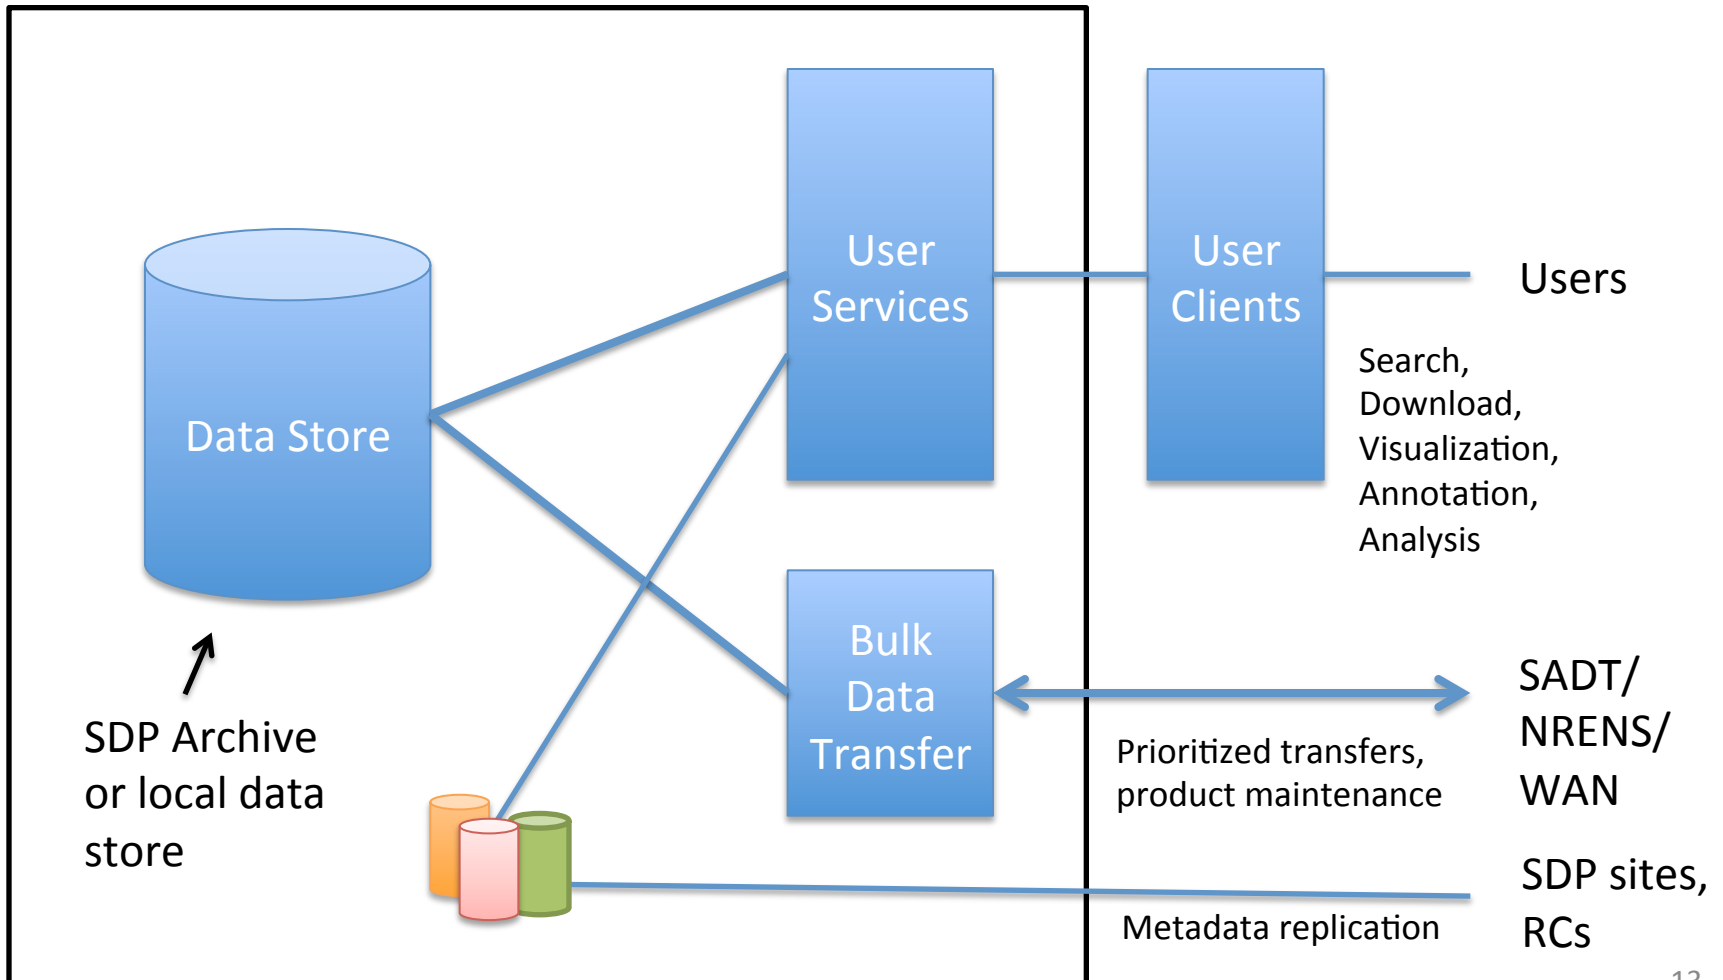
---



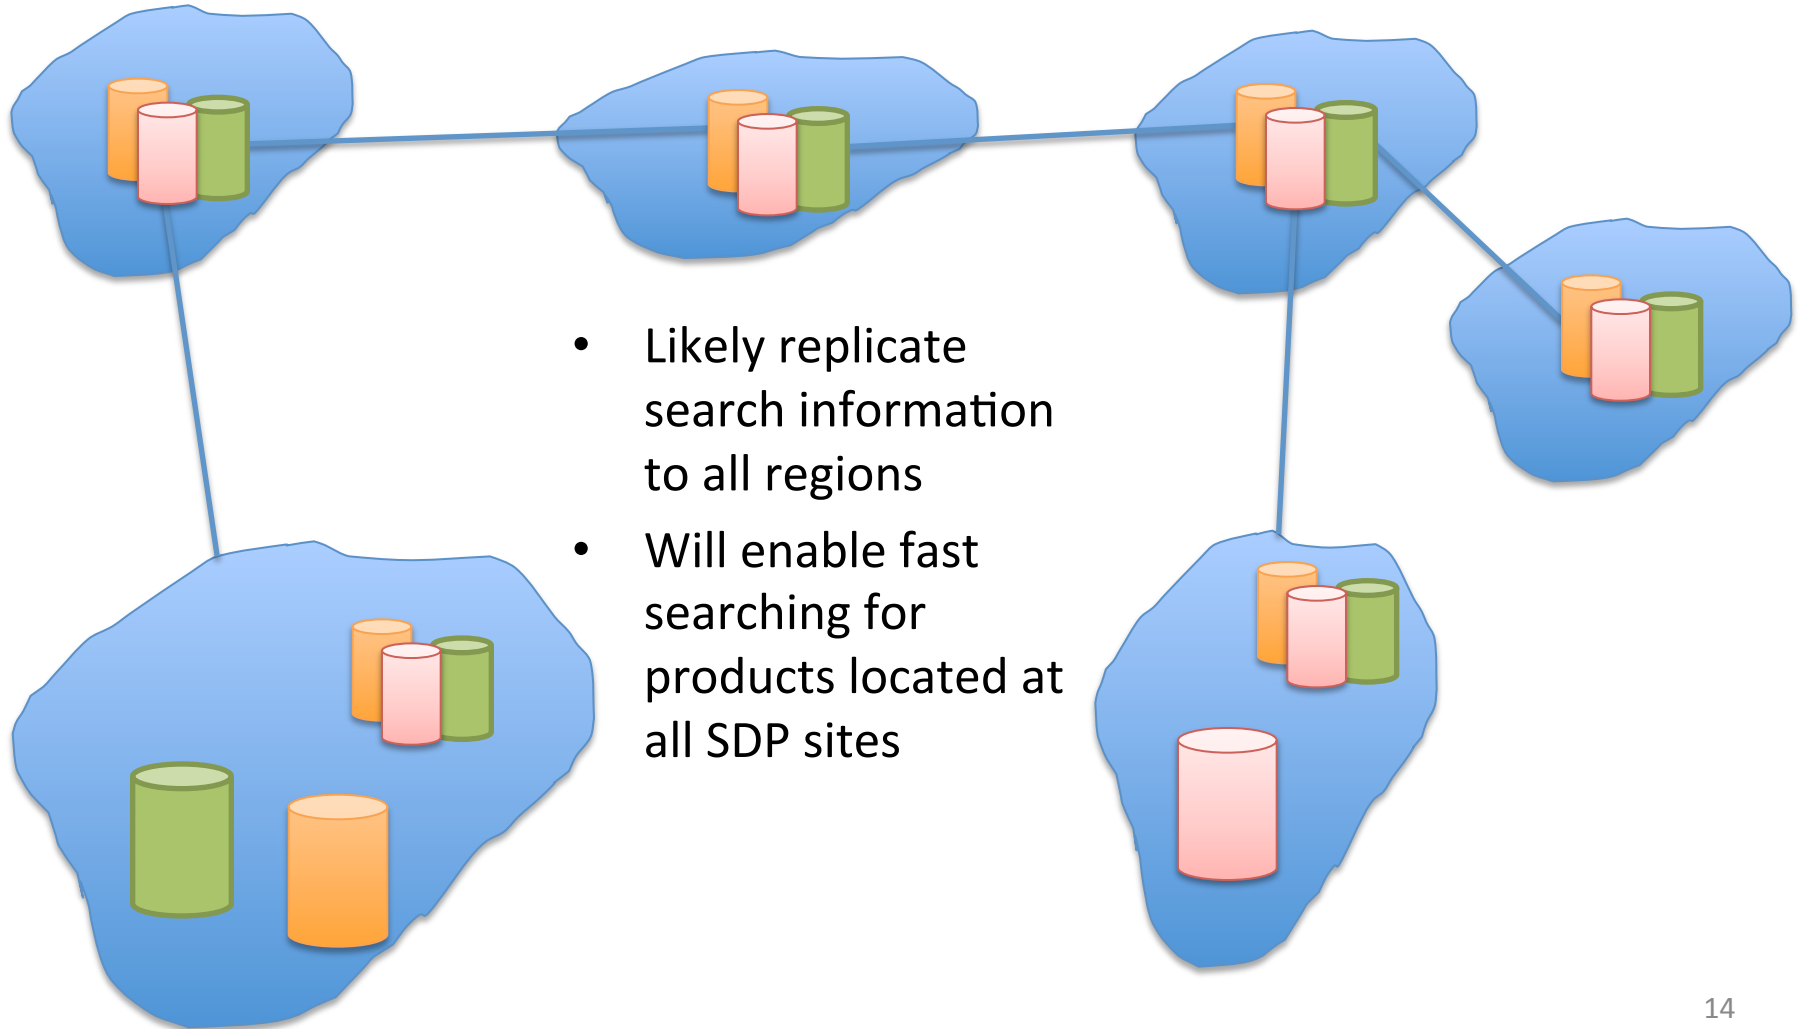
- Design partners:
  - Canadian University Consortium
  - Oxford eScience Centre
  - SKA-SA (input from SAC)
  - Canadian Astronomy Data Centre
  - ASTRON
- Additional prototyping partners:
  - RackForce
  - CHPC
  - iVec
  - University of Cambridge



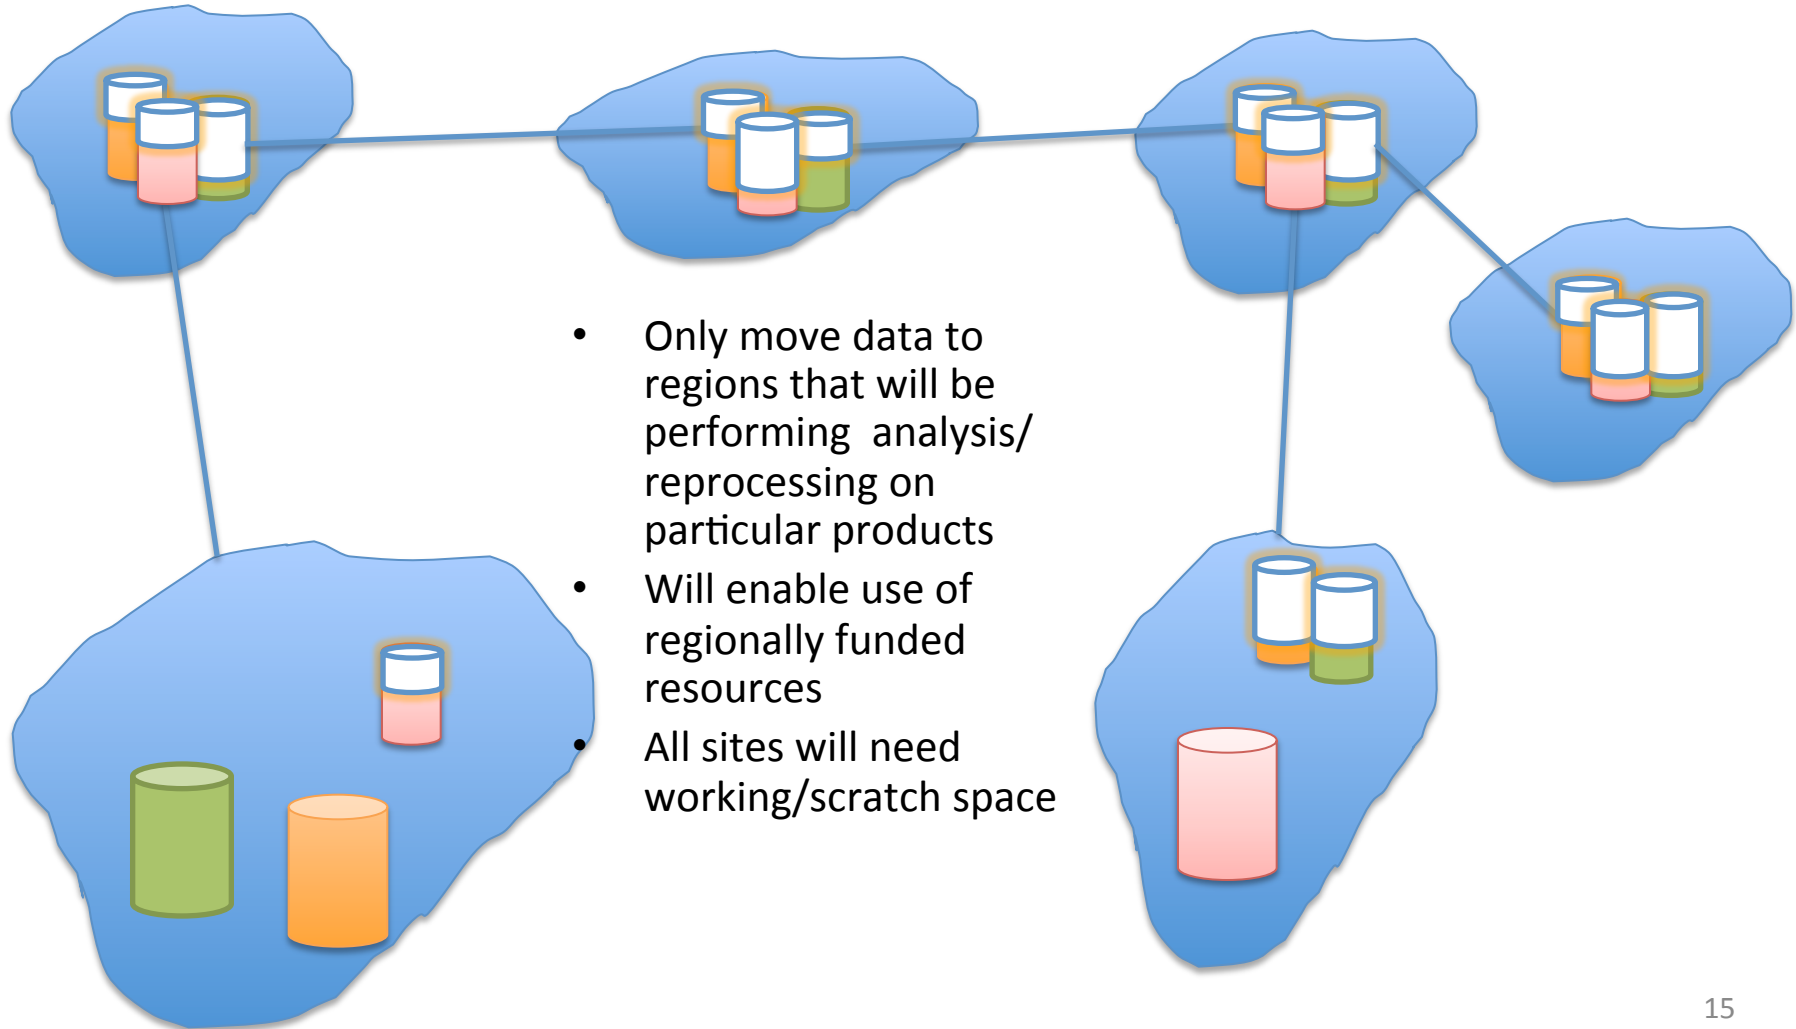
# DELIV – Data Architecture



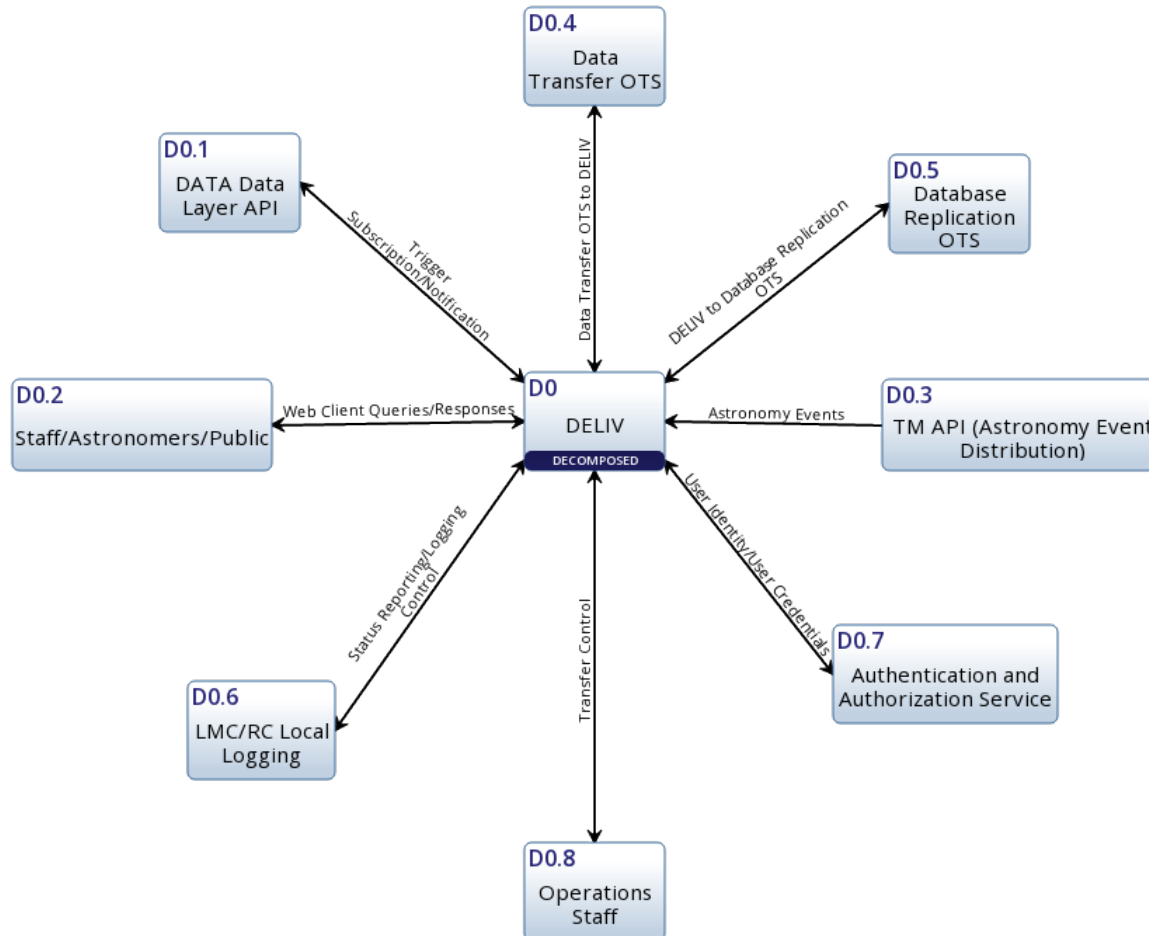
# Search Data Replication



# Data Product Replication

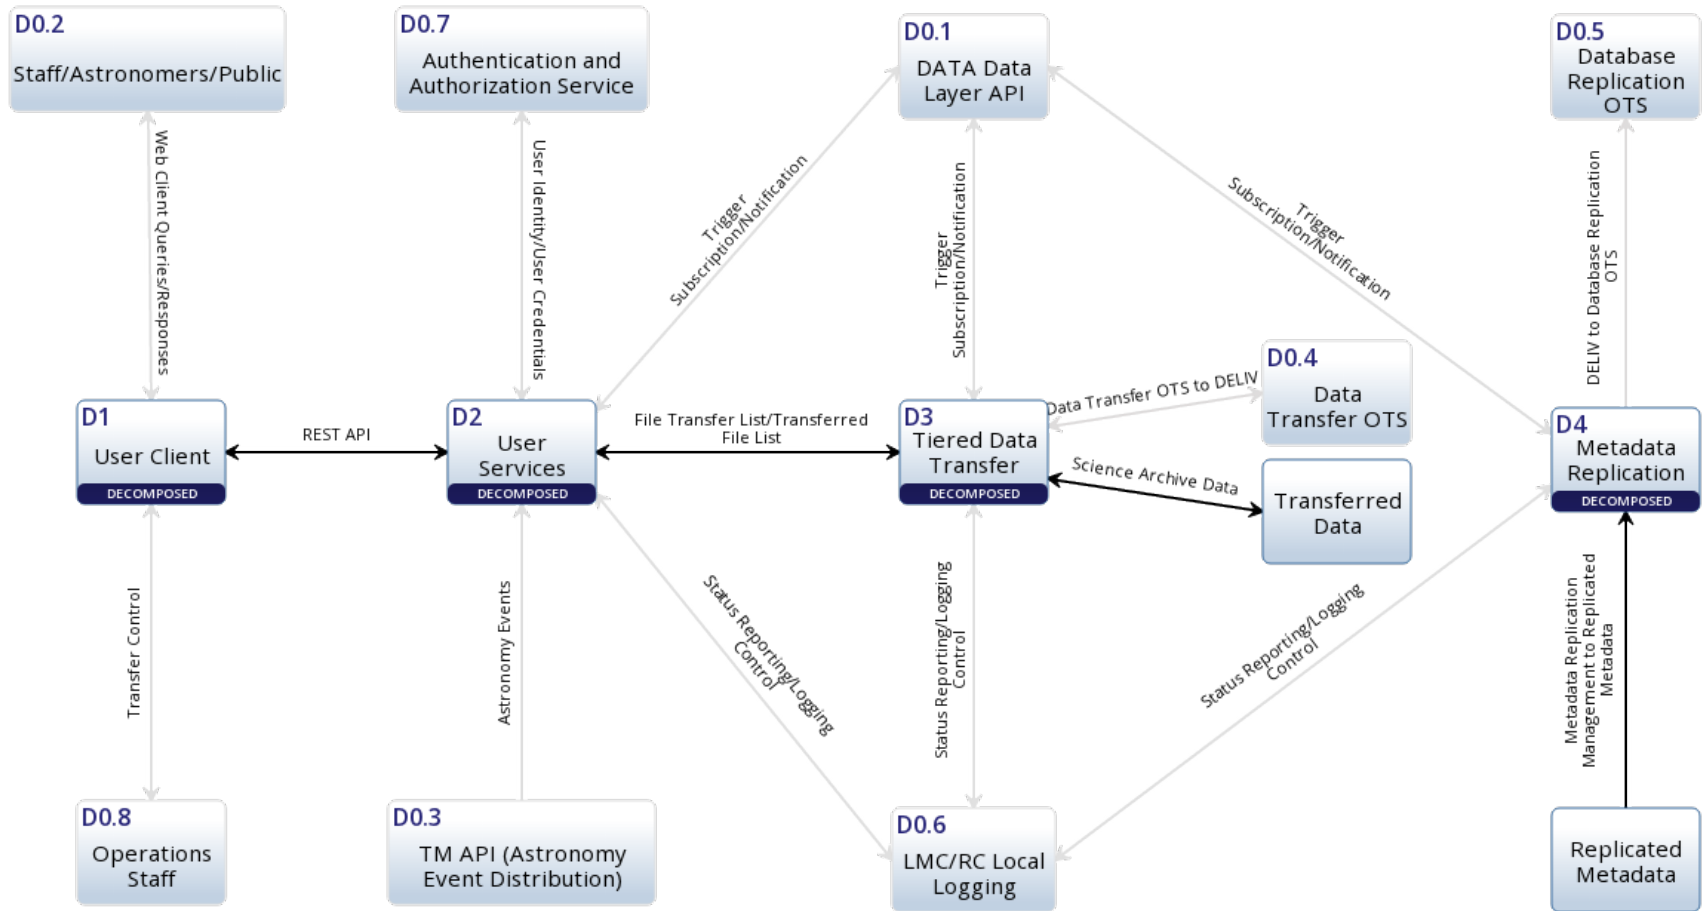


# DELIV Context Diagram

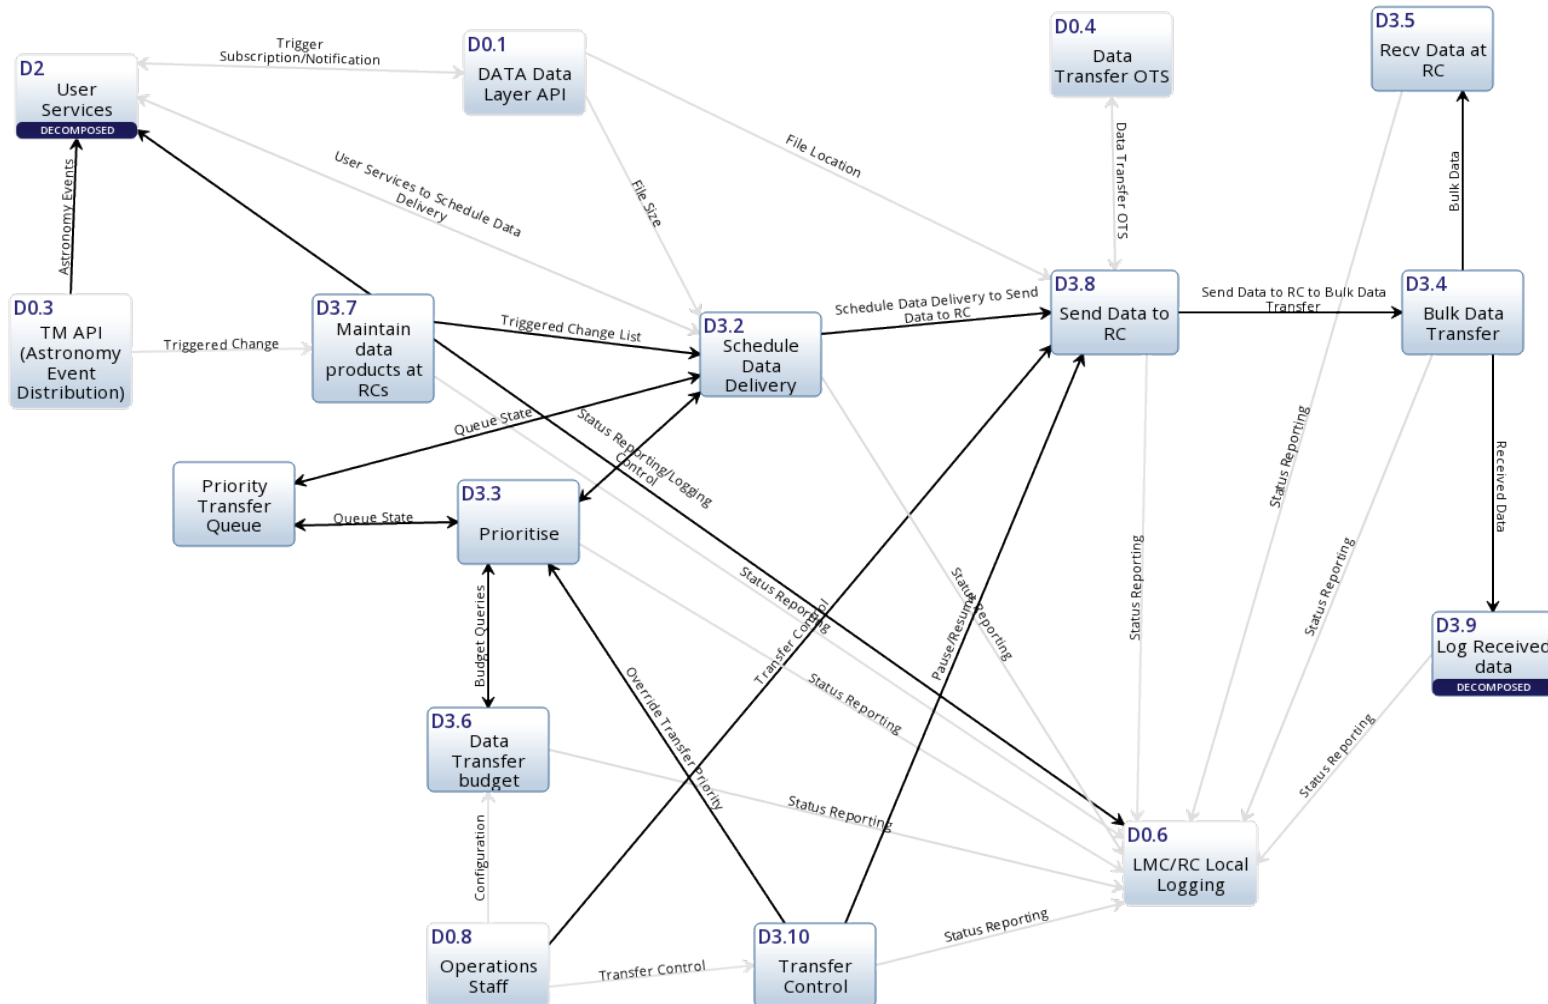




# DELIV: Decomposition



# DELIV : Data transfer



# Regional Centres

---



- Regional Centres not funded within SKA cost cap
- Recommend they provide capabilities to worldwide SKA community
  - Cache/redistribute data, support point to point network links, monitor network links, provide common software stack, enable use of regionally funded systems
- Assume RCs trusted by SKA
  - Enforce access control, log and report data access
- Agreements with SKA office will enable science

# Networks : WAN

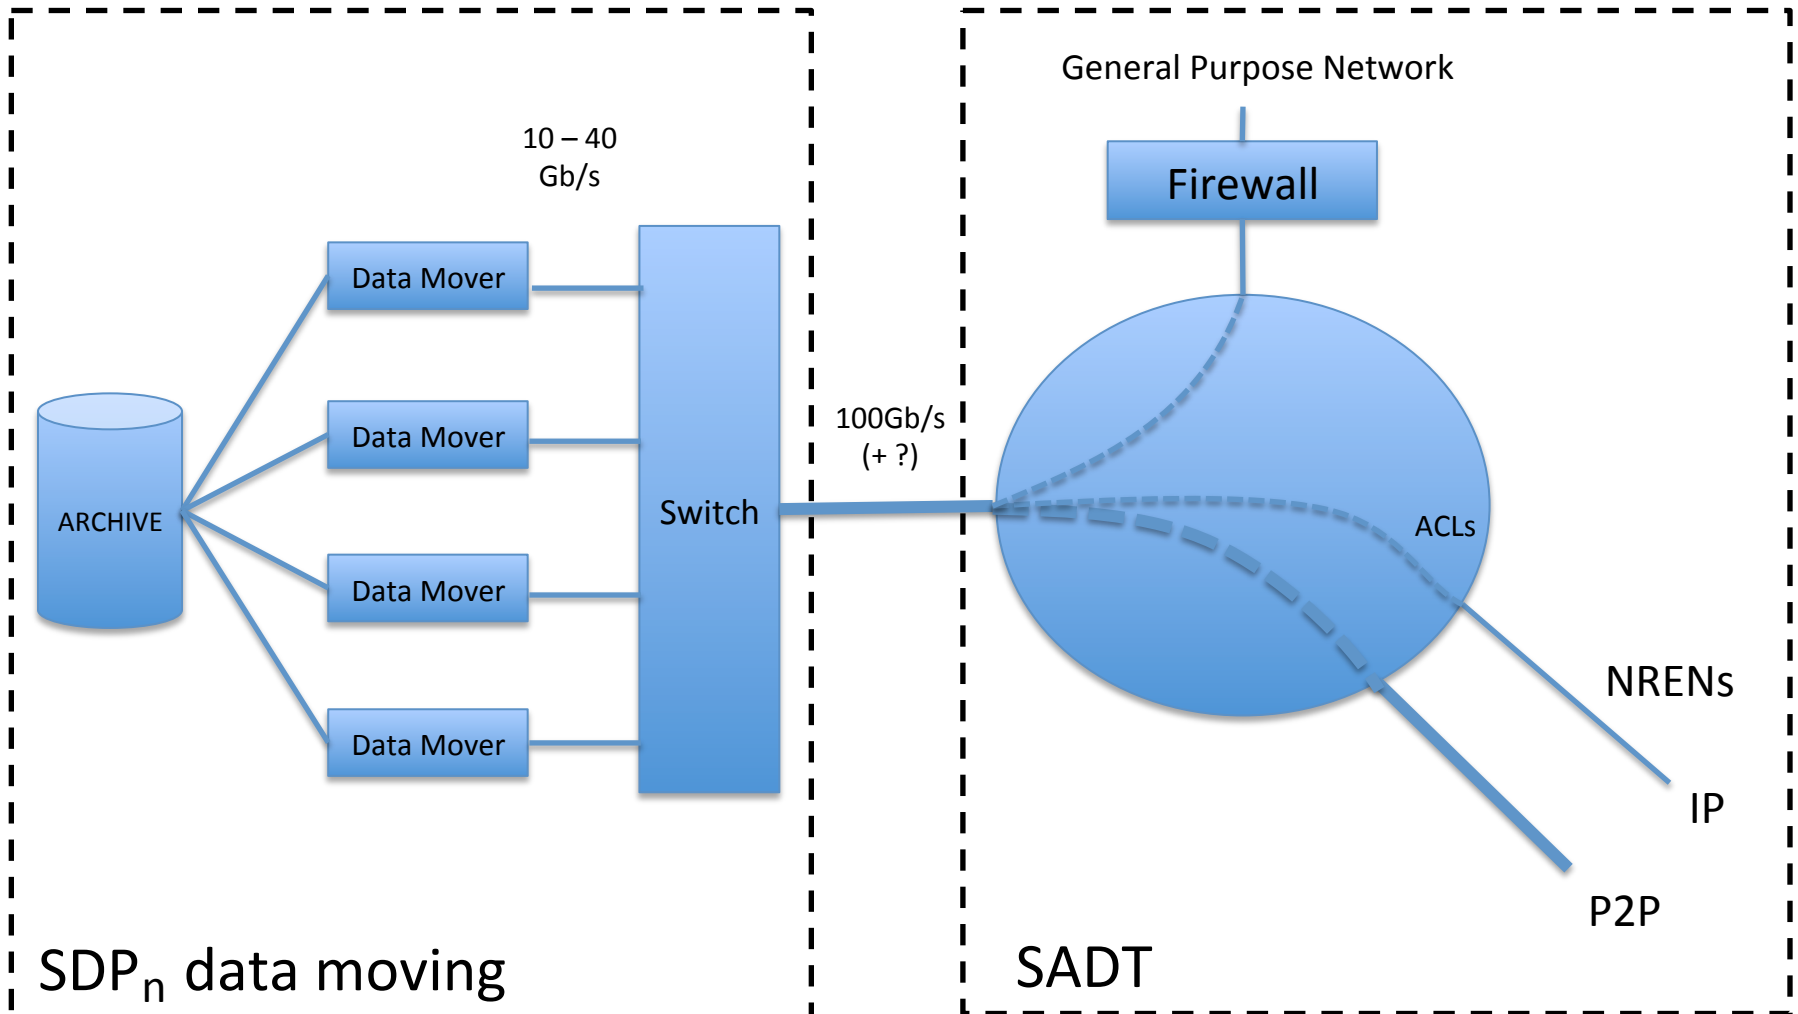
---



- SADT responsible for WAN connectivity to SKA
- DELIV looking at optimizing how WAN is used
- WAN needs to carry :
  - Data replication (bulk) and download
  - Data search information
  - Interactive access to main SDP sites
  - “A&A” information from SKA
  - Data access information to SKA
  - VLBI data (4-10Gb/s)



# Possible WAN Interface



# WAN technologies



- WAN bandwidth improvements
  - Butter's Law – bandwidth doubles every 9 months
  - Price stays the same, bandwidth increases
  - Current research networks deploying 100Gb/s
  - Current technology demonstrators - 1.4Tb/s

## Other :

- Overlay networks
  - Various uses including segment optimization
- Software Defined Networks (e.g., OpenFlow)
  - Redirect P2P paths when needed
  - Deliver to multiple targets concurrently
- Bandwidth-on-demand
  - Assign lightpaths / guarantee bandwidth on shared paths

# Technology exploration



- Level 0 – low level tools and protocols
  - RDMA, SDN, Messaging
- Level 1 – Data Transfer
  - GridFTP, FDT, bbcp, Aspera, ...
- Level 2 – Data Distribution and Management
  - iRods, NGAS, OODT, FhEDEX, CEPH, ...
- Level 3 – User
  - Portals (Elgg, Drupal, Joomla, ..), Remote Visualization (Pureweb, NICE,...), IVOA tools
- Other
  - Data Models (ObsCore), Processing frameworks

DELIV looking at tool use in big data projects

# Web Portal : e.g., Cyberska



- Need interface to search, download and analyze data and enable collaboration

A screenshot of the Cyberska web portal interface. The header features the 'CYBERSKA' logo and the tagline 'A Cyberinfrastructure platform to meet the needs of data intensive radio astronomy on route to the SKA'. Below the header is a navigation bar with links for Home, Profile, Settings, myDashboard, myGroups, Tools, and About Help, along with a search bar and a Log out button. The main content area is divided into several sections: a user profile for Cameron Kiddle with options to subscribe or bookmark; a Contacts grid showing a grid of user avatars; an Event calendar listing upcoming meetings like 'Imaging Science Technical Meeting' and 'CANARIE Site Visit'; a Recent Astro-ph Eprints section with links to arXiv preprints; an Active Users section showing Cameron Kiddle and Samuel George; and an Activity section with recent discussions and comments. A Group membership section on the left lists various subgroups like 'DMS support of Astronomical Data' and 'CyberSKA Sys Admins'.

# CyberSKA Motivation

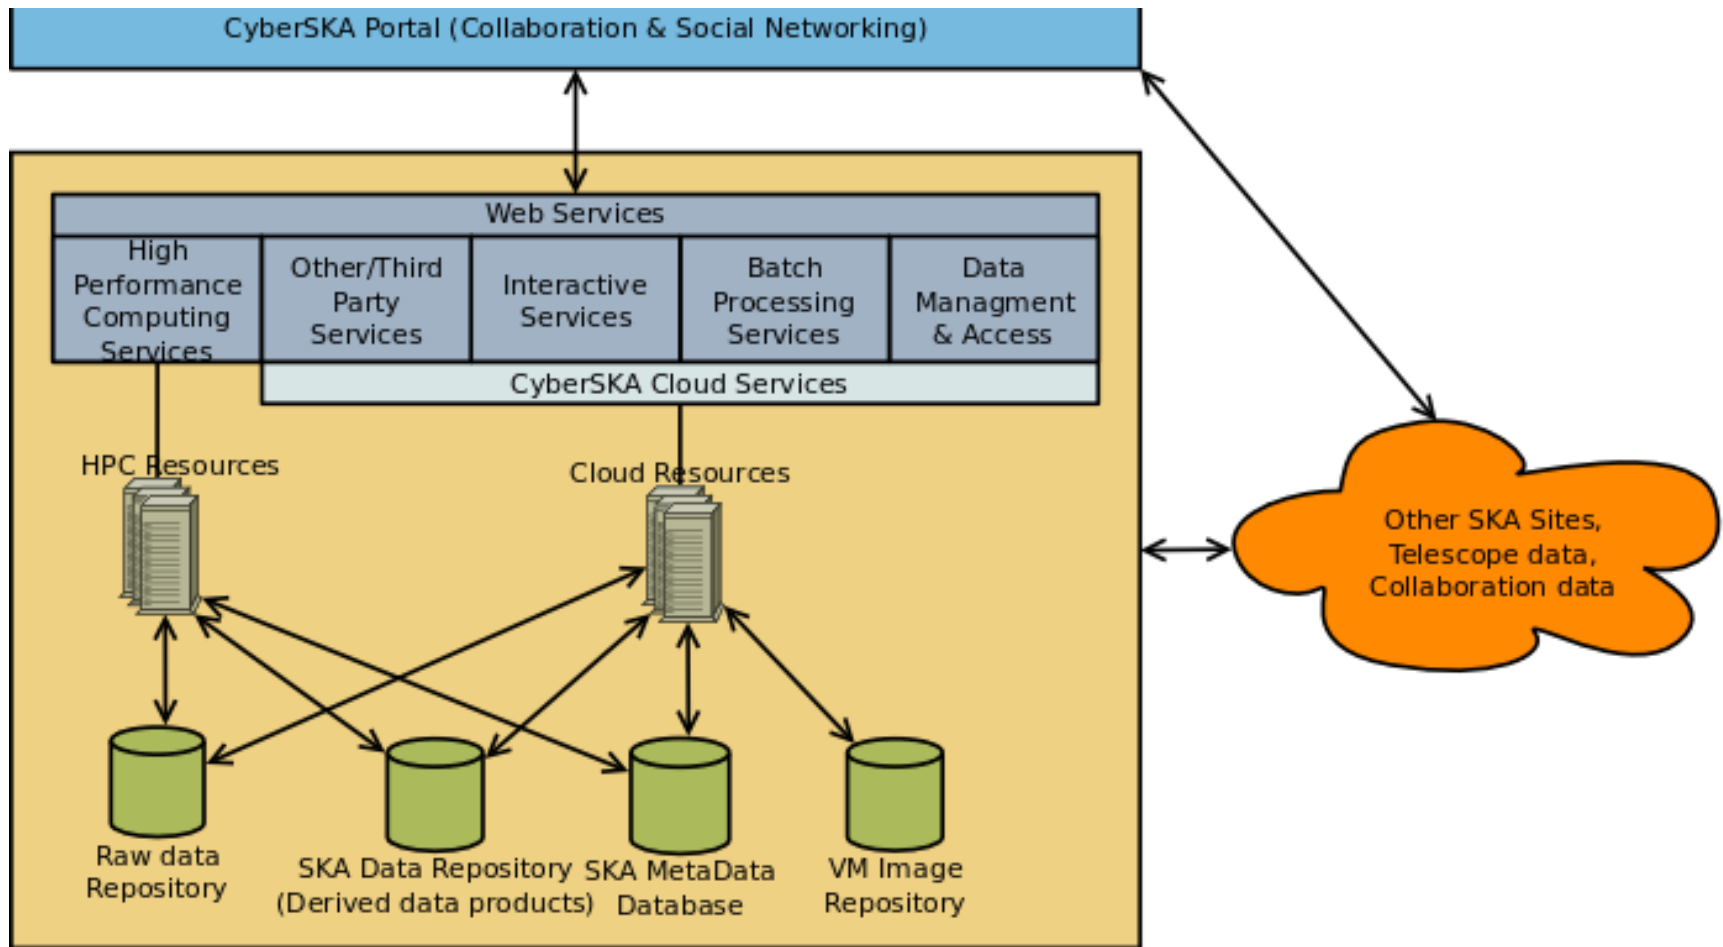
---



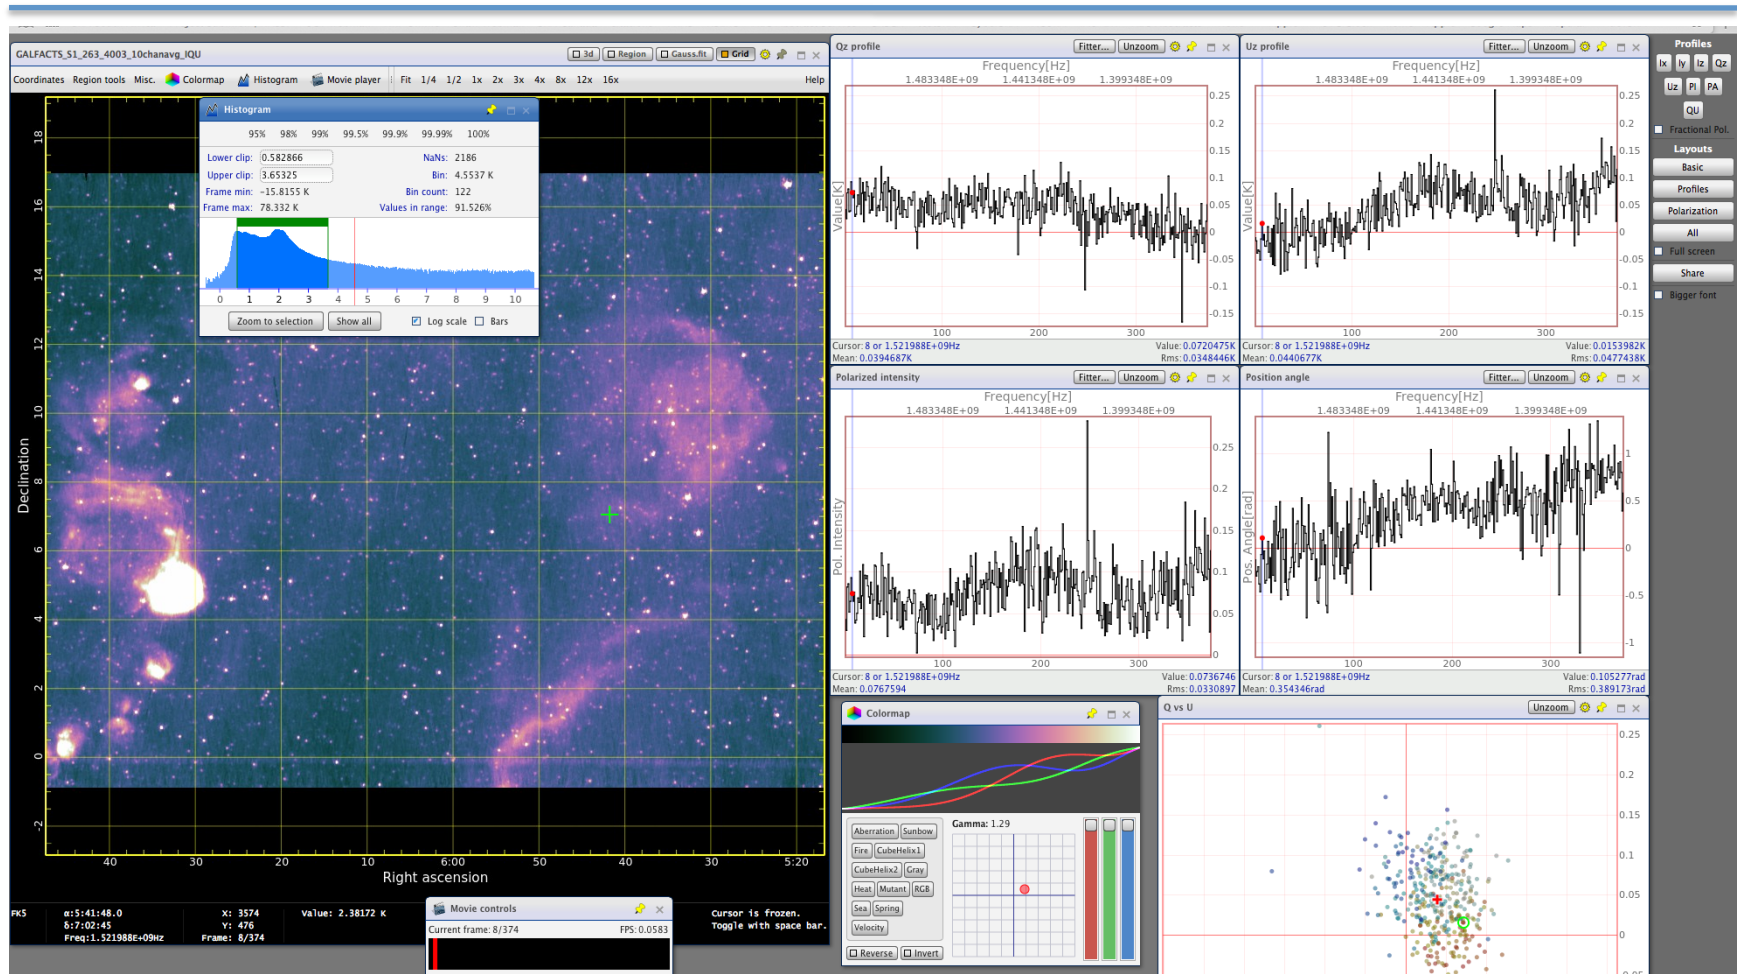
- Most SKA key science goals will be achieved via large-scale survey type observing programs
  - Very high data rates and volumes
  - Complex, multi-purpose, processing and analysis
  - Executed by globally distributed teams of researchers
- Drives the need for cyberinfrastructure solutions for
  - Data storage, management and distribution
  - Data processing, analysis and visualization
  - Collaboration tools

(Tools developed with Calgary Scientific (PureWeb) and IBM Canada)

# Initial CyberSKA architecture




# Remote Viewer



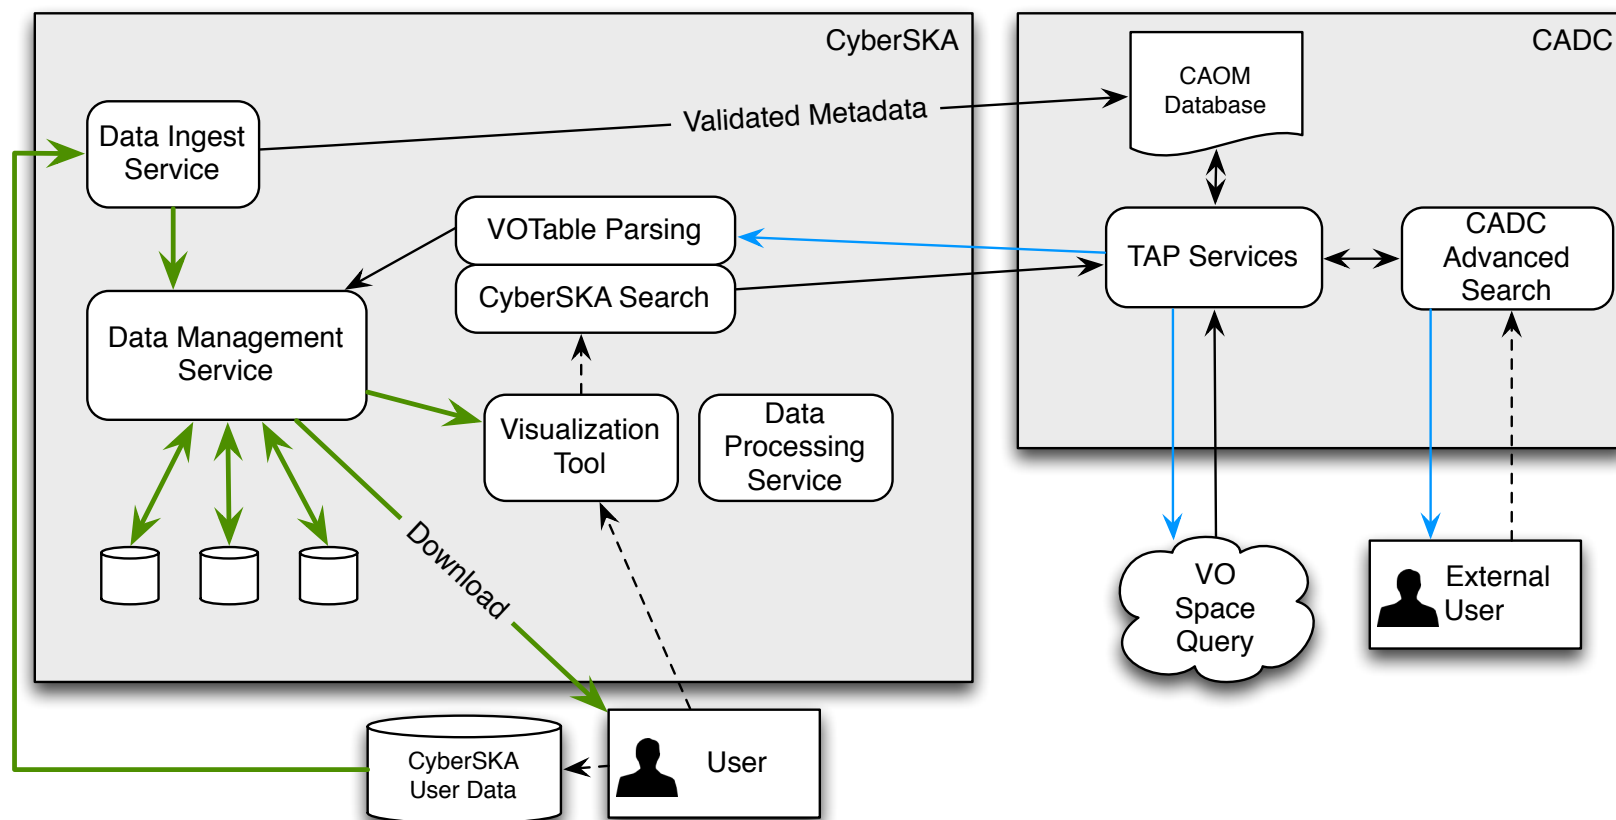


# Visualization demo





# IVOA interface (CADC)



# CyberSKA: Imaging Survey Use Case



Arecibo Observatory



courtesy of the NAIC - Arecibo Observatory, a facility of the NSF

Expanded Very Large Array (EVLA)



courtesy of the NRAO (National Radio Astronomy Observatory)

Giant Metrewave Radio Telescope (GMRT)



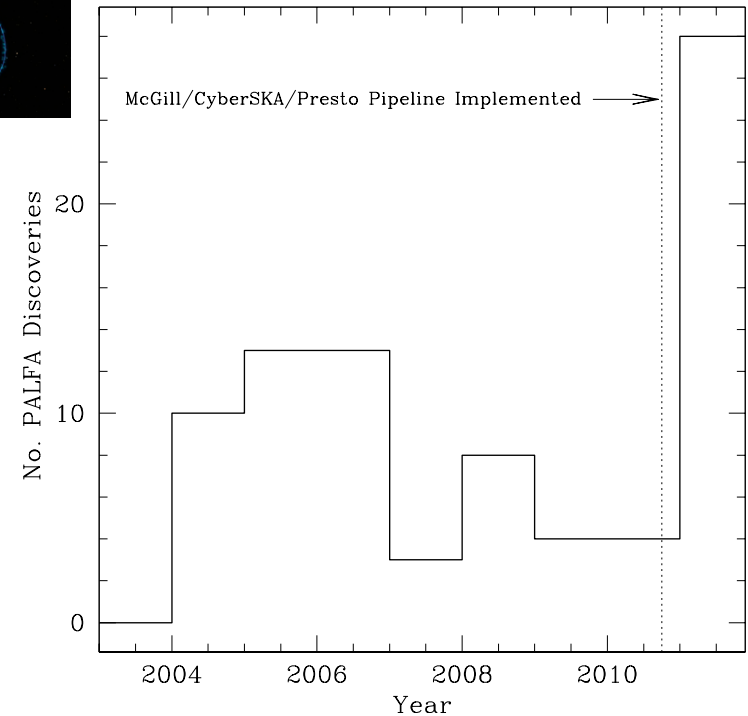
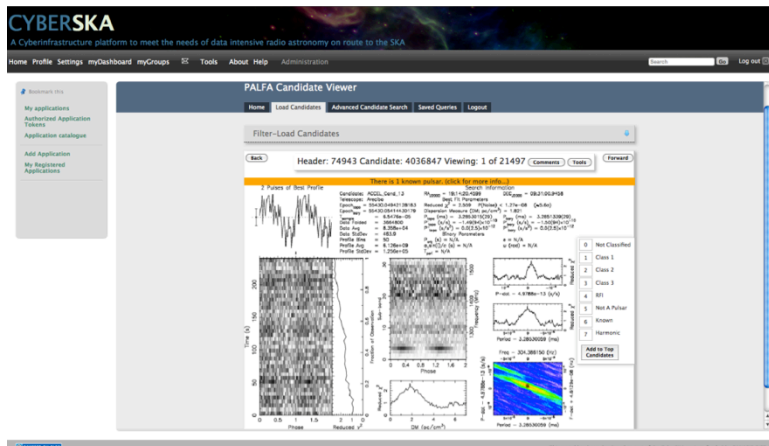

courtesy of the Tata Institute of Fundamental Research

- GALFACTS (Arecibo)
- Deep Polarization Field Surveys (EVLA, GMRT)
- CyberSKA:
  - Used for sharing documents, creating wiki pages, having discussions and bookmarking resources
  - Enables on-line visualization of remote data sets
  - Provides access to GALFACTS survey data and third party applications for running data processing pipelines or source counting

# Enabling Discovery



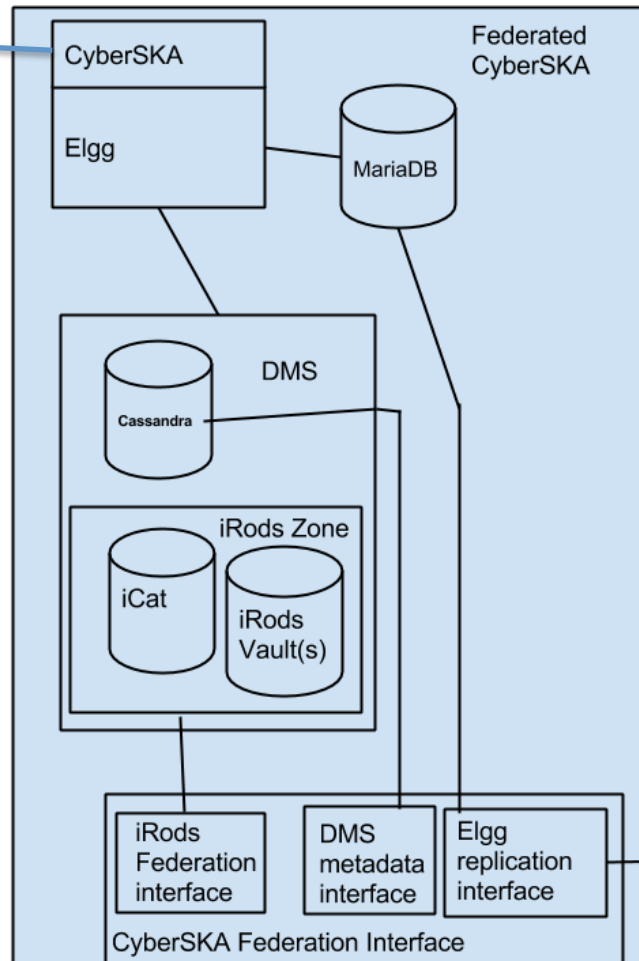
- Pulsar Search Project – 69 global collaborators
  - Collaborative data processing and third-party applications



# Federated CyberSKA



Users



- Experimental/prototype platform for exploring distributed access
- Portal data shared using multi-master MariaDB
- File metadata shared using Cassandra replicated DB
- Separate iRods zone in at each portal instance

WAN to  
other Portal  
Zones

# DELIV: Testbed

---



- Deploying :
  - portal tools for exploration by scientists
  - data movement/management tools to gain understanding of current solutions
    - Aim to integrate best tools in later stages
- Currently running PI tools (federated CyberSKA) in Calgary, at Rackforce in Kelowna and at CHPC in Cape Town.
- Starting testing data movement tools in Calgary, Oxford and at CHPC
- Plan to deploy PI tools at iVec Nov. 2014
- Started discussions with Groningen – host of LOFAR data

# Summary

---



- SKA will produce huge amount of data
- SDP.DELIV defining requirements and architecture for SKA data access by users
- Considering tools and best practices from existing distributed big data projects
- Deploying testbed to gain feedback on interface requirements and explore tools



# Contact Information



Presenter: Rob Simmonds

E-mail: [rwjs@ucalgary.ca](mailto:rwjs@ucalgary.ca)

A screenshot of the CYBERSKA website. The page has a dark background with a starry sky. At the top left, the word 'CYBERSKA' is written in a large, blue, sans-serif font. Below it, a tagline reads 'A Cyberinfrastructure platform to meet the needs of data intensive radio astronomy on route to the SKA'. A navigation bar contains links for 'Home', 'Tools', and 'About'. On the right side of the navigation bar, there are input fields for 'Username' and 'Password', a 'Login' button, and a 'Remember me' checkbox. Below the navigation bar, there is a large image of a desert landscape filled with many white radio telescope dishes. The text 'SQUARE KILOMETRE ARRAY' is overlaid on the bottom of the image, with the subtitle 'Creating the cyberinfrastructure to support what will be the largest radio-telescope ever built'. At the bottom right of the image, it says 'Credit: SPDO / Swinburne Astronomy Productions'.