



REGIONAL
CENTRE
NETWORK

SRCNet project update

Ian Collier
ian.collier@stfc.ac.uk

With thanks to Rosie Bolton and Shari Breen
53rd LHCOPN-LHCONE meeting, IHEP, Beijing

9th October 2024



Summary - the SRCNet Project

- **Background**
- Expected Data rates
- Distribution
- The team of teams
- Implementation of SRCNet0.1
 - Locations
 - Potential implications for network




SKA is a flexible science machine!




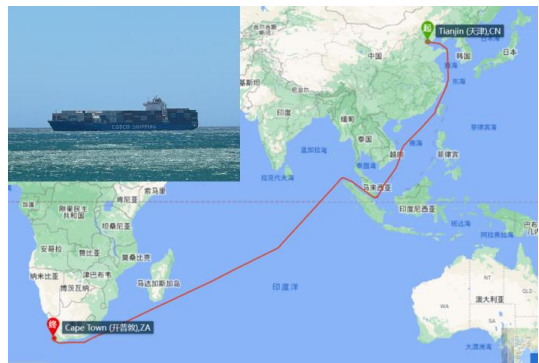
- SKA systems are hugely flexible
 - TWO telescopes to cover a frequency range between 50 MHz and 15.4 GHz
 - Each supports up to 16 subarrays (splitting the 512 stations and 197 dishes into smaller arrays)
 - Very flexible Correlator beam formers (CBFs) but ultimately resource limited
 - Both imaging and non-imaging modes
 - Broad-band continuum, Spectral/zoom, Pulsar and transient search (PSS), pulsar timing (PST), VLBI
 - Commensality supported (data, observing, multiplexed)



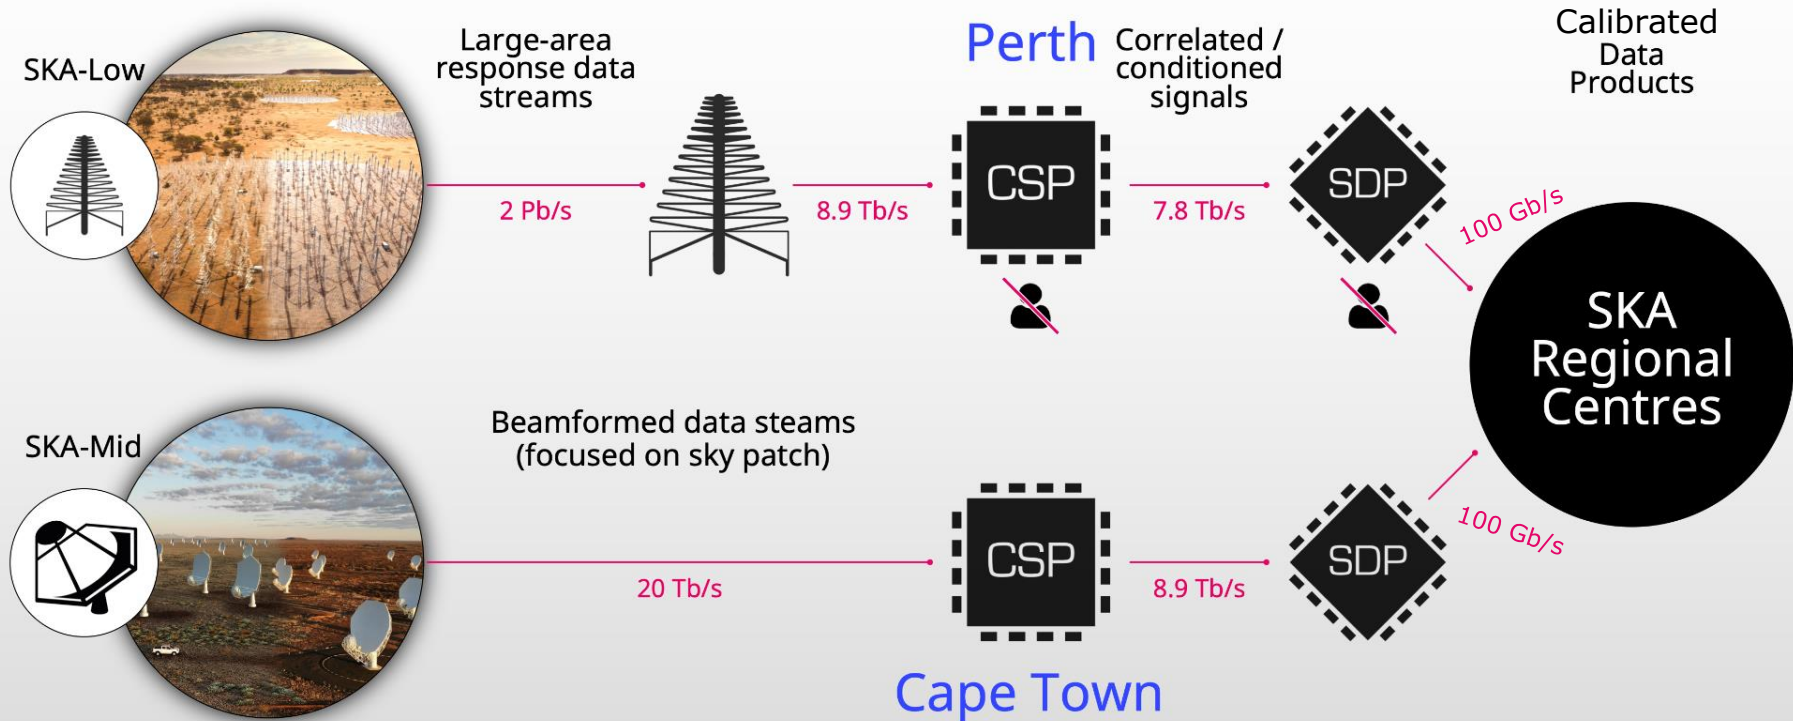
Construction steaming ahead! - low



Construction steaming ahead! - Mid



SKA Regional Centres: SKAO data processing stages



SKA estimated data rates*

*these numbers should be used as a guide only - email Shari.Breen@skao.int for further information about ongoing work

- Numbers refer to data to be delivered to the science community via the SRCNet

Milestone	Year	Primary activity	Estimated data rate	
			Low	Mid
AA2 <ul style="list-style-type: none"> • 64 Mid dishes • 64 Low stations 	2026 - 2027	Science Verification - observed in dedicated ~week long blocks + single observations interspersed throughout. A higher rate of raw data products will be included at this stage.	1.5 PB/week^ 20 Gbps	2 PB/week^ 27 Gbps
AA* <ul style="list-style-type: none"> • 144 Mid dishes • 307 Low stations 	2027 - 2029	Science Verification - observed in dedicated ~week long blocks + single observations interspersed throughout. A higher rate of raw data products will be included at this stage.	5 PB/week^ 66 Gbps	9 PB/week^ 119 Gbps
AA* <ul style="list-style-type: none"> • 144 Mid dishes • 307 Low stations 	2029 +	Operations - Observation cycles, starting with shared risk observing, building to successful science observations ~90% of the time	173 PB/year 44 Gbps	280 PB/year 72 Gbps
Target is to deliver the SKA Baseline Design but the details of this transition between AA* and AA4 are TBD				
AA4 <ul style="list-style-type: none"> • 197 Mid dishes • 512 Low stations 	2030 +	Operations - full SKA baseline design	216 PB/year 55 Gbps	400 PB/year 100 Gbps

^Data rates refer to dedicated Science Verification observing weeks, not an average over a year



SKA estimated data rates*

*these numbers should be used as a guide only - email Shari.Breen@skao.int for further information about ongoing work

- Numbers refer to data to be delivered to the science community via the SRCNet

Milestone	Year	Primary activity	Estimated data rate	
			Low	Mid
AA2 • 64 Mid dishes • 64 Low stations	2026 - 2027	Science Verification - observed in dedicated ~week long blocks + single observations interspersed throughout. A higher rate of raw data products will be included at this stage.	1.5 PB/week^ 20 Gbps	2 PB/week^ 27 Gbps
AA* • 144 Mid dishes • 307 Low stations	2027 - 2029	Science Verification - observed in dedicated ~week long blocks + single observations interspersed throughout. A higher rate of raw data products will be included at this stage.	5 PB/week^ 66 Gbps	9 PB/week^ 119 Gbps
AA* • 144 Mid dishes • 307 Low stations	2029 +	Operations - Observation cycles, starting with shared risk observing, building to successful science observations ~90% of the time	173 PB/year 44 Gbps	280 PB/year 72 Gbps
Target is to deliver the SKA Baseline Design but the details of this transition between AA* and AA4 are TBD				
AA4 • 197 Mid dishes • 512 Low stations	2030 +	Operations - full SKA baseline design	216 PB/year 55 Gbps	400 PB/year 100 Gbps

Aggregate 560PB into entire SRCNet for AA* (1.1EB with second copies)

Aggregate 616PB for AA4 (1.123EB with second copies)

These numbers far from settled



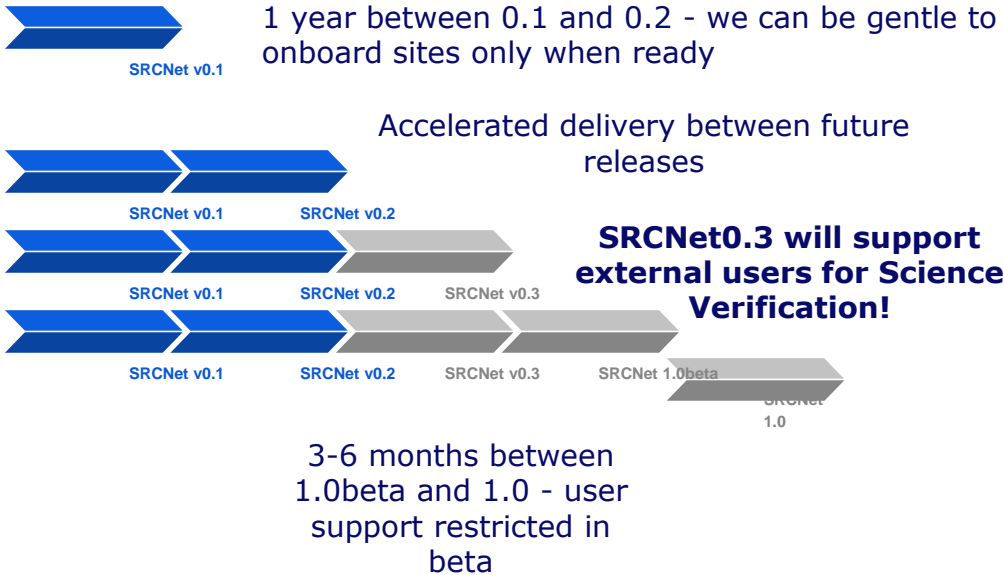
Data Distribution


- Model has data ingested into SRCNet network - and stored on disk at two sites and on tape at one
- Countries are pledging resources in proportion to their contribution to construction
- So e.g. UK takes 16-19% of total (~200PB disk, ~100PB tape per year) - these are VERY approximate figures.
 - Average bandwidth usage to receive 200PB (half from telescopes themselves half from other SRCs) ~60Gbs
 - This (storage and network) is JUST to support basic distribution of data from the instruments.
 - Does not include handling of analysis products
- China, for example) perhaps half those rates
- Affordability is a challenge particularly with recent higher than usual inflation in compute hardware



Staged Delivery and SRCNet releases side by side

Milestone event (earliest)		SKA-Mid (end date)	SKA-Low (end date)
AA0.5	4 dishes 6 stations	2025 May	2024 Nov
AA1	8 dishes 18 stations	2026 June	2025 Nov
AA2	64 dishes 64 stations	2027 May	2026 Dec
AA*	144 dishes 307 stations	2028 Mar	2028 Mar
Operations Readiness Review		2028 Jun	2028 Apr
AA4	197 dishes 512 stations	TBD	TBD

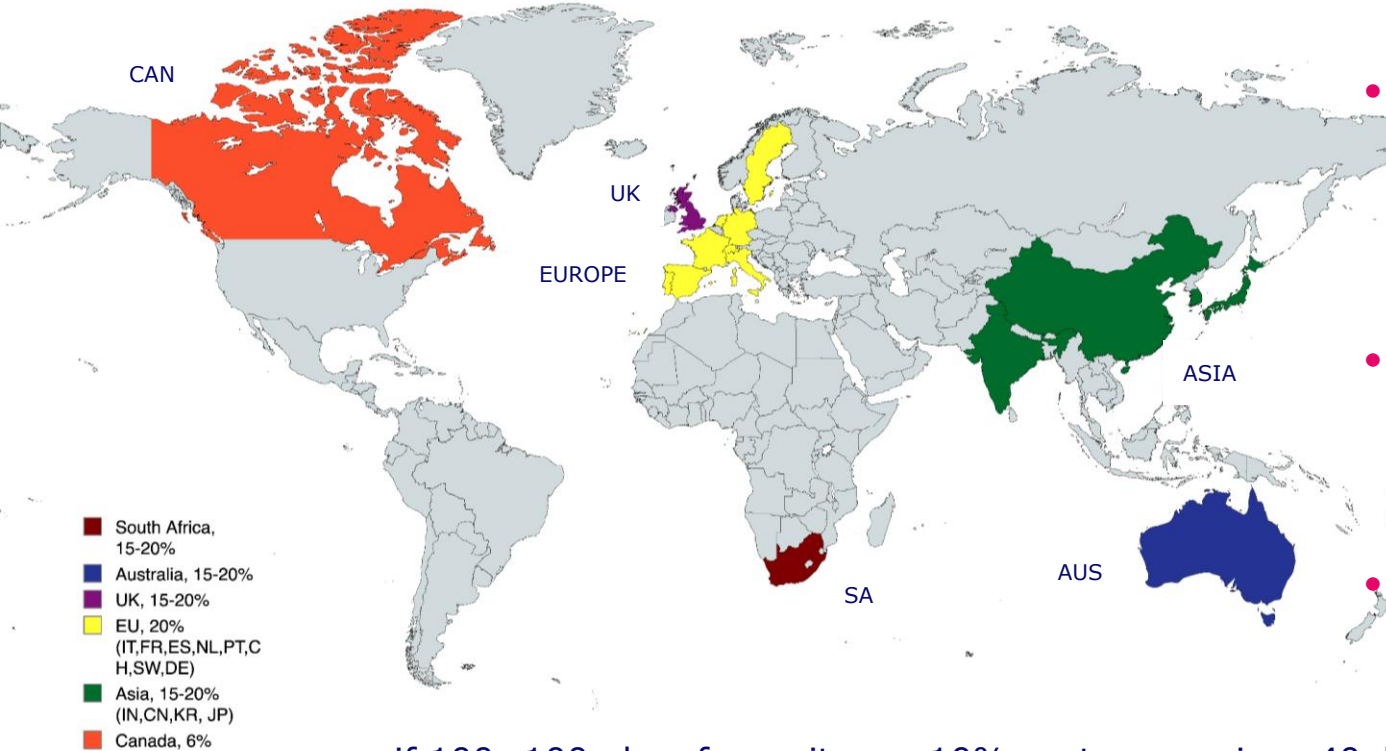




- SKAO Partnership - includes SKAO Member States* and SKAO Observers (as of April 2023)
-  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *
 -  *

- African Partner Countries
- 
 - 
 - 
 - 
 - 
 - 
 - 
 - 

SKA Regional Centre Broad Distribution: Fair Share, AA4 data rates

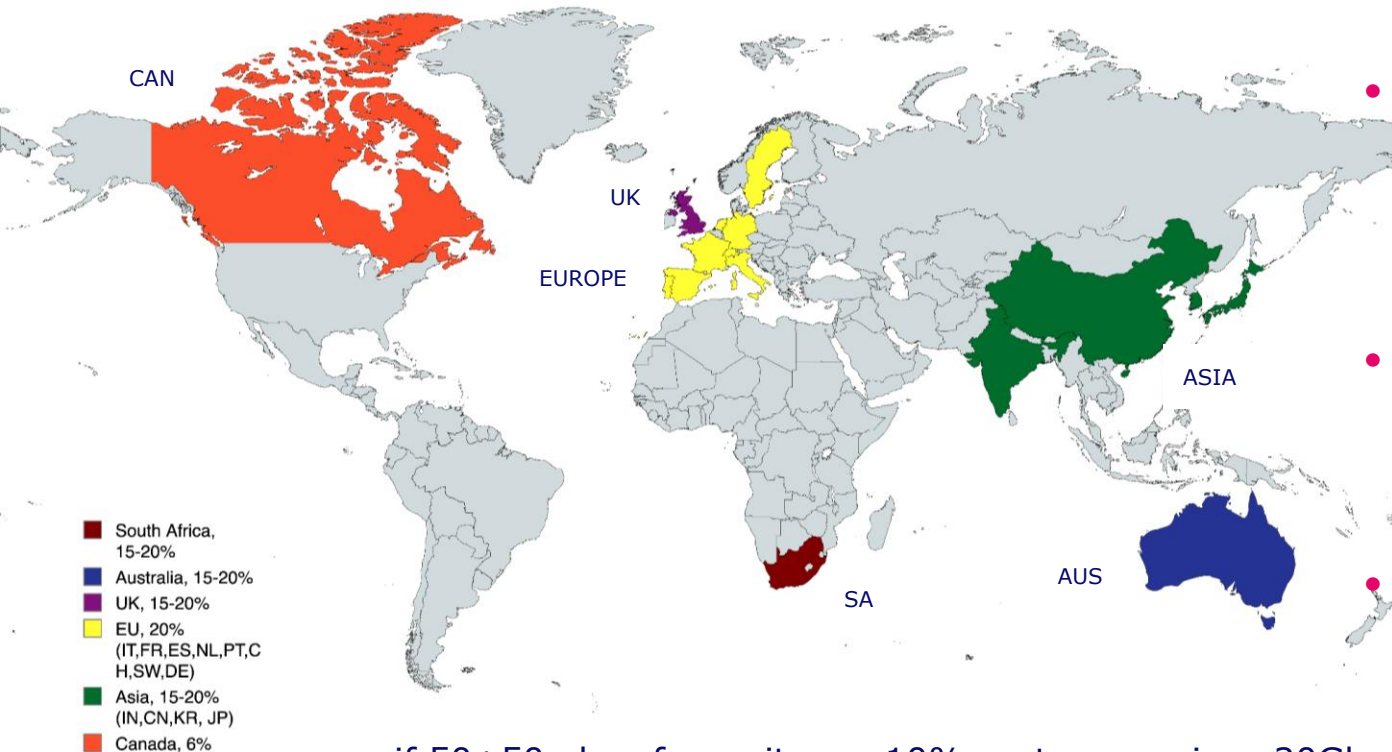


- Roughly, 6 global zones of equivalent size (Canada smaller)
- **Distribute two base copies** of each data product to different countries, and perhaps insist to different regions
- Average incoming rate per (20%) region not more than 2x40 Gbit/s = 80Gbit/s (~2x12 Gbit/s for Canada)
- **Modelling assumes average 100 Gbit/s out of SA and AUS**

e.g. if 100+100 gbps from sites, a 10% partner receives 40gbps data (400 TBytes per day, 140 PBytes per year)



SKA Regional Centre Broad Distribution: Fair Share (if ~50 Gbps per SKAO site)

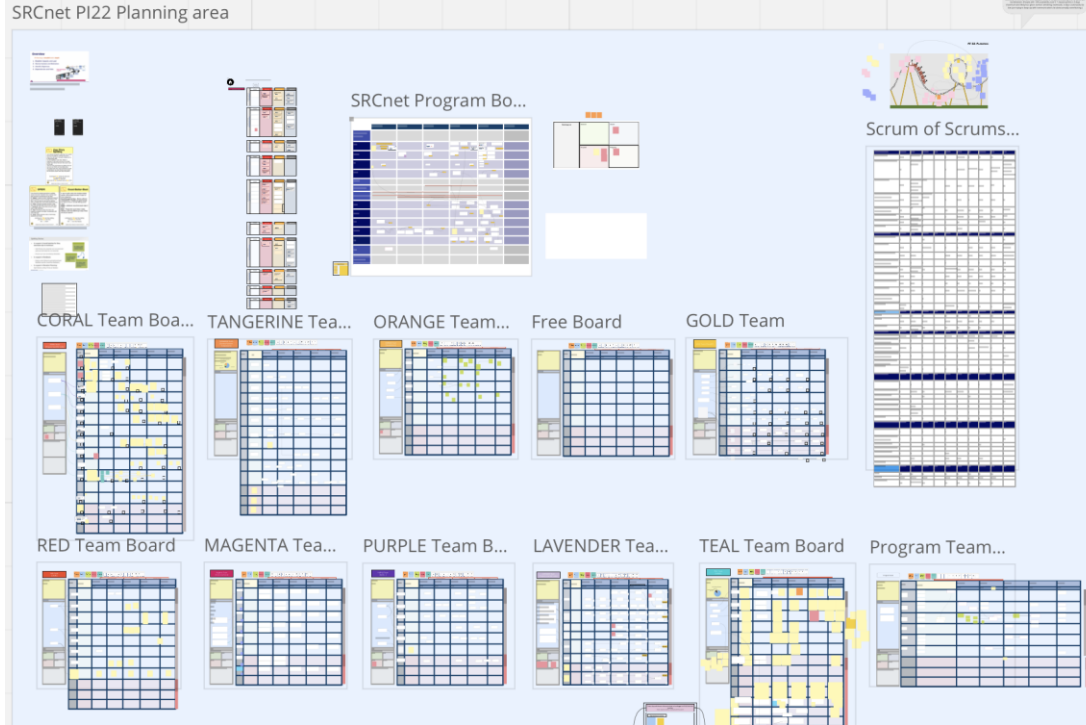


- Roughly, 6 global zones of equivalent size (Canada smaller)
- **Distribute two base copies** of each data product to different countries, and perhaps insist to different regions
- Average incoming rate per (20%) region not more than $2 \times 20 \text{ Gbit/s} = 40 \text{ Gbit/s}$ ($\sim 2 \times 6 \text{ Gbit/s}$ for Canada)
- **Modelling assumes average 100 Gbit/s out of SA and AUS**

e.g. if 50+50 gbps from sites, a 10% partner receives 20Gbps data (200 TBytes per day, 70 PBytes per year)



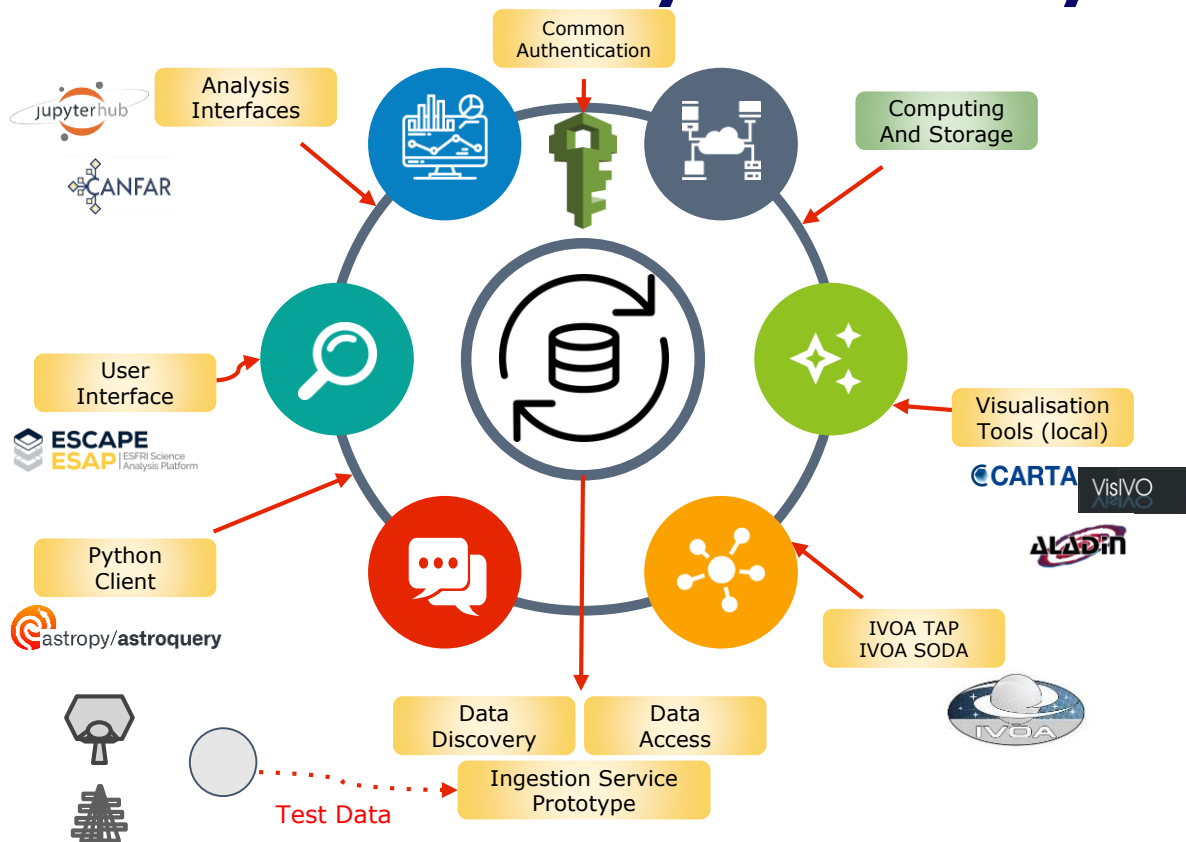
Scaled Agile Framework for SW development



- **Teams Plan together (usually "distributed co-location") once per 3 months**
- **Regular demos (open)**
- **Several Communities of Practice (Identity management, science platform, HPC & Cloud)**
- **Advisory forums**
 - **inc. NREN forum**



Basic Functionality Covered by v0.1

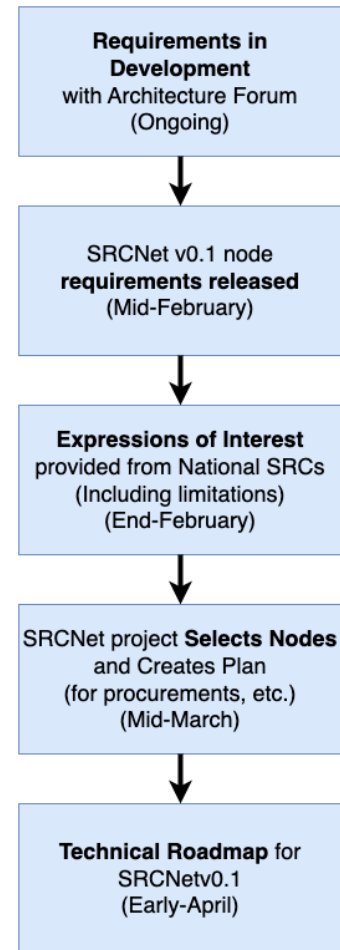


- Common Authentication
 - IAM
- Visualisation Tools (local)
- Data Discovery and Access from Data Lake
- Ingestion Service Prototype
- Python Client
 - Astroquery Module
- User Interface
 - ESAP
 - <https://esap.srcdev.skao.int/>
- Analysis Interfaces
 - JupyterHub
 - CANFAR Science Platform



SRCNet v0.1 deployment

- Deployment of first operational version of v0.1 foreseen for Nov/Dec-2024
- Stable and maintained nodes (target was 4) with enough human and hardware resources
- In the end 9 SRCs participating in first phase
 - Good for inclusivity
 - Somewhat more complex
- Reminder: Focus is on demonstrating global data movement functionality


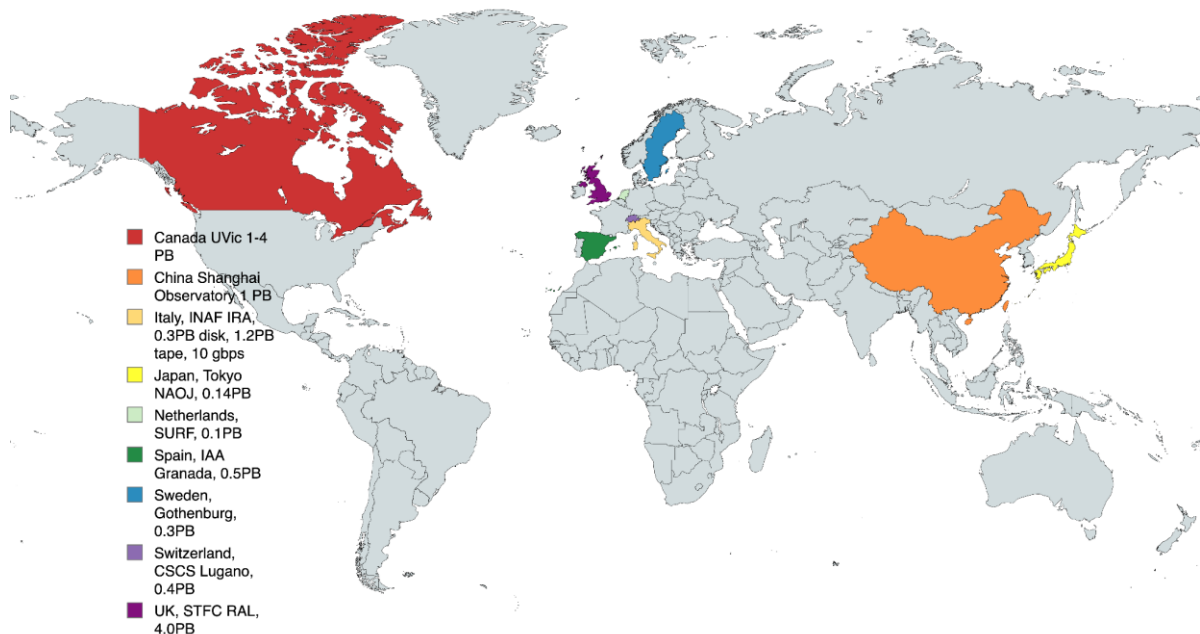


SRCNet0.1 sites

8 PBytes total storage offered for SRCNet0.1 (c.f stated target of 20 PB)

Sites will still have to demonstrate they meet all the resource & operational requirements

Storage (PB)




WLCG experience at some sites (Canada, Netherlands, Sweden, Switzerland, UK)

Some sites quite new and teams will learn by being involved

SRCNet v0.1 Functionality

- Represents first deployment of SRCNet, concentrating on data dissemination, query, movement, and access via 'local' resources

Storage (PB)



Milestone	Description	SRC Net Functionality	Scope (users)
SRCNet v0.1 First quarter of 2025	First version of SRCNet sites deploying common services and connecting via SRCNet APIs. Enable technical tests of the architectural implementation. [Added c.f. document] (Potentially Opportunity to engage SRCNet with AA0.5 data transfer and access.)	<ul style="list-style-type: none"> <u>Test data</u> (and some precursors data) disseminated into a prototype SRC Net <u>Data</u> can be discovered through queries to the SRC Net <u>Data</u> dissemination to SRC nodes <u>Data</u> can be accessed through a prototype data lake <u>Data</u> replication. Data can be moved to a local SRC area where non-connected local interactive analysis portals (notebooks) could allow basic analysis Unified Authentication System for all the SRCs Visualisation of <u>imaging data</u> 	SRC ART members Members of SKA Commissioning team (potentially, but not required)

Considering SRCNet 0.1 use of global data network

Thought experiment (**to trigger conversation only!**)

Imagine SRCNet runs a series of testing challenges, e.g. for a week every 3 months attempt to fill / transfer 20% of offered SRCNet0.1 storage

	Canada SRC	China SRC	Italy SRC	Japan SRC	Netherlands SRC	Spain SRC	Sweden SRC	Switzerland SRC	UK SRC	Total
Data rate if fill 20% storage in 7 days, Gbps	3.2	2.6	0.8	0.4	0.3	1.3	0.8	1.1	10.6	21.0
Current connectivity or usable share	? (but UVic has high connectivity globally, not sure of SRCNet allowance)	1	10	10	100	10	100	10	100	

Based on information from SRC technical leads, all sites could accommodate these tests currently apart from China. Significant share to RAL / UK connections



Automated testing

Screen grab here from the dashboard for the current Rucio prototype

Some science data in Rucio but most dashboard events are from automated test suite (**Rucio task manager**)

- running as a heartbeat currently with all-to-all transfers every few minutes
- Powerful complement lower-level test
- Well suited to running planned data challenges and a key component for SRCNet testing
- Next steps are adding endpoints at each SRCNet 0.1 node and scaling up routine tests



<https://gitlab.com/ska-telescope/src/ska-rucio-task-manager>



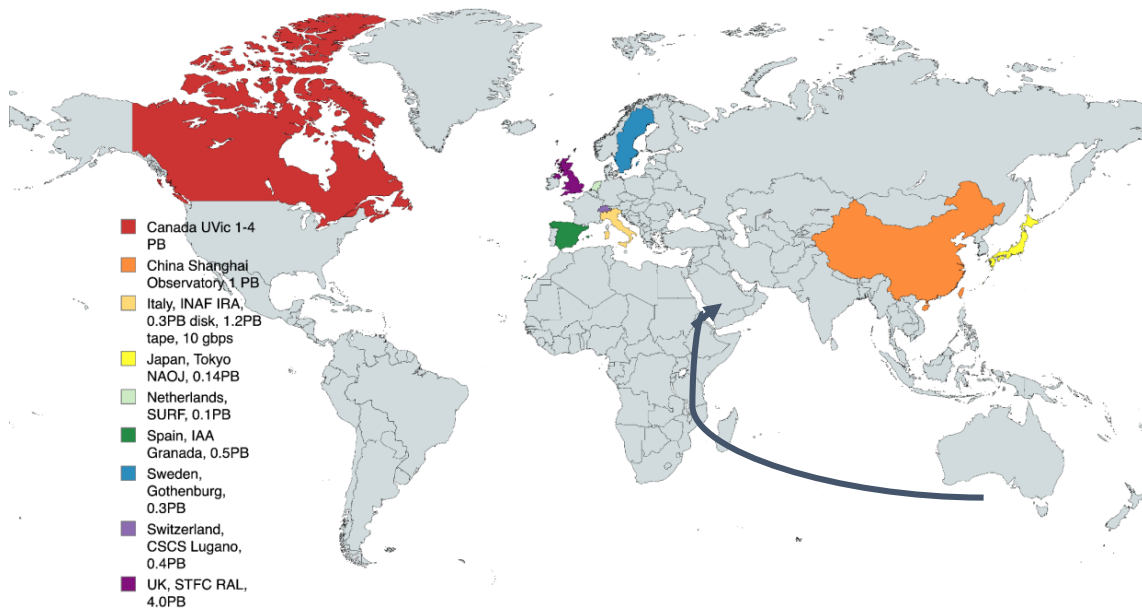
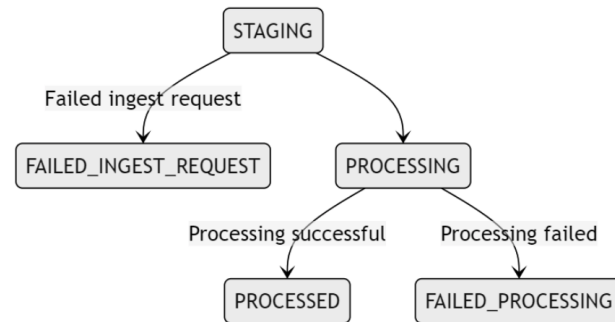
Data Ingestion

Prototype data ingestion service, detects new files and their metadata, adds file to Rucio (triggers replication) and adds metadata to the auxiliary metadata database

In AUS we have a small Rucio storage element and a small team in the SRCNet project.

SRCNet 0.1 will do some testing of long haul data movement/management

<https://gitlab.com/ska-telescope/src/ska-src-ingestion>



Created with mapchart.net



SRCNet0.1 testing

Implementation and testing plans are being developed now

More intensive **data movement challenges** will be a key part of testing the network - and indeed the focus of SRCNet v0.1. (We have been watching WLCG's DC24!)

Planning for 4 periods of focussed data movement testing through calendar 2025 - one per 3 month planning increment



SRCNet Network groups

SRCNet NREN Forum

- Focus on technical activities
- Link testing/transfer tuning
- Implementing Perfsonar mesh
- Monitoring etc.
 - Considering Community of practise

(New) SRCNet Network Strategy Group

- Coordinate discussions with NRENs
- Assuring global connectivity is in place
- Connect architects/CTOs to verify assertion about provision
- Share high level view on plans and requirements
- Coordinate with other communities where appropriate
 - (e.g. continuing to participate in LHCOPN-LHCONE meetings - many overlapping interests)



What's next? SRCNet v0.2 ...

- v0.1 is just the beginning
- Increased functionality and scaling up in future iterations
- For UK, some of this functionality needed / tested earlier
- Science Verification from telescopes =>
~ 1.5 – 2 PB / wk from each telescope for dedicated weeks.
- 2-4x Storage across the SRCNet

SRCNet v0.2 First quarter 2026	AA1 and Commissioning	<ul style="list-style-type: none">● Test data (and some precursors data) disseminated into a prototype SRCNet● Data can be discovered through queries to the SRCNet● Data can be accessed through a prototype data lake● Data replication during processing● Unified Authentication and Authorisation system● Data dissemination using telescopes sites interface● First version of federated execution. Access to remote operations on data using services and the possibility to invoke execution into a relevant SRC● Subset of SDP workflows runnable in the SRCs● First Accounting model implementation. Storage● User storage areas● Visualisation of imaging and time series data through remote operations● Preparation of SRCNet User Support	Selected scientists from community Science Operations and Commissioning teams SRC ART members
---	--------------------------	---	---

Summary

Moving to first operational phase

- Establishing infrastructure management and operations teams across partner sites
- Beginning to procure and deploy and compute and storage resources
- Establishing security policies etc.
- Development work moving on to next phase

Looking forward to

- Significant developments in technical capabilities at SRCNet nodes and across network
- New formal technical roadmap flowing from SKAO's revised deployment roadmap
 - We know 2023 technical roadmap is out of date
- Improving understanding of & monitoring of network between our sites
- Coordination with other communities
 - For example during upcoming WLCG data challenges



Questions / discussion

