

The SRCNet Project

Shari Breen Rosie Bolton
shari.breen@skao.int rosie.bolton@skao.int

Document classification: Public



“This image was taken using the first four completed antenna stations at Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory. Produced using only 1,024 of the planned 131,072 antennas – less than one per cent of the full telescope – it shows an area of sky equivalent to approximately 100 full moons. 85 of the brightest known galaxies in the region can be seen. It’s calculated that the completed SKA-Low will eventually be sensitive enough to show more than 600,000 galaxies in the same frame.”

Announced just two days ago!!<https://www.skao.int/en/news/621/ska-low-first-glimpse-universe>



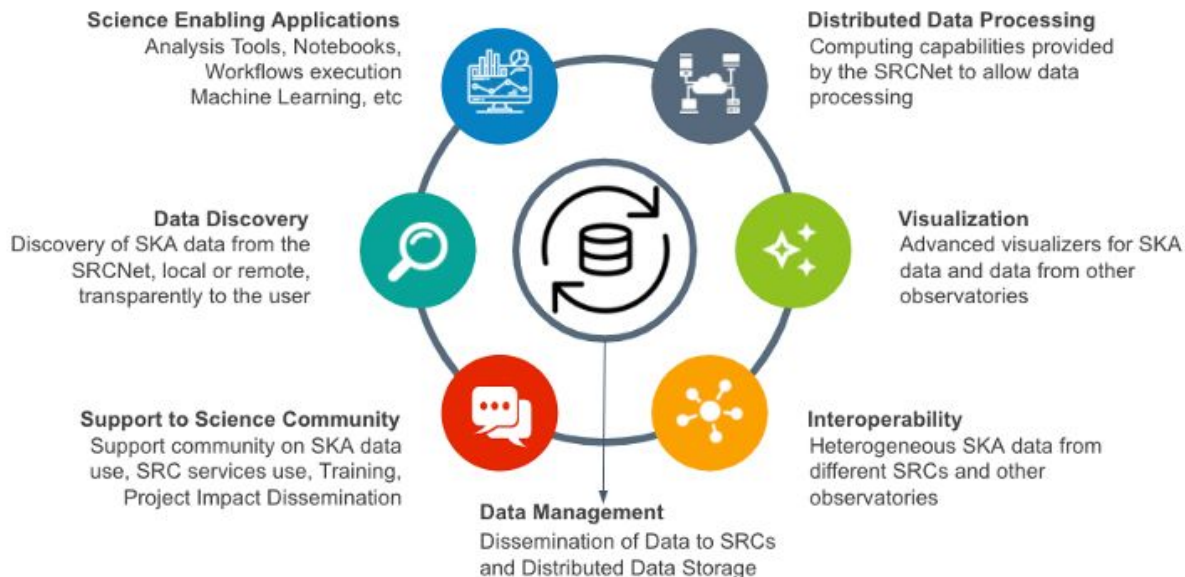
What does the SRCNet need to deliver to support SKA science, and on what timescale?

Shari Breen
Head of Science Operations, SKAO



SRC Network functionality is critical

Delivering SKA data products to scientists, storing SKA data for future use, computer facilities to undertake scientific analysis and local user support all fall outside of the construction budget

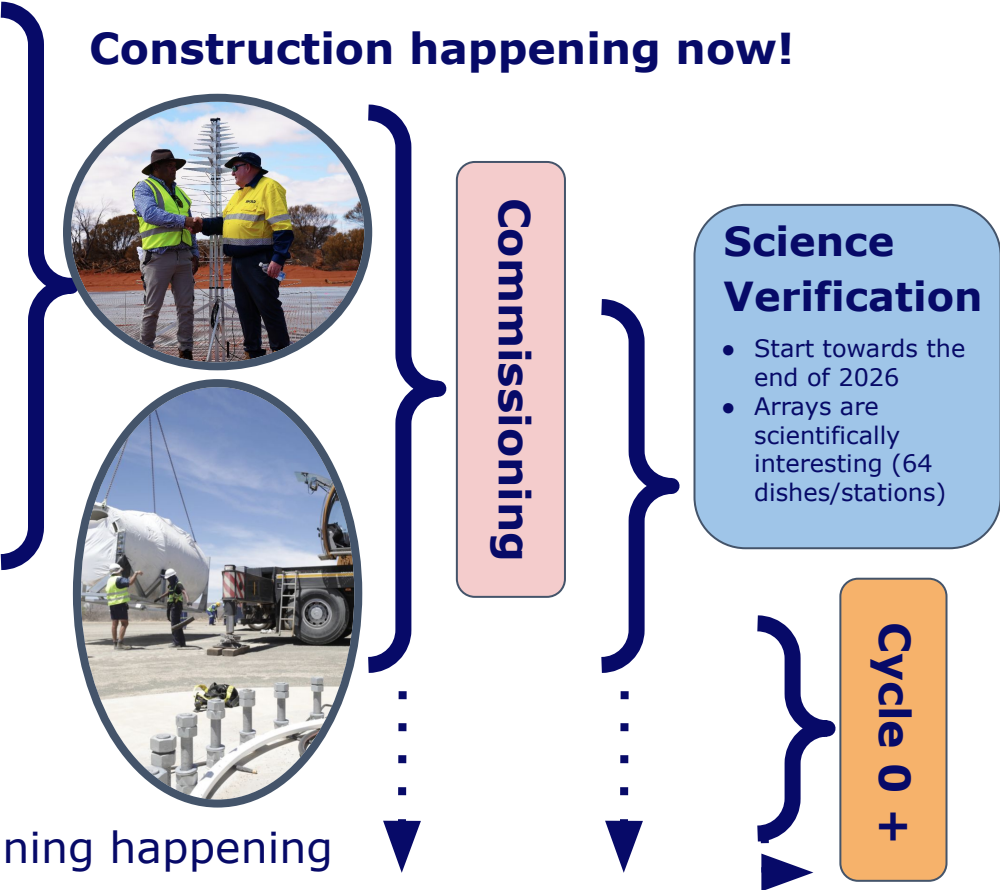


Many of these functions, or at least their scale, is hugely dependant on the SKAO delivery plans → the roadmaps need to be closely aligned!!



At the highest level this is the timeline to the SKA

Milestone event (earliest)		SKA-Mid (end date)	SKA-Low (end date)
AA0.5	4 dishes 4 stations	2026 May	2025 Jul
AA1	8 dishes 18 stations	2027 Jan	2026 Jan
AA2	64 dishes 64 stations	2027 Dec	2026 Nov
AA*	144 dishes 307 stations	2028 Sep	2028 May
Operations Readiness Review		2029 Jan	2028 Jul
Formal end of construction (including contingency)		2029 Mar	
AA4	197 dishes 512 stations	TBD	TBD



...but there is lots of additional planning happening



Year in the Life of the SKA

- What does a year of Operations look like (“standard” year - cycle 2?)
- **Assumes the telescope we are planning to deliver in AA* (in line with the requirements, rollout schedule, pipeline availability [perhaps not always usage]...)**
- Covers as much of the system as we can reasonably include
- **Aiming to be representative not perfect**
- **Why now?**
 - We now know enough detail about systems, science plans etc to develop a quite sophisticated and accurate representation of what a year might look like in full operations
- **Why at all?**
 - operational planning, development of systems, confidence that we can deliver such a science program? Refining to something we can deliver. Setting user expectations.
- **What?**
 - A series of documents. The core being the project lists (and their analysis) which underpin everything. The statistics from these are important inputs for SDP and SRCNet planning. → **this is our best guess at what science we will do!!**



How does this help us plan what SRCNet needs?

- Tells us a lot about what we need to have available in steady state (2030+)
 - Data sizes (~130 PB for Low and ~170 PB for Mid with the currently envisaged AA*)
 - Distribution of types of data products and the kinds of Advanced Data Products our users might want to create on the SRCNet
 - What software and tools might we need to deploy?
 - What workflows might we need to run?
 - What workflow templates might we need to create?
 - How many projects/users might we need to support
 -
- **This is great, but a view 5 years from now is not super useful to guide what needs to be in place in a year or two!**



Year in the Life of the SKA

SKAO-TEL-0000000 Revision 01
Classification: FOR PROJECT USE ONLY
Document type: NOT
Date: 2025-02-10
Status: DRAFT

Role	Name	Designation	Affiliation	Signature	Date
Author	Science Operations		SKAO	_____	
Owner	Shari Breen	Head of Science Operations	SKAO	_____	
Approver	Antonio Chrysostomou	Deputy Director of Operations	SKAO	_____	
Released by	Lewis Ball	Director of Operations	SKAO	_____	



Science delivery value stream

<https://confluence.skatelescope.org/display/SRCSC/SRCNet+Science+Delivery>

- Purpose: Science delivery coordinates, supports and enables the delivery of science via the SRCNet in tandem with the SKAO
 - Needs to represent the needs of the SKA user community and align with the SKAO development
- Three main threads:

User representation

- Ensure that the SRCNet will meet the needs of the astronomer end user
- User flows and requirements linked to the different phases (SV, early cycles etc)

Workflows to support science needs

- What are the workflows needed to support science cases from SV →?
- SKAO pipelines? Efficient user generated?
- Prioritised executable, testable precursor workflows

User support


- What support does the SRCNet need to offer users and when?
- One Helpdesk
- Efficient collaboration with SKAO



Science Delivery value stream - who (at least for now)?



Shari Breen
BO



?
Science Delivery
Lead/Product
Manager



Nichol Cunningham
SRCNet SciOps Lead

Project scientists

- Define user workflows through the scientific milestones
- Requirement generation and validation
- First "real" users

Team formally sits outside the ART

"Support" scientists

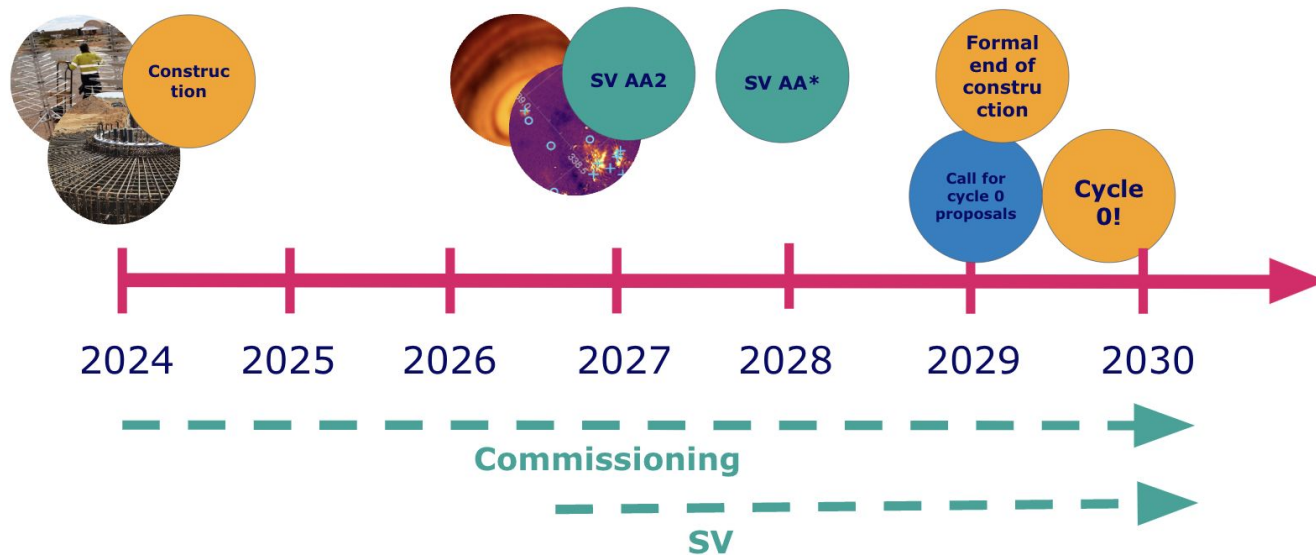
- Develop, adapt, execute workflows as needed through the Scientific milestones
- Technical investigations to inform plans
- Develop workflow templates

Super teal + a few individuals scattered throughout teams



What science do we need to support when?

- Plan to support *real* SRCNet users with *real* SKA data from the end of 2026 in Science Verification
- Modes and data products will be built up over time - don't need to be ready for everything on day 1
- Can define an MPS for the different science milestones!



Science Verification

Science Verification is a full end-to-end test of the system; a “dress rehearsal”, starting with a submitted idea, lightweight assessment, observation, Observatory Data Product creation, quality assurance, data delivery into SRCNet, user access, analysis, Advanced Data Product creation, user support

Science Verification is more effective if we closely mimic the the way SKA users will experience the end-to-end system during observing cycles.

- Early scientific exploitation, allowing users to understand the potential of the SKA and their experience of it as a user
- Get specific feedback about SKAO's scientific capabilities and data quality, including the validation of data reduction workflows, from experts
- Get specific feedback on user experience of the SRCNet - are the right tools available, does the system work?
- Feedback on our policies and their implementation (e.g. user authentication, user support access)
- Public data delivery allows us broad representation of the community in feedback (i.e. reach a broad range of users with representative variety in the things they consider important)



How are we going to meet the needs of the user at SV?

1. Project scientists have taken the expected Observatory capabilities and ODPs in SV
2. Define detailed science cases/investigations
3. Convert into user workflows
4. Define a complete set of user requirements, which evolve in step with the planned SKAO delivery
 - **Defines the user needs in the MPS at AA2+**
 - Aiming to have the full set for SV AA2 and AA* workflows and requirements at the end of PI26
 - **Prioritised roadmap of what capabilities, software, workflows need to be available on what timescale**
 - → **drive/focus the development work of the Support Scientists**

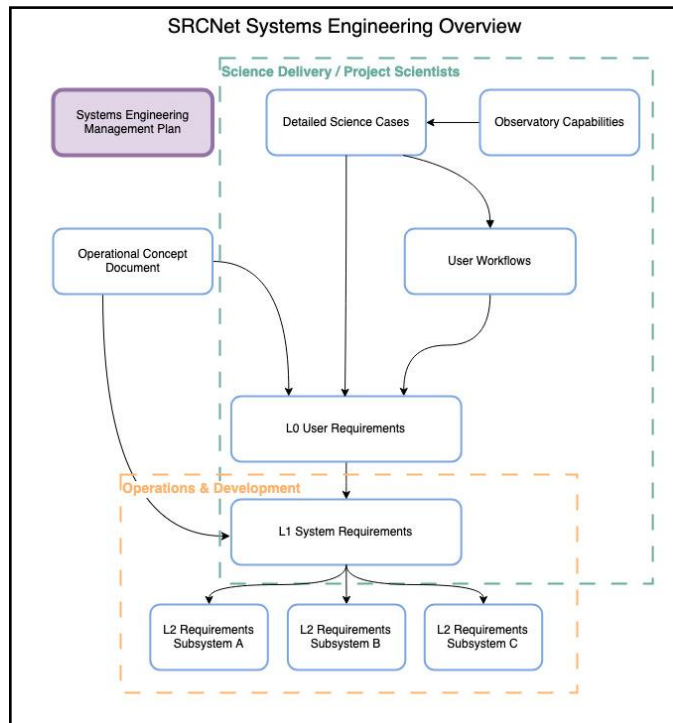


Image credit to Toby Brown!



Updated - SKA expected 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
- Data rates calculated assuming a network uptime of 90% with a network overhead of 20%

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 [^] 29 Gbps	2 PB/week [^] 39 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 [^] 96 Gbps	9 PB/week [^] 174 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	130 PB/year 48 Gbps (1 to 210 Gbps range)	170 PB/year 63 Gbps (0.05 to 300 Gbps range)
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 80 Gbps	400 PB/year 148 Gbps

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



SRCNet project update

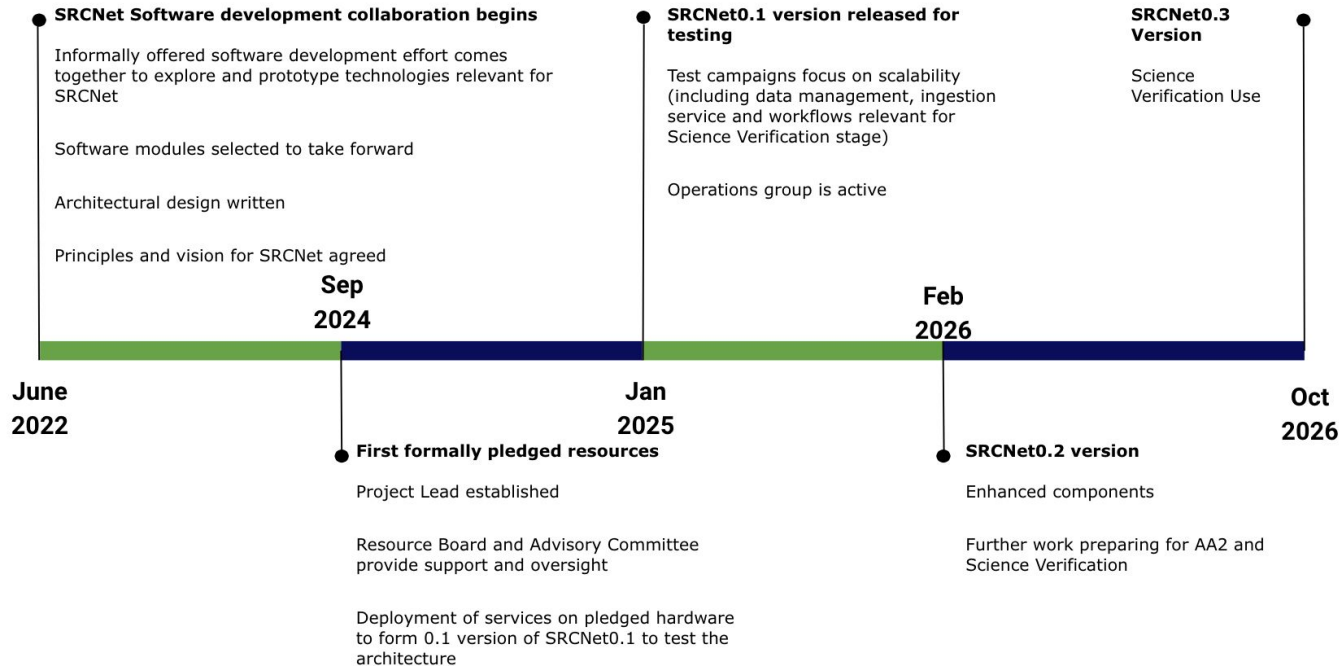
Rosie Bolton
SRCNet Project Lead



SRCNet0.1

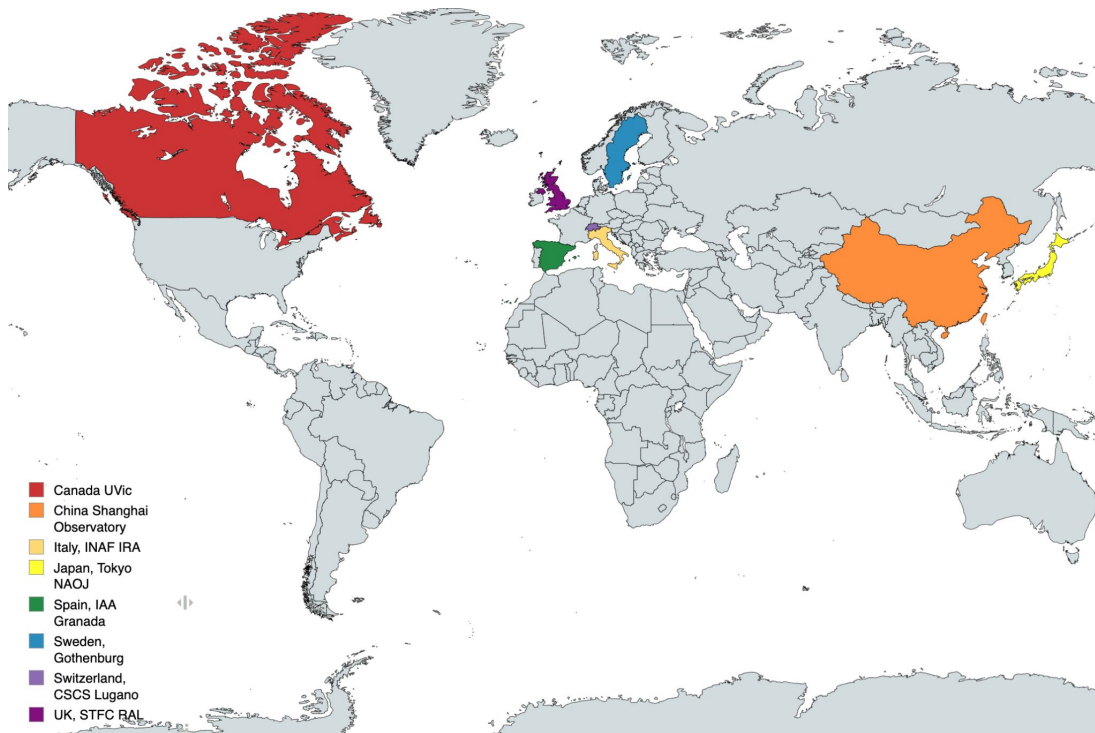
This is an "engineering" version

- **Built to show the architecture and test how it works**
- **Internal only** - no user-facing activities
- Exclusive storage to use in testing
- Compute to use during testing campaigns (may be backfilled when idle)
- Learn how to deploy and operate the services
- Set up of the SRC Operations Group, with limited scope




SRCNet0.1 planned sites


- 8 sites will contribute compute and storage resources to SRCNet0.1, for 2025 testing period
- 7.8 PBytes total storage; 0.5 PFLOPS Compute




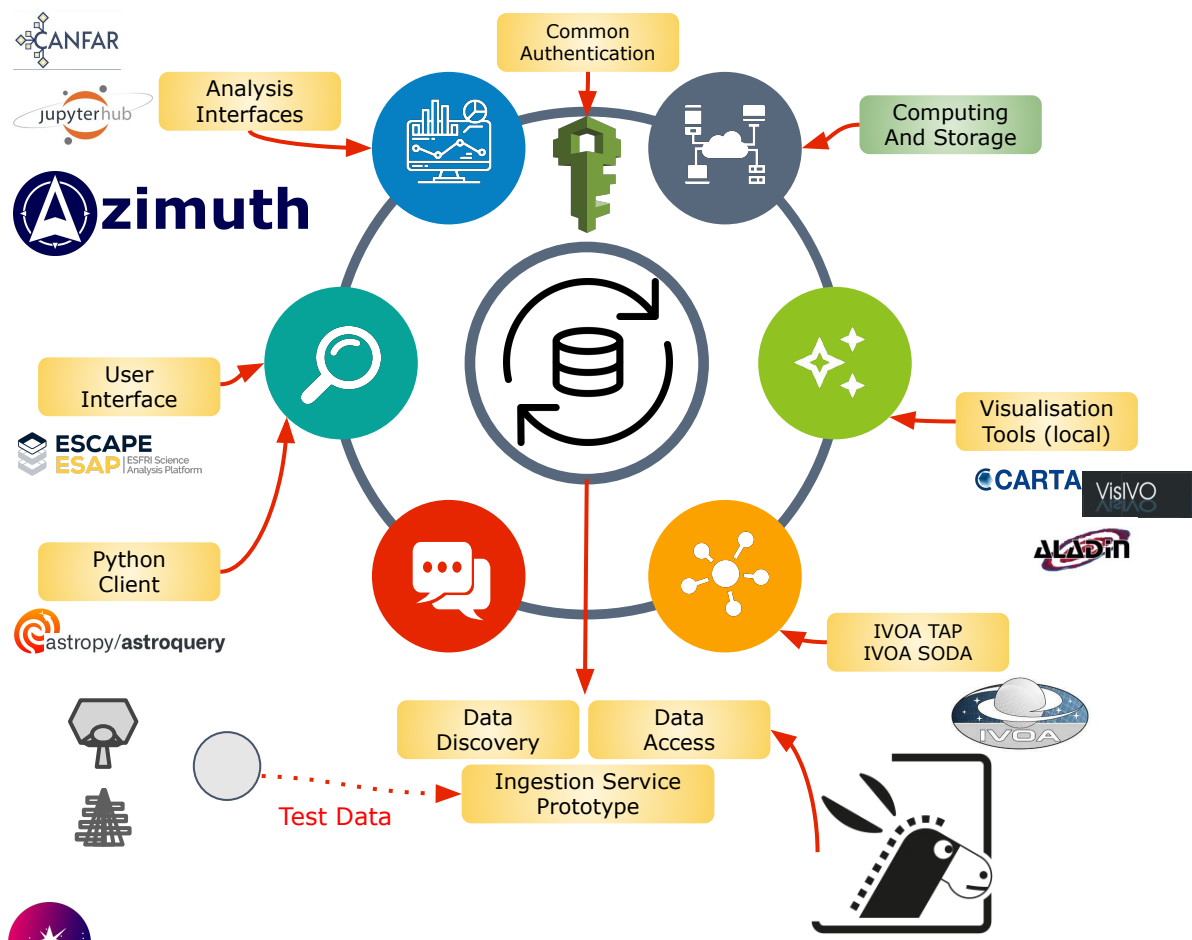
Storage fraction for SRCNet0.1



Compute fraction for SRCNet0.1



Basic Functionality Covered by v0.1



- Common Authentication: IAM
- Visualisation Tools (local)
- IVOA Protocols: TAP, SODA
- Data Discovery and Access from (Rucio)Data Lake
- Ingestion Service Prototype
- Python Client: Astroquery Module
- User Interface: Gateway
<https://gateway.srcdev.skao.int/>
- Analysis Interfaces: JupyterHub (compulsory); CANFAR Science Platform, Azimuth (UK)



Entering the 0.1 Phase



At least 4 nodes with:

- **All services deployed and integrated** (*passing Jesús' test*) - (*3 done. expertise proven, pending*)
- Security point of contact in place (*OK*)
- At least 0.5 FTE to support ongoing operations (*OK*)
- **Pledged resources** (*storage and compute*) to *guarantee access to storage and services during the test campaigns (some gaps, pending)*
- **Support for the OLAs** (*progressing well but pending*)



SRCNet 0.1 progress

8 nodes now intending to participate as nodes in test campaigns. NL stepped back from hosting a node.



Spain, Switzerland and SKAO integration nodes complete

Sweden and UK looking very good

Global Services Status	
Summary	Readiness State %
12 Services	100%

Global Services Status	
Summary	Readiness State %
12 Services	100%

Site Readiness Status (Integration with global services)	
Summary	Readiness State %
espSRC	100%
sweSRC	29%
uksRC	94%
chSRC	88%
cnSRC	12%
canSRC	47%
jpSRC	29%
itsRC	29%

Local Compulsory Services (Deployed)	
Summary	Deployment State %
espSRC	100%
sweSRC	90%
uksRC	90%
chSRC	100%
cnSRC	60%
canSRC	40%
jpSRC	60%
itsRC	70%

Previous plan was to complete this work before December 2024. Now 12 weeks late.

Policies have been redrafted, Operator Level Agreements - agreed ways of working.

Including simple data rules - in SRCNet 0.1 all data used must be public.



Node status (as at 13/3/2025) *(Incident response timescales exception)

Node	% compulsory Services deployed	Integration test passed	Service Operations effort	Security point of contact	OLA endorsement	Pledged Storage <i>and</i> Compute
canSRC	40%		✓	✓	✓	✓
chSRC	✓ 100%	✓	✓	✓	✓	✓
cnSRC	90%		✓	✓	✓	✓
espSRC	✓ 100%	✓	✓	✓	✓	✓
itSRC	70%		✓	✓	✓ *	✓
jpSRC	70%		✓	✓	✓	✓
sweSRC	90%		✓	✓	✓	✓
ukSRC	90%	(happening now!)	✓	✓		
SKAO	✓ 100%	✓	✓	✓	✓	✓

Test Campaigns (James will talk about this!)

Test Coordinator: James Walder (UKSRC)
SRCNet Security officer: Debbie McLachrie (SKAO)



James Walder (UKSRC)

"Engineering" phase

"Experience and function" phase

Ingest and Dissemination

Integration, User Workflows, and Data Lifecycle

Apr/May 2025

June/July 2025

Documenting Results

Sep 2025



Apr 2025

May 2025

July 2025

Network and Storage Performance Dress Rehearsal

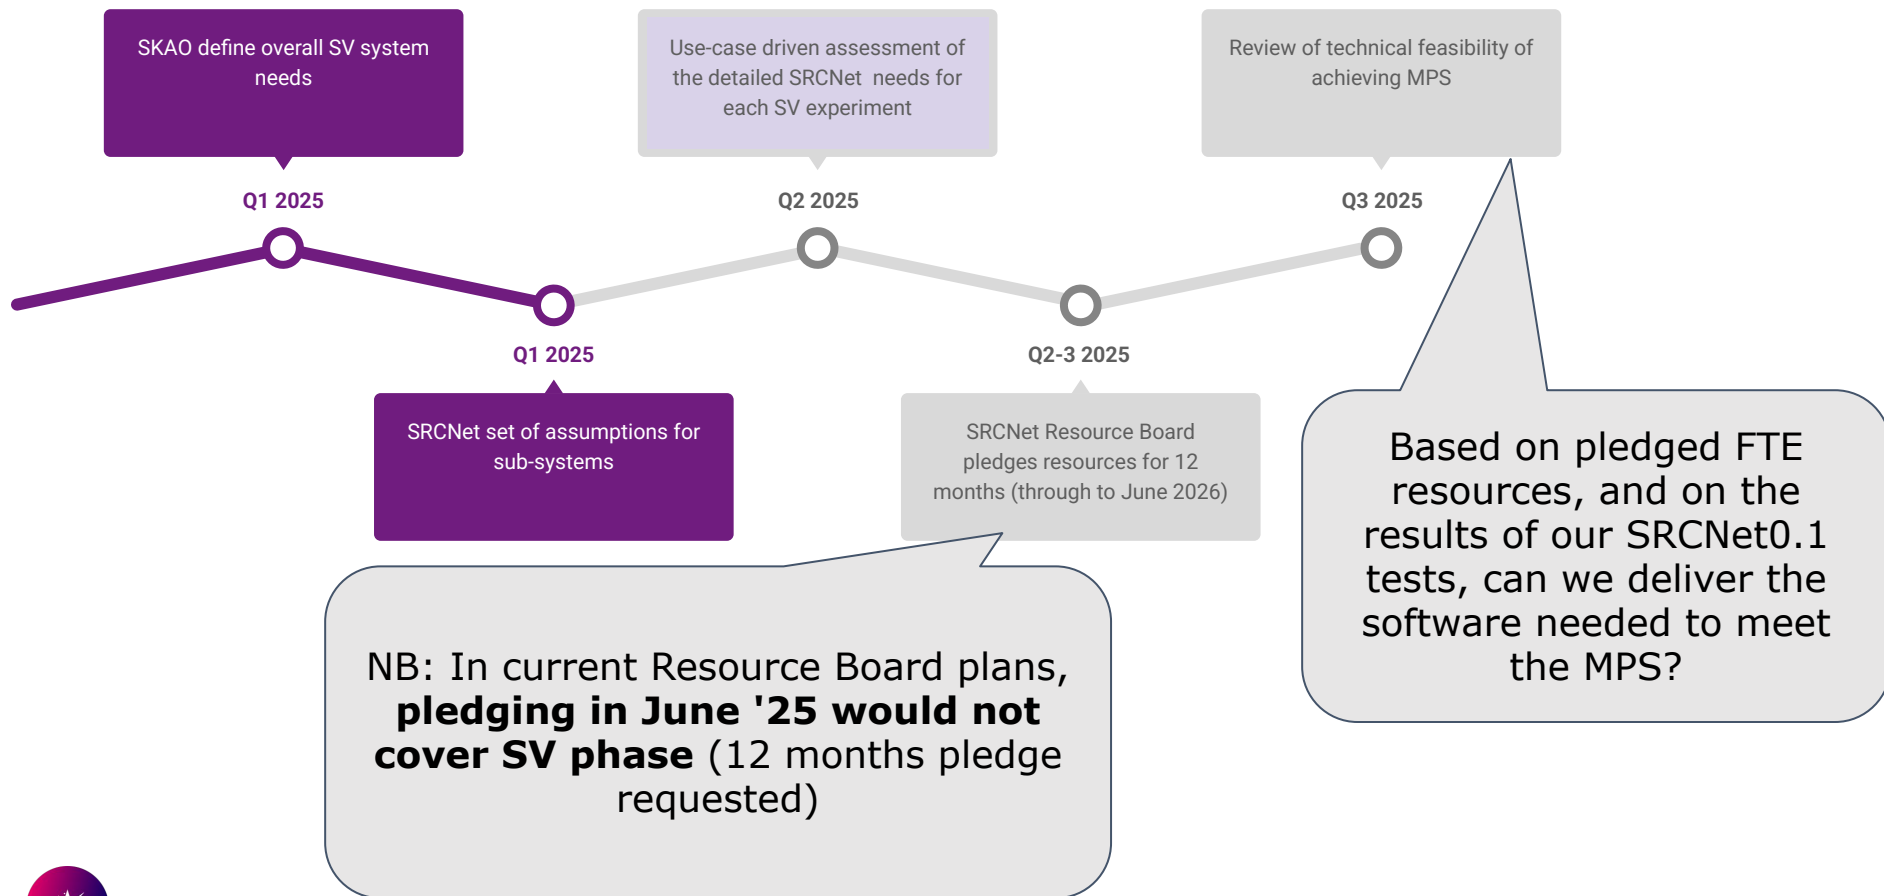
Network and Storage Performance

Scalability and Stress Tests

Only accredited nodes will be involved in the test campaigns



Timeline to MPS review



What comes next?

- Early 2026 - "0.2" version for software, significant scale up of hardware, demonstration of readiness to support Science Verification use
- Late 2026 - "0.3" version / MPS - Integration with SKAO data and users
- Currently, we need to achieve about 100x scale up from now (Q1 2025 to Q3 2028)...

Version	Approximate equivalent expected SKAO status	SRCNet Project Target date	Release date	Status	Access
0.1	-	December 2024	(expect March 2025)	Pending	SRCNet project members
0.2	-	February 2026			Selected scientists
0.3	AA2 Science Verification phase (due Nov 2026 Low)	October 2026			Science verification community (public access)
1.0 Beta		December 2027			Increased Cycle 0 scientists
1.0	AA*, Operations Readiness Review (due July 2028 Low, Jan 2029 Mid)	July 2028			Public

