

SKA Regional Centre Global and Local Network Update

Ian Collier
ian.collier@stfc.ac.uk

GNA-G Meeting
TNC 25

13th June 2025



SRCNet Global and Local Network Provisioning

- **Background**
- Recap SRC Distribution & GREN
- Estimating data volumes and bandwidths
- First data movement tests
- Summary




Square Kilometer Array: Transforming radio astronomy


The Square Kilometer Array (SKA) Observatory (SKAO) is a next-generation radio astronomy facility which will cover the frequency range from 50 MHz to 15 GHz.



Composite image of the SKA telescopes, blending real hardware already on site with artist's impressions. credit: SKA Observatory



Credit: SKA Observatory



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


















* * * * *

African Partner Countries












Regional Centre Nodes & distribution

- Current modelling assumes a Regional Centre share of Resources proportionate to their construction share:
 - ~ 6 similar Global Regions (Canada smaller)
- Australia, South Africa, and the UK dominate as single country entities



The Global Research & Education Network



Latest SKAO Expected data rates

- Until recently $\sim 300\text{PB/year}$ by 2029
 - With mean datarates $\sim 50\text{-}100\text{gbs}$ from each of Australia & South Africa
- New extrapolation from updated construction timelines suggest that will arrive in 2035
- 2029 expecting 25PB of new data
 - Much, much lower mean utilization and in turn much lower load on global R&E network than earlier projections



Latest SKAO Expected storage volumes

- Updated expected storage growth:


- 2029

- 110PB disk
- ~25PB from WA
- ~10gbs average

- 2035

- 1600PB disk
- Archive + user
- ~50gps each

Total Request (PBytes)



New SKAO Expected storage volumes

- Updated expected storage growth:

- 2029


- 110PB disk
- ~25PB from WA
- ~10gbs average

- 2035

- 1600PB disk
- Archive + user
- ~50gps each

New estimates from last week - significantly reduced numbers!

Total Request (PBytes)



Estimating SRC data volumes and bandwidths

- Simple Notebook created to model data rates:
- Telescope volume as inputs
- Data replication parameters
- Network overheads / retries / peak capacity
- Node size (%)
- **Example:**
 - 300 PB telescope output
 - ~ 180 PB of disk required for a 20% share Node
- Peak Ingress rates of O(200)Gb/s may need to be supported
 - ~ O(100) Gb/s for Egress
- This does not account for workflow and intra-node traffic

	Input
Annual Disk Volume [PB]	900.00
ADP size fraction	0.50
Annual increase fraction	0.00
Year count	0.00
Number of Replicas	2.00
Data Consolidation / retry factor	1.50
Network overheads factor	1.25
Peak capacity factor	2.00
Network Bandwidth factor	2.00

Share [%]	1	2	5	10	15	20
Disk [PB]	9.0	18.0	45.0	90.0	135.0	180.0
Ingress Data [PB]	9.0	18.0	45.0	90.0	135.0	180.0
Egress Data [PB]	4.5	9.0	22.5	45.0	67.5	90.0
Ingress Rate [Gb/s]	4.3	8.6	21.4	42.8	64.2	85.6
Ingress Peak Rate [Gb/s]	8.6	17.1	42.8	85.6	128.4	171.2
Ingress Peak Bandwidth [Gb/s]	17.1	34.2	85.6	171.2	256.8	342.5
Egress Rate [Gb/s]	2.1	4.3	10.7	21.4	32.1	42.8
Egress Peak Rate [Gb/s]	4.3	8.6	21.4	42.8	64.2	85.6
Egress Peak Bandwidth [Gb/s]	8.6	17.1	42.8	85.6	128.4	171.2

The combined expected volume of data output from both telescope sites over one year.
This is assumed over one year of continuous average output.

Annual data [PB]

Increase of data (fraction) per year (compounded).
This assumes the total data on disk is increasing slowly over time.

Annual increase in volume fact...

Number of years from the start.
This allows for extrapolation into future years from a given starting point.

Year increment

The average fraction of ADP compared to telescope input.
Assumed average ADP size relative to total telescope data.

ADP size factor

The number of replicas that SRCNet should hold on disk.
(Note this includes the primary copy and additional copies of data.)

Replicas (Disk)

Annual expected data on disk.

Annual Archive Growth [PB]

The percentage of data a given SRC should receive.
This is representative of the contribution of the given SRC.

Percent SRC [%]

Expected volume of data retransmission due to failures and data consolidation activities.

Fraction of re-tries

Additional overhead of network operations (e.g. packet headers).
Network overhead (8:10 enco...

Required SRC Node data disk capacity.

SRC Disk [PB]

Estimating SRC data volumes and bandwidths

- From these totals we can estimate data volumes at individual SRCs based on % share of total
- In turn we can also estimate mean network utilisations
- These mean network utilisations in turn much lower
 - 50gbs from each of SA & WA in 3035
 - Much lower in 2029

SRC Input volume vs Share size

Volume [PB] /Share [15%]	5	10	15
100	13	26	39
250	32.5	65	97.5
500	65	130	195

100

Share [%]	1	2	5	10	15	20
Disk [PB]	3	5	13	26	39	52
Ingress Data [PB]	3	5	13	26	39	52
Egress Data [PB]	1	3	6.5	13	20	26
Ingress Rate [Gb/s]	1	3	6.2	12	19	25
Ingress Peak Rate [Gb/s]	3	5	12	25	37	50
Ingress Peak Bandwidth [Gb/s]	5	10	25	50	74	99
Egress Rate [Gb/s]	1	1	3.1	6.2	9.3	12
Egress Peak Rate [Gb/s]	1	3	6.2	12	19	25
Egress Peak Bandwidth [Gb/s]	3	5	12	25	37	50

250

Share [%]	1	2	5	10	15	20
Disk [PB]	6.5	13	33	65	97.5	130
Ingress Data [PB]	6.5	13	33	65	97.5	130
Egress Data [PB]	3.2	6.5	16	32.5	48.8	65
Ingress Rate [Gb/s]	3.1	6.2	16	30.9	46.4	61.8
Ingress Peak Rate [Gb/s]	6.2	12	31	61.8	92.8	124
Ingress Peak Bandwidth [Gb/s]	12	25	62	124	186	247
Egress Rate [Gb/s]	1.5	3.1	7.7	15.5	23.2	30.9
Egress Peak Rate [Gb/s]	3.1	6.2	16	30.9	46.4	61.8
Egress Peak Bandwidth [Gb/s]	6.2	12	31	61.8	92.8	124

	Input
Annual Disk Volume [PB]	780.00
ADP size fraction	0.30
Annual increase fraction	0.00
Year count	0.00
Number of Replicas	2.00
Data Consolidation / retry factor	1.50
Network overheads factor	1.25
Peak capacity factor	2.00
Network Bandwidth factor	2.00





SRCNet 0.1 testing

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

Sites have demonstrated 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

Test Campaigns

- Test Campaigns separated into:



Data Movement Challenges

- Similar to the (mini) WLCG Data Challenges;
- Tests of both Capability and Capacity

Data-Lifecycle and Science Delivery testing



- Ingestion of data into the Datalake (including the associated metadata)
- Use of data in simple 'analysis' style environments (e.g. Jupyter notebooks, visualisation tooling).

Test Day/date	Mon; Day 1	Tues	Weds	Thurs	Fri
Test Case ID	TC-001	TC-002	TC-003	TC-004	TC-004
Start / end time	0600UTC - 0000UTC (Tues)	0600UTC - 0000UTC (Tues)	0800UTC - 1700UTC	0000UTC - 2359UTC	0000UTC - 2359UTC
Test Case Description	Verify testing code and site readiness	Repeat of Day-1, increased rates	Sustained tests between sites, determined from TP-001-2)	Maximise the throughput within SBCNet	Maximise the throughput within SBCNet
Pre-Conditions	- Basic and critical services running - Data pre-staged at RSEs - Test Plan defined	As Day-1	Rate and matching site characteristics informed from TC-002	As before	As before
Test Steps	Transfers between all links to be used in the tests.	transfers between unique pairs of RSEs for defined rates. Focus on links that can best test each sites Capacity.	Configure full mesh to move data between sites at defined rates	Maximise throughput with strategic set of link transfers	Maximise throughput with strategic set of link transfers (Updated from TC-004)
Expected Results	Successful transfers between participating RSEs, > 70% of possible links covered	Expected volumes of data transferred within the expected period + 20% allowance	Sites can sustain transfers over 3hr periods between links.	75 TB moved in 24hours	150 TB moved in 24hours
Actual Results (summary)	SBCNet Test Campaigns - Progress Monday/April 2025: Bad links identified: ~ 75% of remaining links achieved > 85% data transfer volume in the test window. 6.9% transfer failure rate (including bad links).				
Status	Passed	Failed			
Comments/Observations (summary)	Some instability in the morning of sustained rates, improved in afternoon (cause unknown as yet). No reported issues.				
Defects / Issues raised (Jira's, etc.)	SOG-39 - Failing transfers from SWISS-OSO-T1 and CHES-VER-PROD to STFC_STORM IN PROGRESS				

Test Plan outline



Milestones so far achieved (1)

- Capability tests performed:
 - Replication from Site_A -> Site_B and Site_B->Site_C
 - (Excluding Site_A -> Site_C)
 - Emulating data flow modelling
 - 5GiB files
- Capacity testing:
 - 150TB target in 24hours
 - Target rates for each RSE link-pair
 - Simple modelling emulating movements between larger and smaller RSEs
 - 147TB transferred within the time window
- Improvements in modelling in expected share of resources and long-haul source testing
- SRCNet now visible within the Global FTS Monitoring plots



Milestones achieved (2)

- Test of ability of Rucio / FTS to handle higher submission rates, and tests of enabling services (e.g. monitoring)
- 1M x 1MiB files; target to replicate from a single source RSE to other Sites within 24hours
 - Stored as 1k x 1000 file datasets; submitted throughout the test period
 - Storage at Rutherford Appleton Lab used as source for all transfers



- Nominal submission rate (by rucio) $\sim 50k/\text{hour}$, with increase tested to $100k/\text{hour}$.
Rate comparable to e.g. CMS during the period.

Some things we have we learned

- Clear that mean bandwidths will be lower than foreseen for some time
- BUT for efficient transfers we will want high bandwidth connections to SRC Nodes – want transfers to complete quickly and reliably
- So even a “small” SRC node with maybe only a few 10s of PB of storage will want 100gb links – aim to match output rates of SDPs in South Africa & Australia
 - But these can be shared with other uses
 - Mean utilization may remain 5-10gbs for some SRCs
 - Will be significantly higher for the large SRCs
 - UK, South Africa, Australia
 - As in WLCG we will likely see growth in user/workflow driven transfers



Summary

Data Volumes and Bandwidth estimation

- SKAO is refining its roadmap for how much data will be produced and when
- In SRCNet we can translate that to local bandwidth requirements at SRCNet sites
 - Can map that to the developing requirement
 - **New estimates are much lower than previous ones**
- NRENs and SRCNet sites can negotiate if investment/development is required
 - It is the responsibility of sites to ensure they have the network provision required locally

Data movement tests

- Demonstrating ability to move both large files/datasets and large numbers of files/datasets
- Shaking out network and storage configuration issues

SRCNet Network coordination

- New SRCNet Network Strategy Group
 - To confirm global requirements fit in existing or planned capacity
 - Facilitate discussion to resolve any global or local issues are addressed
 - Meeting at existing international meetings around the world
- Also continuing to collaborate with WLCG via LHCOPN-LHCONE meetings
 - We have many shared areas of interest



Questions

