

SRC | Net

SKAO Regional Centre Network

SRCNet v0.1 Node Requirements

Document Number: SRC-0000008
Classification: UNRESTRICTED
Document type: RSP
Revision: 01
Date: 2024-03-28
Status: Draft
Authors: Garbutt, John; Walder, James; Salgado, Jesús; Joshi, Rohini

<i>Role</i>	<i>Name</i>	<i>Designation</i>	<i>Affiliation</i>	<i>Signature</i>	<i>Date</i>
Author	J. Garbutt	UK SRC Architecture Group	UK SRC	_____	
Owner	J. Salgado	SRCNet Architect	SKAO	_____	
Approver	R. Bolton	SRCNet project Lead	SKAO	_____	
Released by	L. Ball	SKAO Director of Operations	SKAO	_____	

© Copyright 2024 SKA Observatory.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/)

TABLE OF CONTENTS

Status of the Document	2
1. Overview and Scope	3
2. Requirements Summary	4
3. SRCNet v0.1 Software Description	9
4. People and Operations Requirements	11
5. Network Requirements	12
5.1 External Connectivity Requirements	13
5.2 Internal Connectivity Requirements	13
6. Storage Requirements	14
6.1 Tier 2 - Nearline Storage	16
6.2 Tier 1 - Online Storage	16
6.3 Tier 0 - Cache Storage	17
7. Compute Requirements	17
8. SRCNet Node Validation Tests	19
8.1 Access Control Test	19
8.2 Network and Data Transfer Tests	20
8.3 Storage Performance Tests	20
8.4 Compute Tests	20
9. Roadmap to produce a technical plan for the SRCNet v0.1	21
Appendix A: Collaboration from Non-Production Nodes	23



Status of the Document

This document has been produced by the SRNet Architecture Forum. It has been reviewed by SRCNet Members and the SRC Steering Committee and has been endorsed by the SKAO as an SRCNet formal document. It is a stable document and may be used as reference material or cited as a normative reference from another document.

1. Overview and Scope

This document aims to define what is required from an infrastructure provider to support SRCNet v0.1 [AD2]. This is expected to be an iterative process, where we hope to have more clarity when we come to SRCNet v0.2. Where possible, we are informing our initial choices based on work done.

The basic use cases to be covered for this version are:

- Ingestion of relevant data products ingested into the SRCNet nodes, which would include a significant amount of data from precursors, synthetic data from SKA, etc.
- Data available for discovery and access from all the nodes
- Data and services discovery from a simple user interface
- Client command line libraries to allow data access, data discovery and data analysis, e.g. through notebook environments and support for a limited set of science workflows.
- Visualisation of certain formats of imaging data from remote locations
- Common Authentication/Authorisation system

All these tests will allow the SRCNet members to get experience in the operational management of an SRCNet with a relevant size, adding basic science use cases on non-SKA and synthetic SKA data, allowing performance measurements on data latency, network latency and, in general, to prepare the commissioning phase for v0.2 where real science use cases and science verification use cases will appear.

The following document will be used to define in detail the expected requirements per SRCNet node for v0.1 (sections 2 to 7), identify the tests required for the validation of resources, (section 8) and present a timeline with the steps required to produce a technical plan for the SRCNet v0.1 (section 9). Notice that although the SRCNet services could be deployed into a local site as part of a shared architecture, the resources described here are available for the SRCNet and the node's total.



The SRCNet v0.1 will be formed by a limited number of nodes to facilitate the operations and testing in this initial phase. For this reason, the requirements for an SRCNet v0.1 node have been calculated considering that 4 nodes (at a minimum) will be integrated with the network.

Once the requirements per SRCNet v0.1 node, SRCs will be invited to provide expressions of interest, detailing the hardware and human resources that can contribute. Then a deployment plan to build the SRCNet v0.1 will be developed, according to the available resources indicated in the expressions of interest. Certain SRCNet nodes will be designated production nodes for the deployment, while other SRCs will have the opportunity to support the SRCNet v0.1 network in various capacities. Additionally, they may participate in the development network, which operates concurrently with the production network, facilitating testing of changes, new modules, etc. Please refer to the [appendix](#) for a list of contributed tasks.

The number of deployment nodes for future versions will be increased in line with the SRCNet top-level roadmap [AD2], although some possible deviations are expected depending on the SRCs' status.

The storage requirements, compute requirements and expected timelines being used are taken from the SRCNet top-level roadmap [AD2]. As such, the expectation is that v0.1 will be rolled out before the end of 2024, and will be iteratively refined towards v0.2 planned in early 2026. These timelines assume that AA0.5 will be in late 2024, AA1 will be in late 2025, AA2 will be in late 2026 and AA* in early 2028, with science verification and commissioning of production data movement starting in mid-2026. There is an expectation that this will soon need updating in the light of [SKAO-TEL-0002299](#).

2. Requirements Summary

Requirement	Summary	Compliance Level	Section
People and Operational Requirements	Equivalent of 1 FTE to support SRCNet v0.1, including contributions towards building a long-term shared operational team across SRCNet.	Must	People and Operational Requirements



	Maximum 3 working day turnaround on responding to production issues affecting SRCNet	Must	
Network Requirements	Internet connectivity to send and receive data sets between all SRCNet v0.1 sites, including to both South Africa and Australia SKAO sites	Must	Network Requirements
	1Gbps bandwidth from SRCNet v0.1 node and both SKAO sites, South Africa and Australia.	Should	
	Minimum 10Gbps internet connection, likely shared with other users	Should	
	1Gbps bandwidth from SRCNet node and the rest of SRCNet v0.1 nodes	Should	
	IPv6 capable internet connectivity	Should	
Storage Requirements	<p>Online Storage: Minimum of 5PB of online bulk storage for SRCNet use (i.e. comparable with SATA Hard Disks or superior) available by Jan 2025</p> <p>This requirement is defined for an SRCNet deployment composed of only 4 nodes. SRCs should respond to this request by providing information on possible oversubscriptions to compensate other non-available</p>	Should	Storage Requirements



	<p>SRCs in line with the request, subscriptions in line with [AD2] or realistic figures of the resources available for Jan 2025.</p> <p>Plans from different SRCs to join efforts to define a common SRCNet deployment node could also be proposed in the expression of interest.</p>		
	<p>Commitment towards providing 20PB for v0.2 in Jan 2026</p> <p>This requirement is defined for an SRCNet deployment composed of only 4 nodes. SRCs should respond to this request by providing information on possible oversubscriptions to compensate other non-available SRCs in line with the request, subscription in line with [AD2] or realistic figures of the resources available for Jan 2026.</p> <p>Plans from different SRCs to join efforts to define a common SRCNet deployment node could also be proposed in the expression of interest.</p>	Should	
	<p>Scratch/Cache/Computing connected Storage: Minimum of 500TB of high-speed scratch storage accessible from all the Compute resources</p>	Should	
	<p>Assurances around preservation of SRCNet dictated access control policies to all the</p>	Must	



	SRCNet ODPs and ADPs stored at the SRCNet node, following procedures agreed by the Operations Team ¹		
Compute Requirements	<p>0.175 PFLOPS per SRCNet v0.1 node (Rpeak; theoretical peak performance) available by Jan 2025²</p> <p>This requirement is defined for an SRCNet deployment composed of only 4 nodes. SRCs should respond to this request by providing information on possible oversubscriptions to compensate other non-available SRCs in line with the request, subscriptions in line with [AD2] or realistic figures of the resources available for Jan 2025.</p> <p>Plans from different SRCs to join efforts to define a common SRCNet deployment node could also be proposed in the expression of interest.</p>	Should	Compute Requirements
	<p>Commitment of scaling towards providing 0.7PF for v0.2 in Jan 2026</p> <p>This requirement is defined for an SRCNet deployment</p>	Should	

¹ SRCNet v0.1 is not expecting to have proprietary data

² Performance of a general platform could be computed as follows:

Performance (PFLOPS)=Number of Nodes×Cores per Node×Operations per Cycle per Core×Clock Speed (in Hz)/10¹⁵

As the requested values are theoretical, the manufacturer's specifications could be used. A more precise analysis will be obtained later with benchmarking tests



	<p>composed of only 4 nodes. SRCs should respond to this request by providing information on possible oversubscriptions to compensate other non-available SRCs in line with the request, subscriptions in line with [AD2] or realistic figures of the resources available for Jan 2026.</p> <p>Plans from different SRCs to join efforts to define a common SRCNet deployment node could also be proposed in the expression of interest.</p>		
	Performant access to both Online Bulk Storage and appropriate POSIX-like scratch storage ³	Must	
	Minimum of 384GB of RAM per node, ideally at least 8GB per core.	Should	
	Minimum of 25Gb ethernet, ideally capable of RDMA ³	Should	
SRCNet Node Validation Tests	Ability to run all the specified Validation tests, and report the results of those tests.	Must	Validation Tests
SRCNet Node Software	Local deployment software required for an SRCNet node v0.1	Must	3. SRCNet v0.1 Software Description

³ Exact details to be further specified within v0.1 implementation plan is defined



3. SRCNet v0.1 Software Description

SRCNet v0.1 will have a limited distribution for use and testing, only available to SRC ART members and selected members of the SKA Commissioning team. It will be used to deploy a stable version of some of the prototypes already developed by the SRCNet development team, ensuring connectivity between nodes and a higher level of operational maintenance and availability than previous prototypes.

There is a set of SRCNet API services that will serve as a collectively developed and owned translation layer that allows for the interfacing between the different components. Some of them are stateless, can be run as a local service and the responsibility for the same can be shared across 0.1 Sites. Others require to be run centrally.

From the software point of view, the components involved will be, at least:

- A Science Platform Presentation Layer (webserver)
 - Most possible candidate, ESAP
- Common data-related services (discover data, access data, etc)
 - Possible candidates, Rucio prototype or CADC Storage Inventory
 - Data Management API service
- Services metadata-related services (data registration, discovery and access of services of the SRCNet)
 - Most possible candidate, the IVOA TAP server with a dedicated metadata database (PostgreSQL) for discovery
 - If more than one metadata global services available, mirroring services, like CADC CAOM system or open source DB replication software like pgpool-II
 - At least one data registration (ingestion) prototype node, using ingestion prototype server, what includes data import into the SRCNet data lake and a metadata ingestion service
- Interactive analysis interface (notebook interface)
 - Most possible candidate, JupyterHub
- Possibly, a science platform
 - Most possible candidate, CANFAR
- Orchestrator to deploy, manage and scale services
 - Most possible candidate, Kubernetes
- Other local data parsing and visualisation services (IVOA SODA, visualisation services, etc)
- Containerised Visualisation tools running locally
 - CARTA, Aladin, VisIVO,... etc.
- Permissions API service
- Monitoring services
 - Most possible candidates, Perfsonar, Prometheus, and others



- Possible additional services, like Slurm clusters, discoverable through the services discovery service but not fully integrated into the SRCNet
- Where applicable these services will need to interface with or integrate with the global services defined below (IAM, Site Capabilities, etc)

These are examples of services locally running at an SRC site and thus there will be several instances of these services running at any point in time.

In addition to this, there will be certain global services running centrally. Due to the stateful nature of these services, there will likely be only one instance of them running at a time (additional instances for failover/load balancing might be necessary but outside the scope of this document).

These global services include

- The Identity and Access Management service, most possibly candidate Indigo IAM, Keycloak (a solution that supports the OpenID Connect, OIIC standards)
- Site capabilities API - this allows for discoverability of information regarding a site, including but not limited to compute, storage and services available.

This software stack is still under discussion and will be defined in a dedicated implementation plan, although it is expected to be very close to the previously described

From the architectural perspective [AD1], the modules to be deployed and maintained will be:

Version	SRCNet v0.1
Expected release	Nov 2024-Jan 2026 (estimated v0.2)
Software Thread Versions	SRCNet Presentation Tier v0.1 Data Management v0.1 Metadata Management v0.1 AAI preliminary capabilities
Description	Preliminary version of the portal, enabling discovery queries on the data. That would also contain a basic local interactive analysis able to discover and access data from the SRCNet data lake, invoke local replicas and analyse data locally



	<p>A preliminary version of the data management system, to allow the creation of digital object replicas on the SRCNet node selected by a user.</p> <p>A preliminary version of the metadata system to discover data in the data lake</p> <p>Minimum viable product of the authentication/authorisation system, in connection with the data management system to control access policies</p>
--	--

4. People and Operations Requirements

All prototype SRCNet sites will need to have people available to participate in the shared operational needs of SRCNet. There is an expectation that it will need at least 1 FTE locally at each site to help support the operation of SRCNet v0.1 (as specified in the [AD2]).

A non-exhaustive list of tasks to be coordinated by the members of the SRCNet operations team (some already fully applicable for SRCNet v0.1 or to be prepared using this deployment) could be [AD2]:

- Data Distribution: The operations team should be responsible for managing the distribution of the data produced by the SKA telescopes, including storing and, archiving into the SRCNet.
- Data Access: The operations team should be responsible for maintaining the access of the SKA data into the SRCNet following agreed policies.
- Data Processing: The operations team should be responsible for maintaining the data processing modules.
- The operations team should be responsible for the monitoring of the SRCNet and the execution of health and safety procedures of SRCNet components.
- The operations team should be responsible for the coordination of maintenance activities of the SRCNet nodes.
- The operations team should be responsible for the handling of routine operations like the maintenance of users/groups quotas, creation of dashboards and stats, coordination with IT teams on upgrades and security patches, maintenance of software repositories, etc
- Adaptation and execution of SRCNet compatible science workflows, including generation of benchmark metrics



In particular, all sites must be actively contribute to the realization of the SRCNet vision:

- Maximum 3 working day turnaround on responding to production issues affecting SRCNet, and happy working towards globally consistent measures of service availability.
- Sites agreed to use a single harmonised set of policies (e.g. based on <https://wise-community.org/policy-development-kit/>) such that all user access follows the AARC blueprint (AARC Consortium Partners;AppInt members;Nicolas Liampotis, 2019), and flows from a single SRCNet AA service.
- Clear communication of local at-risk windows with all other SRCNet sites, where possible coordinating to ensure minimal disruption to SRCNet as a whole
- Join the SRCNet Operations team responsible for running the SRCNet services across all infrastructure provider sites, centrally reviewing all changes to the systems via GitOps-style processes, and coordinating the upgrade of all sites.
- Ideally, the SRCNet Operations team should have access to SRCNet services' logging and metrics across all the sites, including operational alerts being exposed centrally to the SRCNet Operations team (e.g. via a set of shared Slack channels).
- SRCNet-wide support tickets, triaged centrally between general user support and escalation to local SRCNet infrastructure teams for further information.
- SRCNet resource allocations will be managed globally, with the required usage information collected from each site (likely coming from the centrally operated SRCNet services)
- The SRCNet operations team is responsible for monitoring the operational condition and performance of various services using tools like Perfsonar, Prometheus, and others.

For SRCNet v0.1, some tradeoffs may be required, such as additional groups to manage the acceptance of specific site local policies, and all upgrades being performed by local site staff. But all sites should be helping SRCNet work towards the medium-term vision of a joined-up federation.

5. Network Requirements

In addition to the above acceptance tests, we discuss what hardware is expected to help pass the acceptance tests and any other non-functional requirements that need to be discussed.



5.1 External Connectivity Requirements

Once a final set of sites is chosen, there must be a minimum of 1 Gbps upload and download speed between all chosen SRCNet nodes. All systems must include the ability to expose SRCNet services on the internet and be able to push and pull data to and from all other SRCNet0.1 sites. The site must be equipped with a dedicated minimum 10 Gbps internet connection to ensure optimal performance, which may be shared with other users.

The site must agree to expose detailed metrics around the external connectivity with the central SRCNet operations team. The compilation of data center metrics regarding connectivity and SRCNet services usage for monitoring purposes will not contain any information that may degrade local security measures, and will be aligned with the regulations outlined in GDPR.

The site must demonstrate a roadmap towards a dedicated 100 Gbps link according to the roadmap, with agreements around the required global links with local NRENs, in collaboration with the appropriate SRCNet community of practice, likely including dedicated VRFs by 2030.

In additions all SRCNet sites should:

- Maintain 1 Gbps download bandwidth from both South Africa and Australia SRCNet sites. Ideally, this involves ensuring support for jumbo frames for the full network path.
- Have IPv6-capable internet connectivity
- Have the ability to expose SRCNet services on the internet
- Be able to pull and push data to and from all other SRCNet sites
- Provide a minimum 10 Gbps internet connection, likely shared with other users, with Jumbo frames and IPv6.

5.2 Internal Connectivity Requirements

Currently, the main internal connectivity requirements are around what is required to support good access to the storage. These details are better captured in the storage requirements.

To enable the prototyping of multi-node workloads⁴ using UCX, MPI and similar, ideally, RDMA should be possible between the compute nodes, using a minimum of 25Gb ethernet.

⁴ MPI over 12 nodes example workflow in its current cut down form:
<https://gitlab.com/ska-telescope/src/src-workloads/-/tree/master/workflows/surveys/askap/emucat>



6. Storage Requirements

In addition to the acceptance tests described below, we discuss what hardware is expected to help pass the acceptance tests, and any other non-functional requirements that need to be discussed.

As discussed in [AD2], storage will comprise three tiers with different technologies. On that document, the expectations of resources offered by the different SRCs are:



Document Number SRC-00000
Revision 01
Date 2024-03-28

UNRESTRICTED

SRCNet
Author: John Garbutt
Page 14 of 26

SRCNet v0.1		
Country	Share (%)	Storage (PB)
UK	19	4.03
South Africa	18	3.82
Australia	18	3.82
China	10	2.12
Canada	7	1.48
Italy	6	1.27
India	5	1.06
France	3	0.64
Netherlands	2	0.42
Japan	2	0.42
Spain	2	0.42
Portugal	2	0.42
Switzerland	2	0.42
Sweden	2	0.42
South Korea	1	0.21
Germany	1	0.21
Total	100	21.20

Figure 1: Storage resources of the SRCNet requested for SRCNet v0.1 as per [AD2]

SRCs could respond to this requirements request by oversubscribing to compensate other non-participating SRCs contributions being a main node, subscribe with the proposed resources in the roadmap or provide a figure of the resources available at the v0.1 deployment date so a realistic implementation plan could be prepared.

The storage available for v0.1 will be used to test scaling of scientific use cases using precursor data distributed into the SRCNet and by test data, as close as possible to the SKA future data, for engineering tests on data discovery, access and analysis.



6.1 Tier 2 - Nearline Storage

Nearline (or cold storage) is typically implemented using tape, for efficient long-term storage. Nearline storage is out of scope for v0.1.

Sites must be preparing for around 2028 when it is assumed the global amount of archived data will grow at the rate of around 1000 PB usable per year. There will need to be an effort to define what interfaces are needed to support moving data between Nearline Storage and Online Storage (and Cache Storage).

6.2 Tier 1 - Online Storage

Bulk online storage, today typically implemented using spinning disks, is used for data to be actively analysed.

SRCNet0.1 requires 20 PB across the participating sites. All sites will not be ready for 0.1 so some sites will hopefully provide more resources early on to help bridge gaps. Thus sites provide a minimum amount of storage corresponding to their share of 20PB up to 5PB, if we expect 4 nodes. Site selection based on this requirement will depend on the number of sites and their feedback following this document. Commitment towards providing their corresponding fraction of storage resources for v0.2 by Jan 2026.

Sites must be prepared for around 2028 when we are assuming the sum of all global online bulk storage is expected to be around 1060PB usable (including two copies of the data at the SRCNet for high availability). Data older than one year would be only available from Tier 0 - Nearline Storage [AD2], so Tier 0 will only be actively used after one year of data production (at the time of writing this document, in 2029). See the discussion on possible changes, depending on budget and technology, on the roadmap document.

There is expected to be a single data SRCNet-wide data transfer system to move data between SRCNet sites, and local sites must provide storage that interfaces with that system. There needs to be appropriate resources to support.

For SRCNet v0.1, for expediency, we assume Online Storage will have POSIX-like access.

Before SRCNet v0.2, there needs to be more work on understanding what storage interfaces should be supported to help with scaling the online storage to 200 PB usable capacity and beyond, alongside providing adequate bandwidth for ingesting data from SRCNet, moving data into Nearline storage, moving data into Cache and persisting advanced data products. For example, it is possible standardising around an S3-like object storage interface would be more scalable and improve data movement, if POSIX compatibility layers can be shown to be suitably performant, for pipelines expecting a filesystem interface.



6.3 Tier 0 - Cache Storage

Fast storage attached to computing nodes to be used during data analysis. Also called Scratch storage. For SRCNet v0.1, the expectation is that this storage will be exposed to compute workflows using a POSIX-like file system interface. Some sites may provide an online storage tier that is both performant enough and close enough to all the compute resources, such that there are no workflows that would benefit from staging data into a different cache tier.

Given simulated datasets for SKA might be around 200TB, we need each site to support at least 500TB of high-speed scratch storage that is accessible from the Compute resources.

For more io-intensive and/or compute-intensive pipelines, we expect that data will need to be first staged into faster storage. This is under the assumption that we want a smaller amount of fast storage close to each compute cluster, to improve the efficiency of the pipelines vs stream data from Long-term storage. We should expect the pipelines that need these data pre-staging to change over time.

It should support a minimum of 1GB/s Read and 1GB/s Write when tested with storage benchmark test tools defined in this document.

7. Compute Requirements

The modelling for compute resources needed for analysis of the data to be executed by the scientific community for v0.2 is still under definition.

According to the roadmap document, a gradual deployment of the computing resources will be done and, for v0.1 and considering only a limited number of nodes (4) and only including the major contributors, a reasonable amount of computing would be a combined total of 0.7 PFLOPS across all SRCNet v0.1 sites. Having only 4 SRCNet nodes, the recommendation would be to have around 0.175 PFLOPS per node. These figures would be a clear oversubscription for several countries.

On document [AD2], the expectations of resources offered by the different SRCs are:



SRCNet v0.1		
Country	Share (%)	Computing (PFLOPS)
UK	19	0.13
South Africa	18	0.13
Australia	18	0.13
China	10	0.07
Canada	7	0.05
Italy	6	0.04
India	5	0.04
France	3	0.02
Netherlands	2	0.01
Japan	2	0.01
Spain	2	0.01
Portugal	2	0.01
Switzerland	2	0.01
Sweden	2	0.01
South Korea	1	0.01
Germany	1	0.01
Total	100	0.70

Figure 2: Computing resources of the SRCNet requested for SRCNet v0.1 as per [AD2]

SRCs could respond to this requirements request by oversubscribing to compensate other non-participating SRCs contributions being a main node, subscribe with the proposed resources in the roadmap or provide a figure of the resources available at the v0.1 deployment date so a realistic implementation plan could be prepared.

As stated in [AD2], SRCNet v0.1 will not offer open scientific access (although a minimal number of beta-testers could be appointed). This version will be utilized for testing purposes (such as data logistics tests, data analysis compute tests, stress tests, etc.) and to prepare for future scaled advanced versions (starting from v0.2 and beyond), where external users are expected to run high-load



workflows. In order to achieve this goal, we need to consider the expected workload required for the following scientific use cases:

- Sky Image ODPs
 - Create “flat images” for visualisation
 - Support visualization of ADPs and ODPs, and interactive analysis
- Other ODPs
 - Little data is available at this stage

In terms of a practical baseline each compute node should have:

- Minimum of 384GB of RAM, ideally at least 8GB per core. Examples include:
 - 2 x Intel 6142 with 384GB RAM
 - 2 x AMD 7763 with 1TB RAM
 - 2 x AMD 9654 with 1.5TB RAM
- Minimum of 25Gb ethernet, ideally RDMA capable and ovs offload capable
- Ideally low latency (MPI supported) interconnect, such as the RDMA capable ethernet above
- Access to both Online Bulk Storage and appropriate POSIX-like scratch storage

8. SRCNet Node Validation Tests

There needs to be a suite of SRCNet site validation tests, to confirm a base level of performance and functionality. These tests should validate a site meeting the minimum requirements for joining SRCNet. These tests are expected to evolve as the hardware and software requirements evolve.

While sites need to pass these tests to be accepted, there should also be some ongoing testing to spot any critical performance regressions. For SRCNet v0.1, we have not defined a minimally acceptable level of performance, rather we want to start to collect some baseline results across potential sites, with the intent of defining a baseline level of performance for SRCNet v0.2.

These tests will be containerised such that they can be run using a variety of existing infrastructure interfaces such as a job within an existing HPC cluster with Apptainer/Singularity/Podman, a Kubernetes job, in an OpenStack VM, etc.

8.1 Access Control Test

Access to the infrastructure is expected to be managed using a single central AAI system, as agreed by the global SRCNet operations team.

For SRCNet v0.1, there may be additional local user policies that need to be accepted above and beyond SRCNet centrally managed acceptable user policies,



but the site should be engaged in a process to allow access solely based on acceptance of a global harmonised SRCNet set of policies.

All tests should be simply runnable on a candidate SRCNet site without violating any of the local (or global) policies that must be accepted.

8.2 Network and Data Transfer Tests

First, there will be tests to confirm the wide area network performance to both South Africa and Australia, ideally confirming that both IPv6 and jumbo frames are available.

Secondly, there are tests to ensure data can be moved (“disk to disk”) from both South Africa and Australia into the SRCNet node’s local online bulk storage (i.e. data lake). This test is an end-to-end test of getting data read at the source, sent across the network, and written to the destination, ensuring the data transfer nodes (DTNs) can sustain throughput to both the local storage systems and the wide area network.

Thirdly, there are tests to test how long it takes to move data from the online bulk storage into any faster storage tiers that may be available to speed up compute pipelines.

Not all of these tests have yet currently been containerised and fully automated, so these tests may be delivered as documentation (TBC).

8.3 Storage Performance Tests

There are plans for CASA measurement set tests, on a single node, for serial read, serial write, parallel read and parallel write operations. These CASA tests will make use of a representative set of visibilities and image cubes, likely using a combination of publicly available LOFAR data and simulated SKA data sets.

All compatible storage tiers should be tested, both the online bulk storage (i.e. data lake) and any available faster storage tiers.

The initial version of those tests can be found here:

<https://confluence.skatelescope.org/display/SRCSC/Benchmarking>

8.4 Compute Tests

For SRCNet v0.1 there will be an initial set of workflows that can execute at all candidate SRCNet nodes and can be used to objectively compare the performance between proposed SRCNet nodes⁵. Over time, we can work to

⁵ Previous work in this area:

https://gitlab.com/ska-telescope/src/src-workloads/-/tree/master/bench?ref_type=heads



establish what an acceptable baseline performance looks like and expand the set of representative workloads.

For v0.1, test workflows will be basic science use cases as, e.g.:

- Imaging pipeline, processing raw (public) visibilities into an image cube, (using LOFAR imaging pipelines, WSClean⁶,⁷,..., etc.)
- Visualisation tool like CARTA to visualise a range of image cube sizes, with human verification that the performance is acceptable

Note: There is a clear miss around further pipelines that manipulate images (e.g. source finding) and pipelines that require GPUs. We expect to add these for version 0.2.

These workflows, and others to be compiled, will be close to real science use cases to be executed by users to benchmark the nodes with realistic scenarios. These workflows will be containerised (for portability) and properly documented⁸, so executing them at the SRCs would be as more automatic as possible. Documentation will include the execution interface to run them using public data on any system with a prebuilt OCI-compatible container image, e.g. Apptainer, Podman or similar.

The expectation is the chosen benchmark workflows will work on public data and be publicly available within:

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

9. Roadmap to produce a technical plan for the SRCNet v0.1

After collecting the feedback from the Architecture forum on the SRCNet 0.1 node requirements document, it will be released. Then the SRCs will provide expressions of interest to participate in the deployment of the SRCNet v0.1, detailing the hardware and human resources that can contribute. After compiling these expressions of interest, a process will be initiated to identify:

see also presentation:

<https://docs.google.com/presentation/d/1huLzdZ-OCIJTKMxRtH9ykAAPPdISFmxmcaPnSm19Q5U/>

⁶ <https://wsclean.readthedocs.io/en/latest/>

⁷ Profiling metrics with WSClean and SOFIA previous work at

<https://docs.google.com/presentation/d/1Mk-jqDFop38sK6HDnc48pkFAT32j4OoiVTR1tjjMA7E/>

⁸ Currently, UKSRC is already doing this exercise for their local SRC in a way that could be reused by the SRCNet in general. Other SRCs interested in collaboration in this area are also welcome. Similar activities are also ongoing within Coral and Tangerine Agile teams



- A. SRCs with the capability to provide the resources required to be one of the limited SRCNet v0.1 deployment nodes.
- B. SRCs that while can not provide the high amount of resources required, can participate in the development of the new SRCNet versions, in the support of the deployed services and software as well as in other activities detailed in the [appendix](#) for a list of contributed tasks.

All contributions, whether as a deployment node, as computing resources for testing, or human resources for development, should be acknowledged and rewarded.

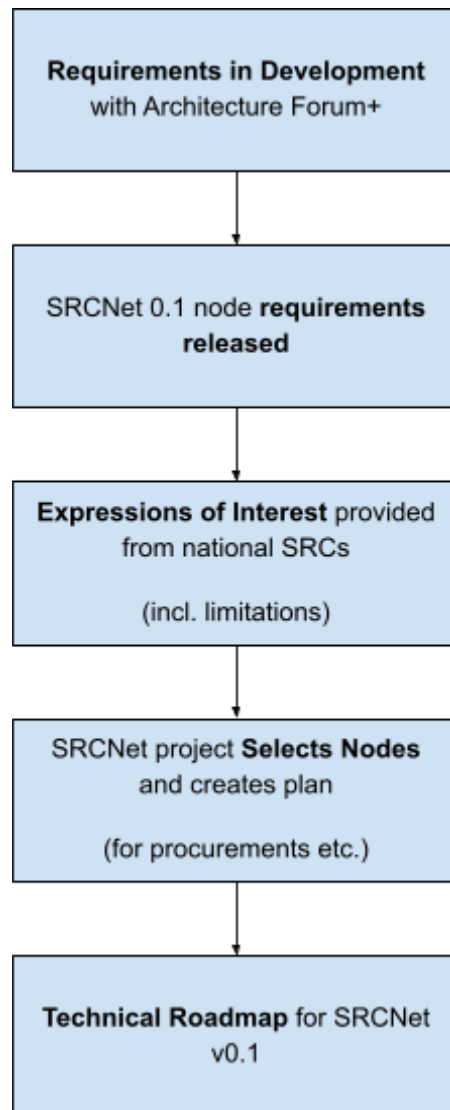


Figure 3: Expected steps for the design of the SRCNet v0.1 deployment plan and roadmap.



Appendix A: Collaboration from Non-Production Nodes

As the first deployment (v0.1) of the SKA Regional Centre Network (SRCNet) approaches, it's essential to ensure that all participating nodes, including those not directly involved in the initial deployment due to resource limitations, can still actively contribute to the network's success. These non-production SRCNet v0.1 nodes possess valuable expertise and capabilities that can significantly enhance the operational activities and overall effectiveness of the network. Here are several ways in which these nodes can collaborate and participate in supporting the SRCNet's operational activities, software development, maintenance, monitoring, and community engagement efforts:

- 1. Operational Procedures and Support:** Non-production SRCNet v0.1 nodes can actively contribute to the development of operational procedures for the SRCNet. This involves documenting best practices, protocols, and guidelines for various operational activities such as data management, services health and safety procedures, and quality control measures.
- 2. Validation and Benchmarking:** Collaborating SRCs can play a crucial role in the validation and benchmarking of nodes within the network. This entails testing the performance and reliability of hardware and software components against established testbeds and benchmarks. Through testing and analysis, non-production SRCNet v0.1 nodes can provide valuable feedback to enhance the overall functionality and efficiency of the SRCNet infrastructure.
- 3. Software Development and Maintenance:** Non-production SRCNet v0.1 nodes can contribute to the development and maintenance of the software stack deployed within the SRCNet. This involves writing code, debugging, optimising performance, and implementing new features or functionalities.
- 4. Monitoring and Support:** Collaborating nodes can assist in monitoring the performance and health of the SRCNet infrastructure. This includes deploying monitoring tools, analysing system metrics, and identifying potential issues or bottlenecks.
- 5. Documentation and Knowledge Sharing:** Non-production SRCNet v0.1 nodes can contribute to the creation of comprehensive documentation covering various aspects of the SRCNet, including system architecture, configuration details, troubleshooting guides, and best practices.



Through active participation in these collaborative efforts, non-deployed SRCNet v0.1 nodes can make meaningful contributions to the SRCNet initiative, ensuring its success and sustainability in advancing radio astronomy research and data management capabilities.

A References

A.1 Applicable documents

The following documents apply to the extent stated herein. In the event of a conflict between the contents of the applicable documents and this document, **the applicable documents** shall take precedence.

[AD1] Salgado, J.; Wicenec, A.; Goliath, S.; Joshi, R.; Swinbank, J.; Bolton, R.; Webster, B.; Oonk, J.; Grainge, K.; Sánchez, S.; Parra, M.; Dack, T.; Hardcastle, M.; Barbosa, D.; Llopis, P.; Fabbro, S.; Beswick, R.; Villote, J.P.; Breen, S.; Yates, J.; Grange, Y.; Gaudet, S.; An, T.; Possenti, A.; Darriba, L.; Holanda, V.; Mendoza, M.; Galluzzi, V.; Svedberg, T.; Lee-Waddell, K.; Vitlacil, D.; Pandey, V.N.; Akahori, T.; Chisholm, L.; Horton, M.; Watson, R.; (2023)

SRC-0000001 SKA Regional Centre Network (SRCNet) Software Architecture

<https://confluence.skatelescope.org/download/attachments/74711238/SRCNet%20Software%20Architecture.pdf?version=1&modificationDate=1693571258908&api=v2>

[AD2] Salgado, J.; Bolton, R.; Swinbank, J.; Joshi, R.; Sánchez, S.; Villote, J.P.; Gaudet, S.; Yates, J.; Barbosa, D.; Taffoni, G.; Frank, B.; van Haarlem, M.; Breen S; Conway, J.; Akahori, T.; Yates, J.; Tolley, E.E.; Wadadekar, Y.; Lee-Waddell, K; de Boer, J. (2023)

SRC-0000002 SRCNet Top-Level Roadmap

<https://confluence.skatelescope.org/download/attachments/74711238/SRCNet%20Top%20Level%20Roadmap.pdf?version=2&modificationDate=1693571286081&api=v2>

A.2 Reference documents

The following documents are referenced in this document. In the event of conflict between the contents of the referenced documents and this document, **this document** shall take precedence.

AARC Consortium Partners; AppInt members; Nicolas Liampotis. (2019, 4 30). *Deliverable*

DJRA1.4: Evolution of the AARC Blueprint Architecture.



LIST OF ABBREVIATIONS

AA	Array Assembly
AARC	Authentication and Authorisation for Research and Collaboration
ADP	Advanced Data Products
API	Application Programming Interface
CASA	Common Astronomy Software Applications
IAM	Identity Access Management
IPv6	Internet Protocol version 6
LOFAR	LOW Frequency ARray
MPI	Message Passing Interface
NREN	National research and education network
NVMe	Nonvolatile memory express
ODP	Observatory Data Products
POSIX	Portable Operating System Interface
RDMA	Remote Direct Memory Access
SATA	Serial AT Attachment
SKA	Square Kilometre Array
SKAO	SKA Observatory
SODA	IVOA Server-side Operations for Data Access
SSD	Solid-state drive
SRC	SKA Regional Centre
SRCNet	SKA Regional Centres Network
UCX	Unified Communication X
VRF	Virtual routing and forwarding



POLICY APPROVAL AND OWNERSHIP DETAILS

POLICY APPROVAL AND OWNERSHIP DETAILS	
Impact Assesment:	Yes/No
Date of Review:	YYYY-MM-DD
Document Location:	[Path link]
Compliance Checks:	Yes/No

DOCUMENT HISTORY

Revision	Date Of Issue	Comments
01.1	2024-01-12	First draft release for internal review
01.2	2024-02-01	Document distributed to SRC SC for comments
01	2024-02-15	Document frozen and distributed for resources gathering

