China's Participation in SRCNet V0.1 Intentions and Planning A strategic Overview

Shaoguang Guo

on behalf of cnSRC

cnSRC contact: shaoska@shao.ac.cn

China SRC prototype: status (2024)

- HW
  - Compute:
    - 35 Intel x86 CPU nodes -> 2240 cores, with 128 cores/node
    - 12 ARM CPU nodes -> 1152 cores, with 96 cores/node
    - 4 GPU nodes -> 16 Nvidia V100, 8 A100
  - Distribute Storage : ~ **9PB**
  - Inter-connection: 100-200Gb/s connection
  - Memory: maximum **4TB/node**, max 36GB/core
  - Tran-continental internet:
    - **10Gbps** international link;
    - 100 Gbps to be considered by 2026;
    - 200Mbps with other SRC nodes
- SW
  - SLURM/MODULE/Container/Instance
- Science
  - MWA/ASKAP/MeerKAT

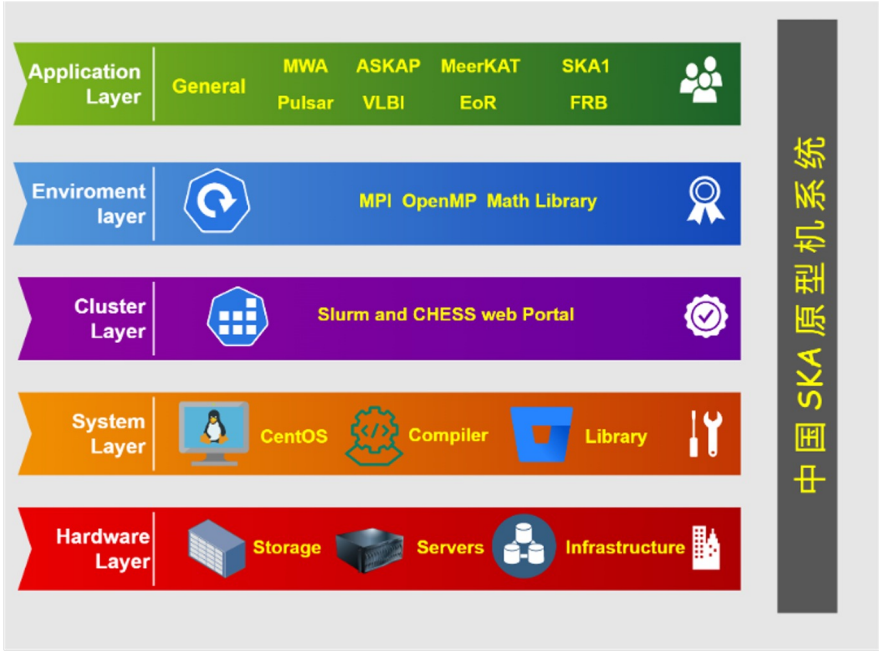

Our Preparation Storage/Computer Resources

cnSRC commit support to 1PB by 2024 (2PB by 2025)/0.175PFLOPS

		SRCNet v0.1	SRCNet v0.2	SRCNet v0.3	SRCNet v1.0b	SRCNet v1.0
		Jan 2025	January 2026	Sep 2026	Nov 2027	Jun 2028
Deployment (%)		2.00	10.00	15.00	50.00	100.00
Country	Share (%)	Computing (PFLOPS)	Computing (PFLOPS)	Computing (PFLOPS)	Computing (PFLOPS)	Computing (PFLOPS)
UK	19	0.13	0.67	1.00	3.33	6.65
South Africa	18	0.13	0.63	0.95	3.15	6.30
Australia	18	0.13	0.63	0.95	3.15	6.30
China	10	0.07	0.35	0.53	1.75	3.50
Canada	7	0.05	0.25	0.37	1.23	2.45
Italy	6	0.04	0.21	0.32	1.05	2.10
India	5	0.04	0.18	0.26	0.88	1.75
France	3	0.02	0.11	0.16	0.53	1.05
Netherlands	2	0.01	0.07	0.11	0.35	0.70
Japan	2	0.01	0.07	0.11	0.35	0.70
Spain	2	0.01	0.07	0.11	0.35	0.70
Portugal	2	0.01	0.07	0.11	0.35	0.70
Switzerland	2	0.01	0.07	0.11	0.35	0.70
Sweden	2	0.01	0.07	0.11	0.35	0.70
South Korea	1	0.01	0.04	0.05	0.18	0.35
Germany	1	0.01	0.04	0.05	0.18	0.35
Total	100	0.70	3.50	5.25	17.50	35.00

Figure 7: SRCNet computing resources. Deployment of the SRCNet computing resources in line

		SRCNet v0.1	SRCNet v0.2	SRCNet v0.3	SRCNet v1.0b	SRCNet v1.0
		Jan 2025	January 2026	Sep 2026	Nov 2027	Jun 2028
Deployment (%)		2.00	10.00	15.00	50.00	100.00
Country	Share (%)	Storage (PB)	Storage (PB)	Storage (PB)	Storage (PB)	Storage (PB)
UK	19	4.03	20.14	30.21	100.70	201.40
South Africa	18	3.82	19.08	28.62	95.40	190.80
Australia	18	3.82	19.08	28.62	95.40	190.80
China	10	2.12	10.60	15.90	53.00	106.00
Canada	7	1.48	7.42	11.13	37.10	74.20
Italy	6	1.27	6.36	9.54	31.80	63.60
India	5	1.06	5.30	7.95	26.50	53.00
France	3	0.64	3.18	4.77	15.90	31.80
Netherlands	2	0.42	2.12	3.18	10.60	21.20
Japan	2	0.42	2.12	3.18	10.60	21.20
Spain	2	0.42	2.12	3.18	10.60	21.20
Portugal	2	0.42	2.12	3.18	10.60	21.20
Switzerland	2	0.42	2.12	3.18	10.60	21.20
Sweden	2	0.42	2.12	3.18	10.60	21.20
South Korea	1	0.21	1.06	1.59	5.30	10.60
Germany	1	0.21	1.06	1.59	5.30	10.60
Total	100	21.20	106.00	159.00	530.00	1060.00

Figure 6: SRCNet storage allocations. Deployment of the SRCNet storage in line with the share per

SRC v0.1 Compliance Summary

SRC v0.1 Compliance Summary

SRC	Spain SRC	Netherlands SRC	Sweden SRC	UK SRC	Switzerland SRC	China SRC	Canada SRC	Japan SRC
Compliance Level	Meets most requirements	Meets most requirements	Meets most requirements	Meets most requirements	Meets most requirements	Meets most requirements	Partially Compliant	Meets most requirer
Summary	Strong commitment, limitations in storage & compute capacity	Strong proposal, that emphasizes scalability & long-term collaboration	Strong proposal, focuses on collaboration & leveraging WLCG	Strong proposal, leverages existing expertise & infrastructure	Strong commitment with existing services, limitations in storage & compute due to funding	Strong foundation, existing deployments, future depends on funding	Partially compliant solution Available storage used for precursors data Possible limitations in operations	Strong proposal, alt limited in operation
Software Stack	Compatible, experience with all listed tools (Mini-SRCNet Demonstrator)	Compatible, experience with all listed tools, containerized deployments	Compatible with proposed stack, experience with WLCG environment	Compatible with proposed stack, experience with large research projects	Compatible, experience with most listed tools, leverages GitOps for deployment	Already have: Rucio, JupyterHub, CANFAR, Kubernetes/Slurm, containerized visualization tools	Already running most required services	To be determined (li compatible)
People (FTEs)	1 FTE dedicated for v0.1, contributes to long-term team	2 FTEs for v0.1, expertise in cloud & distributed systems	1 FTE for v0.1, access to additional WLCG experts	1 FTE for v0.1, dedicated team planned for future	> 1 FTE committed, contributes to long-term operational team	1 FTE (committed to exceeding for long-term team)	Local operations support but not availability for other SRCs support or to receive support from another SRC	0.175 FTEs (agreed u SRCSC)
Operations (Turnaround)	3 days for production issue resolution	2 days for production issue resolution	3 days for production issue resolution	3 days for production issue resolution	3 days for production issue resolution (leveraging CSCS ticketing system)	3 days for production issue resolution	Maybe. IT delegated to Digital Research Alliance of Canada	Possible (3 working
Network (Gbps)	10 (RedIRIS-NOVA, GEANT access)	100 (dedicated line), considering upgrades for future needs	100 (connected to national research network)	10 (minimum), planning for future upgrades	10 (hosted at CSCS), open to participating in bandwidth tests between SRCNet sites	Already have 1Gbps to other SRCNet sites; upgrading to 10Gbps by end of 2025	Confirmed with CANARIE, capacity exists; working with NRENs for SKAO data	Minimum 10 (planne to 100Gbps), 400Gb
Network (IPv6)	Not currently implemented	Implemented	Implemented	Implemented	Implemented	Implemented	Expected	Implemented
Storage (PB)	0.5PB usable by Jan 2025 (funding limitations)	0.1 PB / 100 TB available for the duration of v0.1.	0.29 PB	4PB planned for the procurement of in FY2024/25.	0.40 PB by Jan 2025 (funding limitations), open to oversubscription plans	1 PB by 2024, 2 PB by 2025	1.2 PB (expandable to 4.6 PB if needed) 0.6PB for storage inventory, precursors data 0.6PB for user storage no mention of rucio (but UVic currently does run an RSE for SKA and more for ATLAS and Belle II)	0.144 PB
Compute (PFLOPS)	Planned 0.01 PFLOPS	0.01 PFLOPS by Jan 2025 (existing resources & leveraging cloud)	10.5 TFLOPS 0.0105 PFLOPS	0.175 PFLOPS	0.014 PFLOPS by Jan 2025	0.175 PFLOPS	0.04 PFLOPS (expandable to 0.11 PFLOPS)	0.005 PFlops (FP64)

1Gbps => 10Gbps

CSTNet in the Coming 3 Years from CNIC

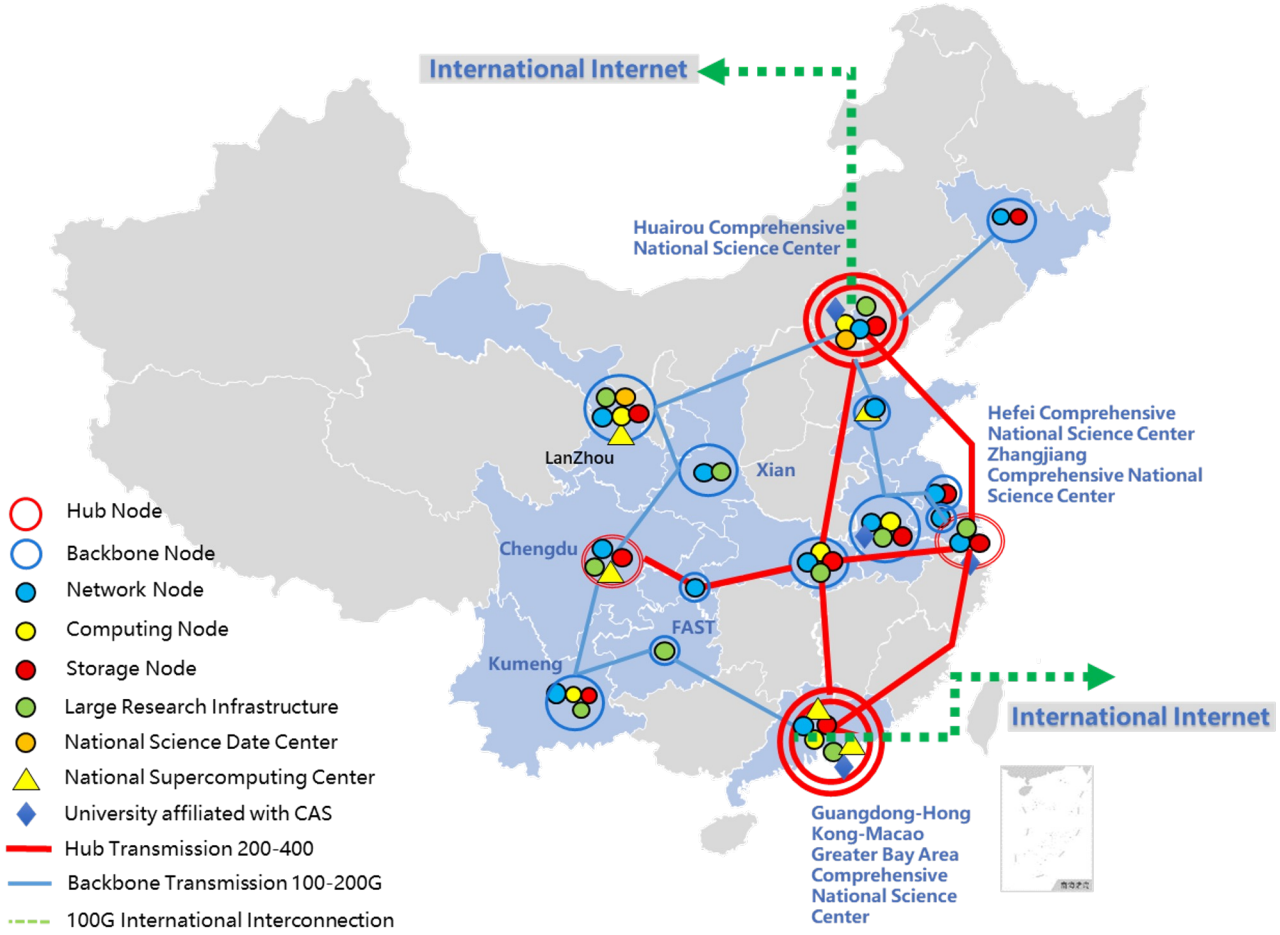

Human-machine-physical integration information infrastructure innovation and test platform

- **16 cities ,140 research facilities**
- SDN , IPv6+, SRv6 Service-oriented Slice
- Multiple 100G international network
- Integrating 1000 Pflops
- 100PB data storage backup
- AIOps platform
- Cloud network integration technology innovation and verification platform

Some of previous work

rucio

carta

aladin

AAI

+70 software

Over 300 versions

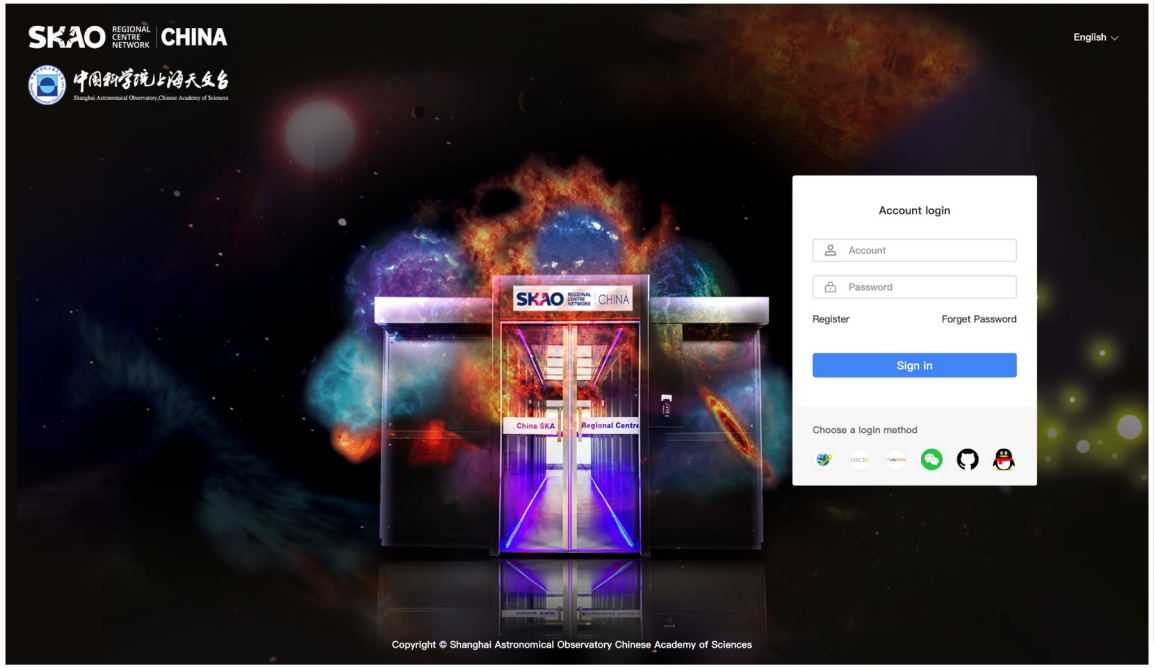

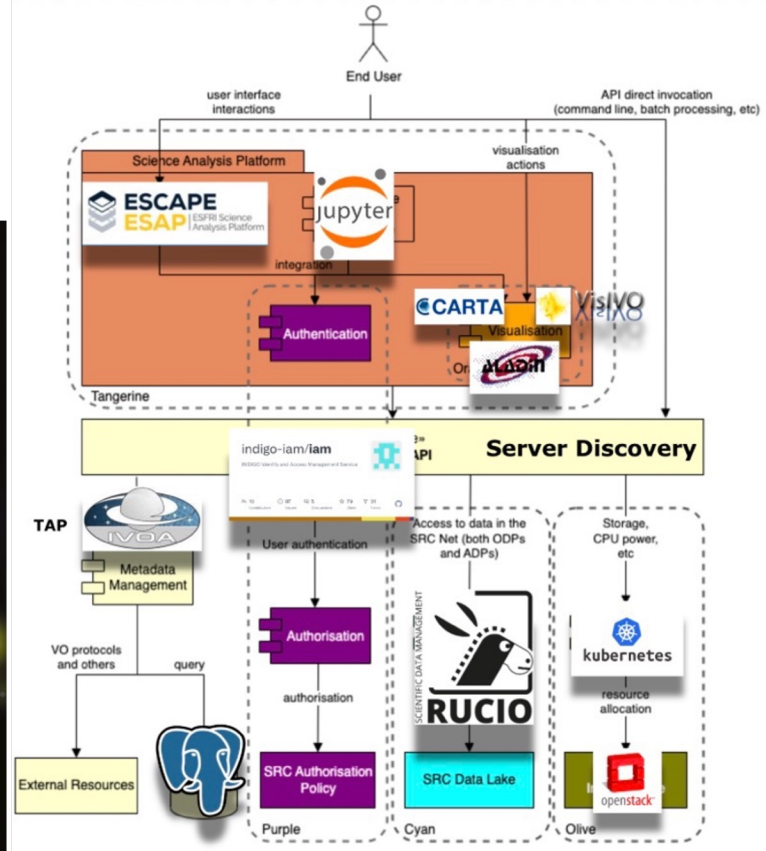

Src\Dst	SPSRC_STORM	JPSRC_STORM	IMPERIAL	CNSRC
STFC_STORM	100%	100%	100%	10
SPSRC_STORM	NO DATA	100%	100%	10
JPSRC_STORM	100%	NO DATA	100%	10
IMPERIAL	100%	100%	NO DATA	10
CNSRC_STORM	100%	100%	100%	NO
CNAF	100%	100%	100%	10
AUSSRC_STORM	100%	100%	100%	10

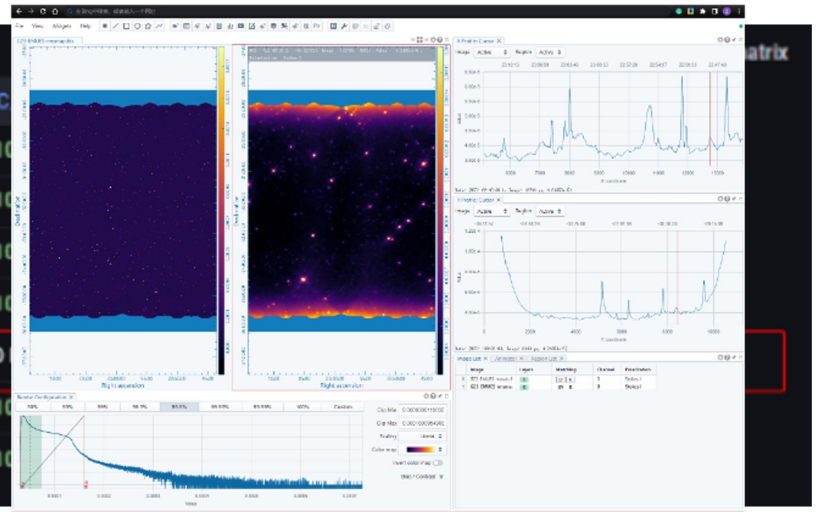

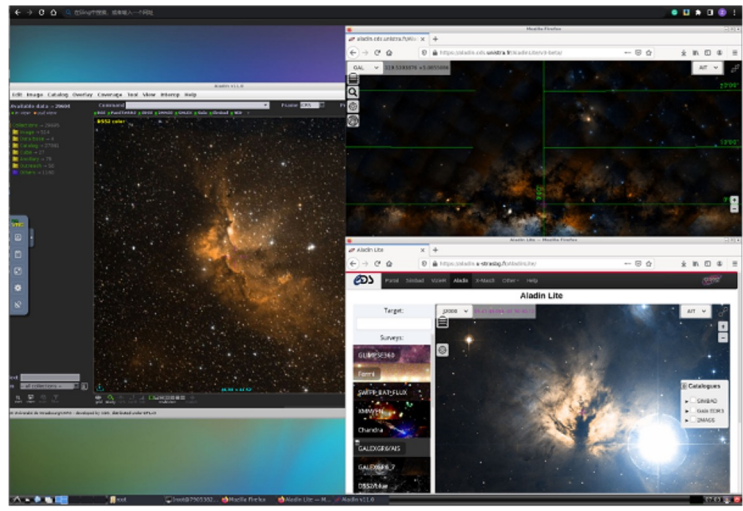

cnSRC is more than a cluster

- Heterogeneous compute resource to explore
  - Intel x86 CPU, Kunpeng ARM CPU, NVIDIA GPU
  - Tiered storage: NVMe online SSD, HDD archive storage
- Service covering SKA beginners to experts
  - Documents for beginners to start the first SRC task
  - Assist experts to port and optimize applications
  - Allow “guru” users to run root apps on bare metal nodes
- Wide interdisciplinary inter-institute collaboration
  - Collaborate with **Shanghai Jiao Tong University’s** HPC Team on BigData, HPC, and Cloud orchestration, who hosts top campus-level compute platform with 100K CPU cores + 65PB storage.

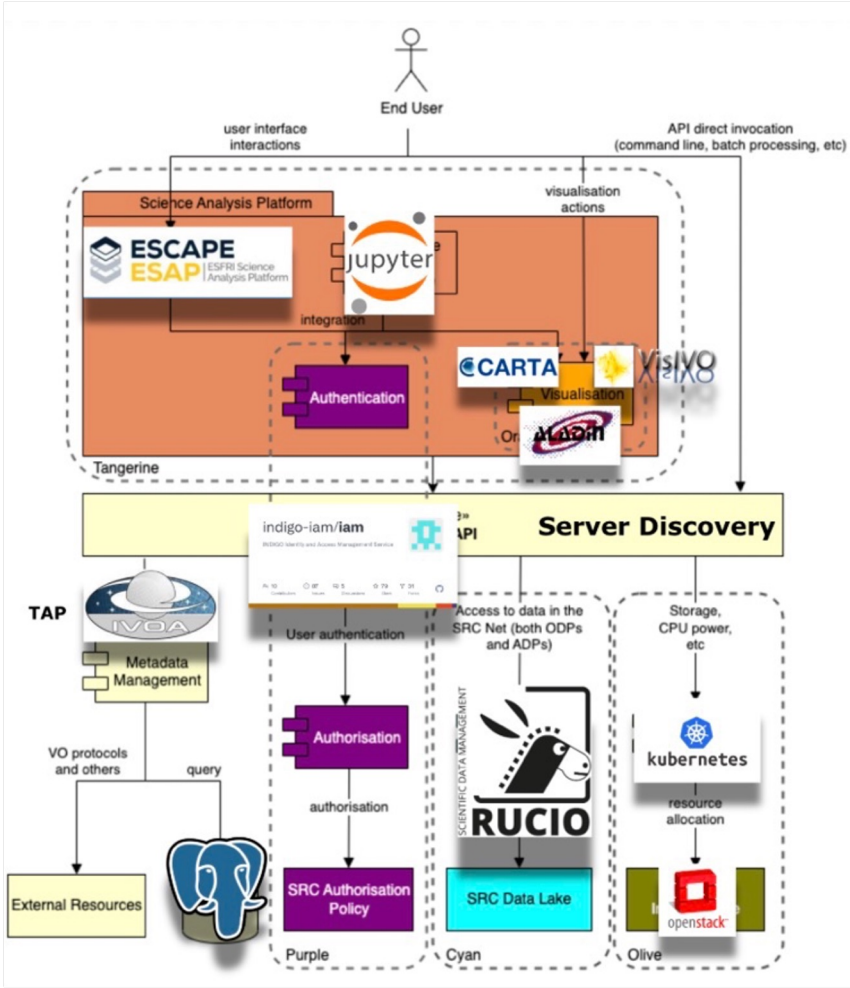

x86 CPU | NVIDIA GPU | ARM CPU

Our focus: Lesson from SKA Pathfinder

- **Lessons:**
  - **MWA /GLEAM**
  - ASKAP /VAST

Hurley-Walker et al. 2022 **Nature** Polarization data processing was partly completed at China SRC

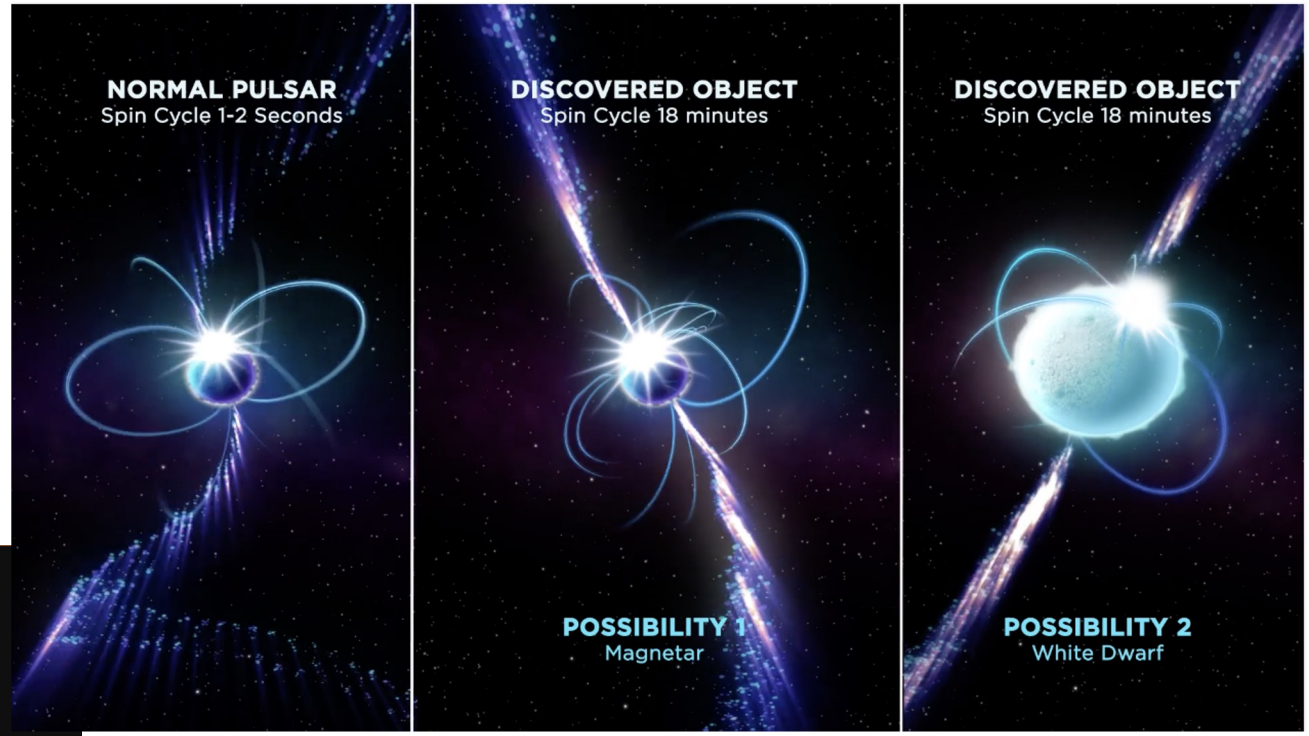

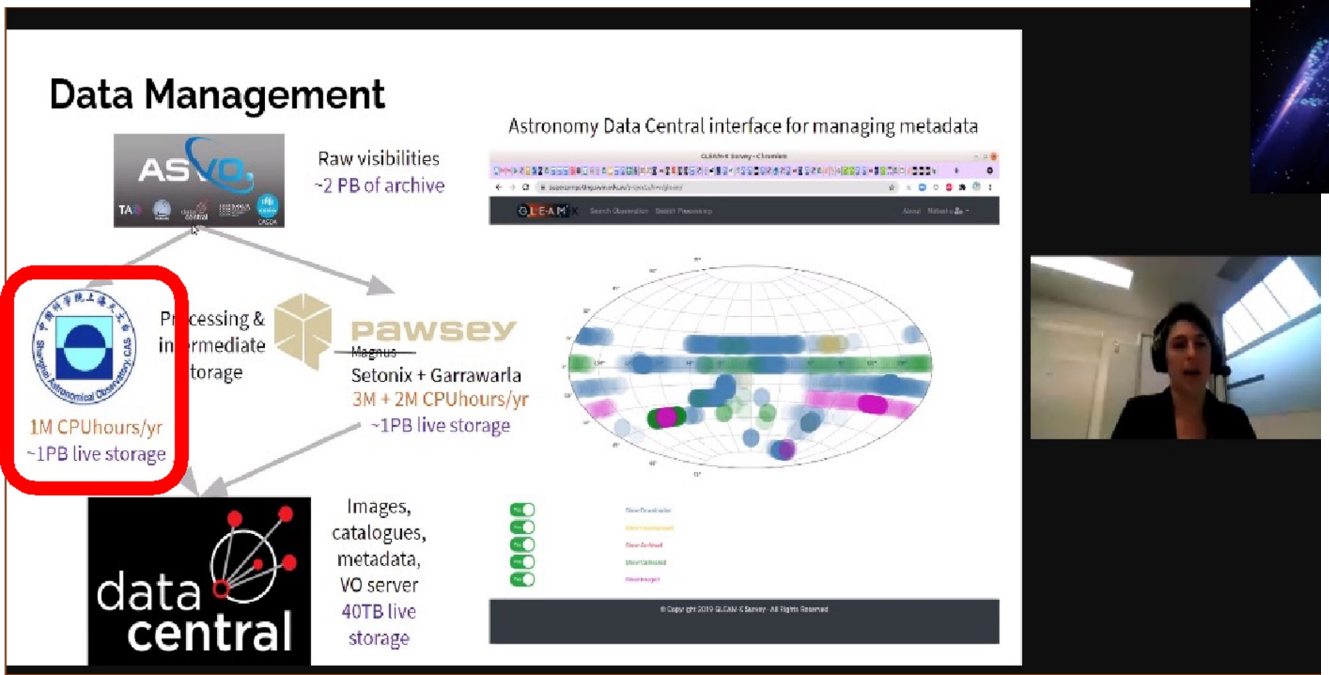

one of the two data processing and storage centres for the GLEAM-X project

Our focus: Lesson from SKA Pathfinder

Tao An, et al. 2023

- **Lessons:**
  - MWA /GLEAM
  - **ASKAP /VAST**

A series of papers using VAST data and the prototype of the China SKA Data Centre, attention by international media. VAST team has signed MoU with China SRC

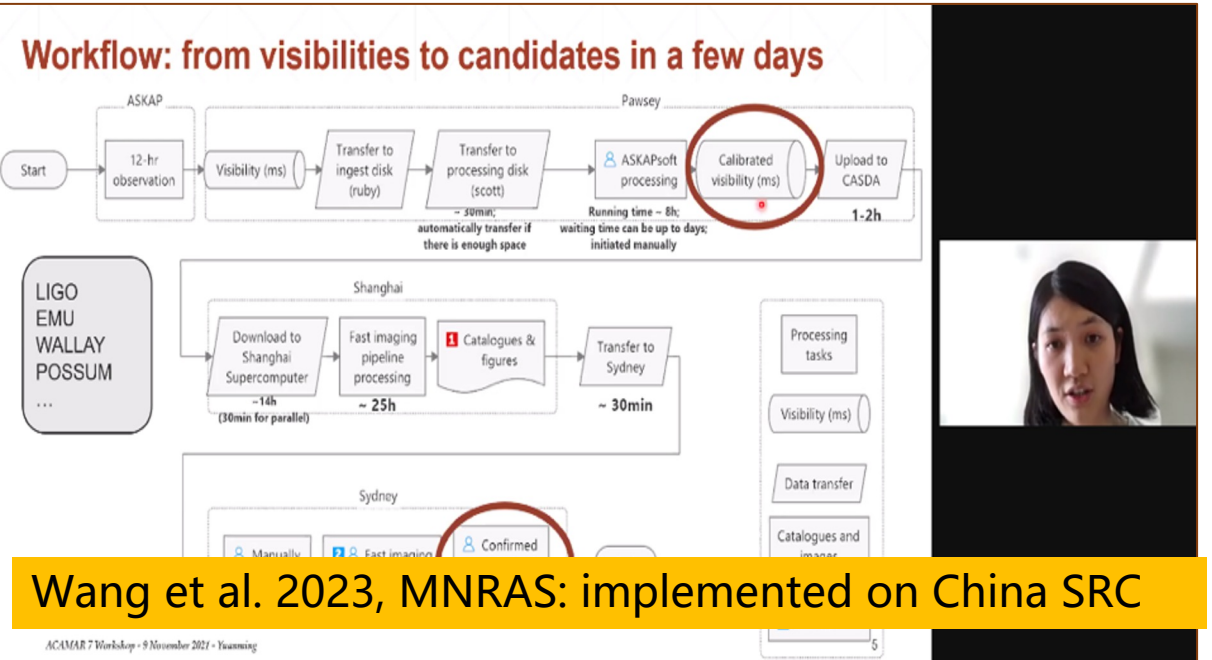

Wang et al. 2023, MNRAS: implemented on China SRC

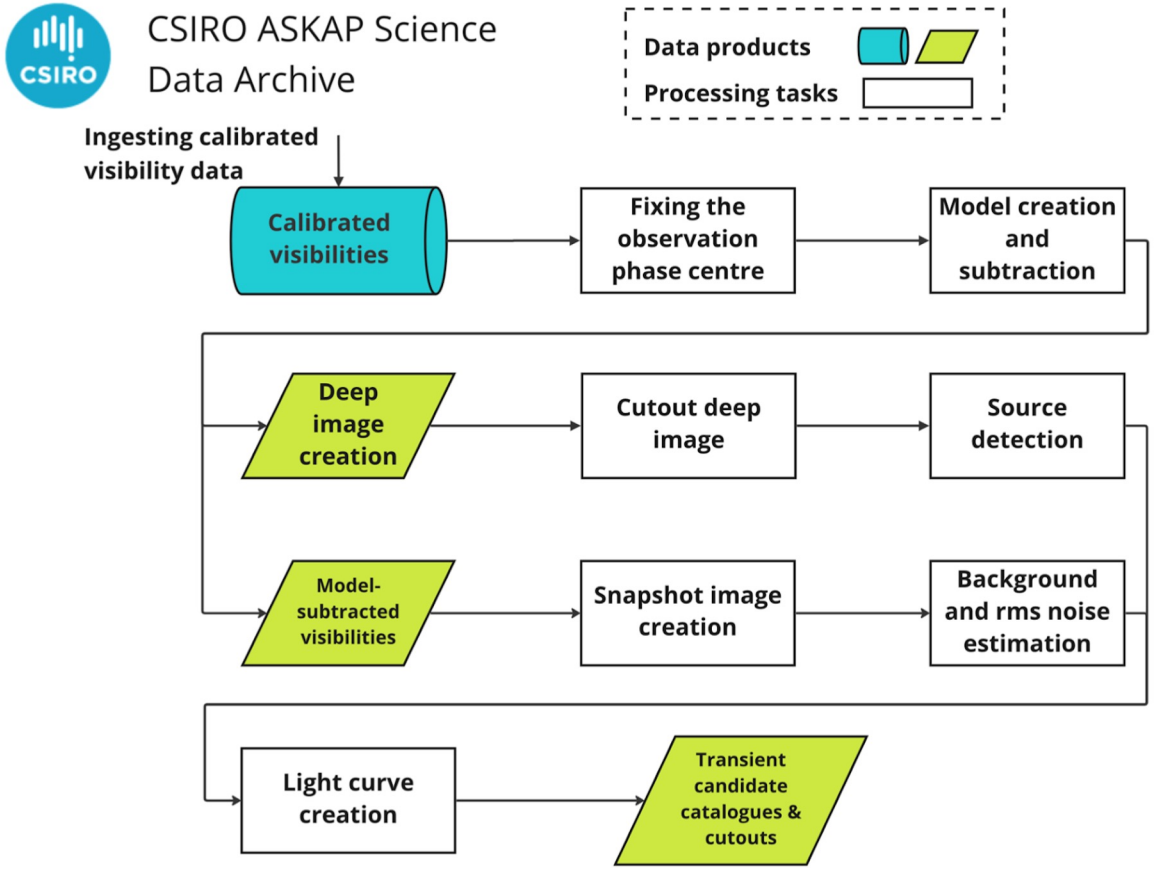

- The optimized pipeline is integrated into DALiuGE, can be used in supercomputers
- The method other ASKAP pipelines.

Our focus: Support SKA1 earliest data

- **Plan:**
  - **cnSRC will support SKA1 earliest data processing**

Construction steaming ahead! - Mid


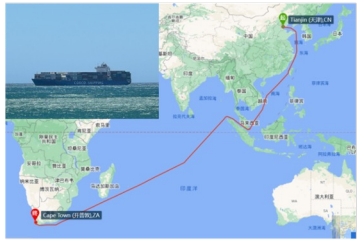


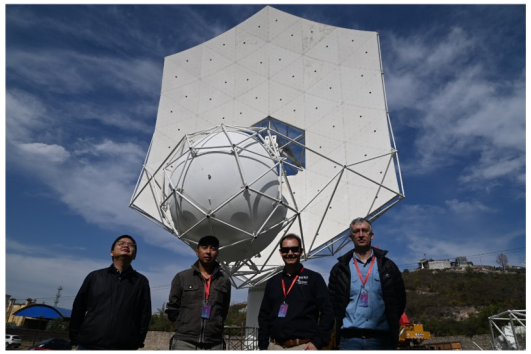

Construction steaming ahead! - Low

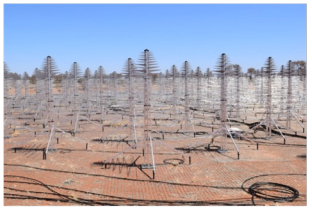

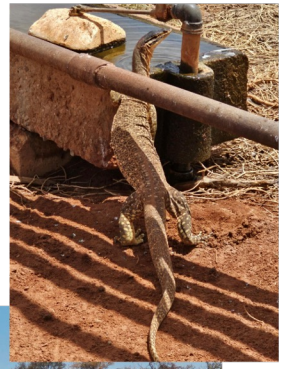


Milestone event (earliest)		SKA-Mid (end date)	SKA-Low (end date)
AA0.5	4 dishes 6 stations	2025 May	2024 Nov
AA1	8 dishes 18 stations	2026 May	2025 Nov
AA2	64 dishes 64 stations	2027 Apr	2026 Oct

Advantage of cnSRC Node

- **Dedicated for SKA prototype, fully controlled**
- **Experience of SKA pathfinder data processing, had generate a lot of science outcomes**
- **Robust team for SKA SRC**
- **Science support for**
  - **EoR/Pulsar/Magnetic/Transient/VLBI/Continuum/SDC**

Challenges and Limitations

- Possible limitations in resource provision in terms of hardware and funding for future expansion.
- Challenges in meeting the minimum requirement of 5PB of online storage and 0.175 PFLOPS compute capability by Jan 2025.
- Resource share plan
  - Storage：full time
  - Computer：peak/partly
- Exposed to several main SRC Node（whitelist)

Collaboration and contribution plans

- **Collaboration with other SRCs to enhance the SRCNet v0.1 deployment.**
  - **SKAO/UKSRC/AUSSRC/JPSRC/KRSRC/SPSRC…**
  - **Institute, University, company, …**
- **Potential areas for contribution beyond hardware resources: Software development, operations monitoring, supporting operational tasks.**