
SKACH
SKAO
CHSRC v0.1
SKACH Winter Meeting 2025, Bern
January 27th, 2025
Presenters: Lukas Gehrig, Pascal Herzog, Rohini Joshi, Pablo Llopis
skach.org
#SKACH
SARAO, Heywood al. (2022) / J.C. Muñoz-Mateos

Overview

Introduction

- SRCNet v0.1, CHSRC v0.1
- Timeline
- v0.1 Test Campaigns

Showcase of the CHSRC v0.1 node

Ongoing developments

SRCNet v0.1

Overview, CHSRC participation, Timeline

SRC Network Vision

We will develop and deploy a collaborative and federated network of SKA Regional Centres, globally distributed across SKA partner countries, to host the SKA Science Archive. The SRC Network will make data storage, processing and collaboration spaces available, while supporting and training the community, ***to maximise the scientific productivity and impact of the SKA.***

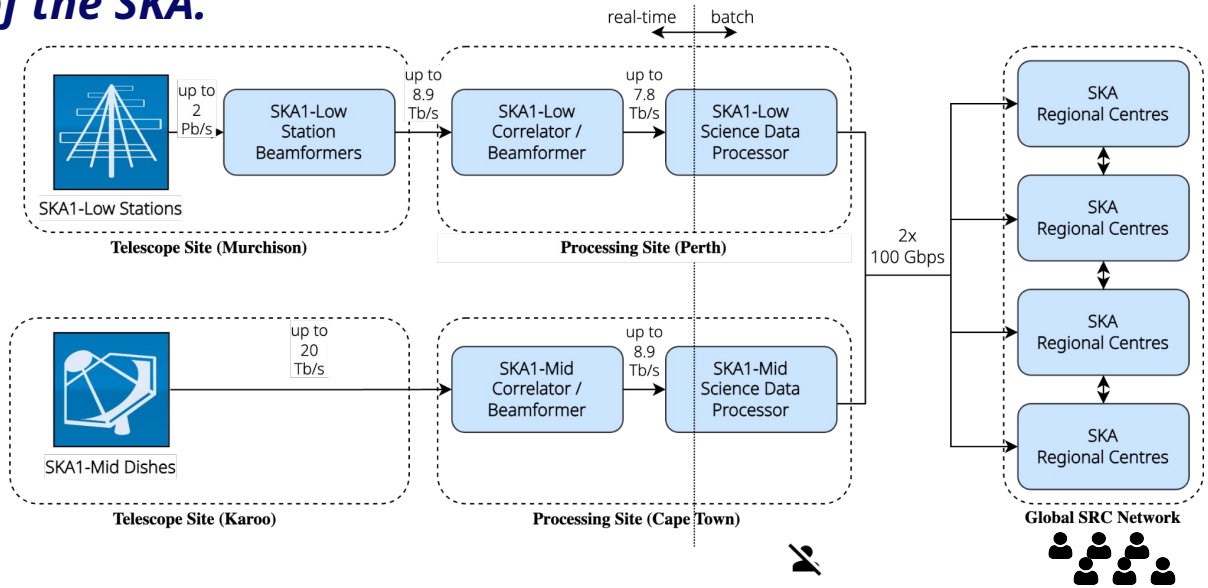

SRC | Net
SKAO Regional Centre Network

SRC Network Vision

We will develop and deploy a collaborative and federated network of SKA Regional Centres, globally distributed across SKA partner countries, to host the SKA Science Archive. The SRC Network will make data storage, processing and collaboration spaces available, while supporting and training the community, ***to maximise the scientific productivity and impact of the SKA.***

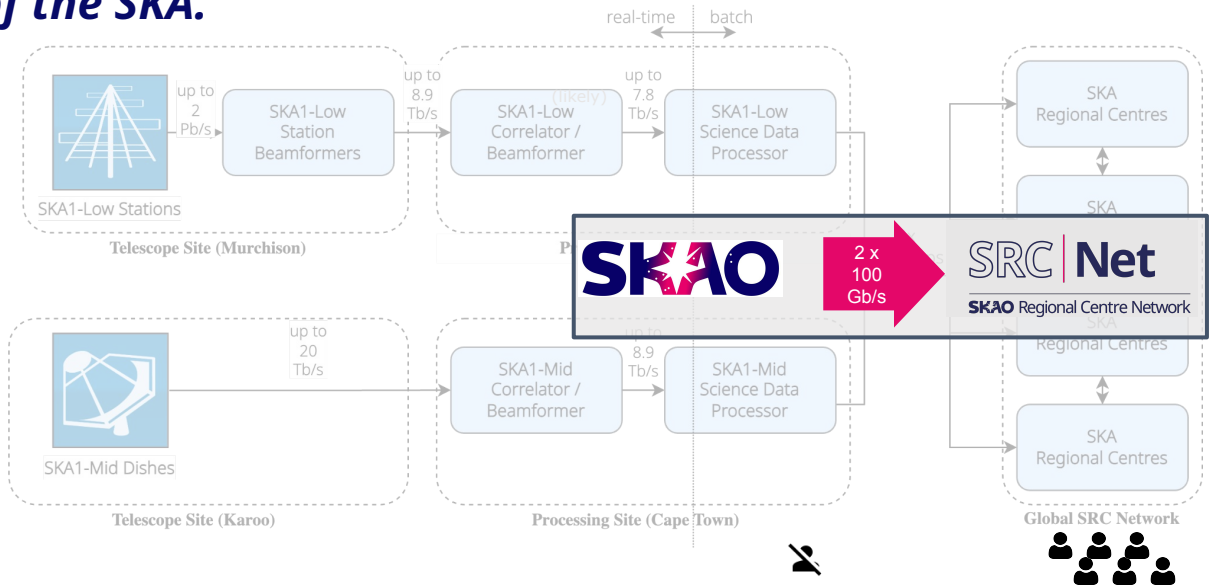

SRC Net
SKAO Regional Centre Network

Basic Functionality Covered by v0.1

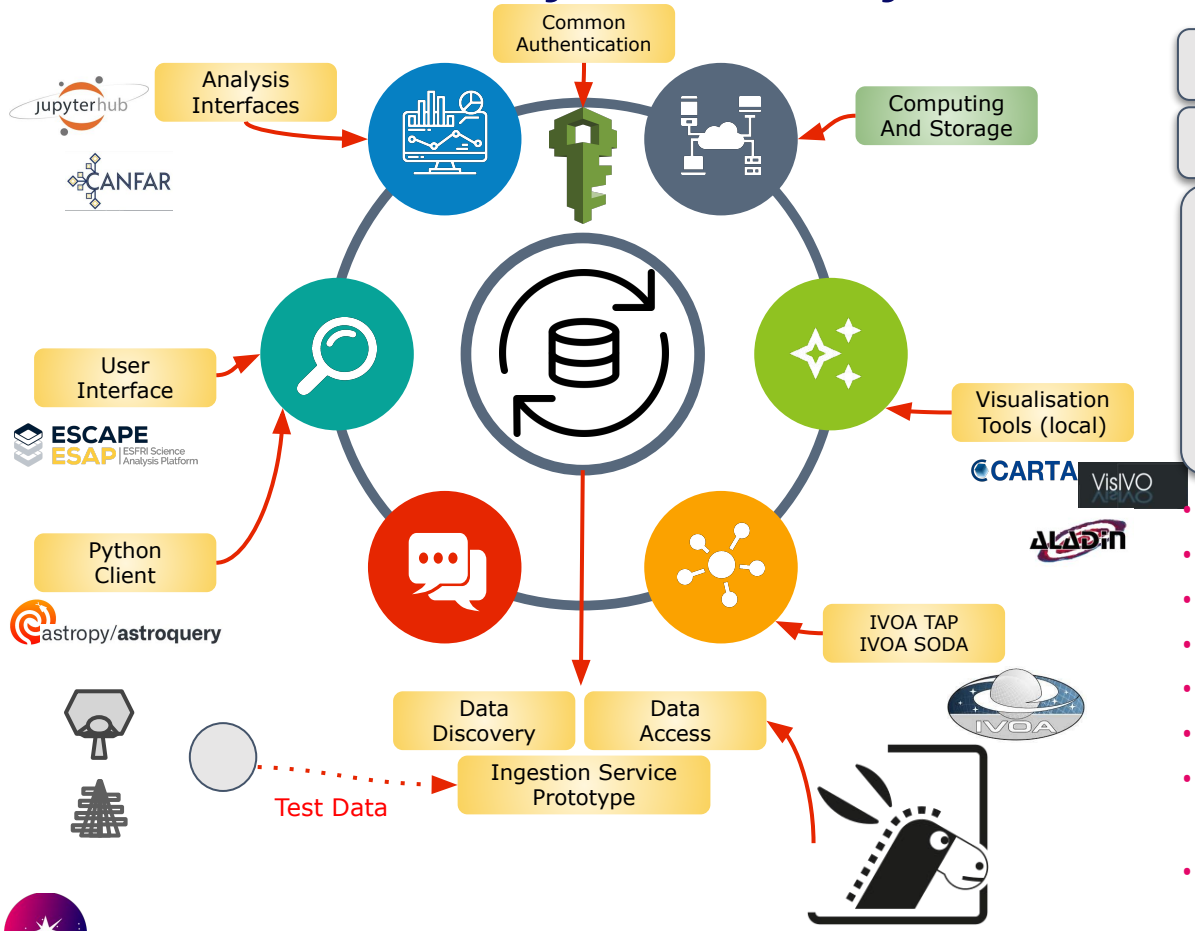

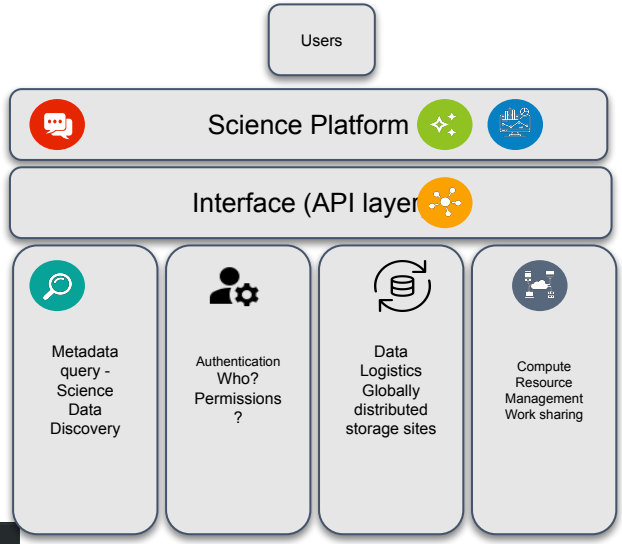

- Common Authentication: IAM
- Visualisation Tools (local)
- IVOA Protocols: TAP, SODA
- Data Discovery and Access from Data Lake
- Ingestion Service Prototype
- Python Client: Astroquery Module
- User Interface: Gateway (ESAP)
  - https://esap.srcdev.skao.int/
- Analysis Interfaces: JupyterHub; CANFAR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

Basic Functionality Covered by v0.1

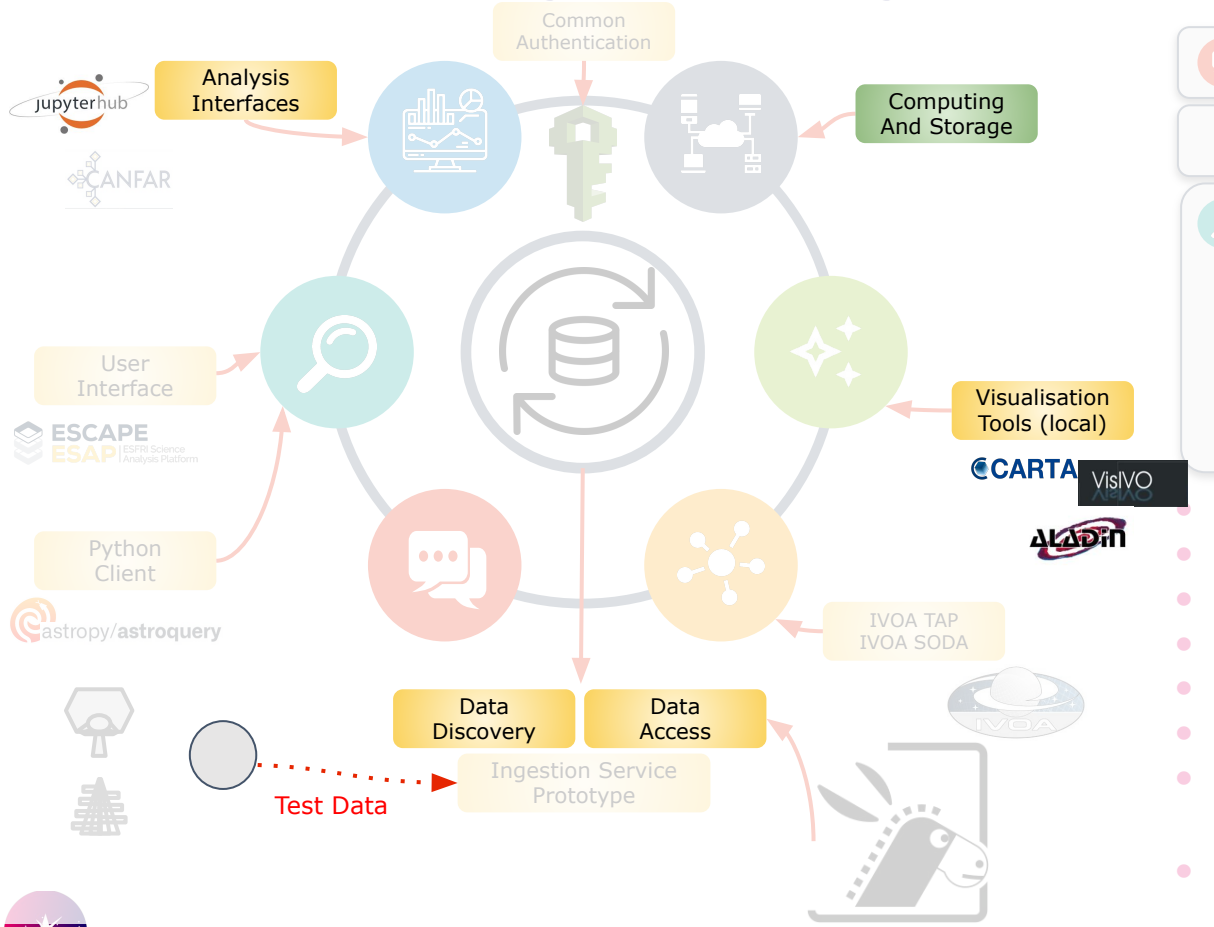

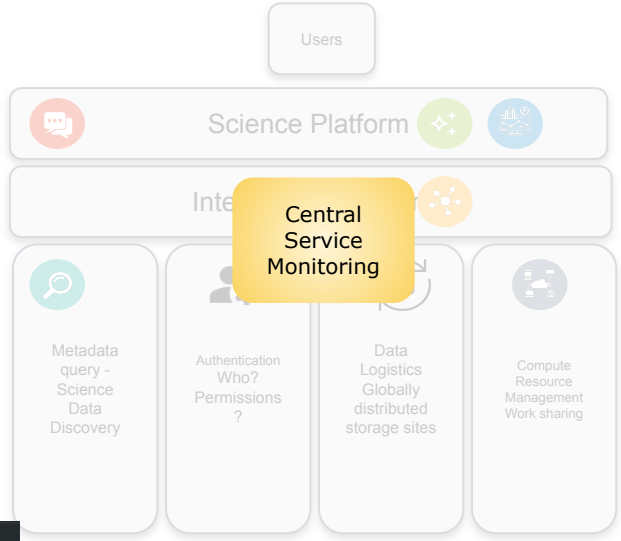

- Common Authentication: IAM
- Visualisation Tools (local)
- IVOA Protocols: TAP, SODA
- Data Discovery and Access from Data Lake
- Ingestion Service Prototype
- Python Client: Astroquery Module
- User Interface: Gateway (ESAP)
  - https://esap.srcdev.skao.int/
- Analysis Interfaces: JupyterHub; CANFAR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

SRCNet v0.1, CHSRC v0.1

SRCNet v0.1 is primarily an engineering exercise

- Validate hardware at distributed sites, interconnected storage
- Common set of local services deployed
- Operations group, establish ways of working across 9 sites

What does v0.1 mean for CHSRC

- 14 TFlops, 400 TB Storage at CSCS via Gornergrat kubernetes cluster(s)
- Services deployed on this infra, integrated, ready for test campaigns

CHSRC collaboration and representation

ch | SRC
SKAO Regional Centre Switzerland

SKACH

Collaborations

- Working with Spanish SRC and SKAO team to build/modify a data staging service
- Karabo data ingestion into SRCNet DM system
- GitOps methodologies adopted by Spanish, Swedish SRC
- Network monitoring setup, perfSONAR, UKSRC
- Aus SRC - MWA Data ingestion using CHSRC node/services

CHSRC representation

- SRCNet Resource Board (Carolyn M Crichton)
- SRCNet Technical Advisory Committee (Elena Gavagnin)
- Chocolate Team (Carolina Lindqvist)
- Operations Group, Security PoC (Pablo Llopis)
- Science Delivery Value stream, Teal Team (Carlo Ferrigno)
- Architecture core team (Rohini Joshi)

Accounting and metrics still a WIP in SRCNet, representative metrics from Sep-Dec 2024

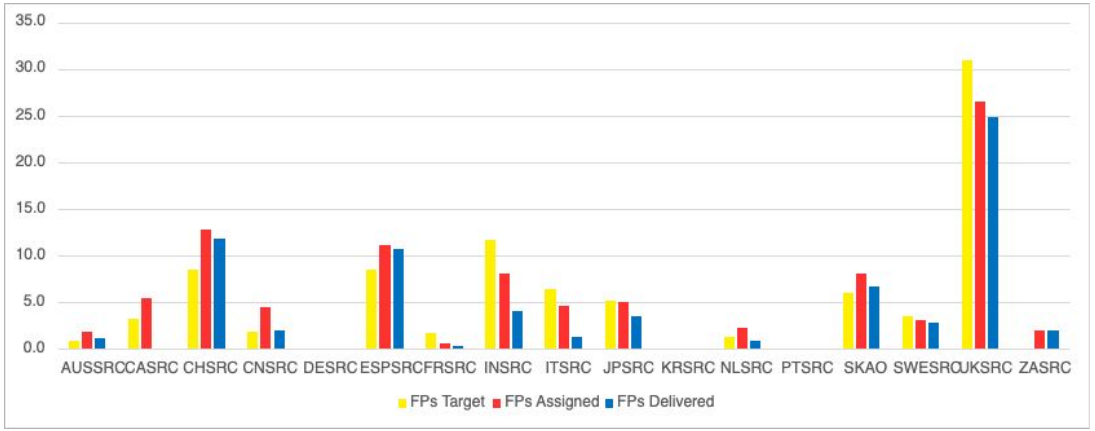

SRCNet Timeline

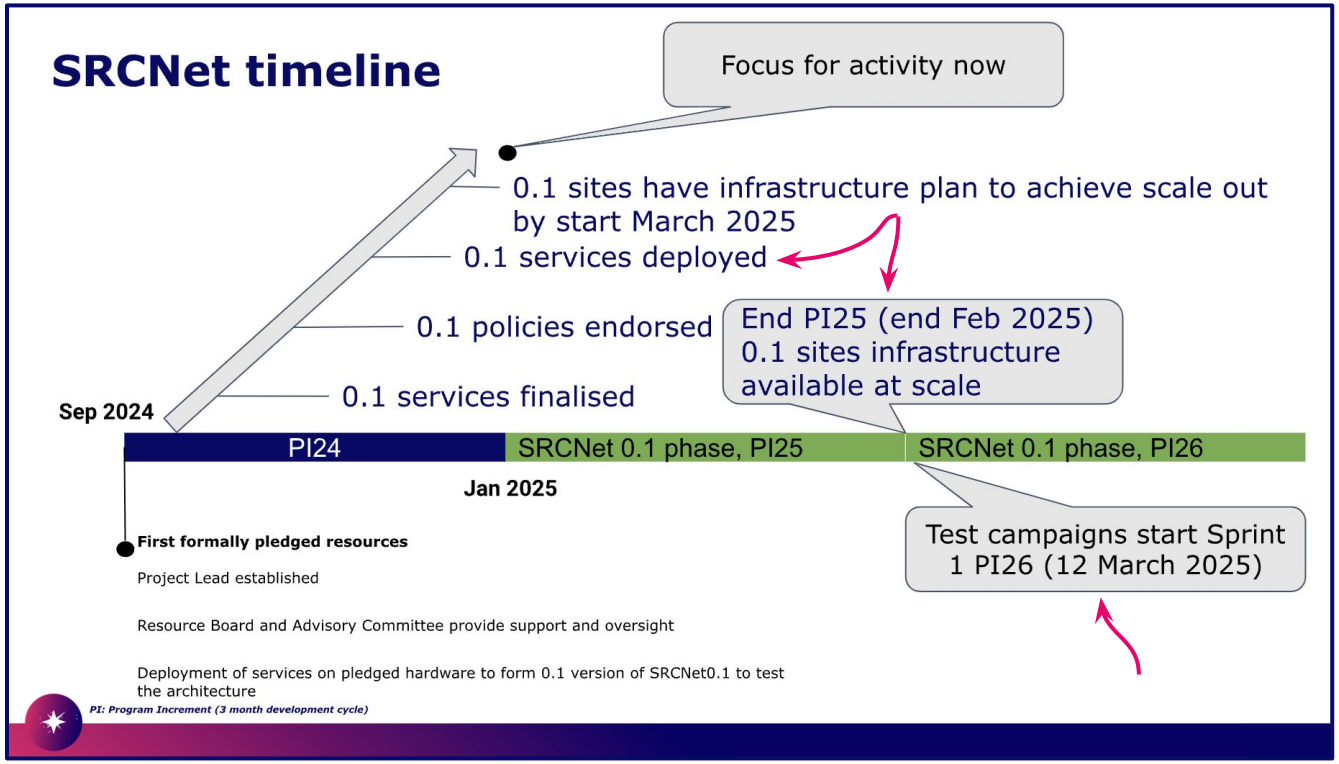

Slide Credit: Rosie Bolton

SRCNet Timeline

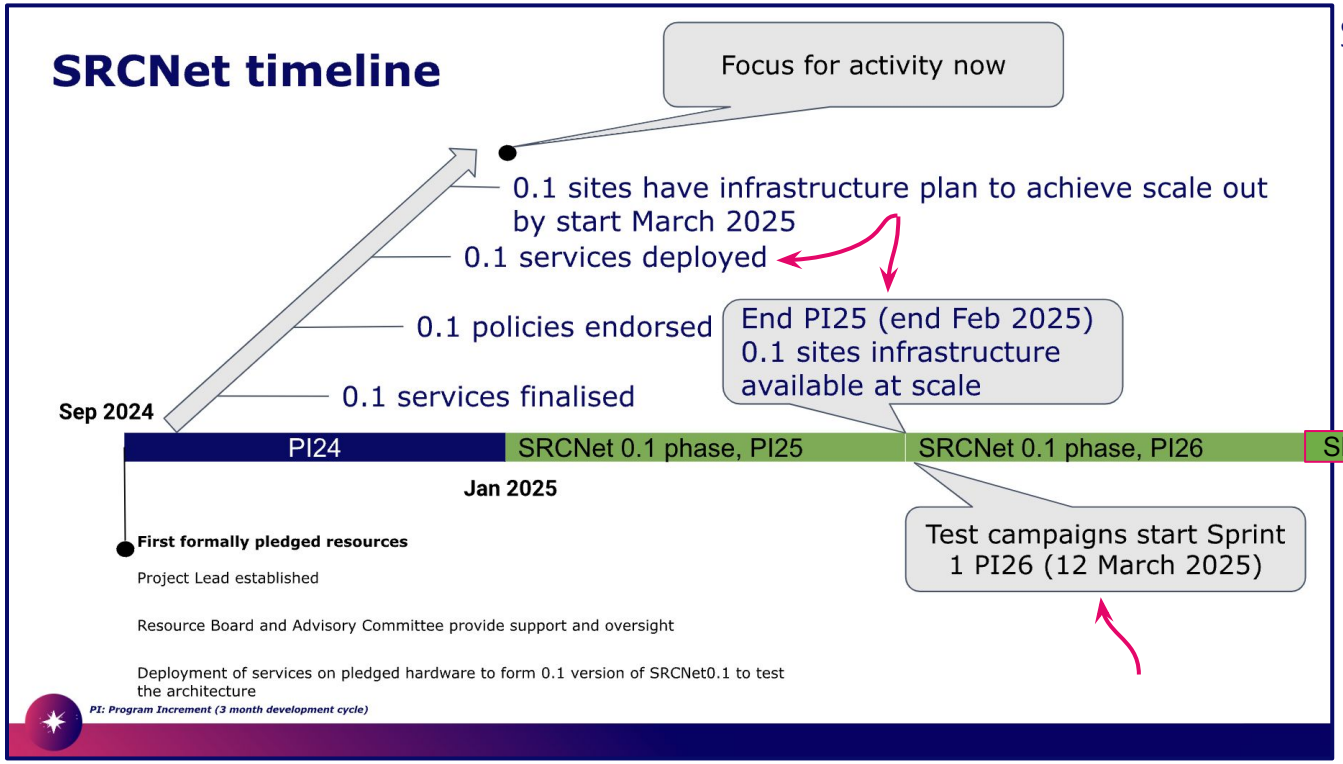

Slide Credit: Rosie Bolton

- SRCNet v0.2 development
  - Federated execution
- Science platform development
- Local CHSRC science users
- Proprietary data?

v0.1 Test Campaigns

Data transfer test campaigns ala WLCG Data Challenges

Data lifecycle tests

- Science representative

Integration tests across services and sites

Operations group - ‘social testing’

Areas of work

‘Staged’ storage for running workflows

- ‘prepareData’ implementation
- Exploring dCache vs xrootd

SRCNet v0.1 readiness at scale

- Gornergrat cluster delivery and migration

	SP SRC	NL SRC	SW SRC	UK SRC	CH SRC	CN SRC	CA SRC	JP SRC	IT SRC	KR SRC	Total
Storage (PB)	0.500	0.100	0.300	4.000	0.400	1.000	1.200	0.651	0.300	0.270	8.711
Compute (PFLOPS)	0.010	0.010	0.011	0.175	0.014	0.175	0.040	0.022	0.100	0.010	0.567
Percentage Storage (%)	5.740	1.148	3.329	45.919	4.592	11.480	13.776	7.473	3.444	3.100	
Percentage Compute (%)	1.765	1.765	1.853	30.891	2.471	30.891	7.061	3.883	17.652	1.765	
Harmonisation Rate *	0.020	0.100	0.036	0.044	0.035	0.175	0.033	0.034	0.333	0.037	

Operations

- Service requirements, docs, support
- Monitoring dashboard

Showcase of the CHSRC v0.1 Node

High level overview of progress made on CHSRC infrastructure

CHSRC Stack Goals (from SKA days 2024)

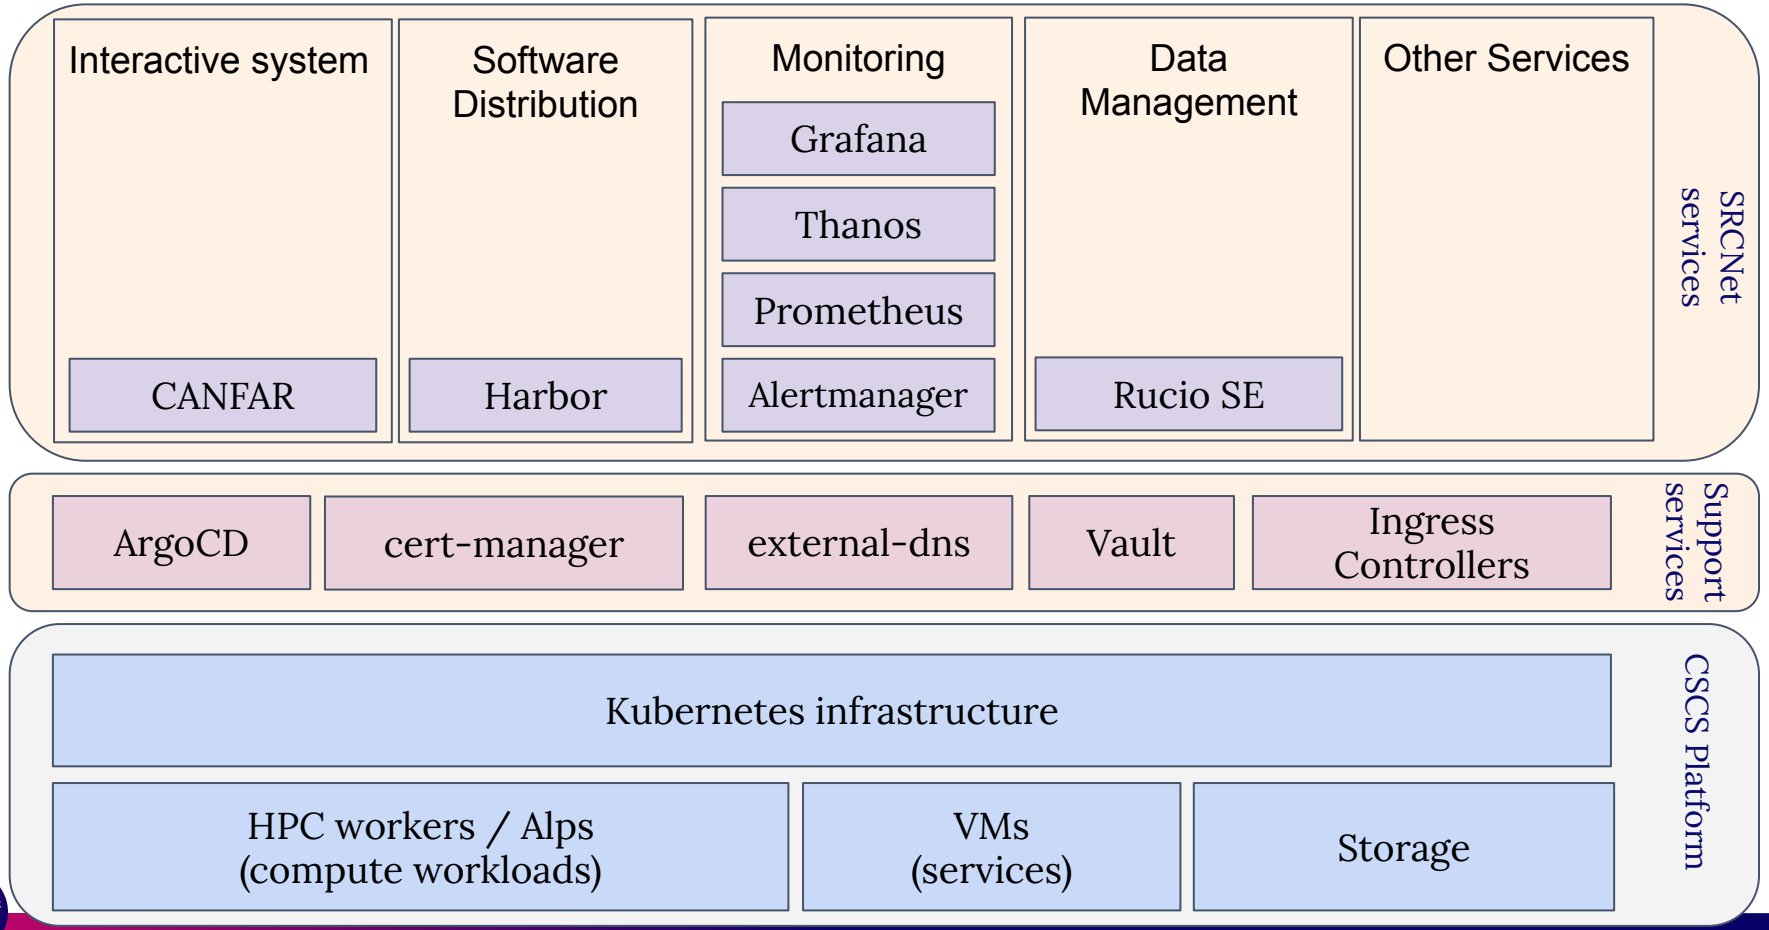

CHSRC Stack (Today) x2

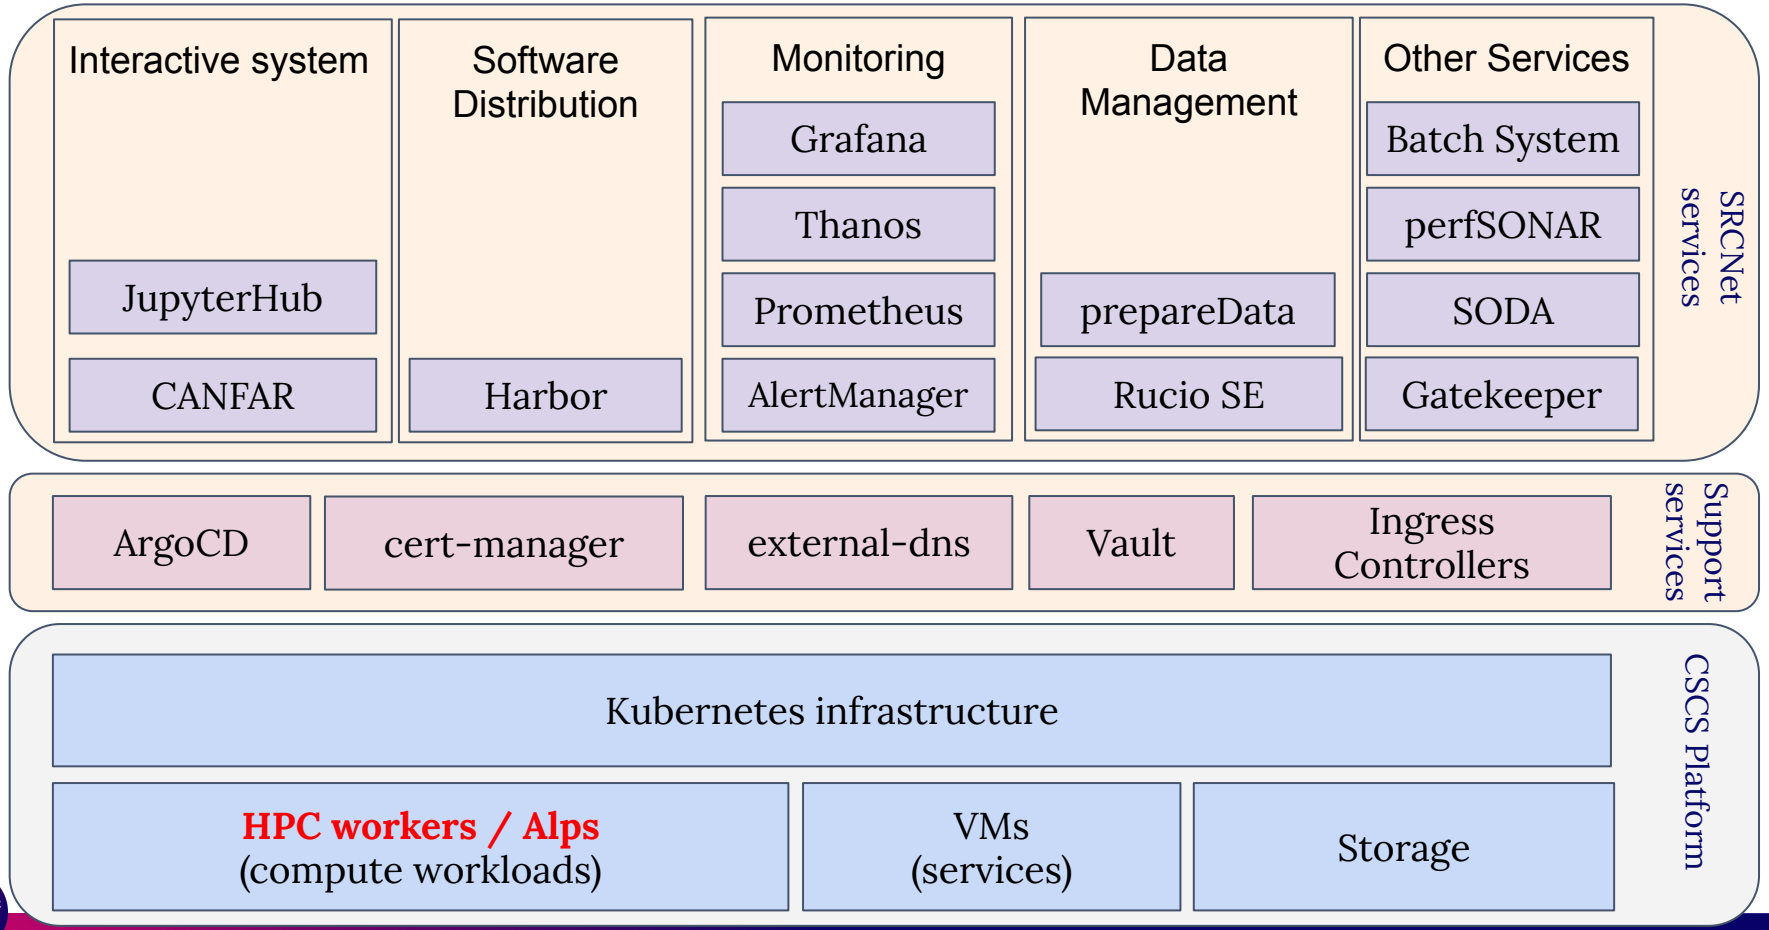

SRCNet service deployment status by participating countries

	espSRC	sweSRC	ukSRC	chSRC	cnSRC	canSRC	jpSRC	itSRC
Meta Services								
Orchestrator Meta-Service	Y	Y	Y	Y	Y	Y	Y	Y
Common Software Repository (GitOps)	Y	Y	Y	Y	N	IP	IP	Y
Compulsory Services								
Local Data - related services: Rucio Storage Element (RSE)	Y	Y	Y	Y	Y	Y	Y	Y
Gatekeeper (to permissions interface)	Y	N	Y	Y	N	IP	IP	Y
SODA - Parsing Local to data for visualisation of remote data (extensions of IVOA SODA services) exposed by Data Management/GateKeeper API	Y	N	Y	Y	N	IP	IP	Y
Containerised visualisation tools to visualise local data (same as Science Platform instances)	Y	Y	Y	Y	N	N	IP	Y
Registering in Monitoring Dashboard	Y	Y	Y	Y	Y	N	Y	IP
JupyterHub or compatible	Y	Y	Y	Y	Y	N	IP	IP
perfSONAR	Y	Y	Y	Y	N	Y	Y	IP
prepareData	IP	B	B	IP	B	B	B	B
Infrastructure Availability								
Storage	Y	Y	Y	Y	Y	Y	Y	IP
Compute	Y	Y	Y	Y	Y	Y	IP	IP
Networking	Y	IP	Y	IP	IP	Y	IP	IP

A day in the life of an SRCNet scientific task

I want to discover Observational Data Products.

I want to visualise data interactively.

I want to carry out computationally expensive workloads to process data products

A day in the life of an SRCNet scientific task

I want to discover Observational Data Products.

I want to visualise data interactively.

I want to carry out computationally expensive workloads to process data products

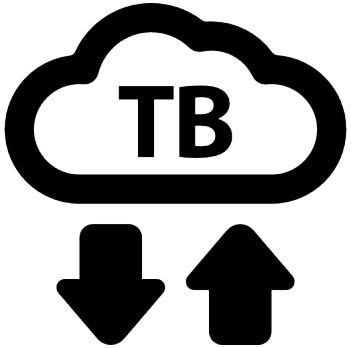

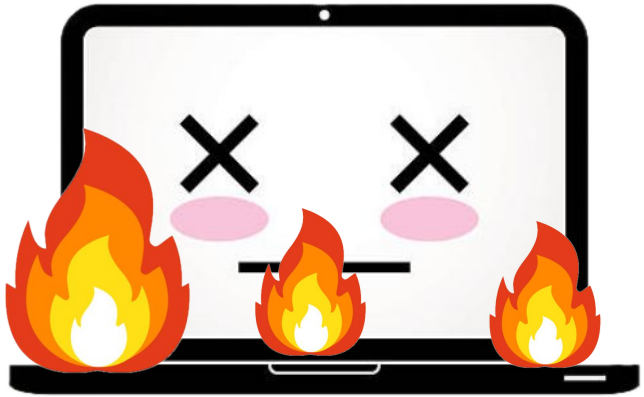

A day in the life of an SRCNet scientific task

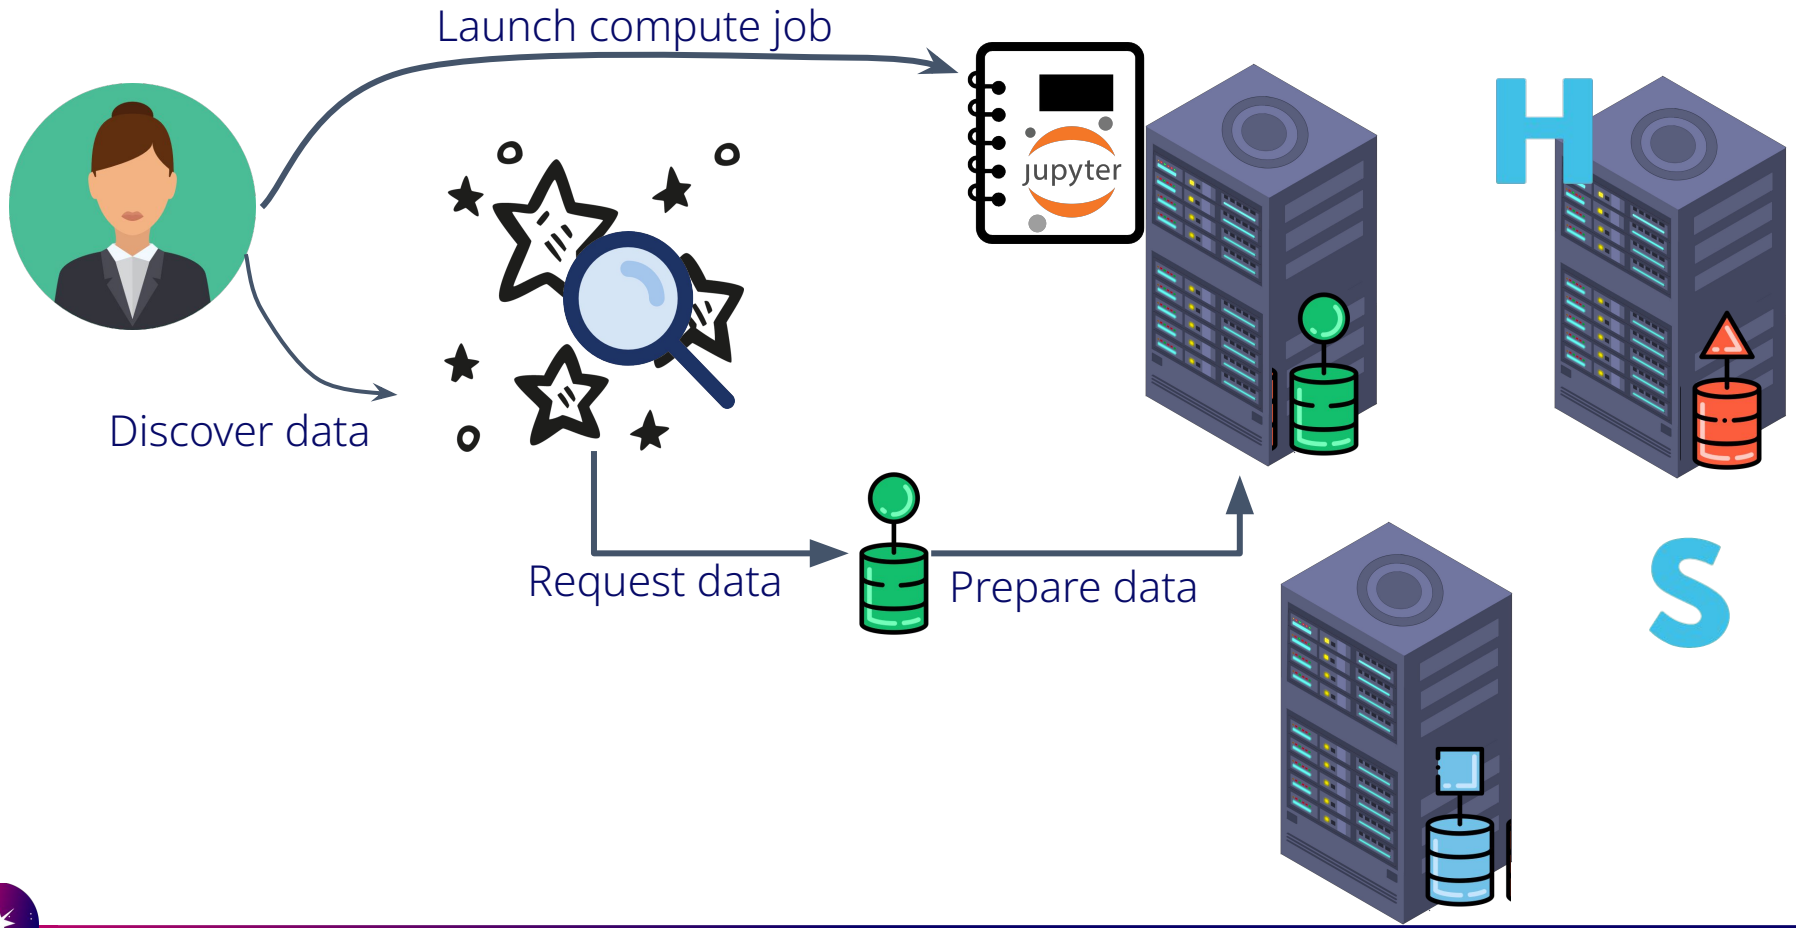

A day in the life of an SRCNet scientific task

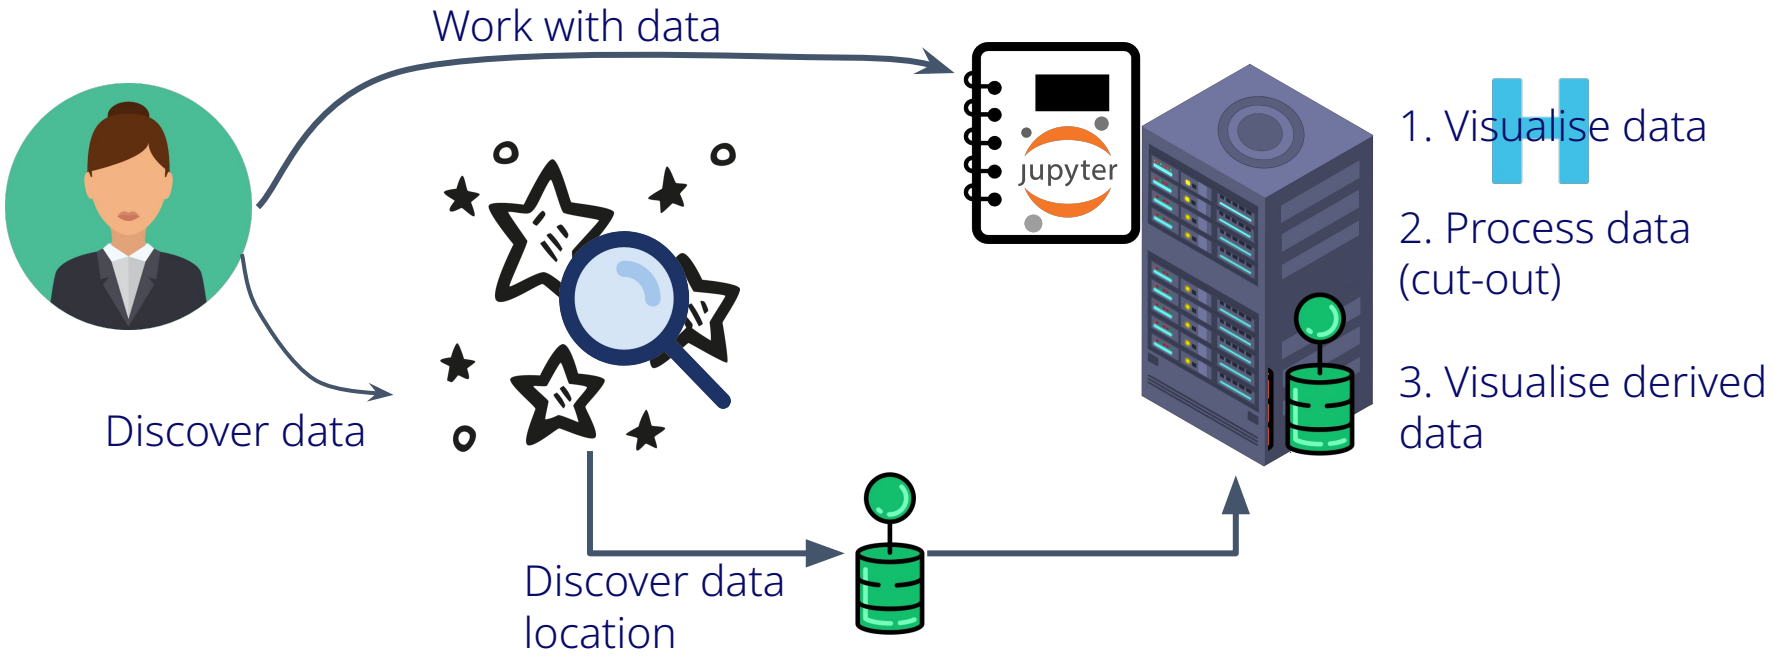

Not Secure https://ivoa.dachs.srcdev.skao.int/rucio/rucio/cone/form

DaCHS Virtual Observatory Publishing

Help

Service info

Metadata

Identifier
ivo://test.skao/rucio/ru

Cite this
Advice on citing this re

Description
SCS query service run

Keywords
subject

Creator
Could be same as cor

Created
2023-01-31T09:00:00

Data updated
2024-11-04T16:39:04

Metadata updated
2024-11-04T16:39:06

SKAO Rucio SCS

SCS query service running against an ObsCore table with a view on the Rucio database.

Position/Name
Coordinates (as h m s, d m s or decimal degrees), or SIMBAD-resolvable object

Search radius
Search radius in arcminutes

Table
Sort by _r ASC
Limit to 100 items.

Output format HTML More output fields

Go

[Result link]

Please report errors and problems to the site operators. Thanks.

Central Monitoring

- Centralized SRCNet service monitoring setup
- View on the local and global services running at the v0.1 SRCNodes
- Store and display historical monitoring data
- Rules-based Alerting

Central Monitoring - Diagram

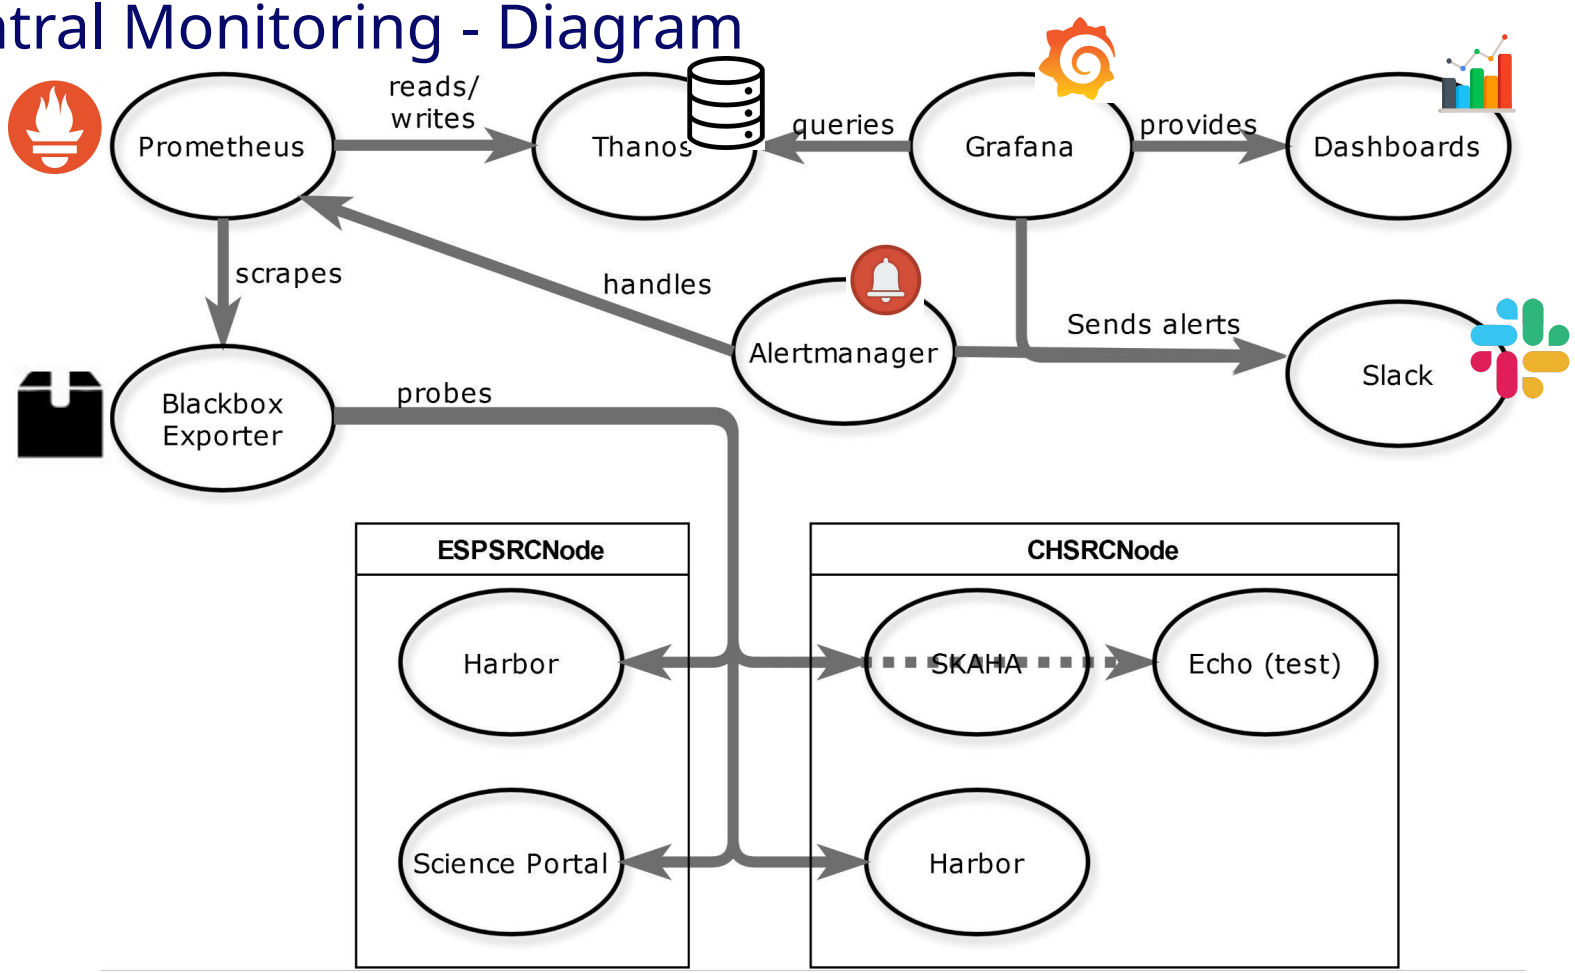

Local services

Overview

location\servicename	perfsonar	rucio-storage-element	cavern	science-portal	gatekeeper	jupyterhub
CANSRC	Online	Offline	N/A	N/A	N/A	N/A
UKSRC	Offline	Online	Online	Online	Online	Offline
ESPSRC	Online	Online	Online	Online	Online	Online
CNSRC	N/A	Online	N/A	N/A	N/A	N/A
CHSRC	Offline	Online	Online	Online	Online	Online
JPSRC	Offline	Offline	N/A	N/A	N/A	N/A
SWESRC	N/A	Offline	N/A	N/A	N/A	N/A

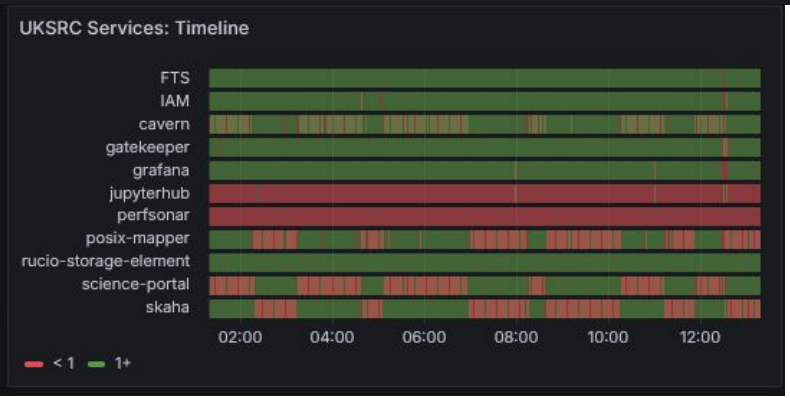

Global services

Adding alerting to the SRCNet monitoring

- Prometheus and Alertmanager alerts
- Grafana alerts
- Routed to Slack
- Further work
  - Define what to monitor
  - Create operator documentation
  - Define who receives the alerts, how they are routed and formatted

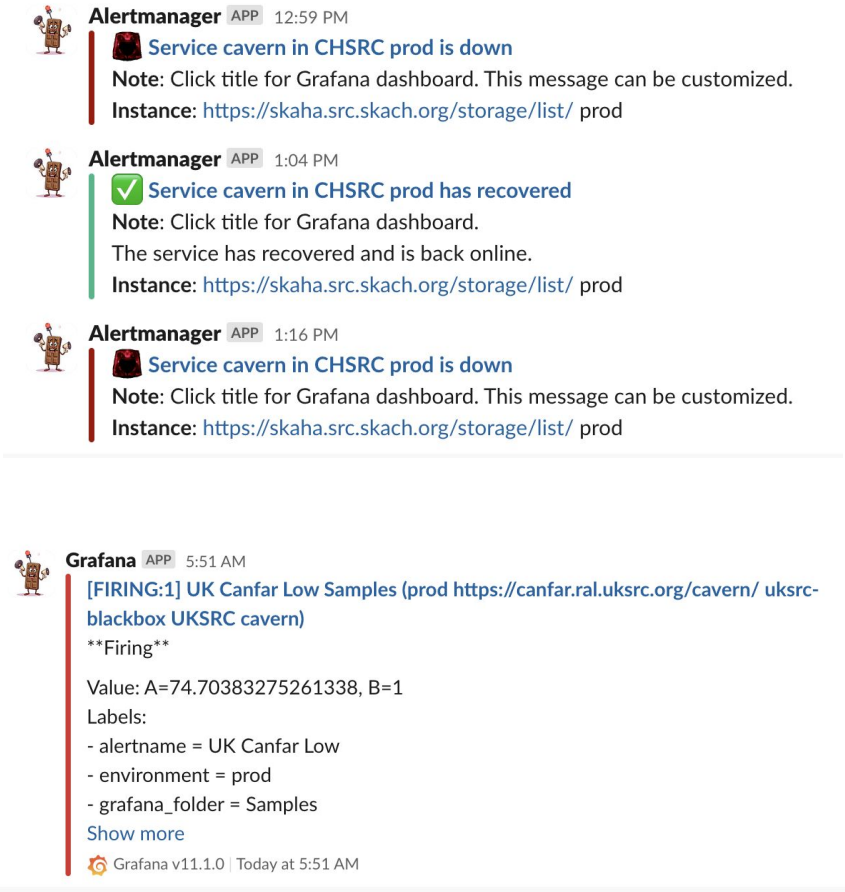

Ongoing developments

prepareData update (data staging), service migration to the new Gornergrat clusters

Data Staging | SRCNet

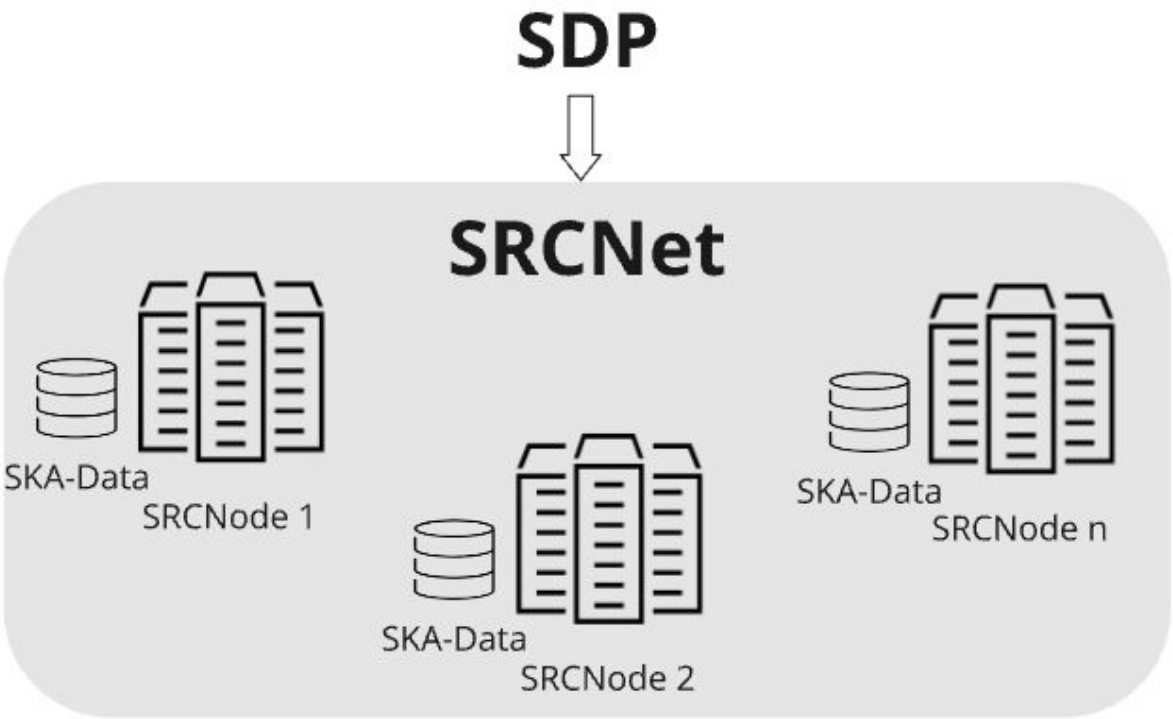

- Data inc. from SKA Science Data Processor (SDP)
- Ingestion into SRCNet landscape in several SRCNodes (using Rucio)
- SRCNodes providing compute, storage & services

Data Staging | Rucio

- Distributed data management system
  - Data ingestion, movement & copying (using FTS)
  - Data recovery & adaptive replication
  - SKA-data & science products
- Unified interface across heterogeneous storage and network infrastructures
- Data organized using Data Identifiers (DIDs) for files, datasets and containers (dataset-groups). All DIDs follow the schema: scope:name (e.g. data2025:sim-dataset-v0.1)
- Integrated authentication and authorization
- Local daemon & global service

https://rucio.cern.ch/

Data Staging | SRCNode

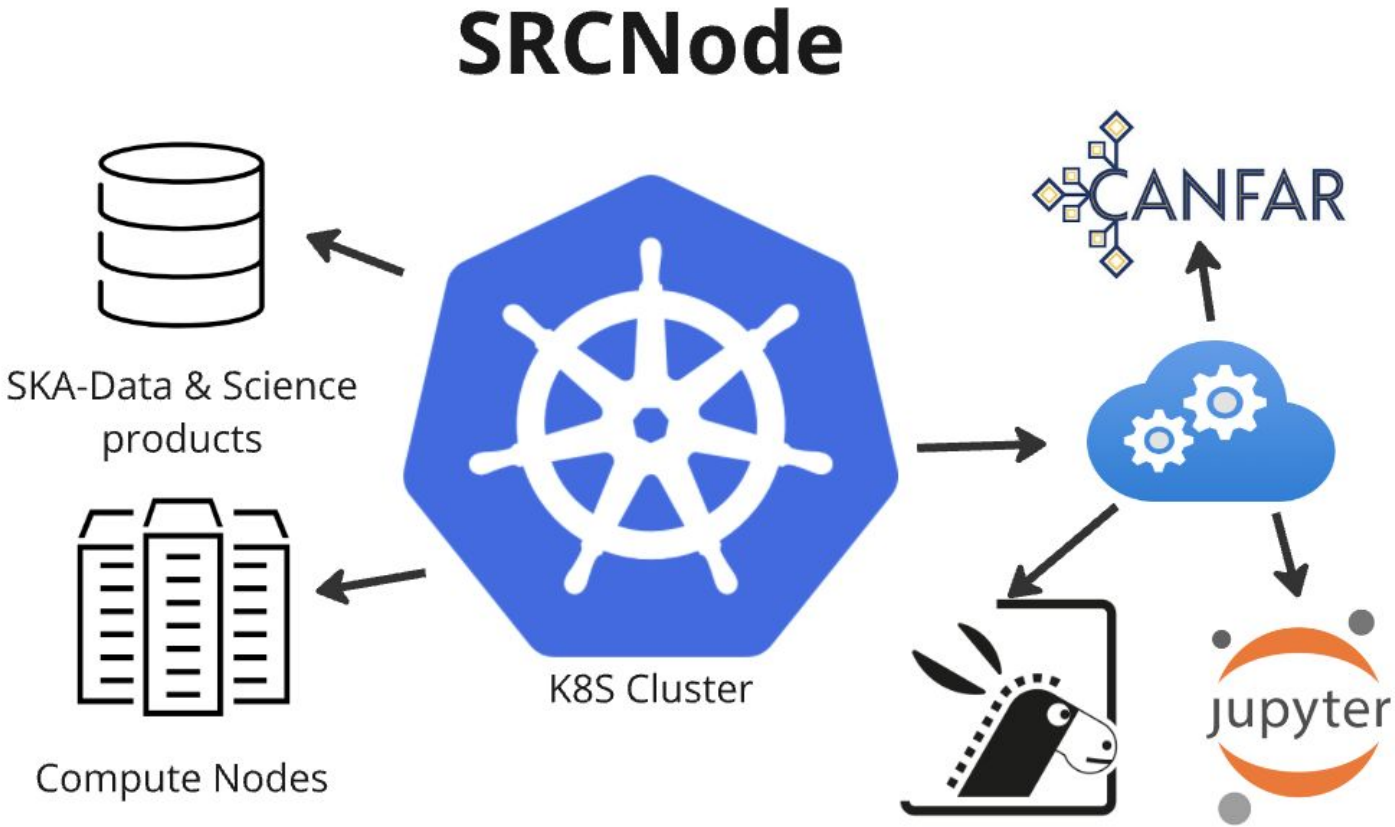

Data Staging | How it should work

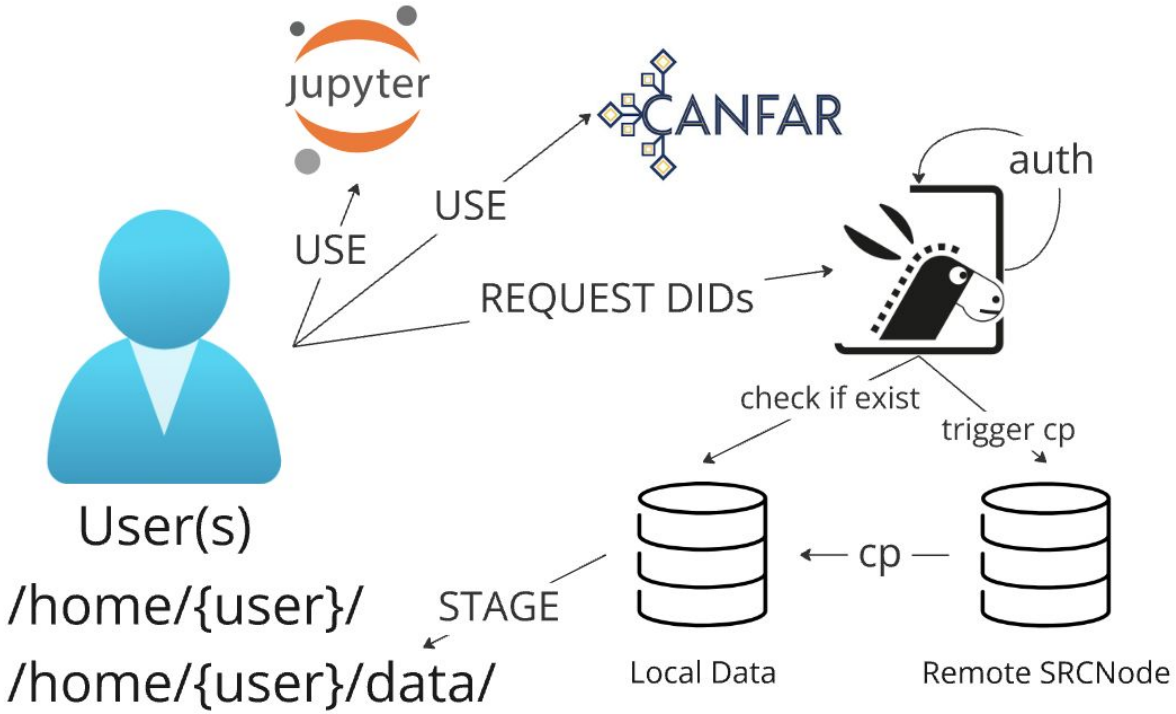

- Data easily available for science
- Request data staging (if permissions allow)
- Copy data from remote SRCNode to local if needed
- Access data in services (e.g. JupyterHub) you need

Data Staging | PrepareData

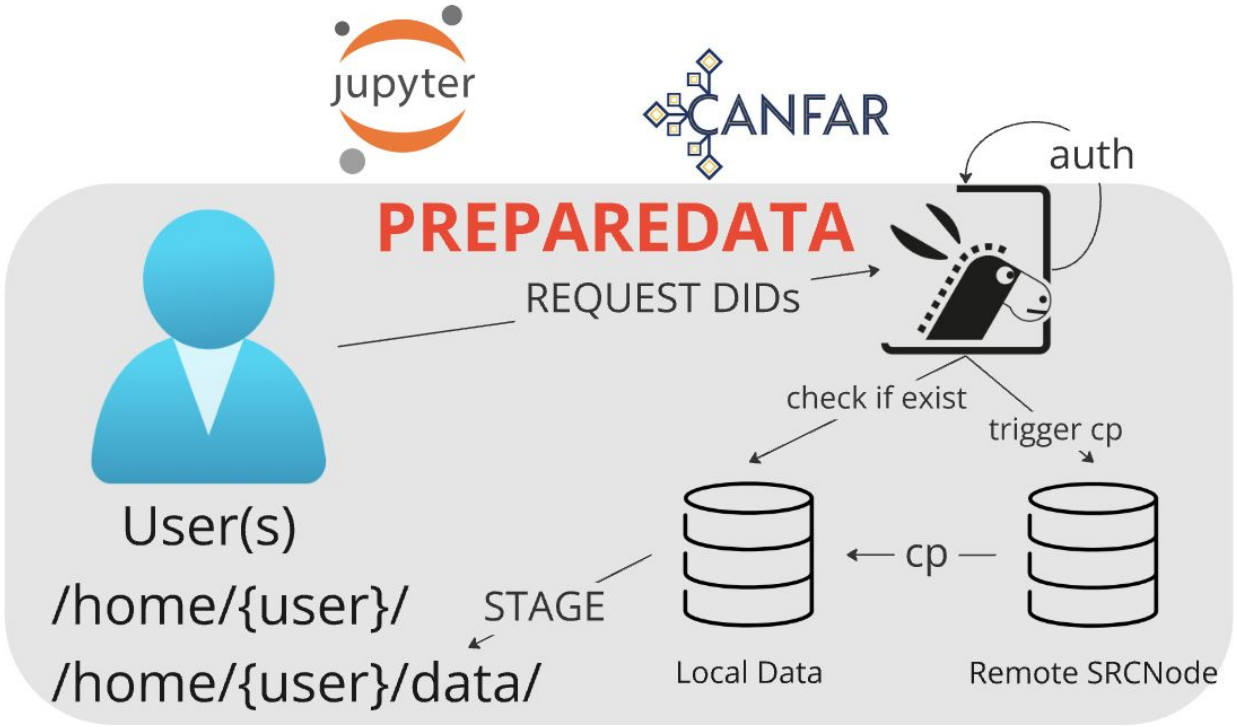

- Rucio doesn't have concept of proprietary data & capability to perform data staging
- Need for service which does that: PrepareData

Data Staging | PrepareData Challenges & WiP

- Different services have different home-mgmt
  - CANFAR/CAVERN using single PV&PVC (/home/{user})
  - JupyterHub using PV&PVC per user (with default uid & gid)
- Stage requested data into user-area without local data-duplication
  - Can’t expose full local Rucio data because of proprietary rights.
- Ideally let users of different services use same home-dir. This makes user-experience easier and we just have to stage data in single storage location.
- Implementing POC ATM: https://gitlab.com/ska-telescope/src/src-dm/ska-src-dm-local-data-preparer/
- Exploring approaches which don’t involve copying data into staged area.
  - Mount only requested part of DIDs (MountPropagation)
  - Symlink requested DIDs into staging area without exposing other DIDs

Thanks!

Any questions ?