The SKA Regional Center Network (SRCNet)

Raymond Oonk (SURF)

19 September, 2024

Outline

1) LOFAR
2) SKA
3) Intro to the SRCNet
4) Data in the SRCNet
5) Challenges & Opportunities

If time allows: Are we getting ready for the SRCNet data challenge ?

Insert SRCNet logo

Your guide today

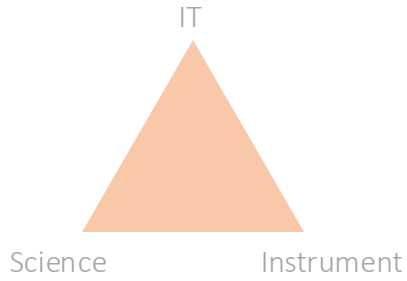

(J. B.) Raymond Oonk

- Sr. Advisor SURF
- Distributed Data Processing (2018) + PM Federation & Interoperability
- PhD Astronomy (Leiden University 2011)
- Postdoc (Leiden University 2014-2018, ASTRON 2011-2018)
- Research: Multi-wavelength

  Intergalactic medium, Galaxy clusters & Cooling flows

  Radio Recombination Lines (GHz down to 10 MHz)
- Contributed to the LOFAR upgrade (LMM, DUPPLO)
- Advanced LOFAR (Grid) processing for a.o. Surveys KSP
- Contributed SKA SDP resolution team (verifying SKA LOW architecture)

SURF

LOFAR ERIC

- LOFAR pathfinder for SKA
- Data distributed to Long Term Archive (SURF, Jülich & Poznan)
- LTA is a federated data lake
- Nearly 15 yrs of operation (12/6/2010)
- LTA largest astronomy archive in the world, but looking forward to SKA ☺

LOFAR ERIC

- LOFAR pathfinder for SKA
- Data distributed to Long Term Archive (SURF, Jülich & Poznan)
- **LTA is a federated data lake**
- **Nearly 15 yrs of operation (12/6/2010)**
- LTA largest astronomy archive in the world, but looking forward to SKA ☺

CAKE 2025 ?

Some of my personal lessons learned from a.o., LOFAR

- It takes time to understand the characteristics of the radio telescope
- Diversity in goals and setups impede delivery of science ready data products
- An observatory can only absorb limited input from the community

- Astronomers are ahead of the observatories in terms of processing algorithms
- Astronomers work in many small teams on many unique observations
- Most of the scientific output is generated by ‘short-lived’ astronomers

- A full compute federation is hard to evolve (science & IT develop fast)
- Astronomers are not IT system administrators and IT knowledge gap is widening
- Young IT specialist are trained in the cloud native and python ecosystem (only)

SURF

=> Ensure a mutually beneficial relation between observatory and astronomers that also gives researchers sufficient freedoms

Introduction to the SKA & SRCNet

Courtesy: (following slides)

- *J. Walder (UK-SRC)*
- *ASTRON*

+ personal observations

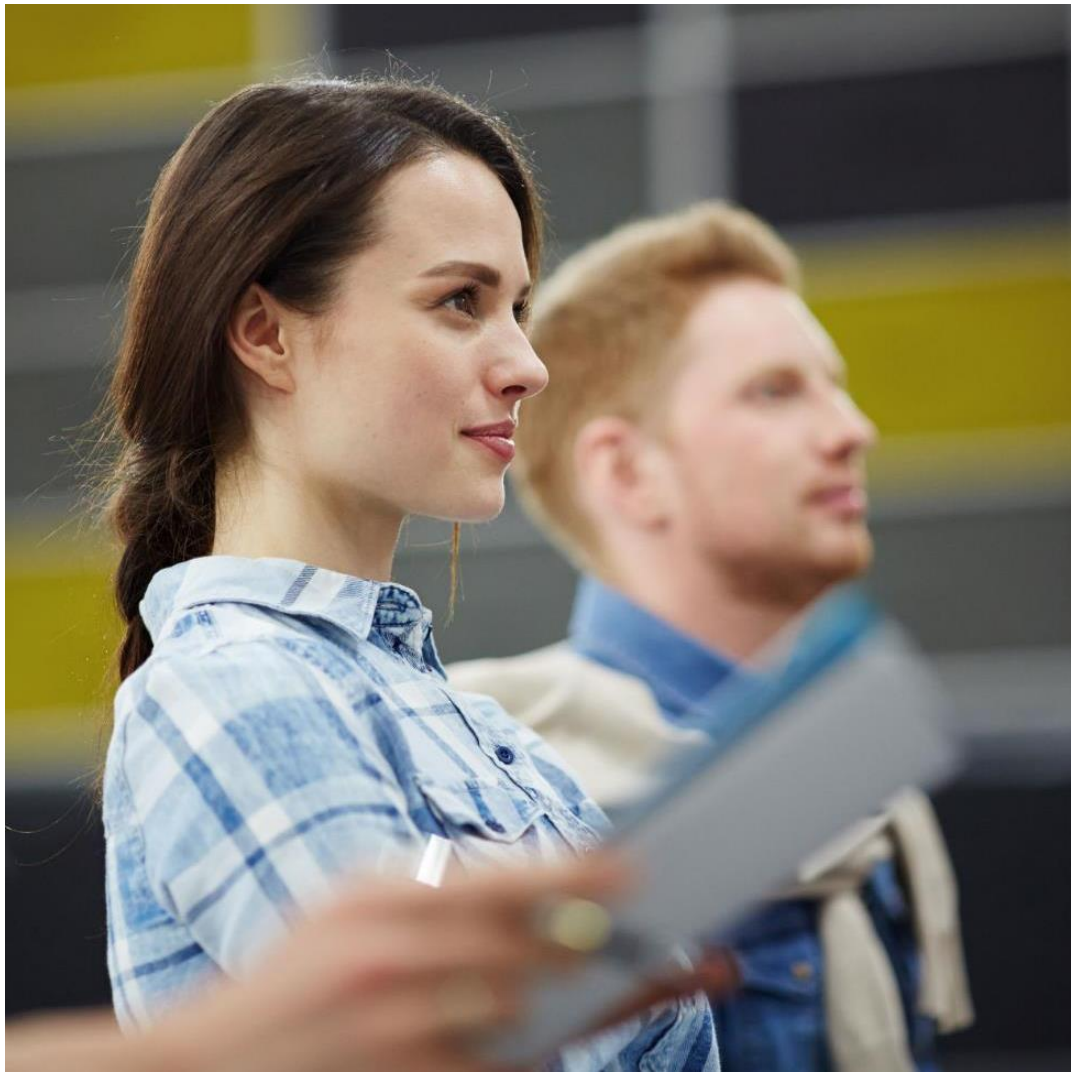

(courtesy: J. Walder)

SKA

The Square Kilometre Array (SKA) is a next-generation radio telescope aimed at being the largest and most powerful ever built

SKAO (Square Kilometre Array Observatory) is the **governing body** and **Intergovernmental Organisation** responsible for managing the global collaboration, scientific goals, and technical aspects of the SKA project

- Global project involving over 10 countries, with construction taking place in Australia (SKA-Low) and South Africa(SKA-Mid).
- > 50 year lifetime

Two Complementary Telescopes:

- SKA-Mid (350 MHz – 15.4 GHz): 197 dish antennas

  located in the Karoo Desert, South Africa.
  Maximum baseline of O(150) km.

- SKA-Low (50 MHz – 350 MHz): Over 131,000 dipole antennas

  grouped into 512 stations in Western Australia, covering a maximum distance of O(60+) km.

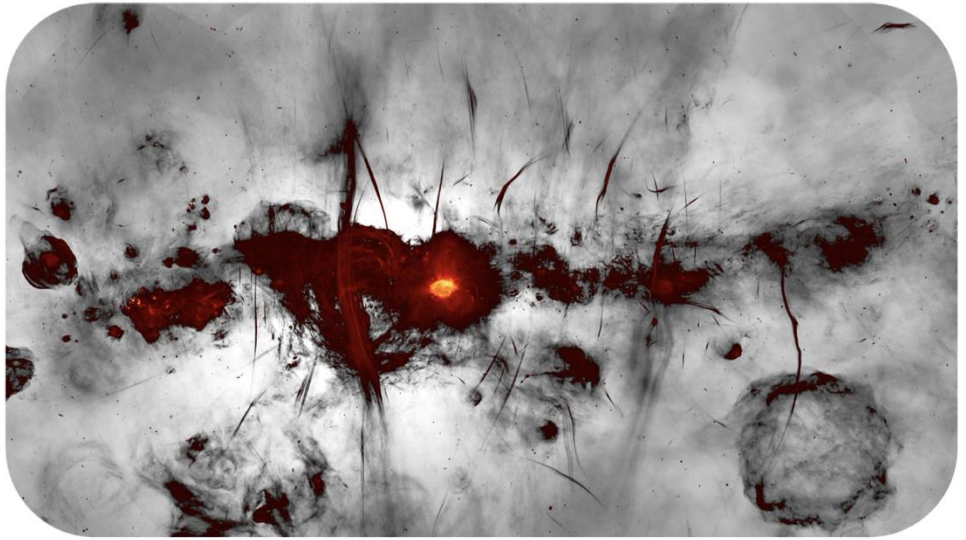

Credit: I. Heywood, SARAO

The SKA Regional Centres (SRCs):

- SRCs will receive data from the SKAO and act as the scientific archive for SKA data.
- Global Distribution: SRCs will be located across the world, creating a global network that provides the computational power, data storage, and tools needed for international collaboration.

(courtesy: J. Walder)

Science goals

Diverse and ambitious set of science goals

- Enabling discoveries across a wide range of fields in astronomy, cosmology, and fundamental physics
- Use of precursor and pathfinder projects to develop and prototype, with own scientific aims

Aspect	Pathfinder Projects	Precursor Projects
Definition	Existing radio telescopes testing technologies relevant to SKA	Prototypes located at SKA sites in South Africa and Australia
Goal	Develop scientific techniques, instrumentation, and methodologies	Conduct early science and engineering directly leading to full SKA
Locations	Worldwide (Europe, Australia, South Africa, etc.)	Located at the SKA sites in South Africa (SKA-Mid) and Australia (SKA-Low)
Notable Projects	LOFAR, ASKAP, MWA, MeerKAT, VLBI	MeerKAT, ASKAP, MWA
Contribution to SKA	Provide foundational data, technology testing, and science exploration	Test SKA-specific designs, contribute to early science, integrate with SKA

Science Goal	Focus
Cosmic Dawn and Epoch of Reionization	First stars, galaxies, and reionization
Galaxy Evolution and Dark Energy	Large-scale structure, galaxy formation, and dark energy
Cosmic Magnetism	Mapping magnetic fields across the universe
Pulsars and Extreme Gravity	Testing general relativity and extreme gravity with pulsars
Fast Radio Bursts (FRBs)	Understanding the origin and mechanisms of FRBs
Search for Extraterrestrial Life (SETI)	Detecting signals or biosignatures from other civilizations
Cradle of Life	Studying planetary systems and conditions for life
Probing the Dark Ages	Observing the universe before the first stars
Transient Events	Real-time monitoring of cosmic explosions and supernovae
Gravitational Waves and Black Holes	Indirect detection of gravitational waves via pulsar timing arrays

Don’t need to be radio astronomy expert to exploit science data outputs

- Multi-wavelength data increasingly needed for transformational science

Property data access and embargo periods for projects

Locations and participating nations

(courtesy: J. Walder)

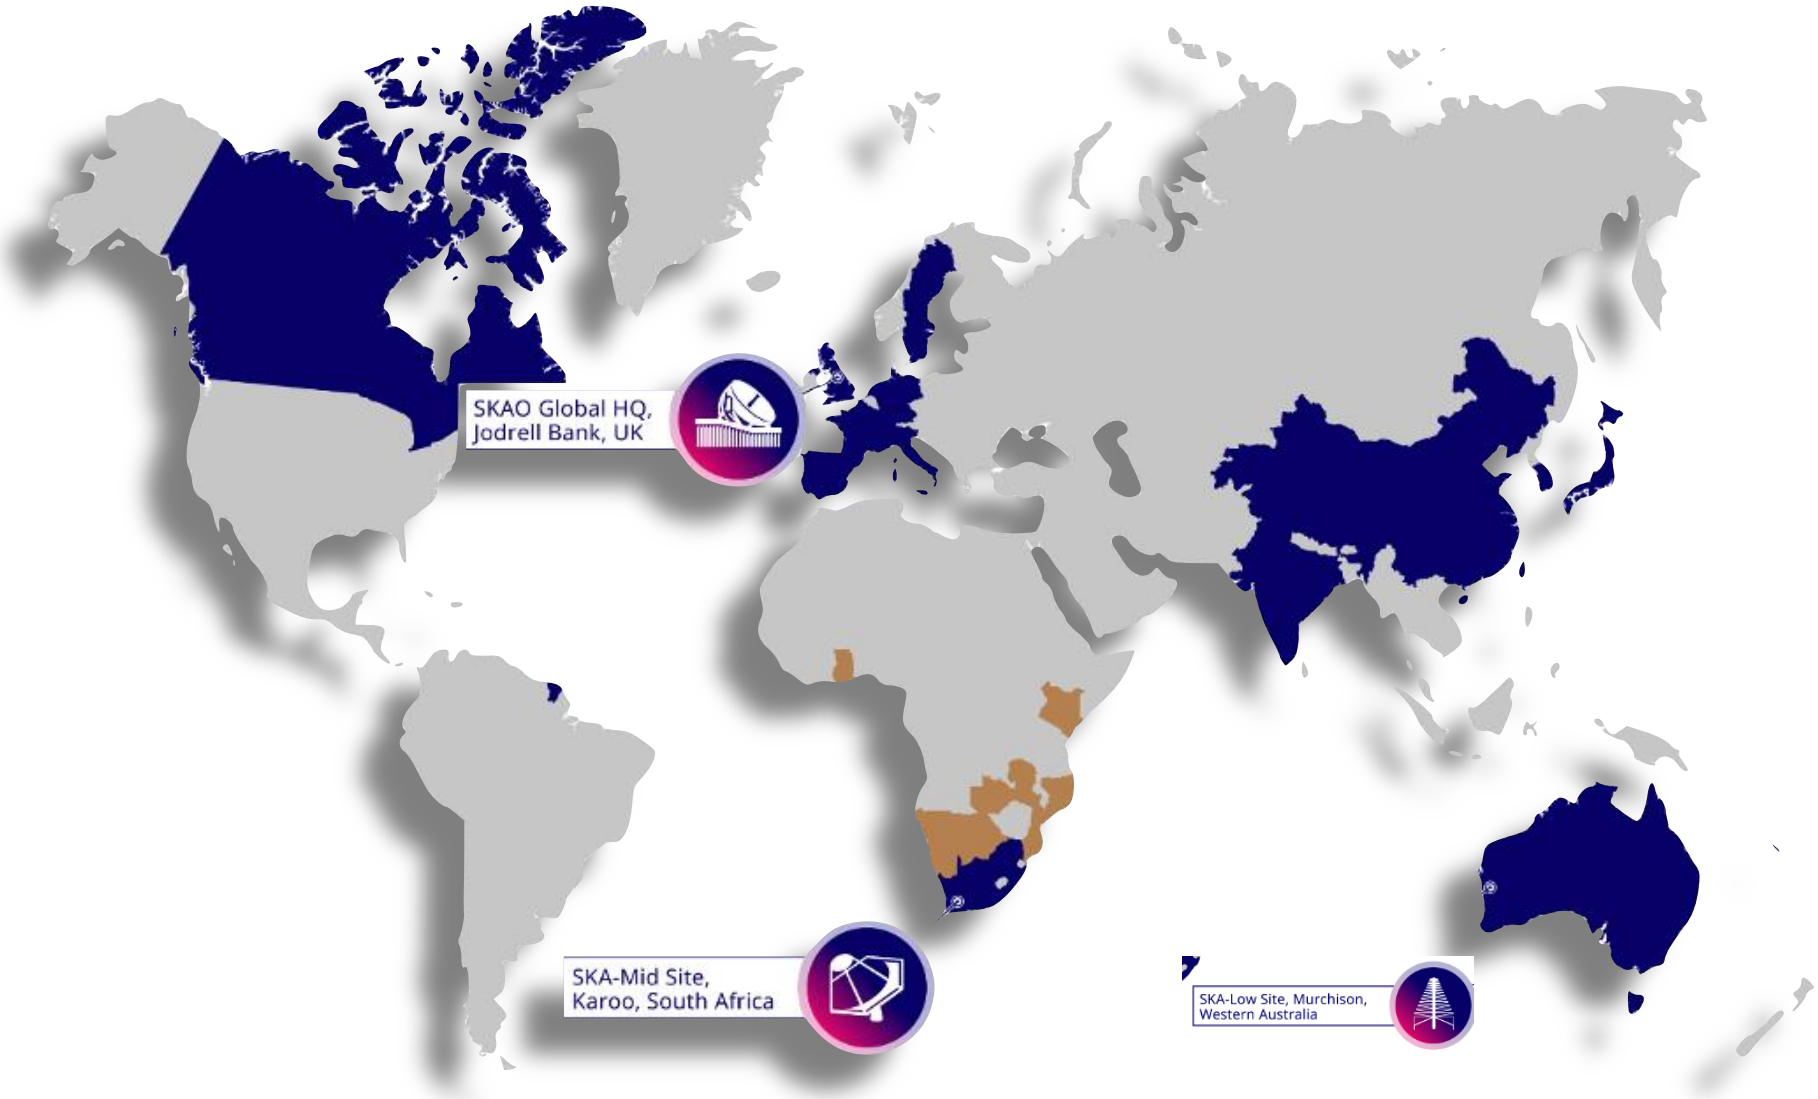

(courtesy: J. Walder)

(courtesy: J. Walder)

SURF

R. Oonk

(courtesy: J. Walder)

First Light! (1 station)

(courtesy: J. Walder)

- This is the first image and video from observations using **one complete SKA-Low station**, known as S-8, produced only **18 months after the start of construction** activities on site, and five months after the first antenna was installed.

- The completion of a station means not only assembling and installing the **256 antennas**, but also integrating them with all the computing systems behind them.

- The video shows a **24-hour observation**, with the Milky Way rising and passing overhead during the night time hours.

  Some other bright radio sources are marked, including the galaxies Centaurus A and M87, and the Sun is also visible during the day.

  => Very reminiscent of early LOFAR (CS1)

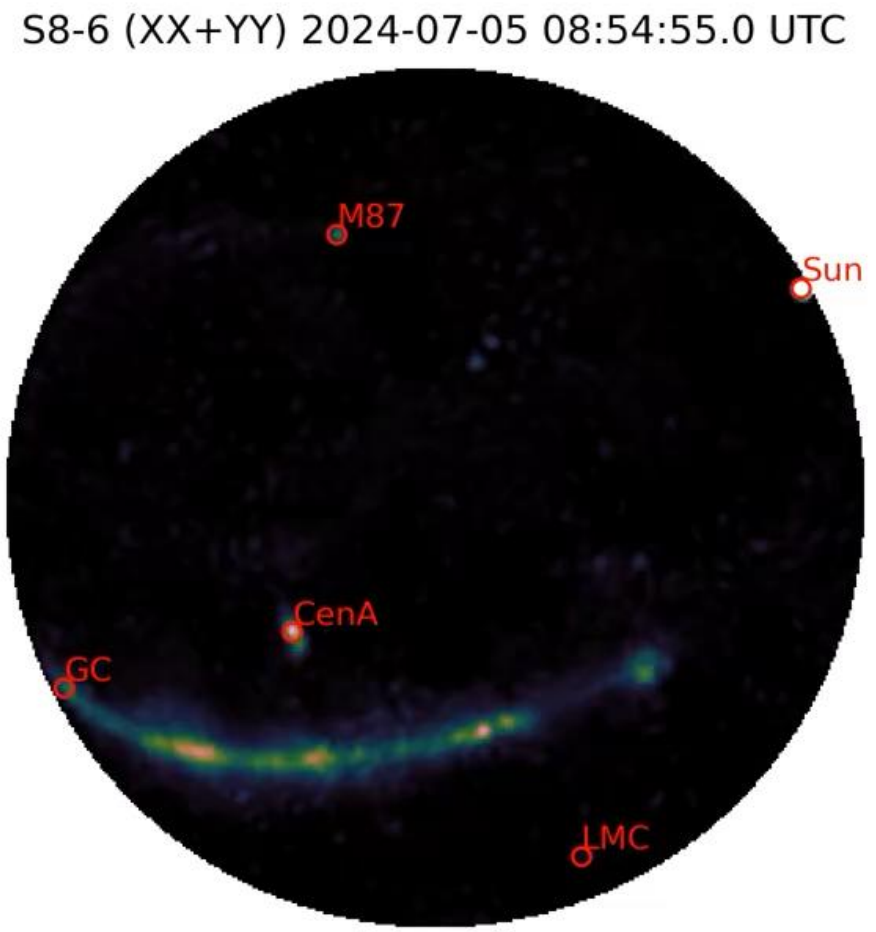

https://.youtube.com/watch?v=zakuQ1-QrGg

First Fringes! (2 stations)

- This shows the phase structure of the correlated signal between **two SKA-Low stations of 256 antennas** each. The phase structure varies with time and frequency and is caused by the projected station separation toward the radio galaxy **Centaurus A**

https://www.skao.int/en/news/570/ska-low-achieves-first-fringes

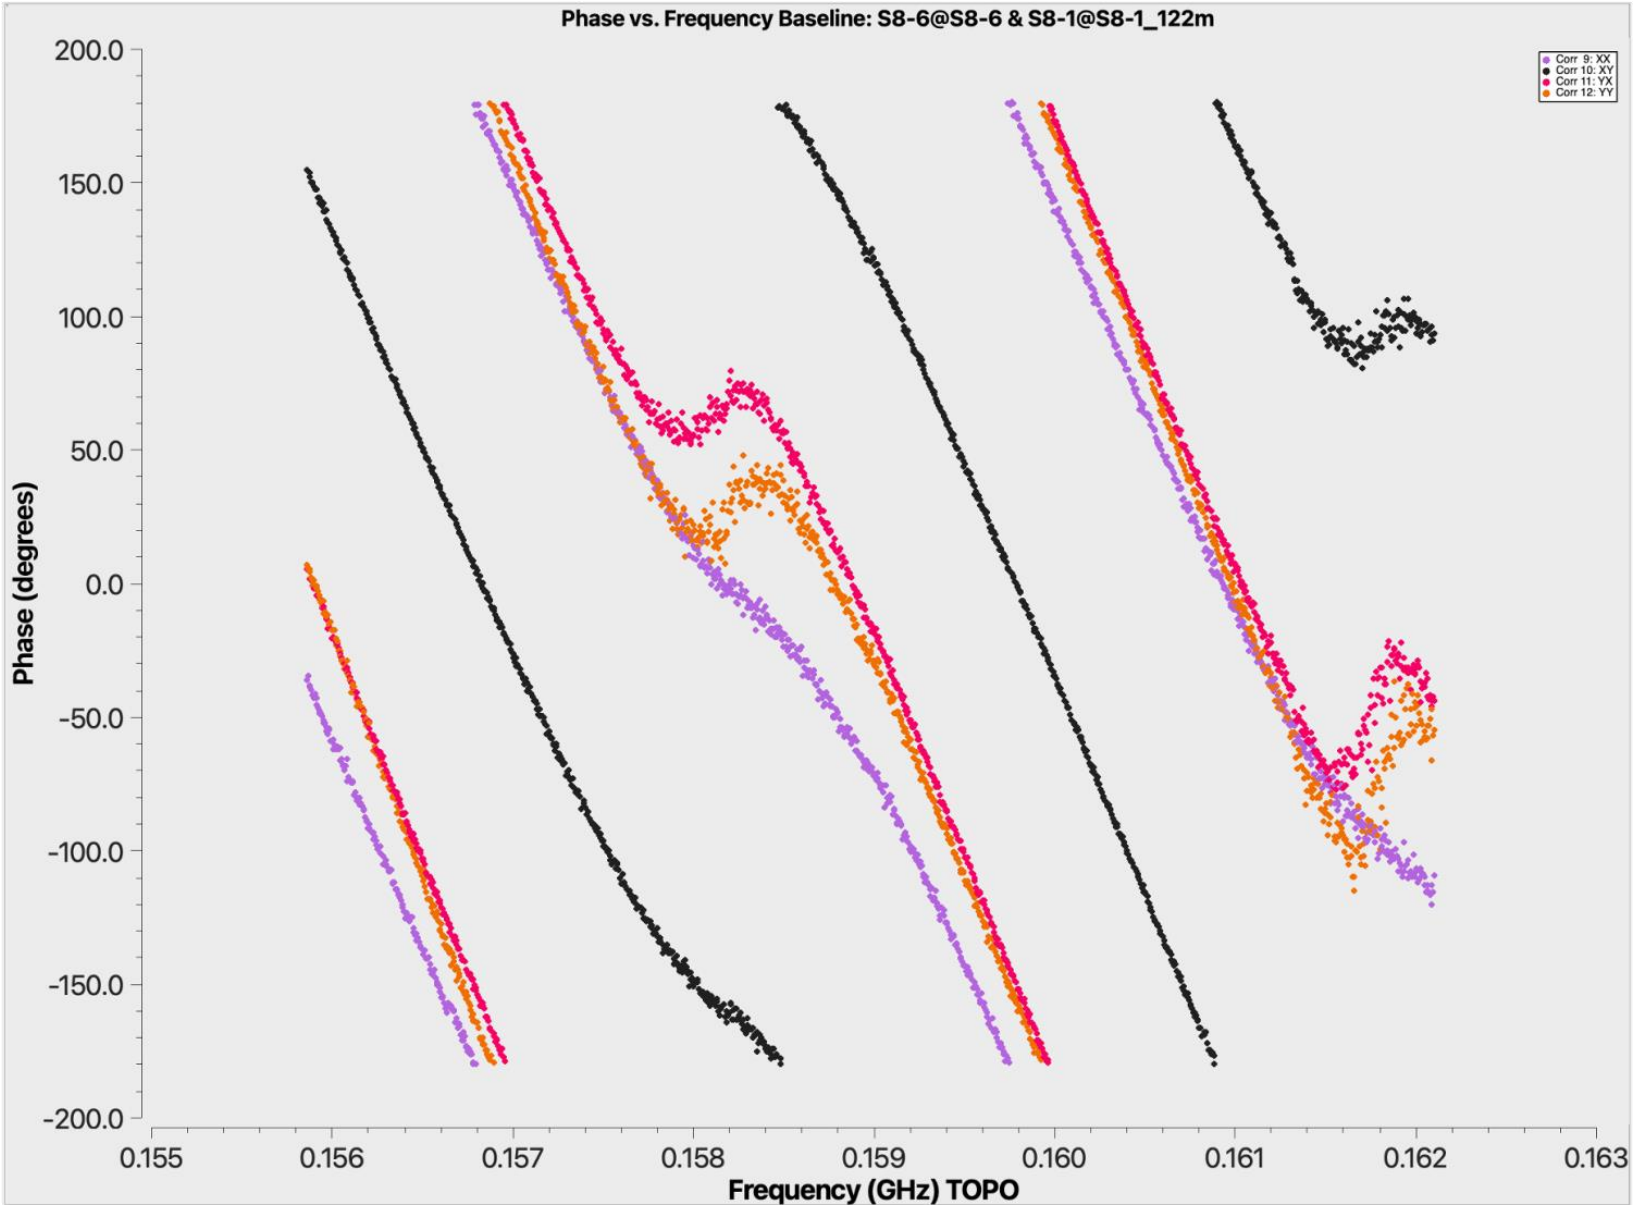

(courtesy: J. Walder)

Construction Timeline

- Construction Strategy
- Target: build the SKA Baseline Design
  (197 Mid dishes; 512 Low stations: AA4)
  - Not all funding yet secured for full AA4 components, following Staged Delivery Plan (AA*)
- Develop the earliest possible working demonstration of the architecture and supply chain (AA0.5).
- Then maintain a continuously working and expanding facility that demonstrates the full performance capabilities of the SKA Design.

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 Apr	2025 Nov
AA2 (64 dishes, 64 stations)	2027 Mar	2026 Oct
AA* (144 dishes, 307 stations)	2027 Dec	2028 Jan
Operations Readiness Review	2028 Apr	2028 Apr
AA4 (197 dishes, 512 stations)	TBD	TBD

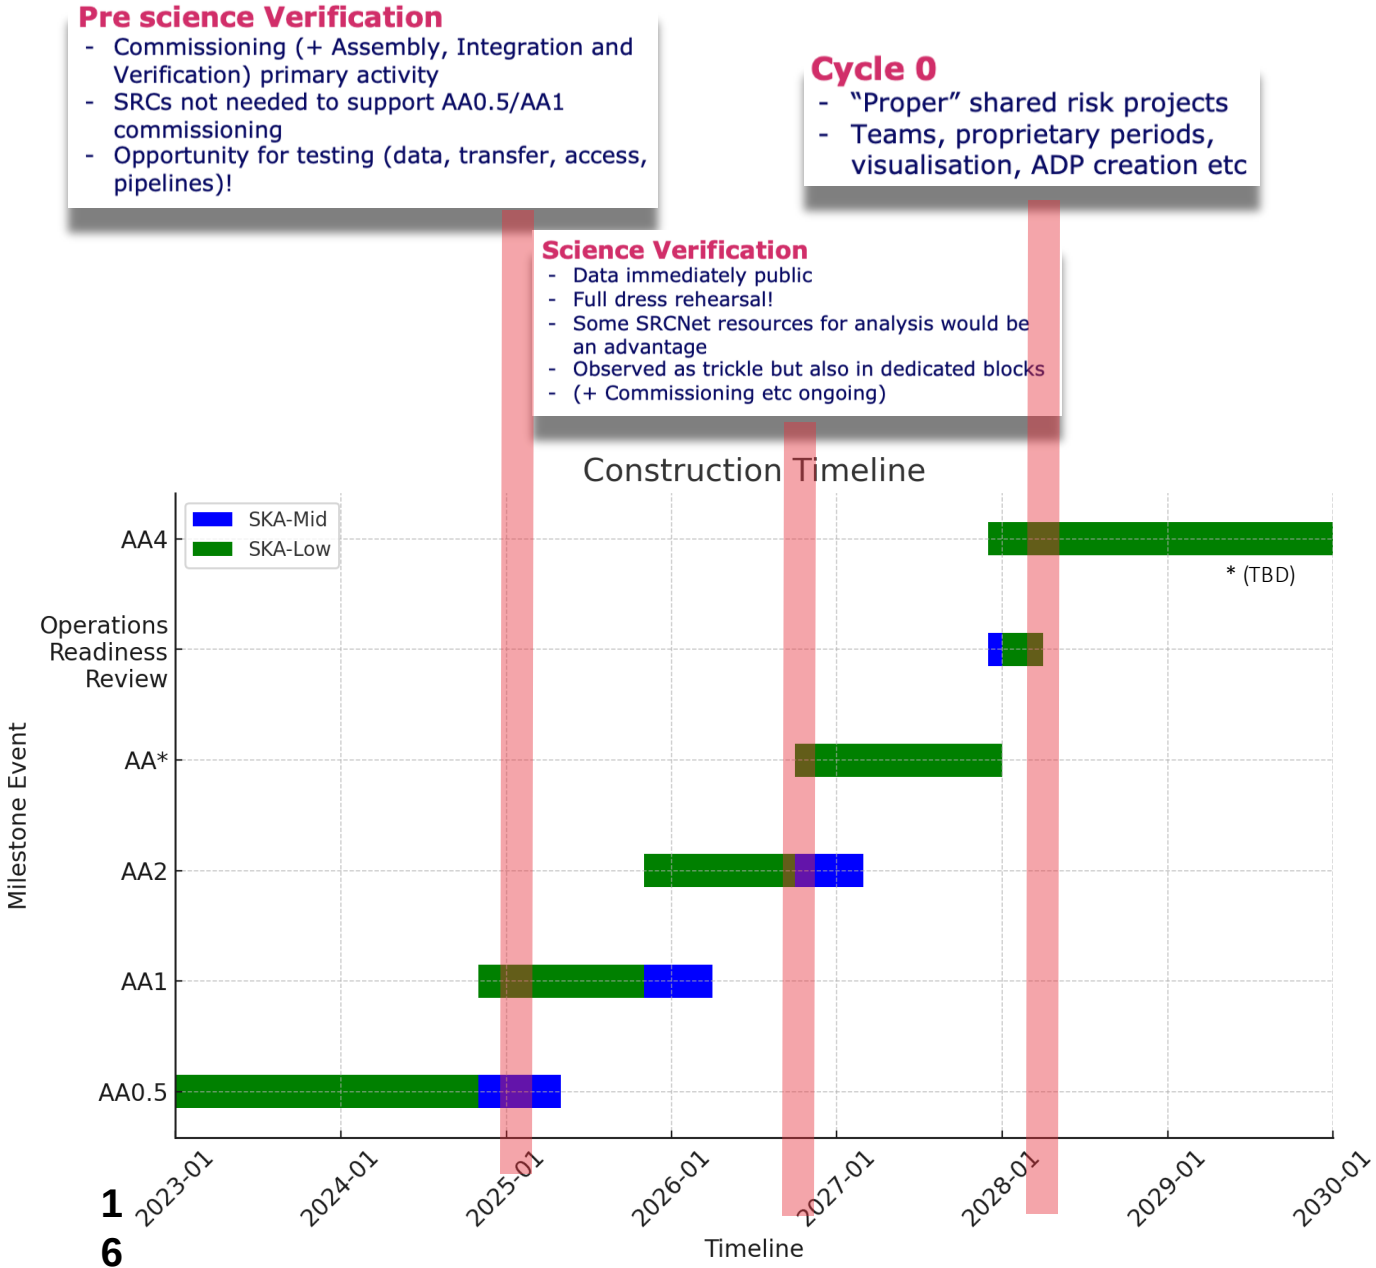

(courtesy: J. Walder)

Data Flows

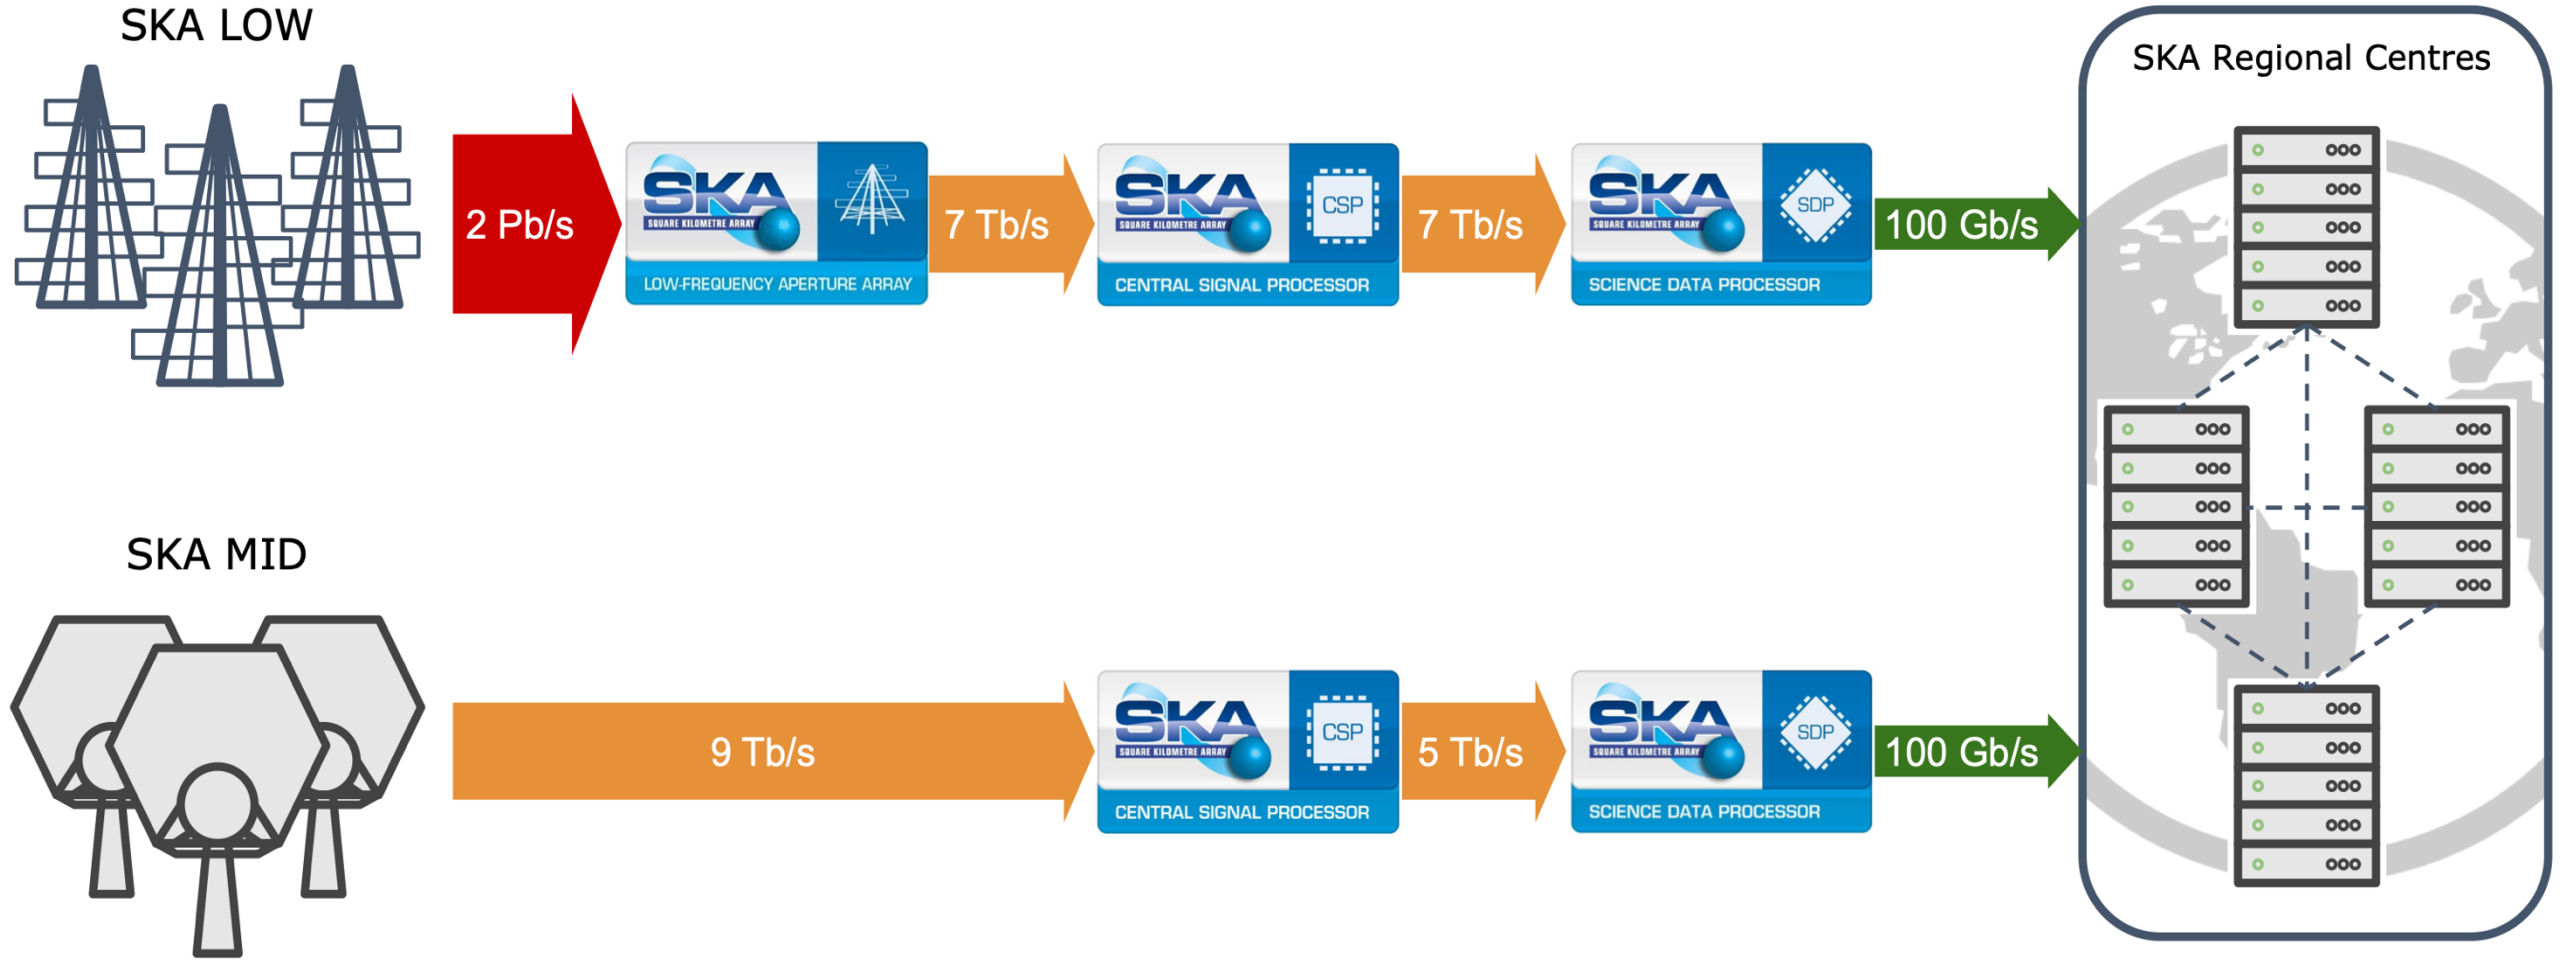

* Data rates approximate

SKA expected data rates*

(courtesy: J. Walder)

***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 (i.e. not data used internally for commissioning etc.)

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^ 20 Gbps	2 PB/week^ 27 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^ 66 Gbps	9 PB/week^ 119 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	173 PB/year 44 Gbps	280 PB/year 72 Gbps
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 55 Gbps	400 PB/year 100 Gbps

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

(courtesy: J. Walder)

SRCNet

* The need for a network of SKA Regional Centres (SRCs) formed around ~ 2018:
  * Distributed compute, storage and expertise to store, process and disseminate data to the communities

Naive mapping between LHC and SKA entities

Don't take too literally

CERN + Experiments	~ SKAO + SKA-MID + SKA-LOW
WLCG	~ SRCNet
GridPP/ …	~ UK-SRC / …

* SRCNet is a global, distributed infrastructure designed to:
  * **Store and manage SKA data:**
    * Centralised data management is impractical given the size of the SKA data output, so SRCNet provides regional centres to manage and store data close to its users.
  * **Enable scientific collaboration:** Astronomers worldwide can access, process, and analyze SKA data in a collaborative environment, regardless of their location.
  *. **Provide computational resources**: Given the vast amount of data, powerful high-performance computing (HPC)resources are needed. SRCNet provides this by distributing the workload across a global network of SRCs.
  * **Support scientific tools and pipelines**: SRCNet offers standardised tools and software pipelines for researchers to work on SKA data, including supporting advanced data analysis, visualisation, and machine learning techniques.

SRCNet Prototyping/Development Teams

(courtesy: ASTRON)

Agile teams (SAFE)

- Cyan - Data Management Technologies
- Tangerine - Science Platform Components
- Orange – Visualisation
- Purple - AAI Prototype
- Coral - European Proto-SRC, Deployment
- Blue-Lavender - Asia-Pacific Proto-SRC, Deployment (split)
- Olive – HPC, Cloud Services, Providers (dormant)
- Evergreen – Hardware & Network (not yet in place)
- Navy – Operations (not yet in place)

Non-Agile groups

- SRCNet Program Team
- Infrastructure Services Management Group
- SRCNet Operations Group
- Architecture Forum
- NREN Forum

- *Architecture doc*
- *Implementation Roadmap doc*
- *Pledges (1st round v0.1 done)*

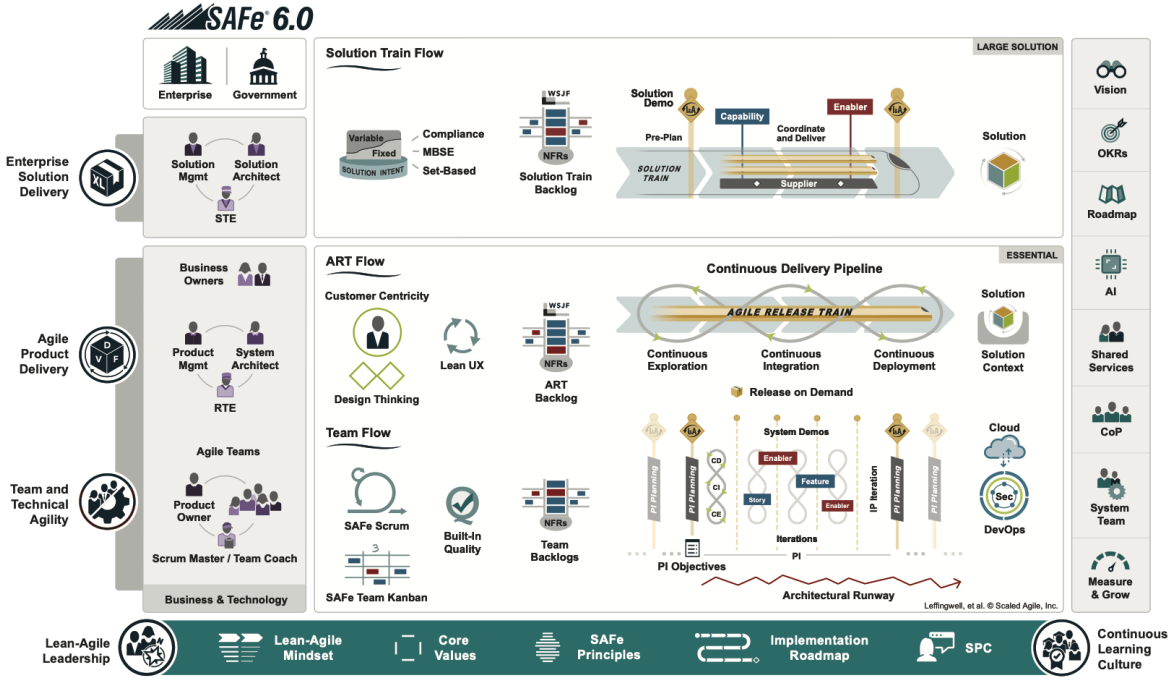

Region Center Capabilities

(courtesy: J. Walder)

Science Enabling Applications
Analysis Tools, Notebooks, Workflows execution Machine Learning, etc

Distributed Data Processing
Computing capabilities provided by the SRCNet to allow data processing

Data Discovery
Discovery of SKA data from the SRCNet, local or remote, transparently to the user

Visualization
Advanced visualizers for SKA data and data from other observatories

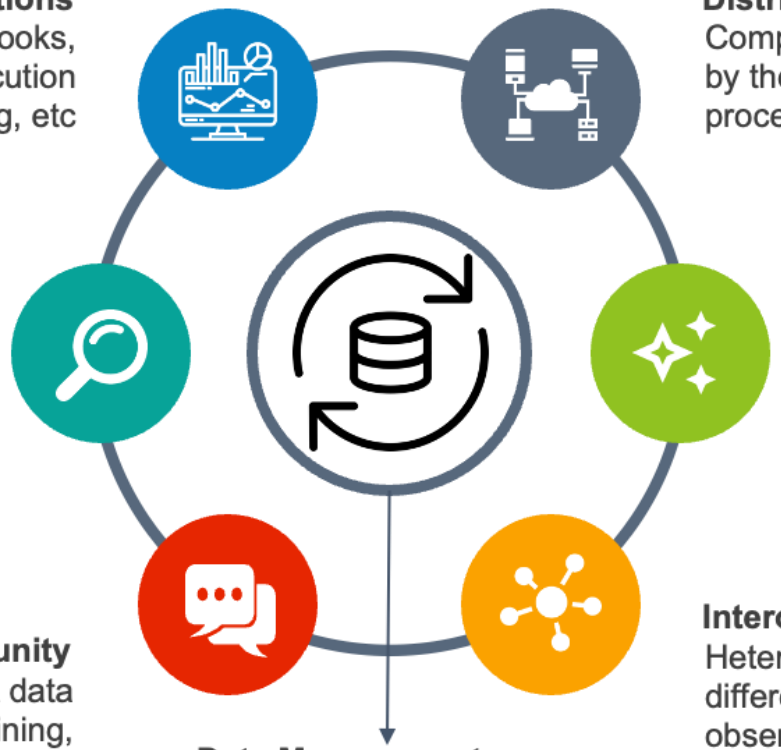

Support to Science Community
Support community on SKA data use, SRC services use, Training, Project Impact Dissemination

Interoperability
Heterogeneous SKA data from different SRCs and other observatories

Data Management
Dissemination of Data to SRCs and Distributed Data Storage

SRCNet Timeline

(courtesy: J. Walder)

* SRCNet timeline:
  * SRCNet v0.1 (2025) :
    First major milestone in SRCNet project

- Built to show the architecture + 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

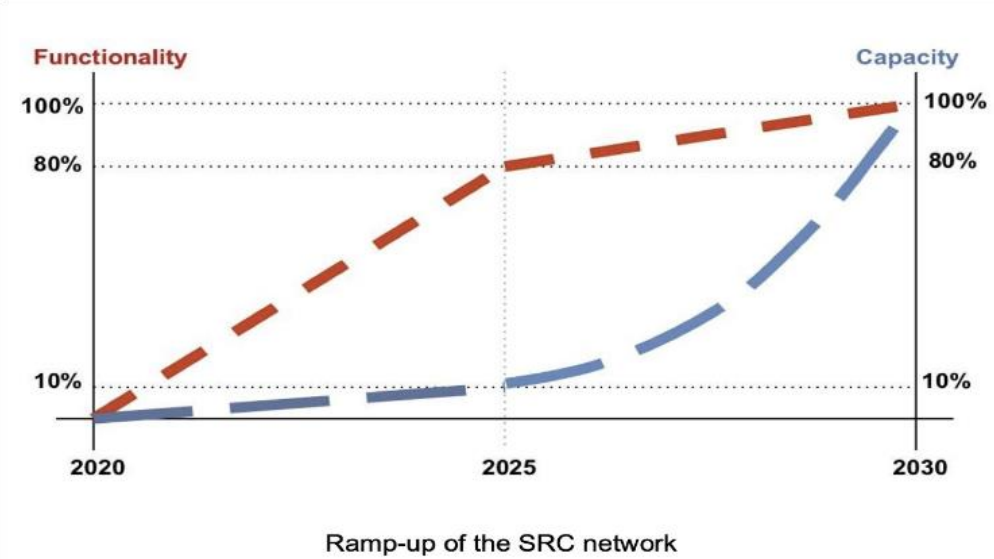

Ramp-up of the SRC network

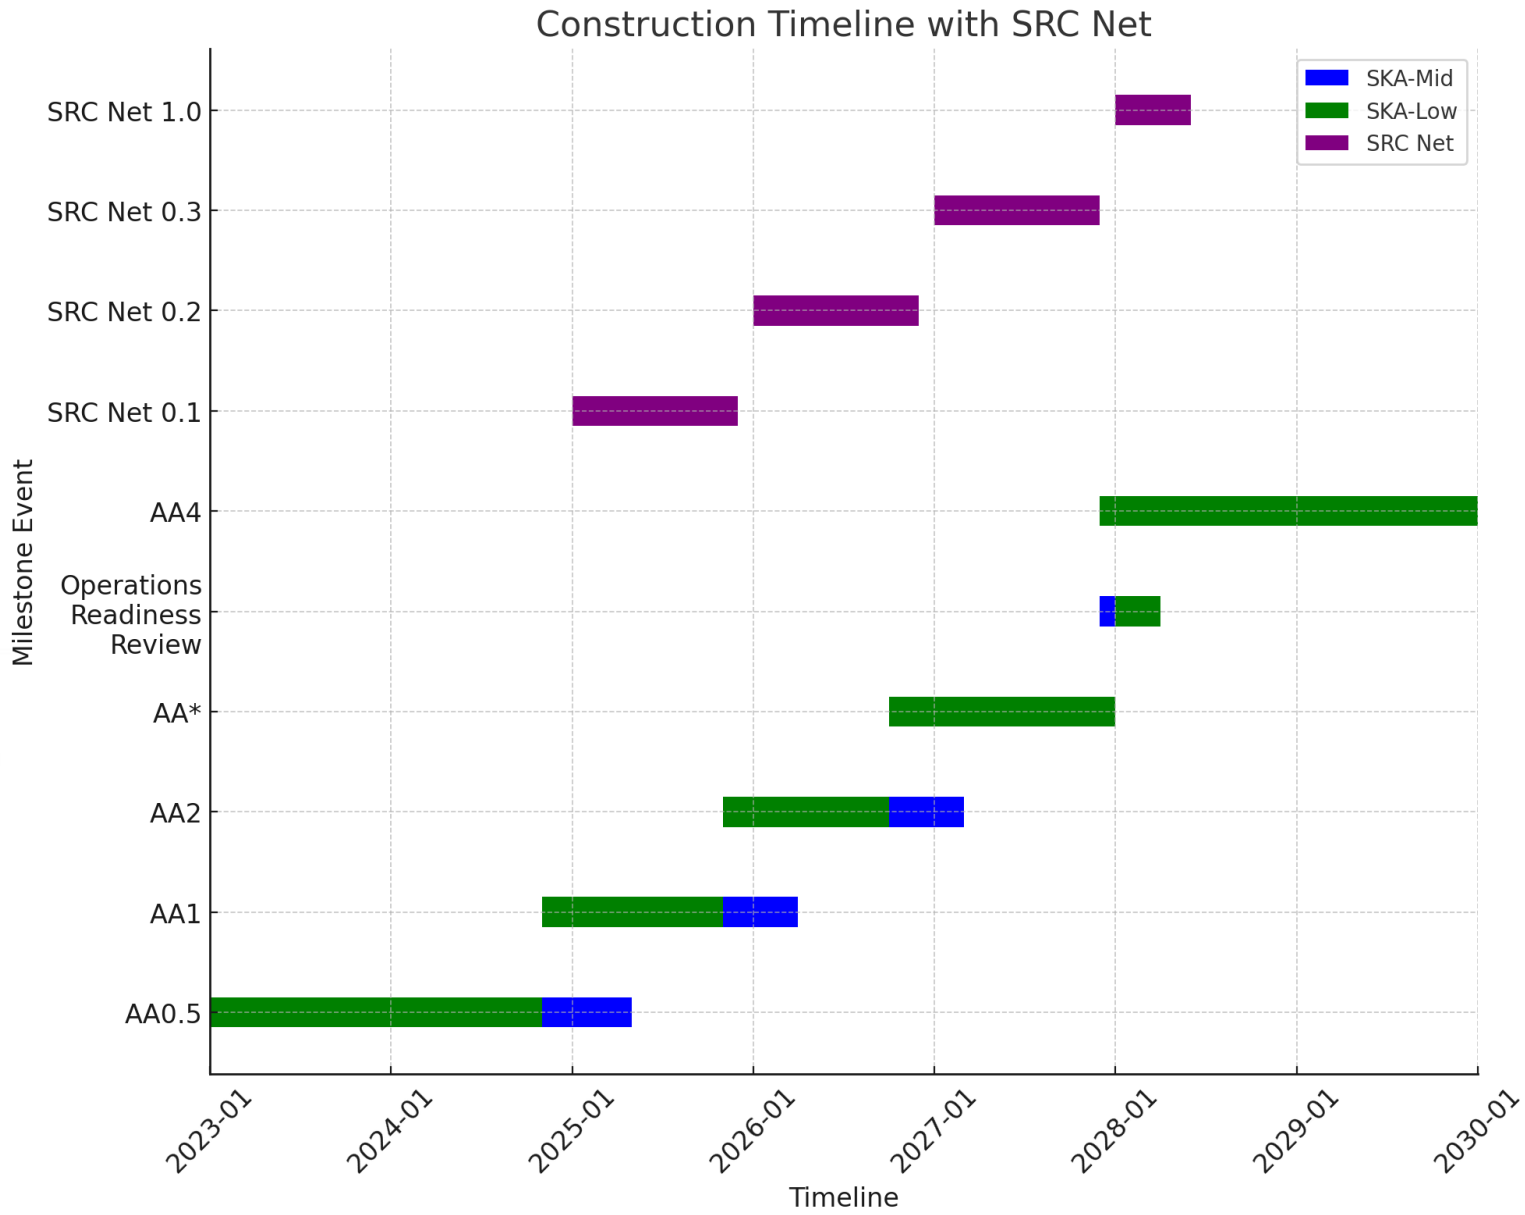

(courtesy: J. Walder)

SRCNet Timeline

- SRCNet timeline:

- SRCNet v0.2 (2026: AA2)
  - Increased Federation
  - Science verification starts
  - Data ingestion from telescopes
  - ca. **182 PB/yr** (20,27 Gbps)

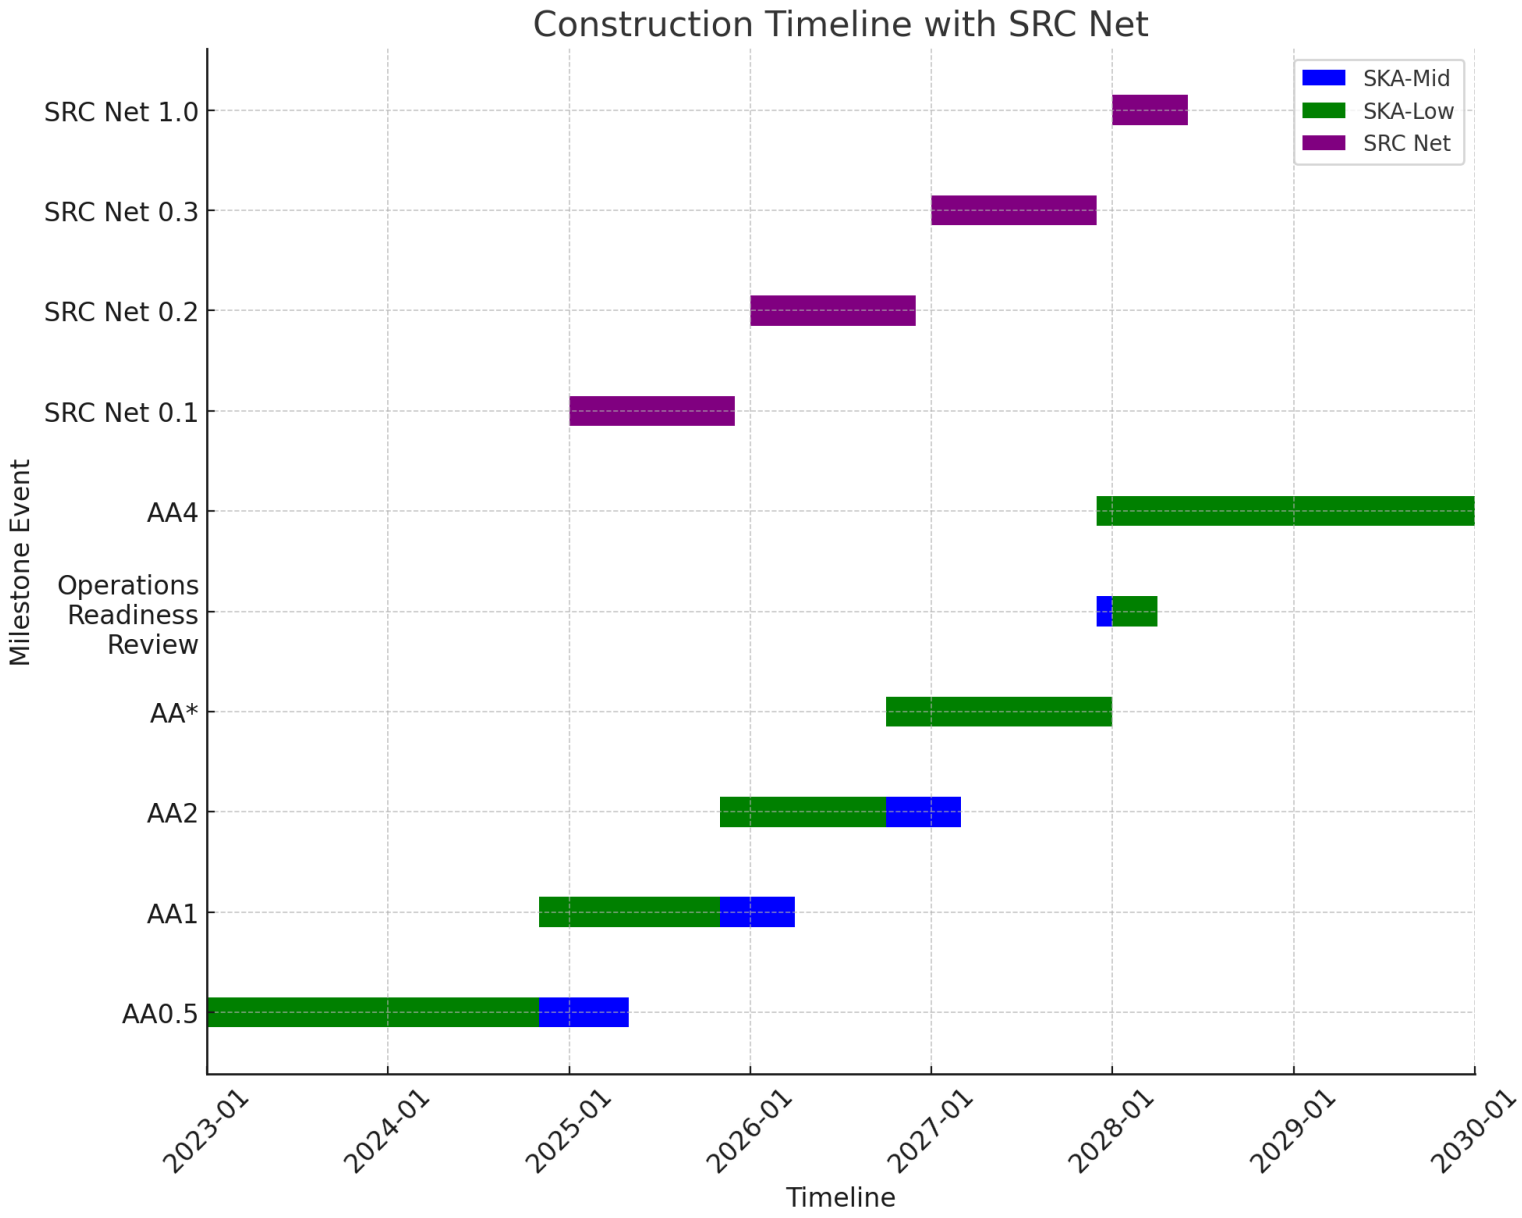

(courtesy: J. Walder)

SRCNet Timeline

- SRCNet timeline:

- SRCNet v0.3 (2027: AA*)
  - Increased science verification
  - Science user testing
  - ca. **728 PB/yr** (66,119 Gbps)

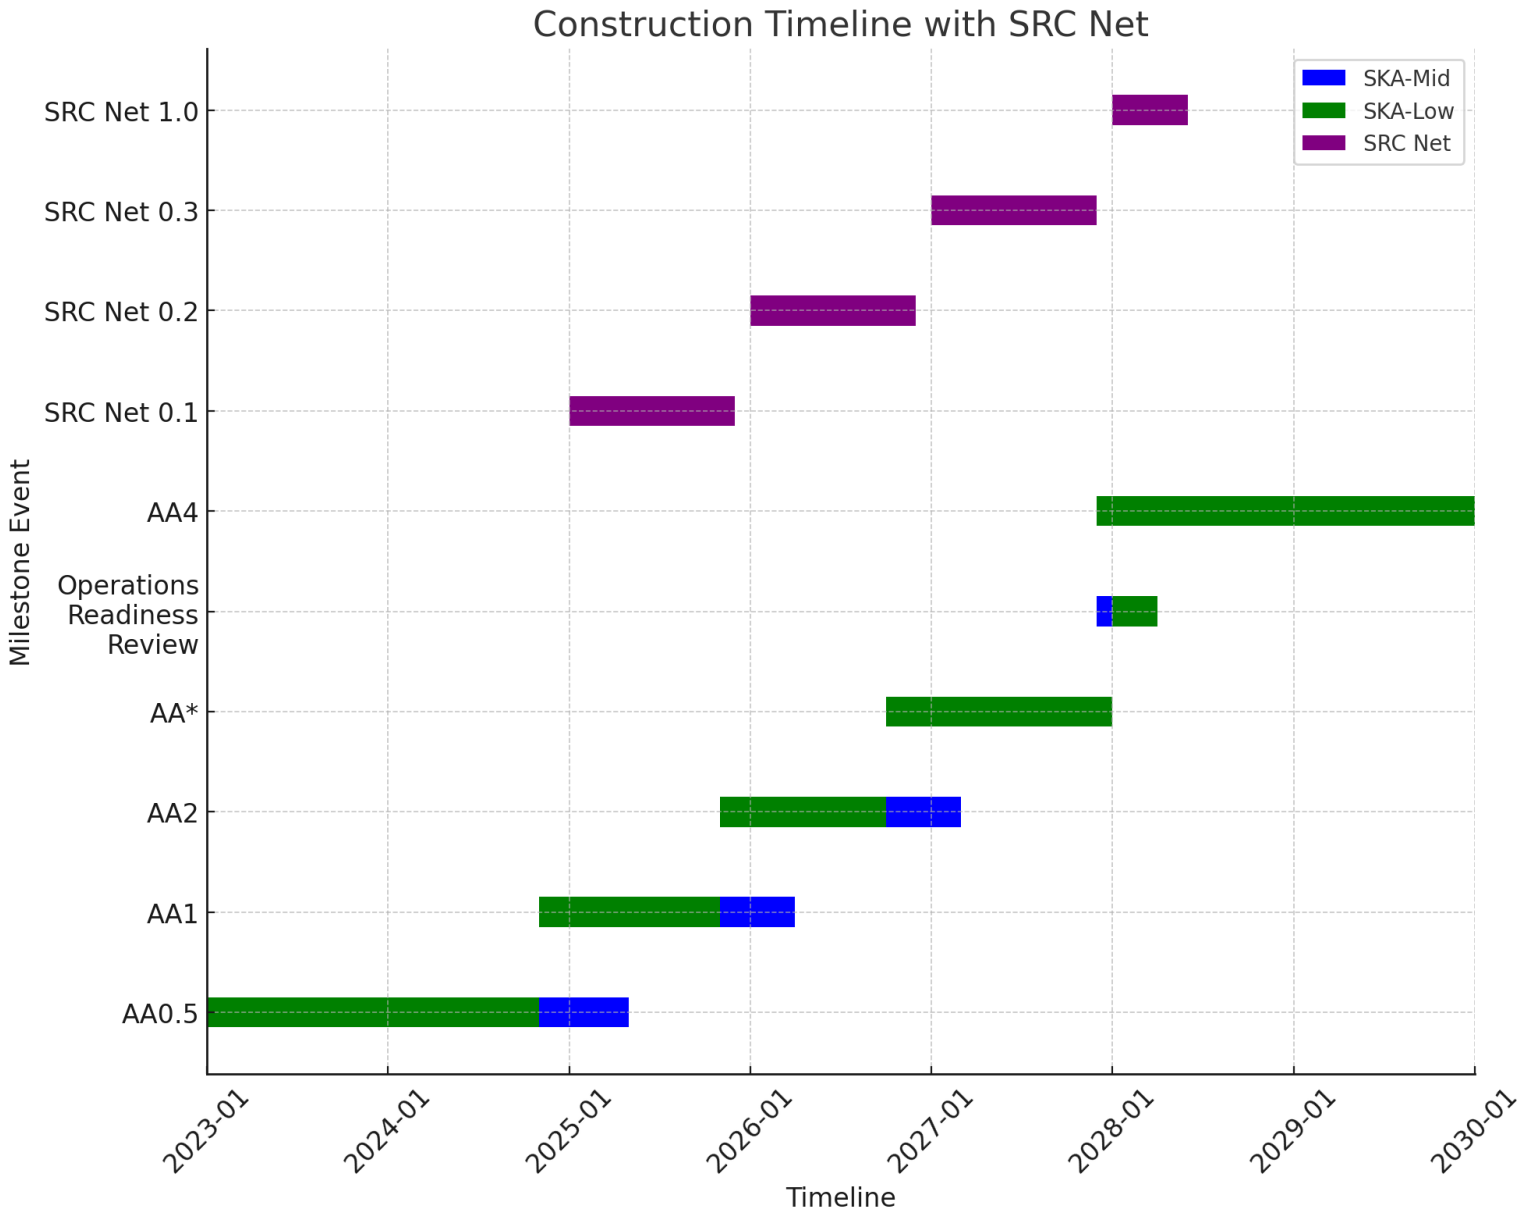

(courtesy: J. Walder)

SRCNet Timeline

- SRCNet timeline:

- SRCNet v1.0 (~2028+ : AA4)
  - Full operations
  - ca. **616 PB/yr** (55,100 Gbps)

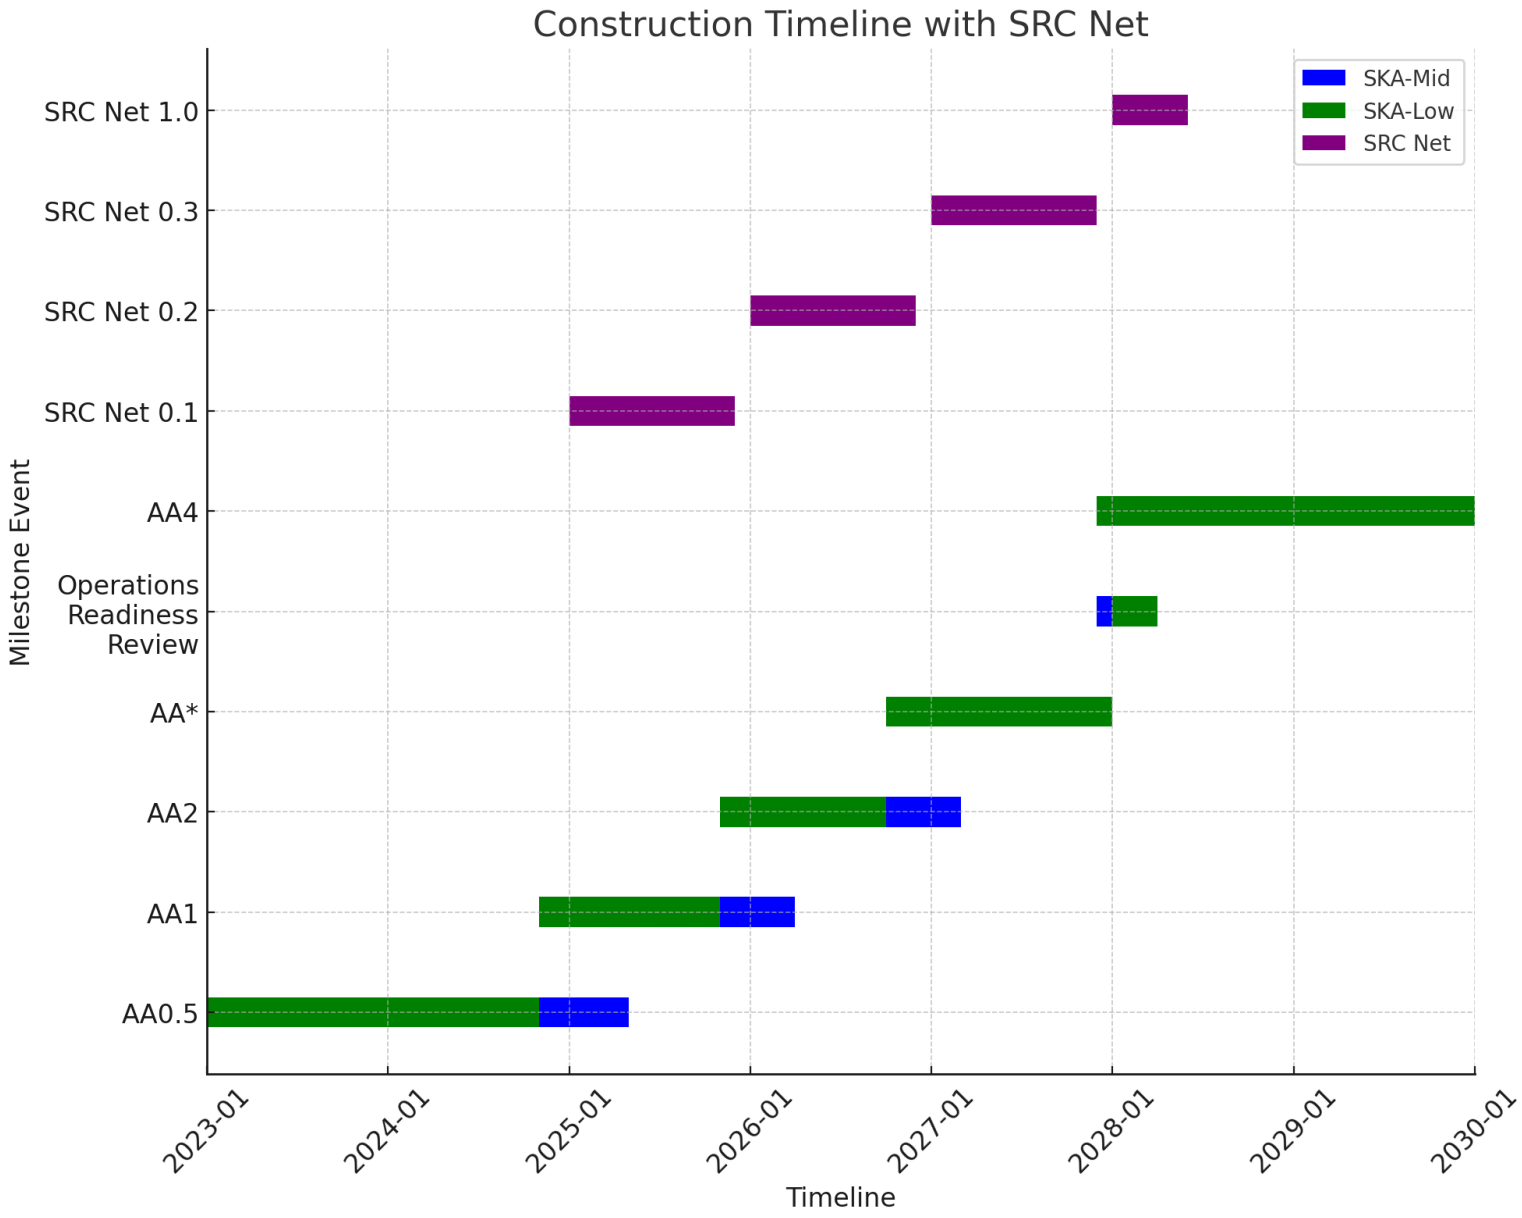

(courtesy: J. Walder)

SRCNet v0.1

Expression of Interest circulated for countries to participate in v0.1.

9 resulting sites; providing resources in line with SKA Top-level roadmap requirements

SRCNet0.1 included sites

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

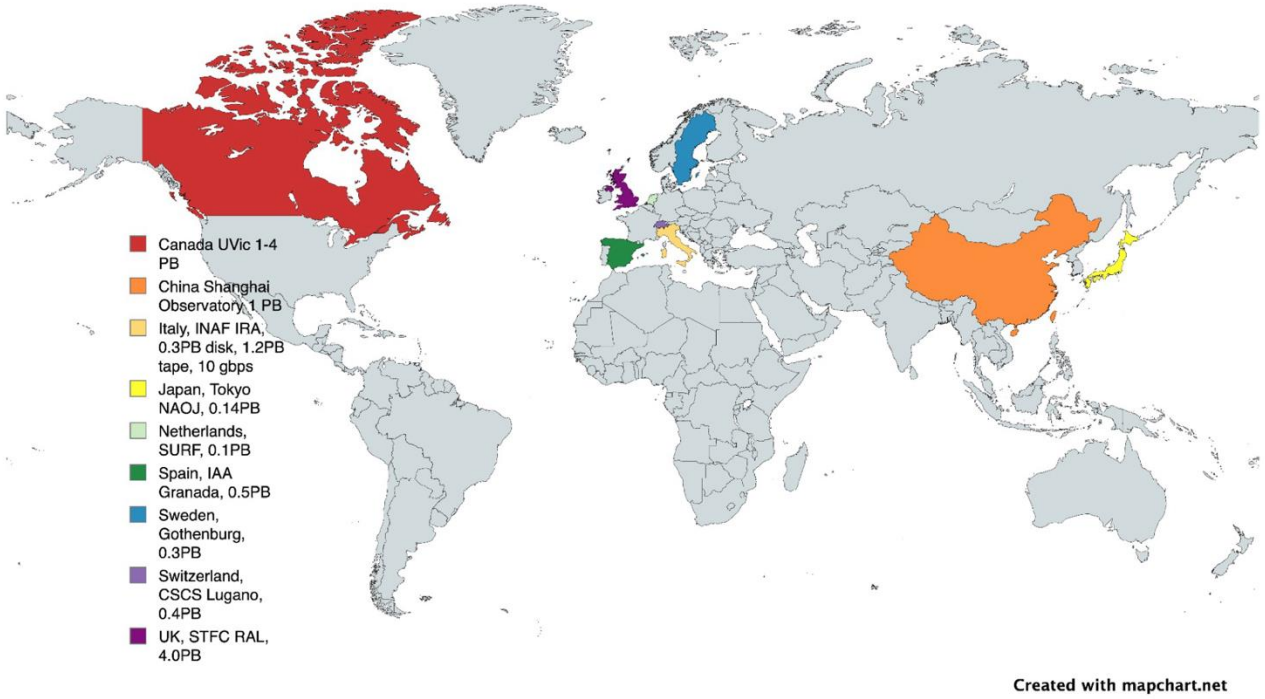

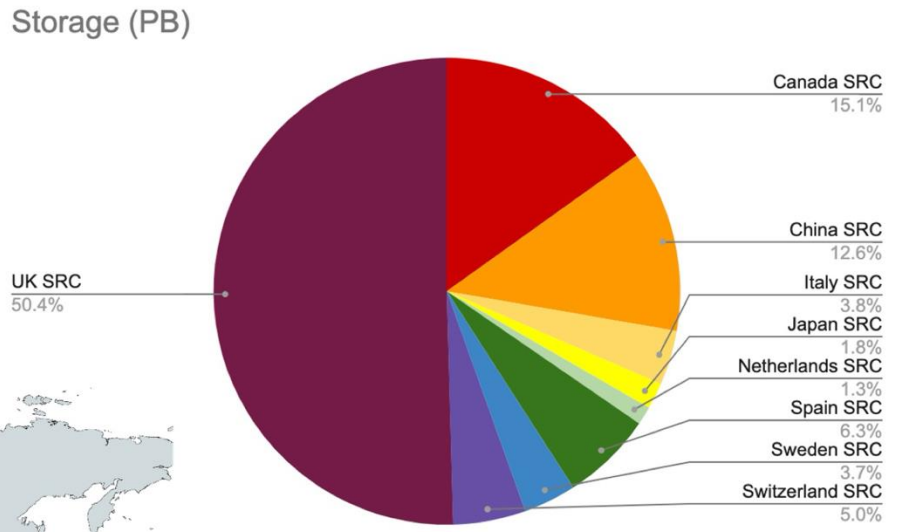

WLCG experience at some sites (Canada, Netherlands, Sweden, Switzerland, UK)

Several new sites and teams will learn by being involved

SRCNet v0.1 Deployment

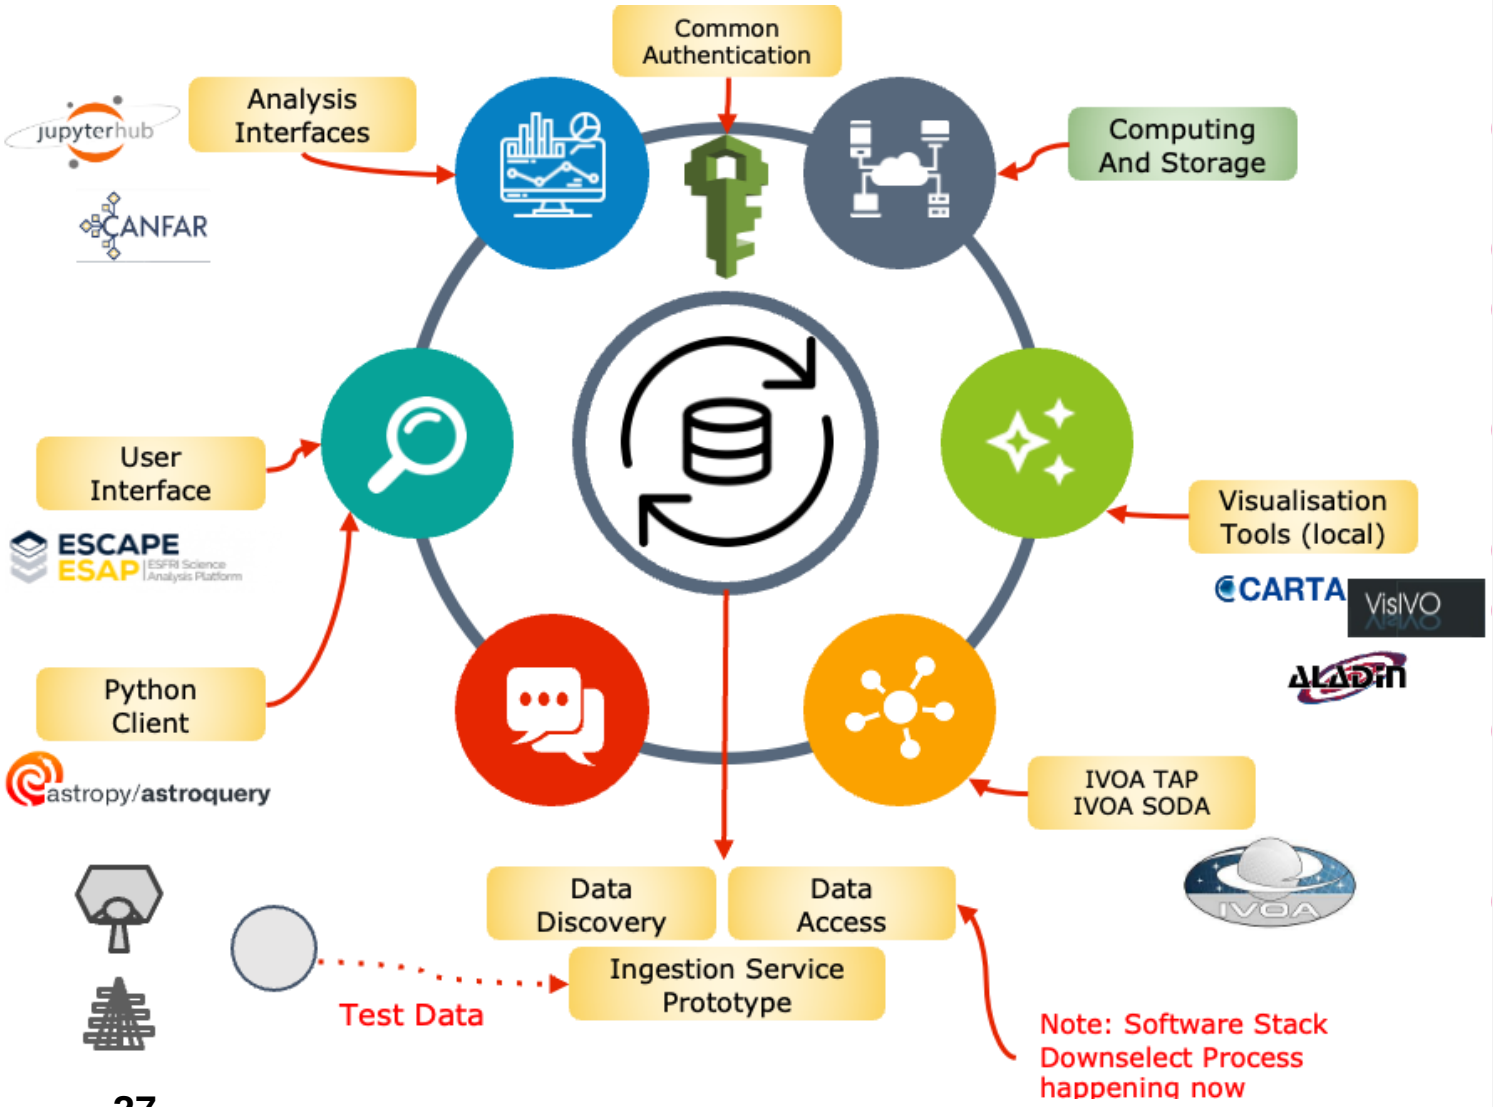

Global (compulsory) services

- Rucio, FTS servers
- SKA IAM (Indigo IAM)
- Software registry
- Global service monitoring

Local (compulsory) services

- Rucio Storage Element
- Local support for DM API (e.g., gate-keeper)
- perfSONAR
- Jupyter(Hub)
- SODA (cut-out service)
- Visualisation (e.g., Aladin, CARTA, VisIVO)
- Compute Platform (e.g., CANFAR, SLURM, Azimuth)
- *Recommended: Orchestrator, Local monitoring*

(courtesy: J. Walder)

Deployment Methodology I (K8S focus)

- Aim for common deployment methods in the SRCNet where possible
  - GitOps style approach recommended (e.g. ArgoCD/FluxCD, k8s);
    - Some difference in deployments due to heterogeneous sites / infrastructure adds complexity. (e.g. for established multi-project sites).
- Target the end goal to decouple site from service and enable deployment / upgrades with minimal complexity and allow central management (within reason …)
- **Tools** from (e.g) HEP, such as **FTS / XRootD** help to **abstract away** site implementation details to allow common designs.

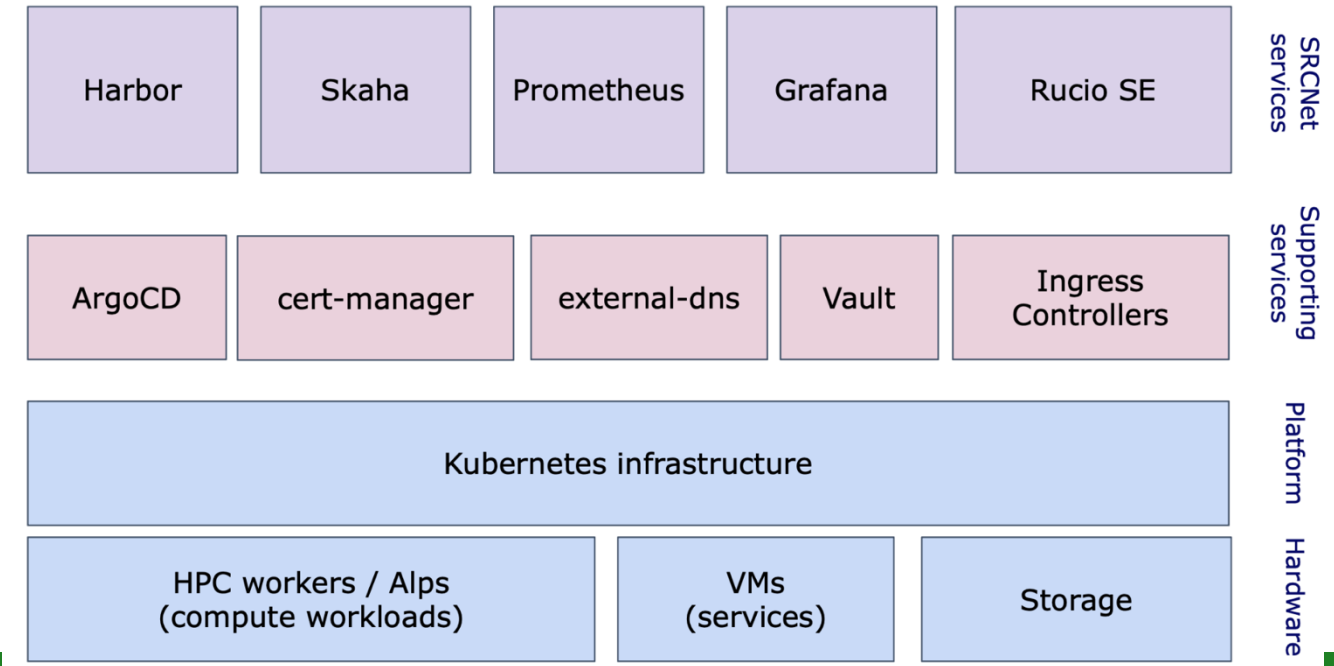

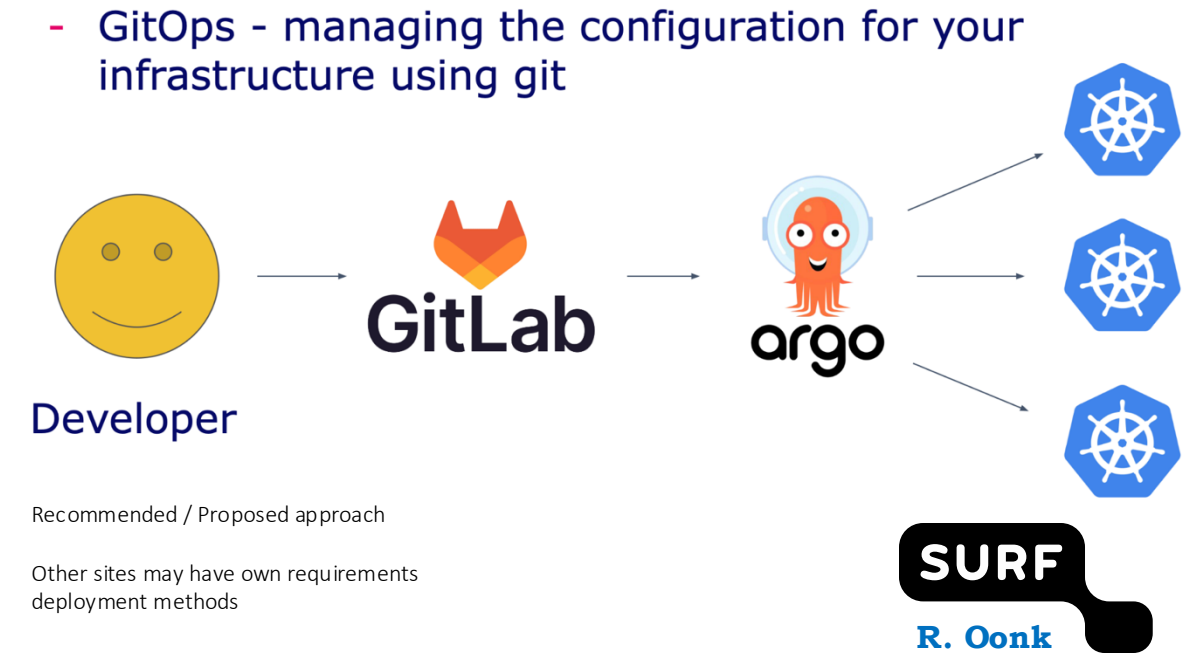

Deployment Methodology II (HPC / HTC focus)

Why:

- SLURM based scheduler with add-on services
- SLURM can support batch jobs and interactive (e.g., Jupyter) use
- Alignment with national and European initiatives
- Aligned with current end-user practices
- Reduce IT service fragmentation (less platforms & more features)
- Managed, persistent services
- Some services (e.g., perfSONAR) require a bare metal setup
- Shared use (optimize utilization)
- Proven scalability (for end-user and node as a whole)
- Hierarchical storage management

Public cloud: overcapacity caters to instantaneous needs

Academic: undercapacity and hence (job) queue-limited

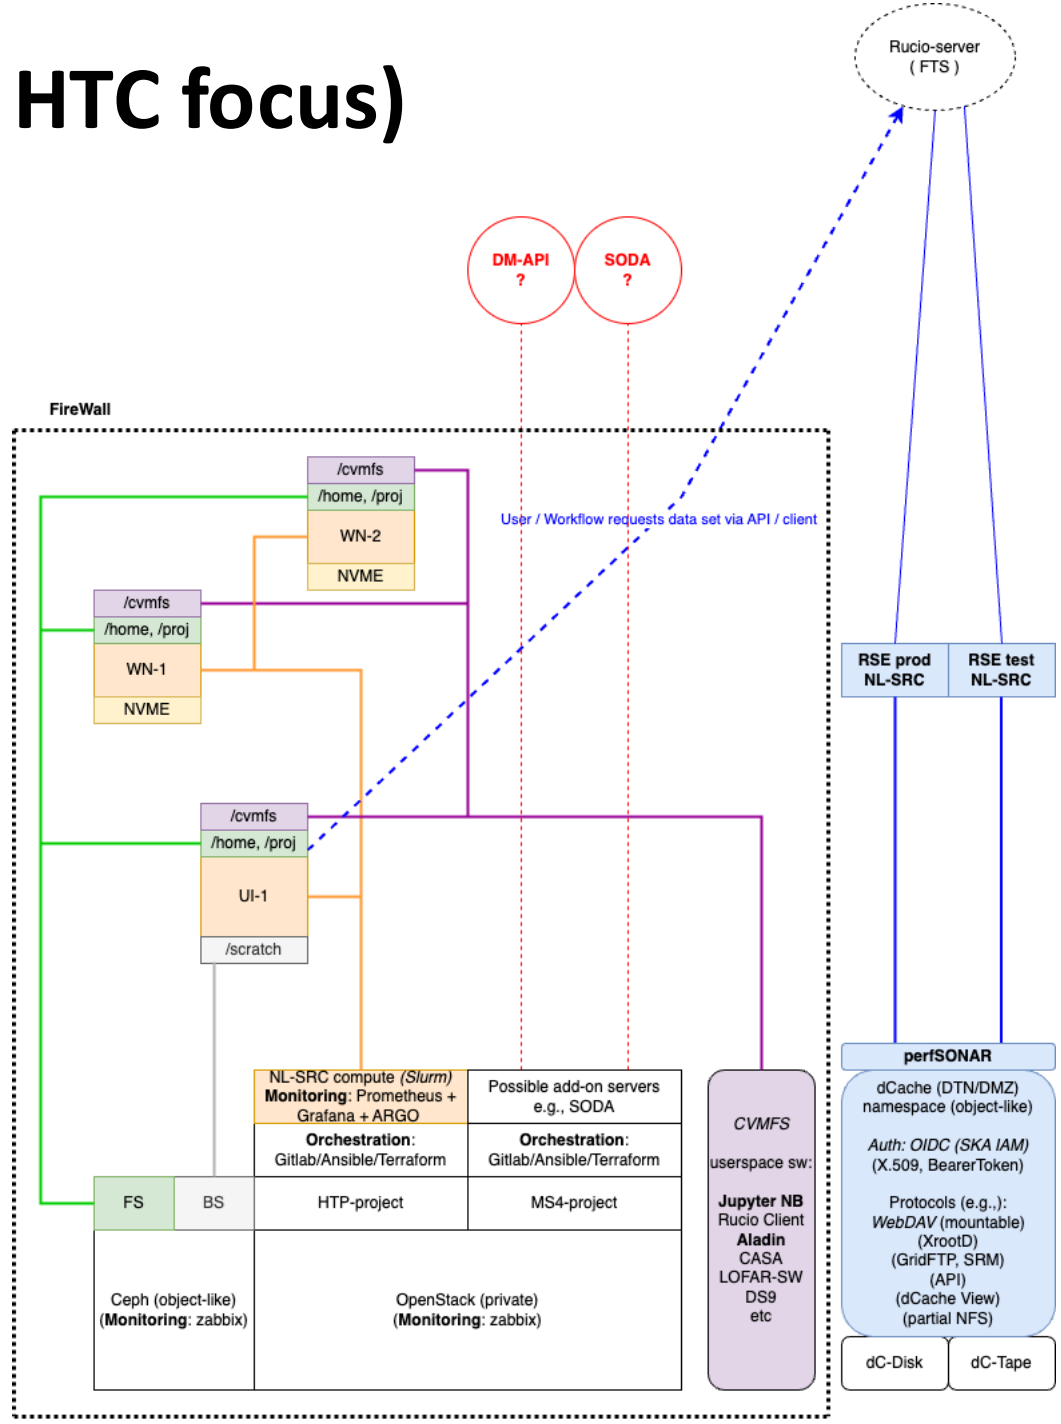

Data lake & network synergies with HEP (WLCG)

(courtesy: J. Walder)

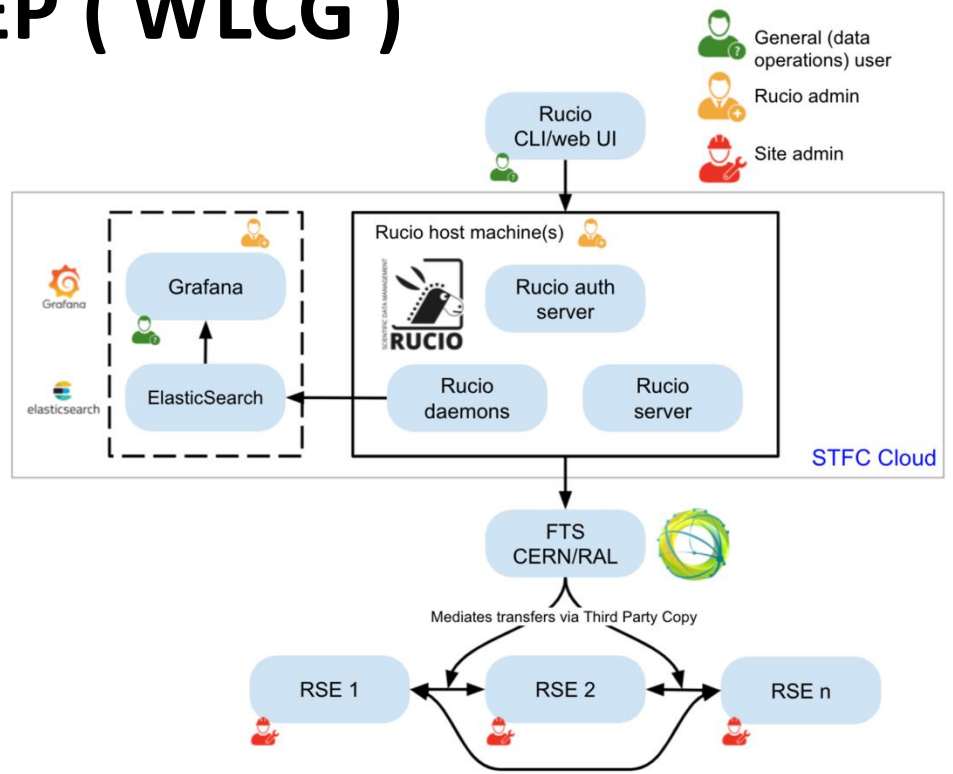

*. Significant prior work undertaken via ESCAPE and with pre-SRCNet v0.1 test deployments.

*. DDM assessment in early 2024:

  *. Conclusion to use **Rucio** (and associated stack) for SRCNet v0.1

    *. Other options, e.g. Storage Inventory considered

    *. Server currently hosted on (via k8s) STFC-Cloud (RAL)

  *. Abstracted via a *DM API service* for interactions with other service

    * Interoperability with Metadata services and other standards (IVOA-VOSpace)

  * **INDIGO-IAM**

    *. No use of X509–VOMS user certificates (token-based)

  *. **FTS**

  *. Sites using **StoRM, dCache, XRootD**

  * **perfSONAR** to be used for network monitoring and diagnosis

James Collinson, et. al

Transfer failure site matrix

Src\Dst	STFC_STORM_ND	STFC_STORM	SPSRC_STORM	NLSRC_PROD_DC	NLSRC_DCACHE	KRSRC_STORM	JPSRC_STORM	IMPERIAL
_XRD_DEVCEPHFS	0%	0%	0%	0%	0%	100%	0%	0%
STFC_STORM	100%	-	0%	0%	0%	100%	0%	0%
SPSRC_STORM	100%	100%	-	0%	0%	100%	0%	0%
RC_PROD_DCACHE	0%	0%	0%	-	0%	100%	0%	0%
NLSRC_DCACHE	0%	0%	0%	0%	-	100%	0%	0%
JPSRC_STORM	100%	100%	0%	0%	0%	100%	-	0%
IMPERIAL	0%	0%	0%	0%	0%	100%	0%	-
CNAF	100%	100%	0%	0%	0%	100%	0%	0%
CASRC_XRD	0%	0%	0%	0%	0%	100%	0%	0%

Data Lifecycle / Movement Challenges 2025

(courtesy: J. Walder)

- SRCNet v0.1 to execute a number of mini Data moving and Data lifecycle challenges.
  - Demonstrate data ingestion (this includes metadata ingestion)
  - Data movement campaigns
  - Data access and usage tests
- Public data and simulated data to be used in campaigns
- Dedicated O(week)s during the year (resources assigned when needed)
- Work ongoing to define the details of the tests
  - Main requirement to verify prototype architecture
  - Need also to explore scalability to prepare for scale of full operations
- Interest in participating in DC2X challenges as experience grows
  - Experience from a number of SRCNet v0.1 sites from WLCG already

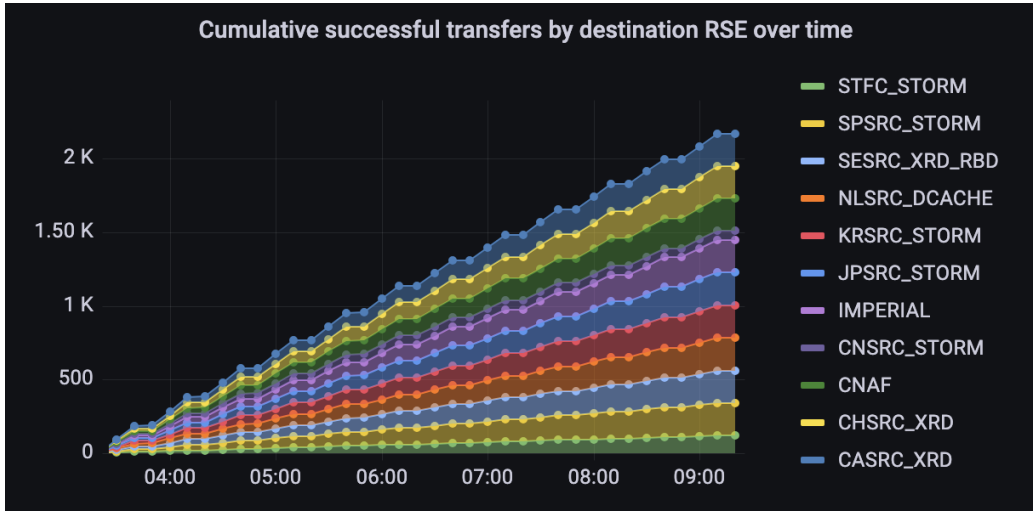

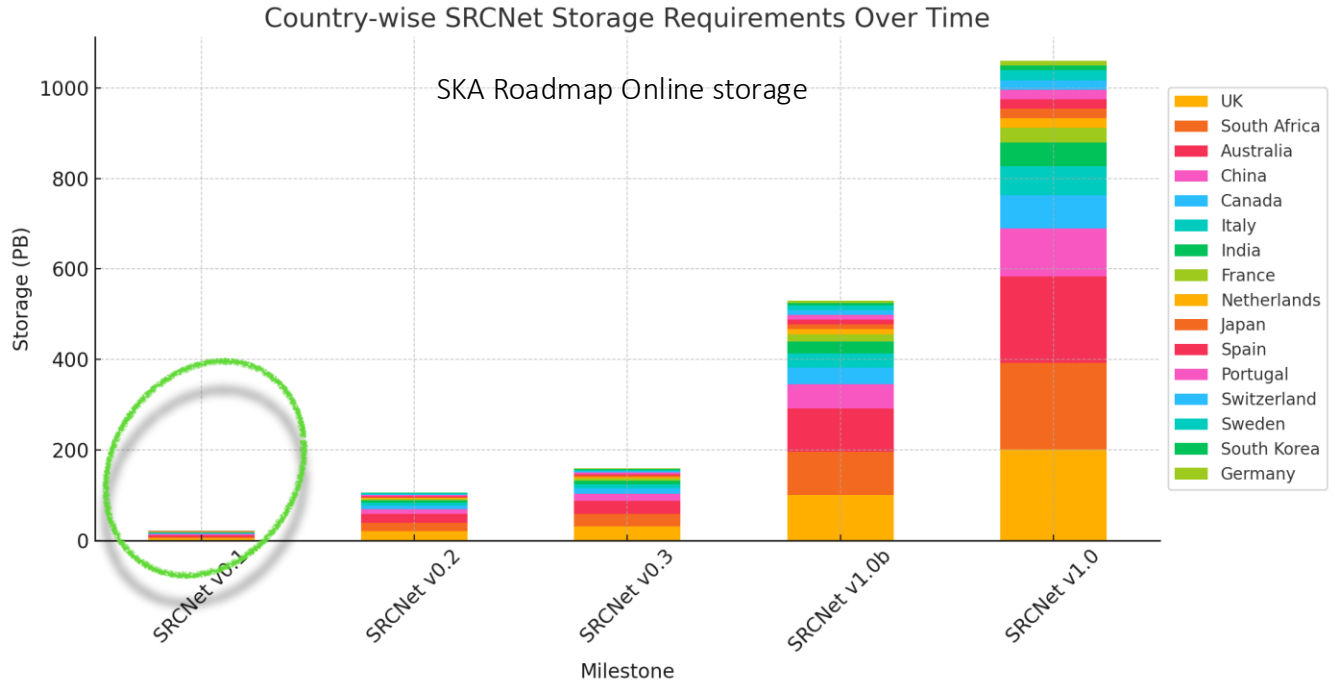

Possible Data Flows from the SKA and between SRCs

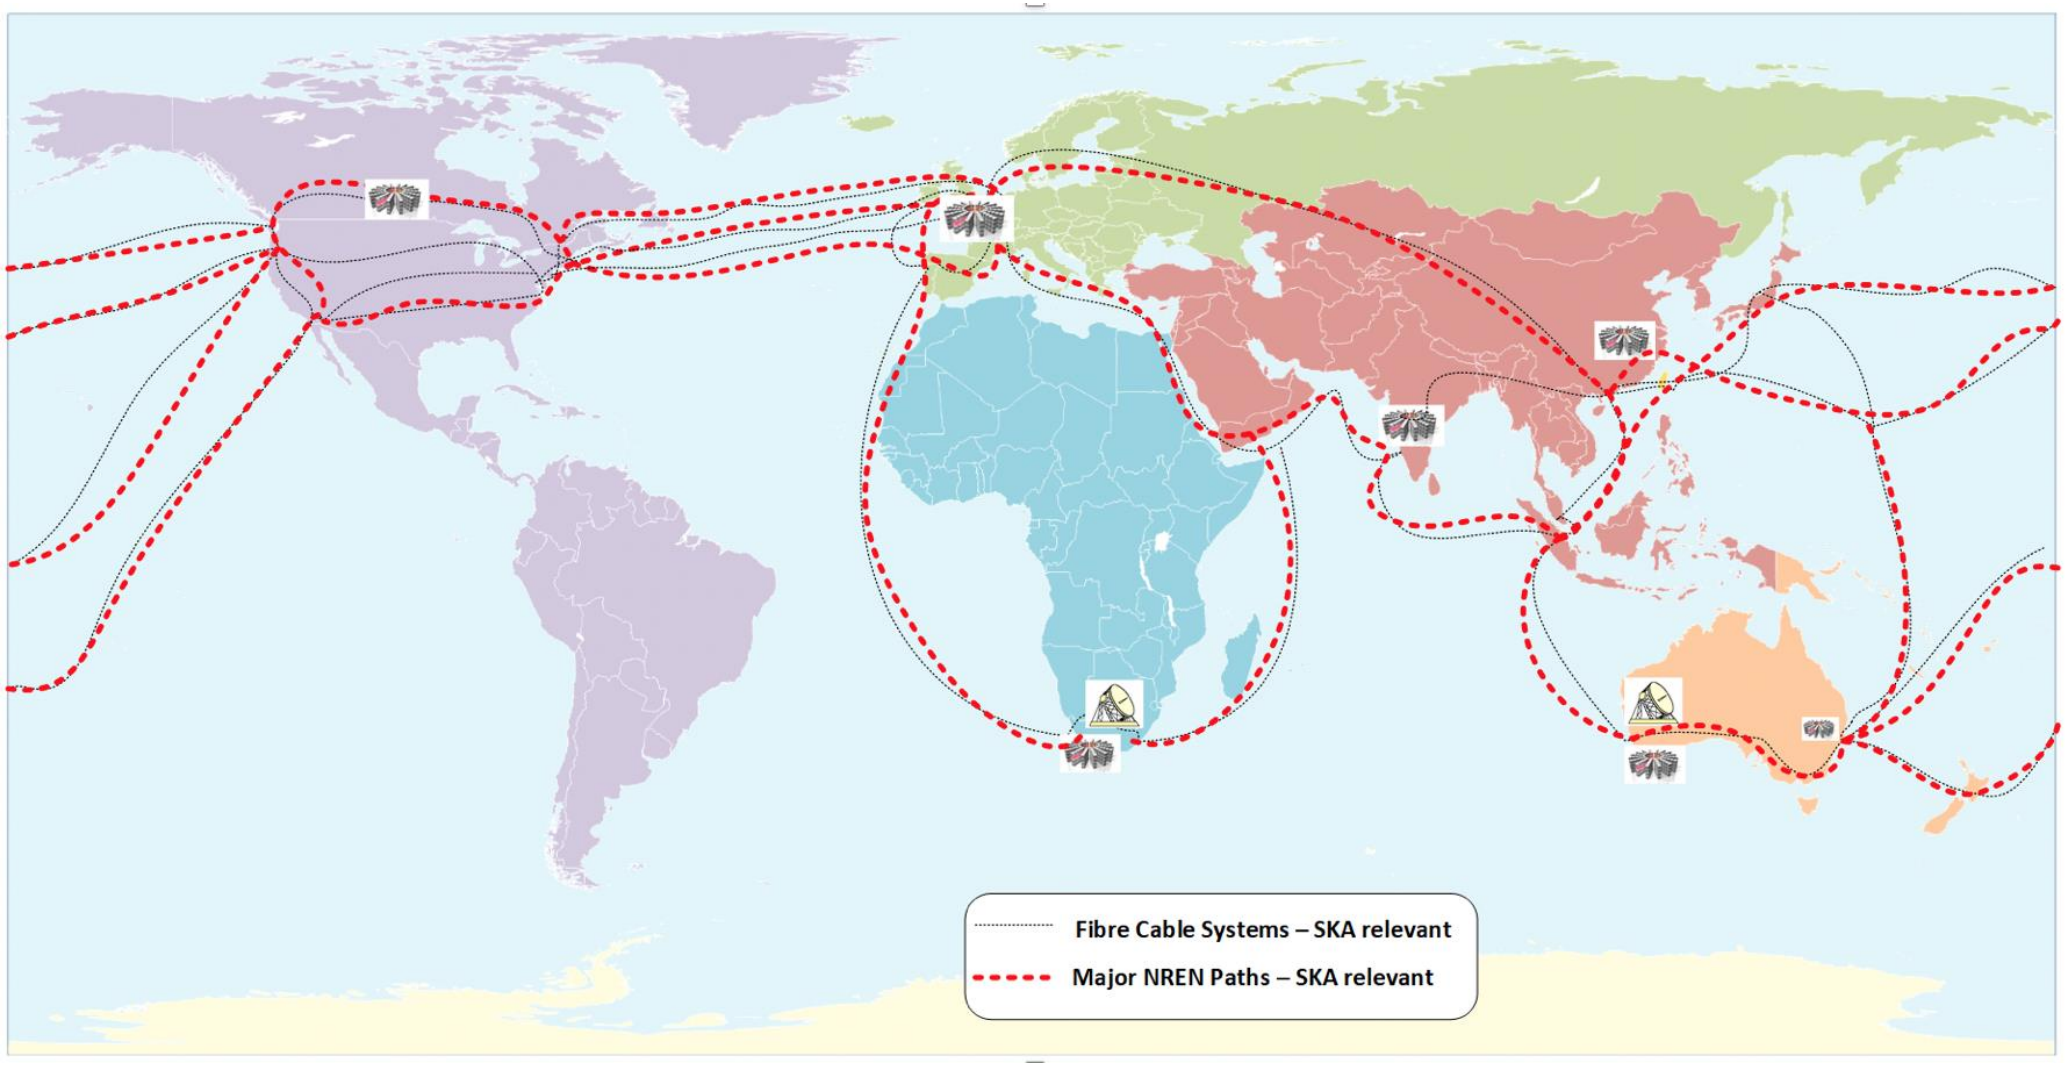

Observatory Data Products flow from the Science Data Processors in Cape Town and Perth into the SRC Network around the globe - and subsequently between SRC Nodes

SURF
R. Oonk

Personal observation: Some Technical Questions

- SRCNet v0.1 services not well defined (+ some scope creep)
  - e.g., some compulsory services not ready (also local vs. global control)
- Many Astronomy codes assume POSIX-like filesystems
  - Scalability to EB scale with POSIX not likely …
  - RSE scalability requires object(-like) e.g., dCache
  - WebDAV mounts , xCache to posix-like ?
  - TPC S3 <–> S3 ? (FTS [auth] not fully ready for S3)
- SRCNet wants to rely only on (OIDC) Tokens …
  - Tokens are in a development state (e.g., issues during DC24)
- Computing models appear to be overly complicated …
  - Compute, Discovery APIs: reinventing… (e.g., DIRAC, CRIC/GOCDB)
  - Monitoring: reinventing… (e.g., ARGO)

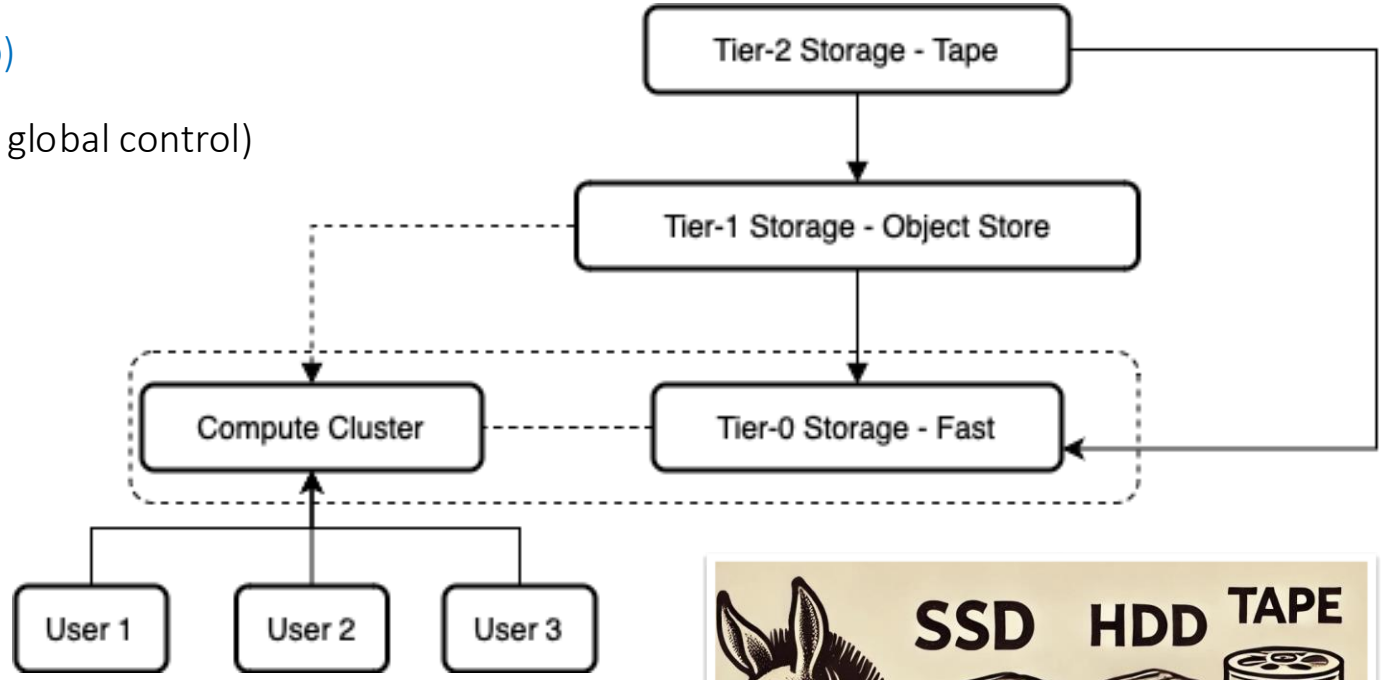

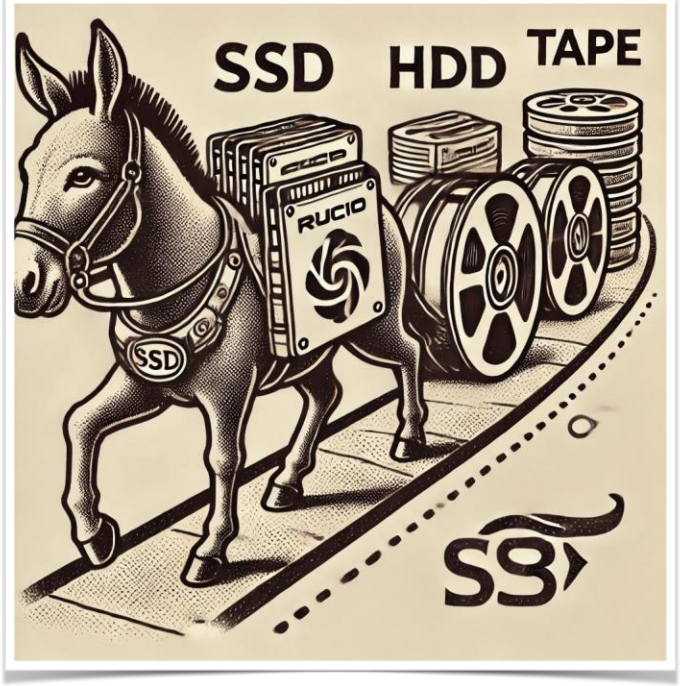

Personal observation: challenges in the current SRCNet development

1) Astronomers (end-users) not visible in the SRCNet architecture and implementation process
2) Lessons from LOFAR and other SKA precursors/pathfinders, although ahead, not (yet) visible within SRCNet
3) Insufficient connection / communication between SDP and SRCNet
4) Focus on adapting infrastructure to (classical) applications (instead of vice versa)
5) Focus on uniform deployment rather than inherent heterogeneity and diverse evolutions
6) Little to no alignment with European-scale (IT) innovation initiatives (e.g., EuroHPC, Dataspaces)
7) SRCNet nodes are funded locally (not SKAO)
8) SRCNet does not (yet) have a governance structure (not SKAO)
9) Policies (incl. security) not (yet) in place
10) Procedures (incl. security) not (yet) in place
11) Service development ‘easy’ but maintenance and evolution less so

** Data replication policy for the SKAO not yet available for SRCNet*

Personal observation: What can LOFAR learn from the current SRCNet developments ?

LOFAR is ahead with in terms of operationally delivering large amounts of data to researchers

LOFAR is ahead with prototyping bulk processing workflows & community/infra engagement

LOFAR is ahead in terms of the data model (access, policies etc)

But,

1) SRCNet abstracts data management at the Rucio/FTS level whereas LOFAR works at the dCache level
2) SRCNet is setting up a functioning perfSONAR mesh to monitor the network
3) SRCNet is setting up a science gateway, global service monitoring, discovery and software distribution
4) SRCNet is trying to understand how posix-like applications (Jupyter, SODA) interact with object-like stores
5) SRCNet is promoting a token-based approach to compute and data (avoid e.g., ssh keys and certificates)
6) SRCNet is trying to find a balance between global releases/interfaces and local deployments
7) SRCNet is investigating harmonized policies for compute (and data) access

EXTRA

Are we getting ready for EB-scale data ?

Are we (getting) ready for the SRCNet ?

Are we ready for the SRCNet - Computing ?

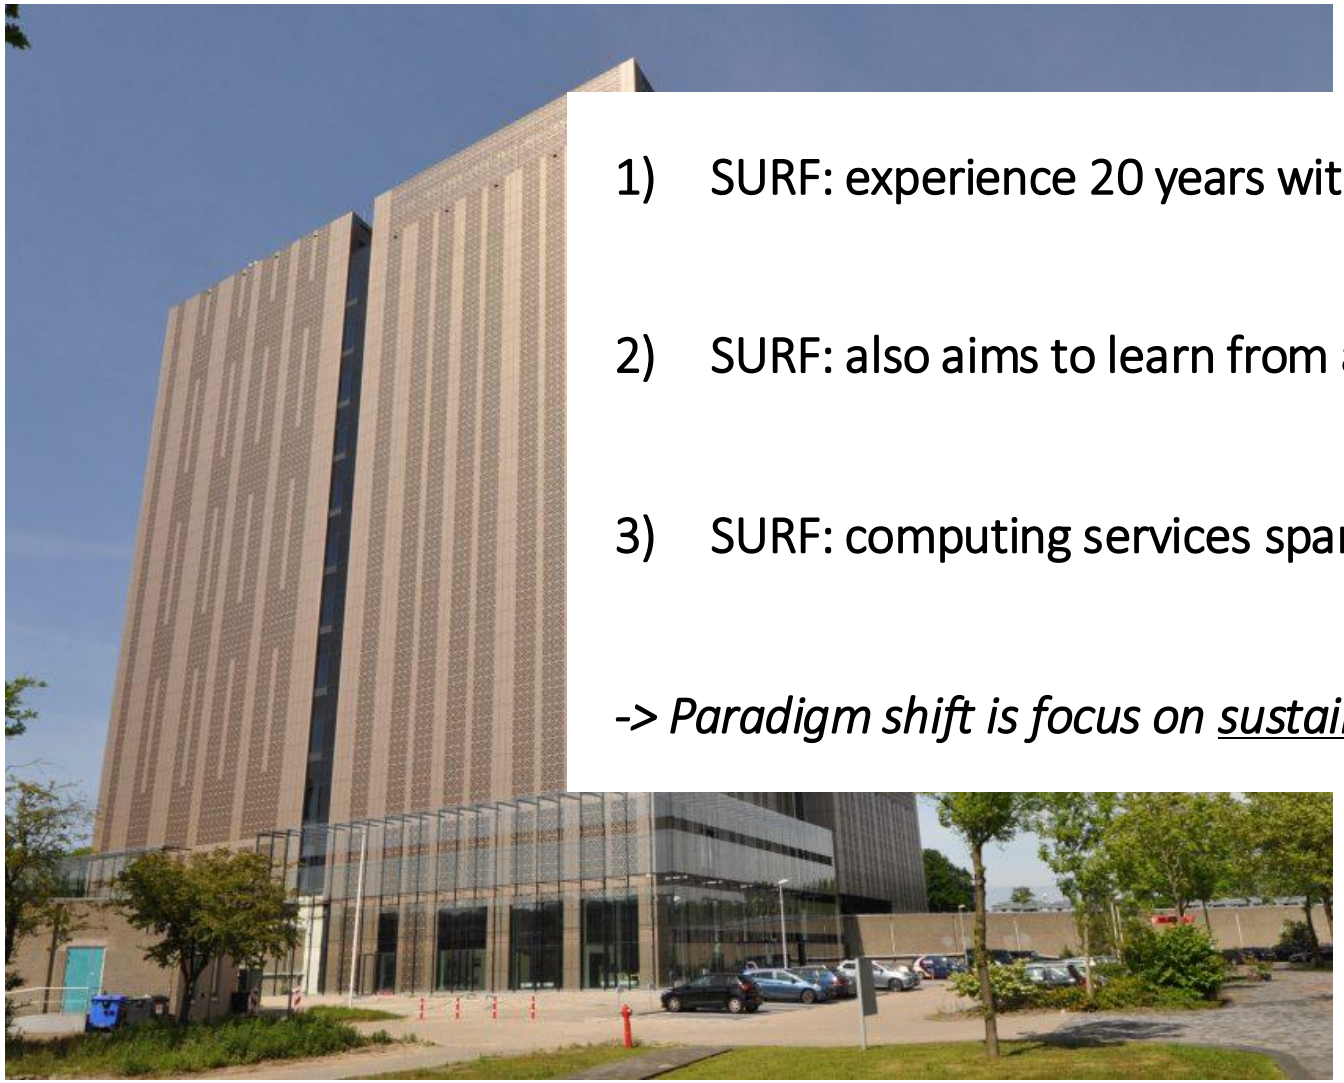

1) SURF: experience 20 years with WLCG (Tier-1) and 15 years with LOFAR

2) SURF: also aims to learn from and apply to other research domains

3) SURF: computing services spanning a range in needs (e.g., HPC, HTC, Cloud)

-> Paradigm shift is focus on sustainability and effective inclusion of HPDA / AI

Are we ready for the SRCNet - Storage ?

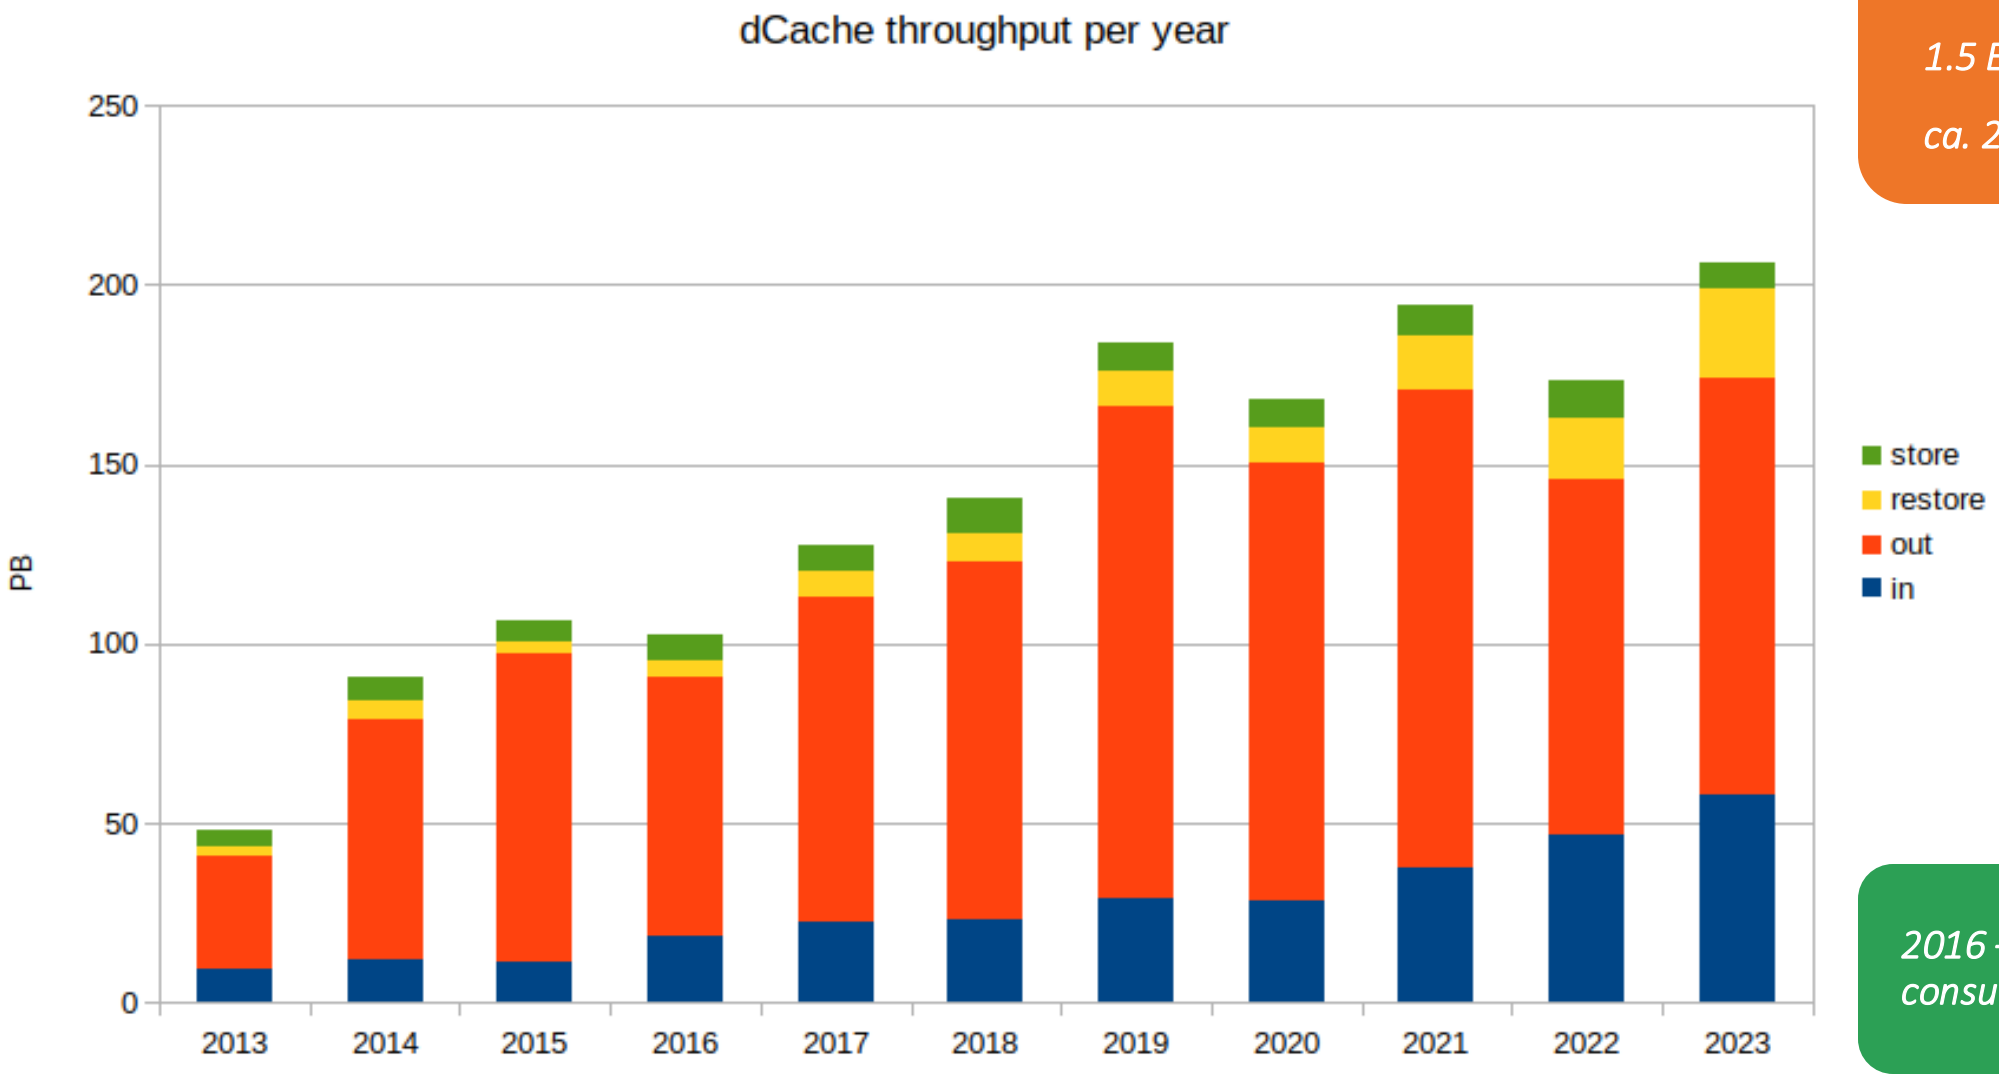

Throughput aggregate 10 yr is 1.5 EB, but manage in 2024 ca. 250 PB of research data

2016 –> 2023 same power consumption for 3x storage

Are we ready for the SRCNet - Network ?

Geneva – Amsterdam (800 Gbps)

Hardware used in the Trial

- SURF's line system with Ribbon hardware
- Nokia PSE-6s Transponders (beta cards)
- Nokia FP5 based SR-1 routers
- SURF's packet canons
  - AMD EPYC 9554P 64C 128T @3,76 Ghz
  - 768GB DDR5 4800 MT/s
  - 4x dual port Mellanox ConnectX-7 connected with 8x100G to the Nokia FP5