LOFAR/ASTRON Status

RADIO Workshop, January 2026

Hanno Holties, André Offringa, John Swinbank

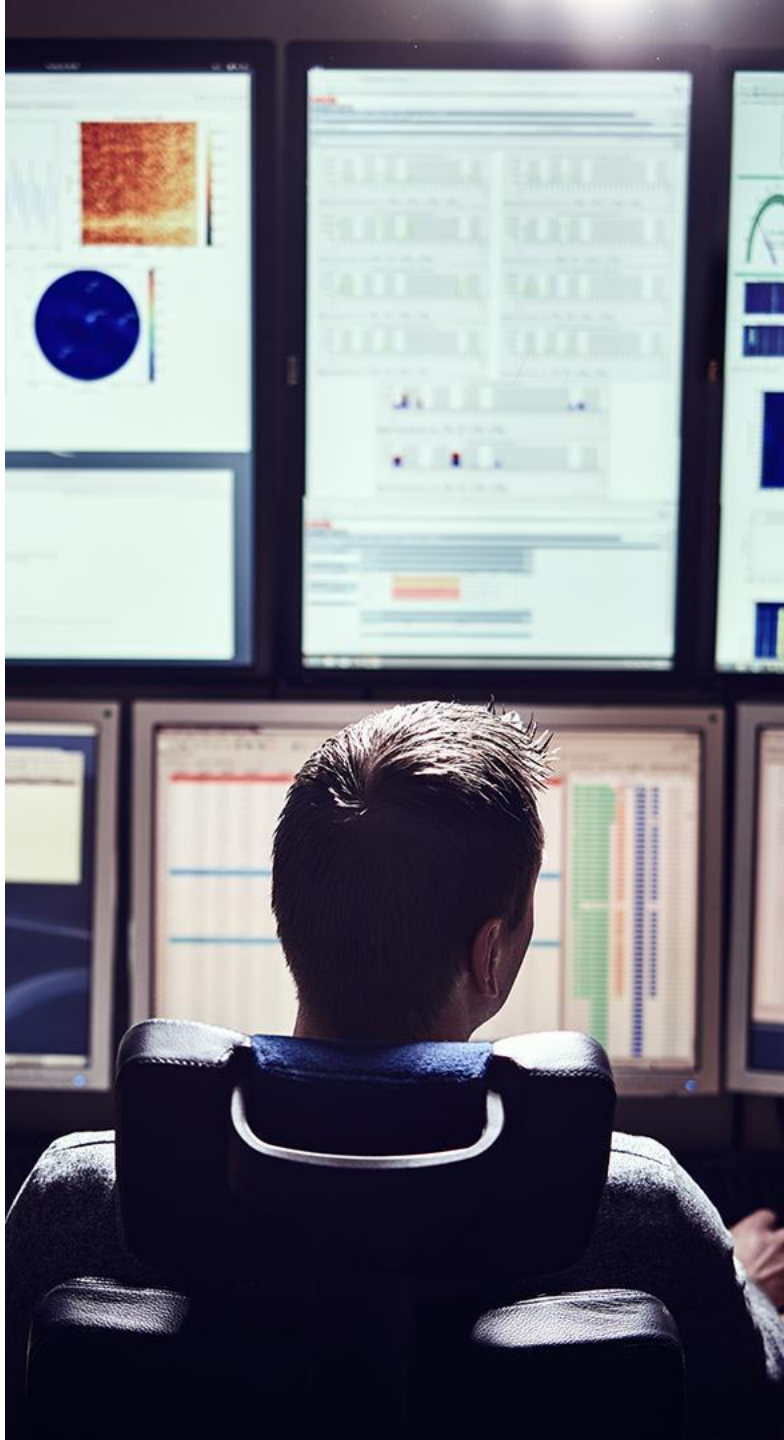

ASTRON

Netherlands Institute for Radio Astronomy

Part One

Introductions & context

Context

- LOFAR is an open-access low-frequency (30–250 MHz) radio telescope spanning Europe from Ireland to Bulgaria.
- The LOFAR Radio Telescope is under the control of LOFAR ERIC (European Research Infrastructure Consortium).
  - This is a consortium of 9 nations: Bulgaria, the Czech Republic, Germany, Ireland, Italy, the Netherlands, Poland, Sweden and the United Kingdom.
  - Policy and decision making power lies with the ERIC Council, representing all of those countries.
- ASTRON has primary responsibility for developing and operating the LOFAR Observatory under the terms of an SLA with LOFAR ERIC.
- Other nations contribute financially, by hosting stations, by in-kind contributions of staff, etc.

Introductions

- The ASTRON **Science Data Centre** is the structure within ASTRON which is responsible to our contributions to (much of) the LOFAR data system, and for providing similar services to other facilities (e.g. SKA SRCNet).
- **Hanno Holties** is the SDC System Architect.
- **André Offringa** (coming tomorrow) is the LOFAR Pipelines Architect.
- **John Swinbank** is the SDC Development Programme Manager.

LOFAR History

- Opened by Queen Beatrix in 2010.
- Twice annual proposal call until 2024.
- Archive of ~50PB of “intermediate” data products (
- Most science produced by teams downloading intermediate data and processing it on their own systems.
  - 1000 refereed publications to date; top 10% of astronomical facilities worldwide.

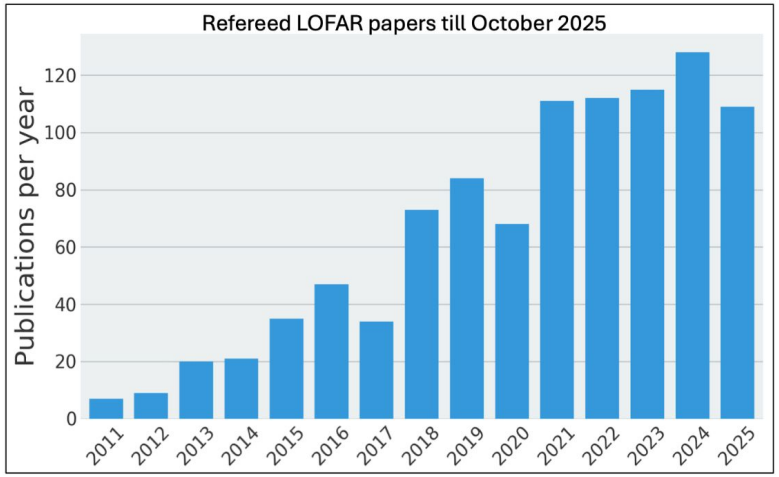

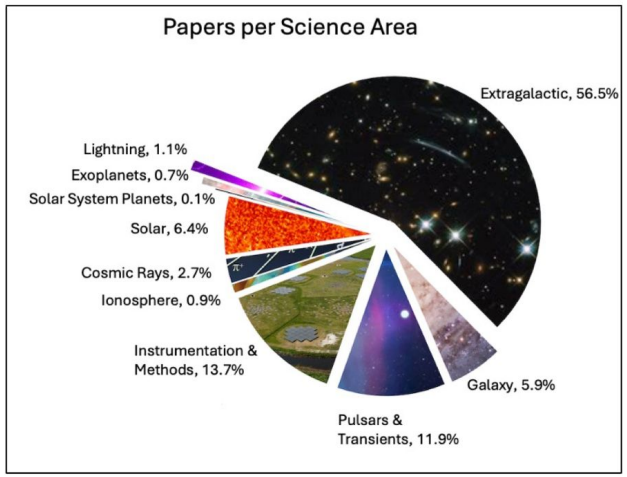

Figure: Roberto Pizzo

LOFAR2.0

- For the last ~18 months, LOFAR has been offline for a major upgrade: **LOFAR2.0**.
  - Simultaneous low & high band observing.
  - More bandwidth.
  - Wider field of view.
  - More flexible control systems.
  - etc etc.

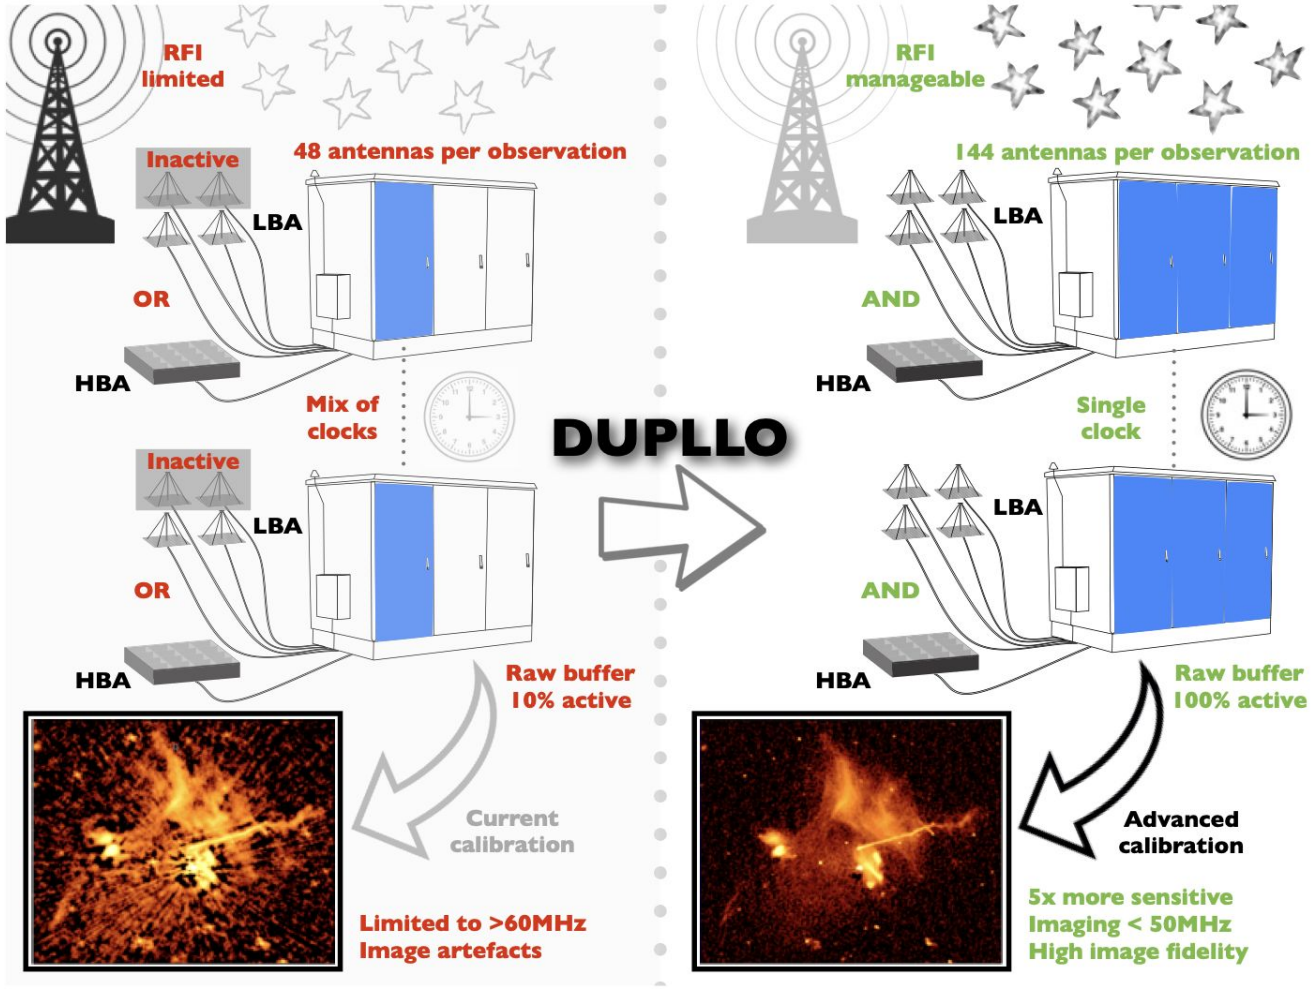

LOFAR Data Handling #1

- The way we handled data for LOFAR1 had some undesirable consequences.
  - The archive accumulated a large volume of intermediate data products which are only occasionally accessed;
  - Final, “science ready” data products are generated outside the Observatory’s control, and may be hard for those outside the responsible science team to access.
  - Generating science ready products requires significant resources and expertise, so the “average” astronomer is excluded.
  - The Observatory cannot offer science-ready products to end users.

LOFAR Data Handling #2

- In LOFAR2 we will address this by:
  - Bringing the generation of (most) advanced data products under Observatory control;
  - Requiring that (most) advanced data products generated outside the Observatory are returned to the archive;
  - Retaining intermediate products for only a limited time;
  - (Eventually) providing online, “science platform” type functionality for end users.
- Aim to increase *accessibility* and *impact*.

LOFAR2: Large Programmes

- The LOFAR2 schedule is based around a series of “large programmes” (L2LPs).
- Large projects; ~thousands of hours of observing per programme.
- Results of a public call a couple of years ago.
- Wide range of science cases: imaging surveys, pulsar timing, transient detection, EoR, lightning, …
- L2LPs help drive design & priority decision making.
- Execution of a large programme is contingent on an agreed data management plan & completing science verification.
- Pending that, exact time allocations are TBD.
- Some time (~20%?) will be available for regular “open skies” calls.

LOFAR2: Schedule

- First version of LOFAR2 – the “minimum production system” – available shortly.
  - Capable of conducting basic operations in a limited number of observing modes, and will low operational efficiency.
- Ongoing development during 2026 expected to enable the first L2LP observing.
- Full operations expected 2027.
- …and then…

LOFAR **E**nhanced **N**etwork for **S**harp **S**urveys

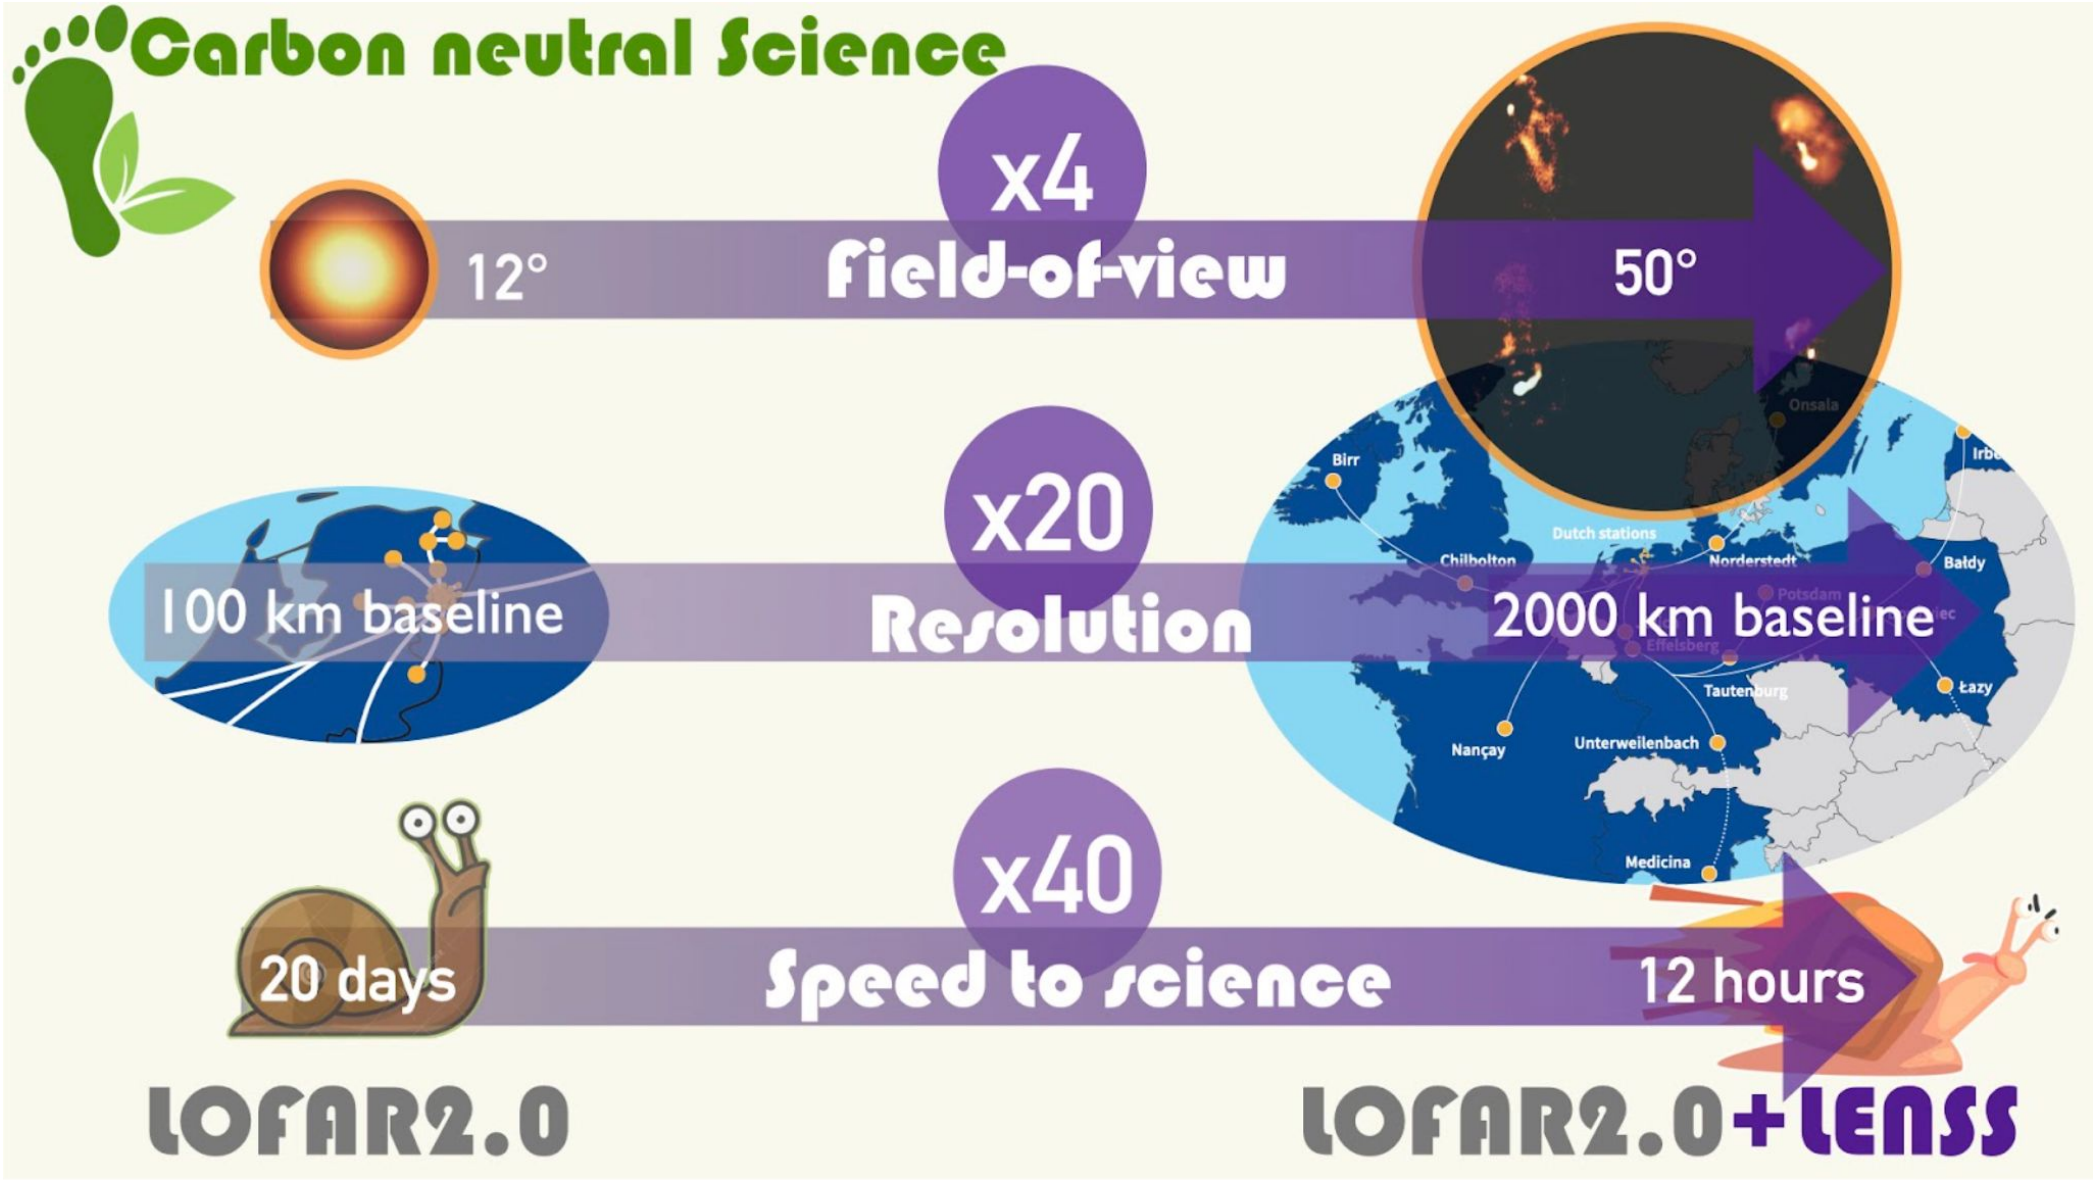

Operational by the early 2030s

Part Two

System overview

Old (LOFAR1 era)

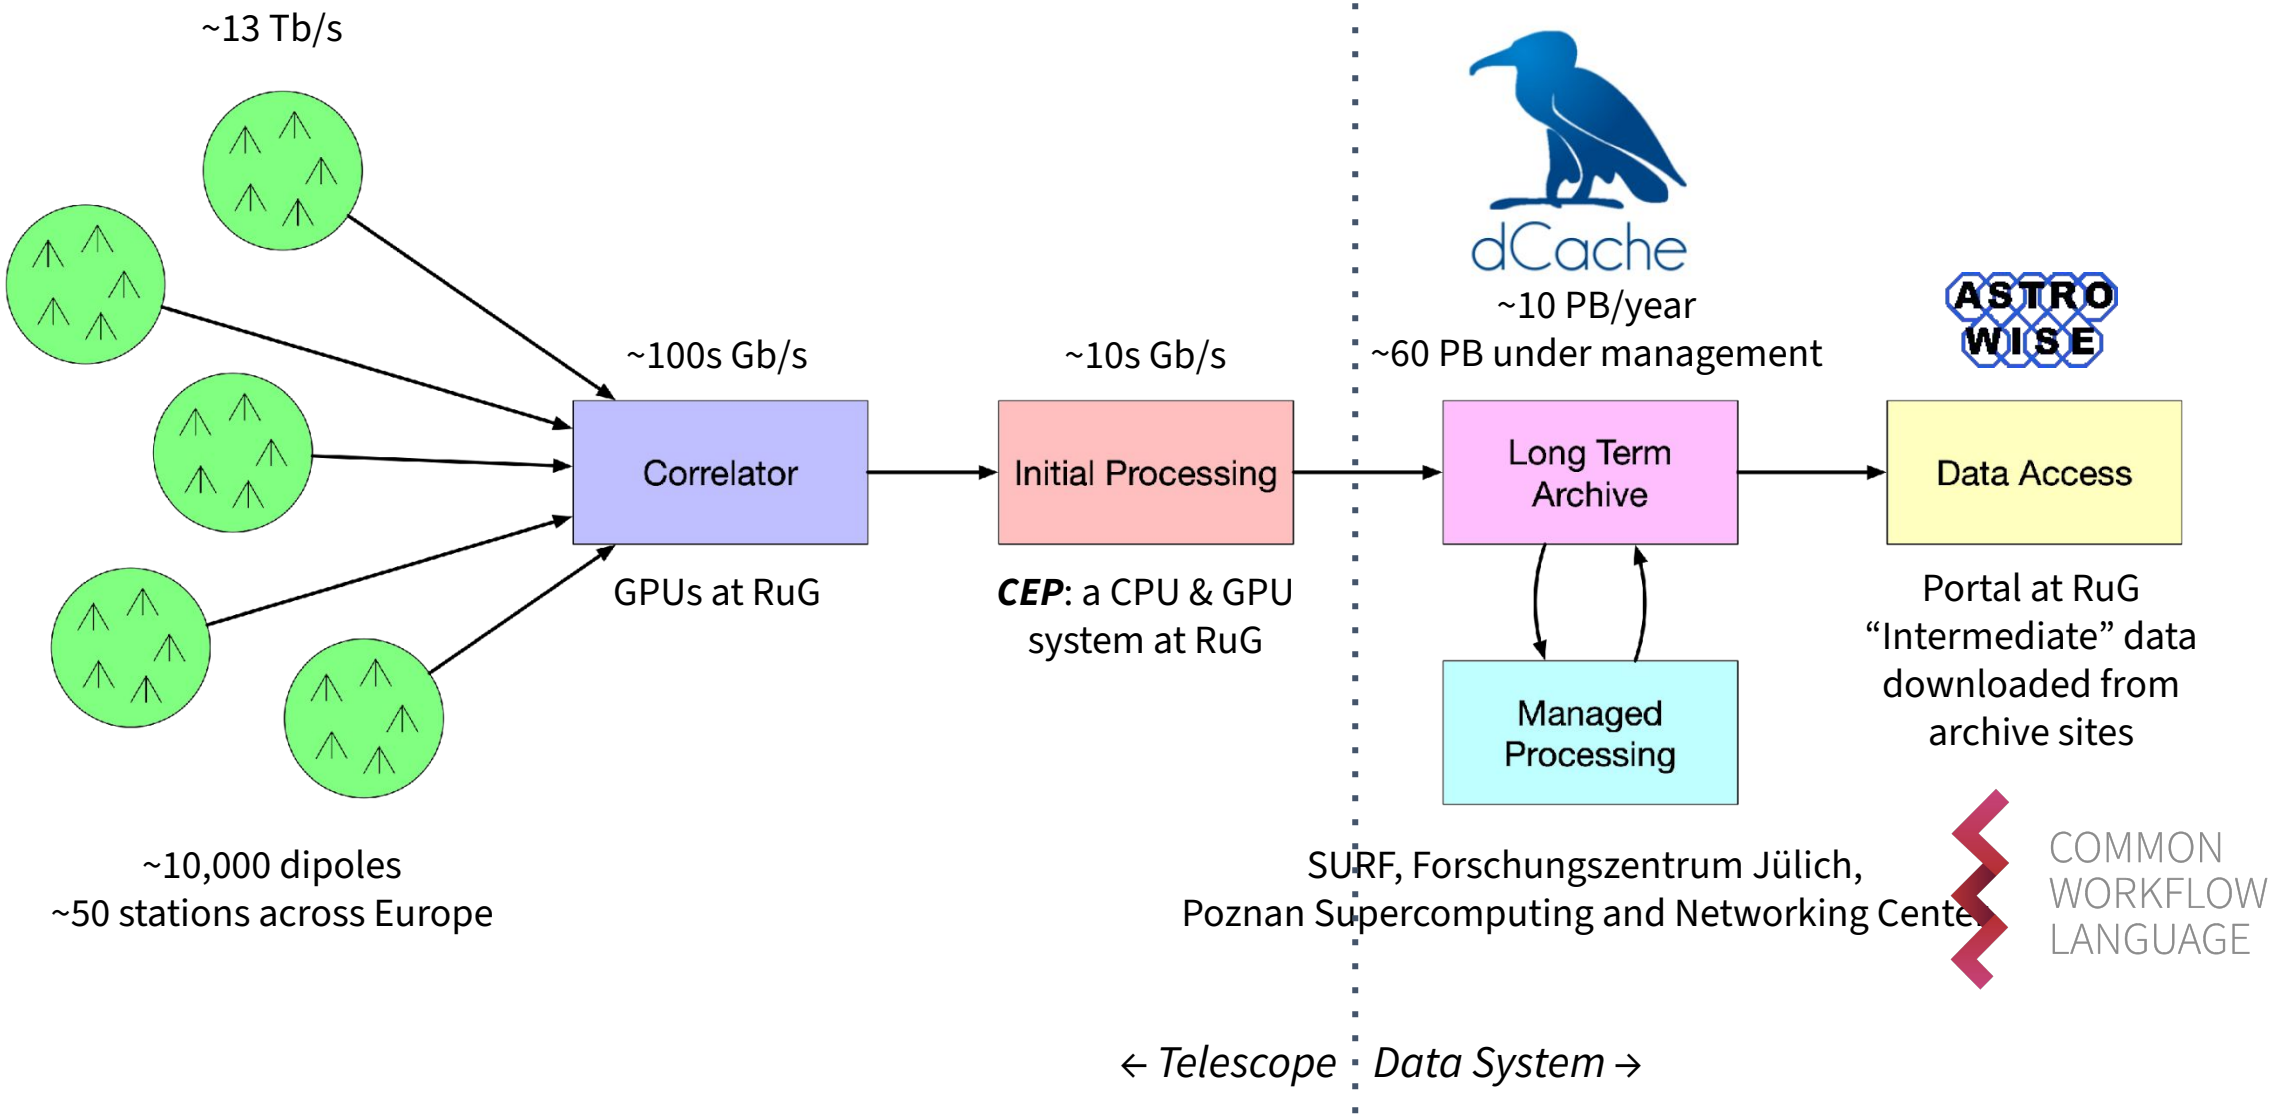

Next (LOFAR2.0 era)

LOFAR2.0 Upgrade (2025/26)

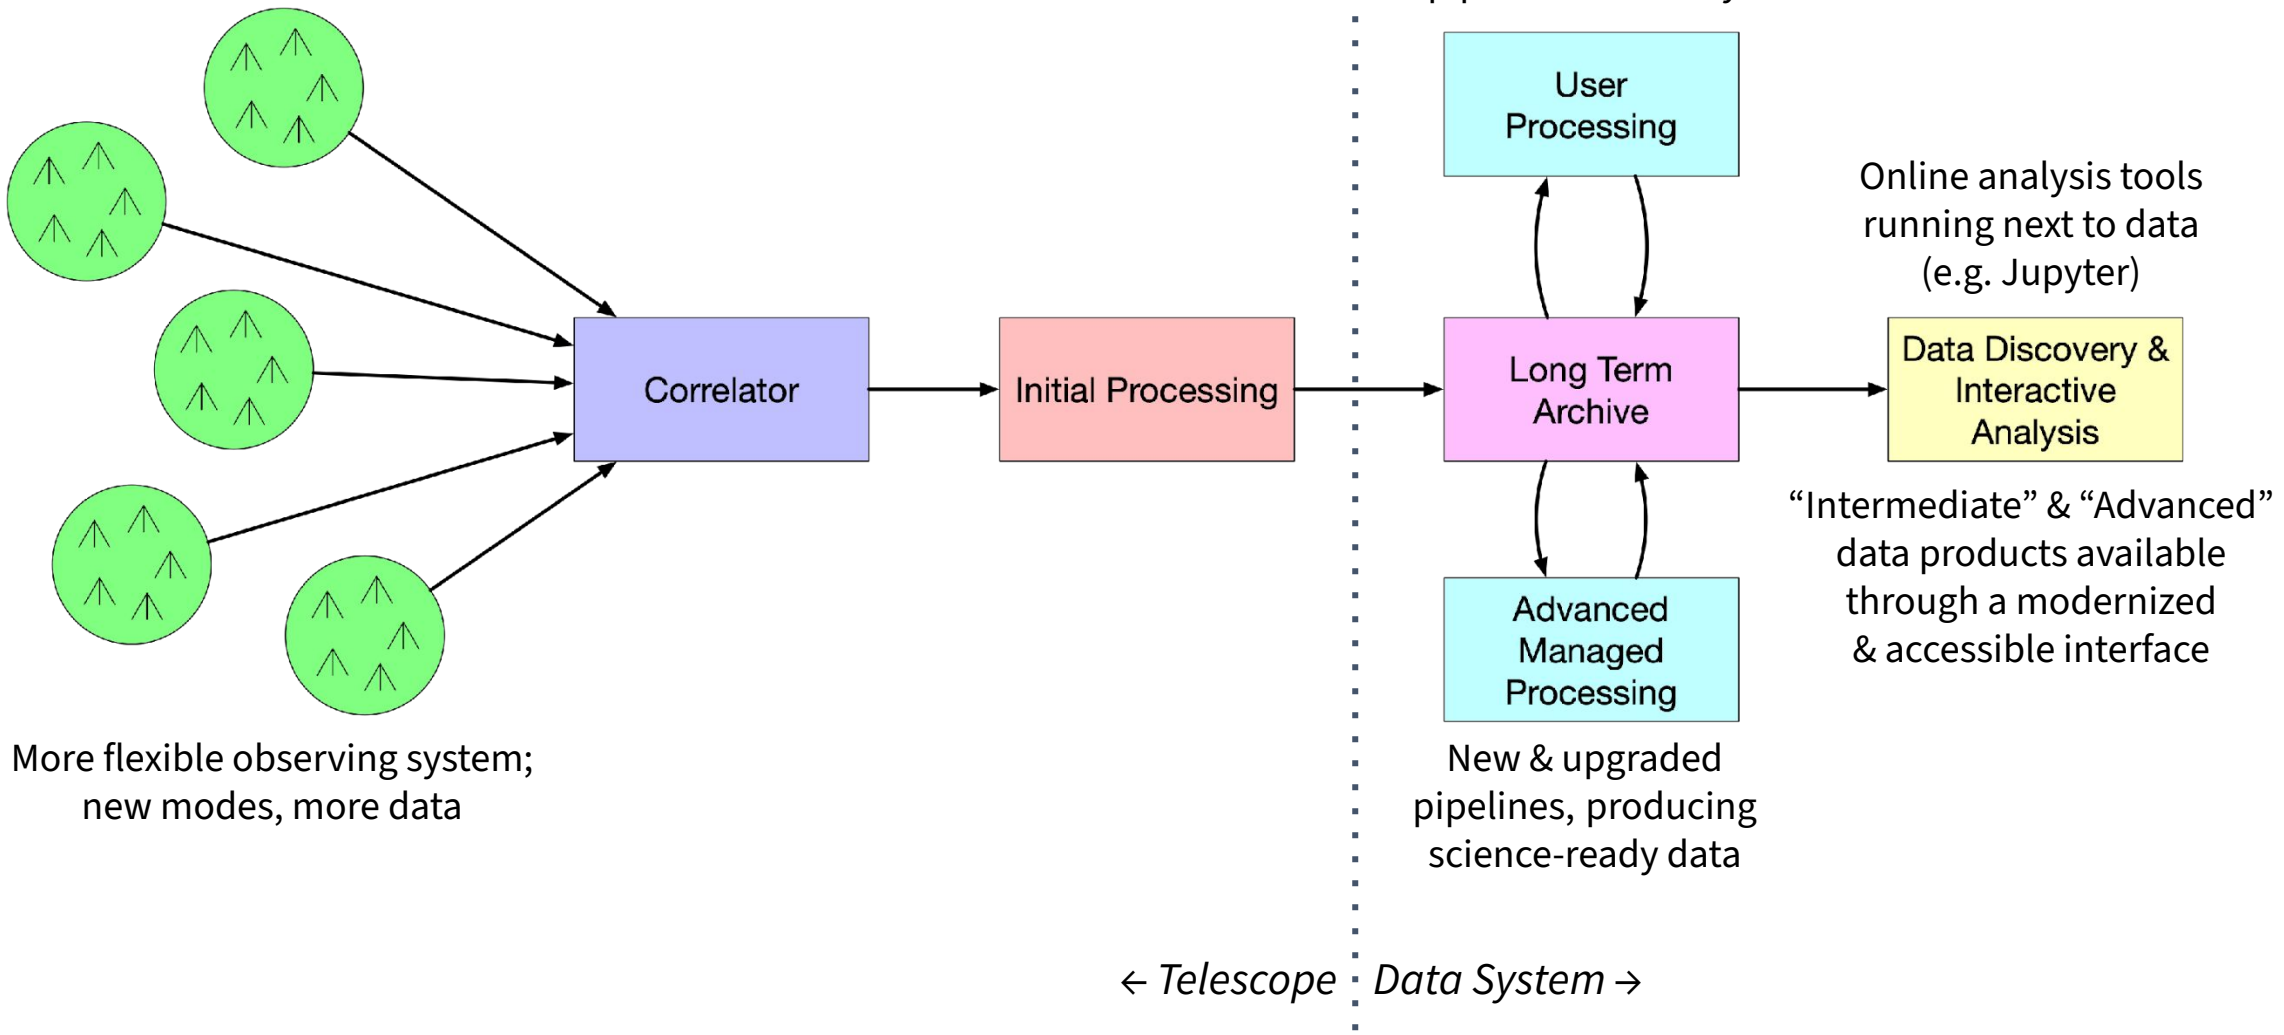

Future (LENSS era; 2030+)

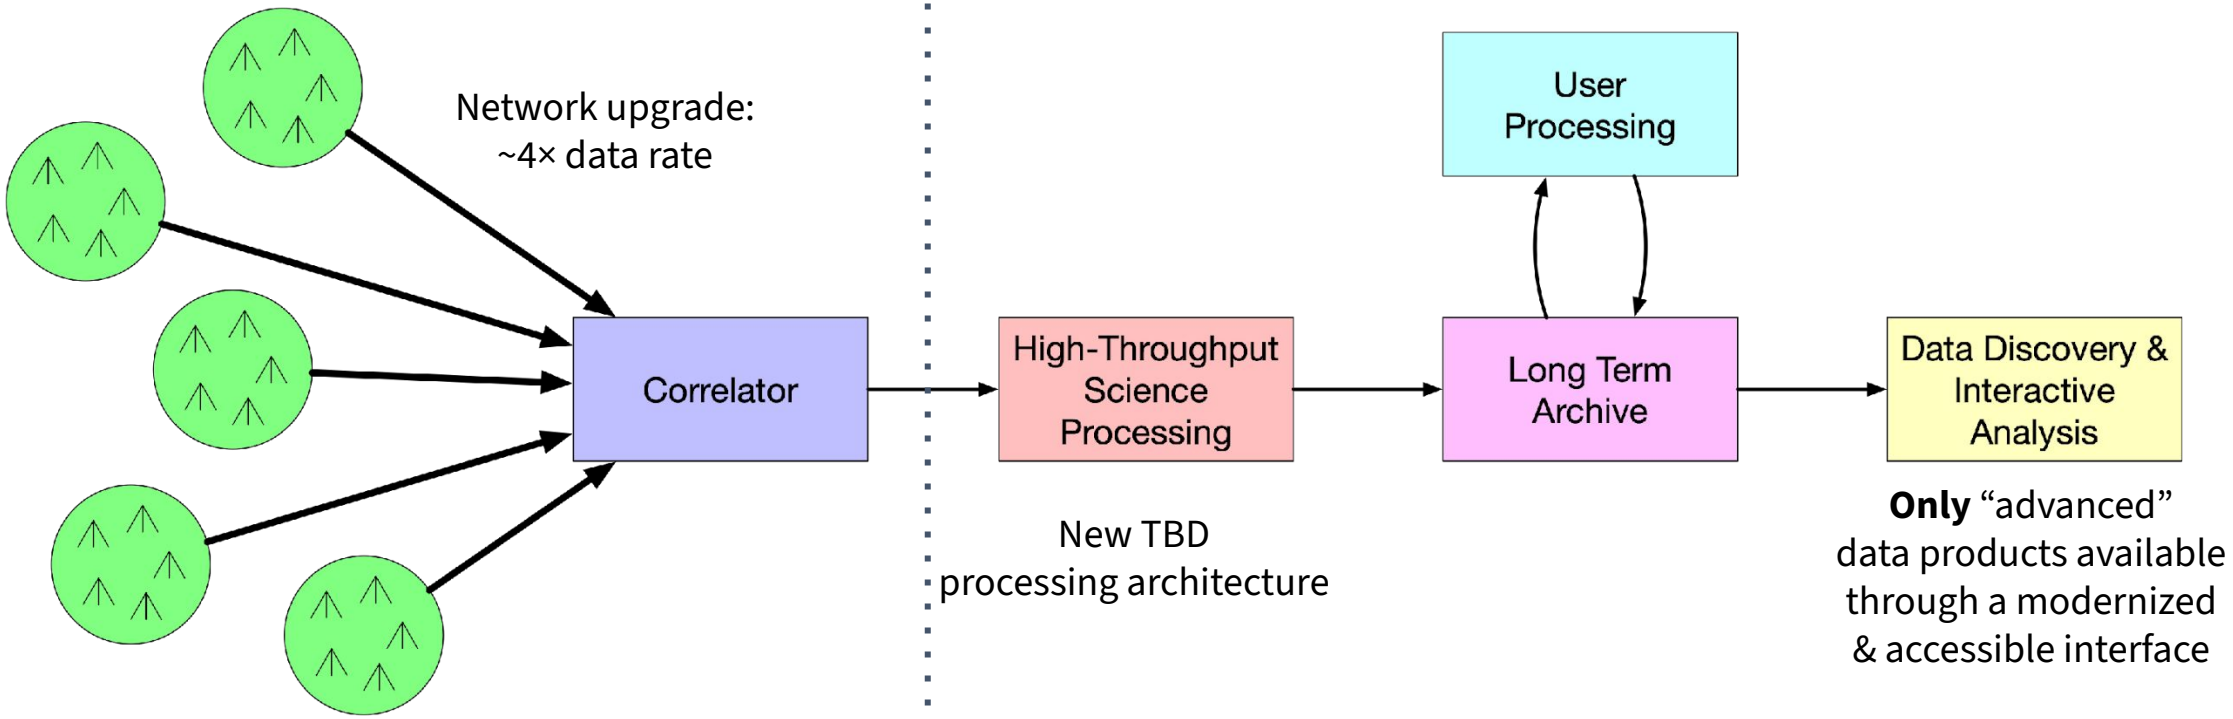


Proposal Management
Archiving & Curation
Scientific Pipelines
ASTRON Science Data Centre
Execution Services
Discovery & Access
Interactive Data Analysis
User Pipeline Execution

Part Three

Archive system

(Hanno’s slides)

Part Four

Data management system

(Hanno’s slides)

Part Five

Data processing system

(Hanno’s slides)

Part Six

Upcoming goals:

- Supporting LOFAR2
- Aligning with SRCNet

Early LOFAR2 Operations

- By the end of 2025 (“the MPS”), SDC functionality encompasses:
  - Observing proposal management (using a new, more flexible and maintainable tool than LOFAR1).
  - Overview of storage/compute resource allocation and use.
  - Execution and management of the LINC pipeline at SURF and Jülich.
  - Ingest of a variety of advanced and ancillary data products to the LOFAR Archive (LTA).
  - Data discovery and access through existing interfaces.

- During 2026, the following functionality is high priority:
  - Establish sustainable and efficient operations.
  - Extend data processing to Poznań.
  - Support a wider range of advanced data products, ultimately including all data produced by large programmes.
  - The ability for external users / science teams to add (supported) data products to the archive (“user data ingest”).
  - Integration of further pipelines (e.g. Rapthor, VLBI, TraP).

- *After this, further work will focus on development in alignment with SKA SRCNet.*

SKA Regional Centre Network

- The SKA Regional Centre Network is the distributed network of data centres being created to support the SKA.
- It will offer functionality equivalent to our ambitions for the LOFAR data system, including “science platform” functionality (advanced data discovery, visualization tools, Jupyter notebooks, user-defined pipelines, etc).
- The SRCNet Project has adopted a broadly “clean sheet” approach to development, building a new system from scratch rather than embracing and extending existing services like LOFAR LTA, ASKAP CASDA, etc.
- Development is undertaken by the SKA partners at the level of around 50 FTE.
- The SRCNet Development Project is nominally scheduled for completion in 2028, but in practice will undoubtedly extend longer than this...
  - ...but does have hard milestones to meet to support SKA Science Verification over the coming years.
- ASTRON has historically (since ~2020) been a major contributor to SRCNet development, but put this work on hold during 2025 to focus on the development of the LOFAR2.0 MPS.
  - Janneke de Boer, an ASTRON employee, remained as SRCNet Project Manager.

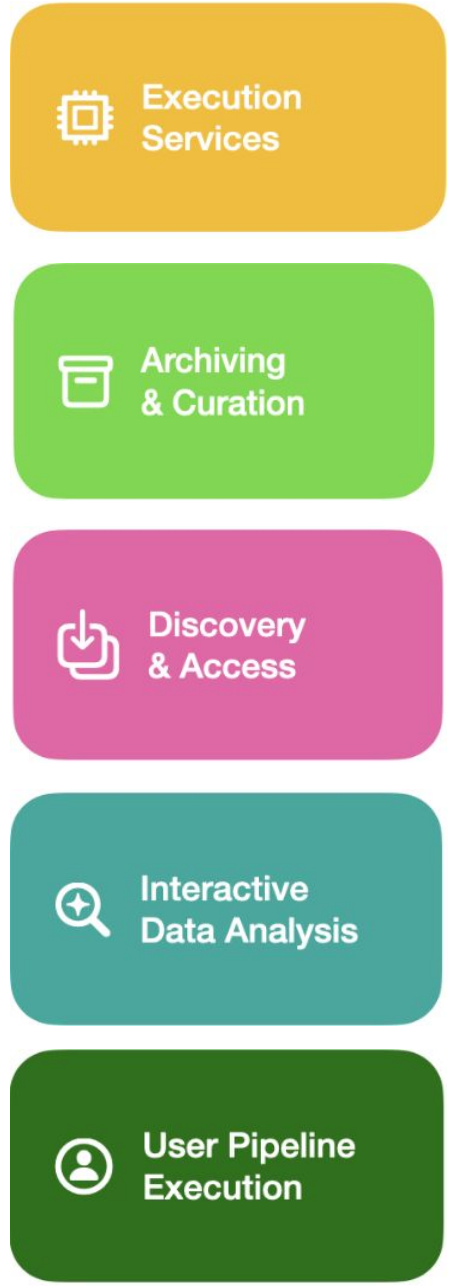

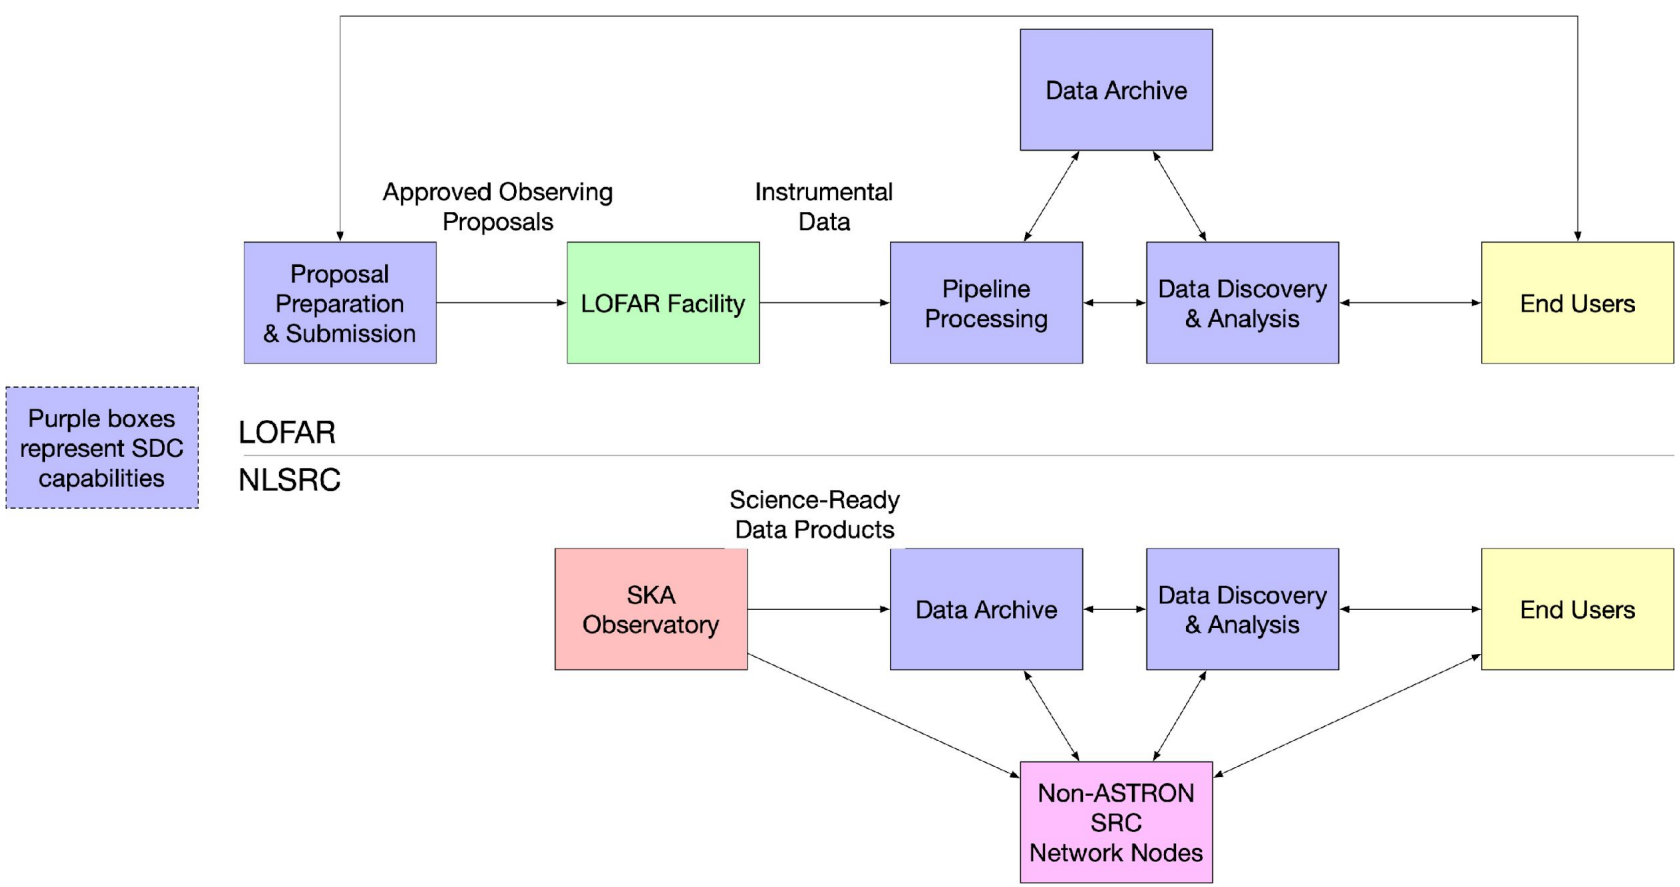

- SRCNet generates science-ready (“observatory”) data products within the telescope and supplies them to SRCNet.
- LOFAR generates lower level (“instrumental”) data products within the telescope and supplies them to the LOFAR Data Services.

- SRCNet is not involved in management of observing proposals.
- The LOFAR Data Services include proposal preparation, submission & management.

- SRCNet is expected to deploy version 1.0 in 2028 (ish)
- LOFAR Data Services need to be ready to support LOFAR2.0 operations in 2026, but development will continue

Execution Services

Archiving & Curation

Discovery & Access

Interactive Data Analysis

User Pipeline Execution

SRCNet & LOFAR Complementarity

- SRCNet is building a system that is aligned to our LOFAR ambitions.
  - Comparatively well resourced (many more FTEs than ASTRON can muster internally).
  - Opportunity to deploy new technologies without managing legacy systems.
- SRCNet is carrying a lot of risk.
  - Clean sheet approach risks repeating mistakes and needlessly reimplementing working systems.
  - Lack of operational experience and test datasets.
- ASTRON will ultimately be responsible for providing a node in the SRCNet network in addition to our responsibilities to LOFAR and other systems.
- *We see aligning LOFAR and SRCNet development as an ideal opportunity.*

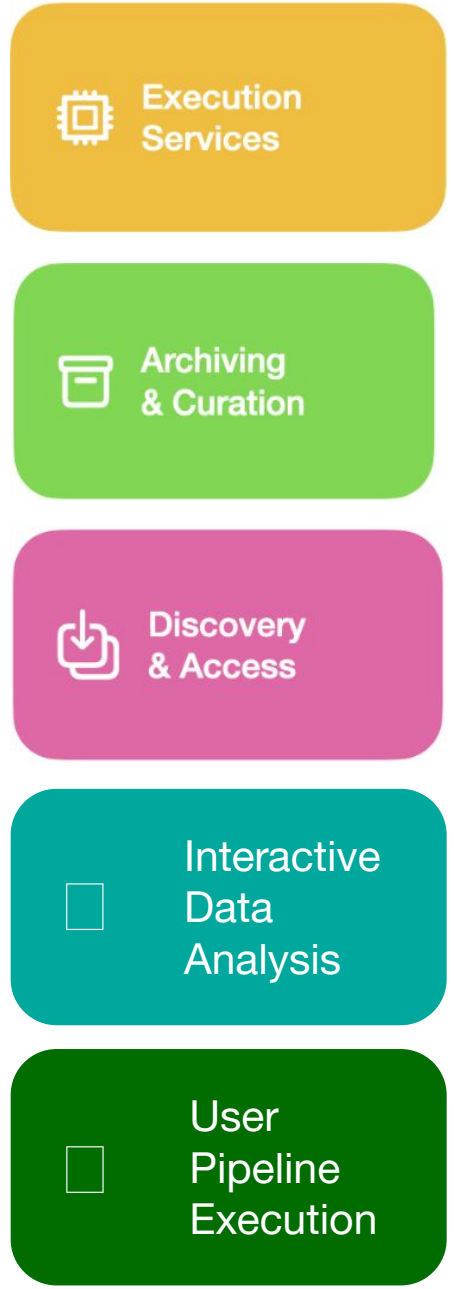

SRCNet & LOFAR Convergence

- During 2026 and beyond, we will re-engage with the SRCNet development process, gradually ramping up our contribution to several FTE.
- Provide support and expertise to the SRCNet Project on testing at-scale using LOFAR data.
- Begin internal efforts to align and adapt applicable technologies for use with LOFAR, and vice versa.
- Our ultimate goal is to arrive at a unified technology stack supporting both the LOFAR Data System and the SRC Network before the end of this decade. This new system will:
  - Provide a range of capabilities (e.g. science platform) that the ASTRON/LOFAR team are not resourced to deliver independently.
  - Modernize the LOFAR infrastructure, bringing in new and more flexible technologies around archive management and data placement (e.g. the Rucio policy-based system).
  - Ensure that end-users can interact with both LOFAR and SKA/SRCNet in a common environment.
  - Make it easier and more appealing for countries supporting *both* LOFAR *and* SKA to roll out compatible data services.
  - Provide scope for expansion to other facilities, e.g. Apertif & e-MERLIN, in addition to LOFAR and SKA.
- This vision has been discussed and agreed with partners including SKAO, UK-SRC, FR-SKA, and the LOFAR archive site at Forschungszentrum Jülich.

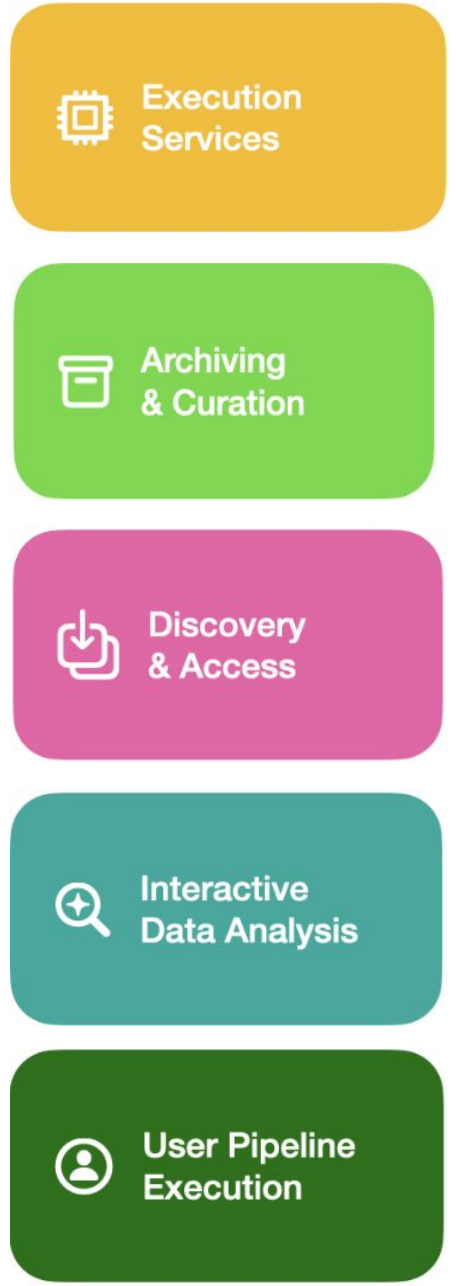

The LOFAR Leftovers

- SRCNet scope does not fully match the services needed for LOFAR. In particular:
  - SRCNet provides no functionality for managing observing or computing proposals (in the SKA model, this is embedded in the telescope rather than the Regional Centres).
  - The LOFAR telescope has its own control system (TMSS) and associated data transmission, etc, interfaces that are not the same as SKA's.
- The delivery of the TULP proposal tool for LOFAR2.0 represents a substantial fraction of this work.
- Future developments on the execution system—including, if applicable, adopting SRCNet's—will require further work to match TMSS.
- This work will be planned and executed internally, while looking for synergy with other observatories and development efforts where possible.

Part Three

Pipeline development

LOFAR2.0 Pipelines

Scientific Pipelines

- ASTRON is internally developing a set of “core” pipelines for LOFAR2.0:
  - Preprocessing (*runs on CEP*)
  - Initial calibration (LINC, C*EP & LOFAR Data System*)
  - Direction-dependent calibration & imaging (Rapthor-HBA, *LOFAR Data System*)
  - Image-plane transient detection & monitoring (TraP, *LOFAR Data System*)
  - Pulsar timing (PULP, *CEP*)
- The ASTRON team is closely involved in development of two further pipelines:
  - VLBI (*LOFAR Data System*)
  - LBA imaging (LiLF, *CEP & LOFAR Data System*)
- We are tracking 19 pipelines developed across the community to complete LOFAR2.0 Large Programmes (a subset of the total).

Core pipelines are primarily C++ executables strung together with CWL and Python.

	Pipeline	Main Contact	Application	Description
1	Pre-processing	Iacobelli	all imaging	Prepare data (average, compress, flag, demix)
2	LINC	Drabent	all HBA imaging	Direction independent calibration (bandpass, clock,…). HBA and LBA versions
3	DDF-pipeline	Hardcastle	6" HBA imaging	NL baseline direction dependent calibration (HBA)
4	RAPTHOR	Rafferty	6" HBA imaging	NL baseline direction dependent calibration (HBA)
5	LOFAR-VLBI pipeline	Morabito	0.3" HBA	ILT baseline direction dependent calibration
6	LiLF	De Gasperin	12" LBA imaging (for now) - int'l baselines in future	NL baseline direction dependent and independent calibration (HBA)
7	Facetselfcal	van Weeren	part of VLBI and Extract pipelines	Direction dependent and independent calibration
8	TrAP	Rowlinson	transients	Time variability
9	Extract	van Weeren	targeted sources within a field	Tailor calibration for a particular target
10	DynSpecMS	Tasse	time + freq spectra of sources; for all imaging pipelines	Time and frequency variability
11	PULP	Kondratiev	pulsar searching	Pulsar processing to get average pulse profiles for further analysis (RFI, dedispersing/folding using PRESTO and dspsr/PSRCHIVE software)
12	AARTFAAC EoR HBA	Offringa		Flagging, averaging, compression, calibration, imaging, foreground removal, power spectrum
13	Lightning	Hare	lightning	Imaging lightning propagation
14	Solar	Zucca	solar	Interferometric imaging beamformed data for dynamic spectra
15	Cosmic ray	Mulrey	cosmic ray	Flagging, calibration and e.g. event reconstruction to determine arrival direction
17	LORDS	Rajwade	Pulsar imaging searching	Image-plane didispersion
18	RMsynthesis	O'Sullivan	linear polarisation	Linear polarisation
19	PyBDSF	Rafferty	source finding	Radio source cataloguing and multi-wavelength crossmatching

Figure: Michiel van Haarlem & Tim Shimwell

Progress: DD imaging

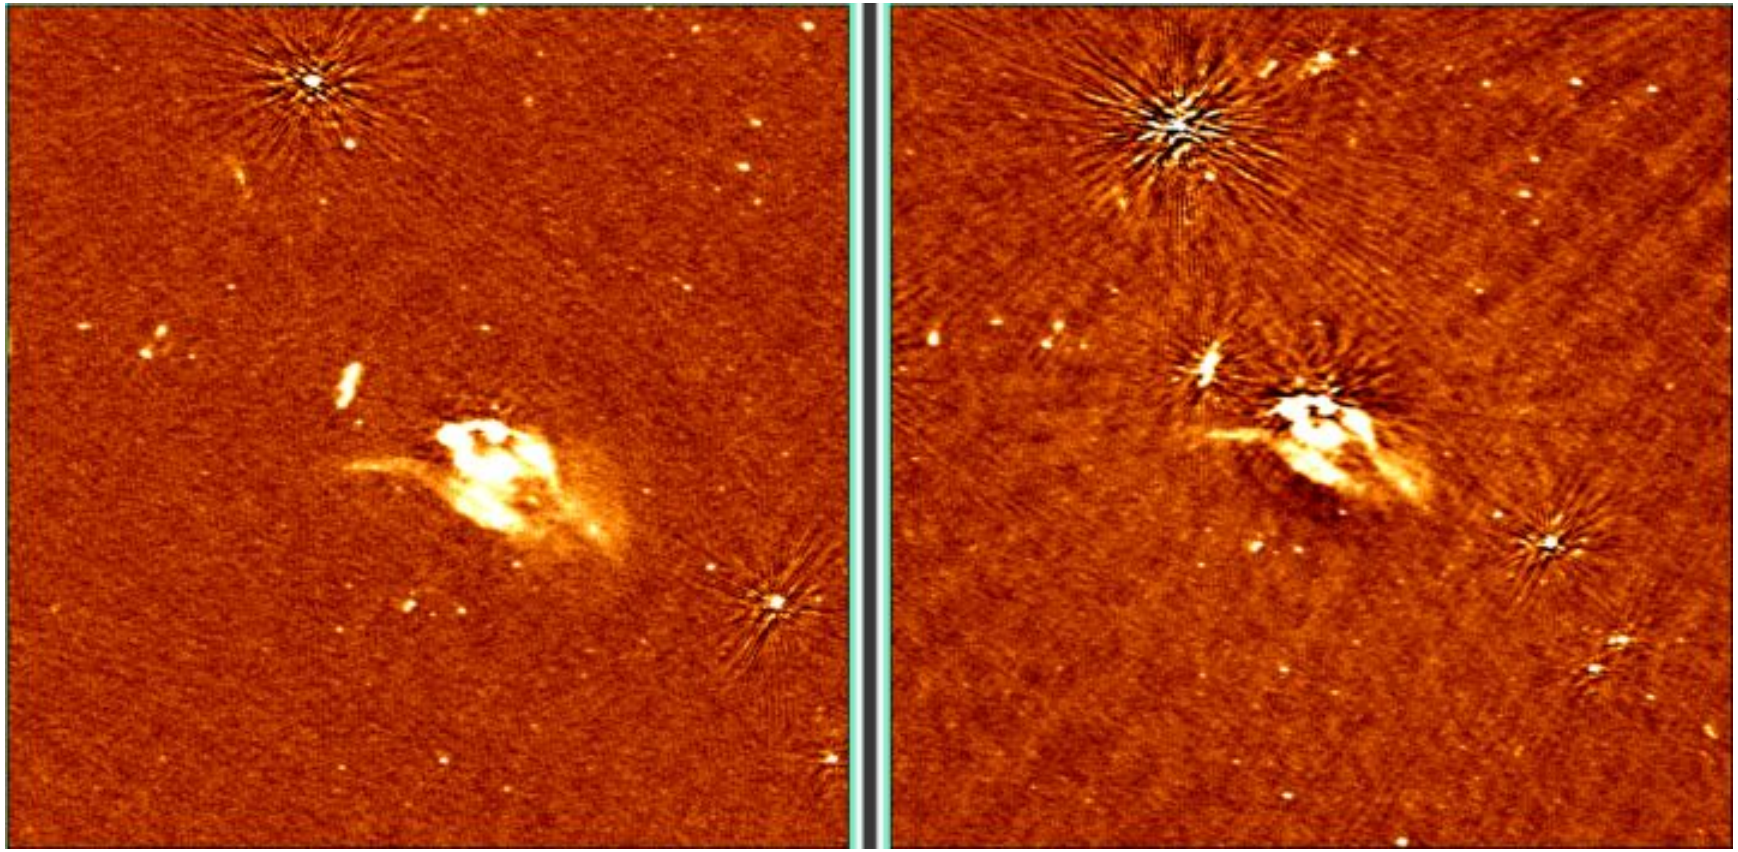

Image:
André Offringa

Rapthor 2.1 Rapthor 2.0

Rapthor 2.0 released in the spring; 2.1 before the end of this year. *Huge* quality improvements, and it's much faster too.

ASTRON
Netherlands Institute for Radio Astronomy

Supporting Community Pipelines

Scientific Pipelines

- The 19 pipelines provide capabilities that are necessary to support both Large Programmes and other users.
- LOFAR/ASTRON policy is that pipelines must (eventually) be under central management, i.e. operated by Observatory staff.
- ASTRON does not have sufficient development effort to support & maintain all of these pipelines.
- We have proposed a set of procedural and technical standards that pipelines must meet in order for them to be “observatory-supported”. This includes *external development teams taking responsibility for maintenance*.
  - Still under discussion with LOFAR ERIC.

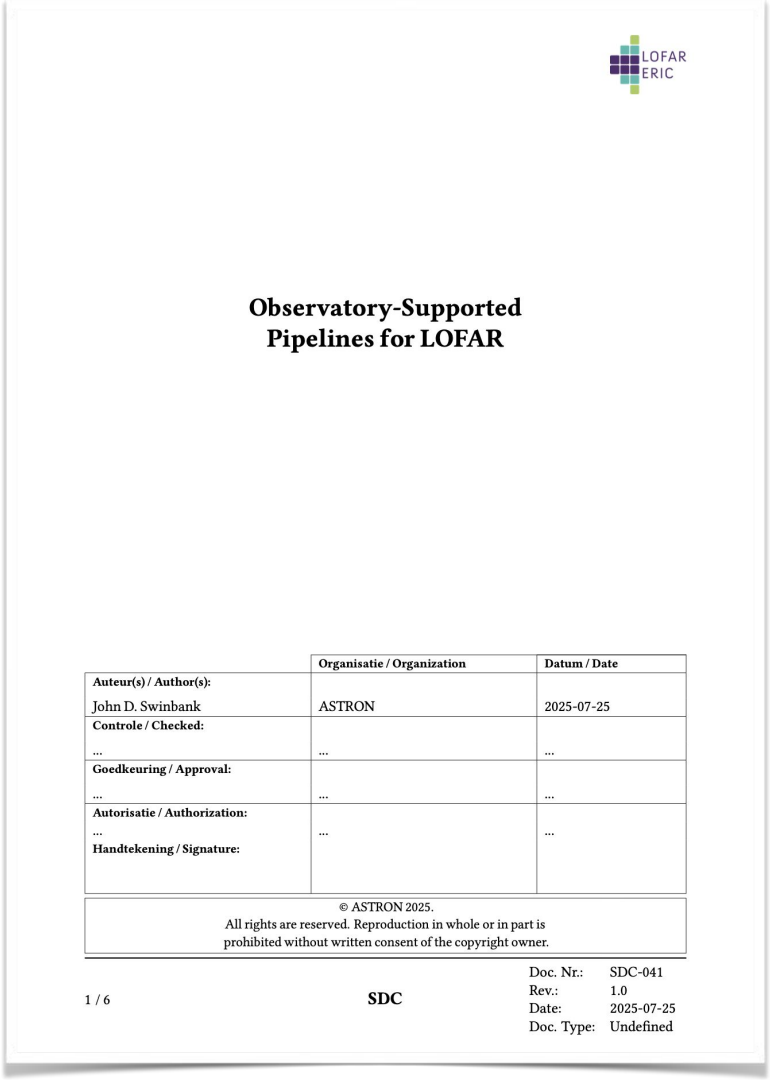

LOFAR ERIC

Observatory-Supported Pipelines for LOFAR

	Organisatie / Organization	Datum / Date
Auteur(s) / Author(s): John D. Swinbank	ASTRON	2025-07-25
Controle / Checked: ...	...	...
Goedkeuring / Approval: ...	...	...
Autorisatie / Authorization: ... **Handtekening / Signature:**	...	...

Future Pipelines: LENSS

Scientific Pipelines

- The LENSS project will provide a wide-field (4× current field of view), high-resolution (European baselines), image-based survey system based on LOFAR. It will become operational in the early 2030s.
- Development of the pipelines for LENSS will build on current efforts within ASTRON and the LOFAR community (LINC, Rapthor, VLBI).
- The data volumes involved mean we must transition from the current system of storing visibilities for later “offline” processing to a system where data is processed as it is taken.
  - This reduces time-to-science and the size of buffer required.
  - Effectively, we match the throughput (but not necessarily the latency) of the data processing system to that of the telescope.
- This implies a pipeline performance increase of ~40× over 2024 VLBI performance.
  - This is plausible due to a) availability of accelerator hardware (GPUs), and b) ongoing improvements to pipeline algorithms.
  - Representing our commitment to sustainability, this pipelines will execute within the same overall energy budget as LOFAR2 processing.
  - This work will be a major focus of research and development for the pipeline teams inside and outside ASTRON over the next several years.
- Development starting during 2026 as a close collaboration between ASTRON and the (community-based) LOFAR VLBI Working Group.

Future Pipelines: SKA-SDP

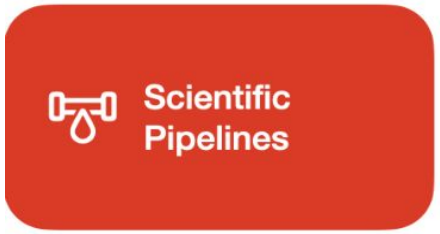

- The institute hosts a team (“Gecko”) working on ICAL pipeline (DD calibration & imaging) for the SKA Science Data Processor.
  - Team Gecko is embedded within the SKA-SDP Agile Release Train.
- The LOFAR and SKA pipelines aren't exactly the same (LOFAR has lower frequencies and longer baselines, SKA has *many* more baselines), but there is a lot of overlap.
- Team Gecko is part of ASTRON's Smart Backend competence group, and works closely with LOFAR pipeline developers to minimise duplication of effort.
- However...

Pipeline Convergence

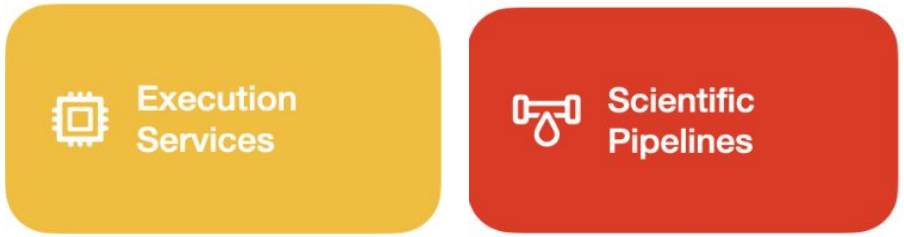

- ...we are unhappy with the current situation for (broadly) two reasons:
  - Despite collaboration within Smart Backend, there are still two distinct projects with different technology stacks and success criteria.
  - End-users who want to interact with both LOFAR and SKA/SRCNet will have to learn to use and/or develop for two different systems.
- We therefore regard it as a priority to converge these systems as far as possible.
  - Same algorithmic components.
  - Same execution environment and systems.
  - Same development/release/test processes.
- SKAO-ASTRON MoU signed in support of this goal in September.
- Initial planning meeting between SKAO and ASTRON teams on 5 November identified first set of development goals.

Part Four

SWOT analysis

Strengths

- Established, scientifically productive Observatory.
- Operational, “battle-tested”, distributed archive.
- Experienced operations team.
- In-house expertise developing pipelines for both LOFAR and SKA.
  - Sharing technology wherever possible.
- Extensive in-house research portfolio in both astrophysics and technology development.
- Engaged and scientifically productive community, often with advanced data processing knowledge.

Weaknesses

- Running at full stretch to meet basic commitments to LOFAR, SKA, etc; struggle to find space to innovate.
- Commitments (both from ASTRON and other LTA partners) to LOFAR ERIC can be unclear.
- Aligning development between LOFAR and SKA adds risk as well as making more efficient use of resources.
- Some of the technologies the archive is based on are 15+ years old, and are showing their age.
- Lack of automation means labour-intensive operations.
- Including ERIC partners in decision making can slow things down.

Opportunities

- Alignment between LOFAR and SKA is potentially a huge force multiplier.
- The LOFAR community is a big asset in developing new functionality (if they can be effectively harnessed).
- LTA partners (SURF, Jülich, Poznan) bring extensive expertise and experience.
- Future LOFAR expansions (LENSS, LOFAR-Low, etc) offer great opportunities.
- Even in the era of SKA, LOFAR has longer baselines, lower frequencies, and a different hemisphere.

Threats

- Costs of data storage and processing becoming prohibitive.
- Environmental sustainability of big compute/storage systems.
- Scientific community sceptical of “black box” data processing.
- Community serves up lots of “stuff” they’d like to see integrated.