PROCESS

PROviding Computing solutions for ExaScale ChallengeS

# Unlocking the LOFAR LTA

Hanno Spreeuw, Souley Madougou, Ronald van Haren,

Berend Weel, Adam Belloum, Jason Maassen

24th Sep 2019, eScience conference – EINFRA, San Diego

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 777533.

# LOFAR: Low Frequency Array

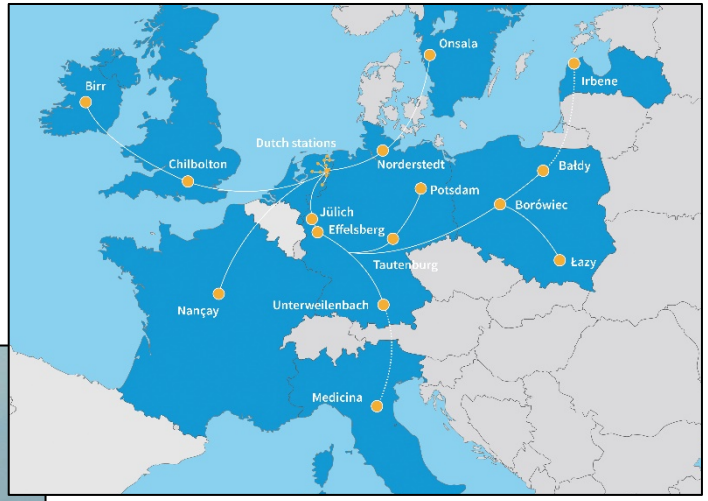


Images courtesy of: ASTRON

LOFAR radio telescope – is a “distributed software telescope” consisting of ~88.000 antennas in ~51 stations scattered over Europe. It produces up to **1.6 TB/s of raw data**, processed real time to combine the signals of multiple antennas (correlation). **This results in up to 35 TB/h of intermediate data** (visibilities) which is stored for further analysis.

# LOFAR: Key Science Projects

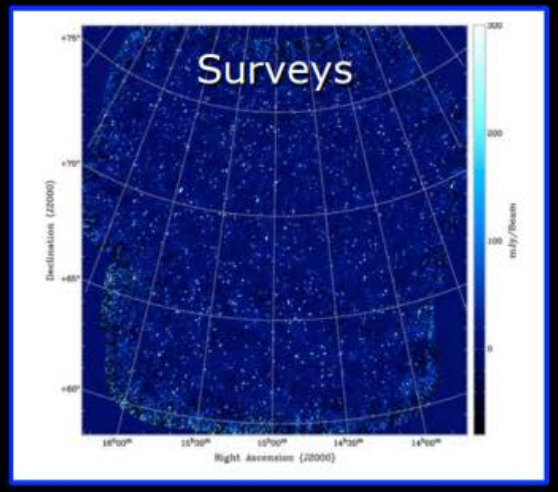


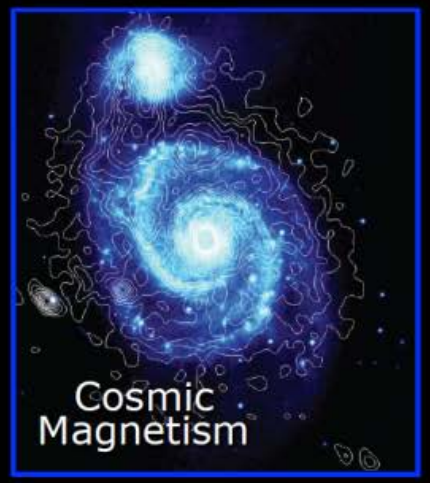


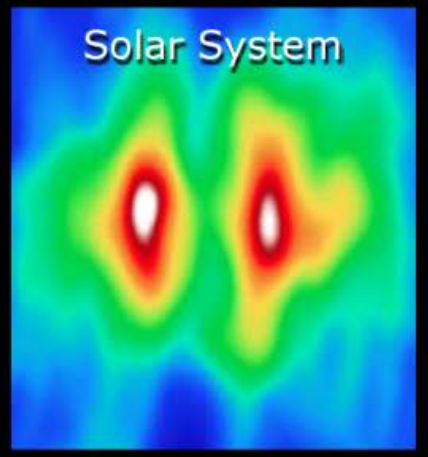


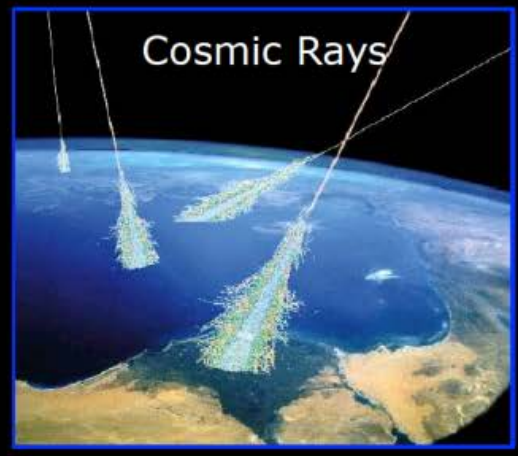



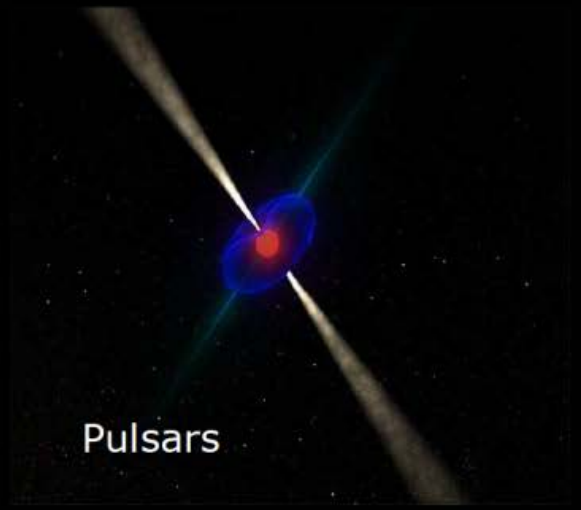


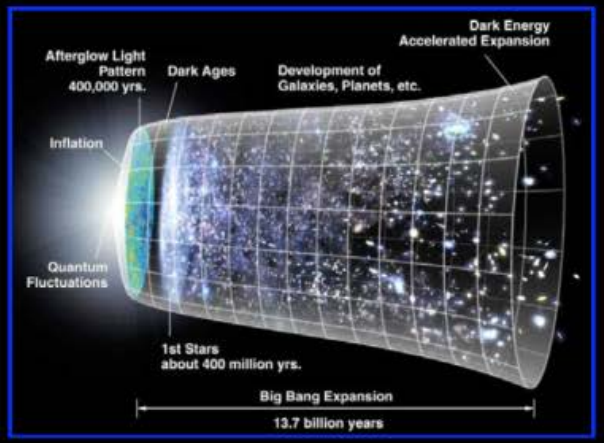


Epoch of Reionization

Slide courtesy of Roberto Pizzo (ASTRON). http://www.tauceti.caltech.edu/science-at-low-frequencies-2016/slides/A/02_Pizzo.pdf

# LOFAR LTA: Long Term Archive

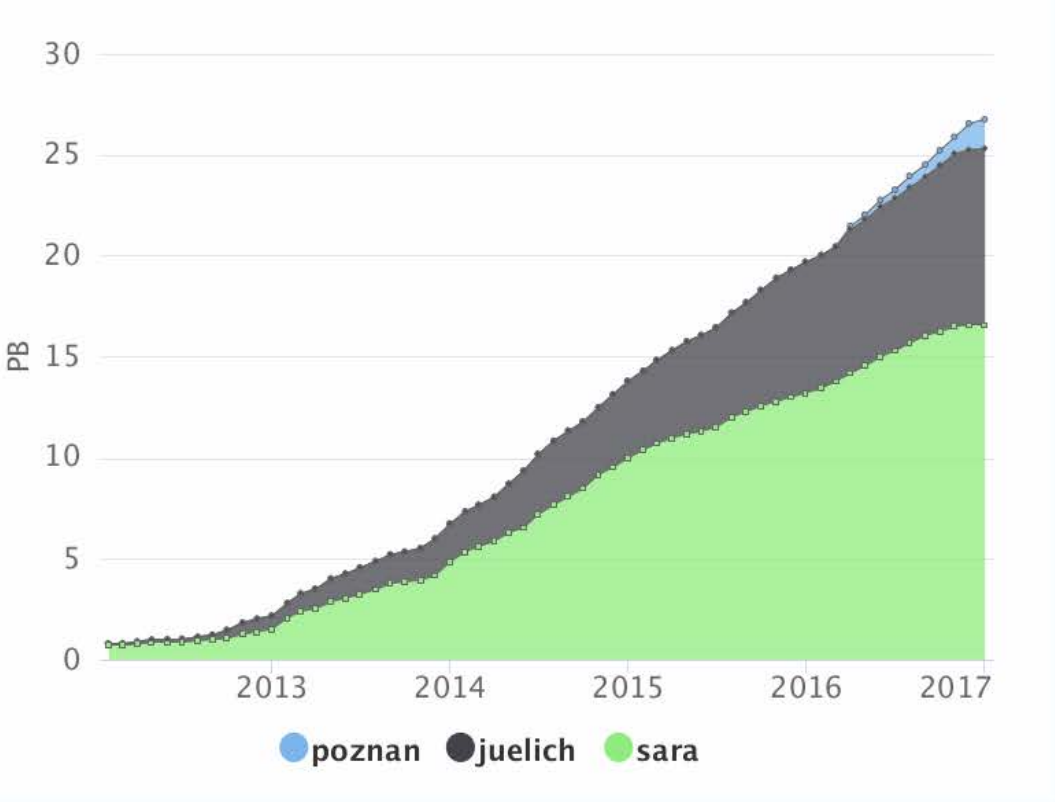


LOFAR Long Term Archive

LOGIN

HOME SEARCH DATA BROWSE PROJECTS HELP

Observation 1 to 100 (showing 100 of total 32940)

edit columns

first previous 1 2 3 4 5 6 7 8 9 10 11 ... next last

| # | Project | Release Date | Instrument Filter | Channel Width [MHz] | Bits per Sample | Number Of SubArray Pointings | Duration [s] | Nr Stations Core | Nr Stations Remote | Nr Stations International | Number Of Stations | Number Of Correlated DataProducts | Number Of BeamFormed DataProducts | Creation Date |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | LC10_024 | | 110-190 MHz | 0.000000 | 8 | 1 | 951.0 | 24 | 14 | 12 | 50 | 243 | 0 | 2018-10-26 11:56:44 |
| 2 | LC10_024 | | 110-190 MHz | 0.000000 | 8 | 3 | 14760.0 | 24 | 14 | 12 | 50 | 486 / 487 | 0 | 2018-10-26 10:07:54 |
| 3 | LC10_024 | | 110-190 MHz | 0.000000 | 8 | 1 | 951.0 | 24 | 14 | 12 | 50 | 243 | 0 | 2018-10-26 08:59:42 |
| 4 | LT10_010 | | 110-190 MHz | 0.000000 | 8 | 1 | 950.0 | 24 | 14 | 12 | 50 | 243 | 0 | 2018-10-25 13:23:14 |
| 5 | LT10_010 | | 110-190 MHz | 0.000000 | 8 | 3 | 29170.0 | 24 | 14 | 12 | 50 | 486 / 487 | 0 | 2018-10-25 10:53:43 |
| 6 | LT10_010 | | 110-190 MHz | 0.000000 | 8 | 1 | 951.0 | 24 | 14 | 12 | 50 | 243 | 0 | 2018-10-25 10:28:31 |
| 7 | LC10_018 | | 110-190 MHz | 0.000000 | 8 | 1 | 960.0 | 24 | 14 | 11 | 49 | 244 | 0 | 2018-10-25 08:13:41 |
| 8 | LC10_018 | | 110-190 MHz | 0.000000 | 8 | 7 | 29180.0 | 24 | 14 | 11 | 49 | 388 | 0 | 2018-10-25 06:27:40 |
| 9 | LC10_018 | | 110-190 MHz | 0.000000 | 8 | 1 | 951.0 | 24 | 14 | 11 | 49 | 244 | 0 | 2018-10-25 05:39:15 |
| 10 | DDT10_006 | 2019-10-22 | 110-190 MHz | 0.000000 | 8 | 1 | 3161.0 | 23 | 0 | 0 | 23 | 0 | 240 | 2018-10-26 17:12:24 |
| 11 | DDT10_006 | 2019-10-22 | 110-190 MHz | 0.000000 | 8 | 1 | 641.0 | 23 | 0 | 0 | 23 | 0 | 240 | 2018-10-26 14:09:42 |
| 12 | LC10_015 | | 110-190 MHz | 0.000000 | 8 | 1 | 2141.0 | 21 | 0 | 0 | 21 | 0 | 1 | 2018-10-25 09:50:23 |
| 13 | LC10_007 | | 30-90 MHz | 0.000000 | 8 | 2 | 29170.0 | 24 | 13 | 0 | 37 | 488 | 0 | 2018-10-22 08:06:45 |
| 14 | DDT10_006 | 2019-10-22 | 110-190 MHz | 0.000000 | 8 | 1 | 641.0 | 23 | 0 | 0 | 23 | 0 | 240 | 2018-10-26 16:33:00 |
| 15 | DDT10_006 | 2019-10-22 | 110-190 MHz | 0.000000 | 8 | 1 | 3161.0 | 23 | 0 | 0 | 23 | 0 | 240 | 2018-10-26 15:55:27 |
| 16 | DDT10_006 | 2019-10-22 | 110-190 MHz | 0.000000 | 8 | 1 | 640.0 | 23 | 0 | 0 | 23 | 0 | 240 | 2018-10-26 14:07:05 |

LOFAR observations are stored in the **long term archive (LTA)** which is distributed over **Amsterdam**, **Juelich** and **Poznan**.
It currently contains **~43 PB of data** and **grows with 5 to 7 PB/year**. Data is stored **mostly on tape**, using hard disks for temporary storage (dCache). A web interface is available to request data to download. Processing needs to be done by the astronomers themselves on local infrastructure → **this is the hard part**!

# LOFAR LTA: Long Term Archive

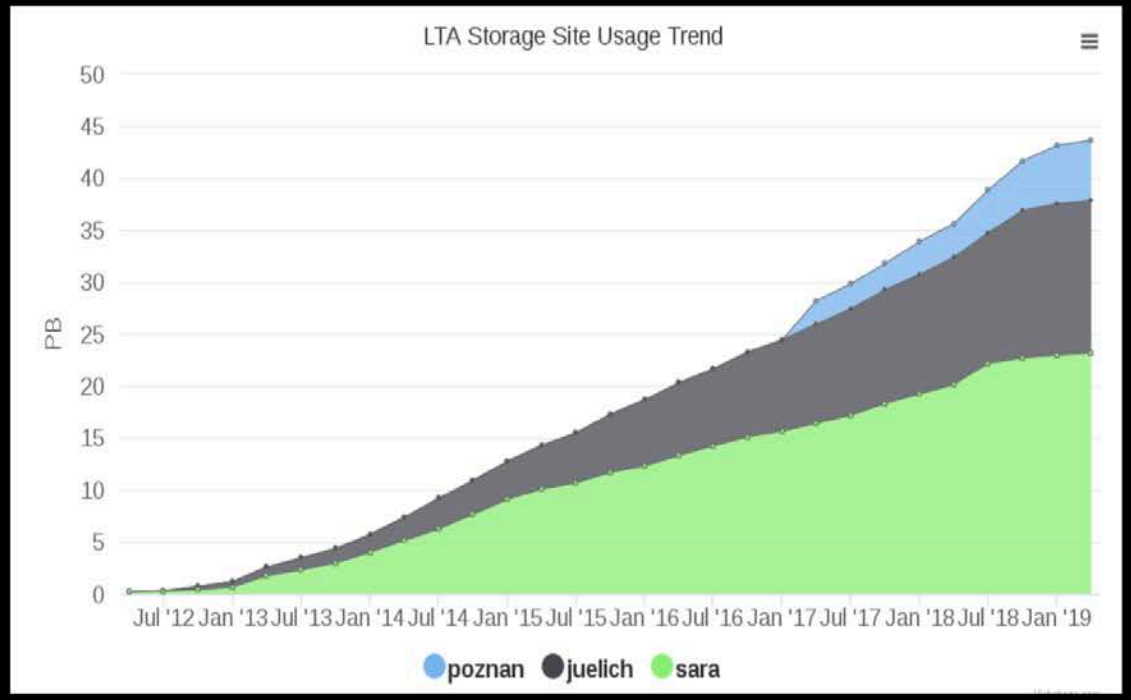


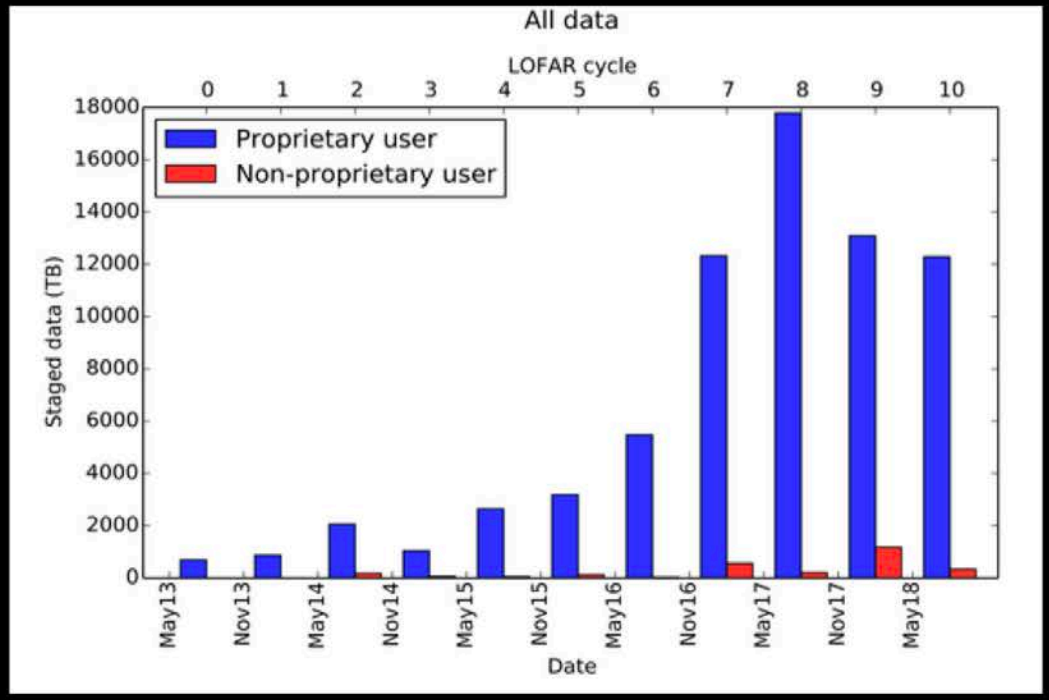


- Manage data growth: compression tool (Dysco) in production since September 2018
- More than a factor of 10 between proprietary vs non-proprietary data downloads: 60 PB vs 5 PB
- Need to make (science-ready) data accessible: LOFAR efficiency + advanced pipelines in production + exposing data

Slide courtesy of Roberto Pizzo (ASTRON). https://www.astron.nl/lofarscience2019/Documents/Monday/LUM_Pizzo.pdf

# Problem: Big Data & Complex Software

The data size for a single observation tend to **be large**, **10's of TBs**, which are non-trivial to download and handle.

There is a lot of software available to process LOFAR observations. Usually several packages are needed for the different aspects of calibration and imaging. These often **require expert knowledge to install and operate**.

The combination of these two makes it **hard** for **astronomers** to make optimal use of the archive.

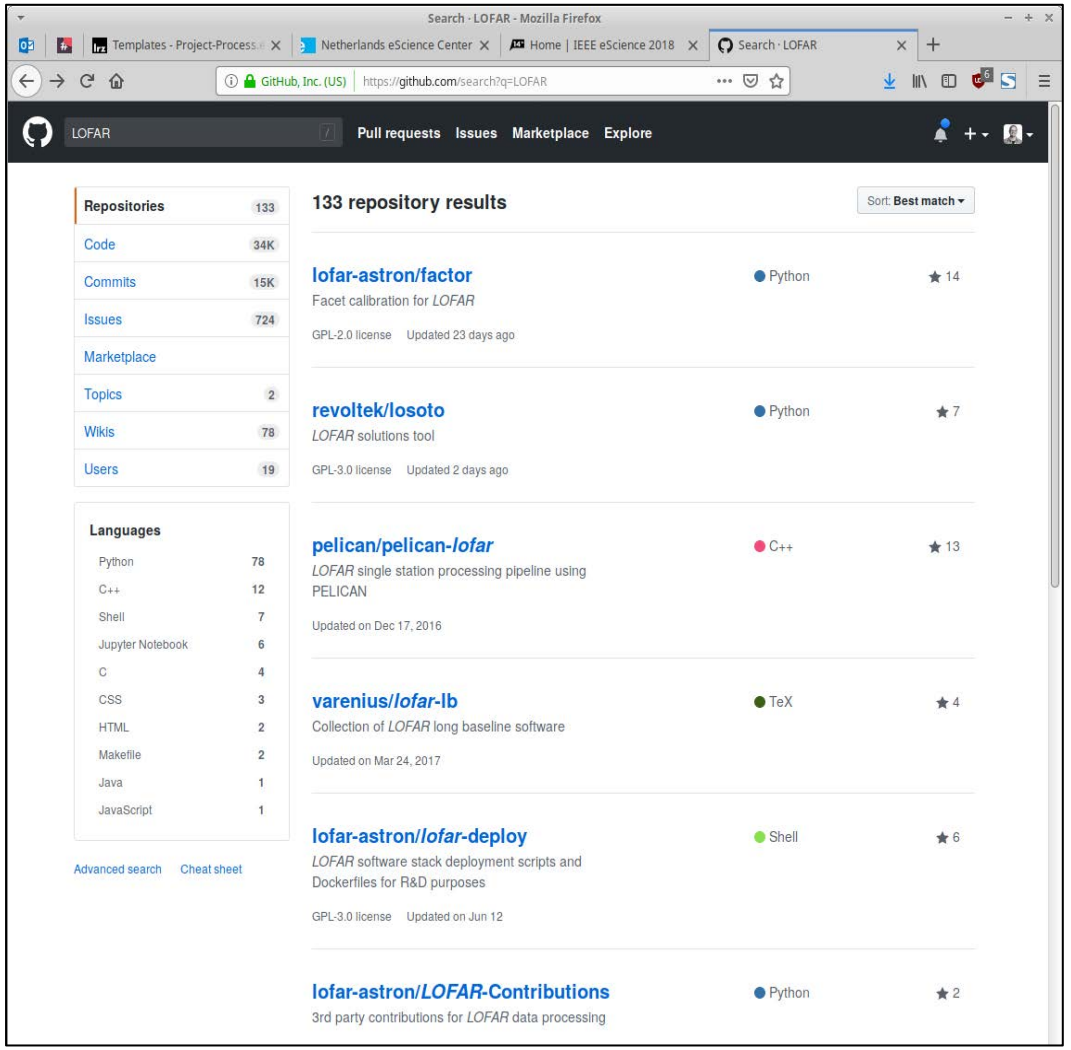

# Goals of PROCESS

**Mature, modular, generalizable open source solutions for user friendly exascale data.**

**UC#1: Exascale learning on medical image data**

**UC#2: Analysis of radio astronomy observations**

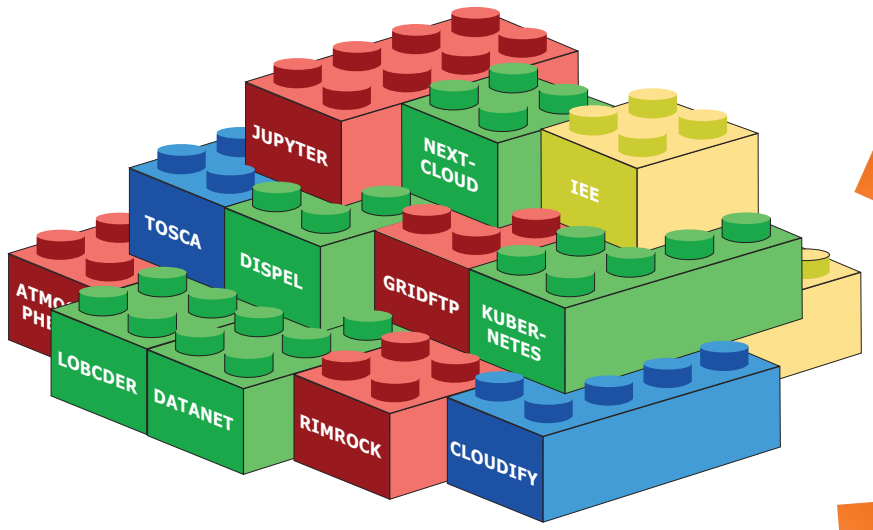


**UC#3: Supporting innovation based on global disaster risk data**

**UC#4: Ancillary pricing/airline revenue management**

# Reusable PROCESS services

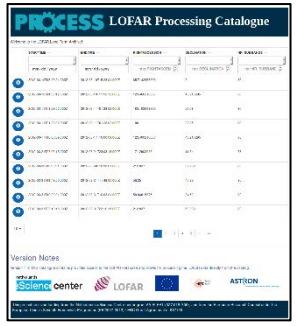


START:
trigger staging
and processing

Astronomer

END:
evaluate
images

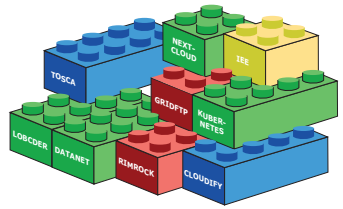


Output Images

~1.2 GB x 244

Tape Archive 28 PB

16 TB

dCache

Step 1: Staging

~1.2GB x 244

pre-process 1
pre-process 2
pre-process 3
pre-process 244

~100MB x 244

calibrate 1
calibrate 2
calibrate 3
calibrate 244

~2MB x 244

Calibrator 1

combine

~100MB

process solutions

~120MB

Calibrator 2

~60 GB x 244

apply 1
apply 2
apply 3
apply 244

average 1
average 2
average 3
average 244

~1.2 GB x 244

~60GB x 244

Target 1

combine 1
combine 10

~12 GB x 24

calibrate 1
calibrate 10

Target 2

Step 2: Direction Independent Pipeline

~12 GB x 24

Imaging and Calibration

Step 3: Direction Dependent Pipeline

# Step 1: Stage in the data

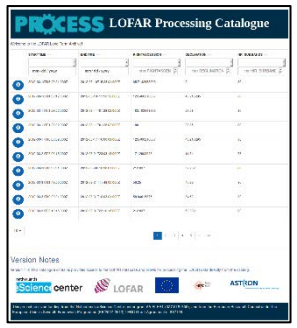


START:
trigger staging
and processing

Astronomer

END:
evaluate
images

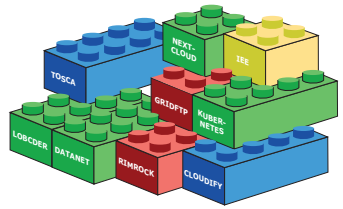


Output Images

~1.2 GB x 244

Data Service

Tape Archive 28 PB

16 TB

dCache

Step 1: Staging

~1.2GB x 244

pre-process 1 → ~100MB x 244 → calibrate 1
pre-process 2 → calibrate 2
pre-process 3 → calibrate 3
pre-process 244 → calibrate 244

~2MB x 244

Calibrator 1

combine → ~100MB → process solutions

~120MB

Calibrator 2

~60 GB x 244

apply 1 → average 1
apply 2 → average 2
apply 3 → average 3
apply 244 → average 244

~1.2 GB x 244

~60GB x 244

Target 1

combine 1 → ~12 GB x 24 → calibrate 1
combine 10 → calibrate 10

Target 2

Step 2: Direction Independent Pipeline

~12 GB x 24

Imaging and Calibration

Step 3: Direction Dependent Pipeline

# Step 2: Launch the processing workflow

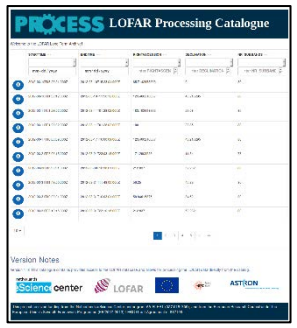


START:
trigger staging
and processing

Astronomer

END:
evaluate
images

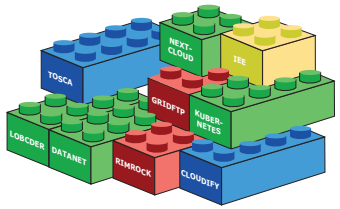


Output Images

~1.2 GB x 244

Data Service

Tape Archive 28 PB

16 TB

dCache

Step 1: Staging

~1.2GB x 244

pre-process 1
pre-process 2
pre-process 3
pre-process 244

~100MB x 244

calibrate 1
calibrate 2
calibrate 3
calibrate 244

~2MB x 244

Calibrator 1

combine

~100MB

process solutions

~120MB

Calibrator 2

~60 GB x 244

apply 1
apply 2
apply 3
apply 244

average 1
average 2
average 3
average 244

~1.2 GB x 244

~60GB x 244

Target 1

combine 1
combine 10

~12 GB x 24

calibrate 1
calibrate 10

Target 2

Step 2: Direction Independent Pipeline

~12 GB x 24

Imaging and Calibration

Step 3: Direction Dependent Pipeline

Compute Service

# Step 3: Stage out results

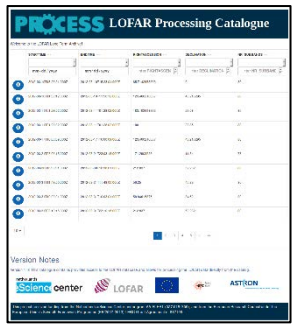


START:
trigger staging
and processing

Astronomer

END:
evaluate
images

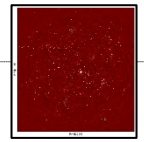

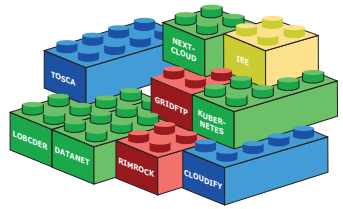


Output Images

~1.2 GB x 244

Data Service

Data Service

Tape Archive 28 PB

16 TB

dCache

Step 1: Staging

~1.2GB x 244

pre-process 1
pre-process 2
pre-process 3
pre-process 244

~100MB x 244

calibrate 1
calibrate 2
calibrate 3
calibrate 244

~2MB x 244

Calibrator 1

combine

~100MB

process solutions

~120MB

Calibrator 2

~60 GB x 244

apply 1
apply 2
apply 3
apply 244

average 1
average 2
average 3
average 244

~1.2 GB x 244

~60GB x 244

Target 1

combine 1
combine 10

~12 GB x 24

calibrate 1
calibrate 10

Target 2

Step 2: Direction Independent Pipeline

~12 GB x 24

Imaging and Calibration

Step 3: Direction Dependent Pipeline

Compute Service

# Effect of Processing

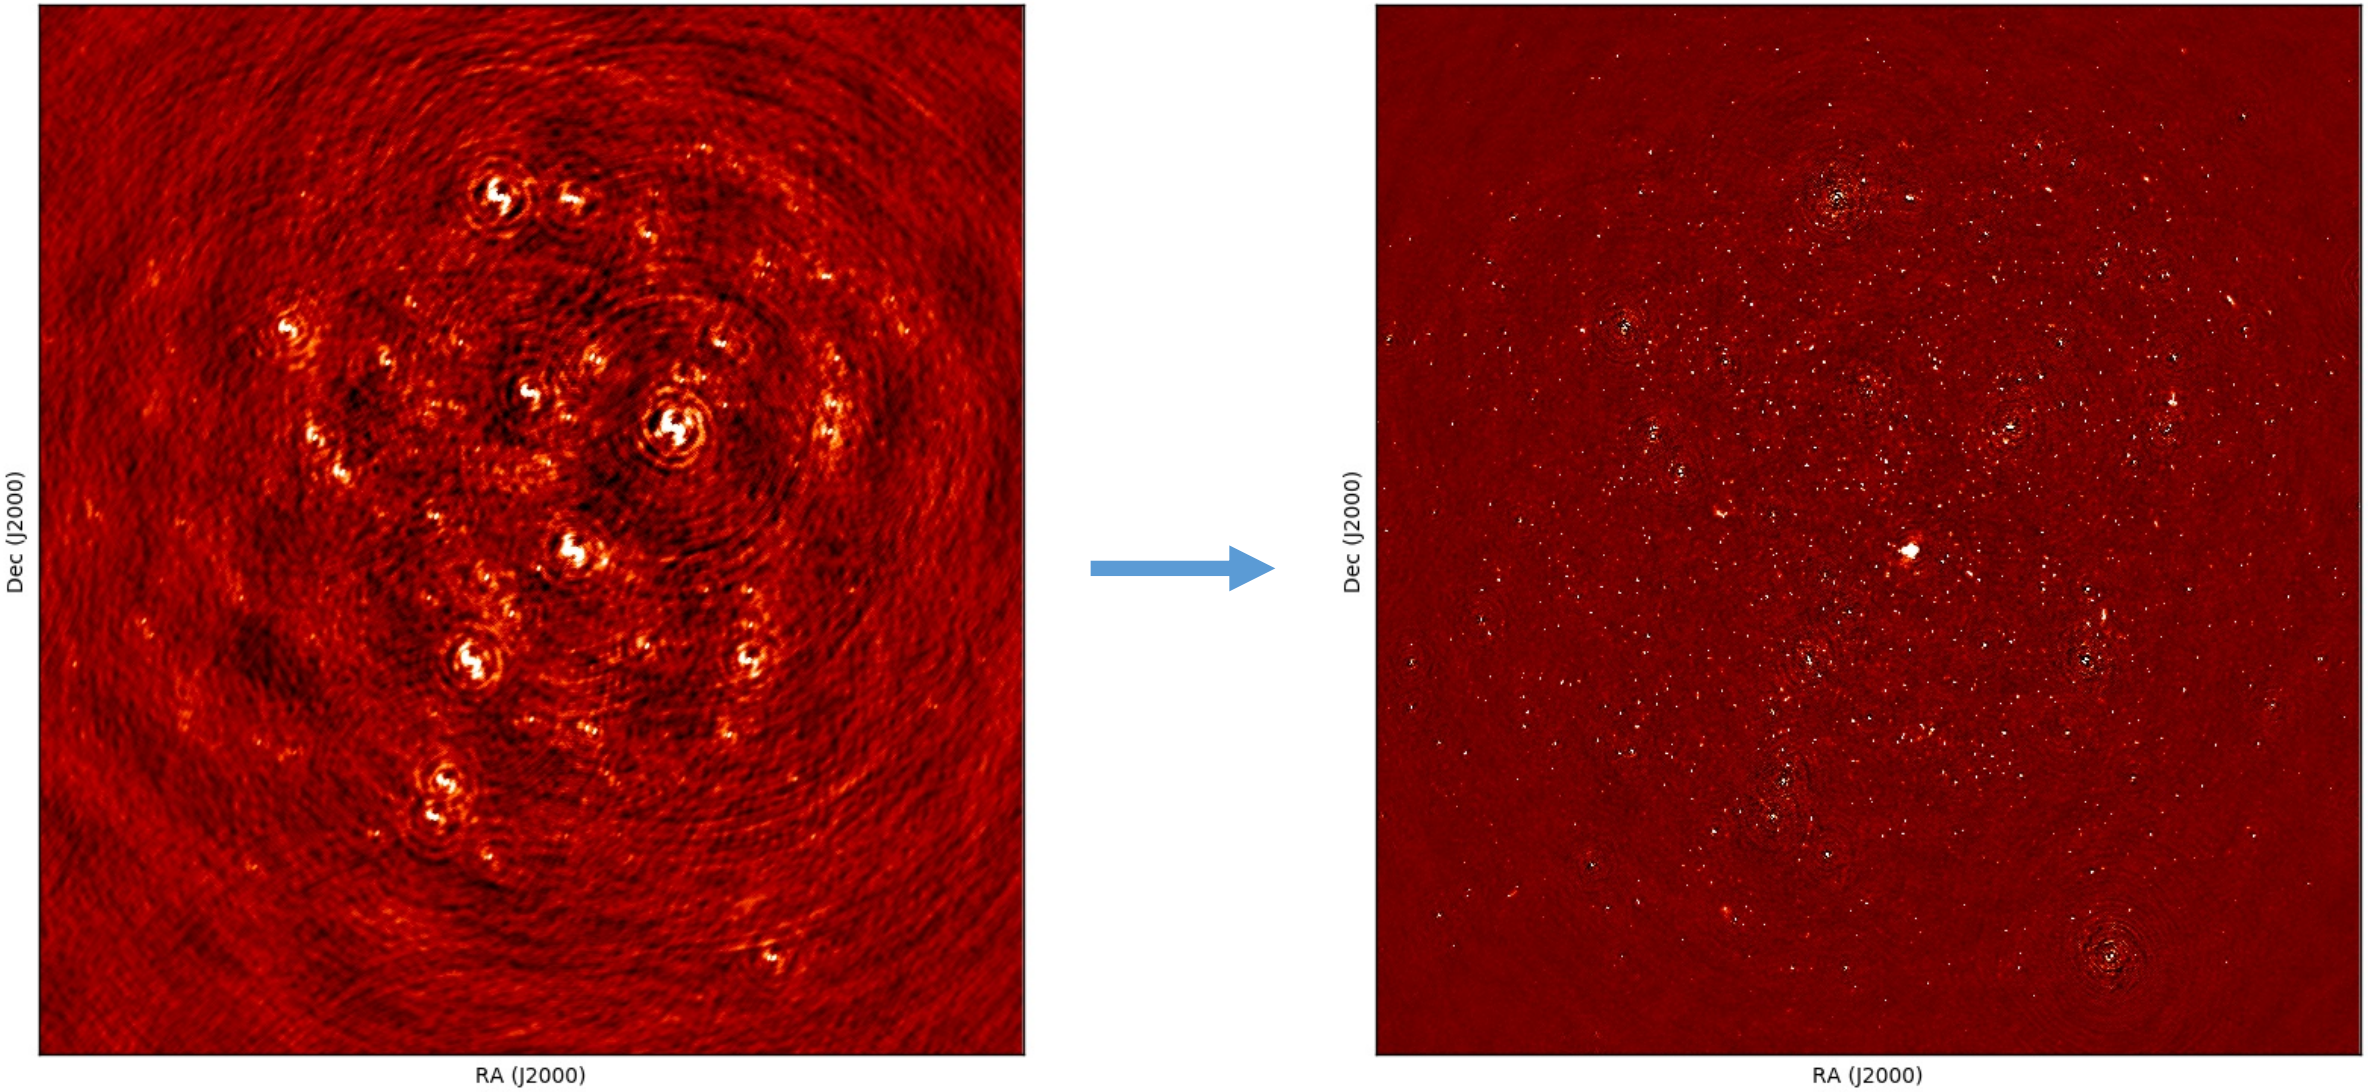

# Overview of Hardware Locations

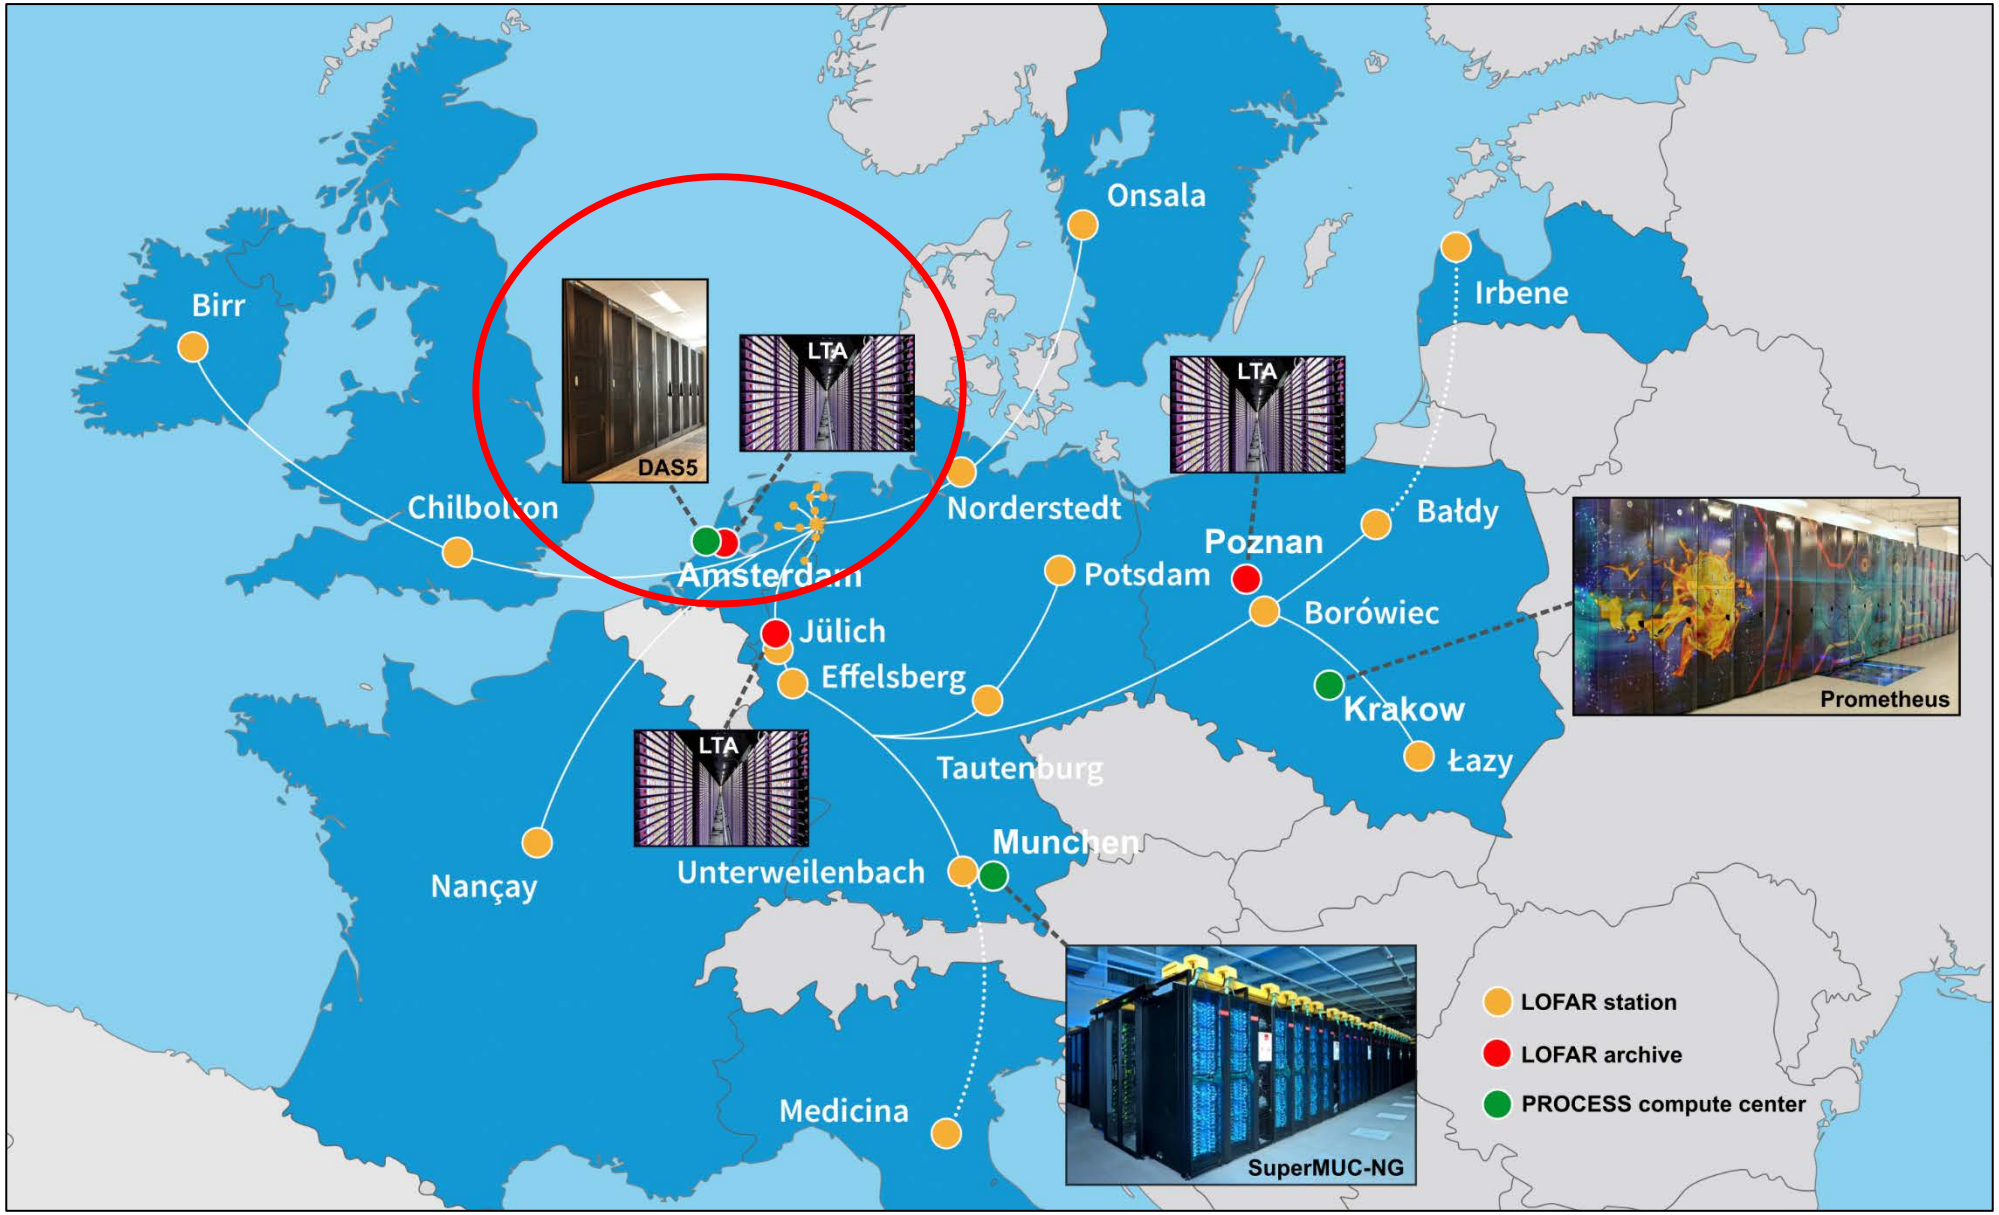

# Video Demonstration 1/3

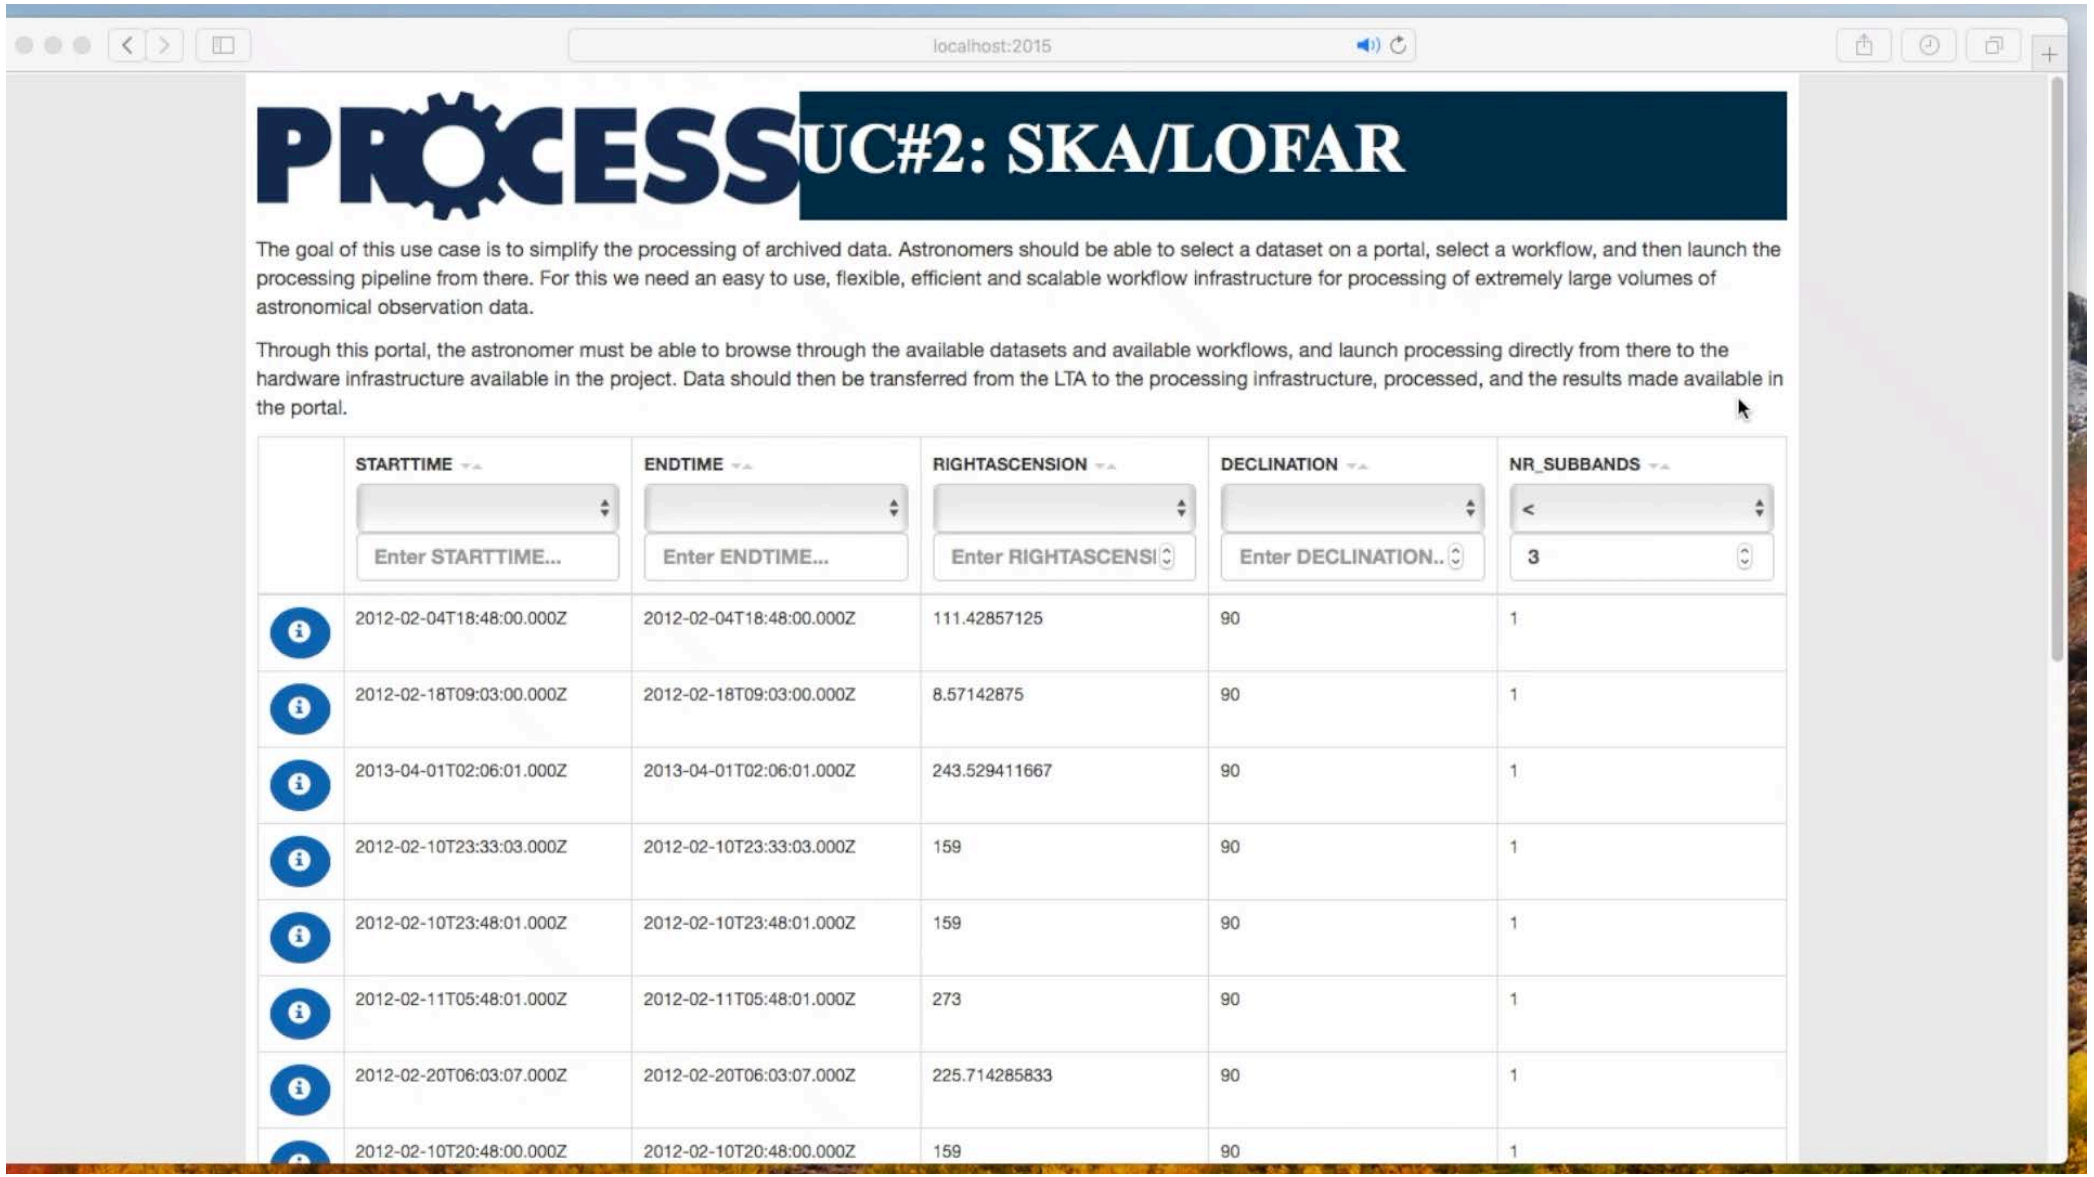

# Video Demonstration 2/3

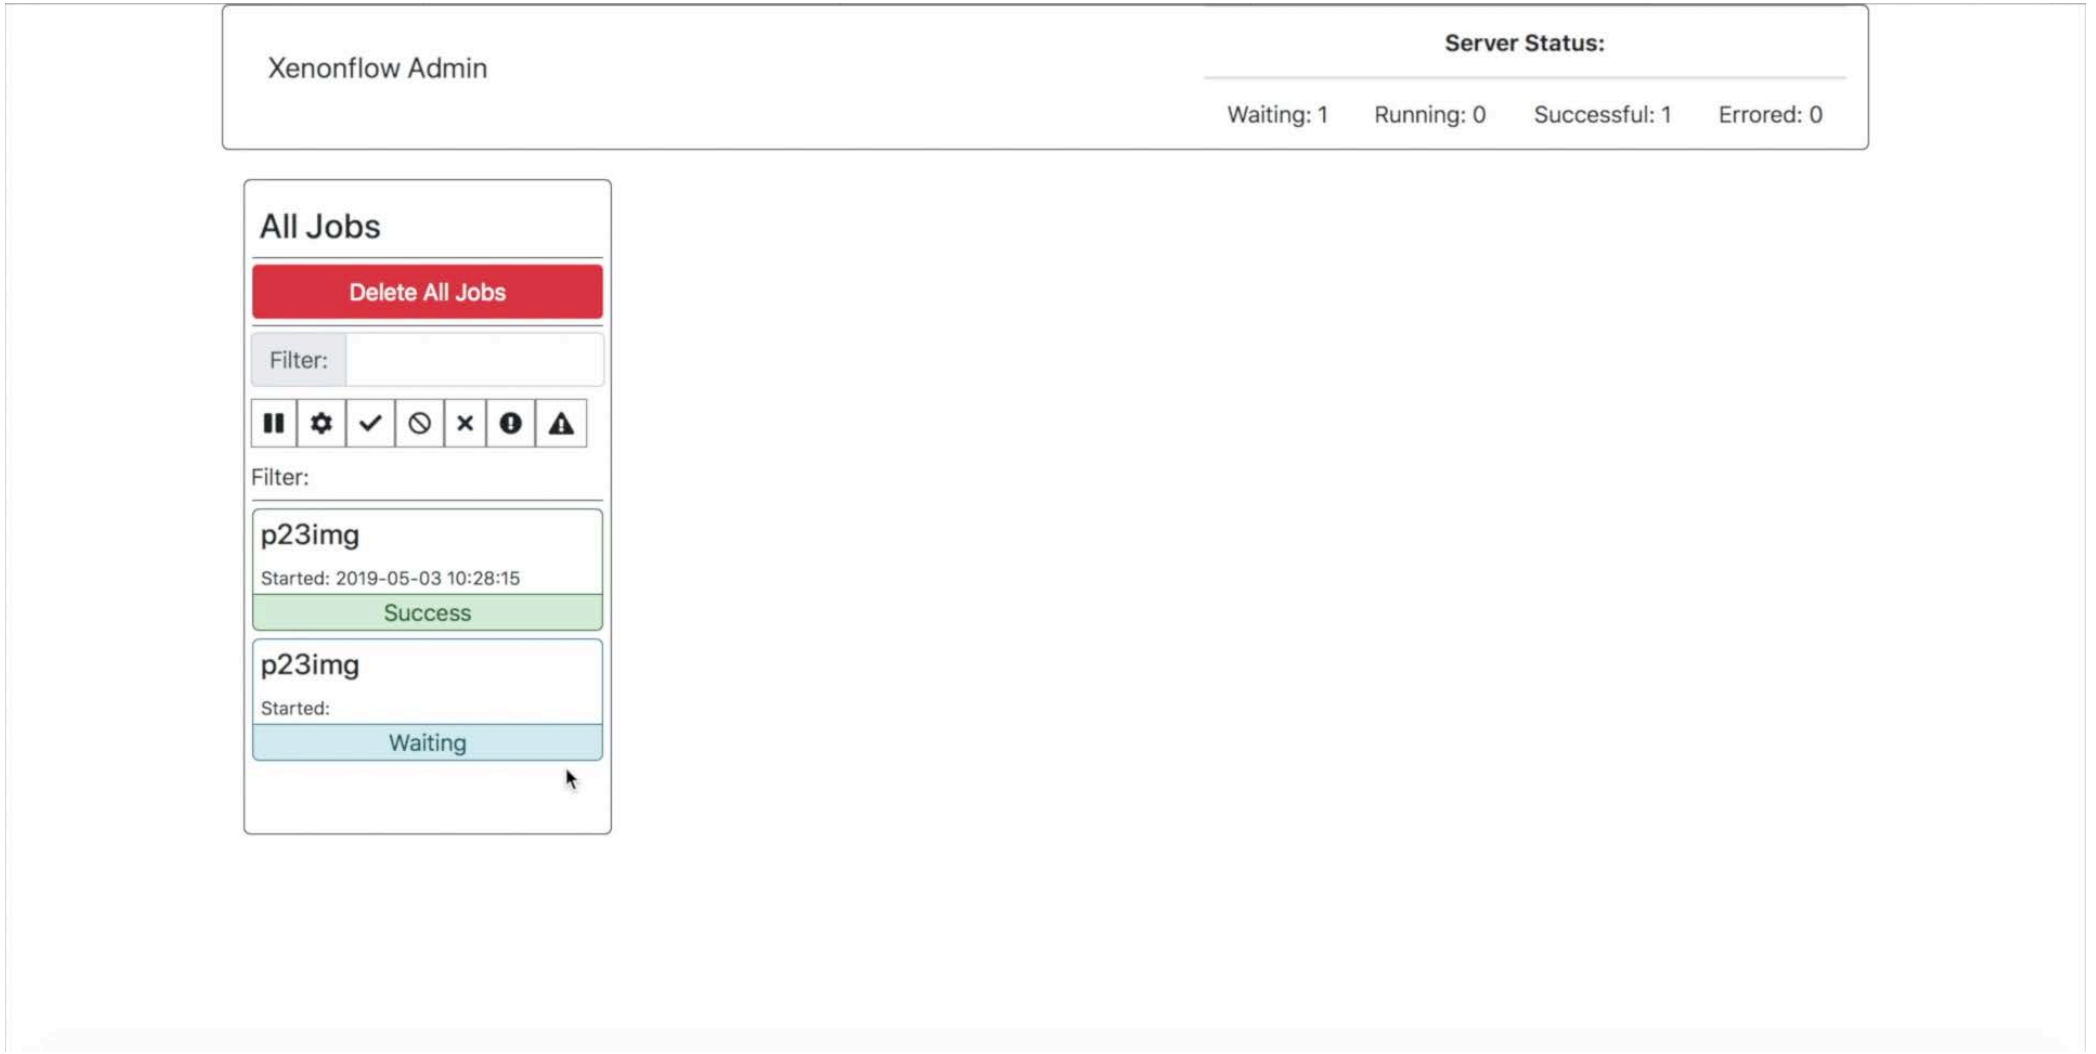

# Video Demonstration 3/3

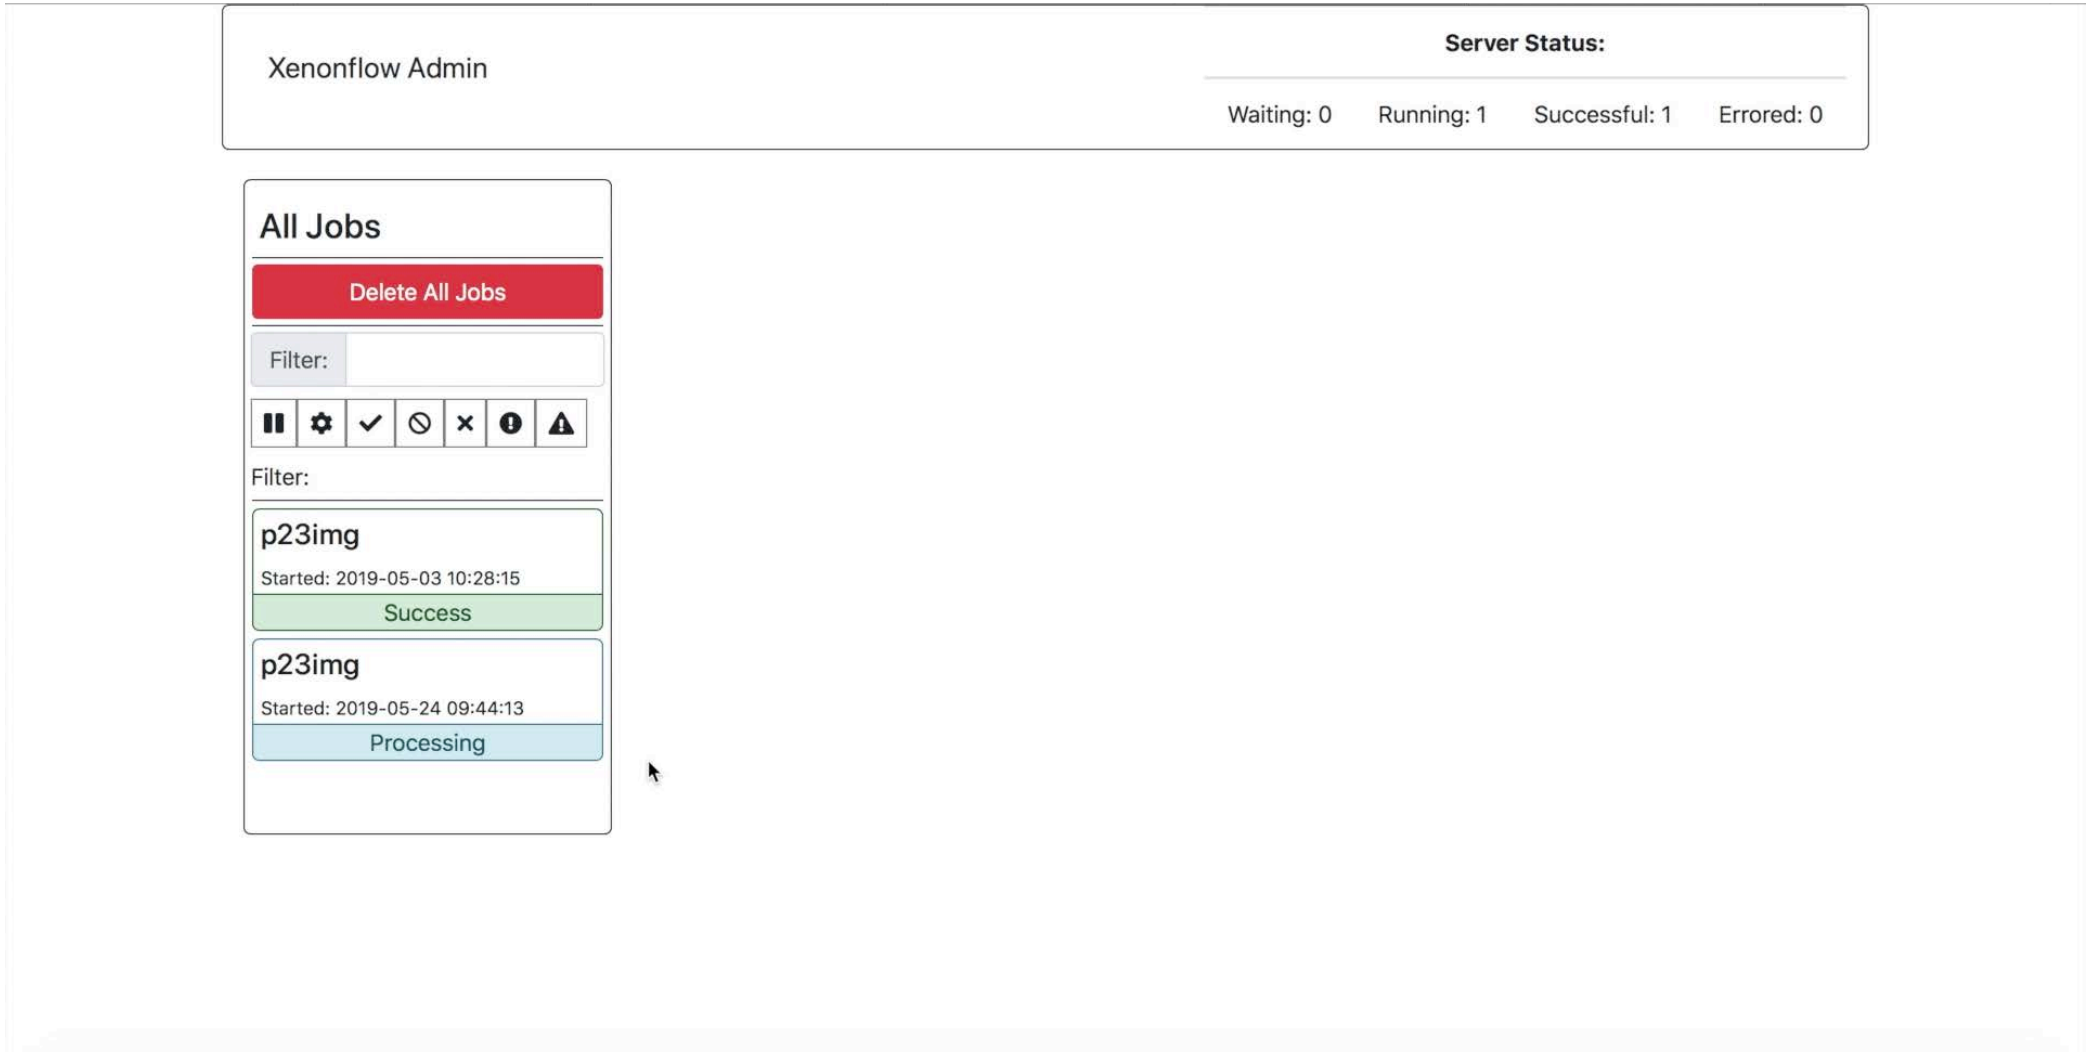

# Conclusions

- There is a need for user-friendly access to the LTA's data and subsequent processing.

"Now I can finally spend my time on actual science!"
Dr. Hanno Spreeuw

- Deliverance of a demo using reusable services.

- **Future work:**
  - Making workflow processing possible on several clusters connected via high-bandwidth networks.
  - Further acceleration of the processing pipeline through parallel computing (direction dependent).

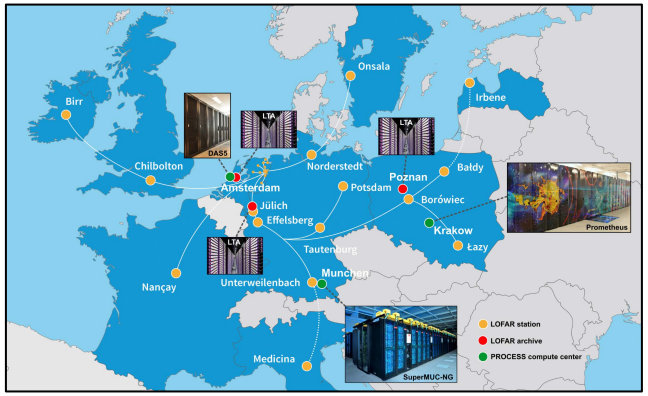

PROCESS

PROviding Computing solutions for ExaScale ChallengeS

# Unlocking the LOFAR LTA

Hanno Spreeuw, Souley Madougou, Ronald van Haren,

Berend Weel, Adam Belloum, Jason Maassen

24th Sep 2019, eScience conference – EINFRA, San Diego

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 777533.