PROCESS

Reference exascale architecture

Martin Bobák, Ladislav Hluchy, Adam S. Z. Belloum, Reginald Cushing, Jan Meizner, Piotr Nowakowski, Viet Tran, Ondrej Habala, Jason Maassen, Balázs Somoskoi, Mara Graziani, Matti Heikkurinen, Maximilian Hob and Jan Schmidt

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

Motivation

- New scientific instruments are producing data at an accelerating pace.
  - Distributed radio telescops (e.g. LOw-Frequency ARray – LOFAR -> contains 30PB of data and grows with 5 to 7 PB/year,
    Square Kilometre Array – SKA -> expected zettabytes/year which will produce 130 to 300PB/year of correlated data, etc.)
  - Space telescops (e.g. Copernicus sentinels -> produce 7.5 PB of raw data each month., etc.)
- Rapid digitization
  - Medical science: medical images of all kinds will soon amount to 30% of all data storage

Storage and Computing Centres

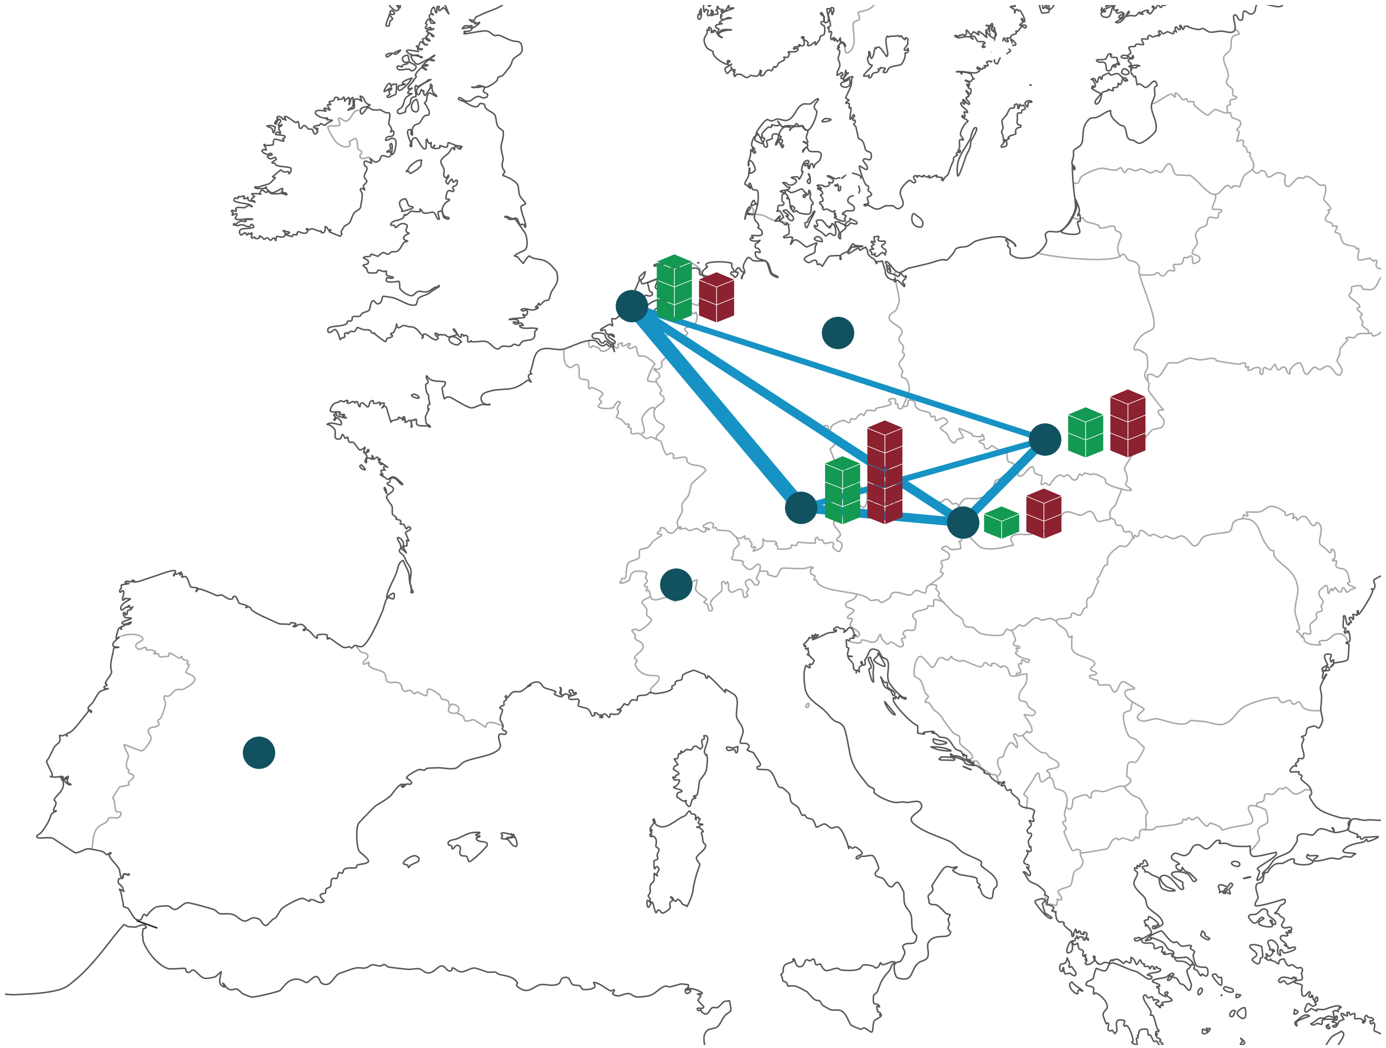

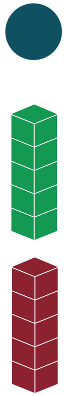

Exascale learning on medical image data

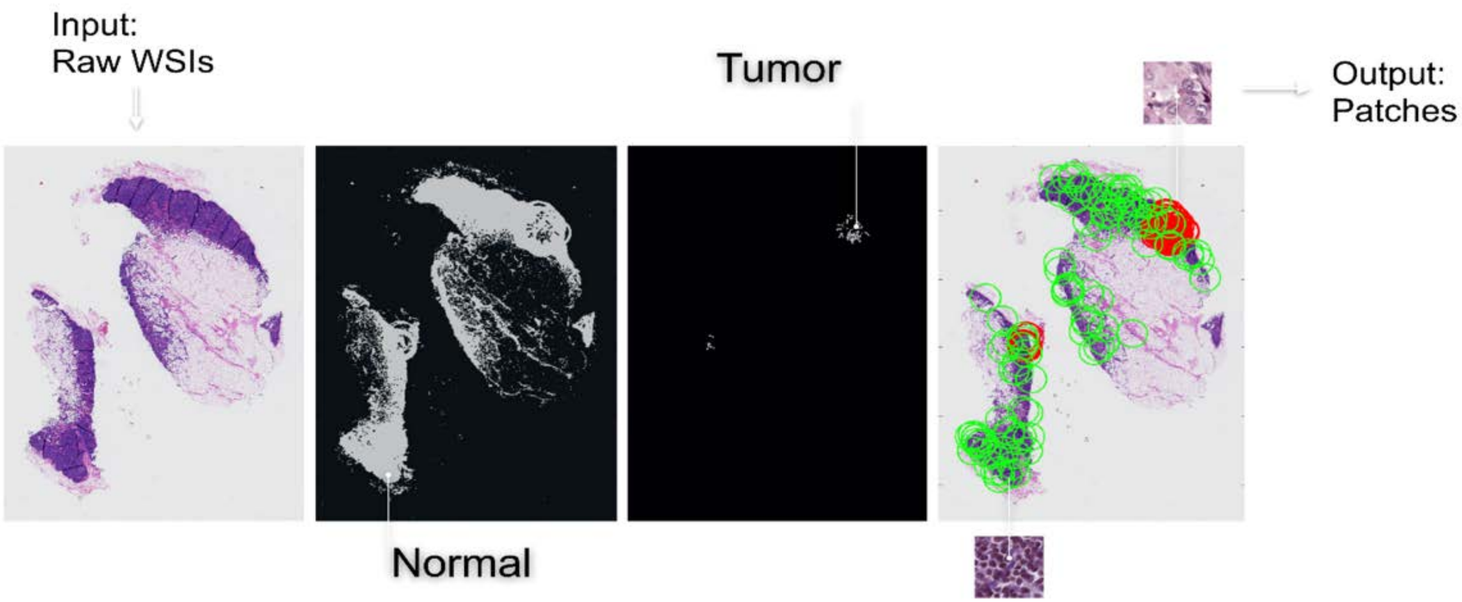

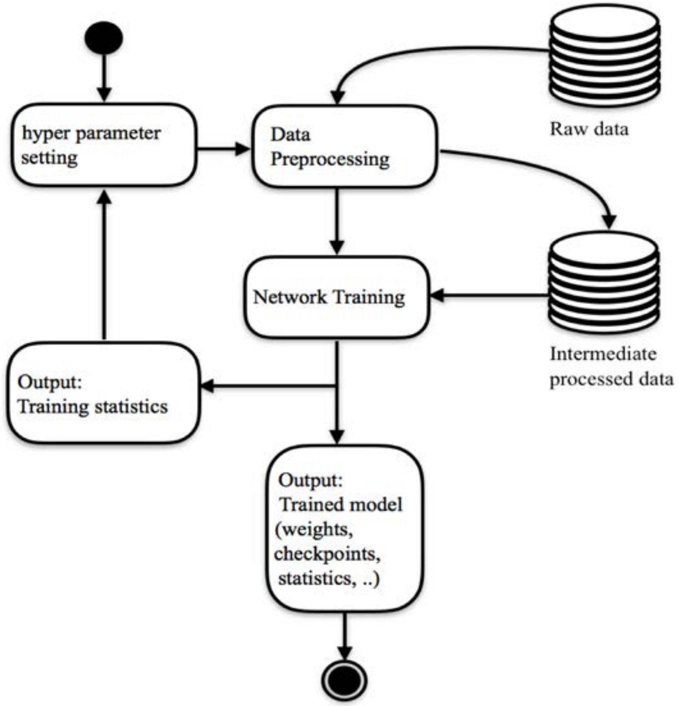

Exascale learning on medical image data

- Extremely high computing demands
- Machine learning and deep learning tools (Tensorflow, Caffe, Theano, Keras, Scikit-learn, PyTorch, Weka)
- GPU acceleration
- Data:
    - 30% of world storage estimated to be medical imaging in 2011
    - In 2010 mammography in the US amounted to 3 PetaBytes
    - This did not include histopathology images!
    - Extremely large: >100,000x100,000 pixels

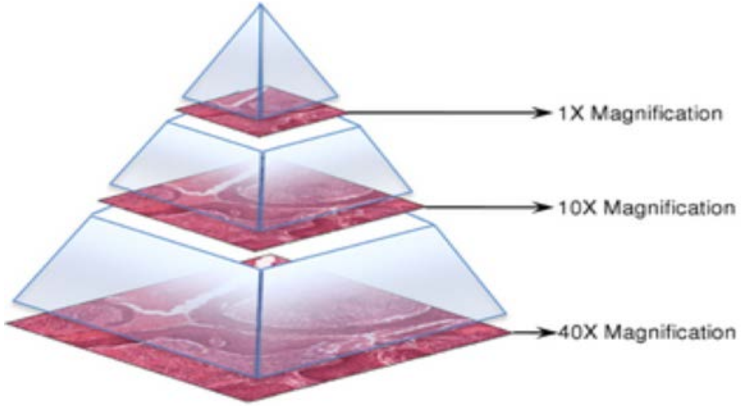

Reduction of radioastronomy observations to sky maps

PROCESS

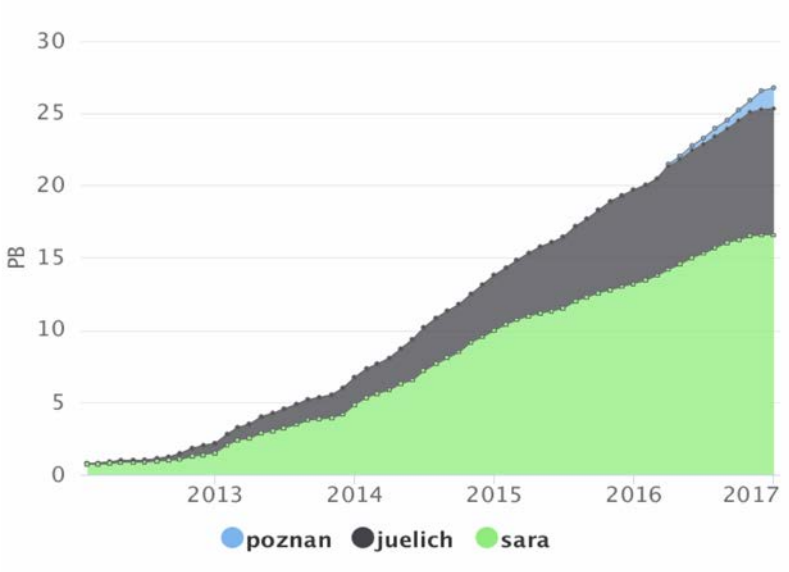

- LOFAR observations are stored in the **long term archive (LTA)** which is distributed over **Amsterdam**, **Juelich** and **Poznan**.
- It currently contains **~30 PB of data** and **grows with 5 to 7 PB/year**. Data is stored on **mostly on tape**, using hard disks for temporary storage (dCache).

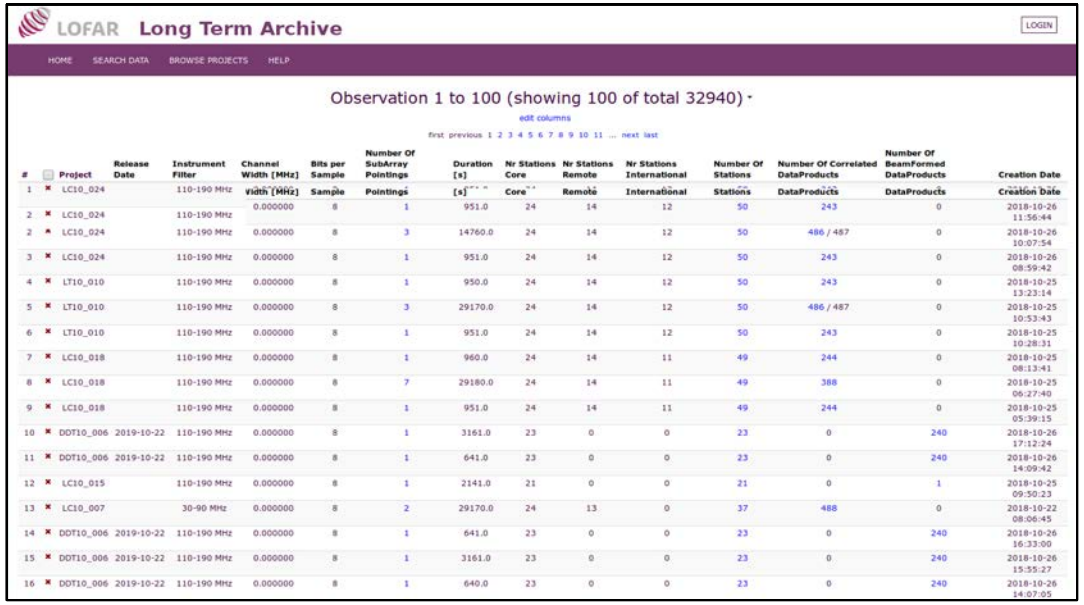

LOFAR Long Term Archive

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											
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

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**!

PROCESS

Reduction of radioastronomy observations to sky maps

- 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 out there 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**.

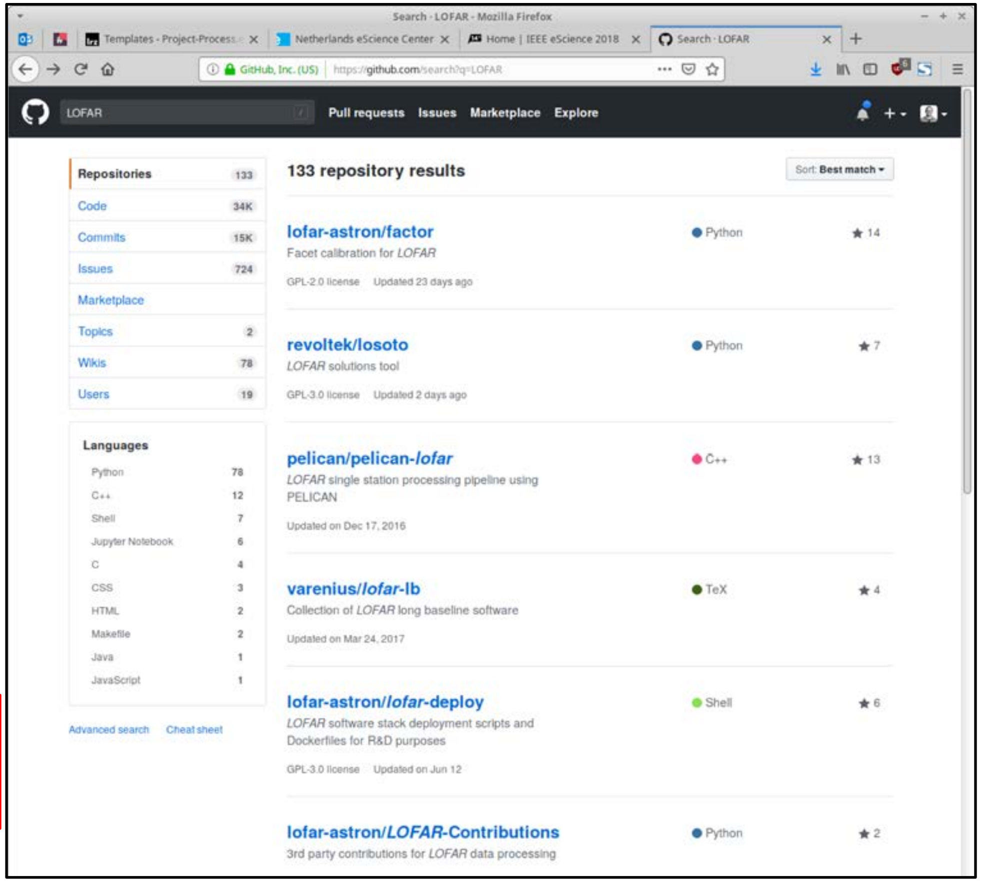

The combination of these two makes it **hard** for **astronomers** to make optimal use of the archive.

Reduction of radioastronomy observations to sky maps

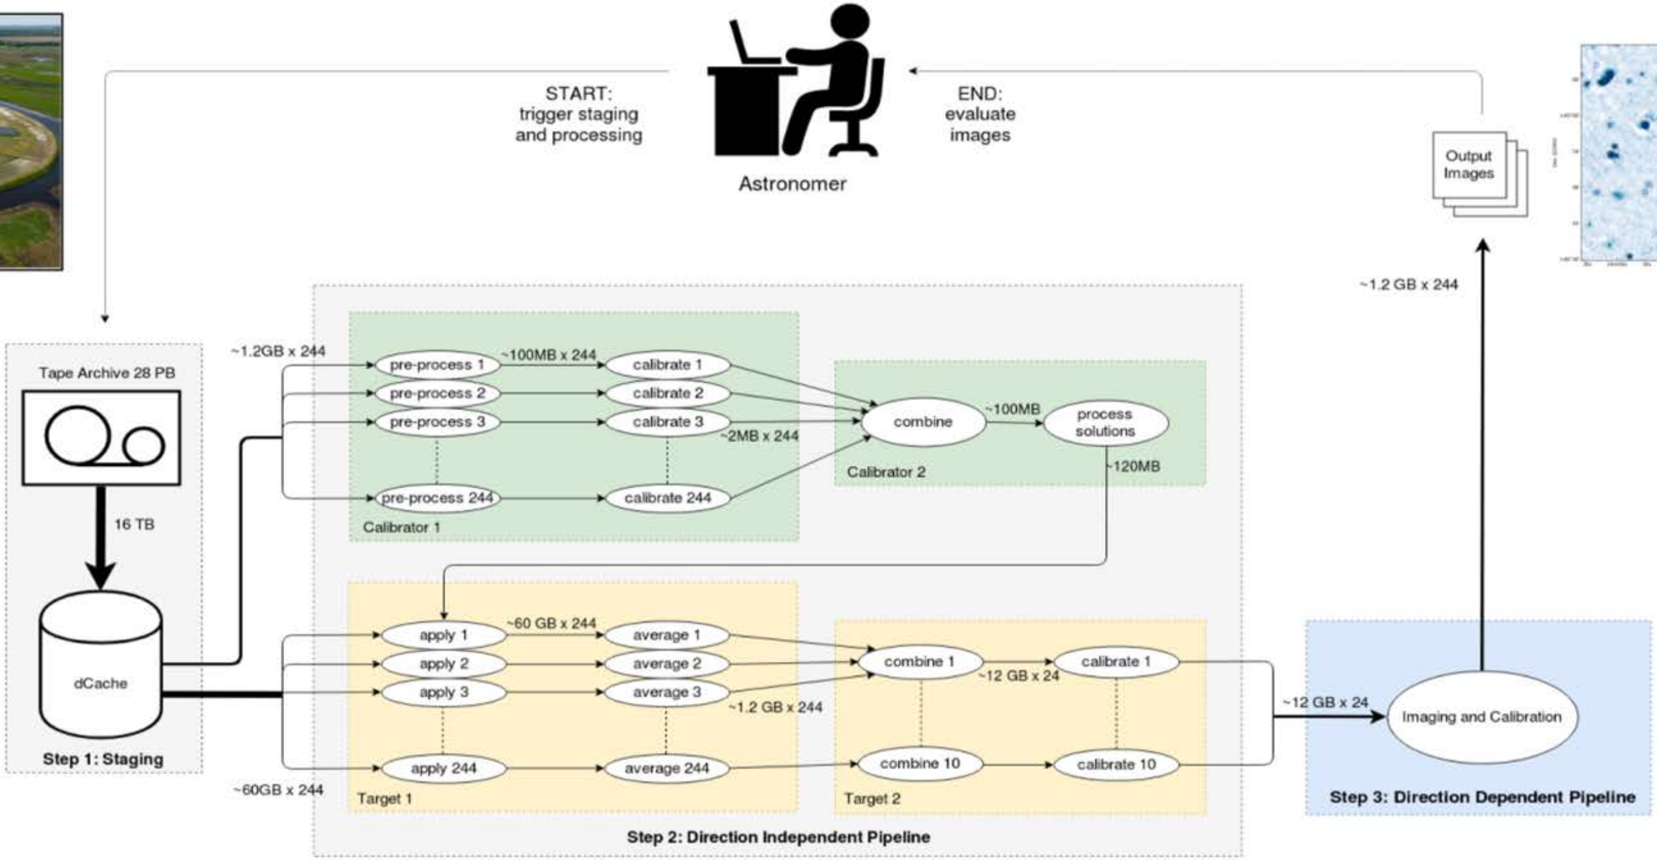

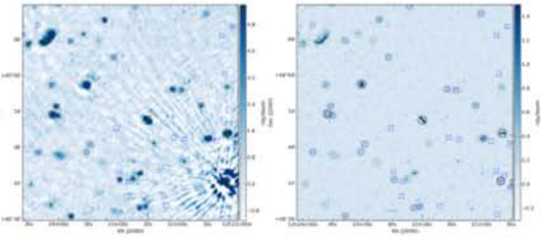

Ancillary pricing for airline revenue management

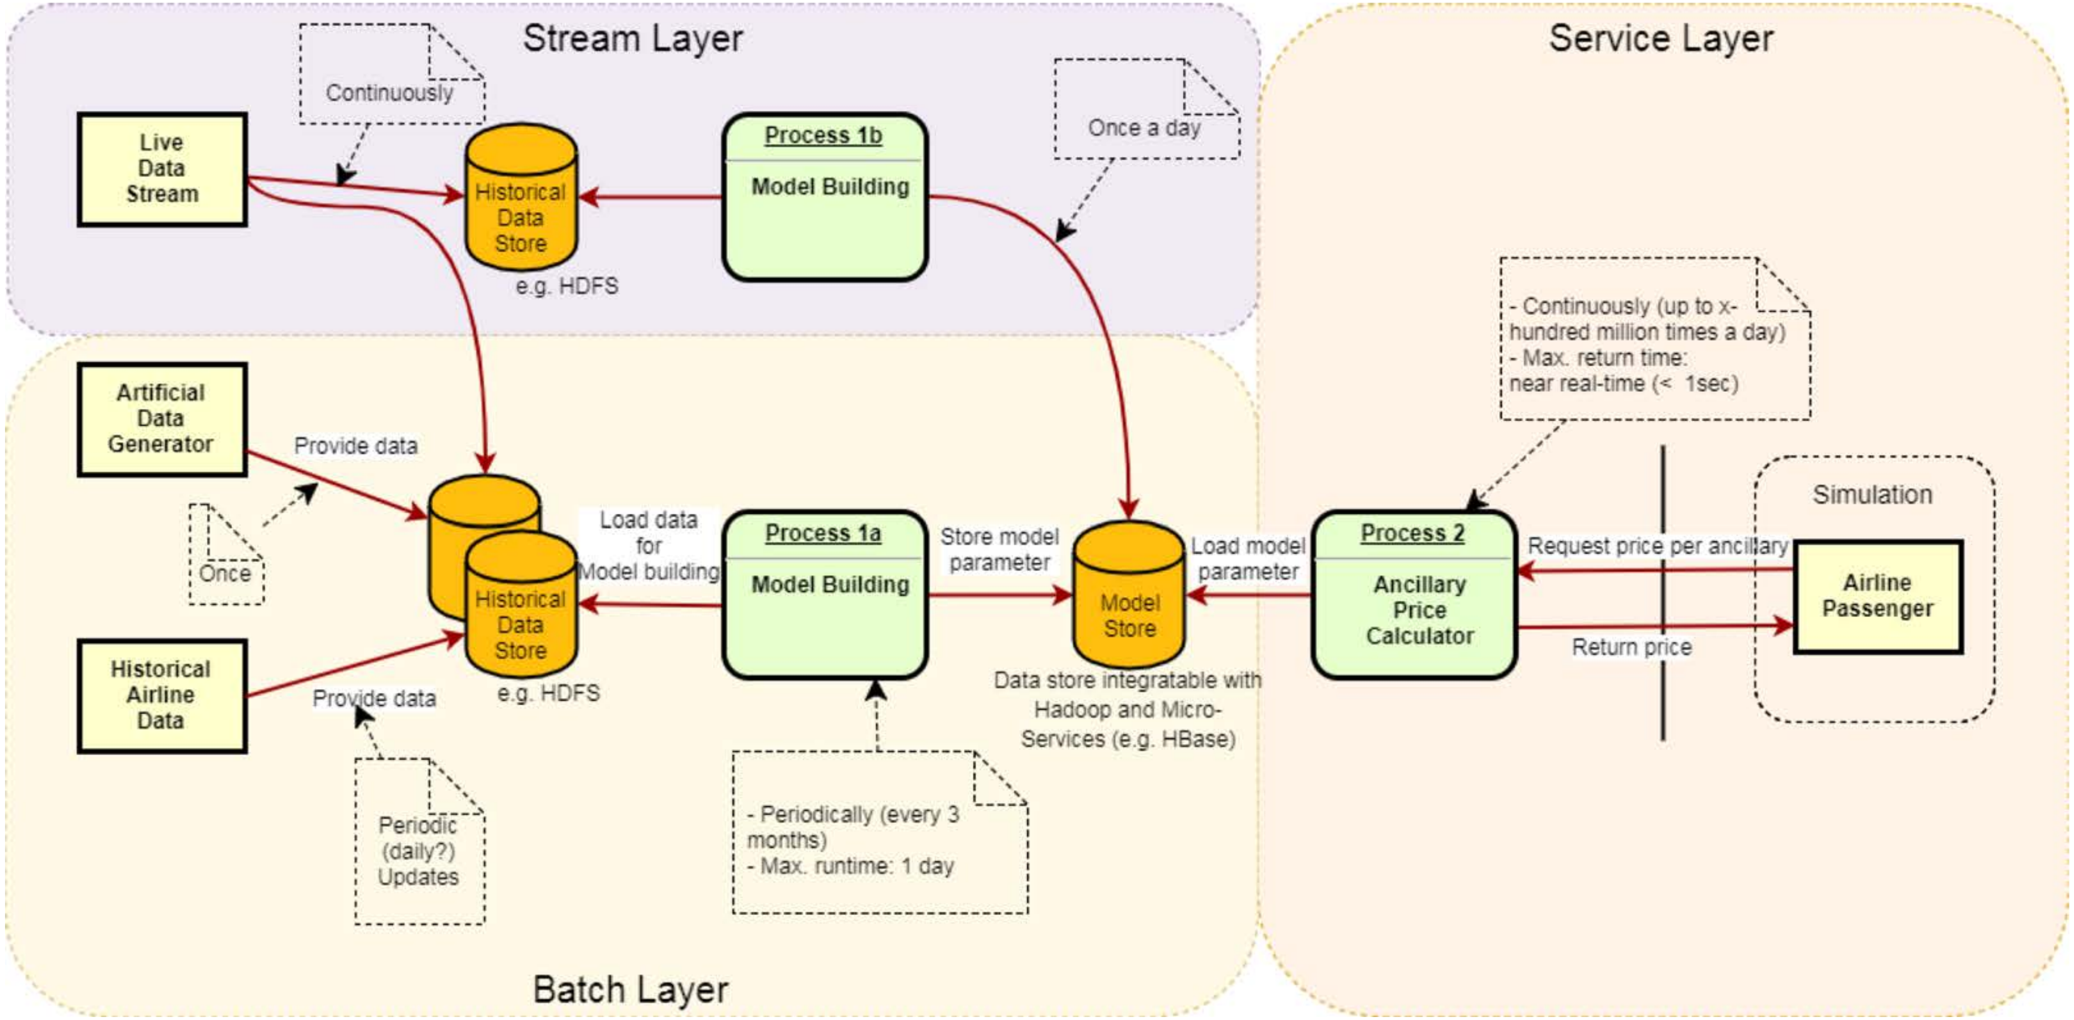

Ancillary pricing for airline revenue management

- The service layer uses state-of-the-art microservice software stacks such as Spring boot, Kubernetes etc. and provides RESTful services.
- The streaming and batch-layer rely on Hadoop/HDFS infrastructure and use HBase as an interface between the big data side and the microservice side.
- The streaming layer is built on Apache Spark.
- The model building part uses Tensorflow and H2O.ai.

Agricultural analysis based on Copernicus data

- Harvest prediction on European scale based on simulation models from cell-level phenomena using Earth observational data (Copernicus satelites produce around 4.5 PB per year).

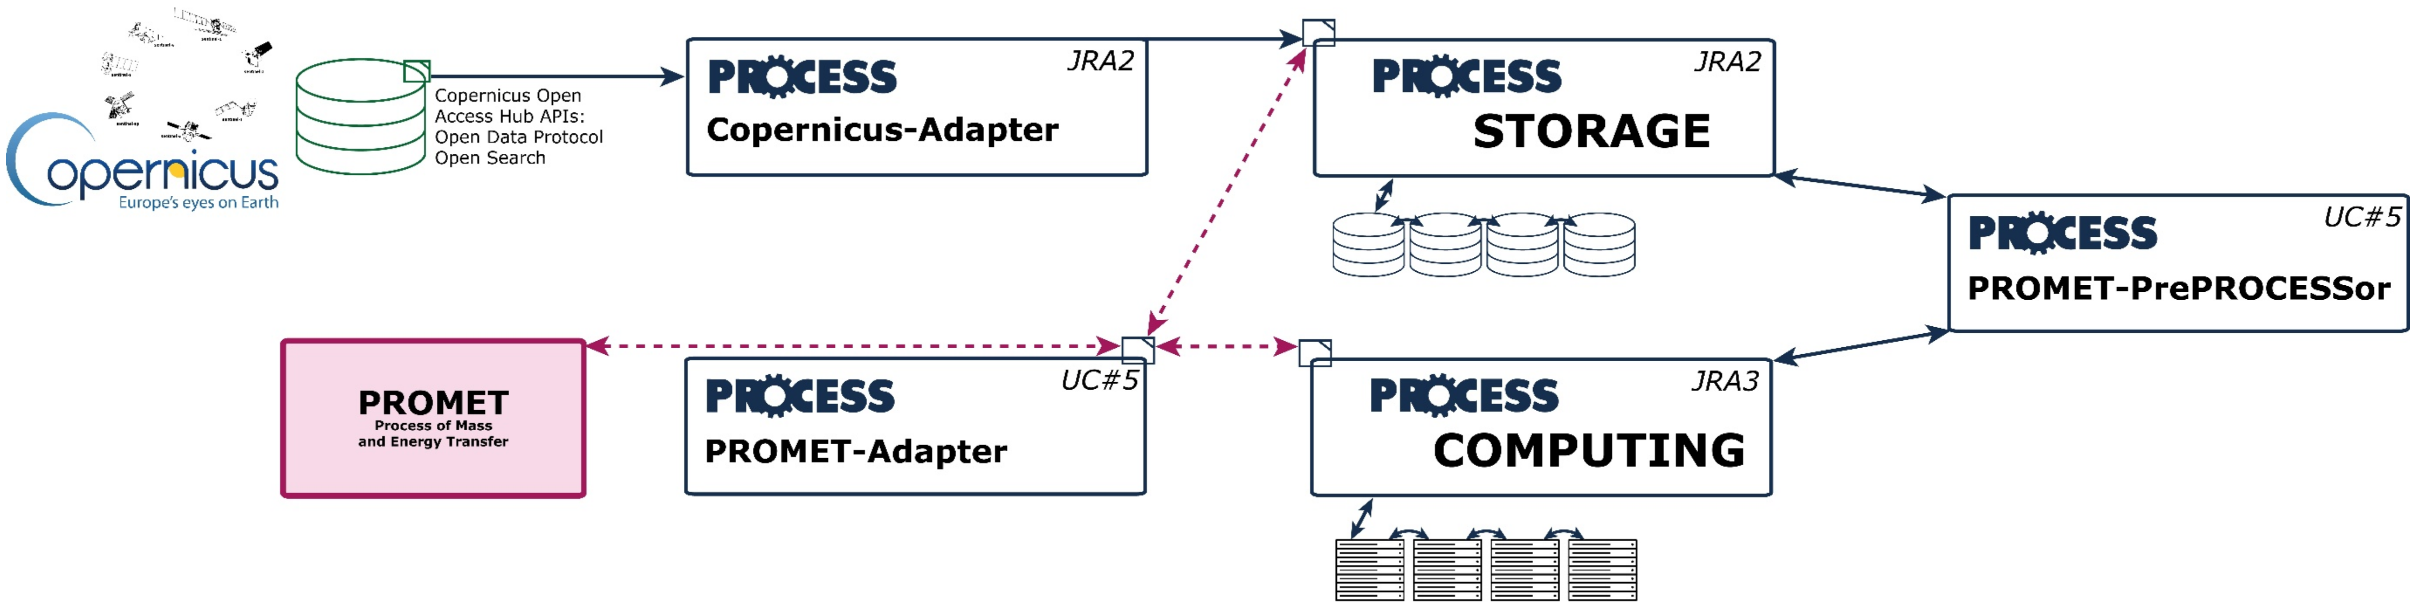

Reference exascale architecture

- High distribution and parallelism -> modularity and scalability
- To extend and adjust the platform, according to the needs of new user communities -> sub-modules which exploits the heterogeneous resources of exascale systems
- The proposed exascale reference architecture is based on containerisation instead of virtual machines.

Reference exascale architecture

- **Virtualization layer:**
  - Support of containers (lightweight virtualization)
  - Easy deployment and maintenance, flexibility, reliability, scalability
  - Exploit various computing infrastructures (portability and interoperability)
- **Data services:**
  - Federated and distributed datasets
  - Data described by meta-data
  - Massive data transfers
- **Computing services:**
  - high distribution across different research computing centers

Reference exascale architecture

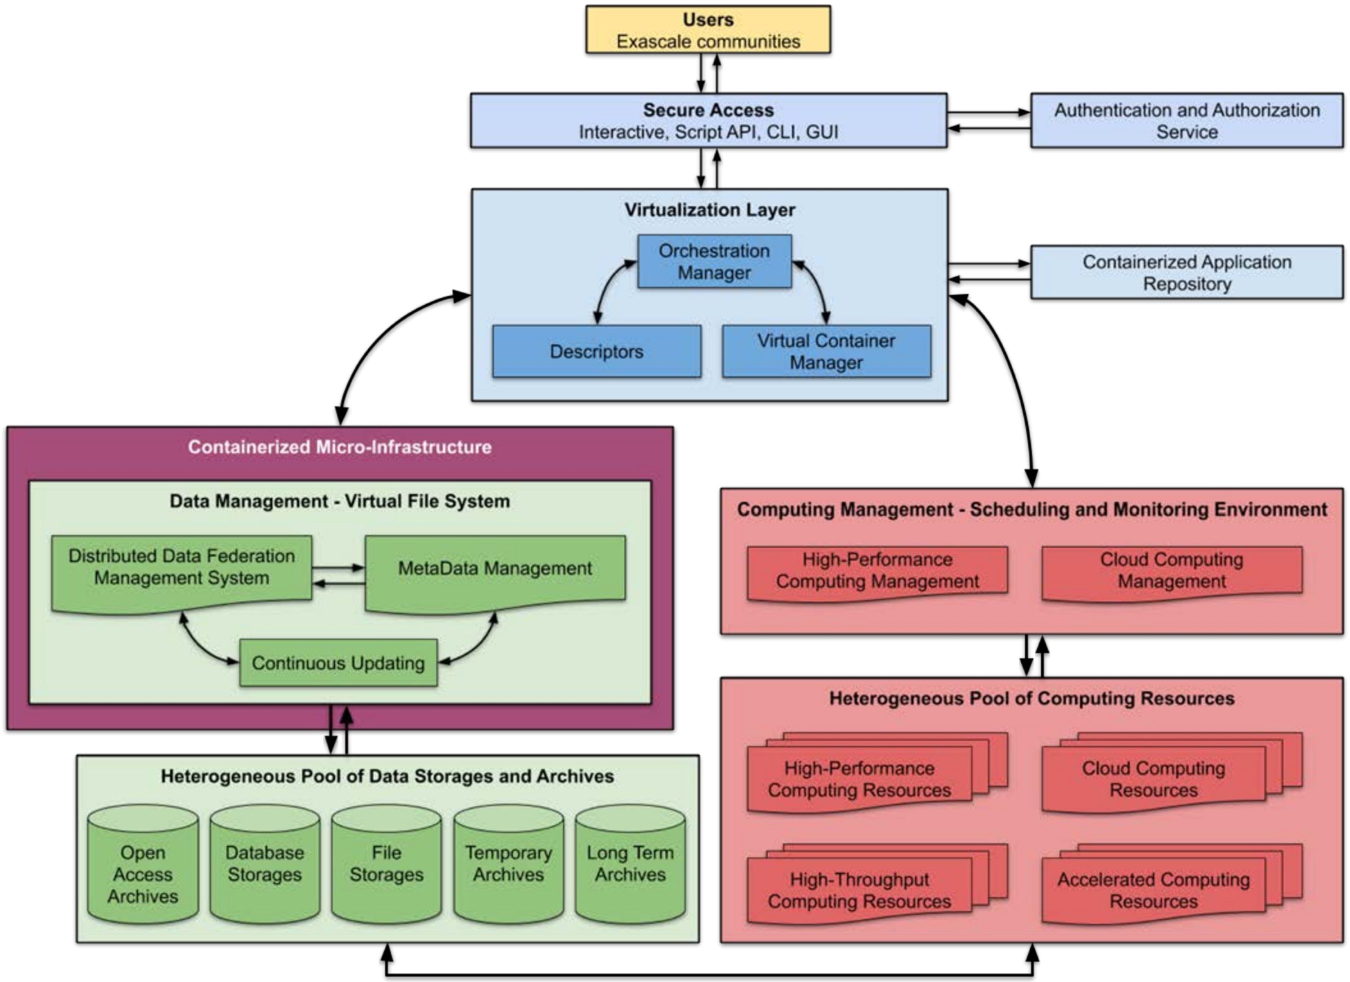

PROCESS architecture

- Modular software platform capable to handle exascale datasets
- Microservice infrastructure
  - services can be customized for the application
  - minimizing global management
  - micro-infrastructures are isolated from each other
- Interaction through RESTfull APIs (mainly)
- Authentication by security tokens

PROCESS architecture

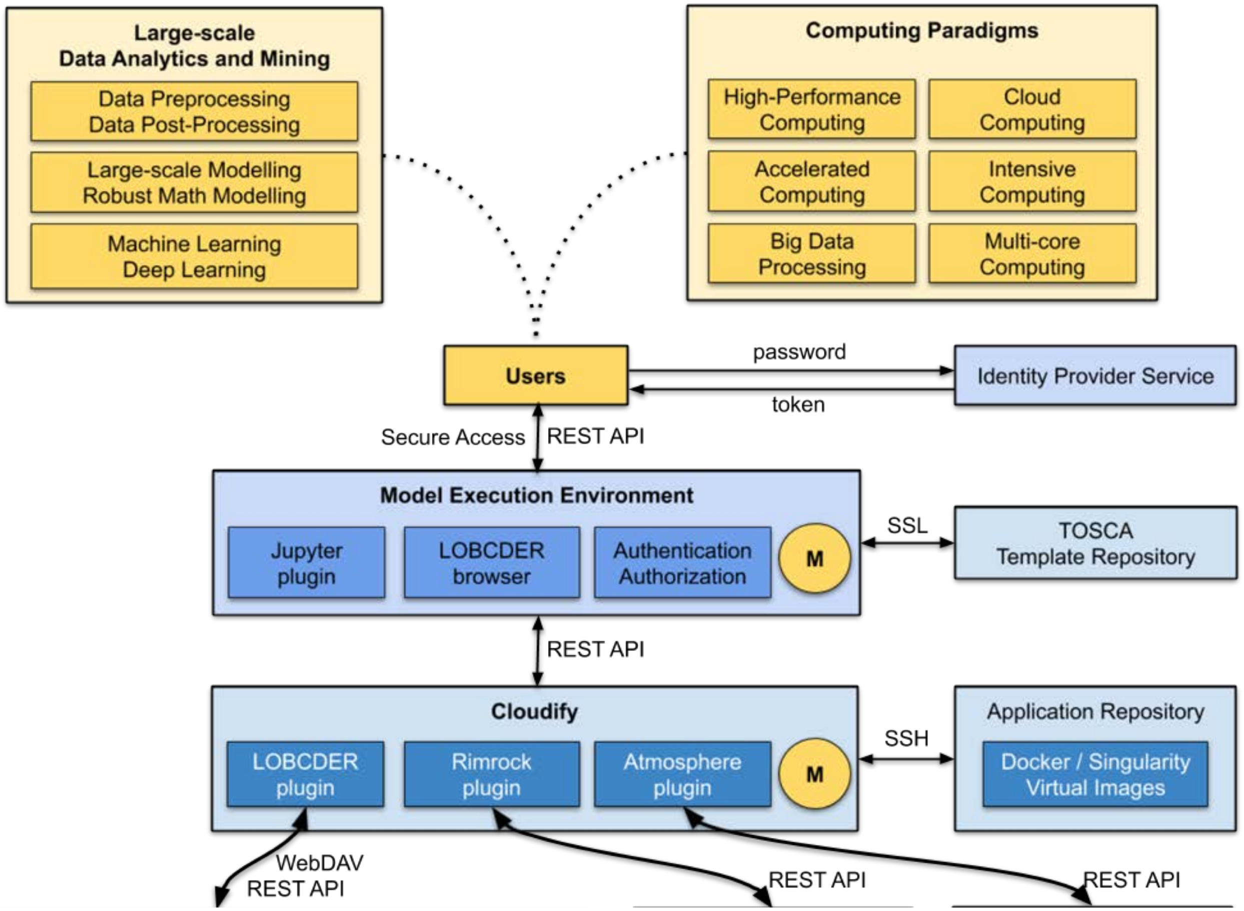

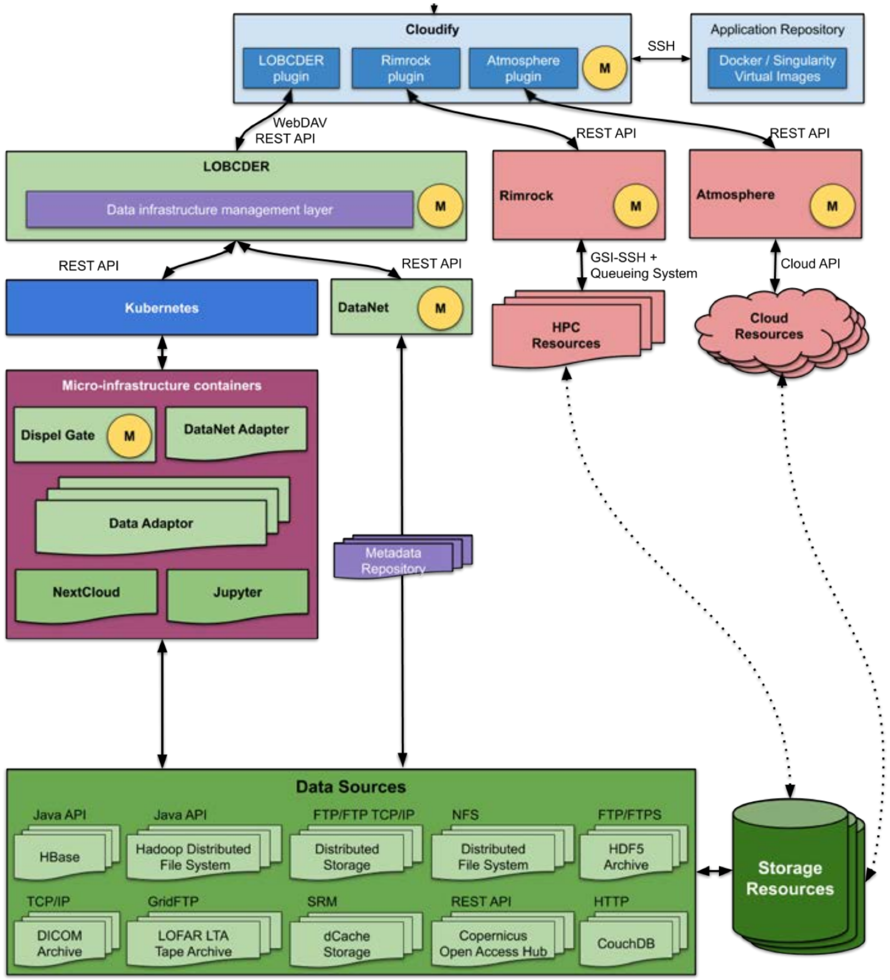

Conclusion

Mature, modular, generalizable Open Source solutions for user friendly exascale data.

UC#1: Exascale learning on medical image data Hes·so

UC#2: Analysis of netherlands eScience center **Radioastronomy Observations**

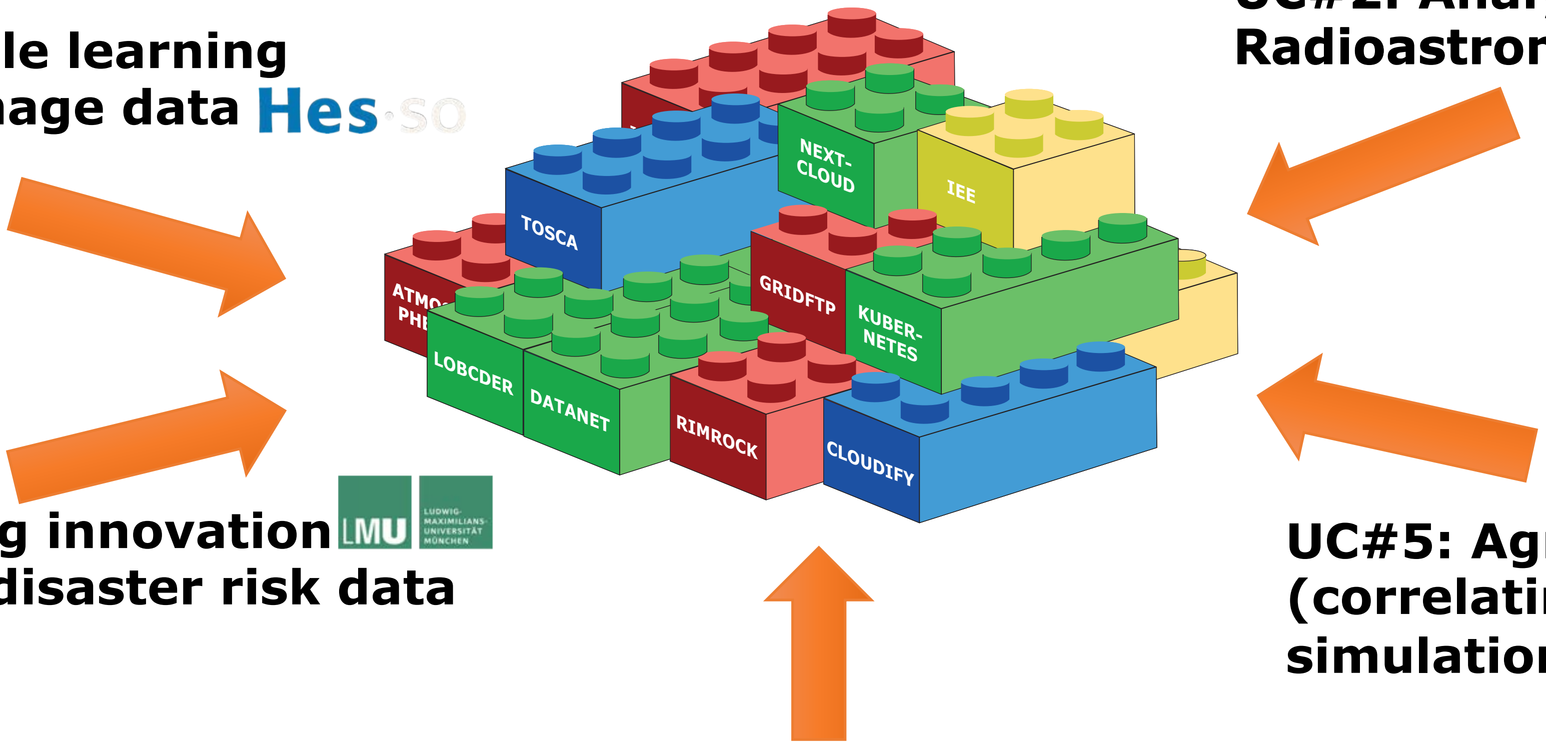

UC#3: Supporting innovation based on global disaster risk data LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

UC#5: Agro-Copernicus (correlating data between simulation and observation) LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

UC#4: Ancillary pricing/airline revenue management Lufthansa Systems

Thank you for your attention.