Data Management strategies at LEAPS facilities: the example of the ESRF

Vincent Favre-Nicolin
Algorithms & scientific Data Analysis

+ A. Götz, J-F. Perrin, A. de Maria, R. Dimper

MORE BRILLIANCE & PHOTONS…TOWARDS BIG DATA

Yearly production Estimates:

2016 – **2.8 PB**

2018 – **8 PB**

2021 – **20 PB**

2025 – **60 PB**

Archiving:

SL8500 tape library Current architecture handles up to **170 PB**

Detector portfolio:

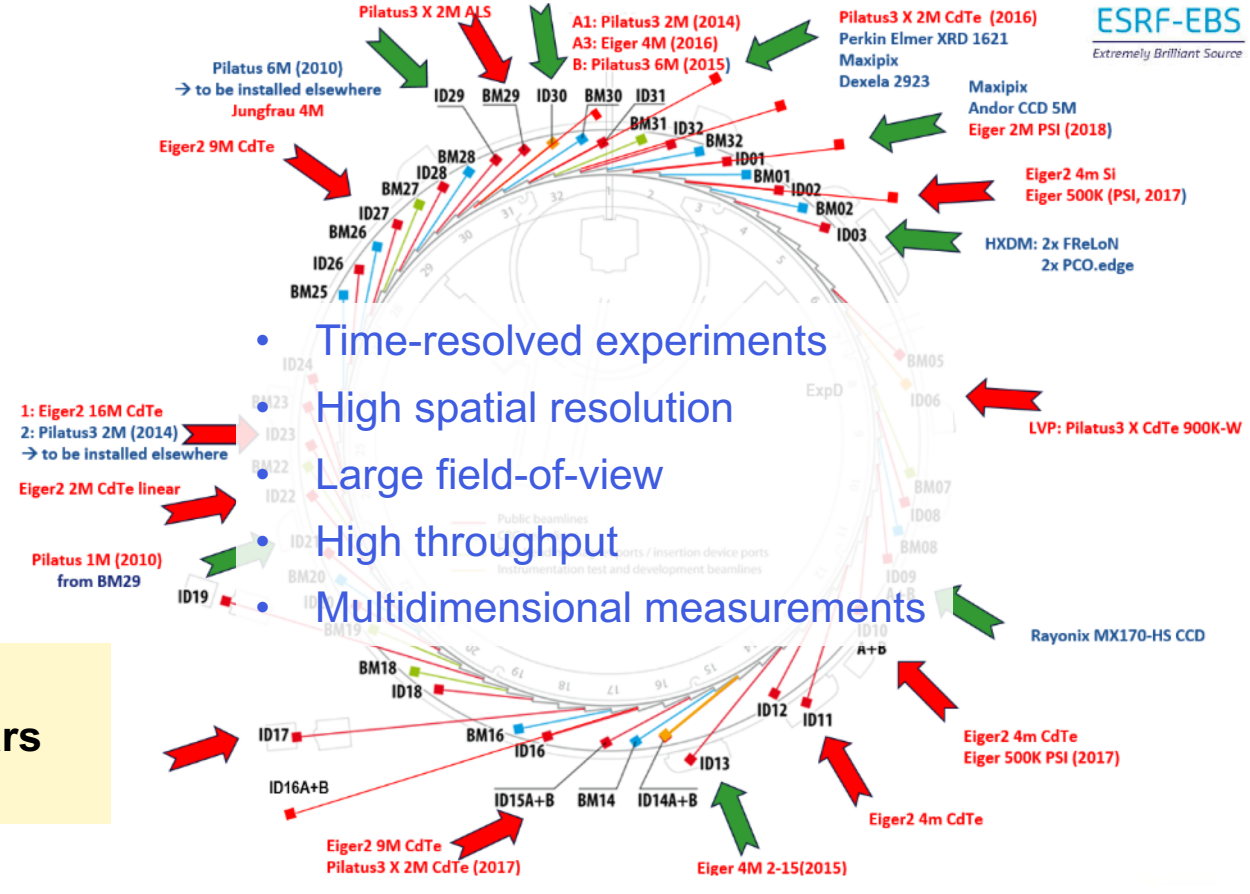

Data Policy: (FAIR principles)

- **archive ‘raw’ data for 10 years**
- **Public (DOI) after 3 years**

DATA MANAGEMENT PILLARS

- Explosion of the data volume
- Significant increase of complexity of the data sets
- limited dedicated resources

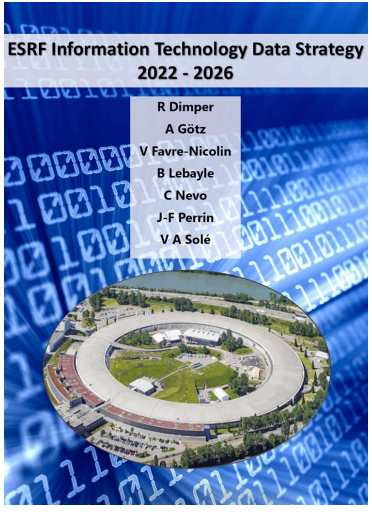

DATA ACQUISITION: SERIAL XTALLOGRAPHY

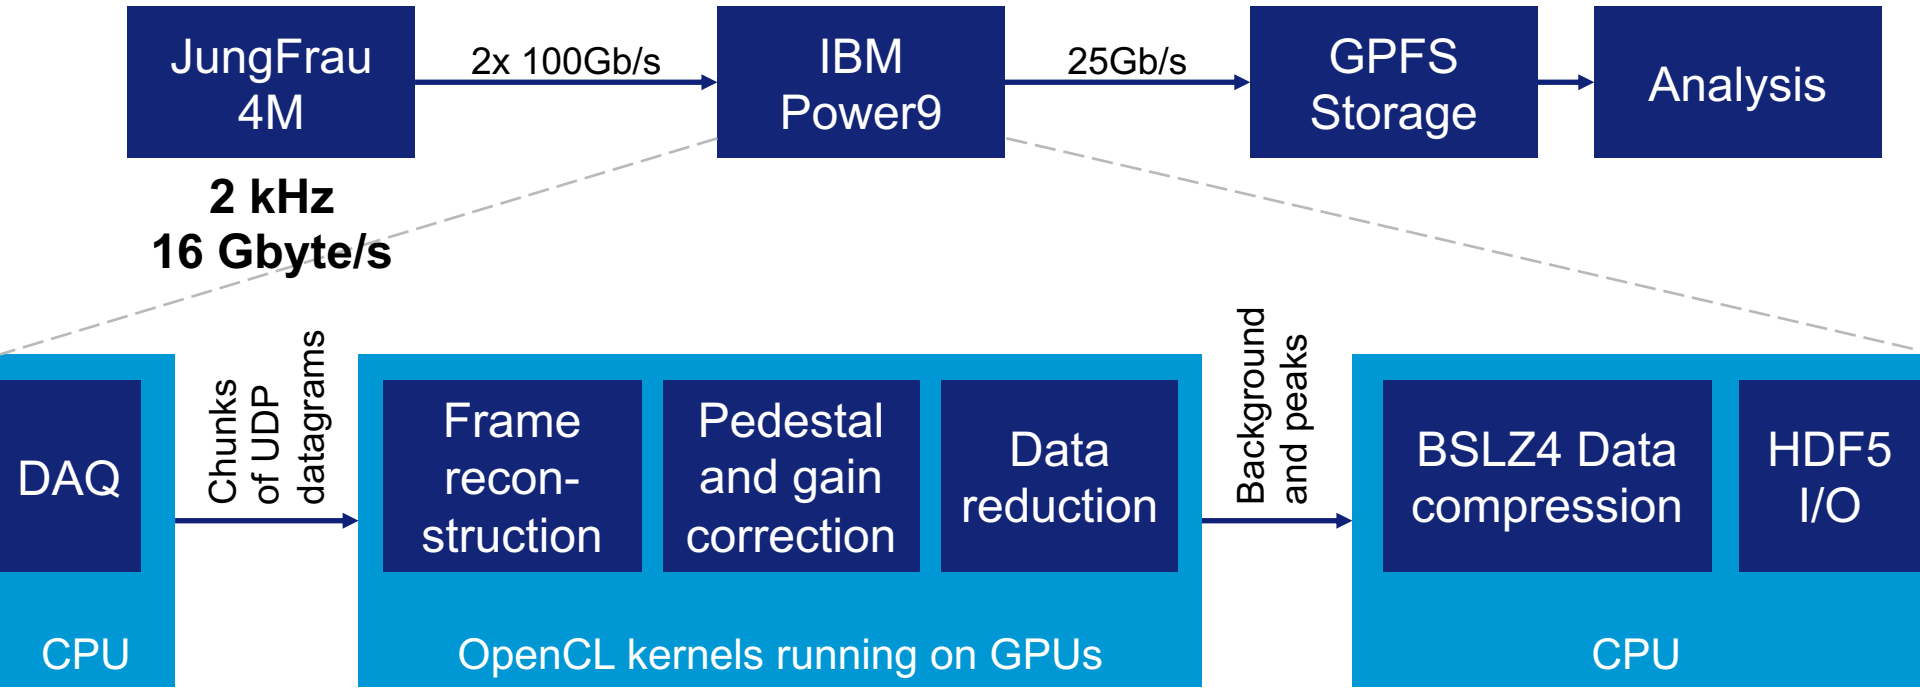

LIMA2 Receivers
multiple instances
collaborating through MPI

Pipeline running @2KHz

STORAGE : GPFS

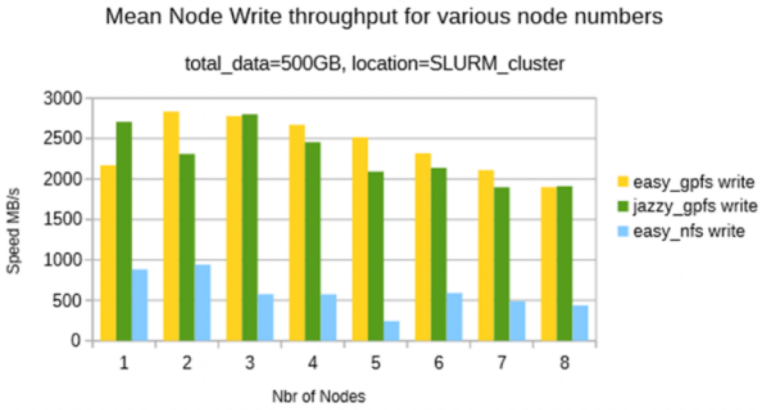

Ethernet

- Most common computer network technology
- Network speed up to 400Gbps for backbone links, up to 100Gbps for cluster nodes
- Cluster nodes come with standard 25Gbps interfaces
- GPFS can fully exploit network bandwidth (i.e. 2.5 GBytes/s from a single thread on a 25Gb/s eth link)

Data is :

- kept on disk for 90 days to enable analysis
- total disk capacity: 15 PB
- then archived

GPFS I/O performance on 25 GB/s ethernet

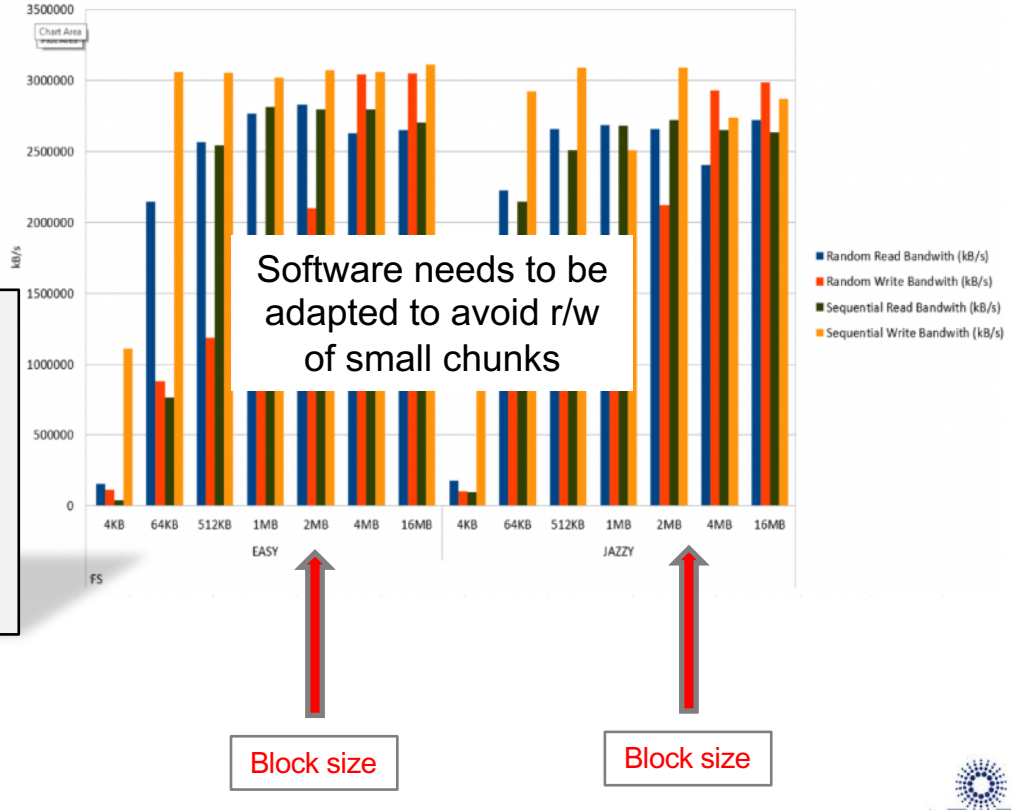

HDF5 / NEXUS – RICH & EXTENDABLE DATA FORMAT

NXroot

Top level. One per file.

NXentry

One group per measurement

NXinstrument

Describe the instrument.
Only one per NXentry

Exclusive **Acquisition** Domain

measurement (@NXcollection)

Flattened view of everything measured

Only one per NXentry

Almost exclusive **Acquisition** Domain

sample (@NXsample)

Define the physical state of the sample during the scan

User/Scientist Domain

NXdata

The default data to be plotted.
One NXdata group per plot

User/Scientist Domain

user (@NXuser)

Details of a user, i.e., name, affiliation, email address, *etc*

Administrative Domain (GDPR? DOI?)

NEXUS: SILX VIEWER

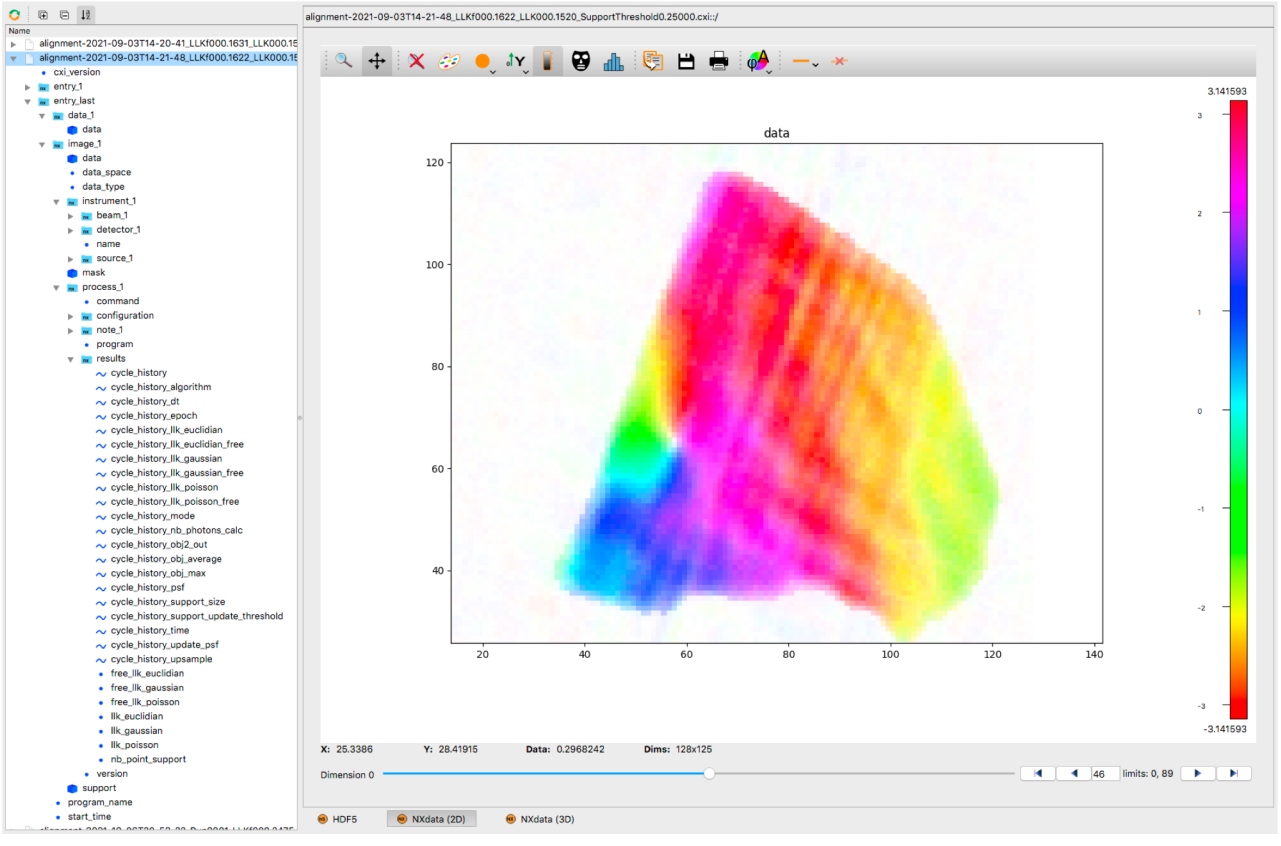

The silx Toolkit

Scientific Library for eXperimentalists

http://www.silx.org/

Standard tools & widgets for data analysis and display.

Automatic rendering of NeXuS files

DATA ACQUISITION: E-LOGBOOK

“The ultimate goal of FAIR is to optimise the reuse of data”

https://www.go-fair.org/fair-principles

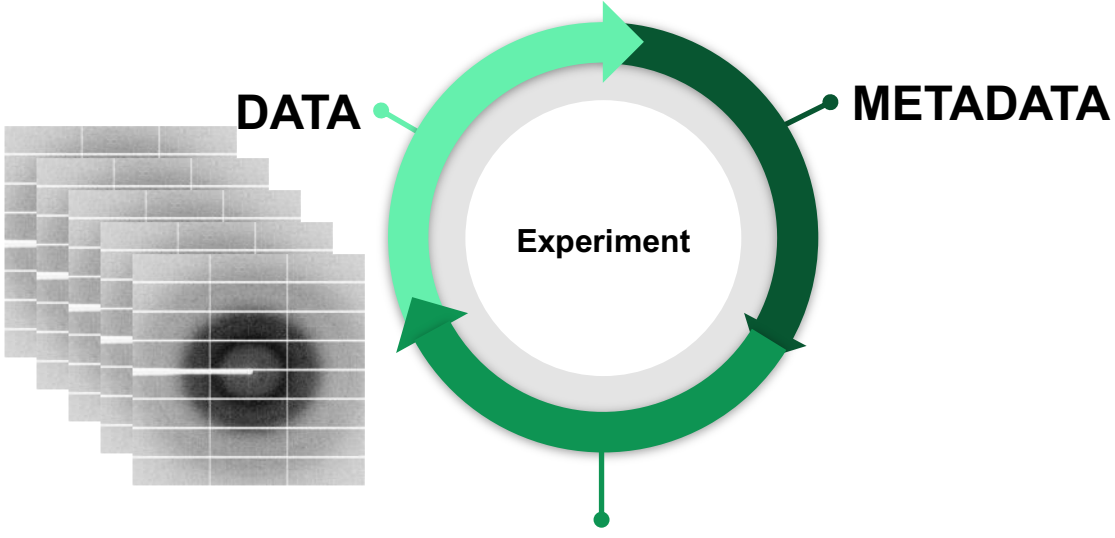

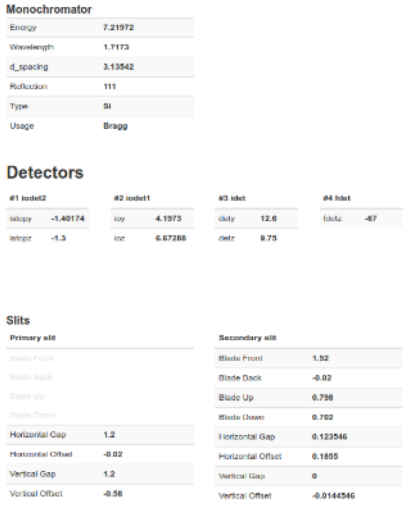

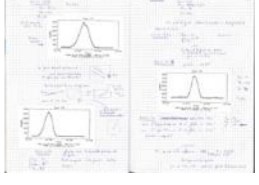

LOGBOOK

- Protocol and notes
- Decision-making
- Observations and/or conclusions

E-LOGBOOK REQUIREMENTS

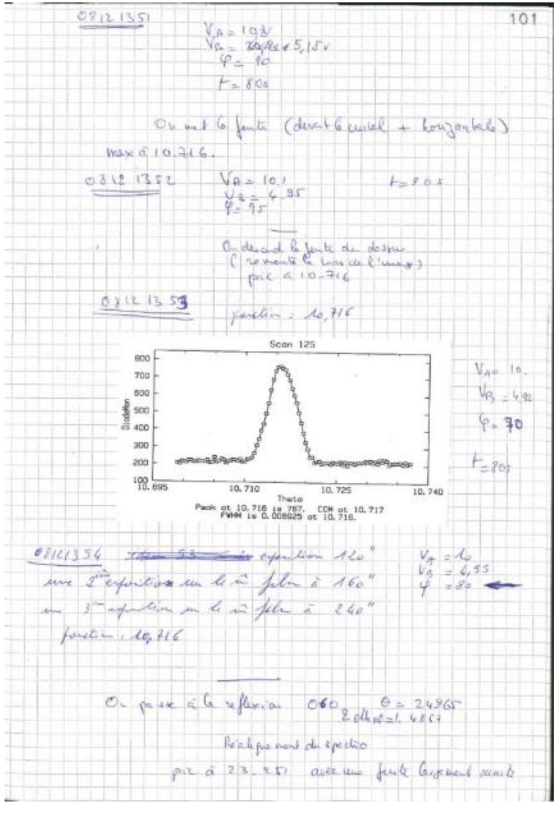

- A logbook is a **list of entries** (logs or events)
- **Time stamped in chronological** order and **sequential**
- Support rich text and **WYSIWYG** editor
- A logbook can be attached to: beamline, proposal, **session**, dataset, sample, datafile
- **Upload** images and files

Plus:

- Control access. User management and security
- API exposed to be used by users or software (SPEC, Bliss, etc..)
- Support different types of logs. For instance: info, error, debug

E-LOGBOOK

00:16:17 OPTICS> New dataset: fe2streptor2_main_root

November 4th 2018

23:59:17 OPTICS> zapxiaimage samy 8.064 7.124 94 samz 25.426 26.056 63 60 0 (zap: #1, spec: #1)

23:58:41 OPTICS> New dataset: fe2streptor2_coarse

23:58:38 OPTICS> New sample: fe2streptor2 ("Fe2 strepto replica 2")

23:58:38 OPTICS> New spec file: /data/visitor/ev280/id21/fe2streptor2/fe2streptor2_spec01/fe2streptor2_spec01.dat

Notify commands

Annotations

- Created by humans

Logbook Usage for period **2021-10-01** and **2021-10-10**.
Change url parameters, startDate and endDate to change manually the range of dates. Example: /manager/starts/logbook?startDate=2020-08-01&endDate=2020-09-01

Number of entries in the logbook per investigation

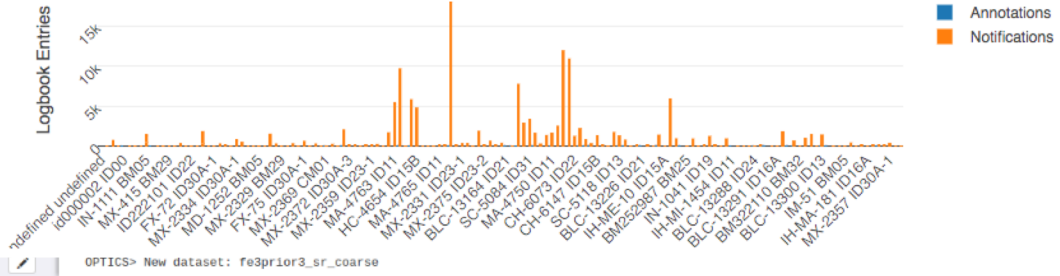

General

Instruments	37
Investigations	153
Average Events/Investigations	980
Max. Annotation/investigation	122
Max. Notifications/investigation	18088
Total Events	149912

by software:

C

y be commented

22:55:29 OPTICS> New dataset: fe3prior3_sr_coarse

European Synchrotron Radiation Facility

ESRF DATA POLICY : GUIDING PRINCIPLES

1. ESRF is the custodian of raw data and metadata
2. ESRF will automatically collect metadata for all experiments
3. ESRF will store metadata in a metadata catalogue (ICAT)
4. High level metadata (Title, Authors, Beamline, Abstract, Experiment Report) will be published as soon as possible
5. Experimental team has sole access to the data during the so-called embargo period of 3 years; can request to extend embargo period
6. After the embargo ESRF will make the data Open Access
7. Users need to create an identifier to get Open Access data
8. Proprietary data belong by default to the PI and are not kept

https://www.esrf.eu/datapolicy

DATA PORTAL

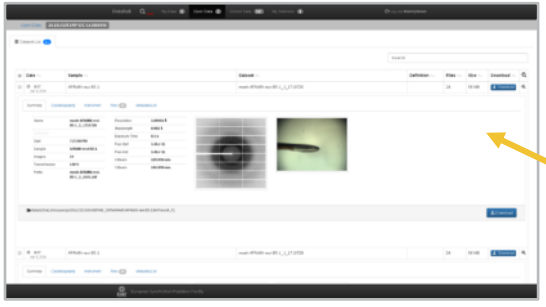

Data Catalog
https://data.esrf.fr

Data Service
Explore and Download data

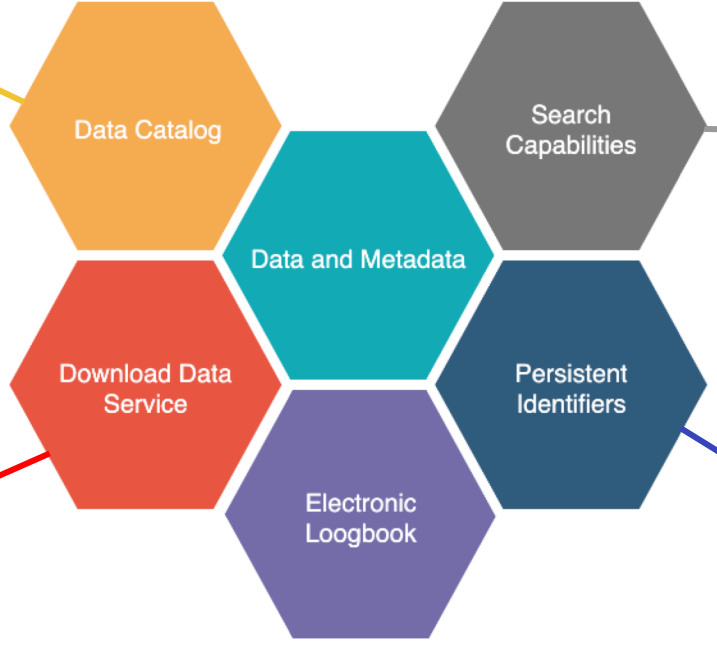

Search Service
Search engine for Big Data

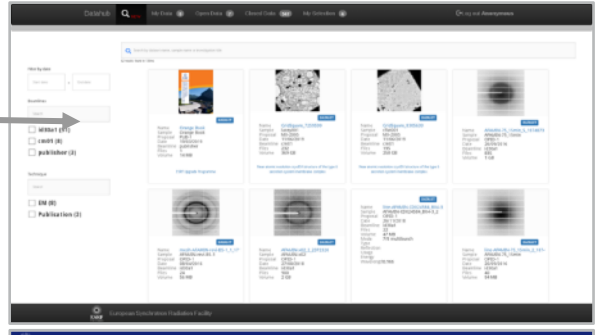

Persistent Identifiers
Make your data findable and searchable

THE DATA PORTAL

https://data.esrf.fr

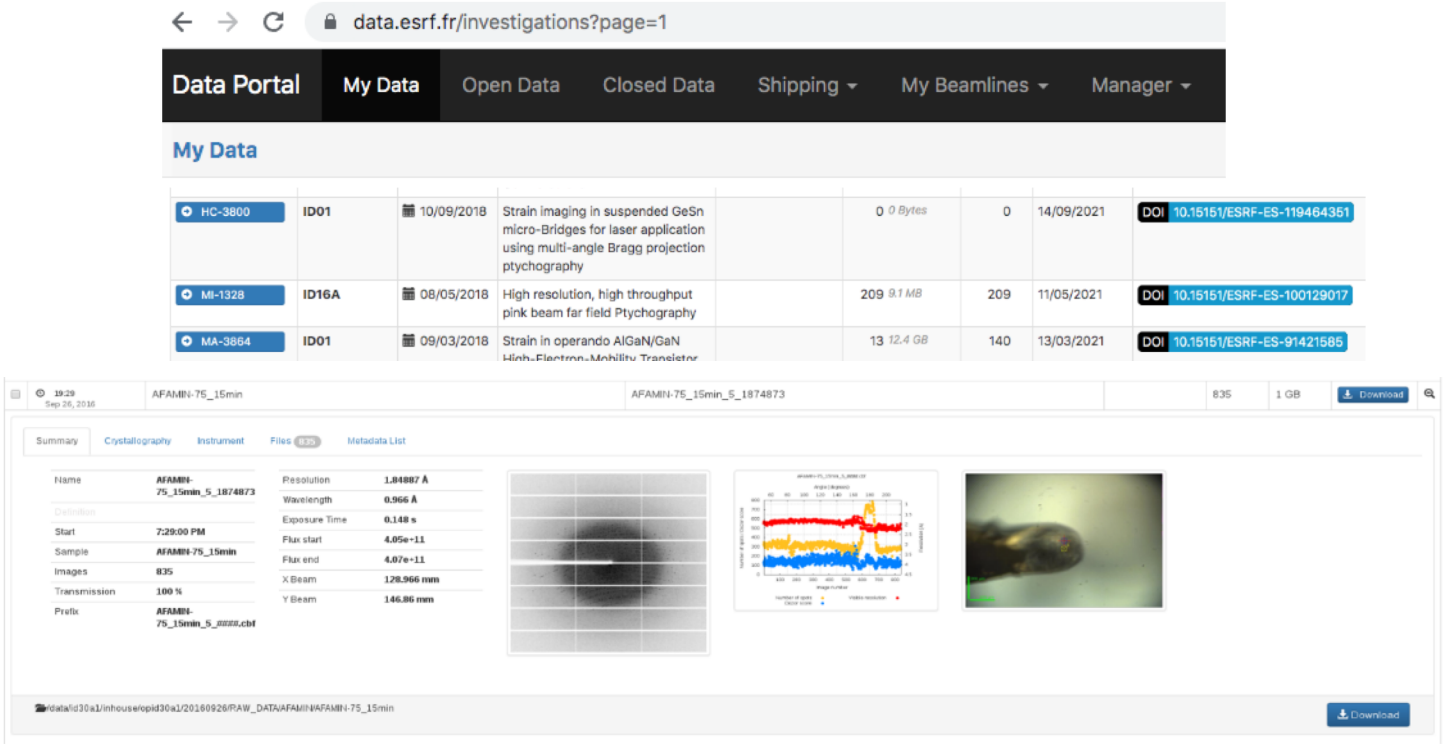

DOWNLOADING LARGE DATA: GLOBUS

For users that want to download large volume of experimental data **(>2GB)**, ESRF users can access the Globus service

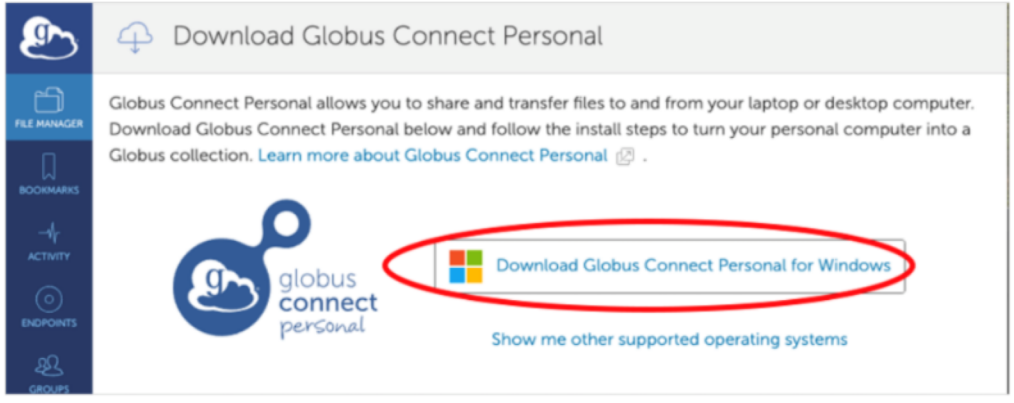

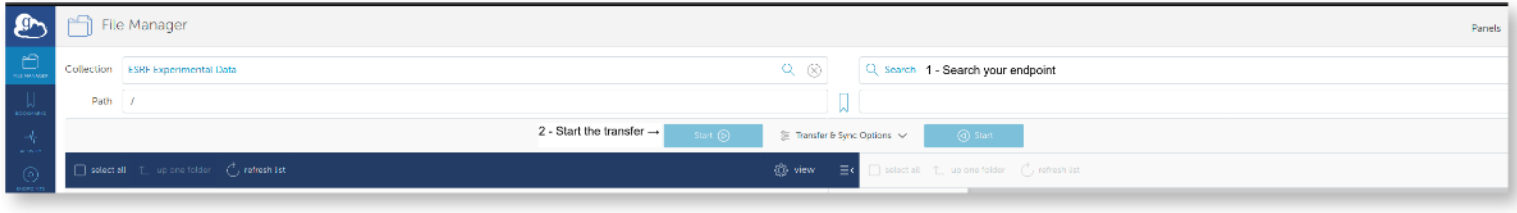

- The service has just opened this week for all users and all data
- This is made possible as data access is protected using Access Control Lists (ACLs) on the storage – users cannot see others data !

DATA PORTAL STATISTICS

Since last 7 days

Summary

Datasets	**930**
Beamlines	**22**
Total Volume	**14.9 TB**
Total Number of files	**854741**

Dataset

Average file count	**919**
Max files	**12683**
Average volume	**16.5 GB**
Max volume	**2.0 TB**
Average metadata	**32.4**

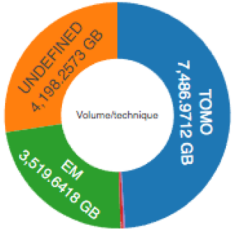

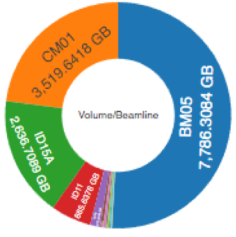

All

Summary

Datasets	**1000853**
Beamlines	**47**
Total Volume	**5.1 PB**
Total Number of files	**373111077**

Dataset

Average file count	**373**
Max files	**200002**
Average volume	**5.4 GB**
Max volume	**8.1 TB**
Average metadata	**26.8**

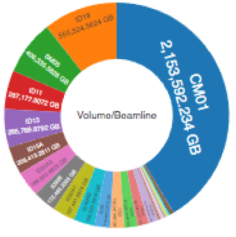

DATA PROCESSING & ANALYSIS

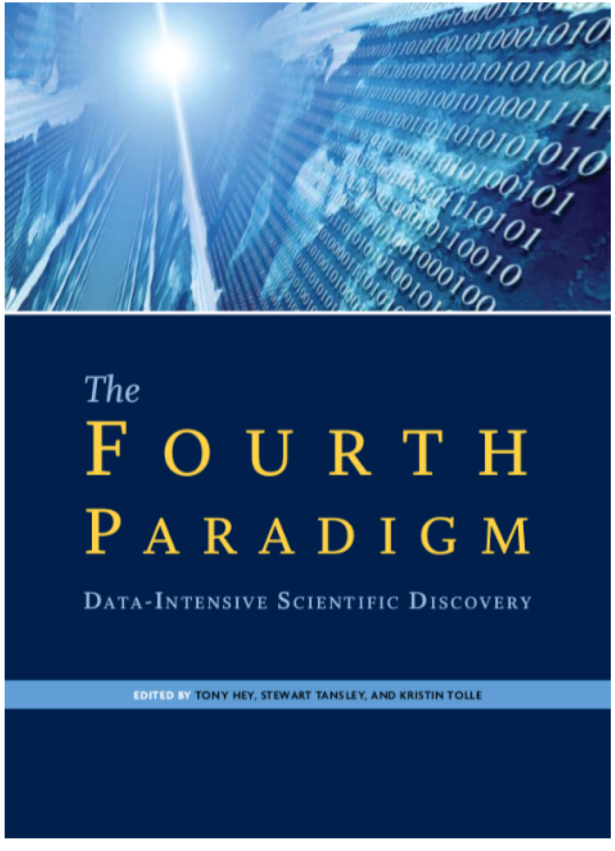

Need to ensure that **data analysis tools**:

- Cover all **data-intensive** techniques
- Are exploited by a **larger user community**
- Are actively **maintained** & improved

Need for **more computing resources**:

- Single beamlines exploiting the ESRF-EBS photon flux can saturate the existing CPU+GPU cluster

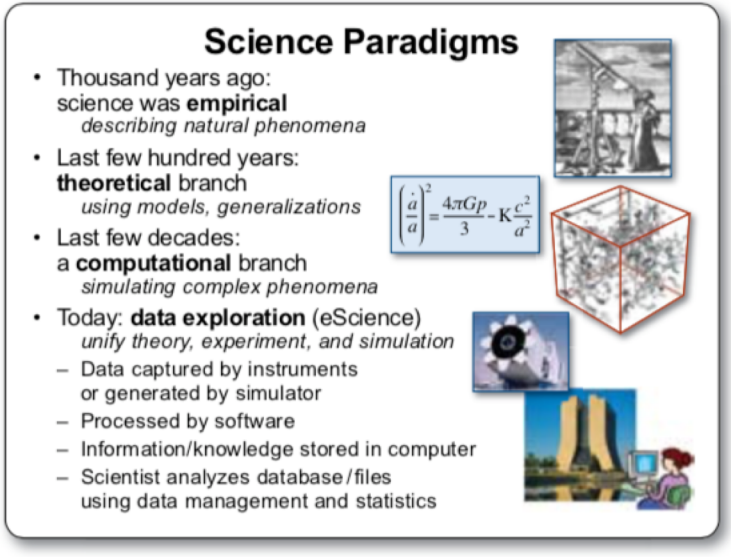

Single-algorithm analysis of of datasets is not enough !

Scientists need tools to **explore & interpret experimental results**.

CHALLENGES: STREAMLINE PROCESSING

ISPyB

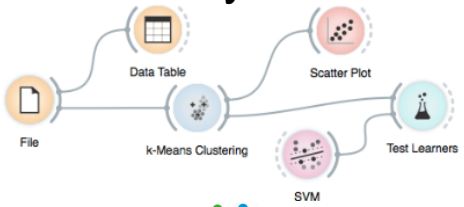

STREAMLINE

Many experimental techniques can be processed by complex but now robust algorithms:

- Macromolecular Crystallography (since early 2000s)
- Tomography
- Coherent imaging
- Scanning fluorescence
- Powder Diffraction
- SAXS
- …

These techniques are already standardised as part of the data policy implementation

A **workflow-based analysis** is often possible:

- Unsupervised & automated data processing
- Users need not be expert in the technique (but know the limits & can handle indicators)

Free users from purely technical aspects of data analysis:

- When the science is in the samples, not the technique
- Faster data analysis… and publication
- Allow a **larger community** to exploit those streamline-capable techniques.
- **Industrial-users** friendly

DATA ANALYSIS AS A 'SERVICE' ?

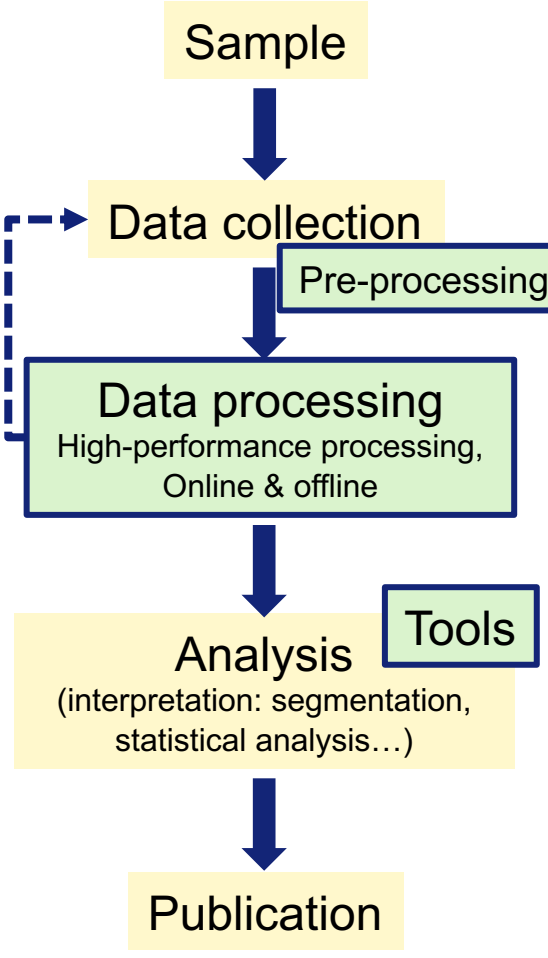

What Data Analysis as a Service <u>is not</u>:

- Doing the most of the analysis for all users
- Users with no understanding the fundamentals & limits of the techniques

What DAaS <u>is</u>:

- Handle data processing <u>when</u> it can be automatized
- Provide tools (software, computers, GPU) for **online data processing** ; enable good decisions during experiments
- Provide tools and computing resources (when possible) for **offline analysis**
- Make data available (data policy)

- For <u>selected</u> techniques: enable users to **focus on interpretation & scientific results** and not be overwhelmed by data handling issues

ONLINE DATA ANALYSIS

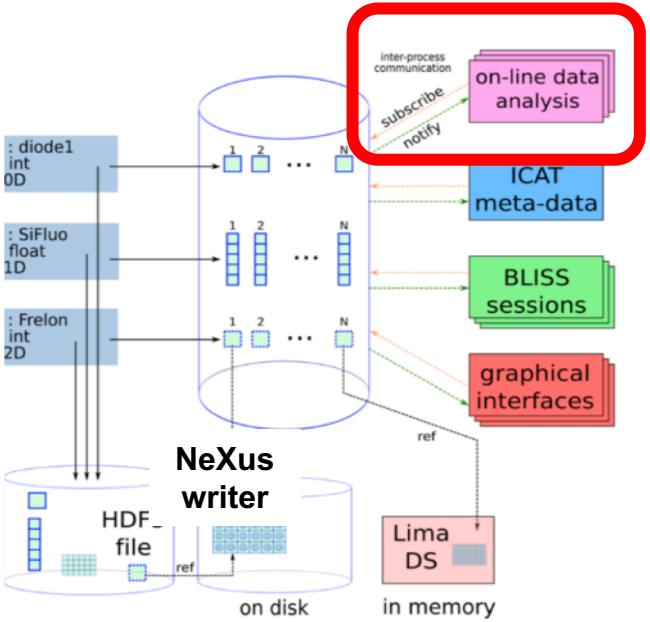

Before:

- '**one-way pipeline**': acquisition, measure, display, store, analyse

Now:

- **Need (more) online analysis** for user decision
- Automated analysis-driven acquisition
- Analysis procedures are not pre-defined except in a few cases (MX): need to have algorithm-development friendly interfaces
- **Scientist community python-educated**: provide standard access so all scientists can easily access & work on data

Development topics (TID + ISDD + Exp.Div):

- **On-the-fly (in-memory) data access** (up to 16 GB/s) for analysis/triage/compression… (see: Memcached, ASAP:O, Bluesky's DataBroker)
- **Workflows**
- **User interfaces** (**web**,..)

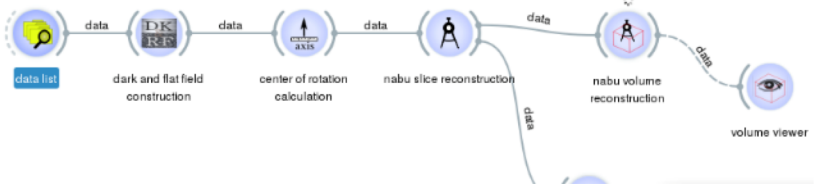

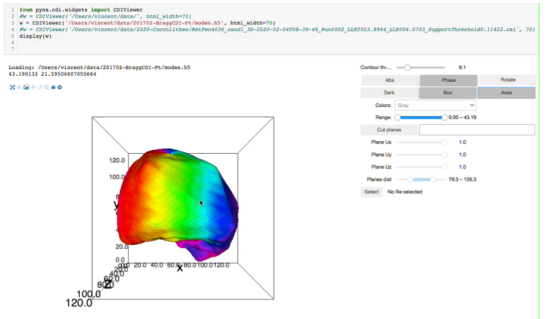

HPC USE CASES: EVOLUTION

- **Machine learning**:
  - **Big data** from scattering / spectroscopy e.g. millions of one-dimensional patterns (from a map or time series)
  - **Automated segmentation of 3D volumes** (up to 16k x 1.6k x 1.6k), or series of volumes (fast tomography can yield *hundreds* of 2000**3 datasets per hour)
- Increasing need for **GPUs** (ML, Fourier & Radon transforms, etc..)

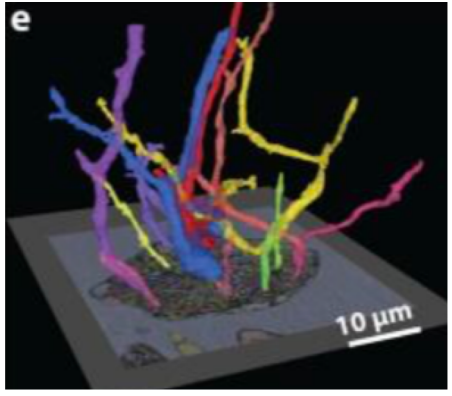

- **Web interface/jupyter notebooks**:
  - Interactive computing (often not an efficient use of resources but extremely practical)
  - Portable remote UI (manage/view jobs)

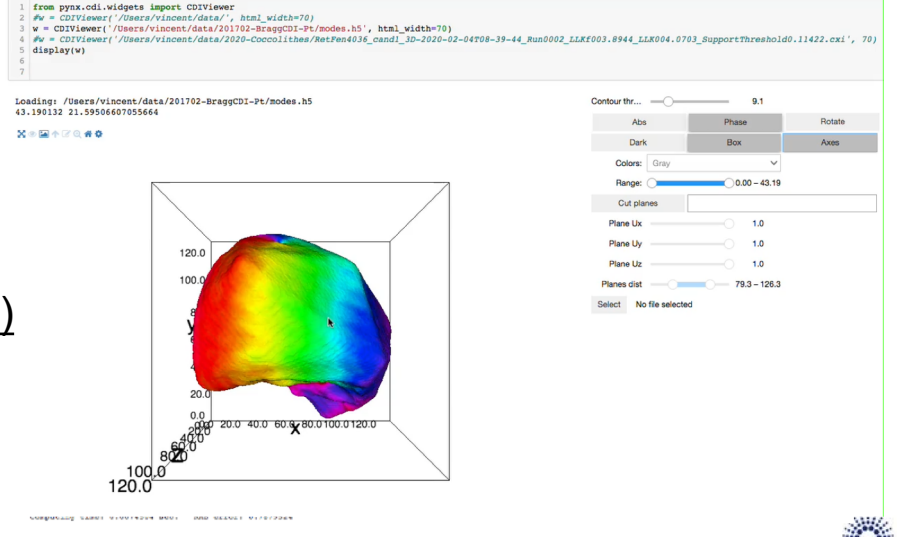

- Future need:
  - remote 3D view (+ segmentation) – learn from gaming platforms ??

DATA ANALYSIS CHALLENGES FOR LEAPS

Facilities can't provide software for *all* analysis.
Community is important (50% of analysis packages) !

Need to train for remote/cloud computing:

- **Users**
- **Software scientists**
  - New interfaces
  - GPU+CPU+slurm

ExPaNDS
panosc
photon and neutron open science cloud

http://jupyter-slurm.esrf.fr
(free access for all ESRF users)

- Also need **standards** (virtual images) for cloud software… or developments won't happen !

VISA

Web/notebook

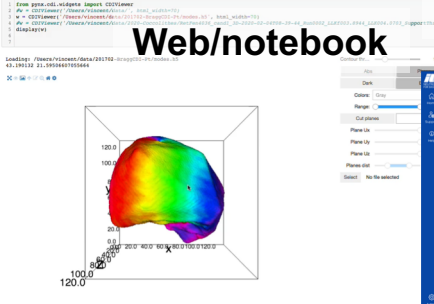

Remote Desktop

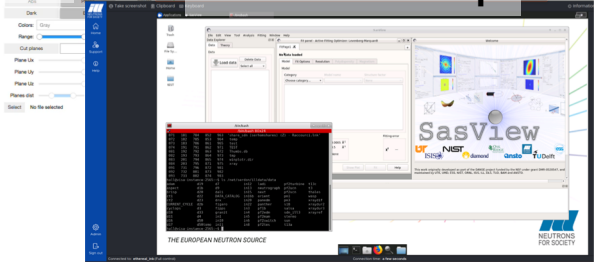

Facilities can't provide compute power for *all* analysis.

Remote/cloud/high-performance computing has become a **basic commodity** for all researchers.

Must users apply for extra compute time, e.g. with EOSC ? More 'paperwork' which will slow down adoption of modern computing.

It should be time to provide a basic set of cloud computing to every researcher – *no questions asked*.

Benefits:

- Faster adoption of cloud resources
- Faster education to modern computing (machine learning..)

Who can pay for this ?