



ESRF Data Policy, Storage, and Services

R. Dimper, A. Götz, A. de Maria, V.A. Solé, M. Chaillet & B. Lebayle

To cite this article: R. Dimper, A. Götz, A. de Maria, V.A. Solé, M. Chaillet & B. Lebayle (2019) ESRF Data Policy, Storage, and Services, Synchrotron Radiation News, 32:3, 7-12, DOI: [10.1080/08940886.2019.1608119](https://doi.org/10.1080/08940886.2019.1608119)

To link to this article: <https://doi.org/10.1080/08940886.2019.1608119>



© 2019 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 23 May 2019.



Submit your article to this journal [↗](#)



Article views: 440



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 2 View citing articles [↗](#)

ESRF Data Policy, Storage, and Services

R. DIMPER, A. GÖTZ, A. DE MARIA, V.A. SOLÉ, M. CHAILLET, AND B. LEBAYLE
 ESRF, Grenoble, France

Introduction

The ESRF was the first synchrotron to adopt a data policy implementing principles that today are known as the FAIR principles (Findable, Accessible, Interoperable, Reusable) [1]. The ESRF Data Policy essentially aims to collect all metadata and keep them forever, and to archive all raw data for 10 years. This places a heavy burden on computing and human resources; before it was proposed, it was considered almost “impossible.” Since the ESRF adopted its data policy, a number of the European synchrotrons have followed, and it is now considered not only good practice, but essential, to implement the FAIR principles for scientific data management. This article retraces the steps ESRF followed to implement the data policy, the current status, and the implications for data storage and services.

Timeline

The job of defining, adopting, and implementing a data policy like that of the ESRF is a long process. The main milestones are summarized below. It is impossible to implement such an ambitious project with wide-ranging political implications without strong backing from top management. Each milestone represents a significant amount of work by dedicated technical staff.

The following milestones and timeline can be identified along the road to implementing the ESRF data policy:

2010: Data policy framework: The PaNdata FP7 project produced a framework for implementing a data policy for photon and neutron

sources. The framework, which was the result of reflection from all of the European photon and neutron sources, was an essential icebreaker in paving the road to defining an ESRF data policy. It addressed the essential but thorny questions of data ownership, metadata, embargo period, etc.

2015: Data policy adoption: The ESRF Council approved the ESRF data policy. The adopted data policy was based on the PaNdata data policy framework, which was tweaked to take into account the discussions with ESRF staff, directors, and the Scientific Advisory Committee. Small modifications, including in-house experiments, data under the Creative Commons CC-BY licence, and special treatment of proprietary research, were added.

2016–2017: Technology choices: A prototype was developed to evaluate different technologies as part of the CRISP FP7 project. As an outcome of this project, it was decided to adopt the ICAT [2] data catalogue developed by STFC (an Oracle database), ActiveMQ queues to batch the metadata, Apache Camel for the workflow, and a Tango device server as interface for the beamline control system.

2017–2018: Data policy implementation: Half of the ESRF beamlines had implemented the data policy [3]. Implementing the data policy on a beamline entails discussing and convincing scientists, defining metadata for the experimental techniques of the beamline, (re)organizing the structure of the raw data on storage, identifying and collecting the metadata on the beamline, defining what a dataset is, and modifying the data acquisition scripts to trigger the collecting of metadata. The pipelines for registering the metadata, curating the raw data, and

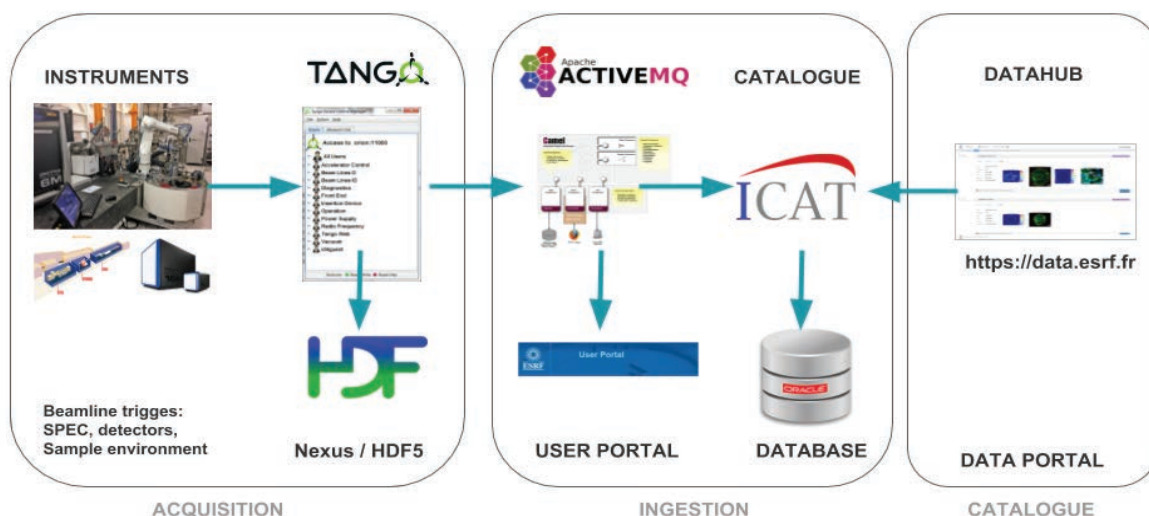


Figure 1: Metadata and data flow at the ESRF.

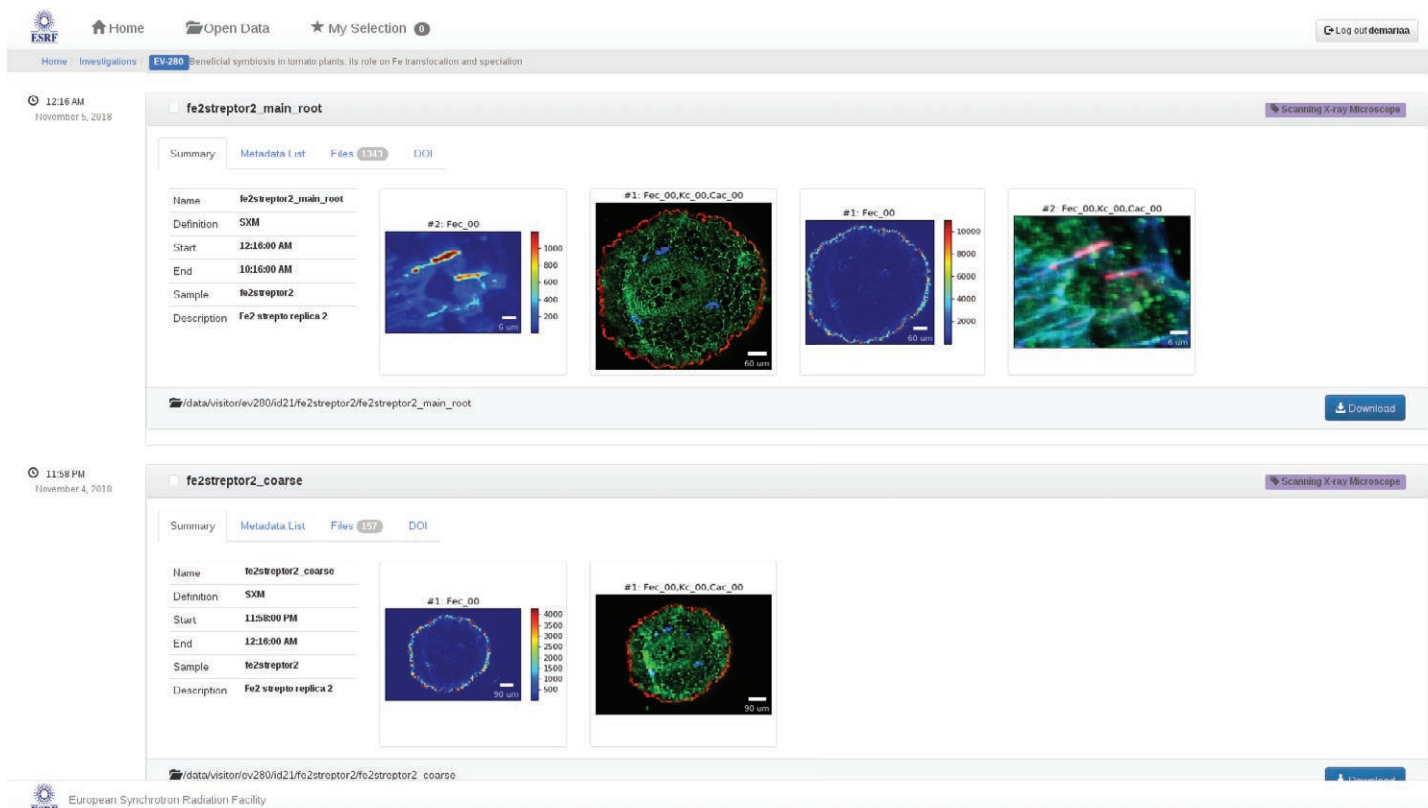


Figure 2: Example of an image gallery for datasets of a cross-section of a tomato plant root (courtesy of Dr. Ana Gonzalez Franco of the Universidad Autonoma de Chihuahua) [16].

archiving were developed during the same period. A first version of the electronic logbook (e-logbook) was also developed.

2019–2020: Full data policy implementation for the Extremely Brilliant Source (EBS) project: In 2020, it is planned to have all beamlines implementing the data policy. The data policy will be implemented in BLISS, the new ESRF beamline control system. The development on the e-logbook and data portal will be completed. All of these developments will be completed just in time for the new ESRF storage ring, EBS.

2021: Data services: From 2021 onwards, ESRF will provide multiple data services for ESRF users and open data for the European Open Science Cloud (EOSC) users at large.

Data policy

Driven by the need to make research data open, the European photon and neutron facilities came together in 2008 for the PaNdata FP7 project to define and, if possible, adopt a common data policy. However, at the end of the PaNdata project, only ISIS and ILL, two of the leading neutron sources in Europe, succeeded in adopting a data policy derived from the PaNdata policy framework published in 2011. It took another four years until the ESRF was able to convince all of its stake-

holders to approve and start the implementation of the first data policy for a photon research infrastructure in Europe, which, in turn, allowed other photon research infrastructures to follow suit.

The main elements of the data policy describe data ownership, data curation, data archiving, and access to open data. The first and most important provision of the data policy is its acceptance conditions prior to the awarding of beam time. The policy covers all types of beam time, on all ESRF and CRG beamlines, with the sole exception of proprietary research. The ill-defined issue of data ownership is left open by declaring that ESRF fulfills the role of custodian and, as such, does its utmost—within reasonable technical and financial limits—to keep all raw data produced in the facility for a duration of 10 years and all metadata forever. Data will be kept in a well-defined data format and only data with metadata generated by ESRF software will be archived. Each experiment and data set will have a unique and persistent digital identifier, which must be quoted when publications are based on open access to the data. Another important provision of the data policy is that raw data and the associated metadata are restricted to the experimental team for an embargo period of three years after the end of the experiments. This embargo period can be further extended by submitting a written request to the ESRF Directors of Research. Otherwise, the data will become openly accessible.

Implementation

At the ESRF, software has been developed to implement the data policy. Each beamline has a TANGO [4] device server installed, which manages the data and metadata related to an experiment and triggers the ingestion after each data acquisition scan (Figure 1).

Users are required to provide the proposal, technique, sample or dataset name and sample or dataset description at the beginning of each acquisition cycle. Once this is done, the capture of metadata and data happens automatically.

Once a dataset has been collected, two main actions are performed. First, an HDF5 [5] master file is created containing the metadata of each dataset. Metadata follow the Nexus conventions [6]. If the raw data are produced in HDF5, the data are linked from the master file using external HDF5 links.

Second, the metadata of the dataset are serialized in XML and sent to a cluster of messaging servers [7] using the protocol Stomp [8]. The ingestion process preserves metadata in a database forever and the raw data are archived for 10 years (see the Storage section for more detail).

In the ingester, messages are processed using Apache Camel [9], which allows us to define different routes and rules for each dataset. Every dataset is enriched with information from the user portal and stored in the ICAT metadata catalogue, allowing us to link proposal, user names, datasets, and samples consistently.

A Digital Object Identifier (DOI) is minted for every peer-review and in-house research proposal for each contiguous beam time session. This is done automatically without user intervention. The DOI landing page with information about the proposal title, authors, and beamline is made public during the session. ESRF uses Datacite [10] for minting DOIs. In addition, users can mint bespoke DOIs for individual or groups of datasets of a session or multiple sessions, thereby making them accessible from publications and search machines (the A in FAIR).

A web user interface based on the React Javascript library [11] has been developed and allows users to follow data production in real time with metadata displayed and data available for downloading. A gallery feature is implemented to associate images—photos, plots, or results—to a dataset (Figure 2). Metadata are stored in Elasticsearch to make them more findable (the F in FAIR).

Rich metadata

In order to make data interoperable (the I in FAIR), metadata are stored following the Nexus conventions [6]. All experiments from all beamlines therefore produce an HDF5 file using a common Nexus schema. The schema is shared among all beamlines. It contains a generic part, which applies to all beamlines and which describes beamlines and environment configurations, and a beamline-specific part, which strongly depends on the technique used for carrying out the experiment.

Currently, this schema supports the following techniques: tomography, fluorescence, kmap, crystallography, electron microscopy, ptychography, and microbeam radiation therapy (Figure 3).

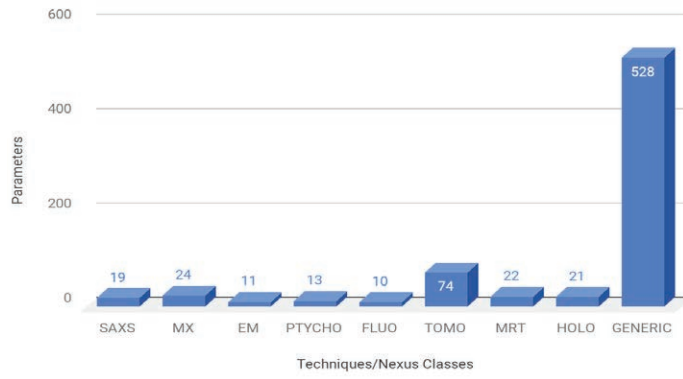


Figure 3: Number of Nexus metadata parameters per technique.

E-logbook

An e-logbook has been developed and deployed at ESRF to make data more FAIR compliant. It keeps track of the context in which data and metadata have been collected, thus enabling a better understanding of how data were gathered. This is essential for members of the experimental team and others to be able to reuse the data (the R in FAIR). The e-logbook contains a list of time-stamped events, which are either sent automatically from beamline devices (errors, debug information, command lines) or written manually by the experimental team in the form of comments (Figures 4 and 5).

In compliance with the ESRF data policy, the e-logbook metadata are permanently stored at ESRF. During the embargo period, the e-logbook entries remain exclusively readable and editable by the experimental team, thus allowing further enrichment or modifications before the logbook content becomes publicly accessible. This enables scientists to curate the logbook before it becomes publicly available. The e-logbook can be accessed with any modern web browser. The web front end, accessible at <https://data.esrf.fr>, is a web GUI (based on the React Javascript library) for all metadata, including the logbook. A back end, the so-called ICAT+, exposes a RESTful api via a NodeJS web server, which stores logbook data (messages and files) in a MongoDB database. ICAT+ was designed as a plugin for ICAT, making use of the existing authentication and authorization mechanisms. The source codes for both the ICAT+ and datahub projects are available on the ESRF gitlab [12].

Data storage

Experimental data need to be stored on disks for processing by scientific software. Because disk storage is expensive, we have chosen to implement large central storage systems, allowing for maximum flexibility in the allocation of disk space while simplifying the system administration. Again for cost reasons, only the last experiments are stored on high-performance storage systems and, as of 2018, experimental data are removed from disks 50 days after the end of the experiment, leaving only a backup copy on tape storage. For safety, data are duplicated on tape libraries located in two separate buildings.

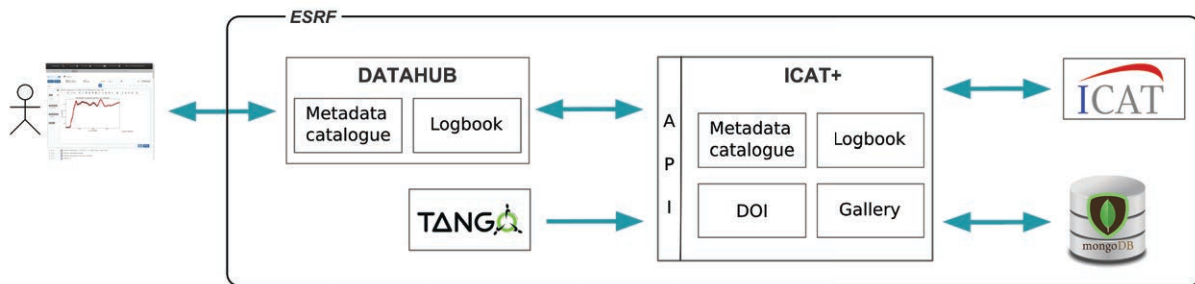


Figure 4: e-logbook architecture.

Figure 5: Example of editing a user comment in the e-logbook.



FalconX8

POWERED BY
SITORO
ACCELERATED ANALYSIS

The World's Most Advanced Digital Pulse Processor for high-rate Spectroscopy

- Up to 8 channels of simultaneous processing
- Over 3 Mcps processed by each channel
- 12% dead time at 1 Mcps
- Exceptional resolution
- Portable, plug & play, minimal setup, intuitive interface

sales@xia.com
Tel: +1-510-401-5760
www.xia.com

XIA
Instruments That Advance The Art

In 2020, because all beamlines will implement the data policy where experimental data are archived automatically to tape storage, the plan is to partially abandon the daily backup to tape.

Disk storage

Disk storage needs to be fast enough to write a single thread of experimental data produced by a fast detector, typically in the range of 1 to 4 GBytes/sec. If the disk storage system is unable to sustain this throughput during the experiment, the data acquisition will fail, leading to loss of data and time. In the case of the ESRF, due to the centralization, the storage systems must be able to deal with 10 to 20 detectors in parallel; i.e., a sustained throughput of at least 20 GBytes/sec.

GPFS (General Parallel File System) has proved to be a file system providing this type of speed, especially for single-thread performance. Since 2018, we operate two independent GPFS systems of 4 PBytes net capacity each, with 40 GBytes/sec bandwidth for each. Tests have shown that single-thread performance can reach up to peaks of 5 GBytes/sec.

With the expected growth of data production once the EBS storage ring and the beamlines become operational, the main issue will be that large disk storage systems are often much faster for writing than for reading data. Data reduction and analysis of experimental data are performed by large compute clusters, resulting in hundreds of parallel read streams. These massive read accesses saturate disk storage systems, and

this will inevitably lead to a situation where write access throughput can no longer be guaranteed and data acquisitions will fail. Until now, the problem has been circumvented by grouping all access from compute clusters via NFS-to-GPFS gateways with limited bandwidth to the storage systems, thus leaving sufficient bandwidth for writing data. A direct consequence of this approach is that compute clusters spend most of their time waiting for data, which in turn slows down data analysis. As a direct consequence, the 50-day limit was often insufficient and had to be extended.

In 2020, we will add an in-memory cache which will act as a write and read buffer to allow fast acquisition and online data analysis. Data will then be migrated in the background with the smallest possible time delay to the general storage systems for analysis and archiving. This option will be thoroughly evaluated before making it available to beamlines in 2020. Another important upgrade will then allow the compute clusters to have much faster access to the data via an Infiniband network.

Tape storage

Tape is still the most economical and reliable data storage media. In data centers, tapes are organized in large libraries, with robotic arms fetching the tapes and transferring them to tape drives for reading and writing data. The number of robotic arms and tape drives determines the degree of parallelism a tape library offers.

**ULTRATHIN
METAL
FOIL**
0.1-25 μ

Filters
Windows
Targets
Photocathodes
Vacuum Tight
Be Windows

LEBOW COMPANY
Goleta, California
Fax & Phone 805/964-7117
www.lebowcompany.com

The two tape libraries used at ESRF are StorageTek (now ORACLE) SL8500s, each capable of holding up to 8375 tapes, 64 tape drives, and eight robotic arms. Tapes are LTO-7 and LTO-8 with 16TB of data per tape and 360 MB/sec read/write access for the latter. The backup software successfully used at ESRF for more than 20 years is Time Navigator (TiNa) from Atempo. TiNa allows accessing data based on time stamps. The archiving functionality is based on in-house software. In 2020, we are considering using the ADA product for backup and archiving.

Two issues need further attention: the very high number of small files still generated by some of the ESRF beamlines must be encapsulated into larger (HDF5) files; otherwise, it will be impossible to keep up with writing them to tape. The recent decision of ORACLE to discontinue the StorageTek product line forces us to review the situation for future investments.

PaNOSC and the EOSC

Science is about reproducibility and being open. ESRF already has an Open Data policy combined with an Open Source software policy. However, besides the scientific goal of reproducibility of results, there are practical considerations that suggest that data services need to go beyond simply downloading data and searching metadata.

ESRF users do not always have the necessary infrastructure to deal with the volume and/or complexity of the data. Ultimately, we would like to be able to provide not only data, metadata, and software, but also appropriate infrastructure for data analysis. Initiatives like the H2020 Calipsoplus [13] and PaNOSC [14] projects go in the direction of implementing an Open Science Cloud and linking up to the European Open Science Cloud (EOSC). The ESRF is coordinating the PaNOSC project and co-leading the Data Analysis as a Service (DAAS) prototype in Calipsoplus. ESRF and ALBA are developing a prototype of a common portal for Data Analysis as Service (DAAS) across multiple synchrotrons in Europe as part of the Calipsoplus project.

PaNOSC is an ambitious 12 M€ project aimed at laying the foundations of an Open Science Cloud for photon and neutron facilities in Europe. It will provide a common API to allow data discovery at different facilities and to share DAAS solutions. For the latter, solutions around JupyterHub, containers, and remote desktops in the browsers (based on Guacamole) have been adopted as the base technology for providing ready-to-use solutions for scientific fields studied at the ESRF.

Conclusion

The adoption and implementation of the data policy are progressing well and constitute a cornerstone for FAIR data management at the ESRF. With the exponential growth of the data produced and the rapidly increasing complexity of the data analysis and the required IT infrastructure, it is essential to help the scientific community with professional data management on which other services can be implemented. Data management will ultimately determine the extent to which the ESRF Extremely Brilliant Source can be used in the most efficient way. With many other synchrotrons following in the footsteps of the ESRF for similar upgrades of their storage rings and beamlines, many developments will be discussed and carried out jointly in Europe and beyond in the frame of the LEAPS [15] initiative. The most visible expression of this collaborative approach will materialize in services accessible through the EOSC. ■

Funding

This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under Grant Agreement No. 823852.

References

1. M. D. Wilkinson et al., *Scientific Data* 3, 160018 (2016). <https://doi.org/10.1038/sdata.2016.18>
2. <https://icatproject.org>
3. <https://www.esrf.eu/datapolicy>
4. <http://www.tango-controls.org>
5. <https://www.hdfgroup.org/solutions/hdf5>
6. <https://www.nexusformat.org>
7. <http://activemq.apache.org>
8. <https://stomp.github.io>
9. <http://camel.apache.org>
10. <https://datacite.org>
11. <https://reactjs.org/>
12. <https://gitlab.esrf.fr/icat/E-DataPortal/>
13. <http://www.calipsoplus.eu/>
14. <https://panosc.eu>
15. <https://www.leaps-initiative.eu/>
16. <https://doi.esrf.fr/10.15151/ESRF-ES-135816585>