ESRF

The European Synchrotron

ESRF

INFORMATION TECHNOLOGY DATA STRATEGY

2022 - 2026

ESRF IT STRATEGY

The ESRF Management charged the Team composed of R. Dimper, A. Götz, V. Favre-Nicolin, B. Lebayle, C. Nevo, J-F. Perrin, V. A. Solé to coordinate and compile the information presented in this paper, which was provided by colleagues from all ESRF divisions, and further developed thanks to the constructive criticism from an international expert panel and the ESRF governing bodies.

TABLE OF CONTENTS

Executive Summary

The ESRF successfully resumed User Operation with the new EBS Storage Ring since 25 August 2020. The largely increased brightness of the new ESRF Storage Ring will allow to carry out new X-ray science. The new scientific capabilities will steeply increase the data volume and must be supported by an ambitious IT investment programme to enable efficient data processing leading to higher throughput and a shorter time span between experiment and publication. The need for better and optimised software and increased data processing capabilities is basically immediate and will further grow when the four additional EBS flagship beamlines become fully operational.

The present ESRF IT data strategy is the outcome of many discussions, internally between IT specialists, the scientific staff and management, and externally with the international Expert Panel who provided a first round of valuable feedback in April, integrated in this paper, and final recommendations in October 2020. It is casted in the activities carried out during the last five years in the frame of the ESRF-EBS IT programme which has focused to enable the ESRF beamlines to start the exploitation of the EBS X-ray source by the optimisation of the existing IT infrastructure, and in parallel study the needs for a forward-looking IT data strategy presented in this paper.

This strategy paper details all required components to considerably enhance the data processing chain, and in particular the:

- IT infrastructure commensurate with the data production, including on-line data analysis (ODA) systems, buffer systems, data network, central disk and tape storage, CPU and GPU clusters, and a new data centre able to host the infrastructure,
- Compression, reduction and analysis software,
- Machine Learning (ML) in addition to and complementing conventional data analysis algorithms, and
- staff to implement the new functionalities.

The funds required over the period 2022 to 2026 to increase the IT capacity represent an increase of 75% to 95% (16 to 20M€) required to provide the ESRF with an adequate new data centre and IT hardware. Beyond 2026, an increase of the IT budget by 70% (+3M€/y) is needed to sustain the IT infrastructure, including the additional electricity consumption.

19 additional engineer positions, hired gradually during the period 2022-2026, representing an increase of 38% of the present staff dedicated to IT activities, are needed for developing, installing and maintaining the necessary software, operating the IT infrastructure and providing high-level user support.

Abbreviations and Definitions

AI Artificial Intelligence - the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, and decision-making

API Application Programming Interface

ARIE Analytical Research Infrastructures of Europe

AWS Amazon Web Services – a comprehensive and broadly adopted cloud platform

BeeGFS (formerly FhGFS) is a parallel file system, developed and optimized for high-performance computing. BeeGFS includes a distributed metadata architecture and is designed for data throughput

BoD Board of Directors - the top-level management of the ESRF, currently consisting of a Director General, Director of Administration, Director of the Accelerator and Sources Division, two Scientific Directors, and the respective Heads of the Instrumentation and Services Development Division and Technical Infrastructure Division

CIO Chief Information Officer - the Director/Head of Division of the IT activities in a company

CISO Chief Information Security Officer - senior-level executive for establishing and maintaining the enterprise's vision, strategy, and programme to ensure information assets and technologies are adequately protected

CRG Collaboration Research Group - an independent third-party contractor operating one or several bending-magnet beamline(s) with an independent research programme. In return for the photons they receive from the ESRF storage ring, CRGs make beamtime available for ESRF's public user programme

DaaS Data as a Service sometimes also used for Data Analysis as a Service (DAaaS)

DOI Digital Object Identifier – a string of numbers, letters and symbols used to permanently identify an article, document, or data set and link it on the web

EC European Commission

EOSC European Open Science Cloud

EPN Campus European Photon and Neutron campus, situated in Grenoble,on the so-called "Polygone Scientifique" and regrouping ESRF, ILL, IBS, and the EMBL Grenoble outstation

ESFRI European Strategic Framework for Research Infrastructures (www.esfri.eu)

Gbps Gigabits per second. Common unit for high performance network connections. 1 GB/s = 8 Gbps

GB Gigabytes

GPFS General Purpose File System (brand name IBM Spectrum Scale) - high-performance clustered file system software developed by IBM and used in many HPC environments among which the Summit Supercomputer at Oak Ridge National Laboratory

TB Terabytes

PB Petabytes

HPC High Performance Computing

ISDD Instrumentation Services and Development Division at the ESRF

IT Anything related to computing technology, such as networking, hardware, software, the Internet, or the people that work with these technologies.

LEAPS League of European Accelerator based Photon Sources (leaps-initiative.eu)

LENS League of European Neutron Sources (lens-initiative.org)

ML Machine Learning is the concept that a computer program can learn and adapt to new data without human interference

NREN National Research and Education Network

NVMe Non-Volatile Memory Express, an open logical device interface specification for accessing non-volatile storage media attached via a PCI-Express bus

ODA On-line Data Analysis - ranging from basic data reduction to image reconstruction, allowing to take good decisions during experiments, does not include scientific interpretation of data

PaN Photon and Neutron - an abbreviation used in European projects indicating the RIs producing photons and neutrons for structure of matter research and working together for common solutions. About 15% of the user community is common between PaNs

PaNOSC Photon and Neutron Open Science Cloud (www.panosc.eu) - a four year project co-financed by the EC and focused on FAIR data management and data analysis services

Raw data The term Raw data in this report refers to data coming from the ESRF beamlines for storage on the central disks storage systems. This data is often already pre-processed for image and noise corrections and may also have been reduced on-line

RI Research Infrastructure

SSD Solid State Disks - devices that use integrated circuits to store data persistently. SSDs have quicker access time and lower latency but are still more expensive than spinning hard disks

TID Technical and Infrastructure Division at the ESRF

UP Acronym to refer to the ESRF Upgrade Programme

Foreword

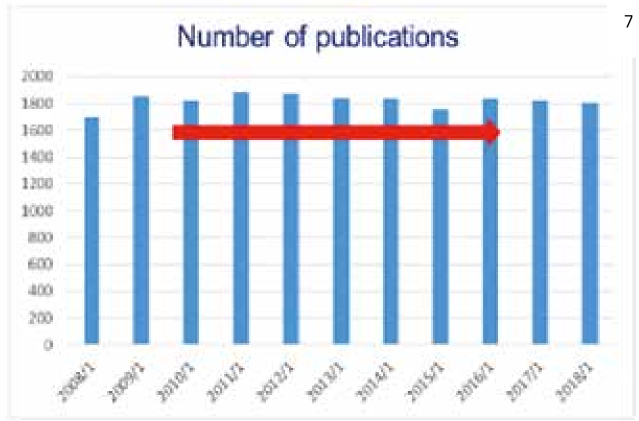

Figure 1: *Beamtime usage and data exploitation over the 2008-2018 period.*

The science case for the ESRF-EBS has been designed on the assumption that the new EBS source delivers X-ray beams of higher brightness and transverse coherence, thus shifting the frontier of synchrotron-based experiments. The five scientific drivers (X-ray imaging, structural and functional biology, science at extreme conditions, time-resolved science and nanoscience & technology), as laid out in detail in both the Purple[1] and Orange Books[2], identified new opportunities in frontier and applied research. A series of EBS workshops, including more than 15 topical workshops, helped to shape the expectations created by the EBS project. The scientific community at large is inspired and excited by the new horizons opening up in X-ray science where new X-ray microscopy techniques will allow to bridge the gap between optical microscopy and electron microscopy.

The prospect is indeed to move towards the characterisation of isolated objects down to the nanoscale thanks to recent developments in nano-beam and coherent scattering techniques such as ptychography and coherent X-ray diffraction imaging. The high photon energies, specifically available at the ESRF, lead to the next exciting objective of multi-scale structural characterisation at all relevant time scales that govern the function of organic or inorganic materials. The record number of research proposals received in April 2020 confirms the enthusiastic interest of our user community.

The optimum use of the increase in peak brightness of a (near-)diffraction-limited source requires new instrumentation including all aspects of X-ray technology from photon beam transport, sample handling and manipulation, detection of photons with the highest efficiency and resolution, to the analysis of large quantities of raw data acquired in a very short time. The upgrade of the ESRF with its Upgrade Programme (UP) Phase I and UP Phase II (EBS) has enabled not only the renewal and upgrade of a large part of the beamline portfolio but also the implementation of four EBS beamlines for cutting-edge science. Already in 2018, the impact of UP Phase I has become apparent by the enlarged user programme (compared to 2008: 42% increase of the number of proposals submitted, 28% increase in the number of experimental sessions, 14% increase in the number of user visits). However, it is worth noting that this increase in the efficiency of beamtime usage did not translate into a similar boost in publications, while the quality level of publications (as measured by the cumulated impacts factors) has indeed increased (figure 1).

This gap between scientific usage of the ESRF and the production of published scientific results has several origins:

First, more than 2/3 of the ESRF beamlines already employ ptychography, coherent diffraction imaging, serial crystallography, diffraction-contrast tomography, etc. Most of these techniques are still very new and unfortunately only very few expert user groups are able to optimally process and analyse the raw data. Data processing is usually time-consuming and software packages for new techniques are generally not user-friendly. Most of the new experiments require therefore personal involvement of the beamline scientists who cannot support, due to lack of time, all the experiments at the same level. To cite an example, a survey at the completely rebuilt ID15A beamline revealed that the beamline scientist can offer support for data pre-processing to only 25% of the experiments involving XRD-CT. Consequently, 60% of the publications from this new beamline involved ESRF contributors or are the result of in-house research. These bottlenecks lead to (i) a long lead-time in publications (average 2-3 years) or (ii) a waste of beamtime with datasets remaining unexploited.

[1] *http://www.esrf.eu/Apache_files/Upgrade/ESRF-SciTechProg2008-2017.pdf*
[2] *http://www.esrf.eu/Apache_files/Upgrade/ESRF-orange-book.pdf*

Second, the increase in brightness or photon flux density, often by several orders of magnitude, and the use of high-speed detectors enable new experiments with enhanced capabilities in terms of throughput, time resolution or size and number of samples. All of those new features result in a dramatic increase of raw data produced by most beamlines. This volume of data is further increased by the integration of imaging modalities at most beamlines as ancillary techniques. Some beamlines will be capable of producing several 100 Terabytes per day and the transfer, processing and archiving of this data is often far beyond the capabilities of the users' home laboratory. Solutions are urgently needed to ensure that data collected at the ESRF beamlines can be fully exploited. The medium and long-term danger is that beamlines have to adapt their capabilities to the existing IT infrastructure and computing bottlenecks and thus limit their full scientific potential.

Third, the output of experiments is scientific data. The beamline control system is designed to sequence and synchronise the data acquisition and to manage the data flow from the detector until it is stored for off-line data analysis. On-line data analysis is key to assess, ideally in real time, the progression of the experiments and the quality and relevance of the data. Its results are indispensable for steering the experiment to a successful conclusion. The design and implementation of the necessary data pipelines, capable of managing hundreds of terabytes, is a challenge in itself and involves all aspects of IT ranging from the data policy, scientific software, data network, storage, to name just a few. The development of generic solutions for data reduction, compression or pre-processing is particularly challenging in X-ray science considering the vast diversity of X-ray techniques and applications. Each ESRF beamline is specialized and highly customisable. The wide range of specialized use cases renders one-size-fits-all solutions impractical and requires a case-by-case approach. This is obviously extremely demanding in terms of resources and coordination. An optimal use of the resources available and the design of a long-term strategy requires not only expertise in computing science but most importantly an in-depth understanding of the science and the users' needs and expectations. We believe therefore that ultimately such a strategy must be science-driven.

In conclusion, the new EBS source properties, in combination with new instrumentation, notably high-throughput 2D detectors, will open new exciting scientific opportunities. However, as anticipated, (*cf.* chapters 4.4 and 4.5 of the Orange Book) the new capabilities will sharply increase the data volume and must be supported by a commensurate IT investment programme. It must be designed to cover all aspects of the challenge in order to enable efficient data processing, which will ultimately lead to higher throughput and a shorter time span between experiment and publication. Since 2016, this question and in particular the concept of "Data Analysis as a Service" has been regularly discussed internally but also at meetings of the SAC and Council. It became clear that a holistic approach, addressing the various facets of the problem, is needed. All the different discussions and considerations are converging to the same conclusion, that the post-EBS development of the ESRF requires an entirely new approach to handle data. This involves IT infrastructure, data policy and tools for on-line and off-line data processing and analysis. This evolution is mission critical for the ESRF, its scientific exploitation and its cooperation with the user community. The challenges are manifold and primarily originate from the explosion of the data volume, the significant increase of complexity of the data sets, and limited dedicated resources. A strategic and successful programme entails therefore not only a long-term effort but also a high level of coordination between all stakeholders.

Summary of Needs

The needs for increased IT infrastructure capacity are driven by the data produced on the ESRF beamlines. A survey was issued to obtain more detailed feedback from the beamline responsibles of the ESRF public and CRG beamlines. The survey not only allowed to collect the expected needs but also to interact with the beamline responsibles to clarify and discuss specific needs. The survey collected feedback on a total of 26 questions about future data rates, data compression rates, data processing and software needs and estimates for the years 2021 and 2025.

Based on this survey, but also on past experience, we expect the data production to increase continuously over the next years as a consequence of automation, increased resolution and speed of detectors, and more flux on the samples. All efforts will be made to reduce the data flow through on-line data reduction and analysis but despite these efforts we expect the data production to increase as given in table 1. This table summarises the raw data written to the central disk storage systems in the period 2016 to 2019 and conservative estimates until 2025. It shall be noted that the low data production in 2019 corresponds to the shutdown of the ESRF to install the new Storage Ring during which only the Cryo-EM was fully operational throughout the year. With the implementation of ten new detectors from commercial companies (*e.g.* DECTRIS) and the start of user operation on BM18 and ID29 in 2021, data production will ramp up significantly. The estimates in table 1 assume that a considerable effort is undertaken to reduce data on-line via data compression and technique specific algorithms. It shall be noted that the raw data production from the beamlines it expected to be five to seven times higher than the estimates in table 1.

Three recent workshops have highlighted the fact that IT will underpin the future of scientific activity: a workshop on deep learning for next generation X-Ray and neutron sources[3] in Copenhagen (September 2019), a workshop organised jointly with the ILL and STFC about Machine Learning and Artificial Intelligence applied to Photon and Neutron Science[4] at the ESRF (November 2019) and the satellite data management workshop organised in February 2020 during the 3-way meeting at the ESRF[5].

On-line and off-line data processing will require powerful software, including software for machine learning, on parallel computing architectures interconnected with high-speed network links, large central disk storage and data archiving systems. Designing the software, purchasing, implementing and operating the IT hardware needs talented IT staff able to constantly adapt to changing needs and new technology. The current staffing level of IT specialists is insufficient to face the numerous challenges generated by the EBS programme and described in this report. In addition, the creation of a small team for ML/AI techniques applied to photon science is proposed to implement algorithms in order to stay at the forefront of science.

The data avalanche generated by the EBS beamlines must be harnessed by on-line and off-line data processing software and hardware. If the IT infrastructure is not upgraded to cope with the data streams and to allow on-line and off-line data reduction and analysis, experiments will have to slow down, frame rates will have to be reduced, resolution reduced, duty rates lowered. The conservative estimate of data production reflected in table 1 will require that most of the beamlines are equipped with dedicated IT hardware for data compression, reduction, and on-line analysis. The overall ESRF data processing capacity needs to be increased by a factor of 10 for the GPU clusters and at least a factor of 3 for the CPU clusters.

Table 1: *Past and expected data storage on central disks*

	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025
Data production (PB)	2.8	3.4	7.9	4.7	10	20	30	40	50	60
Evolution (%)	-	+21	+132	-41	+113	+100	+50	+33	+25	+20

[3] *https://indico.nbi.ku.dk/event/1256/ - Workshop on Perspectives and Applications of Deep Learning for Accelerated Scientific Discovery at next generation X-ray and Neutron Sources*

[4] *https://workshops.ill.fr/event/209/ - Workshop on Artificial Intelligence Applied to Photon and Neutron Science*

[5] *https://indico.esrf.fr/indico/event/39/overview*

MAJOR CHALLENGES OF THE FUTURE

The IT data strategy is aimed at addressing the major known IT challenges facing the ESRF for the next 5 years. It is based on the most recent information known to the authors and which can be managed within the financial and human resource constraints available at the ESRF. This does not exclude that requirements of some beamlines might far exceed the capacity foreseen. As an example, only recently a beamline has foreseen the purchase of a new detector which could generate 65 TB of data per day starting in 2020. This would more than double the total current data production. This just goes to illustrate how quickly the needs can evolve and expectations exceed what is foreseen. The authors realize there are many challenges which will arise. They would like to highlight certain areas where they suspect things might evolve faster than foreseen in the report and will require discussion and the introduction of limits on data production and usage of compute resources and/or even more resources than outlined in this paper. The following areas might evolve faster than foreseen:

1. **Data production** – certain beamlines could generate 100s of TB per day which would require much faster on-line processing and network links to transfer data from detectors to storage
2. **Data analysis** – the quantities of data produced could exceed the storage and compute resources significantly and be a bottleneck for data analysis
3. **Data archiving** – the data volume could exceed the archiving capacity

IT-infrastructure

Overview

The main driver for the IT-infrastructure is the data production on the ESRF beamlines as given in table 1. The data production forces continuous investments in storage systems, CPU/GPU clusters, and network components. This equipment has to be installed in an environment which guarantees reliable and flexible operation with redundant power sources, cooling, fire protection, and easy access for IT-staff. The needs for fast processing of the data streams produced by high-performance detectors, as described in chapter 4, Data Acquisition, has triggered a profound re-design of the IT-infrastructure architecture.

The data centres will host all shared computing equipment such as the compute clusters and central storage, but also most of the beamline dedicated equipment for data buffering and on-line data reduction to avoid installing noisy and power-hungry equipment on the experimental floor. This requirement leads to a significant increase of IT equipment to be accommodated in the on-site data centres. The data buffering between beamline detectors and central storage is mandatory as soon as the detector system itself does not provide a large buffer storage area.

Until 2019 most of the beamlines and foremost the compute clusters were linked to central storage via an NFS gateway (figure 2). This gateway considerably reduced the data exchange and therefore protected the central storage from I/O saturation. However, this also slowed down the compute clusters for accessing data from storage.

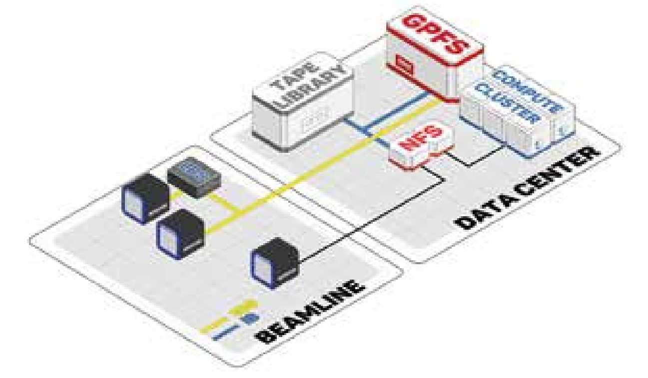

Figure 2: *Block diagram of beamline connection to central IT until 2019*

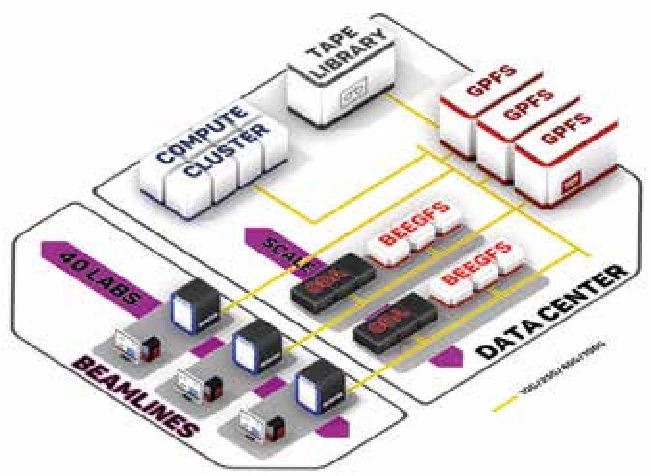

Figure 3: *Block diagram of beamline connection to central IT from 2021 onwards*

IT-infrastructure

As of 2020, and with the need for increased throughput of the compute clusters, it was decided to remove the NFS gateway and to connect new compute cluster nodes directly with multiple 25Gbps Ethernet links to the GPFS storage. This implies having a GPFS client on each of the cluster nodes. The fast access to storage by the computer clusters will unavoidably result in unpredictable bandwidth availability for the beamline detectors. Hence the need to have buffer systems capable of absorbing the data stream from the detectors for several minutes, and ideally for several hours, such that the buffer system can stream the data to central storage with variable throughput whenever the central storage is not solicited entirely by the computer clusters. Figure 3 depicts the new configuration which will be in place as of 2021. The data stream from the detectors will first be transferred to the ODA systems where a first processing in RAM can take place at the "wire speed" of the data stream. The buffer storage will then either be based on large RAM or the BeeGFS file systems with NVMe SSD storage. The dimensioning of these buffer systems has to be optimised as a function of each beamline and detector system with the aim to guarantee that the detector systems are always able to empty their limited buffer space.

Locating the ODA computers and buffer storage systems in the data centre will greatly ease system administration, hardware interventions, and flexibility in dynamically re-assigning systems in an optimal manner between beamlines.

Central storage

The ESRF has a long history of purchasing, configuring and operating RAID storage systems. Over the last two decades we have seen installed storage capacity multiplied by a factor of 1000, with an aggregate read/write performance multiplied by a factor of 700, and a price per TB divided by 1000. We expect that this trend will continue with new technologies reaching maturity for use in the data centre.

At the time of writing the ESRF operates two GPFS systems of 4.2 PB each and a 100TB SSD based storage system from the company NetApp. Before the GPFS based systems, our storage was entirely based on NetApp storage systems. The GPFS systems are used to store experimental data for peer-reviewed experiments for at least 30 days after the end of the experiments and for in-house research projects without time limitation. The storage area for peer-reviewed research is 2.4 PB, leaving 6 PB for in-house research. An order has recently been placed for two additional GPFS storage systems of 2x5 PB net capacity. These systems were initially foreseen to be operational for the restart of user operation, but this is now delayed until November 2020 due to the impact of the COVID-19 situation. The aggregate read/write performance of our GPFS storage systems has been measured at 3.5 GB/s. This has to be multiplied by the number of systems in operation, *i.e.* two right now and four in the last quarter of 2020.

GPFS is proprietary software and not compatible/supported with legacy operation systems and supports a limited number of Linux distributions. For this reason, we have to continue operating a GPFS-to-NFS gateway which allows all legacy computers to access the central storage, although at much lower performance. However, efforts have to continue to replace legacy systems whenever possible.

IT-infrastructure

Storage capacity and performance has to increase in the future such that data can be stored and processed during or shortly after the experiments. As a ballpark figure, it has proven necessary to have at least space to store the equivalent of 50% of the annual ESRF data production, *i.e.* 10 PB for 2021, for the public experimental programme, and roughly the same storage capacity for in-house research.

The choice of GPFS for central storage may be reviewed in the future if the use of BeeGFS proves to be successful for the beamline buffer storage systems. Maintaining and operating two complex clustered file systems will not be economical in the long run and should be avoided.

On-line data compression and reduction will not be sufficient to avoid that the volume of data to be stored on central disks will continue to grow significantly during the next ten years. Low cost storage alternatives have to be evaluated, possibly as tier-2 storage in the background of more expensive high-performance storage. Additional storage capacity will also be needed for restoring data from the tape archive allowing its re-use in the context of the EOSC. Such data will have to be directly visible from the Internet and it might therefore be advisable to separate the storage area from the internal central storage to ensure a high level of IT-security.

SSD storage is already playing an important role for front-end caching in our GPFS systems and is nowadays ubiquitous in PCs and laptops for its superior performance. It is expected that high-density SSD modules will become economically competitive in comparison to the classical spinning hard drive in a few years. Large SSD storage systems will then provide higher density, faster data transfer, lower latency, and reduced power consumption in the data centre.

The management of central storage systems is crucial for the operation of the ESRF and has over time become more complex to harness. In our 24/7 operation it is important that several system administrators are capable of intervening in the case of technical issues or for user assistance. The current situation with a single expert for the management of the storage systems is not sustainable.

CENTRAL STORAGE RECOMMENDATIONS

- Aim at operating only one high-performance cluster file system software at ESRF
- Investigate low-cost, high-volume storage hardware
- Ensure that there is no single point of failure in expertise for system administration by hiring at least two additional system administrator engineers for the operation of the central storage systems

RISKS OF NO CHANGE

- Failing to install buffer systems behind fast detectors means endangering the experiment
- Failing to provide enough central storage space could either result in a blocking situation when disks are full or that data would have to be removed faster, increasing the pressure on data processing

IT-infrastructure

Compute clusters

All efforts will be undertaken to reduce the data while it is produced. However, data sets will increase and will demand substantial computational means before being sufficiently analysed and reduced to be exported for the final data analysis in the institutions of the visiting scientists.

The on-site compute clusters have been regularly modernised and upgraded and with the restart of the EBS Storage Ring and regular user operation at the end of August 2020, two new clusters, one based on GPUs, the other on CPUs, will be operational. At this point the beamlines have access to central computing capacity as summarised in table 2. The column "August 2020" refers to orders of IBM Power9 GPU and DELL Intel servers which are already delivered but not yet put into operation.

	May 2020	August 2020	Needed 2025
Batch cluster CPU-cores	1780	3380	9000
Macromolecular Crystallography (MX) dedicated CPU-cores	728	728	1000
GPUs (V100 or equivalent)	8 (*)	38	300

Table 2: *Batch cluster capacity 2020; (*) 8 V100s in DGX-1 computer, current installation of 28 x Kepler-family GPUs are obsolete.*

According to the feedback obtained through the survey of beamline responsibles, needs may be much higher than the estimates for 2025 in table 2, and it shall be noted that the figures obtained from the survey do not include simulation or ML/AI applications. To cite two examples among the numerous requests expressed in the survey:

- The very successful operation of the CryoEM (Titan Krios) will be further enhanced with an increased use of an automatic data processing pipeline and potentially also a new and faster detector. It is also intended to let users do the re-processing of data at the ESRF. In comparison to today, the computational needs will increase by a factor of three to five over the next five years. The CryoEM has recently started using the DGX-1[6] mini supercomputer for data processing and has seen a significant increase in throughput and capacity to do more downstream processing. The DGX-1 could be interesting for other applications which are adapted to this kind of hardware architecture (compared with Power based systems) *e.g.* to run ML algorithms

- On ID16B (hard X-ray nanoprobe beamline), up to 30-50% of all experiments may in the future be based on holo-tomography, which is very demanding in terms of analysis. Back in 2018, a single scan needed about 360 CPU cores per day for the processing and it was possible to run 10 to 20 scans per day, *i.e.* 3 600 to 7 200 CPU cores per day. At that time only every 10th scan was analysed on-line; a situation far from ideal. Especially for industrial users, this is not acceptable. *In-situ* experiments can now be performed at lower resolution every 7s for up to two hours, *i.e.* 2000 scans, and this can be repeated three to five times a day, possibly up to eight times with improved alignment procedures. Assuming 3 000 scans per day, this would require a peak demand of 250 000 CPU cores. The current analysis code is based on the open-source Octave package and is currently being ported to GPUs. This will improve the situation significantly but it will nevertheless need a large number of GPUs. Ideally the code should not only be ported to GPUs but also improved and optimised. However, this is very time-consuming and will also need extensive validation by the beamline team.

[6] *https://www.nvidia.com/en-us/data-center/dgx-1/*

IT-infrastructure

The on-site processing capacity is generally fully used by scheduled research projects and cannot be sized for peak demands, independently of the pertinence of such demands. Peak demands could be satisfied by scale-out solutions such as cloud computing or access to external HPC resources. The ESRF has tested using cloud services in particular with Amazon AWS.

Tendering for cloud resources is not easy, as the experience with the H2020 projects HNSciCloud[7] and OCRE[8] has shown. It is even more difficult to regulate and monitor the usage of such resources. As of today, it is difficult to imagine free access to cloud resources which are billed at the point of use to the ESRF. A light-weight attribution mechanism would have to be implemented allowing to control and monitor the use of such resources. The cost of cloud resources has been evaluated[9] as being generally more costly than on-site IT infrastructure, even if considering a full-cost model. However, this is not the case if only used for peak demands. For peak demands the cloud model is more competitive and flexible. If the peak demand is not urgent it may in addition be possible to profit from attractive spot market prices.

Very large jobs for simulation and data processing - not necessarily involving the transfer of large data sets - would benefit tremendously from an access to on-demand HPC resources. This would be feasible if a fraction of resources in the European HPC centres is set aside for such usage, similar to what is being established for the DOE BES labs. All past efforts to obtain such access have so far not been fruitful. However, in the frame of EuroHPC, PRACE, and the EOSC, discussions should continue to investigate this further.

COMPUTE CLUSTER STRATEGY RECOMMENDATIONS

- Increase compute capacity within the limit of the capacity of the data centres and/or budget. A temporary solution for housing the equipment may have to be envisaged for 2023+2024 as a function of the advancement of the construction of the new data centre
- Work on the business model for Cloud Computing, make a recurrent budget provision for cloud resources as of 2021
- Continue discussion on national and international levels - also on behalf the entire PaN community - for on-demand access to EOSC and HPC resources
- Hire an additional system administration engineer for supporting the above

RISKS OF NO CHANGE

- Failing to increase compute capacity would mean that batch queues saturate, data would be impossible to process, users would have to transport massive raw data sets home and struggle with data processing in their home lab
- Data analysis would take more time and result in delays of publications or no publications
- Failing to provide access to cloud computing would mean that even peak loads have to be covered by the ESRF infrastructure, or do not get done
- Failing to ensure adequate user assistance for a smooth operation of the cluster would endanger the operation of the ESRF

[7] *www.hnscicloud.eu*
[8] *www.ocre-project.eu*
[9] *https://www.itproportal.com/features/cost-comparison-cloud-vs-on-premise-201819/*

IT-infrastructure

Data Communication Network

The data communication network consists logically of the network links within the data centres, the local network on the beamlines and the interconnecting backbone links between beamlines and the data centres as depicted in figure 4.

A large dedicated network is used for the machine control system which has been completely renewed and recently commissioned before the machine commissioning started. The machine control network is separated from the general network by a firewall.

The backbone network is based on single-mode fibre optic cables from the two data centres to the beamlines and transits through wiring centres. This part of the infrastructure has been renewed and upgraded with additional fibres and new patch panels. The patch panels are now based on Commscope products allowing remote management with the imVision[10] software. The patch panels in the wiring centre can now be used to monitor and report automatically the usage of the fibre links.

Each beamline can now be connected to the data centres with a dual 10Gbps link. For higher bandwidth, dedicated links can be easily deployed on demand using the same fibre infrastructure. It implies changing the transceivers and possibly the beamline network switch and is taking place for the deployment of the EIGER-2 detectors.

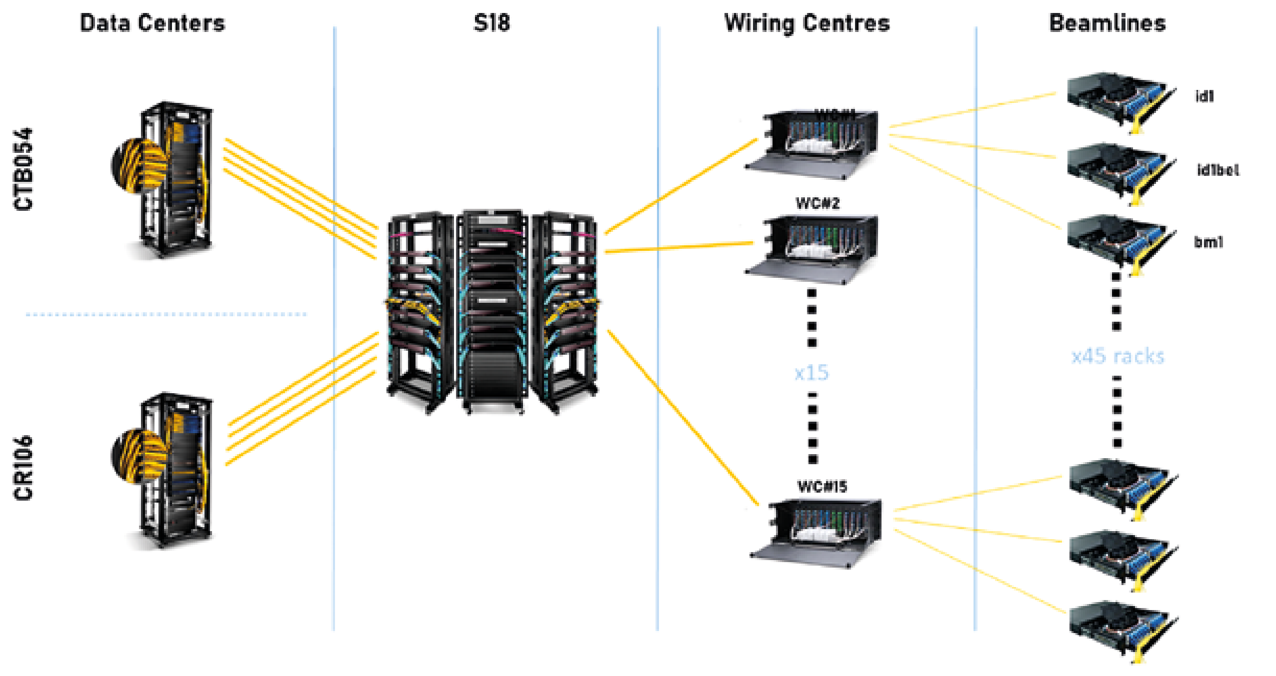

Figure 4: *Fibre optic cabling structure*

[10] *https://www.commscope.com/globalassets/digizuite/2718-imvision-automated-infrastructure-management-br-108978-en.pdf*

IT-infrastructure

In the future it will be possible to implement 400Gbps links, again by exchanging transceivers and upgrading the beamline network switch.

All beamline networks are based on Cat6 copper wiring and in general the beamlines switches are foreseen for 1 Gbps network devices. They have sufficient capacity to connect additional devices and if needed can be upgraded. High-speed detectors are directly connected to the wiring centres unless the number of high-speed devices on a beamline exceeds five.

All beamline switches will be upgraded in two batches in 2022 and 2023. At that time the old switches are likely to be obsolete. This upgrade is planned and budgeted.

The network ports in the data centres are either 40Gbps or even 100Gbps. The overall number of network ports at the ESRF and hence the number of devices connected to the network has increased significantly over the years. The trend of connecting more and more devices to the network (wired or Wi-Fi) will continue. In 2020 we had more than 10 000 active network ports.

NETWORK RECOMMENDATIONS

- Move to faster backbone network connections in function of the needs and the maturity of the technology
- Purchase network components from either CISCO or Extreme Networks according to technical requirements and best value for money
- Hire two additional network engineers

RISKS OF NO CHANGE

- The lack of human resources has already had a detrimental impact on IT projects and has to be avoided in the future

Internet connectivity

The ESRF Internet connectivity is crucial for daily operation. An increasing number of business applications are hosted off-site and most of the communication with the scientific user community takes place via Internet. The Internet link is also used to export scientific data, however, so far at a relatively low rate (<10% of the total data production in 2018). This can be explained by a weak end-to-end bandwidth and a lack of user-friendly tools to transfer big data sets reliably. In 2018 - the last year of full operation with the previous Storage Ring - an average of <10% of the total available bandwidth was used to export data. In addition to data export, two new and important fields will require a high-speed Internet connection in the future: the coupling of experiments at the ESRF to HPC resources in Europe for simulation and data analysis, and the re-use of data in the frame of FAIR[11] data management and the EOSC.

The ESRF Internet link is shared with the three other laboratories present on the campus: ILL, EMBL, and IBS. This has allowed sharing the cost and effort to install and administer a common firewall and DMZ networks. The Internet link has been at 10 Gbps since 2012 and was, overall, used at 10% of its capacity by the laboratories on the campus in 2018. The link feeds into the Grenoble research and education metropolitan area network and on the University Campus into the RENATER French national backbone network for education and research. The schematics of the four network layers, *i.e.* campus, Metropolitan, France, and European are given in figures 5 to 8.

[11] https://ec.europa.eu/info/sites/info/files/turning_fair_into_reality_1.pdf

IT-infrastructure

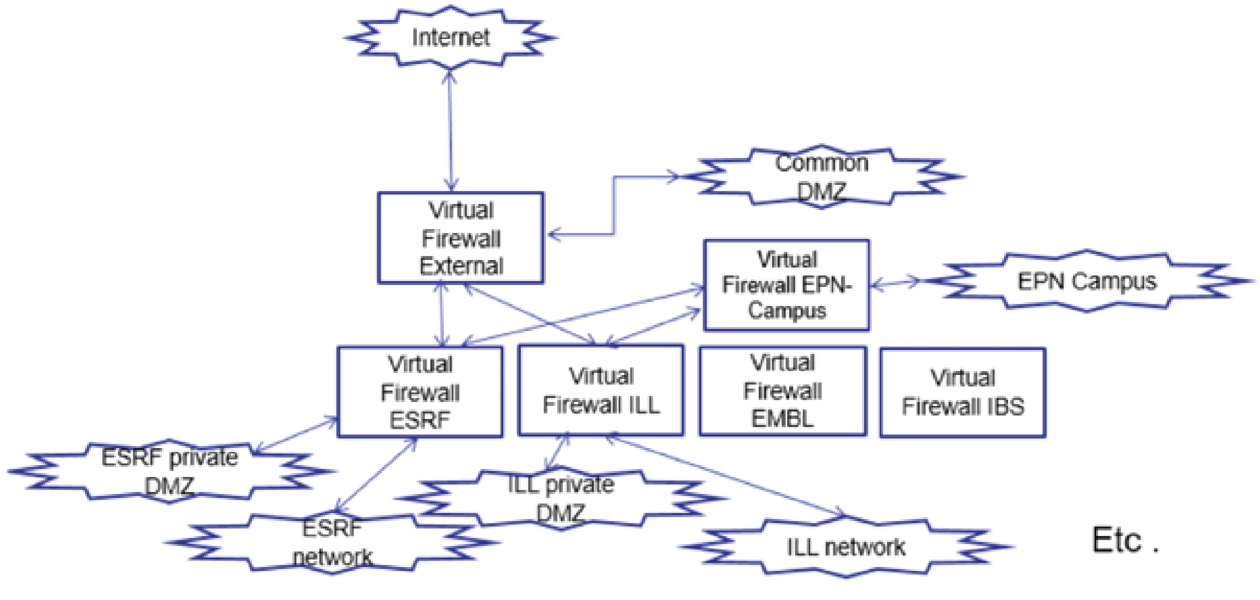

Figure 5: *EPN campus Internet access configuration*

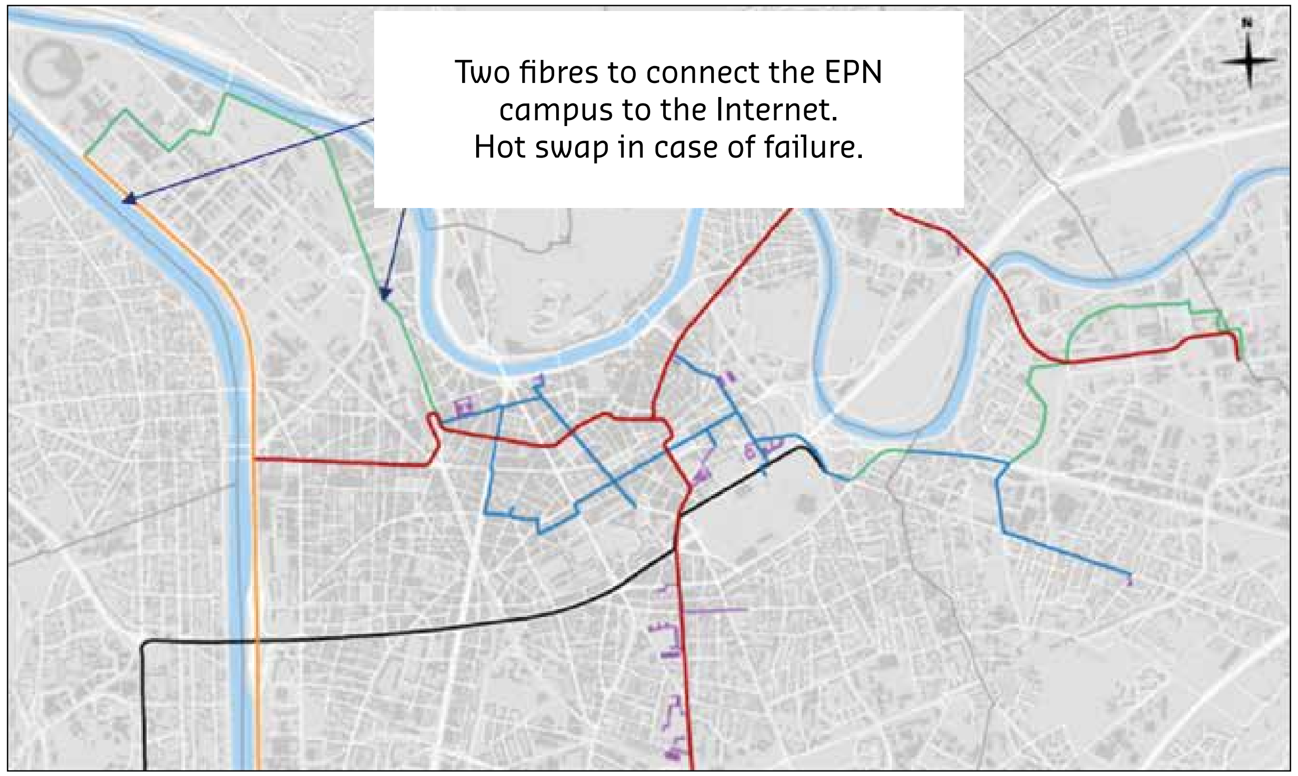

Figure 6: *Grenoble metropolitan network for education and research*

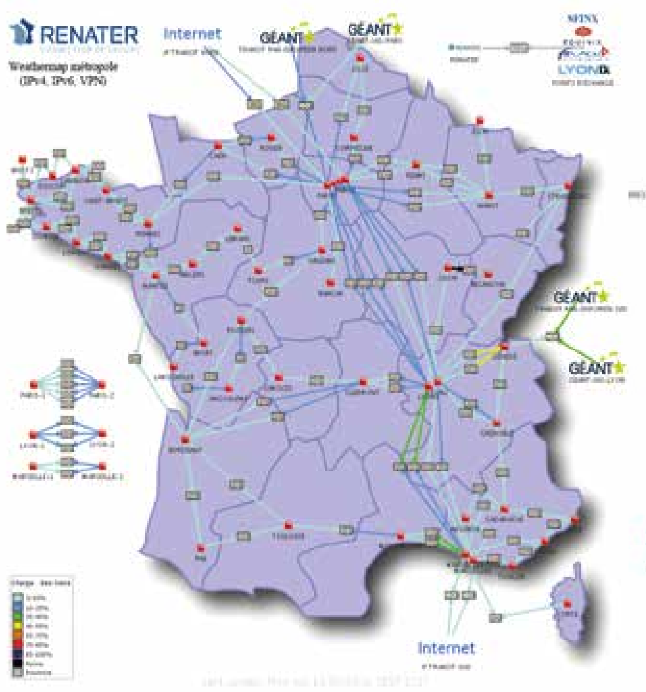

Figure 7: *RENATER "weather map" on June 2019, showing a snapshot of the bandwidth utilisation*

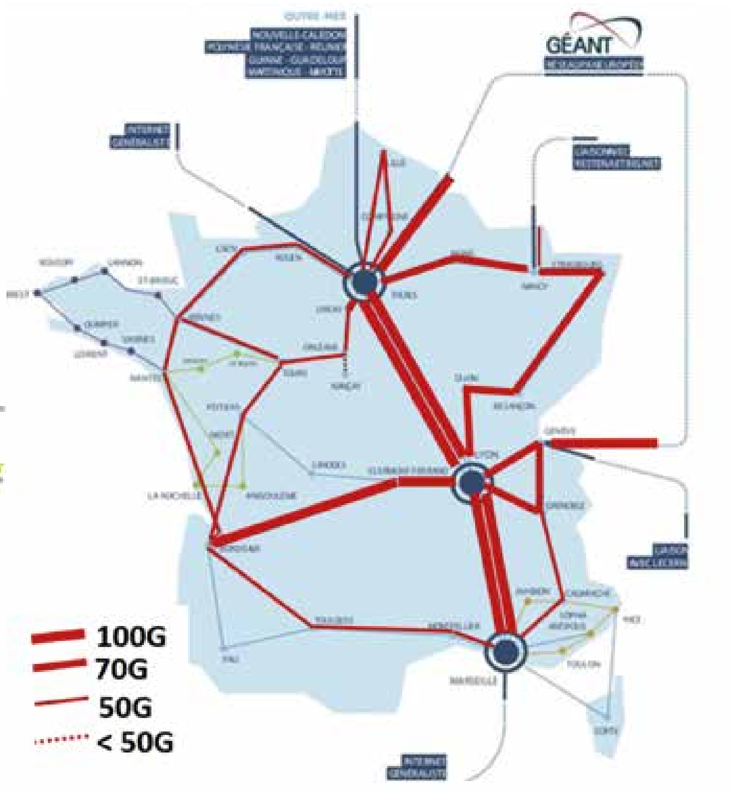

Figure 9: *RENATER's projected IP routing capacity for the end of 2020*

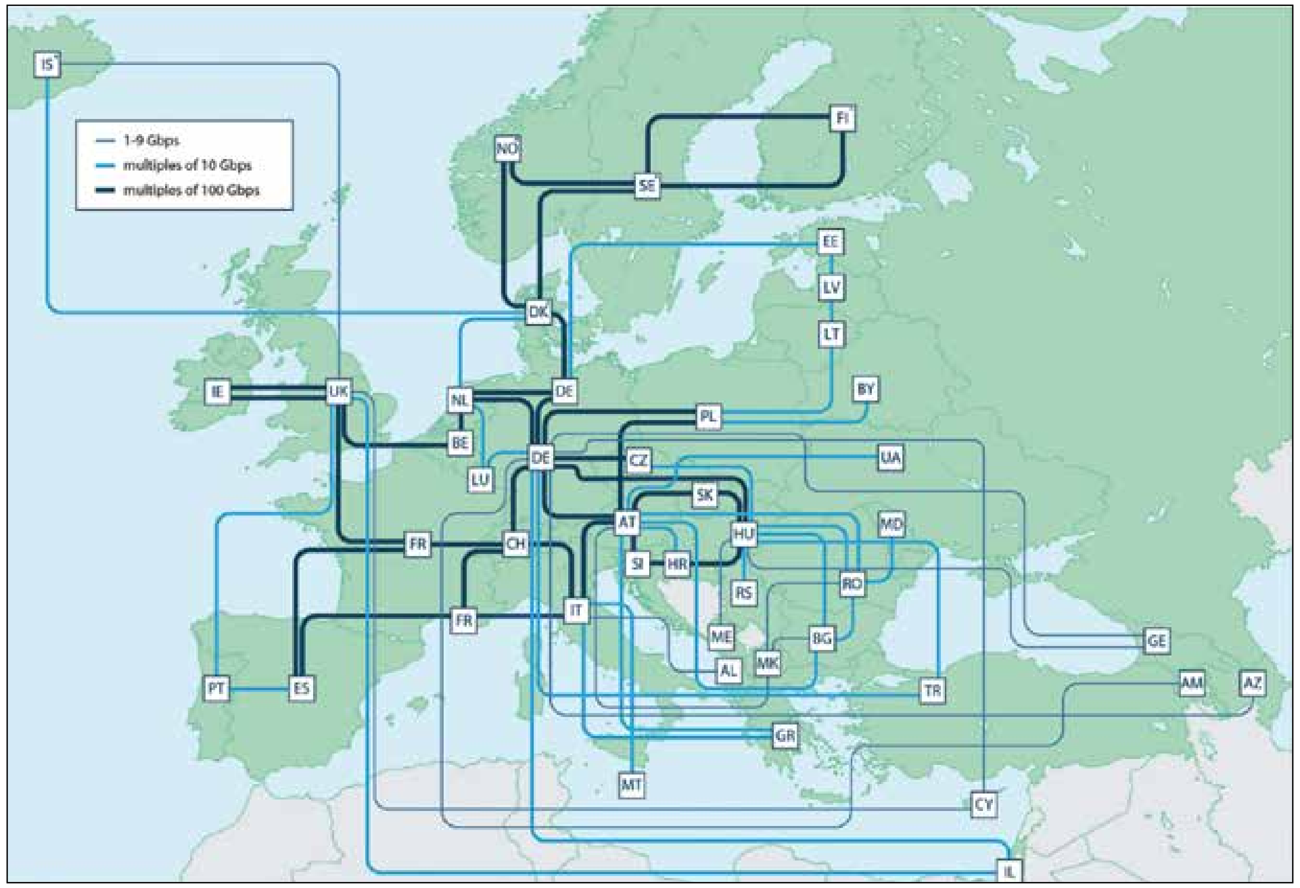

Figure 8: *GÉANT's network interconnecting national research and education network organisations (NRENs) with a high-speed European backbone networkt*

IT-infrastructure

Discussions with the metropolitan network and RENATER took place over the last six months to investigate what could be done to enhance the EPN campus Internet connection. This has resulted in an overall upgrade of the metropolitan network equipment, including the EPN campus router, to 40 Gbps with provision to go to 100 Gbps. However, the EPN campus firewall has not yet been upgraded and it is foreseen - and funds have been allocated - to do this by the end of the year. RENATER presented their plans to upgrade the French backbone network in 2020. This would lead to an IP routing capacity of 70 Gbps on multiple 100 Gbps links between Grenoble and Lyon, 2x100 Gbps to Paris, and 70 Gbps between Grenoble and Geneva. Paris and Geneva provide the transition points between RENATER and the European backbone network GEANT (figure 9). It shall be noted that this capacity is shared with all French higher education and research organisations and that in the past the usage of the backbone capacity between Grenoble-Lyon-Paris has never exceeded 50% of the available bandwidth.

The RENATER upgrade to physical 100 Gbps links will unavoidably lead to a modification of their pricing and therefore to an increase of the ESRF recurrent budget. It shall be noted that the ESRF (and also the ILL and EMBL) is considered as a private French non-profit company and has to pay the RENATER fees. The fees are essentially calculated according to the effective usage of the Internet link (calculated on the 95% percentile of the bandwidth consumed over the last 12 months) and are currently supported at 55% by the ESRF and 45% by the other laboratories on the campus. A likely scenario would be that this fee increases to 150 k€/year in five years from now for an effective usage of 10Gbps, *i.e.* 80 k€/year for the ESRF.

An effective use of HPC resources may require dedicated bandwidth to some selected HPC centres in Europe. This could be achieved with the help of RENATER or private Internet access providers and will need to be investigated and decided when access modalities to the HPC resources are established. A close follow-up of the market evolution has to be pursued to enable ESRF to react quickly knowing that implementing specific Internet connections takes a lot of time.

INTERNET CONNECTIVITY RECOMMENDATIONS

- Upgrade of the site firewall in 2021 for 100 Gbps network speed and beyond
- Foresee a provision of ~80k€/year for the enhanced RENATER bandwidth subscription from 2022 onwards

RISKS OF NO CHANGE

- A lack of bandwidth would severely limit the capacity of users to export data
- Data re-use will also critically depend on available network bandwidth and would be severely limited if we do not pursue our quest for the highest possible bandwidth

IT-infrastructure

Backup and Archiving

With the growing amount of data, the cost of managing and storing data in primary storage, *i.e.* high performance disk storage, adds up fast. Shifting data as quickly as possible to secondary storage, such as tape, can save resources, as long as this is done at the right moment when the likelihood of bringing it back to primary storage is low. Backup and archiving are two distinct methods of keeping one or several copies of data off-line. Backup is intended to keep one or several copies of the data for restoration in case of unintentional erasure or loss of data due to malfunctioning equipment. Backup copies are usually only kept for a limited time span before the storage medium is re-used for new data. Archiving consists in preserving one or several copies of the data for future re-use. Backup and archiving are complementary, but when used for the same data unavoidably means storing the same data several times.

At the ESRF we operate two StorageTek SL8500 robotic tape libraries. StorageTek, a company founded in 1969 was purchased in 2005 by Sun Microsystems and five years later, in early 2010, Sun Microsystems was acquired by Oracle Corporation.

Each SL8500 tape library is situated in a different computer room distant by ~300m and can hold up to 8500 tapes. Backups and archives are made simultaneously on each of the libraries, *i.e.* every dataset is written at least twice to tape. This is done to ensure that we always have a copy of the data even in the worst case scenario of a disaster such as fire or flooding in one of the computer rooms.

Tape storage is by far the cheapest method for long-term data storage, but it has nevertheless a significant cost. The SL8500 tape library itself is a long-term investment, but the tape drives need to be replaced over a time span of five to seven years to newer models allowing the use the latest tape standards. To keep costs down, the ESRF has decided to avoid backup and archiving of the same data. At the restart of the facility in 2020, it is planned that all data produced by visiting scientists at the beamlines will be written exclusively to the tape archive without copying another version to the backup tapes. At a later stage the same strategy will be adopted for in-house research data.

The ESRF uses exclusively LTO tapes and tape drives, and currently has a mixture of LTO-5, LTO-7, and LTO-8 tapes and tape drives. The introduction of LTO-8 was significantly delayed by the patent dispute between Sony and Fujifilm which was only settled in late 2019[12]. The dispute generated a world-wide backlog on tape delivery and it took six months for all orders to be honoured. LTO tape drives feature "built-in" lossless compression, based on the so-called Streaming Lossless Data Compression (SLDC) algorithm. The compression efficiency for our beamline data is observed to be on average 1:1.7, resulting in up to 20 TB capacity for a single LTO-8 tape. Based on LTO-8 tapes and assuming that all lower capacity tapes are replaced, the total storage capacity of the SL8500 tape library is 170PB.

As of 2021, the next generation of LTO tapes (LTO-9) will become available on the market and feature 24TB raw capacity and in our case 40TB of compressed capacity. We usually wait for at least one year before introducing the latest generation tape drives and tapes, which would mean that switching over to LTO-9 will take place in 2022. The evolution of the backup and archive capacity, together with the cost for tape drives and tapes is summarised in Annex 2 for the period 2021 to 2031. Based on the data production estimates in table 1, we will reach the limit of the capacity of our SL8500 tape libraries in 2028. At that point it will be necessary to consider extending or replacing the tape libraries. Provision for the space to locate a second SL8500 in the data centre CTB-54 exists, whereas it is not possible to house a second library in CR106 without relocating a large amount of other equipment. This is foreseen once the new data centre is operational in either 2024 or 2025 according to the scenario retained (see following chapter).

Replacing the two existing SL8500 may have to take place earlier if ORACLE decides to abandon the tape market, leaving us in a situation that maintenance but also new tape drives cannot be purchased anymore for our tape libraries.

For the last 15 years we have used the TiNa (Time Navigator) software from the company Atempo. This software has done an excellent job in the past but has now reached a limit in functionality and performance. The company Atempo has developed specialised software for archiving data, but the pricing is prohibitive. This is why we envisage to use the open source software Bacula[13] instead, initially for archiving and as soon as possible also for backups. Bacula is currently under evaluation.

Backup and archiving rely on one engineer and one technician. The engineer who is in charge of designing the scripts and configuring the server infrastructure for the tape libraries is mission critical. With the sheer volume of data and the size of the infrastructure but also to avoid that this critical infrastructure is entrusted to a single engineer, we propose to foresee a second engineer position.

BACKUP AND ARCHIVING RECOMMENDATIONS

- The software Bacula shall be validated and then used for data archiving and backup as of 2021
- Storage capacity of our existing SL8500 libraries is sufficient for the next five years if the data production evolves like given in table 1 and provided maintenance can be continued by the manufacturer
- Investigate alternative solutions to tape storage and monitor the evolution of the tape market to be able to take the right decisions sufficiently in advance
- Two additional tape libraries shall be budgeted for 2025/26
- Hire one additional engineer as soon as possible for the operation of the backup and archiving facility to alleviate a single point of failure

RISKS OF NO CHANGE

- Tape storage is the most economical way to store large data volumes. Not being able to increase tape storage capacity would mean much more expensive primary disk storage
- The ESRF data policy will not be respected and would have to be revised to reduce our commitment to archive data for 5-10 years.

[12] *https://blocksandfiles.com/2019/08/06/lto-8-tape-media-manufacturing-block-mysteriously-blown-away/*

[13] *www.baculasystems.com/*

IT-infrastructure

New on-site data centre

The very first deliverable of the ESRF Upgrade programme was the extension and refurbishment of the data centre in the Central Building. Delivered in 2011, the new premises host most of the new central computing equipment since then. The evolution of the needs - in particular with the largely enhanced performance of the EBS Storage Ring - has led to the analysis that the current capacity in terms of rack space, power supply and cooling of the two computer rooms will become insufficient within the next few years.

Three different possibilities to extend our IT infrastructure capacity have been investigated:

1. Off-site hosting of storage or compute resources. There are commercial data centres able to provide additional storage or compute resources or host equipment. However, this business model lacks flexibility and is not economical at the scale of the resources needed by the ESRF. Usually this type of operation is ideal for "standard" IT provisioning for which the hosting company deploys standard equipment which is then shared by several clients. The performance, capacity, and reliability requirements of the ESRF do not fit into this business model. Typical suppliers for off-site hosting in the vicinity of the ESRF are EOALAS or HOSTELYON.

2. Cloud services for storage and compute. By definition cloud services imply that the physical location of the equipment is unknown. Cloud services are well suited to respond to peak demands thus allowing to dimension the on-site infrastructure in accordance with the base load. However, the business model for procuring and using cloud resources is difficult to harness for our scientific activity. This is further developed in the paragraph "Cloud Computing" below.

3. Creation of a larger on-site computer room with sufficient capacity to cater for the (basic) needs of the next decade or two. Constructing a new on-site data centre is expensive but provides the highest short and long-term flexibility and reliability in the context of our 24/24h 7d/7d high availability operation.

Scenarios for a new on-site data centre

Building a new on-site data centre has a significant lead time. Typically, four to five years are needed from the time of the decision to the actual delivery of a new or fully refurbished building. This is why we started in the second half of 2019 to look into possible scenarios allowing the ESRF to increase its on-site IT hosting capacity.

Current Situation

The ESRF currently operates two data centres, one located in the Central Building (room CTB054) with 300 m² of floor space for computing equipment and a second room located some 300 m away in the Control Room (room CR106) building with 145 m² floor space. Operating two distinct rooms offers the possibility to physically separate the disk storage from the tape storage thus ensuring that a copy of the data is preserved in case of a fire in one of the rooms. However, we have not been able to deploy all computing equipment symmetrically between the two rooms because of their different floor sizes and power/cooling capacities. Figure 10 depicts the current rack configuration in the two rooms.

IT-infrastructure

Table 3 below summarises the main characteristics of the two existing rooms. The room CR106 needs major refurbishment. The power distribution of the room is not adequate anymore and the UPS system, currently located on the roof of the booster in sector 30, would need to be located closer to CR106 and upgraded.

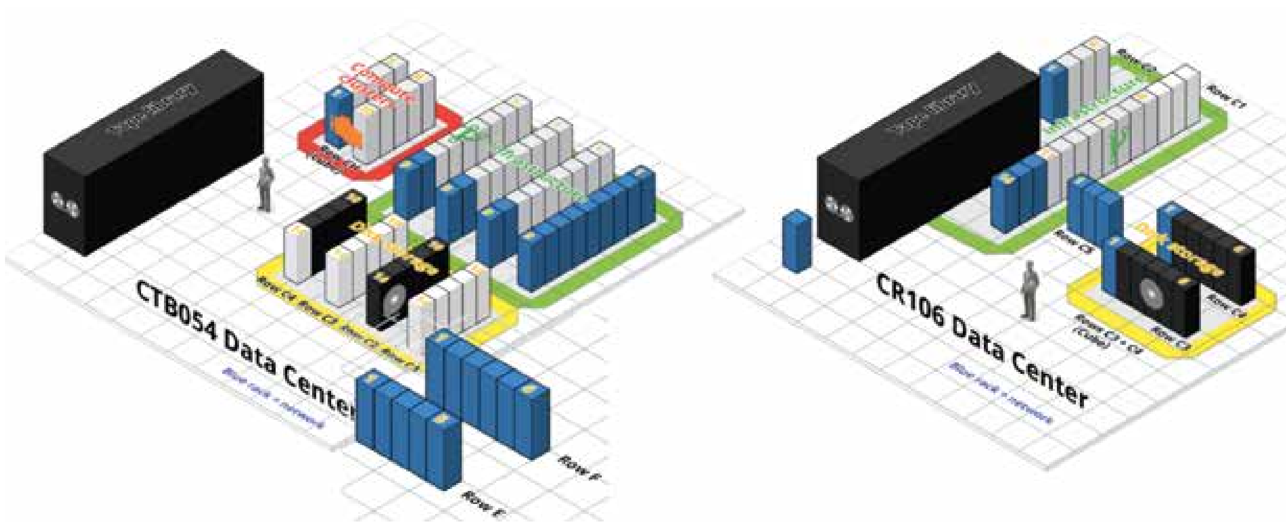

Figure 10 : *Layout of the two ESRF data centres CTB054 and CR106)*

	Room CTB054	Room CR106
Net surface for computer equipment	297 m²	145 m²
Maximum power + cooling	370 kW	150 kW
Current power usage	125 kW	65 kW
19" racks	67	24
Water cooled racks for high-power equipment	6	3

Table 3 : *Main characteristics of the CTB054 and CR106 rooms*

The equipment in the data centres is subject to constant evolution driven by the needs to modernise existing equipment and to satisfy the additional needs from the beamlines for more storage and compute capacity. One of the key challenges is to anticipate physical limitations which would block deployment of additional equipment. Physical limits are the available surface and rack space, cooling capacity, or power supply. Based on the projected data production of our beamlines (table 1) and the needs derived for on-line and off-line storage and processing capacity, we will reach the power and cooling limit of CTB054 and CR106 and rack space at the end of 2023, and we would require to double the tape robotics by 2028.

IT-infrastructure

The only way to anticipate these limitations consists in planning for a new on-site data centre which must be available as soon as possible and not later than 2023/24. If this timing proves impossible we would have to resort to costly temporary solutions, such as installing one or several containers on a parking lot.

A new on-site data centre would enable the relocation of all computing equipment of CR106 to the new room. Optionally it is possible to leave the StorageTek tape library in CR106. This equipment is difficult and expensive to move and does not need much power and air conditioning. If all other computing equipment is removed from CR106, there would be sufficient space for a second tape library in this room.

Need for a new on-site data centre

Five different scenarios have been studied starting from the assumption that within ten years the increasing needs might require a surface in the order of 1000 m² and 1MW of power. From the five scenarios, three have been rejected at an early stage because of strong technical limitations. They concern:

1. The modification of the Science Building, which is very limited because of the numerous structural walls of the building making it impossible to create larger contiguous surfaces. Implementing an efficient cooling system for the heat generated by the computer systems in this building would also be extremely difficult.

	Scenario 1 New Building in front of the Central Building	Scenario 2 Conversion of SBGS building
Surface	1500 m² with 680 m² net surface for IT racks	**Phase 1:** two floors of 390 m² with 282 m² net surface for IT racks **Phase 2:** relocating of Stores adding another two floors of 390 m²
Power/cooling	1.1 MW	**Phase 1:** 570 kW **Phase 2:** 500 kW
Specific technical aspects	▪ would allow easy implementation of solar panels on roof ▪ would allow to envisage cooling with ground water	▪ No building permit ▪ Requires relocation of mechanical workshop, ▪ Requires an independent support structure for the 2nd floor ▪ Proximity to the PGUT building makes power and cooling connection easier
Cost	10.5 to 11 M€	**Phase 1:** 5.5 - 6 M€ **Phase 2:** 3.5 - 4 M€
Earliest possible delivery	2025	2024 for phase 1

Table 4 : *Comparison of the two possible scenarios for a new on-site data centre*

IT-infrastructure

2. The use of the accelerator assembly building ESRF01. This building has a total surface of 1100 m² and could easily be retrofitted and converted into a data centre. However, it would require moving the newly installed Insertion Device Laboratory to another place.

3. The liberated space of the former Insertion Device Laboratory in the Central Building. With a total surface of 365 m² it could provide at most a usable space of 200 m². The height is insufficient to construct a second floor and it would be necessary to reinforce the slab. Increasing the power consumption of the Central Building further might be risky and being in the direct vicinity of the Central Building data centre would increase the risk to lose the entire IT infrastructure in case of fire or flooding.

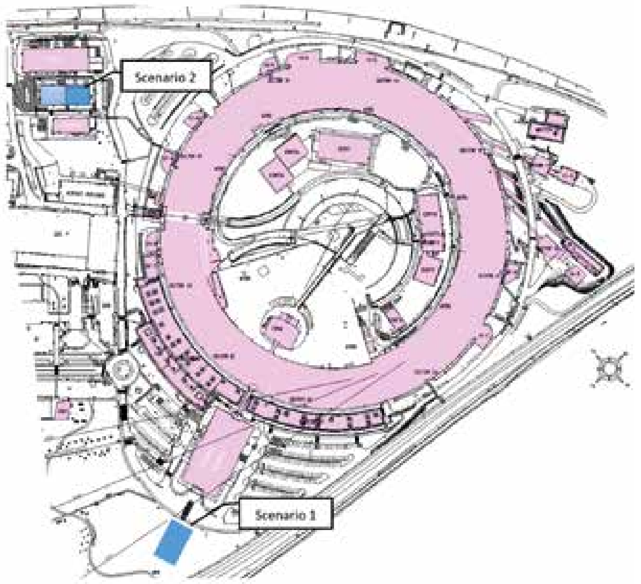

Figure 11 : *Locations of the potential new data centre building in front of the Central Building (Scenario 1) and the conversion of half of the SBGS building (Scenario 2)*

The two other scenarios are the construction of a new building in front of the Central Building or the conversion of the central workshop area in the SBGS building. Figure 11 shows the physical location of these two scenarios.

The construction of a new building has the obvious advantage that it offers the possibility to create an optimal technical environment without impacting other activities on the site. Converting half of the SBGS building has the advantage of making the premises available on a shorter time scale and leaves the option to convert the second half of the building - currently occupied by the common Stores - at a later stage if the needs materialise as forecasted. Table 4 shows the main characteristics of the two scenarios. It shall be noted that the cost estimates for Scenario 2 do not include costs related to moving the activities currently in the SBGS building to other locations.

DATA CENTRE RECOMMENDATIONS

- Construction of a new on-site data centre with enough floor space, power and cooling to respond to the needs for the next decade
- Move all IT equipment, except for the SL8500 tape library, from the CR106 room to the new data centre and dedicate CR106 to the tape library with sufficient space for a future second tape library

RISKS OF NO CHANGE

- If a new data centre cannot be financed, an intermediate solution could become permanent and result in less reliable operation conditions of the IT equipment. Alternatively the amount of IT equipment could be downscaled to the detriment of data processing

Software

Software

Introduction

Software is the brain behind the body of IT infrastructure. Without the software the IT infrastructure would be practically unusable. For this reason, it is vital to include software as part of the IT data strategy. The main group dedicated to developing software at the ESRF is the Software Group within the ISDD. It is made up mostly of software engineers working on accelerator and beamline control and data acquisition. A smaller unit of scientists and software engineers inside the group work on data analysis and processing, data visualisation, and data management. In addition to the Software Group there are a large number of scientists in the Experiments Division and Accelerator Division who develop software on a part-time basis. This section will cover the software strategy for Accelerator Control, Beamline Control and Data Acquisition and Scientific Software taking into account the needs of the EBS. The developments carried out in the Experiments Division and Accelerator Division will need to be aligned with the software and IT data strategy to ensure an optimal use of the infrastructure.

Accelerator Control

The control system for the EBS machine was completed in time for the commissioning of the EBS Storage Ring which started in December 2019. The control system fulfilled all the requirements to start the commissioning of the EBS Storage Ring. Some developments of the control system are still necessary to reach full design specification by the end of 2020. This section looks into the future evolution of the EBS control system beyond 2020.

The accelerator control system is based on Tango Controls (https://tango-controls.org) which was developed at the ESRF as an open source project, starting in 1999 and continued until today. The EBS Storage Ring runs on the latest version V9 which is under Long Term Support (LTS). Long term means it will be maintained for at least 5 years after the next major release of Tango (V10). Tango has been adopted by a number of sites and is maintained and developed by a consortium of 11 institutes[14]. The consortium members contribute financially to the maintenance of Tango Controls which is managed by the ESRF. The major need of the EBS regarding Tango Controls is that it is maintained and enhanced for the next 20 years. This also holds true for the community as some of the projects using Tango Controls will come on-line in 2025 and need to run for another 30 years (*e.g.* the SKA telescope - the world's largest telescope, the new SOLEIL and ELETTRA sources, the ELI laser facility and other sites in Europe and Russia). Preparatory work has started on the next major version of Tango Controls which will have a pluggable protocol layer to allow it to adapt to new network protocols.

The major areas of development for the EBS control system are:

- **graphical user interfaces** - develop new graphical applications using web-based technologies like WebJive[15] to gradually replace Java Swing applications so as to be independent of the deprecated Swing toolkit.
- **archiving** - archived data are the memory of the accelerator behaviour and equipment. The need is to archive thousands of control data (attributes) at 10 Hz or more. A new archive back-end was developed for the EBS which uses decimation and time-to-live (TTL) to limit the amount of data stored on disk to a few 10s TB/year. Archive data from the accelerator are kept forever therefore the increase in production of archived data will require an increase in the storage capacity for the accelerator. These data are essential as training data for developing ML algorithms in the future.
- **timing** - the EBS Storage Ring is the first synchrotron to use the White Rabbit (WR) timing system for synchronisation. WR can provide event correlation via absolute time stamps when connected to a GPS clock standard. This will be implemented to allow correlating events across the accelerator with nanosecond precision.
- **Tango V10** - ESRF will continue to play a leading role in developing the new version of Tango and test it with the Tango community to ensure that Tango continues to fulfil the needs of the EBS and is maintained for another 20 years.
- **virtualisation** - continue replacing physical hosts with virtualised IT infrastructure and use containers and orchestrators to manage software process.
- **insertion device (ID)** - many EBS beamlines and especially those employing spectroscopy need real-time synchronisation between the monochromator and the IDs. The synchronisation of experiments with IDs for energy scans movements will be improved to provide faster and more precise energy scanning.
- **digital twin simulator** - the EBS simulator has played an essential role in preparing the EBS commissioning. A live instance of the EBS simulator will be maintained and run permanently in parallel with the real operation. It will be used to detect discrepancies or drift between the model and the physical machine at an early stage as well as test new lattices. The live model is referred to as the "digital twin". The digital twin will be enriched by a set of tools inspired by machine learning algorithms to anticipate any unexpected drift and anticipate a solution before the problem shows up to the users.
- **live data streaming** - generalise the streaming of live big data (hundreds and thousands of attributes) to operator consoles and to machine learning algorithms. The use of streaming architectures *a la* Kafka[16] will be explored.
- **machine learning** - exploit archived and streamed data from both the live and simulated accelerator using machine learning algorithms to detect and correct machine anomalies. Recent publications[17] on machine learning applied to accelerators show the promise of ML applied to accelerators to maintain the beam stable during insertion device movements.
- **Python** - generalise the use of Python and Jupyter notebooks as an alternative to Matlab, complete the ongoing replacement of Matlab for calculating the beam behaviour in the simulator, and provide a Python interface for searching and extracting data from the archive.

Human resources for developing and maintaining the accelerator control system are spread across the ASD and ISDD divisions. The current staff for both hardware and software are considered sufficient for the future. The need to manage software obsolescence by introducing new technologies means staff need to continuously learn new technologies.

[14] *https://www.tango-controls.org/about-us/#steering%20committee*

[15] *https://developer.skatelescope.org/projects/ska-engineering-ui-compose-utils/en/latest/overview.html*

[16] *https://kafka.apache.org/*

[17] *doi:10.1103/PhysRevLett.123.194801* "Demonstration of Machine Learning-Based Model-Independent Stabilization of Source Properties in Synchrotron Light Sources" *by Leeman et. al.*

Software

Software

ACCELERATOR CONTROL STRATEGY RECOMMENDATIONS

- Continue improving the interface between the beamlines and the accelerator (specifically the control of insertion devices)
- Enhance the performance of the accelerator by introducing machine learning algorithms. Increase the capacity of data archiving to ensure adequate machine learning training data
- Maintain a modern control system, by continuing to lead the Tango Controls Collaboration

RISKS OF NO CHANGE

- Experiments will not be optimized for certain beamlines which need to move IDs often, *e.g.* spectroscopy beamlines
- ESRF will lag behind and not profit from ML/AI for improving the accelerator performance and operation
- If Tango is not maintained it will eventually need to be replaced at a very high cost in terms of human resources

Beamline Control

The beamlines are where the experiments are conducted and scientific data produced at the ESRF. Their optical, mechanical, hardware, detector and software features determine what experiments are possible and what data is provided to users who come to the ESRF. The ESRF beamline portfolio is constantly being improved through the construction of new beamlines and upgrades of existing beamlines. The same is true for the beamline control and data acquisition system.

The ESRF beamline control system plays a key role in the sequencing of experiments. The ESRF has used SPEC[18] to sequence experiments at the ESRF since 1992. SPEC is not adapted to performing very complex data acquisition sequences, especially rapid continuous scans. The decision to develop a new control system (BLISS) was described in the Orange Book[19] and started in 2015. The first version was used to control experiments in 2017 and presented at the ICALEPCS conference in Barcelona[20]. The core functionality of BLISS has been developed over the last 4 years and after being initially deployed on MX beamlines, BLISS has now reached a level of maturity which makes it possible to deploy it on beamlines for more complex experiments *e.g.* tomography and full field spectroscopy. The team developing the core of BLISS was for a certain period of time made up of 5 full-time developers. Although BLISS is mature now, the core will constantly require human resources to continue to evolve according to the needs of new experiments driven by the evolution of the beamline portfolio. The resources required for developing BLISS are separate from the resources required to convert beamlines from SPEC to BLISS and need to be ring-fenced for the long-term.

The conversion of beamlines from SPEC to BLISS has started on 13 ESRF and 1 CRG out of 42 beamlines. These will be commissioned with the new control system, BLISS, and the EBS Storage Ring in 2020. The conversion is a time and human resource intensive process due to the large number and complexity of experimental techniques on almost all ESRF beamlines. The effort needed to fully convert a beamline can be estimated at anything from 3 to 6 months or more for very complex beamlines. Conversion is carried out on running beamlines and therefore testing can only be done during the short shutdowns for maintenance. This means conversion will take longer in elapsed time and could be up to 18 months (with zero downtime for a running beamline). By conversion, it is meant having the full range of experimental techniques which the beamline uses converted to BLISS. Converting a beamline to BLISS is an opportunity to streamline and improve the data acquisition sequences by implementing continuous scans and more efficient data taking scenarios using the new features in BLISS. It is also an opportunity to migrate all data files to HDF5. HDF5 is a hierarchical data format which supports large multi-image data files. Large files are more efficient for data storage and transfer. The ESRF data policy has adopted HDF5 as its main format. All beamlines need to move to producing and supporting HDF5 files. The human resources to convert beamlines from SPEC to BLISS needs to be secured so that the conversion can be planned over the next few years.

Beamline control areas which have been identified as requiring effort are:

- **graphical user interfaces** (GUIs) - empower users to be more efficient during their beamtime. The goal is to offer more “turn-key” type user experiences on a maximum number of beamlines. GUIs hide the complexity of the underlying system and help the user concentrate on the essential tasks and support users to run experiments in remote access mode. Human resources have been foreseen to boost the GUI effort. The adaptation of GUIs (web and desktop) on top of BLISS for visualising the sample and preparing the data acquisition strategy are being developed. These solutions need to be generalised and further developed for all beamlines which require them.
- **mechatronics** - mechatronics is a combination of high precision mechanics, electronics and real-time acquisition and feedback to do high precision control. The need to integrate mechatronics in the beamline control is increasing due to the generalised use of nano-beams (<50 nm beam). The mechatronics solution adopted for the DCM on ID21 needs to be extended to other applications *e.g.* continuous scanning of nano sample stages (ID16A, ID31 and other beamlines in the future) and made sustainable.
- **timing and synchronisation** - the EBS will provide more photons on the sample which will allow faster experiments. With the increased flux, topping up and new timing modes, a better synchronisation of the beamlines with the accelerator timing is needed. The solution is the integration of faster acquisition electronics and the new timing system for event correlation. The whole chain for high-speed data acquisition needs to be re-designed for the EBS to be fully exploited.
- **on-line data reduction** (ODR)- automate online data reduction and compression to reduce data volumes. It is essential that data processing is completely integrated into the beamline control. This task is critical for the EBS in order to deal with the data volumes which will be produced (see section on data acquisition). The beamline control has to ensure that the ODA algorithms developed by the data scientists can be executed in real-time by adopting the right architecture and providing the necessary metadata.

[18] *https://www.certif.com/content/spec/*
[19] *https://doi.esrf.fr/10.15151/ESRF-DC-186715732*
[20] *http://accelconf.web.cern.ch/AccelConf/icalepcs2017/papers/webpl05.pdf*

Software

- **data policy** (DP) - ESRF has been a leader in the synchrotron community in implementing an open data policy[21]. The strategy for the future is to complete the implementation of the data policy on all beamlines and for all techniques. The next steps in the medium term include generalising the use of electronic logbooks on all beamlines, consolidating and extending the definitions of metadata and adopting Nexus/ HDF5 on all beamlines. The long-term goal is to make data FAIR[22] and available to the community.

- **new beamlines** (EBSL) - the EBS foresees to build 4 new beamlines[23] and upgrade a number of existing ones. The beamline control for the new EBS beamlines will need to be designed and implemented. The IT infrastructure for each beamline will need to be foreseen and installed. Each new beamline requires 1 FTE for 1 year dedicated to construction and commissioning. Extra human resources are required to cover these needs in the short term to avoid impacting the improvements needed to the existing beamlines to profit from the EBS source.

- **SPEC to BLISS** - the EBS will start up with 15 beamlines operating with BLISS. The remaining beamlines need to be converted from SPEC to BLISS over the next 5 years. This is a major task which will need significant human resources for beamline control. The initial estimates are 6 months on average per beamline. With 25 beamlines still to convert this implies 12.5 FTEs minimum over the next 5 years. Extra resources are required before 2025 to reduce the impact of the conversion on the operation.

BEAMLINE CONTROL STRATEGY RECOMMENDATIONS

- Continue converting beamlines to BLISS, the new Python based beamline control system, while enhancing the core functionality of BLISS based on the requirements of the beamlines for the EBS
- Train beamline scientists in Python programming so they can develop their own experiment sequences
- Additional human resources are required for the EBSL beamlines and conversion of SPEC beamlines to BLISS before 2025

RISKS OF NO CHANGE

- Conversion of beamlines to BLISS will be slowed down and beamlines under SPEC will be frozen
- Beamlines will obtain less operational support in the programming of experimental sequences

[21] https://doi.org/10.1080/08940886.2019.1608119, "ESRF Data Policy, Storage and Services", by R.Dimper, A.Götz, A.de Maria, V.A.Solé, M.Chaillet, B.Lebayle

[22] https://ec.europa.eu/info/sites/info/files/turning_fair_into_reality_1.pdf

[23] https://ebs.esrf.fr/2017/06/29/four-new-beamlines-for-ebs-2/

Software

Data Acquisition

Data acquisition is one of the most critical aspects of experiments especially for high-speed data acquisition. High-speed acquisition is required for high time resolution experiments (so-called time-resolved experiments) and for high throughput experiments. ESRF uses the LIMA software[24] framework for integrating detectors. LIMA is an open source software developed at the ESRF and is used by a number of other sites. LIMA has plug-ins for over 34 detectors[25] and is used to integrate all detectors used on ESRF beamlines.

The main challenges of integrating detectors and which impact the IT infrastructure are (1) **high data rates** (a few *Gigabytes/s*) which need to be sustained during the experiment, (2) **big data volumes** (tens of *Terabytes/ hour*), (3) **multiple detectors** are producing data **simultaneously**, and (4) the fact that most **detectors** are **connected** to a **single PC** which makes it very challenging to parallelise data processing and storing. Table 5 gives a summary of the number of high data-rate detectors for all beamlines in 2020.

In principle all beamlines can work simultaneously in parallel and therefore the IT infrastructure has to be dimensioned to be able to absorb up to N times the data rates and volumes produced by detectors listed above. In practice not all the detectors will work in parallel, nonetheless we can expect up to 10 high-speed detectors acquiring data simultaneously and all of the high-speed detectors working on average every 4 weeks of beam time. In the past it has not been possible to have a constant writing speed sustained above 700 kB/s. The IT infrastructure could not sustain writing constant data rates at speeds higher than 500 kB/s due to the fact that data were written directly to central storage in the majority of cases and the central resources are shared among many beamlines and the data analysis cluster. Single client performance was not guaranteed which led to data acquisition scans at high-speed crashing after an unpredictable time. With the new detectors (8 x Eiger2) procured, the required data rate to be sustained for each one is 2 GB/s. The data rates will only increase with new detectors like the Jungfrau (see below) which produces data at 9 GB/s and in the future 46 GB/s. In view of the many high-speed detectors required for the EBS it is clear that providing the correct IT infrastructure to acquire and store the data is one of the major challenges for the future

Detector	Type	Maximum data rate	# of beamlines
	PCO EDGE	1 GB/s	10
	PCO DIMAX	8 GB/s	2
	Eiger1 /Eiger2	4GB/s	2+8
	PSI Eiger	2 GB/s	2
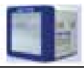	Pilatus	2 GB/s	7
	Frelon	1 GB/s	8
	Jungfrau	9 GB/s	1
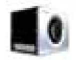	Maxipix	1GB/s	15

Table 5 : *Inventory of high data-rate detectors in 2020*

[24] https://github.com/esrf-bliss/Lima

[25] https://lima1.readthedocs.io/en/latest/supported_cameras.html

Software

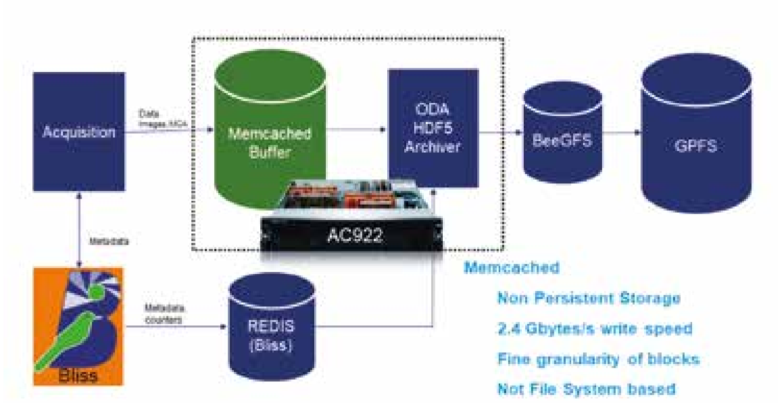

Figure 12 - *Diagram of buffering levels of online data acquisition + data processing*

The strategy proposed is to use a memory cache or local disk as a buffer dedicated to each beamline with high data rate detectors. This corresponds to 20+ local buffers. A memory cache has the following advantages (1) higher bandwidth than a disk cache, (2) clients can do online data analysis on the data in the buffer without data hitting the disk (3) data triage and compression can be done "on-the-fly" thereby reducing the data to be stored and archived. The minimum "online processing" case is considered to be combining the data in memory with metadata and writing it to Nexus/HDF5 files. The only disadvantage of a memory cache might be the need to have a very large cache due to the very high-speed of incoming data and the duration it has to be sustained. The latest generation of high-speed memory, DDR4, has a measured speed of 15+ GB/s. This is faster than the current high-speed links (100 Gbps) used to transfer data from the detector computer to the buffer computer. In case higher bandwidth is required the memory cache could be distributed across multiple computers. This would be necessary for future detectors (see Jungfrau example below). The proposed architecture for the memory cache solution is depicted in figure 12 above.

In the case where a beamline needs a larger buffer than can be stored in memory (for technical or financial reasons) the strategy is to provide a dedicated file system as a buffer based on a low-cost distributed file system like BeeGFS as depicted above. This would allow 10s or even 100s of TB to be stored locally and processed locally. This option will be proposed to those beamlines which need them.

A new major release called LIMA2 is under development. LIMA2 will be distributed (see block diagram figure 13).

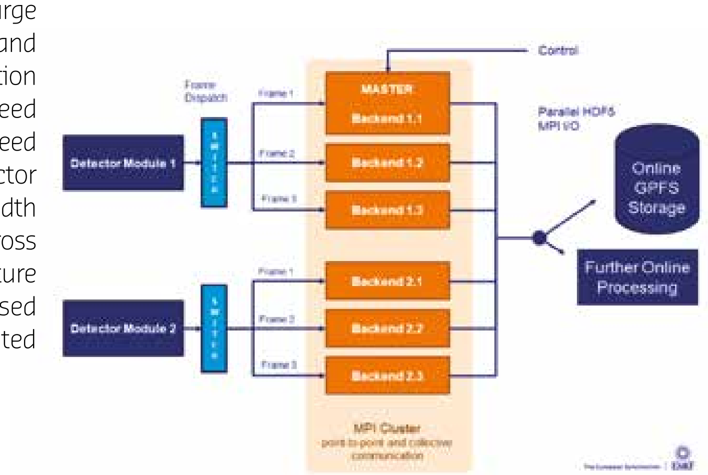

Figure 13 - *LIMA2 distributed acquisition and processing architecture*

Software

JUNGFRAU DETECTOR AN EXAMPLE OF THE NEXT GENERATION VERY HIGH DATA RATE DETECTORS

The Jungfrau detector [1] has been chosen as the main detector for the EBS ID29 serial crystallography beamline. This detector offers unprecedented performance for high-speed large area data acquisition. The first version which will be commissioned on ID29 in 2021 has 4 Mpixel at 1.1 kHz frame-rate and generates raw data at 9 GB/s. The current IT infrastructure cannot store data at this rate. To address this problem a solution based on the Power 9 computer where all data processing can be carried out on one computer has been chosen. The Power 9 offers 3 main interfaces: NVLink for CPU-GPGPU communication (up to 75 GB/s in each direction), OpenCAPI for CPU-FPGA communication (up to 25 GB/s in each direction), and PCIe 4.0 (up to 32 GB/s in each direction, double that of the PCIe 3.0) [2] which can be combined together to implement a pipeline to calibrate, reject empty frames, and compress the data to < 2 GB/s so they can be stored and further processed by the analysis programs. It is clear that we will need to reject up to 90% of the incoming frames in order to have acceptable data volumes to store and archive. A lossy compression will be used to further reduce the data volume. The proposed solution is aligned with the approach being adopted by SLS@PSI for the next generation of the Jungfrau which has 10 Mpixels @ framerate of 2.2 kHz generating data at 49 GB/s [3]. ID29 plans to upgrade to the next generation Jungfrau in the future. The architecture described here (calibrate, reject, lossy compression on Power9) implemented in Lima2 (see below) will also be used to address the very high data rates from other next generation detectors.

[1] F. Leonarski, et al, "Fast and accurate data collection for macromolecular crystallography using the JUNGFRAU detector," Nat. Methods 15(10), 799–804 (2018).
[2] S. Roberts, P. Ramanna, and J. Walthour, "AC922 data movement for CORAL," in IEEE High Performance Extreme Computing Conference (HPEC) (2018), pp. 1–5.
[3] doi: 10.1063/1.5143480: F.Leonarski et al. "JUNGFRAU detector for brighter x-ray sources: Solutions for IT and data science challenges in macromolecular crystallography ", Struct. Dyn. 7, 014305 (2020)

Online Data Reduction and on-the-fly data compression

In view of the very high data volumes which will be produced, it is essential to perform online data reduction to reduce the data to be stored and archived and to provide feedback to the users about the quality of the data to guide the experiment. Real-time feedback ensures better use of beamtime and more useful data.

Online data reduction means treating data at the speed they are produced i.e. in real-time. With high data rates it is essential to treat the data before they are written to disk i.e. in live streaming mode or via a memory cache. This implies a strong coupling between the data acquisition, experiment sequencer and data reduction algorithms. It also implies a powerful dedicated computer for the online reduction. After comparing different architectures and computers the current strategy is to use the IBM Power System AC922 as an online data processing workhorse. The AC922[26] is the building block of the world's two fastest super computers in 2019 (Summit and Sierra).

[26] *https://www.ibm.com/products/power-systems-ac922*

Software

The architecture of the AC922 includes the Power9 RISC CPU, an FPGA, PCIe4 and up to 6 NVIDIA Tesla V100 GPUs. The AC922 is the only computer currently supporting the NVIDIA NVLink high-speed transfer between CPU and GPU and system memory as an extension of GPU memory. Benchmarks demonstrate a 5.6X increase in transfer speed from CPU to GPU. Such a computer can be used to crunch data at high-speeds using either the CPU, GPU, CPU+GPU and/or FPGA based algorithms. The AC922 is equipped with an FPGA based gzip compression algorithm in order to compress data at very high rates (quasi-real time for current data rates from detectors). Very close collaboration between the beamline control, data acquisition, data processing teams and beamline scientists will be required to ensure that the data reduction is done online. A number of algorithms exist already which could be applied to reduce the data but require a modification of the experiment protocol to ensure that the calibration and required metadata are available before the data acquisition starts.

A second level of data reduction will be provided using workflows. Workflows are high level processing pipelines which can be tuned by beamline staff to suit the experiment's needs. The strategy concerning workflows is described below.

Data compression and triage will be applied as much as possible. Currently the following algorithms for compression will be evaluated: lossless ones such as gzip, lz4+bit-shuffle, jpeg2000, and blosc2 and lossy ones such as jpeg2000 and zfp. They have to be applied on-the-fly without slowing down the data acquisition process and will therefore be accessed in memory cache access before data are written to disk.

New algorithms[27] using machine learning will also be studied. ESRF is an active partner in the work package on data compression in the LEAPS proposal which has been submitted to the EC for funding and, if funded, should start in 2021.

DATA ACQUISITION AND ON-LINE DATA REDUCTION STRATEGY RECOMMENDATIONS

- Develop a distributed version of Lima - called Lima2 - for the next generation of high data rate detectors with data rates > 9 GB/s
- Integrating high data rate detectors will be based on a first level memory cache to ensure sustained high-speed data acquisition, a second level dedicated file system cache, both allowing to run online data reduction. The detector computers will be housed in the data centre to benefit from the high-speed interconnect with online computers and data storage via Infiniband
- Algorithms for data triage, compression and data reduction will be ported to the Power 9 CPU+GPU+FPGA platform to reach real-time performance. The output data will be stored in HDF5 using the Nexus conventions
- Additional highly skilled human resources are needed to implement and maintain Lima2 and to develop calibration, rejection and compression schemes on GPU+FPGA

RISKS OF NO CHANGE

- Fast and high data rate detectors will not be able to be used at full speed
- High-speed / large data volume experiments will be difficult or impossible
- Next generation of detectors (*e.g.* Jungfrau) will not be integrated in time

[27] *https://arxiv.org/pdf/2004.02872*

Software

Software Standardization

Standards are necessary to enable collaborative efforts and software projects need to have clear directives concerning licenses, target platforms, development practices, coding tools and software libraries.

Licenses

Science needs reproducibility and this is incompatible with software in black boxes. Scientists must be able to verify the results obtained, using a program that can be trusted. All scientific software has to be open source (not necessarily free). ESRF software developments are carried out under open source licenses (MIT, GPL, LGPL, BSD ...) and any other license choice must be duly justified.

Platforms

Linux is the de facto standard operating system of the ESRF software group. Windows and Mac-OS compatibility is a secondary objective but can be mandatory for software libraries and offline data analysis tools.

Coding Practices

Software must be reliable and maintainable. This implies following a set of software engineering best coding practices: version control, code review, continuous integration, and continuous documentation. These concepts are not new. What is relatively recent is the emergence of tools making it simpler to go from good intentions to actual practice. The commonly accepted standard today for version control is Git, and tools like GitLab and GitHub provide the appropriate environment to implement all the coding best practices while ensuring a fluid communication among the developers irrespective of their physical location. The ESRF software group uses Git and the forking workflow.

Computing languages

While there is no computing language that can claim to be ideal for solving all problems, the combination of Python, C and C++ is the *de-facto* solution adopted by most synchrotrons. With its syntax simplicity, its scripting capabilities and its seamless integration with C and C++, Python is becoming the favourite user front end to data acquisition and data analysis. These features have propelled Python to be a main entry point to access relatively new fields like machine learning.

The ESRF policy is to encourage data acquisition and data analysis software developments in Python, C and C++, while making sure these developments can be easily accessed from Python.

Javascript will be used for front-end web developments.

Software libraries

The use of common libraries shares the maintenance effort among the different projects instead of accumulating them. Phase 2 of the ESRF Upgrade financed the development of the SILX-kit (https://github.com/silx-kit) with the aim of providing a common set of tools on which to base further developments. The use of different tools should be duly justified and approved by the Software Group.

As a rule of thumb, we can expect that a software product needs a continuous investment in maintenance of about 30% of the time required to develop it. For any successful software project, this maintenance effort can continue for years. This is incompatible with software developments of beamline staff recruited on time-limited contracts if the developments are made independently of the Software Group and its coding practices.

Software

Data formats
Data are the core product of synchrotron radiation facilities. While ESRF made a visionary choice by starting to introduce Python in 2000, the situation concerning data formats have been far from ideal. It is not unusual to have to search for all the pieces of relevant information in different places, in different formats and even on hand-written paper notes.

The hierarchical data format HDF5[28] offers the versatility of a filesystem in a file. In simple words, if information can be stored on disk, it can be stored in an HDF5 file. This fact, combined with the availability of libraries to read HDF5 files on multiple operating systems and using multiple computing languages, make HDF5 a safe data format choice.

To enhance interoperability with other facilities, additional conventions on data layout inside the HDF5 file can be adopted. The ESRF will follow the NeXus conventions (http://www.nexusformat.org) whenever appropriate.

High Performance Computing
To achieve the best use of compute clusters, it is necessary to ensure that all programs used for large-scale data analysis (MX, Imaging...) are fully optimised. This requires specific developments and testing (such as profiling code) which can only be carried out by trained software scientists. Optimisations which make best use of CPUs and GPUs have to take into account memory and network bandwidths and can easily result in 2x to 10x gains in execution time. The difference can be the need of 10 instead 100 cluster nodes for a given analysis code. This also implies that inefficient software should be phased out to ensure the efficient use of the compute clusters.

Among software developed at ESRF during the last 4 years, several already provide standard algorithms as a toolkit which should be re-used in other programs: SILX for image processing and data reduction, pyFAI for azimuthal integration, Nabu for tomography, PyNX or coherent imaging, etc...

SOFTWARE STANDARDIZATION RECOMMENDATIONS

- All ESRF software shall be open source and accessible via Git repositories and follow best practices in software development
- The main development platform remains LINUX
- The use of libraries or computing languages other than those recommended by the Software Group has to be duly justified and approved by the latter
- Staff whose primary goal is to develop software should integrate the appropriate ESRF software-related-unit
- All software used for large scale data analysis should be fully optimised to run either on CPU or GPU clusters. Inefficient software shall be phased out to avoid wasting expensive compute resources

RISKS OF NO CHANGE

- The recommendations above are considered minimum Best Practices and have to be applied to ensure optimal use of resources

[28] https://www.hdfgroup.org/solutions/hdf5

Software

Scientific Software and communities

The ESRF welcomes more than 7000 external scientists every year, corresponding to 1800 experimental sessions. These numbers are expected to grow due to the higher photon flux from the new EBS Storage Ring. The research performed at the ESRF can be divided into 12 scientific areas: Chemistry, Earth Science, Environment, Hard Condensed Matter, Cultural Heritage, Life Sciences, Applied Material Science, Medicine, Engineering, Methods and Instrumentation, Structural Biology and Soft Condensed Matter.

Within each user community there is a wide range of expertise, ranging from expert users who are able to carry out new experiments and can analyse the data, to users who will come (or just ship samples) to exploit techniques and require an automated analysis of the data. In between we can find the majority of users who will use established techniques, and will need varied levels of help to analyse the data. Most importantly when it comes to data analysis, there are a number of software developments which are carried out by the user communities and are indispensable on ESRF beamlines.

Evolution of techniques & user expectations
While the increasing volume of data produced will be a clear challenge, another one will be to adapt to the evolution of the techniques: when the ESRF began operation, a lot of experiments were at the cutting edge of science and technology, and practically all users needed to be capable to analyse data on their own. Since then, a number of techniques have become standardised and are mature enough to rely on almost completely automated data analysis. The most outstanding example of this evolution is in the Macromolecular Crystallography community (MX) which benefited from

- the standardisation of experiments, with almost all data being collected in the same configuration (rotation single crystal diffraction near the Selenium K-edge), and
- the development of strong collaborative computational projects (initially around CCP4, then others like phenix) within the user community.

As a result, during the last 20 years the data collection of 30% of diffraction MX experiments has shifted to remote data collection (with samples being mailed in), and analysis is mostly automated.

One consequence however is that MX users today *expect* the data collection and analysis to be automated, which puts significant pressure on synchrotron facilities to provide a seamless experience. This is something that other experimental techniques will face in the near future, as many of them are becoming standardised in their data collection strategy and which will in principle move towards automated data processing, provided the necessary software is available. Among techniques which will face this challenge we can cite: tomography (absorption or phase-contrast), coherent imaging, EXAFS spectroscopy, scanning X-ray fluorescence, single crystal or powder diffraction, small-angle scattering, etc...

Making sure that software is readily available for the data produced by these mature techniques is not only important for existing users, but also to *enlarge the user community* to scientists who are not interested in the experimental or data analysis details, but only want to focus on information for their material and applications. Making synchrotron experiments accessible to a larger community is thus an existential goal – and this also includes industrial users who are demanding more streamlined data analysis.

Software

This trend will put a strong emphasis on all data analysis services in the future. It is nevertheless important to mention that there is a limit to what can be done: the processing of data can be provided by ESRF as long as it can be considered a straightforward process, with only a few parameters to adapt. However when it comes to interpreting and understanding the results – where the actual research on the materials comes into play - it has to remain entirely the user's responsibility. Thus, a continued effort to identify techniques which are mature enough should be made, and existing software should then be adapted to automated workflows.

Working with scientists and user communities to strengthen software development

In the survey on software currently used at ESRF, the beamline responsibles listed a total of ~100 unique software packages covering more than 60 techniques used for data reduction and analysis, not counting many in-house scripts mostly written in Matlab or Python. Among those, about half came from in-house developments (either from the software group or the beamline scientists), and half from the community, with only a few from other facilities. The survey also listed more than 60 different techniques used on the beamlines.

One outstanding example of an application for which updating the software is essential is holo-tomography (3D near field imaging): in this technique (on ID16B) a single tomography dataset can now be collected every 7 seconds, and requires 360 CPU cores for 1 day for its analysis. With *in situ* data acquisition running continuously for two hours, three times a day, a top-500 supercomputer would be needed for the data analysis! The correct approach, however, demands the development of an optimised GPU-based software which will lower the computing requirements by orders of magnitude. Many applications are now facing the same challenge, with the increased data throughput demanding high performance algorithms and software.

Less than half of the staff for scientific programming are available for developing, maintaining and supporting new data *analysis* software. The majority of the staff is focused either on core generic tools (SILX, workflows, cluster schedulers...), data management or maintaining existing data analysis software. Facing the data volume challenge will require a tighter integration of data reduction software with beamline control. This will require closer collaboration between the data analysis and beamline control units and therefore further reduce the staff available for developing data analysis algorithms. There is thus not enough dedicated personnel to develop new tools to cover all the needs of ESRF beamlines and users.

This analysis was already carried out a few years ago and it was decided to develop core tools to standardise and facilitate software development [SILX, see Software Standardization] and to recruit 5-year computing scientists on the following subjects: coherent imaging, core spectroscopy, diffraction microscopy, serial crystallography and tomography. These positions were recruited (from 2016 to 2019) in the Experiments Division and work in close collaboration with the Data Analysis Unit in the Software Group. The first two of these contracts will finish by the end of 2020. While these positions have been very successful so far, allowing both to produce new software but also develop new algorithms, this effort was temporary as part of the EBS upgrade plan.

Software

To move towards a sustainable plan for data analysis the following points should be taken into account:

- Existing permanent staff will not have the time to continue projects which were developed with the 5-year positions, so these positions should be maintained or expanded, either as permanent or as 5-year scientist positions on a specific subject.
- Data analysis software is primarily written by scientists in or coming from the scientific community. It is thus extremely important to collaborate as closely as possible with the Experiments Division to include such scientists in the overall development of software. This is already the case regarding coding standards (*e.g.* use of HDF5 and the SILX toolkit) but it will be important to have a stronger coordination to avoid beamline-developed software which cannot be maintained. The same applies to developments on CRG beamlines which can also be used on ESRF beamlines.
- User communities already contribute half of the software used at ESRF. It would be beneficial to coordinate and encourage these developments while following best practices. This can include scientists involved in long-term projects or externally funded projects hosted at ESRF (*e.g.* European Research Council grants). They should be encouraged to include software development as part of their project. This can be done by hosting visiting scientists for a few months to a few years. Similarly, durable relationships with applied mathematics departments should be encouraged (either from the University of Grenoble or from other European Universities).
- There are only a few data analysis projects which are co-developed by multiple photon & neutron facilities in Europe. As more experimental techniques reach maturity (stable data format and analysis algorithms), software collaborations should be encouraged on targeted subjects, while avoiding too large collaborations which are difficult to coordinate (*e.g.* take into account the number of person months devoted to the projects to give a voting weight in the steering committee).
- Training in scientific computing: this should be expanded both to train users on specific software packages, but also to train computing scientists in mastering the details of the software and toolkits developed and coordinated by ESRF so that they can contribute software based on good standards. An effort to provide training material is already included in the PaNOSC project.

Enable access to cloud and HPC computing for users

The data deluge is a challenge which puts a significant strain on the ESRF but also on our users who have to extract knowledge from the data. It is therefore essential to make sure that users have access both to adequate software and computing resources to interpret their results. This now often means having access to high-performance computing, whether this is only a few GPUs or larger clusters. As these are not always available in each user institution, it is necessary to enable access to cloud computing resources.

Currently the access to computing resources is not straightforward for synchrotron users. Most HPC centres are focused on simulation and do not encourage experimental data analysis. Moreover, the actual use of these computing centres is usually complicated and reserved for many-core applications running for many hours. It implies the need for a new peer review of the project, delay in obtaining computing time, impractical access modes (*e.g.* from a single IP address, in-person authentication needed,

Software

console only interface) and often involves a sophisticated process to install software on a cluster frontend. While this is slowly evolving, *e.g.* with the EOSC (see chapter 8), in practice only a few users will have the technological know-how to exploit cloud or HPC resources.

To improve this:

- The ESRF is coordinating the PaNOSC project (see chapter 8) which will provide access to on premise or remote cloud computing resources for Photon and Neutron facilities via a web portal. The ultimate goal of the portal is to provide data analysis as a service (DAaS). These services will enable users to run data reduction and analysis software on their data remotely. This is essential for remote experiments (see chapter 5) but also for all users who need help reducing and analysing data. Such a service naturally requires extra resources. ESRF management needs to decide how much to invest in term of human resources and hardware for such a service. A minimum service running Jupyter notebooks is already in place and is proving to be very successful.

- Training on how to use these cloud computing resources (not just publicly-funded ones but also commercial resources such as Amazon, Google, OVH) should be done, both for users and software scientists. For the latter, this should include instructions on how to package software to be readily installed in computing centres, provide automated testing, etc.

- ESRF IT staff shall prepare setting up cloud resources such that users can easily access and use those resources.

- We need to collaborate with computing centres to standardise software installation (*e.g.* via containers or virtual environments), as well as data synchronisation services.

SCIENTIFIC SOFTWARE & COMMUNITY RECOMMENDATIONS

- There is a critical need to maintain human resources devoted to data analysis by renewing the 5 scientist positions and creating others to work on Artificial Intelligence-based methods (see also chapter 4)
- Involve the scientific community (both from the experiments division and external) to contribute to writing data analysis software. This will require a strong coordination and collaboration between computing engineers and scientists
- Scientific computing training sessions will be needed to ensure that best practices are followed (to ensure maintainability of software) and enable the efficient use of computing resources (*e.g.* GPU, cloud resources)
- Evolve increasingly towards automated workflows for those techniques which have reached enough maturity, to enable expanding our academic and industrial user communities

RISKS OF NO CHANGE

- Optimization of algorithms to GPUs will be limited and more compute resources will be needed which will lead to bottlenecks in data analysis and publishing
- Data compression and triage will be limited which will lead to a storage bottleneck
- Data storage needs will escalate which will require more budget for storage

Coupling Algorithms, Machine Learning, and Artificial Intelligence

Machine Learning (ML) and Artificial Intelligence (AI) are revolutionizing many aspects of science and technology and are evolving extremely rapidly. New applications are showing up daily and have a tremendous impact across industry and science, largely due to the availability of data sets on which algorithms can train to produce effective models. It is of utmost importance that the ESRF not only embraces ML techniques in order to stay at the forefront of science but also for simply carrying out certain experiments which will otherwise be impossible. ML/AI has the potential to improve experiments and save beamtime, reduce data volumes, quickly derive insights from data and in general enhance the impact of our user facility. In addition, ML/AI can help unmask the complexity hidden in problems in high-dimensional spaces by finding connections elusive to human observations.

The three workshops held at the end of 2019[3,4] and beginning of 2020[5] have highlighted the vast potential but also the specialised know-how required to couple ML/AI to traditional techniques. The presentations and discussions which took place during these workshops suggest that ML/AI can greatly accelerate the quest to probe and understand fundamental phenomena across a vast range of lengths and time-scales, potentially leading to transformational advances across scientific disciplines. ML/AI is not only confined to data analysis, but can also contribute to optimising the control of the accelerator, experimental systems, and real-time data acquisition capabilities.

ML/AI is already routinely applied to image analysis at the ESRF in particular for data from the Cryo-EM and neuroimaging. Other experiments and methods would greatly benefit from ML algorithms for data compression, denoising, pattern recognition, or optimised data acquisition strategies. Although still under development, promising results have been obtained by automatically segmenting large volumes using a small subset of manually segmented data for training by the ASEMI project (financed by ATTRACT)[29]. Supervised machine learning is particularly well suited to experiments where the purpose is to distinguish among several phases (microscopy, speciation, segmentation...).

With the adoption and implementation of the data policy, the ESRF has laid an important foundation on which ML/AI techniques can be built: access to structured and open data. The H2020 project PaNOSC, and before that the FP7 series of PaNdata projects, have prepared the Photon and Neutron facilities in Europe to adopt similar policies and to create data archives which can now be the basis for sharing the data for machine learning driven discovery.

ML techniques are now technologically mature enough to be applied to particle accelerators[30] to meet new demands for beam energy, brightness, and stability. ML/AI methods can also be applied to automate beamline control and experimentation or to reduce dramatically tuning time and improve performance of beamline operation[31]. Deep neural networks are also natural candidates for lossy or lossless compression and could be designed to combine compression and data reduction[32].

[29] *https://attract-eu.com/selected-projects/automated-segmentation-of-microtomography-imaging-asemi/*
[30] *https://arxiv.org/pdf/1811.03172.pdf*
[31] *https://www.newswise.com/doescience/?article_id=718303&returnurl=aHR0cHM6Ly93d3cubmV3c3dpc2UuY29tL2FydGljbGVzL2xpc3Q=&sc=c57*
[32] *Sofia Vallecorsa "Comments to the ESRF Information Technology Strategy draft"*

Coupling Algorithms, Machine Learning, and Artificial Intelligence

In terms of software, the basic building blocks are readily available libraries: TensorFlow (with and without Keras, a high-level API for building and training deep learning models), PyTorch and SciKit Learn are different solutions with their own strengths and weaknesses.

ML/AI requires extensive computation as models may require to calculate and update millions of parameters in run-time for a single iterative model like deep neural networks. The platform of choice for such computations are GPUs. This implies that the need for GPUs will increase.

Adapting, implementing and operating ML algorithms requires a tight collaboration between ML specialists, mathematicians and statisticians capable of understanding the scientific problems, computer science experts and software engineers geared towards practical implementations, and domain specialists to ensure relevant and useable results.

ML specialists are in high demand on the market and constituting a small expert team of ML engineers/scientists is not easy. ML/AI projects will have to be carefully selected in the beginning to start with the most obvious problems for which there is immediate "return on investment". Such projects will automatically be of general interest in the photon community and many of them can and should be done jointly with other laboratories or at least "exported". ML/AI is still a relatively new area for which collaborations with academic groups or specialised laboratories such as INRIA[33] or CAMERA[34] shall be launched. The recent DOE programme on "Data, Artificial Intelligence, and Machine Learning at DOE Scientific User Facilities" funds a number of projects for a total of $30M and offers also an opportunity to join forces on ML/AI developments applied to light sources.

The more challenging problems will be most suitable for collaborations with Universities in the frame of PhD thesis projects.

MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE RECOMMENDATIONS

- Constitute a team of 4-5 ML engineers and data scientists to work on ML/AI projects in close collaboration with our beamline scientists. The team shall be reinforced with PhD students and external experts
- Collaborate with the community on ML/AI projects such that the results can be easily applied to national labs for similar beamlines/methods
- Update the Data Policy to make data collected at the ESRF available internally for machine learning purposes prior to the end of the embargo period
- Initiate a PhD programme dedicated to ML/AI

RISKS OF NO CHANGE

- With no additional human resources, ML techniques can only be applied to very few applications and some experiments may simply not be possible or totally inefficient
- Optimization of algorithms to GPUs will be limited, *i.e.* more compute resources will be needed
- Data compression and triage will be limited and more IT resources will be needed

[33] *www.inria.fr*
[34] *https://camera.lbl.gov/*

Remote access

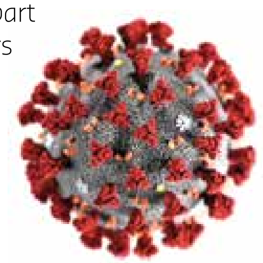

Remote access has always been part of the MX beamline strategy. Users send their samples by courier service and get the results back via internet. A lot of effort has gone into developing databases like ISPyB and ICAT, software and user interfaces which enable users to follow and even conduct their experiments remotely.

The COVID-19 pandemic has forced many people to work remotely via Internet. Many people discovered that this is in fact possible and that the tools and Internet bandwidth exist for remote work today. In the current COVID-19 phase, travel is restricted and remote access is required for as many experiments as possible. The ESRF has taken up the challenge and accelerated and extended the implementation of remote access foreseen in the STREAMLINE[35] project to all beamlines. This meant developing the appropriate mail-in and sample tracking solutions (database and web interface) and developing web-based interfaces for new techniques following the example of MX. The IT infrastructure has taken this into account by implementing a remote desktop solution and ensuring there is sufficient Internet bandwidth and server capacity for remote access. The tools developed for remote access and for data analysis and data transfer are part of the PaNOSC[36] project. This project came at the right time to help achieve 100% remote access. The H-2020 STREAMLINE project will extend these solutions further for certain high-throughput techniques to take full advantage of the EBS. COVID-19 has increased the priority on the developments in this area and the need to have dedicated IT infrastructure for remote access.

Web access

The current trend towards web-based access for most if not all user interfaces will continue *i.e.* there is no evidence that this will change in the near future. The ESRF has recognised this by moving certain applications to web portals. In addition, the ESRF is developing and/or adopting toolkits for web applications and displaying scientific data in the web as part of PaNOSC and other projects. These efforts will continue over the next five years and even accelerate as the tools mature. The next generation user interfaces for the BLISS control system are web-based and will be remotely accessible. This development takes into account that users can be remote. This means building support for user authentication and authorisation into the web applications.

Databases

Databases are essential for sample preparation, mail-in and tracking. The ISPyB database has proven the usefulness of this for MX experiments. ISPyB is very MX specific, it will be necessary to develop new databases and web interfaces to support more techniques. Dedicated manpower will be required to ensure the long-term sustainability of these solutions.

Data Portal

With the adoption of the PaNdata data policy[37], the ESRF is committed to act as curator/custodian of all data collected on site for a minimum of 5 years while aiming at 10 years. The PaNdata data policy is being updated as part of the H-2020 EOSC PaNOSC project, and effort is being invested to making FAIR[38] data available to the EOSC. With remote access being implemented, the ESRF data portal[39] is now an essential tool for users to follow their experiments via electronic-logbooks, browse and download their data. Data coming from beamlines get DOIs automatically

[35] *https://streamline.esrf.fr/*
[36] *https://www.panosc.eu/*
[37] *https://doi.org/10.5281/zenodo.3738497*
[38] *Findable, Accessible, Interoperable and Reusable (FAIR) first described in https://doi.org/10.1038/sdata.2016.18, Wilkinson, M., Dumontier, M., Aalbersberg, I. et al. "The FAIR Guiding Principles for scientific data management and stewardship.", Sci Data3, 160018 (2016).*
[39] *https://data.esrf.fr*

Remote access

minted and an electronic logbook keeps track of data acquisition and user input. Users can extend the logbook content beyond the allocated experimental time thus allowing to enrich the logbook even during the analysis phase.

Once the EOSC is a reality the data from ESRF will be exposed to a much larger community as it will be made available via the EOSC. ESRF has chosen ReactJS as the recommended front-end web development framework. The ESRF data portal (https://data.esrf.fr) follows that recommendation.

Downloading data will be essential for remote access but also for all users. Currently many users are challenged by exporting the data from their experiment. The approach of using USB disks does not scale to exporting Terabytes of data. Recently a robust solution for downloading data based on Globus[40] has been installed. Other solutions are being explored in the PaNOSC project and EOSC.

Downloading data is only possible for small to medium sized data volumes. For very large data volumes the data will stay on site at the ESRF and remote access will be the preferred or only solution. The PaNOSC project isdeveloping a data analysis as a service portal for providing remote access to generic tools like Jupyter as well as specific software for running bespoke analysis workflows. Dedicated IT infrastructure will be needed in the future to ensure the remote access users do not impact experiments while they are running. The remote data analysis as a service portal will enable easy access to the EOSC once it exists.

REMOTE ACCESS RECOMMENDATIONS

- Further develop and deploy solutions for sample mail-in and tracking for all beamlines
- Further develop and deploy the data analysis portal for remote access with appropriate load balancing and monitoring tools to avoid saturation problems
- Hire permanent resources to further develop and maintain databases and remote access tools and support users

RISKS OF NO CHANGE

- The ESRF will do fewer remote access experiments
- Users will have only basic access to data, software, and services
- Databases will not be extended/maintained for remote access for new techniques

[40] *https://www.globus.org/*

Organisational aspects

IT governance

A commonly referenced definition of IT governance states:

> *Enterprise governance of IT is an integral part of corporate governance exercised by the board and addresses the definition and implementation of processes, structures and relational mechanisms in the organisation that enable both business and IP people to execute their responsibilities in support of business/IT alignment and the creation of value from IT investments[41].*

Effective IT governance should encourage and leverage the creativity and expertise of all IT personnel, while ensuring compliance with the company's overall vision and principles. Good IT governance should be able to achieve a management paradox of simultaneously empowering and controlling. An effective IT operation does not stem merely from technology but also from effective IT governance. When conflicting interests between various stakeholders arise, effective governance provides a framework for resolving these conflicts. There is also a growing need to comply with increasing regulatory, legal, and security requirements, calling for a strong involvement of top-level management in IT governance.

IT is underpinning every activity at the ESRF. An important, if not the most important, ingredient for an effective IT governance is the constructive and efficient interaction with all stakeholders. Feedback, communication and staff training has to be more than ever at the heart of the future IT governance.

The ESRF has gone through two different organisational phases of its IT governance in the past, both of them addressed a particular situation, initially the construction phase and from 2010 onwards a focus on scientific instrumentation. The very demanding and ambitious needs to adapt to the rapid evolution due to the introduction of new and disruptive technologies and the demands of data intensive applications enabling the highest possible scientific productivity, calls for an efficient and streamlined IT governance.

The reliable but also safe and trustful operation of our IT infrastructure is of highest importance. Up to now, IT security has been approached as a team effort, coordinated by the Head of the Systems and Communication Group. It is felt that this is insufficient for the future operation of the ESRF. More time needs to be devoted to analyse weaknesses, propose solutions, train staff and stay tuned with respect to ever evolving threats and regulations. This is why we recommend creating an engineer position for a CISO (Chief Information Security Officer). This person has to be directly under the responsibility of a Director or Head of Division. The CISO would also be the key person following up the implementation of the GDPR in accordance with the recommendations of the Internal Auditor. A CISO for the ESRF will be very difficult to hire because of the high-level requirements: fluent French and English, excellent IT knowledge and communication skills, diplomatic and persuasive.

[41] *Steven De Haes, A. J. (2017). Exploring how corporate governance codes address IT governance. ISAC Journal.*

Organisational aspects

IT GOVERNANCE RECOMMENDATIONS

- Management to address the IT organization coherently with the implementation of the future IT data strategy at the ESRF
- Creation of a CISO engineer position under the direct responsibility of a Director/Head of Division

RISKS OF NO CHANGE

- The complexity of the development and operation challenges calls for a tight integration and coordination of all IT resources. Failing to do so might create unnecessary frictions in the organization and lead to inefficiency
- Not hiring a CISO will lead to a less systematic approach to cyber security and increase the exposure to security threats

Hiring and retaining IT staff

The very specific scientific environment of the ESRF requires the cultivation of talents in-house. We must hire people who can learn and adapt: IT talents must be more capable and adaptable than ever before. Traditional training and talent development approaches will not be enough, because in the decade to come we will need unprecedented flexibility and adaptability. Soft skills – things like the ability to communicate and collaborate efficiently – is now a must-have throughout IT.

Many of the changes associated with digital transformation — such as the increased role of data and analytics; the advent of ML/AI, the cloud, and mobile technologies, increasingly complex digital architectures, and the growth in personalization and user-level control, social network use, and storage capacities — not only are having an impact on the day-to-day work of IT professionals but also are creating the need for new skills and competencies. In addition, socio-political changes are introducing new expectations of the current and entering workforce at the same time that they are bringing their own shifting expectations of the workplace. All these changes are creating new opportunities and problems and demand an adaptation of human resource management.

Hiring IT staff has become a major difficulty at ESRF, but also in general in France, Europe, and beyond. In France alone, the deficit of IT staff is estimated at 200 000 persons over the next two years[42]. The number of vacancies for IT specialists in Germany is rapidly increasing and represented 82 000 open positions in 2018[43]. This deficit in IT talents concerns in particular systems and network administrators, IT security engineers, software developers

[42] *France, C. d. (2020). Le Rapport Public Annuel 2020 - Tome II - Le numérique au service de la transformation de l'action publique. Cours des Comptes.*

[43] *https://www.bitkom.org/Presse/Presseinformation/82000-freie-Jobs-IT-Fachkraeftemangel-spitzt-sich-zu#item-4086--2*

Organisational aspects

and system integrators, data scientists, and IT managers. The strong competition for talented IT staff in France and Europe makes it necessary for organisations such as the ESRF to streamline the recruitment process, reach out to schools and universities, establish a strong internal training programme, seek methods for promoting the advantages of working in a scientific environment and to adapt remuneration conditions. Particular attention has to be paid to the increasing difficulty to hire female IT staff.

Our recruitment process is not competitive and leaves the advantage to organisations able to identify, interview, and conclude work contracts in less than two weeks. Active headhunting is now the norm for finding and recruiting young and gifted people.

Over the last years we are seeing a new generation of highly educated persons entering the work place. Generation Z, born after 1995, seems to be more diverse with a different approach to authority, increased ethical standards and searching a work place providing a "cool" environment while being well-paid. The ethical considerations and flexible work environment, including measures such as working remotely for a fraction of the working time, open space office configurations, collaborative team management and regular changes in their work assignment, can be elements in favour of the ESRF if well promoted during the recruitment process.

If the difficulties for hiring IT staff persist over the years to come, it may be necessary to contract out to freelancers or self-employed workers. This may also be necessary for work which can only be done by highly specialised persons.

HIRING AND RETAINING IT STAFF RECOMMENDATIONS

- Shorten the recruitment process, develop an agile recruitment process
- Promote and develop a flexible work environment
- Establish close contacts with universities and engineering schools in Europe for sandwich courses and diploma projects
- Adapt starting salary conditions for IT staff related to the market evolution

RISKS OF NO CHANGE

- Not implementing all possible measures to hire highly competent IT staff will have a direct impact on the capacity of the ESRF to work on the challenges outlined in this report

National and international projects

The ESRF has a long tradition of working with other Research Infrastructures in Europe and beyond. This collaborative approach covers a wide range of scientific, technical, and organisational topics. It results in a constant funding stream and allows to hire additional staff whenever funding is sufficient to finance more than a fraction of a post.

The European Open Science Cloud (EOSC) is a European Commission initiative aiming at developing an infrastructure providing its users with services promoting open science practices. Besides being open science oriented, the envisaged infrastructure is built by aggregating services from several providers following a system of systems approach. The European Commission funded a large number of projects contributing to the development of EOSC, among which:

- EOSCpilot (January 2017 – May 2019)[44]
- EOSC-hub (January 2018 – December 2020)[45]
- OpenAIRE-Advance (January 2018 – December 2020)[46]
- EOSC-enhance (December 2019 – January 2022)[47]
- EOSCsecretariat (January 2019 – June 2021)[48]
- FAIRsFAIR (March 2019 – February 2022)[49]
- the thematic ESFRI Science Clusters (ENVRI-FAIR, EOSC-Life, ESCAPE, PaNOSC, SSHOC)[50] (typically from 2019 to 2022/23)
- the so-called regional Clusters (EOSC-Nordic, EOSC-Pillar, EOSC-Synergy, ExPaNDS, NI4OS-Europe)[51] (typically from 2020 to 2022/23)

The ESRF is coordinating the PaNOSC project. The partners of the PaNOSC project are CERIC-ERIC, EGI, ELI, ESRF, ESS, EU-XFEL and ILL. PaNOSC focuses on the implementation of FAIR data management, data catalogues, simulation, implementation and delivery of data analysis services, and training. The project started in December 2018 with seven partners, runs over a period of 48 months, and has a total budget of 12 M€, out of which 2M€ are allocated to the ESRF. PaNOSC works in close collaboration with the ExPaNDS project which is very similar to PaNOSC for the national photon and neutron RIs.

The five ESFRI Science Clusters participate, together with the European e-infrastructures, in the next phase of the implementation of the EOSC as defined by the H2020 INFRAEOSC-03 call. The Science Clusters are expected to bring data and scientists to the EOSC. The project proposal has been submitted and, if funded, the project will start early 2021.

The implementation of EOSC will also depend on the seamless integration of national initiatives such as DAPHNE NFDI[52] in Germany or GRICAD[53] in France.

[44] *www.eoscpilot.eu*
[45] *www.eosc-hub.eu*
[46] *www.openaire.eu*
[47] *www.eosc-portal.eu/enhance*
[48] *www.eoscsecretariat.eu*
[49] *www.fairsfair.eu*
[50] *envri.eu/home-envri-fair, www.eosc-life.eu, projectescape.eu, www.panosc.eu, sshopencloud.eu*
[51] *www.eosc-nordic.eu, www.eosc-pillar.eu, www.eosc-synergy.eu, expands.eu, ni4os.eu*
[52] *www.dfg.de/en/research_funding/programmes/nfdi*
[53] *www.france-grilles.fr/home*

National and international projects

The ESRF is member of LEAPS. LEAPS could become the platform for a structured and sustainable collaboration between all European light sources through which we could join forces for many future IT projects. The pilot project on data compression will hopefully set the scene for future common projects.

The ESRF is member of EIROforum and actively participates in the EIROforum IT technical working group. This group meets twice per year to exchange strategic information and best practices. It is also a forum aligning the positions of the eight members (CERN, ILL, EMBL, ESA, ESO, ESRF, Eurofusion, Eu-XFEL) towards the Commission.

Being present, participate and sometimes lead in international collaborations is part of ESRF's mission. However, it has always been difficult to do this, because our human resources are very much focused towards the operation of the ESRF and the ESRF has no dedicated European Office. Little time is left for projects which have no immediate and direct link to developing and operating the facility. When projects are funded (*e.g.* by the Commission), it is difficult to involve staff needed on a part-time basis, indispensable for their specific competencies. The same persons are always heavily solicited by their daily duties. It is important that once a collaborative project is approved that the top- and middle-management of all concerned groups come together to discuss the needs and implications for staff participation.

The ESRF is situated in the neighbourhood of many high-tech laboratories, industry, and the Grenoble Universities (such as the CEA, CNRS, INRIA, UGA, Grenoble Alpes Data Institute). This proximity could be an asset for many collaborations, in particular for software developments and ML/AI projects, and could also help to find some of the talents the ESRF desperately needs. However, this shall not be done to the detriment of initiatives with European laboratories and universities. A balanced national and international approach has to be maintained.

NATIONAL AND INTERNATIONAL PROJECT RECOMMENDATIONS

- Maintain a strong presence in the EOSC related projects as long as these projects align with the needs of the ESRF and/or the PaN community
- Explore local, national, and international opportunities for collaborations on software developments and ML/AI
- Continue an active and pro-active participation in the EIROforum IT-TWG and LEAPS
- Create a European Project Office at the ESRF

RISKS OF NO CHANGE

- A lack of European and international involvement of the ESRF in IT related initiatives and collaborative projects would mean missing funding opportunities and constitute a risk that developments are not aligned with other major players in the field.

Resource requirements

The financial analysis of the implementation of the IT data strategy foresee a major capital investment for the construction of a new data centre and for the procurement of required components, equipment and data storage infrastructure, which will be carried out over a five-year period (2022-2026). Beyond this period, the operation of the increased computing infrastructure will involve increased maintenance and electricity costs, as well as additional dedicated staff.

The implementation of the ESRF's scientific data management strategy is under discussion with the ESRF's Governing Bodies, with the objective of launching a ten-year resource-loaded plan in 2022.

ESRF IT STRATEGY

ESRF
The European Synchrotron

ESRF, THE EUROPEAN SYNCHROTRON

71 avenue des Martyrs, F-38043 GRENOBLE Cedex 09
@esrfsynchrotron
www.esrf.eu

DESIGN : S. ROUSSIN - PHOTO : SHUTTERSTOCK