Collaboration Agreement
in the fields of
Information and Communication Technologies

Reference KN_____/IT

(The "Agreement")

BETWEEN:

CNAF is the national center of INFN (Italian Institute for Nuclear Physics) dedicated to Research and Development on Information and Communication Technologies.
having its seat at Bologna, Italy, duly represented by its Director General, Gaetano Maron

AND:

THE EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH ("CERN"), an Intergovernmental Organization having its seat at Geneva, Switzerland, duly represented by Frédéric Hemmer, Head of CERN's IT Department.

Hereinafter "Party" and collectively "Parties".

CONSIDERING:

That CERN, an Intergovernmental Organization, is a leading global laboratory in particle physics, including in information and communication technologies, providing for collaboration of a pure scientific and fundamental character, with participation by scientific institutes from all over the world;

That CNAF, the central computing facility of INFN, is historically involved in the management and evolution of the most important information and data transmission services in Italy, in support of INFN activities at national and international level. Moreover, CNAF hosts the Italian Tier-1 data center for the high-energy physics experiments at the Large Hadron Collider, providing the resources, support and services needed for data storage and distribution, data processing and analysis, and Monte Carlo production;

That the Parties wish to collaborate in the exchange of expertise related to the operation of IT services by extending their mutual cooperation regarding information research collaboration;

The mutual benefit that the Parties would derive from such collaboration between them;

AGREE AS FOLLOWS:

Article 1: Purpose

This Collaboration Agreement establishes the framework for collaboration between the Parties for their mutual benefits in the field of information and communications technology, and in any

other area of mutual interest. The implementation of this Agreement by the Parties shall be subject to the availability of resources at the Parties. The Parties shall use the results of their collaboration for non-military purposes only.

Article 2: Areas of collaboration

2.1 This agrees aims at stablishing a framework for collaborative projects that address joint challenges in the area of operation of IT services. In this framework, the Parties shall:

- 2.1.1 Hold periodic meetings to monitor the progress and discuss the strategic direction of the collaboration.
- 2.1.2 Organise technical co-ordination meetings, liaison meetings (either face to face, or remotely by electronic mail, videoconference or telephone) for each area of collaboration as required to make progress
- 2.1.3 Propose small demonstrators or prototypes where relevant to each area

2.2 Each Party's contribution to a specific collaboration ("Areas of collaboration"), including, where applicable, the required resources, the duration of the activities and any deliverables, milestones, acceptance procedures and the management of the Project shall be set out in Annex 1 of this Agreement (as amended and updated from time to time by the mutual written agreement of the Parties). Each area shall be subject to the provisions of this Agreement, varied, where applicable, by reference to new provisions in Annex 1.

2.3 Except as agreed otherwise by the Parties in writing, each Party shall bear the cost of its participation in the collaboration and the Project(s).

Article 3: Experts

3.1 Each Party shall ensure the selection of experts with the necessary skills and competence to execute each Project on its behalf, taking into account the nature and the environment of the activities.

Article 4: Conduct and safety

4.1 The experts shall comply with the rules of conduct and safety in force at the host Party.

4.2 Any activity, equipment or other item contributed by a Party to the collaboration shall conform to the safety rules, including any specific safety requirements, in force at the host Party where such activity will be performed or such equipment or other item will be installed and operated.

Article 5: Intellectual Property

5.1 The disclosure of information under this Agreement does not create any proprietary right for the receiving Party.

5.2 Title in intellectual property developed by a Party in the execution of this Agreement shall be vested in that Party, who shall grant the other Party a free, non-exclusive license for the use of such intellectual property in the execution of its scientific programme by itself or through its partners and contractors, and for commercial exploitation, subject to any conditions set by original licensor(s) (if applicable) and/or any pre-existing rights which third party(ies) may have.

5.3 Where intellectual property is developed jointly by the Parties and title is therefore vested in them jointly, they shall grant each other a free, non-exclusive license for the use of such intellectual property in the execution of their scientific programmes by themselves or through their partners and contractors, and for commercial exploitation, subject to any conditions set by original licensor(s) (if applicable) and/or any pre-existing rights which third party may have.

5.4 As for software, the granting Party may provide software under MIT license. The receiving Party's use of such software shall be subject to the terms and conditions of such MIT license.

5.5 The granting Party provides no warranty in respect of intellectual property made available by it under this Agreement. Use or exploitation of such intellectual property is solely at the receiving Party's risk and responsibility. The receiving Party shall hold the granting Party free and harmless from any liability arising from the receiving Party's use and/or exploitation (including, where applicable, by its partners and contractors) of such intellectual property.

Article 6: Publications

6.1 The Parties shall strive to jointly publish the results of the Projects as Open Access publications.

6.2 Insofar as the Parties do not jointly publish the results of a Project, publications by one Party involving results developed by the other Party shall be subject to the latter's prior written approval, which shall not be withheld unreasonably.

6.3 Publications shall acknowledge the collaboration between the Parties including, whenever appropriate, the experts having taken part in the development of the results covered by the publication.

Article 7: Confidentiality

The Parties agree to execute the Projects in a spirit of openness. However, where, exceptionally, confidentiality is required, the following provisions shall apply:

7.1 Each Party shall treat as confidential any information provided to it by the other Party that is designated as confidential. Except as agreed otherwise in writing, this obligation shall continue for a period of five (5) years from the date of termination of this Agreement.

7.2 The receiving Party shall:

(i) not use confidential information for any other purpose than for the execution of this Agreement;
(ii) limit the circle of recipients of such confidential information on a need-to-know basis and ensure that the recipients are aware of and comply with the obligations as specified in this Article.

7.3 No confidentiality obligation shall apply to information which:

(i) the receiving Party demonstrates was in the public domain prior to its communication by the disclosing Party;
(ii) became part of the public domain after such communication but not through any fault of the receiving Party;
(iii) was already in possession of the receiving Party at the time of signature of this Agreement;
(iv) has been lawfully received by the receiving Party from a third party without any confidentiality obligation; or
(v) has been developed by the receiving Party independently and outside the scope of this Agreement.

Article 8: Liability

8.1 Except as provided in Articles 3.2, 5.5 and in this Article 8, each Party shall bear its own loss and damage in connection with this Agreement.

8.2 Subject to Article 8.3, the responsible Party shall indemnify the other Party for its loss and damage resulting from gross negligence or wilful misconduct by the responsible Party, or a violation by the responsible Party of the rules of conduct and safety in force at the host Party.

8.3 Notwithstanding the foregoing, the Parties shall in no event be liable to each other for any consequential loss or damage, such as loss of income or of availability of data or installations.

Article 9: Entry into force, duration and termination

9.1 This Agreement shall enter into force upon its signature by the Parties. It shall remain in force for the duration of the collaboration, unless terminated by joint written agreement, or by one Party giving at least six (3) months prior written notification to the other Party. Except as otherwise agreed by the Parties, termination of this Agreement shall be without prejudice to the completion of outstanding Projects set out in Addenda to this Agreement.

9.2 In case of a substantial breach by a Party of its obligations under an Addendum to this Agreement, the other Party may terminate that Addendum in whole or in part if no corrective action satisfactory to the other Party is taken within one (1) month of the issue of a letter of notice by the other Party to the breaching Party.

9.3 Articles 3.2, 5, 7, 8, 9.4 and 10 of this Agreement shall survive its termination, howsoever caused.

Article 10: Governing law and dispute resolution

10.1 The terms of this Agreement shall be interpreted in accordance with their true meaning and effect and as a consequence of CERN's status as an Intergovernmental Organization, independently of national and local law. If this Agreement does not expressly stipulate, or any of its terms is ambiguous or unclear, then in those circumstances only and not in respect of this Agreement as a whole, reference shall be made to Swiss substantive law.

10.2 The Parties shall settle any difference concerning this Agreement amicably. Where this is not possible, the Parties shall resort to arbitration in accordance with a procedure to be

specified by the Parties. Notwithstanding reference of the dispute to arbitration, the Parties shall continue to perform their obligations under this Agreement.

Article 11: Amendments

Any amendment to this Agreement shall be made in writing and signed by the authorized representatives of the Parties.

Signed in two copies, in English:

Date: Gaetano Maron Director General CNAF	Date: Frédéric Hemmer Head of the Information Technology Department European Organization for Nuclear Research

Annex 1: Areas of Collaboration

Listing and description of Areas of Collaboration:
https://docs.google.com/document/d/1Q8LtmjCORjQSDO1ngVd80v-PZESuCpTK3HiUUXqpJNk/edit

Topic 1: Large scale infrastructures management, including KPIs and monitoring (Tim Bell; Andrea Chierici and Cristina Duma)

For all the topics there is interest in creating a common knowledge base regarding the installation and configuration of the different services, build or contribute to already existing repositories of images (container registries), CMs (configuration management systems such as puppet, ansible). The main areas of interest are:

1. Kubernetes and virtualisation (D. Michelotto, A. Ceccanti)

Deploying cloud services such as VMs and Kubernetes for science creates some unique challenges in areas such as scheduling, performance optimization, auto scaling and resource tracking/accounting. There are many potential improvements in these areas with integration of common solutions such as identity management solutions (like Keycloak or IAM), CEPH, EOS and CVMFS through the Kubernetes interface.
Specific topics of interest:

- Deployment and operations of an K8S production level cluster
  - Pros and cons of the different approaches of deployment, on top of OpenStack vs. bare-metal, based on experiences of CERN & CNAF (CNAF experience lies in the deployment on top of OpenStack, not leveraging the Magnum service).
  - Large scale deployments of K8S cluster – based on experiences of CERN large deployment, strengthen the deployment of the production K8S cluster, regarding its high availability, security, scalability, resilience, monitoring and resource management

b. Central management of small K8S clusters
   - Many use cases require the creation of K8S clusters on which to deploy applications or workflows. It could be interesting to provide a Kubernetes as a Service, with a central point of management of multiple clusters, hence easing the management, monitoring, and the security configuration. CNAF has been testing Rancher for such a solution in cloud environments (INFN-Cloud and Cloud@CNAF) and would be interested in co-evaluating pros and cons, as well as adding features like authentication through Keycloak, other OpenID Connect providers.

c. Utilize high-level templating languages (e.g. TOSCA, Ansible) to express user requirement and workflows. This might include both end-users and community-specific services (e.g. CVMFS; squid proxies). The templating language, coupled with appropriate resource and container orchestration will then deploy the applications into one or more

data centers (see e.g. DODAS). Some use cases that could be supported in this way are:

- Big Data Pre-Post processing – automated deployment of processing engines and solutions (such as HDFS, Spark).
- Exploitation of Machine Learning technologies – building of on-demand training facilities, deployment of inference engines (see Topic 1.3.b).
- Batch System as a Service – deployment of HTCondor batch systems, HTCondor federation solutions for HPC/HTC extensions (see Topic 4).
- Data Caching as a Service – e.g. XRootd/http federations on demand (see topic 3.2)

2. Large scale deployments

Handling the lifecycle of bare metal from its original delivery to retirement requires a set of workflows and procedures. At CNAF we still rely on traditional manual approach and feel the need to better handle the processes.

1. Many operations such as inventory, benchmarking, firmware management, remote access and cleaning should be improved and possibly automated
    - Procedures put in place at CERN may significantly help our handling of machine lifecycles.
    - OpenStack Ironic may be of help as well, but we primarily feel the need to handle the post procurement process.

2. Accounting of cloud resources usage: An important aspect that seems lacking in community support is the accounting of cloud resource usage. At this moment both CNAF and CERN are using custom, in-house developed, systems. We can foresee a collaboration in defining and adopting a common solution, sharing experiences and eventually make potential contributions to guide the development directions of the upstream open source community

3. Building a big data platform in order to collect, analyze and store logs, metrics and system messages
    - serve system administrators for monitoring and alerting
    - extrapolate KPIs values
    - do predictive maintenance for hardware
    - support final users for on demand services for big data analysis and data stream processing

3. GPUs

a. The use of GPU accelerators potentially provides significant performance improvements. The cost of these devices, however, means that efficiency of usage is key. This area would investigate different use cases from batch to notebooks, allocation policies such as quotas, accounting, benchmarking and potential use of vGPUs.

b. Another interesting aspect to investigate is the exploitation of GPU nodes attached to Kubernetes clusters in particular for ML and DL applications. CNAF has experience on this kind of setup (a use-case in the DEEP-HDC project) and would like to continue the investigation in order to be able to propose production clusters, possibly integrated with IAM AAI.

4. Network

CNAF is designing from scratch the entire LAN of the new datacenter. A collaboration with CERN on network design could be very useful to understand how to build a data center network, able to match the performance requirements for the HL-LHC era and the dynamicity that an orchestrated and virtualized computing infrastructure needs to be efficiently deployed. The network should enable different scientific communities to share the core and access infrastructure maintaining the necessary segregation for security and to differentiate the WAN access when needed by the collaborations.

A good understanding of the integration between the physical network elements and the different overlay technologies that will be used in the data center is crucial in the network design.

A methodology for automatic deployment and management is an important factor we need to study to build a scalable and dynamic network infrastructure (e.g. Ansible or proprietary solutions if more efficient and affordable).

We are interested in evaluating the feasibility of a potential CERN-CNAF DCI. CNAF adopts different network extensions for DCI, some realized using overlay technologies like L3 MPLS VPN (CNAF-BARI) and the Terabit DCI extension to CINECA using an underlay transponder approach. It could be interesting to test a long distance and high capacity DCI between CERN and CNAF using a Packet/Optical Transponder based approach involving GARR and GEANT.

Within a multi-tenant environment, software defined networks can provide flexibility to provision custom network configurations such as container environments like Kubernetes. Advanced functions such as load balancing can be applied to create high availability. Interactions between bare metal servers, physical switches and the cloud environments need to be understood.

CNAF is also interested in exchanging views on how to provide VPN-like external access to the data center.

Topic 2: AAI and security (Hannah Short and Paolo Tedesco with input from the CERN Security team; Andrea Ceccanti and Vincenzo Ciaschini)

The objective of the collaboration is to share knowledge and common approaches on AAI and security. A list of possible topics for collaboration is the following.

AAI:

- **Keycloak/FreeIPA**: Keycloak and FreeIPA represent core components for the CERN AAI. CNAF is looking at these technologies to evolve its AAI and implement unified centralized management of users/groups/policies for all data center services. CNAF is also moving the IAM codebase (to be used by WLCG) to Keycloak. The common objective here is to share experience and adopt common

approaches to Keycloak/FreeIPA deployment, configuration, operations and custom extensions development in order to join forces and avoid duplication of efforts. Common security assessment of the chosen approaches is also of key importance.

- **Support IAM deployment for LHC VOs**: LHC VOs will soon deploy IAM instances to support the WLCG transition to token-based AAI. CNAF can help in setting up the deployments at CERN, including integration with CERN's HR Database, and provide a privileged support channel to CERN for IAM operations.
- **CLI tools for token-based AAI**: while Web-based token-based authentication flows are well understood, not much interest from the industry focused on CLI usage. CERN and CNAF will join forces in evaluating current tools (e.g., oidc-agent) or developing new ones to address WLCG and other communities' use cases.
- **Secret management**: as token-based AAI is deployed to WLCG sites, new approaches to secret management, in particular to support long-running jobs, have to be investigated. The collaboration will focus on evaluating how standard industry solutions like Hashicorp Vault could be used to support WLCG and other communities use cases.
- **User provisioning and de-provisioning**: harmonised approaches to user registration, life cycle management, provisioning and deprovisioning will be investigated, possibly converging, in the long term, on a co-developed open source solution.

Security:

- **Incident management, training and security knowledge sharing**: wider sharing of security-related information, like compromises or attacks, can only lead to better security for both centers. Exchange expertise in incident detection, handling, and response. On the same line of reasoning, both centers have experts on security related issues. Organizing recurring workshops and cross-trainings among them, each dedicated to a single issue, can improve knowledge and preparation on both sides.

Topic 3: Storage and data management (Simone Campana and Jakub Moscicki; Vladimir Sapunenko and Alessandro Costantini)

1. CEPH and EOS as possible GPFS replacement (storage team CERN, storage team CNAF)

GPFS-replacement exploration. Taking into account the current scope and objectives of EOS and CEPH testing at CNAF identify the critical storage areas and compare functionality, resilience, performance, usability (user's point of view) as well as costs (TCO). Collaborate on portable EOS deployment and coordinated effort in testing and validation of GPFS-replacement scenarios. Data flow between HPC storage and general user storage. HPC storage systems are specifically difficult to access from outside the HPC cluster. Explore possibilities to use cloud technology to make it easier to export/import user files, while keeping the workspace for HPC-workloads in dedicated HPC storage.

CEPH and EOS are part of the storage services portfolio at CERN. At CNAF, CEPH is seen at the moment as a promising candidate to replace our current storage solution based on IBM Spectrum Scale (GPFS).

CNAF also is looking at CEPH for its object storage capabilities and for the support to POSIX use cases (through CephFS), required by some communities. To test the different features of CEPH, a limited testbed has been deployed at CNAF from which we have learned many important things. A much bigger installation is needed to better understand its features.

The following activities are proposed:

- Advanced testing of CEPH. CNAF expectations from such activities are focused on comparing functionalities and performance at scale. Interest here will be also devoted to understanding the real need of POSIX standard for the supported communities.
- Evaluate the possibility of interaction among the CEPH services and other storage resource manager solutions adopted at CNAF (e.g. the tape system).
- Evaluate resiliency of the stored data.
- Evaluate the usability both from admin's and user's point of view.

The experience at CERN in the deployment and management of different CEPH installations can really speed-up the above-mentioned activities and they are a matter of collaboration among the partners involved.

CNAF has also gained some experience, in the past months, in the deployment and configuration of EOS storage service. Unfortunately, we have found some difficulties with the documentation, which has prevented further exploration of this solution. Access to an external testbed located at CERN would likely simplify testing of this solution.

The activities and the related results and considerations should be used to define the Total Cost of Ownership (TCO) of the proposed solutions (for production purposes) confronting them with other technologies.

2. Remote access in the context of DOMA and ESCAPE (Simone, Xavi, Daniele Cesini)

A quite accepted vision is the enabling the (transparent) access to data, stored in a DataLake (as being implemented by ESCAPE project) from a computing infrastructure without permanent storage such as the cloud (e.g. INFNCLOUD for INFN). This requires on the one hand the adoption of distributed caches and, on the other, tools to facilitate federation and usage of storage resources. Projects such as the INFN IDDLS project can also be used to test and verify dynamic connection to data lakes.

The solutions used to create, manage and access DataLakes for huge collaborations (DOMA, ESCAPE, XDC) will be investigated for the transparent creation of redundant federated storage resources, i.e. automatic replication in different, geographically distributed, availability zones done at the infrastructure level, without the user intervention. The approach of storage events notifications that trigger RUCIO-orchestrated data movements could be investigated to realize a quasi-real-time replication, with the possibility of specifying different QoS for the replicas. This activity is tightly linked to those described in section 3.3 for the creation of easy-to-use data management and sync-and-share services for "smaller" collaborations and for the long tail of science.

The following activities are proposed:

- Evaluate the adoption of a community-independent, storage management orchestrator (based on RUCIO).
- Evaluate the possible integration of RUCIO-based data placement policies with orchestrator-based application allocation policies.
- Integration of the currently adopted storage resource manager (StoRM) at CNAF with the storage management orchestrator.
- Evaluate different mechanisms to implement the DataLake solution.
- Evaluation of the storage events notifications approach for quasi-real-time replication through different storage solutions.

3. FTS general purpose service for general science

This project also targets the non-LHC experiments and the "long tail of science". Those scientific communities would not invest nor need a very sophisticated data management system, but benefit from a lightweight point-to-point data transfer service. The project aims at providing an enriched data management Web service, addressing the functionalities needed by those communities with a simple interface and monitoring. Ultimately, the product could be integrated with sync-and-share services and become a free-for-use competitor of Globus Online. Examples of similar integrated service environments include end-user Cloud storage for CMS analysis [1] and multi-cloud deployment with ScienceBox [2]. A candidate for the present activity is represented by FTS where CERN has both the development and operational expertise of the FTS service.

The definition of a simple interface for data management is an interesting activity for CNAF. The long-time experience of CNAF providing resources and support to many and diverse scientific experiments in the "long tail of science", which are not familiar with Grid technologies, can be used in the project to evaluate different sync-and-share solutions to be adopted for production purposes. In this schema, the evaluation of the FTS service, including its integration with higher level orchestration tools, can be also of interest.

The following activities are proposed:

- Evaluation of the available data management solutions to be offered to not-experienced user communities and selection of the most suitable one.
- Integration and/or development of a common sync-and-share solution for end-users
- Evaluation of the proposed solution and testing of the integrated services
- Feedback loop among users and developers
- Integration with RUCIO and FTS as backend services to steer data movements based on policies in a transparent way for the end-user.

The project could potentially attract EU funding that should be invested both in the development and commissioning efforts.

[1] “Integration of end-user Cloud storage for CMS analysis” https://doi.org/10.1016/j.future.2017.04.021
[2] “ScienceBox: Kubernetes + Storage + Jupyter Notebooks” http://cern.ch/sciencebox
[3] “Declarative Big Data Analysis for High-Energy Physics: TOTEM Use Case” https://dx.doi.org/10.1007/978-3-030-29400-7_18

4. CTA

The CNAF tape system is currently managed through IBM Spectrum Protect (formerly Tivoli Storage Manager - TSM) that interacts with GPFS, StoRM and GEMSS (Grid Enabled Mass Storage System), an in-house developed software layer that manages migration and recall queues. This is considered a suitable and reliable solution. At the time of the present document, the evaluation of other solutions is not foreseen: this will be reconsidered if EOS will be chosen as a storage system.

Topic 4: Integrating HPC in a WLCG site in a transparent way (Maria Girone; Tommaso Boccali and Stefano Dal Pra)

1. Method of use:

a. A preferred way to access remote computing resources by the experiments is via an elastic extension of existing sites. Tier-1 sites are clearly more trusted, have closer links with WLCG and that’s enough to guarantee that transient or opportunistic resources accessed through WLCG sites should be considered reliable and valid.

b. Even though each experiment may require a specific policy, the following solutions can satisfy them:

i. Access HPC resources through a specific dedicated End Point (managed by a Computing Element). In this case it is up to the experiment to push specific workflows (like low mem, low IO, or combinations of these) known to run on the computing resources related to that endpoint. The drawback, from the experiment point of view, comes exactly from the need to “know” and manage more and more site particular cases, and dispatch there only reasonable workflows. This easily produces several specific workflows, with loss of generality.

ii. Having the HPC resources at the same level of other site resources, and delegate the decision process of "what to run" to some late binding match making
    1. For example, if using pilots, via a cherry picking of jobs having a well crafted set of ClassAd requests
    2. Or via pre-exec scripts
    3. Or via the job inserting specific resource requests in the LSF queue (for example)
    4. In general ii-1 is the best for experiments since does not need central experiments actions

c. HPC will not normally have persistent storage. Data access at scale by HPC sites will need to be done via cache, remote access, and/or dynamic placement

. Exercise the infrastructure with the **Data Access Demonstrator**
    1. Deliver a multi-petabyte data sample on the time scale faster than the expected processing time
    2. Exercise intelligent caching systems at the scale needed to feed HPC scale workflows

2. Accounting needs to be looked on

. Should be easy if we use a CE, otherwise it has the same problematics of running with systems like VAC: the worker node is in charge to report

. Problem: how? Can the WN reach the collector?

i. In that case, only rely on the experiment accounting ? (not nice)

ii. In order to have meaningful accounting we need an objective measure of the performance. The **Benchmarking workflow demonstrator** is a standalone containerized workflow suite to measure the performance of different computing resources for our workflows.

3. Authorization, authentication

. We need to make sure the site trusts the external mechanism (AAI). To establish trust at the HPC sites in the authorization and authentication model we will need to demonstrate access and auditing of decisions.

. Authenticated Workflow Demonstrator to show that the workflows performed by VO users can be securely supported by HPC, a demonstrator on authenticated and authorised job submission would be needed. For the future OAuth2.0 based models, envisaged in the next 5 years, HPC sites should support standard OAuth2.0 flows

4. Portability libraries

. The effective use of heterogeneous architectures which are available at HPC centers and Cloud providers will largely depend on addressing the problem of efficiently programming machines with increasingly heterogeneous computational resources. Programming tools to best map programs to processors and memories are being investigated. CUDA, Alpaka/cupla, Kokkos, SYCL, Data Parallel C++ (DPC++) are examples of such tools of interest to our community.

. **OneAPI Demonstrator:** Intel oneAPI is a toolkit based on SYCL and its extension DPC++, bundled with performance analysis and debugging tools, that promises a unified, standards-based programming model that allows to produce software that could be deployed on heterogeneous platforms. We propose to carry on a common investigation of the portability of the reconstruction and simulation workflows of the LHC experiments using Intel oneAPI.

Topic 5: Quantum simulation (Alberto Di Meglio; Francesco Giacomini)

New computing architectures are emerging that could offer interesting opportunities of acceleration for various HEP workloads. Examples include more traditional accelerators like GPUs and FPGAs, new classes of Machine Learning/Deep Learning accelerators, and non-Von Neumann architectures such as neuromorphic and quantum devices.

To fully assess and harness their potential, these architectures require extensive investigations to concretely understand their role and impact of HEP workflows. It is also important to understand where effort is needed in terms of algorithms, tools, programming languages, and hardware resources.

This collaboration targets longer-term joint investigations on these novel architectures, mainly quantum computing, although the same approach could be extended to other architectures such as neuromorphic computing.

The areas of R&D that could be covered in this collaboration are:

1) **Identification of HEP workflow elements** and test use cases that could benefit from these new technologies. Notable among those are

    a. Online computing, trigger and raw data anomaly detection;
    b. Simulation;
    c. Reconstruction / tracking;
    d. Data analysis, classification and filtering;
    e. Detector alignment;
    f. Simulation of quantum effects;

2) **Survey and Investigations** of the algorithms (quantum and neuromorphic) applicable to the identified use cases and their requirements in terms of tools and computing resources. Identification of the areas where novel algorithms may be needed, how they are supported by simulators and development kits, and of the resources needed to develop them. Some of the algorithms for example to be considered can be

    a. Variational Eigensolvers;
    b. qSVM algorithms;
    c. qGAN algorithms;
    d. Other qML/qDL algorithms;
    e. Distributed quantum algorithms;
    f. Quantum Graph Networks;
    g. Quantum Homomorphic Encryption algorithms;
    h. Direct simulation of quantum systems;

3) **Design of a proof-of-concept architecture** of distributed resources where quantum computing simulators can be evaluated and initial benchmarks and development can be performed, making it easier for developers to set up the required tools and environments

    a. Set up a shared pool of hardware resources with cloud access

i. On-premise: mainly CPUs to start with as support for GPUs in current quantum simulators is still in its infancy (but it could be investigated later)
ii. Cloud: integration of current public resources such as IBM Experience or Amazon Braket

b. Set up a shared library of code, examples, tools
c. Design access interfaces: Jupyter-based interface to create shared notebooks for initial development tests, training, building expertise and giving access to the library of code and tools and the different types of available resources

Topic 6: Data Centre Certification (Alberto Di Meglio; Barbara Martelli)

Use cases where one would seek certification. Medical use cases that require a certified environment. The project would be dedicated to the study of technologies and organizational measures useful to manage experiments and research communities dealing with personal data (e.g. genomic data, laboratory analysis, nutrition, etc.) or in general data with strong confidentiality and privacy requirements.

The activity won't be addressed to gain some form of certification. The focus will be on procedures and technicalities able to solve the problems. In this context the certificate is a corollary, not the main objective.

Two, complementary, approaches will be investigated:

- Organizational approach:
    - Investigation of ISO/IEC 27001, 270017, 27018, 27701 frameworks and implications of their adoption in INFN and CERN projects.
    - Study and definition of de-facto anonymization procedures in the context of genomic data management
    - Threat and risk management: ISO/IEC 31000, other frameworks?
    - Integration of Privacy-by-design principle in our context:
        - 1. Proactive not Reactive; Preventative not Remedial
        - 2. Privacy as the Default Setting
        - 3. Privacy Embedded into Design
        - 4. Full Functionality – Positive-Sum, not Zero-Sum
        - 5. End-to-End Security – Full Lifecycle Protection
        - 6. Visibility and Transparency – Keep it Open
        - 7. Respect for User Privacy – Keep it User-Centric

- Technical approach:
    - Application of encryption techniques such as homomorphic or functional encryption allowing to share data and perform analysis over public/shared resources while preserving the privacy of the content (securing data in transit)
    - Federated learning approaches, differential privacy (in the context of machine learning)
    - Encrypted file systems (LUKS or similar) (securing data at rest)
    - Data pseudonymization (for clinical applications)

[1] “Pseudonymisation techniques and best practices”, Enisa, ISBN 978-92-9204-307-0, DOI 10.2824/247711

[2] “Privacy by design in Big Data”, Enisa, ISBN: 978-92-9204-160-1, DOI: 10.2824/641480

[3] “Big Data Security”, Enisa, ISBN 978-92-9204-142-7, DOI 10.2824/13094