

# **International Symposium on Grids & Clouds (ISGC) 2026**

## **Report of Contributions**

Contribution ID: 96

Type: **Poster Presentation**

## from Open Access to FAIR

(The proposed topics of FAIR mainly focus on technology issues, such as making data FAIR in a machine-actionable way or generating automated metadata, not only by humans but also by algorithms.

Nevertheless, how to impose FAIR principles on the researchers and research institutions to make their data easily findable and accessible is indeed a policy issue, and a governance or a guideline is necessary for facilitating FAIR.

If you kindly allow me to focus on the imposition of the FAIR principles and the policy issues, I would like to prepare a draft of my talk for the session on FAIR. The following is an abstract for my talk.)

The freedom in cyberspace is essential these days for interactive scientific research. The freedom includes Open Access, which allows the researchers to search, access, and download any articles which are published and uploaded on internet. More or less, the access has been open to a limited people (members of a particular academic society with annual memberships, or who pay subscription fees for the journals, or who are willing to pay for each download). This kind of closedness led to the accident of WASET.

Nowadays, Open Access is widely accepted and established by many academic societies and the relevant publishers as well.

And Open Access has been evolved to Open Data, and to Open Science. In between, the principle of FAIR came in. Many governments and international organizations are engaged in establishing and imposing the FAIR principles to the scientific data trade and distributions in cyberspace.

I would like to introduce the flowchart from Open Access to FAIR principles, and compare the OECD version of the FAIR principles with those of UNESCO. Also, I would like to report the stance of Korean government toward the FAIR principles.

**Primary author:** OH, Sun Kun (Suranaree University of Technology, Thailand)

**Presenter:** OH, Sun Kun (Suranaree University of Technology, Thailand)

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data

Contribution ID: 97

Type: **Oral Presentation**

## Boosting AI for Science by Highly Available Cloud Services in the Era of Agentic AI

*Wednesday, 18 March 2026 15:06 (22 minutes)*

The shift from conventional computing, networking, and storage to AI-driven scientific discovery (AI4S) calls for a new generation of intelligent infrastructure. In this era of agentic AI, scalable access to models, tools, data, and agents has become critical—serving as essential utilities powering next-generation research. To meet these demands, the HepAI team has developed Qionwu, a core technology that supports unified integration, dynamic authorization, and high-concurrency API services for large models, scientific tools, and diverse AI agents. The system reliably enables various AI applications—including data agents, physics analysis assistants, and conversational bots. This foundation paves the way for large-scale autonomous exploration and collaborative research involving thousands of intelligent agents. In this talk, we will share the design and implementation of our highly available cloud service infrastructure based on Qionwu, with a special emphasis on how it powers physics analysis agents such as Dr.Sai. We will demonstrate its capabilities through real use cases in scientific domains and discuss future directions for agentic AI in accelerating scientific breakthroughs.

**Primary author:** ZHANG (✉), Zhengde (Institute of High Energy Physics (IHEP), CAS, China)

**Co-authors:** QI, Fazhi; GANG, Chen

**Presenter:** QI, Fazhi

**Session Classification:** Artificial Intelligence (AI) - I

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 98

Type: **Oral Presentation**

## Doing Science in the Cloud: Collaborative, Reproducible Workflows with DECTRIS CLOUD

*Friday, 20 March 2026 09:00 (22 minutes)*

Modern instruments generate enormous datasets, but scientific insight depends on how quickly and collaboratively those data can be processed. DECTRIS CLOUD is shaping a future where researchers no longer move data, but move science into the cloud. The platform provides ready-to-use environments where experimental data, processing pipelines, and computational resources are instantly available.

Researchers can launch, share, and re-run their analyses within minutes, ensuring full reproducibility and transparent provenance. By encapsulating jobs as reusable templates, DECTRIS CLOUD turns every analysis into a living, shareable workflow instead of a one-off script.

The service fosters collaboration in the clouds: teams can co-develop and execute jobs, inspect results in real time, and reuse validated pipelines across facilities and disciplines. This approach lowers the entry barrier for advanced techniques, accelerates the path from data to insight, and promotes open, FAIR-compliant science.

We present the vision and ongoing developments that position DECTRIS CLOUD as a collaborative platform where scientists focus on doing science, not on managing infrastructure.

**Primary author:** VAN DIJKEN, Daphne (DECTRIS Ltd.)

**Co-author:** THERE ARE NO CO-AUTHORS

**Presenter:** VAN DIJKEN, Daphne (DECTRIS Ltd.)

**Session Classification:** Virtual Research Environment (VRE) - I

**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 99

Type: **Poster Presentation**

## Fair-IHEFT:A fairness-aware algorithm for multi-workflow scheduling in High Performance Computing Environment

Devising a multi-workflow scheduling algorithm is paramount to explore high performance from High Performance Computing Environment. In this article, we propose a new list scheduling algorithm taking fairness into consideration for assigning multi-workflow to heterogeneous processors. The proposed algorithm, a Fairness-aware Improved HEFT (Fair-IHEFT) algorithm is devised to schedule multiple workflows to optimality while minimizing the scheduling length by calculating task prioritization based on workflow structure. The experiments based on the workflows of four real-world applications show that the Fair-IHEFT algorithm significantly outperforms previous list scheduling algorithm in terms of makespan, fairness, makespan standard deviation, and frequency of best results. Additionally, job runtime is a crucial parameter for scheduling algorithms. To obtain accurate job runtime, we propose a job runtime prediction framework that can integrate a variety of machine learning algorithms to predict job runtime.

**Primary authors:** CAN WU (✉); Mr HAILI XIAO (✉✉); Ms XIAONING WANG (✉✉); Ms SHASHA LU (✉✉); Ms RONG HE (✉)

**Presenter:** CAN WU (✉)

**Session Classification:** Converging High PerformanComputing Infrastructures: Supercomputers, clouds, accelerators

**Track Classification:** Track 9: Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

Contribution ID: 100

Type: **Oral Presentation**

## High-Quality AI4S-Oriented Scientific Data Ecosystem for Large Scientific Facilities

*Thursday, 19 March 2026 12:06 (22 minutes)*

Large Scientific Facilities such as synchrotron radiation facility (e.g., BSRF, HEPS) and spallation neutron sources (e.g., CSNS), are generating massive, complex, and heterogeneous datasets continuously during routine operations and scientific experiments. Managing and utilizing the diverse experimental data, along with simulation results and literature-derived information, is presenting a critical challenge not only in cutting-edge scientific research, and also for industrial applications in the current AI-driven era. We designed and developed a comprehensive data ecosystem aimed at maximizing the scientific data utilization from such facilities, automatically generating high-quality and AI4S-oriented scientific datasets through AI technologies.

First, this paper presents a data ecosystem designed to enhance data quality, accessibility, and reusability, with the ultimate goal of generating high-quality scientific datasets and database. The data ecosystem comprises three core components: data policies & standards, data software & tools, data fusion & provision. Data policies and standards form the foundational element, establishing top-level guidelines and operational frameworks for data collection, storage, sharing, and management. They ensure that data generated across different facilities are interoperable and reusable, enabling cross-facility data sharing. Data software and tools support the entire lifecycle of scientific data, spanning acquisition, management, processing, and analysis. These tools help maintain data quality right from the source of data acquisition. Data fusion and provision utilize data AI agents to align experimental and simulated datasets according to specific research objectives, producing AI-ready datasets. This component also offers versatile APIs and interfaces, facilitating flexible and efficient data access and utilization.

Second, we will report the progress in each component of the ecosystem. In the area of data policies and standards, we are collaborating with multiple scientific facilities to develop a unified data policy and metadata standard, with the goal of establishing it as a national standard (GB) in China. Efforts are also underway to promote and deploy this metadata standard across facilities. For data software and tools, we have developed a suite of frameworks including Mamba for data acquisition, DOMAS for data management, and DAISY for data processing. These tools are tightly integrated with the metadata standards to ensure normative and consistent data handling throughout the entire data lifecycle. In data fusion and provision, data agents have been developed to automate data cleaning, fusion, and alignment. For instance, our synchrotron radiation X-ray diffraction data agent can simulate diffraction data from crystal structure files, automatically process experimental data, and further perform intelligent refinement and deep integration of simulated and experimental data. This process enables the generation of aligned, fused datasets that are AI-ready for model training.

Finally, we summarize the key contributions and outcomes of this work. The proposed data ecosystem significantly improves data management and utilization efficiency, providing a sustainable supply of high-quality datasets that support both the AI4S research paradigm and facility-driven scientific innovation.

**Primary author:** HU, Hao (Institute of High Energy Physics)

**Presenter:** HU, Hao (Institute of High Energy Physics)

**Session Classification:** Artificial Intelligence (AI) - III

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 101

Type: **Oral Presentation**

# The Helmholtz Model Zoo: Enabling AI Model Sharing and Inference in the Helmholtz Cloud

*Friday, 20 March 2026 09:22 (22 minutes)*

The Helmholtz Model Zoo (HMZ) is a cloud-based platform that provides remote access to deep learning models within the Helmholtz Association. It enables seamless inference execution via both a web interface and a REST API, lowering the barrier for scientists to integrate state-of-the-art AI models into their research.

Scientists from all 18 Helmholtz centers can contribute their models to HMZ through a streamlined, well-documented submission process on GitLab. This process minimizes effort for model providers while ensuring flexibility for diverse scientific use cases. Based on the information provided about the model, HMZ automatically generates the web interface and API, tests the model, and deploys it. The REST API further allows for easy integration of HMZ models into other computational pipelines.

With the launch of HMZ, researchers can now run AI models directly within the Helmholtz Cloud, ensuring that all data remain within the association and that our data sovereignty is preserved. The platform imposes no strict limits on the number of inferences or the volume of uploaded data, while Helmholtz Virtual Organizations (VOs) enable fine-grained access control for specialized models. External researchers can also access HMZ through Helmholtz VOs upon invitation by a Helmholtz representative, facilitating collaborative research beyond the association's boundaries. Data uploaded for inference is stored within HIFIS dCache InfiniteSpace and remains under the ownership of the uploading user.

HMZ is powered by GPU nodes, hosted as part of the DESY Hamburg HPC cluster. Model inference is managed through the NVIDIA Triton Inference Server, ensuring efficient GPU utilization. The development and maintenance of HMZ are led by the Helmholtz Imaging Support Team at DESY, with support from Helmholtz Federated IT Services (HIFIS) and the Helmholtz AI platform. Hardware and implementation have been supported by funds from the Haicore initiative.

Our presentation will provide an overview of HMZ architecture and its integration into a professional HPC environment. It will also address the scientific foundations of selected models and emphasise the benefits of operating them entirely within the Helmholtz infrastructure.

**Primary author:** Dr FUHRMANN, Patrick (DESY/dCache.org)

**Co-authors:** WERNERS, Hans (Deutsches Elektronen-Synchrotron (DESY)); Dr EREN, Engin (Deutsches Elektronen-Synchrotron (DESY)); Dr HEUSER, Philipp (DESY)

**Presenter:** Dr FUHRMANN, Patrick (DESY/dCache.org)

**Session Classification:** Virtual Research Environment (VRE) - I

**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 102

Type: **Oral Presentation**

## Development and Practice of AI Agents and Framework for Scientific Data Processing

*Wednesday, 18 March 2026 17:06 (22 minutes)*

2025 is widely recognized as the Year of the AI Agent. Large language models have moved beyond conversational interfaces to become callable tools that boost productivity—evident in the rapid adoption of systems like Manus, Claude-Code, and Cursor. AI Agent technologies are also increasingly being applied in scientific research to assist in data analysis and literature exploration, as demonstrated by systems such as SciMaster, FutureHouse, Machine Chemist, and SciToolAgent. The Computing Center of the Institute of High Energy Physics (IHEP) initiated research on scientific AI Agents in 2023 and developed Dr.Sai, an intelligent agent for BESIII (Beijing Spectrometer) physics analysis. Building upon this experience, we present OpenDrSai—a scientific AI agent framework designed to accelerate the development and deployment of AI agents for scientific data processing.

OpenDrSai integrates core capabilities including self-learning and reflection, real-time human-agent interaction, long task management, and multi-agent collaboration. The framework offers modular components for multimodal scientific data perception, knowledge and memory management, scientific tool orchestration, and complex workflow execution. It also features a flexible multi-agent architecture, a scalable backend system, an interactive human-machine interface, and standardized APIs. These features address key challenges in scientific AI development, such as integrating complex tools, managing long-running tasks, and handling domain-specific data and knowledge.

OpenDrSai is already deployed or planned for use in several large-scale scientific experiments, including the China Spallation Neutron Source, Beijing Synchrotron Radiation Facility, Large High Altitude Air Shower Observatory (LHAASO), JUNO Neutrino Experiment, and the Deep-Sea Neutrino Telescope. Some specialized agents—such as DataAgent, RongZai Agent, and BOSS8 Assistant—have been developed to support tasks including neutron diffraction and PDF refinement, as well as data processing for large-scale experimental facilities.

**Primary authors:** Dr XIONG, Dongbo (Institute of High Energy Physics of the Chima Academy of Sciences); QI, Fazhi (Institute of High Energy Physics,CAS); ZHANG (X), Zhengde (Institute of High Energy Physics (IHEP), CAS, China)

**Presenter:** Dr XIONG, Dongbo (Institute of High Energy Physics of the Chima Academy of Sciences)

**Session Classification:** Artificial Intelligence (AI) - II

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 103

Type: **Oral Presentation**

## Dr. Sai assistant for particle physics based on LLMs

*Wednesday, 18 March 2026 16:44 (22 minutes)*

With the widespread adoption of large language models (LLMs) like ChatGPT, AI has emerged as a transformative productivity tool across human industries. LLM-based agents—capable of autonomous task planning, tool utilization, and result explanation—have consequently become a hot point of recent research.

High-energy physics analysis presents a compelling AI4Science scenario. It features established, AI-adaptable processes like data processing and parsing, yet demands specialized domain knowledge and involves intricate, multi-step workflows.

Building on this, we introduce Dr. Sai, a multi-agent system designed for BESIII high-energy physics experiment analysis. By decomposing complex tasks into manageable subtasks, specialized agent teams collaboratively execute core workflows—including algorithm generation, job script construction/submission, and result visualization.

Currently, Dr. Sai is able to automate physics data pre-processing steps. Future development will prioritize intuitive human-computer interaction and adaptive learning to continuously enhance efficiency in complex scientific research tasks.

**Primary author:** ZHANG, bolun (IHEP)

**Presenter:** ZHANG, bolun (IHEP)

**Session Classification:** Artificial Intelligence (AI) - II

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 104

Type: **Oral Presentation**

## How Open Science, Data Science, and Scientific Computing Revitalize and Make Hard-to-reach Population Visible: A Two-decade Research on Taiwan Indigenous Peoples

*Tuesday, 17 March 2026 15:00 (30 minutes)*

Taiwan indigenous peoples (TIPs) are a branch of Austronesians or Polynesians. Persistent lack of data on TIPs made TIPs become isolated, marginalized, underdeveloped, and thus “invisible” hard-to-reach populations in the real world. I will address how scientific computing, data science, and open science are applied to build TIPD (<https://osf.io/e4rvz>) big data based on Taiwan’s household register. Scientific computing methods and technologies, data science, and open science are three crucial dimensions that are utilized to overcome the aforementioned challenges. My talks will highlight theoretical foundation, implementation process, challenges, and ways to overcome research barriers in synthesizing, processing, enriching, managing, and sharing big data. The efforts in building open data not only enable us to access HRP and build insights into various issues they encounter, but also allow us to design effective policy measures to empower and revitalize HRPs. The contributions of my research on TIPs and HRPs are as follows. First, a contribution of moving from “closed” to “open” data access. Second, a contribution of moving research from “the elite” to “the ordinary” people. Third, a contribution of moving from “local” to “global” perspective. Fourth, a contribution of enabling TIPs research from “macro and static” to “micro and dynamic” data (<https://doi.org/10.1007/s43545-025-01049-1>).

Keywords: data science, hard-to-reach population, open science, open data, scientific computing, Taiwan indigenous peoples

**Primary author:** LIN, Ji-Ping (Academia Sinica Taiwan)

**Presenter:** LIN, Ji-Ping (Academia Sinica Taiwan)

**Session Classification:** FAIR, Sovereign & Trusted Data - I

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data

Contribution ID: 105

Type: **Oral Presentation**

## Data Agent-Driven Construction of High-Quality Cross-Facility Diffraction Datasets

*Tuesday, 17 March 2026 14:30 (30 minutes)*

Experimental data generated by scientific facilities such as synchrotron radiation and neutron sources provide a fundamental resource for scientific research. With the deepening integration of massive scientific data and artificial intelligence (AI) technologies, the paradigm of scientific inquiry is undergoing a transformative shift. This evolution places higher demands on the integrated management of data systems in large-scale facilities and the AI-ready datasets processing of simulated and experimental data. In the field of materials science, X-ray diffraction (XRD) and neutron powder diffraction (NPD) each offer unique and complementary advantages in resolving crystal structures. However, effectively integrating these two types of heterogeneous data remains a significant technical challenge. To address this issue, we have developed a cross-facility data agent designed to unify and coordinate diffraction data streams from synchrotron radiation and neutron sources. This agent establishes a closed-loop workflow that encompasses simulated data generation and joint refinement of experimental data. The core mechanism involves systematically minimizing the discrepancy between simulated and experimental data, thereby achieving effective alignment of simulated data with experimental observations. The high-quality cross-facility diffraction datasets constructed through this approach provide a reliable data foundation for training AI models in crystal structure prediction, while also opening up a new technical pathway for accelerating research on structure–property relationships in materials.

**Primary authors:** LI, QINGMENG (Institute of High Energy Physics, Chinese Academy of Sciences); XIONG, Dongbo (Institute of High Energy Physics, Chinese Academy of Sciences); ZHANG, Zhengde (Institute of High Energy Physics, Chinese Academy of Sciences)

**Co-authors:** DU, Rong (China Spallation Neutron Source); WANG, Hao (China Spallation Neutron Source); DONG, Kang (Institute of High Energy Physics, Chinese Academy of Sciences); LI, Zhenbang (Institute of High Energy Physics, Chinese Academy of Sciences)

**Presenter:** LI, QINGMENG (Institute of High Energy Physics, Chinese Academy of Sciences)

**Session Classification:** FAIR, Sovereign & Trusted Data - I

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data

Contribution ID: 106

Type: **Poster Presentation**

## **Preliminary Practice of Proxmox VE and A Simple Management System**

Proxmox VE (PVE) is a light-weighted decentralized virtualization platform based on KVM and LXC. PVE is suitable for small and medium-sized data centers. We deployed a small PVE cluster combined with NVME-based Ceph cluster to evaluate PVE's reliability and robustness, and it has been operating stably for over a year since its launch. PVE And We also developed a simple management system to implement multi-tenant management.

**Primary authors:** LI, Haibo (Institute of High Energy Physics, Academy of Sciences); Ms HOU, Siqi (Institute of High Energy Physics, Chinese Academy of Sciences); BI, Yujiang (IHEP); Mr YAN, Xiaofei (Institute of High Energy Physics, Chinese Academy of Sciences)

**Presenters:** Ms HOU, Siqi (Institute of High Energy Physics, Chinese Academy of Sciences); BI, Yujiang (IHEP)

**Session Classification:** Infrastructure Clouds & Virtualisation

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 107

Type: **Oral Presentation**

# Research on Anomaly Detection Technology for Raw Data Transmission Performance of JUNO Experiment

*Thursday, 19 March 2026 12:20 (20 minutes)*

The JUNO Experiment officially started data acquisition at the end of 2024, producing approximately 8TB of raw data daily, all of which must be transmitted to the offline computing platform in Beijing in real time. Nevertheless, in production, during data transmission, unstable transmission performance may occur, and existing monitoring methods cannot quickly identify the root causes. Therefore, there is an urgent need to develop a system capable of rapidly locating issues affecting raw data transmission performance. This report proposes a network performance anomaly detection method for high-energy physics raw data transmission. Through network session collection, collation, and analysis, an anomaly detection model based on ensemble attention and temporal memory identifies abnormal sessions. Combined with the iterative in-depth analysis method and IHEP's OpenDrSai, an intelligent agent for network anomaly classification is implemented. Experiments show that this method can achieve accurate identification and diagnosis of network performance issues in JUNO's raw data transmission.

**Primary authors:** Mr LI, Cheng (IHEP); ZENG, SHAN (IHEP); QI, Fazhi (Institute of High Energy Physics,CAS)

**Presenter:** ZENG, SHAN (IHEP)

**Session Classification:** Networking, Security & Operations - I (11:00 - 12:40)

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: **108**Type: **Oral Presentation**

## NACO, the "N" Attribute Conformity Checker

*Thursday, 19 March 2026 11:00 (20 minutes)*

NACO is a comprehensive Python-based validation tool designed to ensure compliance with attribute requirements, such as NFDI or AARG-G056. The system provides dual interfaces - a command-line tool (naco) and a web service (naco-web) - for validating OIDC tokens and SAML assertions against configurable attribute specifications.

The tool supports flexible attribute checking through JSON-based configuration files that define both mandatory and optional attribute sets (basic and extended). NACO can validate against multiple OpenID Connect providers simultaneously, with each provider having its own configuration section including credentials, mytokens, and administrative contacts.

Key capabilities include external and community scope validation, and comprehensive email notification systems for administrators. The architecture is built around core components including a spec checker for validation logic, token handling utilities for OIDC operations, and a robust configuration system. This makes NACO suitable for research data infrastructure environments where consistent attribute validation across federated identity providers is critical for access control and compliance.

**Primary author:** HARDT, Marcus (KIT)

**Presenter:** HARDT, Marcus (KIT)

**Session Classification:** Networking, Security & Operations - I (11:00 - 12:40)

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 109

Type: **Oral Presentation**

## Research and Application of Object Storage Systems for AI-based Computing in High Energy Physics

*Friday, 20 March 2026 12:02 (18 minutes)*

With the deepening application of large artificial intelligence (AI) models in high energy physics (HEP) data analysis, both data access patterns and storage architectures are undergoing profound transformation. Traditional file systems face performance and scalability bottlenecks when dealing with unstructured data, high-concurrency access, and cross-site data sharing. Object storage, with its advantages of high scalability, cost efficiency, strong consistency, and native API access, is increasingly becoming the new data infrastructure for AI computing platforms.

This report analyzes typical application scenarios of object storage in AI model training, inference, and data management within HEP experiments, and introduces a newly developed cross-wide-area-network distributed object storage system, JWanFS. The system's applications in training dataset management and cross-center collaboration are examined in detail.

By comparing JWanFS with mainstream object storage systems, the report proposes key optimization strategies for deploying object storage in HEP data centers, including multi-level caching, wide-area data synchronization, and metadata acceleration mechanisms. The study demonstrates that AI- and science-oriented object storage systems like JWanFS can significantly improve data access efficiency and operational flexibility in AI computing platforms, providing a solid foundation for future large-scale scientific computing.

**Primary authors:** LI, Haibo (Institute of High Energy Physics, Academy of Sciences); CHENG, Yaodong (IHEP, CAS); BI, Yujiang (IHEP); Mr WEI, lichang (IHEP)

**Presenter:** LI, Haibo (Institute of High Energy Physics, Academy of Sciences)

**Session Classification:** Networking, Security & Operations - II

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 110

Type: **Oral Presentation**

## Federation Manager: a New Component for the Federated Resource Management in INFN Cloud

*Thursday, 19 March 2026 11:00 (22 minutes)*

As part of the strategic refactoring and modernization of the INFN Cloud orchestration system, the Federation Manager has been developed to enhance the flexibility, scalability, and interoperability of the distributed DataCloud infrastructure. This initiative represents a key step in the long-term evolution of INFN Cloud toward a more modular, service-oriented architecture capable of supporting hybrid and multi-cloud orchestration across heterogeneous resource providers. Within this context, the Federation Manager plays a central role by automating the federation of new providers and managing resource access requests from scientific communities in a unified and secure procedure.

The service architecture is built on a Python backend, leveraging the FastAPI framework to implement RESTful APIs that follow best practices and relevant RFC standards. This ensures high performance, consistency, and ease of integration with other components of the INFN Cloud.

A modern, responsive web interface, developed with React and Next.js, provides users with a clear and intuitive experience for managing provider federations and resource allocations. Designed using Figma, the interface focuses on scalability and maintainability, aligning with the shared design principles adopted across the PaaS ecosystem.

Authentication and authorization mechanisms are based on OAuth2 and OpenID Connect (OIDC), ensuring secure, standards-compliant identity management. Centralized and flexible policy enforcement is achieved through Open Policy Agent (OPA), while Apache Kafka handles asynchronous communication with other PaaS components, guaranteeing reliability and scalability in data exchange.

A dedicated Python-based monitoring component complements the Federation Manager by periodically executing Rally benchmarking tests on the federated cloud providers. This module automates the evaluation of performance indicators such as compute, storage, and network efficiency, collecting metrics that are stored and analyzed to detect anomalies or degradation in service quality. The integration of this testing component provides continuous insight into the health and reliability of the federated infrastructure, supporting proactive maintenance and capacity planning.

By defining reusable API specifications, adopting shared UI frameworks, and integrating modern security and policy technologies, the Federation Manager contributes to a coherent and sustainable technical framework for INFN Cloud. This unified approach simplifies the integration of new services, improves maintainability, and strengthens the overall interoperability of the platform. Ultimately, the Federation Manager supports INFN Cloud's mission to provide a robust, extensible, and standards-aligned infrastructure for data-intensive scientific research in federated and multi-cloud environments.

**Primary author:** SAVARESE, Giovanni (INFN Bari)

**Co-authors:** CASALE, Alessandra (INFN LNGS); COSTANTINI, Alessandro (INFN-CNAF); DONVITO, Giacinto (INFN); GASPARETTO, Jacopo (INFN); GIOMMI, Luca (INFN CNAF); SERRA, Ettore (INFN)

**Presenter:** SAVARESE, Giovanni (INFN Bari)

**Session Classification:** Infrastructure Clouds and Visualisations - I (11:00 - 12:40)

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 111

Type: **Oral Presentation**

# High Performance Plasma Physics Simulation Output Management

*Wednesday, 18 March 2026 14:00 (22 minutes)*

Models of physical systems simulated on HPC clusters often produce large amounts of valuable data that need to be safely managed both during the research projects ongoing activities and afterwards. To help to derive the most benefit for scientific advancement, we use results of the Horizon Europe Project EXA4MIND applying the tools to managing Particle In Cell simulation data from research conducted at the LMU physics department. We consider interoperation of HPC compute and file system with databases and object stores. We evaluate post-processing workflows for physics simulations run on supercomputing systems at LRZ (Garching b.M./DE) in collaboration with LMU Munichs Chair of Plasma and Computational Physics and Experimental Physicists from University of Jena, using the Plasma Simulation Code (Ruhl et al., Ludwig Maximilian University of Munich/DE) to produce TBs of particle movement data of set-ups that investigate the acceleration of particles in laser plasma interaction. We apply our workflow tools to post-processing and visualizing to discuss the parameter studies of simulated models.

Our focus includes improving the performance of data queries and processing steps with the chosen execution methods. We aim to build an ecosystem at the computing center where users can enjoy fast and flexible access to both raw and post-processed data by offering different storage and database systems. This includes facilitating a research data management adhering to the FAIR (findable, accessible, interoperable, reusable) principles.

This research received the support of the EXA4MIND project, funded by a European Union's Horizon Europe Research and Innovation Programme, under Grant Agreement N° 101092944. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

**Primary authors:** HACHINGER, Stephan (Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities); PAUW, Viktoria (Leibniz Rechenzentrum)

**Presenter:** PAUW, Viktoria (Leibniz Rechenzentrum)

**Session Classification:** Physics and Engineering Applications -I

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 112

Type: **Oral Presentation**

## Secure Cloud-Native Infrastructure for AI-Based Tumor Tissue Analysis: A Collaborative Approach

*Thursday, 19 March 2026 11:22 (22 minutes)*

The integration of artificial intelligence (AI) into biomedical research is transforming the analysis of complex datasets such as high-resolution images of tumor tissues. As part of a collaboration between the Italian EOSC and BBMRI-ERIC nodes, INFN and BBMRI-ERIC have launched a joint initiative to define and deploy a secure and scalable infrastructure capable of supporting AI-driven workflows for histopathological image analysis. This effort has been supported by the CERIT-SC team at Masaryk University, which successfully hosts a small-scale, secure, AI-ready computing environment known as SensitiveCloud.

INFN contributed by designing and deploying a cloud-native platform tailored to the specific requirements of medical data handling within an Information Security Management System (ISMS). The ISMS infrastructure is built on a Kubernetes (K8s) cluster composed of virtual machines provisioned via OpenStack. Each VM runs a hardened operating system, and the cluster is orchestrated using RKE2 with Center for Internet Security (CIS) benchmarks enforced, ensuring compliance with best practices for secure configuration.

To facilitate controlled access to applications and data, the platform integrates Keycloak as an external identity and access management system. This setup enables federated authentication and fine-grained authorization policies, supporting multi-institutional collaboration while preserving data privacy and integrity. Integration with the LifeScience Authentication and Authorization Infrastructure (AAI) is currently under evaluation to further enhance interoperability and user trust. This system underpins authentication in the BBMRI-ERIC EOSC Node and is fully compatible with the EOSC AAI, including support for government-backed eID identities.

The infrastructure hosts AI pipelines designed to analyze digitized tumor tissue samples, leveraging deep learning models for feature extraction, classification, and pattern recognition. These workflows are containerized and deployed within the RKE2 cluster, benefiting from the elasticity and isolation provided by Kubernetes. The entire infrastructure was developed as part of the BioMedAI project at Masaryk University and transferred to INFN as Docker containers for the relevant components (mlFlow, JupyterHub with custom images, xOpat viewer). The platform also supports reproducibility and traceability through integrated logging and monitoring tools.

These activities demonstrate the synergy between infrastructure providers and domain experts, paving the way for scalable, secure, and privacy-preserving AI applications in the health and life sciences.

**Primary authors:** SINISI, Francesco (INFN); BRÁZDIL, Tomáš (Faculty of Informatics, Masaryk University); COSTANTINI, Alessandro (INFN); GASPARETTO, Jacopo (INFN); HEJTMÁNEK, Lukáš (Institute of Computer Science, Masaryk University); HOLUB, Petr (BBMRI-ERIC & Institute of Computer Science, Masaryk University); HORÁK, Jiří (BBMRI-ERIC & Faculty of Informatics and Institute of Computer Science, Masaryk University); MARTELLI, Barbara (INFN); MUSIL, Vit (Faculty of Informatics, Masaryk University); SERGI, Giusy (INFN); TUERK, Andreas (BBMRI)

**Presenter:** SINISI, Francesco (INFN)

**Session Classification:** Infrastructure Clouds and Visualisations - I (11:00 - 12:40)

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 113

Type: **Oral Presentation**

## **The WeNMR Case Study: Research Software as a Service for Demonstrating Impact, Securing Funding, and Serving a Global Community**

*Friday, 20 March 2026 09:44 (24 minutes)*

Sustained investment in research software requires moving beyond downloads/citations to demonstrate impact. This presentation uses the WeNMR platform – a VRE for structural biology with over 70,000 users – as a case for the “Research-Software-as-a-Service” (RSaaS) model. We will detail how the RSaaS model provides visible and quantifiable impact. By moving software into a centralized, managed service, we capture crucial metrics: over 12 million computational jobs processed per year, representing over 4,800 CPU-years per year of scientific output. These concrete usage statistics, alongside user growth and support requests, provide an irrefutable narrative of widespread adoption and scientific utility that is otherwise invisible when software is simply distributed. This data has been our primary tool to successfully convey importance to funding agencies and secure resources. We will share strategies for translating research software into services that we can be used into compelling arguments for institutional support, national grants, and international funding. The talk will cover technical aspects needed to create and operate RsaaS and discuss our integrations with EOSC/EGI for HTC and the development of a containerized architecture. This presentation will provide an overview on how to build, and sustain research software services by demonstrating its indispensable role in the modern research ecosystem.

**Primary authors:** BONVIN, Alexandre M.J.J. (Utrecht University); Dr HONORATO, Rodrigo V.

**Presenter:** BONVIN, Alexandre M.J.J. (Utrecht University)

**Session Classification:** Virtual Research Environment (VRE) - I

**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 114

Type: **Oral Presentation**

## Hybrid Cloud Architecture for Astronomical Observation Data: Case Studies on ALMA Data Archive and Analysis

*Thursday, 19 March 2026 11:44 (22 minutes)*

An astronomical observatory requires not only state-of-the-art telescopes but also robust computing infrastructure to archive and analyze the vast amounts of astronomical observation data. Consequently, optimizing the operation of these computing systems is a crucial issue. Adopting public cloud services is expected to reduce the Total Cost of Ownership (TCO) and allow the use of cutting-edge technologies (e.g., the latest GPUs). However, a systematic methodology for designing an optimal hybrid architecture to fully realize these advantages is currently lacking.

The National Astronomical Observatory of Japan and the National Institute of Informatics have been conducting case studies to demonstrate best practices for designing and implementing a hybrid cloud architecture dedicated to storing and analyzing ALMA radio telescope data. We collected and analyzed both system operation data (e.g., storage usage, analysis execution time) and application data (e.g., observed data from ALMA 12m and 7m antennas) to conduct two primary experiments.

First, we estimated the storage cost for petabyte-scale ALMA observation data, which requires high-speed access for hot data (frequently accessed data for analysis) and low-cost, long-term archiving for cold data (infrequently accessed data), based on the actual statistics such as data capacity increase of each year and the relationship between data age and access frequency. Although the cost of public cloud storage is often considered a barrier, our studies demonstrated that a three-tier storage architecture, comprising on-premises storage, cloud object storage, and cloud cold storage, with optimized data life-cycle management, reduces the overall cost.

Second, we developed an optimization method for selecting optimal server instances on public cloud services. Choosing the right instances requires considerable domain knowledge of both the application programs and the available service instances. Our method utilizes machine learning models to estimate the required CPU performance and memory capacity for ALMA data analysis based on the metadata of observations, such as the telescope resolution and data size. Crucially, the method adaptively chooses different server instances for different computing phases; for example, assigning a single-core instance for the calibration phase and switching to a multiple-core instance for the imaging phase. We also developed a procedure to execute the switching of those instances automatically by utilizing the cloud service API. Our experimental results indicate a 60% cost reduction accompanied by only a slight increase in execution time compared to the unoptimized execution, particularly for 7m antenna observation data.

Our case studies demonstrate that a hybrid architecture improves the operational efficiency of computing systems for astronomical observatories. While the cost of public cloud services remains a significant concern, our optimization methods reduce this barrier while preserving the necessary performance. Our presentation will detail the systematic design, methodologies, and quantitative experimental results that validate our findings.

**Primary author:** YOSHIDA, Hiroshi (National Institute of Informatics)

**Co-authors:** AIDA, Kento (National Institute of Informatics); KOSUGI, George (National Astronomical Observatory of Japan); MORITA, Eisuke (National Astronomical Observatory of Japan); NAKAZATO,

Takeshi (National Astronomical Observatory of Japan); SHIZUGAMI, Makoto (National Astronomical Observatory of Japan); MIEL, Renaud (National Astronomical Observatory of Japan); TAGAWA, Hiroaki (National Astronomical Observatory of Japan); SUGIMOTO, Kanako (National Astronomical Observatory of Japan)

**Presenter:** YOSHIDA, Hiroshi (National Institute of Informatics)

**Session Classification:** Infrastructure Clouds and Visualisations - I (11:00 - 12:40)

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 115

Type: **Poster Presentation**

## Design and Implementation of an Intelligent Network Management Platform Based on NetDisco

The Institute of High Energy Physics (IHEP) currently has over 2,000 network devices. Traditional network management methods are time-consuming, labor-intensive, and unable to quickly locate problems in network operations and maintenance. This report introduces an intelligent and visualized network management platform for IHEP based on the open-source network management tool Netdisco. Netdisco uses protocols such as SNMP and LLDP to achieve automatic network device discovery, topology generation, asset statistics, and device inventory. Building upon this, several functional enhancements and extensions have been designed and implemented based on data acquisition, integration, storage, and analysis technologies. These include automated discovery of large-scale subnets and unified visualization of multi-device network topology and link traffic. Furthermore, the platform plans to introduce Large Language Model (LLM) and toolcall technologies, allowing users to query device information, port status, and network topology using natural language, and providing statistical analysis capabilities. In the future, this system will be applied to the daily network operations and maintenance of IHEP, assisting frontline staff in troubleshooting and reporting problems.

**Primary authors:** SUN, jialiag ( the Institute of High Energy Physics, Chinese Academy of Sciences); ZENG, shan ( the Institute of High Energy Physics, Chinese Academy of Sciences)

**Presenter:** SUN, jialiag ( the Institute of High Energy Physics, Chinese Academy of Sciences)

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 116

Type: **Oral Presentation**

## Open Computing Platform: A Self-service Remote Access System for Computing Continuum

*Friday, 20 March 2026 10:08 (22 minutes)*

Advanced computing infrastructures offer unparalleled computing capabilities and effectively support a multitude of computing requirements across diverse fields such as scientific research, big data analysis, artificial intelligence training and inference, and many more. Secure Shell (SSH) is a widely used method for accessing remote computing resources. To aim for efficient and secure access to computing resources, this paper proposed the establishment of an authentication chain consisting of a CSTNET Passport and an SSH lightweight certificate. To enhance security, the model restricts the time window, narrowing it from an unrestricted period to a specific time range designated by the user. The system changes access from anytime and anywhere to on-demand access with various restrictions specified by users, and also transfers controls from administrators to users. The experiments show that, the system is easy to deploy, occupies fewer resources, does not introduce extra hardware costs, and can effectively increase the usability and security of computing resource.

**Primary author:** CAO, Rongqiang (Computer Network Information Center, Chinese Academy of Sciences)

**Co-authors:** Prof. JUE, Wang (Computer Network Information Center, Chinese Academy of Sciences.); Mr WAN, Meng (Computer Network Information Center, Chinese Academy of Sciences.); Prof. WANG, Yangang (Computer Network Information Center, Chinese Academy of Sciences.)

**Presenters:** CAO, Rongqiang (Computer Network Information Center, Chinese Academy of Sciences.); Prof. JUE, Wang (Computer Network Information Center, Chinese Academy of Sciences.)

**Session Classification:** Virtual Research Environment (VRE) - I

**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 117

Type: **Oral Presentation**

# Revisiting Multivariate Time Series Forecasting via Multi-Scale Selective Enhancement and Cross-Level Attention

*Friday, 20 March 2026 10:06 (22 minutes)*

Multivariate time series forecasting often suffers from noise interference, inconsistent dynamics across variables, and limited capacity to capture both short-term fluctuations and long-term trends. This paper proposes a novel framework that addresses these challenges through three coordinated modules. First, a channel-wise modulation mechanism selectively filters anomalous patterns by assigning learnable weights to individual time points. Second, a multi-scale temporal pooling module extracts coarse- to-fine features within each sequence, enabling the model to capture diverse temporal structures. Third, a cross-level attention mechanism bridges low-level signals and high-level abstractions to enhance semantic integration across timescales. Together, these components allow the model to focus adaptively on informative patterns while maintaining computational efficiency. Experiments on seven public datasets demonstrate that the proposed method achieves superior accuracy compared to existing approaches, particularly in scenarios with strong periodicity or irregular variance. The design also offers low memory overhead, making it suitable for deployment in resource-constrained environments.

**Primary author:** Prof. WANG, Jue (Computer Network Information Center, Chinese Academy of Sciences.)

**Co-authors:** CAO, Rongqiang (Computer Network Information Center, Chinese Academy of Sciences); WANG, Yangang (Computer Network Information Center, Chinese Academy of Sciences.)

**Presenters:** Prof. WANG, Jue (Computer Network Information Center, Chinese Academy of Sciences.); CAO, Rongqiang (Computer Network Information Center, Chinese Academy of Sciences)

**Session Classification:** Artificial Intelligence (AI) - IV

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 118

Type: **Oral Presentation**

## Agile Development with Agentic AI

*Wednesday, 18 March 2026 14:00 (22 minutes)*

Commercial software development today happens almost exclusively in agile teams, using either the SAFe or Scrum framework. AI agents can communicate via MCP and take on different roles, including developer, scrum master, product owner and QA roles. The presentation discusses how to set up agile teams of agentic AIs, giving human developers the opportunity to let a group of AIs develop a product from start to finish. The approach seems appropriate also for scientific development, which is traditionally “starved” of developers.

**Primary author:** BERLICH, Ruediger (Gemfony scientific UG (haftungsbeschraenkt))

**Presenter:** BERLICH, Ruediger (Gemfony scientific UG (haftungsbeschraenkt))

**Session Classification:** Artificial Intelligence (AI) - I

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 119

Type: **Oral Presentation**

## A phased phase-out plan for IPv4 at a WLCG Tier-1 site

*Friday, 20 March 2026 10:50 (18 minutes)*

Universities and research institutes have been early adopters of IPv4, which have served scientific research infrastructure well in the past. But now the time has come to let go of the legacy protocol with awkward limits, and phase it out in favour of IPv6.

The World-wide LHC Computing Grid (WLCG) is half-way through the transition from IPv4 to IPv6, with almost all services now being dual-stack with both IPv4 and IPv6. Now the time has come to plan for the rest, where we discard the complexity of dual stack in favor of IPv6-only operations.

The driver for doing this in the Nordic Tier-1 site (NT1) sooner rather than later is that we foresee a significant risk of running out of IPv4 addresses when scaling storage servers horizontally in order to handle the High Luminosity LHC (HL-LHC) data rates. We expect to have a data rate of 10-20 times when HL-LHC comes online in 2030, and the most cost-effective way to serve this is to have a larger number storage servers than today. And in order to prove that we are ready for HL-LHC data taking in 2030, it would be good to finish the bulk of the phase-out of IPv4 by Data Challenge 2027

This move comes with lots of constraints though. Since it is only “almost all” services that understand IPv6 today, we cannot completely shut IPv4 down without considerations on how the legacy systems can access data. There might also be unknown dependencies on IPv4 in access or management of services, that we will only detect in testing or production. Individual scientists might want to access the data outside of the grid, for instance from their own laptop which might not have IPv6 yet. There are even reasons that the physics experiments might want to run legacy software for reproducibility, some of it too old for IPv6 support.

Together this indicates a phased approach, and this talk will concentrate on the planning and current status of this effort, with steps towards the end goal and tentative timing of them.

**Primary author:** WADENSTEIN, Mattias (NeIC)

**Presenter:** WADENSTEIN, Mattias (NeIC)

**Session Classification:** Networking, Security & Operations - II

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 120

Type: **Oral Presentation**

## Cloud-Edge Machine Learning in the Swiss Energy Prosumer Pilot: Enhancing Social Acceptability

*Thursday, 19 March 2026 14:00 (30 minutes)*

The O-CEI project is a forward-thinking initiative aimed at addressing key challenges in the European supply chain by creating an open and cross-sector Cloud Edge DOT (CEI) continuum platform. O-CEI uses and upgrades technological innovations from previous successful projects (i.e. existing cloud technology spectrum, including meta-operating systems, cognitive cloud technologies, and decentralized swarm intelligence). It will provide a comprehensive framework in order to be integrated, deployed and validated in 8 multidimensional pilots in key strategic sectors ranging from electromobility, software-defined vehicles, logistics, smart agriculture and agri-food, smart urban environments and multiple electricity grids management. Uptake of the proposed technology will pave the way for a more sustainable and resilient CEI ecosystem, fostering a smoother transition towards a cleaner energy future.

Energy Transmission Service Operators have expressed that achieving flexibility in the energy market may quickly overburden computing infrastructures. Hence, the concept of the computing continuum as a single manageable entity, breaking the silos of edge and cloud computing has gained momentum. Thus, several cross-disciplinary actions are needed to minimize the related ecological impact (e.g., in renewable energy, social education). Being able to exploit those resources, having them cooperating in the matters of anticipation and prediction in energy flexibility is pivotal in O-CEI and in Europe.

The goal of the Swiss pilot to foster prosumers into the increasingly liberalized energy market pursuing high social acceptability, enabling them to reap new forms of electricity production and consumption. The pilot is deployed at three locations in Switzerland, that engage users and energy flexibility in three types of buildings: (i) urban apartments (Geneva), (ii) mixed residential/industrial (Fribourg) and (iii) ski villages (Valais). This pilot will demonstrate how innovative Smart IoT solutions (already) in place can become part of an integral (O-CEI) platform, increasing such acceptance and user engagement while deploying Smart Energy Services promoted by advanced Cloud-Edge-IoT utilities and technologies.

The three scenarios that roughly coincide with the geographic locations are presented.

Scenario 1: Flexibility applications for Distribution System Operator (DSO) or Local Energy Communities (LEC): this scenario deals with flexibility applications that enable DSOs to ensure Grid stability and LECs to optimize their consumption costs and minimize their dependence on DSOs.

Scenario 2: Optimizing Charging Costs of electric vehicles (EVs): This scenario focuses on charging stations for EVs enabling owners to optimize vehicle charging through smart contracts.

Scenario 3: Tourism scenario: this scenario is related to touristic resorts enabling hotels, restaurants and ski resorts to optimize their energy consumption.

This contribution describes the current and future machine learning required to support these three scenarios.

As these scenarios involve gathering data that is potentially personal, the work is presented that ensures that the data privacy regulations of Switzerland and the EU are followed.

**Primary author:** Dr WHITE, John (HES-SO, Geneva, CH)

**Presenter:** Dr WHITE, John (HES-SO, Geneva, CH)

**Session Classification:** Infrastructural Clouds and Virtualisation - II

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 121

Type: **Oral Presentation**

## The NFDI National Node and EOSC Beyond Pilot Node in the EOSC federation

*Tuesday, 17 March 2026 14:00 (30 minutes)*

The European Open Science Cloud (EOSC) aims to create a federated, interoperable ecosystem for research data and services, enabling seamless access and reuse across disciplines. The components of this infrastructure are the EOSC nodes: distributed, compliant, and interoperable service providers from research infrastructures that make up EOSC as a whole. This work presents the design, deployment, and operational challenges of setting up nodes compliant with the EOSC Federation Handbook.

A key initiative in this context is the EOSC Beyond project, which is currently establishing pilot nodes as technical demonstrators to validate the operational and interoperability requirements for the interoperability in the EOSC federation. These pilot nodes serve as prototypes for the broader ecosystem, providing a reference implementation for future candidate nodes – independent deployments currently being set up by various actors across the EOSC landscape. By addressing real-world challenges in federation, security, and service integration, the pilot nodes offer a scalable blueprint for future service providers in EOSC, ensuring consistency with its evolving technical and governance frameworks.

The German National Research Data Infrastructure (NFDI), in collaboration with DESY (Deutsches Elektronen-Synchrotron), KIT (Karlsruhe Institute of Technology), and Forschungszentrum Jülich (FZJ), plays a central role in both the development of the pilot nodes within EOSC Beyond and the support of existing and future nodes. This collaboration leverages the partners' expertise in large-scale research data management, high-performance computing, and federated infrastructures to ensure robust, production-grade implementations. Key contributions include:

- Compliance with the EOSC Federation Handbook and its implementation,
- Standardized deployment workflows for node operators, reducing barriers to deliver services to a wide audience,
- Interoperability testing between pilot and candidate nodes to validate cross-provider functionalities and portability,
- Alignment with NFDI's domain-specific requirements, ensuring the nodes cater to diverse scientific communities (e.g., photon and neutron sciences, climate research, or particle physics).

Through this work, the project demonstrates how coordinated efforts between pan-European initiatives (EOSC Beyond) German initiatives (NFDI) and research infrastructures (DESY/KIT/FZJ) can accelerate the adoption of EOSC nodes, fostering an interoperable federation for open science. The insights gained from the pilot and candidate nodes will inform EOSC's long-term roadmap, particularly in areas such as automated compliance checking, cross-border data flows, and hybrid cloud-HPC integration – topics of direct relevance to the ISGC community.

In the talk we will present our findings with a focus on common goals for the nodes, technical and organisational obstacles and how issues are overcome. An overview of how we plan to advance the NFDI pilot and national node will conclude the presentation.

**Primary authors:** FUHRMANN, Patrick (DESY/dCache.org); WETZEL, Tim (Deutsches Elektronen-Synchrotron DESY)

**Presenter:** FUHRMANN, Patrick (DESY/dCache.org)

**Session Classification:** FAIR, Sovereign & Trusted Data - I

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data

Contribution ID: 122

Type: **Poster Presentation**

## False Positive Reduction in Intrusion Detection System Based on Transfer Learning

In network intrusion detection systems, alert logs generated by intrusion detection devices contain a large number of false positive alert logs, which seriously impair the accuracy of security incident analysis. Thus, filtering false positive alert logs is of great significance. The essence of false positive alert filtering is a classification task: each alert log is labeled to indicate whether it is a false positive, and supervised learning can accurately identify such false positive alert logs. However, the availability of logs with the aforementioned labels is very limited, resulting in unsatisfactory performance of traditional supervised learning methods. To solve this problem, this paper proposes a false alert filtering method based on transfer learning, which aims to realize intelligent discrimination of alert logs by utilizing a small number of logs with the aforementioned labels and a large number of logs without such labels. In this method, the aforementioned labels are defined as source domain labels, and all logs are assigned a target domain label indicating whether the log possesses a source domain label. Subsequently, a Domain-Adversarial Neural Network (DANN) is introduced, comprising a feature extractor, a label predictor, and a domain classifier. The feature extractor conducts feature extraction on the original logs, the label predictor identifies false positive logs, and the domain classifier determines whether a log has a source domain label based on the log's target domain label and its extracted features. The Gradient Reversal Layer (GRL) ensures that the feature distributions of logs with source domain labels and those without become similar—rendering the domain classifier unable to distinguish between the two—ultimately achieving accurate discrimination of alert logs.

**Primary authors:** [redacted], [redacted] ([redacted]); [redacted], [redacted] ([redacted]); [redacted], [redacted] ([redacted]); [redacted], [redacted] ([redacted])

**Presenter:** [redacted], [redacted] ([redacted])

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 123

Type: **Oral Presentation**

## PHOTONIC - Data access at DESY made easy

*Tuesday, 17 March 2026 11:00 (20 minutes)*

German large scale research facilities like DESY with its PETRA III synchrotron beamlines are strong partners for universities to conduct cutting-edge experiments, which generate vast amounts of data that often exceed the universities' storage and compute capacities for analysis. Transferring data over wide area networks is becoming ever more impractical and expensive with the growing amount of data produced. In the project PHOTONIC, DESY and its university partners are developing a model for data access tailored to scientists' (post-)beamtime needs in order to address the aforementioned challenge. At the same time, we hope to enable faster integration of data into analysis workflows and easier access in general by also offering live data access and visualisation possibilities.

The service that will be deployed and integrated with DESY's storage infrastructure is planned to be capable of understanding and interpreting data formats, offer transcoding functionalities and stream partial data by providing API-based access to photon science data. These functionalities will optimize data transfers by delivering only the needed subsets of data, reducing bandwidth and storage demands while accelerating analysis workflows. They also allow for more focused views on data without the need for transferring all available data.

The project is led by DESY (IT & Photon Science) and carried out with partners from universities over 3 years. Its' four positions to implement the functionalities described above are funded by the German Federal Ministry of Research, Technology and Space (BMFTR) and will take care of the backend and client integration tasks as well as training, dissemination and communication.

In this talk we will present a brief overview of the current data-access methods and their limitations followed by an outline of the project's goals, methods & deployment architecture as well as expected outcomes.

**Primary authors:** Mr BURKE, Devin (Deutsches Elektronen-Synchrotron DESY); FUHRMANN, Patrick (DESY/dCache.org); WETZEL, Tim (Deutsches Elektronen-Synchrotron DESY); Dr PITHAN, Linus (Deutsches Elektronen-Synchrotron DESY); MILLAR, Paul (DESY); SERVAN, Sophie (Deutsches Elektronen-Synchrotron DESY)

**Presenter:** FUHRMANN, Patrick (DESY/dCache.org)

**Session Classification:** FAIR, Sovereign & Trusted Data - I

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data



Contribution ID: 125

Type: **Oral Presentation**

## Modeling a Multitier AI Framework for Secure Processing of Confidential Legal Documents

*Friday, 20 March 2026 11:50 (30 minutes)*

This paper presents a multilayer architectural model designed to support secure, AI-driven processing of confidential legal documents. The platform integrates strong authentication and authorization mechanisms, ensuring controlled access to sensitive information in accordance with privacy and security requirements. A key component is an automated contract-generation module based on predefined legal templates and enhanced by advanced language models such as LegalBERT and GPT. This module is complemented by a risk-assessment system that protects personal data and enforces compliance with applicable legislation.

The architecture includes a continuous feedback environment that enables iterative improvement and optimization of system performance. Its modular and extensible design supports the development of additional intelligent components for specialized tasks, including testing and refining legal document processing with advanced NLP techniques.

Beyond document generation, the platform provides comprehensive tools for legal contract analysis. These include automated identification of key provisions, assessment of contractual favorability, and verification of compliance with mandatory statutory rules. The system also incorporates methods for evaluating the reliability of contractual partners using publicly available data sources. To support effective decision-making, it generates clear visualizations and statistical reports summarizing relevant contractual insights and risk indicators.

Overall, the proposed multilayer architecture offers a secure, extensible, and data-driven foundation for AI-assisted legal document processing, enhancing efficiency, accuracy, and regulatory conformity in legal and compliance workflows.

**Primary author:** NGUYEN, Giang (Slovak University of Technology)

**Presenter:** NGUYEN, Giang (Slovak University of Technology)

**Session Classification:** Artificial Intelligence (AU) - V

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 126

Type: **Oral Presentation**

## **Building Trust in Aerospace Infrastructures: Resilient Operations and Provenance-Driven Data Sovereignty**

*Tuesday, 17 March 2026 11:20 (20 minutes)*

The proliferation of distributed space systems, encompassing large-scale satellite constellations and federated ground segments, is generating unparalleled opportunities for global open science, particularly in the domains of earth observation, intelligence, security, and defense. However, this federated model poses significant challenges to operational integrity and data trustworthiness. The increasing reliance on AI-driven discovery underscores the necessity for not only the accuracy of underlying data, but also the transparency of its origin and processing. Addressing this issue necessitates a holistic approach that unifies the operation of secure infrastructure with the principles of FAIR, sovereign, and trusted data.

This contribution presents a pioneering architectural framework derived from ongoing research at the German Aerospace Center (DLR), encompassing projects for resilient and trustworthy networking of space systems) and for trustworthy, responsive mission control). The proposed approach establishes end-to-end trust in a federated e-infrastructure by addressing the system and data layers in a coordinated manner.

The organization's areas of expertise include infrastructure integrity, network security, and operational resilience. The development of resilient networking protocols and operational concepts for distributed ground segments is underway. These protocols and concepts are intended to ensure service continuity during component failures or targeted attacks. This encompasses the implementation of robust federated identity management and threat models (e.g., STRIDE), which are adapted to address specific threats pertinent to space, including signal spoofing and ground station compromises. By ensuring the security of the communication and computational infrastructure, a trusted foundation is established for all subsequent data operations.

The foundation's credibility is leveraged to ensure data governance and cultivate a reliable data ecosystem. Standardized, machine-actionable data provenance (W3C PROV) is integrated throughout the entire data lifecycle, from sensor acquisition to final processing. Automated provenance capture is a process that ensures all data products are traceable, verifiable, and transparent. This fulfills a critical component of the FAIR data principles (findable, accessible, interoperable, and reusable). This verifiable data integrity is imperative for the development of trustworthy AI models and the facilitation of secure, sovereign research collaboration across international borders.

The findings of our research indicate that the attainment of data sovereignty and secure collaboration necessitates more than mere data management; it demands a networking and operations infrastructure that is both resilient and reliable, and which has been proven to be secure.

**Primary author:** Mr SCHREIBER, Andreas (German Aerospace Center (DLR))

**Presenter:** Mr SCHREIBER, Andreas (German Aerospace Center (DLR))

**Session Classification:** FAIR, Sovereign & Trusted Data - I

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data

Contribution ID: 127

Type: **Poster Presentation**

## Scenario reconstructions of the 1920 and 0403 Hualien tsunami events and tsunami hazard assessment of Hualien area

On 3 April 2024, a moment magnitude ( $M_w$ ) 7.2 earthquake struck Hualien, Taiwan, with a focal depth of 22.5 km. The event generated a tsunami that reached 89 cm at the Wushi tide gauge and produced a run-up height of approximately 2.5 m along the coastal area of Yanchao Village in Shoufeng Township. Fortunately, the tsunami occurred during low tide, resulting in no casualties. Nevertheless, this event demonstrated that the coastal geomorphology from Hualien to Yilan is conducive to tsunami amplitude amplification. Therefore, a larger tsunami source could pose significant hazards to this region. The earthquake occurred near the boundary between the Philippine Sea Plate and the Eurasian Plate, where the fault and trench systems are complex. Although the mechanism of this event remains under discussion, it is widely believed to be associated with rupture along the Coastal Range or Central Range fault systems. Due to its proximity to the Ryukyu Trench, the  $M_w$  8 earthquake in the Ryukyu subduction zone in 1920 has also drawn renewed attention.

This study uses the 3 April 2024 Hualien earthquake tsunami as a starting point. By integrating extensive coastal tide gauge observations and, for the first time, seafloor cable seismic data from the earthquake's vicinity, tsunami simulations and analyses are conducted using the COMCOT model to infer the tsunami source of the 0403 Hualien event. Furthermore, the 1920 Ryukyu subduction zone earthquake tsunami is reconstructed and analyzed. Given that subduction-zone earthquakes have a higher likelihood of reaching magnitudes of  $M_w$  8 or greater, additional scenario analyses are performed to assess the potential impacts of large-scale subduction-type tsunamis on the Hualien–Yilan region and the entirety of Taiwan.

A notable feature of the 0403 Hualien earthquake is that, despite its epicenter being located on land, it still generated a tsunami along the eastern coast—an uncommon characteristic among tsunami events. To investigate this type of nearshore tsunami, this study conducts a qualitative analysis of the relationship between epicentral distance and coastline proximity. Threshold criteria for issuing tsunami alerts for such nearshore earthquakes are further identified, along with evaluations of the resulting tsunami amplitudes and their potential impacts.

**Primary author:** , 

**Presenter:** , 

**Track Classification:** Track 3: Earth & Environmental Sciences & Biodiversity Applications

Contribution ID: 128

Type: **Oral Presentation**

## (Remote) Automating Distributed Cloud Infrastructures for AI-Driven Earth Observation Services

Thursday, 19 March 2026 15:00 (30 minutes)

The **European Centre for Medium-Range Weather Forecasts** (ECMWF) operates as both a world-leading provider of global numerical weather predictions and a major research organization advancing Earth system science.

Within its mission to deliver reliable data, forecasts, and insights for the benefit of society, ECMWF is actively involved in several **European initiatives** such as **EO4EU** [1], **MediTwin** [2], **CLIMRES** [3], and **BUILDSPACE** [4]. These projects aim to create advanced services and applications that simplify access to **Earth Observation** (EO) data and leverage **AI/ML techniques** to deliver high-level analytics, decision-support tools, and user-oriented services across multiple domains.

To support these projects, we have developed a **fully automated, cloud-agnostic infrastructure** capable of managing multiple **Kubernetes clusters** deployed across **geographically distributed cloud environments**. This setup ensures that the deployment of services and applications is **secure, reproducible** and **scalable**. Automation tools such as **Git**, **Terraform**, **Ansible** and **Fleet** are integrated into a **GitOps-driven workflow**, enabling consistent cluster management and streamlined application delivery.

The infrastructure also supports **AI/ML-oriented workloads**, including the automated deployment of **JupyterHub environments with GPU-enabled nodes** managed through the **NVIDIA GPU Operator**, allowing for efficient resource partitioning and dynamic scaling of GPU resources.

Building on this foundation, we are now exploring the integration of **Cluster API**, **Crossplane**, and **Sveltos** to enable **declarative cluster provisioning** and **dynamic add-on management**, with the goal of making the platform even more automated, adaptive, and interoperable through **GitOps-based workflows**.

The oral contribution will provide implementation details of this approach and discuss how these technologies are being leveraged to support AI-driven scientific workflows across distributed cloud infrastructures.

References:

- [1] <https://www.eo4eu.eu/>
- [2] <https://meditwin-project.eu/>
- [3] <https://climres.eu/>
- [4] <https://buildspaceproject.eu/>

**Primary authors:** ANTONACCI, Marica (ECMWF); ALBUGHDADI, Mohanad (ECMWF); BAOUSIS, Vasileios (ECMWF); FORNARI, Federico (ECMWF); KAPROL, Tolga (ECMWF); PISA, Claudio

**Presenter:** ANTONACCI, Marica (ECMWF)

**Session Classification:** Infrastructural Clouds and Virtualisation - II

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 129

Type: **Oral Presentation**

## AI-Assisted DevSecOps for Secure Software Delivery in Earth Observation Projects

*Wednesday, 18 March 2026 16:00 (22 minutes)*

Modern Earth Observation (EO) platforms integrate diverse distributed components and scientific workflows across heterogeneous cloud environments. Ensuring software security, maintainability and rapid delivery within such complex systems represents a major operational challenge. To address this, we developed an AI-assisted DevSecOps framework that augments continuous integration and deployment (CI/CD) pipelines with Large Language Model (LLM) capabilities for automated vulnerability detection, remediation and testing.

The framework extends a GitLab–Dagger–FluxCD toolchain with an agentic AI layer composed of three coordinated components: (i) container-level vulnerability analysis using Trivy [1] combined with LLM-generated hardening suggestions; (ii) source-level security and quality inspection using SonarQube [2] enriched with LLM-based corrective patches; and (iii) automated enhancement of unit tests through LLM-guided generation and refinement. All inference is performed on-premises on an NVIDIA GPU-accelerated cloud instance, where a local deployment of LM Studio [3] exposes an OpenAI-compatible interface secured via NGINX, ensuring full confidentiality of code and security findings.

This approach enables autonomous remediation cycles: vulnerabilities detected during CI trigger the generation of secure Dockerfile updates, code fixes or new test cases, which are validated and reintegrated into the workflow. The solution has demonstrated a substantial reduction in remediation time, improved test coverage and increased consistency of security practices across distributed platform components. The system is deployed using Dagger [4] for reproducible pipeline logic and FluxCD [5] for GitOps-based cluster reconciliation.

Future developments include extending support to additional programming languages, integrating unified security insights across all platform services, optimising on-prem GPU inference performance and exploring policy-driven hardening and runtime anomaly detection techniques.

This work demonstrates how AI-enhanced DevSecOps can strengthen the reliability and security of scientific platforms while maintaining full data sovereignty, aligning with ISGC's focus on AI-enabled scientific workflows and advanced operational practices.

### References:

- [1] <https://aquasecurity.github.io/trivy/>
- [2] <https://www.sonarsource.com/products/sonarqube/>
- [3] <https://lmstudio.ai/>
- [4] <https://dagger.io/>
- [5] <https://fluxcd.io/>

**Primary authors:** PISA, Claudio; Dr FORNARI, Federico (ECMWF); ANTONACCI, Marica (ECMWF (on leave from INFN)); ALBUGHDADI, Mohanad (ECMWF); KAPROL, Tolga (ECMWF); BAOUSIS, Vasileios (ECMWF)

**Presenter:** Dr FORNARI, Federico (ECMWF)

**Session Classification:** Artificial Intelligence (AI) - II

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 130

Type: **Oral Presentation**

## Quantum Reservoir Learning on Pasqal neutral-atom Quantum Computer

*Wednesday, 18 March 2026 16:44 (22 minutes)*

The rapid progress in quantum technologies and the advent of Noisy Intermediate-Scale Quantum (NISQ) devices have inaugurated a new computational paradigm. Concurrently, Artificial Intelligence continues to drive transformative advances across research and industry. Recent studies have demonstrated that the integration of these two paradigms can yield mutual benefits, fostering the exploration of hybrid quantum-classical schemes to mitigate the inherent limitations of current quantum hardware. Among the most extensively investigated approaches are Variational Quantum Algorithms (VQAs), which exploit the complementary capabilities of quantum and classical processors through variational optimization. Despite their early success, VQAs exhibit significant scalability challenges, particularly due to convergence to local minima and the barren plateau phenomenon, which severely impairs gradient-based training.

A promising alternative is offered by Quantum Reservoir Computing (QRC), which circumvents the need for costly classical optimization by leveraging the quantum reservoir dynamics coupled with a linear classical layer. This design intrinsically avoids barren plateaus and provides an efficient framework for machine learning applications on NISQ devices. In this work, we present the implementation of a QRC pipeline on the Pasqal neutral-atom quantum platform, including both classical emulation and real-hardware execution. Developed within the Pulser-Pasqal environment, our approach encodes classical inputs into global or local detuning waveforms and extracts embeddings from quantum measurements for a supervised learning task. The obtained QRC embeddings demonstrated a robust improvement over classical baselines, achieving higher test accuracies across different data encoding strategies and using a small-scale neutral-atom setup as the quantum reservoir.

We validate the effectiveness of our implementation through systematic comparisons between different encoding schemes, assessing their impact on both model performance and computational efficiency. This analysis provides crucial insights into the optimal strategies for quantum feature extraction on current hardware. Finally, we discuss the transition from emulated to physical QPU execution, performed through Cineca's HPC infrastructure, highlighting the impact of hardware noise and shot statistics. Our results outline future directions for enhancing QRC performance on real quantum processors and demonstrate the viability of quantum machine learning on near-term quantum devices.

**Primary authors:** BONACORSI, Daniele (University of Bologna); ANGELOZZI, Simone; GASPERINI, Simone (University of Bologna & INFN)

**Presenter:** ANGELOZZI, Simone

**Session Classification:** Physics and Engineering Applications - II

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 131

Type: **Oral Presentation**

## **Integrating False Positive Reduction and LLM-Based Multi-Agent System into IHEP's SOC to Enhance Network Security Operation Intelligence**

*Friday, 20 March 2026 11:44 (18 minutes)*

Network security operations at the Institute of High Energy Physics (IHEP) face severe challenges, including massive data volumes, high alarm complexity, and low manual processing efficiency. While the current Security Operations Center (SOC) system at IHEP has improved cybersecurity operational efficiency to a certain extent through big data platforms and automated workflows, its intelligence level remains insufficient and requires further enhancement. This work proposes integrating False Positive Reduction (FPR) and Large Language Model (LLM)-based multi-agent technology to upgrade IHEP's existing SOC, endowing it with autonomous decision-making, collaborative reasoning, closed-loop task execution, and accurate false positive filtering capabilities. For FPR, a transfer learning-based false alert filtering method is developed to achieve intelligent discrimination of alert logs using limited labeled samples and a large number of unlabeled samples. For the LLM-based multi-agent system, a collaborative mechanism is designed, integrating key technologies such as Retrieval-Augmented Generation (RAG), Text-to-SQL, Text-to-Command, and Chain-of-Thought (CoT) reasoning. This research provides a scalable and accurate technical pathway for the integration and application of LLMs, multi-agent systems, and transfer learning in the field of network security operations.

**Primary authors:** WANG, Jiarong (Institute of High Energy Physics); Mr LONG, Futao (IHEP); Mr ZHOU, Jingkai (IHEP); YAN, Tian (IHEP); QI, Fazhi (Institute of High Energy Physics, CAS)

**Presenter:** WANG, Jiarong (Institute of High Energy Physics)

**Session Classification:** Networking, Security & Operations - II

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 132

Type: **Oral Presentation**

## EGI's Path Towards Democratising AI for Science

The EGI Federation has been a cornerstone of European and global research for over 15 years, providing a federated e-infrastructure that empowers 150,000+ researchers across all scientific disciplines. The recently approved “EGI Federation Strategy 2026–2030” sets out an ambitious plan for the next 5 years to ensure that EGI remains an accelerator for science.

AI became a broadly applied technique in science over the past few years. Recent EGI flagship projects, such as iMagine (Imaging data and services for aquatic science) and InterTwin (Co-designing and prototyping an interdisciplinary Digital Twin Engine) demonstrated the use of AI frameworks, models and data across the EGI e-infrastructure. iMagine used EGI's federated cloud and EGI members' AI4OS technology to deploy an AI platform for aquatic imaging, and to support 17 scientific teams in AI-powered, large-scale image analysis relating to marine and sweet water ecosystems. InterTwin leveraged HPC services from EGI to prototype an open-source Digital Twin Engine for interdisciplinary science and to support the setup of AI-driven digital twins in physics and environmental sciences.

Current EGI flagship initiatives further strengthen the AI trajectory: The EOSC Data Commons project establishes a trusted Research Commons with AI tools that provide seamless access and integration of research outputs, applications and services. RI-SCALE develops and deploys Data Exploitation Platforms that enable peer-partnerships between Research Infrastructure data holdings and HPC/cloud providers to jointly offer scalable, AI-driven analysis and processing environments for scientists.

The new EGI Federation Strategy aims to strengthen EGI's value proposition in AI, and to make a bold contribution to the European and global AI landscapes. The new EGI strategy is based on three integrated pillars: (1) Expand EGI's sovereign Compute-Data Federation with additional compute resources (for AI GPUs and HPC sites), and with scientific datasets ready for exploitation; (2) eliminate and work around policy and funding barriers that limit cross-border access and delivery of compute, data and application services; (3) Strengthen EGI as an R&I ecosystem where innovations persist beyond project lifecycles.

RAISE (Resource for Artificial Intelligence Science in Europe) is Europe's recent initiative to strengthen the continent's position on science and AI. RAISE aims to establish and run a pan-European virtual institute that pools compute, data and talent to operate on two dimensions: 1) “Science for AI” (advancing safe frontier models) and 2) “AI in Science” (applying AI to scientific disciplines). The new EGI strategy, and a recently acquired EC project will help EGI funnel its AI results into the RAISE programme.

This presentation will provide an overview of EGI's new 5-year strategy, and offer a deep dive into its AI-specific implementation, detailing the contributing and related European and national projects and initiatives.

**Primary authors:** SERGIO, Andreozzi (EGI Foundation); SIPOS, Gergely (EGI); RORRO, Marco (EGI Foundation); FERRARI, Tiziana (EGI Foundation)

**Presenter:** SERGIO, Andreozzi (EGI Foundation)

**Session Classification:** Keynote - II

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 133

Type: **Oral Presentation**

## The Impact of Language Models on Academic Evaluation: Rethinking Exam Design and Integrity

*Friday, 20 March 2026 09:44 (22 minutes)*

The rapid spread of large language models (LLM) in higher education has intensified discussions about their promise as instructional support tools and their risks as enablers of academic misconduct. Depending on how they are used, LLMs can assist instructors in developing more efficient learning and evaluation materials, as well as students to prepare for a test, or they can undermine assessment integrity when students or even educators rely on them uncritically.

This work investigates that tension through a quantitative study of two university courses - spanning software/computing for particle physics at a Master in Physics and applied machine learning at a Master in Bioinformatics, both at University of Bologna (Italy) - taught by the author over multiple years. A large archive of instructor-generated questions, routinely used to produce randomised multiple-choice exams, is compared against simulated exam sessions created by collecting answers produced by different LLMs to the same items. Contrasting human and synthetic performance across several academic years reveals distinctive statistical trends, offering insight into how closely LLMs mirror student behaviour, where they fail, and how their presence should influence future assessment design. The analysis further explores whether it is possible to construct LLM-resistant examinations, and how the models themselves may help shape more robust, learning-oriented evaluation strategies. The study ultimately underscores how generative AI is reshaping the landscape of academic responsibility.

As the name ChatGPT drifts phonetically between TeachGPT and CheatGPT, it neatly captures the tension explored in this work: the pedagogical promise and the integrity risks of language models are, quite literally, only a syllable apart.

**Primary author:** BONACORSI, Daniele (University of Bologna)

**Presenter:** BONACORSI, Daniele (University of Bologna)

**Session Classification:** Artificial Intelligence (AI) - IV

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 134

Type: **Poster Presentation**

## Quantum and Digital Annealing for the Subset Sum Problem

*Tuesday, 17 March 2026 16:50 (25 minutes)*

The Subset Sum Problem (SSP) asks whether there exists a subset of a given set  $S$  of integers such that the sum of the elements in the subset corresponds to a predefined target  $t$ . The SSP is one of Karp's NP-complete problems and has several real-life applications. In previous works, we have shown that the SSP can be easily mapped onto a Quadratic Unconstrained Binary Optimization (aka QUBO) formulation, opening the doors for experimental solutions based on alternative computing approaches. We present our computational investigations where hard instances of the SSP are solved by the D-Wave quantum annealer, as well as by more recent quantum-inspired digital technologies. We present a study on the encoding limitations inherent in both quantum and digital annealers when applied to the SSP, with the final aim of understanding how approximated integer input values can influence the optimality and the feasibility of the final solutions returned by the annealers.

**Primary authors:** MUCHERINO, Antonio (IRISA, University of Rennes); Prof. GOLÇALVES, Douglas S. (Federal University of Santa Catarina); Mrs WU, Fang-Tzu (Academia Sinica); Prof. LIN, Jung-Hsin (Academia Sinica)

**Presenter:** MUCHERINO, Antonio (IRISA, University of Rennes)

**Session Classification:** Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

**Track Classification:** Track 9: Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

Contribution ID: 135

Type: **Oral Presentation**

## INFN EPIC Cloud: Building ISO/IEC 27001-Certified Secure Cloud Regions for Trustworthy AI and Open Science

*Thursday, 19 March 2026 11:20 (20 minutes)*

The rapid growth of data-intensive life science research requests infrastructures and services that guarantee security, compliance, and interoperability across federated environments. EPIC Cloud (Enhanced Privacy and Compliance Cloud) represents the highly secure region of INFN DataCloud, representing the backbone of the Italian EOSC national node. Designed to meet stringent privacy and data protection requirements while ensuring FAIRness of scientific data, EPIC Cloud Information Security Management System is certified under ISO/IEC 27001, 27017, and 27018, ensuring a process-based approach to information security, cloud service governance, and personal data protection.

This contribution describes the organizational, architectural and operational principles underpinning the INFN EPIC Cloud, highlighting how ISO-certified information security processes enable trustworthy infrastructures to manage sensitive biomedical data and AI-driven workflows. EPIC Cloud supports critical use cases such as the Italian Health Ministry founded Health Big Data project, addressing the creation of a secure data lake for medical research in Italy, several use cases coming from NRRP-funded projects such as ICSC-Spoke8 and DARE (Digital Lifelong Prevention), the BBMRI-ERIC use case aimed to hosting AI pipelines implemented to analyse digitized tumour tissue samples and exploring the federated authentication, and the BOSCO computational genomics platform, powering large-scale analysis in compliance with GDPR and FAIR principles. By embedding security and compliance into the infrastructure lifecycle, EPIC Cloud advances data sovereignty, fosters secure research collaboration, and aligns with EOSC's vision for global open science.

A distinctive feature of EPIC Cloud is its advanced process-oriented governance model, inspired by ISO/IEC 27022:2021 guidelines for information security process management and strategically aligned with Porter's Value Chain framework. This approach goes beyond compliance, embedding security, and privacy as integral components of INFN's organizational and operational ecosystem.

Security and compliance are not treated as isolated functions but as value-generating activities embedded throughout the chain:

- **Primary Activities:** Data ingestion, AI-driven analytics, federated authentication, and secure data sharing are reinforced by ISO-driven controls, ensuring trust and reproducibility.
- **Support Activities:** Infrastructure management, personnel training, and regulatory alignment provide the foundation for operational resilience and scalability.

We will describe how this integration allows INFN to identify critical dependencies, optimize resource allocation (skilled personnel, operational time, and financial resources), and enhance the overall value delivered to stakeholders in EOSC, in life sciences and open science communities.

Moreover, we'll present the ongoing evolution of INFN EPIC Cloud towards a multiregion cloud that today already includes three INFN sites located in Bologna, Bari and Catania and how EPIC process-based governance model delivers tangible benefits: improved operational resilience, enhanced transparency, and readiness for multi-region scalability.

By embedding security into the value chain, EPIC Cloud establishes a replicable blueprint for scientific infrastructures committed to openness without compromising trust.

**Primary authors:** MARTELLI, Barbara (INFN - CNAF); BRUNENGO, Alessandro (INFN Genova); BELLUOMO, Patrizia (INFN Catania); BETTONI, Diego (INFN Governing Council ); BOVINA, Stefano (INFN CNAF ); BOVO, Eleonora (INFN Central Administration ); CARBONE, Luca (INFN Milano Bicocca ); CARLETTI, Renato (INFN Central Administration ); CARLINO, Gianpaolo (INFN Napoli); CESINI, Daniele (INFN-CNAF); CHIERICI, Andrea (INFN CNAF ); COSTANTINI, Alessandro (INFN-CNAF); D'ANTONIO, Domenico; D'ISIDORO, Antonella (INFN Central Administration ); DELL'AGNELLO, Luca (INFN CNAF ); DONVITO, Giacinto (INFN); ENRICO, Pasqualucci (INFN Roma1); FALABELLA, Antonio (INFN); FOGGETTI, Nadina (INFN Bari and CNR ); GASPARETTO, Jacopo (INFN); GIUGLIANO, Carmen (INFN CNAF ); GRANDI, Claudio (INFN Bologna); IENTILE, Valentina (INFN Catania); LANZI, Leandro (INFN Firenze); LONGO, Stefano (INFN CNAF ); MANZARI, Vito (INFN Governing Council); MAURO, Patano (INFN Bari); MICHELOTTO, Diego (INFN-CNAF); MONFORTE, Salvatore (INFN Catania); MORETTA, Valeria (INFN Central Administration ); PELLEGRINO, Carmelo (INFN-CNAF); ROTONDO, Riccardo (INFN Catania); SCARPONI, Luigi (INFN CNAF); SEQUI, Attilio Gaetano (INFN Central Administration ); SERGI, Giusy (INFN CNAF); SINISI, Francesco (INFN-CNAF); SPINOSO, Vincenzo (INFN Bari); TELESCA, Giuseppe (INFN Central Administration ); TRICOMI, Alessia Rita (INFN Governing Council); VALSECCHI, Veronica (INFN Central Administration ); VINCENZO, Ciaschini (INFN CNAF ); VISTOLI, Cristina (INFN CNAF ); ZANI, Stefano (INFN CNAF ); ZOTTI, Stefano (INFN CNAF)

**Presenter:** MARTELLI, Barbara (INFN - CNAF)

**Session Classification:** Networking, Security & Operations - I (11:00 - 12:40)

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 136

Type: **Poster Presentation**

## LightPacker: An Adaptive Tool for Data Compression and Serialization in HEPS

The experimental data generated by the High Energy Photon Source (HEPS) is featured with massive scale and high diversity, which imposes severe challenges on the real-time efficiency and long-term storage of data processing pipelines. As efficient data serialization and compression are critical to data transmission and storage, traditional fixed strategies fail to adapt to the diverse characteristics of HEPS data and application scenarios. To address these critical issues, this paper designs and implements an adaptive serialization and compression decision-making system, namely LightPacker, which automatically selects the optimal combination of serialization and compression algorithms for specific data through five core modules: a Feature Analysis Module, an Offline Analysis Module, an Online Evaluation Module, a Decision Execution Module, as well as a supporting Algorithm Library and Historical Information Repository. Specifically, the Algorithm Library integrates various serialization and compression methods, providing a unified calling interface and standardizing the integration rules for new algorithms. The Feature Analysis Module extracts key features of real-time data (e.g., metadata, structure, distribution, and sparsity) and calculates the matching degree between real-time data and historical data. The Offline Analysis Module conducts comprehensive benchmark tests of algorithm combinations on accumulated historical data, storing the optimal solutions corresponding to different feature patterns in the Historical Information Repository to support rapid online decision-making. For real-time data, the Online Evaluation Module first retrieves the optimal algorithm combination from the Historical Information Repository based on the feature matching degree; in the absence of matching records, it performs rapid benchmark tests on candidate algorithms to generate preliminary screening results. Finally, the Decision Execution Module invokes the selected algorithm combination to complete data processing. By adopting strategies such as sampling, the system effectively controls additional performance overhead while ensuring the accuracy of decision-making. Experimental results demonstrate that in typical HEPS scenarios, LightPacker significantly outperforms traditional fixed strategies in both data transmission and storage tasks. This system can effectively reduce network bandwidth occupancy, improve preprocessing efficiency, and minimize storage resource requirements, thereby providing an efficient and adaptive solution for data processing pipelines in HEPS and other similar large-scale scientific facilities.

Keywords

data serialization; data compression; algorithm selection; performance testing; intelligent optimization; HDF5

**Primary author:** LIU, Dian (XXXXXXXX)

**Co-authors:** HU , Yu (IHEP); FU , Shiyuan (IHEP); SUN , Haokai; LIU , Jianli (IHEP); WANG, Lei

**Presenter:** LIU, Dian (XXXXXXXX)

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 137

Type: **Oral Presentation**

## Confidential Computing for Sensitive Data Analysis: The EOSC-SIESTA Approach

*Thursday, 19 March 2026 12:06 (22 minutes)*

The FAIR principles provide a foundational framework for ensuring that scientific data is accessible and reusable, and their implementation is a central objective of the European Open Science Cloud (EOSC). However, enabling access to sensitive or confidential data while simultaneously preserving privacy, confidentiality, and usability for researchers remains an open challenge. Existing approaches—such as safe rooms, safe pods, and data safe havens—often hinder the development of reproducible research and can appear counter-intuitive in the context of open science and FAIR-compliant data practices.

The SIESTA project addresses this challenge by developing trusted, cloud-based environments designed for the secure management and sharing of sensitive data. These environments are created using reproducible methodologies and are complemented by a suite of services and tools that simplify the secure exchange of sensitive data within the EOSC, leveraging state-of-the-art anonymization techniques. The overarching goal is to enhance the EOSC Exchange services by delivering cloud-based trusted environments capable of supporting the analysis of sensitive data while demonstrating that the FAIR principles can be effectively upheld in such contexts.

At the core of EOSC-SIESTA lies a distributed, cloud-based computing platform built on Trusted Execution Environments (TEEs)—hardware-backed secure enclaves that isolate sensitive code and data from the surrounding operating environment (e.g., AMD SEV-SNP, Intel TDX, ARM TrustZone). The platform supports both Confidential Computing, where entire virtual machines operate as secure enclaves, and Confidential Containers, a Kubernetes-based secure computing solution in which containers serve as enclaves through technologies such as Kata Containers.

In addition, the DevSecOps toolchain will include a comprehensive set of security capabilities, including continuous vulnerability scanning and tracking to identify and monitor risks throughout the software lifecycle; automated misconfiguration detection and prevention to ensure infrastructure and application settings adhere to security best practices; cryptographic verification and signing of build artifacts to guarantee their integrity and authenticity; and robust access and secret management solutions to securely handle credentials, tokens, and other sensitive information across development, deployment, and operational environments.

**Primary author:** TRAN, Viet (Institute of Informatics, Slovak Academy of Sciences)

**Presenter:** TRAN, Viet (Institute of Informatics, Slovak Academy of Sciences)

**Session Classification:** Infrastructure Clouds and Visualisations - I (11:00 - 12:40)

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 138

Type: **Oral Presentation**

## **Towards Generative Archival Intelligence: Piloting Agentic AI Workflows for Generative Feature Discovery and Dynamic Knowledge Graph Construction from Historic Document Collections**

*Friday, 20 March 2026 09:22 (22 minutes)*

This study investigates how LLMs can act as generative reasoning engines to interpret and re-structure complex archival document collections. Building upon earlier work with the President's Personal File (PPF 9: Gifts) from the Franklin D. Roosevelt Presidential Library, the research explores how an LLM can infer relationships, sequences, and contextual features from textual and descriptive inputs derived from scanned historical pages. The project seeks to understand whether such models can perform higher-order interpretation, recognizing the tacit organization, document groupings, and semantic cues embedded in archival materials. These archival materials are composed of thousands of sequentially arranged pages, often grouped into file folders, and contain a mixture of printed and handwritten textual elements that signal relationships not captured in existing cataloging systems.

The experimental framework employs a multi-stage agentic workflow in which models analyze scanned documents to infer boundaries, sequences, and associations among individual pages, and the LLM serves as both interpreter and orchestrator. The AI workflow process is guided by a baseline knowledge graph that encodes known entities, relationships, and attributes provided by the archival team. With this knowledge, the agentic AI workflow process autonomously analyzes the extracted text and descriptive metadata to hypothesize how documents are bound or sequenced within each folder. As it reasons through patterns in dates, names, annotations, and file markings, the model proposes new candidate features adding to the novel entities, relationships, or attributes that enrich the evolving graph structure. These features are iteratively proposed, validated, and refined through human review, forming a continuous cycle of automated discovery and curatorial oversight.

Over multiple cycles, the knowledge graph becomes a living, self-updating representation of the archival corpus, capturing both its explicit and inferred organizational logic. By analyzing how the system interprets sequentiality, detects tacit groupings, and generates candidate features for approval, the research seeks to evaluate generative AI's capacity for understanding archival logic beyond explicit text recognition. A prototype conversational interface demonstrates how the evolving graphs can be queried through text and image inputs, offering a multimodal means of scholarly engagement with historical data.

This study contributes to the fields of digital humanities and computational archival science by proposing a new generative approach to archival interpretation, one that couples autonomous feature discovery with human interpretive judgment. The resulting framework offers a pathway toward dynamic, explainable, and continuously learning archival knowledge systems that reimagine how context and meaning can be reconstructed from complex historical document collections.

**Primary authors:** GNANASEKARAN, Rajesh Kumar (the University of Maryland); Prof. MARCIANO, Richard (University of Maryland, College Park)

**Presenter:** GNANASEKARAN, Rajesh Kumar (the University of Maryland)

**Session Classification:** Artificial Intelligence (AI) - IV

**Track Classification:** Track 10: Artificial Intelligence (AI)



Contribution ID: 140

Type: **Oral Presentation**

## **JUNO Distributed Computing Infrastructure: Commissioning and Early Data Taking**

*Thursday, 19 March 2026 14:30 (30 minutes)*

The Jiangmen Underground Neutrino Observatory (JUNO) is a 20 kton liquid scintillator detector located in South China. Its construction as completed on November 2024. JUNO aims at an unprecedented 3% energy resolution at 1 MeV to determine neutrino mass ordering and study oscillation phenomena with high precision.

To support this ambitious physics program, JUNO relies on a Distributed Computing Infrastructure (DCI) jointly developed by IHEP and several European data centers, leveraging WLCG technologies. The DCI is expected to handle about 3 PB of data per year, enabling storage, distribution, and coordinated analysis across three continents.

This contribution reports on the deployment and operational experience of the DCI during water and LS filling, commissioning phase and the first data-taking period. We describe the integration of computing resources, data management workflows, and performance metrics observed in early operations. Lessons learned from commissioning activities and initial physics runs are discussed, highlighting the readiness of the infrastructure to meet JUNO's long-term computing requirements.

**Primary author:** ANDRONICO, Giuseppe (INFN Sez. CT)

**Co-authors:** Dr ZHANG, Xiaomei (IHEP); ATHAYDE MARCONDES DE ANDRE, Joao Pedro (IPHC CNRS)

**Presenter:** ANDRONICO, Giuseppe (INFN Sez. CT)

**Session Classification:** Infrastructural Clouds and Virtualisation - II

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 141

Type: **Oral Presentation**

## Supercomputing, Databases and Data Ecosystems: The Best of Three Worlds for Efficient Data-Driven Workflows

*Tuesday, 17 March 2026 16:00 (25 minutes)*

As supercomputing facilities increasingly run data analytics and artificial intelligence (AI) workloads, an efficient handling of external data sources, of storage and of data flows becomes paramount. However, often this is just addressed via optimization of parallel file systems, network connectivity and I/O libraries. In industry, it is common practice to support data-driven applications with optimized database-management systems, object storage, and data-streaming and caching techniques. The respective backend technologies are also increasingly important for data exchange and sharing via standardised methods (e.g., the IDSA Dataspace protocol) which give a response to data availability challenges in AI.

We thus argue that the IT behind data-driven workflows must be evolved at supercomputing centres – which is one core idea of our Extreme Data Analytics project “EXA4MIND”. On the computing-infrastructure side, the centres have shifted from exclusively running High-Performance-System (HPC) clusters to offering a mixed landscape of HPC, Cloud-Computing, GPU-computing, quantum-computing and other systems.

EXA4MIND makes best usage of such infrastructure and provides flexibly deployable modules for boosting and managing data-driven workflows. For piloting, it involves IT4Innovations (CZ) and LRZ (DE) as HPC centres and treats application cases in molecular dynamics (MD), autonomous driving, smart vineyards and public-health data. Within this scope, the project produces the following modules for Extreme Data processing: (i) storage submodules for the management of temporarily-instantiated or longer-term data stores optimised for each application case (object storage middleware, SQL and noSQL databases including vector databases); (ii) the Advanced Query and Indexing System (AQIS), where Extreme Data and AI workflows are orchestrated and data queries across backends can be executed via natural-language query and through extensions to the LEXIS Platform 2 enabling efficient orchestration across computing systems; (iii) toolboxes for data transfer, validation, preprocessing and analytics; (iv) connectivity modules enabling FAIR research data management and connectivity to European or international data ecosystems such as Data Spaces according to the IDSA Dataspace protocol..

In this contribution, we present the EXA4MIND approach and architecture, giving users from science, SMEs and industry a direction for realising Extreme Data workflows.

*This research received the support of the EXA4MIND (“EXtreme Analytics for MINing Data Spaces”) project, funded by the European Union’s Horizon Europe Research and Innovation Programme, under Grant Agreement N° 101092944. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them. The authors gratefully acknowledge the IT resources provided by IT4Innovations and by LRZ. This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254).*

**Primary authors:** Dr HACHINGER, Stephan (Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities); Dr GOLASOWSKI, Martin (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava); Mr HAYEK, Mohamad (Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities); Prof. KARAGOZ, Pinar (De-

partment Of Computer Engineering, Middle East Technical University); Prof. TOROSLU, Ismail Hakki (Department Of Computer Engineering, Middle East Technical University); Prof. OTYEPKA, Michal (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava); Mr CIZ, David (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava); Dr HURYCH, David (valeo.ai); Dr HARSH, Piyush (Terraview GmbH); Mrs CULELL, Maria Ignacia (AUSTRALO); Dr SLANINOVA, Katerina (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava); Mr FURMANEK, Radek (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava); Mr SWIATKOWSKI, Jan (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava); Dr MARTINOVIC, Jan (IT4Innovations National Supercomputing Center (IT4I), VŠB - Technical University of Ostrava)

**Presenter:** Dr HACHINGER, Stephan (Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities)

**Session Classification:** Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

**Track Classification:** Track 9: Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

Contribution ID: 142

Type: **Oral Presentation**

## Evolution of IFARM: A Multi-Schedd HTCondor Cluster for Korean Scientific Communities at GSDC

*Tuesday, 17 March 2026 16:25 (25 minutes)*

IFARM is a computing farm operated by GSDC/KISTI that supports computational experiments for several Korean scientific communities. IFARM has been in service since 2019 by integrating the previously separate computing-farm services that had been provisioned for each scientific community. As of the end of 2025, the service targets are the CMS Tier-3 and ALICE Tier-3 communities. By mid-2025, the BIO community (a Korean bioinformatics group) had also been hosted on this farm, but it is currently spun off into a separate service in July, 2025.

IFARM is a multi-schedd HTCondor cluster, where each community possesses 1 HTCondor user access point node. There are shared hosts for all communities : 1 HTCondor central manager node, 1 XRootD frontend node, 5 XRootD backend nodes, a CVMFS cache node and 76 worker nodes with around 5000 logical cores in total. IFARM is currently being utilized for interactive analysis, local analysis, and code development and testing by the users of CMS Tier-2 and ALICE Tier-1, respectively, at GSDC.

By the end of 2025, the introduction of the new servers for substituting the old servers and enlargement of XRootD storage capacity of CMS Tier-3 are planned at IFARM. We want to share our current challenges regarding HEP computing services and discuss the direction for improving IFARM's system architecture in relation to its upgrade. The application of system virtualization is actively pursued for all computing services, including CMS Tier-3 and ALICE Tier-3, as well as FCC (a newly launched computing service by the end of 2025), considering the current system users' behaviour patterns. We also plan to build a reliable system by duplicating key service nodes, including UI nodes, to ensure high availability of scientific computing service. Additionally, we would like to discuss here, the methodologies that enhances the current system operation methods or system monitoring method for proactively mitigating failures and for streamlining incident response.

**Primary author:** JEONG, Heeseok (Korea Institute of Science and Technology Institute)

**Co-authors:** Dr YOON, Heejun (Korea Institute of Science and Technology Information); Dr RYU, Geonmo (Korea Institute of Science and Technology Information)

**Presenter:** JEONG, Heeseok (Korea Institute of Science and Technology Institute)

**Session Classification:** Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

**Track Classification:** Track 9: Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators

Contribution ID: 143

Type: **Oral Presentation**

## Data and Computing Infrastructure for KM3NeT: Towards Full Deployment

*Wednesday, 18 March 2026 16:22 (22 minutes)*

The KM3NeT Collaboration is building two Cherenkov neutrino detectors in the depths of the Mediterranean Sea to study both the intrinsic properties of neutrinos and cosmic high-energy neutrino sources. Neutrinos are elementary particles with no electric charge and almost no mass, interacting only through the weak force. These characteristics allow the neutrino particles to travel straight across vast cosmic distances, on the contrary, this makes it extremely difficult to detect them. In the rare case that interaction occurs, neutrinos can create charged secondary particles, such as muons. These particles may travel faster than light moves in the medium, and when this occurs they emit a cone of photons known as Cherenkov radiation. Neutrino detectors are thus exploiting this characteristic finding the best conditions in large, transparent media, typically deep water or ice, where this light can be recorded by optical sensors.

KM3NeT addresses these challenges with two detectors: ORCA (Oscillation Research with Cosmics in the Abyss), optimised for atmospheric neutrinos in the GeV range to study neutrino oscillations and determine the neutrino mass ordering; and ARCA (Astroparticle Research with Cosmics in the Abyss), designed for TeV–PeV astrophysical neutrinos, targeting cosmic accelerators and the sources of high-energy cosmic rays.

Both detectors share the same fundamental component, the Digital Optical Module (DOM), a glass sphere instrumented with 31 photomultipliers and the necessary electronics to transmit detected light signals. The DOMs are then arranged in vertical strings of 18 units called Detection Units (DUs), which are anchored to the seabed. The communication between the underwater detectors and shore-based computing centers is ensured via electro-optical cables. At completion, ORCA will consist of a single modular block of 115 DUs, while ARCA will comprise two such blocks, instrumenting a volume of water of about 1Km<sup>3</sup>.

Thanks to the modular architecture of the detector, data acquisition, processing, and analysis are already in progress.

**This contribution focuses on the KM3NeT data and computing infrastructure, describing the current computing model and its evolution toward distributed Grid-based resources as the experiment approaches full deployment.** The expected data volume grows roughly linearly with the number of DUs, reaching hundreds of terabytes per year at full deployment. To handle this load, Rucio file catalogue will manage distributed data storage and replication, and the DIRAC workload manager will orchestrate large-scale data processing and Monte Carlo simulations. This infrastructure allows the collaboration to scale efficiently with detector size while supporting real-time data acquisition, processing, and analysis workflows.

**Primary authors:** Dr VAZQUEZ DE SOLA FERNANDEZ, Francisco (Nikhef); BRUNO, Riccardo (INFN)

**Presenter:** BRUNO, Riccardo (INFN)

**Session Classification:** Physics and Engineering Applications - II

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 144

Type: **Poster Presentation**

## Enhancing Virtual Research Environments for Environmental Research Using Advanced AI

Virtual Research Environments (VREs) constitute an essential aspect in attaining globally impactful research. With rapid environmental changes worldwide, including global warming, significant challenges that require collaboration between the environmental sub-domains need to be addressed. While workflows from individual Research Infrastructures (RIs) exist, environmental change is a global phenomenon that requires massive collaboration on a global scale. To achieve this level of collaboration and knowledge exchange at scale, federated access to computing resources, shared datasets or metadata, and shared analysis codes are vital assets for research infrastructures.

The ENVRI-Hub NEXT project aims to address these collaboration challenges by providing seamless access to environmental data and services across multiple European Environmental Research Infrastructures and by addressing key challenges in environmental data discovery, access, composition, processing, and usability. It enhances the collaboration and interoperability of RIs by providing tools for secure authentication, automated data processing, and search, while supporting virtual research environments to facilitate reproducing research analyses.

The ENVRI-Hub provides searching capabilities to locate environmental content, including data, metadata, or notebooks, in three different ways. It offers a classical search system powered by a Large Language Model (LLM), which uses content from ENVRI-Hub to provide users with a summarised response followed by content categorised into four classes: i) Web Content, ii) Metadata (supplemented by Catalogue of Services for expert users), iii) APIs, and iv) Jupyter Notebooks. Furthermore, it offers a dialogue-based search system, which complements classical search with a natural, conversational style, allowing users to discuss and analyse content from ENVRI-Hub. Finally, it provides a virtual research environment to run and analyse the notebooks retrieved by the search systems, allowing exploration of datasets, metadata, and APIs. Overall, to consolidate all services, a federated authentication system is implemented, enabling secure authentication for any user and allowing RIs to index their data while adhering to FAIR principles seamlessly.

The ENVRI-Hub primarily provides a unified search feature across 13 RIs. This feature supports researchers affiliated with these research infrastructures and independent scholars in searching for environmental data and metadata. It helps them learn about dataset attributes and obtain downloadable links to datasets and associated notebooks for analysis.

By implementing the FAIR principles end-to-end, the ENVRI-Hub transforms heterogeneous RI assets (datasets, metadata, APIs, web contents, and notebooks) into Findable (enriched indexing across 13 RIs), Accessible (federated AAI), Interoperable (cross-disciplinary index mappings), and Reusable resources (reusable Jupyter notebooks), powered by generative LLMs. This FAIR-by-design enables researchers to seamlessly integrate multimodal content, including text, data, and Jupyter notebooks, thereby enabling reproducible research analyses and tackling environmental challenges globally. In summary, the ENVRI-Hub realises the VRE vision by enabling scientific knowledge through secure, federated access to heterogeneous resources, while hiding infrastructure complexity and accelerating reproducible and collaborative environmental science.

**Primary authors:** ISLAM, Nafis Tanveer (University of Amsterdam); DRAGO, Federico (EGI Foundation); GUTIERREZ, Marta (EGI Foundation); DOLBER, Delphine (Euro-Argo ERIC); BUNDKE, Ulrich (F. Juelich); ZHAO, Zhiming (University of Amsterdam)

**Presenters:** DRAGO, Federico (EGIFoundation); ZHAO, Zhiming (University of Amsterdam)

**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 145

Type: **Oral Presentation**

## Open Cloud Mesh (OCM) - Towards IETF standardization of trusted resource sharing

*Tuesday, 17 March 2026 11:40 (20 minutes)*

“Open Cloud Mesh (OCM) is a server federation protocol that is used to notify a receiving party that they have been granted access to some resource.”

This is the simple, yet effective introduction to the IETF draft of the Open Cloud Mesh (OCM) specification. The proliferation of cloud storage solutions has led to a fragmented landscape, where seamless and secure sharing of resources across organizational and platform boundaries remains a significant challenge. The Open Cloud Mesh (OCM) protocol addresses this gap by providing a standardized server federation mechanism that enables interoperable sharing of cloud resources between independent service providers.

With OCM being one of the key requirements of the EOSC federation, it ultimately aims to contributing to a ‘web of FAIR (Findable, Accessible, Interoperable and Reuseable) data and services’ for science in Europe, upon which a wide range of value-added services can be built. During the EOSC Symposium 2025, EOSC Node candidates for a trusted OCM federation have been announced, with participants from the EU Node, CERN, SURF, Finland, and EUDAT, implemented by underlying OCM-compatible technologies from ownCloud, Nextcloud, CERNbox, and openCloud. This presentation will detail the architecture, core protocol flows, and security considerations of OCM, highlighting its potential to enable a truly open and interoperable cloud ecosystem. Attendees will gain insights into implementation strategies, integration scenarios, and the broader impact of OCM on collaborative research and cross-cloud interoperability.

**Primary authors:** Mr LO PRESTI, Giuseppe (CERN); Mr NORDIN, Micke (SUNET); FREITAG, Richard (SUNET)

**Presenter:** FREITAG, Richard (SUNET)

**Session Classification:** FAIR, Sovereign & Trusted Data - I

**Track Classification:** Track 6: FAIR, Sovereign & Trusted Data

Contribution ID: 146

Type: **Oral Presentation**

## Building a Virtual Research Environment for the Einstein Telescope Project

*Friday, 20 March 2026 11:15 (25 minutes)*

In this contribution, we present the development of a Virtual Research Environment (VRE) for the Einstein Telescope (ET) project, specifically in the Bologna research unit, designed to support collaborative, high-performance, and reproducible research within the ET community. The Einstein Telescope is a next-generation underground gravitational-wave observatory that will explore the Universe throughout its cosmic history. Achieving its ambitious scientific objectives - ranging from probing black-hole physics and neutron-star matter to studying dark energy and the early Universe - requires advanced computational and data-analysis infrastructures.

The ET Bologna VRE is built upon the so-called BETIF/DIFAET computing infrastructure, funded by the Italian NRRP (National Recovery and Resilience Plan): it adopts a modular, cloud-native architecture based on open-source technologies such as Docker, Kubernetes, Jupyter, and Rucio/Reana developed at CERN. This design enables both interactive analyses and large-scale computations within an orchestrated and containerized environment. The platform is fully customizable, supporting multiple software stacks via the CERN Virtual Machine File System (CVMFS) and providing seamless integration with external Rucio Storage Elements for distributed data management. Authentication and authorization are handled through Indigo-IAM, ensuring compliance with the ET federation's identity and access management policies.

Furthermore, the system supports heterogeneous computing resources, including CPU- and GPU-accelerated environments, and allows dynamic scaling according to workload and user requirements. Through its Python-friendly interface and integration with common scientific frameworks, the VRE lowers the entry barrier for analysis development while guaranteeing portability and reproducibility of workflows across the collaboration.

Beyond its immediate application to data analysis and algorithm prototyping, the ET Bologna VRE serves as a testbed for future computational strategies within the broader ET project. It demonstrates how local resources can be orchestrated into a flexible, cloud-native environment, paving the way for a distributed, sustainable, and collaborative data-analysis model essential for the next era of gravitational-wave astronomy.

**Primary authors:** BONACORSI, Daniele (University of Bologna); Dr LORUSSO, Marco (INFN - CNAF); DAL PRA, Stefano (Istituto Nazionale di Fisica Nucleare); DIOTALEVI, Tommaso (INFN and University of Bologna)

**Presenter:** Dr LORUSSO, Marco (INFN - CNAF)

**Session Classification:** Virtual Research Environment (VRE) - II

**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 147

Type: **Oral Presentation**

## Advancing Fundamental Research and Space Economy in Italy through the ICSC large-scale national initiative

*Wednesday, 18 March 2026 16:00 (22 minutes)*

Forty months after its inception under the Italian National Recovery and Resilience Plan (NRRP), the National Centre for High-Performance Computing, Big Data, and Quantum Computing - established and managed by the ICSC Foundation - has reached a mature and productive phase, consolidating its role as a strategic infrastructure for research, innovation, and industrial competitiveness. The Centre's mission to strengthen national computational capabilities and promote cross-sectoral collaboration between academia, research institutions, and industry has led to significant progress across its network of hubs and spokes.

In this contribution, we present the main results achieved during this first phase, with emphasis on frontier research in theoretical and experimental physics. Various activities have advanced the data-intensive and computational frontiers of high-energy, astroparticle and gravitational-wave physics. By integrating high-performance computing, cloud and distributed infrastructures, and advanced data analytics, the Centre has supported an expanding portfolio of scientific collaborations, developed novel tools for large-scale simulation, and fostered an interoperable ecosystem linking national and European research infrastructures. The deployment of heterogeneous computing architectures, including GPU- and FPGA-based systems, has also accelerated analysis pipelines and reinforced synergies across scientific domains.

As we move towards the future of the Centre, efforts will focus on consolidating these achievements and ensuring the long-term sustainability of the developed infrastructures, while expanding collaborations with European and industrial partners and exploring emerging opportunities in exascale computing and artificial intelligence.

**Primary authors:** DIOTALEVI, Tommaso (INFN and University of Bologna); BONACORSI, Daniele (University of Bologna); BOCCALI, Tommaso (INFN); MALVEZZI, Sandra (INFN)

**Presenter:** DIOTALEVI, Tommaso (INFN and University of Bologna)

**Session Classification:** Physics and Engineering Applications - II

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 148

Type: **Oral Presentation**

## Exploiting the Model-Context Protocol for Earth Observation Data Access and Workflow Orchestration

*Thursday, 19 March 2026 11:44 (22 minutes)*

The EO4EU project [1] democratises access to Earth Observation (EO) data by providing a comprehensive platform that caters to a wide spectrum of stakeholders, from researchers to policymakers. The EO4EU platform facilitates the seamless retrieval of EO data and the orchestration of complex computational and machine learning workflows. To this aim, the EO4EU platform integrates a semantic knowledge graph to enhance EO metadata discoverability, supports data fusion across heterogeneous sources, and enables advanced data visualisation techniques, including immersive VR/XR environments.

The Model-Context Protocol (MCP) [2] is a lightweight, standards-based protocol designed to support the dynamic coupling of Large Language Models (LLMs) and contextual data in distributed computing environments. MCP is built on top of JSON-RPC, a stateless, transport-agnostic remote procedure call protocol that enables structured communication between clients and services. MCP introduces a set of abstractions, namely tools, resources, and prompts, tailored to the operational logic of LLMs. These abstractions allow clients to discover the functionalities of local and remote platforms. Tools represent executable components; resources encapsulate data or metadata; and prompts define suggested context to instantiate or run a tool. This design promotes modularity, reusability, and interoperability across heterogeneous systems. It enables the integration of existing platforms with LLMs regardless of their accessibility model, whether proprietary, open weight or open source, deployed locally or accessed remotely.

As part of the EO4EU project, we have developed a prototype MCP server for the EO4EU platform. This server leverages the EO4EU REST APIs and encapsulates them into MCP-compliant tools. Each tool is designed to perform a goal-oriented operation, such as data discovery or workflow management, exposing to MCP clients a description which is interpretable and actionable by LLMs. This approach enables seamless interaction between language models and the EO4EU platform, facilitating the discovery of datasets tailored to specific use cases and the management of workflows from a single user interface.

Lessons learned include the importance of managing LLM context window length: verbose API responses can quickly exhaust the available context window, but the MCP server can mitigate this by filtering out irrelevant information. Additionally, a one-to-many mapping between MCP tools and API endpoints can help in reducing the length of the context. Moreover, in general, several challenges remain, such as LLM hallucinations which may result in infinite MCP tool invocation loops. Finally, we observed that swapping the underlying LLM can significantly alter the system's behaviour, stressing the importance of the specific implementation of MCP abstractions.

### References:

[1] <https://www.eo4eu.eu/>

[2] <https://modelcontextprotocol.io/>

**Primary authors:** PISA, Claudio; FORNARI, Federico (ECMWF); ANTONACCI, Marica (ECMWF (on leave from INFN)); ALBUGHDADI, Mohanad (ECMWF); KAPROL, Tolga (ECMWF); BAOUSIS, Vasileios (ECMWF)

**Presenter:** PISA, Claudio

**Session Classification:** Artificial Intelligence (AI) - III

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 149

Type: **Oral Presentation**

## Offloading CMS data analysis on a distributed high-throughput platform with RDataFrame

*Thursday, 19 March 2026 16:00 (30 minutes)*

The ability to ingest, process, and analyze large datasets within minimal timeframes is a cornerstone of modern big data applications. In High Energy Physics (HEP), this need becomes increasingly critical as the upcoming High-Luminosity phase of the LHC at CERN is expected to produce data volumes approaching 100 PB per year. Recent advancements in resource management and open-source computing frameworks - such as Jupyter, Dask, and HTCondor - are driving a shift from traditional batch-oriented workflows toward interactive, high-throughput analysis environments.

Within this context, and leveraging the computing resources of the Italian “National Center for High-Performance Computing, Big Data, and Quantum Computing (ICSC)”, a scalable analysis platform has been developed. Such system allows users to dynamically distribute workloads across local Kubernetes resources or offload them to remote infrastructures through interLink, a technology that extends the Virtual Kubelet concept to federate heterogeneous resources, like High-Throughput Computing (HTC), High-Performance Computing (HPC), and Cloud systems, under a unified orchestration layer.

The platform’s performance has been then evaluated using a representative use case: the study of the CMS Drift Tubes (DT) muon detector performance, in phase-space regions driven by analysis needs. By exploiting the declarative model of ROOT RDataFrame (RDF) and its distributed execution via Dask, the study demonstrates significant improvements in scalability and speed-up compared to traditional serial workflows. These results confirm the effectiveness of the proposed distributed analysis approach, addressing the computational challenges posed by the High-Luminosity LHC era.

**Primary authors:** DIOTALEVI, Tommaso (INFN and University of Bologna); BATTILANA, Carlo (University of Bologna and INFN); FANFANI, Alessandra (University of Bologna and INFN); ROSSI, Elvira (University of Naples and INFN); SPIGA, Daniele (INFN-PG); TEDESCHI, Tommaso (University and INFN, Perugia (Italy)); CIANGOTTINI, Diego (INFN Perugia); DORIA, Alessandra (INFN); PARDI, Silvio (INFN-Napoli); SPISSO, Bernardino (INFN); STELLACCI, Simona Maria (INFN)

**Presenter:** DIOTALEVI, Tommaso (INFN and University of Bologna)

**Session Classification:** Infrastructure Clouds and Virtualisations - III

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 150

Type: **Oral Presentation**

## Deploying the INDIGO PaaS Orchestration System Components on a Kubernetes-Based Infrastructure Using ArgoCD

*Thursday, 19 March 2026 16:30 (30 minutes)*

The National Institute for Nuclear Physics (INFN) manages the INFN Cloud, a federated cloud platform offering a customizable portfolio of IaaS, PaaS, and SaaS services tailored to the needs of the scientific communities it supports. The PaaS services are defined through an Infrastructure as Code approach, combining TOSCA templates to model application stacks, Ansible roles for automated configuration, Docker containers for packaging software and runtimes, and Helm charts for deploying applications on Kubernetes clusters. This approach allows the platform to provide flexible, reproducible, and consistent environments for a wide range of scientific workloads.

The platform's federation middleware is based on the INDIGO PaaS Orchestration system, which integrates several open-source microservices. Among these, the INDIGO PaaS Orchestrator is responsible for managing high-level deployment requests from users and coordinating the provisioning process across multiple federated IaaS platforms, ensuring efficient utilization of distributed resources and streamlined application delivery. Over the past year, development efforts have focused on replacing legacy components with new, modular services to extend system functionalities and mitigate security vulnerabilities. To achieve long-term maintainability, Python was adopted as the primary programming language, chosen for its readability, ease of maintenance, and compatibility with modern development practices, ensuring that the system can be efficiently updated and extended in the future.

This contribution presents the activity aimed at deploying the microservices composing the INDIGO PaaS Orchestration system on a Kubernetes-based infrastructure using ArgoCD. The migration is expected to improve resource management, simplify service deployment, and enable continuous integration and delivery through ArgoCD's GitOps approach.

Kubernetes plays a central role by providing resilience, self-healing capabilities, load balancing, and automated recovery, ensuring high availability and fault tolerance for critical orchestration services. ArgoCD manages and synchronizes microservice deployments directly from Git repositories, guaranteeing configuration consistency, version control, and traceability of all changes.

The ultimate objective is to achieve a more resilient, maintainable, and automated deployment environment, fully aligned with modern cloud-native and DevOps best practices.

**Primary author:** GIOMMI, Luca (INFN CNAF)

**Co-authors:** SINISI, Francesco (INFN-CNAF); SAVARESE, Giovanni (INFN Bari); COSTANTINI, Alessandro (INFN-CNAF); DONVITO, Giacinto (INFN)

**Presenter:** GIOMMI, Luca (INFN CNAF)

**Session Classification:** Infrastructure Clouds and Virtualisations - III

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 151

Type: **Oral Presentation**

## A Big Data Platform for Scalable Monitoring and Analytics at INFN-CNAF

*Thursday, 19 March 2026 17:00 (30 minutes)*

An heterogeneous data management and analytics platform has been deployed over the years at INFN-CNAF in Bologna, with the goal to support both system administrators and user-support teams in monitoring, diagnostics, and operational decision-making across the data center, that serves multiple particle and astroparticle experiments (and beyond) in Italy and internationally. The platform is built around a general-purpose message-handling backbone based on Apache Kafka, enabling scalable ingestion of high-volume, high-velocity data streams. On top of this backbone, a log-analysis stack - comprising Logstash, OpenSearch, and OpenSearch Dashboards, the open-source fork of the ELK suite maintained by AWS - provides powerful indexing, search, and visualization capabilities, including advanced authentication and authorization plugins. This contribution presents the architecture of the platform and the design choices that guided its implementation, complemented by a use case on farming services that demonstrates the platform's applicability in real-world scenarios. We also report on the extensions and enhancements introduced to transform the platform into a robust environment for Big Data analytics and AI applications aimed at improving data-center monitoring, alerting, and operational control. Finally, we illustrate the range of user-facing capabilities that make the platform a flexible and accessible tool for advanced analytics.

**Primary authors:** MICHELOTTO, Diego (INFN-CNAF); LORUSSO, Marco (University of Bologna and INFN)

**Presenter:** LORUSSO, Marco (University of Bologna and INFN)

**Session Classification:** Infrastructure Clouds and Virtualisations - III

**Track Classification:** Track 8: Infrastructure Clouds and Virtualizations

Contribution ID: 152

Type: **Oral Presentation**

## **Integrating AI into Data-Center Operations: Methods for Detection, Prediction, and Adaptive Control**

*Thursday, 19 March 2026 11:22 (22 minutes)*

The increasing complexity and scale of modern data centers generate operational environments where the ability to detect anomalies, anticipate failures, and optimize resource usage is becoming critically important. Recent advances in machine learning and artificial intelligence offer powerful techniques for extracting actionable insights from heterogeneous monitoring data, ranging from logs and metrics to event streams and security signals. In this contribution, we explore a set of AI-driven approaches that can be applied to a specific large-scale computing facility hosted at INFN-CNAF in Bologna. We will discuss methods for anomaly detection, predictive alerting and fault classification. Particular attention is given to techniques capable of leveraging high-volume, real-time data pipelines, including deep learning models for temporal analysis, clustering algorithms for behaviour profiling, and hybrid systems combining statistical baselines with learned representations. We also outline how such tools could support operational workflows, improve reliability, and reduce downtime, ultimately enhancing the overall efficiency of data-center operations. This study also establishes the methodological foundations for future integration with the INFN-CNAF Big Data Platform, which is expected to provide the unified data backbone for advanced operational analytics. The goal of this contribution is to highlight promising research directions and practical use cases where AI can provide measurable value to large distributed computing infrastructures.

**Primary authors:** BONACORSI, Daniele (University of Bologna); LORUSSO, Marco (University of Bologna and INFN); MICHELOTTO, Diego (INFN-CNAF); DELL'AGNELLO, Luca (INFN); MORGANTI, Lucia (INFN - CNAF)

**Presenter:** LORUSSO, Marco (University of Bologna and INFN)

**Session Classification:** Artificial Intelligence (AI) - III

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 153

Type: **Oral Presentation**

## Quantum Invertible Neural Networks for Fast Calorimetry as Quantum Generative Model with high interpretability at Future Colliders

*Wednesday, 18 March 2026 14:22 (22 minutes)*

Future collider experiments will require fast, scalable, and highly accurate calorimeter simulation to cope with unprecedented event rates and detector granularity. While machine-learning-based simulation has become a central strategy, the next step may come from quantum-native generative models capable of learning expressive, bijective mappings between physics parameters and detector observables. In this work we explore the feasibility of such an approach through a fully quantum-inspired generative framework based on Quantum Invertible Neural Networks (qINN), an architecture that remains largely unexplored in high-energy physics.

The qINN model provides a reversible transformation between input kinematic variables and calorimeter shower observables, enabling explicit likelihood evaluation and enhancing interpretability, an increasingly valuable feature for, and not only, detector design studies. The framework is developed and validated using ATLAS fast calorimeter simulation (fastCaloSim) software and runs on ATLAS open data, ensuring full reproducibility and allowing easy adaptation to different detector concepts in sight of future colliders application. Moreover, its structure accords to the current best-performing algorithm in fastCaloSim which takes advantage of a classical INN core.

We present methodology, software integration, and initial performance studies on both energy deposition and shower shape generation, focusing on fidelity, robustness, and scalability with respect to detector geometry and granularity. Early results demonstrate the promise of qINN-based generative models as a foundation for quantum-enhanced fast calorimeter simulation, opening a pathway toward simulation tools tailored for next-generation collider detectors.

**Primary authors:** FRANCHINI, matteo; Mr CORCHIA, Federico Andrea Guillaume (University of Bologna and INFN-Bo)

**Presenter:** FRANCHINI, matteo

**Session Classification:** Physics and Engineering Applications -I

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 154

Type: **Oral Presentation**

# The AARC Compendium: A Practical Guide to Federated Identity Management for Research Infrastructures

*Friday, 20 March 2026 11:08 (18 minutes)*

The Authentication and Authorisation for Research and Collaboration (AARC) community, funded by the AAR-TREE project is releasing the AARC Compendium, a comprehensive introductory guide to implementing federated identity management for research infrastructures and their communities. Building on the AARC Blueprint Architecture (AARC BPA), the Compendium bridges the gap between technology, policy, and community needs.

The Compendium provides practical guidance across the full lifecycle of Authentication and Authorisation Infrastructure (AAI) implementation. A glossary of key terms and extensive FAQ section help demystify technical concepts and address common challenges. The guide covers implementation scenarios—from hosted AAI platforms to self-hosted proxy architectures—and presents an overview of the landscape of existing AAI solutions, including commonly used software, services, and complete hosted solutions already deployed across research infrastructures.

Key sections address technical requirements such as harmonised identity representation, authorisation and access control, and interoperability architecture. The Compendium also tackles critical policy aspects, including security, data protection, and practical use of the AARC Policy Development Kit. Importantly, it demonstrates how to bridge legal, policy, and technology considerations. The final version includes real-world use cases from research communities, infrastructures, and service providers offering concrete implementation examples and proven patterns that can be adapted to different contexts, helping to inspire and guide adoption across the research ecosystem.

**Target Audience:**

This presentation is directed at research communities and infrastructure providers offering services in federated environments

**Aims of the presentation:**

- Raise awareness of the AARC Compendium and its practical value
- Demonstrate how the guide addresses real-world implementation challenges through use cases
- Encourage engagement with the AARC community and promote uptake of federated identity management practices

**Primary authors:** SHORT, Hannah (CERN); KREMERS, Maarten (SURF); GROEP, David (Nikhef and Maastricht University); KELSEY, David (STFC-RAL)

**Presenter:** KREMERS, Maarten (SURF)

**Session Classification:** Networking, Security & Operations - II

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 155

Type: **Oral Presentation**

## Modeling Multivariate Sensor Dynamics with STGNNs: An Anomaly Detection Study at the INFN CNAF Tier-1 Facility

*Friday, 20 March 2026 11:20 (30 minutes)*

Modern data centres are cyber-physical infrastructures whose reliable operation depends on the continuous interaction of electrical and mechanical subsystems. Detecting anomalous behaviour in these environments is essential for ensuring operational continuity, improving energy efficiency, and enabling early fault prediction.

We present an anomaly-detection case study using Spatio-Temporal Graph Neural Networks (STGNNs) [1] to analyse signals collected at the INFN CNAF Tier-1 data centre, a high-throughput computing facility in Bologna, Italy. The Building Management System records electrical measurements from transformers and DRUPS units, as well as cooling-plant data such as chiller and pump power consumption, refrigerant pressures and temperatures, and water and air inlet/outlet temperatures.

Traditional statistical and machine-learning approaches often struggle to model the nonlinear dependencies and temporal dynamics present in such heterogeneous multivariate data. To address this challenge, we represent the sensor measurements as a dynamic graph in which nodes correspond to individual variables and edges capture evolving correlations or causal relationships. Such spatial-temporal graph can be described as a series of attributed graphs, which effectively represent (multivariate) time series data in conjunction with evolving structural information over time.

Our approach applies STGNNs to jointly encode spatial and temporal structure. The architecture combines graph-convolutional layers that aggregate information across related variables with recurrent temporal components that track system evolution over time. The spatial and temporal layers are trained jointly because the fused models are end-to-end differentiable. Factorized and coupled model variants exist, depending on the combination style of the spatial and temporal modules. Our unified modelling strategy is implemented through the PyTorch Geometric Temporal framework [2] and provides a robust foundation for anomaly detection for continuous monitoring of high-dimensional industrial sensor networks.

[1] Seo, Y., Defferrard, M., Vandergheynst, P., & Bresson, X. (2018, November). Structured sequence modeling with graph convolutional recurrent networks. In *International conference on neural information processing* (pp. 362-373). Cham: Springer International Publishing.

[2] Rozemberczki, B., Scherer, P., He, Y., Panagopoulos, G., Riedel, A., Astefanoaei, M., ... & Sarkar, R. (2021, October). Pytorch geometric temporal: Spatiotemporal signal processing with neural machine learning models. In *Proceedings of the 30th ACM international conference on information & knowledge management* (pp. 4564-4573).

**Primary authors:** Dr ZURLO, Giovanni (INFN CNAF); RONCHIERI, Elisabetta (INFN CNAF); Dr SCARPONI, Luigi Benedetto (INFN CNAF); Dr CESINI, Daniele (INFN CNAF); Dr COSTANTINI, Alessandro (INFN CNAF)

**Presenter:** Dr ZURLO, Giovanni (INFN CNAF)

**Session Classification:** Artificial Intelligence (AU) - V

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 156

Type: **Oral Presentation**

# Feasibility Study of QNN-Based Latent Space Generation for Fast Calorimeter Simulation in ATLAS

*Wednesday, 18 March 2026 14:44 (22 minutes)*

The increasing demands on simulation statistics for HL-LHC analyses challenge the scalability of traditional calorimeter simulation within all LHC collaborations. Fast simulation techniques based on machine learning have proven effective, yet further improvements may arise from quantum-inspired models.

In this study we investigate the feasibility of integrating Quantum Neural Network (QNN) components into the ATLAS fast calorimeter simulation framework, to evaluate potential gains in generative performance and generalisation while preserving practical inference times.

We focus on a hybrid pipeline in which a quantum Generative Adversarial Network (qGAN) is employed during training to generate a compressed latent space of calorimeter shower representations. The resulting latent vectors are then passed to classical neural networks for sample inference within the fast simulation framework. This approach allows quantum-inspired elements to be embedded where they can have the largest impact - namely in learning compact, expressive latent manifolds - while retaining full compatibility with ATLAS production workflows.

We report on implementation details, benchmarking against established fast simulation components, and performance metrics including fidelity, stability, and computational cost. Preliminary results indicate that QNN-augmented latent space models are technically feasible within ATLAS software and can achieve competitive performance. This work outlines a realistic roadmap for incorporating quantum-inspired techniques into fast simulation for HL-LHC applications.

**Primary authors:** CORCHIA, Federico Andrea Guillaume (University of Bologna and INFN Bologna); FRANCHINI, Matteo (University of Bologna and INFN Bologna); APICELLA, Emilio (University of Bologna and INFN Bologna)

**Presenter:** CORCHIA, Federico Andrea Guillaume (University of Bologna and INFN Bologna)

**Session Classification:** Physics and Engineering Applications -I

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 157

Type: **Oral Presentation**

## Point-Cloud Machine Learning for Detector Data Integration in ATLAS at the LHC

*Friday, 20 March 2026 10:50 (30 minutes)*

Large-scale scientific experiments increasingly depend on trustworthy data infrastructures and AI-driven workflows to extract reliable insights from complex sensor systems. In the ATLAS experiment at CERN's LHC, each collision generates heterogeneous information across multiple sub-detectors, making the accurate integration of these data streams essential for ensuring reconstruction quality and downstream scientific validity.

A key challenge is the association of particle trajectories from the tracking system with corresponding energy deposits in the calorimeter. Traditional rule-based methods perform this step sequentially and rely on hand-crafted, physics-based criteria, which can struggle in high-density or rare events and do not fully exploit the complementary nature of the underlying data.

We explore a machine learning approach that treats detector hits as point clouds and learns associations directly from spatial and contextual patterns. The talk will highlight the practical training considerations and challenges encountered, including data representation choices, training strategies and performance metrics. In this way, we aim at offering insights relevant for developing trustworthy AI methods in large-scale scientific applications.

**Primary authors:** Prof. VIVARELLI, Iacopo (University of Bologna); CLISSA, Luca (University of Bologna & INFN); Dr SWIATLOWSKI, Maximilian (TRIUMF)

**Presenter:** CLISSA, Luca (University of Bologna & INFN)

**Session Classification:** Artificial Intelligence (AI) - V

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 158

Type: **Oral Presentation**

# The Dutch EOSC Node: Federating National Research Infrastructure with European Open Science

*Friday, 20 March 2026 10:50 (25 minutes)*

The European Open Science Cloud (EOSC) creates a federated environment where researchers across Europe can publish, find, and reuse data, tools, and services. Implemented as a 'system of systems', EOSC connects data repositories, research infrastructures, and e-infrastructures through a network of national and thematic nodes. The Dutch EOSC Node, operated by SURF, is one of the initial candidate nodes working to bridge national infrastructure with the broader European ecosystem.

While SURF has been running established services like SURF Research Cloud (computing resources), ResearchDrive (file sync and share), and FAIR data repositories for years, the SURF EOSC Node itself is still in active pilot phase. A demonstration at the EOSC Symposium in November 2025 showcased progress in federating these services within the EOSC Federation. We're building out the node infrastructure, leveraging services we already operate reliably for Dutch research institutions and now making them available through a federated EOSC node approach—extending their reach and impact across Europe while maintaining interoperability with other nodes.

The Dutch pilot demonstrates how national research infrastructure providers can serve their research communities while contributing to pan-European open science. By federating proven services using AAI standards and federation protocols, SURF is addressing practical challenges of cross-border service delivery—maintaining institutional control while enabling seamless collaboration that truly advances open science in Europe and globally.

This presentation shares our current status, lessons learned during the build-up phase, and our confident path forward based on solid operational foundations. We'll discuss what's working, remaining challenges, and how our federated approach might inform other nations developing their EOSC nodes.

**Target Audience:**

This presentation is directed at research communities and infrastructure providers as well as the general audience with an interest in Open Science

**Aims of the presentation:**

- Share honest from the Dutch EOSC Node pilot phase
- Demonstrate how federated infrastructure advances open science in practice
- Encourage collaboration, knowledge exchange and promote uptake of federated management practices

**Primary author:** KREMERS, Maarten (SURF)**Presenter:** KREMERS, Maarten (SURF)**Session Classification:** Virtual Research Environment (VRE) - II**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 159

Type: **Oral Presentation**

## Weather and Landslide Data Analysis Using Spatio-Temporal and Categorical Data Mining

*Wednesday, 18 March 2026 14:44 (22 minutes)*

Spatio-temporal data mining is effective for extracting useful information from the occurrence frequencies and patterns of real-world physical phenomena. The author has previously proposed a spatio-temporal and categorical data mining method that not only extracts occurrence frequencies and patterns from spatio-temporal features, but also performs semantic interpretation of relationships between two different events by incorporating categorical information.

In this study, with the aim of examining whether this method is applicable beyond the commerce domain, the author investigate whether future landslide occurrences can be predicted by computing correlation measures between weather data and landslide occurrence data. A distinctive feature of the proposed approach is that prediction is performed not by conventional statistical calculations, but by searching for past events similar to the current situation and utilizing the subsequent temporal evolution of those events. In this presentation, the author describe the details of the method, report experimental results, and discuss its potential applicability to other domains.

**Primary author:** HAYASHI, Yasuhiro (Musashino University)

**Presenter:** HAYASHI, Yasuhiro (Musashino University)

**Session Classification:** Artificial Intelligence (AI) - I

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 160

Type: **Oral Presentation**

## Conformal Prediction for Reliable Uncertainty Quantification in Scientific AI Models

*Wednesday, 18 March 2026 14:22 (22 minutes)*

Reliable uncertainty quantification is essential in scientific applications, where predictive results must be supported by a transparent assessment of confidence. Among the many approaches proposed for this purpose, Conformal Prediction (CP) is especially compelling because it offers finite-sample, distribution-free coverage guarantees and can calibrate uncertainty on top of any trained model without requiring retraining.

Using the Higgs Uncertainty Dataset as a benchmark, we illustrate how CP can produce prediction intervals with guaranteed coverage levels while making no assumptions about the underlying data distribution. We also compare CP with traditional likelihood-based inference and common ML-driven uncertainty estimation techniques to highlight their respective strengths and limitations. Taken together, these results show that CP provides a competitive and flexible approach that integrates seamlessly with existing ML workflows, making it a promising building block for the development of trustworthy and reproducible AI in scientific research.

**Primary authors:** BONACORSI, Daniele (University of Bologna); Dr PLINI, Leonardo (University La Sapienza); CLISSA, Luca (University of Bologna & INFN)

**Presenter:** CLISSA, Luca (University of Bologna & INFN)

**Session Classification:** Artificial Intelligence (AI) - I

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 161

Type: **Oral Presentation**

## **MLOps for Scientific Applications: Building Reliable, Reproducible, and Transparent AI Workflows**

*Wednesday, 18 March 2026 16:22 (22 minutes)*

Modern scientific workflows increasingly rely on Machine Learning (ML) models whose development, deployment, and validation must meet high standards of reliability, transparency, and reproducibility. Yet many scientific ML pipelines still lack robust engineering practices, making experiments difficult to track, compare, and replicate. In this contribution, we present a structured MLOps methodology tailored for scientific data analysis, integrating data exploration and versioning, experiment tracking, and model lineage within a unified workflow.

We illustrate the methodology with a representative use case inspired by a High-Energy Physics application, showing how a fully instrumented pipeline built on MLOps principles can improve reproducibility, enable transparent comparison across models, and support effective collaboration in multidisciplinary teams.

Our results highlight how well-designed MLOps workflows can bridge the gap between experimental science and modern ML engineering, providing a scalable foundation for deploying reliable AI models in scientific applications.

**Primary authors:** BONACORSI, Daniele (University of Bologna); CLISSA, Luca (University of Bologna & INFN); PLINI, Leonardo (University La Sapienza)

**Presenter:** CLISSA, Luca (University of Bologna & INFN)

**Session Classification:** Artificial Intelligence (AI) - II

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 162

Type: **Oral Presentation**

## Transfer Learning to Overcome Domain Shift in Football Analytics and Beyond

*Thursday, 19 March 2026 11:00 (22 minutes)*

Domain shift occurs when the distributions of features, underlying behaviours or operational conditions differ between source (training) and target (test) domains, causing models to struggle when applied to data from a different context than the one used for training. To mitigate this, several transfer learning approaches have been proposed to reuse and adapt knowledge acquired in the source domain, thereby enhancing model performance in the target domain.

In this work we apply transfer learning to the domain of football analytics, where models developed for one league, season or team often degrade when transferred to another due to differences in playing style, data distribution, formation dynamics or sensor setups. By leveraging publicly available event and tracking data in football, we explore how transfer learning techniques can reduce this degradation. Compared to training a model from scratch on the target domain, our approach shows improved robustness and generalisation under domain shift. Using standard machine learning models and targeted transfer learning steps, we present a workflow that is effective in sport and broadly applicable to other scientific fields facing similar domain-shift challenges.

**Primary authors:** Dr MACALUSO, Antontio (German Research Center for Artificial Intelligence (DFKI)); CLISSA, Luca (University of Bologna & INFN)

**Presenter:** CLISSA, Luca (University of Bologna & INFN)

**Session Classification:** Artificial Intelligence (AI) - III

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 163

Type: **Oral Presentation**

## Development of an LLM-Based System for Automatic Code Generation from HEP Publications

*Friday, 20 March 2026 09:00 (22 minutes)*

Ensuring the reproducibility of physics results is a critical challenge in high-energy physics (HEP). In this study, we aim to develop a system that automatically extracts analysis procedures from HEP publications and generates executable analysis code capable of reproducing published results, leveraging recent advances in large language models (LLMs).

Our approach employs open-source LLMs to accurately extract event selection criteria, definitions of physical quantities, and other relevant information described in scientific papers. The system also traces referenced publications when necessary, enabling the construction of a selection list that remains faithful to the original analysis. Based on this extracted information, the system automatically generates analysis code, which is executed on ATLAS Open Data to validate the reproducibility of the published results.

Specifically, we utilize proton-proton collision data recorded in 2015 and 2016 by the ATLAS experiment and released as open data. This dataset allows direct comparison between our automatically generated analysis results and those reported in the literature. By comparing with manually developed analysis code used as a baseline, we evaluate the current performance of open-source LLMs in terms of code quality, computational efficiency, and physics validity.

This research contributes to the advancement of reproducible science in HEP and explores the potential of LLM-driven automation in physics analysis workflows. In the long term, our study envisions research environments where physicists can perform data analysis through natural-language interaction, and where automated verification and review support improves the reliability and accessibility of scientific publications. In this presentation, we report the status of our prototype system and initial performance evaluation results.

**Primary authors:** TANAKA, Junichi (University of Tokyo); SAITO, Masahiko (ICEPP, The University of Tokyo); KISHIMOTO, Tomoe (KEK)

**Presenter:** SAITO, Masahiko (ICEPP, The University of Tokyo)

**Session Classification:** Artificial Intelligence (AI) - IV

**Track Classification:** Track 10: Artificial Intelligence (AI)

Contribution ID: 164

Type: **Oral Presentation**

## The successful removal of IPv4 from WLCG wide area network links

*Thursday, 19 March 2026 11:40 (20 minutes)*

The Worldwide Large Hadron Collider Computing Grid (WLCG) community's deployment of IPv6 on its worldwide storage infrastructure is very successful and has been presented by us at earlier ISGC conferences. The campaign to deploy IPv6 on CPU services and all worker nodes is progressing well. Dual-stack IPv6/IPv4 is not, however, a viable long-term solution; the HEPiX IPv6 Working Group has focused on studying where and why IPv4 is still being used, and how to change such traffic to IPv6. The agreed end goal is to turn IPv4 off wherever possible and for all wide area data transfers to happen over IPv6, to simplify both operations and security management.

This paper will report on our work since the ISGC2025 conference. Firstly, we will report on our ongoing campaign to encourage the deployment of IPv6 on CPU services and Worker Nodes. Then, we will present further work to identify and correct the use of IPv4 between two dual-stack endpoints. The Research Networking Technical Working Group has identified marking the IPv6 packet header as one approach for understanding complex large data flows. This provides another driver for full transition to the use of IPv6 in WLCG data transfers. Several WLCG sites have now stated their wish to move all their services to IPv6-only in the near future. We are now close to being able to support this, should the experiments they support agree.

We present the working group's proposed plans and timescale for moving WLCG wide area network links on LHCOPN and LHCONE to "IPv6-only". In the summer of 2025, ESnet and the Tier1 centres in the USA successfully removed IPv4 from their two LHCOPN links to CERN. We will present how this was done and our work and plans for the next LHCOPN links to remove IPv4. Our proposal is to complete this work on all LHCOPN links during 2026 in time for the WLCG 2027 data challenge (DC27). The long-term aim is to also remove IPv4 from WLCG data transfers over the LHCONE network before the start of HL-LHC run 4 in 2030. This would require completion of much of that transition before the data challenged DC29. We will present the steps towards making this possible.

**Primary author:** KELSEY, David (STFC-RAL)

**Presenter:** KELSEY, David (STFC-RAL)

**Session Classification:** Networking, Security & Operations - I (11:00 - 12:40)

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 165

Type: **Oral Presentation**

## An introduction to SCDF and a case study of sPhenix compu+ng

*Wednesday, 18 March 2026 15:06 (22 minutes)*

Scientific Computing and Data Facilities (SCDF) at Brookhaven Lab began in 1997 when the Relativistic Heavy Ion Collider (RHIC) and ATLAS Computing Facility was established. The full-service scientific computing facility has since supported some of the most notable physics experiments, including Broad RAnge Hadron Magnetic Spectrometers (BRAHMS), Pioneering High Energy Nuclear Interaction eXperiment (PHENIX), PHOBOS Collaboration, and Solenoidal Tracker at RHIC (STAR), by providing dedicated data processing, storage, and analysis resources for these diverse, expansive experiments with general computing capabilities and support for users.

Today, this history of providing useful, resilient, and large-scale computational science, data management, and analysis infrastructure has grown, and the SCDF (formerly SDCC) has evolved to support additional facilities and experiments. These include the National Synchrotron Light Source II (NSLS-II) and Center for Functional Nanomaterials (CFN) at Brookhaven Lab, as well as other DOE Office of Science User Facilities; the ATLAS experiment at CERN's Large Hadron Collider in Europe; and Belle-II at KEK in Japan and DESY in Germany. SCDF also is planning for its role in future experiments with the sPHENIX, Deep Underground Neutrino Experiment, and Electron-Ion Collider.

sPHENIX is a radical makeover of the PHENIX experiment, one of the original detectors designed to collect data at Brookhaven Lab's Relativistic Heavy Ion Collider. It includes many new components that significantly enhance scientists' ability to learn about quark-gluon plasma (QGP), an exotic form of nuclear matter created in RHIC's energetic particle smashups.

The SCDF has been providing sPhenix with compu+ng, processing, storage and networking since 2023. We want to cover how it was accomplished in SCDF.

**Primary author:** CHOU, Tim (Brookhaven National Laboratory)

**Presenter:** CHOU, Tim (Brookhaven National Laboratory)

**Session Classification:** Physics and Engineering Applications -I

**Track Classification:** Track 1: Physics and Engineering Applications

Contribution ID: 166

Type: **Oral Presentation**

## From Batch to Interactive: Interactive Analysis Workbench for HEP Data Analysis at IHEP

*Friday, 20 March 2026 11:40 (25 minutes)*

The IHEP computing platform is facing new demands in data analysis, including restricted access to login nodes, increasing needs for code debugging tools, and more efficient data access for collaborative workflows. To address these challenges, we have developed INK, a web-based “Interactive aNalysis worKbench” that enables users to access IHEP login nodes, cluster computing resources, and data directly through their browsers. INK also allows users to create both general-purpose analytical web applications and experiment-specific customized applications, making it easier to utilize IHEP’s computing resources. The INK system adopts a decoupled front-end/back-end architecture and provides multi-source user authentication and authorization mechanisms for heterogeneous computing and data resources. A set of unified interfaces is exposed to the front-end, ensuring consistent environments while enabling seamless integration with interactive analysis tools.

**Primary author:** ZHANG, Xuantong (IHEP)**Presenter:** ZHANG, Xuantong (IHEP)**Session Classification:** Virtual Research Environment (VRE) - II**Track Classification:** Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)

Contribution ID: 167

Type: **Oral Presentation**

## Deploying Security Operations Centre capabilities

*Thursday, 19 March 2026 12:00 (20 minutes)*

Supported by the work of the Security Operations Centre Working Group that has been reported on at ISGC previously we have built experience across the Research & Education community deploying security operations centre capabilities. This work has used the reference design of the working group, developed to provide guidance on the elements required to deploy a Security Operations Centre, in various ways. A recent addition to the working group has been a focus on people and process.

In this presentation we will report on the lessons learned from the experience of deploying these capabilities from the perspective of different organisations, to support others in developing their own activities. While the technology stack is an important element of this work, the sustained staffing and processes needed are a vital part of this work.

**Primary authors:** CROOKS, David (UKRI STFC); VALSAN, Liviu (CERN)

**Presenters:** CROOKS, David (UKRI STFC); VALSAN, Liviu (CERN)

**Session Classification:** Networking, Security & Operations - I (11:00 - 12:40)

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 168

Type: **Oral Presentation**

## IT Security training for multi domain/multi stakeholder infrastructures

*Friday, 20 March 2026 11:26 (18 minutes)*

In this presentation we will share our experience in providing training for security personnel providing operational security for different types of large distributed infrastructures.

Depending on the target audience and the topics to be addressed, these were developed in two categories, technical hands on training and table top exercises.

In the first category either a technical training infrastructure is required, here also the existing infrastructure can be used, but poses extra challenges on the trainers. Here the focus is on developing technical skills, and ideally also covering the higher level aspects of incident response communication and incident response coordination. The latter aspect is of particular interest in distributed infrastructures and is usually not covered in existing similar training activities.

Table top exercises, as they were developed by us, on the other hand rather focus on the higher level security posture of the whole organisation. Based on the existing policies and procedures also escalations to press and legal can be addressed, making an active involvement of the management possible/wanted. This usually only happens when IT security incidents have a high impact on the organisation itself, and experience with this situation is usually not available, and needed communication channels either do not yet exist (for example involving external security provider) or are not regularly used. Still a "tested" procedure for these cases is extremely important, since any delay here can easily get very costly.

Here we present our experience developing and running security exercises in both of the above mentioned categories of security exercises which also are part of the Thematic CERN School of computing (security) and led to the security training organised here at ISGC-2026.

**Primary authors:** KRASOVEC, Barbara (Jožef Stefan Institute (JSI) /EGI); KOURIL, Daniel (CES-NET / Masaryk University); CROOKS, David (UKRI STFC); GABRIEL, Sven (Nikhef/EGI); GROEP, David (Nikhef and Maastricht University); KELSEY, David (STFC-RAL)

**Presenter:** GABRIEL, Sven (Nikhef/EGI)

**Session Classification:** Networking, Security & Operations - II

**Track Classification:** Track 7: Network, Security, Infrastructure & Operations

Contribution ID: 170

Type: **Oral Presentation**

## The Internet of FAIR data and Services

*Tuesday, 17 March 2026 09:10 (40 minutes)*

### ABSTRACT:

The 'next generation of the Internet' will have to serve machines much like the current Internet serves people. Preferably we should avoid some of the apparent downsides of the current situation, where a few major players dominate the Internet and the Web. This is why we started the Leiden Initiative for FAIR and Equitable Science, LIFES. In this not for profit, globally oriented association, we connect international members who are 'serious about FAIR and Equitable access to information'. LIFES is a public private partnership composed of Application and Service Providers, Users and Recognised Expert Communities. The latter category is a very important one as it allows existing and new communities with a strong mandate for a particular topic or domain, but also with a regional or national mandate to 'take care' of the relevant services, standards, tooling and infrastructure in their domain or region. The growing LIFES community formally 'recognises' the leading role of a Recognised Expert Community in its domain or region. This minimises the risk for duplication of efforts and reinvention of wheels.

LIFES as an association, connects its members to jointly implement and innovate on outstanding challenges. Next to a rigorous commitment to the FAIR principles for machine actionable data stewardship, LIFES members commit to strive towards global 'data visiting', which leaves data at the source wherever possible and makes these data FAIR and visitable by algorithms. This approach mitigates many of the GDPR, privacy and security issues associated with more traditional data sharing and it also allows scientists and innovators from low connectivity areas to fully participate in the Internet of FAIR data and Services.

A discussion will be initiated on how Taiwan and the entire region can optimally participate in this global effort, which also relates to other regional and global opens science commons initiatives such as EOSC, AOSP and GOSC.

### BIOGRAPHY:

Barend Mons (born 1957, The Hague) is a molecular biologist and a FAIR data specialist. The first decade of his scientific career he spent on fundamental research on malaria parasites and later on translational research for malaria vaccines. In the year 2000 he switched to advanced data stewardship and (biological) systems analytics. He is most known for innovations in scholarly collaboration, especially nanopublications, and knowledge graph based discovery.

In 2012 Barend was appointed full Professor in biosemantics in the Department of Human Genetics at the Leiden University Medical Center (LUMC) in The Netherlands.[1][2][3] In 2014 he organised the seminal FAIR conference at the Lorentz centre that led to the FAIR data initiative and GO FAIR. In 2015 he was appointed chair of the High Level Expert Group on the European Open Science Cloud.

From 2018 to 2023 Barend was the elected president of CODATA, the affiliated organisation on research data related issues of the International Science Council. He has also been the European representative in the Board on Research Data and Information (BRDI) of the National Academies of Sciences, Engineering, and Medicine in the USA. In 2023 he was also appointed professor at the Leiden Academic Centre for Drug Research.

In 2024, he was appointed as Fellow of the International Science Council. At his retirement in 2024 he was Knighted by the Dutch King in the 'Order of the Netherlands Lion', the oldest and highest reward for cultural and scientific contributions to the international society.

Barend is a frequent keynote speaker about FAIR and open science around the world, and continues to participate in various national and international boards.

To know more details of Prof. Barend Mons, please refer to [https://en.wikipedia.org/wiki/Barend\\_Mons](https://en.wikipedia.org/wiki/Barend_Mons)

**Presenter:** Prof. MONS, Barend

**Session Classification:** Opening Ceremony & Keynote Session I

Contribution ID: 171

Type: **Oral Presentation**

## **Title (TBC)**

*Tuesday, 17 March 2026 09:50 (40 minutes)*

**Presenter:** Prof. CHEN, Chii-Dong (Academia Sinica)

**Session Classification:** Opening Ceremony & Keynote Session I

Contribution ID: 172

Type: **not specified**

## Opening & Welcome Remarks

*Tuesday, 17 March 2026 09:00 (10 minutes)*

**Presenters:** Prof. MATYSKA, Ludek; Prof. WANG, Song-Ming

**Session Classification:** Opening Ceremony & Keynote Session I

Contribution ID: 173

Type: **Oral Presentation**

## **Machine Learning Physics” – an emergent new arena of research unifying AI and quantum physics**

*Wednesday, 18 March 2026 09:00 (45 minutes)*

### Abstract

Machine learning and physics have long been deeply intertwined, and there have been eras when their relationship came to the forefront. Even in today’s revolutionary AI development, physics has played a significant role—for example, in diffusion models. From a physics standpoint as well, an integrative perspective across various specialized domains is provided by innovative new mathematical frameworks, and machine learning serves as one such framework. Launched in fiscal year 2022, the MEXT Scientific Transformation Area Research (A) initiative “Foundation of Machine Learning Physics” was established to forge a new interdisciplinary field merging machine learning and quantum physics in Japan. Now in its fourth year, it has produced a wide range of research outcomes and functions as a central hub where many researchers gather. In this talk, I will introduce the goals of this initiative, illustrating them with specific research examples, and discuss the future relationship between machine learning and quantum physics. The research examples include quantum simulations accelerated by machine learning methods in particle physics, computational physics and condensed matter physics, as well as novel approaches to AI architectures such as quantum interpretation of diffusion models and gravitational symmetries within neural networks.

### Biography

Professor Koji Hashimoto is a leading theoretical physicist specializing in string theory, quantum gravity, and the interface between machine learning and fundamental physics. He earned his Ph.D. in Physics from Kyoto University in 2000 and has held research positions at the University of California, Santa Barbara; the University of Tokyo and RIKEN. In 2012 he was appointed as a full professor at Osaka University, and since 2021, he has been a professor at the Graduate School of Science, Kyoto University. Professor Hashimoto’s research encompasses string theory, holographic principle, and the application of deep learning to theoretical physics. His contribution includes analyses of quantum chaos using out-of-time-order correlators, quantum computational complexity and black holes, and neural networks representing quantum gravity spacetimes. He has authored several influential books, including “D-branes” (Springer, 2012) and “Deep Learning and Physics” (Springer, 2021). As a director he leads the MEXT “Machine Learning Physics” initiative, aiming to integrate AI methodologies with theoretical physics. In addition to his academic work, Professor Hashimoto has contributed to science communication by novels and essays, public lectures, art performance, musical collaboration, and supervising the physics in the film “Shin Ultraman” and the subtitles for “Oppenheimer”. His interdisciplinary approach continues to influence both scientific and public understanding of physics.

**Presenter:** Prof. HASHIMOTO, Koji (Department of physics, Kyoto University, Japan)

**Session Classification:** Keynote II

Contribution ID: 174

Type: **Oral Presentation**

## EGI's path towards an AI-ready Compute and Data Continuum for Research

*Wednesday, 18 March 2026 09:45 (45 minutes)*

### Abstract:

Artificial intelligence is transforming the practice of science. From protein structure prediction to climate modelling and high-energy physics, AI techniques are accelerating discovery, enabling researchers to extract insights from datasets of unprecedented scale and complexity. At the same time, the infrastructure demand of AI-driven research (large-scale GPU computing, access to curated scientific datasets, and interoperable platforms) is growing faster than any single institution or country can address alone.

Europe is responding with a coordinated policy framework, the AI Continent Action Plan, and has launched RAISE (Resource for Artificial Intelligence Science in Europe) to establish a pan-European virtual institute that pools compute, data, and talent across two dimensions: advancing safe frontier models ("Science for AI") and applying AI to scientific disciplines ("AI in Science"). Together, these initiatives form an emerging ecosystem where sovereign infrastructure, open science, and AI converge. In this context, the SCIANCE project will drive the definition of a Strategic Research and Innovation Agenda (SRIA) to identify AI-in-science research and innovation priorities and related key infrastructure needs.

Within this broader context, the EGI Federation is evolving from a compute federation into Europe's AI-ready compute and data continuum for research. The recently published EGI Federation Strategy 2026–2030 defines three pillars for this transformation: 1) expanding the Compute Federation with AI-capable resources and scientific datasets; 2) removing the policy and funding barriers that limit cross-border service delivery; and 3) strengthening EGI as an innovation ecosystem. Recent and ongoing EGI flagship projects, including iMagine for AI-powered aquatic imaging, interTwin for AI-driven digital twins, the EOSC Data Commons and RI-SCALE's Data Exploitation Platforms, implement this strategy in practice. Moreover, through the SPECTRUM project, EGI and leading research communities are now co-designing a Technical Blueprint and a Strategic Research, Innovation and Deployment Agenda (SRIDA) for the future compute and data continuum.

This keynote will provide an overview of how AI is reshaping scientific research, outline the European policy response, and offer a deep dive into EGI's strategy and concrete contributions to democratising AI-enabled scientific computing across Europe and beyond.

**Bio:** Sergio Andreozzi is Head of Strategy, Innovation and Communications at the EGI Foundation, where he leads strategic planning, governance, and the development of EGI's business and innovation agenda. He recently co-authored the EGI Federation Strategy 2026–2030 and is lead editor of the Strategic Research, Innovation and Deployment Agenda (SRIDA) in the SPECTRUM project. He also contributes to the SCIANCE project, where EGI leads work on infrastructure roadmapping and communications for AI in science. Sergio holds an Executive Master in Management of Research Infrastructures (University of Milano-Bicocca), a PhD in Computer Science (University of Bologna), and an MSc in Computer Science Engineering (University of Pisa).

**Presenter:** Dr ANDREOZZI, Sergio

**Session Classification:** Keynote II

Contribution ID: 175

Type: **Oral Presentation**

## **e-Science Activities in Asia-Pacific (Japan)**

*Wednesday, 18 March 2026 11:00 (15 minutes)*

**Presenter:** Dr AIDA, Kento (National Institute of Informatics)

**Session Classification:** e-Sciences Activities in Asia-Pacific

Contribution ID: 176

Type: **Oral Presentation**

## **e-Science Activities in Asia-Pacific (Taiwan)**

*Wednesday, 18 March 2026 11:15 (15 minutes)*

**Presenter:** TBC

**Session Classification:** e-Sciences Activities in Asia-Pacific

Contribution ID: 177

Type: **Oral Presentation**

## **e-Sciences Activities in Asia-Pacific (Australia)**

*Wednesday, 18 March 2026 11:30 (15 minutes)*

**Presenter:** TBC

**Session Classification:** e-Sciences Activities in Asia-Pacific

Contribution ID: 178

Type: **Oral Presentation**

## **e-Sciences Activities in Asia-Pacific (China)**

*Wednesday, 18 March 2026 11:45 (15 minutes)*

**Presenter:** Dr CHEN, Gang (Institute Of High Energy Physics)

**Session Classification:** e-Sciences Activities in Asia-Pacific

Contribution ID: 179

Type: **not specified**

## **e-Sciences Activities in Asia-Pacific (Thailand)**

*Wednesday, 18 March 2026 12:00 (15 minutes)*

**Presenter:** TBC

**Session Classification:** e-Sciences Activities in Asia-Pacific

Contribution ID: **180**

Type: **Oral Presentation**

## **e-Sciences Activities in Asia-Pacific (Korea)**

*Wednesday, 18 March 2026 12:15 (15 minutes)*

**Presenter:** TBC

**Session Classification:** e-Sciences Activities in Asia-Pacific

Contribution ID: **181**

Type: **Oral Presentation**

## **Title (TBC)**

*Thursday, 19 March 2026 09:00 (45 minutes)*

**Presenter:** Dr DART, Eli

**Session Classification:** Keynote - IV

Contribution ID: 182

Type: **Oral Presentation**

# From Terabytes to Insights: The Data and Computational Challenges of Distributed Acoustic Sensing

*Thursday, 19 March 2026 09:45 (45 minutes)*

## Abstract

Distributed Acoustic Sensing (DAS) is revolutionizing the Earth and environmental sciences by repurposing fiber-optic infrastructure into ultra-dense seismic networks that provide unprecedented spatiotemporal resolution of the Earth's subsurface and environment. This technology is driving transformative scientific breakthroughs, ranging from precision earthquake localization and early warning systems to broad-scale land-to-sea environmental monitoring, and critical borehole applications in geothermal energy and carbon capture and storage (CCS).

However, this high-fidelity sensing capability generates massive, continuous data streams that frequently exceed terabytes per day, creating a formidable "Big Data" challenge that outpaces the capacity of traditional processing workflows to extract timely insight. This keynote highlights the tension between these rich scientific opportunities and the logistical hurdles of data management, ultimately serving as an open call for the high-performance computing and artificial intelligence communities to collaborate on scalable architectures and next-generation algorithms that can unlock the full scientific potential of vast DAS datasets.

Prof. Hsin-Hua Huang

Position: Associate Research Fellow

Affiliation: Institute of Earth Sciences, Academia Sinica

Email: [hhuang@earth.sinica.edu.tw](mailto:hhuang@earth.sinica.edu.tw)

Website:

<https://sites.google.com/view/hsinhuahuang/home>

## Biography

Dr. Hsin-Hua Huang is currently an Associate

Research Fellow at the Institute of Earth Sciences, Academia Sinica, and the Executive Secretary of the Taiwan Earthquake Research Center. A leading expert in observational seismology, earth imaging, and geohazards, Dr. Huang specializes in developing advanced 3D/4D seismic tomography and ambient noise interferometry to resolve complex subsurface structures. He holds a Ph.D. from National Taiwan University (2013) and conducted postdoctoral research at Caltech and the University of Utah in the U.S.

Dr. Huang's pioneering work has yielded significant breakthroughs, including unveiling hidden magma reservoirs beneath the Yellowstone and Tatun volcanoes, deciphering complex fault structures through machine-learning catalogs, and enhancing earthquake early warning systems by incorporating rupture directivity effects. Recently, he has been at the forefront of introducing Distributed Acoustic Sensing (DAS) technology to Taiwan, applying fiber-optic sensing to revolutionize seismic observation and environmental monitoring.

Recognized for his academic excellence, Dr. Huang is a recipient of the NSTC Ta-You Wu Memorial Award (2021) and the Academia Sinica Presidential Scholars Program (2025). He plays a pivotal role in the global geosciences community, serving as an Associate Editor for top-tier journals including *Geology* and *Journal of Geophysical Research*, and frequently delivers invited talks at major international conferences such as AGU and AOGS.

**Presenter:** Prof. HUANG, Hsin-Hua (Academia Sinica)

**Session Classification:** Keynote - IV