

CERN School of Computing on IT Services 2025

Report of Contributions

Contribution ID: 1

Type: **not specified**

Opening Session

Monday, 3 November 2025 09:15 (30 minutes)

Presenter: PACE, Alberto (CERN)

Contribution ID: 2

Type: **not specified**

Opening Lecture: The need for IT Service in accelerator and particle physics

Monday, 3 November 2025 09:45 (1 hour)

Starting from accelerator and particle physics, we'll try to see what are the needs of the experiments and accelerator people in terms of IT services.

Presenter: PONCE, Sebastien (CERN)

Contribution ID: 3

Type: **not specified**

Student self presentation

Monday, 3 November 2025 17:00 (1 hour)

Contribution ID: 4

Type: **not specified**

Creating and managing websites at CERN

Tuesday, 4 November 2025 10:45 (1h 30m)

An introduction to website provisioning and hosting at CERN with the services provided by IT. We will focus on creating websites with WordPress and go into the details of the WordPress service at CERN.

Presenters: DEL PIANTA, Lorenzo (CERN); SHARMA, Vasvi

Contribution ID: 5

Type: **not specified**

Core compute services (part 1 of 2)

Tuesday, 4 November 2025 13:30 (1 hour)

An in depth set of use case where IT services are heavily used for physics, analysis and engineering applications.

In part I of the series we will show use cases for Openstack, Linux and virtual-machine based configuration management.

Presenter: CASTRO LEON, Jose (CERN)

Contribution ID: 6

Type: **not specified**

Database Services (part 1 of 4) - Introduction to DBoD

Monday, 3 November 2025 11:15 (1 hour)

- Short intro about databases
- Presentation of the DBoD service
- How to create a database in DBoD
- How to connect to my DBoD

Presenter: NOWICKI, Andrzej (CERN)

Contribution ID: 7

Type: **not specified**

Project Management and documentation

Monday, 3 November 2025 16:00 (1 hour)

In this session we will explore the solutions for project management and effective software development.

We'll start by demonstrating how to plan and track project progress using Jira, available at <https://its.cern.ch>.

Next, we'll highlight how GitLab Pages can be utilized to store and share information, whether for internal use or a broader audience, showcasing instances of technical documentation delivered to end-users.

Finally, we'll explore Confluence for documentation storage, providing a walkthrough of its features and real-world examples of its use.

Presenter: BORGES AURINDO BARROS, Francisco (CERN)

Contribution ID: **8**

Type: **not specified**

Core compute services (part 2 of 2)

Wednesday, 5 November 2025 08:45 (1h 30m)

An in depth set of use case where IT services are heavily used for physics, analysis and engineering applications.

In part 2 of the series we will explore HTCondor, the high-throughput compute platform used for batch computing.

Authors: DELGADO PERIS, Antonio (CERN); JONES, Ben (CERN); GKONIS, Panagiotis (CERN)

Presenters: DELGADO PERIS, Antonio (CERN); GKONIS, Panagiotis (CERN)

Contribution ID: 9

Type: **not specified**

Core compute services (part 3 of 4)

An in depth set of use case where IT services are heavily used for physics, analysis and engineering applications.

In part 3 of the series we will explore the Slurm, the technology underlying the HPC platform at CERN.

Presenter: HØIMYR, Nils (CERN)

Contribution ID: **10**

Type: **not specified**

Using Openstack, Batch and Storage (provisional)

An in depth set of use case where IT services are heavily used for physics, analysis and engineering applications

Presenter: TENAGLIA, Giacomo (CERN)

Contribution ID: **11**

Type: **not specified**

Configuration & Infrastructure as Code

Wednesday, 5 November 2025 10:45 (1h 30m)

IT central configuration management services for virtual and physical hosts.

Presenter: TENAGLIA, Giacomo (CERN)

Contribution ID: 12

Type: **not specified**

Creation and maintenance of a website (part 2 of 2)

Overview of the Web Services Portal, and how it can be used for website creation and management at CERN. Website hosting and management services - Drupal/WordPress, WebEOS, GitLab pages etc - will be highlighted and their specific use cases outlined. The use of Matomo for web analytics will be demonstrated.

Presenter: SHARMA, Vasvi

Contribution ID: 13

Type: **not specified**

Database Services (part 2 of 4) - DBoD maintenance exercises

Wednesday, 5 November 2025 13:30 (1 hour)

- What are the typical tasks to be performed as DBoD owner Exercises on:
- Cloning mechanism
- Upgrades
- TLS certificates

Presenter: NOWICKI, Andrzej (CERN)

Contribution ID: 14

Type: **not specified**

Database Services (part 4 of 4) - DBoD maintenance exercises

Friday, 7 November 2025 11:15 (1 hour)

- What are the typical tasks to be performed as DBoD owner Exercises on:
- Cloning mechanism
- Upgrades
- TLS certificates

Presenter: NOWICKI, Andrzej (CERN)

Contribution ID: 15

Type: **not specified**

Database Services (part 3 of 4) - Oracle Database

Thursday, 6 November 2025 11:45 (30 minutes)

- Introduction of the Oracle Database service
- Resource portal as a way to manage Oracle database users - needed e-groups
- Other tooling provided by the team - Session Manager
- How to connect to the Oracle database
- Where to get the client?
- What is the tnsnames.ora file?

Presenter: NOWICKI, Andrzej (CERN)

Contribution ID: **16**

Type: **not specified**

Application development with DB BackEnd (part 5 of 5)

Dbod, MySql, Postgress, ...

Presenter: NOWICKI, Andrzej (CERN)

Contribution ID: 17

Type: **not specified**

Modern Application Development & Deployment

Tuesday, 4 November 2025 08:45 (1h 30m)

“I need to develop an application X for user community Y, which will need to be run and maintained over time”

1. Application development a. This part of the session will explore the multiple types of applications, how to leverage version control system Gitlab and its CI to have modern application deployment. This is organized as a workshop and will include a hands-on experience covering the best practices to develop containerized applications and strategies for deploying them.

Participants will begin by exploring various application types and learning how to leverage GitLab's version control and CI pipelines for efficient deployment. Through hands-on exercises, attendees will develop a simple application, while learning key concepts such as:

- Best practices for developing containerized applications.
- Writing DockerFiles for application deployment.
- Setting up continuous integration (CI) workflows to automate testing and build and publish Docker images.

Presenter: BORGES AURINDO BARROS, Francisco (CERN)

Contribution ID: 23

Type: **not specified**

Modern Application Development & Deployment (part 2 of 2)

“I need to develop an application X for user community Y, which will need to be run and maintained over time”

1. Application development a. This part of the session will explore the multiple types of applications, how to leverage version control system Gitlab and its CI to have modern application deployment. This is organized as a workshop and will include a hands-on experience covering the best practices to develop containerized applications and strategies for deploying them.

Presenter: BORGES AURINDO BARROS, Francisco (CERN)

Contribution ID: 24

Type: **not specified**

Authentication and authorization

Wednesday, 5 November 2025 16:00 (1 hour)

(Part of “Software development and hosting” track)

In this class, we will see how to:

- Enable authentication with CERN SSO
- Define role-based authorization for our applications
- Get tokens, and use them access APIs

Presenter: SHORT, Hannah (CERN)

Contribution ID: 25

Type: **not specified**

Deploying applications (part 1 of 2)

Tuesday, 4 November 2025 16:00 (1 hour)

In this lecture, we will understand the difference between IaaS, PaaS and SaaS.
Then, we will learn how to deploy custom and off-shelf applications to OKD PaaS.

Presenter: PIMPO, Alberto

Contribution ID: **26**

Type: **not specified**

Deploying applications (part 2 of 2)

Tuesday, 4 November 2025 17:00 (1 hour)

Exercises regarding how to deploy custom and off-shelf applications to OKD PaaS.

Presenter: PIMPO, Alberto

Contribution ID: 27

Type: **not specified**

Software development and hosting (part 7 of 7) - Deploying applications (exercises)

Presenter: PIMPO, Alberto

Contribution ID: 28

Type: **not specified**

Data Analysis Techniques using SWAN and REANA (part 1 of 3)

Wednesday, 5 November 2025 14:30 (1 hour)

In this first session, we will give an overview of the SWAN service. This will include the following points:

- Interface: classic and JupyterLab
- Creation of projects, notebooks and terminals
- Integration with CVMFS for software provisioning
- Integration with EOS for storage and CERNBox for sharing
- Use of GPUs
- Connection to Spark clusters

Moreover, we will give a live demo that participants will be able to follow along and get familiar with the basic features of SWAN.

Presenters: CASTRO, Diogo (CERN); ESTEVES MAXIMINO, Pedro Miguel; HENRIQUES SOBRAL, Rodrigo Fernando

Contribution ID: 29

Type: **not specified**

Data Analysis Techniques using SWAN and REANA (part 2 of 3)

Friday, 7 November 2025 08:45 (1 hour)

In the second session of this series, we shall present REANA reusable and reproducible analysis platform. REANA allows researchers to structure their data analyses by means of declarative workflow languages (CWL, Snakemake, Yadage) and run containerised data analysis pipelines on remote compute clouds (Kubernetes, HTCondor, Slurm).

In the first part of this session, we shall discuss the notions of computational reproducibility and reusability, underlying the importance of encapsulating the original computing environments by means of containers and documenting the steps necessary to arrive at results. We shall provide a brief introduction to declarative workflow languages and discuss its pros and cons when compared to imperative analysis code programming.

In the second part of this session, the participants will familiarise themselves with the REANA platform by means of running a simple analysis example. We shall use the <https://reana.cern.ch> instance at CERN to run a RooFit demo example.

Presenters: MCCLYMONT, Cameron Duncan; SIMKO, Tibor (CERN)

Contribution ID: 30

Type: **not specified**

Data Analysis Techniques using SWAN and REANA (part 3 of 3)

Friday, 7 November 2025 09:45 (1 hour)

In the third session of this series, we will propose short exercises using SWAN and REANA to cover more data analysis examples and use cases. The session will be split in two parts, one for each tool, where participants will be able to work on the exercises and get assistance from the lecturers.

Presenters: MCCLYMONT, Cameron Duncan; CASTRO, Diogo (CERN); ESTEVES MAXIMINO, Pedro Miguel; HENRIQUES SOBRAL, Rodrigo Fernando; SIMKO, Tibor (CERN)

Contribution ID: 31

Type: **not specified**

Services for Machine Learning applications (part 1 of 3)

Thursday, 6 November 2025 08:45 (1 hour)

This session will introduce the different phases in a ML lifecycle, and how IT services can help in each of the parts. In particular, it will:

- Overview of ML and use cases, containerization and how it helps out in defining single units of computation, isolate custom software environments, and ensure sustainability for reproducible results
- Demo how cloud native environments (Kubernetes and its ecosystem) can help manage those units of computation and scale them out to large amounts of resources
- Provide an example of how to scale out using both on-premises and public cloud resources, and when this might be useful and cost effective

Presenter: CHIORESCU, Raulian-Ionut

Contribution ID: 32

Type: **not specified**

Services for Machine Learning applications (part 2 of 3)

Thursday, 6 November 2025 10:45 (1 hour)

This session will focus on the infrastructure and low level tools required to efficiently deploy machine learning applications. In particular, it will cover:

- The different data types and how they can impact ML workloads, as well as support in different types of hardware and software libraries
- Key differences between CPUs and GPUs and how they impact ML workloads (training and serving)
- The available techniques in IT services for GPU sharing and partitioning. In particular, it will cover how applications can build on the existing Kubernetes service to simplify these operations
- Hands-on exercises on using GPUs for different types of workloads

Presenter: HANSEN, Hannes Jakob

Contribution ID: 33

Type: **not specified**

Services for Machine Learning applications (part 3 of 3)

Thursday, 6 November 2025 13:30 (1 hour)

This session will focus on available ML techniques for distributed training of models, hyperparameter optimization and model service. In particular, starting from a well known use case it will demonstrate:

- How to go from a script, to a docker image training on a single node, to a distributed training setup with multiple nodes
- How to do hyperparameter optimization, which kind of optimizers are available, how to monitor the workloads and how to publish the models
- How to serve models in production, at scale, with a simple http endpoint or embedding the model in an application

Presenter: LAHOUEL, Amine (CERN)

Contribution ID: 34

Type: **not specified**

Lunch

Contribution ID: 35

Type: **not specified**

Closing Session

Friday, 7 November 2025 16:45 (1 hour)

Presenter: PACE, Alberto (CERN)

Contribution ID: 36

Type: **not specified**

Storage (part 1 of 2) and re-requisites

Monday, 3 November 2025 13:30 (1 hour)

This two-part lecture series provides an overview of the various storage services at CERN. We will look into the motivation behind our large scale storage systems, cover some fundamentals and the design principles of the storage systems we've developed and use. We will also look into some practical use cases covering the ecosystem of the many storage systems we run in the IT department. This should serve as a basis for choosing the correct storage services for the applications you would develop and practical considerations into utilizing storage effectively.

Presenters: LEKSHMANAN, Abhishek (CERN); DEL MONTE, Gianmaria (CERN)

Contribution ID: 37

Type: **not specified**

Storage (part 2 of 2)

Monday, 3 November 2025 14:30 (1 hour)

Presenters: LEKSHMANAN, Abhishek (CERN); DEL MONTE, Gianmaria (CERN)

Contribution ID: 38

Type: **not specified**

Application security

Friday, 7 November 2025 13:30 (1h 30m)

Short introduction to best practices for secure development, testing and deployment

- Three golden rules for system security
- Software security, typical vulnerability types
- How security analysis tools can help
- Introduction to penetration testing
- Deployment security best practices

Presenter: BRILLAULT, Vincent (CERN)

Contribution ID: 39

Type: **not specified**

Authentication and authorization (Excercises)

Wednesday, 5 November 2025 17:00 (1 hour)

(Part of “Software development and hosting” track)

In this class, we will see how to:

- Enable authentication with CERN SSO
- Define role-based authorization for our applications
- Get tokens, and use them access APIs

Presenter: SHORT, Hannah (CERN)

Contribution ID: **40**

Type: **not specified**

Welcome address from the IT department head

Monday, 3 November 2025 09:00 (15 minutes)

Author: PORCARI, Enrica Maria (CERN)

Presenters: PORCARI, Enrica Maria (CERN); CAMPANA, Simone (CERN)

Contribution ID: 41

Type: **not specified**

Lightning talks

Thursday, 6 November 2025 09:45 (30 minutes)

- Evangelos Katralis - GPU acceleration of a Particle in Cell (PIC) code
- Jorge Belmonte - File Conversions and Batch Jobs Through Cluster Computing for the PLM Platform
- Andrzej Nowicki - AI Meets Databases: How to build an image search using pgvector extension of PostgreSQL

Presenters: NOWICKI, Andrzej (CERN); KATRALIS, Evangelos (Aristotle University of Thessaloniki (GR)); VELASQUEZ BELMONTE, Jorge Antonio

Contribution ID: 42

Type: **not specified**

Introduction to registry.cern.ch and Kubernetes

Tuesday, 4 November 2025 14:30 (1 hour)

Presenters: GAPONCIC, Diana (CERN); MUNDAY, Jack Charlie