Digital Solutions @ CERN

Nick Ziogas

Knowledge Transfer Group
CERN

Overview

- Innovation drivers – context
- Fast inference and Machine/Deep Learning
- Large data set analysis tools
- Large scale storage, management tools and integration
- Industrial controls and monitoring
- Robotics
- Simulation software
- Collaboration tools
- Conclusions

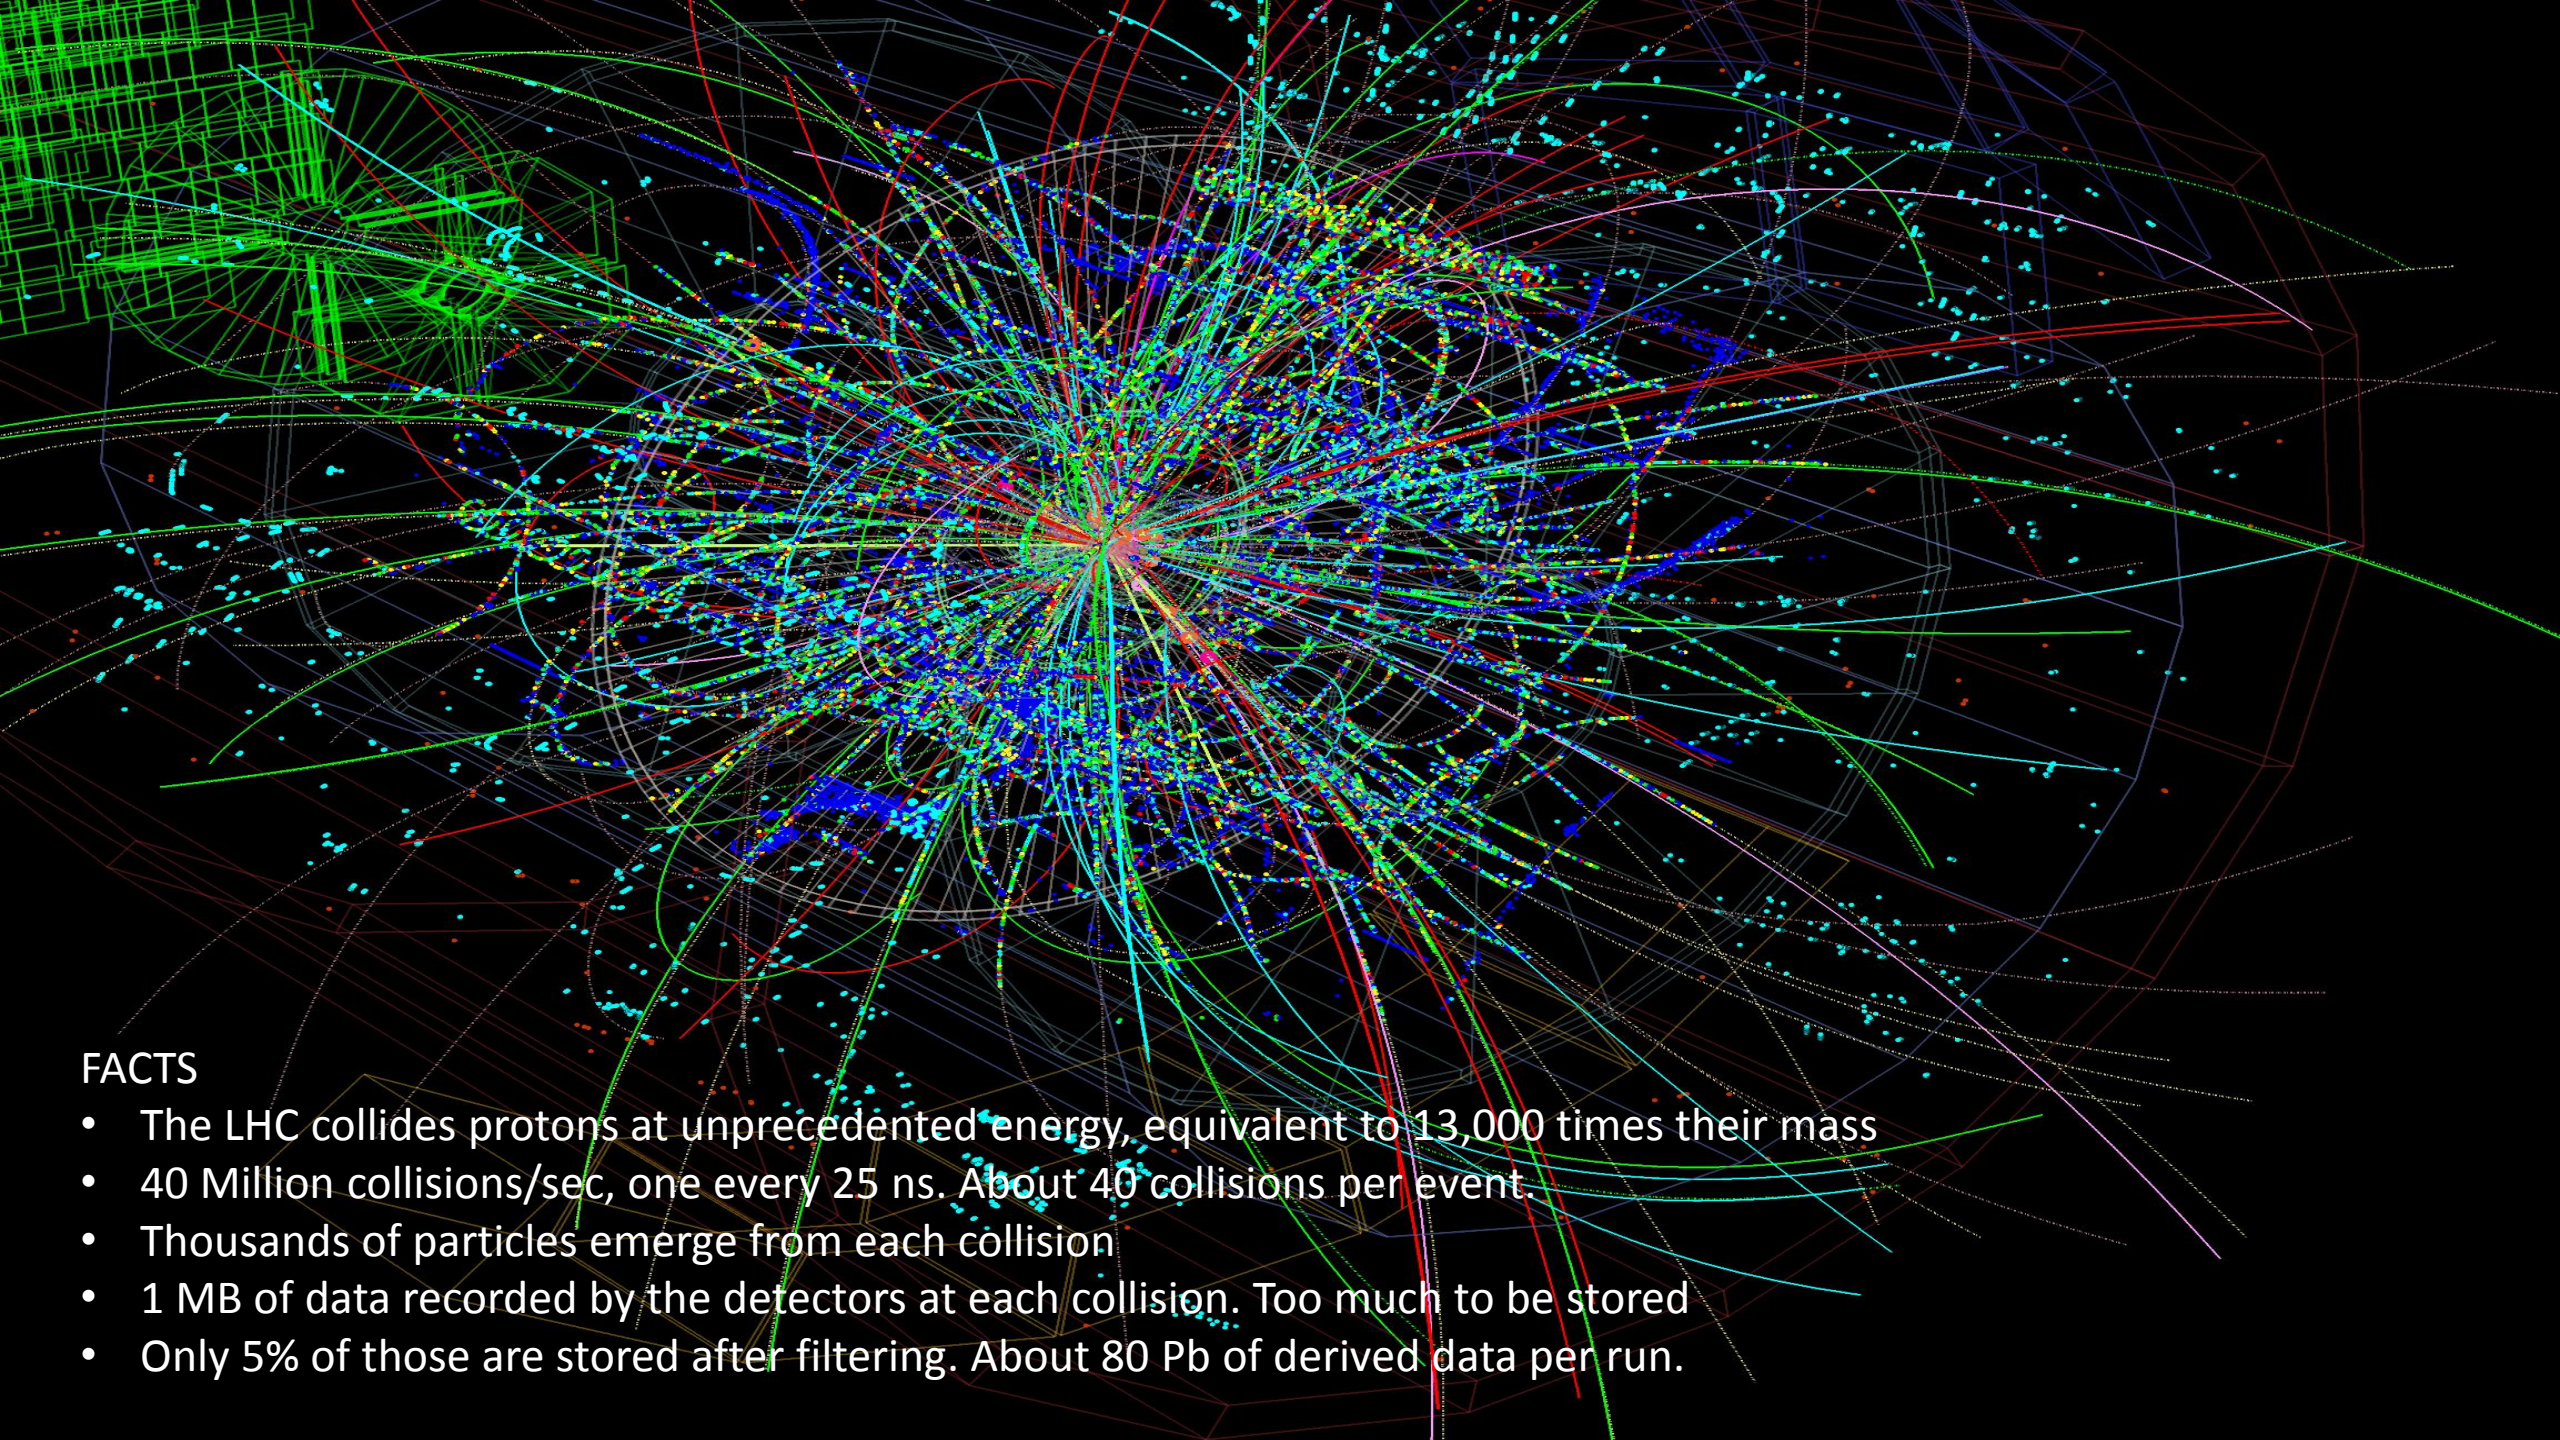

FACTS

- The LHC collides protons at unprecedented energy, equivalent to 13,000 times their mass
- 40 Million collisions/sec, one every 25 ns. About 40 collisions per event.
- Thousands of particles emerge from each collision
- 1 MB of data recorded by the detectors at each collision. Too much to be stored
- Only 5% of those are stored after filtering. About 80 Pb of derived data per run.

The NEXT Challenge

- Upgrade of the LHC the CERN flagship accelerator which led to the discovery of the Higgs boson.
- Increase of instantaneous luminosities by a factor of five larger than the LHC nominal value
- Enlarge experiment data sample by one order of magnitude compared with the LHC baseline programme.
- Operational around (2026)-2027

The LHC Big Data Challenge – HL LHC

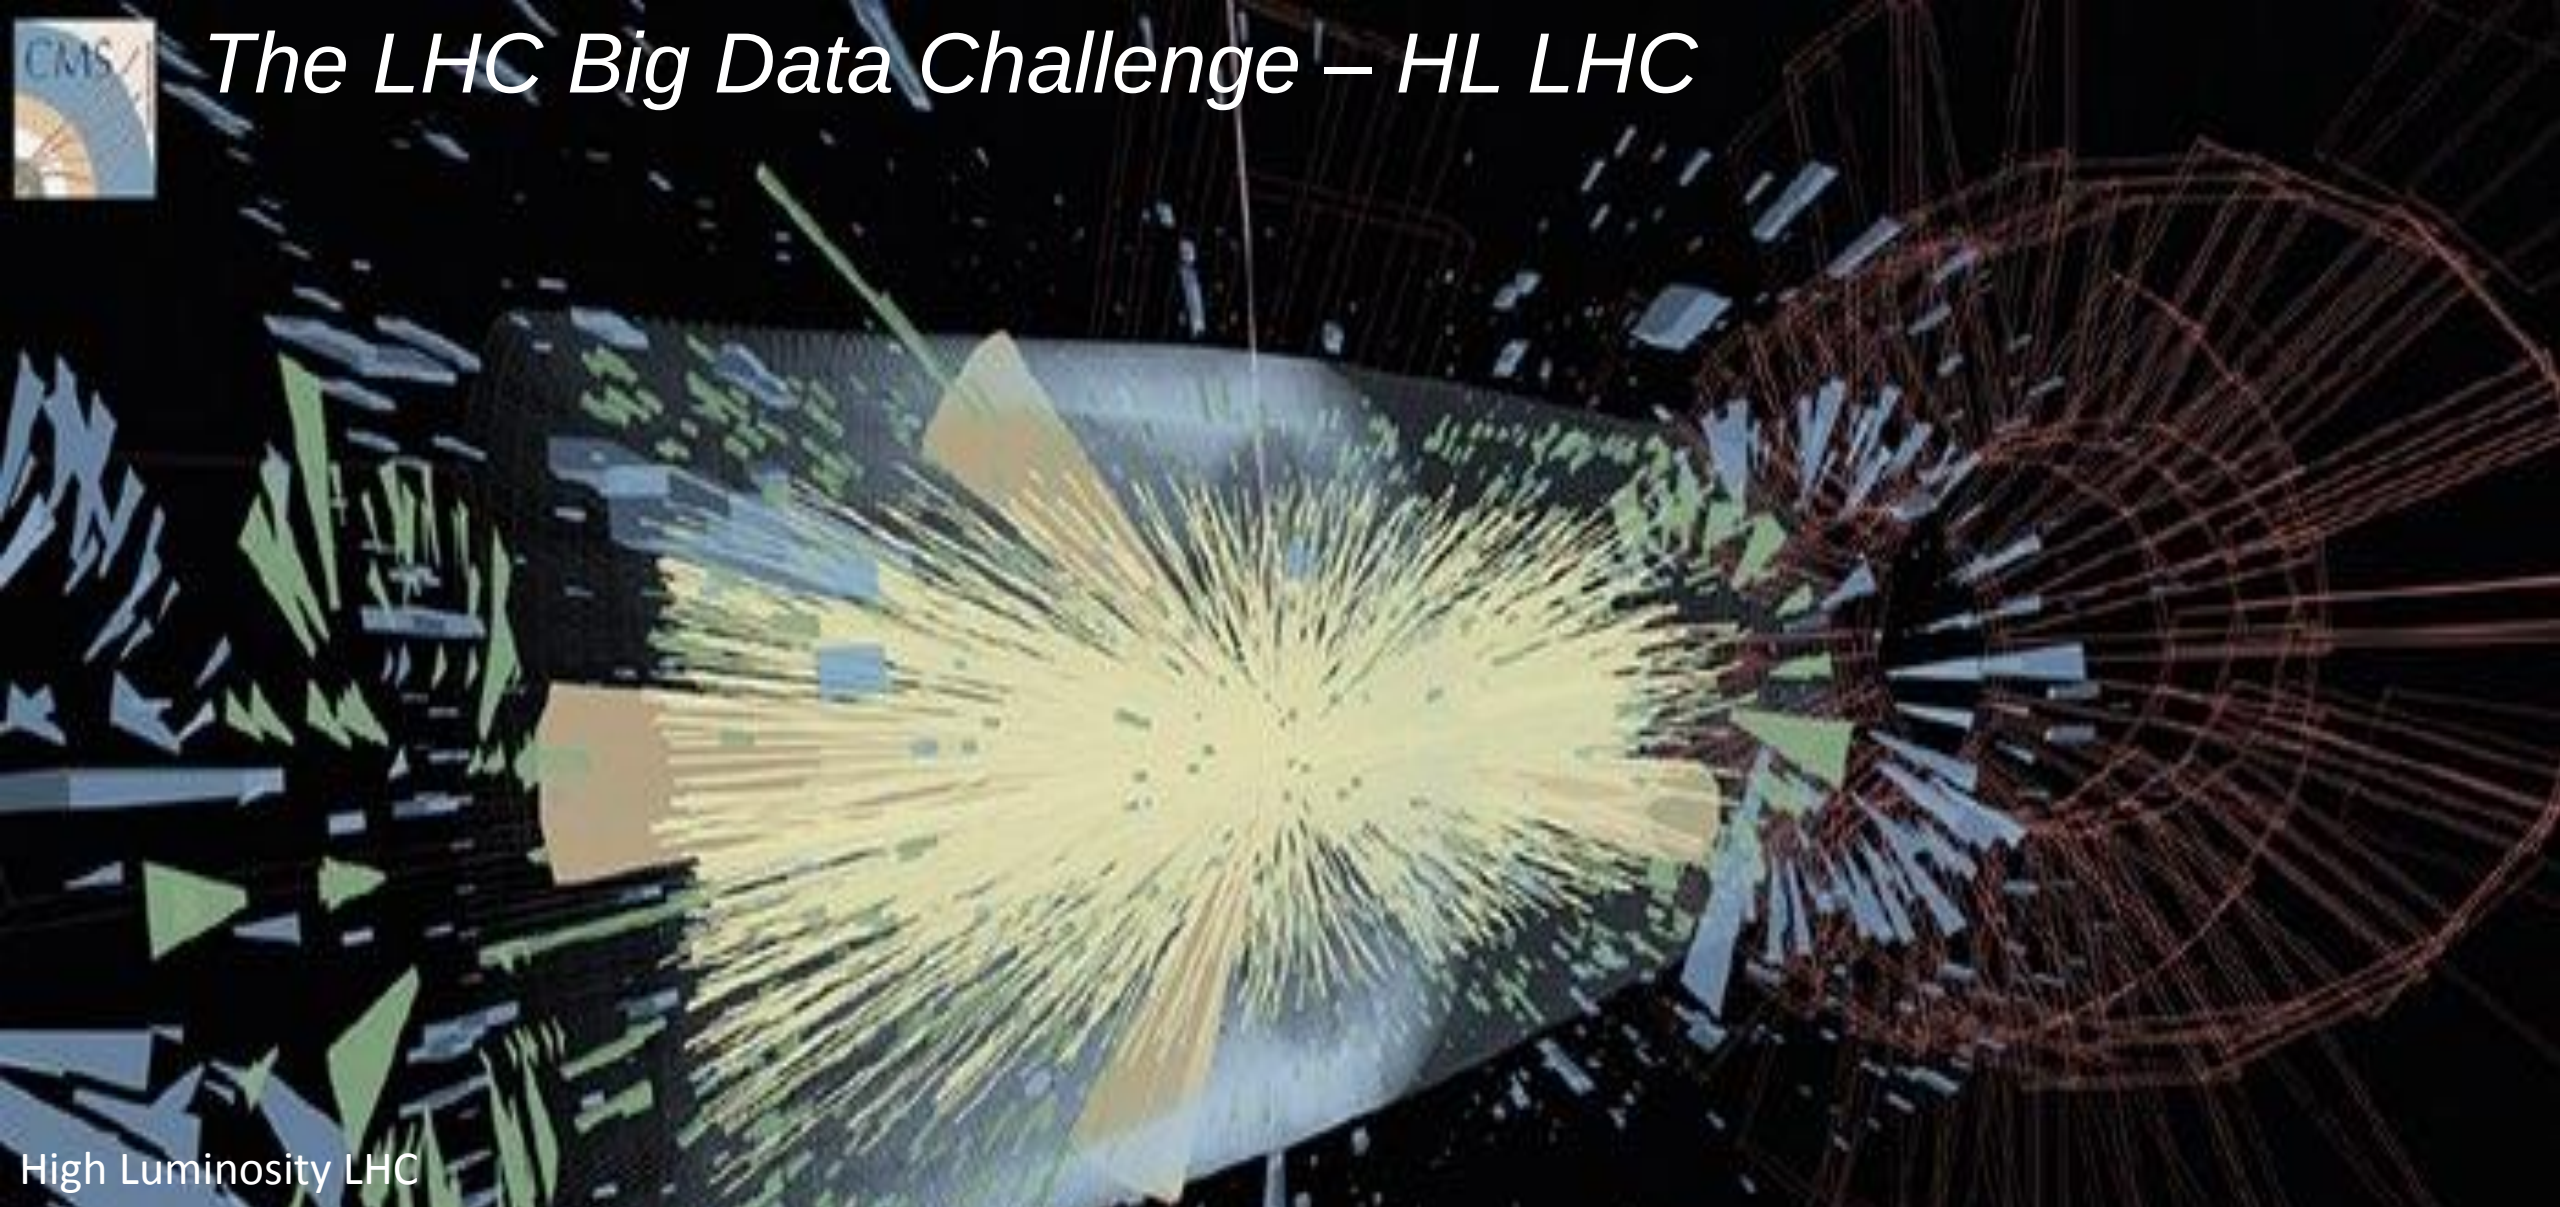

High Luminosity LHC

- 200 collisions per event vs 40 today. Need to disentangle 200 collisions happening at once.
- Event complexity grows non linearly
- A HL-LHC run would need to store about 900 Pb of derived data. A data deluge!
- Even taking into account HW progress (storage & processing), we are off by a factor of 10, projecting to 2027

- ✓ *Physics always drives CERN Innovation but we need*
  - ✓ *Accelerators, Detectors & Computing to underpin the Physics exploration*
- ✓ *R&D is happening now, in many domains*
- ✓ *In Computing*
  - *Data Analytics*
  - *Machine/Deep Learning at detector output & many other areas*
  - *Quantum Computing*
- ✓ *Some of the promising areas to address our challenge*

Technology Focus

Fast inference - HLS4ML Software

Tool to deploy Neural Networks to FPGAs

- reads as input models trained on standard Deep Learning libraries
- Uses HLS libraries to deliver a firmware implementation of a given network on FPGA
- Could also be used to design AI-specific ASICs for future experiments

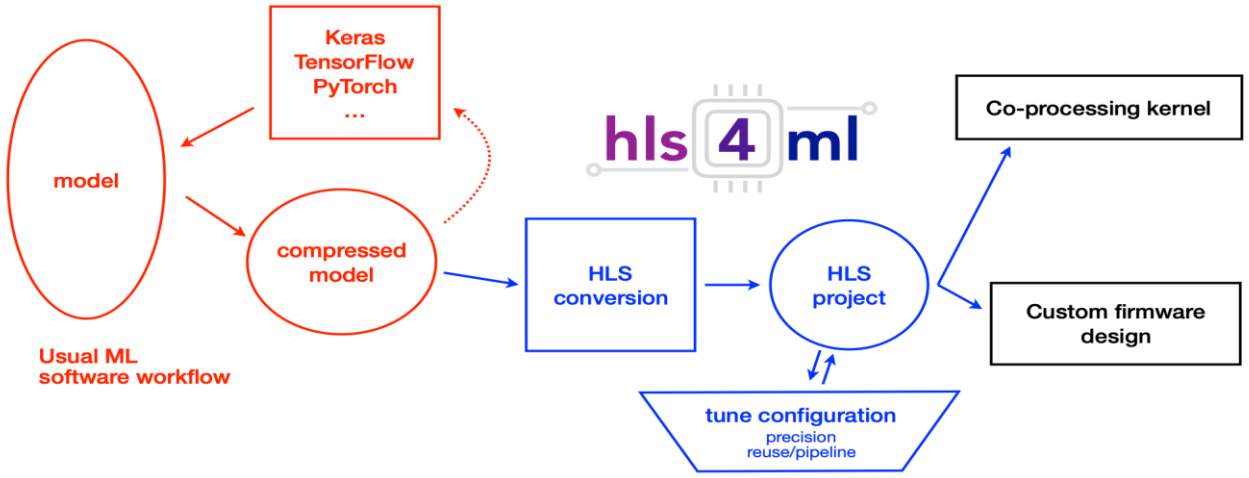

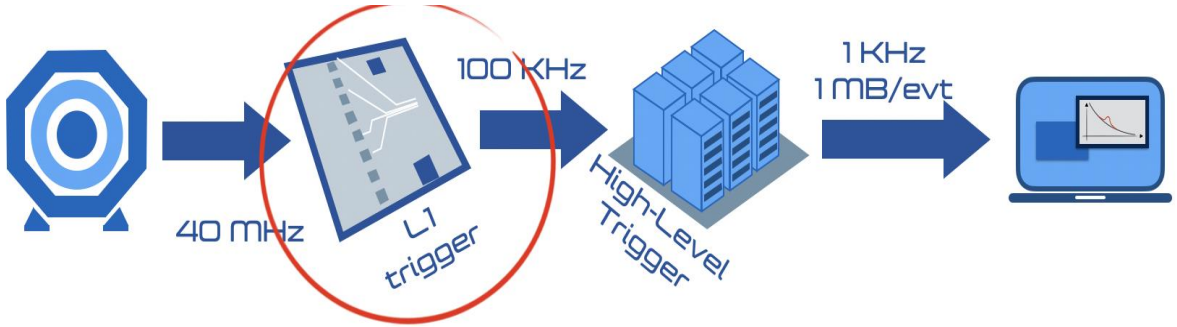


~~Zenuity~~ Zenseacat (Volvo Cars) teams up with CERN on fast machine learning using FPGAs .
PFT 701
PHH 388

Overview: Opportunities for adoption and joint R&D in ML/DL

Proven CERN capability

Use case specific

Fast ML

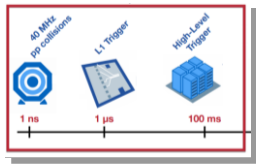

Ultra-fast on-edge inference under strict latency constraints

Anomaly detection

Object identification, classification, anomaly detection in big and noisy data sets

Industrial controls

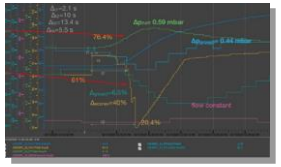

Machine efficiency and predictive maintenance with industrial control systems

Distributed computing

Optimization of distributed computing, storage, and networks; fast I/O for large files

Cross use case

Large scale, science grade data analytics and visualization

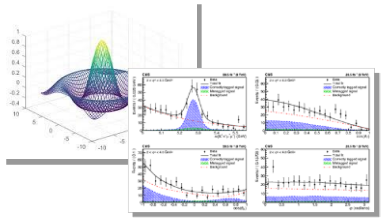

- Optimization and evaluation for science-grade precision of large data sets using advanced data analytics
- Data visualization, interactive plotting (e.g., statistical visualizations, uncertainties, distributions), model visualization
- Large-scale, quality-controlled CERN data as testbed/ benchmark (e.g., single data set with 100m examples, >1TB)

In development, opportunity for joint R&D

Simulation

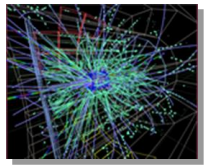

Simulation and reconstruction with generative DL for efficient computation

Graphs

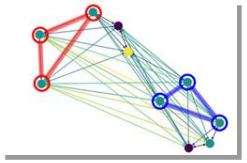

Exploring Graph NNs for high-multiplicity problems with non-linear distances

Machine design

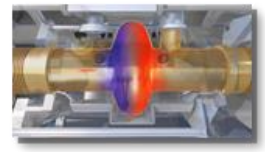

Determining optimal machine design and component configuration

ML in Robotics

Remote maintenance and safety with autonomous robots and computer vision

Quantum ML

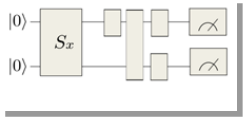

Research quantum algorithms to solve pattern recognition, classification and generation problems

Computing parallelization

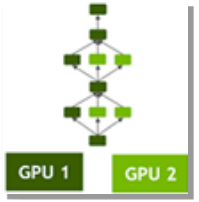

Training and optimization of complex NNs on parallelized GPU infrastructure

Key technology: Robust Big Data Analysis Framework

ROOT
Data Analysis Framework

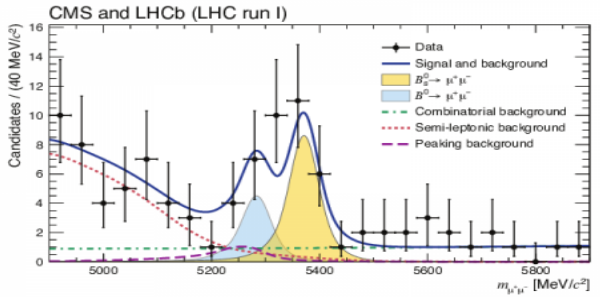

what

ROOT / TMVA is a modular big data software framework, providing the functionalities needed to deal with big data statistical analysis, visualisation and storage. It is mainly written in C++ but integrated with other languages such as Python and R. Integrated machine learning environment (bindings for Python is provided).

apps

Good for analysis of extremy large sets of structured data.
Used in industry, physics, biology, finance and insurance fraud analysis.

Possible application in processing and analysis of large medical datasets, for example genomics data, EEG/ECG data, biosensor data.

tech specs

- Rectangular cut optimisation
- Projective likelihood estimation
- Multidimensional estimations
- Linear discriminant analysis
- Function discriminant analysis
- Boosted/bagged decision trees
- Predictive learning
- Support Vector Machine
- Neural Networks

added value

- More than 1 Exabyte (10^9 Gb) of data are stored in ROOT file format.
- → **Proven capability to handle and analuse very large datasets**
- Fits and parameters' estimations for discoveries (e.g. the Higgs)
  → **Proven capability for extreme analytic applications**
- Thousands of ROOT plots in scientific publications
  → **Capable of making advanced graphs and visualizations**

Collaboration with WUR to support national banks and regulators to detect trading anomalies in financial markets

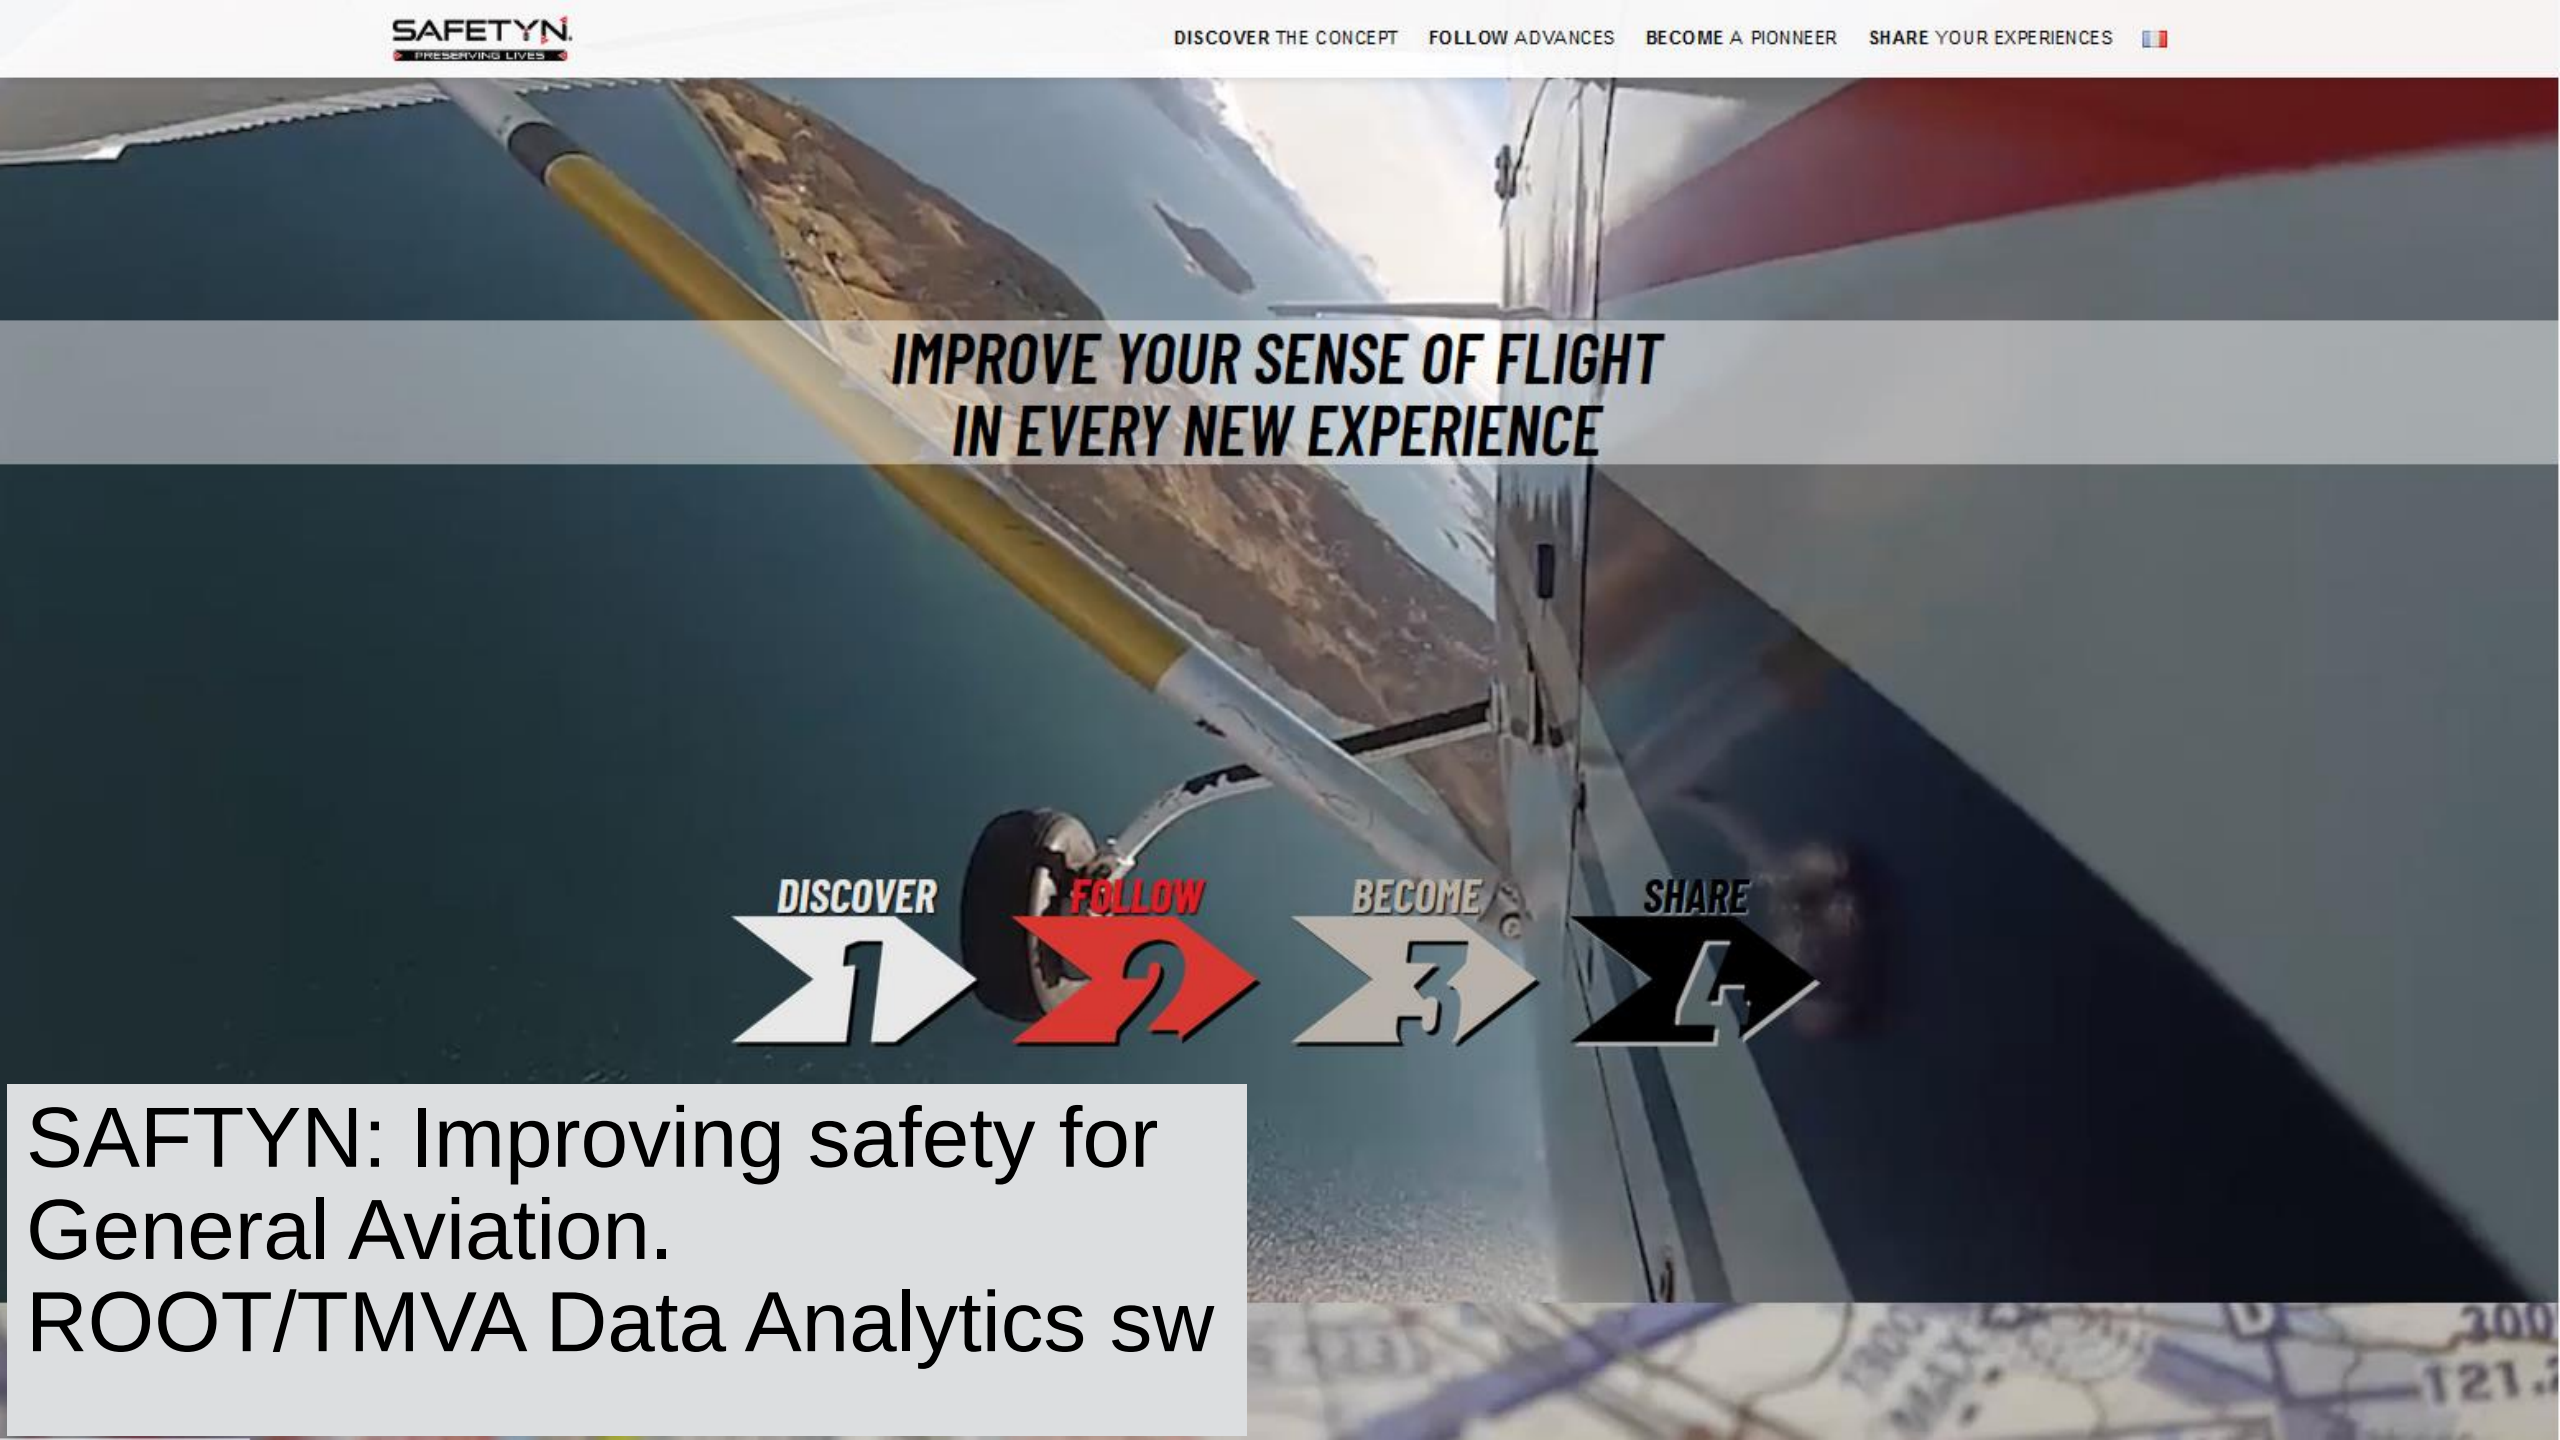

SAFTYN: Improving safety for General Aviation. ROOT/TMVA Data Analytics sw

Key technology: highly scalable cloud storage engine

what

CERN needs to store vast amounts of data. Therefore, CERN developed a reliable cloud storage engine where all the LHC data and more, is stored. This multi storage cloud software platform (EOS) is characterized by very low latency, high flexible and high scalability.

apps

- Data centers with wide range of quality of service requirements.
- Cloud service providers.
- Research infrastructures generating large data quantities.

tech specs

- Open Source Software with long term support at CERN.
- Low latency. Very fast simple namespace with no database dependency.
- Unlimited scalability. Tunable file layout per directory.
- High distribution of data across nodes makes EOS highly resilient to hardware failures.
- Multi protocol support. Main access via xroot, but also http, fuse, gridftp.

added value

- On demand quality of service.
  → **Allows to quickly adapt and meet changing performance and reliability requirements.**
- Highly scalable, multi protocol.
  → **From 2 Pb to 256 Pb in 6 years. No inherent limit.**
- Cost effective, reliable data storage.
  → **Based on commodity hard disks. Just a Bunch Of Disks layout.**

Key technology: Scalable, reliable, high volume sw distribution service CernVM-FS

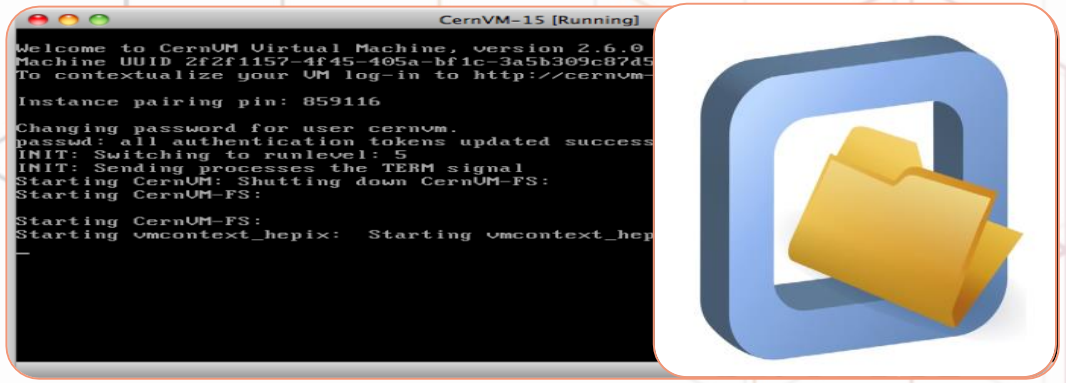

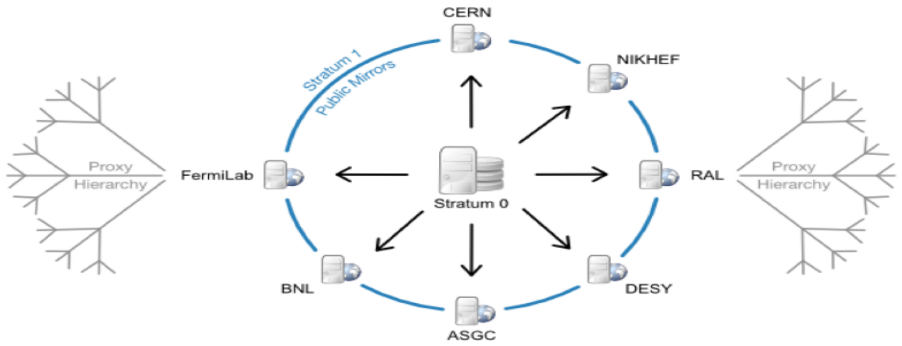

what

High Energy Physics (HEP) collaborations need to frequently and reliably deploy a lot of software on the worldwide-distributed computing infrastructure used to run data processing applications. Such service to HEP community is build on CernVM-FS OSS.

apps

- Reliable, large number of (relatively) small file distribution across the globe.
- Large scale computing and big data processing.
- Research infrastructures generating large data quantities.

tech specs

- Open Source Software with long term support at CERN.
- Implemented as a POSIX read-only file system in user space (a FUSE module)
- Outgoing HTTP connections only, avoiding most of the firewall issues of other network file systems.
- High volume file distribution capability. Over 100 Mi files and directories of LHC experiment software
- Data transfers on demand.
- Data integrity secured by cryptographic hashes.

added value

- Web-based truly global, scalable, versioning distributed file system
  → **Proven technology used for all LHC (and other) software distribution.**
- Highly scalable, low latency. Capability to work offline.
  → **No more issues with different software versions between teams.**
- Uses standard technologies
  → **Can be used everywhere, grid, cloud, local clusters..**
- Low maintenance, large community behind, committed support
  → **Mature software with large base and active development team.**

ROCHE is using CernVM-FS for application and library distribution worldwide. Strong interest in the financial services sector.

Key competence: Designing Data Analytics Infrastructure

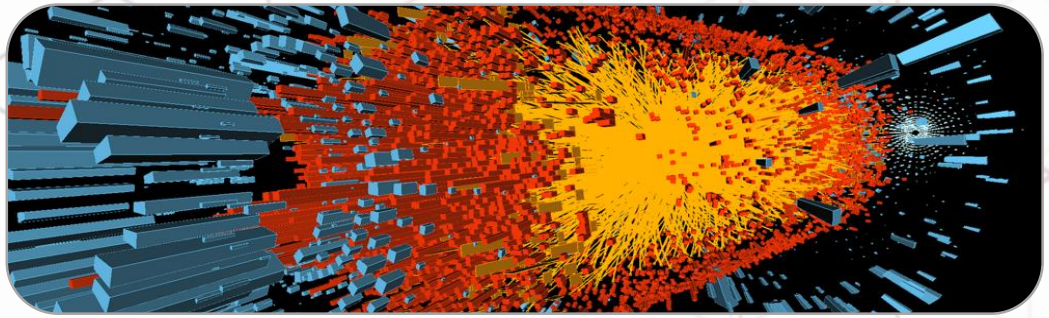

what

In order to process and analyze the vast amounts of data generated by the experiments at CERN, a data infrastructure was designed for distributed analytics. This infrastructure is made of various layers and allows 1000 clients to access the data for analysis, handling >5 million data transaction per day.

tech specs

Unique knowhow in structuring big data sets for efficient analysis.
Components used for big data and related analytics:

- User Interface: Notebooks, SWAN (developed by CERN)
- Data analysis: ROOT / TMVA (developed by CERN)
- Uses EOS, CernVM-FS, FTS (reliable large FT)
- Apache Hadoop clusters with YARN and HDFS (also HBase, Impala, Hive,…)
- Apache Spark for analytics and Apache Kafka for streaming

apps

- Any big data / data analytics application, especially those with
  - Distributed data (multiple data centers)
  - Users with different analysis needs and applications
  - Huge volumes of data stored (data lakes)

added value

- CERN has proven capability to handle huge amounts of data
  → **Know how in setup of 'big data' analytics infrastructure**
- Experience with data analysis across large complex datasets
  → **Know how on deployment at global distributed systems**
- Infrastructure build with different open source analytics components
  → **In depth knowledge of all state-of-the-art open source analytics tools**

Key competence: Open continuous process controls

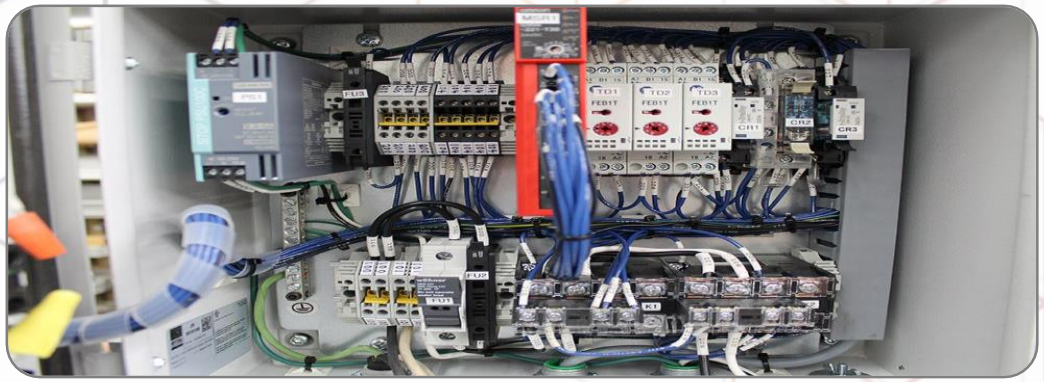

what

Expertise in designing and implementing strategies for highly complex industrial controls environments incorporating multi-vendor devices. Automation of the controls processes in an open, vendor agnostic environment.

apps

- Large, inhomogeneous industrial controls environments.
- Complex applications, mixing PLCs, OPC-UA servers, FEC, virtual and other devices.
- Any kind of Industrial controls applications form Cryogenics, HVAC to Interlocks.

tech specs

- Open architecture for distributed control and automation, applying object oriented techniques to process control.
- Coherent approach & integration for all three controls layer.
- Standards based approach. ISA-88 / IEC-61512 / IEC-61499
- Use of Formal Methods for PCL code verification
- Modeling & Simulation techniques. Virtual commissioning. Digital twin.
- Machine learning and analytics for predictive maintenance of controls systems.

added value

- Broad experience accumulated over decades with complex, wide range multi-vendor, multi-devices controls environments.
  → **Proven capability to design processes and tools to orchestrate highly sophisticated controls environments.**
- IIOT scenarios, virtual PLCs, software re-configurable processes and other technologies for the future CERN projects.
  → **Access to state of the art technologies and developments**
- Collaboration with one of the world's leading research institutes
  → **Possibility of using CERN labels for your branding and marketing**

Key technology: Unified industrial control system

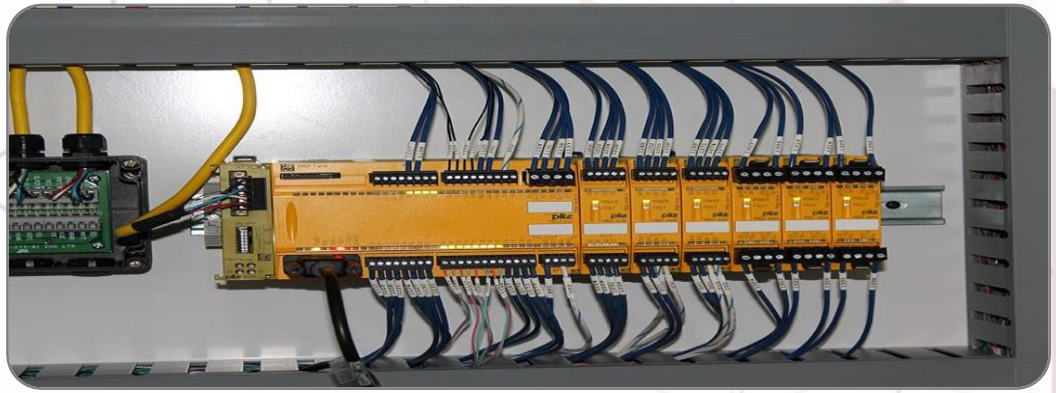

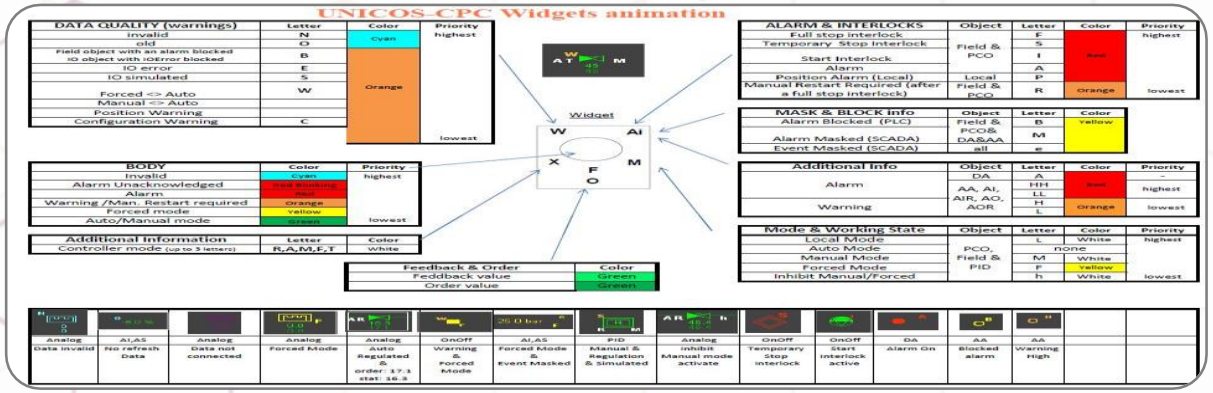

what

UNICOS (UNified Industrial Control System) is a CERN-made framework to develop industrial control applications. It is a standards based, vendor independent framework for controls applications. Provides Supervision and Controls layers for highly complex heterogeneous device controls applications

apps

- Any type of industrial control application, form Vacuum to Cryogenics to electrical distribution to magnet controls.
- Multi vendor, multi device complex controls environments.
- Virtual commissioning and controls simulation.

tech specs

- Tools to build applications based on Continuous Process Control principles.
- Object oriented approach for Modularity, Complexity reduction and code Re-use.
- Support for the full engineering lifecycle. Decomposition, Specification, Automatic code generation, Application specifics, Deployment, Reverse engineering.
- Uses WinCC OA (ETM) for the SCADA part.

added value

- Comprehensive. Supports the development, commissioning, operation, diagnostics and maintenance phases.
  → **Full lifecycle support allowing changing topology & processes**
- Simple drag & drop interface to create controls applications
  → **Empowers operators & reduces development time, while preserving coherent controls environment**
- Open framework with APIs to many peripheral systems.
  → **Avoids vendor locking and allows integration with other infrastructure systems (alarms, archiving, asset management etc)**

Key technology: Control and monitoring platform

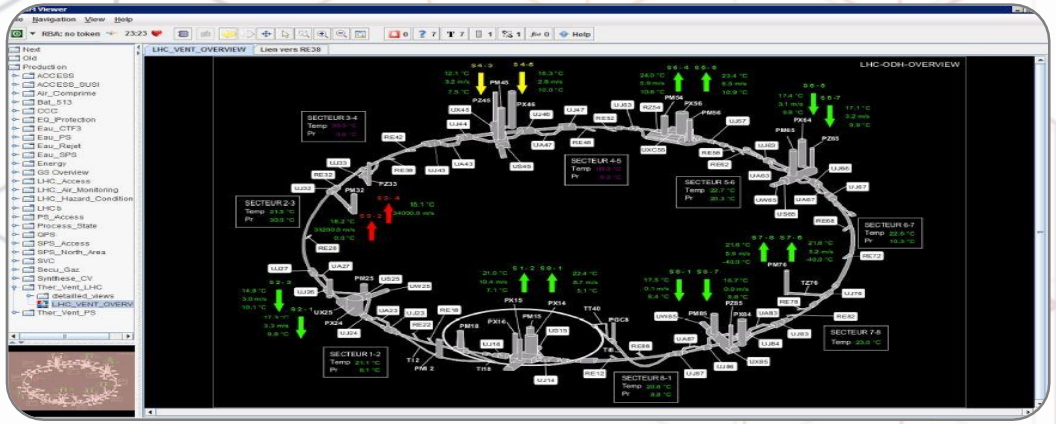

what

A modular Java framework called C2MON for large-scale industrial monitoring and control solutions, with all core functionalities adaptable to a wide variety of monitoring systems. A highly scalable heterogeneous platform for many Big Data and IoT scenarios.

apps

Industrial control applications with complex monitoring requirements.

- Large and complex control & monitoring environments.
- Grid operations, oil & gas industry, manufacturing industry.
- Energy production monitoring.
- Healthcare applications like patient monitoring.

tech specs

- Three tier modular design, allowing custom extensions. Horizontal scalability at all layers
- Sophisticated filtering & alarm mechanism for setting meaningful alerts.
- Ability to handle high throughput and millions of different sensors.
- Supports arbitrary sensor value objects (JSON serialisation in planning)
- Load-balanced server clustering capability.
- Made to handle sudden and unforeseen machine breakdowns
- Integrated history browsing for industrial dashboards
- Modern HTML5 web interface for easy navigation
- High reliability applications (>99.9% uptime)

added value

- Open Source Software
  → **No vendor locking. Highly extendable to suit specific environment needs. Community support.**
- Supports heterogeneous vendor environments. Can act as technology bridge for IoT scenarios
  → **Enables communication between non compatible devices**
- Built-in rule engine to express complex data relations.
  → **Reduced complexity when building monitoring applications**

SecurAxis: Real Time Analysis, Reporting and Localization with Smart Acoustic Sensors C2MON monitoring software

CHECK YOUR AIR QUALITY

PlanetWatch

About us ▾ Projects ▾ Buy ▾ Explorer Blog Support **Become a planetwatcher**

Look after the air you breathe!

Earn rewards and help us build a global air quality monitoring network to detect pollution hotspots and protect everyone's health.

Join our global community
Become a PlanetWatcher

PlanetWatch: Building a global network of sensors to monitor air quality
C2MON monitoring software

Key technology: Controls middleware for vital services

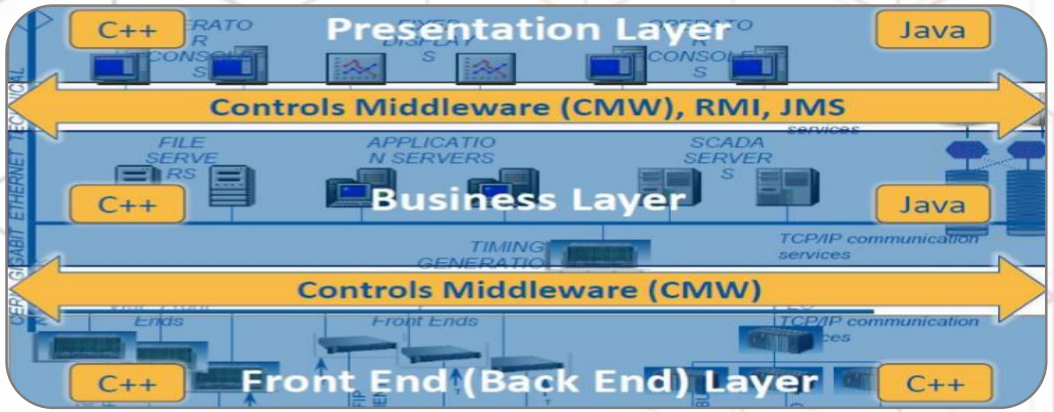

what

Controls Middleware (CMW) is software that controls CERN's accelerator complex. Its function is to provide vital services and interoperability between heterogeneous complex distributed systems whilst hiding the complexity of such systems from the users.

apps

- Manufacturing Execution Systems (MES) for Industry.
- Factory automation applications.
- Controlling large complex heterogeneous environments like for example large energy plants, manufacturing or production sites.

tech specs

- Software layer between the Operating System (OS) and the applications.
- Hides complexity & heterogeneity of distributed system.
- Handles issues related to OS, network protocols & hardware platforms.
- Support for Naming, Location, Service discovery, Protocol handling, Synchronisation, Concurrency, Failover, Scalability, Role based access control and Authentication.
- Uses ZeroMQ for distributed messaging.

added value

- Low latency middleware services. Deployed at the CERN accelerator complex since 2002.
  → **4'000 CMW Servers, 85'000 devices, >2 million IO points. Proven robust software for highly demanding applications.**
- Based on ridable and generic RDA3 (Remote Device Access v.3)
  → **Can be used in any environment. It is not CERN specific.**
- Vital software and part of CERN's accelerator complex infrastructure.
  → **Long term future and support assured.**

CERN CMW (controls middleware) for Industrial Manufacturing Execution System

Key competence: Mechatronics and custom robotic solutions

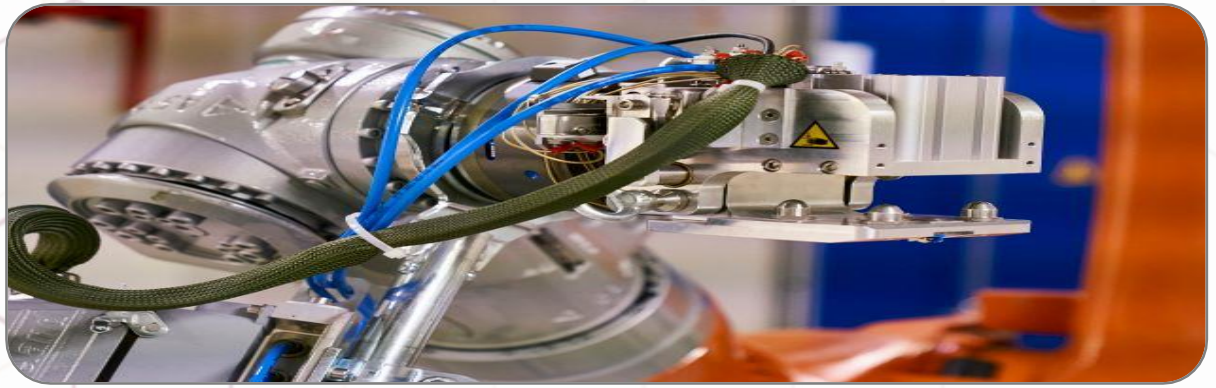

what

Expertise in designing and implementing robotic solutions and use of such solutions for monitoring, remote inspections and tele-manipulation to execute complex tasks in harsh environments reducing human exposure to ionizing radiation.

apps

- Development of specific solutions for demanding environments with real time and strong safety requirements.
- Tele manipulation for delicate and high risk operations.

tech specs

- Proficiency in software development and integration of third party components.
- Development of augmented reality techniques for operator training.
- Experience in real time applications and support for Safety Integrity Level (SIL).
- Competence to handle complex mechanical design and integration requirements.

added value

- Many highly complex interventions performed in CERNs experimental areas.
  → Proven capability to train man & machine to perform in harsh environments
- Design of robotic solutions using standard components to fit the purpose and environment.
  → Affordable, easily maintained robotics for highly demanding applications. Continuous research and evolution.

Key technology: Modular and flexible robotic platform

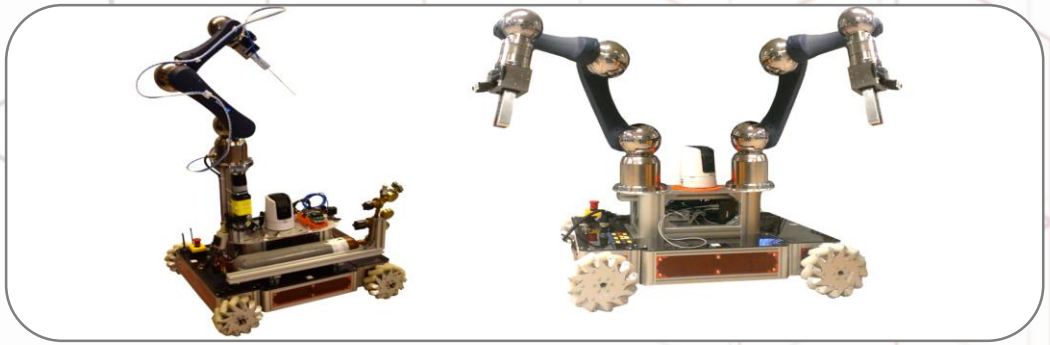

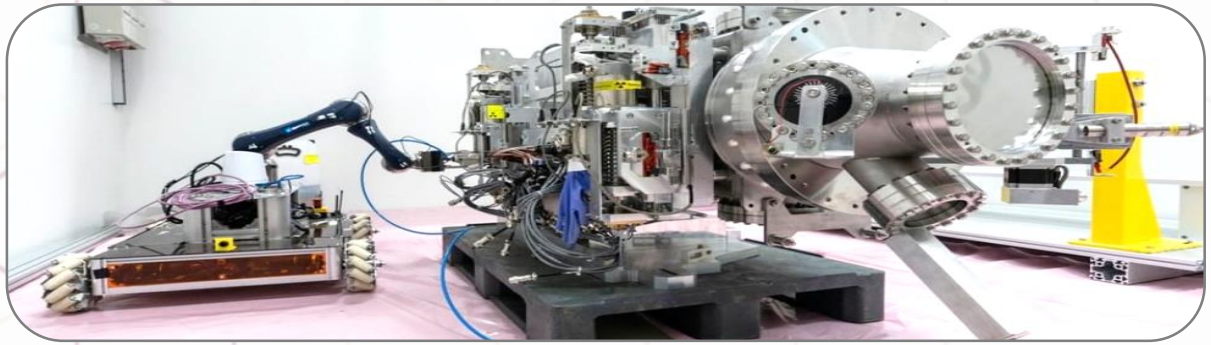

what

CERNBot is a modular and flexible robotic platform capable of delicate interventions in presence of ionization radiation. Primarily aimed for indoors interventions. Used across the CERN accelerator complex for critical interventions.

apps

- Autonomous and tele-operated interventions in presence of ionizing radiation.
- Complex interventions involving coordination of two robotics arms.
- Search and rescue tasks.

tech specs

- Two robotics arms that can be mounted on a custom made elevator platform for interventions of up to 3m in height.
- Autonomous operation.
- Payload up to 250 kg.
- Very stable platform.
- Seamless integration of different arms controls due to deep software integration.

added value

- Flexible platform for a wide variety of use cases.
  → **Adaptable topology to meet the task requirements**
- Solution based on standard industry components with predicable upgrade path.
  → **Low cost comparing to commercial platforms with similar capabilities**

Key technology: Autonomous robotic inspection platform

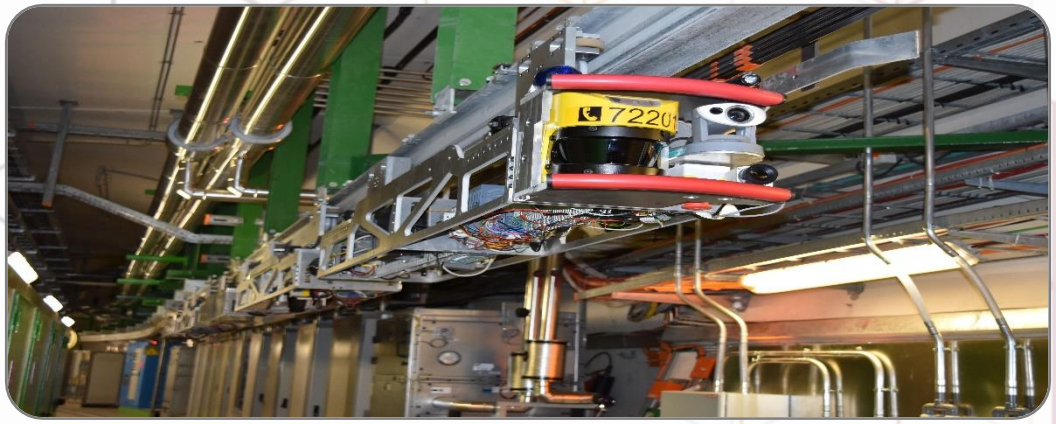

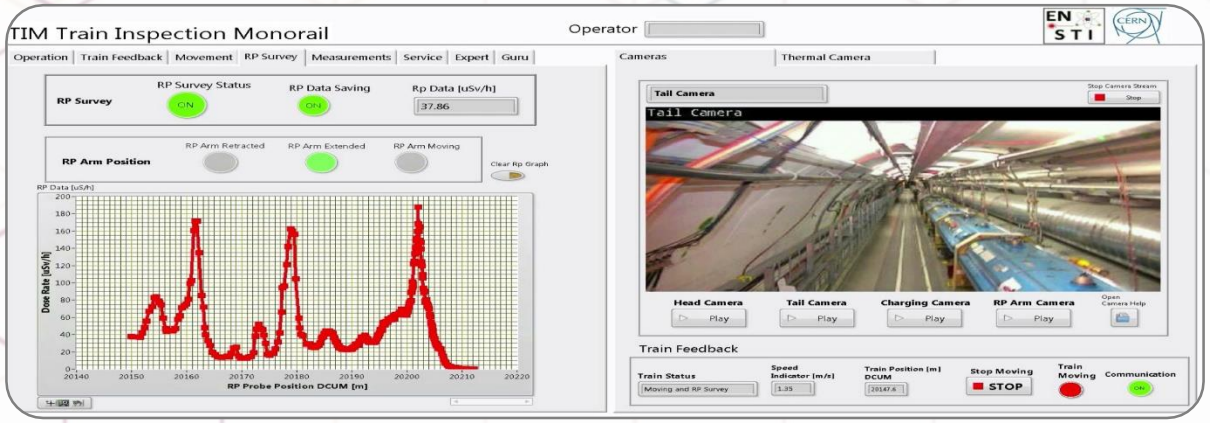

what

Train Inspection Monorail (TIM) is a unique modular, extensible, robotic platform capable to accomplish autonomously a variety of different missions. Including regular inspections, safety tasks, monitoring, complex interventions and others.

apps

- Industrial asset inspection and safety
- Large infrastructure tunnels such as those of utility companies
- Safety & monitoring of tunnels involving people like train or road tunnels
- Photogrammetry applications

tech specs

- Autonomous vehicle control
- Modular design
- Automated visual inspection
- Different sensors packages
- Handling robotics on board

added value

- Many hours of successful interventions in the LHC tunnel
  → **Proven track record and live demonstration of TIM in action**
- Four TIMs deployed in LHC with more to come
  → **Very long term support and development of TIM assured.**
- Highly modular and autonomous.
  → **Wagons providing dedicated functionality can be added/adapted. Clever software orchestrates autonomous operations**

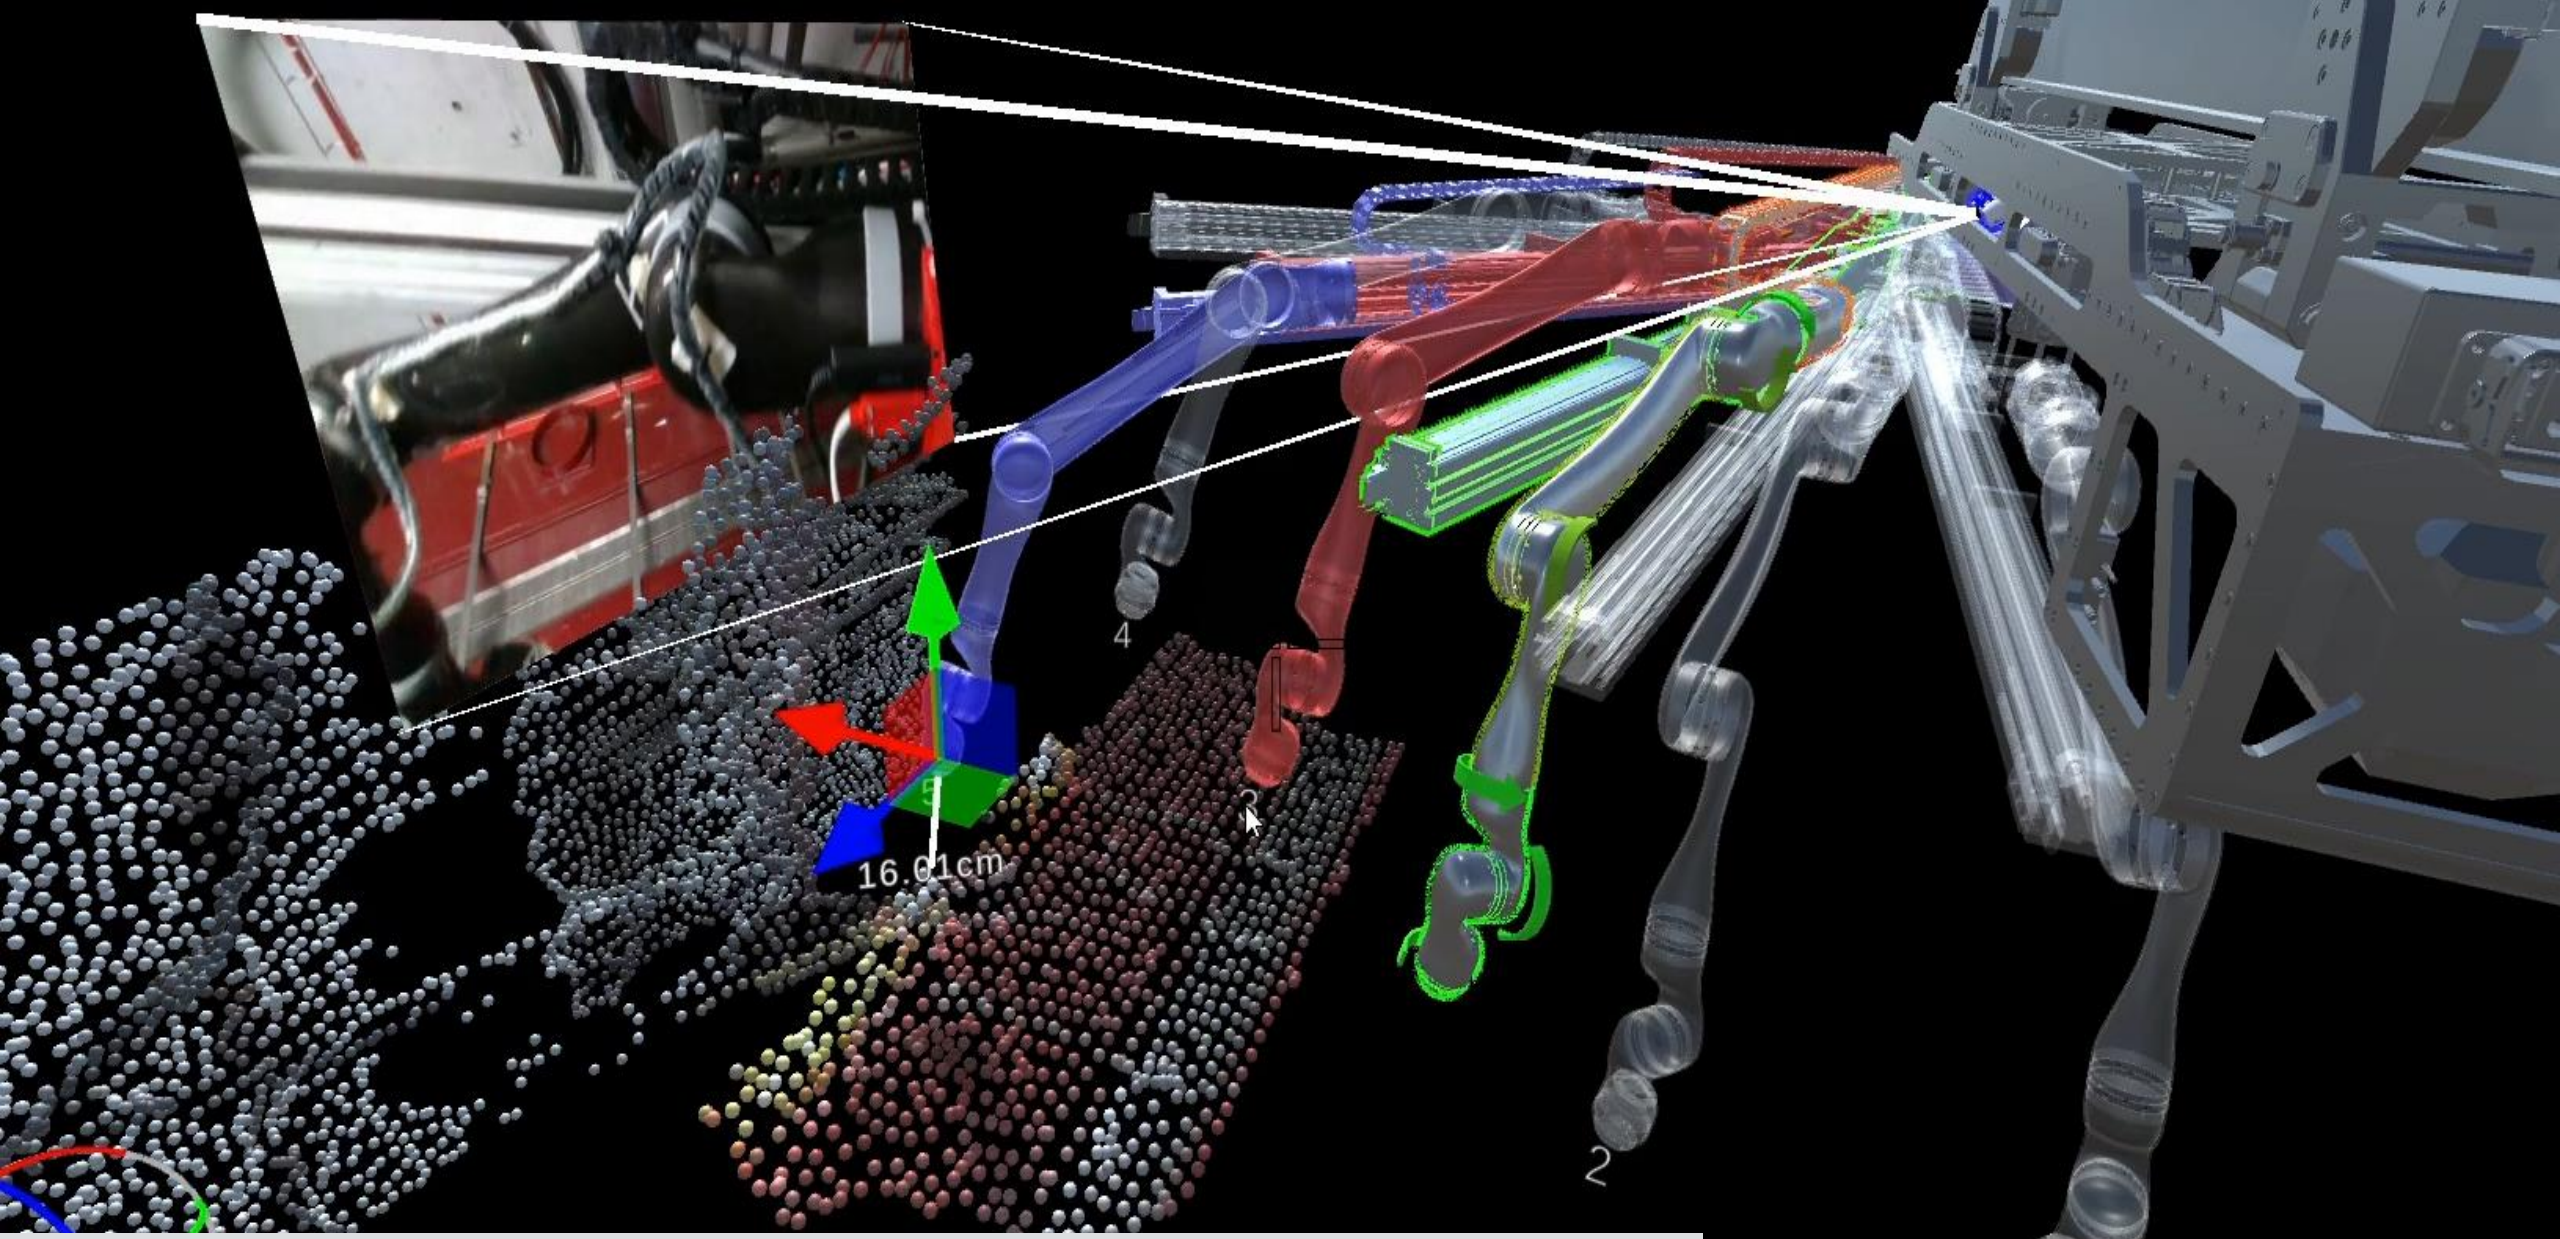

Mixed reality for Robotic teleoperation
BLM project

Technology Focus

Simulation Software

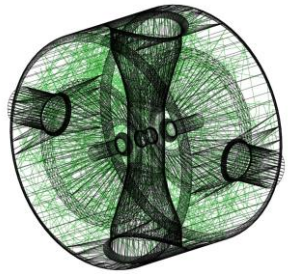

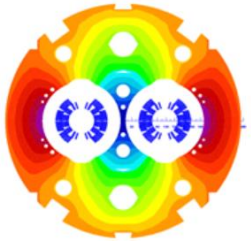

- **Molflow+** : MC simulation SW to calculate the steady-state pressure in an arbitrarily complex geometry under ultra-high vacuum conditions.
- **ROXIE**: Electromagnetic simulation and optimization of magnets. From concept to quench post-mortem analysis
- **Geant4**: is a toolkit for the simulation of the passage of particles through matter. Application areas: High energy, nuclear and accelerator physics, Medical and Space science.

Key technology: Simulation software for particle transport

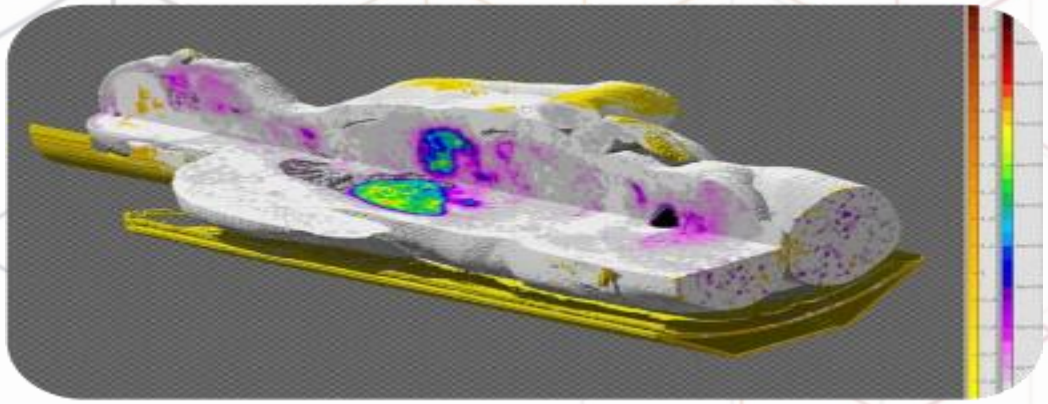

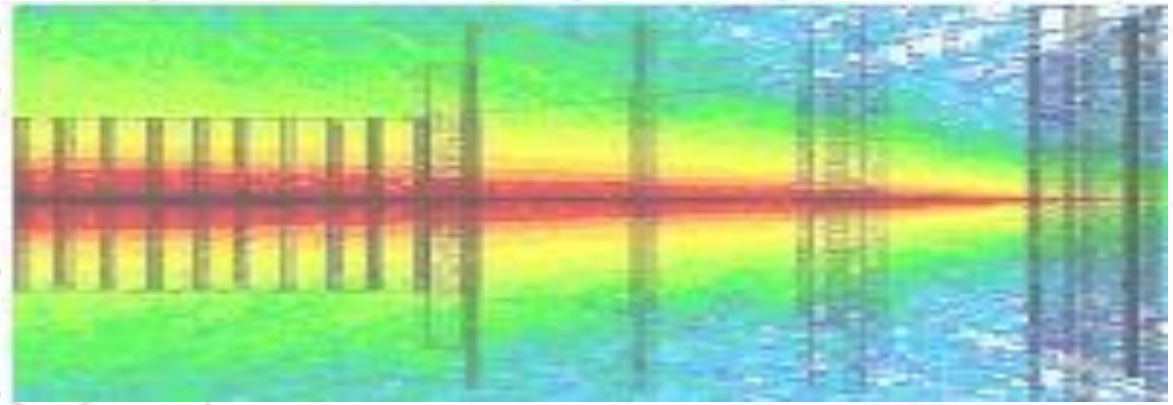

what

FLUKA (Fluctuating Cascade) is a general purpose tool for calculations of particle transport and interactions with matter. FLUKA can simulate the interaction and propagation in matter of about 60 different particles with high accuracy. FLUKA can handle very complex geometries and yields very accurate simulations.

apps

- Ionizing radiation shielding.
- Radiotherapy and treatment planning systems.
- Detector design.
- Radiation protection and dosimetry.
- Accelerator driven systems.

tech specs

- High accuracy simulation of the interaction and propagation of about 60 different particles, including photons and electrons from 1 keV to thousands of TeV, neutrinos, muons of any energy, hadrons of energies up to 20 TeV.
- This includes all corresponding antiparticles, neutrons down to thermal energies and heavy ions. It can also transport polarised photons (e.g. synchrotron radiation) and optical photons.
- FLUKA can handle very complex geometries, via an improved version of the well-known Combinatorial Geometry (CG) package. The FLUKA CG has been designed to track correctly charged particles (even in the presence of magnetic or electric fields)

added value

- Very high accuracy simulations of a large number of particles.
  → **Suitable for complex environments that need regulatory approval.**
- For most applications, no programming is required from the user.
  → **FLAIR distributed with FLUKA is a friendly interface that greatly enhances user experience.**
- Large user base and active user community and forum.
  → **widely available support, documentation and training courses.**
- High reliability. Proven technology with a wide use base in Academia but also Industry.

Key competence: Software for empowering collaboration

what

Large global scientific collaborations are the foundation of CERN achievements. Scientists need to exchange documents, publications, photos, videos and more. They need to meet, physically or virtually to debate and exchange ideas. CERN has expertise in building software tools empowering such collaborations.

apps

- Digital repositories for large organizations.
- Provision of custom made digital repository services.
- Solutions for empowering sector specific collaborations.
- Large scale event management.

tech specs

- Scalable software supporting large volumes of data.
- Long term data preservation is a key requirement.
- Open source software avoiding vendor lock-in and securing long term future for the software stack.
- Compliance with standards for web content interoperability.

added value

- Large global collaborations are driving CERN's successes.
  → **Expertise put to test every day by the needs of our scientists.**
- Wide range of applications adaptable to the specific use case.
  → **CERN expertise is used to provide support to a wide variety of projects, from museum support to publication data storage.**
- Open collaborations are CERN's driving force.
  → **Joining a wide community facing similar challenges, benefiting from each other's experiences.**

Key technology: Large scale digital repositories

what

Software framework for large scale digital repositories. The CERN Document Server, CERN's institutional repository containing more than 2 million records, including research publications, audiovisual material, images, is powered by INVENIO.

apps

- Used at CERN, UN, EPFL, CalTech, INSPIRE, …
- Institutional multimedia digital content server systems.
- Specialized repositories with dedicated functionality.
- Back end for large collaboration distributed information systems.

tech specs

- Enables each content producer to establish his own visual and functional identity.
- Invenio v3 framework grants increased control on the data model and lifecycle of the content.
- Powerful search with additional options of combined meta data, citation and full text search.
- Advanced file management and organization of documents in community collections with precise access control.
- Long term preservation.

added value

- Framework architecture allowing for specific application development.
  → **Fully tested at CERN. Example the videos.cern.ch application.**
- Years of experience managing large digital libraries and publications.
  → **Invenio software has evolved due to direct feedback from our highly demanding user community.**
- Open source framework benefiting from a large community.
  → **CERN's driving role ensures the framework's long term future.**

spin-off

TIND

HOME | TIND ILS | TIND RDM | TIND DA | TIND IR | BLOG | NEWS | Get in touch

REIMAGINING LIBRARY TECHNOLOGY

TIND ILS
Management of print and electronic resources.

TIND RDM
Capture and preserve datasets of any size and format.

TIND DA
Digital archive for digitized and special collections.

TIND IR
Manage publications and other traditional research output.

TIND, CERN Spin-off based in Norway selling digital library services across the world

SUPREME COURT LIBRARY QUEENSLAND

Supreme Court Library Queensland will migrate to a combined TIND ILS and DA instance, moving away from Millennium.

Read press release

Key technology: Software service for open science data deposits

what

Zenodo is a general-purpose open-access repository developed and operated by CERN. It allows researchers to deposit data sets, research software, reports, and any other research related digital artifacts. Zenodo code is itself open source, and is built on the foundation of the Invenio digital library.

apps

- Technology demonstrator for INVENIO based Software as Service.
- Solid platform for communities to create their own digital repositories.
- Provides state of the art digital repositories otherwise unaffordable for small projects and labs.

tech specs

- Provides Digital Object Identifiers (DOIs) making submissions easily citable.
- Allows uploads of files up to 50 GB.
- Closed and restricted content support while the research is ongoing, ready to be openly shared at later stage.
- Powered by Invenio the CERN document server software and EOS, the LHC low latency data storage system. Hosted at CERN's computer center.
- Strong data protection and secure user access.

added value

- Funded by EU FP7 and H2020 projects (OpenAIRE and OpenAIRE2020)
  → **Multidisciplinary repository welcoming any open research data.**
- Hosted at CERN, combines robust CERN technologies.
  → **Production service managed by CERN teams.**
- Based on CERN's open source software for the benefit of Open Science.
  → **CERN's experience with digital repositories and large data set management benefits the wider scientific community.**

Orvium

Researchers Publishers Reviewers About Blog Join for free

One platform for **all** your publishing needs

From submitting your research to manage your Open Access peer reviewed journal

Subscribe Join for free

Orvium, aims to revolutionise scientific publishing.
Uses ZENODO as its repository & back end

Conclusions

- ✓ CERN offers a very wide range of digital solutions to solve our challenges.
- ✓ These solutions are transferable to Industry to address wider societal challenges.
- ✓ Some are more readily transferable while others require adaptation via closer collaboration.
- ✓ We have the expertise, You have the domain knowledge and related domain expertise.
- ✓ Knowledge Transfer has the means to put the two together and create fruitful collaborations.

Thank you!

Nick.Ziogas@cern.ch

cern.ch/kt

www.cern.ch

Key competence: Fast neural network inference in FPGAs

hls4ml

what

CERN needs ultra fast machine learning interference (execution in µsec), requiring compact code for FPGAs. A companion compiler package for this work is developed based on High-Level Synthesis (HLS) called **hls4ml** to build machine learning models in FPGAs, allowing for fast prototyping and shorter time to results.

apps

For any ultra fast ML challenge including

- Fast triggering and classification of events
- Search engines (FPGAs already used in BING)
- Chatbots (FPGAs already used in SIRI)
- Automated driving

tech specs

- Automatically translating a trained neural network, specified by the model's architecture, weights, and biases, into HLS code
- Reads as input models trained with standard DL libraries
- Uses Xilinx HLS software (accessible to non-expert, engineers resource not common in HEP)
- Comes with implementation of common ingredients (layers, activation functions, binary NN …)
- ML code support: Keras/Tensorflow, PyTorch, scikit-learn (planned xgboost)

added value

- CERN has built a system to adapt Neural Networks to FPGA architecture
  → **hls4ml (open source) decreases drastically the firmware development time**
- Efficient Neural Network design for FPGAs
  → **Optimization of FPGA design through compression/quantization/parallelization**
- Custom Neural Networks tested in the CERN experiments
  → **Advice/consulting on ML & FPGA development from world class experts**

Key competence: Designing & Training Neural Networks

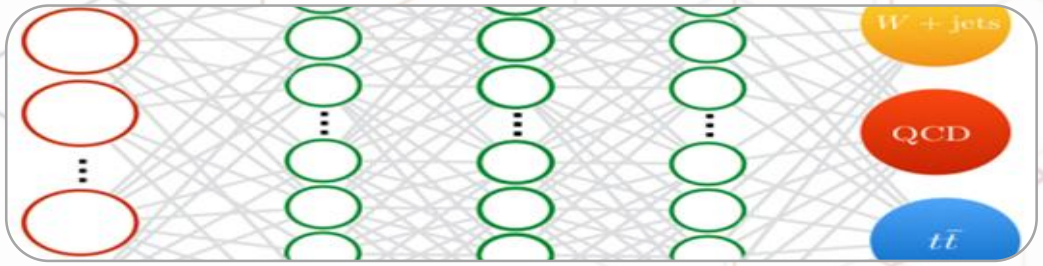

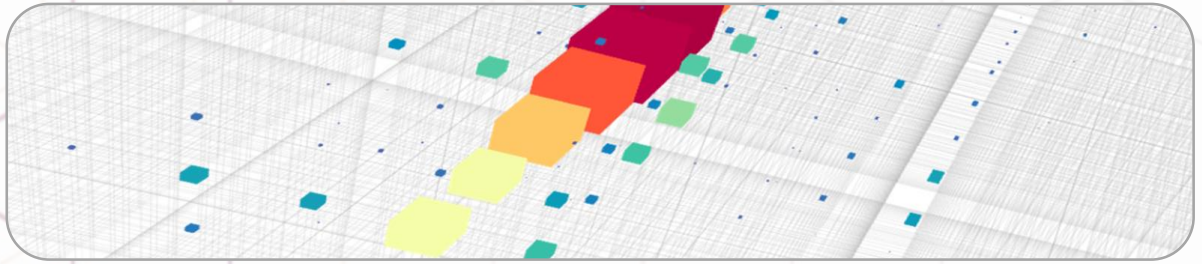

what

CERN has a long history in the design and training of neural networks in for example classification, filtering, event and particle detection, regression, clustering and anomaly detection. Most of the ML/DL codes are tailor made using C++, Phyton, TensorFlow and Keras and applied in software or hardware (FPGAs).

apps

- Creating and training custom made neural networks
- Analysis of very large datasets (both structured and unstructured)
- Very fast (<50µsec) selection and filtering of data information
- Search for exceptional data points or events (online or offline)
- Classification of events, images or datasets

tech specs

Experience with design, training and executing applications of:

- 2D and 3D Convolutional Neural Networks (CNNs)
- Deep Neural Networks (DNNs)
- Recurrent Neural Networls (RNNs)
- Graph Neural Networks (GNNs)
- Graph Convolutional Networks (GCNs)
- Generative Adversarial Networks (GANs)
- Boosted Decision Trees (BDTs)
- Variational Auto Encoders (VAEs)

added value

- CERN has built a large variety of neural networks from scratch
  → **We can design ML/DL algorithms for a wide variety of use cases**
- Hands on experience with large scale training of neural networks
  → **Help in defining and executing optimal training / learning protocols**
- CERN experts have created custom made neural networks since decades
  → **Provide independent advice on ML/DL strategy**

Key competence: Storing & managing very large data sets

what

Data storage and management of very large data sets in highly distributed, global environment. Rapidly responding and adaptable to the LHC experiments frequently changing requirements for fast, reliable and affordable storage.

apps

Any large volume data storage and management application

- Cloud Services
- Data Centers
- Large distributed data environments
- Long term data preservation policy

tech specs

- Fine tuned, adaptable data storage policies to accommodate conflicting requirements between speed, reliability and affordability.
- Expertise in very high rate, reliable, sustained data transfers
- Running a highly complex service integrating different types of storage devices (tapes, ssds, memory etc) combined to provide an infinitely scalable and highly flexible service.
- Made to measure services (storage pools) with on-demand "Quality of service"

added value

- Experience accumulated over decades of handling very large, multi-user, multiplatform-access distributed data sets.
  → **Proven capability to manage reliable 'big data' operation**
- Near future extreme data storage & management requirements feed continuous research investment
  → **Preview of future tech for even larger (x10) data sets**
- Collaboration with one of the world's leading research institutes
  → **Possibility to leverage from CERN expertise and knowhow**