A successful public-private partnership

Alberto Di Meglio
CERN openlab Head

CERN openlab

DAQ

- **Targeting Run 3 and Run 4**
  - Common requirements from all experiments, with different timelines and scales
  - LHCb, ALICE → LS2/Run 3 (2019-2023)
  - ATLAS, CMS → LS3/Run 4 (2024-2029)
- **Main areas of investigation**
  - High data rate network fabric
  - Software optimization and acceleration
  - Very large memory/NVRAM
- **Possible change of paradigm**
  - Bring offline and online closer
  - Not sequential anymore, but task-based, leveraging the availability of the technologies above
- Interest in working together on a more “holistic” view of how to do DAQ within the respect of each experiment specific needs

Networks

- Working very close with open source community
- Many challenges both for the distributed infrastructure for HL-LHC and the local infrastructure
  - Also depends on what strategy is adopted to provide more capacity
- Move towards a distributed data network of fewer, bigger data centers → how does this impact the network architecture?
- Need to modernize the local infrastructure
  - FTTO, seamless roaming between wi-fi and LTE
- Looking forward at the needs of FCC
  - IoT? Wireless connectivity at the detector sensor level?
- Many questions:
  - Looking at network end-to-end instead of service by service
  - What about global services like nameservers or metadata?
  - Experiments looking at techniques (CDN, NDN, dynamic network circuits) to integrate the network and their data management systems. Essential for computing models evolution for HL-LHC

Computing Provisioning

- Working very close with open source community
- Effort towards integrating different types of platforms behind the same operation management process (VMs, containers, bare metal)
- Looking at rack disaggregation (Rack Scale?) as a potentially interesting approach to improve flexibility and decrease TCO
- Work on deploying hybrid clouds
  - How to integrate public and private clouds? Technical aspects and business aspects
  - Spot/elastic market Vs. long-term dedicated resources market
  - Large scale tests done on spot market by the LHC Experiments (CMS, ATLAS) have shown we are very close to be able to use public cloud for the entire reprocessing flow at interesting costs

Storage and databases

- Need to have systems able to support many different use cases and QoS levels
- “Productization” is very important to allow other sites to deploy the systems efficiently and go beyond HEP
- Investigating new technologies to reduce complexity and decrease the cost/MB ratio (kinetic, MSD), balance across EHDD/CHDD/SDD as price evolves
- Understand the impact of large NVRAM pools (3D XPoint)
- Data ingestion (long term view): assess both enterprise and open source possibilities
- **Questions:**
  - Look at storage and network together, what are the implications of the change in distributed data cloud on network needs?
  - HEP approach towards storage is that it is essentially decoupled from computing. Maybe there are cases where it would be efficient to have computing and storage together

Code Modernization

- Code modernization and software optimization are needed in many different areas
- Many activities ongoing in both HEP and in collaboration with other scientific disciplines
- Still software optimization and hardware acceleration might not be sufficient in the long term
  - Looking also at new approaches like applications of Machine Learning on fast simulation
- Need to look at new and different computing platforms
  - What is the possible role of specialized hardware (e.g. for machine learning applications?)
  - Ensure the code is unbiased and can exploit benefits offered by different platforms as they evolve, code quality, resources, etc. But need to watch how the market evolves

Software Challenges

- Challenges for the future
  - CPU consumption and specifically its scaling with event pile-up
    - Algorithms and code are 20-year old
    - Code clean-up, parallelization, algorithm redesign, machine learning
    - No hotspot lefts, more R&D in this area is needed
  - Understanding performance is key to addressing the challenges
    - Tools, compilers, overall system efficiency, development, modelling

Data Analytics/Machine Learning

- Many diverse applications
  - Engineering (control systems, monitoring, anomaly detection, preventive maintenance, design, etc.)
  - Infrastructure (service analytics and optimization, data storage and transfers)
  - Physics (event indexing, classification, anomaly detection)
- Considerable effort ongoing on understanding the toolchain, integration with existing and new services (e.g. Spark+ROOT integration), workflows, etc.
- OpenAccess data critical for collaboration with third-parties
  - What is the right way of providing it (just the data? Packaged in a ready-made environment?)

Data Analytics/Machine Learning

- CERN openlab collaborations around DA/ML are rapidly growing
  - Existing projects with Siemens (control systems), Yandex (anomaly detection), Oracle (anomaly detection)
  - A main top-level topic for Phase VI
- Setting up a general platform and service
  - Currently working on underlying infrastructure, platforms and tools
  - Enable different use cases from physics (CMS), engineering (Controls), software tools (integration/benchmarking with ROOT), infrastructure, medical applications, etc.
- Exploit the general interest in using and providing technology in this field to build expertise within the HEP community and across communities

Data Analytics/Machine Learning

- ML opportunities in the LHC Experiments
  - Analysis, reconstruction, fast simulation, data taking
  - Bottom-up approach (replace blocks) or top-down approach (rethink)?
  - 3D-imaging approaches?
- Ongoing discussions to set up R&D collaborations within openlab from 2017 onwards
  - Talk to experts in industry and other communities

Collaborations Beyond HEP

- Growing interest in working on big societal problems with other communities having similar computing and data challenges
  - Distributed infrastructures
  - Fast simulation
  - Large-scale applications of Data Analytics/Machine Learning
- No formal mechanism in place so far
  - Pilot projects can provide a useful base for understanding difference and similarities, create and take part in experts communities, exchange ideas, look at policies, …

Towards CERN openlab VI

- In 2017, CERN openlab will prepare for the next phase
- **This workshop** is the start of a journey to collect requirements and needs from the community
- **Collaboration Board 2017** on March 22nd
- We will have **3 thematic 1-day workshops** in Spring 2017 (end February, end March (23rd?), end April, final dates TBC) open to industry, CERN groups and experiments
  - Infrastructure and Data Center Technologies (Clouds, Data Storage, Networks, Hardware Platforms, ...)
  - Compute architectures, platforms, software performance
  - Data Analytics and Machine Learning
- CERN openlab white paper, expected in Fall 2017
- CERN openlab Open Day in September 2017

Towards CERN openlab VI

- No major changes foreseen in the Framework model
  - Existing FA to be renewed for 3 more years
- Running Phase V projects will continue in Phase VI as necessary

Feedback

- The growth of CERN openlab in the past two years requires a new approach in the way we work and interact with the community
- We welcome feedback on this event, the next planned events and any additional activity where openlab can support you and your activities

EXECUTIVE CONTACT

Alberto Di Meglio, CERN openlab Head

alberto.di.meglio@cern.ch

TECHNICAL CONTACTS

Maria Girone, CERN openlab CTO

maria.girone@cern.ch

Fons Rademakers, CERN openlab CRO

fons.rademakers@cern.ch

COMMUNICATION CONTACTS

Andrew Purcell, CERN openlab Communication Officer

andrew.purcell@cern.ch

Mélissa Gaillard, IT Dep. Communication Officer

melissa.gaillard@cern.ch

ADMIN/FINANCE CONTACT

Kristina Gunne, CERN openlab Administration Officer

kristina.gunne@cern.ch