CERN openlab: Engaging Industry for Innovation in the LHC Run3-4 R&D Programme

Maria Girone, CERN IT, CERN openlab CTO, maria.girone@cern.ch

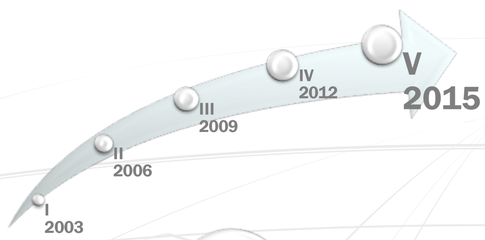

Mission

CERN openlab has been created to support the computing and data-management goals set by the LHC
- 15 years of innovative projects between CERN and leading IT companies

In phase V, CERN openlab is working to solve some of the key technical challenges facing the LHC in Run3 and Run4
- Mutual benefit for industry and research communities

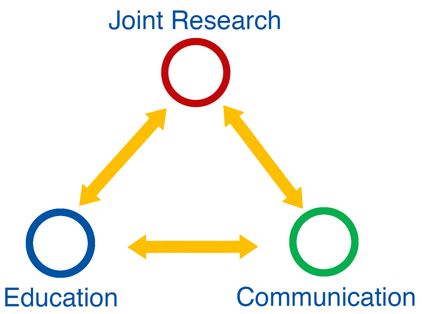

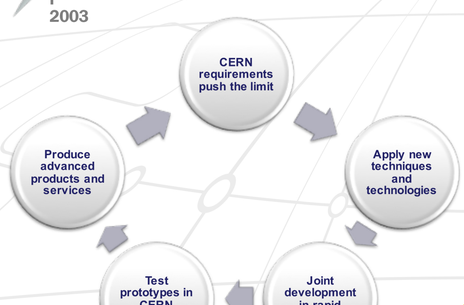

Research Activities

Data acquisition and filtering
Collecting data

Networks and connectivity
Connecting resources

Data storage architectures
Storing and serving data

Compute management and provisioning (cloud)
Managing resources for processing

Computing platforms, data analysis, simulation
Improving processing and code efficiency

Data analytics/Machine Learning
Extracting information

High Luminosity - LHC Challenges

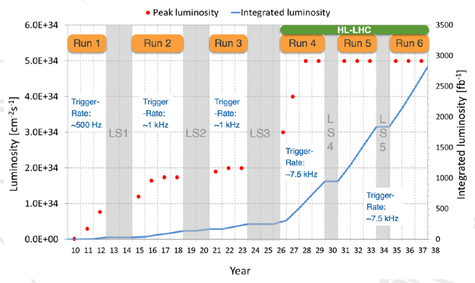

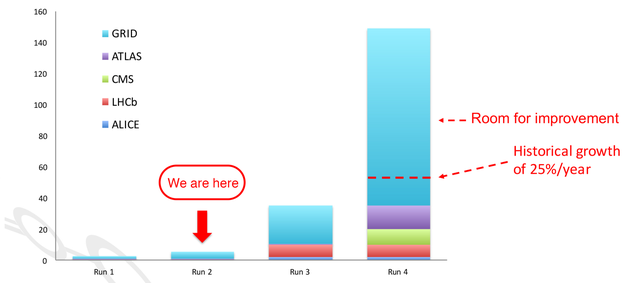

Assuming we need to save a factor-10 cost over what we may expect from:
- Moore's law expectations
- 1/3 from reducing infrastracture cost
- 1/3 from software performance
- 1/3 from more intelligence — write less data, move processing closer to experiment

Compute Management and Provisioning

Computing Platforms and Code Optimization

Data Analytics

The other CERN openlab research areas are more directly linked to experiments and IT services.

Data Storage

Network and Connectivity

Data Acquisition

Collaborations with Companies

Partners

intel ORACLE SIEMENS

Contributors

rackspace SEAGATE BROCADE IDT

Associates

ComTrade HUAWEI Yandex

Research

Kazan Federal University, Newcastle University, GSI, EMBL-EBI

Collaboration with Communities

WLCG Worldwide LHC Computing Grid, IT, EN, EP, ATLAS, ALICE, CMS, LHCb

Education

- Regular workshops
- Topical workshops and lectures
- Training courses
- CERN openlab Summer Student Programme

Networking and Connectivity

- **openlab collaboration with Brocade**
  - Enhance and generalize the Brocade Flow Optimizer (BFO) architecture to meet CERN requirements
- **Fellow actively participating in product development**
  - Strong integration with Brocade's development team based in San Jose, California
  - Preparing functional specifications, code developments, testing
  - BFO 1.3 released in Aug 2016 with essential features and enhancements
  - CERN contributions: Bro IDS & Palo Alto firewall plugins
- **Use case identified: BFO-enabled Intrusion Detection System**
  - Scalable, programmable IDS setup
  - Featuring BFO, OpenFlow and Brocade MLXe hardware
  - Symmetric load-balancing, mirroring suspicious traffic to a dedicated PCAP server, dynamic IDS offloading of bulk data transfers
  - Promising perspective of production deployment

Rackspace and CERN Collaboration

- **Federation project is now complete with all code integrated into the upstream OpenStack community**
- **Now working on use of containers on OpenStack**
  - Project 'Magnum'
  - Kubernetes, Swarm and Mesos support
  - 37 improvements from CERN in 1st year
  - CERN is the second largest contributor to the project
  - Scale test for 7 million requests/s using 1000 servers successful

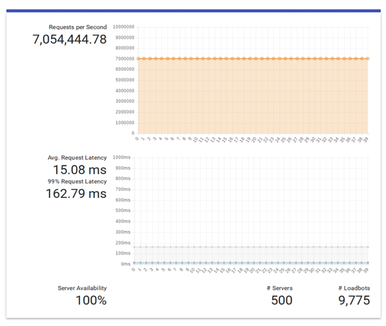

Code Modernization Project

- **The increasing need for computing has prompted an effort to optimize scientific codes for the new computing architectures**
  - Possible to achieve enormous improvements in code performance using modern techniques
  - One of the few areas with enough potential for improvement to close the resource gaps in the upgrade program
- **The Code Modernization Project is an umbrella for addressing several use cases in different disciplines**
  - Possible extensions currently under discussion

Geant The best Geant ever

BioDynaMo The Biology Dynamic Modeller

FAIR

HTCondor High Throughput Computing

Data Analytics and Machine Learning

- **Investigating ways to more efficiently select events from the stream of data using "big data" techniques**
- **Looking at ways to use machine learning for**
  - Resource optimization and anomaly detection
  - Event categorization
  - Identification of physics objects
  - Reconstruction through computer visualization

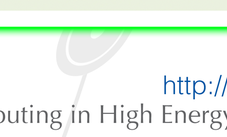

High Throughput Computing Collaboration

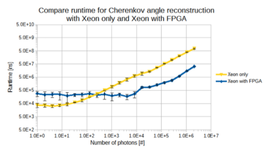

Complex inverse ray-tracing ~ factor 30 speed—up FPGA vs Xeon
(Note: this FPGA has no floating-point support, significant improvement expected for next generation (Arria10) with floating point support

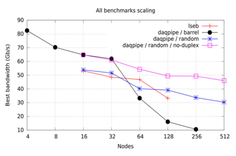

Future high-speed DAQ will use 100 Gbit/s LAN. *Measurement* of up to 1024 simultaneous 100 Gbit/s communications on Intel OmniPath Achieve > 37 Tbit/s aggregated throughput (Note: best performance achieved by leaving scheduling of traffic to intelligence on Intel switches)

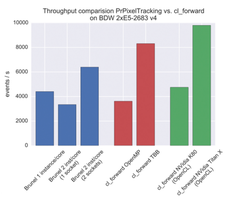

A comparison of different thread parallel implementations (OpenMP, TBB, OpenCL) of a track reconstruction algorithm with the production implementation. Using TBB a speedup of 30% can be achieved on the same hardware. The same code is now being tested and ported to Intel Xeon Phi KNL.

IDT Collaboration on RapidIO

- **ROOT and DAQPIPE continuously improved**
- **ROOT, DAQPIPE and Hadoop performance metrics in place**
- **Contribution to Realtime Conference 2016**
- **Scale up: 16 standard servers mounted with RapidIO cards**
- **Preparing for CHEP – performance and scalability**

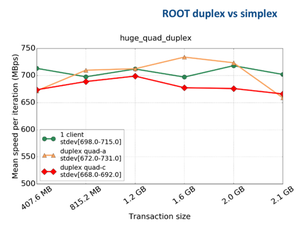

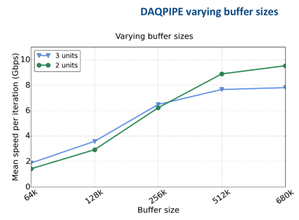

Storage Technology R&D

- **Object Disks are available and offer**
  - More scalable semantics matched to shingled recording
  - Open Standard Process with all main vendors (Seagate, WD. Toshiba)
- **Joint project with Seagate**
  - Transparent integration with EOS system achieved

- **Next Steps**
  - Install second 1PB system with second-generation Kinetic disks
  - Release high-available multi-path access to Kinetic disks in EOS
  - Exercise failure recovery and repair modes
  - Optimize and benchmark read performance for CERN analysis/archive use cases
  - Evaluate TCO gain with EOS prototype system

Deployment: Multi-Path Capability

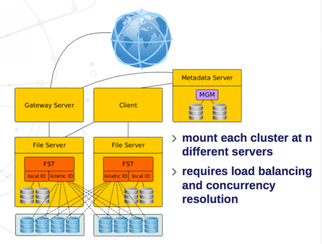

› mount each cluster at n different servers

› requires load balancing and concurrency resolution

CERN openlab Management Team

Alberto Di Meglio
Maria Girone
Fons Rademakers
Kristina Gunne
Sotirios Pavlou
Andrew Purcell
Maria-Athanasia Pachou

openlab-V-management@cern.ch
www.cern.ch/openlab

Poster Abstract Number 318

The 22nd International Conference on Computing in High Energy and Nuclear Physics, CHEP 2016