The Challenges of Scientific Computing

Alberto Di Meglio
CERN openlab CTO

DOI: 10.5281/zenodo.7116

HUAWEI

CERN openlab

Outline

- What is CERN and How it Works
- Computing and Data Challenges in HEP
- New Requirements and Future Directions
- Scientific Collaborations

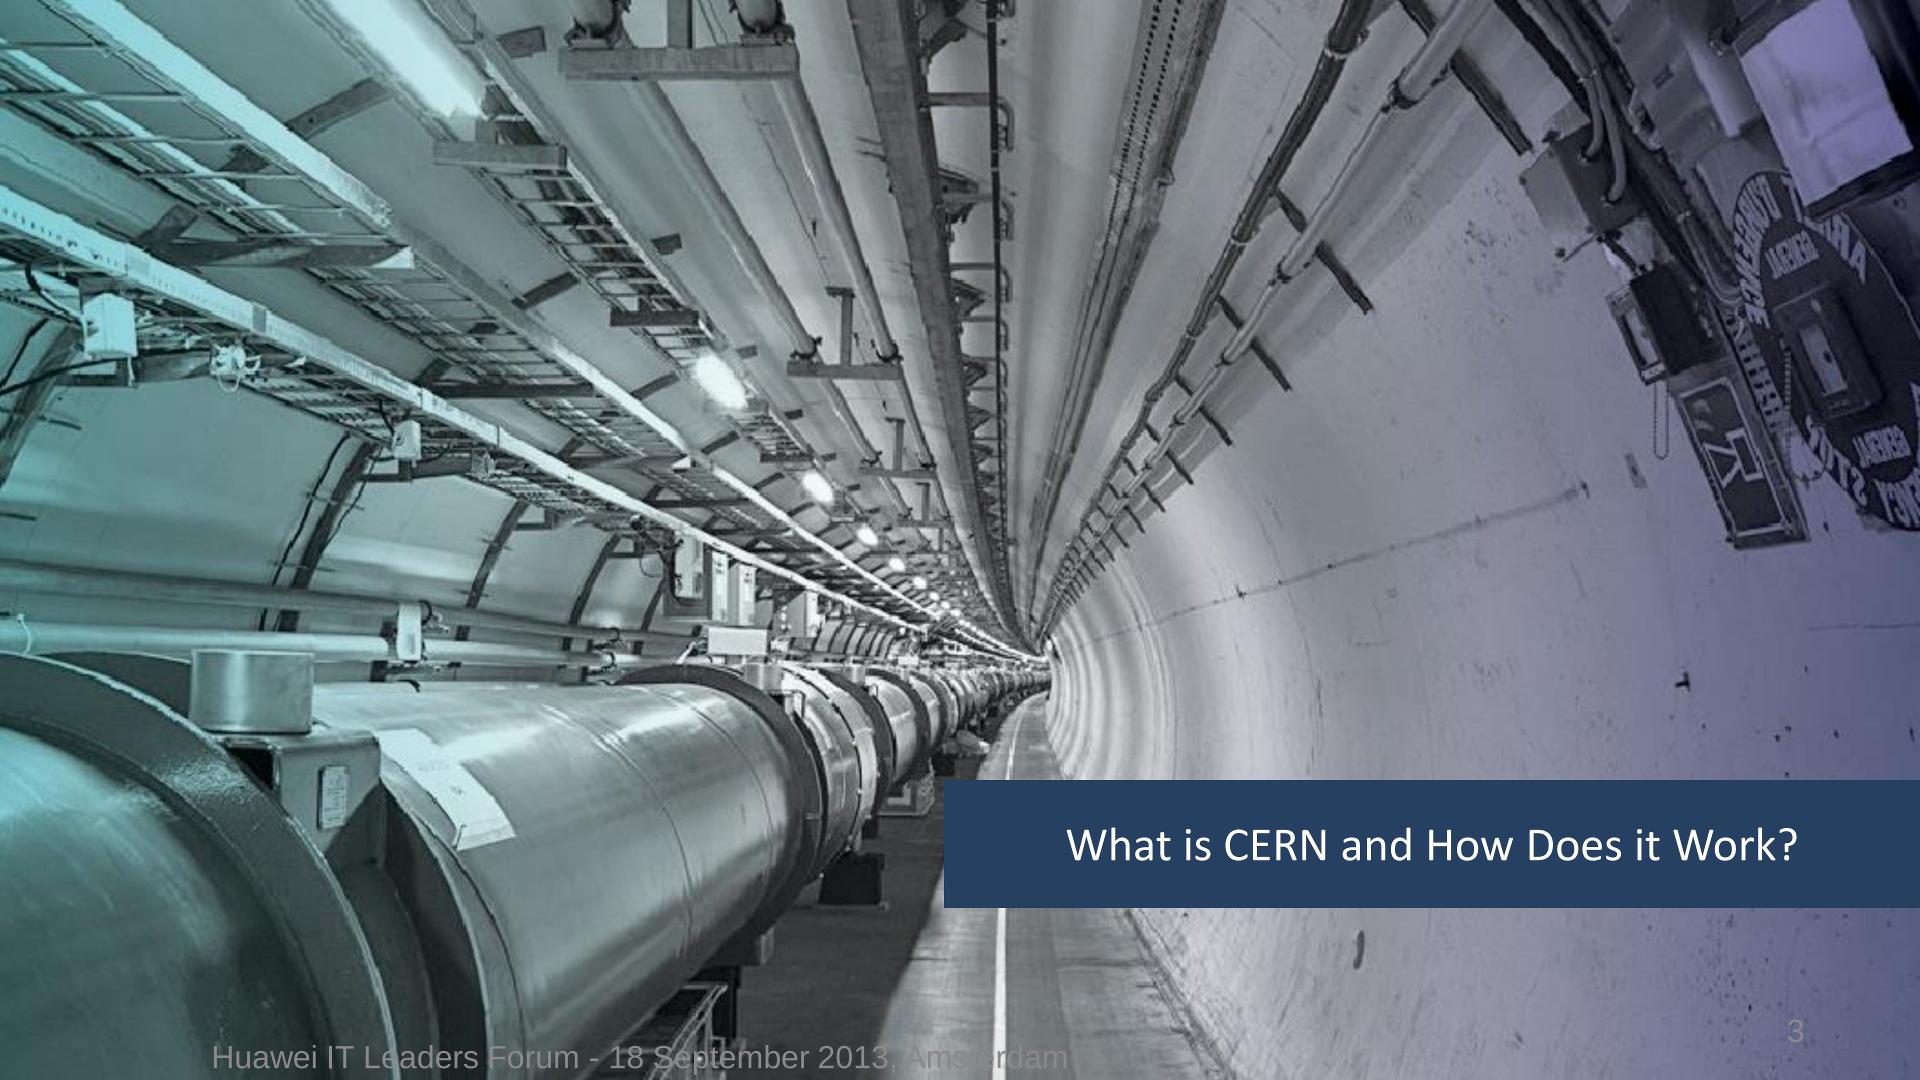

What is CERN and How Does it Work?

What is CERN?

European Organization for Nuclear Research

~ 2300 staff

~ 1050 other paid personnel

~ 11000 users

Budget (2012) ~1100 MCHF

Founded in 1954

20 Member States: Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, the Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom

Candidate for Accession: Romania

Associate Members in the Pre-Stage to Membership: Israel, Serbia, Cyprus

Applicant States: Slovenia, Turkey

Observers to Council: India, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and UNESCO

What is the Universe made of?

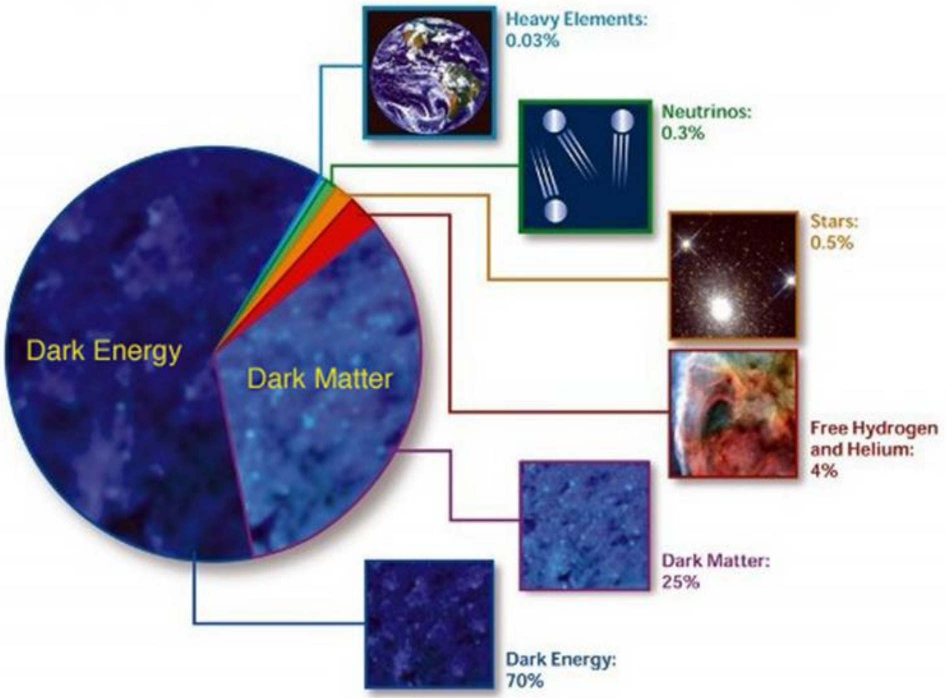

- What gives the particles their masses?
- How can gravity be integrated into a unified theory?
- Why is there only matter and no anti-matter in the universe?
- Are there more space-time dimensions than the 4 we know of?
- What is dark energy and dark matter which makes up 95% of the universe ?

The Large Hadron Collider (LHC)

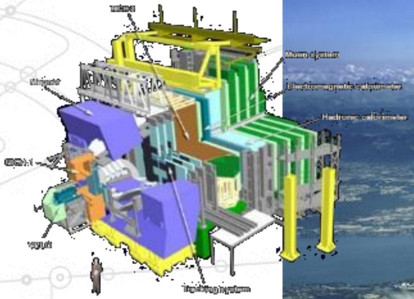

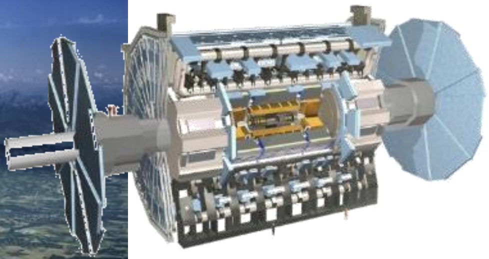

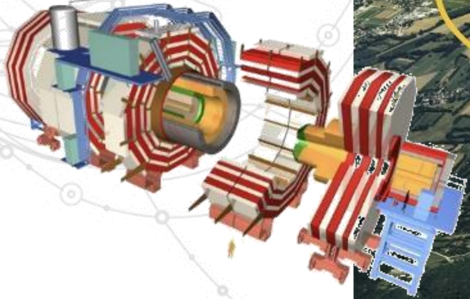

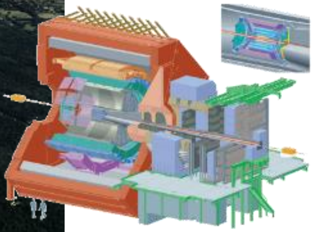

LHC Facts

- **Biggest accelerator (largest machine) in the world**
- **Fastest racetrack on Earth**
  - Protons circulate at 99.9999991% the speed of light
- **Emptiest place in the solar system**
  - Pressure 10^{-13} atm (10x less than on the moon)
- **World's largest refrigerator -271.3 °C (1.9K)**
- **Hottest spot in the galaxy**
  - temperatures 100 000x hotter than the heart of the sun
  - 5.5 Trillion K
- **World's biggest and most sophisticated detectors**
- **Most data of any scientific experiment**
  - 20-30 PB per year (as of today we have about 75 PB)

Collisions in the LHC

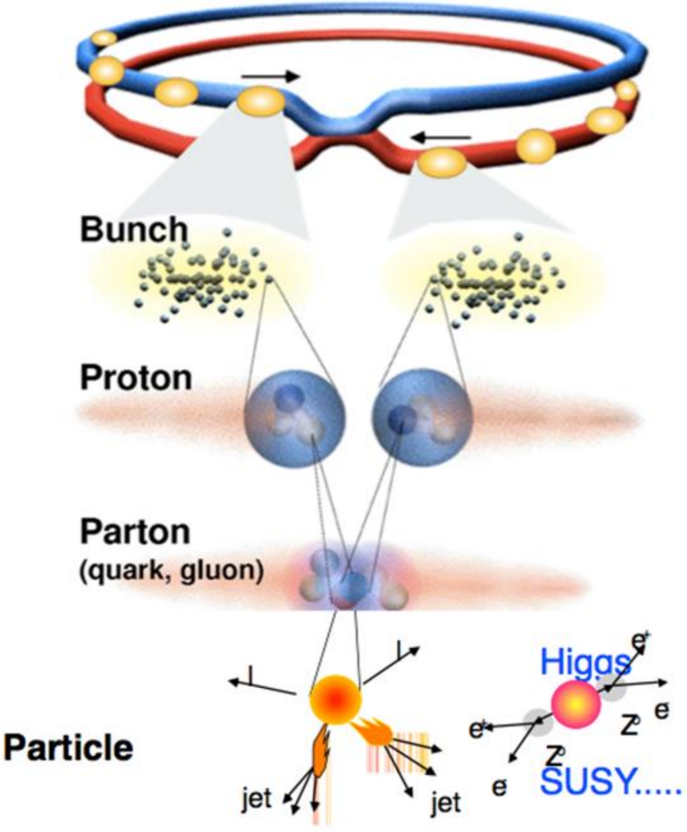

Proton - Proton	2808 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	$10^{34} cm^{-2}s^{-1}$
Crossing rate	40 MHz
Collision rate ≈	10^7-10^9

New physics rate ≈ .00001 Hz

Event selection:
1 in 10,000,000,000,000

Computing and Data Challenges in HEP

The LHC Challenges

- Signal/Noise: 10^{-13} (10^{-9} offline)
- Data volume
  - High rate * large number of channels * 4 experiments
  - ➔ ~30 PB of new data each year
- Compute power
  - Event complexity * Nb. events * thousands users
  - ➔ 300 k CPUs
  - ➔ 170 PB of disk storage
- Worldwide analysis & funding
  - Computing funding locally in major regions & countries
  - Efficient analysis everywhere
  - ~1.5M jobs/day, 150k CPU-years/year
  - ➔ GRID technology

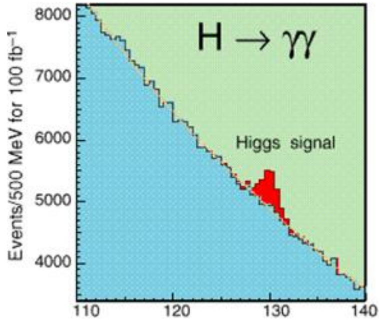

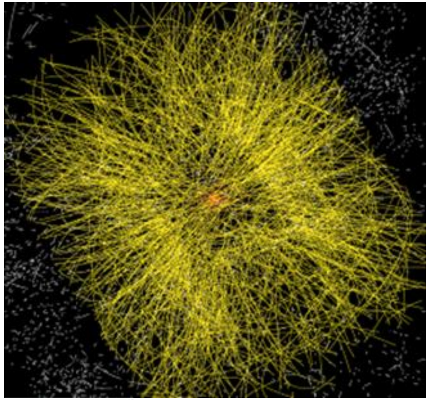

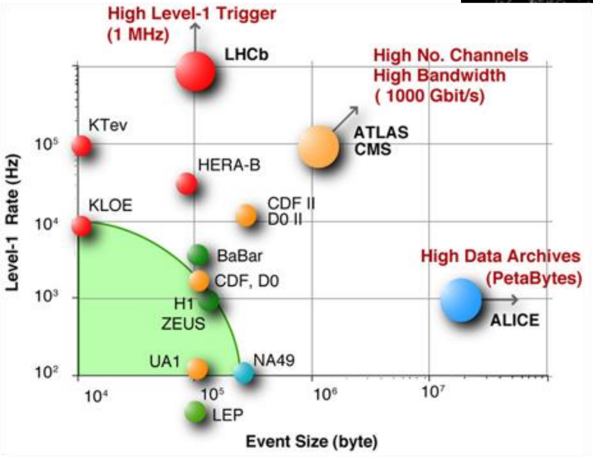

Data Handling and Computation

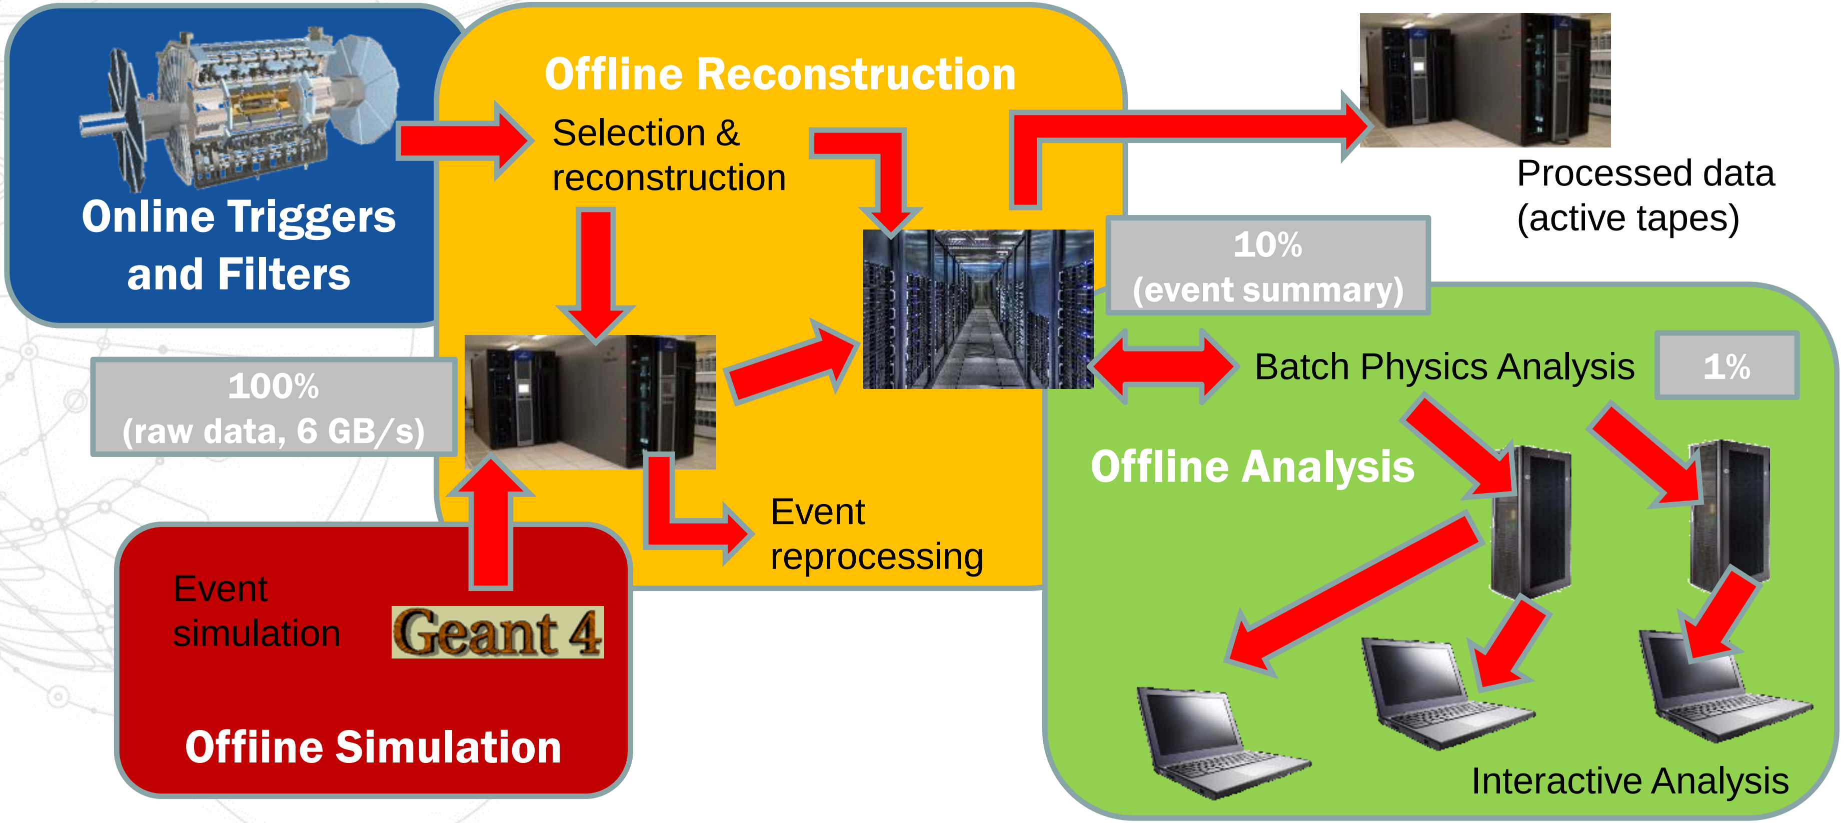

Data Storage

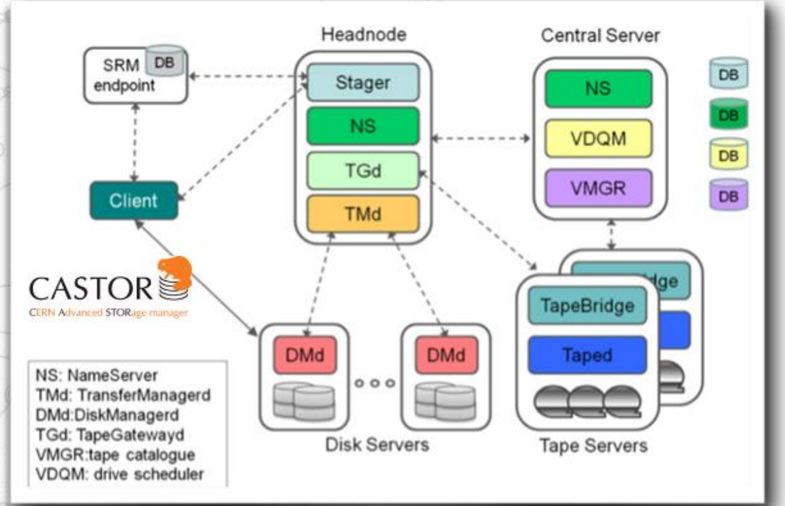

CASTOR and EOS
Storage Systems developed at CERN

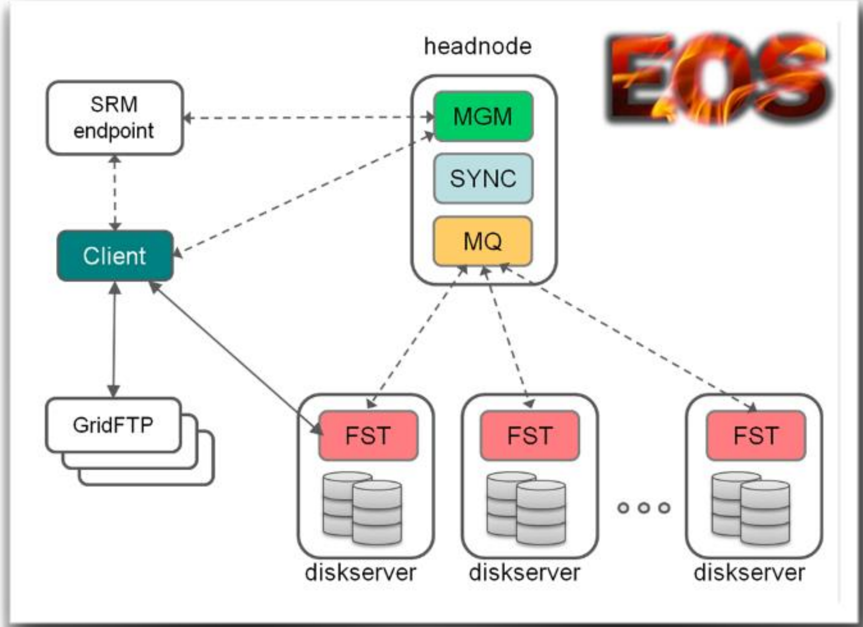

CASTOR and EOS are using the same commodity disk servers

- With RAID-1 for CASTOR (2 copies in the mirror)
- JBOD with RAIN for EOS
- Replicas spread over different disk servers
- Tunable redundancy

The Grid

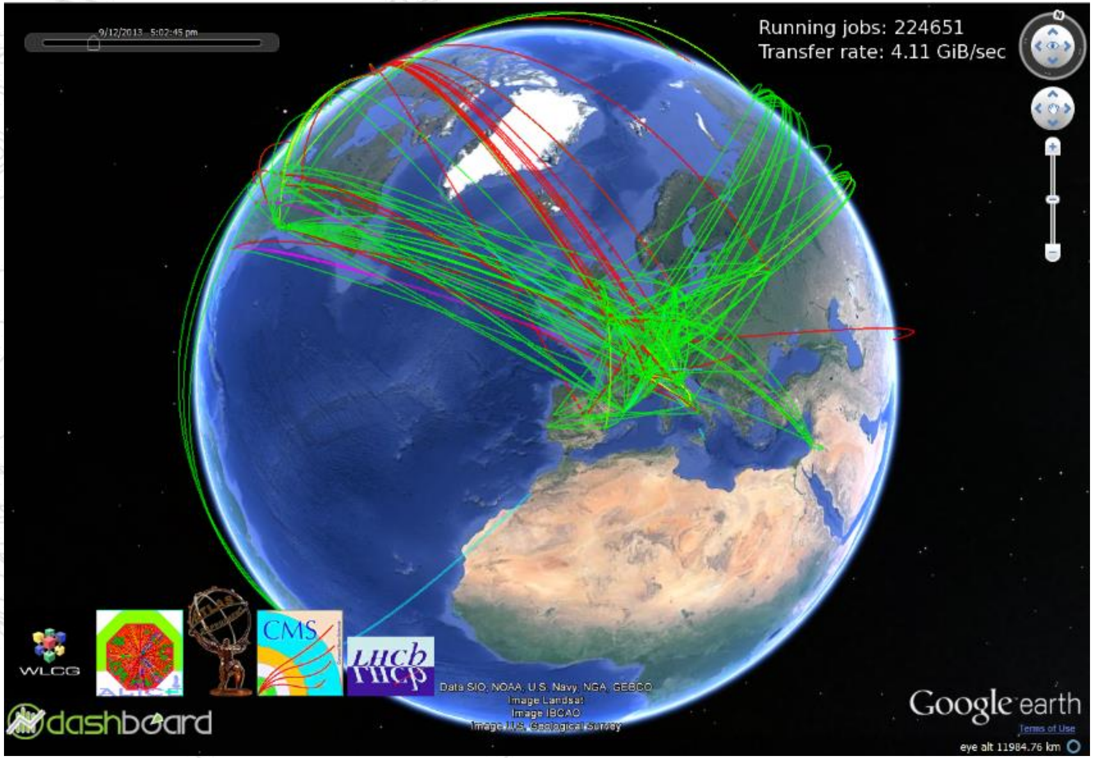

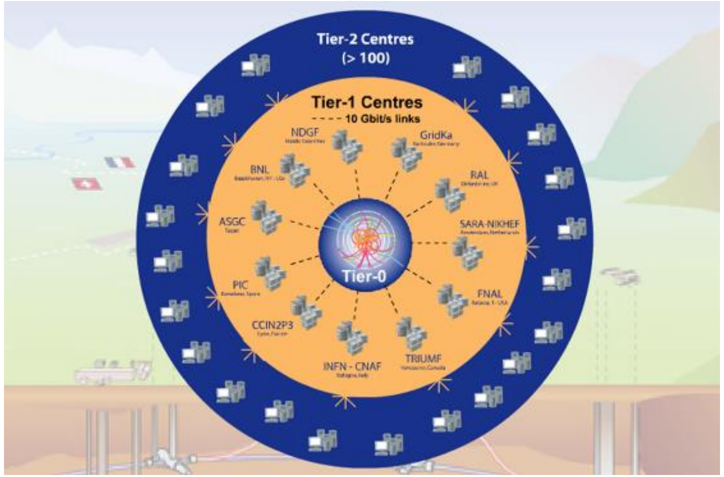

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (11 centres):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (~130 centres):
- Simulation
- End-user analysis

New Requirements and Future Directions (Big Data and Data Analytics)

LHC Schedule

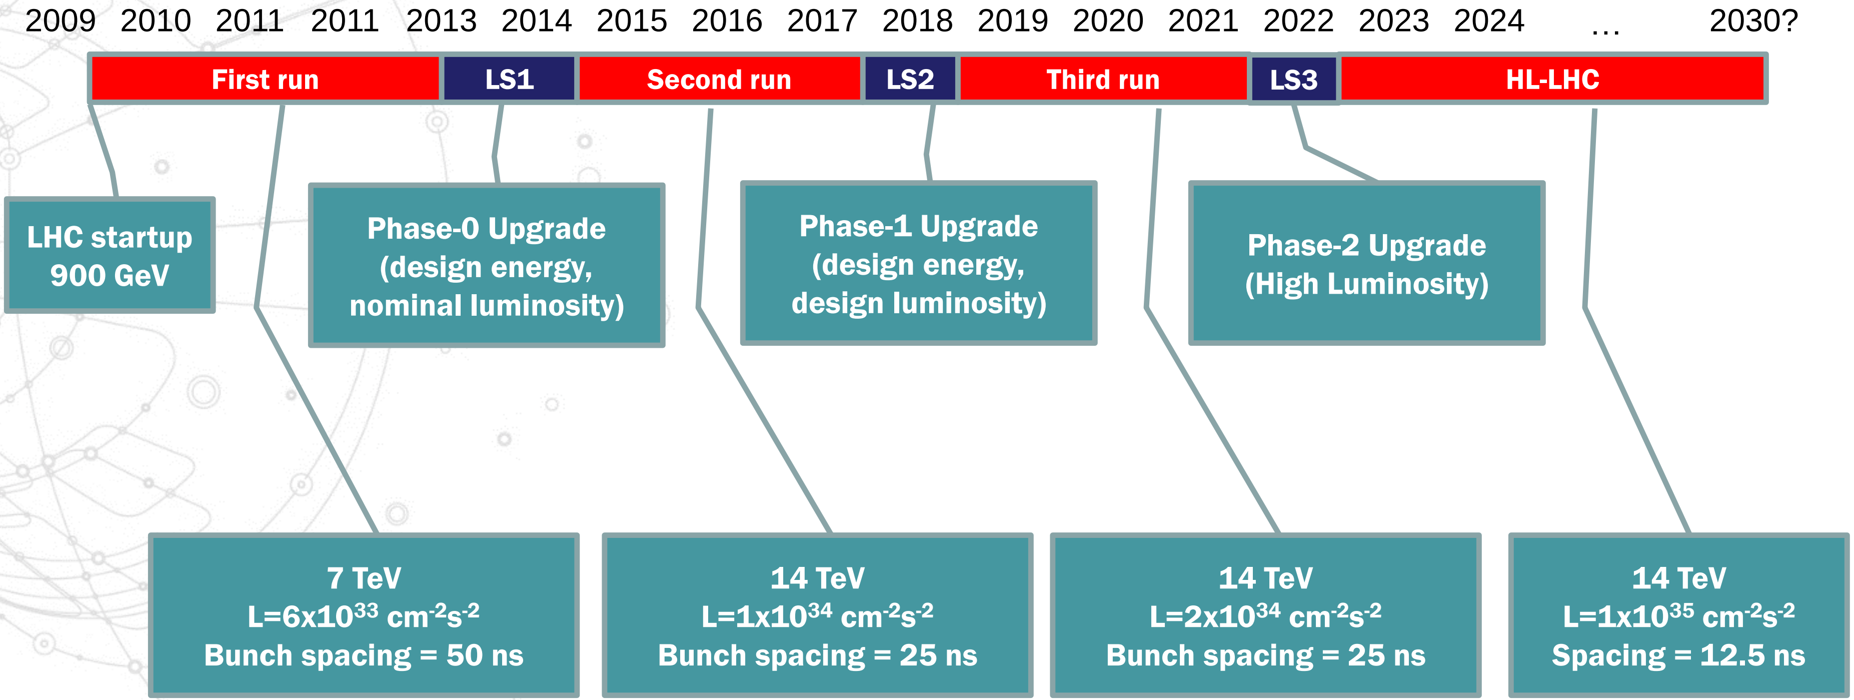

Online Triggers and Filters

Online DAQ and Triggers

- Most events produced in the detectors describe known physics
- Only 1 in 10^{13} events are considered interesting now, but discarded raw data is lost forever
- Filtering is done using two-level triggers, the first in HW (L1), the second in SW (HLT)
- L1 triggers use complex custom processors, difficult to program, maintain and reconfigure
- LS1, LS2 and LS3 require higher and higher data rates
- Can the L1 run at the same rate as LHC?

- Implement L1 in software
- Replace HW triggers with commodity processors/co-processors (GPGPUs/XeonPhis?) → easier to program and upgrade/maintain
- Long-term, merge L1 and HLT, need lots of fast, low-power, low-cost links
- DAQ evolution
  - multi Tbit transfers
  - Close integration of network and CPUs
- Rackscale architecture?

CERN openlab

Event simulation Geant 4

Offline Simulation

Multi-Core Platforms

- Very high-level of parallelism, both online and offline
- Parallelism exploited sending multiple independent computational jobs on separate physical or virtual nodes
- Multiple physical cores partitioned across virtual machines
- Current software not written to exploit multi-cores
- Very important issue with detector simulations, very CPU intensive tasks

- Exploit multi-core platforms using vectorization techniques
- Build experience with data and task parallelism using Cilk and Haswell (CilkPlus and GCC)
- Requires rewriting the software → Geant 5

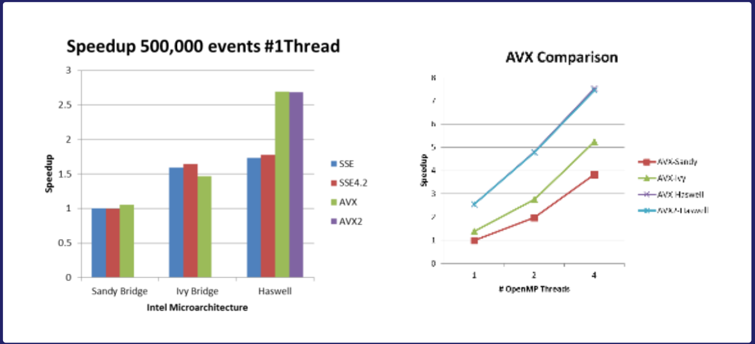

Cloud Infrastructure

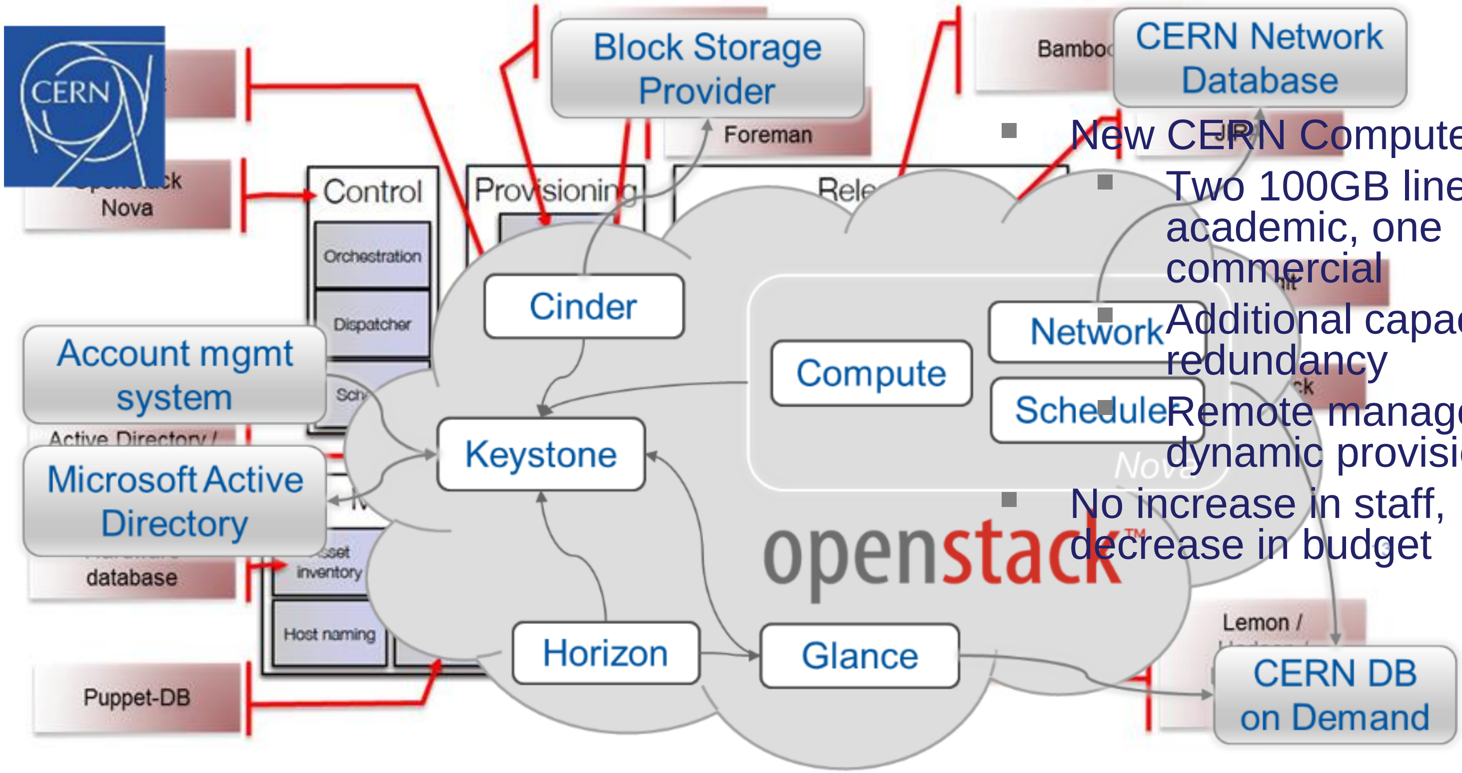

- New CERN Computer Centre
  - Two 100GB lines, one academic, one commercial
  - Additional capacity, redundancy
  - Remote management, dynamic provisioning
- No increase in staff, decrease in budget

Cloud Storage

- **Data storage is critical**
  - Currently 75/100 PB tapes, 25/40 PB disk, getting close to max capacity
- Need to be able to **scale** both predictably (planning) and dynamically (peak times)
- Can **industrial cloud storage** deliver the required quality of service maximizing Reliability, Performance and Cost at the same time?

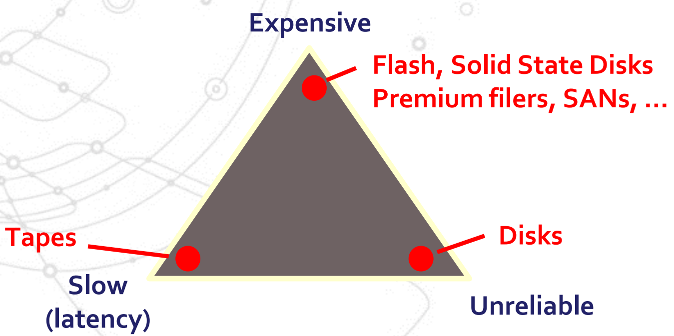

- **Partership CERN/IT & Huawei (S3 Storage Appliance, 0.8 PB, 1 year)**
- **Shared areas of investigations**
  - Reliability **(100% despite disk failures)**
  - Functional tests on Amazon S3 protocol and interoperability with Grid production environments **(tested ok)**
  - Performance measurements with sparse data read **(comparable to existing LHC production systems)**
  - Total Cost of Ownership (TCO): clear model, **no hidden costs found**
- **References used for comparison**
  - Open Source self assembled solutions (based on openstack / swift)
  - Traditional Intel-based file server based on Linux operating system and locally attached disks

Data Analytics

- Massive amounts of data from
  - The LHC data processing
- But also
  - The LHC subsystems
    - Magnets
    - Cryogenics
    - Control systems
    - Electrical systems
  - Logging data and alerts
- How can the data be analyzed efficiently to find patterns or detect problems at early stages?

- Investigations of
  - Physics data analysis in databases exploiting columnar DBs (Oracle Exadata)
  - Hadoop and MapReduce techniques
  - Physics event indexing, filtering and searching
    - Machine Learning techniques

Scientific Collaborations

International Scientific Collaborations

- Many scientific projects are global collaborations of 100s of partners
- Efficient computing and data infrastructures have become critical as the quantity, variety and rates of data generation keep increasing
- Funding does not scale in the same way
  - Optimization and sharing of resources
- Collaboration with commercial IT companies increasingly important
  - Requirements are not unique anymore

CERN Biomedical Facility

Scientific Computing as a Service

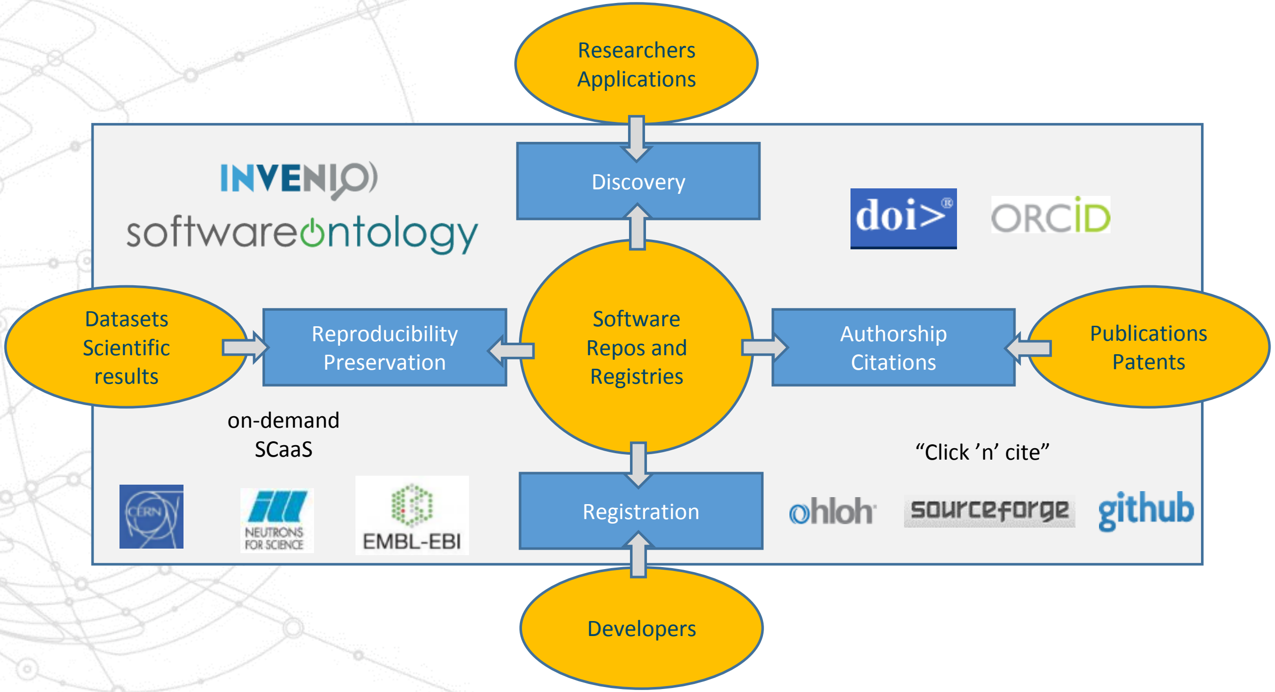

- An on-demand computing and data analysis service
- Unique IDs and formal relationships among digital objects (publications, datasets, software, people identities)
- Reproducibility and preservation
- Results attribution and recognition

CERN openlab in a nutshell

CERNopenlab

- A science – industry partnership to drive R&D and innovation with over a decade of success
- Evaluate state-of-the-art technologies in a challenging environment and improve them
- Test in a research environment today what will be used in many business sectors tomorrow
- Train next generation of engineers/employees
- Disseminate results and outreach to new audiences

PARTNERS

hp

HUAWEI

intel

ORACLE

SIEMENS

CONTRIBUTOR

rackspace

ASSOCIATE

Yandex

CERN openlab recipe: key ingredients

- The extreme demands from CERN’s scientific programme
- The alignment of goals between partners
- Trust
- Young researchers with their talents, expertise and energy
- Efficient and lightweight structure
- Regular checkpoints/reviews
- Outreach/communications/training

Conclusions

- CERN and the LHC program have been among the first to address “big data” challenges
- Solutions have been developed and important results obtained
- However, not unique anymore in both scientific and consumer applications
- Need to exploit emerging technologies and share expertise with academia and commercial partners
- LHC schedule will keep it at the bleeding edge of technology, providing excellent opportunities to companies to test ideas and technologies ahead of the market

GET IN TOUCH www.cern.ch/openlab

EXECUTIVE CONTACT

Bob Jones, Head of CERN openlab

bob.jones@cern.ch

TECHNICAL CONTACTS

CERN openlab Chief Technology Office

Alberto Di Meglio, alberto.di.meglio@cern.ch

Sverre Jarp, sverre.jarp@cern.ch

COMMUNICATION CONTACT

Mélissa Gaillard, CERN openlab Communications Officer

melissa.gaillard@cern.ch