
CERN

CERN Data Services for Science

Massimo Lamanna / CERN

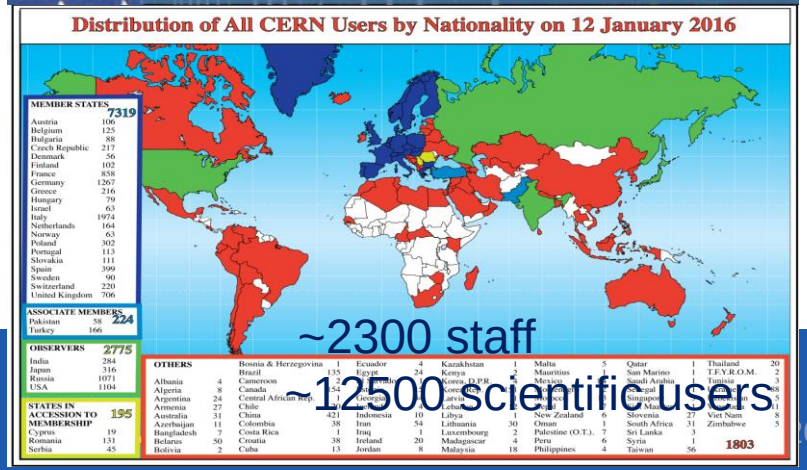

CERN

- Mandate:
  - Push back the frontiers of knowledge
  - Develop new technologies for accelerators and detectors
    - E.g. the Web
  - Train scientists and engineers of tomorrow
  - Unite people from different countries and cultures

Large Hadron Collider

Quantity	number
Circumference	26 659 m
Dipole operating temperature	1.9 K (-271.3°C)
Number of magnets	9593
Number of main dipoles	1232
Number of main quadrupoles	392
Number of RF cavities	8 per beam
Nominal energy, protons	7 TeV
Nominal energy, ions	2.76 TeV/u (*)
Peak magnetic dipole field	8.33 T
Min. distance between bunches	~7 m
Design luminosity	10^{34} cm^{-2} s^{-1}
No. of bunches per proton beam	2808
No. of protons per bunch (at start)	1.1×10^{11}
Number of turns per second	11 245
Number of collisions per second	600 million

(*) Energy per nucleon

Event rate to tape o(100) Hz
(Online filtering)

- 10,000 magnets
  - 1232 19-metre dipoles
- Largest superconductive installation in the world
  - 8000 km super-conducting cables
  - 120 t of liquid Helium

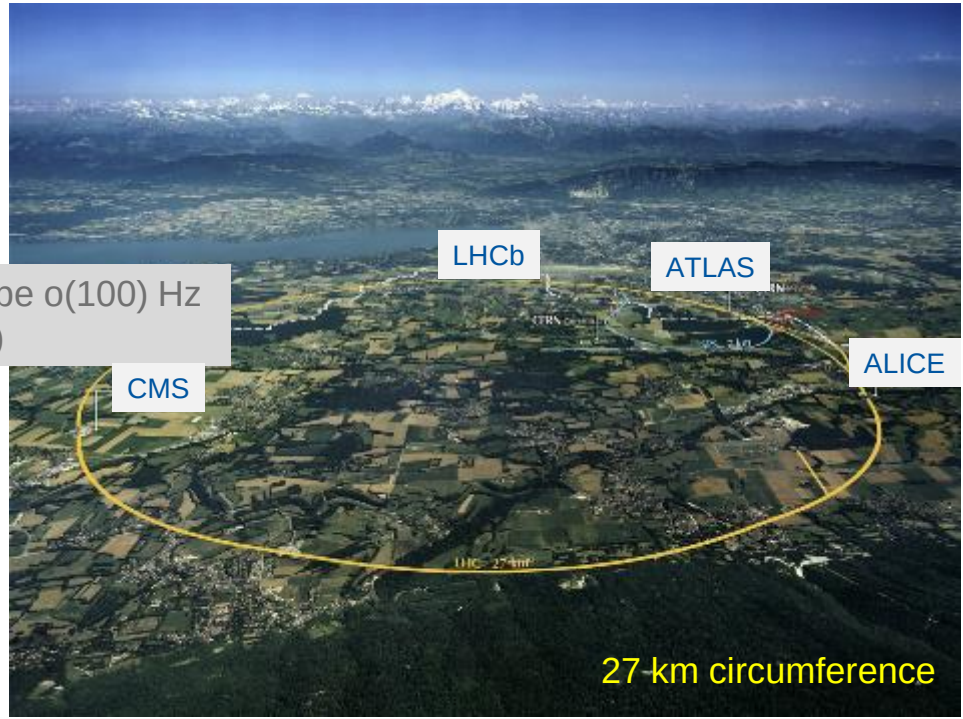

ATLAS

ATLAS is one of the 4 main LHC experiments (together with ALICE, CMS and LHCb).

It is a general-purpose detector designed to cover the widest possible range of physics at the LHC, from the search for the Higgs boson to supersymmetry (SUSY) and extra dimensions. The main feature of the ATLAS detector is its enormous doughnut-shaped magnet system. This consists of eight 25-m long superconducting magnet coils, arranged to form a cylinder around the beam pipe through the centre of the detector. ATLAS is the largest-volume collider-detector ever constructed. The collaboration is the largest among the LHC experiments and consists of more than 3000 physicists (~1000 students) from 171 institutes in 37 countries (May2013).

Detector: 46 m long, 25 m high and 25 m wide; 7000 t
http://atlas.ch/

Data rate (proton-proton): ≈ 1GB/s

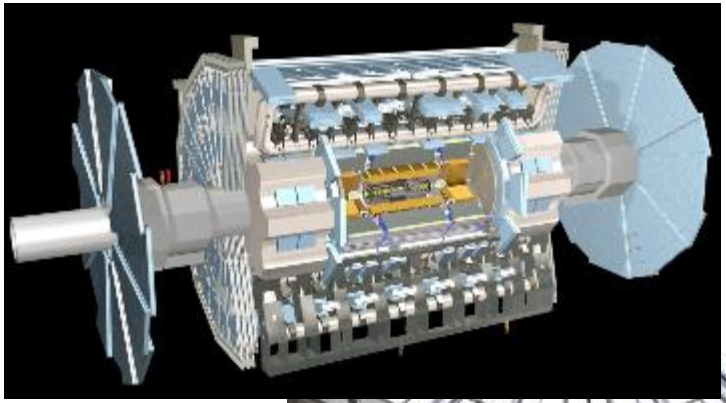

~ 100 million electronic channels
~ 3 000 km of cables

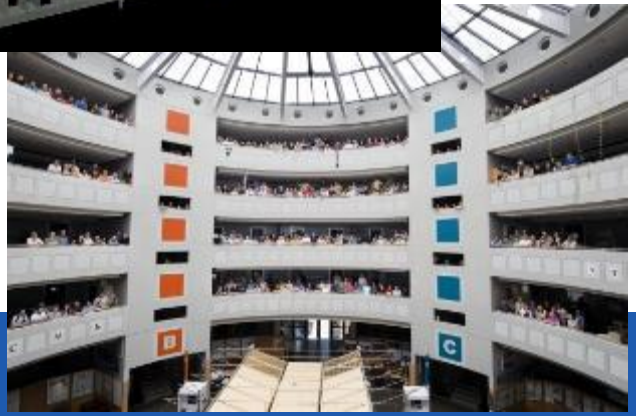

Part of the ATLAS collaboration

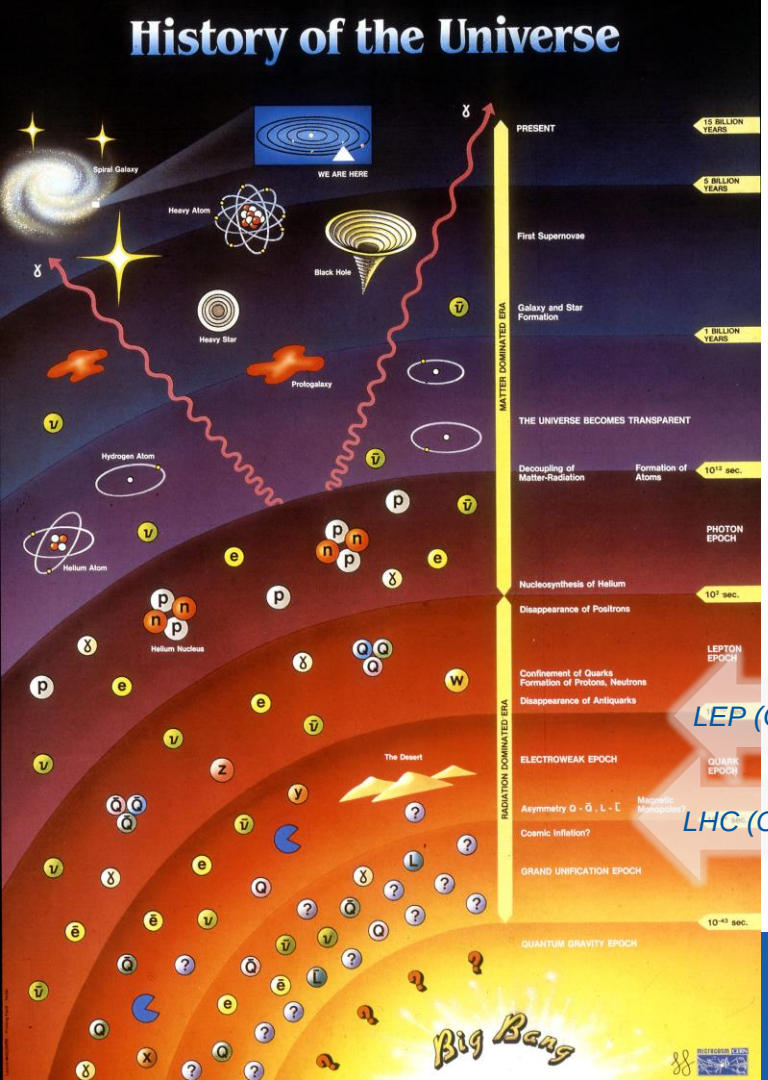

Fundamental research

Higher energies

↔

smaller scale

Higher temperature (energy)

↔

Big Bang earlier points in time

Microscopic phenomena

↔

Large-scale structure of the universe

LEP (CERN) 1989-2000

LHC (CERN)

Higgs Boson Discovery

2012

Higgs to 4μ candidate event

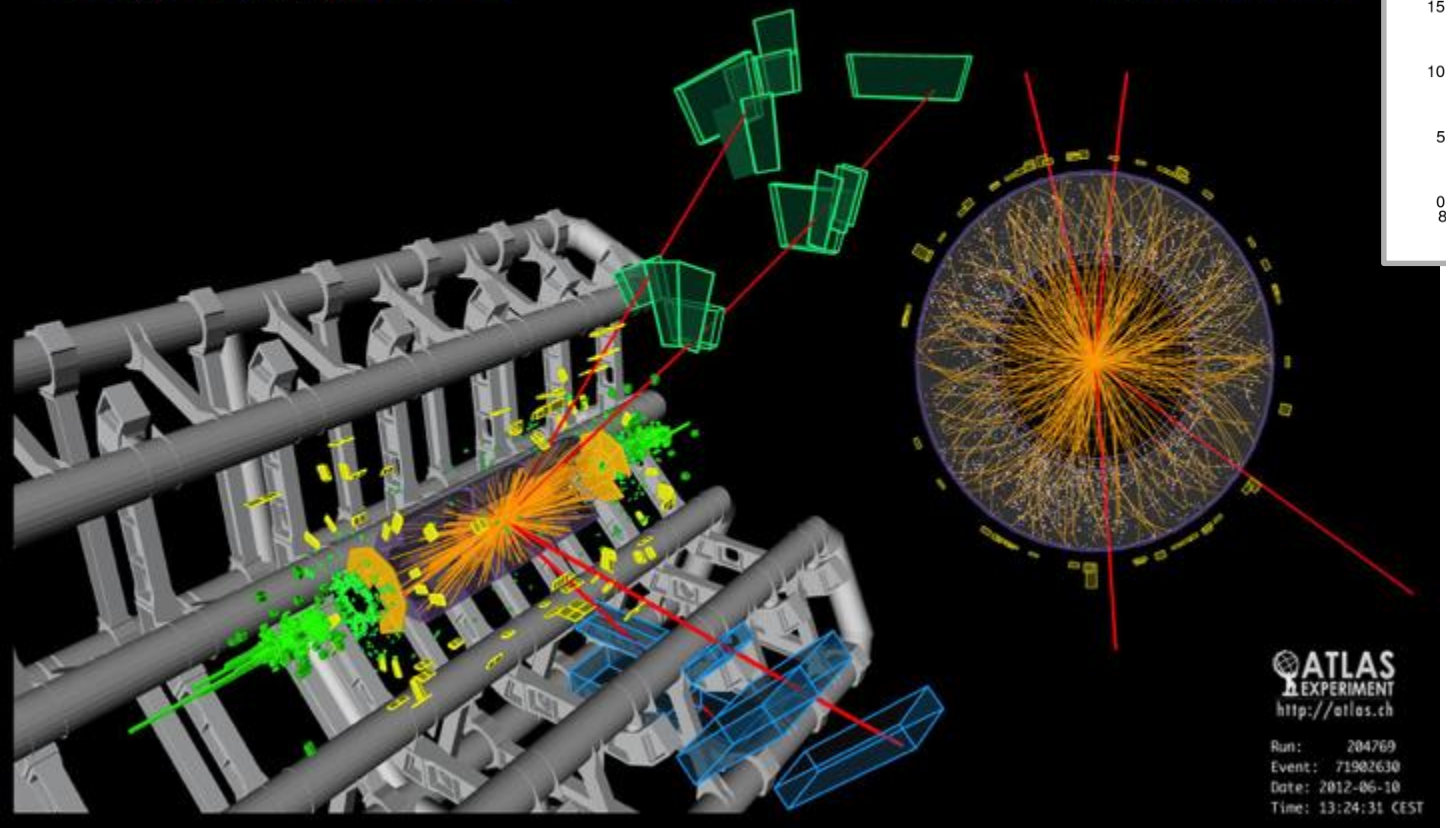

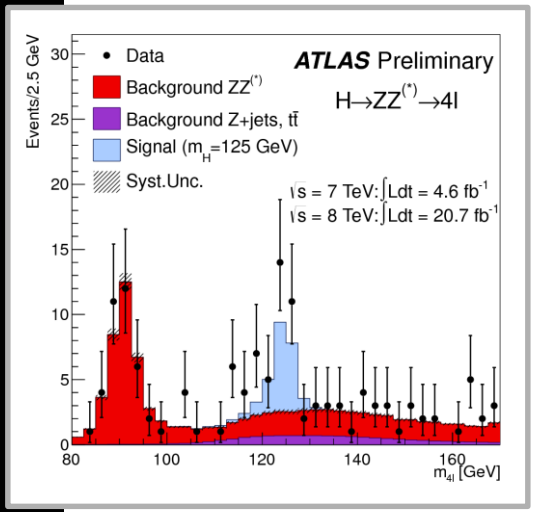

Data science

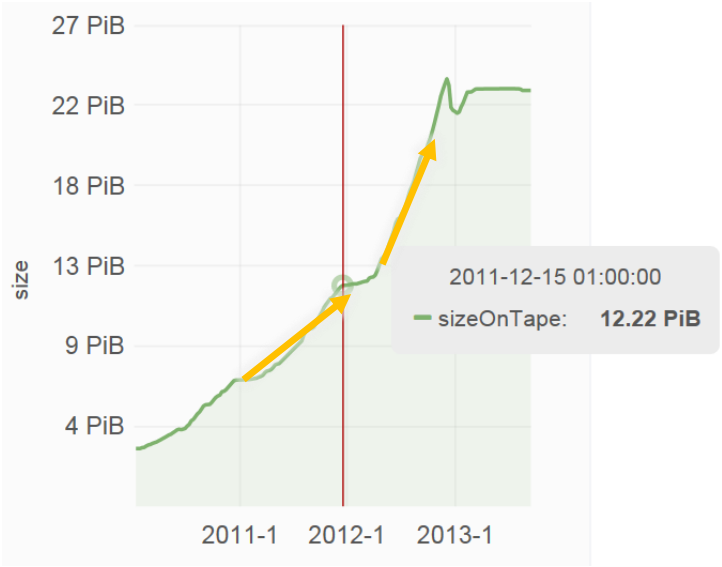

ATLAS data in Run 1 (2011 and 2012)

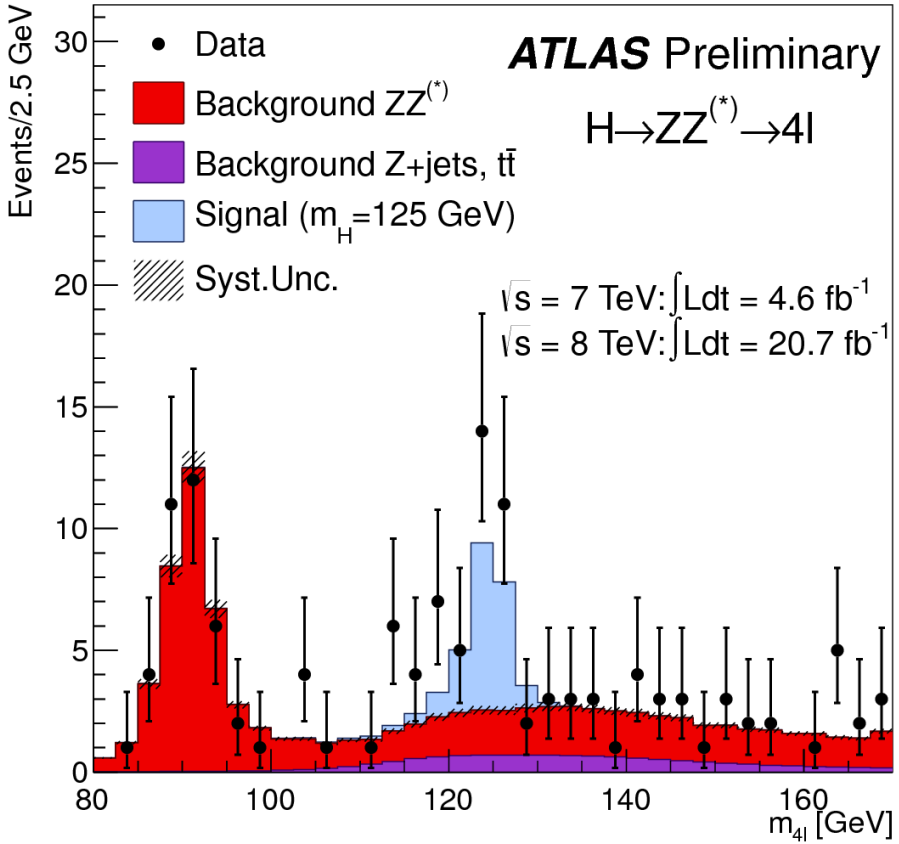

Steady requirements?

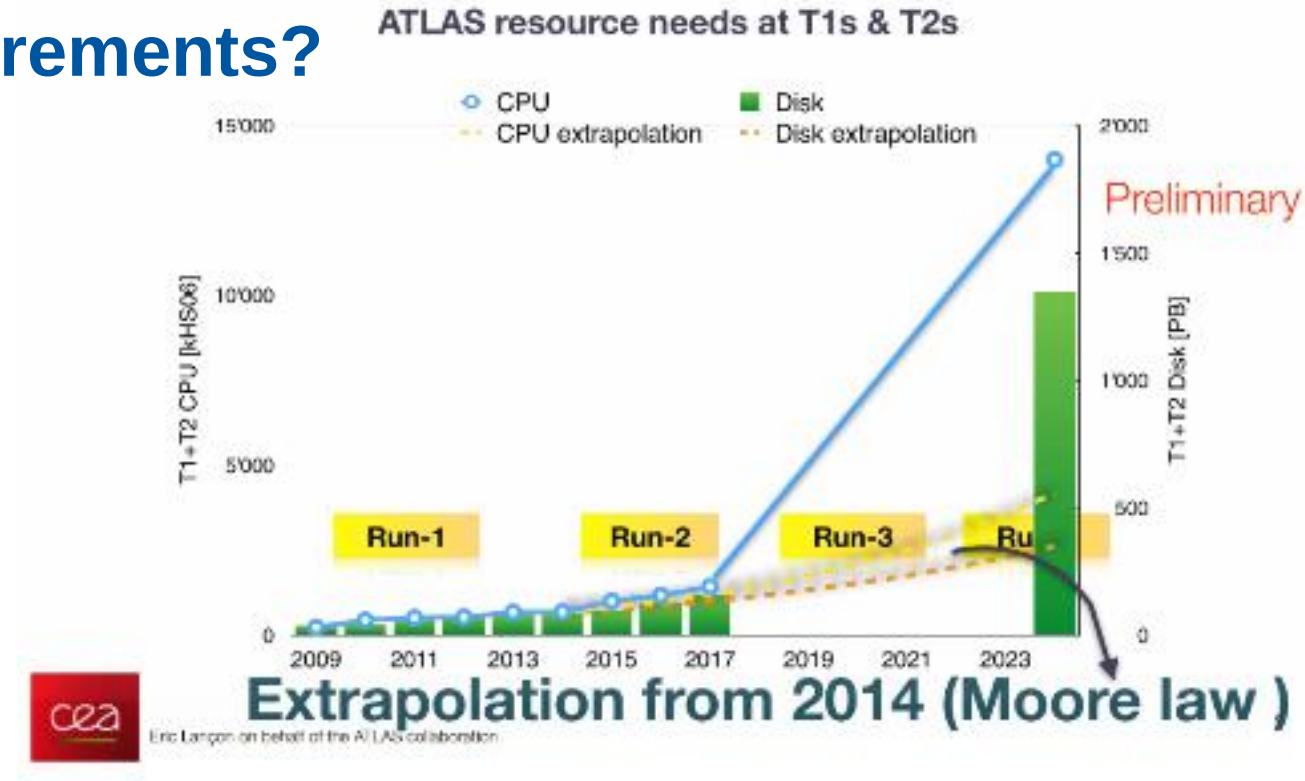

M.Krzewicki, E. Lancon et al., ECFA HL-LHC Computing Workshop, 2014

WLCG

Worldwide LHC Computing Grid

Make LHC computing possible

Worldwide infrastructure (collaboration) open to all LHC physicists

Computing/storage resources at CERN: ~ 20%; 80% across about 200 sites worldwide

Data Reconstruction

Goals: data quality and immediate access for analysis

Organised activity dominated by heavy processing and replication (each expt: 1-8 GByte/s)

Data Analysis

Goals: extract physics quantities (discovery)

Individual activities dominated by event selection and sharing (thousands of physicists)

(Detector) simulation

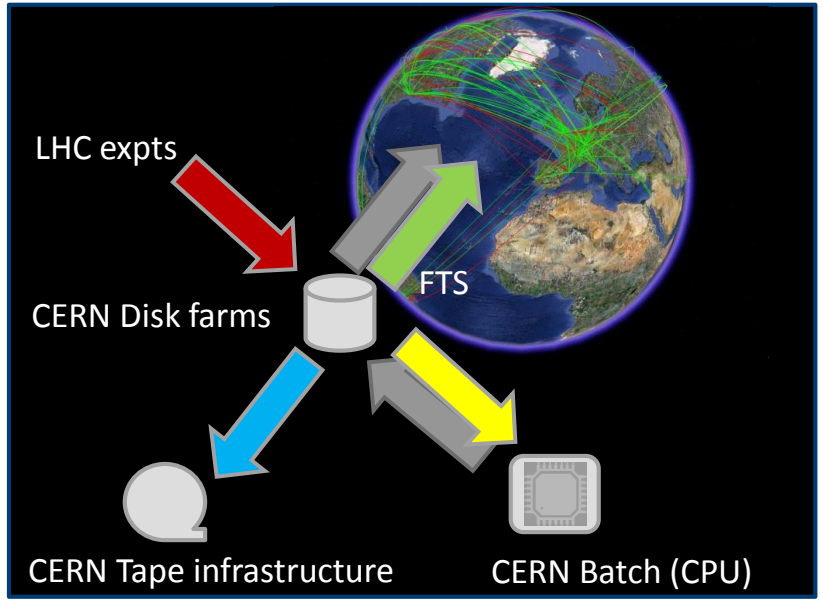

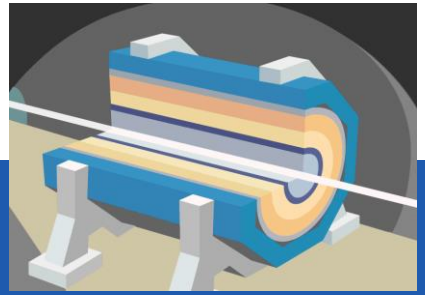

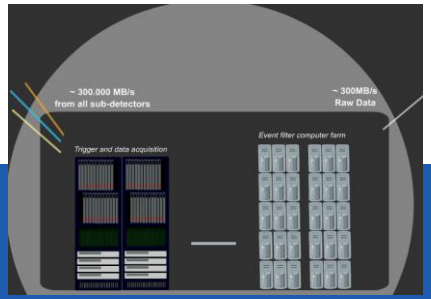

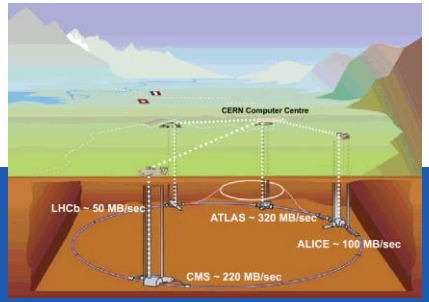

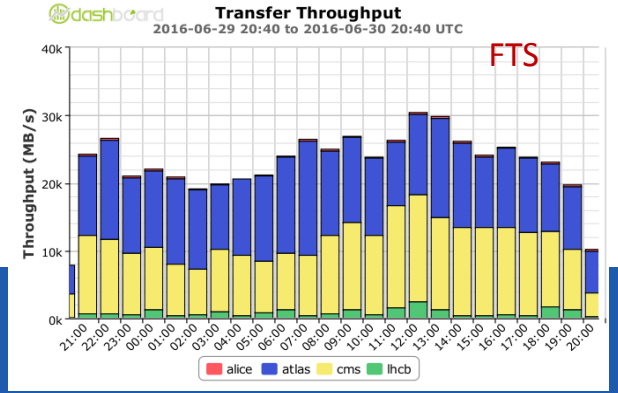

Worldwide LHC Computing Grid

Tier-0:
data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, user analysis

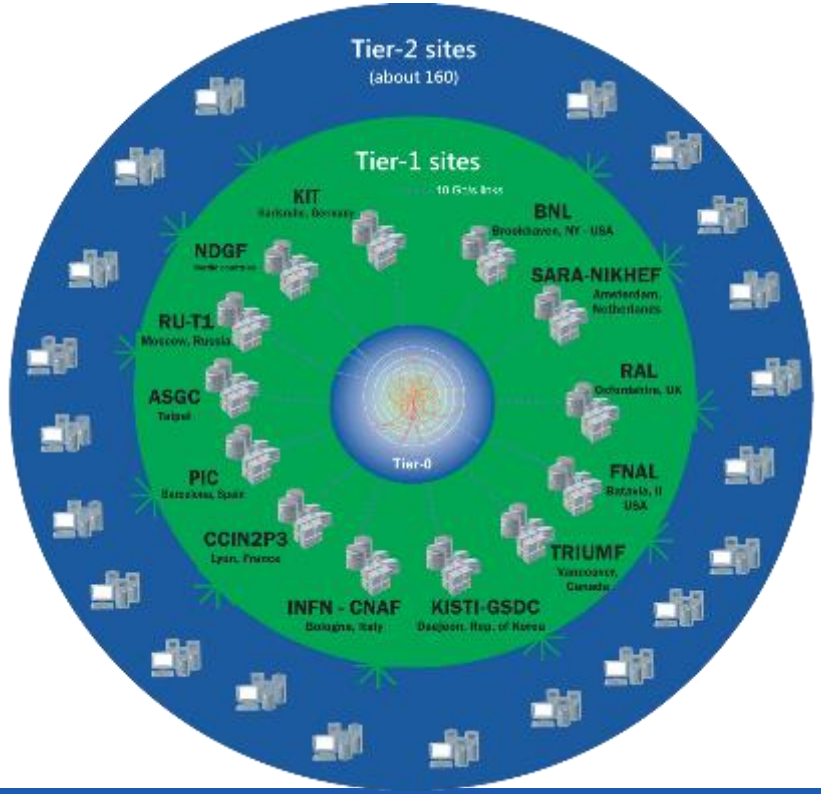

~170 sites, 40 countries

~500k CPU cores

500 PB of storage

2+ million jobs/day

Multiple 10-100 Gb links

LCG:
Initial description: 2001
Tech. Design Report: 2005

CERN data centre (main room)

CERN data centre

Overview: Data Centre

MEYRIN DATA CENTRE

	last_value
Number of Cores in Meyrin	146,156
Number of Drives in Meyrin	83,433
Number of 10G NIC in Meyrin	8,763
Number of 1G NIC in Meyrin	22,258
Number of Processors in Meyrin	24,629
Number of Servers in Meyrin	13,116
Total Disk Space in Meyrin (TB)	163,652
Total Memory Capacity in Meyrin (TB)	593

WIGNER DATA CENTRE

	last_value
Number of Cores in Wigner	56,000
Number of Drives in Wigner	29,694
Number of 10G NIC in Wigner	2,981
Numer of 1G NIC in Wigner	6,579
Number of Processors in Wigner	7,002
Number of Servers in Wigner	3,504
Total Disk Space in Wigner (TB)	97,315
Total Memory Capacity in Wigner (TB)	221

NETWORK AND STORAGE

	last_value
Tape Drives	104
Tape Cartridges	20,517
Data Volume on Tape (TB)	150,321
Free Space on Tape (TB)	33,600
Routers (GPN)	140
Routers (TN)	30
Routers (Others)	106
Switches	3,695

BATCH JOBS (#)

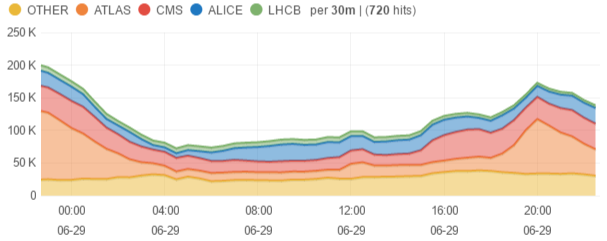

EOS ACTIVE DATA TRANSFERS (#)

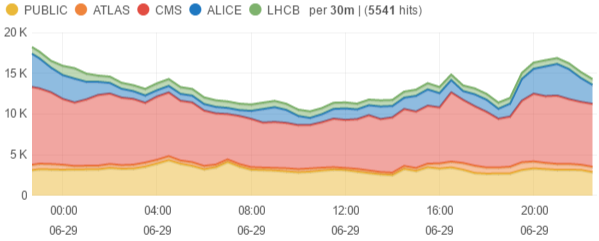

FILE TRANSFER THROUGHPUT (GB/S)

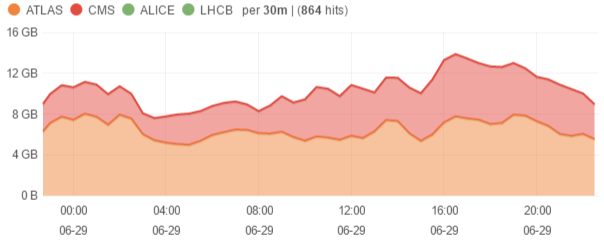

DATABASE TRANSACTIONS PER SECOND

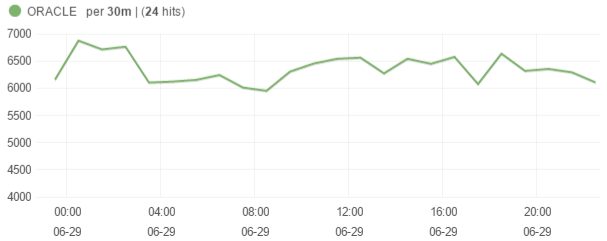

VIRTUAL MACHINES CREATED (#)

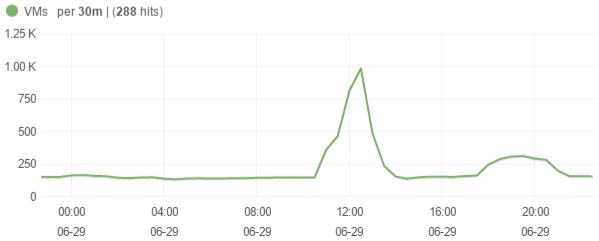

WIGNER NETWORK LINKS (GBIT/S)

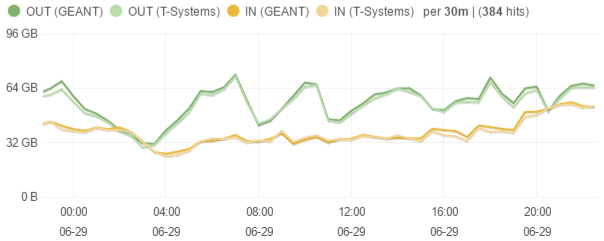

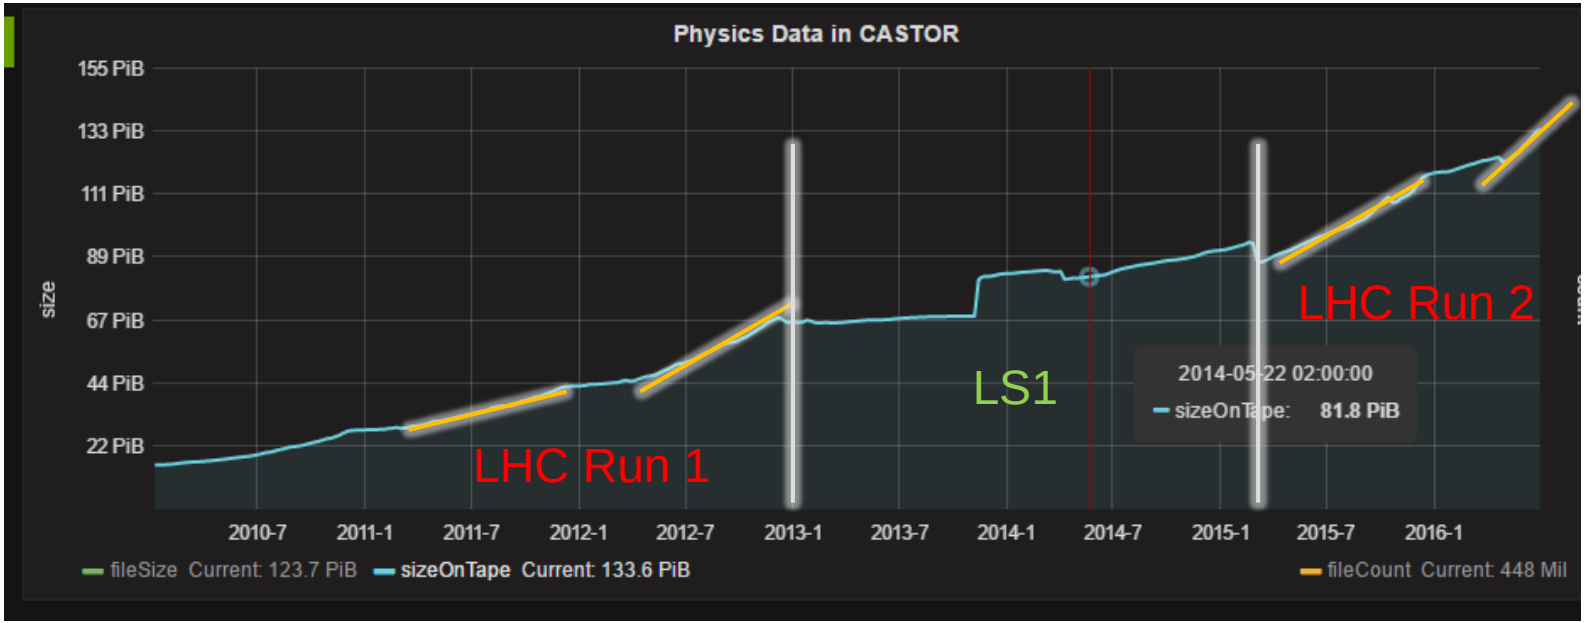

CASTOR
CERN Advanced STORage manager

N.B. Last month LHC
Recorded 10 PB (in a month!)

EOS

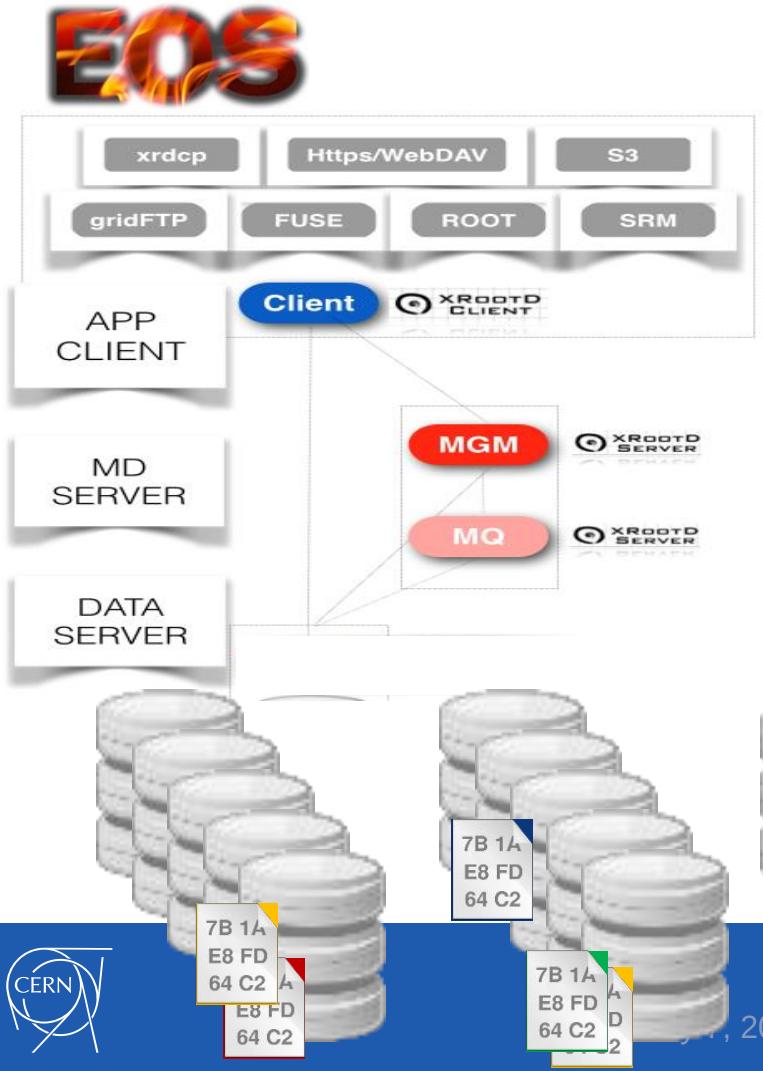

- **EOS: Large disk farms for physics and beyond**
  - Developed at CERN
  - LHC: PBs for 100s/1000s independent scientists
  - 200 PB JBOD installed (CERN installations)

- **Strategic points**
  - Distill **20+ years** of experience data management
  - Ultra-fast name space
  - Arbitrary level of data durability: cross-node file replication or RAIN on **commodity** hardware
  - Optimised **protocols**

Our “20-ms-large” computer centre

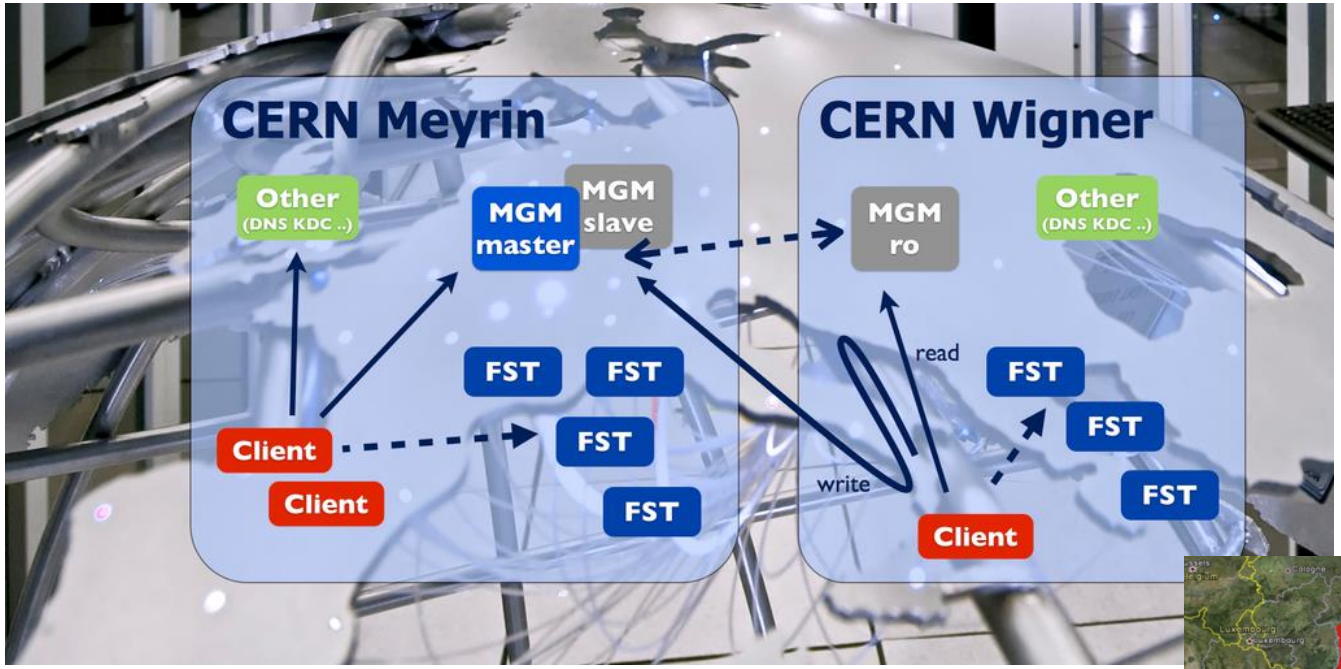

Autonomic,
Locality,
Disaster recovery/
business continuity

MGM = NameSpace/Metadata

FST = Disk servers

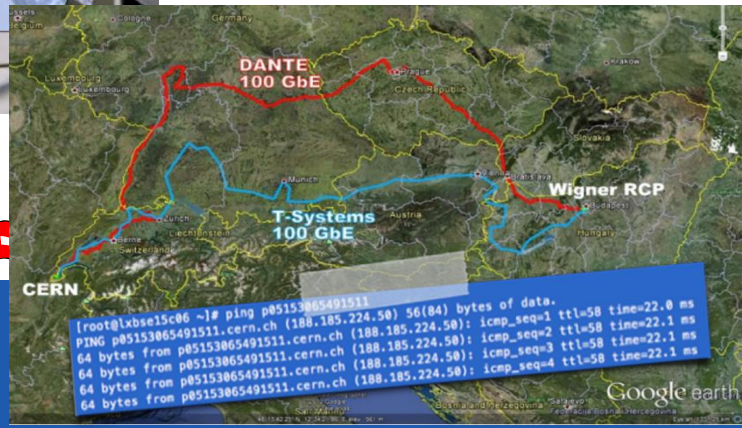

Linking the CERN Data Centres

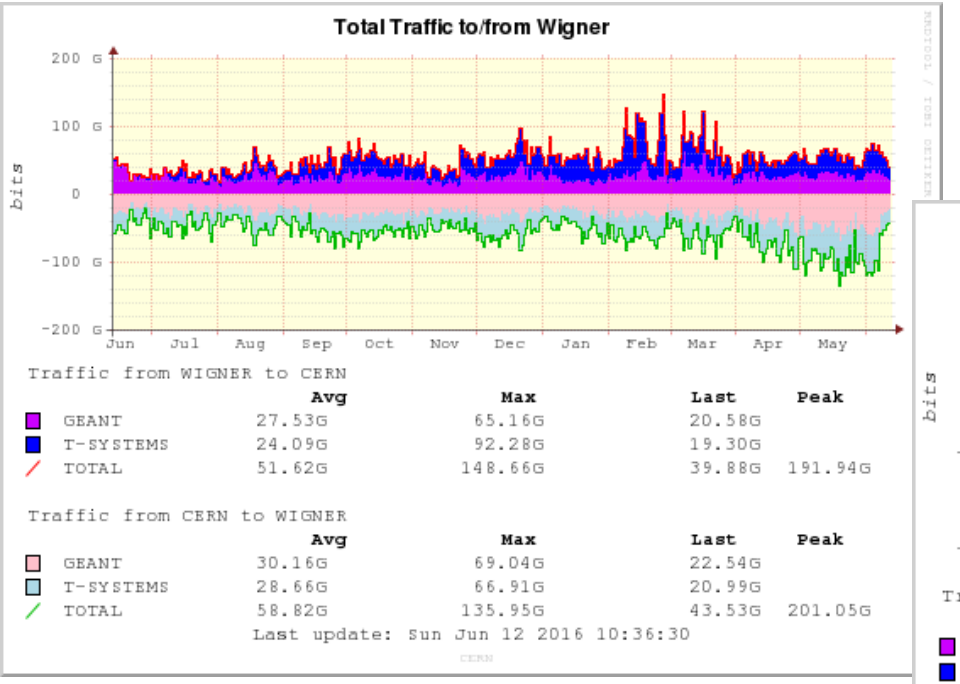

Traffic from WIGNER to CERN

	Avg	Max	Last	Peak
GEANT	27.53G	65.16G	20.58G	
T-SYSTEMS	24.09G	92.28G	19.30G	
TOTAL	51.62G	148.66G	39.88G	191.94G

Traffic from CERN to WIGNER

	Avg	Max	Last	Peak
GEANT	30.16G	69.04G	22.54G	
T-SYSTEMS	28.66G	66.91G	20.99G	
TOTAL	58.82G	135.95G	43.53G	201.05G

Last update: Sun Jun 12 2016 10:36:30

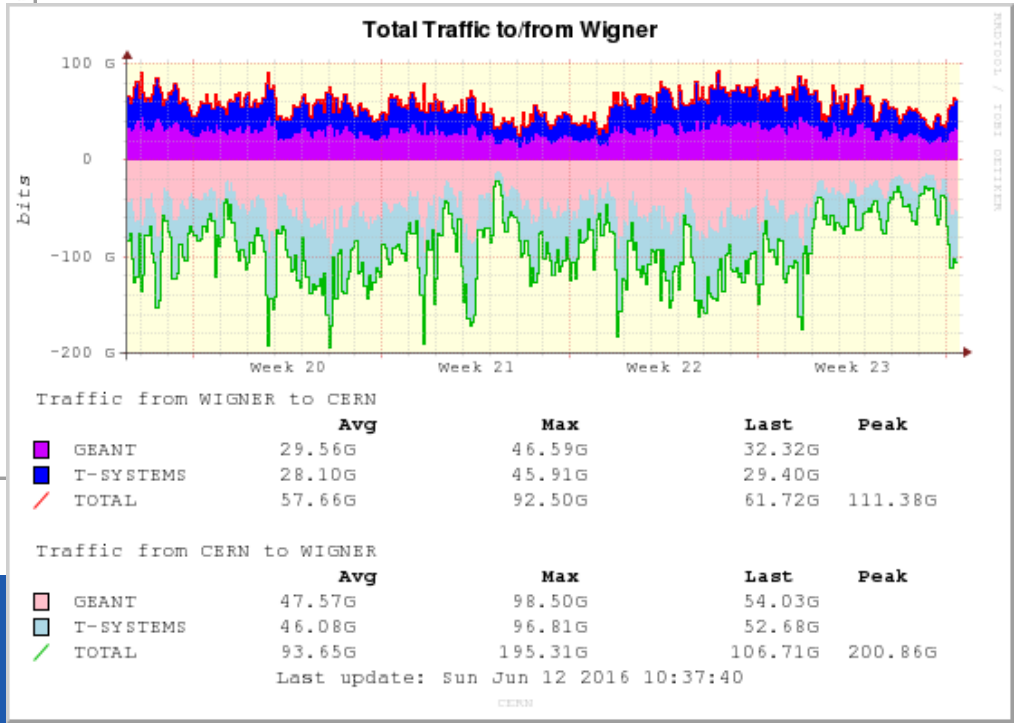

Traffic from WIGNER to CERN

	Avg	Max	Last	Peak
GEANT	29.56G	46.59G	32.32G	
T-SYSTEMS	28.10G	45.91G	29.40G	
TOTAL	57.66G	92.50G	61.72G	111.38G

Traffic from CERN to WIGNER

	Avg	Max	Last	Peak
GEANT	47.57G	98.50G	54.03G	
T-SYSTEMS	46.08G	96.81G	52.68G	
TOTAL	93.65G	195.31G	106.71G	200.86G

Last update: Sun Jun 12 2016 10:37:40

Data Rates during Run2

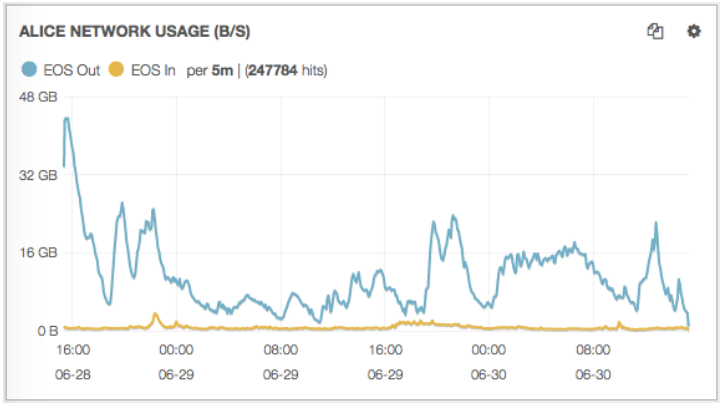

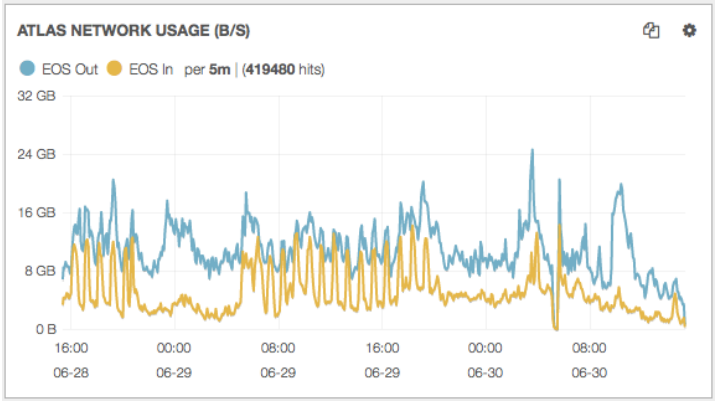

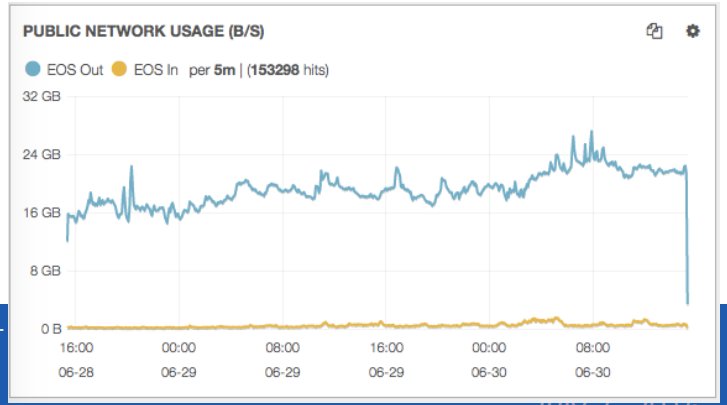

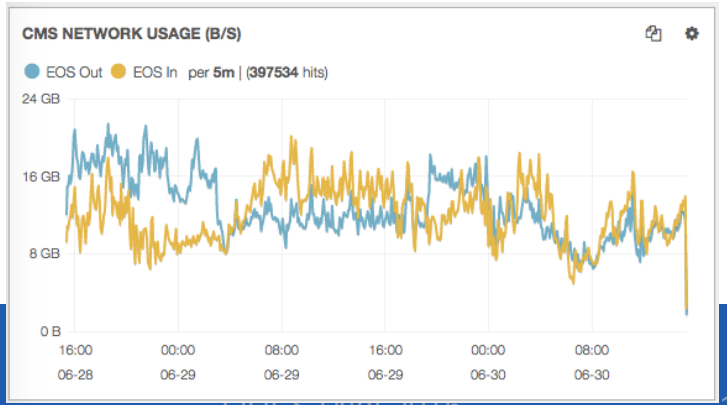

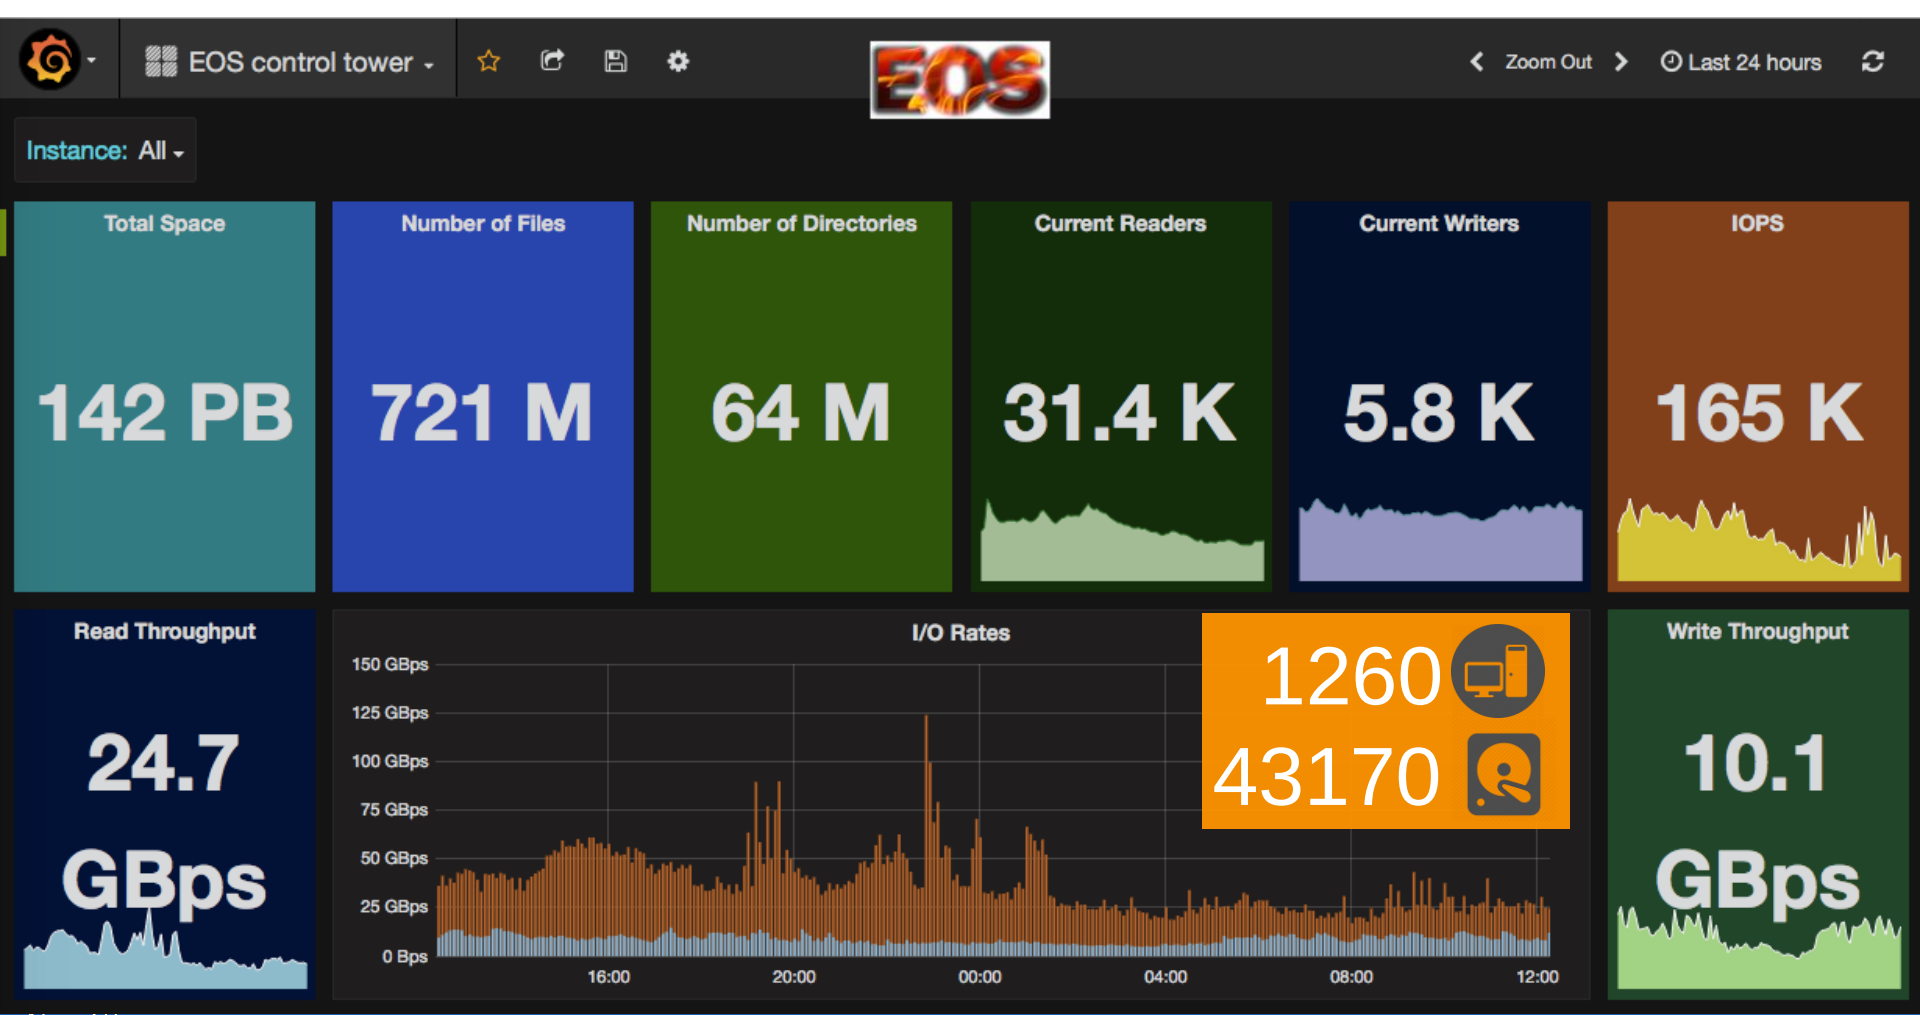
EOS control tower
Zoom Out
Last 24 hours
Instance: All
Total Space
142 PB
Number of Files
721 M
Number of Directories
64 M
Current Readers
31.4 K
Current Writers
5.8 K
IOPS
165 K
Read Throughput
24.7
GBps
I/O Rates
150 GBps
125 GBps
100 GBps
75 GBps
50 GBps
25 GBps
0 Bps
16:00
20:00
00:00
04:00
08:00
12:00
1260
43170
Write Throughput
10.1
GBps

CERNBox

Rationale behind cloud storage

- Existing professional usage of public cloud storage (Dropbox) at CERN
  - CERN (IT)
    - Harmonically extend our service portfolio
  - Today users expect:
    - Cross-platform, easy install and setup (BYOD)
    - Modern web interfaces
    - Mobile access

- Selected ownCloud as starting point
  - Open software
  - Data on premise
  - General public use and deployments in edu/research

- **Scale up in size**: interesting challenge (leverage on EOS)
- **Extend in the use-case phase space**: innovation

Access Methods: Sync

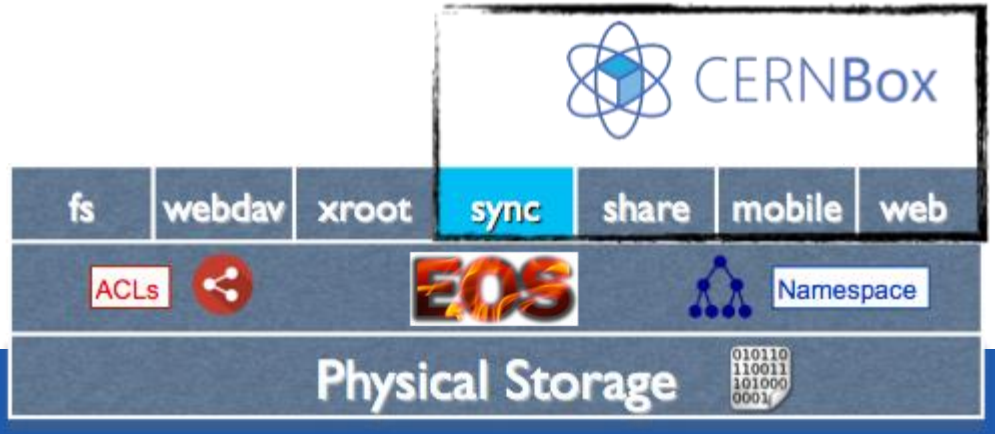

Access Methods: Sharing

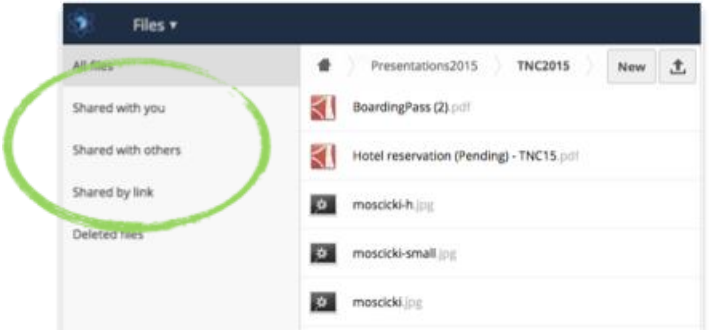

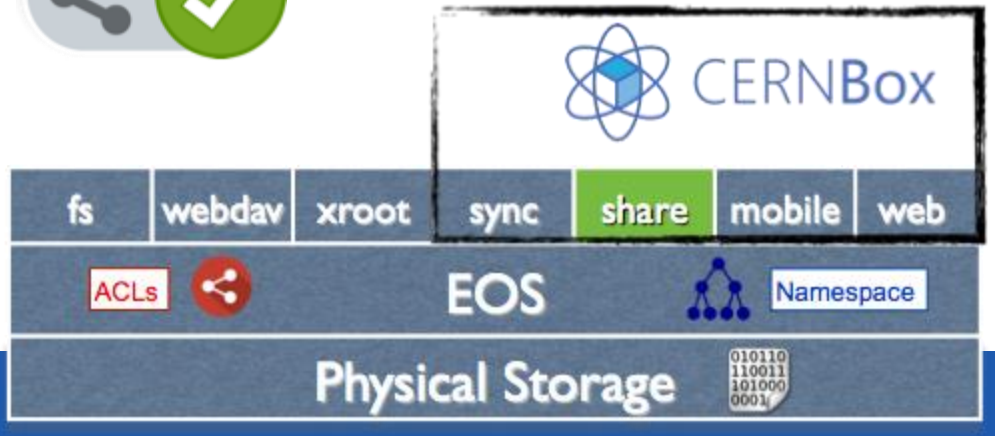

Access Methods: Mobile & Web

Access Methods: WebDAV

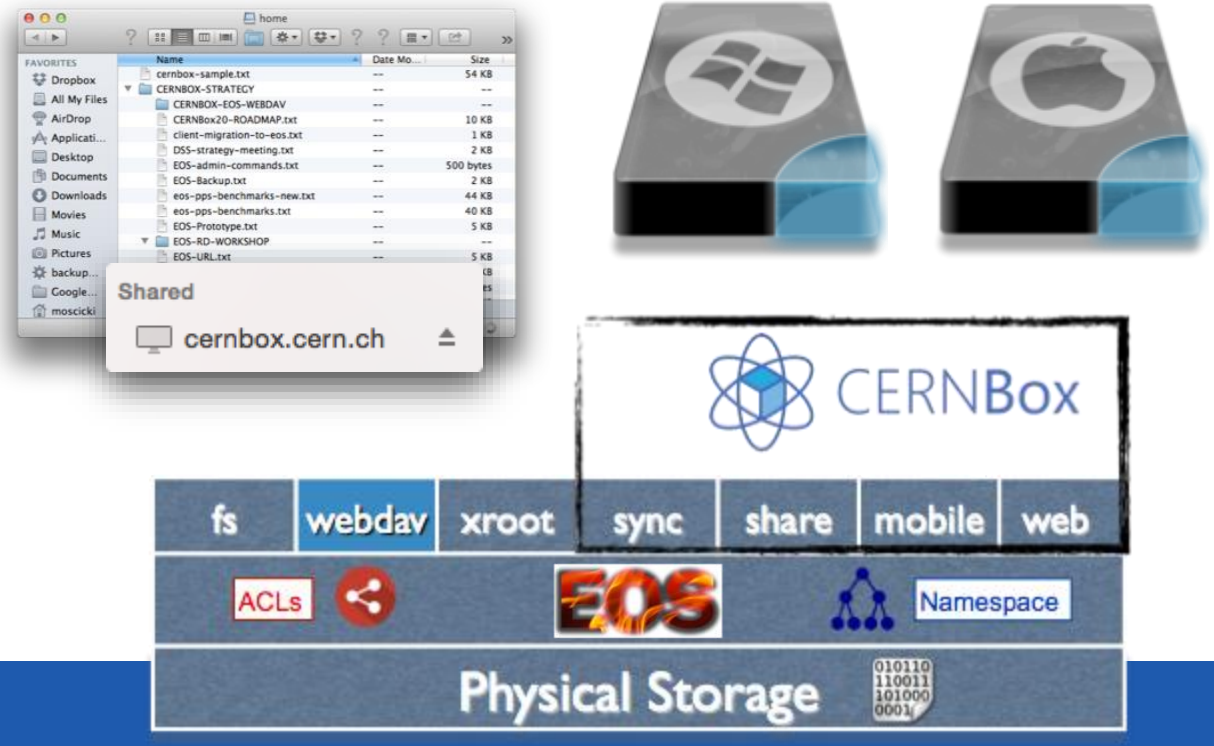

Access Methods: FUSE


```
[lmascett@lxplus2015 ~]#
[lmascett@lxplus2015 ~]# df -H -t fuse
Filesystem      Size  Used Avail Use% Mounted on
eosuser         506T   70T  437T  14% /eos/user
eosatlas         36P   17P   20P  45% /eos/atlas
eosalice         20P   11P  8.5P  57% /eos/alice
eoscms           28P   14P   15P  49% /eos/cms
eoslhcb          13P  7.6P  4.6P  63% /eos/lhcb
eospublic        16P  5.8P   11P  36% /eos/public
[lmascett@lxplus2015 ~]#
[lmascett@lxplus2015 ~]# ls -lc /eos/user/l/lmascett/
total 6644
drwx------. 1 lmascett c3       5 Dec 10 15:58 CERN
drwx------. 1 lmascett c3       0 Jan 26 18:18 debug
drwx------. 1 lmascett c3       0 Dec 11 09:43 download
drwx------. 1 lmascett c3       0 Oct 31 18:24 pdf
drwx------. 1 lmascett c3       1 Dec 11 09:44 personal
drwx------. 1 lmascett c3       8 Dec 10 12:11 pictures
```

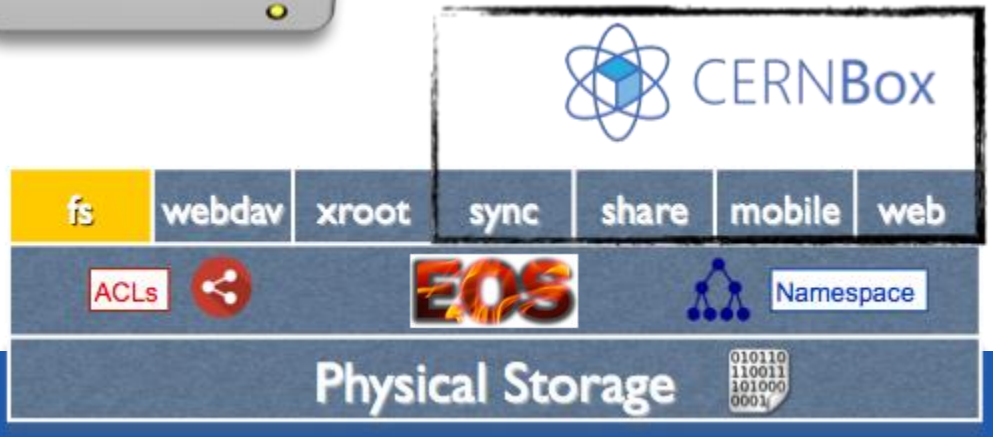

Optimised access

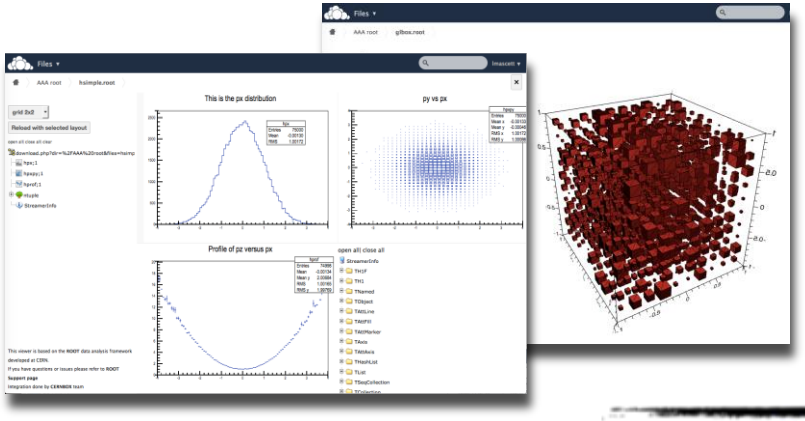

Embedded ROOT viewer
in CERNBox browser

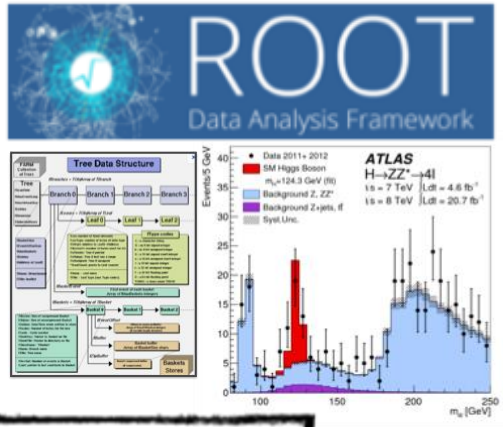

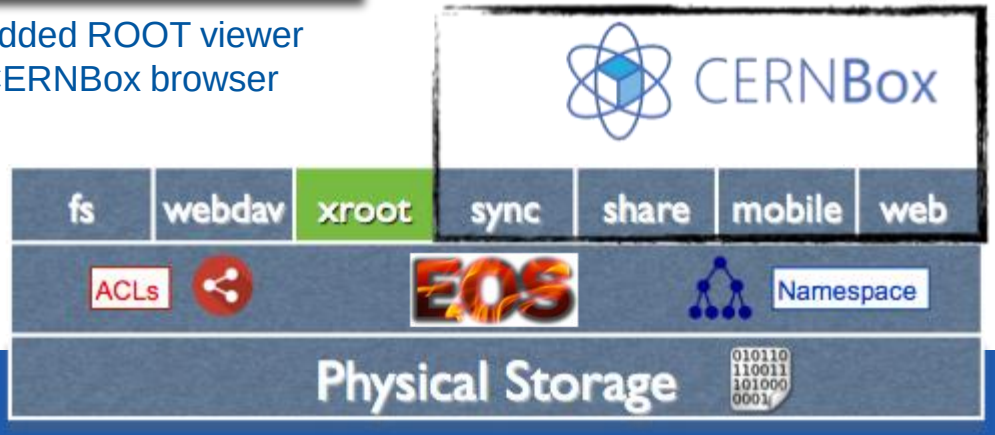

EOSUSER/CERNBox Numbers

EOS offers “virtually unlimited” cloud-storage for our end-users

Users	***5612***
# files	***83 Million***
# dirs	***11 Million***
Quota	***1TB/user***
Used Space	***173 TB***
Deployed Space	***1.3 PB***

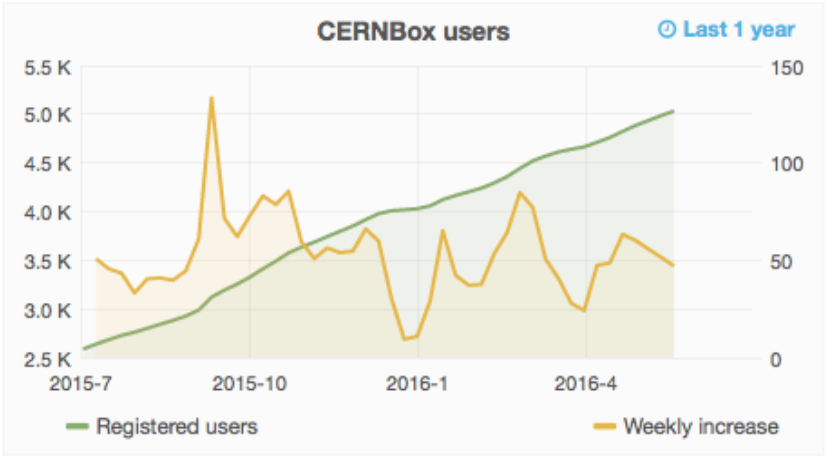

20% 60% 20%

NB: all batch is Linux

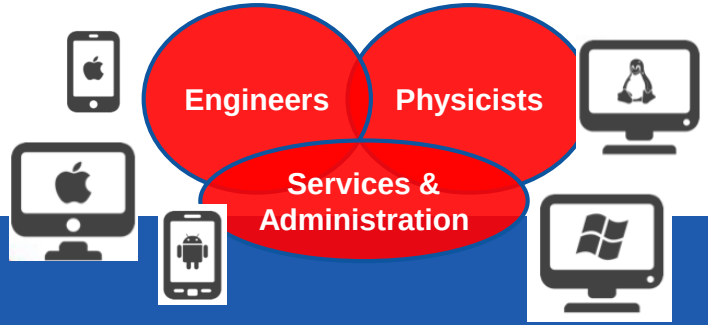

EOS/CERNBox HTTP Operations

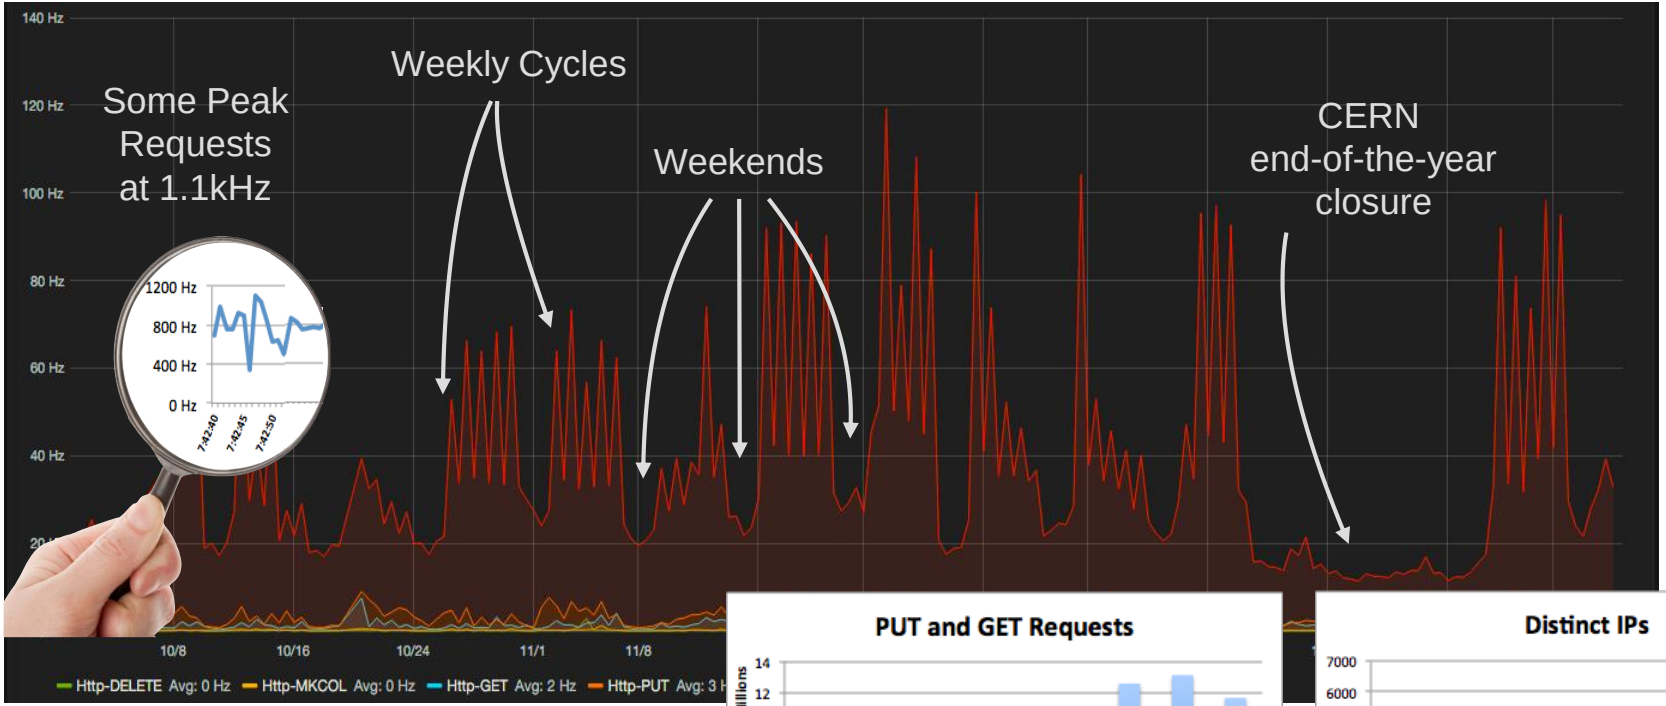

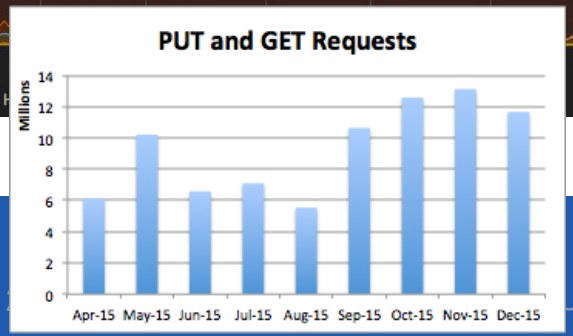

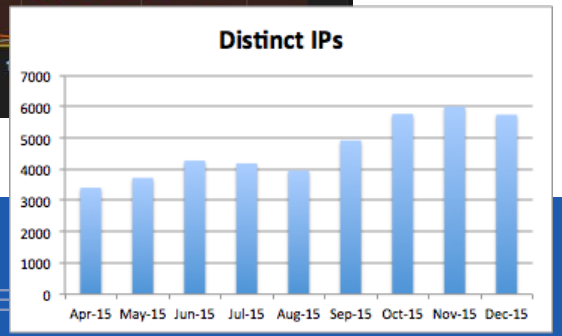

Dec 2015: Geolocation Active Users

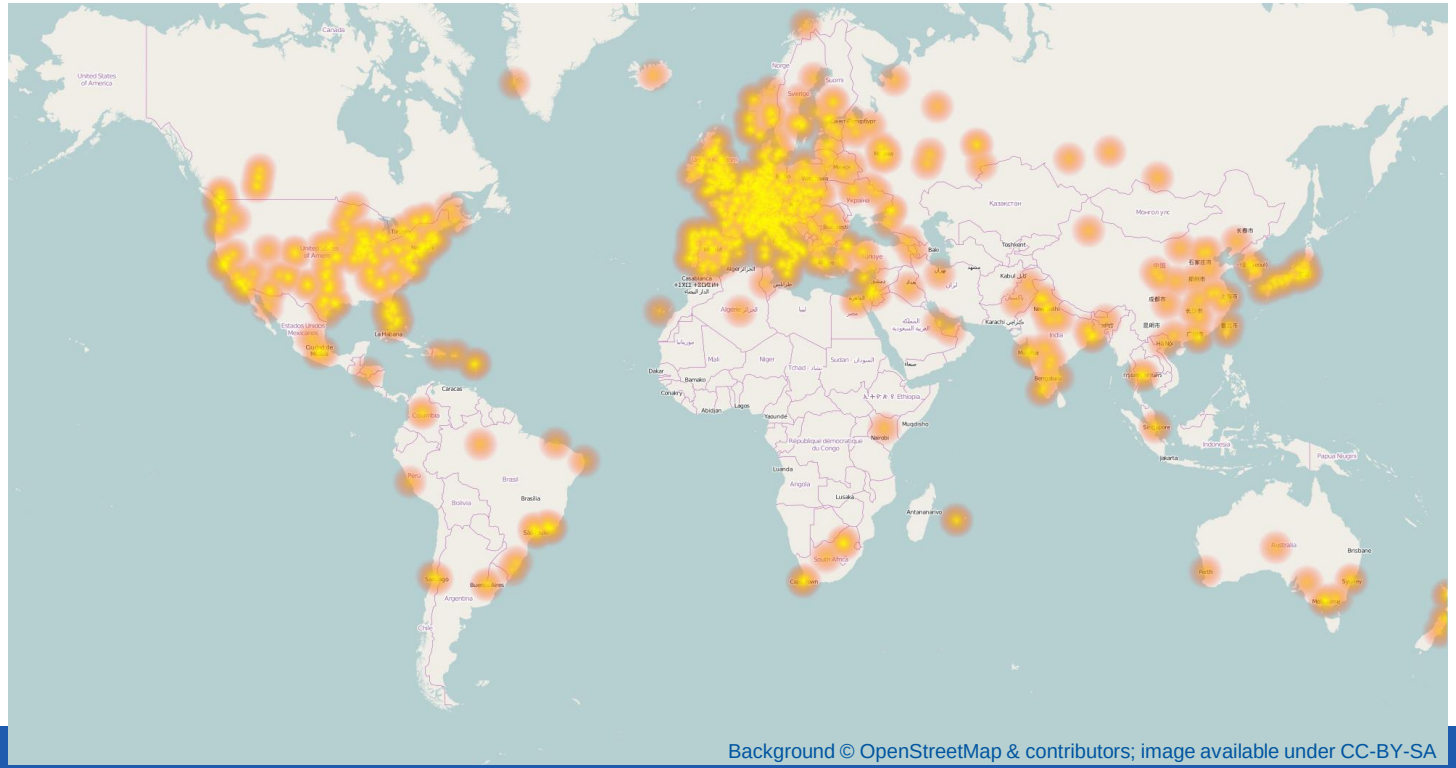

Collaborations

CERN data at your fingertips

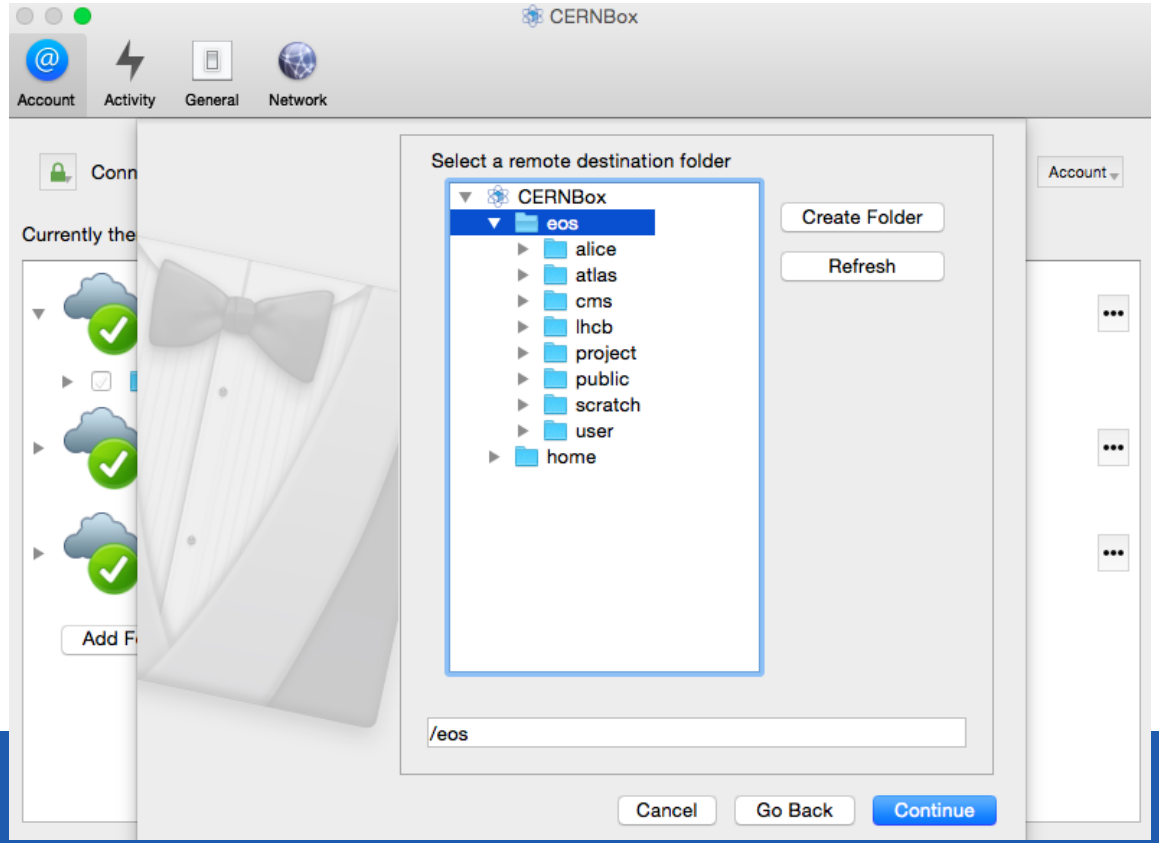

AARNET collaboration

More EOS deployments: “Exploring the 300 ms region”…

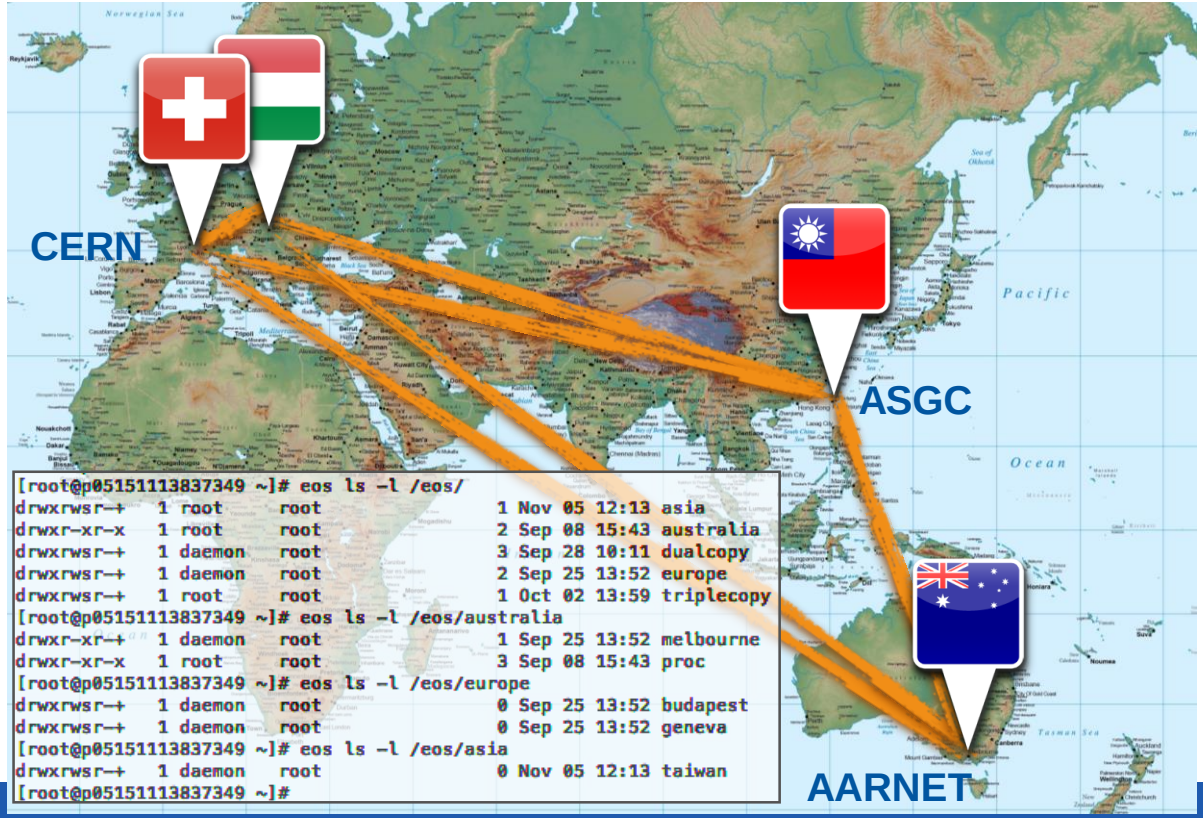

D. Jericho (AARNET),L. Mascetti (CERN),
Asa Hsu (ASGC Taipei)

R&D - EOS World-Wide Deployment

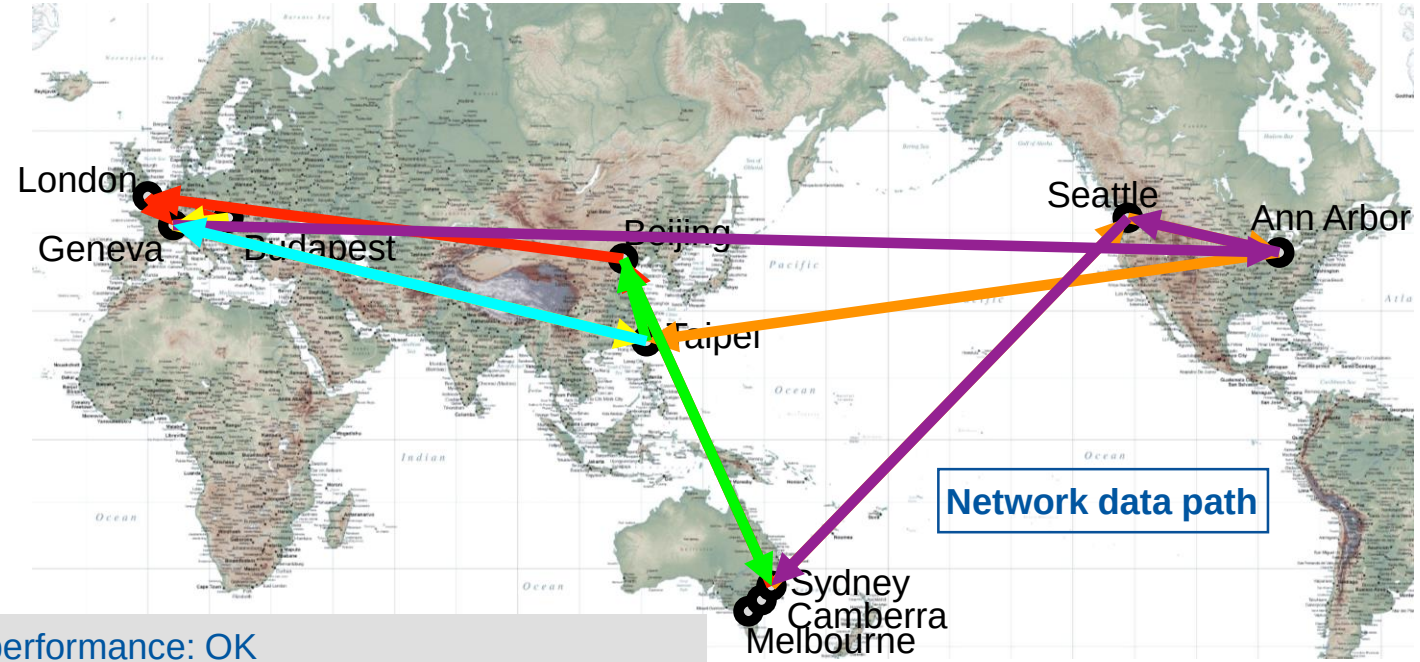

- Streaming performance: OK
- Possible problems in case of packet drops (tcp window)
  - TCP settings could be optimised
- Latency in read hidden by the read-only NS
- Latency in write to contact the read-write NS

JRC collaboration

The Joint Research Centre (JRC)

- JRC is the science service of the European Commission
- JRC provides independent scientific support to EU policy making
- Wide usage of **Earth Observation [EO] data** as basis for research and policy support

"Earth Observation & Social Sensing Big Data Pilot Project"

- The EU ***Copernicus*** Programme with the ***Sentinel*** fleet of satellites acts as a game changer by bringing EO in the Big Data era:
  - *expected 10TB/day of **free and open** data*
  - *Requires new approaches for data management and processing*
- Pilot project launched in January 2015
- Major goal: set up a central infrastructure for storing and processing of Earth Observation and Social Sensing data at JRC

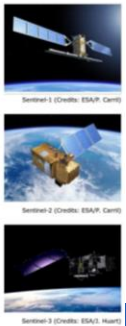

Joint Research Centre

European Commission

Proposal for a "JRC Earth Observation Data Processing Platform" (JEO-DPP)

- Main focus on **satellite image** data
- Shall support existing processing workflows and environments (C/C++, Python, Matlab, Java)
- Provide different processing levels:
  - *Low-level batch processing*
  - *High-level interactive processing*
- Project timeline:
  - *Prototype development: end 2015 – mid 2017*
  - *Scaling-up in 2017/18: JRC Data Centre vs a public cloud solution*

EOS

Joint Research Centre

European Commission

United Nations Institute for Training and Research

UNOSAT
UNITAR Operational Satellite Applications Programme

RESEARCH FOUNDATION
cima
OBSERVE TO PREDICT, PREDICT TO PREVENT

CERNBox

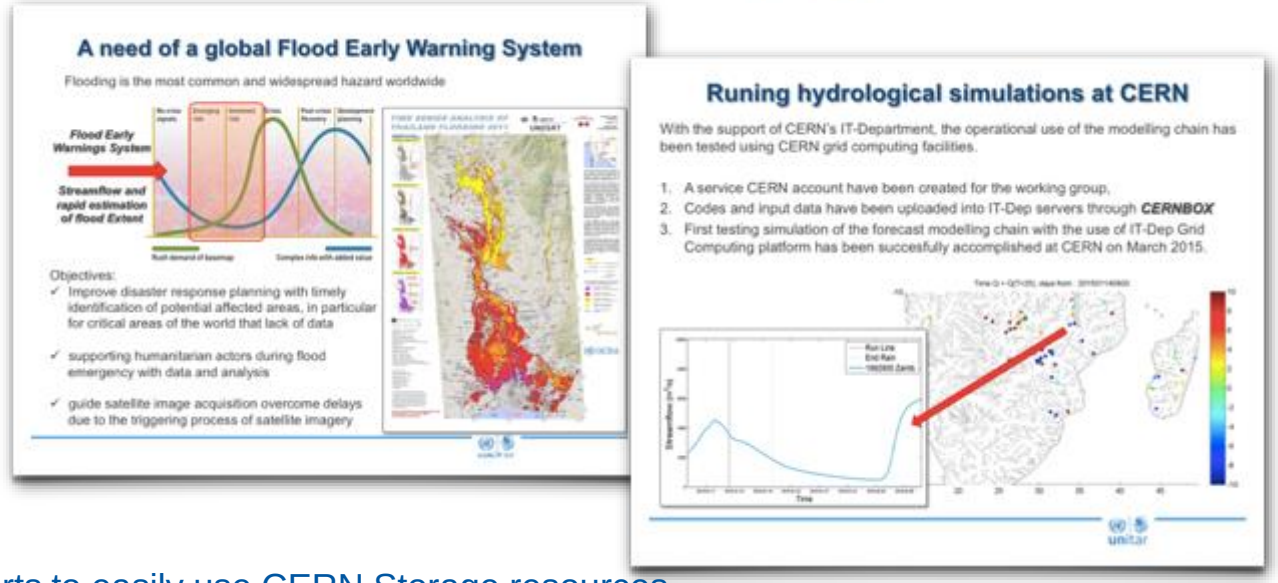

- Enable non-experts to easily use CERN Storage resources
- Powerful integration with the batch system
- Simple to share result with collaborators

Thanks to Mauro Arcorace, members of **UNITAR/UNOSAT** and **CIMA foundation** for the material provided

EOS and COMTRADE

- Streamline EOS
  - More platforms
  - More “generic” installations
    - Remove CERN dependencies
  - More documentation

- Collaboration with Comtrade
  - 24-month project (started in July 2015)
  - http://storage.comtrade.com
  - Company based in Serbia
  - Technology interest
  - Support opportunity for EOS
    - N.B. EOS is open software

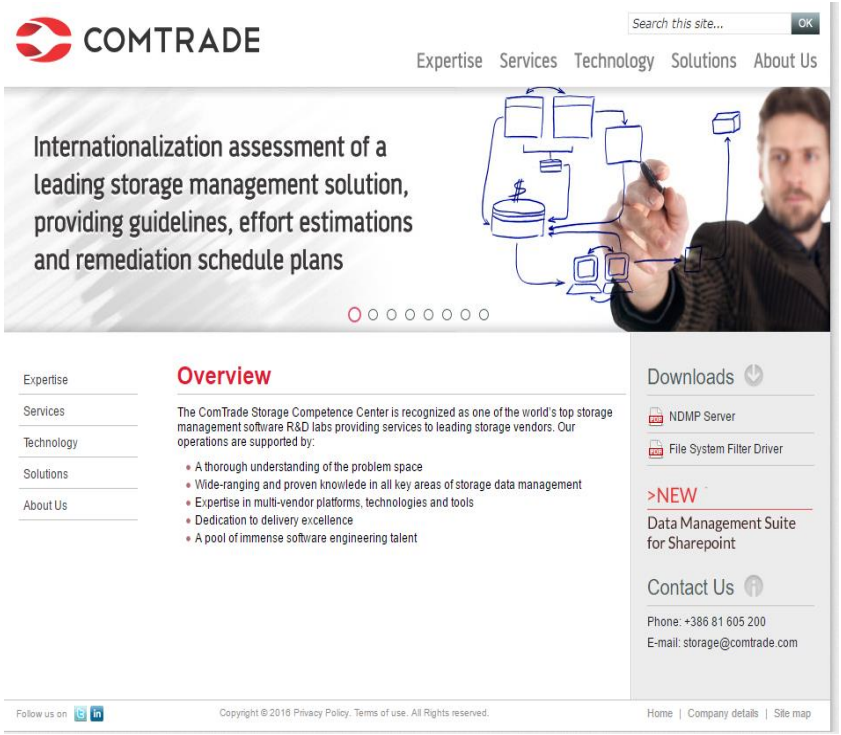

Interconnected Private Clouds for Universities and Researchers

CS3 Workshops

Cloud Services for Synchronisation and Sharing

CS3 in ETH Zurich (Jan 2016):

- 90 participants
- 40 institutions
- 33 contributions
- 16 countries

See you at the 3rd CS3 in Amsterdam!

https://cs3.surfsara.nl

Novel applications?

Notebook analysis
SWAN project
with CERN Physics Department

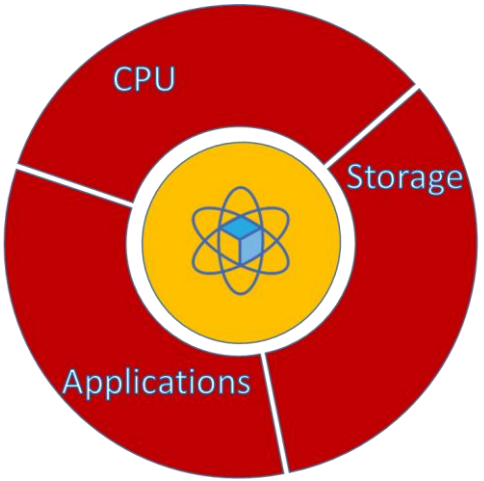

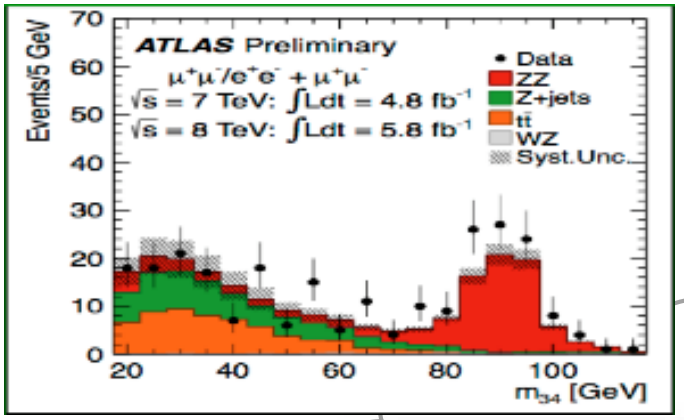

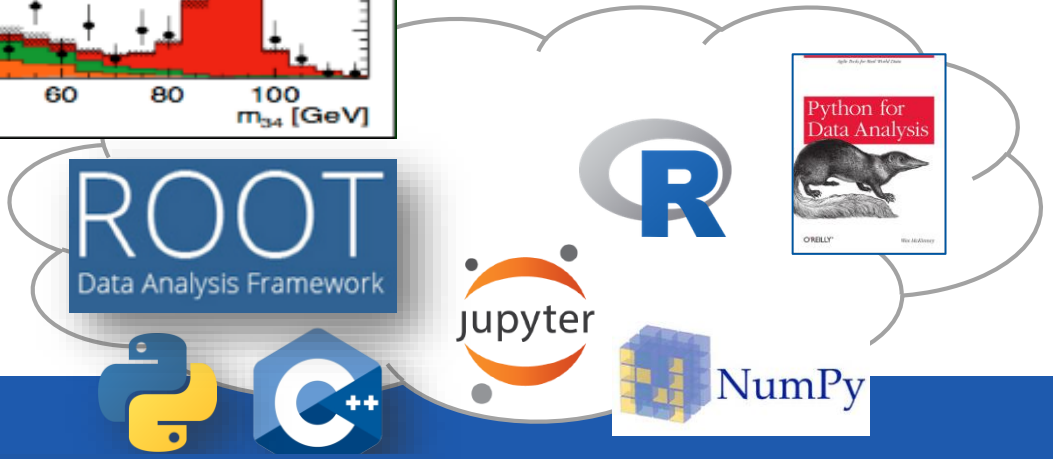

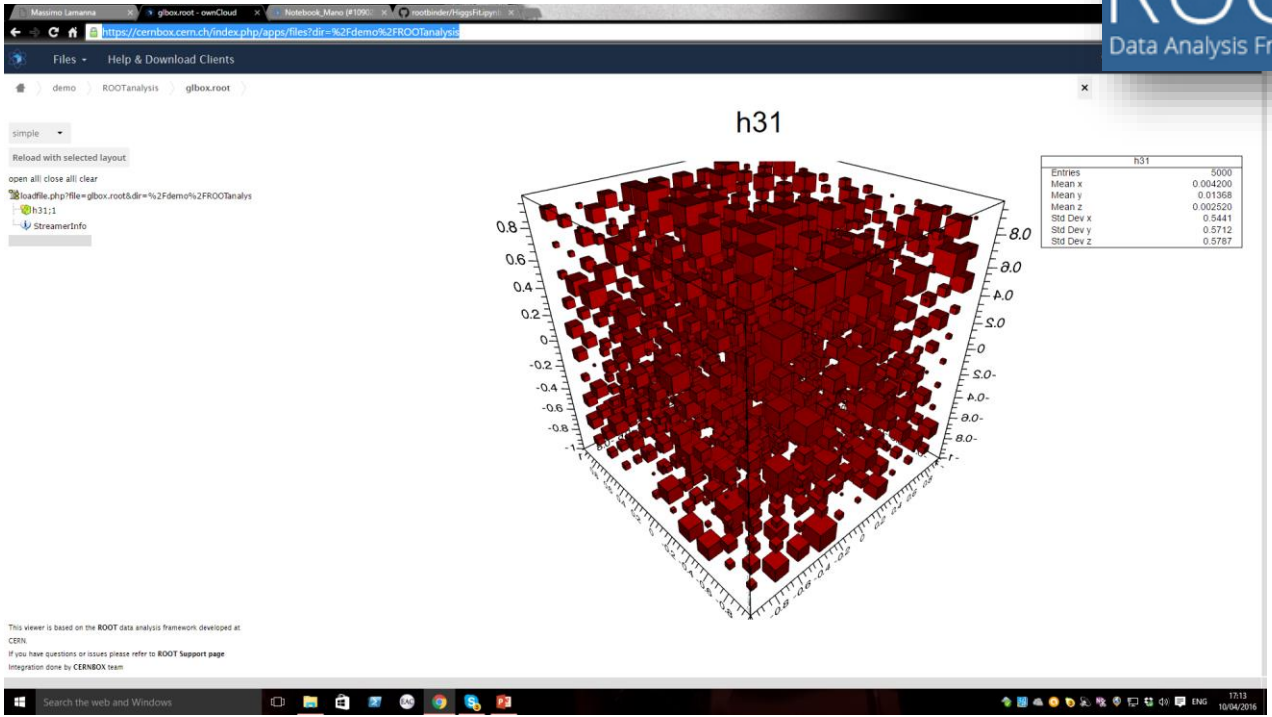

Traditional data analysis → Web-based analysis

Massimo Lamanna | glbox.root - ownCloud | Notebook_Mano (#10902 | rootbinder/HiggsFit.ipynb

GitHub, Inc. [US] https://github.com/cernphsft/rootbinder/blob/master/notebooks/HiggsFit.ipynb

Full ecosystem with CERN tools/initiatives like INDICO (collaborative meeting), Zenodo (publication) and OpenData (public acces to scientific data)

Visualisation of the Result

In [8]:
```
auto plot = x->frame();
data.plotOn(plot);
model->plotOn(plot);
model->plotOn(plot, RooFit::Components("bmodel"),RooFit::LineStyle(kDashed));
model->plotOn(plot, RooFit::Components("smodel"),RooFit::LineColor(kRed));
model->paramOn(plot);

TCanvas c;
plot->Draw();
c.Draw();
```

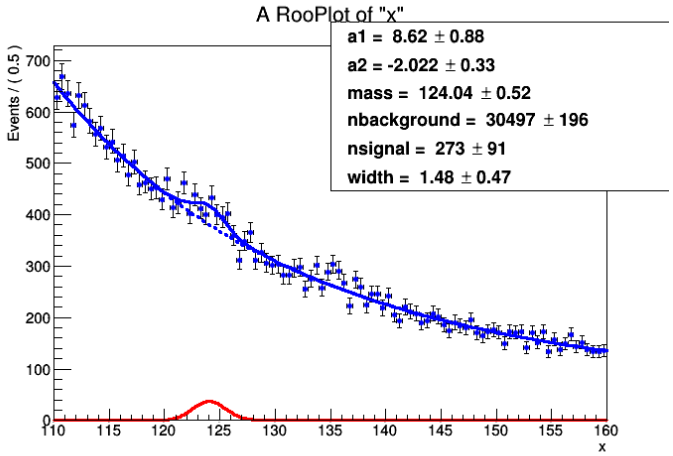


```
[#1] INFO:NumericIntegration -- RooRealIntegral::init(bmodel_Int[x]) using numeric integrator RooIntegrator1D to calculate Int(x)
[#1] INFO:Plotting -- RooAbsPdf::plotOn(model) directly selected PDF components: (bmodel)
[#1] INFO:Plotting -- RooAbsPdf::plotOn(model) indirectly selected PDF components: (z)
[#1] INFO:NumericIntegration -- RooRealIntegral::init(bmodel_Int[x]) using numeric integrator RooIntegrator1D to calculate Int(x)
[#1] INFO:Plotting -- RooAbsPdf::plotOn(model) directly selected PDF components: (smodel)
[#1] INFO:Plotting -- RooAbsPdf::plotOn(model) indirectly selected PDF components: ()
[#1] INFO:NumericIntegration -- RooRealIntegral::init(bmodel_Int[x]) using numeric integrator RooIntegrator1D to calculate Int(x)
```

New environment for physics research!

Enabled by cloud storage (CERNBox)

e.g. http://nbviewer.jupyter.org/github/dpiparo/swanExamples/blob/master/notebooks/CMSDimuon_py.ipynb

Massimo Lamanna | glbox.root - ownCloud | Notebook_Mano (#1090 | rootbinder/HiggsFit.ipynb

https://cernbox.cern.ch/index.php/apps/files/?dir=%2FNotebook_Mano%20(%2310902349)

Files | Help & Download Clients

```
import math
PI_Y.append(math.pi) #On defini la valeur de Pi pour l'axe y
PI_Y.append(math.pi) #Et pour l'axe y

c1.cd().SetLogx()

gr = TGraph( n, x, y )
gr.SetLineColor( 1 ) #On choisi la couleur de la ligne
gr.SetLineWidth( 1 ) #Son epaisseur
gr.SetMarkerColor( 4 ) #La couleur des points
gr.SetMarkerStyle( 18 ) #Et le style de point
gr.SetTitle( 'Methode de Monte Carlo' ) #On choisi le nom du graphique
gr.GetXaxis().SetTitle( 'Nombre de Points' ) #Et le nom des axes x et y
gr.GetYaxis().SetTitle( 'Valeurs estimee de Pi' )

gr.Draw( '' )

gr2=TGraph(2, PI_X, PI_Y)
gr2.SetLineColor( 2 )
gr2.SetLineWidth( 1 )

gr2.Draw("Same")

c1.Draw()
```

```
TCanvas::Constructor:0: RuntimeWarning: Deleting canvas with same name: c1
```

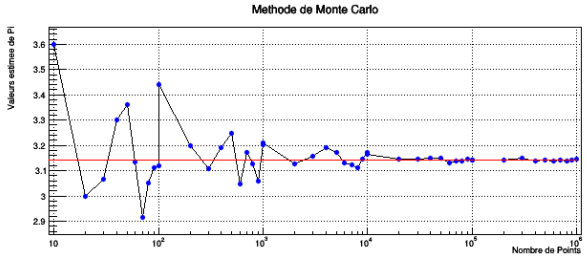

On sauvegarde la liste 1 comme ca on peut renouveler l'experience et comparer les resultats.

```
In [25]: list2 = list[:]

In [26]: list2

Out[26]: [[10, 3.6, 0.458407346410207],
          [20, 3.0, -0.14159265358979312],
          [30, 3.066666666666667, -0.07492598692312624],
```

Search the web and Windows

Training tool to get used with computing, mathematics, …

Massimo Lamanna | glbox.root - ownCloud | Notebook_Mano (#1090 | rootbinder/HiggsFit.ipynb

GitHub, Inc. [US] https://github.com/cernphsft/rootbinder/blob/master/notebooks/HiggsFit.ipynb

Higgs decay to two photons

The Standard Model predicted the decay of the Higgs bosons into photons. THe process is depicted by the diagrams below:

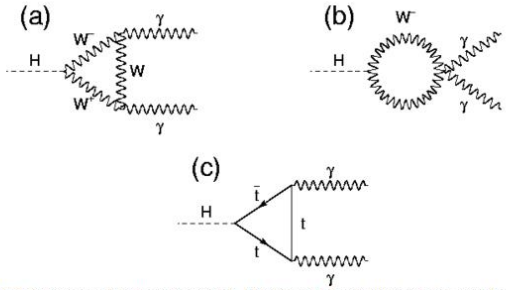

At the Large Hadron Collider, this process has been measured. This figure shows how an Higgs boson decay looks in the CMS detector:

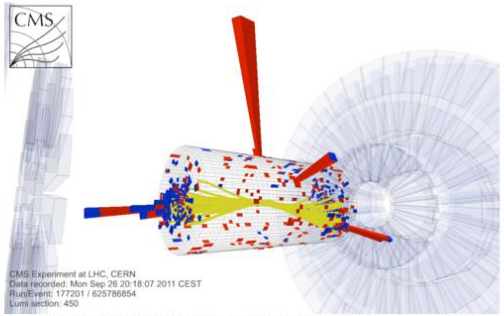

This ROOTbook illustrates a simplified fitting procedure aiming to identify the peak due to the Higgs boson decay over the exponentially falling background.

Importing input data into a ROOT file

First of all we import the input data, here simplistically stored into a text file, into a ROOT file.

In [1]:
```
TTree tree("HiggsTree","The tree cont");
auto nevt = tree.ReadFile("Hgg.txt","x");
if (nevt <= 0) {
```

Outreach to explain our research…

Search the web and Windows

Summary

- **Solid foundations**
  - 200 PB LHC disk infrastructure
    - Steadily growing!
  - Data management experience in innovation and service provision
- **Cloud storage enables new use cases**
  - and new ways to work and to collaborate
- **CERNBox/EOS**
  - Home for innovative applications
  - CS3 workshop: visit cs3.ethz.ch

DOI 10.5281/zenodo.44783


EOS
101
011

www.cern.ch