

# ISEF 2013 Students @ CERN

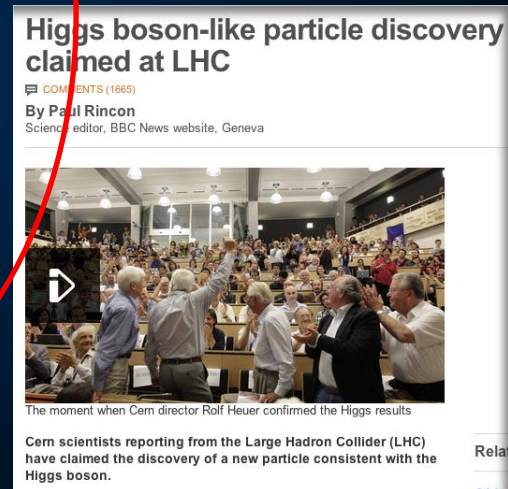

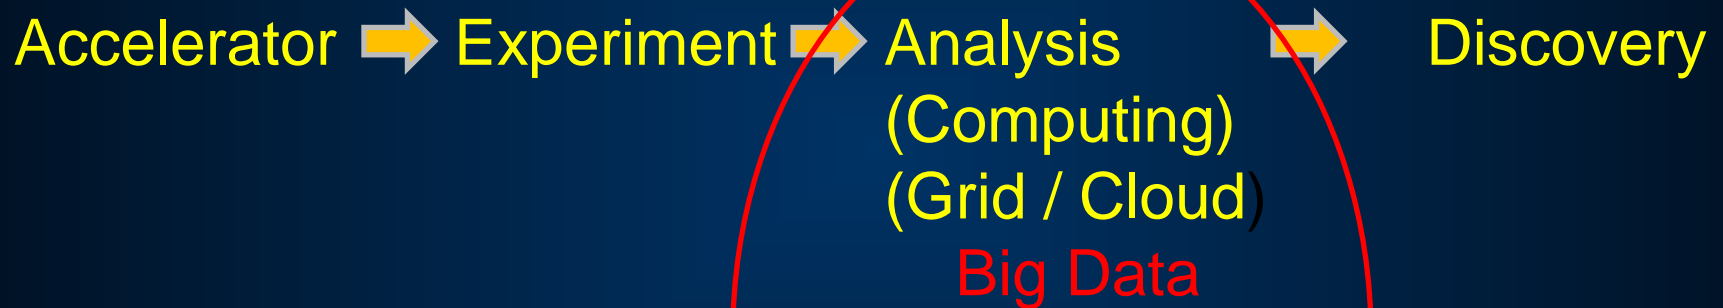
## Introduction to the CERN Data Centre(s)

Frédéric Hemmer  
Information Technology Department Head  
CERN


27 June 2013



# How to make a discovery ?













# Tier 0 at CERN: Acquisition, First pass reconstruction, Storage & Distribution





# The LHC Data Challenge

- The accelerator will run for 20 years
- Experiments are producing about **25 Million Gigabytes** of data each year (about 3 million DVDs – 850 years of movies!)
- LHC data analysis requires a computing power equivalent to **~100,000 of today's fastest PC processors**
- Requires many cooperating computer centres, as CERN can **only provide ~20% of the capacity**



# Solution: the Grid

- Use the Grid to unite computing resources of particle physics institutes around the world

The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations


The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe





# WLCG – what and why?

- A distributed computing infrastructure to provide the production and analysis environments for the LHC experiments
- Managed and operated by a worldwide collaboration between the experiments and the participating computer centres
- The resources are distributed – for funding and sociological reasons
- Our task was to make use of the resources available to us – no matter where they are located



## 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

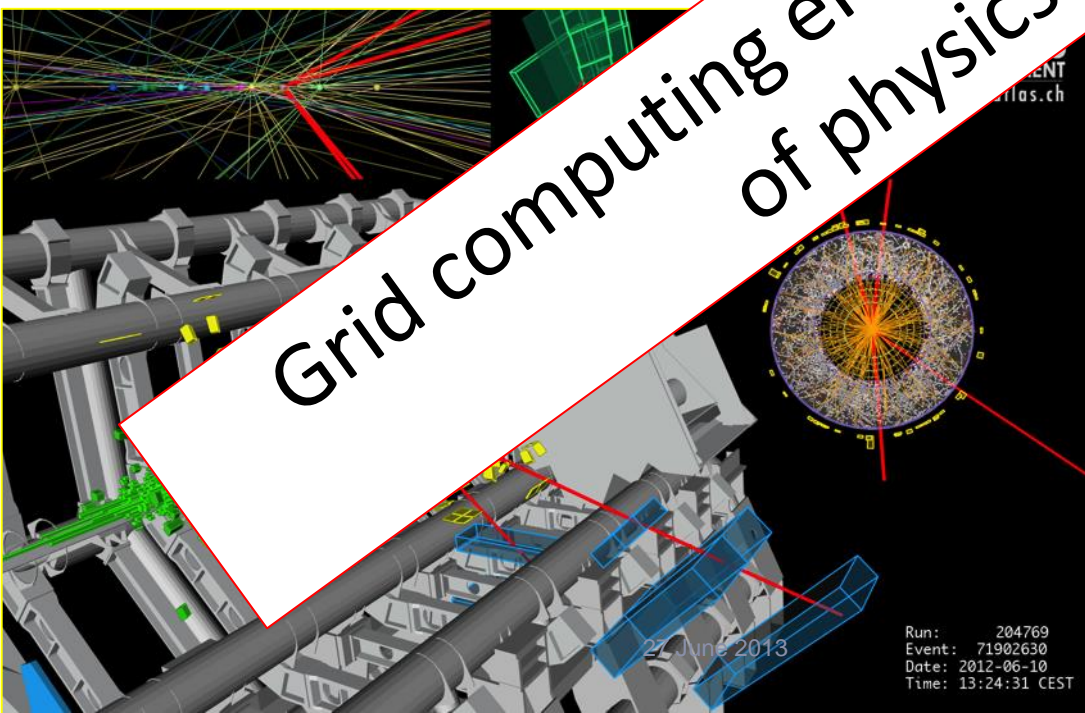
- ~ 160 sites, 35 countries
- 250000 cores
- 173 PB of storage
- 2 Million jobs/day
- 10 Gbps links



$H \rightarrow \gamma\gamma$   
candidate



Grid computing enables the rapid delivery  
of physics results



Local p<sub>0</sub>

27 June 2013

Run: 204769  
Event: 71902630  
Date: 2012-06-10  
Time: 13:24:31 CEST

## Higgs boson-like particle discovery claimed at LHC

COMMENTS (1665)

By Paul Rincon

Science editor, BBC News website, Geneva



The moment when Cern director Rolf Heuer confirmed the Higgs results

Cern scientists reporting from the Large Hadron Collider (LHC) have claimed the discovery of a new particle consistent with the Higgs boson.

Relat

Q&A



# The Data

# Some history of scale...

Date	Collaboration sizes	Data volume, archive technology
Late 1950's	2-3	Kilobits, notebooks
1960's	10-15	kB, punchcards
1970's	~35	MB, tape
1980's	~100	GB, tape, disk
1990's	700-800	TB, tape, disk
2010's	~3000	PB, tape, disk

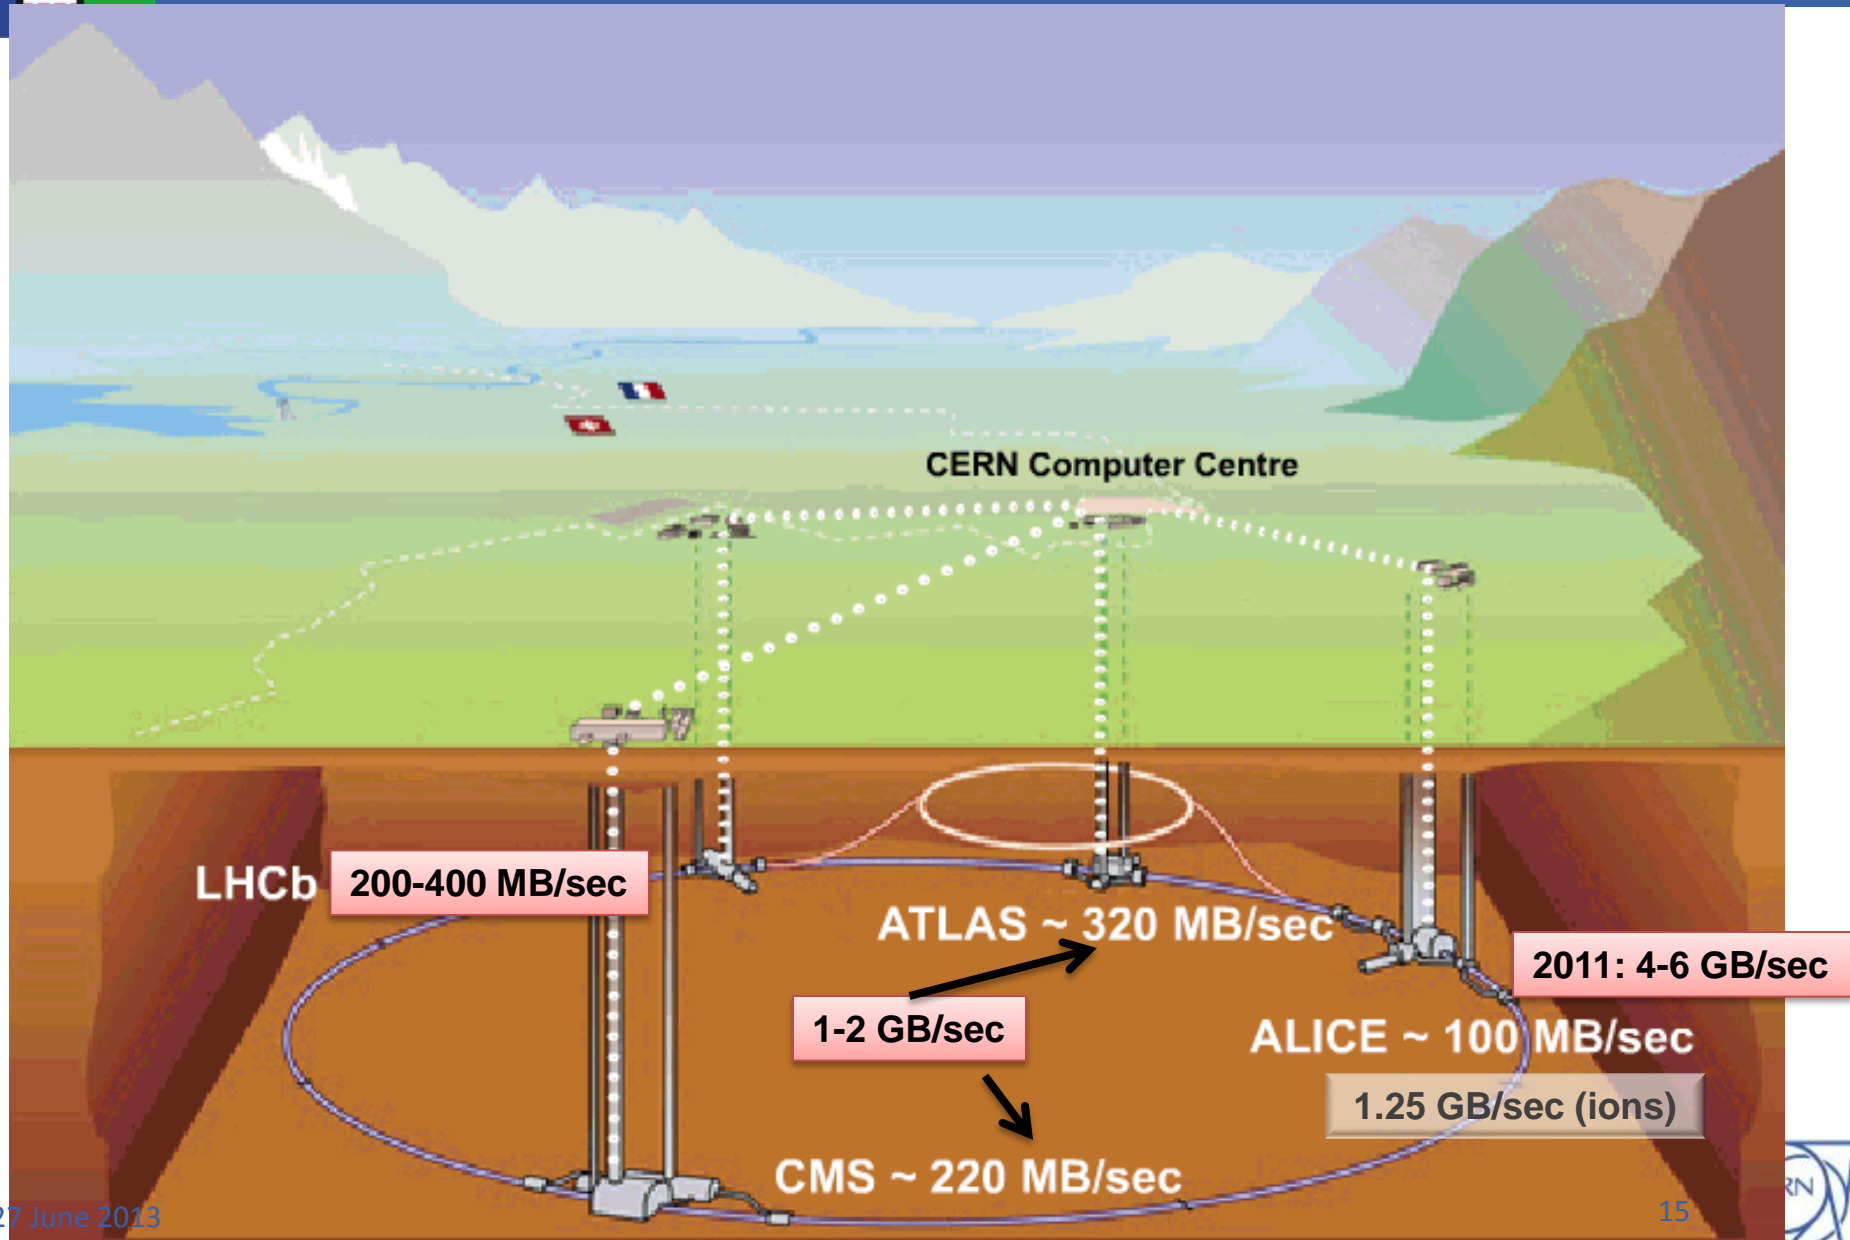
For comparison:

1990's: Total LEP data set ~few TB  
Would fit on 1 tape today


Today: 1 year of LHC data ~25 PB

## Where does all this data come from?






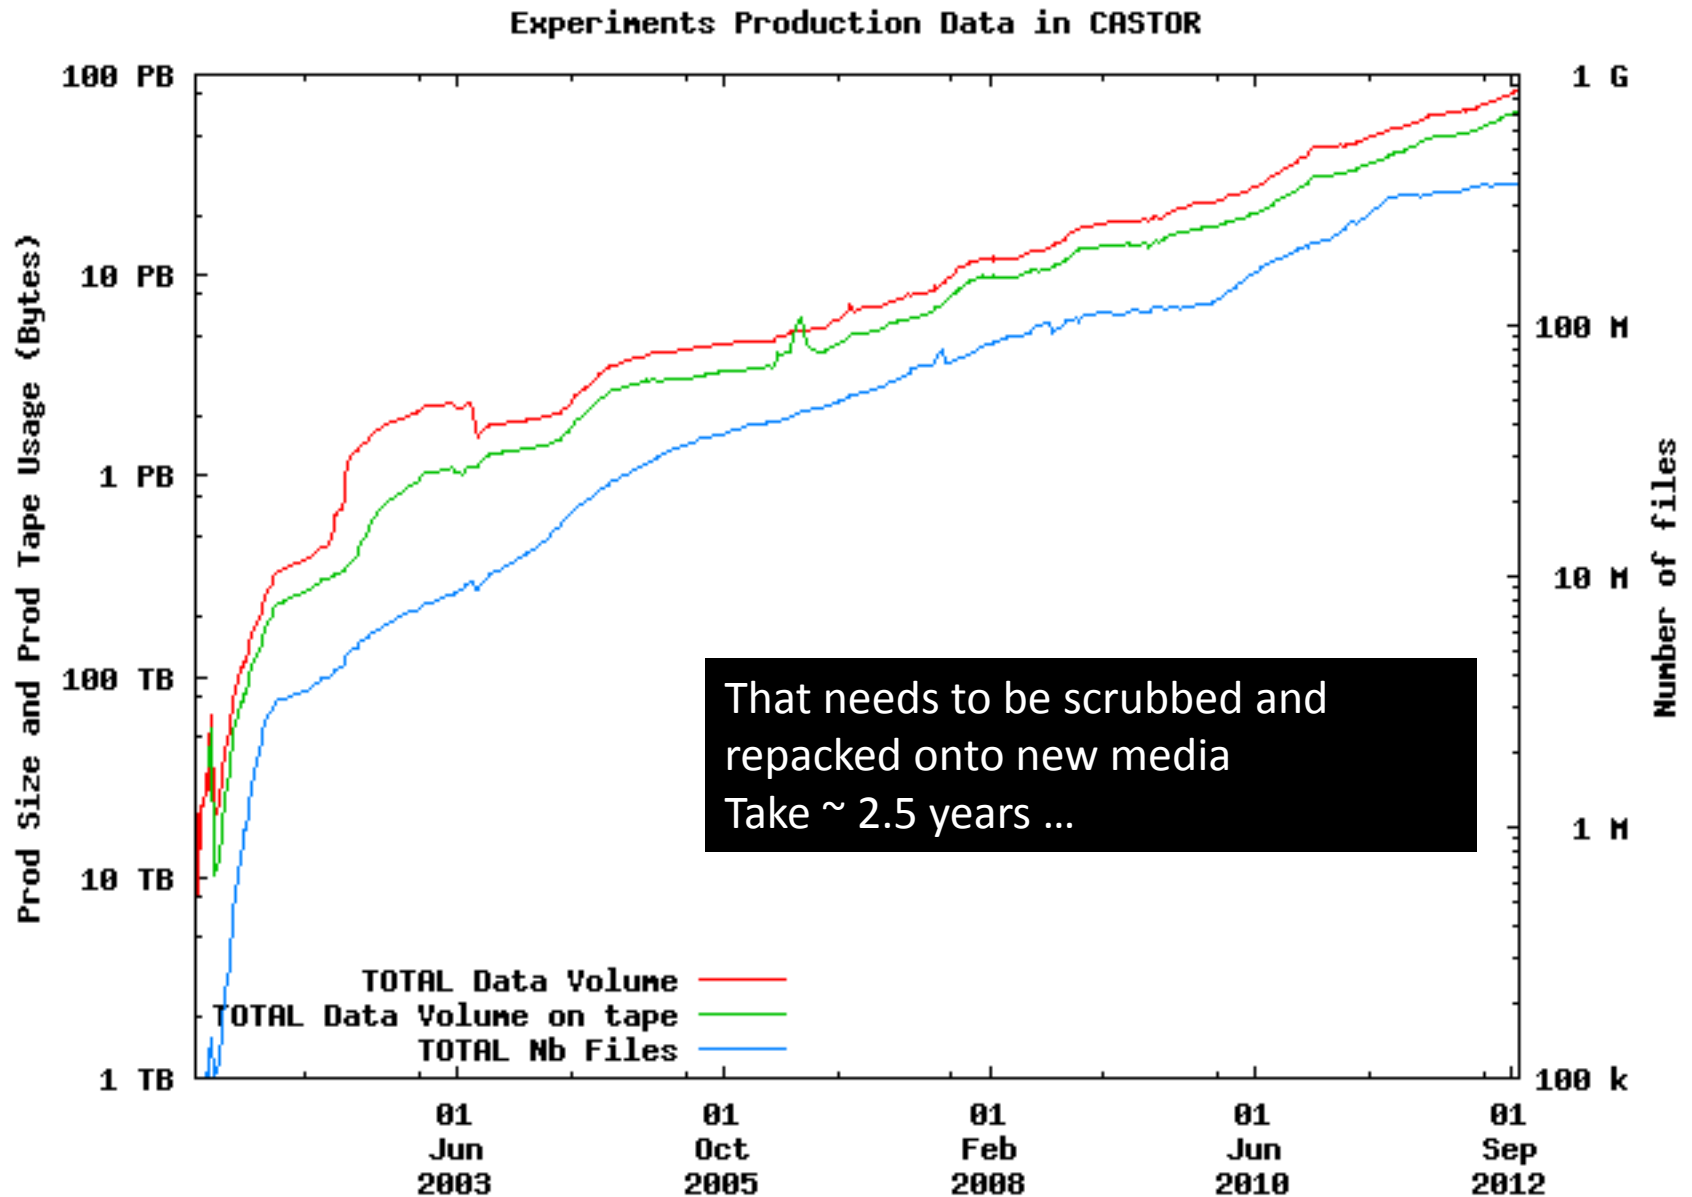
# Data written 2008-12



Data rates increased  
 3-4 GB/s input  
 ~15 GB/s output



# CERN alone has ~100 PB archive


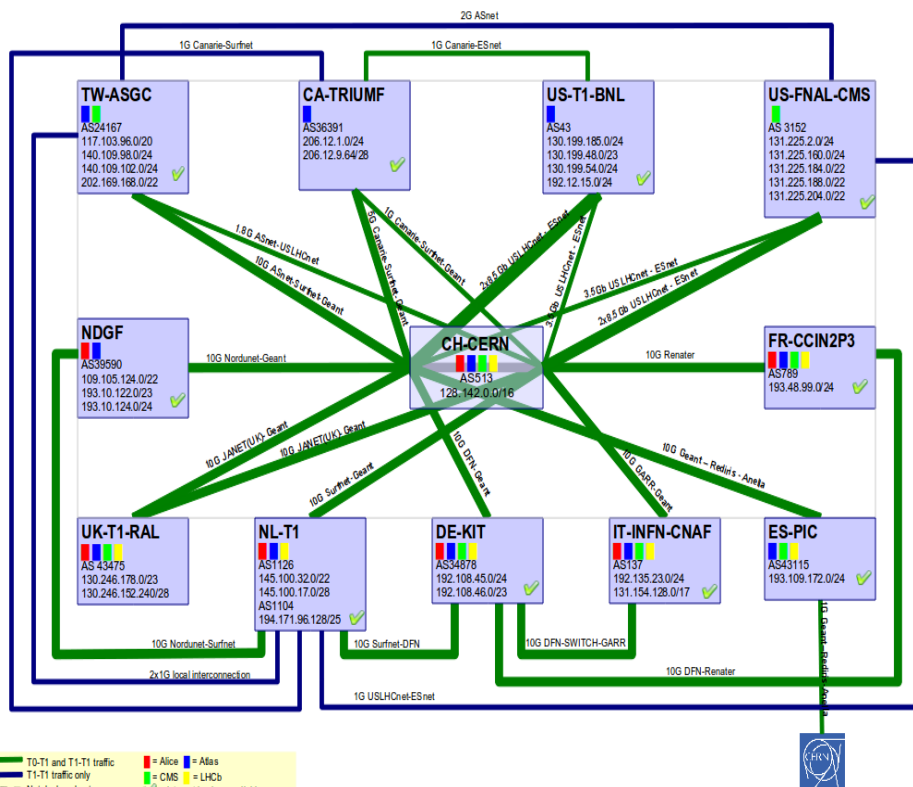




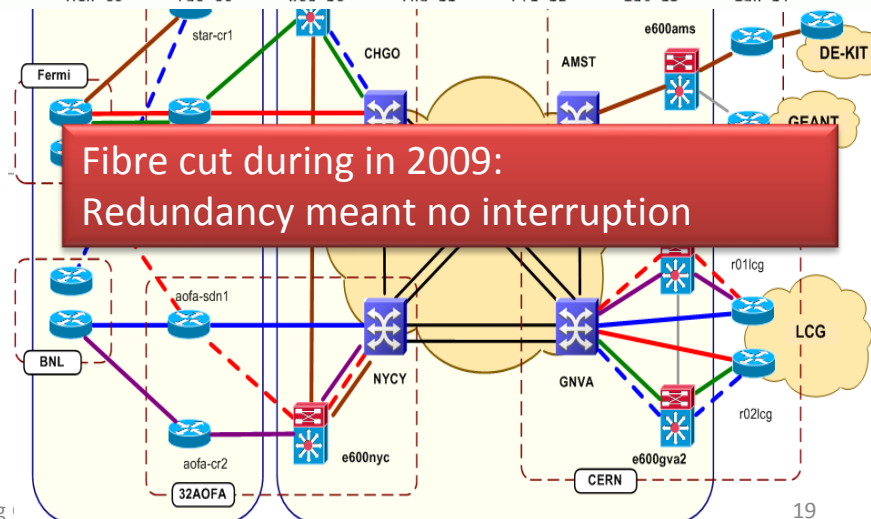

# The Network(s)

# LHC Networking

## LHCOPN



LHCOPN Total Traffic




Relies on

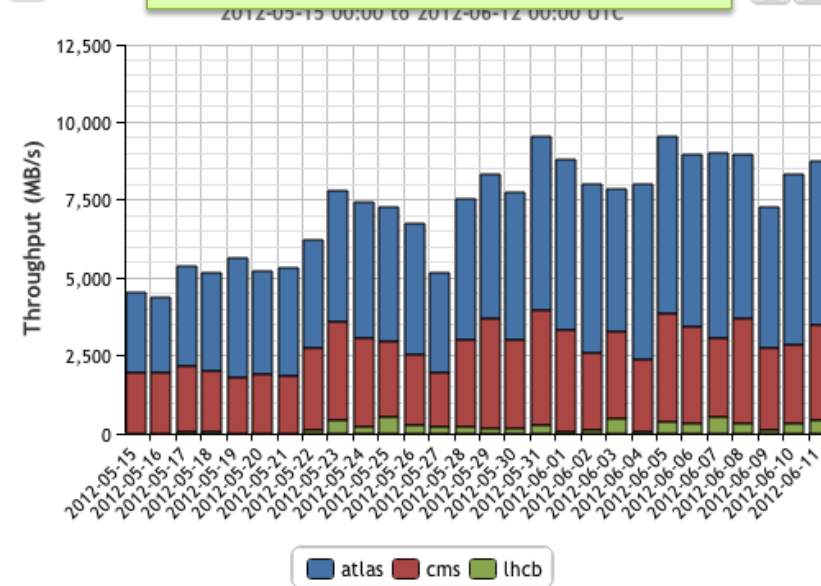
- OPN, GEANT, US-LHCNet
- NRENs & other national & international providers

# Data transfers


Global transfers > 10 GB/s (1 day)



Global transfers (May-June)



CERN → Tier 1s (with LHC on): 2 GB/s







Planned rates:

Tier 0 – Tier 1: 650 MB/s + margin




## Data Transfers

A key technology for  
LHC



## Total data movement inside CERN



A typical month:

4 PB, 2.6M files exported from CERN to Tier 1s  
20 PB, 33M files moved worldwide

Data movement inside Tier 0:

- Data in: 4.3 GB/s (peak 13 GB/s)
- Data out: 11.5 GB/s (peak 25 GB/s)
- ~1.5 PB/day data movement at CERN

# Wigner Status (April 2013)

- Contract Adjudicated in March 2012 to the Wigner Research Center for Physics in Budapest
- Two 100 Gbps circuits adjudicated to two different providers
  - At reasonable & comparable costs
  - Operational in 2012
- Works progressing well
  - Completion scheduled for June 2013
  - Workers were peaking at 100-120 in October 2012
  - Equipment delivered and running
- Some teething problems
  - Wigner had assumed they could preinstall the networking
  - Customs/Taxes issues still not sorted out
  - Networking physical paths



June 2012




October 2012



October 2012



April 2013

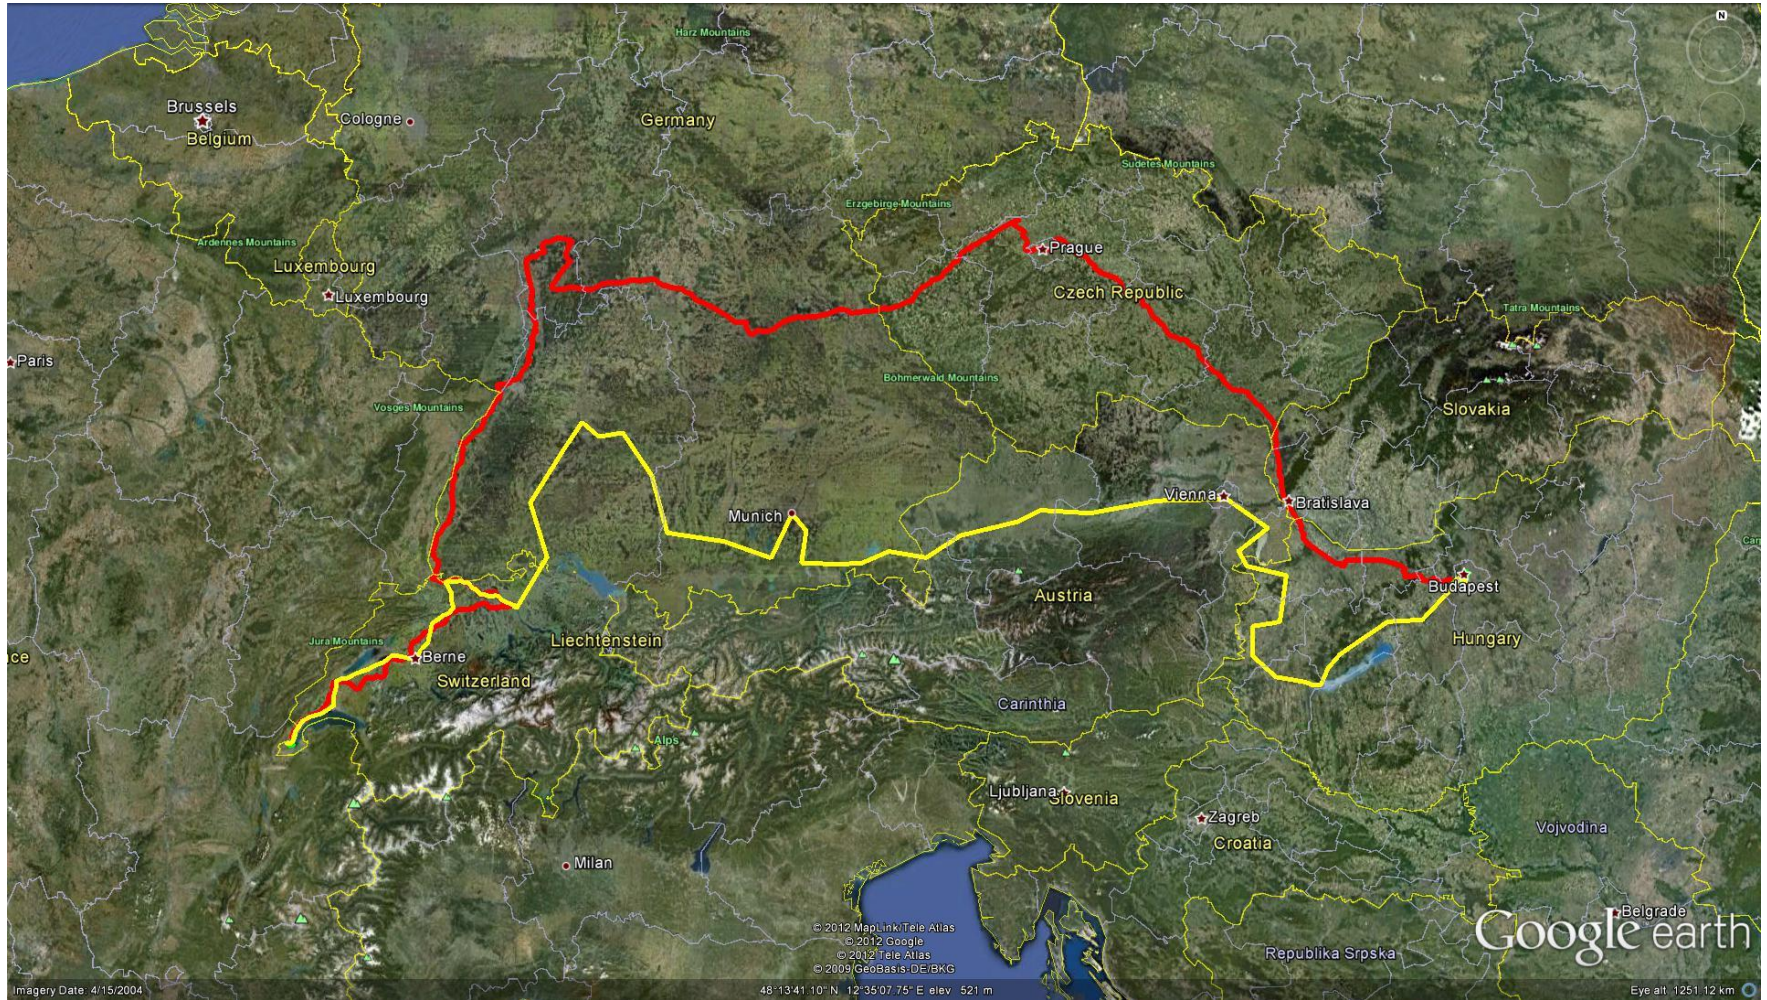


January 2013





# Connectivity (100 Gbps)







# Processing

# Processing on the grid

Usage continues to grow...


- # jobs/day
- CPU usage

~ 150,000 years of CPU delivered each year



**This is close to full capacity  
We always need more!**

CPU delivered January 2011




CERN  
BNL  
CNAF  
KIT  
NL  
LHC/Tier-1  
RAL  
FNAL  
CC-IN2P3  
ASGC  
PIC  
NDGF  
TRIUMF  
Tier2

# CPU – around the Tiers

- The grid really works
- All sites, large and small can contribute
  - And their contributions are needed!


Tier2 CPU delivered by country January 2011



USA	UK
France	Germany
Italy	Russian Federation
Spain	Canada
Poland	Switzerland
Slovenia	Czech Republic
China	Portugal
Japan	Sweden
Israel	Romania
Belgium	Austria
Hungary	Taipei
Australia	Republic of Korea
Norway	Turkey
Ukraine	Finland
India	Pakistan
Estonia	Brazil
Greece	

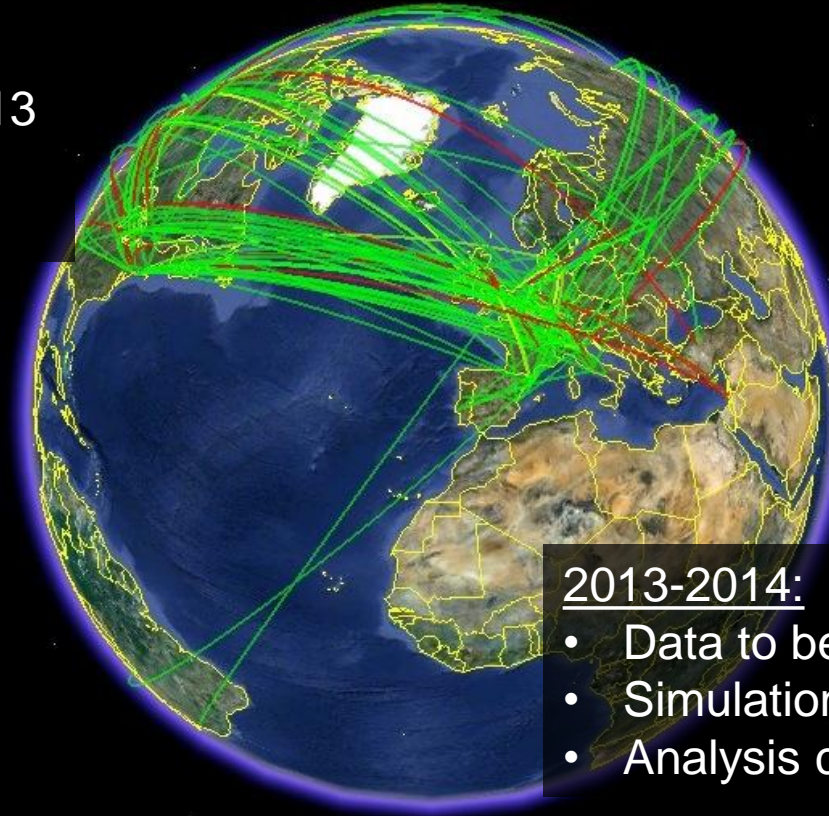


# No stop for the computing !



Running jobs: 246791  
Transfer rate: 13.98 GiB/sec

Activity on 1 January 2013  
Running Jobs: 246791  
Transfer rate: ~14 GiB/s



## 2013-2014:


- Data to be reprocessed
- Simulation for Higher Energy
- Analysis continuing

Data SIO, NOAA, U.S. Navy, NGA, GEBCO  
© 2012 Google  
US Dept of State Geographer  
© 2009 GeoBasis-DE/BKG



Google

Eye alt 22

# Electrical Consumption



**Assuming CERN is 20 %  
→ 25 MW Worldwide  
24x7...**



# Collaborating with Other Sciences

# Broader Impact of the LHC Computing Grid

- WLCG has been leveraged on both sides of the Atlantic, to benefit the wider scientific community

- Europe:

- Enabling Grids for E-science (EGEE) 2004-2010
- European Grid Infrastructure (EGI) 2010--


- USA:

- Open Science Grid (OSG) 2006-2012 (+ extension?)

- Many scientific applications →

**Archeology**  
**Astronomy**  
**Astrophysics**  
**Civil Protection**  
**Comp. Chemistry**  
**Earth Sciences**  
**Finance**  
**Fusion**  
**Geophysics**  
**High Energy**  
**Physics**  
**Life Sciences**  
**Multimedia**  
**Material Sciences**

...







Information Society Technologies



### Measurement of Pulmonary Trunk

Dynamic measurements of the Pulmonary Trunk

- Pulmonary trunk morphology classification based on the diameters measured at 50% and 100% of the pulmonary trunk length

### RV and LV Automatic Modelling

Automatic

### Surgery Planning

Virtual Volume Reduction Surgery

### Personalised Simulation

Personalised Simulation: Results

Simulated parameters

Simulated results

### Semantic Browsing

Biomechanical Model

### Similarity Search

Similarity Search

### Temporal Modelling

Temporal Modeling

- Each visit's internal dependencies are represented as edges between nodes of the same layer (intra slice topology)
- Temporal dependencies between visits are represented as edges between nodes belonging to different layers (inter slice topology)

### Visual Data Mining

Visual Data Mining

Visual data

- Coloring
- Size
- Labels
- Hierarchy

Ontological hierarchy is reflected automatically

### Genetics Profiling

PA - Cerebrum vs. Cerebellum

Siemens

### Treatment Response

Example: treatment response

Follow-up (1 year)

### Inferring Outcome

Temporal reasoning: Inferring outcome

### Biomechanical Models

Biomechanical Model

# Sequence Production & IT Infrastructure at EMBL


4 x Illumina HiSeq2000



2 x Illumina GAllx



25 TB data  
each week




Compute Power:  
2000+ CPU Cores, 6+ TB RAM



Storage:  
1+ PB High Performance Disk

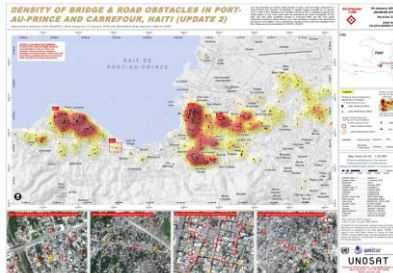
# Problem - Technology Explosion with NGS

**Bases Sequenced / Sample / Run @ EMBL  
(Illumina)**



# Did you know...?

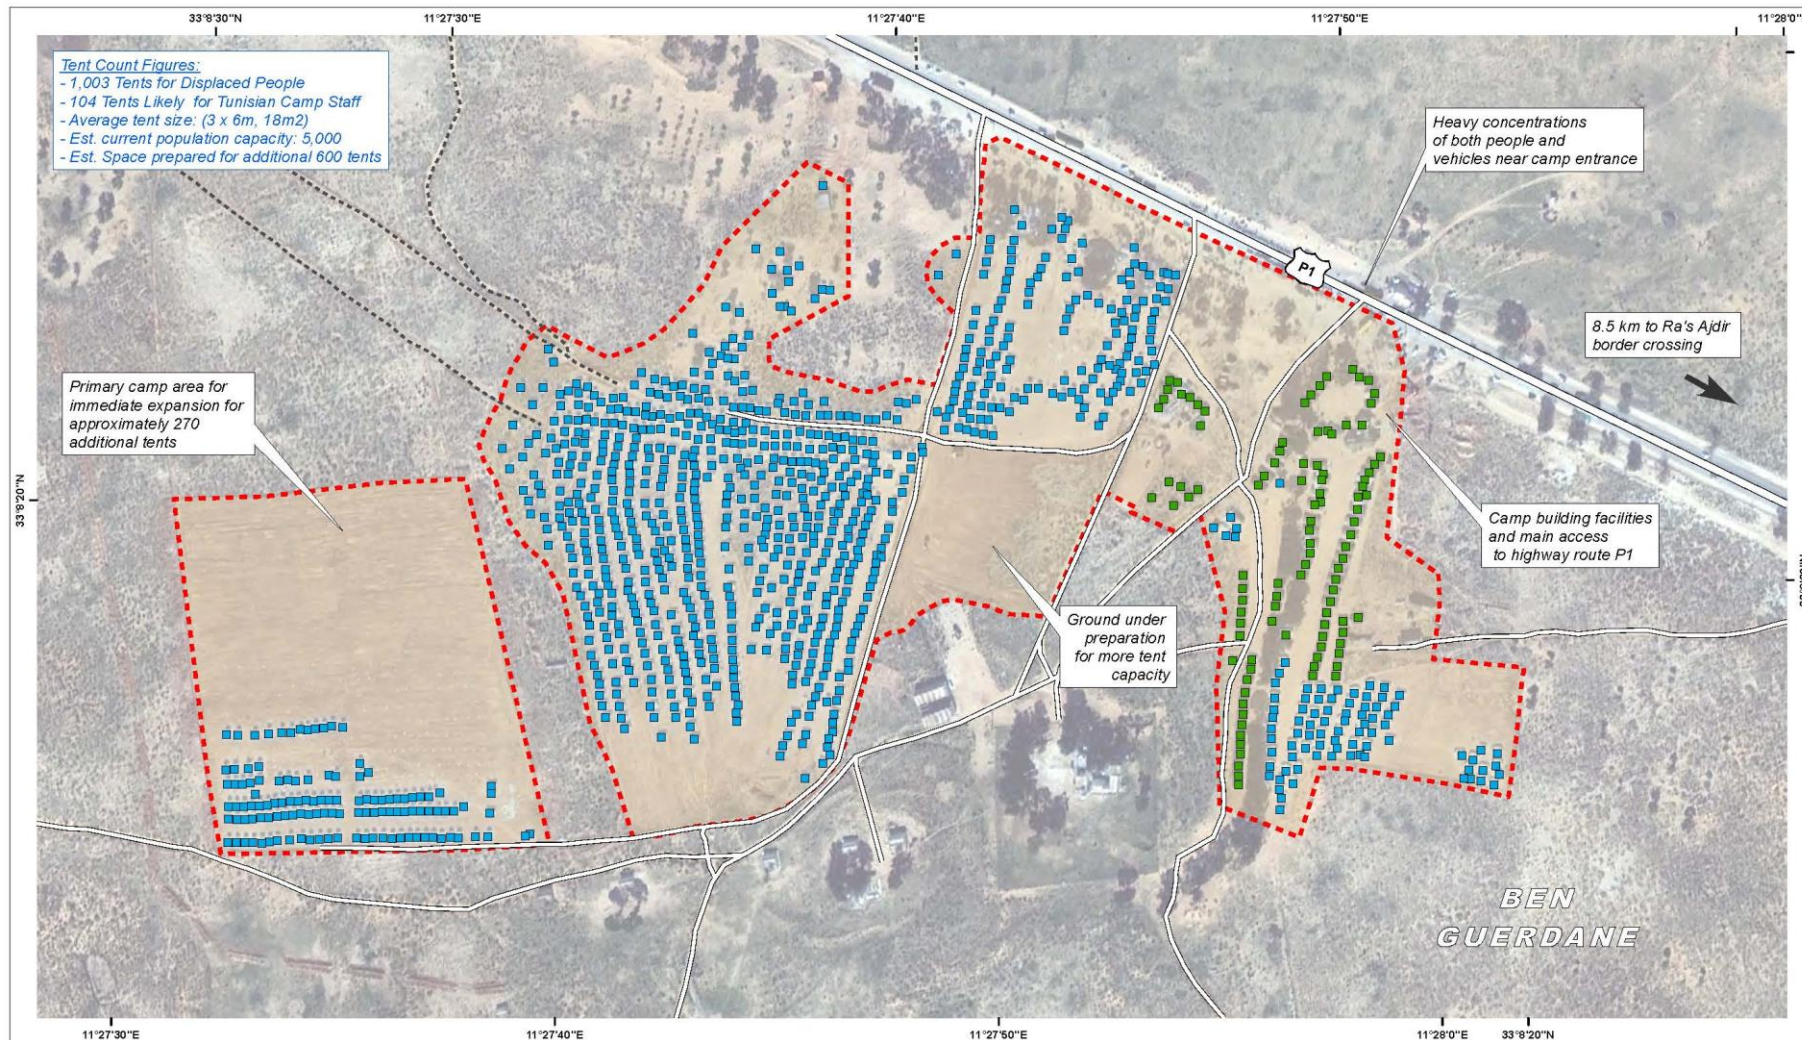
WHO is connected via  
CERN to the Commodity  
Internet



Cancer	Cardiovascular diseases	Diabetes
<b>7.6 million</b> deaths from cancer worldwide in 2008.	<b>29%</b> of global deaths due to cardiovascular diseases.	<b>220 million</b> people have diabetes worldwide.
<a href="#">Read the cancer fact sheet</a>	<a href="#">Read the cardiovascular diseases fact sheet</a>	<a href="#">Read the diabetes fact sheet</a>

UNOSAT staff and  
servers hosted at  
CERN





# Helping developing countries

Invenio is an open source software suite enabling you to run your own digital library or document repository on the web. The technology offered by the software covers all aspects of digital library management from document ingestion through classification, indexing, and curation to dissemination.

Invenio is the tool used for LHC publications workflow approval at CERN.



UNESCO & CERN have sponsored digital library schools in Africa (Rwanda 2009, Morocco 2010, Senegal 2011)

CERN has donated (WLCG Tier-0) computers to Africa (Morocco, Ghana, Egypt soon).





# IT Department Services

- **E-Mail and Distribution Lists**
  - Up to 250000 incoming messages/day, 80% detected as spam
  - 27'000 mailboxes
- **AFS (worldwide) Distributed File System**
  - 1.5 B files, 125 TB, 4 B access/day, 15000 clients
- **Backup Services (files, databases)**
  - 7.4 PB, 2 B files
- **Network, Telephony**
  - More than 400 star points and ~80 000 UTP sockets
  - ~450 wireless base stations
  - GSM (Site, Tunnels & Experimental halls) – 4300 subscriptions
  - Multi-10Gb/s External Networking Capacities
  - TETRA replaced VHF communications for Fire Brigade
- **Web Services**
  - 12000 Web sites
  - IIS hosted, Drupal, Sharepoint
- **Active Directory & SSO**
  - Central authentication service for Linux and Windows computers and (Web) applications
  - Online X509 Certificate Authority (~8'000 User certificates, 40'000 host certificates)
  - Multifactor authentication introduced
  - Federated authentication coming
- **Windows Services**
  - 200 TB of DFS workspaces
  - ~ 7'000 active PCs managed by CMF
- **Virtualisation Infrastructure**
  - 3500 VM's (80% Linux, 20% Windows)
  - 355 Hypervisors

# 2012 Statistics and Numbers - OS & Devices

<b>Active Computers</b>	<b>2012</b>	<b>2011</b>	<b>Δ %</b>
Linux	20759	18452	12.50
Windows	10407	9637	7.99
Mac	2351	2173	8.19

<b>Active Mobile devices</b>	<b>2012</b>	<b>2011</b>	<b>Δ %</b>
Android	1129	366	208.47
iPhone	976	621	57.17
iPad	89	---	---
iPod	39	64	-39.06
Symbian	43	59	-27.12
Win Mob	19	20	-5.00


<b>Windows Desktops</b>	<b>2012</b>	<b>2011</b>	<b>Δ %</b>
Total	7051	6607	6.72
Win 7	4676	3014	55.14
Win XP	2375	3593	-33.90
Win7 64 bit	2188	915	139.13
Win7 64/32 [%]	46.79	30.36	54.13



# 2012 Statistics and Numbers - Web Services


## Total 11803 sites (+7%) on ~55 Web servers

	Nov 2012	Nov 2011	$\Delta$ %
centrally hosted	5298	4760	11.3
AfsGateway	3389	3008	12.7
Sharepoint	2115	1814	16.6
Drupal	342	163	109.8
Twiki	294	260	13.1
J2EEApp	202	161	25.5
NtGateway	126	778	-83.8
Frontpage	37	81	-54.3
Total	11803	11025	7.1



## Web traffic:

- October 2012: **1.04 billion hits**
- May 2012: 7.5 TByte / day



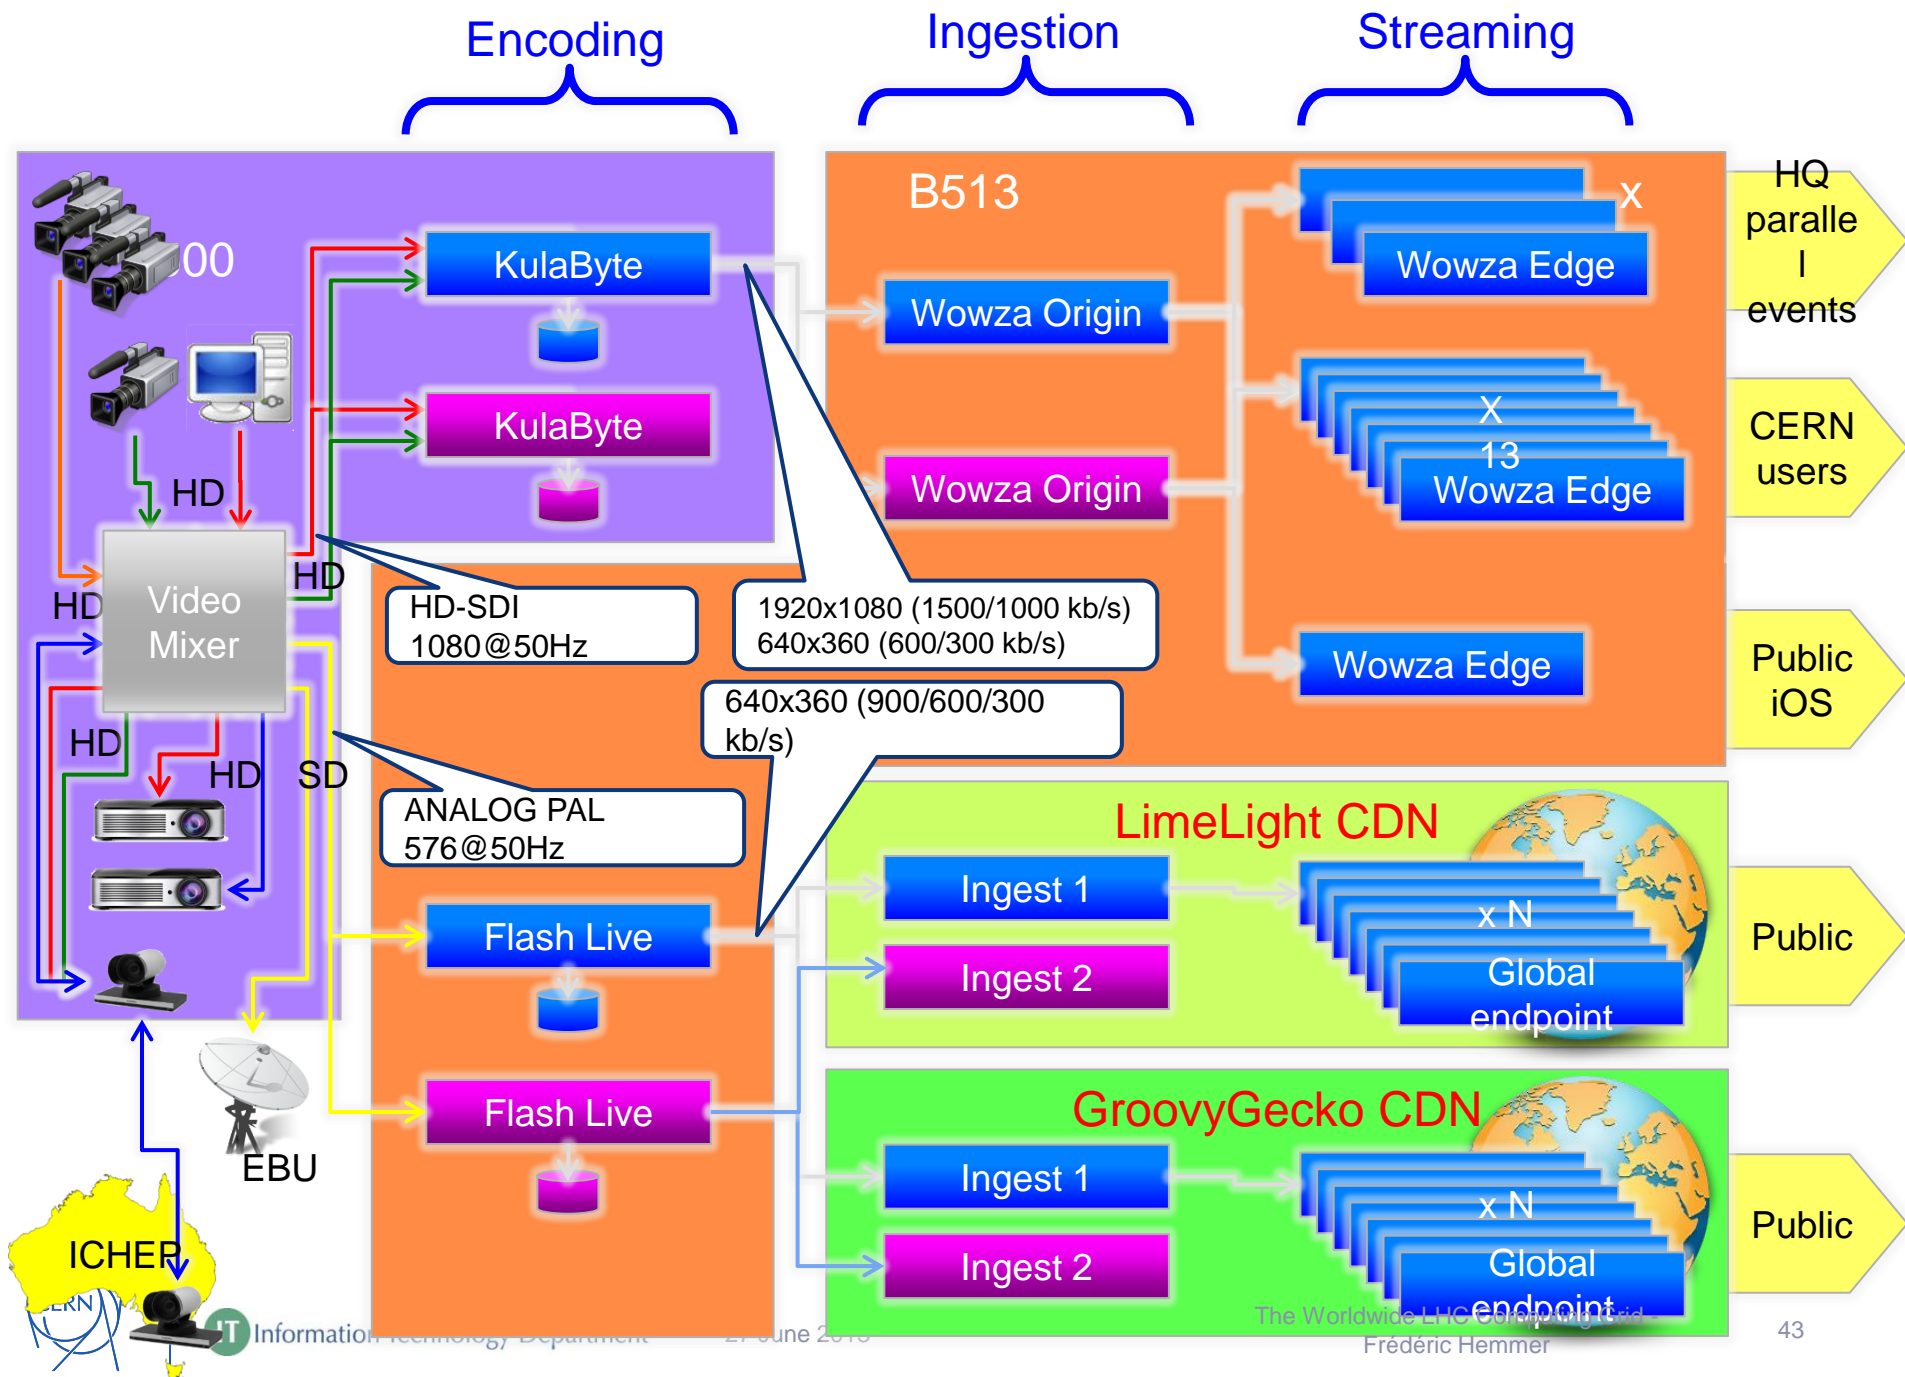
# IT Department Services

- **Database and Application Deployment Services**
  - Mainly based on Oracle software
  - AIS DBs and Applications, EDMS, Accelerator DBs, IT DBs, CASTOR DBs, Physics databases (Calibration, Alignment, etc...), Public J2EE Service, etc...
  - 120 General Purpose Databases, 500 TB of NAS storage
  - 130 Web /Application Servers with 700 virtual hosts
  - 50 Terabytes of worldwide replicated Physics databases
  - Database on Demand Service (Oracle, MySQL)
- **Engineering and Software Development Services**
  - Mechanical and electronic CAE, field calculations, structural analysis, simulations, mathematics, etc
  - 50 packages, 1000 users
  - Twiki Service
  - 12'000 users, 60'000 pages updated per month
  - Version Control Services (CVS/SVN)
  - 2'500 users, 400 projects
- **Audiovisual Service**
  - support, record and archive official committees and events
- **Video Conference Service**
  - video conferencing in rooms across site
- **Conference Management System (Indico)**
  - Distributed and used worldwide
- **CDS-Invenio, a Digital Library Open Source Software produced, used and maintained at CERN**
  - free support via mailing lists
  - commercial-like support via a maintenance contract
- **Computer Security**

# Higgs Update Seminar 4.7.2012

- Very short notice given
- Worldwide visibility
- A/V Team (and others) worked days & nights!
- Estimated 1 billion people reached
  - Including Antarctica ...
  - <http://avc-dashboard.web.cern.ch/node/3>
- Main auditorium upgraded with HQ just, just in time ...







# The CERN Data Centre in Numbers

- Data Centre Operations (Tier 0)
  - 24x7 operator support and System Administration services to support 24x7 operation of all IT services.
  - Hardware installation & retirement
    - ~7,000 hardware movements/year; ~1800 disk failures/year
  - Management and Automation framework for large scale Linux clusters



Racks	828
Servers	10,070
Processors	17,259
Cores	88,414
HEPSpec06	482,507

Disks	79,505
Raw disk capacity (TiB)	124,660
Memory modules	63,326
Memory capacity (TiB)	205
RAID controllers	3,749

Tape Drives	160
Tape Cartridges	45000
Tape slots	56000
Data on Tape (PiB)	65

High Speed Routers (640 Mbps → 10 Tbps)	25
Ethernet Switches	500
10 Gbps/100Gbps ports	1200/50
Switching Capacity	6 Tbps
1 Gbps ports	16,041
10 Gbps ports	1,884


IT Power Consumption	2392 KW
Total Power Consumption	3929 KW



# Summary - scaling CERN Data Centre(s) to anticipated Physics needs

CERN Data Centre dates back to the 70's

- Upgraded in 2005 to support LHC (2.9 MW)
- Still optimizing the current facility (cooling automation, temperatures, infrastructure)



Renovation of the “barn” for accommodating 450 KW of “critical” IT loads (increasing DC total to 3.5 MW)

Exploitation of 100 KW of remote facility down town

- Understanding costs, remote dynamic management, improve business continuity



Exploitation of a remote Data centre in Hungary

- Max. 2.7 MW (N+1 redundancy)
  - Improve business continuity
- 100 Gbps connections





Thank You !



[www.cern.ch](http://www.cern.ch)