Transfer rate: 1.62



The CERN Computer Centre and the LHC Computing Grid

1.

Università degli Studi di Trieste Corso di Laurea in Fisica

Saturday 25th of February 2012

Massimo Lamanna IT Data Storage Services

© 2011 Europa Technologies US Dept of State Geographer © 2011 MapLink/Tele Atlas © 2011 Google





I keep six honest serving-men (They taught me all I knew); Their names are **What** and **Why** and **When** And **How** and **Where** and **Who** | R. Kipling







CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

RC

CERN

Department



1960: IBM 709 at Geneva airport





CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it







1983: UA1 discovers Ws and Z0 at the CERN SppS

.....

1983: Computer centre mainframes (IBM, Siemens, CDC) Bd.513

01



www.cern.ch/it

LHC experiments (ATLAS) CERNIT

ATLAS: 7,000 tons, 150 million sensors generating data 40 millions times per second i.e. a petabyte/s (1 million GB/s)

Another view of ATLAS

CERN**IT** Department



The LHC Computing Grid - April 2011

The role of the CERN Computer CERNIT









The LHC Computing Grid - April 2011

Department

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

RC

CERN Computer Centre: Storage, Distribution CERN and Processing (Reconstruction and Analysis) D



Department

Computer Centre By Numbers



full name: Computer Centre By Numbers short name: CCBYNUM

group: IT-CF-FPP

site: CERN

email: imre.szebenyi@cern.ch

manager: Imre Szebenyi 😔



Additional service information (more)				
Number of processors:	14,972			
Number of cores:	64,623			
Memory capacity (TiB):	165			
Memory modules:	55,729			
Raw HDD capacity (TiB):	62,660			
Number of HDD's:	62,023			
Number of systems:	7,975			
Number of RAID controllers:	3,607			
Number of enclosures:	1,554			
SPEC CPU2006:	503,637			
Number of racks:	1,070			
Number of virtual machines:	3,908			
Number of Fibre channel ports:	742			
Number of 1G ports:	16,773			
Number of 10G ports:	622			
Current power consumption (kW):	2,186			
Current power consumption (kVA):	2,305			

24x7 operator and system admin support

٠

Management and Automation framework for large scale Linux clusters

- Hardware installation & retirement ~7,000 hardware
- movements/year; ~1000 disk failures/year

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

The LHC Computing Grid - February 2012



What are all these computers for?









CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it





Electrons $p_T = 34.0, 31.9 \text{ GeV/c}$ Inv. mass = 91.2 GeV/c²



The LHC Computing Grid – February 2012

From Physics to Raw Data

CERN

Department



CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

The LHC Computing Grid – February 2012

From Raw Data to Physics

CERN

Department



The LHC Computing Grid – February 2012

Analysis flow (user view)

CERN

Department



But how this is done *in practice*? Of course we need CPUs, disks, networks etc.. We cannot rush to the solution yet...

Dataset concept = collection of files. Only a small fraction of data in real DBs (e.g calibrations)

The HEP Data Challenge



LHC will run for 20 years
Experiments *are* producing about 15 Million Gigabytes of data each year (about 20 million CDs!)
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

A challenge for physics... ... and a challenge for technology research and industry as well





CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

"Historical" example...

The Economist

The data deluge

CERN**IT** Department

1990s:

The web was invented at CERN! The machine used by Tim Berners-Lee in 1990 to develop and run the first WWW server, multi-media browser and web editor.

R

COMPASS proposal (1996)

- Use a Linux PC farm
 - Instead of the "usual super computers"
- C++
 - Instead of good ol' Fortran IV
- All data in a data base (Objectivity/DB)
 - Instead of "plain" files
- First experiment using CASTOR
 - Instead of writing tapes "by ourself"

- ... a quark, confined in the proton.
- Data challenge

A high-energy

collision course with ...

electron on

- Semi inclusive → reconstruct "the rest" of the deep inelastic event
- Cool, but:
 - 2 times less data 15 years (~2⁷) before LHCb (LHC experiment)
 - 64x factor against us

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it





www.shutterstock.com · 3748

ENFORCED

RC

CERN**IT** Department



Heavy Ion run (2012) Experiment CMS Red = Expt → Computer Centre Green = Computer Centre to clients (internal and external)

RC

CERN IT Department CH-1211 Genève 23

Switzerland www.cern.ch/it



www.cern.ch/it

A new 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

Department

CH-1211 Genève 23

Switzerland www.cern.ch/it



The LHC Computing Grid – February 2012

Department

WLCG Tiers Organization



Tier-0 (CERN):

Data recording Initial data

CERN

Department

- reconstruction
- Data distribution

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

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

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

How does it work?

ERN**IT** Department

ATLAS

Not substantially different for the other HEP experiments Heavily simplified...

What do we want to achieve

The user wants to specify a subset of the data and run applications on it (chain of programs reading intermediate outputs)

Only at the end of the chain data sizes and computational complexity this can be (possibly) done on a laptop

1000+ of physicists worldwide after the same data

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

25-NOV-09



Distributed analysis





CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

THE REPORT OF A



Behind the scenes...

CERN



Impact of the LHC Computing Grid in Europe Department

Enabling Grids for E-sciencE

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

21:13:50 UTC

LCG has been the driving force for the European multi-science Grid EGEE (Enabling Grids for EsciencE) EGEE was a global effort, and the largest Grid infrastructure worldwide

Co-funded by the European Commission (Cost: ~170 M€ over 6 years, funded by EU ~100M€) Now moving to a permanent European Infrastructure: EGI >300 sites

>300 sites
48 countries
>200,000 cores
>20 PetaBytes
>10,000 users
>150 VOs
>150,000 jobs/day

Scheduled = 21539

Running = 25374



The LHC Computing Grid - April 2011



More info: INFN (Istituto Nazionale Fisica Nucleare): http://www.infn.it IGI (Italian Grid Initiative): http:// www.italiangrid.org/



CH-1211 Genève 23 Switzerland www.cern.ch/it The LHC Computing Grid – February 2012

ITU conference (2006)

40° S

Figure 1

The problem: Assign frequencies for digital radio and television (international treaty)

Critical point: Need on dependability: verify (iteratively) the compatibility between radio stations

Solution: Use the EGEE grid + asystem used in ATLAS and LHCb to increase the reliability of the Grid



CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

Questions?



CERN

Department

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

RC

The LHC Computing Grid – February 2012^{35}

	CERN			-	
	https://ert.cern.ch/browse_www	v/wd_pds?p_web_site_id=1	☆11▼)	• (G * twiki computer cen	trQ)
I GCal Docs News ▼	Bank - RW - Dict - Mix - Old - Ub	untu and Free S 🔻 Latest Headlines	৯ Apple ▼ Google Ma	aps YouTube Popular -	Pytho
e-Rec	ruitment	● Vacancies ○ All CERN C	HR Department	Site map Contact us CERN Ho	ome th
HR General Information	Recruitment Training	Staff Career	Services	CERN Official Docume	nts
Welcome Page					
Intranet	CERN is the European Organization fo Franco-Swiss border near Geneva (mo	r Nuclear Research, based on the ore).	News		
Posictor in o.PT			News		
Login to o-PT	We are at the forefront of technologie opportunities for both working and lea	Welcome to our e-Recruitment website! We rely on			
Search Vacancies	graduate programmes, as well as vaca electricity mechanics electronics and	your feedback to continue improving this site. Problems can be reported by mail.			
Sedicit vacancies		comparing, etci			
Full Search	Please use the left-hand menu to sear opportunities or look for further inform	CERN staff please click on "Internal posts" in the menu to access internal vacancy notices (AIS login			
	and programmes, recruitment events	or to contact us.	required)! (more info)		
Recently Published		Important technical information.			
By Reference		Turiguis			
Employment Conditions			10		
Information for			Deadlines		
Staff					
Fellows			Technical &	07-AUG-09	
Associates			Doctoral Students	07.050.00	
Students		Com Jak	reliows	07-SEP-09	
Marie Curie Actions			Scientific Associates	19-MAR-09	
Special Programs		Do not wait until the last day to send your application as additional information will be			
Apprentices			requested by CERN one	ce your application form has	
Your Feedback		12 Star Avenue	been received! More In	ifo	
Contact Us					_
FAQ		CARL HAT	Focus on		
			Are you an undergradu: State nationality in a te practical training period project? CERN has a Te that could interest you. least 18 months of your studies, and your cours	ate student of a CERN Memb schnical field looking for a l or a place to do your final chnical Student programme If you have completed at r technical undergraduate e requires a practical trainin	er



CERN options for students

University level (BS/Master)

- Summer student
- OpenLab summer students

Master thesis

Technical student (non physicist)

PhD students

Doctoral students

Young scientists/engineers

Fellowship

¥

*

ert.cern.ch

period of 6 to 12 months, which you wish to spend at CERN, apply to the Technical Student Programme.

late that students of theoretical and experimenta

Other programmes

CERN IT Department CH-1211 Genève 23 Switzerland www.cern.ch/it

Done

36