



# CERN, GRID and E-Science

## ***Contents:***

- ***Introduction***
- ***Computer intensive science***
- ***Particle physics and the LHC***
- ***The LHC data challenge***
- ***LCG – the LHC Computing Grid***
- ***The CERN IT Openlab***

Nils Høimyr, CERN IT

***Includes presentation contents from Frédéric Hemmer, Bob Jones and the CERN IT Openlab***


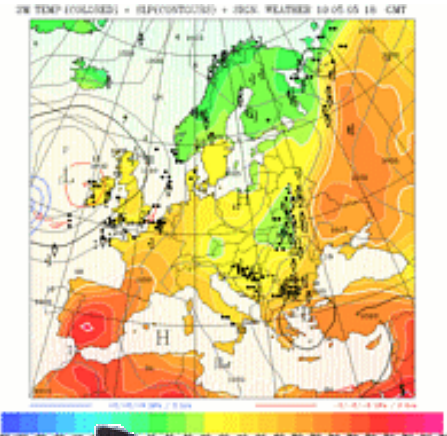



# IT at CERN – more than the Grid

- **Physics computing – Grids (this talk!)**
- **Administrative information systems**
  - Financial and administrative management systems, e-business...
- **Desktop and office computing**
  - Windows, Linux and Web infrastructure for day to day use
- **Engineering applications and databases**
  - CAD/CAM/CAE (Autocad, Catia, Cadence, Ansys etc)
  - A number of technical information systems based on Oracle, MySQL
- **Controls systems**
  - Process control of accelerators, experiments and infrastructure
- **Networks and telecom**
  - European IP hub, security, voice over IP...

**More information:** <http://cern.ch/it>

- Science is becoming increasingly **digital** and needs to deal with increasing amounts of data
- **Simulations** get ever more detailed
  - Nanotechnology – design of new materials from the molecular scale
  - Modelling and predicting complex systems (weather forecasting, river floods, earthquake)
  - Decoding the human genome
- **Experimental Science** uses ever bigger **sensors** to make precise measurements
  - Compute a lot of statistics
  - Huge amounts of data
  - Serves user community around the world






# Particle Physics (I)




- **CERN: the world's largest particle physics laboratory**
- **Particle physics requires special tools to create and study new particles: accelerators and detectors**
- **Large Hadron Collider (LHC):**
  - most powerful instrument ever built to investigate elementary particles
  - four experiments: ALICE, ATLAS, CMS, LHCb
  - 27 km circumference tunnel
  - First beam 10<sup>th</sup> September 2008



- **Physicists smash particles into each other to:**


- identify their components
- create new particles
- reveal the nature of the interactions between them
- create an environment similar to the one present at the origin of our Universe




- **A particle collision = an event**

- need to count, trace and characterize all the particles produced and fully reconstruct the process


- **Among all tracks, the presence of “special shapes” is the sign for the occurrence of interesting interactions.**



Starting from  
this event




Looking for  
this “signature”




→ **Selectivity: 1 in  $10^{13}$**

(Like looking for a needle in 20 million haystacks)

## View of the ATLAS detector (under construction)



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





- **Simulation**

  - compute what the detector should have seen


- **Reconstruction**


  - transform signals from the detector to physical properties (energies, charge of particles, ...)

- **Analysis**

  - use complex algorithms to extract physics

→ **LHC data analysis requires a computing power equivalent to ~ 100,000 of today's fastest PC processors!**





# CERN Computing – Tier 0 in numbers

- Computing – CPU:
  - 6300 systems / 39k cores (+ planned 2400 systems, 16k cores)
  - Used for CPU servers, disk servers, general services
- Computing – disk:
  - 14 PB on 42.5k disk drives (+ planned 19 PB on 20k drives)
- Computing – tape:
  - 34 PB on 45k tape cartridges
  - 56k tape slots in robots, 160 tape drives
- Computer centre:
  - 2.9 MW usable power, + ~1.5 MW for cooling

# 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





# LHC Computing Grid project (LCG)


- More than 140 computing centres
- 12 large centres for primary data management: CERN (Tier-0) and eleven Tier-1s
- 38 federations of smaller





# WLCG Collaboration

- **The Collaboration**
  - 4 LHC experiments
  - ~140 computing centres
  - 12 large centres (Tier-0, Tier-1)
  - 38 federations of smaller “Tier-2” centres
  - ~35 countries
- **Memorandum of Understanding**
  - Agreed in October 2005
- **Resources**
  - Focuses on the needs of the four LHC experiments
  - Commits resources
    - each October for the coming year
    - 5-year forward look
  - Agrees on standards and procedures
- **Relies on EGEE and OSG (and other regional efforts)**




Open Science Grid

EGEE

Enabling Grids for  
E-science in Europe






# Preparation for accelerator start up

- Since 2004 WLCG has been running a series of challenges to demonstrate aspects of the system; with increasing targets for:
  - Data throughput
  - Workloads
  - Service availability and reliability
- Recent significant challenges
  - May 2008 – Combined Readiness Challenge
    - All 4 experiments running realistic work (simulating data taking)
    - Demonstrated that we were ready for real data
  - June 2009 – Scale Testing
    - Stress and scale testing of all workloads including massive analysis loads
  - Real data – cosmic rays many months of 2008/9; collision data Nov/Dec 2009
    - Successful data acquisition and distribution /analysis at full data rates
- In essence the LHC Grid service has been running for several years

Tier-2s and Tier-1s are inter-connected by the general purpose research networks

### LHCOPN Total Traffic



	Min	Avg	Max (avg)	Max (max)
ASGC	0.00	1.77 G	6.65 G	7.97 G
CNAF	0.00	1.49 G	5.03 G	6.01 G
KIT	0.00	2.72 G	7.75 G	10.27 G
IN2P3	204.57 M	3.27 G	9.46 G	9.87 G
NDGF	0.00	435.28 M	1.75 G	2.07 G
PIC	5.39 M	3.34 G	9.53 G	12.12 G
RAL	153.39 M	3.62 G	9.35 G	10.56 G
NLT1	26.03 k	1.93 G	7.20 G	8.29 G
TRIUMF	0.00	713.47 M	2.72 G	4.28 G
BNL	50.32 M	3.12 G	8.82 G	11.21 G
FNAL	166.72 M	1.67 G	5.80 G	8.41 G
Total Bits	4.67 G	24.08 G	45.10 G	64.74 G



SPECTRUM Report Gateway


Updated: Mon Jun 15 08:00:32 2009 MEST

PROTOCOL / TOBI OETIKER

ded is  
ain  
r 1  
cover  
n  
ded  
ed  
d their

# LHCOPN Total Traffic

RRD2TOOL / TOBI OETIKER




	Min	Avg	Max (avg)	Max (max)
ASGC	0.00	1.77 G	6.65 G	7.97 G
CNAF	0.00	1.49 G	5.03 G	6.01 G
KIT	0.00	2.72 G	7.75 G	10.27 G
IN2P3	204.57 M	3.27 G	9.46 G	9.87 G
NDGF	0.00	435.28 M	1.75 G	2.07 G
PIC	5.39 M	3.34 G	9.53 G	12.12 G
RAL	153.39 M	3.62 G	9.35 G	10.56 G
NLT1	26.03 k	1.93 G	7.20 G	8.29 G
TRIUMF	0.00	713.47 M	2.72 G	4.28 G
BNL	50.32 M	3.12 G	8.82 G	11.21 G
FNAL	166.72 M	1.67 G	5.80 G	8.41 G
<b>Total Bits</b>	<b>4.67 G</b>	<b>24.08 G</b>	<b>45.10 G</b>	<b>64.74 G</b>




SPECTRUM Report Gateway

Updated: Mon Jun 15 08:00:32 2009 MEST






# Grid Activity – distribution of CPU delivered




- Distribution of work across Tier0/Tier1/Tier 2 really illustrates the importance of the grid system
  - Tier 2 contribution is ~ 50%;
  - >85% is external to CERN

## Tier 0 + Tier 1 sites



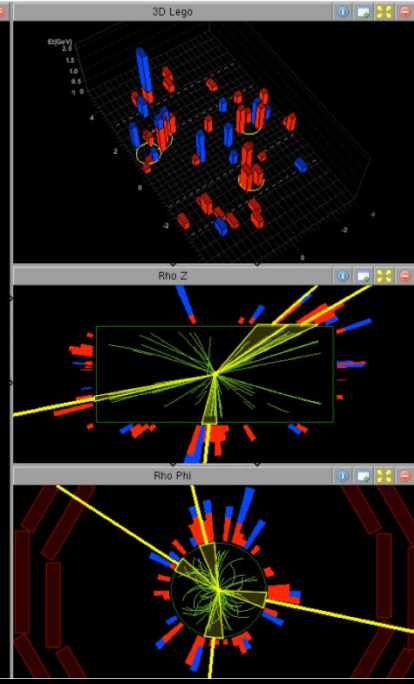

## Tier 2 sites




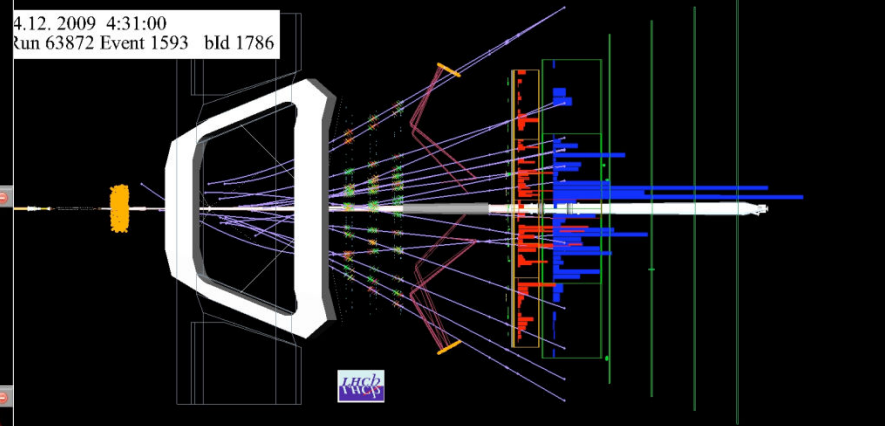
- |                       |                           |
|-----------------------|---------------------------|
| ■ T2_US_Purdue        | ■ RU-RDG                  |
| ■ DE-DESY RWTH CMS T2 | ■ US-MW2                  |
| ■ FR-GNF              | ■ FR-IN2P3-CC-T2          |
| ■ ATLAS               | ■ T2_US_UCSD              |
| ■ T2_US_MII           | ■ T2_US_Wisconsin         |
| ■ UK4-London-Tier2    | ■ US-SWT2                 |
| ■ JP-Tokyo-ATLAS-T2   | ■ UK-NorthGrid            |
| ■ RO-LCG              | ■ IT-ATLAS-federation     |
| ■ US-AGIT2            | ■ SF-SMIL-T2              |
| ■ T2_US_Florida       | ■ T2_US_Caltech           |
| ■ US-SWT2             | ■ SI-SIGNET               |
| ■ PL-TIER2-WLCG       | ■ US-IBEL2                |
| ■ UK-SouthGrid        | ■ UK-ScotGrid             |
| ■ TW-TT2              | ■ TW-TT2                  |
| ■ ES-ATLAS            | ■ IT-CRIS-federation      |
| ■ ES-CMS-T2           | ■ PT-LIP-LCG-Tier2        |
| ■ FR-IN2P3-APC        | ■ FR-ALICE                |
| ■ IT-AUCL-federation  | ■ DL-IILIB/IGW/UPP/ITALIA |
| ■ FR-IN2P3-SURAT/FCM  | ■ DE-DESY-ATLAS-T2        |
| ■ CA-WFST-T2          | ■ CH-CHIPP-CSCS           |
| ■ HO-HORDGRID-I2      | ■ FR-IN2P3-LAPP           |
| ■ DE-MCAT             | ■ C7-Prague-T2            |
| ■ CH-HEP              | ■ BE-TIER2                |
| ■ IJU-IGCC-T2         | ■ KR-KISTI-T2             |
| ■ IT-LHCb-federation  | ■ ES-LHCb-I2              |
| ■ TR-Tier2-federation | ■ IN-INDIA-CMS-TIER       |
| ■ AT-HEP1-VIENNA-UBK  | ■ IN-DAC-KOLKATA-TIER2    |
| ■ PK-CMS-I2           | ■ DE-GSI                  |
| ■ CA-CAST-T2          | ■ FI-ITP-T2               |
| ■ IT-HEP-Tier2        |                           |






CMS Experiment at the LHC, CERN  
Date Recorded: 2009-12-14 05:41 CET  
Run/Event: 124120/16701049  
Candidate Multijet Event at 2.36 TeV



4.12.2009 4:31:00  
Run 63872 Event 1593 bld 1786



### 2-Jet Event at 2.36 TeV



**ATLAS**  
EXPERIMENT

26 January 2010  
2009-12-08, 21:40 CET  
Run 142065 Event 116969

Frédéric Hemmer



# CERN openlab in brief

- A science – industry partnership to drive R&D and innovation
- Started in 2002, now in phase 3

*Motto: “you make it – we break it”*

- Evaluates state-of-the-art technologies in a very complex environment and improves them
- Test in a research environment today what will be used in industry tomorrow
- Training:
  - openlab student programme
  - Topical seminars
  - CERN School of Computing

- Covers 2009-2011
- Status
  - Partners: HP, Intel, Oracle and Siemens
- New topics
  - Global wireless coverage for CERN (HP)
  - Power-efficient solutions (Intel)
  - Performance Tuning (Oracle)
  - Control systems and PLC security (Siemens)
  - Advanced storage systems and/or global file system (partner to be identified)
  - 100Gb/s networking (partner to be identified)

# openlab people: students in 2009





# More information



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



**GridCafé**  
[www.gridcafe.org](http://www.gridcafe.org)



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



[www.nordugrid.org](http://www.nordugrid.org)



[www.eu-egi.org/](http://www.eu-egi.org/)



[www.eu-egee.org](http://www.eu-egee.org)