e-FISCAL Summer Workshop

# Opportunities

# A Realistic Study of ~~Costs~~ Associated to Datacenter Installation and Operation *in a Research Institute*

*can we do EVEN better?*

*Samos, 3rd July 2012*


Jesús Marco de Lucas

*CSIC Research Professor* @ **Instituto de Física de Cantabria**

marco@ifca.unican.es



**thanks to: R.Marco, I.Cabrillo, P.Orviz, A.Lopez, L.Cabellos, M.A.Nuñez, I.Campos**

# Outlook

Where we do come from?

Where are we?

- Computing in a Research Institute
- Operating a Grid
- HPC and Supercomputing
- What about Cloud?

Where do we go?

- The real BUSINESS MODEL

# *Where we do come from?*

We have walked for long…

- >200.000 years Homo Sapiens
- >2.500 years ***μάθημα***
- >500 years "modern" math/phys

We have also computed for long…but not so long…

- Alan Turing Centennial!

# Computing in a Research Institute

## Instituto de Física de Cantabria, Santander, SPAIN

Joint center CSIC (National Research Council)-University of Cantabria

Around 100 researchers (25 senior staff + postdocs, contracts, fellows)

## What are our research fields?

- High Energy Physics (LEP,Tevatron, now mainly LHC)
- Astrophysics (XMM, Planck, next EUCLID)
- Statistical Physics and Meteorology

## How do we compute?

- Individual workstations, small clusters
- Datacenter
- Integrated in a e-Infrastructure

# Datacenter at IFCA

## The “complete” experience

- Built from “scratch” in 2004 to host medium size clusters
  - Conditioning of around 120 m2 in basement: floor, painting
  - Basic cooling (40KW), Basic UPS (15 KW)
  - Electrical boards (80 KW)
- Upgrade to host supercomputing node in 2006
  - Refrigeration Unit (+65kW), Second UPS (15 KW)
  - Technical floor , new electrical board (+ 80 KW)
  - Fire and alarm systems
- Upgrade to host large cluster (for Grid computing) in 2008
  - Second Refrigeration Unit (+65kW)
- Improvements to cooling “crisis”
  - Direct external air

# GRID enabled clusters and storage

- CLUSTERS: ASTRO (6x48) + CMS (36x12) + GRID-CSIC (182x8) ~ 2200 cores
- STORAGE: CMS Tier-2 (1 Pb) + GRID-CSIC (300 Tb) + IFCA-CSIC (1 Pb) ~2,3 Pb
- NETWORK: 10Gb backbone+ with direct link to RedIris Nova (dark fiber)
- Cluster/Storage/Grid managers: Iban Cabrillo, Pablo Orviz, Alvaro López
- Funding: CSIC (50%), University (25%), HEP project Tier-2 CMS (25%)
- TIME SHARING: 25% local & institutional, 75% Ibergrid –EGI including LHC-CMS
- Important operational points:
  - Services virtualized (requiring around 10% extra servers, including login nodes)
  - Unified storage under GPFS
  - Single batch system using GridEngine, both for local cluster and Grid users

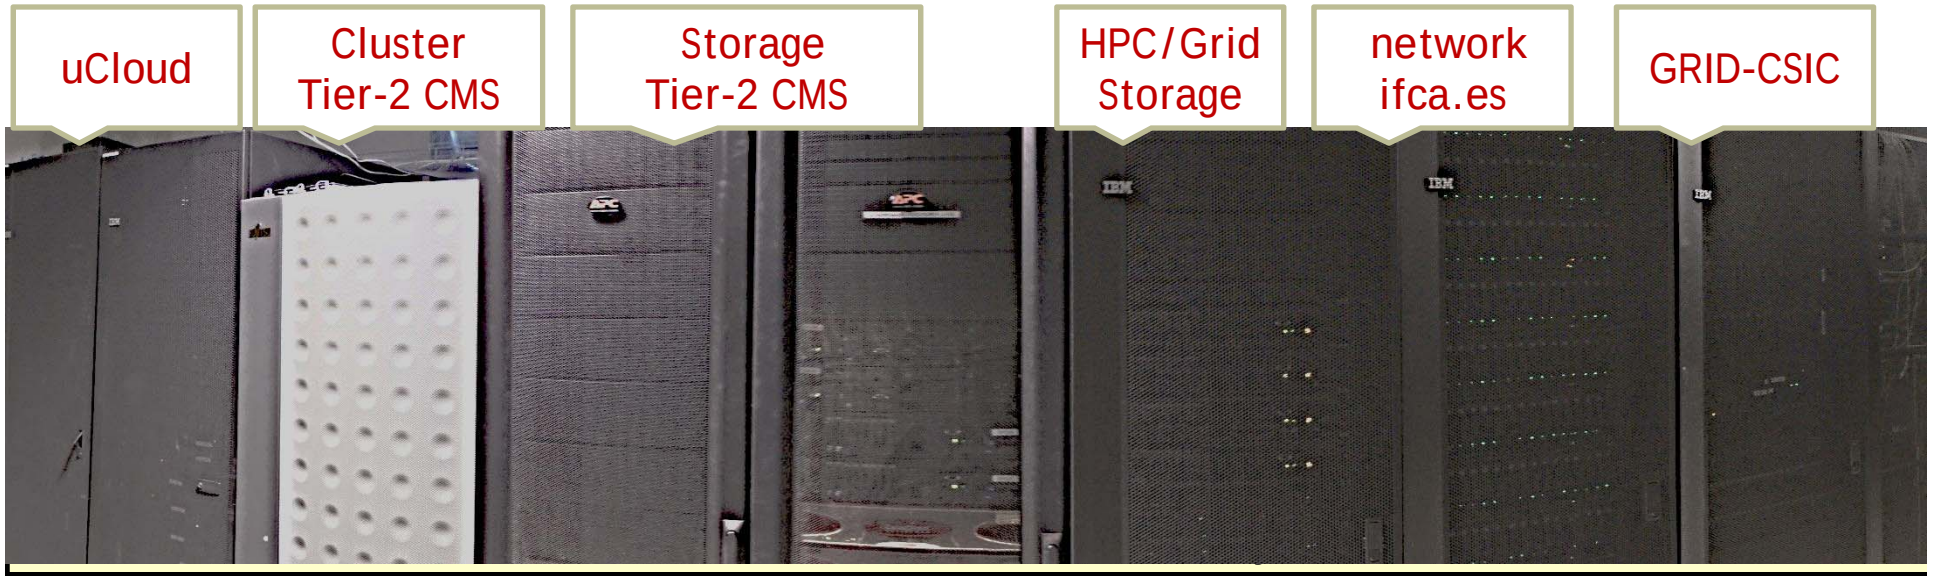

# Supercomputing node at UC

- CLUSTER: INTEL SB (160x16) + POWER7 (11x16) + GPU (512x10) ~ 7800 cores
- POWER >60 Teraflops (52 Tflop + 4 Tflop + 5 Tflop ) (top500.org with IH)
- NETWORK: Infiniband FDR10 (40Gbps) both for MPI and for GPFS
- Cluster/Storage/SC managers: Luis Cabellos, Iban Cabrillo
- Funding: University (100%) through national call for Campus of Excelence
- TIME SHARING: 75% local & institutional, 25% Spanish Supercomputing Network (RES)

- REMARKS:
  - Unified Storage with GRID/Cluster
  - Flexibility to join GRID
  - NO HT nor Virtualization
  - **Supercomputing BIOS and Network**
  - Periodic call for proposals

# Detailed Costs

- Datacenter conditioning (for 10 years)
  - 120K€ + 100K€ + 80K€ ~ 300K€
- Datacenter energy (per year)
  - 150K€ (150KW 24h 365 days)
- Clusters GRID (5-years life)
  - 200K€+500K€+80K€+120K€ = 900K€
- Storage (7-10 years life)
  - 320K€+ 250K€ + 150K€ = 720K€
- Supercomputer node(5-years life)
  - 900K€ + 150K€ + 50K€ = 1100K€
- Network backbone(5-years life)
  - 60 K€+30K€+ 120K€ = 210 K€

- GRID managers
  - 130K€/year
- Storage manager
  - 60K€/year
- Supercomputer manager
  - 60K€ / year
- Datacenter support (2 people, part time)
  - 100 K€

YEARLY COST: 600K€ + 350K€ = 950K€

CORE COST PER HOUR @ 80% eff: **0,03 €**
(5000 cores, GPU not included, **storage included**)

# What about Cloud

## Not so “complete” experience…

Compose Reply Reply All Forward Delete Add Addresses Previous Next Close Move message to folder:

From Amazon Web Services <no-reply-aws@amazon.com>
Sent Tuesday, July 3, 2012 12:48 pm
To "marco@ifca.unican.es" <marco@ifca.unican.es>
Subject Amazon Web Services Billing Statement Available

Greetings from Amazon Web Services,

This e-mail confirms that your latest billing statement is available on the AWS web site. Your account will be charged the following:

Total: $1.89

Please see the Account Activity area of the AWS web site for detailed account information:

https://portal.aws.amazon.com/gp/aws/developer/account/index.html?action=activity-summary&statementTimePeriod=1338508800

Did you know? You can now receive billing alerts to be notified via e-mail when estimated charges reach a threshold that you choose. Visit your Account Activity page to enable or learn more at: https://portal.aws.amazon.com/gp/aws/developer/account/index.html?ie=UTF8&action=billing-alerts

Thank you for using Amazon Web Services.

Sincerely,

Amazon Web Services

This message was produced and distributed by Amazon Web Services LLC, 1918 8th Avenue, Seattle, Washington 98101-5210

# *KEYS in our "business model"*

- % use
- % "effective" use
- **% use exploiting HPC/GRID**
- **IMPACT on RESEARCH**

# *Searches in Particle Physics*

- Particle physics: studying the basic constituents of all matter around!

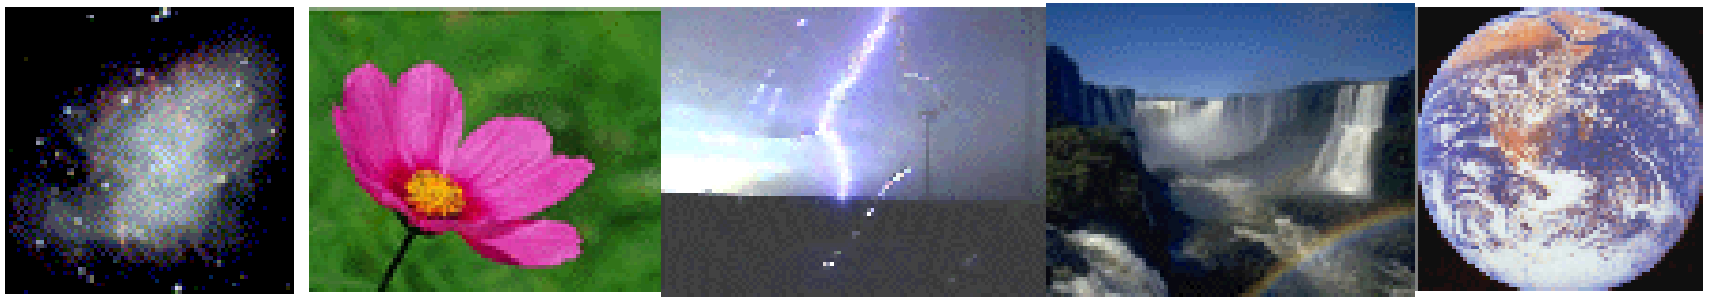

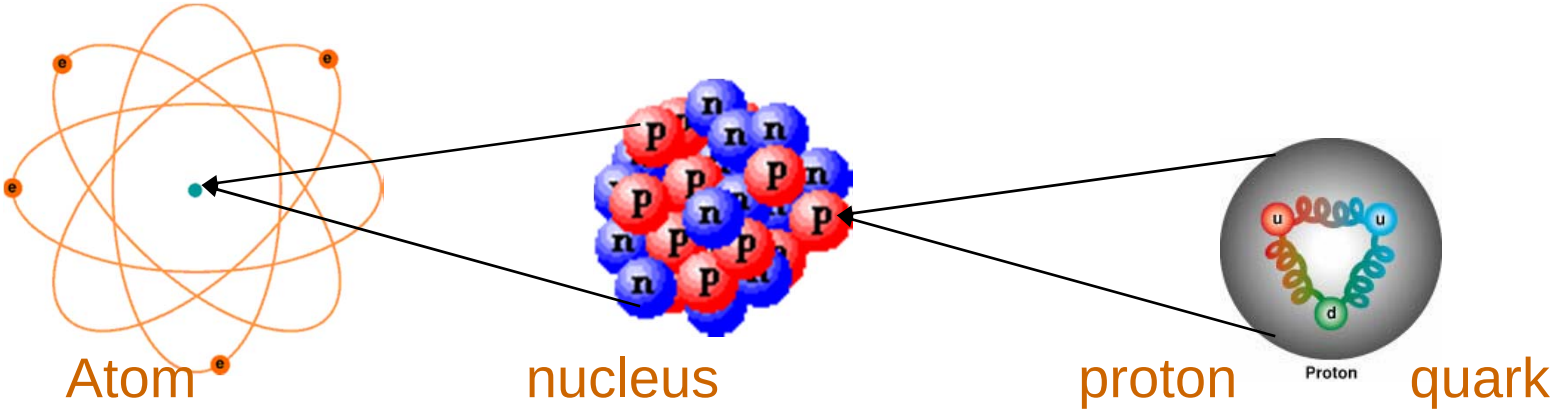


The origin of the *"mass"* of all particles is linked to a fundamental particle predicted but not yet discovered: **the Higgs boson**

# Higgs into 4 leptons

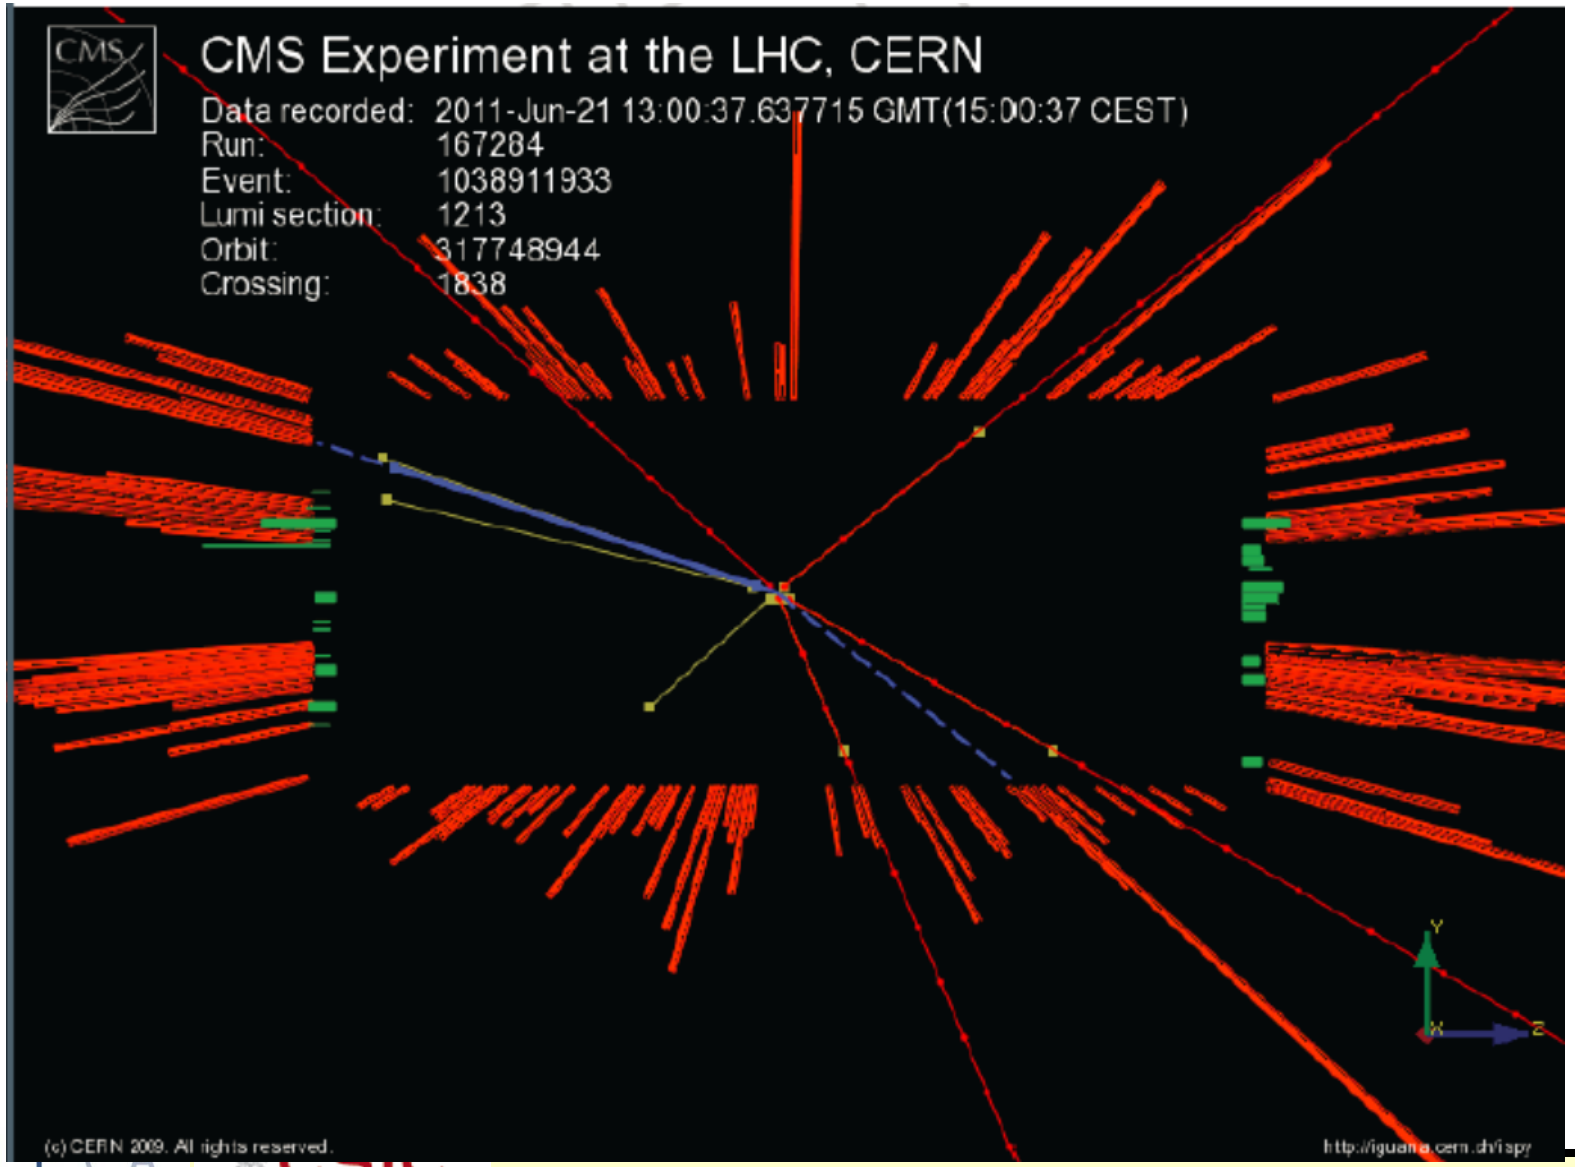

# European Grid Infrastructure
(June 2012)

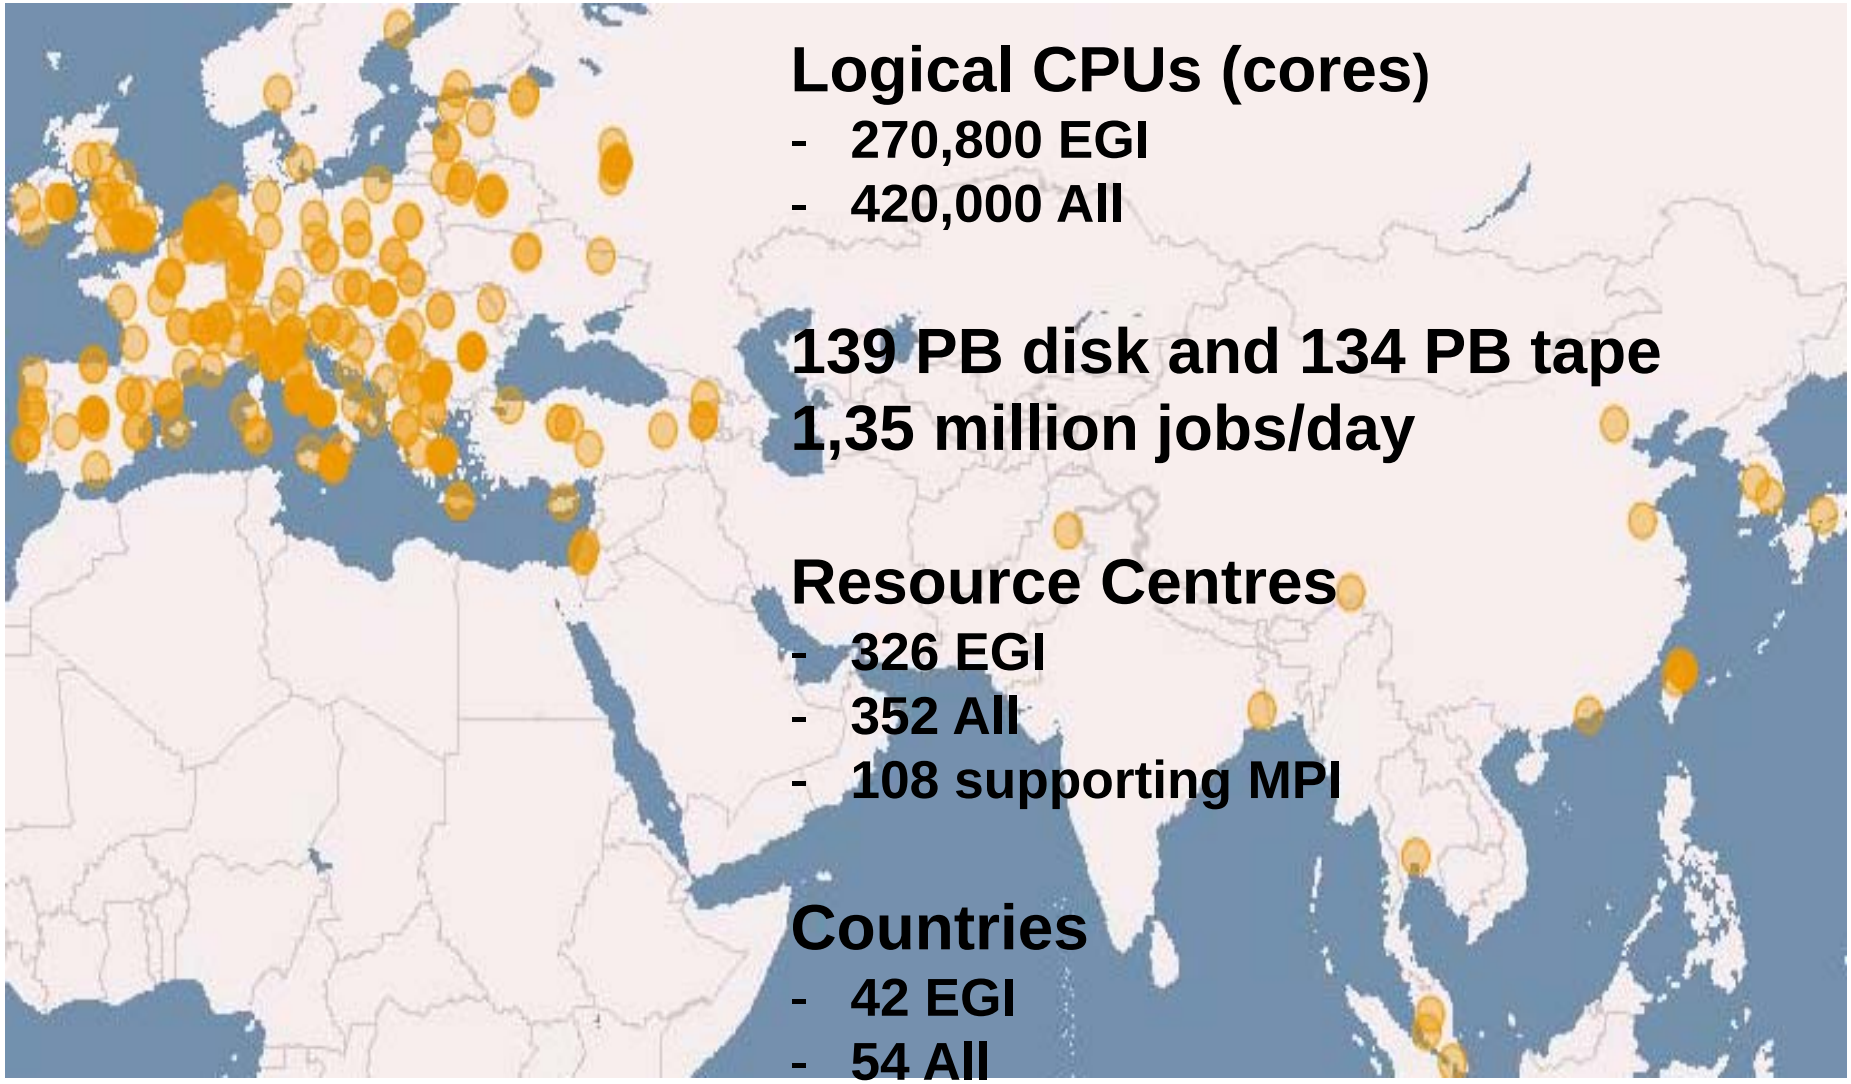

**Logical CPUs (cores)**
- 270,800 EGI
- 420,000 All

**139 PB disk and 134 PB tape**
**1,35 million jobs/day**

**Resource Centres**
- 326 EGI
- 352 All
- 108 supporting MPI

**Countries**
- 42 EGI
- 54 All

# Processing CMS data
## looking for Higgs boson events!

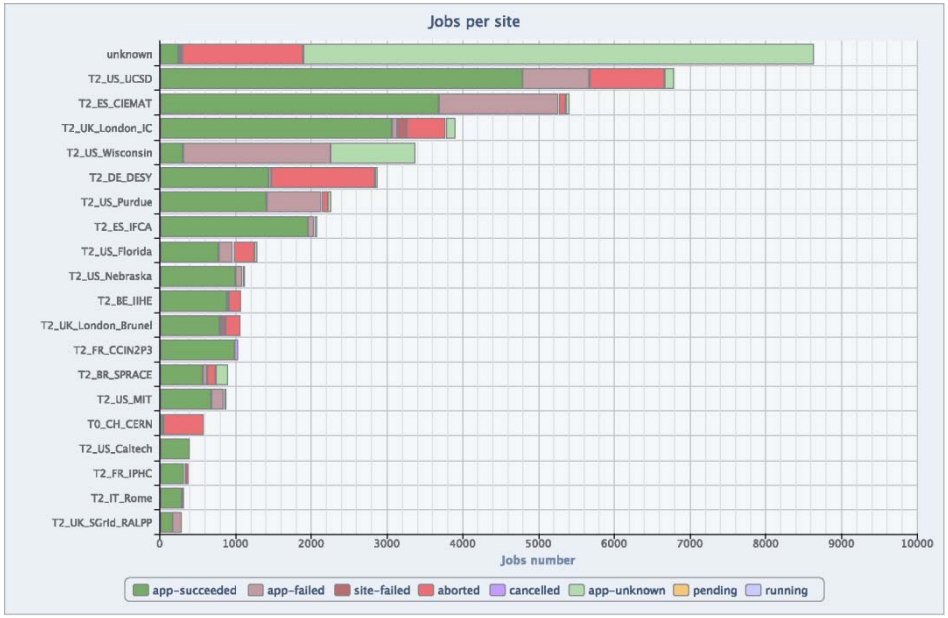


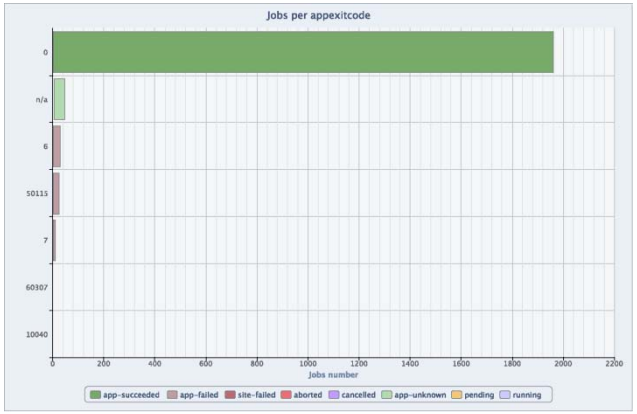


@IFCA: thousands of jobs last weeks (skimming+analysis)>99 % eff.
Executing on 60TB input, italian + spanish team

**Key support to groups in the Higgs into WW channel**

***2012: DOUBLE luminosity: DOUBLE pressure on our systems!!!***

# What next?

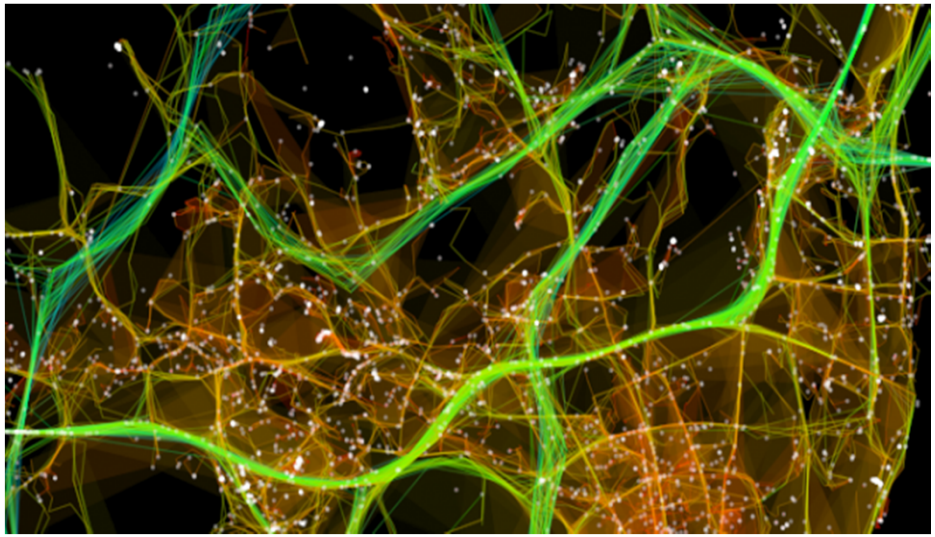

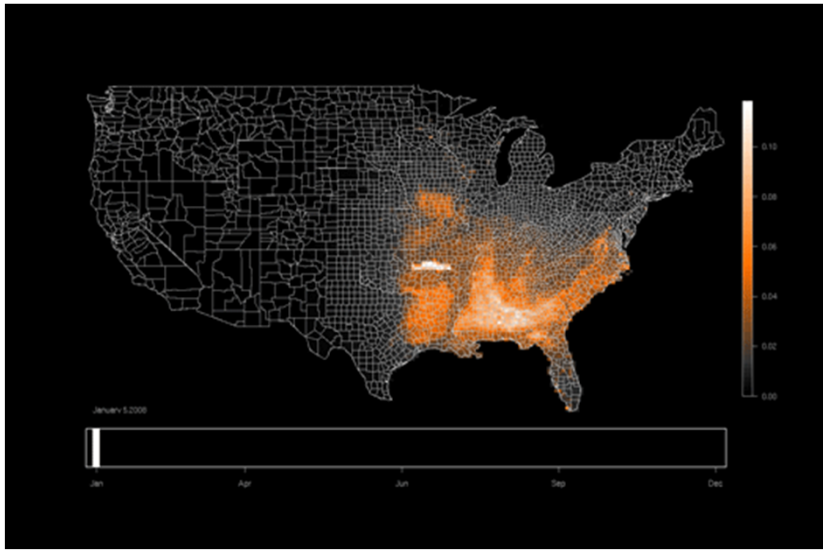

- Global objective:

**pattern matching in a context**

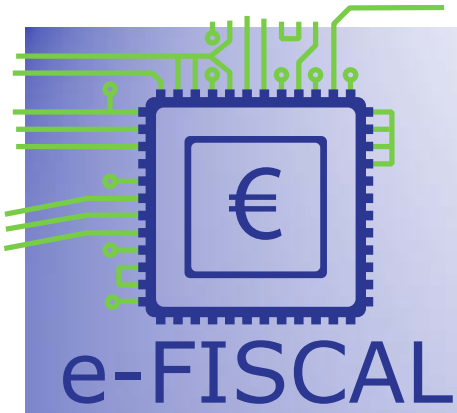


e-infrastructure CAPACITIES

# e-Infrastructure for research and science: owned, leased or hybrid approaches

SAMOS 3rd JULY 2012

e-FISCAL project, contract number RI-283449

e-FISCAL

# Provocative Statements

- SETUP A GOOD TEAM TO SUPPORT YOUR e-INFRASTRUCTURE FOR **RESEARCH**
- **GET INVOLVED IN/ ORGANIZE KEY PROJECTS**
- FORGET EXTERNAL PRIVATE CLOUDS

e-FISCAL

# Low Hanging Fruits

- MAKE IT SIMPLE (but I almost give up!)

- **MARKETING / TRAINING**
  - Did we include these “costs” in the analysis?
  - Oriented to the FINAL USERS: **RESEARCHERS**!