


Computing @ LIP-Lisbon

LIP workshop 2014

Jorge Gomes, Gonalo Borges et al

Computing @ LIP

- 0 Lisbon computing team provides IT services to LIP and its research groups:
 - Integrated management of scientific computing resources
 - Typical IT services for scientific, technical and administrative users
 - Support LIP physics research projects
- 0 e-Science and e-Infrastructures:
 - Grid Computing (WLCG and LIP-Tier-2/3)
 - Services for the scientific community
 - R&D and innovation



Competences

GRID

Batch systems (SGE, Torque, MAUI)

Parallel computing

Infrastructure planning
deployment and
coordination

High Performance
Filesystems

Middleware management

Datacenter management

CLOUD

Implementations
(OpenNebula,
OpenStack)

Federated
Clouds

General IT

Systems and
fabric
management
(Linux, FreeBSD,
Windows)

Network
Management and
Network Security

Virtualisation
(KVM and Xen)

Security, Incident
and Vulnerability
management,
GRID CSIRT

Authentication,
Authorization,
digital certificates
(IGTF and
EugridPMA)


Databases and
information
systems

Web applications
and design

Open Source
Software and
Solutions


LIP Computing Facilities

IT services




IT facilities

Total CPU ~2K cores
Total Storage ~500TB



NCG

- National service supported by the Portuguese Science Foundation (FCT)
- **Created by LIP** in partnership with FCCN (NREN) and LNEC (national civil engineering lab)



NCG services

0 High Throughput Computing cluster

- Mostly grid computing:
WLCG, IBERGRID and EGI

0 High Performance Computing cluster

- Infiniband (low latency) and GbE interconnects
- Parallel applications (MPI and OpenMP)

0 High performance distributed storage

- Lustre for high throughput computing (grid)
- Lustre for parallel applications
- NFS for software and home directories

0 Cloud Computing pilot service

- OpenStack




NCG and LNEC



LABORATÓRIO NACIONAL
DE ENGENHARIA CIVIL

○ LNEC HPC cluster (Medusa)

- ~ 60 servers, ~ 280 Cores
- Seamlessly integrated in the NCG computing platform
- Managed by the LIP team



IT based services activities (operations, support and R&D)

Resource and application support

- 0 Automatic installations using KICKSTART / PXE
 - SL profiles (for HEP research activities): interactive services, computing cluster, desktops...
 - CentOS profiles (for services): storage services, frontends, central services, ...
 - Fedora profiles (for users): desktops
- 0 Application porting
- 0 Software configuration based on “Modules”
 - Empower the user

Manage heterogeneous resources and multidisciplinary applications on an evolving environment

LIP WEB / Design


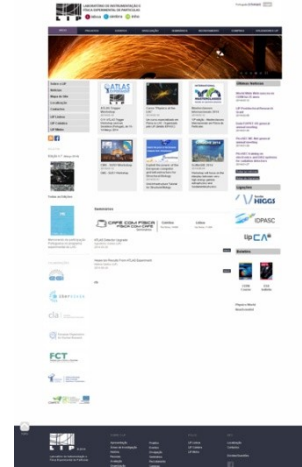
0 Institutional Web Pages

- 5 sites (Portal, LIP Users, 3 sites)
- Centralized data in Projects DB
- + 20.000 Unique Visitors (2013)
- +210.000 Pages Views (2013)

0 Events

- Home Made Management System
 - Web pages, Registration
 - Secretariat Area (Badges, receipts, Certificates, ...)
- + 50 Events LIP/Projects, 20 Events IDPASC, Posters Design

Institutional Web Pages



Events



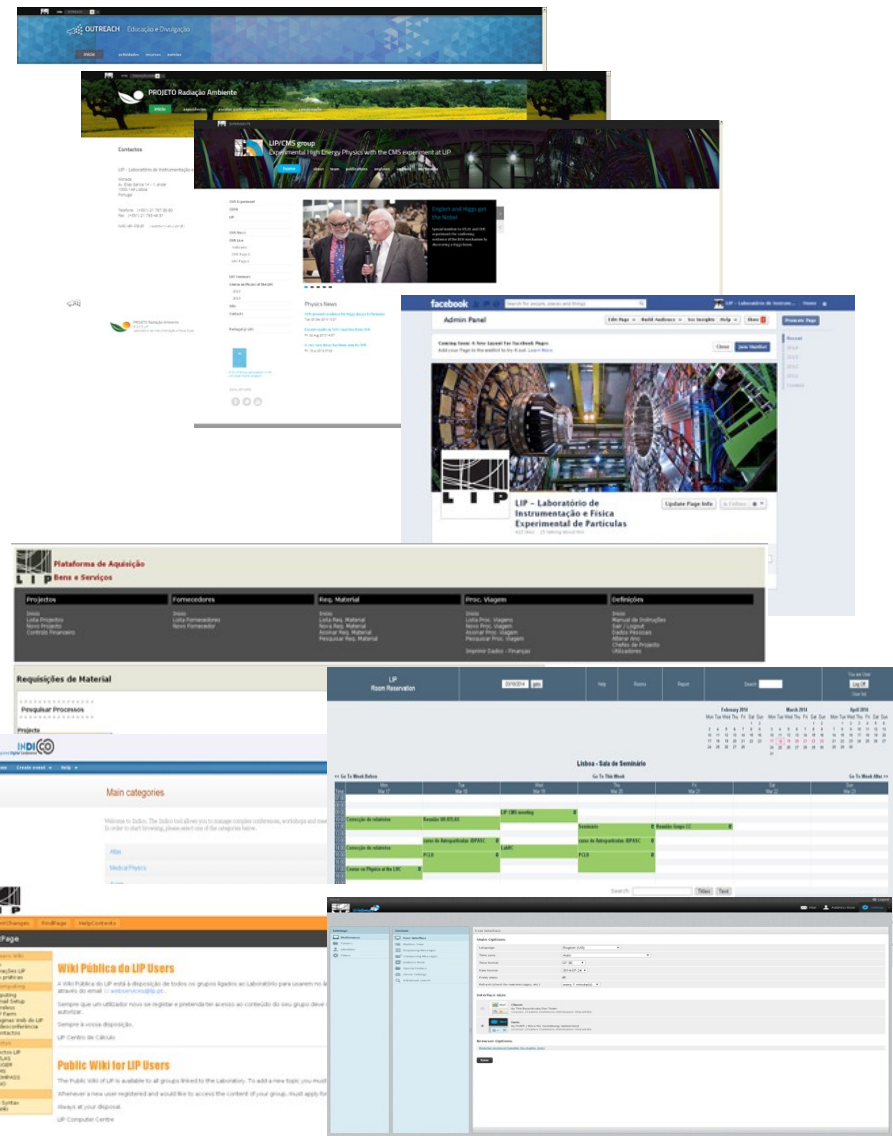
LIP WEB / Design

Corporate Image

- Experiments and Outreach Websites Reformulation
- Logos, Documents
- Facebook

User Support Tools

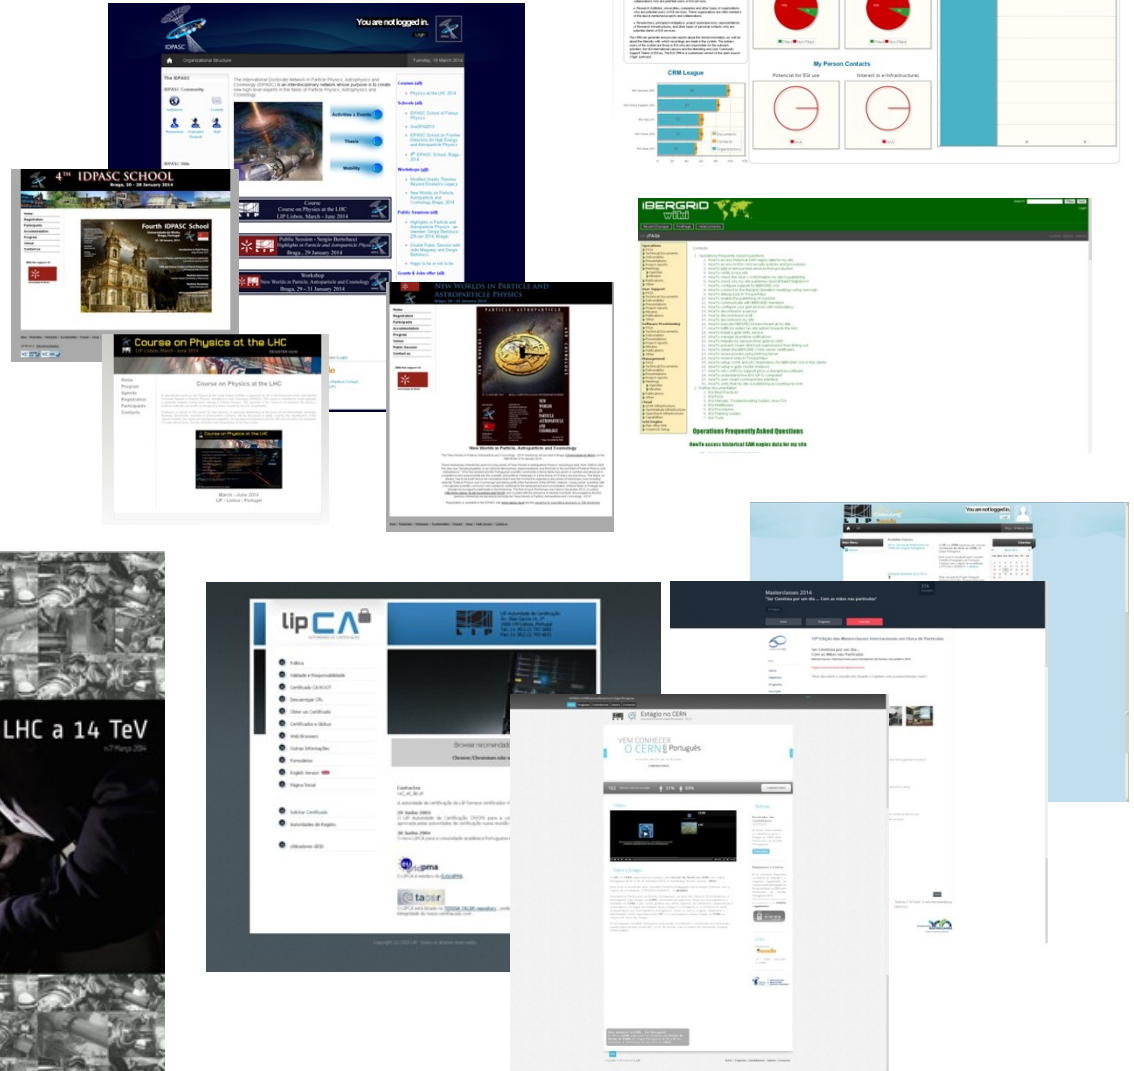
- Conference Room Booking
- Requisitions Material/Missions
- Secretariat support
- Users Wiki's, Webmail



LIP WEB / Design

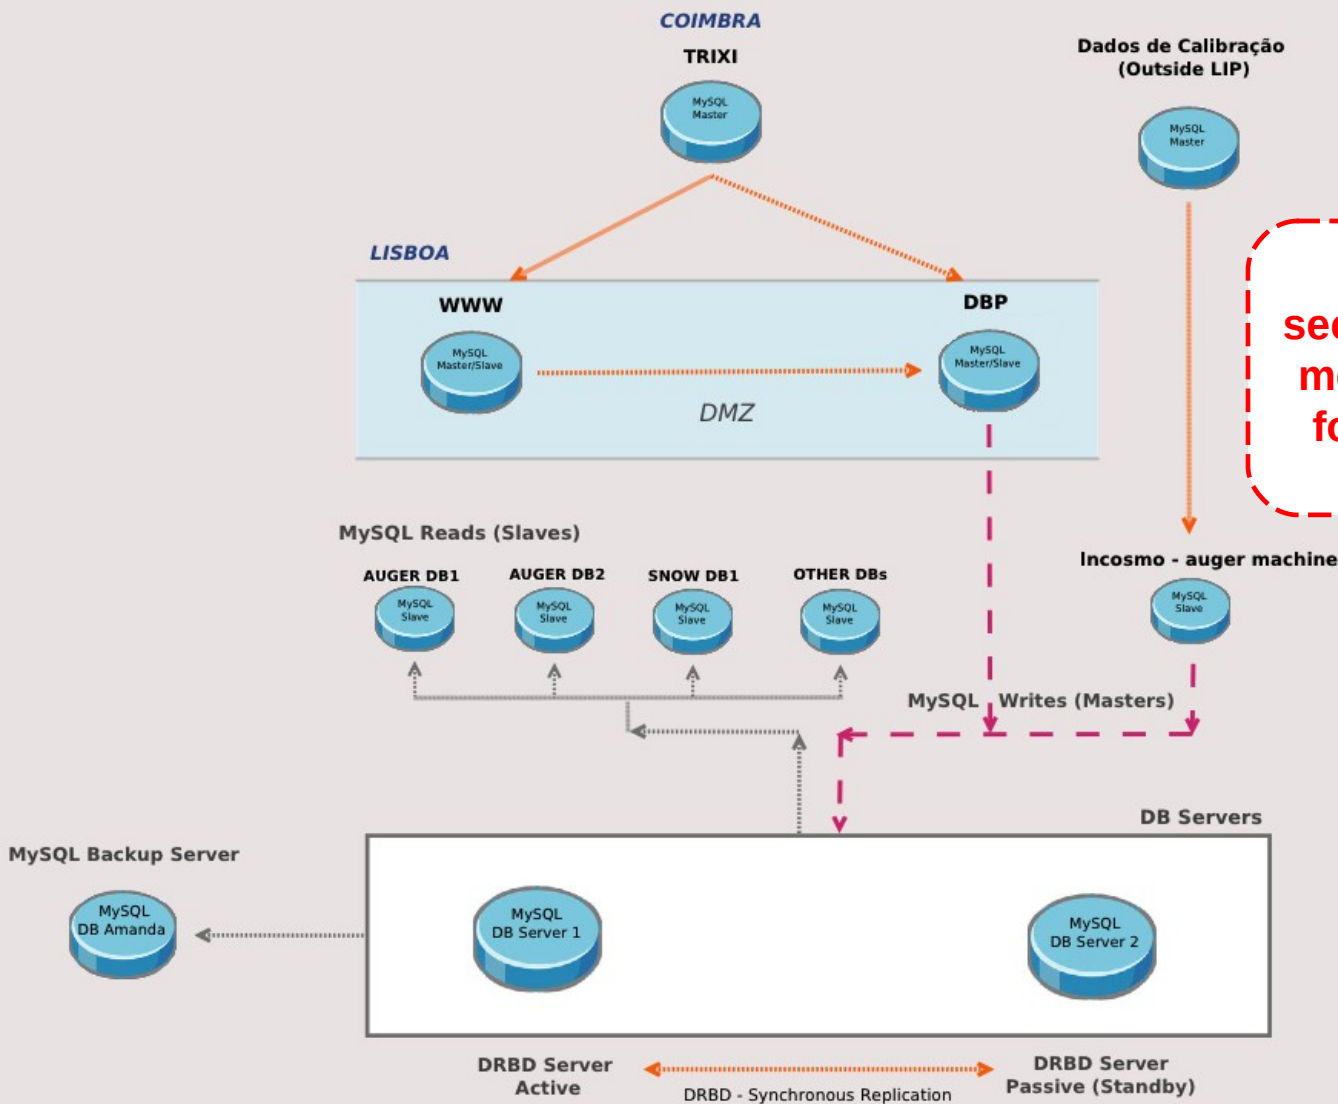
Other Activities

- Bulletin
- IDPASC Network
- International Tasks
- LIP CA, EGI CRM



MySQL DB failover

Transparent failover,
security and load balancing
mechanism for DB access
for experiments and web
applications




WAN Networking

○ LIP centres in a single cloud

- 10 Gbits/s between LIP sites (Layer 2)
- 3 Gbits/s to the academic network (Layer 3)


Highest consumer of the Portuguese academic network bandwidth



Virtualisation and DRBD replication


- Heavy use of virtualization
 - Started with XEN
 - Migrated to KVM
 - Managed through libvirt
 - Majority of critical services running as VMs
 - Images replicated (via DRBD)

**Increase services uptime
and resilience to hardware
failures**




R&D

- Storage consolidation (avoid the proliferation of servers)
- Understand the usage profile and take educated decisions regarding the next technologies to adopt



R&D

LIP solution to power-off FARM idle resources



10% power consumption
reduction over 1 year
of operation

Accepted for publishing
in Computing and
Informatics journal

Grid Computing

WLCG Tier-2 / Tier-3 summary

- 0 In 2013 the Tier-2 delivered 104% of the pledged HS06 processing capacity (ATLAS+CMS)
- 0 A reduction of Tier-2/Tier-3 storage capacity was applied to both experiments (~ 12%)
 - Equipment is getting older (more frequent hardware failures)
 - Reserve hardware to perform repairs
 - Need to fit equipment within NCG available rack space
- 0 The LIP-Lisbon Tier-2 and Tier-3 services have been migrated
 - Most to NCG (the CMS T2/T3 is now fully housed at NCG)
 - Some ATLAS resources were moved to Coimbra
- 0 Coimbra datacenter services were fully reorganized during August


Tier-2 and Tier-3

X – migrated in 2013

CPU	ATLAS		CMS		Total
Sites/Tiers	T2	T3	T2	T3	T2+T3
Lisbon	X	570 HS	X	X	570 HS
Coimbra	950 HS	300 HS	-	-	1250 HS
NCG	2250 HS	-	3200 HS	870 HS	6320 HS

Storage	ATLAS		CMS		Total
Sites/Tiers	T2	T3	T2	T3	T2+T3
Lisbon	X	10 TB	X	X	10 TB
Coimbra	95 TB	30 TB	-	-	125 TB
NCG	95 TB	-	170 TB	60 TB	325 TB


Tier-2 Reliability




Tier-2 CPU Accounting

Developed by CESGA EGL View: / normcpu-HEPSPEC06 / 2010-6-2014-2 / SITE-DATE / lhc (x) / ACCBAR-LIN / x

2014-03-11 01:56




European Tier-2s in 2010-2013




Tier-2 support project

- 0 Project funded by FCT
 - Operation and maintenance of the Tier-2
 - Started in 2013 duration 3 years
 - Got the maximum funding from FCT (500.000 Euro)
 - Both the team and the proposal evaluated as excellent
- 0 Status:
 - Development of tools to collect and analyse data access
 - Conducting a study of the Tier-2 storage access patterns
 - Market and technology surveys
- 0 On-going:
 - Prepare a tender for the partial renewal of the Tier-2 storage

Portuguese National Grid Activity 2013




IBERGRID non-LHC Grid Usage 2013

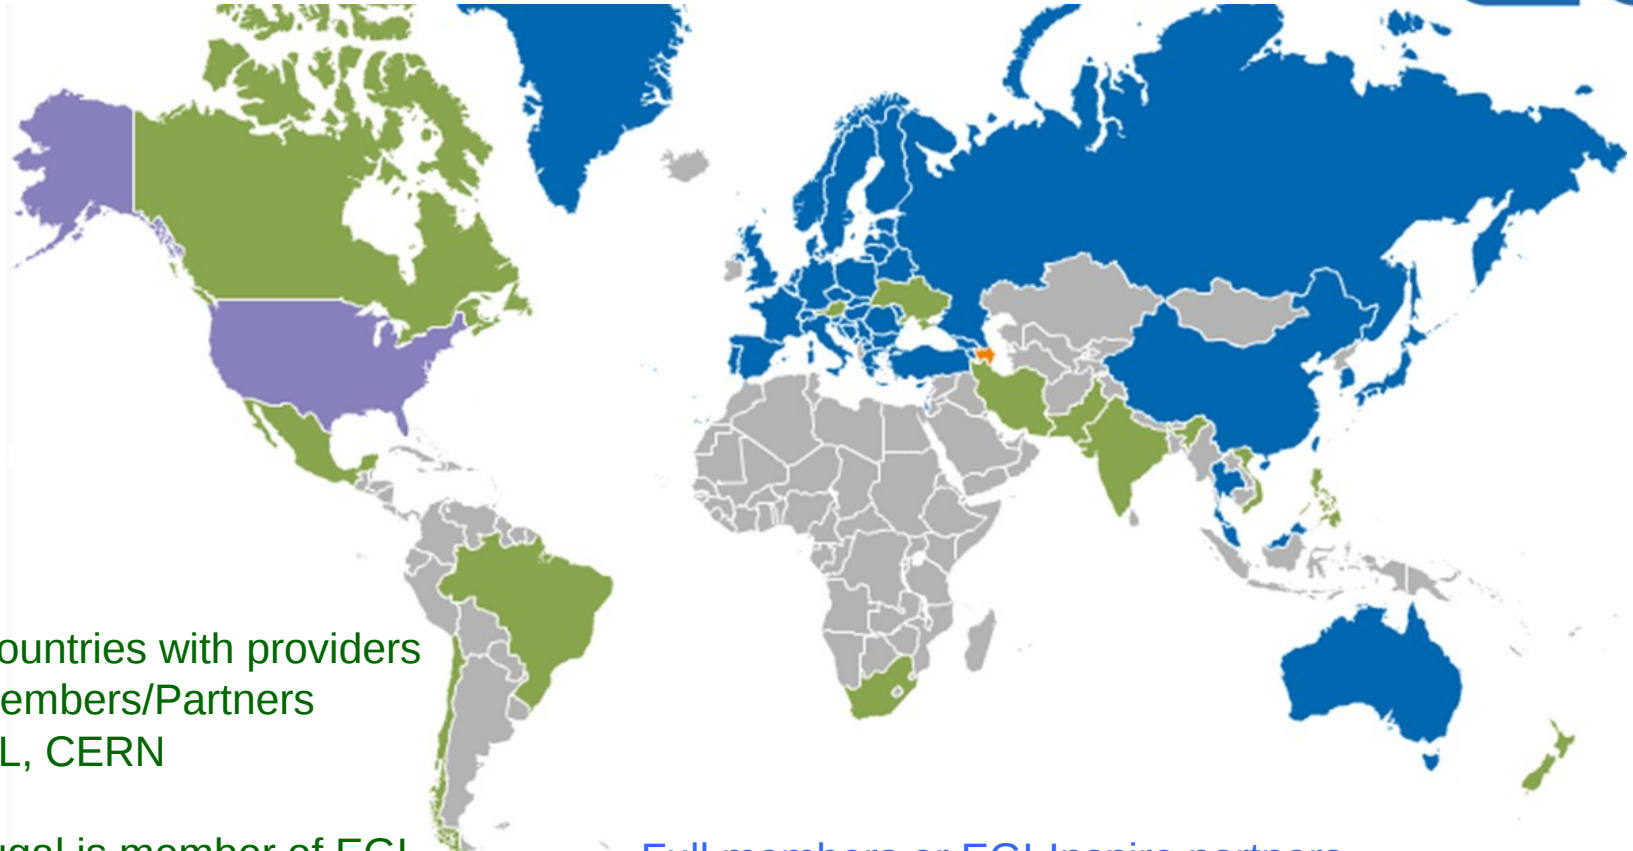


Distributed Computing Projects and Infrastructures

History of Projects and Initiatives



European Grid Initiative (EGI)



56 Countries with providers
42 Members/Partners
EMBL, CERN

Portugal is member of EGI
through FCT

Full members or EGI-Inspire partners
External integrated resource providers
Peer resource providers

EGI

0 Federated infrastructure:

- Resource providers:
 - **347 sites**
 - **Universities**
 - **Research centres**
- Grouped in NGIs
- Capacity:
 - **462,000 CPU cores**
 - **235PB disk storage**
 - **176PB tape storage**

0 User communities:

- **233 virtual organizations**
- **ATLAS, CMS, AUGER, SNO+,**
have virtual organizations
(WLCG uses EGI services)




42 million jobs/month

LIP in the European Grid Initiative



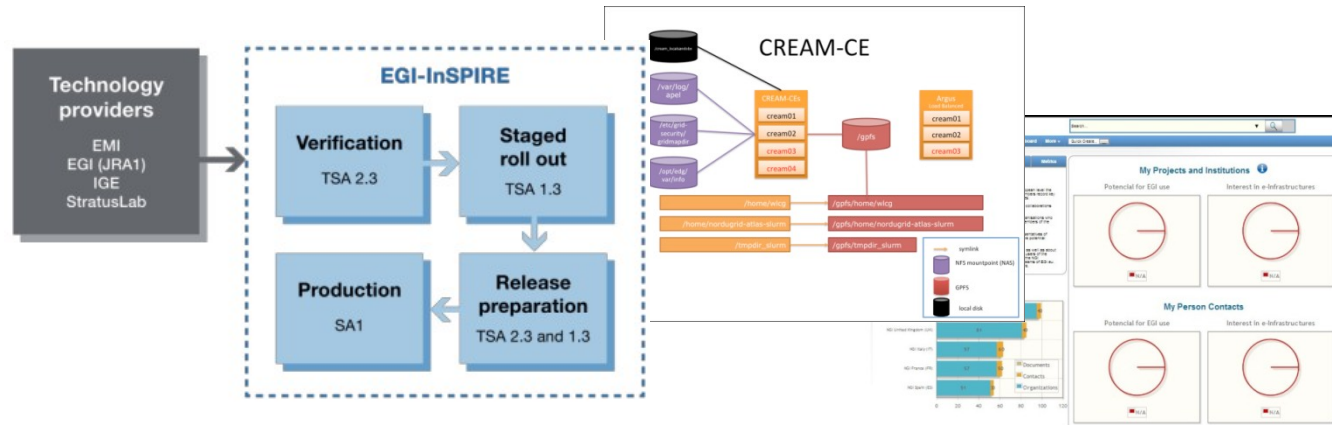
- LIP is partner of the EGI-Inspire project (2010-2014)
- Liaison between EGI and the Iberian grid infrastructure:
 - Infrastructure coordination
 - Policy, training, dissemination and relations with user communities
- Tasks (Iberian and National):
 - Regional support for users and sites
 - Regional Operations on Duty (ROD) shifts
 - Operation of core infrastructure services
 - Core grid computing services
 - Development and improvement of infrastructure tools and services
 - Certification authority for Portugal X.509



LIP in the European Grid Initiative



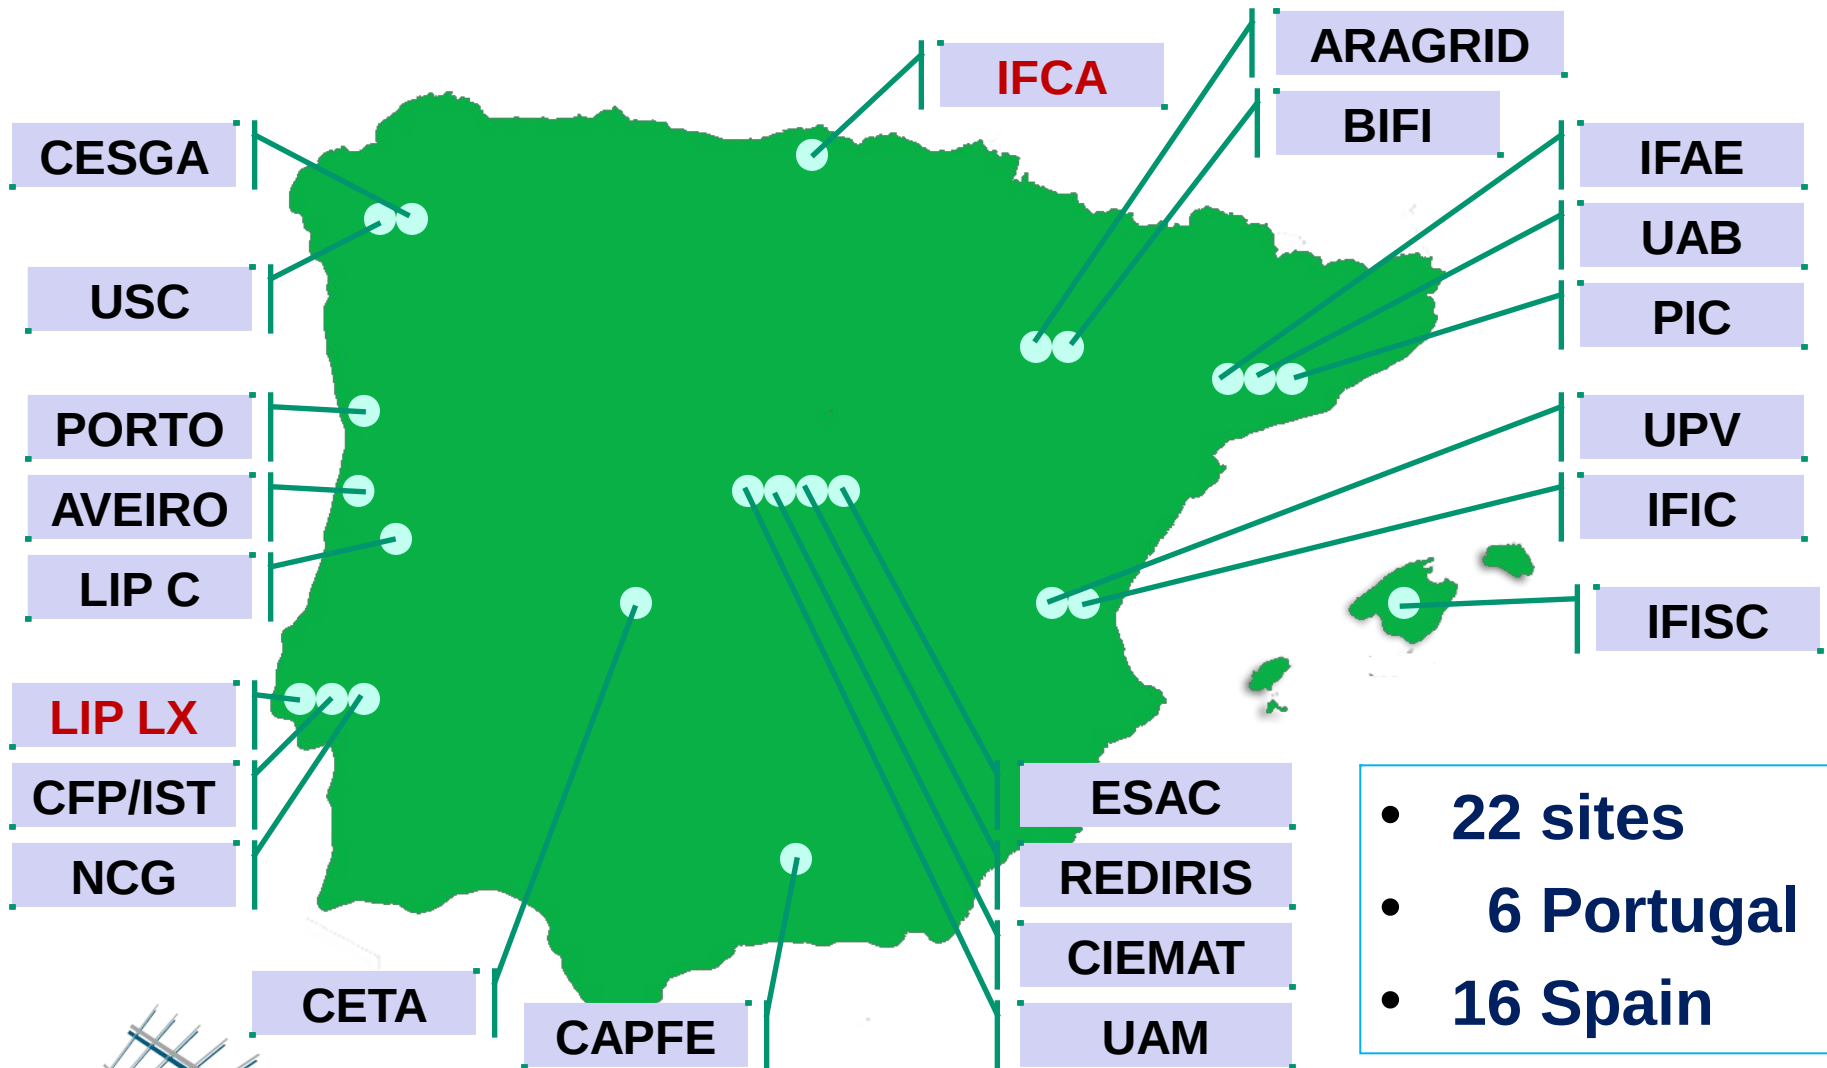
- Tasks (International / Global):
 - Global Coordination of the middleware rollout process
 - Global quality criteria definition and validation of the middleware
 - Coordination/management vulnerabilities and intrusions both in the region and at European level
 - Global support for GE-CREAM-CE
 - Development of a CRM solution for EGI based on vTiger CRM



IBERGRID

- 0 Federated Portuguese and Spanish grid infrastructure
 - **Very successful collaboration**
 - Seamless sharing of resources, expertise and services
 - Successfully supporting international scientific user communities
 - Iberian WLCG sites are part of IBERGRID
- 0 IBERGRID and EGI:
 - **Joint operational participation in EGI (has been fundamental)**
 - Coordination of the Portuguese and Spanish sites in EGI is done through IBERGRID

IBERGRID



- 22 sites
- 6 Portugal
- 16 Spain

IBERGRID in EGI governance

- EGI foundation (EGI.eu) Executive Board (EB)
 - Highest EGI management body composed of 5 to 7 members
 - Restricted membership through council election
 - Jorge Gomes (director since 2012 appointed by EGI council)
- EGI Inspire Project Management Board (PMB)
 - Iberian representation in the core EGI-Inspire EC project mgmt
 - Jorge Gomes (until 2012 representing IBERGRID)
 - Isabel Campos (IFCA/CSIC since 2012)
- EGI Council
 - Strategic body (Jorge Gomes as FCT representative)
 - Both countries are independently represented with voting rights
 - Always pursuing a common Iberian policy and positions

IBERGRID and the EGI global tasks

○ Current Global tasks:

- Middleware rollout (LIP/IFCA)
- Middleware criteria definition & validation (IFCA/CESGA/LIP)
- VO technical services (LIP)
- EGI accounting (CESGA)

○ Near future (2014-2016):

- Middleware rollout
- Middleware acceptance criteria, definition and validation
- EGI accounting portal
- EGI 1st level support
- EGI 2nd level support

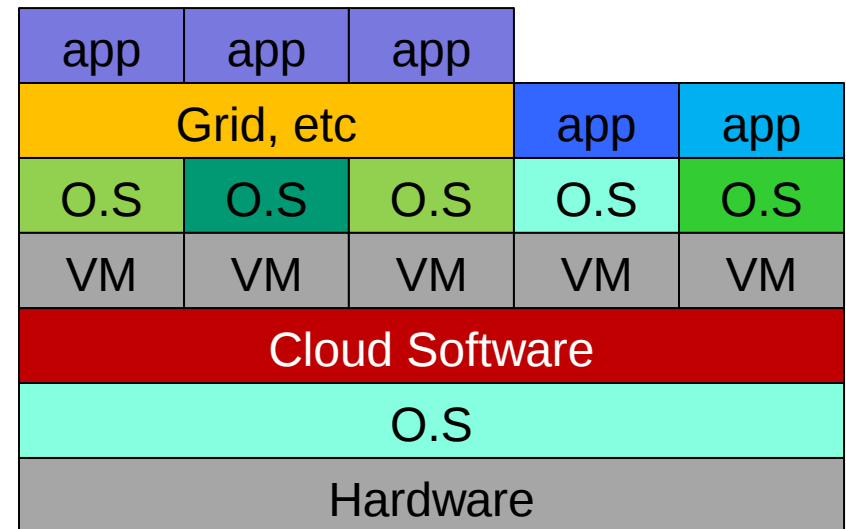
○ Global tasks:

- They are the coordination of EGI international activities
- Strategic activity and leadership at European level
- IBERGRID won again the international bid for 2014-2016
- Important positioning for future funding

Cloud Computing

Cloud experiences @ LIP

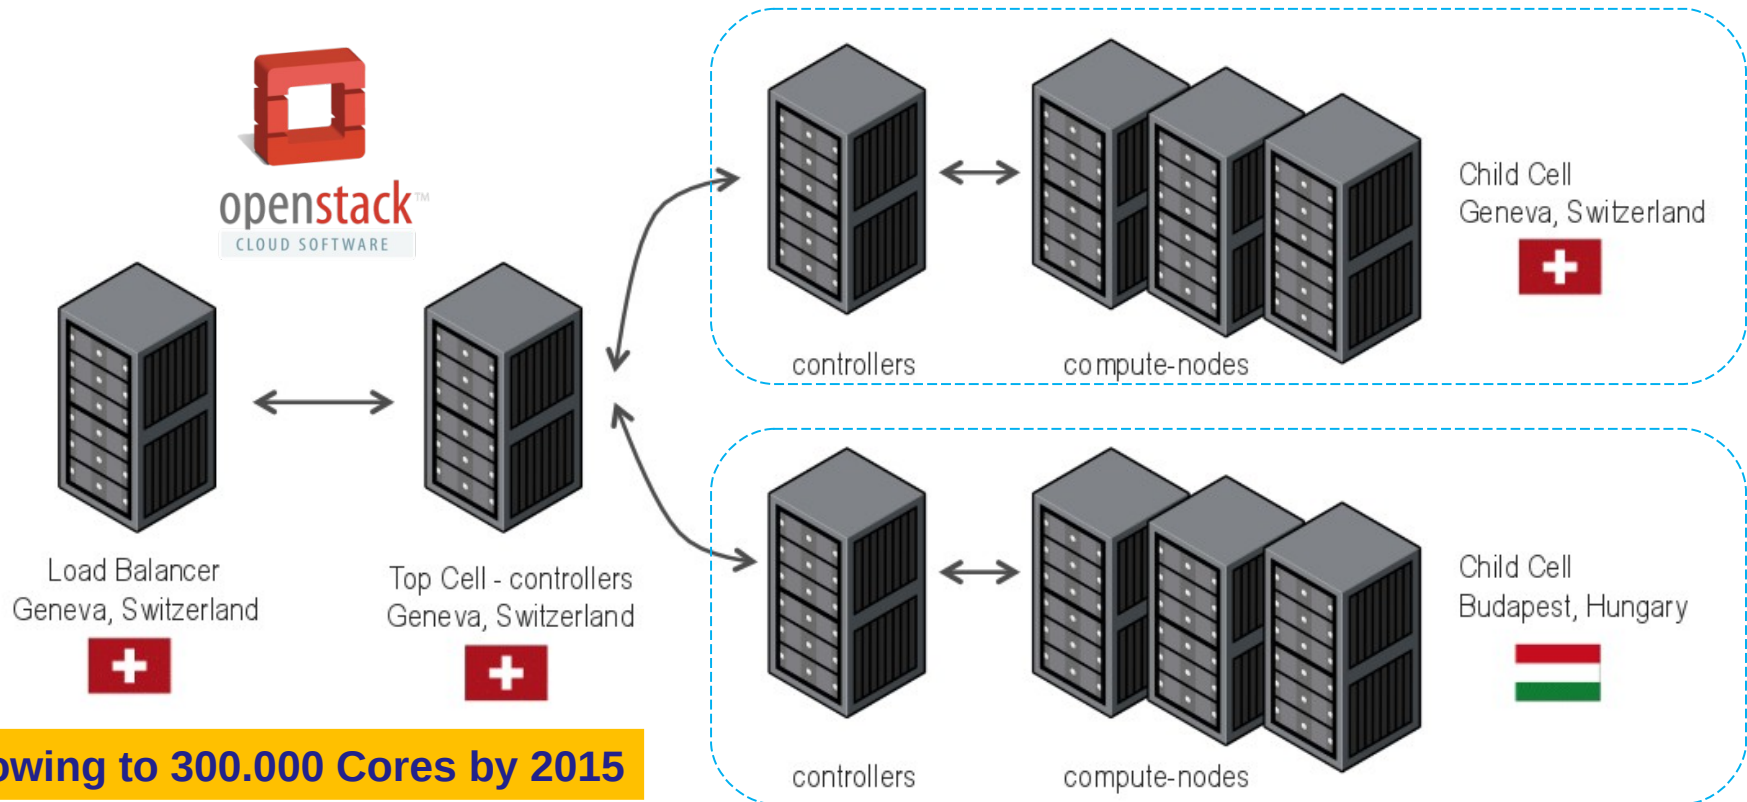
- Started with OpenNebula some years ago
- Lack of human resources
- LIP is deploying openstack and working hard to catch-up in this domain
- Very complex framework that needs to be tailored and customized
- Started by NASA and rackspace
- Backed by industry
- Hypervisors: KVM, Xen, HyperV, Baremetal, Docker, LXC



CERN Cloud





- Manage large IT platforms with highly diversified requirements
- CERN has deployed a private **IaaS cloud** infrastructure



Growing to 300.000 Cores by 2015



Wider services portfolio

farm





- Batch job scheduler
- Single farm
- HPC + HTC

grid



- Distributed batch
- Use multiple farms
- Data management
- Mostly HTC

**IaaS
cloud**



- On-demand VM allocation
- Highly flexible environment
- Empower users

Challenges and opportunities

EGI FedCloud

0 **Federated** IaaS cloud within the EGI broader distributed infrastructure:

- Aggregate institutional/private IaaS clouds
- Based on standards and open source solutions
- Support multiple cloud implementations at the site level
- Leverage EGI experience in federating distributed resources



EGI FedCloud

Communities exploiting the FedCloud Pre-Production

Ecology – BioVeL: Biodiversity Virtual e-Laboratory

Structural biology – WeNMR: a worldwide e-Infrastructure for NMR and structural biology

Linguistics – CLARIN: ‘British National Corpus’ service (BNCWeb)

Earth Observation – SSEP: European Space Agency’s Supersites Exploitation Platform for volcano and earthquakes monitoring

Software Engineering – SCI-BUS: simulated environments for portal testing

Software Engineering – DIRAC: deploying ready-to-use distributed computing systems

Software Engineering – Catania Science Gateway Framework

Musicology – Peachnote: dynamic analysis of musical scores

Earth Observation – ENVRI: Common Operations of Environmental Research infrastructures

Geology – VERCE: Virtual Earthquake and seismology Research

Ecology – LifeWatch: European Infrastructure for Biodiversity and Ecosystem Research

High Energy Physics – CERN ATLAS: ATLAS processing cluster via HelixNebula



CLARIN
Common Language Resources and Technology Infrastructure



IBERGRID towards cloud computing

- Ongoing R&D activity

- Partners: IFCA, UPV, CESGA, **LIP** and BIFI

- OpenStack: IFCA, **LIP**, BIFI

- OpenNebula: UPV and CESGA

- **Federation:**


- Currently openstack authentication federated via LDAP (LIP)

- Access to all OpenStack sites from a single entry point using modified interfaces

- Best live demo at the EGI Technical Forum 2013

- Working on federated authentication SAML and grid integration via VOMS

- Working with the EGI federated cloud task force



GRID > INCD

- 0 Infraestrutura Nacional de Computação Distribuída (INCD)
 - Proposal submitted in 2013 to the FCT infrastructures roadmap
 - Proposal recently approved (some funding expected in 2014)
- 0 Initial consortium:
 - LIP as leading partner together with FCCN and LNEC
 - With the support of Universities of Minho, Porto and Aveiro
- 0 Six year program objectives:
 - Maintenance and development of the national infrastructure
 - Enlarge the infrastructure services and users scope (cloud and data)
 - Equipment renewal looking at new requirements and technologies
 - Establish competence centres: support, development, porting, training, education, new technologies

LifeWatch

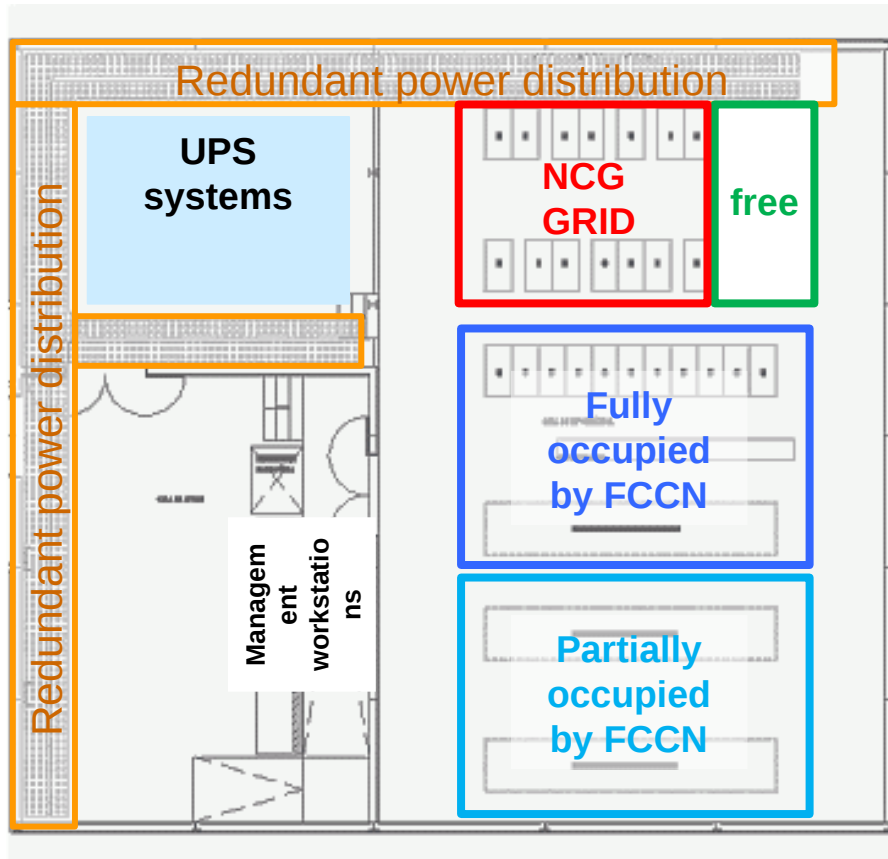
- E-science European Infrastructure for Biodiversity and Ecosystem Research (ESFRI) expected to become an ERIC
- Spain is deeply engaged in LifeWatch
 - Ministerio de Economía y Competitividad of Spain (MINECO)
 - CSIC is the technical coordinator (European HQ in Doñana Park)
 - Engagement from our Spanish colleagues in IBERGRID in the IT core infrastructures and services for LifeWatch
- Meeting at FCT last December
 - Discuss LifeWatch adhesion + reinforce and consolidate IBERGRID
 - FCT, MINECO, LIP and CIBIO
 - FCT very open regarding Portugal joining LifeWatch
 - Portuguese participation through CIBIO and INCD (in IT)

Opportunities (1)

- 0 Horizon2020, EGI and LifeWatch
 - Develop an European Federated Cloud
 - Virtual research platforms/environments on top of the cloud infrastructure
- 0 INCD
 - Further engage and reinforce collaboration at national level
 - Maintain and improve the national infrastructure
- 0 LIP
 - Working hard to prepare a pilot for a national cloud service in the INCD context together with FCCN (student from FCT UNL on cloud storage)
 - The INCD cloud service will be integrated in the EGI federated cloud
 - LIP as EGI Distributed Competence Centre supporting cloud applications and data intensive applications (LifeWatch, HEP, life sciences, ...)
 - Engage in H2020 cloud and community oriented projects
 - Opportunity to secure support to keep the LIP IT and the Tier-2/3 operational

Opportunities (2)

- 0 Further consolidation of computing services at NCG (Tier-2/3)



- Area 370 m²
- 3 years of uptime without failures
- Fully redundant power system
- Diesel backed power generator
- 1 MW capacity (500kW + 500kW)
- Water cooled redundant HVAC
- Recently ranked by AMA as one of the best datacentres in the public administration

**Scientific computing
services managed by
LIP**

IBERGRID conference series

0 IBERGRID 2013 and EGI Technical Forum 2013

- collocated both events
- took place in Madrid last September
- both events were jointly organized by IFCA/CSIC, UPV and LIP
- 500 participants




0 IBERGRID 2014

- 8th to 10th September
- will take place in Aveiro
- will be organized by Univ of Aveiro with LIP support



Computing Team @ LIP-Lisbon

- Mix of engineers, technicians and PhDs based in Lisbon
- Extensive experience in e-Science, advanced computing and distributed computing
- Boosted by grid computing R&D activities for the CERN LHC experiments
- Large reduction of team members and PhDs that are leaving



Thank you

<http://www.ingrid.pt/>