CERN-IT-GT

WLCG
Worldwide LHC Computing Grid

Markus Schulz

August 2010
Openlab Summer Students

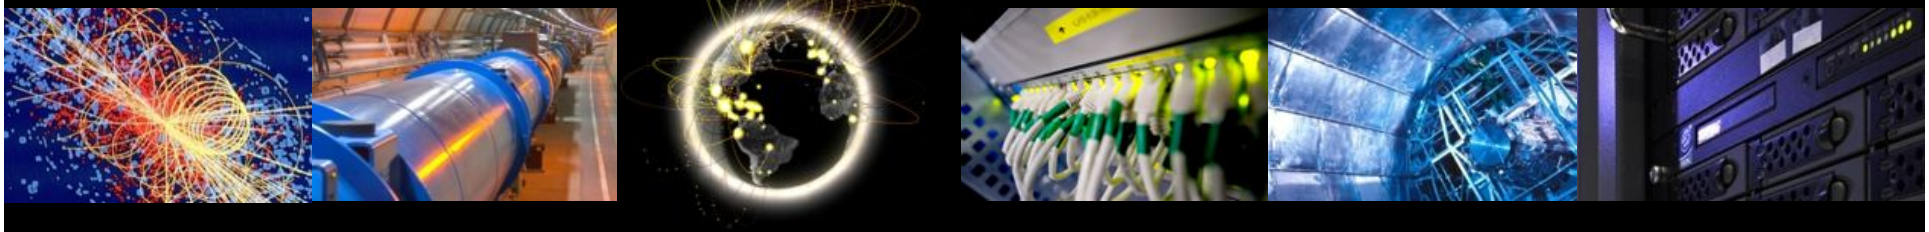

WLCG
Worldwide LHC Computing Grid

Overview

- Background
- The Infrastructure
- Usage
- Evolution

Not covered today

- Grid Computing Technology
- gLite Middleware

WLCG Worldwide LHC Computing Grid

CERN in Numbers

- 2328 staff*
- 711 Fellows and Associates*
- 9535 users*
- Budget (2009) 887 MCHF (595M€)

*6 January 2009

- Member States: Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom.
- Observers to Council: India, Israel, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and Unesco

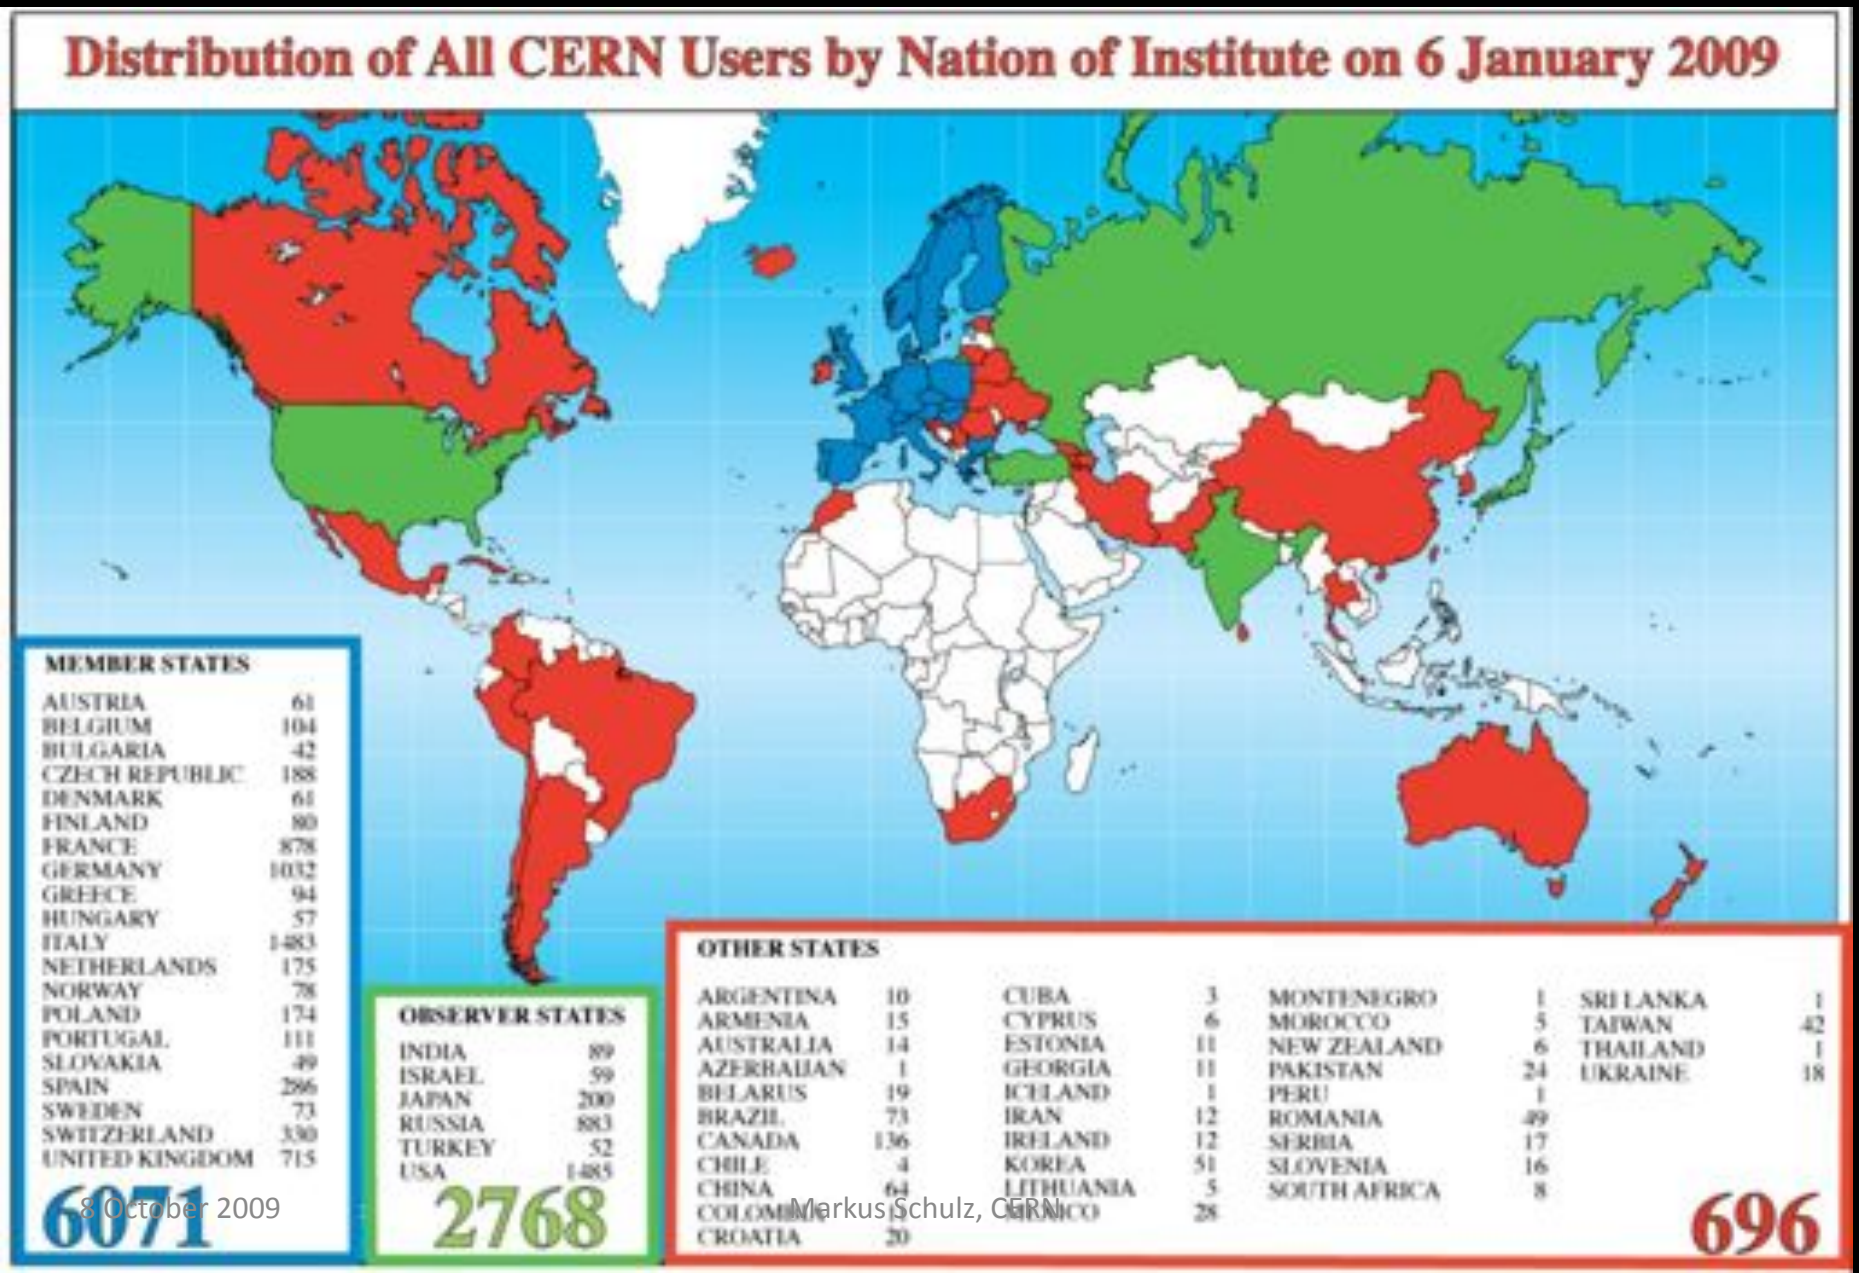
Distribution of All CERN Users by Nation of Institute on 6 January 2009
MEMBER STATES
AUSTRIA 61
BELGIUM 104
BULGARIA 42
CZECH REPUBLIC 188
DENMARK 61
FINLAND 80
FRANCE 878
GERMANY 1032
GREECE 94
HUNGARY 57
ITALY 1483
NETHERLANDS 175
NORWAY 78
POLAND 174
PORTUGAL 111
SLOVAKIA 49
SPAIN 286
SWEDEN 73
SWITZERLAND 330
UNITED KINGDOM 715
6071
OBSERVER STATES
INDIA 89
ISRAEL 59
JAPAN 200
RUSSIA 883
TURKEY 52
USA 1485
2768
OTHER STATES
ARGENTINA 10
ARMENIA 15
AUSTRALIA 14
AZERBAIJAN 1
BELARUS 19
BRAZIL 73
CANADA 136
CHILE 4
CHINA 64
COLOMBIA 5
CROATIA 20
CUBA 3
CYPRUS 6
ESTONIA 11
GEORGIA 11
ICELAND 1
IRAN 12
IRELAND 12
KOREA 51
LITHUANIA 5
MEXICO 28
MONTENEGRO 1
MOROCCO 5
NEW ZEALAND 6
PAKISTAN 24
PERU 1
ROMANIA 49
SERBIA 17
SLOVENIA 16
SOUTH AFRICA 8
SRI LANKA 1
TAIWAN 42
THAILAND 1
UKRAINE 18
696

One of our data generators: ATLAS

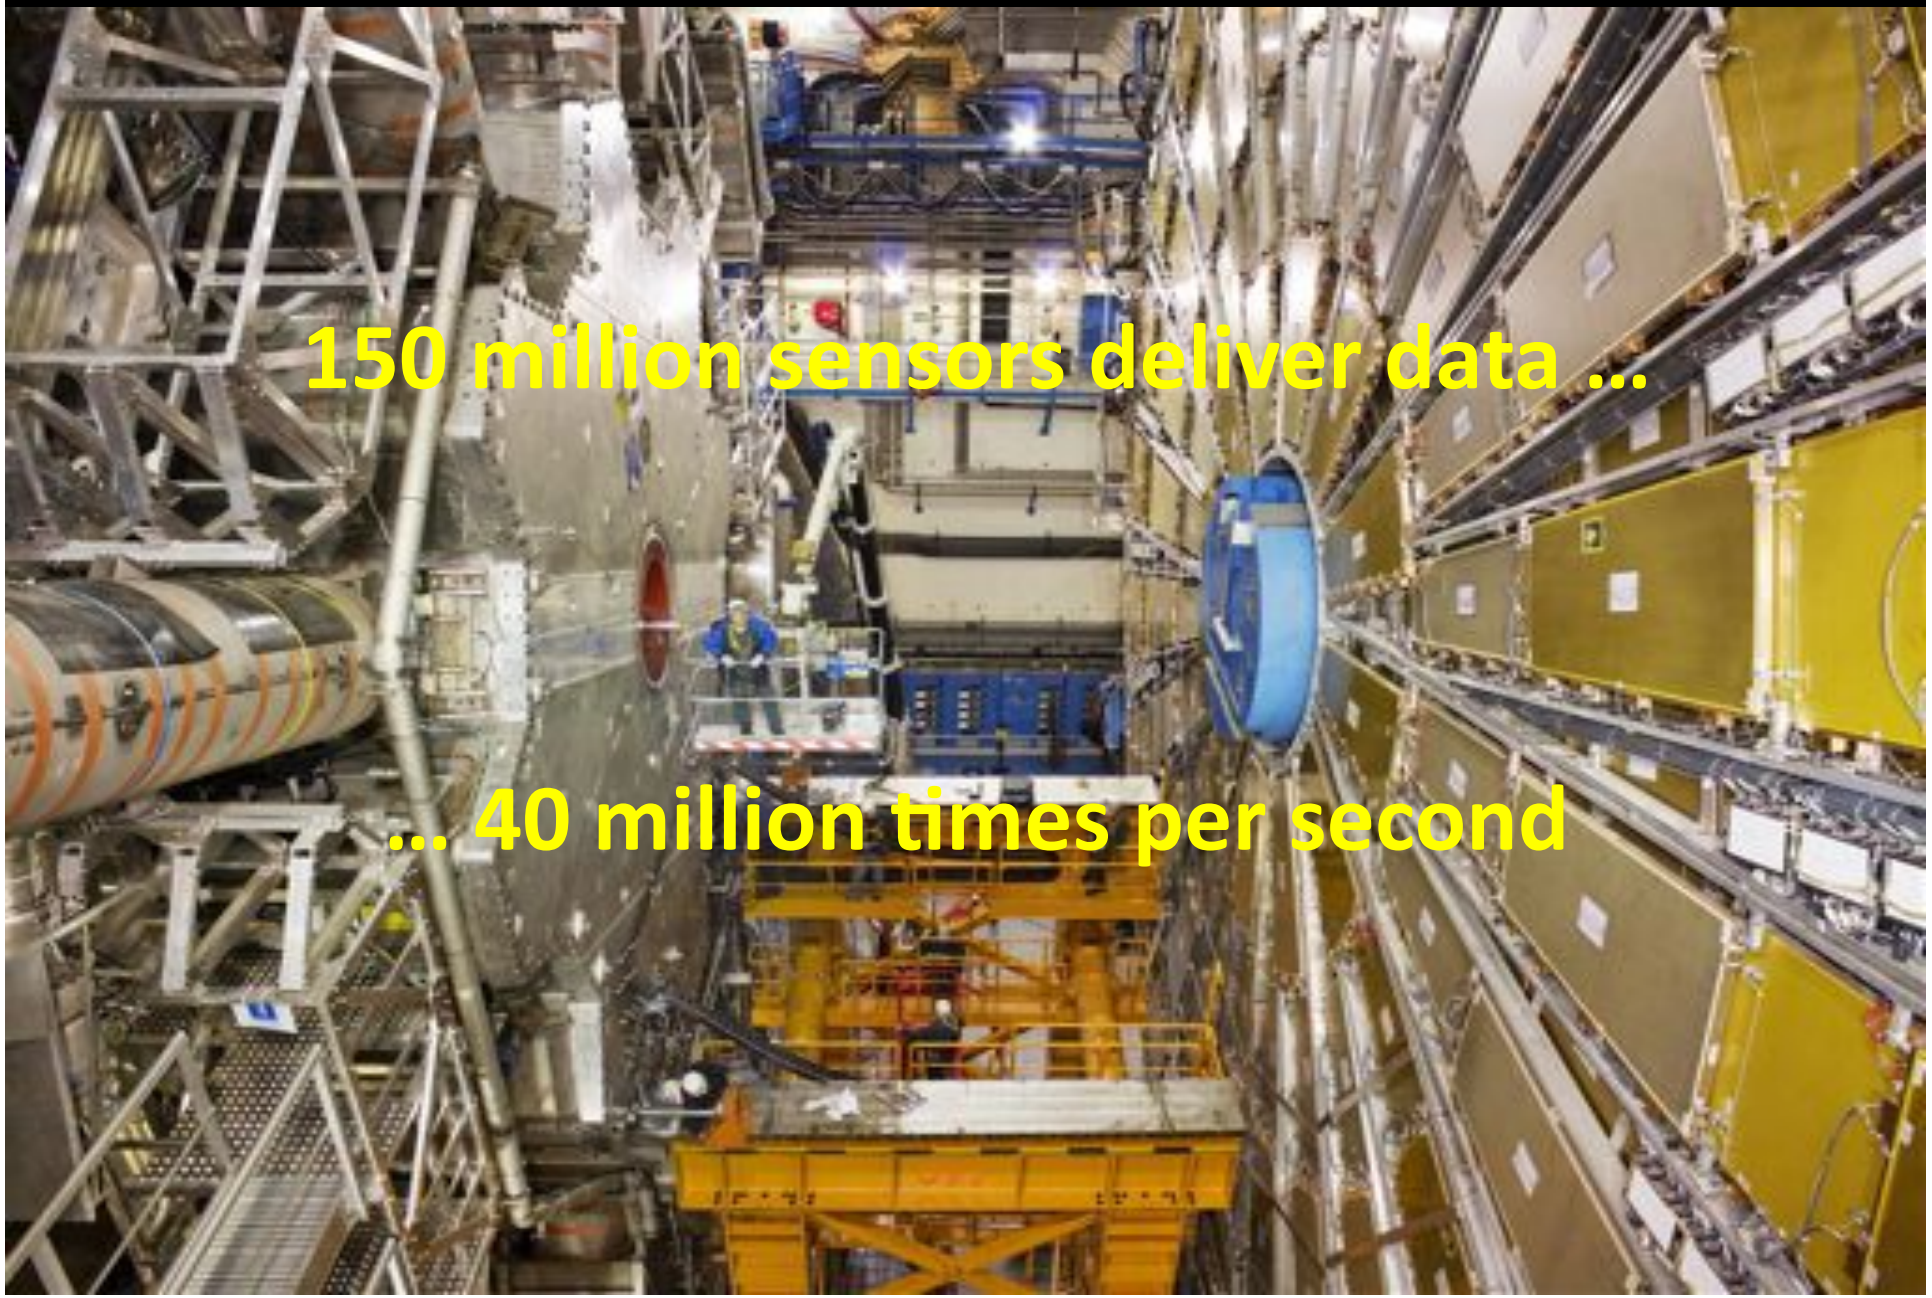

Tier 0 at CERN: Acquisition, First pass processing

Storage & Distribution

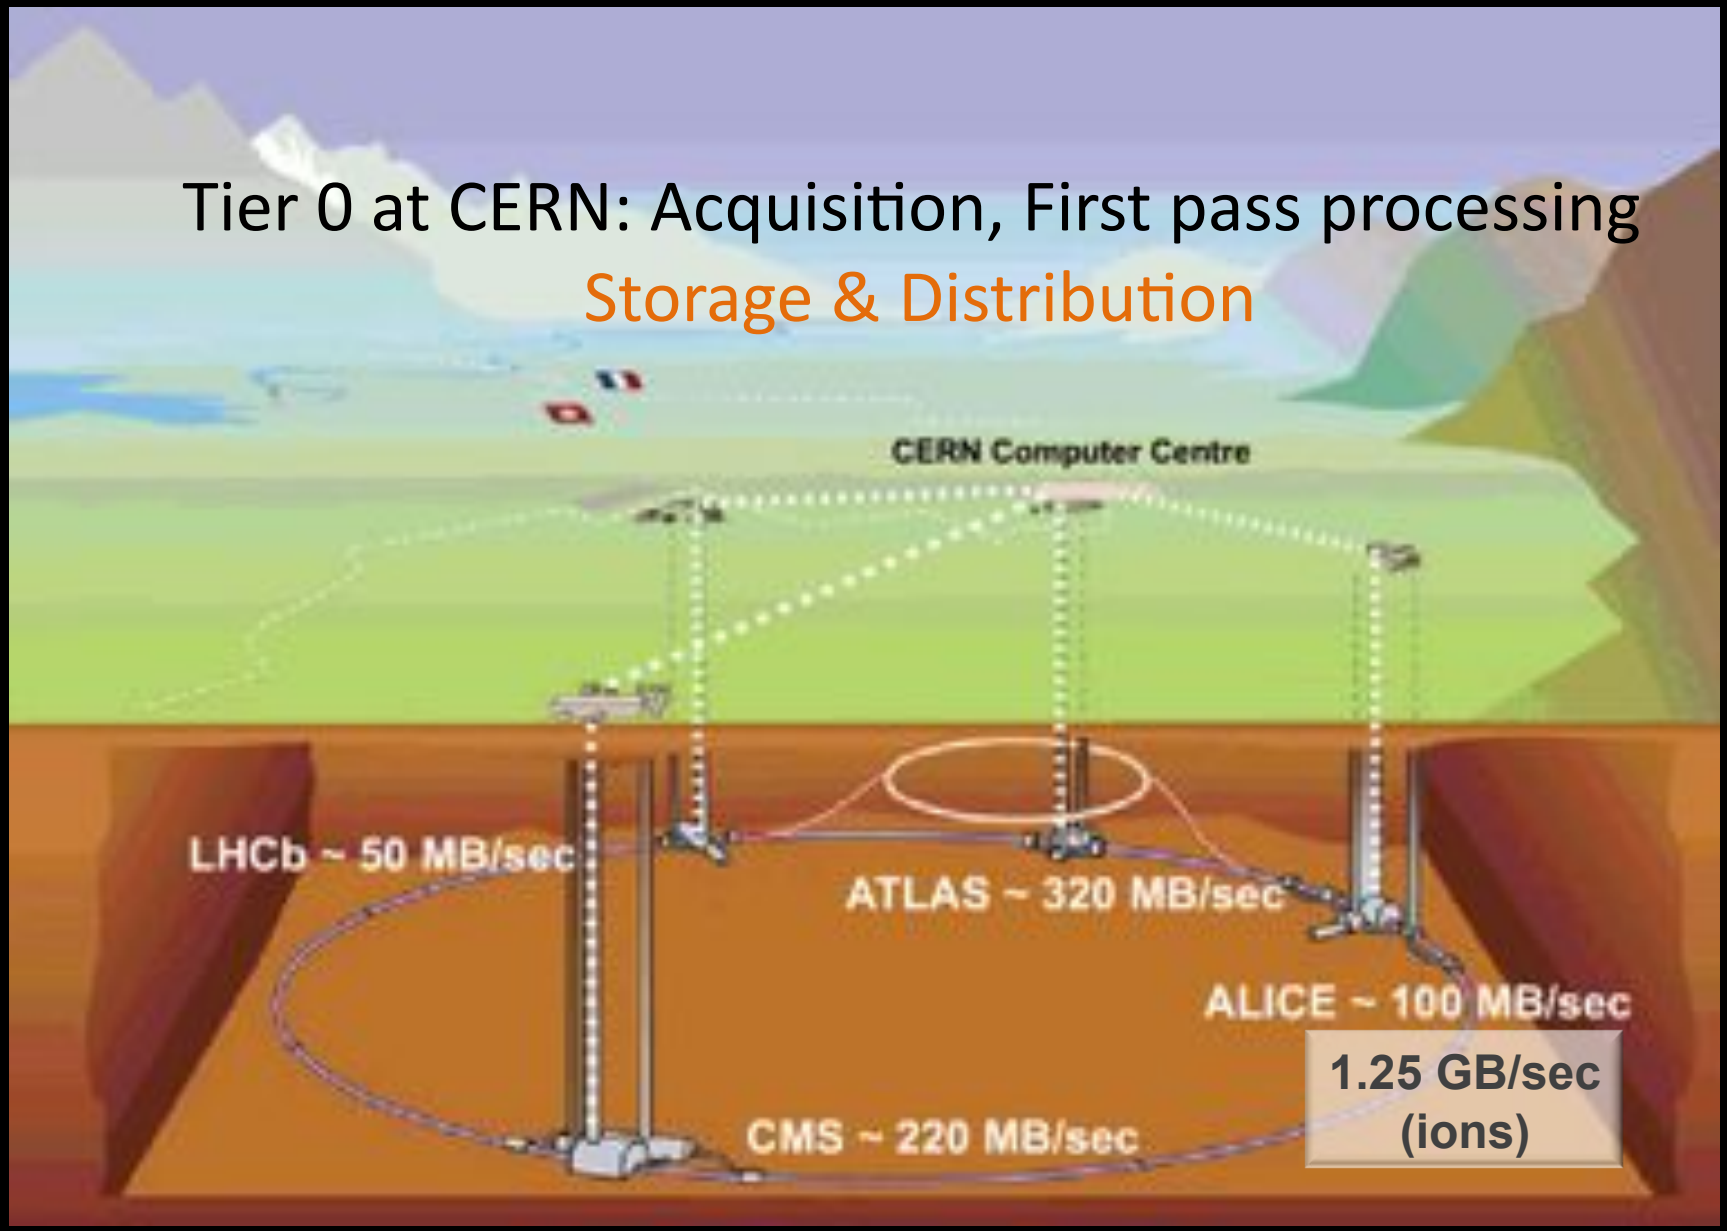

Flow in and out of the center

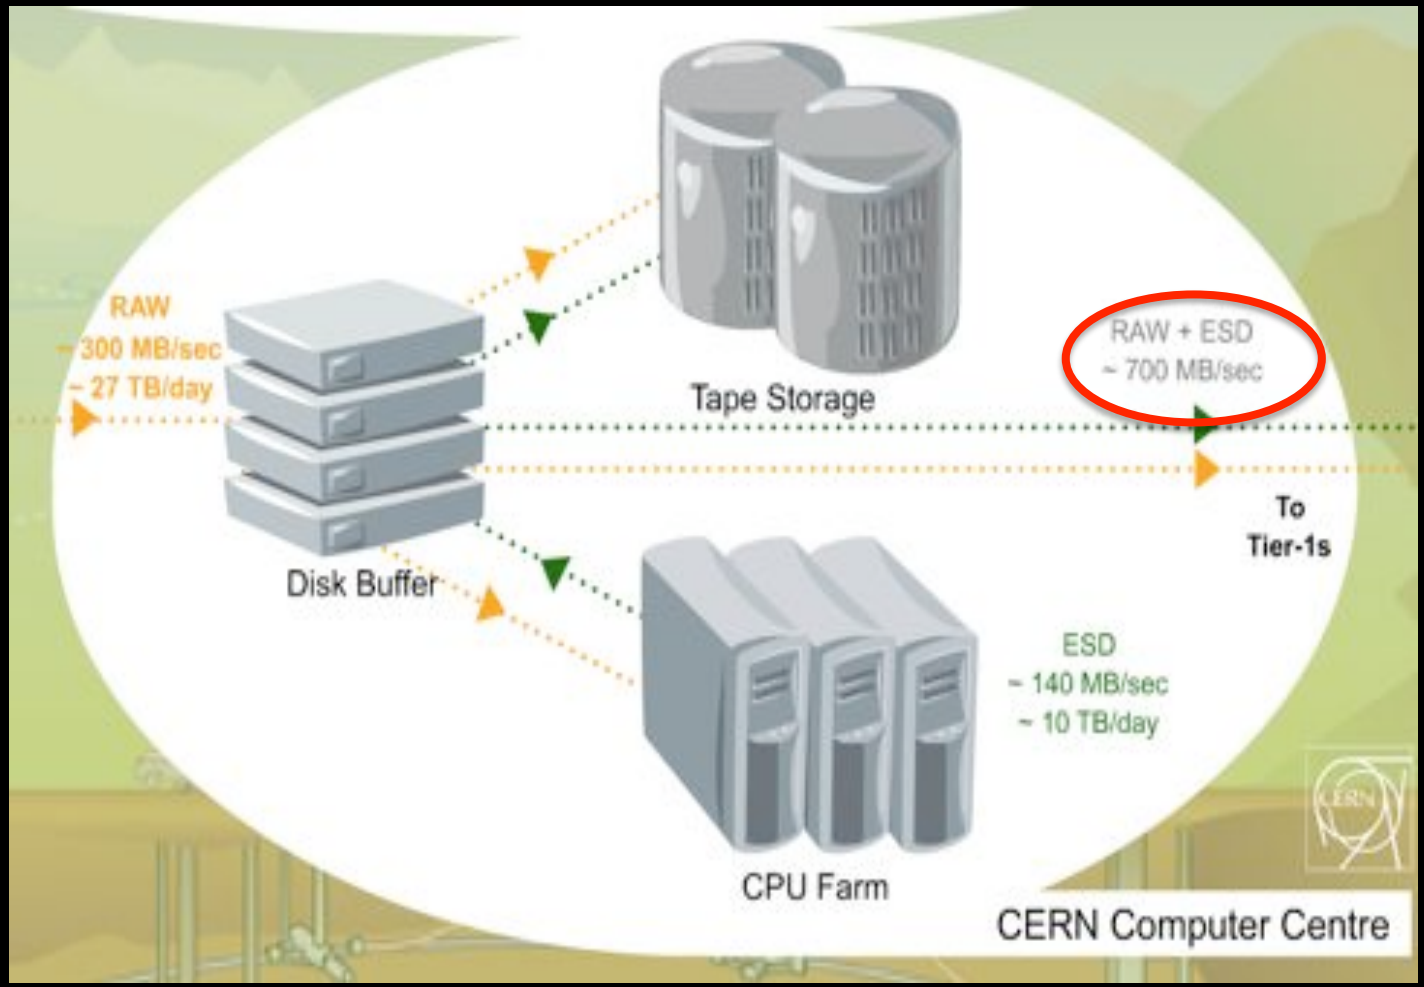

Data and Algorithms

- HEP data are organized as *Events* (particle collisions)
- Simulation, Reconstruction and Analysis programs process "one Event at a time"
  - Events are fairly independent
    → Trivial parallel processing
- Event processing programs are composed of a number of Algorithms selecting and transforming "raw" Event data into "processed" (reconstructed) Event data and statistics

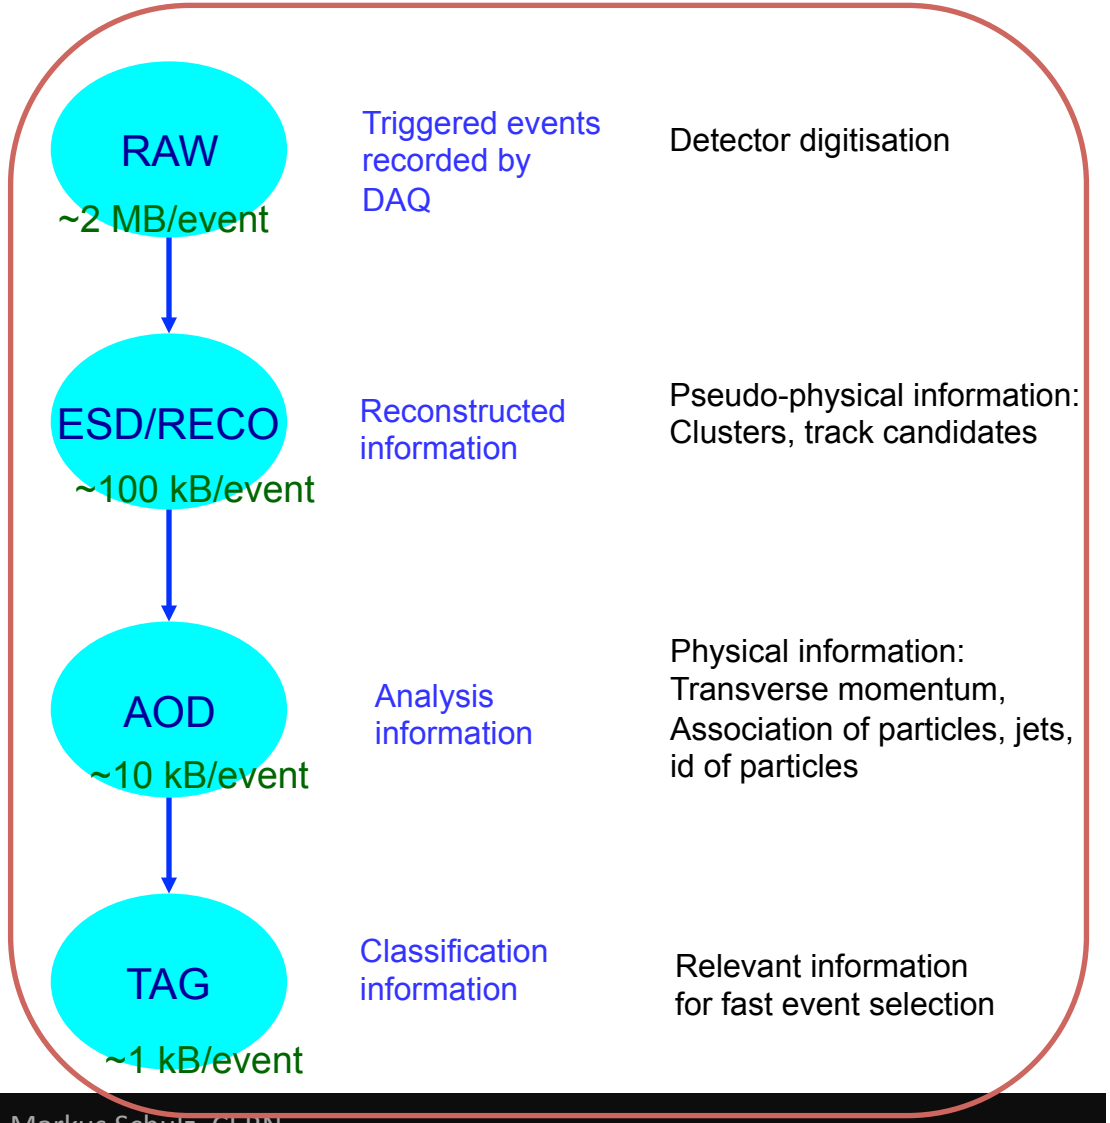

pp collisions at 14 TeV at 10^{34} $cm^{-2}s^{-1}$

A very difficult environment …

How to extract this:
(Higgs → 4 muons)

From this:

With:

20 proton-proton collisions overlap

And this repeats every 25 ns…

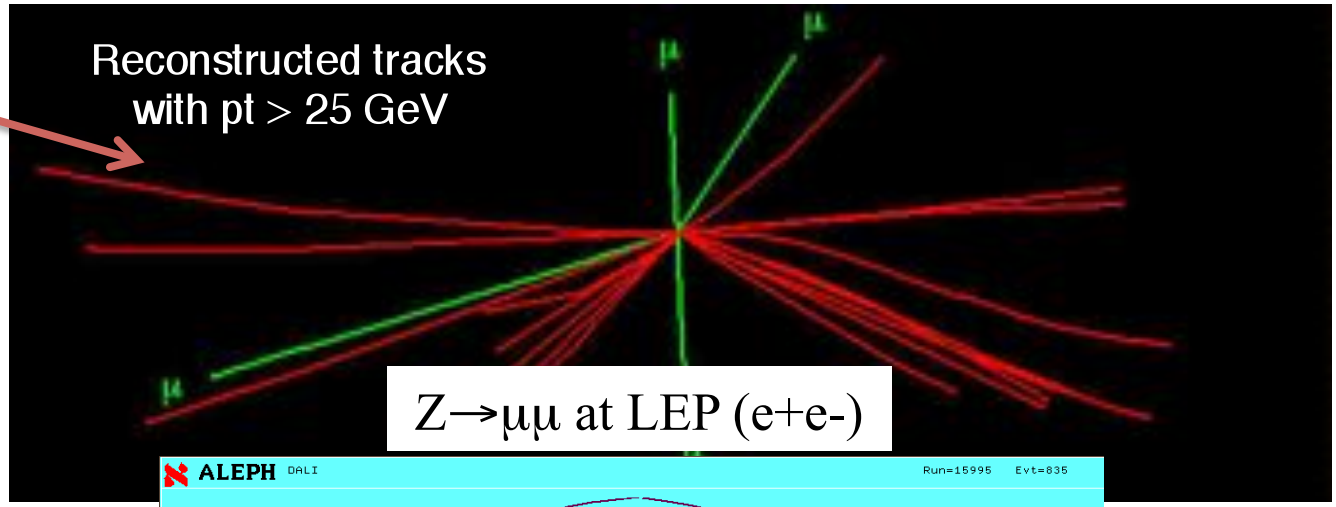

Z→μμ at LEP (e+e-)

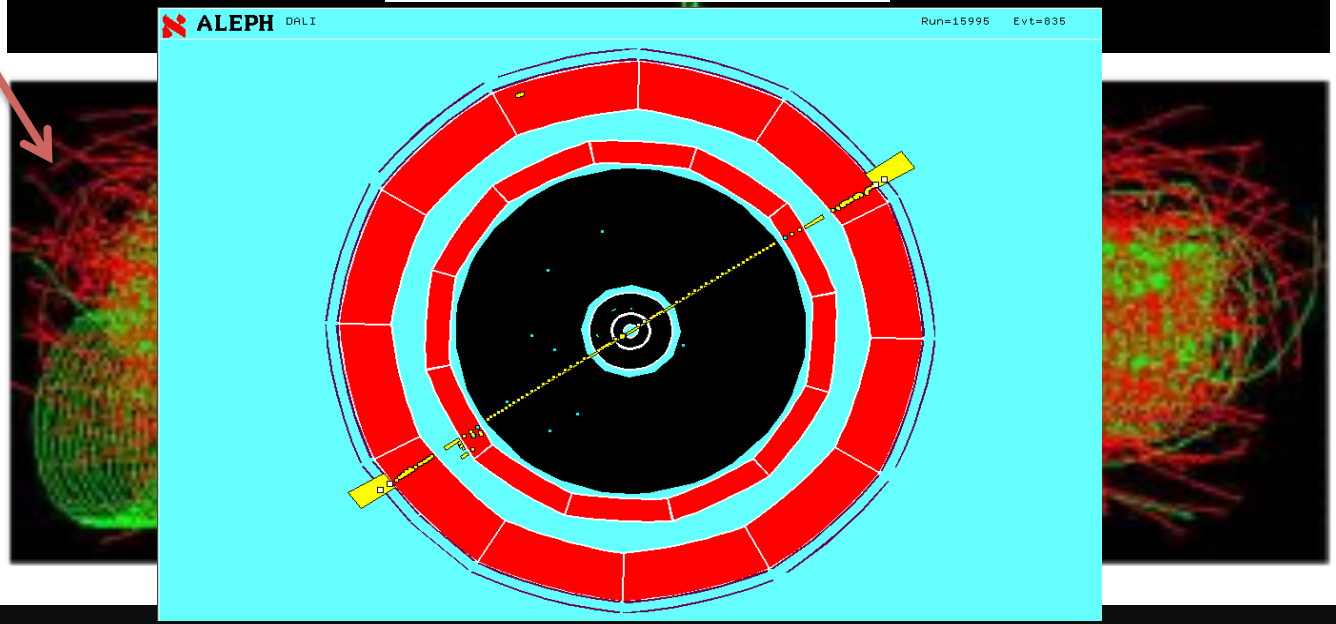

The LHC Computing Challenge

- Signal/Noise: 10^{-13} (10^{-9} offline)
- Data volume
  - High rate * large number of channels * 4 experiments

  ➔ **15 Peta Bytes of new data each year**
- Compute power
  - Event complexity * Nb. events * thousands users

  ➔ **200 k of (today's) fastest CPUs**

  ➔ **45 PB of disk storage**
- Worldwide analysis & funding
  - Computing funding locally in major regions & countries
  - Efficient analysis everywhere

  ➔ **GRID technology**

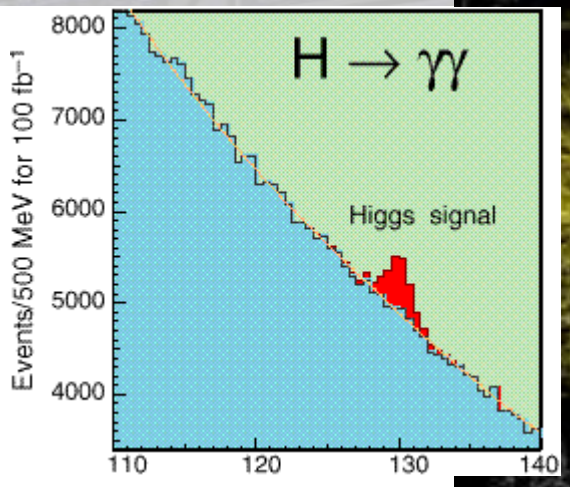

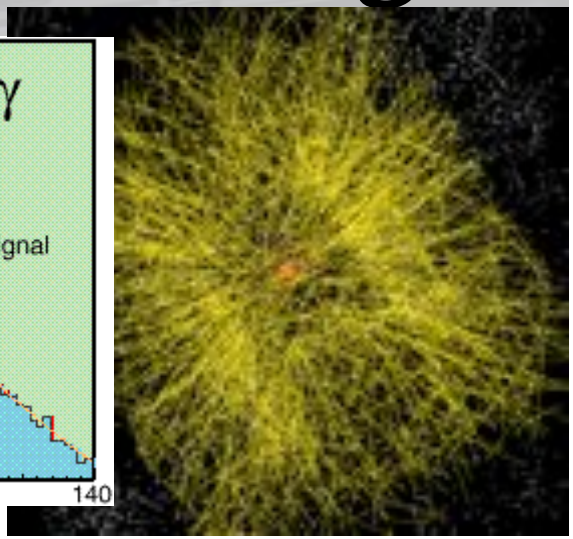

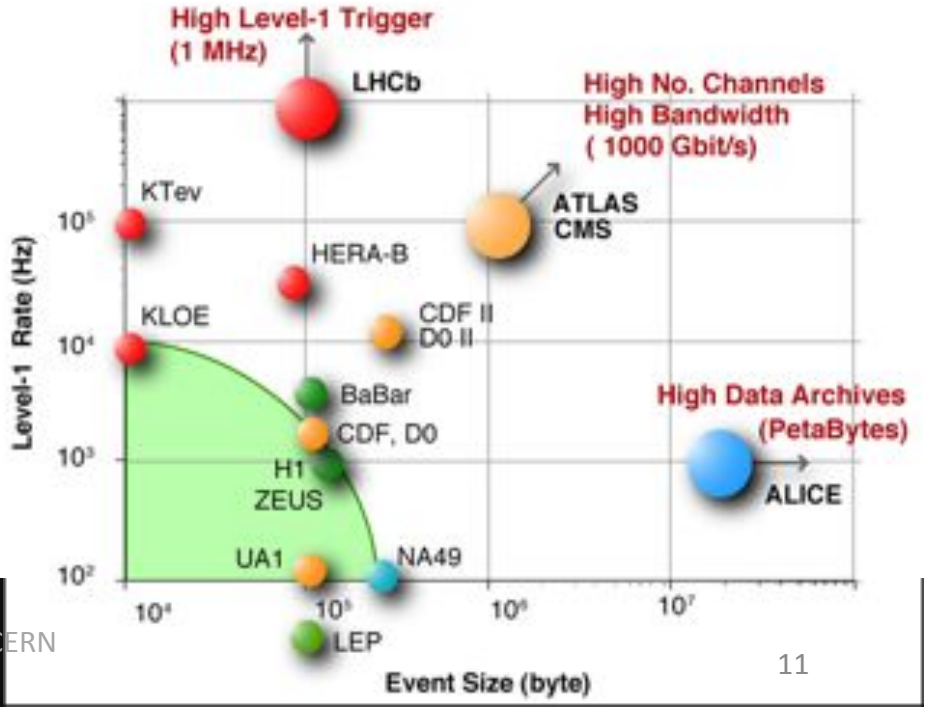

So what is a Grid?

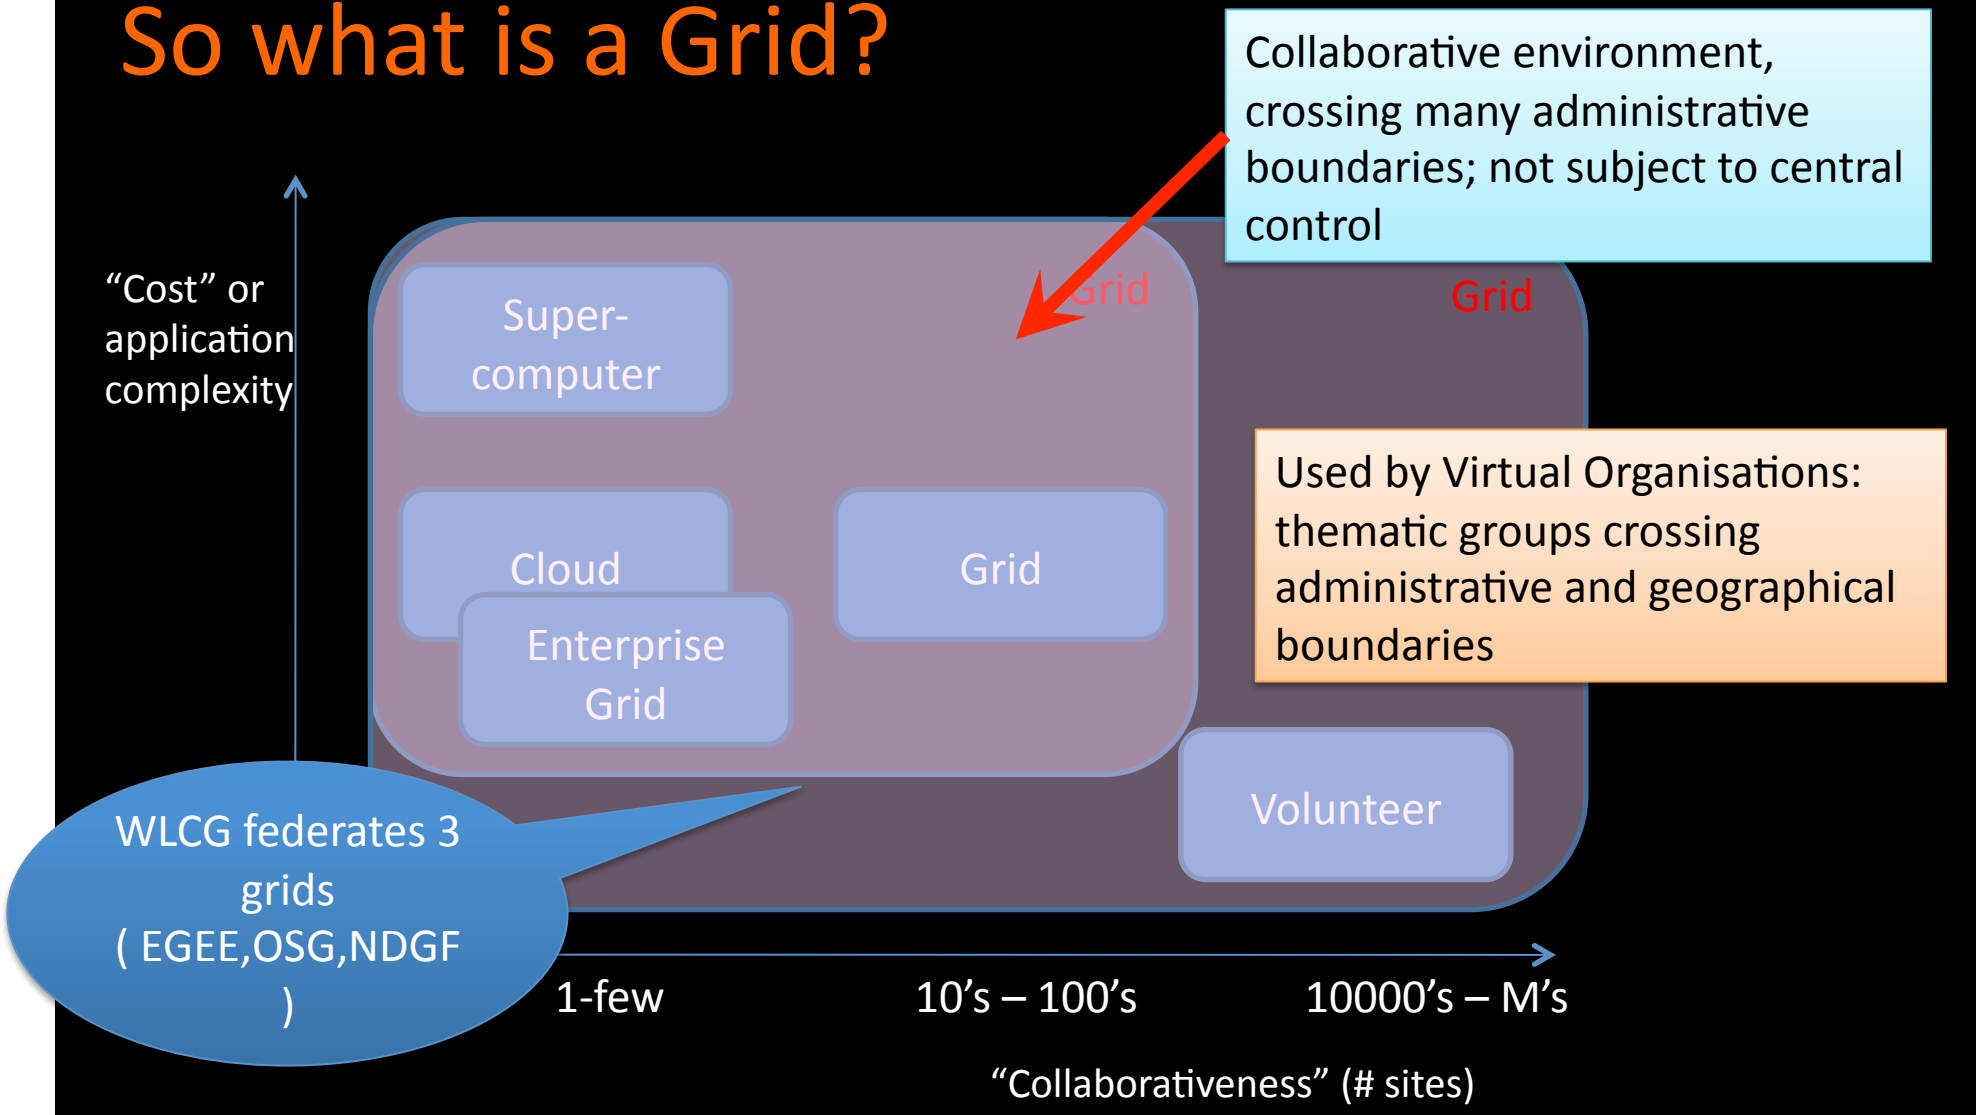

What is Grid Middleware?

- For today:
- The glue that creates the illusion that a distributed infrastructure is a single resource

- If it works, no one will notice it

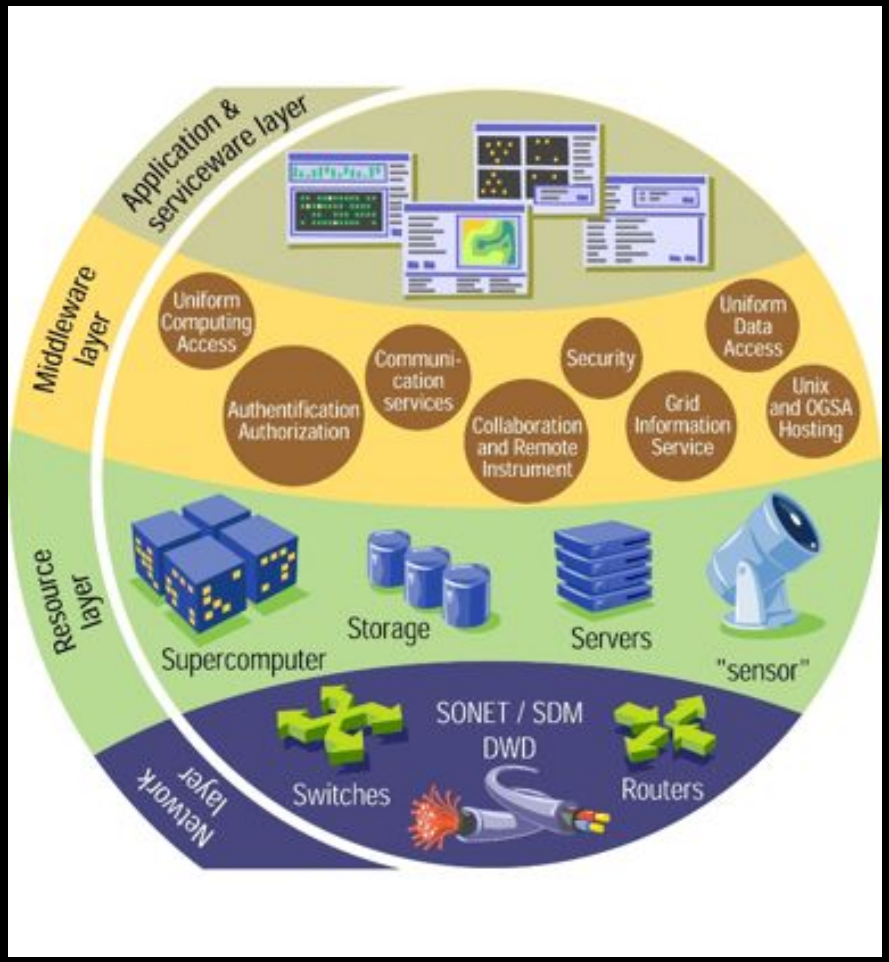

Tier 0 – Tier 1 – Tier 2 the Service Hierarchy

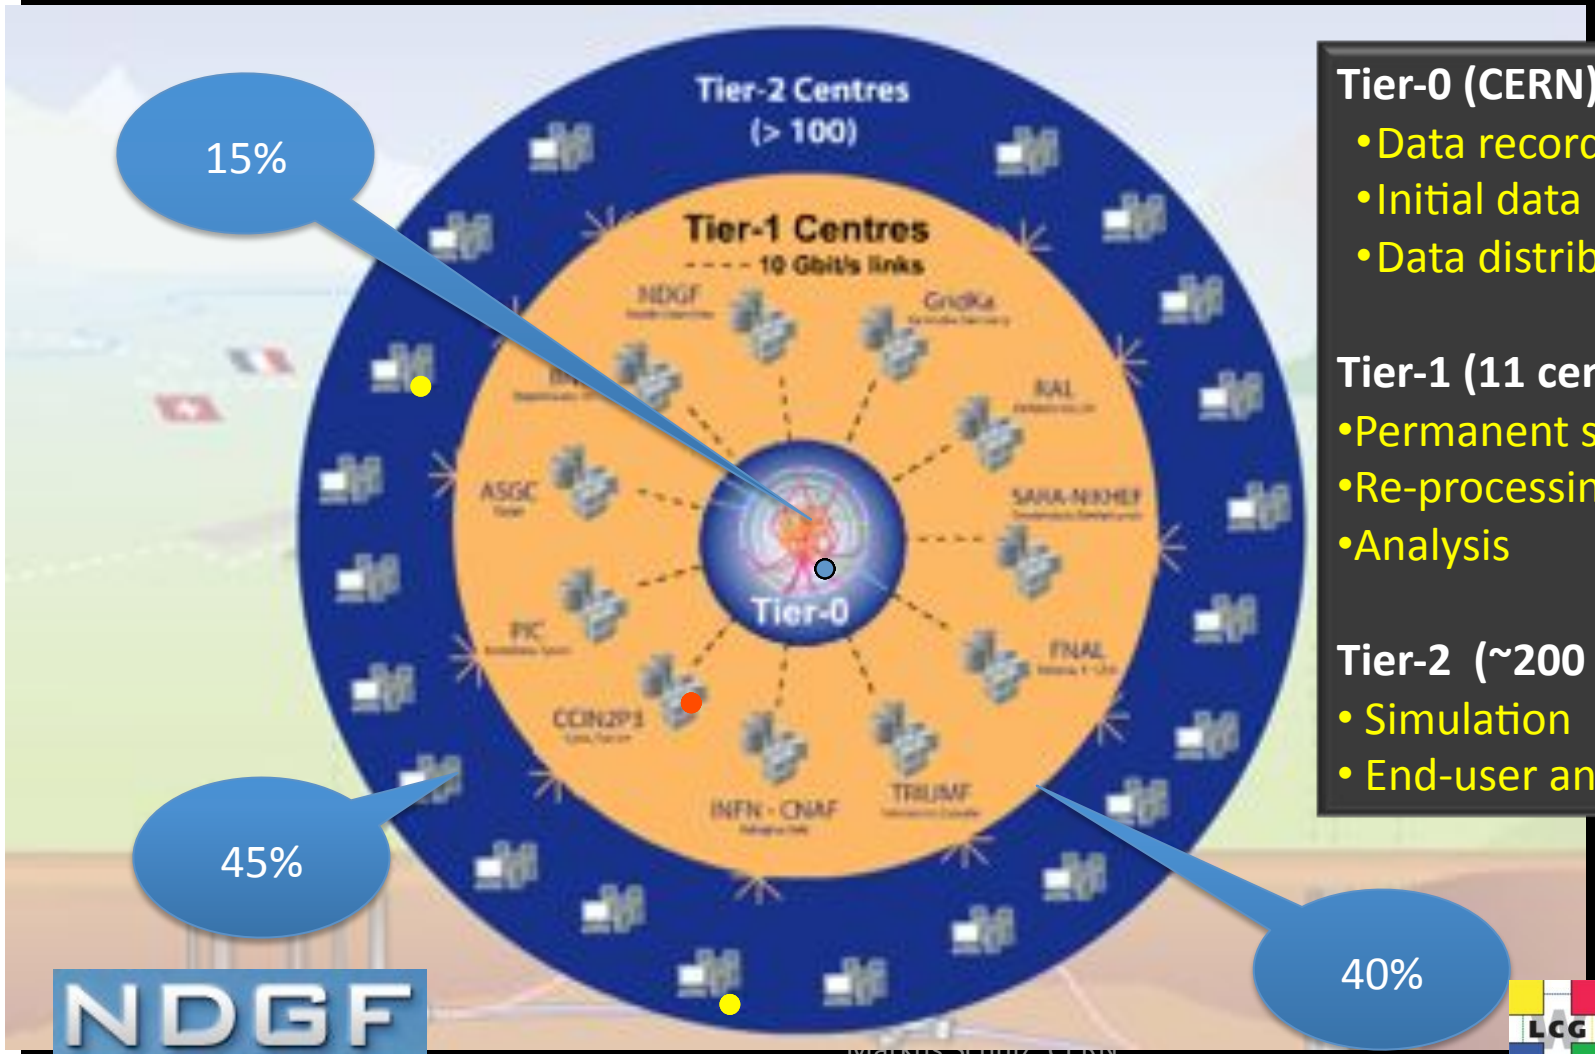

Tier-0 (CERN): (15%)
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (11 centres): (40%)
- Permanent storage
- Re-processing
- Analysis

Tier-2 (~200 centres): (45%)
- Simulation
- End-user analysis

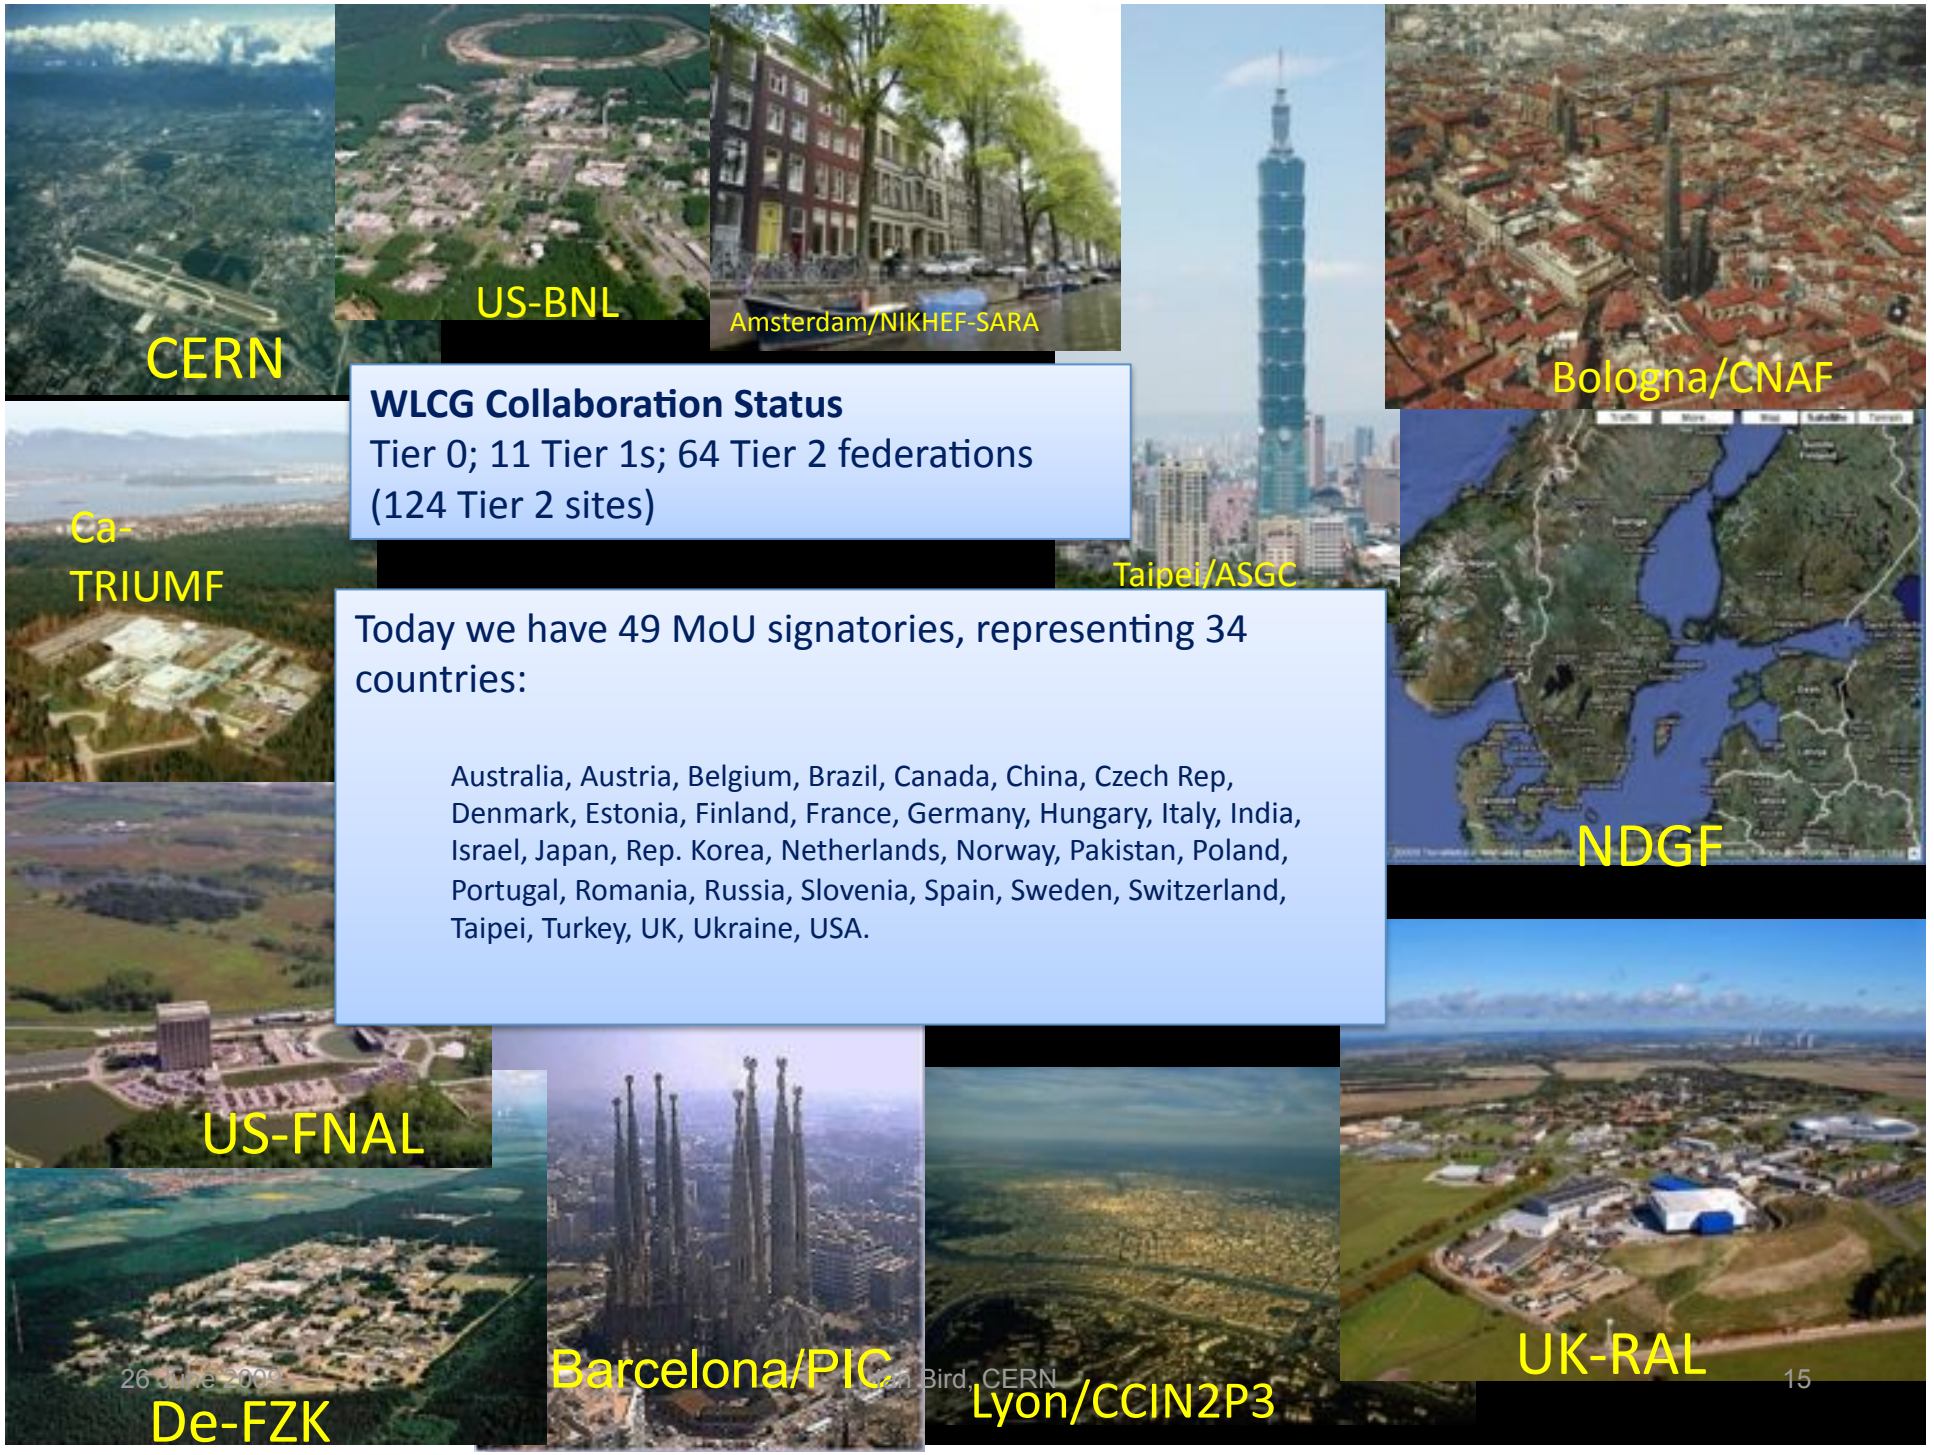

WLCG Collaboration Status
Tier 0; 11 Tier 1s; 64 Tier 2 federations
(124 Tier 2 sites)

Today we have 49 MoU signatories, representing 34 countries:

Australia, Austria, Belgium, Brazil, Canada, China, Czech Rep, Denmark, Estonia, Finland, France, Germany, Hungary, Italy, India, Israel, Japan, Rep. Korea, Netherlands, Norway, Pakistan, Poland, Portugal, Romania, Russia, Slovenia, Spain, Sweden, Switzerland, Taipei, Turkey, UK, Ukraine, USA.

History

- 1999 - MONARC project
  - First LHC computing architecture – hierarchical distributed model
- 2000 – growing interest in grid technology
  - HEP community main driver in launching the DataGrid project
- 2001-2004 - EU DataGrid project
  - middleware & testbed for an operational grid
- 2002-2005 – LHC Computing Grid – LCG
  - deploying the results of DataGrid to provide a production facility for LHC experiments
- 2004-2006 – EU EGEE project phase 1
  - starts from the LCG grid
  - shared production infrastructure
  - expanding to other communities and sciences
- 2006-2008 – EU EGEE project phase 2
  - expanding to other communities and sciences
  - Scale and stability
  - Interoperations/Interoperability
- 2008-2010 – EU EGEE project phase 3
  - More communities
  - Efficient operations
  - Less central coordination
- 2010 – 201x EGI and EMI
  - Sustainable infrastructures based on National Grid Infrastructures
  - Decoupling of middleware development and infrastructure

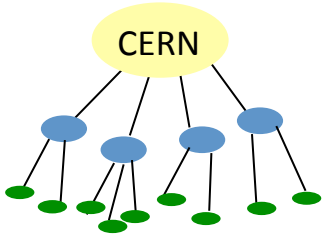

Evolution of (W)LCG

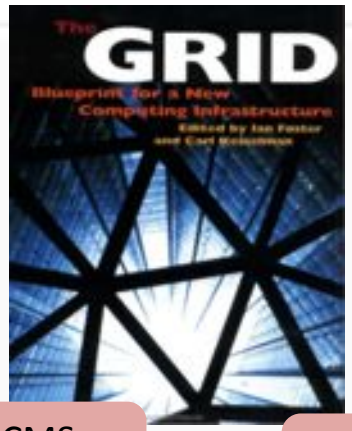

LHC approved

GriPhyN, iVDGL, PPDG

GRID 3

OSG

WLCG

EU DataGrid

EGEE 1

EGEE 2

EGEE 3

LCG 1

LCG 2

ATLAS&CMS CTP

10^7 MIPS

100 TB disk

"Hoffmann" Review

$7x10^7$ MIPS

1,900 TB disk

Computing TDRs

$55x10^7$ MIPS (140 MSi2k)

70,000 TB disk

1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008

Service Challenges

Cosmics

First physics

Data Challenges

ATLAS & CMS approved

ALICE approved

LHCb approved

Ever increasing requirements

Production Grids

- WLCG relies on a *production quality* infrastructure
  - Requires standards of:
    - Availability/reliability
    - Performance
    - Manageability
  - Is used 365 days a year ... (has been for several years!)
  - Tier 1s must store the data for at least the lifetime of the LHC - ~20 years
    - Not passive – requires active migration to newer media
- Vital that we build a fault-tolerant and reliable system
  - That can deal with individual sites being down and recover

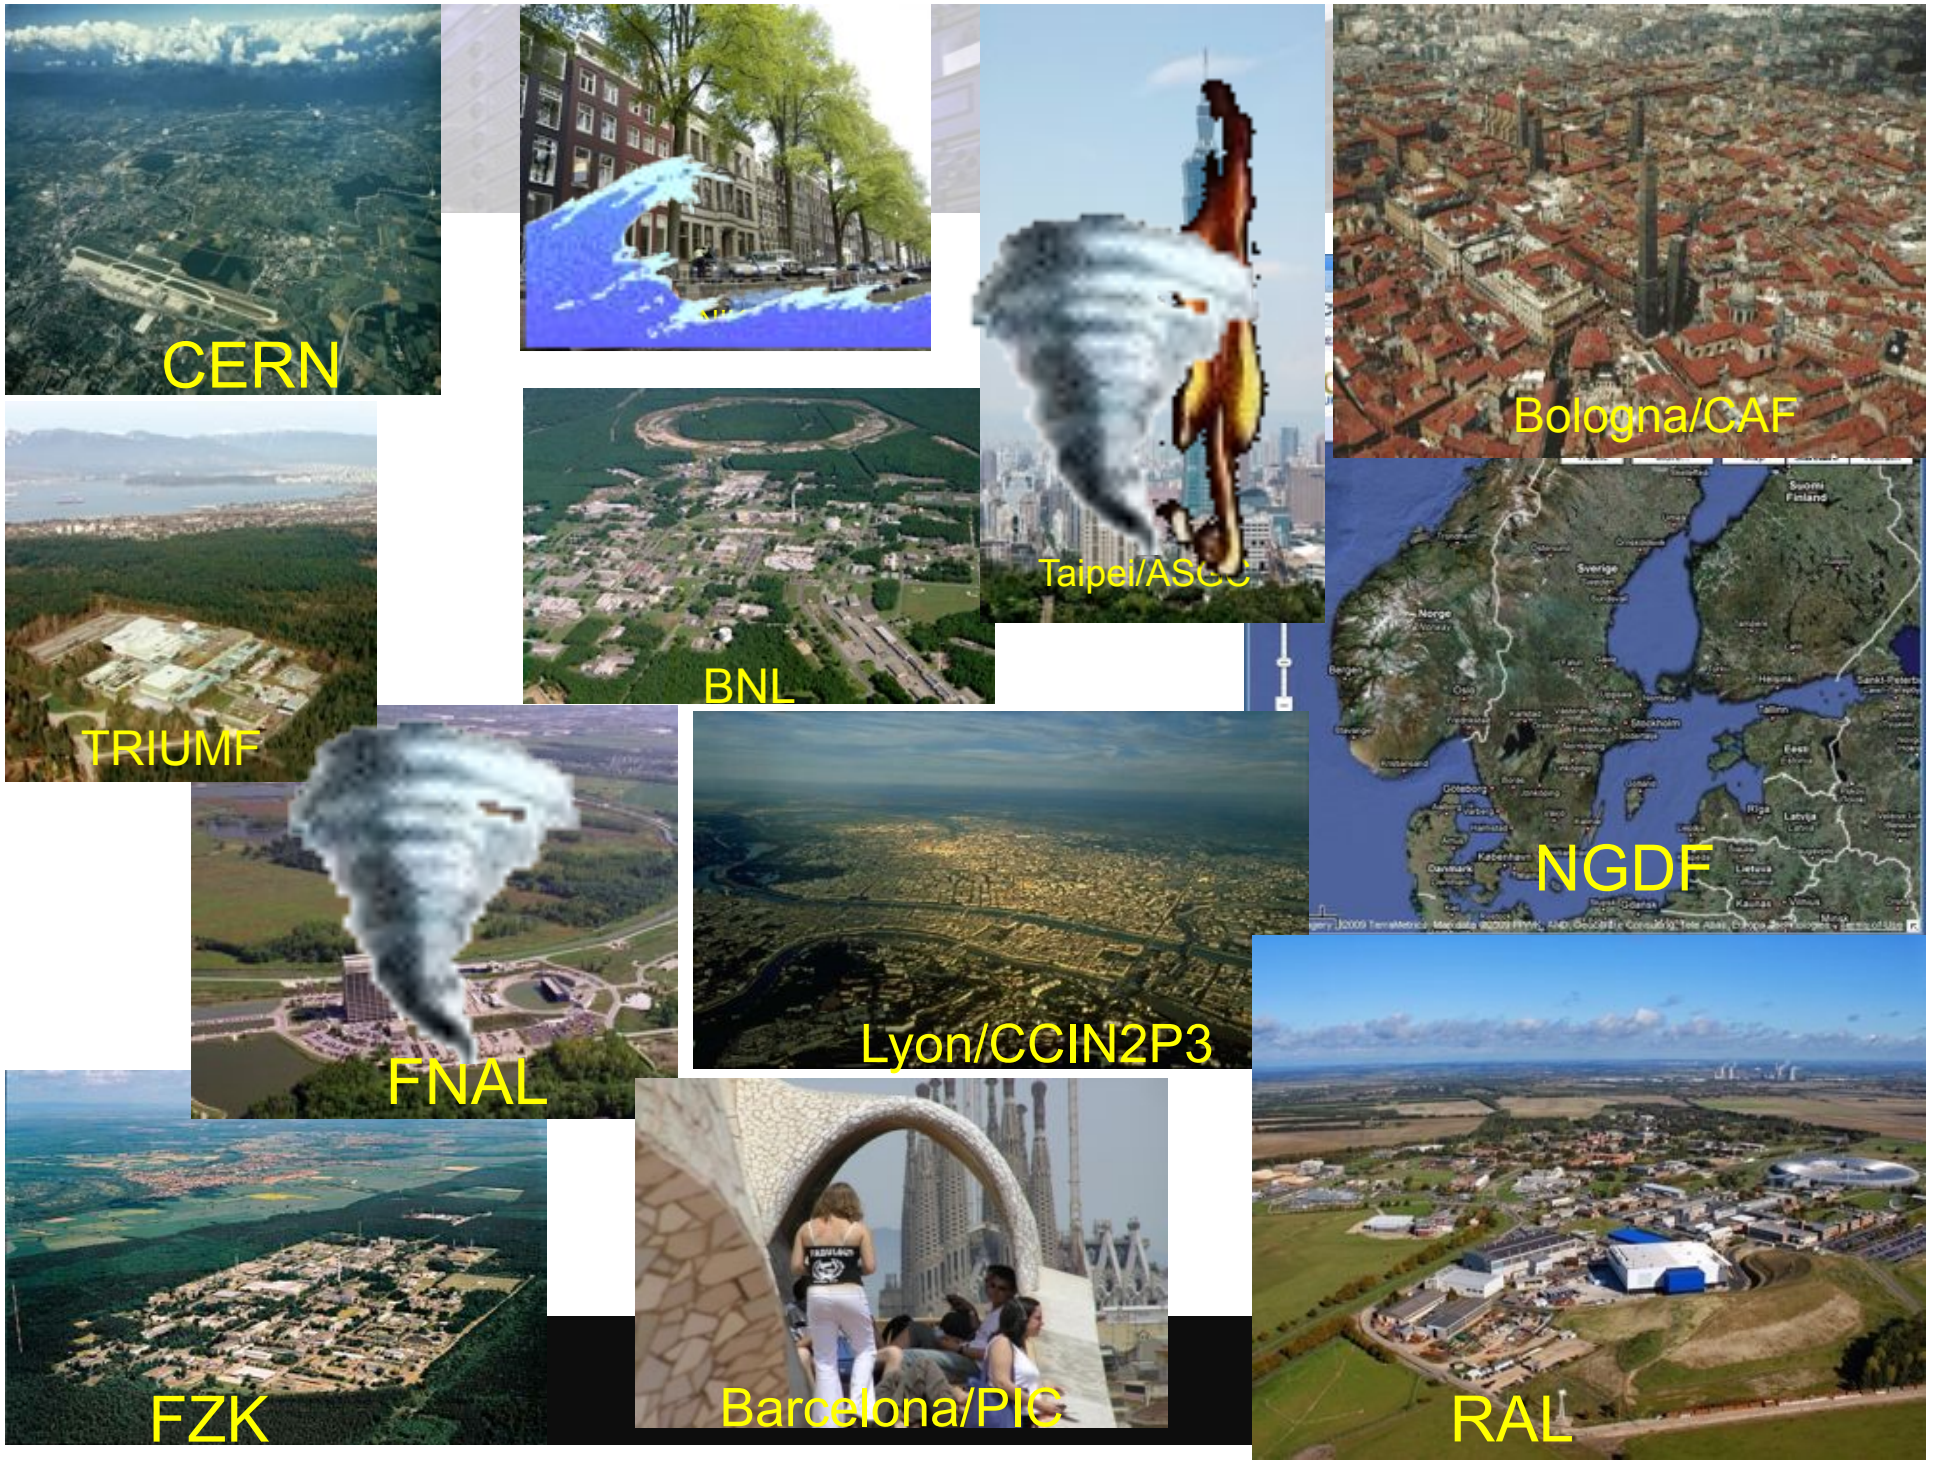
CERN
Bologna/CAF
TRIUMF
BNL
Taipei/ASGC
NGDF
FNAL
Lyon/CCIN2P3
FZK
Barcelona/PIC
RAL

What is needed to make it work?

- Apart from Middleware
- Apart from Computer Centers

WLCG Worldwide LHC Computing Grid

Everything you need in a Computer Center!

- Management
- Fabric
- Networking
- Security
- Monitoring
- User Support
- Problem Tracking
- Accounting
- Service support
- SLAs…..

- But now on a global scale
  - Respecting the sites' independence
  - Linking the different infrastructures
    - NDGF, EGEE (EGI), OSG

WLCG Worldwide LHC Computing Grid

Worldwide LCG Organisation

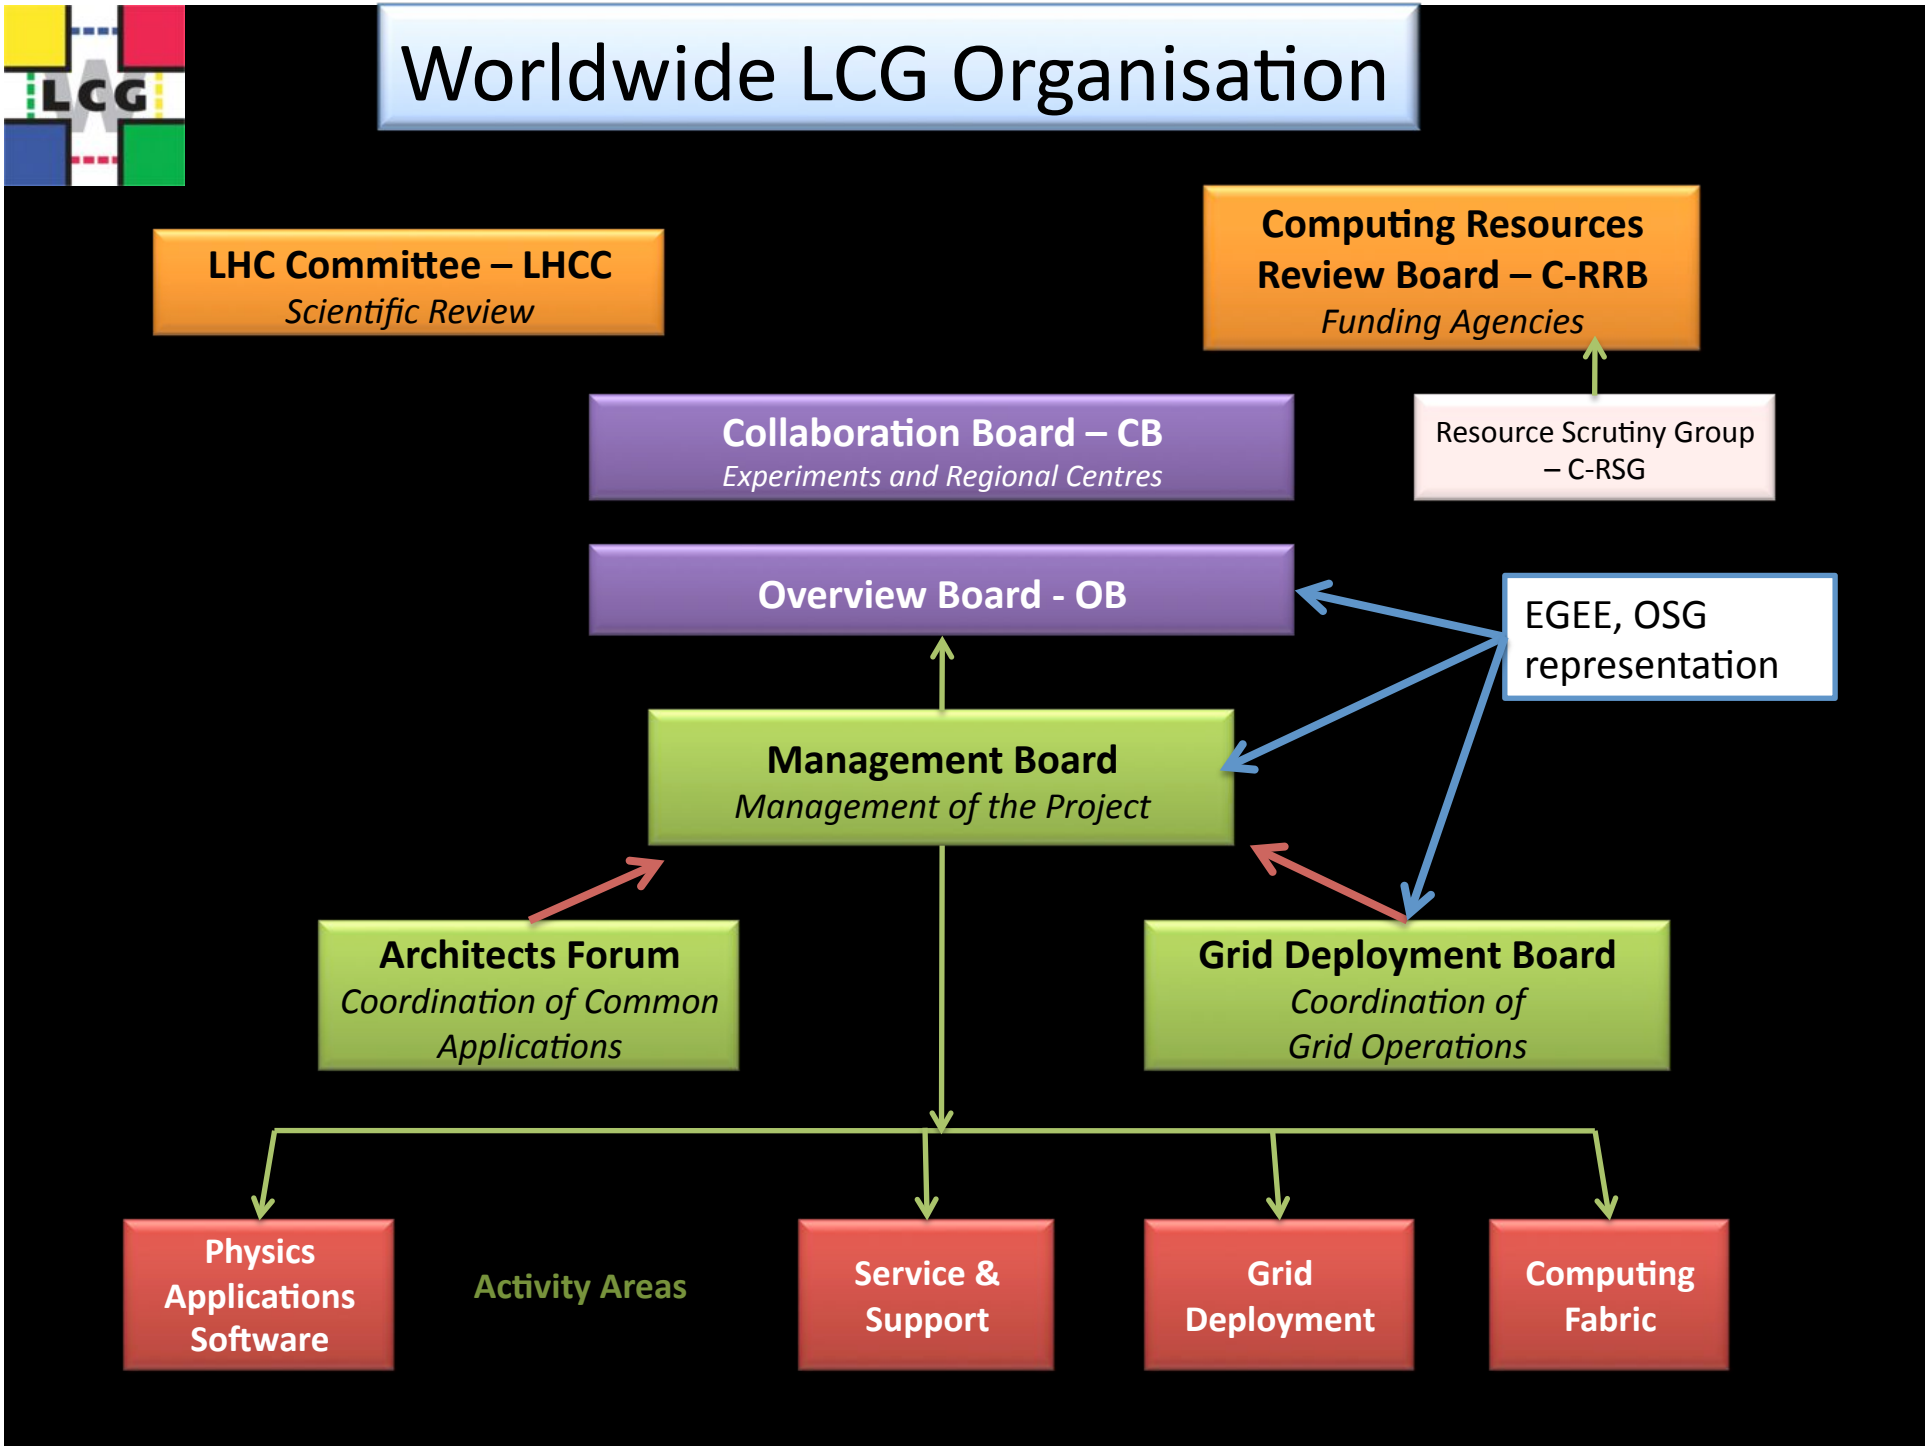

The EGEE Infrastructure

Support Structures & Processes

- Operations Coordination Centre
- Regional Operations Centres
- Global Grid User Support
- EGEE Network Operations Centre (SA2)
- Operational Security Coordination Team
- Training activities (NA3)

Test-beds & Services

- Production Service
- Pre-production service
- Certification test-beds (SA3)
- Training infrastructure (NA4)

Security & Policy Groups

- Joint Security Policy Group
- EuGridPMA (& IGTF)
- Grid Security Vulnerability Group
- Operations Advisory Group (+NA4)

Operations Infrastructure

- Regional Operations Centers (ROCs)
  - Little central coordination, policy driven
  - CIC portal as a common information point
- Connect to users and admins
  - Problem tracking tools
    - GGUS (Global Grid User Support)
      - Web interface, mail interface
      - Handles ~1000 tickets every month
      - Interfaced to local tools -→ acceptance
- Accounting: APEL
  - Central DB + Portal (UK & Spain)
  - Interfaces for other systems: DGAS, OSG, NDGF

In addition to EGEE/EGI Operations

- Daily WLCG Operations Meetings
  - 30 minutes
  - Follow up on current problems
- Every two weeks WLCG T1 Service Coordination meeting
  - Operational Planning
  - Incidents followup
- Detailed monitoring of the SLAs.

Global Grid User Support

- GGUS: Web based portal
  - about 1000 tickets per months
  - Grid security aware
  - Interfaces to regional/national support structures

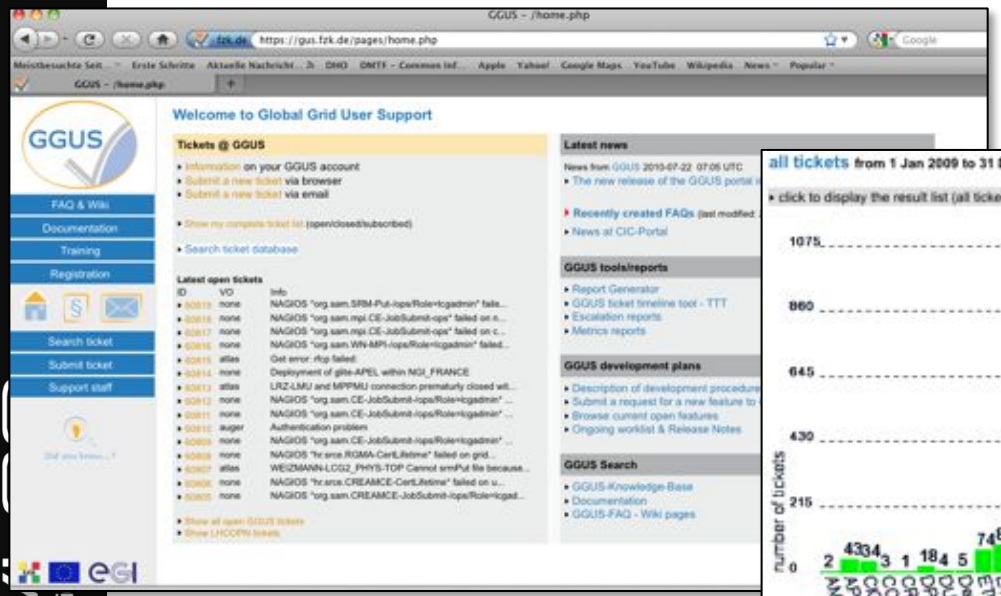

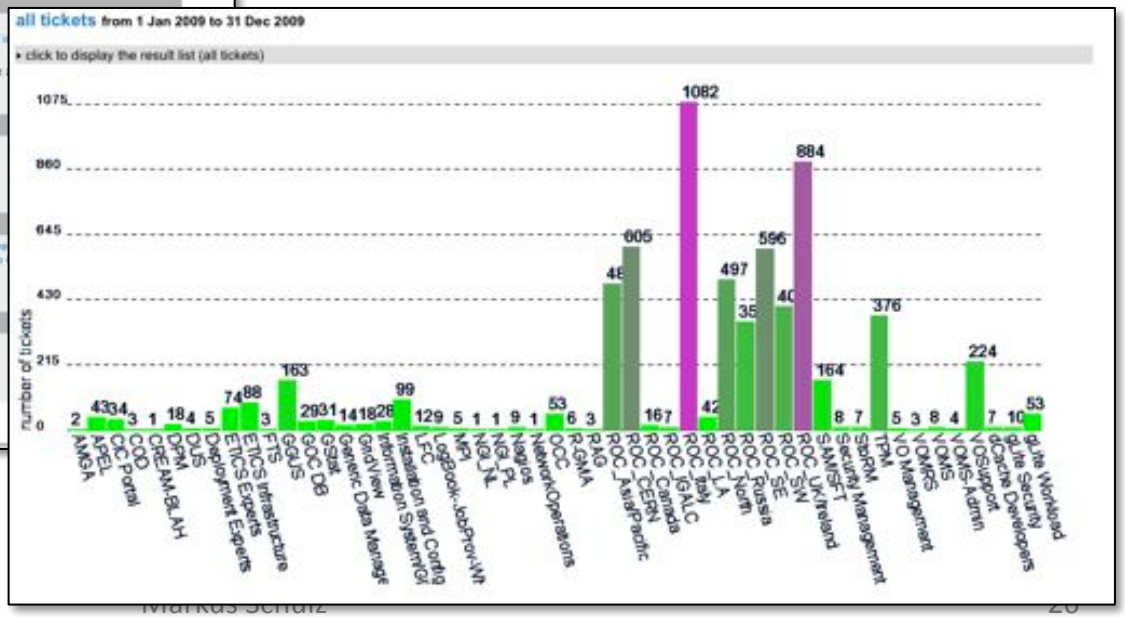

Reliabilities

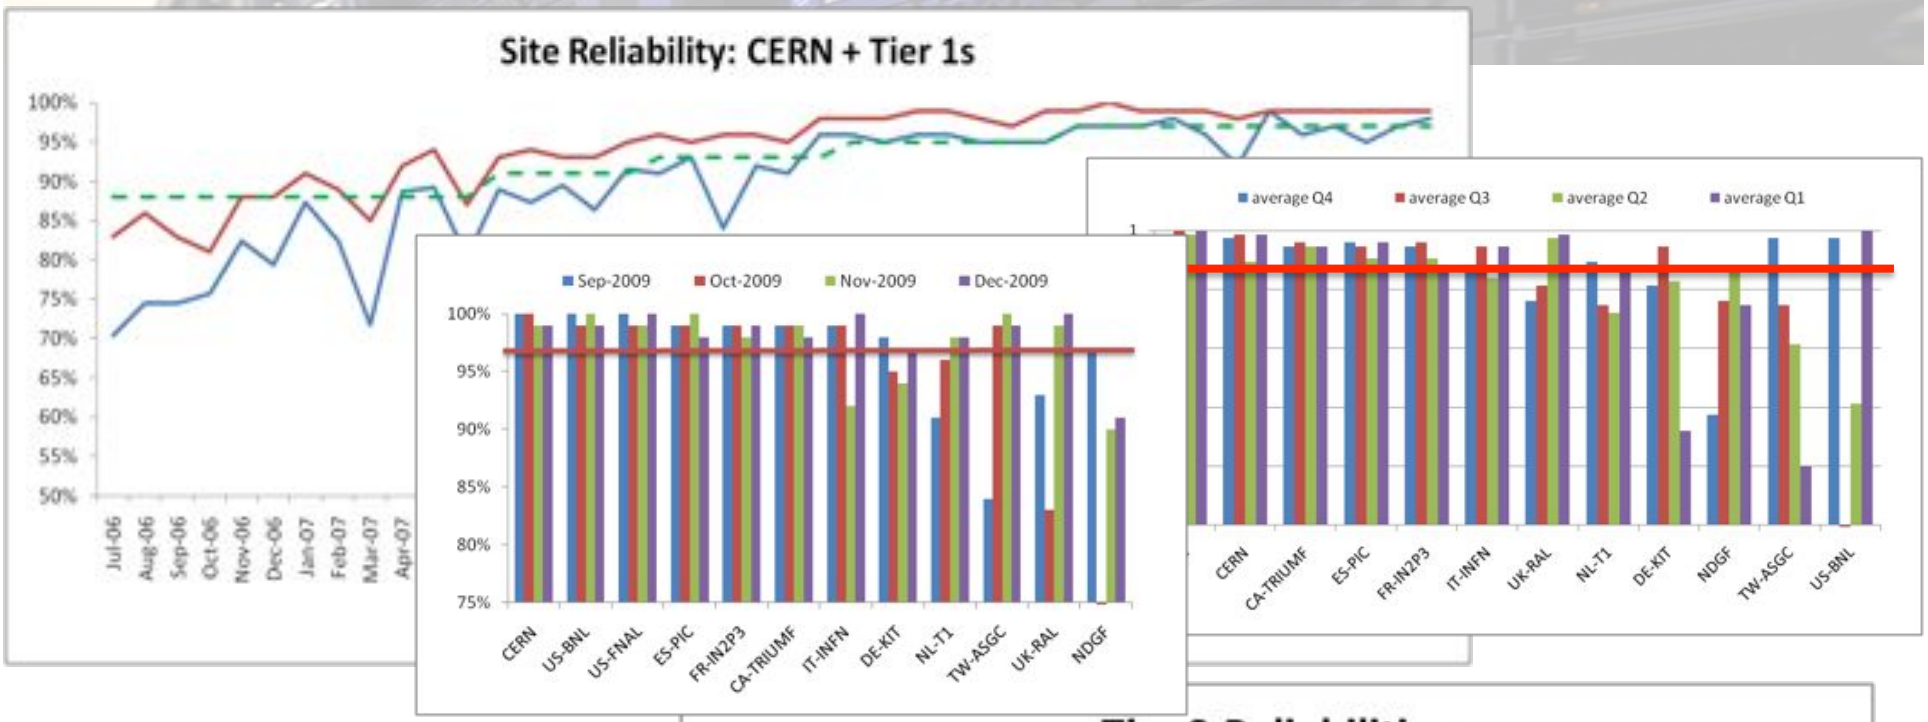

- This is not the full picture:
- Experiment-specific measures give complementary view
- Need to be used together with some understanding of underlying issues

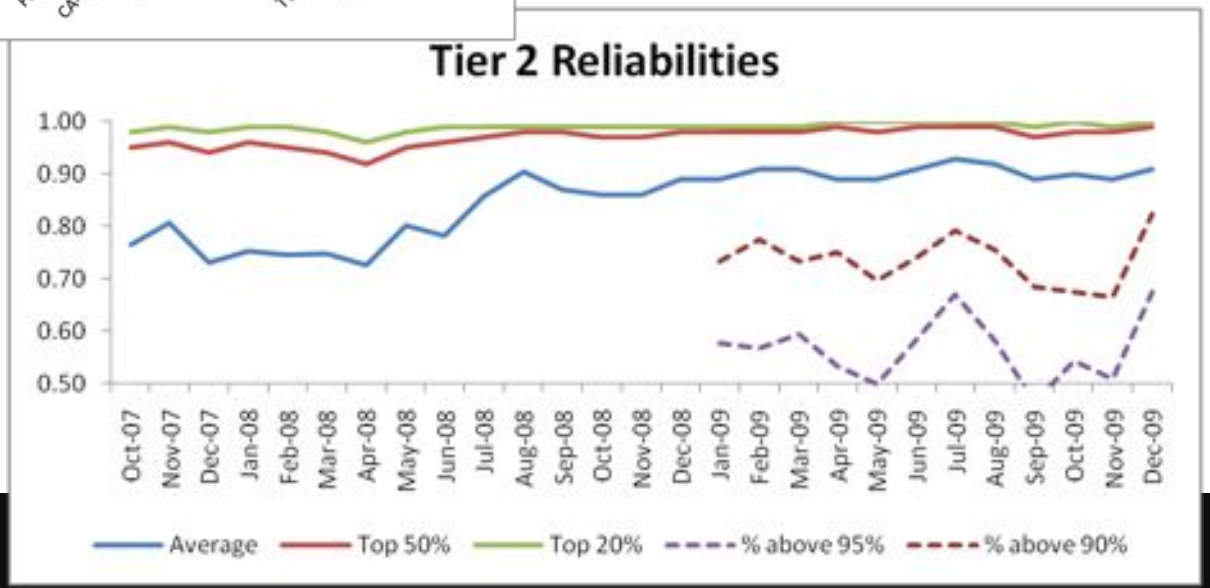

Availability metrics - GridView

Overall Service Availablity for site CERNPR : Daily Report

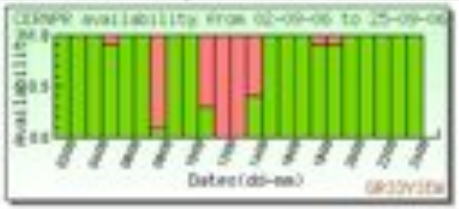

Individual Service Availablity for site CERNPR : Daily Report

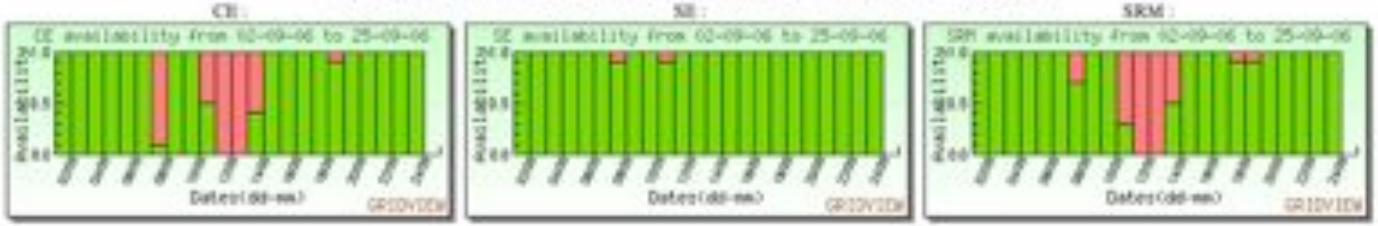

Service Instance Availablity for site CERNPR : Daily Report

Site Services

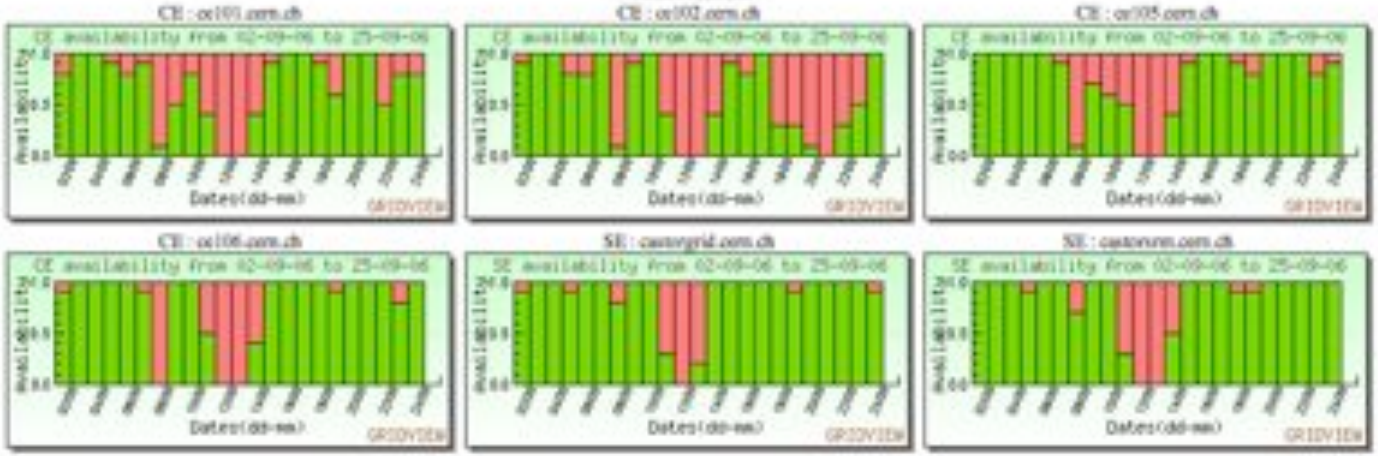

Monitoring to Improve Reliability

- Monitoring
- Metrics
- Workshops
- Data challenges
- Experience
- Systematic problem analysis
- Priority from software developers

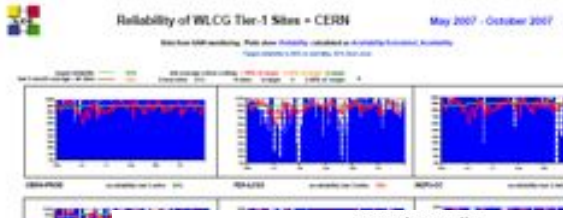

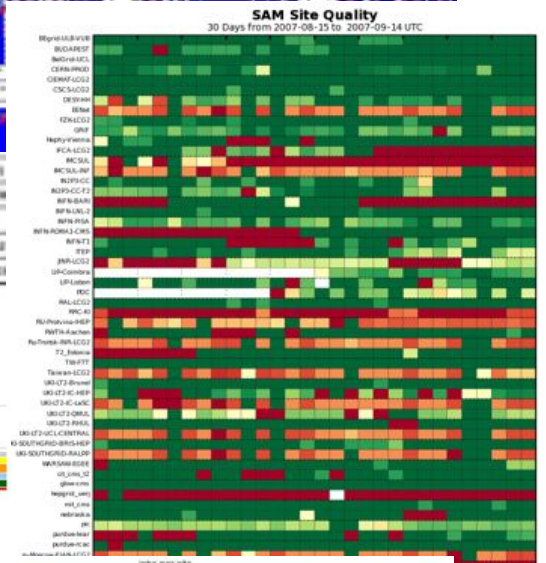

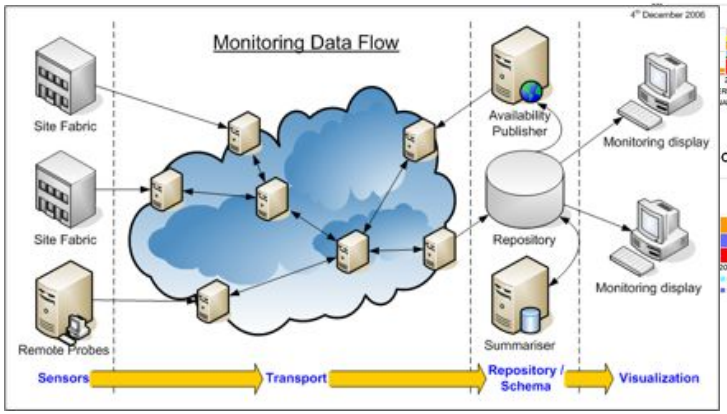

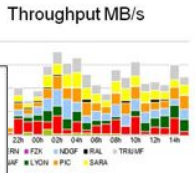

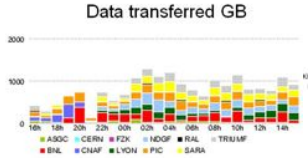

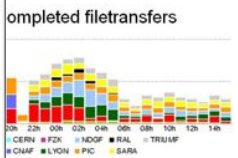

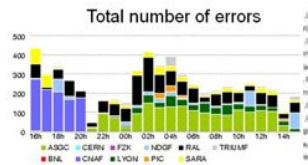

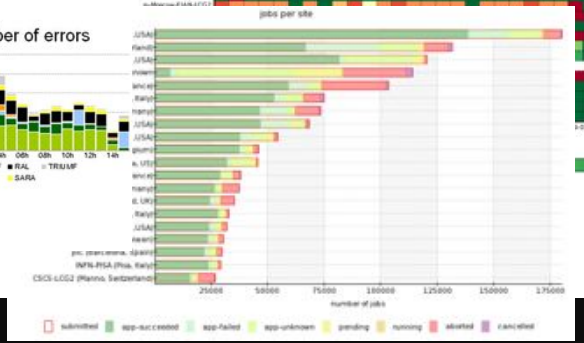

ATLAS M4 Data Monitoring – August 31

GridMap Visualization

Link: http://gridmap.cern.ch

Drilldown into region by clicking on the title

Grid topology view (grouping)

Metric selection for **size** of rectangles

Metric selection for **colour** of rectangles

VO selection

Overall Site or *Site Service* selection

Show SAM status

Show GridView availability data

Context sensitive information

Description of current view

Colour Key

Grid Monitoring

- The critical activity to achieve reliability

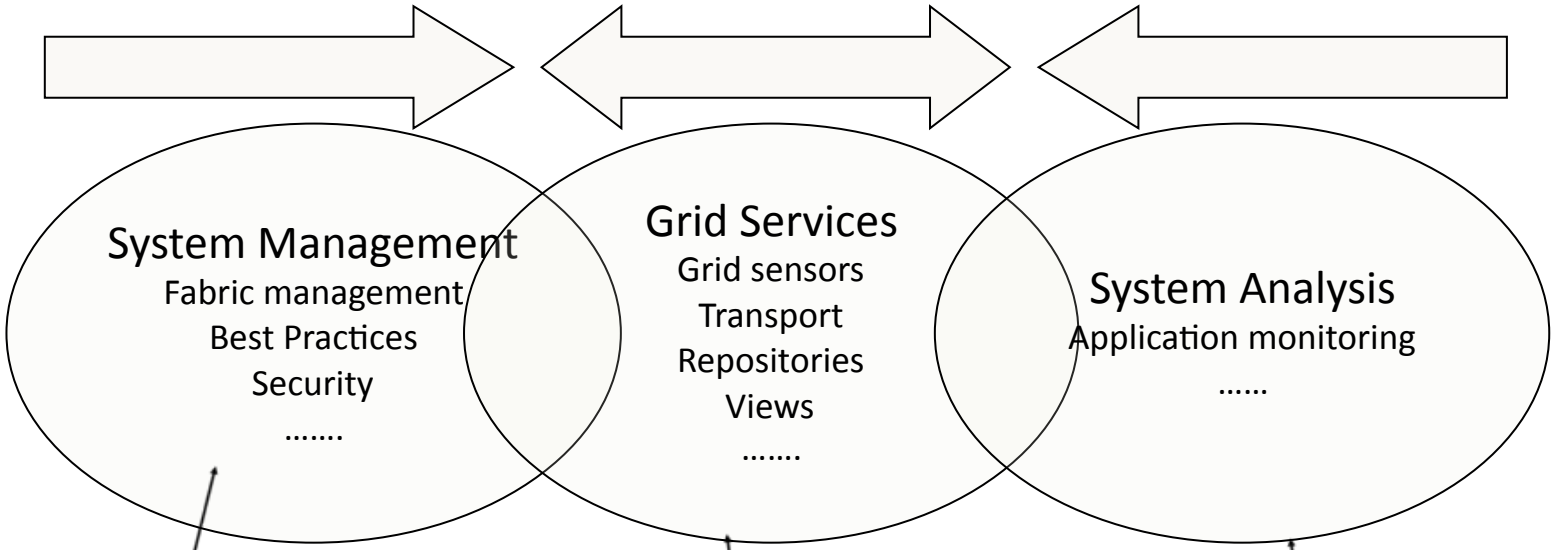

- " … improving system management practices,
- Provide site manager input to requirements on grid monitoring and management tools
- Propose existing tools to the grid monitoring working group
- Produce a Grid Site Fabric Management cook-book
- Identify training needs

- "… To help improve the reliability of the grid infrastructure …"
- " … provide stakeholders with views of the infrastructure allowing them to understand the current and historical status of the service …"

- " … to gain understanding of application failures in the grid environment and to provide an application view of the state of the infrastructure …"

Monitoring

ActiveMQ

- Availability/Reliability monitoring
  - SAM tests and infrastructure
    - Now migrated to NAGIOS based system, decentralized
  - Visualization: GridView, GridMap, dashboards……
  - Solid foundation: Monitoring Infrastructure

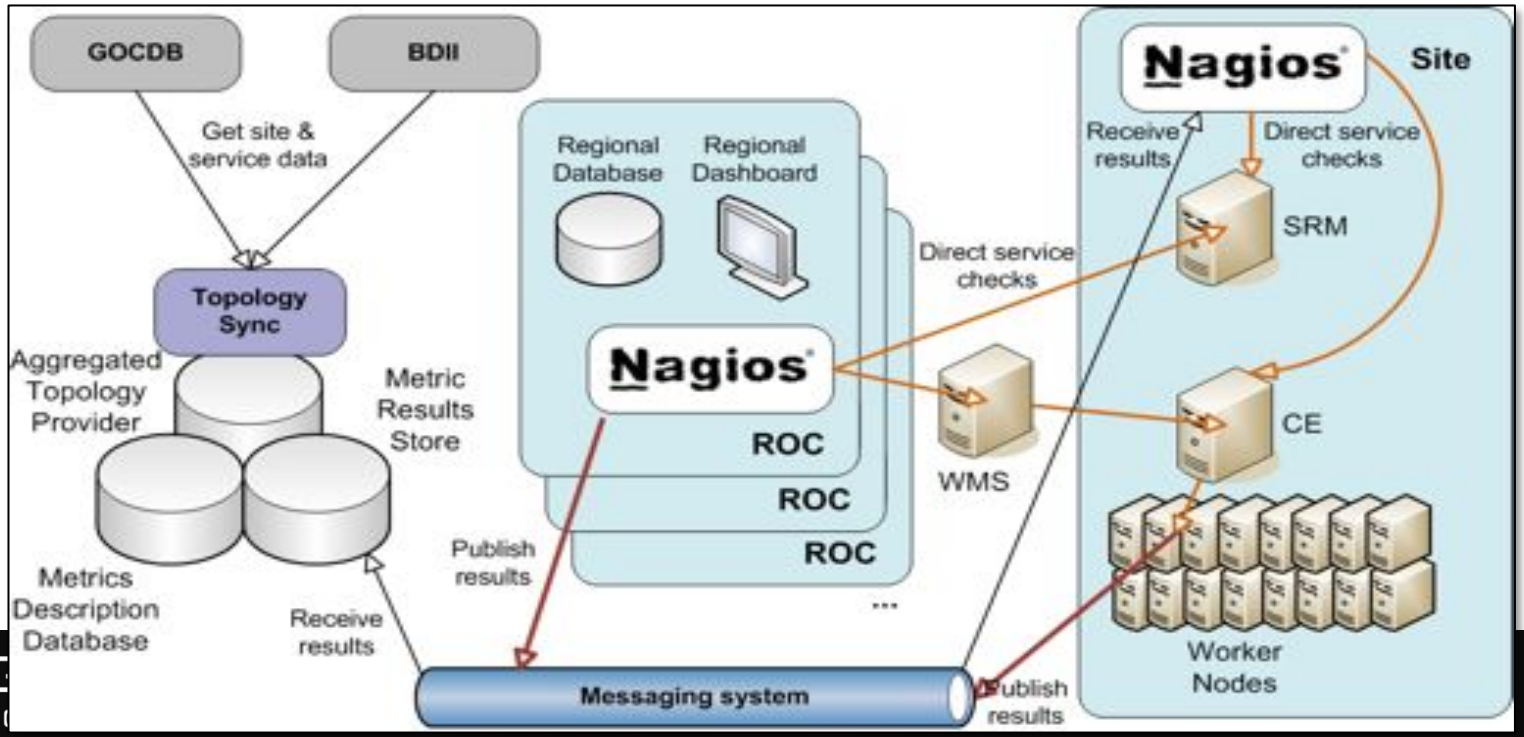

Security & Policy

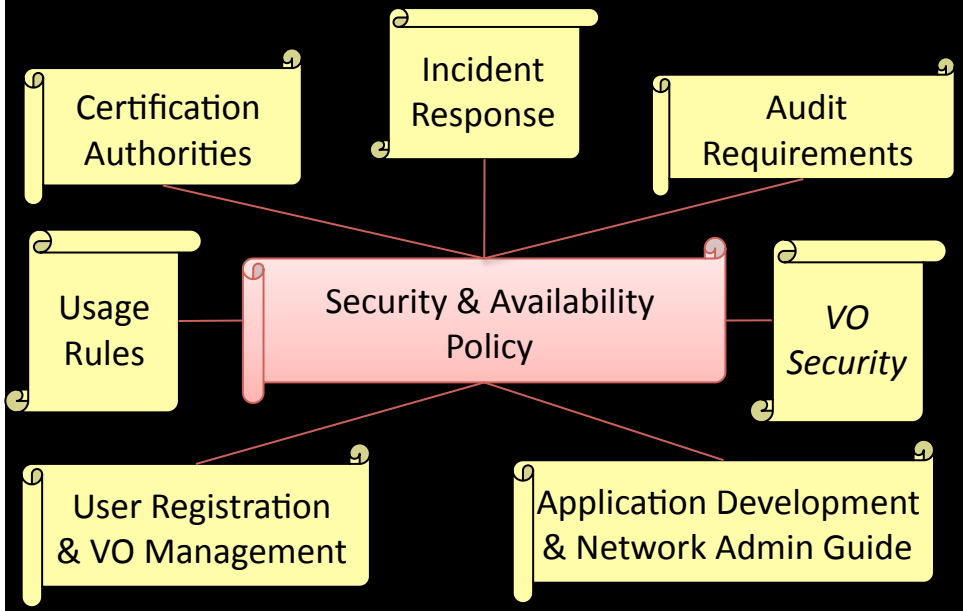

Collaborative policy development
- Many policy aspects are collaborative works; e.g.:

- Joint Security Policy Group
- Certification Authorities
  - EUGridPMA → IGTF, etc.
- Grid Acceptable Use Policy (AUP)
  - common, general and simple AUP
  - for all VO members using many Grid infrastructures
    - EGEE, OSG, SEE-GRID, DEISA, national Grids…
- Incident Handling and Response
  - defines basic communications paths
  - defines requirements (MUSTs) for IR
  - not to replace or interfere with local response plans

Security & Policy Groups

Joint Security Policy Group	EuGridPMA (& IGTF)
Grid Security Vulnerability Group	Operations Advisory Group (+NA4)

Security groups

- Joint Security Policy Group:
  - Joint with WLCG, OSG, and others
  - Focus on policy issues
  - Strong input to e-IRG
- EUGridPMA
  - Pan-European trust federation of CAs
  - Included in IGTF (and was model for it)
  - Success: most grid projects now subscribe to the IGTF
- Grid Security Vulnerability Group
  - Looking at how to manage vulnerabilities
  - Risk analysis is fundamental
  - Hard to balance between openness and giving away insider info
- Operational Security Coordination Team
  - Main day-to-day operational security work
  - Incident response and follow up
  - Members in all ROCs and sites
  - Frequent tests (Security Challenges)

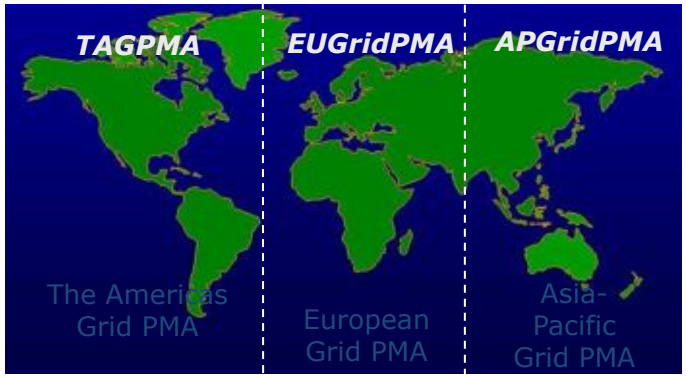

The new European Network Backbone

- LCG working group with Tier-1s and national/ regional research network organisations
- New GÉANT 2 – research network backbone

  → Strong correlation with major European LHC centres (Swiss PoP at CERN)
  → Core links are fibre
- Two 622 Mbps circuits to Israel

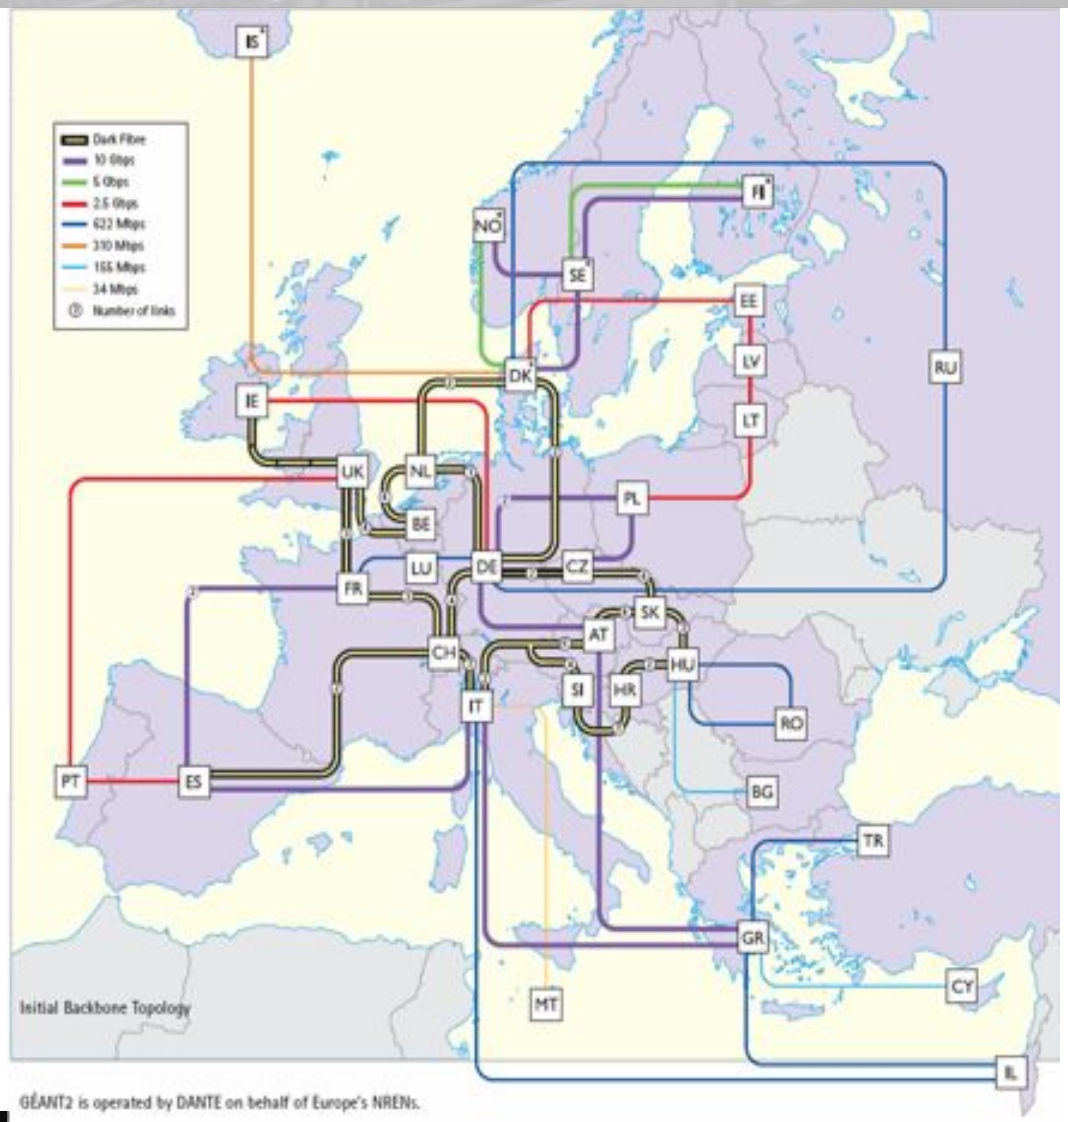

Usage

WLCG Worldwide LHC Computing Grid

Global Multi Science Infrastructure mission critical for many communities

eGee
Enabling Grids for E-scien

Open Science Grid

NDGF
Nordic DataGrid Facility

Arche
Astro
Astro
Civil
Comp
Earth
Finan
Fusio
Geop
High
Life S
Multi
Mater
…

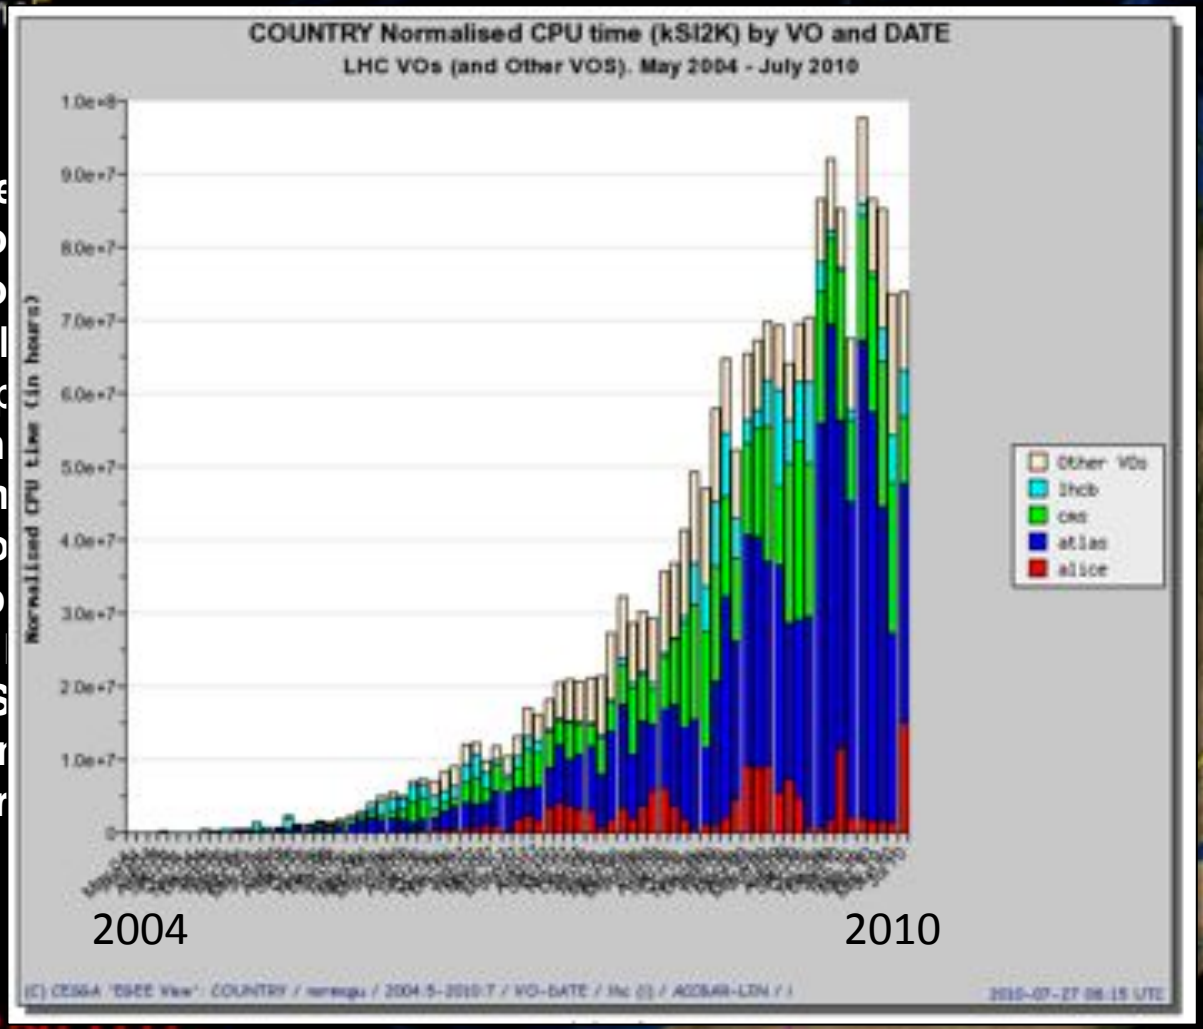

>340sites
48 countries
>170,000 CPUs
>25 PetaBytes disk
>10,000 users
>170 communities
>500,000 jobs/day

GridPP
UK Computing for Particle Physics

21:13:

Grid Activity – delivered CPU

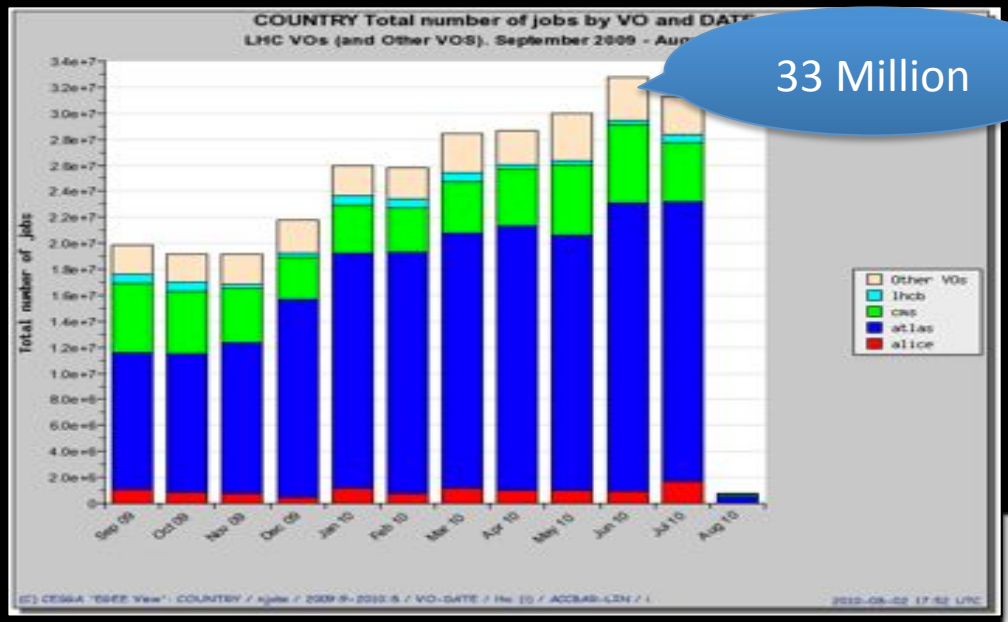

33 Million

- 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 2 sites

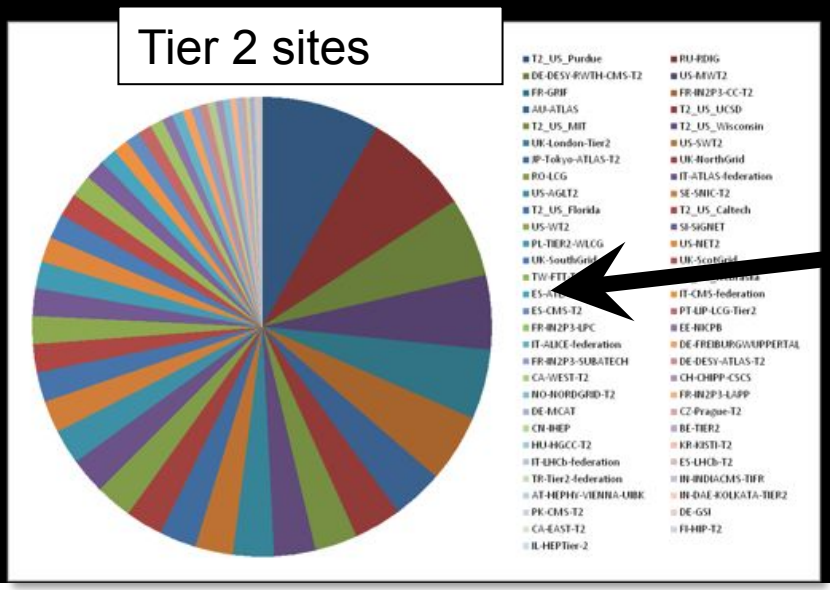

Tier 0 + Tier 1 sites

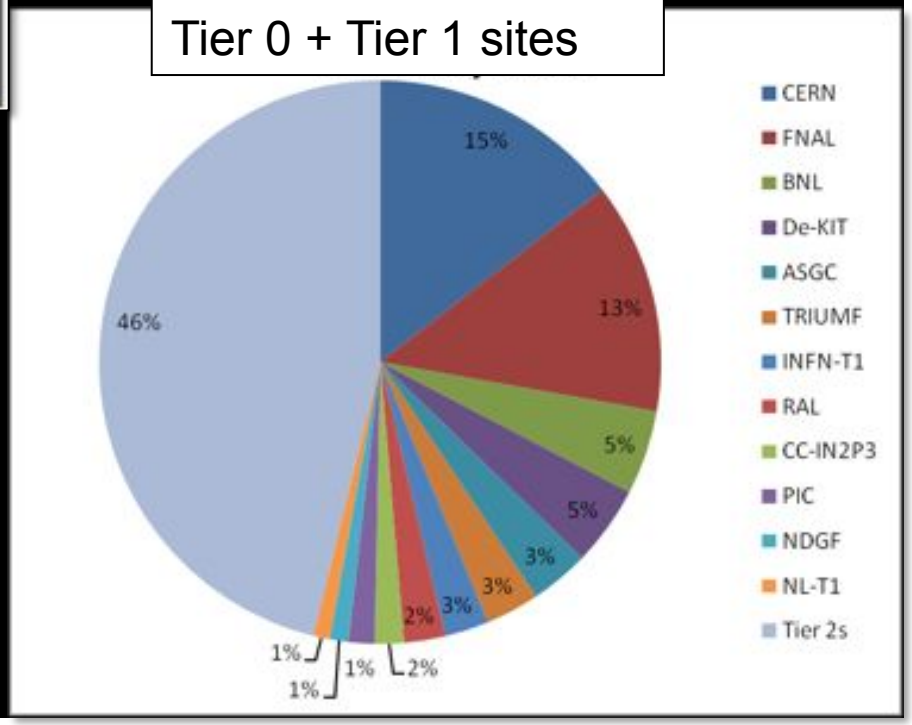

Data Transfers

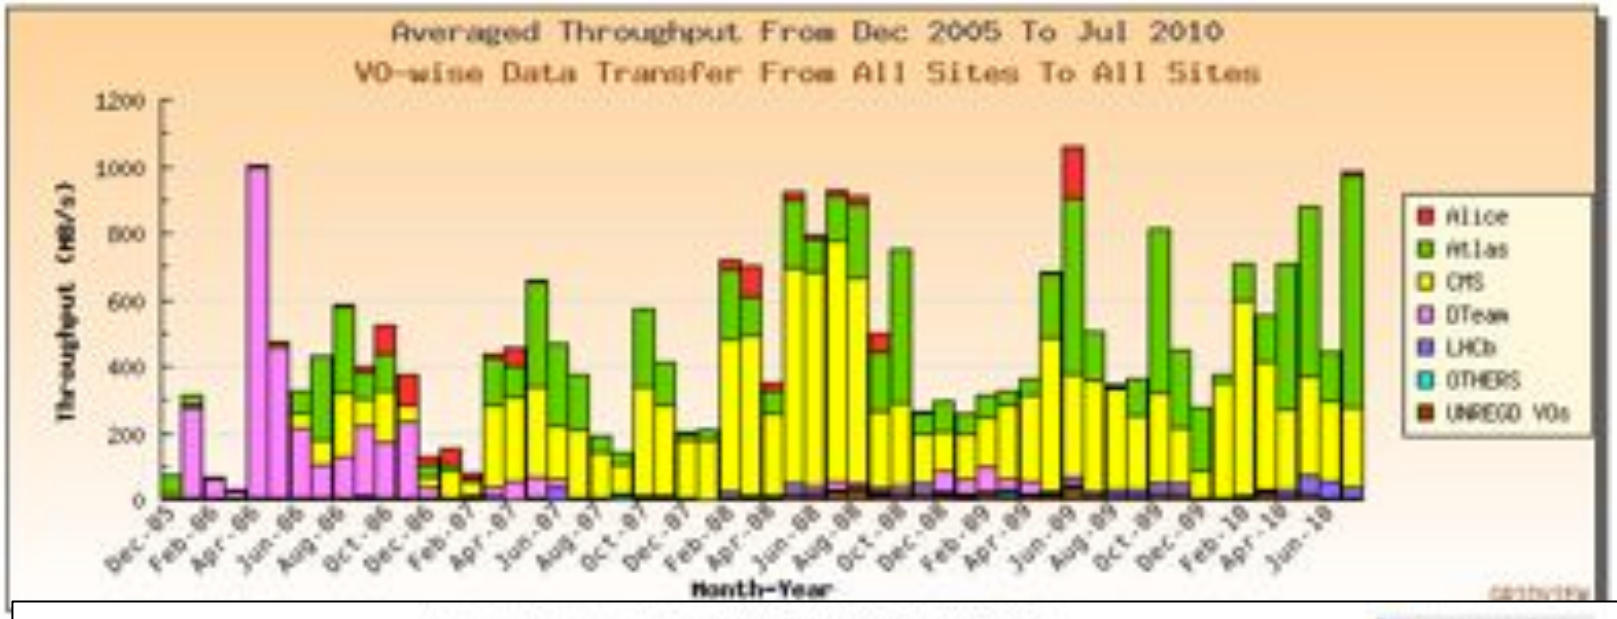

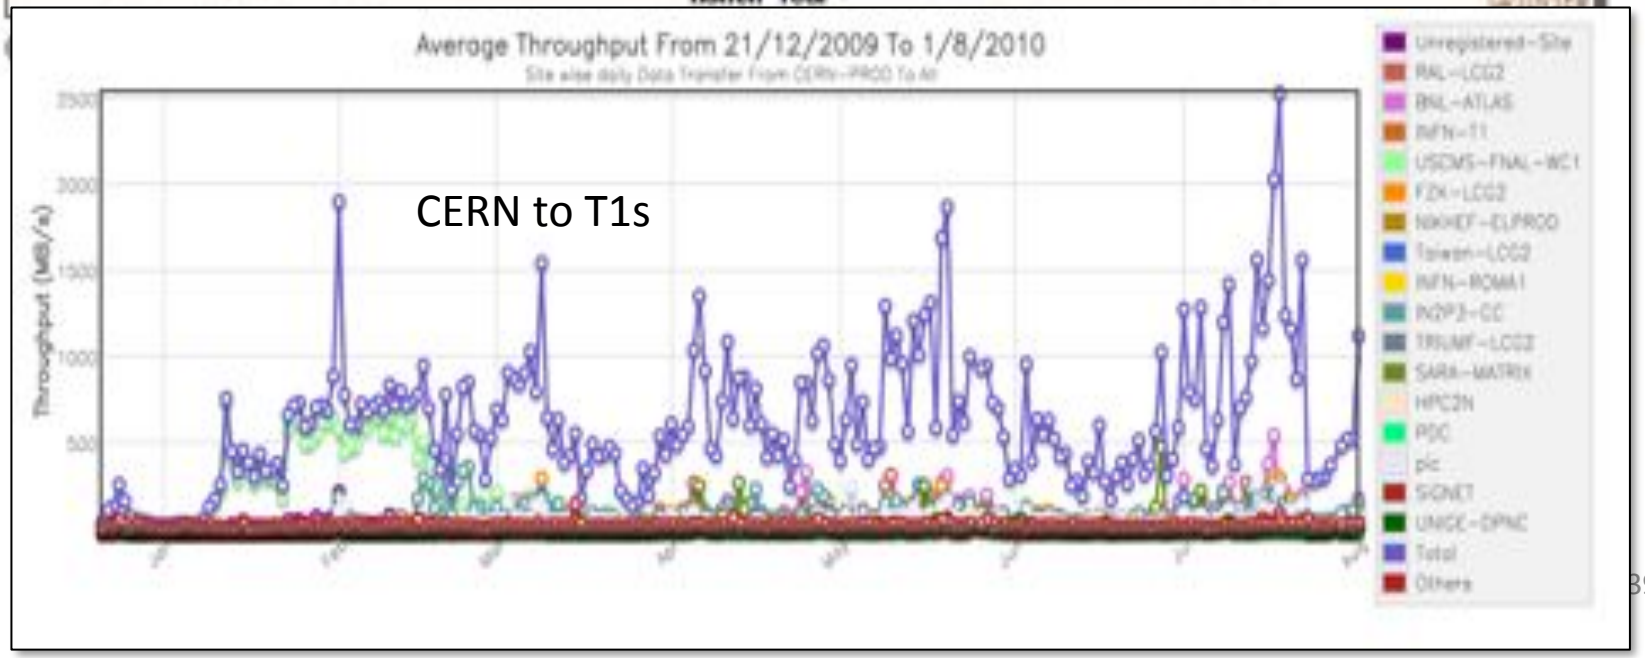

EGEE Achievements - Applications

- >270 VOs from several scientific domains
  - Astronomy & Astrophysics
  - Civil Protection
  - Computational Chemistry
  - Comp. Fluid Dynamics
  - Computer Science/Tools
  - Condensed Matter Physics
  - Earth Sciences
  - Fusion
  - High Energy Physics
  - Life Sciences
- Further applications under evaluation

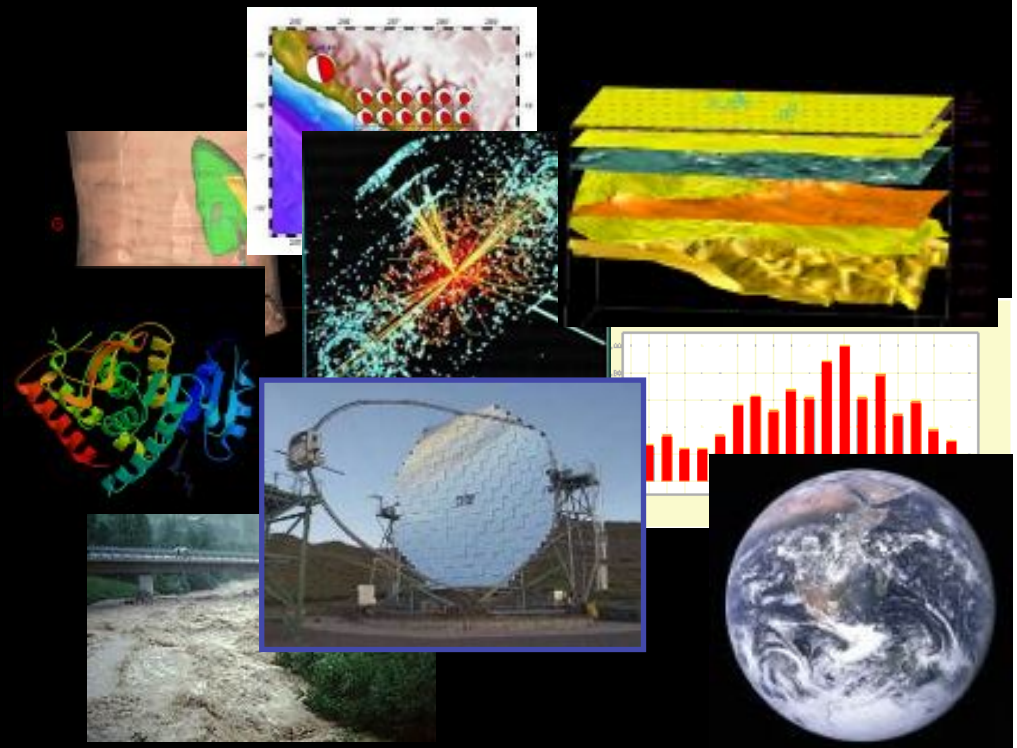

Applications have moved from testing to routine and daily usage

EGEE Registered Collaborating Projects

web page

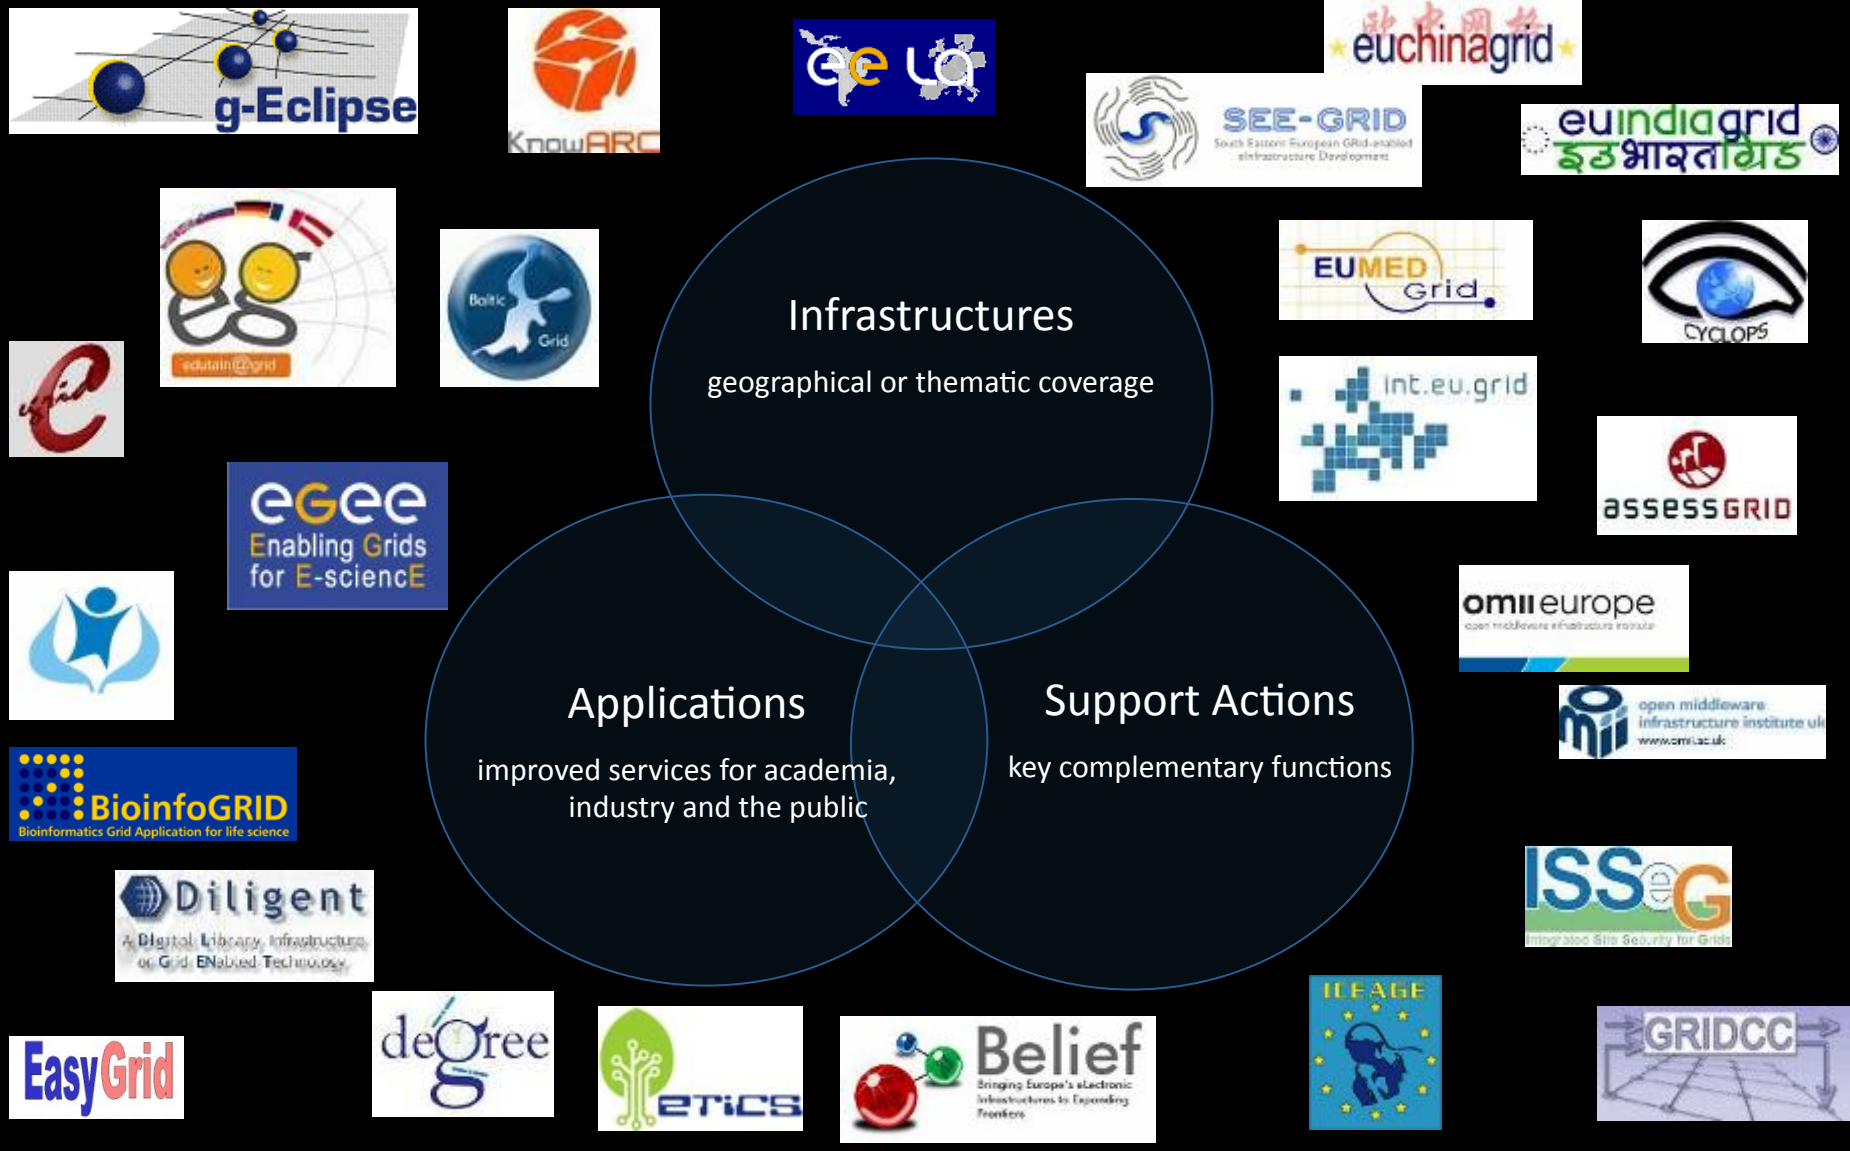

Future

European e-Infrastructure

- European Data Grid (EDG)
  - Explore concepts in a testbed
- Enabling Grid for E-sciencE (EGEE)
  - Moving from prototype to production
- European Grid Infrastructure (EGI)
  - Routine usage of a sustainable e-infrastructure

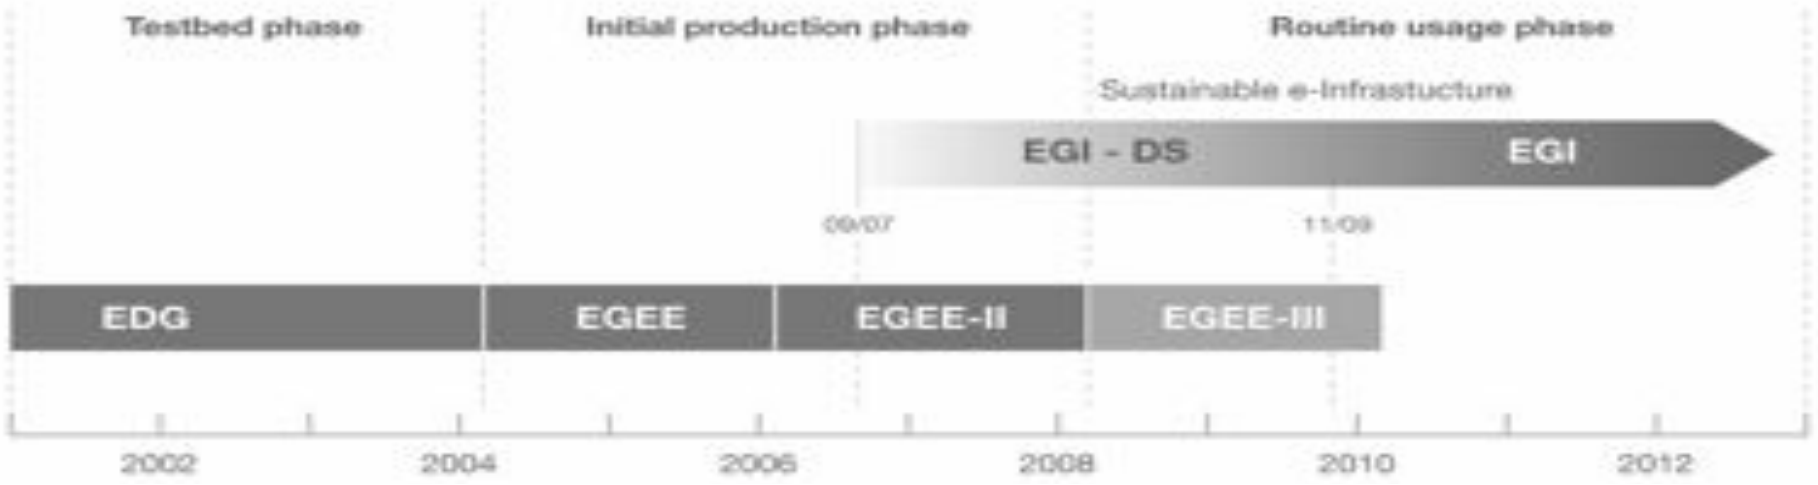

NGIs in Europe

www.eu-egi.eu

What is an e-Infrastructure?

- Resources linked by high speed networks
  - Compute, Storage, Instruments, ...
- Controlled access to shared resources
  - Authentication, Authorisation, Accounting, ...
- Dependable services for others to use
  - Driven by availability and reliability metrics
- Services that are there for the long-term
  - Supporting experiments lasting decades

Moving from EGEE to EGI

- What is different?
  - EGEE did 'everything'
  - EGI focuses just on infrastructure operations
- What is the same?
  - Running a 24/7 production quality infrastructure
  - Providing a support framework for the users
  - Collaboration to drive European DCI forward

What other changes?

- Virtualization (fabric)
- Integration with Clouds
  - Commercial
  - Community
- T0-T1-T2 connection reaches a new quality
  - Less hierarchical system
- Addressing Interoperability
- The end of Moore's law?
- ……..

WLCG Worldwide LHC Computing Grid

Summary

- WLCG delivers the computing resources needed by the experiments
- The infrastructure grew exponentially for several years
  - While reducing support effort and improving reliability
- The main challenges are:
  - Scaling
  - Changes in technology
  - Organizational changes