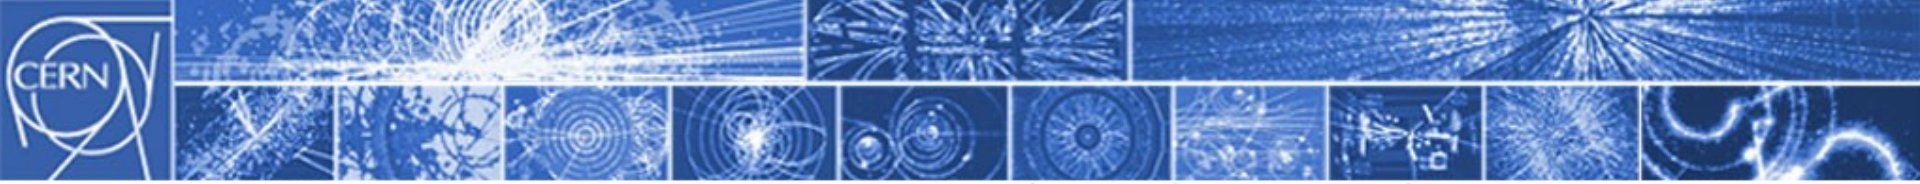

CERN

European Organization for Nuclear Research
Organisation Européenne pour la Recherche Nucléaire

Massive Computing at CERN and lessons learnt

Bob Jones

CERN

Bob.Jones <at> CERN.ch

WLCG – what and why?

- A distributed computing infrastructure to provide the production and analysis environments for the LHC experiments
- Managed and operated by a worldwide collaboration between the experiments and the participating computer centres
- The resources are distributed – for funding and sociological reasons
- Our task is to make use of the resources available to us – no matter where they are located

What is WLCG today?

Collaboration

- Coordination & management & reporting
- Coordinate resources & funding
- Coordination with service & technology providers
- Common requirements
- Memorandum of Understanding

Framework

- Service management
- Service coordination
- Operational security
- Support processes & tools
- Common tools
- Monitoring & Accounting
- World-wide trust federation for CA's and VO's
- Complete Policy framework

Distributed Computing services

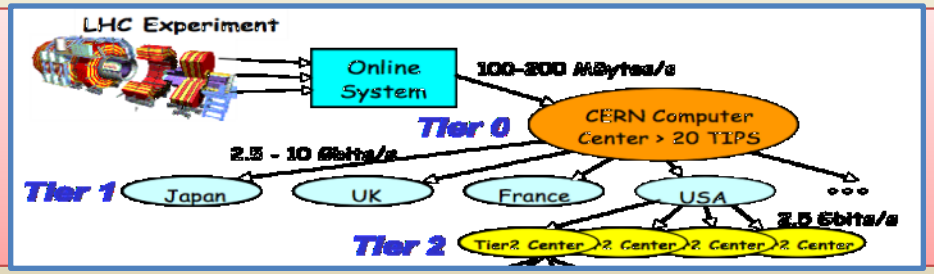

Physical resources: CPU, Disk, Tape, Networks

WLCG data processing model

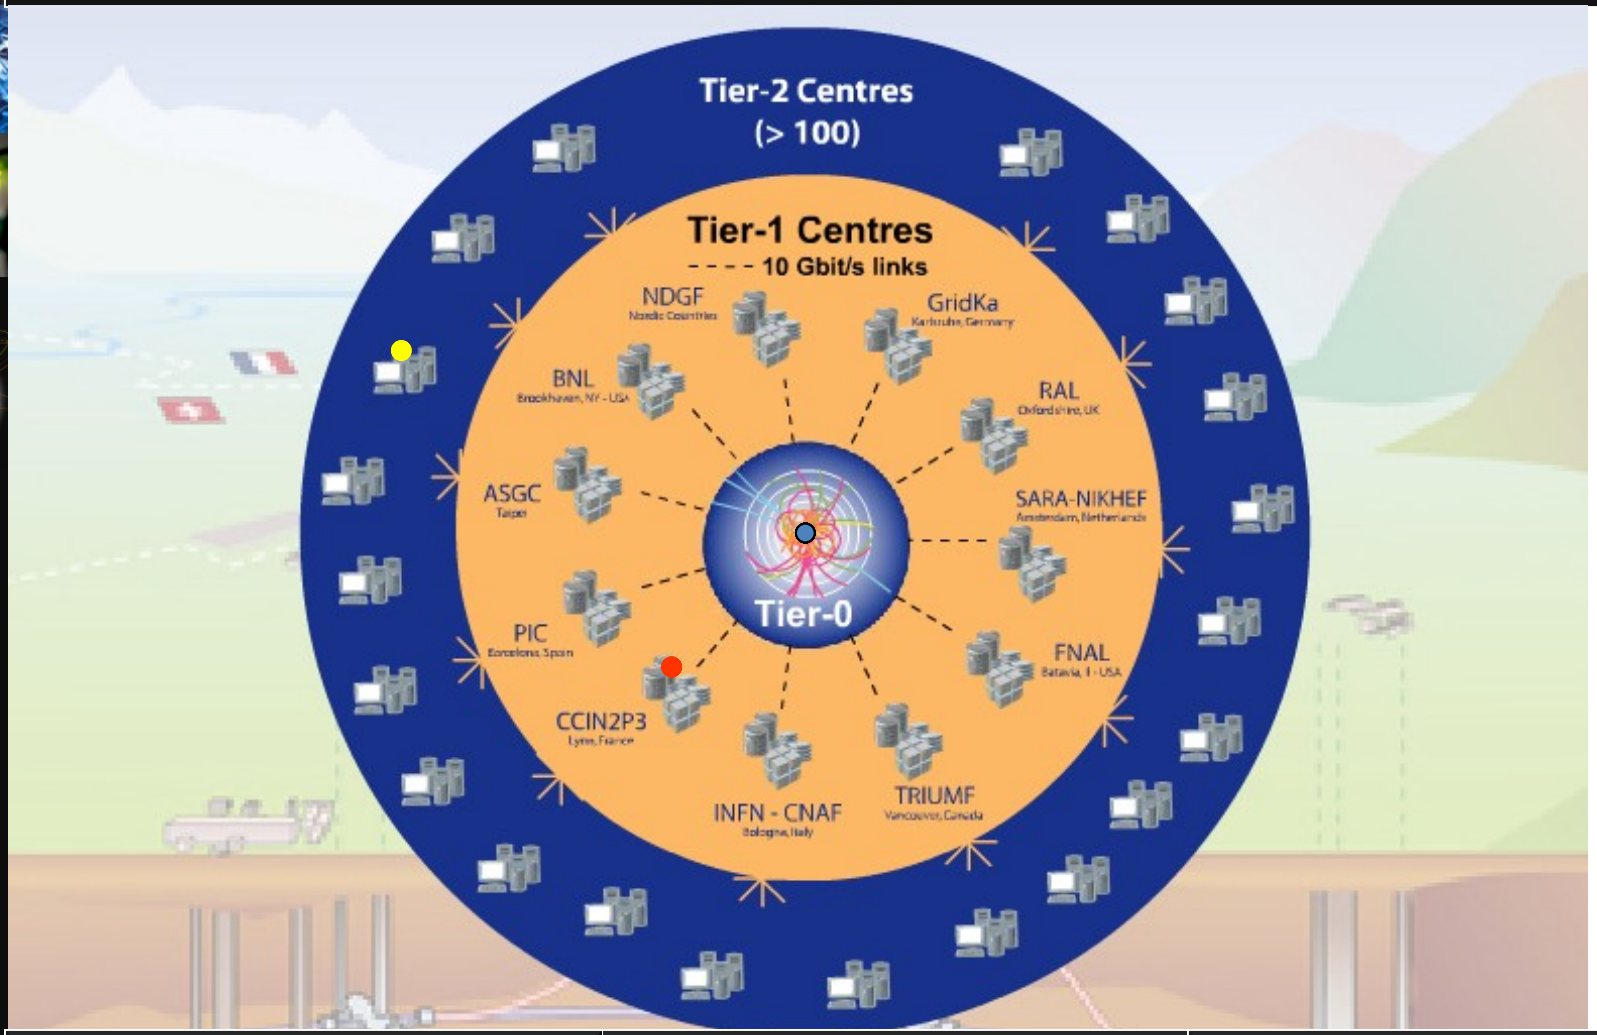

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

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

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

WLCG Worldwide LHC Computing Grid

WLCG Collaboration Status
Tier 0; 11 Tier 1s; 64 Tier 2 federations

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.

LHCOPN – current status

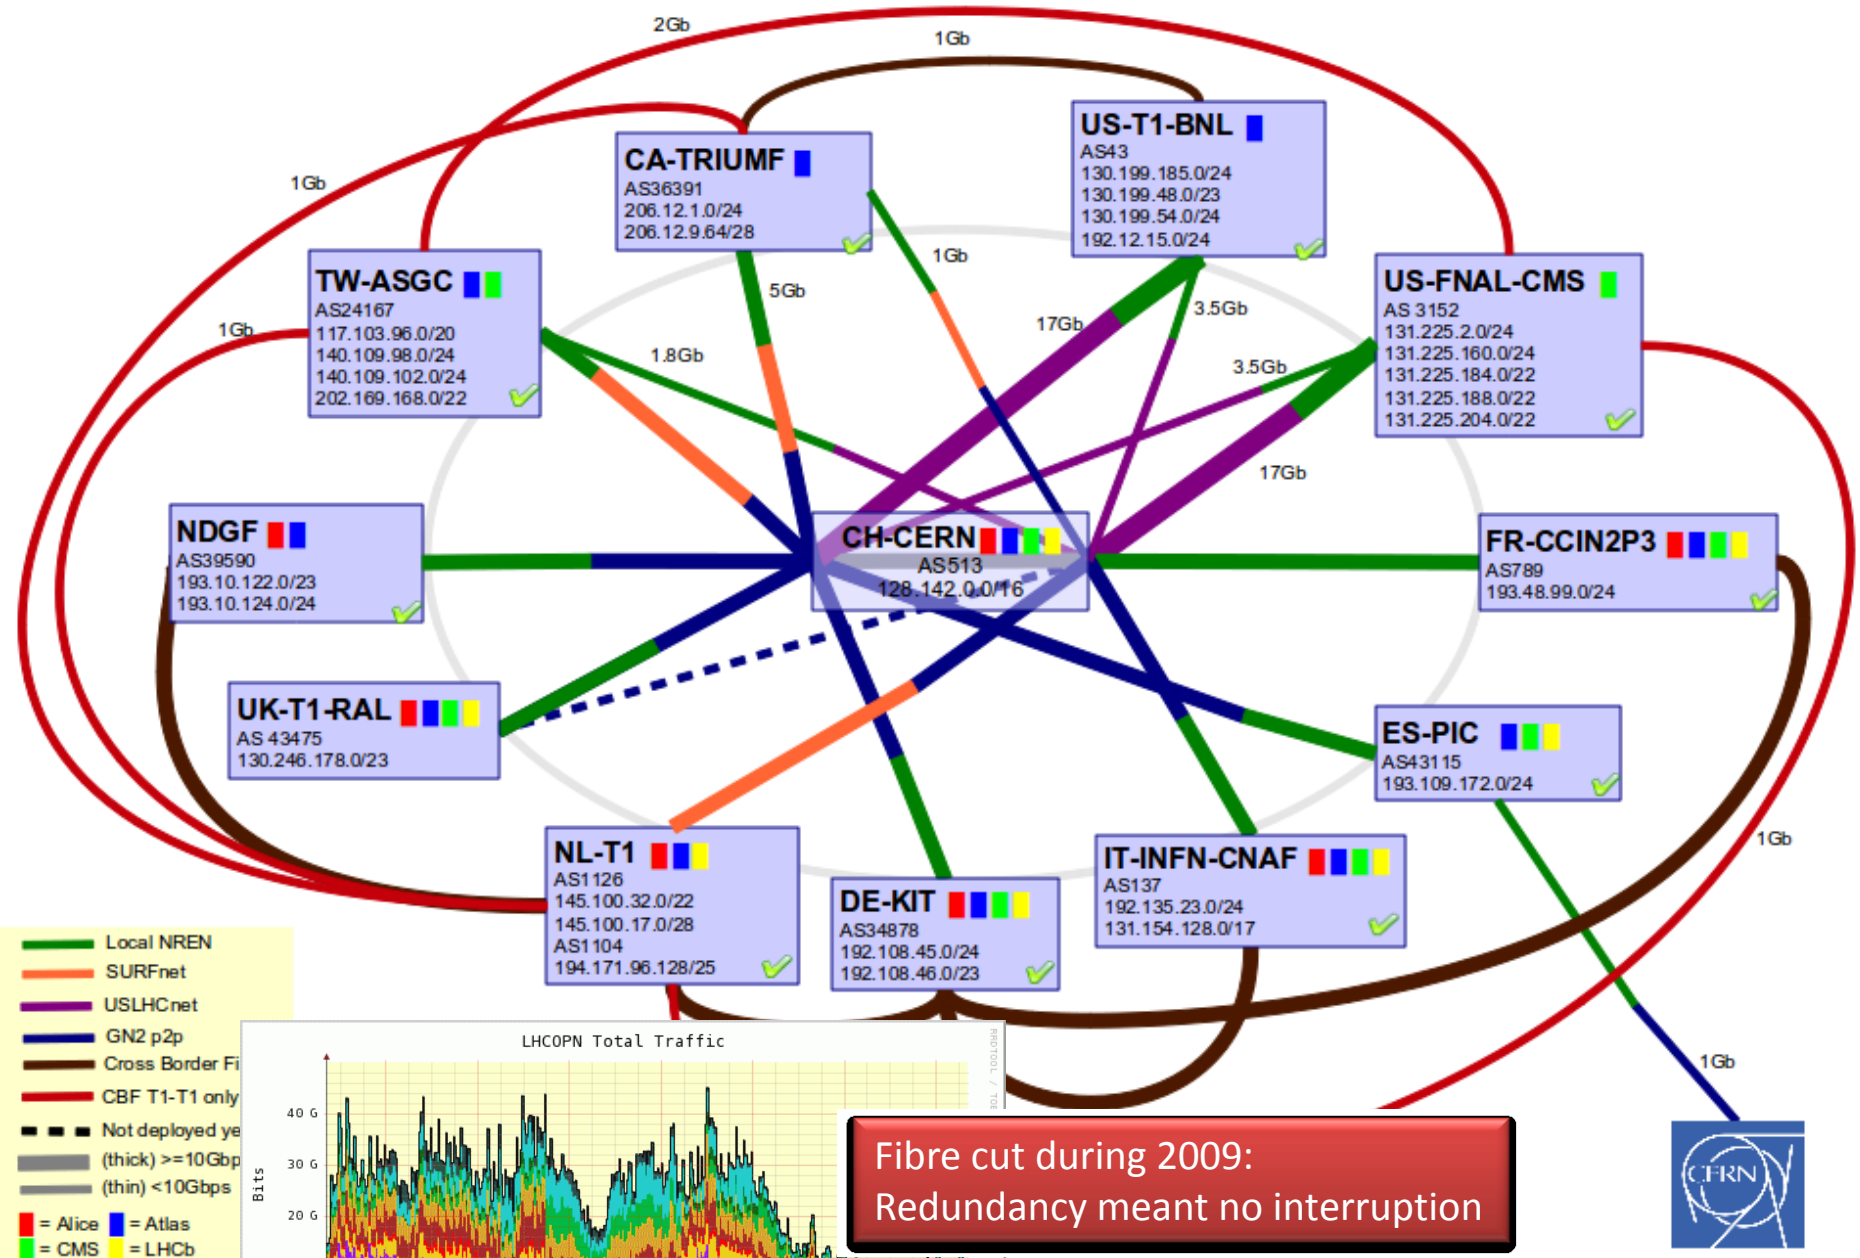

edoardo.martelli@cern.ch 20100205

Worldwide resources

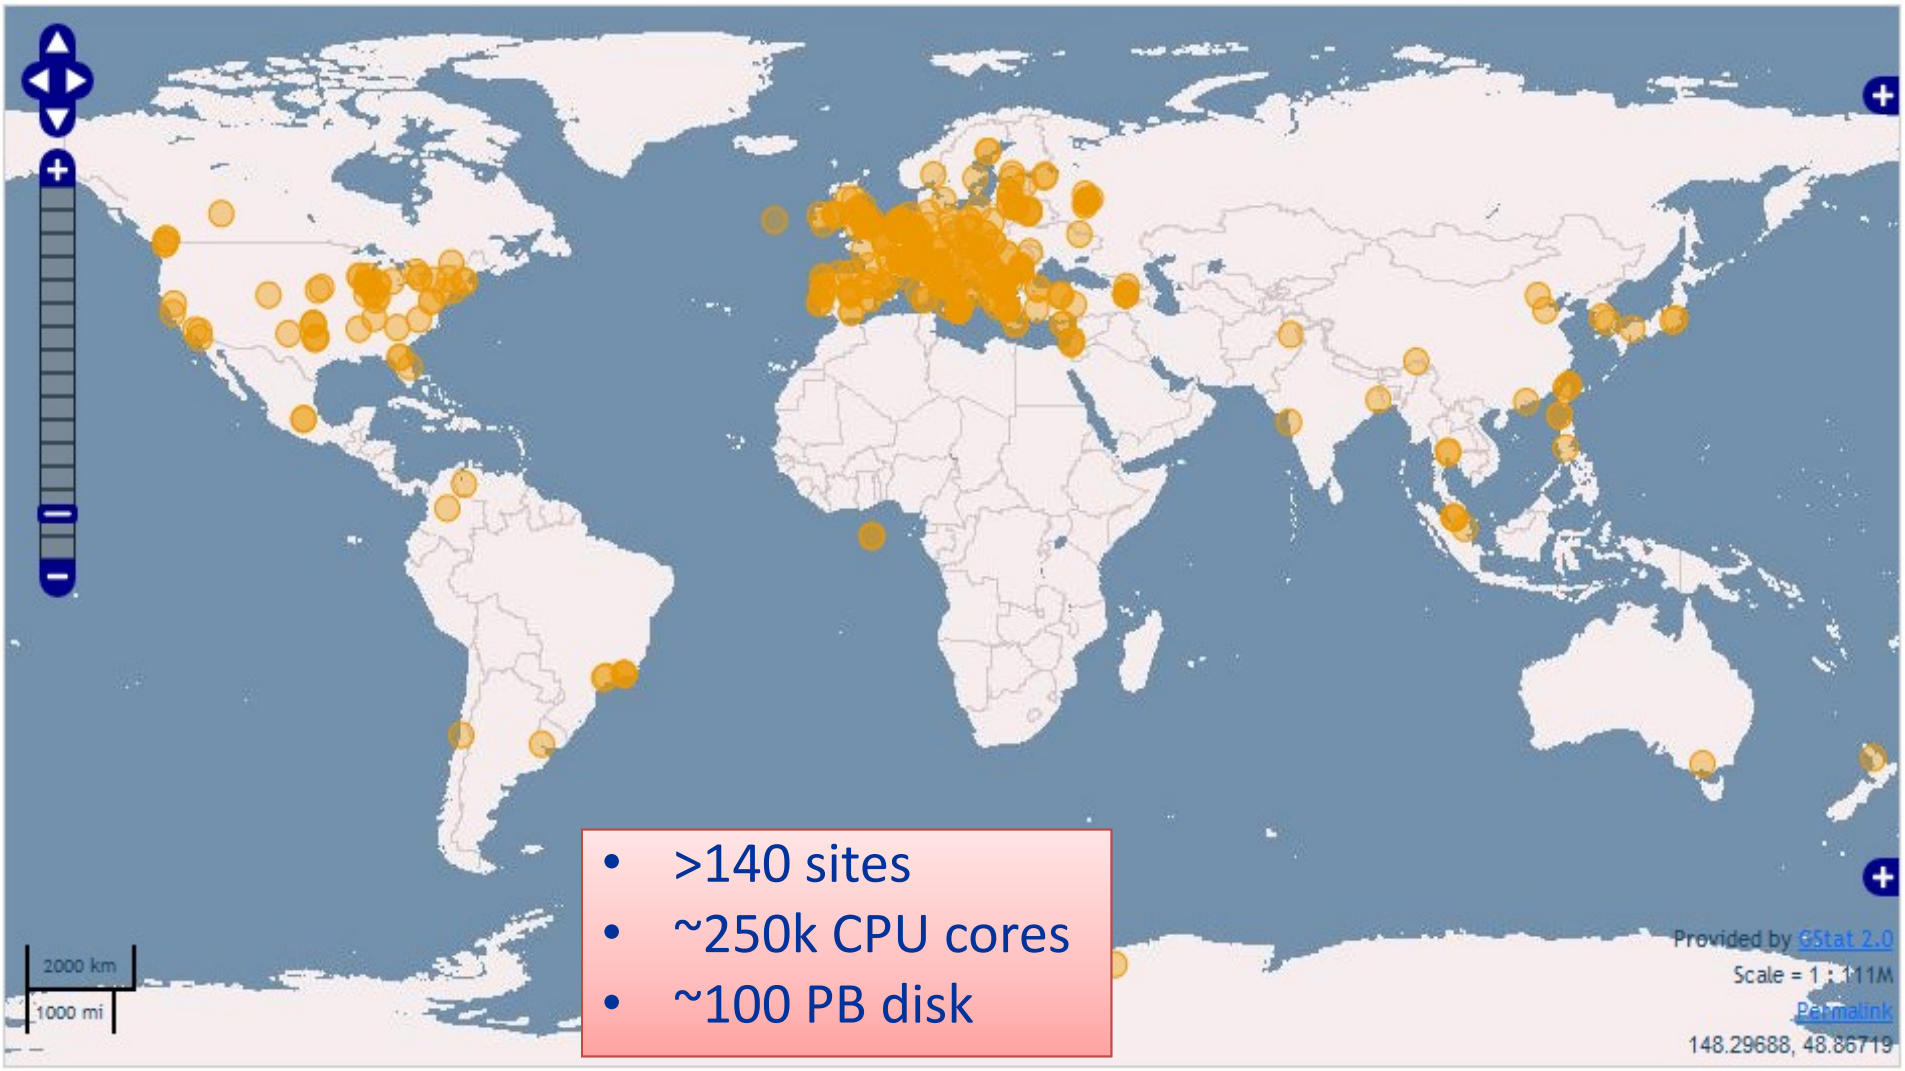

- >140 sites
- ~250k CPU cores
- ~100 PB disk

Service quality: defined in MoU

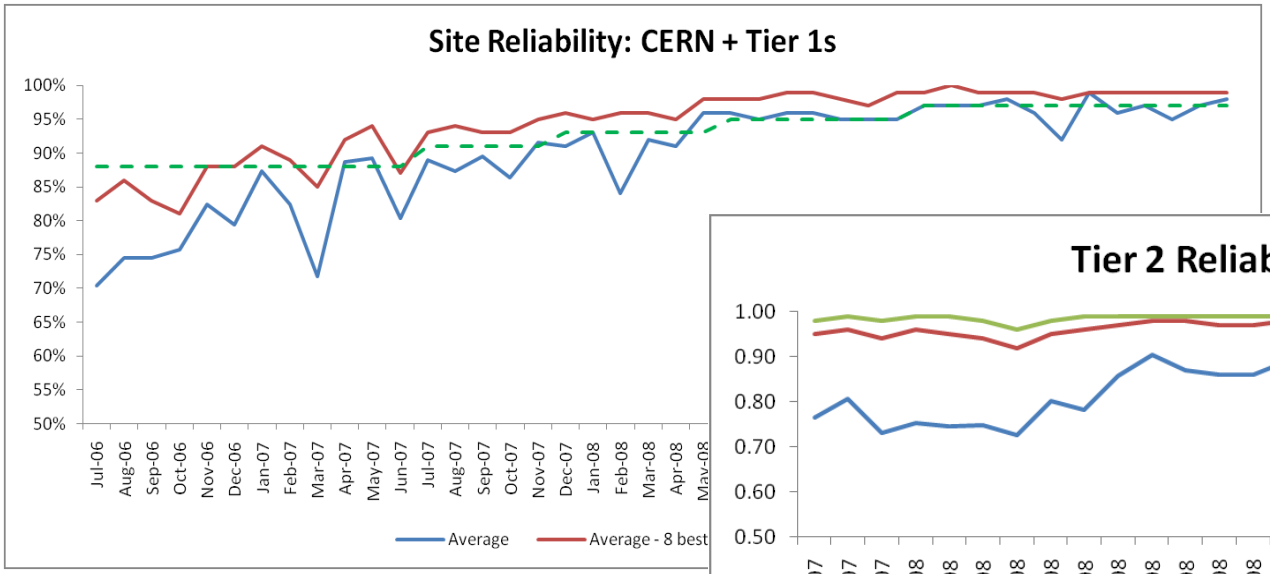

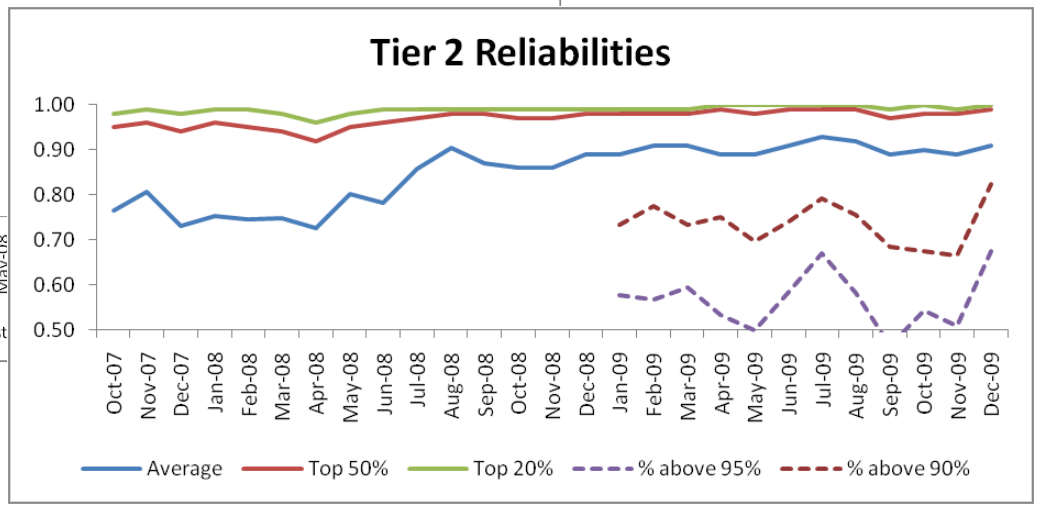

- MoU defines key performance and support metrics for Tier 1 and Tier 2 sites
  - Reliabilities are an approximation for some of these
  - Also metrics on response times, resources, etc.
- The MoU has been an important tool in bringing services to an acceptable level

From testing to data:

Independent Experiment Data Challenges

Service Challenges proposed in 2004

To demonstrate service aspects:

- -Data transfers for weeks on end
- -Data management
- -Scaling of job workloads
- -Security incidents ("fire drills")
- -Interoperability
- -Support processes

- Focus on real and continuous production use of the service over several years (simulations since 2003, cosmic ray data, etc.)
- Data and Service challenges to exercise all aspects of the service – not just for data transfers, but workloads, support structures etc.

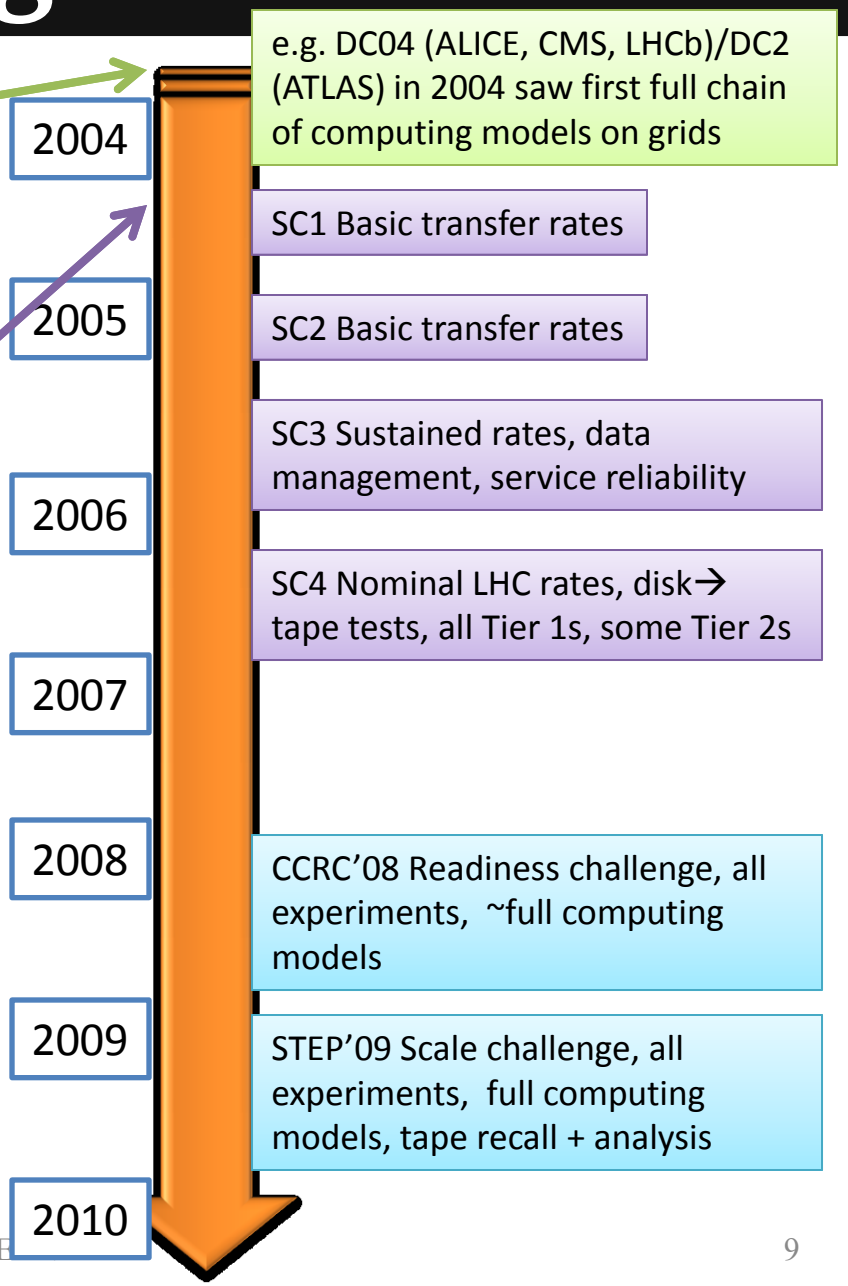

Large scale = long times

- LHC, the experiments, & computing have taken ~20 years to build and commission
- They will run for at least 20 years
- We must be able to rely on long term infrastructures
  - Global networking
  - Strong and stable NGIs (or their evolution)
    - That should be eventually self-sustaining
  - Long term sustainability - must come out of the current short term project funding cycles

Grids & HEP: Common history

- CERN and the HEP community have been involved with grids from the beginning
- Recognised as a key technology for implementing the LHC computing model
- HEP work with EC-funded EDG/EGEE in Europe, iVDGL/Grid3/OSG etc. in US has been of clear *mutual* benefit
  - Infrastructure development driven by HEP needs
  - Robustness needed by WLCG is benefitting other communities
  - Transfer of technology from HEP
    - Ganga, AMGA, etc used by many communities now

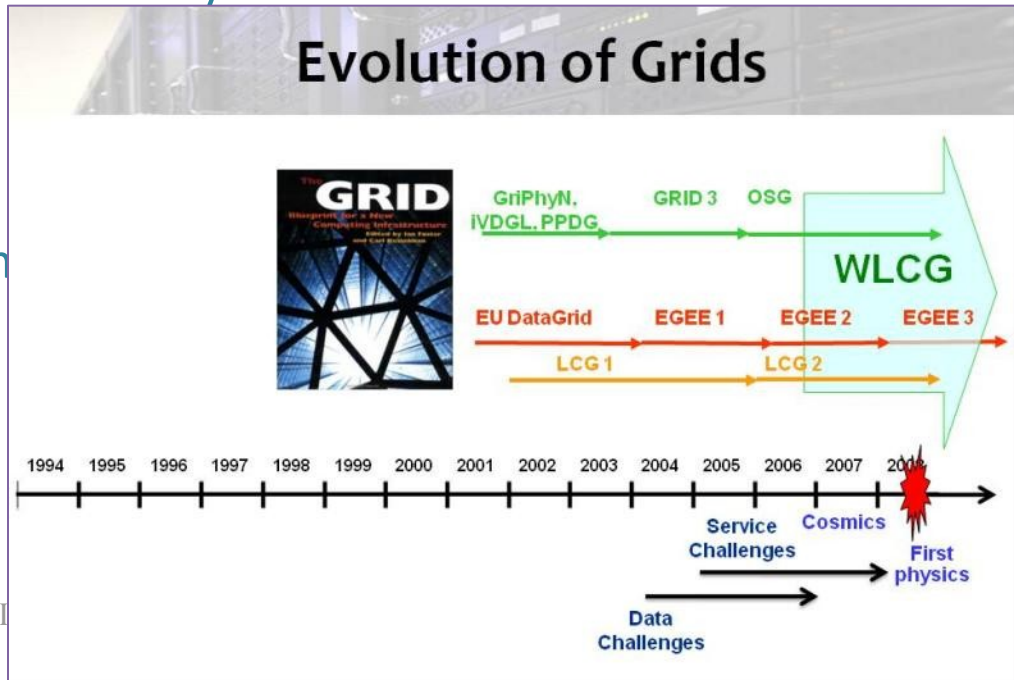

European Grid Infrastructure

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

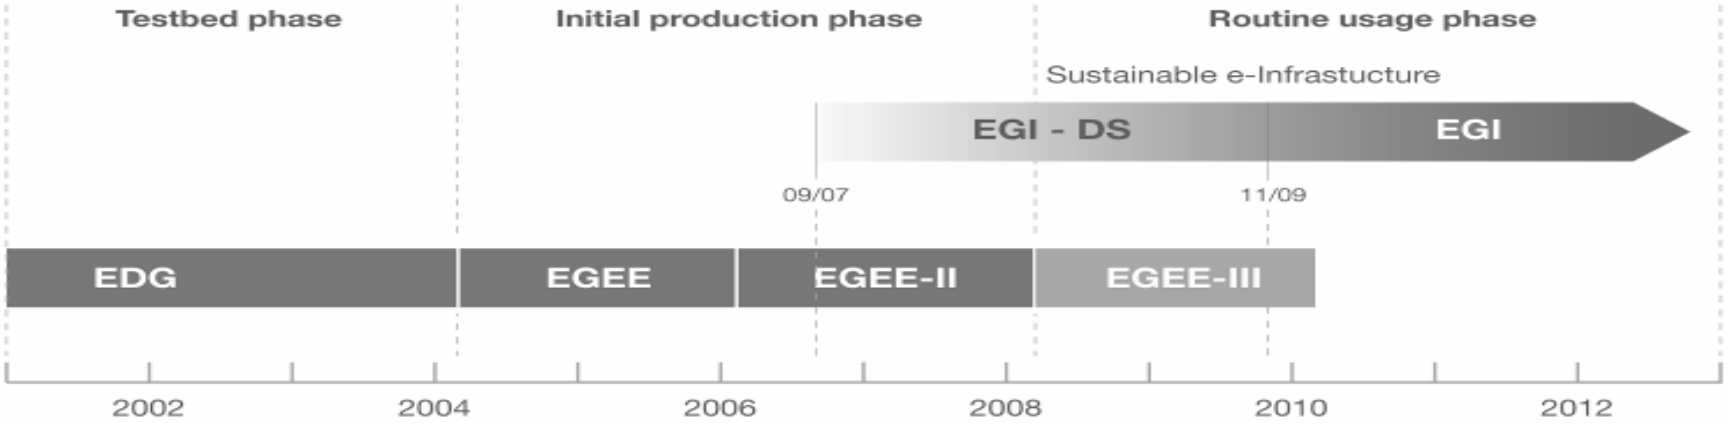

European Grid Infrastructure
(Status April 2011 – yearly increase)

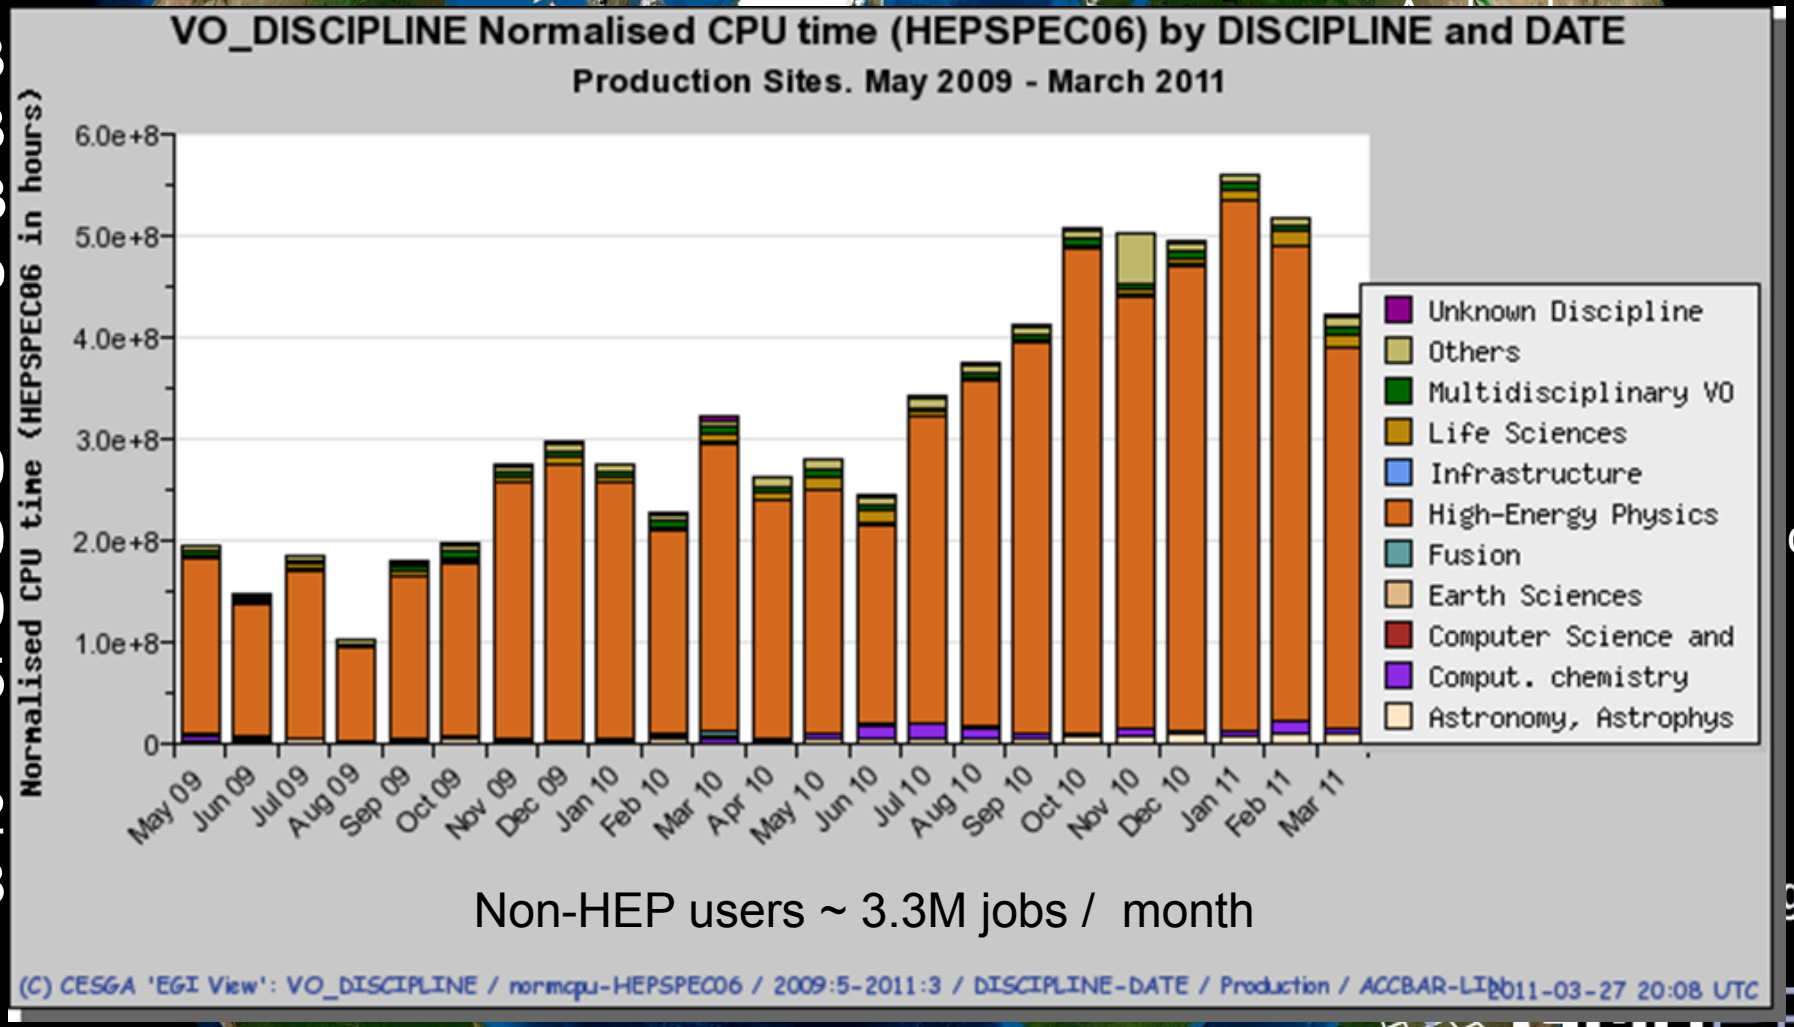

Non-HEP users ~ 3.3M jobs / month

Grids, clouds, supercomputers, etc.

Grids
- Collaborative environment
- Distributed resources (political/sociological)
- Commodity hardware
- (HEP) data management
- Complex interfaces (bug not feature)
- Communities expected to contribute resources

Supercomputers
- Scarce
- Low latency interconnects
- Applications peer reviewed
- Parallel/coupled applications
- Also SC grids (DEISA/PRACE, Teragrid/XD)

Clouds
- Proprietary (implementation)
- Economies of scale in management
- Commodity hardware
- Pay-as-you-go usage model
- Details of physical resources hidden
- Simple interfaces

Volunteer computing
- Simple mechanism to access millions CPUs
- Difficult if (much) data involved
- Control of environment ➔ check
- Community building – people involved in Science
- Potential for huge amounts of real work

Collaboration with the General Public: Citizen Cyberscience Centre

CERN IT Department

Philosophy: promote web-based citizen participation in science projects as an appropriate low cost technology for scientists in the developing world.

Partners: CERN, UN Institute for Training and Research, University of Geneva

Sponsors: IBM, HP Labs, Shuttleworth Foundation

Technology: open source platforms for internet-based distributed collaboration

Projects:

- **Computing for Clean Water** optimizing nanotube based water filters by large scale simulation on volunteer PCs
- **AfricaMap** volunteer thinking to generate maps of regions of Africa from satellite images, with UNOSAT
- **LHC@home** new volunteer project for public participation in LHC collision simulations, using VM technology

Plans: Training workshops in 2011 in India, China, Brazil and South Africa

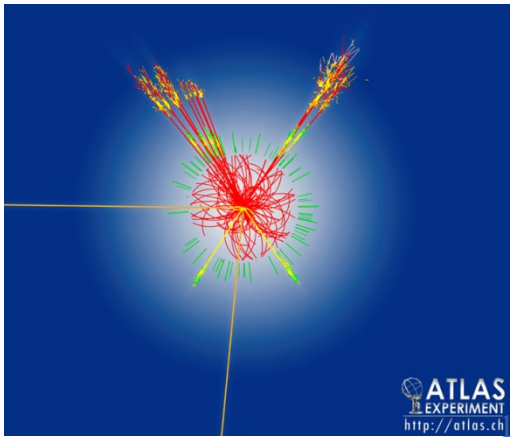

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Frédéric Hemmer

Some more questions to be answered

- Computing model
  - How many computing models exist in the community and can they all use the same computing infrastructure?
- Continuous load or periodic campaigns?
  - How intensely and frequently will the community use the computing infrastructure?
- Manpower
  - Do you have enough geeks to port the code and support it?
- How committed is the community?
  - Are you prepared to contribute and share computing resources?