Computing for the LHC
Worldwide LHC Computing Grid

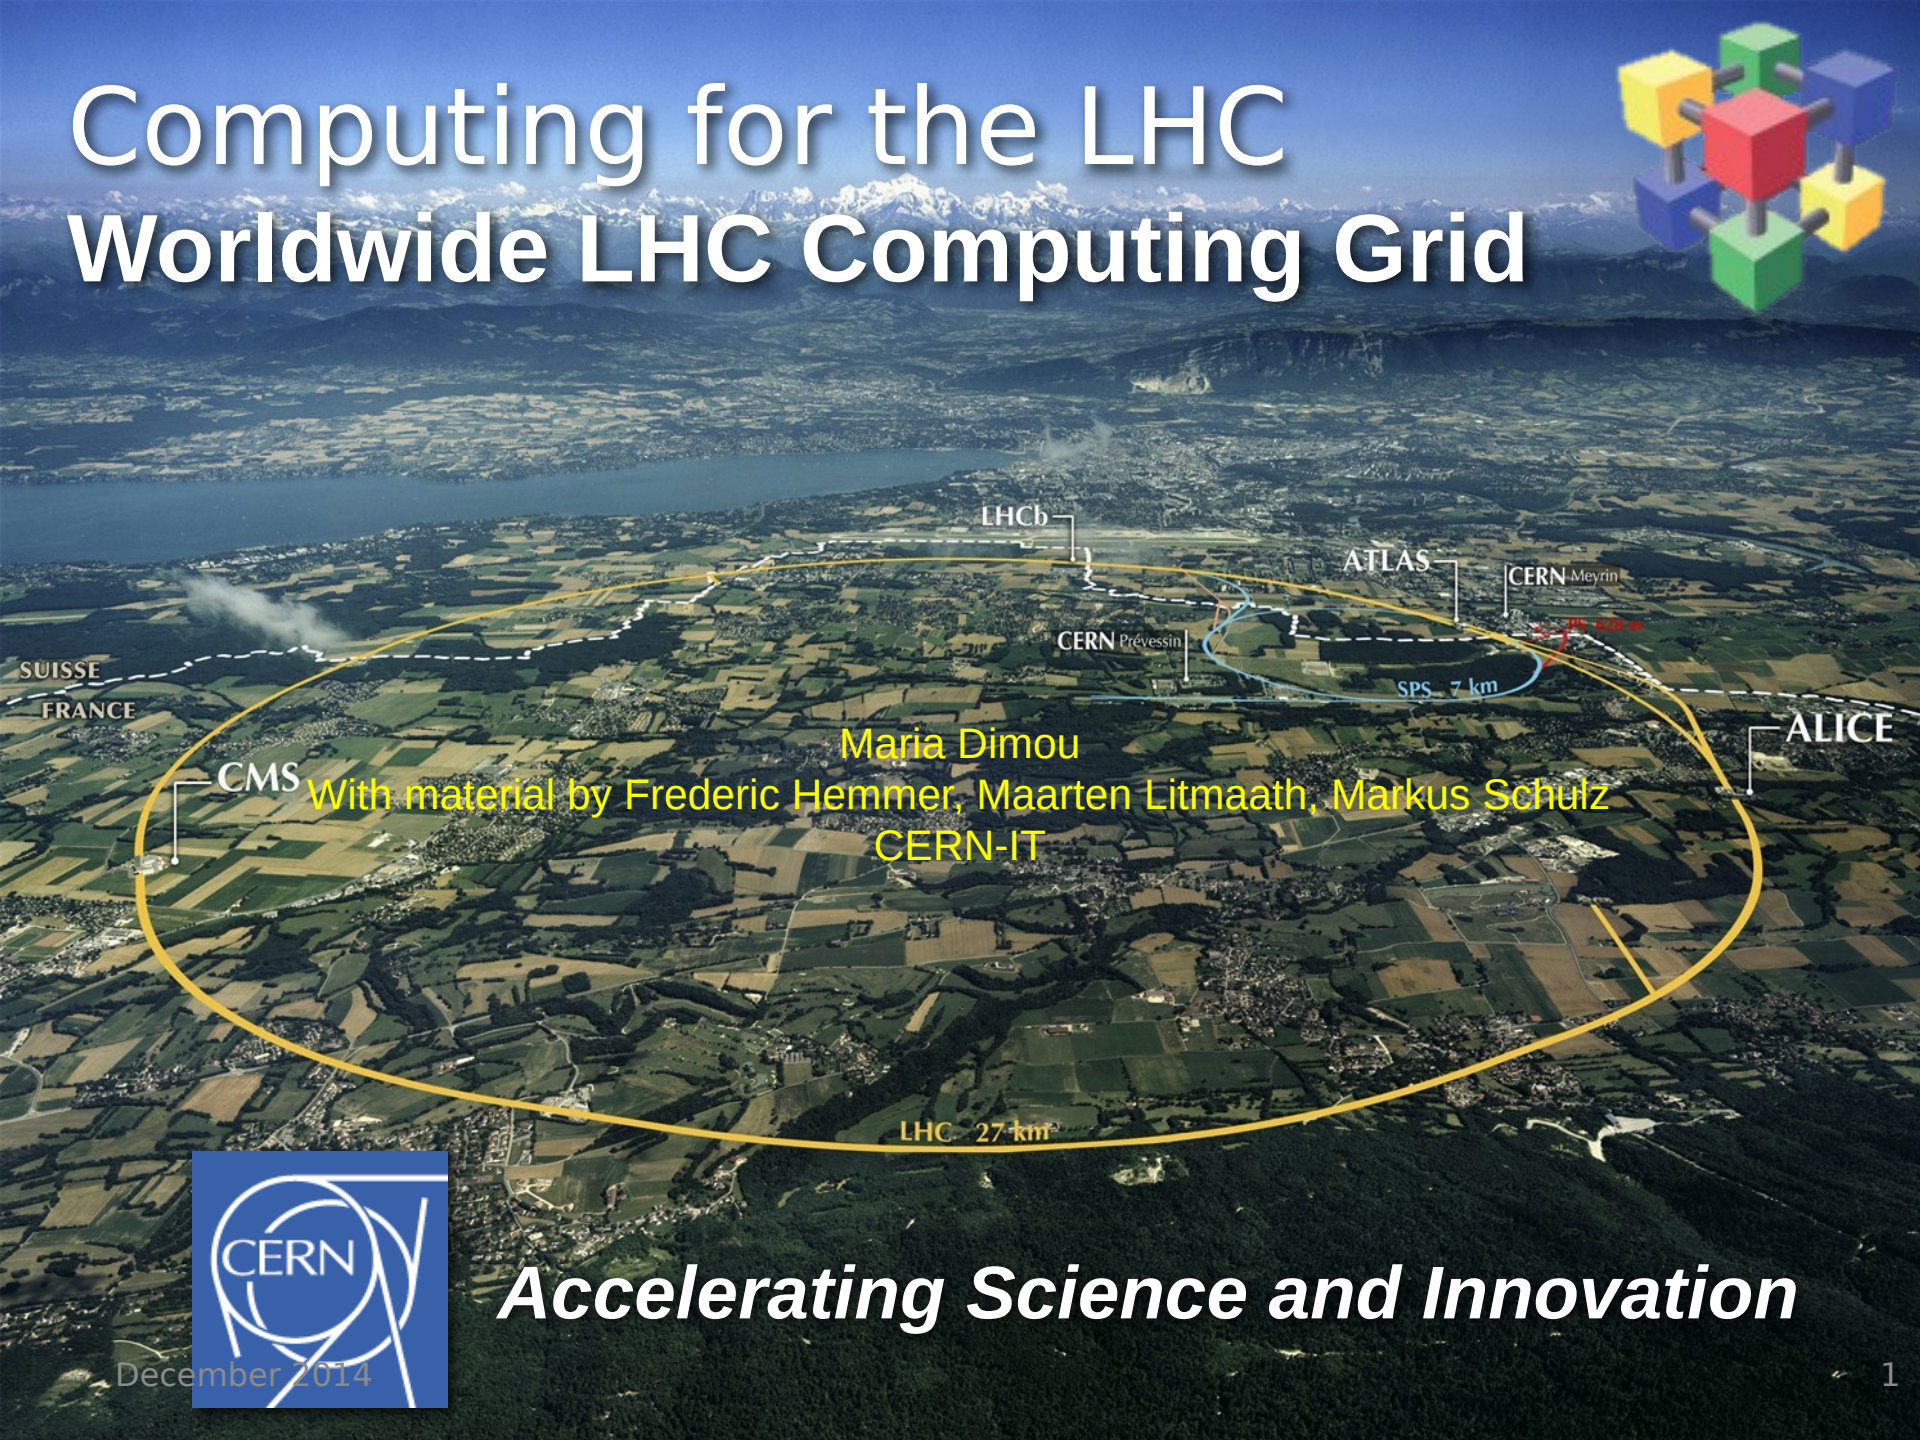

Maria Dimou
With material by Frederic Hemmer, Maarten Litmaath, Markus Schulz
CERN-IT

Accelerating Science and Innovation

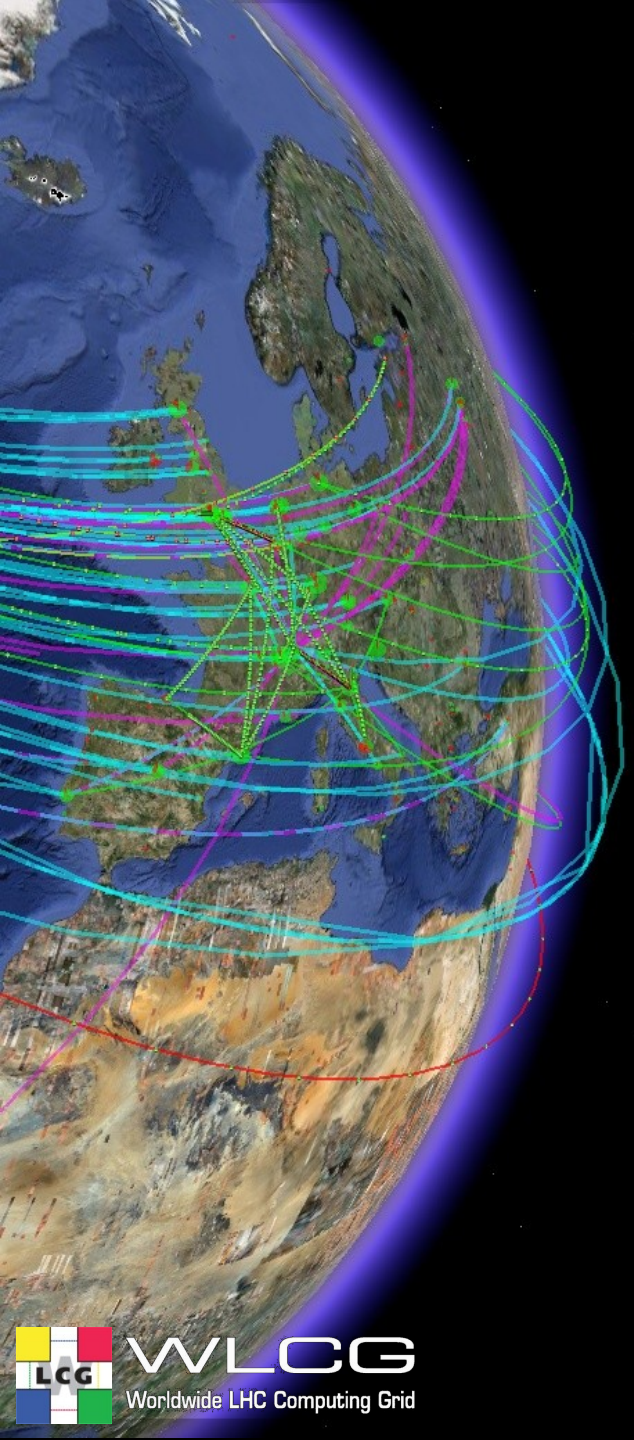

Overview

- The LHC experiment data
- Why do we need the Grid?
- WLCG, what is this?
- What is a Tier centre?
- What is the Grid Middleware?
- How is the Grid operated?
- Summary

WLCG
Worldwide LHC Computing Grid

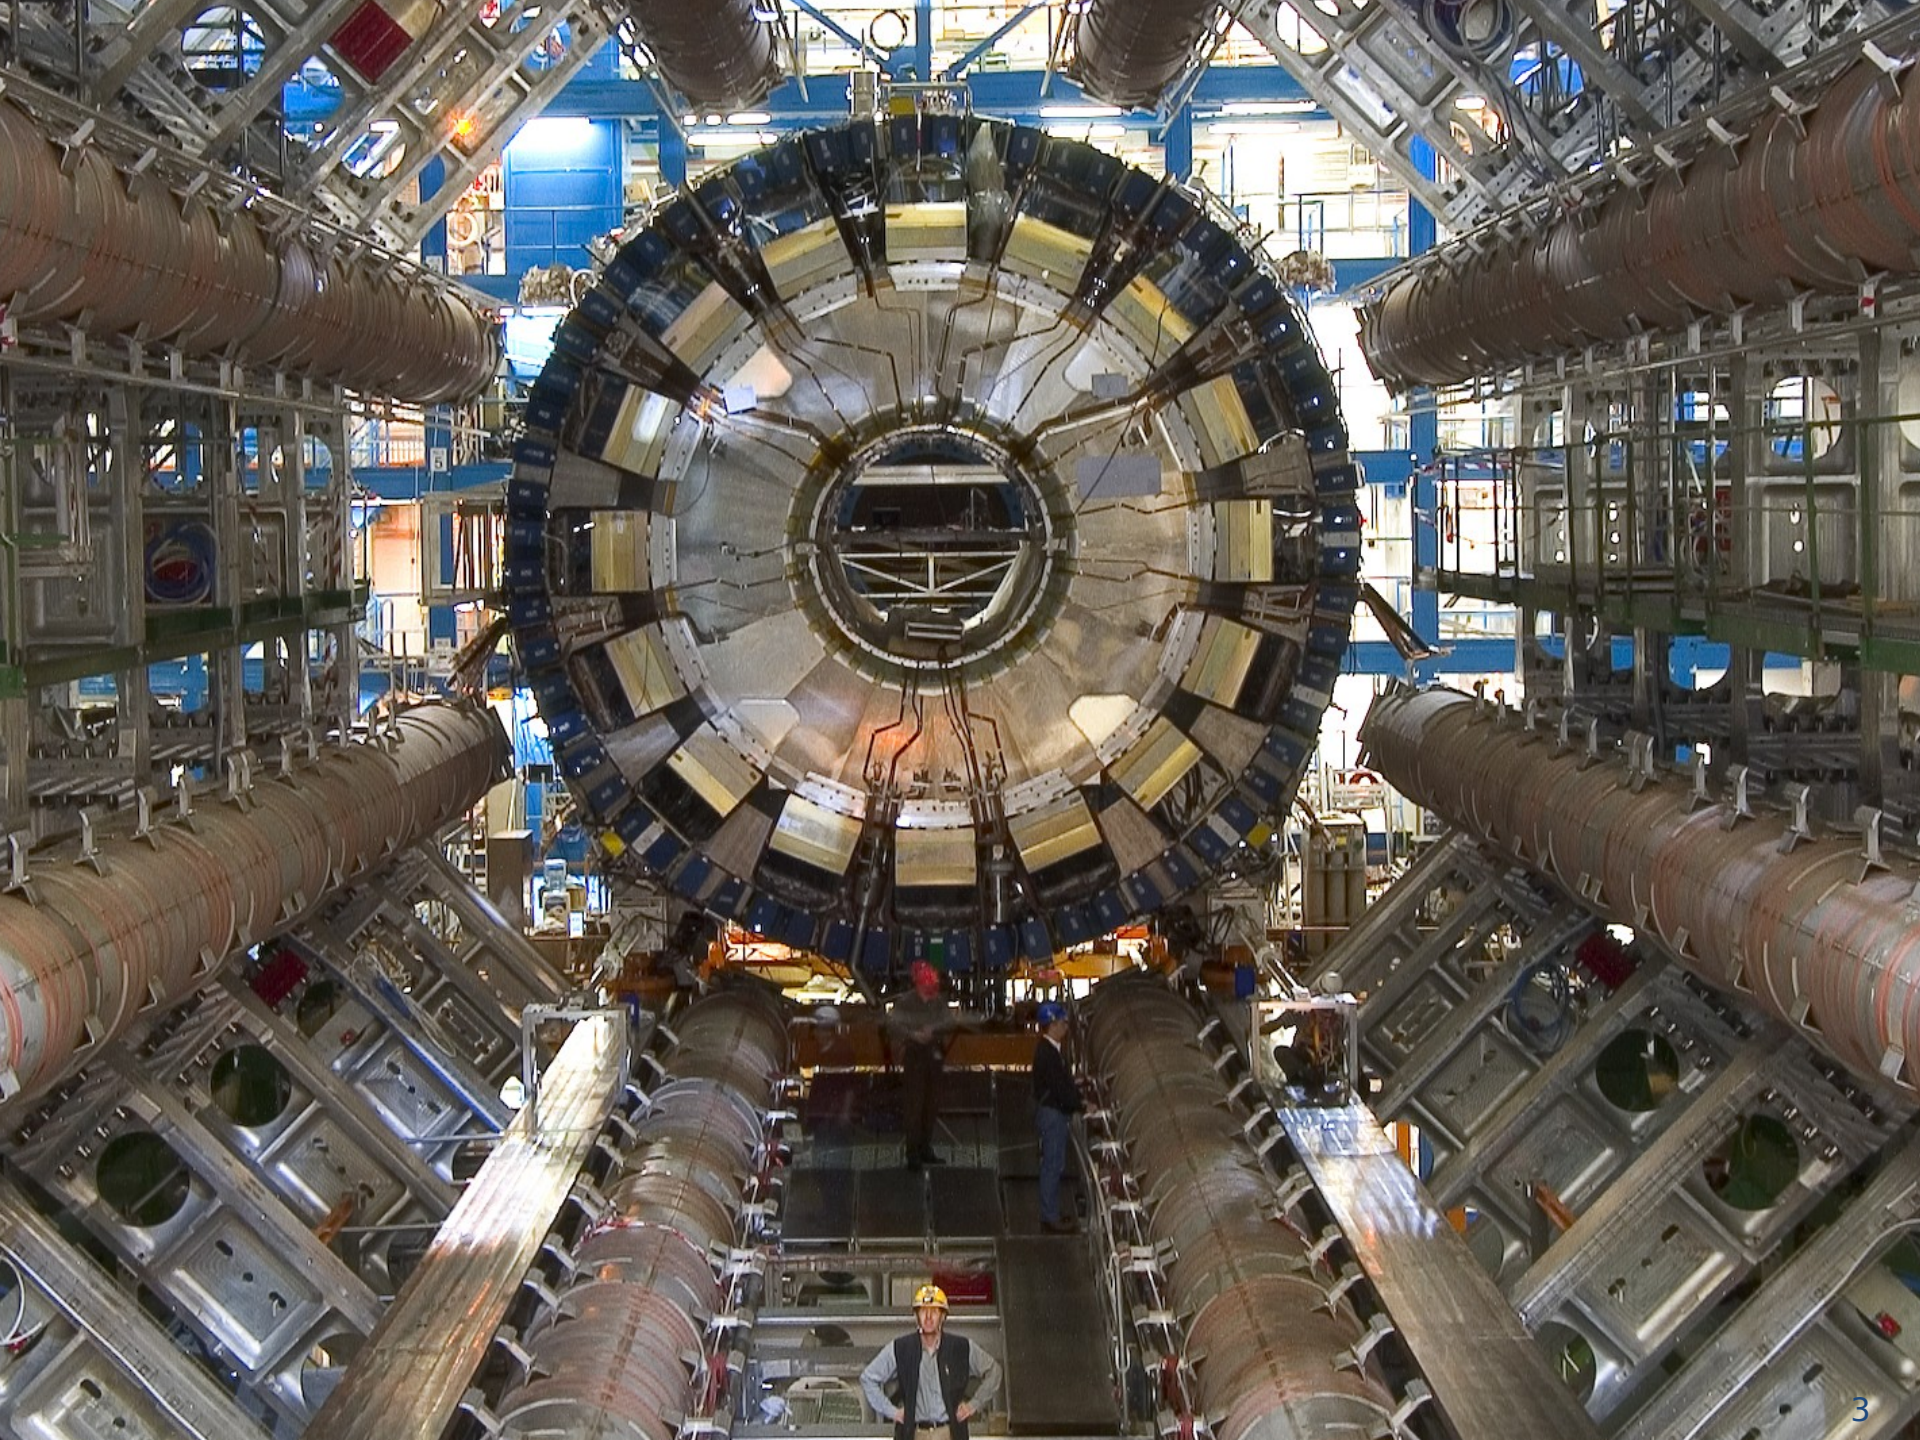

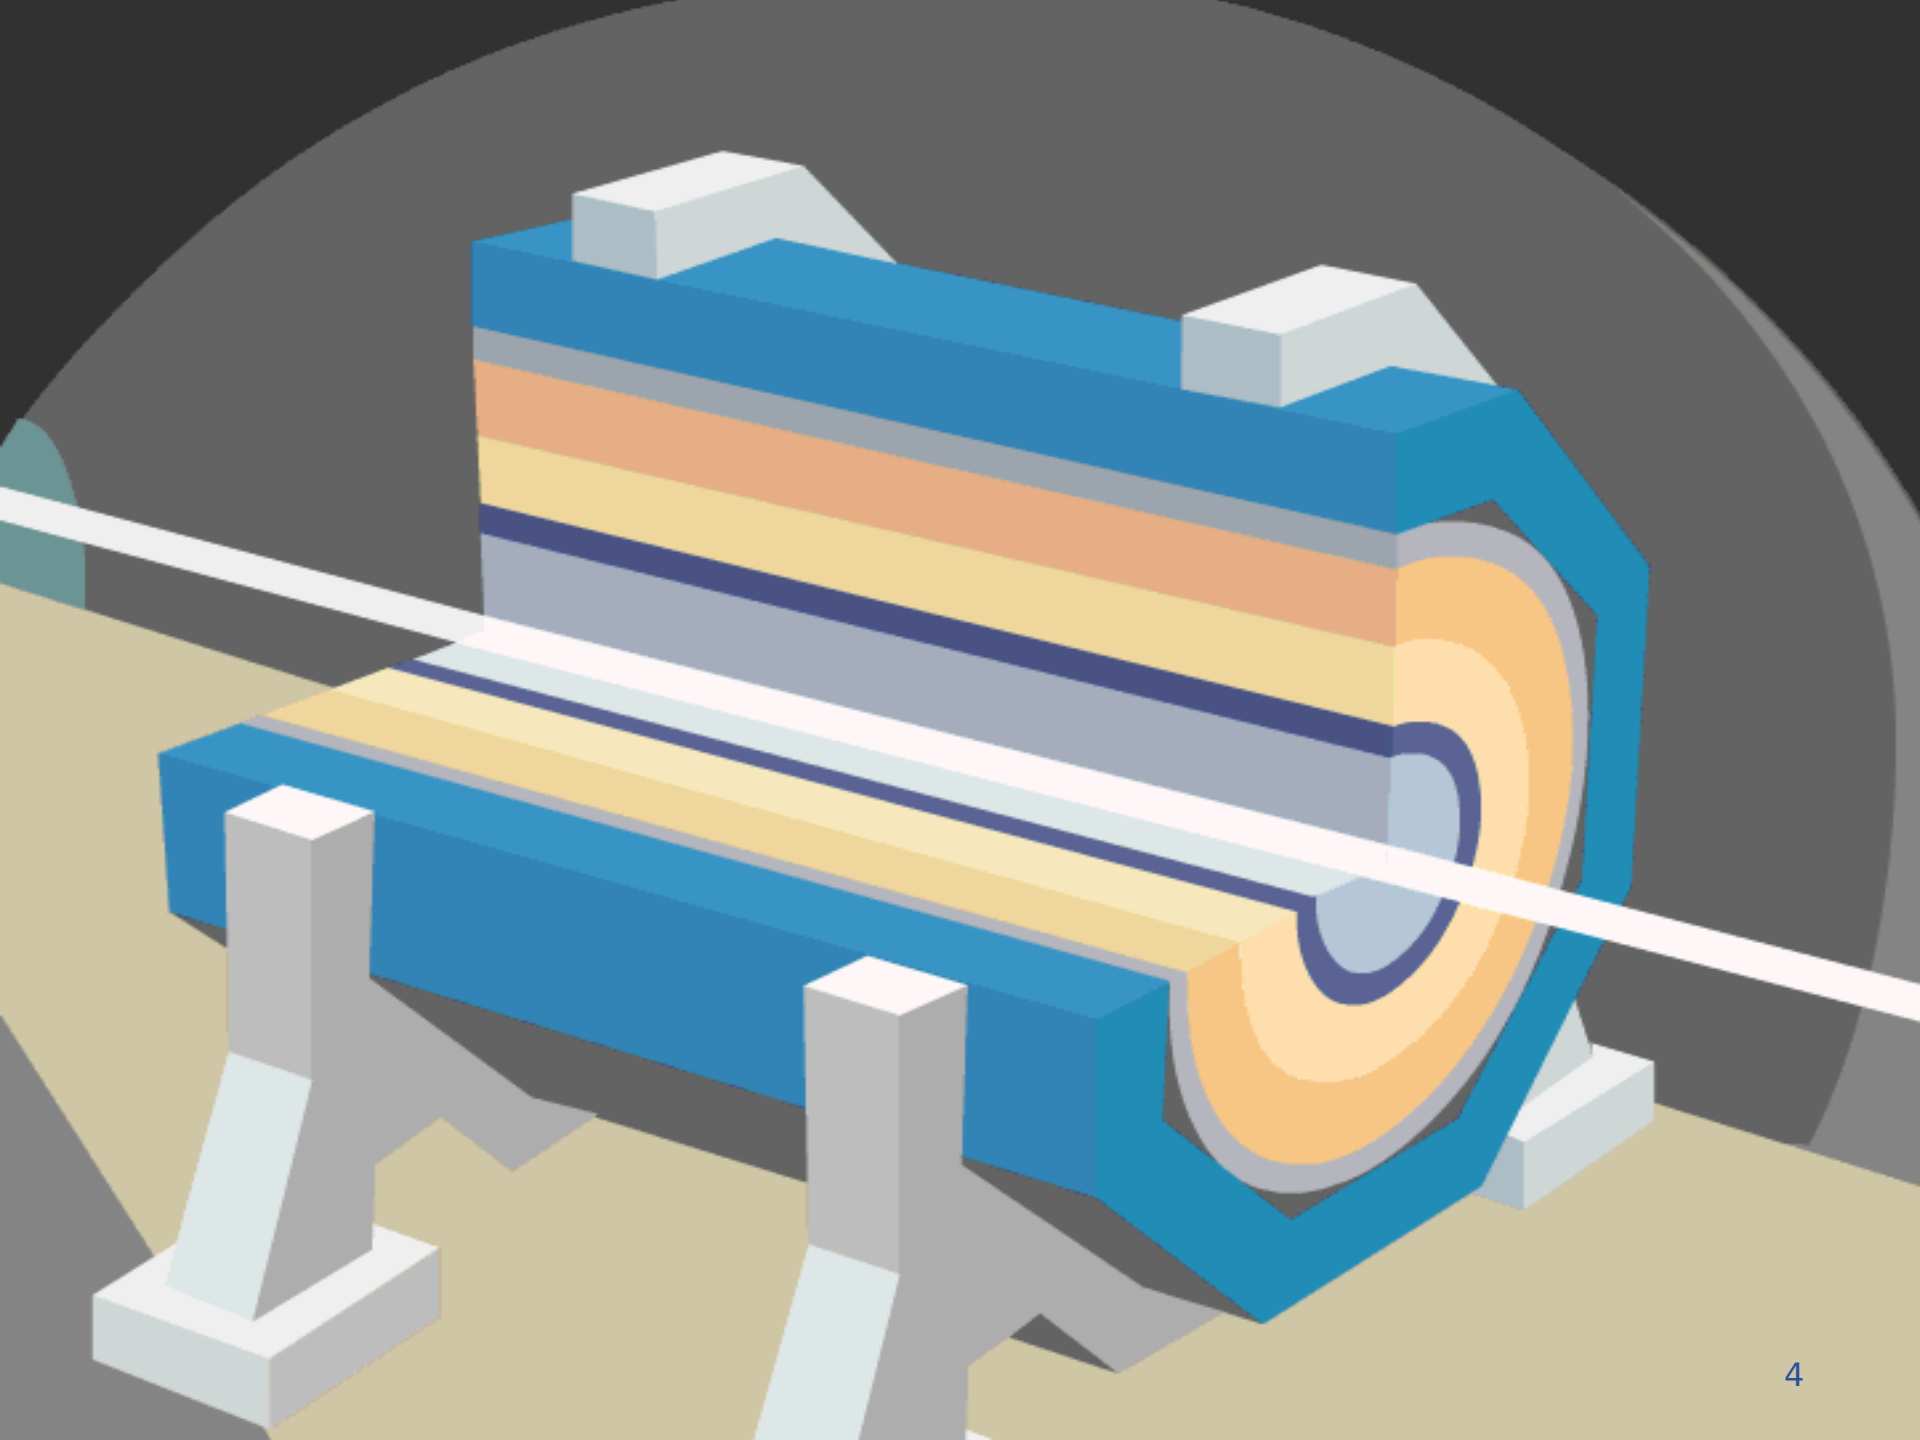

~ 300.000 MB/s
from all sub-detectors

~ 300MB/s
Raw Data

Trigger and data acquisition

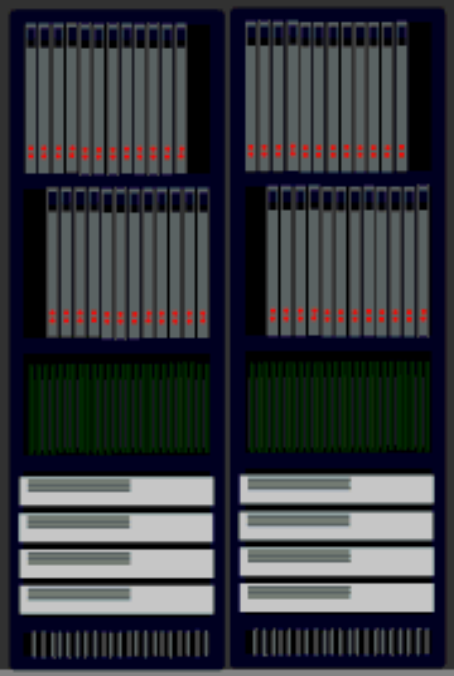

Event filter computer farm

Data transfer from the experiments to the CERN Computer Centre

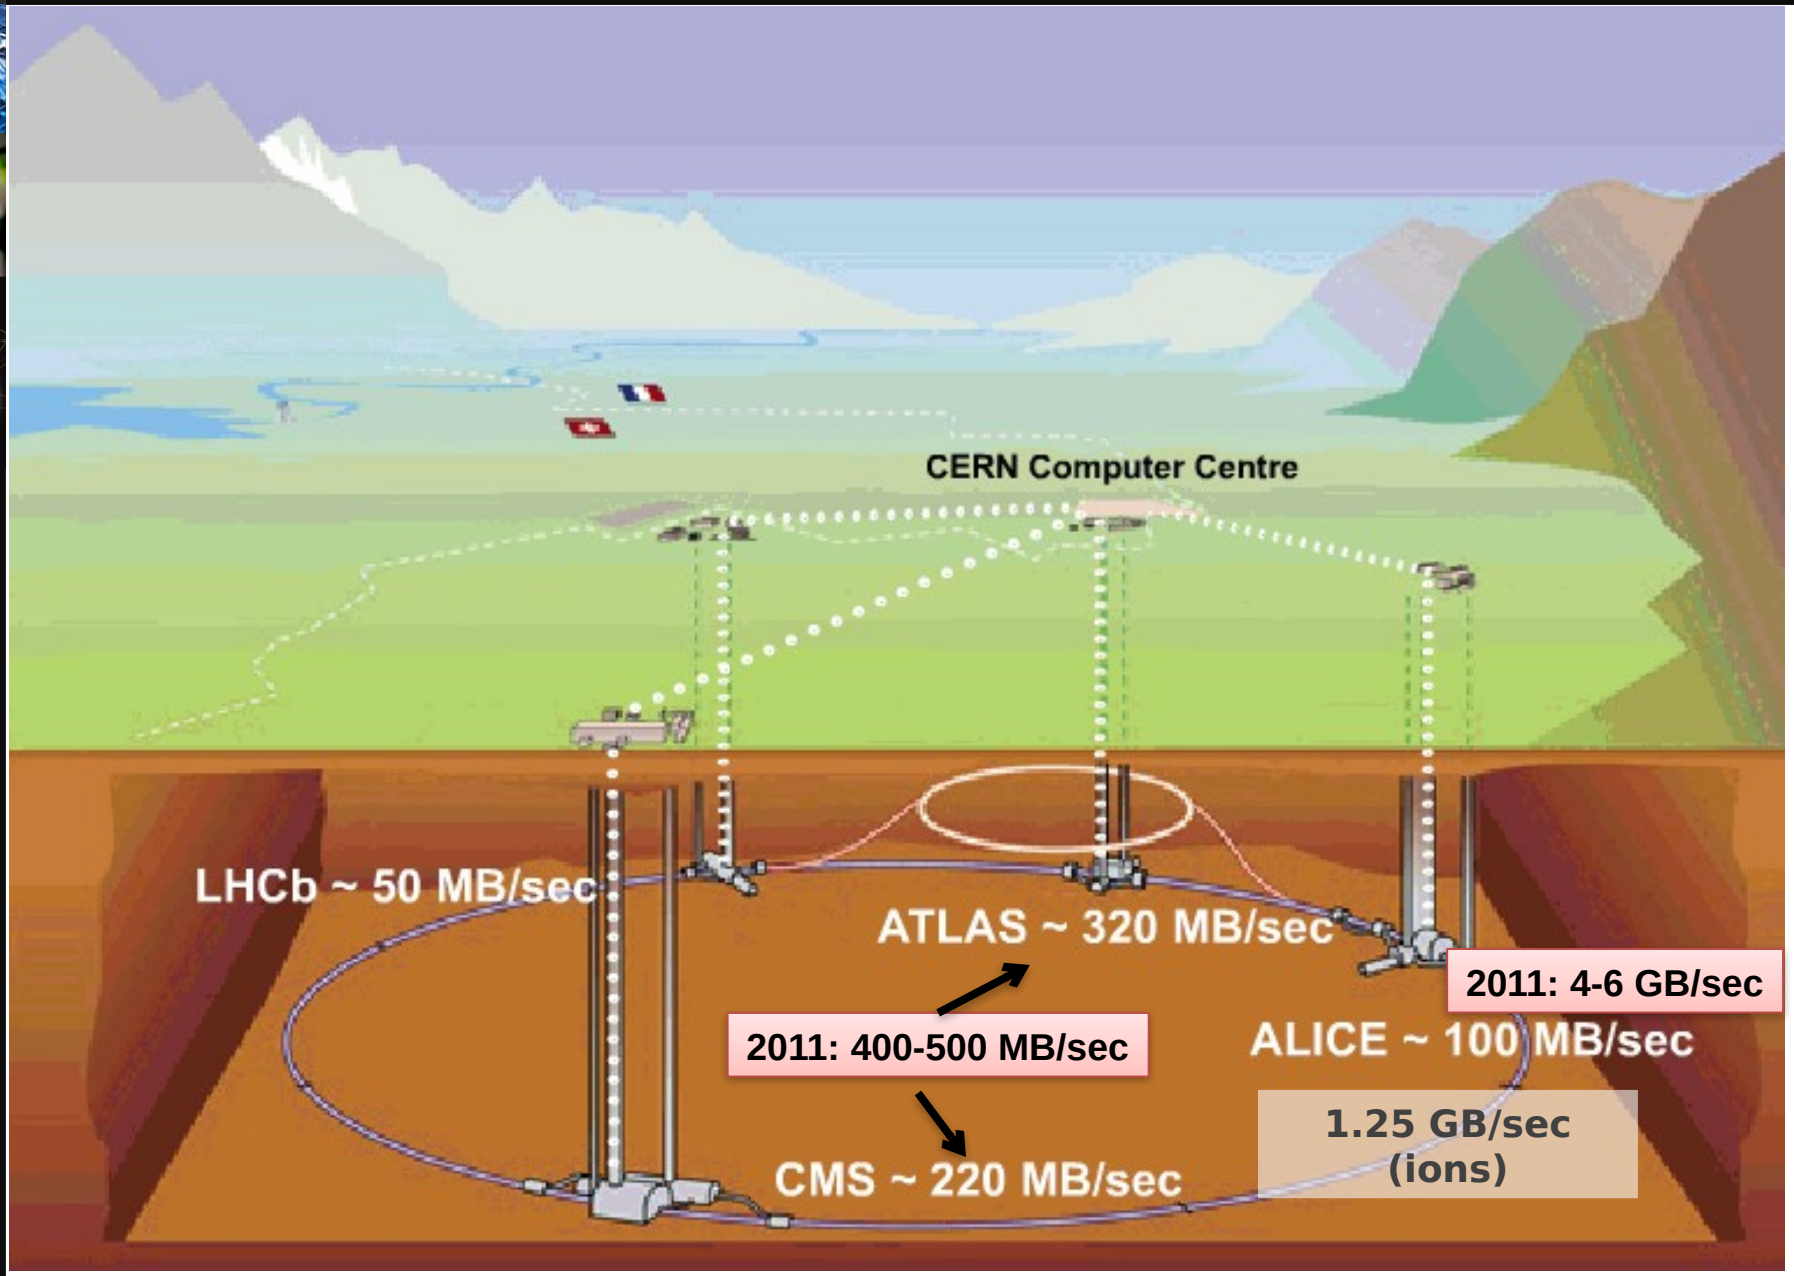

WLCG Worldwide LHC Computing Grid

The CERN Data Centre in Numbers

These numbers will get out of date soon!

Servers	12 k
Processors	21 k
Cores	120 k
Disks	80 k
Total disk space	111 PiB
Memory modules	77 k
Total memory	430 TiB
1-Gbit NICs	21 k
10-Gbit NICs	5 k
Total tape space used	78 PB
Power	3.9 MW

The Tier centres

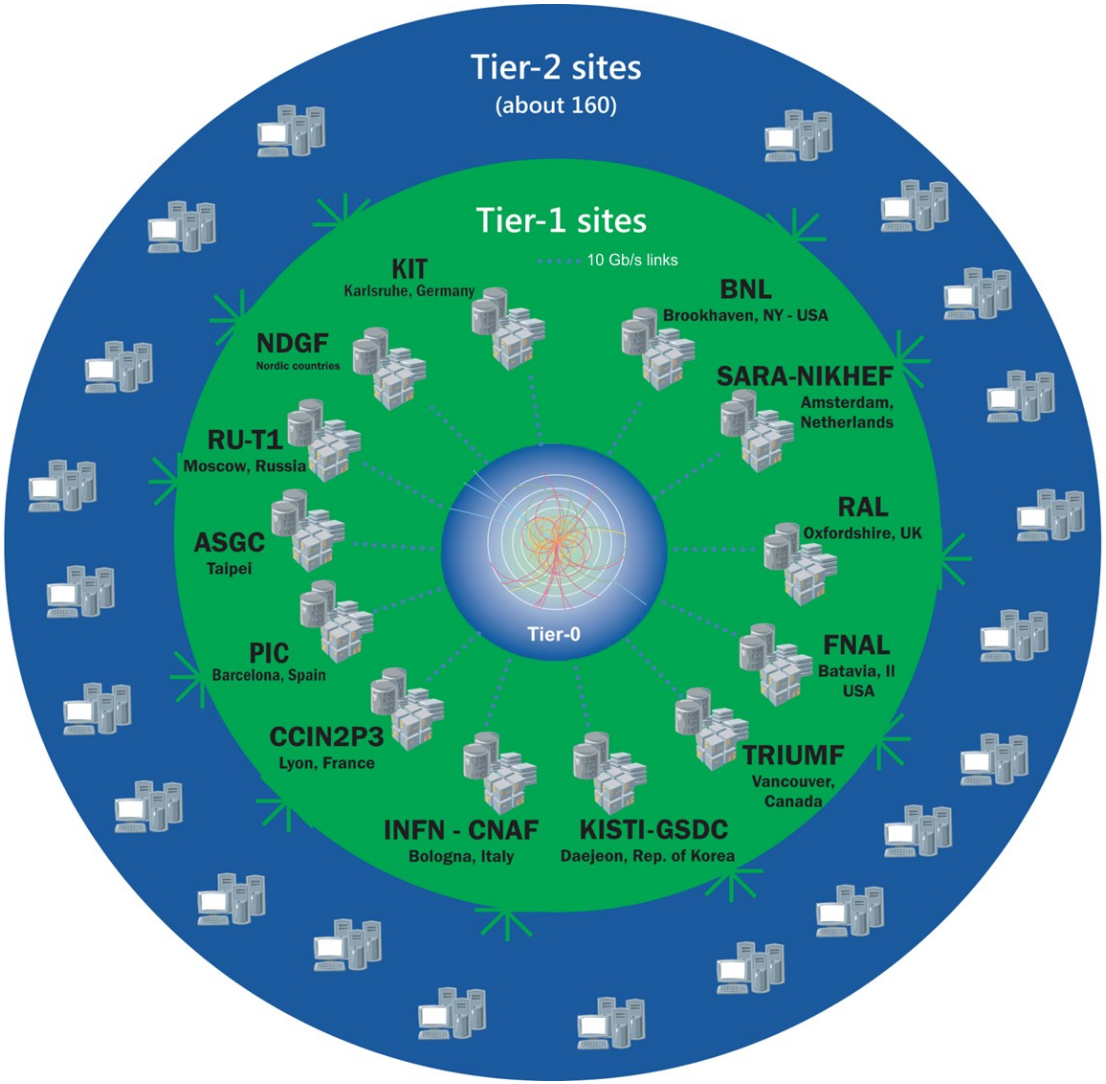

- ~ 160 sites, 35 countries
- 300000 cores
- 200 PB of storage
- 2 Million jobs/day
- 10 Gbps links

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

Tier-1 (13 centres):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (~140 centres):
- Simulation
- End-user analysis

LHC Networking

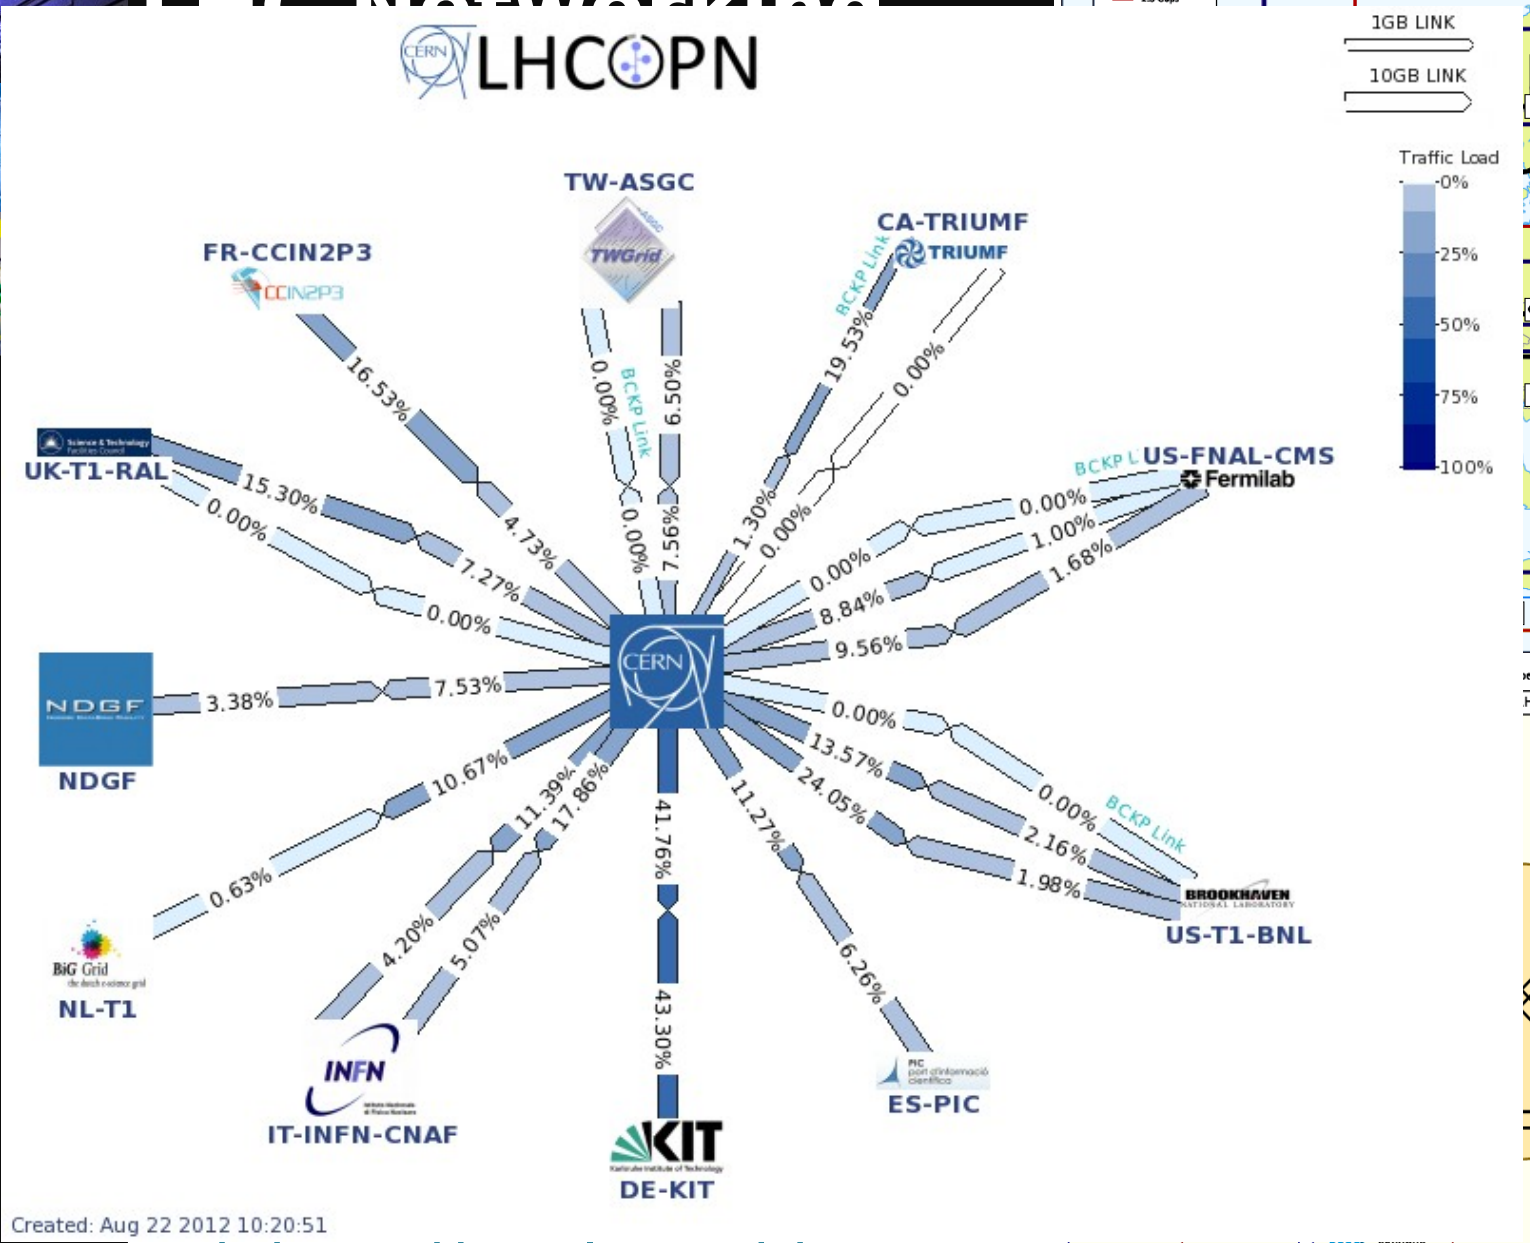

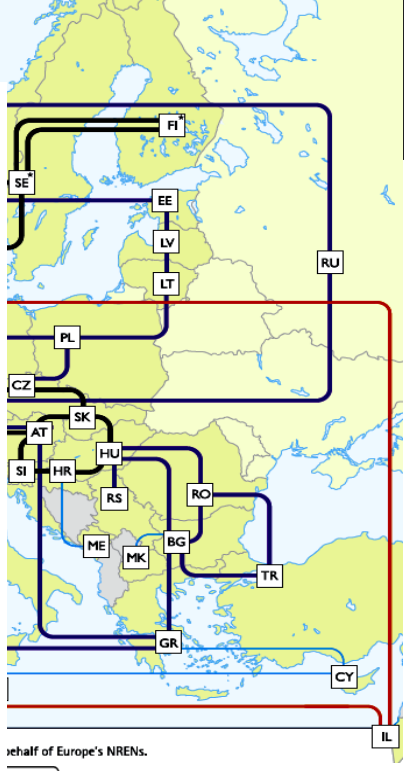

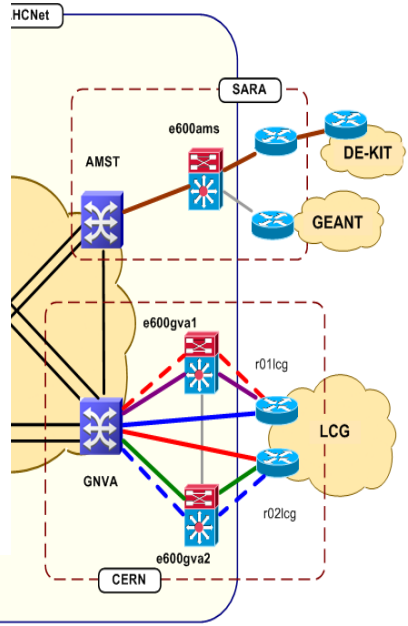

international providers

WLCG Transfers

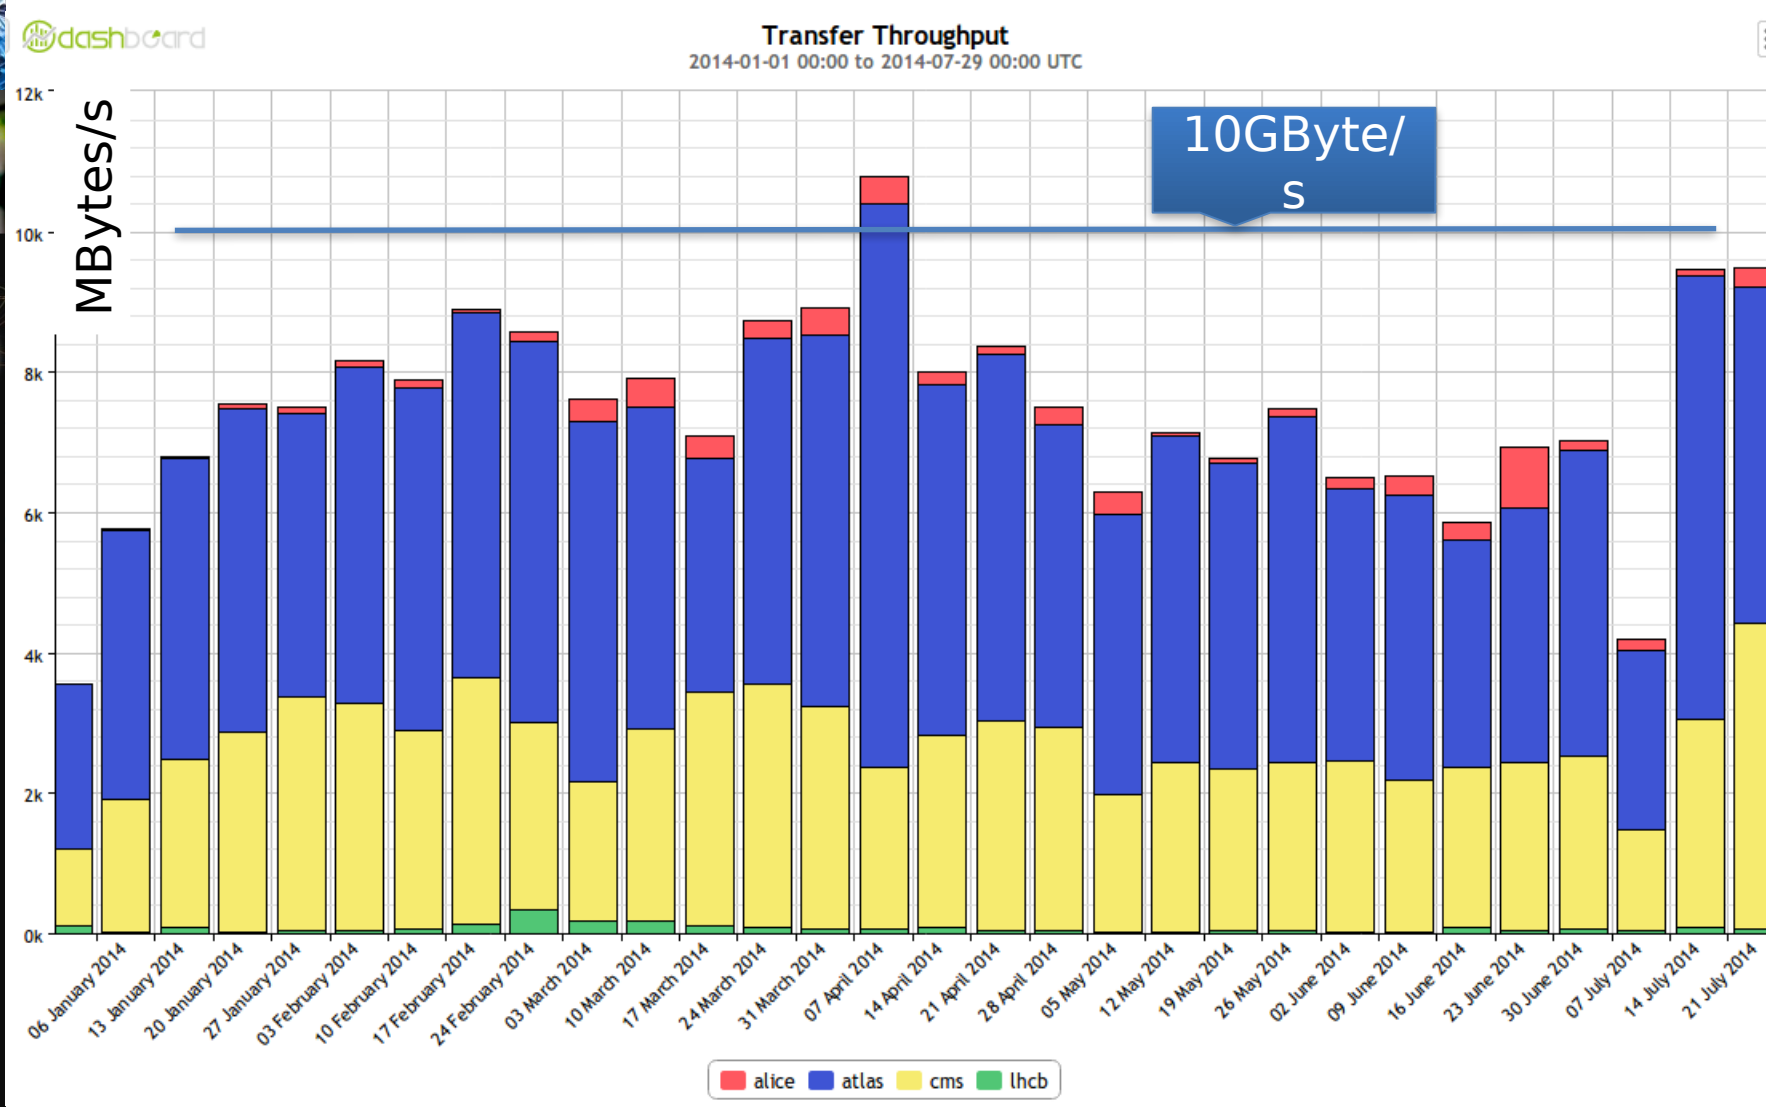

Connectivity over 4 hours

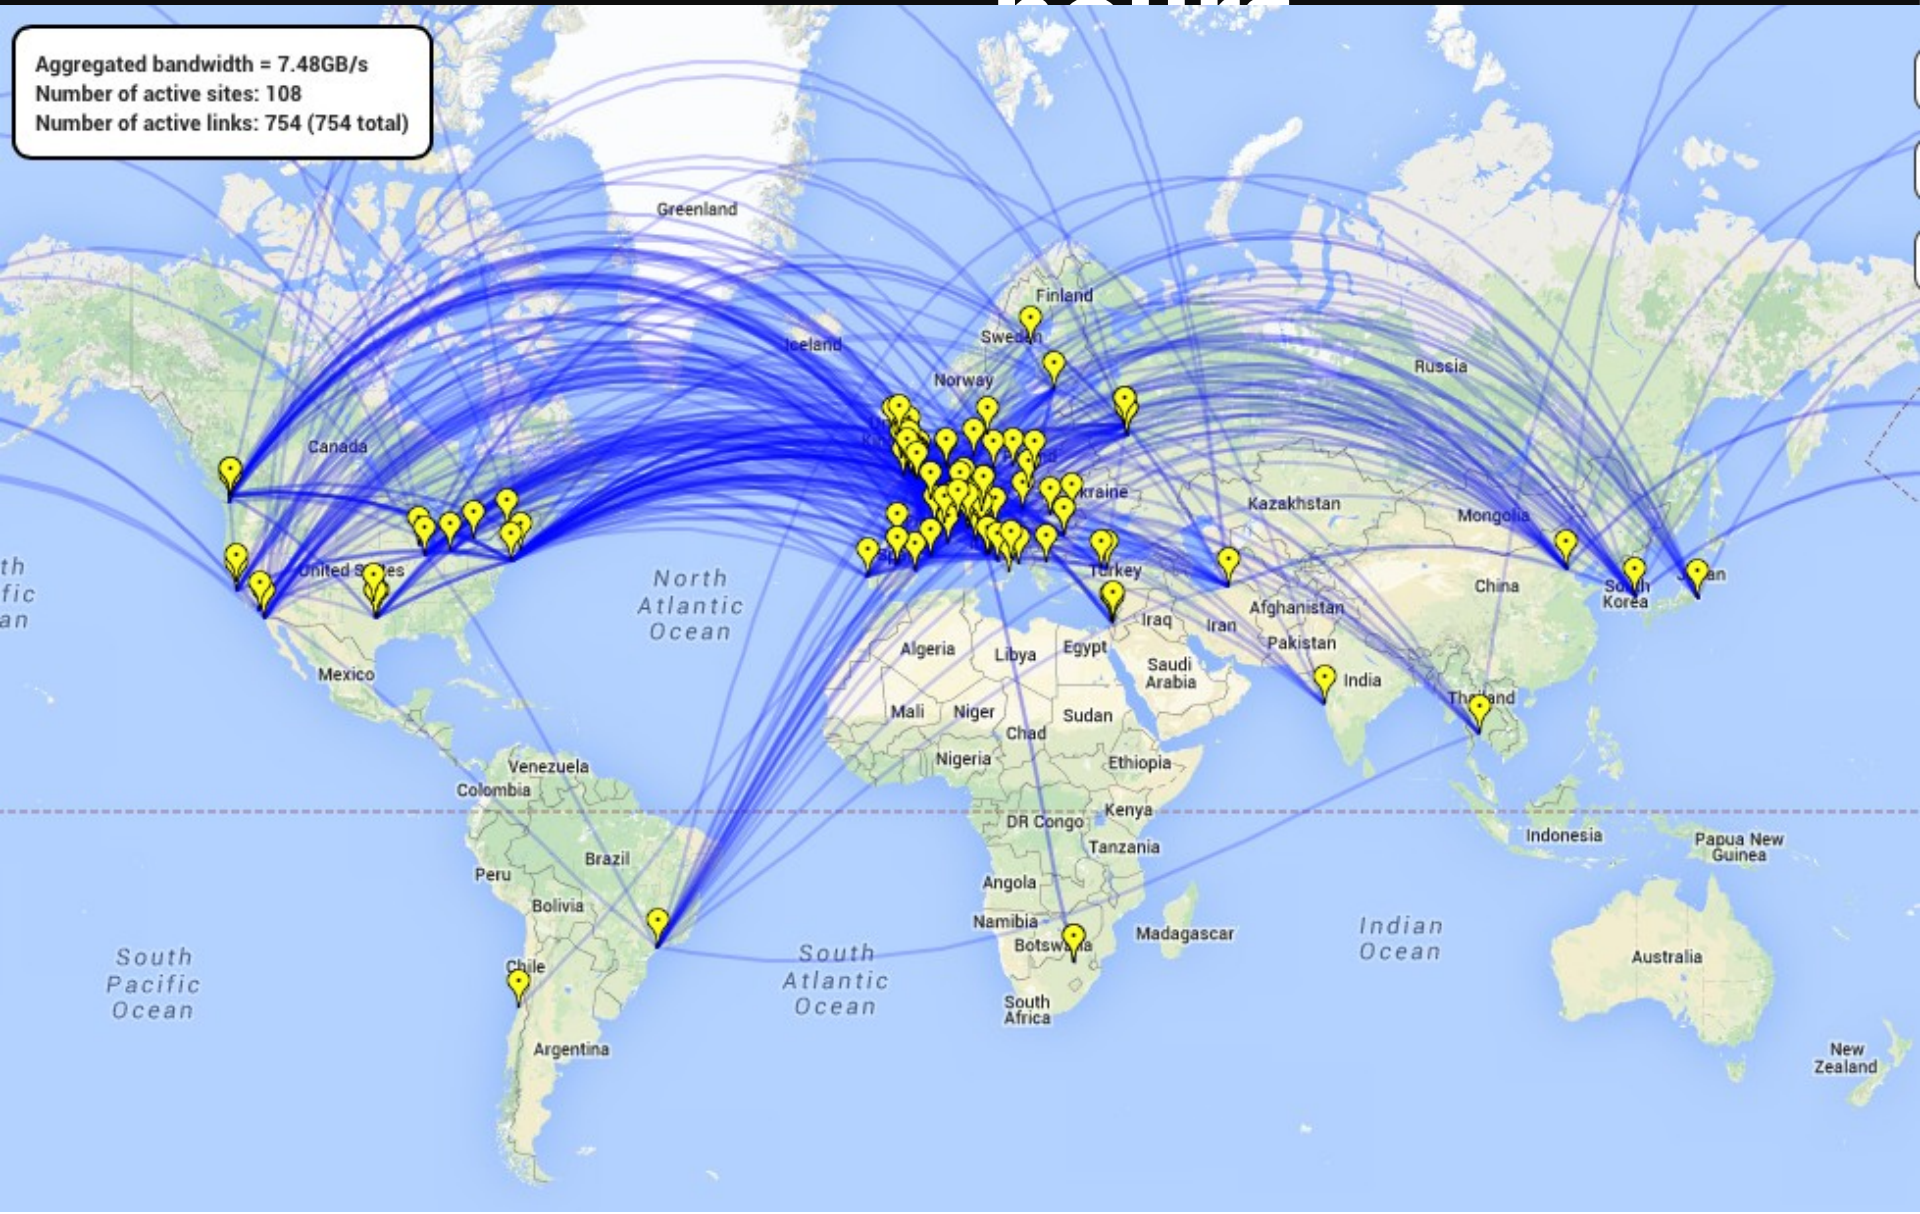

Connectivity over 4 hours

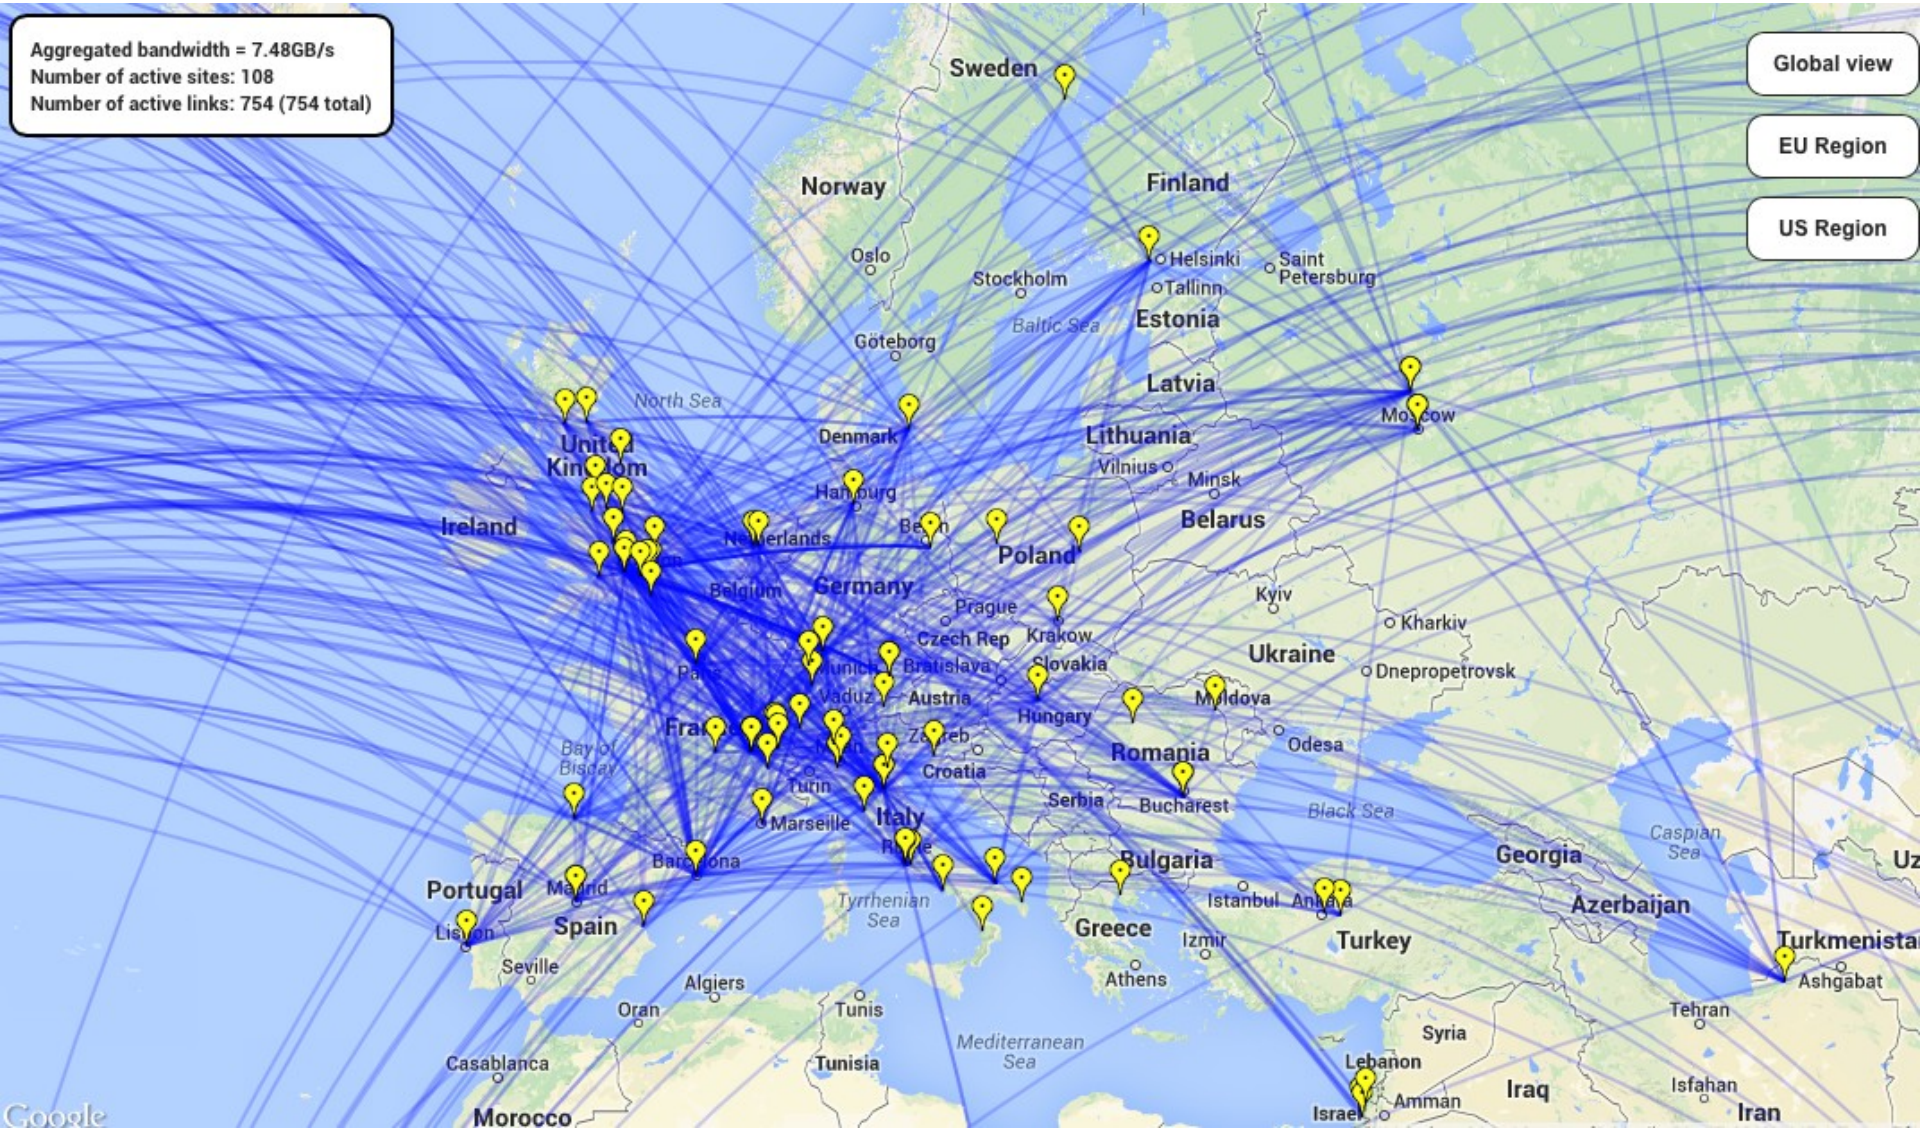

T0/T1/T2 Split

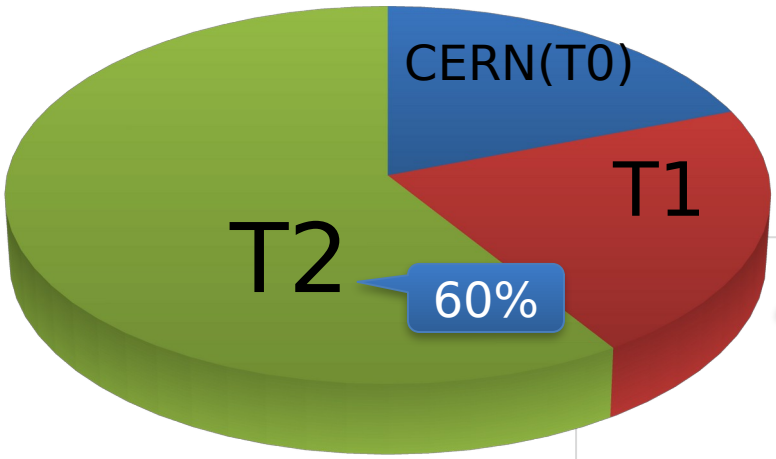

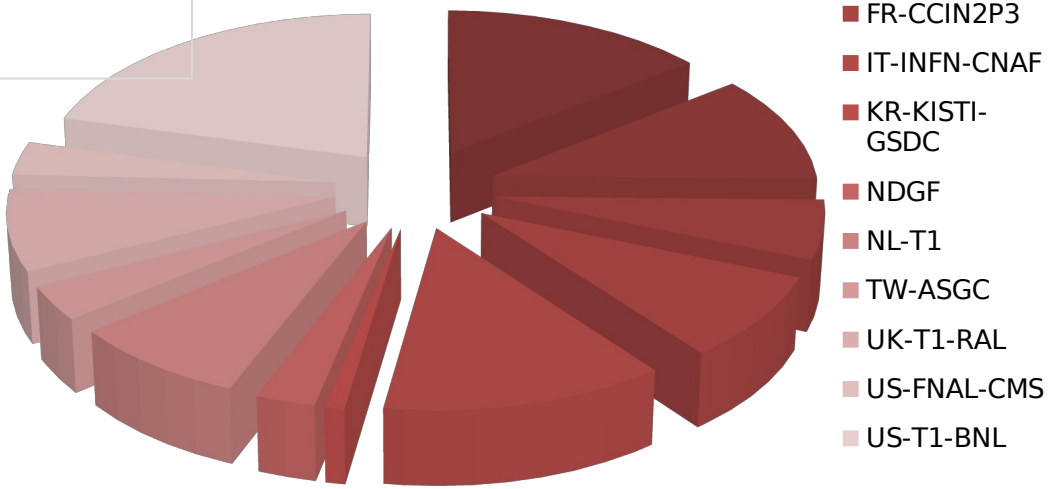

Distribution of jobs in WLCG 2012-2013

T2

T2s are grouped by country or region

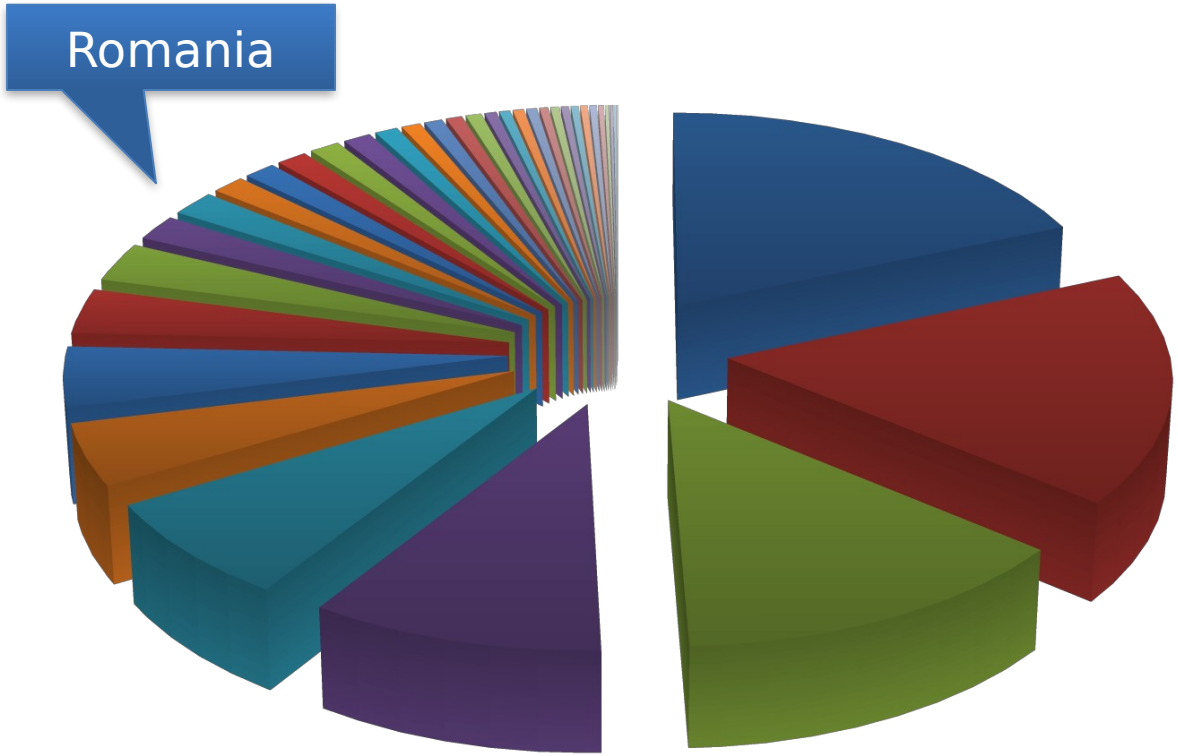

Distribution of jobs in WLCG 2012-2013

Some history of scale…

Date	Collaboration members	Data volume, archive technology
Late 1950's	2-3	Kilobits, notebooks
1960's	10-15	kB, punchcards
1970's	~35	MB, tape
1980's	~100	GB, tape, disk
1990's	700-800	TB, tape, disk
2010's	~3000	PB, tape, disk

For comparison:
1990's: Total LEP data set ~few TB
Would fit on 1 tape today

Today: 1 year of LHC data ~25 PB

WLCG provides about 250 PB of reliable storage

A Very Big Community

Distribution of All CERN Users by Location of Institute on 14 January 2014

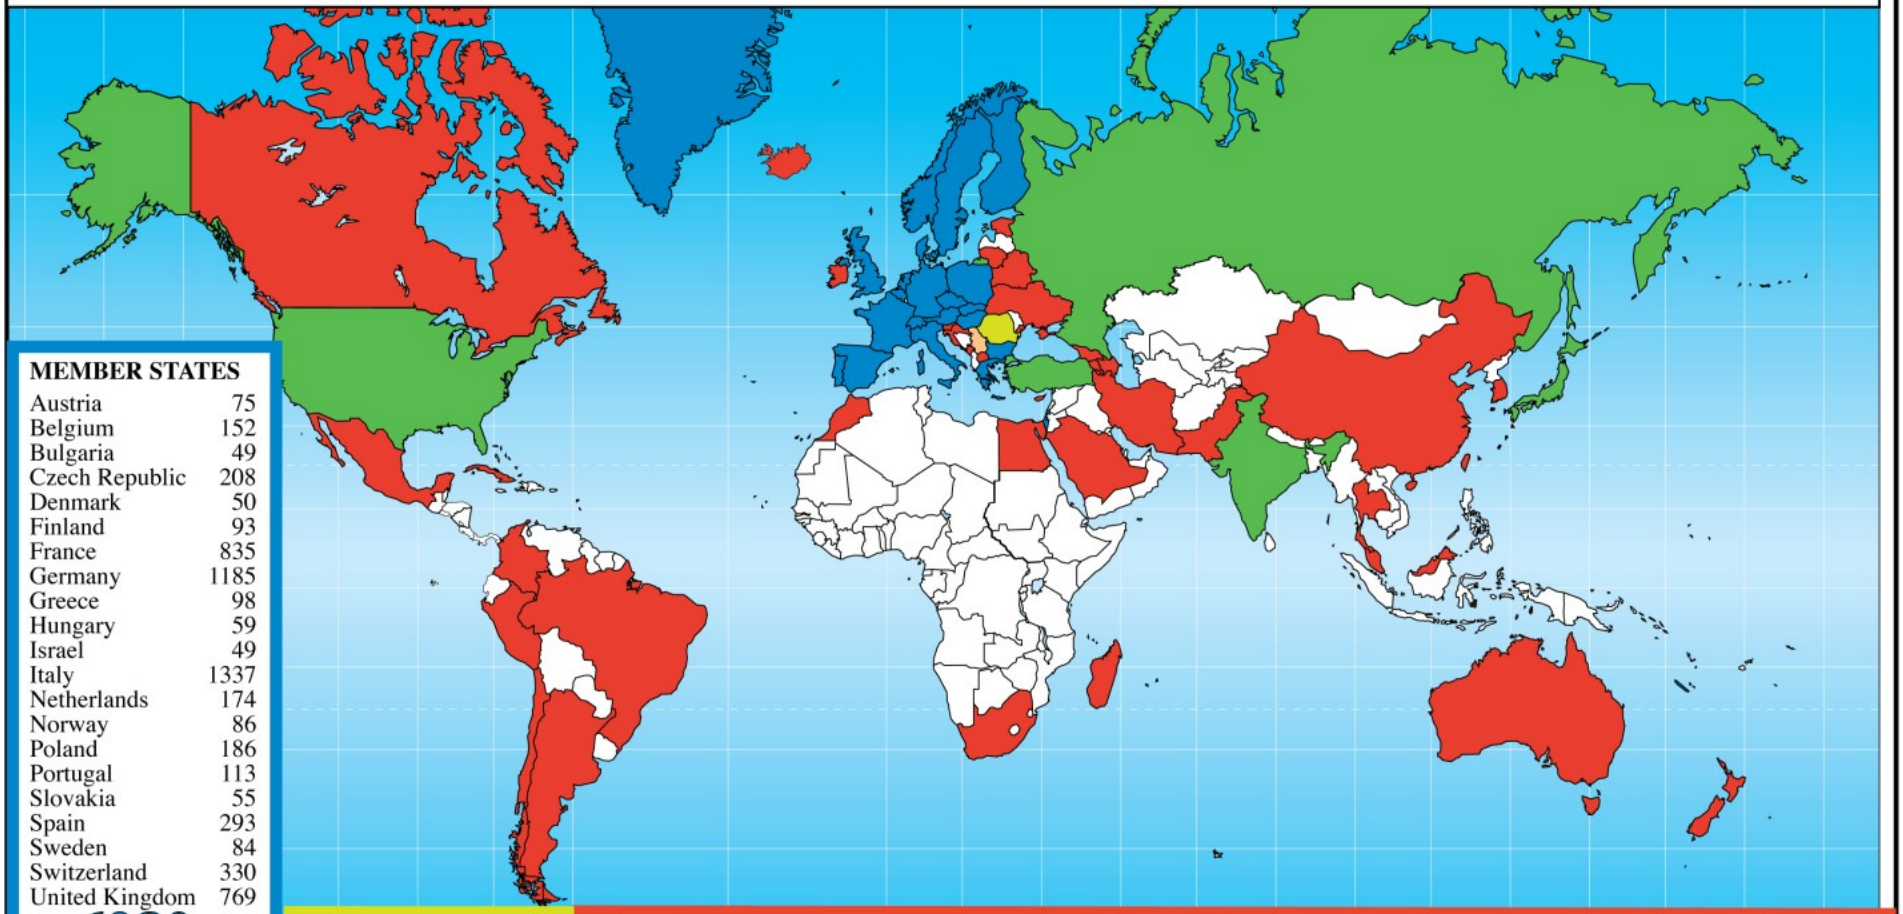

MEMBER STATES	
Austria	75
Belgium	152
Bulgaria	49
Czech Republic	208
Denmark	50
Finland	93
France	835
Germany	1185
Greece	98
Hungary	59
Israel	49
Italy	1337
Netherlands	174
Norway	86
Poland	186
Portugal	113
Slovakia	55
Spain	293
Sweden	84
Switzerland	330
United Kingdom	769

6280

OBSERVERS	
India	153
Japan	217
Russia	890
Turkey	110
USA	1724

3094

CANDIDATE FOR ACCESSION	
Romania	86

ASSOCIATE MEMBER IN THE PRE-STAGE TO MEMBERSHIP	
Serbia	30

OTHERS									
Argentina	13	China	122	Iran	20	Pakistan	18		
Armenia	16	China (Taipei)	71	Ireland	5	Peru	2		
Australia	39	Colombia	10	Korea	105	Saudi Arabia	3		
Azerbaijan	2	Croatia	23	Lithuania	13	Slovenia	25		
Belarus	24	Cuba	3	Madagascar	3	South Africa	32		
Brazil	116	Cyprus	13	Malaysia	8	Thailand	8		
Canada	147	Egypt	18	Mexico	46	T.F.Y.R.O.M.	1		
Chile	8	Estonia	17	Montenegro	1	Ukraine	24		
		Georgia	11	Morocco	6				
		Iceland	4	New Zealand	5				

982

December 2014

What is cloud computing?

- Transparent use of generic computing resources off-site
  - Dynamically pro
  - Metered

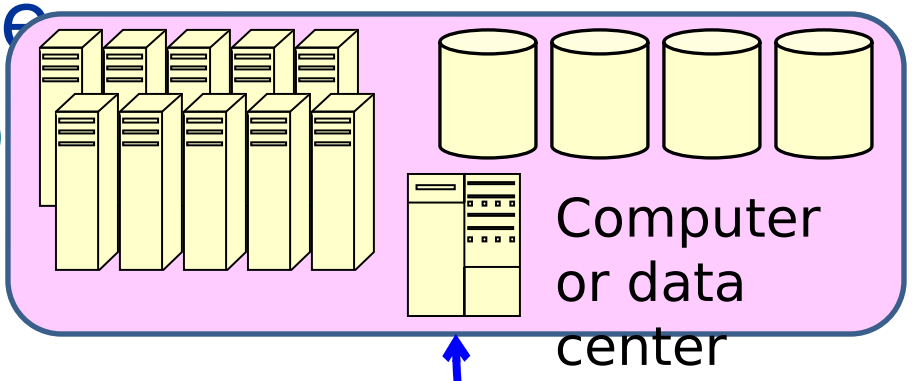

- Neutral to applications
  - Rent-a-center

- Amazon EC2, S3
- Oracle, Sun
- Google
- IBM
- HP
- Microsoft
- …we investigate but we do grid computing.

What is grid computing?

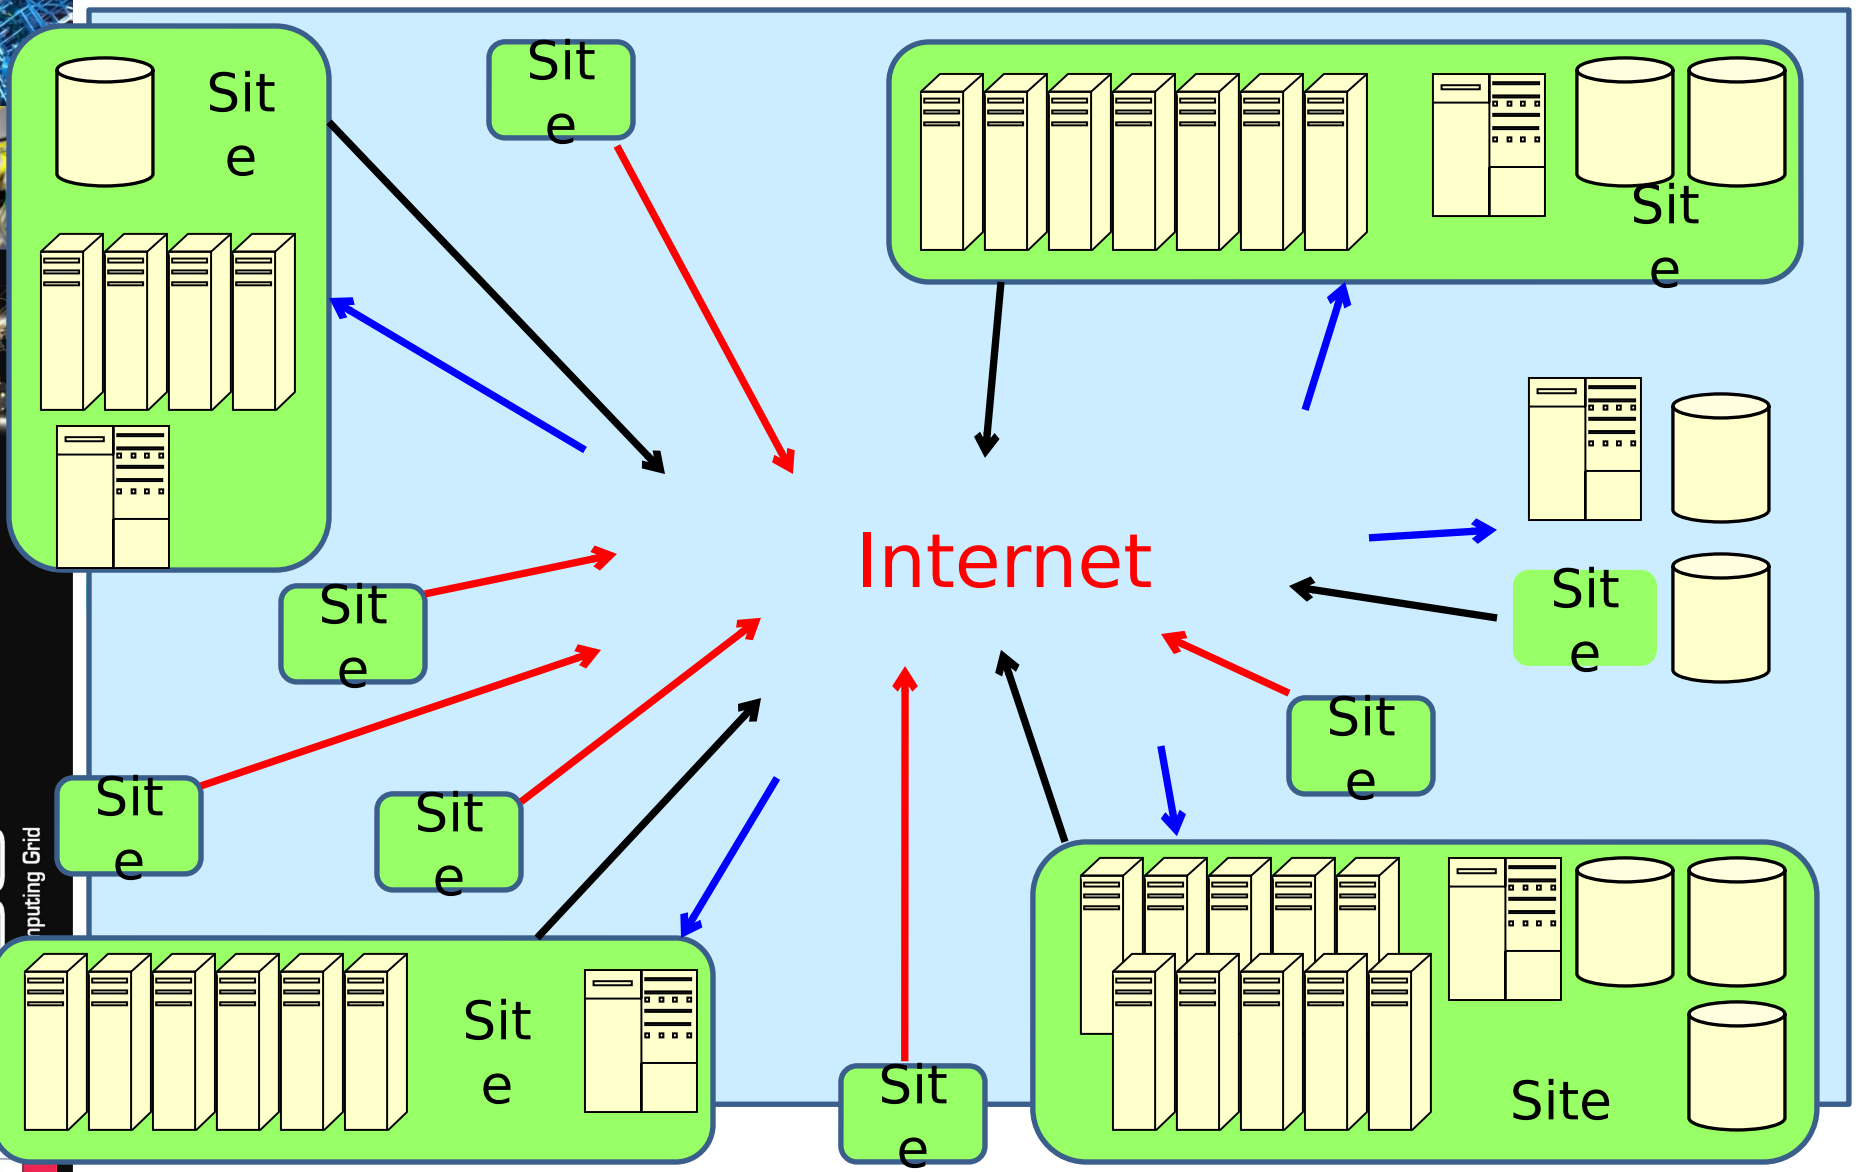

Grid projects' history

- 1999 - MONARC project
  - First LHC computing architecture – hierarchical distributed model, focus on network control
- 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 – 2013 EGI and EMI
  - Sustainable infrastructures based on National Grid Infrastructures
  - Decoupling of middleware development and infrastructure
  - Merging middleware stacks in Europe

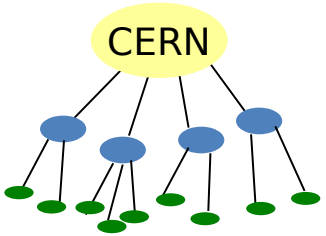

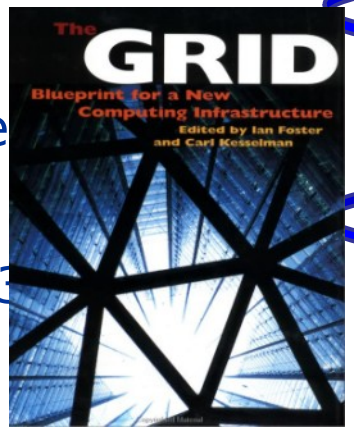

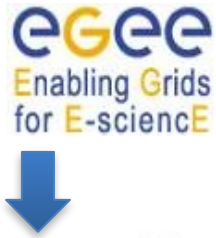

How does a grid work?

- Middleware is key
- It makes multiple computer centres look like a single system
- Not easy though
- Sites have different infrastructures and other user communities

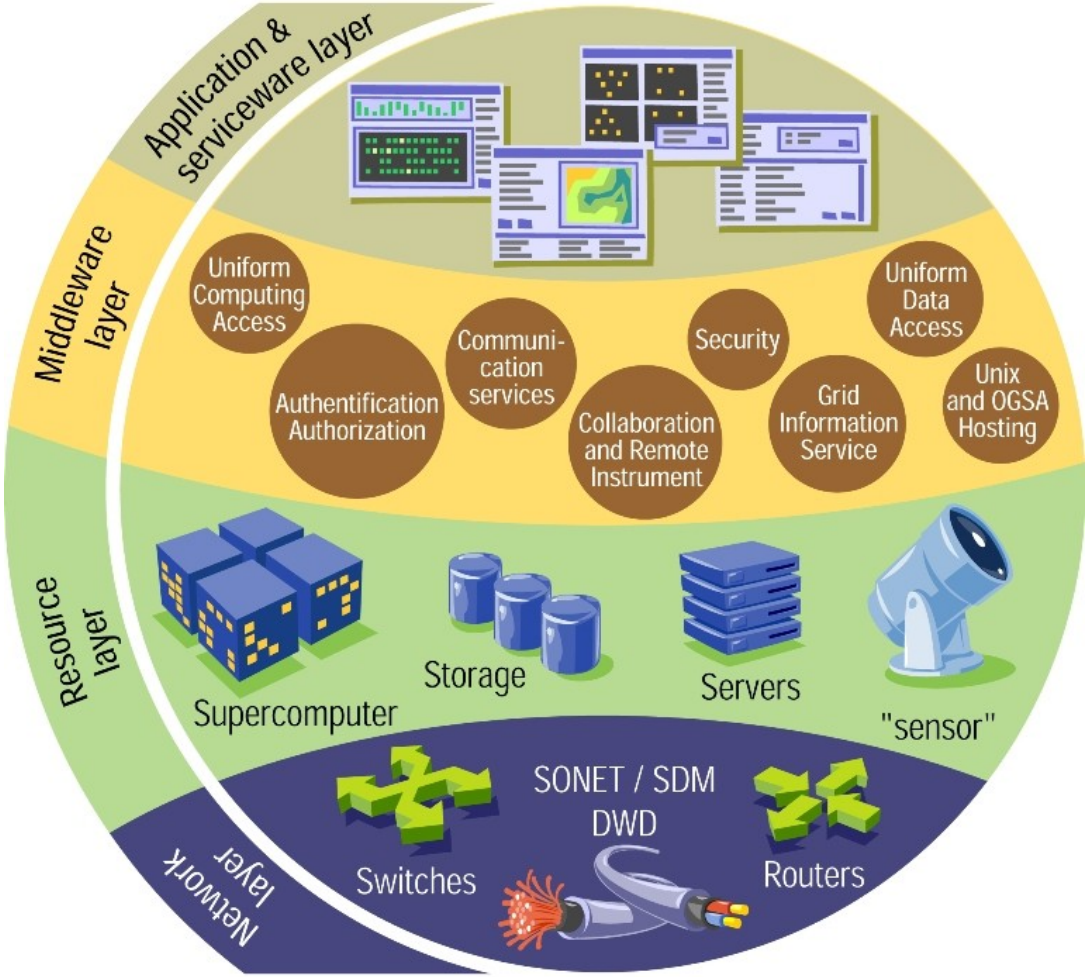

What is Grid Middleware


Data Management Services

- Storage Element
- File Catalogue Service
- Grid file access tools
- File Transfer Service
- GridFTP service
- Database and DB Replication Services
- POOL Object Persistency Servi

Security Services

- Certificate Management Service
- VO Membership Service
- Authentication Service
- Authorization Servic

Job Management Services

- Compute E ent
- Wo ement ce
- Agent Service
- Application Software Install Service

Inform Services

- Info System
- Messaging Service
- ounti ng Service
- Site Availability Monitor
- Monitoring tools: experiment dashboards; site monitoring

All those come in different flavours

Experiments invested considerable effort into integrating their software with grid services; and hiding complexity from users

How is the Grid operated?

- Not all is provided by WLCG directly
- WLCG links the services provided by the underlying national or cross-country infrastructures. Examples:
  - EGI (European Grid Infrastructure) links computer centres across Europe supporting scientific research.
  - NDGF (Nordic Data Grid Facility) is a common e-Science infrastructure provided by Denmark, Finland, Norway, Sweden and Iceland for scientific computing and data storage.
  - OSG (Open Science Grid) is a consortium of researchers, service & resource providers across the USA.
- Indispensable paraphernalia are operated for:
  - Support
  - Accounting
  - Monitoring

Collaboration & Education offsprings

- CERN openlab
  - Intel, Huawei, Oracle, Rackspace, Siemens
  - http://cern.ch/openlab

- CERN School of Computing
  - http://cern.ch/csc

- UNOSAT
  - http://cern.ch/un

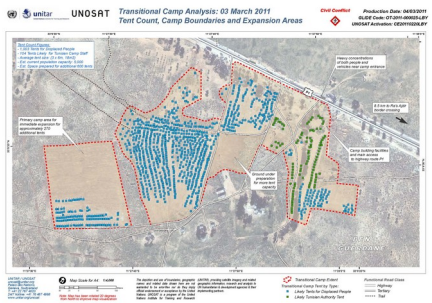

- Citizen Cyber Science Collaboration
  - Involving the General Public

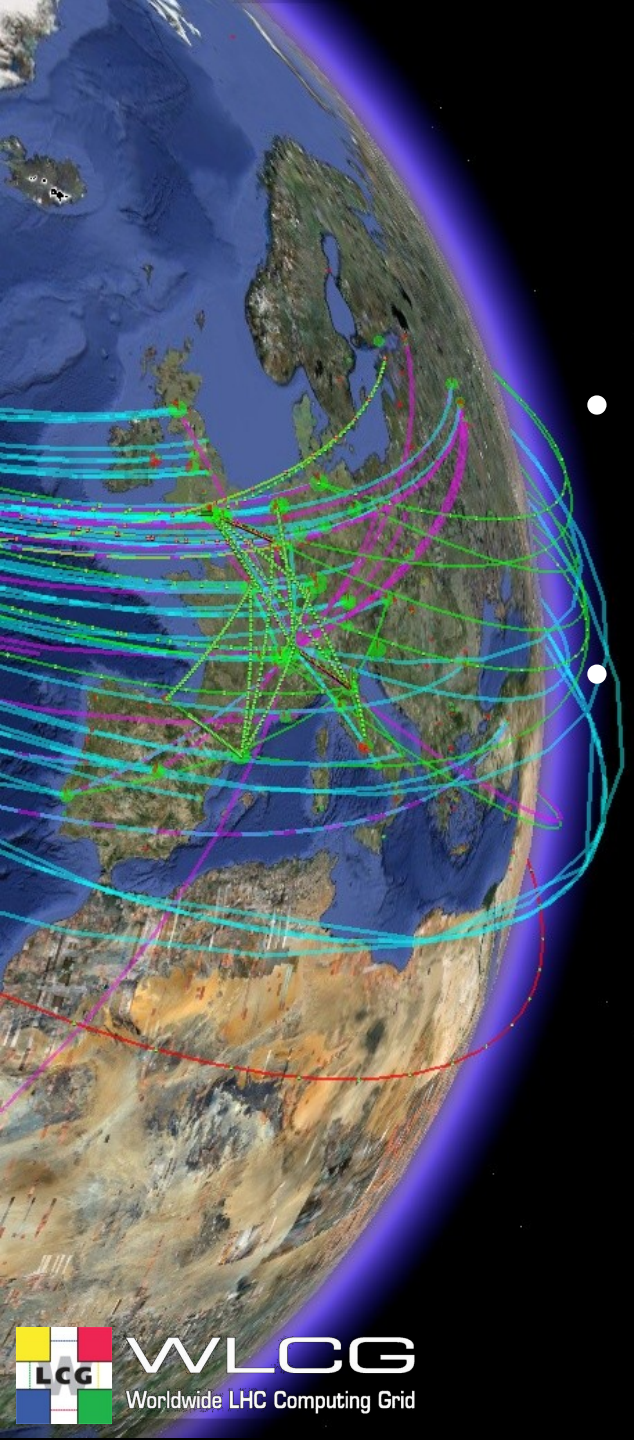

Summary

- WLCG was so far and still remains the only option to handle the LHC data volumes.
- Future LHC runs at higher energy and luminosity will present computing challenges that will need more R&D and further investigation of alternative resources (e.g. cloud & volunteer computing).

WLCG
Worldwide LHC Computing Grid