The Worldwide LHC Computing Grid (WLCG)

From Grids to Clouds

Dr. Markus Schulz
Support for Distributed Computing
CERN

Accelerating Science and Innovation

Tools: LHC and Detectors

pp, B-Physics, CP Violation
(matter-antimatter symmetry)

General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

Exploration of a new energy frontier in p-p and Pb-Pb collisions

Heavy ions, pp
(state of matter of early universe)

What is the data we collect?

- 150 million sensors deliver data …40 million times per second

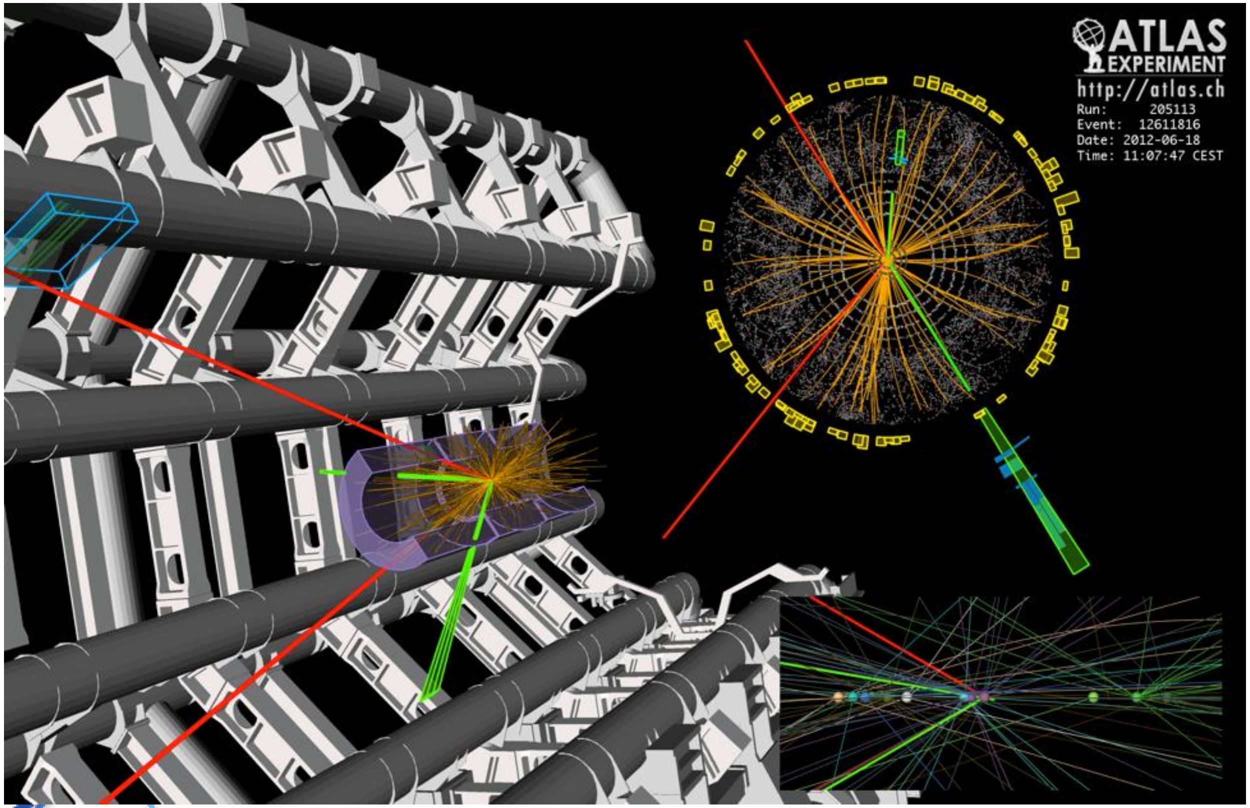

Up to 6 GB/s to be permanently stored after filtering

Data Collection and Archiving at CERN

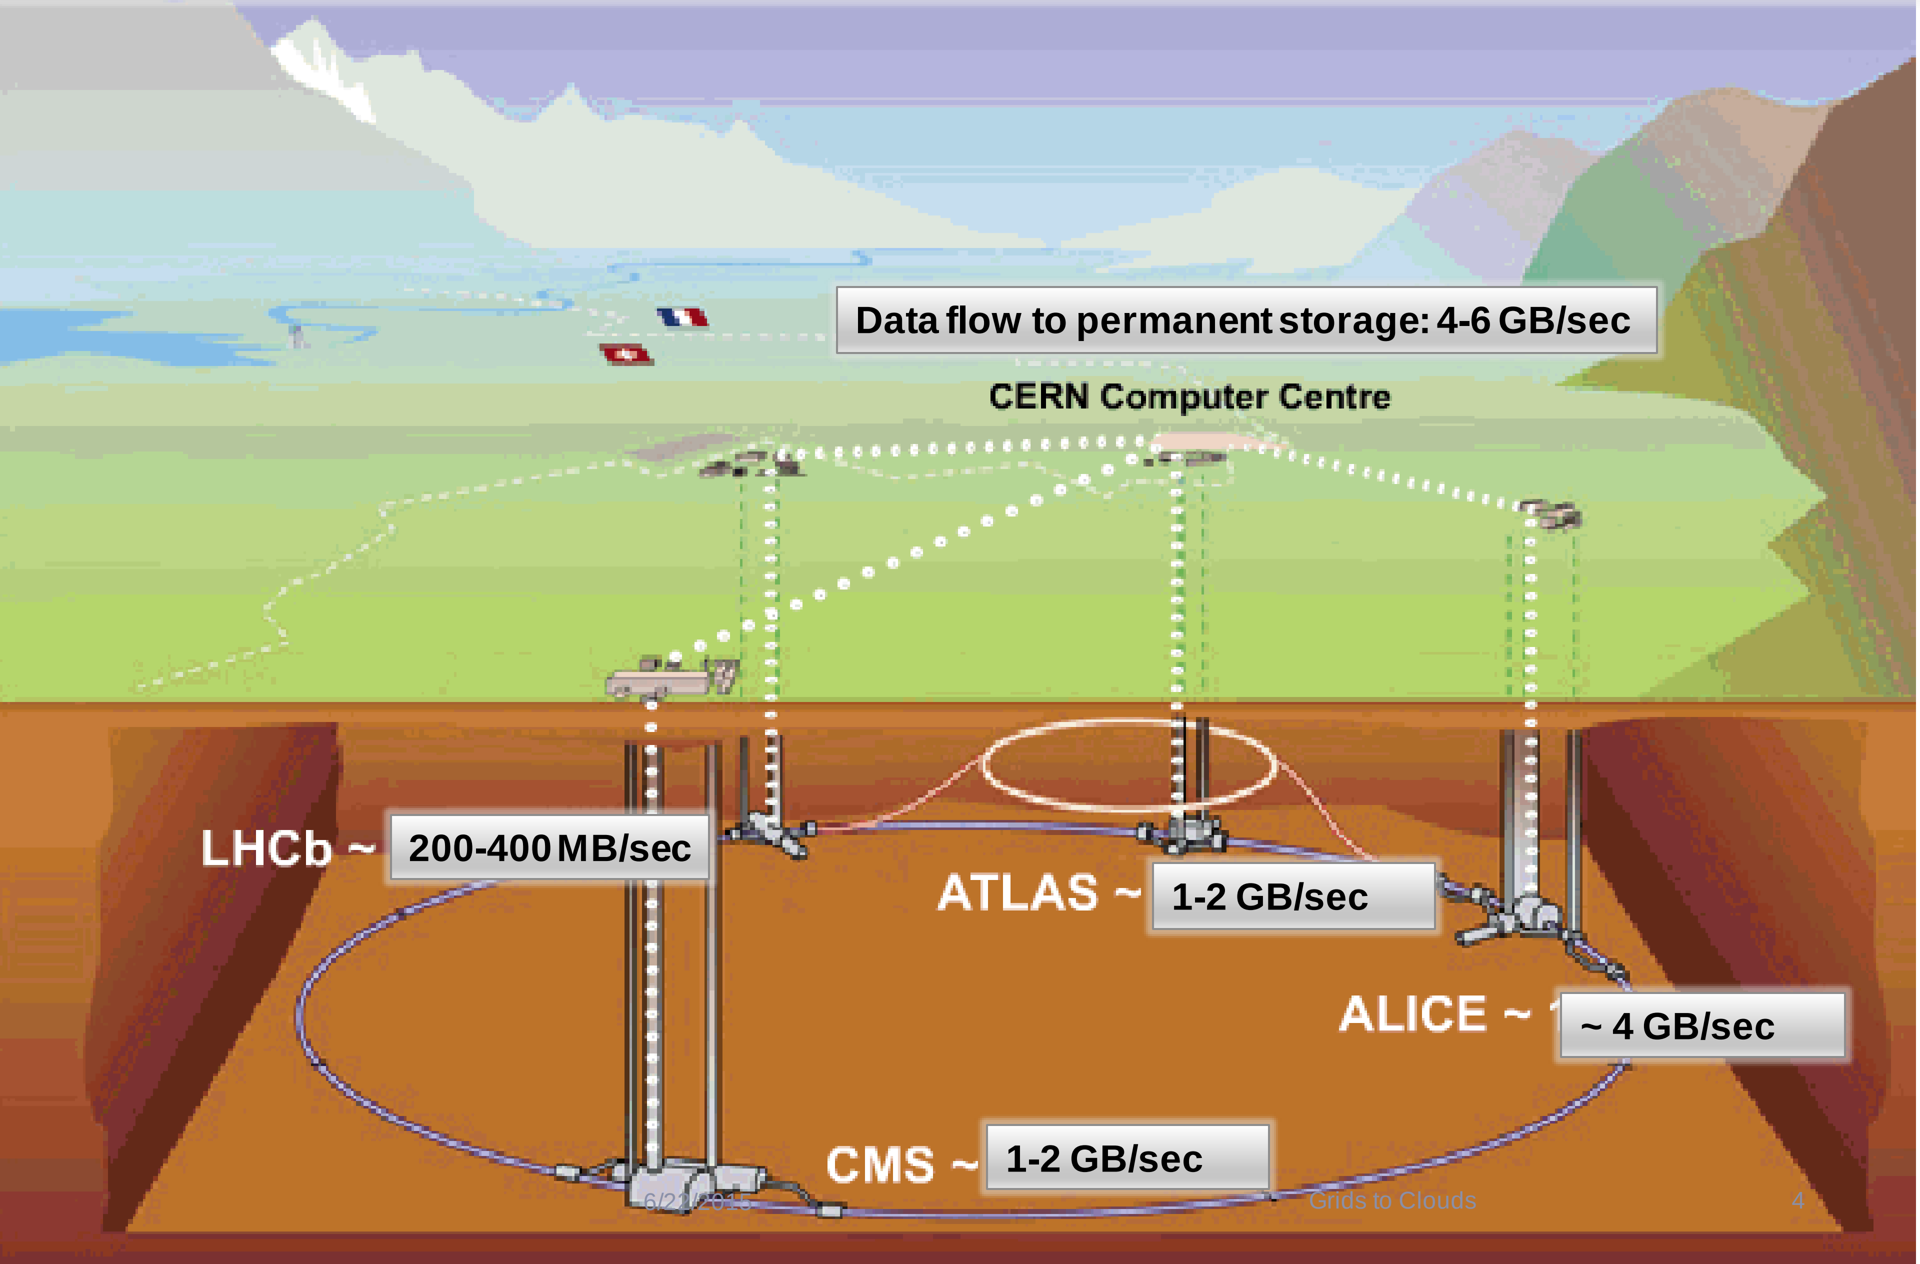

The Worldwide LHC Computing Grid

WLCG

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis

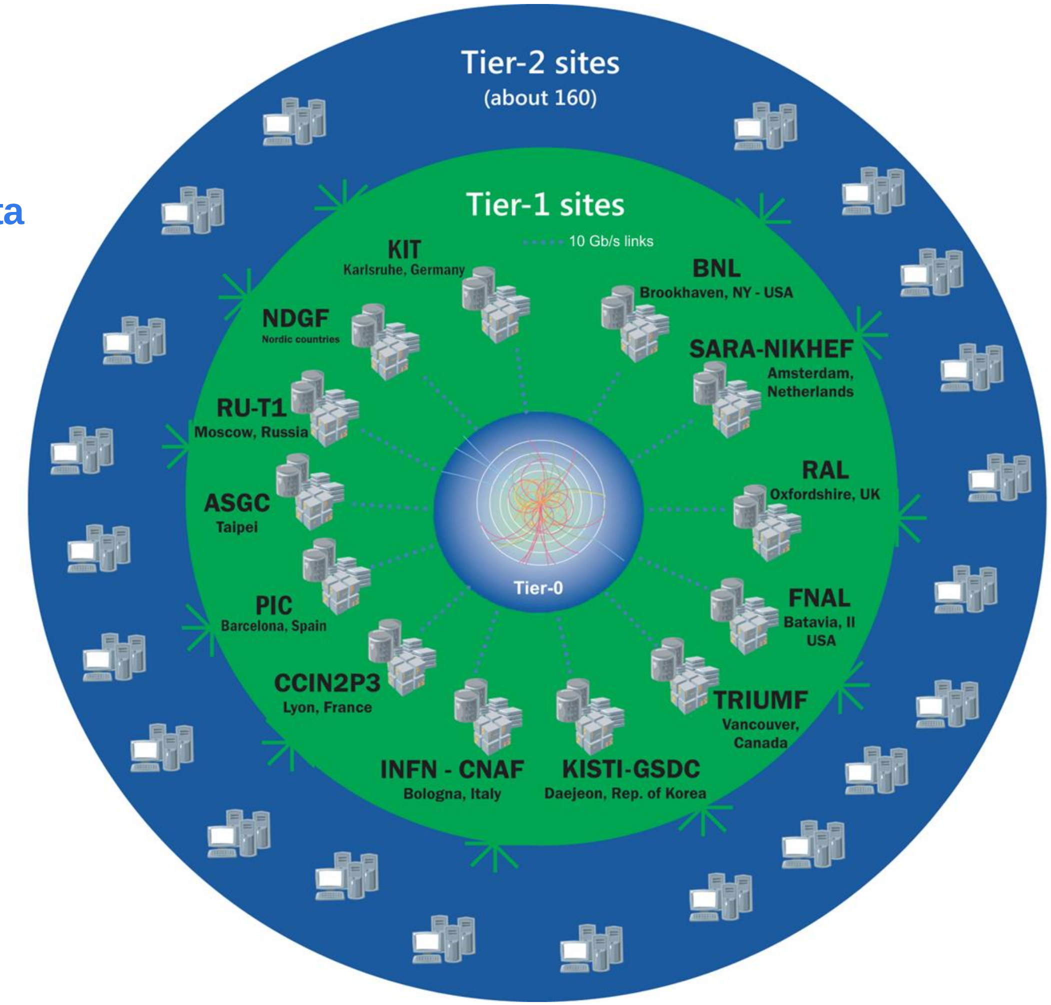

nearly 170 sites, 40 countries

~350'000 cores

500 PB of storage

>2 million jobs/day

10-100 Gb links
>2 million file transfers/day

WLCG:
An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

YEARS / ANS CERN

Original LHC computing model ~1999

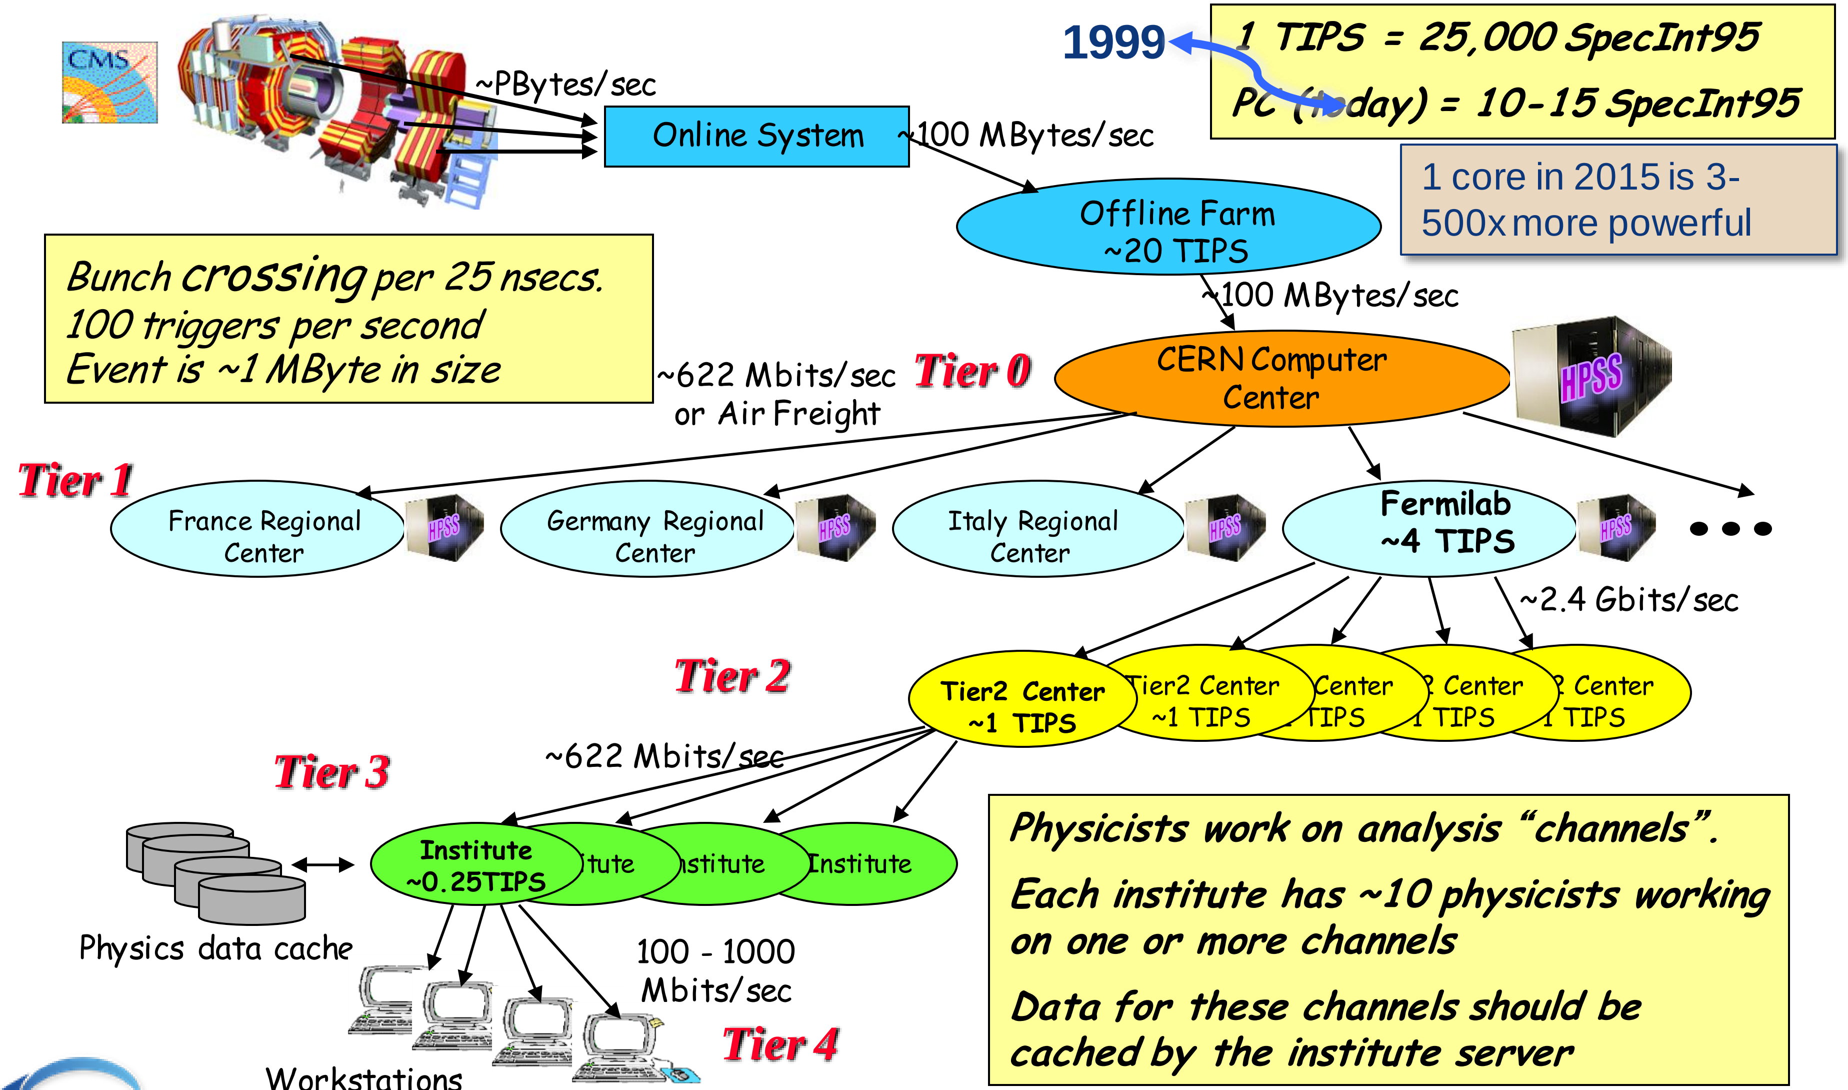

World-wide infrastructure

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 – L
  - deploying the results of DataGrid to provide a production facility for LHC experiments
- 2004-2006 – EU EGEE project phase
  - starts from the LCG grid
  - shared production infrastructure
  - expanding to other communities and sciences
- 2006-2008 – EU EGEE project phase
  - 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
  - Merging middleware stacks in Europe

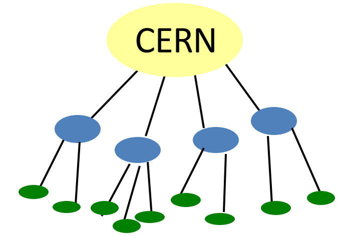

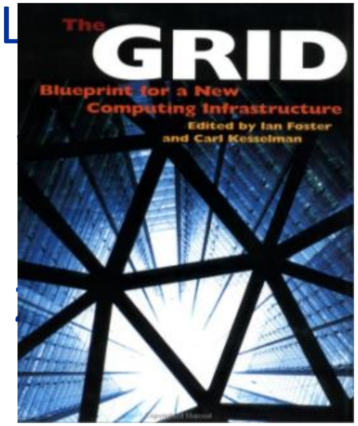

DataGRID

LCG

egee
Enabling Grids for E-sciencE

EUROPEAN MIDDLEWARE INITIATIVE

egi

Security Model - overview

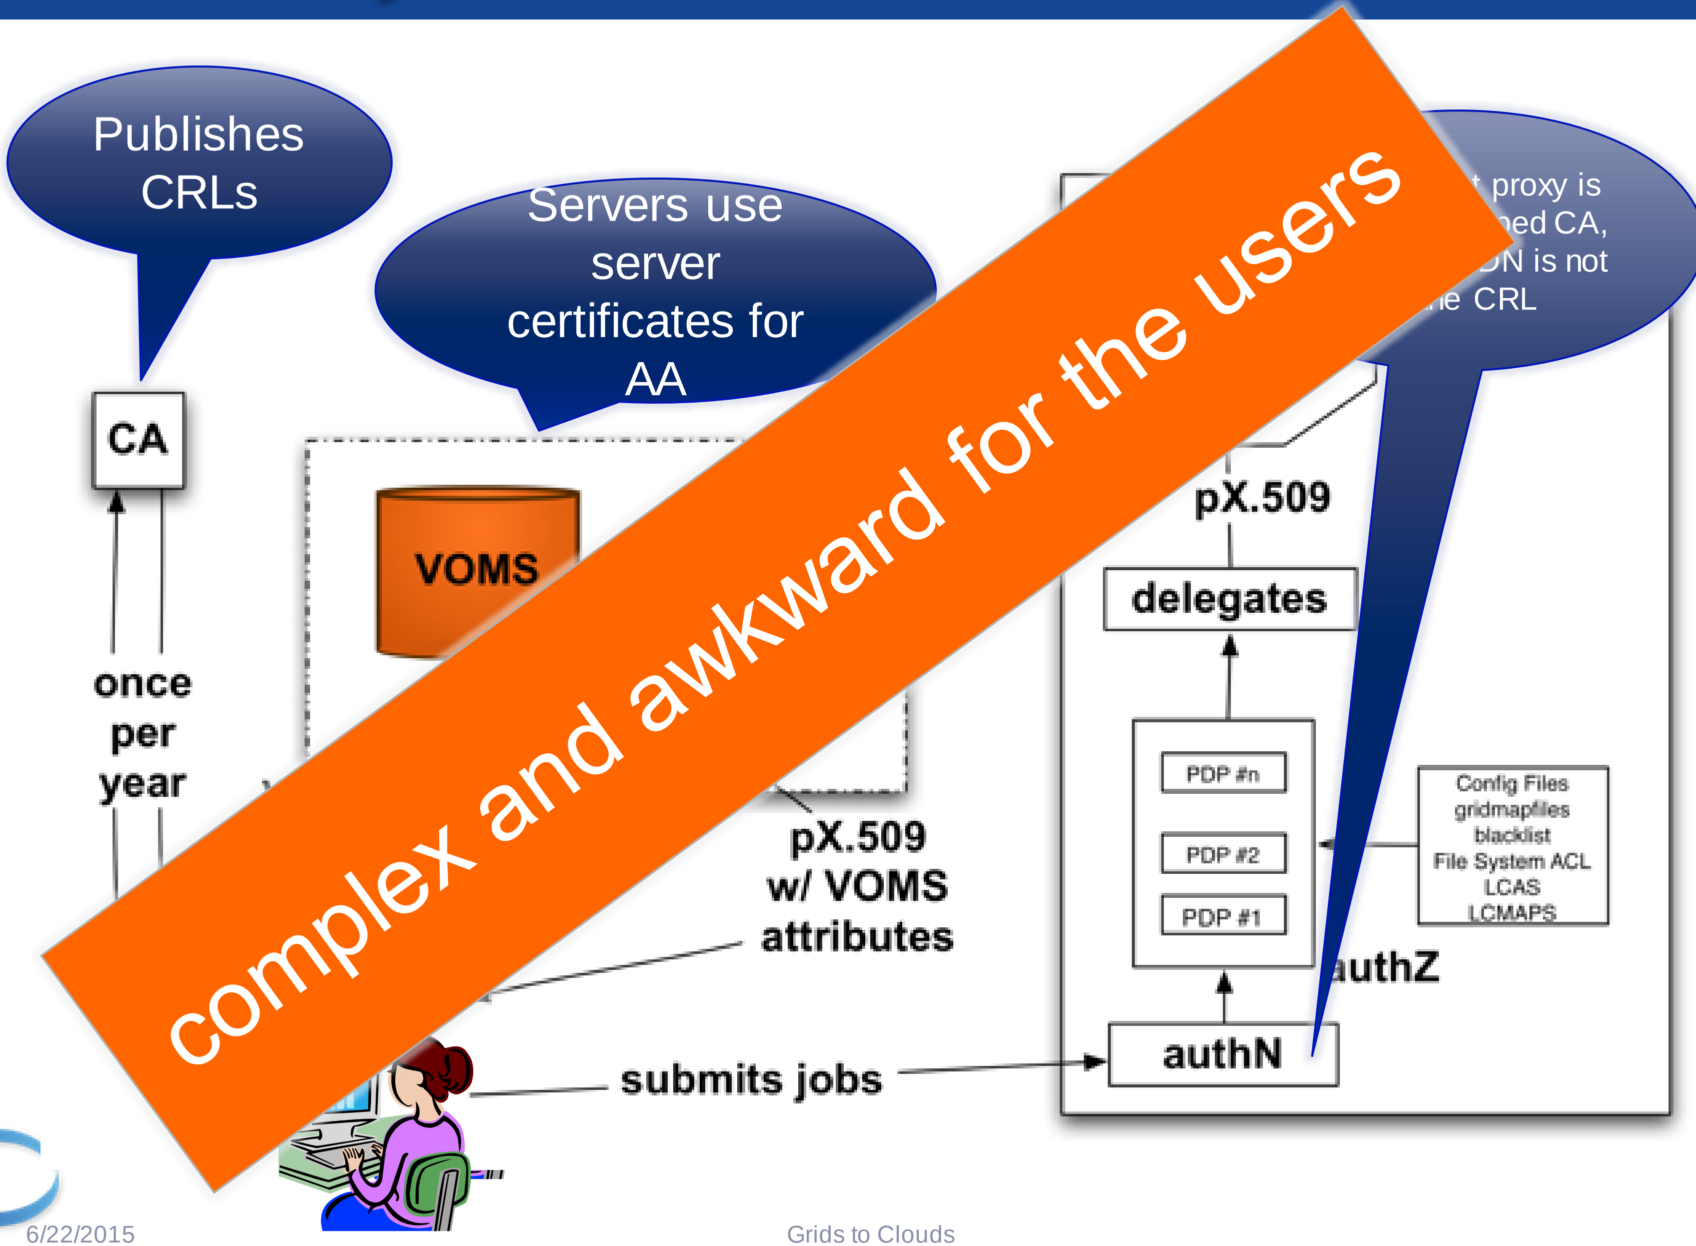

Public Key Based Security (X509)

- Public Key Infrastructure for Virtual

Moving to a system based on
Federated Identities
(close to what is used by eduroam)
Fundamental change!!!

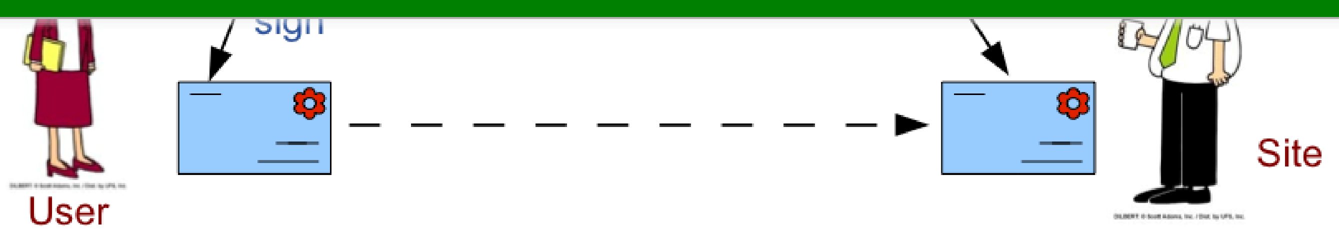

WLCG Worldwide LHC Computing Grid

LCG

CERN

eduGAIN

- Built on existing federations and infrastructures

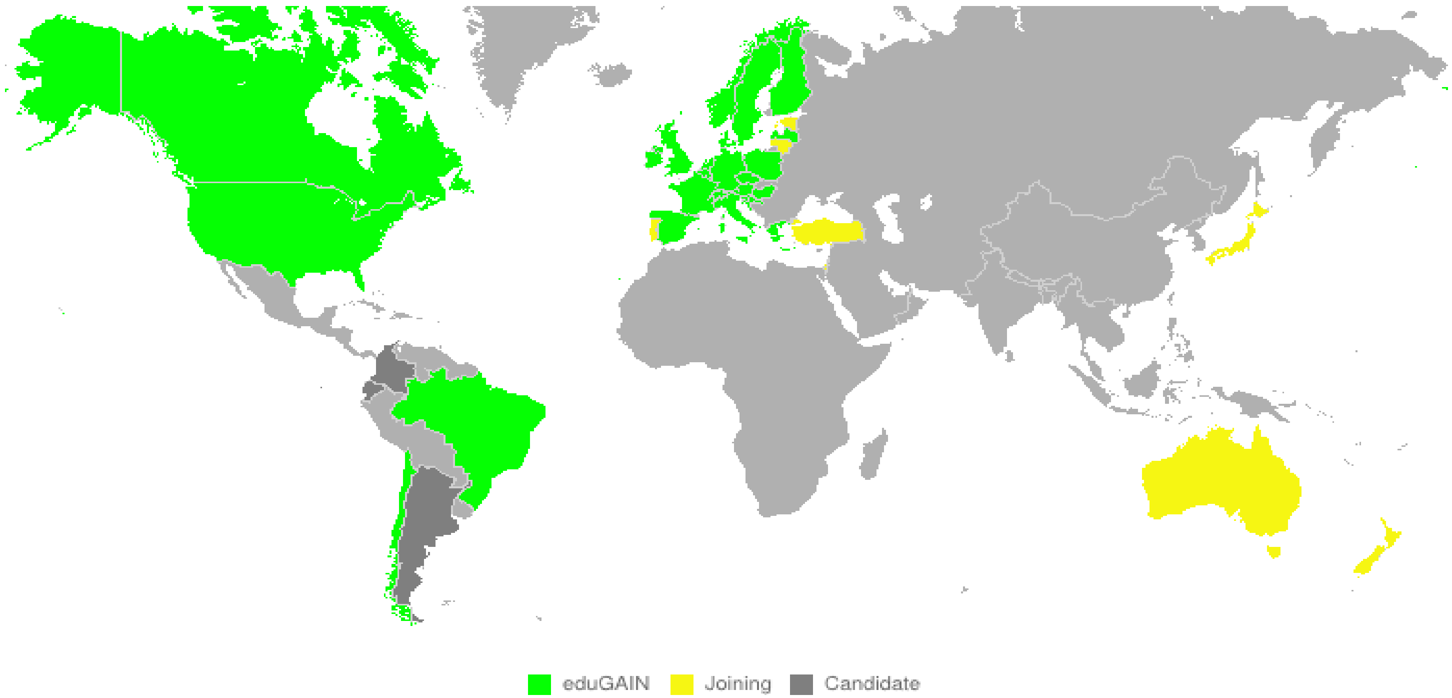

- CERN participates in eduGAIN via SWITCHaai
- Many NRENs participate in eduGAIN too

T0-T1s-T2s

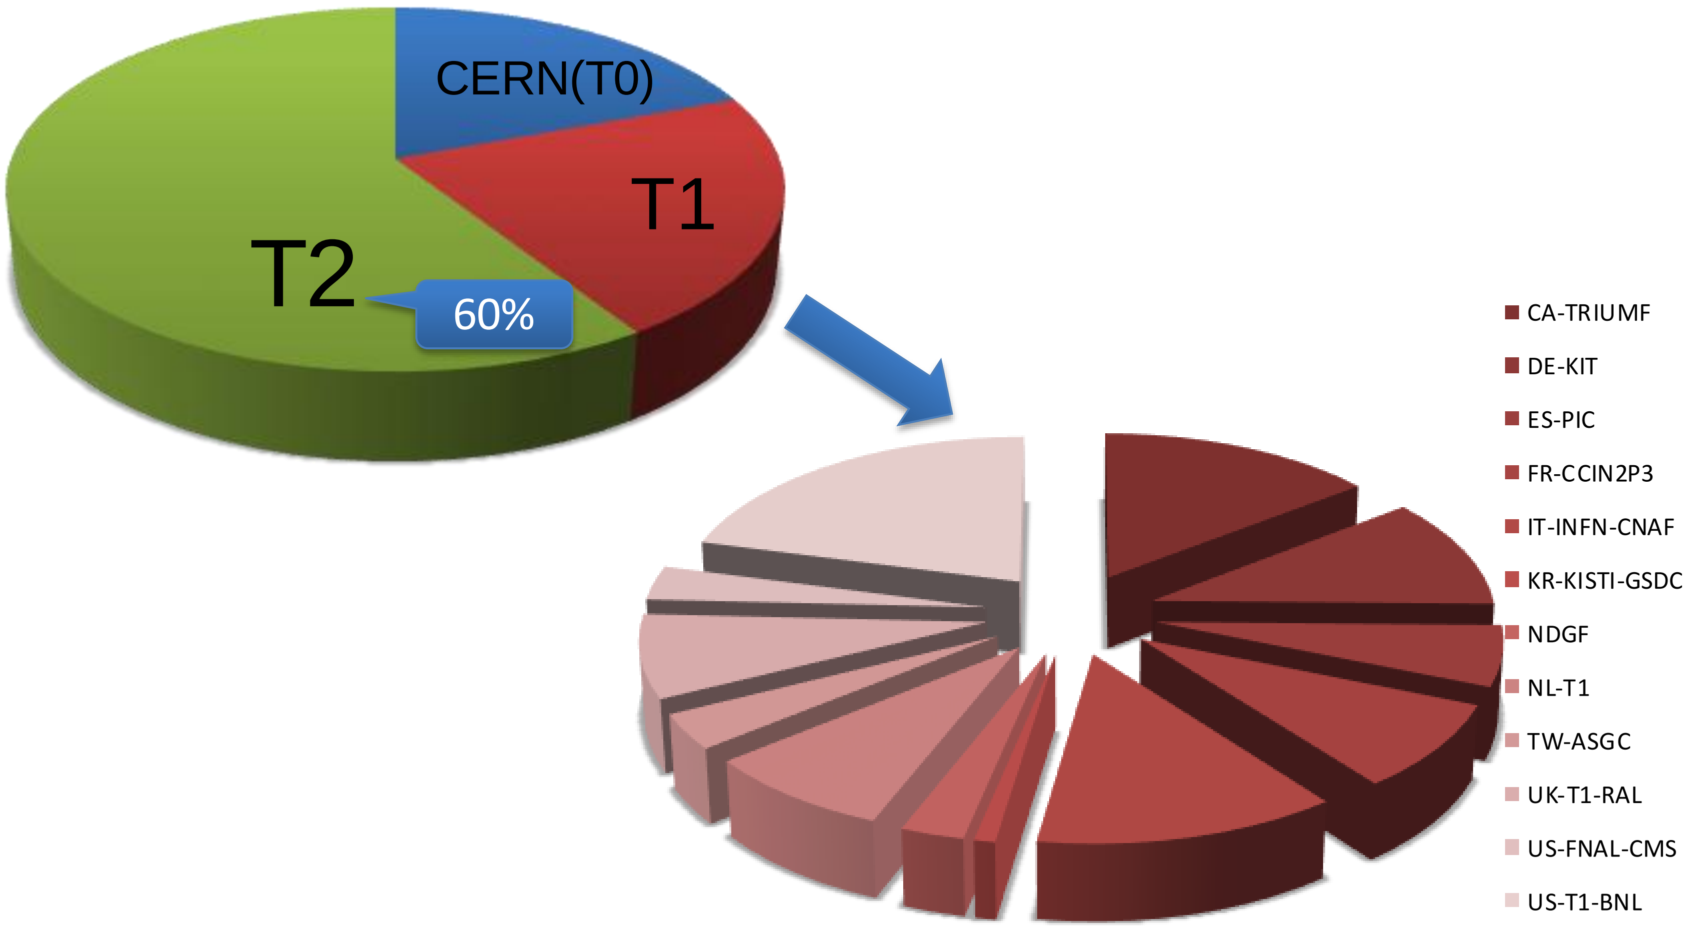

Distribution of jobs in WLCG

T2s are grouped by country or region

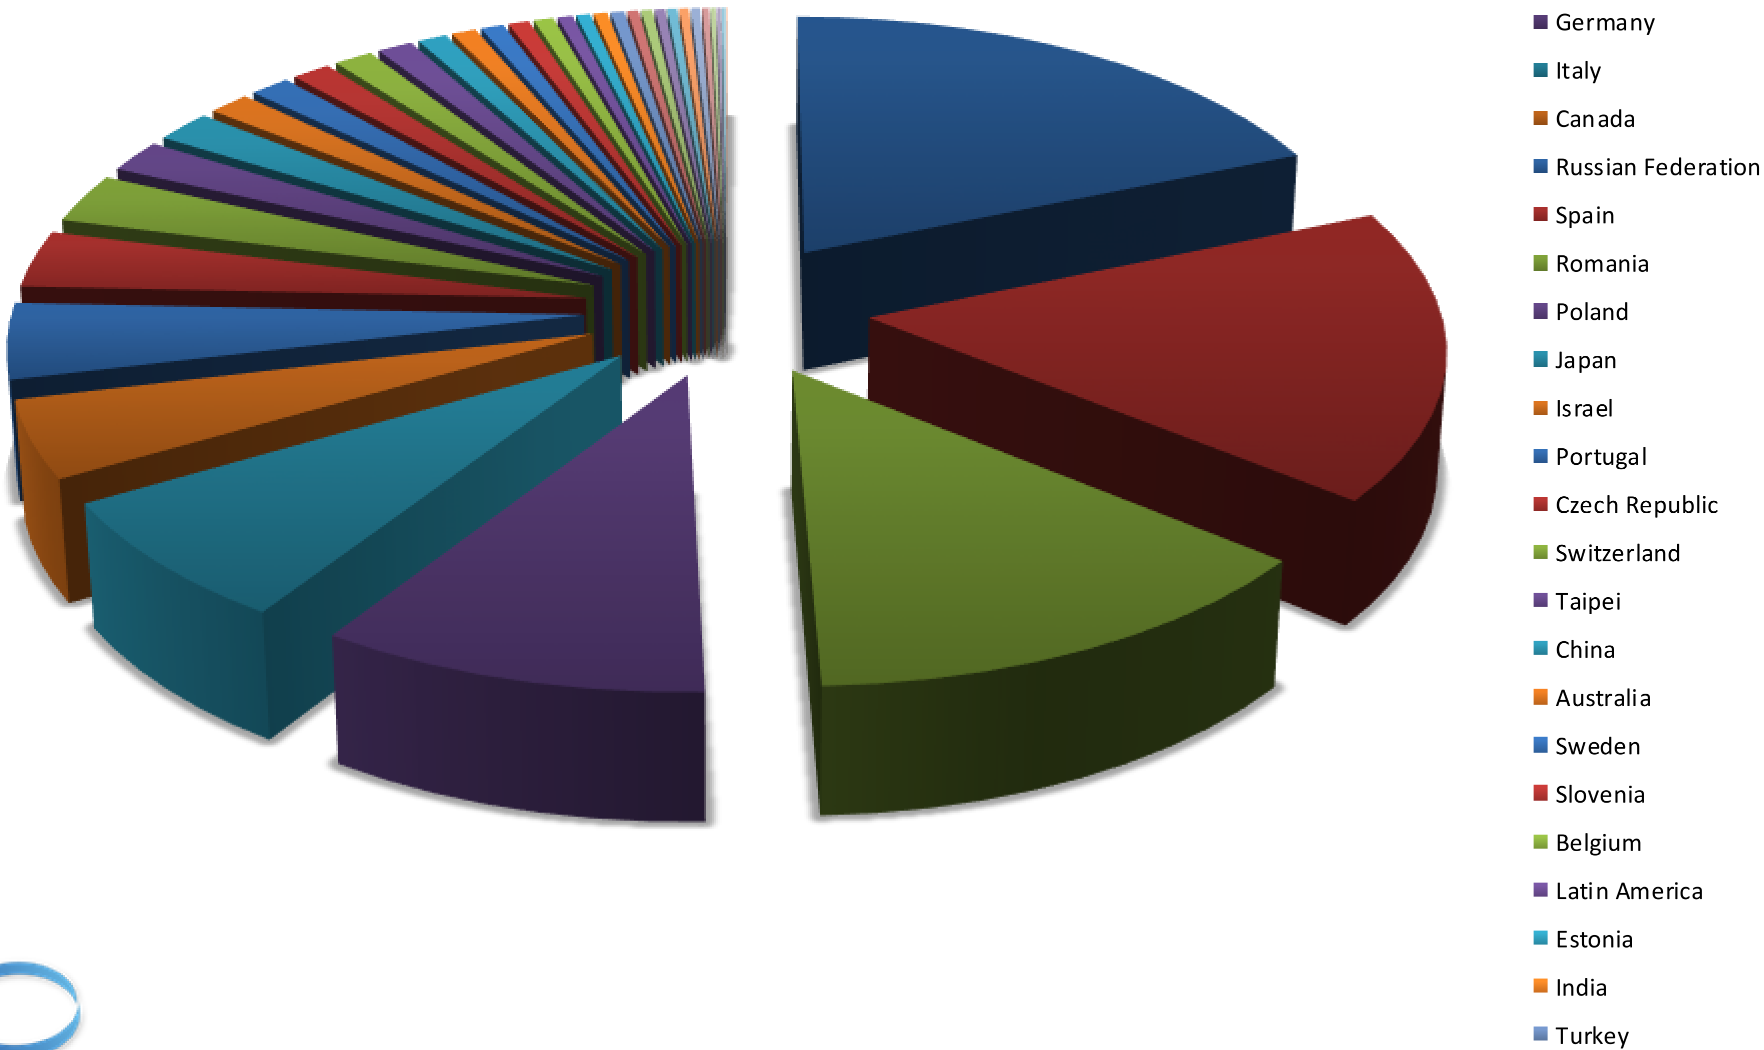

WLCG Worldwide LHC Computing Grid

HepSpec06 Hours

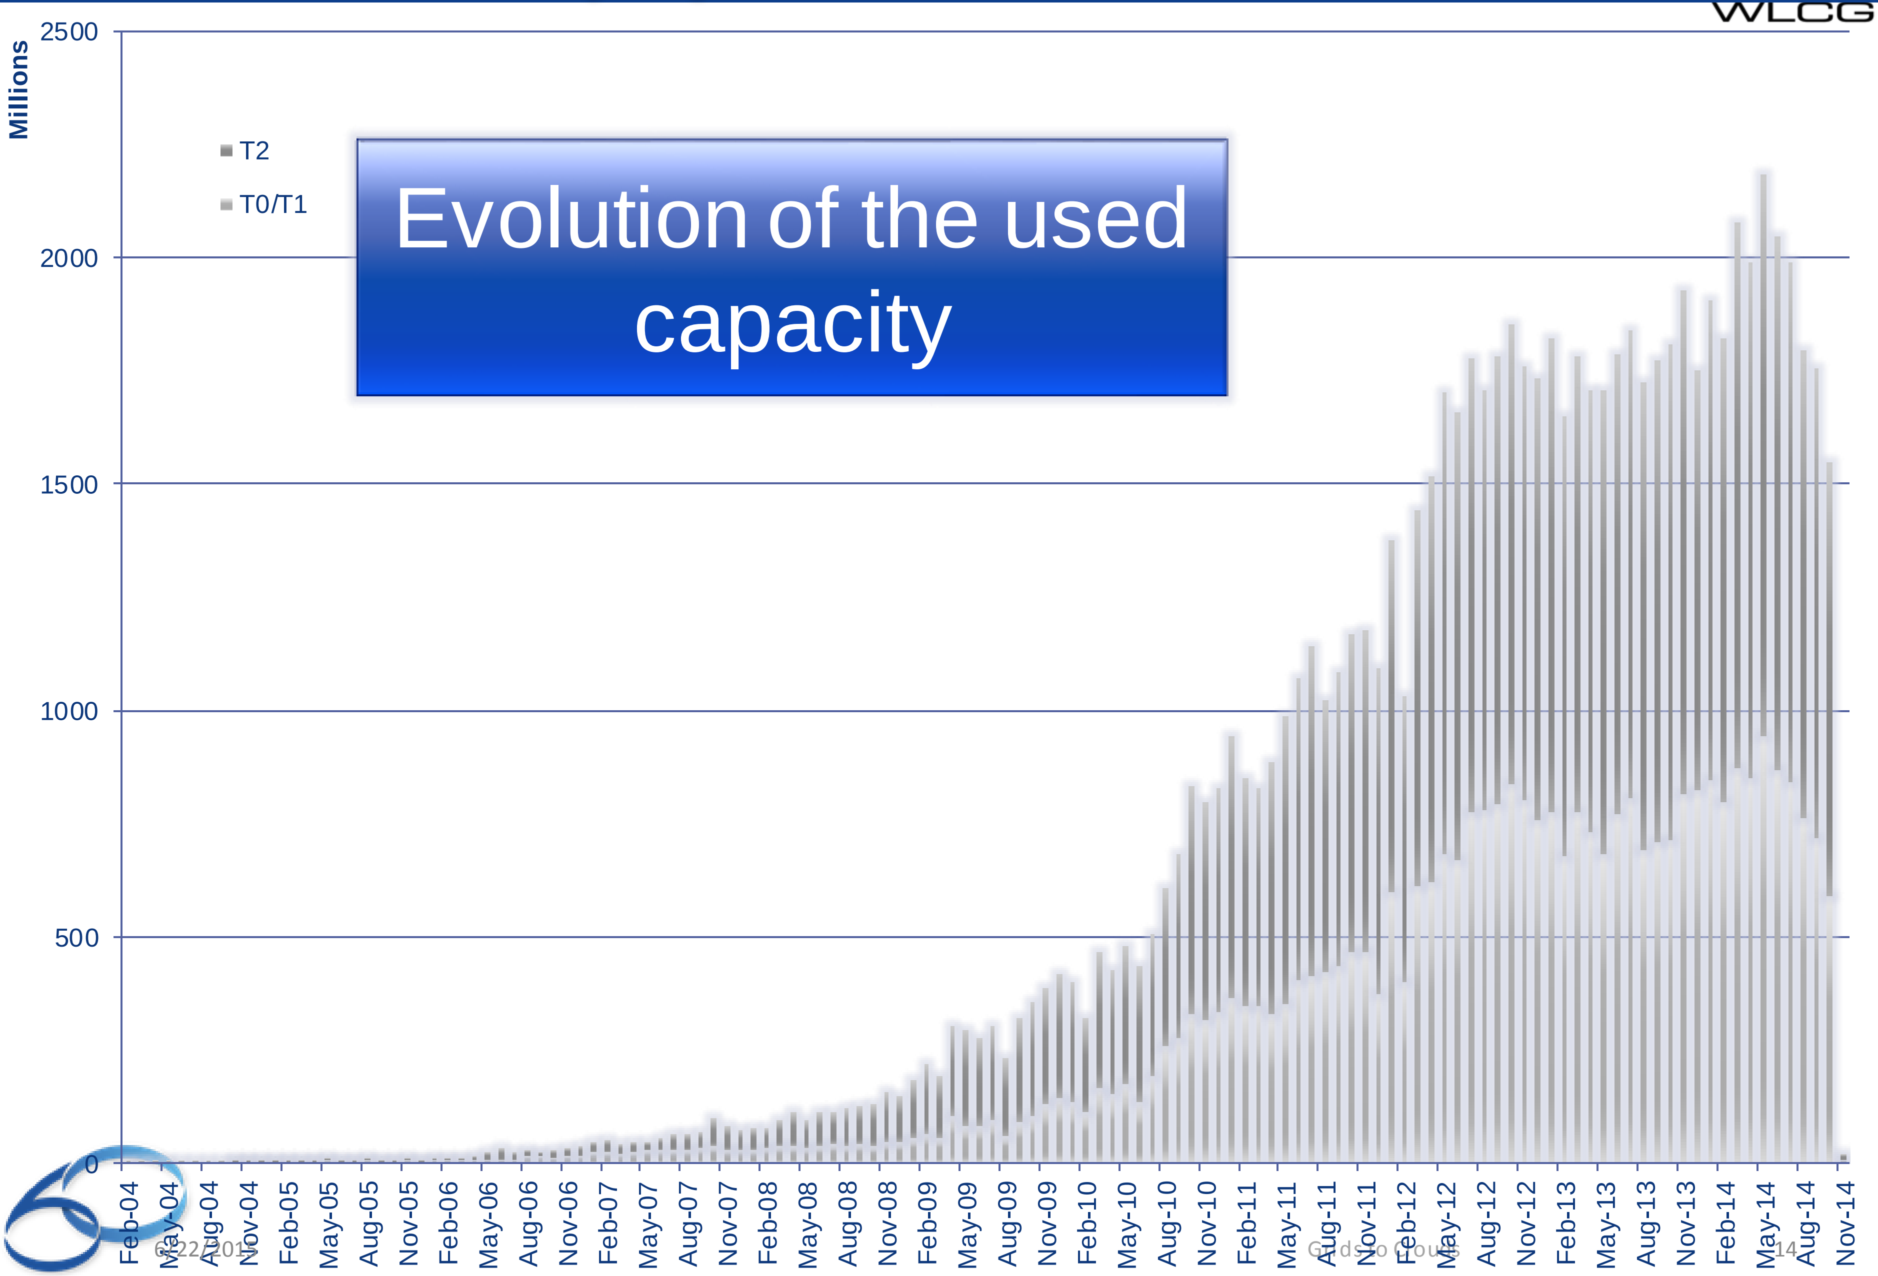

YEARS / ANS CERN

All Contributions

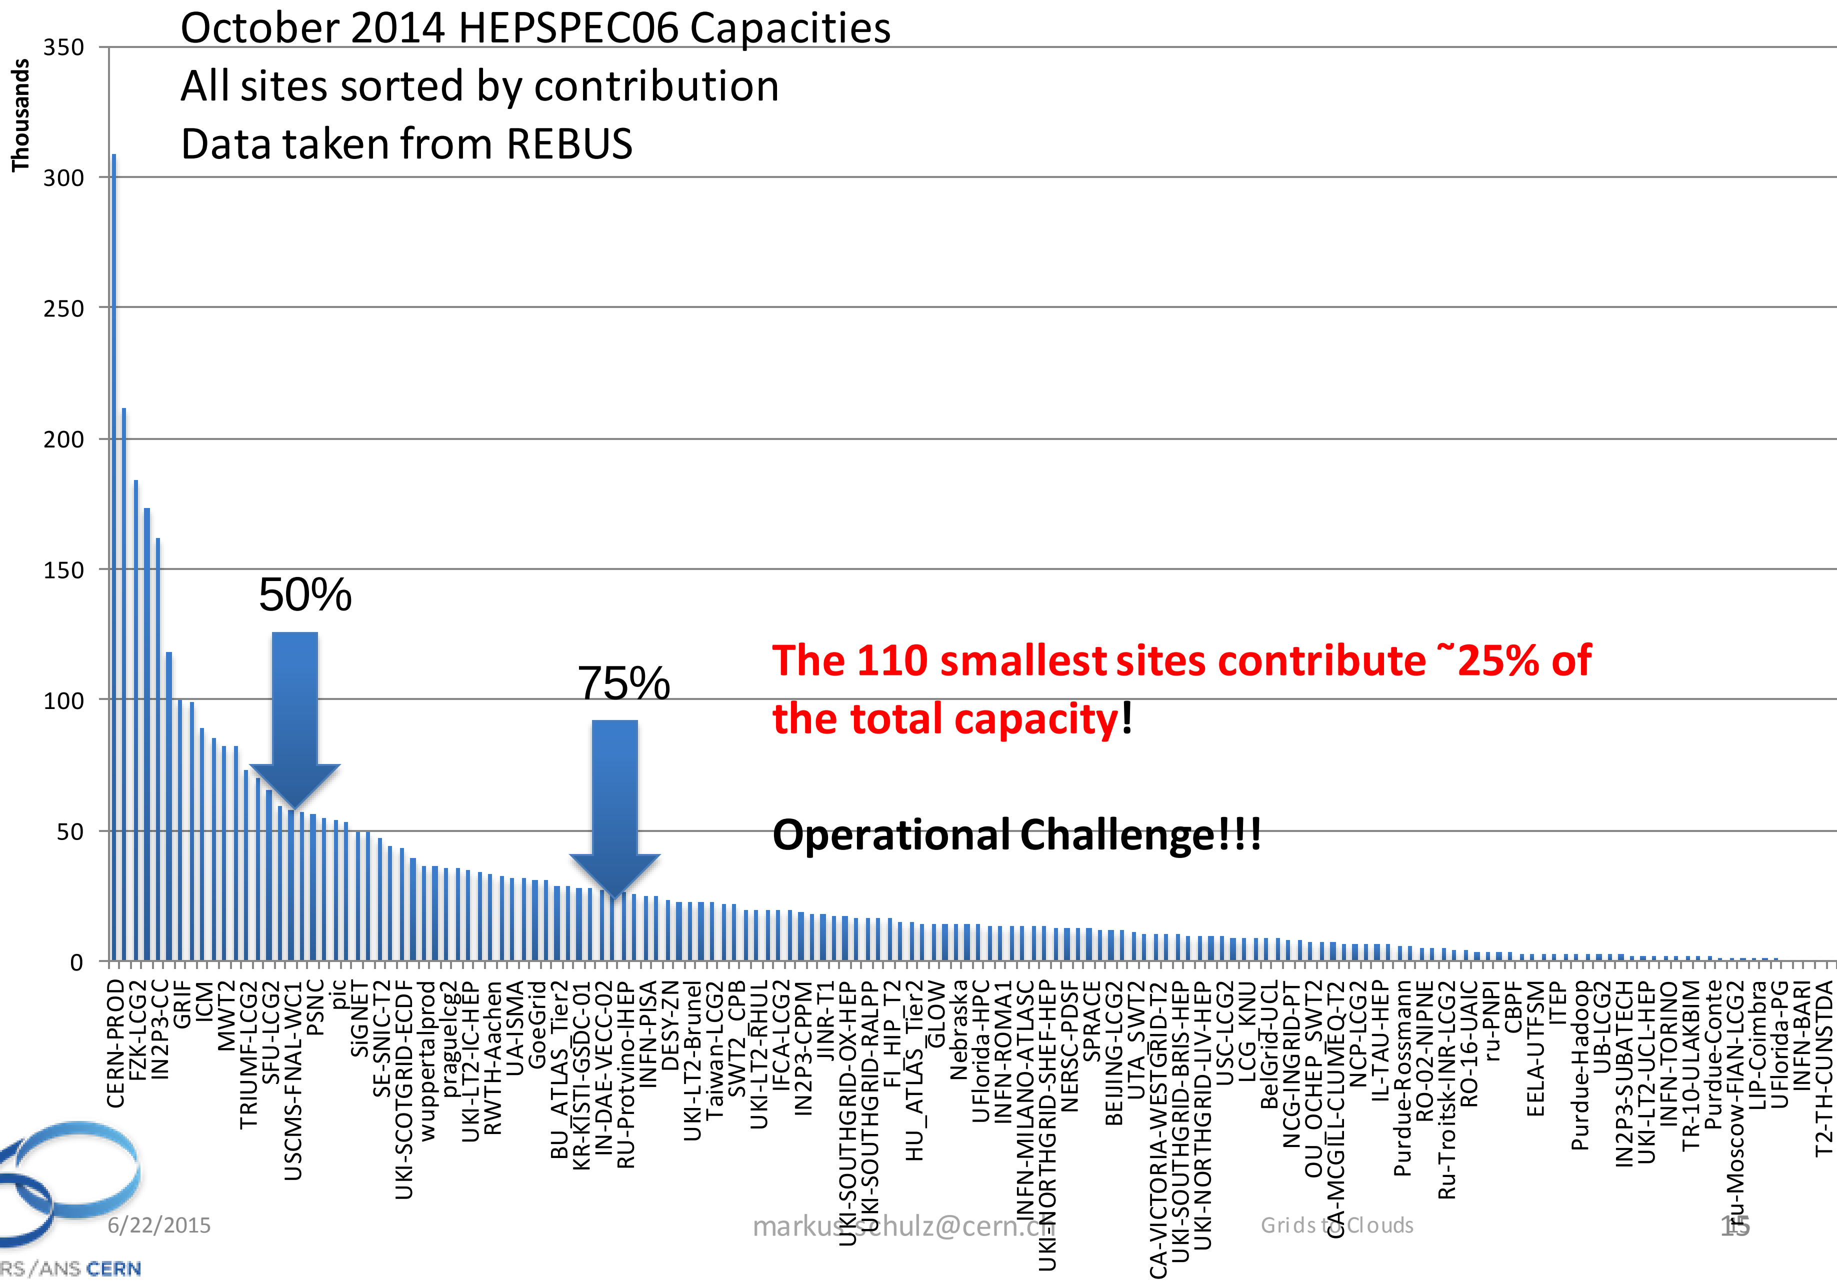

Scale of data today …

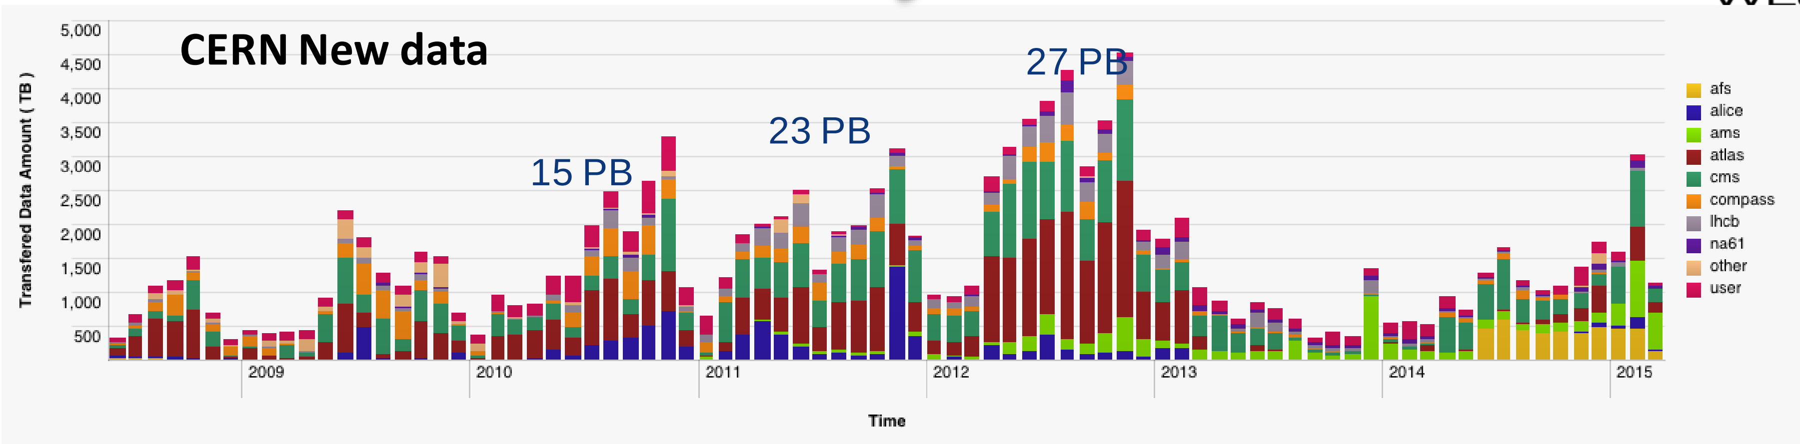

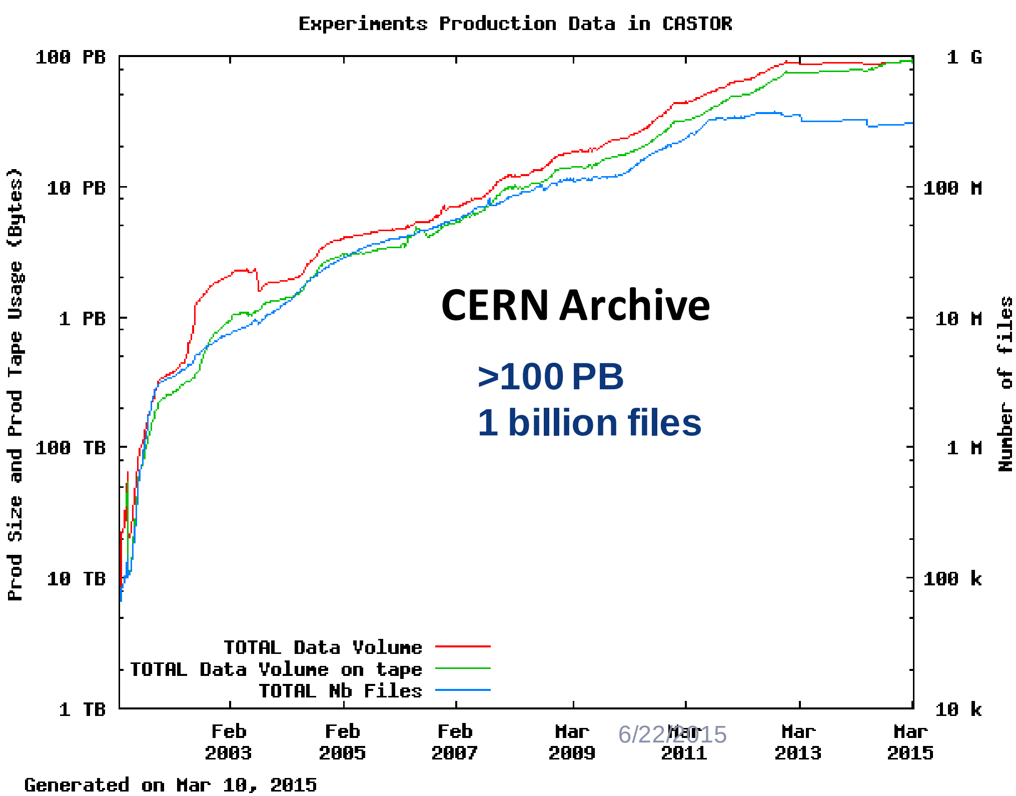

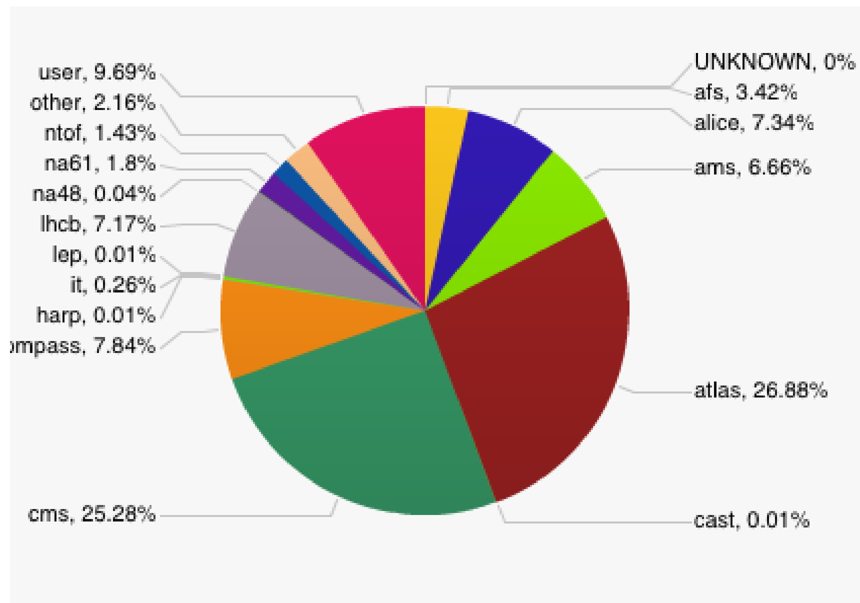

LHC – Big Data…

Few PB of raw data becomes ~100 PB! →

- Duplicate raw data
- Simulated data
- Many derived data products
- Recreate as software gets improved
- Replicate to allow physicists to access it

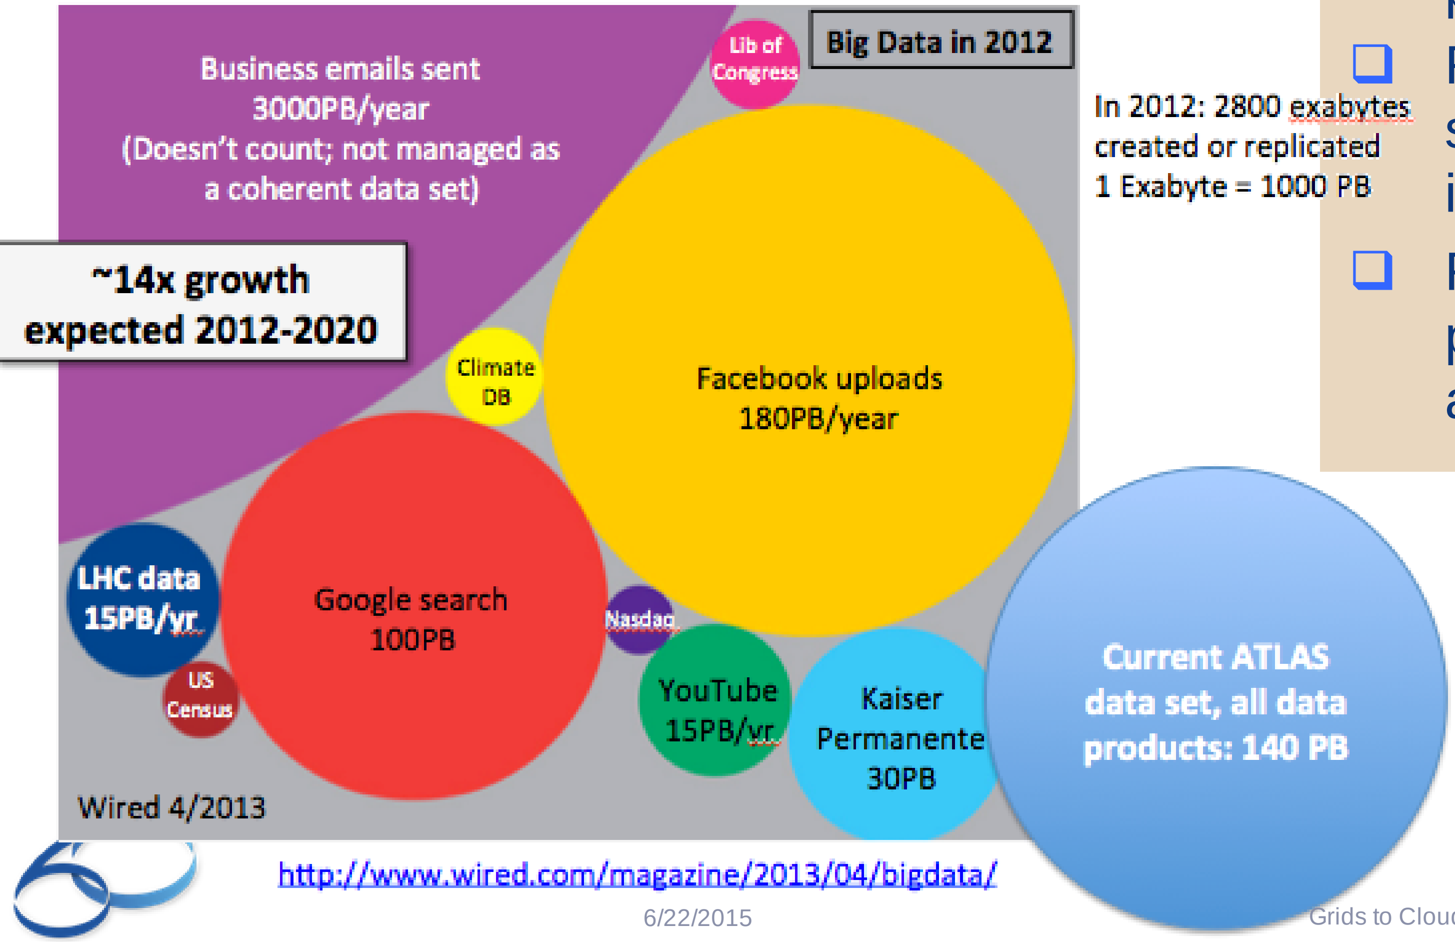

In 2012: 2800 exabytes created or replicated
1 Exabyte = 1000 PB

http://www.wired.com/magazine/2013/04/bigdata/

LHC OPN

- Optical Private Network
- Support T0 – T1 transfers
- Some T1 – T1 traffic
- Managed by LHC Tier 0 and Tier 1 sites

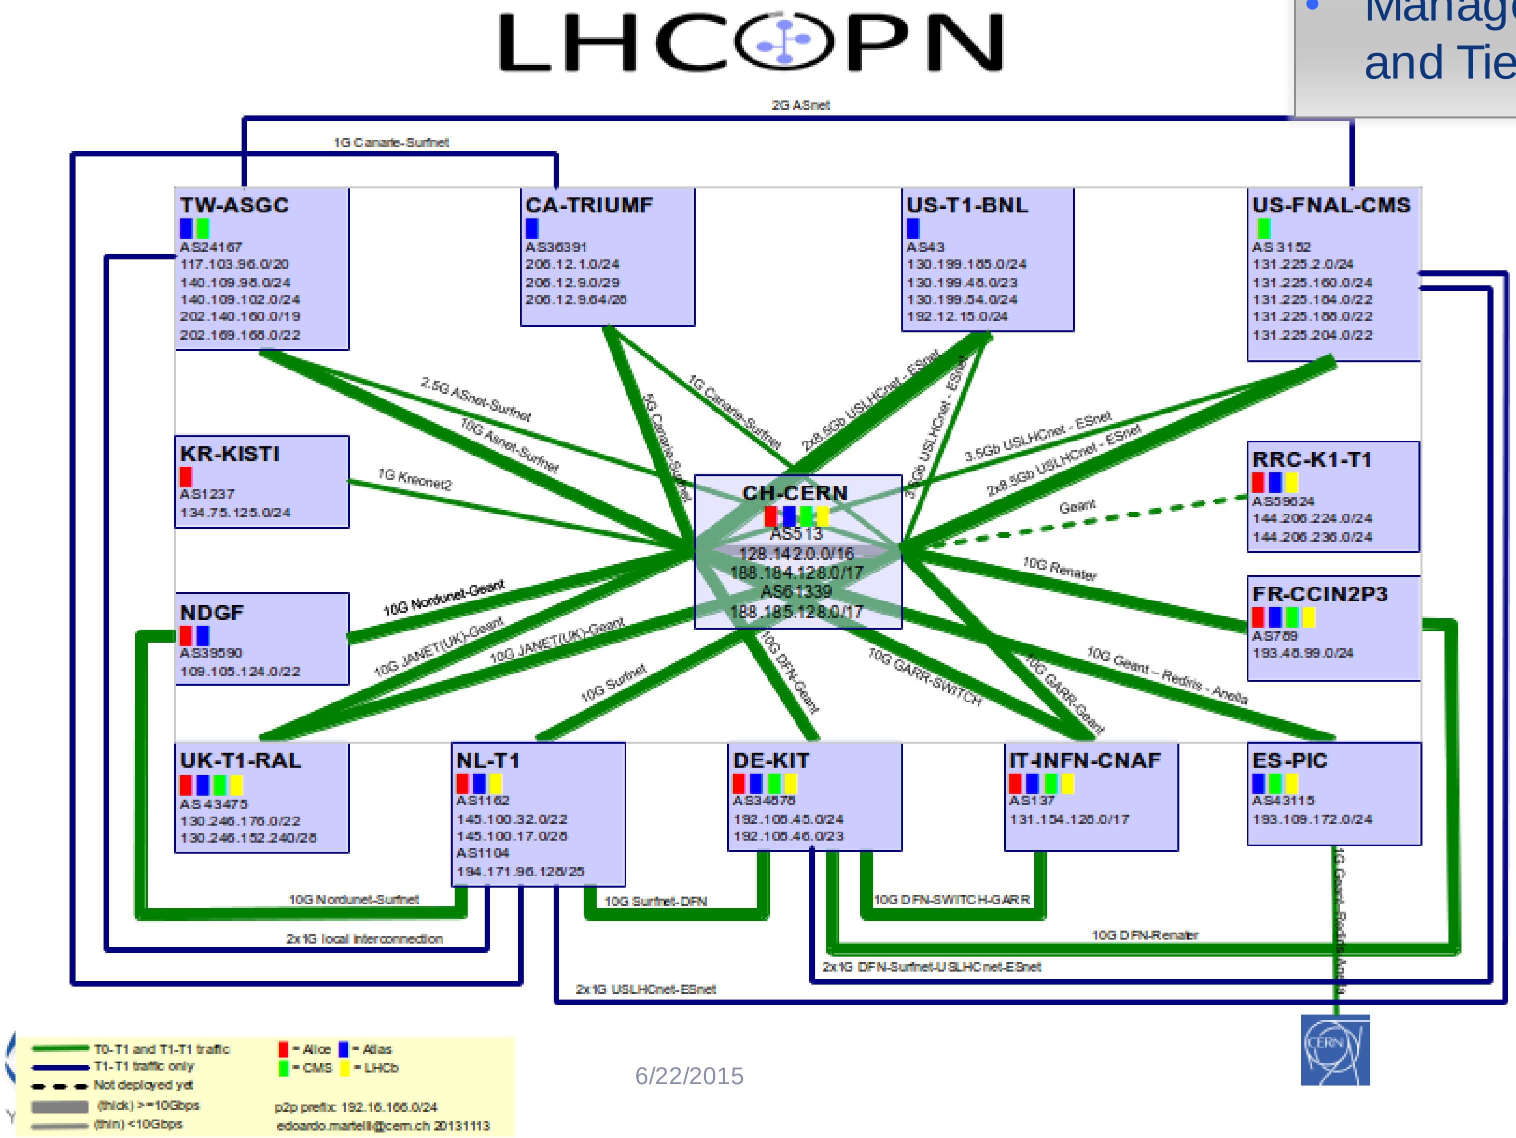

LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management

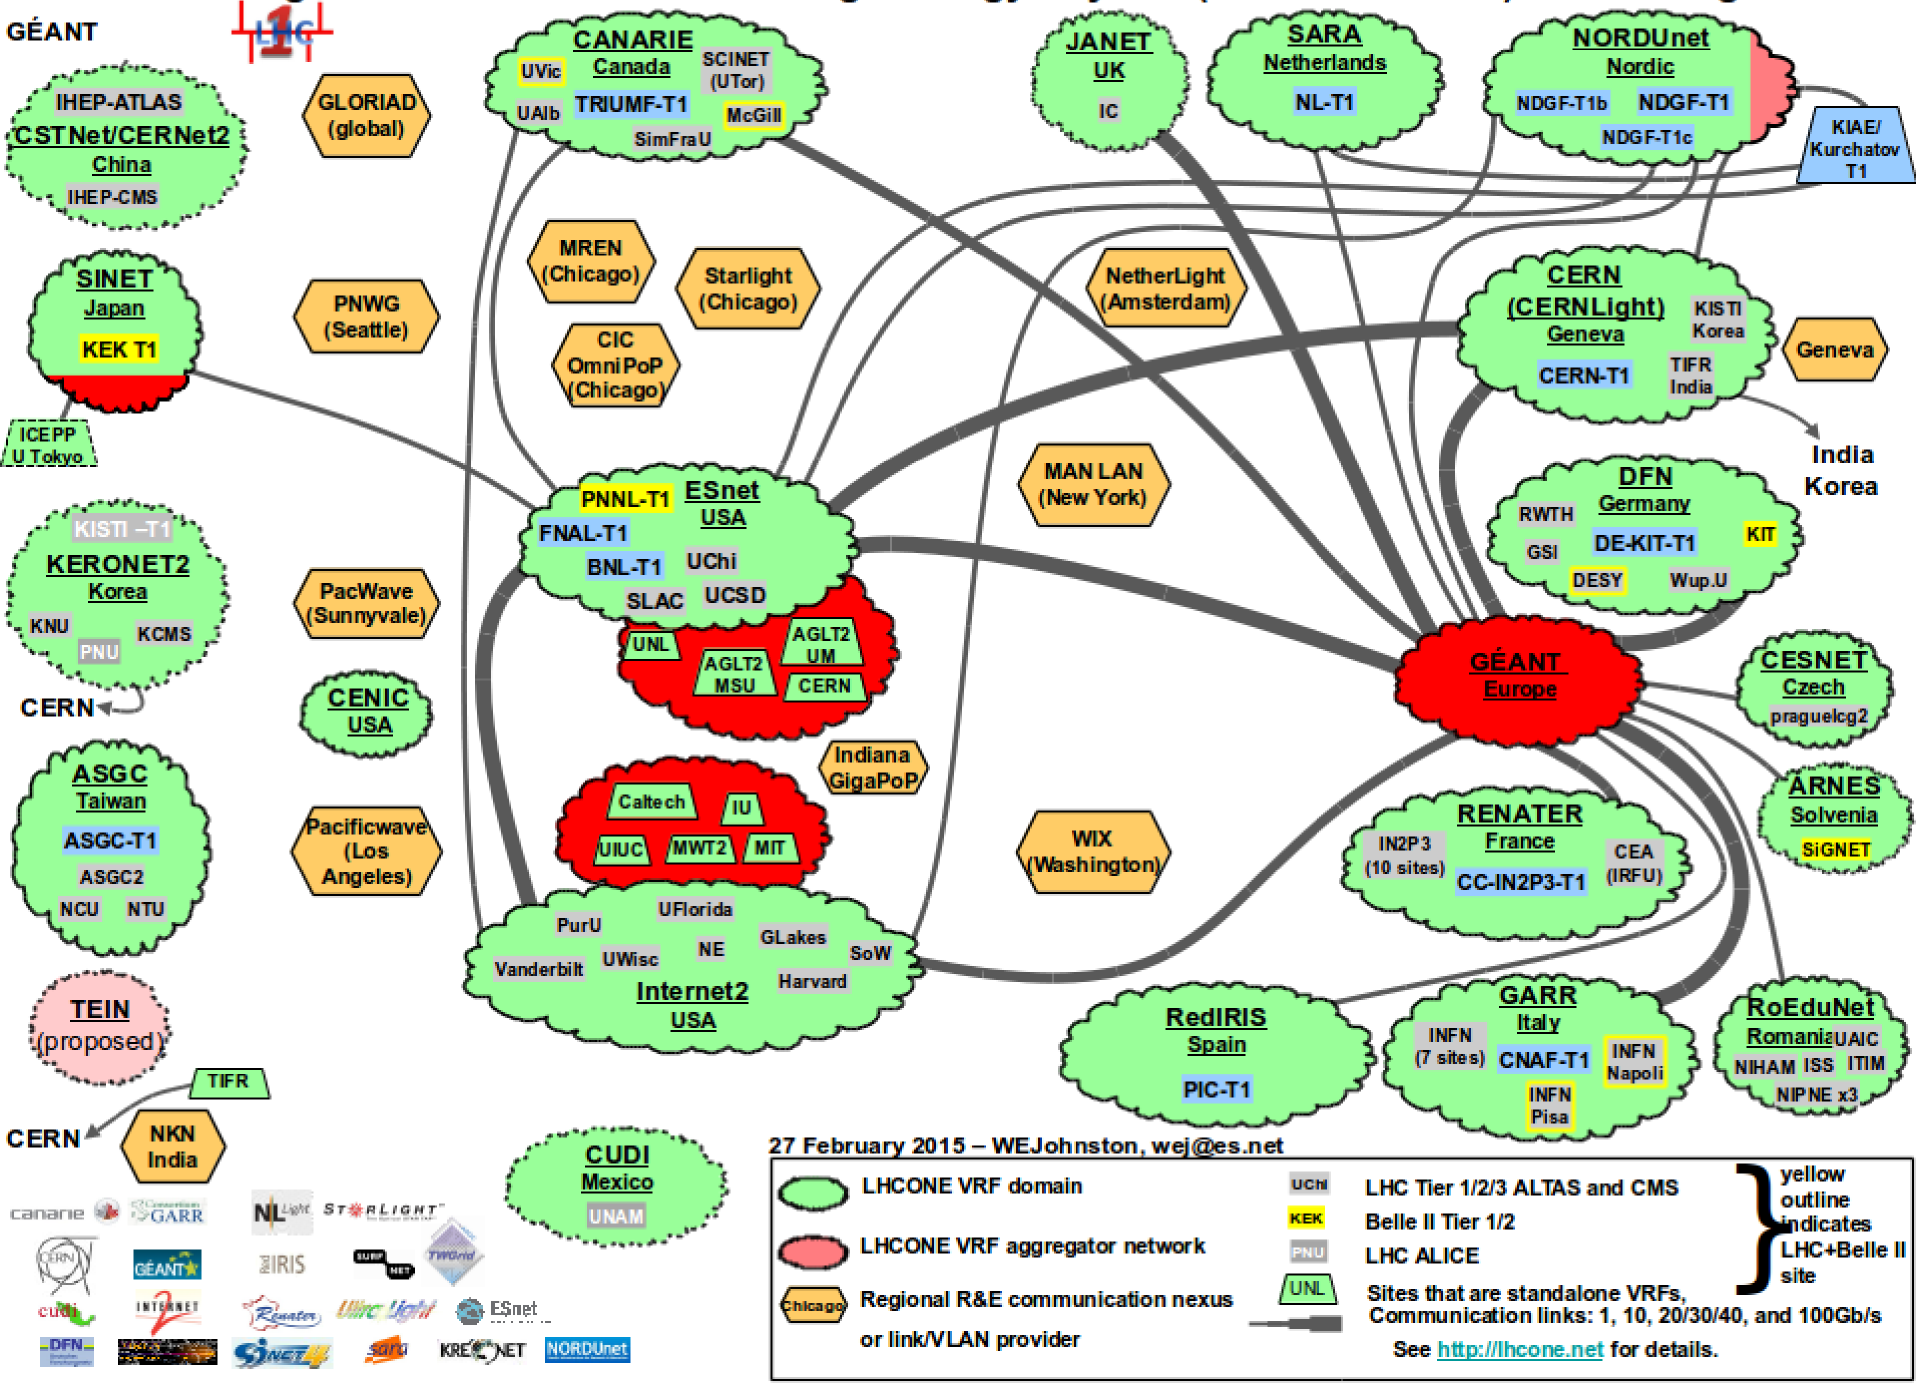

Data Transfers

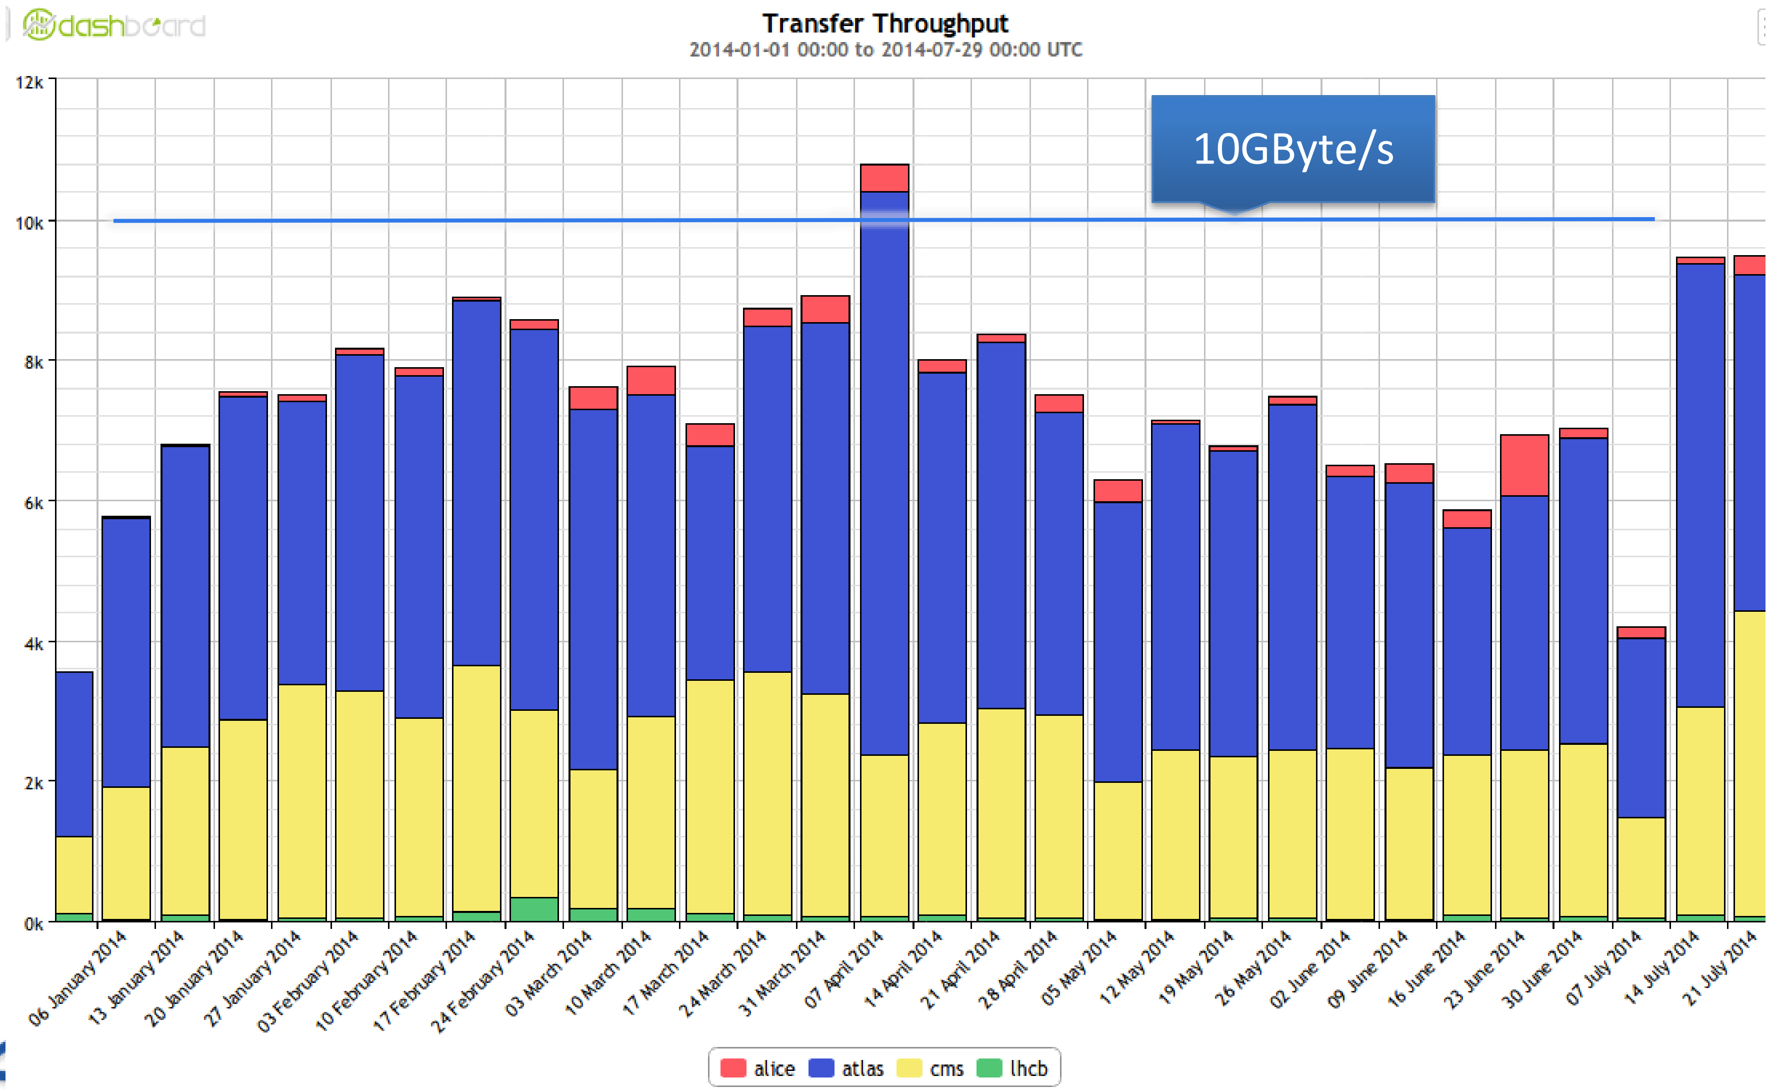

Connectivity over 4 hours

Connectivity over 4 hours

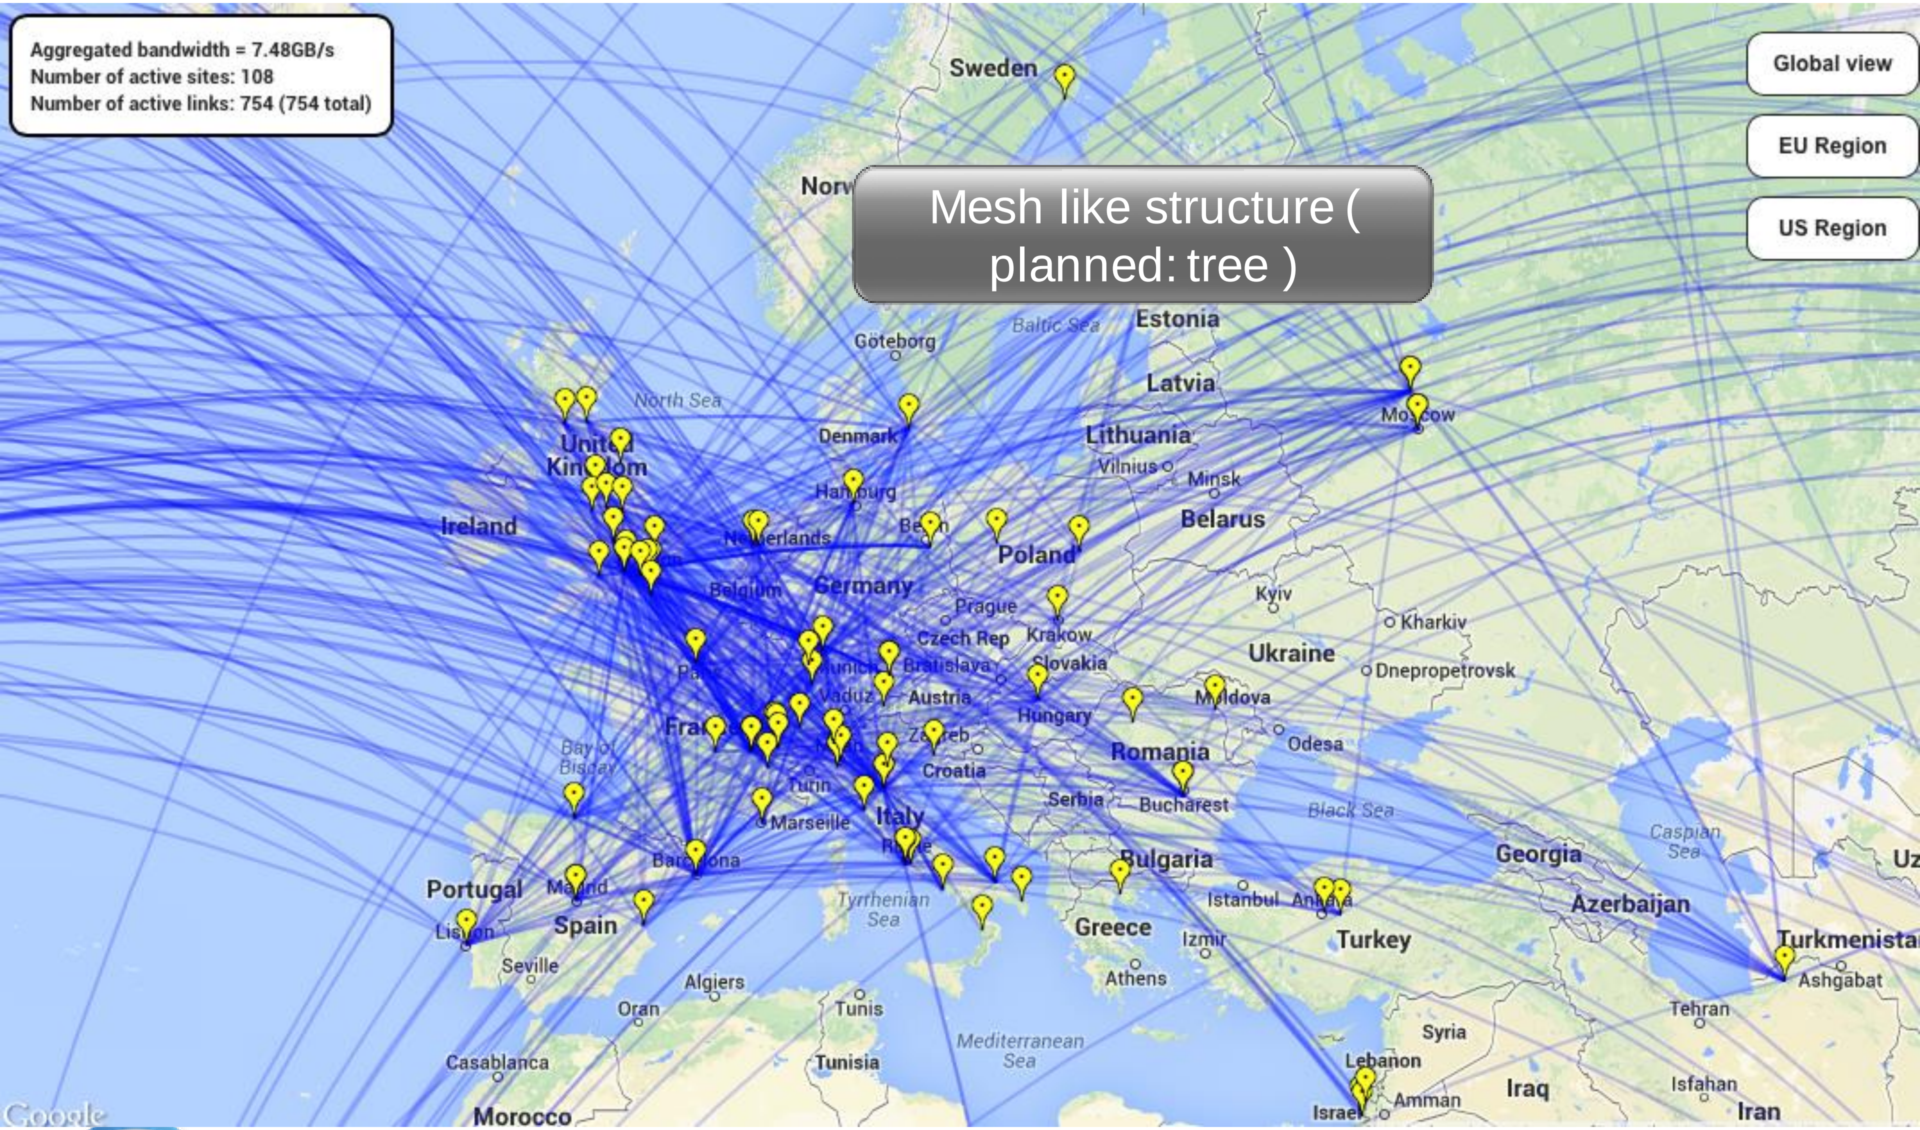

What's next?

The LHC timeline

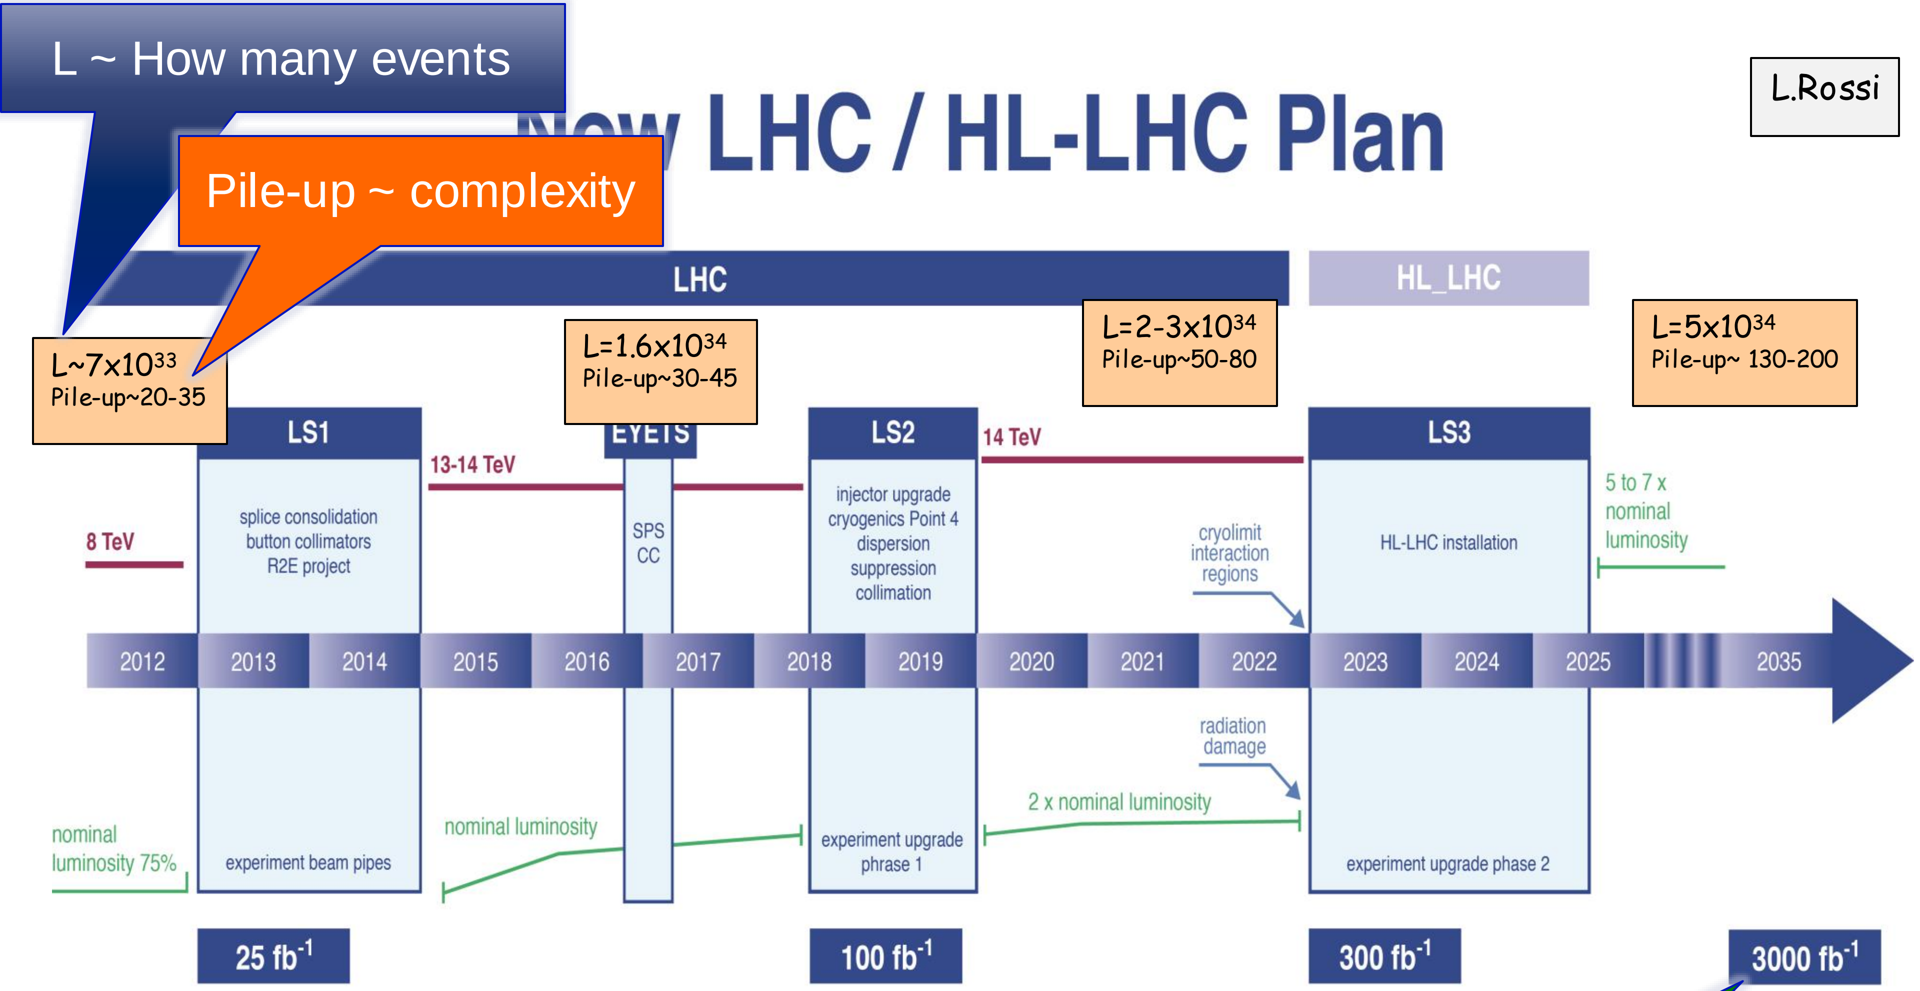

Computing needs increase nonlinear with complexity

~ Accumulated Data

Data: Outlook for HL-LHC

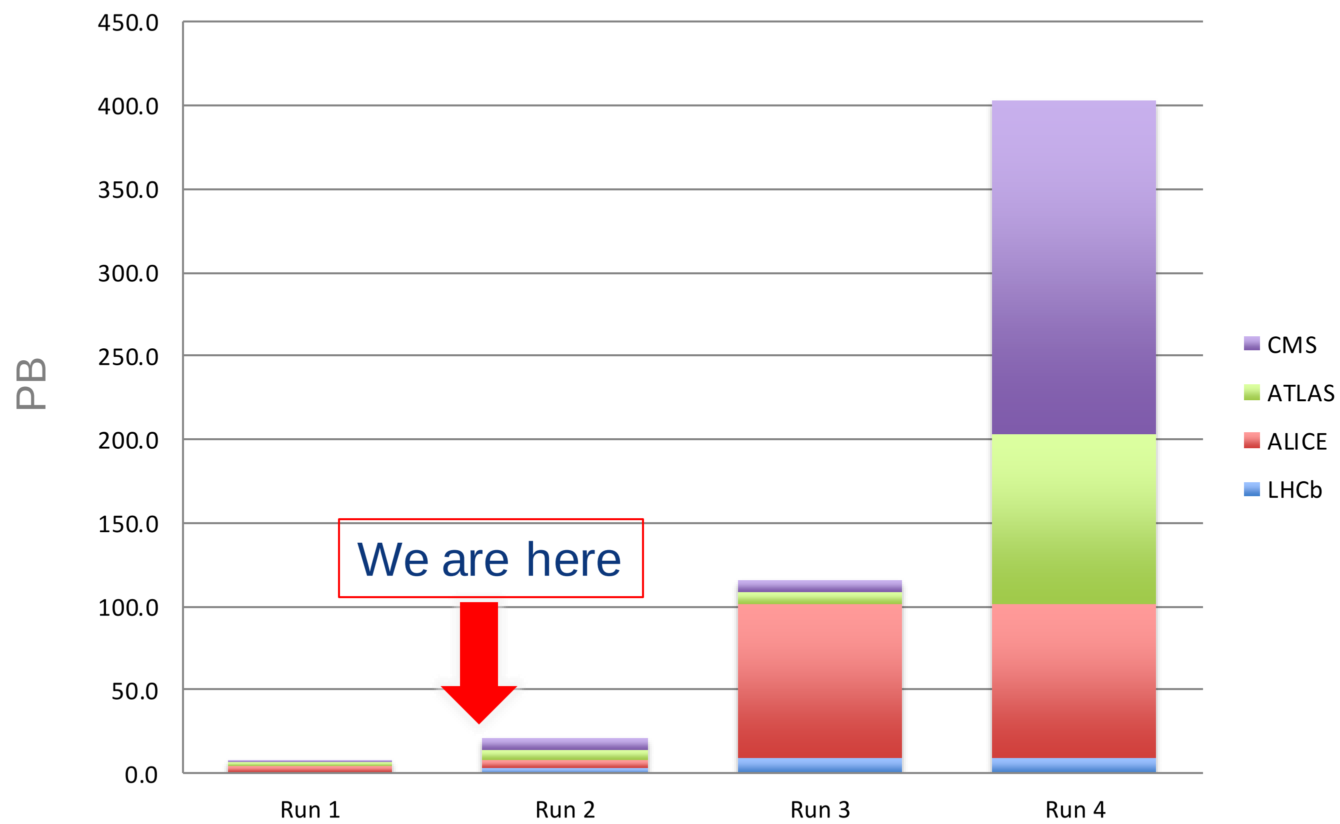

- Very rough estimate of a new RAW data per year of running using a simple extrapolation of current data volume scaled by the output rates.
  - To be added: derived data (ESD, AOD), simulation, user data…

CPU: Online + Offline

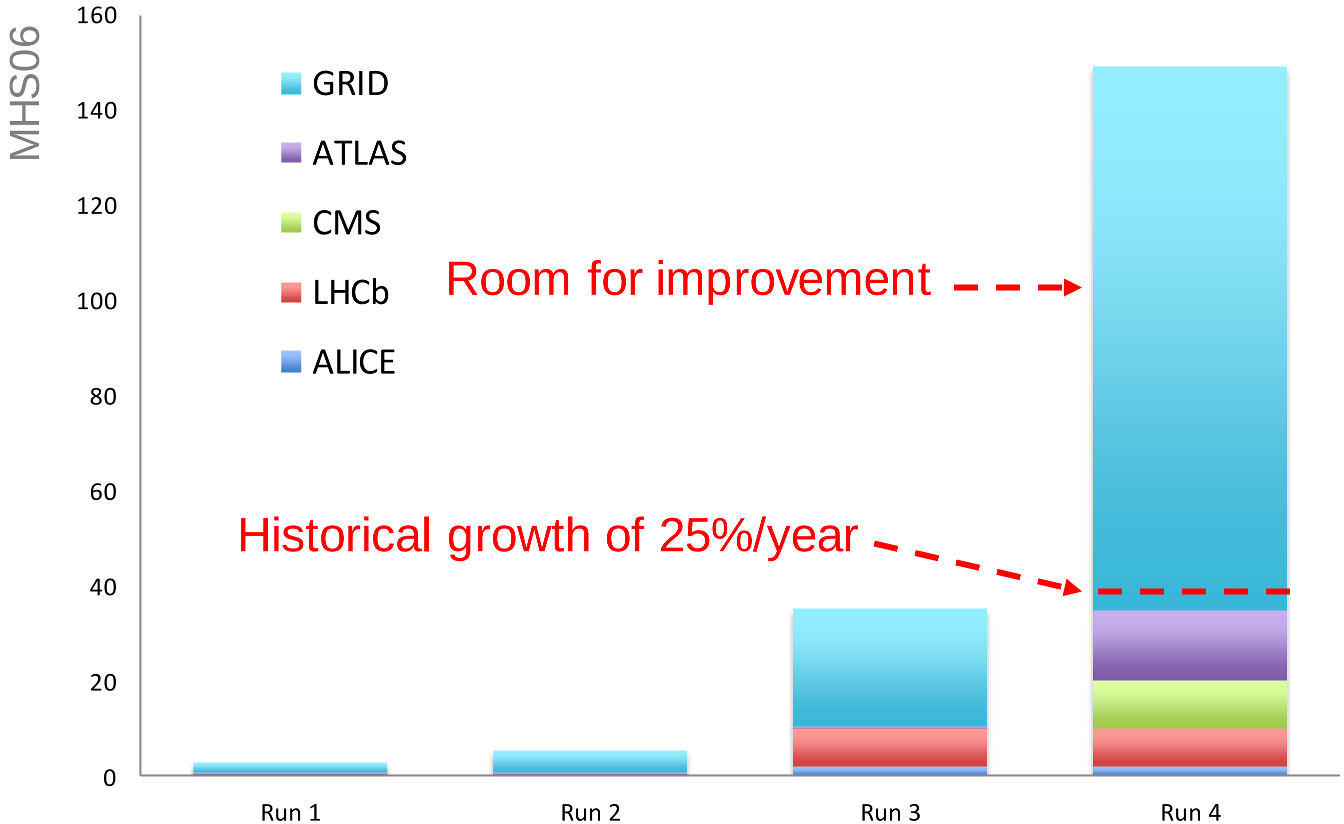

- Very rough estimate of new CPU requirements for online and offline processing per year of data taking using a simple extrapolation of Run 1 performance scaled by the number of events.

Not only physics

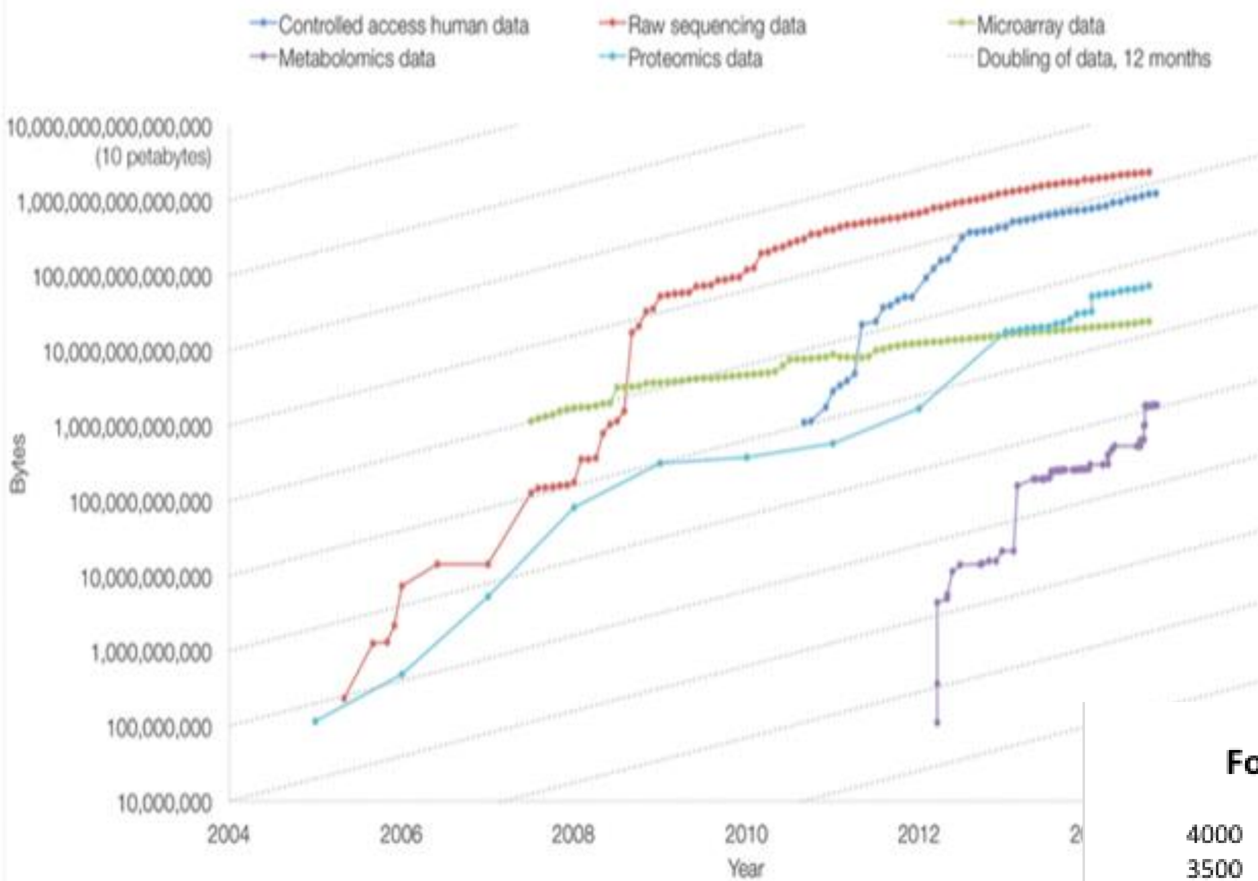

Growth of EBI repositories, lines are 12 month doubling

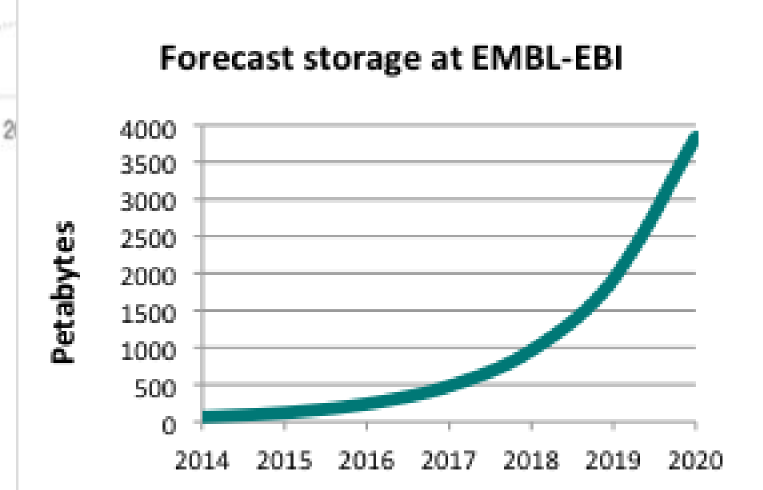

Evolution of the computing models

Evolution of Computing models – enabled by networks

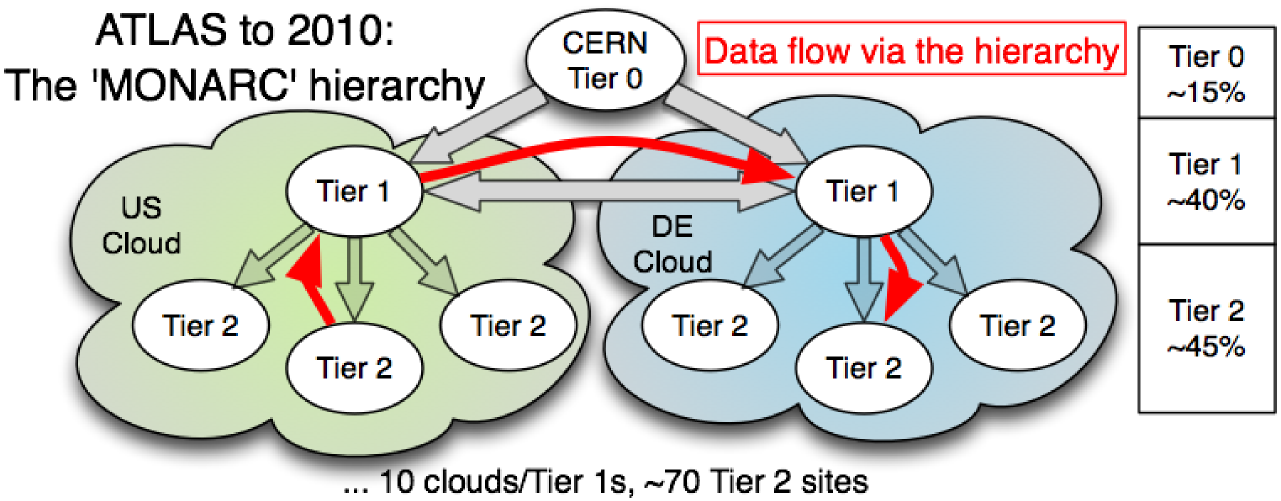

Original model:

Static strict hierarchy

Multi-hop data flows

Lesser demands on Tier 2 networking

Virtue of simplicity

Designed for <~2.5 Gb/s within the hierarchy

Today:

Bandwidths 10-100 Gb/s, not limited to the hierarchy

Flatter, mostly a mesh

Sites contribute based on capability

Greater flexibility and efficiency

More fully utilize available resources

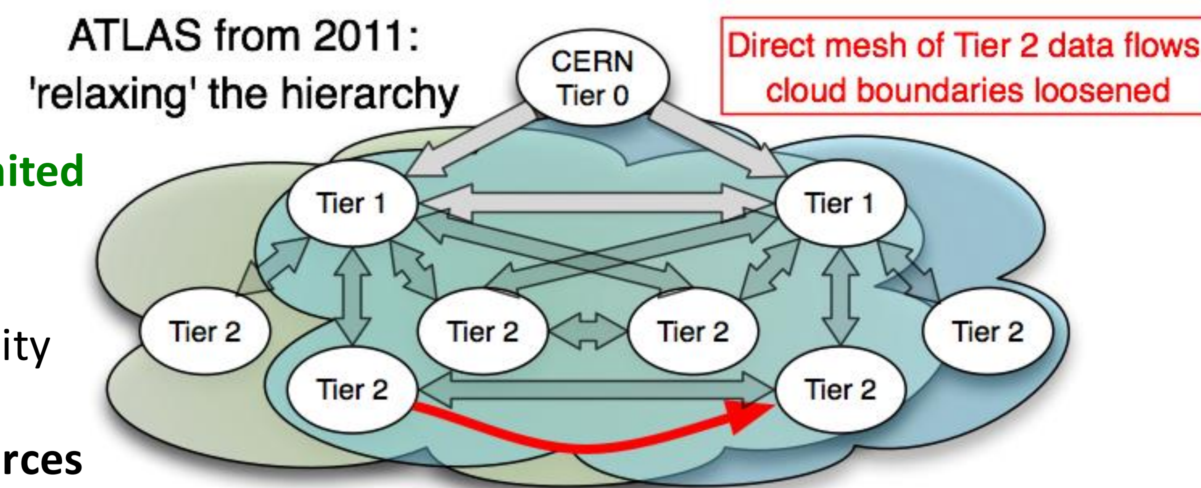

Networking growth has been dramatic

US ESnet as an example

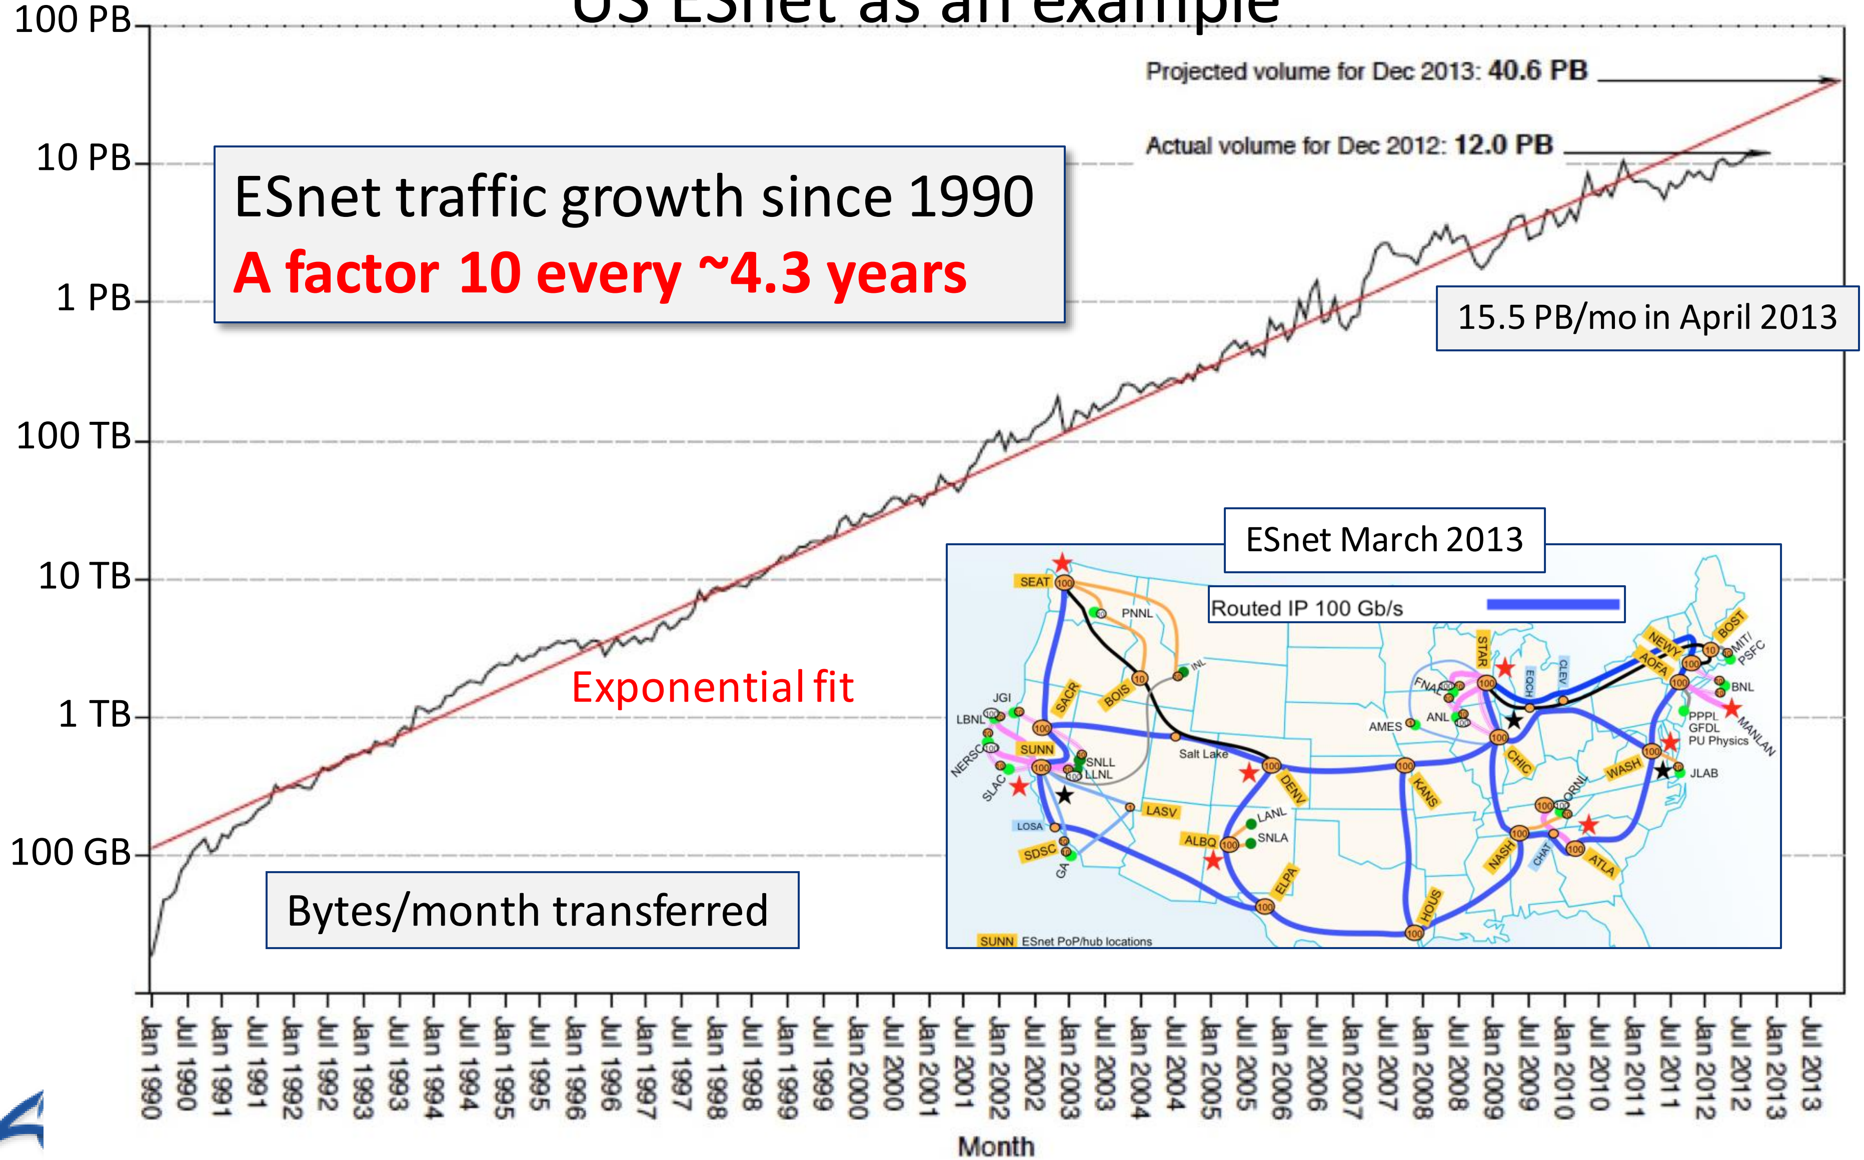

Data federations

- Wide area access to data
  - Access any data from any site without the need to first copy it
  - Optimizing data access from jobs: remote access, remote I/O
  - Performance is good – more intelligence and better caching at all levels
- More intelligent data placement/caching; pre-placement vs dynamic caching
- Federated storage important for big data
  - Distributed management and uniform access – preserves administrative autonomy and is inherently scalable (scale-out)

ATLAS FAX Infrastructure

(From Rob Gardner)

WLCG

- Provides a global namespace
- Unifies dCache, DPM, Lustre/GPFS, Xrootd storage backends
- Xrootd an efficient protocol for WAN access
- Main Fall-back use case **in production** at many sites
- Regional redirection network provides lookup scalability

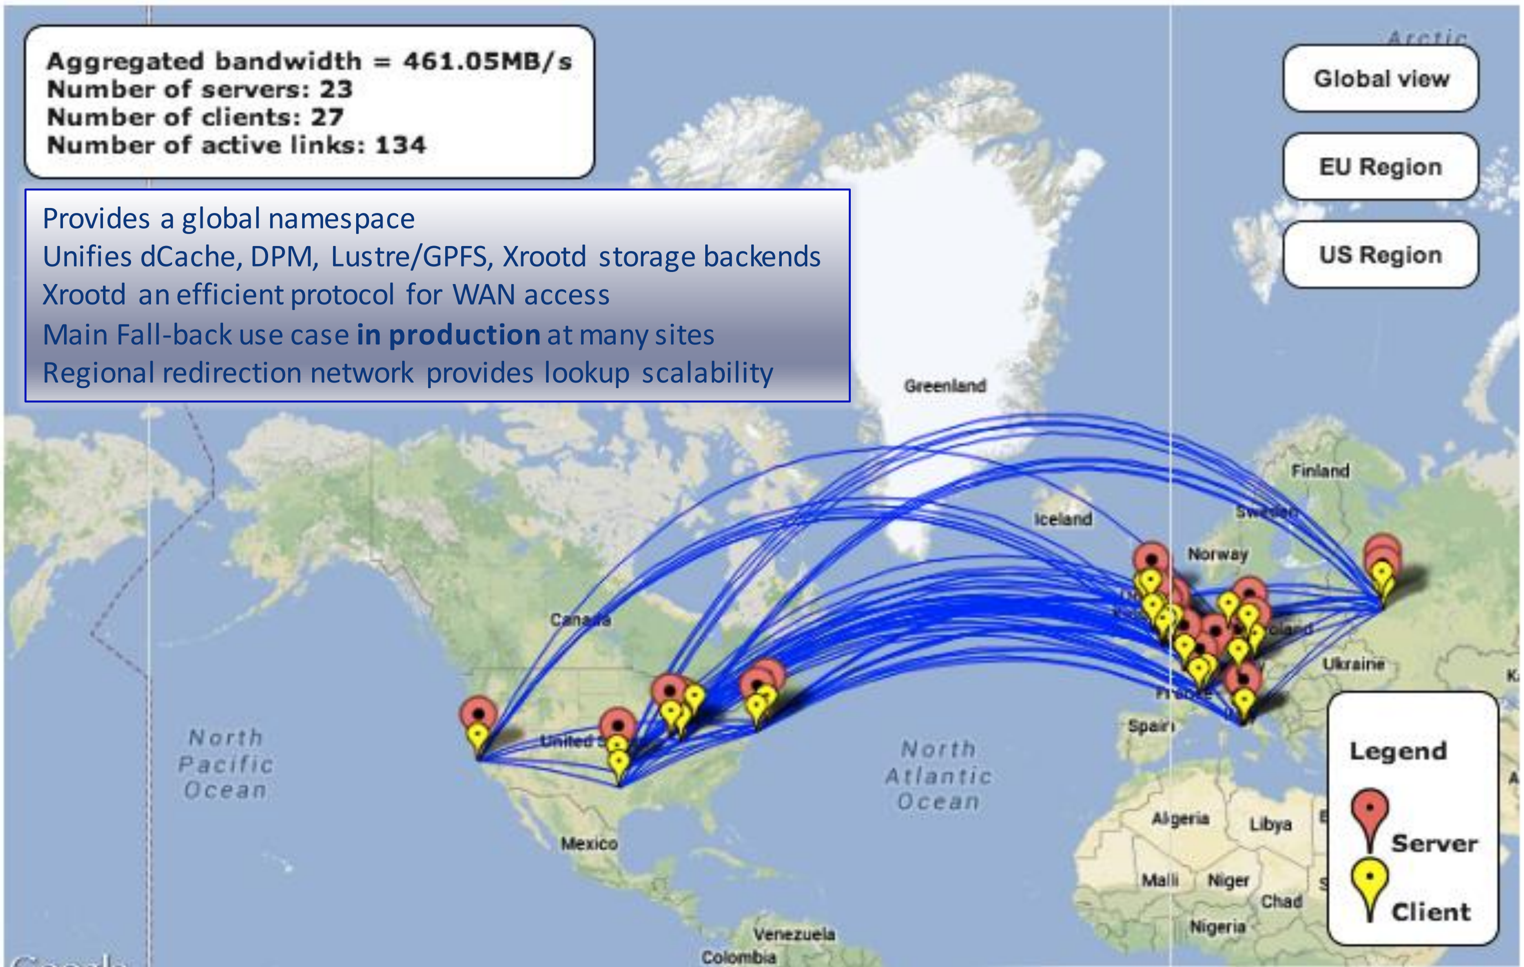

Today uses "xrootd" protocol – (and HTTP)

Drivers of change

WLCG

- Must reduce the (distributed) provisioning layer of compute to something simple, we need a hybrid and be able to use:
  - Our-own resources
  - Commercial resources
  - Opportunistic use of clouds, grids, HPC, volunteer resources, etc.
- Move towards simpler site management
  - Reduce operational costs at grid sites
  - Reduce “special” grid middleware support cost
- The remote data capabilities of the data federation allows us to separate the use of opportunistic compute from the need to distribute data everywhere
- Today (2015) it is cheaper for us to operate our own data centres
  - We use 100% of our resources 24x365
- We also get a large synergistic set of resources in many Tier 2s – essentially for “free” – over and above the pledged resources
- However, commercial pricing is now getting more competitive
  - Large scale hosting contracts, commercial cloud provisioning

Compute resources – Grid

- LHC has a **federated**, globally **distributed**, computing system
  - A "Grid" by definition; autonomous resource provisioning and operation
- Until now the middleware used to implement that has been mostly developed and supported by HEP, and grid projects funded by national and international funding agencies
  - No industrial take up, no global support, etc
    - When we started there was no large scale computing infrastructure or tools (no Amazon, Google, Facebook,…)
- **Federated use of Cloud technologies give us an alternate implementation of compute provisioning**
  - Huge support community
  - Industrial-strength tools tested at scales larger than ours (hmmm … well mostly…)

Why Clouds?

- Clouds offer flexibility
  - user workloads and system requirements are decoupled
  - dynamic allocation of resources
  - commercial and non-commercial providers
- Based on established, open technology and protocols
  - expertise is widely available
  - products and tools evolve rapidly
  - commercial and non-commercial users
- Proven scalability
  - small in-house systems to world wide distributed systems

Clouds in LHC

- CERN and many WLCG sites are now using cloud technologies to provision their compute clusters
  - Together with "devops" toolchains to manage the scale we are now at
  - Many are deploying Openstack – global community
- Cloud provisioning
  - Better cluster management and flexibility
  - Can run existing grid services on top – but don't really need to
- LHC experiments also manage HLT farms with Openstack
  - Allows them to switch between DAQ and processing

Openstack - Federation

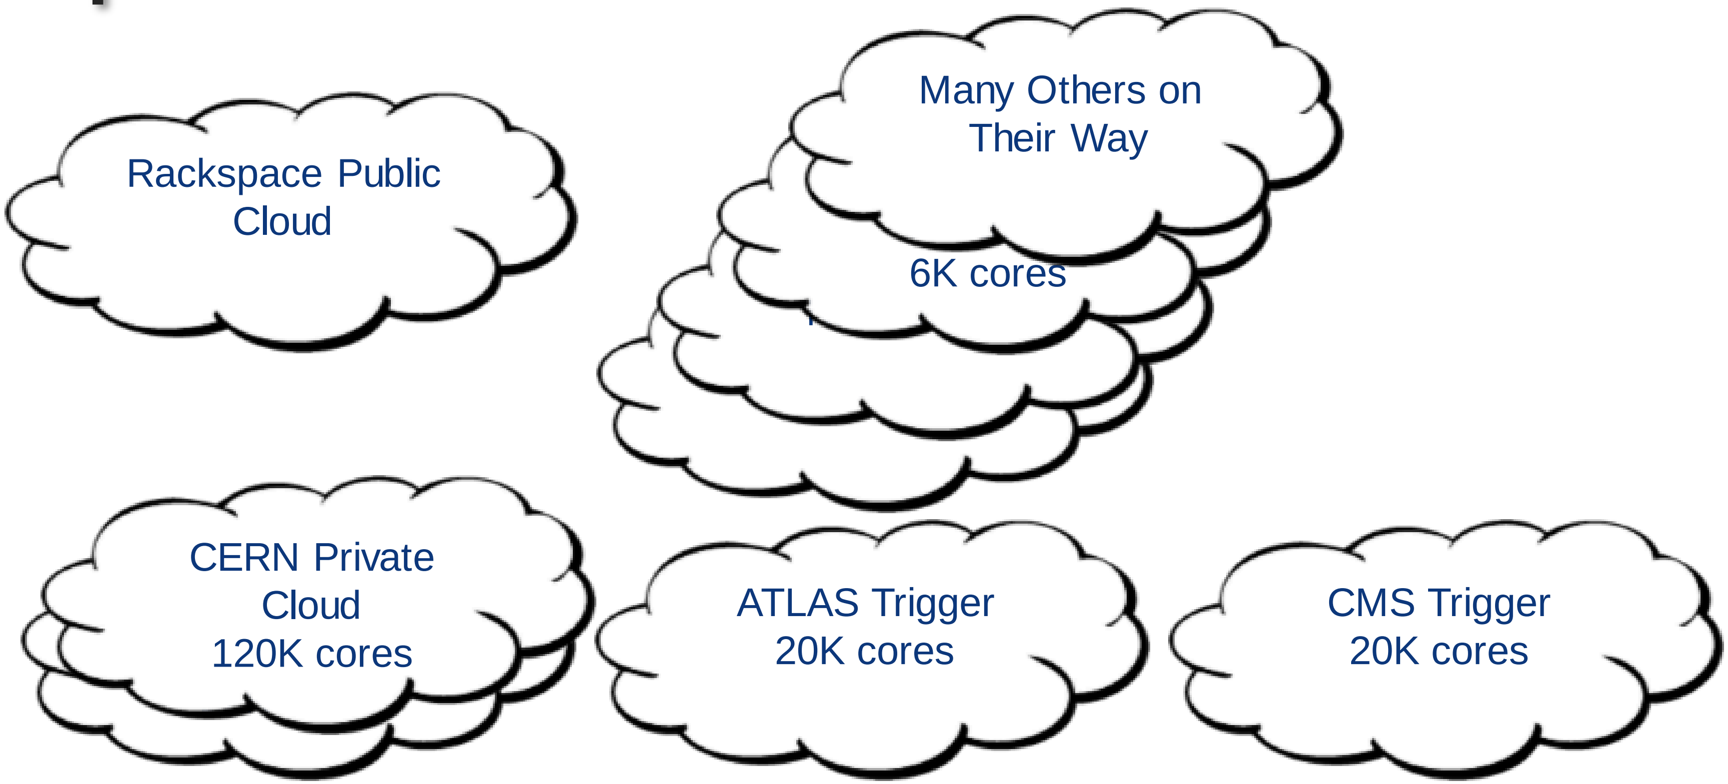

- Federate clouds; based on EduGain identities
- In collaboration with Rackspace in CERN-openlab
- All contributions are to OpenStack upstream so will appear in all OpenStack clouds at all the sites

A European Cloud Computing Partnership: big science teams up with big business

Strategic Plan

- Establish multi-tenant, multi-provider cloud infrastructure
- Identify and adopt policies for trust, security and privacy
- Create governance structure
- Define funding schemes

CERN: To support the computing capacity needs for the ATLAS experiment

EMBL: Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity

ESA: To create an Earth Observation platform, focusing on earthquake and volcano research

Email: contact@helix-nebula.eu Twitter: HelixNebulaSC Facebook: HelixNebula.TheScienceCloud

Cloud

VMs on demand

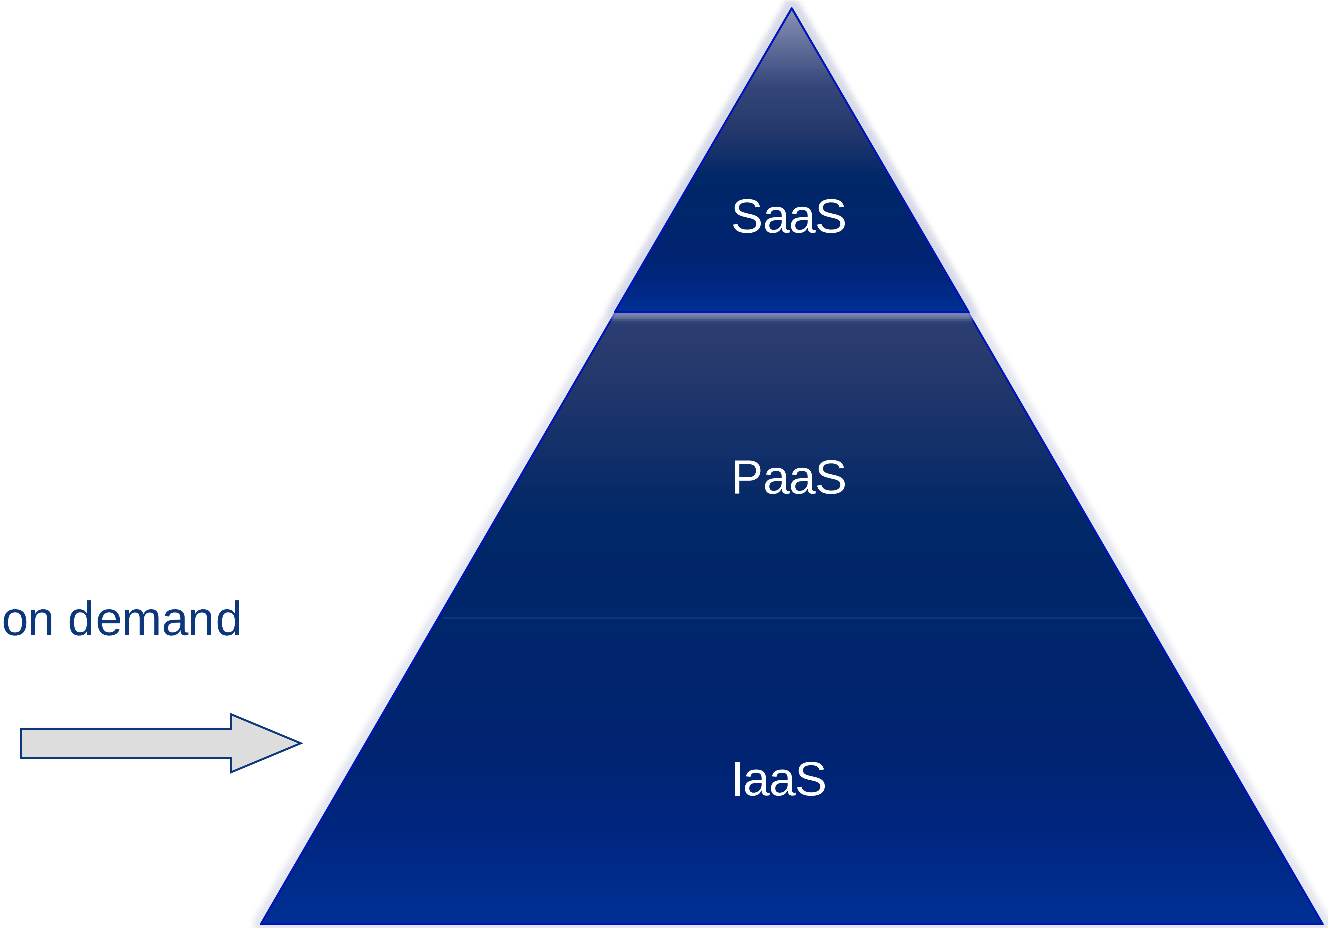

Grid vs Clouds

- Grids
  - abstraction for **Services**
    - Batch, Storage...
    - high level, huge variety of services
  - management of communities
    - **Virtual Organisations (VO)**
  - **Provider Centric**
    - **monitoring, accounting, security model, quotas**
- **Clouds**
  - abstraction for Infrastructure (IaaS)
  - low level services
    - CPU, object store,....
  - no management of communities
  - high level services VO centric
    - Workflow, accounting, quotas, security
  - User Centric!
    - users have to organise workflows, accounting, conceptualisation, monitoring, sharing.....

Areas to be addressed

- Image Management
- Capacity Management
- Monitoring
- Accounting (Provider and Client)
- Pilot Job Framework
- Data Access and Networking
- Quota Management
- Supporting Services

Image Management

- Provides the job environment
  - Software
    - CVMFS
    - PilotJob
  - Configuration
  - Contextualization
- Balance pre- and post-instantiation operations
  - Simplicity, Complexity, Data Transfer, Frequency of Updates
- Transient
  - No updates of running machines
    - Destroy (gracefully) and create new instance

Capacity Management

- Managing the VM life cycle isn't the focus
  - It is about ensuring there are enough resources (capacity)
- Requires a specific component with some intelligence
  - Do I need to start a VM and if so where?
  - Do I need to stop a VM and if so where?
  - Are the VMs that I started OK?
- Existing solutions focus on deploying applications in the cloud
  - Different components, one cloud
  - May manage load balancing and failover
  - One configuration, many places, enough instances?
- Developing our own solutions
  - Site centric
    - The VAC model
  - VO centric

Monitoring

- Fabric management
  - The responsibility of the VO
  - Basic monitoring is required
- The objective is to triage the machines
  - Invoke a restart operation if it not ok
    - Detection of the not ok state maybe non-trivial
- Other metrics may be of interest
- Spotting dark resources
  - Deployed but not usable
- Can help to identify issues in other systems
  - Discovering inconsistent information through cross-checks
- A Common for all VOs
  - Pilot jobs monitoring in VO specific

Provider Accounting

- Helix Nebula
  - Pathfinder project
    - Development and exploitation
      - Cloud Computing Infrastructure
  - Divided into supply and demand
  - Three flagship applications
    - CERN (ATLAS simulation)
    - EMBL
    - ESA
- **FW: New Invoice!**
  - *Can you please confirm that these are legit?*
  - Need to method to ***record*** usage to cross-check invoices
  - Dark resources
    - Billed for x machines but not delivered (controllable)

Consumer-Side Accounting

- Monitor resource usage
  - Course granularity acceptable
    - No need to accurately measure
- What, where, when for resources
  - Basic infrastructure level
    - VM instances and whatever else is billed for
- Report generation
  - Mirror invoices
    - Use same metrics as charged for
- Needs a uniform approach
  - Should work for all VOs
    - Deliver same information to the budget holder

Provisioning & monitoring chain

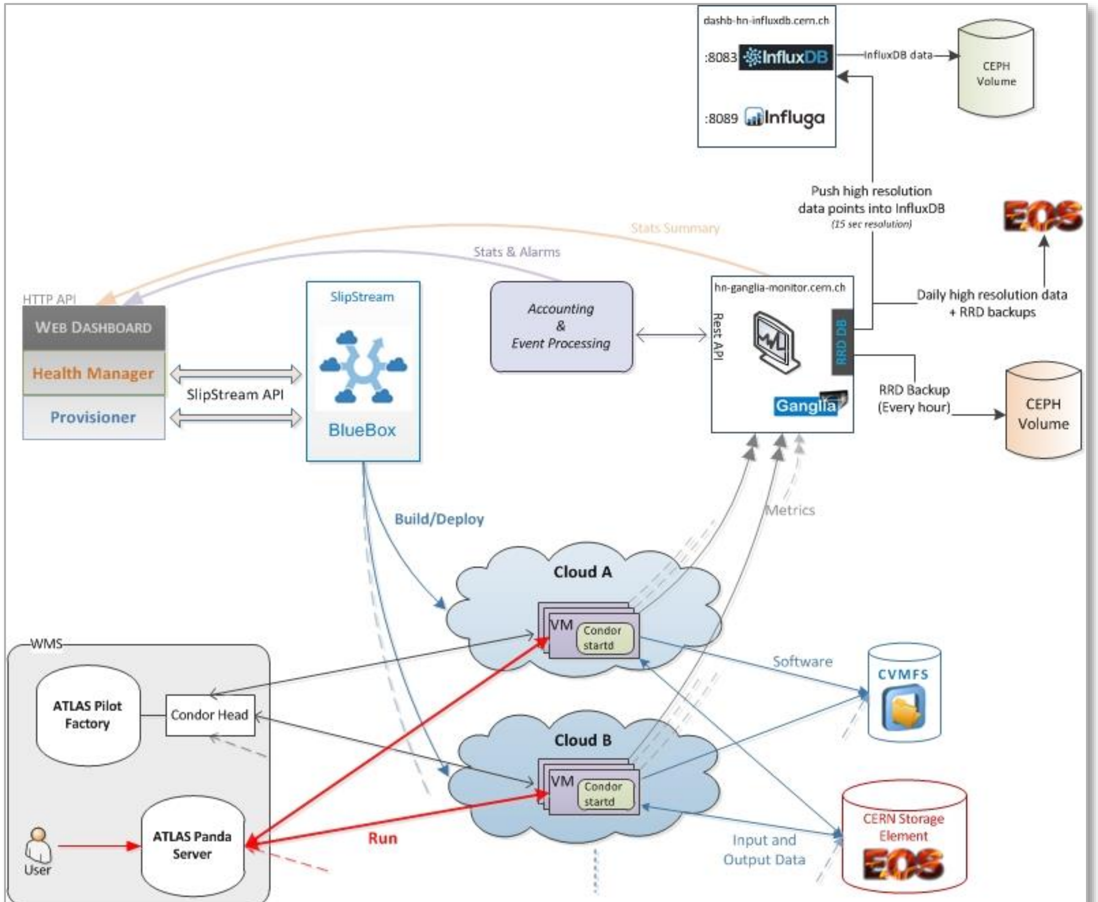

- Key role of VM monitoring
  - Real-time monitoring
  - Alarming
  - Accounting
  - Benchmarking
- Strategy
  - Ganglia data preserved with **15s** time resolution
  - **Benchmark** each VM at startup

Benchmarking

- Each created VM has been benchmarked using ATLAS KitValidation
  - ~30,000 VM benchmark performed
  - 100 Single Muon events simulated (**~2 min** to run)
- Results
  - CPU performance uniform within **15%** spread
  - Benchmark profile consistent over time

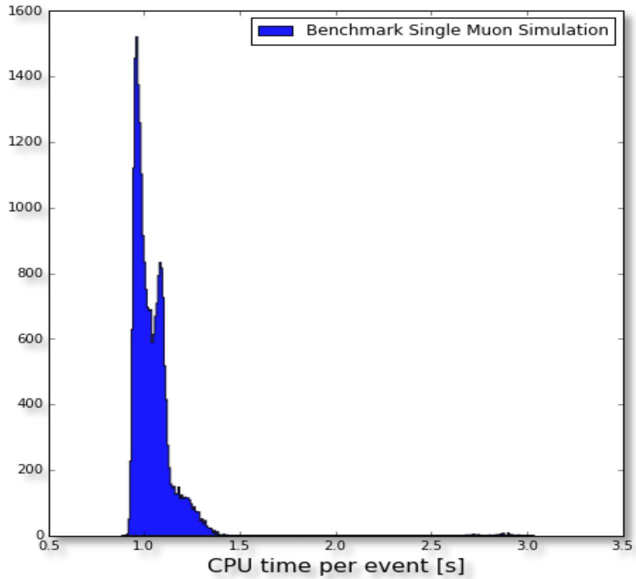

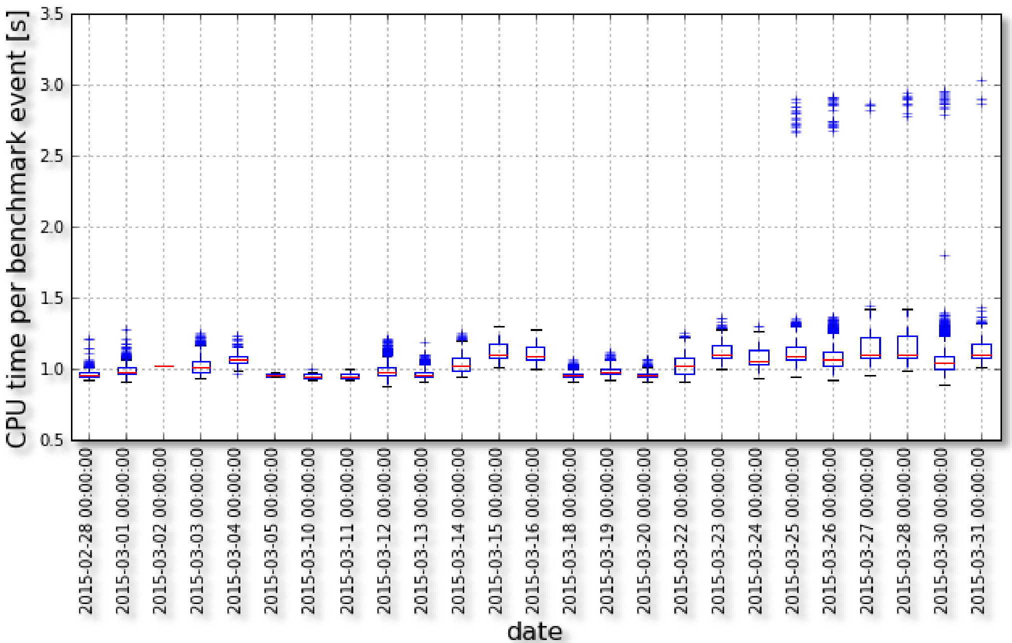

Benchmarks vs Job performance

- Consistent job CPU performance and benchmark
  - Correlated behavior
- Outliers detection
  - KV bmk (2') is a *prompt and effective solution* to identify VMs with poor performance

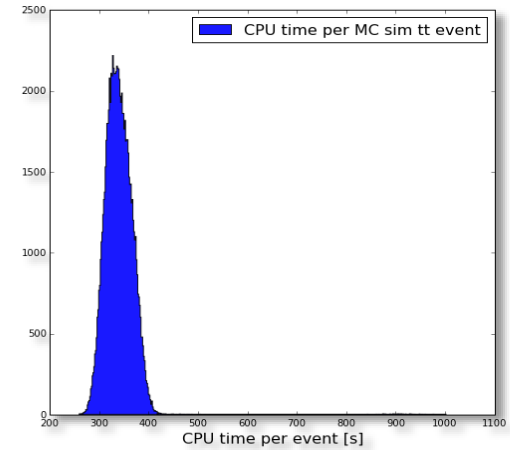

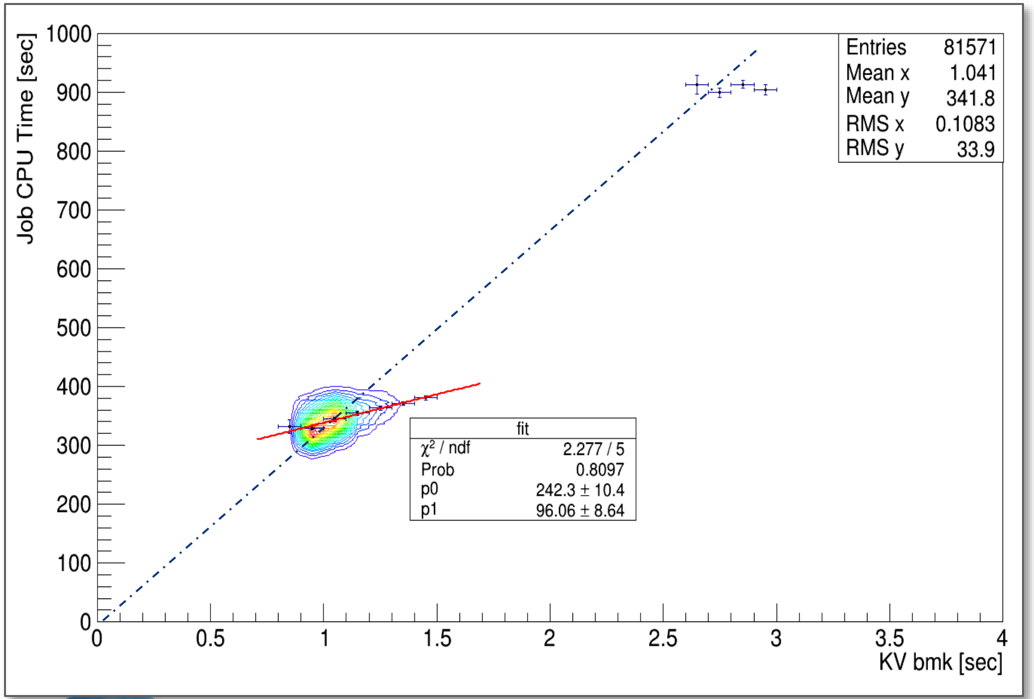

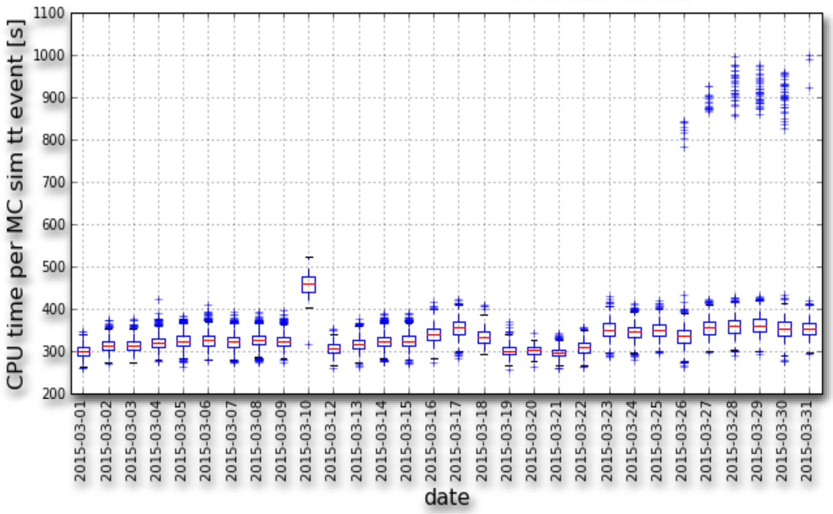

Other Considerations

- Data access and networking
  - Have so far focus on non-data intensive workloads
- Quota Management
  - Currently mostly fixed limits
    - Leading the partitioning of resources between VOs
      - How can the sharing of resources be implemented?
- Supporting Services
  - What else is required?
  - Eg squid caches in the provider
  - How are these managed and by who?
- Non-Virtualized approaches
  - Instantiation of a pilot job
    - Without CE

Commercial Clouds

- Helix Nebula
  - A public-private partnership
    - Between research organizations and IT industry
- Microsoft Azure Pilot
  - Preliminary discussions with CERN OpenLab
- Amazon
  - BNL RACF for ATLAS and CMS
  - With new Scientific Computing group at AWS
- Deutsche Börse Cloud Exchange AG
  - Beta testing platform
  - Will go live beginning of May
- PICSE
  - Procurement Innovation for Cloud Services in Europe
- European Science Cloud Pilot
  - Pre-Commercial Procurement (PCP) proposal
    - Buyers group public organizations that are members of the WLCG collaboration

Example: Commercial Cloud Test

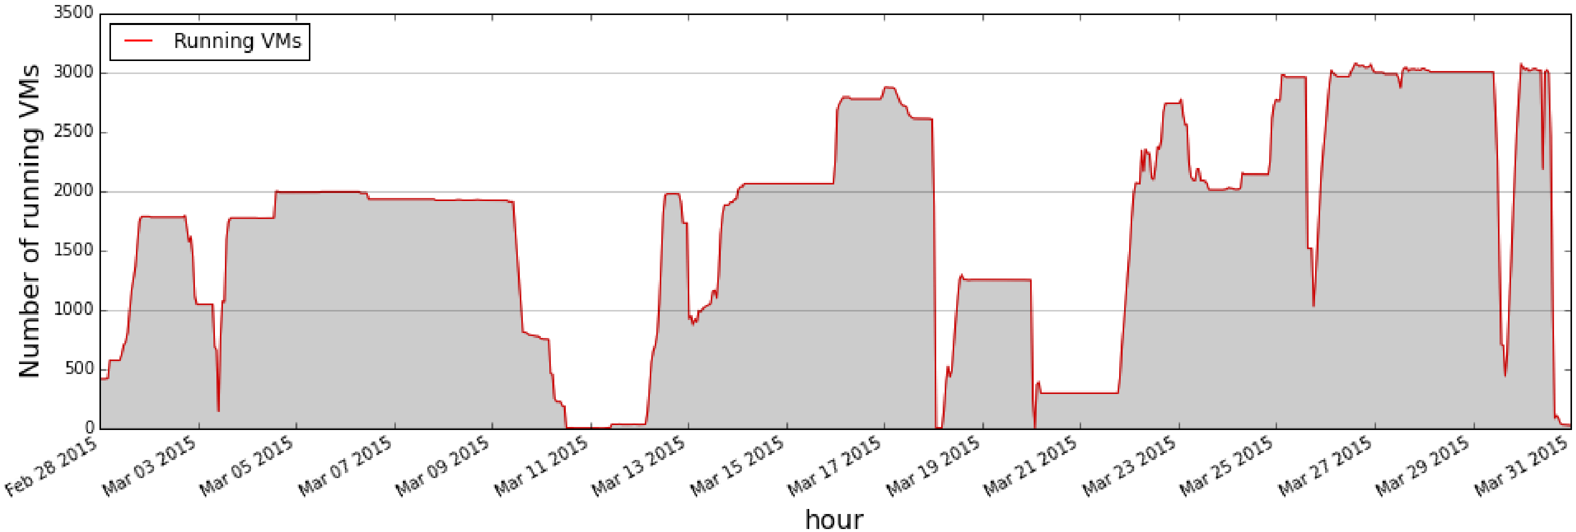

- Up to **3,000** concurrent running VMs
  - 4 (+1) weeks of production
  - ~1.2 million CPU hours of processing

- ATLAS GEANT4 Simulation of $t\bar{t}$ events
  - ~11.5 million events processed ⇔ ~160,000 files produced
  - ~93% CPU/Wall time ratio
    - ~9 hours single job duration
  - ~97% job wall time used for successful runs
    - Lost heartbeat is the main source of failures (~81%)

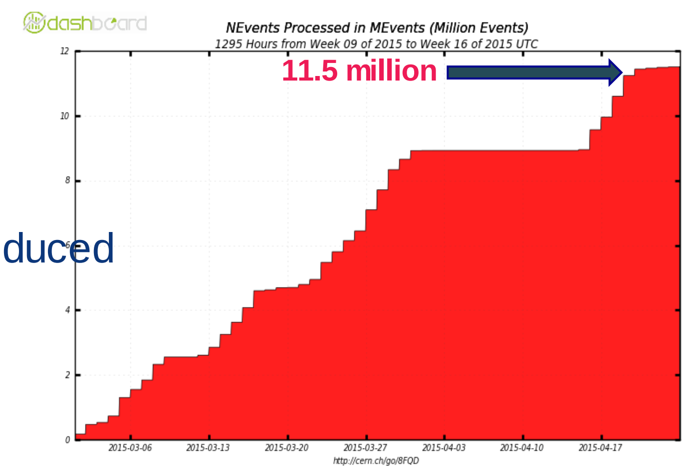

Compared with other ATLAS cloud sites (March)

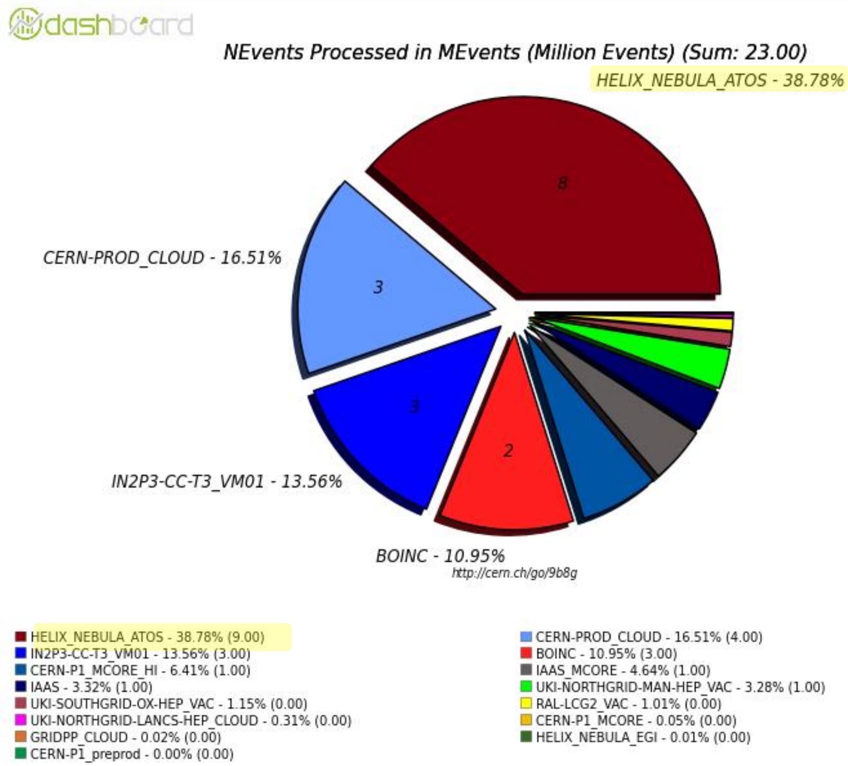

- Significant contribution compared with other ATLAS cloud sites running simulation
  - **Largest** # of processed events
  - **Longest** wall time consumption
  - **High** wall time efficiency

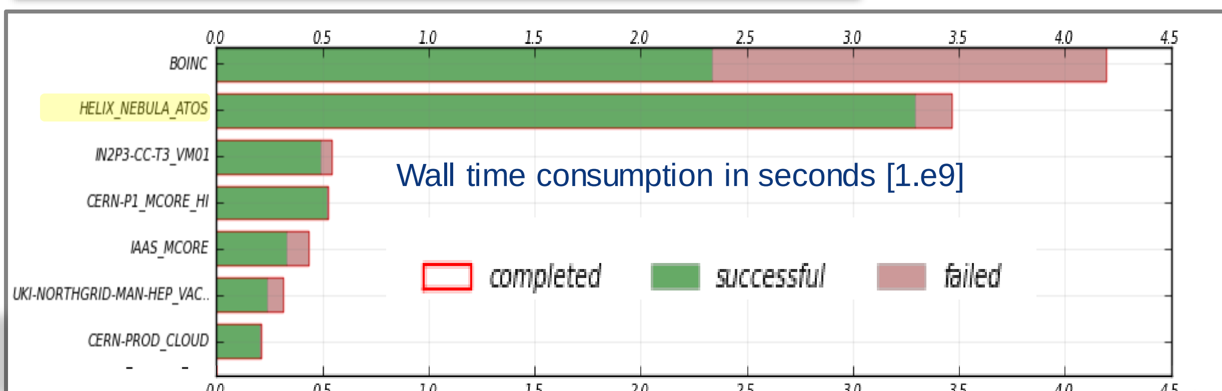

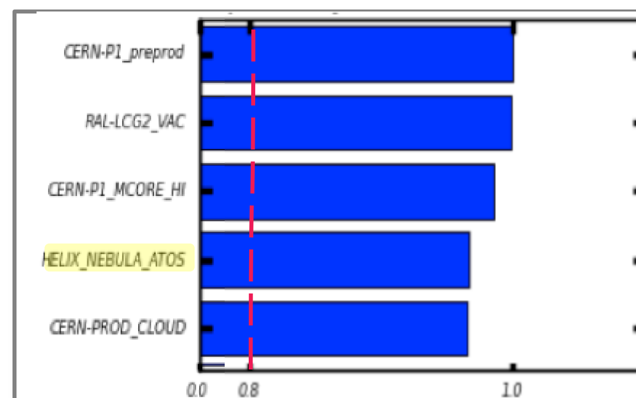

The monthly activity

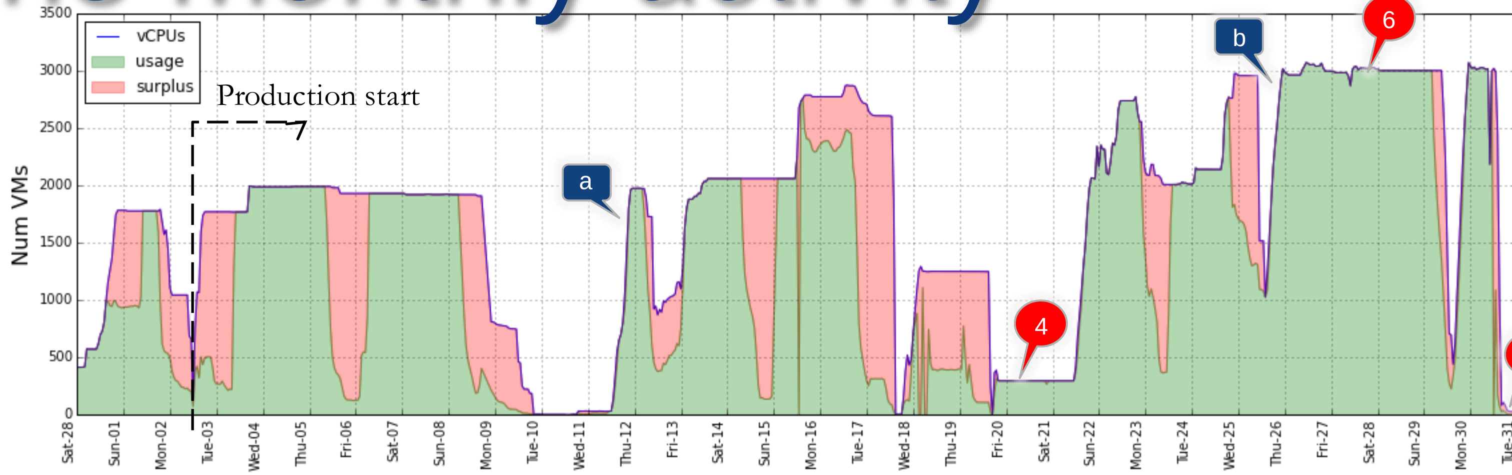

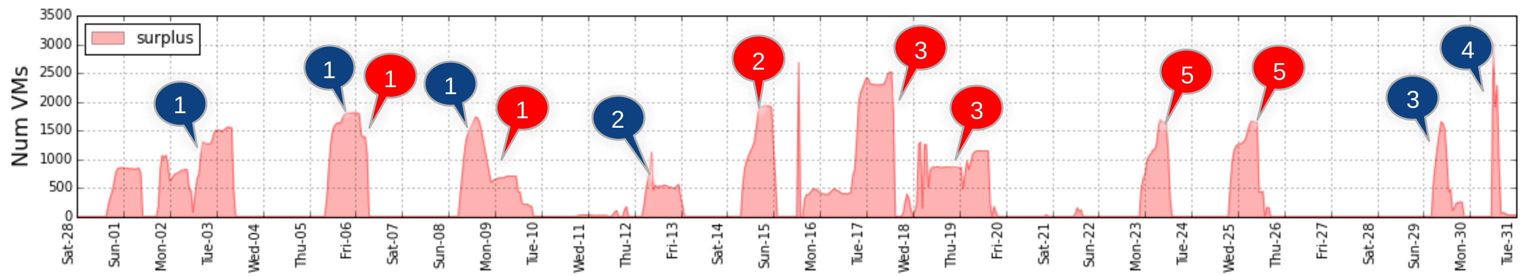

Reasons

CERN

1. Task completed (no more jobs)
2. CERN network issues
3. Agent auth. cache not renewed
4. Task abruptly terminated

a Improvement: auto-scaling (up/down) based on load

b Improvement: Orchestrator-less single-VM runs

Resource usage:

- Effective: **77%**
  - Surplus causes 10% CERN,13% IaaS
- After improvement "b":
  - Effective: **93%**

IaaS

1. VMs stuck in provisioning
2. Cloud layer reports zero VM running
3. Read-only file system
4. Stuck orchestrators: 4 faulty KVMs
5. Stuck runs: missing status from few VMs/run
6. VMs still alive after run deletion

Opportunistic resources

- Today this has become more important
  - Opportunistic use of:
    - HPC's (backfill)
    - Large cloud providers
    - Other offers for "off-peak" or short periods
    - Etc.
    - All at very low or no cost (for hardware)
  - But scale and cost are unpredictable
- Also growing in importance:
  - Volunteer computing (citizen science)
    - BOINC-like (LHC@home, ATLAS/CMS/LHCb@home, etc)
    - Now can be used for many workloads – as well as the outreach opportunities

Volunteer Computing

Volunteer Computing

- A type of distributed computing
- Origins in mid 1990s
- Computer owners *donate* computing capacity
  - To a *cause* or project
- Not necessarily only spare cycles on desktops
  - Idle machines in data centers
  - Home clusters
- SETI@home and Folding@home
  - Launched 1999
- CERN runs LHC@home
  - http://lhcathome.web.cern.ch/

BOINC

- Berkeley Open Infrastructure for Network Computing
  - Started in 2002
  - Funded by the National Science Foundation (NSF)
  - Developed by a team based at the Space Sciences Laboratory
    - University of California, Berkeley
    - Led by David Anderson
- Provides the middleware for volunteer computing
  - Client (Mac, Windows, Linux, Android)
  - GUI
  - Application runtime system
  - Server software
  - Project Web site

SETI@HOME

- Search for Extra-Terrestrial Intelligence
- Analyses radio signals
  - Arecibo Observatory in Puerto Rico

- Supporting scientific work
  - Detection intelligent life outside Earth
    - Yielded no conclusive results
      - No evidence for ETI signals
- Viability and practicality of volunteer computing
  - 120K Active Users
  - **180K Active Hosts**

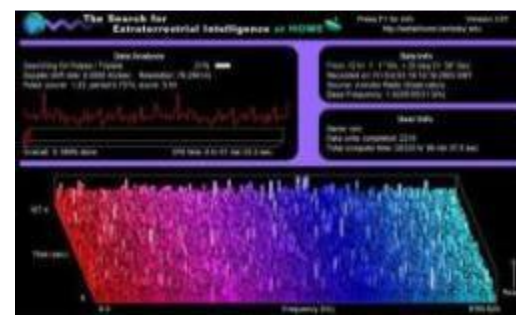

BOINC With Virtualization for LHC

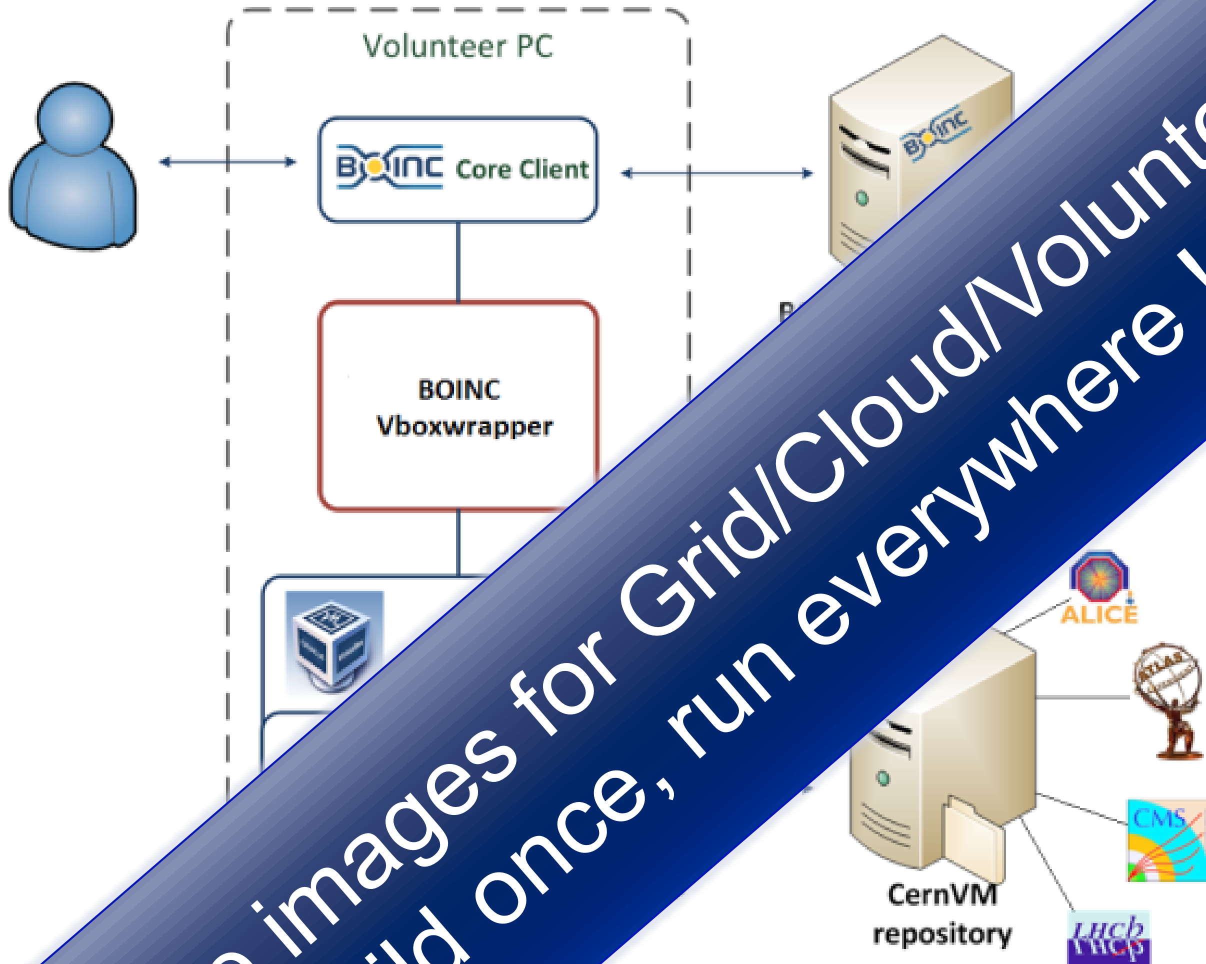

Same images for Grid/Cloud/Volunteer
Build once, run everywhere !

The WLCG DataBridge

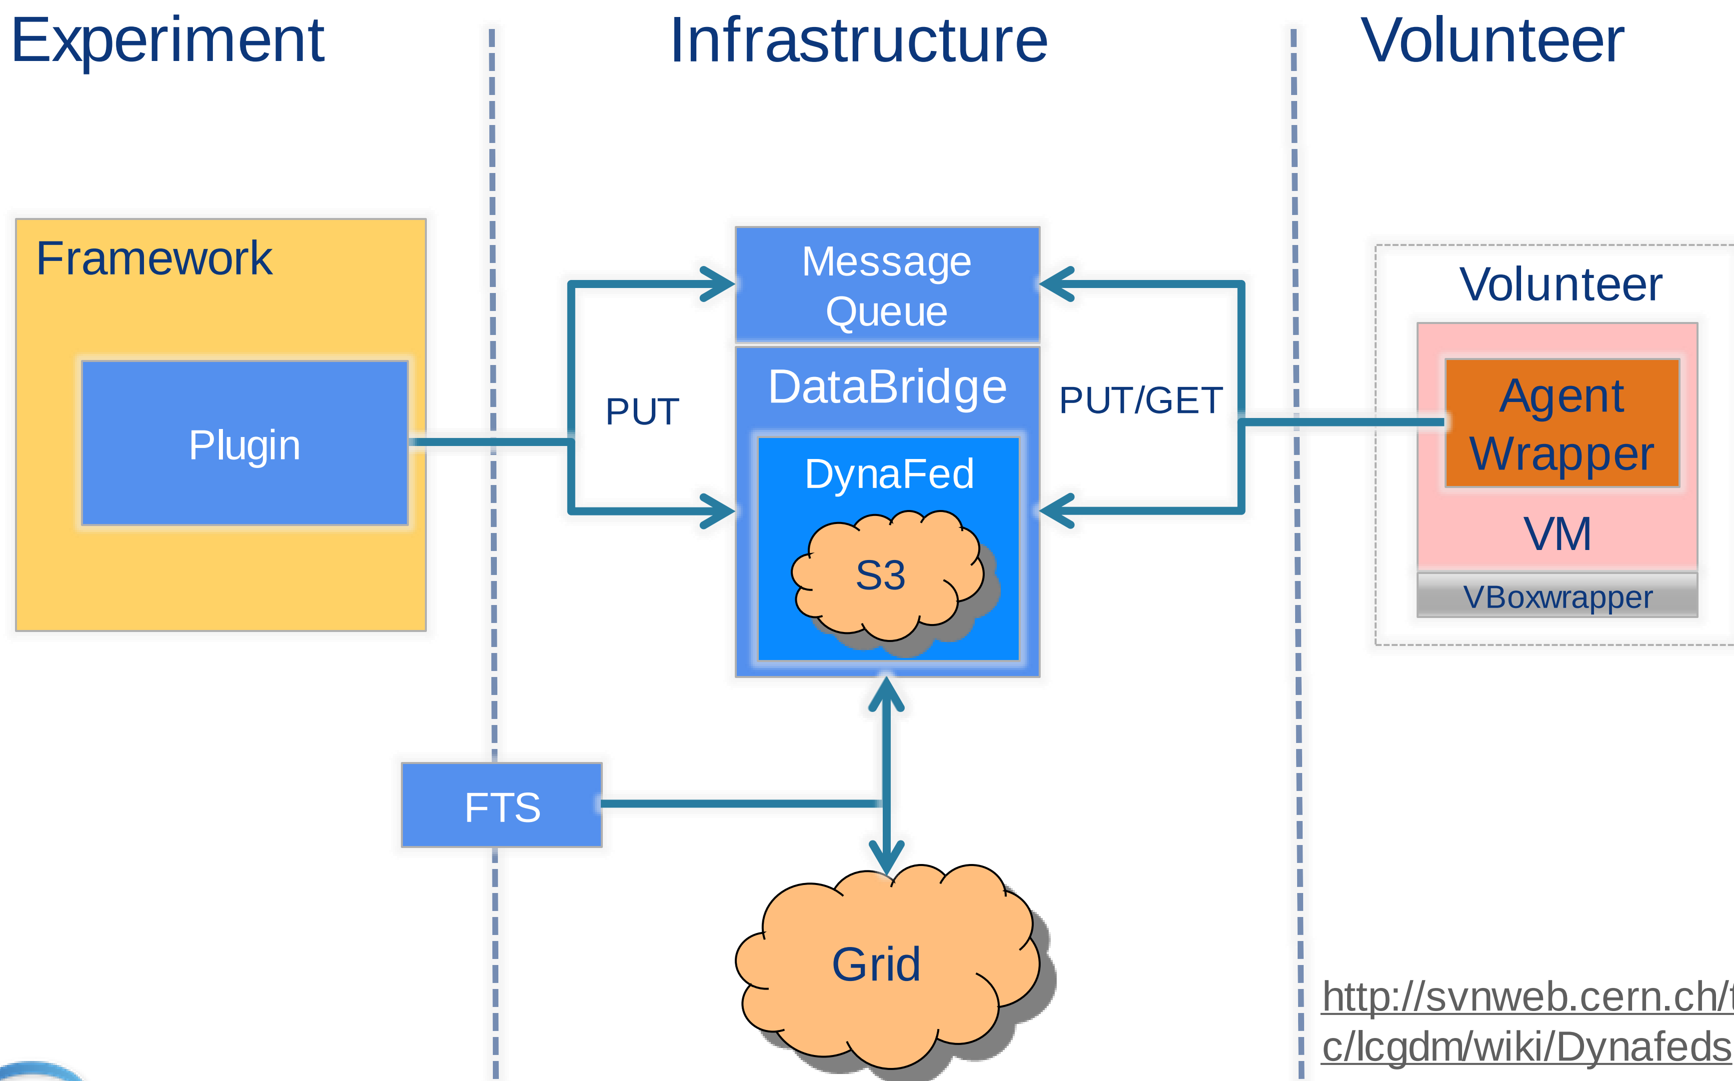

http://svnweb.cern.ch/trac/lcgdm/wiki/Dynafeds

CERN BOINC Service Monitor

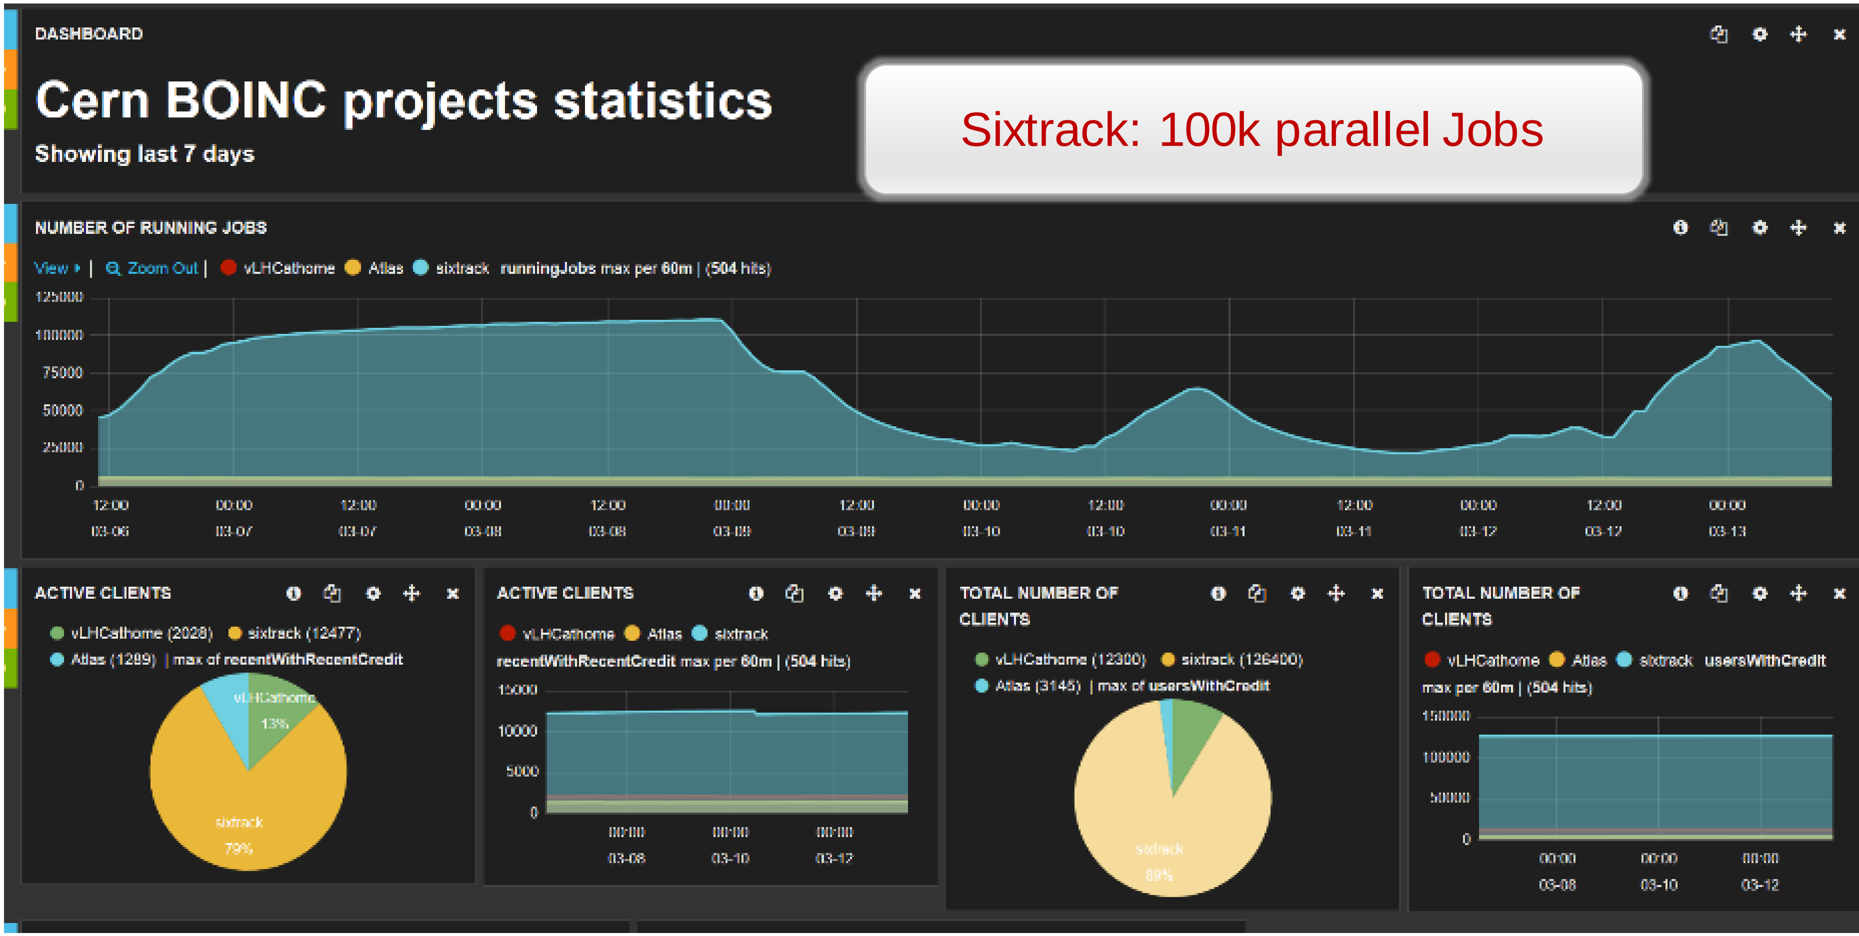

IT-PES http://cern.ch/go/9nRz

SixTrack (LHC machine)

- Original classic BOINC project for beam simulations
  - Calculates stability of proton orbits in the LHC accelerator
    - Simulates particle trajectories
- Based on experience from the Compact Physics Screensaver (CPSS)
  - Ran SixTrack on desktop computers at CERN
- Outreach project for CERN's 50th anniversary 2004
  - Also Year of Physics (Einstein Year) 2005
- Application written in FORTRAN
  - Runs on Linux, Mac and Windows platforms
- Renewed effort for LHC upgrade studies (HL-LHC)
  - 12K Active Users
  - 19K Active Hosts
  - 35 TeraFLOPS

Test4Theory

- Theoretical fitting of all past experimental data
  - Including from the LHC
  - Using Monte Carlo simulation based on Standard Model
- Launched 2011
  - In partnership with the Citizen Cyberscience Centre (CCC)
- Pioneered use of Virtualization with BOINC
- Uses recent developments from CERN's PH-SFT Group
  - CernVM
  - CernVMFS
  - CoPilot
- Wide range of potential (physics) applications
  - In 2014 changed name to Virtual LHC@home

Test4Theory Usage

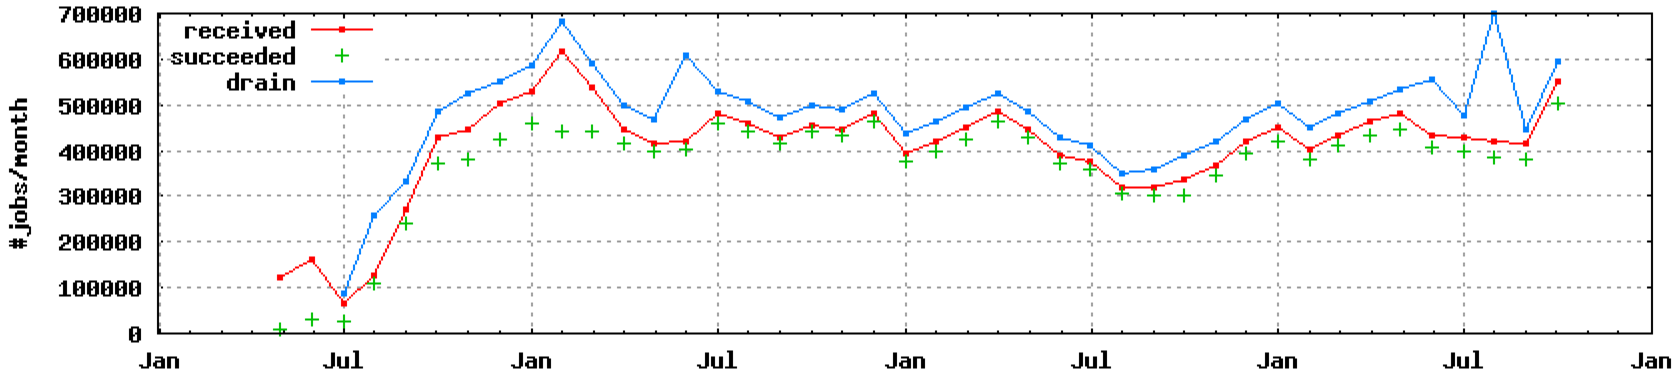

- Total of 1.7 trillion events simulated since 2011
- Source: MC Plots (http://mcplots-dev.cern.ch/production.php)
- See also: http://cern.ch/go/9nRz

ATLAS@home

- Started as pilot beginning of 2014
  - Now open to the public
    - http://atlasathome.cern.ch
- Also using CernVM and virtualization
  - Classic BOINC model
- ARC CE used to interface with BOINC
  - PanDA for job management
- Supports simulations
  - Potentially other types of ATLAS workloads
- Job size and 64bit image limits to “hardcore” volunteers
  - Already significant CPU contribution
- Integrated with LHC@home environment
  - BOINC server hosted by CERN’s IT-PES group
  - ARC-CE and BOINC sharing data via NFS

YEARS/ANS CERN

ATLAS@home Usage

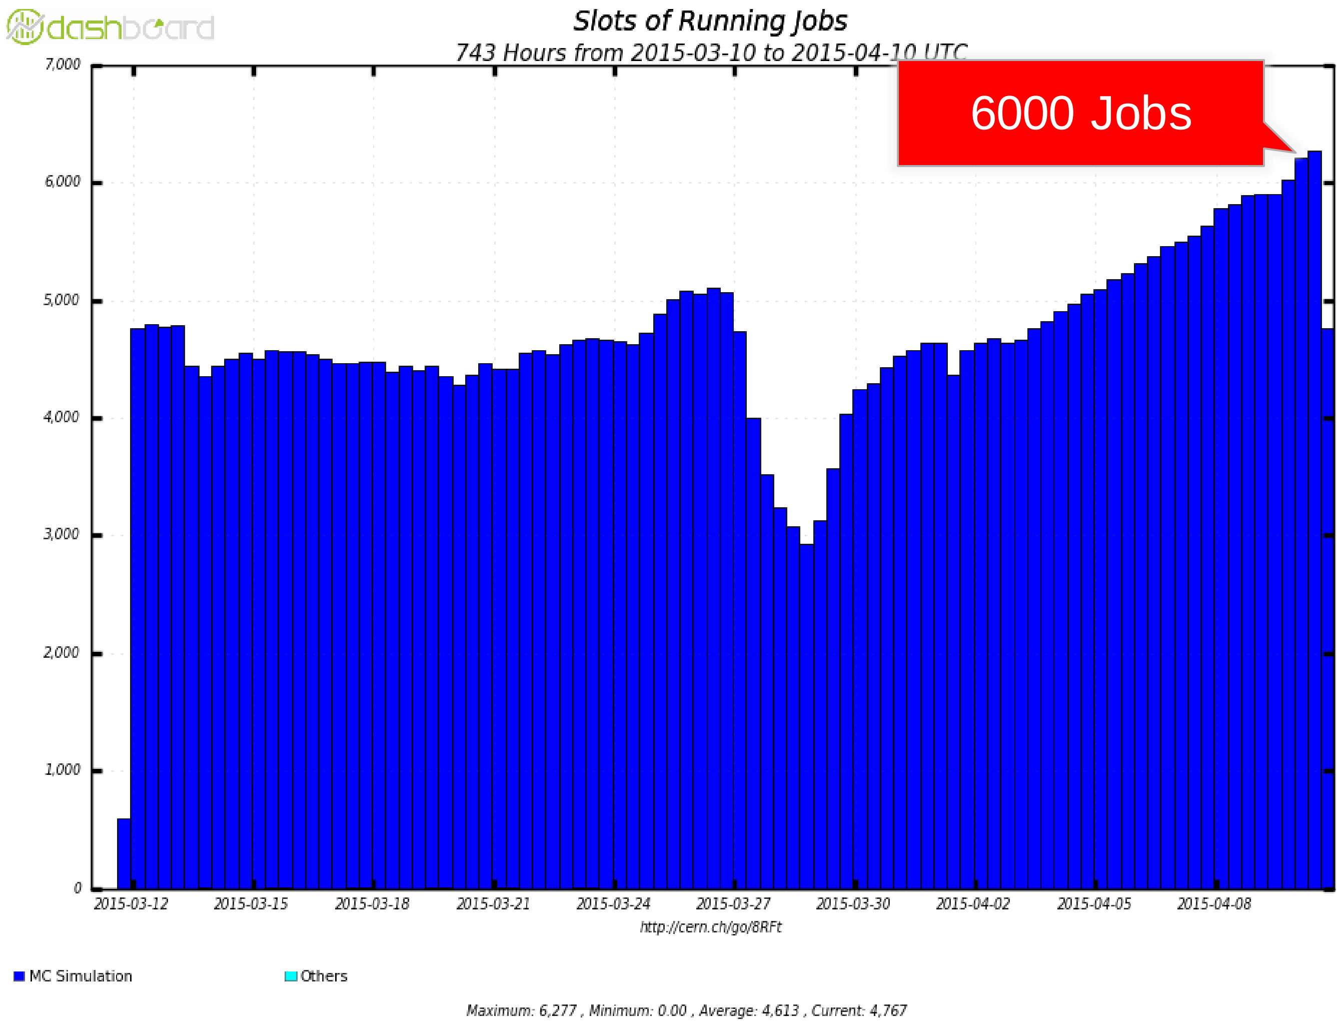

ATLAS@home Contribution

- 2nd largest simulation site
- Running 4-5k parallel jobs
- 20M events simulated
- 5M CPU hours

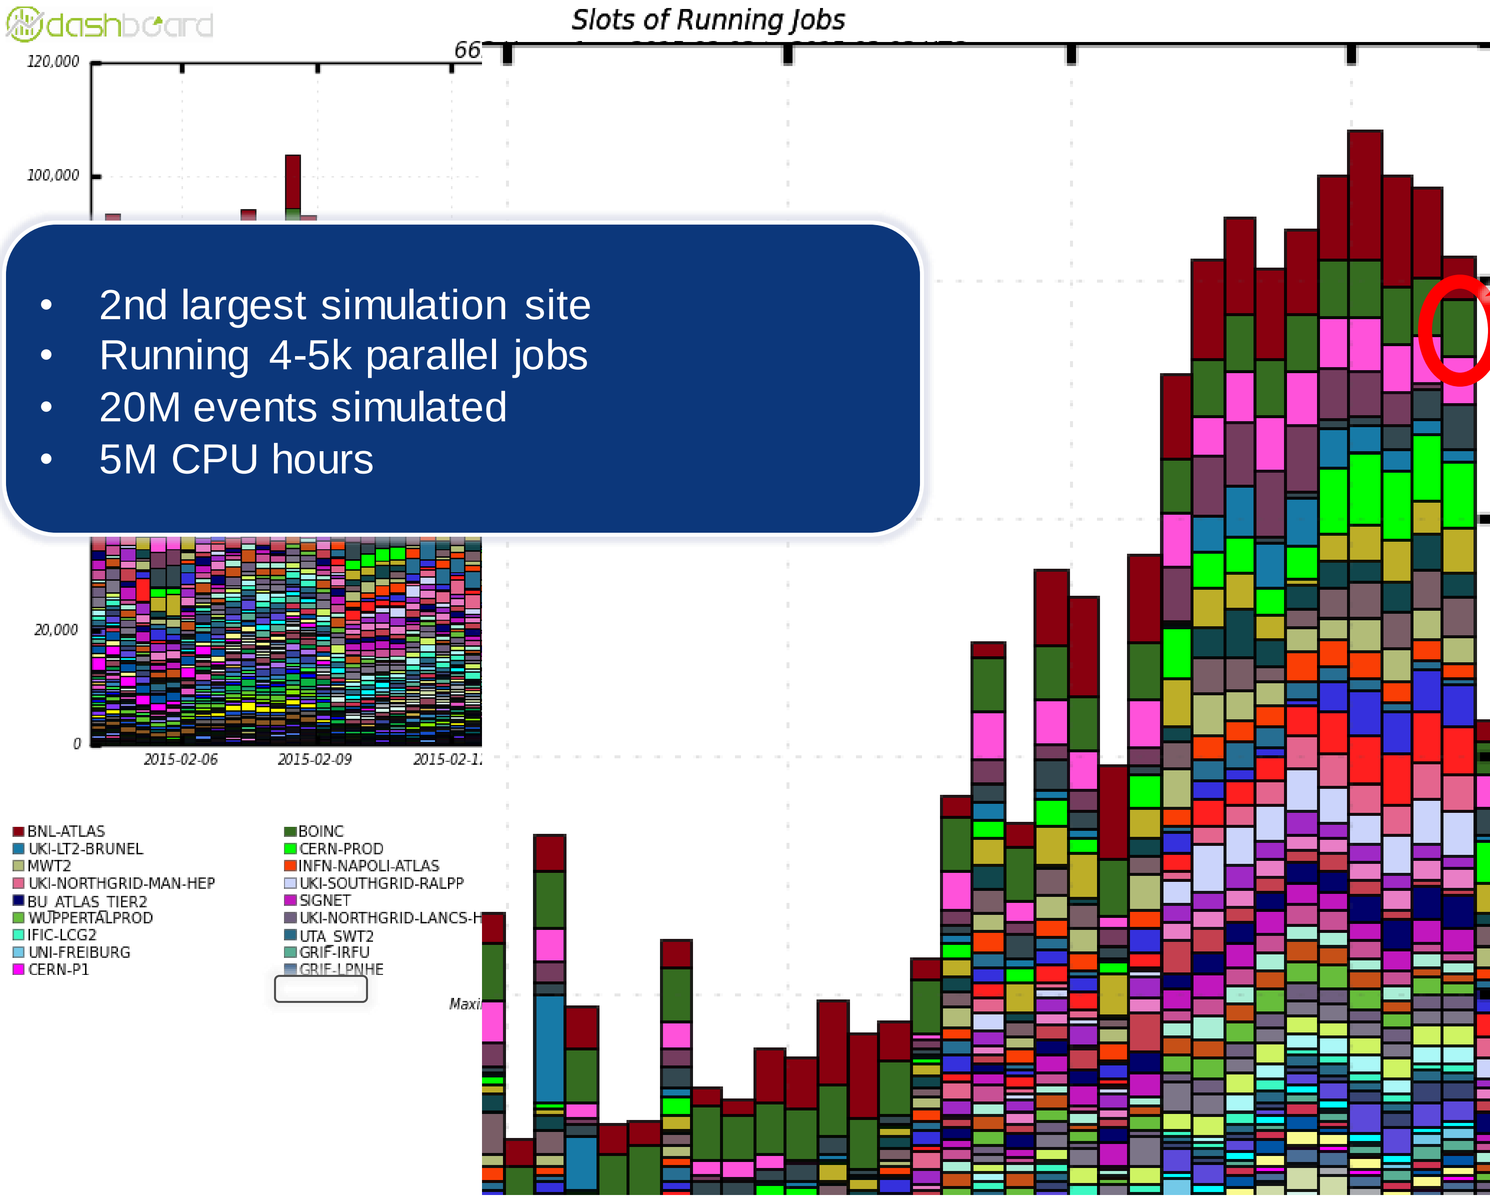

http://atlasathome.cern.ch/atlas_job.php

YEARS / ANS CERN

CMS@home prototype

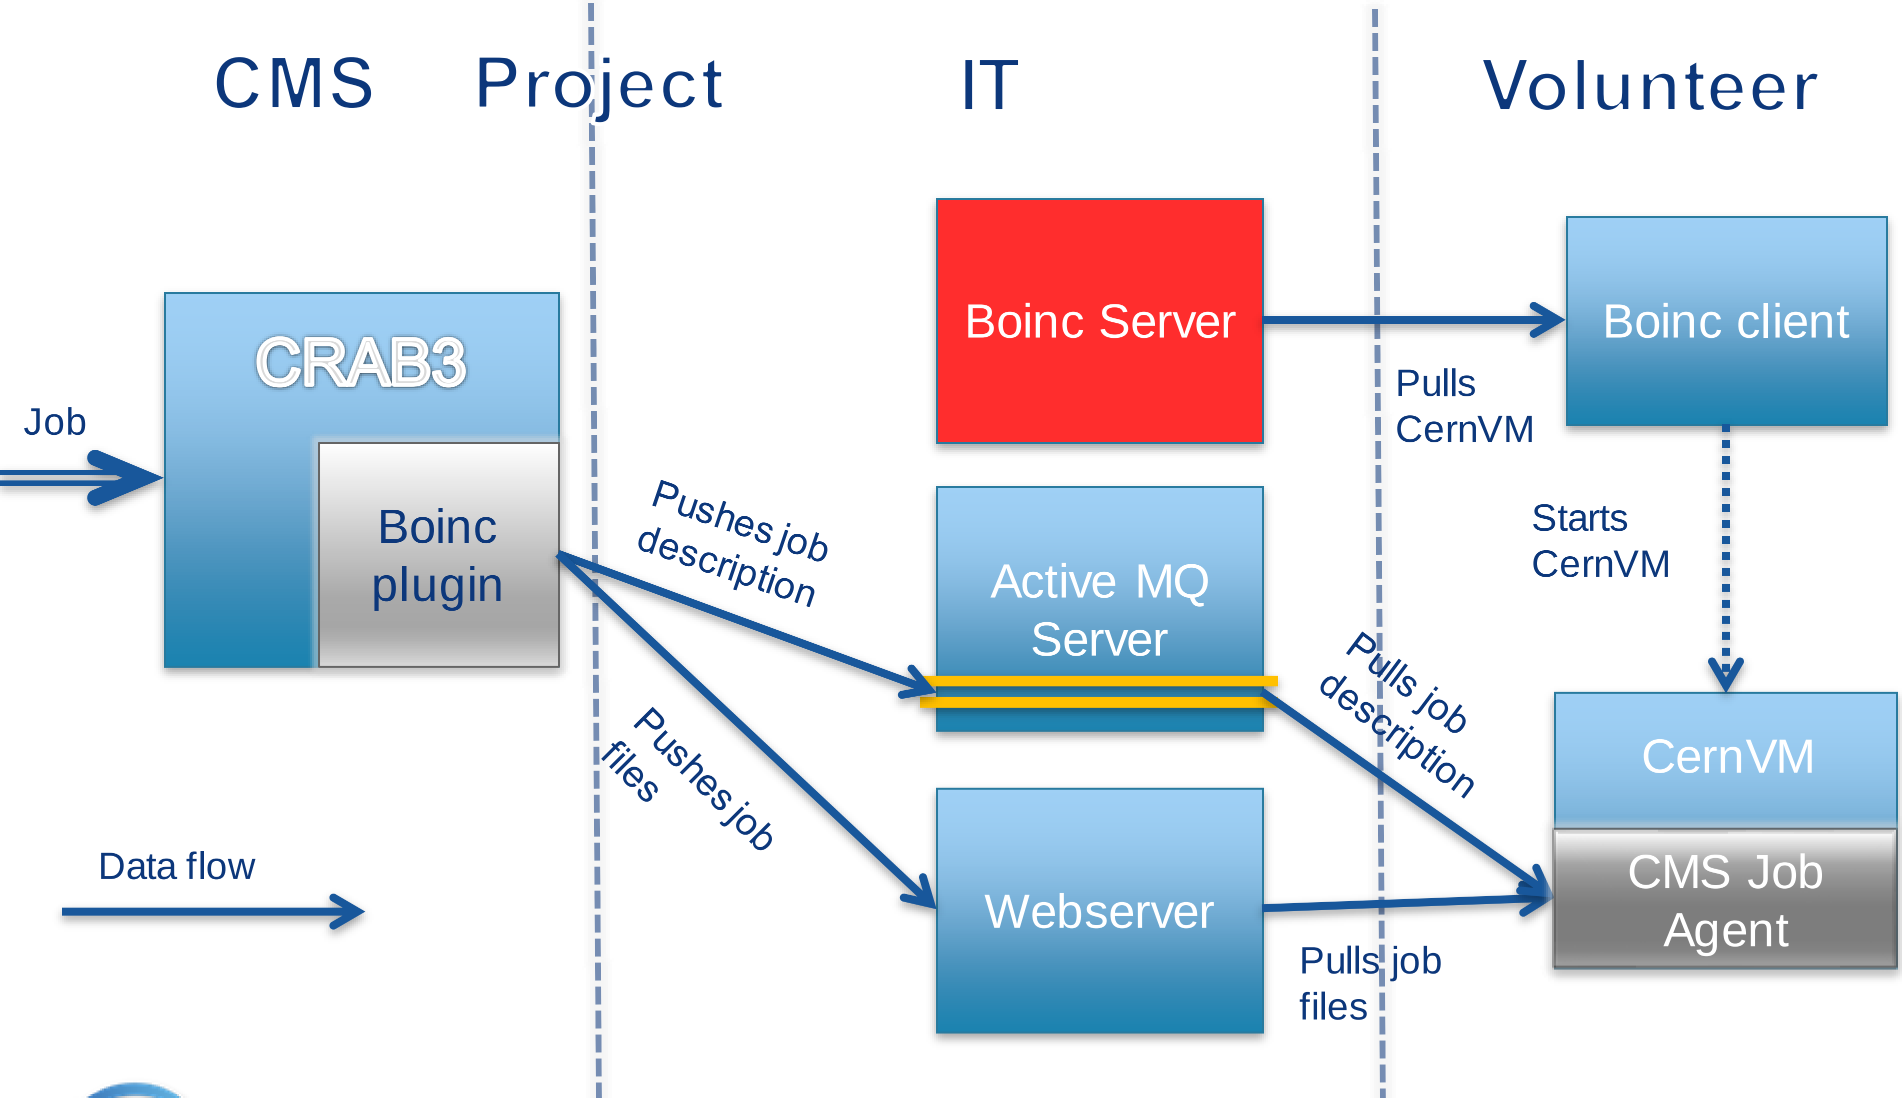

Summary

WLCG

- Grid computing worked for WLCG
- Model changed from Tree to Mesh structure
  - networks improved much faster than CPUs
- Shift from resource provider to user community
  - new tasks, new responsibilities, new tool-chains
- Focus now:
  - Lower operations costs → Clouds
  - Common technologies
  - Private/Commercial Clouds
  - Opportunistic resources
  - Optimization of code and workflows
    - need ~ factor 10-20 improvement!

YEARS / ANS CERN

www.cern.ch