



Status of LHC computing activities at CCIN2P3

Renaud Vernet – Nov. 2013





News from the LHC



The Nobel Prize in Physics 2013

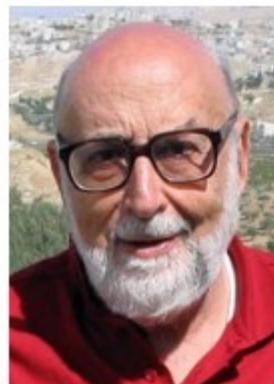


Photo: Pnicolet via
Wikimedia Commons

François Englert

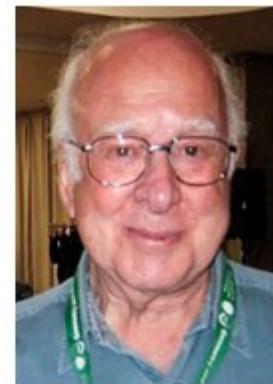


Photo: G-M Greuel via
Wikimedia Commons

Peter W. Higgs

The Nobel Prize in Physics 2013 was awarded jointly to François Englert and Peter W. Higgs "for the theoretical discovery of a mechanism that contributes to our understanding of the origin of mass of subatomic particles, and which recently was confirmed through the discovery of the predicted fundamental particle, by the ATLAS and CMS experiments at CERN's Large Hadron Collider"

TECH
Technology Consortium
of the Ohio Board of Regents

CC-IN2P3, a French Computing Center
For Nuclear Physics and Particle Physics

work for them

Who we are

They send us their data

Our services

The IN2P3 Computing Center (CC-IN2P3) is a service and research unit belonging to CNRS. A major French research infrastructure, it provides computing and data storage resources for researchers involved in corpuscular physics experiments.

The main services offered by CC-IN2P3 are the **storage and processing** of large volumes of data and their transfer over very high-speed international networks. Experiments in high-energy and astroparticle physics require the analysis of billions of interactions between particles. More than 50 international collaborations in physics use the resources of CC-IN2P3.

Higgs Boson Discovery

The Nobel Prize in Physics 2013 was awarded to F. Englert and P.W. Higgs for the discovery of an essential part of the Standard Model of particle physics: the **Higgs Boson**.

The search for the Higgs at the Large Hadron Collider was an **international effort** involving thousands of people. The Worldwide LHC Computing Grid gave a community of over 8000 researchers near **real-time access to LHC data**.

France played an active role in LHC data processing with the CC-IN2P3, one of the biggest computing centers on the WLCG.

cnrs

OUR SECOND DATACENTER
For Nuclear Physics and Particle Physics

Modularity

- From rack to building
- Capacity planning up to 2019 3 phases 9 MW on site 3.6 MW in room
- Multi-tier architecture
- Movable wall between computing and spare room

Ease of deployment

- All distribution from ceiling Confined hot aisle 15 kW mean per 11 rack
- 6 aisles of 2x20 racks Inflammable parts away
- Delivery and test room
- Preparation decreases deployment time

Secured power lines

- Dedicated and primary 9 MW
- 2 MW backup line no power generator
- First UPS chain capacity 2 MW
- Transparent switch to main power line
- Small battery capacity 20 minutes full phase I

Cooling system

- Resilient
- Backup for 20 minutes
- Capacity to provide 70 m³/hour at 50°C (122°F)
- PUE: 1.47

First phase completed in April 2013

230

CCIN2P3 @ SC'13

OUR SECOND DATACENTER

For Nuclear Physics and Particle Physics

First phase completed in April 2013

Modularity

- From rack to building
- Capacity planning up to 2019
- 3 phases 9 MW on site
- 3.6 MW in room
- Multi-tier architecture
- Movable wall between computing and spare room

Ease of deployment

- All distribution from ceiling
- 6 aisles of 2*20 racks
- Delivery and test room
- Preparation decreases deployment time

Secured power lines

- Dedicated and primary 9 MW
- 2 MW backup line
- First UPS chain capacity 2 MW
- Transparent switch to main power line
- Small battery capacity 20 minutes full phase 1

Cooling system

- Resilient
- Backup for 20 minutes
- Capacity to provide 70 m³/hour at 50°C (122°F)
- PUE: 1.47

LHC

The world's largest particle accelerator

LHC main goals

Beyond revealing a new world of unknown particles, the LHC experiments could explain why those particles exist and behave as they do. With the discovery of the Higgs boson, one of the main goals has been achieved. Until last year, the Higgs boson mechanism was the last remaining piece of the physics standard model to be experimentally verified. But the LHC experiments could also shed light on dark matter, uncover hidden symmetries of the universe, and possibly find extra dimensions of space.

LHC in numbers

- 27-kilometer ring of superconducting magnets
- Circular tunnel 100 meters beneath the Swiss/French border at Geneva
- Produces 15 Petabytes of data annually
- A community of more than 8,000 physicists around the world

Worldwide LHC Computing Grid

The WLCG project is a worldwide grid infrastructure that provides mutualized computing resources to store, distribute, and analyze LHC data. CC-IN2P3 is one of the biggest computing centers of the WLCG.

CC-IN2P3 Contribution

- 10,000 logical CPUs for LHC
- 2 PB storage capacity for LHC
- Fast access DAS for High-Throughput Computing (HTC)
- Heterogeneous solutions for heterogeneous needs
- Robust and performant networking solutions
- Hierarchical vs. data custody level vs. data access speed

Strategic Storage Organization

LHC DATA
COMPUTING (Compute power)
STORAGE (Data)
SLG LinexA
Heterogeneous Solutions for high-throughput needs

Each region of the sky will correspond to an individual database in the Osserv database system. The regions are slightly overlapping to optimize the queries on objects located close to a boundary.

<http://wlcg.web.cern.ch> <http://cc.in2p3.fr>

LSST

Large Synoptic Survey Telescope

LSST main goals

The LSST project aims at mapping the mysterious dark matter and characterizing the properties of the even more mysterious dark energy thanks to its large aperture and giant camera. With a light-gathering power among the largest in the world, it can detect faint objects with short exposures.

LSST in numbers

- 3.2 Gigapixels high sensitivity CCD camera
- 15 TB of data per night
- Entire sky covered twice a week

Total dataset over the ten years of operation:

- ~60 PB for the raw data
- ~34 PB for the catalog database

Total data volume after processing:

- several hundreds PB
- requires 500 Tflops for catalog generation

LSST Data Management

A set of production pipelines will process the images to produce a number of catalogs containing the detected astronomical sources and resolved astrophysical objects.

These catalogs will be stored in a large scale database system distributed on thousands of nodes using a "shared nothing" technology. This distributed system, named Gserv, is specifically developed within the LSST collaboration. The whole dataset will be jointly reprocessed yearly at CC-IN2P3 in France and at the National Center for Supercomputing Application (NCSA) in Illinois.

CC-IN2P3 Contribution

By 2022, the CC-IN2P3 will have deployed 15 PB of disk space, 8 PB of mass storage and 20,000 cores ready to process its share of the first data coming out from the telescopes. This capacity will be pushed to more than 100,000 cores, 37 PB of disk and 82 PB of mass storage by the end of the project. CC-IN2P3 is also contributing to the LSST database development effort by providing a 900-node test platform.

Applications

Science Software has been built in the framework of the EGI-CHARGE and CHARMED projects to demonstrate how the Science Gateway paradigm can be used to realize the core scientific infrastructure of LSST, based on distributed management and architecture, using EGI Cloud or using a local cluster, transparently using each other, at user request level.

<http://www.lsst.org>

JSAGA

Simple API for Grid Applications

JSAGA main goals

JSAGA is developed at the CC-IN2P3 since 2008. It was initially funded by the GTFD grant from the French National Research Agency.

JSAGA in numbers

- JSAGA
- Grid
- Java
- API

SAGA: Simple API for Grid Applications

Specification defined by the Open Grid Forum. SAGA is a specification for generic access to different grid infrastructures. This specification defines a middleware-independent and programming-language-independent interface to submitting and monitoring jobs, and managing logical and physical resources.

SAGA is independent of existing middleware packages and lets it communicate.

Others major implementations worldwide:

- SAGA++ (C++)
- SMURF (Python)

JSAGA

- is independent of operating system (Linux/Mac/Windows)
- can execute via plugins
- is designed for efficient reuse of resources (one workflow per job)
- offers seamless usage of middleware native features
- is insulating unpatched sources

Workflow of plugins chain

A LSST pipeline is described as a workflow starting from raw data, processed by several filters, then by a reduction step, followed by a final analysis step.

Streaming mode

The stream of LSST events is processed through three stages to generate 100% low-latency through high data rates.

<http://grid-in2p3.fr/jsaga>

LAVOISIER

Data aggregation framework

Lavoisier main goals

Lavoisier is a framework to retrieve, transform, merge and query heterogeneous data sources. Lavoisier can really be extended through new plug-ins.

Plugins for input data

Input generic technologies

Workflow of plugins chain

A LSST pipeline is described as a workflow starting from raw data, processed by several filters, then by a reduction step, followed by a final analysis step.

Streaming mode

The stream of LSST events is processed through three stages to generate 100% low-latency through high data rates.

<http://grid-in2p3.fr/lavoisier>

CC-IN2P3, a French Computing Center

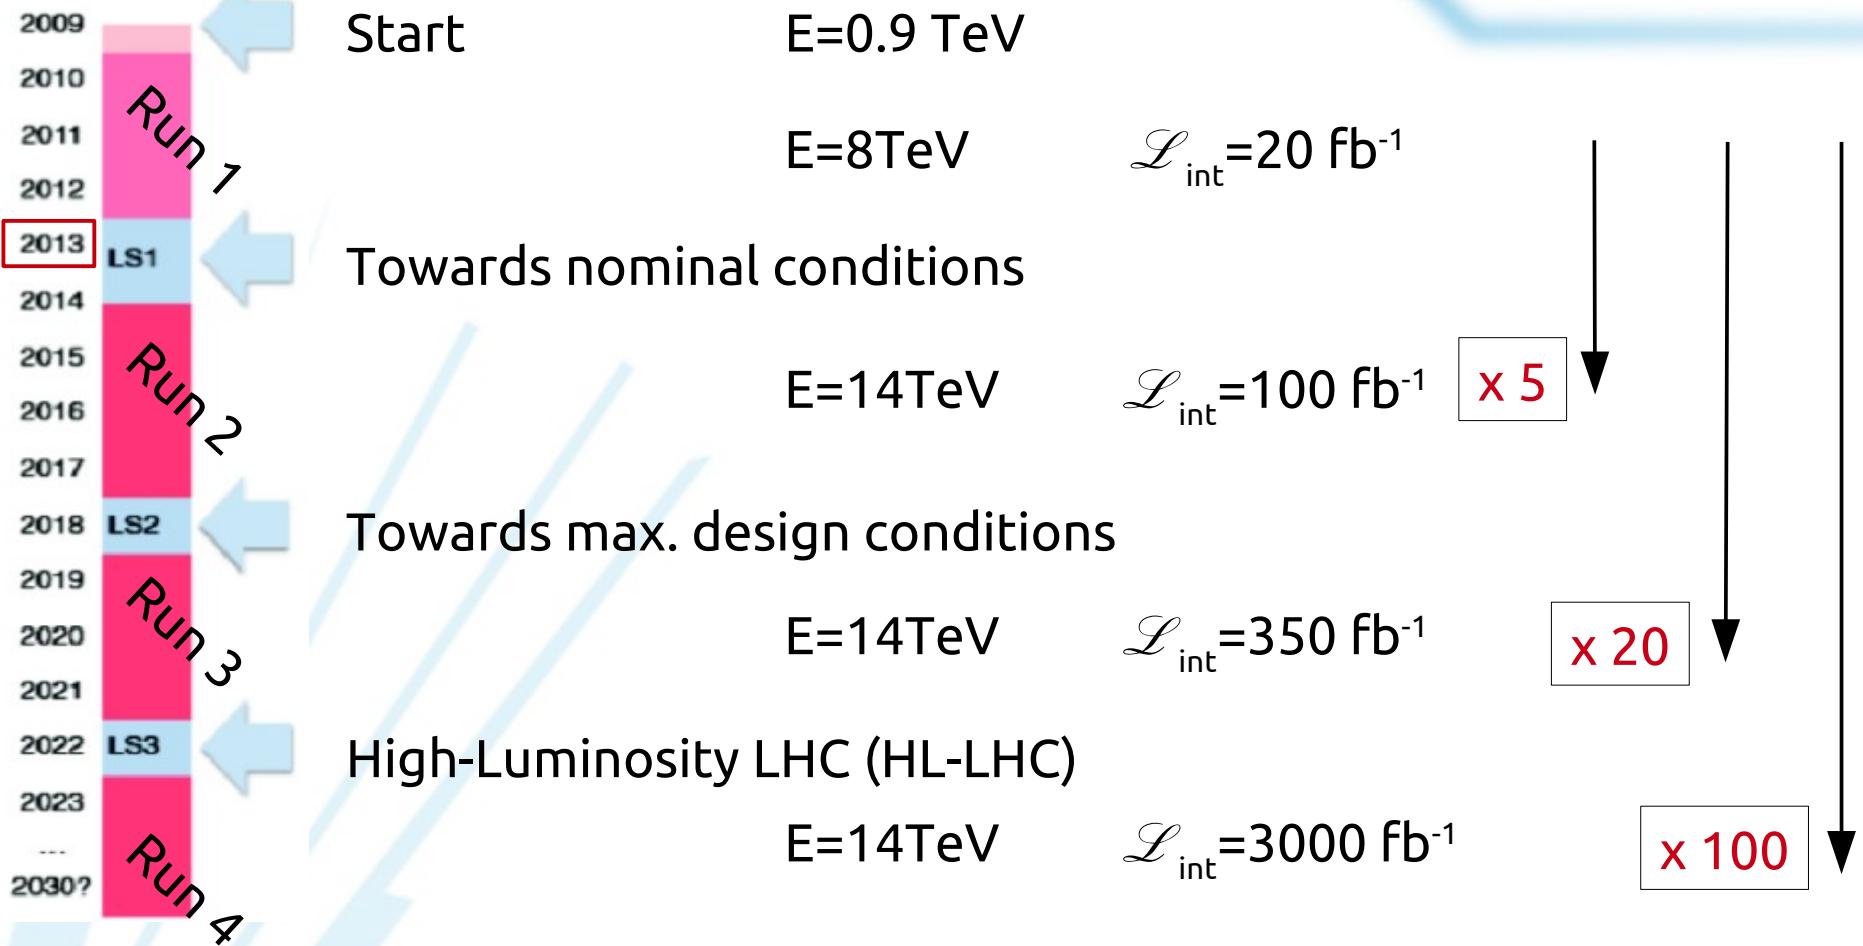
For Nuclear Physics and Particle Physics

We work for them

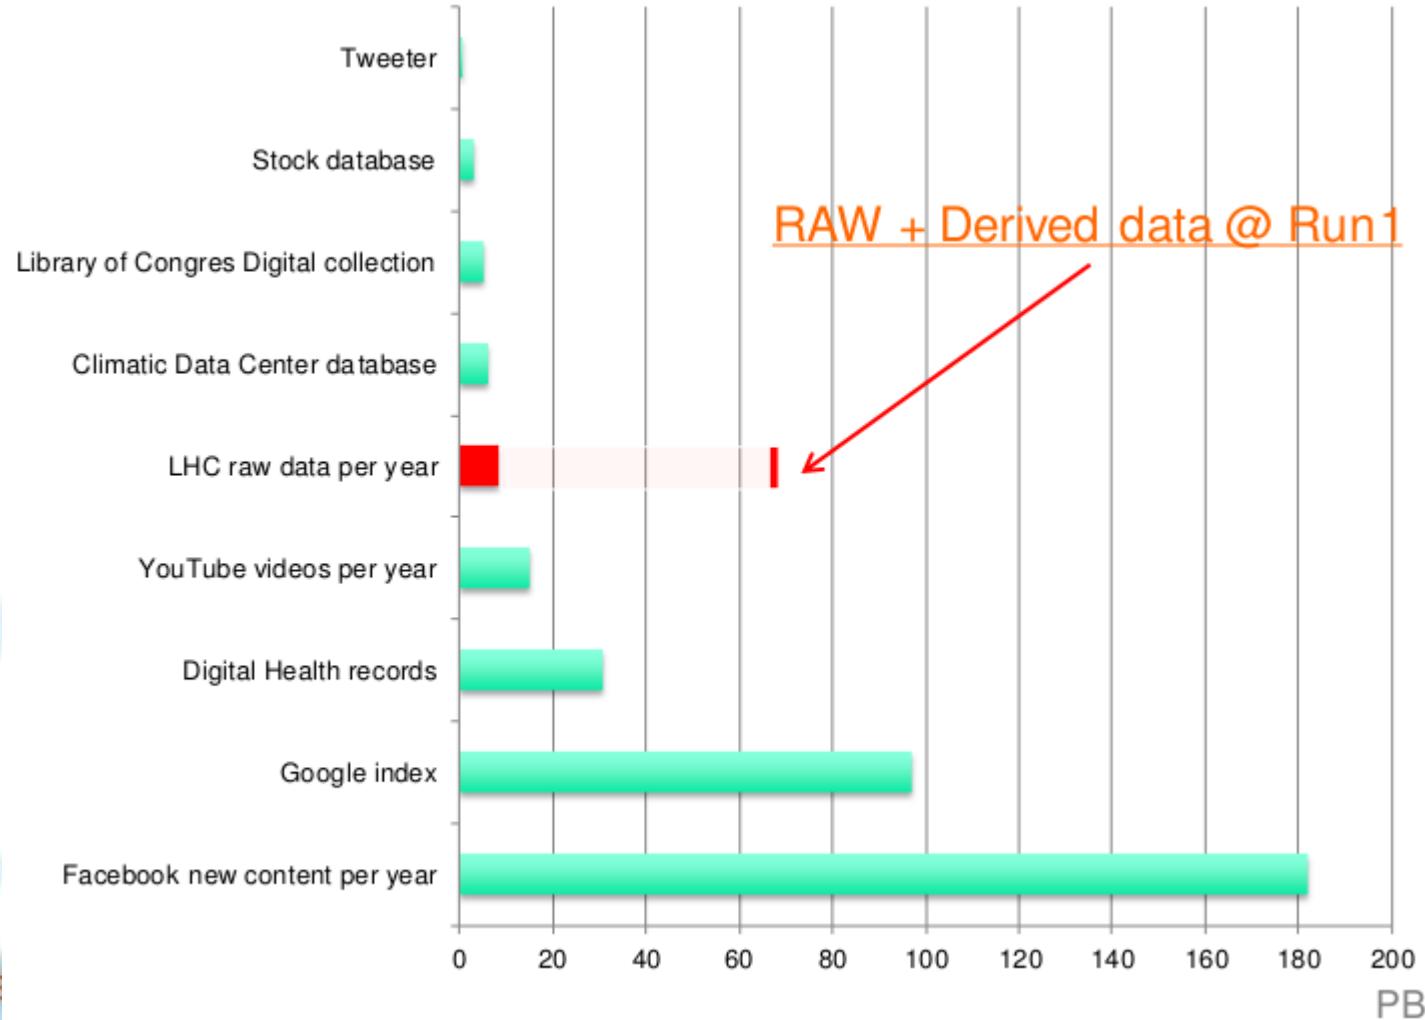
Who we are



LHC planned schedule



Data volume : where we stand today

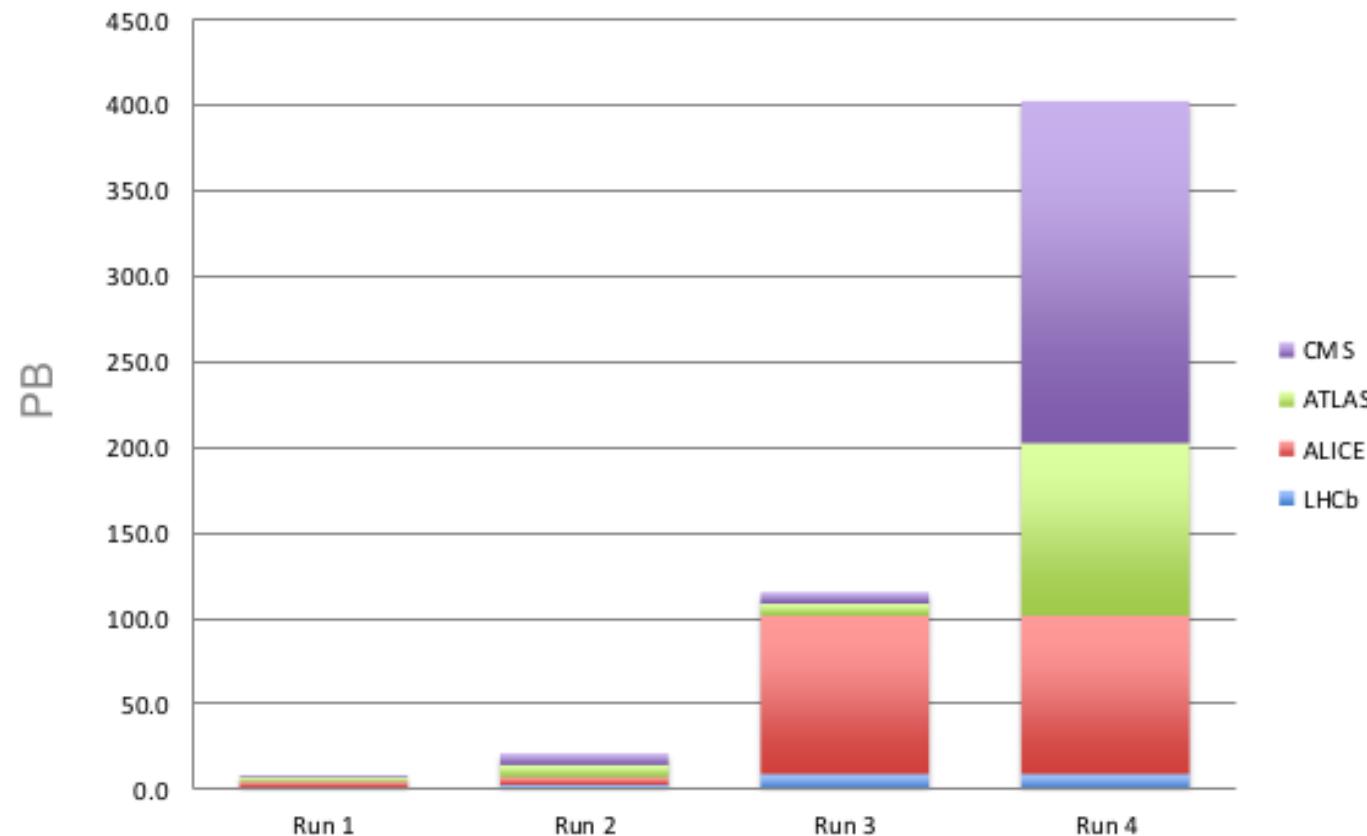


P. Buncic @ EFCA High Lum. LHC 2013

Data volume : where we (plan to) go



■ Raw data volume per year



P. Buncic @ EFCA High Lum. LHC 2013

Outline

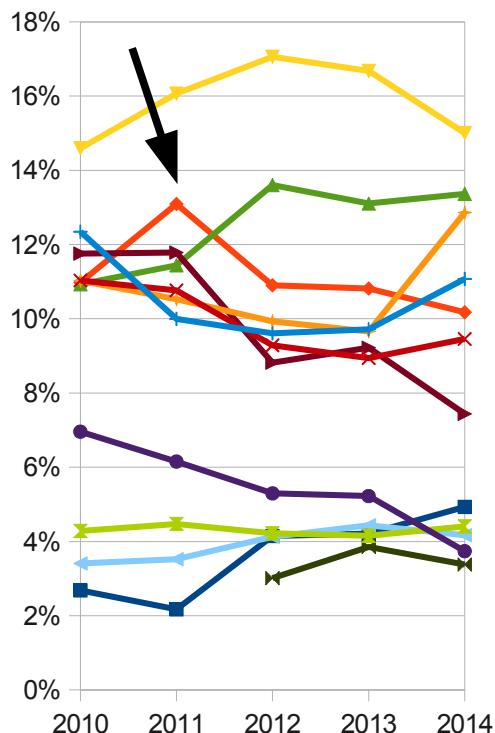


- Resource status and usage
- Recent news (upgrades)
- Site performance
- Benchmarks
- Conclusions

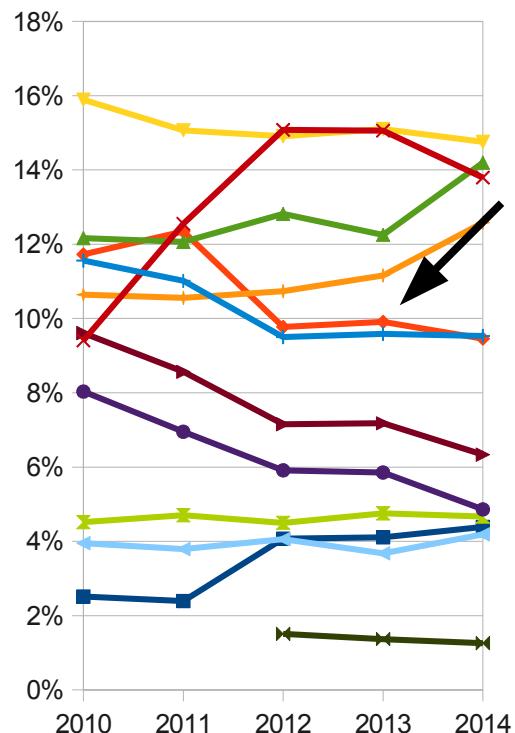
Our weight among all T1's



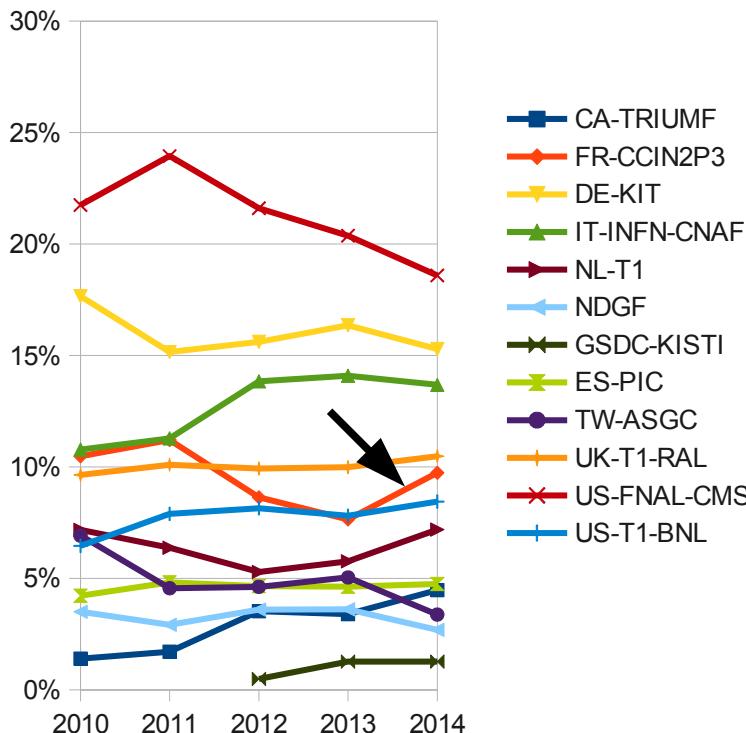
T1 shares (CPU)



T1 shares (Disk)



T1 shares (Tape)

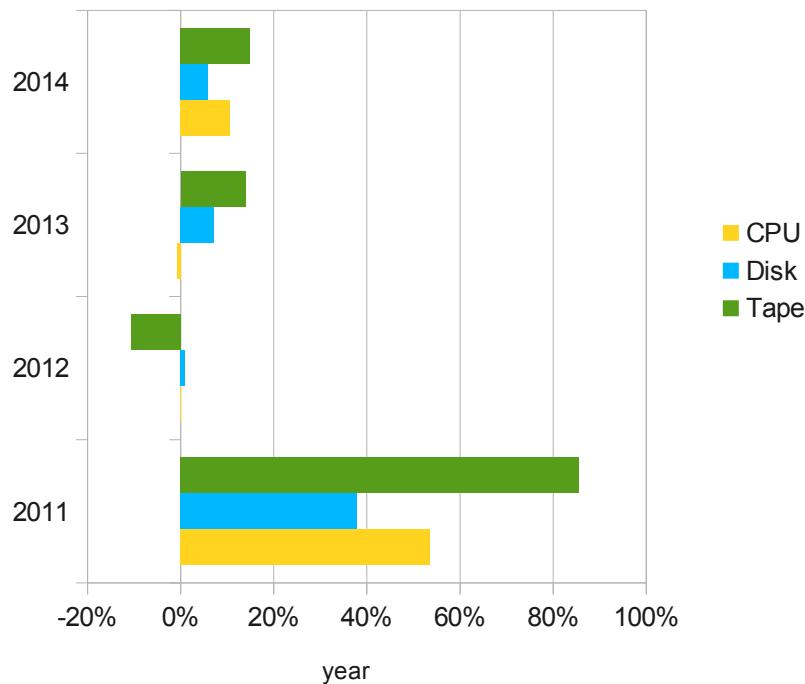


LHC resources growth at CCIN2P3



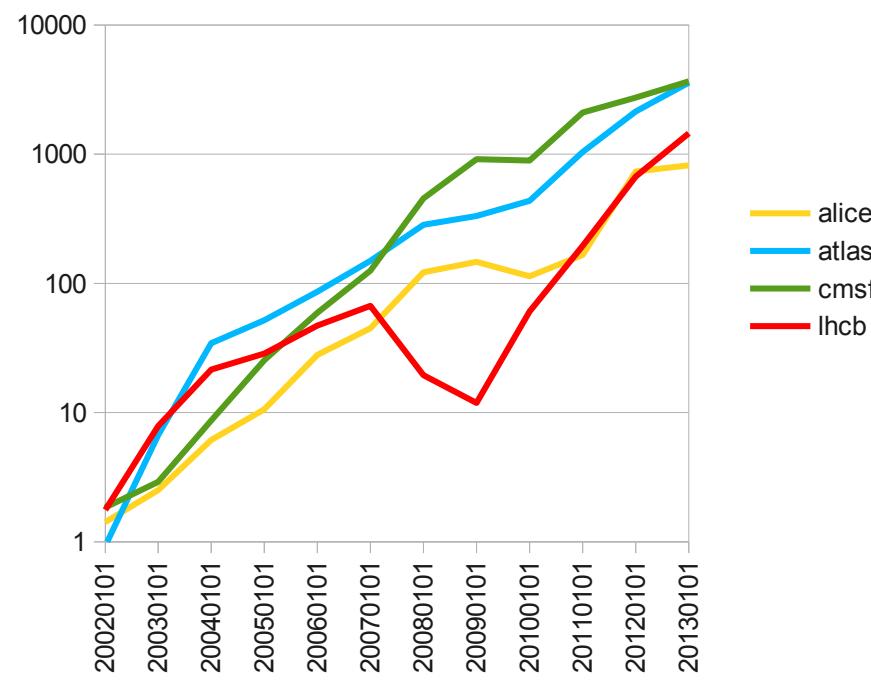
Resource Growth

T1+T2



Space used by VOs on HPSS

(in TB)

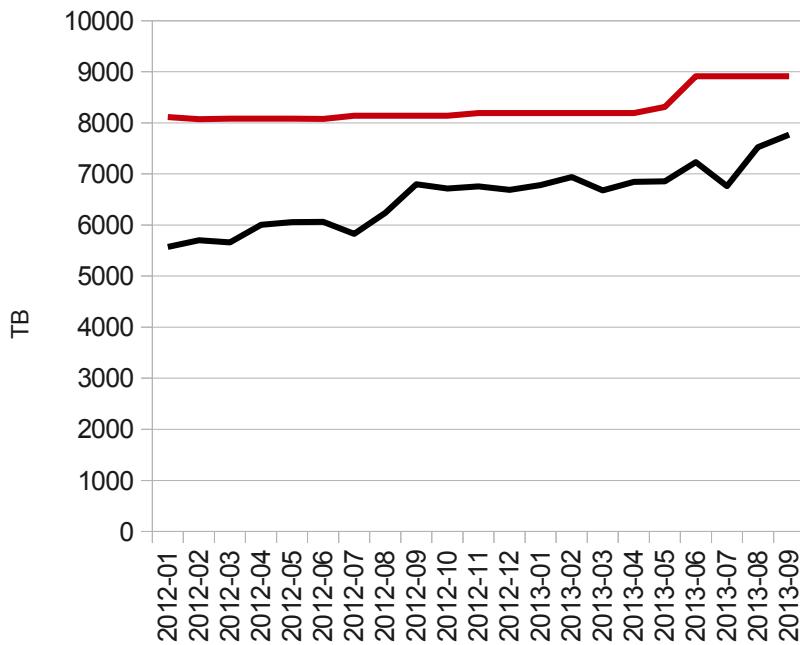


Storage usage



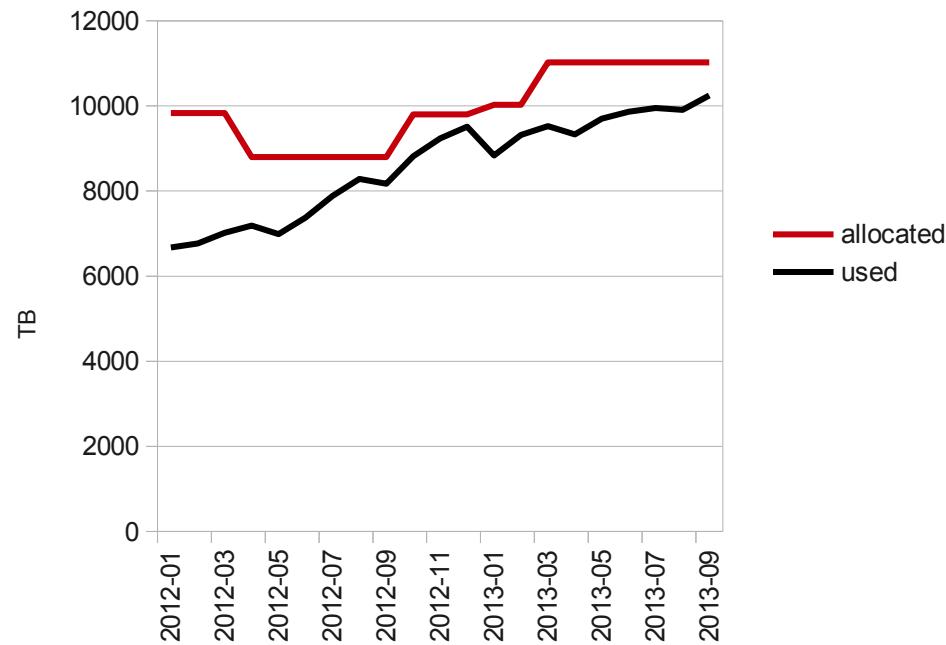
Disk usage - LHC

CCIN2P3



Tapes used vs allocated (TB) : LHC exp.

CCIN2P3

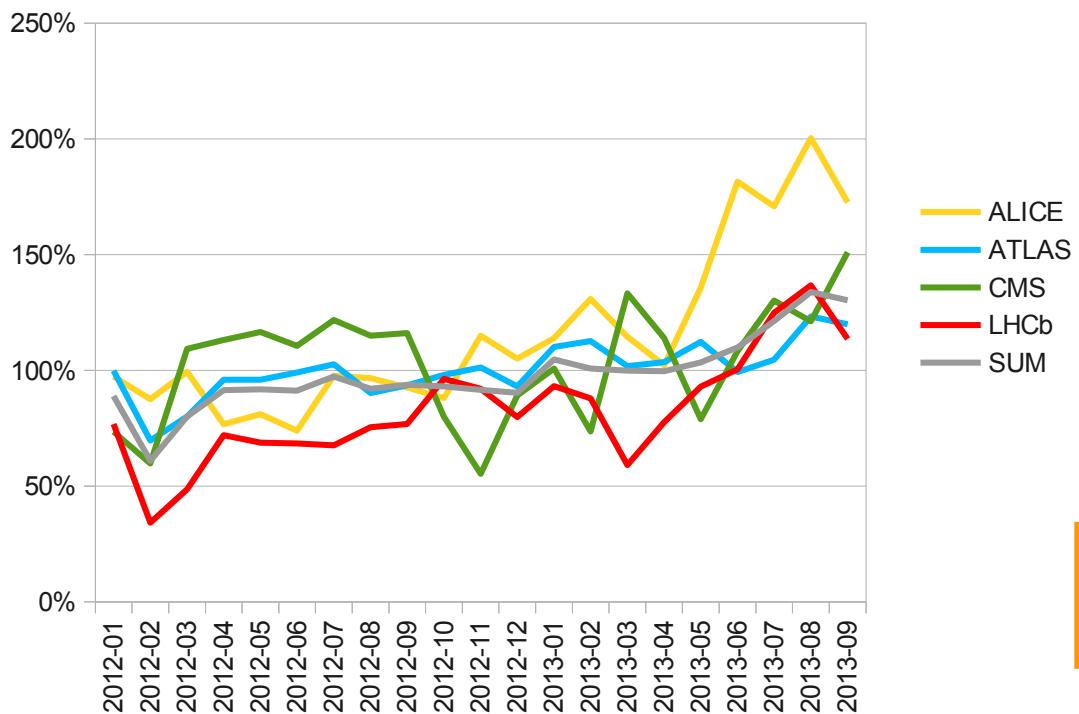


CPU usage



Wall / pledged

all jobs



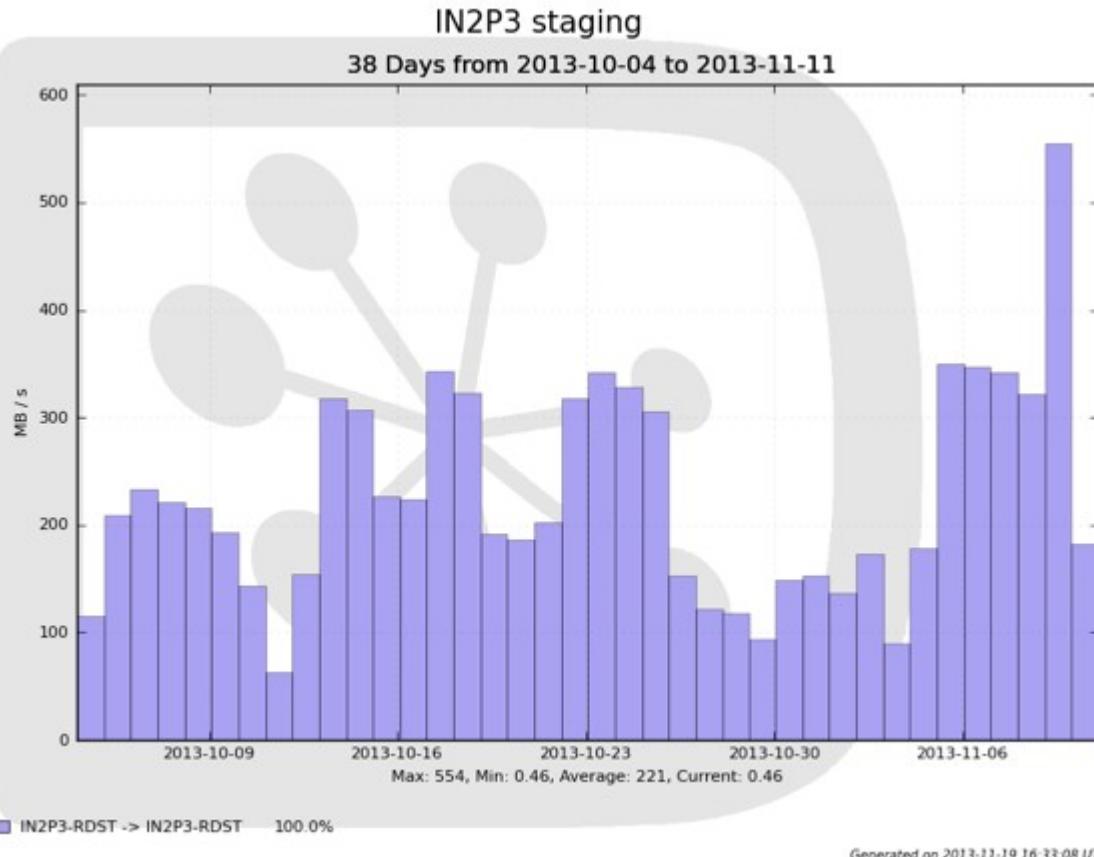
Very good usage of
CPU ressources this year

Noticeable changes since 6 months



- Migration CVMFS (ALICE)
- Into CMS xrootd federation
- Migration SL6 ~achieved
- Migration from Oracle to Univa Grid Engine
- Migration SHA-2
- Installation FTS3 instance

Performance : staging



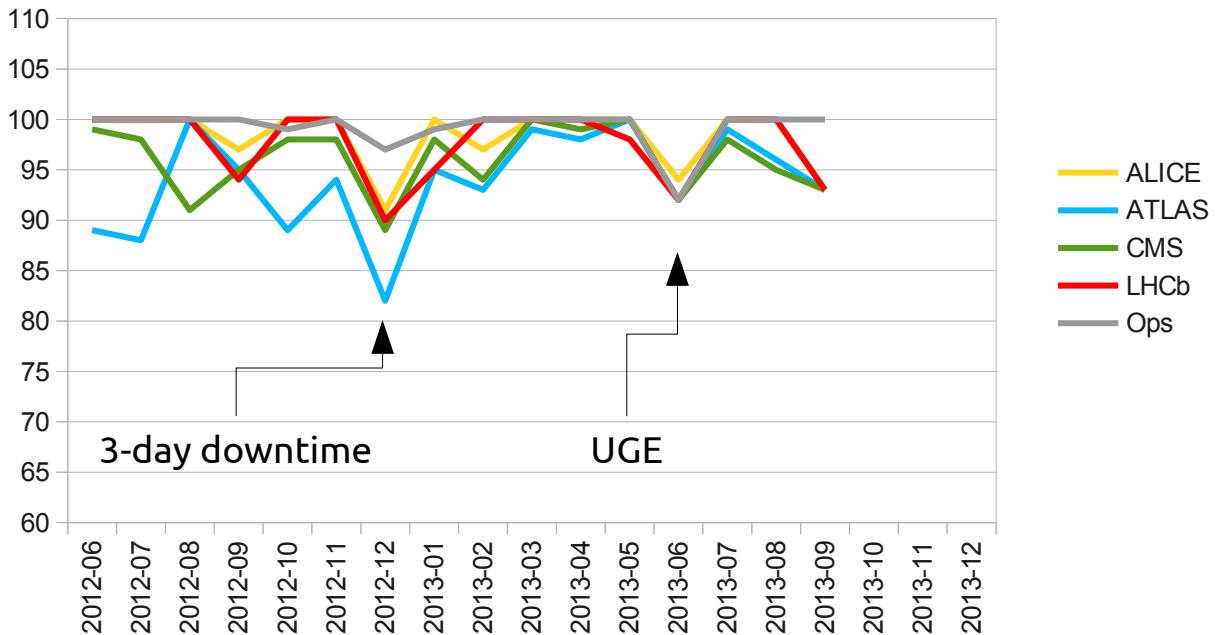
LHCb tape staging :
220 MB/s observed
138 MB/s required

Tier1-site availability



T1 availability

CC-IN2P3



Average on the last 12 months :

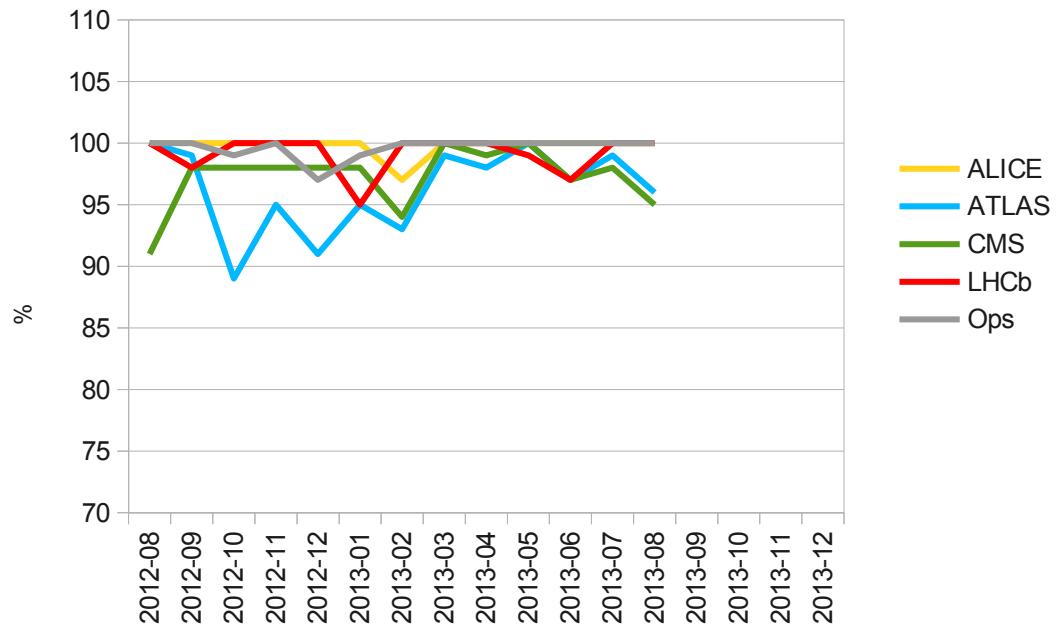
	ALICE	ATLAS	CMS	LHCb	Ops
	97.7	94.6	96.0	97.1	98.9

Tier1-site reliability



T1 reliability

CCIN2P3



We must try to be better !

Average on the last 12 months :

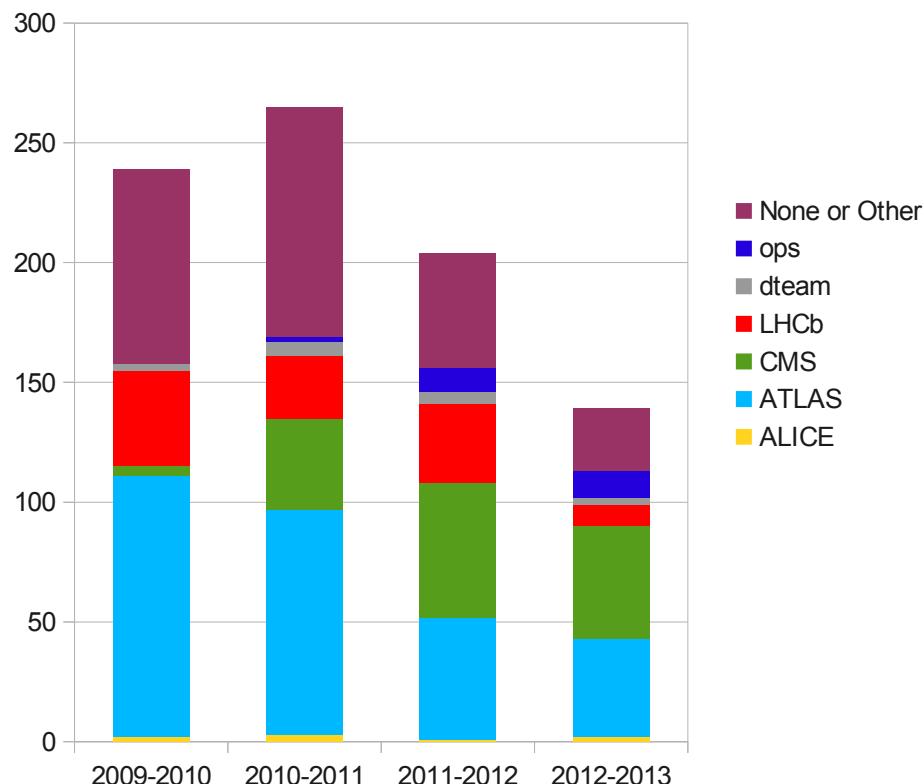
	ALICE	ATLAS	CMS	LHCb	Ops
	99.7	96.3	97.7	99.1	99.6

GGUS tickets



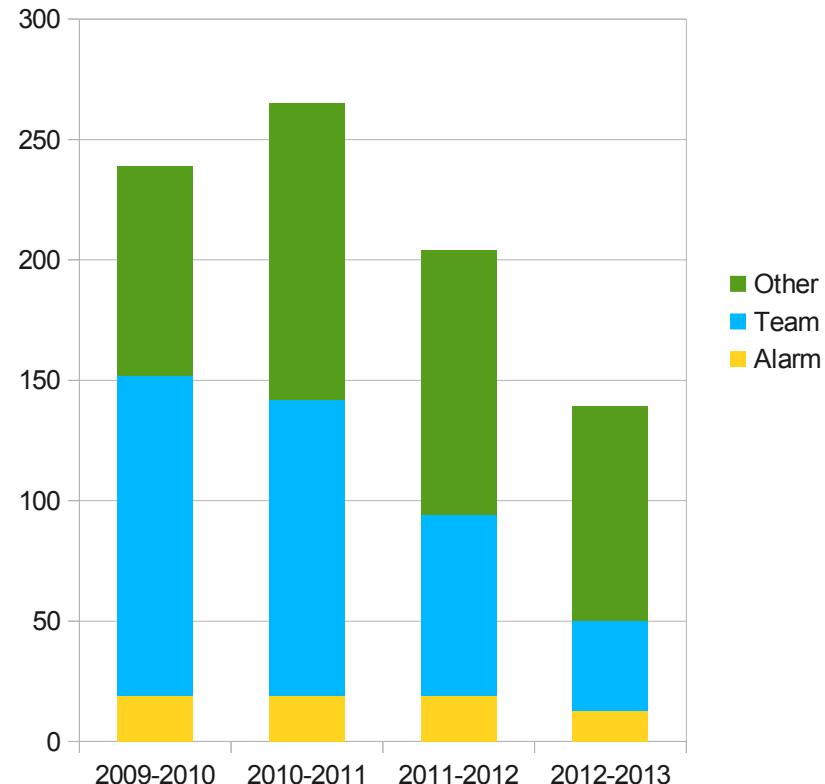
Number of GGUS tickets opened per VO

nov. 1st - oct. 31st



GGUS Ticket Type

nov. 1st - oct. 31 st



Benchmarks



■ Discussion 😊

Conclusions



- Site performance improves with time
- Continuous decrease of incidents (GGUS tickets)
- Availability and reliability to improve !
 - Has to be investigated deeper
 - VOs manage to use walltime > pledges , however
- Up-to-date
 - Middleware, CVMFS, SHA-2, SL6