

DIRAC Distributed Computing Services

*A. Tsaregorodtsev,
CPPM-IN2P3-CNRS*

CHEP 2013, Amsterdam, 17 October 2013

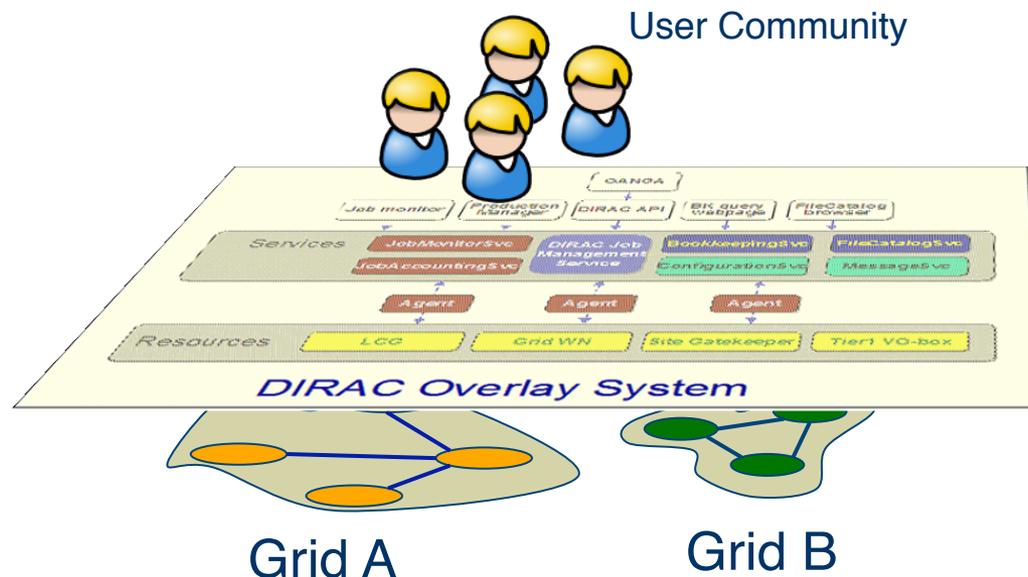


- ▶ Motivation and brief history
 - ▶ From LHCb to a general purpose project
 - ▶ From community to a general purpose service
- ▶ DIRAC as a Service
 - ▶ France-Grilles and other services
 - ▶ Resources available
 - ▶ Services available
- ▶ Conclusions

- ▶ LHC experiments pioneered the massive use of computational grids
 - ▶ 10s of PBytes of data per year
 - ▶ 100s of thousands CPUs in 100s of centers
 - ▶ 10s GB/sec network transfers
 - ▶ 100s of users from 100s of institutions
- ▶ Other domains are catching up quickly with the HEP experiments
 - ▶ Life sciences, earth sciences, astrophysics, social sciences, etc

- ▶ The computing expertise level in non-HEP scientific domains is relatively lower
 - ▶ Grouped around well known applications and scientific portals
 - ▶ Moving existing applications to run in distributed environments is still difficult
- ▶ Convenient tools for small research groups with no local gurus are clearly needed
- ▶ All LHC experiments developed their own middleware
 - ▶ PanDA, AliEn, glideIn WMS, PhEDEx, DIRAC, ...
 - ▶ WMS with pilot jobs, intelligent data management, software distribution, ...
- ▶ Experience of the LHC experiments in using distributed computing infrastructures should now be made available for non-LHC user communities

- ▶ DIRAC provides all the necessary components to build ad-hoc distributed computing infrastructures interconnecting resources of different types, allowing interoperability and simplifying interfaces. This allows to speak about the DIRAC *interware*.



- ▶ Several new experiments expressed interest in using this software relying on its proven functionality
- ▶ In 2009 the core DIRAC development team decided to generalize the software to make it suitable for any user community.
 - ▶ Separate LHCb specific functionality into a set of extensions
 - ▶ Introduce new services to make it a complete solution
 - ▶ Support for multiple small groups by a single DIRAC installation
 - ▶ General refurbishing of the code, code management, deployment, documentation, etc
- ▶ This work made it possible to offer general-purpose DIRAC services to any scientific community

DIRAC Community Installations

- ▶ LHCb stays the most important client of the DIRAC system

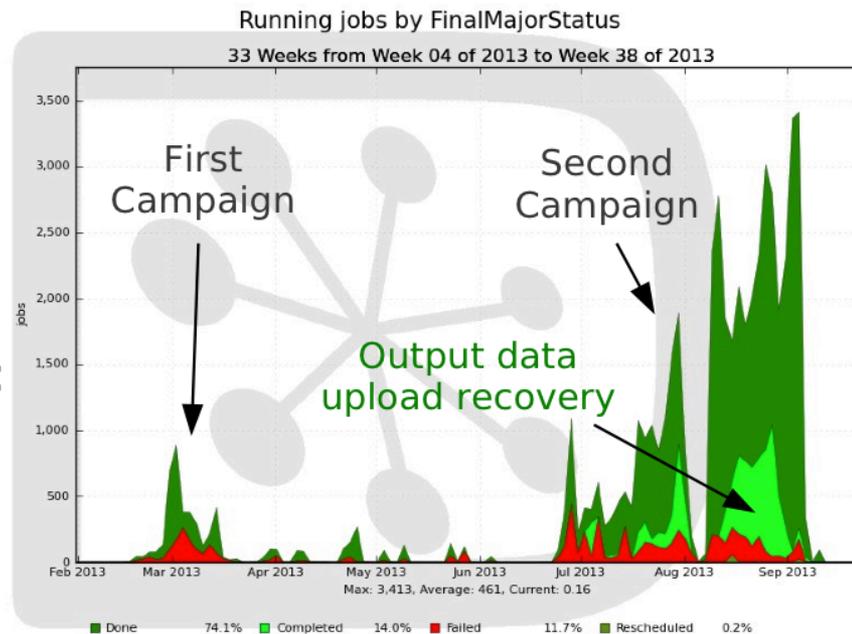


- ▶ Using DIRAC for all the computing tasks:
 - ▶ WMS, DMS, Data Production Management, Accounting, etc
- ▶ See presentations [here](#), [here](#), and others

- ▶ Belle II



- ▶ Combination of the non-grid, grid sites and (commercial) clouds is a requirement
- ▶ 2 GB/s, 40 PB of data in 2019
- ▶ Belle II grid resources
 - ▶ WLCG, OSG grids
 - ▶ KEK Computing Center
 - ▶ Amazon EC2 cloud
- ▶ First production run is done
 - ▶ See *T.Kurh's* [presentation](#)





- ▶ ILC/CLIC detector Collaboration
 - ▶ MC simulations
 - ▶ DIRAC File Catalog was developed to meet the ILC/CLIC requirements
 - ▶ See [poster](#)



BESIII Experiment

- ▶ BES III, IHEP, China
 - ▶ Using DIRAC DMS: File Replica and Metadata Catalog, Transfer services
 - ▶ Dataset management developed for the needs of BES III



- ▶ CTA
 - ▶ CTA started as France-Grilles DIRAC service customer
 - ▶ Now is using a dedicated installation at PIC, Barcelona
 - ▶ Using complex workflows
 - ▶ See [poster](#)
- ▶ DIRAC evaluations by other experiments
 - ▶ Fermi-LAT, LSST, Auger, TREND, Daya Bay, Geant4, ...
 - ▶ Evaluations can be done with general purpose DIRAC services

DIRAC as a Service

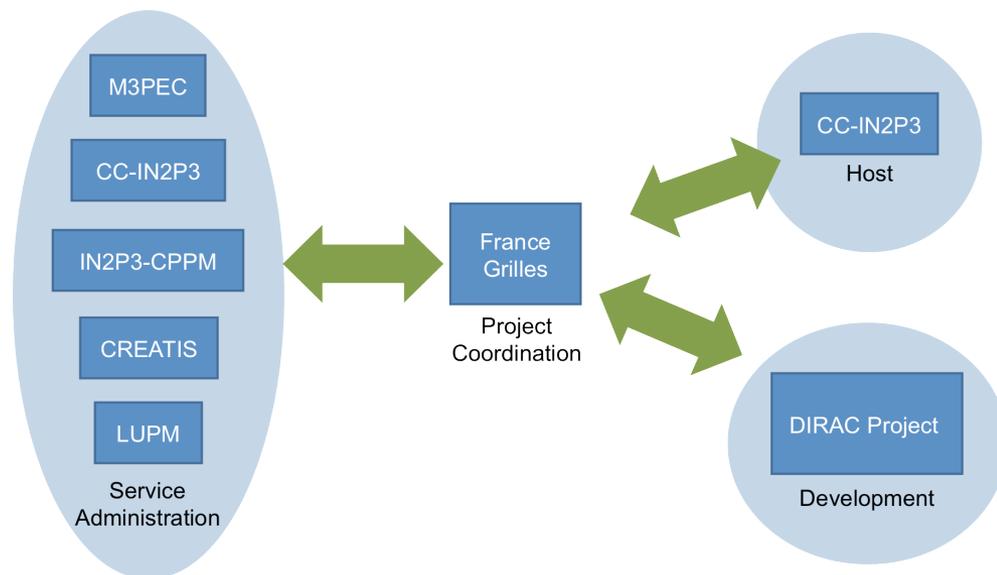
- ▶ **DIRAC client is easy to install**
 - ▶ Part of a usual tutorial

- ▶ **DIRAC services are easy to install but**
 - ▶ Needs dedicated hardware for hosting
 - ▶ Configuration, maintenance needs expert manpower
 - ▶ Monitoring computing resources is a tedious every-day task

- ▶ **Small user communities can not afford maintaining dedicated DIRAC services**
 - ▶ Still need easy access to computing resources

- ▶ **Large grid infrastructures can provide DIRAC services for their users.**

- ▶ Several regional and university campus installations in France
 - ▶ Complex maintenance
- ▶ Joint effort to provide France-Grid DIRAC service
 - ▶ Hosted by the CC/IN2P3, Lyon, T1 center
 - ▶ 6 virtual servers, MySQL server
 - ▶ Distributed team of service administrators
 - ▶ 5 participating universities



<http://dirac.france-grilles.fr>



▶ France-Grilles users

▶ 15 VO's, ~100 registered users

▶ astro, auger, biomed, esr, euasia, gilda, glast.org, prod.vo.eu-eela.eu, vo.cta.in2p3.fr, vo.formation.idgrilles.fr, vo.france-asia.org, vo.france-grilles.fr, vo.msfg.fr, vo.mcia.fr

▶ robot users

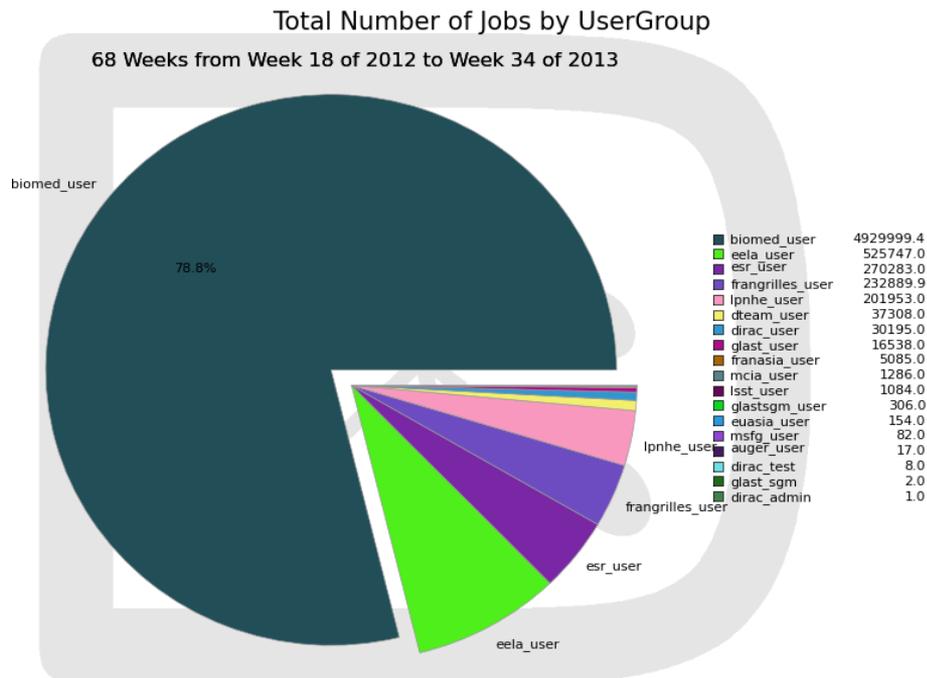
- VIP/GateLab Biomed
- Science Gateway

▶ More VO's and users can be added as necessary

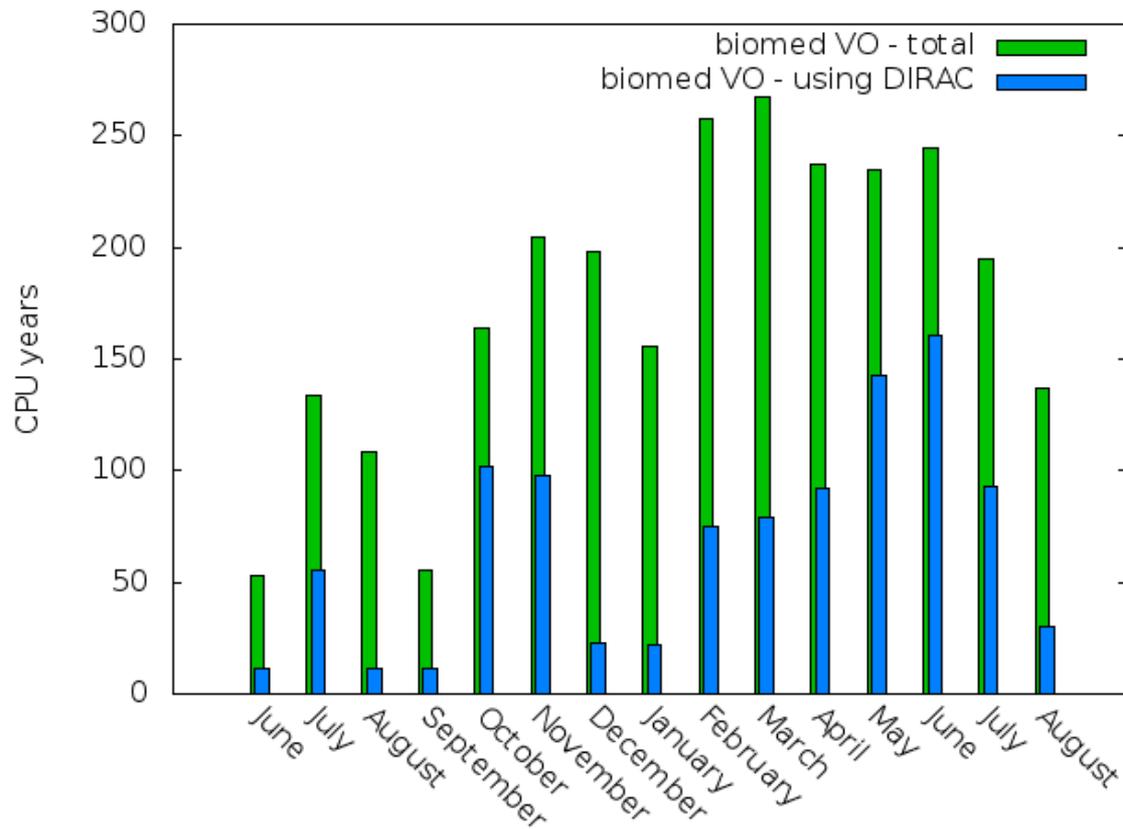
▶ In production since May 2012

▶ ~7 millions jobs went through the system

- Mostly biomed applications



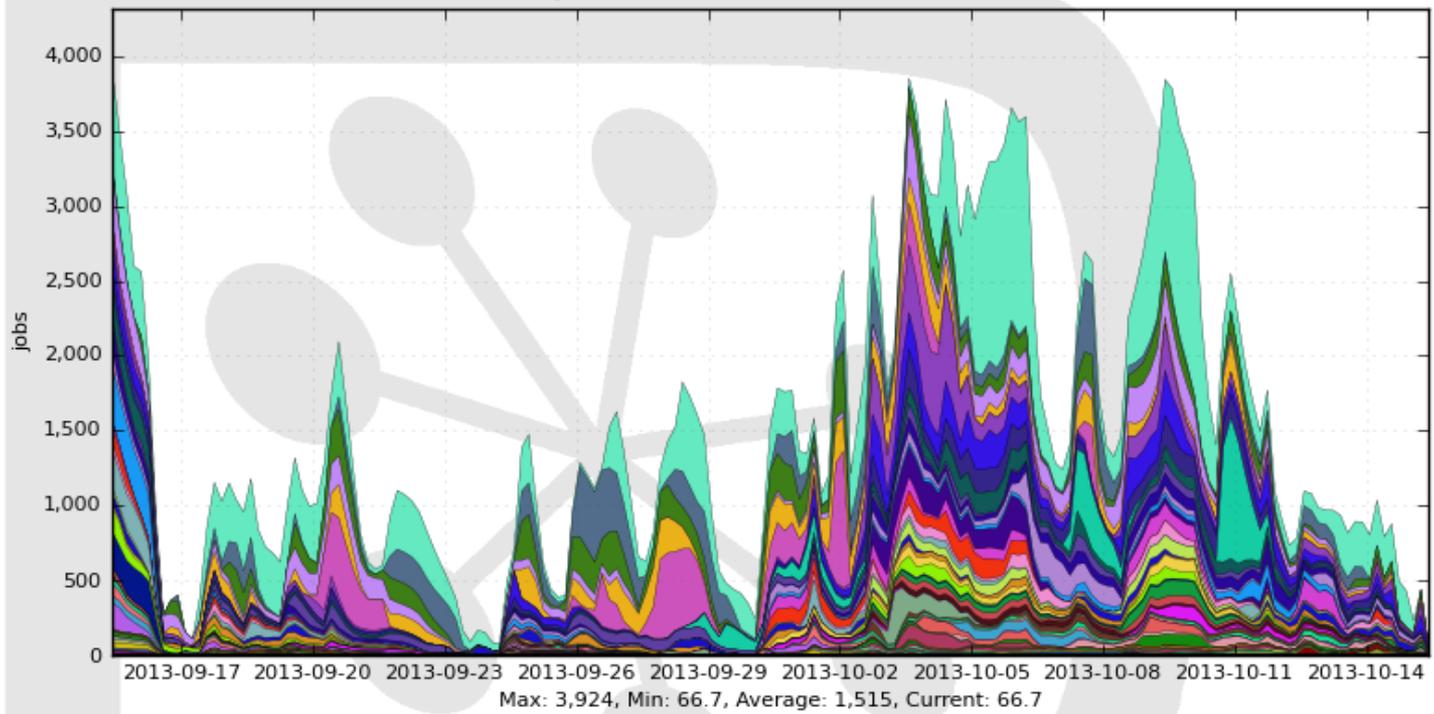
- ▶ Use of computing resources in the biomed grid community
 - ▶ DIRAC instance provided by France-Grilles since June 2012



*Tristan Glatard,
CREATIS*

Running jobs by Site

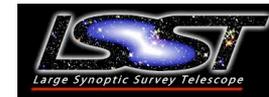
30 Days from 2013-09-15 to 2013-10-15



■ LCG.SBG.fr	22.0%	■ LCG.IPNL.fr	3.0%	■ LCG.CIEMAT.es	1.3%
■ LCG.LAL.fr	6.4%	■ LCG.UKIAC.uk	2.3%	■ LCG.UKIR.uk	1.3%
■ LCG.LPNHE.fr	6.0%	■ LCG.MSFG.fr	2.2%	■ LCG.RUG.nl	1.2%
■ LCG.BARI.it	5.4%	■ LCG.CREATIS.fr	2.2%	■ LCG.OBSPM.fr	1.1%
■ LCG.M3PEC.fr	5.1%	■ LCG.UNINA.it	2.2%	■ LCG.LILLE.fr	1.1%
■ LCG.DATAGRID.fr	4.7%	■ LCG.CNAF.it	2.1%	■ LCG.PADOVA.it	1.1%
■ LCG.IN2P3.fr	4.5%	■ LCG.UKI.uk	1.7%	■ LCG.EFDA.or	1.1%
■ LCG.SARA.nl	3.9%	■ LCG.NIKHEF.nl	1.5%	■ LCG.BEIJING.cn	0.9%
■ LCG.ROMA3.it	3.5%	■ LCG.PISA.it	1.4%	... plus 41 more	

- ▶ Heavily used for the grid tutorials
 - ▶ Using resources of the *VO france-formation*
- ▶ Support for users, applications
 - ▶ Forum for experience dissemination
 - ▶ Help in porting applications to the grid
 - ▶ Help new communities to try out DIRAC for their production systems

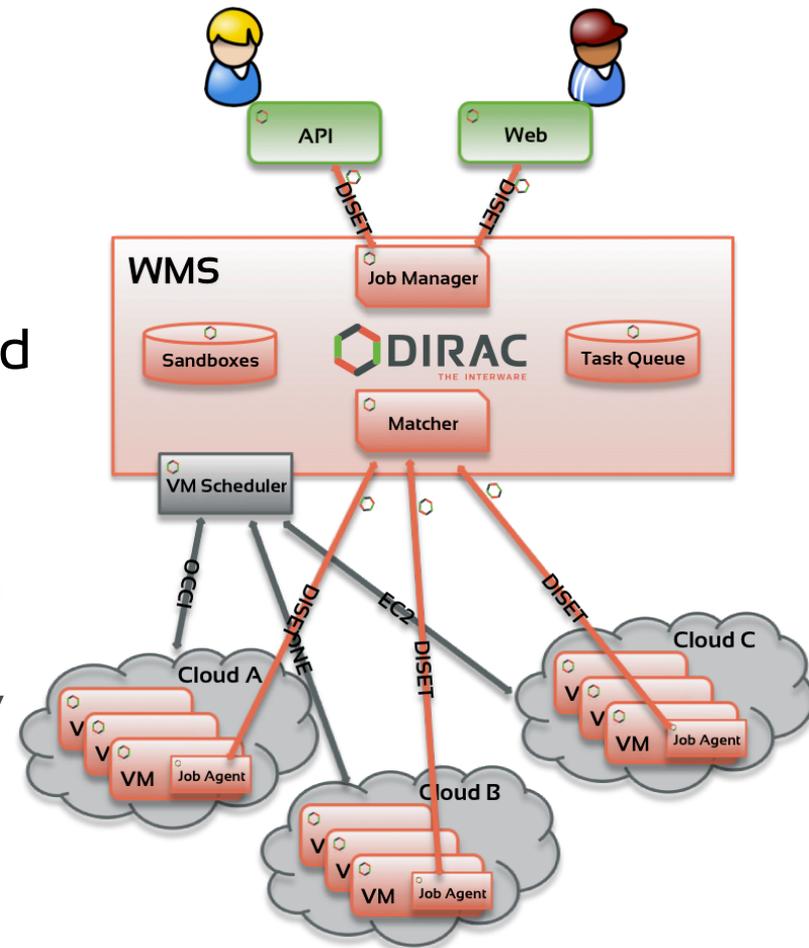
- Fermi-LAT
- LSST



Resources available via the DIRAC service

- ▶ Support for grids with based on different middlewares
 - ▶ gLite: EGI, GISELA, etc
 - ▶ VDT: OSG
 - ▶ ARC: NDGF sites, RAL, ...
 - ▶ Other types of grids can be supported
 - ▶ As requested by customers

- ▶ VM Scheduler originally developed for the Belle MC on Amazon EC2
 - ▶ Dynamic VM spawning taking into account the Task Queue state
 - ▶ Discarding VMs automatically when no more needed
 - ▶ Multiple Cloud managers/APIs are now supported
 - ▶ OpenStack, OpenNebula, CloudStack, EC2, OCCl
 - ▶ See *M. Ubeda's talk*



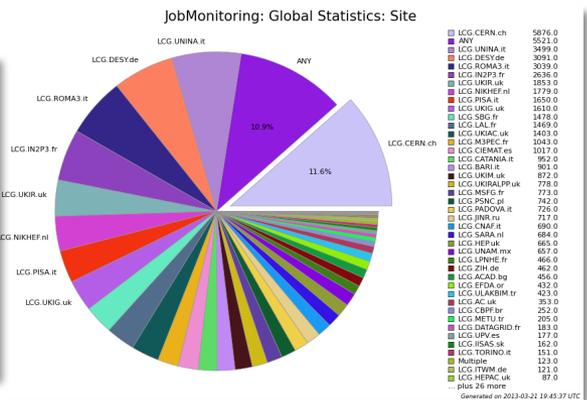
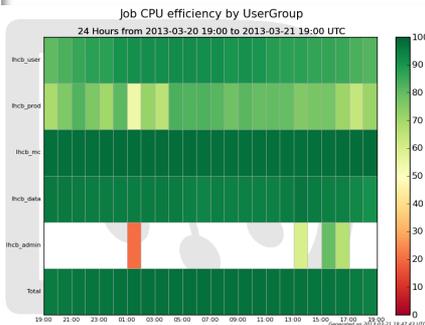
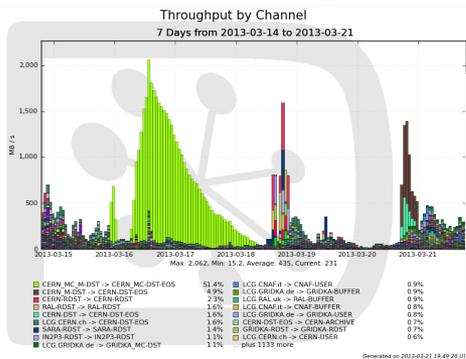
- ▶ **Non-grid sites accessed through an SSH tunnel**
 - ▶ No grid middleware installation needed on site
 - ▶ Examples:
 - ▶ DIRAC.Yandex.ru
 - 1800 cores, Torque batch system, access by SSH
 - Second largest LHCb MC production site
 - ▶ LRZ Computing Center, Munich
 - SLURM batch system, GRAM5 CE service, gateway access by GSISsh
 - Considerable resources for biomed community (work in progress)
 - ▶ HPC mesocentre Aix-Marseille University
 - OAR batch system, access by SSH
 - Open to multiple communities

- ▶ **Volunteer resources**
 - ▶ European Desktop Grid Initiative (EDGI)
 - ▶ access through a special CREAM CE service
 - ▶ BOINC based solution with virtualized client nodes

- ▶ **All the computing resources are combined transparently for the users**

Services

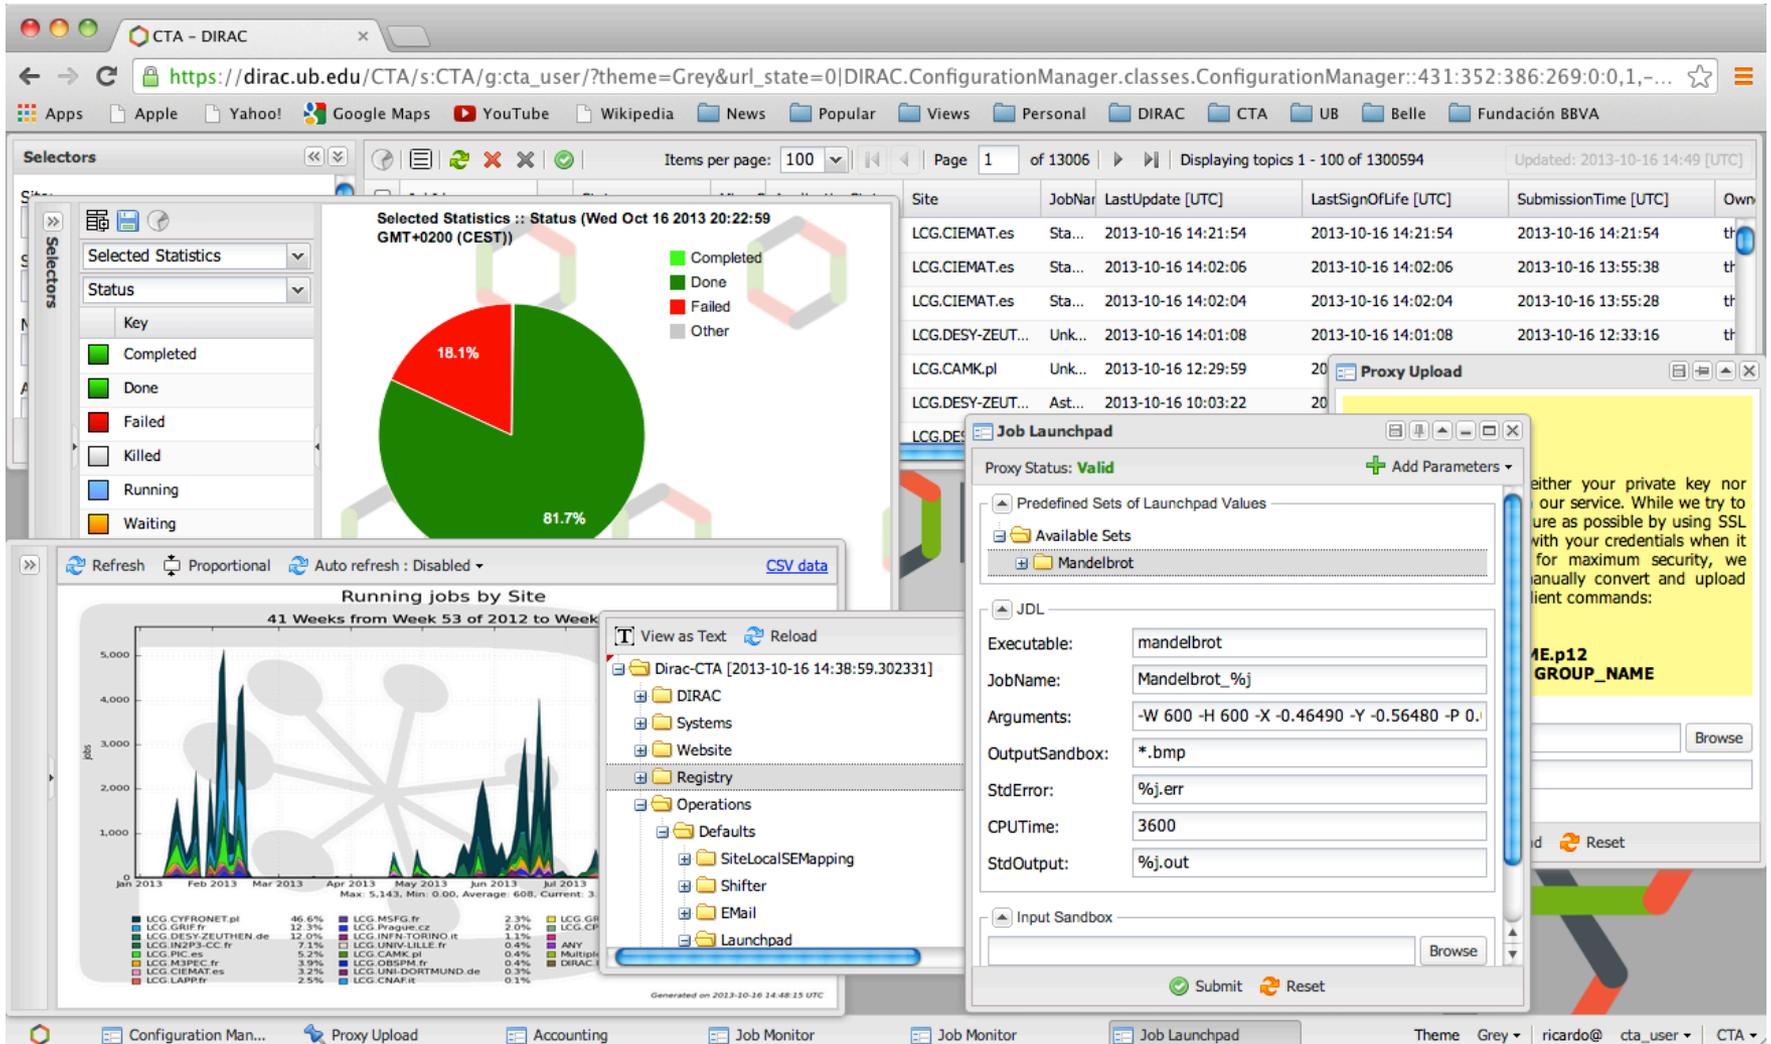
- ▶ Resources description and monitoring
- ▶ WMS – pilot based management of user jobs
 - ▶ Job submission, monitoring, retrieval
- ▶ DMS – managing user data basic tasks
 - ▶ Access to standard Grid Storage Elements
 - ▶ Providing DIRAC Storage Element in CC/Lyon
 - ▶ DIRAC File Replica and Metadata Catalog
 - ▶ “LFC+AMGA” equivalent
 - ▶ Several LFC services configured in DIRAC DMS
- ▶ Accounting
 - ▶ CPU and storage resources consumption, data transfers



- ▶ **Web Portal**
 - ▶ Support of most of the user tasks (jobs, data, monitoring, management)
 - ▶ Secure with X509 certificates

- ▶ **DIRAC RESTful interface**
 - ▶ Language neutral
 - ▶ E.g. Scala client is available (OpenMOLE Project)
 - ▶ JSAGA DIRAC plug-in is available
 - ▶ evaluated by the Science Gateway portal project

- ▶ **Other interfaces include**
 - ▶ Extensive Python API
 - ▶ E.g. used by GANGA user front-end
 - ▶ A rich set of command line tools (>200 commands)



The screenshot displays the DIRAC web portal interface, showing various monitoring and configuration tools. The browser address bar indicates the URL: `https://dirac.ub.edu/CTA/s:CTA/g:cta_user/?theme=Grey&url_state=0|DIRAC.ConfigurationManager.classes.ConfigurationManager::431:352:386:269:0:0,1,-...`

Selected Statistics :: Status (Wed Oct 16 2013 20:22:59 GMT+0200 (CEST))

Key:

- Completed (Green)
- Done (Dark Green)
- Failed (Red)
- Other (Grey)

81.7% Completed, 18.1% Failed

Running jobs by Site

41 Weeks from Week 53 of 2012 to Week 3 of 2013

Max: 5,143, Min: 0.00, Average: 608, Current: 3

Site	JobName	LastUpdate [UTC]	LastSignOfLife [UTC]	SubmissionTime [UTC]	Own
LCG.CIEMAT.es	Sta...	2013-10-16 14:21:54	2013-10-16 14:21:54	2013-10-16 14:21:54	th
LCG.CIEMAT.es	Sta...	2013-10-16 14:02:06	2013-10-16 14:02:06	2013-10-16 13:55:38	th
LCG.CIEMAT.es	Sta...	2013-10-16 14:02:04	2013-10-16 14:02:04	2013-10-16 13:55:28	th
LCG.DESY-ZEUT...	Unk...	2013-10-16 14:01:08	2013-10-16 14:01:08	2013-10-16 12:33:16	th
LCG.CAMK.pl	Unk...	2013-10-16 12:29:59	2013-10-16 12:29:59		
LCG.DESY-ZEUT...	Ast...	2013-10-16 10:03:22	2013-10-16 10:03:22		

Job Launchpad

Proxy Status: **Valid**

Predefined Sets of Launchpad Values

- Available Sets
- Mandelbrot

JDL

Executable: mandelbrot

JobName: Mandelbrot_%j

Arguments: -W 600 -H 600 -X -0.46490 -Y -0.56480 -P 0.

OutputSandbox: *.bmp

StdError: %j.err

CPUtime: 3600

StdOutput: %j.out

Input Sandbox: [Browse]

Submit Reset

4E.p12 GROUP_NAME

Reset

Configuration Man... Proxy Upload Accounting Job Monitor Job Monitor Job Launchpad Theme Grey ricardo@ cta_user CTA

- ▶ **More advanced services can be made available**
 - ▶ Following the user demands
 - ▶ Transformation Service (automated, data driven job submission)
 - ▶ Replication Service (automated data replication)
 - ▶ Data integrity inspection
 - ▶ User storage consumption accounting
 - ▶ Support for MPI jobs
 - ▶ ...
- ▶ **Hosting Community DIRAC services**
 - ▶ Specific services developed in the DIRAC framework can be hosted in the same infrastructure

- ▶ **GISELA Latin American grid**
 - ▶ In production since 2010
 - ▶ Since 2012 GISELA DIRAC services are provided by France-Grid
- ▶ **Ibergrid Spanish/Portugal NGI**
 - ▶ Magic, Gaussian experiments, ...
- ▶ **DIRAC services in an evaluation/start-up phase**
 - ▶ GridPP, DIRAC installation in Imperial College
 - NA62, T2K, LondonGrid, ...
 - ▶ IGI, CNAF
 - ▶ CNGrid, IHEP, Beijing
 - ▶ BOINC, ex-GOS sites, IHEP supercomputing centre
 - ▶ TREND, Daya Bay Neutrino Experiment
- ▶ **ILC/CLIC+CALICE multi-VO installation at CERN**
 - ▶ Considering GEANT4 VO to join this service
- ▶ **More projects in testing and/or discussion:**
 - ▶ Ukraine, Russia, ...



- ▶ Distributed computing is no more something exotic, it is used in a daily work by users in various scientific domains
- ▶ LHC experience and tools should now be shared with other user communities
- ▶ DIRAC provides a framework for building distributed computing systems and a rich set of ready to use services.
- ▶ There is an increasing number of regional and national DIRAC service projects
- ▶ DIRAC can help users to get started in the world of distributed computing and discover its full potential

