

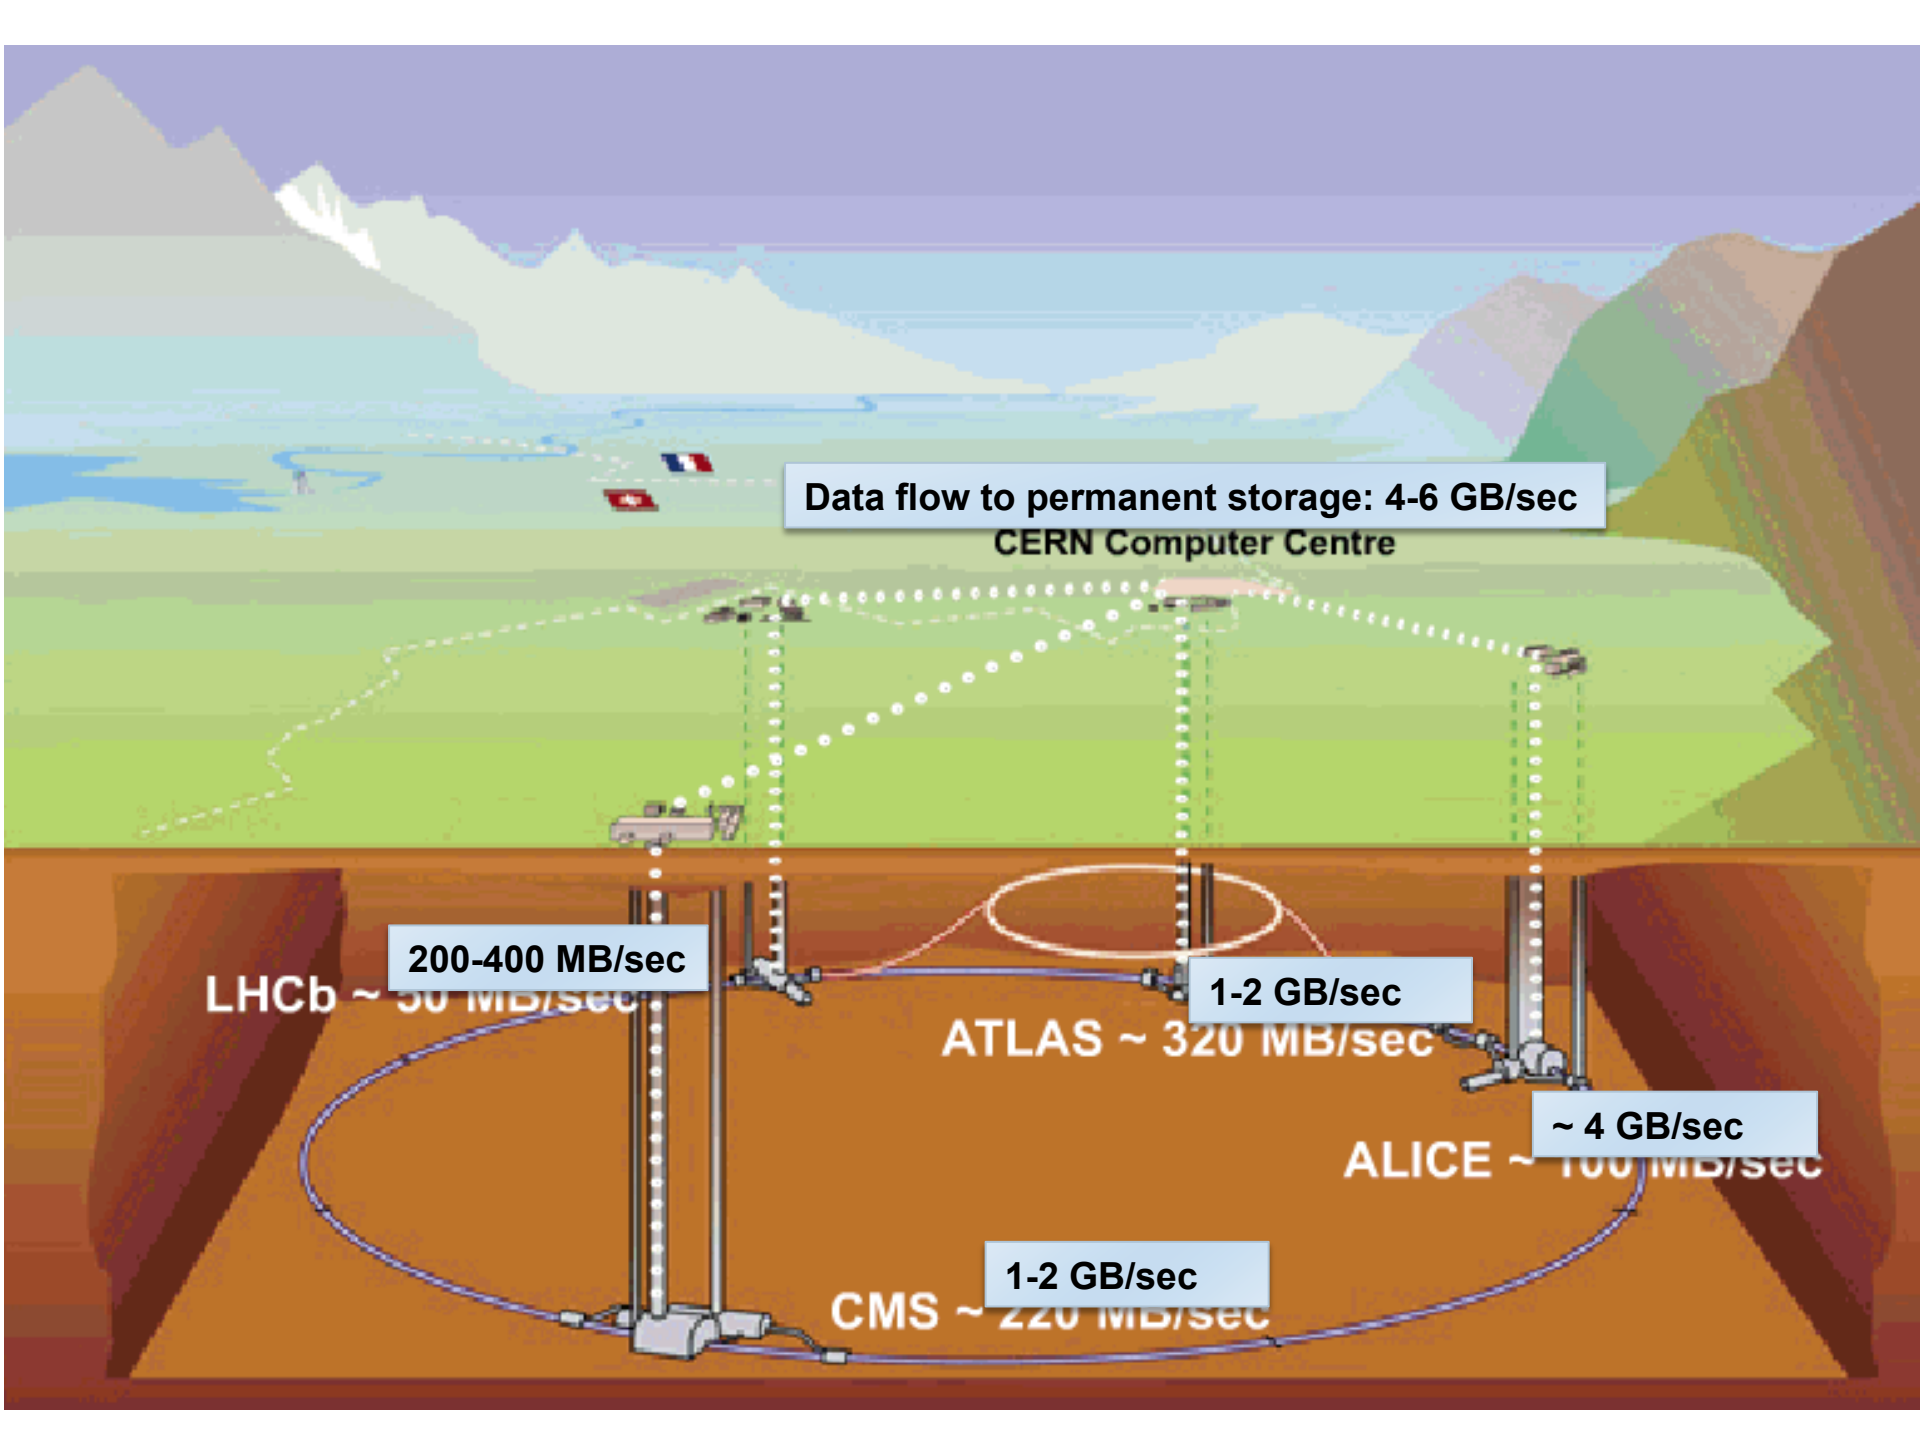
DIRAC Services for EGI

*A. Tsaregorodtsev,
CPPM-IN2P3-CNRS*

Grid 2014, Dubna, 1 July 2014



- ▶ DIRAC Project
- ▶ DIRAC as a Service
- ▶ DIRAC Service for EGI
- ▶ Conclusions



Data flow to permanent storage: 4-6 GB/sec
CERN Computer Centre

200-400 MB/sec
LHCb ~ 50 MB/sec

1-2 GB/sec
ATLAS ~ 320 MB/sec


~ 4 GB/sec
ALICE ~ 100 MB/sec

1-2 GB/sec
CMS ~ 220 MB/sec

- ▶ 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 has all the necessary components to build ad-hoc grid infrastructures **interconnecting** computing 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

- ▶ Belle II

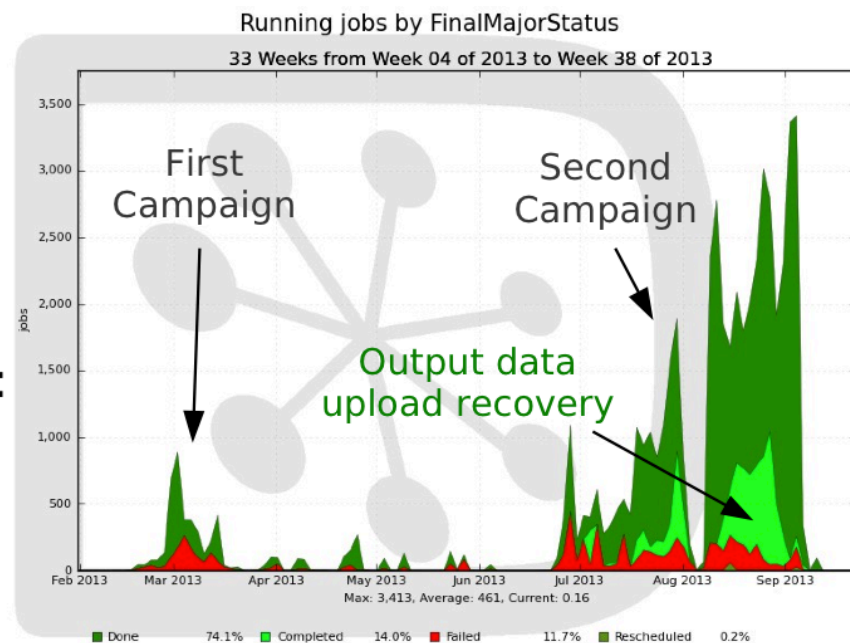
- ▶ Combination of the non-grid, grid sites and (commercial) clouds is a requirement

- ▶ 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



Thomas Kuhr, Belle II



- ▶ ILC/CLIC detector Collaboration, Calice VO
 - ▶ MC simulations
 - ▶ DIRAC File Catalog was developed to meet the ILC/CLIC requirements



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
- ▶ 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 VOs, ~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
 - ▶ >10 millions jobs went through the system
 - Mostly biomed applications



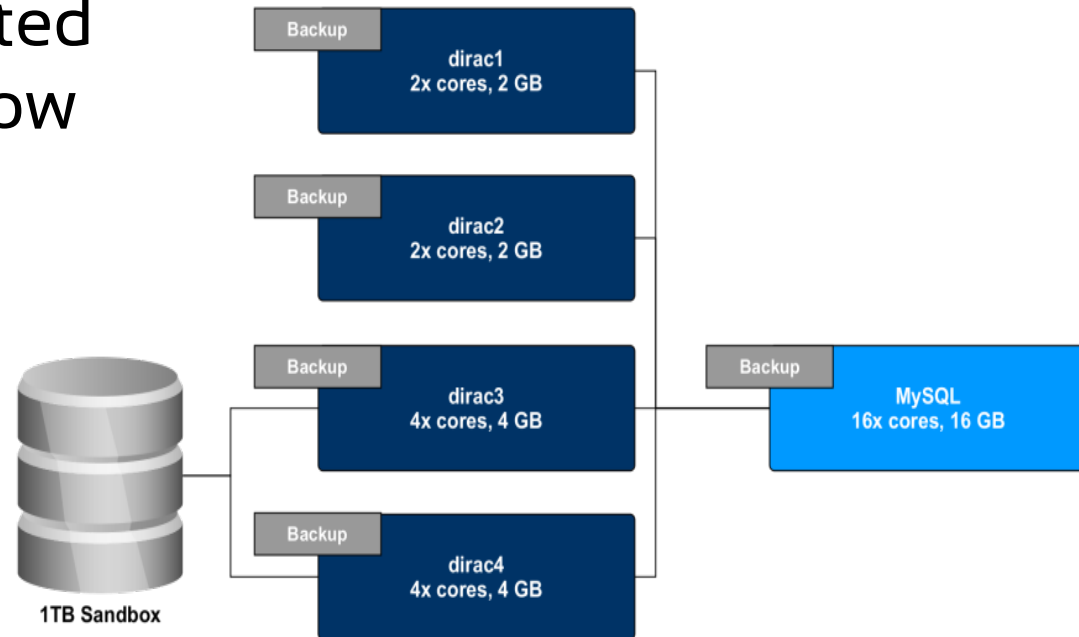
- ▶ **GISELA Latin American grid**
 - ▶ In production since 2010
 - ▶ Since 2012 GISELA DIRAC services are provided by France-Grid
- ▶ **Ibergrid Spanish/Portugal NGI**
 - ▶ Magic, Gaussian experiments, ...
- ▶ **DIRAC services in an evaluation/start-up phase**
 - ▶ GridPP, DIRAC installation in Imperial 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, ...



DIRAC service for EGI

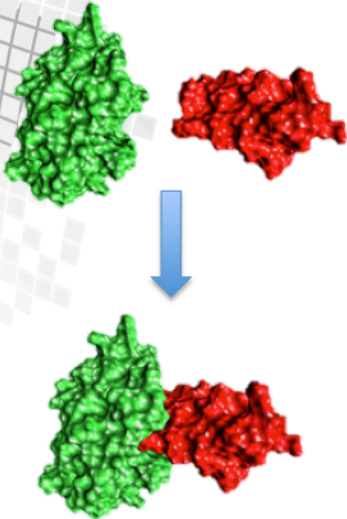
- ▶ The DIRAC4EGI service is largely inspired by the success of the FG-DIRAC Project
 - ▶ Built following similar principles
 - ▶ Joint project of several NGL's
- ▶ The DIRAC4EGI is operated by the EGI.eu Project
 - ▶ Overall coordination, negotiations with the client VO's
- ▶ Maintenance and administration coordinated by France-Grilles
 - ▶ Operations shifts, training of administrators
 - ▶ User support
 - ▶ Documentation, wiki space, etc
- ▶ Team of administrators from institutions in France, Poland, Italy, etc.

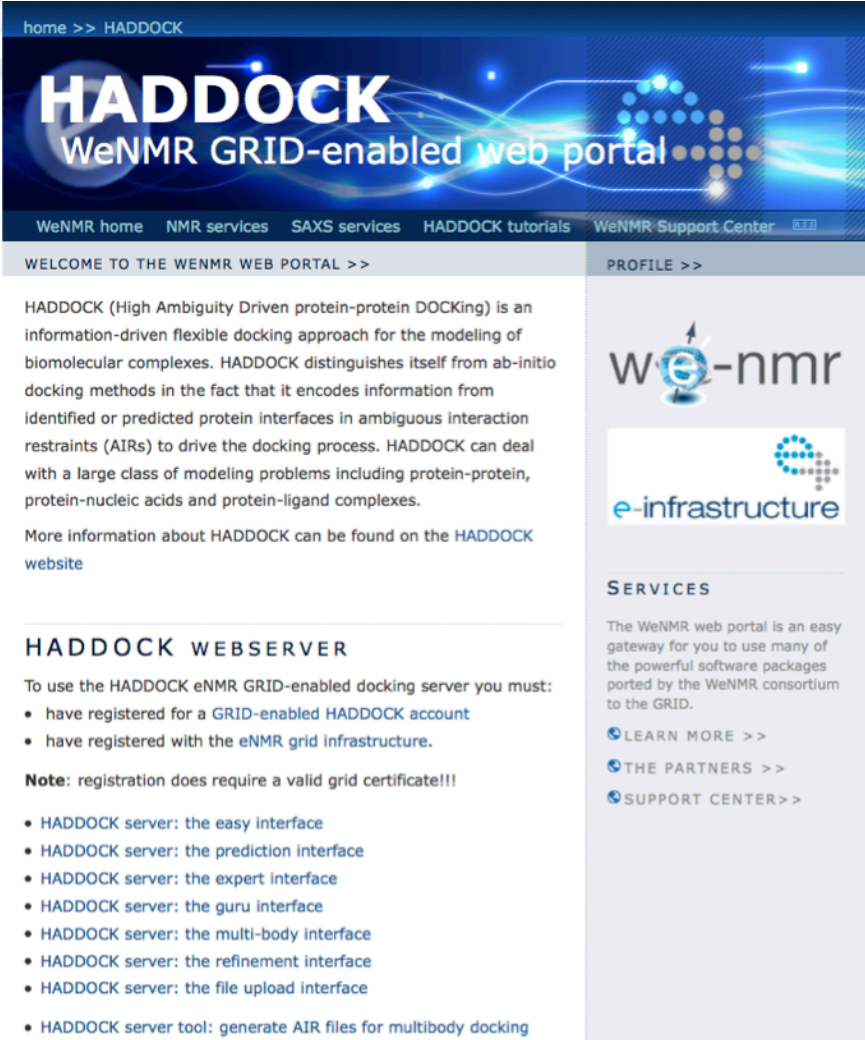
- ▶ The services are hosted by CYFRONET, Krakow
 - ▶ 4 virtual servers
 - ▶ Special choice of powerful hardware
 - ▶ MySQL DB services (regular backups)
 - ▶ 1TB DIRAC Storage
 - ▶ Sandboxes, DIRAC SE
 - ▶ Web portal <https://dirac.egi.eu>



- ▶ The service was launched in February 2014 in a prototype mode
 - ▶ First VO enmr.eu, WeNMR Haddock application
 - ▶ Setting up the services, describing resources and users
- ▶ Starting from May 2014 the service is in production
 - ▶ Reported in the EGI CF:
 - ▶ More users to come
- ▶ Currently serving VOs
 - ▶ enmr.eu, vlemmed, biomed
 - ▶ More to come
- ▶ Main contact points
 - ▶ dirac@mailman.egi.eu
 - ▶ dirac-admin@mailman.egi.eu for operations

Haddock web portal





home >> HADDOCK

HADDOCK

WeNMR GRID-enabled web portal

WeNMR home NMR services SAXS services HADDOCK tutorials WeNMR Support Center

WELCOME TO THE WENMR WEB PORTAL >>

HADDOCK (High Ambiguity Driven protein-protein DOCKing) is an information-driven flexible docking approach for the modeling of biomolecular complexes. HADDOCK distinguishes itself from ab-initio docking methods in the fact that it encodes information from identified or predicted protein interfaces in ambiguous interaction restraints (AIRs) to drive the docking process. HADDOCK can deal with a large class of modeling problems including protein-protein, protein-nucleic acids and protein-ligand complexes.

More information about HADDOCK can be found on the [HADDOCK website](#)

HADDOCK WEBSERVER



To use the HADDOCK eNMR GRID-enabled docking server you must:

- have registered for a **GRID-enabled HADDOCK account**
- have registered with the **eNMR grid infrastructure**.

Note: registration does require a valid grid certificate!!!

- HADDOCK server: the easy interface
- HADDOCK server: the prediction interface
- HADDOCK server: the expert interface
- HADDOCK server: the guru interface
- HADDOCK server: the multi-body interface
- HADDOCK server: the refinement interface
- HADDOCK server: the file upload interface
- HADDOCK server tool: generate AIR files for multibody docking

PROFILE >>






SERVICES

The WeNMR web portal is an easy gateway for you to use many of the powerful software packages ported by the WeNMR consortium to the GRID.

- [LEARN MORE >>](#)
- [THE PARTNERS >>](#)
- [SUPPORT CENTER >>](#)

HADDOCK goes DIRAC

- **DIRAC submission enabled at minimum cost!**
 - In one afternoon, thanks to the help of Ricardo  and Andrei 
 - Clone of the HADDOCK server on a different machine
 - No root access required, no EMI software installation required
- **Minimal changes to our submission and polling scripts**
 - Requirements and ranking no longer needed, only CPUTime


```
JobName = "dirac-xxx";
CPUTime = 100000;
Executable = "dirac-xxx.sh";
StdOutput = "dirac-xxx.out";
StdError = "dirac-xxx.err";
InputSandbox = {"dirac-xxx.sh", "dirac-xxx.tar.gz"};
OutputSandbox = {"dirac-xxx.out", "dirac-xxx.err", "dirac-xxx-result.tar.gz"};
```

- **Very efficient submission (~2s per job – without changing our submission mechanism), high job throughput**



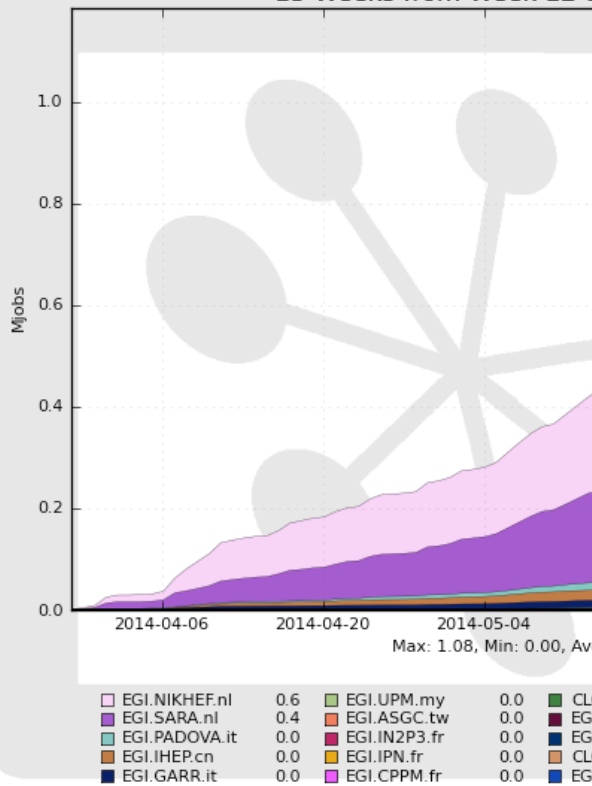
Some WeNMR statistics

HADDOCK server grid jobs (total 2014 ~2.5M / DIRAC~1.35M)




Cumulative Jobs by Site

13 Weeks from Week 12 of 2014 to Week 26 of 2014



Jobs by FinalMajorStatus

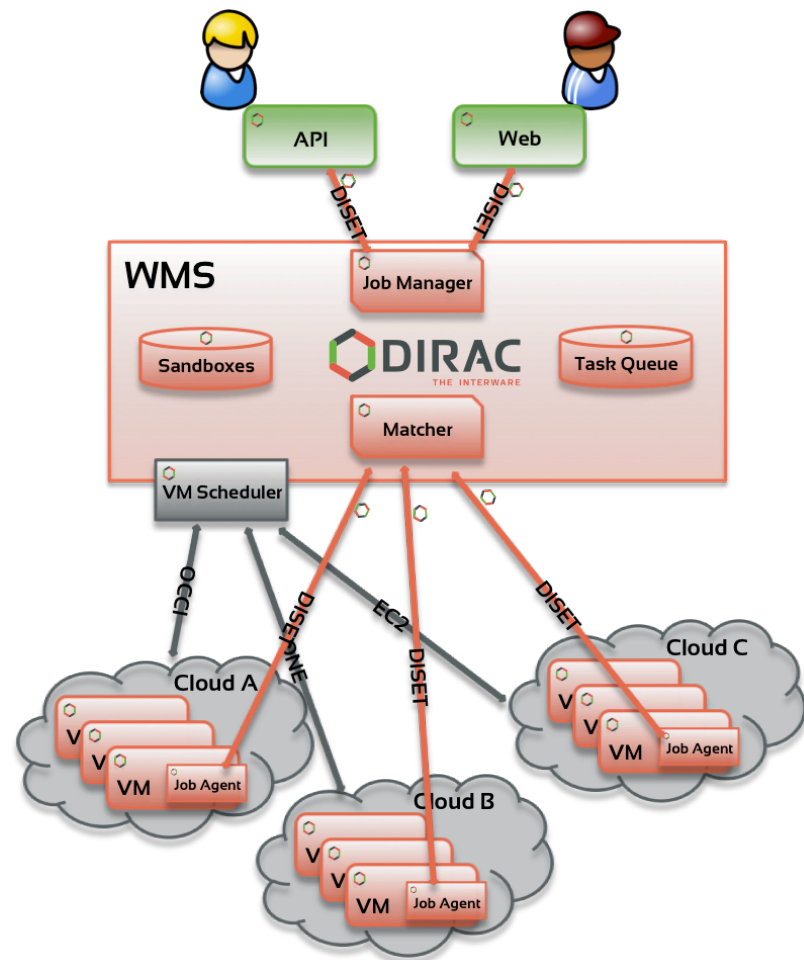
30 Days from 2014-05-31 to 2014-06-30



Resources available via the DIRAC service

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

- ▶ Support for cloud resources will be added
 - ▶ FedCloud
 - ▶ 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




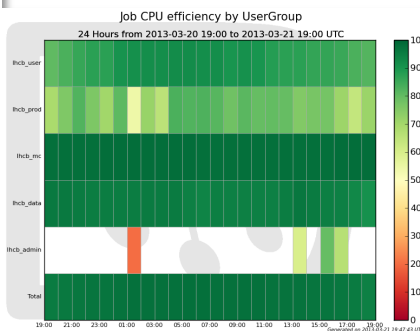

- ▶ Non-grid sites accessed through an SSH tunnel
 - ▶ No grid middleware installation needed on site
 - ▶ Examples:
 - ▶ DIRAC.Yandex.ru
 - 1800 cores, Torque batch system, access by SSH
 - Second largest LHCb MC production site
 - ▶ HPC mesocentre Aix-Marseille University
 - OAR batch system, access 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

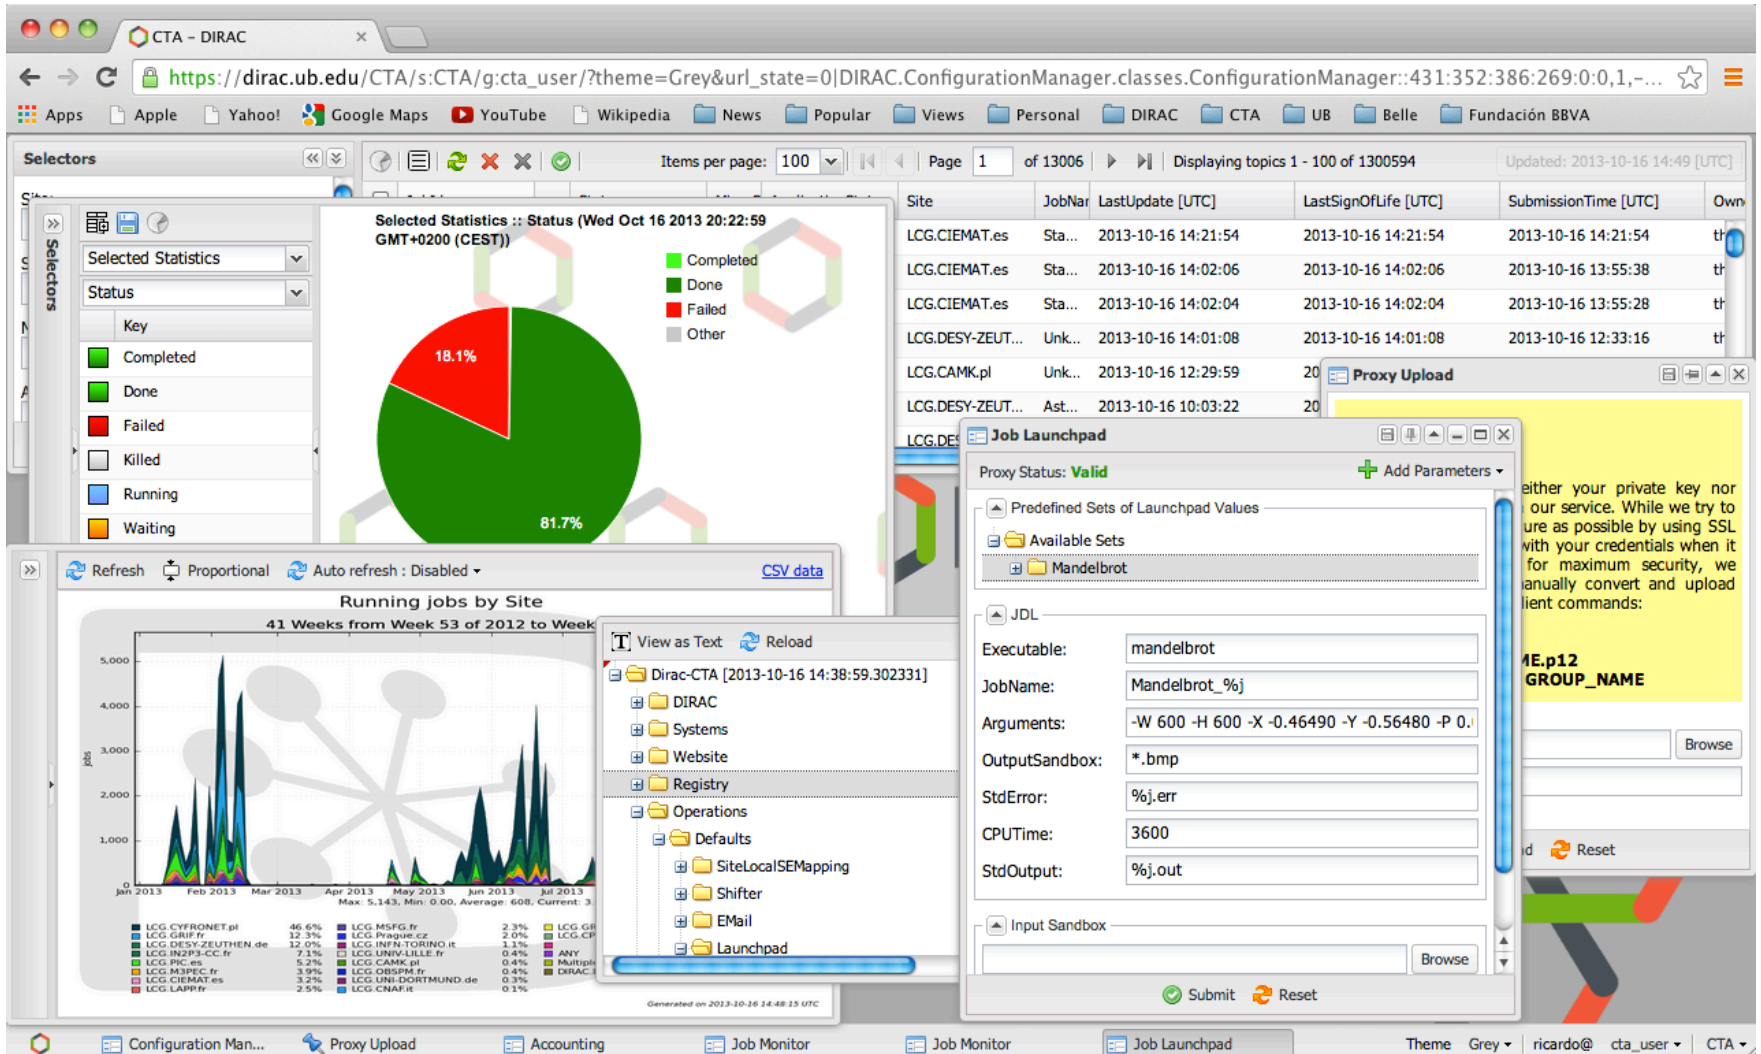
- ▶ Resources description and monitoring
- ▶ WMS – pilot based management of user jobs
 - ▶ Job submission, monitoring, retrieval
- ▶ DMS – managing user data basic tasks
 - ▶ Access to standard Grid Storage Elements
 - ▶ Providing DIRAC Storage Element in CYFRONET
 - ▶ DIRAC File Replica and Metadata 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,-...`

The main interface is divided into several sections:

- Selected Statistics :: Status (Wed Oct 16 2013 20:22:59 GMT+0200 (CEST))**: A pie chart showing job status distribution:
 - Completed: 81.7%
 - Failed: 18.1%
 - Other: 0.2%
- Running jobs by Site**: A line graph showing job counts over 41 weeks from Week 53 of 2012 to Week 3 of 2013. The graph shows significant peaks in job counts, particularly in late 2012 and early 2013. A legend below the graph lists various sites and their percentages:

LCG.CYFRONET.pl	46.6%	LCG.MSFG.fr	2.3%	LCG.GR
LCG.GRIF.fr	12.3%	LCG.Prague.cz	2.0%	LCG.CP
LCG.DESY-ZEUTHEN.de	12.0%	LCG.INFN-TORINO.it	1.1%	LCG.CP
LCG.BRIP3-CC.fr	7.3%	LCG.UNIV-LILLE.fr	0.4%	ANY
LCG.PIC.es	5.2%	LCG.CAMK.pl	0.4%	Multiple
LCG.M3PEC.fr	3.9%	LCG.OBSPM.fr	0.4%	DIRAC
LCG.CIEMAT.es	3.2%	LCG.UNI-DORTMUND.de	0.3%	
LCG.LAPP.fr	2.5%	LCG.CNAF.it	0.1%	
- Job Launchpad**: A configuration window for a job named "Mandelbrot". It shows the executable, job name, arguments, output sandbox, and CPU time. The status is "Valid".
- Proxy Upload**: A window for uploading a proxy key, with a warning about SSL and security.
- Job Launcher**: A table listing job details such as Site, JobName, LastUpdate [UTC], LastSignOfLife [UTC], SubmissionTime [UTC], and Owner.

- ▶ **More advanced services can be made available**
 - ▶ Following the user demands
 - ▶ Transformation Service (automated, data driven job submission, workflows)
 - ▶ 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

- ▶ LHC experience and tools for using distributed computing resources 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
- ▶ DIRAC4EGI service can help users to get started in the world of distributed computing and discover its full potential

