

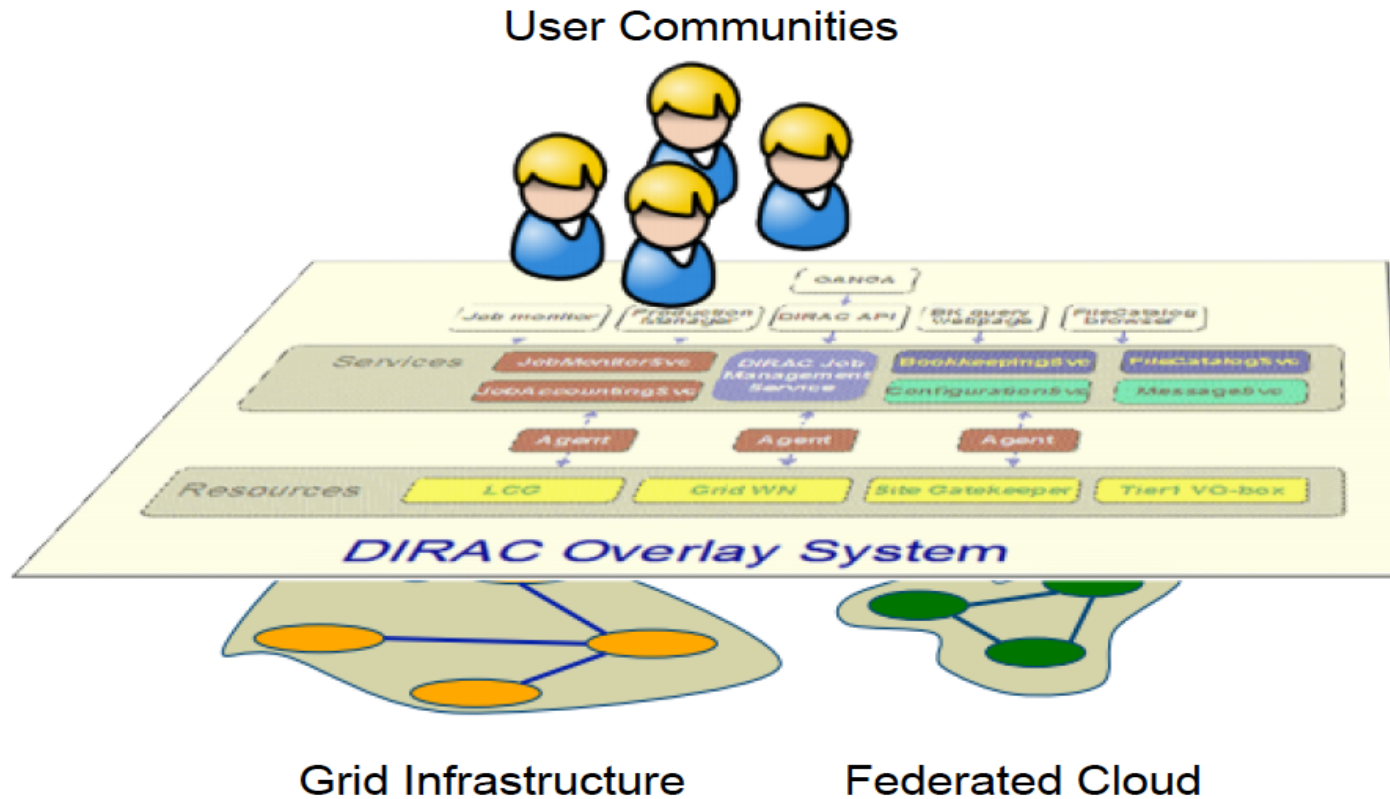
Running Scientific Applications with DIRAC in Federated Clouds


*Víctor Méndez Muñoz
PIC, UAB*

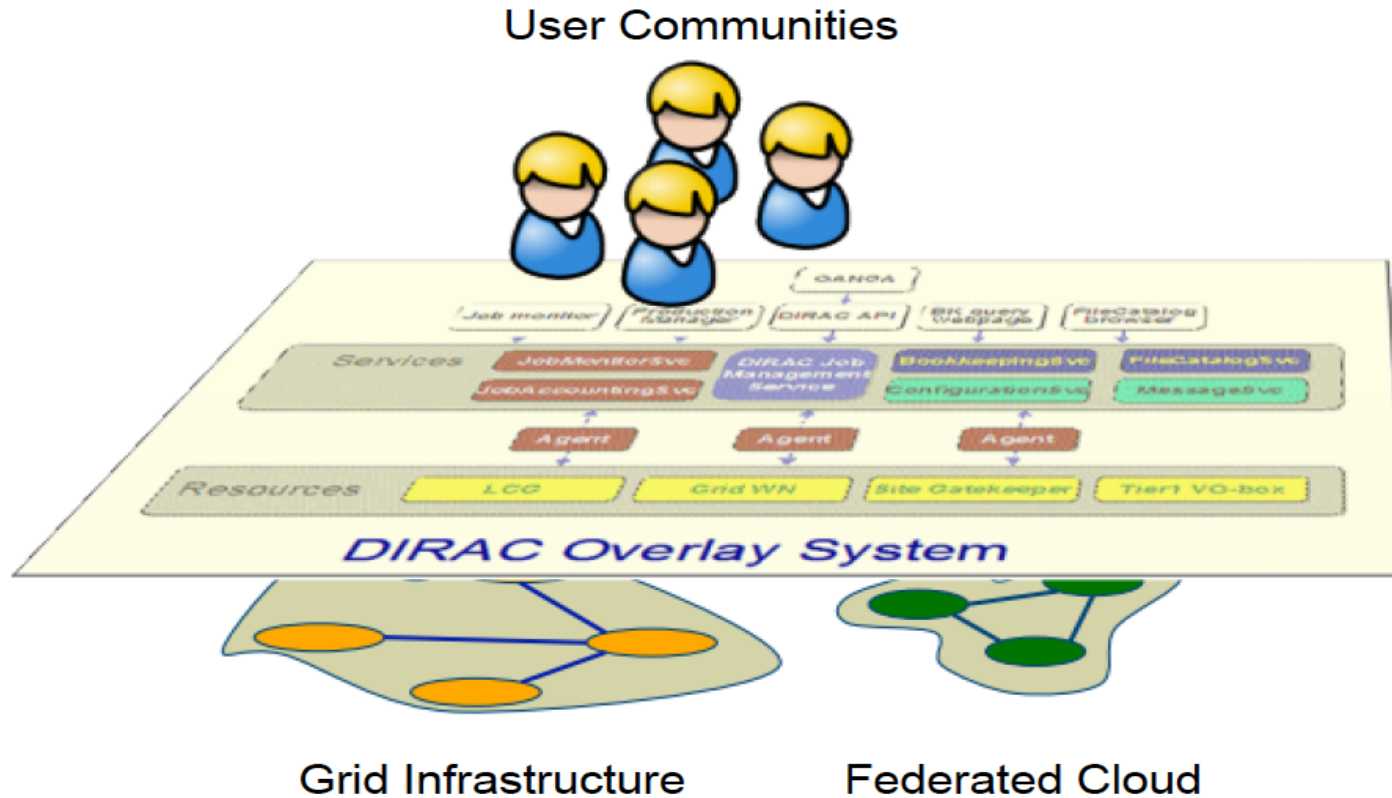
ADVCOMP, 30 September 2013



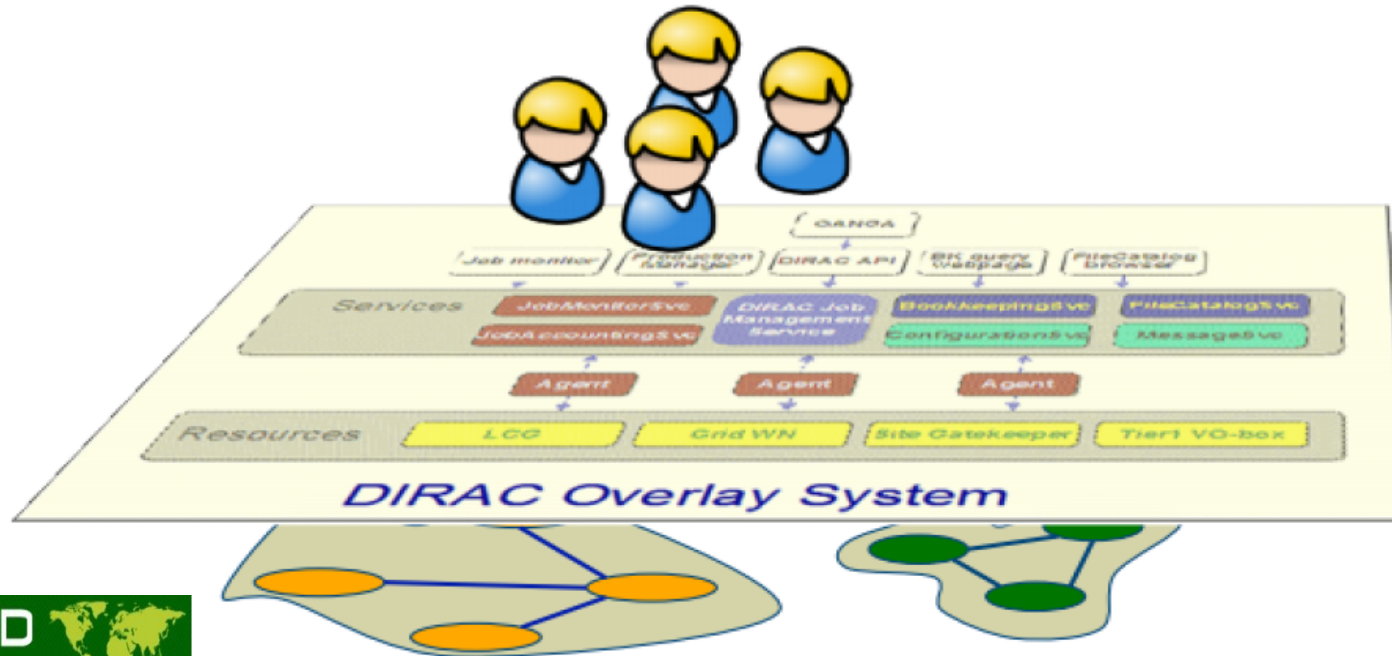
- DIRAC overview
- Federated hybrid cloud computing model for scientific communities
- VMDIRAC: the Cloud extension of DIRAC
- VMDIRAC v0r9 main features
- Next milestones of VMDIRAC roadmap
- Current use cases of VMDIRAC
- Conclusions







User Communities




Grid Infrastructure

Federated Cloud



User Communities



Grid Infrastructure


Federated Cloud



LHCb/WLCG Cloud
 EGI Federated Cloud
 FG Federated Cloud



User Communities



Grid Infrastructure


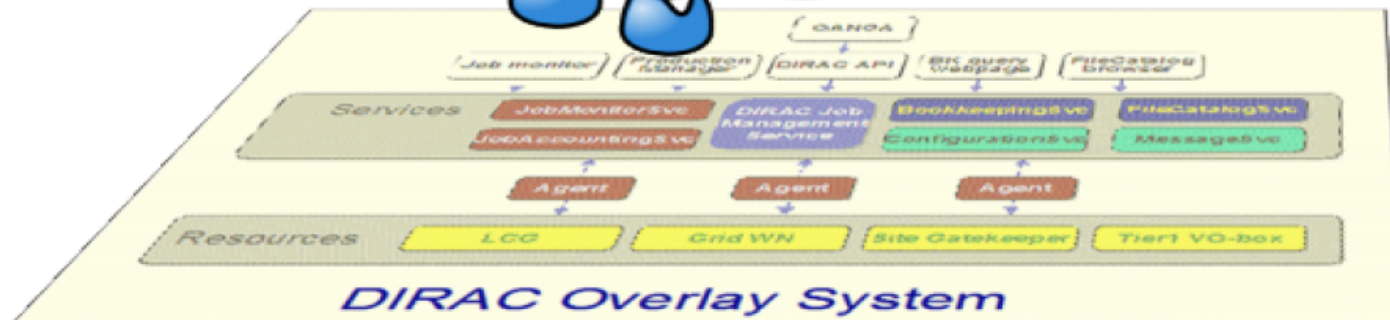
Federated Cloud



LHCb/WLCG Cloud
 EGI Federated Cloud
 FG Federated Cloud



User Communities



Grid Infrastructure

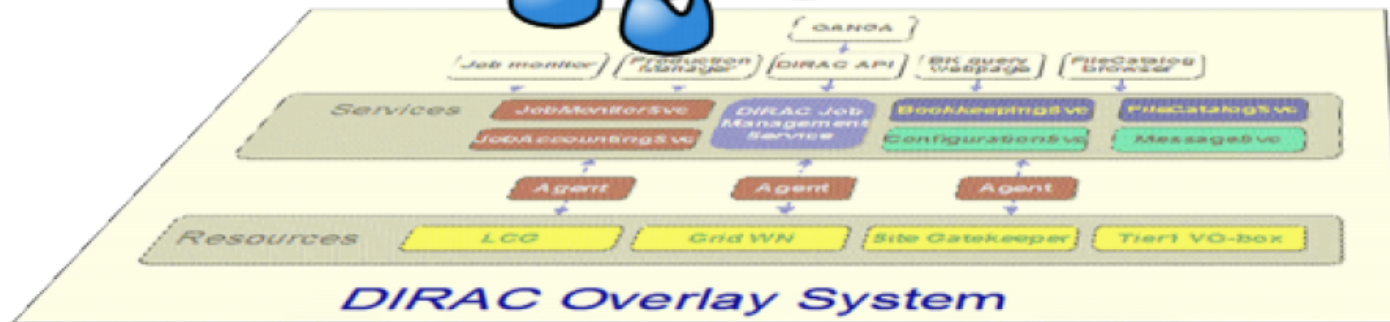
Federated Cloud



LHCb/WLCG Cloud
 EGI Federated Cloud
 FG Federated Cloud



User Communities



Grid Infrastructure

Federated Cloud



LHCb/WLCG Cloud
 EGI Federated Cloud
 FG Federated Cloud

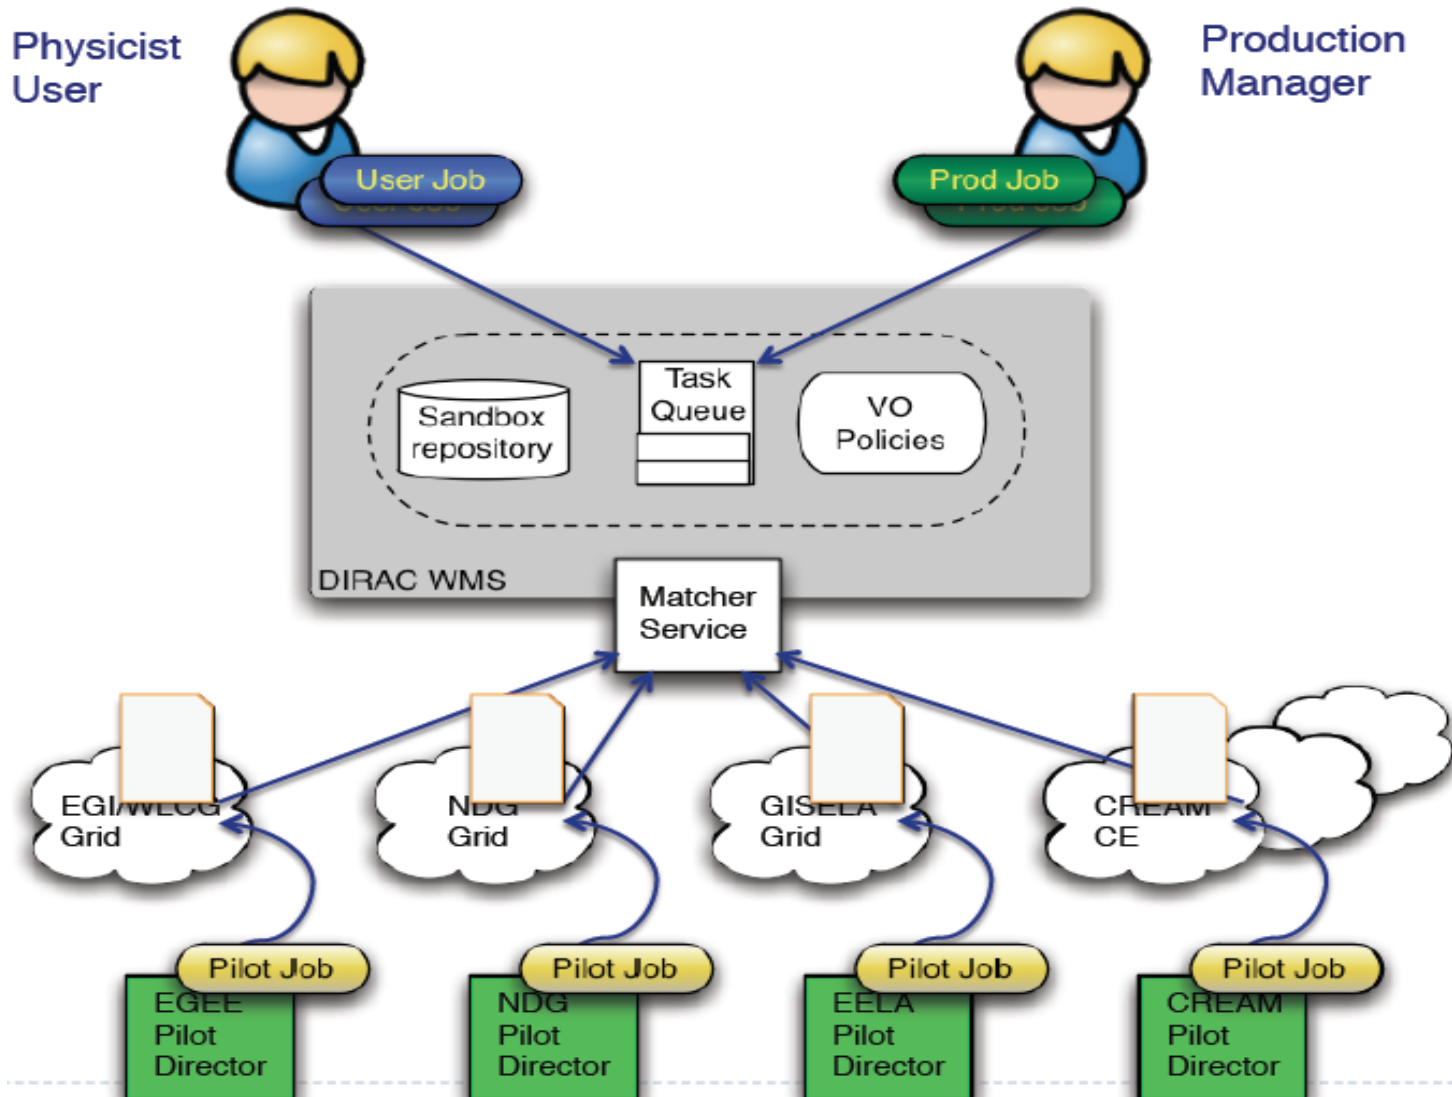
- DIRAC has all the necessary components to build ad-hoc distributed infrastructures interconnecting computing resources of different types

- DIRAC has all the necessary components to build ad-hoc distributed infrastructures interconnecting computing resources of different types

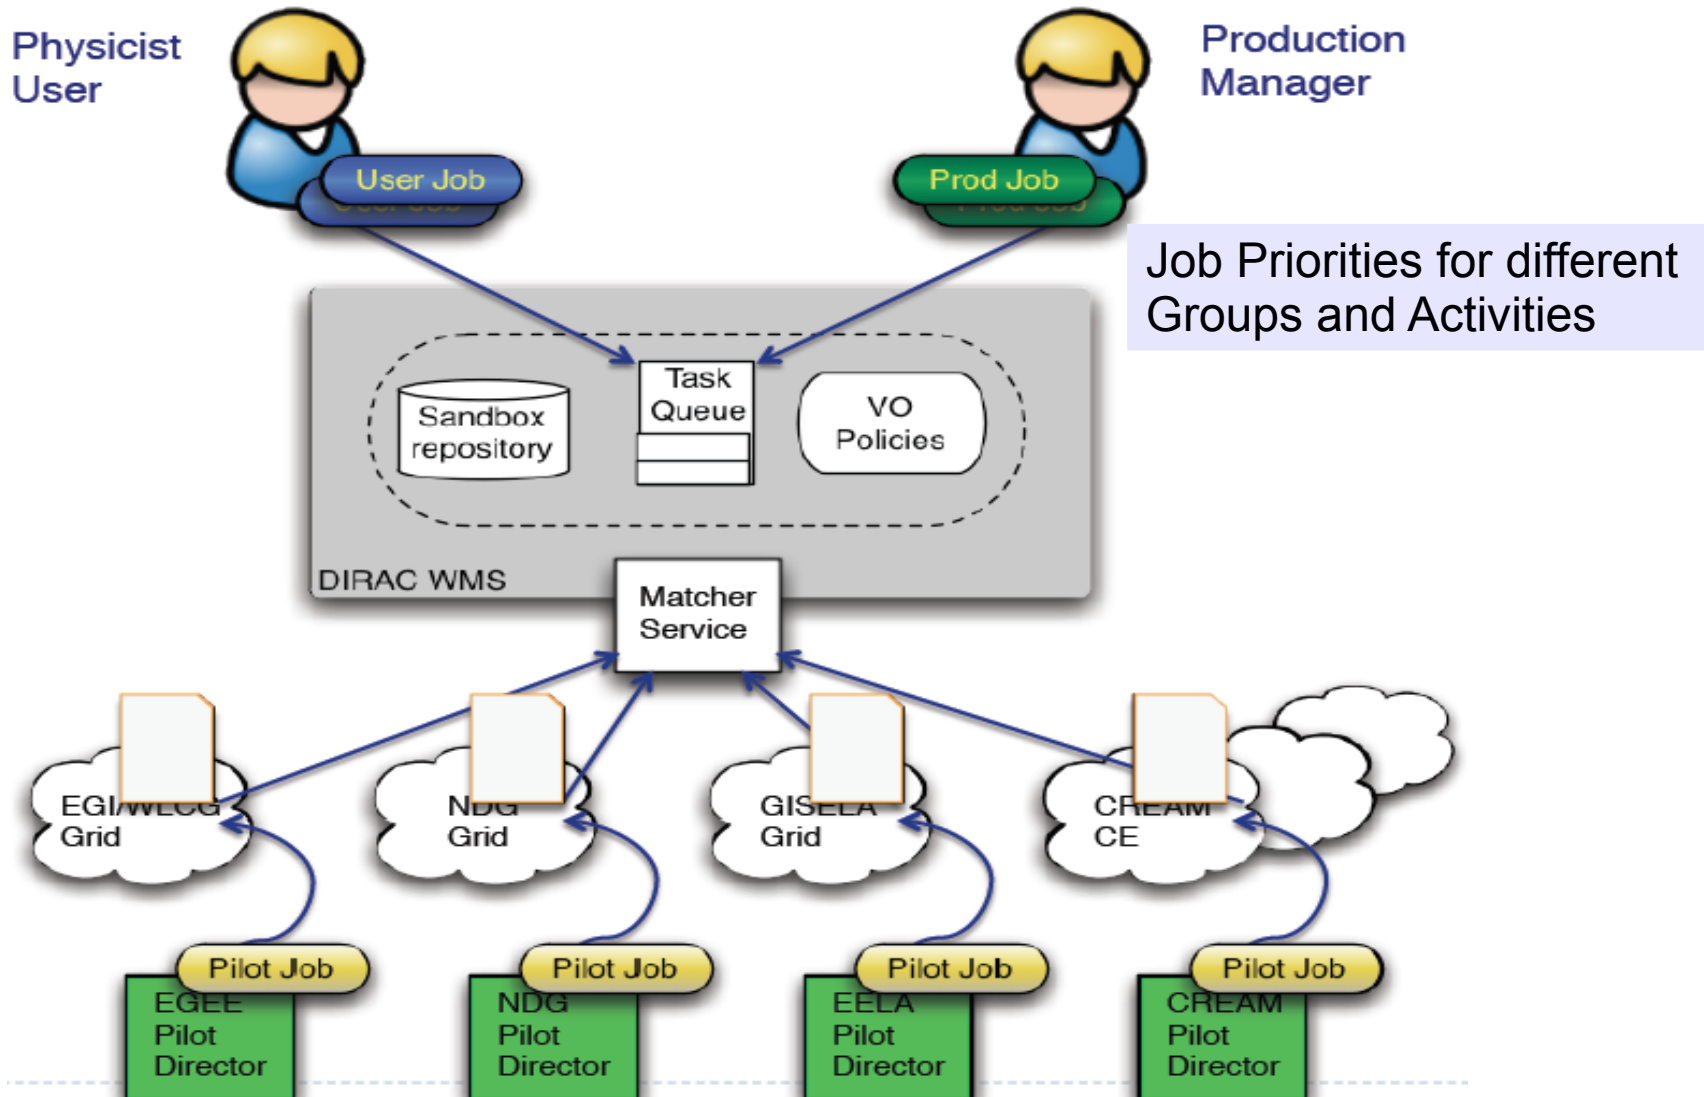
THE INTERWARE

- Furthermore of grid and cloud resources, DIRAC also integrates:
 - Standalone computing clusters:
 - LSF, BQS, SGE, PBS/Torque, Condor"
 - More to come: OAR, SLURM, LoadLeveler. Etc
 - Volunteer grids based on BOINC + virtualization technology
 - Third party Virtual Clusters (KVM) with *vacuum* resource allocation, from the grid site job queue instead of DIRAC job queue.

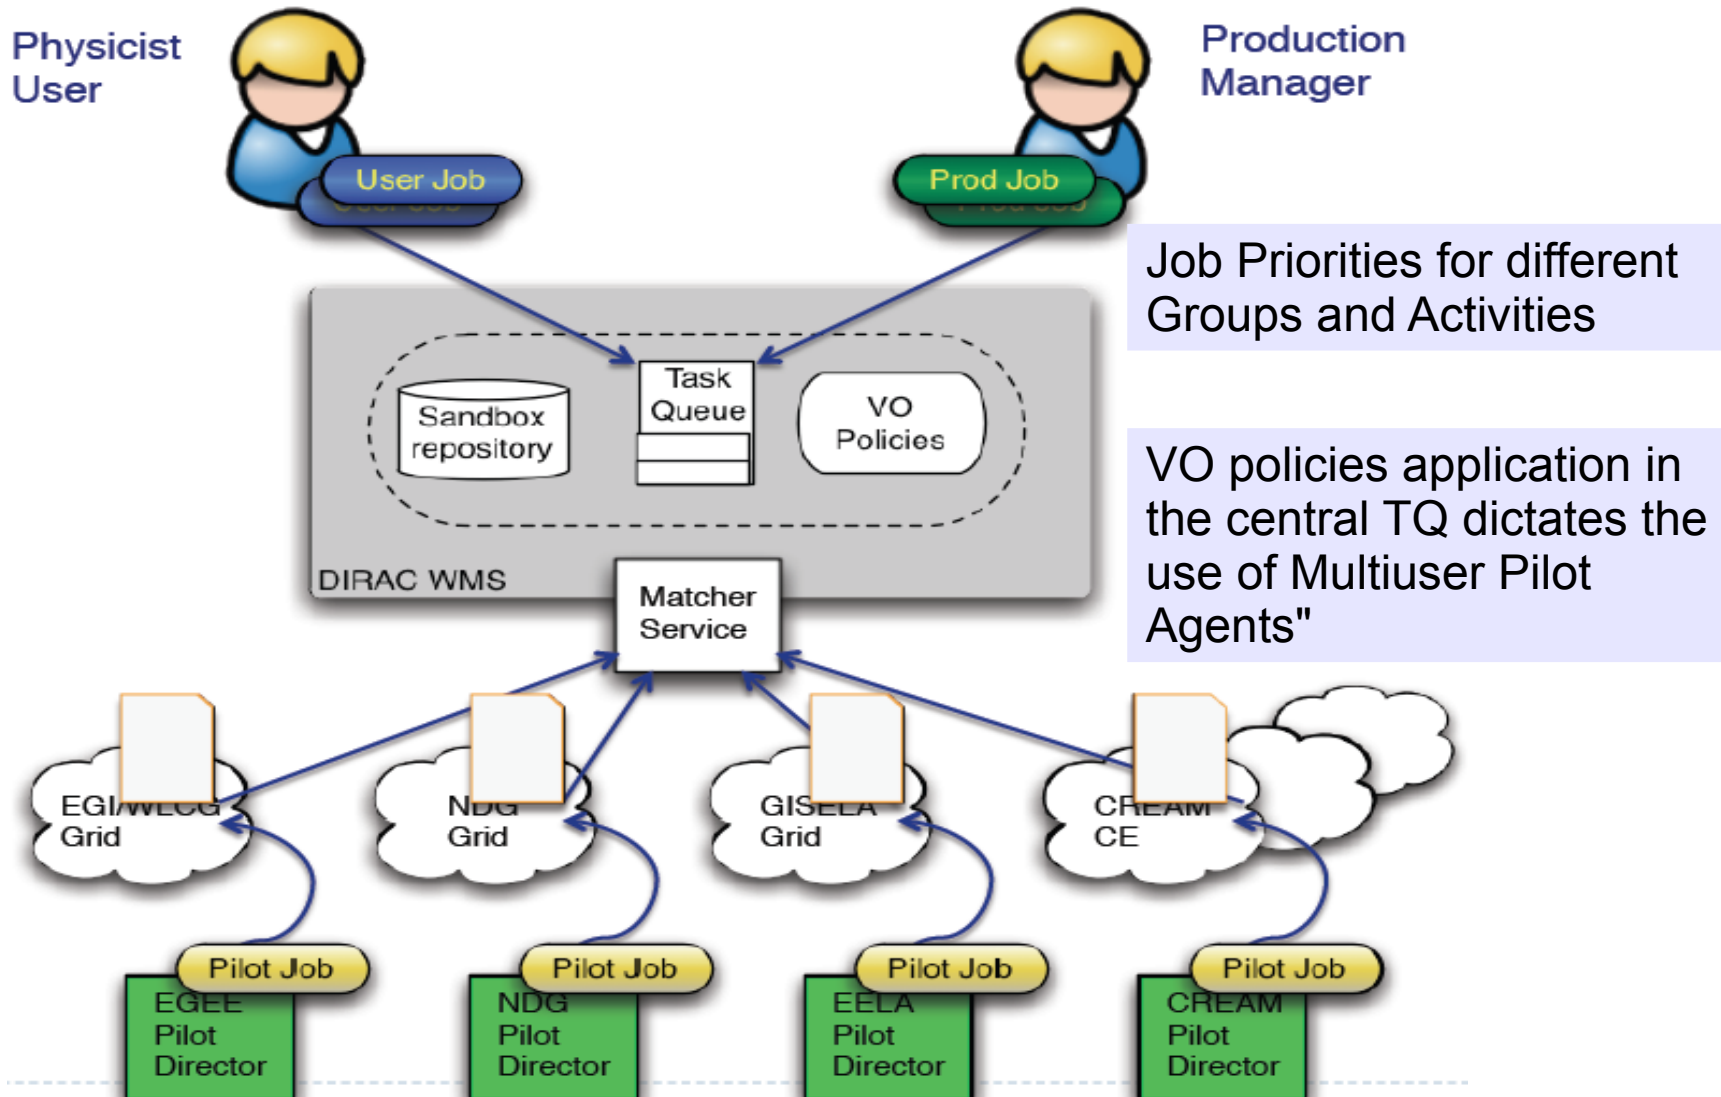
DIRAC Workload Management



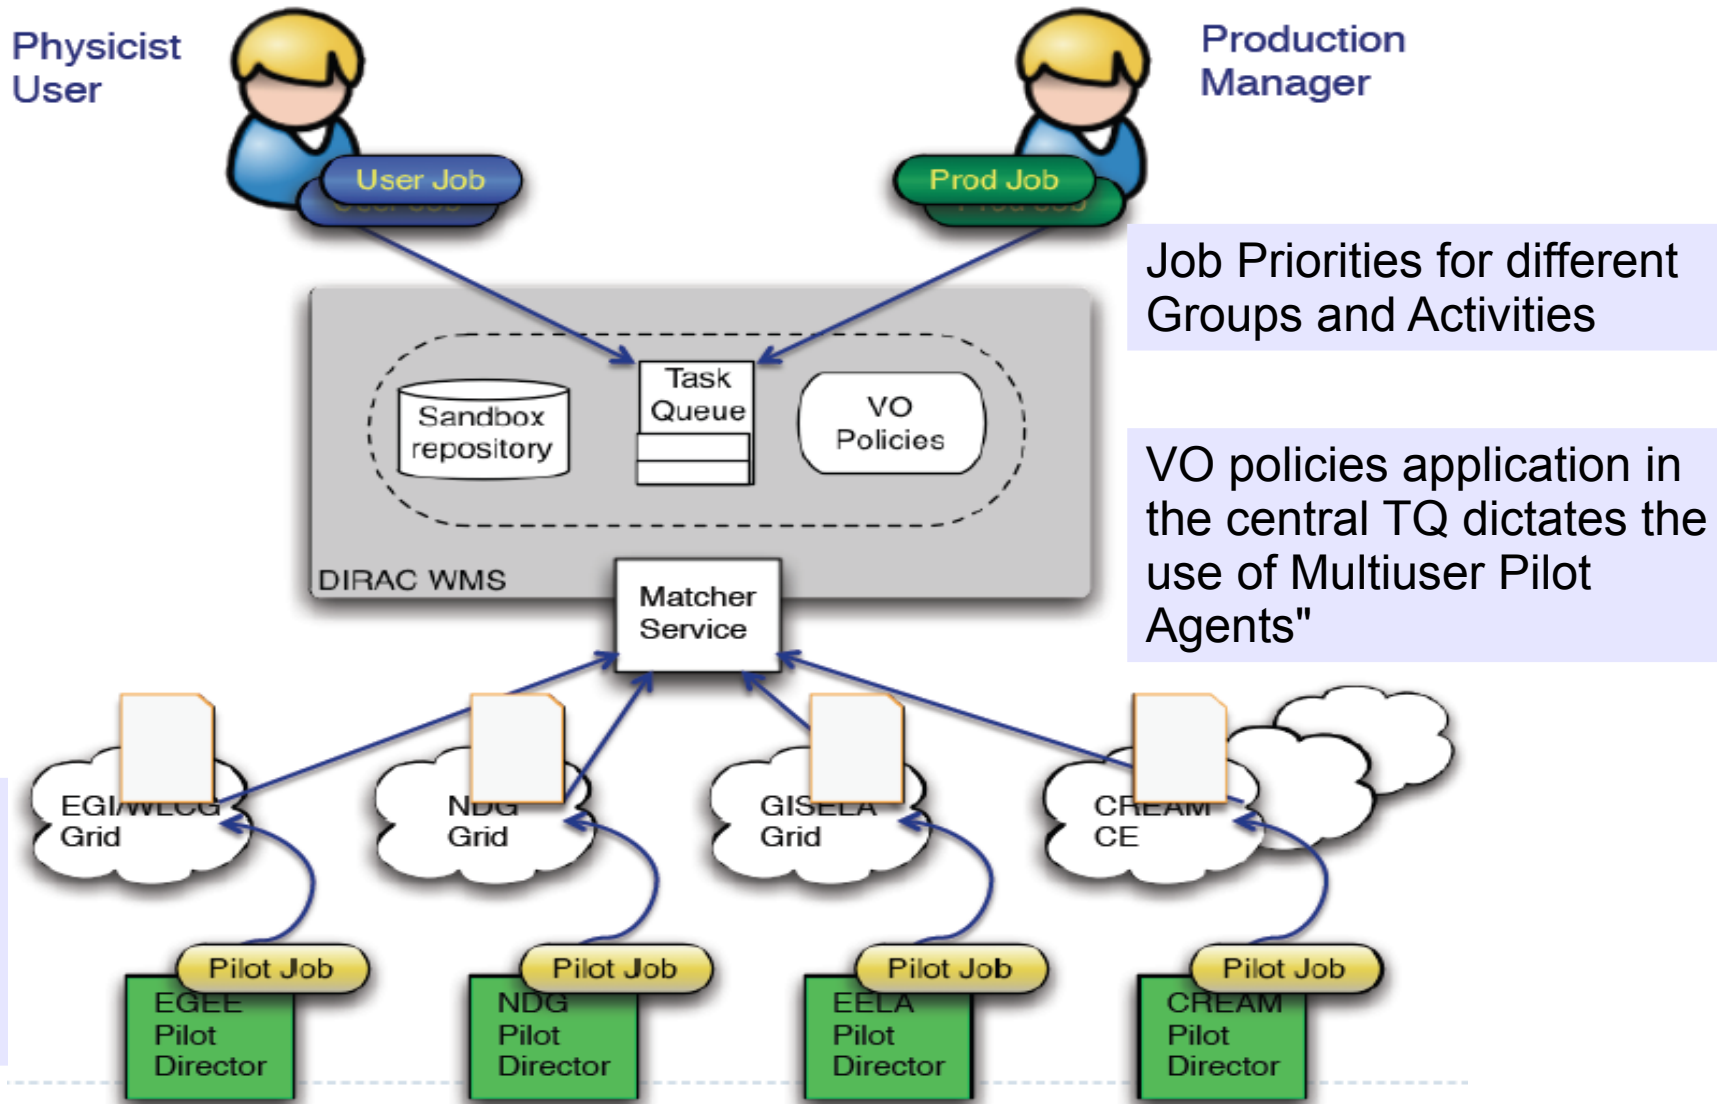
DIRAC Workload Management



DIRAC Workload Management



DIRAC Workload Management

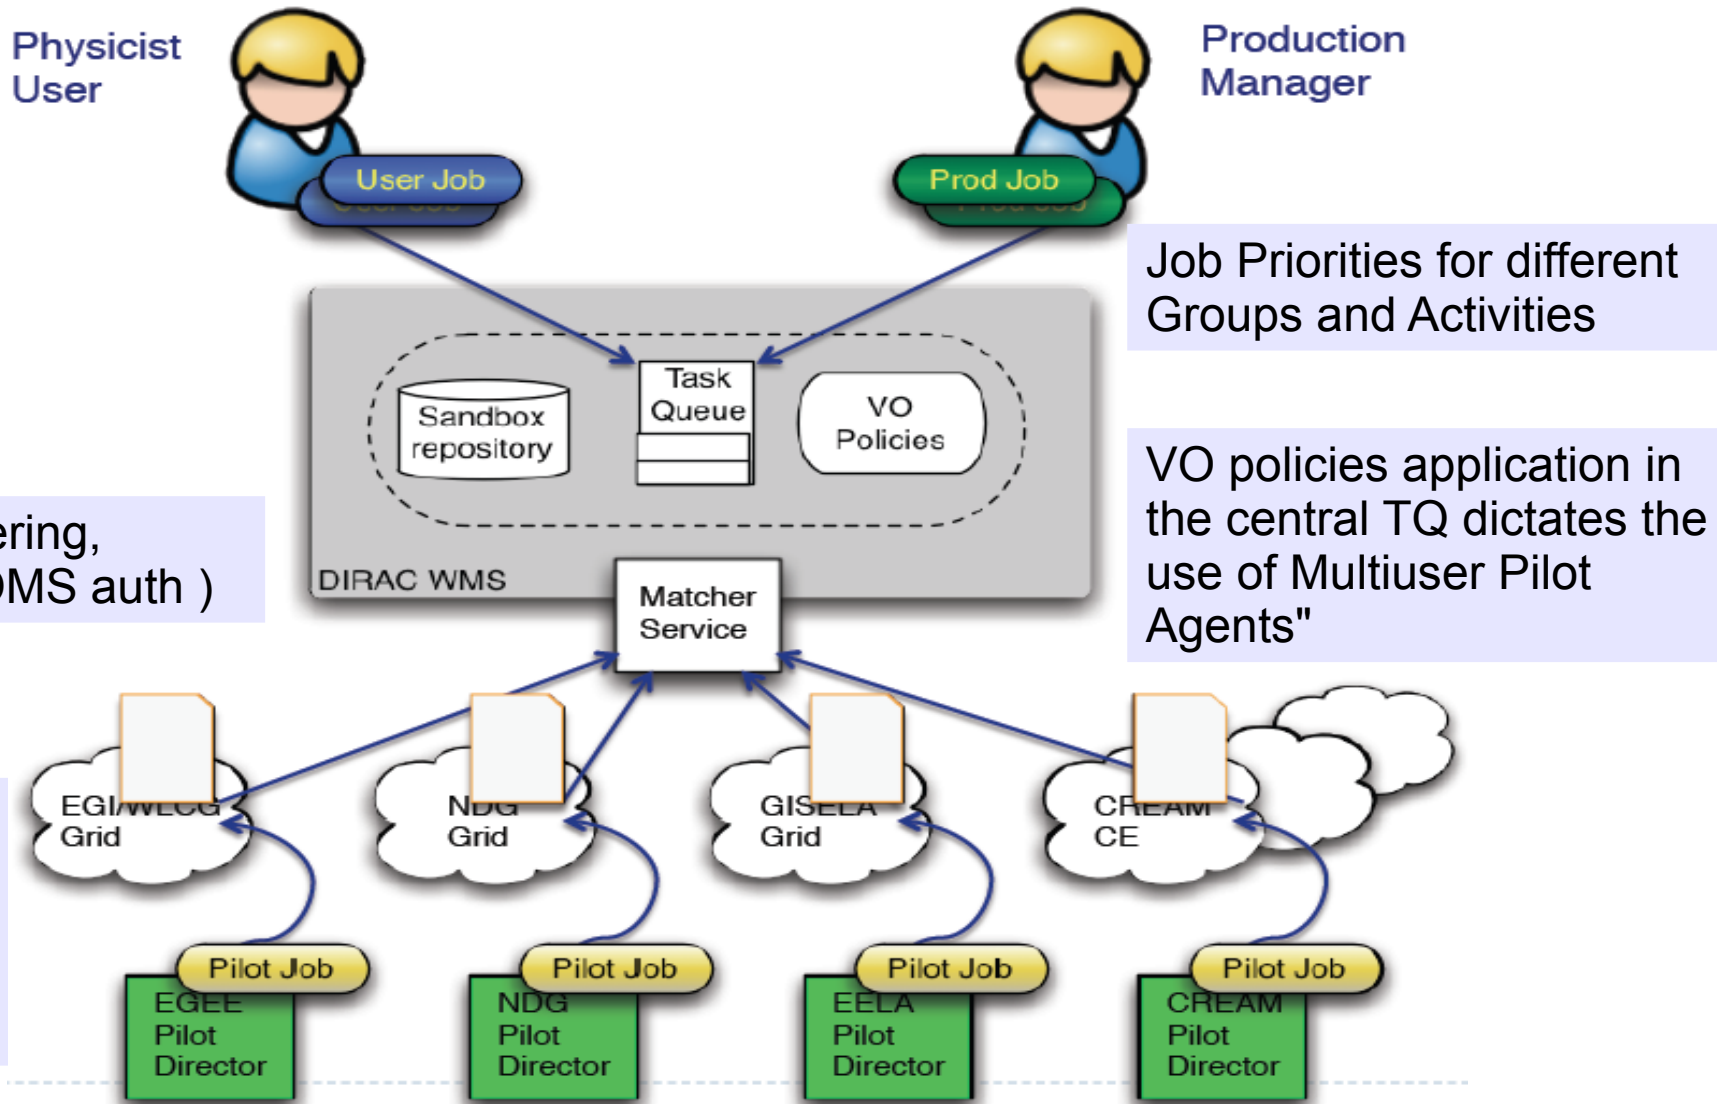


Job Priorities for different Groups and Activities

VO policies application in the central TQ dictates the use of Multiuser Pilot Agents"

Pilot pre-allocates the computing resources (Pilot auth)

DIRAC Workload Management



Job Brokering,
(User VOMS auth)

Job Priorities for different
Groups and Activities

VO policies application in
the central TQ dictates the
use of Multiuser Pilot
Agents"

Pilot pre-allocates the
computing resources
(Pilot auth)

DIRAC Workload Management

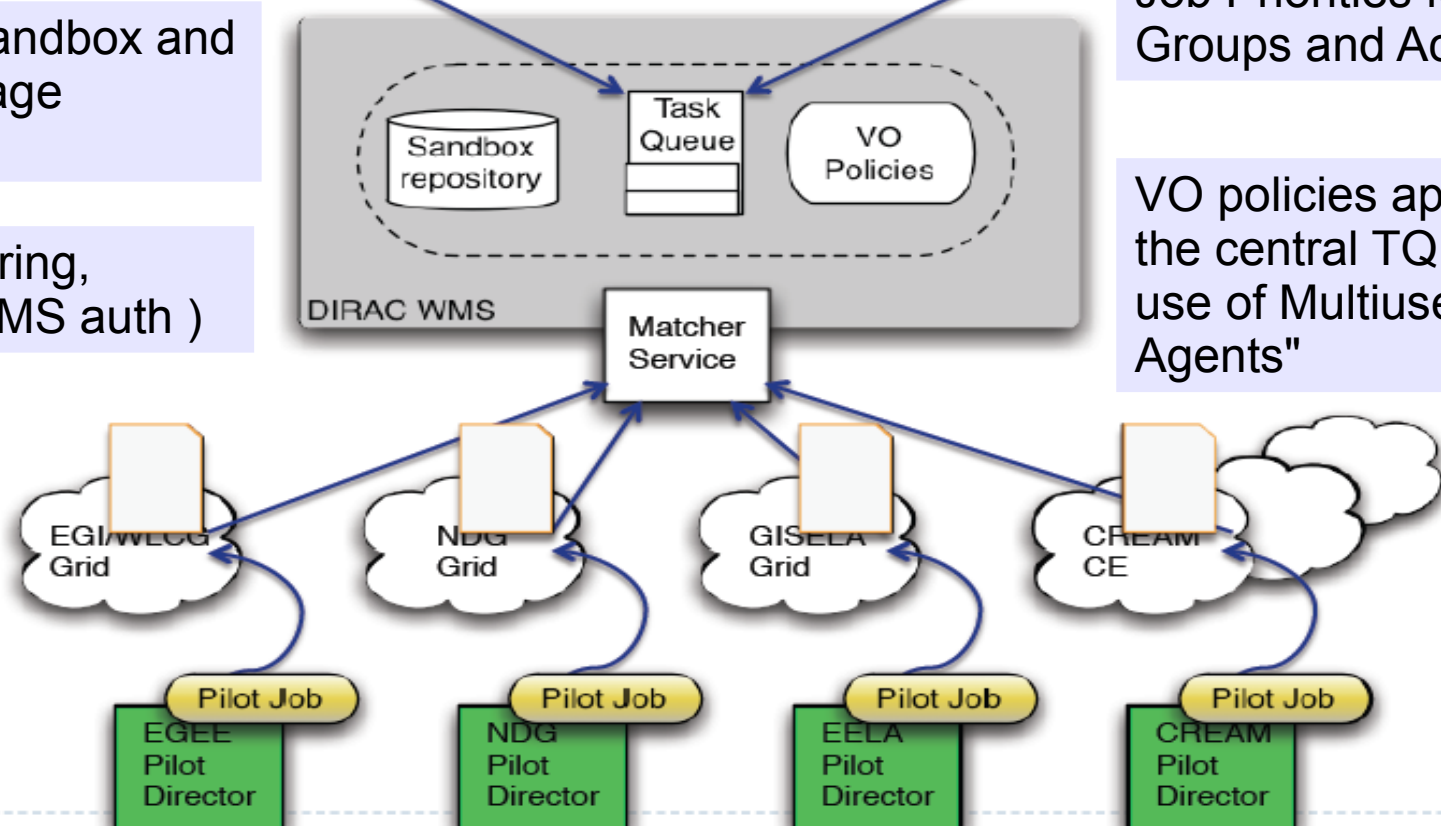


Input/Output Sandbox and third party storage management

Job Brokering, (User VOMS auth)

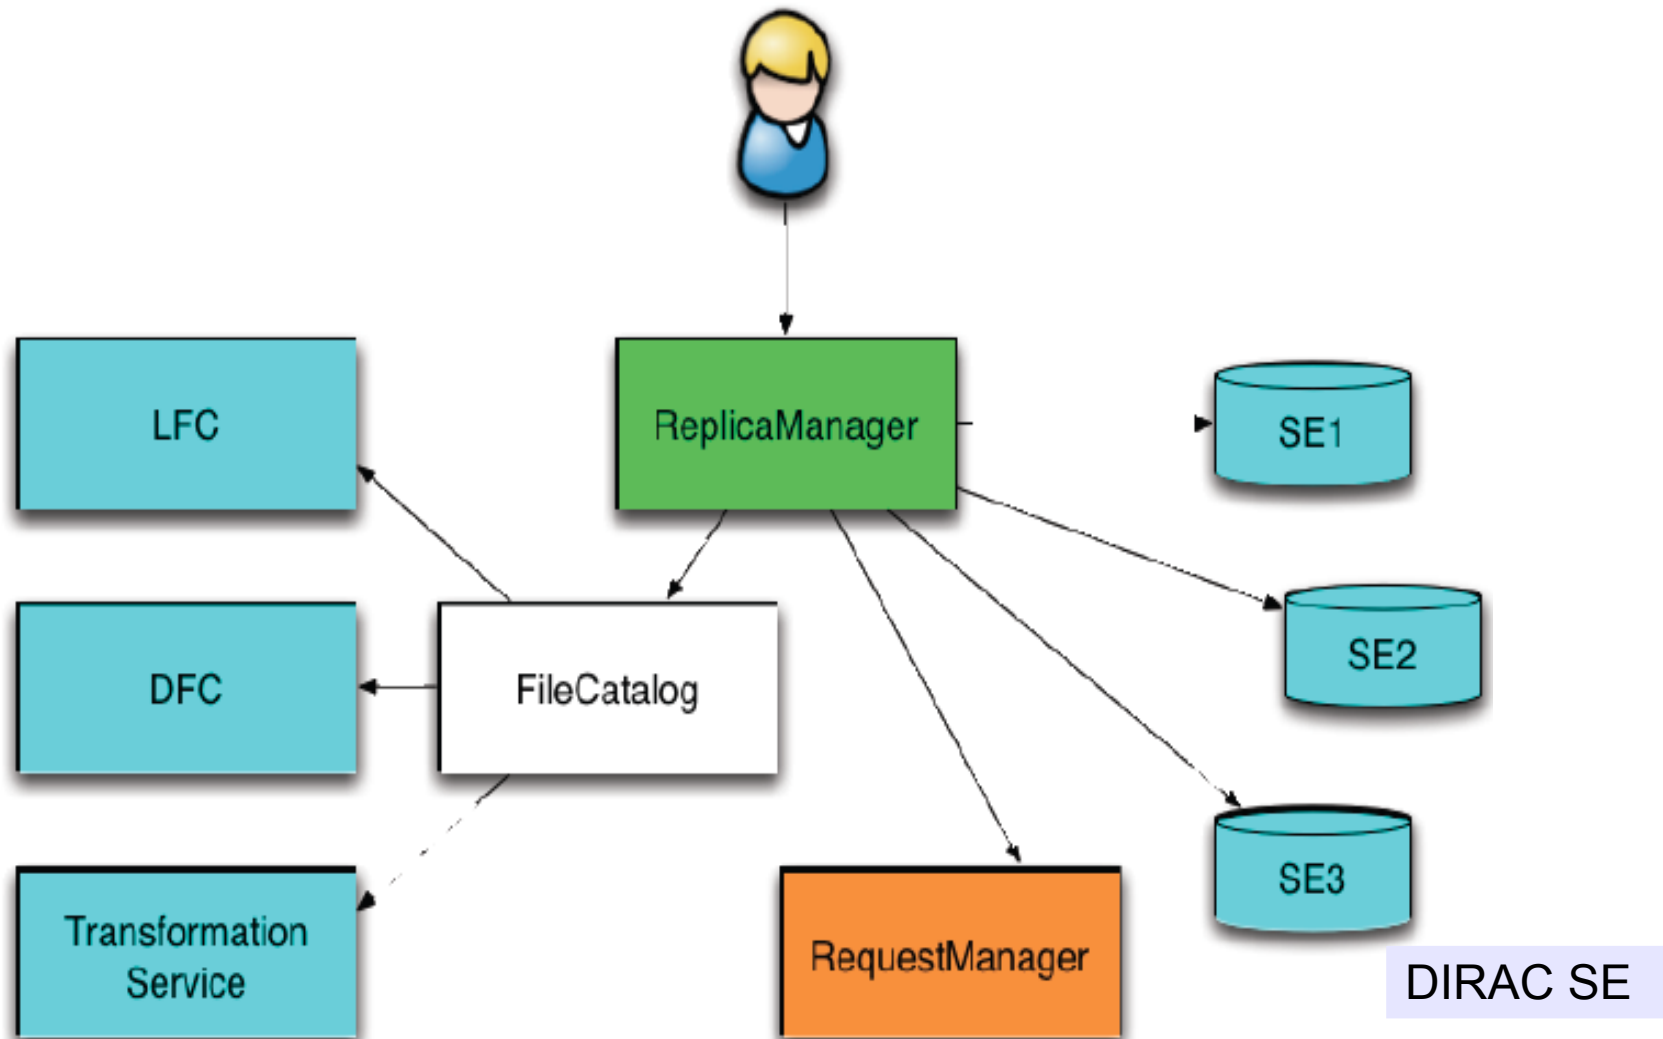
Job Priorities for different Groups and Activities

VO policies application in the central TQ dictates the use of Multiuser Pilot Agents"

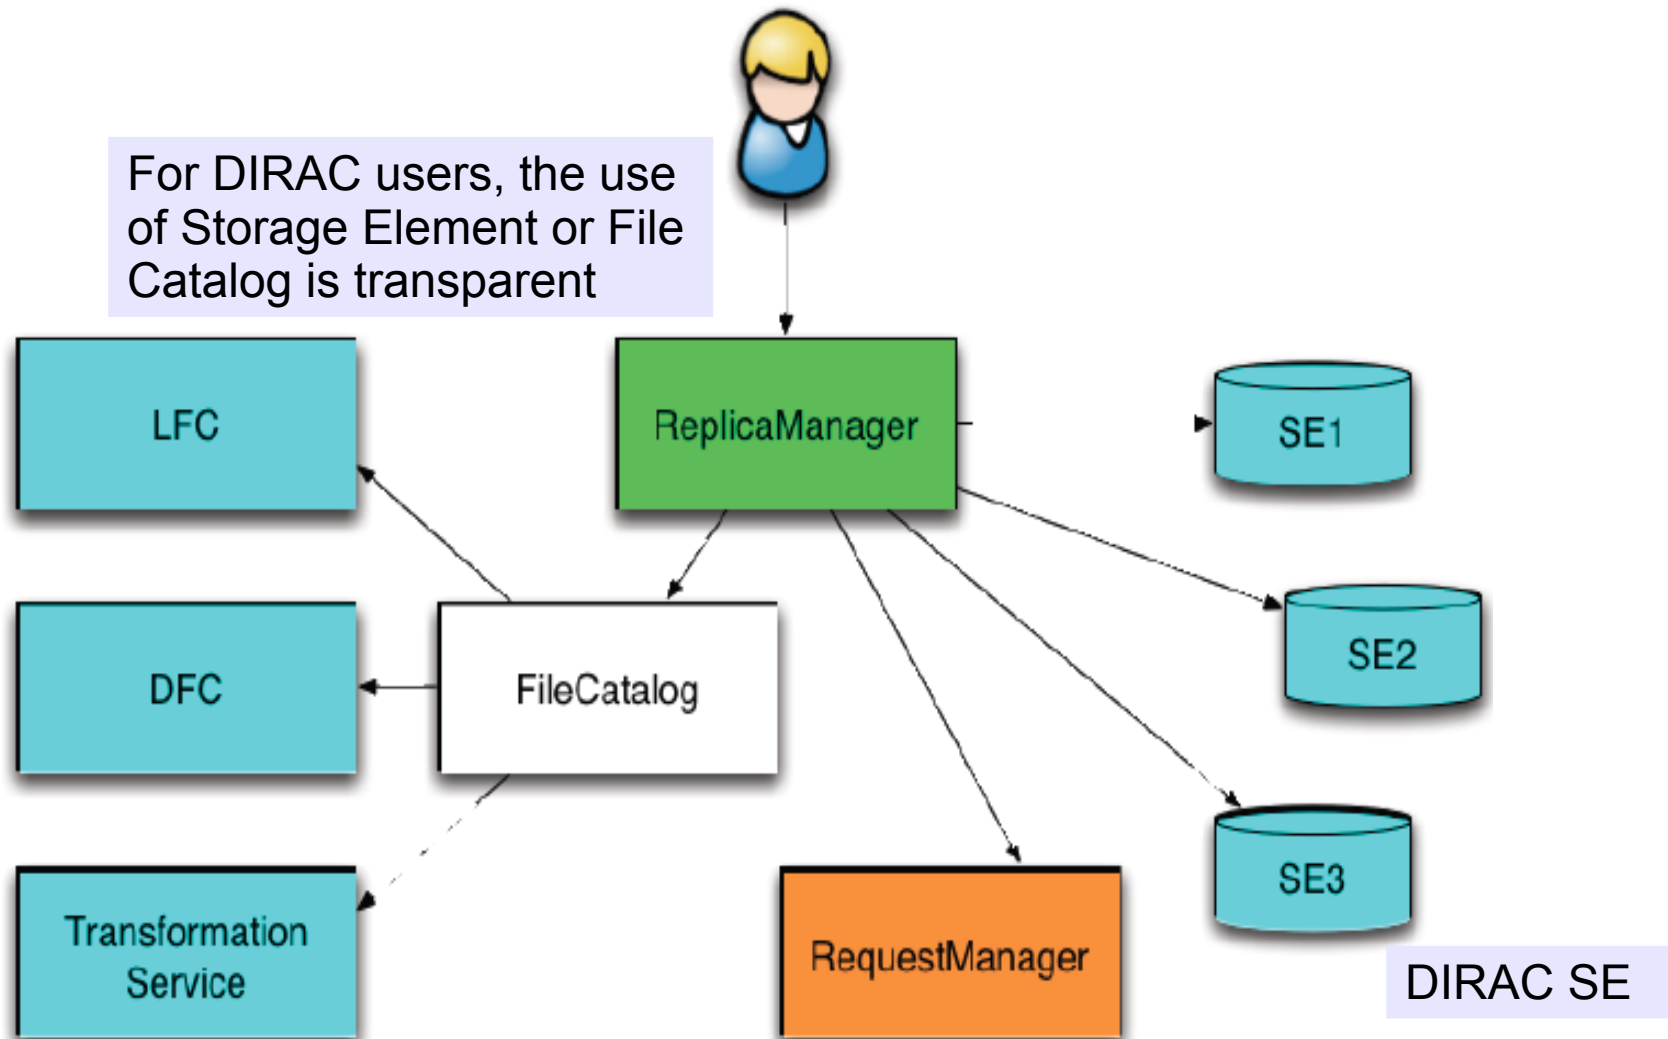


Pilot pre-allocates the computing resources (Pilot auth)

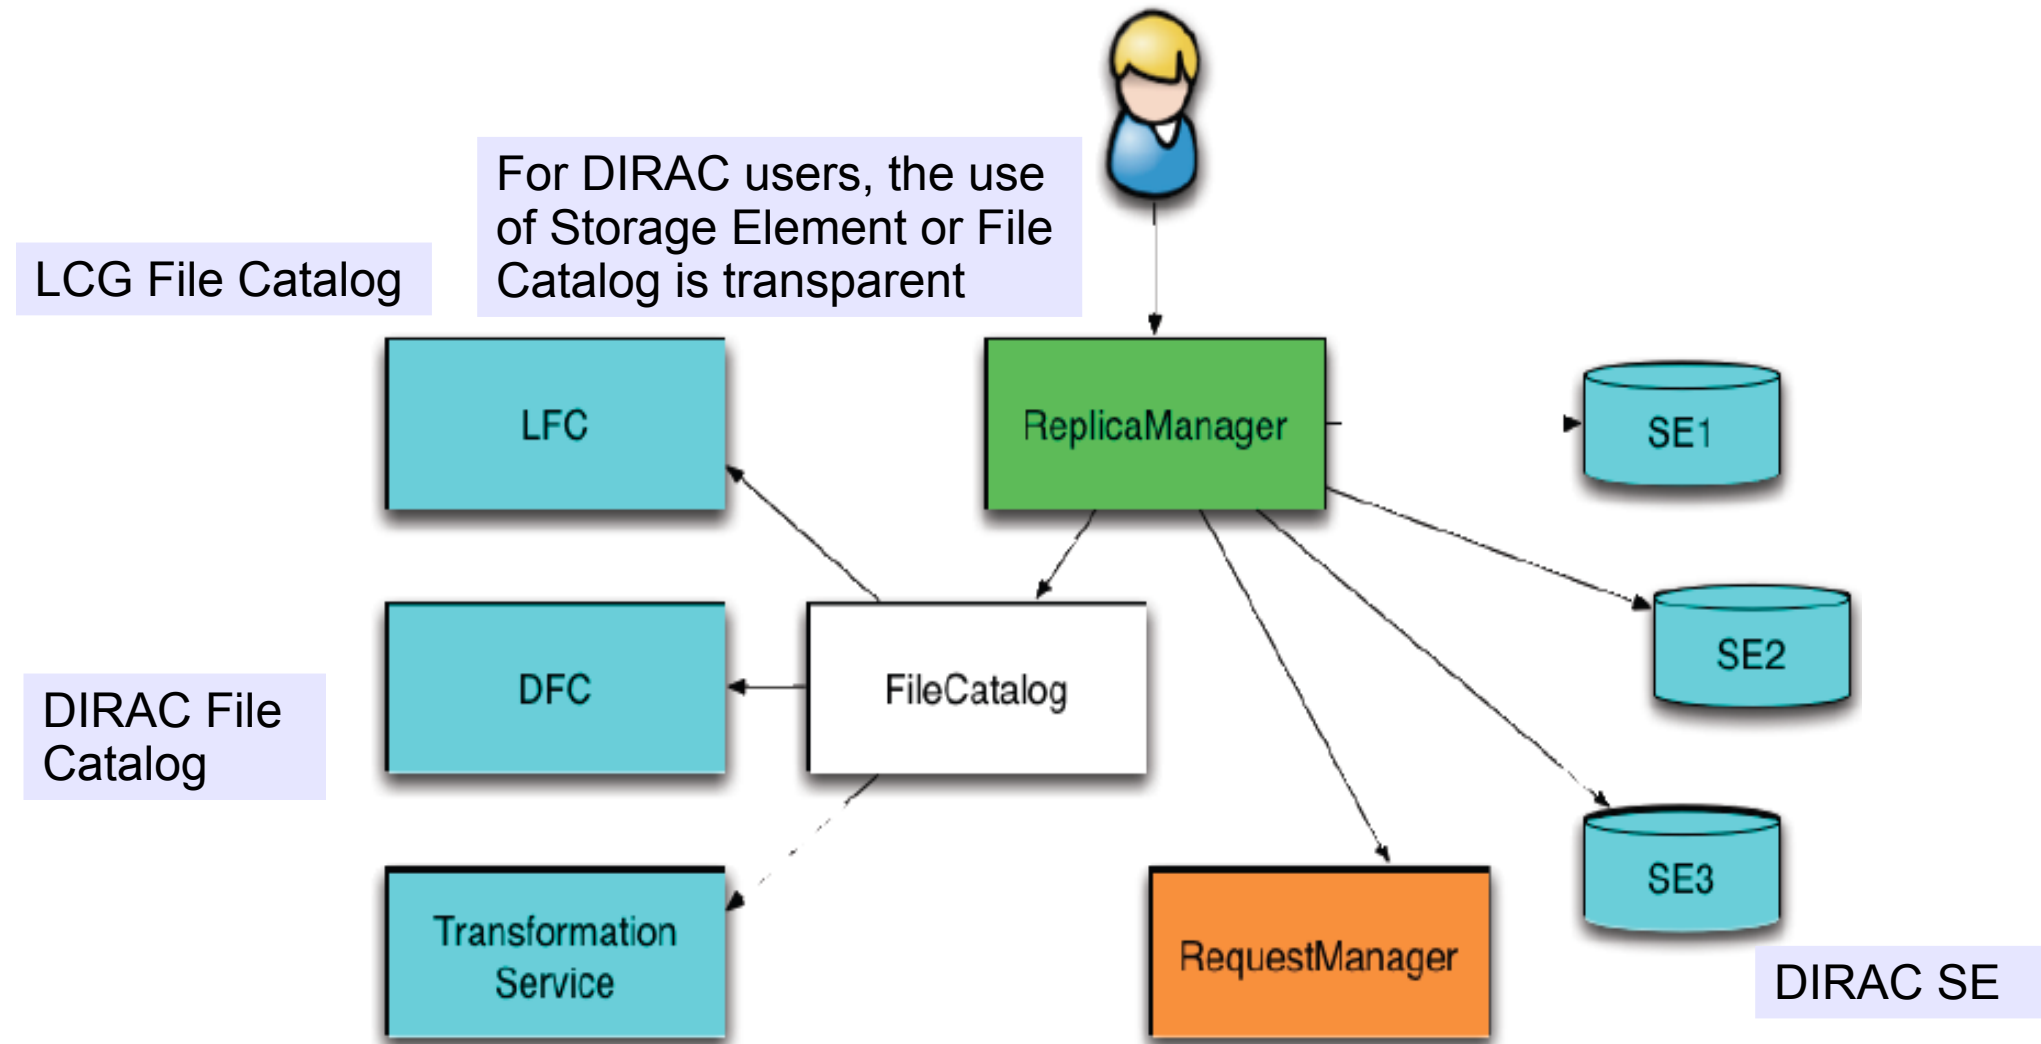
DIRAC Data Management



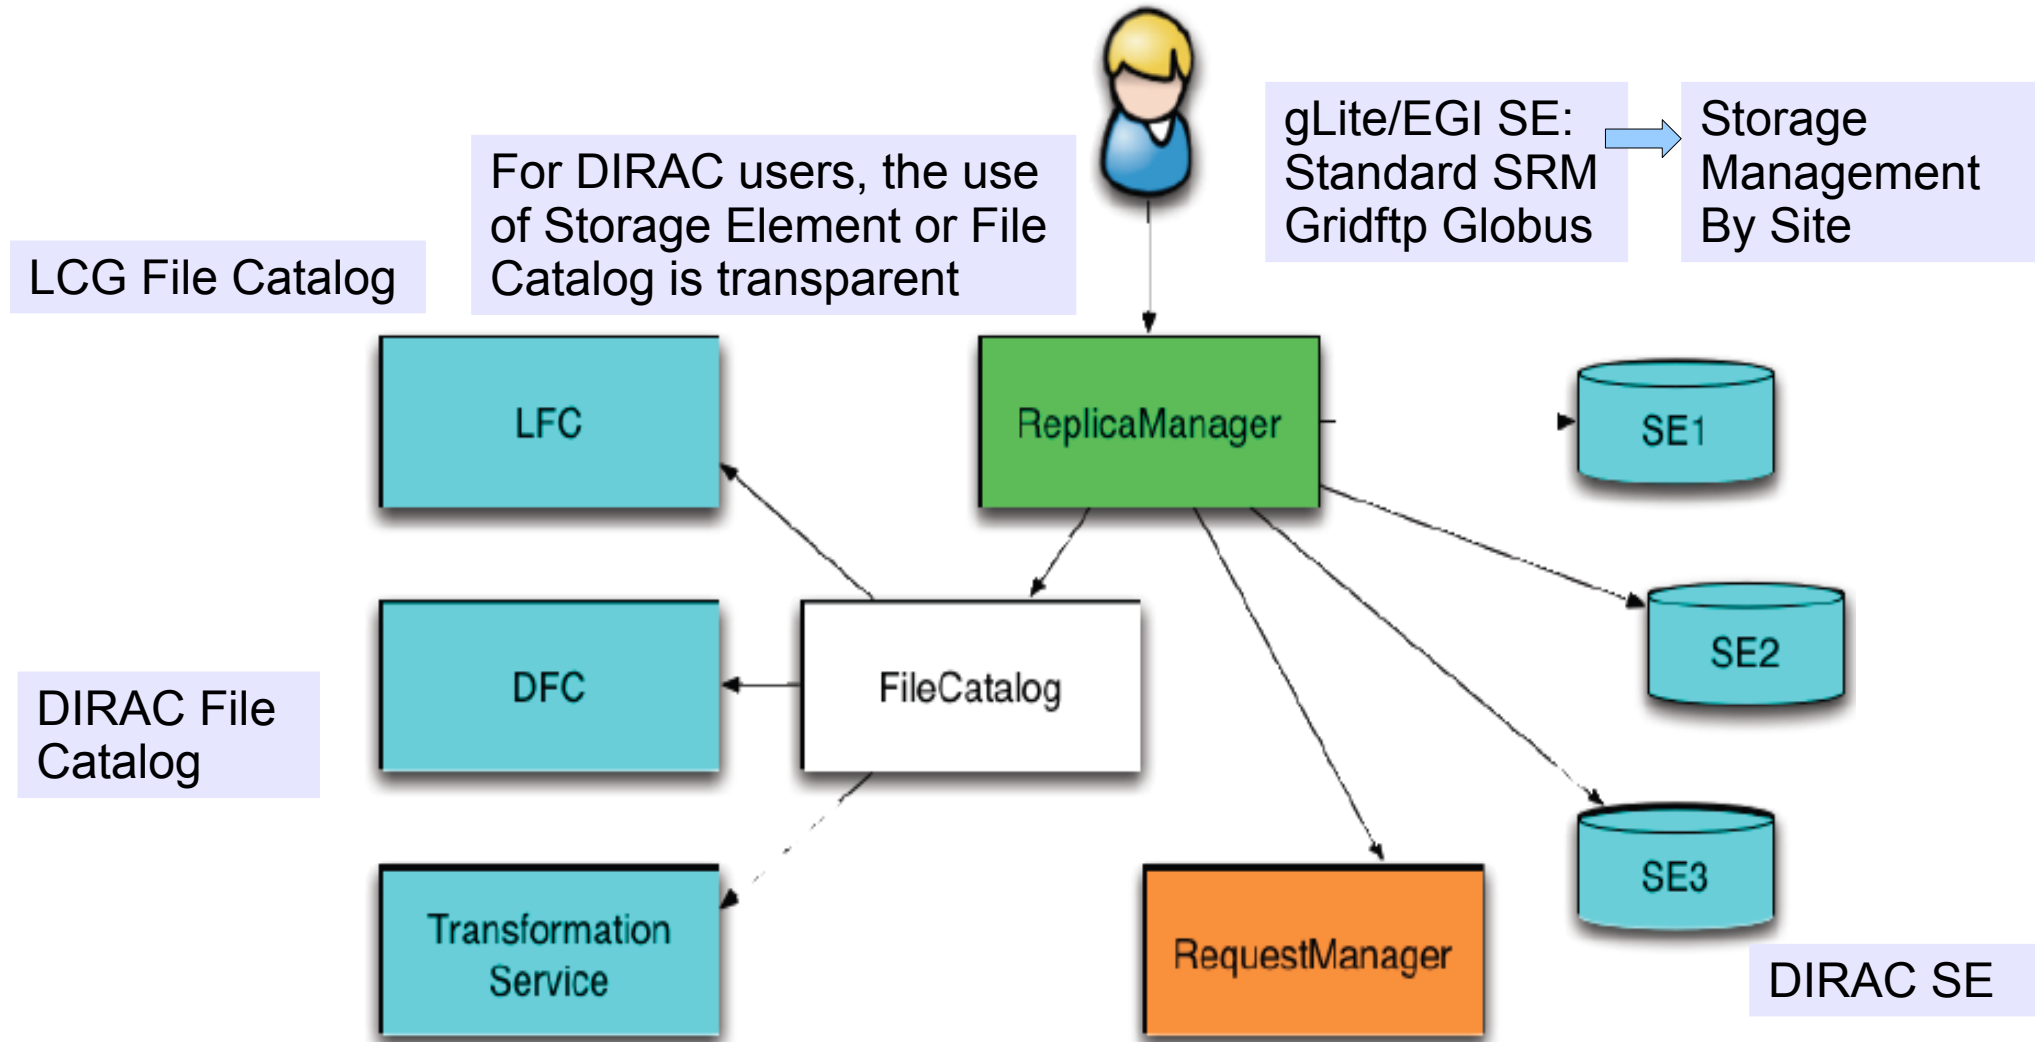
DIRAC Data Management



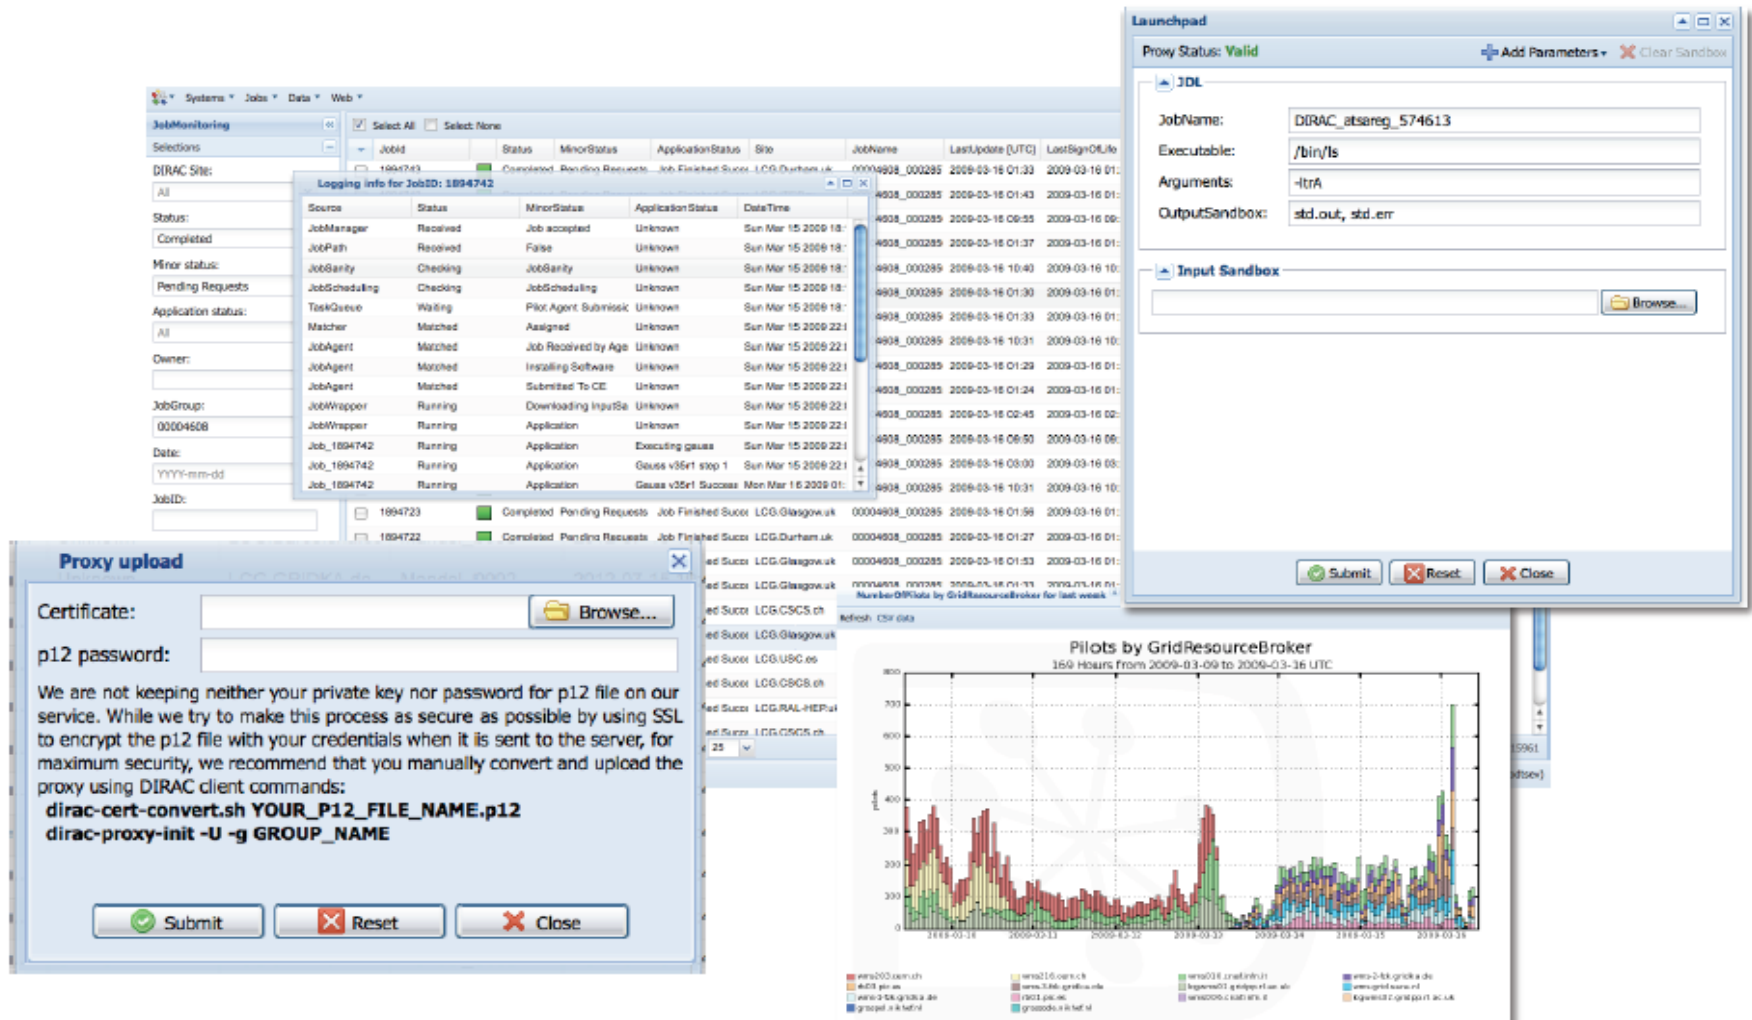
DIRAC Data Management



DIRAC Data Management



DIRAC Web Portal

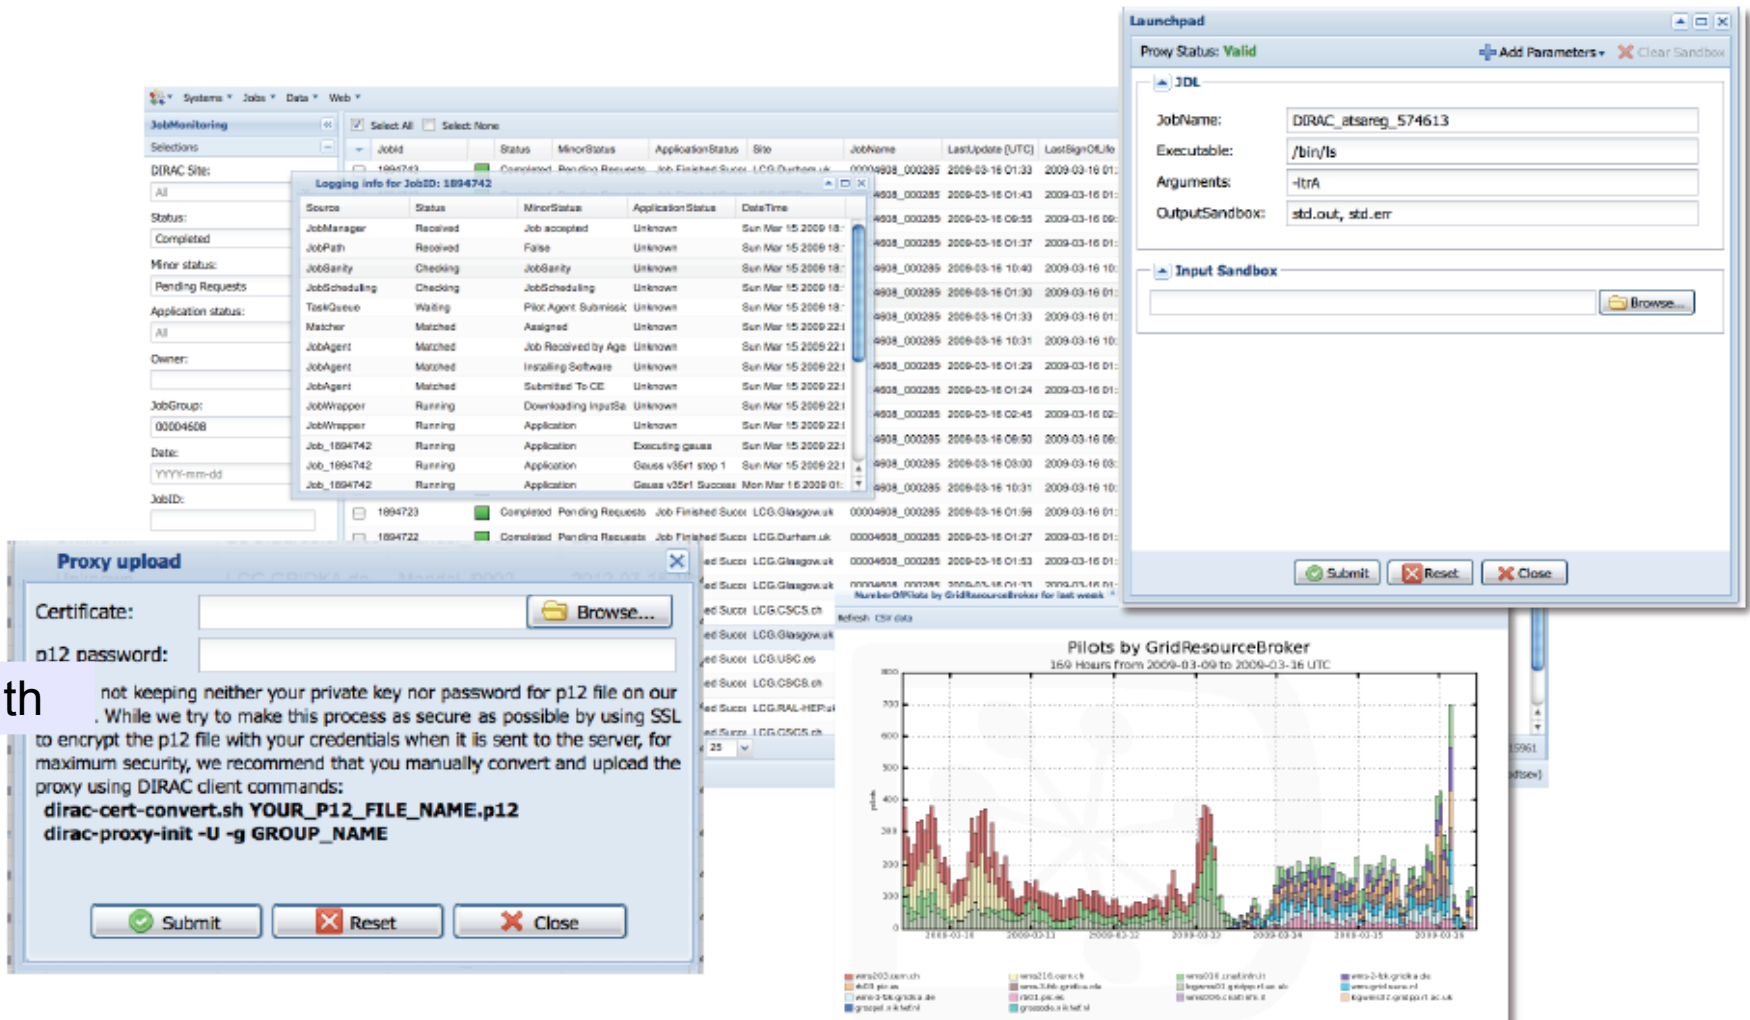


The screenshot displays the DIRAC Web Portal interface with several key components:

- Job Monitoring:** A central table showing job details for JobID: 1894742. The table includes columns for Source, Status, MinorStatus, ApplicationStatus, and DateTime. The job is currently in a 'Running' state.
- Proxy Upload:** A dialog box for uploading a proxy certificate. It includes a 'Certificate' field with a 'Browse...' button and a 'p12 password' field. Below the fields, there is a security notice and two command-line instructions:


```
dirac-cert-convert.sh YOUR_P12_FILE_NAME,p12
dirac-proxy-init -U -g GROUP_NAME
```
- Launchpad:** A configuration window for a job. It shows 'Proxy Status: Valid' and fields for 'JobName' (DIRAC_atsreg_574613), 'Executable' (/bin/ls), 'Arguments' (-ltrA), and 'OutputSandbox' (std.out, std.err). There is a 'Browse...' button for the sandbox and 'Submit', 'Reset', and 'Close' buttons at the bottom.
- Pilots by GridResourceBroker:** A bar chart showing pilot activity from 2009-03-09 to 2009-03-16 UTC. The chart displays the number of pilots over time, with a legend identifying various GridResourceBroker (GRB) instances such as wms202.cern.ch, wms216.cern.ch, and others.

DIRAC Web Portal



The screenshot displays the DIRAC Web Portal interface with several key components:

- Job Monitoring:** A table showing job details for JobID: 1894742. The table includes columns for Source, Status, MinorStatus, ApplicationStatus, and DateTime. The job is currently in a 'Running' state.
- Proxy upload:** A dialog box for uploading a proxy certificate. It includes a 'Certificate' field with a 'Browse...' button and a 'p12 password:' field. Below the fields, there is a text box with instructions and commands:

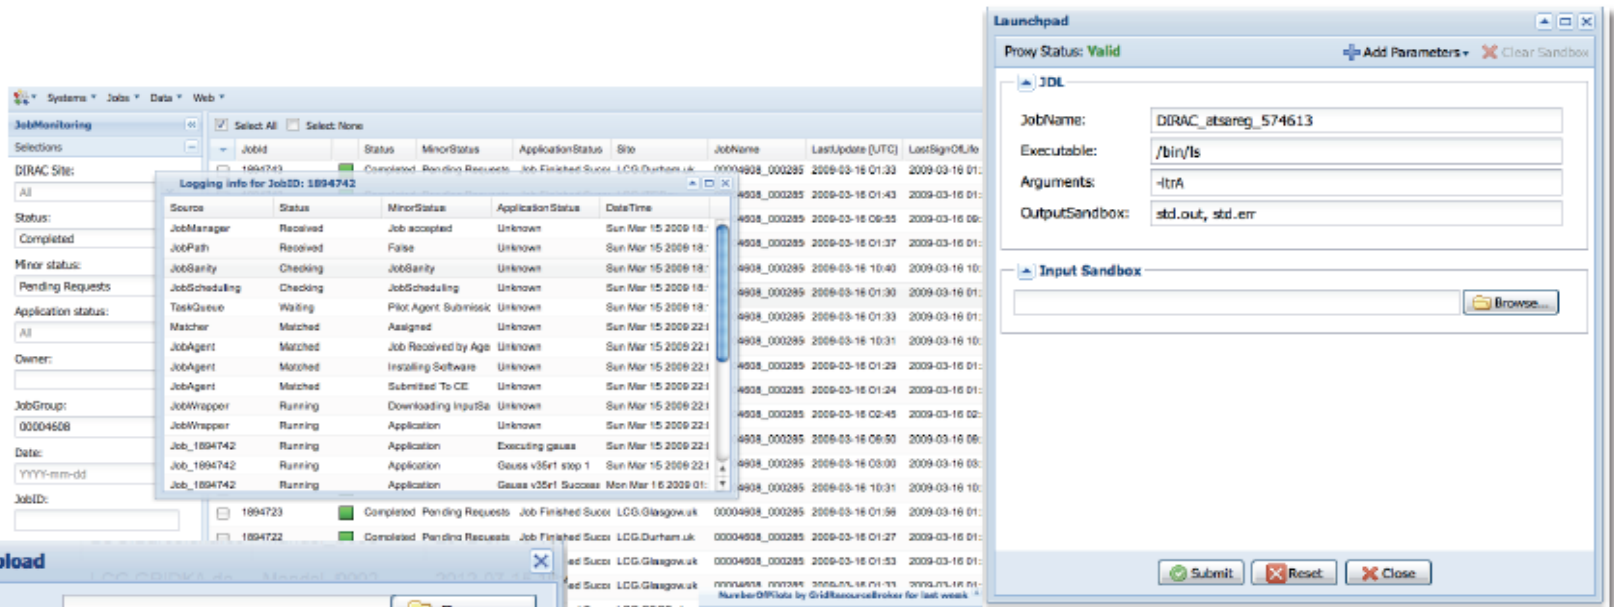

```

      not keeping neither your private key nor password for p12 file on our
      . While we try to make this process as secure as possible by using SSL
      to encrypt the p12 file with your credentials when it is sent to the server, for
      maximum security, we recommend that you manually convert and upload the
      proxy using DIRAC client commands:
      dirac-cert-convert.sh YOUR_P12_FILE_NAME,p12
      dirac-proxy-init -U -g GROUP_NAME
      
```
- Launchpad:** A window for configuring job execution. It shows 'Proxy Status: Valid' and fields for 'JobName' (DIRAC_atsreg_574613), 'Executable' (/bin/ls), 'Arguments' (-ltrA), and 'OutputSandbox' (std.out, std.err). There is an 'Input Sandbox' field with a 'Browse...' button and 'Submit', 'Reset', and 'Close' buttons at the bottom.
- Pilots by GridResourceBroker:** A line graph showing pilot activity over time from 2009-03-09 to 2009-03-16 UTC. The Y-axis represents the number of pilots (0 to 800). The graph shows multiple colored lines representing different GridResourceBroker (GRB) instances, with a significant peak in activity around 2009-03-15.

Proxy Auth

DIRAC Web Portal

Job Launch Pad



The screenshot shows the DIRAC web portal interface. On the left, there's a 'JobMonitoring' sidebar with various filters and job details. The main area displays a table of jobs with columns for JobName, LastUpdate (UTC), and LastSignOff. A 'Logging info for JobID: 1894742' window is open, showing a detailed log of job events. On the right, the 'Launchpad' window is visible, showing fields for JobName (DIRAC_atsreg_574613), Executable (/bin/ls), Arguments (-ltrA), and OutputSandbox (std.out, std.err). There are 'Submit', 'Reset', and 'Close' buttons at the bottom of the Launchpad.




Proxy upload

Certificate:

p12 password:

Proxy Auth

not keeping neither your private key nor password for p12 file on our . While we try to make this process as secure as possible by using SSL to encrypt the p12 file with your credentials when it is sent to the server, for maximum security, we recommend that you manually convert and upload the proxy using DIRAC client commands:
dirac-cert-convert.sh YOUR_P12_FILE_NAME,p12
dirac-proxy-init -U -g GROUP_NAME



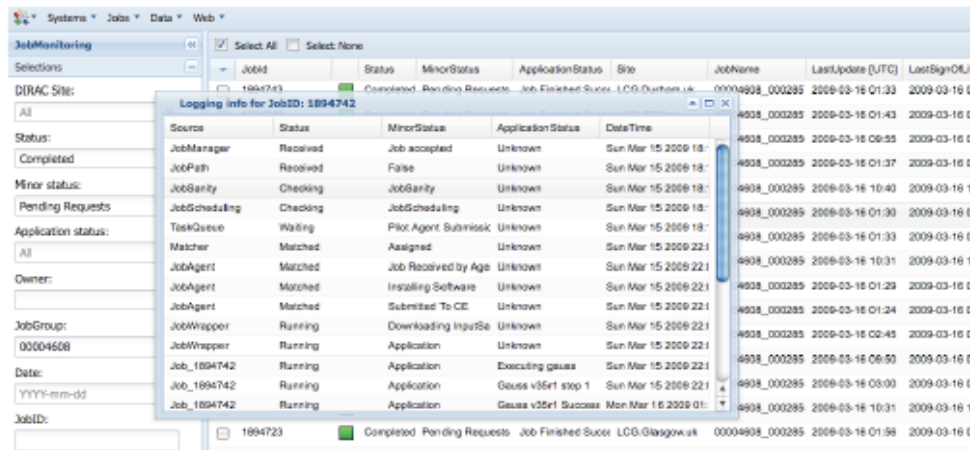
DIRAC Web Portal

Job Launch Pad

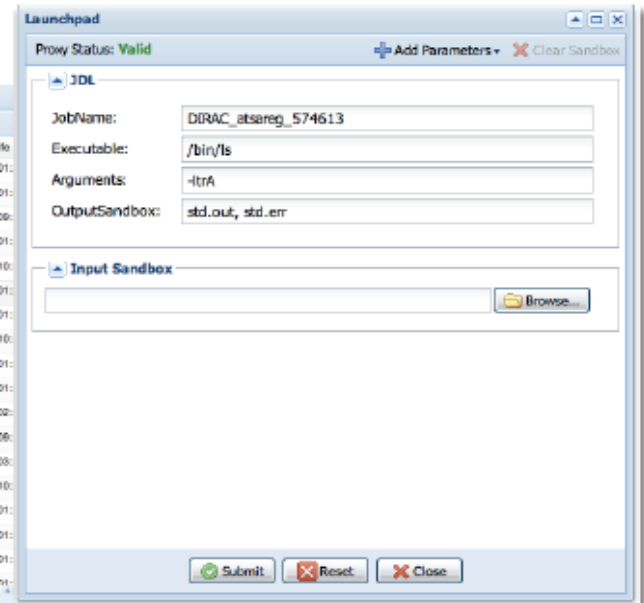
Pilot Management

Job Management

Data Management



Source	Status	MinorStatus	ApplicationStatus	DateTime
JobManager	Received	Job accepted	Unknown	Sun Mar 15 2009 18:...
JobPath	Received	False	Unknown	Sun Mar 15 2009 18:...
JobGarity	Checking	JobGarity	Unknown	Sun Mar 15 2009 18:...
JobScheduling	Checking	JobScheduling	Unknown	Sun Mar 15 2009 18:...
TaskQueue	Waiting	Pilot.Agent.Submitsk	Unknown	Sun Mar 15 2009 18:...
Matcher	Matched	Assigned	Unknown	Sun Mar 15 2009 22:1...
JobAgent	Matched	Job Received by Age	Unknown	Sun Mar 15 2009 22:1...
JobAgent	Matched	Installing Software	Unknown	Sun Mar 15 2009 22:1...
JobAgent	Matched	Submitted To CE	Unknown	Sun Mar 15 2009 22:1...
JobWrapper	Running	Downloading InputSe	Unknown	Sun Mar 15 2009 22:1...
JobWrapper	Running	Application	Unknown	Sun Mar 15 2009 22:1...
Job	Running	Application	Executing gauss	Sun Mar 15 2009 22:1...
Job	Running	Application	Gauss v35r1 stop 1	Sun Mar 15 2009 22:1...
Job	Running	Application	Gauss v35r1 Success	Mon Mar 16 2009 01:0...



Proxy Status: Valid

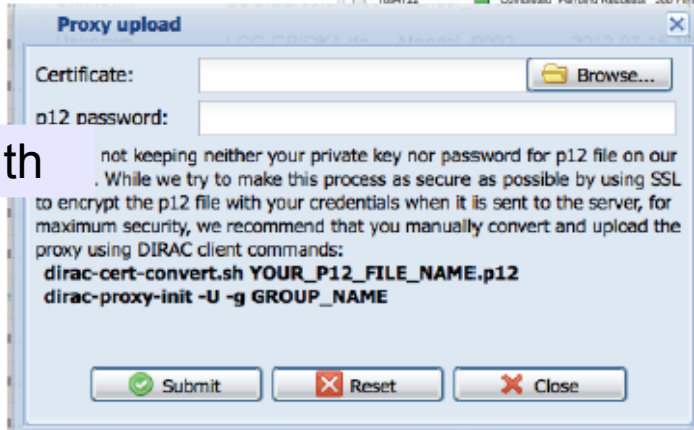
JobName: DIRAC_atsreg_574613

Executable: /bin/lis

Arguments: -lra

OutputSandbox: std.out, std.err

Submit Reset Close



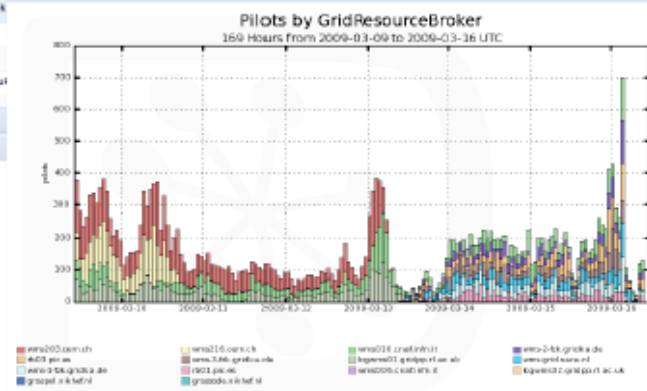
Certificate: [Browse...]

p12 password: []

not keeping neither your private key nor password for p12 file on our . While we try to make this process as secure as possible by using SSL to encrypt the p12 file with your credentials when it is sent to the server, for maximum security, we recommend that you manually convert and upload the proxy using DIRAC client commands:
dirac-cert-convert.sh YOUR_P12_FILE_NAME,p12
dirac-proxy-init -U -g GROUP_NAME

Submit Reset Close

Proxy Auth

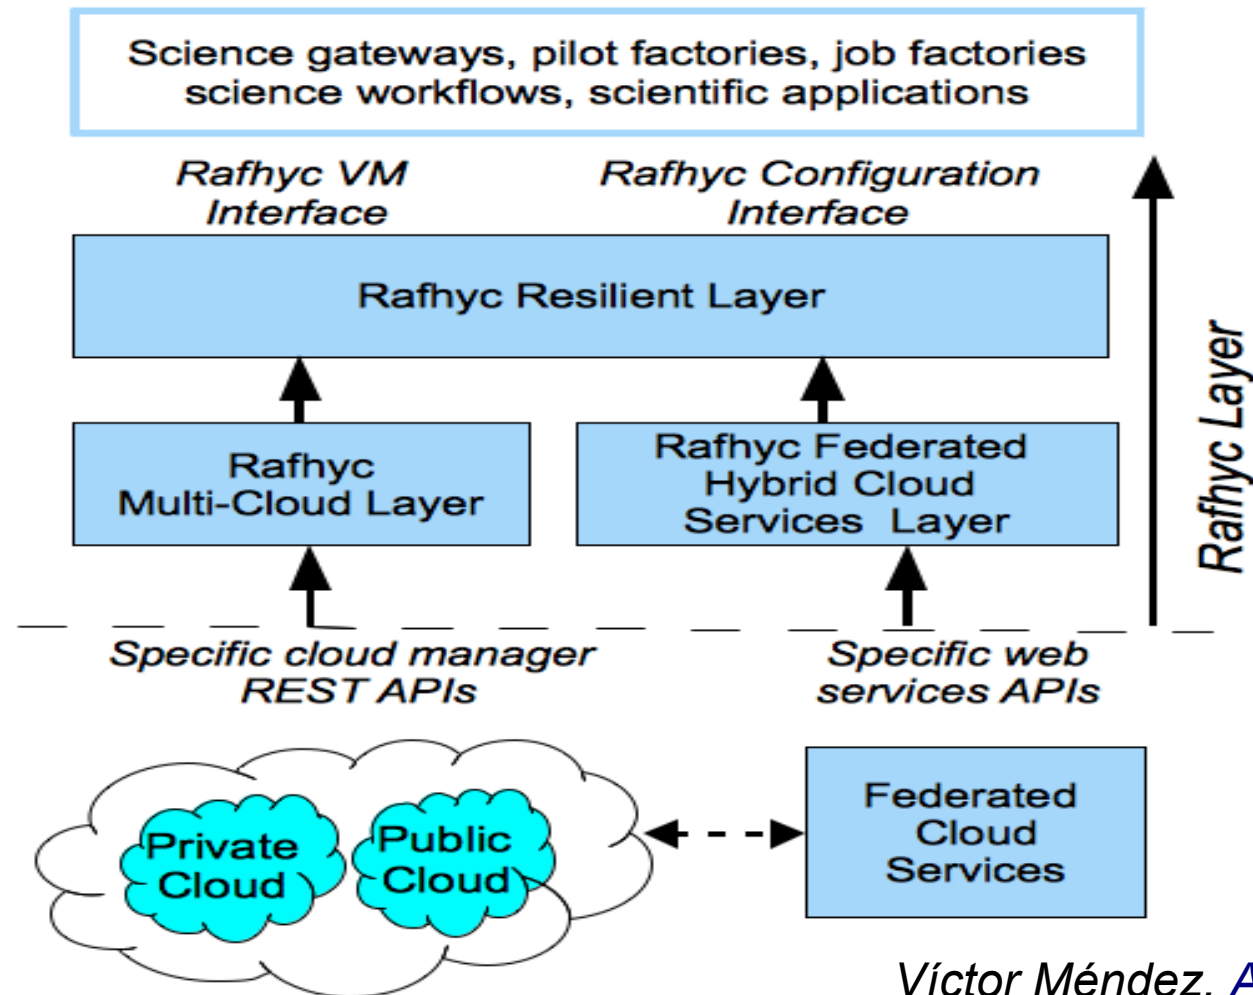


Accounting Statistics

Federated hybrid cloud computing model for scientific communities

- Federated Cloud requires SaaS:
 - Integrating different IaaS providers
- Hybrid Cloud requires:
 - Community Clouds IaaS providers
 - Commercial Clouds IaaS providers
- Scientific communities requires
 - A minimal set of functionalities for scientific computing communities:
 - Information, Monitoring, Auth, Image Metadata, Accounting

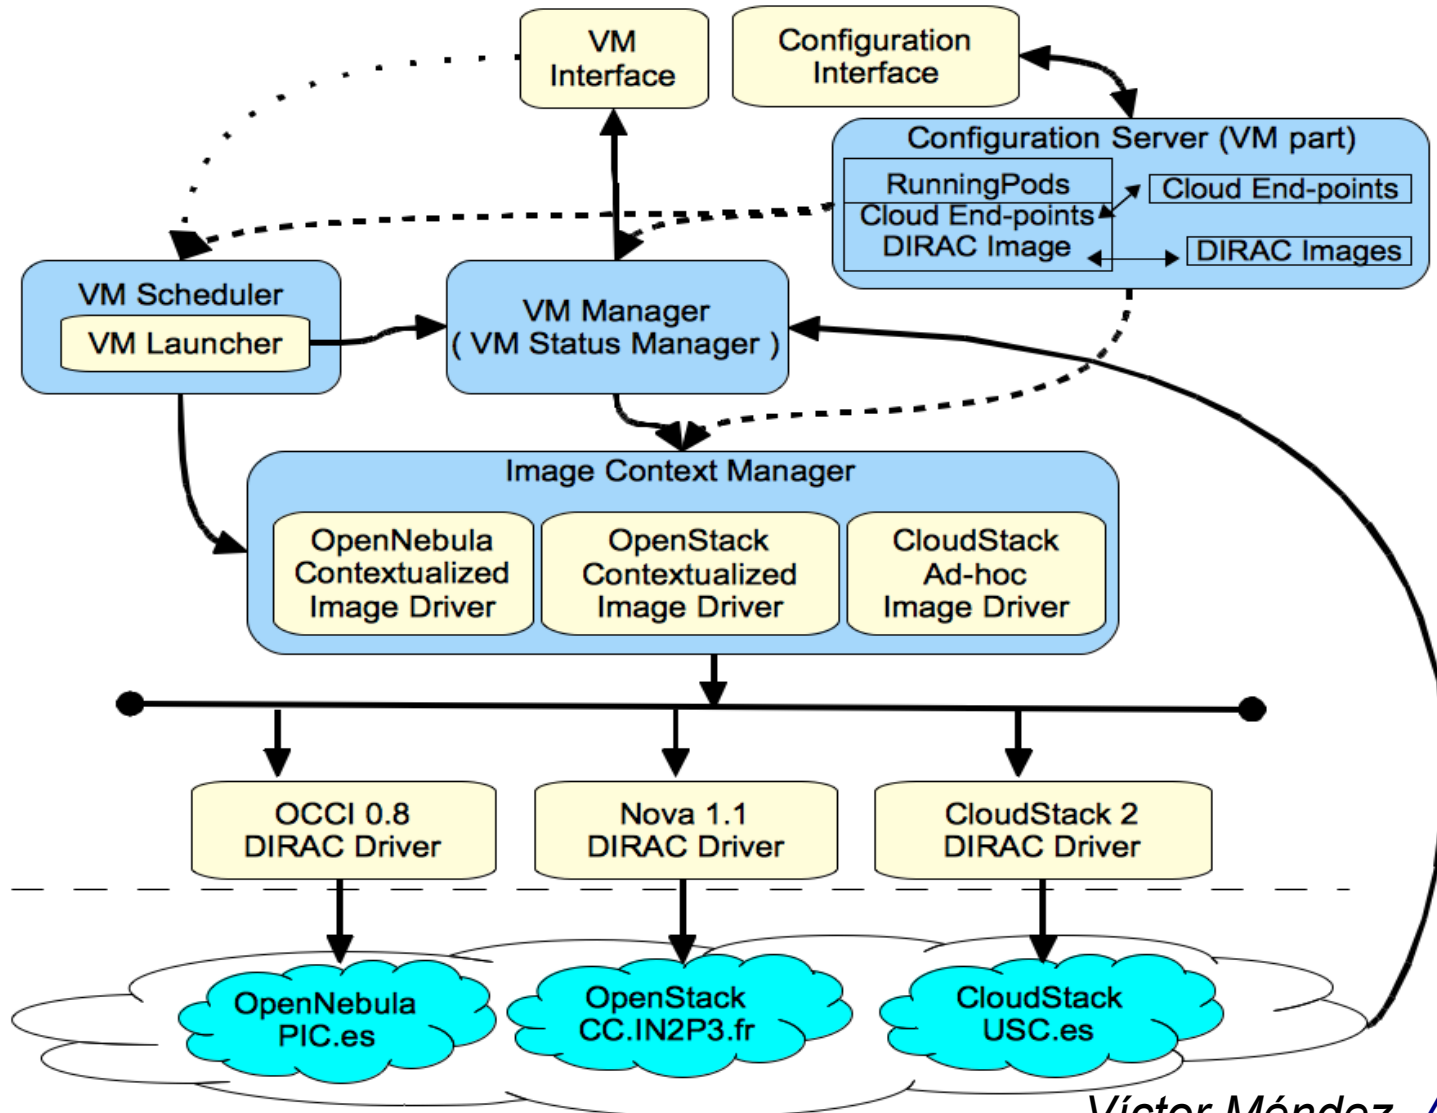
Rafhyc Overview: Resilient Architecture of Federated HYbrid Clouds




- Rafhyc adoption in DIRAC I
 - Resilient High-level Layer to provide service delivery and management (VMs):
 - Persistent Configuration: An extension of DIRAC Configuration System
 - VM Manager: Taking statistics of VM status
 - VM Scheduler: Integrated with DIRAC TQ job brokering, or with third party job brokers
 - VM Interface: An extension of the DIRAC Web portal
 - Multi-Cloud Low-level layer:
 - OpenNebula OCCl 0.8 and rOCCl 1.1, CloudStack 2, OpenStack nova 1.1 driver, EC2 Amazon
 - Image Context Manager

- Rafhyc adoption in DIRAC II
 - Federated Hybrid Cloud Services Low-level layer:
 - V0r9: Federated static info managed by admins at DIRAC Configuration System
 - Image Metadata, IaaS Information
 - Third party VOMS authentication is integrated in DIRAC (OpenNebula, OpenStack)
 - DIRAC VM monitoring is provided for VOs
 - Third party ganglia monitoring system has been tested for IaaS monitoring of the VMs
 - Next: Accounting

VMDIRAC Architecture

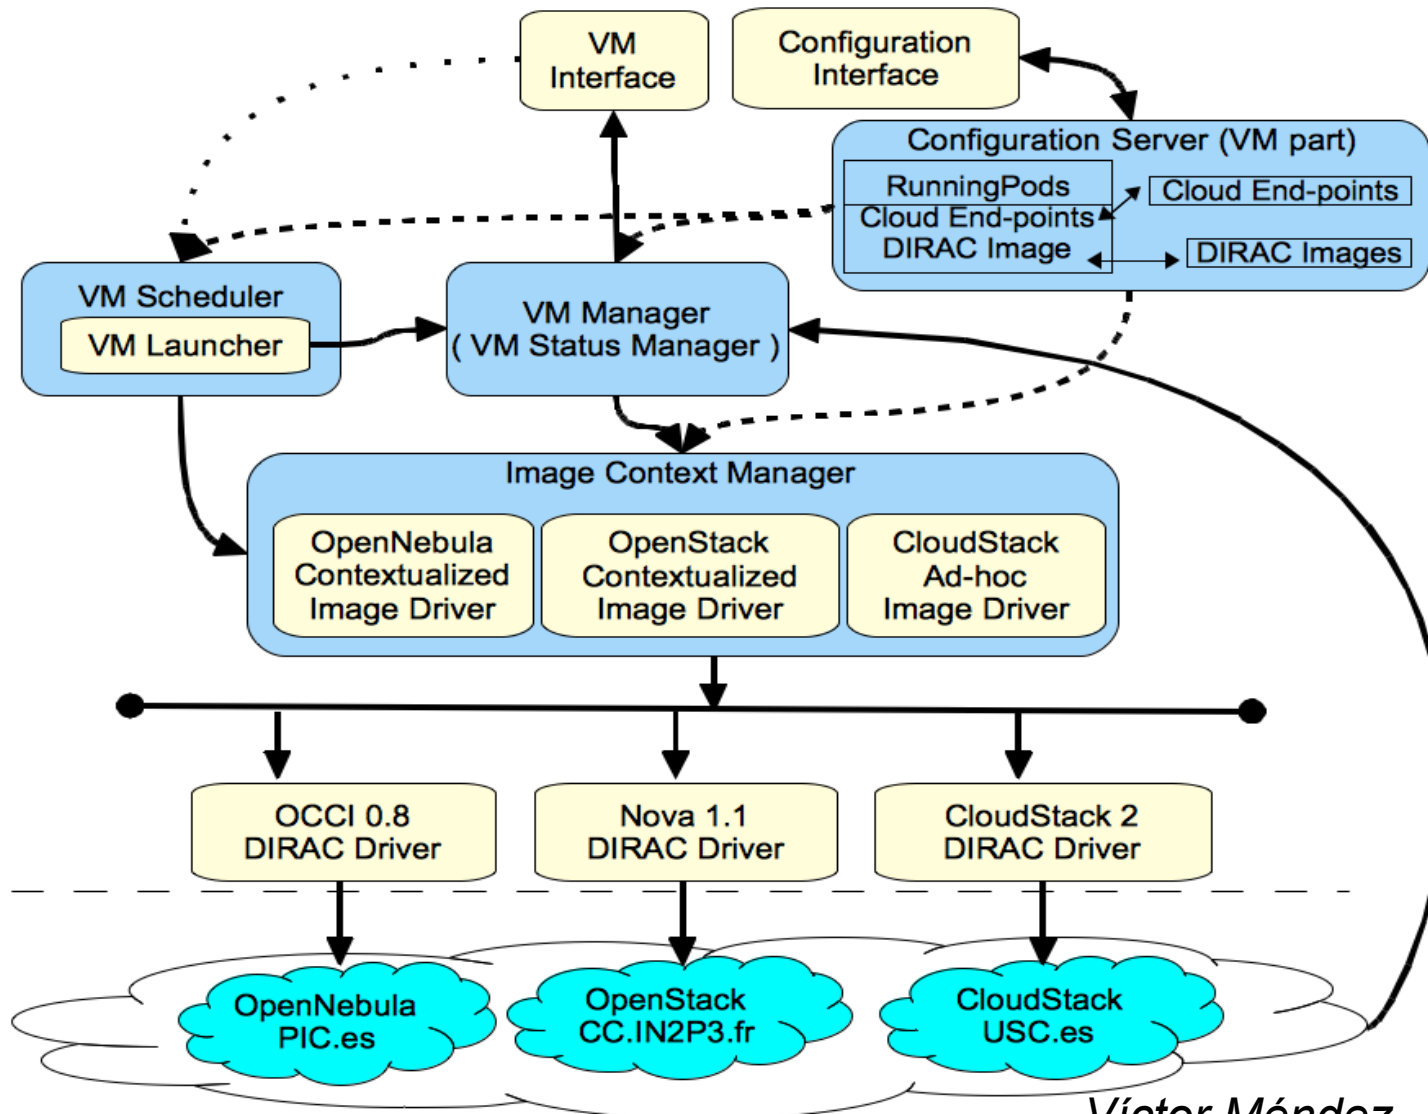


VMDIRAC Architecture



Running Pod: logical abstraction particular running conditions.

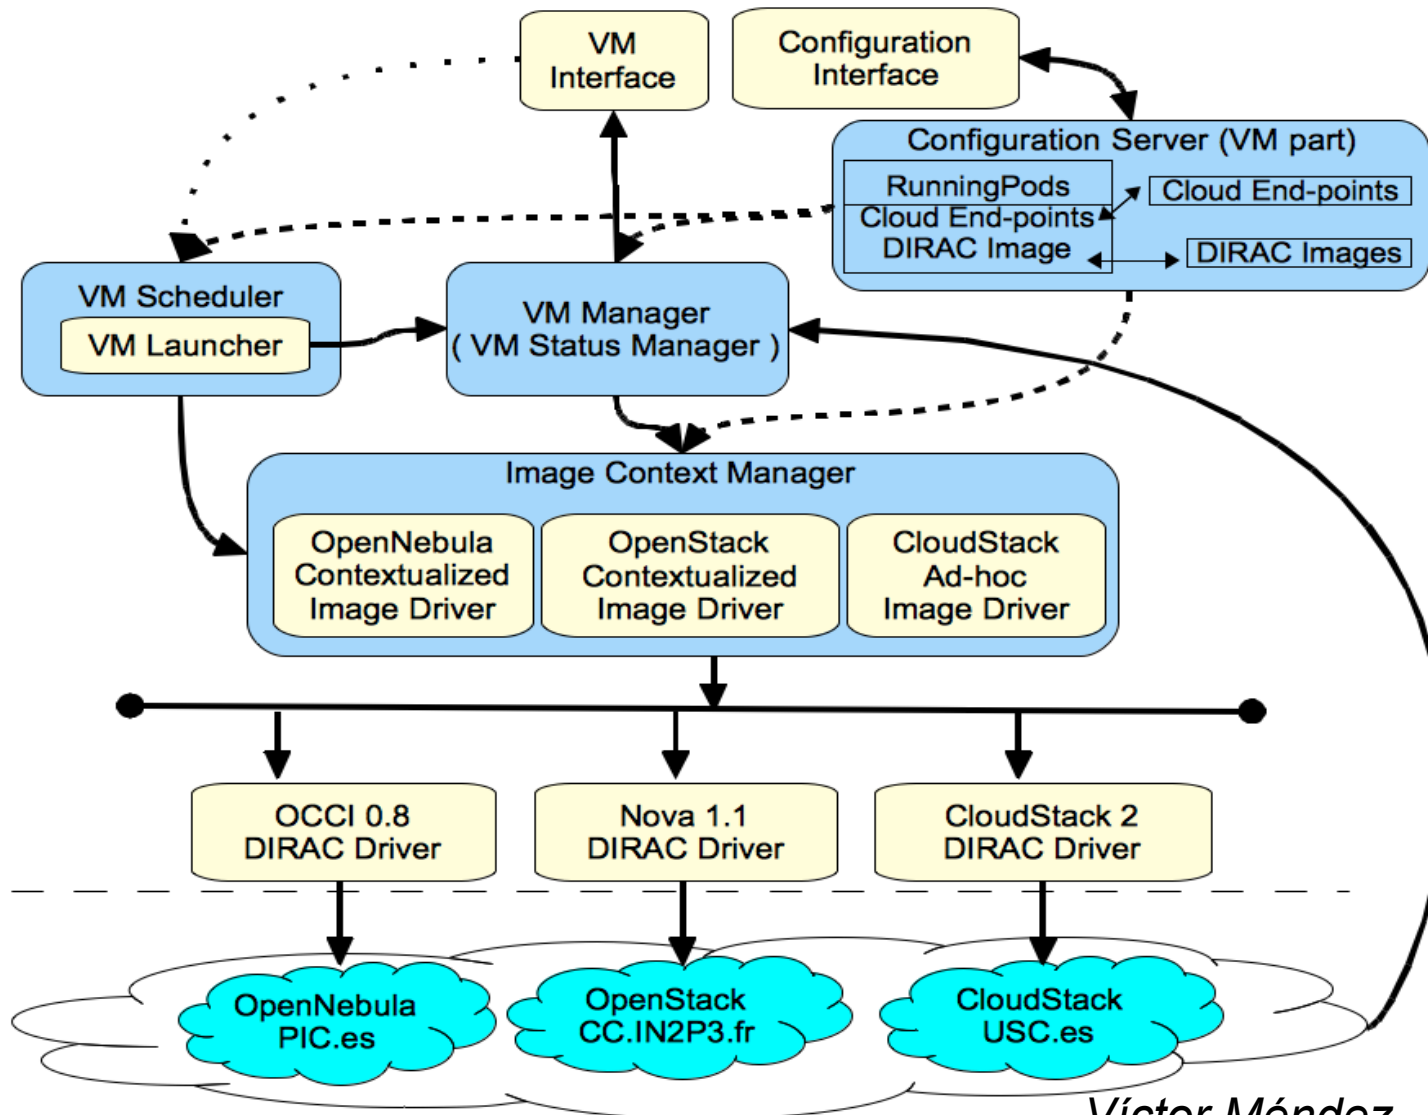
VMDIRAC Architecture



Running Pod: logical abstraction particular running conditions.

DIRAC image: boot image and Context

VMDIRAC Architecture



Running Pod: logical abstraction particular running conditions.

DIRAC image: boot image and Context

Cloud End-points: Specific IaaS API and Context

VMDIRAC Web Portal

System ▾ Jobs ▾ Views ▾ Virtual machines ▾ Tools ▾ Selected setup: VMDIRAC-Production ▾

Select all Select none Stop

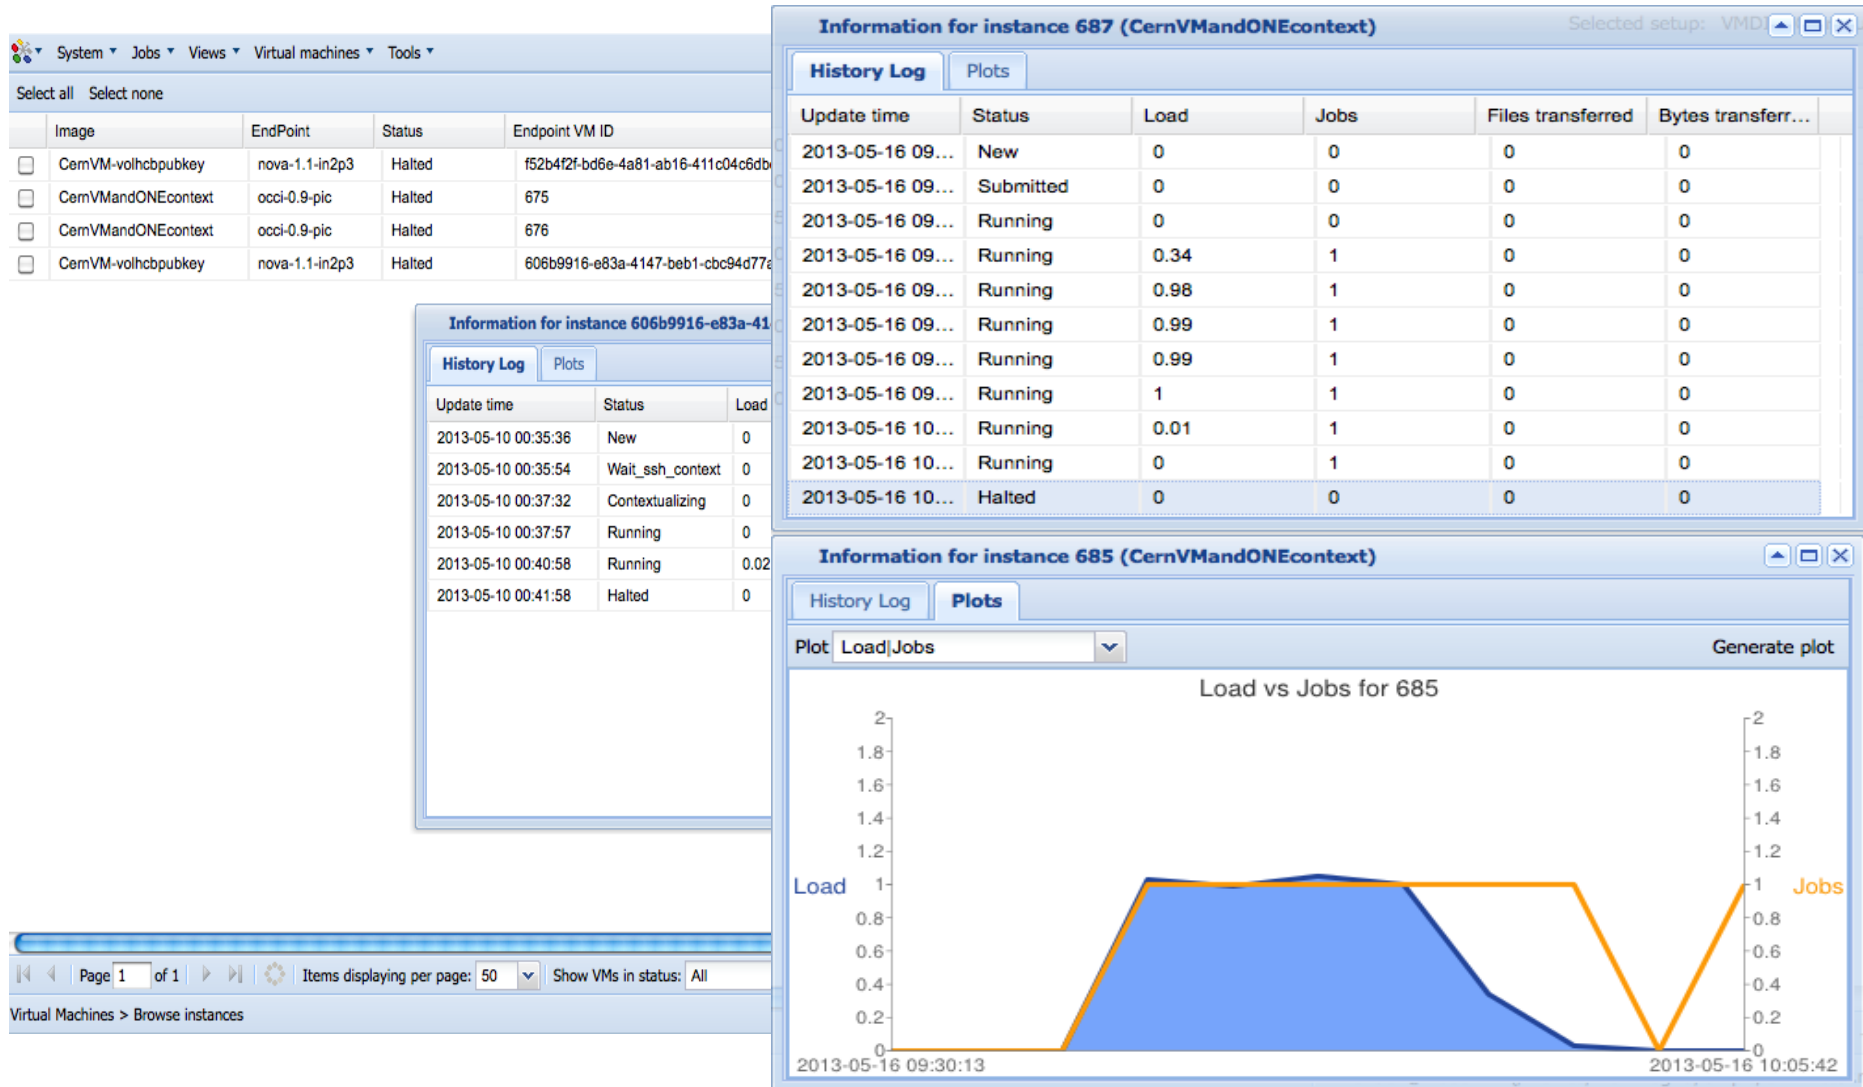
	Image	EndPoint	Status	Endpoint VM ID	IP	Load	Uptime	Jobs	Last Update (UTC) ▾	Error
<input type="checkbox"/>	CernVM-volhcbpubkey	nova-1.1-in2p3	Halted	f52b4f2f-bd6e-4a81-ab16-411c04c6dbd2	134.158.245.4	0.09	0:05:28	0	2013-05-10 01:02:01	
<input type="checkbox"/>	CernVMandONEcontext	occi-0.9-pic	Halted	675	193.109.173.210	0.06	0:05:48	5	2013-05-10 01:01:28	
<input type="checkbox"/>	CernVMandONEcontext	occi-0.9-pic	Halted	676	193.109.173.211	0.11	0:05:02	5	2013-05-10 01:00:42	
<input type="checkbox"/>	CernVM-volhcbpubkey	nova-1.1-in2p3	Halted	606b9916-e83a-4147-beb1-cbc94d77a711	134.158.245.3	0.02	0:05:01	2	2013-05-10 00:41:58	

Information for instance 606b9916-e83a-4147-beb1-cbc94d77a711 (CernVM-volhcbpubkey) [▲] [□] [×]

History Log Plots

Update time	Status	Load	Jobs	Files transfe...	Bytes transf...
2013-05-10 00:35:36	New	0	0	0	0
2013-05-10 00:35:54	Wait_ssh_context	0	0	0	0
2013-05-10 00:37:32	Contextualizing	0	0	0	0
2013-05-10 00:37:57	Running	0	0	0	0
2013-05-10 00:40:58	Running	0.02	2	0	0
2013-05-10 00:41:58	Halted	0	0	0	0

VMDIRAC Web Portal



The screenshot displays the VMDIRAC Web Portal interface. At the top, there are navigation menus for System, Jobs, Views, Virtual machines, and Tools. Below these is a table listing virtual machines with columns for Image, EndPoint, Status, and Endpoint VM ID.

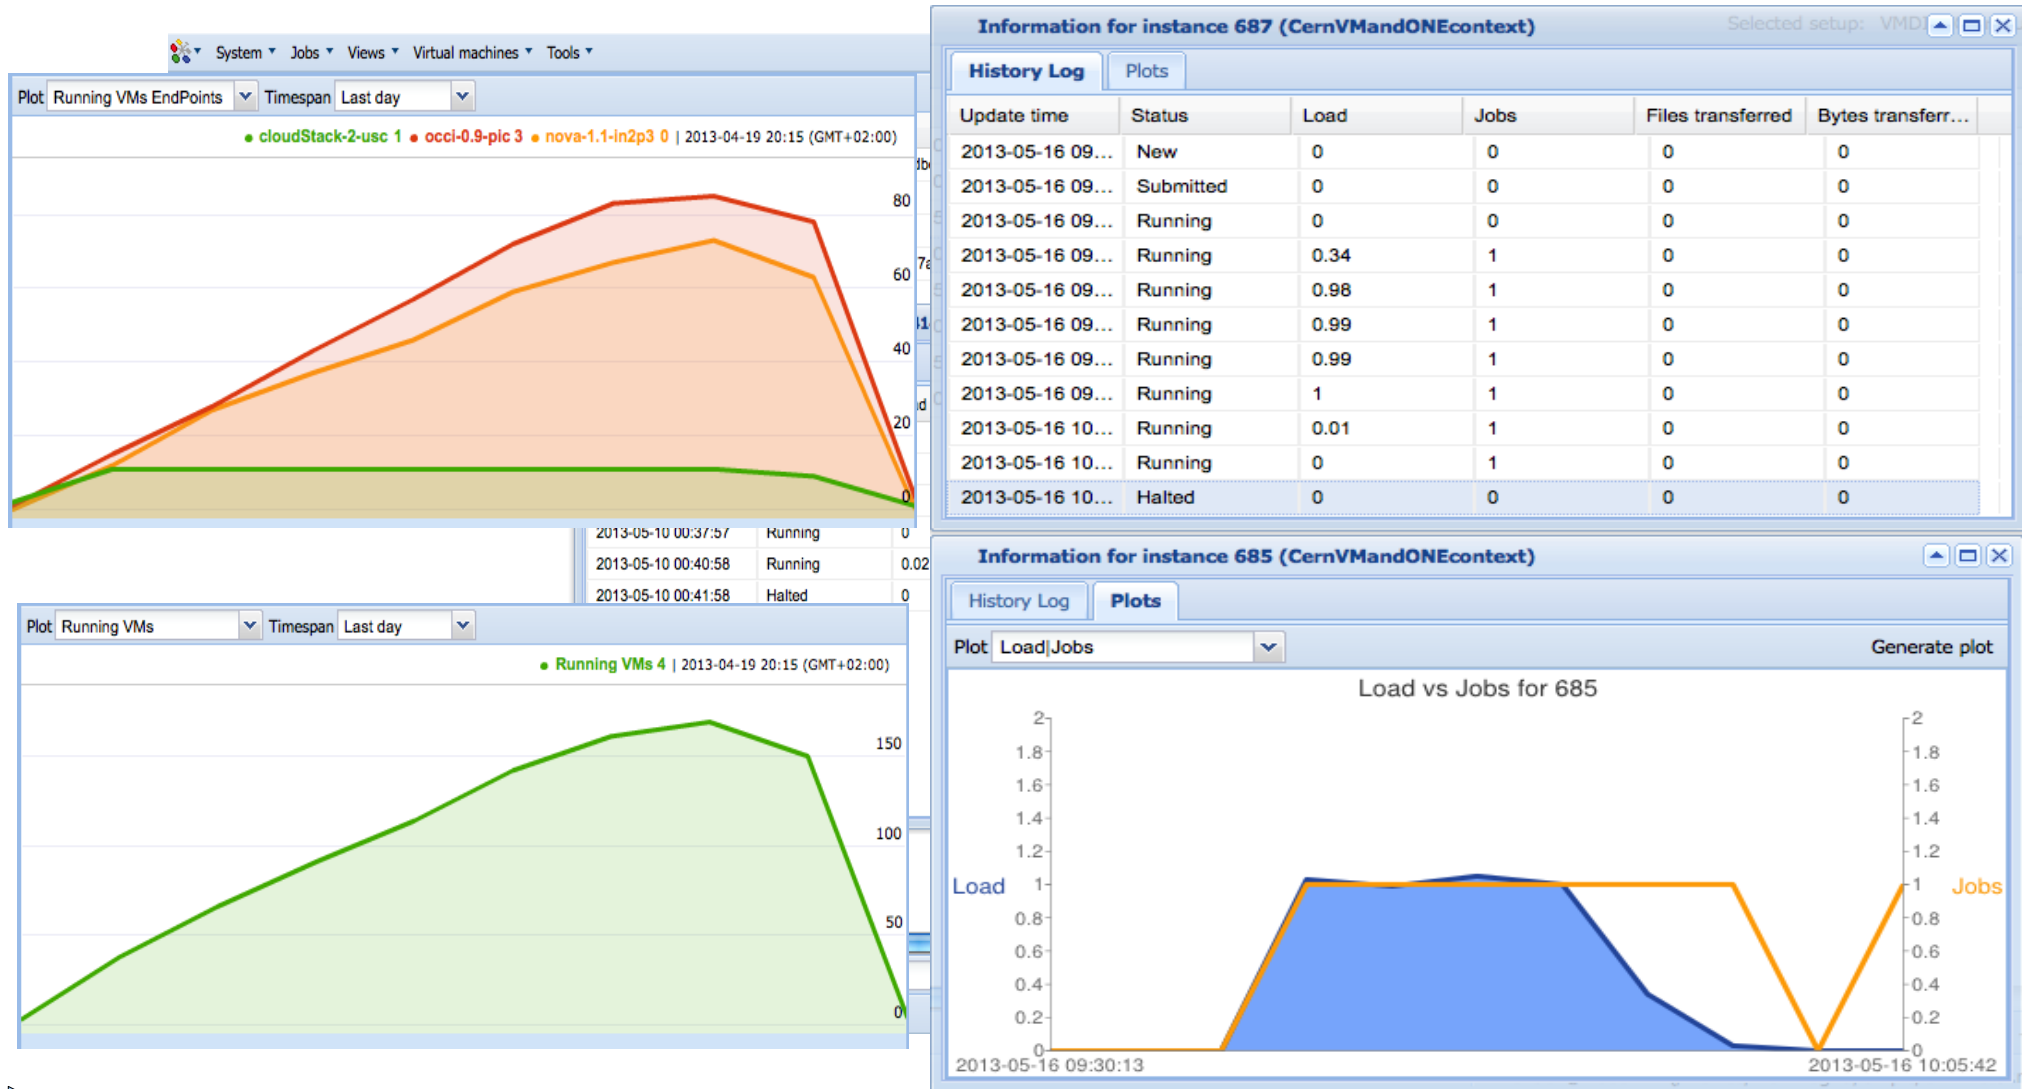
Two detailed windows are open, showing information for specific instances:

- Information for instance 687 (CernVMandONEcontext):** This window has tabs for History Log and Plots. The History Log table shows the following data:

Update time	Status	Load	Jobs	Files transferred	Bytes transferr...
2013-05-16 09...	New	0	0	0	0
2013-05-16 09...	Submitted	0	0	0	0
2013-05-16 09...	Running	0	0	0	0
2013-05-16 09...	Running	0.34	1	0	0
2013-05-16 09...	Running	0.98	1	0	0
2013-05-16 09...	Running	0.99	1	0	0
2013-05-16 09...	Running	0.99	1	0	0
2013-05-16 09...	Running	1	1	0	0
2013-05-16 10...	Running	0.01	1	0	0
2013-05-16 10...	Running	0	1	0	0
2013-05-16 10...	Halted	0	0	0	0
- Information for instance 685 (CernVMandONEcontext):** This window also has tabs for History Log and Plots. The Plot section shows a graph titled "Load vs Jobs for 685". The plot displays Load (blue area) and Jobs (orange line) over time. The x-axis represents time from 2013-05-16 09:30:13 to 2013-05-16 10:05:42. The y-axis represents Load and Jobs, ranging from 0 to 2.0. The graph shows a period of high load (around 1.0) and 1 job running, followed by a drop in both metrics.

At the bottom of the main interface, there is a pagination control showing "Page 1 of 1" and "Items displaying per page: 50".

VMDIRAC Web Portal



- X509 proxy authentication and authorization transparent for user by the DIRAC user/group system integrating third party VOMS
- DIRAC image setup to run VMs:
 - Ad-hoc image
 - Golden image and dynamic contextualization
 - HEPiX contextualization (OpenNebula, OpenStack)
 - SSH generic contextualization

- VM horizontal auto-scaling setup I:
 - VM allocation policy
 - Elastic: Depending in the work of the DIRAC TQ:
 - CPUPerInstance: VM Scheduler option to balance efficiency and wall time:
 - a) Zero to submit a new VM with no minimal CPU in the jobs of the tasks queue.
 - b) A longer value could be the average required CPU of the jobs as a compromise solution between VM efficiency and total wall time.
 - c) A very large value to maximize the efficiency in terms of VM creation overhead, for the cases where the production total wall time is not a constrain.
 - Static: A constant number of VMs defined by IaaS provider

- VM horizontal auto-scaling setup II:
 - VM stoppage policy:
 - Elastic: VM is automatically stopped by DIRAC if there are no more jobs running in the last VM halting margin time, which is an option to be setup.
 - Never: VMs is not automatically stopped by DIRAC
 - VM is requested to stop by the VO operator or by IaaS provider using DIRAC interface (Web or HEPiX machine features)
 - VM Monitor Agent manages to orderly shutdown

- v1r0: Cloud VM allocation by resource use accounting and top-down resource usage policy
- v1.r1: Integration of the Federated Cloud Accounting Service
- v1.rX: Integration of other Federated Cloud Services: Information System, Federated Market Place
- v1.rX: Cost/Price efficiency integration of commercial and community Clouds

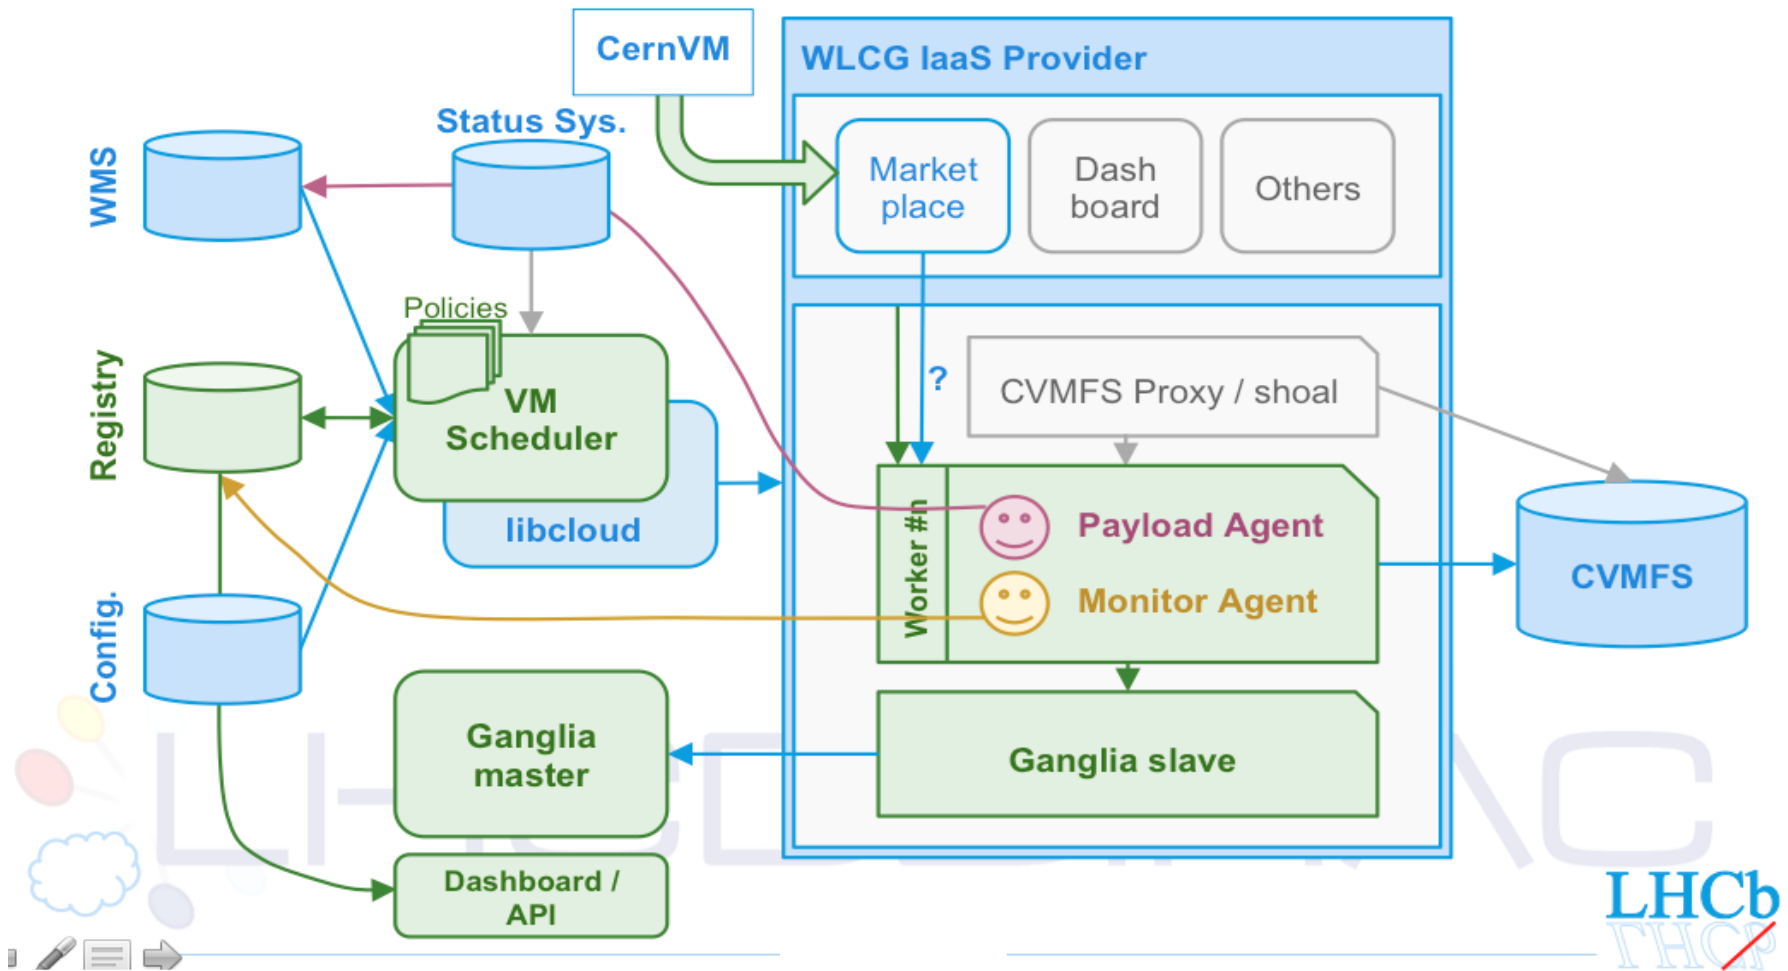
Cloud Resources in the LHCb computing

- Running on Production at:
 - CERN (OpenStack)
 - PIC (OpenNebula)
- Jobs Types:
 - LHCb Monte Carlo Simulations
 - Data processing (LHCb offline)

Cloud Resources in the LHCb computing

- Integrating Cloud end-points as computing resources of DIRAC sites lcg.cern.ch and lcg.pic.es by HEPiX context
- VM multicore optimization on memory consumption of [Nathalie Rauschmayr work](#) with GaudiMP, the LHCb software framework. Integration with DIRAC in [CHEP2013](#)

Cloud Resources in the LHCb computing



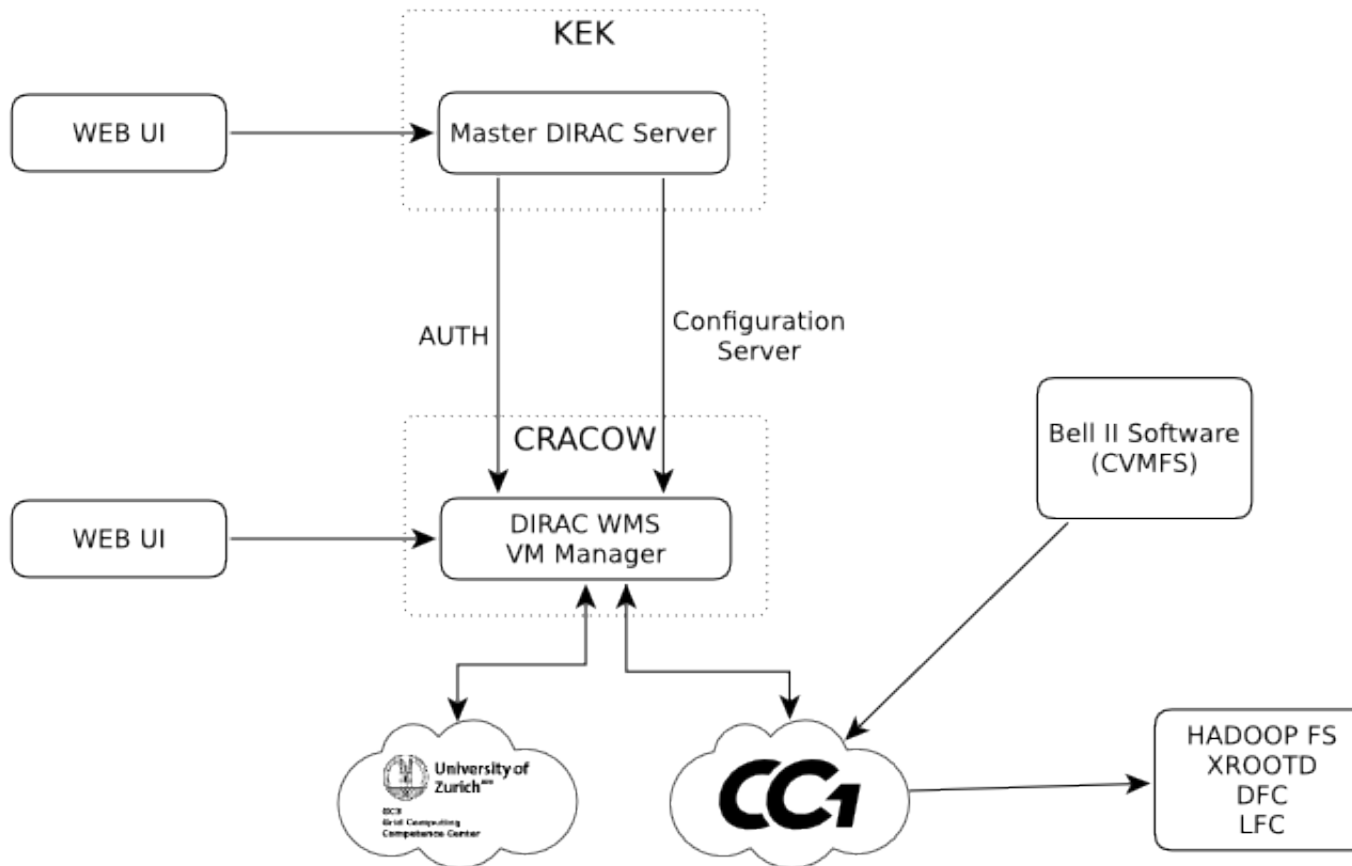
Cloud Resources in the Belle II computing

- Testbed DIRAC WMS with VMDIRAC located in Krakow, but Configuration, Accounting and other systems are fetching from master Belle II DIRAC server located in KEK (Japan)
- Testbed is in simulation jobs, it is not using input data for jobs, but in future will be

Cloud Resources in the Belle II computing

- Output ROOT files are stored straightly on grid SEs and registered in LFC and Amga metadata catalog (in KEK)
- VMDIRAC is connected to CC1 cloud now by EC2 interface and soon to OpenStack cloud in Zurich
- Testbed successful, going to production in October 2013

Cloud Resources in the Belle II computing



Belle II Cloud scheme courtesy of Rafał Grzymkowski in [CHEP2013](#)

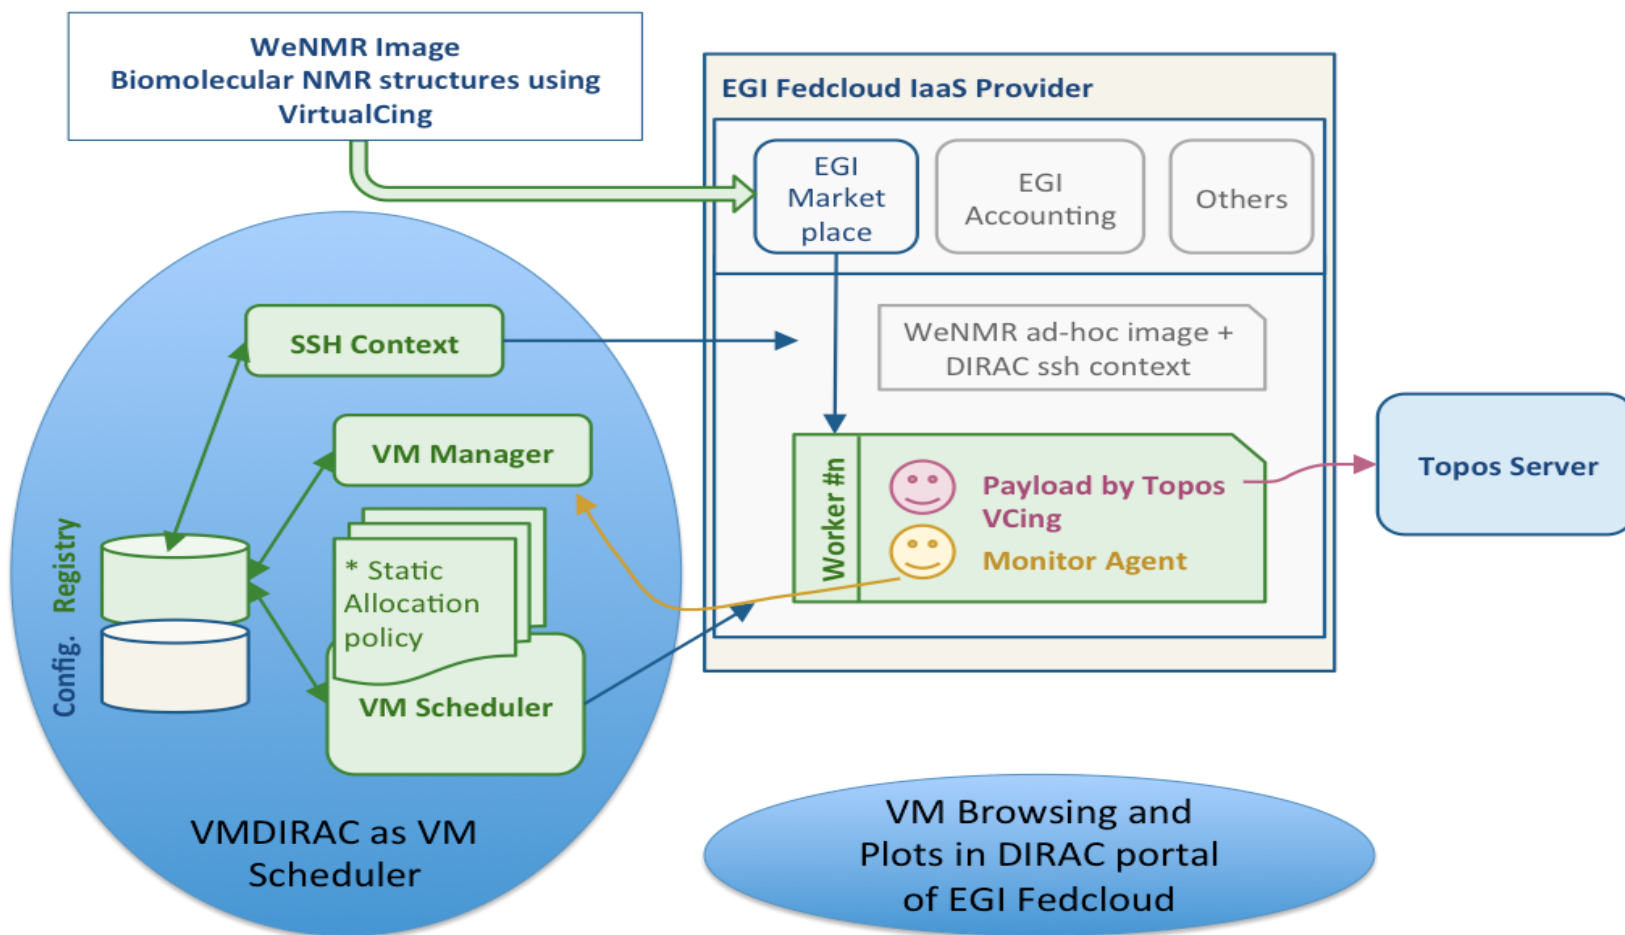
WeNMR-VMDIRAC EGI Fedcloud Use Case

- Validating and improving **biomolecular NMR** structures using VCing, a suite of ~25 programs
- Demonstrated using **EGI Fedcloud infrastructure**
- **EGI Fedcloud computing model** is moving beyond the frontier, integrating multiple IaaS providers, and deploying necessary services for eScience communities

WeNMR-VMDIRAC EGI Fedcloud Use Case

- **WeNMR with VMDIRAC** taking advantage of automatic VM management in VMDIRAC
- VMDIRAC portal as VM scheduler
- DIRAC broker is not involved, job payload for the VCing VM is provided by the ToPoS server
- A mix of *ad-hoc* VCing image and DIRAC ssh contextualization, which runs only the VM Monitor to manage VM status and stoppage

WeNMR-VMDIRAC EGI Fedcloud Use Case



France Grilles Cloud and FG-DIRAC portal

- FG DIRAC portal is supporting different communities, mainly of life sciences
- FG Cloud aggregates multiple IaaS providers, using the FG-DIRAC portal with VMDIRAC server, successful in testbed and now is going to production
- Contextualization is ssh for generic image management, HEPiX with Cernvm can also be used

- VMDIRAC is a proved tool to aggregate IaaS providers in the level of NGIs supporting multiple VOs, and also in medium and big scientific communities
- VMDIRAC provides solutions to SaaS deployment and management in a wide range of scientific communities
- Small communities can be federated to face the DIRAC portal operations (EGI Fedcloud, FG Cloud)

- Medium and big communities may exploit their own DIRAC portal with Cloud resources (LHCb WLCG Clouds, Bell II private Cloud)
- VMDIRAC strategy is addressing sustainability through *industrial concentration* of SaaS management in Federated Hybrid Clouds, at the same time allowing *local development* by the aggregation of distributed IaaS resources
- Detailed instructions for potential IaaS providers and users in the “How to” paper in Proc. ADVCOMP 2013



News in
<https://github.com/DIRACGrid/VMDIRAC/wiki>

Thanks