



Enabling Grids for E-scienceE

# ARDA Experiment Dashboard

*Ricardo Rocha (ARDA – CERN)  
on behalf of the Dashboard Team*

[www.eu-egee.org](http://www.eu-egee.org)

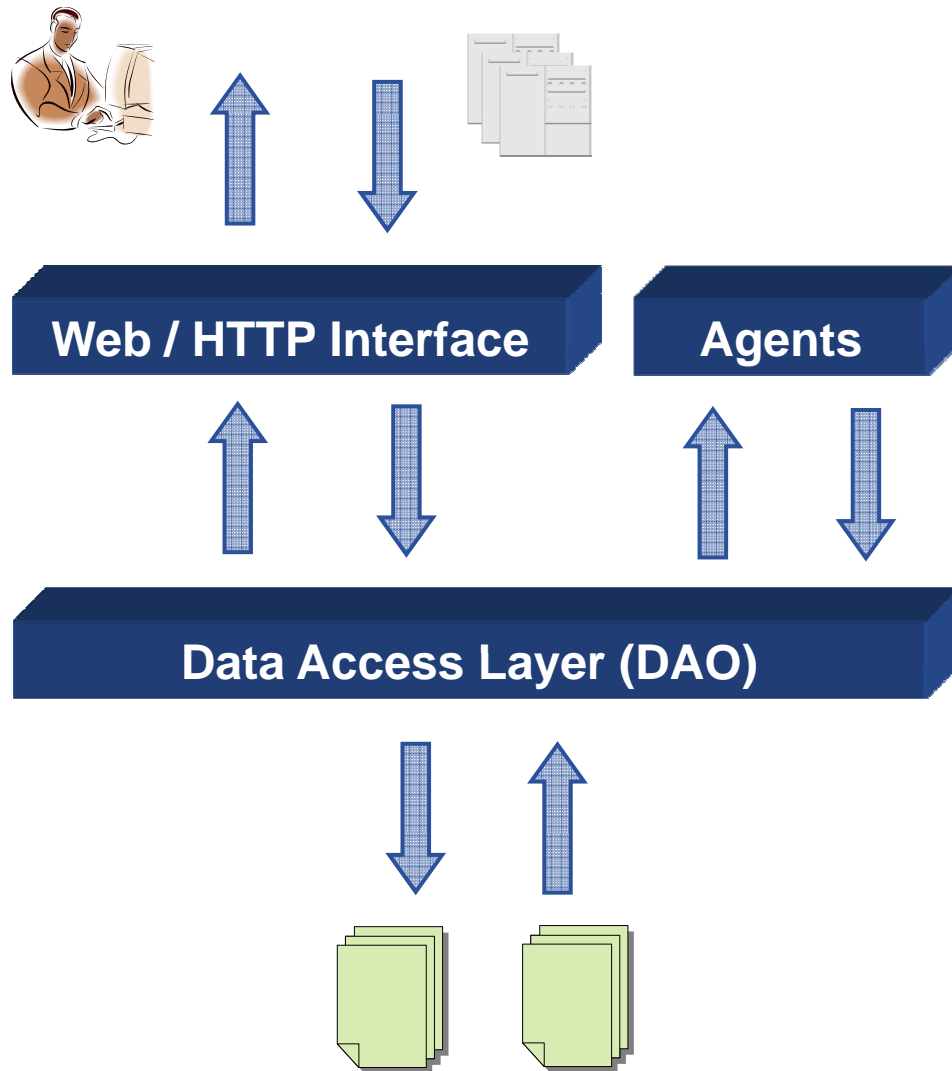


Information Society



- **Background**
- **Dashboard Framework**
- **VO Monitoring Applications**
  - Job Monitoring
  - Site Monitoring / Efficiency
  - Data Management Monitoring
  - Additional Applications
- **Conclusion and Future Work**

- **Started in 2005 inside the EGEE/ARDA group**
- **First application: Grid Job / Application Monitoring for the CMS experiment**
  - implementation in PHP / Python
- **Redesign early 2006**
  - fully python based solution
  - more modular / structured approach
  - easily extensible
- **Additional application areas: data management (ATLAS DDM), site efficiency monitoring, ...**

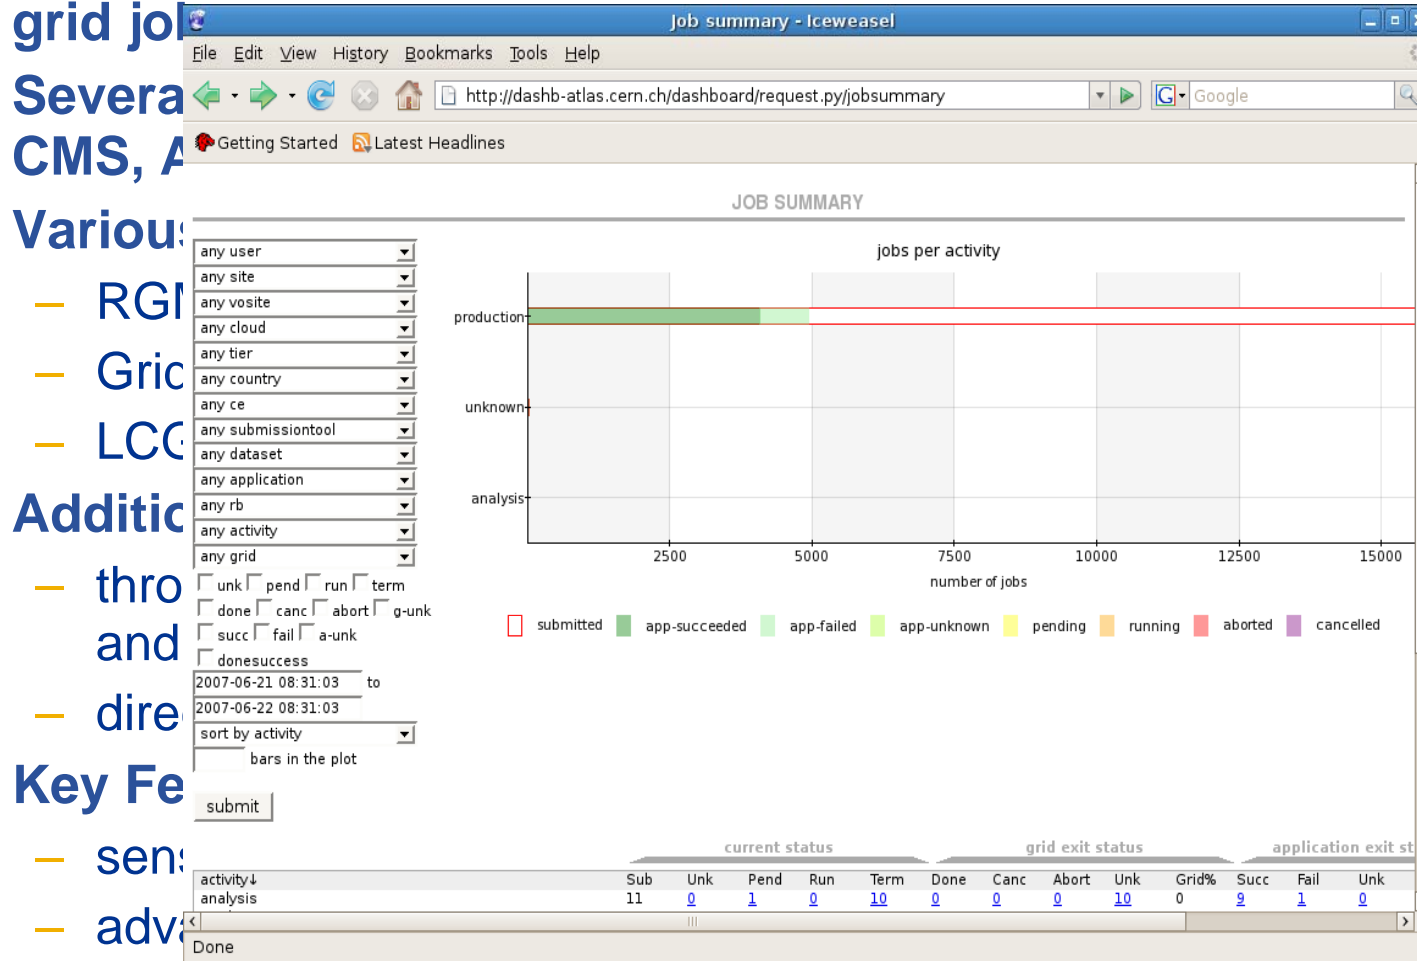


- **Dashboard Clients**
  - Scripts: pycurl, ...
  - Command line tools (optparser + pycurl)
  - Shell based: curl, ...
- **Web Application**
  - Apache + mod\_python
  - Model View Controller (MVC) pattern
  - multiple output formats: plain text, CSV, XML, XHTML
  - GSI support using gridsite
- **Agents**
  - collectors: RGMA, ICXML, BDII, ...
  - stats generation, alert managers, ...
  - *Service Configurator* pattern
  - common configuration (XML file) and management: *stop, start, status, list*
  - common monitoring mechanism
- **Data Access Layer (DAO)**
  - interfaces available to different backends (Oracle and PostgreSQL mainly, easy to add additional ones)
  - connection pooling

- **Build and development environment**
  - based on python distutils (with several extensions)
  - covers code validation, binaries and documentation generation, unit testing and reports
  - automatic build for each of the release branches
  - packaging uses RPMs – APT repository available
- **Release procedure**
  - three main branches: nightly, unstable, stable
  - releases per component
  - enforced versioning scheme (no manual versioning or tagging, all done via distutils command extensions)
- **Interesting links**
  - Developers guide:
    - <http://dashb-build.cern.ch/build/nightly/doc/guides/common/html/dev/index.html>
  - Savannah Project
    - <http://savannah.cern.ch/groups/dashboard>

- **Real time and summary views over the virtual organization (VO) grid jobs**
- **Several instances in production serving different communities: CMS, ATLAS, LHCb, Alice, VL MED**
- **Various grid information sources used:**
  - RGMA
  - GridPP XML files collection
  - LCG BDII
- **Value added: VO specific information**
  - through job instrumentation. Using Monalisa's ApMon (CMS), Panda and Ganga monitoring (ATLAS)
  - directly querying VO databases (ex: ATLAS production database)
- **Key Features:**
  - sensible merging of information from different sources
  - advanced filtering for different usages (VO manager, site admin, community user)

- Real time and summary views over the virtual organization (VO)



activity	current status									grid exit status				application exit status		
	Sub	Unk	Pend	Run	Term	Done	Canc	Abort	Unk	Grid%	Succ	Fail	Unk			
analysis	11	0	1	0	10	0	0	0	10	0	9	1	0			

- Sever...
- CMS, A...

- Variou...

- RGI
- Grid
- LCC

- Additic...

- thro
- and
- dire

- Key Fe...

- sen
- adv
- commu

unities:

), Panda

base)

lmin,

- **Task Monitoring**
  - deployed and used in CMS
- **Integration with SAM tests**
  - already using the new LCG standards
  - prototype in place
- **Alert mechanism**
  - in development
- **HTTP API for publishing job information**
  - very easy to integrate with existing tools
  - similar to the mechanism used for data management



- **Built on top of the job monitoring data**
- **Main goal: identify reasons for job failures in sites**
- **Uses the information coming from RGMA**
- **Available today for the same set of communities: ATLAS, CMS, LHCb, Alice, VL MED**
- **Provides both summary and detailed information**
- **Current ongoing work**
  - provide generic (non VO) specific view over the data

- 
- 
- 
- 
- 
- 
- 

Site Efficiency - Iceweasel

File Edit View History Bookmarks Tools Help

http://lxarda04.cern.ch/dashboard/request.py/SiteEfficiency

Getting Started Latest Headlines

CNR-ILC-PISA	585	428	57.75%
INFN-NAPOLI-PAMELA	520	344	60.19%
UKI-SOUTHGRID-OX-HEP	1232	753	62.07%
INFN-CNAF	689	401	63.21%
RU-Protvino-IHEP	11225	6460	63.47%
HG-05-FORTH	1290	734	63.74%
INFN-PISA	945	533	63.94%
UKI-SCOTGRID-DURHAM	836	445	65.26%
FZK-LCG2	82716	37586	68.76%
INFN-NAPOLI	495	194	71.84%
griditce01.na.infn.it:2119/jobmanager-lcgpbs-alice	495	194	71.84%

Jobids	# jobs	Successful?	Error message
<a href="#">See all the jobids...</a>	28	Ignored	Job MaxRetryCount hit
<a href="#">See all the jobids...</a>	1	Ignored	Job MaxRetryCount hit
<a href="#">See all the jobids...</a>	1	Ignored	Job MaxRetryCount hit
<a href="#">See all the jobids...</a>	107	No	local batch system problem
<a href="#">See all the jobids...</a>	13	No	job successfully submitted to Globus
<a href="#">See all the jobids...</a>	10	No	local batch system problem

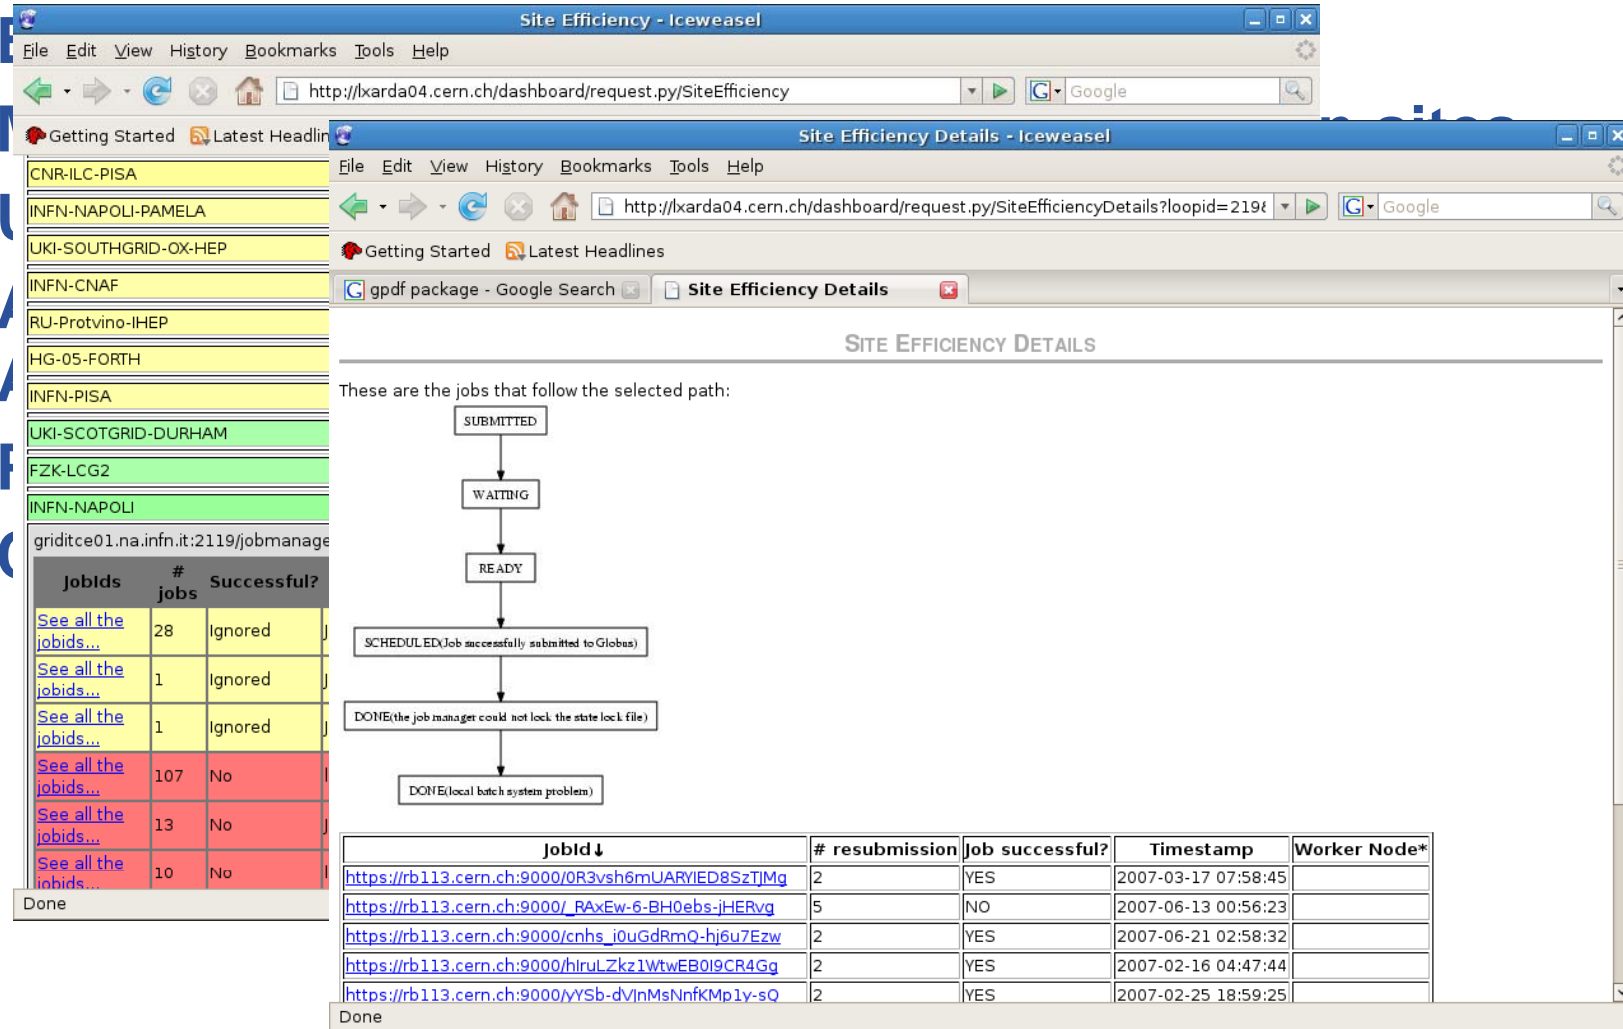
Done

n sites

ilities:

nation

ata




The screenshot shows the 'Site Efficiency - Iceweasel' web application. On the left is a sidebar with a list of sites: CNR-ILC-PISA, INFN-NAPOLI-PAMELA, UKI-SOUTHGRID-OX-HEP, INFN-CNAF, RU-Protvino-IHEP, HG-05-FORTH, INFN-PISA, UKI-SCOTGRID-DURHAM, FZK-LCG2, and INFN-NAPOLI. Below the list is a table with columns 'Jobids', '# jobs', and 'Successful?'. The main content area is titled 'SITE EFFICIENCY DETAILS' and contains a flowchart showing the job lifecycle: SUBMITTED -> WAITING -> READY -> SCHEDULED (Job successfully submitted to Globus) -> DONE (the job manager could not lock the state lock file) -> DONE (local batch system problem). Below the flowchart is a table with columns: 'jobid ↓', '# resubmission', 'job successful?', 'Timestamp', and 'Worker Node\*'. The table contains five rows of job data.

Jobids	# jobs	Successful?
<a href="#">See all the jobids...</a>	28	Ignored
<a href="#">See all the jobids...</a>	1	Ignored
<a href="#">See all the jobids...</a>	1	Ignored
<a href="#">See all the jobids...</a>	107	No
<a href="#">See all the jobids...</a>	13	No
<a href="#">See all the jobids...</a>	10	No

jobid ↓	# resubmission	job successful?	Timestamp	Worker Node*
<a href="https://rb113.cern.ch:9000/0R3vsh6mUARYIED8SzTJMg">https://rb113.cern.ch:9000/0R3vsh6mUARYIED8SzTJMg</a>	2	YES	2007-03-17 07:58:45	
<a href="https://rb113.cern.ch:9000/_RAXEw-6-BH0ebs-jHERvg">https://rb113.cern.ch:9000/_RAXEw-6-BH0ebs-jHERvg</a>	5	NO	2007-06-13 00:56:23	
<a href="https://rb113.cern.ch:9000/cnhs_i0uGdRmQ-hj6u7Ezw">https://rb113.cern.ch:9000/cnhs_i0uGdRmQ-hj6u7Ezw</a>	2	YES	2007-06-21 02:58:32	
<a href="https://rb113.cern.ch:9000/hirUL_Zkz1WtwEB0I9CR4Gg">https://rb113.cern.ch:9000/hirUL_Zkz1WtwEB0I9CR4Gg</a>	2	YES	2007-02-16 04:47:44	
<a href="https://rb113.cern.ch:9000/yYSb-dvInMsNnfKmp1y-sQ">https://rb113.cern.ch:9000/yYSb-dvInMsNnfKmp1y-sQ</a>	2	YES	2007-02-25 18:59:25	

- **Tied to the ATLAS Distributed Data Management (DDM) system**
- **Used successfully both in the production and Tier0 test environments**
- **Data sources:**
  - DDM site services: the main source, providing all the transfer and placement information
  - SAM tests: for correlation of DDM results with the state of the grid fabric services
  - Storage space availability: from BDII but soon including other available tools
- **Views over the data:**
  - Global: site overview covering different metrics (throughput, files / datasets completed, ...); summary of the most common errors (transfer and placement)
  - Detailed: starting from the dataset state, to the state of each of its files, to the history of each single file placement (all state changes)



ent (DDM) system  
Tier0 test

all the transfer and  
the state of the grid  
including other available

throughput, files /  
common errors (transfer

and placement)

- Detailed: starting from the dataset state, to the state of each of its files, to the history of each single file placement (all state changes)

## ent (DDM) system Tier0 test

### ATLAS Distributed Data Management

The dashboard displays 'Overview' and 'Dataset Info' tabs. The 'Overview' section includes two stacked bar charts showing activity over time (08h to 18h) for various clouds. The 'Activity Summary' table provides a detailed breakdown of transfers and errors for each cloud.

Cloud	Transfers				Errors			
	Efficiency	Avg Throughput	Files Done	Datasets Done	Transfer	Local	Remote	Central
ASGC	0%	0 MB/s	0		0	0	0	
BNL	73%	110 MB/s	3638		1341	0	0	
CERN	37%	0 MB/s	220		367	0	0	
CNAF	10%	35 MB/s	960		8684	0	0	
FZK	34%	35 MB/s	1177		2305	0	0	
LYON	66%	48 MB/s	1473		771	0	0	
NDGF	64%	19 MB/s	564		317	0	0	
PIC	8%	13 MB/s	358		4187	0	0	
RAL	0%	0 MB/s	0		0	0	0	
SARA	17%	10 MB/s	358		1797	0	0	
TRIUMF	79%	33 MB/s	979		261	0	0	

Legend for activity status: CRITICAL (red), WARNING (yellow), NORMAL (blue), GOOD (green), NO\_ACTIVITY (grey).

http://dashb-atlas-data-test.cern.ch/dashboard/request.p

and placeme

- Detailed: sta to the history

- **Other features**

- periodic site behavior reports (sent by email)
- alerts (on specific errors, when a site goes below a certain threshold, ...)

- **Coming soon**

- user specific views (authentication via X509 certificates) – “my datasets”
- better site summary data: overview of dataset / file states in the site (radar plots), average time in each placement step, additional error summaries
- python *query* API module
- python *publish* API module (open the tool to other applications / communities)



- **The Dashboard monitors the grid from the point of view of its communities**
  - and focuses on the different user's interests (managers, admins, end users)
- **Grid information is not enough (additional VO information is invaluable)**
- **Framework**
  - Flexible and stable: proven by the variety of applications available in production
  - Effort put into install / packaging paid off: first external installation has already been done (VLMED)
- **Future work**
  - integration with local monitoring systems (feed summaries back to the site admins)
  - improved alert system
  - adapt to recently defined data exchange / query standards
- **<http://dashboard.cern.ch>**