application level monitoring for LHC experiments with the experiment dashboard

Benjamin Gaidioz

Hepix fall 2007

the experiment dashboard project

- a monitoring project, showing the activity of VOs on the grid,
  - integration of several informations. For example:
    - grid info: jobs, computing/storage resources, topology,
    - VO info: type of job, application exit code, datasets.
- a project from EGEE/ARDA (CERN),
- a framework for collecting and showing information.

▸ "home page"

outline

dashboard

outline

dashboard

the dashboard framework

- most of the applications are implemented using a common Python framework.
  - clear distinction between information storage, information queries and rendering:

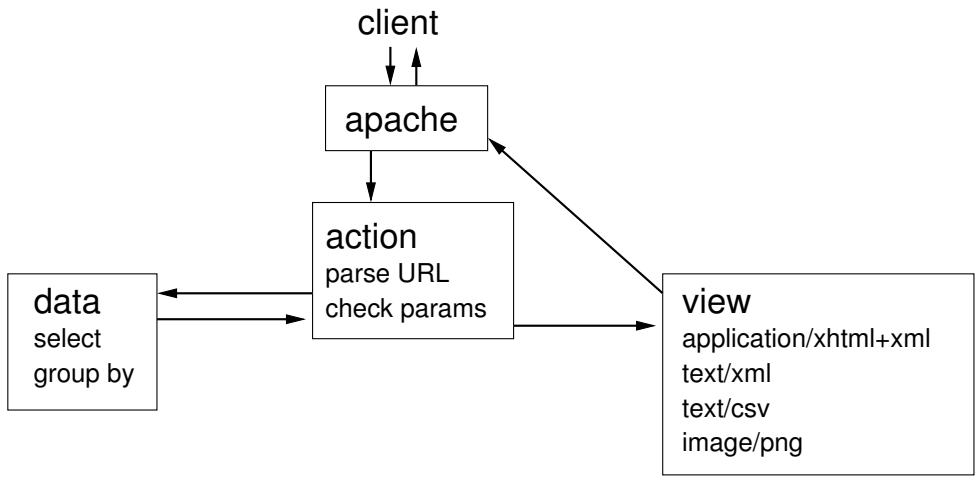

  - support for API, CLI.
  - performance: thread safety + session pooling + apache/mod_python,
- even a developer's guide is available, with examples.

the dashboard framework

- most of the applications are implemented using a common Python framework.
  - clear distinction between information storage, information queries and rendering:

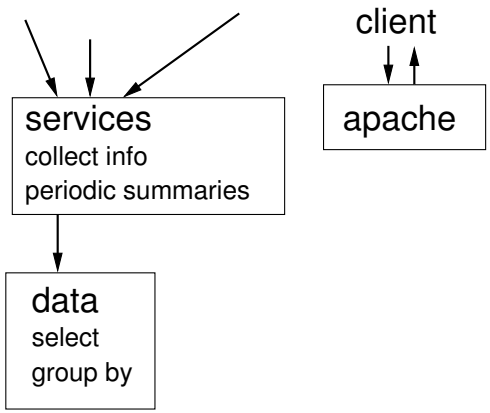

  - support for API, CLI.
  - performance: thread safety + session pooling + apache/mod_python,
- even a developer's guide is available, with examples.

features

- based on HTTP,
- multiple output formats (text/xml and text/csv coming for free),

```
curl -H 'Accept:  text/xml' http://...
```

- framework for API and CLI: based on HTTP, man pages, build of the command-line tool, common options,

```
from dashboard.api.production.ProductionQuery import ProductionQuery
query = ProductionQuery('dashb-atlas-prodsys-test.cern.ch', 80)
sites = query.errors(error='WRAPLCG_STAGEOUT_LCGCR',
    grouping='site', grid='LCG')
```

- framework for permanent *services* (info collectors, computation of summaries, etc.): monitoring, babysitting, common configuration, etc.
  - status exported to our main web server,
  - simple alerts in case of warning status: e-mail or SMS.

dashboard

releasing and build system

- a project module has a stable/unstable/nightly release,
  - RPMs are currently distributed using apt,
  - we also distribute some external RPMs (javascript toolkits, python packages not available in SLC4),
- the build system is based on the python distutils,
  - automatic building of SLC4 RPMs running at night,
  - also builds and installs the latest documentation (docbook),
  - (doesn't run the unit tests yet.)
  - possibility to trigger the build at anytime if needed.

operations

- applications using the common framework:
  - same tools maintenance, same log files, similar services, etc.
  - maintenance guides, documentation.
- hosts are SLC4 and are quattor managed,
- about maintenance, some recent developments:
  - dashboard services and Apache now running as “dashbop” UNIX user,
  - definition of an operator quattor role,
  - operators can run a limited set of commands: restart httpd, restart dashboard services,
  - permits to implement some simple and safe maintenance by non-experts.
- but still some legacy applications (no framework, SLC3, etc.).

operations (coming)

- most of our applications don't generate much load on the same host:
  - but it's not practical to run several per host,
  - applications A and B may require at some point different releases of a common dashboard RPM,
  - maintenance of A impacts B (restart httpd, same configuration files, same log files.),
  - it's a limit of our framework (not easy to have two configurations).
- dashboard applications are probably good candidates for hardware virtualization,
  - it could probably be practical for testing purposes (check the installation on a new host or maintenance procedure, run full functionality tests.

outline

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

  - grid info taken from RGMA, GridPP XML files, LCG BDII,
  - VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
  - installed for all four main experiments, plus the "vlemed" VO.

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

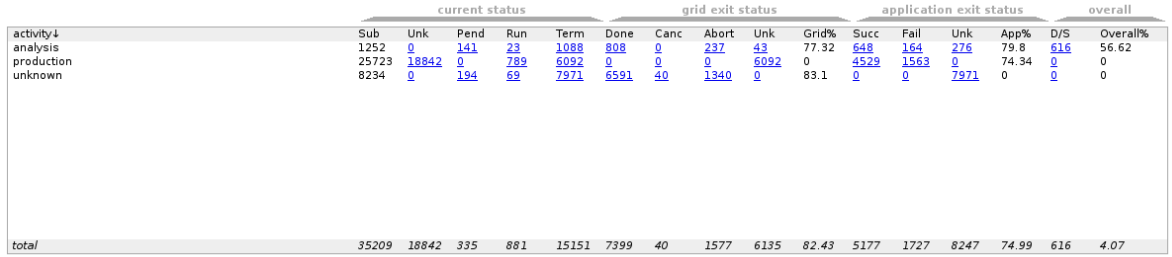

	current status					grid exit status					application exit status				overall	
activity↓	Sub	Unk	Pend	Run	Term	Done	Canc	Abort	Unk	Grid%	Succ	Fail	Unk	App%	D/S	Overall%
analysis	1252	0	141	23	1088	808	0	237	43	77.32	648	164	276	79.8	616	56.62
production	25723	18842	0	789	6092	0	0	0	6092	0	4529	1563	0	74.34	0	0
unknown	8234	0	194	69	7971	6591	40	1340	0	83.1	0	0	7971	0	0	0
total	*35209*	*18842*	*335*	*881*	*15151*	*7399*	*40*	*1577*	*6135*	*82.43*	*5177*	*1727*	*8247*	*74.99*	*616*	*4.07*

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

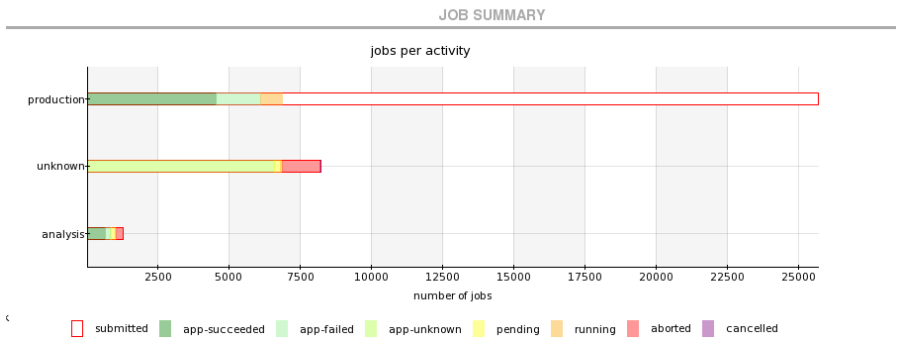

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

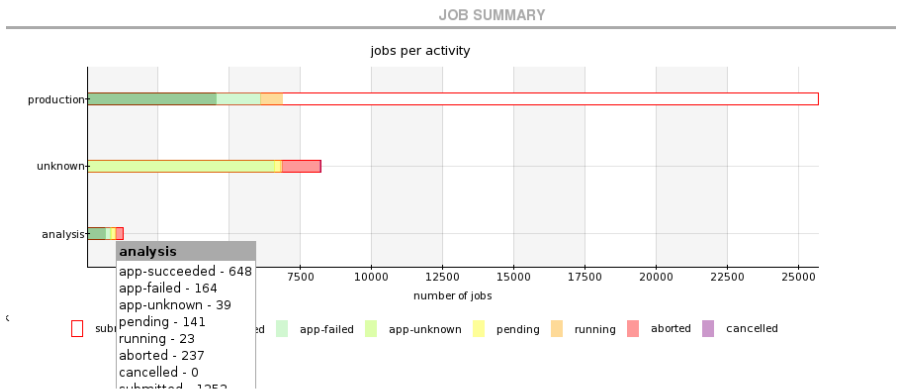

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

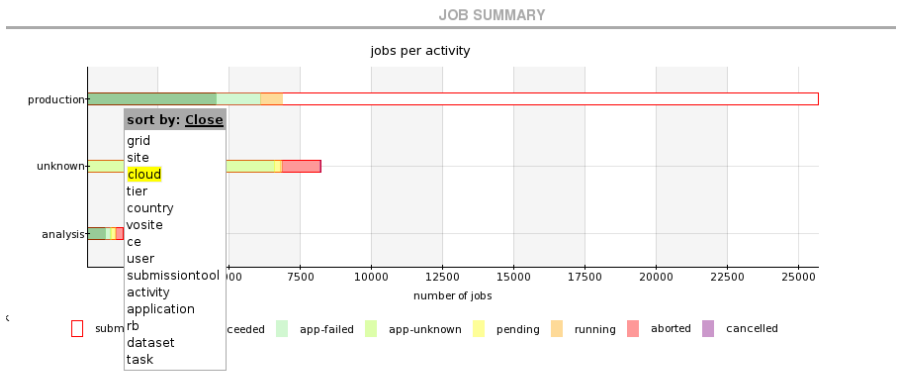

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

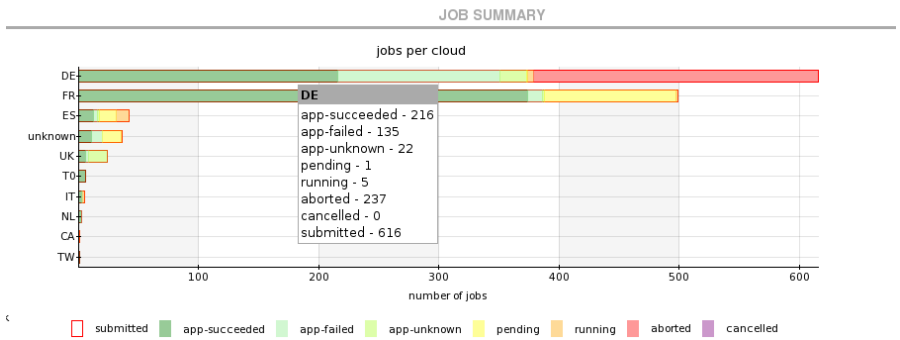

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

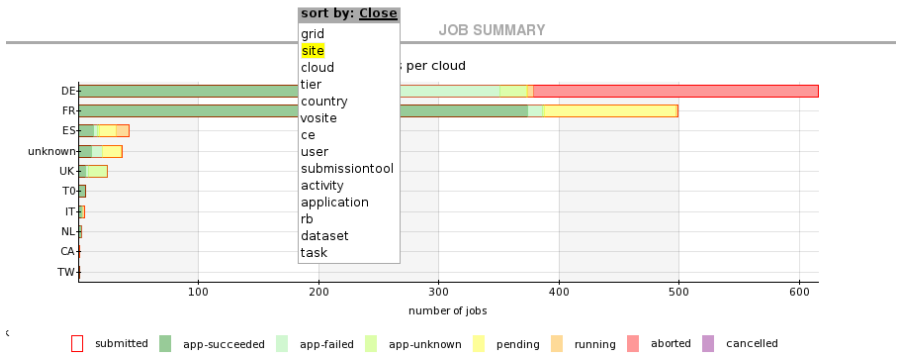

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

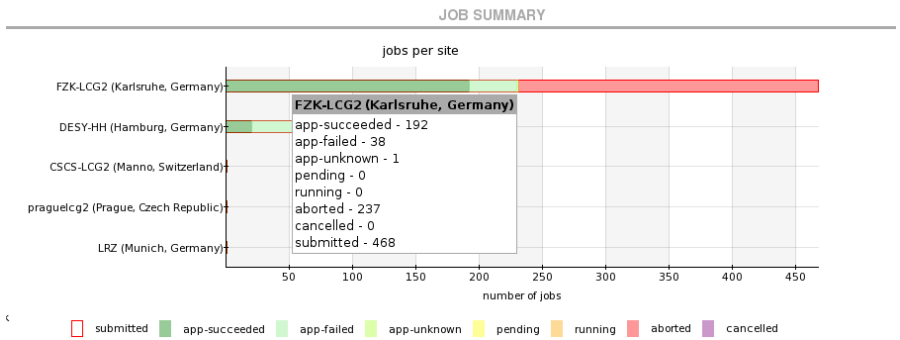

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

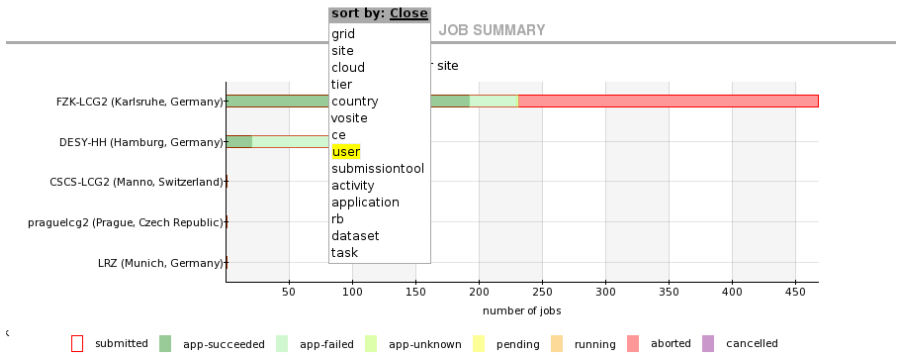

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

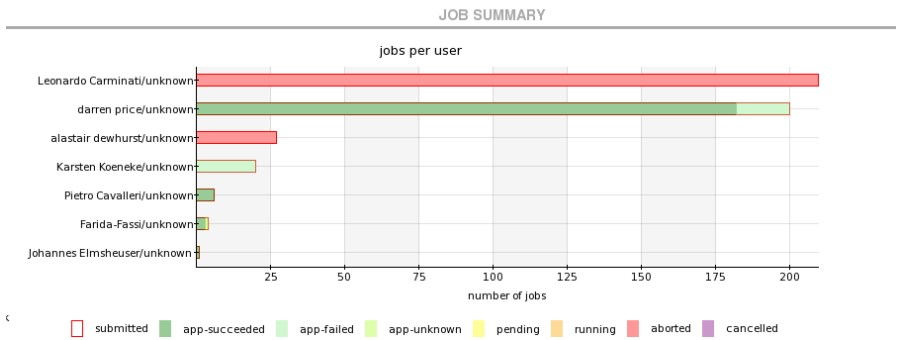

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

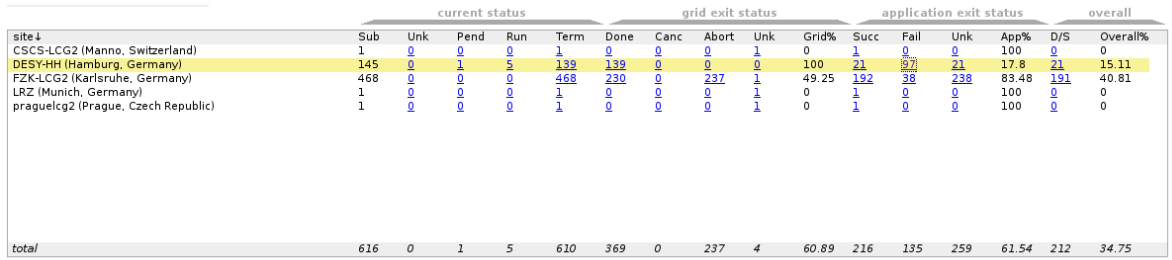

	current status					grid exit status					application exit status				overall	
site↓	Sub	Unk	Pend	Run	Term	Done	Canc	Abort	Unk	Grid%	Succ	Fail	Unk	App%	D/S	Overall%
CSCS-LCG2 (Manno, Switzerland)	1	0	0	0	1	0	0	0	1	0	1	0	0	100	0	0
DESY-HH (Hamburg, Germany)	145	0	1	5	139	139	0	0	0	100	21	97	21	17.8	21	15.11
FZK-LCG2 (Karlsruhe, Germany)	468	0	0	0	468	230	0	237	1	49.25	192	38	238	83.48	191	40.81
LRZ (Munich, Germany)	1	0	0	0	1	0	0	0	1	0	1	0	0	100	0	0
praguelcg2 (Prague, Czech Republic)	1	0	0	0	1	0	0	0	1	0	1	0	0	100	0	0
total	*616*	*0*	*1*	*5*	*610*	*369*	*0*	*237*	*4*	*60.89*	*216*	*135*	*259*	*61.54*	*212*	*34.75*

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the “vlemed” VO.

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

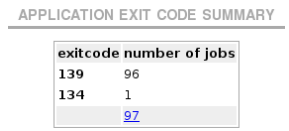
APPLICATION EXIT CODE SUMMARY

exitcode	number of jobs
139	96
134	1
	97

- grid info taken from RGMA, GridPP XML files, LCG BDII,
- VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
- installed for all four main experiments, plus the "vlemed" VO.

dashboard

grid job monitoring

- show what the jobs are doing on the Grid from the *VO point of view*:

  - grid info taken from RGMA, GridPP XML files, LCG BDII,
  - VO info: sent by jobs (using Monalisa, or read bookkeeping DB,
  - installed for all four main experiments, plus the "vlemed" VO.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

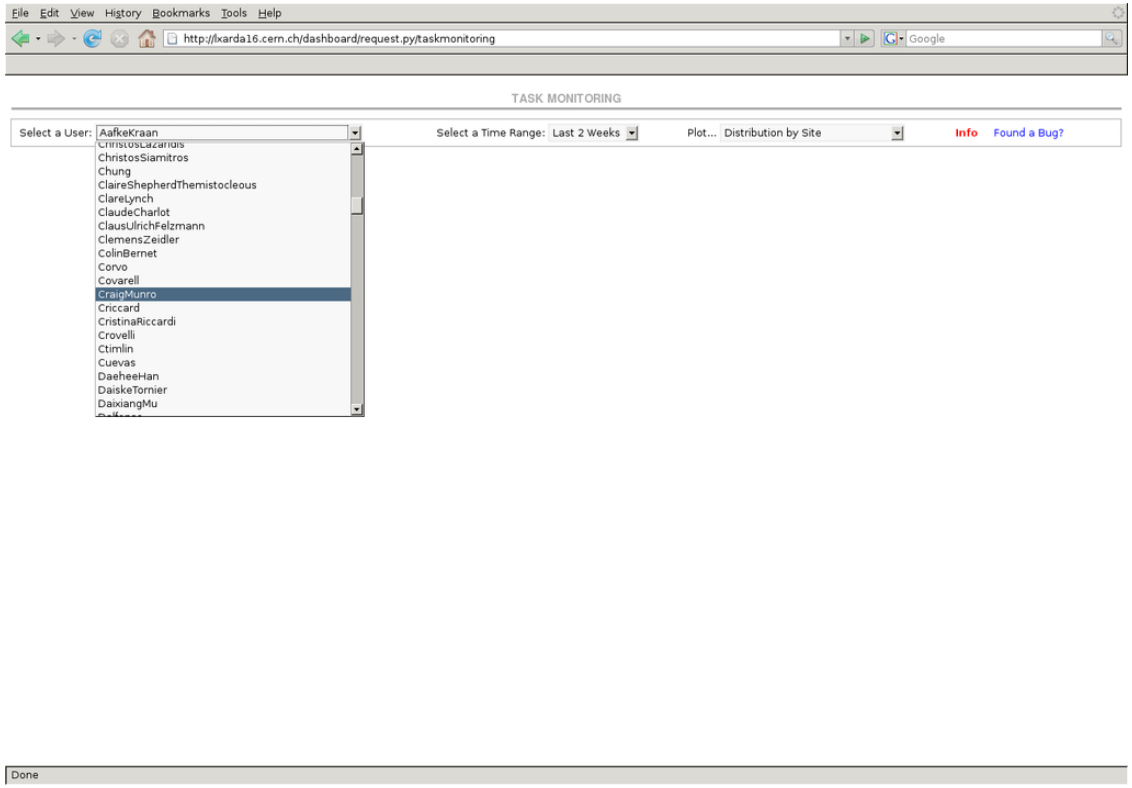

- show “my tasks”,
- progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

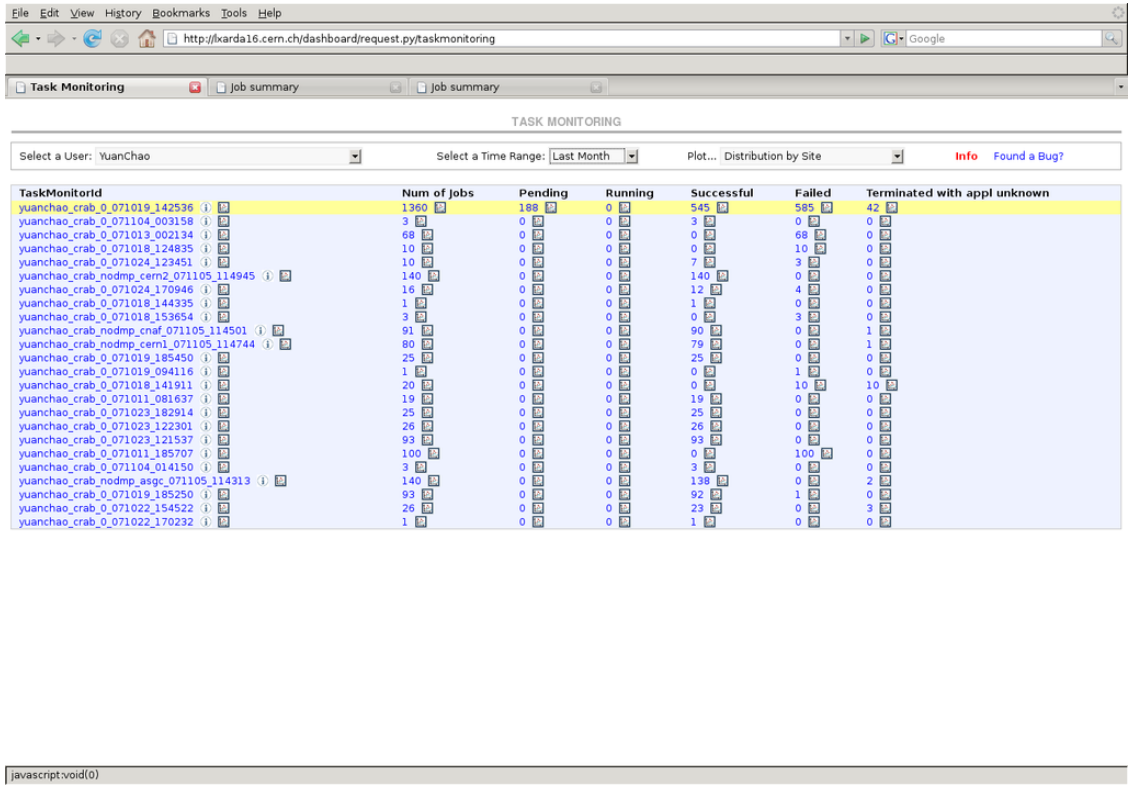

  - show “my tasks”,
  - progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

Pending	Running	Successful	Failed	Terminated with appl
188	0	545	585	42
0	0	3	0	0
0	0	0	68	0

- show “my tasks”,
- progress of each, investigate the reasons of the failures.

dashboard

grid job monitoring: user tasks

- view of the same data from a user running analysis:

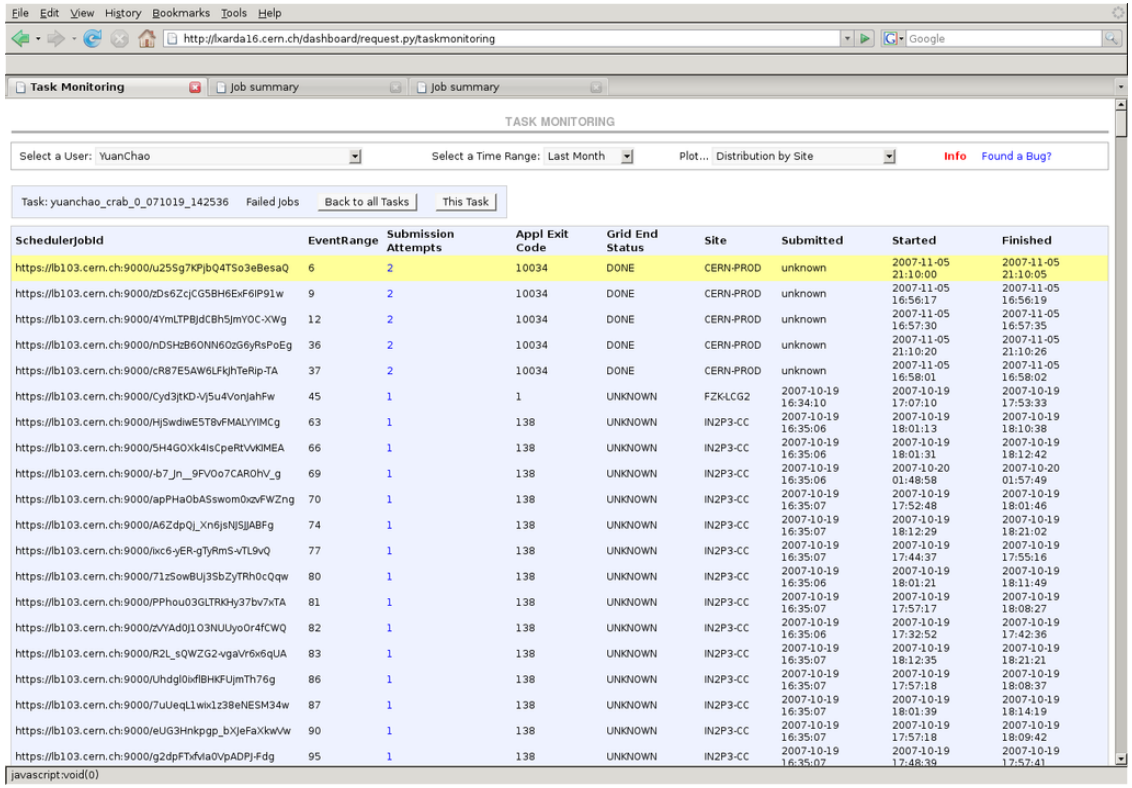

  - show “my tasks”,
  - progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

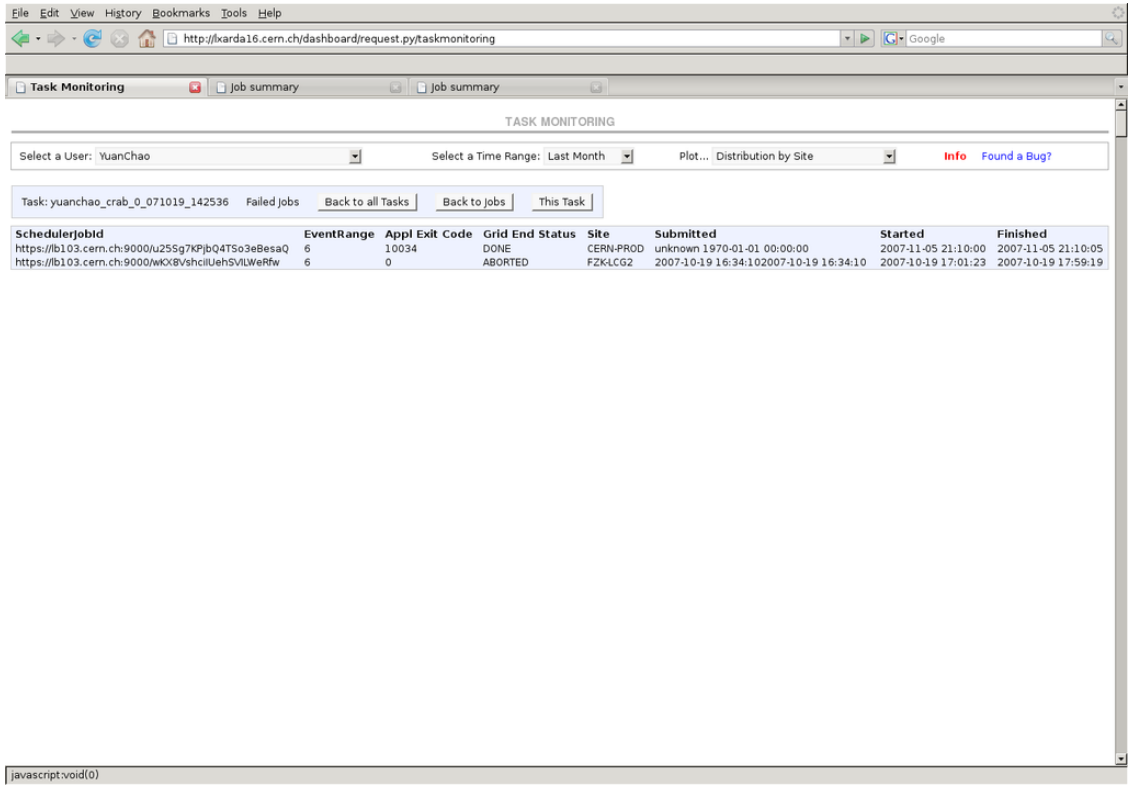

  - show “my tasks”,
  - progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

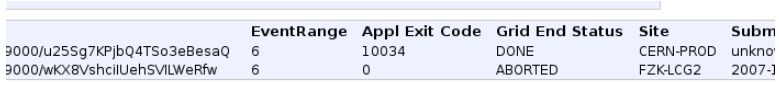

	EventRange	Appl Exit Code	Grid End Status	Site	Subm
9000/u25Sg7KPjbQ4TSo3eBesaQ	6	10034	DONE	CERN-PROD	unkno
9000/wKX8VshciIUehSVILWeRfw	6	0	ABORTED	FZK-LCG2	2007-1

- show “my tasks”,
- progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

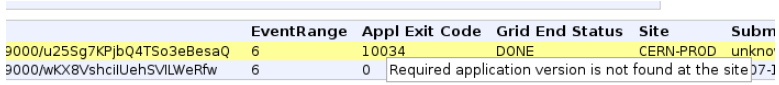

	EventRange	Appl Exit Code	Grid End Status	Site	Subm
9000/u25Sg7KPjbQ4TSo3eBesaQ	6	10034	DONE	CERN-PROD	unkno
9000/wKX8VshciIUehSVILWeRfw	6	0	Required application version is not found at the site		07-J

- show “my tasks”,
- progress of each, investigate the reasons of the failures.

dashboard

grid job monitoring: user tasks

- view of the same data from a user running analysis:

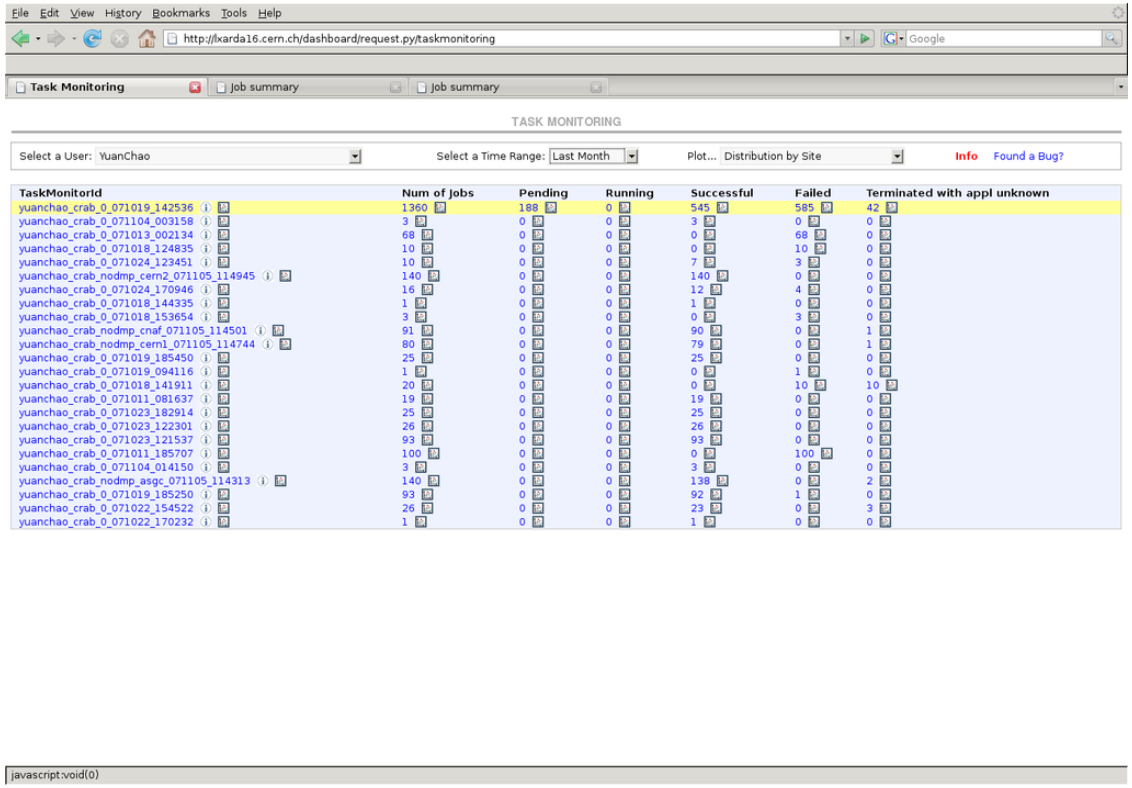

  - show “my tasks”,
  - progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

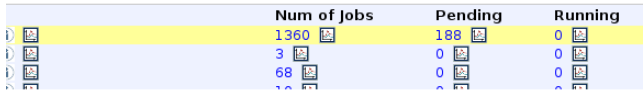

	Num of Jobs	Pending	Running
	1360	188	0
	3	0	0
	68	0	0

- show “my tasks”,
- progress of each, investigate the reasons of the failures.

dashboard

grid job monitoring: user tasks

- view of the same data from a user running analysis:

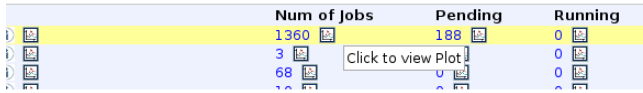

- show “my tasks”,
- progress of each, investigate the reasons of the failures.

grid job monitoring: user tasks

- view of the same data from a user running analysis:

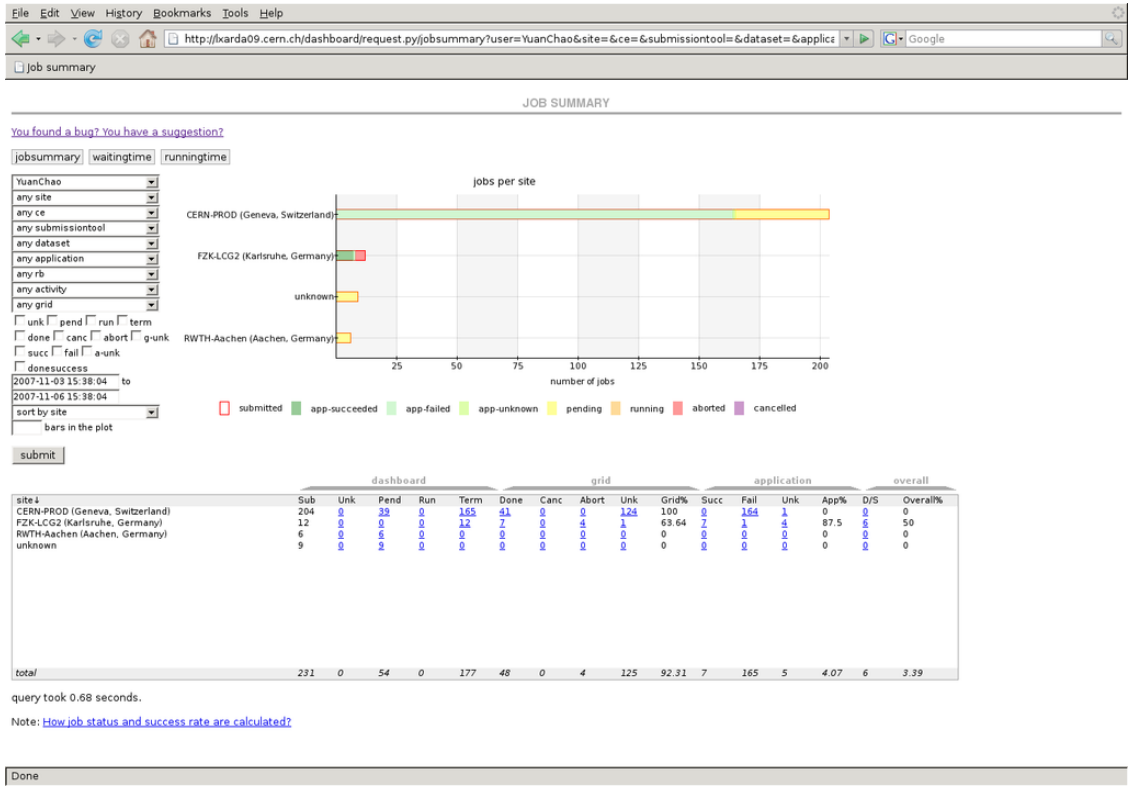

  - show “my tasks”,
  - progress of each, investigate the reasons of the failures.

grid job monitoring: summaries

- permits to show precompute summaries (faster queries),

  - for application exit codes, VO activities, etc.
  - same information as before.

grid job monitoring: summaries

- permits to show precompute summaries (faster queries),

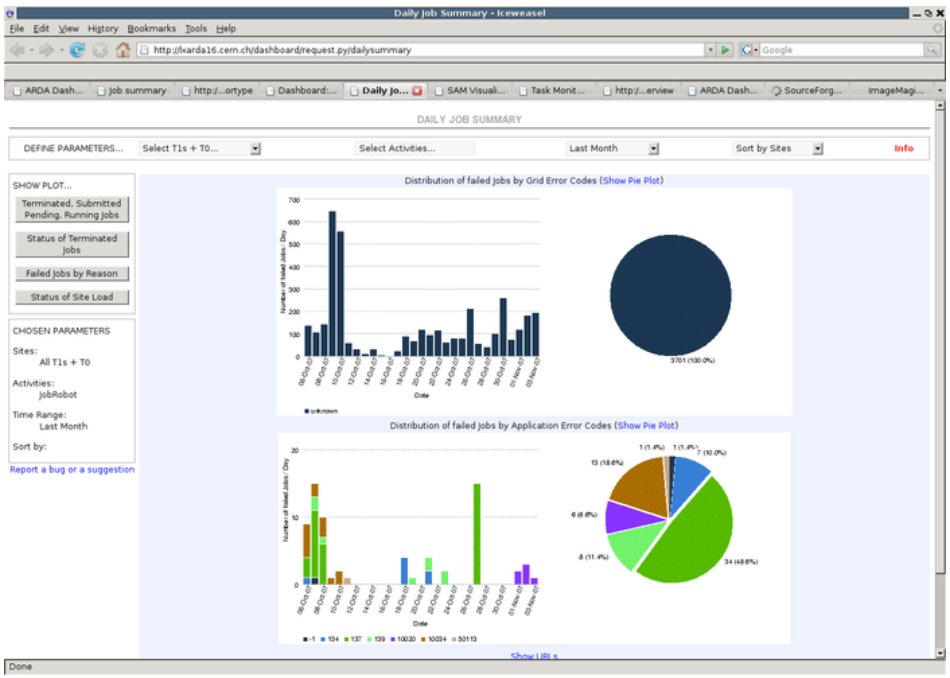

  - for application exit codes, VO activities, etc.
  - same information as before.

grid site reliability

- uses RGMA and IC-XML info for computing efficiency of sites,
- aware of *middleware resubmissions* to sites A, B and C,

- computation of daily rates, possibility to dig in the information.

grid site reliability

- uses RGMA and IC-XML info for computing efficiency of sites,
- aware of *middleware resubmissions* to sites A, B and C,

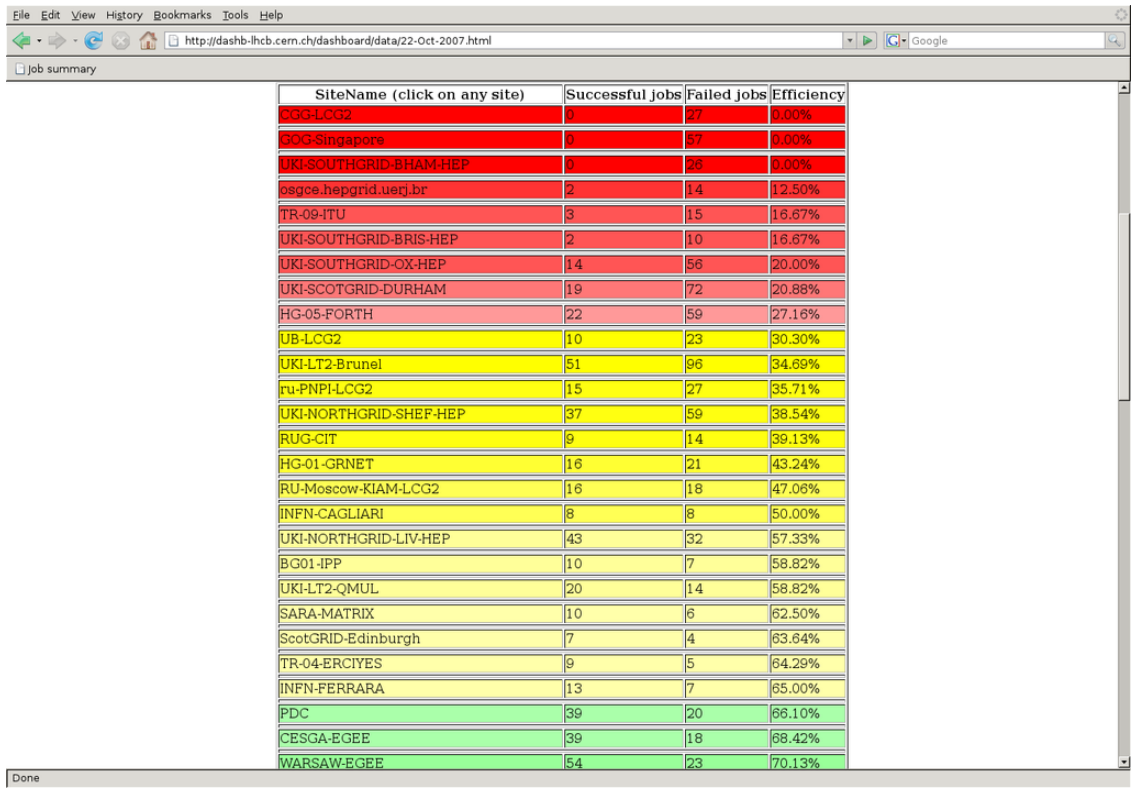

SiteName (click on any site)	Successful jobs	Failed jobs	Efficiency
CGG-LCG2	0	27	0.00%
GOG-Singapore	0	57	0.00%
UKI-SOUTHGRID-BHAM-HEP	0	26	0.00%
osgce.hepgrid.uerj.br	2	14	12.50%
TR-09-ITU	3	15	16.67%
UKI-SOUTHGRID-BRIS-HEP	2	10	16.67%
UKI-SOUTHGRID-OX-HEP	14	56	20.00%
UKI-SCOTGRID-DURHAM	19	72	20.88%
HG-05-FORTH	22	59	27.16%
UB-LCG2	10	23	30.30%
UKI-LT2-Brunel	51	96	34.69%
ru-PNPI-LCG2	15	27	35.71%
UKI-NORTHGRID-SHEF-HEP	37	59	38.54%
RUG-CIT	9	14	39.13%
HG-01-GRNET	16	21	43.24%
RU-Moscow-KIAM-LCG2	16	18	47.06%
INFN-CAGLIARI	8	8	50.00%
UKI-NORTHGRID-LIV-HEP	43	32	57.33%
BG01-IPP	10	7	58.82%
UKI-LT2-QMUL	20	14	58.82%
SARA-MATRIX	10	6	62.50%
ScotGRID-Edinburgh	7	4	63.64%
TR-04-ERCIYES	9	5	64.29%
INFN-FERRARA	13	7	65.00%
PDC	39	20	66.10%
CESGA-EGEE	39	18	68.42%
WARSAW-EGEE	54	23	70.13%

- computation of daily rates, possibility to dig in the information.

dashboard

grid site reliability

- uses RGMA and IC-XML info for computing efficiency of sites,
- aware of *middleware resubmissions* to sites A, B and C,

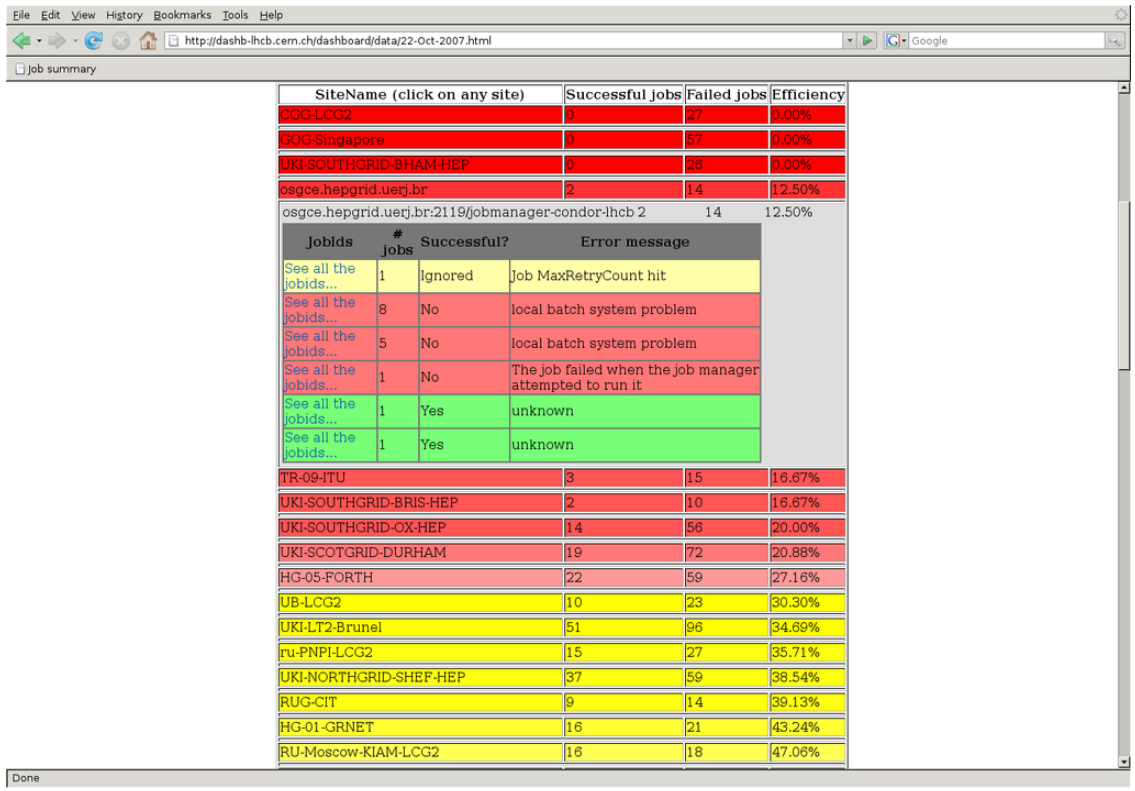

SiteName (click on any site)	Successful jobs	Failed jobs	Efficiency
CGG-LCG2	0	27	0.00%
GOG-Singapore	0	57	0.00%
UKI-SOUTHGRID-BHAM-HEP	0	26	0.00%
osgce.hepgrid.uerj.br	2	14	12.50%
osgce.hepgrid.uerj.br:2119/jobmanager-condor-lhcb	2	14	12.50%

JobIds	# jobs	Successful?	Error message
See all the jobids...	1	Ignored	Job MaxRetryCount hit
See all the jobids...	8	No	local batch system problem
See all the jobids...	5	No	local batch system problem
See all the jobids...	1	No	The job failed when the job manager attempted to run it
See all the jobids...	1	Yes	unknown
See all the jobids...	1	Yes	unknown

TR-09-ITU	3	15	16.67%
UKI-SOUTHGRID-BRIS-HEP	2	10	16.67%
UKI-SOUTHGRID-OX-HEP	14	56	20.00%
UKI-SCOTGRID-DURHAM	19	72	20.88%
HG-05-FORTH	22	59	27.16%
UB-LCG2	10	23	30.30%
UKI-LT2-Brunel	51	96	34.69%
ru-PNPI-LCG2	15	27	35.71%
UKI-NORTHGRID-SHEF-HEP	37	59	38.54%
RUG-CIT	9	14	39.13%
HG-01-GRNET	16	21	43.24%
RU-Moscow-KIAM-LCG2	16	18	47.06%

- computation of daily rates, possibility to dig in the information.

dashboard

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

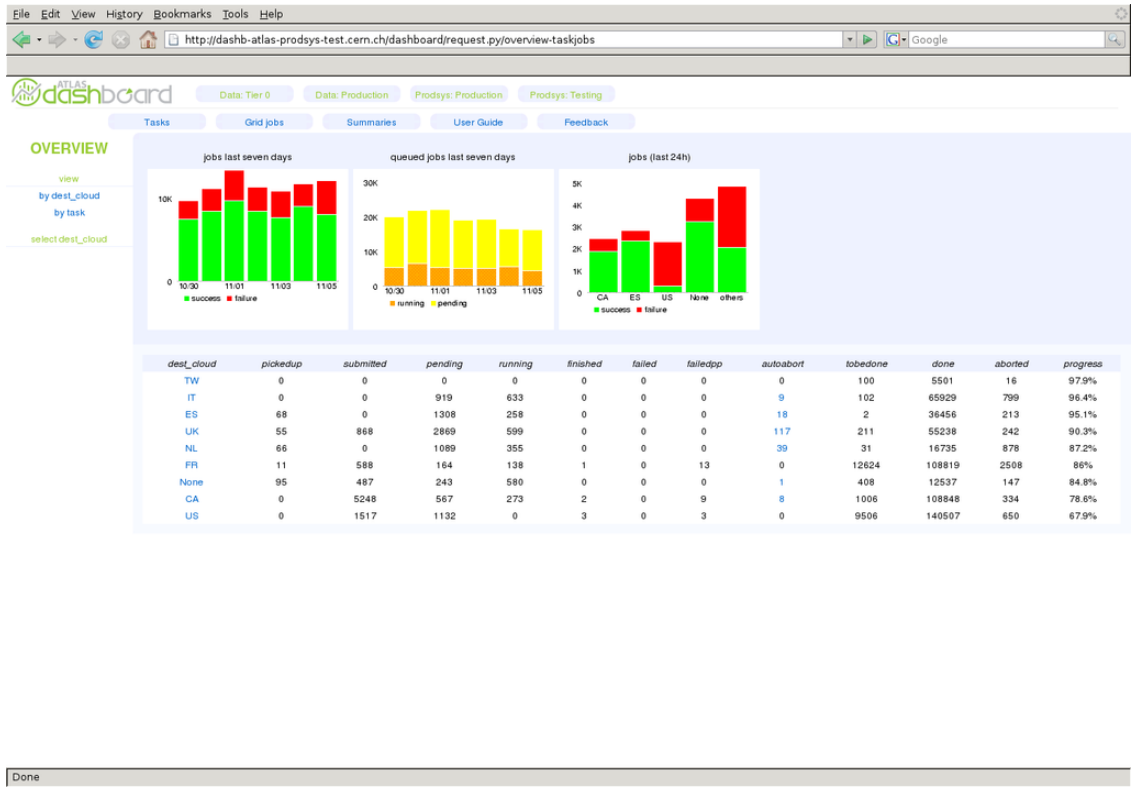

dest_cloud	pickedup	submitted	pending	running	finished	failed	failedpp	autoabort	tobedone	done	aborted	progress
TW	0	0	0	0	0	0	0	0	100	5501	16	97.9%
IT	0	0	919	633	0	0	0	9	102	65929	799	96.4%
ES	68	0	1308	258	0	0	0	18	2	36456	213	95.1%
UK	55	868	2869	599	0	0	0	117	211	55238	242	90.3%
NL	66	0	1089	355	0	0	0	39	31	16735	878	87.2%
FR	11	588	164	138	1	0	13	0	12624	108819	2508	86%
None	95	487	243	580	0	0	0	1	408	12537	147	84.8%
CA	0	5248	567	273	2	0	9	8	1006	108848	334	78.6%
US	0	1517	1132	0	3	0	3	0	9506	140507	650	67.9%

- under “review” by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

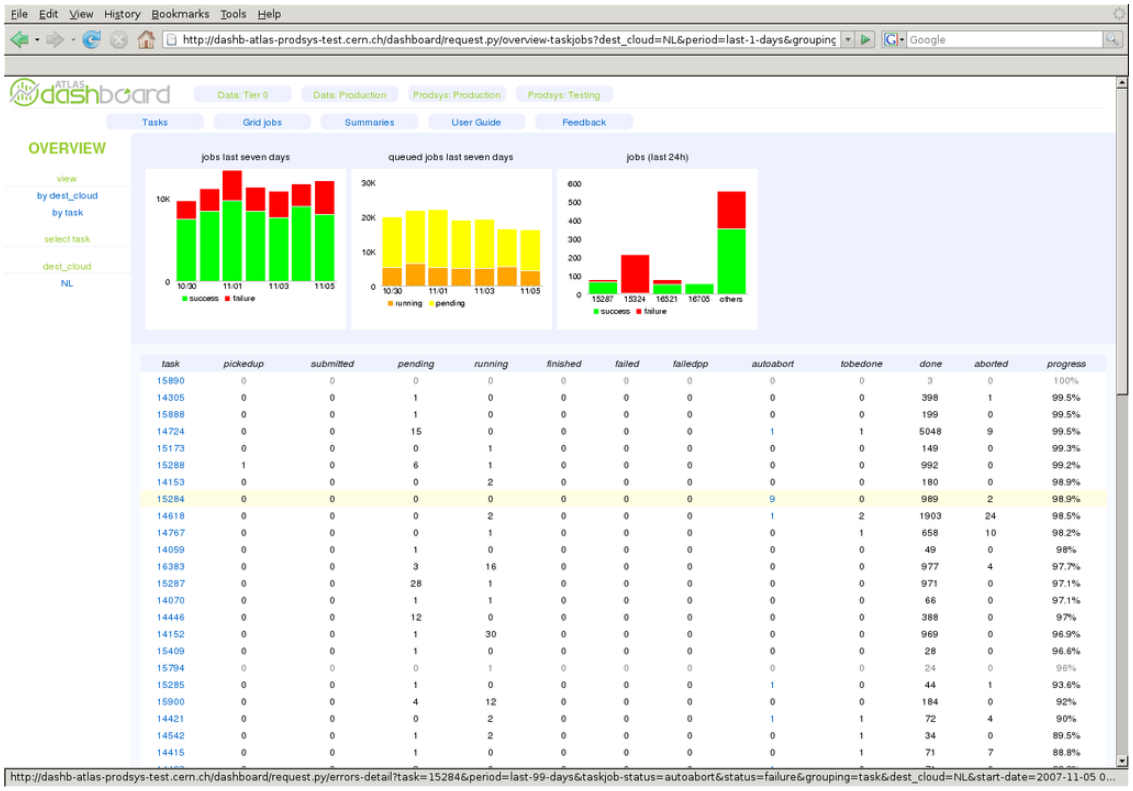

- under “review” by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

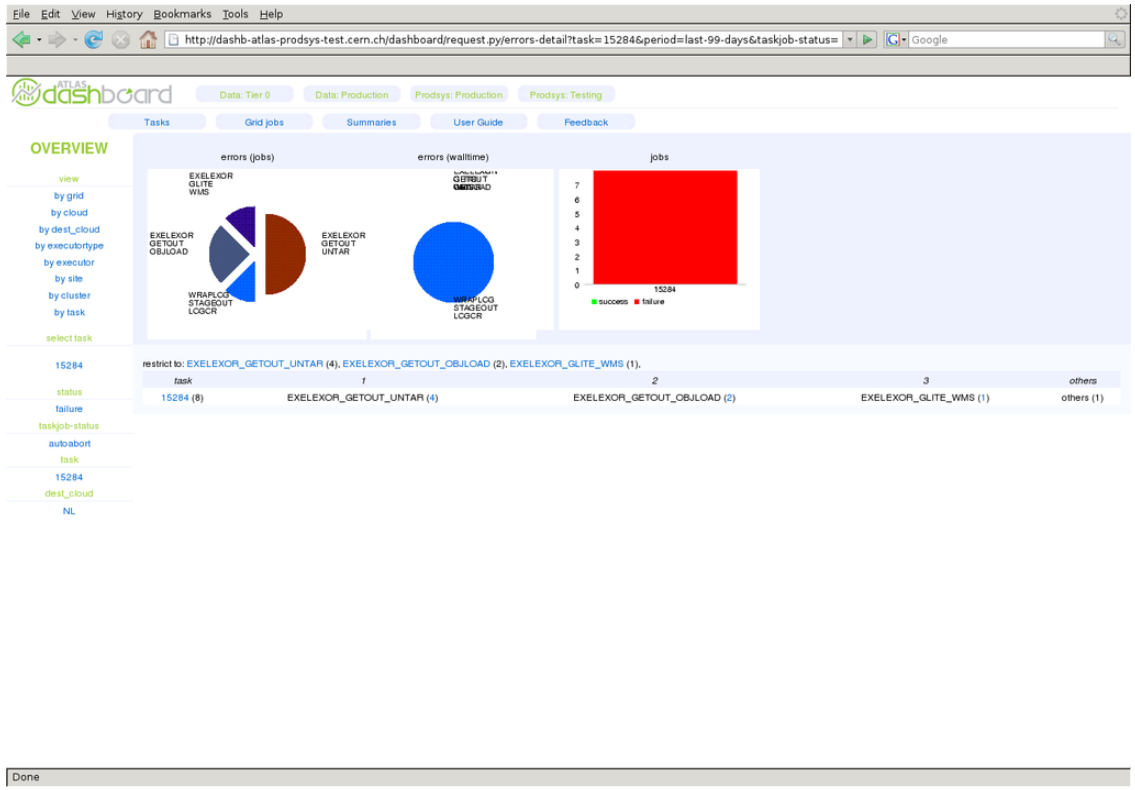

- under “review” by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

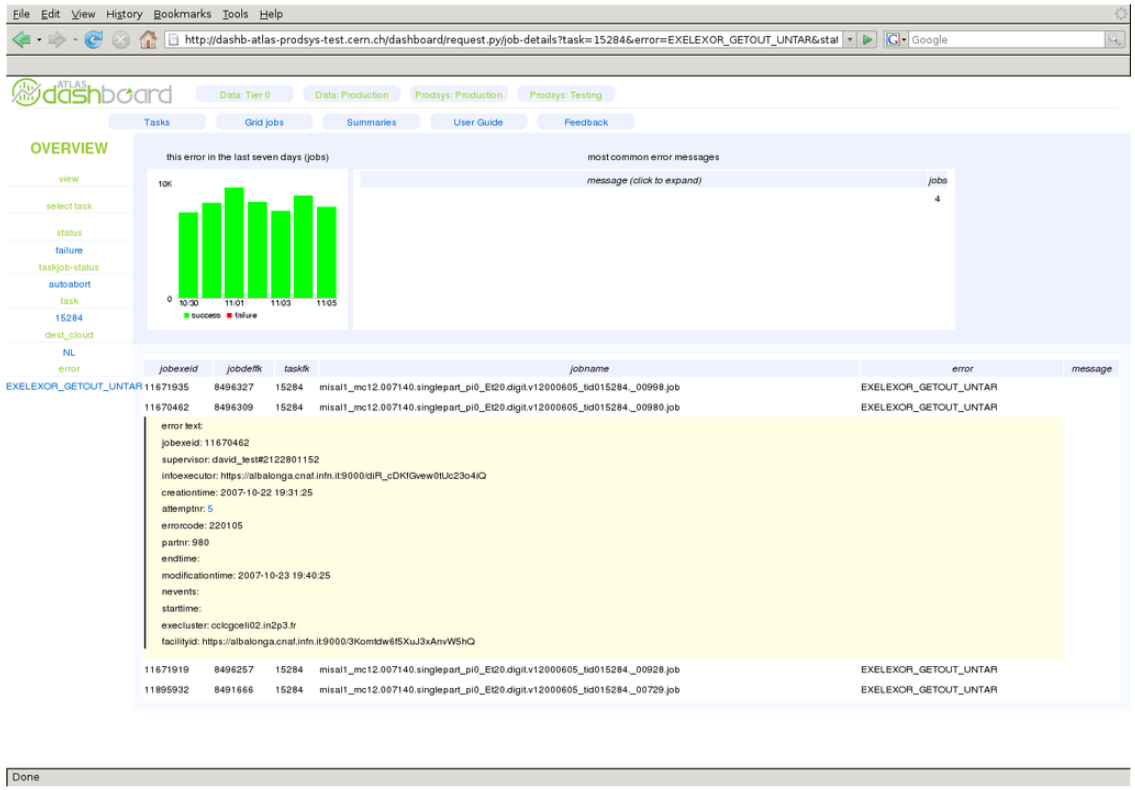

- under “review” by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

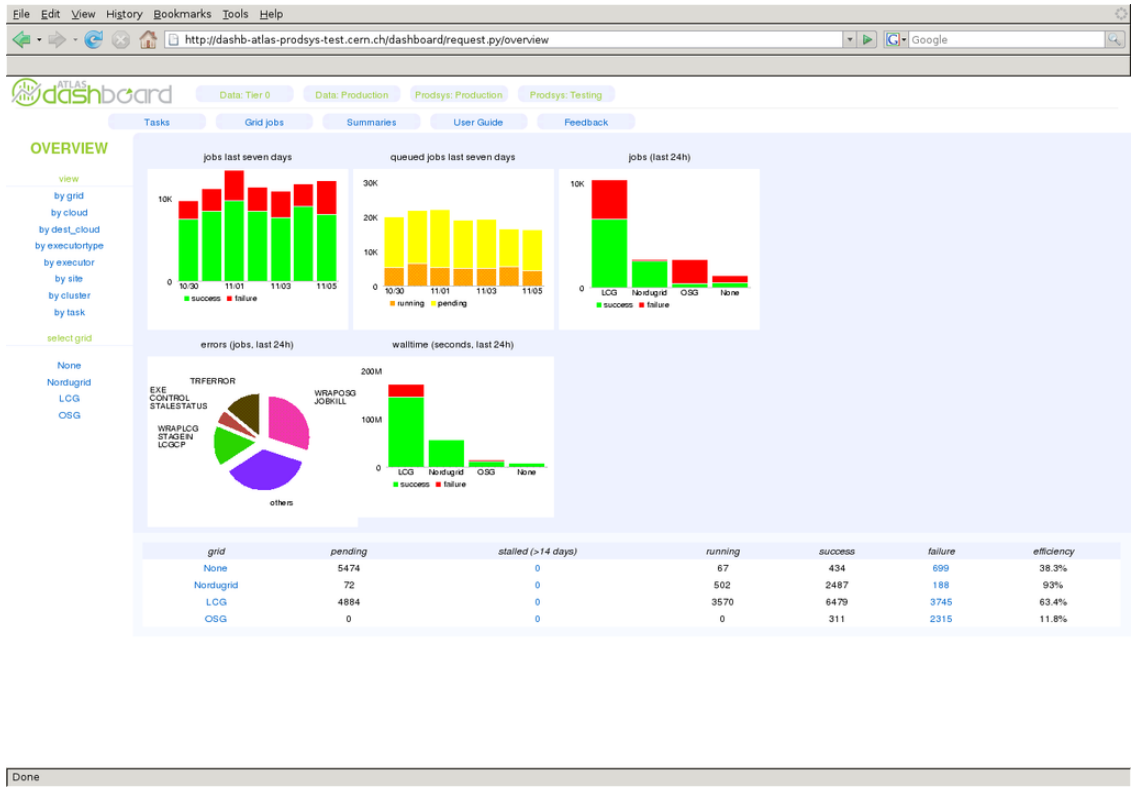

- under “review” by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: "shifters",

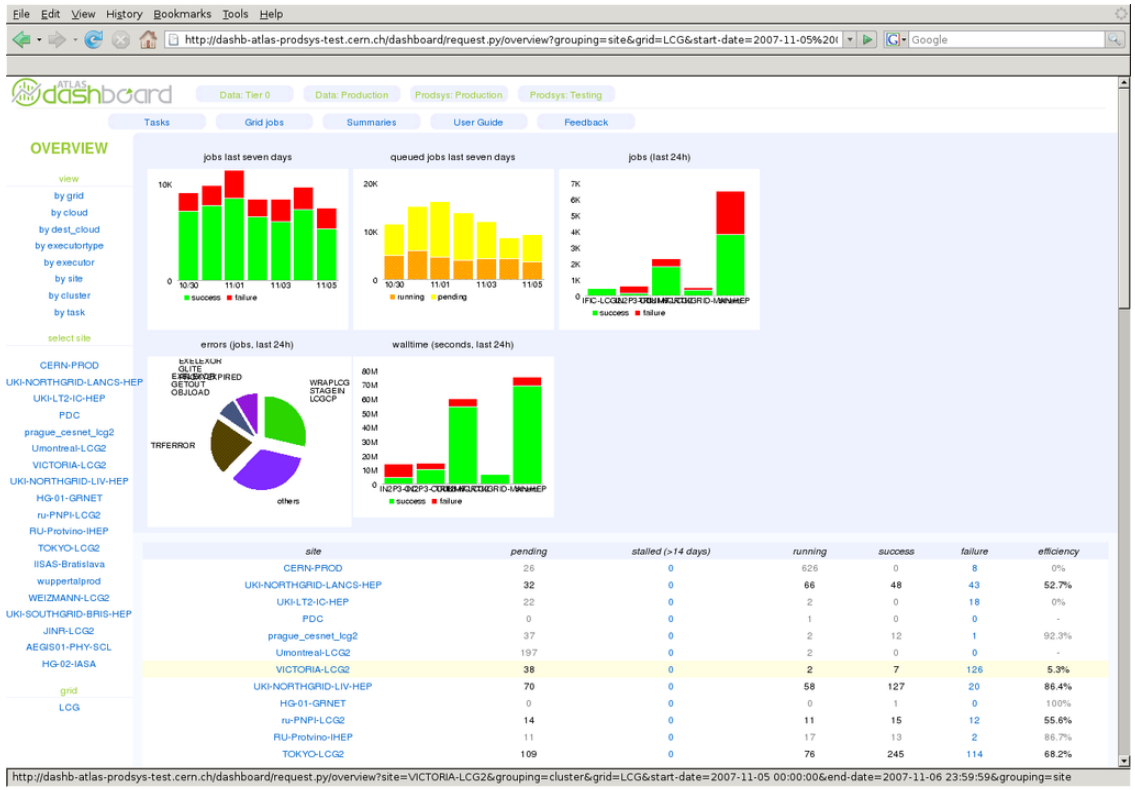

- under "review" by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

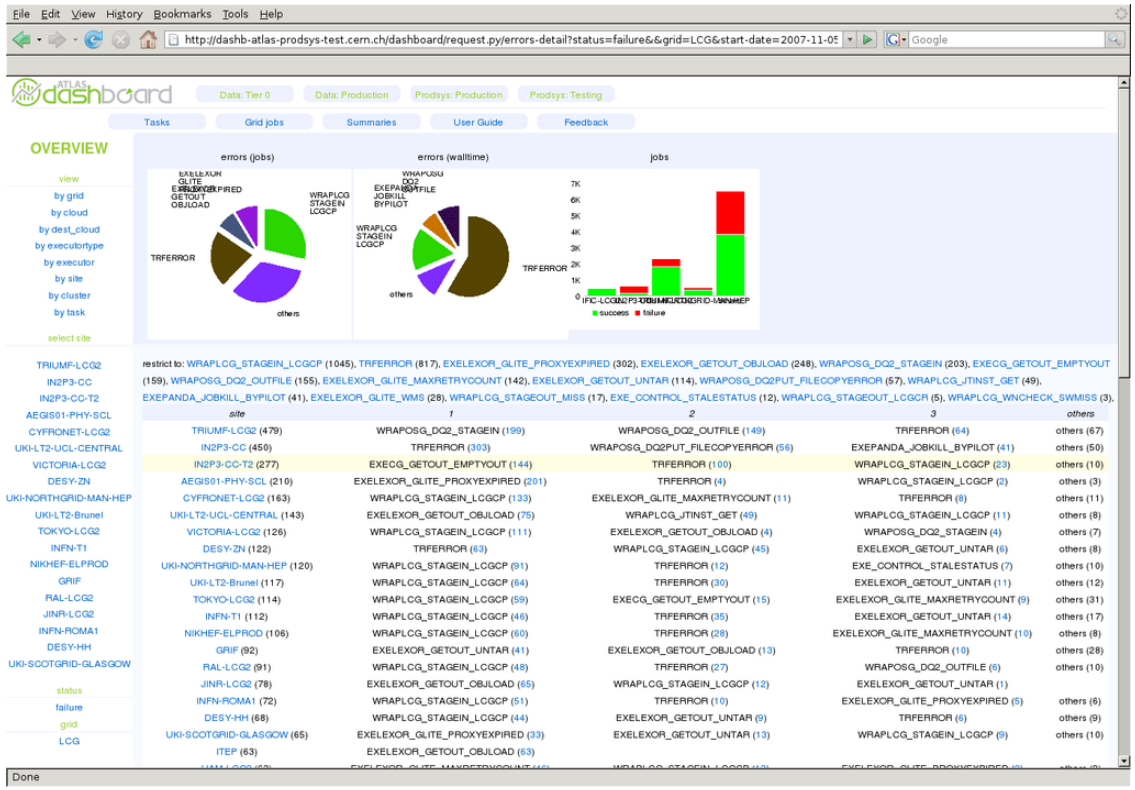

- under “review” by master-shifters.

grid job monitoring: ATLAS production

- monitoring offloaded to the dashboard,
- target: “shifters”,

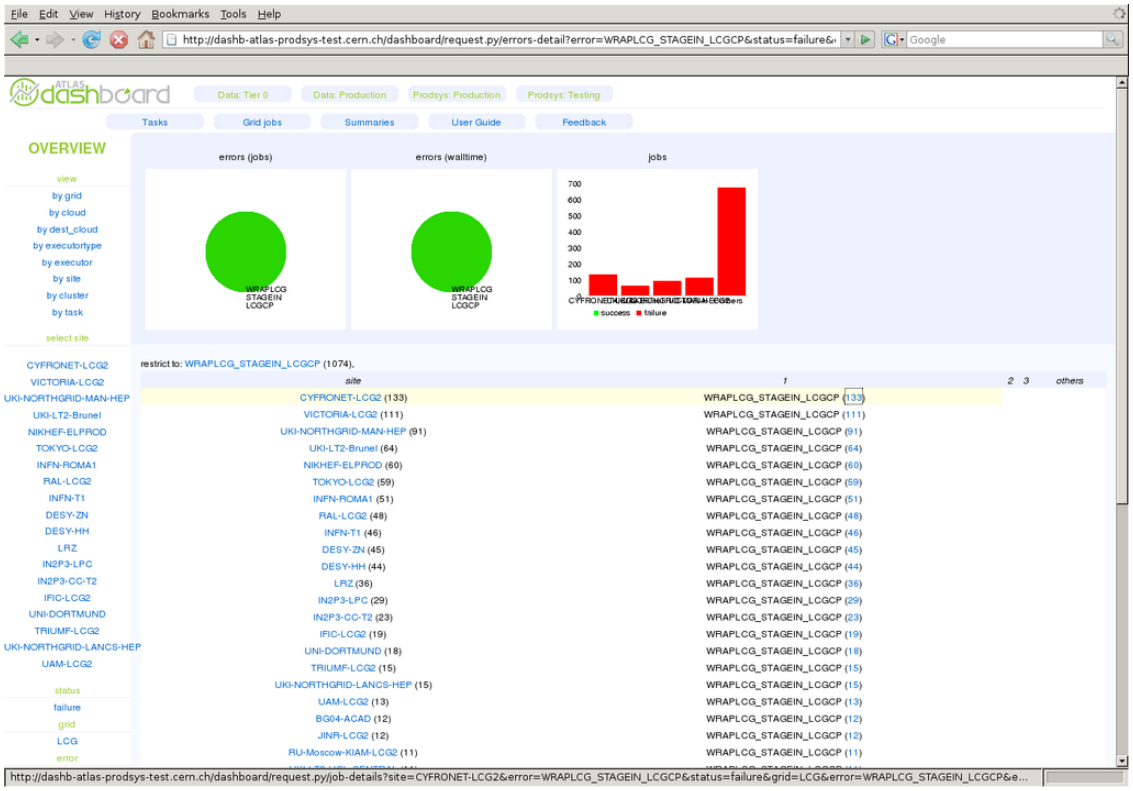

- under “review” by master-shifters.

grid job monitoring: CMS production

- close collaboration between CMS and the dashboard,
- the CMS production agents publish their status in the dashboard (HTTP POST) in real time,
- dashboard services are computing periodic statistics,
- the *view* is *not* implemented with the dashboard:
  - info is retrieved using the dashboard interface,
  - CMS will make their own interface on top.

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

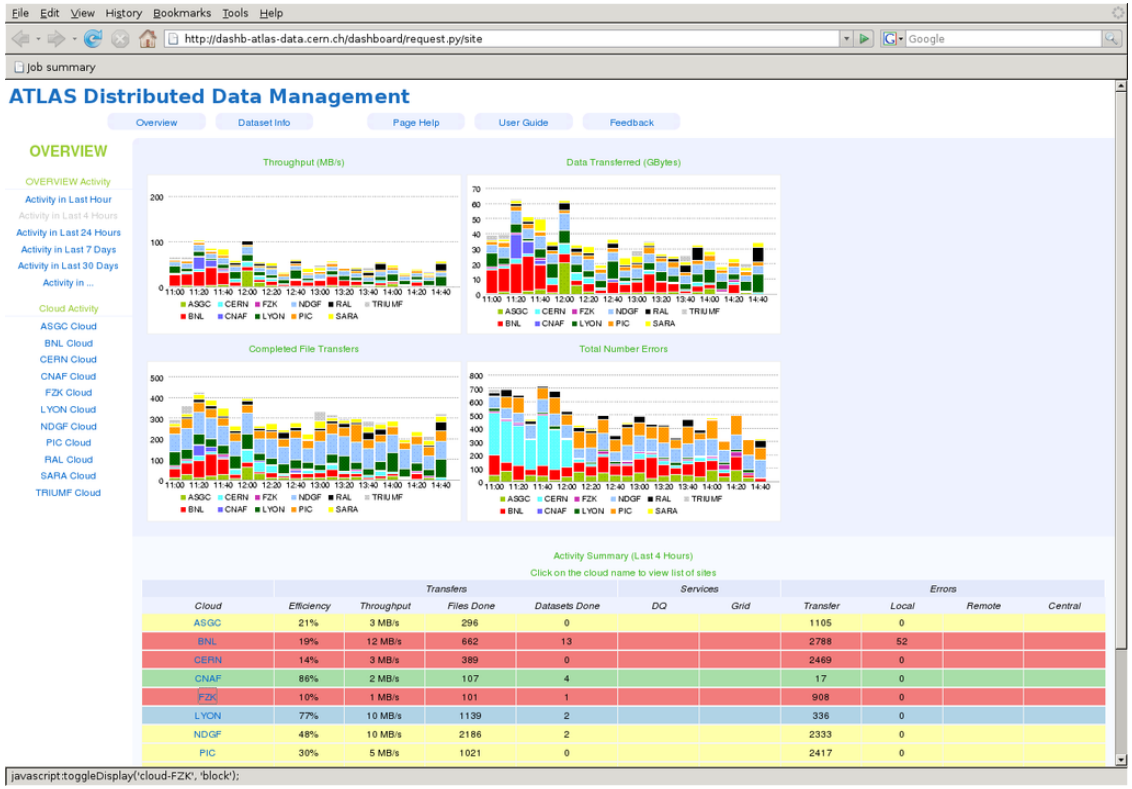

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

Activity Summary (Last 4 Hours)

Click on the cloud name to view list of sites

	Transfers				Services		Errors			
Cloud	Efficiency	Throughput	Files Done	Datasets Done	DQ	Grid	Transfer	Local	Remote	Central
ASGC	21%	3 MB/s	296	0			1105	0		
BNL	19%	12 MB/s	662	13			2788	52		
CERN	14%	3 MB/s	389	0			2469	0		
CNAF	86%	2 MB/s	107	4			17	0		
FZK	10%	1 MB/s	101	1			908	0		
LYON	77%	10 MB/s	1139	2			336	0		
NDGF	48%	10 MB/s	2186	2			2333	0		
PIC	30%	5 MB/s	1021	0			2417	0		
RAL	36%	4 MB/s	345	1			621	0		
SARA	66%	5 MB/s	530	1			274	4		
TRIUMF	64%	1 MB/s	224	9			126	0		

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

dashboard

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

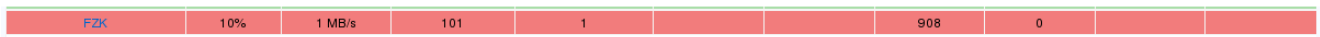

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

FZK	10%	1 MB/s	101	1			908	0		
Click on the site name to go to the site page										
FZKDISK	0%	0 MB/s	0	0	OK		56	0		
FZKTAPE	0%	0 MB/s	0	0	OK		0	0		

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

FZK	10%	1 MB/s	101	1			908	0		
Click on the site name to go to the site page										
FZKDISK	0%	0 MB/s	0	0	OK		56	0		
State from FTS: Failed; Retries: 1; Reason: DESTINATION error during PREPARATION phase: [CONNECTION] failed to contact on remote SRM [httpg://gridka-dcache.fzk.de:8443/srm/managerv1]. Givin' up after 3 tries										56
FZKTAPE	0%	0 MB/s	0	0	OK		0	0		

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

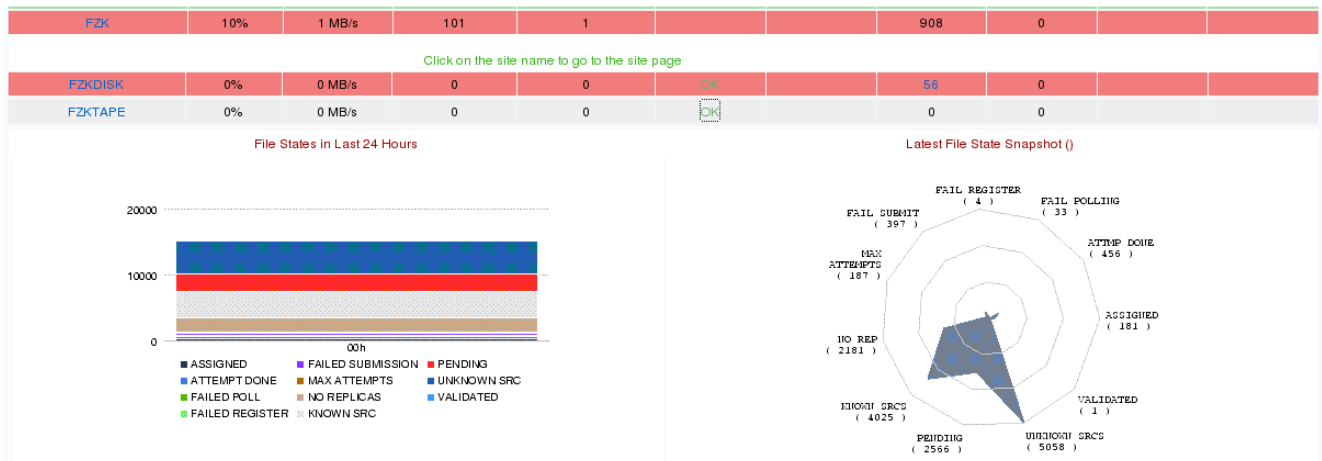

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

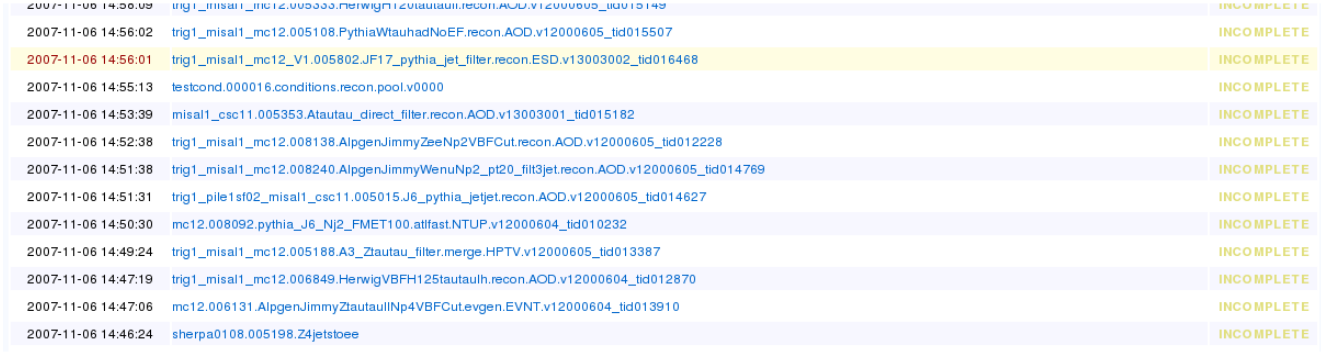

2007-11-06 14:56:02	trig1_misal1_mc12.005108.PythiaWtauhadNoEF.recon.AOD.v12000605_tid015507	INCOMPLETE
2007-11-06 14:56:01	trig1_misal1_mc12_V1.005802.JF17_pythia_jet_filter.recon.ESD.v13003002_tid016468	INCOMPLETE
2007-11-06 14:55:13	testcond.000016.conditions.recon.pool.v0000	INCOMPLETE
2007-11-06 14:53:39	misal1_csc11.005353.Atautau_direct_filter.recon.AOD.v13003001_tid015182	INCOMPLETE
2007-11-06 14:52:38	trig1_misal1_mc12.008138.AlpgenJimmyZeeNp2VBFCut.recon.AOD.v12000605_tid012228	INCOMPLETE
2007-11-06 14:51:38	trig1_misal1_mc12.008240.AlpgenJimmyWenuNp2_pt20_filt3jet.recon.AOD.v12000605_tid014769	INCOMPLETE
2007-11-06 14:51:31	trig1_pile1sf02_misal1_csc11.005015.J6_pythia_jetjet.recon.AOD.v12000605_tid014627	INCOMPLETE
2007-11-06 14:50:30	mc12.008092.pythia_J6_Nj2_FMET100.atlfast.NTUP.v12000604_tid010232	INCOMPLETE
2007-11-06 14:49:24	trig1_misal1_mc12.005188.A3_Ztautau_filter.merge.HPTV.v12000605_tid013387	INCOMPLETE
2007-11-06 14:47:19	trig1_misal1_mc12.006849.HerwigVBFH125tautaulh.recon.AOD.v12000604_tid012870	INCOMPLETE
2007-11-06 14:47:06	mc12.006131.AlpgenJimmyZtautaulNp4VBFCut.evgen.EVNT.v12000604_tid013910	INCOMPLETE
2007-11-06 14:46:24	sherpa0108.005198.Z4jetstoee	INCOMPLETE

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

2007-11-06 14:56:02	trig1_misal1_mc12.005108.PythiaWtauhadNoEF.recon.AOD.v12000605_tid015507	INCOMPLETE
2007-11-06 14:56:01	trig1_misal1_mc12_V1.005802.JF17_pythia_jet_filter.recon.ESD.v13003002_tid016468	INCOMPLETE
DUID:	8e958bf0-838f-11dc-b26c-00034731f394 VUID: 8e958fd9-838f-11dc-bc3a-00034731f394	
CREATION:	2007-10-26 16:35:27 VERSION: 1	
2007-11-06 14:55:13	testcond.000016.conditions.recon.pool.v0000	INCOMPLETE
2007-11-06 14:53:39	misal1_csc11.005353.Atautau_direct_filter.recon.AOD.v13003001_tid015182	INCOMPLETE
2007-11-06 14:52:38	trig1_misal1_mc12.008138.AlpgenJimmyZeeNp2VBFCut.recon.AOD.v12000605_tid012228	INCOMPLETE
2007-11-06 14:51:38	trig1_misal1_mc12.008240.AlpgenJimmyWenuNp2_pt20_filt3jet.recon.AOD.v12000605_tid014769	INCOMPLETE
2007-11-06 14:51:31	trig1_pile1sf02_misal1_csc11.005015.J6_pythia_jetjet.recon.AOD.v12000605_tid014627	INCOMPLETE
2007-11-06 14:50:30	mc12.008092.pythia_J6_Nj2_FMET100.atlfast.NTUP.v12000604_tid010232	INCOMPLETE
2007-11-06 14:49:24	trig1_misal1_mc12.005188.A3_Ztautau_filter.merge.HPTV.v12000605_tid013387	INCOMPLETE
2007-11-06 14:47:19	trig1_misal1_mc12.006849.HerwigVBFH125tautaulh.recon.AOD.v12000604_tid012870	INCOMPLETE

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

2007-11-06 14:53:45	ESD.016468._06338.pool.root.2	PENDING
2007-11-06 14:53:45	ESD.016468._06341.pool.root.2	PENDING
2007-11-06 14:31:35	ESD.016468._06263.pool.root.4	PENDING
2007-11-06 14:31:35	ESD.016468._06075.pool.root.4	PENDING
2007-11-04 16:07:59	ESD.016468._06443.pool.root.2	FILE_DONE
2007-11-04 14:55:01	ESD.016468._06422.pool.root.2	FILE_DONE
2007-11-04 12:46:06	ESD.016468._06448.pool.root.2	FILE_DONE
2007-11-04 07:45:00	ESD.016468._06154.pool.root.3	FILE_DONE
2007-11-04 04:19:31	ESD.016468._06322.pool.root.2	FILE_DONE
2007-11-04 04:19:31	ESD.016468._06303.pool.root.4	FILE_DONE
2007-11-04 04:19:31	ESD.016468._06288.pool.root.3	FILE_DONE
2007-11-04 04:19:31	ESD.016468._06440.pool.root.2	FILE_DONE

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

grid data transfers monitoring

- activity conducted with ATLAS DDM (distributed data management),
- show the status and performance of the DDM system:

2007-11-06 14:53:45	ESD.016468._06338.pool.root.2	PENDING
2007-11-06 14:53:45	ESD.016468._06341.pool.root.2	PENDING
2007-11-06 14:31:35	ESD.016468._06263.pool.root.4	PENDING
2007-11-06 14:31:35	ESD.016468._06075.pool.root.4	PENDING
2007-11-04 16:07:59	ESD.016468._06443.pool.root.2	FILE_DONE
GUID:	A454B508-2389-DC11-937F-00E081401061 CREATION: 2007-11-04 14:38:30	
FILE SIZE:	1016093104 bytes	
SURL:	srm://gridka-dcache.fzk.de/pnfs/gridka.de/atlas/disk-only/trig1_misal1_mc12_V1/ESD/trig1_misal1_mc12_V1.005802.JF17_pyt hia_jet_filter.recon.ESD.v13003002_tid016468/ESD.016468._06443.pool.root.2__DQ2-1194187402	
CHECKSUM:	md5:56feb6125a0edbbf0491b42c79b56903	
2007-11-04 14:55:01	ESD.016468._06422.pool.root.2	FILE_DONE
2007-11-04 12:46:06	ESD.016468._06448.pool.root.2	FILE_DONE
2007-11-04 07:45:00	ESD.016468._06154.pool.root.3	FILE_DONE

- topology and names are specific to ATLAS, ATLAS datasets,
- info is published *directly* by the DDM servers, using a specific API (a dashboard API based on HTTP).

outline

conclusion & future plans

dashboard home page

- framework and operations:
  - installation, maintenance procedures,
  - systematic testing (unit tests).
- applications:
  - grid jobs: more info sources and support for pilot jobs,
  - integrate more our applications with themselves (!),
  - develop the integration with external tools (python API),
  - alert systems.

dashboard