

ATLAS Production System Monitoring

John Kennedy
LMU München
CHEP 07
Victoria BC



06/09/2007

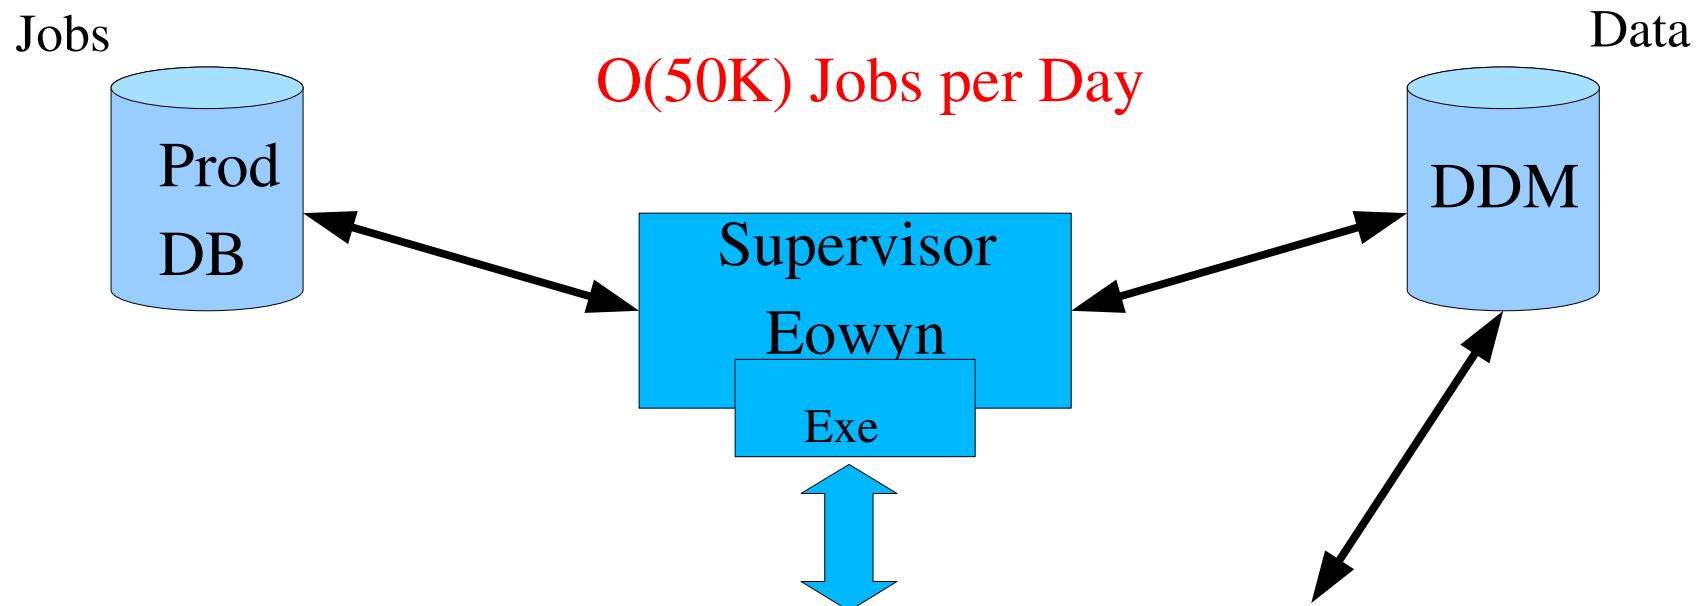
Overview

- Production System – quick intro
- Monitoring Goals
- Design
- Functionality - examples
- Usage
- Conclusion + outlook

Prod-sys

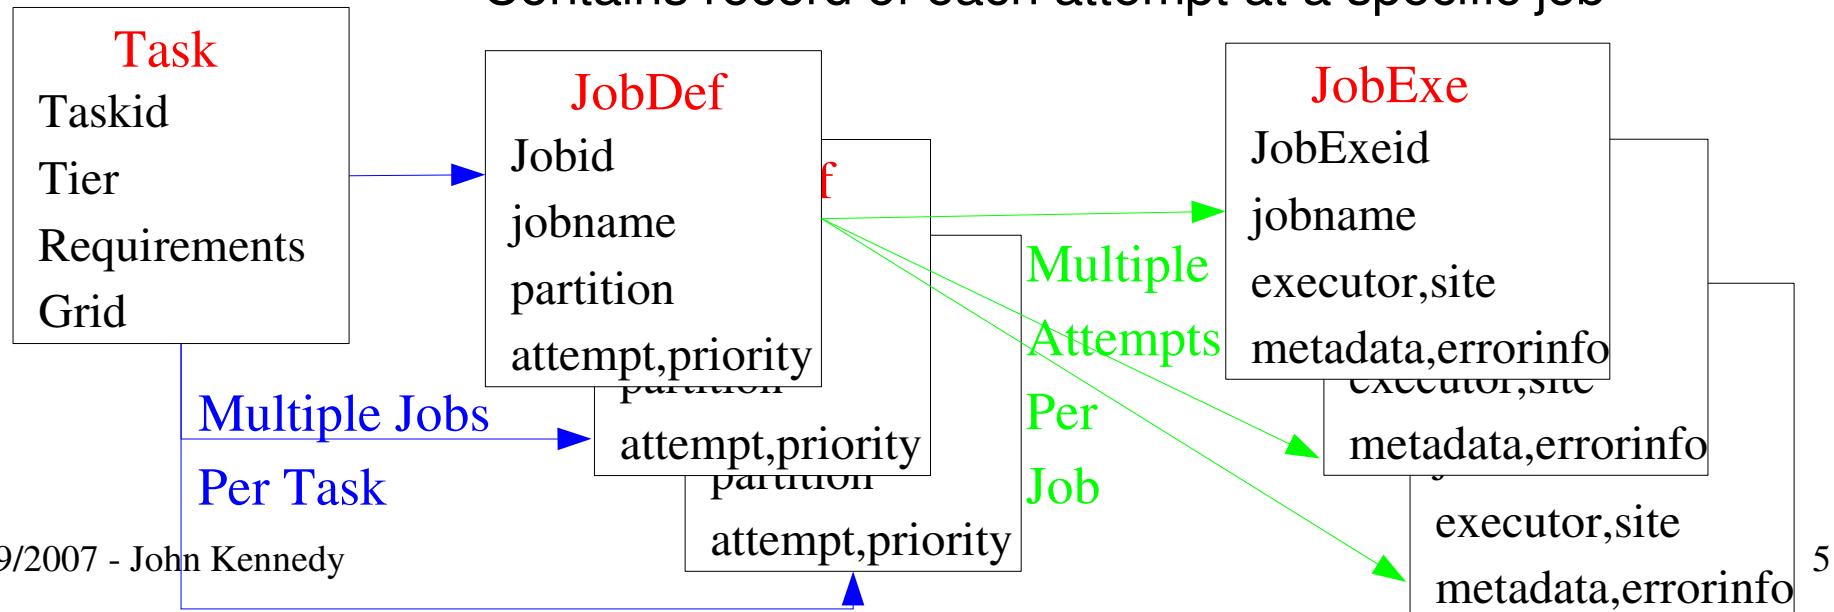
- Central Database with
 - Tasks – Groups of jobs
 - Jobs – job definitions
 - Attempts – attempt info for jobs
- Supervisor(Eowyn)
 - Retrieves jobs from DB
 - Passes to GRID through Executor
 - Manages job attempt info
- Executor(per grid flavour)
 - Create wrappers etc for submission
 - Submit job
 - Manage Job + interpret output (failure reasons etc)

Prod-sys



Production Database

- Hierarchy of tables
 - Task Table - 11,018
 - Contains task definition data
 - JobDef Table - 6,619,954
 - Contains definition of each job within a task
 - JobExe Table - 9,457,938
 - Contains record of each attempt at a specific job



Monitoring Goals

- Provide easy access to monitoring and accounting info
- **Avoid load on central database**
- Provide Views for
 - Recent jobs to allow problem spotting
 - Configurable time windows for # jobs/walltime and errors
 - Higher level views to give an understanding of job states and available jobs
- **Provide High and Low Level Views!**

Design

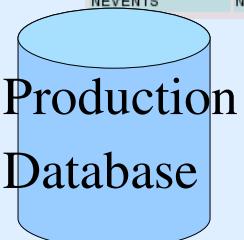
- Monitoring tables defined in specific monitoring account
- Oracle Stored procedures used to fill monitoring tables from Prodsys tables
- Monitoring tables used in combination with smaller proddb tables to provide web pages
- PHP pages used to interface to tables
- Cron jobs to give periodic views

PRODUCTION

Field	Type	Length
JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORFILE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXEQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
METADATA	VARCHAR	4000
CPUCOUNT	NUMBER	20
CPUUNIT	VARCHAR	128
TRANSEXITCODE	VARCHAR	128
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128
STARTEPOCH	NUMBER	12
ENDEPOCH	NUMBER	12
WHENACRONYM	VARCHAR	128
WHENCODE	NUMBER	9
IGNOREIDENTICAL	NUMBER	1
UPDATETIME	DATE	7
TRFCODE	NUMBER	9
TRFACRONYM	VARCHAR	128
STAGEIN	NUMBER	9
STAGEOUT	NUMBER	9
CPUMODEL	VARCHAR	256
NEVENTS	NUMBER	9

All Job Execution Records

10M



Design

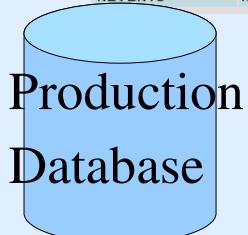
PRODUCTION

Field	Type	Length
JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
METADATA	VARCHAR	4000
CPUCOUNT	NUMBER	20
CPUUNIT	VARCHAR	128
TRANSEXITCODE	VARCHAR	128
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128
STARTEPOCH	NUMBER	12
ENDEPOCH	NUMBER	12
WHENACRONYM	VARCHAR	128
WHENCODE	NUMBER	9
IGNOREIDENTICAL	NUMBER	1
UPDATETIME	DATE	7
TRFCODE	NUMBER	9
TRFACRONYM	VARCHAR	128
STAGEIN	NUMBER	9
STAGEOUT	NUMBER	9
CPUMODEL	VARCHAR	256
NEVENTS	NUMBER	9

All Job Execution Records

10M

Snapshot filled with
Last 24hrs

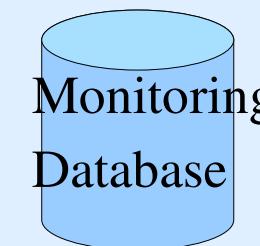


Design

Field	Type	Length
JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
METADATA	VARCHAR	4000
CPUCOUNT	NUMBER	20
CPUUNIT	VARCHAR	128
TRANSEXITCODE	VARCHAR	128
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128
STARTEPOCH	NUMBER	12
ENDEPOCH	NUMBER	12
WHENACRONYM	VARCHAR	128
WHENCODE	NUMBER	9
IGNOREIDENTICAL	NUMBER	1
UPDATETIME	DATE	7
TRFCODE	NUMBER	9
TRFACRONYM	VARCHAR	128
STAGEIN	NUMBER	9
STAGEOUT	NUMBER	9
CPUMODEL	VARCHAR	256
NEVENTS	NUMBER	9

24hr Window of Execution Records

100K



Summary Table

Filled from snapshot

Summary

Field	Type	Length
PROCESSINGDAY	DATE	7
NROFRECORDS	NUMBER	9
CPUTIME	NUMBER	12
WALLTIME	NUMBER	12
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
ERRORACRONYM	VARCHAR	128
JOBSTATUS	VARCHAR	64
EXECLUSTER	VARCHAR	128
EXQUEUE	VARCHAR	128

PRODUCTION

Field	Type	Length
JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
METADATA	VARCHAR	4000
CPUCOUNT	NUMBER	20
CPUUNIT	VARCHAR	128
TRANSEXITCODE	VARCHAR	128
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128
STARTEPOCH	NUMBER	12
ENDEPOCH	NUMBER	12
WHENACRONYM	VARCHAR	128
WHENCODE	NUMBER	9
IGNOREIDENTICAL	NUMBER	1
UPDATETIME	DATE	7
TRFCODE	NUMBER	9
TRFACRONYM	VARCHAR	128
STAGEIN	NUMBER	9
STAGEOUT	NUMBER	9
CPUMODEL	VARCHAR	256
NEVENTS	NUMBER	9

All Job Execution Records

10M

Snapshot filled with
Last 24hrs

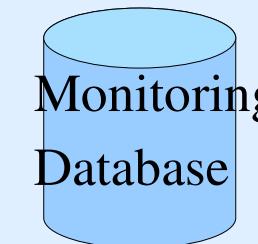
Production
Database

Design

Field	Type	Length
JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
METADATA	VARCHAR	4000
CPUCOUNT	NUMBER	20
CPUUNIT	VARCHAR	128
TRANSEXITCODE	VARCHAR	128
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128
STARTEPOCH	NUMBER	12
ENDEPOCH	NUMBER	12
WHENACRONYM	VARCHAR	128
WHENCODE	NUMBER	9
IGNOREIDENTICAL	NUMBER	1
UPDATETIME	DATE	7
TRFCODE	NUMBER	9
TRFACRONYM	VARCHAR	128
STAGEIN	NUMBER	9
STAGEOUT	NUMBER	9
CPUMODEL	VARCHAR	256
NEVENTS	NUMBER	9

24hr Window of Execution Records

100K



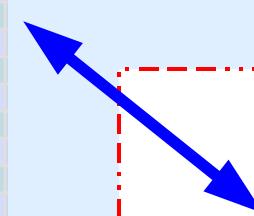
MONITORING

Field	Type	Length
PROCESSINGDAY	DATE	7
NROFRECORDS	NUMBER	9
CPUTIME	NUMBER	12
WALLTIME	NUMBER	12
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
ERRORACRONYM	VARCHAR	128
JOBSTATUS	VARCHAR	64
EXECLUSTER	VARCHAR	128
EXQUEUE	VARCHAR	128

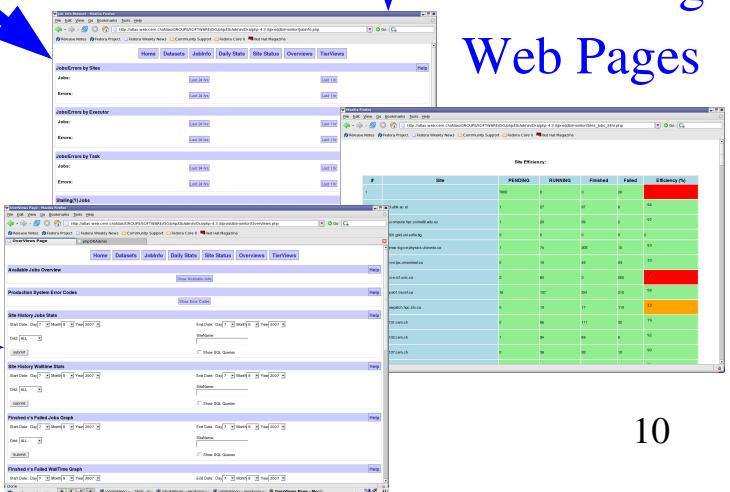
Summary

Summary Table

Filled from snapshot



Monitoring
Web Pages



Last 24hr View

- Snapshot table in the monitoring database contains 24hr sliding window of the job execution records.
- Pages generated from this table to give views of jobs/errors grouped by
 - Task
 - Site
 - Executor
- This 24hr view is aimed towards people on shift + spotting recent problems

Last 24hr View

The screenshot shows a Mozilla Firefox browser window titled "Job Info Browser - Mozilla Firefox". The address bar contains the URL <http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/jobInfo.php>. Below the address bar is a toolbar with links to "Release Notes", "Fedora Project", "Fedora Weekly News", "Community Support", "Fedora Core 6", and "Red Hat Magazine". The main content area features a navigation menu at the top with tabs: Home, Datasets, JobInfo (which is selected), Daily Stats, Site Status, Overviews, and TierViews. The "JobInfo" section is divided into several sub-sections: "Jobs/Errors by Sites", "Jobs/Errors by Executor", "Jobs/Errors by Task", "Stalling(?) Jobs", "Running Tasks", and "Detailed Executor info". Each sub-section includes buttons for "Last 24 hrs" and "Last 1 hr". The "Detailed Executor info" section has a "Executor ID:" input field and a "Done" button.

Views group info based on
Tasks – Sites - Executors
Aids problem solving

Last 24hr View

The screenshot shows the 'Job Info Browser - Mozilla Firefox' window. At the top, there's a menu bar with File, Edit, View, Go, Bookmarks, Tools, and Help. Below the menu is a toolbar with icons for Back, Forward, Stop, Refresh, Home, and a search bar. The address bar shows the URL: http://atlas.web.cern.ch/Atlas/GROUPS/ SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/jobInfo.php. The main content area has several tabs: Home, Datasets, JobInfo, Daily Stats, Site Status, Overviews, and TierViews. The 'JobInfo' tab is active. Under 'Jobs/Errors by Sites', there are sections for 'Jobs' and 'Errors', each with 'Last 24 hrs' and 'Last 1 hr' buttons. Under 'Jobs/Errors by Executor', similar sections are shown. Under 'Jobs/Errors by Task', sections for 'Jobs' and 'Errors' are present. At the bottom left, there's a 'Stalling(?) Jobs' section with 'By Site' and 'By Executor' buttons, and a 'Running Tasks' section with a 'Running Tasks' button. On the right, there's a 'Detailed Executor info' section with a 'Executor ID:' input field and a 'Done' button.

Details for sites giving
pend/run/fin/fail jobs
and colour coded efficiency
spot sinners!

06/09/2007 - John Kennedy

Views group info based on
Tasks – Sites - Executors
Aids problem solving

Click to get breakdown

The screenshot shows a table titled 'Site Efficiency:' from the 'Sites_jobs_24hr.php' page. The table has columns for #, Site, PENDING, RUNNING, Finished, Failed, and Efficiency (%). The data is as follows:

#	Site	PENDING	RUNNING	Finished	Failed	Efficiency (%)
1		7860	0	0	28	0
2	grid.uibk.ac.at	1	27	97	6	94
3	lcg-compute.hpc.unimelb.edu.au	1	20	59	2	97
4	ce001.grid.uni-sofia.bg	3	0	0	0	0
5	bigmac-lcg-ce.physics.utoronto.ca	1	74	205	15	93
6	lcg-ce.lps.umontreal.ca	0	10	46	93	33
7	lcg-ce.rcf.uvic.ca	2	60	3	285	1
8	lcgce01.triumf.ca	18	107	291	210	58
9	snowpatch.hpc.sfu.ca	0	13	17	110	13
10	ce101.cern.ch	0	66	111	30	79
11	ce102.cern.ch	1	34	89	8	92
12	ce107.cern.ch	0	39	93	10	90

Last 24hr View

#	Site	PENDING	RUNNING	Finished	Failed	Efficiency (%)
1		7860	0	0	28	0
2	grid.uibk.ac.at	1	27	97	6	94
3	lcg-compute.hpc.unimelb.edu.au	1	20	59	2	97
4	ce001.grid.uni-sofia.bg	3	0	0	0	0
5	bigmac-lcg-ce.physics.utoronto.ca	1	74	205	15	93
6	lcg-ce.lps.umontreal.ca	0	10	46	93	33
7	lcg-ce.rcf.uvic.ca	2	60	3	285	1
8	lcgce01.triumf.ca	18	107	291	210	58
9	snowpatch.hpc.sfu.ca	0	13	17	110	13
10	ce101.cern.ch	0	66	111	30	79
11	ce102.cern.ch	1	34	89	8	92
12	ce107.cern.ch	0	39	93	10	90

Last 24hr View

#	Site	PENDING	RUNNING	Finished	Failed	Efficiency (%)
1		7860	0	0	28	0
2	grid.uibk.ac.at	1	27	97	6	94
3	lcg-compute.hpc.unimelb.edu.au	1	20	59	2	97
4	ce001.grid.uni-sofia.bg	3	0	0	0	0
5	bigmac-lcg-ce.physics.utoronto.ca	1	74	205	15	93
6	lcg-ce.lps.umontreal.ca	0	10	46	93	33
7	lcg-ce.rcf.uvic.ca	2	60	3	285	1
8	lcgce01.triumf.ca	18	107	291	210	58
9	snowpatch.hpc.sfu.ca	0	13	17	110	13
10	ce101.cern.ch	0	66	111	30	
11	ce102.cern.ch	1	34	89	8	
12	ce107.cern.ch	0	39	93	10	

Site	EXELEXOR GLITE WMS	WRAPLCG JTINST GET	WRAPLCG STAGEIN LCGCP	EXELEXOR GETOUT UNTAR	EXELEXOR GETOUT OBJLOAD
snowpatch.hpc.sfu.ca	6	17	28	15	44

Click on efficiency to get breakdown of errors

Last 24hr View

#	Site	PENDING	RUNNING	Finished	Failed	Efficiency (%)
1		7860	0	0	28	0
2	grid.uibk.ac.at	1	27	97	6	94
3	lcg-compute.hpc.unimelb.edu.au	1	20	59	2	97
4	ce001.grid.uni-sofia.bg	3	0	0	0	0
5	bigmac-lcg-ce.physics.utoronto.ca	1	74	205	15	93
6	lcg-ce.lps.umontreal.ca	0	10	46	93	33
7	lcg-ce.rcf.uvic.ca	2	60	3	285	1
8	lcgce01.triumf.ca	18	107	291	210	58
9	snowpatch.hpc.sfu.ca	0	13	17	110	13
10	ce101.cern.ch	0	66	111	30	
11	ce102.cern.ch	1	34	89	8	
12	ce107.cern.ch	0	39	93	10	

Click on efficiency to get breakdown of errors

Site	EXELEXOR GLITE WMS	WRAPLCG JTINST GET	WRAPLCG STAGEIN LCGCP	EXELEXOR GETOUT UNTAR	EXELEXOR GETOUT OBJLOAD
snowpatch.hpc.sfu.ca	6	17	28	15	44

Click # to get logfiles (egee)

#	JOBEXID	CREATIONTIME	MODIFICATIONTIME	JOBDEFFK	TASKFK	JOBNAME	PARTNR	ATTEMPTNR	SUPERVISOR
1	9275947	01-AUG-07	05-AUG-07	6510352	11066	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	25715	1	silvia_nu#9944341
2	9257623	31-JUL-07	05-AUG-07	6523943	11065	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	17912	2	silvia_nu#9944341
3	9264016	31-JUL-07	05-AUG-07	6515818	11065	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	15954	1	silvia_nu#9944341
4	9274086	01-AUG-07	05-AUG-07	6499380	11063	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	6604	2	silvia_nu#9944341
5	9256401	31-JUL-07	05-AUG-07	6534325	11065	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	19896	1	silvia_nu#9944341
6	9259134	31-JUL-07	05-AUG-07	6534405	11065	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	19976	1	silvia_nu#9944341
7	9268994	01-AUG-07	05-AUG-07	6500270	11065	misall_mc12_005801.JF7_pythia_loosejet_filter.digit.v12000605	12574	2	silvia_nu#9944341

06/09/2007 - Job

Click on error to get detailed info job execution info

Tasks Status

DataSet View

Dataset Status: All

Grid: All

Match Dataset

Dataset name: PythiattH160WW

Submit

How to Find Datasets Help

Select Task by
Grid
State and Name

Tasks Status

DataSet View

Dataset Status: All Grid: All

Match Dataset

Dataset name: PythiattH160WW

[How to Find Datasets](#) [Help](#)

Select Task by
Grid
State and Name

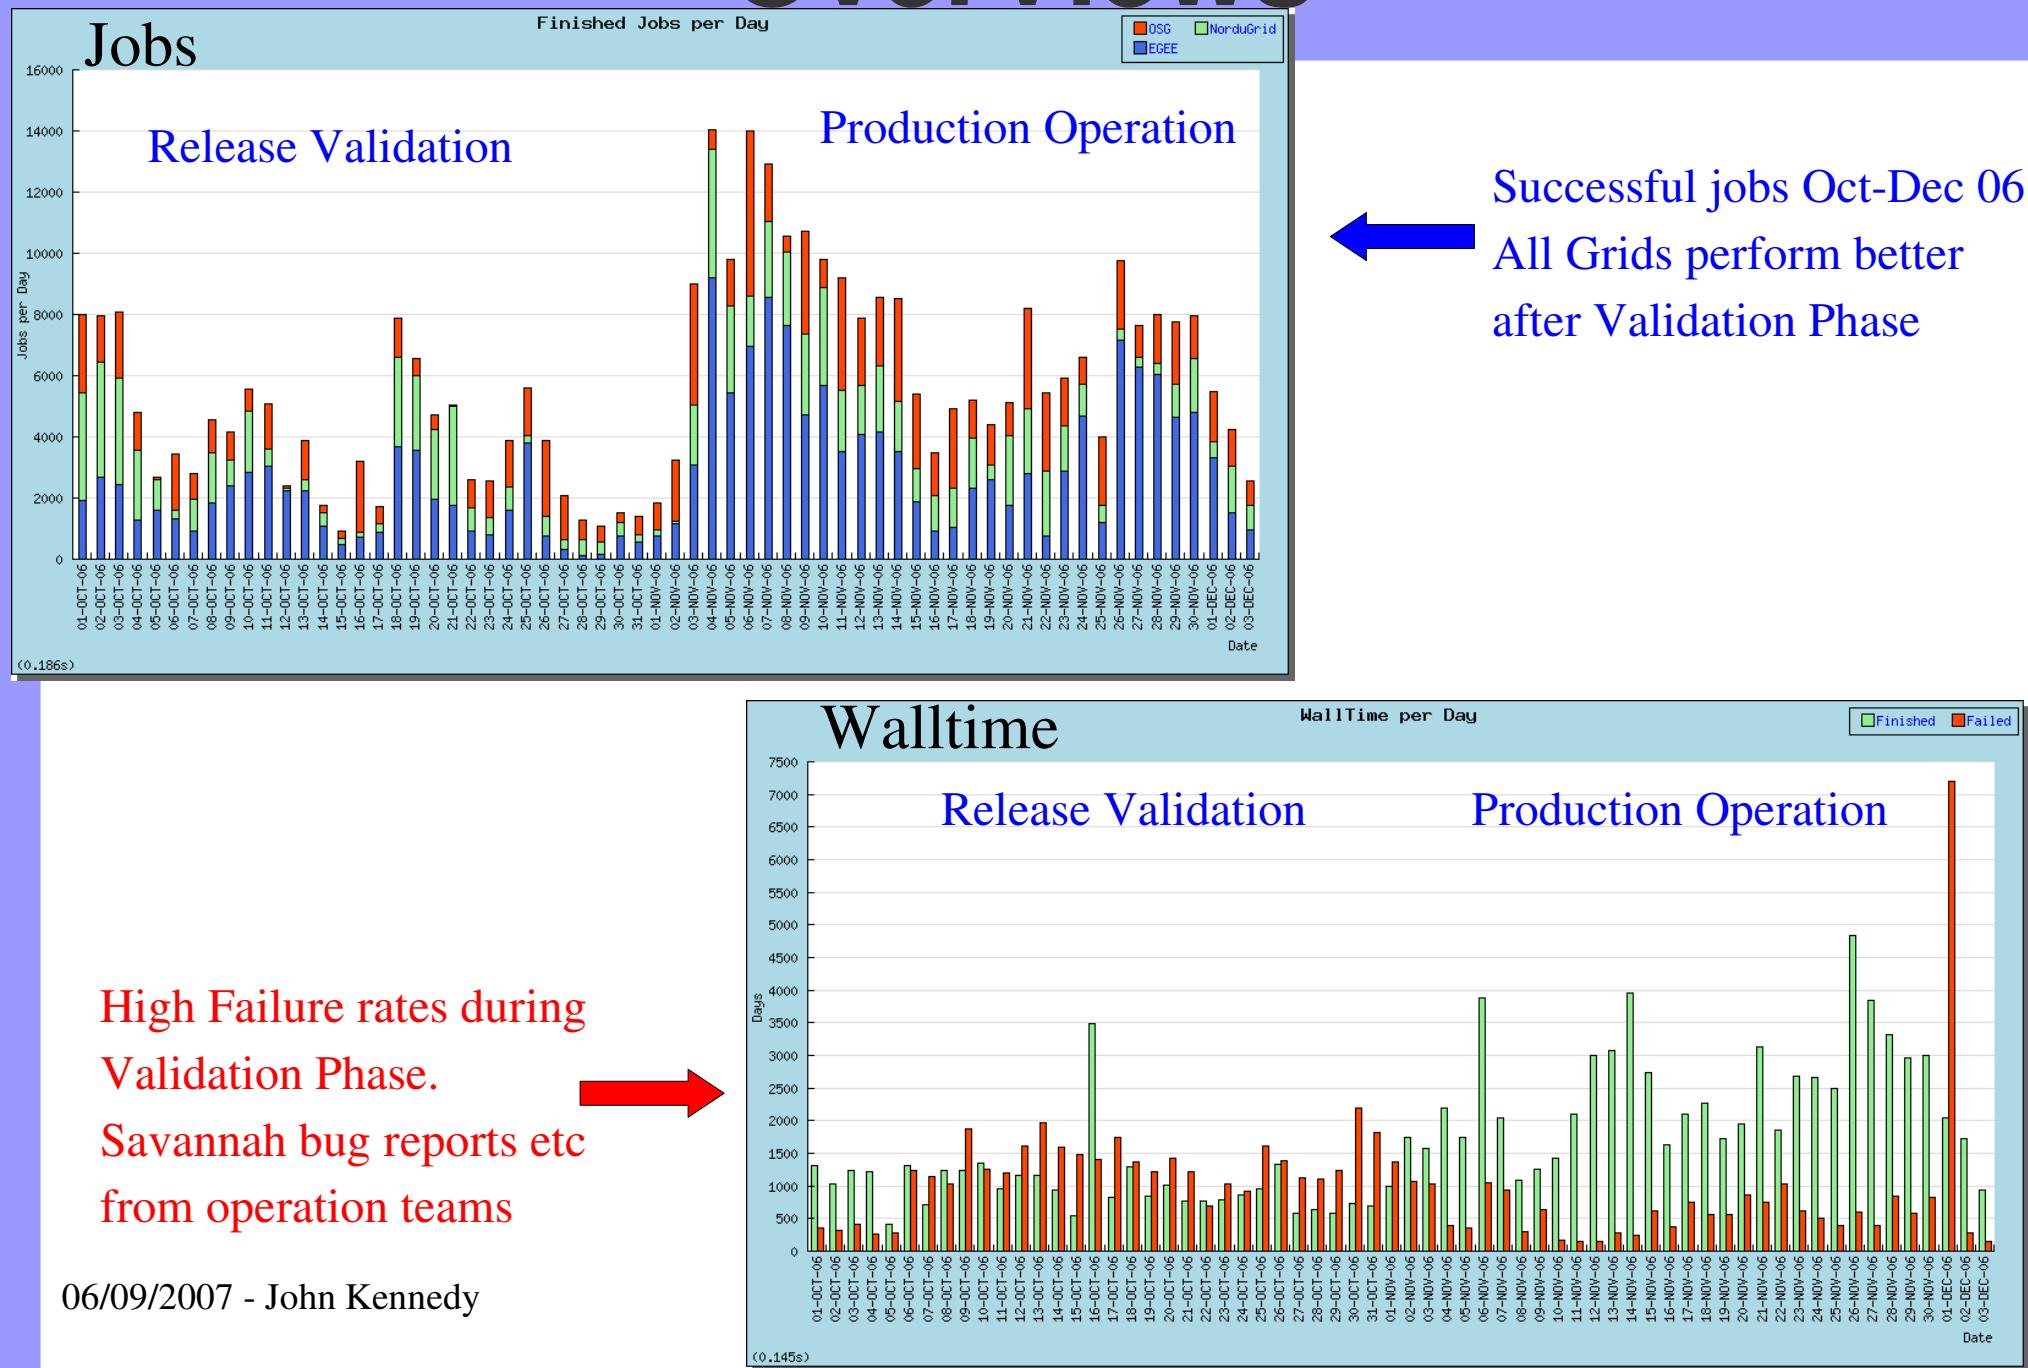
Click to get
Details

Dataset/Task	TaskID	Done	Running	To be done	WaitInp	Aborted	Completed(%)	Grid	Atlas ver	Tier
misal1_mc12.005345.PythiattH160WW3l3v4j_neglep_H1T2.digit.v12003107.task	4696	99	0	1	0	0	99.00	LCG-DQ	Atlas-12.0.31	ES
misal1_mc12.005347.PythiattH160WW3l3v4j_neglep_H2T1.digit.v12003107.task	4698	99	0	1	0	0	99.00	LCG-DQ	Atlas-12.0.31	UK
trig1_misal1_mc12.005345.PythiattH160WW3l3v4j_neglep_H1T2.recon.v12000601.task	6448	19	0	1	0	0	95.00	LCG-DQ	Atlas-12.0.6	ES
trig1_misal1_mc12.005346.PythiattH160WW3l3v4j_poslep_H2T1.recon.v12000601.task	6449	18	0	2	0	0	90.00	LCG-DQ	Atlas-12.0.6	DE
trig1_misal1_mc12_V1.005346.PythiattH160WW3l3v4j_poslep_H2T1.recon.v12000601.task	7855	19	0	1	0	0	95.00	LCG-DQ	Atlas-12.0.6	IT
trig1_pile1sf05_misal1_mc12.005346.PythiattH160WW3l3v4j_poslep_H2T1.recon.v12000605.task	9861	21	1	0	0	8	95.45	LCG-DQ	Atlas-12.0.6	IT
pile1sf05_misal1_mc12.005346.PythiattH160WW3l3v4j_poslep_H2T1.digit.v12000605.task	9860	32	0	3	0	5	91.43	LCG-DQ	Atlas-12.0.6	IT
misal1_mc12.005345.PythiattH160WW3l3v4j_neglep_H1T2.digit.v12003103.task	7833	99	0	1	0	0	99.00	LCG-DQ	Atlas-12.0.31	ES
trig1_misal1_mc12_V1.005345.PythiattH160WW3l3v4j_neglep_H1T2.recon.v12000601.task	9740	13	2	0	0	5	86.67	LCG-DQ	Atlas-12.0.6	ES
trig1_misal1_mc12_V1.005346.PythiattH160WW3l3v4j_poslep_H2T1.recon.v12000601.task	9741	18	0	1	0	1	94.74	LCG-DQ	Atlas-12.0.6	IT
misal1_mc12.005345.PythiattH160WW3l3v4j_neglep_H1T2.digit.v12003103.task	12144	145	55	0	0	0	72.50	LCG-DQ	Atlas-12.0.31	ES
trig1_misal1_mc12_V1.005345.PythiattH160WW3l3v4j_neglep_H1T2.recon.v12000601.task	12145	0	0	5	0	0	0.00	LCG-DQ	Atlas-12.0.6	None
misal1_mc12.005346.PythiattH160WW3l3v4j_poslep_H2T1.digit.v12003107.task	4697	100	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	DE
trig1_misal1_mc12.005344.PythiattH160WW3l3v4j_poslep_H1T2.recon.v12000601.task	6362	39	0	0	0	1	100.00	OSG	Atlas-12.0.6	None
mc12.005344.PythiattH160WW3l3v4j_poslep_H1T2.evgen.v12003102.task	3706	8	0	0	0	2	100.00	LCG-DQ	Atlas-12.0.31	DE
mc12.005345.PythiattH160WW3l3v4j_neglep_H1T2.evgen.v12003102.task	3707	10	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	ES
misal1_mc12.005345.PythiattH160WW3l3v4j_neglep_H1T2.digit.v12003107.task	6481	100	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	ES
misal1_mc12.005346.PythiattH160WW3l3v4j_poslep_H2T1.digit.v12003107.task	6482	100	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	IT
misal1_mc12.005347.PythiattH160WW3l3v4j_neglep_H2T1.digit.v12003107.task	6483	100	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	UK
mc12.005346.PythiattH160WW3l3v4j_poslep_H2T1.evgen.v12003105.task	4228	10	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	IT
mc12.005347.PythiattH160WW3l3v4j_neglep_H2T1.evgen.v12003105.task	4229	10	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.31	UK
misal1_mc12.005344.PythiattH160WW3l3v4j_poslep_H1T2.digit.v12003108.task	5453	200	0	0	0	0	100.00	OSG	Atlas-12.0.31	None
trig1_misal1_mc12.005347.PythiattH160WW3l3v4j_neglep_H2T1.recon.v12000601.task	6450	20	0	0	0	0	100.00	LCG-DQ	Atlas-12.0.6	UK
mc12.005344.PythiattH160WW3l3v4j_poslep_H1T2.evgen.v12000401.task	4399	10	0	0	0	0	100.00	OSG	Atlas-12.0.4	None
trig1_misal1_mc12_V1.005347.PythiattH160WW3l3v4j_neglep_H2T1.recon.v12000601.task	7856	19	0	0	0	1	100.00	LCG-DQ	Atlas-12.0.6	ES

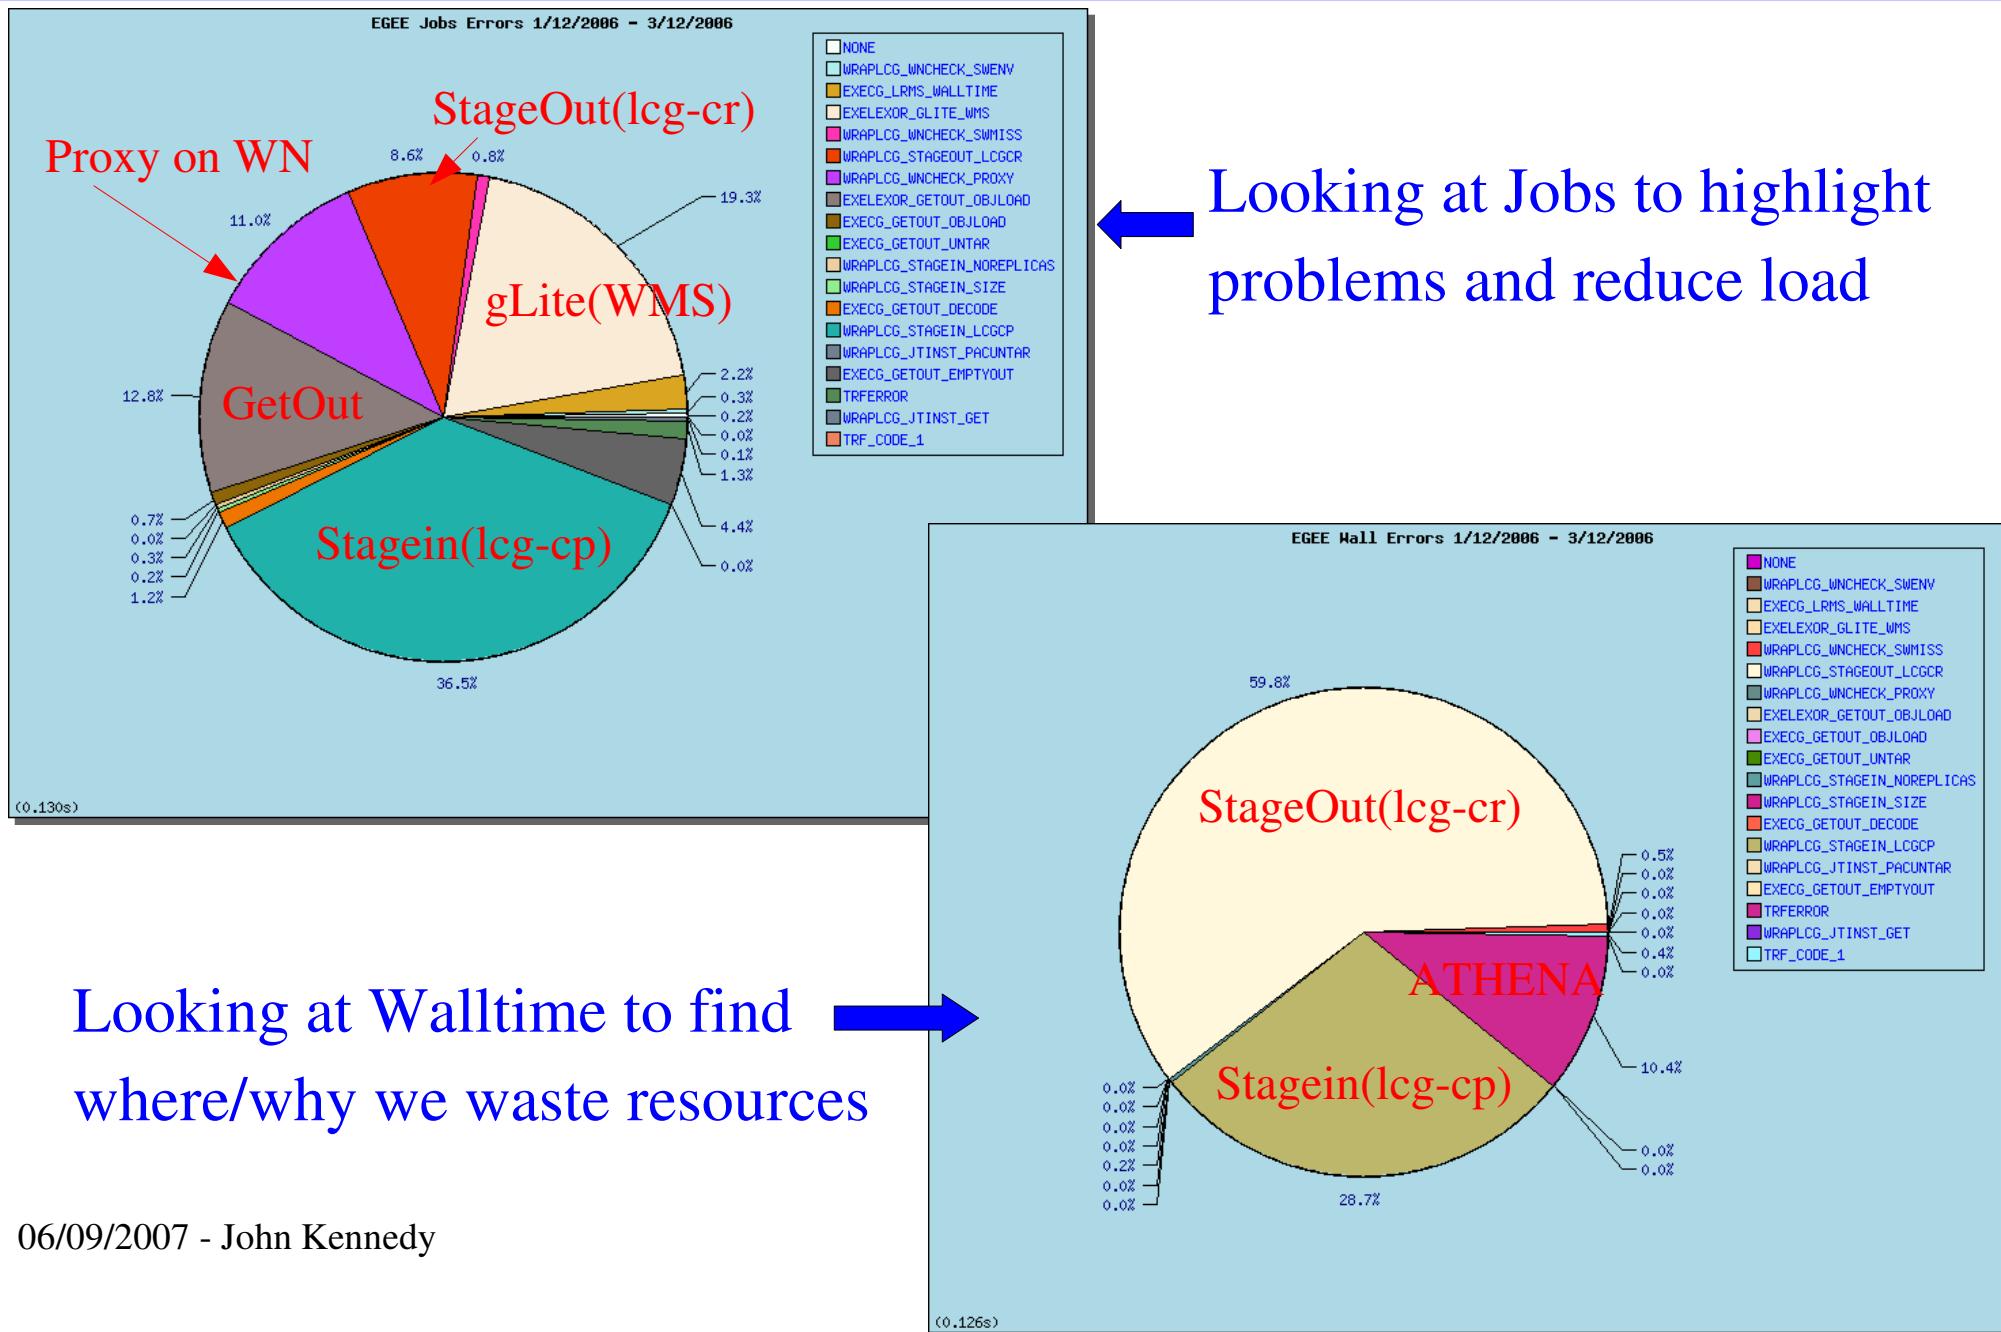
OverViews

- Allow Views to be generated over defined time windows
- Uses Summary table
 - Ensures quick response
 - Low load on database
- Provide Views of
 - Finish/Failed Jobs/Walltime
 - Error reports, pie charts
 - Allocated Tasks/jobs – ensure grids have jobs to run
 - Allow split by grid flavour and sites

OverViews



OverViews



OverViews

Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/00/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/available_jobs.html
Release Notes Fedora Project Fedora Weekly News Community Support Fedora Core 6 Red Hat Magazine

Available Jobs Overview Ran at 13:25:09 on 06-08-2007

SW Version	Grid	Status	Num Jobs
None	TEST-LCG	TOBEDONE	1.0
Atlas-13.0.10	NORDUGRID	TOBEDONE	11.0
Atlas-13.0.10	OSG	TOBEDONE	57.0
Atlas-12.0.95	LCG-DQ	TOBEDONE	1.0
Atlas-12.0.7	LCG-DQ	TOBEDONE	61.0
Atlas-12.0.7	LCG-DQ	WAITINGINPUT	211.0
Atlas-12.0.7	OSG	TOBEDONE	44.0
Atlas-12.0.7	OSG	WAITINGINPUT	20.0
Atlas-12.0.6	CRONUS	WAITING	9.0
Atlas-12.0.6	CRONUS_NORDIC	WAITING	9.0
Atlas-12.0.6	LCG-DQ	TOBEDONE	44563.0
Atlas-12.0.6	LCG-DQ	WAITING	174.0
Atlas-12.0.6	LCG-DQ	WAITINGINPUT	5.0
Atlas-12.0.6	NORDUGRID	TOBEDONE	37.0
Atlas-12.0.6	OSG	TOBEDONE	1053.0
Atlas-12.0.6	OSG	WAITING	78000.0
Atlas-12.0.6	OSG	WAITINGCOPY	4610.0
Atlas-12.0.6	OSG	WAITINGINPUT	10.0
Atlas-12.0.4	OSG	WAITINGINPUT	
Atlas-12.0.31	LCG-DQ	TOBEDONE	

Done

Definitions of Error Codes/Acronyms with Meaning
- helps shifters to interpret problems

Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/00/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/error_codes.html
Release Notes Fedora Project Fedora Weekly News Community Support Fedora Core 6 Red Hat Magazine

Error Acronyms Ran at 03:07:35 on 06-08-2007

ErrorCode	ErrorAcronym	Meaning
0	ALLOK	no errors occurred
100000	SUP	unspecified supervisor error
100100	SUP_DDM	unspecified supervisor DDM error
100101	SUP_DDM_TARGETEXISTS	attach (rename) failed because target file already exists
100102	SUP_DDM_NOSOURCE	attach (rename) failed because source file does not exist
200000	EXE	unspecified executor error
200100	EXE_INPUTS	unspecified input file error
200101	EXE_INPUTS_NOTINDS	An input file is not in the dataset as specified in job pars
200102	EXE_INPUTS_NOSURLS	An input file has no surls after successful LRC lookup
200103	EXE_INPUTS_TMPNOSURLS	An input file has no surls but an LRC lookup failed
200104	EXE_INPUTS_TMPNOTINDS	An input file is not in the dataset, but one DDM interaction failed
200201	EXE_CONTROL_STALESTATUS	Job removed by executor because status was unchanged for N days
200300	EXE_JOBDEF	Job definition leads to error
200301	EXE_JOBDEF_LOGREGEEXP	Log tgz filename cannot be regexped to get log file name. possible invalid character in log name.
210000	EXECG	unspecified CG executor error
210100	EXECG_GETOUT	Unspecified problem in getOutput: stdout is parsed, uudecode, and untarred.
210101	EXECG_GETOUT_MISROUTERR	stdout and stderr both do not exist
210102	EXECG_GETOUT_MISSOUT	stdout files does not exist
210103	EXECG_GETOUT_EMPTYOUT	stdout has zero length
210104	EXECG_GETOUT_DECODE	uudecode of stdout failed
210105	EXECG_GETOUT_UNTAR	Untar/zip of extracted out.tar.gz file failed
210106	EXECG_GETOUT_OBJMISS	Wrapper output object is missing
210107	EXECG_GETOUT_OBJLOAD	Wrapper output object exist but could not be unpickled.
210200	EXECG_LRMS	Unspecified local batch system problem
210201	EXECG_LRMS_WALLTIME	Job exceeded waltime limit
220000	EXELEXOR	unspecified Lexor executor error

Done

EGEE Shift System

- During the last 12 months an EGEE shift system has been put in place
- People responsible for
 - Assigning Tasks to Tiers
 - Monitoring Jobs
 - Monitoring Data movement
 - Ensuring Tasks are progressing
 - Reporting problems to sites/physicists
- Strong users of Monitoring + lots of feedback
- Pages Developed to help shifters

EGEE Shift Pages

Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/site_problems.html
Release Notes Fedora Project Fedora Weekly News Community Support Fedora Core 6 Red Hat Magazine

Site Problems View Ran at 13:45:33 on 06-08-2007

Cluster	Host	Count
ce01.ariagni.hellasgrid.gr	wn028.ariagni.hellasgrid.gr	26.0
epgce2.ph.bham.ac.uk	epbf081.ph.bham.ac.uk	31.0
epgce2.ph.bham.ac.uk	epbf082.ph.bham.ac.uk	25.0
epgce2.ph.bham.ac.uk	epbf084.ph.bham.ac.uk	13.0
epgce2.ph.bham.ac.uk	epbf085.ph.bham.ac.uk	2.0
leg-ce.ref.uvic.ca	c01b12	1.0
leg-ce.ref.uvic.ca	c02b04	1.0
leg-ce.ref.uvic.ca	c02b10	1.0
leg-ce.ref.uvic.ca	c05b12	1.0
leg-ce.ref.uvic.ca	c05b14	1.0
leg-ce.ref.uvic.ca	c06b10	1.0
leg-ce.ref.uvic.ca	c06b12	1.0
leg-ce.ref.uvic.ca	c06b13	1.0
leg-ce.ref.uvic.ca	c06b14	1.0
leg-ce.ref.uvic.ca	c12b09	1.0
tnb20.nikhef.nl	wn-lui2-007.farm.nikhef.nl	3.0

Done

Identify sites with ATLAS SW installations problems

Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/running_egee_tasks.php
Release Notes Fedora Project Fedora Weekly News Community Support Fedora Core 6 Red Hat Magazine

Running EGEE Tasks

#	TaskName	TaskID	CreationTime	ModificationTime	Tier	Ndone	Ntotal	Fraction done (%)
1	trig1_misall_mc12.005807.JF35_pythia_jet_filter.recon.v12000604.task	9372	19-05-2007	06-08-2007	IT	3108	3983	78.0
2	misall_mc12.005212.MeAtNlo_ttbar_plus_jets.digit.v12000604.task	9692	23-05-2007	03-08-2007	FR	21970	22000	99.9
3	trig1_misall_mc12.006311.HerwigVBFH140gamgam.recon.v12000604.task	9733	24-05-2007	28-07-2007	IT	65	80	81.2
4	trig1_misall_mc12_V1.005346.Pythia1ttH160WW3l3v4j_poslep_H2T1.recon.v12000601.task	9741	24-05-2007	23-07-2007	IT	18	20	90.0
5	pile1sf05_misall_mc12.005175.AcerMC_Zbb_2l.digit.v12000605.task	9805	24-05-2007	03-08-2007	IT	330	399	82.7
6	pile1sf05_misall_mc12.005348.Pythia1ttH120inv.digit.v12000605.task	9833	25-05-2007	02-08-2007	IT	21	40	52.5
7	pile1sf05_misall_mc12.005346.Pythia1ttH160WW3l3v4j_poslep_H2T1.digit.v12000605.task	9860	25-05-2007	26-07-2007	IT	32	40	80.0
8	pile1sf05_misall_mc12.006315.MeAtNlo_Jimmy_H170Wp_WWem.digit.v12000605.task	9864	25-05-2007	21-07-2007	IT	12	20	60.0
9	trig1_misall_mc12.006379.PythiaPhotonJet_AsymJetFilter.recon.v12000603.task	9704	25-05-2007	03-08-2007	IT	3250	3968	81.9
10	trig1_misall_mc12.005801.AcerMC_schan.recon.v12000604.task	9896	25-05-2007	21-07-2007	IT	40	100	40.0
11	trig1_misall_mc12.008091.pythia_15_NJ2_FMET100.merge.v12000605.task	9913	25-05-2007	02-08-2007	IT	24	28	85.7
12	trig1_misall_mc12.005720.Pythia_B_bb_lpsie2e2Xmu6.recon.v12000604.task	10033	30-05-2007	03-08-2007	NL	161	200	80.5
13	trig1_misall_mc12.005212.MeAtNlo_ttbar_plus_jets.recon.v12000605.task	10093	31-05-2007	06-08-2007	FR	4468	4394	101.7
14	trig1_misall_mc12.005212.MeAtNlo_ttbar_plus_jets.merge.v12000605.task	10094	01-06-2007	29-07-2007	FR	534	540	98.9
15	pile1sf05_misall_mc12.005212.MeAtNlo_ttbar_plus_jets.digit.v12000605.task	10227	02-06-2007	03-08-2007	FR	5096	4392	116.0
16	misall_mc12.005802.JF17_pythia_jet_filter.digit.v12000601.task	10154	05-06-2007	31-07-2007	FR	10093	10000	100.9
17	misall_mc12.005802.JF17_pythia_jet_filter.digit.v12000601.task	10155	05-06-2007	03-08-2007	FR	10115	10000	101.1
18	misall_mc12.005802.JF17_pythia_jet_filter.digit.v12000601.task	10432	08-06-2007	03-08-2007	FR	100	200	80.5

Done

Monitor Running Tasks

Identify old Tasks which need action

Identify nearly completed Tasks

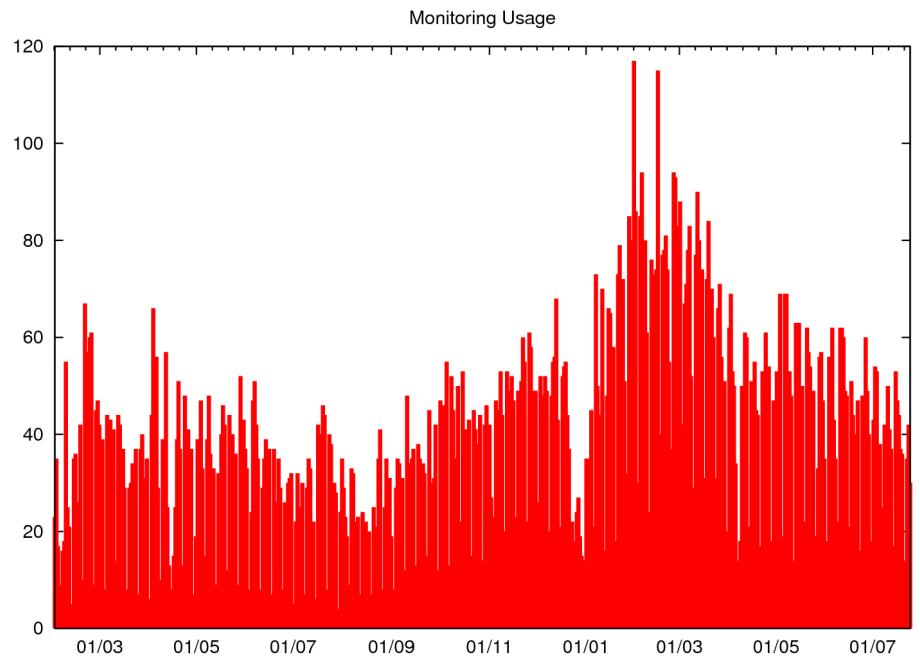
06/09/2007 - John Kennedy

Monitoring Usage

- Monitor the Monitoring
- Monitor usage of main web pages
 - Total # users
 - Busy pages
- Log SQL query times of Cron Jobs
 - Spot increasing load and problems when crons run
- Monitoring account which can be tracked

Monitoring Usage

- Monitor the Monitoring
- Monitor usage of main web pages
 - Total # users
 - Busy pages
- Log SQL query times of Cron Jobs
 - Spot increasing load and problems when cron jobs run
- Monitoring account which can be tracked



Currently O(40-50)
Individual users per day

Conclusion

- The model of using cooked down monitoring tables has proved successful
 - Quicker queries
 - Reduced load on database
- Both High and low level views have been implemented allowing for
 - Problem Identification
 - Accounting
- EGEE shift pages are well developed due to good communication with users – I am part of the shift team!
- Much more functionality provided than originally expected
- From a personal point of view - It's been fun and it's useful

Outlook

- Much has been learned during the project
 - By developers
 - Snapshot/summary table use
 - What users want
 - By users
 - What's useful
 - What's not
 - What would be nice
- Prod-sys monitoring being integrated into ARDA dashboard
 - Functionality based on experience gained here
 - Database snapshot tables etc will remain
 - PHP replaced by new framework
- Migration has started – expected to be complete by start of

THANKS!

Thanks for you time and attention!